

SEEKING CLEANER AIR – ACCESSIBLE AIR QUALITY MONITORING
AND EVENT DIAGNOSIS

by

James Allen Hall Jr.

A dissertation

submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy in Electrical and Computer Engineering

Boise State University

May 2014

© 2014

James Allen Hall Jr.

ALL RIGHTS RESERVED

BOISE STATE UNIVERSITY GRADUATE COLLEGE

DEFENSE COMMITTEE AND FINAL READING APPROVALS

of the dissertation submitted by

James Allen Hall Jr.

Dissertation Title: Seeking Cleaner Air – Accessible Air Quality Monitoring
and Event Diagnosis

Date of Final Oral Examination: 08 April 2014

The following individuals read and discussed the dissertation submitted by student James Allen Hall Jr., and they evaluated his presentation and response to questions during the final oral examination. They found that the student passed the final oral examination.

Sin Ming Loo, Ph.D.	Chair, Supervisory Committee
Jim Browning, Ph.D.	Member, Supervisory Committee
Dale Stephenson, Ph.D.	Member, Supervisory Committee
Vishal Saxena, Ph.D.	Member, Supervisory Committee
Clifford Weisel, Ph.D.	External Examiner

The final reading approval of the dissertation was granted by Sin Ming Loo, Ph.D., Chair of the Supervisory Committee. The thesis was approved for the Graduate College by John R. Pelton, Ph.D., Dean of the Graduate College.

ACKNOWLEDGEMENTS

I owe this work to the boundless love and support of my wife Melinda. Her patience and encouragement continue to sustain me. I owe my thanks to my family. My daughters Sarah and Rachel have given me strength with their energy and enthusiasm. My mother and father, Myrtle and Jim Hall, have helped us all through good and bad times. My wife's parents, Barbara and Dennis Schrecongost, have cared for us and carried us through. My sister, Julie Battermann, stood by me when I was weakest. My family is my life, and I thank them for this time of learning.

I would like to thank my advisor, Dr. Sin Ming Loo, for his support and guidance, and for his patience with my long road. I owe much to the team at the Harman Systems Integration Lab, for their friendship and their work on the wireless sensor system. I would also like to thank my committee, Dr. Jim Browning, Dr. Vishal Saxena, and Dr. Dale Stephenson, for their help and their time.

I would like to thank my friends and colleagues at my former workplace for their understanding in letting me go, and my friends and colleagues at my current workplace for their patience in letting me finish.

Finally, I would like to thank those in the medical profession who brought me back from illness to health. Dr. Natalie Shiflett was the first to diagnose me, and helped my family come to terms with my illness. Dr. Dan Zuckerman, Dr. Stephen Smith, and the staff at MSTI cared for me over those long months of treatment.

ABSTRACT

Airborne particulate matter has been shown to be associated with morbidity and mortality, and can have a detrimental impact on sensitive materials and processes. Understanding the levels and movements of particulate matter in an enclosed space can lead to a reduction in the negative consequences of poor air quality. A system of environmental sensors including particulate matter, selected gases, and other atmospheric factors can be used to provide a real-time assessment of air quality. This assessment can be used to assist in the diagnosis of the source of particulate matter. This dissertation describes the creation of a framework for the monitoring and diagnosis of air quality events that will assist in the move towards cleaner air. To achieve this goal, a low-cost optical particle counter is created for incorporation into a wireless sensor network (WSN), including the exploration of novel methods for improving the quality of such a sensor. A WSN is developed using the aforementioned sensors, along with a system for data collection and visualization. Finally, a distributed event detection framework is formulated to facilitate human interaction with the system. Test results from the prototype system are presented for real-world environments where air quality is of significant concern. The results of this work may be leveraged into multiple applications that assist in the diagnosis and treatment of air quality issues that have tangible impact on health and material quality.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iv
ABSTRACT	v
LIST OF TABLES	xiv
LIST OF FIGURES	xv
LIST OF ABBREVIATIONS	xxxv
CHAPTER ONE: INTRODUCTION	1
Motivations	5
Contribution	8
Health and Material Impact	8
Accessibility through Lower Cost	8
Simplifying Analysis	9
System for a New Paradigm	9
Work Summary	10
System Design and Prototype	11
System Simulation	11
Particle Detection and Sizing.....	11
Component and System Optimization	12
Wireless Sensor Network Platform.....	12
Detection Framework.....	12

Demonstration.....	12
CHAPTER TWO: PREVIOUS WORK	14
CHAPTER THREE: WIRELESS SENSOR NETWORK SYSTEM.....	17
System Overview	18
System Deployments	18
Sensor Node Hardware	20
Microcontroller	21
Power Management	22
Communication.....	24
Data Storage.....	25
User Interface.....	26
Sensors	26
Sensor Node Firmware	28
Tasks	29
Device Drivers	30
Scheduler.....	31
Sensor Management.....	31
Storage	32
Communications	32
Data Manager.....	32
Application.....	34
Wireless Sensor Network.....	35
Network Layer	35

Time Synchronization.....	36
Server Software.....	38
Connector Server	38
WAQM Database.....	39
Web Graphing Interface.....	40
Visualization Software.....	41
Wireless Sensor Network System Conclusions	42
CHAPTER FOUR: OPTICAL PARTICLE COUNTER.....	44
Analog Front End Design	44
AFE Power Supply	47
Operational Amplifier Choice.....	48
Baseline AFE Circuit	50
Power Line Noise Coupling.....	68
Bootstrap AFE Circuit	76
Analog-to-Digital Conversion	88
Laser Source.....	91
Laser Module	92
Power Output	94
Laser Beam Profile	96
Reliability.....	98
Laser Output Stability	99
Mechanical Design.....	102
Airflow	104

Incident Laser Light.....	106
Reflected Laser Light.....	108
Ambient Light.....	110
Particle Counter Hardware Cost	110
Particle Counter Hardware Conclusions.....	111
CHAPTER FIVE: SIMULATION	114
Simulation Flow.....	114
Coordinate System.....	116
Segmenting Sensors	117
Particle Sets.....	120
Particle Vectors.....	120
Laser Intensity and Polarization.....	122
Scattering Irradiance	126
Irradiance Components	129
Angle of Incidence Scaling.....	132
Solid Angle Calculations	133
Segment Integration	135
Basic Scattering Simulations	137
Sensor Irradiance Simulations	140
Single Sensor Irradiance	152
Dual Sensor Irradiance.....	158
Tandem Sensors	161
Orthogonal Sensors.....	168

Operational Simulations.....	175
Simulation with Other Particle Types.....	177
Modified Laser Wavelength	182
Simulation Conclusions	186
CHAPTER SIX: DETECTION AND SIZING.....	189
Particulate Generation and Measurement Comparisons	190
Single Sensor Detection and Sizing.....	192
Particle Detection.....	193
Particle Sizing	198
Example Data.....	205
Dual Sensor Detection and Sizing	213
Particle Detection.....	215
Particle Sizing	218
Example Data.....	220
Misfit Particles	239
Sizing Implementation	242
Detection and Sizing Conclusions	247
CHAPTER SEVEN: PARTICLE COUNTER TESTING AND CALIBRATION.....	250
OPC Error Sources.....	250
Air Flow	252
Sampling Loss.....	254
Laser Variability	257
Analog Front End Variability	259

Digitizing Error	260
Coincidence Error	262
Error Mitigation	264
OPC Calibration System	267
Calibration Methods.....	274
Cold and Hot Calibration	274
PSL Particle Calibration	281
Testing and Calibration Conclusions	287
CHAPTER EIGHT: EVENT DETECTION FRAMEWORK.....	289
Dynamic Event Triggering	291
Sensor Node Event Triggers	291
WSN-Level Event Triggers	294
User Interaction.....	296
User Notification.....	296
Event Labeling	299
Event Detection Conclusions.....	300
CHAPTER NINE: SYSTEM DEMONSTRATION	302
Controlled Environment Testing.....	303
Simulated Aircraft Cabin	304
Sensor Node Configuration.....	306
Data Generation and Processing	307
Two-Dimensional Testing	308
Three-Dimensional Testing	319

Controlled Environment Testing Conclusions.....	334
Field Environment Testing Inside an Operating Room	335
Operating Room Layout	336
Operating Room Sensor Configuration	337
Operating Room Data Access	338
Operating Room Air Supply	339
Operation Room Data	341
Background Levels of Particulate Matter	350
Operating Room Air Movement	351
Humidifier Impact on Particulate Matter Levels	354
False Particle Peaks.....	356
Summary of Operating Room Testing.....	358
Field Environment Testing in Residential Homes	360
CHAPTER TEN: FUTURE WORK.....	366
Wireless Sensor Network System Improvements.....	366
Mesh Network Software	366
Flash Buffering	367
Preemptive Multitasking.....	369
Particle Counter Hardware Areas for Improvement	370
Improvements to the AFE.....	370
Laser Source.....	370
Airflow Improvements.....	372
Light Management Improvements.....	372

Simulation Areas for Improvement	373
Front-End Improvements	373
Laser Beam Intensity Assumptions	374
Particle Shape Assumptions.....	376
Simulation of X-Axis Movement.....	377
Detection and Sizing Areas for Improvement	377
Sizing Larger Particles.....	378
Further Testing of the Orthogonal Sensor Arrangement	380
Other Algorithms and Metrics	381
Particle Counter Calibration System Improvements.....	383
Large Particle Aerosol Generation.....	383
Monodisperse Aerosol	383
Monitoring Equipment.....	385
Cold and Hot Testing	385
Aerosol Mixing	386
Event Detection Future Work	386
Analytics	387
CHAPTER ELEVEN: CONCLUSIONS.....	389
REFERENCES	392

LIST OF TABLES

Table 1	WAQM node sensors and their key parameters.	27
Table 2	Operational amplifier selection parameters for the WAQM OPC. The top four choices are displayed, with positives in green, negatives in yellow, and strong negatives in orange.....	49
Table 3	RMS noise values for the baseline transimpedance amplifier in volts. ...	63
Table 4	RMS noise values for the bootstrap transimpedance amplifier in volts. ..	87
Table 5	Particle counter approximate component cost.	111
Table 6	Root mean square error comparison of the single sensor and orthogonal sensor arrangements using the PMS device providing the expected value. The orthogonal sensor arrangement outperforms the single sensor arrangement for large and small particle categories.	238
Table 7	Variance comparison of the single sensor and orthogonal sensor arrangements. The orthogonal sensor arrangement outperforms the single sensor arrangement in most cases, with the exception of small particles for the group size of 0.3 μm	239
Table 8	Particles per Liter limits for coincidence errors between 10 and 50 percent for four different particle sizes.....	263
Table 9	Event trigger condition types.	292
Table 10	Controlled environment testing sensor configuration.....	306
Table 11	Two-dimensional testing sequence of events.	309
Table 12	Three-dimensional testing sequence of events.	320

LIST OF FIGURES

Figure 1	Sensor system diagram. Two wireless sensor mesh networks are shown, each connected to a server through a different connection path to the Internet. WSN-A connects to the Internet through Wi-Fi and a local access point. WSN-B connects to the Internet through a bridge device using a serial connection. Visualization devices also attached to the Internet can access sensor data and send commands to the WSN systems through the server.	19
Figure 2	BSU Wireless Air Quality Monitor node. The monitor measures 150x150x115 mm and weighs 700g.	20
Figure 3	WAQM sensor node system diagram. Major system components are shown, including the sensors, microcontroller, and communication ports. Coordinator nodes may also optionally support Wi-Fi or cellular modem modules.	21
Figure 4	WAQM node firmware architecture diagram. A layered, modular architecture is used to facilitate rapid component changes and multiple developer environments.	29
Figure 5	Sensor system data flow. Data at each sensor node typically flow from the sensor task, through the data manager input buffer list, and onto the tag list of each data manager target.	34
Figure 6	BSU Sensor Monitor application, showing temperature data from three different sensor nodes simultaneously.	42
Figure 7	Analog front end block diagram. The circuit consists of a photodiode, transimpedance amplifier, a pair of amplifiers feeding the small and large channels, and an analog to digital convertor. The analog front-end is powered by a dedicated linear regulator to reduce system noise.	46
Figure 8	Circuit diagram for the baseline analog front end design. The photodiode current is amplified by a typical transimpedance amplifier arrangement. The output of the transimpedance amplifier feeds high-gain and low-gain second stage amplifiers to produce the small and large channel outputs respectively.	50

Figure 9	Photodiode capacitance as a function of reverse bias. The red line marks the point where the reverse bias equals 1.65V.....	52
Figure 10	Current to voltage gain model for the baseline transimpedance amplifier. The photodiode is broken down into a current source and capacitance C_D . The amplifier's input common mode and differential capacitances are C_{icm} and C_{id} , respectively. The output voltage e_o is developed from the voltage across the amplifier inputs e_o/AOL , which is the output voltage divided by the open loop gain of the amplifier.....	54
Figure 11	Noise analysis model for the baseline transimpedance amplifier. The amplifier's input current and voltage noise sources are shown as i_{ni} and e_{ni} , respectively. The feedback resistor's noise contribution is shown as e_{nR} . The amplifier's common mode and differential input capacitances are shown as C_{icm} and C_{id} . The photodiode is broken down into current, resistance, and capacitance components with shot noise and thermal noise i_{ns} and i_{nr} , respectively.....	55
Figure 12	Transimpedance amplifier gain for the baseline circuit. The open loop gain AOL for the OPA2314 is shown in dark blue, the gain demand $1/\beta$ is shown in green, and the noise gain A_{ne} is shown in red. The cyan line shows the current to voltage gain curve, which drops at 40dB/decade. ...	58
Figure 13	Noise density plot for the baseline transimpedance amplifier circuit. The open loop gain scaled by the noise floor density is shown in blue. The gain demand scaled by the noise floor density is shown in green. The noise density for the circuit is shown in red.....	60
Figure 14	Second stage amplifier gain for the small channel, high-gain amplifier including the DC blocking capacitor separating the transimpedance amplifier and the second stage. The open loop gain of the amplifier is shown in blue, the DC blocking capacitor gain is shown in green, and the small channel amplifier gain is shown in red.....	66
Figure 15	Power line noise seen on the small particle concentration plot for three WAQM OPC units. Unit 74 in blue has a high particle detection threshold, making it more immune to this noise. Units 72 and 75 in red and green have lower thresholds.	69
Figure 16	Power line noise on the device and at the AC power connection. The yellow trace is connected to the +3.3V power rail of the WAQM OPC and the blue trace is connected to the hot side of AC power where the wall transformer is connected. The oscilloscope ground is not attached to the WAQM OPC, which is powered by a two-prong wall transformer. Noise likely caused by an SCR is visible on the rising edge of both waveforms.	70

Figure 17	Baseline and bootstrap circuit interaction with noise coupling. The noise current i_n develops a voltage e_n across the feedback impedances. In the case of the baseline circuit, only the inverting input of the amplifier sees the voltage. The bootstrap circuit sees the voltage on both inputs, converting the noise to a common mode signal..... 72
Figure 18	Testing noise coupling on a grounded copper sheet. The WAQM OPC unit is powered through a two-prong wall transformer with isolated ground.. 74
Figure 19	Impact of bootstrap circuit on noise coupling. The four waveforms are taken from trace data captured by the WAQM OPC analog to digital converter using a sample rate of 4k samples per second. At top-left is the baseline circuit, at top-right is the bootstrap circuit with $R_{F2} = 0.1 \times R_{F1}$, at bottom-left is the bootstrap circuit with $R_{F2} = 0.5 \times R_{F1}$, and at bottom-right is the final bootstrap circuit with $R_{F2} = R_{F1}$ 75
Figure 20	Power line noise impact on the small particle concentration plot for three WAQM OPC units. Unit 72 now has the bootstrap circuit and no longer shows susceptibility to power line noise..... 76
Figure 21	Circuit diagram for the bootstrap analog front end design. The photodiode is connected across the transimpedance amplifier inputs. The non-inverting input of the transimpedance amplifier is biased to $\frac{1}{2} V_{CC}$ by a pair of $10M\Omega$ resistors, which in parallel equal the impedance of the $5.0M\Omega$ feedback resistor..... 78
Figure 22	Current to voltage gain model for the bootstrap transimpedance amplifier. The photodiode is broken down into a current source and capacitance C_D . The amplifier's input common mode and differential capacitances are C_{icm} and C_{id} , respectively. The output voltage e_o is developed from the voltage across the amplifier inputs e_o/AOL , which is the output voltage divided by the open loop gain of the amplifier. 80
Figure 23	Noise analysis model for the bootstrap transimpedance amplifier. The amplifier's input current and voltage noise sources are shown as i_{ni} and e_{ni} , respectively. The feedback resistor's noise contribution is shown as e_{nR} . The amplifier's common mode and differential input capacitances are shown as C_{icm} and C_{id} . The photodiode is represented as a capacitance only, since it is unbiased and carries no dark current. 82
Figure 24	Transimpedance amplifier gain for the bootstrap circuit. The open loop gain AOL for the OPA2314 is shown in dark blue, the gain demand $1/\beta$ is shown in green, and the noise gain A_{ne} is shown in red. The cyan line shows the current to voltage gain curve, which drops at 40dB/decade. ... 84

Figure 25	Noise density plot for the bootstrap transimpedance amplifier circuit. The open loop gain scaled by the noise floor density is shown in blue. The gain demand scaled by the noise floor density is shown in green. The noise density for the circuit is shown in red. 86
Figure 26	Block diagram of the sampling process. The AFE signals are multiplexed into the ADC, which stores digital values in the conversion register after being triggered by the timer. The DMA then moves the digital values into system memory. 89
Figure 27	Laser modules evaluated during WAQM OPC development. At top are three of the modules considered for prototype development. At bottom is a dissected module of the type deployed in the WAQM OPC. 93
Figure 28	Example laser beam profile as seen at the laser beam waist. Over 80% of the measurable beam power is within a 10 μm diameter spot. Note that the power measurement is uncalibrated to attenuation on the front-end laser beam profiler. 98
Figure 29	Laser module input current and output power plotted against supply voltage. Two different modules are shown, Apinex #41 and Apinex #13, each producing a different level of output power. 101
Figure 30	Example laser output noise. Normal amplifier output with a particle induced pulse at left. Transient laser output noise at right. Each case was measured with the same oscilloscope set to 100 mV per division. 102
Figure 31	WAQM OPC Flow Channel Design. The flow channel module base is shown in brown. The laser turret holds the laser, which shines through the light blocks and into the light trap. The air flow is shown in blue, starting at the inlet and passing by the laser beam in the flow channel. The grey pieces are separate parts which can be replaced with different designs supporting other configurations. The laser turret can rotate, allowing fine positioning of the laser beam. 103
Figure 32	Light blocks for the WAQM OPC light path. The light block on the left has a cylindrical hole and the light block on the right has a conical hole. 107
Figure 33	Laser illumination in the light trap of the WAQM OPC flow channel. Light emitted from the laser module at far right passes through the light blocks and crosses the flow channel to impact the deflecting wall of the light trap. The light trap is angled to prevent most of the light from reflecting back into the flow channel. 109

Figure 34	General simulation flow diagram. A set of particles is created and then iterated through to find the power received at each segment of the sensors.	115
Figure 35	Simulation coordinate system. Three-dimensional Cartesian coordinates are used, with (0, 0, 0) located at the laser beam waist. X is in the direction of airflow, or particle movement. Y is in the direction opposite the normal vector into the surface of the floor of the flow channel. Z is in the direction of propagation of laser light.	116
Figure 36	The surface of each sensor is divided into an m by n array of sub-rectangles. The center of each sub-rectangle will be the terminating endpoint of a scattered light vector from each particle.	118
Figure 37	Convergence of 10 μm diameter particle simulation as the number of sub-rectangles on the sensor surface increases. There are large differences in results as the sensor grid points increase from 5 x 5 through 30 x 30. Beyond this point, increases in the number of points show minimal improvement in accuracy.	119
Figure 38	Scattering vectors to two different sensors from a single particle. Each sensor is segmented into a 3 x 3 array of sub-rectangles. The blue vectors define the path to each corner of each sub-rectangle. The red vectors define the path to the center of each sub-rectangle.	121
Figure 39	Laser irradiance profile for the WAQM OPC. The irradiance peaks at the beam waist at the center of the flow channel and falls off to either side.	124
Figure 40	Mie scattering framework, with incident light propagating in the positive z direction. The scattering plane is defined by the incident light and scattered light vectors, with ϕ the azimuthal angle and θ the elevation angle of the scattering vector.	126
Figure 41	Translation from scattering polarization components V_r and V_l into polarization components relative to the surface of the sensor V_p and V_s	130
Figure 42	Transmittance through the sensor surface by polarization component of the incident scattered light over the angle of incidence. S_{rel} is the directional characteristic of the sensor, and the adjusted transmittances are T_p and T_s scaled by S_{rel}	134
Figure 43	Solid angle calculation for the light intercepting a single rectangular sensor segment. Each rectangular segment is split into a pair of triangles, which map to a solid area on the surface of the sphere surrounding the particle.	135

Figure 44	Mie angular scattering irradiance across 180 degrees of polar angle θ , calculated for a particle of water with diameter of 0.525 μm illuminated with 525 nm light. The red and blue lines show scattering irradiance of light polarized parallel and perpendicular to the scattering plane respectively.	137
Figure 45	Linear plot of Mie angular scattering irradiance in polar form for a particle of water with diameter of 0.525 μm illuminated with 525 nm light. Scattering perpendicular to the scattering plane is shown in the range from 0-180 degrees. Scattering parallel to the scattering plane is shown in the range from 360-180 degrees. The red line magnifies by 10 the data shown in the blue line to show some of the detail of the backscattered lobes. ...	138
Figure 46	Logarithmic plot of Mie angular scattering irradiance in polar form for a particle of water with diameter of 0.525 μm illuminated with 525 nm light. Scattering perpendicular to the scattering plane is shown in the range from 0-180 degrees. Scattering parallel to the scattering plane is shown in the range from 360-180 degrees.	139
Figure 47	Logarithmic plot of Mie angular scattering irradiance in polar form for a particle of water with diameter of 0.621 μm illuminated with 650nm red light as used in the WAQM system.	141
Figure 48	Physical arrangement of the sensor irradiance simulations. The sensor is located 2.0 mm directly below the center of the beam waist. Laser light may be polarized either parallel or perpendicular to the surface of the sensor as shown.....	142
Figure 49	Side view of the single photodiode arrangement, showing a cut-away of the center line of the sensor. The 2 mm distance from the beam waist to the surface of the photodiode captures a 67 degree section of the scattered light.	143
Figure 50	Normalized angular scattering from a 0.3 μm diameter PSL particle presented in polar form on a logarithmic scale. The range of angles from 0–180 degrees represents light polarized perpendicularly to the scattering plane, while the range of angles from 180–360 degrees represents light polarized parallel to the scattering plane.	144
Figure 51	Simulation of received power across a sensor from a 0.3 μm diameter PSL particle passing through the laser beam waist with polarization parallel to the sensor surface, plotted on a logarithmic scale. In this simulation, the beam waist is positioned over the center of the sensor at 2.0 mm from the surface. The scattering pattern produces a single maximum on the sensor surface.	145

Figure 52	Simulation of received power across a sensor from a 0.3 μm diameter particle passing through the laser beam waist with polarization perpendicular to the sensor surface, plotted on a logarithmic scale. In this simulation the beam waist is positioned over the center of the sensor at 2.0 mm from the surface. The scattering pattern shows a maximum toward the forward scattering direction and a single deep minimum. 148
Figure 53	Normalized angular scattering from a 2.5 μm diameter PSL particle presented in polar form. The range of angles from 0–180 degrees represents light polarized perpendicularly to the scattering plane, while the range of angles from 180-360 degrees represents light polarized parallel to the scattering plane. 150
Figure 54	Simulation of received power across a sensor from a 2.5 μm diameter PSL particle passing through the laser beam waist with polarization parallel to the sensor surface, plotted on a logarithmic scale. The beam waist is positioned over the center of the sensor at 2.0 mm from the surface. The scattering pattern produces multiple maxima on the sensor surface. 151
Figure 55	Simulation of received power across a sensor from a 2.5 μm diameter particle passing through the laser beam waist with polarization perpendicular to the sensor surface, plotted on a logarithmic scale. In this simulation, the beam waist is positioned over the center of the sensor at 2.0 mm from the surface. The scattering pattern shows a significantly different signature compared to parallel polarization. 152
Figure 56	Power received at the WAQM photodiode sensor for a set of particle sizes, plotted as a function of the particle position along the z-axis of the coordinate system. The laser is oriented with polarization parallel to the surface of the sensor. The particle intercepts the center of the beam, $y = 0$, for any given z-axis position. 154
Figure 57	Normalized angular scattering from a 0.4 μm diameter PSL particle presented in polar form on a logarithmic scale. The minimum at 120 degrees results in a lower intensity backscattering profile for incident light polarized parallel to the sensor surface. 155
Figure 58	Power received at the WAQM photodiode sensor for a set of particle sizes, plotted as a function of the particle position along the z-axis of the coordinate system. The laser is oriented with polarization perpendicular to the surface of the sensor. The particle intercepts the center of the beam, $y = 0$, for any given z-axis position. 157
Figure 59	Normalized angular scattering from a 0.1 μm diameter PSL particle presented in polar form on a logarithmic scale. The range of angles from 0–180 degrees represents light polarized perpendicularly to the scattering

	plane, while the range of angles from 180-360 degrees represents light polarized parallel to the scattering plane.	158
Figure 60	Configurations for dual-sensor optical particle counters. At left is the tandem sensor arrangement with the two sensors side by side. At right is the orthogonal sensor arrangement with one sensor horizontal and the other vertical relative to the floor of the flow channel. In each case, the beam waist remains over the center of the channel.	159
Figure 61	Side view of the tandem photodiode arrangement, showing a cut-away of the center line of the sensors. The 2 mm distance from the beam waist to the plane of the photodiode surfaces results in capturing 40 degree sections of the scattered light.	161
Figure 62	Tandem sensor simulation of received power across a sensor from a 0.3 μm diameter PSL particle passing through the laser beam waist with polarization parallel to the sensor surface, plotted on a logarithmic scale. The plot on the left is the upstream sensor and the plot on the right is the downstream sensor.	163
Figure 63	Tandem sensor simulation of received power across a sensor from a 2.5 μm diameter PSL particle passing through the laser beam waist with polarization parallel to the sensor surface, plotted on a logarithmic scale. The plot on the left is the upstream sensor and the plot on the right is the downstream sensor.	164
Figure 64	Power received at the tandem photodiode sensors for a set of particle sizes, plotted as a function of the particle position along the z-axis of the coordinate system. The laser is oriented with polarization perpendicular to the surface of the sensors. The particle intercepts the center of the beam, $y = 0$, for any given z-axis position. Solid lines represent the power received by the upstream sensor, and dashed lines represent power received by the downstream sensor.	165
Figure 65	Tandem arrangement power ratio received at the upstream sensor to the downstream sensor plotted as a function of the particle position along the z-axis of the coordinate system, on a logarithmic scale. The laser is oriented with polarization perpendicular to the surface of the sensors. The particle intercepts the center of the beam, $y = 0$, for any given z-axis position.	167
Figure 66	Top view of the orthogonal photodiode arrangement, showing the positions of the horizontal and vertical sensors. The vertical photodiode surface is 6 mm from the beam waist at 45 degrees from the direction of incident light propagation, resulting in the capture of a 25 degree section	

	of the scattered light. The horizontal photodiode has identical position to that of the single sensor arrangement.....	168
Figure 67	Orthogonal sensor simulation of received power across a sensor from a 0.3 μm diameter PSL particle passing through the laser beam waist with incident light polarization parallel to the sensor surface, plotted on a logarithmic scale. The plot on the left is the horizontal sensor and the plot on the right is the vertical sensor.	170
Figure 68	Orthogonal sensor simulation of received power across a sensor from a 2.5 μm diameter PSL particle passing through the laser beam waist with incident light polarization parallel to the sensor surface, plotted on a logarithmic scale. The plot on the left is the horizontal sensor and the plot on the right is the vertical sensor.	171
Figure 69	Power received at the orthogonal photodiode sensors for a set of particle sizes, plotted as a function of the particle position along the z-axis of the coordinate system. The laser is oriented with polarization perpendicular to the surface of the horizontal sensors. The particle intercepts the center of the beam, $y = 0$, for any given z-axis position.	173
Figure 70	Orthogonal arrangement power ratio received at the horizontal sensor to the vertical sensor plotted as a function of the particle position along the z-axis of the coordinate system, on a logarithmic scale. The laser is oriented with polarization perpendicular to the surface of the sensors. The particle intercepts the center of the beam, $y = 0$, for any given z-axis position. .	174
Figure 71	Example operational simulation of 1 μm PSL particles passing through the channel at random locations for the orthogonal sensor arrangement. Received power from each particle at each sensor has been converted into an ADC level, with the horizontal sensor on the horizontal axis and the vertical sensor on the vertical axis.	176
Figure 72	Comparison of PSL (left) and average rural air pollution (right), showing the received power from the respective particle types against the position in the flow channel. The solid lines represent the responses from the horizontal sensor, and the dashed lines represent the responses from the vertical sensor. The two material types result in very similar responses, though the magnitude response is reduced for the rural air pollution.....	178
Figure 73	Comparison of PSL (left) and rural air pollution (right) power ratios between the horizontal and vertical sensors on the orthogonal arrangement. The responses are very similar with the response for small particles being greater on the horizontal sensor than the vertical for much of the channel width.	179

Figure 74	Comparison of PSL (left) and soot (right), showing the received power from the respective particle types against the position in the flow channel. The solid lines represent the responses from the horizontal sensor, and the dashed lines represent the responses from the vertical sensor. The two material types result in somewhat different responses depending on the particle size. The peak response values for the beam waist are in general reduced for the soot particles.	180
Figure 75	Comparison of PSL (left) and soot (right) power ratios between the horizontal and vertical sensors on the orthogonal arrangement. The responses differ significantly for larger particles, with the response ratio for 5 μm particles being the highest of the group.	181
Figure 76	Mie intensity plotted against particle size for three different laser wavelengths at the beam waist, assuming that the intensity at the beam waist is constant across the three laser sources. The intensity is plotted assuming a 7 mm^2 photodiode sensor placed 2 mm under the beam waist. The 405 nm laser appears to offer some advantage at the smallest intensities since the transition into the Rayleigh regime occurs at smaller particle sizes for shorter wavelengths.	183
Figure 77	Mie intensity plotted against particle size for three different laser wavelengths at the beam waist, assuming that the intensity at the beam waist scales with wavelength. The intensity is plotted assuming a 7 mm^2 photodiode sensor placed 2 mm under the beam waist.	184
Figure 78	Relative spectral sensitivity curve for the photodiode used in the WAQM OPC design, reproduced from the Osram BPW34 datasheet. The photodiode response peaks at 850 nm.	185
Figure 79	Mie intensity including spectral sensitivity scaling plotted against particle size for three different laser wavelengths at the beam waist, assuming that the intensity at the beam waist scales with wavelength. The intensity is plotted assuming a 7 mm^2 photodiode sensor placed 2 mm under the beam waist.	186
Figure 80	Magnitude and phase responses for infinite impulse response filters with a range of scaling factors. Solid lines represent the magnitude responses and dashed lines the phase responses. The blue line represents a scaling factor of 256, which is the implemented value for the WAQM OPC. The -6dB cutoff frequency for this filter is just over 100 Hz.	195
Figure 81	Particle trace data captured by a WAQM OPC unit. The graph plots the ADC value against the 10 μs samples. The red line is the baseline, blue is the small channel output, and green is the large channel output. Note that the baseline only changes slightly due to the particle pulse.	197

Figure 82	An arbitrary intensity threshold drawn on the single-sensor arrangement received power curves across the flow channel. The threshold intersects the curves for multiple particle sizes. A very small particle intercepting the laser beam near the waist may produce enough scattered irradiance to cross the threshold into the large space. Conversely, large particles passing through the beam near the edge of the channel may not..... 199
Figure 83	Sensing area plotted on a logarithmic scale against particle size. Since the amount of light scattered by a particle is dependent on its size, and since the incident beam irradiance varies across the channel with the focused beam, the effective sensing area will also vary with particle size. The small channel lines denote the detection threshold for the single sensor arrangement WAQM OPC, and the large channel lines denote the typical large particle threshold..... 201
Figure 84	Air velocity along the centerline of the laser across the width of the flow channel. The air velocity is relatively flat across the center portion of the channel, but falls off on each side near the edges..... 203
Figure 85	Durations for several particle sizes for the WAQM OPC single sensor arrangement, plotted against position across the flow channel. This is the duration seen by a particle traversing the beam in the $y=0$ plane, which is the highest intensity portion of the laser beam for any given z -axis position..... 204
Figure 86	Large ($>1 \mu\text{m}$) particle count comparison between single sensor arrangement WAQM OPC and PMS device data. The WAQM OPC is using intensity from the single sensor to size particles. Several different PSL particles sizes are used, which are listed in magenta along the bottom of the graph. There is significant over-counting for particles in the $0.6 \mu\text{m}$ range due to the wide range of intensities produced by monodisperse particles in the focused laser beam of the WAQM OPC. 206
Figure 87	Particle intensity over time for the example PSL particle test, plotted as the ADC level from the horizontal sensor large channel. Test periods for $0.3 \mu\text{m}$, $0.6 \mu\text{m}$, $0.8 \mu\text{m}$, $1.0 \mu\text{m}$, and $1.6 \mu\text{m}$ are labeled in white. More high-intensity particles are generated as the size of the particles grows. 207
Figure 88	Large ($>1\mu\text{m}$) particle count comparison between single sensor arrangement WAQM OPC and PMS device data, using intensity and duration for WAQM OPC size classification. The duration limitation significantly reduces the number of particles that are available to count. Compensating by dropping the intensity threshold greatly increases the over counting of particles that are below the desired size threshold, as is the case with the $0.6\mu\text{m}$ test particles. 209

Figure 89	Small (<1 μm) particle count comparison between single sensor arrangement WAQM OPC and PMS device data, using pulse height only for WAQM OPC size classification. While the 0.3 μm particles fit well, there is significant over counting for larger particles. 211
Figure 90	Small (<1 μm) particle count comparison between single sensor arrangement WAQM OPC and PMS device data, using pulse height and duration for WAQM OPC size classification. While there is still significant over counting for particles larger than 0.3 μm , the level is reduced from that of sizing using pulse height alone. 212
Figure 91	Particle trace data captured by a WAQM OPC unit with orthogonal sensor arrangement. The graph plots the ADC value as a sequence of samples. The red, blue, and green lines are the baseline, small channel, and large channel for the horizontal sensor, respectively. The yellow, cyan, and magenta lines are the baseline, small channel, and large channel for the vertical sensor. The data from the large channels are captured and shown even though they are not necessary for the implementation..... 217
Figure 92	Particle sizing with the orthogonal sensor arrangement. Using the pulse duration and pulse height ratio, the beam can be conceptually divided into three distinct sections. The narrow center waist is most useful for small particle detection. The wider sections to the sides, upstream near the laser source and downstream near the vertical sensor, are used to detect and size larger particles..... 218
Figure 93	Comparison of orthogonal sensor arrangement simulated (left column) and experimental (right column) results for 0.3 μm (top row), 0.8 μm (middle row), and 1.6 μm (bottom row) PSL particles. The horizontal and vertical axes are the horizontal and vertical sensor ADC levels, respectively. The simulations assume a monodisperse aerosol, which is not possible to achieve with the actual experiments. The general trends in distribution do match according to particle size..... 221
Figure 94	Large (>1 μm) particle count comparison between orthogonal sensor arrangement WAQM OPC and PMS device data. The WAQM OPC uses the intensity ratio between the two sensors and pulse duration to size particles. Several different PSL particles sizes are used, which are listed in magenta along the bottom of the graph..... 225
Figure 95	Medium (0.5 μm to 1 μm) particle count comparison between orthogonal sensor arrangement WAQM OPC and PMS device data. The WAQM OPC uses the intensity ratio between the two sensors and pulse duration to size particles. Several different PSL particles sizes are used, which are listed in magenta along the bottom of the graph. There is significant over-counting for particles in the 0.8 μm to 1.6 μm range. 227

Figure 96	Medium (0.5 μm to 1 μm) particle count comparison between orthogonal sensor arrangement WAQM OPC and PMS device data. The WAQM OPC uses the intensity ratio between the two sensors and pulse duration to size particles, with correction for large particle count. The over-counting for particles in the 0.8 μm to 1.6 μm range is greatly reduced. 228
Figure 97	Small (<0.5 μm) particle count comparison between orthogonal sensor arrangement WAQM OPC and PMS device data. There is significant over-counting for particles in the 1.0 μm to 1.6 μm range due to the intensity ratio overlap between particles in this size range and true small particles. 229
Figure 98	Scatter plot showing orthogonal arrangement particles received by the WAQM OPC. The horizontal axis is the ADC level for the horizontal sensor, and the vertical axis is the ADC level for the vertical sensor. Each particle is plotted as a small circle according to its horizontal and vertical ADC level. Particle colors indicate sorting into various groups. Red, green, and blue represent large, medium, and small particles, respectively. Black, magenta, cyan, and yellow represent particles that are rejected due to high horizontal intensity, low vertical intensity, high intensity threshold, and duration, respectively. 231
Figure 99	Three-dimensional scatter plot showing orthogonal arrangement particles received by the WAQM OPC. The ADC levels for the horizontal and vertical sensors are plotted against the vertical sensor duration in samples. This representation shows the distribution of particle durations, and the relationship between duration and large, medium, and small particles. . 235
Figure 100	Classification scatter plots for 0.3 μm (top row), 0.8 μm (middle row), and 1.6 μm (bottom row) PSL particles. The horizontal and vertical axes of the plots on the left are the horizontal and vertical sensor ADC levels, respectively. The plots on the right are for the same data as on the left, rotated to view vertical sensor duration on the third axis. 237
Figure 101	Oscilloscope trace of amplifier output for a large, low velocity particle intercepting the laser across multiple peaks on the orthogonal sensor arrangement. The yellow trace is from the horizontal sensor and the green trace is from the vertical sensor. 240
Figure 102	Flowchart for the single sensor sizing classifier. Only three basic decisions are made including intensity and duration of the pulse. The intensity differentiates the particle as either large or small at decision point 1, and then the particle may be rejected based on the pulse duration being too long. The duration used as a threshold in decision point 2 for small particles is typically shorter than that for decision point 3 for large. 242

Figure 103	Flowchart for the dual sensor sizing classifier used with the orthogonal sensor arrangement. The decision tree is more complex than that for the single sensor due to the additional data available from the second sensor. Particles are sized based on the horizontal to vertical sensor intensity ratio at decision point 2 and the vertical sensor duration at decision point 6. The remaining decision points filter particles to refine the counts. 244
Figure 104	Inertial deposition in a bend in the WAQM OPC flow channel. This bend occurs after the photodiode, but illustrates particulate impaction when the air velocity changes direction. The particulate matter builds up in the channel after running continuously for many weeks. 255
Figure 105	Measurement of digitization error for a typical small particle, taken from the small channel output. At left is the full pulse; at right is a zoomed in view of the peak. The span of the trace at right is 10 μ sec, or the time between two samples at 100 ksp. The trace is centered on the peak of the pulse, with the horizontal markers bounding the difference between peak and sampled voltage. The difference is just over 50 mV, or approximately 16 out of 1024 ADC levels. 261
Figure 106	Summary of WAQM OPC error sources and their impacts on sensitivity and absolute particle count. While the digital sampling error also impacts sensitivity and absolute particle count, it is treated as uncertainty error since compensation for it cannot be achieved through calibration. 265
Figure 107	Sensitivity differences between two units and the resulting impact on absolute particle count. The test aerosol generated between 10 and 35 minutes is composed of smaller particles than the aerosol generated between 40 and 60 minutes. The two units agree on the larger particle count, but the more sensitive Unit 60 detects and counts a larger number of smaller particles. 267
Figure 108	Particle generation system developed for the testing of the WAQM OPC. The system is fed by the building air supply, generating a particulate matter aerosol using a Collison Nebulizer as shown. Options exist for heating the aerosol, drying the aerosol, or taking the aerosol directly into the test chamber. The particulate matter is filtered through HEPA filters as air exits the test chamber..... 268
Figure 109	View of the OPC calibration system showing the air supply, source dryer and filter, nebulizer, and test chamber inlet. 270
Figure 110	View of the OPC calibration system showing the heater and diffusion drier. 272

Figure 111	Relative sensitivity and scale factor plotted for Unit 60 against the trigger set point when calibrated to match Unit 72. The relative sensitivity falls past zero when the trigger for Unit 60 moves past 25 ADC levels. The scale factor is approximately 1.75 at this trigger, meaning the particle count for Unit 60 must be multiplied by this value to normalize the two units.....	278
Figure 112	Scaled results for the two units after calibrating sensitivities. The two units now agree closely for cold and hot particle counts.....	279
Figure 113	Calibration results from a run of 9 units with a TSI Aerotrak 9303 used as the golden standard. The plot on the left shows the absolute particle counts of the 9 WAQM OPC units. The plot on the right shows the calibrated, scaled results. A diverse set of sensitivities and absolute particle counts can be corrected to perform similarly to the TSI unit.....	280
Figure 114	Particle size distribution for a set of PSL particles generated by the OPC calibration system as measured by the PMS device. Concentration is plotted on a logarithmic scale over time for each of the six different channels of the PMS device. The particle size tested is shown in red along the time axis. The shift in size distribution can be seen among the different channels as the tested particle size increases, however many particles of sizes outside the desired range are detected as well.	284
Figure 115	Relative sensitivity and scale factor plotted against the trigger set point for a single sensor arrangement WAQM OPC unit when calibrated to match the PMS device for particles larger than 1.0 μm . The relative sensitivity for the large channel crosses zero when the trigger for the unit moves past 128 ADC levels. The scale factor is approximately 22.8 at this trigger, meaning the particle count for the unit must be multiplied by this value to normalize the two units.....	286
Figure 116	Node trigger structure. Each node has a set of triggers. Each trigger is comprised of one or more conditions that lead to an event. Transitions between conditions are configurable, allowing the creation of small state machines to reach an event.	291
Figure 117	Example event trigger. The trigger starts with a rising edge condition. A counting level condition ensures the concentration remains high. Once the level count expires, the event message is generated.....	293
Figure 118	Event notification dataflow. The unit detecting the event notifies the server, which then generates an email to the user. The user may then directly view the relevant data by following a link to the server, which will return a formatted view of the event.	297

Figure 119	Example event data web page returned from the server from an event notification email hyperlink. The default format returns 1 hour of data before and after the event. The event is marked with a vertical line, and labeled with a type number.	298
Figure 120	Interior of the aircraft cabin mockup with WAQM sensors in place. The blue LEDs on the sensors are visible on the seat headrests.	304
Figure 121	KSU aircraft cabin particulate matter generation. At left is the particle size distribution for smoke created with the fog generator. At right is the distribution for the powder generator. These data were taken from a pair of test runs using a model 3321 Aerodynamic Particle Sizer from TSI Inc.	306
Figure 122	WAQM sensor node layout for two-dimensional test.	308
Figure 123	Two-dimensional test particulate matter concentration in particles per liter for all sixteen sensor nodes on a logarithmic scale.	310
Figure 124	Two-dimensional test CO ₂ concentration in parts per million for all sixteen sensor nodes on a linear scale.	311
Figure 125	Two-dimensional tests with particulate and CO ₂ contour plots just prior to Run 1.	313
Figure 126	Two-dimensional tests with particulate and CO ₂ contour plots at the beginning of substance dispersal.	313
Figure 127	Two-dimensional tests with particulate and CO ₂ contour at the peak of substance dispersal.	314
Figure 128	Two-dimensional tests with particulate and CO ₂ contours as the particulate dilutes.	315
Figure 129	Two-dimensional tests with particulate and CO ₂ contours showing human activity in the aircraft cabin.	316
Figure 130	Two-dimensional testing humidity profile. The dehumidifier was turned off and humidifier turned on between Run 2 and Run 3.	316
Figure 131	Two-dimensional testing with the dehumidifier off and the humidifier on.	317
Figure 132	Two-dimensional testing temperature and pressure profiles.	318
Figure 133	WAQM Sensor Node Layout for Three-Dimensional Test.	319

Figure 134	Aircraft cabin with three-dimensional setup in place.	320
Figure 135	Three-dimensional test particulate matter concentration in particles per liter for all sixteen sensor nodes on a logarithmic scale.	321
Figure 136	Three-dimensional test CO ₂ concentration in parts per million for all sixteen sensor nodes on a linear scale.	322
Figure 137	Three-dimensional smoke and CO ₂ test with particulate and CO ₂ contour plots just prior to run 1.	324
Figure 138	Three-dimensional smoke and CO ₂ test with particulate and CO ₂ contour plots at the beginning of substance injection.	325
Figure 139	Three-dimensional smoke and CO ₂ test with particulate and CO ₂ contour plots as substance injection continues.	326
Figure 140	Three-dimensional smoke and CO ₂ test with particulate and CO ₂ contour plots as CO ₂ begins to appear.	326
Figure 141	Three-dimensional smoke and CO ₂ test with particulate and CO ₂ contour plots at the height of substance injection.	327
Figure 142	Three-dimensional smoke and CO ₂ test with particulate and CO ₂ contour plots as substance injection tapers off.	328
Figure 143	Three-dimensional smoke and CO ₂ test with and CO ₂ contour plots near the end of substance injection.	329
Figure 144	Three-dimensional talcum powder test with particulate and CO ₂ contour plots at the start of substance injection.	331
Figure 145	Three-dimensional talcum powder test with particulate and CO ₂ contour plots as the substance spreads.	331
Figure 146	Three-dimensional talcum powder test with particulate and CO ₂ contour plots near the end of substance injection.	332
Figure 147	Three-dimensional testing humidity profile. The dehumidifier was turned off during run 4.	333
Figure 148	Three-dimensional testing temperature and pressure profiles.	334
Figure 149	Operating room air quality monitoring configuration. WAQM nodes are shown in green and the coordinator node is purple.	336

Figure 150	Data flow diagram for the WAQM OR deployment. Data are sent via TCP connection from the coordinator to the BSU server. Data are viewed over HTTPS using a web browser to connect to the BSU server. 338
Figure 151	Hospital ventilation system, showing the supply and return paths. A return air damper regulates the amount of fresh vs. recirculated air in the system. 340
Figure 152	Submicron particulate matter generated during the example surgical procedure starting at 08:33 and ending at 10:47. The three sensors in the OR (#67, #70, and #71) are exposed to increased particulate levels during the procedure. The sensors in the hallway (#60) and substerile room (#78) do not register much of an increase. 342
Figure 153	Particulate matter greater than 1 μm in size during the example surgical procedure starting at 08:33 and ending at 10:47. The concentration of large particulate matter is much less than that for submicron particulate, with peaks that do not generally coincide with the submicron peaks. Two of the highest amplitude peaks occur at the beginning and end of the procedure..... 343
Figure 154	Carbon dioxide measurement during the example surgical procedure starting at 08:33 and ending at 10:47. The deviations above average levels for the units inside the OR (#67, #70, and #71) are likely the result of exhalations from the patient and surgical staff. 344
Figure 155	Temperature measurement in Celsius during the example surgical procedure starting at 08:33 and ending at 10:47. Temperatures in the hallway (#60) and substerile room (#78) appear to track in similar fashion, as do the temperatures of the two nodes in the OR outside the laminar flow area (#70 and #71). The unit in the laminar flow area (#67) shows marked increase during the procedure, possibly due to heat produced by lighting, equipment, and/or human bodies in the area. 346
Figure 156	Pressure measurements in Pascals during the example surgical procedure starting at 08:33 and ending at 10:47. The data tend to follow atmospheric conditions rather than any significant activity inside the hospital. 347
Figure 157	Relative humidity measurement during the example surgical procedure starting at 08:33 and ending at 10:47. The downward trend of Unit 67 may be due to the rising temperature in the laminar flow area. The spikes in humidity after the end of the procedure for the units inside the OR (#70 and #71) may be due to cleaning activity. 348
Figure 158	Carbon monoxide measurement in PPM during the example surgical procedure starting at 08:33 and ending at 10:47. 349

Figure 159	Correlation between a BSU laboratory (red – unit 66) and the operating room. The general trends are very similar, suggesting that the background levels of submicron particulate are due to outdoor air quality.....	350
Figure 160	Particulate event propagation. This submicron particulate matter event appears to trigger all particulate matter sensors nearly simultaneously.	352
Figure 161	Zooming in on the submicron particulate matter event that has highest intensity in the hallway is inconclusive in determining the source due to the low, 1-minute sampling resolution of the particulate matter concentration samples.....	353
Figure 162	The particulate matter event seen with the raw submicron particle count output data.....	354
Figure 163	Humidifier valve position data for the operating room over a period of several days. The humidifier is active during the last half of May 11 th , and for several hours toward the end of May 13 th .	355
Figure 164	Relative humidity readings over the period corresponding with hospital humidifier valve data. The humidifier activity appears to start when the relative humidity is falling sharply.	356
Figure 165	Submicron particulate matter readings over the period corresponding with hospital humidifier valve data. The time periods corresponding to humidifier activity do not show increases in particulate matter.	357
Figure 166	False particles on WAQM node number 67, inside the laminar flow area of the OR.....	358
Figure 167	Submicron particulate matter in a residential home with two sensors. Unit 75 is placed in the garage, and Unit 76 is in the main living space. Large changes in particulate levels in the garage occur when the door is opened and vehicles enter or exit.	361
Figure 168	Particulate matter over 1 μm in size in a residential home. Unit 75 is placed in the garage, and Unit 76 is in the main living space. The large peak in garage particulate occurs when an automobile enters from outside.	362
Figure 169	Temperature measurements in a residential home. Unit 75 is placed in the garage, and Unit 76 is in the main living space. The opening of the garage door causes rapid temperature changes at 08:00, 16:20, and 18:00.....	363
Figure 170	Carbon dioxide levels in a residential home. Unit 75 is placed in the garage, and Unit 76 is in the main living space. CO ₂ levels are elevated	

	during the night with all occupants in the home. Levels in the home fall steadily during the course of the day while the occupants are gone for the day.....	364
Figure 171	Carbon monoxide levels in a residential home. Unit 75 is placed in the garage, and Unit 76 is in the main living space. CO levels in the garage spike when an automobile starts, then remain high when the garage door closes, trapping the CO. The level in the home can be seen to rise in response, as the trapped garage CO leaks into the home.....	365
Figure 172	Example laser beam profile taken from a WAQM OPC laser module. The elliptical profile of the beam is apparent, with the minor axis demonstrating a reasonable Gaussian fit and the major axis showing multiple peaks.	376
Figure 173	Simulation for 2.5 μm PSL particles with the orthogonal sensor arrangement. The resulting output pattern is similar to other particles larger than 1.0 μm	379

LIST OF ABBREVIATIONS

ACGIH	American Conference of Governmental Industrial Hygienists
ADC	Analog to Digital Convertor
AFE	Analog Front End
AQWSN	Air Quality Wireless Sensor Network
BSU	Boise State University
COPD	Chronic Obstructive Pulmonary Disease
DI	Deionized
DMA	Direct Memory Access
EPA	Environmental Protection Agency
ESD	Electrostatic Discharge
FAA	Federal Aviation Administration
GPIO	General Purpose Input-Output
HSIL	Hartman Systems Integration Laboratory
IIR	Infinite Impulse Response
I2C	Inter-Integrated Circuit
KSU	Kansas State University

LCD	Liquid Crystal Display
LSB	Least Significant Bit
NAAQS	National Ambient Air Quality Standards
NIH	National Institutes of Health
NIOSH	National Institutes for Occupational Safety and Health
OPC	Optical Particle Counter
OR	Operating Room
PID	Proportional-Integral-Derivative
PMS	Particle Measuring Systems
PSL	Polystyrene Latex
RTC	Real-Time Clock
SCR	Silicon Controlled Rectifier
SD	Secure Digital
SOC	System-On-Chip
TEOM	Tapered Element Oscillating Microbalance
WAQM	Wireless Air Quality monitor
WHO	World Health Organization
WSN	Wireless Sensor Network

CHAPTER ONE: INTRODUCTION

Wireless sensor network (WSN) research for environmental monitoring has dramatically increased in the last decade. The convergence of improvements in sensor, battery, and semiconductor technology has made the design and deployment of these networks feasible. Sensors designed to measure environmental parameters such as gas and vapor concentrations, pressure, humidity, and temperature are now readily available in small form factors and at reasonably affordable prices. Yet one key component in determining air quality is missing from this suite, a sensor designed to monitor the real-time concentration of airborne particulate matter.

Airborne particulate matter consists of chemically and physically diverse solid or liquid particles suspended in air. Particulate matter exists as discrete particles, and originates from a variety of natural and anthropogenic sources. The particles may be emitted directly from a primary source through a chemical or physical process, or may form from the transformation of secondary components such as sulfur oxides, nitrogen oxides, or volatile organic compounds [1]. Natural sources include bacteria, pollen, fungal spores, plant and animal debris, dust, and ash. Anthropogenic sources include combustion by-products from the burning of wood and fossil fuels, tobacco smoke, cooking exhaust, and cleaning activities [1]–[3].

Particulate matter has long been studied in an attempt to gain a better understanding of its effects on human health. A significant number of studies conducted

over the last several decades suggest connections between concentrations of ambient particulate matter and increases in morbidity and mortality, including potential connections with chronic obstructive pulmonary disease (COPD), asthma, lung cancer, and cardiovascular disease [4]–[6]. In the United States, it is estimated that 22,000 – 55,000 deaths per year are caused by particulate matter [7]. A recent study specific to California estimated annual premature particulate matter related deaths at 9,200 in that state alone[8]. The World Health Organization (WHO) estimates more than 2 million premature deaths worldwide are caused by indoor and outdoor air pollution [9]. Short-term increases in particulate matter concentration have been shown to correlate with increases of 1-8% in deaths per 50 $\mu\text{g}/\text{m}^3$ increase in outdoor particulate concentration [10]. Recent studies also suggest that reducing long-term exposure to particulate air pollution can increase life expectancy [11].

Particulate matter can also have a significant impact on property, plant-life, and industrial, experimental, and medical processes, and is much studied in terms of particle deposition onto sensitive surfaces. For example, particle deposition onto the reflective surfaces of sensitive optical equipment can degrade the images produced [12]. Particulate matter in semiconductor manufacturing can greatly impact process yield since circuit feature sizes can be much smaller than suspended particulate [13]. Elemental carbon particles in the form of soot can produce perceptible soiling of museum artwork over time frames that are relatively short in relation to the desired lifetime of the art [14]. Particulate in medical operating theatres has the potential for negative health impacts on the patient and the operating staff [15]. Particulate matter deposition can also impact the

vigor, competitive viability, and reproductive fitness of individual plants [16], potentially affecting agriculture and natural ecosystems.

In the United States, the Clean Air Act of 1970 authorized the establishment of the National Ambient Air Quality Standards (NAAQS). These standards for outdoor ambient air quality are promulgated by the United States Environmental Protection Agency (EPA), and, in most instances, primacy of compliance enforcement and air monitoring has been delegated to individual US States. In 1987, EPA expanded the NAAQSs and promulgated a new standard limiting exposure to particulate matter having a diameter less than 10 μm , termed PM_{10} . Particles of this size are commonly referred to as “thoracic particles” and, when inhaled, have the ability to penetrate to the tracheobronchial region of the lung. In 1997, evidence of serious human health effects prompted EPA to further subdivide its regulation of particulate matter to include a NAAQS limiting exposure to particles having diameters less than 2.5 μm , termed $\text{PM}_{2.5}$. This was due, in part, to the ability of these “respirable particles” to penetrate to the alveolar region of the lung and adversely affect gas-exchange during human respiration.

EPA does not currently regulate indoor air quality, though it does provide guidance in protecting indoor air from contaminants such as mold, radon, and formaldehyde. In an effort to protect the health of the US workforce, the American Conference of Governmental Industrial Hygienists (ACGIH), the Occupational Safety and Health Administration (OSHA), and the National Institutes for Occupational Safety and Health (NIOSH) have established sampling conventions, provided recommendations, and promulgated standards concerning exposure to particulate matter [17]–[19]. The American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE)

publishes standards for building ventilation with the intent of maintaining acceptable air quality [20], but these are not federally mandated in the U.S.

The impact of indoor air quality on human health is also significant given that most people in the U.S. spend 90% or more of their time indoors [21], resulting in a greater exposure to indoor particulate matter than direct exposure to outdoor levels [22]. When indoor sources of particles are not present, indoor levels of particulate matter tend to show similar behavior over time as outdoor levels. However, when indoor sources are active, the correlation no longer exists, making it impossible to estimate indoor levels from outdoor measurements [23]. In short, depending solely on outdoor measurements may be misleading in estimating personal exposure to particulate matter.

Given the importance of understanding personal exposure levels, it becomes interesting to explore what can be done to provide more of this information through sensor networks. While sensors for measuring real-time particulate matter concentrations have been available for decades, most of the commercially available devices have serious drawbacks for incorporation in a WSN. These devices tend to be relatively large, making them cumbersome during personal monitoring applications when they must be worn on a person's body. They are also intended for use as stand-alone units, and are generally not appropriate for integration within other compact devices. Most of the commercially available devices are also relatively expensive, with typical starting prices greater than two thousand U.S. dollars. These factors can make it difficult to include multiple devices into an environmental study, and prohibit inclusion of a particulate matter sensor in each wireless sensor node.

Motivations

The creation of an Air Quality WSN (AQWSN) that includes particulate matter sensors opens many possibilities for new applications that provide benefits spanning areas from health and wellness to material concerns. These applications provide motivation for the work behind this dissertation, as they underline the possibilities of what may be achieved given the availability of such an AQWSN. This section discusses some of these applications and the benefits that they may provide.

The assessment of personal exposure is perhaps one of the most compelling arguments for understanding the impact of particulate matter on individual human health. Most individuals have access only to outdoor air quality data provided by government agencies such as the EPA. While outdoor air quality is important, the general lack of correlation between it and personal exposure makes reliance on these data problematic. For individuals suffering from health problems such as chronic obstructive pulmonary disease (COPD), asthma, or severe allergies, the assessment and control of personal exposure to particulate matter can lead to improvements in quality of life and even increases in life expectancy. A suite of particulate matter sensors installed in a personal space has the potential to identify and assist in the control of particulate sources, leading to an improvement in air quality and a resultant reduction in human exposure.

Even for individuals without health problems, raising awareness of air quality and its correlation to personal exposure can have long-term benefits. Providing effective tools for visualization of air quality can lead to action that benefits the individual and those nearby [24]. A WSN monitoring air quality, along with the appropriate visualization hardware and software, has the potential to positively affect human behavior. For

example, a smoker may not fully understand the impact that this activity has on others. Providing information about the quantity of particulate generated from smoking could help alter the smoker's behavior to the benefit of those around them.

Many studies that include the measurement of personal exposure to particulate matter rely on single measurements within a space with the assumption that concentrations associated with the introduction of new particulate matter will quickly achieve equilibrium in ambient air. This assumption may not always hold true, since local concentrations of particulate matter can be significantly impacted by the activity of an individual [25] and can take minutes to disperse depending on airflow and topographic features in the room [26]. A WSN containing particulate matter sensors would enable the visualization of dynamic particulate matter movement within personal spaces, helping to track contaminants and identify locations where particulate matter concentrations might reach levels that could negatively impact health. For example, there has been recent interest in understanding diesel particulate matter exposures received by children on school buses [27]. Outfitting a school bus with such a WSN could indicate how particulate matter enters the bus and where it lingers, and sensor nodes positioned near the bus entrance could provide exposure levels for children waiting to board the bus. Thus, the information gained by application of the WSN could help identify countermeasures for controlling exposures to acceptable levels.

The investigation of airborne contaminant movement within the highly dynamic environments of aircraft cabins has been facilitated by the use of WSNs. Past studies have deployed WSNs containing suites of gas and other environmental sensors [28]. The addition of a particulate matter sensor to this suite would provide for the assessment of

another important environmental factor determining the air quality of an aircraft cabin [29]. For example, there is significant interest in the measurement of the microbial content of cabin air [30]. A particulate matter sensing WSN would aid in the understanding of the movement and concentrations of these and other particulates in the aircraft environment.

Electronic equipment failures can cause significant loss in productivity, which can multiply when not quickly discovered. Equipment failures causing partial oxidization of electronic components could be detected by particulate matter sensors deployed to a local area. A particulate matter sensor equipped WSN could potentially detect increases in particulate released by the oxidization of a component, providing an early failure warning. This could be particularly useful in large data centers commonly built for today's cloud computing needs. A WSN installed in such a facility could help pinpoint equipment failures and increase the uptime of electronic equipment.

The quest for energy efficient homes has led to improvements in insulation to minimize heat and air exchange with the outside environment. Unfortunately, this can have the undesired impact of not allowing fresh air to naturally infiltrate the building and dilute indoor pollutants. In these newer buildings, mechanical ventilation may be required to bring in fresh air [20]. Work has been done on intelligent systems that take advantage of the difference between indoor and outdoor air to cool a home [31]. Future smart homes may include the additional ability to detect air quality inside and outside the home and thus determine when to perform an air exchange or when to circulate indoor air through a filtering device. Affordable particulate matter and gas sensing WSNs can help to enable this vision.

Contribution

The work behind this dissertation strives to produce a reference system that can be used as a basis for many of the applications discussed in the previous section. The key contribution of the work is in its enablement of these new applications. The end goal is to make a significant impact on the quality of health and the protection of material value. This is achieved through the accessible, low-cost air quality monitoring framework that this dissertation describes.

Health and Material Impact

Many of the target applications have potential benefits related to the improvement of health and wellness, and for the understanding and mitigation of material impact. The work behind this dissertation will open the door for enabling these applications, which moves us toward the ability to live and work with cleaner air in our personal and industrial environments. The applications that follow will have significant potential to impact the lives of many who are now lacking the information about their personal air quality.

Accessibility through Lower Cost

In the past, understanding particulate matter concentrations at a local, personal level has typically been an expensive endeavor, putting many applications out of reach. The critical cost component restricting more widespread use is the optical particle counter (OPC). While this type of device has been well-studied and documented in the past, a low-cost version has not been created that can be included in wireless sensor networks. Many challenges exist in producing quality data using low-cost components. This

dissertation explores and proposes methods for correcting many of these issues. The expectation is not to produce an instrument that can replace a high-cost, high-quality scientific device, but to provide a level of quality that gives a respectable view into aspects of air quality not typically monitored.

Simplifying Analysis

With any human-machine interface, the desire is to reduce the amount of time required to interact with the machine. A simple WSN collecting data might require frequent analysis and the manual fusion of data by the human user. For many, the burden of this activity might preclude the use of the system. For many of the listed applications, the situation would be much improved if the system itself were able to identify and alert the user to events of interest. This problem can pose a significant challenge, as environments will vary widely in terms of air quality. For accurate environment-specific event recognition, the system must be trained. While the development of an intelligent, learning system is left to future work, this dissertation describes an event detection framework that provides freedom from direct monitoring of output and enables flexibility in defining triggering mechanisms.

System for a New Paradigm

Connecting these components will provide a framework and system for the measurement of air quality monitoring and event diagnosis that has not previously been available. This combination will enable the use of multiple OPC-enabled air quality measurement devices to perform measurements of unprecedented node density in diverse

environments with minimal preparation. The low-cost aspect will facilitate penetration into areas previously unreachable by commercial air quality devices.

Work Summary

This dissertation takes a step toward enabling a new set of applications through the creation of an Air Quality WSN, combining a cost-effective airborne particulate matter sensor with other commercially available sensors for air quality factors. The dissertation focuses on three areas in particular: creating the particulate matter sensor and WSN, making the WSN usable for a wide range of situations through the development of an event detection framework, and demonstrating the system through testing in real-world environments.

The creation of a particulate matter enabled WSN is a novel combination and extension of existing work across multiple disciplines. The dissertation work builds on the following areas of embedded systems and environmental factor sensing:

- Optical particulate matter sensing based on light scattering from single particles,
- WSN development including hardware, firmware, and algorithmic development, and
- Event detection and notification in sensor networks.

The work performed for this dissertation has been in some cases the result of a collaboration of several students working together to build components of the overall system. This has been true of some of the base platform work, some of the software work, and in areas where expertise is required that lies outside of the author's field such

as mechanical engineering. The work performed for the basis of this dissertation is listed below.

System Design and Prototype

Low-cost, small form-factor particulate matter sensors are created that approach the data quality of more costly commercially available devices. The specific approach uses a focused laser and direct sensing with a photodiode, avoiding the use and expense of additional optics to enhance the scattered light. The optical measurement area encompasses only a portion of the airflow, allowing the use of a low static pressure axial fan. While sensors with similar characteristics have been created or proposed previously, none go to this extreme of cost and size. Sensor arrangements not previously seen in low-cost designs are also explored. The bulk of this work has been performed by the author, with the assistance of mechanical engineering students on the mechanical design and electrical engineering students in the development and construction of prototype systems.

System Simulation

While basic scattering calculations can predict some of the results delivered by the particulate matter sensor, a more comprehensive system simulation provides a powerful tool for design analysis. A system simulation is developed that is used in conjunction with physical prototypes to predict and analyze the behavior of the particulate matter sensor. This work has been performed by the author.

Particle Detection and Sizing

New methods to improve particle detection and sizing are developed, with low-cost single particle sensing through the combination of data from two photodiode sensors.

This is compared to single sensor results in laboratory testing. This work has been completed by the author with minimal assistance.

Component and System Optimization

Novel methods are developed to correct for component and system variability related to low-cost sensing. This work has been completed by the author with minimal assistance.

Wireless Sensor Network Platform

A supporting WSN platform is developed. The base platform contains many elements common to other platforms found in research and industry, yet also contains some unique components. This work has been completed in cooperation with a larger team. Key pieces of the platform have been developed by the author.

Detection Framework

A framework is developed for event detection to ease deployment of air quality WSNs across a wide range of applications. The framework builds upon many common themes in WSN research with aspects that are novel to particulate matter sensing. This work has been completed by the author with assistance on server platform development from a larger student team.

Demonstration

The system is demonstrated through operation in several different environments. This is the first known demonstration of a particulate matter equipped WSN with this

level of cost and density. This work has been completed by the author with assistance from a larger student team.

Specific details of each of these activities can be found in the following chapters of this dissertation. The remainder of this dissertation discusses previous work (Chapter 2), air quality sensor system (Chapter 3), optical particle counter hardware (Chapter 4), optical particle counter simulations (Chapter 5), particulate matter detection and sizing (Chapter 6), particle counter testing and calibration (Chapter 7), event detection framework (Chapter 8), and system demonstrations (Chapter 9). This is followed by chapters discussing future work (Chapter 10) and conclusions (Chapter 11).

CHAPTER TWO: PREVIOUS WORK

One might say that the measurement of airborne particulate matter has been performed since early humans first noticed the haze from dust and smoke. Equipment specifically designed for this purpose has been developed that employs multiple methods with various tradeoffs for each. The U.S. Federal Reference Method (FRM) for sampling PM_{10} and $PM_{2.5}$ involves inertial particle size separation, capture on a filter, and weighing the filter before and after sampling to determine the mass gain [32]. This method is capable of providing very good mass measurement but does not provide real-time data. Other mass measurement methods such as the Tapered Element Oscillating Microbalance (TEOM) [33] can provide more timely data but are typically large and must use a selective input to resolve particle size.

Devices using optical methods to measure particulate matter can be split into two main categories: those that operate on the scattering and extinction of light from single particles, and those that operate on the scattering and extinction of light from an ensemble of particles [34]. Devices operating on an ensemble of particles are typically unable to resolve particle size. Single particle devices are commonly referred to as optical particle counters. These devices operate by illuminating a sample volume with intense light, where particles intersecting the light will scatter some portion of the light. The scattered light is sampled with a photodetector, producing a signal that can be analyzed to detect particles. This technology has been developed over multiple decades, with many

commercial devices available or described in patents [35], [36]. This includes devices utilizing laser illumination to detect single particles [37], [38]. Most of the commercially available devices are relatively expensive and have form factor limitations that make them impractical for inclusion in a WSN.

Some options for relatively low-cost optical particle counters do exist. A low-cost optical particle counter is available from Dylos Corporation [39], but this unit is also relatively large, does not include networking capabilities, and still costs approximately US\$350 for a single unit that allows data extraction. Components described as optical dust sensors are manufactured and sold by Sharp Microelectronics [40] and others, but these have very poor sensitivity to respirable particles relative to a laser-based OPC design.

Several new technologies under development have the potential to reduce the cost and size of currently available equipment. These include thin-film bulk acoustic wave resonator (FBAR) [41] and microelectromechanical resonator [42] particulate sensors. These technologies may be of significant interest in future developments of particulate matter sensing WSNs, but at this point are unavailable for deployment.

Most modern optical particle counters will depend on photoelectric conversion through the use of one or more photodiodes. The typical methods for counting single particles recognize that the electrical signal has reached a minimum threshold, and then measure the following pulse. Various methods are documented for calculating the particle size based on the maximum pulse amplitude, integration of the pulse, or measuring the pulse duration [43], [44]. Other techniques have been published that attempt to improve sensitivity through noise reduction and particle recognition with digital filters [45], [46].

Multiple instances of research WSNs targeted at monitoring air quality do exist. Some of these do incorporate some form of particulate monitoring [47], [48]. Each of these WSNs use the optical dust sensors mentioned above, which are only able to detect the coarse fraction of particulate. As such, these sensor networks will not give an accurate indication of the harmful respirable fraction of suspended particulate matter. Commercial systems do exist that include network capability, but these systems are expensive and not suitable for incorporation into an affordable WSN system. One such system described in a multiple room study uses a central measurement device that samples multiple points through a network of sampling tubes [49]. This type of system centralizes the cost of the measurement devices, but can be difficult to deploy relative to a WSN and cannot sample from each point simultaneously.

Some of the air quality monitoring systems described in the literature include diagnostic capabilities through the use of expert systems [50]. Environmental event detection through sensor networks has been explored for the determination of building occupancy [49]. WSN use for monitoring of patients in assisted-living has included the generation of alerts for healthcare providers [51]. Each of these systems appears to be tuned to the specific environment of installation or a specific application, and do not focus on providing a mechanism for real-time user feedback that could be extensible to multiple scenarios. Other air quality systems have been developed for monitoring air quality that provide web-based interfaces [24], [52], yet do not provide notification interfaces outside of directly monitoring the web site.

CHAPTER THREE: WIRELESS SENSOR NETWORK SYSTEM

The wireless sensor network (WSN) system forms the basic platform for the work performed in support of this dissertation. This WSN system is an evolutionary work, having been developed over multiple years by the Hartman Systems Integration Laboratory (HSIL). The latest generation of the WSN system includes the optical particle counter (OPC) developed as part of this dissertation. Much of the work on the evolutionary sensor system has been performed in partnership with the Federal Aviation Administration (FAA) under their Airliner Cabin Environment Research program. The inspiration and funding for the original design of the OPC came from work done for the National Institutes of Health's (NIH) National Children's Study. The Wireless Air Quality Monitor (WAQM) device produced for this study is the basic sensor node from which is built the larger WSN system. This basic sensor node developed for the NIH is targeted at monitoring the environments of children in schools and homes. The WSN system and OPC have evolved together, with many aspects of the system development being performed by the larger team working in the HSIL. Given that the WSN system continues to evolve, this chapter will describe the state of the system as it existed near the end of the work supporting this dissertation.

This chapter will start with an overview of the WSN system and will describe example deployments of the system. The WAQM sensor node will then be described, followed by the server and visualization software used to complete the system.

System Overview

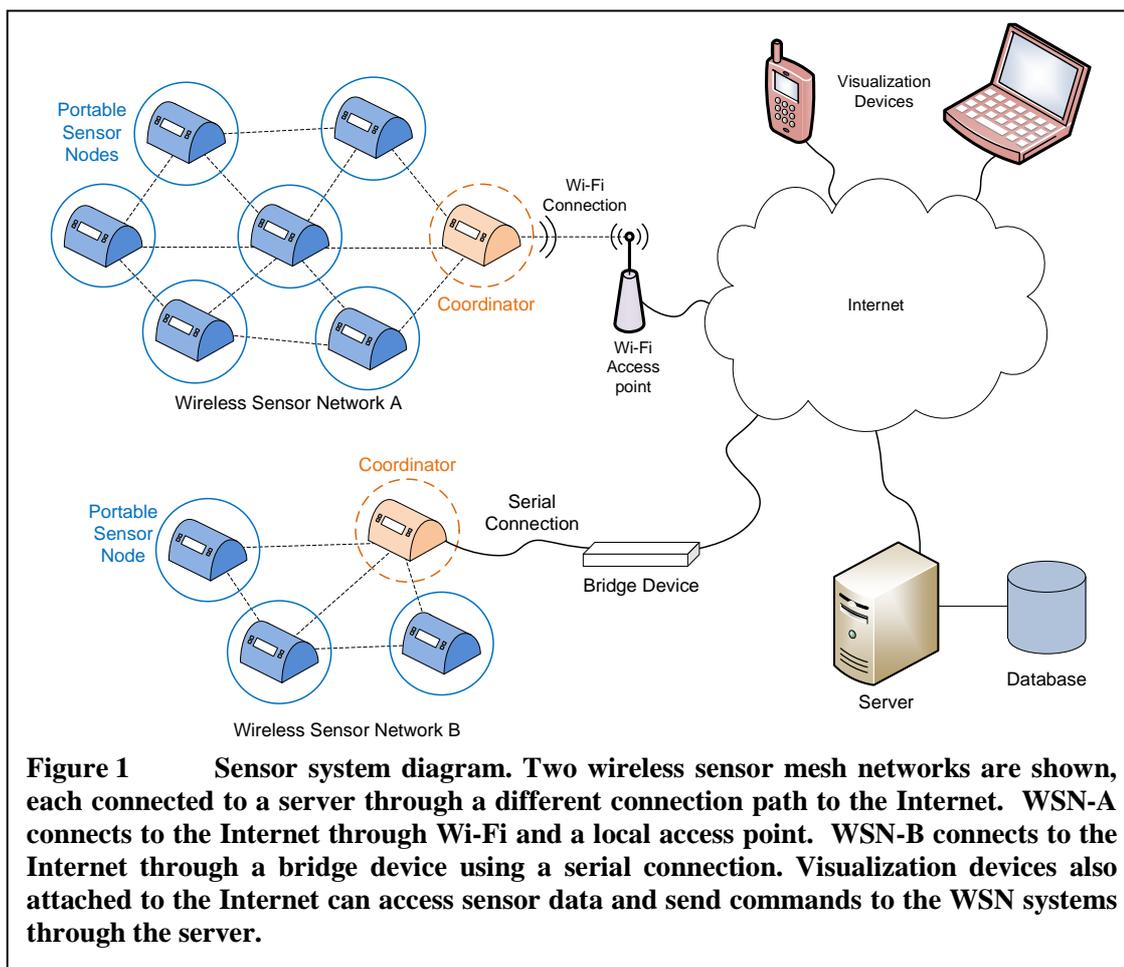
The WSN system incorporates many parts to create an entire solution. This includes sensor nodes, mesh network coordinators, servers, databases, and visualization devices. Data originates at the sensor nodes taking environmental measurements. These nodes are arranged into wireless mesh networks that are controlled by the mesh network coordinators. The coordinators bridge the mesh networks to the outside world, connecting to the Internet through several available connection methods. Data are accepted by the servers, and either passed on to a storage database or directly to a visualization device. The visualization devices are the interface points for the end users of the system, providing a platform for software that interprets the data generated by the system.

Each of these pieces must work together to achieve the sensor system's mission. At the highest level, this mission is to simplify the collection of high-quality environmental data from a distributed set of points and present it to the user in an understandable form. The individual pieces must together constitute a cost effective and robust system to perform this mission.

System Deployments

There are many possible deployments of the WSN system, with multiple wireless mesh networks connected to the Internet in several different ways. Data can flow across the Internet to the server and database, or can flow directly to a visualization device. Similarly the visualization devices can pull data either from the server or directly from the sensor nodes in the mesh networks.

Figure 1 shows an example deployment of the air quality sensor system. In the example arrangement, two WSNs are shown with communication links to a server over



the Internet. WSN-A connects to the Internet through a Wi-Fi access point and WSN-B through a serial connection to a bridging device. In the first case, the sensor data from the WSN are passed from the coordinator directly to the server over the Internet. In the second case, a Wi-Fi connection is unavailable and the coordinator connects to the Internet through a bridging device using its serial connection. The bridging device forwards the data to the server across the Internet.

Alternatively, the server does not need to be used at all. Data may be passed directly to an end user's PC for real-time viewing with the appropriate visualization software. The server software may also split the data out to multiple endpoints, targeting

the database and a set of real-time viewing visualization devices. If the mesh network cannot be used, data may also be collected only on the end nodes with no transmission. The data are retrieved from each sensor's local storage at a later time to be merged into a single view. The overall system provides flexibility in communication and data collection to enable a number of possible configurations that will fit many different sensing scenarios.

Sensor Node Hardware

The main hardware building block of the WSN system is the sensor node. The most recent implementation of the HSIL sensor node architecture is the Wireless Air Quality Monitor (WAQM), which has been developed for the NIH National Children's Study. The WAQM sensor node, shown in Figure 2, was designed to operate as a standalone unit and as part of a larger sensor system. As such, it has components that may not be typically seen in a WSN. In particular, each sensor node has a liquid crystal (LCD) display for viewing current measurement data. While this adds to the cost of each sensor node, it can be seen as a superset of the functionality required for any given implementation of a sensor node. Subsequent implementations include sensor nodes that have only a subset of the functionality of the current design,

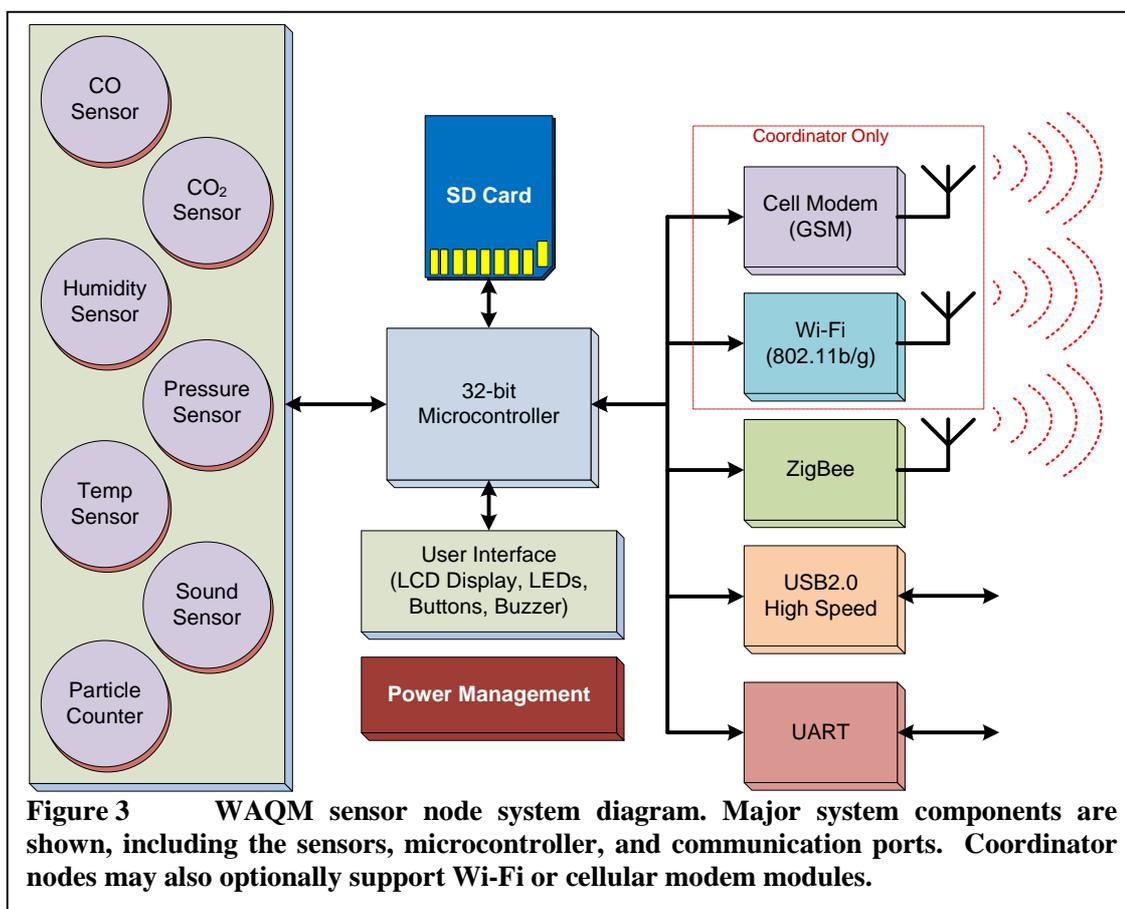


omitting the unnecessary sensors and components as the needs of the application dictate.

The architecture of the WAQM sensor node hardware is shown in Figure 3. The major electronic components of the sensor node are the microcontroller, sensors, communication, data storage, user interface, and power management. These components are described briefly in the following section. The optical particle counter is described in detail in later chapters.

Microcontroller

The WAQM sensor node is architected around an Atmel AVR32 microcontroller, which contains a 32-bit RISC CPU that is capable of 1.51 Dhrystone MIPS per MHz [53].



The part includes 256 kB of flash memory, 128 kB of SRAM, and many other features that are useful for the development of an embedded system. An integrated 10-bit analog to digital convertor (ADC) is critical for use with the optical particle counter. The part also includes a variety of interconnects including four USARTs, two Serial Peripheral Interface ports, two Inter-Integrated Circuit (I2C) ports, a Secure-Digital (SD) interface, and a High-Speed USB interface.

The processor also contains a real-time clock (RTC) function. The processor does not retain the time in the event of power loss, so an external, coin-cell battery powered RTC connected to one of the processor's I2C ports is used for this purpose. On initialization, the external RTC time is loaded into the processor RTC, which has a much faster access time. The internal RTC is then used for all measurement timestamps.

This particular microcontroller was chosen for the implementation of the sensor node due to its rich peripheral set, relatively large internal memories, good performance/power ratio, and available software framework. The feature set allows a large portion of the sensor system to be implemented with the microcontroller alone.

Power Management

The design goals for the sensor node dictated that the device be able to run in stationary positions for extended periods of time, and in mobile environments for short periods of time. With this in mind, the sensor node was designed to run off of either an external wall transformer or an internal rechargeable battery. One category of applications requiring battery operation is air quality measurement during transportation activities, such as in busses, automobiles, bicycles, or even directly carried by pedestrians. The battery affords the additional flexibility of quickly deploying a

temporary sensor setup without worrying about power distribution, such as outdoors or in a large meeting hall. The sensor node uses a two-cell lithium-ion battery with a 1,400 mAh capacity. The sensor node consumes approximately 1.3 Watts of power when operating with sensors and the ZigBee network active. With a full charge, the unit can run for approximately 8 hours from the internal battery.

The sensor node power circuit can run from a voltage range of 7.3V – 20V, and is typically powered from a 12V wall transformer when not running from the battery. For battery charging, the input voltage must be at least 11.7V to meet the operating requirements of the battery charging circuit. To prevent over-discharge of the battery, a voltage supervisor circuit automatically shuts down the system voltage regulators when the battery voltage falls below 6.3V.

Most components operate from a +3.3V supply, including the microcontroller, analog-to-digital converters, much of the analog circuitry, and the particle counter laser. Some components such as the electro-chemical sensors and their support circuitry require +5V and -5V supplies. The particle counter fan also requires a +5V supply. To maximize efficiency when running from the battery, switching regulators are used to provide power at the positive voltages. The parts enable conversion efficiencies of well over 80% at the relevant system loads. The parts run at 1.6 MHz, allowing the use of relatively small inductors and capacitors. Since only a small amount of current is required from the -5V supply, a switched capacitor regulator is used. This is sufficient to generate the negative voltage supply for the sensor amplifier circuits.

Due to the particle counter amplifier's high gain, the circuit is very susceptible to power supply noise. Any disturbance on the +3.3V supply for this circuit can be

amplified into signals that appear very similar to the photodiode current produced by a detected particle. For this reason, a separate linear regulator supply is used to generate a quiet +3.3V supply dedicated to the particle counter amplifiers. The quiet power supply allows the particle detection threshold to be lowered substantially, making the system more sensitive to small particles.

Communication

Communication of sensor measurements and status information is a critical function of any sensor network. For the sensor node design, multiple communication options have been provided for wired and wireless connection. All sensor nodes are configured with a UART connection, a USB port, and a ZigBee radio. The UART connection operates at 115.2 kbaud, and is typically used only for debugging during development due to the lack of serial ports on most modern computers. The USB connection offers a readily available connection option through a mini-USB port on the device. On a personal computer, the USB connection emulates a COM port, allowing communication with standard terminal programs or the BSU Sensor Monitor application.

Several options exist for building the ZigBee radio connection. There are multiple vendors offering System-on-Chip (SOC) and module solutions. Module solutions have the benefit of incorporating most of the radio frequency sensitive components and layout, often coming with pre-certification for modular regulatory approval, which can simplify the testing requirements for a design. The WAQM system uses such a module design.

The behavior of the ZigBee radio connection depends on the configuration of the sensor node, particularly in whether the node is a set up to be a coordinator or a router. For router nodes, the ZigBee connection is used to transmit sensor measurements and

status to a coordinator node. The coordinator node receives sensor measurements from other nodes over the ZigBee network and forwards them on to another connection or its own internal storage.

Beyond the communication methods discussed above, the coordinator nodes may also be configured with either an 802.11b/g Wi-Fi transceiver or a cell modem. These connections enable remote data collection and the extension of the sensor network beyond a single ZigBee mesh. While the Wi-Fi transceiver is fully supported and commonly deployed with test systems, the cell modem has only been prototyped. In both cases, radio modules are used to simplify the radio frequency design.

Data Storage

Many sensor deployment situations require the integrity of data regardless of the state of the communication connections. A sensor node's connection to the rest of the system may be lost or interrupted, leading to the loss of transmitted data that may be critical to the application. For this reason, all sensor nodes are equipped with local data storage in the form of an SD card slot. With the current availability of relatively low cost multiple-gigabyte SD cards, a significant number of measurements can be stored directly on the sensor node. Furthermore, the coordinator nodes participating in a ZigBee network may optionally log all network data into their local SD card for a consolidated data collection point.

Data retrieval from a sensor node can be performed either directly by physically removing the SD card from the unit, or over the sensor node USB port. For the latter case, a special mode exists that allows the sensor node to enumerate as a mass-storage

device, appearing to the connected computer as an external storage drive similar to a USB flash drive.

User Interface

Visualization can be a key component in modifying human behavior. Often, sensor nodes will provide little in the way of visual feedback to indicate the sensor measurement values. Early prototypes of the WAQM sensor node provided little in the way of direct visual feedback to indicate their status, and gave no indication of the sensor measurement values. To improve on this situation, the current version of the sensor node implements a backlit LCD panel that can give indication of multiple sensor values simultaneously. This allows the casual user to know the current measurement level at a glance, giving some indication of environmental status without requiring the use of a computer.

The display can scroll through several sets of measurements with a set of four buttons located around the edge of the display. These buttons are also used in special sequence to set the state of the LCB backlight and to start a new set of log files on the device SD storage.

Sensors

Besides the optical particle counter, the WAQM node includes sensors that measure CO, CO₂, humidity, atmospheric pressure, temperature, and sound pressure level. These particular sensors were chosen to meet the needs of the NIH National Children's Study program, and allow monitoring of various environmental conditions related to air quality, health, and comfort. Table 1 lists the sensors incorporated in the

Sensor	Technology	Range	Resolution	Accuracy	Response Time
CO	Electro-chemical	0-500 ppm	0.1 ppm	2 ppm	40 seconds
CO ₂	Infrared	0-2000 ppm	1 ppm	10% of reading or 75 ppm	< 120 seconds
Humidity	Capacitive	0-100% RH	0.03%	+/- 3%	5 seconds
Pressure	Diaphragm, capacitive	30-120 kPa	3 Pa	+/- 150 Pa	0.55 seconds
Sound	Electret microphone	48-110 dBA	0.02 dBA	--	< 1 second
Temperature	NTC thermistor	-40oC - +125oC	< 0.01oC	5% of reading	1.2 seconds

design, along with some of their key datasheet parameters. The sensors in the WAQM monitor were chosen to measure conditions in a range relevant to human comfort and safety. In particular, the gas sensors cover a range near the conditions expected to be seen during normal activity. The CO sensor also covers ranges that are dangerous to human health, but the intent is to provide relatively high accuracy at low exposure levels.

Interfacing to a diverse set of sensors can be one of the more challenging aspects of designing a sensor system. Sensor outputs can range from a standard digital interface that can be directly connected to a microcontroller, to a simple current requiring amplification, level shifting, and digitization. Sensors may also need controlled stimulation to provide output, such as commands sent through a digital interface or timed waveform generation. It is the job of the sensor node to “normalize” these requirements to produce sensor measurements in a common form that can be handled by the communication and processing components that make up the larger system.

From a hardware standpoint, the objective is to provide reasonably stable digital data to the system processor. In some cases, as with the CO sensor, the output signal is a current proportional to the gas concentration. A potentiostat circuit drives the sensor,

producing the output current. This is converted into a voltage by a transimpedance amplifier, which is then level shifted by a final non-inverting amplifier stage. The resulting voltage is fed into an analog-to-digital convertor, which is connected to the microcontroller using an SPI bus. Other sensors, such as the pressure sensor, integrate the digitization process and present conditioned values directly to the processor across a digital bus.

The sensors in the WAQM monitor are calibrated in a number of ways depending on the physical parameter measured. For the gas sensors, fixed gas concentrations are purchased and used in an enclosed chamber to calibrate the sensors at multiple concentration points. Humidity, pressure, temperature, and sound are calibrated against commercial equipment at multiple points.

Sensor Node Firmware

The development model for sensor systems can dictate rapid changes to system components, including basic sensors, communication interfaces, and even the system microcontroller. Systems may require customization for a particular application or environment, leading to multiple unique instantiations of the same base system. Such a model often requires multiple developers working on the same code base simultaneously, creating challenges with division of labor and change coordination.

The HSIL team has developed a layered, modular firmware architecture for WSN nodes to support multiple different deployments, including the WAQM sensor node. The firmware code base is predominantly developed using the C programming language. The lowest level of the sensor node firmware is the device driver, which is responsible for interfacing with the system hardware. The device drivers abstract the details of the

underlying hardware, presenting a concise interface to upper layers and enabling a simplified programming model. An interrupt-based scheduler sits on top of the device drivers, providing critical high-level timing for the system and all firmware modules. The firmware architecture for sensor management provides a framework for quickly and easily adding new sensors to the system. The sensor manager collects data from each sensor, passing it on to the storage and communication modules through a data manager. At the logical top of the firmware architecture, the application layer controls the overall functionality and behavior of the system.

Architecting a set of embedded system firmware to handle these conditions requires solid interface definitions and a high level of modularity. The embedded firmware architecture for the WAQM sensor node is shown in Figure 4. The layered, modular architecture is designed to meet the challenges posed by the sensor system development model.

Tasks

To aid in the modularity of the architecture, the firmware for the sensor node is

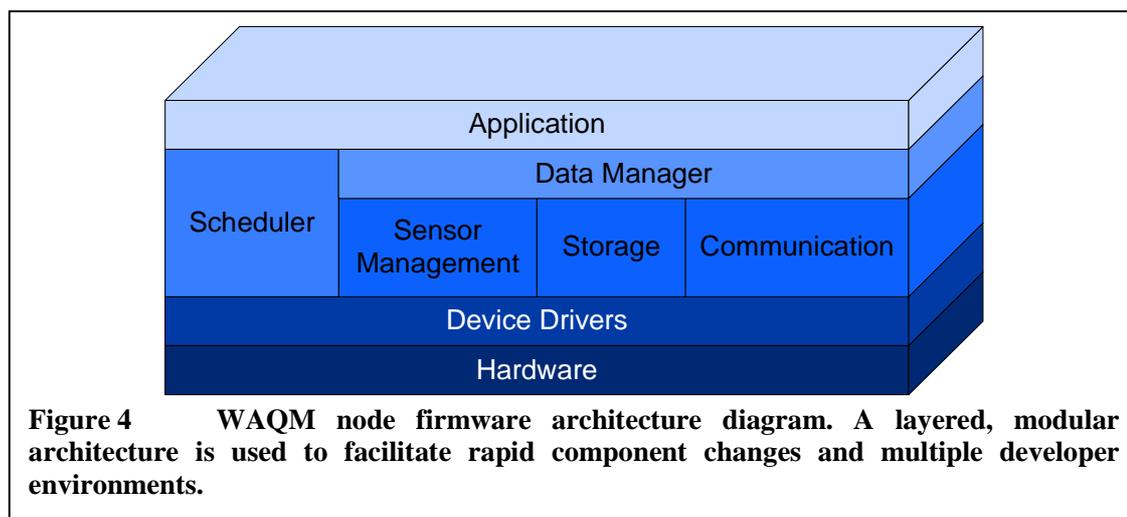


Figure 4 WAQM node firmware architecture diagram. A layered, modular architecture is used to facilitate rapid component changes and multiple developer environments.

split into multiple tasks with each task handling a different part of the system functionality. This cooperative multitasking system runs as a simple loop with tasks executing sequentially, using an unconstrained amount of processor time. To avoid starvation of other tasks in the system, all tasks are designed to be non-blocking and to use a limited amount of processor time when executed.

Tasks may be configured at system run-time through the use of configuration information on the SD card, including which tasks are executed in the run loop. This allows a common code base to be compiled for multiple sensor node configurations, with only the configuration file determining the specifics of operation at run-time. For instance, the firmware may be compiled to support a superset of sensor and I/O modules. A set of sensor nodes with various subsets of these sensors and I/O modules may all share this firmware executable, disabling firmware modules for unavailable hardware through their respective SD card configuration settings.

Device Drivers

The lowest level of the sensor node firmware is the device driver, responsible for interfacing with the system hardware. The device drivers abstract the details of the underlying hardware, presenting a concise interface to upper layers, enabling a simplified programming model. This abstraction makes the interface between the device drivers and higher layers a key component of firmware portability, as a well-written interface can allow the higher layers to be used on multiple hardware platforms with little modification. The device drivers on the WAQM system cover a wide range of hardware, including system interfaces and blocks internal to the microcontroller.

Scheduler

An interrupt-based scheduler provides critical high-level timing for the system and all firmware modules. The scheduler firmware provides a method for registering call-back functions for modules that require timed code execution. Methods are also provided for timed signal generation on General purpose I/O (GPIO) pins, and timed sampling of the microcontroller's internal analog-to-digital converter.

Sensor Management

The firmware architecture for the sensor subsystem provides a framework for quickly and easily adding new sensors to the system. Firmware functionality for managing the sensors in the system is handled by a single sensor task. The sensor task maintains a list of sensors currently available to the system, along with key parameters for sensor identification and data access. Each sensor has a specific task that is called by the main sensor task, allowing the execution of sensor specific functions at a priority determined by the main sensor task. Currently the sensor task simply runs each sensor function in a round-robin loop, though the system is extensible to more complex schemes.

Data are passed to the main sensor task by each sensor through a shared data structure. Notification of new data is indicated using one of two methods. The first method indicates that data is ready for consumption by the system, and should be passed through to each data target such as ZigBee or the log file on the SD card. The second method indicates that data are ready for sampling by functions internal to the system, such as the front LCD display. This second method allows features such as a display that

updates at a rate faster than the official sample rate of the sensor, providing for a more interesting human-machine interface.

Storage

An important consideration for internal storage for a sensor node is how it will be accessed. For the WAQM system, the internal storage takes the form of a removable SD card. To ease the retrieval of data from internal storage, a FAT file system is used on the card that allows access by any PC with an SD card reader. The sensor node firmware implements multiple file handles to allow the storage of sensor data and system events concurrently. Additionally, a USB mass storage interface has been implemented to allow internal storage access from a PC using only a USB cable.

Communications

Several communication interfaces have been developed to enable the WAQM sensor to communicate with the outside world, including serial, ZigBee, Wi-Fi, and cell modem. Each communication task implements the specifics of its communication protocol, interfacing to the data manager for sending and receiving data in the system. The modular nature of the firmware architecture allows the addition of new interfaces with minimal impact to the rest of the system.

Data Manager

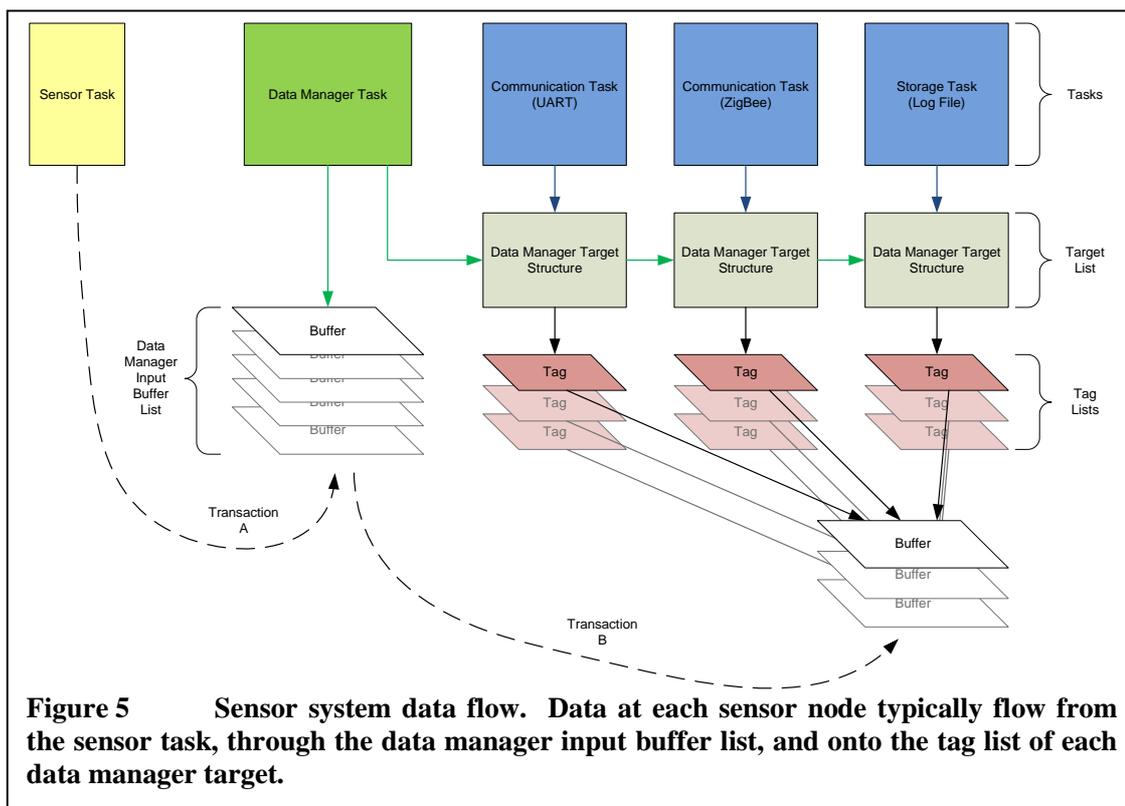
Data flow in the sensor system is handled by the data manager. The data manager provides a general framework for passing data through the system with minimal resource utilization. The data manager can handle a large number of input sources and output targets, supporting data replication to multiple targets on an individual buffer basis. Input

sources all feed their generated data buffers into the data manager's input buffer list. Target destinations are tracked through a list of data manager targets. Each target is associated with a particular system task, typically a communication or storage task such as ZigBee or the system log file.

Data buffer allocation is handled by each source, allowing optimization of the amount of memory used for each source application. Each data buffer also carries a priority, allowing the data manager to perform functions such as quickly passing through critical information or discarding low-priority information during a period of congestion.

The data manager uses a pool of small "tag" structures to track ownership of each data buffer, linking the data buffer with each of the relevant data manager targets. As a new data buffer is received into the data manager's input buffer list, its destination targets are determined and a new tag is set up for each association. Information in each data buffer tracks how many tags are currently attached to the data. As the data are consumed by each data manager target task, the associated tag is returned to the data manager. Once all tags associated with a particular data buffer are returned, the data manager returns the buffer to the originating source for re-use.

Figure 5 shows the operation of the data manager. In this representation, the sensor task is generating data buffers that contain sensor information. These buffers are passed to the data manager by writing target information to the header of each data buffer, and then inserting them in the data manager input buffer list. In this case, the data buffers are each destined for the UART, ZigBee, and log file tasks. The data manager parses the target information in each buffer, and then assigns a tag for each unique target. Each tag is added to the appropriate tag list associated with a data manager target. The



data manager target tasks use the information stored in their associated data manager target structure to access their tag lists, and retrieve the data.

All sensor, control, and status data that pass through the data manager are represented in a custom binary format. The format has been optimized for size efficiency, positively impacting the amount of data storage required in the microcontroller and the amount of bandwidth required to transmit the data.

Application

The application layer of the sensor node firmware controls the overall functionality and behavior of the system. At present, most of the sensor system runs without direct management from the application layer. The critical functions of the

application layer are registration of nodes across the communication interfaces, time synchronization, and command communication from the network back to the nodes.

Wireless Sensor Network

Multiple options exist for building a wireless sensor network [54], including using standards-based implementations or creating a custom network. A standards-based implementation brings many benefits, including behaviors that are typically well-defined and tested as well as off-the-shelf hardware and software support. For the WAQM WSN implementation, the ZigBee/IEEE 802.15.4 standards were chosen. These standards define a robust, low-power wireless mesh networking solution. Several functions built on top of the commercial ZigBee stack allow the sensor system to meet some of its critical goals, including a reliable network layer and time synchronization between sensor nodes.

Network Layer

Within a WSN, the flow of sensor data often moves from the sensor nodes across one or more wireless hops to a single coordinator node acting as a data sink. It is important to note that the generation of data may never cease. Sensors may be set up to periodically poll one or more environmental parameters and transmit this data into the network, creating a steady stream of information that must be passed on to a final destination. Within any network exists the possibility that communication will be interrupted through congestion or physical disruption of a link. When this occurs, the WSN must determine what to do with the data that are still being generated.

Further complicating matters, to keep costs and power consumption down, a WSN node may be constructed using a relatively small amount of memory. Even when

data can be buffered during a communication interruption, care must be taken not to congest the network with pent-up data when the disruption is resolved. In short, there will be a finite amount of time in which the data stream can be buffered before something must be lost.

The WAQM sensor nodes implement a network layer on top of the WSN that performs best effort delivery with a relatively small amount of buffering. The buffering is intended to only handle short periods of disruption or congestion. Future enhancements to this algorithm may include increased buffering with delayed delivery, or data rarefaction, with adaptive pacing to avoid congestion.

The WAQM network layer creates a virtual socket interface between an end node and the coordinator. This allows the coordinator to recognize and manage the connections to multiple sensor nodes, and provides a framework for reliable delivery. The network layer handles fragmentation of long messages across multiple ZigBee packets, and can aggregate multiple short messages into a single ZigBee packet to optimize the number of transmissions. The network layer also supports an aggregate packet type that will pack multiple sensor readings into the same packet, reducing overall network overhead with the cost of additional delay between measurement and data delivery. The network layer for the mesh network uses a binary packet format, while messages sent outside the mesh network are converted into a readable, ASCII text-based format.

Time Synchronization

A key aspect of any WSN is the ability to temporally correlate measurements between nodes. Each measurement must be time stamped in some way to indicate its location in the data chronology. This could be accomplished by time stamping the data at

the final destination, but this method would not take into account the time delay for data to transit the network. This delay could be highly variable depending on current network conditions, such as the presence of congestion or an interrupted link. Mesh networks with the potential for multiple hops and dynamic route changes between a sensor node and the coordinator are very susceptible to this variability. Destination time stamping would also preclude the ability for the sensor to store data locally with reference to a known time.

Resolving these issues can be accomplished by keeping a local time at each node, and then synchronizing that local time with a master in the sensor system. Rather than worrying about delay variation for each measurement packet as it crosses the sensor network, the problem is isolated into one of estimating delays for a controlled exchange during a clock synchronization event. Many well-tested clock synchronization algorithms are already in existence that can produce very good results [54].

The WAQM sensor system currently synchronizes times by sending a request from the sensor node to the coordinator. The coordinator replies with its current time, and the sensor node uses this along with the round trip time to calculate the current time at the node. The round trip time is measured by the sensor node from request transmission to response receipt. For relatively small test deployments, this method has typically resulted in synchronization with sub-second accuracy. With a large number of nodes, the round trip time can increase significantly with the possibility of miscalculating flight times if one hop is disproportionately delayed. This algorithm might be modified in the future to include multiple round trip estimates to produce a more accurate and consistent synchronization.

Server Software

With a functioning sensor network, it quickly becomes apparent that the data never stops coming. Dealing with the copious amounts of data is a significant challenge in any such system. A set of server software has been developed as part of the WAQM system to deal with the transport, storage, and display of collected data. This section will briefly cover these functions and their implementation in the WAQM system.

Connector Server

In the initial implementations of the WAQM system, the WSN coordinators would connect directly to the WAQM database. This made it difficult to access the data streaming from the WSN in real time. Often, real-time access is desired when directly monitoring a deployment, especially when performing laboratory experiments or when setting up a new deployment. To provide a tap into this stream without disrupting the database connection, the Connector Server software was created. This is a program written in Perl that performs bi-directional routing between the WSNs in the system and one or more server or visualization nodes on the network.

The choice of Perl for the Connector Server was made for portability, allowing the software to be run on Linux or Windows-based machines. During the time frame for development of the WAQM system, the popularity of the low-cost ARM-based Raspberry Pi computer [55] took off. These systems run Linux and can be easily deployed with the Connector Server running on them to provide a tap into the data stream local to the WSN deployment. The Connector Server can also be run on the same machine hosting the database or on one's own PC, but a Raspberry Pi deployment allows

one to run the server and leave it running uninterrupted as an independent piece of the system.

The Connector Server may also be used as a bridge between a coordinator and the Internet, as shown for WSN-B in Figure 1. This variant of the Connector Server allows bridging to a wired network when Wi-Fi is not an option.

In addition to its routing functionality, the Connector Server has been the platform for implementing notifications for event detection. This will be described in further detail in the event detection chapter of this dissertation.

WAQM Database

Long-term storage of the data from the WAQM system is a critical function, as much of the value in a sensor network comes from access to historical data to examine long-term trends. To provide this functionality, a dedicated Linux server has been provisioned with a MySQL database. A program written in Perl provides a connection into the database through TCP/IP, accepting incoming connections from WSN coordinators or instances of the Connector Server. Data streaming into the connection is parsed for errors, and then written to the database.

The database is structured with multiple tables. The largest of these is the measurement table, which contains all measurements from every device in the system. Measurements are stored along with a sensor unit identification number and a sensor type. Both the time of measurement reported by the unit and the time of insertion are stored. Two separate data fields allow for storage of a raw measurement and a formatted measurement. The use of these two fields is dependent on the particular sensor, but the

former is typically used to store a raw digital reading and the latter for a scaled, presentable value. A location field is also updated with the reporting unit's current deployed location as stored in the database.

Other tables in the database include a table of units associated with the system, a table of sensor types used in the system, a table of defined locations for deployment, and a table of events. The unit table holds a WAQM sensor node's unique unit ID along with a description and the current location of the unit. The sensor table holds a list of sensor IDs along with sensor descriptions, models, and reporting units to be referenced by entries in the measurement table. The location table holds a list of deployment locations and descriptions, which are referenced by the unit and measurement tables. The events table holds events reported by units in the system, including the event type, time, and related sensor ID.

The WAQM database is currently managed through a manual process. A persistent problem with this and any sensor network database is its propensity to grow to a very large size over time. With multiple units and multiple sensors, each reporting values multiple times per minute, the WAQM database can grow to hundreds of Gigabytes over several months of operation. This requires management of the data through periodic purges of the database. This may be automated in the future to avoid the ongoing effort in data management.

Web Graphing Interface

A final piece of server software spans the domains of server and visualization. To facilitate viewing on mobile devices, a java servlet was developed to create on-demand plots of sensor data pulled from the WAQM database. These are created and served as

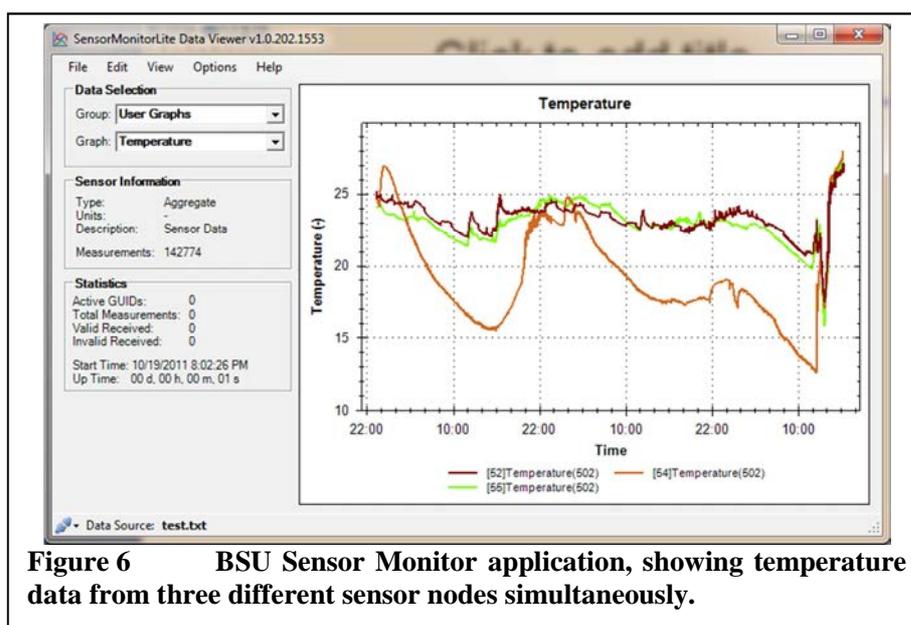
Portable Network Graphics (PNG) images embedded in an HTML page. The servlet will respond to a set of inputs specifying the units, sensors, time frame, and other relevant factors, creating the plot on the fly. More will be discussed regarding this part of the server in the event detection chapter of this dissertation.

Visualization Software

For many applications, the ability to collect sensor data is useless without the ability to store and visualize it effectively. With this in mind, a software application has been developed to provide these functions for the WAQM monitor and other HSIL sensor systems.

The Boise State University (BSU) Sensor Monitor application is the primary software used to visualize data generated by the sensor system. The application is capable of receiving data directly from the network via a TCP/IP connection, from a standard PC COM port, or alternatively can load stored data from a flat file. The sensor data can be plotted by unit, by sensor, or with custom groups that allow simultaneous visualization of multiple types of data at once.

Figure 6 shows a screenshot of the Sensor Monitor application in use, with temperature data from three different nodes displayed simultaneously. A key advantage of this software when used with the WAQM system is its ability to update plots automatically based on real-time data. In the WAQM system framework, this allows an instance of the Sensor Monitor software to connect directly to a Connector Server, providing real-time visualization of the data stream.



Wireless Sensor Network System Conclusions

The development of the WSN system was a significant portion of the work required to form a basis from which this dissertation could be written. This work was a cooperative effort shared among multiple engineering students working in the Hartman Systems Integration Lab. While the work was shared, much of the sensor framework, data management, and mesh networking layer were developed by the author. This provided a base platform for the development of the optical particle counter and the air quality wireless sensor network.

As it currently exists, the WAQM system has performed admirably, with its longest uninterrupted deployments lasting over three months without intervention. The evolution of the WSN system may look to move to more advanced methods of managing the data flow through the system, such as taking control of the mesh networking stack software and running a preemptive multitasking operating system. These might lead to

more reliable data movement for future platforms. More attention might also be paid to system cost and power consumption, though for prototype quantities these issues may be of secondary concern.

CHAPTER FOUR: OPTICAL PARTICLE COUNTER

This chapter describes the design of the WAQM optical particle counter hardware. The description is split into several sections. In the first section, the design of the analog front end (AFE) is described, which converts scattered light into an analog signal presented to the system analog to digital convertor. The second section covers the analog-to-digital conversion process, specifically detailing the use of hardware in the microcontroller to offload the system firmware. The third section discusses considerations in selecting the laser source and some of the resulting issues impacting the OPC. The fourth and final section covers mechanical issues, including the airflow carrying particulate matter past the sensor and structures for light management.

Analog Front End Design

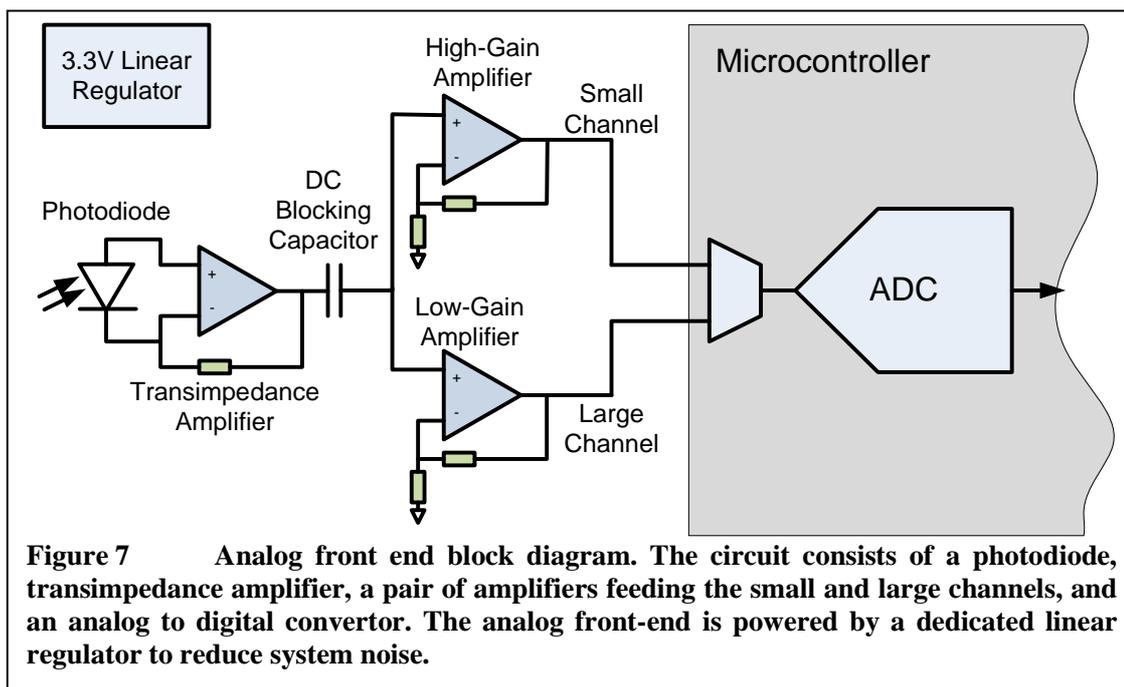
The analog front end (AFE) of the particle counter converts photons scattered from particles in the airflow into voltages that can be sampled by the system's analog to digital converter. Particles passing through the OPC laser will cause pulses of photons, resulting in pulses of voltage on the outputs of the AFE. Inside the AFE, this process involves the conversion of photons into a current using a silicon photodiode, and the subsequent conversion of this current into a pair of voltage waveforms that can be sampled by the ADC. The two waveforms have different amounts of amplification. The first is a high-gain channel that is used to detect particles, with the large amount of

amplification allowing the detection of small particles. The second is a low-gain channel that is used to size particles.

The key metrics for the AFE design are to have good performance in producing pulses from particulate light scattering at a relatively low cost. With this in mind, the AFE is designed using commonly available components, including a silicon PIN photodiode, single-rail amplifiers, and a low-cost, low-dropout 3.3V linear regulator. The resulting design provides a reasonable level of performance using readily available, low-cost components.

Two different circuits are described in this section. The first is the baseline circuit that was initially used in the development of the WAQM OPC. The transimpedance amplifier in this baseline circuit was found to be susceptible to noise coupled from the AC power source through the wall transformer. The second circuit contains a redesigned transimpedance amplifier that addresses the noise issue. This second circuit is believed to be a novel approach to reducing noise coupling into a single-rail photodiode amplifier circuit when using a low-cost wall transformer without a direct earth ground connection. The improved circuit is referred to as the “bootstrap” AFE circuit, since the voltage swing across the photodiode is removed. The baseline and bootstrap circuits are described in detail, since the bootstrap circuit suffers from a higher total RMS noise than the baseline and in some cases may not be required for power line noise rejection.

Figure 7 shows a block diagram of the AFE for the WAQM OPC. Light enters the AFE at the photodiode, which converts the photons into a current. This current is then converted into a voltage by a transimpedance amplifier. A DC blocking capacitor follows the transimpedance amplifier to remove the positive voltage bias from this amplifier’s



signal. At this point, the AFE splits into two channels. The high-gain amplifier produces the small particle channel and the low-gain amplifier produces the large particle channel. These amplifiers are necessary to boost the signal out of the transimpedance amplifier, which is relatively low in amplitude due to the small amount of light received at the photodiode from particle scattering. The outputs from the pair of amplifiers are fed into the ADC contained in the WAQM OPC microcontroller, which then produces a digital representation of the signals.

The circuits described are used for two different configurations of WAQM OPC. The OPC may be either a single sensor arrangement or a dual sensor arrangement. The single sensor arrangement uses a single photodiode with a single AFE circuit, and the dual sensor arrangement uses two photodiodes with two AFE circuits in parallel. As such, the single sensor arrangement requires two ADC inputs to support the small and large channels. The dual sensor arrangement as first designed requires four, as each AFE

outputs a small and large channel. The dual sensor arrangement was shown to not require the large channel for sizing, and its final implementation only requires two ADC inputs to support the two parallel small channels. More will be discussed about the different sensor arrangements in following chapters.

AFE Power Supply

Both of the described AFE implementations are single rail designs. This means that only a positive voltage and ground are used to power the circuit, rather than a set of positive and negative voltages seen in many amplifier circuits. The main reason for designing with a single rail is cost, as removing the need for a negative voltage supply reduces overall system cost. Additionally, many low-cost amplifier integrated circuits are now available that support single-rail operation. For the following discussions regarding the AFE design, the positive supply rail is labeled as VCC, which is nominally 3.3V. This voltage was chosen to match that of the ADC inside the system microcontroller, which allowed the AFE to be connected directly to the ADC without further voltage conversions.

The very first designs of the WAQM OPC did attempt to operate from a 3.3V switching supply that was shared with the rest of the WAQM system. From a cost standpoint, this was ideal, since the relatively low current requirements of the AFE appeared to allow the use of the existing supply without modification, adding no additional components to the system. Unfortunately, the high-gain of the AFE makes it very susceptible to power supply noise, and even small disturbances caused by normal system activity would cause deviations on the AFE output that were difficult to distinguish from actual particles. These “false particles” would be detected and counted

by the OPC system. For example, an increased current draw of approximately 30mA on the 3.3V supply when the system's SD card was written could cause enough of a supply disturbance to cause the detection of a false particle. Adding significant charge storage and bypass capacitance around the system did reduce but not completely solve the issue.

The solution to this problem was to provide a dedicated 3.3V linear regulator for the AFE circuit. On the WAQM OPC system, this supply draws current from the 5V power rail and outputs a relatively clean 3.3V supply for use with just the analog AFE circuit. With cost in mind, the 1117 series low-dropout regulator [56] was chosen for this role. While not the latest technology, the part is well-suited for this application, being low cost and sourced from several different vendors. The parts will typically show greater than 50dB of input ripple rejection across a fairly wide range of frequencies, removing nearly any supply related noise of concern that would generate false particles in the AFE.

The requirement for a dedicated supply for the AFE is system dependent. If a system is able to provide a supply of sufficient quality that is not disturbed by other activity, then the AFE may be able to run from this directly. The safest path in designing a new system with the WAQM OPC is to provide at least an option for a dedicated supply, so that any supply noise may be addressed by adding the regulator if it is found to be necessary.

Operational Amplifier Choice

Before discussing the details of the AFE circuit implementation, the choice of operational amplifiers is discussed briefly. To keep system costs down, it was desired to build the AFE from single-supply amplifiers to remove the requirement for a negative voltage supply, and to operate from 3.3V to allow sharing of the system power supply if

possible. It was also seen as advantageous to use the same operational amplifier for the transimpedance amplifier and second stage amplifiers. This would allow the use of dual-amplifier and quad-amplifier packages, which would reduce space and lower cost for the circuit. The amplifier should also be low-cost to meet the program requirements.

With this in mind, four main choices for operational amplifier were selected. Table 2 shows these four devices, along with the critical parameters in making the selection. Parameters for a particular device that are in a good range are shown in green, in a concerning range are in yellow, and in a poor range are orange. The LMP7716 wins on many of the operational parameters, but has higher cost. The high common mode input capacitance may also limit bandwidth for some circuit configurations. The LMC6482A and LTC6084 both appear to be reasonable choices, but their low gain bandwidth product will limit their responses at higher frequencies. The OPA2314 was seen as the best choice due to its low cost coupled with a relatively good gain bandwidth

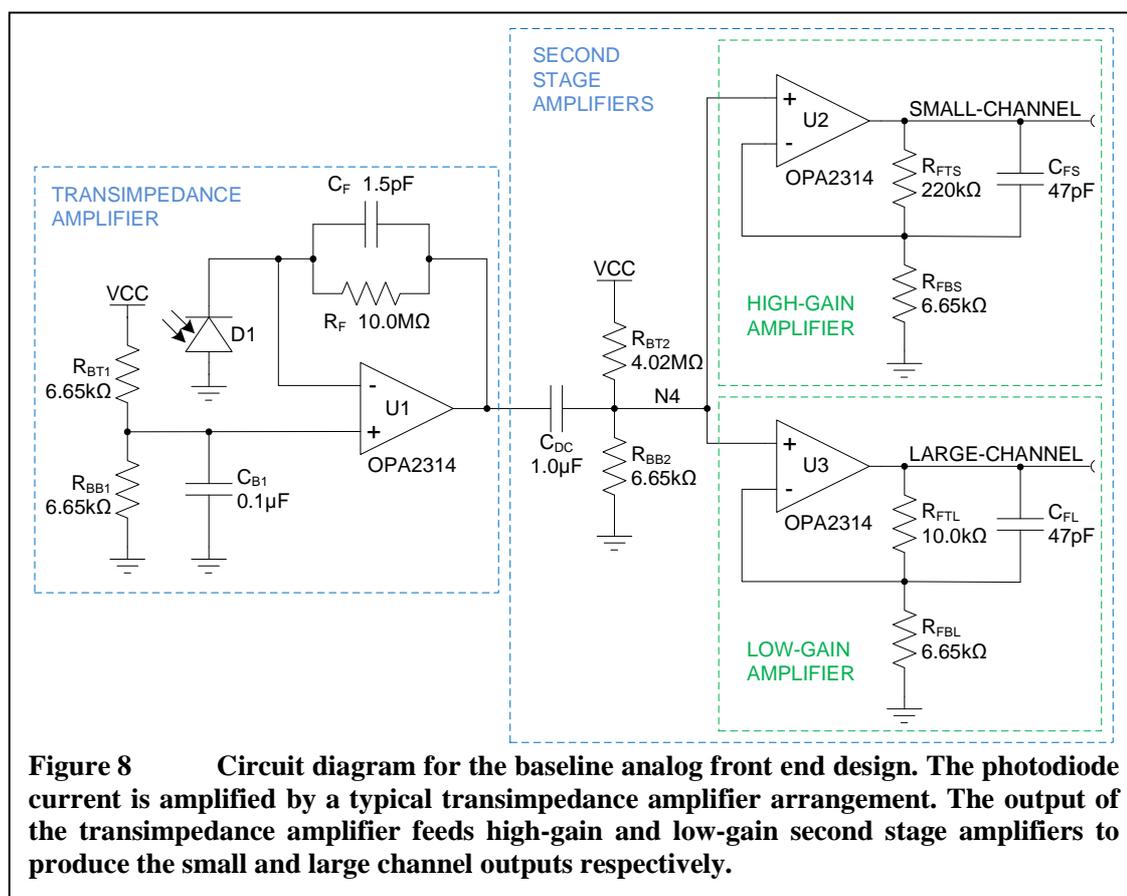
Table 2 Operational amplifier selection parameters for the WAQM OPC. The top four choices are displayed, with positives in green, negatives in yellow, and strong negatives in orange.

Part			LMC6482A	LMP7716	LTC6084	OPA2314
Input Offset Voltage	V_{os}	Typical	0.9 mV	20 μ V	0.30 mV	0.5 mV
		Max	2.7 mV	180 μ V	0.75 mV	2.5 mV
Input Offset Current	i_{os}	Typical	0.01 pA	0.006 pA	0.5 pA	0.2 pA
		Max		0.5 pA	30 pA	10 pA
Input Bias Current	I_b	Typical	0.02 pA	0.05 pA	1 pA	0.2 pA
		Max		1 pA	40 pA	10 pA
Common Mode Rejection Ratio	CMRR	Typical	74 dB	100 dB	80 dB	86 dB
		Limit	64 dB	83 dB	64 dB	70 dB
Power Supply Rejection Ratio	PSRR	Typical	80 dB	100 dB	115 dB	92 dB
		Limit	68 dB	85 dB	94 dB	78 dB
Gain Bandwidth	GBW		1.0 MHz	14 MHz	1.5 MHz	2.7 MHz
Common Mode Input Capacitance	C_{in}		3 pF	15 pF	9 pF	5 pF
Input-Referred Voltage Noise	e_n	@ 1kHz	37 nV/ \sqrt Hz	5.8 nV/ \sqrt Hz	31 nV/ \sqrt Hz	14 nV/ \sqrt Hz
Input-Referred Current Noise	i_n	@ 1kHz	0.03 pA/ \sqrt Hz	0.01 pA/ \sqrt Hz	0.56 fA/ \sqrt Hz	5 fA/ \sqrt Hz
Conditions			$V_+ = 3.0V, V_- = 0V$	$V_+ = 2.5V, V_- = 0V$	$V_+ = 2.5V, V_- = 0V$	$V_s = 1.8V$ to 5.5V
Digikey Cost		1 Unit	\$2.14	\$3.01	\$1.93	\$1.73
Digikey Cost		100 Units	\$1.55	\$1.28	\$1.10	\$0.67
Packages			SO-8	MSOP-8	MSOP-8	SO-8, MSOP-8

product. The input-referred voltage noise is also relatively low, which is of critical importance in keeping the total RMS noise low in the AFE as will be shown in the following subsections.

Baseline AFE Circuit

The baseline AFE circuit is a straightforward implementation of a photodiode transimpedance amplifier followed by a pair of second stage non-inverting amplifiers. The circuit diagram for the baseline is shown in Figure 8. Light enters the AFE at photodiode D1 on the left side of the circuit. Photons impacting the depletion region in the device are converted into hole/electron pairs, which are accelerated across the region by its electric field. This creates a reverse current i_p from cathode to anode in the device.



The amplifier U1 is forced to attempt to match this current through feedback resistor R_F to keep the bias at its inverting input equal to that at its non-inverting input. Ideally, for a constant photodiode current, this will create a voltage at the amplifier output equal to $R_F \times i_p$.

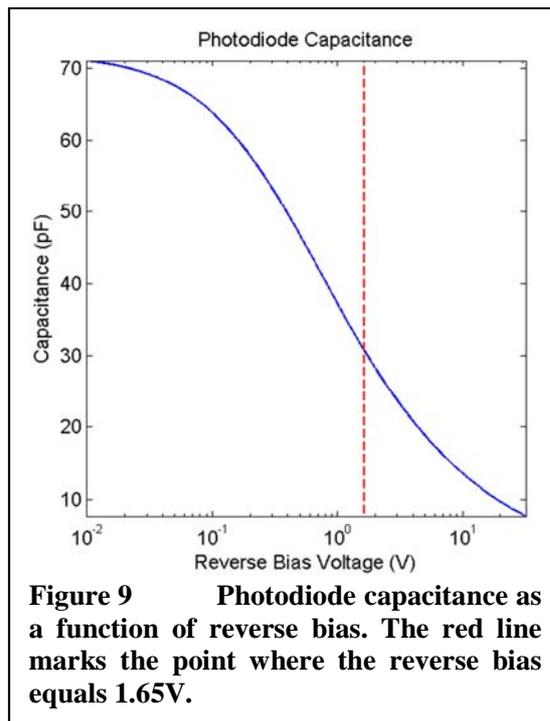
The non-inverting input of the transimpedance amplifier U1 is biased to $V_{CC}/2$, which is nominally 1.65V for the WAQM OPC. A minimum bias at this node is recommended for a single-rail amplifier, which may not operate consistently very near to either of the supply rails. In this case, the non-inverting input is biased at the center of the supply range by the $6.65k\Omega$ R_{BT1} and R_{BB1} . This is done to strike a balance between providing a reverse bias on the photodiode while allowing enough output range for particle-induced pulses. A capacitor C_{B1} is also attached to the non-inverting input with the intent of creating an AC ground at this point, approximating the behavior of a dual-rail amplifier with the non-inverting input attached to ground.

The reverse bias is applied to the photodiode to improve the bandwidth of the circuit. As will be shown below, the photodiode capacitance impacts the transimpedance amplifier bandwidth. Applying a reverse bias to the photodiode reduces the junction capacitance of the photodiode by increasing the size of the depletion region. This has the secondary benefit of increasing the amount of volume in the photodiode available to capture incident photons, increasing the sensitivity of the device. The downside of the reverse bias is that a dark current is established through the photodiode, introducing increased noise into the system. The photodiode's junction capacitance C_j is defined as

$$C_j = \frac{C_{j0}}{\sqrt{1 + \frac{V_R}{\phi_B}}} \quad 4.1$$

where C_{j0} is the photodiode capacitance at zero bias, ϕ_B is the built-in voltage of the diode junction, and V_R is the reverse bias voltage.

Figure 9 shows a plot of the Osram BPW34 photodiode [57] used in the WAQM OPC. The zero bias capacitance is 72pF for this device, with a built-in voltage of 365mV. The capacitance falls with increasing reverse bias, reaching a value of approximately 31pF at 1.65V, which is the target bias point for the baseline AFE. This reduces the junction capacitance to $\frac{1}{2}$ the value of the zero bias capacitance.

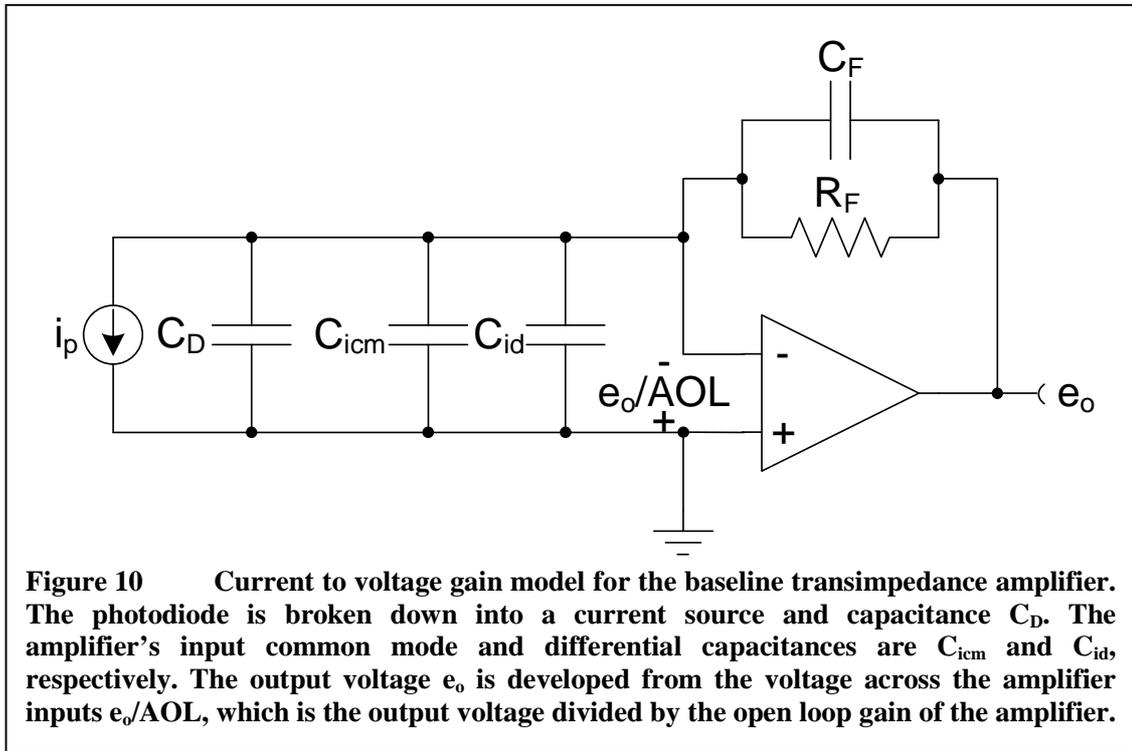


While the reverse bias could be increased above 1.65V, doing so would reduce the range of the transimpedance amplifier. Since the bias at the non-inverting input translates directly to the output of the amplifier, the available output range with reverse

bias is $V_{CC} - V_R$, or 1.65V. This means that the output can only swing between V_{CC} and 1.65V, effectively halving the intensity range of possible particle pulses that can be recreated on the output of the transimpedance amplifier. As will be shown below, this reduction does not impact the performance of the AFE, since the range is further limited by the second stage amplifiers.

The feedback compensation capacitor C_F is required to stabilize the amplifier. The photodiode and amplifier input capacitances act as an impedance to ground that begins to roll off at higher frequencies. This impedance acts to create a voltage divider with the feedback resistor R_F , effectively inserting a zero into the response of the system. Left unchecked, this would cause the circuit to oscillate. Inserting the compensation capacitor rolls off the feedback impedance, inserting a pole into the response that balances the input capacitance zero and stabilizes the system.

To find the gain of the system, each of the amplifiers is examined separately. The transimpedance amplifier current to voltage (I-to-V) gain can be found through node analysis of the circuit. The model for gain analysis is shown in Figure 10. The photodiode is broken down into a current source and capacitance C_D . The amplifier's common mode and differential input capacitances are C_{icm} and C_{id} , respectively. The non-inverting input and differential input capacitance are shown attached to ground, since the bypass capacitance C_{B1} approximates AC ground at the non-inverting input. The output voltage e_o is divided by the open loop gain AOL to find the voltage across the amplifier inputs e_o/AOL .



The node analysis sums the currents at the node attached to the inverting input of the amplifier as

$$\frac{e_o + \frac{e_o}{AOL}}{Z_F} + \frac{e_o}{Z_i} - i_p = 0 \quad 4.2$$

where the first term represents the current through the feedback network, the second term the current through the photodiode capacitor, and the third term the photodiode current.

The impedance of the feedback network is represented by

$$Z_F = R_F \parallel Z_C = \frac{R_F}{1 + sR_FC_F} \quad 4.3$$

where $s = j\omega$, and the impedance to ground at the node is represented by

$$Z_i = \frac{1}{sC_i} \quad 4.4$$

where C_i is the sum of the photodiode and amplifier input capacitances $C_i = C_D + C_{icm} + C_{id}$. Solving for the current to voltage gain leads to

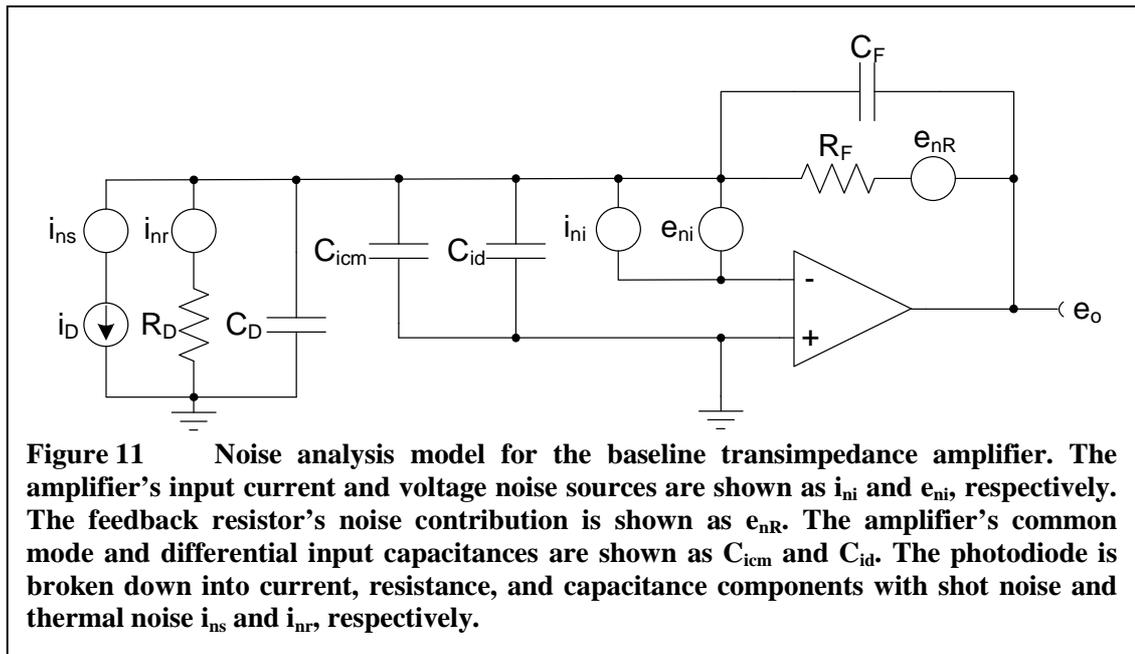
$$\frac{e_o}{i_p} = \frac{1}{\frac{1}{Z_F} + \frac{1}{AOL} \left(\frac{1}{Z_F} + \frac{1}{Z_i} \right)} \quad 4.5$$

The open loop gain AOL can be approximated by the single pole response $\omega_c/j\omega$ for most of the useful frequency range, where ω_c is the unity gain crossover frequency of the amplifier. Substituting in the resistances and capacitances for the impedances gives the I-to-V gain in terms of the individual components in the circuit as

$$\frac{e_o}{i_p} = \frac{R_F}{1 + s \left(R_F C_F + \frac{1}{\omega_c} \right) + s^2 \left(\frac{R_F (C_F + C_i)}{\omega_c} \right)} \quad 4.6$$

Note the second order component of the I-to-V gain, which will cause the gain to roll off at 40dB per decade at higher frequencies.

For noise analysis of the transimpedance amplifier, the methods discussed by



Graeme [58] are employed. The model used for noise analysis is shown in Figure 11. The photodiode is modeled as a current source, shunt resistance R_D , and capacitance C_D . Photodiode noise sources due to reverse biasing are included as i_{ns} and i_{nr} for the shot and thermal noise, respectively. The amplifier common mode and differential input capacitances are shown as C_{icm} and C_{id} , respectively. As with the gain model, the non-inverting input is connected to ground as the bypassing at this node approximates an AC ground. Amplifier noise sources i_{ni} and e_{ni} are the input current and voltage noises, respectively. The thermal noise of the feedback resistor R_F is shown as E_{nR} .

The feedback factor β is defined as the fraction of the amplifier's output signal that is fed back to the input. This relationship is established by voltage divider created by the feedback network and the input capacitance consisting of the photodiode and amplifier input capacitances. For an arrangement without the feedback capacitance C_F , the feedback factor is

$$\beta = \frac{X_{Ci}}{R_F + X_{Ci}} = \frac{1}{1 + sR_F C_i} \quad 4.7$$

The reciprocal of this, $1/\beta$, will be the gain demand on the amplifier, and contains a zero at frequency

$$f_z = \frac{1}{2\pi R_F C_i} \quad 4.8$$

which is the cause for instability and the need for the feedback capacitance C_F . Adding the feedback capacitance results in a feedback factor

$$\beta = \frac{X_{Ci}}{R_F \parallel X_{CF} + X_{Ci}} = \frac{1 + sR_F C_F}{1 + sR_F (C_i + C_F)} \quad 4.9$$

The reciprocal gain demand now has a zero at frequency

$$f_z = \frac{1}{2\pi R_F(C_i + C_F)} \quad 4.10$$

and a pole to counteract it at

$$f_p = \frac{1}{2\pi R_F C_F} \quad 4.11$$

The noise voltage gain A_{ne} will follow $1/\beta$ until it is limited by the open loop gain of the amplifier. Ignoring the open loop gain limit, the noise voltage follows

$$A_{ne} = \frac{1}{\beta} = \frac{1 + sR_F(C_i + C_F)}{1 + sR_F C_F} \quad 4.12$$

Including the open loop gain into the equation adds an additional pole in the gain equation as

$$A_{ne} = \frac{[1 + sR_F(C_i + C_F)]}{(1 + sR_F C_F)(1 + s/\omega_i)} \quad 4.13$$

where ω_i is the intersection frequency of the $1/\beta$ gain demand curve and the open loop gain. This additional pole causes the noise voltage gain to fall off at 20dB/decade with the open loop gain.

The gain curves discussed are plotted together in Figure 12. The dark blue curve is the open loop gain of the OPA2314 amplifier, which crosses the unity gain point near the 2.7MHz gain-bandwidth product of the part. The green line is the $1/\beta$ gain demand curve. The red line is the noise gain, which is a combination of the gain demand curve and the open loop gain. The cyan line is the current to voltage gain, which falls at 40dB/decade soon after the open loop gain limit is reached. Some of the key frequencies are marked with vertical red lines and are named at the bottom of the plot. The frequency f_z shows where the photodiode and amplifier input capacitance creates a zero in the noise gain. Likewise, f_p shows where the addition of the feedback capacitor adds the pole that rolls off the noise gain. The frequency f_i is the point where the gain demand $1/\beta$ crosses the open loop gain. Finally, f_c marks the unity gain crossing point of the amplifier's open loop gain.

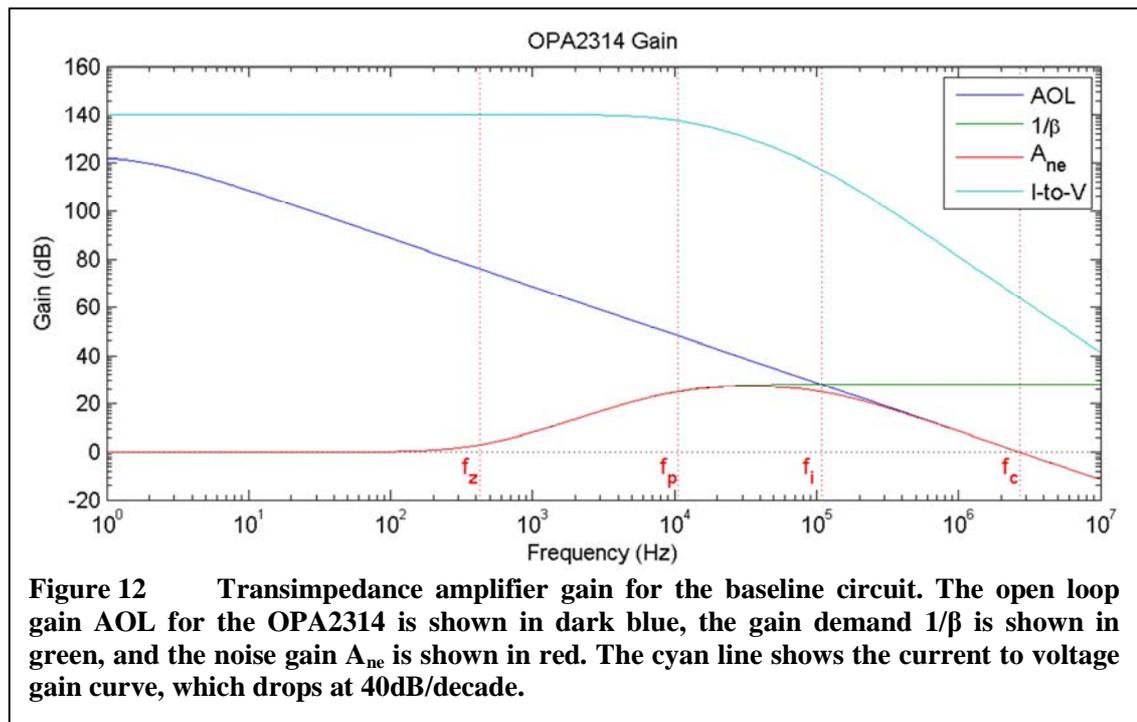


Figure 12 Transimpedance amplifier gain for the baseline circuit. The open loop gain AOL for the OPA2314 is shown in dark blue, the gain demand $1/\beta$ is shown in green, and the noise gain A_{ne} is shown in red. The cyan line shows the current to voltage gain curve, which drops at 40dB/decade.

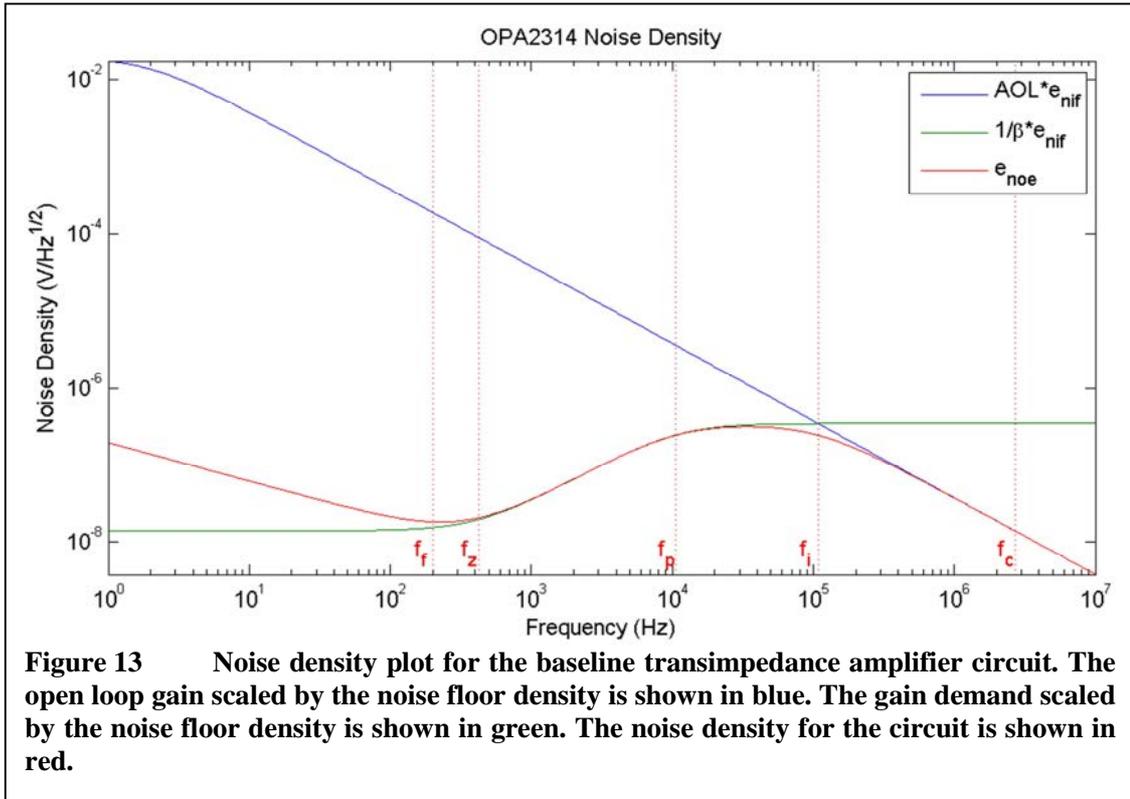
With the noise gain in hand, the noise density follows by multiplying the noise gain by the amplifier's input voltage noise density e_{ni} . This yields the equation

$$e_{noe} = A_{ne} e_{ni} = A_{ne} e_{nif} \sqrt{(\omega_f + s)/s} \quad 4.14$$

where ω_f is the frequency of the 1/f noise floor corner of the amplifier and e_{nif} is the input voltage noise density at the floor. Expanding A_{ne} leads to the final noise density equation

$$e_{noe} = \frac{[1 + sR_F(C_i + C_F)] \sqrt{(\omega_f + s)/s}}{(1 + sR_F C_F)(1 + s/\omega_i)} e_{nif} \quad 4.15$$

Figure 13 shows the voltage noise density e_{noe} for the baseline transimpedance amplifier circuit is plotted in red. The impact of the 1/f noise can be seen at lower frequencies, where the curve falls at 20dB/decade until reaching the floor corner at f_f . At this point, the voltage noise density is at the floor value of e_{nif} . Shortly afterwards, the zero in the noise gain at f_z causes the density to rise at 20dB/decade, until leveling off at f_p . The density falls again once the open loop gain is reached, with a corner near f_i .



Calculating an RMS noise for the transimpedance amplifier circuit would require evaluating the integral of the square of Equation 4.15. Graeme [58] suggests a method for simplifying this task by splitting the voltage noise density curve into five regions divided by f_f , f_z , f_p , and f_i . This results in region 1 from DC to f_f , region 2 from f_f to f_z , region 3 from f_z to f_p , region 4 from f_p to f_i , and region 5 from f_i to infinity. The RMS noise for each region can be calculated, and then combined in a final result.

Following the method described in Graeme, the first region captures the RMS noise where the amplifier's $1/f$ noise is dominant. The RMS noise for region 1, E_{noe1} , is

$$E_{noe1} = e_{nif} \sqrt{f_f \ln \frac{f_f}{f_1}} \quad 4.16$$

where f_l is ideally a DC value of 0 Hz. Practically, this leads to an indefinite response and a reasonably small value must be chosen instead. In the calculations for the WAQM OPC, a value of 0.01Hz is used.

Region 2 between f_f and f_z sits at the noise floor level e_{nif} , and its RMS noise E_{noe2} is

$$E_{noe2} = e_{nif} \sqrt{f_z - f_f} \quad 4.17$$

Region 3 between f_z and f_p includes the portion of the noise gain where the response rises at 20dB/decade, and its RMS noise E_{noe3} is

$$E_{noe3} = \left(\frac{e_{nif}}{f_z} \right) \sqrt{\frac{(f_p^3 - f_z^3)}{3}} \quad 4.18$$

Region 4 between f_p and f_i include the plateau where the feedback capacitance has leveled the response, and its RMS noise E_{noe4} is

$$E_{noe4} = \left(1 + \frac{C_i}{C_f} \right) e_{nif} \sqrt{f_i - f_p} \quad 4.19$$

The final region covers everything above the interception point of the gain demand and the open loop gain at f_i . The RMS noise of this region, E_{noe5} , is

$$E_{noe5} = e_{nif} f_c \sqrt{\frac{1}{f_i}} \quad 4.20$$

Additional poles in the amplifier's response above the unity gain crossover point are ignored since the declining level of the noise gain minimizes their impact on the RMS noise.

Combining the RMS values for each region will yield the total RMS noise due to the amplifier's input voltage noise as

$$E_{noe} = \sqrt{E_{noe1}^2 + E_{noe2}^2 + E_{noe3}^2 + E_{noe4}^2 + E_{noe5}^2} \quad 4.21$$

This can be combined with other noise sources to create the total RMS output noise for the transimpedance amplifier. Other noise sources found in the model of Figure 11 include the thermal noise of the feedback resistor, the amplifier's input current noise, and the photodiode current noise. The feedback resistor thermal noise E_{noR} can be calculated with

$$E_{noR} = \sqrt{2KTR_F\pi BW_t} \quad 4.22$$

where K is the Boltzmann constant 1.38×10^{-23} , T is the temperature in Kelvin, and $BW_t = 1/2\pi R_F C_F$, which rolls off the noise component according to the feedback network. Likewise, the amplifier's input current noise E_{noi} can be calculated with

$$E_{noi} = R_F \sqrt{q\pi BW_t I_B} \quad 4.23$$

where I_B is the input bias current of the amplifier and q is the electron charge 1.602×10^{-19} . Similar to the amplifier's input current noise, the noise component due to the photodiode shot noise E_{nops} can be calculated with

$$E_{nops} = R_F \sqrt{q\pi BW_t I_p} \quad 4.24$$

where I_p is the photodiode dark current. The photodiode thermal noise E_{nopt} can be found with

$$E_{nopt} = R_F \sqrt{\frac{2KTBW_t}{R_D}} \quad 4.25$$

where R_D is the photodiode shunt resistance. Note that both current noise sources and thermal noise sources are applied directly to the output through the feedback network. Finally, the noise sources may be combined into a single term for the total RMS output noise for the transimpedance amplifier circuit E_{no} as

$$E_{no} = \sqrt{E_{noe}^2 + E_{noR}^2 + E_{noi}^2 + E_{nops}^2 + E_{nopt}^2} \quad 4.26$$

For the WAQM OPC baseline circuit, the RMS noise calculations for the various noise components are shown in Table 3. Each of the components is listed, along with the total RMS noise E_{no} . The magnitude of the total RMS noise is relatively small, at $168\mu\text{V}$. Note that nearly all of the noise is attributed to the amplifier's input voltage noise, with the next nearest contributor being the thermal noise due to the feedback resistor. This is an interesting conclusion as the literature debates the merits of having a reverse bias on the photodiode due to the tradeoff between capacitance and increased noise [58], [59]. In the case of the WAQM OPC, the RMS noise due to the photodiode dark current is an order of magnitude less than that due to the amplifier's input voltage noise.

The output of the transimpedance amplifier is fed into the second stage amplifiers through a high-pass filter consisting of DC blocking capacitor C_{DC} . The DC blocking capacitor is required to remove the DC bias on the output of the transimpedance

E_{noe}	E_{noR}	E_{noi}	E_{nops}	E_{nopt}	E_{no}
$1.60 \times 10^{-4} \text{V}$	$5.01 \times 10^{-5} \text{V}$	$2.31 \times 10^{-6} \text{V}$	$1.63 \times 10^{-5} \text{V}$	$1.56 \times 10^{-6} \text{V}$	$1.68 \times 10^{-4} \text{V}$

amplifier. Since the transimpedance amplifier's non-inverting input is biased to 1.65V, the output is also at this voltage when no current is flowing through the photodiode. Without removing this bias, the second stage amplifiers would saturate with a gain of only 2.

The far side of the DC blocking capacitor is connected to the non-inverting inputs of the two second stage amplifiers, and is biased by a pair of resistors, R_{BT2} and R_{BB2} . Most of the interesting signal coming from particle-caused pulses will have a positive voltage relative to the bias, so setting this bias as low as possible will deliver the most usable range from the amplifier. However, the bias must be high enough to keep the non-inverting input of the second stage amplifiers set to a level that is above their maximum offset voltages. If the bias is set to a point lower than this, considerable unit-to-unit variation in the response of the second stage amplifiers may be seen due to changes in response specific to the unit's offset voltage. In the case of the WAQM OPC AFE, this bias is set to just over 5mV, which is comfortably higher than the maximum offset voltage of 2.5mV for the OPA2314 operational amplifier.

The gain of the DC blocking capacitor A_{CDC} can be found from node analysis as

$$A_{CDC} = \frac{\frac{1}{Z_{DC}}}{\frac{1}{Z_{DC}} + \frac{1}{R_{BT2}} + \frac{1}{R_{BB2}}} = \frac{R_{BT2}R_{BB2}}{R_{BT2}R_{BB2} + \frac{R_{BT2}}{sC_{DC}} + \frac{R_{BB2}}{sC_{DC}}} \quad 4.27$$

where Z_{DC} is the impedance of the DC blocking capacitor C_{DC} , R_{BT2} is the bias resistor tied to VCC, and R_{BB2} is the bias resistor tied to ground. The gain of the DC blocking circuit will rise to a pole at a relatively low frequency and will flatten to zero dB for higher frequencies.

The two second stage amplifiers operate as non-inverting arrangements for the small and large channels, and will apply further gain to the output of the transimpedance amplifier. The small channel employs a higher gain than the large channel, since it is attempting to detect the smallest particles. The large channel employs a lower gain to provide a large range for sizing particles.

Looking at the small, high-gain channel, the gain A_{VOS} can be approximated by the voltage divider formed by R_{FTS} and R_{FBS} at low frequencies as

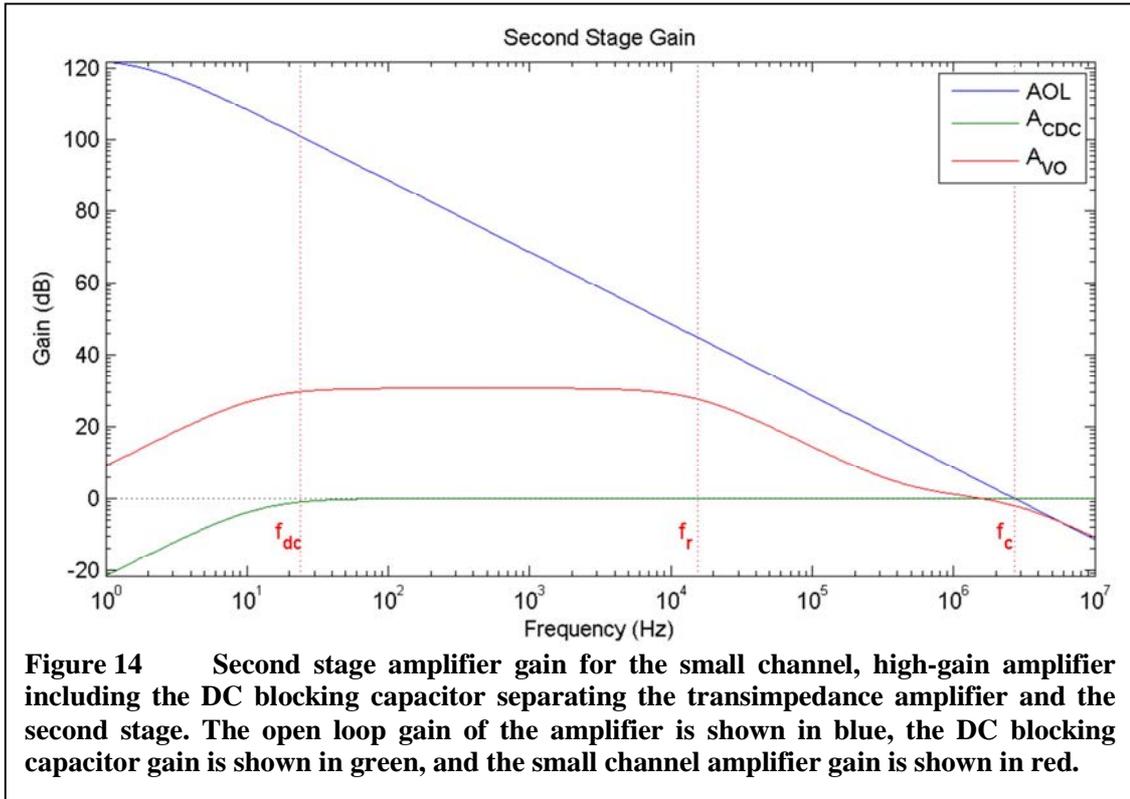
$$A_{VOS} = \frac{R_{FTS} + R_{FBS}}{R_{FBS}} \quad 4.28$$

For the small channel amplifier with $R_{FTS} = 220\text{k}\Omega$ and $R_{FBS} = 6.65\text{k}\Omega$, this gives a value of $A_{VOS} = 34.1$. This particular gain value was chosen through experimentation, balancing the channel's noise with a desire to maximize the gain for the smallest particles. Taking into account the capacitance of the compensation capacitor C_{FS} and the amplifier common mode input capacitance C_{icm} yields

$$A_{VOS} = \frac{Z_{FTS} + Z_{FBS}}{Z_{FBS}} = 1 + \frac{R_{FTS} + sR_{FTS}R_{FBS}C_{icm}}{R_{FBS} + sR_{FTS}R_{FBS}C_{FS}} \quad 4.29$$

where Z_{FTS} is the impedance of the resistor R_{FTS} in parallel with the feedback compensation capacitor C_{FS} , and Z_{FBS} is the impedance of resistor R_{FBS} in parallel with the amplifier input common mode capacitance C_{icm} .

The second stage gain for the small channel, high-gain amplifier is shown in Figure 14. The blue line is the open loop gain of the operational amplifier, the green line is the gain of the DC blocking capacitor, and the red line is the gain of the small channel amplifier including the impact of the DC blocking capacitor. Looking at the figure, it is



clear that the amplifier is overcompensated with the value of 47pF for the compensation capacitor. While it is possible that a higher bandwidth could be achieved with a smaller compensation capacitor, this particular circuit arrangement has functioned well and was not changed in the latest iteration of the WAQM OPC. Further study of the amplifier circuits may conclude that more sensitivity can be wrested from the circuit by lowering or removing the compensation capacitance altogether. The noise response of such a modification should be carefully studied to make sure that an increase in gain is not just amplifying noise.

In the figure, f_{dc} represents the corner frequency at which the DC blocking capacitor becomes a low impedance to the system. The frequency f_r indicates the point at which the compensation capacitor rolls off the gain. The open loop gain crosses the unity gain point at f_c . The gain of the small channel second stage amplifier drops at 20dB per

decibel after f_r , then flattens out as the gain approaches unity. The response rolls off again as it approaches the open loop gain of the amplifier near f_c .

The gain A_{VOL} of the large channel, low-gain amplifier can be approximated at low frequencies by

$$A_{VOL} = \frac{R_{FTL} + R_{FBL}}{R_{FBL}} \quad 4.30$$

For the large channel amplifier with $R_{FTL} = 10.0\text{k}\Omega$ and $R_{FBL} = 6.65\text{k}\Omega$, this gives a value of $A_{VOL} = 2.5$. As with the small channel, this particular gain value was chosen through experimentation, attempting to maximize the range of the channel while ensuring that particles large enough to be considered for sizing would produce a reasonable amount of signal.

The total RMS noise of the full AFE may now be considered, taking the value at the output of the transimpedance amplifier and applying the second stage gain. Only the small channel is considered since the gain is much larger than that of the large channel. While the frequency response of the second stage will help to limit output noise, the worst case is approximated as the low frequency gain of Equation 4.28. The input noise sources of the second stage are also ignored, since they will be orders of magnitude lower than the amplified total RMS noise out of the transimpedance amplifier.

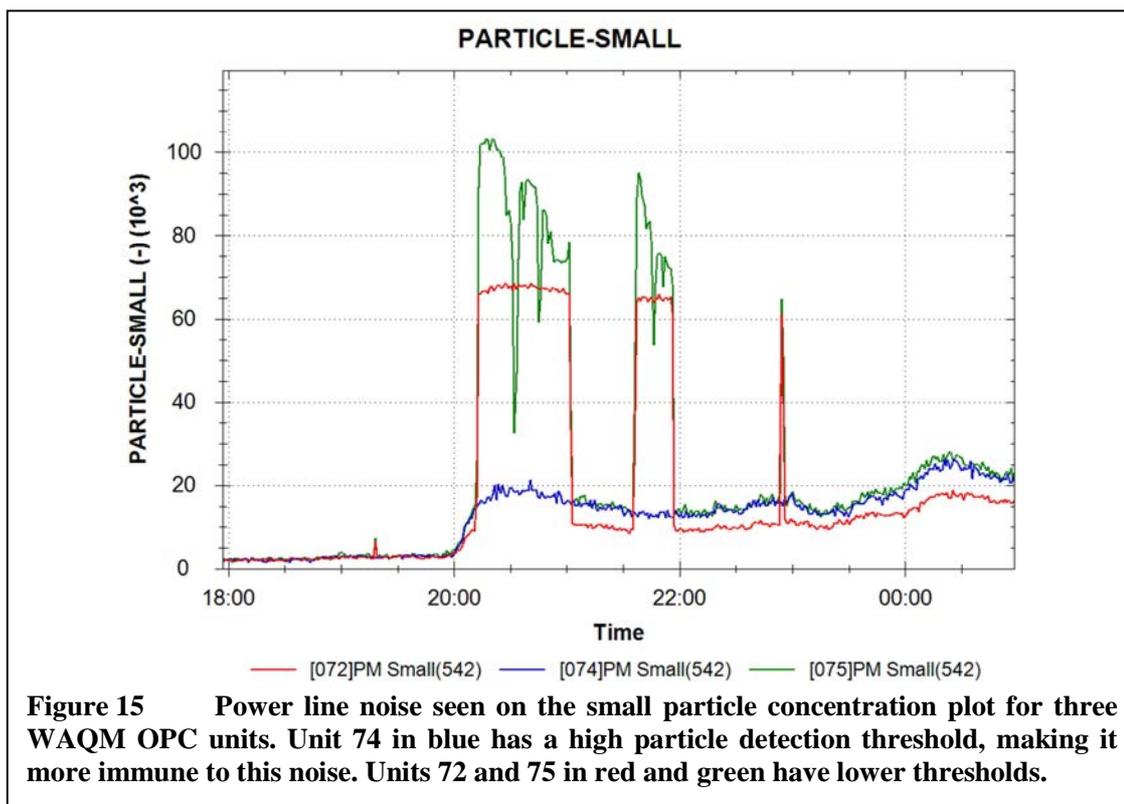
Applying the small channel's low frequency gain of 34.1 to the total RMS noise of the transimpedance amplifier yields a value of 5.73mV of total RMS noise for the system. Converting this to a peak-to-peak value, the total RMS noise is assumed to be Gaussian and is multiplied by 6.6 to cover 99.9% of the noise. This yields a peak-to-peak noise of 37.8mV. Since the main concern for detection of small particles is in being able

to distinguish deviation above an average value, this peak-to-peak noise is halved to 18.9mV. While undesirable, the value does represent less than 3 least significant bits of error to the system with the 10 bit ADC used in the WAQM OPC. This error is well below the typical threshold set for detecting particles by the system microcontroller.

Power Line Noise Coupling

Shortly after building the first few WAQM OPC units, testing and trial runs began in the laboratory and the field. Soon afterwards, it was noticed that the OPC appeared to be susceptible to some form of environmental electrical noise that would cause elevated particle counts for varying periods of time. Some events could last for an hour or more, but would start and stop suddenly. Other events appeared to be very short duration, lasting perhaps a minute or less. The event severity appeared to be dependent on location, though at first it could not be tied directly to any particular activity by the unit or other electrical devices in the area. Susceptibility to the noise appeared to be dependent on the method used to power the unit. Units powered by battery did not show evidence of the noise. This was also the case for units that had a direct connection to earth ground, say through a debug connection to a grounded computer.

Figure 15 shows such events in a plot of small particle concentration, as captured by three WAQM OPC units located in the same residential household. Units 72 and 75 appear to change particulate concentration rapidly, at times more than doubling the concentration between successive readings. Unit 74 appears to be unaffected by these events, and continues to report a relatively smooth profile of particle concentration over time. As will be discussed in detail in later chapters, each unit detects individual particles based on a calibrated threshold value that is dependent on many factors in the system.



Units with higher power lasers or better laser alignment tend to have a higher threshold than units with poorer hardware quality. The three units have different particle detection thresholds, with Unit 74 having a relatively high threshold and Units 72 and 75 having relatively low thresholds. As a result, Unit 74 does not appear to be susceptible to the noise that causes the particulate concentration readings of the other two units to jump up and down by large magnitudes.

The WAQM OPC, as with many low-power systems, was designed to run from a low-cost AC to DC wall transformer. The source of the noise appeared to be on the AC mains powering the wall transformer and was somehow making its way into the transimpedance amplifier of the OPC. The noise would cause a pulse on the transimpedance amplifier, which would then be counted as a particle. Figure 16 shows an example of this noise on both the hot side of the AC mains in blue and on the +3.3V rail

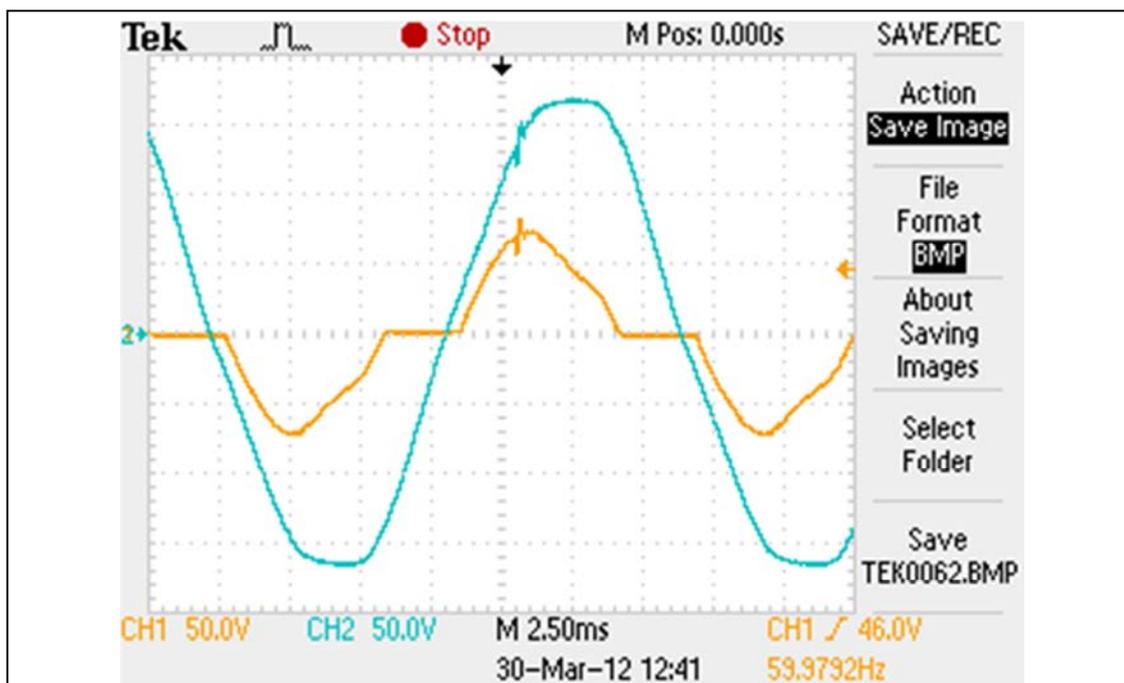


Figure 16 Power line noise on the device and at the AC power connection. The yellow trace is connected to the +3.3V power rail of the WAQM OPC and the blue trace is connected to the hot side of AC power where the wall transformer is connected. The oscilloscope ground is not attached to the WAQM OPC, which is powered by a two-prong wall transformer. Noise likely caused by an SCR is visible on the rising edge of both waveforms.

of the WAQM OPC in yellow. The small, high-frequency pulse near the positive peak of both waveforms is likely caused by a silicon controlled rectifier (SCR) or some other misbehaving load on the AC mains.

This particular measurement was taken without a ground connection between the oscilloscope and the WAQM OPC, since grounding the unit would make the noise issue disappear. In this state, the WAQM OPC has no direct connection to earth ground through its two-pronged wall transformer. Its isolated ground appears to oscillate at 60Hz with the voltage on the AC mains relative to the earth ground of the oscilloscope. Note that the voltage seen at the +3.3V rail relative to the oscilloscope is approximately 150V peak-to-peak, though the voltage only appears to deviate significantly at the peaks of the

voltage on the AC mains. This is likely due to the deviation only occurring when the rectifier circuit in the wall transformer is conducting. While the exact mechanism for this signal reaching the WAQM OPC device was not found, it is assumed that parasitic properties of the components in the wall transformer may be passing enough current to charge and discharge the isolated ground of the system relative to earth ground.

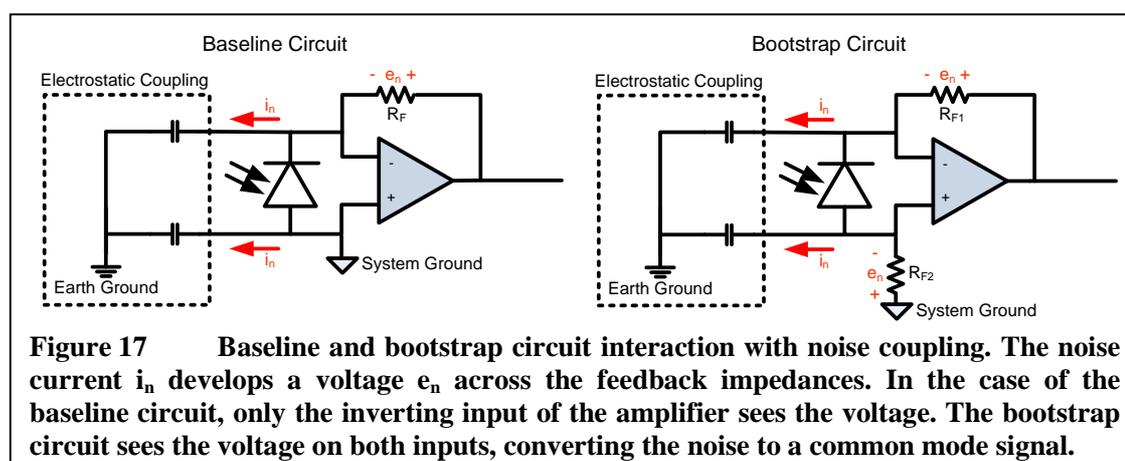
While the original noise sources were not identified directly, two different behaviors were seen. What is referred to as SCR noise tends to cause an elevated particle count for a relatively long duration of time, presumably as long as the SCR is turned on. The original SCR noise source was not located, but similar behavior was reproduced in the laboratory. A simple dimmer switch used to control a heater in the original OPC calibration system would cause an overwhelming amount of this noise, rendering devices running from wall transformers unusable for accurate particle detection.

The second type of noise is much shorter in duration, but can be seen periodically throughout the day. Based on persistent monitoring of the AC mains, it is believed that this noise is due to smart metering communication on the AC mains. In particular, the power meter at the residence where the noise was first seen uses a form of communication known as TWACSTM [60]. Relatively high frequency signaling can be observed from time to time on the power mains at this home that matches the signature described by this communication method.

Given that the noise sources on the power mains appear to be relatively common, it was desired to somehow defeat this noise so that it would not pollute the results collected by the WAQM OPC. The obvious solutions to the problem had negative implications to the system. Shielding the OPC provided some reduction in the noise seen

at the transimpedance amplifier, but did not remove it all. This also added cost and complexity to the unit, which ran against the main goals of the program. Raising the detection threshold provided some relief to the noise since it would typically appear as relatively small magnitude pulses, but this also meant throwing away the smallest particles that the OPC was capable of detecting. Grounding the system cleaned up all of the noise issues, but meant providing a separate connection to earth ground. Even wall transformers with three prongs do not connect the earth ground through to the DC ground for safety reasons. Adding a separate wire to a device meant to be easily connected to power was not acceptable.

Looking deeper into the mechanism for the generation of the pulses in the transimpedance amplifier provided a solution. As described in Graeme [58], the high impedance to ground at the inverting input of the transimpedance amplifier provides a source for noise susceptibility. While this reference describes noise coupled into the system, it is clear that noise may also be coupled out of the system. Figure 17 shows a graphic representation of the system. The circuit on the left shows a simplified view of the baseline transimpedance amplifier as described in the preceding section. A noise



voltage on the system couples electrostatically to an external earth ground. This effectively forms capacitors between the various points on the circuit and the external earth ground through which a noise current i_n will flow. In the case of the high impedance inverting input of the baseline system, this current develops a voltage e_n across the large feedback resistor R_F . Even for very small amounts of electrostatic coupling, the voltage e_n may reach hundreds of microvolts. For example, a 1 square millimeter plate capacitor at 1 centimeter will have a capacitance of approximately 1×10^{-15} farad, or one femtofarad. At 60 Hz, this will have an impedance of $2.7T\Omega$, a very high impedance. The voltage divider formed by this impedance with the $10M\Omega$ resistor R_F when 150V is placed across it will result in approximately $600\mu V$ potential difference between the two amplifier inputs. This would saturate the output of the AFE. Given the much smaller amount of noise seen on the AFE output, the electrostatic coupling must be even less than this.

Balancing the voltages seen by the amplifier inputs will allow the circuit to take advantage of the high common mode rejection of the amplifier. The circuit on the right of Figure 17 shows the bootstrap configuration, with separate current to voltage conversion resistors R_{F1} and R_{F2} on the two amplifier inputs. Assuming the two resistances are perfectly matched, the voltage developed across each by equal noise currents will be identical. This common mode voltage will then meet the common mode rejection of the amplifier, which is 70dB at a minimum for the OPA2314 amplifier used in the WAQM OPC.

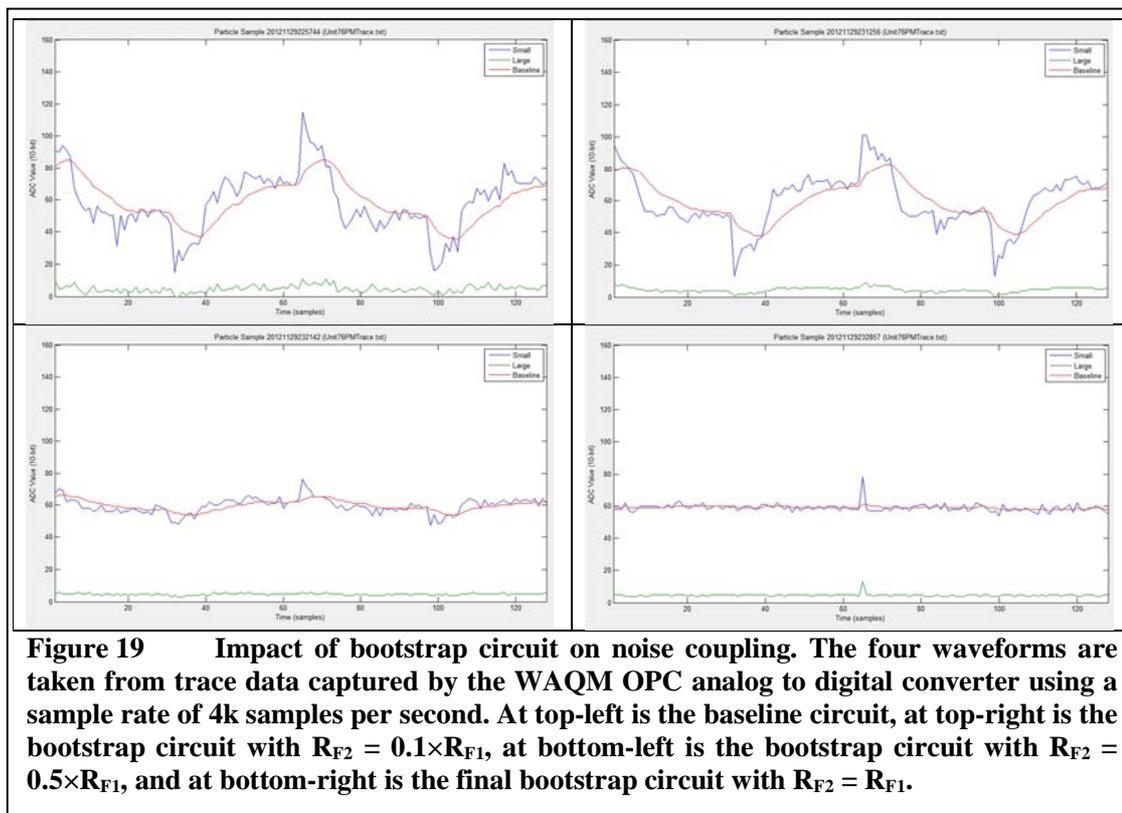
Testing this theory in the laboratory was accomplished by placing a WAQM OPC on top of a grounded copper sheet while it was powered from a wall transformer. This arrangement is shown in Figure 18, and was perhaps the situation that generated the most



Figure 18 Testing noise coupling on a grounded copper sheet. The WAQM OPC unit is powered through a two-prong wall transformer with isolated ground.

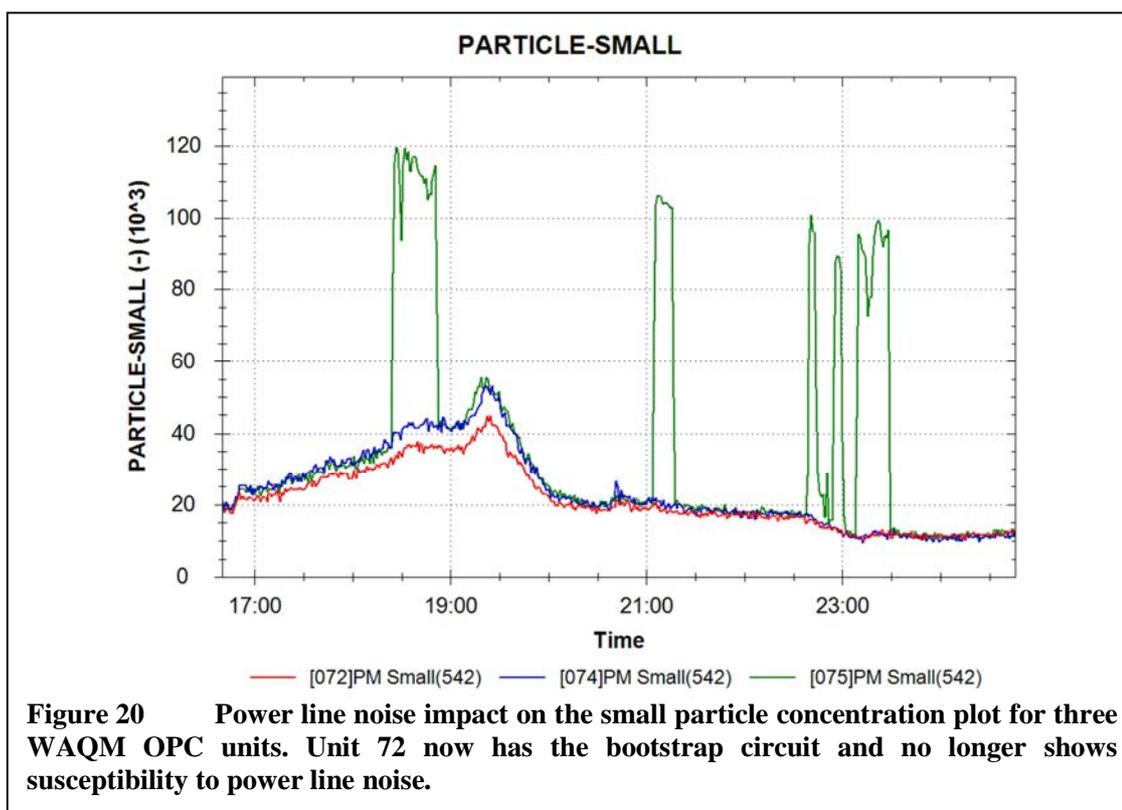
noise on the WAQM OPC transimpedance amplifier. Similar situations such as placing the device on top of a metal shelf or on a refrigerator showed poor noise performance in field testing.

Test results for a migration from the baseline circuit to the bootstrap circuit are shown in Figure 19. These four traces were taken by the analog to digital converter on the WAQM OPC, operating at a sampling frequency of 4k samples per second. This frequency was chosen to allow the viewing of nearly two full cycles of 60 Hz noise. The blue line is the output of the small channel, the green line is the output of the large channel, and the red line is a running average of the small channel used in particle detection. The plot at top-left is taken from the baseline circuit. The other three plots are taken from bootstrap configurations, but with different ratios of R_{F2} relative to R_{F1} . The plot at top-right has $R_{F2} = 0.1 \times R_{F1}$, the plot at bottom-left has $R_{F2} = 0.5 \times R_{F1}$, and the plot at bottom-right has $R_{F2} = R_{F1}$. The noise on the small channel seen in the baseline circuit



is quite bad, with a peak-to-peak amplitude of over 200mV. As the ratio between R_{F1} and R_{F2} approaches unity, the noise disappears. It is clear, however, that the resistances must match fairly closely for this to work well. Note that the pulse in the bottom-right plot is from an actual particle passing through the laser.

Testing the bootstrap circuit in the field shows marked improvement in rejecting power line noise from SCRs or smart power monitoring systems. Figure 20 shows the results after changing the transimpedance amplifier of Unit 72 to the bootstrap configuration. While Unit 75 continues to show noise events, Unit 72 tracks with the high-threshold Unit 74. Since implementing the bootstrap circuit on a large number of WAQM OPC systems, this type of noise susceptibility has no longer been observed. As will be shown in the following section, the coupled power line noise immunity does come



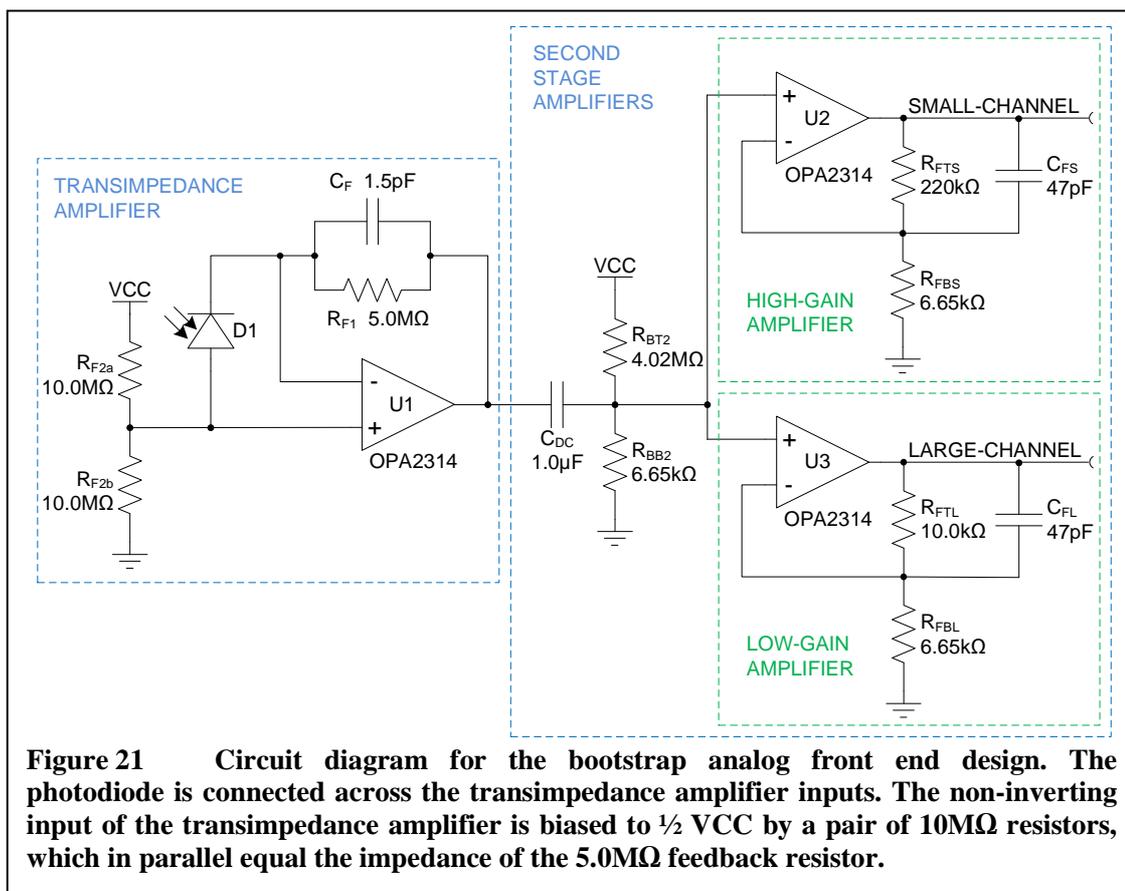
at a cost of increased total RMS noise in the amplifier, but this is a tradeoff decision easily made in the case of the WAQM OPC.

Bootstrap AFE Circuit

The bootstrap version of the AFE circuit is a modification of the baseline circuit intended to defeat most noise coupling described in the previous section. The circuit changes are relatively minor, and involve only the transimpedance amplifier. These changes are shown in the bootstrap circuit in Figure 21. The term “bootstrap” refers to the connection of the photodiode across the terminals of the amplifier, using the amplifier output to remove the voltage swing across the device.

The objective for the bootstrap transimpedance amplifier is to equalize the impedance to AC ground at the inputs of the operational amplifier. To accomplish this,

the relatively low impedance connection to AC ground at the non-inverting input of the amplifier of the baseline circuit is removed. The photodiode is now connected directly between the two inputs of the operational amplifier, rather than being reverse biased as it is in the baseline circuit. A bias at the non-inverting input is still required, since the amplifier is single-rail. This is formed with a pair of high-impedance resistors, R_{F2a} and R_{F2b} , that are each double the size of the feedback resistor R_{F1} from the output of the amplifier to the inverting input. VCC is considered to be an AC ground, since at even moderate frequencies the impedance between VCC and ground is low. At these frequencies, R_{F2a} and R_{F2b} are effectively connected in parallel to AC ground, forming an impedance to AC ground that is equivalent to the impedance R_{F1} . The connection that R_{F1} makes between the inverting input and the amplifier output may also be considered to be AC ground due to the relatively low impedance of the operational amplifier output.



The two operational amplifier inputs now have impedances to AC ground that are nearly equivalent. Coupling between the inputs and an external AC ground sink now develop similar voltages on the two inputs, allowing the high common mode rejection of the operational amplifier to remove this noise from the output. The OPA2314 operational amplifier used in the WAQM OPC has a common mode rejection ratio of 75dB at a minimum, which is very effective at removing the coupling noise seen in the previous section.

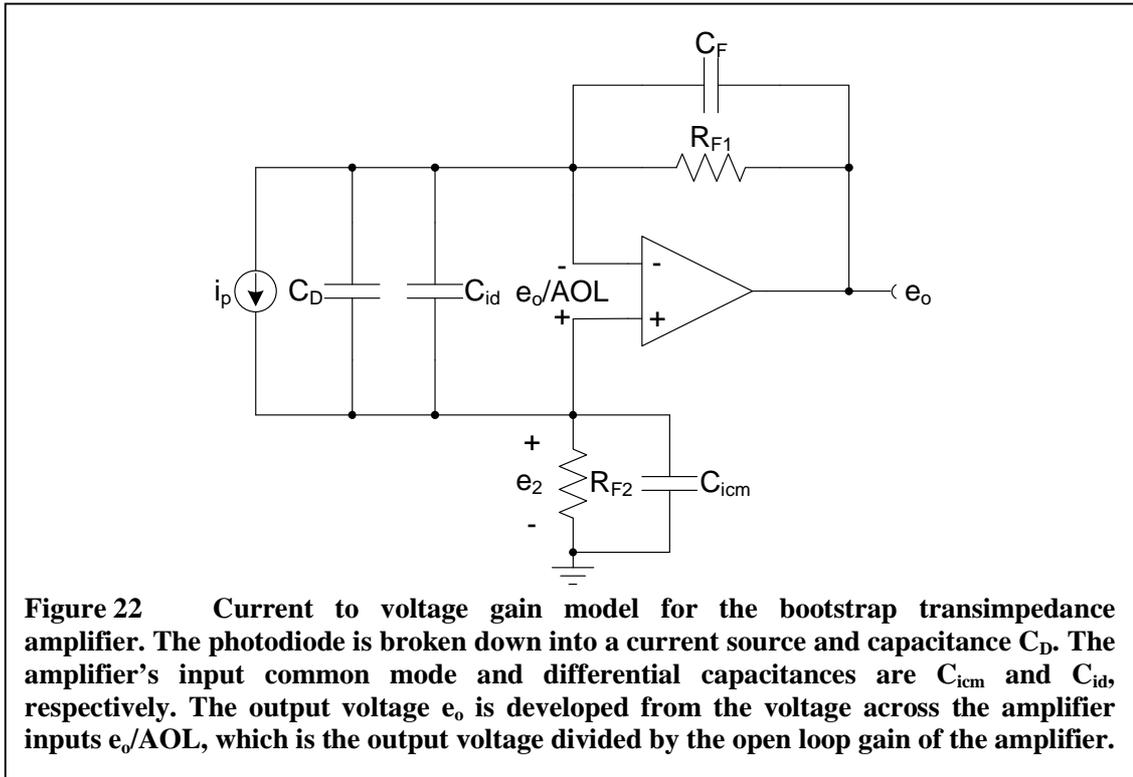
This circuit is similar to one described in Graeme [58], but with the modification to work with a single-rail amplifier. This provides the reduction in susceptibility to power line noise, while still allowing the use of a low cost single-rail design. As will be shown, the modified design has similar gain performance while suffering from an increase in

total RMS noise. This tradeoff is certainly acceptable given the improved resistance to the couple noise seen in the baseline AFE circuit.

As with the baseline circuit, the transimpedance amplifier current to voltage (I-to-V) gain can be found through node analysis of the circuit. The model for gain analysis is shown in Figure 22, and is different from that of the baseline circuit in the addition of R_{F2} and the bootstrap connection of the photodiode. R_{F2} is the parallel combination of the two bias setting resistors R_{F2a} and R_{F2b} in the circuit diagram of Figure 21, which are both taken to be connected to AC ground for the gain analysis. The amplifier's common mode input capacitance C_{icm} now shunts resistor R_{F2} since the non-inverting input is no longer connected to an AC ground. The output voltage e_o is divided by the open loop gain AOL to find the voltage across the amplifier inputs e_o/AOL .

The node analysis first sums the currents at the node attached to the non-inverting input of the amplifier as

$$\frac{e_2}{Z_2} + \frac{e_o}{Z_i AOL} - i_p = 0 \quad 4.31$$



where Z_2 is the impedance of the parallel combination of R_{F2} and C_{icm} , and Z_i is the impedance of C_{id} and C_D . Voltage e_2 is developed across impedance Z_2 , and the first term is the current through Z_2 . The second term is the current through the impedance of capacitors C_{id} and C_D , which appear across the amplifier inputs. The third term i_p is the photodiode current. These currents sum to zero, and solving for the voltage e_2 gives

$$e_2 = Z_2 \left(i_p - \frac{e_o}{AOL} \right) \quad 4.32$$

The currents into the node attached to the inverting node of the amplifier are

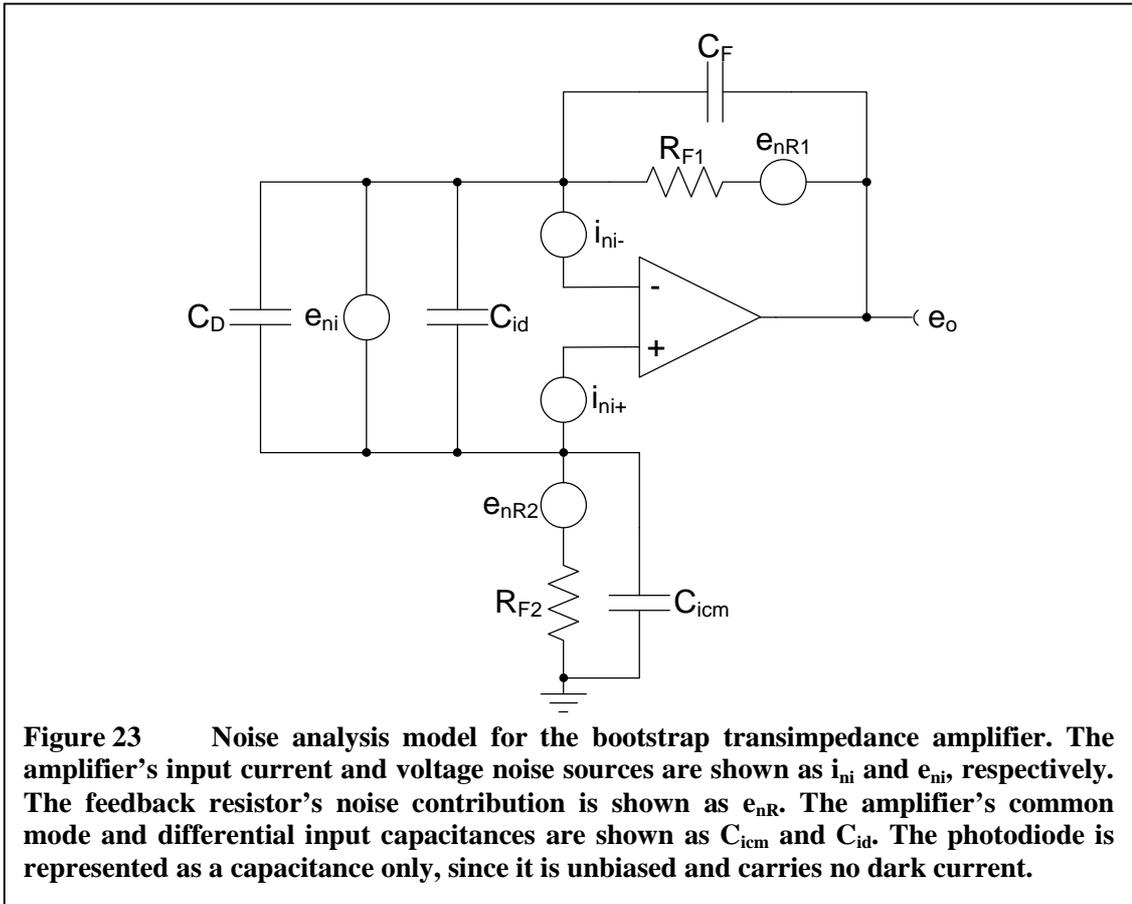
$$\frac{e_o - e_2 + \frac{e_o}{AOL}}{Z_1} + \frac{e_o}{Z_i} - i_p = 0 \quad 4.33$$

where Z_1 is the impedance of the parallel combination of R_{F1} and C_F . The first term is the current through Z_1 , the second term is the current through the capacitance across the amplifier inputs, and the final term is the photodiode current. Combining Equations 4.28 and 4.29 and solving for the current to voltage gain yields

$$\frac{e_o}{i_p} = \frac{(Z_1 + Z_2)Z_i AOL}{Z_1 + Z_2 + Z_i(1 + AOL)} \quad 4.34$$

Substituting in the components of the various impedances leads to a relatively complex equation and is not shown. As with the baseline circuit, second order components cause the I-to-V gain to roll off at 40dB when limited by the amplifier's open loop gain.

As with the baseline circuit, the noise analysis of the bootstrap transimpedance amplifier uses the methods discussed by Graeme [58]. The model used for noise analysis is shown in Figure 23. With the removal of the AC ground at the non-inverting input, the amplifier's input voltage noise e_{ni} is applied directly across the inputs, and amplifier input current noise sources i_{ni+} and i_{ni-} apply to both inputs. Resistor thermal noise sources e_{nR1} and e_{nR2} apply in series with each of the current to voltage conversion resistors. The amplifier common mode and differential input capacitances are shown as C_{icm} and C_{id} , respectively. As with the gain model, C_{icm} shunts resistor R_{F2} to AC ground. The photodiode is modeled as just capacitor C_D . Photodiode noise sources are no longer included since the reverse bias has been removed in the bootstrap configuration, eliminating the dark current that was present in the baseline model.



The feedback factor β can be found by performing a loop analysis on the noise model. The output voltage e_{noe} will be equal to the sum of voltages on the input side of the amplifier as

$$\frac{e_{ni}}{Z_i} Z_1 + e_{ni} + \frac{e_{ni}}{Z_i} Z_2 = e_{noe} \quad 4.35$$

where the first term is the voltage across impedance Z_1 , the second term is the voltage across the amplifier inputs, and the third term is the voltage across impedance Z_2 . As with the gain model, Z_1 is the parallel combination of R_{F1} and C_F , Z_2 is the parallel combination of R_{F2} and C_{icm} , and Z_i is the impedance of C_{id} and C_D . Solving for the feedback factor leads to

$$\beta = \frac{e_{ni}}{e_{noe}} = \frac{Z_i}{Z_1 + Z_2 + Z_i} \quad 4.36$$

The reciprocal $1/\beta$ will be the gain demand. This will have a zero at a frequency of approximately

$$f_z = \frac{1}{2\pi(R_{F1} + R_{F2})C_i} = \frac{1}{4\pi R_{F1}C_i} \quad 4.37$$

when $R_{F1} = R_{F2}$ as is the case with the bootstrap configuration.

The impedances Z_1 and Z_2 will roll off at different frequencies, since C_F is not equal to C_{icm} . Note that the two capacitances could be equalized, though matching C_{icm} may not be an easy task over operational corners. This would also require increasing C_F , which would lower the circuit bandwidth unnecessarily as no negative impact has been observed relative to the noise coupling issue discussed in the previous section. The larger capacitance C_{icm} begins to roll off impedance Z_2 at a frequency

$$f_{p2} = \frac{1}{2\pi R_{F2}C_{icm}} \quad 4.38$$

and the smaller capacitance C_F begins to roll off impedance Z_1 at a frequency

$$f_{p1} = \frac{1}{2\pi R_{F1}C_F} \quad 4.39$$

The noise voltage gain A_{ne} will follow $1/\beta$ until it is limited by the open loop gain of the amplifier. Ignoring the open loop gain limit, the noise voltage follows

$$A_{ne} = \frac{1}{\beta} = \frac{Z_1 + Z_2 + Z_i}{Z_i} \quad 4.40$$

Including the open loop gain into the equation adds an additional pole in the gain equation as

$$A_{ne} = \frac{Z_1 + Z_2 + Z_i}{Z_i(1 + s/\omega_i)} \quad 4.41$$

where ω_i is the intersection frequency of the $1/\beta$ gain demand curve and the open loop gain. As with the baseline circuit, the additional pole causes the noise voltage gain to fall off at 20dB/decade with the open loop gain.

The gain curves are plotted together in Figure 24. The dark blue curve is the open loop gain of the OPA2314 amplifier, the green line is the $1/\beta$ gain demand curve, the red line is the noise gain, and the cyan line is the current to voltage gain. As with the plot of the baseline circuit in Figure 12, the key frequencies are marked with vertical red lines and are named at the bottom of the plot. For the bootstrap design, the pole frequency is

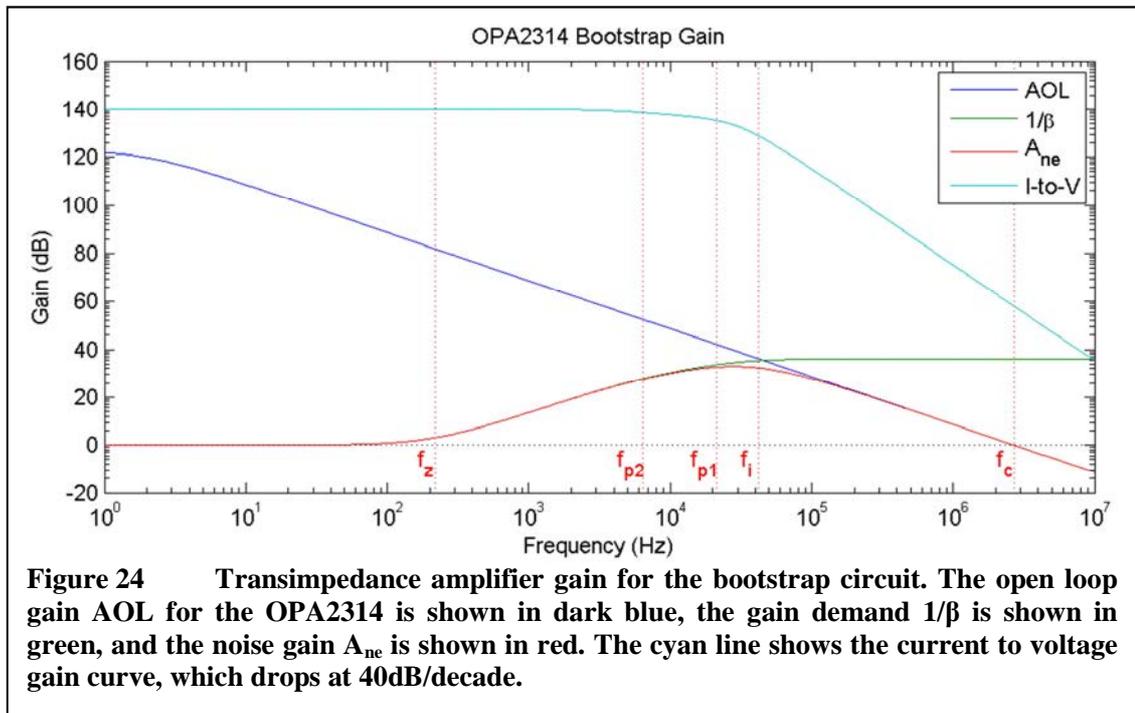


Figure 24 Transimpedance amplifier gain for the bootstrap circuit. The open loop gain AOL for the OPA2314 is shown in dark blue, the gain demand $1/\beta$ is shown in green, and the noise gain A_{ne} is shown in red. The cyan line shows the current to voltage gain curve, which drops at 40dB/decade.

split into two, with f_{p1} being the point where Z_1 roll off begins and f_{p2} the point where Z_2 roll off begins. This impacts the noise and current to voltage gains in that the lower frequency pole begins to roll off the gains at a relatively low frequency, but the roll off doesn't reach 20dB per decade until the higher frequency pole is reached. This can be seen most clearly with the current to voltage gain, where the gain starts to drop prior to f_{p2} , and then starts to fall off at a greater rate after f_{p1} . Comparing the current to voltage gain to the baseline circuit, f_{p2} occurs at a lower frequency than the baseline f_p , but f_{p1} occurs at a higher frequency. The combination produces a roughly similar result in overall gain.

Compared to the baseline circuit, the frequency of the initial zero f_z has fallen from 424 Hz to 218 Hz. This change is mostly due to the loss of the reverse bias on the photodiode. The unbiased photodiode has a higher capacitance value C_D , more than double that of the 1.65V biased photodiode in the baseline circuit. As a result, the noise gain rises at this lower frequency and reaches a higher point of 44dB at 27.7 kHz versus only 22.8dB at 33.6 kHz for the baseline circuit.

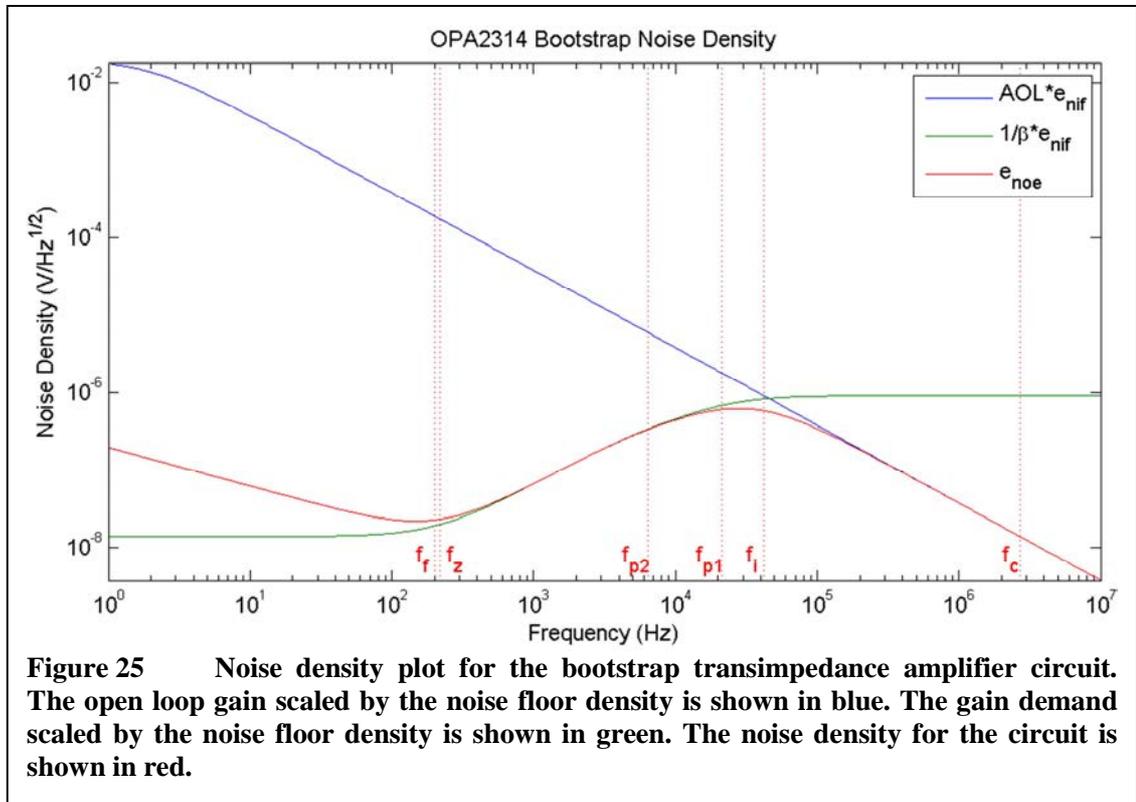
The noise density may now be found by multiplying the noise gain by the amplifier's input voltage noise density as shown in Equation 4.14. Expanding A_{ne} leads to the final noise density equation

$$e_{noe} = A_{ne}e_{ni} = \frac{(Z_1 + Z_2 + Z_i)\sqrt{(\omega_f + s)/s}}{Z_i(1 + s/\omega_i)}e_{nif} \quad 4.42$$

where ω_f is the frequency of the 1/f noise floor corner of the amplifier and e_{nif} is the input voltage noise density at the floor.

Figure 25 shows the voltage noise density e_{noe} for the bootstrap transimpedance amplifier circuit is plotted in red. The impact of the $1/f$ noise can be seen at lower frequencies, where the curve falls at 20dB/decade until reaching the floor corner at f_f . With the bootstrap circuit, this frequency is very close to the lower corner frequency of the zero f_z due to the photodiode capacitance. The earlier rise in noise density allows a higher peak of $617\text{nV}/\text{Hz}^{1/2}$ at 27.7 kHz compared to just $319\text{nV}/\text{Hz}^{1/2}$ at 33.6 kHz for the baseline circuit.

Finding the RMS noise for the bootstrap transimpedance amplifier circuit will allow the quantification of the difference in noise performance. Following the same method used to calculate the baseline RMS noise due to the input noise voltage, the response is divided into five sections. Though the boundary frequencies used in the



equations differ, the basic Equations 4.16-4.21 may still be used to calculate the RMS noise E_{noe} .

The bootstrap circuit changes do impact other RMS noise calculations used to find the total RMS noise. The thermal noise E_{noR} due to the two current to voltage conversion resistors R_{F1} and R_{F2} with $R_{F1} = R_{F2}$ may be calculated by

$$E_{noR} = \sqrt{2KT\pi R_{F1}(f_{p1} + f_{p2})} \quad 4.43$$

which rolls off according to the poles at f_{p1} and f_{p2} . Likewise, the amplifier's input current noise E_{noi} can be calculated with

$$E_{noi} = R_{F1} \sqrt{q\pi I_B (f_{p1} + f_{p2})} \quad 4.44$$

Finally, the noise sources may be combined into a single term for the total RMS output noise for the transimpedance amplifier circuit E_{no} as

$$E_{no} = \sqrt{E_{noe}^2 + E_{noR}^2 + E_{noi}^2} \quad 4.45$$

For the bootstrap circuit, the RMS noise calculations for the various noise components are shown in Table 4. Each of the components is listed, along with the total RMS noise E_{no} . The magnitude of the total RMS noise is relatively small, at $239\mu\text{V}$, but has nearly doubled from that of the baseline circuit. While this is an undesirable effect of going to the bootstrap circuit, the benefit of power line noise rejection may outweigh this

Table 4 RMS noise values for the bootstrap transimpedance amplifier in volts.

E_{noe}	E_{noR}	E_{noi}	E_{no}
2.39×10^{-4}	5.71×10^{-5}	1.86×10^{-6}	2.46×10^{-4}

negative in certain circumstances.

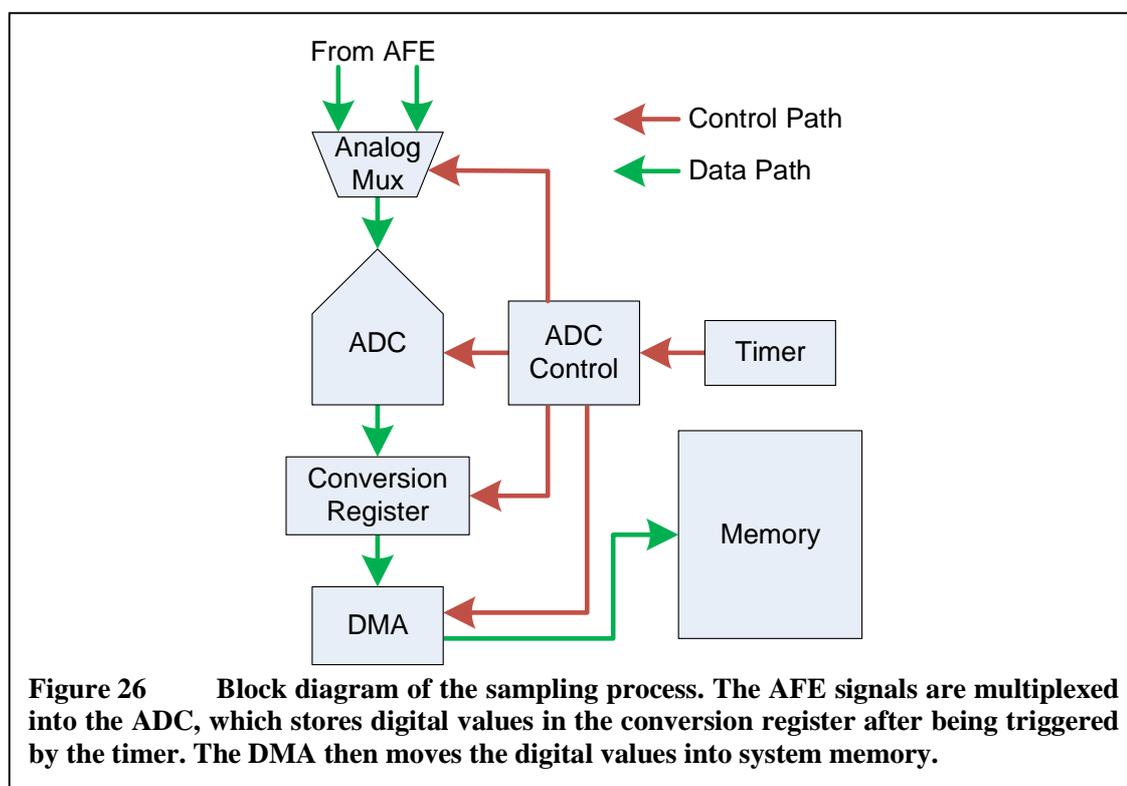
The second stage amplifiers are not modified from the baseline for the bootstrap circuit. The second stage gain curves will be the same as those shown in Figure 14. Applying the small channel gain to the transimpedance amplifier output as was done with the baseline circuit yields a total RMS noise for the AFE of 8.38mV. Converting this to a peak-to-peak value, the total RMS noise is assumed to be Gaussian and is multiplied by 6.6 to cover 99.9% of the noise. This yields a peak-to-peak noise of 55.3mV. Since the main concern for detection of small particles is in being able to distinguish deviation above an average value, this peak-to-peak noise is halved to 27.6mV. While higher than the total RMS noise of the baseline circuit, this still represents just over 3 least significant bits of the system's 10 bit ADC.

Analog-to-Digital Conversion

The analog-to-digital conversion samples the output of the AFE and produces a stream of digital values to be processed by the WAQM system's processor. Within modern microcontrollers, a number of hardware structures exist to assist with the sampling process, reducing load from the processor. This is true of the WAQM system's microcontroller, which allows a hardware sampling system to be built from multiple blocks in the device.

Figure 26 shows a flow diagram of the WAQM OPC's analog-to-digital conversion and sampling process. Though only a single analog-to-digital converter exists in the microcontroller, multiple analog inputs can be fed into the system and automatically sampled. The front end of the ADC includes an analog mux that is driven by the ADC control block. This mux will cycle through each of the active analog inputs, sampling each one when the ADC is triggered. With this mechanism, a single trigger will cause sampling of all AFE channels. In the case of a single sensor arrangement, this will include both the small and large channels.

The ADC can be set up to trigger based on one of the timer/counter blocks inside the microcontroller. This timed triggering allows for very precise sample timing in that no software intervention is required to start each sample. This removes most of the



potential for sample timing variation in the system. In the case of the WAQM OPC, the timer is set to trigger the ADC every 10 microseconds, resulting in a sampling frequency of 100k samples per second for each AFE channel. The timer block will automatically reload itself once triggering, keeping the sampling process running until disabled by the embedded software. The timer is the control point for enabling and disabling the entire AFE sampling process.

As the ADC samples each channel, the digital value is transferred into a conversion register. When this happens, the ADC control block indicates that a sample has completed. This sample complete signal may be used to trigger a direct memory access (DMA) block in the microcontroller. When triggered, the DMA automatically moves the sampled data into the main system memory without software intervention. The DMA is set up to move blocks of sampled data. In the current implementation of the WAQM OPC, the block size is 512 samples per channel. The DMA adds sample data to each block, transferring it as the ADC converts and stores digital values in the conversion register. The data for the channels is interleaved, since the DMA increments its target address counter for each transfer and cannot separate the channels into separate blocks. This is not really an inconvenience since often it is desired to have the large and small data together, and skipping over the large data to access consecutive small channel values only requires pointer addition.

As each block is filled, the DMA interrupts the processor to signal completion. Given the size of the blocks, this means the processor is only interrupted every 5.12 milliseconds to update the DMA. This is much better than requiring an interrupt each time the ADC completes, which for a two-channel system would require interaction with

the ADC twice every 10 microseconds. As an added benefit, the microcontroller used in the WAQM system has a reload function that will start a subsequent DMA automatically. This removes much of the urgency to service the DMA, since the system has another 5.12 milliseconds to reload the DMA before the system starves.

The sampling firmware in the WAQM system maintains a circular list of sample blocks that can be targeted by the DMA. As each block is filled, the system firmware responsible for processing the data parses the block for particle signatures. When complete, the firmware returns the empty block to the circular list and moves to the next available block. Failsafe functionality has been built into the buffer processing algorithm in the case that the circular block list overflows. This may happen if some part of the system software inhibits processing of the blocks for an extended period of time, such as has been seen with the SD card write process from time to time. If this happens, the DMA interrupt may still load empty blocks into the DMA process until all blocks are filled. If this happens, the system will recover and return to sampling once the blocking situation has cleared. The particle detection algorithms that are performed as part of the block processing will be discussed in later chapters.

Laser Source

The laser in the optical particle counter is one of the most critical pieces of the system. This component directly impacts the sensitivity of the OPC, and is the one component that requires careful alignment and attachment to ensure correct operation. Depending on the quality of the component, the laser may encounter failures during the life of the OPC that render the device useless. Great care should be taken when designing in a laser source so that the product quality and endurance can be maximized. This

section discusses the requirements for the WAQM OPC laser, and the lessons learned during the development of the OPC.

Laser Module

Early in the design of the OPC, it was decided for simplicity to use a prepackaged laser module that could be purchased at a relatively low cost. This decision was made to ease the acquisition of the various components required to create a focused laser beam. At the time, it was seen as a distraction to the main project goals to spend the effort putting together a laser diode, driver circuit, and lens into a reliable package that could be used for the OPC. This was reinforced by the availability of relatively small and inexpensive laser modules that had become very common due to their use in laser pointers.

One parameter required to make the OPC function is not common in low-cost laser modules, and greatly limits the selection of acceptable devices. This parameter is the ability to adjust the focus of the beam. Most available modules have a fixed lens that focuses the laser into a collimated beam, meant to have very little divergence over a long distance. The OPC requires that the laser be focused to a narrow beam waist just over the horizontal sensor, dictating that the lens be adjustable. During the course of the project, several different modules were sampled for suitability, but only a single source was found to be acceptable with relatively small size and adjustable focus lens.

Some of the laser modules used in the development of the WAQM OPC are shown in the top image in Figure 27. The laser chosen for the WAQM OPC was the Apinex 03008B focusable red laser module [61]. The laser outputs a 3mW beam of 650 nm light, and has an adjustable plastic lens. The laser is housed in a brass tubular enclosure that is 23 mm long and 8 mm in diameter. The module nominally pulls 30mA at 3.0V across power and ground wires that extend from the rear of the module. A dissected view of the module is shown in the bottom image of Figure 27, showing the adjustable lens holder, the acrylic lens, the module housing, and the laser diode assembly that sits inside the housing.

While the module itself met the needs of a limited production prototype system, several features make the module less attractive for larger scale use. One of these was the beam alignment accuracy. The manufacturer places the laser diode assembly into the module housing and then glues it into place. This process appears to leave a great deal of opportunity for less than perfect alignment. Modules were routinely found with up to 5 degrees of angular deviation from the housing centerline. This precluded having a module that would be fixed in place relative to the sensor, since the module needed



Figure 27 Laser modules evaluated during WAQM OPC development. At top are three of the modules considered for prototype development. At bottom is a dissected module of the type deployed in the WAQM OPC.

physical adjustment to correct for the angular deviation of the laser.

The output power specification for the Apinex module was also not well constrained, stated as being less than 3 mW. Measured module-to-module variation ranged from 2 mW up to nearly 5 mW. To achieve some level of similarity in OPC sensitivity between units, the modules had to be sorted after purchase, to select modules inside an acceptable performance range. Typically this meant eliminating the lowest power modules, since higher power modules could be normalized through calibration. Since many modules were measured at power output near 4 mW, modules with power less than 3 mW were rejected.

Another concern during the development of the WAQM OPC was the assurance of supply for the module. In particular, there was only one source, with a non-standard form factor. Furthermore, the vendor is not the manufacturer, which appears to be a small Chinese corporation with which there was no direct contact. If the vendor or its supplier for some reason stopped selling the module or module components, the program could be jeopardized. Fortunately the vendor continues to sell the module, and availability has not impacted development.

Power Output

For the Apinex laser module used in the WAQM OPC, the power output was not adjustable, being controlled by the laser diode assembly. As is the case with many commercially available laser diodes, the output is controlled with feedback from a photodiode placed behind the rear facet of the diode. While the rear facet is coated with a reflective material, enough light still passes through to allow this photodiode to be used for feedback in an intensity control circuit. The photodiode is also exposed to light

entering the front of the module, which can greatly impact the output power in the case of reflected laser light. Issues with the alignment of the diode in the module housing often caused reflections in the lens assembly that could impact output power through the photodiode feedback.

The power output of each module was measured prior to integration into a WAQM OPC unit. Power measurements were made using a Coherent Incorporated photodiode-based laser power meter, the PowerMax-USB [62]. The measurement process was performed by hand, with several constraints placed on the operator to ensure a repeatable measurement. Producing repeatable results when taking power measurements can be challenging due to multiple factors beyond just the condition of the power meter itself. When measuring the laser output power, if the laser is aligned orthogonally to the surface of the power meter, a significant amount of light can be reflected directly back into the feedback photodiode, causing a drop in the output power. More consistent results are achieved by slightly angling the incident laser path relative to the surface of the power meter.

Angling the laser's incident path into the power meter will also have an impact on the measured power. The light emitted from the laser diode is highly polarized, with polarization parallel to the long axis of the emitting facet of the diode. Due to the strong polarization of the laser diode output, the orientation of the laser diode output facet relative to the power meter surface will dictate how much light is reflected from the sensor surface. As mentioned previously, light exiting the laser diode will be polarized parallel to the long axis of the output facet. Just as with the photodiode in the particle counter, the difference in indices of refraction between air and the sensor will cause

variation in transmitted irradiance depending on the percentage of incident light that is polarized perpendicular to the plane of incidence. In practical terms, this means that rotating the laser module will impact the power reading when the beam is oblique to the sensor surface.

Care must also be taken when measuring the power of a beam that is focused to a narrow waist rather than collimated. If the laser is measured with the focal point on the sensor itself, depending on the type of sensor it may saturate or even be damaged by the high intensity at the beam waist. The beam should be measured either in front of or behind the waist, while taking care that the beam spot fits inside the active sensor area.

Additionally, depending on the quality of the module, there may be reflections of significant intensity coming from the module's lens assembly. These reflections can produce artifacts outside the normal beam path that would normally not illuminate particulate matter inside the particle counter. Care should be taken that these artifacts not artificially increase the measured power of the beam, since they will not contribute to the intensity of scattered light in the particle counter itself.

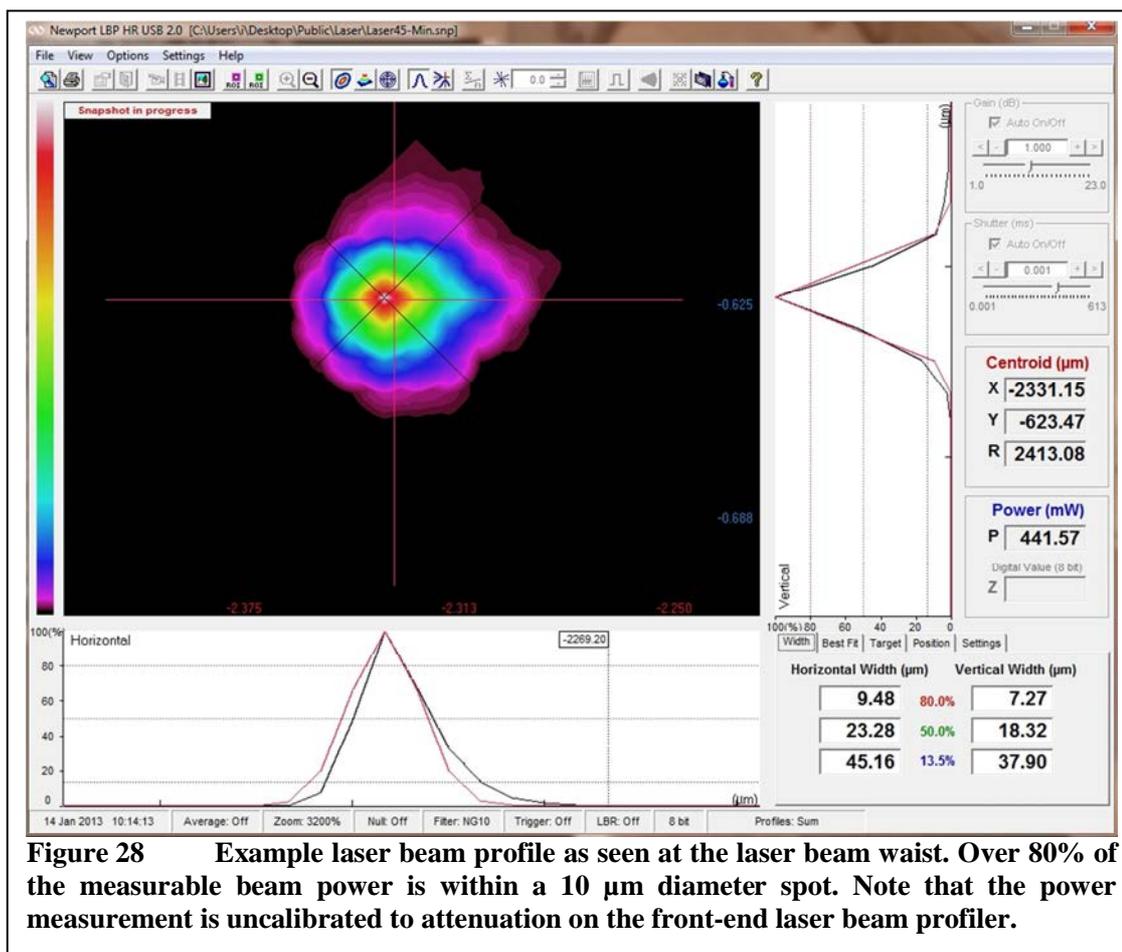
Laser Beam Profile

The shape of the laser beam significantly impacts the performance of the OPC. To achieve an incident irradiance with enough power per unit area to cause detectable scattering from small particles, the laser beam must be focused to a narrow waist. For example, assuming a flat profile, a 3 mW laser with a collimated output and beam diameter of 3 mm will have a relatively low irradiance of 420 W/m^2 . Focusing the same beam to a waist with diameter of $10 \mu\text{m}$ yields an irradiance of more than 38 MW/m^2 , a

significant improvement. This allows the detection of very small submicron particles. A collimated beam can only cause detectable scattering from very large particles.

For the WAQM OPC, a laser beam profiler was procured to assist in examining the quality of the beam produced by the Apinex laser modules used in the system. The Newport LBP HR laser beam profiler [63] is a charge-coupled device (CCD) based imaging system that allows the insertion of attenuation filters on the front end, which was critical to measuring a the beam waist of the focused laser. With the beam focused to a high irradiance at the waist, the sensor will be saturated without attenuation. It is also possible for damage to occur to the sensor without attenuation, and care must be taken when setting up such a measurement. To measure the beam waist of the focused Apinex module, three attenuators were stacked to reduce the irradiance to acceptable levels. The Newport sensor would saturate at $20\text{W}/\text{m}^2$, and potentially take damage at $500\text{W}/\text{m}^2$. The three attenuators reduced the irradiance at the sensor by a factor of 3×10^{-7} , reducing the irradiance of the beam at the waist to levels less than the saturation point of the sensor.

An example laser beam profile taken at the beam's waist is shown in Figure 28. The beam is measured with more than 80% of its output power within a $10\ \mu\text{m}$ diameter circle. Note that the power measurement is uncalibrated to the attenuating filters that have been added to the input of the profiler. The profiler is also very near to its resolution limit, having pixel sizes of $4.65\ \mu\text{m}$. The profiler did reveal that variation in the beam shape occurred between modules, but not in a way that would significantly impact the beam intensity. One key revelation had to do with the shape of the laser beam away from the beam waist. The focused beam exhibited intensity rings with the potential to cause



interesting pulse profiles as a particle passes through them. These “misfit” particles will be discussed in the chapter on particle detection and sizing.

Reliability

The reliability of the laser used in an OPC is an important factor to consider when choosing a laser source. Failure of the laser module will render the OPC useless. Even worse, degradation of the laser’s power output over time may lead to inaccurate particle counting results. Ideally, the laser source will be purchased with a guarantee of a minimum amount of power change over the expected life of the OPC. For the WAQM OPC, this was certainly not the case.

The reliability of the laser modules used in the WAQM OPC was an ongoing issue for the program. As previously described, the choice of laser modules was driven mostly by the availability of an adjustable lens, with other considerations being secondary. The module found to meet the basic minimum requirements for the WAQM OPC did not have a failure rate specification. Furthermore, as the module's primary use was in laser pointers, continuous operation was likely not the intended use model. Given this, it is not surprising that issues were encountered.

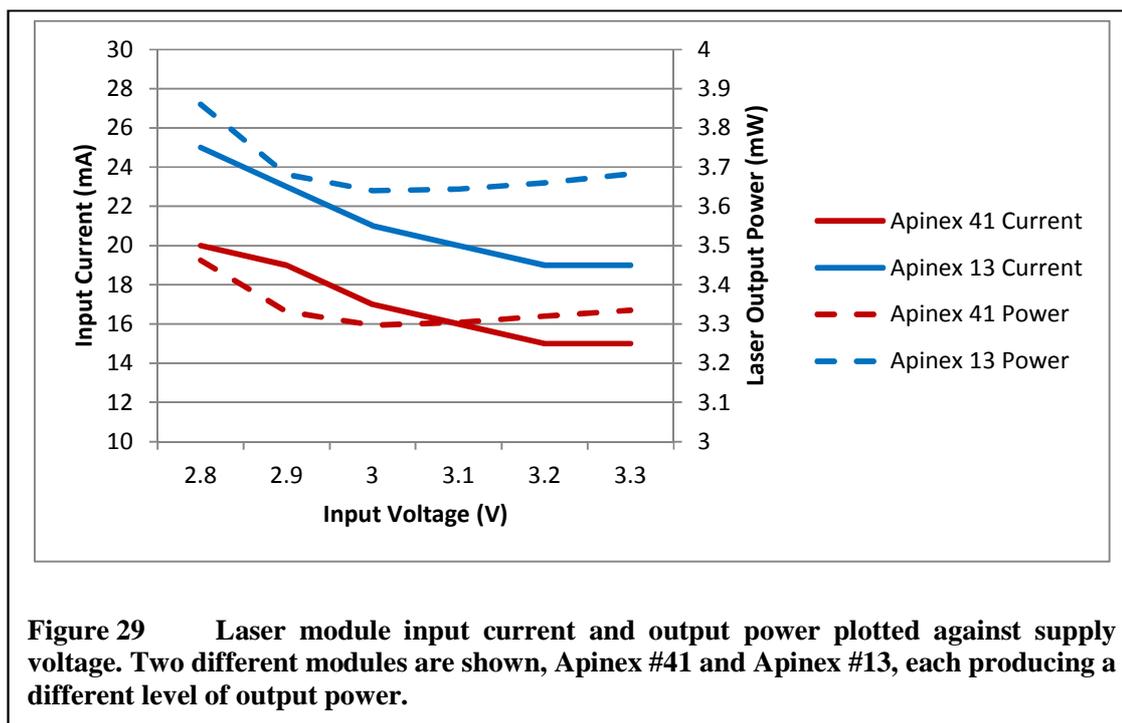
Over the course of the WAQM OPC development, 81 modules were purchased, tested, and in most cases incorporated into an OPC design. Of these 81 modules, nearly 10% fell out at pre-screening due to low output power. A further 26% of the modules failed in operation or had power output that degraded significantly. The typical failure mode seen involved a significant drop in output power with a corresponding increase in current pulled from the power rail. While the exact mechanism was not investigated, the cause did not appear to be due to overheating of the module housing. The vendor did suggest that supplying the module with 3.3V may have caused some of the failures, but further testing with a 2.9V supply did not bear this out. In fact, dropping the laser module supply voltage caused other issues, as will be discussed in the next section.

Laser Output Stability

A key issue that was encountered late in the testing phase of the WAQM system was related to laser stability. This is of particular importance to mention due to the impact that it had on some of the test results that will be covered in the system demonstration chapter.

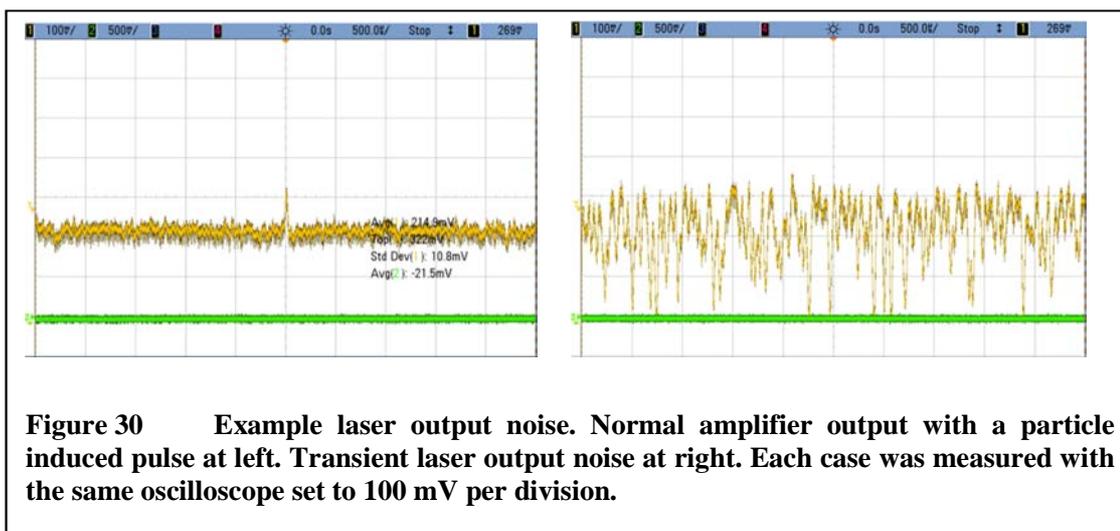
In an attempt to address some of the reliability issues experienced with the laser module, the nominal supply voltage for the module was lowered from 3.3V down to 2.9V. This was accomplished by inserting a small linear regulator into the design, taking power from the WAQM unit's 5V rail through the regulator rather than directly from the 3.3V rail. This was done at the suggestion of the module vendor, who stated that the module failures being seen were likely due to the high 3.3V supply voltage.

For the most part, this change appeared to have little impact on either the particle counter sensitivity or the failures seen with the modules. Figure 29 shows the current consumption and power output of two different Apinex laser modules across a voltage range of 2.8V to 3.3V. The output power does dip slightly near the nominal supply voltage of 3.0V, and rises with decreasing voltage. Note that even though the modules were operated with a lowered supply voltage, there was no reduction in the rate of early-life failures.



Though the lowered supply voltage did not at first appear to have any detrimental impact on performance, it was discovered that some units were seeing occasional large peaks in particle count. On further investigation, it was found that the laser output appeared to be generating noise that caused false particle counts. Figure 30 shows an example of a unit when the laser is relatively noise-free, and when it is unstable. The unstable period generates a significant number of false small particle counts.

It is unclear at this point why the laser oscillates periodically when operated at voltages below 3.0V. The module specification lists a recommended operating voltage of 3.0V, but the vendor's support contact stated that the module would see longer life if operated between 2.8V and 3.0V. The instability may vary with temperature, since often the oscillations will only occur just after power up, and will die out over the course of several minutes. The exact cause was not determined, since the module is basically a black-box and the focus of future work will likely include the design of a laser source



using discrete components for the diode and lens. Operating at higher voltages does appear to eliminate the instability problem.

Mechanical Design

The mechanical design of an OPC is a critical component in getting good results in real-world operating environments. The airflow through the device must be designed well to allow particulate matter to enter the channel without significant loss. The ambient light from sources external to the OPC must be managed. The laser light in the sensing area must be controlled, to keep a dark environment for the photodiode to detect the low intensity pulses from particle scattering. All of this must be handled in a small form factor to enable a portable, low-cost device.

Much of the work in generating the mechanical design for the WAQM OPC air flow was performed by the mechanical engineer on the development team, and as such only an overview of the design and related issues are covered. The models for the WAQM enclosure were designed using SolidWorks modeling tools and constructed using a stereo lithographic printer. Simulations of the air flow were performed with

computational fluid dynamics tools. The part of the mechanical design specific to the OPC is the flow channel, which is a modular part of the enclosure that screws into the enclosure base.

The flow channel design for the WAQM OPC is shown in Figure 31. Air enters the flow channel at a wide inlet that rapidly narrows down to the main flow channel. The main flow channel is 6 mm wide by 10 mm in height, and contains the photodiode and

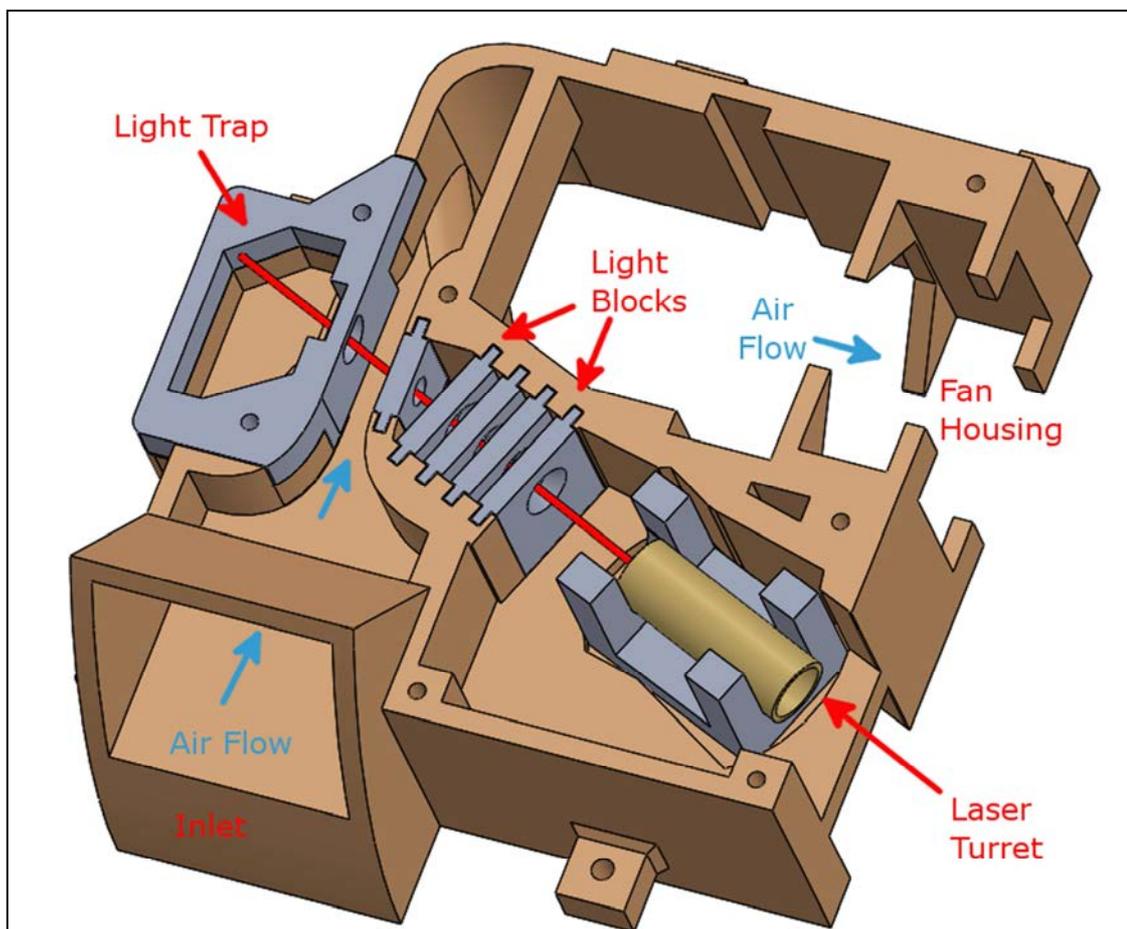


Figure 31 WAQM OPC Flow Channel Design. The flow channel module base is shown in brown. The laser turret holds the laser, which shines through the light blocks and into the light trap. The air flow is shown in blue, starting at the inlet and passing by the laser beam in the flow channel. The grey pieces are separate parts which can be replaced with different designs supporting other configurations. The laser turret can rotate, allowing fine positioning of the laser beam.

laser path. Past the main flow channel, the airflow opens up into a larger space at the end of which is a DC fan that creates the static pressure driving the airflow. The path of the airflow is highlighted with blue arrows in the figure. Note that the fan is not shown, but sits in the location labeled as the fan housing.

There are several pieces to the flow channel, including the main body shown in brown in the figure, and several detachable pieces shown in gray. These have all been fabricated using stereo lithography, and are made from photopolymer. The gray pieces make up the light path, starting with the laser turret, which holds the laser module for the OPC. The other pieces in the light path are the light blocks, which sit between the laser module and the flow channel, and the light trap, which is on the far side of the flow channel from the laser. When fully assembled, the top of the flow channel is covered with the OPC's main circuit board, which provides the fourth side of the flow channel. The circuit board is not shown in the figure.

Airflow

The airflow through the device delivers particulate matter to the sensing area, passing individual particles through the laser beam past the photodiode and out through the fan. The fan creates the pressure differential that draws air through the system. This fan is one significant difference between the WAQM OPC and a typical commercially available OPC. Most of the more expensive OPCs will use a vacuum pump to move air through the system at a fixed rate. While this provides a relatively stable and known flow rate through the system, the vacuum pumps are typically expensive and quite noisy. While not a primary concern, the noise of a vacuum pump is a significant concern given

that the WAQM OPC may be deployed in homes and businesses where a quiet environment is expected.

To approximate the stable flow rate of a vacuum pump, the fan speed is controlled to a fixed number of rotations per minute. While this is not as good as a vacuum pump at dealing with ambient air flow conditions, it does provide at least a stable fan speed from unit to unit, which should provide a fairly uniform flow rate between units given matching fan sizes. The fan speed is controlled with a tachometer feeding speed into a control algorithm in the WAQM OPC microcontroller. This control algorithm is a proportional-integral-derivative (PID) controller, which runs the fan slightly slower than its fastest speed to allow unit-to-unit normalization.

While the PID controlled fan approximates a stable flow, it is really only stable when the air being sampled is relatively stable. Significant air motion may impact the flow rate through the OPC, which in turn will impact the measured concentration. Care must be taken when using the WAQM OPC in such an environment to make sure that ambient air flow does not impact the measured concentration. This can typically be done by shielding the OPC in some way to prevent significant external air motion from entering either the inlet or outlet of the flow channel.

There are many other significant issues in creating the flow channel through which the air moves, but most of these are mechanical and aerodynamic issues that are not discussed in detail in this document. Care should be taken in designing the mechanics of the flow channel such that particulate matter is accelerated gently, avoiding sudden turns that might cause particles to impact the sides of the flow channel. Much work could yet be done to optimize the mechanical aspects of the WAQM OPC flow channel.

Incident Laser Light

The OPC is set up to be as sensitive as possible to relatively small numbers of photons scattering from sub-micron particles. To aid in this sensitivity, keeping stray light from illuminating the photodiode is critical. Incident laser light coming from the direction of the laser module must be managed to keep it off of the photodiode surface. This incident light may be directly from the laser module or reflected off of light blocks in between the laser module and photodiode.

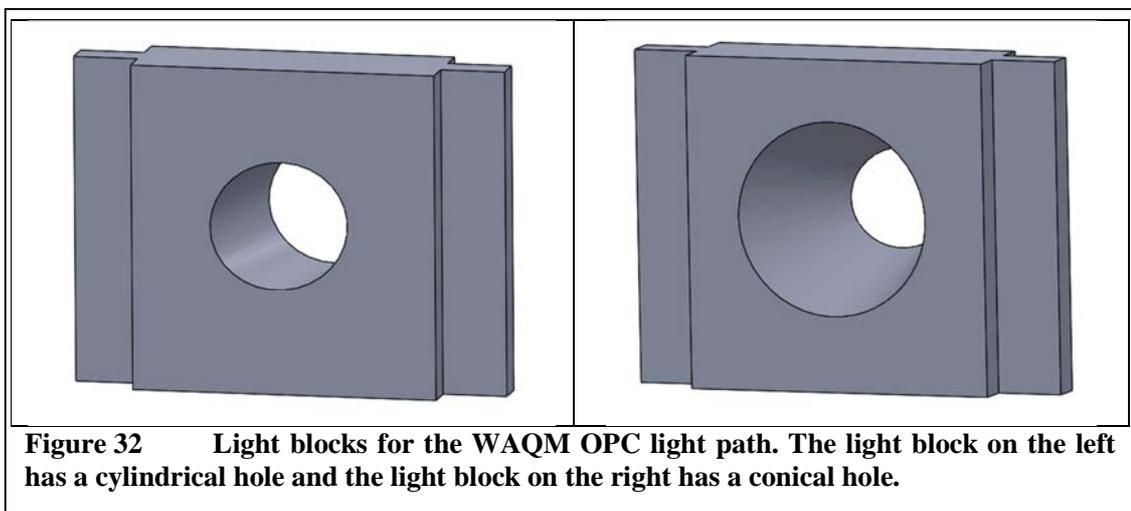
Unwanted incident light directly from the laser module is light that does not make up the main beam. A fair amount of laser light is emitted from the laser module outside of the main path of the beam. This is light that impacts the inside surface of the laser module housing or lens holder, scattering from the laser module lens at angles outside of the main beam. With nothing between the laser module and the photodiode, this light would strike the photodiode directly. The light blocks are inserted between the laser module and photodiode with the sole purpose of stopping this light before it reaches the flow channel.

Various iterations of the WAQM OPC design had between two and five light blocks between the laser module and the photodiode. With experimentation, it was found that more light blocks provided better shielding of the photodiode from unwanted incident laser light. It was also found that the shape of the holes through the light block had some impact on their effectiveness. If a cylindrical hole is punched through the light block, light may illuminate the inside surface of this hole and reflect through to the photodiode. Shaping the light blocks with a conical hole with the wide opening toward the laser module has the effect of reflecting this light back towards the laser module rather than onward to the photodiode. Light blocks with conical holes tended to work

better than those with cylindrical holes, provided the hole diameters were equal at their narrowest points.

This last requirement is critical when shaping the conical hole for a stereo lithographic process, since the narrow end of the conical hole may not form exactly as intended. If the conical hole comes to a sharp edge at its narrow end, the stereo lithographic process will not be able to build the edge, leaving the hole larger than intended. This necessitates leaving at least some width of the hole with a cylindrical shape at the narrow end of the cone. Figure 32 shows a pair of light blocks, one with a cylindrical hole on the left and one with a conical hole on the right.

The first designs of the OPC had fixed emplacements for the laser module, but issues with the angular deviation of many laser modules were causing them to be discarded because they could not be aligned with the light path. Misalignment with the light path typically causes at least some portion of the beam to impact the inside edges of the light blocks, sending reflections down the light path as unwanted incident light that might reach the photodiode. The turret design allows the laser to move side to side



slightly, with the turret swivel point directly under the point where light leaves the laser diode in a typical module. This allows an adjustment in the horizontal alignment of the housing. The module can also slide up and down slightly in the U-shaped mounts that hold it, allowing a vertical adjustment. Between the two adjustments, most laser modules can be accommodated regardless of deviation from the module housing centerline.

This adjustment allowance is particularly important due to the polarization of the laser beam. As will be discussed in subsequent chapters, the polarization will impact the scattering intensity for a given particle, and making sure that the laser module polarization aligns correctly with the flow channel is critical to match expected results. This means that laser angular deviation from the housing centerline cannot be overcome by rotating the laser module in the U-shaped mounts inside the flow channel. With rotation constrained, the ability to adjust the laser module both horizontally and vertically becomes important in correcting angular deviation of the diode.

Reflected Laser Light

Once the laser crosses the flow channel, it is no longer required for its useful purpose of illuminating particles. The laser light must be absorbed without reflecting back into the flow channel and into the photodiode. The purpose of the light trap is to fill this role, deflecting and capturing the main incident beam to keep it from reflecting back into the photodiode. The laser beam will impact the deflecting wall of the light trap directly, which both reflects and scatters the incident light. The deflecting wall is angled with the intent of reflecting light deeper into the light trap rather than back out into the flow channel. Light that scatters from the deflecting wall is more problematic. To keep this scattered light from impacting the photodiode, the impact point of the light is hidden

from the photodiode by the top of the light trap inlet. If the impact point is not hidden, this very intense spot of light can scatter directly back into the photodiode.

Figure 33 shows an image of the laser passing through the light path. The laser originates at far right in the laser module, passing through the light blocks and across the flow channel. The bright point of laser light is the beam's impact point on the deflecting wall of the light trap. If this bright point was not hidden from the photodiode, a significant amount of scattered light would reach its surface. The illumination of the wall near the laser module shows some of the light coming off of the laser module outside of the main beam, necessitating the light blocks between the laser and photodiode.

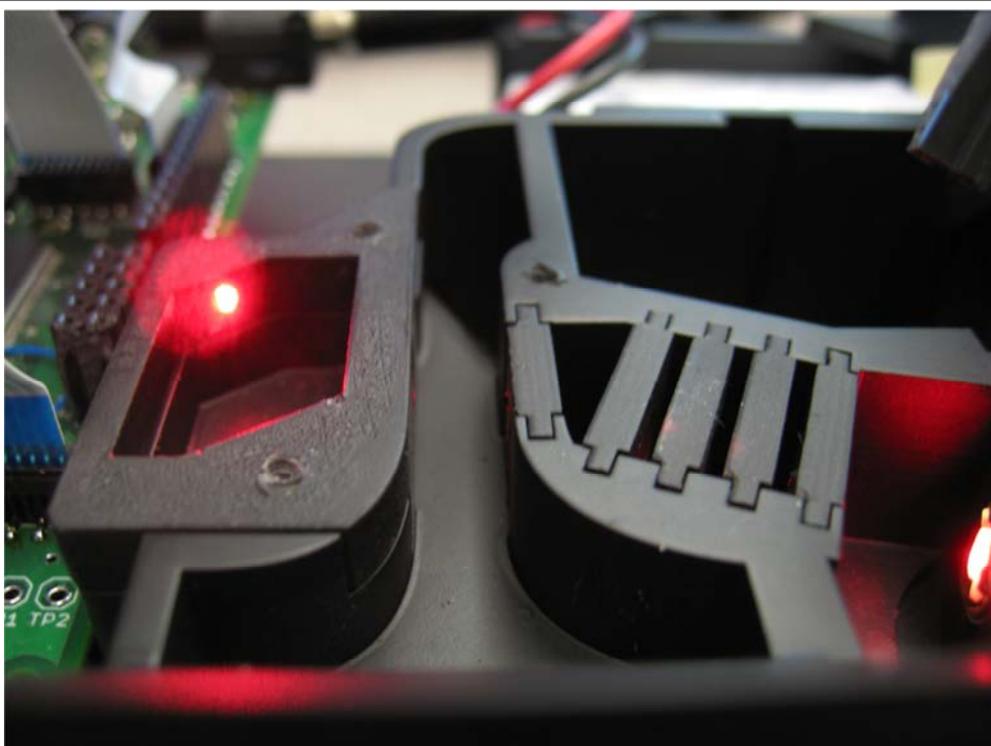


Figure 33 Laser illumination in the light trap of the WAQM OPC flow channel. Light emitted from the laser module at far right passes through the light blocks and crosses the flow channel to impact the deflecting wall of the light trap. The light trap is angled to prevent most of the light from reflecting back into the flow channel.

Ambient Light

While stray laser light can impact the sensitivity of the OPC, ambient light can render the OPC inoperable. This is especially true when the WAQM OPC is used outdoors. Sunlight entering the inlet of the flow channel at the wrong angle can reflect off of the floor of the flow channel and into the photodiode. When this happens, it can saturate the photodiode, dropping the measured particle count to zero immediately. Light can also enter the flow channel through the outlet. When this happens, the fan may act as a chopper, creating pulses of light that reach the photodiode and generate a large number of false particles.

Countermeasures to help block ambient light have been implemented at the inlet and outlet of the flow channel. These consist of a series of louvers covering the openings, which are designed to block much of the ambient light while allowing air to pass with minimal blockage. While the countermeasures are typically effective with indoor lighting, even indirect sunlight may still cause the issues described above. There is just not enough distance between the inlet and the photodiode to deal with sunlight. Use in an outdoor environment should consider shielding the inlet and outlet from sunlight.

Particle Counter Hardware Cost

The starting premise for the particle counter design was that it had to be low cost to enable inclusion in a WSN. To reach the lowest price points, the components needed to be simple off-the-shelf parts, preferably those that were already being mass produced. Most of the electronics for the detector system, dedicated power supply, and fan control were readily available and sourced by multiple vendors, including the large area photodiode and amplifier components. By far the greatest expenses come from the DC

fan and the laser module. Single-quantity prices for even small, sleeve bearing fans are on the order of \$10. Small, low power laser modules typically intended for laser pointer use can be purchased for very low prices, but they are typically fixed with a collimated focus. It is much more difficult to find small, low cost laser modules that have adjustable focus, but they can be found for prices on the order of \$15.

The approximate component cost for adding the particle counter to an existing sensor system is shown in Table 5, and is on the order of \$30-\$35 when purchasing components in small, prototype quantities. This does not include the cost of additional PCB area for the circuit or the enclosure plastic for the airflow, since these factors are

highly dependent on the rest of the sensor system. This cost also does not include the two analog-to-digital converters used for the two size channels, since these are often included in common microcontrollers that might be used for a sensor system.

Component	Approximate Cost
Laser Module	\$15.00
DC Fan	\$10.00
Detector Circuit	\$5.00
Power Supply	\$1.50
Other	\$1.50
Total	\$33.00

Were this system to be mass produced, it is highly likely that improvements both in lowering cost and increasing quality of components could be achieved. The laser module in particular might be custom sourced or built from discrete components, with tighter tolerance on the lens assembly and output power that would improve the quality of the system.

Particle Counter Hardware Conclusions

The design of the particle counter hardware was one of the more challenging aspects of the WAQM OPC development. This effort is a complex, multi-disciplinary

exercise in balancing electrical, optical, and aerodynamic performance against system size, measurement accuracy, and a host of other factors. The effort for the WAQM OPC was overall successful in producing a design that with calibration yielded relatively consistent results, allowing the creation of multiple units that can be used together to create a multi-node particle measuring network. As will be seen in the system demonstration chapter, 16 units were used simultaneously to map the particulate matter flow through a simulated aircraft cabin.

The AFE design presented a significant challenge in reducing the impact of power line noise coupled into the high impedance node of the transimpedance amplifier. Solving this issue with the single-ended bootstrap circuit was a major step forward in improving the measurement results from this low-cost wall transformer powered OPC. Further work could be done on this circuit to lower the total RMS noise, which is the most serious issue remaining in the AFE design.

The mechanical and optical portions of the design were evolved as the development of the WAQM OPC proceeded. Much was learned relative to how to handle unwanted laser and ambient light, as well as how to balance this with the construction of a good airflow. Moving from a monolithic plastic design to a modular design helped to speed up the prototyping and evaluation process for the system. This allowed several different designs to be tested at once while minimizing solid model development effort and prototype material costs.

From a holistic view point, the decision to use a low-cost laser pointer as the laser source was the correct one to make for a prototype. Using a module removed the complexity of combining a lens, diode, and driver circuit together into a workable

solution. This choice does continue to have implications for the design constraints on new versions of the OPC, as the form factor of the module can significantly influence the dimensions of the system in which it is deployed. There is also significant risk in that the assurance of supply for the module is bound to a single source, and loss of this source could derail further production of WAQM OPC units. Future versions of the OPC may move to a discrete design, taking control of many design and performance factors that are wrapped up in the current module.

CHAPTER FIVE: SIMULATION

The WAQM optical particle counter depends on the detection of light that is scattered by particles moving through a focused laser beam. To determine the optimal sensor arrangements and detection algorithms, a set of simulations was created to predict the response for particles of given size and composition. At the heart of the simulations are algorithms that depend on calculations of the solution for light scattering by homogeneous spheres, developed and published by Gustav Mie in 1908 [64]. These Mie scattering calculations allow us to approximate the intensity and direction of light scattered by particulate matter passing through a laser.

This chapter will cover the simulation flow along with several examples of its use in analyzing and designing the WAQM OPC. In general, the simulations will employ polystyrene latex material parameters to mimic the calibration particles typically used by OPC systems. Results are also explored for particulate matter that the WAQM optical particle counter might be expected to encounter when in use.

Simulation Flow

The general simulation flow is shown in Figure 34, and is the model used for many of the specific simulations used in this work. The basic system is assumed to consist of a laser beam with a waist centered in a simulated particle sampling volume. Particles will then move through the volume orthogonal to the beam, scattering laser light when they intercept the beam. Scattered light will fall on the surfaces of one or more

sensors that are either on the surface directly below and parallel to the axis of the beam, or standing vertically to the side of the beam at some angle off of the axis. The purpose of the simulation is to calculate the power of the light received at each sensor for any given particle intercepting the beam at an arbitrary point.

The simulation first divides each of the sensors into segments, splitting the area of the sensor into multiple equally sized rectangles. The number of rectangles will determine the granularity of the simulation and the overall accuracy of the resulting

irradiance calculations, which must be balanced against the computation time required for each simulation. A set of particles is then created, with one or more sizes and one or more indices of refraction across a set of intercept vectors within the sampling volume.

For each particle, a set of vectors is created from the particle to the center and corners of all of the segments on the surfaces of the sensors. These vectors will represent the direction of scattered light and the corners of a spherical rectangle that the light intercepts. The laser intensity at the point the particle intercepts the beam is then

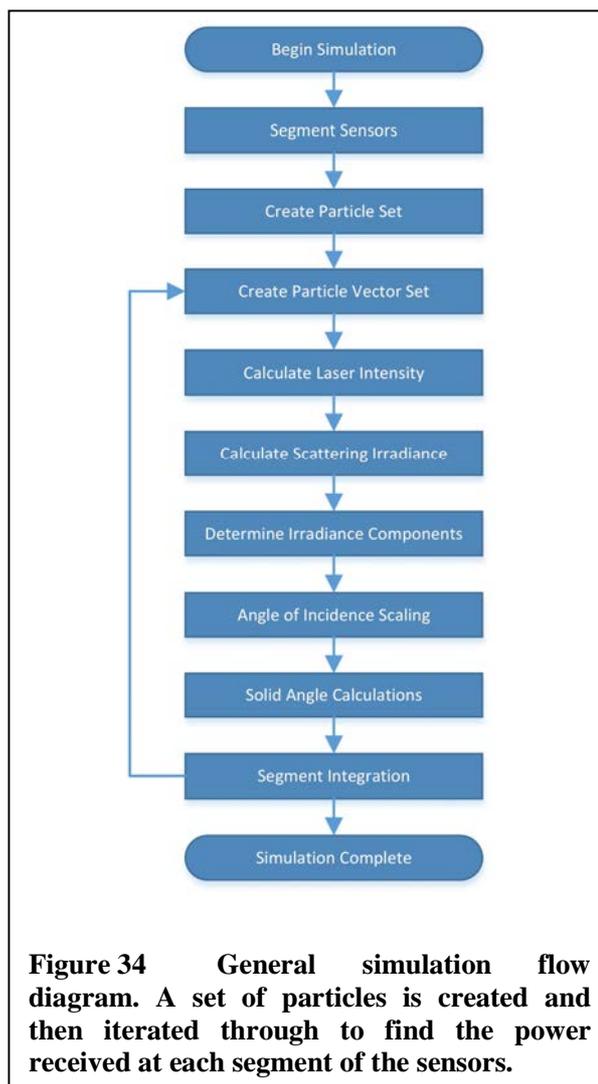


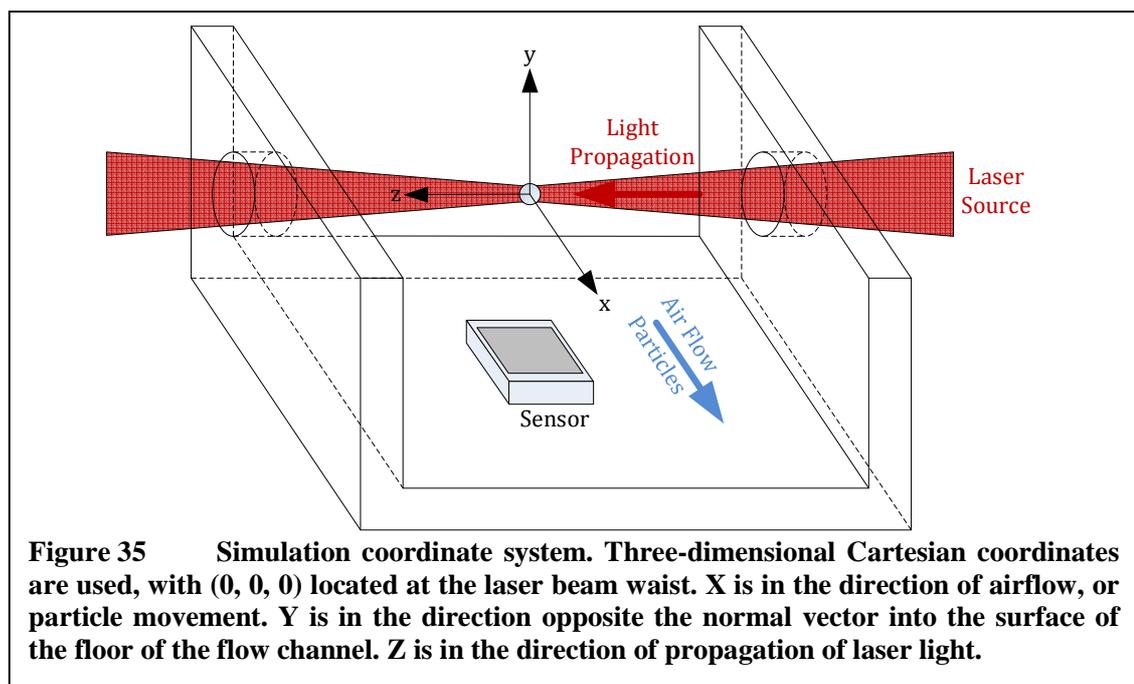
Figure 34 General simulation flow diagram. A set of particles is created and then iterated through to find the power received at each segment of the sensors.

calculated, and from this the Mie scattering values are found to determine the irradiance at the particle. As polarization plays a strong role in scattering theory, the polarization components relative to the scattering plane are found, and are then used to scale the light intensity based on the angle of incidence with the sensor surface. The power received at each rectangular sensor subsection is then scaled by the solid angle calculation. The subsections are then integrated to find the total power received at the sensor.

Each section of the simulation will be discussed in detail below. The simulations were developed and coded using MATLAB. Most of the plots shown in this section were directly created using the simulations with MATLAB plotting functions.

Coordinate System

The simulation uses a Cartesian coordinate system for most calculations within the flow channel. Figure 35 shows the coordinate system centered with the origin $(0, 0, 0)$



on the laser beam waist. In the figure, the laser light is propagating across the flow channel from right to left between two holes in the flow channel walls. The air flow through the channel is directed out of the page.

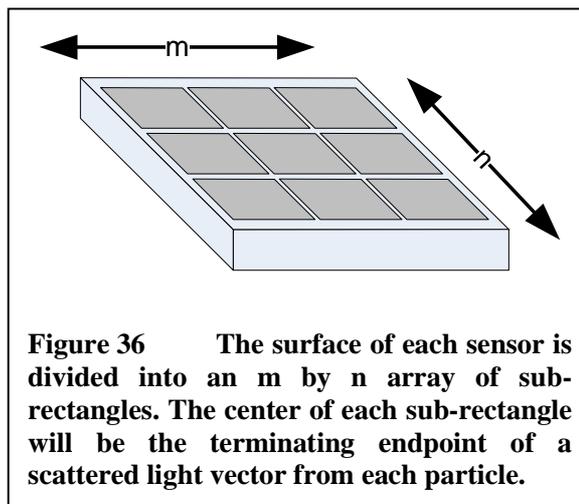
The positive X-axis of the coordinate system is aligned with the air flow through the channel, which is also the direction of particle movement. The positive Y-axis is opposite the normal of the floor of the flow channel. The positive Z-axis is aligned with the direction of laser light propagation. This convention was chosen to align with the typical coordinate systems used in Mie scattering calculations, as will be shown in subsequent sections of this chapter.

Segmenting Sensors

The simulation was performed on two types of sensor configurations. The first assumed a sensor lying flat on the floor of the flow channel, with the normal to the surface of the sensor being opposite the Y-axis of the coordinate system. The position of the sensor in X, Y, and Z directions might change, but the normal to the surface of the sensor was always in the opposite direction to the Y-axis. This type of sensor position fits the surface mount photodiode used for much of the particle counter development.

The second type of sensor configuration assumed a sensor standing vertically in the flow channel, with the normal to the surface of the sensor being orthogonal to the Y-axis. The direction of the sensor surface normal could be any arbitrary arrangement as long as it was orthogonal to the Y-axis. This type of sensor position fits the through-hole photodiode used in the “orthogonal” dual-sensor arrangements.

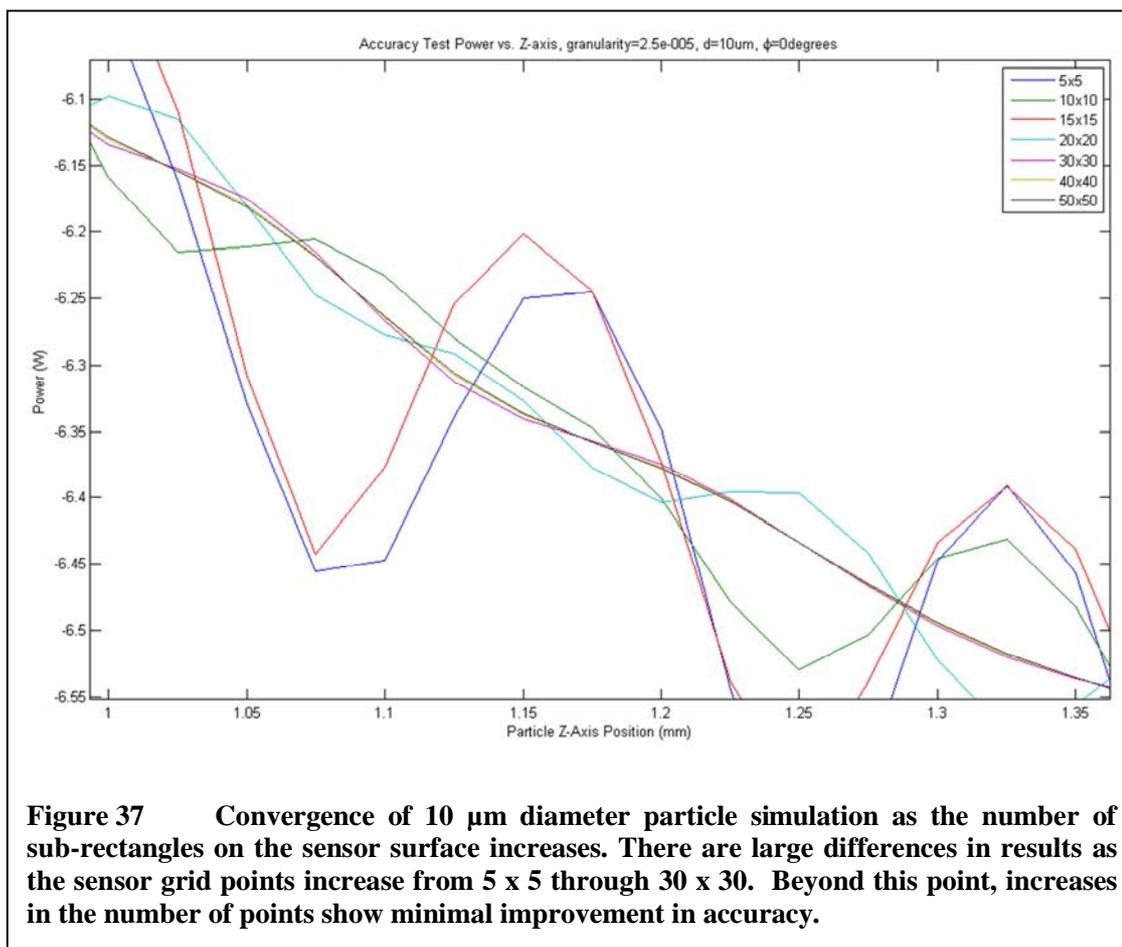
Each type of sensor was divided into a two-dimensional array of rectangular areas as shown in Figure 36. The center of each sub-rectangle in the m by n array will be the endpoint of a vector from a particle intercepting the laser and the sensor surface. This vector will be the direction of light scattered



from the particle to the sensor sub-rectangle. The subsections are rectangular to facilitate the calculation of solid angles intercepting the scattered light. These rectangles will be divided into two symmetrical triangles for the purpose of the solid angle calculations.

The number of sub-rectangles in the array determines the number of individual scattering irradiance calculations for each particle. This in turn impacts the accuracy of the overall calculation of the amount of scattered light that is intercepted by the sensor, since the granularity of the simulation will constrain the ability to reconstruct the scattered light irradiance on the sensor.

To find the optimal grid granularity, the simulation output was examined with an increasing number of sub-rectangles. Figure 37 shows the results of this search in the form of the power captured at the sensor from light scattered by a $10\ \mu\text{m}$ particle in the center axis of the laser beam. The particle is moved parallel to the beam in $25\ \mu\text{m}$ increments, remaining in the center of the beam the entire time. While this is not the motion expected of a typical particle in the WAQM optical particle counter, it does show the variation due to the simulation's granularity. A particle with $10\ \mu\text{m}$ diameter was



chosen for the simulation since the complexity of the scattered light pattern increases with increasing particle size, and a 10 μm particle was at the upper end of the size of interest for particles detected by the OPC device.

In Figure 37, it can be seen that the variation across the sensor surface differs significantly as the granularity goes from 5 x 5 up through 30 x 30. Beyond this point, the variation appears to have converged on a characteristic signature, with very little difference between 30 x 30 and finer granularities. Since the number of calculations increases by the square of the grid dimension, the simulation was limited to a grid of 30 x 30 in most cases to avoid increased simulation durations and memory requirements.

Particle Sets

The simulations were typically run on sets of particles, rather than just single particles intercepting the beam. The particle sets determined the size, composition, and position in the flow channel of each particle. These were the variables most often modified between runs to determine their impact on scattering intensity.

The composition of particles was defined in terms of their complex indices of refraction. The real part of the refractive index describes refraction, and the complex portion describes the absorption. For the most part, each simulation was run with a constant index of refraction for all particles, which would then be compared against a similar run with particles of a different index of refraction.

In some cases, the particle set defined a line of particles across the length of the laser beam in the channel to calculate the various scattering intensities of particles with positions across the channel. In other cases, the positions might vary randomly to create a simulation close to what might be seen during actual operation. Each type of particle set might require a different type of back-end processing to display the data in a way that was most visually meaningful. Examples of this will be given in following results sections.

Particle Vectors

For each simulated particle, a set of vectors is generated to each sensor in the system. Two different categories of vectors are created: vectors from the particle to the center of each sub-rectangle on the surface of the sensor, and vectors from the particle to each corner of each sub-rectangle on the surface of the sensor. The first type vector is required to calculate the direction of scattered light from the particle to the sensor

surface. This is an important factor since the scattering direction can greatly impact the intensity of the scattered light. The light intensity at the center point will be used across the surface of a solid angle to calculate the radiated power received by each sensor sub-rectangle. The second type of vector is required to calculate the solid angle intercepted by each sub-rectangle as seen by a sphere of radiated energy centered on the particle. As will be discussed below, the sub-rectangle is further split into two triangles to perform the solid angle calculation.

Figure 38 shows two example sets of vectors going from the particle to a pair of sensors arranged in an “orthogonal” scheme with the sensor surfaces in different planes. The blue vectors define the path from the particle to the corner of each sub-rectangle of a sensor surface. The red vectors define the path from the particle to the center of each sub-rectangle.

This example shows vectors for a 3 x 3 grid of rectangles on the surface of each

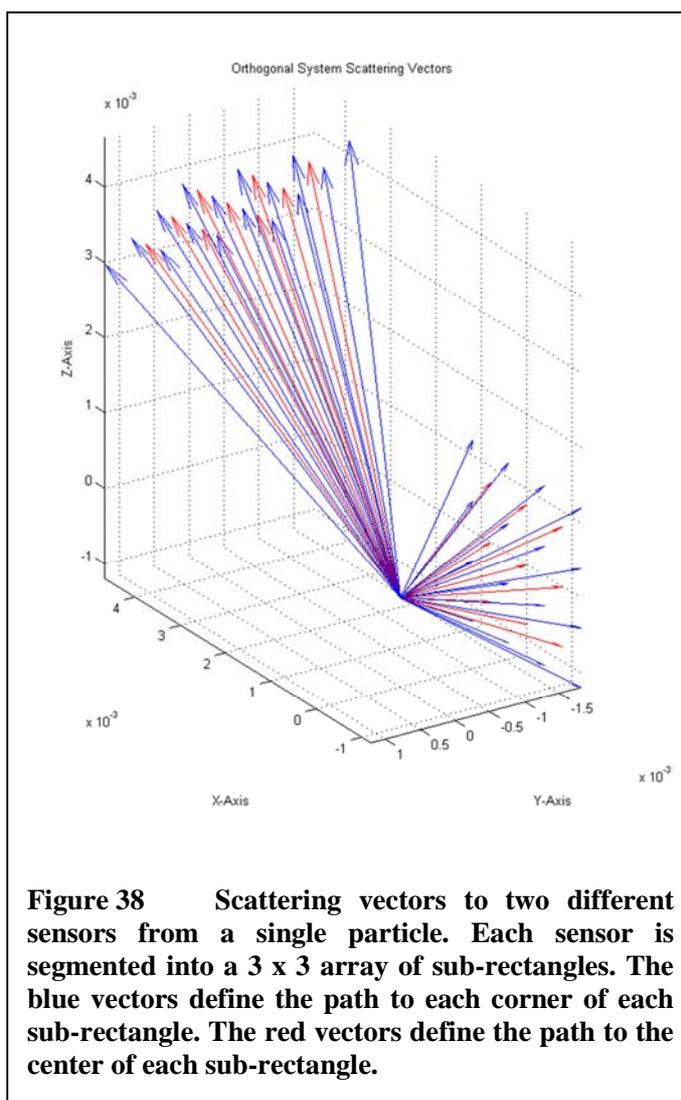


Figure 38 Scattering vectors to two different sensors from a single particle. Each sensor is segmented into a 3 x 3 array of sub-rectangles. The blue vectors define the path to each corner of each sub-rectangle. The red vectors define the path to the center of each sub-rectangle.

sensor. As mentioned previously, the number of sub-rectangles impacts the accuracy of the simulation and the runtime and memory requirements of the simulation. A grid of 30 x 30 rectangles was commonly used for all of the simulations, including the results shown below.

Laser Intensity and Polarization

The laser is the most critical component of the WAQM OPC system, since its power output, beam shape, and stability can greatly impact the performance of the device. With the low-cost constraints on the system, the laser will typically be relatively low power. The laser used in the WAQM OPC is nominally a 3 mW output device with an adjustable focus. The focus is necessary to create a high-intensity sampling volume that will create a detectable amount of scattering from even very small submicron particles. For the WAQM OPC, the laser is focused at the center point of the flow channel, creating a high-intensity beam waist at that point. The laser intensity falls off from the beam waist out to the edges of the flow channel in both directions.

The simulation flow will allow for the insertion of an arbitrary model for the laser intensity across the channel. For the WAQM OPC simulation, the laser is treated as a simple Gaussian beam, which follows a Gaussian function across the cross-section of the beam. Other models for the beam intensity may be inserted according to the needs of the simulation. While more accurate models may be used for diode lasers [65], the simple Gaussian model was deemed sufficient to meet the needs of the WAQM OPC design effort. Further study and analysis of the characteristics of the laser might lead to improvements in modeling accuracy for future projects.

Assuming an ideal Gaussian beam, the radius of the beam at the waist ω_0 may be determined from the wavelength of the laser λ , the focal length F , and the diameter of the aperture D .

$$\omega_0 = \left(\frac{2\lambda}{\pi}\right) \left(\frac{F}{D}\right) \quad 5.1$$

This is the radius at the beam waist measured to the point where the beam irradiance is equal to $1/e^2$ of the peak beam irradiance. For the WAQM laser module with wavelength of 650 nm, an aperture of 3 mm, and a focal length of 3.5 mm, the ideal beam waist will be less than 10 μm in diameter.

The radius of the beam to the $1/e^2$ irradiance at any arbitrary point along the beam axis $\omega(z)$ can be found with the equation

$$\omega(z) = \omega_0 \sqrt{1 + \left(\frac{\lambda z}{\pi \omega_0^2}\right)^2} \quad 5.2$$

The irradiance at the center of the Gaussian beam, at an arbitrary point along the beam axis z , is

$$I_0 = \frac{2P_0}{\pi \omega^2(z)} \quad 5.3$$

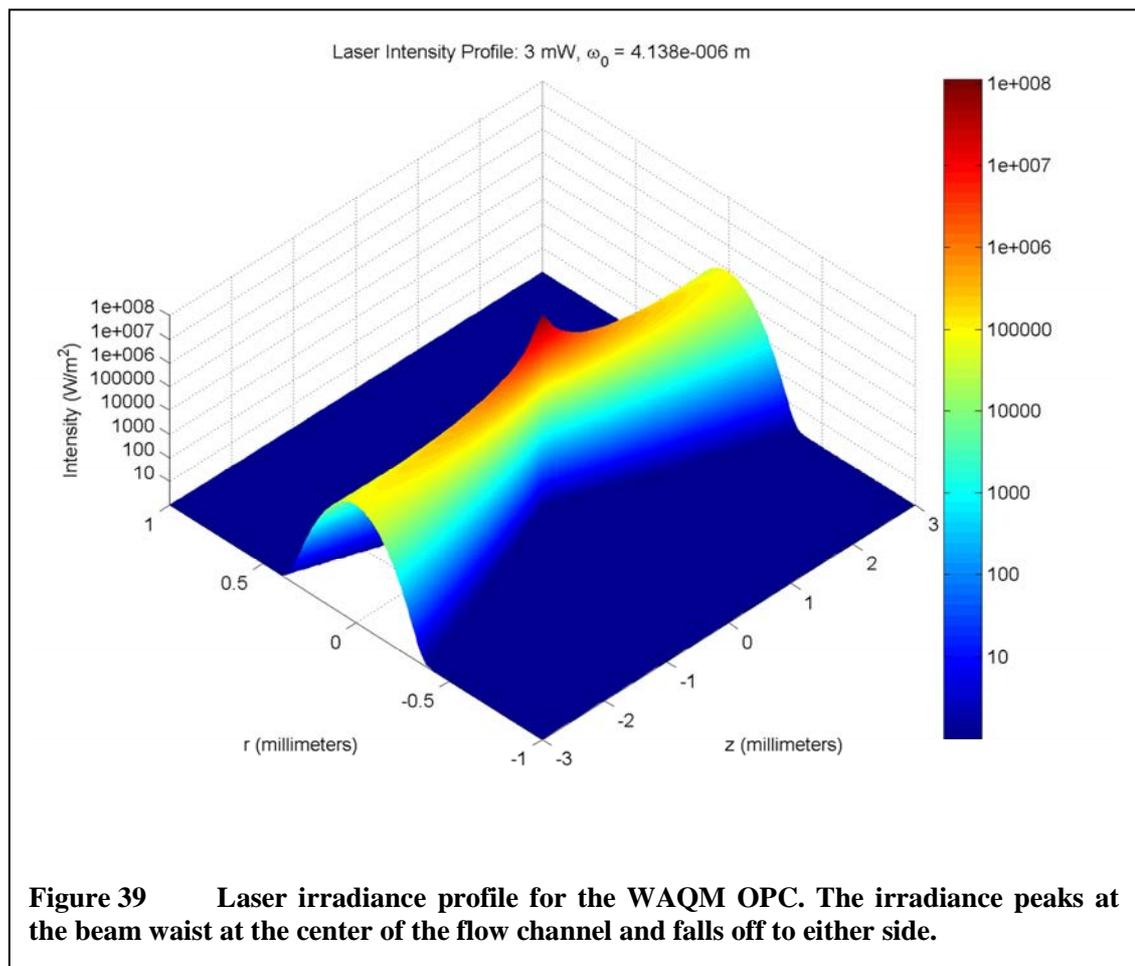
where P_0 is the total power transmitted by the beam. The small ideal waist size of the WAQM laser will produce an irradiance at the center of the waist that is nearly 8.2×10^7 Watts/m². This high intensity is what allows the WAQM OPC to detect very small particles with an inexpensive laser.

The irradiance of the Gaussian beam relative to the direction of light propagation z and distance from the beam axis r , is

$$I(r, z) = I_0 \left(\frac{\omega_0}{\omega(z)} \right)^2 e^{-\frac{2r^2}{\omega^2(z)}} \quad 5.4$$

which is the equation used in finding the irradiance for an arbitrary particle location. A plot of the irradiance for the WAQM OPC laser is shown in Figure 39. Note how the irradiance peaks at the waist, and falls off on either side. This irradiance is then used in the subsequent scattering calculations to determine the scattered irradiance in any given direction.

The degree of polarization of the laser is a significant factor in the determination of scattering irradiance from a particle, as will be described in subsequent sections. The degree of polarization can be measured with a polarizer, measuring the maximum and



minimum transmitted irradiances across the range of polarizer orientation. The degree of polarization V can be defined as

$$V = \frac{I_{max} - I_{min}}{I_{max} + I_{min}} \quad 5.5$$

with I_{max} and I_{min} being the maximum and minimum irradiance respectively, measured through the polarizer.

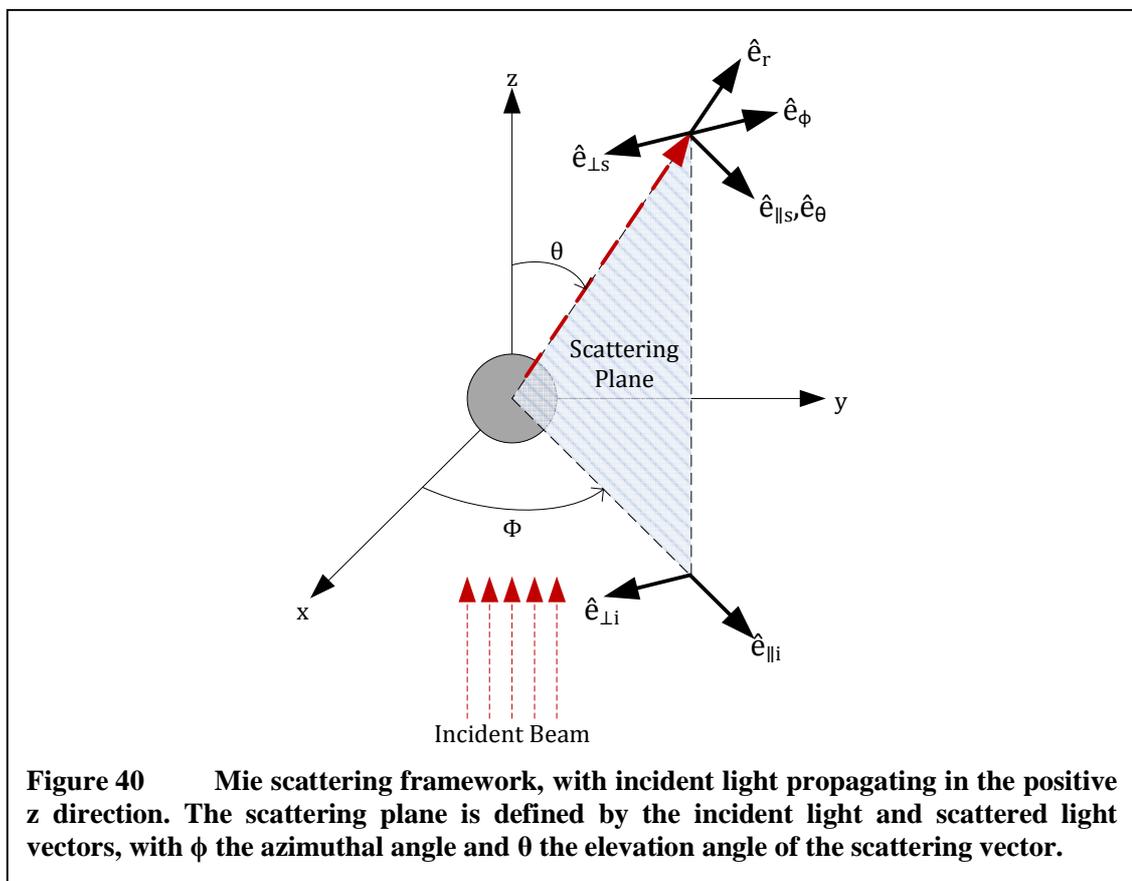
For the laser diode used in the WAQM OPC, the degree of polarization was measured using a sheet polarizer with 6000:1 extinction ratio and a photodiode-based power meter. The degree of polarization was found to be 98.8%, meaning that nearly all of the light emitted from the laser is polarized in a single direction. With such a high degree of polarization, it might be assumed that one could ignore the effect of the 1.2% of light that is orthogonally polarized. However, it should be kept in mind that at certain points in a given scattering pattern one might see a difference of 20dB or more between the parallel and perpendicular components of scattered light. For narrow ranges of angles, the contribution of orthogonally polarized light may become significant.

For the purposes of these simulations, it is assumed that all of the light has a single polarization. Given the high degree of polarization of the laser and the relatively wide angle of intercept of the sensor, this is a reasonable approximation for the purposes of exploring configurations and matching with actual results. It should be kept in mind however that a large variation in the polarization-dependent radiation patterns may cause this to be a significant factor for certain configurations.

Scattering Irradiance

The scattering irradiance calculations are the heart of the simulation flow. These calculations follow the work of Gustav Mie [64] in calculating the light scattering from small spheres. This work has been presented by Bohren and Huffman [66] as a set of Fortran routines. A conversion of these routines to MATLAB by Matzler [67] was used as a starting point for this portion of the simulation flow. Further optimizations were made to the algorithms to facilitate the calculation of scattering from multiple particles of the same size and refractive index simultaneously, without recalculating the scattering parameters for each.

The basic framework for Mie scattering is shown in Figure 40. The incident light



propagates in the positive z direction. The scattered light direction is defined by the azimuthal angle ϕ and elevation angle θ . The scattering plane is defined by the incident light and scattered light vectors. These terms and definitions will be used in the scattering calculations below.

The key components of the scattering calculations are a set of Mie coefficients a_n , b_n , c_n , and d_n . The first two coefficients are used to calculate the scattered field, and the second two are used to find the field internal to the sphere. For the purposes of this simulation, only the scattering coefficients are of interest, and are given as

$$a_n = \frac{m\psi_n(mx)\psi'_n(x) - \psi_n(x)\psi'_n(mx)}{m\psi_n(mx)\xi'_n(x) - \xi_n(x)\psi'_n(mx)} \quad 5.6$$

$$b_n = \frac{\psi_n(mx)\psi'_n(x) - m\psi_n(x)\psi'_n(mx)}{\psi_n(mx)\xi'_n(x) - m\xi_n(x)\psi'_n(mx)} \quad 5.7$$

where m is the refractive index of the sphere, x is the size parameter, and ψ_n and ξ_n are the Riccati-Bessel functions. The size parameter x is the radius of the sphere a , times the wavenumber k .

$$x = ka = \frac{2\pi}{\lambda} a \quad 5.8$$

The Riccati-Bessel functions are defined as

$$\psi_n(\rho) = \rho j_n(\rho) \quad 5.9$$

$$\xi_n(\rho) = \rho h_n^{(1)}(\rho) \quad 5.10$$

where j_n and h_n are the spherical Bessel and spherical Hankel functions. The spherical Hankel function is a combination of spherical Bessel functions.

$$h_n^{(1)}(\rho) = j_n(\rho) + iy_n(\rho) \quad 5.11$$

The spherical Bessel functions relate to Bessel functions as

$$j_n(\rho) = \sqrt{\frac{\pi}{2\rho}} J_{n+1/2}(\rho) \quad 5.12$$

$$y_n(\rho) = \sqrt{\frac{\pi}{2\rho}} Y_{n+1/2}(\rho) \quad 5.13$$

where J_n and Y_n are Bessel functions of the first and second kind, respectively.

The scattering coefficients are used to calculate a set of scattered field amplitude parameters S_1 and S_2 .

$$S_1(\theta) = \sum_n \frac{2n+1}{n(n+1)} (a_n \pi_n + b_n \tau_n) \quad 5.14$$

$$S_2(\theta) = \sum_n \frac{2n+1}{n(n+1)} (a_n \tau_n + b_n \pi_n) \quad 5.15$$

The functions π_n and τ_n are the recurrence relations

$$\pi_n = \frac{2n-1}{n-1} \pi_{n-1} \cos\theta - \frac{n}{n-1} \pi_{n-2} \quad 5.16$$

$$\tau_n = n \pi_n \cos\theta - (n+1) \pi_{n-1} \quad 5.17$$

where $\pi_0 = 0$ and $\pi_1 = 1$. The recurrence relations are limited to a number of iterations that is dependent on the size parameter x , as recommended by Wiscombe [68].

The scattered electric field can then be found from the scattering parameters

$$E_{s\theta} = -E_0 \frac{e^{-ikr+i\omega t}}{ikr} \cos\phi S_2(\theta) \quad 5.18$$

$$E_{s\phi} = E_0 \frac{e^{-ikr+i\omega t}}{ikr} \sin\phi S_1(\theta) \quad 5.19$$

where $E_{s\theta}$ is the component of the scattered field with polarization parallel to the scattering plane defined by the directions of incident and scattered light, and $E_{s\phi}$ is the component of the scattered field orthogonal to this.

The irradiance can be found directly from the scattering parameters S_1 and S_2 , with the irradiance split into components polarized parallel and perpendicular relative to the scattering plane.

$$I_{\perp} = \sin^2(\phi - \phi_L) \frac{I(r, z) |S_1|^2}{k^2 r^2} \quad 5.20$$

$$I_{\parallel} = \cos^2(\phi - \phi_L) \frac{I(r, z) |S_2|^2}{k^2 r^2} \quad 5.21$$

The scattering is scaled by the incident light intensity $I(r, z)$ at the particle and the azimuthal angle of scattering ϕ relative to the incident light polarization ϕ_L . Since the laser used in the OPC is highly polarized, the scaling by azimuthal angle relative to the laser polarization is critical. The polarization of the incident light will greatly impact the irradiance in any given direction.

Irradiance Components

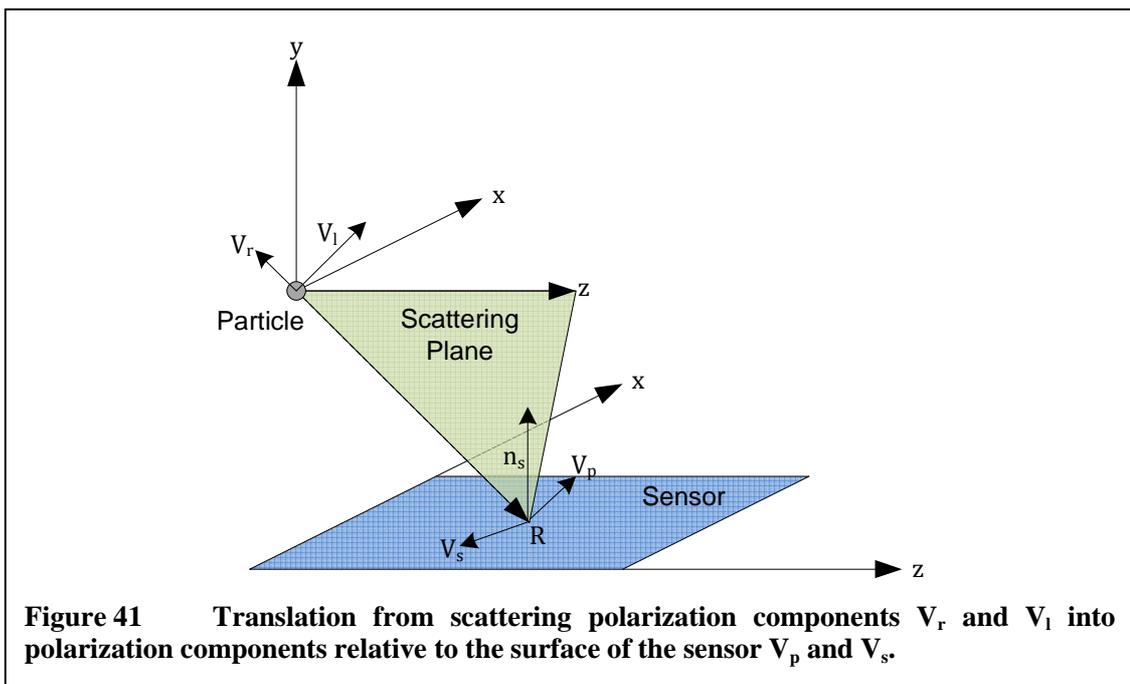
Now that the irradiance components parallel and perpendicular to the scattering plane have been found, they must be translated into components relative to the sensor surface. Once this is done, the impact of polarization on angle of incidence to the sensor

can be calculated. Figure 41 shows the basic translation for this step. Vectors representing direction of polarization for the scattered light are calculated from the direction of the scattered light, \mathbf{R} , the direction of incident laser light, \mathbf{z} , and the normal to the surface of the sensor, \mathbf{n}_s . The scattering plane is defined as the plane that contains both \mathbf{R} and \mathbf{z} . These vectors are represented in a Cartesian coordinate system, with the origin at the scattering particle. The magnitude of each vector is unimportant, as it is only the final vector directions that are important for this step.

\mathbf{V}_r is defined as the scattered component with polarization perpendicular to the scattering plane and \mathbf{V}_l as the scattered light component with polarization parallel to the scattering plane. The direction of these two vectors will correspond to the direction of the irradiance components I_{\perp} and I_{\parallel} .

$$\mathbf{V}_r = \mathbf{R} \times \mathbf{z}$$

5.22



$$\mathbf{V}_l = \mathbf{V}_r \times \mathbf{R} \quad 5.23$$

Relative to the sensor, the polarization components are parallel and perpendicular to the plane of incidence, defined as the plane containing the scattered light propagation vector \mathbf{R} and the normal to the surface of the sensor \mathbf{n}_s .

$$\mathbf{V}_s = \mathbf{R} \times \mathbf{S}_n \quad 5.24$$

$$\mathbf{V}_p = \mathbf{V}_s \times \mathbf{R} \quad 5.25$$

\mathbf{V}_s is the scattered component with polarization perpendicular to the plane of incidence and \mathbf{V}_p is the component with polarization parallel to the plane of incidence.

Only the angles between the vectors are of interest, so each vector is normalized to simplify the angle calculation.

$$\hat{\mathbf{V}}_r = \frac{\mathbf{V}_r}{\|\mathbf{V}_r\|}, \hat{\mathbf{V}}_l = \frac{\mathbf{V}_l}{\|\mathbf{V}_l\|}, \hat{\mathbf{V}}_s = \frac{\mathbf{V}_s}{\|\mathbf{V}_s\|}, \hat{\mathbf{V}}_p = \frac{\mathbf{V}_p}{\|\mathbf{V}_p\|} \quad 5.26$$

The cosine of the angle between each pair of normalized vectors may then be calculated using the dot product.

$$\cos \theta_{pr} = \hat{\mathbf{V}}_r \cdot \hat{\mathbf{V}}_p, \cos \theta_{pl} = \hat{\mathbf{V}}_l \cdot \hat{\mathbf{V}}_p, \cos \theta_{sr} = \hat{\mathbf{V}}_r \cdot \hat{\mathbf{V}}_s, \cos \theta_{sl} = \hat{\mathbf{V}}_l \cdot \hat{\mathbf{V}}_s \quad 5.27$$

The incident irradiance for each polarization component relative to the sensor surface is then calculated by Malus's law [69].

$$I_{pi} = I_{\perp} \cos^2 \theta_{pr} + I_{\parallel} \cos^2 \theta_{pl} \quad 5.28$$

$$I_{si} = I_{\perp} \cos^2 \theta_{sr} + I_{\parallel} \cos^2 \theta_{sl} \quad 5.29$$

These irradiance components will be used in the next step to scale the received light by the angle of incidence through the sensor surface material.

Angle of Incidence Scaling

Since light from the laser is polarized, the effect of this on the transmission onto the photodiode surface becomes important. This is especially true since the close proximity of the laser to the photodiode can create some very acute angles between the light and the surface.

The surface of the photodiode package will transmit different irradiances to the semiconductor based on the polarization components of the incident light. The transmittance T , or ratio of transmitted to incident irradiance, is found using the Fresnel equations [69]

$$T = \frac{I_t \cos \theta_t}{I_i \cos \theta_i} = \frac{n_t \cos \theta_t}{n_i \cos \theta_i} \left(\frac{E_{0t}}{E_{0i}} \right)^2 = \frac{n_t \cos \theta_t}{n_i \cos \theta_i} t^2 \quad 5.30$$

where t is the transmission amplitude coefficient specific to the polarity of the incident light. Substituting the coefficient specific to the parallel and perpendicular components yields two equations for the intensity components

$$\begin{aligned} T_s &= \frac{n_t \cos \theta_t}{n_i \cos \theta_i} t_s^2 = \frac{n_t \cos \theta_t}{n_i \cos \theta_i} \left(\frac{2n_i \cos \theta_i}{n_i \cos \theta_i + n_t \cos \theta_t} \right)^2 \\ &= \frac{4n_i n_t \cos \theta_i \cos \theta_t}{(n_i \cos \theta_i + n_t \cos \theta_t)^2} \end{aligned} \quad 5.31$$

$$\begin{aligned} T_p &= \frac{n_t \cos \theta_t}{n_i \cos \theta_i} t_p^2 = \frac{n_t \cos \theta_t}{n_i \cos \theta_i} \left(\frac{2n_i \cos \theta_i}{n_i \cos \theta_t + n_t \cos \theta_i} \right)^2 \\ &= \frac{4n_i n_t \cos \theta_i \cos \theta_t}{(n_i \cos \theta_t + n_t \cos \theta_i)^2} \end{aligned} \quad 5.32$$

where T_p is the transmittance for light polarized parallel to the sensor and T_s is the transmittance for light polarized perpendicular to the sensor. Scaling the incident

irradiance components by the transmittance and summing the results leads to the calculation of the final intensity seen at the sensor.

$$I = I_{si}T_s + I_{pi}T_p \quad 5.33$$

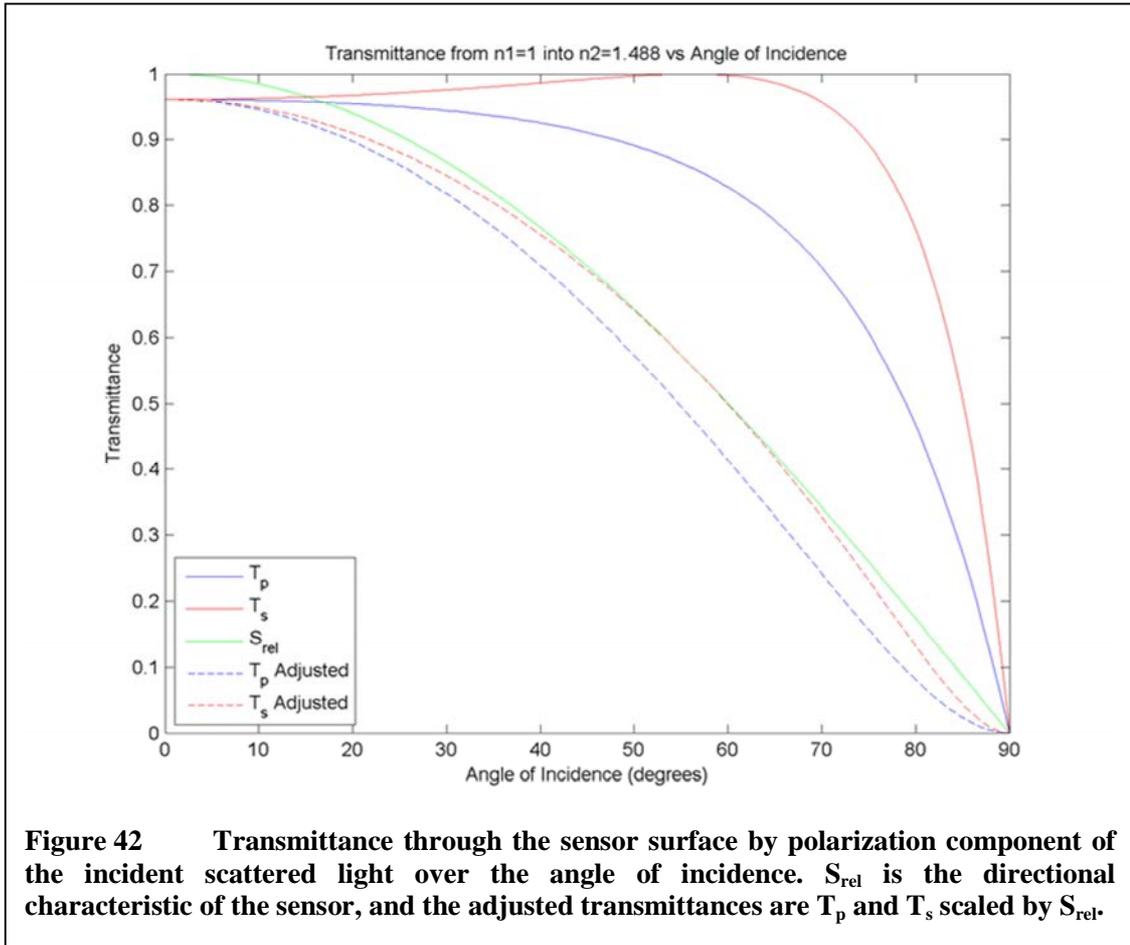
This intensity may be further scaled by a sensor-specific directional characteristic. For example, the photodiode used in the WAQM OPC has a relative irradiance scaling that is close to a cosine of the angle between the incident light and the sensor normal vector. Figure 42 shows the curves for transmittance through an acrylic surface material T_p and T_s . The curve S_{rel} is the directional characteristic of the sensor. T_p *Adjusted* and T_s *Adjusted* are the final values to be used as the irradiance captured by the sensor.

Solid Angle Calculations

Now that the irradiances have been found for the grid of points on the sensor surface, each irradiance value must be applied to its respective segment of the sensor. The irradiance in Watts per square meter must be multiplied by the effective area intercepted by the segment. The physical arrangement is a spherically expanding field intercepting a group of rectangular segments on a plane, and the solid angle intercepted by each segment will vary. To find the power received at each segment, the solid angle subtended by each rectangle on the surface of the sensor must be calculated.

Van Oosterom and Strackee [70] have derived a relationship between the solid angle Ω subtended by a plane triangle and the vectors from the center of the sphere to the corners of the triangle as

$$\tan\left(\frac{1}{2}\Omega\right) = \frac{\mathbf{R}_1 \cdot (\mathbf{R}_2 \times \mathbf{R}_3)}{R_1R_2R_3 + (\mathbf{R}_1 \cdot \mathbf{R}_2)R_3 + (\mathbf{R}_1 \cdot \mathbf{R}_3)R_2 + (\mathbf{R}_2 \cdot \mathbf{R}_3)R_1} \quad 5.34$$



which can be used to directly calculate the solid area of the triangle on the surface of the sphere. This solid angle is given in the following relation.

$$\Omega = 2 \tan^{-1} \left(\frac{\mathbf{R}_1 \cdot (\mathbf{R}_2 \times \mathbf{R}_3)}{R_1 R_2 R_3 + (\mathbf{R}_1 \cdot \mathbf{R}_2) R_3 + (\mathbf{R}_1 \cdot \mathbf{R}_3) R_2 + (\mathbf{R}_2 \cdot \mathbf{R}_3) R_1} \right) \quad 5.35$$

The triangle is described by the three Cartesian vectors \mathbf{R}_1 , \mathbf{R}_2 , and \mathbf{R}_3 . The numerator is the scalar triple product of the three vectors. The surface of the sensor is divided into a number of equally sized rectangles, the number chosen to match a desired accuracy of the numeric summation of intensities across the sensor surface. Each rectangle is split into two equally sized triangles, each described by three vectors from

the origin of the particle to the triangle corners. The solid angle is calculated from these vectors, and used to scale the irradiance in the calculation of the power received by the rectangular portion of the sensor. The solid angle is related to the surface area S on the sphere by the square of the sphere's radius.

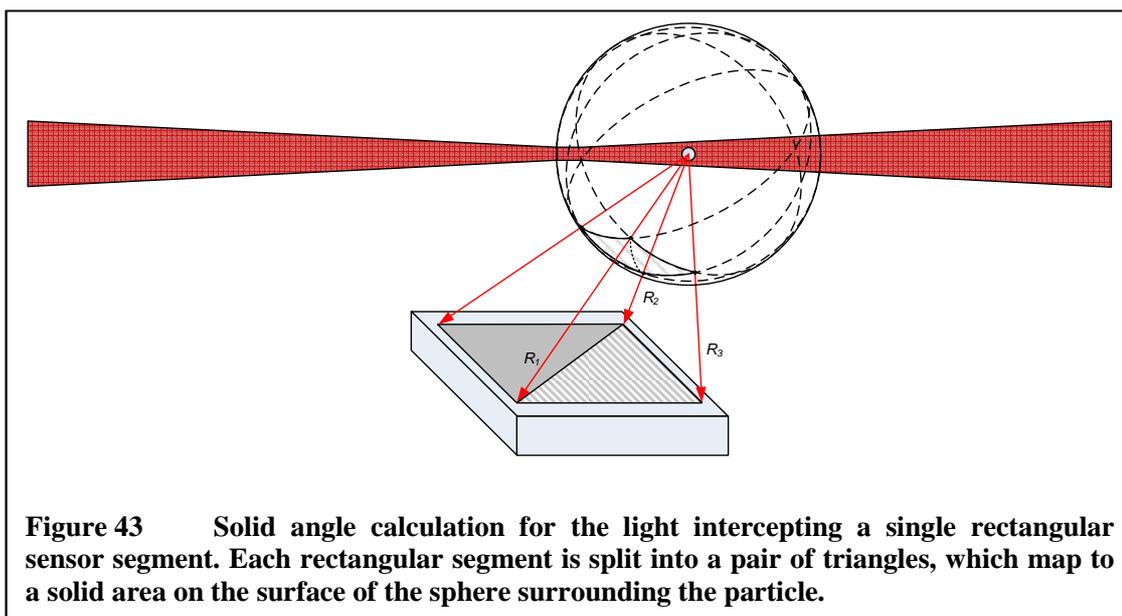
$$S = \Omega r^2 \quad 5.36$$

This can now be used to calculate the power received by the sensor.

Figure 43 shows a graphical representation of the solid angle calculation for one sub-triangle of a single sensor grid rectangle. The grid rectangle is split into two triangles, and the solid angle is found for each. This is repeated across the surface of the sensor until all grid points are covered. Note that the two triangles may not cover the same solid angle, depending on the angle of incidence to the sensor.

Segment Integration

As a final step in the simulation of the irradiance intercepted by the sensor, the



total irradiance across the sensor is calculated from a summation of the power intercepted by the solid area of each sub-rectangle. This received power can be found with

$$P = \sum_{i=1}^n \Omega_i r_i^2 I_i \quad 5.37$$

where Ω_i is the solid angle for each sub-area, r_i is the radial distance from the particle to the center of the sub-area, and I_i is the irradiance at the center of each sub-area. This final result can then be used to approximate the current produced by the photodiode sensor through its rated spectral sensitivity in Amperes per Watt. This current may then be used in models of the final amplifier to approximate the signal produced by the OPC transimpedance amplifier.

The simulation stops at this point in terms of taking the analysis of received power any deeper into the physics of the photodiode device, or other impacts of the following amplification circuitry. Further additions of these effects into the simulation may provide more accurate results, but are left to future work at this point in time. It is felt that for the purposes of the WAQM OPC, the level of accuracy produced by the above process is sufficient to provide relevant guidance on device construction.

Basic Scattering Simulations

To validate the basic scattering simulation, the correctness of the output scattering data from the simulation was compared to published data. This simulation isolates just the basic scattering irradiance step of the simulation flow. The text by Bohren and Huffman [66] gives example scattering irradiance plots for a water particle with complex refractive index of $1.33 + j10^{-8}$, illuminated with a 525 nm light source. Using the same parameters, the plot shown in Figure 44 was produced. This plot shows the scattering irradiance of light polarized both parallel and perpendicular to the scattering plane,

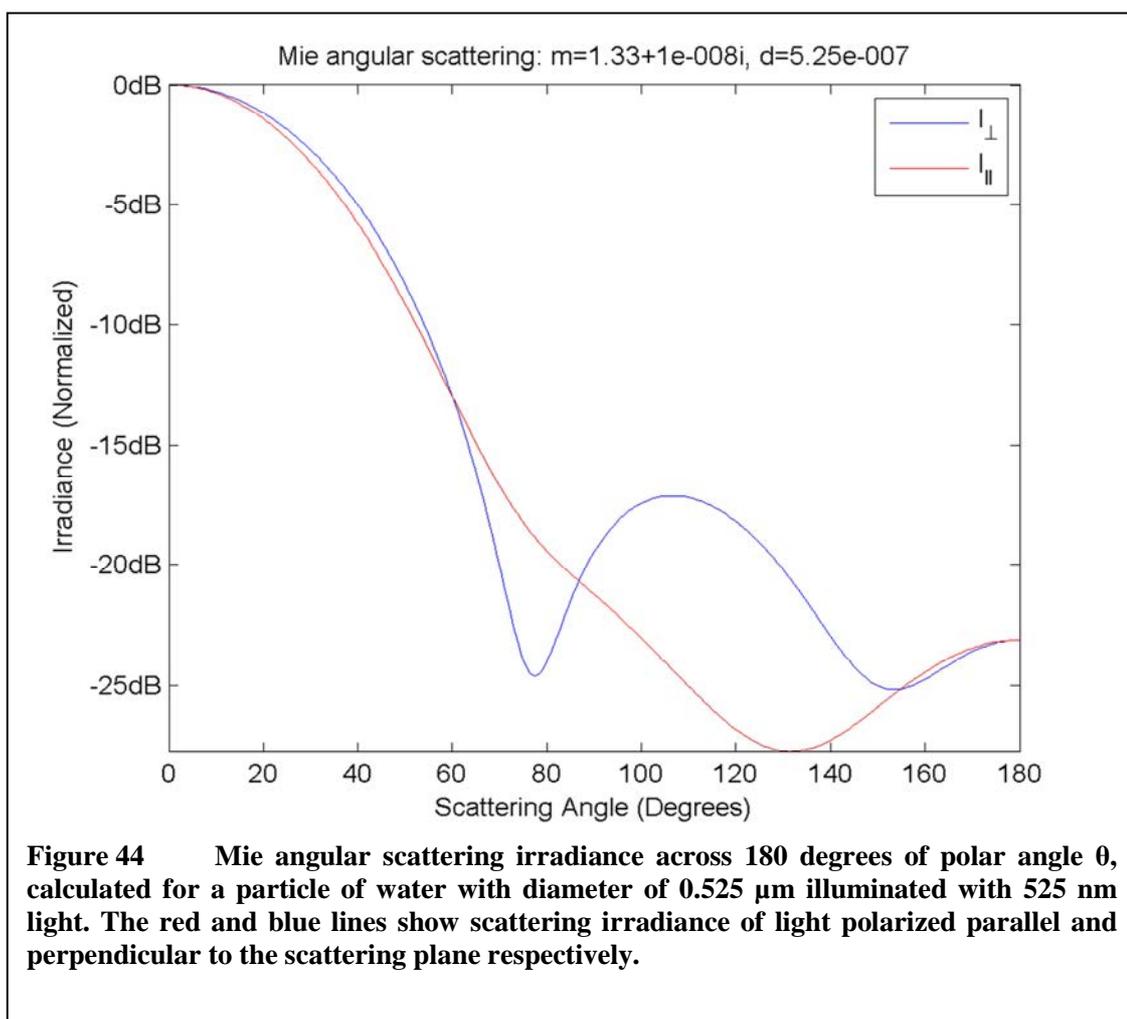
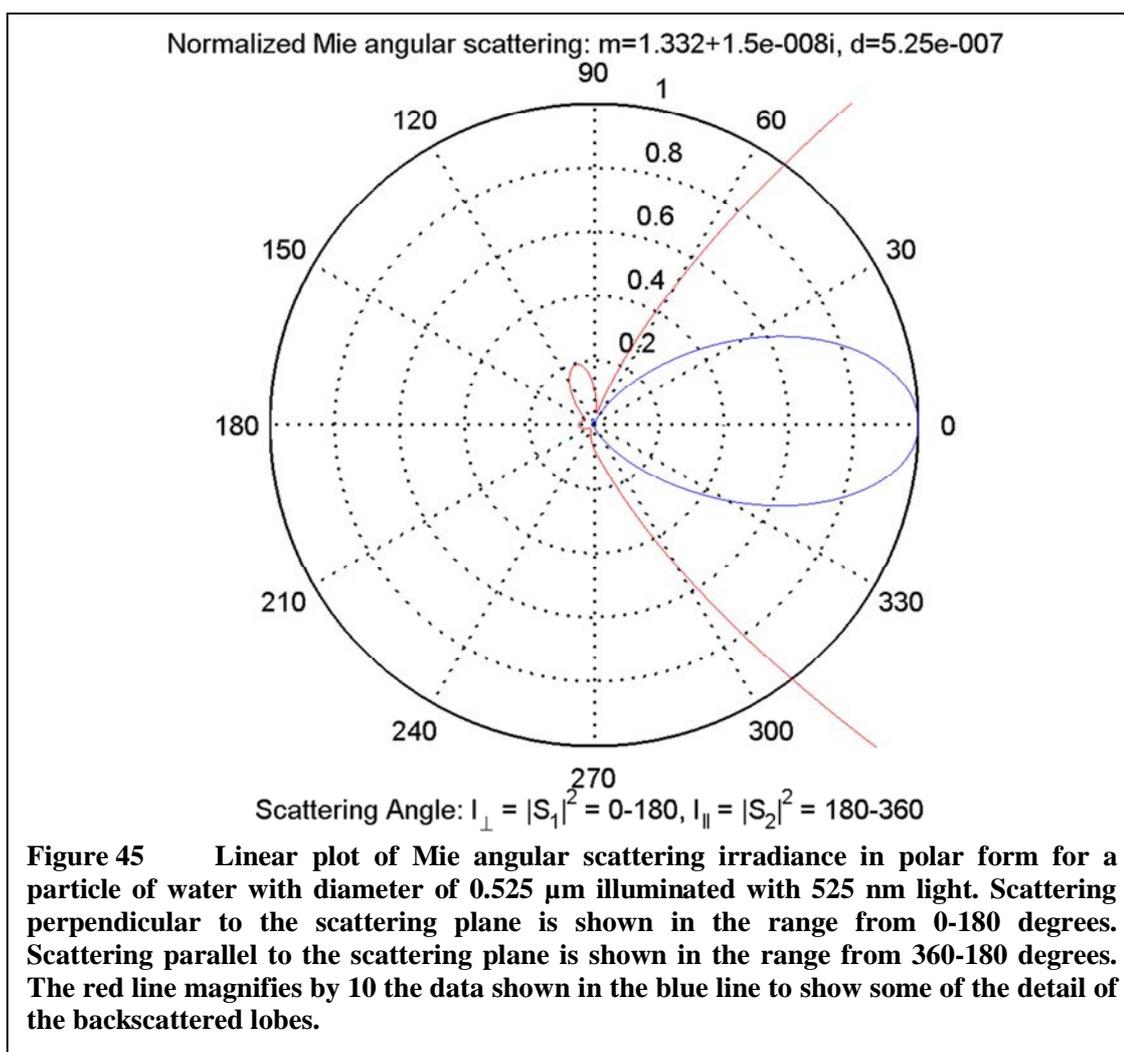


Figure 44 Mie angular scattering irradiance across 180 degrees of polar angle θ , calculated for a particle of water with diameter of 0.525 μm illuminated with 525 nm light. The red and blue lines show scattering irradiance of light polarized parallel and perpendicular to the scattering plane respectively.

varying the polar angle θ from 0 to 180 degrees. The result is identical to that shown in the text, with the strongest scattering in the forward, 0 degrees direction. Back-scattering in the 180 degree direction is more than 100 times less intense.

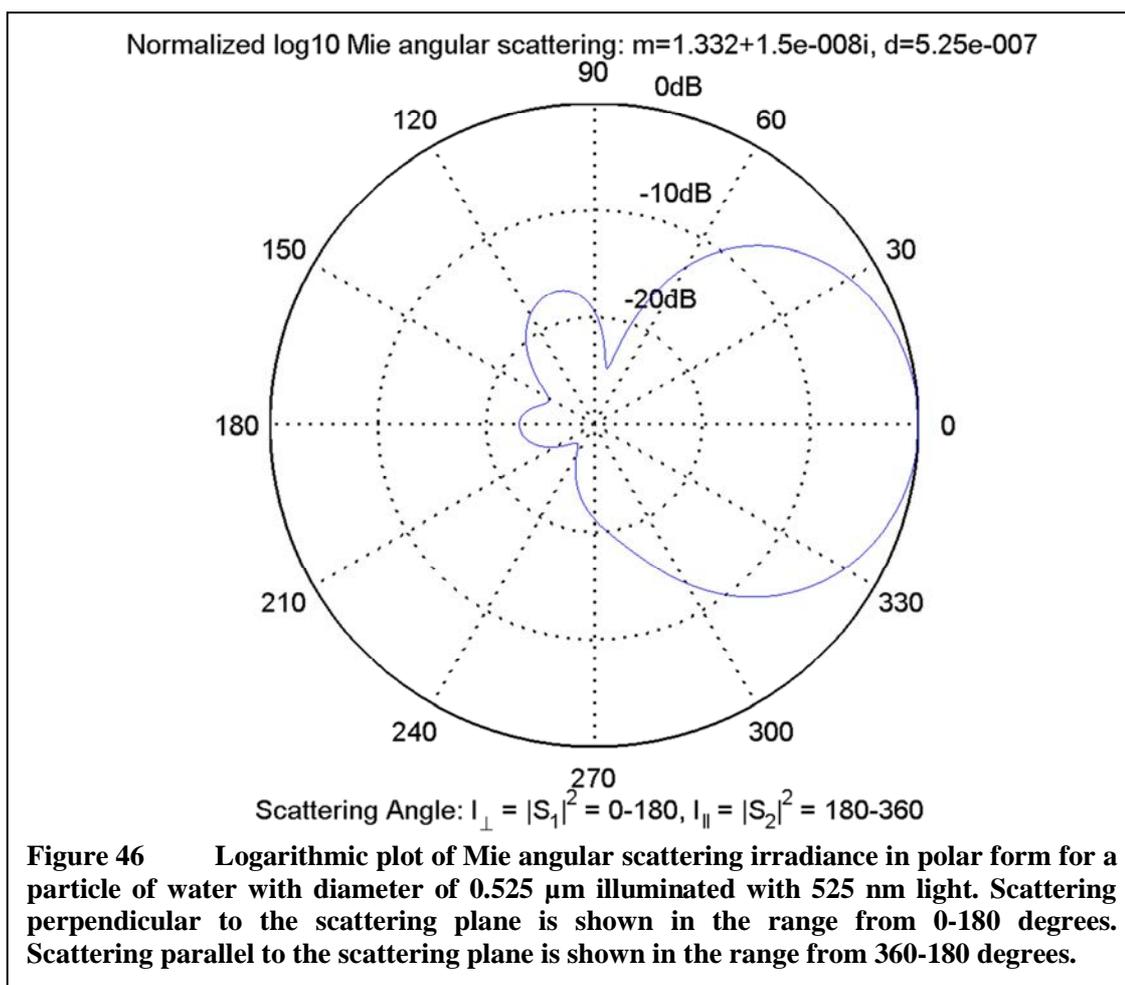
The data can also be represented in polar form, as shown in the linear plot of Figure 45. The light scattered with polarization perpendicular to the scattering plane is plotted from 0–180 degrees, and light scattered with polarization parallel to the scattering plane is plotted from 360–180 degrees. The 0 degree point corresponds to forward scattering, and is the same direction as incident light propagation in the system. The 180



degree point is in the direction of the light source, opposite to its direction of propagation.

Since the plot is linear, the low-intensity backscattering lobes don't show up in the normalized plot represented by the blue line, as they are less than $1/100^{\text{th}}$ the magnitude of the main forward scattering lobe. The red line shows the same linear data magnified by 10, which just barely shows the backscatter lobe structure.

Representing the polar plot on a logarithmic scale provides a more easily readable display of the scattering irradiance across the polar angle range. Figure 46 shows the scattering from the same $0.525 \mu\text{m}$ particle on such a logarithmic scale. Note that the data



from 0-180 degrees corresponds to the blue line in Figure 44, and the data from 360-180 degrees corresponds to red line in the same plot.

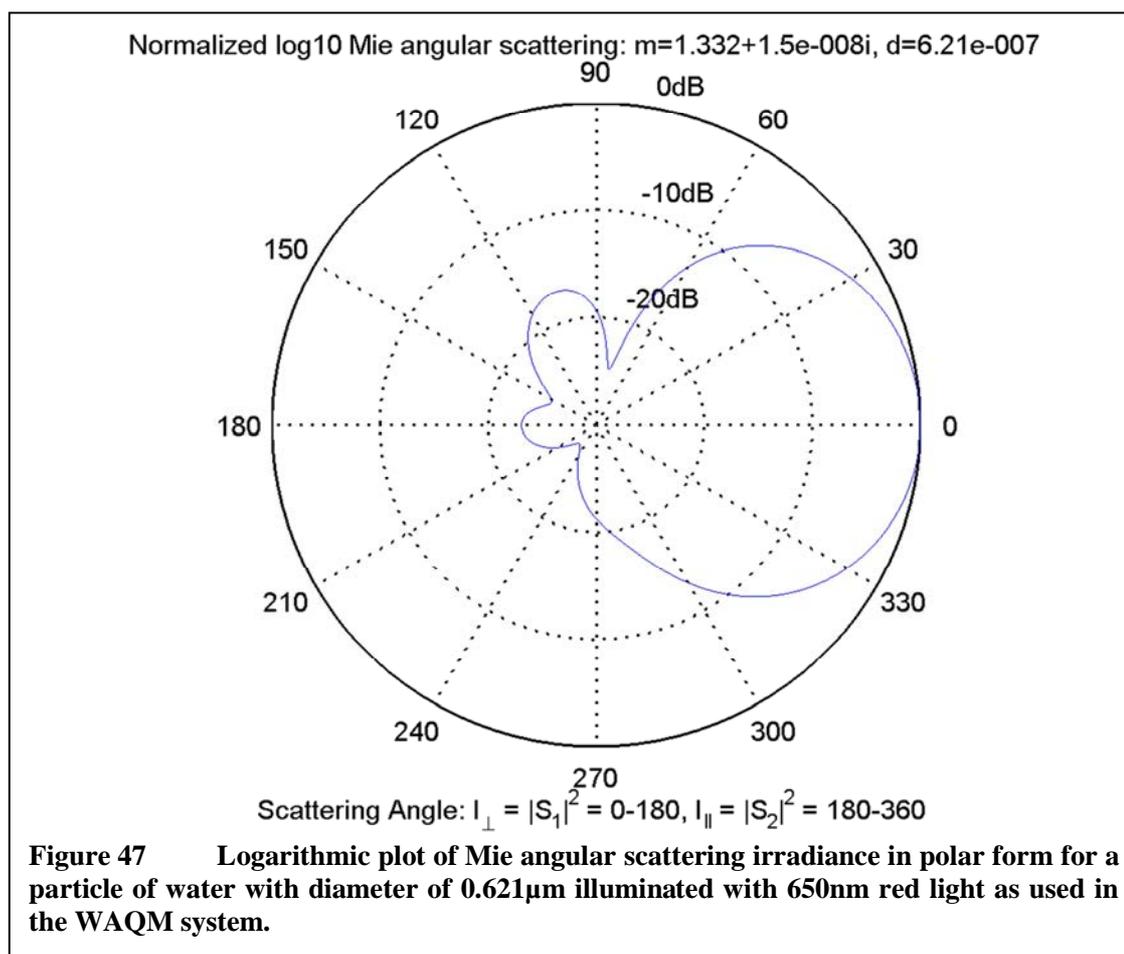
The illumination used in the preceding example was a greenish 525 nm, which does not correspond to the laser used in the WAQM system. The different wavelength of the red 650 nm laser used in the design will impact the scattering pattern and intensity as calculated in the scattering coefficients by the factor x , where

$$x = \kappa a = \frac{2\pi}{\lambda} a \quad 5.38$$

with a the radius of the particle. To check for similar output for the longer wavelength light, the radius that yields the same value of x for 650 nm is found, giving a particle diameter of 0.621 μm . The resulting polar plot with logarithmic scale is shown in Figure 47, which appears identical to the plot for the text example of Figure 46. This confirms that scaling the simulation to the WAQM system parameters will yield the correct results based on Mie calculations.

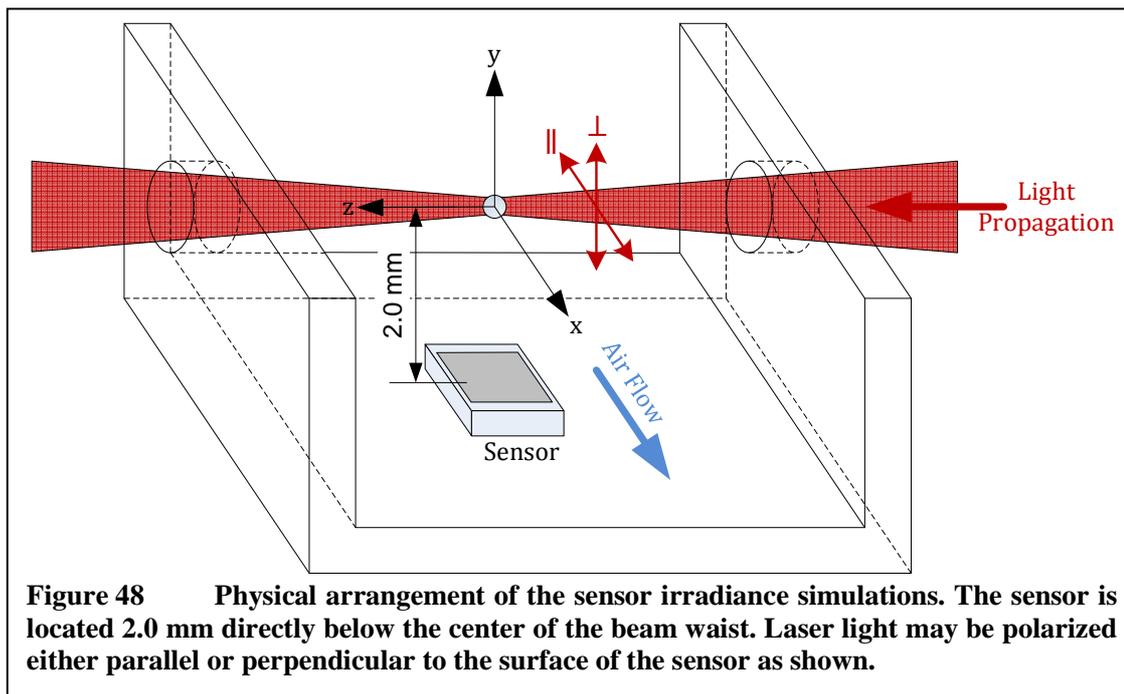
Sensor Irradiance Simulations

As a next step in validating the simulations, the irradiance levels across the sensor surface are displayed as surface plots and compared to the polar plots of normalized Mie scattering produced for the same particle. This type of simulation will provide a visualization of the pattern received at the sensor surface, which will allow a sanity-check comparison with the basic Mie scattering plots. The data displayed in this simulation are taken just prior to the segment integration step in the overall simulation flow.



For many of the following simulations, the material used for the particle composition is Polystyrene Latex (PSL). This material is commonly used in the calibration of optical particle counters, and is provided as concentrated particulate in solution that can be nebulized into an aerosol. PSL is used in the calibration of the WAQM OPC as well, and as such it is convenient to simulate with the same material for the purpose of comparing simulation data to actual data collected from the OPC in the calibration chamber. The refractive index used for this material is $1.59 + j5 \times 10^{-4}$ as shown by Ma et al [71].

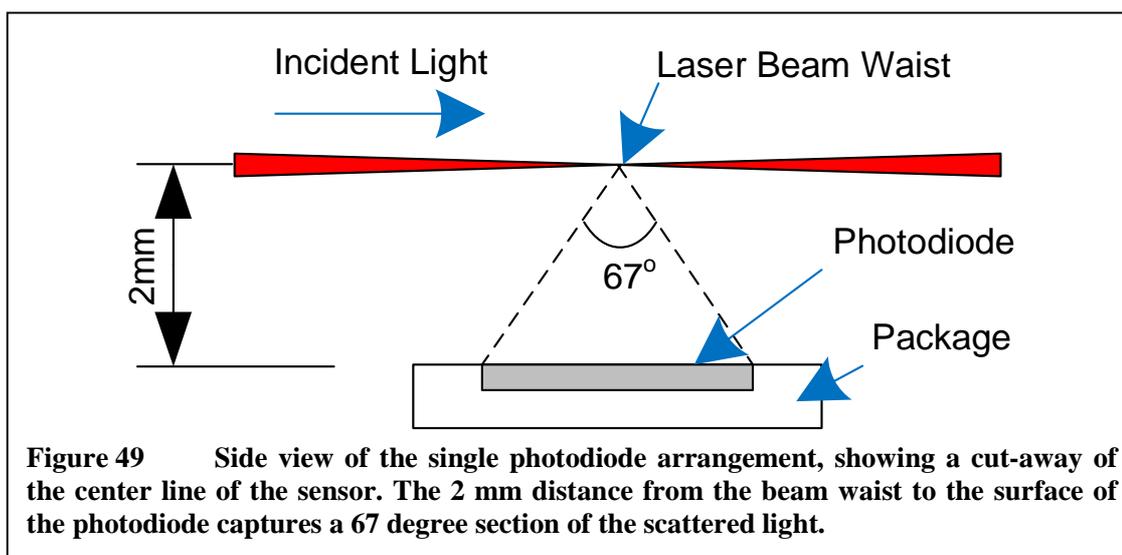
The physical arrangement of the system for this set of simulations is shown in Figure 48. This arrangement is similar to the design of the single-sensor WAQM OPC



with the beam waist located near to the sensor. More specifically, the sensor is positioned 2.0 mm below the center of the beam waist. The sensor itself is dimensioned to match the component used in the WAQM OPC, being 2.65 mm on each side with an area of 7 mm².

When a particle passes through the beam waist, this physical arrangement captures approximately 67 degrees of the scattered light along the center of the photodiode directly under the laser as shown in Figure 49. Note that this two-dimensional angle is important only in comparing the power received on the centerline of the photodiode to the polar plots of scattering intensity. The power received is calculated from the solid angle that is intercepted by the sensor, which takes the three-dimensional situation into account.

The polarization of the incident light is also set to a single orientation relative to the surface of the sensor. For the semiconductor diode used in the WAQM OPC, the emitted light is highly polarized. For the purposes of the simulations, it is assumed that



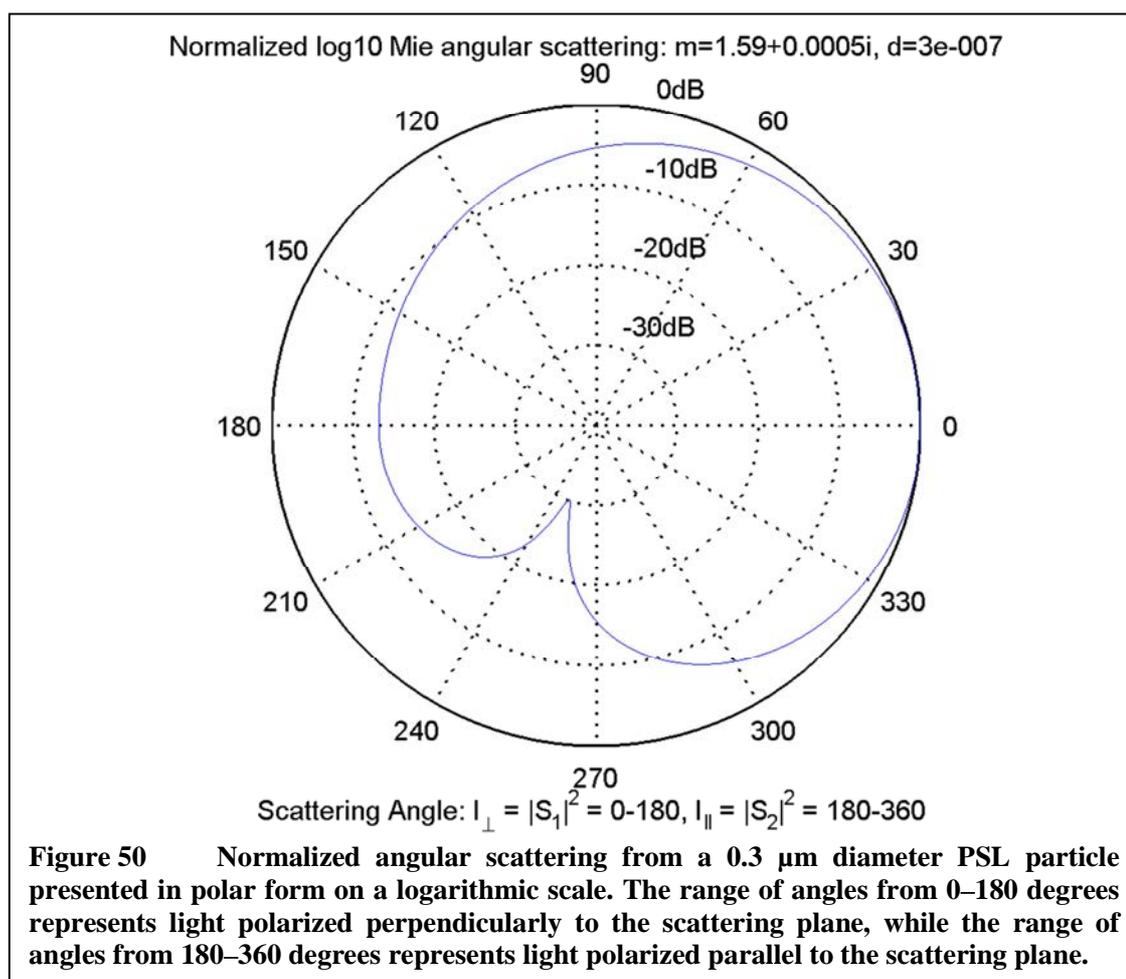
the incident light is completely polarized at a single orientation, typically either parallel or perpendicular relative to the surface of the sensor. The illumination wavelength was set to 650 nm to match that of the red laser used in the WAQM OPC, and is the wavelength used for all following simulations unless otherwise mentioned.

It is important to note that the sensor's position relative to a particle passing through the center of the beam waist will result in interception of scattered light with a scattering plane that is at or near perpendicular to the surface of the sensor. This means that incident light polarized parallel to the sensor surface scattering down into the sensor will be predominantly polarized perpendicular to the scattering plane. This relationship is critical to understand due to the strong impact of polarization on the irradiance seen at the sensor.

For an initial set of data, a small $0.3\mu\text{m}$ PSL particle is chosen due to its relatively smooth scattering pattern and clear differentiation in behavior between the two incident light polarization angles. The polar plot for normalized angular scattering for this scenario is shown in Figure 50 on a logarithmic scale. This is the basic Mie scattering

calculation for the particle, showing the scattering of light polarized both parallel and perpendicular to the scattering plane. Normalized scattering from light polarized perpendicular to the scattering plane is shown in the range from 0 to 180 degrees and from light polarized parallel to the scattering plane in the range from 360 to 180 degrees.

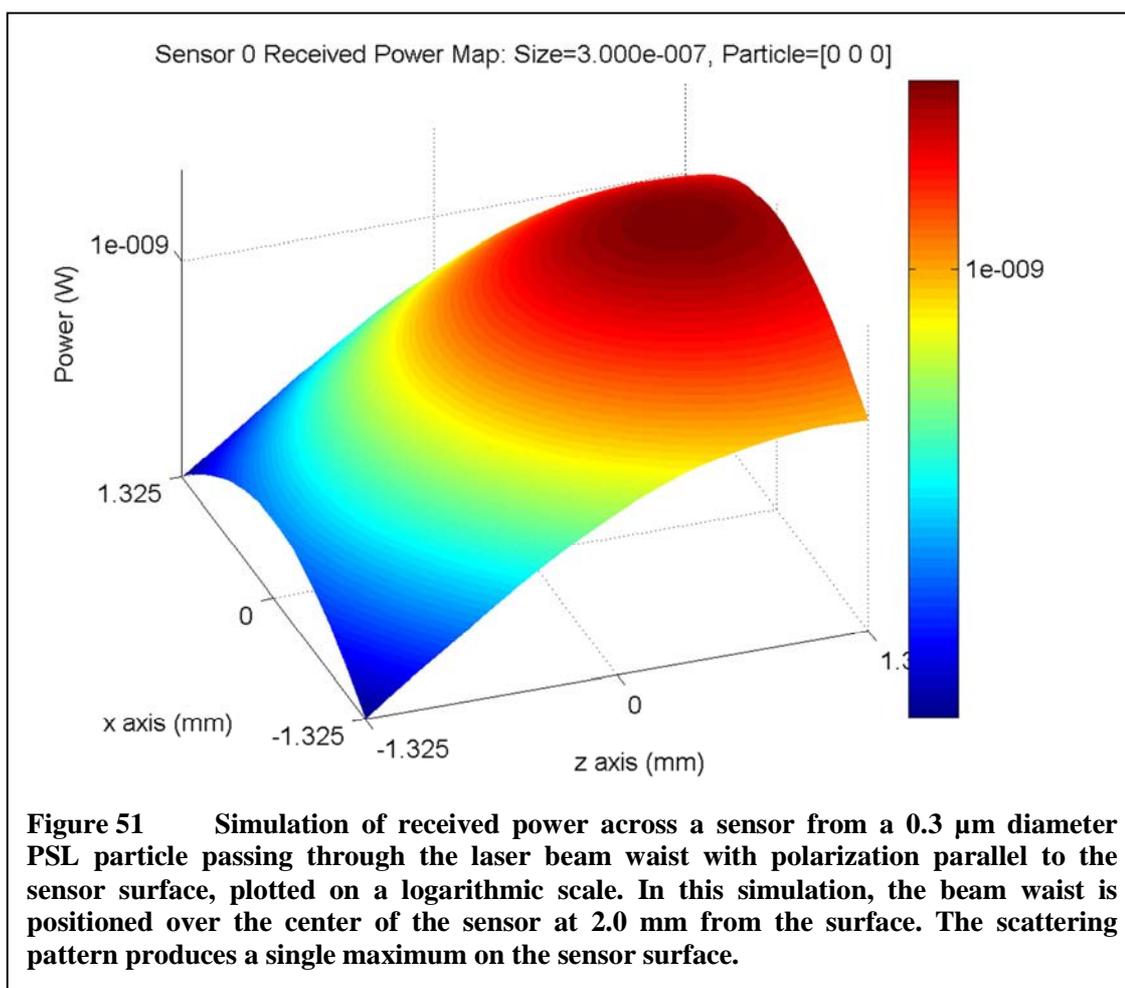
The highest intensity is in the forward direction where the normalized plot is 0dB, aligned with the propagation of incident light. Light that is polarized perpendicular to the scattering plane tapers off gradually to be approximately -13dB in the backscattering direction 180 degrees from the propagation direction of the incident light. Light that is polarized parallel to the scattering plane differs significantly in that there is a -30dB



minimum at approximately 110 degrees from the forward direction (250 degrees on the plot).

Figure 51 shows the surface plot of power received at the sensor for the same 0.3 μm PSL particle in the center of the laser beam waist, plotted on a logarithmic scale. As described previously, the sensor is positioned with its center directly 2.0 mm under the beam waist. In this case, the laser polarization has an orientation parallel to the sensor surface.

Since the scattering does not change the polarization of the light relative to the



physical arrangement, the light received at the sensor also has a polarization that is parallel to its surface. Furthermore, the scattering plane for light scattered directly from the beam waist to the line parallel to the z-axis of the sensor where x equals 0 will be perpendicular to this polarization. Moving away from this line in the positive or negative x direction will increase the amount of light reaching the sensor that is polarized parallel to the scattering plane. This means that the profile of the power received at the sensor along the $x = 0$ line should match the polar plot of Figure 50 for 67 degree section of the Mie scattering plot centered at the 90 degree point, with the negative end of the z-axis of the surface plot corresponding to 123.5 degrees on the polar plot and the positive end of the z-axis corresponding to 56.5 degrees. Moving away from the $x = 0$ line will begin to mix in a contribution from light polarized parallel to the scattering plane in the 360-180 degree range of the polar plot.

As might be expected from the polar plot, the power surface plot has a single maximum with a magnitude less than 10nW. The peak power is also centered on the z-axis of the simulation coordinate system previously shown in Figure 35. The peak is shifted forward in the positive-z direction, indicating that intensity increases in the forward scattering direction. The peak is not at the edge of the sensor, even though scattering closer to the forward direction should be higher than any other direction intercepted by the sensor surface. As the edge is approached, the increased scattering intensity is outweighed by the decrease in intensity due to both the larger distance from the particle to the sensor surface and the shallower angle of incidence between the scattered light and the sensor normal vector.

If the light source is rotated 90 degrees about the propagation axis, the incident light will have a polarization that is perpendicular to the surface of the sensor. This results in the scattered light having an orientation mostly parallel to the scattering plane. The scattered light intercepted at the $x = 0$ line of the sensor should most closely match the polar plot of Figure 50 for the 67 degree section of intensity centered at 270 degrees, or roughly 303.5 to 236.5 degrees. Note that for both of the polarization scenarios the forward scattering direction is toward the positive end of the z -axis. The surface plot with perpendicular polarization relative to the scattering plane corresponds to the polar plot from roughly 60 degrees at the positive end of the z -axis to 120 degrees at the negative end, the plot with parallel polarization relative to the scattering plane corresponds to 300 degrees at the positive end of the z -axis to 240 degrees at the negative end.

The surface plot of the power received at the sensor for incident light with polarization perpendicular to the sensor surface is shown in Figure 52. As with the case of light polarization parallel to the sensor, this plot shows a maximum toward the forward scattering direction. This plot has the additional feature of a deep minimum in the backscattering direction. This matches the expected pattern as seen on the polar plot of Figure 50 for the 67 degree section centered at 270 degrees, with the strongest scattering in the forward direction and the -30dB minimum at approximately 110 degrees from the direction of incident light propagation. With a separation between the particle and sensor

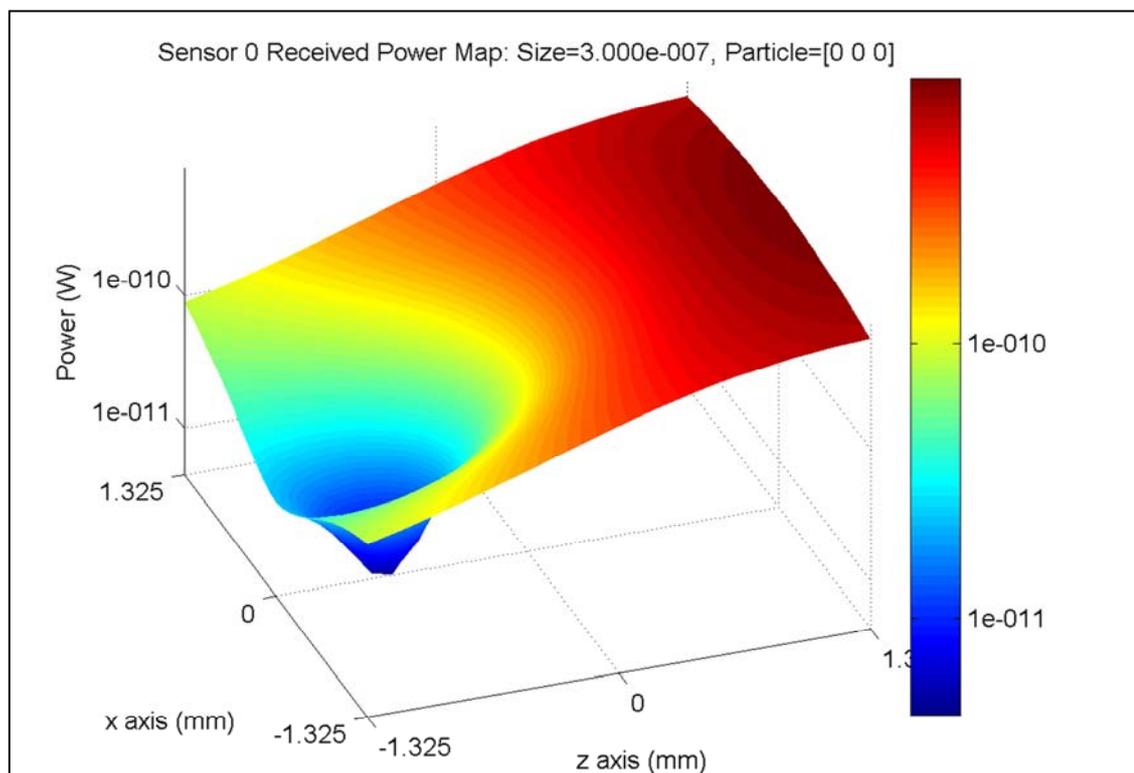
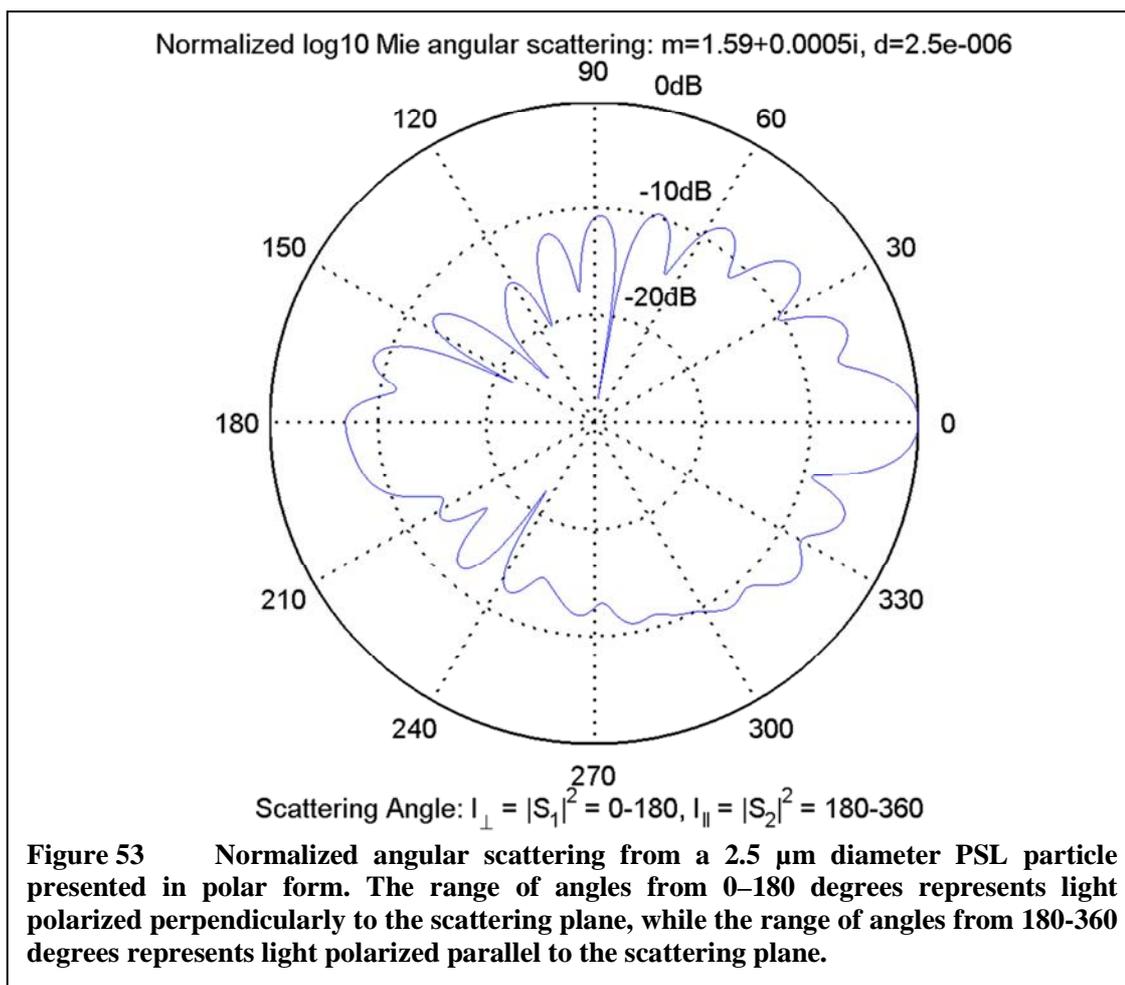


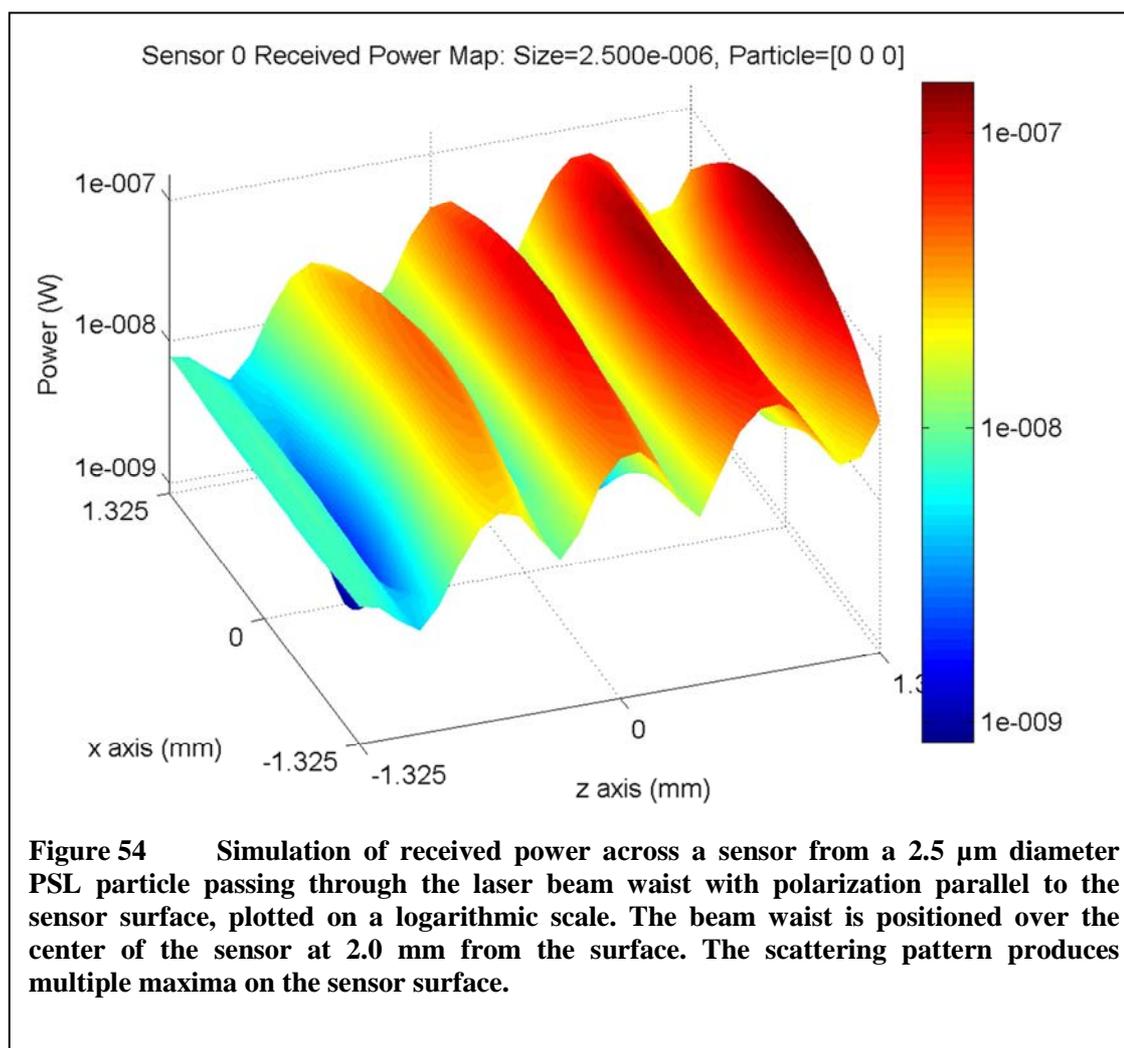
Figure 52 Simulation of received power across a sensor from a $0.3 \mu\text{m}$ diameter particle passing through the laser beam waist with polarization perpendicular to the sensor surface, plotted on a logarithmic scale. In this simulation the beam waist is positioned over the center of the sensor at 2.0 mm from the surface. The scattering pattern shows a maximum toward the forward scattering direction and a single deep minimum.

surface of 2.0 mm, the minimum would be expected to be located approximately halfway between the center of the sensor and the edge, which matches the pattern seen on the surface plot.

Larger particles tend to exhibit more complex scattering patterns as the size of the particle exceeds the wavelength of scattered light. A larger particle is chosen for a second example of the simulation of the power received at the sensor surface. In this case, a 2.5 μm diameter PSL particle is used. The Mie scattering pattern for this particle is shown in Figure 53. Note the much more complex scattering pattern in comparison with the 0.3 μm particle simulated previously. The scattering is strongest in the forward direction, with multiple lobes for both polarizations of scattered light.

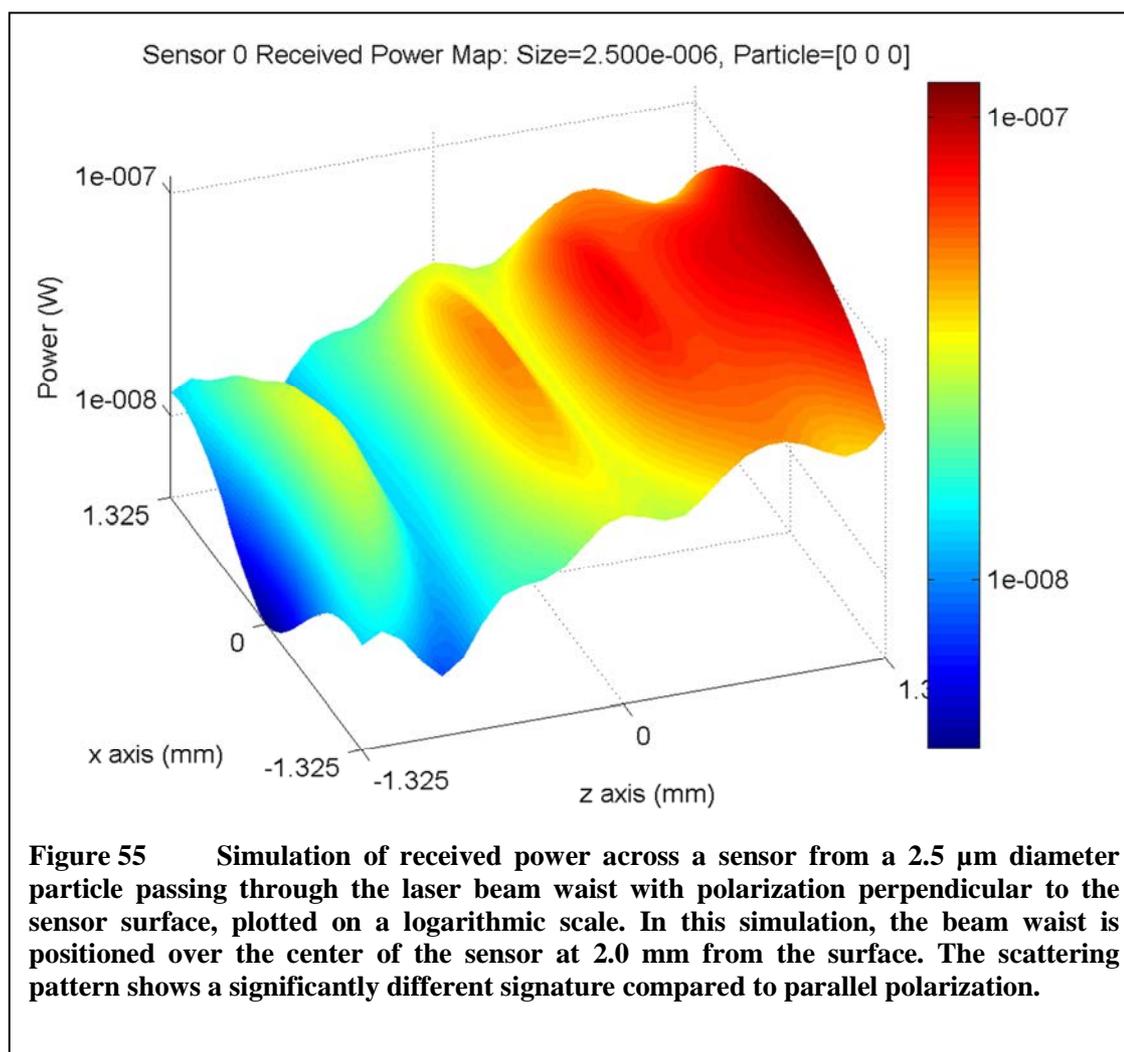
The corresponding surface plot for the power received at the sensor is shown in Figure 54. As before, the profile of the power received at the sensor along the $x = 0$ line should match a 67 degree section of the Mie scattering plot centered at the 90 degree point of Figure 53. Looking at this Mie scattering polar plot, this would cover three main lobes with the edges rising toward 4th and 5th lobes in the forward and backward scattering directions between approximately 60 and 120 degrees on the polar plot. This matches the pattern seen in the surface plot of Figure 54.





Note also that the ridges on the surface plot appear to bend as they move from the $x = 0$ centerline out to the edge of the sensor. This is an expected consequence of the interception of a spherically radiating maximum with a flat plane, since the intersection of the two will be a curve.

Simulating the same 2.5 μm PSL particle with incident light polarized perpendicular to the sensor surface shows a significantly different signature. Figure 55 shows this plot with all other parameters the same. Note that the surface plot is much smoother in this case, as would be expected from the relatively smooth pattern seen



between 240 and 300 degrees in the Mie scattering plot of Figure 53. The relatively deep minimum at just less than 240 degrees is visible as the lowest point on the left side of the surface plot at the $x = 0$ point.

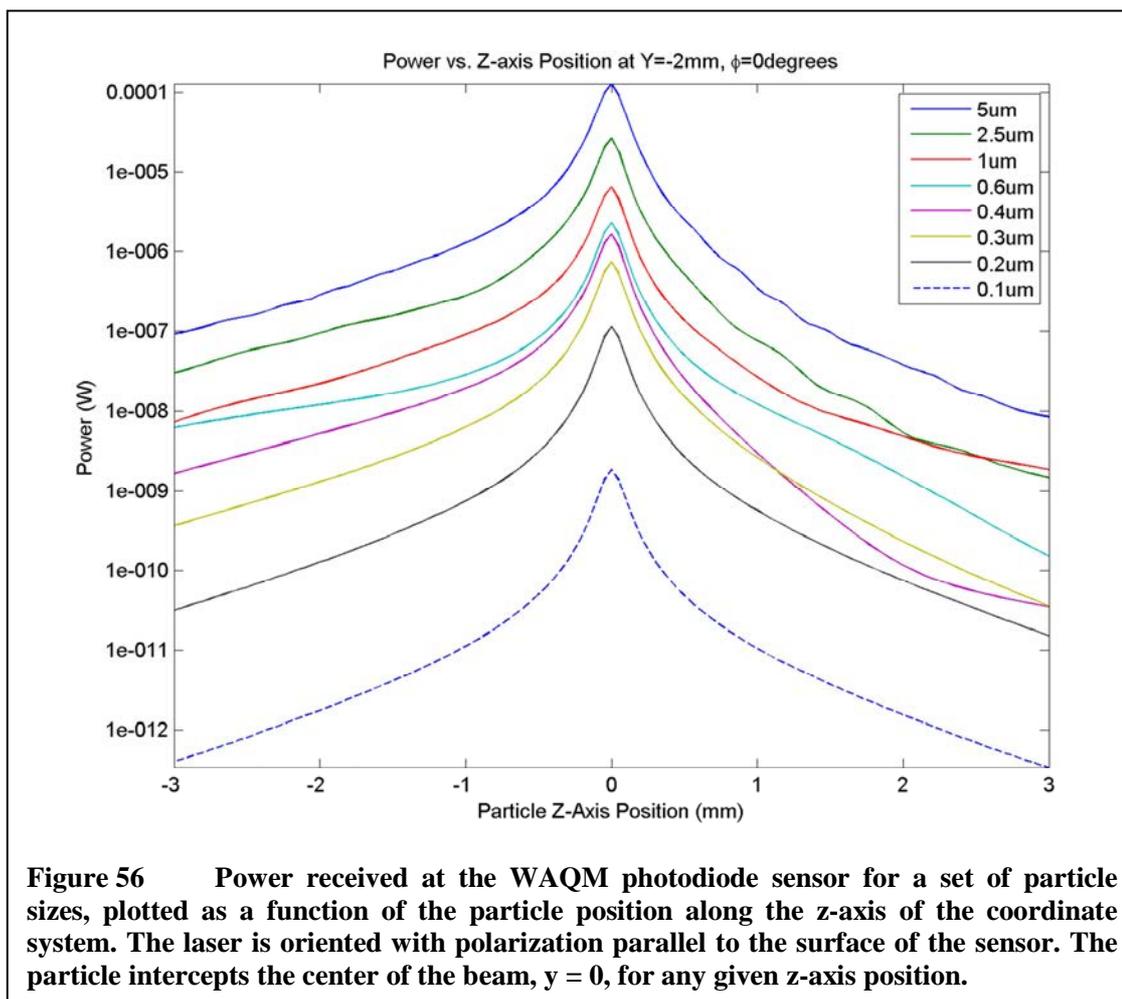
Single Sensor Irradiance

Using the simulation framework, an important question that can be explored relates to the irradiance one might expect to see from a given particle passing through an arbitrary point in the laser beam. Since the flow channel is wide relative to the geometry of the laser beam, a particle may intercept the beam at a very high intensity point at or

near the waist, or at a point with a much lower intensity. The question is then raised as to how particle sizing might be accomplished based solely on the intensity of the scattered light.

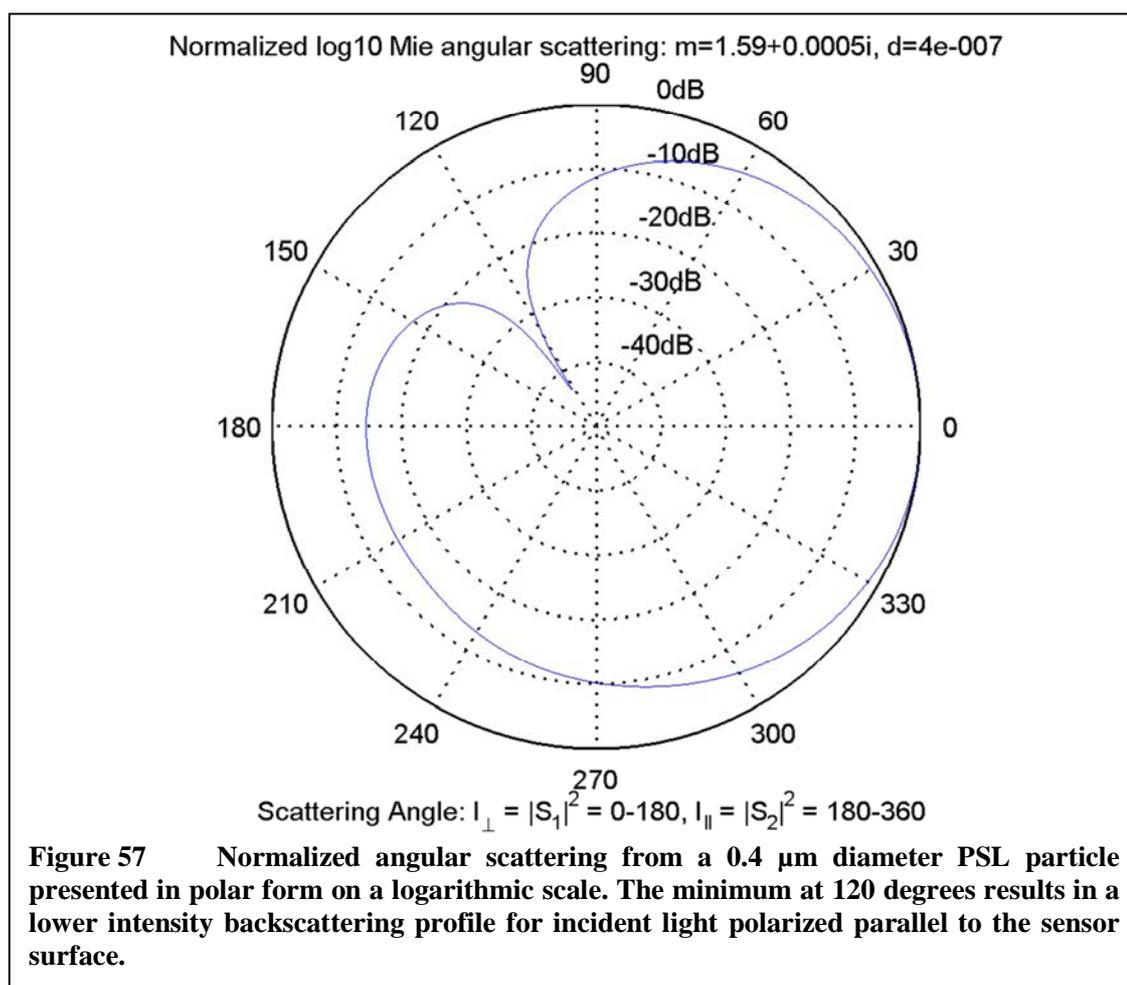
The set of curves in Figure 56 quantifies this for a range of PSL particles with diameters from 0.1 μm to 5.0 μm when the laser is oriented with polarization parallel to the sensor surface. Each curve peaks at the beam waist, where $z = 0$, which is also the center of the 6 mm wide flow channel. The left side of the plot, with $z < 0$, is the region of the beam between the beam waist and the wall of the flow channel closest to the laser source. Likewise, the right side of the plot with $z > 0$ is the region between the beam waist and the flow channel wall furthest from the laser source. The peak at the beam waist occurs due to the high intensity at the waist relative to the rest of the beam. The power distribution of the beam is then symmetrical between the $z < 0$ and $z > 0$ sides of the flow channel, with the intensity falling off the further the particle gets from the beam waist.

One of the first characteristics of Figure 56 to notice is that the power for any given particle may vary more than three orders of magnitude from the edge of the flow channel to the beam waist. The cause of this variation is dominated by the change in intensity of the laser across this space. This is not a surprise given that the laser irradiance can vary three orders of magnitude along the beam axis across the flow channel.



It is also interesting to note that on average the power received at the sensor is higher for a particle crossing between the waist and source compared with that for a particle crossing between the waist and the far side of the channel. This variation is due to the dominance of forward scattering over that occurring in other directions. This is particularly true for particles large enough to be outside the Rayleigh regime, which is approximately $0.3 \mu\text{m}$ and larger for the 650 nm laser of the WAQM OPC.

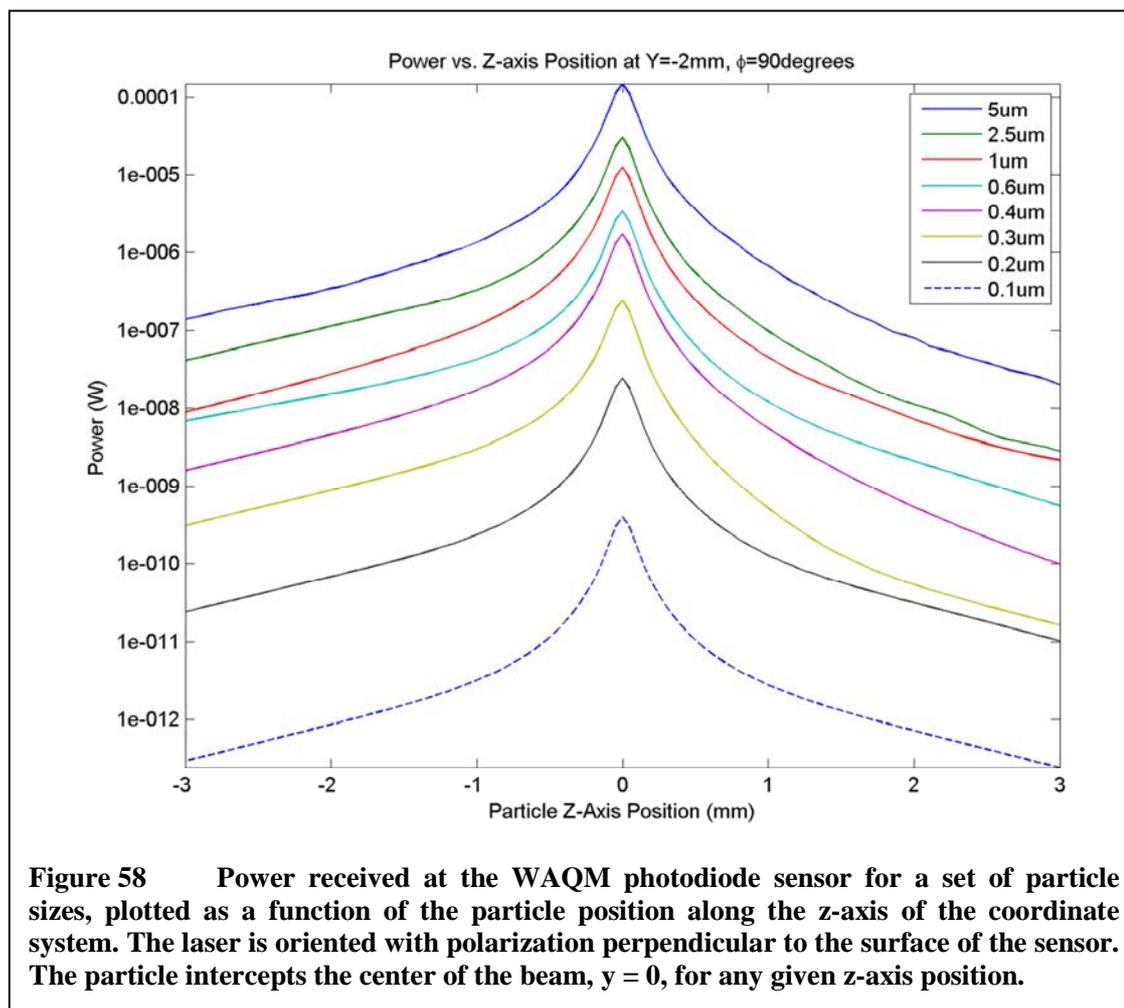
The received power curve shapes vary somewhat depending on the particle size, which is due to the shapes of the scattering patterns for any given particle diameter. Some particle sizes will show a lower intensity backscattering profile than others, resulting in a curve that is lower in magnitude than typical on the right half of the plot. This can result in an even greater variation in overall magnitude. For example, the 0.4 μm particle in the plot shows more than four orders of magnitude variation between the peak and the lowest point on the right side of the plot. The Mie scattering plot for this particle is shown in Figure 57. Note the deep minimum at an angle of 120 degrees in the polar plot. This is the cause of the lower magnitude backscattering profile for this particular particle size.



From Figure 56, it can be seen that each individual particle size can generate a very wide range of possible intensities received at the sensor. It should be noted that the curves only depict the maximum intensity at $y = 0$, the center of the laser beam, for any point across the z -axis of the channel. The plotted lines assume that the particle intercepts the laser at its maximum in the y -direction, but there is no guarantee of this actually occurring. In fact, particles passing through the flow channel may intercept this maximum, but are much more likely to strike an arbitrary point in the y -direction, which would result in a lower intensity scattering.

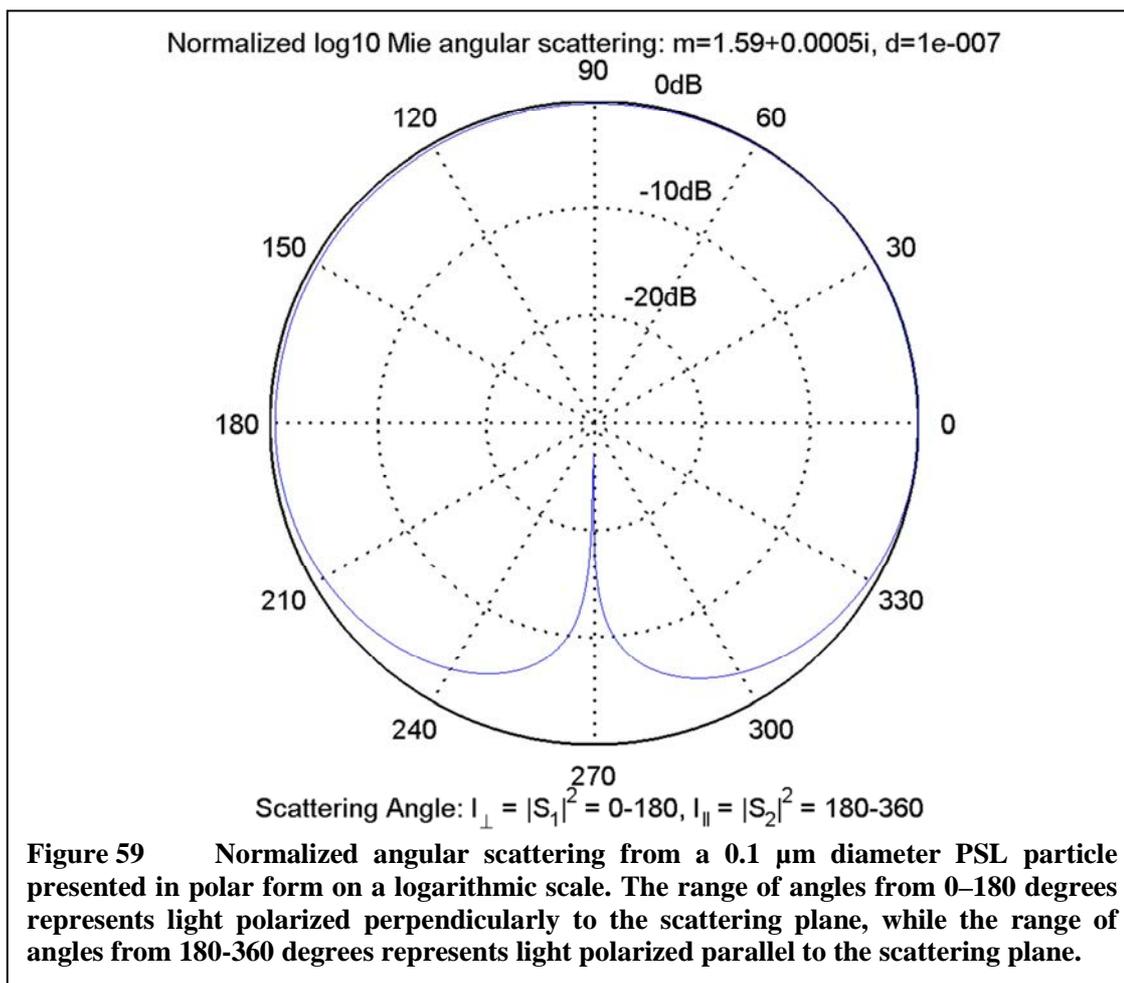
Given this, it is obvious that the ranges of intensities for particles of fairly dissimilar sizes can overlap depending on where they might cross the laser inside the flow channel. For example, a $0.2 \mu\text{m}$ particle passing through the beam waist will result in more than 10 nW being received at the sensor. At the same time, a $5.0 \mu\text{m}$ particle passing through a less intense section of the beam near the edge of the flow channel farthest from the laser source might cause an equivalent power to be received at the sensor. Clearly, it will be difficult to distinguish particle sizes based on received intensity alone.

Rotating the laser to a position where the incident light is polarized perpendicular to the sensor surface yields the set of curves in Figure 58. The curves look very similar to those seen in Figure 56, but with some noticeable differences. The most important difference is in the peak power captured from the smallest particles. The particles that have diameters smaller than $0.4 \mu\text{m}$ show a significant decrease in peak intensity when the incident light is polarized perpendicular to the sensor. This decrease is due to the nature of scattering as the ratio of the particle radius to the wavelength of incident light



becomes small. The resulting scattering pattern begins to exhibit a strong minimum for light polarized parallel to the scattering plane, which for this physical arrangement corresponds to light polarized perpendicular to the sensor surface.

This minimum can be seen developing in Figure 50, which shows a $0.3 \mu\text{m}$ particle. With small enough particles, the minimum becomes oriented at 90 degrees from the propagation direction of the incident light, as can be seen for a $0.1 \mu\text{m}$ particle in Figure 59. For scattered light polarized parallel to the scattering plane, this corresponds to the 270 degree point on the polar plot. With the sensor centered directly below the laser



beam waist, the minimum is then aimed directly at the center of the sensor when a particle passes through the waist.

This important result then leads one to conclude that for detection of the smallest particles, the laser polarization should be parallel to the sensor surface, and perpendicular to the scattering plane. This will yield the highest intensity scattering towards the sensor when small particles pass through the beam waist.

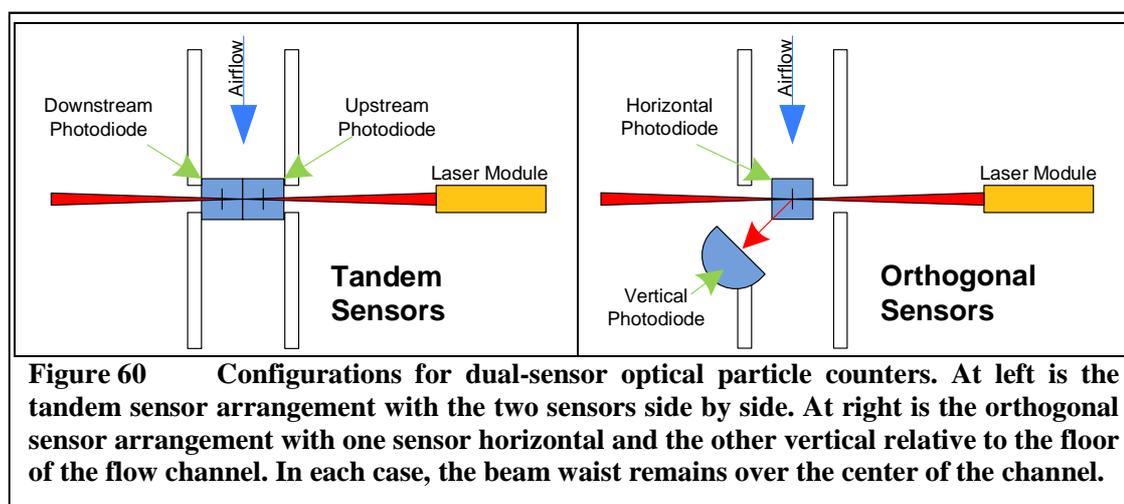
Dual Sensor Irradiance

As described above, a single sensor is limited in its ability to size particles due to the large range of scattered light intensities that may result from a particle crossing the

beam at an arbitrary point. This is primarily due to the wide variation in beam intensity across the flow channel coupled with a lack of information about the location of the particle in the beam. Position dependent scattering intensity also plays a role in the intensity variation, with forward scattering being more intense in general than backscattering.

Adding a second photodiode to the system may provide some significant advantages in particle sizing without adding a great deal of cost to the system. To start, having two sensors may provide some indication of position within the beam based on the ratio of intensities received at each sensor. This ratio may be used in normalizing the scattering intensity, helping to remove some of the scattering intensity's dependence on the particle's position in the beam.

Figure 60 shows two different dual-sensor arrangements that are feasible to create within the framework of the WAQM OPC. The diagram on the left shows a “tandem” arrangement, with two horizontal surface-mount sensors placed side-by-side under the beam waist. The beam waist in this case is focused over the point immediately between



the two sensors. The sensor closest to the laser source is referred to as the “upstream” sensor, and the sensor furthest from the laser source is referred to as the “downstream” sensor. The intent of this arrangement is to explore the differences in positioning across the flow channel with sensors that have the same basic orientation relative to the source laser polarization. Some sensitivity may be lost since the beam waist is no longer directly over a sensor center, though it is now relatively near to two sensors.

The diagram on the right half of Figure 60 shows an “orthogonal” arrangement, where one horizontal sensor is positioned directly below the beam waist and one vertical sensor is placed at an angle to the side of the beam waist. This arrangement has the benefit of keeping the horizontal sensor as close as possible to the beam waist as in the single sensor arrangement. The vertical sensor provides a second position that is both off-angle and aligned to capture scattering that is polarized orthogonally to that captured by the horizontal sensor.

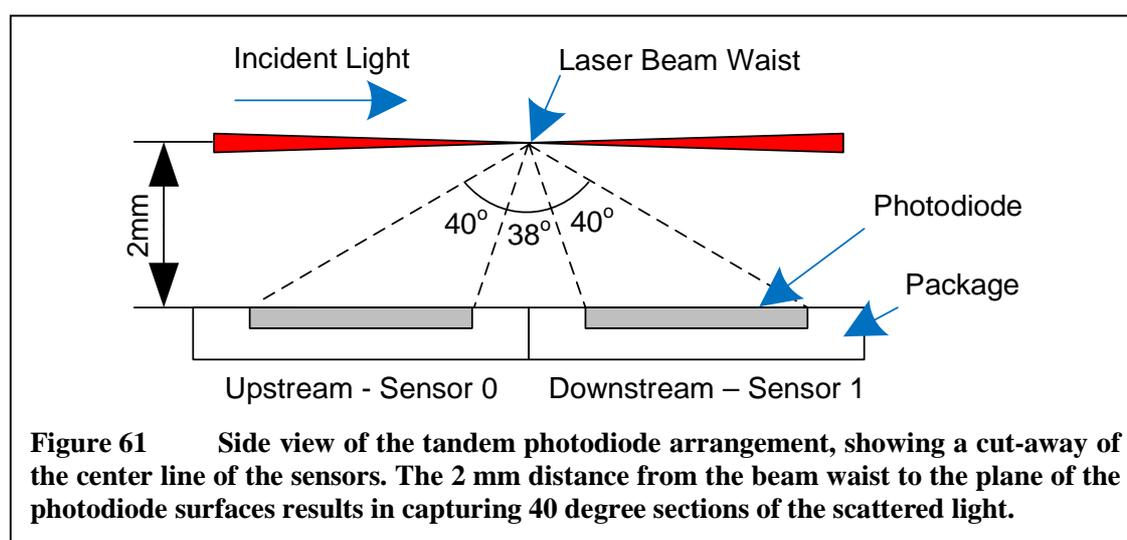
Other dual-sensor arrangements were explored in the process of determining which would be best suited for implementation in the WAQM OPC. These were for the most part permutations of the tandem and orthogonal designs described above, with the sensors positions shifted relative to center of the beam waist. These other permutations were for the most part no better at providing discriminating data for particle detection and sizing.

A significantly different physical arrangement was also contemplated, with two vertical through-hole sensors arranged in such a way as to have the air flow pass between them with the laser beam waist centered equidistant from the sensor surfaces. This particular arrangement would allow the sensors to detect scattered light that is polarized

both parallel and perpendicular to the sensor surfaces. Unfortunately, to accommodate the air flow channel between the sensors, the separation between the sensors and the beam waist would have been significantly larger than is the case for the tandem or orthogonal arrangements. This greater separation would have negatively impacted the OPC's ability to detect very small particles, so the design was not pursued further.

Tandem Sensors

Simulation of the dual-sensor arrangements is very similar to that for the single sensor, but with data for both sensor positions returned simultaneously. A particle passing through the beam waist results in a significantly different geometry for the captured scattered light than that of the single sensor. Figure 61 shows a cut-away view of the arrangement, resulting in each sensor capturing approximately 40 degrees of the scattered light when considering this two-dimensional surface defined by the center of the two sensors and the beam waist. It is important to note that the upstream sensor is receiving backscattering from a particle passing through the beam waist, and the downstream sensor is receiving forward scattering. For incident light polarization perpendicular to the

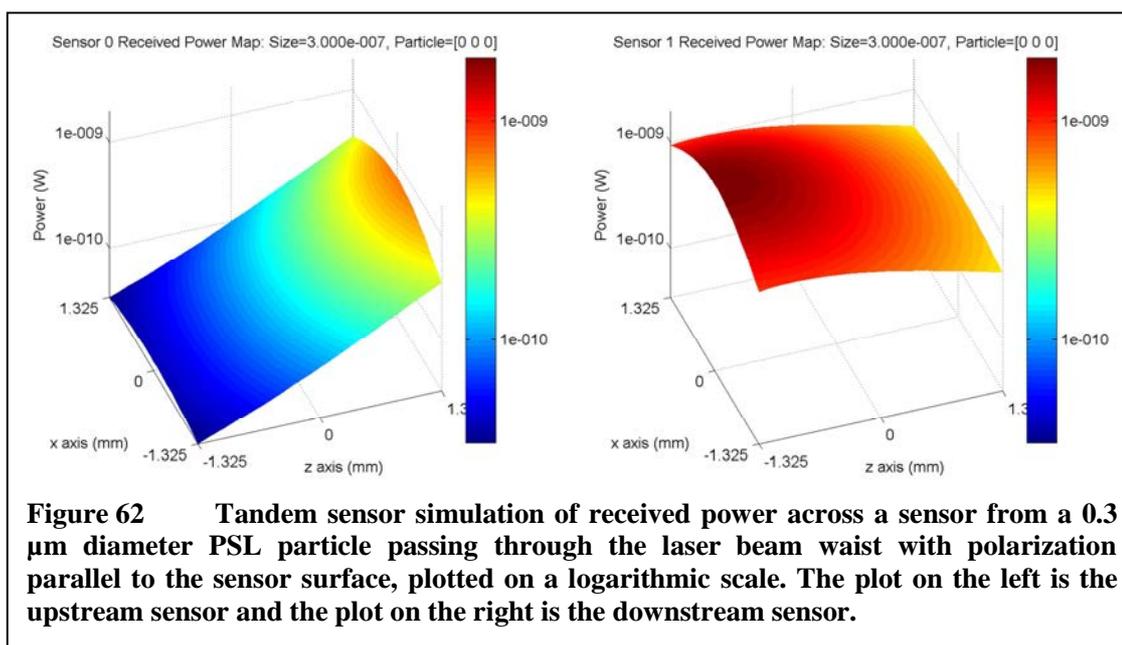


scattering plane, the upstream sensor will see scattering corresponding to 109-149 degrees on the Mie scattering polar plot and the downstream sensor will see scattering corresponding to 31-71 degrees. For incident light polarization parallel to the scattering plane, the upstream sensor will see scattering corresponding to 211-251 degrees on the Mie scattering polar plot and the downstream sensor will see scattering corresponding to 289-329 degrees. As with the single sensor arrangement, these two-dimensional angles are important only in comparing the power received on the centerline of the photodiodes to the polar plots of scattering intensity. The power received is calculated from the solid angle that is intercepted by each sensor, which takes the three-dimensional situation into account.

With the tandem arrangement there is a “dead space” directly under the beam waist where the gap between photodiodes is created by the device packages forcing some separation. Unfortunately for the tandem arrangement, this is the highest sensitivity point for small particles, which can impact the sensitivity of the device. The impact of this arrangement will be examined in the following section.

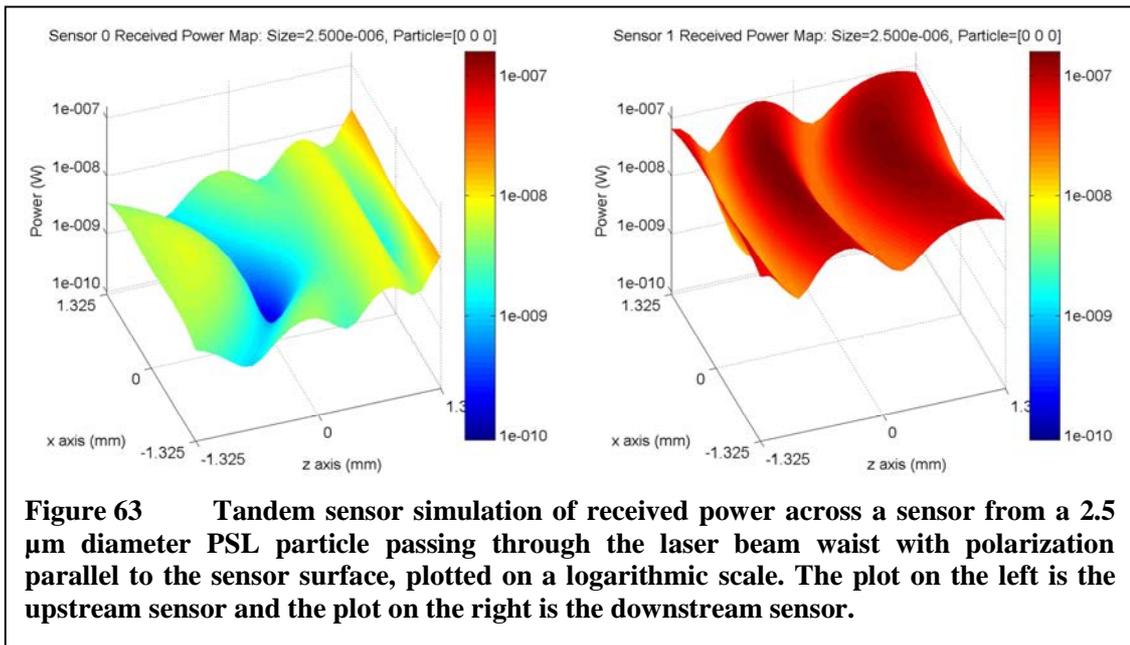
Figure 62 shows the results from the tandem arrangement when a 0.3 μm PSL particle passes through the beam waist with incident light polarized parallel to the sensor surface. The surface plot on the left is for the upstream sensor nearest the laser source, and the surface plot on the right is the downstream sensor. The two surface plots use the same color scale, which is set to cover the aggregate range between the two plots.

As described above, these plots will correspond to 109-149 degrees on the Mie scattering polar plot for the upstream sensor, 31-71 degrees for the downstream sensor. Comparing to the Mie scattering plot of Figure 50, it is clear that the downstream sensor



should be receiving more scattered light irradiance than the upstream, with a difference of about 10dB between the two sensors. Other factors such as the variation in distance from the scattering point and the variation in angle of incidence between the light and sensor surface cause the curving shapes of each individual surface plot. These factors also cause the downstream sensor irradiance to fall off with increasing z-axis position, which is contrary to the steadily increasing Mie scattering irradiance from the particle when these factors are not taken into account.

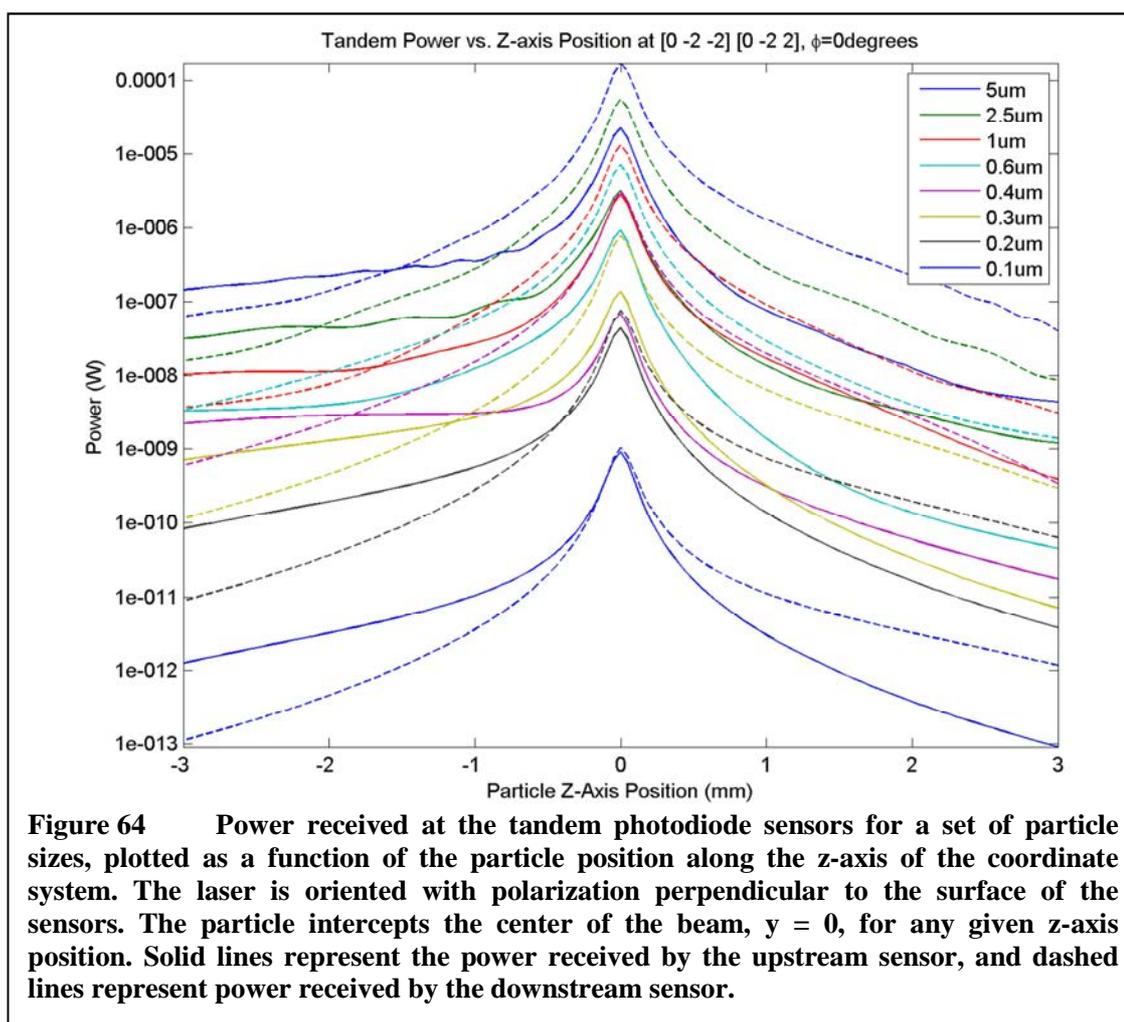
A more complex scattering pattern is seen in Figure 63 with a 2.5 μm particle passing through the beam waist with the upstream sensor shown in the surface plot on the left, and the downstream sensor on the right. The surface plots correspond to the Mie scattering polar plot shown in Figure 53. The upstream sensor is clearly capturing the three lobes between 109-149 degrees on the Mie scattering polar plot, with the deepest null showing just to the left of center of the sensor. The downstream sensor likewise



captures the three lobes from 31-71 degrees. The downstream sensor dominates in received irradiance, as expected from the polar plot.

As with the single sensor, the maxima and minima show a definite pattern of curving from the center out towards the edges of the sensor. This is an expected effect of the interception of a spherically radiating maximum and a plane, as is the case with a Mie scattering maximum radiating out from a point source being intercepted by the planar sensor surface.

As with the single sensor arrangement, one would like to know the power received at the tandem sensors by particles intercepting the laser at arbitrary points across the flow channel. The set of curves in Figure 64 quantifies this for a range of PSL particles with diameters from 0.1 μm to 5.0 μm when the laser is oriented with polarization parallel to the sensor surfaces. Data from both of the sensors in the tandem arrangement are represented, with the upstream sensor shown with solid lines and the downstream sensor shown with dashed lines.



As with the single sensor, each curve for both of the tandem sensors peaks at the beam waist where $z = 0$. The high peak in incident light intensity at the beam waist outweighs the differences in position of the two sensors and any impact on scattering pattern when the particle passes at or near the waist. The curves show that the downstream sensor in general shows a higher peak for the beam waist than the upstream sensor, which tends to match the stronger forward scattering patterns. This effect is lessened for very small particles that begin to show a more uniform scattering pattern for incident light polarized parallel to the sensor surfaces.

Looking at the response for each particle at the $z = 3$ mm point, or furthest from the laser in the flow channel, all sizes show a stronger response from the downstream sensor than from the upstream. This occurs since the particle will be significantly closer to the downstream sensor than to the upstream, and since backscattering intensity is in general much weaker than forward scattering.

The response for each particle size at $z = -3$ mm, or closest to the laser, shows more inconsistency. For the most part, particles at this point will show a higher intensity at the upstream sensor than the downstream. The particle's proximity to the upstream sensor alone drives much of this, since both sensors would be dominated by detection of forward scattering light. However, there are exceptions to this behavior in the range of the $0.6 \mu\text{m}$ particle. This particular particle size exhibits a scattering signature that greatly favors the downstream sensor, causing the intensity received there to be stronger than the nearby upstream sensor.

The power data from the individual sensors are difficult to reconcile into a useful pattern. Combining data from the two sensors simultaneously can lead to something more interesting. In Figure 65, the ratio of the upstream sensor power to the downstream sensor power is plotted across the width of the channel on a logarithmic scale. There is a fairly clear trend of a shrinking ratio moving from the point closest to the laser source at $z = -3$ mm to the opposite side of the channel. Furthermore, a ratio might be chosen that would indicate whether or not the particle was toward the left side or right side of the channel. This line would leave a fair amount of uncertainty in the middle of the channel, but at least some position differentiation could be made.

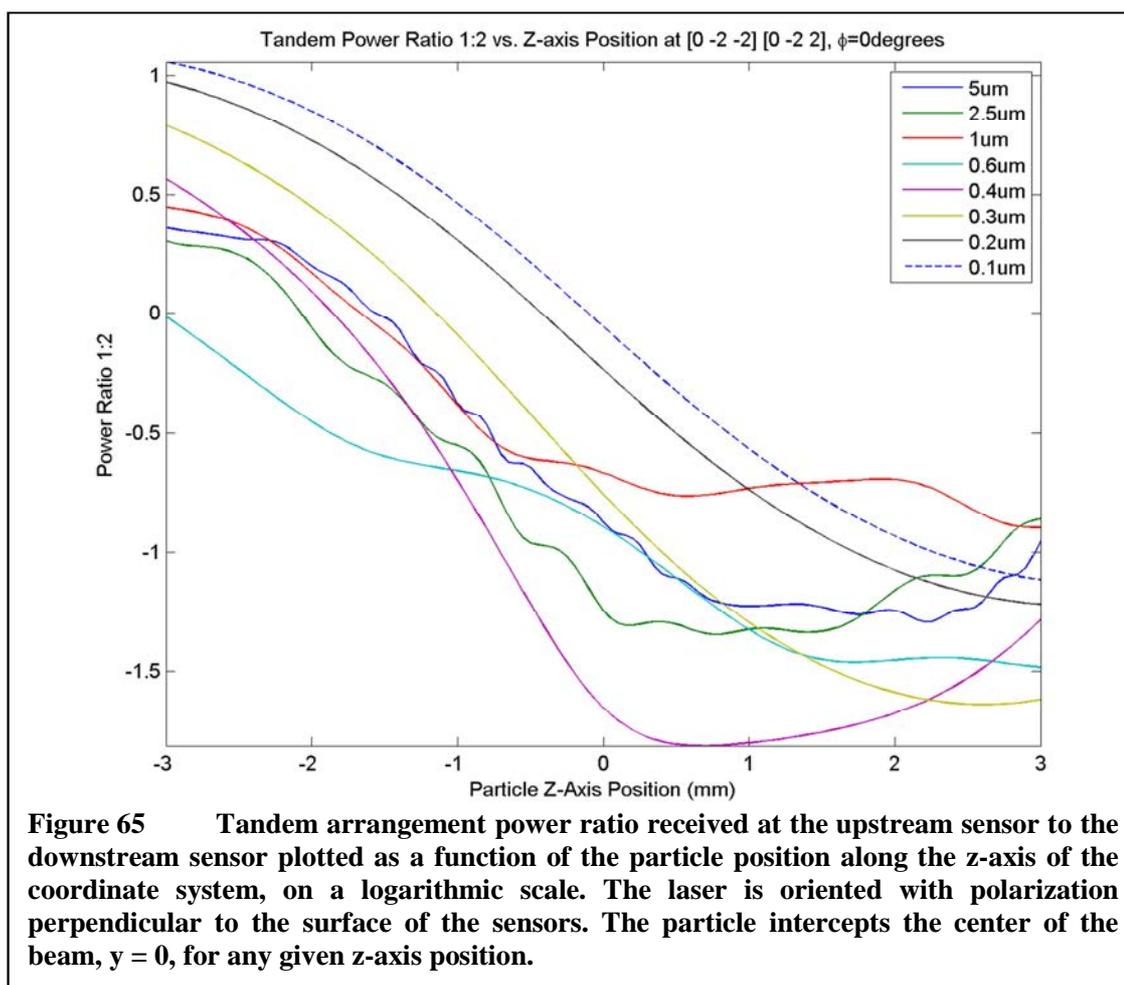


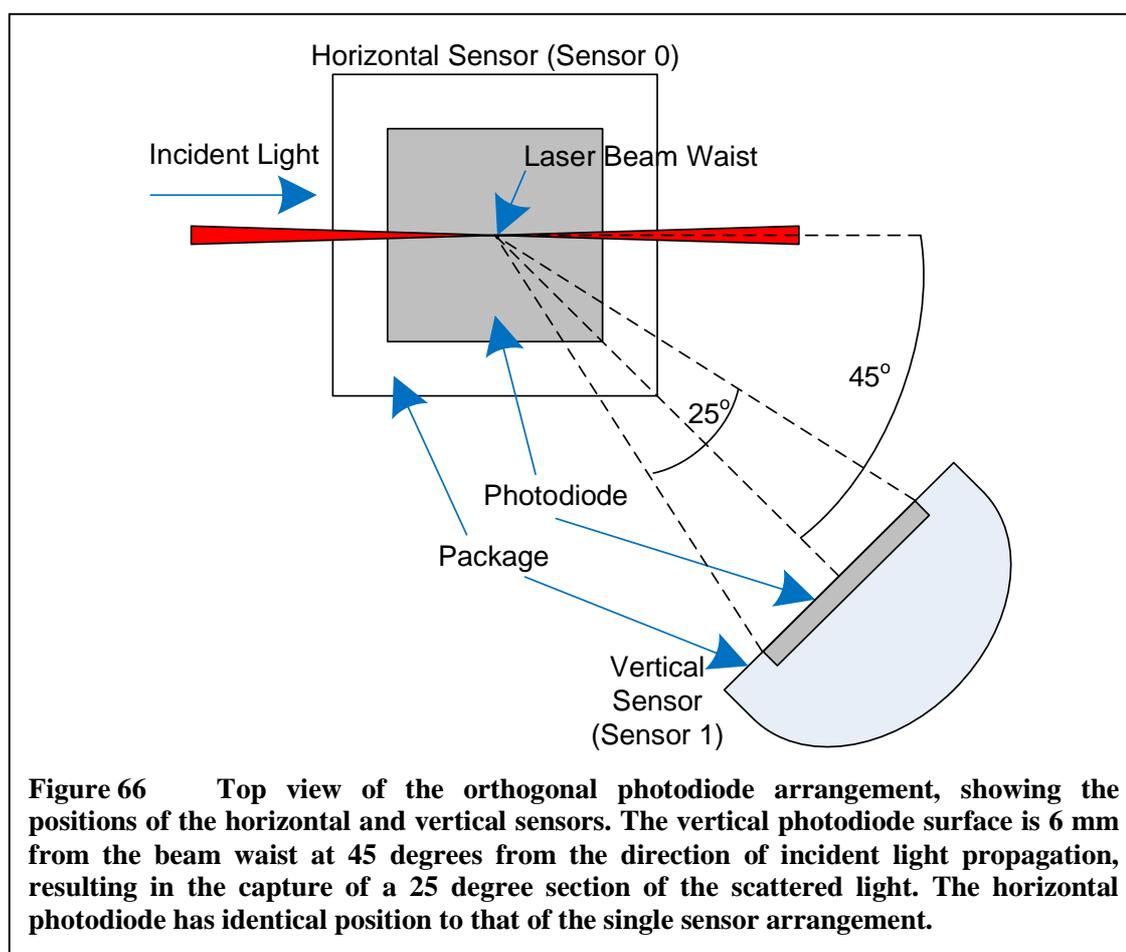
Figure 65 Tandem arrangement power ratio received at the upstream sensor to the downstream sensor plotted as a function of the particle position along the z-axis of the coordinate system, on a logarithmic scale. The laser is oriented with polarization perpendicular to the surface of the sensors. The particle intercepts the center of the beam, $y = 0$, for any given z-axis position.

One significant downside of the tandem sensor arrangement can be seen by comparing the peak values for the smallest particles detected in the tandem arrangement to those detected by the single sensor arrangement. The peak power values for particles at $0.3 \mu\text{m}$ and below are reduced for the tandem sensor, due to the increased distance between the beam waist and the sensors and the shallower angle of incidence between the scattered light and the sensor surface. This difference will impact the range of particles that the OPC is able to detect, increasing the minimum detectable size. This might be overcome by moving the sensors relative to the beam waist, with the upstream sensor directly under the beam waist. This pushes the downstream sensor further away from the

beam waist, which ends up putting the sensor mostly in the light trap area of the OPC when dealing with a channel width of 6mm. This may increase the difficulty of shielding the downstream sensor from the ambient laser light coming from the light trap.

Orthogonal Sensors

The orthogonal sensor arrangement is significantly different from both the single sensor and tandem sensor arrangements in that it includes a photodiode oriented vertically, “orthogonal” to the surface of the horizontal sensor under the beam waist. Figure 66 shows a top view of the orthogonal sensor arrangement, with the horizontal sensor in the spot identical to the single sensor arrangement. The vertical sensor is set up



at a 45 degree angle to the direction of incident light in the system, with a distance of 6 mm from the beam waist. This specific geometry was chosen for the WAQM OPC to position the vertical sensor as close as possible to the beam waist while allowing the airflow in the channel to pass by with minimum disruption. The arrangement also allows the laser to pass by the vertical sensor with enough light blocking to keep the photodiode and device package from being directly illuminated by the laser source.

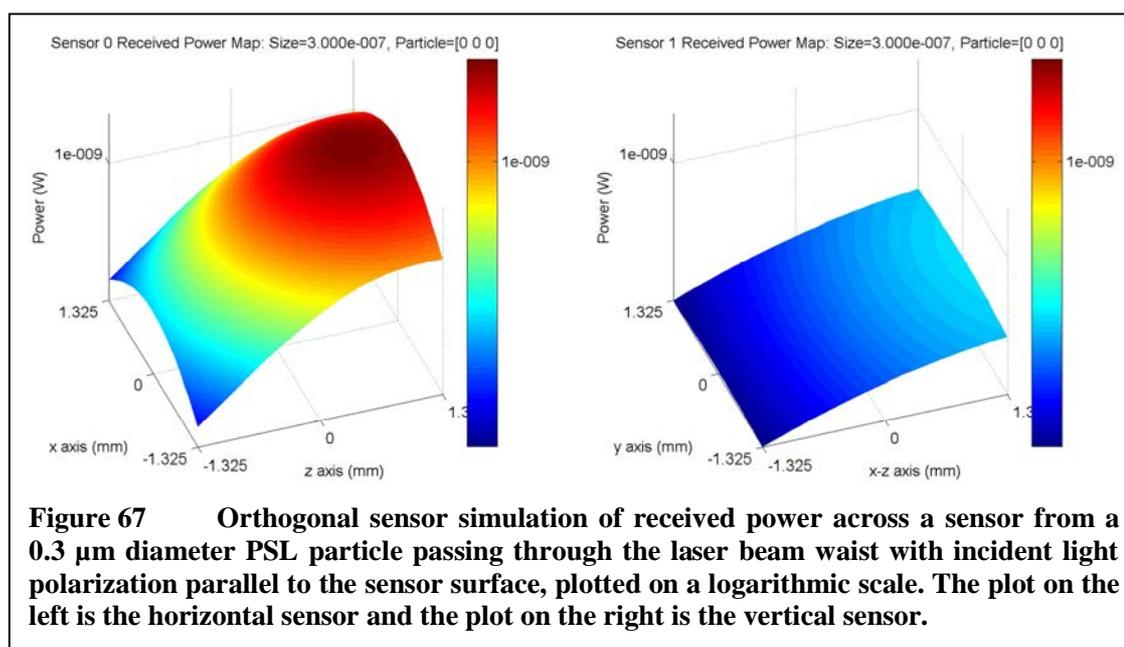
While the horizontal sensor is a surface mount device with the sensor surface parallel to the mounting PCB, the vertical sensor is a through-hole device that orients the sensor surface at 90 degrees to the PCB. The device packages are different but the photodiodes themselves are identical in specification and performance. With this particular orientation, the vertical sensor will mostly capture light that is polarized parallel to the scattering plane, which is also “orthogonal” to most of the light captured by the horizontal sensor.

The intent of this arrangement is to maintain the horizontal sensor in the position of highest sensitivity for small particles that are approaching the Rayleigh regime, yet position the vertical sensor in a way to better capture forward scattering. The arrangement should allow a better differentiation of particle size while maintaining the optimal sensitivity possible with the inexpensive photodiodes. It may be desirable to attempt to position the vertical sensor where the angle between the sensor normal and the incident light is decreased to better capture forward scattering, but this was not possible due to the physical constraints of the WAQM enclosure.

As with the single sensor design, the horizontal sensor will capture scattered light mostly polarized perpendicular to the scattering plane. This geometry corresponds to a 67

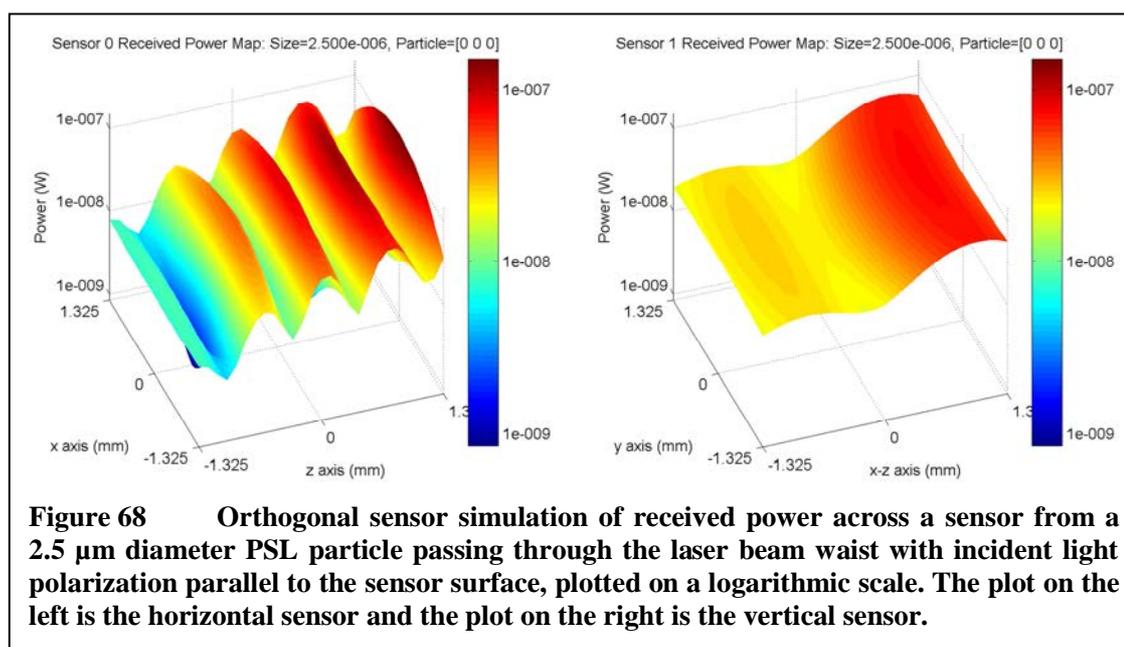
degree section of the Mie scattering polar plot centered at the 90 degree point, with the negative end of the z-axis of the surface plot corresponding to 123.5 degrees on the polar plot and the positive end of the z-axis corresponding to 56.5 degrees. The vertical sensor will capture scattered light that is mostly polarized parallel to the scattering plane over a 25 degree section of the Mie scattering polar plot centered at the 315 degree point. The negative end of the z-axis of the surface plot will correspond to 302.5 degrees on the polar plot and the positive end of the z-axis corresponding to 327.5 degrees.

As with the other sensor arrangements, Figure 67 shows the results from the orthogonal arrangement when a $0.3\ \mu\text{m}$ PSL particle passes through the beam waist with incident light polarized parallel to the sensor surface. The surface plot on the left is for the horizontal sensor, and the surface plot on the right is the vertical sensor. As with previous dual sensor data, the two surface plots use the same color scale, which is set to cover the aggregate range between the two plots.



The difference in received power is immediately noticeable between the two sensors. The horizontal sensor receives a greater amount of irradiation, mostly due to the proximity to the particle as it scatters light. Looking at the Mie scattering diagram of Figure 50 would lead one to believe that the scattering intensity would be greater for the vertical sensor, which is centered at 315 degrees on the polar plot while the horizontal sensor would be centered at 90 degrees. The irradiation is stronger in the direction of the vertical sensor, but the 6 mm distance to the sensor puts it 3 times further away from the particle than the horizontal sensor at 2 mm.

Figure 68 shows the received power from a larger $2.5\ \mu\text{m}$ particle passing through the beam waist. As before, this produces a more complex scattering pattern. The horizontal sensor shown in the surface plot on the left and the vertical sensor is on the right. The surface plots correspond to the Mie scattering polar plot shown in Figure 53. The horizontal sensor is capturing the multiple lobes between 56.5 and 123.5 degrees on

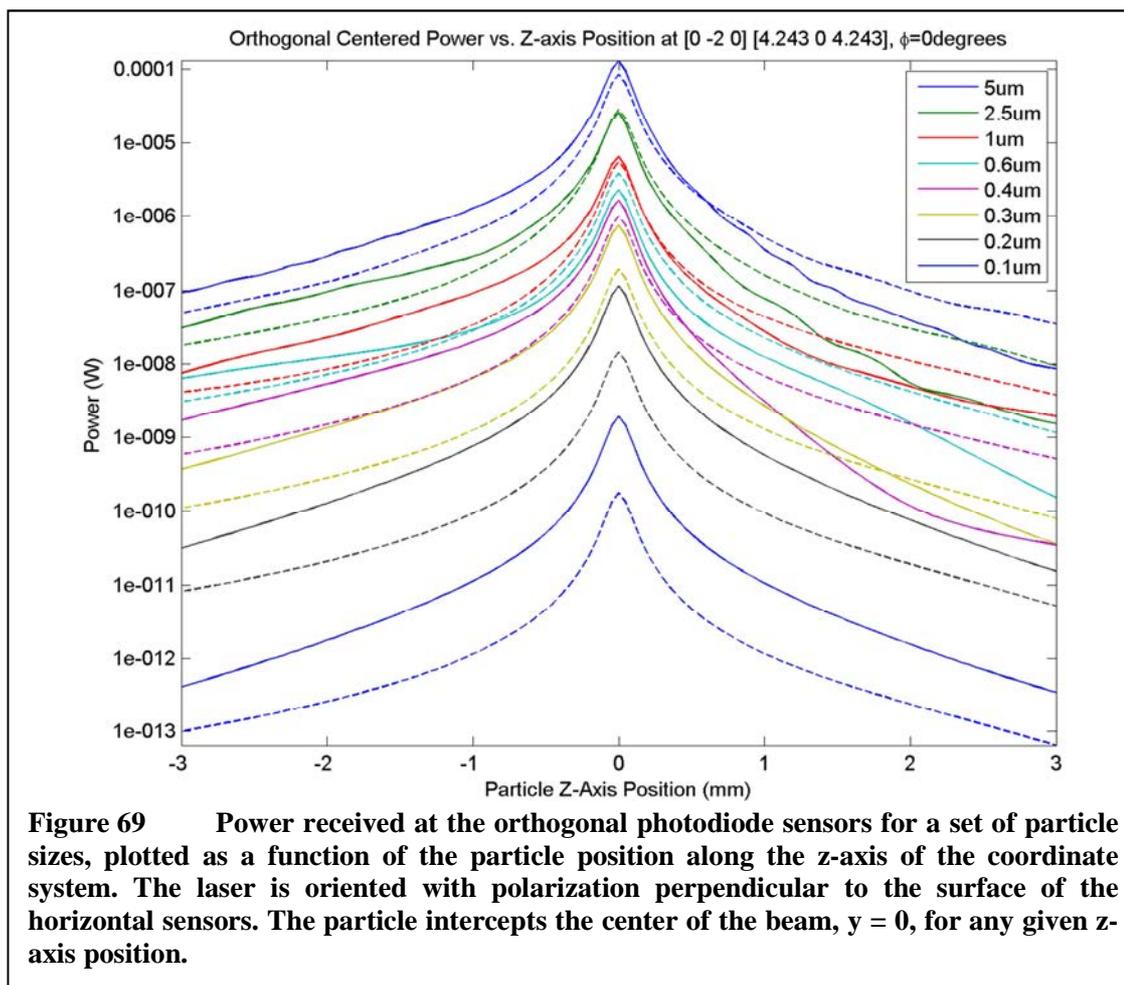


the Mie scattering polar plot. The vertical sensor captures the much more uniform region from 303.5-327.5 degrees.

From the Mie scattering polar plot, the light scattered toward the vertical sensor peaks approximately 8dB higher than the light scattered toward the center of the horizontal sensor. The surface plots show that the power received at the vertical sensor is near to that received by the horizontal sensor. This is due to the difference in distance between the two sensors. Since the vertical sensor is three times more distant than the horizontal sensor, the received power at the vertical sensor is reduced by a factor of nine.

The power magnitudes received at the horizontal and vertical sensors from particles intercepting the laser at arbitrary points across the flow channel are plotted in Figure 69. The data are shown for a range of PSL particles with diameters from 0.1 μm to 5.0 μm when the laser is oriented with polarization parallel to the horizontal sensor surface. Solid lines show data from the horizontal sensor, while dashed lines are data from the vertical sensor.

Similar to the plots for previously described sensor arrangements, these curves show the highest magnitude received power when the particle passes through the highest intensity point of the laser beam waist. This is true for both sensors, and as previously indicated has to do more with the intensity of the laser at the beam waist than the position of the particle relative to the sensors. In particular, even though it is possible for a particle to travel through the beam at a point much closer to the vertical sensor, the scattered light intensity will still be much lower due to the reduced intensity of the laser beam at that point in the flow channel.



The curves for the smallest particles in the 0.1 to 0.2 μm range show a higher intensity of light received at the horizontal sensor than the vertical, across the entire width of the flow channel. Larger particles consistently show higher intensity at the vertical sensor for at least a portion of the flow channel that is furthest from the laser source. The smallest particles that are near the Rayleigh scattering regime exhibit a scattering minimum for light polarized parallel to the scattering plane, which is the dominant mode received by the vertical sensor. This scattering feature keeps these particles from showing greater intensity at the vertical sensor than the horizontal. Larger particles have a more consistent response for light polarized in both directions when

averaged across the surface of the sensors. Certainly for PSL particles of $0.4\ \mu\text{m}$ and larger, there is a stronger response at the vertical sensor for when the particles pass through the side of the channel furthest from the laser source.

The ratio of the power received at the horizontal sensor to the power received at the vertical sensor is plotted in Figure 70 for the same range of particle sizes. As would be expected, the smallest particles show a fairly high ratio across the entire width of the channel, since the power received at the horizontal sensor is always greater than that received at the vertical sensor for these particles. Larger particles show a relatively tight grouping across the width of the channel, with ratios greater than 1 on the side of the

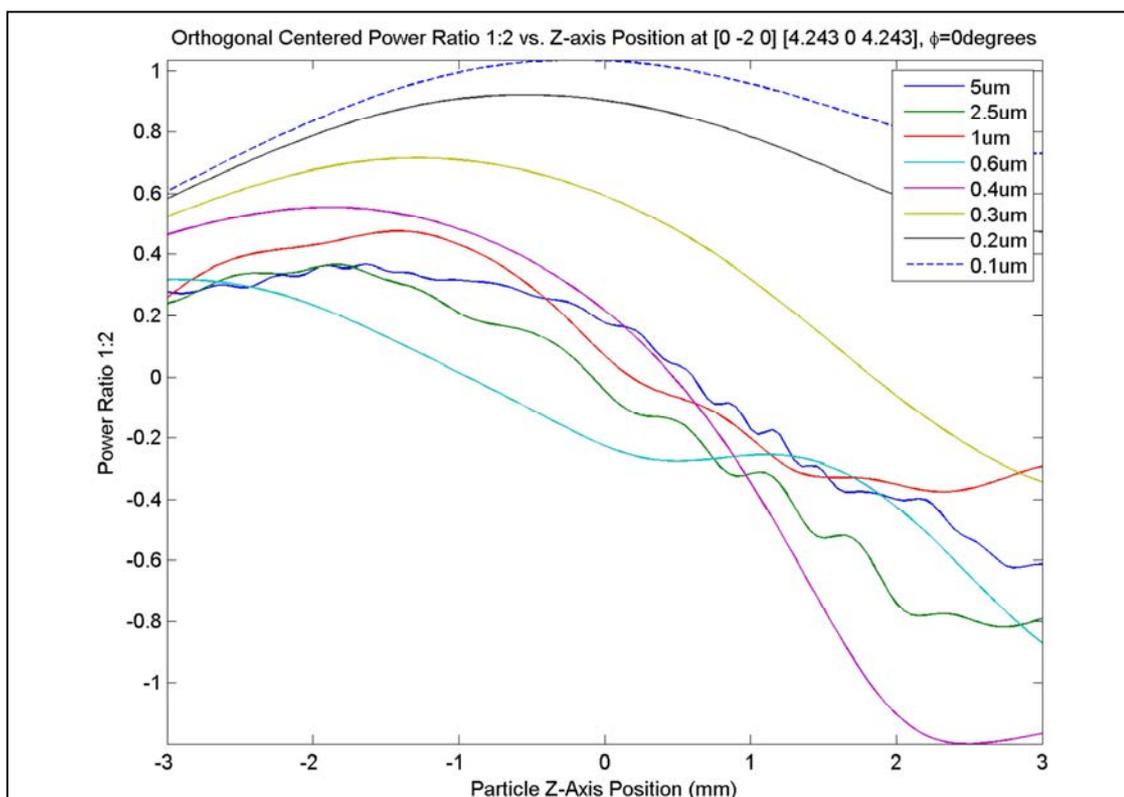


Figure 70 Orthogonal arrangement power ratio received at the horizontal sensor to the vertical sensor plotted as a function of the particle position along the z-axis of the coordinate system, on a logarithmic scale. The laser is oriented with polarization perpendicular to the surface of the sensors. The particle intercepts the center of the beam, $y = 0$, for any given z-axis position.

channel nearest the laser, and ratios less than 1 on the side furthest from the laser. As with the tandem arrangement, this may provide some indication of the position of a particle relative to the z-axis as it passes through the laser beam.

At first glance, the behavior of the smallest particles may be seen as a disadvantage for the orthogonal arrangement, since these particles have a received power ratio that is greater than 1 across the entire width of the channel. These smallest particles then do not show a clear pattern for detecting which side of the channel they are on relative to the horizontal sensor. However, when the absolute intensity of scattered light is taken into consideration along with the sensitivity of the OPC circuitry, this can be turned into an advantage for recognizing these smallest particles. This will be shown in the next chapter as the sensor arrangements are tested and compared with random particle simulations.

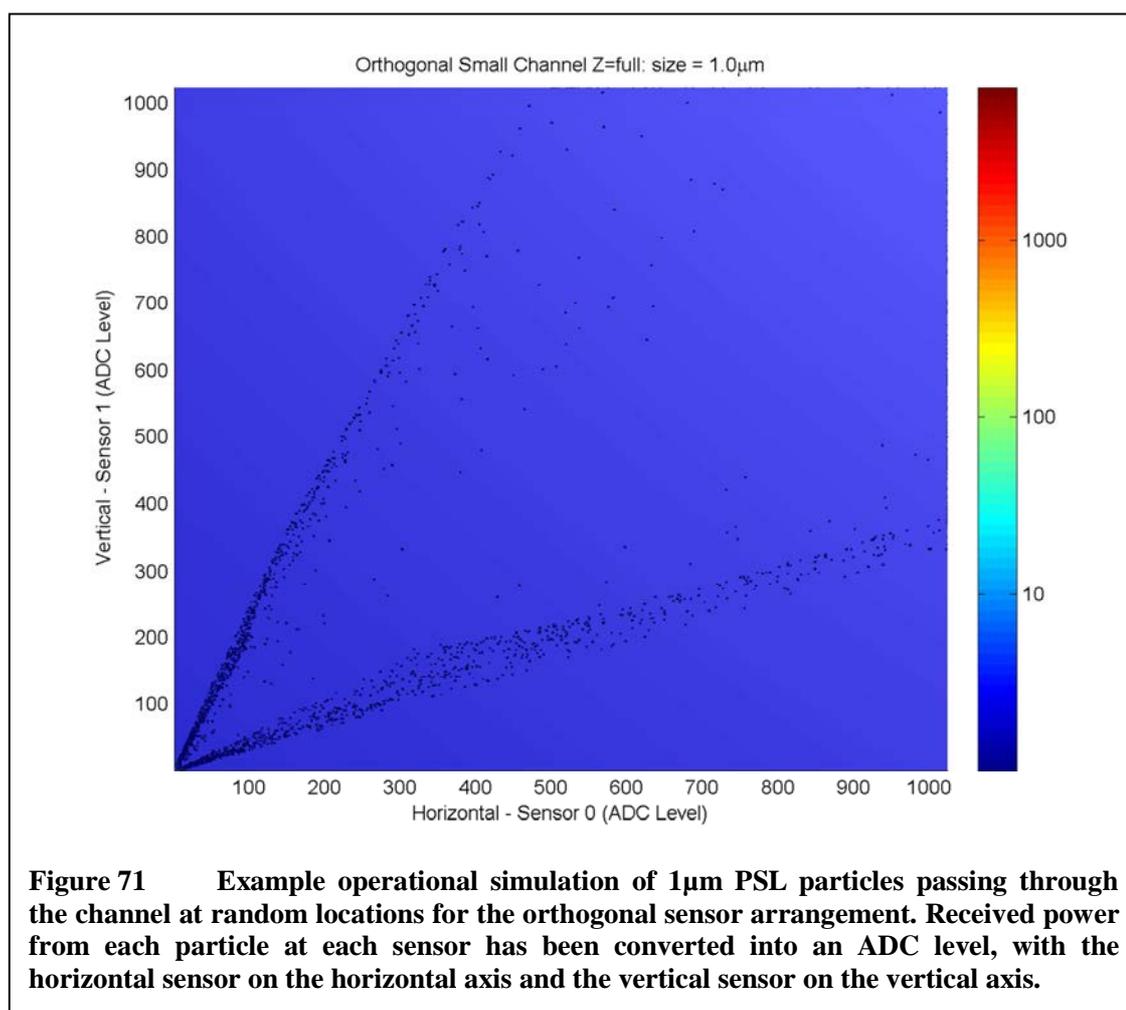
Operational Simulations

The end goal of the construction of the simulation environment is to model the actual operation of an OPC. Towards this goal, the simulation flow was used to create a set of simulations of multiple random particle positions in the flow channel. The simulations were typically run with a single particle size and refractive index at any one time, since it was desired to understand the system response to a given particle.

The main differences in process for the operational simulations compared to those previously discussed are with the particle position generation, and with the display of results. The particle position generation is typically done as a random process, creating a set of particles passing through the flow channel as they would in the actual device. The simulations also typically look only at the peak irradiance produced by a particle, and

would position each particle at the $x = 0$ point in the channel, with the y and z coordinates chosen randomly. This makes the assumption that the peak scattering irradiance occurs at the point where the incident light is most intense. This may not always strictly be true, as the scattering pattern and random position might result in higher intensities elsewhere.

Figure 71 shows an example of the output of an operational simulation of the orthogonal sensor arrangement. In this case, the received power from each $1\ \mu\text{m}$ PSL particle at each sensor has been converted into an analog-to-digital converter (ADC) level by applying the appropriate gain from the transimpedance and small-channel amplifiers.



A surface plot is then created with each particle plotted on the horizontal and vertical sensor axes to give a visual representation of how the particles map into the dual-sensor space. Note that in this case, there are two main groupings of particles that tend to show stronger intensity at either the horizontal or vertical sensor.

Further examples of operational simulations will be given in the next chapter exploring detection and sizing algorithms. These will be presented alongside actual testing data comparing expected and actual results.

Simulation with Other Particle Types

The simulation results shown above were all generated using the physical parameters for polystyrene latex (PSL) material with a refractive index of $1.59+5j \times 10^{-4}$. While PSL is generally used in the calibration of OPC devices, in actual use it is unlikely that PSL will be encountered in the natural environment. It is well documented that optical particle counters will measure an “optical diameter” of a particle, which can differ significantly from the physical diameter depending on the measurement technique and particle properties such as shape and refractive index [34], [72]. It is important to be aware that the optical properties of the particulate matter being measured will impact OPC performance, and as such some other particle types are examined in simulation as well. Data for this section are only presented for the orthogonal sensor arrangement. This arrangement was the focus of much of the work for the detection and sizing, and has the added benefit of including a sensor that is in the same relative location as the single sensor arrangement.

Since pollution and its impact on the health and wellness of individuals is of primary importance to the development of the WAQM OPC, it is important to look at

particles with characteristics similar to common pollution. Various studies have been performed to sample air pollution and determine average particulate matter characteristics, such particulate matter found in urban and rural environments [73], [74].

A comparison of the response for the orthogonal sensor arrangement for PSL and average rural air pollution as described in these studies is shown in Figure 72. A refractive index of $1.53 + j6 \times 10^{-3}$ is assumed for the rural air pollution. Note that the imaginary portion of the refractive index is the absorption coefficient, which though small for the rural air pollution, is still larger than that for PSL. The comparison shows a similar response on both sensors for the two material types, which might be expected given their similar refractive indices. The rural air pollution does show a slightly decreased response, which corresponds to the more absorptive nature of the material.

The differences between responses can be better seen in the logarithmic plot of ratios between the horizontal and vertical sensors in the orthogonal arrangement. These

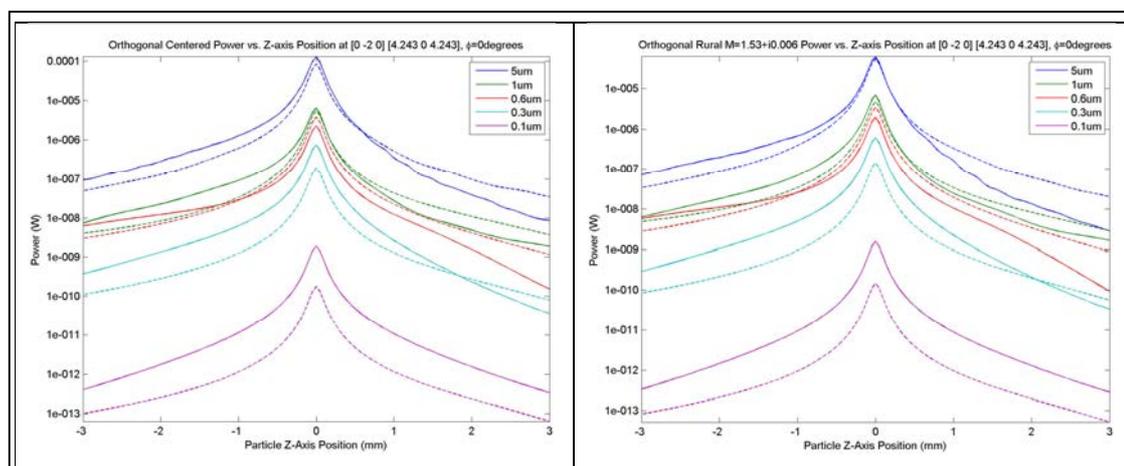


Figure 72 Comparison of PSL (left) and average rural air pollution (right), showing the received power from the respective particle types against the position in the flow channel. The solid lines represent the responses from the horizontal sensor, and the dashed lines represent the responses from the vertical sensor. The two material types result in very similar responses, though the magnitude response is reduced for the rural air pollution.

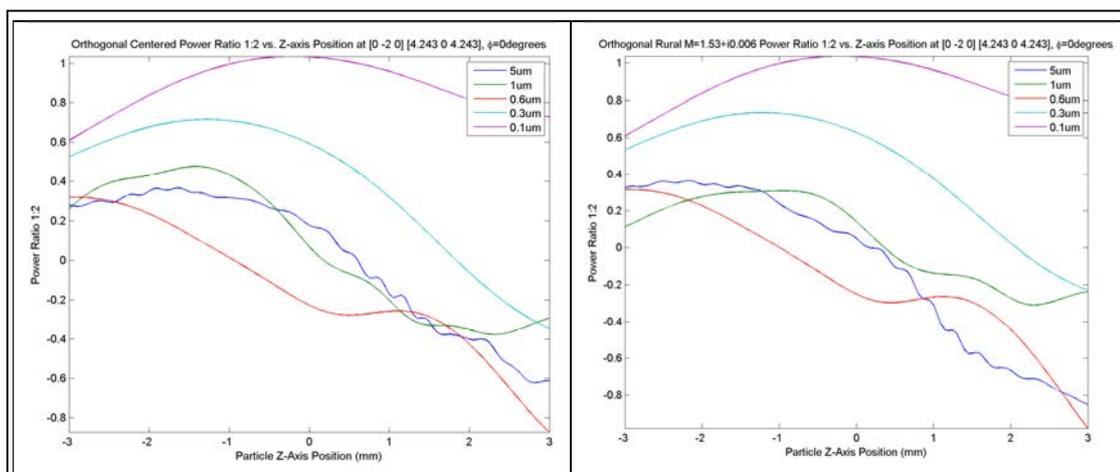
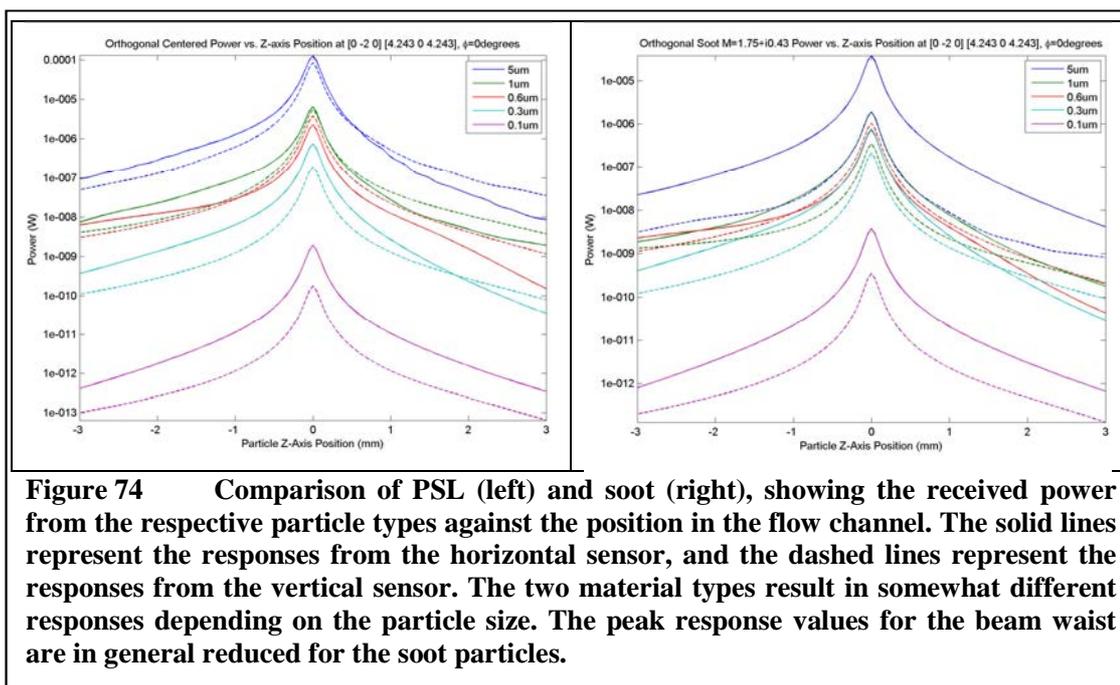


Figure 73 Comparison of PSL (left) and rural air pollution (right) power ratios between the horizontal and vertical sensors on the orthogonal arrangement. The responses are very similar with the response for small particles being greater on the horizontal sensor than the vertical for much of the channel width.

data are shown in Figure 73, with the power ratio for PSL on the left, and for rural air pollution on the right. The differences between the material types are relatively small, especially for the smallest particle sizes. For the small sizes, the response from the horizontal sensor remains higher than the vertical sensor response for most of the width of the channel. The larger particles tend to have higher responses for the horizontal sensor when on the side of the channel closest to the laser source, and higher responses for the vertical sensor when on the opposite side.

Compared to rural air pollution, the refractive index of soot has a significantly different refractive index. With a refractive index of $1.75+0.43j$, this might be considered on the extreme end of what might be encountered when measuring air pollution. A comparison of the response for the orthogonal sensor arrangement for PSL and soot is shown in Figure 74. The differences between PSL and soot are much more pronounced. In general, the response for soot is reduced significantly from the PSL response,



especially for the larger particles at the vertical sensor. The decreased scattering intensity from the soot particles matches what would be expected for a more absorptive material.

The power ratio plots in Figure 75 show more detail in the differences between the two particle types. In this figure, the ratio between the power received at the horizontal and vertical sensors is plotted for a set of particles, with the ratios for PSL on the left and soot on the right. Note that while the ratios for the small particles do not appear to change drastically, the ratios for the larger particles do. The largest particle plotted, 5 μm , has a higher ratio than even the smallest particles. Even the ratio for the 2.5 μm particle size is higher than that for the 0.3 μm particle. The implications for use of the power ratio for particle sizing are significant, since the change in refractive index has now made larger particles produce a power ratio that may be indistinguishable from the smallest particles.

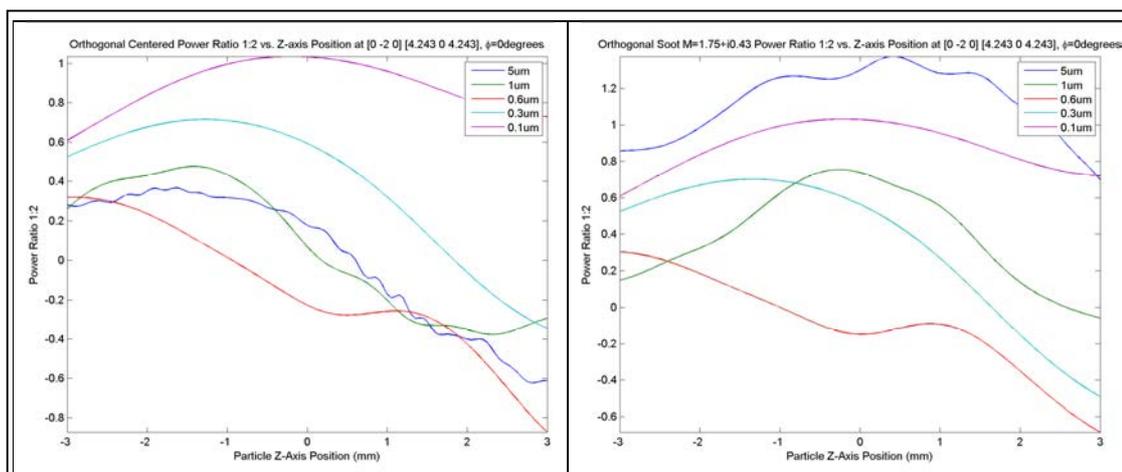


Figure 75 Comparison of PSL (left) and soot (right) power ratios between the horizontal and vertical sensors on the orthogonal arrangement. The responses differ significantly for larger particles, with the response ratio for 5 μm particles being the highest of the group.

For example, consider a 0.3 μm particle passing through the laser beam waist. For both the PSL and soot particles, this will produce a response with a moderate intensity on the horizontal sensor and a relatively low intensity on the vertical sensor, resulting in a ratio greater than one between the two. For 5 μm particles, the PSL particle passing near the beam waist would produce a ratio near unity, but the soot particle would produce a ratio significantly greater than one. Furthermore, if the 5 μm particle did not pass through the most intense portion of the waist, it might produce an absolute intensity very similar to that of a much smaller particle.

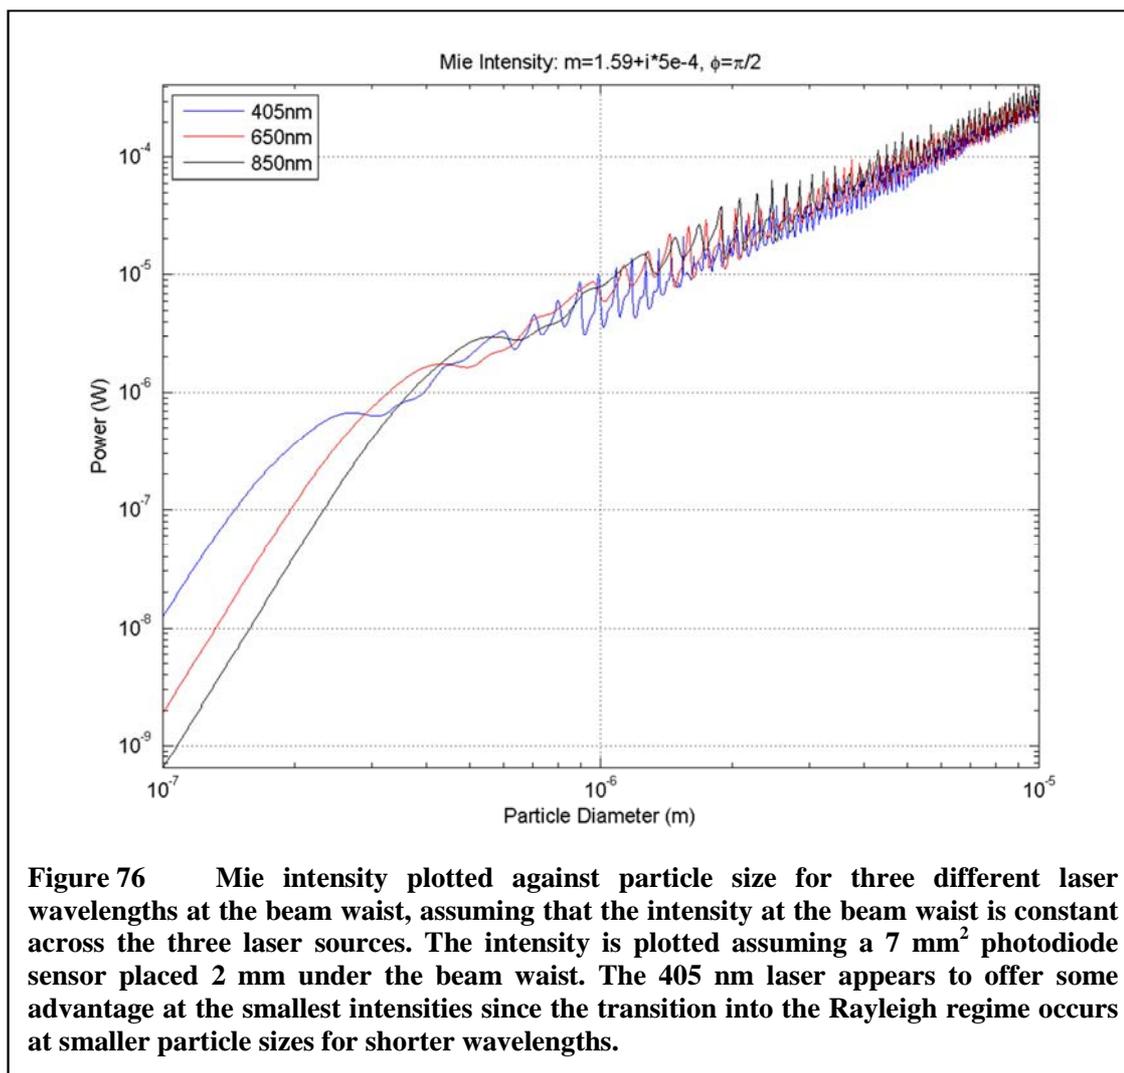
This goes to highlight the issues in designing an OPC that can work effectively across a large range of particle sizes and compositions. When employing an OPC, one must understand what types of particles are being measured and balance that against the type of particle counter being used. The response for a particular size may or may not match across a range of refractive indices, potentially skewing results in ways unexpected by the user.

Modified Laser Wavelength

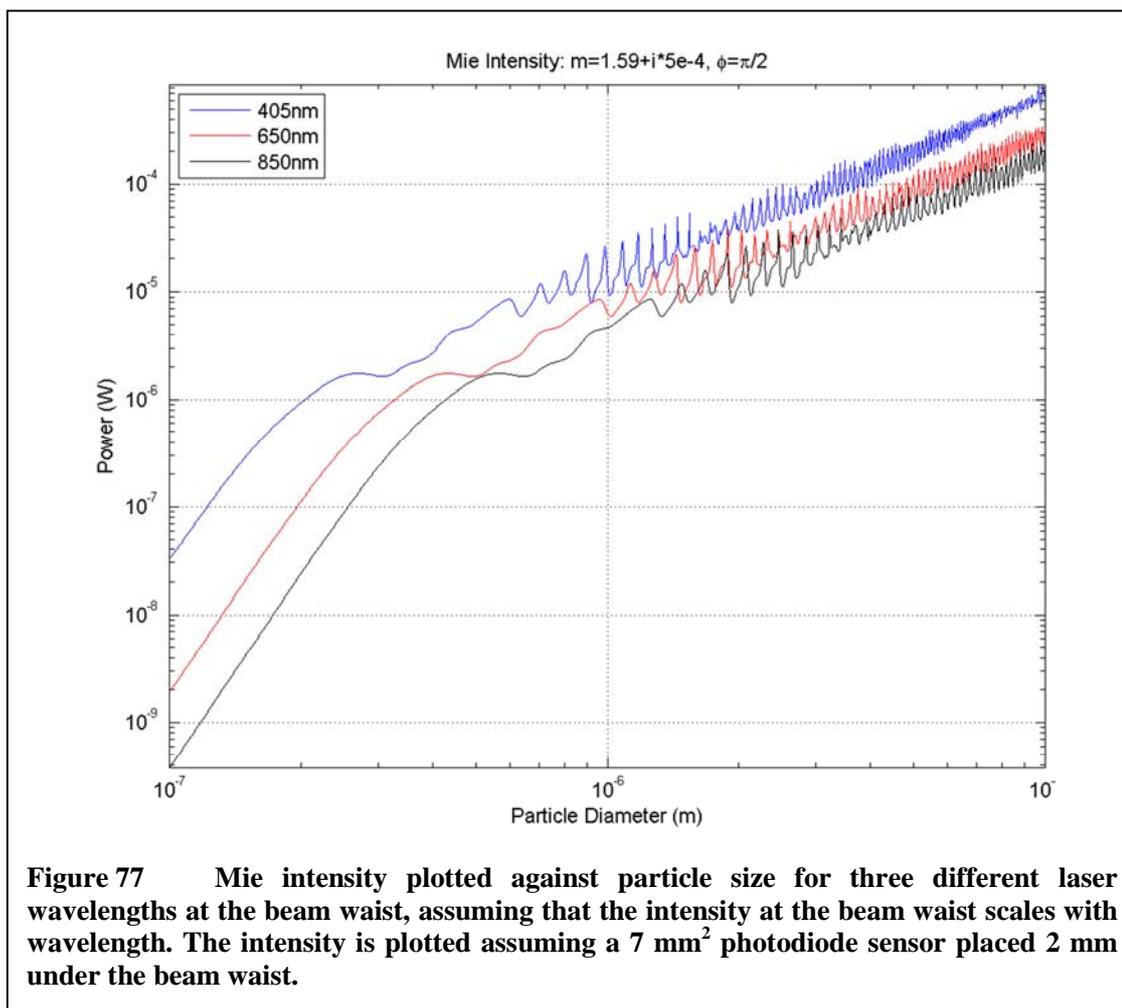
One of the main tenets of the WAQM OPC design is low cost, which pushes one towards use of the least expensive 650 nm red laser as an illumination source. With the advent of optical disk readers employing 405 nm lasers, there is some opportunity for a low cost implementation using this shorter wavelength. There are also options available for infrared lasers, though these may be more expensive and also introduce issues in safety and visual alignment. The question becomes one of which would provide the best performance given the tradeoffs of cost and ease of use.

Some assumptions must be made regarding what physical parameters will change when a different laser source is used. For the purposes of this analysis, it is assumed that the aperture size and focal length of the laser do not change, and that the performance of the laser optics does not change across different wavelengths. These assumptions will most certainly not hold when comparing two different laser sources, so care must be made when applying this study to the real world.

If the intensity at the beam waist does not change, the resulting scattering intensities appear to mostly overlap. Figure 76 shows this case, with the scattered intensities plotted against particle size, captured by a 7 mm² photodiode sensor placed 2 mm under the beam waist as found in the WAQM OPC single sensor arrangement. This is the scattering intensity when the particle passes through the center of the beam waist, and scattering is mostly at a 90 degree angle to the direction of light propagation. Note that the shorter wavelength 405 nm laser appears to have some advantage over the other laser wavelengths at the smallest particle sizes on the plot. The shorter wavelength causes the transition into the steeply sloped Rayleigh regime to occur at smaller particle sizes.



If aperture size of the laser source does not change and the output power of the laser source is constant, then the beam waist for a shorter wavelength will be smaller and produce a corresponding increase in intensity at the waist when compared with longer wavelength sources. This change in intensity will cancel with the wavelength-dependent factor in the Mie scattering calculations. This can be seen in Figure 77, where the beam waist intensity is allowed to change with the change in wavelength. Note that the curves look identical, with just a lateral translation depending on the laser wavelength. The



shorter wavelength 405 nm laser now appears to have a significant advantage across the entire range of particle sizes.

While the performance increase may be attractive, it is important to note that the decrease in the size of the beam waist will also decrease the size of the sampling volume. This decrease will directly impact the number of particles that impact the laser beam, which in turn will reduce the absolute particle count seen by the device. This tradeoff may be desirable if the desire is to sense the smallest possible particle sizes.

Another critical factor in the choice of laser wavelength is the behavior of the photodiode. Photodiodes will have an intrinsic spectral sensitivity curve, typically with a

single peak frequency of highest sensitivity falling off as the incident light moves to higher or lower frequencies. The relative spectral sensitivity for the photodiode used in the WAQM OPC is shown in Figure 78. This curve is reproduced from the Osram BPW34 datasheet [75], and represents the device's sensitivity to different wavelengths of light. The peak sensitivity occurs at 850 nm, which is in the near infrared range. The 650 nm laser used in the WAQM OPC falls at approximately 70% relative sensitivity. Following the curve to shorter wavelengths, the 405 nm response can be seen to be approximately 12%.

This significantly reduced response must be taken into account for the overall system performance. The impact of applying the spectral sensitivity to the Mie scattering curves can be seen in Figure 79. This curve still has the beam waist intensity scaled with wavelength, which benefits the smaller wavelength sources. The curve for the 405 nm laser is now greatly reduced due to the poor photodiode sensitivity at that frequency. The 850 nm laser has the benefit of being at the peak relative sensitivity, but this does not raise it past the level of performance of the 650 nm laser.

One might consider finding a photodiode that has a relative spectral

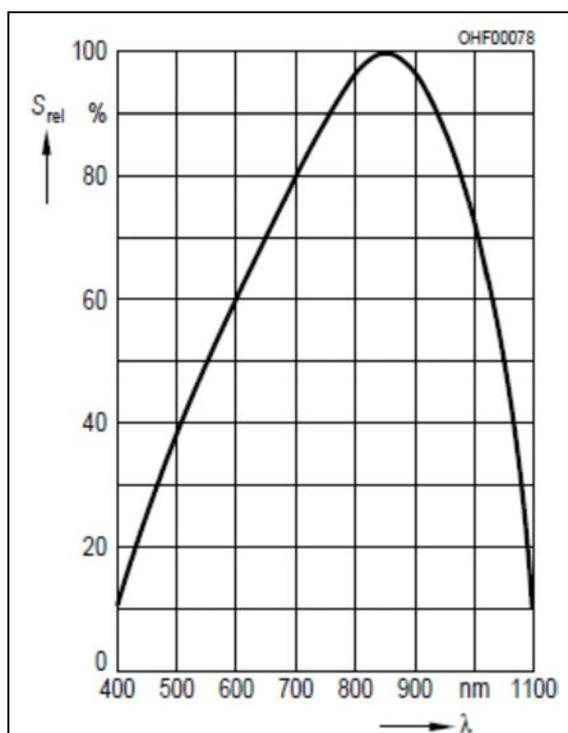
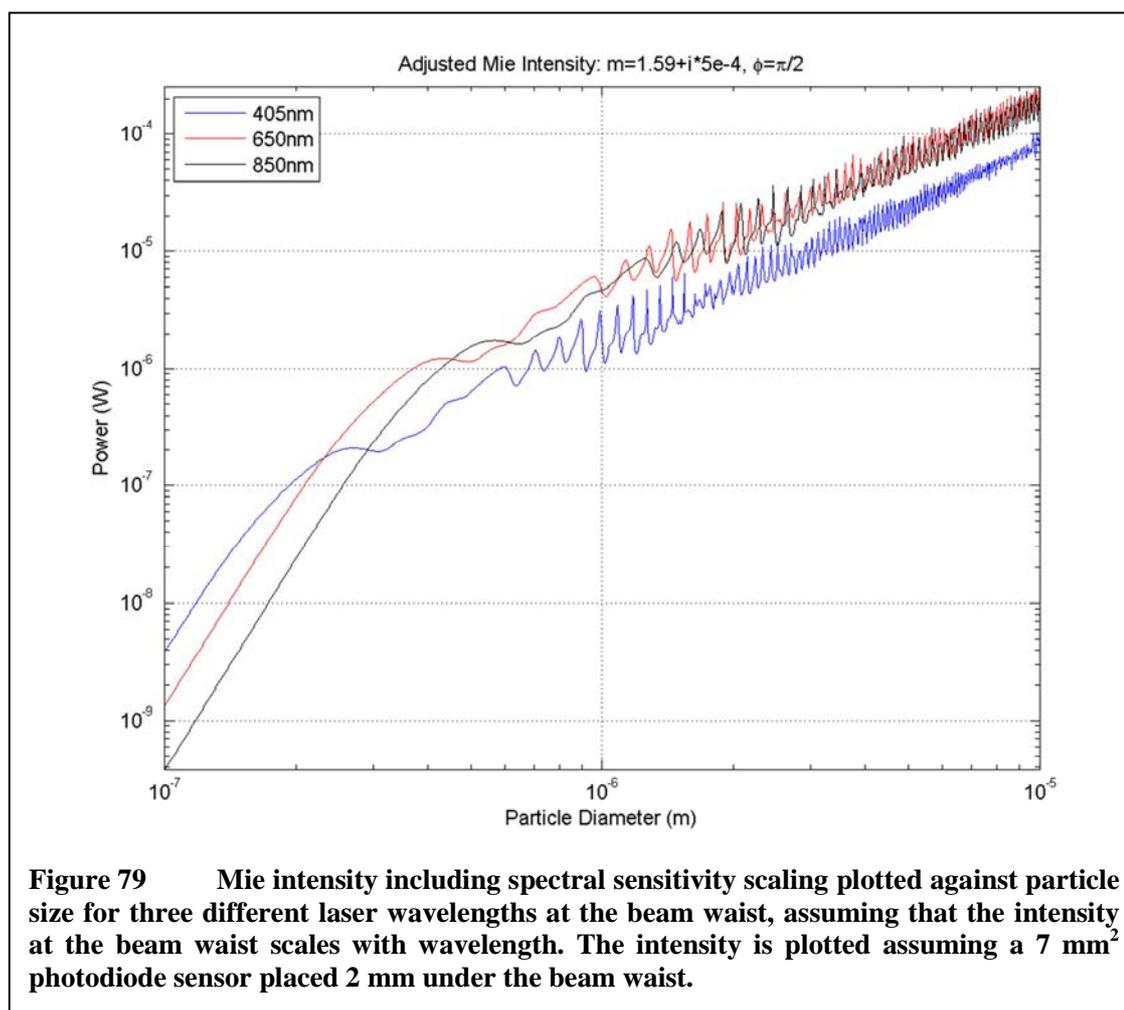


Figure 78 Relative spectral sensitivity curve for the photodiode used in the WAQM OPC design, reproduced from the Osram BPW34 datasheet. The photodiode response peaks at 850 nm.



sensitivity that is more favorable to shorter wavelengths. Moving away from a standard photodiode will most likely increase the cost of the system, and may not produce a result that is any better than the 650 nm laser. For example, photodiodes that are enhanced to be more sensitive to blue/violet light may only get to a 30% relative sensitivity at 405 nm and will be significantly more expensive than the standard photodiode.

Simulation Conclusions

Putting together a modeling tool for the WAQM OPC was a very useful endeavor for multiple reasons. Of primary importance is the simulations ability to convey a sense of the physics behind the OPC operation. Being able to calculate and view the ideal

behavior of the system allows one to quickly test different ideas relative to sensor location and orientation prior to physically constructing a device. For example, several different scenarios of dual sensor positioning relative to the laser beam waist were tested in simulation without ever constructing any but the most promising actual systems. The simulator also allows one to delve into the reasons why certain behaviors are exhibited by particle counters, such as why it is difficult to sense particles smaller than $0.1\ \mu\text{m}$ with an OPC.

The simulation was critical in guiding the development of the WAQM OPC. For example, the simulation helped to solidify the choice of a red laser rather than moving to a 405 nm laser by showing that there wasn't any tangible gain to be had without increasing laser power significantly. The increase in coverage over smaller particles by the shorter wavelength light was outweighed by the loss in photodiode response at those frequencies. The simulation also showed the optimal polarization alignment for the laser module to sense the smallest particles possible. This guided the design towards orienting the laser polarization direction to be parallel to the surface of the horizontal sensor under the beam waist, which would produce scattering that was mostly perpendicular to the scattering plane. This avoided the large null in this direction that occurs at opposite polarization, and gave the best sensitivity results overall.

A final reason for performing modeling with the simulation was to provide verification of behavior of the WAQM OPC. As will be shown in the next chapter, when testing with PSL particles, the simulation was used to correlate with the data taken from the actual OPC design. This helped instill confidence in the developers that the OPC was

functioning as expected, and that decisions made through the simulation would be translate to valid improvements in the OPC implementation.

For all of the above reasons, modeling the OPC has been a positive experience. Improving the accuracy of the simulation as future work may show some merit, whether through better laser models, incorporating airflow and x-axis movement, or allowing for different particle shapes. As it exists, the simulation has been an invaluable tool in understanding and guiding the WAQM OPC development.

CHAPTER SIX: DETECTION AND SIZING

This chapter deals with the implementation of the WAQM optical particle counter (OPC), focusing on the optimization of the particle detection and sizing methods. The WAQM OPC design presents several challenges to accurately detecting and sizing particles. The main issue is in the definition of the sensing volume, or the volume defined by the intersection of the particulate flow and the laser beam. Most of the OPC devices that are commercially available will define the sensing volume by passing a small stream of particles through a relatively uniform area of illumination. This provides a level of predictability in terms of the amount and direction of scattering for a given particle. With the WAQM OPC, the sensing volume is defined with the focused laser passing through a constrained airflow. The laser intensity varies greatly across the width of the flow channel, and particles can pass through the laser at any point. This results in a large variation in possible pulse intensities and durations produced by any given particle, depending on its path through the OPC.

The detection and sizing methods for single and dual sensor arrangements are detailed in this chapter. For the dual sensor implementation, the focus is on the orthogonal sensor arrangement as described in previous sections. This arrangement shows the most promise for improved particle sizing while preserving the maximum sensitivity for detection of the smallest particles.

Particulate Generation and Measurement Comparisons

Before beginning, it is important to note that it can be difficult to produce agreement between multiple optical particle counters. This often holds true even for two optical particle counters of the same model from the same manufacturer. Calibration is critical, and various factors will cause output values to shift over time. Different models of OPC may employ different sensing and sizing methods, and may calibrate in different ways, all of which can lead to mismatch in output when measuring the same aerosol. While this is not studied extensively in this text, it should be kept in mind that achieving results that are a precise match to commercial OPC devices is a problematic endeavor at best.

When creating a new OPC, it is important to be able to evaluate its performance relative to some known standard. Many methods exist for performing this evaluation, such as creating an aerosol of known concentration and particle size, or using a gold-standard device against which to compare. Methods for generating an aerosol of known composition may be preferable for this task, as being able to control the aerosol composition precisely allows a greater degree of determinism in testing a new OPC. Unfortunately, the options open to the team creating the WAQM OPC did not allow for the expenditures required to build such a system. Equipment such as a differential mobility analyzer was outside the range of funding available when creating the WAQM OPC.

The remaining option, to use a commercial OPC as a golden-standard, did fit into the available budget. With this in mind, two different systems were purchased. The first was a handheld device: the TSI 9303 Aerotrak [76] optical particle counter. This device

is able to measure concentrations up to 2 million particles per cubic foot with a flow rate of 0.1 cubic feet per minute. The device measures three particle sizes simultaneously between 0.3 μm and 25 μm , with the middle range adjustable in cutoff size. This device was relatively inexpensive for a commercial OPC, being just under \$3,000 when purchased new. The second device purchased as a gold-standard was the Particle Measuring Systems (PMS) Lasair II-110 [77]. This device is able to measure up to 500,000 particles per cubic foot with a flow rate of 1.0 cubic feet per minute. The device has five sizing ranges from 0.1 μm up through 5.0 μm . This was a much more expensive device than the TSI OPC, at nearly \$20,000 for one unit.

Most of the data shown in this chapter utilizes the PMS device to compare to the WAQM OPC. In general, this device is considered to produce higher quality measurements than the TSI handheld units, and is used as the golden standard for the purposes of the WAQM OPC development. The TSI handheld units are used as calibration standards when a higher concentration of particles is desired.

Particulate matter is generated for testing the WAQM OPC using a Collison nebulizer from BGI, Incorporated. The nebulizer is capable of producing an aerosol from a liquid solution such as mixtures of deionized water and polystyrene latex (PSL). This is the type of solution used to produce the test data shown in this chapter. The PSL particles themselves are specifically made for the calibration of OPC systems. This system performed well when generating particles in the 0.3 μm to 2.0 μm range, but was not able to successfully generate particles of larger size. Further details of the calibration and particle generation processes will be discussed in the chapter on OPC calibration.

Single Sensor Detection and Sizing

The most basic implementation of the WAQM OPC employs a single photodiode sensor to detect and size particles. The basic concept of using a focused laser with a sensor near to the beam has been derived from the work by Dylos Corporation to create a low-cost commercial OPC [39], [78]. The WAQM OPC employs a transimpedance amplifier to convert the photodiode current into a voltage, as described above in Chapter Four. The single sensor implementation uses two amplifiers to create a pair of channels, one for particle detection and one for particle sizing. The particle detection channel uses a high-gain to allow the sensing of very small particles. The particle sizing channel has a much lower gain to cover a higher intensity range of particles. These channels are also referred to as the “small channel” and “large channel,” respectively, as the large channel is used to determine whether an individual particle is above the threshold to be considered “large.”

The number of logical sizing channels is somewhat arbitrary, as a particle could be placed in one of many size groupings based on the information obtained from the photodiode. Given the multitude of issues in sizing with the WAQM OPC laser and sensor arrangement, it was decided to keep the sizing channels to two: large and small. The actual size of the large determination is somewhat arbitrary, but for the purposes of the WAQM OPC this has been set to 1.0 μm . There are multiple reasons for choosing this cutoff point. For one, this size is near to the 2.5 μm respirable fraction cutoff as defined by the US Environmental Protection Agency. For another, most bacteria are larger than 1 μm , which is an important consideration for some WAQM OPC target applications. Finally, the particle generation equipment available for this project was somewhat limited

in its ability to aerosolize particulate matter in sizes much larger than 2 μm . Keeping the large cutoff below this range facilitated the development of sizing algorithms within acceptable budget constraints.

Particle Detection

As described previously, the WAQM OPC system consists of a laser beam in near proximity to a photodiode. Light scattered by particles passing through the beam is captured by the photodiode and converted into a voltage by a set of amplifiers. The voltage is then converted into a digital value by an analog to digital convertor (ADC) in the system's main microcontroller. This ADC produces a stream of digital values that represent the amount of light seen by the photodiode at any given time. The microcontroller must then parse this stream, looking for the characteristic signature of particles passing through the laser.

Due to the biasing requirements of the amplifiers, the steady-state input to the ADC is a non-zero value. This value is typically about 180 mV with no ambient light, which translates to an ADC value of approximately 60 out of the 1024 levels for the 10-bit ADC. This value may fluctuate for various reasons. First, component tolerances from the photodiode, amplifiers, microcontroller, or power supplies may change the level produced by the ADC from unit-to-unit, or even for the same unit as component performance drifts over time. Secondly, undesired ambient light may be intercepted by the photodiode, from either the laser source or from a source external to the unit. Finally, the OPC circuitry may be affected by electromagnetic noise sources that are typically external to the unit. The level seen by the ADC without any particulate light scattering

may change slowly or rapidly, and it is desired that these changes should not affect the performance of the OPC.

Changes to the ADC level due to component tolerances, ambient laser light, and some sources of external ambient light (e.g., sunlight) typically happen very slowly, and may be considered an unchanging bias to the level. Other disturbances to the ADC level such as electromagnetic noise and some types of external ambient light pollution (e.g., fluorescent lights) may cause rapid and repetitive changes to the level. If possible, these disturbances to the ADC level should be removed so as not to affect the operation of the OPC.

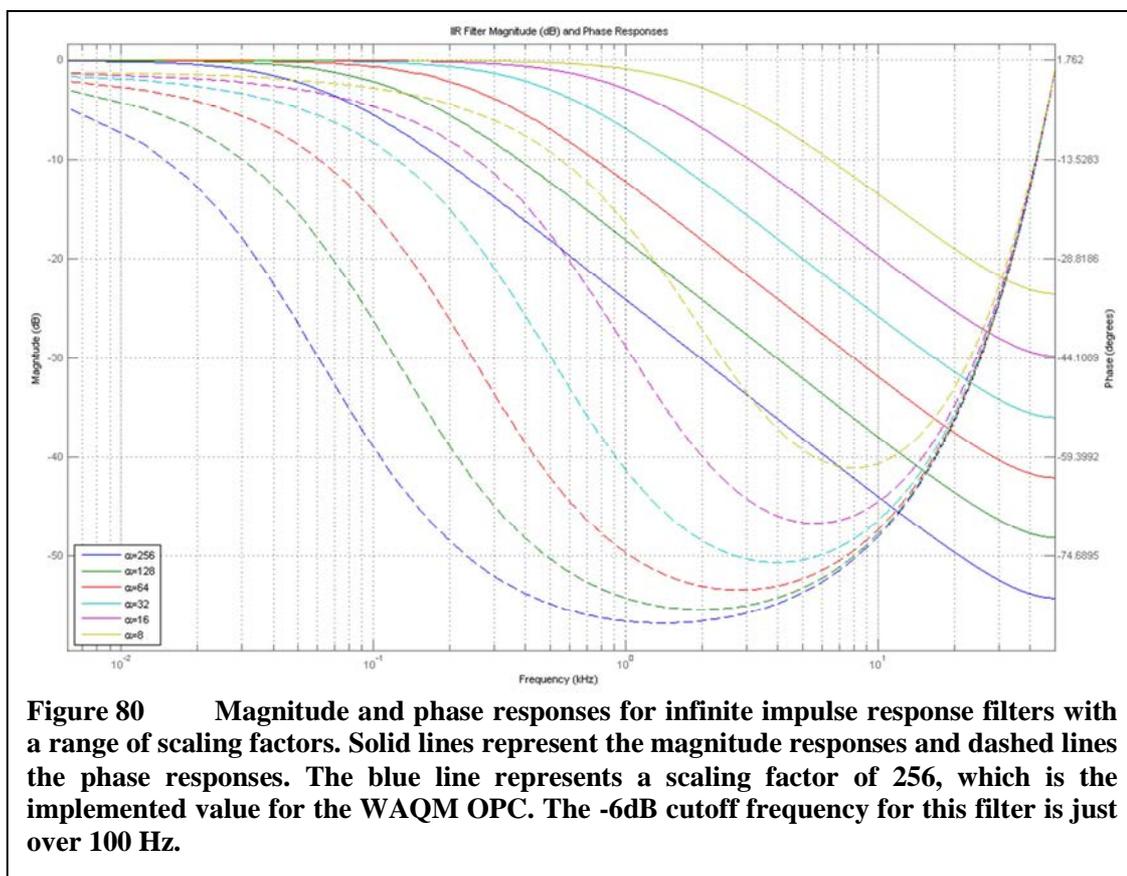
To help mitigate these issues, a virtual average is created for the baseline output of the ADC. This baseline will be used as the “zero” value against which high frequency changes, such as pulses due to particulate light scattering, will be compared. To create this baseline, a low-pass filter is implemented in firmware to remove the bias changes and some of the low frequency noise from the desired signal. The filter implemented is an infinite impulse response (IIR) autoregressive moving average filter, which has been selected specifically to present minimal load to the system processor. This IIR filter takes the form of

$$y[n] = \frac{\alpha - 1}{\alpha} y[n - 1] + \frac{1}{\alpha} x[n] \quad 6.1$$

where $y[n]$ and $x[n]$ are the output and input respectively at time step n , $y[n-1]$ is the output at the previous time step, and α is a scaling factor. Choosing the scaling factor α to be a power of 2 allows the calculation to be performed with shifting and subtraction rather than with multiplication and division.

Figure 80 shows the magnitude and phase responses for the selected IIR filter with various scaling values. These values are all powers of 2, from 256 down to 8. For the final implementation, a scaling value of 256 was chosen, which has a -6dB cutoff frequency of just over 100 Hz. This particular frequency provides some attenuation of 60 Hz AC line noise without significantly impacting legitimate particle-caused pulses.

All values are represented as integers in the firmware implementation of the IIR filter, since floating point operations take significantly longer to process with the WAQM microcontroller. To perform the IIR calculation with shifting and subtraction, a right shift of 8 bits is required. For typical ADC values in the range of 60, this shift would reduce the intermediate calculated values to zero, which is detrimental to the correct operation of



the filter. To get around this, the input and output values are created from the ADC values by first left-shifting by 8 bits, effectively multiplying the values by 256. Subsequent intermediate 8-bit right-shifts then do not remove low-order bits that are important to the averaging calculation. This does add a set of right and left shift operations to the burden on the processor, but the load is minimal as they are performed as single instructions that are implemented in hardware.

With a baseline implemented, basic particle detection occurs by finding values that deviate in the positive direction from the baseline by a set number of ADC levels. This set amount is referred to as the “trigger,” and is something that can be calibrated to match a specific design. Once the trigger value is exceeded, the detection algorithm then looks for a negative transition below a second value, called the “return,” to mark the end of the particle pulse. The return magnitude is typically less than the trigger value to ensure that a particle pulse is not prematurely marked as having ended. This return value is not typically calibrated per unit, since the return signifies a lack of scattered light and other factors such as ambient light are removed by the baseline IIR filter.

Figure 81 shows a graphical representation of a particle pulse as captured by a WAQM OPC. The WAQM OPC contains a debug function that allows the capture of particle trace data directly from the detection algorithm. This trace data are stored and exported into a file on the WAQM system SD card. The trace data are then post-processed by a MATLAB script that builds the graphical representation as shown in the figure. This graph shows three lines. The red line is the baseline as constructed by the IIR moving average filter. The blue and green lines are the output of the small and large channels respectively. Note that the large channel output is lower magnitude than the

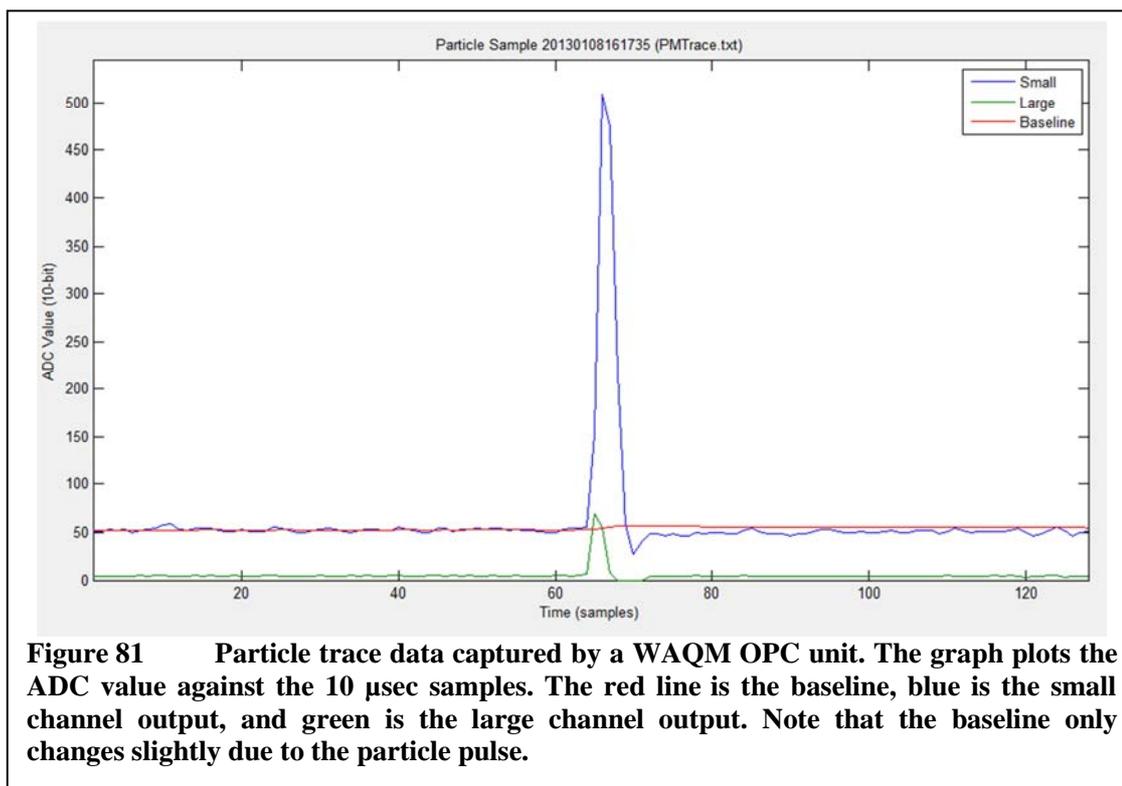


Figure 81 Particle trace data captured by a WAQM OPC unit. The graph plots the ADC value against the 10 μ sec samples. The red line is the baseline, blue is the small channel output, and green is the large channel output. Note that the baseline only changes slightly due to the particle pulse.

small channel since the large channel has less gain. The small channel also shows the bias mentioned previously, which is resting at just above 50 ADC levels. This is also the level that the baseline moves to, and is the starting point for the trigger and return calculations.

Further processing is performed on the particle once the return is detected, which marks the end of the particle pulse. As each sample is processed past the start of the particle, the maximum ADC level, or pulse height, reached during the particle pulse is stored for the small and large channels. The maximum for the small channel is calculated relative to the baseline, and the maximum for the large channel is taken from the zero ADC level. The bias is ignored for the large channel, since the channel's expanded range makes the bias a much less significant portion of the overall measurement, typically 0.5% of the large channel range. The count of samples between trigger and return is also

stored, indicating the duration of the particle pulse. These factors are then used to determine the validity and size classification of the particle.

For the WAQM OPC single-channel arrangement, the validity of a particle is based solely on the maximum ADC level reached on the small channel. A minimum small channel threshold must be met or the particle will be rejected. This threshold will vary from unit to unit based on component and system variability, and must be calibrated on a per-unit basis. This is different from the trigger, which is typically determined once for a particular design. The only limitation from the standpoint of the firmware implementation is that the small channel threshold be greater than or equal to the trigger.

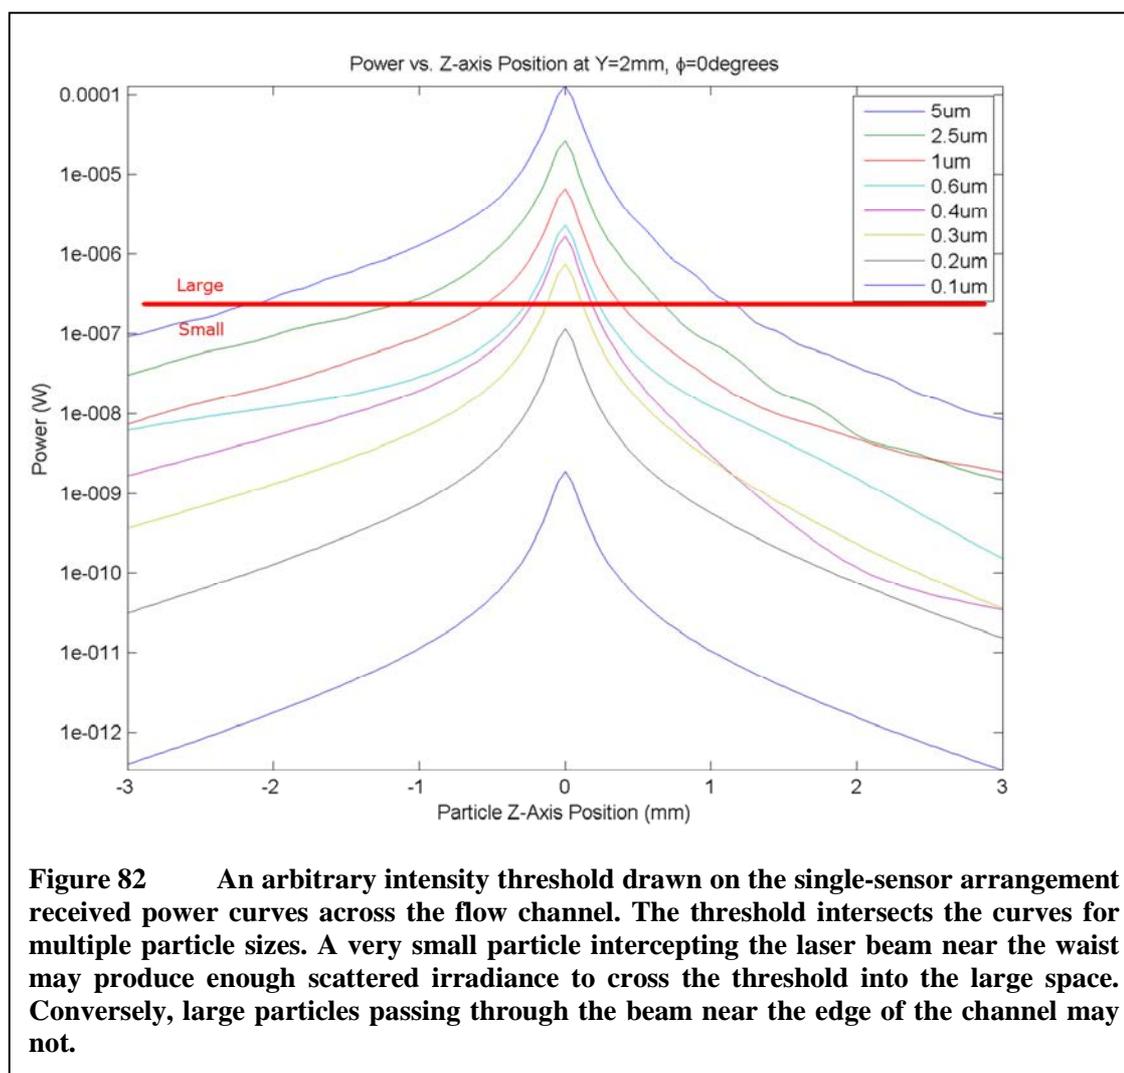
The trigger may be thought of as the indication for the unit to perform processing on a particle. The small channel threshold then, based on this processing, determines whether or not the particle is accepted. The trigger is typically kept at a fixed value per design when calibrating, so that each unit can be compared against a common basis. The small channel threshold is then set at a level greater than or equal to the trigger, depending on the unit-specific sensitivity. It may be noted that having the trigger lower than the small channel threshold will result in the unnecessary processing of particles that fall into the gap. This is true for many applications, and to conserve processor bandwidth and/or power, the trigger may also be raised to the small channel threshold on a per-unit basis.

Particle Sizing

The single sensor arrangement only attempts to categorize accepted particles into two different size ranges. While finer differentiation may be possible, sizing into smaller ranges will suffer from accuracy issues. For the focused beam arrangement of the

WAQM OPC with a single sensor, even just the single differentiation into two size ranges has significant accuracy issues, as will be explored below.

The single channel arrangement performs sizing based on the maximum intensity of the large channel for a given particle. This overlays a sizing on the received power profile as shown in Figure 82. The set of curves shows the received power across the width of the flow channel for a set of particle sizes. The threshold drawn in red is an arbitrary threshold to demonstrate the methodology. This single intensity limit has

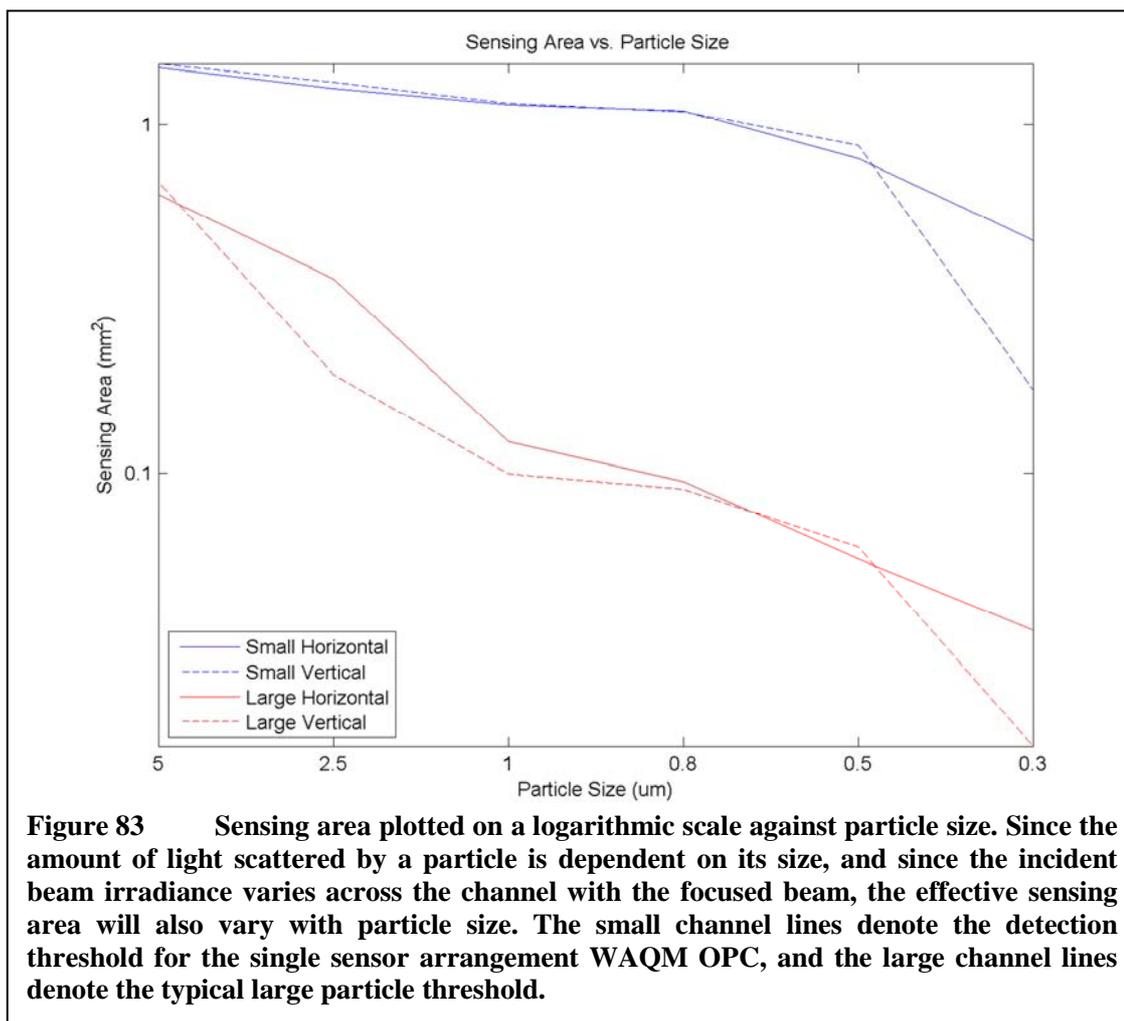


significant issues that are apparent from the graph. The most obvious issue is in the position dependence of the size determination. For example, a 0.6 μm particle intercepting the laser near the beam waist may cross the large threshold, while the same particle passing closer to one of the edges of the flow channel will not.

Another major issue is with scaling of the sensing area. The higher the threshold is set, the smaller the overall sensing area in terms of the available area where the laser is of high enough incident irradiance to cause scattering that will exceed the threshold. For example, if a threshold is set to avoid detecting all particles of a certain size by being higher than the peak scattering irradiance at the beam waist, then particles sizes slightly larger than this will only cause positive large determination when passing through a very small portion of the beam. Thus, counting for these near-threshold particles will be very nearly zero simply due to the low probability that a particle of this size would pass through such a small portion of the beam.

This means that to reliably detect some portion of a certain particle size passing through the flow channel, the threshold must be set low enough to include a relatively large area of the beam for sensing. This will inherently allow some much smaller particles to be detected and identified as large, since they may pass through a higher intensity portion of the beam. No reasonably effective way has been discovered to avoid this with the focused laser arrangement of the WAQM OPC and a single sensor.

This phenomenon is illustrated in Figure 83 where the effective sensing area is plotted against particle size. This plot was produced by simulation, finding the sensing area by calculating the maximum received scattering irradiance by the different particle sizes across the $x = 0$ plane of the flow channel. Note that the sensing area is particle size



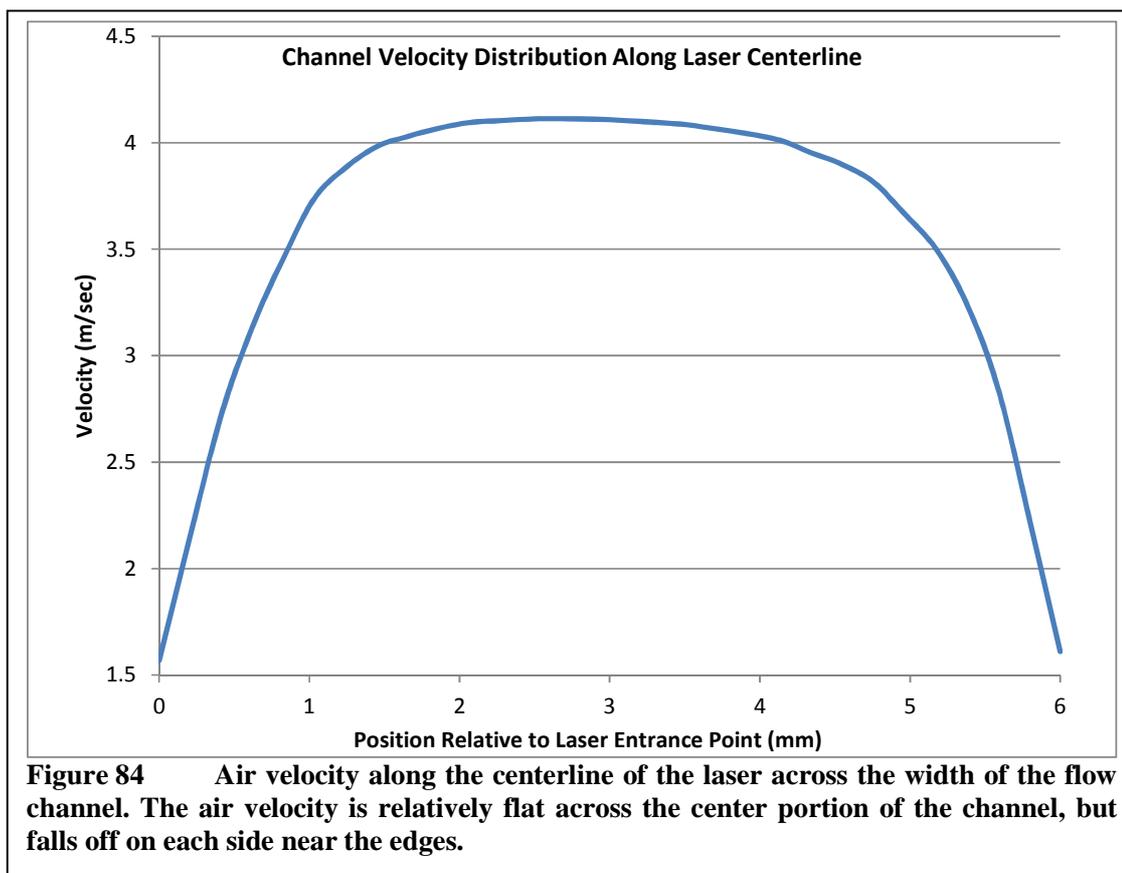
dependent due to the variation in laser beam irradiance across the channel. The variation in laser irradiance is due to the focusing of the beam in the center of the channel. Large particles, scattering more light than small particles, will cross the thresholds for a larger cross-section of the laser beam than will small particles. Note that the vertical sensor falls off dramatically for very small particles with diameters under $0.5 \mu\text{m}$ due to lessened forward scattering.

This feature of the WAQM OPC also impacts the accuracy of scaling a given particle size range. For example, if the small particle range is taken as $0.3 \mu\text{m} - 1.0 \mu\text{m}$, one can see that there is a smaller effective sensing area for the small-sized end of the

range than there is for the large-sized end. When measuring the number of particles passing through the flow channel for a given unit of time, the smaller particles will be undercounted relative to large particles due to the reduced sensing area. This reduced sensing area decreases the total volume of air that will be sampled for smaller particles for a given unit of time.

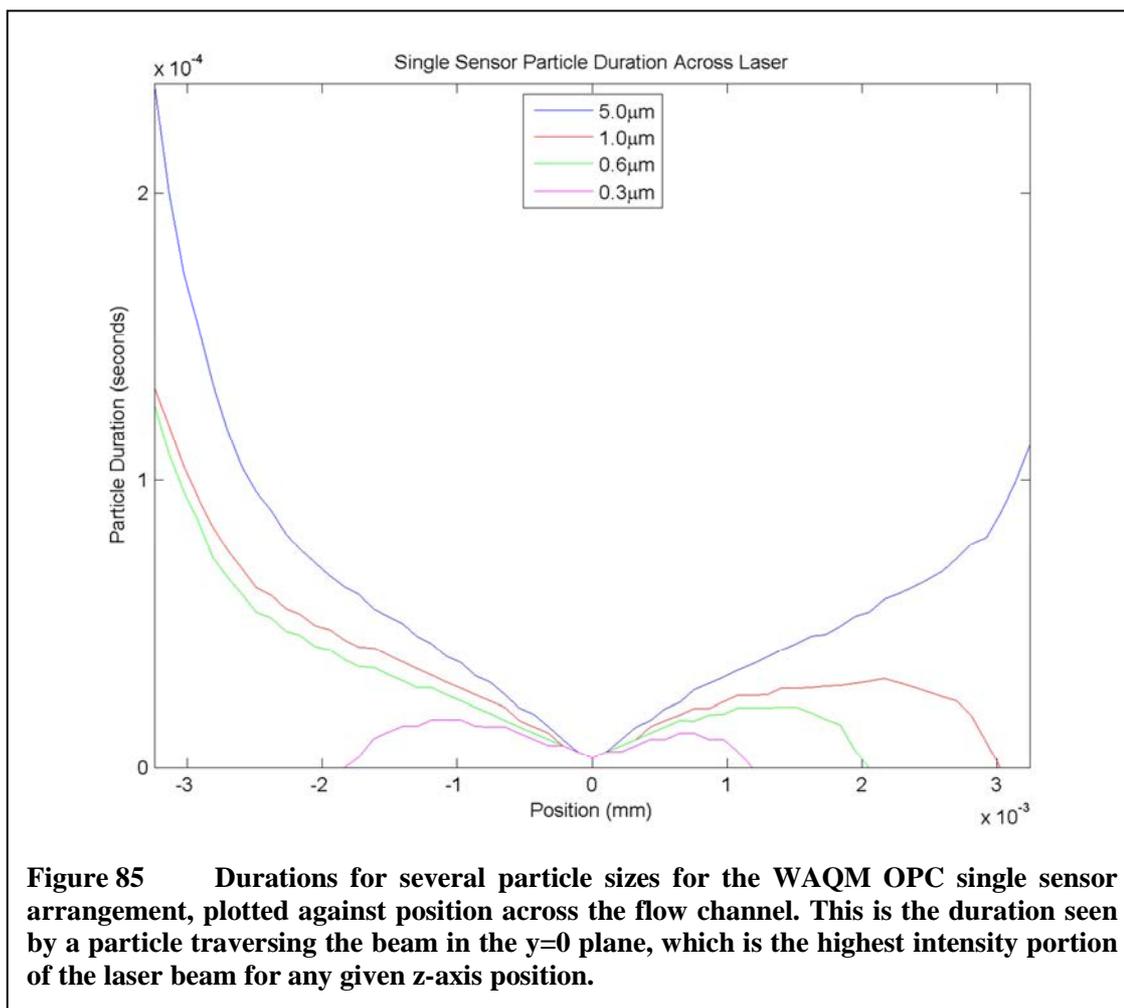
The duration of a particle may also be used in sizing. The duration of a particle may depend on the portion of the laser beam that is traversed by the particle, and on the maximum irradiance generated by the particle. The portion of the beam traversed will impact the duration, since the beam will be relatively wide near the edges of the flow channel and narrow at the waist. The wider portions of the beam will naturally produce longer duration pulses of scattered light due to the longer time to transit the beam. This must be balanced against the reduced incident light irradiance in the wider portions of the beam. This reduced irradiance will work to narrow the portion of the particle-caused pulse that is above the detection threshold.

A simplistic model might assume that the air velocity across the channel is uniform, but in fact it is not. The air velocity for the WAQM OPC flow channel is shown in Figure 84. This plot was produced using an ANSYS computational fluid dynamics simulation of the flow channel. This curve shows the air velocity across the flow channel at the center axis of the laser beam. The air velocity is highest at the center of the channel, remaining somewhat flat for a significant portion of the width and then falling off towards the edges. The air velocity profile amplifies the duration extending effect of the beam geometry across the flow channel, since the air velocity is highest where the beam is narrowest, and the velocity is lowest toward the edges where the beam is widest.



Note that the air velocity across the flow channel has not been measured directly due to the small size of the channel and the lack of equipment on a small enough scale to perform such a measurement.

Incorporating the air velocity into a calculation of the effective sensing area results can provide insight into the particle duration that may be seen across the width of the flow channel. Figure 85 shows the durations expected for four different particle sizes as they traverse the laser beam when the air velocity is applied to the distance through the beam that the particle will scatter light above the detection threshold. This is produced for the WAQM OPC single sensor arrangement. In this case, the duration is measured only in the $y=0$ plane, where the particle will transit through the center of the beam. Particles passing through the channel above or below the $y=0$ plane would likely produce shorter



duration pulses due to the off-axis narrowing of the beam. As with other simulation plots, the position is relative to the z -axis, with the zero point at the beam waist, negative positions closest to the laser source, and positive positions furthest from the laser source.

As would be expected, the larger particles produce longer durations pulses in general, and are able to create detectable pulses even out to the edges of the flow channel. Due to stronger forward scattering, there is a tendency for particles closest to the laser source to generate detectable pulses over larger areas than those furthest from the source. Particles transiting the laser on the far side of the waist from the source will rely on scattering at angles greater than 90 degrees, closer to backscattering than forward

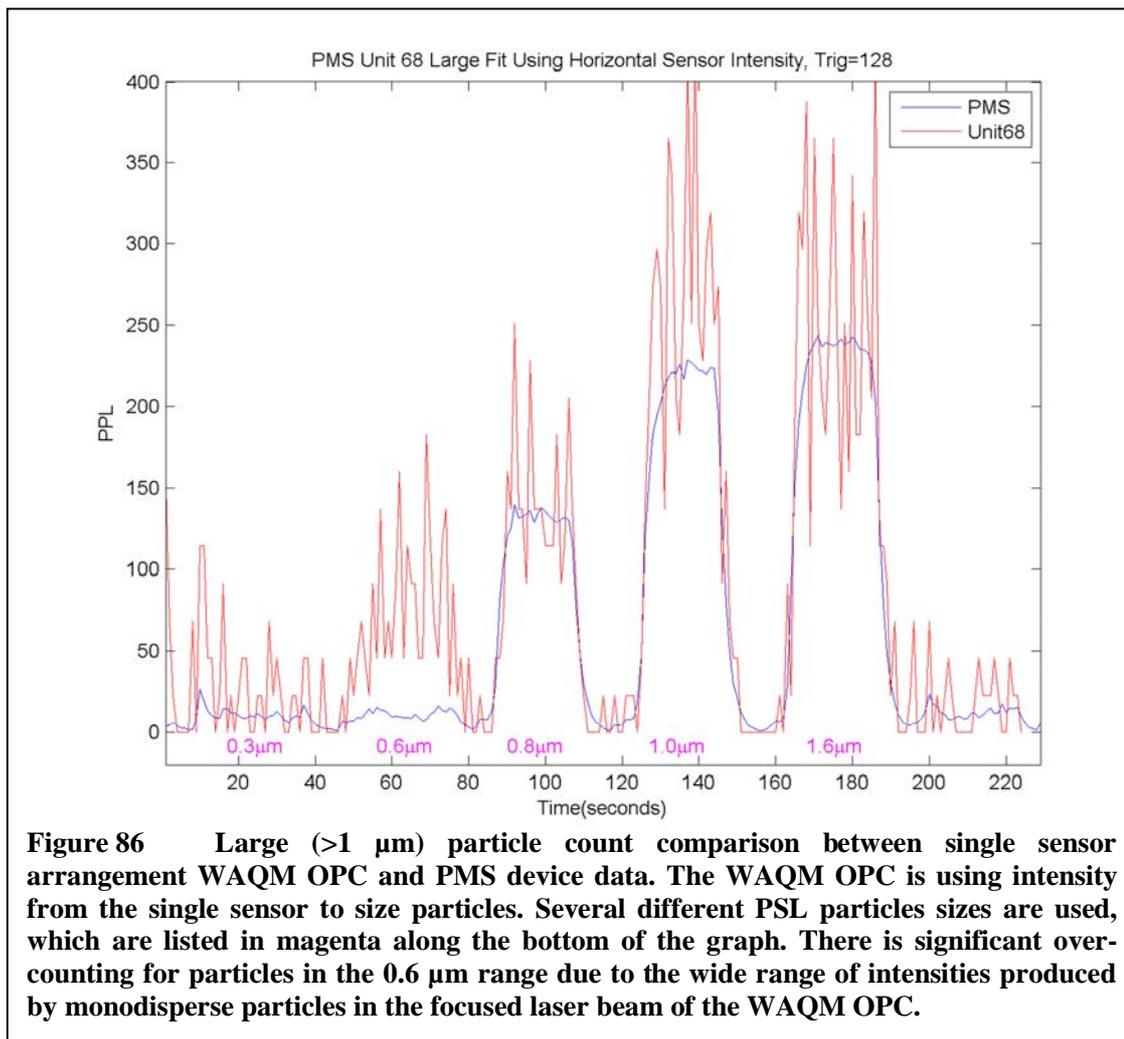
scattering. As a result, particles in these positions may not be detected at all. Particles moving close to the edge of the flow channel nearest the laser source can generate very long pulses due to the slow air velocity, wide beam width, and stronger forward scattering. It is important to note that for the smallest particles, particles moving through the areas of the laser beam near to the edges of the flow channel will not be detected. This will be important to the subsequent examination of duration as a sizing metric.

Example Data

To demonstrate the detection and sizing of the WAQM OPC single sensor arrangement, a series of particles of several different sizes is generated and injected into a test chamber. The WAQM OPC is compared against the PMS OPC running simultaneously in the same test chamber. Further details of the testing setup will be discussed in the chapter on calibration.

The test set of particles includes 0.3 μm , 0.6 μm , 0.8 μm , 1.0 μm , and 1.6 μm PSL particles. These particles are atomized and fed into the test chamber sequentially, clearing the chamber with clean air in between each particle type. Figure 86 shows the large particle results from the test with threshold set to 1.0 μm . The sizing performed using only pulse height in terms of the maximum ADC level seen for the particle. The results for the WAQM OPC are shown in red, and results for the PMS OPC are shown in blue. Measurements for the both devices are averaged over a 60 second period and reported in particles per liter.

It is immediately apparent that the WAQM OPC data has much greater variability than that of the PMS OPC. This is due mostly to the difference in the sampling volumes between the two devices. The PMS device samples 1 $\text{ft}^3/\text{minute}$, or 28.3 L/minute, and

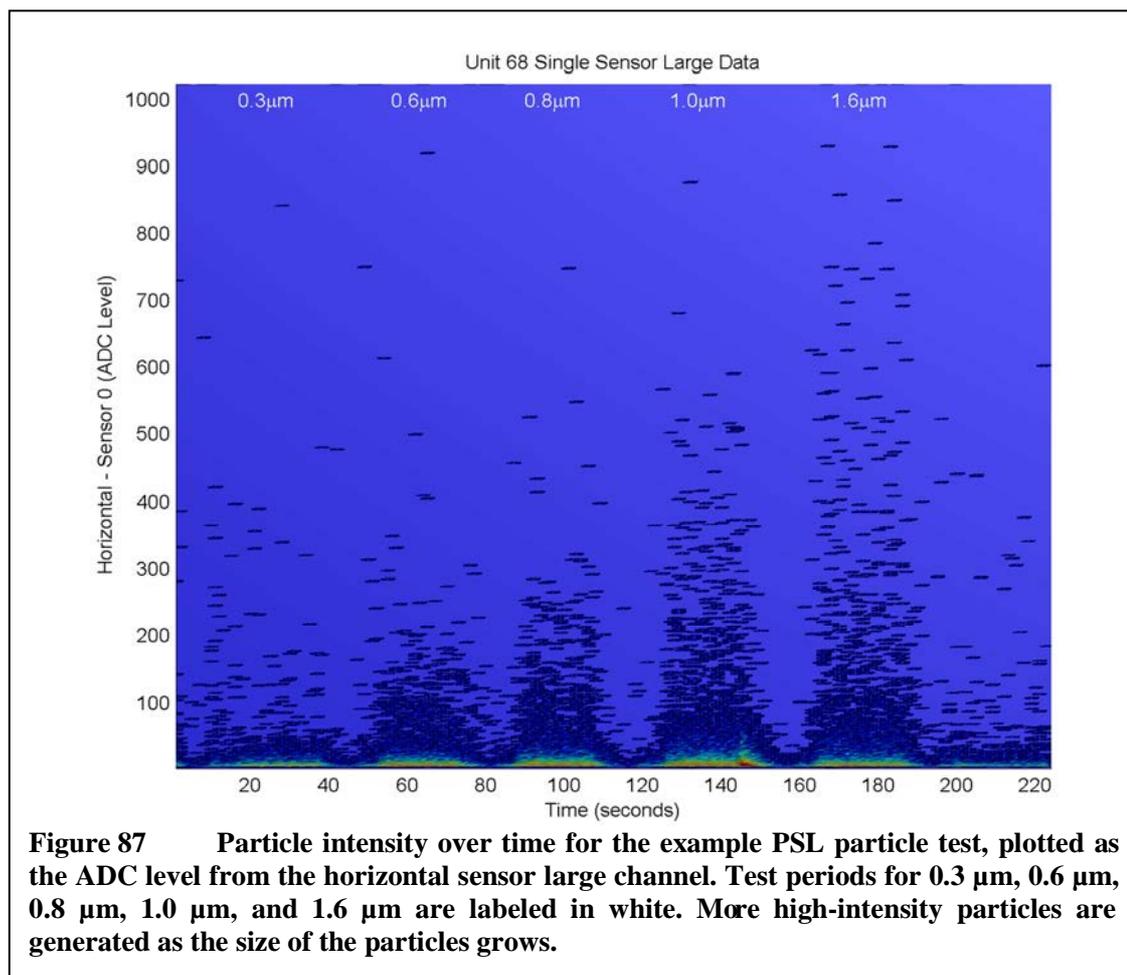


the WAQM OPC with its small sensing area and low flow rate samples only 24 mL/minute for 1µm particles at a typical large particle threshold. This is a difference of more than 1000x the sampling volume. In short, the WAQM OPC has fewer particles with which to work when making a concentration determination. This is especially apparent when the particle concentrations are relatively low as is the case when testing with generated PSL particles.

In general, the WAQM OPC does a fair job of indicating when large particles are present or not, and some indication of the concentration that is similar to that produced by

the PMS device. The major exception is for particles under the large threshold in the 0.6 μm range. The WAQM OPC is falsely reporting a peak in large particles that do not appear on the PMS device. As shown previously in Figure 82, this is due to the crossing of multiple particle sizes by a single threshold. There is a definite peak of particles in this small range that cross the threshold into the large range when they pass through a higher intensity portion of the laser beam.

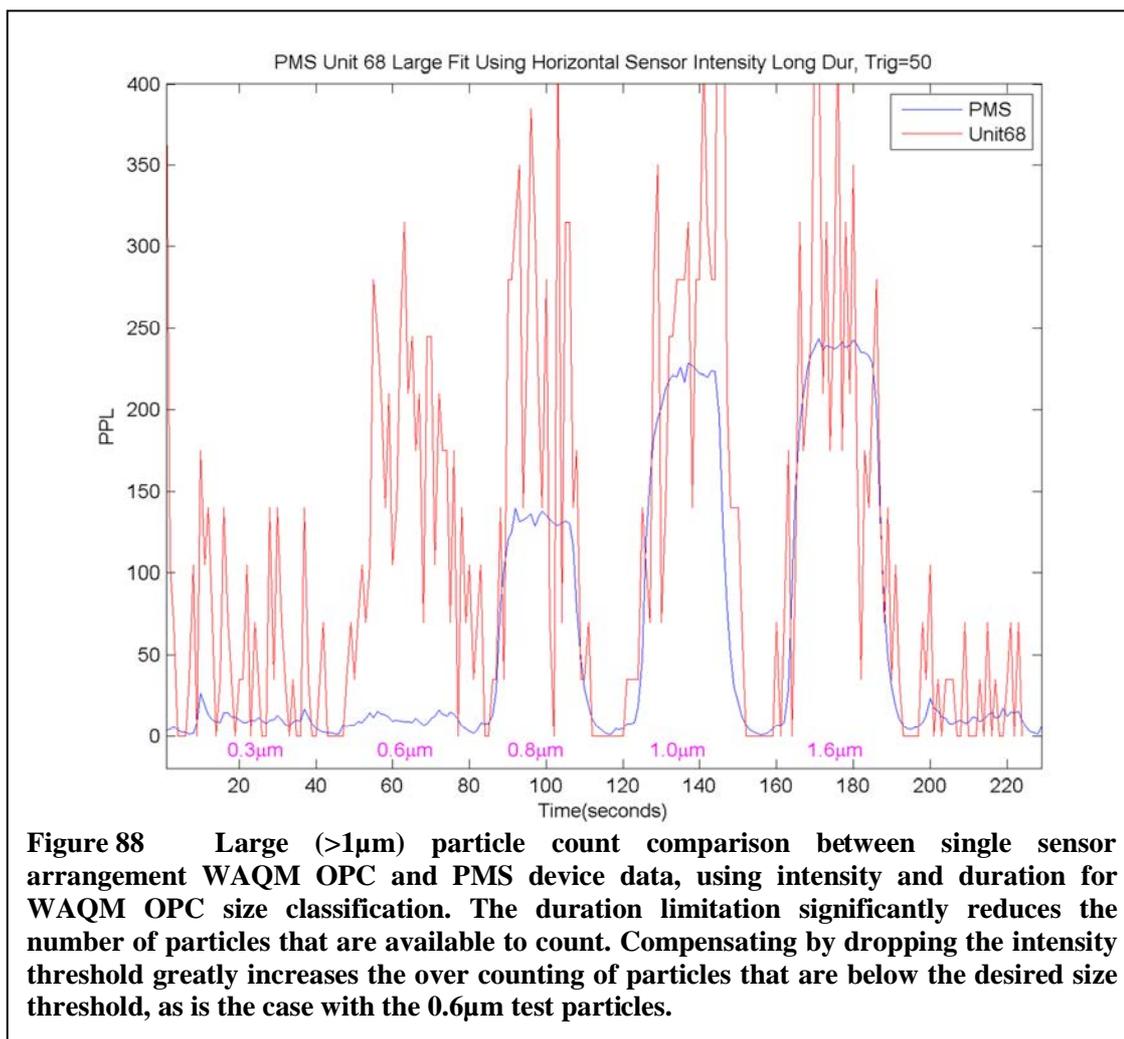
Figure 87 gives some indication of the counting issues with depending on a single sensor and a pulse height determination of particle size. This figure shows the time series of particles for the same test run, with the pulse height in terms of the peak ADC level



seen by the single sensor on the vertical axis. For this unit, the large particle threshold was set at an ADC level of 128. The bulk of the particles seen by the OPC are very small compared to the threshold, yet it is clear that particles in the 0.6 μm group will exceed a level of 128. However, if the intensity is raised to avoid most of the 0.6 μm particles, say to a level of 200, then there will be even fewer of the larger 0.8 μm – 1.6 μm particles to count during the time those groups are present. Picking a threshold that is too high will result in a very noisy data set, with very few particles to count. Too low and peak pulses from smaller particles will leak into the large particle range.

Some effort may be made to discriminate based on pulse duration by combining this with the pulse height. For the single sensor arrangement, this shows minimal promise. The premise would be to allow the lowering of the pulse height threshold to include more particles into the determination, and then filter out smaller particles based on duration. Unfortunately this does not work well for the single sensor arrangement, as there is not enough distinction between the pulse duration for particles of different sizes when passing on either side of the laser beam waist. In other words, a smaller particle passing on the laser source side of the beam waist may appear very similar to a larger particle passing on the opposite side of the waist. A lack of distinction in position makes the duration discrimination method very unreliable.

In the end, reducing the pulse height threshold and filtering on duration increases the pollution from smaller particles into the large range, and reduces the overall number of large particles to count. Figure 88 shows the results from filtering the same test series based on pulse height and pulse duration. Note the increased variation in the data compared to Figure 86, and the increased size of the peak for the 0.6 μm particles.



Clearly the best results for large particle sizing with the single sensor arrangement come from pulse height threshold alone. This means living with some amount of pollution into the large range from smaller particles, but this may be acceptable for some applications.

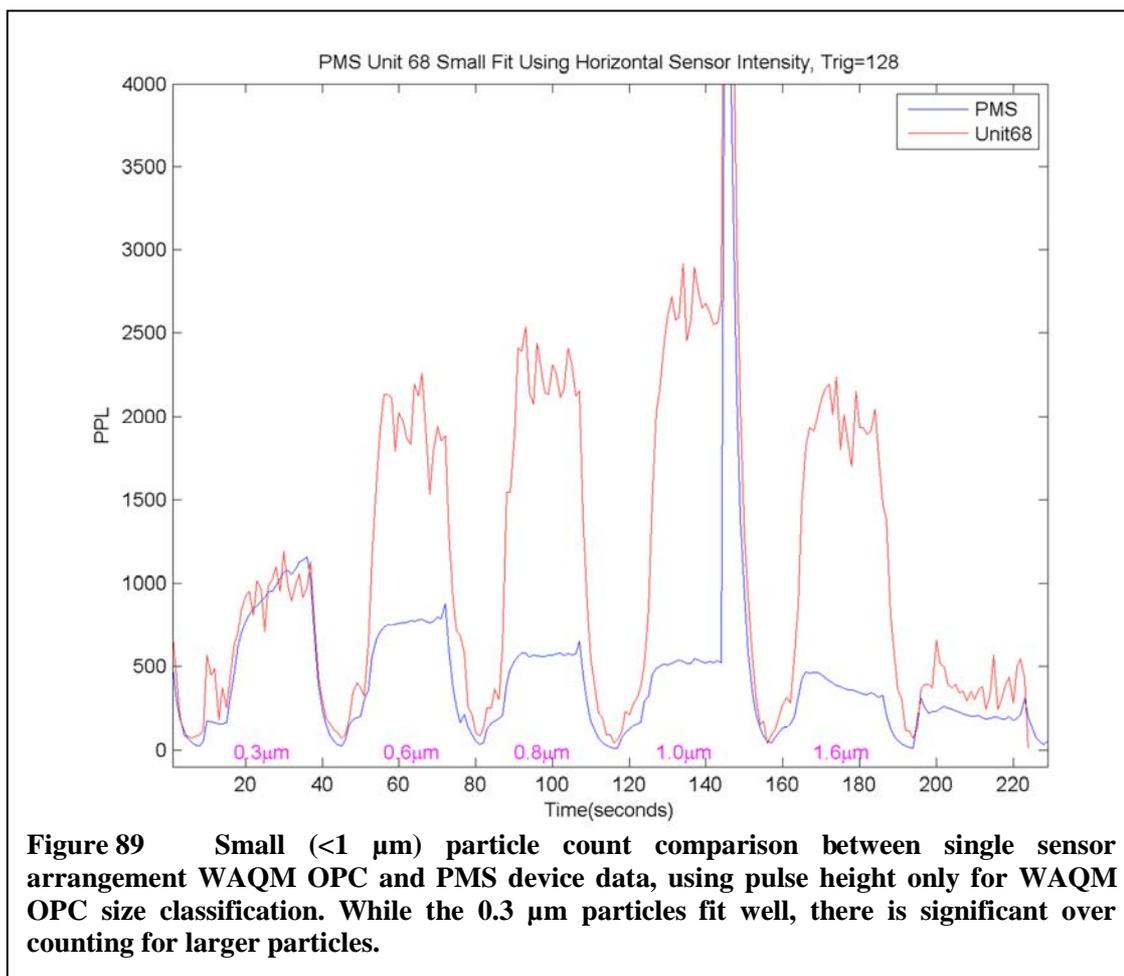
When classifying particles into two size categories, there are four possible outcomes for any given particle. The particle can actually be large and be classified as large, or can actually be small and classified as small. These are the desired outcomes, but particles may also be classified incorrectly. Figure 86 shows the case where particles that are actually small are being classified as large, as is the case with a portion of the 0.6

μm particle group. In this case some smaller particles passing through the more intense areas of the laser beam are passing the threshold for pulse height with their increased scattering irradiance. This problem is somewhat limited in scope, since the cross section of the laser that will produce high pulse heights for small particles is smaller than the cross section that will produce similar pulse heights for large particles. The physics of the situation work in favor of limiting the number of false large particles.

The final situation, where particles that are actually large are classified as small particles is a more serious issue for a test environment. In this case, the cross section of the laser that can produce low pulse heights for large particles may be greater than that for small particles. Given an equal number of particles of different sizes in two separate runs, the group of larger particles is likely to produce a higher small count than the group of small particles. This is because large particles passing through a low intensity portion of the laser may produce a small pulse height, high enough to be detected but not high enough to be classified as a large particle. A small particle passing through the same spot may not be detected at all.

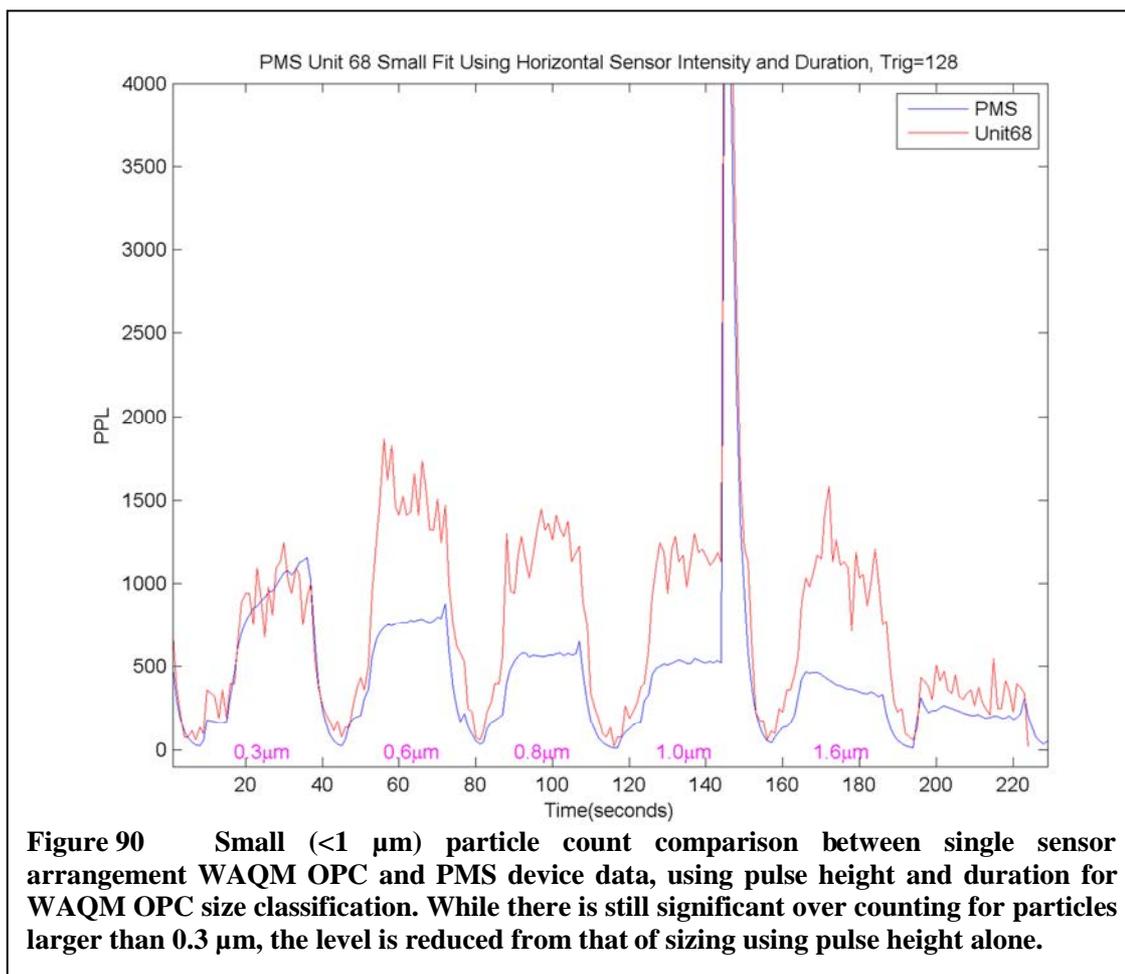
Figure 89 shows the small particle count for the same test run, with sizing based on pulse height alone. The issue just discussed is immediately apparent, with some fairly significant mismatch between the WAQM OPC and the PMS device for particle sizes of $0.6 \mu\text{m}$ and above. For these cases, there are many large particles that are incorrectly classified as small, raising the count well above that reported by the PMS device.

Some improvements may be made by including particle duration data into the particle sizing. Since small particles will typically only be detected when passing through the more intense, narrow portions of the beam, they should not typically produce pulses



that are of long duration. Large particle passing through a less intense, wider portion of the beam may produce pulses of long duration. This provides an opportunity to reject some large particles that might otherwise be classified as small if looking at pulse height alone.

Figure 90 shows small particle counts for the same test run with sizing based on pulse height and pulse duration. The over-counting for particles of size 0.6 μm and greater is clearly reduced from that of sizing based on pulse height alone. Unfortunately, even with this improvement the over-counting is significant, in some cases appearing to account for more than half of the total particle count for some of the test cases. There is



only so much that can be done with the single sensor arrangement to take the focused laser system and produce results similar to systems of more than an order of magnitude greater cost.

That being said, even though the physics of the situation make this issue appear to be serious for reliable small particle counting with the single sensor WAQM OPC, the particulate matter typically measured in an unconstrained environment can trivialize the problem. Particulate matter in ambient air most often follows a lognormal distribution for total particle concentration that peaks at very small particle sizes and falls off with increasing size [3], [34], [79]. A typical indoor or outdoor environment will have small

particle counts that are one or more orders of magnitude higher than the large particle count. This tends to put the issues with over-counting of false small particles into the noise of the actual small particle counts. It is typically only when testing in a laboratory environment that the small particle count can be reduced enough to show the issue as being significant. Even though the false small count may not show up in typical use, it should be well understood that this occurs when using the device in environments where the large particle concentration approaches that of the small particles.

As a side note on the small data, one might notice the peak of high particle count near the end of the 1.0 μm test group. These peaks occur from time to time as the testing system is adjusted, particularly when the system goes between states of particle generation to flushing the chamber. When this happens, some of the air routing lines that transition from inactive to active air movement may contain particulate that has not cleared from a previous run, or has moved into the line from another source when inactive. These data are typically ignored for the purposes of calibrating the WAQM OPC. The calibration system and process will be further explored in subsequent chapters.

Dual Sensor Detection and Sizing

Given the limitations with the single sensor arrangement explored above, the desire is to create something that performs better with respect to particle sizing and at least maintains the same detection performance. This must also be done with minimal cost addition in terms of hardware added to the system, and must not greatly impact the size of the implementation. Adding a second photodiode to the system to create a dual sensor arrangement meets these goals. The photodiode and amplifiers are relatively inexpensive, and a second set can be added without increasing the system cost greatly.

Furthermore, the two sensor design is able to perform sizing without using large channels. This keeps the cost addition down to a photodiode and a single amplifier. Depending on amplifier packaging, this single amplifier may be free. With two photodiodes driving two single amplifier channels, only two analog-to-digital conversion channels are needed as is the case with the single sensor design.

Note that even though the single sensor large/sizing amplifier channel is no longer needed, it has remained inside the prototypes that have been built for this project. This has been done merely as an insurance policy, to allow fallback to the original single sensor design and to allow comparison testing. For example, much of the data taken in the preceding section on the single channel design was actually taken with a unit that implements the orthogonal set of sensors, but only the horizontal sensor was employed with its detection/small and sizing/large channels.

In the simulation chapter above, two different designs were explored in detail: the tandem sensor arrangement and the orthogonal sensor arrangement. In this chapter, the focus for the dual sensor design is strictly on the orthogonal sensor arrangement. While the tandem sensor arrangement was built and tested alongside the orthogonal sensor arrangement, early on it was deemed that the orthogonal arrangement was superior in performance to the tandem arrangement. This was mostly due to the proximity of the horizontal photodiode to the beam waist, allowing the detection of the smallest particles possible. The tandem arrangement ends up with two horizontal photodiodes with the beam waist positioned between them, increasing the distance between the waist and the sensor surfaces relative to the single horizontal sensor of the orthogonal arrangement. The orthogonal arrangement was also chosen due to the positioning of the vertical sensor.

The simulations of the received power ratios appeared to offer better sizing differentiation with this orientation than with the tandem arrangement. This is not to say that there is no merit in the tandem arrangement, or perhaps a variant that positions one of the two tandem sensors directly below the waist. The tandem sensor arrangement was not carried forward mostly to allow a focus on the performance of the more promising orthogonal arrangement.

The orthogonal sensor arrangement has been described in previous sections, and can be seen in Figure 66. This arrangement has one horizontal sensor directly below the laser beam waist, and a second, vertical sensor further from and at a 45 degree angle to the waist. The position of the horizontal sensor is identical to that of the single sensor arrangement, and captures scattered light that is polarized mostly perpendicular to the scattering plane. The vertical sensor is positioned to capture light that is polarized mostly parallel to the scattering plane, though its angle in the forward scattering hemisphere provides it some significant component of perpendicular polarization as well.

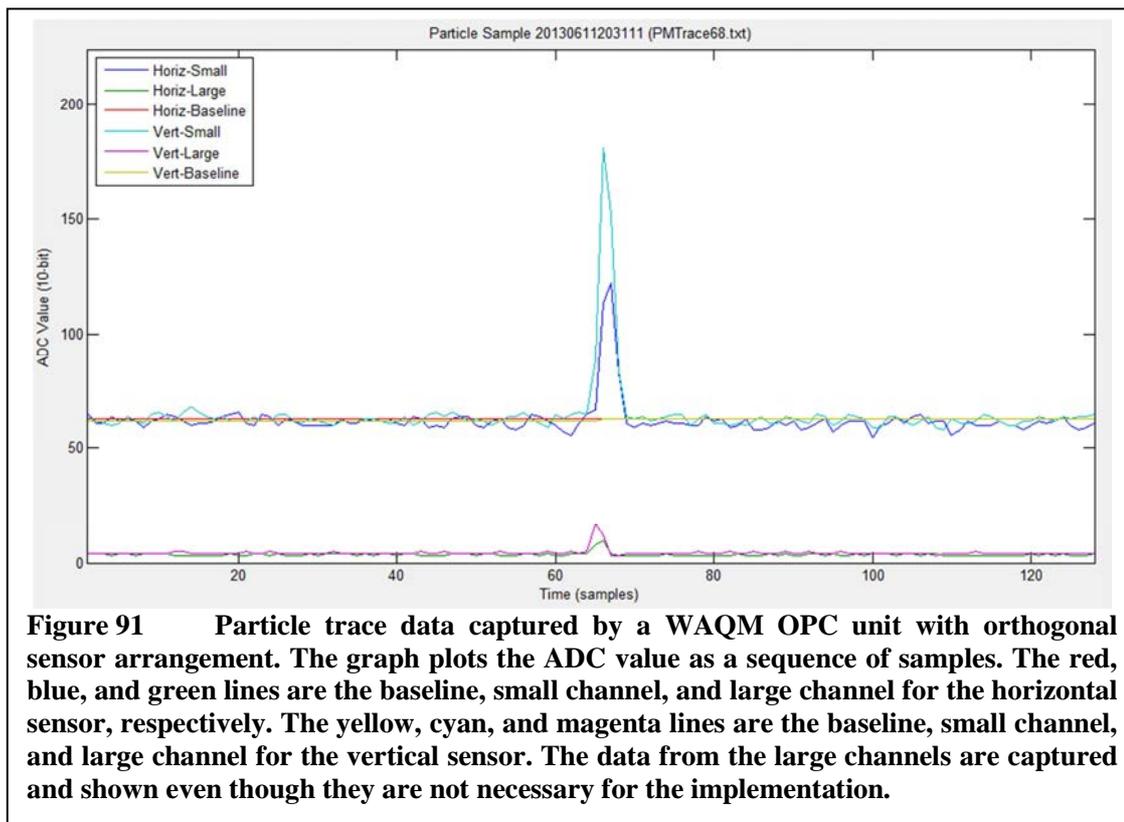
Particle Detection

Particle detection for the orthogonal sensor arrangement is very similar to that for the single sensor arrangement. However, in this case, there are two sensors that can be used to detect the particle. As before, a baseline is created from the ADC levels coming from the detection channels to remove the amplifier bias using an IIR moving average filter as described by Equation 6.1. For this implementation, it is performed for each sensor. Using the baseline as a starting point, data from both of the sensors is parsed looking for deviations above a certain trigger threshold to indicate the start of a particle.

Once the particle start is detected, the data from both sensors are parsed to determine when the pulses fall below the return threshold.

The start of a particle is observed when either the horizontal or vertical channel exceeds the trigger threshold. The return of a particle is observed when the horizontal and vertical channels fall below the return threshold. This is the most permissive form of detection, in that scattered light must only be above the threshold on a single sensor to start processing the particle. This is not necessarily the most efficient method for particle detection, as one may only be interested in particles that trigger based on the horizontal sensor data. If this is the case, triggering on both channels may process many particles that are not of interest, and could be ignored for all practical purposes. Since the early implementation of the dual sensor OPC has been targeted at exploration of the space, the most permissive form of detection was used to capture the most particles possible. As will be seen, for the final sizing algorithm, most of the particles that are detected by only the vertical sensor will be discarded. The vertical triggering could likely be discarded for an optimized design.

An example of the raw data seen by the orthogonal arrangement is shown in Figure 91. These are data taken by the WAQM OPC directly from the particle detection algorithm and stored on the internal SD card for display with a MATLAB tool. The data shows the response to a particle passing through the laser and producing a medium-sized pulse. The vertical range has been reduced by the display tool to zoom in on the particle. The baselines for the horizontal and vertical channels are shown in red and yellow, respectively. Note that the two baselines nearly overlap, as would be hoped for two independent channels measuring the same light levels. The baselines are offset as with



the single sensor arrangement due to amplifier biasing. The horizontal and vertical small channels are shown in blue and cyan, respectively. This particle shows a more intense response on the vertical sensor than the horizontal sensor. This is mirrored in the horizontal and vertical large channels shown in green and magenta, respectively. Note that the large channels are not necessarily needed for detection and sizing, but are included since they were available in the prototype orthogonal sensor design.

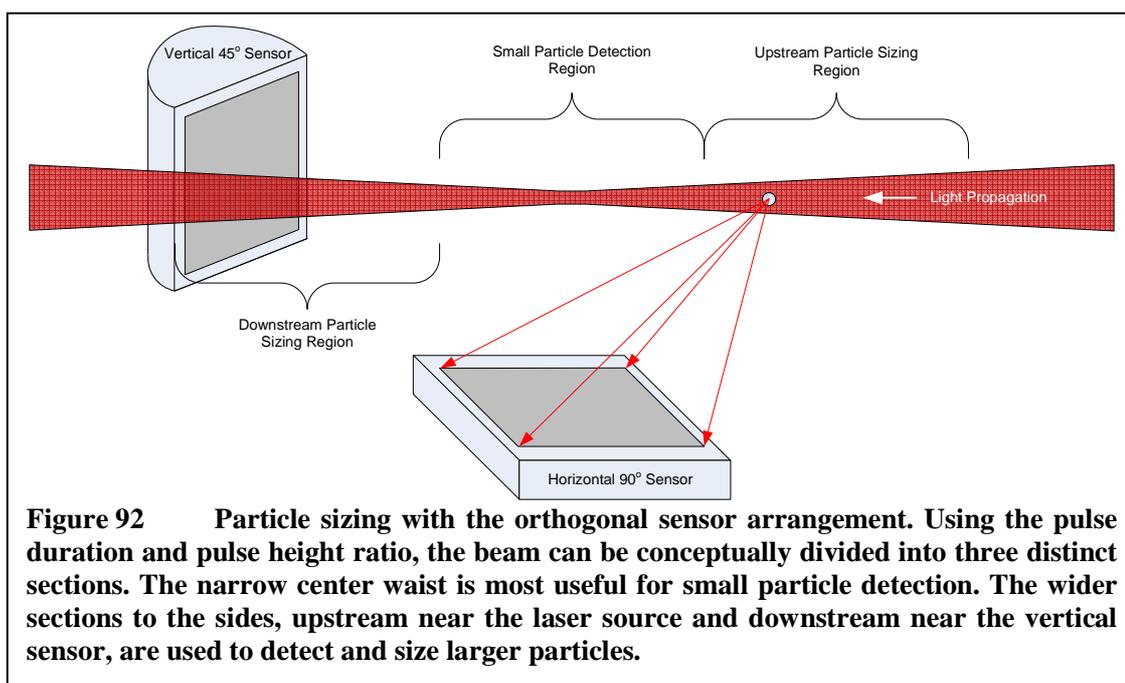
Once the two channels fall below the return levels, further processing is performed on the particle. The particle is typically kept or rejected based on a minimum small channel threshold. This may be implemented on one or both of the channels in either an AND or an OR type of function. With permissiveness in mind, the current implementation of the WAQM OPC will accept particles based on a minimum threshold

being achieved on either the horizontal or vertical channel. At this point, the maximum ADC level reached and pulse duration for both channels is stored and passed on to the sizing algorithm.

Particle Sizing

The orthogonal sensor arrangement provides a new dimension in the data available for sizing particles. With the single sensor arrangement, issues were encountered in using pulse duration effectively due to the inability to determine through which part of the beam the particle passed.

Figure 92 shows the concept of using different portions of the laser beam as it passes through the air flow channel. In this case, the beam is conceptually split into three effective regions. The central region is the highest intensity portion of the laser containing the beam waist. This portion of the beam may be thought of as the small



particle detection region as it is most effective at producing detectable pulses from the smallest particles due to the high intensity and close proximity to the horizontal sensor. The regions on either side of the center are more useful for the sizing of particles, though they are also used for the detection of larger particles. The region closest to the laser source is referred to as the “upstream” particle sizing region, and the region furthest from source is referred to as the “downstream” particle sizing region.

With the two sensors, a ratio of the horizontal received irradiance to the vertical received irradiance may be created. As was discussed in the simulation chapter, this ratio can provide some information as to the location of the particle in the flow channel relative to the laser beam waist. As shown in Figure 70, the ratio between the horizontal and vertical sensors for most of the larger particles is greater than one on the upstream side of the waist, and less than one on the downstream side. This allows the determination of position relative to the beam waist for larger particles, which will be useful in determining their sizes.

Small particles in the range of $0.3\ \mu\text{m}$ also have horizontal to vertical irradiance ratios larger than one, and have this ratio on both sides of the beam waist. The weak scattering from these smallest particles helps to keep them from being confused with larger particles. The scattering irradiance is typically too weak from these particles when passing through the portions of the laser outside the central region to be detected by either sensor. That leaves these particles producing detectable pulses only when passing through this smaller central region.

Large particles passing through the central region of the beam may produce very high intensity pulses, as was seen with the single sensor arrangement. For the single

sensor arrangement, the low-gain large channel was used to size pulses of high intensity since they would cause the high-gain small channel to saturate. For the orthogonal arrangement, the large particles are taken from the upstream and downstream regions of the laser beam, and these high intensity central region pulses are rejected. This allows the use of just the high-gain small channel for sizing. Any pulses that saturate the small channel are rejected, since the ratio and duration of the small channel are used to size the particles. The rationale for taking this tact is to avoid the peak intensities in sizing particles, which has a great deal of variation in the central region of the laser beam. As seen in Figure 82, attempting to draw a threshold in the central region leads to more pollution between the particle sizes.

Example Data

The two sensors in the orthogonal arrangement allow for some new representations of the particle data. The six plots in Figure 93 show one of these representations, with the intensity in terms of the ADC level displayed on the two axes: the horizontal axis for the horizontal sensor intensity, and the vertical axis for the vertical sensor intensity. The particle count at each position is represented as a color indicated in the bar on the right of each plot. Of the six plots, the left column contains simulated data for the orthogonal sensor arrangement and the right column contains actual data taken from a prototype unit. The top, middle, and bottom rows are 0.3 μm , 0.8 μm , and 1.6 μm PSL particles, respectively. The simulated data are of course an ideal spherical PSL particle. The actual data are taken from calibration PSL spheres mixed in solution with distilled water and atomized with a Collison nebulizer.

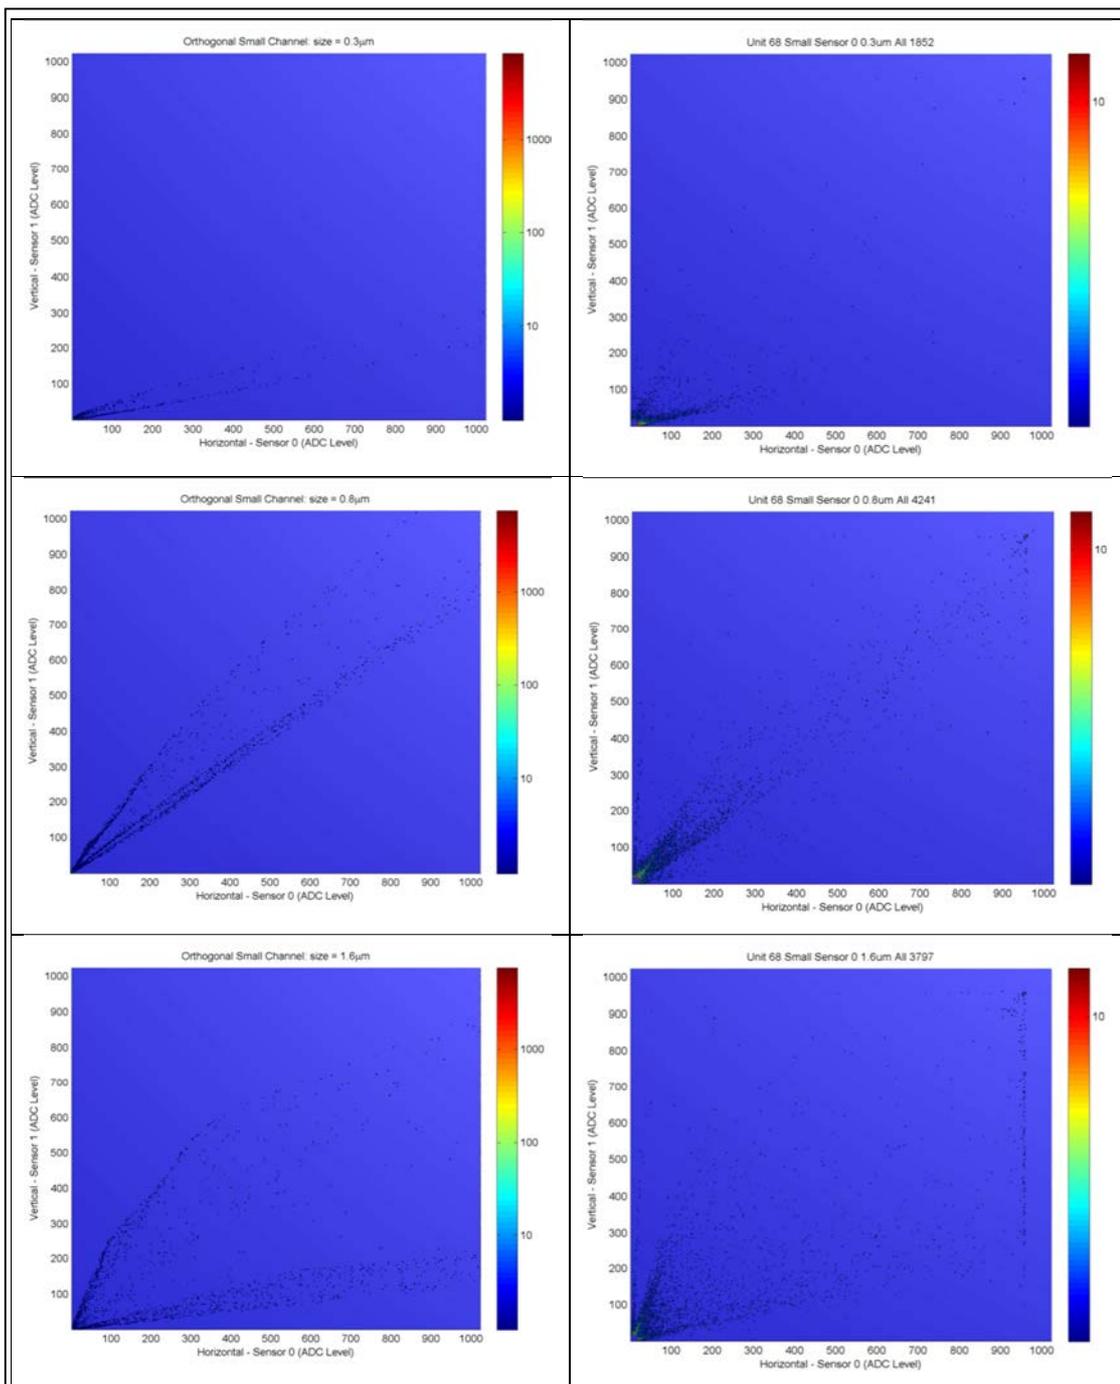


Figure 93 Comparison of orthogonal sensor arrangement simulated (left column) and experimental (right column) results for 0.3 μm (top row), 0.8 μm (middle row), and 1.6 μm (bottom row) PSL particles. The horizontal and vertical axes are the horizontal and vertical sensor ADC levels, respectively. The simulations assume a monodisperse aerosol, which is not possible to achieve with the actual experiments. The general trends in distribution do match according to particle size.

Looking at the data overall, one can see that the general trends match relatively well between the simulated and actual data. The 0.3 μm particles show a strong trend to have higher intensity on the horizontal sensor than the vertical, or a high horizontal to vertical intensity ratio. This would be expected from the plot of intensity ratios in Figure 70. The actual test data do show a grouping of particles with a low horizontal to vertical intensity ratio that does not appear on the simulated data. As might be expected, the actual testing is not ideal, and as will be seen, contains some particles that register as larger sizes than 0.3 μm , even with the PMS OPC.

The middle row shows the results for 0.8 μm particles, which have a trend towards a horizontal to vertical intensity ratio of one. There are two distinct groupings with different ratios in the simulated and actual data. These correspond to particles passing through the upstream and downstream portions of the beam. The upstream particles generate ratios that are larger than those passing through the downstream portion of the beam. Upstream particles generally tend to cause a higher intensity on the horizontal sensor than the vertical, since they are significantly closer to the horizontal sensor than the vertical.

The bottom row shows the results from 1.6 μm particles, which show greater diversity between particles passing through the upstream and downstream portions of the beam. The two groupings of particles, one with high horizontal to vertical intensity ratio and the other with low intensity ratio, correspond to the particles passing through the upstream and downstream portions of the beam. As with the other particle sizes in the figure, the actual test data show a more diverse set of intensity ratios than the simulated

data. This is to be expected as the actual test particulate matter is by no means comprised of ideal particles with 1.6 μm diameters.

Note that the saturation point is relatively easy to see in the actual data for the 1.6 μm particles. This appears as the strong line of particles running along the right edge of the plot, leaving a band that contains no particles along the edge. The line also exists along the top, though it is less easy to detect without more particle data. The actual data are plotted after the baseline has been subtracted from the particle intensity data. This empty band results from the range that is above the saturation point of the amplifiers in the system.

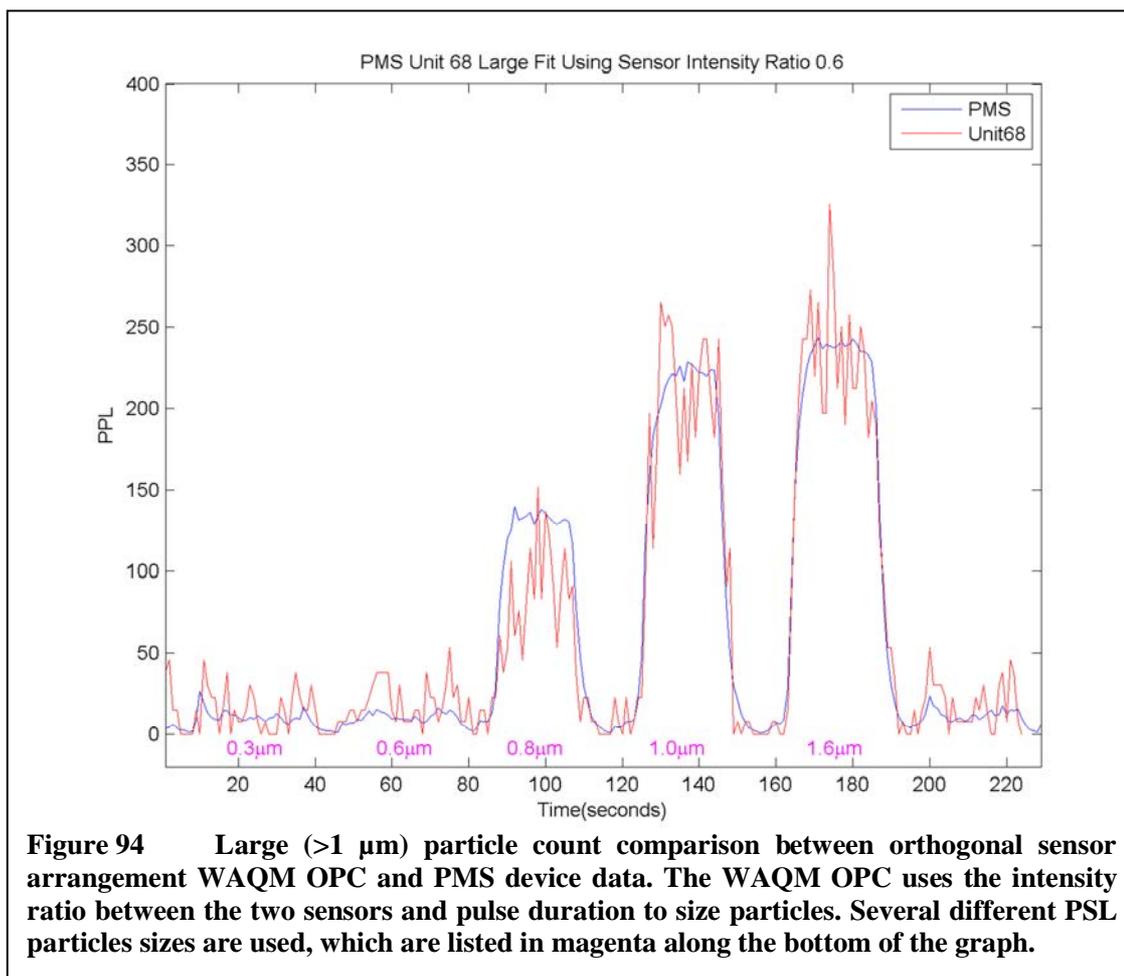
Another line of particles also appears along the left edge of the plot for the 0.8 μm and 1.6 μm particles. These are clearly particles that have relatively high intensity at the vertical sensor, but almost no intensity at the horizontal sensor. These particles are not predicted by the simulation, and may be showing some secondary effect inherent in the physical arrangement of the laser and photodiode, or the electronic system of photodiode and amplifiers, that is not taken into account in the simulation. These particles of very low horizontal to vertical intensity ratio are discarded for the purposes of particle counting and sizing.

When considering the sizing of particles based on this data, the small particles may be taken from the area where there is a high horizontal to vertical intensity ratio, as with the 0.3 μm particles. Larger particles may be taken from the low intensity ratio area. Unlike the case with the single sensor arrangement, pollution from small particles into the large particle category is not much of an issue, since small particles do not generate small horizontal to vertical intensity ratios. However, pollution from large particles into the

small category will still be a problem, as is evident by the large band of high intensity ratio particles generated in the 1.6 μm data. To decide whether a particle is large or small, a linear ratio threshold may be used to divide the intensity plot into two portions. The high horizontal to vertical intensity ratio will be small particles, and the low ratio will be large particles.

Classifying particles as large based only on the intensity ratio will only work for certain ranges of particle sizes. Smaller particles in the 0.6 μm range may also produce pulses that have a low horizontal to vertical intensity ratio when passing through the downstream portion of the laser beam. Based on the intensity ratio alone, these are indistinguishable from larger particles. If the large threshold is desired to be at a size of 1.0 μm , something else must be done to distinguish between particles of this size and those that are smaller but still produce a low intensity ratio. This is where the pulse duration can help with differentiation. The duration is more effective as a differentiator when isolated to one side of the laser beam only. Using the intensity ratio to first decide which side of the laser is generating the pulse, and then the duration to separate sizes, a relatively good classification of particles over 1.0 μm can be created. Pulses of long duration are captured as large particles, and pulses of shorter duration are small.

Using this method of classification for large particles, the data from the WAQM OPC can be compared to those from the PMS OPC for the test set of particles. Figure 94 shows the large particle counts with the large threshold set at 1.0 μm . As before, the PSL particle sizes for each grouping are listed in magenta along the bottom of the plot. Comparing to the data from the single sensor arrangement shown in Figure 86, the data generated by the orthogonal arrangement have much less variation. The data also do not

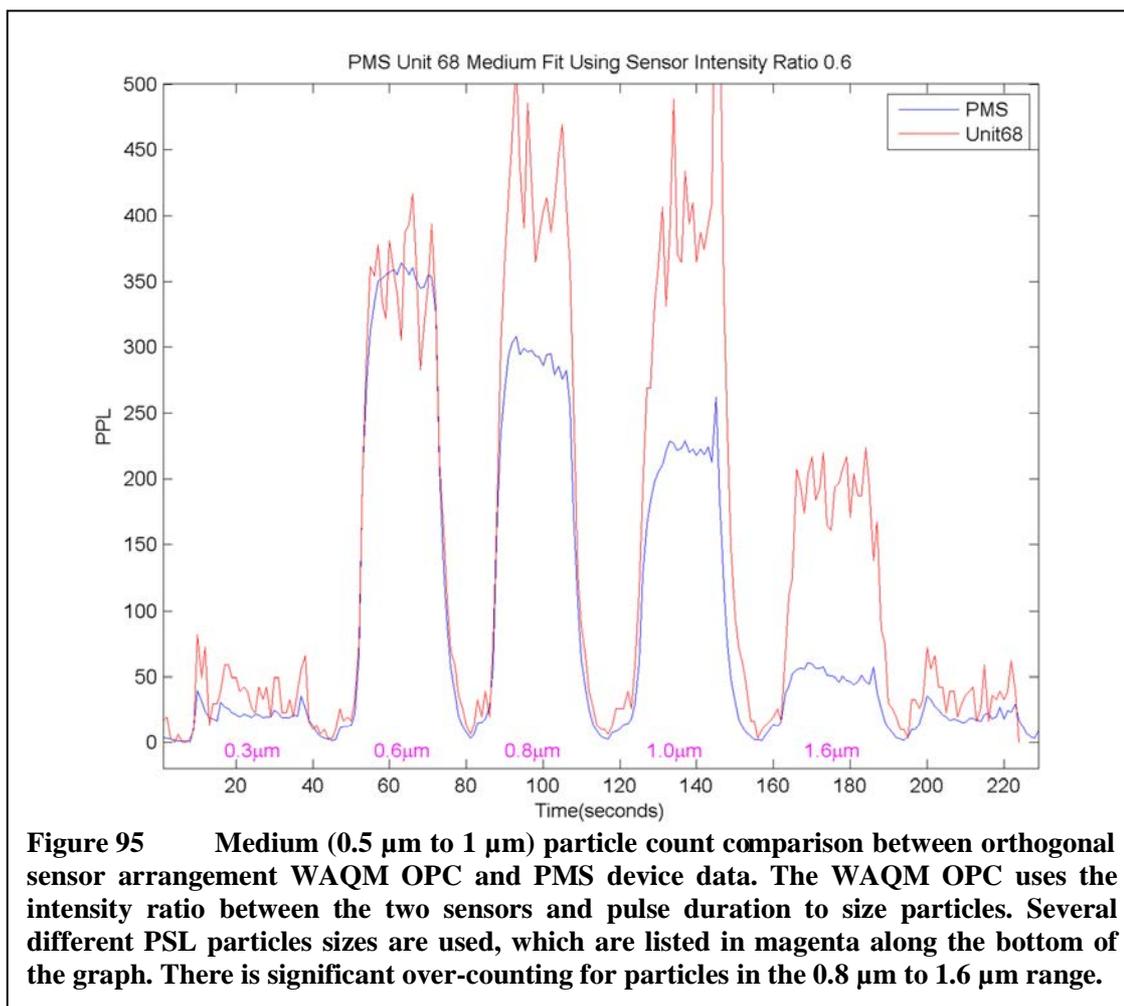


have nearly the level of over-counting in the smaller particle groupings, which was a significant issue for 0.6 µm in the single sensor arrangement. There is some undercounting in the 0.8 µm group, which is somewhat of a balance between over-counting in smaller particles, and undercounting in larger particles. This large count is a function of the ratio threshold, and moving it to higher horizontal to vertical intensity ratios leads to an increase in over-counting for the smaller particle groupings. Moving the ratio to lower values increases the under counting for larger particles. The figure shows a balance that is somewhat weighted toward a high value, causing some undercounting in the 0.8 µm group with a relatively flat response for the 0.6 µm group.

Another likely contributor to the lower count for the 0.8 μm group is the inaccuracy of applying a single scaling value to particles over a range of sizes. As previously described, particles near the low end of the size range will see a smaller effective detection area than particles at the large end of the range. The difference in sampling area will translate to count differences that cannot be corrected exactly by a single scaling value.

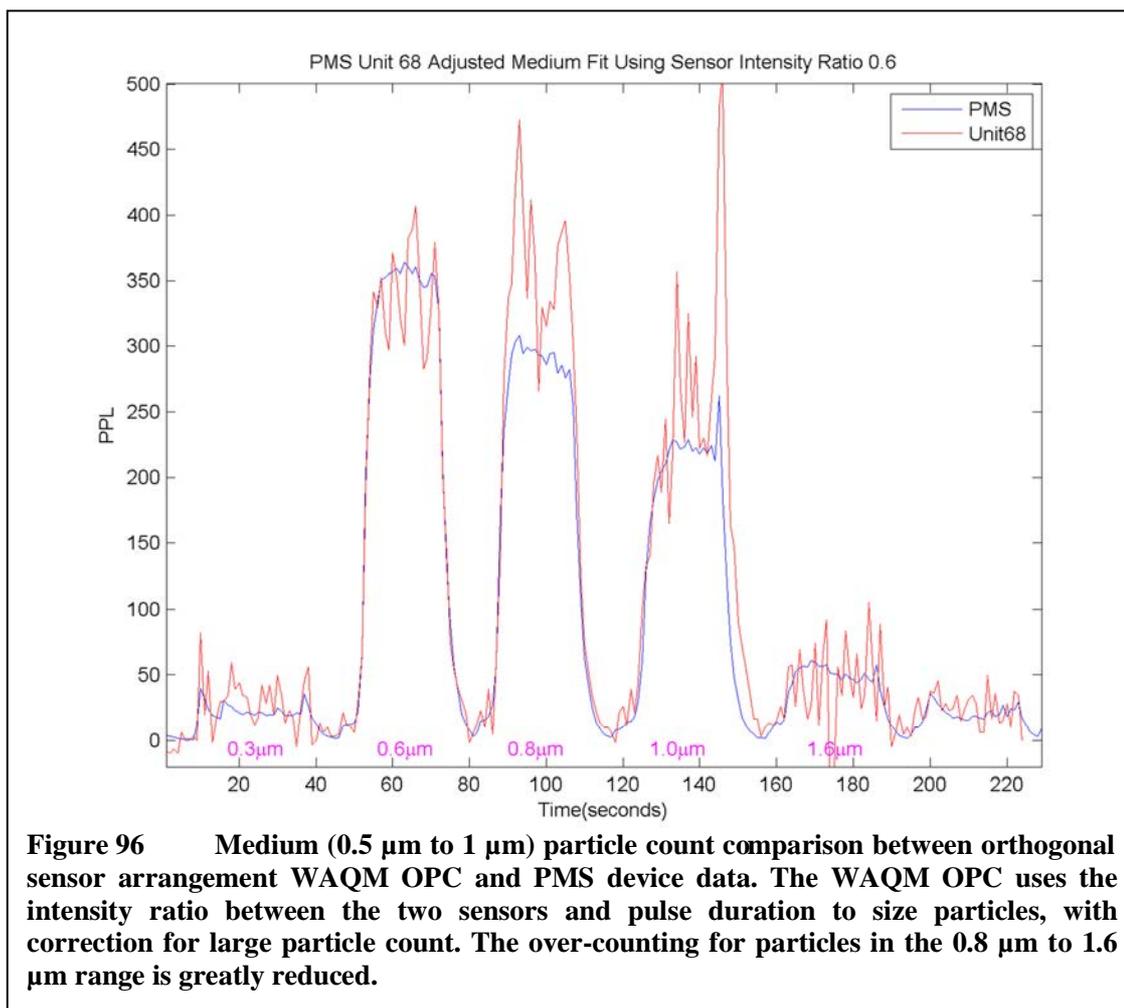
An added benefit of this classification method is that the particles of shorter duration can be classified as medium size particles in the range of 0.5 μm -1.0 μm . These particles are of relatively short duration, but create pulses with low horizontal to vertical intensity ratios. Figure 95 shows the results of taking those particles that are of shorter duration, placing them into the medium category, and plotting them alongside the same range of particles reported by the PMS device.

It is immediately apparent that there is significant over-counting for groups of particles that are larger than 0.6 μm . Using the longer pulse durations as a limit is really only effective in one direction. Medium-sized particles may be unable to generate long pulses, but large particles may still generate shorter pulses based on the intersection path through the laser beam. These shorter pulses will be rejected from the large particle count, but incorrectly included into the medium count. There appears to be no effective way to correct for this phenomenon on a particle-by-particle basis, but it is possible to perform a correction on the aggregate of counted particles over a reporting period. Since the particle counts are typically taken for a period of time, say 60 seconds, and then reported as a concentration, there is a chance to use the various counts to perform corrections prior to reporting the concentration.



In the case of the medium particles, a scaled subtraction of the large particle count from the medium count can be performed. The results from performing this type of correction on the medium particle count are shown in Figure 96. The improvement is significant, with the over-counting greatly reduced for particles larger than 0.6 μm . The inability of medium particles to generate longer pulses can be used to treat the large particles as a known quantity, and use it to correct the unknown quantity of medium particles to a more accurate value.

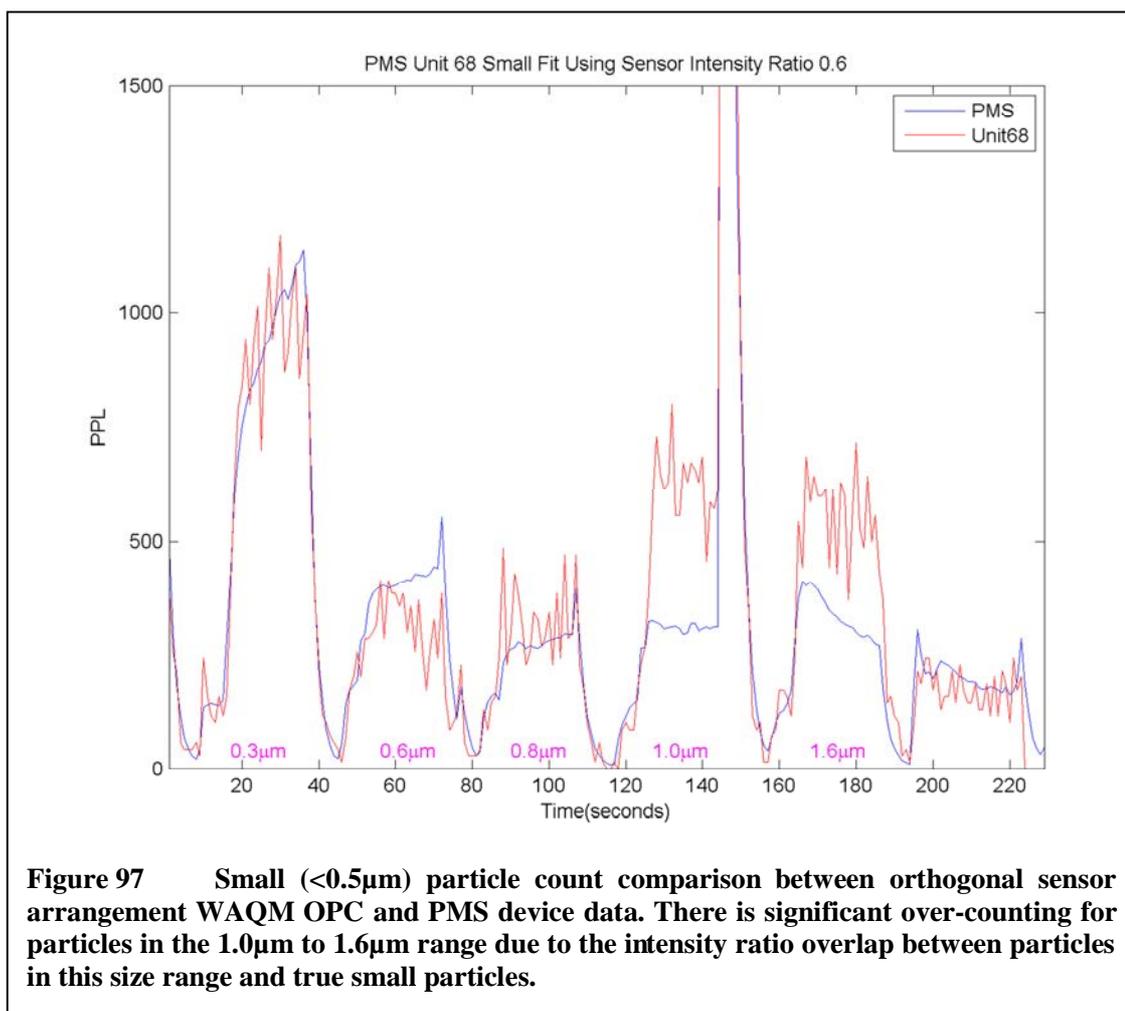
Finally, the remaining particles to examine are the smallest particles in the range from 0.3 μm to 0.5 μm . These particles are those with high horizontal to vertical intensity



ratios and relatively short pulse durations. Unfortunately, there is some overlapping with larger particle sizes in intensity ratio and duration that cannot be avoided. This can be seen in Figure 93 when comparing the data from the 0.3 μm and 1.6 μm particles. The smaller particle range has a high intensity ratio grouping without a low intensity ratio grouping. The larger particles have both groupings, and these larger particles will overlap with the small particles, causing over-counting for small particles when they are present. Some filtering can be done based on pulse height and pulse duration, since small particles will be limited in the maximum height and duration that can be produced. This cannot

remove all large particle size over-counting, since large particles passing through weaker portions of the laser beam may still produce short pulse heights and durations.

Figure 97 shows the results from the small particle classification. This has been created by taking those particles with a large horizontal to vertical intensity ratio, and then applying pulse height and pulse duration filters in an attempt to reduce larger particles. The pulse height filter removes high intensity particles from the group, since those particles with a high scattering irradiance will predominantly be large particles. Similarly, the pulse duration filter removes long duration particles that are generated only



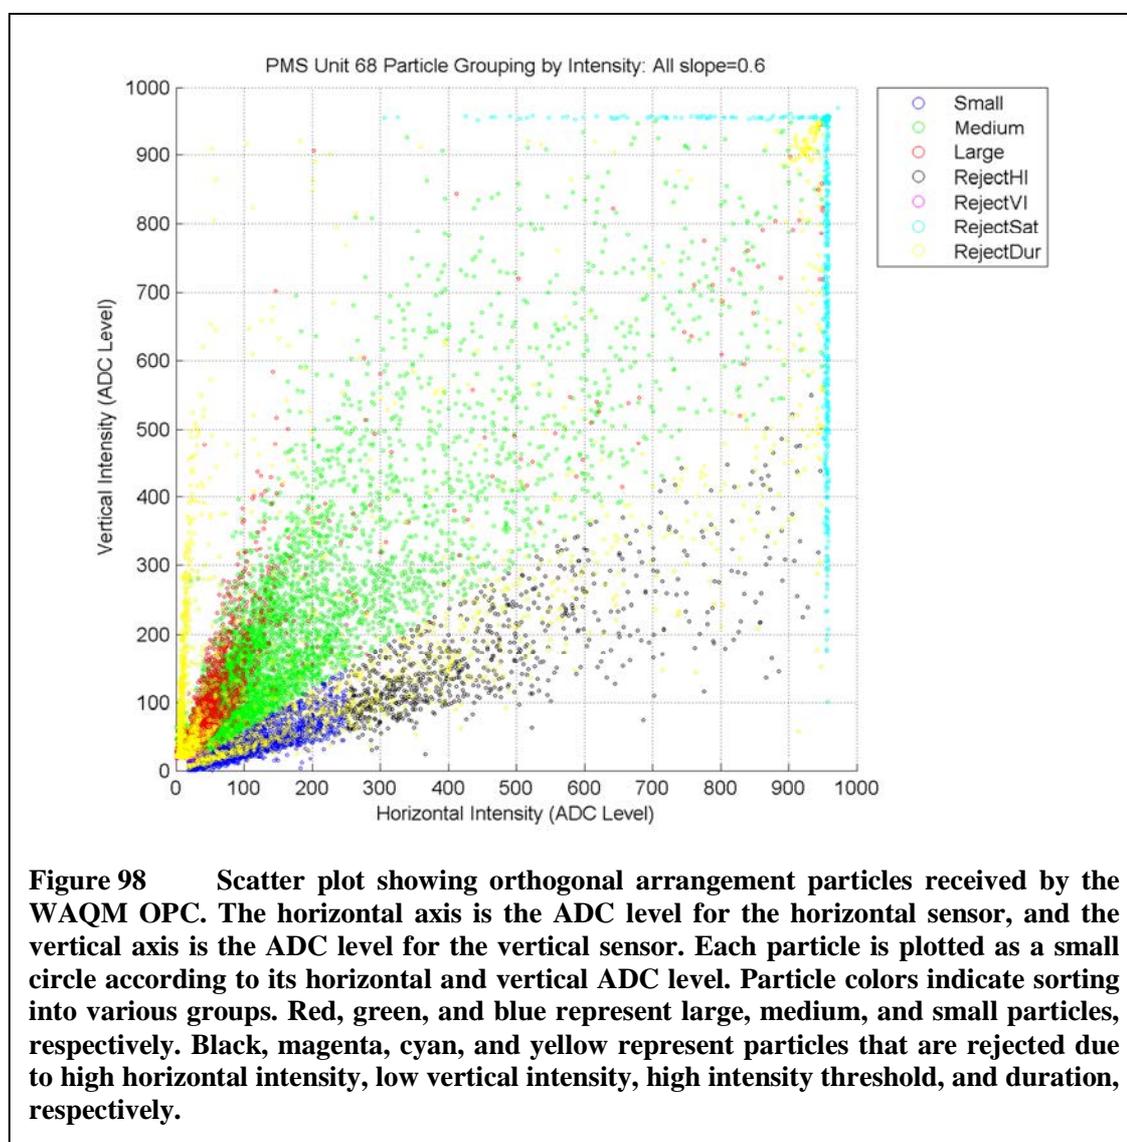
by larger particles, since small particles cannot produce detectable irradiance from the wider portions of the laser beam. There is still some remaining over-counting in the 1.0 μm and 1.6 μm particle groups where there is a significant high intensity ratio group for these particle sizes that cannot be completely distinguished from truly small particles, but performance is still much improved from the single sensor arrangement.

It is possible that the same method of correcting the medium particle count using the scaled subtraction of large particles could be applied. However, there is a range of particles near 1.0 μm that do not produce the same high intensity ratio grouping that causes the over-counting. This can be seen in the 0.8 μm particle group, where the small count is close to correct. In Figure 94, it is clear that there is a significant large count in this group, and subtracting it from the small count would result in undercounting for this range of particle sizes. It should be noted that this may be due to slightly incorrect sizing by the PMS device. It is interesting that the 0.8 μm particles do produce a significant large count on the PMS device, but fail to pollute the small count on the WAQM OPC. This discrepancy could indicate a possible over-sizing of these particles by the PMS device. This might be verified in the future if other more accurate measurement methods become available to the WAQM OPC development effort.

As with the single sensor arrangement, there is some pollution of the small count caused by large particles. While this is undesirable, some consolation may be had by considering that typical aerosols measured by the WAQM OPC will have a lognormal size distribution. This typical distribution will make the small count much higher than the typical large count. As such, the resulting pollution into the small count from large

particles passing through low-intensity portions of the laser beam should be a minor factor in the accuracy of the small particulate concentration.

To visualize the application of the sizing algorithm to the data, the particles may be displayed on scatter plot with colorization to indicate how sizing has been selected. Figure 98 shows such a scatter plot, with the horizontal sensor intensity on the horizontal axis, and the vertical sensor intensity on the vertical axis. The particle data plotted are from the entire test run shown in the previous figures, including all particle groups from



0.3 μm up to 1.6 μm and all spaces in between. Particles categorized as small are shown in blue, which lie predominantly in a sector with high horizontal to vertical intensity ratio. Particles categorized as medium are green, and lie in the space with low horizontal to vertical intensity ratio. Large particles are red, and are also in this area of low horizontal to vertical intensity ratio. For this particular unit, the small and medium/large particles are separated by a line with slope of 0.6 drawn from the origin.

Several categories of rejected particles are also shown. Particles filtered from the small count due to high intensity are shown as black. These lie to the right end of the horizontal intensity scale. These particles are considered to have too high a horizontal sensor intensity to be true small particles, as they are mostly due to particles with size greater than 1.0 μm passing through the laser beam near to the laser source. Some particles in the small range from 0.3 μm to 0.5 μm are also lost in this filtering, but the number is confined to a small percentage of the particles passing through the most intense portion of the laser beam.

Particles filtered from the large and medium counts based on vertical sensor intensity are shown in magenta. Most of these filtered particles have a low horizontal to vertical sensor intensity ratio, but also have a low vertical intensity. Particles in this class include a number of particles that are on the large end of the small size range, but not large enough to have high vertical intensity. These are rejected to remove pollution from the small range into the medium and large categories.

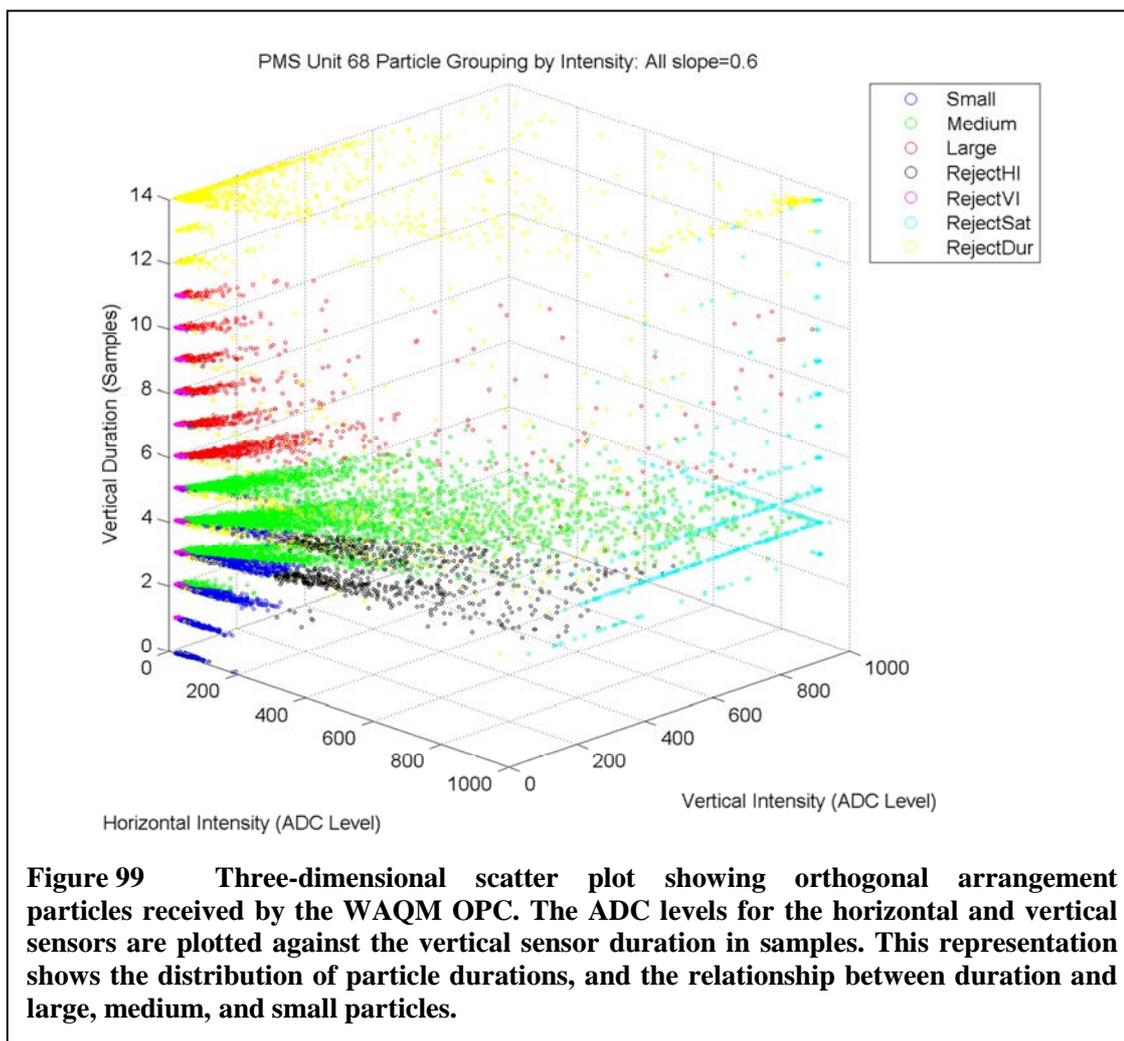
Particles that saturate either sensor channel are rejected and show as cyan in the plot. Clearly particles that saturate both sensor channels will be difficult to size correctly. Even though they are most likely to be large particles, these points could also be due to

smaller particles passing through the most intense portion of the beam. Particles saturating the horizontal sensor channel but not the vertical are also most likely large particles, but could again be small particles passing through the beam near the waist. Particles saturating the vertical sensor channel but not the horizontal are likely either large or medium particles, but differentiating between the two is more difficult when the absolute intensity is lost. In general, only a small portion of the particles produce scattering that is intense enough to saturate either of the sensor channels, so rejecting them is not of much consequence.

Particles rejected due to long duration are shown as yellow, and are filtered from the small and large groupings. These are scattered across the plot, but also have two significant clusters in the low-intensity and high-intensity corners of the plot. These particles are most likely generated by large particles passing through the laser beam near the edge of the flow channel. Particles passing through these portions of the beam are rejected for sizing as they may exhibit multiple peaks due to the shape of the beam and their slow transition through it. More will be discussed regarding these “misfit” particles towards the end of this chapter. This group also includes particles with very low horizontal to vertical intensity ratio, which are together in a narrow strip along the left edge of the plot. These particles are interesting since the simulation does not predict their appearance. Further investigation may reveal the source of this group, but since the count is small and the duration is relatively long, they are rejected at present.

Looking more closely at the distribution of particle durations relative to this plot requires a three-dimensional scattering. Figure 99 shows the scatter plot of Figure 98 turned at an angle, with the vertical sensor duration in number of samples plotted on the

third axis. The vertical duration is used instead of the horizontal duration, since this is the metric used in distinguishing between large and medium particles. The vertical duration is used since the large and medium particles that are counted pass through the downstream side of the laser beam, and are relatively near to the vertical sensor. These particles also are scattering in the forward direction toward the vertical sensor, and backscattering towards the horizontal sensor. Since forward scattering is typically the stronger direction of scattering for larger particles, duration from the vertical sensor is used as the more reliable sizing metric for these particles.



Looking at the data in Figure 99, the bulk of the particles with relatively high intensity have vertical sensor durations of at least 3 samples. Some are smaller, and even have vertical durations of zero, meaning that they are only detected by the horizontal sensor. This lack of high intensities at low durations is mostly a function of the circuit response, since the pulse height by its nature will impact the pulse duration as a minimum rise and fall time for the particle is inherent in the circuit response. The separation between large and medium particles is set to a duration of 6 samples, as can be seen by the change in color of the low intensity ratio particles from green to red. The maximum duration accepted is 11 samples; particles colored yellow are rejected as too long in

duration. Note that most of the small particles in blue are of relatively short vertical sensor duration, whereas the large particles in red have durations that extend up many more samples than this.

Figure 100 shows scatter plots for multiple different particle groups from the test series. The top, middle, and bottom rows are 0.3 μm , 0.8 μm , and 1.6 μm PSL particles, respectively. The left column shows the two-dimensional plot with the horizontal and vertical sensor intensities on the horizontal and vertical axes, respectively. The right column shows the same data, but in three dimensions with the vertical sensor duration shown on the third axis. For the 0.3 μm set in the top row, the main group of particles falls in the small area with relatively short duration. These particles are colored blue in the plot. The 0.8 μm particles in the middle row are mostly in the medium sized area with relatively short duration. These particles are colored green in the plot. The 1.6 μm particles in the bottom row have more particles in the large area with relatively long duration. These particles are colored red. Note that the 1.6 μm particle group also has a fair number of particles that are categorized as small. These are particles contributing to the over-counting show in the small concentration plot of Figure 97.

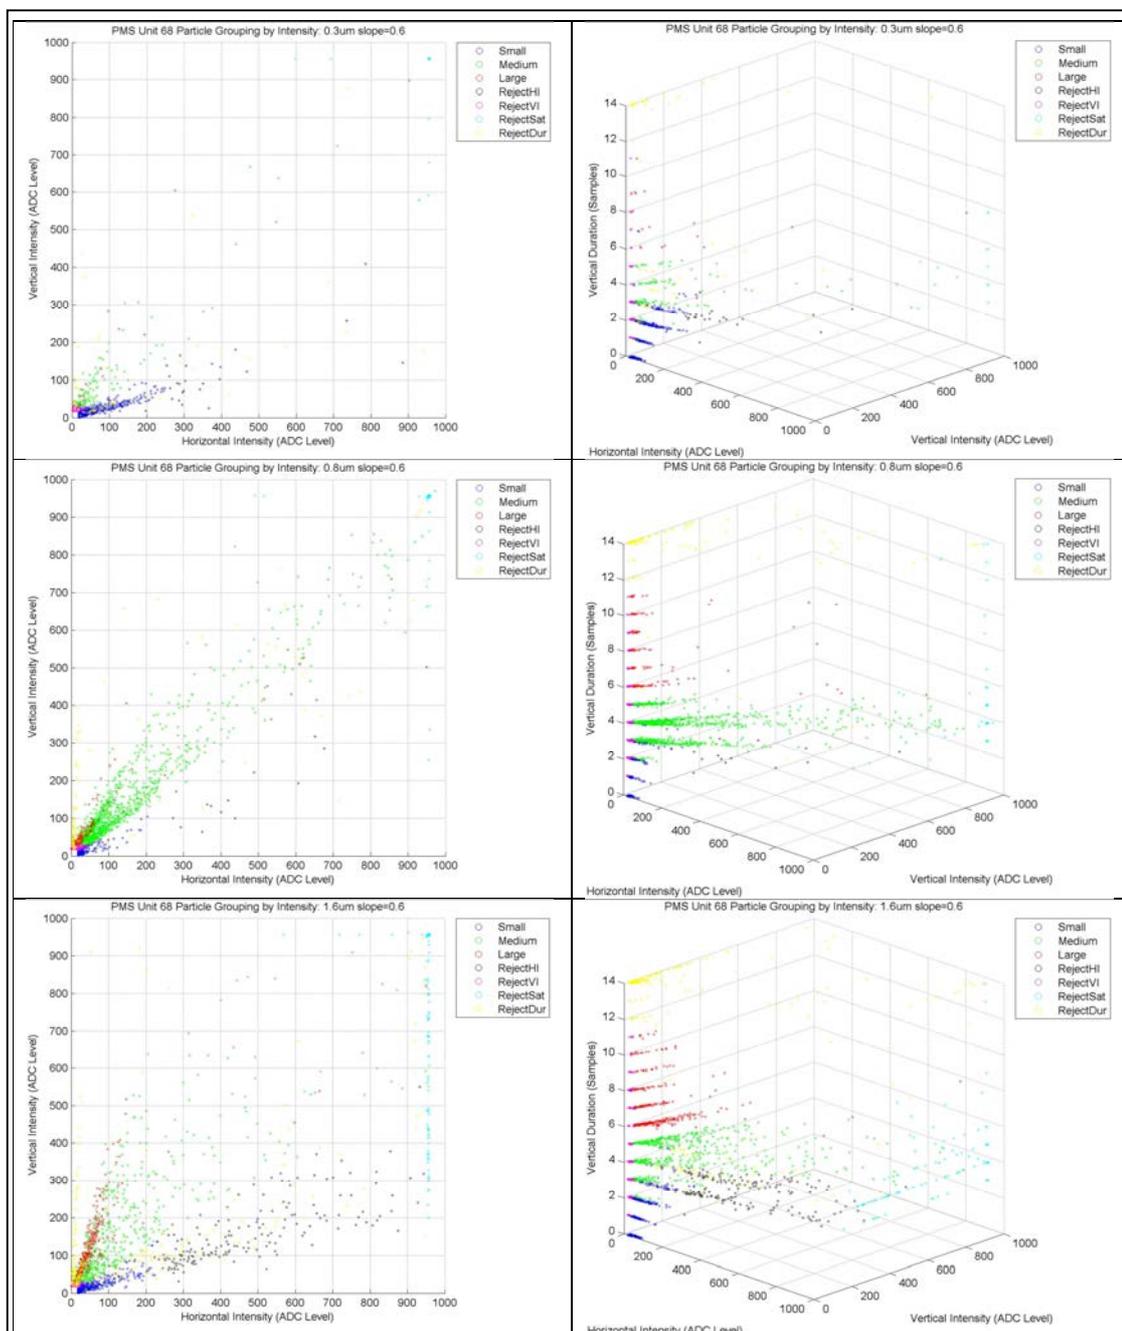


Figure 100 Classification scatter plots for 0.3 μm (top row), 0.8 μm (middle row), and 1.6 μm (bottom row) PSL particles. The horizontal and vertical axes of the plots on the left are the horizontal and vertical sensor ADC levels, respectively. The plots on the right are for the same data as on the left, rotated to view vertical sensor duration on the third axis.

Looking at the single sensor and orthogonal sensor arrangements quantitatively, one might be most interested in how the two arrangements compare to the PMS device. Using the PMS device as the expected value, the root mean square errors (RMSE) for both devices can be calculated for the various particle size groupings. Table 6 shows the RMSE values for the two different arrangements in this way. The means of each particle group are used as the values entered into the RMSE calculation. This compares the basic performance of the sensor arrangement in hitting the expected mean values coming from the PMS device.

Table 6 Root mean square error comparison of the single sensor and orthogonal sensor arrangements using the PMS device providing the expected value. The orthogonal sensor arrangement outperforms the single sensor arrangement for large and small particle categories.		
Particle Size	Single Sensor RMSE	Orthogonal Sensor RMSE
Small	3.40E+03	2.61E+02
Medium	NA	8.06E+01
Large	1.13E+02	3.55E+01

The orthogonal arrangement consistently shows better RMSE values for the small and large particle size categories. For the case of the small particles, this is mostly due to the over-counting of large particles that have relatively low intensity. In the case of the large particles, this is mostly due to the over-counting of small particles that have relatively high intensity.

Comparing the variances of the results of the two sensor arrangements is also useful in examining where each performs well. Table 7 shows the variances for the small and large classifications for the two sensor arrangements, with the variance shown for each particle size group tested with the sensors. For most of the particle size groups, the

Table 7 Variance comparison of the single sensor and orthogonal sensor arrangements. The orthogonal sensor arrangement outperforms the single sensor arrangement in most cases, with the exception of small particles for the group size of 0.3 μm .

Group Size	Single Sensor Small	Orthogonal Small	Single Sensor Large	Orthogonal Large
0.3um	1.37E+04	1.59E+04	4.73E+02	1.18E+02
0.6um	4.00E+04	2.67E+03	1.88E+03	2.04E+02
0.8um	1.54E+04	4.29E+03	1.50E+03	6.87E+02
1.0um	1.47E+04	7.58E+03	6.29E+03	1.28E+03
1.6um	1.84E+04	9.19E+03	5.43E+03	1.39E+03

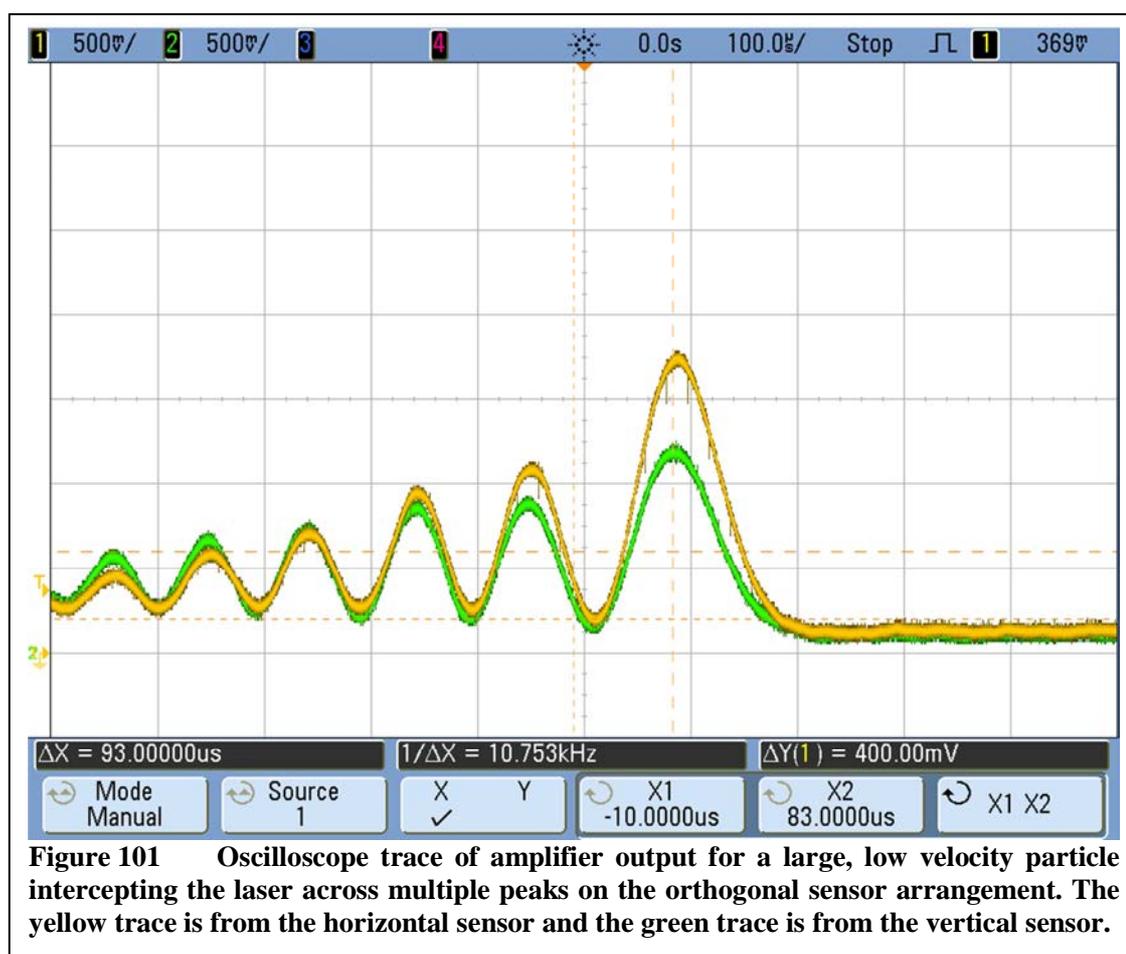
orthogonal arrangement performs better than the single sensor arrangement. The one exception to this is the 0.3 μm group for the small size classification. This is mostly due to the smaller number of particles that the orthogonal sensor arrangement has to work with in comparison to the single sensor arrangement. As implemented, the algorithm used with the orthogonal sensor arrangement tends to reject a portion of the smallest pulse height particles due to the use of the horizontal to vertical sensor intensity ratio. While the algorithm results in an overall more accurate set of particle counts as shown with the RMSE data, this is one area where the results are somewhat worse. This is a fair tradeoff to make, since the typical particulate matter environments measured by the WAQM OPC tend to follow a lognormal distribution in particle sizes. This results in many more small particles to work with for counting than large particles, which will help offset this issue with rejecting some of the small particles.

Misfit Particles

As shown in Figure 84 and Figure 85, the air flow along the edge of the flow channel can lead to durations that are quite long. This leads to some unintended consequences when combined with artifacts in the laser beam at its widest points in the

flow channel. The laser beam may contain multiple peaks and nulls. When a particle of sufficient size moves through these regions of varying intensity, they can produce a series of pulses. If these pulses are not recognized and handled appropriately, they can skew the results of one or more of the particle sizes.

An example of such a particle passing through a region of the laser with varying intensity is shown in Figure 101. This particle generates multiple pulses that individually might be seen as single particles passing through the laser in close sequence. While somewhat rare, these types of particles can be detected and rejected by the algorithm by the requirement of a minimum recovery time after the falling edge of a pulse. If a



subsequent rising edge violates the recovery time, the original particle duration is continued. This typically leads to the particle being rejected due to the duration limits on both small and large particles.

The rare nature of these particles makes them somewhat unimportant for the purposes of counting, but their ability to generate multiple false particles with a single pass through the laser beam necessitates their removal. They will mostly impact the counts in situations where there are very few particles, such as when performing size calibrations. A single large particle passing through the beam near the edge of the flow channel might produce many pulses that are recognized as small particles, skewing the results.

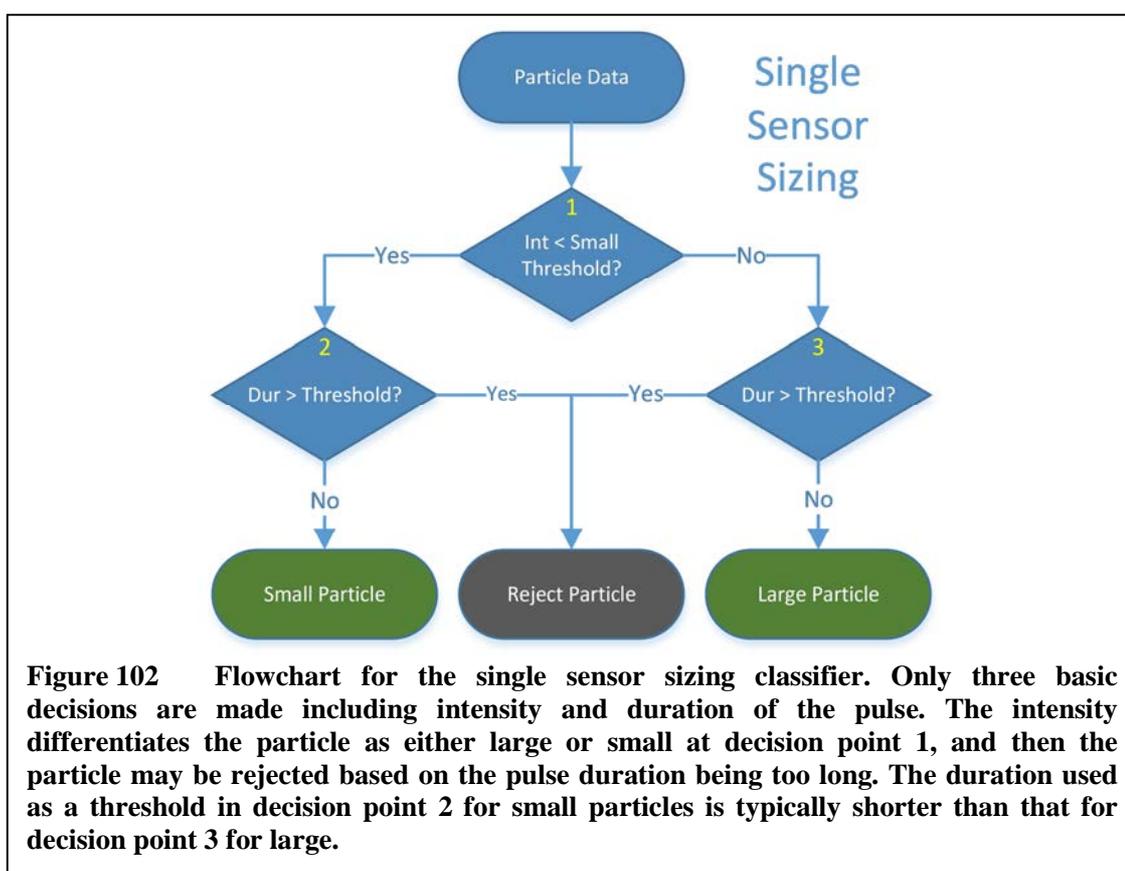
One negative outcome resulting from the removal of these misfit particles through the use of a recovery time is the increase in coincidence error. Coincidence errors occur when two particles pass through the beam close enough together in time that they cannot be distinguished as individual particles. In this case, the two particles may be counted as one. While this is the desired outcome with particles that are truly misfits, it is not the desired outcome when two particles are incorrectly grouped. The addition of the recovery time increases the probability that two individual particles might be counted as one. This is typically only a significant issue when the particle concentrations being measured are high, and the occurrence of multiple particles through the beam in close succession happens frequently. This issue will manifest itself as an increasing undercounting error as the particulate matter concentration increases.

For high concentration measurements, especially those with a low large particulate concentration, it may be desired to remove the misfit particle rejection step to

avoid the increased coincidence error. Though not implemented in the WAQM OPC, an automatic disabling of the misfit particle rejection step could be added to the functionality of the device. This could enable the rejection at low concentrations when it would most greatly impact the small count values, and disable it at high concentrations when it would mostly increase undercounting.

Sizing Implementation

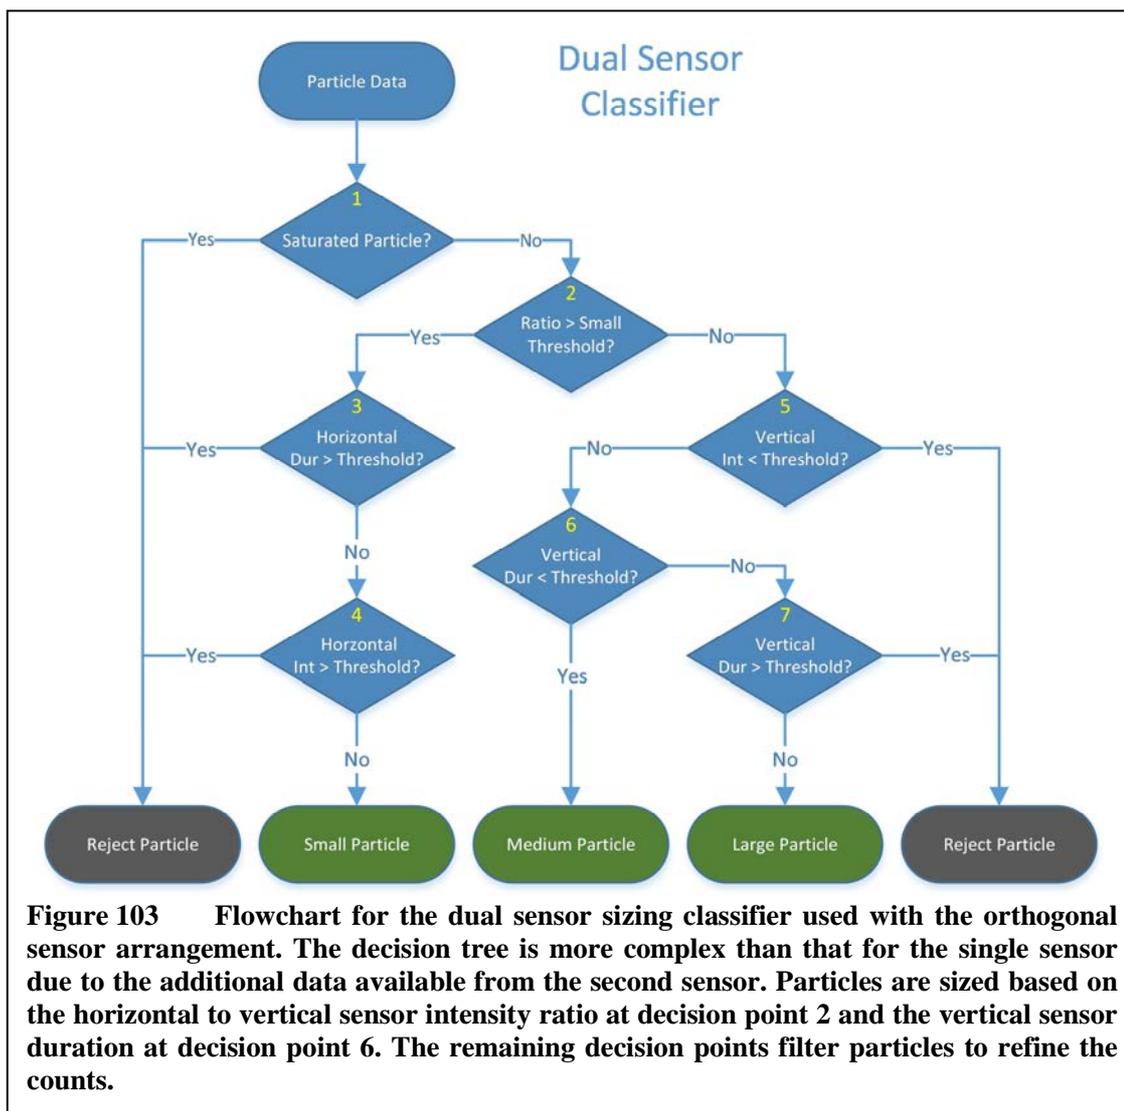
The sizing algorithms for the WAQM OPC are implemented in firmware running on the processor in the system's main microcontroller. For the single sensor implementation, this algorithm is straightforward, involving just a pair of simple decisions to classify a particle as either large or small. These are coded as inequalities in



the firmware, taking the form of simple comparison operations from the processor instruction set. Figure 102 shows the flowchart for the single sensor sizing classification. The first decision point separates the particle into either small or large categories based on the pulse height received by the sensor as passed through the large, low-gain channel. The second and third decision points filter out particles to improve the accuracy of the count based on pulse duration.

The thresholds for the second and third decision points are typically different, with the threshold for small particles being less than that for the large particles. The large particle duration threshold is typically set relatively high, at a value greater than 10 samples. This is to allow for the detection of large particles passing through a wider portion of the laser beam, while rejecting misfit particles as described above. The small particle duration threshold is typically set to a value near 5 samples, which will be just higher than the duration of the bulk of all small particles. As shown previously in Figure 90, this threshold helps to remove large particles with low intensity that might otherwise pollute the small particle count.

The flowchart for the orthogonal sensor arrangement is shown in Figure 103. This algorithm differs significantly from that of the single sensor. The algorithm is more complex, since data from an additional sensor is available to use for decision points and the additional category of medium particles is formed. This algorithm runs from data taken only from the small, high-gain channels of the two sensors, and does not depend at all on the large, low-gain channels used with the single sensor. This important detail allows the orthogonal arrangement to be implemented with very minor additional hardware cost in comparison with the single sensor arrangement.



The algorithm starts at decision point 1 by rejecting any particles that saturate either of the horizontal or vertical channels. Decision point 2 is the key sizing check, comparing the horizontal to vertical intensity ratio to a set threshold. Particles with intensity ratios greater than the threshold will take the small particle path, being further filtered to refine the particle counts. Decision point 3 removes particles with long duration from the small particle path, since these long duration particles are most likely to be large particles passing through a wide, low intensity portion of the beam. The small particle count is further refined by removing particles with high intensity at decision

point 4. These are particles with pulse heights greater than those produced by the vast majority of the small particles, which are mostly large particles that were not rejected by the filter at decision point 3.

Particles with horizontal to vertical intensity ratios less than the threshold at decision point 2 will fall down the large particle sizing path. Decision point 5 rejects particles that are low vertical intensity, since in this range there is a mixture of small, medium, and large particles that are not easy to categorize. Decision point 6 is the second sizing point, where large and medium particles are separated based on their pulse durations. Shorter duration particles are immediately classified as medium-sized particles. Longer duration particles are further filtered at decision point 7, with very long duration pulses being rejected to remove misfit particles as described above.

A further modification to the algorithm that has shown some promise in increasing the number of small particles available to count might be inserted after decision point 5. Particles that are under the vertical intensity threshold contain a significant number of particles that are actually small particles. These particles are difficult to separate from larger particles based on a single pulse duration threshold, as the durations for small and medium particles appear to overlap in this range of intensities. Some promise has been shown in separating out the small particles in this range based on a ratio of sensor durations. This method attempts to take advantage of the increased forward scattering from large particles, which would result in vertical sensor durations that are greater than their equivalent horizontal sensor durations. Small particles with less intense forward scattering would have longer horizontal than vertical sensor durations, especially since these particles tend to be detected at points closer to the horizontal

sensor. Effectively implementing this modification might improve the variance seen with the orthogonal sensor arrangement as shown in Table 7. However, not enough work has been done to show that the additional gains to the small particle count by this method are effectively removing larger particles from this range of intensities.

Most of the decision points in the orthogonal sensor arrangement sizing algorithm are simple comparisons implemented as inequalities in firmware on the microcontroller. The first sizing decision at decision point 2 is different, in that it contains a ratio of intensities that must be compared against a threshold.

$$\frac{I_{Horiz}}{I_{Vert}} > Threshold \quad 6.1$$

Calculating the ratio using division is not an ideal situation for the WAQM OPC. The microcontroller in the system implements integer division using a multi-cycle algorithm that can take nearly thirty times longer to compute than an integer multiplication. If the ratio comparison is rearranged, it can be converted to a simple multiplication.

$$I_{Horiz} > Threshold \times I_{Vert} \quad 6.2$$

The microcontroller does contain a hardware multiplier that completes in a single cycle through the execution pipe, keeping the overhead for this rearranged formula to the order of the single multiply to scale the vertical intensity by the threshold. Unfortunately, this formula with an integer multiply does not allow the threshold to have a resolution of less than 1. This can be changed by adding another multiply to the equation, to scale the horizontal and vertical intensity and achieve an effective higher resolution on the ratio threshold.

$$Scale \times I_{Horiz} > Scale \times Threshold \times I_{Vert} \quad 6.3$$

The threshold can be pre-multiplied by the scale to keep the number of per-particle multiplications down to two. For example, the scale value can be the number 10, which would allow threshold values with resolution of 0.1. In many of the examples above, the intensity ratio threshold used was 0.6. With a scale of 10, the per-particle calculation becomes

$$10 \times I_{Horiz} > 6 \times I_{Vert} \quad 6.4$$

which requires two integer multiplications and a simple comparison to complete. With intensity values from the 10-bit ADC, the 32-bit integers used in the WAQM OPC will not overflow. In fact, with 32-bit integers, much finer resolution thresholds could be supported without overflowing by increasing the size of the scale value if so desired.

Note that the algorithm does not show any post-processing steps, such as the correction of the medium particle count using the large particle count. This subtraction of a scaled large particle count from the medium particle count is performed only at the time the concentrations are reported, and not when each particle arrives. As such, the overhead for performing these types of post-processing corrections is much lower than that for the per-particle steps taken in the algorithm shown in Figure 103.

Detection and Sizing Conclusions

Producing accurate particulate matter concentration readings with the WAQM OPC is a function of how well the detection and sizing algorithms compensate for the deficiencies of the system. The desire to produce a small, low cost OPC device removes the freedom to use high powered lasers, quality beam shaping optics, light gathering

reflectors, or other costly hardware that might produce more predictable results. The simplicity of defining a sensing volume with a focused laser beam in an air flow channel meets the cost objectives but produces results that are difficult to use to create accurate information particulate matter concentrations.

The problem of detection might at first seem to be relatively simple, with the pulse produced by the analog front end of the system being fairly easy to distinguish from the background low-light condition. The less obvious issue with detection is the question of how to reliably convert this into a concentration that indicates the number of particles in a given volume of air. This is a straightforward operation if the sensing area and the velocity of the air are known. In this case, the volume of air sampled will just be the product of these two factors, and the concentration will be the particle count divided by this volume. However, as shown in Figure 83, the effective sensing area depends on the size of the particle. Without knowing the size of the particle detected, and the associated effective sampling area, converting the number of particles detected per unit time into a concentration will produce an accuracy error that might not be acceptable to a given application. Thus, for the WAQM OPC, it is not possible to separate detection and sizing without giving up accuracy in reported concentration. Better accuracy in particle sizing leads to improved accuracy in the reported concentration.

The single sensor arrangement has been shown to perform reasonably well, but with serious issues related to sizing accuracy and pollution of counts across sizing categories. This is especially true for large particles generating pulses that are counted as small particles due to passage through a low-intensity portion of the laser beam. The single sensor arrangement leaves the sizing of a particle as a function of the received

intensity at the sensor, which is converted to pulse height by the analog front end and the microcontroller ADC. This arrangement can provide a rough idea of the size and concentration of the airborne particulate matter, but leaves much to be desired in terms of concentration accuracy.

The orthogonal sensor arrangement improves upon this by taking advantage of the second sensor to roughly determine the position of the particle in the flow channel relative to the laser beam waist. Knowledge of this position can help in performing sizing of the particle, especially in determining if the particle is not small in size. This determination results in having more large particles with which to work in developing a count, and helps in removing large particle count pollution into the small category. Further differentiation in particle size can then be done using the duration of the pulses. The orthogonal sensor allows the sizing of an intermediate medium-sized particle category, which further aides in matching to an accurate sensing volume for all particle sizes.

While the WAQM OPC leaves much to be desired in reproducing the results of much more expensive equipment, it does provide valuable information regarding particulate matter sizes and concentrations. When targeting a particular application, the inaccuracies and issues associated with any OPC system must be considered, and this is no different for the WAQM system. When cost, size, and power are issues, the WAQM OPC performs well against other options. In situations where single point sampling accuracy is critical and size and cost factors are less important, other options may be a better fit.

CHAPTER SEVEN: PARTICLE COUNTER TESTING AND CALIBRATION

The development of an optical particle counter is not possible without a method for its testing and calibration. In this section, the construction and use of a testing and calibration system for the WAQM OPC is discussed. This system has been critical in the evaluation and normalization of the WAQM OPC design, along with the numerous derivative instances produced by the Hartman Systems Integration Laboratory.

In this chapter, the recognized sources of error in the WAQM OPC design are discussed along with their impact to OPC performance. Proposed methods for mitigating these errors are then identified. The details of the OPC calibration system are revealed, followed by the actual calibration methods used with the WAQM OPC.

Note that only calibration of the optical particle counter is discussed. Other sensors in the WAQM system do require calibration, but they are typically commercial sensors with established calibration procedures. Calibration of these other sensor is not covered by this work.

OPC Error Sources

Before discussing the calibration of the units, an analysis of the sources of error in the WAQM OPC is in order. In previous chapters, some of the issues with the WAQM OPC design are discussed. The issues are collected and analyzed in terms of their impact on the quality of the OPC output.

The quality of the OPC will be evaluated through two different metrics. The first metric is the sensitivity of the OPC. The sensitivity determines how small a particle can be detected when passing through the laser beam. Higher sensitivity will detect smaller particles, and a larger range of particles overall. The second metric is the absolute particle count, which is the number of particles counted in a given unit of time when measuring an aerosol of known concentration. Units with a higher absolute particle count will detect more particles per unit of time in a given aerosol than units with a lower absolute particle count.

The two metrics can easily be confused, since a unit with low sensitivity may also appear to have a low absolute particle count when measuring particulate matter with very small diameters. In general, the sensitivity of a unit must first be addressed, to normalize the unit's ability to detect the smallest particles. Once this is done, the absolute particle count may be scaled to match the actual particle concentration. If sensitivity is ignored, units will show a great deal of variation when measuring particulate matter with different size distributions.

There are a number of factors in the construction of the WAQM OPC that can impact the sensitivity, the absolute particle count, or both. The factors considered are the air flow, sampling loss, laser variability, analog front end variability, digital sampling error, and coincidence error. Each factor will be discussed along with a determination of impact on the two metrics of interest.

The intent of the calibration process is to remove the impact of unit-to-unit differences in sensitivity and absolute particle count on the output of the OPC. For the most part, the issues impacting sensitivity and absolute particle count can be corrected

through calibration. However, there are other factors in the design that might change over time, or are not consistent from particle-to-particle. These factors are grouped into a set of repeatability errors and are not addressable by calibration at this point in time.

Air Flow

The air velocity through the flow channel of the OPC directly impacts the sampling volume of the OPC. The sampling area inside the flow channel may be thought of as a two-dimensional area aligned orthogonally to the direction of air flow. This sampling area is defined by the laser, and for any single particle size it is the area where the laser is of sufficient intensity to cause detection when a particle passes through it. The size of this sampling area is particle-size dependent, as discussed previously and shown in Figure 83.

For this single particle size, the volume sampled will be determined by the amount of air traveling through the sampling area in a given unit of time. Higher velocity will result in a larger volume sampled per unit time. Assuming a uniform distribution of particles in the sampled air, the volume sampled, and by transitivity the air velocity, will be directly related to the resulting absolute particle count. Changes in the air velocity should not impact the sensitivity of the unit, as long as the air velocity does not increase significantly over the design point. A significant increase in air velocity might cause some particles to be missed due to shortening of the time the particle is inside the laser illumination, which may cause pulse durations that are short enough to be missed by the sampling of the ADC.

The single factor controlling the air velocity in the WAQM OPC is the DC fan at the outlet of the flow channel. As discussed in the chapter on the OPC design, this fan is

controlled by a proportional-integral-derivative (PID) controller that is based on the electrical measurement of the fan speed. Assuming the correct operation of the PID controller, the fan speed should remain constant. Furthermore, since the PID controller is a digital control system based on the microcontroller clock frequency, the unit-to-unit variation should be mostly dependent on the system clock frequency. This frequency is dependent on the crystal oscillator, which in turn is on the order of 100 PPM. Errors of such low magnitude are negligible in the overall error of the OPC. For the purposes of this analysis, the fan is assumed to run at a constant speed, and that this speed is constant across units.

If the fan speed is assumed to be constant, the air velocity will then be dependent on the unit-to-unit variation in fan construction, the fan's positioning inside the unit, how well the flow channel is sealed, and the deposition of dust and other particulate matter in the channel. Each of these factors will contribute to losses in the air velocity through the flow channel when compared with an ideal unit. While the first factor will likely remain constant, the other three factors may change over time with the aging of the OPC. The second and third factors may change due to thermal cycling of the enclosure and use-dependent mechanical stress. The final factor will depend on the environment in which the OPC is used. These factors may require periodic cleaning, adjustment, and recalibration, but a detailed analysis of these changes over time is outside the scope of this document.

A factor ignored in this analysis but important to understand is the condition of the ambient air velocity. The WAQM OPC is designed without any significant consideration to ambient air movement. If an external source of air movement is incident

upon the OPC, the air velocity through the channel may change enough to impact the absolute particle count. As such, the OPC should be protected from external air velocity due to wind or movement of the OPC. Without such protection, the WAQM OPC lacks the means to correct for changes in external air velocity.

Sampling Loss

The physical transport of airborne particulate matter from the ambient air in to the OPC sensor is fraught with opportunities to modify the particulate concentration in undesirable ways. This modification is referred to as sampling loss, in that some particulate matter may be lost as the ambient air is taken into the OPC and accelerated into the flow channel. The first issue related to sampling loss has to do with aspiration efficiency, which is the process of taking ambient air into the flow channel. The fan at the end of the flow channel creates a pressure differential that pulls air into the inlet of the OPC, causing the sampled aerosol to change velocity in joining the air flow moving into the OPC. This change in velocity will have an impact on the particle size distribution in the airflow. Depending on the severity of the change, large particles may not be drawn into the OPC at a rate equal to that of smaller particles due to inertia.

Once in the flow channel, the particulate matter may not move directly along a path parallel to the air flow. Depending on any given particle's initial velocity outside the OPC, the particle may follow a curved path that possibly impacts an edge of the flow channel. The particle may adhere to the flow channel wall, effectively removing it from those particles that are available for counting. This effect may be exacerbated by the shape of the flow channel, or the addition of ambient light blocking features at the inlet. Any redirection of air in the flow channel may form an inertial impaction point that

removes particles from the sampled air. Depending on the particular feature, this may selectively remove particles that are of a particular size or larger.

Figure 104 shows the state of a bend in the WAQM OPC flow channel after running for many weeks. Particulate matter has impacted the wall of the flow channel due to the velocity changes in the air flow at this point. The bend occurs after the photodiode in the flow channel, so is not a serious concern for sampling loss.

Other types of loss may occur in the flow channel as the aerosol moves from the OPC inlet to the sampling area. This transport loss may result from a variety of causes. Gravitational settling may occur, which tends to pull particles out of the air flow over



Figure 104 Inertial deposition in a bend in the WAQM OPC flow channel. This bend occurs after the photodiode, but illustrates particulate impaction when the air velocity changes direction. The particulate matter builds up in the channel after running continuously for many weeks.

time. Turbulence in the flow channel may also cause particles to impact the flow channel walls, removing them from suspension. While these effects have not been measured, it is assumed they are very small due to the relatively short length of the flow channel in the WAQM OPC. Electrostatic deposition may occur if a charge differential exists between particles and the flow channel walls. Given the nature of the environments in which the WAQM OPC is expected to be used, it is not likely that the sampled aerosol will be highly charged, and not much loss is expected from this phenomenon. Though not strictly accurate, for the purposes of the WAQM OPC these effects are assumed to have a small constant impact on the absolute particle count performance. Further study may be warranted in an attempt to quantify the impact of these phenomena on future designs of the OPC flow channel.

In general, analysis of sampling loss was not a significant focus of the development of the WAQM OPC. The mechanical and aerodynamic nature of the study were outside the main focus of the effort, and were only peripherally studied in an attempt to minimize the impact of sampling loss on the OPC. The main method for mitigating sampling loss was to keep the inlet relatively large with a somewhat gradual narrowing into the main flow channel at the sampling area without significantly curving the flow channel. For the purposes of error analysis, the sampling loss is expected to impact the absolute particle count for any given size range. This is addressed with adjusting the scaling values for a given size. It is expected that there will be some mismatch between the largest and smallest particles in a given size category since the sampling loss will tend to favor the removal of larger over smaller particles. Further study of the mechanical and aerodynamic characteristics of the WAQM OPC flow

channel may be warranted to better understand the impact of these losses on the overall error.

Laser Variability

Variation in the power output by the laser source is most certainly the largest source of unit-to-unit difference in the WAQM OPC design. The laser used does not accommodate any tuning of the power output, but is a fixed value that can vary widely. Some of this variability was removed by sorting of the laser modules prior to installation, removing the poorest performing modules among each lot used. Even after sorting, it was common to end up with a range of laser modules with powers that varied by plus or minus 33%. A finer sorting would be possible to remove so much variability in output power, but the cost of rejecting lasers that fall outside a narrower range was deemed too high for the WAQM OPC goals.

The laser power is impacted not only by the accuracy of the driving circuit and the quality of the laser diode, but also by the physical alignment of the diode inside the module. The power output of the laser module is regulated by a feedback photodiode located at the rear of the laser diode package. Light escaping from the rear of the laser diode will impact the surface of the photodiode, which is then used to regulate the diode power. The low-cost module employed in the WAQM OPC suffered from poor consistency in laser alignment, showing deviations that were off from the central axis of the laser module by as much as 5 degrees. Some of this light may strike the edge of the brass housing for the module, reflecting back into the laser diode. Light from such reflections impacting the feedback photodiode can reduce the overall output of the laser diode.

The alignment of the laser source within the OPC also plays a factor in variability. Due to the poor quality of the laser module, the module must be aligned inside the OPC to compensate for misalignment of the laser diode in the module housing. This alignment process may not provide a consistent placement of the laser beam in the flow channel relative to the photodiodes. A laser that is aligned with its beam waist further from the photodiode will produce received power levels at the photodiode that are lower than those for a beam that has its waist nearer the photodiode. As shown in Chapter 5, the scattering irradiance received by the photodiode will fall off as $1/r^2$, where r is the distance from the particle to the photodiode.

Finally, the shape of the laser beam in the flow channel is also a source of unit-to-unit variability. Variation in the laser module lens, diode, and positioning may impact the starting aperture size of the beam. A smaller aperture size will produce a beam that is narrower where it enters and exits the flow channel than a larger aperture size, but will produce a wider waist. An increase in waist diameter will result in a less intense waist, impacting the sensitivity of the OPC. The change in beam diameter will also impact the effective sampling area for different particle sizes in different ways. For example, the narrower beam stretches out the highest intensity portions of the waist across a wider portion of the flow channel in comparison to the wider beam, increasing the effective sampling area for a range of smaller particles. At the same time, the area where large particles may be detected decreases as the beam narrows, reducing the size of the low-intensity portion of the beam in the flow channel.

The output power of the laser diode directly impacts the incident irradiance of the laser beam in the flow channel, and the resulting magnitude of the irradiance from a

particle scattering this light. This in turn directly impacts the sensitivity of the OPC, as particles that might scatter enough irradiance to be detected with a higher power laser may not with a lower power laser. The output power will also impact the effective sampling area, which in turn will impact the absolute particle count.

The variability in laser alignment with its resulting potential to impact the received power at the photodiode will also impact the OPC sensitivity, as will the shape of the laser beam and its impact on the size and intensity of the laser beam waist. These factors will also affect the absolute particle count due to their impact on the effective sampling areas intersecting the flow channel.

Analog Front End Variability

The analog front end consists of the system of electronics from the photodiode to the input to the analog to digital converter. Each component in the system contains unit-to-unit variability that will impact the sensitivity and absolute particle count of the OPC. This variability is discussed in detail in the section of Chapter 4 on the construction of the OPC analog front end.

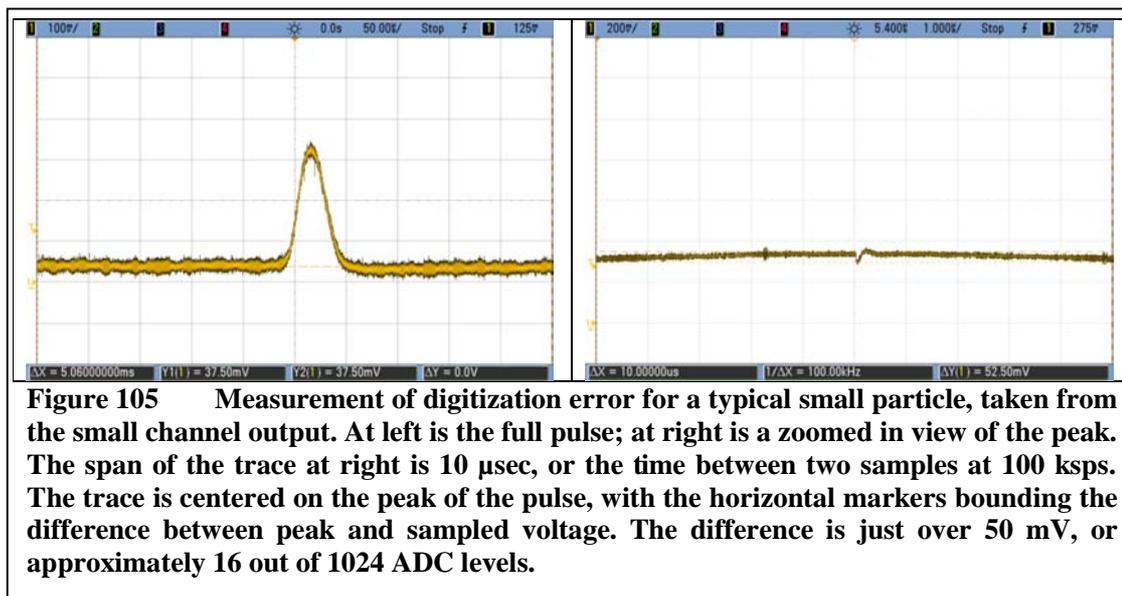
At present the WAQM OPC does not attempt to compensate for any temperature related impact on the performance of the analog front end, or for that matter, the laser module. Further study of this impact may warrant the inclusion of a compensation factor based on temperature feedback. For the prototyping activities up to this point, temperature related impact to the OPC sensitivity and absolute particle count has been assumed to be negligible. This may be an arguable assumption for indoor monitoring, but may not hold true for use in very hot or cold environments.

Digitizing Error

The process of converting the analog signal to digital values introduces further error to the signal beyond the analog front end. Two different errors are examined. The first is introduced by the analog-to-digital converter (ADC), in the accuracy of the device itself. The second is the peak detection error, which is a function of the sampling and pulse frequencies.

An ADC typically measures its difference from ideal through a set of error parameters. These commonly include the differential nonlinearity error, integral nonlinearity error, gain error, and offset error. These errors may be combined as an overarching absolute accuracy error for the device, in lieu of providing more detailed descriptions of each error. This is the case with the ADC in the AVR32 microcontroller used in the WAQM OPC. The absolute accuracy is given by the AVR32 datasheet [53] in units of least significant bits (LSB) as a maximum of 3 LSB when operated with 10-bit resolution. One might assume that this error could be calibrated out of the system, but since the lumped accuracy includes errors that may behave differently over the sampled signal range, or may change with temperature or voltage, the ADC error is treated as a repeatability error.

The peak detection error is an artifact of the sampling method. The ADC in the system is configured to sample at a rate of 100,000 samples per second for the single sensor and orthogonal sensor arrangements. The pulses produced by the analog front end typically last for at least 50 μ sec, and since there is no sample-hold circuit in the system, the peak must be determined from the sampling across the width of the pulse. Figure 105 shows a measurement of this error on a typical small particle pulse produced by the



analog front end. The image at left is the full pulse, displayed with a span of 500 μ sec. The image at right is a zoomed-in view of the pulse's peak with a span of 10 μ sec. Assuming that the worst case sampling would position two samples equidistant from the actual pulse peak, the measured difference between sampled value and actual peak value is just over 50 mV. With a full-scale range of 0V to 3.3V, this equates to 16 out of 1024 ADC levels. The best case sampling would be to hit the actual peak with a sample, so the measured ADC sampling uncertainty ranges from 0 to 16 ADC levels for typical particle pulses.

Summing the two sources of error in the analog-to-digital conversion process results in an uncertainty of 19 ADC levels, or just less than 2% uncertainty on the full scale of the ADC. This error will impact the OPC in two ways. First, the detection of very small particles is dependent on the pulse height exceeding a minimum threshold. The uncertainty will mean that particles generating pulses with heights that are within this uncertainty range above the threshold may or may not be detected. This will result in a gradual falling off of the absolute particle count as the particle size decreases towards the

threshold. The second impact to the OPC will occur at thresholds for particle sizing. Where decisions in the classification algorithm are made based on the pulse height and duration, the uncertainty in the analog-to-digital conversion process will result in some blurring of the classification compared to what might occur with the ideal sampling of the pulse height and duration.

Coincidence Error

There is a finite limit to the number of particles that can be detected by the OPC per second. For a single particle OPC, there is first a fundamental limit in how many particles may pass through the sensing area without overlapping in time. Particles that overlap may be counted as a single particle, or may be rejected as particles that exceed the duration limit. A typical small particle passing through the WAQM OPC takes approximately 50 μsec . In addition, the OPC may have a recovery time after the falling edge of the pulse caused by the particle. In the case of the WAQM OPC, there is a misfit particle rejection period after the main particle pulse during which no subsequent particle may be counted. This period is typically set to 50 μsec to ensure the rejection of particles passing through the fringes of the flow channel.

This sums to an effective particle transition and recovery time of 100 μsec through the laser beam. Given that particles will have random arrival times and will not pass through serially to maximize the available counting time, the number of particles that can be counted cannot be found just from the inverse of the transition and recovery time. Current OPC standards [80] calculate the coincidence loss using the formula

$$L = [1 - e^{-qtC_{max}}] \quad 7.1$$

where L is the coincidence loss, q is the flow rate in m^3/sec , t is the transition and recovery time, and C_{max} is the maximum particulate matter concentration in $\text{particles}/\text{m}^3$.

Some values for the WAQM OPC coincidence error are calculated in Table 8. The maximum particulate matter concentration for a set of particle sizes is calculated for various coincidence error rates. Note that the table uses particles/liter, since this is the typical standard of measure used for the WAQM OPC. As an example, to keep the coincidence error for $0.3 \mu\text{m}$ particles below 10%, the WAQM OPC should not sample concentrations greater than 569k particles/liter. As the coincidence error rises, the maximum particle concentration rises as well. A somewhat unique feature of the focused laser beam used in the WAQM OPC is that the maximum particulate matter concentration allowed for each coincidence error rate falls as the particle size grows. This is due to the dependence of the effective sampling area on the particle size. Since the effective sampling area grows with increasing particle size, the opportunity for coincidence error grows as well.

Given the types of environments in which the WAQM OPC will be used, with typically lognormal particle size distributions, it is most likely that the $0.3 \mu\text{m}$ particles

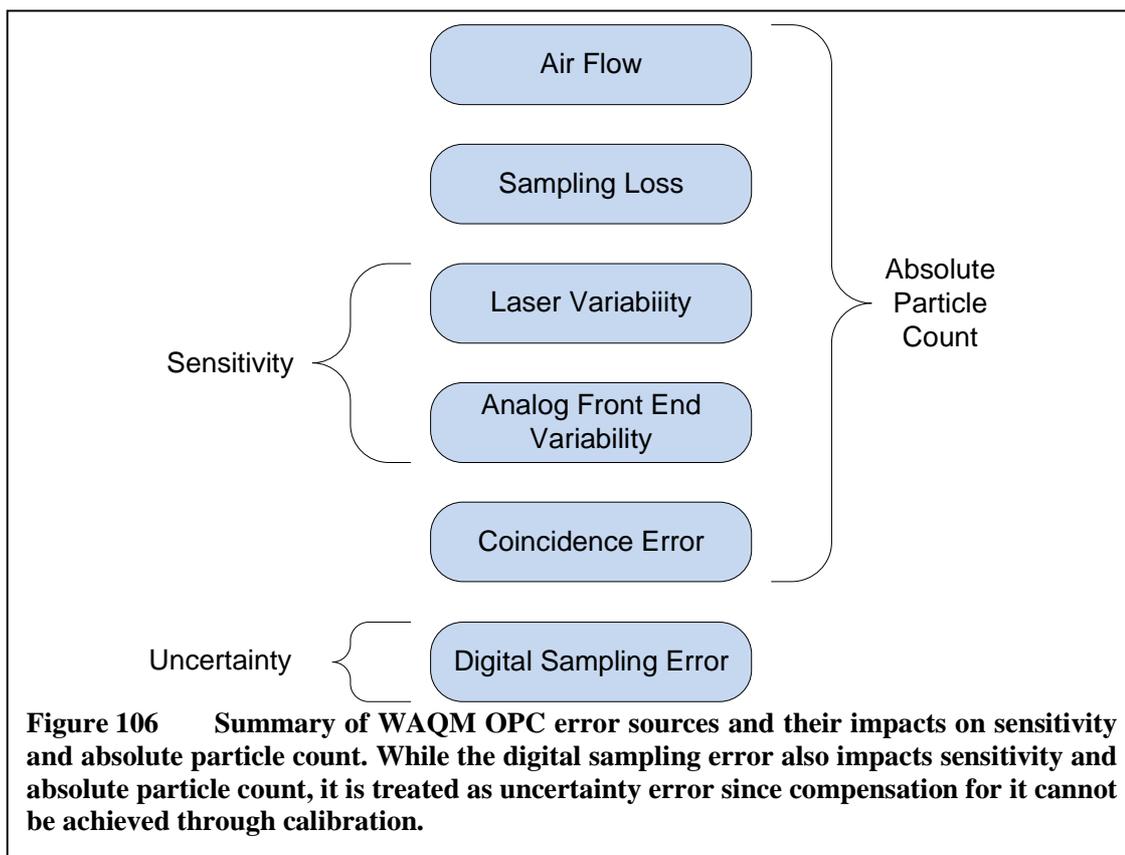
Coincidence Error	Particles per Liter			
	0.3um	0.5um	1.0um	5.0um
10%	5.69E+05	3.31E+05	1.93E+05	1.20E+05
20%	1.20E+06	7.01E+05	4.08E+05	2.55E+05
30%	1.93E+06	1.12E+06	6.52E+05	4.07E+05
40%	2.76E+06	1.60E+06	9.34E+05	5.83E+05
50%	3.74E+06	2.18E+06	1.27E+06	7.91E+05

will be the limitation for coincidence error. Larger sized particles are typically found in much lower concentrations than the smallest particles that the WAQM OPC can detect. This coincidence error will appear as an increasing error in the absolute particle count as the sampled concentration increases. At present, the WAQM OPC does not attempt to correct for coincidence error, as this is not performed on the commercial devices against which the WAQM OPC has been compared in testing.

Error Mitigation

Given the error sources discussed above, the objective of calibration becomes one of mitigating as many of them as possible to normalize the response of the units relative to each other and to match as closely as possible the output of other commercial OPC devices. Figure 106 summarizes the error sources and their impact on sensitivity and absolute particle count. The more mechanical processes of air flow and sampling loss tend to impact just the absolute particle count. The electrical portions of the system will affect sensitivity and absolute particle count. Coincidence error will impact the absolute particle count, but only at higher particulate matter concentrations. Finally, the digital sampling error is effectively uncorrectable, impacting sensitivity and absolute particle count due to inherent issues with the system design.

The calibration process will address sensitivity and absolute particle count. The former is handled by adjusting the detection trigger level, and the latter by adjusting a scaling factor. The trigger level will move the minimum ADC level that must be reached when detecting a particle before particle processing can occur. Raising this level reduces sensitivity by requiring a higher magnitude pulse height at the ADC for particle detection to occur. This compensates for factor such as a higher intensity laser beam, beam position



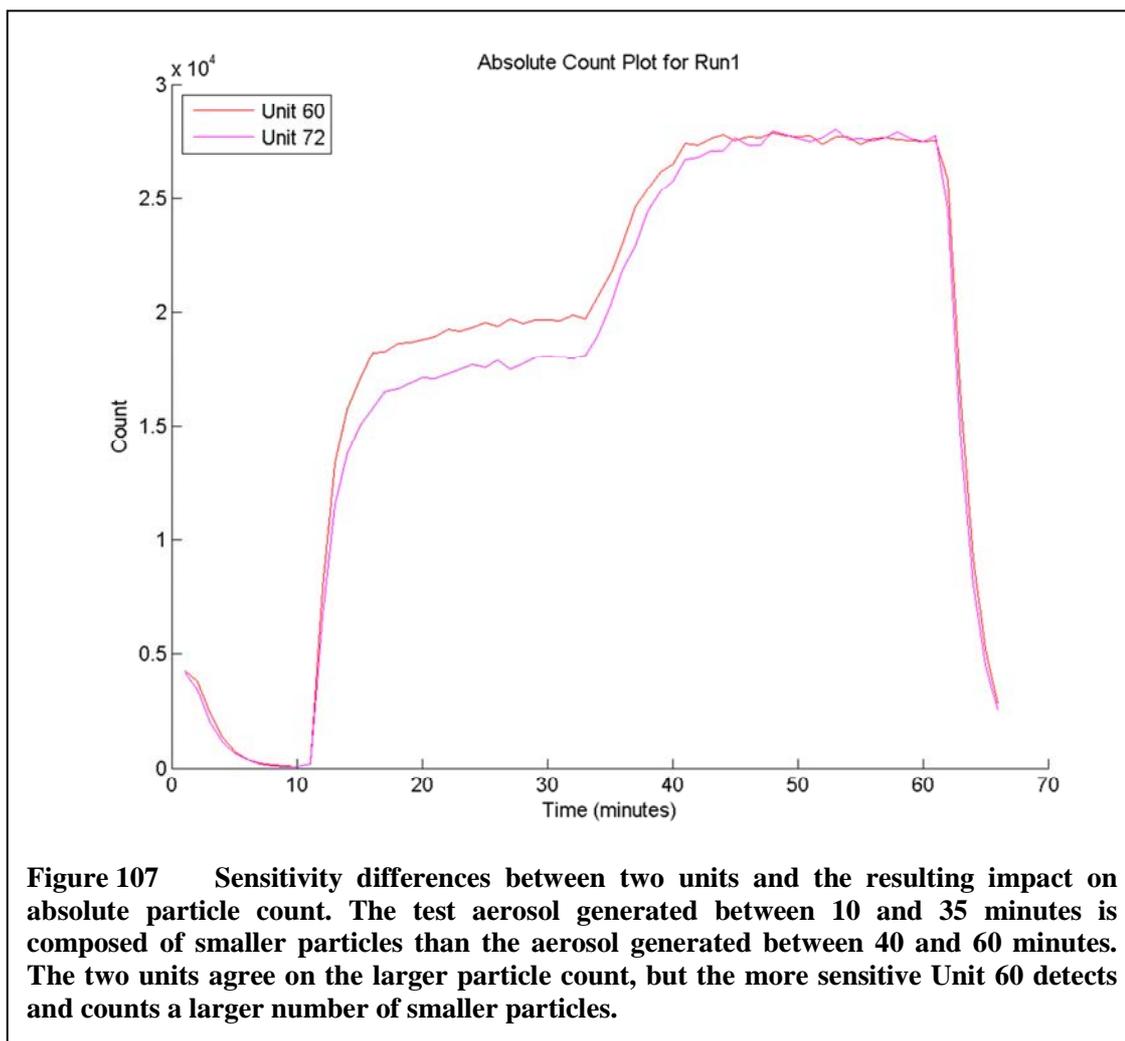
resulting in a reduced distance from the laser to the sensor, or variability in the analog front end.

Once the sensitivity of a unit is calibrated, the absolute particle count can be addressed. Note that it is critical to address the sensitivity first, especially with a calibration system that is not capable of producing a highly monodisperse aerosol. This is due to the impact of small particles at the detection threshold. A more sensitive device will include very small particles in its absolute count that would not appear in the absolute count of a less sensitive device, skewing the results. For example, the two uncalibrated units in Figure 107 show a relative difference between particle counts when sampling small particles as compared to sampling larger particles. The aerosol generated from 10 to 35 minutes is small, near the threshold of the two devices. The aerosol

generated from 40 to 60 minutes is larger. While the units appear to agree on an absolute count for the larger particles, Unit 60 reports a higher count than Unit 72 for the smaller particles. In this case, Unit 60 is more sensitive than Unit 72, allowing it to detect more small particles that pass undetected through Unit 72. Scaling the counts at this point would lead to errors when small particles are present due to the ability of Unit 60 to detect particles that Unit 72 cannot. To normalize the outputs from the two units, Unit 60 must first have its sensitivity adjusted to reject particles that Unit 72 cannot detect.

Once the sensitivity is addressed, the absolute particle counts can be matched through scaling. This is typically performed with a reference unit, which for the WAQM OPC must start with a commercial device as a golden standard. As mentioned previously, the devices available for the WAQM OPC development are the TSI Aerotrak 9303 and the Particle Measuring Systems (PMS) Lasair II-110. Once a calibrated WAQM OPC is available, it may also be used in the calibration of other units, though it is advisable to regularly check against one of the commercial devices.

The scaling will apply a multiplier to the absolute count to meet the concentration reading of the commercial device. This scaling is applied according to the type of sensor arrangement being calibrated. In the case of the single sensor arrangement, there are only two groups of absolute particle counts, large and small. In the case of the orthogonal arrangement, a medium size count must also be scaled. More specifics regarding the process of calibration will be given in the following sections.

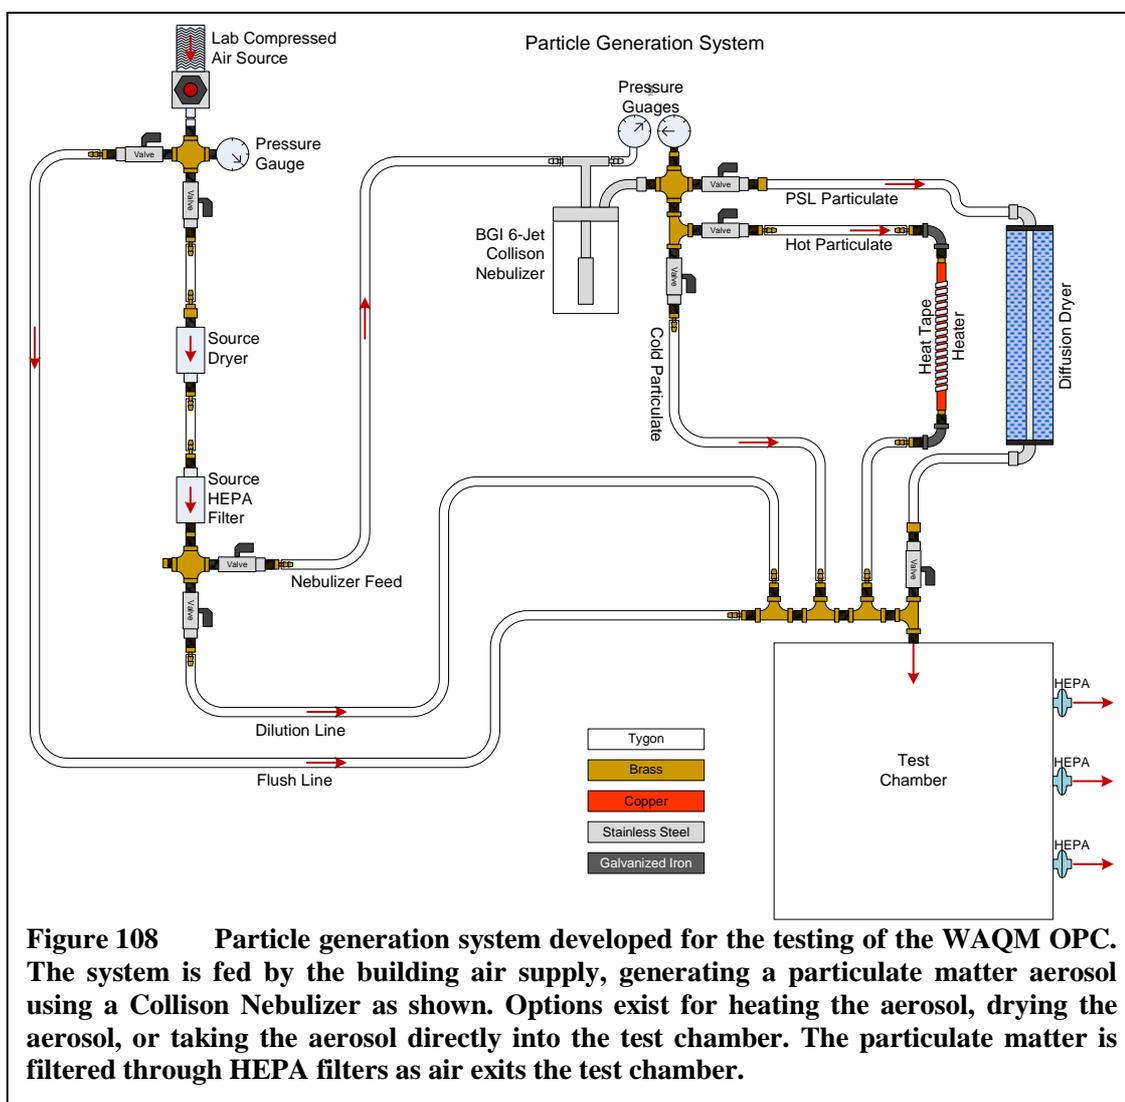


OPC Calibration System

The goal of the OPC calibration system is to produce aerosols that can be used to test and calibrate the WAQM OPC. Calibration is achieved through two separate methods. The first of these methods focuses on the sensitivity and absolute particle count of the small channel through the use of common tap water to produce relatively small particles in two different size categories. The second method generates larger particles using polystyrene latex (PSL) particles for the calibration of large and medium particle

concentrations. This section will focus on the system used to create the two types of aerosols.

The basic system was initially developed by the Boise State University Health Sciences department for respiratory studies based on a method described in a paper by Grassian et al [81]. This effort provided the chamber and nebulizer for the system, which was further developed for use as the OPC calibration system. A schematic representation of the OPC calibration system in its current state is shown in the diagram in Figure 108.



Red arrows in the diagram indicate the direction of air movement when the system is pressurized. Tubing and connection materials are indicated by color according to the key in the diagram. Most of the system connections are made using Tygon tubing, with intersections constructed of brass, stainless steel, or galvanized iron.

The calibration system depends first and foremost on a clean pressurized air supply shown in the upper left of the diagram as the compressed air source. This was provided in the lab hosting the calibration system by a tap into the building air supply. The air coming from this supply is mostly free of particulate matter and generally low in humidity. For consistency, the source air is still conditioned through a separate drier and filter prior to feeding the nebulizer. The source desiccant drier and HEPA filter are shown on the left side of the diagram being fed by the compressed air source of the diagram. A set of valves near the compressed air source select whether air will flow through the source filter and drier to the nebulizer, or directly into the test chamber. This second option is used to flush the chamber with air directly from the source.

Downstream from the source drier and filter is another set of valves. One path leads to the nebulizer feed, and the other path is a dilution line going directly to the test chamber. The valves are used to balance the flow into the nebulizer to set the desired pressure through the nebulizer jets. The nebulizer feed is typically run with its valve fully open, and the dilution line is then closed off to a point where the nebulizer inlet pressure rises to the desired level. This inlet pressure is monitored by the gauge shown on the top of the Collison nebulizer in the diagram. This pressure is normally run at 20 psig for the purposes of calibrating the WAQM OPC, producing 12 liters per minute of aerosol into the test chamber. Note that the dilution line can also be used to flush the test chamber, if

the compressed air source is not clean enough to provide a baseline of particle-free air. This was not necessary for the WAQM OPC development, since the compressed air supply was very clean in comparison with the generated aerosols.

The key piece of the OPC calibration system is the six-jet Collison nebulizer [82] produced by BGI Incorporated, shown near the top-center of the diagram. This device consists of an inlet for compressed air, a jar to hold the fluid to be nebulized, and a jet that draws in fluid and sprays it out toward the periphery of the jar. The resulting aerosol is pushed up through the lid of the device into an outlet that attaches to the rest of the calibration system. For the purposes of calibrating the WAQM OPC, the nebulizer is used with either tap water, or PSL particles mixed with deionized (DI) water. The compressed

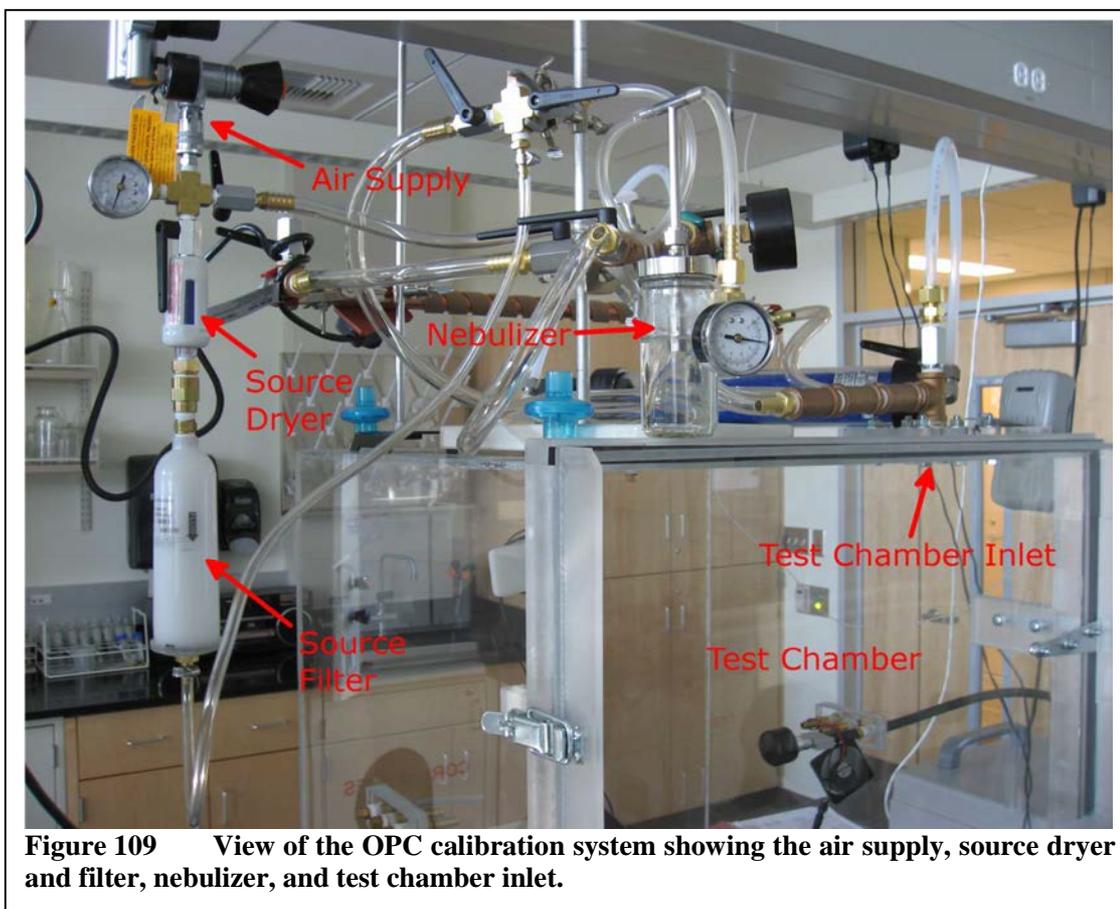


Figure 109 View of the OPC calibration system showing the air supply, source dryer and filter, nebulizer, and test chamber inlet.

air source and Collison nebulizer are shown in the picture of the calibration system in Figure 109.

Downstream from the Collison nebulizer is a set of three valves that select the path into the test chamber. Typically only one of these three paths is open at any time. Note that it is very important that at least one path is open when the nebulizer is pressurized; otherwise, the glass jar on the nebulizer may be destroyed. The path labeled “cold particulate” leads directly into the test chamber from the nebulizer. The path labeled “hot particulate” passes through a heated pipe before entering the test chamber. The final path labeled “PSL particulate” passes through a diffusion dryer before entering the test chamber. The cold and hot particulate paths are used for cold/hot testing, which is used to set the sensitivity of the units to be calibrated. The PSL particulate path is used for calibration of particle sizing. These testing methods will be described in the following sections of this chapter.

The heater in the calibration system is constructed from a copper pipe wrapped with an electric heat tape. The heat tape is controlled by a thermostat containing a temperature sensor for feedback that is pressed up against the copper tube. Copper tubing is used for its high thermal conductivity, which does a good job of distributing thermal energy from the heat tape along the length of the tube. To help avoid issues with heating the Tygon tubing that connects the heater to the rest of the system, a set of galvanized iron elbows are used to isolate the copper tube from the Tygon. The relatively low thermal conductivity of the iron helps prevent overheating of the Tygon material. The heat tape is capable of reaching temperatures of over 400 degrees Fahrenheit, but typical operation during cold/hot testing only requires the copper tube to reach 100 degrees

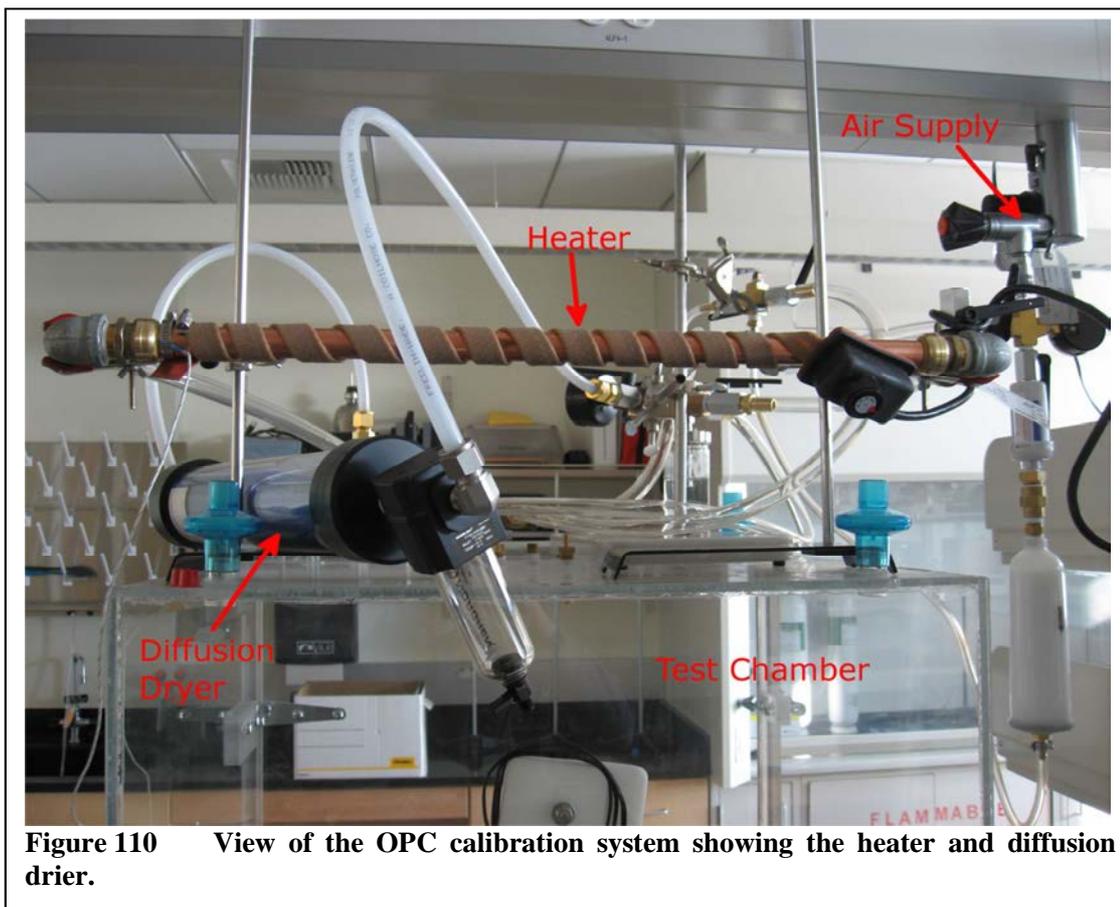


Figure 110 View of the OPC calibration system showing the heater and diffusion drier.

Fahrenheit. Care must be taken to not overheat the system with the heat tape as the Tygon tubing and other plastic components may be damaged. The calibration system heater is shown in Figure 110.

The main function of the PSL particulate path is to dry the aerosol generated from solutions of PSL particles and deionized water. These aerosols are used in the testing and calibration of the sizing algorithms in the WAQM OPC, and the desire is to have particulate matter that approximates a monodisperse aerosol. The desiccant drier will remove a portion of the water from the aerosol. This in turn helps to remove water from the surfaces of aerosolized PSL particles, moving them closer to their intended sizes. The diffusion dryer used in the calibration system is a TSI Incorporated 3062 Diffusion Dryer

[83]. This device uses desiccant particles surrounding a screen mesh tube through which the aerosol will pass. The desiccant particles will absorb water from the aerosol, and must be replaced regularly to keep the dryer effectiveness at a high level. The particles themselves may be dried in an oven and reused in the diffusion dryer multiple times. A valve at the output of the diffusion dryer prevents any flow back into the dryer when it is not in use, preserving the desiccant material from exposure to ambient moisture.

The outputs from the cold, hot, and PSL particulate paths join at the test chamber. The test chamber is a Plexiglas box with a footprint of 2.0 by 2.0 feet, and a height of 2.5 feet. One side of the chamber is a hinged door that latches shut, with a foam gasket around the edge to prevent air movement where it joins the chamber walls. The inlet from the particle generation system is in the ceiling of the test chamber, allowing test aerosols to flow downwards towards the floor where units under test are positioned. A small fan positioned under the inlet directs air in a circular pattern around the periphery of the test chamber, helping to mix the aerosol within the chamber. To prevent pressurizing the test chamber, a set of openings allows air to flow out into the room from the chamber. The chamber must not be pressurized to any significant level above the room air pressure, as this would in turn cause pressure in the nebulizer, possibly destroying the nebulizer jar. The room outlets pass through small HEPA filters to remove particulate from the air.

Up to nine WAQM OPC devices can be situated on the floor of the test chamber at any time, along with a handheld commercial device such as the TSI Aerotrak 9303. More devices might fit in the chamber physically, but limiting this to nine allows a comfortable amount of space around the inlet and outlet of each WAQM OPC.

Attempting to fit more devices in might constrict the airflow through a device, skewing the test results.

The test chamber also includes a sampling inlet for the PMS Lasair II 110 to allow this device to reside outside the chamber. Since the device is relatively large, it cannot be housed inside the chamber as can be done with the handheld TSI Aerotrak units. The PMS inlet passes through a hole in the chamber wall, using tubing supplied with the PMS device. The inlet may be covered when the PMS device is not used, to prevent air from passing through the PMS device when it is powered off.

Calibration Methods

Two different calibration methods are discussed, which may be used separately or in conjunction with each other. The first method, referred to as “Cold and Hot” calibration is a novel method of low-cost calibration for the sensitivity and absolute particle count of the small channel of the WAQM OPC. The second method employs calibrated polystyrene latex (PSL) particles that can be used to calibrate the small, medium, and large channels of the OPC.

Cold and Hot Calibration

Cold and Hot calibration is a low-cost method of calibrating small channel sensitivity and particle concentration. Common tap water is used as the starting fluid to be nebulized for the test aerosol, producing particles seeded by the ion content of the water. The particles, composed of mostly water, can then be used to produce two different size distributions depending on the application of heat. Particle count results from the two distributions are then used in the calibration of the WAQM OPC sensitivity

and small channel particle concentration. This method was developed alongside the WAQM OPC, and was discovered to be an effective for calibration during the initial testing of the device. The method is similar to the concept of particle size amplification [84] used in many condensation particle counters, which increases the size of particles through supersaturation based condensation.

The original particle generation system developed by the BSU Health Sciences department included a relatively small heater between the nebulizer and the test chamber that was intended to dry the particles produced by the nebulizer. This heater consisted of a cast iron pipe wrapped with a heat tape through which the nebulized aerosol would pass. The heater did not work as intended. In fact, engaging the heater actually increased the size and number of detectable particles in the chamber. This was due to the heater acting as a “growth tube” through the creation of supersaturated water vapor. Moisture collecting inside the pipe caused the partial pressure of water vapor near the pipe wall to be very near to the equilibrium vapor pressure at the wall temperature. When cooler air from the nebulizer passed through this tube, water vapor would diffuse from the pipe wall into the aerosol at a rate faster than the aerosol could heat to the pipe temperature. This created a supersaturation in the aerosol, which started heterogeneous nucleation, condensing water onto the particles. This would grow undetectable particles into the detectable size range, increasing the total number of particles counted. It also grew detectable particles into larger particles, moving up the total size distribution.

Even though the heater did not perform as expected in drying particles, it did provide a simple way to produce two aerosols with distinct differences in size and concentration. This is the key to calibrating the sensitivity of the WAQM OPC. The

measured particle counts for two units can be compared at the two different size distributions using Equation 7.2 for the relative sensitivity

$$S_{rel} = \frac{\sum_{T_{cold}} C_{unitB}}{\sum_{T_{cold}} C_{unitA}} - \frac{\sum_{T_{hot}} C_{unitB}}{\sum_{T_{hot}} C_{unitA}}, \quad 7.2$$

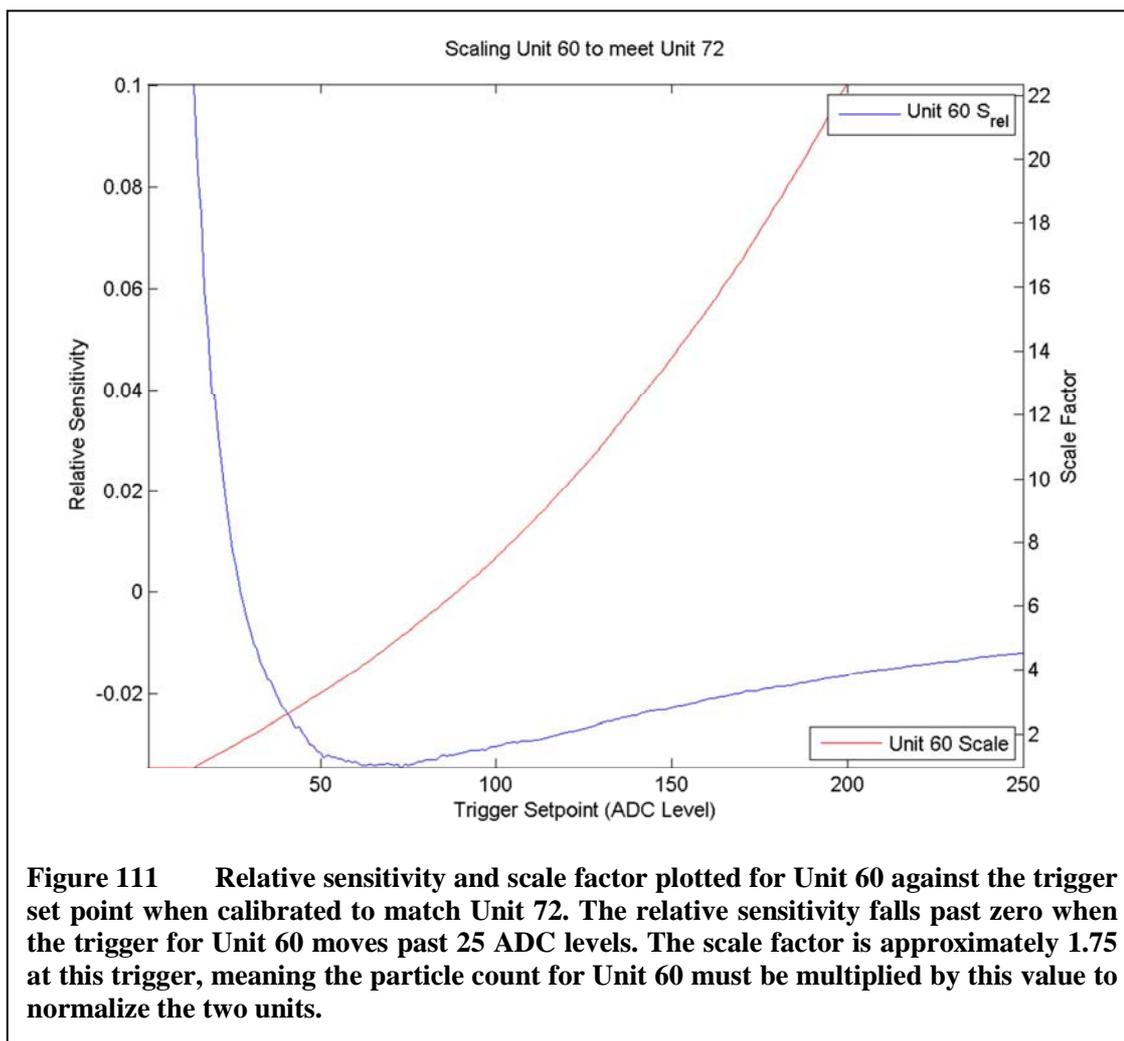
with C_{unitA} and C_{unitB} being the particle counts for the two units and T_{cold} and T_{hot} the time slices for the cold and hot particle generation respectively. As described previously, the cold particles will have a smaller average size than the hot particles. If S_{rel} is positive, the cold, small particle ratio is larger than the hot, large particle ratio. This indicates that Unit B is detecting more small particles than Unit A, which means that Unit B more sensitive than Unit A. If S_{rel} is negative, the hot, large particle ratio is larger than the cold, small particle ratio. This indicates that Unit B is detecting fewer small particles than Unit A, which means that Unit A more sensitive than Unit B. If S_{rel} is zero, the sensitivities for the two units match.

Using this relation, the more sensitive of the two units can have its trigger threshold raised until the relationship is at or near zero. Typically hitting a value of zero is not possible due to the discrete nature of the number of particles in the counts, so the trigger value that minimizes the magnitude of S_{rel} is chosen. The WAQM OPC firmware used for the calibration process stores an array of counts relative to the range of pulse heights every minute. These vectors of 1024 counts enable the programmatic determination of the correct threshold once the calibration run has completed. Note that all units undergoing calibration will have their trigger thresholds set to the same value, which will typically be more than 10 ADC levels. This minimum trigger is required to

avoid counting particles that are due to system or external noise passing through the analog front end.

One might ask why this type of calibration is not named “hot and cold,” to match the commonly used phrasing. The answer is rooted in the thermal properties of the system and time efficiency in performing the calibration. The heater is able to come to temperature quickly, so moving from cold to hot is a rapid process. Cooling the system down from hot to cold requires the heated pipe to cool into the ambient air and any air moving through the pipe, which is a much longer process. The cold particles are always generated first and the hot particles second, since there is inevitably time between successive runs of units to switch out devices, examine data, eat lunch, or perform other maintenance tasks. This time can be used to flush air through the heater to cool the system down for the next run. This is much more important when the air for cold and hot particles passes through the heater’s pipe. The latest version of the calibration system has a separate path for the heater, but the cold-then-hot rule is still observed to avoid inadvertent heating during the cold particle generation.

In Figure 107, two units were compared that had different sensitivities, resulting in counts that differed for cold particles but matched for the hot particles. In this example, Unit 60 detected more of the cold, small particles, indicating that the unit’s sensitivity was better than that of Unit 72. Applying the sensitivity calibration process, Unit 60’s sensitivity may be reduced to match that of Unit 72. Figure 111 shows a plot of the relative sensitivity for Unit 60 compared to Unit 72 and the resulting scale required to match the count of Unit 60 to Unit 72, both plotted against the trigger threshold set point. The blue line is the relative sensitivity, which can be seen to start at a value larger than

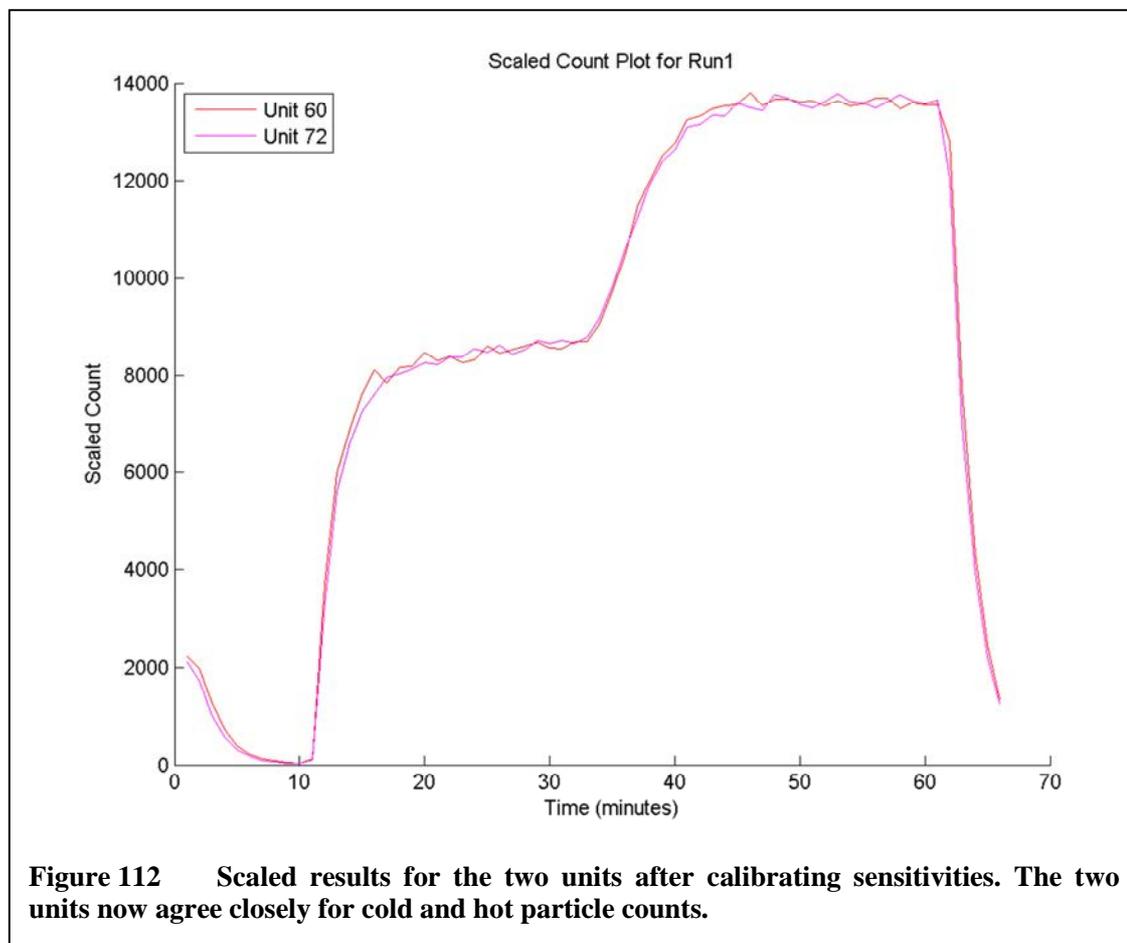


zero at the left of the plot. This indicates that when the trigger threshold for Unit 60 is at its minimum value the unit is more sensitive than Unit 72.

For this run, the minimum trigger threshold used for all units was 12 ADC levels, which is why there are no data for the relative sensitivity below this point. As the trigger threshold for Unit 60 increases, its relative sensitivity compared to Unit 72 falls, crossing zero at a value of 25 ADC levels. At the same time, the scale factor required for Unit 60 to reach the count of Unit 72, shown in red, increases. This increase in scale factor occurs since raising the trigger threshold for Unit 60 effectively throws away particles that

would otherwise have been counted. At the point that the relative sensitivity is nearest to zero, or 25 ADC levels, the scale factor for Unit 60 is approximately 1.75. This means that the count for Unit 60 must be multiplied by 1.75 to equal that of Unit 72. It should be noted that the scale factor required to match a commercial device is typically much larger due to the relatively small effective sampling area of the WAQM OPC.

Applying these values determined by calibration to the data from the two units results in the scaled particle count plot in Figure 112. Comparing this to Figure 107, the improvement in matching counts can be seen. The two units now match for cold and hot particles, improving the consistency in results between the units across a range of particle



sizes.

This particular example is somewhat extreme, in that the trigger thresholds for the two units found from calibration are quite different. Reducing the sensitivity of Unit 60 to match Unit 72 results in the loss of a range of particle sizes that could otherwise be detected by Unit 60. This also results in a relatively large scale factor for Unit 60, which typically results in a noisier output for the unit. In fact, Unit 72 might be considered to be outside the acceptable range of sensitivities for the targeted application. This determination must be made on a case-by-case basis, weighing the inclusion of less sensitive devices against the loss in sensitivity for other, better performing units. The sensitivity of Unit 72 might also be increased by realignment of the laser module, or replacement of the laser module with one that produces a higher power output.

This method can also be applied to multiple WAQM OPC units with a single commercial device as the golden reference. As mentioned previously, up to 9 WAQM

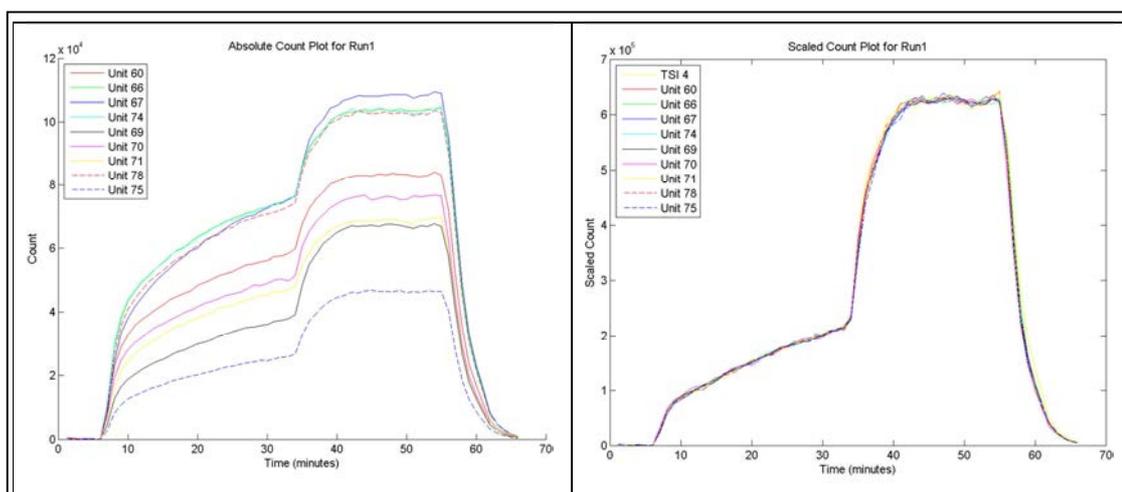


Figure 113 Calibration results from a run of 9 units with a TSI Aerotrak 9303 used as the golden standard. The plot on the left shows the absolute particle counts of the 9 WAQM OPC units. The plot on the right shows the calibrated, scaled results. A diverse set of sensitivities and absolute particle counts can be corrected to perform similarly to the TSI unit.

OPC devices can be fit into the test chamber at one time, along with a commercial device to be used as a golden standard. The pair of plots in Figure 113 shows just such a calibration run using Cold and Hot calibration. The plot on the left shows the absolute counts for the nine WAQM OPC devices, which have a large amount of variation in sensitivity and absolute particle count. The plot on the right shows the units after calibration to the TSI Aerotrak. The units show good correlation with the Aerotrak and with each other. Note that the cold particle portion of the test is first, and ramps up gradually over the course of the test. This is typical behavior for the calibration system when it is first used, as the system temperature and moisture levels stabilize. There does not appear to be any significant impact to the accuracy of the calibrations if cold particle data are taken from this ramp period.

PSL Particle Calibration

The second type of calibration and test activity conducted with the OPC calibration system uses polystyrene latex (PSL) particles in an attempt to generate an aerosol with a known size distribution. A solution of PSL particles and deionized water is created and nebulized, passing through the calibration system's diffusion dryer to remove water from the aerosol. As will be shown, the results are far from a monodisperse aerosol, but are useful nonetheless in calibrating the WAQM OPC.

The creation of the PSL solution is a step that must be performed with care. Pollution of the PSL solution can lead to erroneous results that will impact the accuracy of the calibration. The first step in the creation of a PSL-based aerosol must be a clean nebulizer jar. Any contaminants in the jar must be removed by washing with a cleanser and deionized or distilled water. All the cleanser should be removed with thorough

rinsing prior to adding the dilution water to the jar. In the development of the WAQM OPC, it was found that the use of a cleanser between PSL runs made a large difference in the amount of contaminant in the solution. Even the generic hand cleanser at the laboratory sink dispenser was effective in this use. A PSL solution will typically contain a surfactant agent to prevent particle clumping, which will remain on the sides of the nebulizer jar after a run. Rinsing the jar with water alone will not remove this material, which will then impact the amount of pollution in the subsequent run.

The nebulizer nozzle and stem should also be rinsed before starting a new run. While it is difficult to remove the nozzle stem and the jar lid from the OPC calibration system, a nebulizer jar that is nearly full of deionized water may be attached to the nebulizer lid to flush the nozzle and clean the stem. A flow of air should be passed through the nebulizer while the clean water is in place to flush the nozzle.

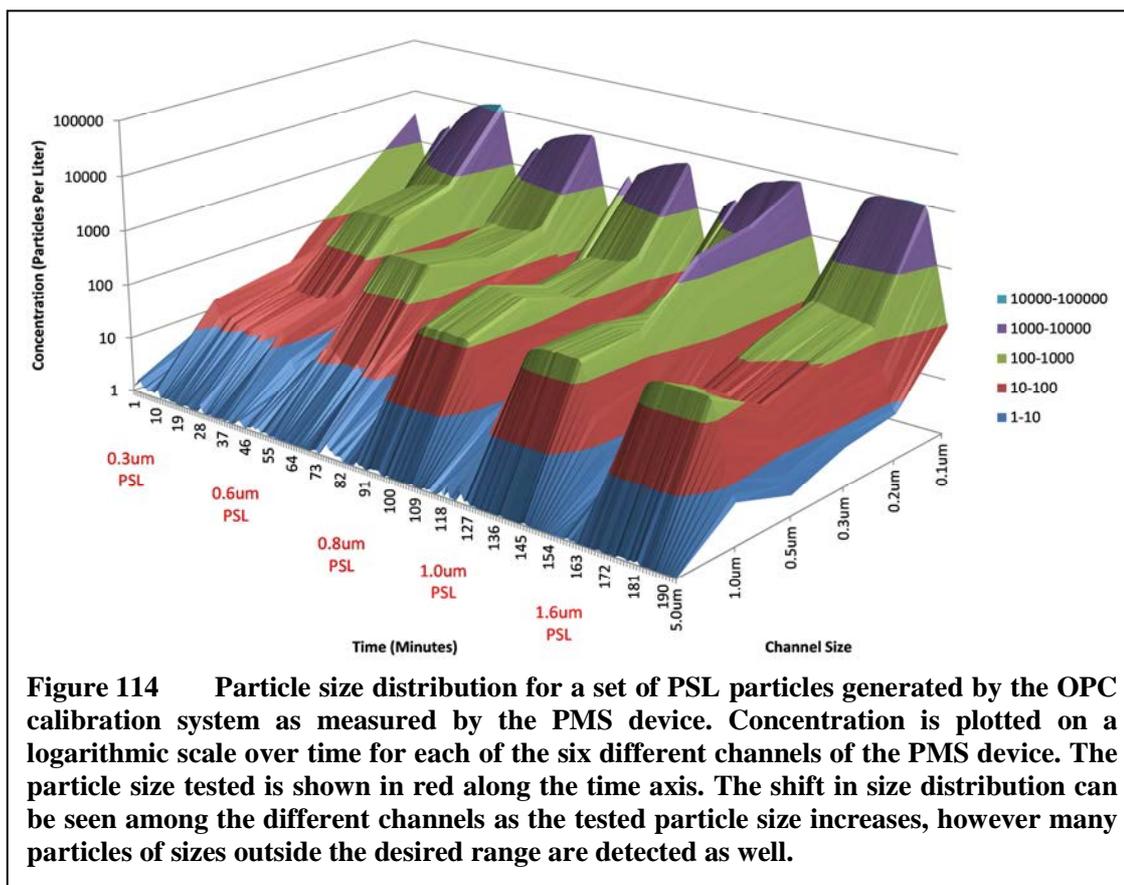
As a basis for the calibration, the water used to dilute the PSL particles must be free from any contaminants. Ions or other matter in the diluting water will cause the formation of non-PSL particles that may consist of water nucleated onto the seed material. As a rule, the first part of any test should be to nebulize the diluting water alone prior to mixing in PSL. This should result in an aerosol with very little particulate matter. If this is not the case, the water should be discarded, the nebulizer hardware washed, and the process restarted. In the development of the WAQM OPC, it was found that the laboratory provided deionized water was very clean, and resulted in an aerosol with very little contaminant. For dilution of PSL, 60 mL of this deionized water would be used to fill the bottom of the nebulizer jar. This provided enough liquid to allow the nebulizer to

operate for approximately 30 minutes while keeping the PSL particle concentration high without using too much PSL solution.

After testing the diluting water alone to check for contaminants, the PSL particles should be added to the nebulizer. For the development of the WAQM OPC, Thermo Scientific 3K/4K Series [85] particles were used with sizes from 0.3 μm to 1.6 μm . These monodisperse spherical particles come in a suspension with an approximate concentration of 1 billion particles per mL. The PSL material is applied to the nebulizer jar with a dropper built into the PSL container. Applying five drops of PSL solution to the 60 mL of clean diluting water will typically create aerosols with concentrations of several hundred particles per liter of air at the desired particle size.

Unfortunately, the OPC calibration system does not create a monodisperse aerosol as desired. The resulting aerosol covers a broad spectrum of particle sizes as reported by the PMS device. Figure 114 shows a sample run of multiple PSL particle sizes including 0.3 μm , 0.6 μm , 0.8 μm , 1.0 μm , and 1.6 μm . Each of these PSL particle sizes is tested individually as indicated in red along the time axis. The concentration measured for each of the six PMS channels is shown plotted on a logarithmic scale. It can be seen that the particle size shifts upwards for each set as the tested particle size increases. For example, the 0.3 μm test contains very little particulate in the 1.0 μm channel but the 1.6 μm test shows a significant peak of more than 100 particles per liter.

There is, however, a significant amount of undesired small particulate matter that comes along with the larger PSL particles. For each test run, there is a large amount of particulate matter detected by the 0.1 μm and 0.2 μm channels. This may be due to nebulization of the surfactant in the PSL solution combined with residual matter in the



deionized water and nebulizer jar. There is also a fair amount of pollution between the larger channels. For example, for the 0.8 μm test, there is a significant amount of particulate matter detected by the 1.0 μm channel. For a truly monodisperse aerosol, this should not happen. For this larger inaccuracy, it is likely that the cause is water residing on the particles left over from nebulizing the PSL solution.

While it would be most desirable that these issues with the creation of a monodisperse aerosol could be fixed, the cost and effort in doing so was not feasible in the budget and time frame of the WAQM OPC development. Further, a truly monodisperse aerosol is not strictly needed when calibrating against a commercial OPC as a golden reference. For the WAQM OPC development, the desire is to match the operation of the commercial OPC, which can be tested even with a polydisperse aerosol.

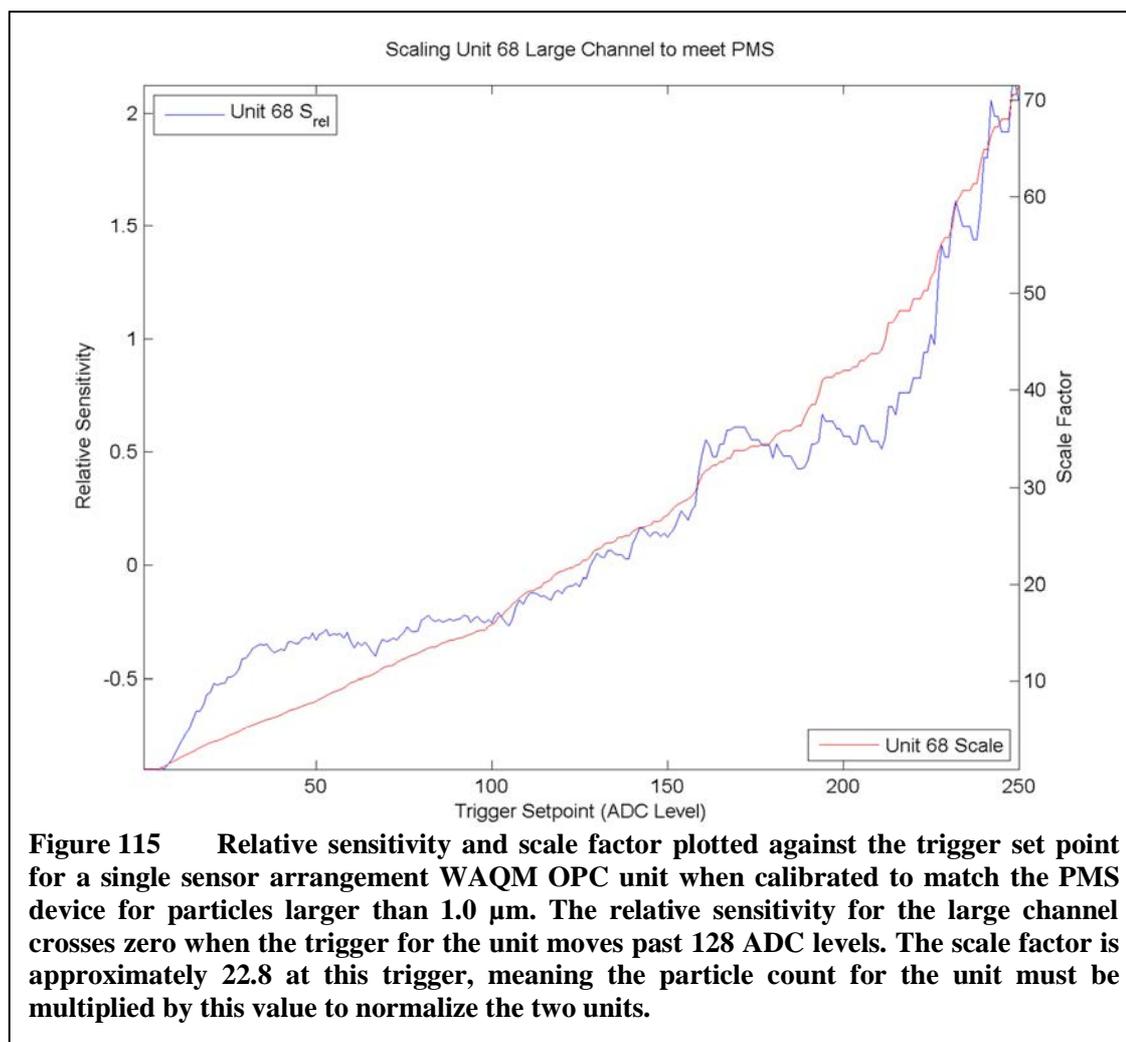
It is most important that at least some measurable component of the particulate matter is of the target size, but it is not absolutely necessary that the aerosol be free of other sizes. The mixture of sizes can still be matched between the commercial OPC and the WAQM OPC.

One final deficiency to mention with the OPC calibration system is the generation of aerosols containing PSL particles that are greater than 2.0 μm in size. In testing with larger particles, the system was unable to produce an aerosol with a measurable content of particles in the desired size range. There are several possible reasons that the current particle generation system is failing to produce particles of these sizes. To start, the nebulizer itself may be unable to aerosolize these larger particles at the level of air pressure used. To correct this issue, the air pressure should be run at a higher level. To test this theory, a higher pressure source may be required, or at the least the pressure available at the nebulizer should be increased through some means. The amount of tubing connecting the nebulizer to the test chamber may also be an issue in that a longer path might be causing the loss of most of the larger particles through deposition on the air flow path. This could be caused by gravitational settling or possibly electrostatic charge on the particles themselves. A final possible cause is in the PSL solution itself. The larger particles have a lower concentration in solution and may require significantly more PSL solution to be added to create an aerosol with a recognizable concentration at the desired size. Further study and experimentation will be required to address this issue.

To calibrate the WAQM OPC medium and large particle counts, a similar method is used as with Cold and Hot testing. Testing is performed with two different aerosols that contain PSL particles with nominal sizes greater than and less than the desired size

threshold. This provides the two size distributions that can be used with Equation 7.2 to calibrate the specific size threshold. For thresholds involving intensity or duration, a plot can be created showing the relative sensitivity against the threshold value, picking the zero crossing point as was done in Cold and Hot testing. This also applies to thresholds involving the ratio between sensor intensities, such as for the orthogonal sensor arrangement.

As an example, Figure 115 shows the relative sensitivity and scale factor for the large particle channel of a single sensor arrangement unit when calibrated to the PMS



device for particles greater in size than 1.0 μm . This relative sensitivity data are taken from the test run shown in Figure 86. The relative sensitivity crosses zero at an ADC level of 128, producing a scale factor of approximately 22.8. The large scale factor is required due to the relatively small effective sampling area of the WAQM OPC. Note that the relative sensitivity curve is not very smooth and is not monotonically increasing. This is due to the relatively small number of particles being used to calibrate the large channel. For example, in the case of this particular run, only 843 particles were classified as large over the 11 minute period in which 1.6 μm particles were being sampled. With a low number of particles, the entire spectrum of possible intensities is not filled out during this time period, especially for the more intense particles. The higher the threshold is raised, the fewer particles meet the mark, leading to increasing noisiness in the data.

Testing and Calibration Conclusions

The development of the OPC calibration system has been a critical piece of the WAQM OPC development. Without a system to measure and calibrate the WAQM OPC effectiveness, the unit could not have been properly designed or constructed. The initial system was constructed for use by the Boise State University Health Sciences department for respiratory studies, and was available as a starting point for WAQM OPC development. Multiple enhancements have been made to this system including better heating, drying, and aerosol measurement. Along with these improvements, processes for the reliable creation of aerosols and the test procedures for conducting calibration were developed. A rich set of tools and scripts were also produced to extract data from the WAQM OPC devices and analyze it in an automated fashion.

In many ways, the development of the calibration system was as difficult as the creation of the WAQM OPC itself. One might say that the two went hand-in-hand, with improvements in the calibration system leading to improvements in the OPC design. This bootstrapping process was instructive and incredibly frustrating at times. Without having a known good system for comparison, the significance of results often required a fair amount of correlation with simulation and bench testing to understand.

While many improvements could be made to the calibration system for future work, the system as it exists performs relatively well for prototype development and experimentation. The system was built for relatively low cost, with the most expensive portions being the PMS optical particle counter, the diffusion dryers, and the nebulizer. All told, the system was constructed for less than US\$25,000 in material.

The system should be relatively portable, in that it is not bound to the laboratory in which it was constructed. With a sufficiently clean compressed air source and the availability of deionized water, the system should be able to exist in almost any laboratory. Maintenance should also be minimal, with replacement of the desiccant in the diffusion dryer and regular calibration of the PMS OPC the most demanding pieces. As long as the nebulizer is kept clean and the heater is not misused, the calibration system should last for many years.

CHAPTER EIGHT: EVENT DETECTION FRAMEWORK

The goal of the event detection framework is to provide a basic platform for recognizing and alerting users to air quality episodes that may be unique to the particular environment. Other air quality systems have explored the implementation of expert systems for diagnosing events in specific situations, such as for building environmental control. With the WAQM system, a framework is provided that is extensible to many different situations. The framework does not attempt to provide a trained diagnostic system, but does facilitate the development of such a system. The framework is extensible to allow for the inclusion of such a system should it be desired in the future.

The goal of the framework is to ease the diagnosis of air quality events, removing the need to manually scan data for issues of significance. The system may be used across a wide array of applications, from home air quality issues for asthmatics, to the monitoring of aircraft cabins for bleed air issues. As such, the framework should have the following features.

- Detection of events of interest local to each node,
- Detection of events of distributed interest across the WSN,
- Dynamic programmability of event triggers from a user visualization device or server,
- Forwarding of event indications from the WSN through to a system server,

- Server generation of notification messages to the user, including a link to a web-viewable dataset, and
- Server acceptance of user input labeling a particular event.

While much if not all of the detection work could be accomplished at the server level, there are several reasons why it is compelling to perform as much as possible at the sensor nodes themselves. Perhaps the most compelling reason is flexibility in deployment. Sensor nodes may not always have a connection back to a server, and providing some form of notification locally, such as an audible alarm, may be the only choice. In addition to this, pushing the work of identifying events out to the node distributes the workload, making the system more easily scalable compared to processing at a single point of the server. On the downside, the processing power of the sensor nodes limits what can be accomplished. Even so, providing the option for sensor node processing gives the most flexibility, in that nothing stops the server from performing additional processing on the data to detect events that are out of reach of the sensor node capability.

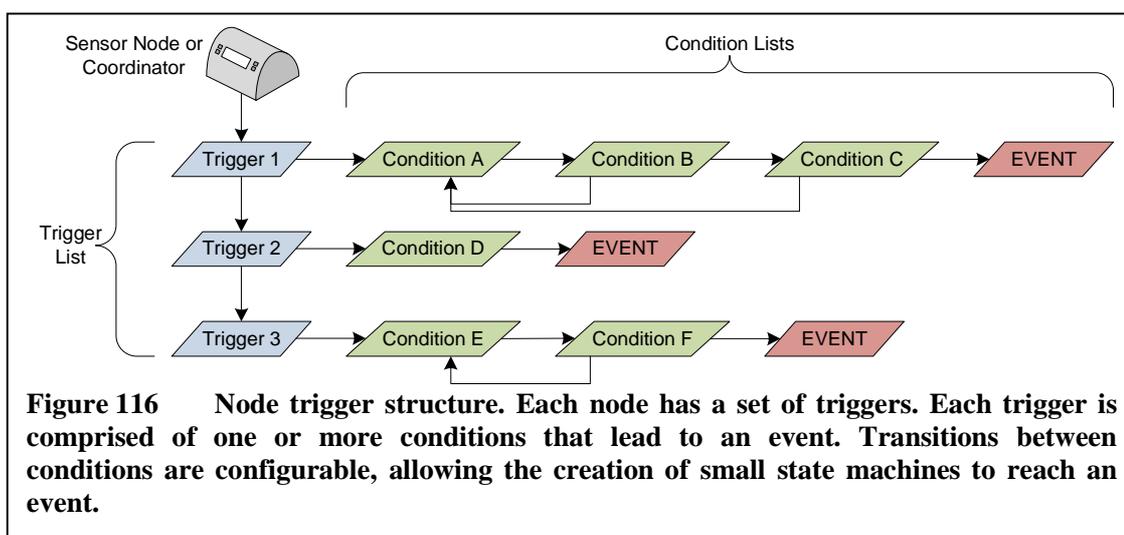
This development was started and employed in a partial form as part of the WAQM system development prior to the writing of this dissertation. Some parts have not yet been developed including the WSN distributed event detection, dynamic programmability, and user labeling of events. Though the development is incomplete, the design is complete and may be implemented in future versions of systems developed by the Hartman Systems Integration Lab.

Dynamic Event Triggering

Event triggers are a set of programmable trigger conditions that will result in sending an event message from the WAQM node meeting the trigger conditions. Each event trigger chains together a set of conditions to form a sequence that must be met before the event trigger fires, sending the message. This allows a complex series of conditions to define a trigger. Rather than just being able to trigger on a single condition, the trigger can be composed of multiple conditions that must be met in sequence before the event message is generated.

Sensor Node Event Triggers

Figure 116 shows the basic structure of the sensor node event triggers. The triggers are linked together in a list, and are processed one at a time each time a measurement is passed to the main sensor task from a specific sensor. The trigger tracks a current condition state, which indicates the condition on which the trigger is waiting. The measurement is tested against the current condition for each trigger, and an action is taken based on the result. The action may move the trigger to the next condition in the



chain, return to a previous condition, or return to the start of the condition chain, resetting the trigger. Once the end of the trigger's condition chain is reached, the event message is generated, signifying that the sequence was successfully completed.

Each condition has a success path and a failure path. If the condition is true, the success path is taken. If false, the failure path is taken. This allows a rich set of potential trigger definitions to be constructed, where a failure may reset the trigger, or move it back to a prior state

At the sensor node level, a basic set of trigger conditions is defined as shown in Table 9. These consist of level, edge, and delay conditions. Level conditions occur if the sensor's absolute level is above or below the level associated with the condition. Edge conditions occur if the sensor's rate of change exceeds the rate of change associated with the conditions. Delay conditions provided timing events, pausing the processing of conditions for a specified period of time.

Delays in timing may also be achieved through the level and edge conditions, as both types of conditions include counters. These counters cause the state of the trigger to remain at the counting condition until the count has been reached. This forces the condition to remain true for the duration of the count, otherwise the failure path is taken. If it is not desired to monitor the level or edge condition for the length of the count, a

Table 9 Event trigger condition types.	
Condition	Description
Level Above	The sensor data have reached a level above the defined threshold.
Level Below	The sensor data have reached a level below the defined threshold.
Rising Edge	The sensor data are rising at a rate greater than the defined rate.
Falling Edge	The sensor data are falling at a rate greater than the defined rate.
Delay	A specified amount of time has passed.

delay condition can be chained onto the end of the level or edge condition to provide a delay that is not dependent on measurements.

An example of an event trigger is shown in Figure 117. The trigger's purpose is to capture a carbon monoxide event in the garage of a residential home. The desire is to trigger on a case where an automobile starts in the garage, and then the garage is shut, trapping the carbon monoxide. To build the trigger, two conditions are used. The first condition is a rising edge condition, which detects the sharp increase in the carbon monoxide concentration. The second condition is a level condition that ensures that the carbon monoxide concentration doesn't drop quickly, as would happen if the garage remained open. Once the event message is sent, the trigger is reset. In this case, the rising

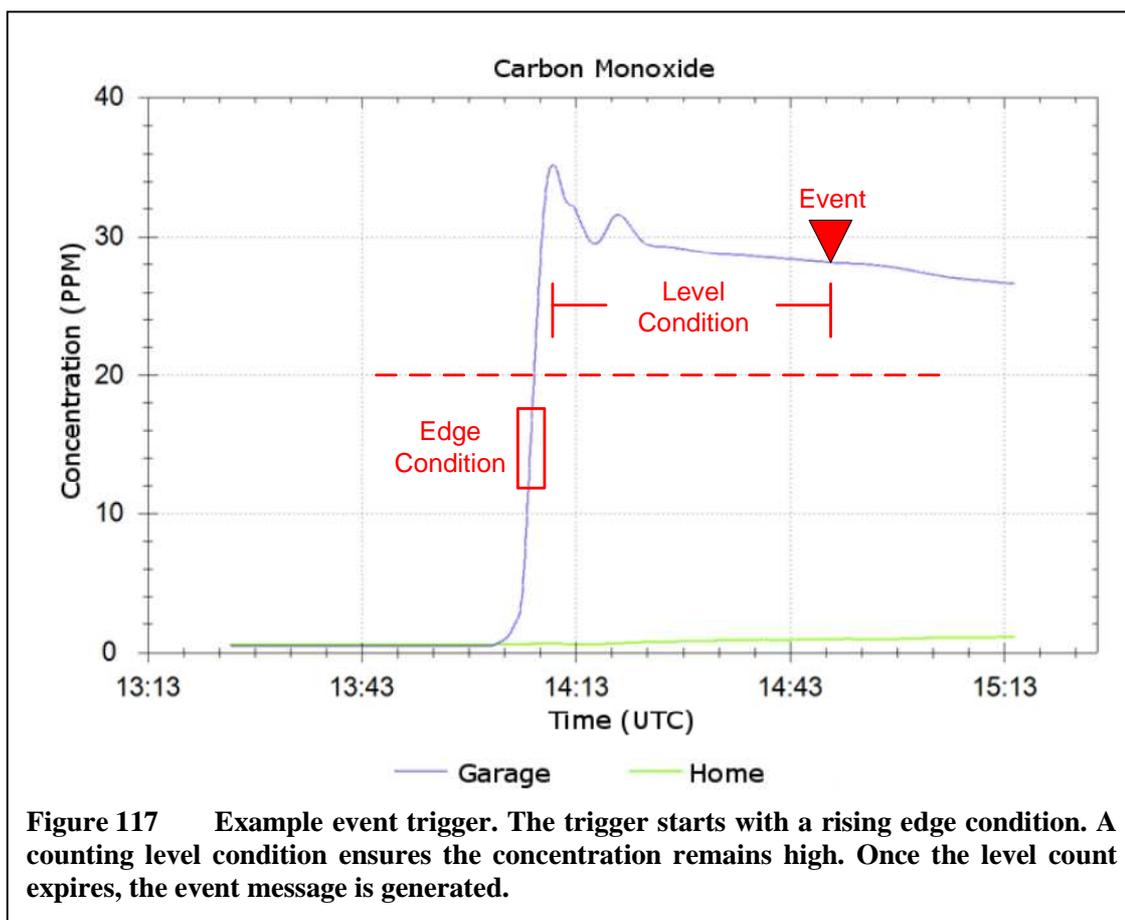


Figure 117 Example event trigger. The trigger starts with a rising edge condition. A counting level condition ensures the concentration remains high. Once the level count expires, the event message is generated.

edge condition prevents a second message from being generated, even though the carbon monoxide level remains high.

The event detection triggers and conditions are intended to be programmable at the sensor node, and will be set up to allow modification by the server based on a continued learning process. While it is possible to move all of the event detection to the server, this would not be as scalable for large networks, and would preclude operation of the framework without a connection to the server. The server is important to the notification of the user in this case, but event messages are also implemented to provide notification without a server. The framework currently issues an audible alarm from the sensor node issuing the event message, allowing local awareness of the event without server involvement.

As mentioned previously, dynamic event triggering has only been partially prototyped and not fully implemented. A key piece that has not yet been implemented is the ability to build event triggers dynamically, pushing the triggers from the server to the sensor node. This will require extensions to the existing messaging protocols used to pass data and control information between the sensor nodes and the server. The task is not seen as critical to the description of the framework, and will likely be completed at a later date as part of the ongoing work of the Hartman Systems Integration Lab.

WSN-Level Event Triggers

While single sensor node events are definitely interesting, using the WSN as a single logical sensor can provide further value in event detection. For example, while it may be interesting to know that one or more sensor nodes are detecting a particular value on a particular sensor, the end user may not want to be notified until all sensors in the

area see the same value. For instance, an end user may be employing the sensor network to monitor particulate matter in a home with three sensor nodes placed throughout the building. Once a peak in particulate occurs somewhere in the house, the user may want to run the central air conditioning fan to filter the particulate out of the air. The user may then want to be notified when the particulate matter reaches a low level at all sensor nodes, allowing him to shut off the fan. The potential also exists to add a control system for ventilation that is based on particulate matter concentration. Normal ventilation control systems are based on temperature, but with this system a high particulate matter concentration could trigger the fan to keep the air clean.

There are several methods for the WAQM system to detect WSN-level events. The brute force approach would be for the coordinator of the WSN to process the data streams from all the nodes in the WSN, processing a set of event triggers that have characteristics similar to those described for the sensor nodes. The coordinator would also need to have conditions that are sensor unit dependent to build inter-node dependencies. This approach is not very scalable, since it depends on the coordinator processing an amount of data that grows at least linearly with each added node in the WSN and possibly exponentially depending on the desired set of triggers.

A lighter-weight version would make decisions at the coordinator based on event messages from the sensor nodes. Rather than processing the data streams, the coordinator relies on the sensor nodes to process their own streams, and then collects event messages from each node. These event messages are then used to follow a slightly more complex version of the event trigger structures shown in Figure 116, where instead of conditions

that are dependent only on single sensor edges or levels, the conditions are dependent on a collection of unit states.

While the WSN-level event detection has not yet been prototyped, the current design intent is to have the coordinator employ a light-weight event message version of event detection. This will require some modification to the sensor node event trigger processing, as it will require some event triggers to be specific to intra-WSN notification of the coordinator. This will be implemented as a simple flag on the trigger, causing a message to be generated that will not leave the confines of the WSN.

User Interaction

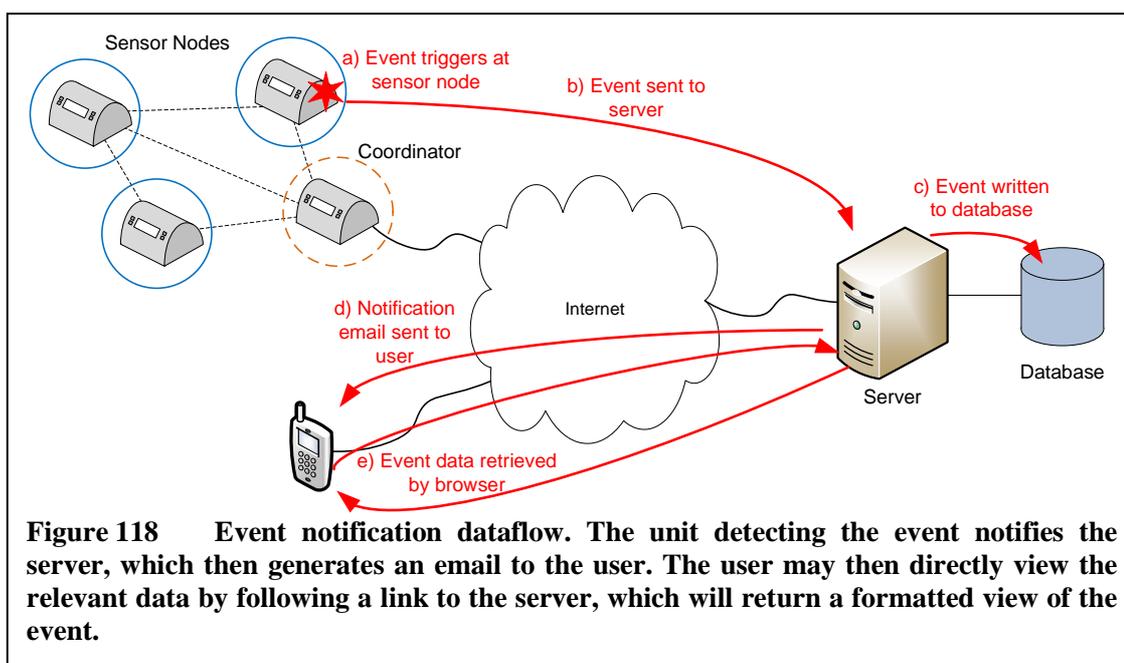
The event detection framework is meant to provide an interactive interface to the user. This interface notifies the user when an event of interest occurs, and should allow the user to participate in the event by investigating and labeling it in real-time. While some of this functionality has not been implemented, the design intent for all features is described.

User Notification

The end goal of event detection for the WAQM system is to alert a user of the system that something interesting has occurred. As such, user notification of an event is of prime importance. With the proliferation of smart phones, it is very common for a person to have a data connection into their personal device, allowing the receipt of emails and the viewing of web pages. Following this paradigm, the WAQM system implements an event notification path that takes advantage of these technologies. This allows events

to be pushed to the user in an email, which then provides a pointer to a web page where the event data can be used.

This flow of information is shown in Figure 118. The event trigger fires at the sensor node, causing a message to pass through the coordinator, across the Internet, and to the server. The server recognizes the event message and writes the event into the system database. The server then generates an email notification to a list of users that have been previously configured in the system. This email contains text indicating that the event has occurred, and a hypertext link that the server builds dynamically based on the event message. The user receives the email, either on a computer or a mobile device, and may then select the hypertext link to retrieve the data. This link will typically open a web browser on the user's device, which then accesses the server using parameters built into the link. These parameters cause the server to format a special page specifically for the event, containing a snapshot of the sensor data around the event. This is returned to



the user's device where it is displayed.

An example of data returned from an event message email is shown in Figure 119. In this example, the WAQM sensor node is located in a garage of a residential dwelling. When the garage door opens, the relatively high submicron particulate matter concentration causes a trigger on the small channel of the particle counter. The event causes the server to generate the notification email containing a link to the web page shown in the figure. The default format of the returned web page contains the sensor data for one hour before and after the event. In many cases, the link will be followed immediately after the event email is received, and in this case as much data as possible after the event is returned.

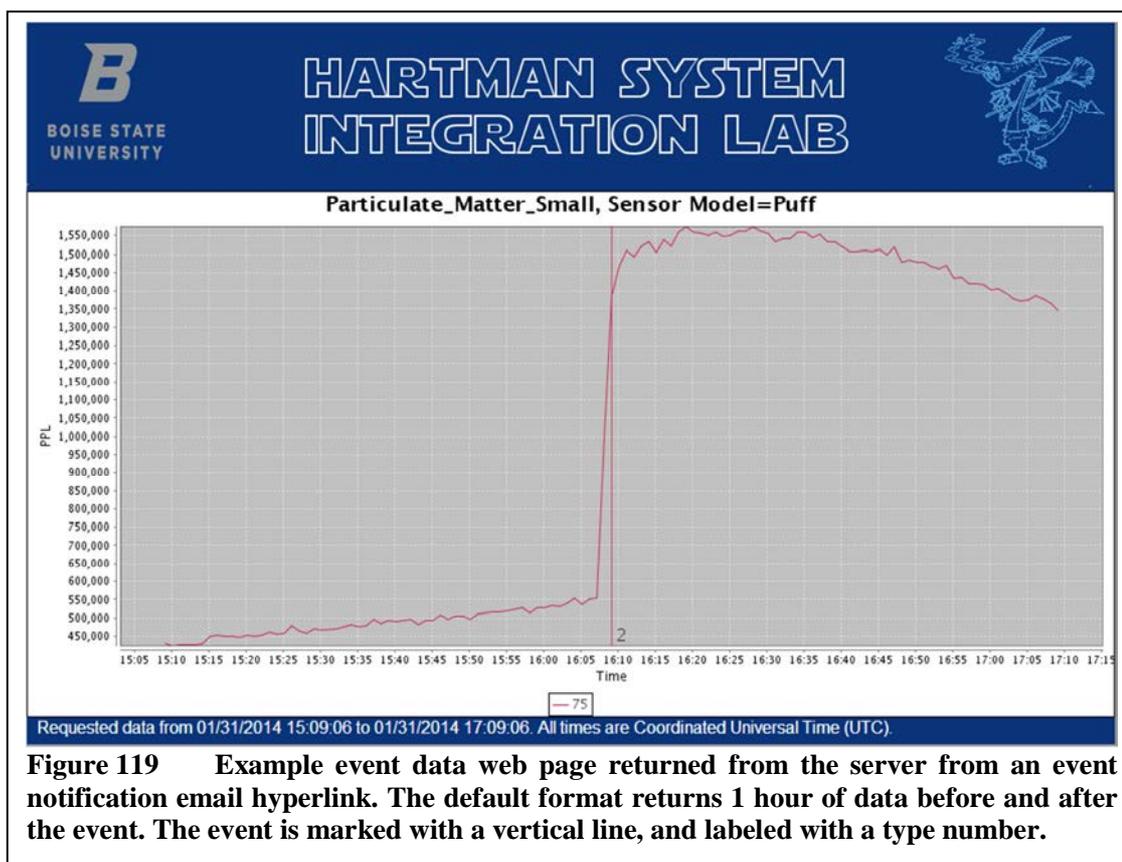


Figure 119 Example event data web page returned from the server from an event notification email hyperlink. The default format returns 1 hour of data before and after the event. The event is marked with a vertical line, and labeled with a type number.

The event is displayed along with an event type number, as can be seen near the bottom of the vertical line indicating the event in the figure. The number is meant to allow correlation between the event and the type of trigger that caused it. Currently, only a number is displayed, which is linked to the trigger in the sensor node. Future versions of the WAQM system may implement a more descriptive message, allowing easier correlation by the end user.

The user notification functionality in the WAQM system is spread between three different pieces of the server system. These pieces are described in the Wireless Sensor Network System chapter of this dissertation. The initial receipt of the event message and generation of the email notification is handled by the Connector Server. The event is stored by the WAQM Database. Finally, the user viewable data are generated by the Web Graphing Interface, which pulls the data from the WAQM Database to generate the data plots and event indicators as shown in Figure 119.

Event Labeling

Beyond user notification, the intent of the WAQM event detection framework is to allow users to tag events with a description that can be used to identify the event and possibly incorporate it into better event detection algorithms. This is important to reliable identification of events, as it can be difficult to keep track of correlation between events and actual circumstances manually. For instance, if the response method were inconvenient, a user may attempt to classify multiple events over an extended time period when they are in a convenient position to do so. The convenience factor can significantly increase the accuracy of reporting by giving the user immediate opportunity to quickly classify events.

The original goal for the WAQM event detection framework was to incorporate event labeling into the framework, which would give the user an opportunity to directly label the event through the web interface when the event message occurs. While event labeling has not yet been implemented, it is a fairly straightforward task to add it to the framework. The event table in the WAQM database is extendable to holding a description field. The Java servlet implementing the web graphing interface could return a field that would allow the user to enter a description that would then be stored in the event table of the database. This will be pursued in future versions of the systems developed by the HSIL.

Event Detection Conclusions

The event detection framework is a useful construct built on top of the WAQM system. This construct was originally envisioned to help a researcher who is setting up a new sensor network recognize and label events as they occur. The framework is then extensible to a learning algorithm that could modify triggers to better tune the system event notifications based on user feedback.

While the complete framework as envisioned was not implemented in full, even the parts that were implemented proved useful in field testing. The generation of event messages allows the researcher to focus on other activities while a test is underway, paying serious attention to the data only when alerted by the system. If the researcher does not have the patience or time to directly monitor the data, critical events may not be noticed until some later time when the data are perused. At that time, the event may have long since passed, and the opportunity to investigate the cause may have vanished.

Adding event detection alerts the researcher into action, allowing near immediate observation of the event as it is occurring.

Much more work could be done in this area to add value on top of the WAQM system. Intelligence added to the system will be somewhat dependent on the end application, and as the WAQM system is leveraged forward multiple opportunities will likely present themselves. With the event detection framework in place, it should be straightforward to tune system intelligence to fit a new scenario.

CHAPTER NINE: SYSTEM DEMONSTRATION

Two types of testing have been performed to demonstrate the OPC enabled system in action: controlled environment testing and field environment testing. The purpose of the controlled environment test is to show the capability of the system in tracking the motion of particulate matter through the use of a relatively dense set of sensor nodes. Field environment testing is meant to show the system operating in real-world environments, typically with just a few nodes located in a single site.

One of the more interesting applications of an OPC enabled WSN is the real-time tracking of particulate matter as it moves in an enclosed space. Though this type of experiment has been carried out before, to demonstrate the capabilities of the system an experiment was performed with unprecedented sensor density and previously unobtainable equipment cost. The large number of nodes also enabled the experiment was to be run with sensors arranged in two and three dimensional arrays. The experiment involved the use of sixteen sensor nodes operating in a controlled environment, in this case an aircraft cabin environment simulator. The simulated aircraft cabin provided for the release of particulate matter and carbon dioxide from the same emission point, allowing simultaneous substance tracking with the sensor array. The advantages of the battery-operated WSN are particularly useful, as a relatively large grid of devices can be set up and taken down quickly and with very little effort.

Several examples of field environment testing are presented to show the system in operation in real-world situations. The first example is a presentation of data taken from a typical operating room in a surgical hospital. The sensor units were deployed inside the operating room and in the adjoining hallway and substerile room. Other examples are provided of deployment in residential homes, showing the impact of daily activities on the air quality.

Controlled Environment Testing

Controlled environment testing was performed inside the Kansas State University (KSU) Boeing 767 mock-up cabin section. This environment provides several features that provide a good proving ground for the particulate tracking capabilities of the WAQM sensor network. First, the aircraft cabin has a dedicated ventilation system that continually flushes the space with outside air that is filtered, temperature controlled, and can be selectively dehumidified. The cabin has a dedicated system for injecting CO₂ into the cabin, and two different methods for creating particulate matter. The injection of gas and particulate occurs through the same single-point inlet in the cabin, allowing simultaneous observation of the movement of gas and particulate matter. The space is also relatively small, allowing a dense distribution of sensing points using a constrained population of sensor nodes.

For testing in the aircraft cabin, two different node arrangements were exercised. First the sensor nodes were set out in a two-dimensional array on the tops of the seat backs in the cabin. The goal of this test was to cover most of the cabin area with sensors to show large-scale movement of gas and particulate matter within the space. The second arrangement concentrated the nodes at two vertical levels across two rows near the front

of the cabin. This more dense arrangement attempts to show finer scale movement within a smaller area in three dimensions.

Simulated Aircraft Cabin

The KSU aircraft cabin is designed to simulate the interior conditions of a section of a Boeing 767. The cabin contains two aisles and eleven rows of seven seats arranged in a 2-3-2 configuration. The interior space is approximately 9.6 meters long by 4.7 meters wide by 2.0 meters high and is modeled in shape similar to the actual aircraft. Each seat in the cabin is occupied by a simulated human in the form of a mannequin. These mannequins include heating elements to mimic the body heat produced by an actual passenger. Figure 120 shows a view from the rear of the interior of the aircraft cabin with the WAQM sensor nodes in place on the seatbacks.

Ventilation for the aircraft cabin is provided by an air supply system that takes in air from outside the test facility. This air is first taken through HEPA filters to remove ambient particulate matter and a dehumidifier controls the amount of moisture in the air.



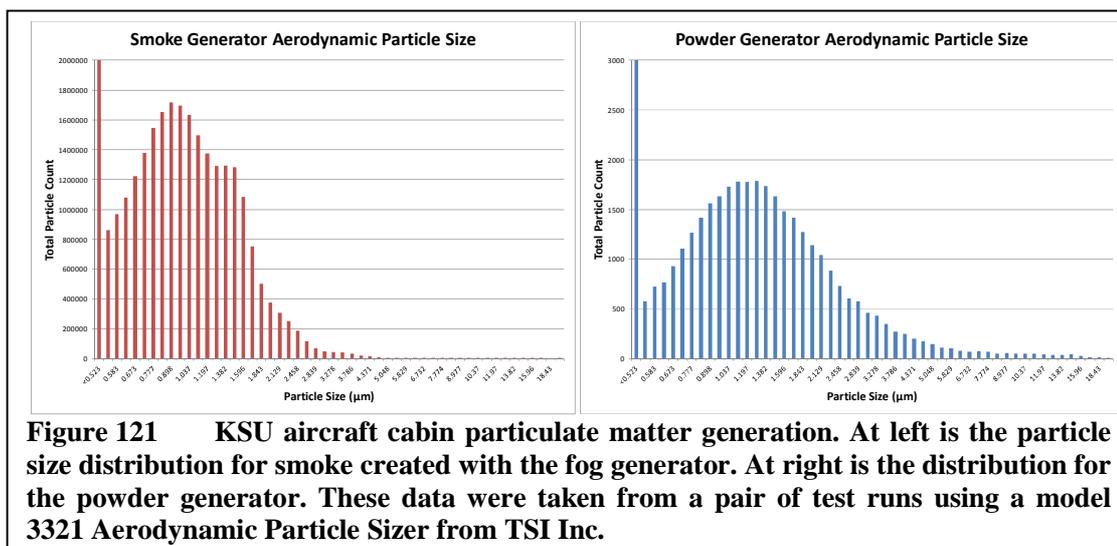
Figure 120 Interior of the aircraft cabin mockup with WAQM sensors in place. The blue LEDs on the sensors are visible on the seat headrests.

The air is also conditioned for temperature with heating and cooling elements. Air enters the cabin from a set of diffusers in the ceiling arranged above the center section of seats. Gaps along the floor on both sides of the cabin allow air to exit as it would on an actual aircraft. The air is not recirculated once it exits the cabin, and is instead replaced by fresh air from the ventilation system.

Two different types of substances were available to be released into the cabin: CO₂ and particulate matter. The CO₂ source was a cylinder of compressed gas with a regulated output. When active, the CO₂ source releases 7 liters/minute of undiluted gas. The CO₂ is piped into the aircraft cabin to a single point at the front of the cabin where it is then released into the air.

The KSU system has two different methods for generating particulate matter. Figure 121 shows the typical size distribution for each in terms of total count over a run. The first method generates smoke through a Chauvet Hurricane 1050 commercial fog generator using an aerosolized mixture of propylene glycol and glycerin. This produces a large amount of fine particulate under 2.5 um in diameter, as can be seen in the histogram on the right of Figure 121. The particulate from the fog generator is introduced into the cabin using the same pipe that carries the CO₂. During each run that it is used, the fog generator remains active for five minutes and is then shut down, allowing the particulate to clear from the cabin.

The second particulate generation method uses talcum powder and a series of seven air nozzles that disperse the powder from small containers placed at each seat across a single row. The talcum powder particulate is released in a single burst of air. Compared to the fog generator this method produces a much smaller total amount of



particulate matter with a larger distribution in sizes, with some particles reaching diameters greater than 10 μm . The histogram on the right side of Figure 121 shows a typical distribution for this particulate generation method.

Sensor Node Configuration

Sixteen sensor nodes and one coordinator node were used in the controlled environment testing in the aircraft cabin. Each of the sensor nodes was configured identically, with sensors set to report data according to Table 10. Note that some of the sensors in each node were not used, either because the particular substance was not of interest in this testing environment, or to conserve wireless network bandwidth. Particulate matter count was set to a relatively high sample rate to provide good temporal resolution. The CO_2 sensor was set to a longer period since that particular sensor could not meet a 2-second sample rate.

Table 10 Controlled environment testing sensor configuration.

Sensor	Sample Period
CO	Disabled
CO ₂	5 seconds
Humidity	5 seconds
Particulate concentration	60 seconds
Particulate count	2 seconds
Pressure	5 seconds
Radio-status	30 seconds
Performance	Disabled
Sound	Disabled
Temperature	5 seconds
Battery voltage	30 seconds

Bandwidth was of particular concern due to the number of nodes involved and the amount of data being sampled. The ZigBee WSN offers a relatively low-bandwidth connection, and overtaxing it could lead to data loss at the coordinator. A backup system of local node logging was in place in case this occurred, but in the end no significant data loss was experienced.

Each of the sensor nodes was powered from its internal lithium polymer battery during the tests. This greatly simplified the setup of the system, but caused some issues with data loss towards the end of testing. One unit in particular had a battery that performed much worse than the other units due to its age, causing the unit to power down before the end of testing.

Data Generation and Processing

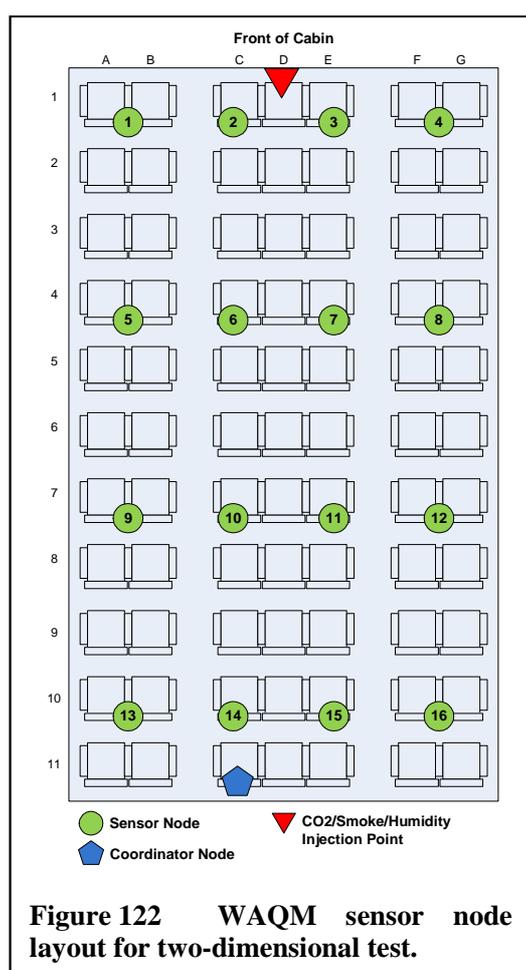
The data presented in this section have been collected using the BSU Sensor Monitor program and exported into a comma-delimited format that preserves the timestamp for each piece of sensor data. From there, the data was imported into MATLAB for further processing and plotting. Custom scripts with a graphical user interface were generated using the MATLAB GUIDE feature. This allowed the data to be viewed and plotted interactively. Though the contour data presented consist only of plots of single slices of time, the script created also allowed the export of a time sequence of contour plots in video form.

To create visually stimulating contour plots, the data were plotted on a logarithmic scale. This helped to bring out the low-level changes that would otherwise have been washed out on a linear scale. This was done to the particulate matter and CO₂ data used in the plots. The particulate matter data were also scaled during post-processing

to normalize the particle count data used for the contour plots. The scaling was required since the particle count data are output from the sensors as raw data, and the scaling is necessary to equalize the performance of the various nodes. The scaling used was identical to the calibration scaling used by the nodes on the particulate matter concentration sensor. Note that particulate matter concentration was not used directly since this sensor has a period of one minute, which did not meet the desire for fast sampling in this dynamic environment.

Two-Dimensional Testing

The first set of tests arranged nodes in a two-dimensional 4x4 array across the



body of the simulated cabin as shown in Figure 122. The goal of this arrangement was to cover as much of the cabin as possible at a level near the typical head-height of the passengers. Each node is shown as a green circle with the unit's position reference number indicated inside the circle. A set of four nodes was placed on the top of the seat-backs every three rows. Spacing between the units was approximately 130 cm laterally across the row and 250 cm between each row of units down the length of the cabin. Particulate matter and CO₂ were injected into

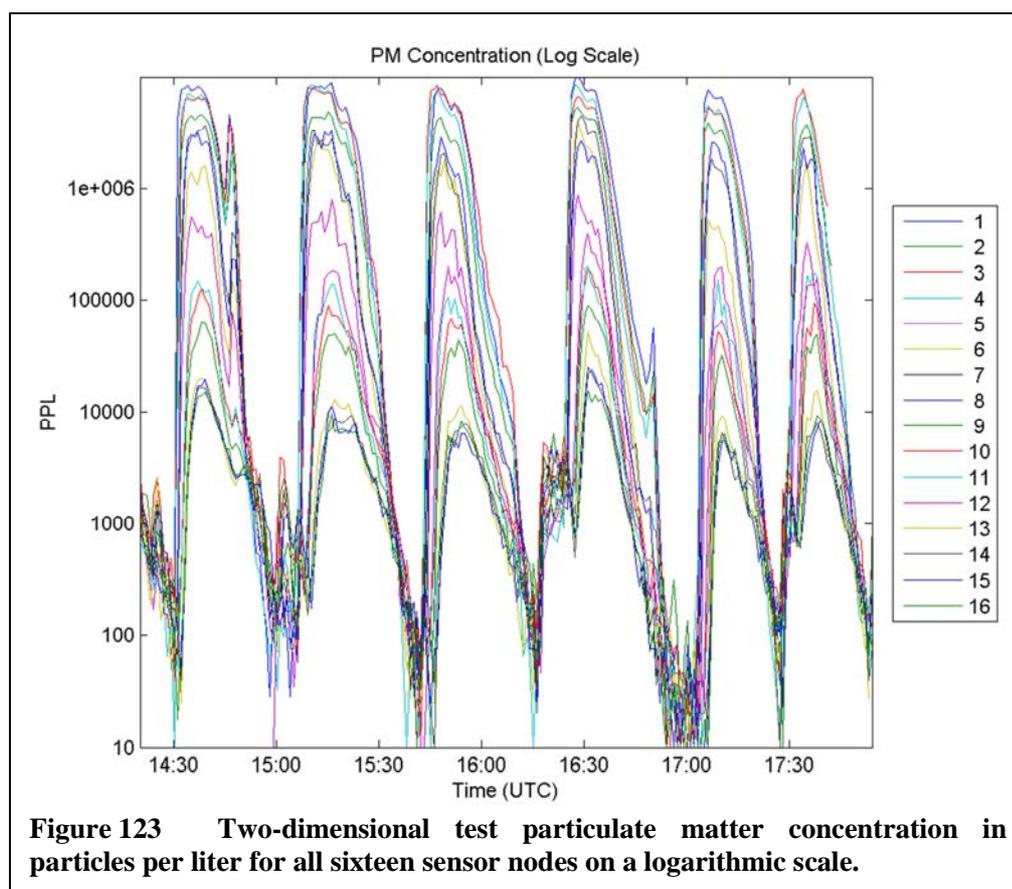
the cabin from a single substance injection point approximately 10 cm above the top of the seat backs as indicated by the red triangle in Figure 122. A coordinator collecting data for the WSN was located near the back of the cabin as indicated by the blue pentagon. A serial cable connected the coordinator node to a monitoring computer outside the cabin.

The CO₂ and particulate matter were released into the cabin concurrently six times over the course of three and one-half hours. Table 11 lists the tests conducted for the two-dimensional setup, with references for each of the six tests listed in the leftmost column of the table.

Reference	Variables	Log (UTC Time)	Comments
Run 0	Particulate (smoke) CO ₂ Dehumidifier active Humidifier inactive	14:30: Start CO ₂ , smoke release 14:47: Stop CO ₂	Initial test to verify sensor network formation and general operation. Position 12 (Unit 74) found to have malfunctioning particle counter.
Run 1	Particulate (smoke) CO ₂ Dehumidifier active Humidifier inactive	15:08: Start CO ₂ , smoke release 15:28: Stop CO ₂	First full run. Unit 74 replaced with Unit 79.
Run 2	Particulate (smoke) CO ₂ Dehumidifier active Humidifier inactive	15:43: Start CO ₂ , smoke release 16:09: Stop CO ₂ 16:11: Stop dehumidifier	Second full run.
Run 3	Particulate (smoke) CO ₂ Dehumidifier inactive Humidifier active	16:25: Start CO ₂ , smoke release, humidifier on 16:50: Stop CO ₂	Dehumidifier turned off. Humidifier appears to modify air currents around CO ₂ /particulate cabin input.
Run 4	Particulate (smoke) CO ₂ Dehumidifier inactive Humidifier active	17:03: Start CO ₂ , smoke release 17:18: Stop CO ₂	Repeat of Run 3 conditions. Humidifier remained active since the start of Run 3.
Run 5	Particulate (smoke) CO ₂ Dehumidifier inactive	17:26: Stop Humidifier 17:29: Start CO ₂ , smoke release 17:29: Battery died on Unit 60 17:45: Stop CO ₂	Same conditions as Run 4, with humidifier inactive. Confirm humidifier modifies CO ₂ , likely due to air current changes.

The first test, run 0, was performed as a check of the system operation. The sensor unit hardware had been shipped to the testing facility, and needed to be tested for correct operation after unpacking. During this run, it was discovered that the unit at position 12 had a malfunctioning particle counter and was replaced with a backup unit before continuing with run 1.

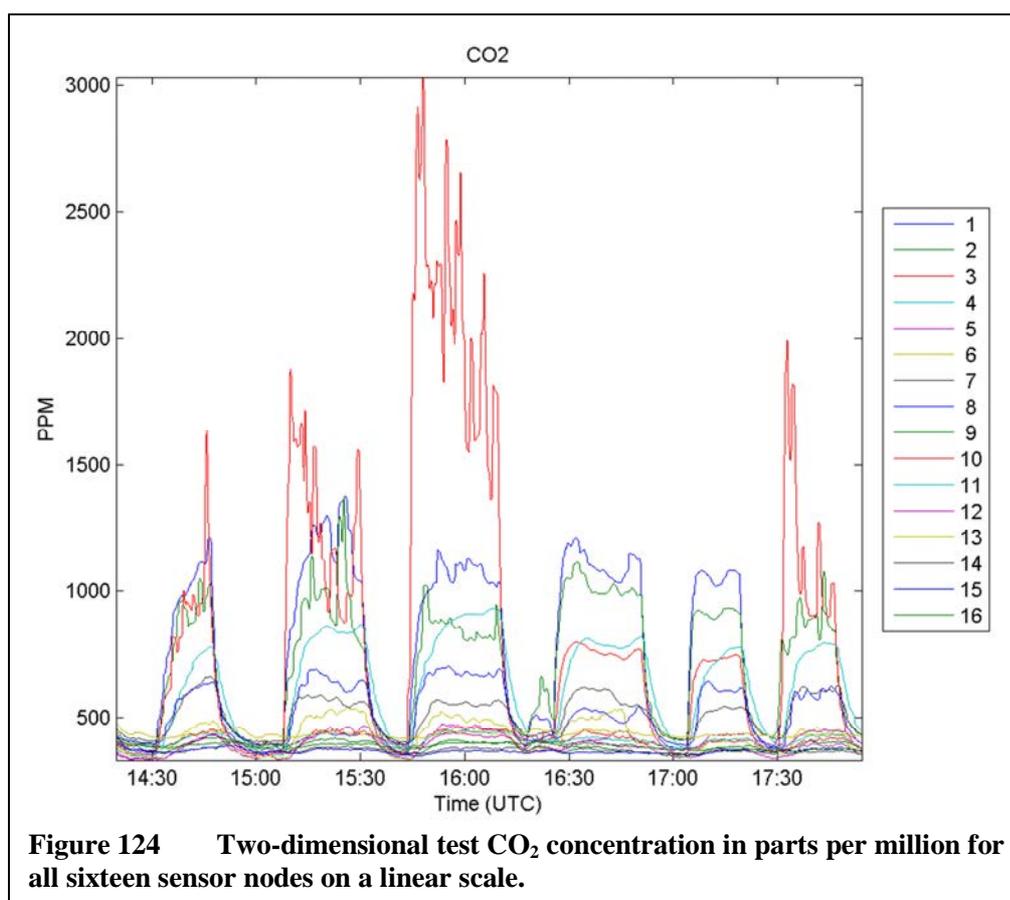
Figure 123 shows the time-series data for the particulate matter concentration plotted on a logarithmic scale. All sixteen units are shown concurrently to give an idea of the distribution seen across the different sensing positions. The concentrations tend to be higher towards the front of the aircraft cabin where the particulate matter is injected into the environment, and fall off moving towards the rear. The largest peaks in particulate



matter concentration approach 10 million particles per liter when smoke is being actively injected into the cabin, and fall off to 1000 particles per liter or lower in between testing.

Note that the particulate matter level between test runs tends to reflect activity in the cabin. For example, the humidifier was set up in the cabin at approximately 16:20 prior to run 3. The doors to the cabin were opened and several people entered, raising the particulate matter concentration as existing particulate was stirred up and new particulate entered through the doors. In contrast, no one entered the cabin in between run 3 and run 4 at approximately 17:00 hours. In this case, the particulate falls to a very low level as the particulate-laden air from testing is replaced by clean air from the ventilation ducts.

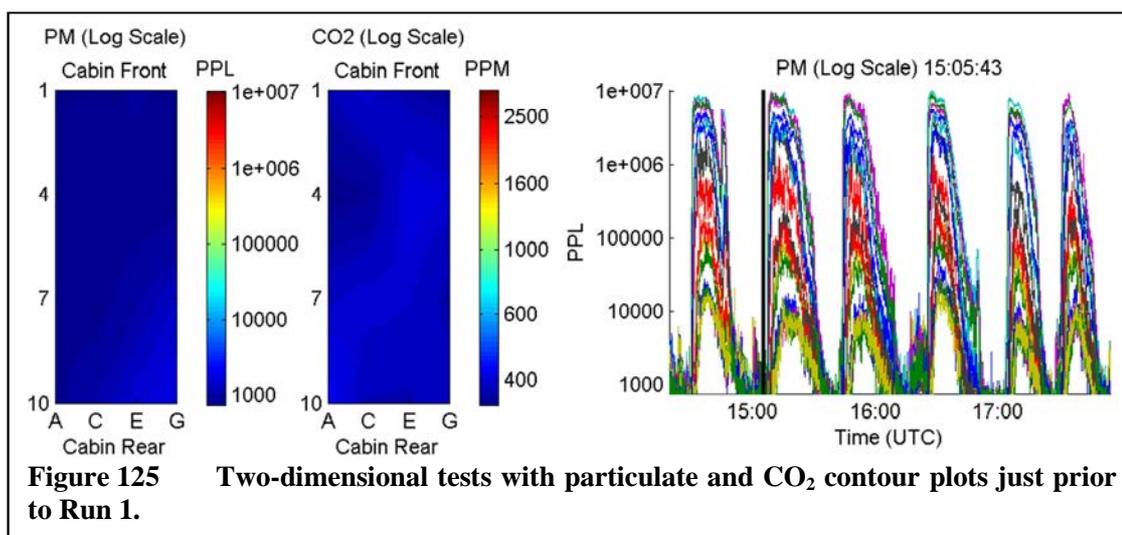
The time-series data for the CO₂ concentration are shown in Figure 124. As with



the particulate matter, concentrations tend to be most intense at the front of the cabin near the substance injection point and fall off toward the rear. Of particular interest is the behavior of the sensor node at position 3. This node was located at the front of the cabin very near to the substance injection point and apparently received high concentration doses of the gas as it flowed into the space. With the introduction of the humidifier at the front of the cabin prior to run 3, the CO₂ concentration seen by this node fell off drastically. This appeared to be due to the forced air from the humidifier shifting the air currents in the cabin, deflecting the high levels of CO₂ from this sensor node when running. This theory was tested by turning off the humidifier for run 5, which caused the high concentrations at position 3 to return.

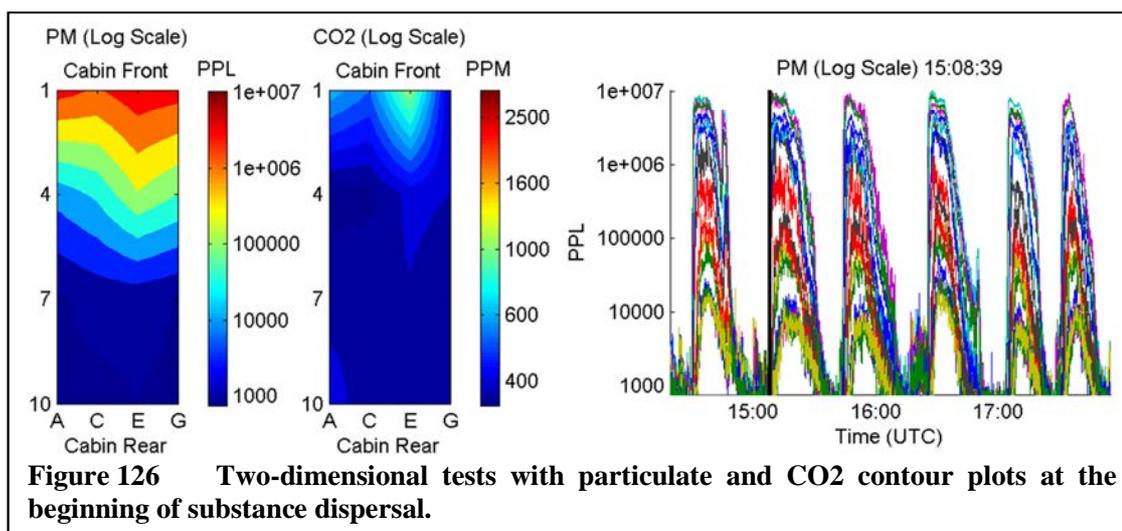
As noted in the discussion on the particulate matter time-series plot, at approximately 16:20 the humidifier was installed in the cabin. Two individuals were working on this activity for several minutes in the aisle between positions 2 and 3. In Figure 124, the CO₂ exhalations of the two individuals involved in this activity can be observed as the small peak in concentration at positions 2 and 3.

Looking at the contour data across the cabin provides further insight into the distribution and movement of the particulate matter and CO₂. Figure 125 shows two contour plots along with a time-series plot. The contour plot on the left shows the particulate matter count across the area of the aircraft cabin, with row numbers along the left and seat letters along the bottom. Similarly, the contour plot on the right shows the CO₂ concentration in the cabin. Both contour plots use a logarithmic scale to better highlight the concentrations across the entire range that was seen during testing. The two contour plots show the data from the time indicated with the vertical black line in the



time-series plot on the far right. This time-series plot of particle counts corresponds to the particulate matter concentration plot of Figure 123. The contour plots both show relatively low concentrations of particulate matter and CO₂ in the cabin just before the substances are released into the cabin.

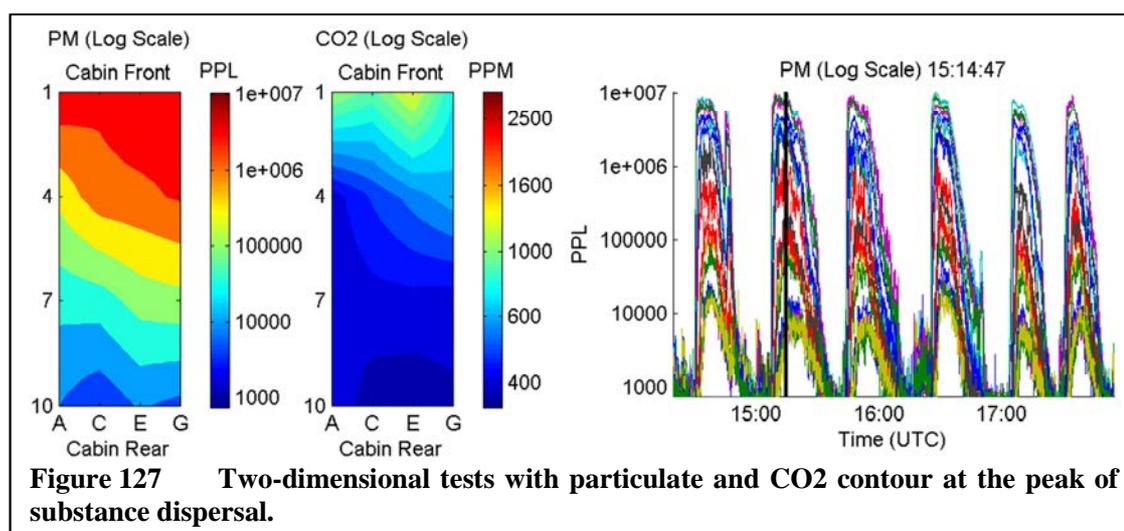
The time just after the substances are introduced to the cabin is shown in Figure 126. It is immediately apparent that there is a large difference between the spread of the particulate matter and the CO₂ concentrations. While this may appear to be significant, it may be due to a difference in response times of the two sensors. The particulate matter



sensor has a forced air system and will register an increase in particulate matter in as little as two seconds from the time it is pulled into the sensor's air intake. The CO₂ sensor on the other hand has a response time that is specified as being less than two minutes.

Figure 127 shows the concentrations at the peak of the substance dispersal during run 1. The particulate concentration is relatively high, and spreads down the length of the cabin. It is apparent that the cabin airflow moves the substances across the width of the cabin much more effectively than down the length. Even more than six minutes from the first appearance of the particulate matter in the cabin, it has still not equalized down the length. The evacuation of air from the vents at the sides of the cabin must be removing the particulate before it can spread. Similarly, the CO₂ concentration spreads across the width of the cabin much more strongly than down the length.

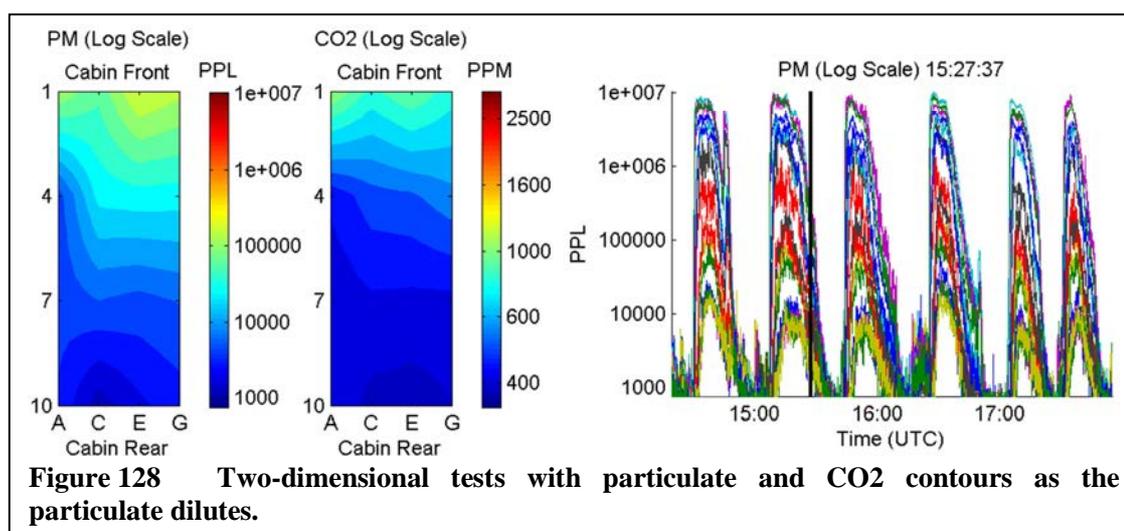
Comparing the concentrations of the two substances, one can see a similar shape to the flow down the length of the cabin. It appears that there is more movement of substance along the right side (seats F and G) compared to the left (seats A and B) towards the rear of the cabin. The correlation between the two different sensors may

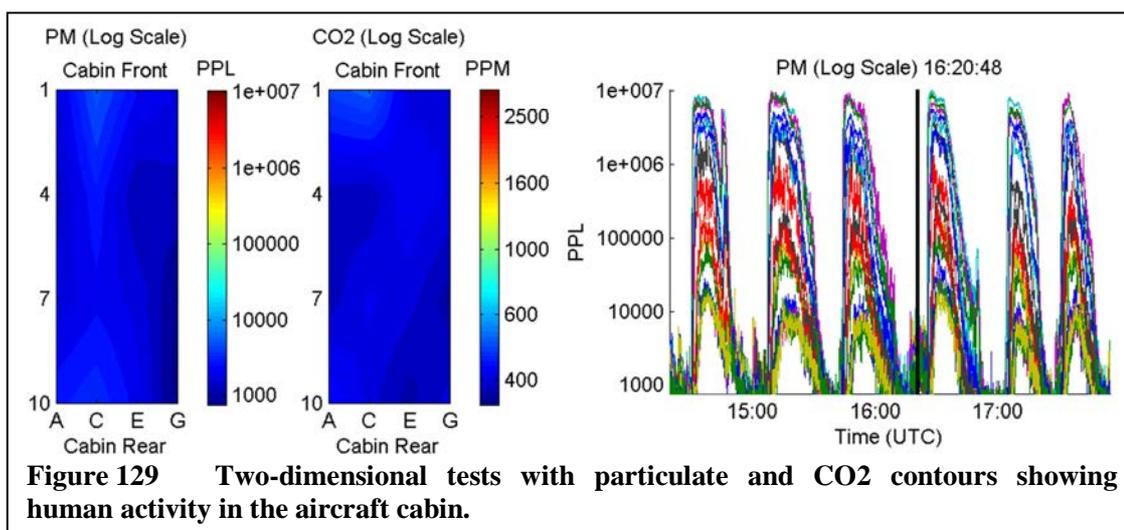


indicate an actual difference in airflow.

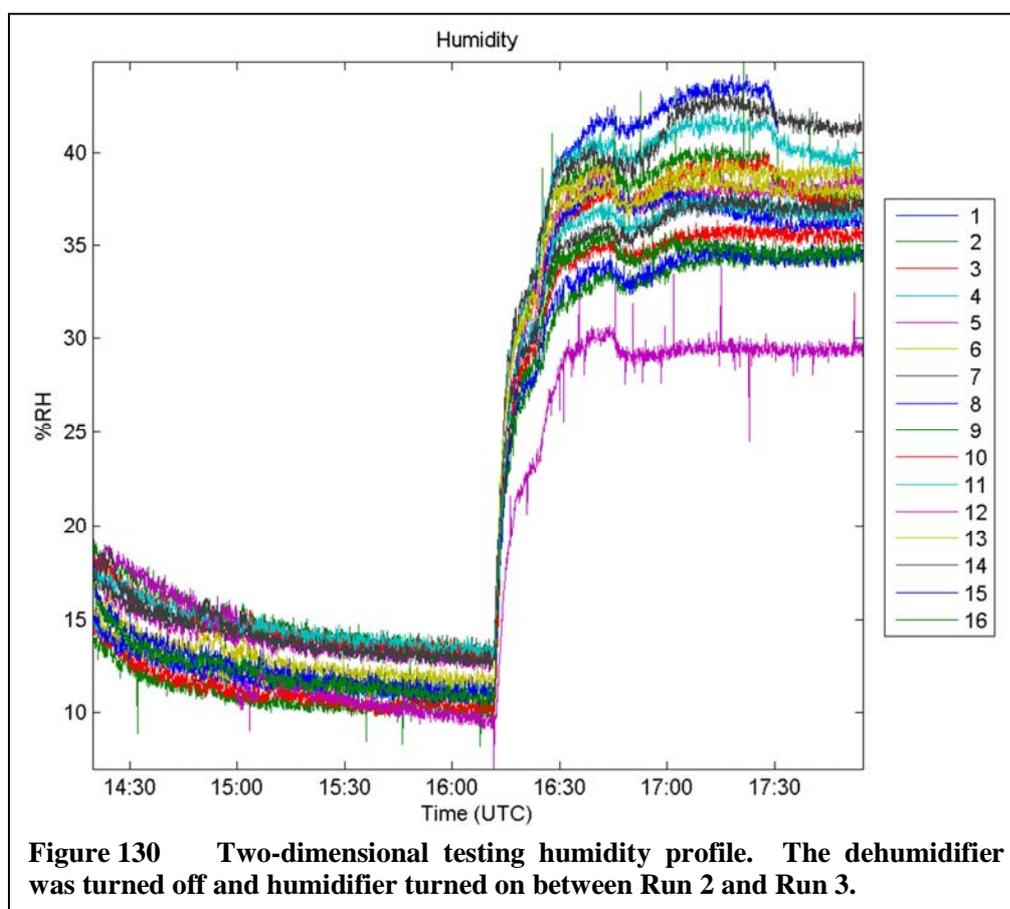
After a period of time, the smoke from the fog generator starts to dissipate and the CO₂ is turned off at the source. Figure 128 shows this condition at the end of run 1. Eventually, the substance is cleared from the cabin and conditions return to those similar to what is shown in Figure 125.

From the time-series plots of Figure 123 and Figure 124 can be seen evidence of human activity in the cabin between runs. This can also be seen in the contour plots from the same time frame. The contour plot on the left of Figure 129 shows the introduction of particulate matter along the aisle on the left side of the aircraft cabin as the individuals move up and down this aisle while bringing the humidifier into the cabin. Note that the concentration is much lower than an event caused by the fog generator. The individuals were active at the front-left of the aircraft cabin as they worked to activate the humidifier. The contour plot on the right of Figure 129 shows this as an increase in CO₂ at this location as the individuals exhale. As with the particulate concentration, the CO₂ increase caused by the individuals is much lower than that of the pure CO₂ injection.



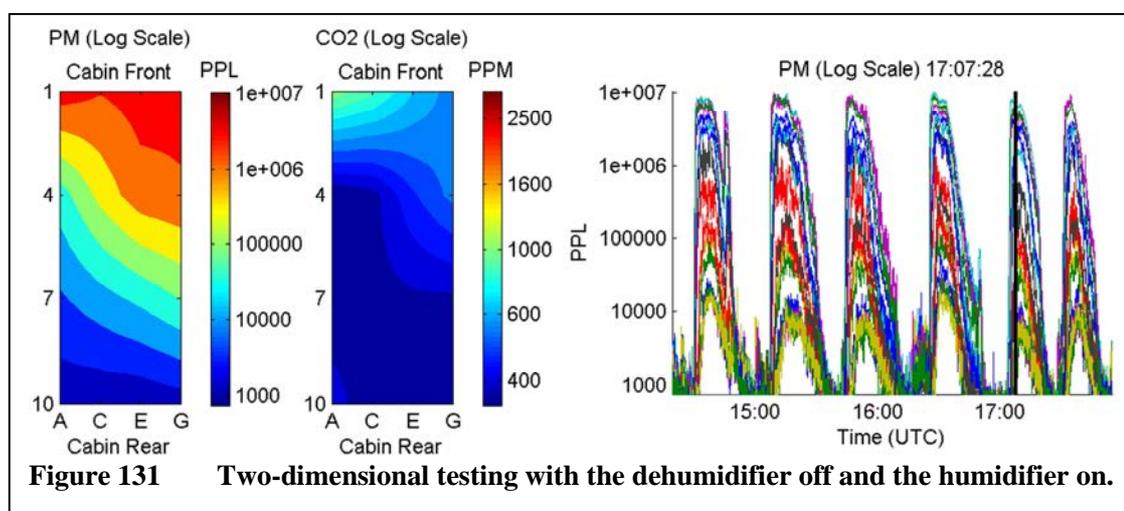


The impact of humidity on the substance movement in the cabin was also tested by turning off the dehumidifier in the ventilation system and adding a humidifier at the front of the cabin. Figure 130 shows the time-series plot of the humidity in the cabin. The



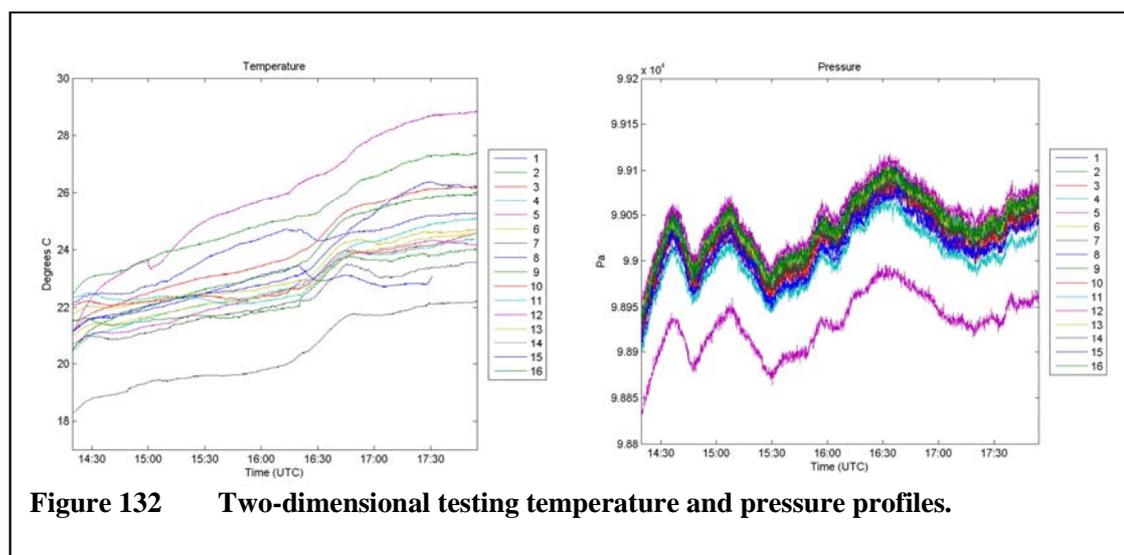
change in humidity is clearly visible, occurring between run 2 and run 3 when the humidifier was installed. Note that most of the humidity in the cabin occurred due to the deactivation of the dehumidifier. This can be seen prior to run 5 at approximately 17:30 when the humidifier was deactivated. There is only a slight drop in the humidity seen by some of the sensor nodes toward the front of the cabin, indicating that most of the humidity contribution is due to the moisture content of the ventilation air.

The biggest impact of the humidity change appears to be more due to the forced airflow of the humidifier than any impact of the moisture content of the air. This shows up mostly in the CO₂ concentration as it appears to remove the large peaks seen by the sensor node at position 3. This is presumably due to the change in airflow pushing the CO₂ away from the sensor at position 3. Figure 131 shows the contour data from run 4. The intense peak of CO₂ is now missing from position 3, but there does not appear to be much else that is different from runs with low humidity. There is still the same increase in concentration on the right side of the cabin in comparison to the left with both substances. Run 5, the final run of the test period, was executed with the humidifier



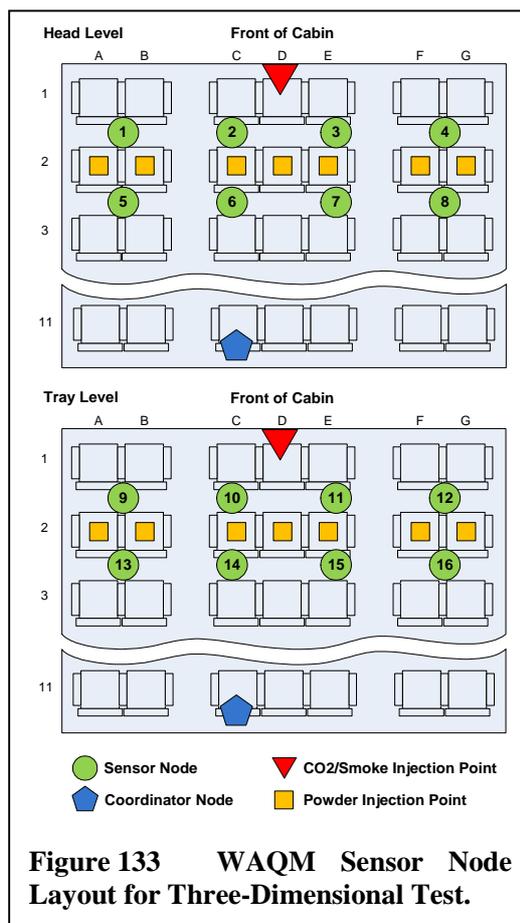
turned off to verify that the CO₂ peaks would return to the sensor node at position 3. This was the case, which can be most clearly seen in Figure 124.

During testing both temperature and pressure were measured by the WAQM sensor nodes. Figure 132 shows the temperature profile on the left and the pressure profile on the right. During testing the temperature steadily rose in the aircraft cabin. The rate of temperature change increases across most of the nodes at approximately 16:30, which appears to correlate with the change in the ventilation system dehumidifier settings. Since the cabin is not pressurized, the pressure profile reflects the ambient atmospheric pressure.



Three-Dimensional Testing

The second set of tests arranged sixteen nodes in a 4x2x2 three-dimensional array at the front of the cabin with two vertical layers covering two rows of seats. The goal of this arrangement was to cover an area near the substance injection point with a dense matrix of nodes in an attempt to look at the three-dimensional movement of particulate and CO₂ in the aircraft cabin. Figure 133 shows the layout of the sensor nodes in the cabin for this test, with the head-height layer at top and the tray-height layer at the bottom.



the bottom. The top layer of sensor nodes was suspended from the ceiling at a height above head level for a seated passenger and the bottom layer was placed on the tray table of each seat. Spacing between the units was approximately 130 cm laterally across the row, 80 cm between the two rows of units, and 80 cm between the two vertical layers. As with the two-dimensional test, CO₂ and particulate matter from a fog machine were injected at the front of the cabin. Additionally, the talcum powder dispersion system was tested that had injection points across the second row of seats in the cabin. The WSN coordinator node remained at the rear of the cabin, passing data to an external computer over a serial cable. A picture of the three-dimensional setup with sensors in place can be

seen in Figure 134. The head-level units are suspended from the ceiling of the cabin with string.

The CO₂ and particulate matter from the fog machine were released into the cabin concurrently



Figure 134 Aircraft cabin with three-dimensional setup in place.

three times during the first part of the testing period. During the latter portion of the testing period, the talcum powder dispersion system was used to release particulate matter into the cabin twice without the injection of CO₂. Table 12 lists the tests conducted for the three-dimensional setup, with references for each of the five tests listed in the leftmost column of the table. The dehumidifier was on for all tests except for run 4, and the humidifier remained off for the entire set.

Reference	Variables	Log (UTC Time)	Comments
Run 1	Particulate (smoke) CO ₂ Dehumidifier active Humidifier inactive	20:00: Start CO ₂ , smoke release 20:15: Stop CO ₂	First test of 3-D arrangement.
Run 2	Particulate (smoke) CO ₂ Dehumidifier active Humidifier inactive	20:28: Start CO ₂ , smoke release 20:45: Stop CO ₂	Second test of 3-D arrangement. Units at positions 11 and 16 swapped.
Run 3	Particulate (smoke) CO ₂ Dehumidifier active Humidifier inactive	20:55: Start CO ₂ , smoke release 21:14: Stop CO ₂	Third test of 3-D arrangement.
Run 4	Particulate (talcum) Dehumidifier inactive Humidifier inactive	21:25: Door open (powder load) 21:31: Door closed 21:33: Powder released	First talcum powder based particulate test.
Run 5	Particulate (talcum) Dehumidifier active Humidifier inactive	21:45: Door open (powder load) 21:49: Door closed 21:09: Powder released	Second talcum powder based particulate test.

Figure 135 shows the time-series plot of the particulate matter concentration for the three-dimensional testing plotted on a logarithmic scale. The first three large peaks correspond to runs 1, 2, and 3 in Table 12, in which a fog generator was used for the injection of particulate matter. These tests exhibit curves similar to what was seen in the two-dimensional testing. The last three large peaks correspond to testing with talcum powder. Note that the first of these talcum powder peaks was a demonstration using just a single of the seven talcum powder dispersal locations, and is not considered a formal test. The remaining two peaks correspond to runs 4 and 5 in Table 12. Note that the fog generator creates peaks in particulate matter that are nearly two orders of magnitude higher in concentration than the talcum powder dispersal system. The talcum powder peaks are also much shorter in duration, since they are released in a burst and have no

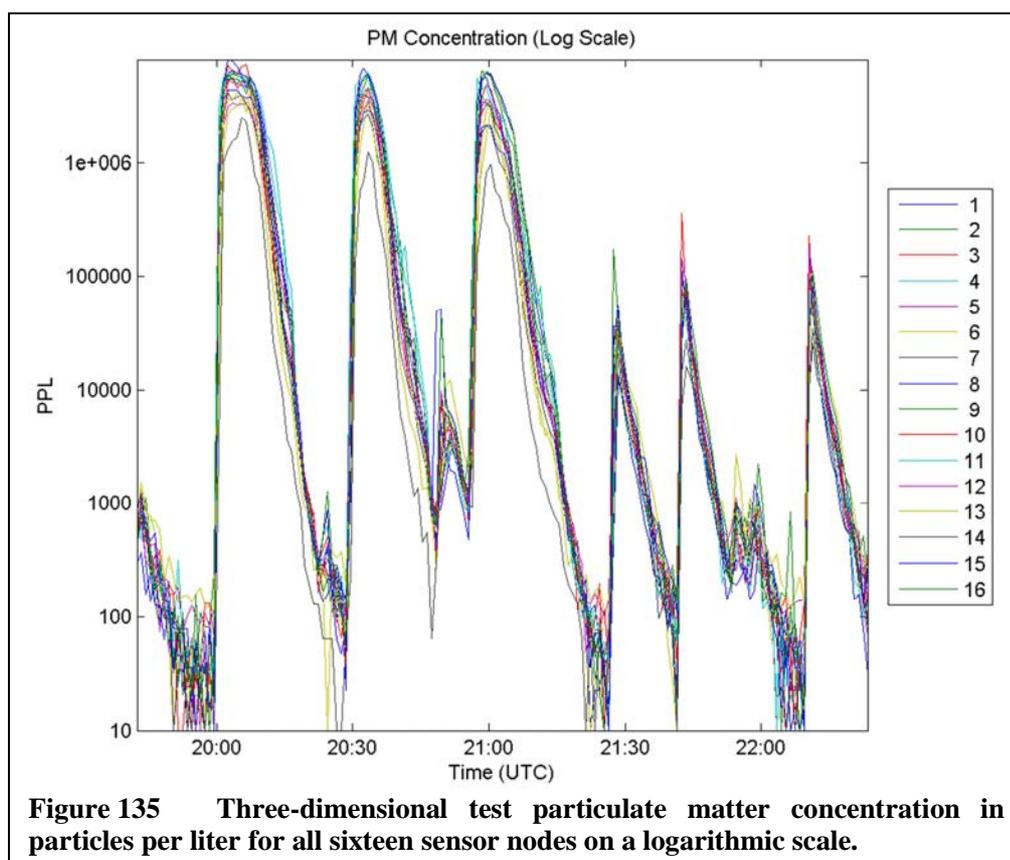


Figure 135 Three-dimensional test particulate matter concentration in particles per liter for all sixteen sensor nodes on a logarithmic scale.

sustained source of generating material.

Compared with the two-dimensional testing, the particulate matter concentrations for the three-dimensional tests are much closer in magnitude across the set of sensor nodes. This is likely due to the close proximity of the nodes in the cabin for the three dimensional testing, especially in the direction of the axis of the airplane. Much of the variation between sensor nodes in the two-dimensional tests came from the change in concentration down the length of the cabin. Since the layout of the nodes in the three-dimensional test only covers two rows in this direction, one might expect that the nodes would observe a smaller difference in concentration.

The time-series data for the CO₂ concentration are shown in Figure 136. As with

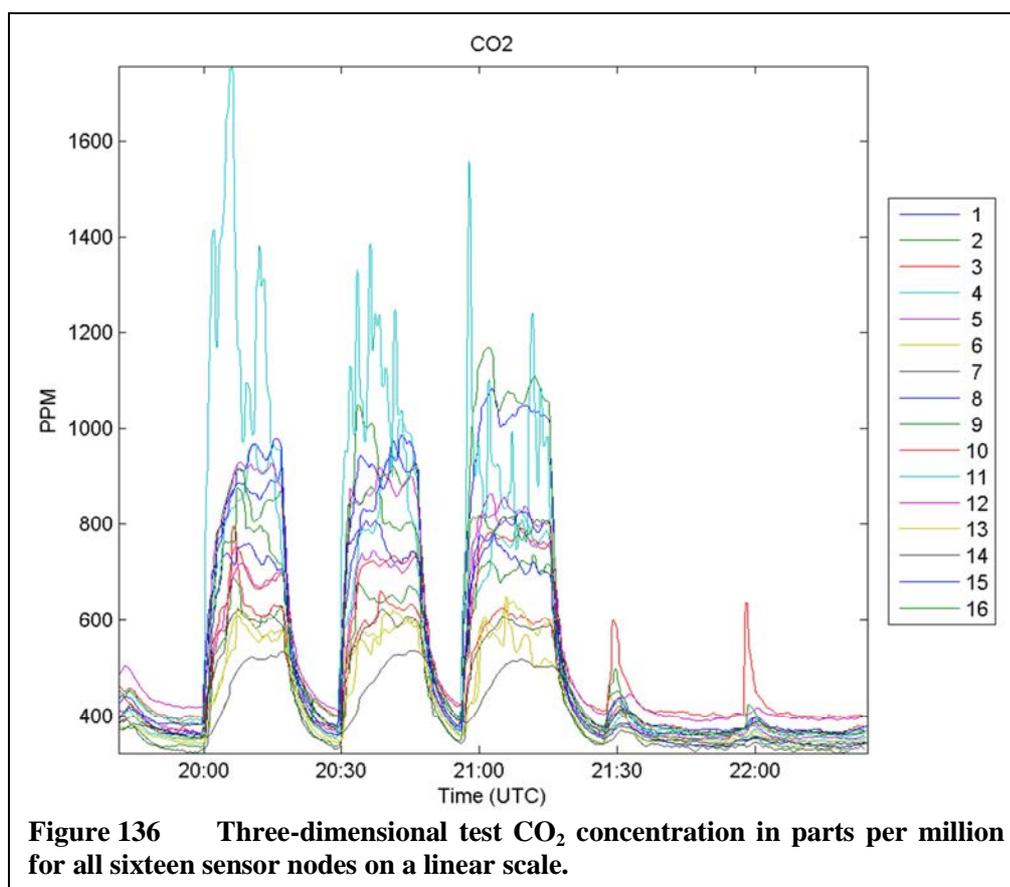


Figure 136 Three-dimensional test CO₂ concentration in parts per million for all sixteen sensor nodes on a linear scale.

the particulate data, there is less of a distribution of concentrations across the sensor nodes, likely due to the dense clustering near to the substance injection point. The node at position 11 exhibits very high spikes of CO₂ during the first three test runs. This is similar to what was seen with the node at position 3 during the two-dimensional testing, and corresponds to roughly the same position though at the tray table level rather than at the top of the seat back. To make sure that this was not a phenomenon specific to the sensor node at this point, the sensors at positions 11 and 15 were swapped after run 1. The high concentration peaks followed the position and not the specific sensor node, verifying that this was likely due to proximity to the substance injection point.

The two smaller peaks in CO₂ concentration at approximately 21:30 and 22:00 were due to human activity in the aircraft cabin, as no CO₂ was released during the talcum powder testing. The highest concentrations came from positions 2 and 10 at the tray table and head height units at the front of the left aisle in the cabin. This corresponds to locations where individuals were working in the cabin in between test runs.

The contour data for the three-dimensional testing are somewhat challenging to present. Figure 137 shows a set of four contour plots and two time-series plots. The plots on the left side of the figure show the particulate matter concentration, and those on the right show the CO₂ concentration. The two contour plots at the top of the figure show the particulate matter and CO₂ concentrations for head-level sensors, and the two contour plots in the middle show concentrations for the tray-level sensors. The black vertical line on each of the time-series plots shows the point in time from which the contour plot data are taken. As with the two-dimensional contour data, the plots use a logarithmic scale to better highlight the concentrations across the entire range that was seen during testing.

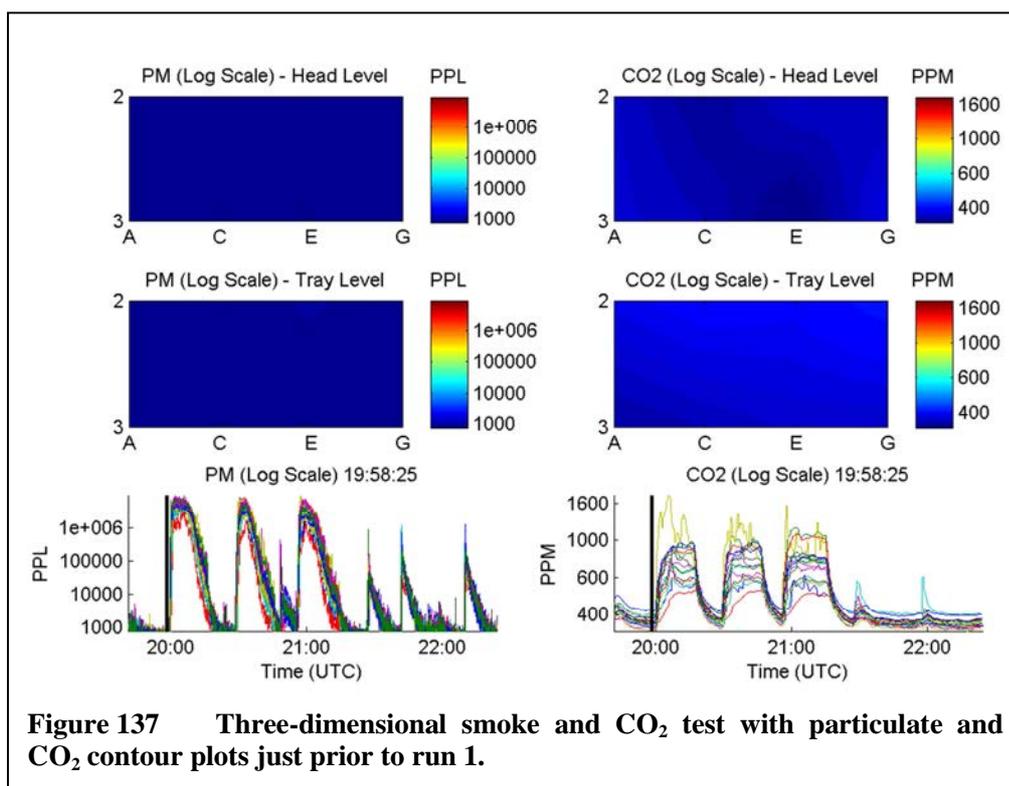
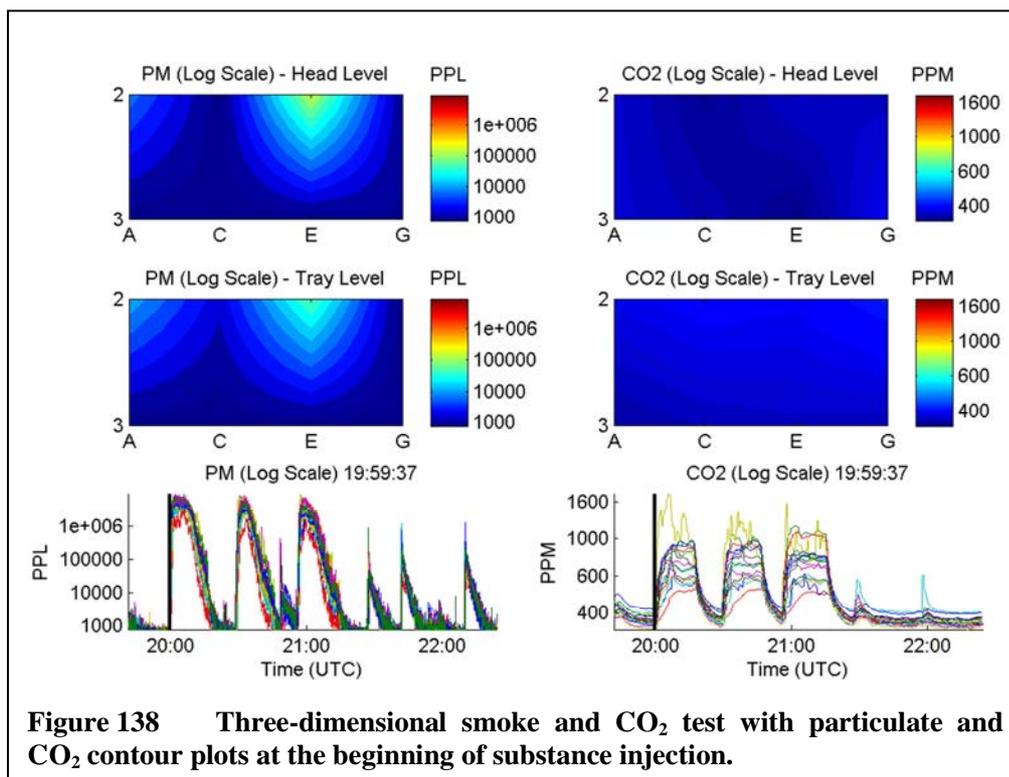


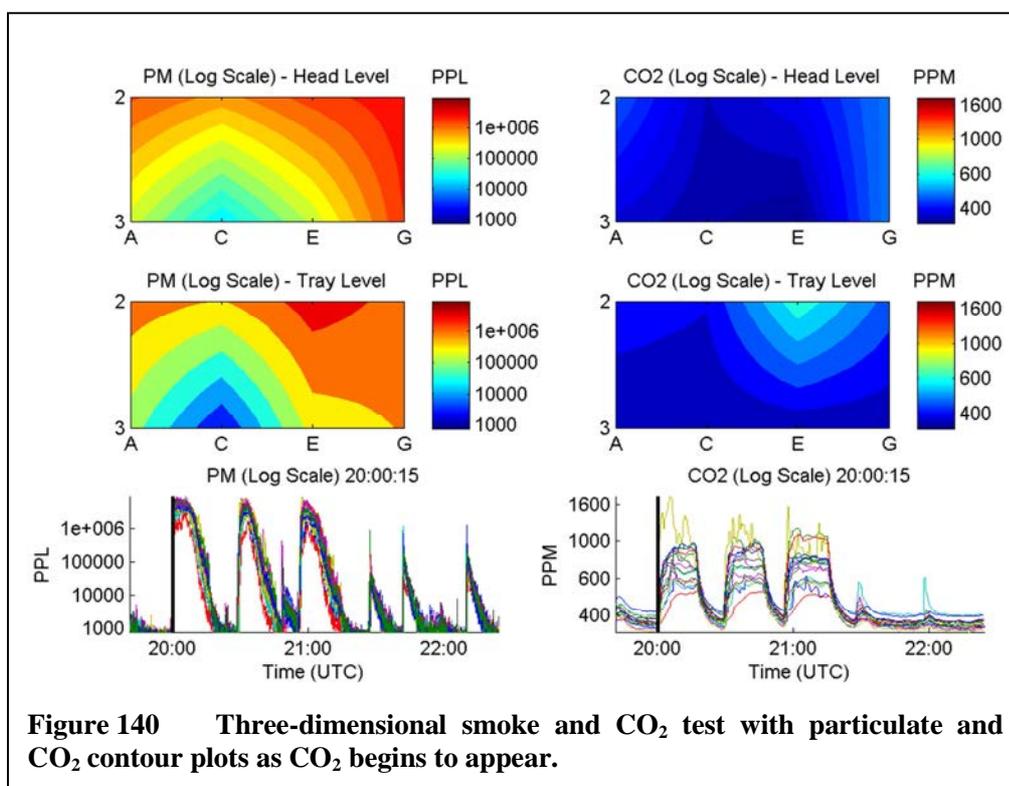
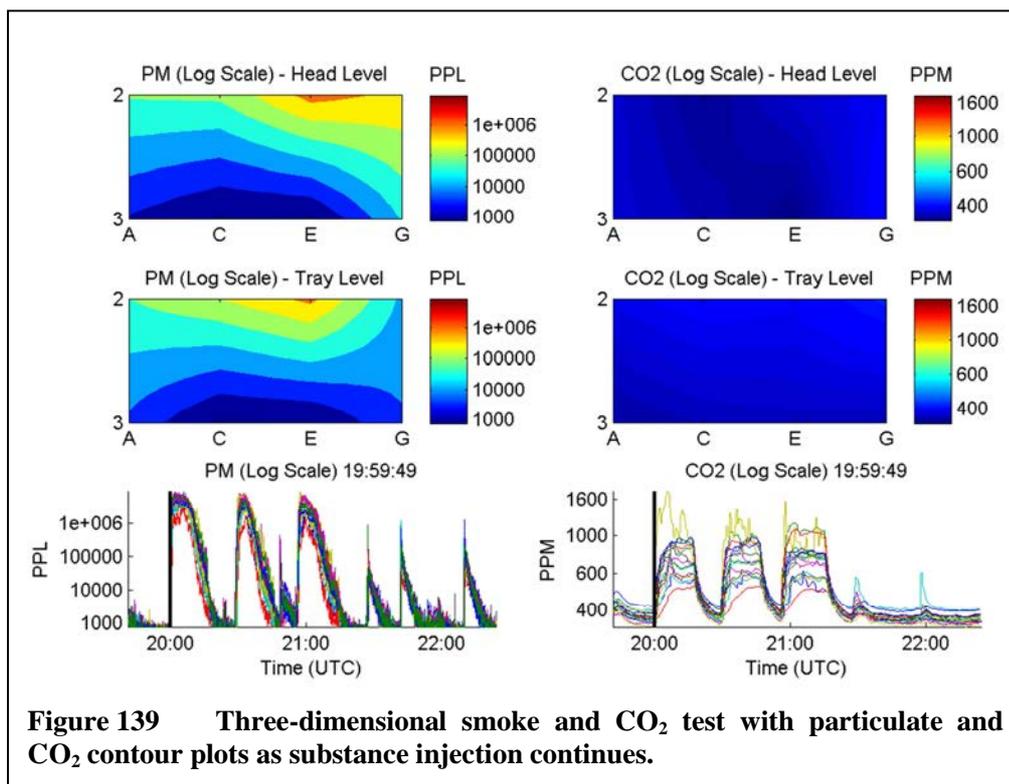
Figure 137 shows the concentrations of particulate and CO₂ just before the start of the substance release for run 1. Both concentrations are relatively low, as the cabin ventilation system has cleared most of the particulate matter and equalized most of the peaks in CO₂.

The start of substance injection can be seen in Figure 138. As with the two-dimensional testing, the particulate matter substance begins to appear earlier than the CO₂. This is likely due to sensor differences as explained above. The particulate matter first appears at a sensor position nearest to the injection point at the front-center of the cabin. It is not clear why there is some initial response in particulate matter from the two sensors at the front-left of the cabin, but this appears to happen in runs 2 and 3 as well. It is possible that the airflow in the cabin is forcing some particulate into this corner early in the cycle, bypassing the sensors immediately to the left of the outlet.



The increase in particulate matter substance shows some interesting spatial trends early in the run. Figure 139 shows the particulate building up to higher concentrations toward the front of the cabin at head and tray levels. The concentration has spread more quickly to the sides than across the seats into the third row. Also, the substance appears to have moved further to the right side of the cabin at the head-level, yet stays more concentrated around the injection point at the tray-level. This might be due to the ventilation inlet diffusers at the centerline of the ceiling pushing the particulate outwards nearer to the ceiling. The sensors are still not detecting the increase in CO₂ at this point in time.

Figure 140 shows that as the test run progresses, the sensors begin to register the increase in CO₂ concentrations in the cabin. The particulate matter has built up to



relatively high concentrations in the cabin, tending to move more strongly to the right side of the cabin than the left, and spreading laterally more quickly than down the length of the cabin. This increased concentration down the right side of the cabin matches what was seen across a larger area in the two dimensional testing. Also notable is the larger difference between highest and lowest concentrations at the tray-level in comparison to head-level. This may be due to the seats inhibiting the airflow at the tray-level. The CO₂ concentration does not build at the head-level nodes in the same way that is seen with particulate matter when it first appears. The sensor at position 3 that saw an early peak in particulate matter does not register much of an increase in CO₂. This could be due to the lack of forced airflow in the CO₂ sensor, or even the differences in height between the two sensor types within the sensor node.

Figure 141 shows a view of the cabin at the peak of substance injection for run 1.

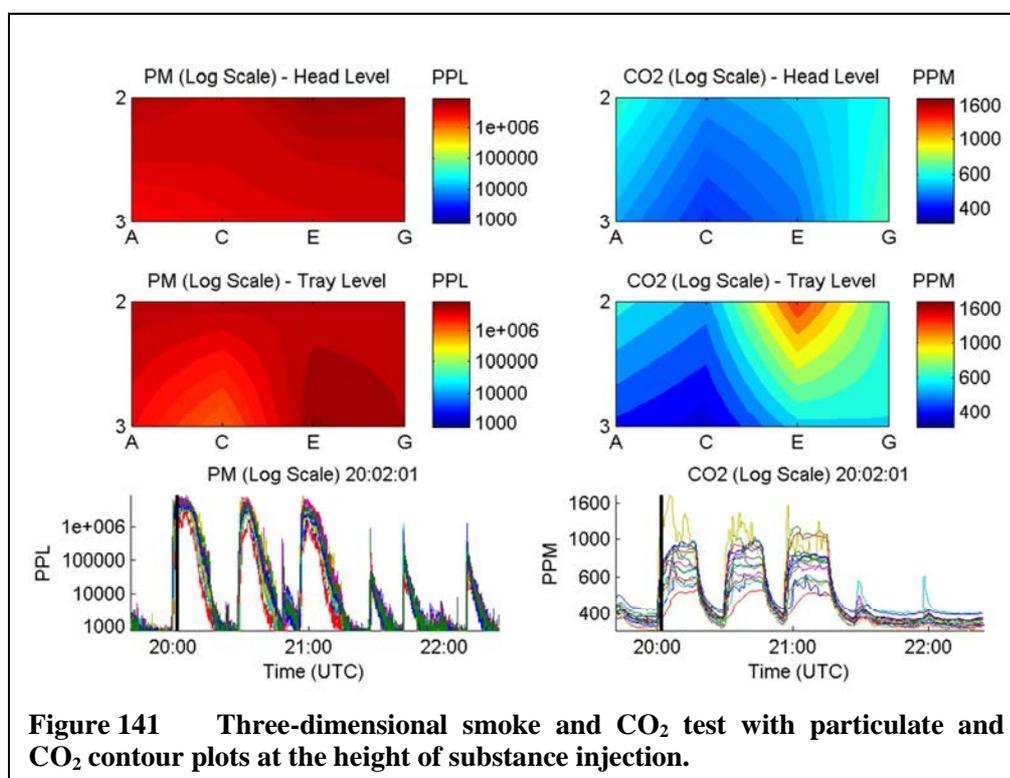
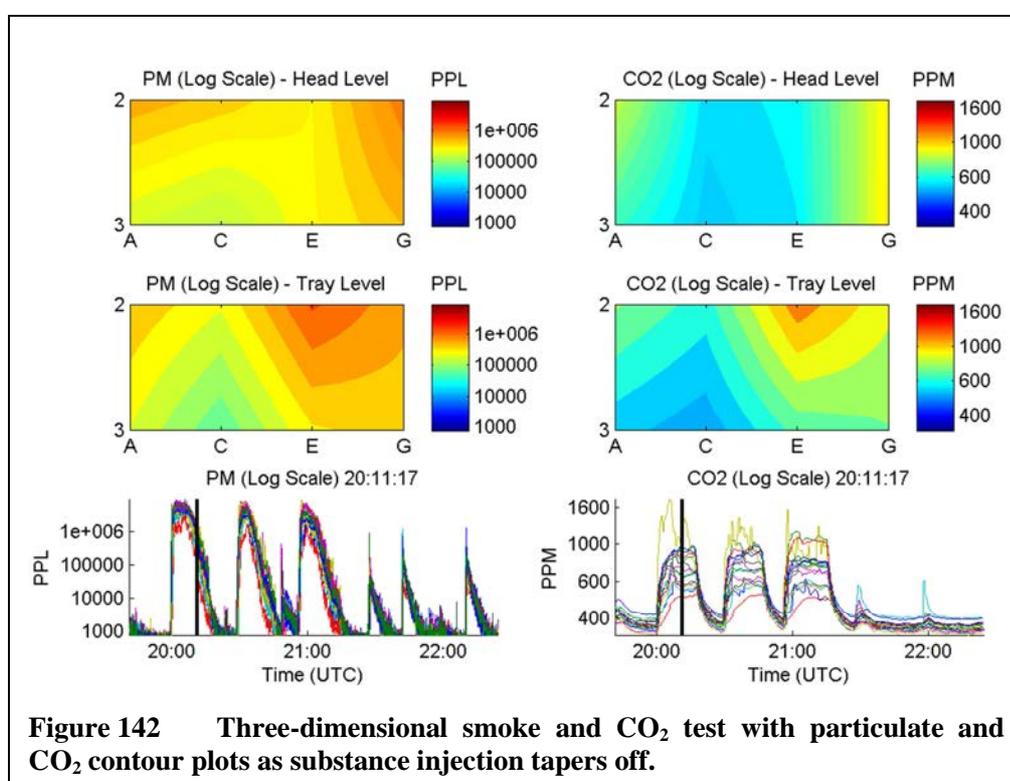


Figure 141 Three-dimensional smoke and CO₂ test with particulate and CO₂ contour plots at the height of substance injection.

The particulate matter has spread through the cabin, with marked differences between the left and right sides of the cabin. The concentration differences appear larger at the tray-level than at head-level. The minimum in particulate matter at position 14 in row 3 on the left side of the cabin appears to be similar to the minimum seen in the CO₂ data. While the large peaks in CO₂ at position 11 tend to push the other contours down in scale, there does still appear to be a concentration that is more intense along the right side of the cabin. The CO₂ at head-level does appear to concentrate more toward the sides of the cabin, with the contour lines running closer to parallel to the cabin centerline as opposed to what is seen with the particulate matter.

Figure 142 shows data from a point in time after the particulate matter concentration has peaked and is starting to be cleared out of the cabin by the ventilation system. The concentration of this substance appears to move away from the centerline,



with the exception of a local maximum near the substance injection point. This peak may be due to residual smoke flushing from the injection system, or due to air being constrained by the seats in the cabin. The CO₂ shows a trend at head-level to move toward the sides of the cabin away from the centerline. The largest peak remains at tray-level near the injection point, with a minimum that matches the particulate matter on the left side of the cabin at position 14. At this point in time, there is a fairly good match in concentration distribution at the tray-level for particulate matter and CO₂. This may suggest that at this level the airflow constraints caused by the cabin seating may be overcoming any differences between the sensor types within each node.

Figure 143 shows the tail end of the run, with smoke and CO₂ injections complete. The peaks in both substances near position 11 have disappeared with particulate matter being significantly down across the entire measurement area. The CO₂

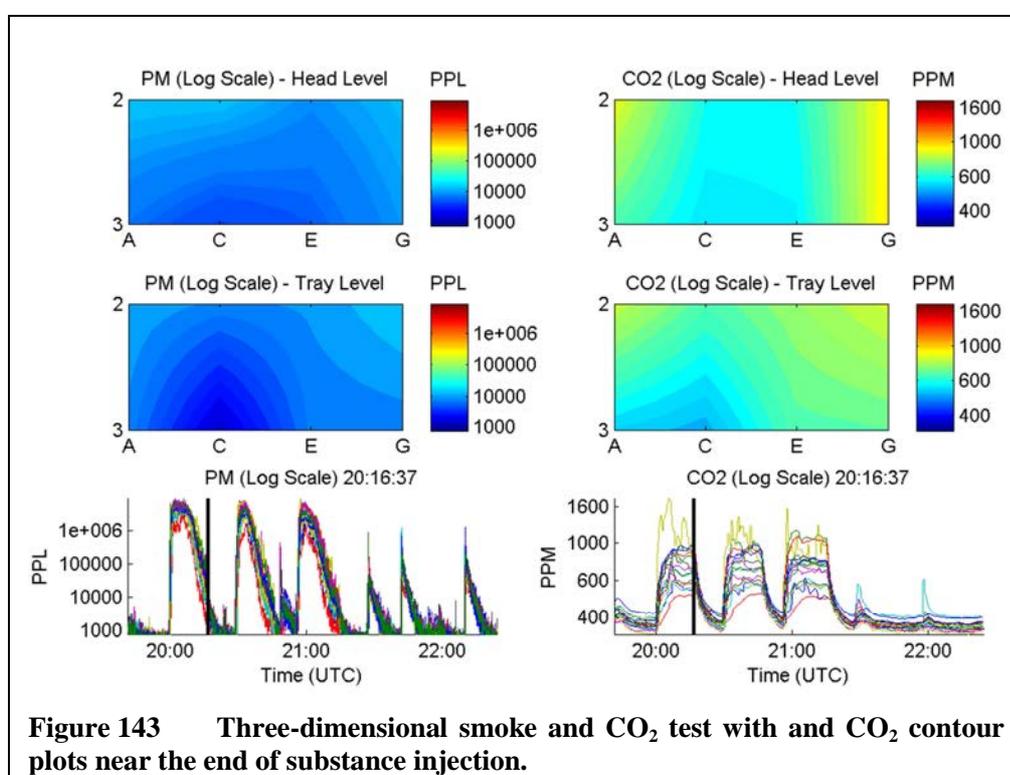


Figure 143 Three-dimensional smoke and CO₂ test with and CO₂ contour plots near the end of substance injection.

concentration is still relatively high at head-level near the sides of the cabin, having apparently been pushed there from the centerline by the ventilation diffusers. The minimum in particulate and CO₂ is still visible at position 14. Eventually the concentrations of both substances return to levels similar to what is seen in Figure 137, with clean ventilation air pushing out the substances. Similar results were seen with runs 2 and 3.

The KSU aircraft cabin simulator included the ability to test particulate matter using a talcum powder injection system that released particulate in a very short burst. The last two runs of the three-dimensional testing used this method for particulate matter injection without the use of CO₂. The injection points consisted of 7 nozzles across row 2 of the cabin, with one nozzle centered in each seat at a level just above the seat armrests. Figure 144 shows the start of run 4, which is the first of the two talcum powder tests. The particulate matter first appears at the sensor nodes in row 2, which is where the talcum powder injection nozzles are located. It appears that for both talcum powder runs the sensors at positions 3 and 10 were the first to pick up the increase in particulate. The substance then moves over the seats to the sensors at head-level along row 3. This can be seen in Figure 145 where the concentrations are relatively high at head level along row 3 but remain lower at the tray level. The minimum appears at position 14 as it did when testing with smoke. Note that CO₂ never increases in the measurement area, since it is not injected into the cabin during the talcum powder tests.

After the initial injection of particulate, the concentrations rapidly spread and begin to be removed by the ventilation system. Figure 146 shows the contour data as the concentrations abate. The particulate does appear to move outwards from the center to

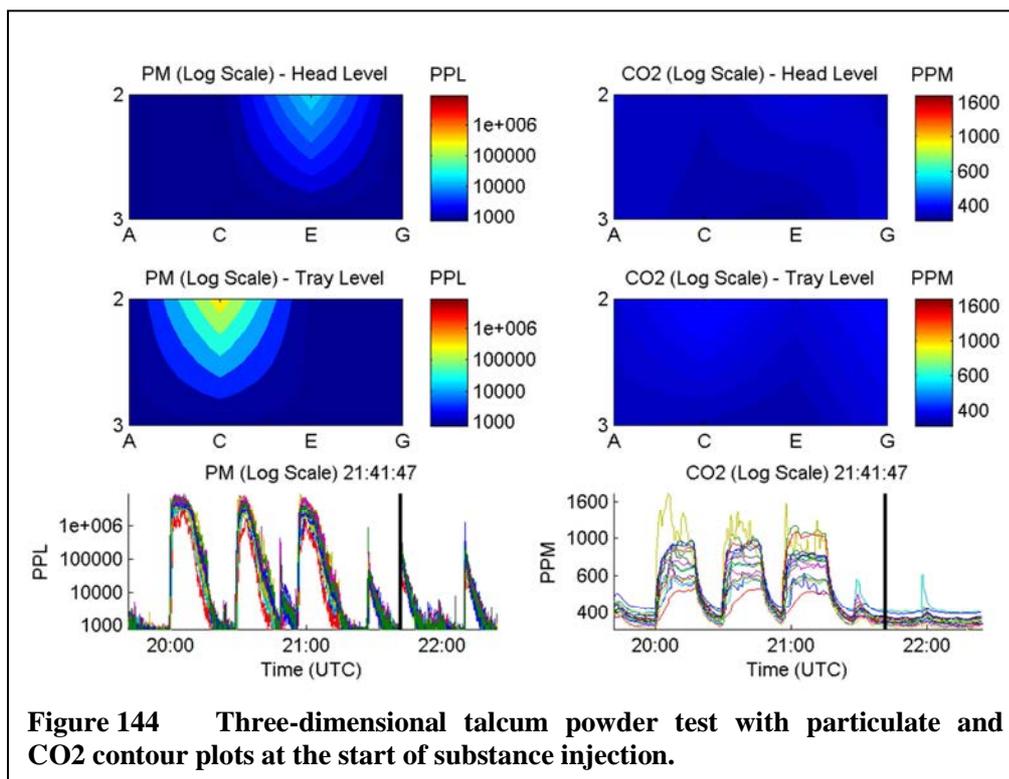


Figure 144 Three-dimensional talcum powder test with particulate and CO₂ contour plots at the start of substance injection.

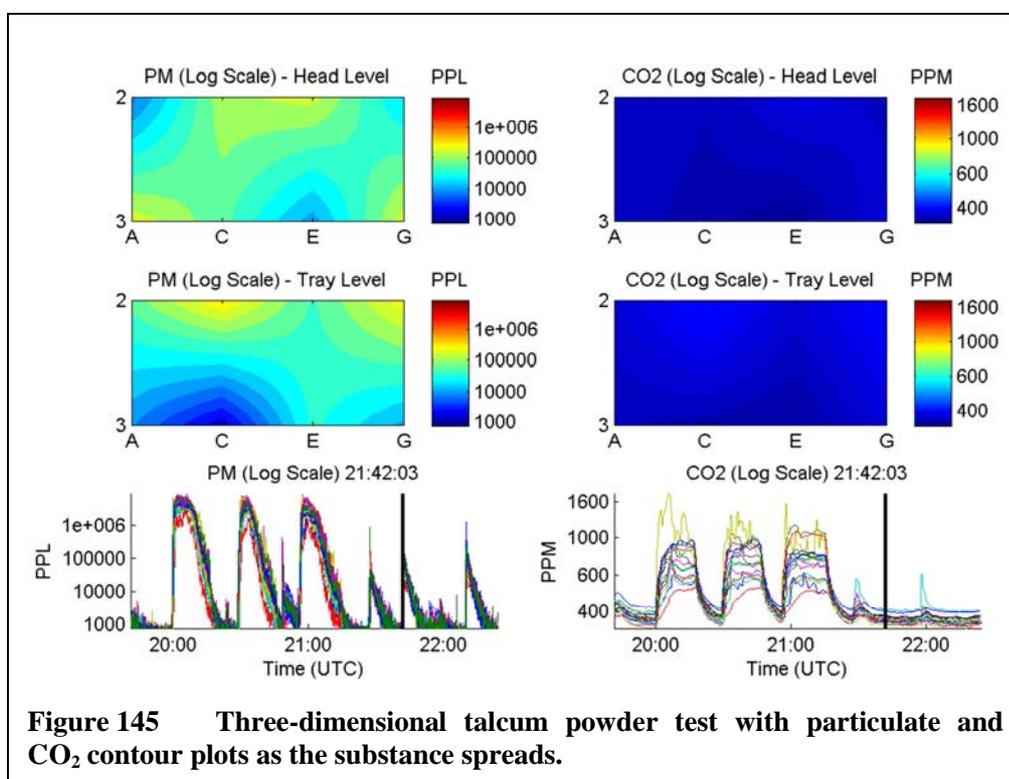
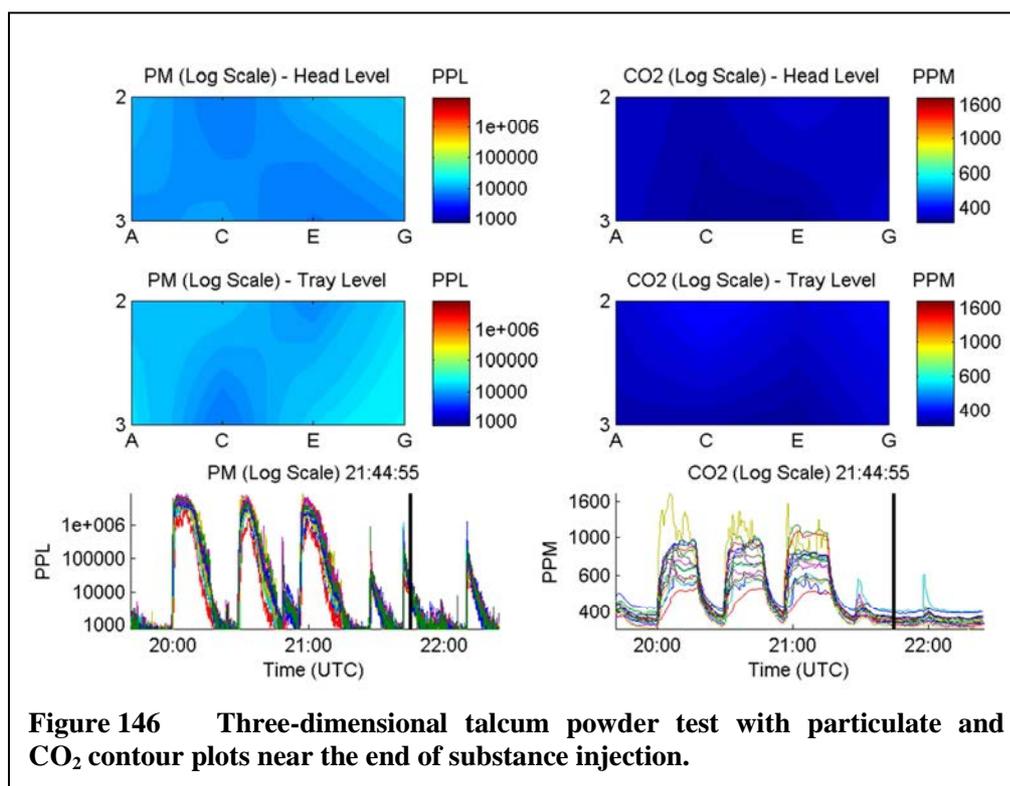


Figure 145 Three-dimensional talcum powder test with particulate and CO₂ contour plots as the substance spreads.



the sides as it is replaced by clean air from the ventilation system, likely being pushed outward from the diffusers along the centerline.

It is difficult to make much of a qualitative comparison of the two different types of particulate matter substance used in the testing. The amounts and distributions of the two substances were quite different, and resulted in concentration peaks that were different by nearly two orders of magnitude. While there may be differences in the behavior of the two materials, further testing with similar concentrations and injection points would be required to make an attempt at any definitive statements along these lines.

The humidity profile for the three-dimensional testing is shown in Figure 147. The deactivation of the dehumidifier during run 4 can be seen as the rise in humidity

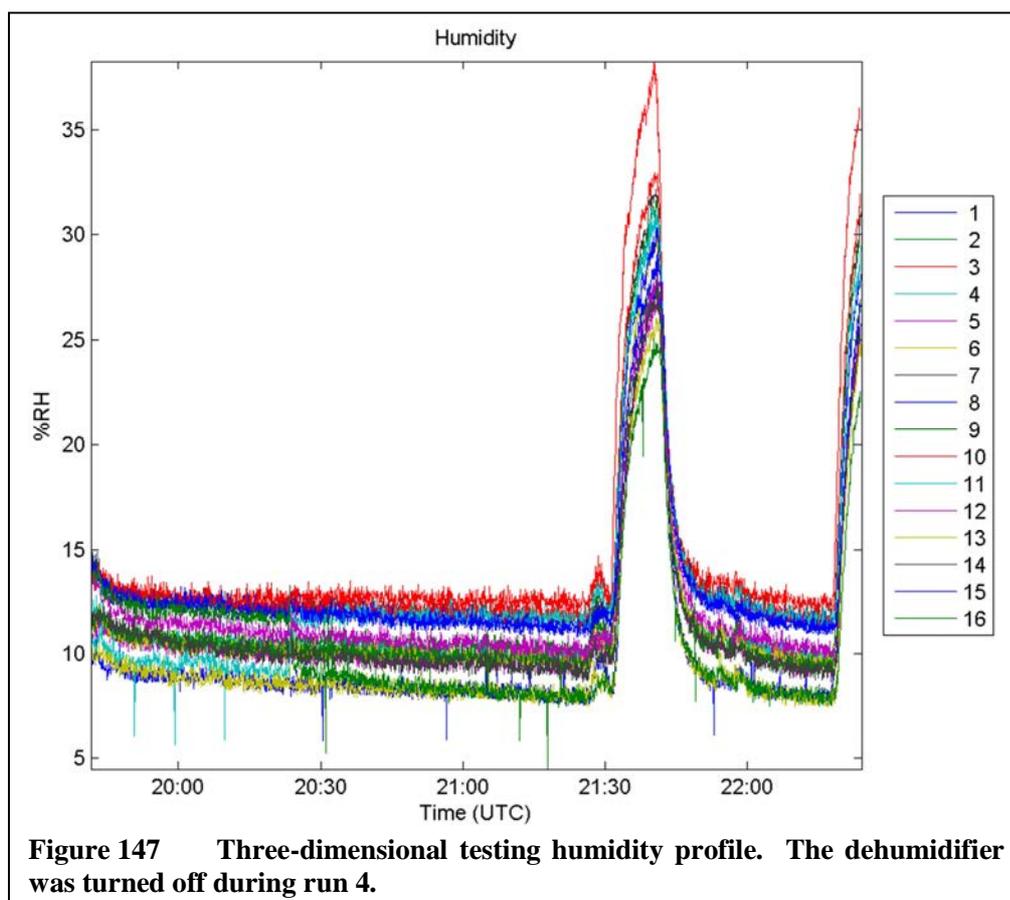


Figure 147 Three-dimensional testing humidity profile. The dehumidifier was turned off during run 4.

throughout the cabin after 21:30. The humidity also rises at the end of the plot, since the dehumidifier was shut down in preparation for closing the testing facilities for the day.

The temperature and pressure profiles measured during the three-dimensional tests are shown in Figure 148. The temperature profile continued the rising trend seen during the two-dimensional testing earlier in the day, though at a slower rate. The rapid changes in readings from the sensor nodes at positions 11 and 16 are due to the swapping of these units to check the CO₂ peaking positional dependence. Since the cabin is not pressurized, the pressure profile reflects the ambient atmospheric pressure.

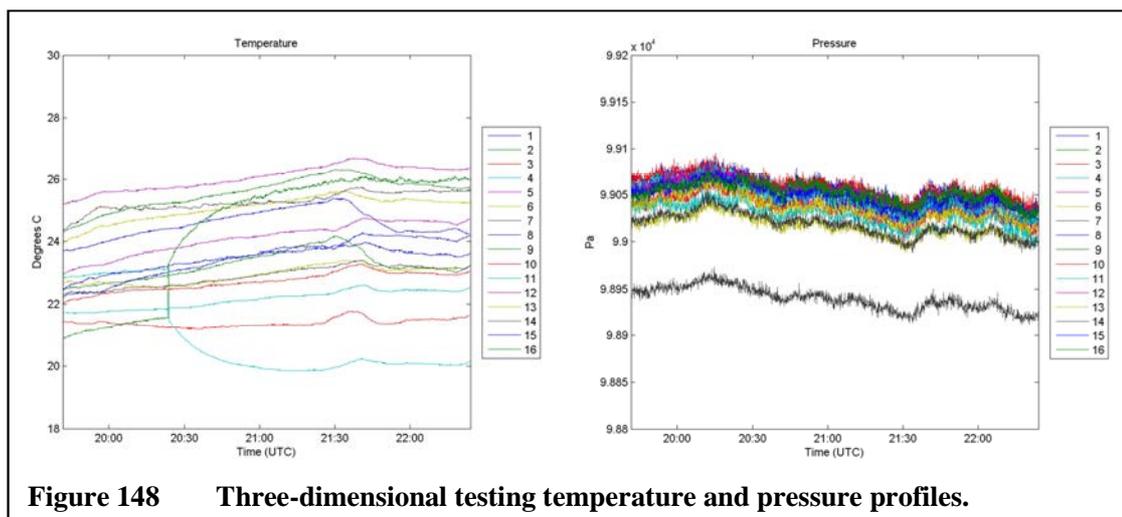


Figure 148 Three-dimensional testing temperature and pressure profiles.

Controlled Environment Testing Conclusions

This testing has shown that a relatively dense set of sensor nodes can be used to track the motion of particulate matter across an enclosed space. This can be done in two and in three dimensions. This can be used to identify source locations and watch the eventual disposition of substances as they move through the space. The coordinated use of sensors for different substance factors can be used to formulate theories regarding airflow and the simultaneous movement of multiple substances.

The coordination of multiple sensor types becomes problematic when dealing with movements that are at the temporal sensing limits of those sensors. This was shown in the data above when measuring the simultaneous injection of particulate matter and CO₂ substances. The particulate matter sensor responds much more quickly to changes in substance, as it is directly sensing particles that are pulled through the system with active airflow and reporting the particle counts on two-second intervals. The CO₂ sensor on the other hand has a listed maximum response time of two minutes. While the typical response time seen in testing is significantly shorter than this, there is still a very

noticeable delay between the appearance of particulate matter and CO₂ in the sensor data output. This discrepancy in data must be recognized and accounted for to meet the requirements of the particular application.

Field Environment Testing Inside an Operating Room

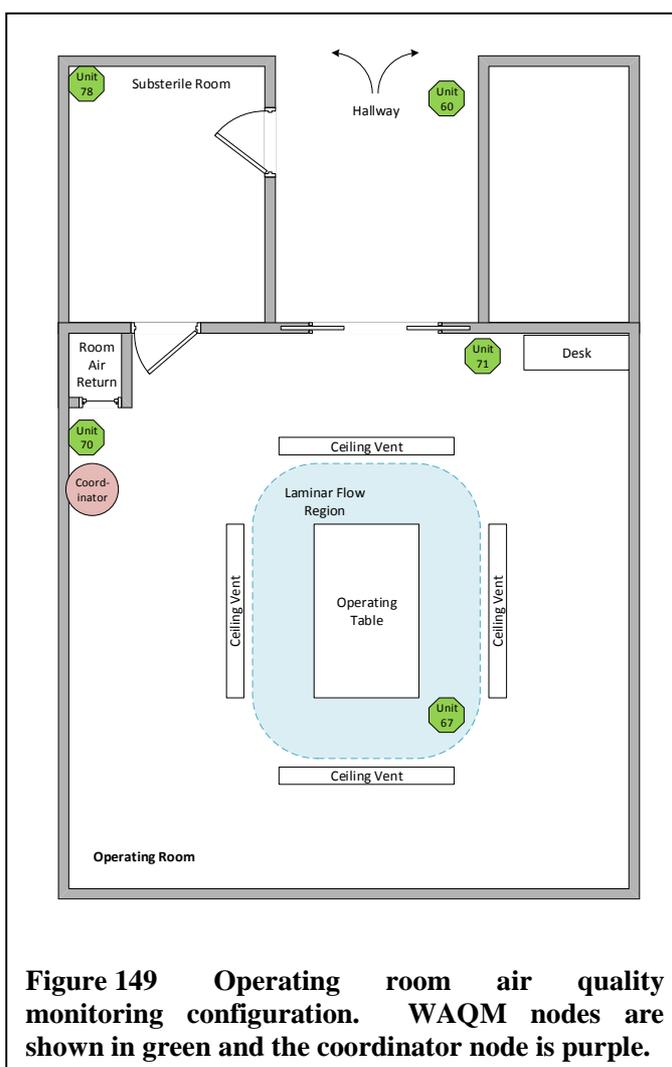
Field environment testing was performed to evaluate the operation of the WAQM system in uncontrolled or partially-controlled environments similar to those that would be seen with an actual system deployment. Over the course of the development of the WAQM system field environment testing was performed in a variety of locations including office spaces, residential housing, a mountain cabin, and a recreational vehicle. One of the more interesting field environments was a hospital orthopedic surgery operating room at a hospital, the name of which is being omitted from this text for reasons of confidentiality.

The testing at this hospital came about after the HSIL team presented a poster covering the WAQM system at a research symposium held at the aforementioned hospital. The poster caught the attention of one of the orthopedic surgeons who happened to be part of a project to reduce the hospital's infection rate during operations. From the ensuing discussion, a collaborative effort was launched between BSU and the hospital to investigate the particulate matter load of the operating rooms at the hospital. Of primary interest was the particulate matter concentration during an operation, and the effects of opening both the main operating room doors and the smaller door through the adjoining substerile room. The hospital's team desired to monitor orthopedic surgeries due to the increased difficulty and cost of managing infections related to surgical implants [86], [87].

Operating Room Layout

A single active operating room (OR) at the hospital was chosen for the experiment. The particular room was chosen due to the frequent orthoscopic surgeries performed there. The room was also very convenient for the test setup, as AC connections were available to power the units in the OR, its adjoining substerile room, and the hallway outside.

The OR field environment test setup is shown in Figure 149. This OR configuration is typical of those at the hospital. The main OR room contains the operating table in the center, with various life support, monitoring, and surgical equipment surrounding it. An attached substerile room contains support equipment. The OR is entered through either the main double doors, or through the adjoining substerile room. According to the surgeons at the hospital, the double doors should only be used to bring the patient into the OR, with all other traffic going through the substerile room to lessen the exposure to air from the hallway.



The operating table in the OR is surrounded by ceiling vents that deliver air into the room. The region in the center of the vents is referred to as the “laminar flow region,” since the air flow is designed keep a sheet of fresh, clean air surrounding the patient with air movement from the center of the room outwards to the periphery. A single, large air return located in the lower-corner of the OR exhausts air from the room. In an attempt to prevent airborne substance from entering through the doors, positive pressure is maintained in the OR relative to the hallway. This forces air out the main double doors and substerile room door when opened, and causes air to flow through any spaces around the edges of the doors when closed.

Operating Room Sensor Configuration

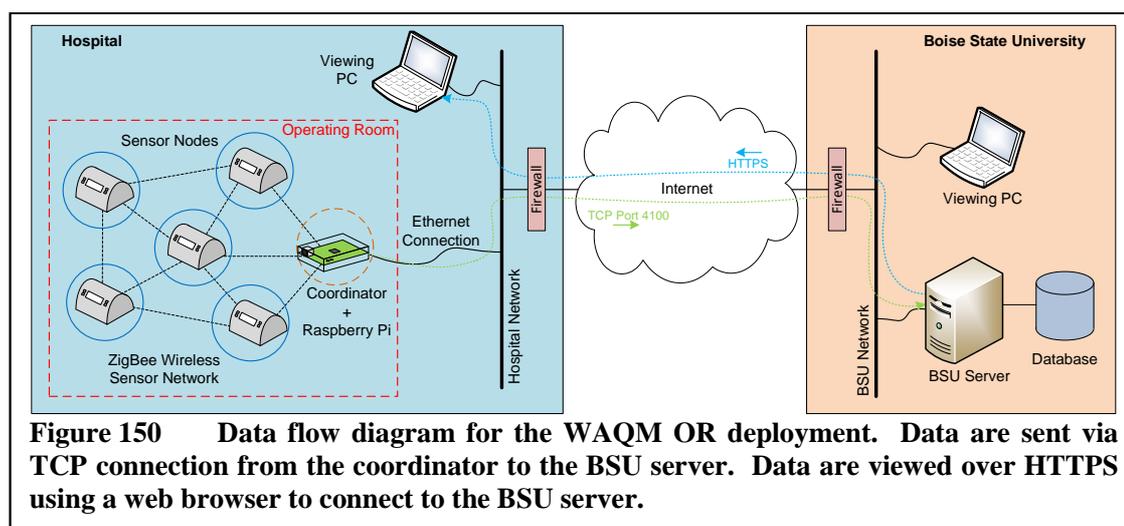
To monitor the OR, five WAQM sensor nodes were employed along with a dedicated coordinator unit. The arrangement of the sensor nodes and the coordinator can be seen in Figure 149. Three of the sensor nodes were strategically positioned inside the OR to cover specific areas of interest. One unit (Unit 67) was located on a power and gas delivery boom inside the laminar flow area, above the level of the patient. The desire was to have this unit as close as possible to the patient without interfering with operating procedures. The second unit (Unit 71) was located near the main OR doors, with the intent of capturing any increase in substance entering when the doors were opened. The third unit in the OR (Unit 70) was located low on the wall near the air return for the room. This unit was meant to sample the OR air substance levels as it left the room through the intended exhaust outlet.

Outside of the main OR room, one unit (Unit 78) was positioned in the corner of the substerile room near the ceiling to sample substance moving through this space

between the main OR and the hallway. The location of the unit within the substerile room was driven mostly by access to power, and to keep the unit out of the way of the OR staff. The final unit (Unit 60) was suspended from the ceiling outside the main OR doors, at the entrance to the short stub connecting the OR to the main hallway. This unit was positioned to provide correlation between the air in the hallway and the air in the OR.

Operating Room Data Access

When setting up the WAQM sensor network in the OR, it was desired to enable real-time viewing of the data and direct streaming of data into the BSU database. To support this, the sensor network would need to be connected to the hospital local area network, which could then provide a connection to the BSU network across the Internet. The data flow that was created is shown in Figure 150. At the time, the WAQM framework only supported a Wi-Fi Ethernet connection using Wi-Fi Protected Access II (WPA2), and the hospital network required a version of the Challenge-Handshake Authentication Protocol (CHAP). A guest login to the Wi-Fi was also available, but needed to accept a user a usage agreement and would time out after approximately one

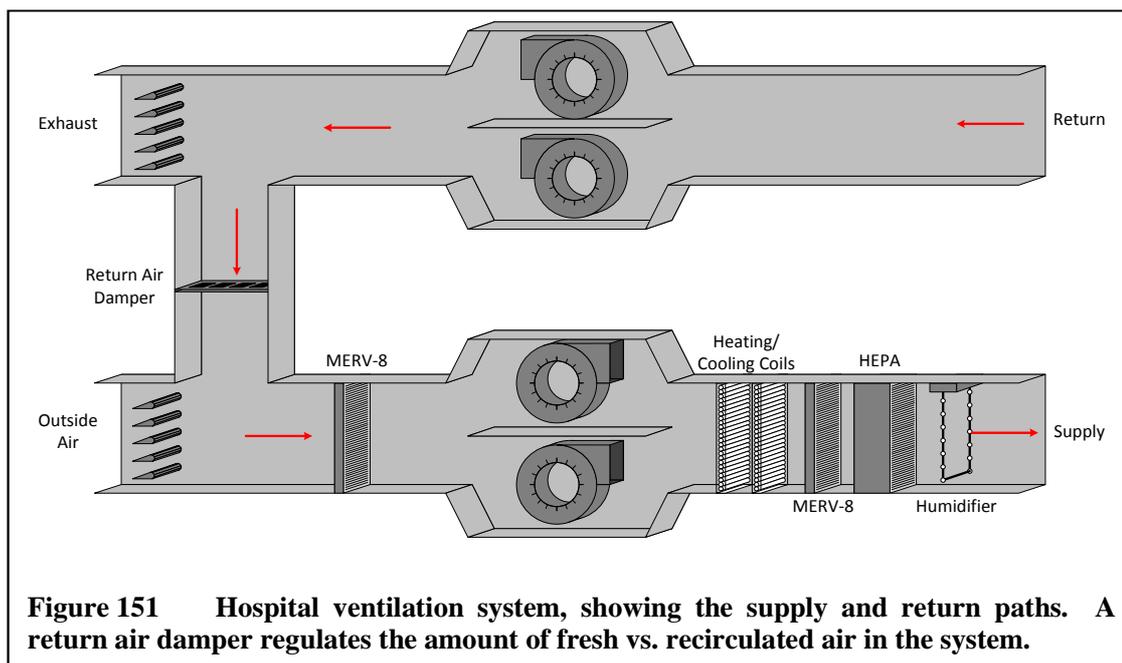


hour. Wired 10/100 Ethernet ports were available in the OR, but unfortunately none of the HSIL hardware supported this connection type.

Since it was infeasible to support one of the Wi-Fi connection methods prior to deployment, it was decided to bridge the WAQM network to the hospital network using wired 10/100 Ethernet. To accomplish this, a Raspberry Pi [88] was used to connect the UART of the coordinator to Ethernet. A Perl script was created that routed UART measurement and command messages to a TCP/IP stream socket on the Ethernet port that was connected to the Connector Server at BSU over the Internet. The hospital IT staff opened a hole in the hospital firewall specifically for this connection, and for Network Time Protocol queries coming from the Raspberry Pi. Using the Raspberry Pi allowed a solution to this connection problem to be created and deployed in less than a day.

Operating Room Air Supply

Air coming into the operating room is supplied by a forced air ventilation system that can mix fresh, outside air with recirculated air. A diagram of the system is shown in Figure 151. Air is forced through the system with blowers directing air through both the supply and exhaust sides of the flow. Air from outside the building enters the system at bottom left, and is passed through a MERV-8 class filter before reaching the supply-side blowers. After the supply-side blowers, the air passes through heating and cooling coils, another MERV-8 filter, a 12-inch thick box HEPA filter, and a humidifier before entering the operating room. The mix of recirculated and fresh air is controlled with dampers in the supply inlet, exhaust outlet, and a duct connecting the supply and exhaust paths. The entire system feeds multiple operating rooms in the hospital simultaneously.



Of particular interest in the ventilation system are the mechanisms intended to remove particulate matter from the air. This is directly accomplished by the three filter stages used in the supply-side of the ventilation path. The two MERV-8 filters surrounding the blowers and heating/cooling coils are only effective for relatively large particles, being rated at greater than 70% efficient for particles in the 3-10 μm range. The MERV-8 filters are not rated for particles smaller than 3 μm . The box HEPA filter further downstream has an efficiency rating of 99.97% for particle sizes of 0.3 μm .

Of potential importance to the sensor system testing is the humidifier in the supply-side of the ventilation path. Humidity control is important in hospital environments to reduce the occurrence of electrostatic discharge (ESD). Air mixtures with high oxygen concentrations are often used in hospitals, especially in operating rooms, and can pose a significant fire hazard in the presence of ignition sources such as ESD. The humidifier in this particular hospital ventilation system is a feedback-controlled device that injects steam into the supply-side air. The return-air relative

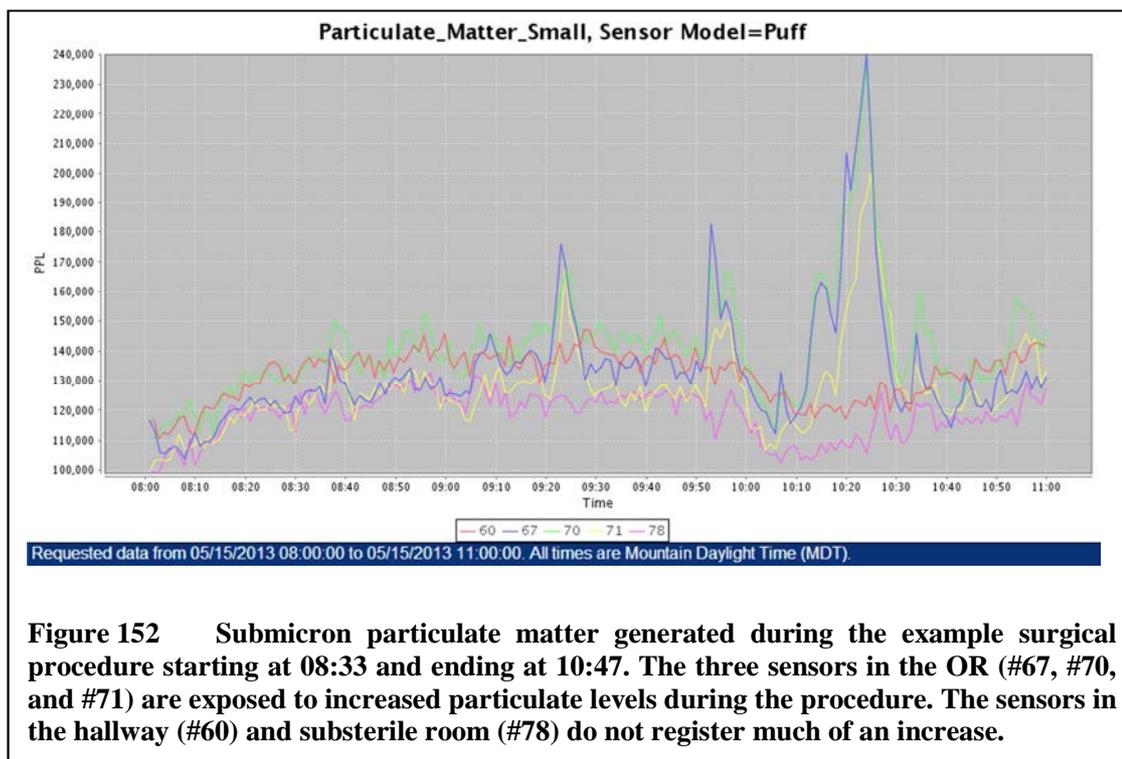
humidity is sampled and used to control the flow of steam into the humidifier. The steam is produced from water source fed through a reverse-osmosis filter that is rated to remove a minimum of 98% of all sodium chloride from the input water.

Operation Room Data

The WAQM sensor network was installed and running in the selected operating room for approximately 12 weeks. During the course of the testing, there was no physical access to the OR itself, making correlation with actual activity a significant challenge. Daily reports on OR activity were provided, but were limited to the titles of the performed procedures along with their start and stop times.

While 190 recorded procedures were covered during the course of the test period, a single procedure is presented as a representative example of the data collected. This particular procedure was a robotic assisted hysterectomy, which started at 08:33 and ended at 10:47 on the morning of May 15th, 2013. The example surgical procedure was chosen for its relative clarity in representing some of the typical characteristics of many procedures seen during the test period. In the next several figures, data from each of the WAQM node environmental sensors are shown for the same surgical procedure.

The submicron particulate matter concentration during the course of the procedure is shown in Figure 152. Three distinct peaks in particulate matter can be seen in the plot, with all sensors inside the OR registering some coincident increase. The units in the hallway and substerile room do not show the same peaking, indicating some isolation from the events. Since the events occurred during the course of the procedure with the first occurring nearly 50 minutes after the start, it is very likely that the doors were closed when the particulate matter was introduced into the OR.



While there is no conclusive proof as to the origin of the particulate matter peaks, it does appear likely that it is generated inside the OR itself. Some portions of the procedure, such as the use of electrosurgical apparatus or bone saws are known to generate significant amounts of particulate matter [89]. Unfortunately due to access issues the study was not able to precisely correlate the particulate events with any set of activities in the OR.

Measurements of large particulate matter over 1 μm in size appear to show profiles that are mostly independent of the submicron particulate variations. Figure 153 shows the large particulate matter concentration for the example surgical procedure. The highest amplitude peaks do not correspond with those in Figure 152 for submicron particulate. Upon closer examination, it is clear that the peaks are occurring very near to

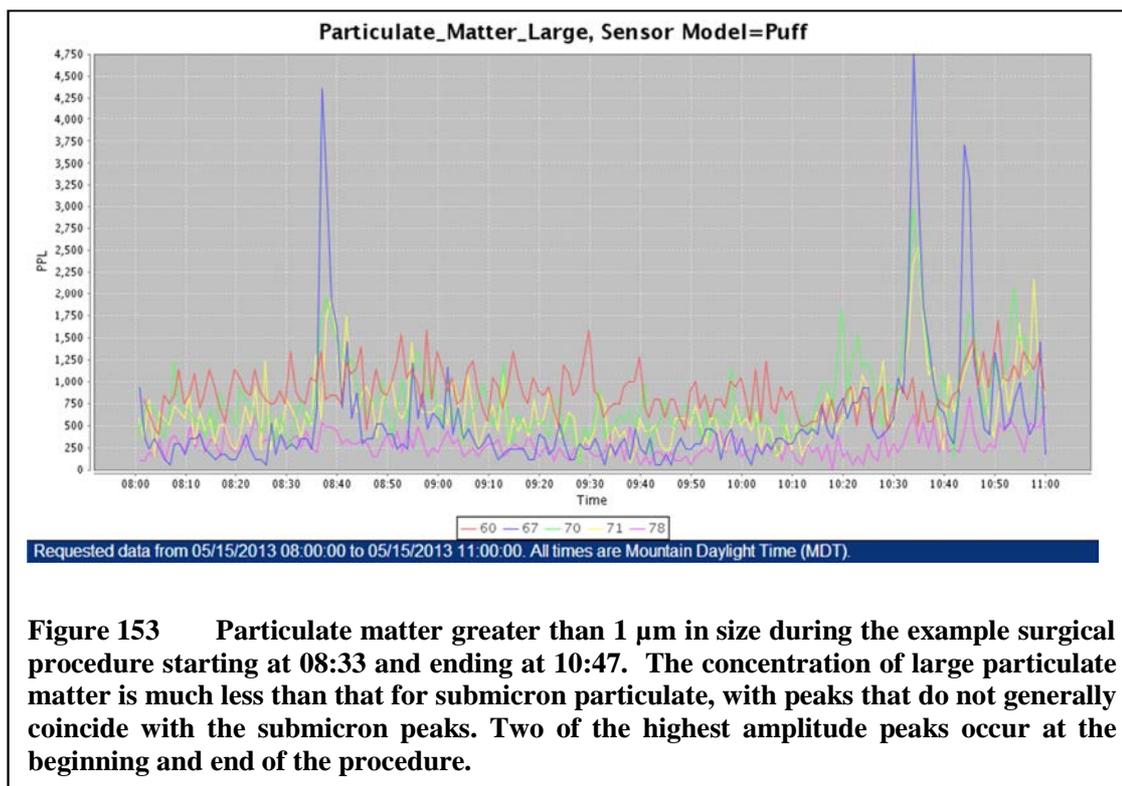


Figure 153 Particulate matter greater than 1 μm in size during the example surgical procedure starting at 08:33 and ending at 10:47. The concentration of large particulate matter is much less than that for submicron particulate, with peaks that do not generally coincide with the submicron peaks. Two of the highest amplitude peaks occur at the beginning and end of the procedure.

the start and end of the procedure, possibly coinciding with large movements of equipment and personnel in the OR.

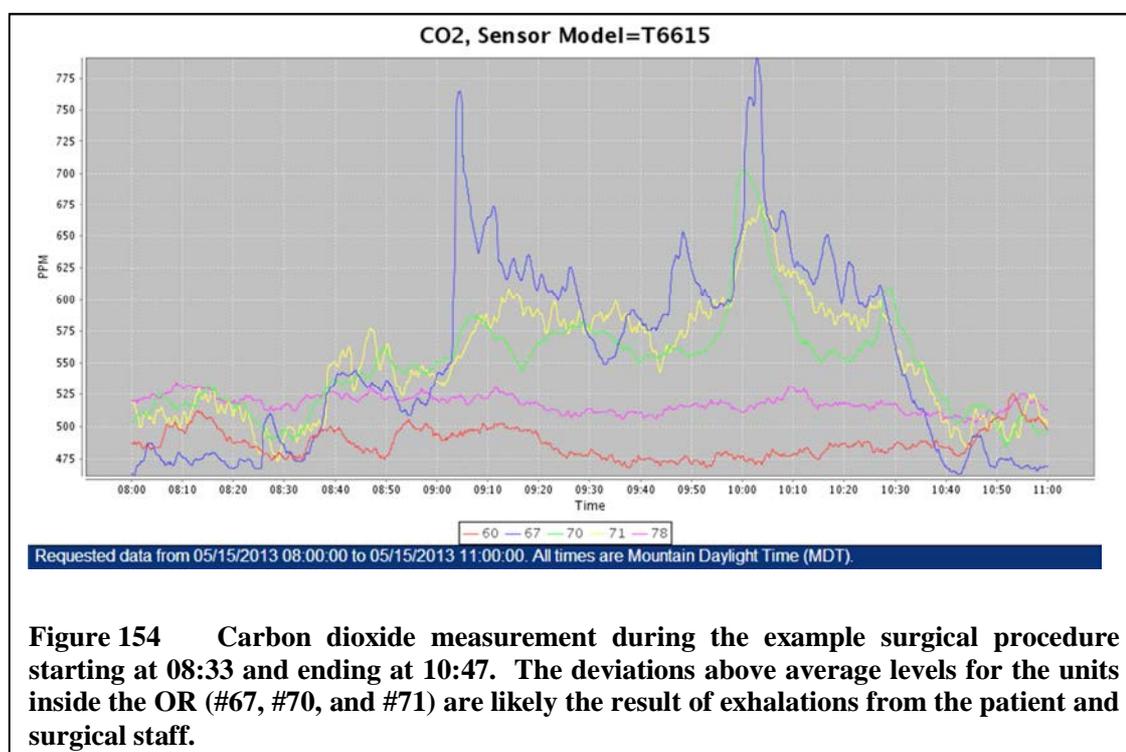
The large particulate matter peaks may also coincide with increases in movement through the door into the hallway and substerile room. However, it does not appear that the particulate is airborne when it enters the OR, since the sensors in the hallway and substerile room do not show increases in concentration. It appears most likely that the particulate is being generated in the room, possibly by resuspension of particles that have settled on equipment or personnel already in the room.

Many of the surgical procedures tracked during the study showed similar large particulate matter profiles, with peaks near the start and end of the logged time. This large particulate movement may be more significant to infection risk in the OR

environment than particulate of submicron sizes, since many bacteria are larger than 1 μm in size [1].

Carbon Dioxide levels in the OR in general appear to provide a good indicator of room usage. Levels tend to stay near to the 400-500 PPM range when the OR is not in use, and then spike to higher levels during a procedure. Figure 154 shows the CO_2 levels during the example surgical procedure. Note that there is clear spiking in the OR, but not in the hallway or substerile room. It is also of interest that the sensor node inside the laminar flow area sees the largest magnitude spikes. This may be due to the concentration of surgical staff near the patient during the procedure.

It is possible that with calibration against actual usage data, the CO_2 measurements could be used to track the number of staff and activity in the OR. As mentioned previously, this particular deployment of sensor nodes did not allow the direct

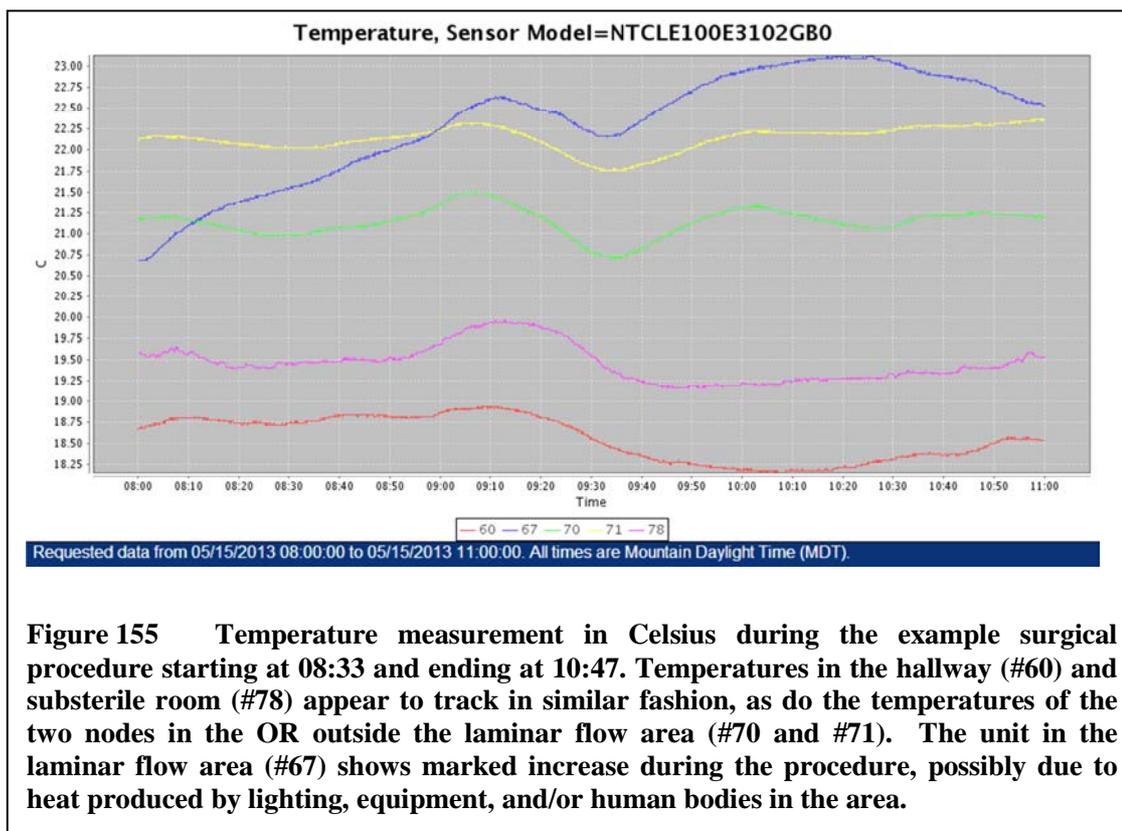


observation of the OR while in use. It is also unclear at this time whether or not human respiratory expiration is the only source of increased CO₂ concentration in the OR, or if there is another chemical or process source involved. Further investigation and correlated testing would be required before these data could be used to recognize staffing levels or activities.

Temperature levels measured by the sensor nodes tended to track in groups based on location. The three units in the OR tended to track each other much more closely than the units in the hallway and substerile room. This might be expected given the common environment of the three OR units. The unit inside the laminar flow area also tended to show significant positive excursions during a procedure. This might be due to an increase in heating due to OR lighting above the operating table, additional equipment active during the procedure, or contribution of body heat by the patient and surgical staff. It is also possible that the power and gas delivery arm to which the laminar flow area unit was attached was warming during the procedure, directly conducting heat to the sensor node.

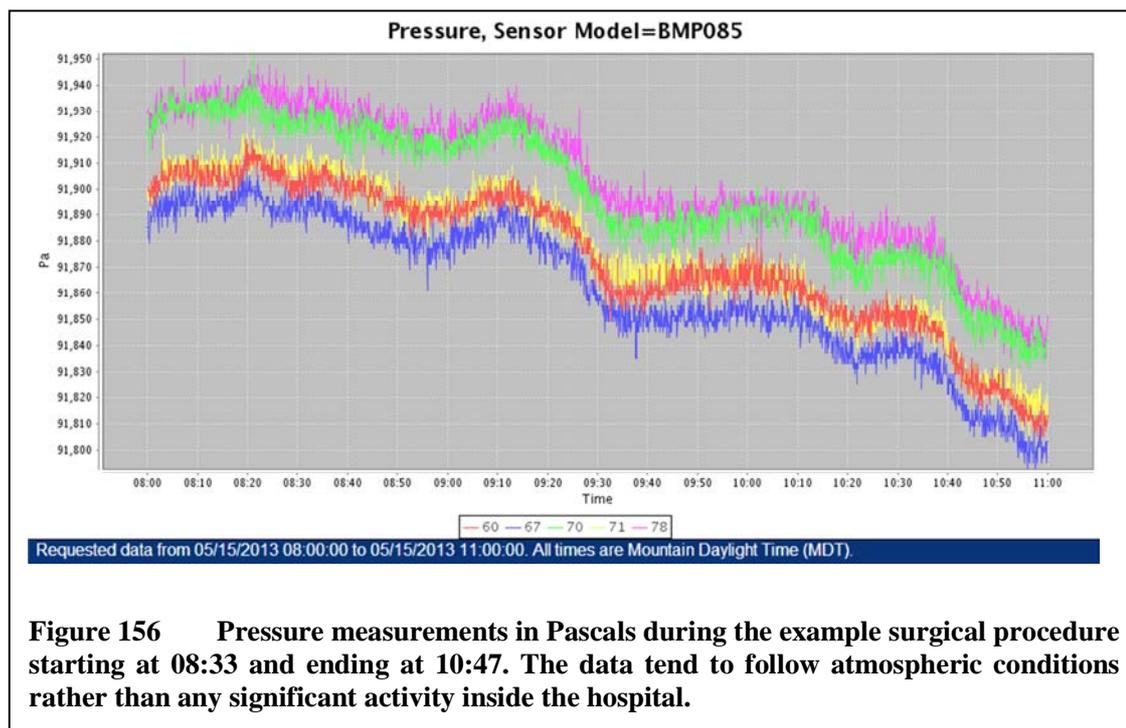
Figure 155 shows the temperature measurements from the OR during the example procedure. Note that the units in the hallway and substerile room tend to follow the same general trends. The units inside the OR also follow a common trend, but the unit inside the laminar flow area shows significant increase in temperature during the procedure that is not exactly reflected in the temperature data collected by the other two units.

The positive temperature deviation by the laminar flow area sensor may also be an indication of effective airflow separation between the laminar flow area and the outer perimeter of the OR. If the air temperature remains elevated inside the laminar flow area without greatly impacting the other sensors in the OR, it could mean that much of the



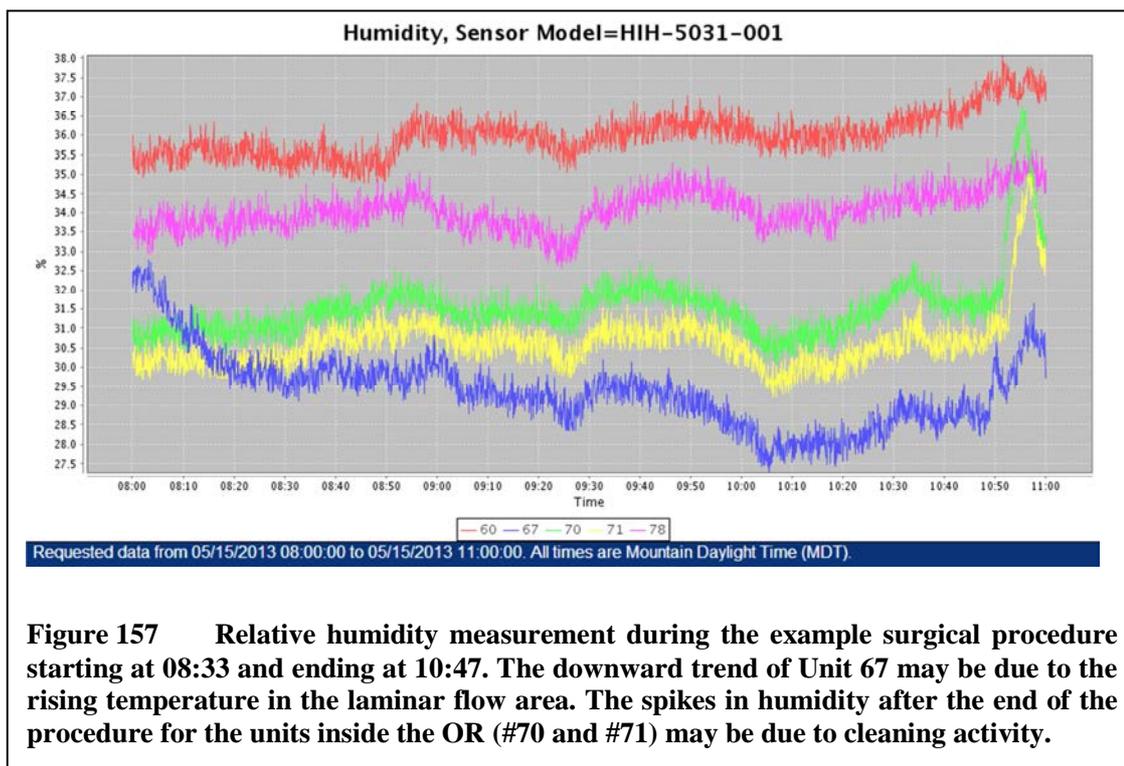
heated air is remaining inside the laminar flow area. Of course this type of restricted flow was not greatly observed in the data from the particulate matter sensors, which saw spiking on all sensors inside the OR when particulate matter was generated. Given this disparity between the two types of measurement, one might also conclude that the sensor unit itself was heated by an external source, rather than seeing a direct increase in air temperature. One might speculate that the black paint on the sensor nodes was effective at absorbing radiant energy from the OR lights, causing a temperature increase inside the sensor node.

The air pressure in the OR tended to follow atmospheric conditions when compared against units operating at other locations. Figure 156 shows the pressure measurements from the example procedure. Differences seen between the five different



sensor units were well within the absolute accuracy specifications of ± 250 Pa for the Bosch BMP085 sensor used. There was some hope at the beginning of the experiment that it would be possible to detect the opening and closing of the main OR doors with the pressure sensors due to the positive pressure bias inside the OR, but this small pressure change was not detectable.

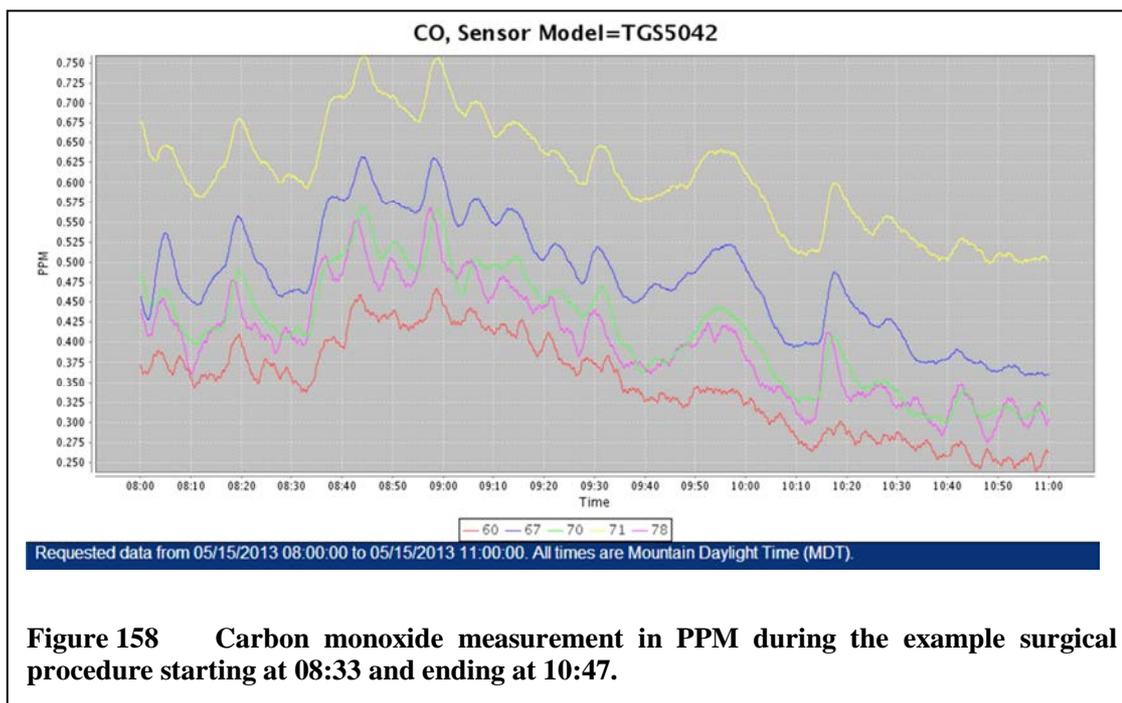
Relative humidity inside the OR can display a great amount of variation due to its dependence on the air's absolute water content as well as the air temperature and pressure. Comparing the three measurements can lead to some understanding of actual moisture content in the OR. For example, Figure 157 shows the relative humidity for the example surgical procedure. As with temperature, the units in the hallway and substerile room tend to track each other. This is also true for two of the units inside the OR, but not so for the unit in the laminar flow area. This unit shows a steady decline in relative humidity during the course of the procedure. This is likely due to the rising temperature



in the laminar flow area rather than a decrease in the water content of the air, since the relative humidity reading is directly impacted by the change in temperature.

The spikes in relative humidity detected by the units in the OR after the end of the procedure are not reflected in the temperature data, which likely indicates that they are actual increases in air moisture content. The hospital procedure of cleaning each OR after it is used could be a possible cause of the spikes, if the cleaning process is briefly increasing the air's water content.

The carbon monoxide readings in the OR were particularly interesting, since it was expected that none would be observed. In general, the CO readings appeared to peak during active procedures, and flatten out during periods of disuse. Figure 158 shows the CO measurements during the example surgical procedure. Note that the levels peak upwards during the procedure, then fall off at the end. The levels seen are very low

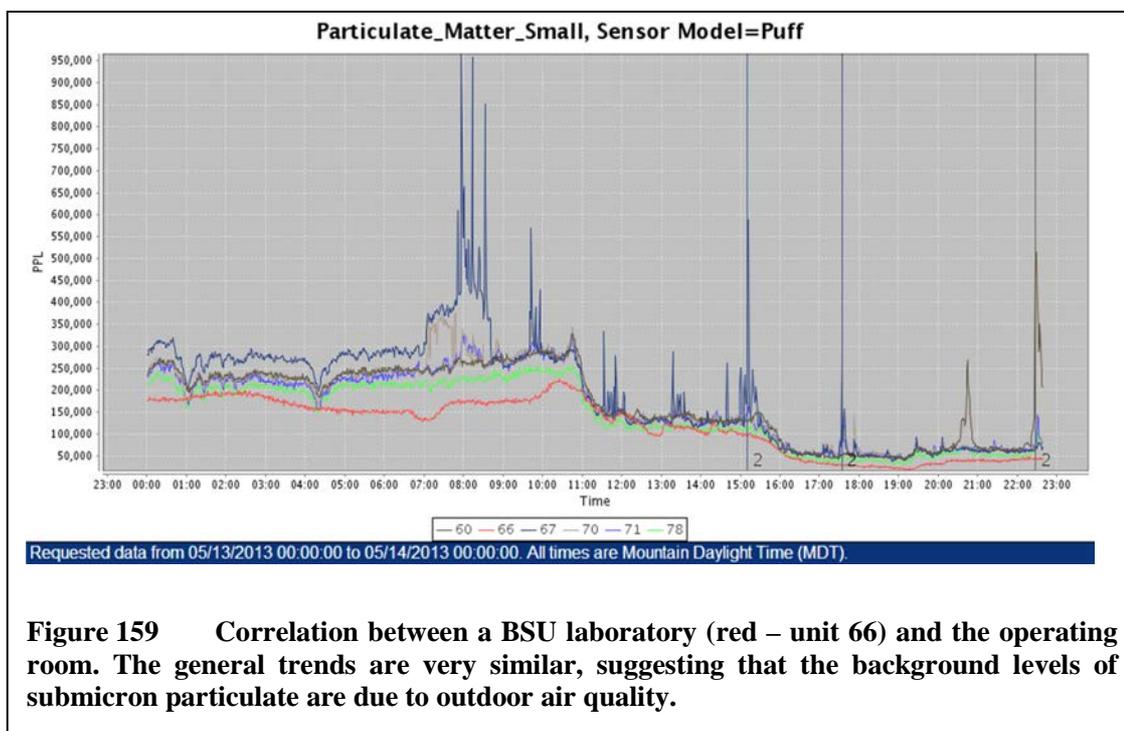


relative to the sensor range of 0-10,000 PPM, but the correlation in readings between the units appear to indicate that the actual levels are changing during the procedure.

It is very possible that the readings are not due to carbon monoxide at all, and are instead due to some other chemical in the OR that the Figaro TGS5042 carbon monoxide sensor is sensitive to. The cross-sensitivity list for this particular sensor includes a long list of chemicals including hydrogen, methane, isopropyl alcohol, toluene, and ethylene. Use of isopropyl alcohol in the OR could cause the carbon monoxide sensor to react. While the common OR anesthetics such as isoflurane, enflurane, and halothane are not directly listed on the cross-sensitivity list, it is also possible that these gases are causing the carbon monoxide sensor to react.

Background Levels of Particulate Matter

The OR levels of submicron particulate were relatively high, even when no surgical procedure was taking place. There could be many sources for this particulate, including fresh air from the outdoors, recirculated air, or some other source in the air flow such as the humidifier. The experiment did not include a WAQM unit outdoors near the fresh air intake, so it was difficult to directly compare the indoor and outdoor levels. It was, however, possible to compare background levels with units located in the Boise State University laboratory. Figure 159 shows one such unit in comparison with the units in the OR. The red line is the unit at BSU. While there is not a perfect match, the background trends tend to match on a macro-level, even though the two buildings are more than 1.5 km apart. This matching of background levels was seen consistently during the course of the 12 week experiment, strongly indicating that the levels are mostly due to ambient outdoor particulate concentrations.



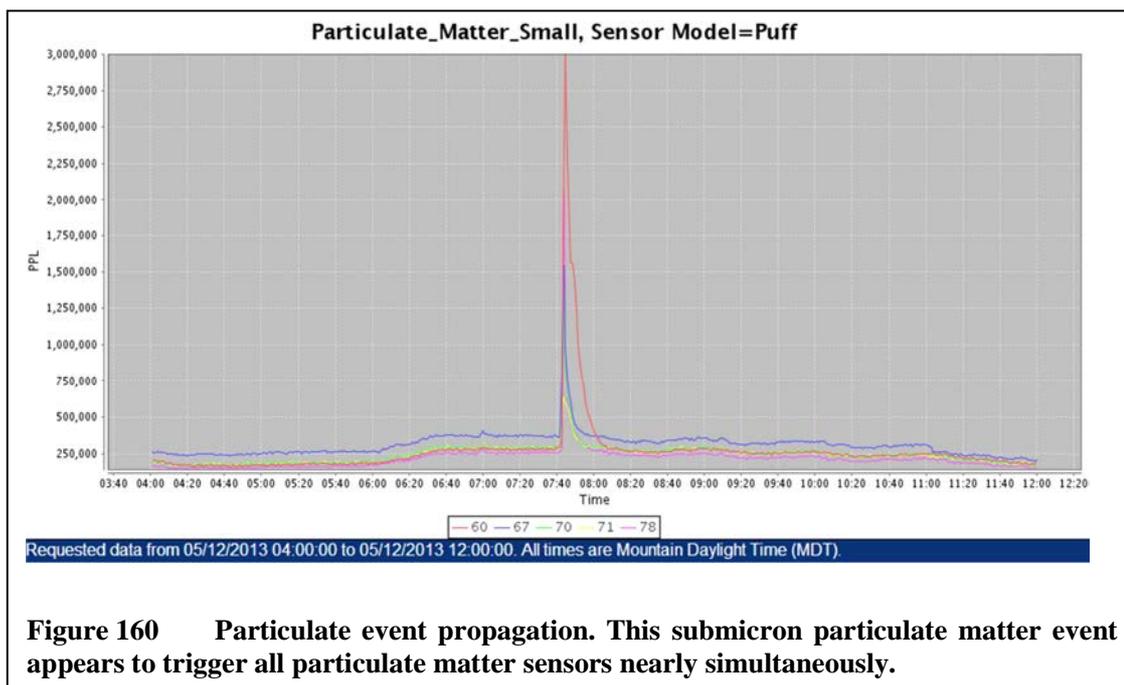
These data are curious given the filtering in the OR supply-side airflow. While the MERV-8 filters would not be expected to impact the submicron levels, the HEPA filter certainly should. The filter's rating of >99.97% efficiency for removal of 0.3 μm particles would lead one to expect a much cleaner air supply. On querying the hospital staff about the filter schedule, this HEPA is apparently only switched out once per year. It is possible that the filter had lost effectiveness when the experiment was run.

Since the WAQM nodes don't detect particles much smaller than 0.3 μm , it does not appear likely that the particulate is smaller than the filter's rated size. However, if the particle sizes are increasing at some point past the filter stage, the WAQM might detect particles that are not filtered. Particle growth might occur through heterogeneous nucleation if the humidity and temperature of the air reaches a supersaturated state [72]. This may occur if warm, humid air containing particulate matter cools, resulting in the adsorption of water molecules onto the particles.

Operating Room Air Movement

One of the key desires of setting up the particulate monitoring system was to determine whether or not particulate was entering the OR from the hallway. The data from the monitoring period was analyzed for such events, but none were found that clearly indicated migration in this direction. In fact, it consistently appears that particulate events are sourced in the OR itself, and then propagate out into the hallway. Some events appear to trigger all WAQM nodes simultaneously, but on closer examination of the data reveal a source within the OR.

Figure 160 shows one such event detected by the submicron particulate matter sensor. This event at first appears to have a source in the hallway simply due to the



magnitude of the particulate matter peak at the sensor there relative to those inside the OR. Note the red line indicating the concentration at Unit 60 in the hallway reaches nearly 3 million particles per liter, while the next largest magnitude is seen at Unit 78 inside the substerile room.

Figure 161 shows a zoomed-in view of the submicron particulate matter concentration for the event. In this view, it is clear that it is difficult to determine the source due to the low resolution reporting of the particulate matter concentration. The units outside the OR have the largest magnitude readings, yet appear to rise just after two of the sensors in the OR. Fortunately the sensor network is also set up to output raw particle counts from each node at a higher frequency than the particulate matter concentration. The absolute magnitudes of the counts are not calibrated, but the readings do provide a more accurate indication of timing with a 5-second resolution. This is much improved over the 60-second resolution of the particulate matter concentration readings.

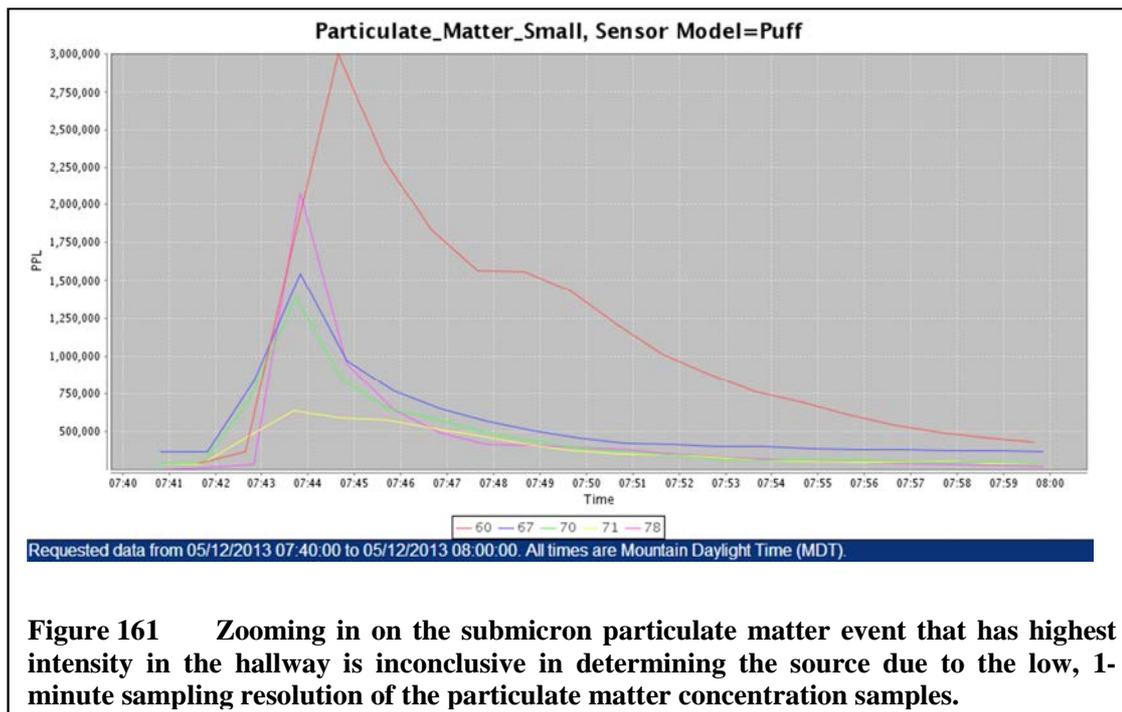


Figure 162 shows a view of the same event looking at the raw submicron particle count data. From this plot one can see that the particle counts rise inside the OR more than 30 seconds before they start to rise in the hallway and substerile room. While it is still difficult to surmise the exact location of the particulate generator causing the event, these earliest peaks would seem to indicate that the event did not originate in the hallway or substerile room. The rapid movement of particulate into the hallway and substerile rooms may indicate that the event was near to the two doors, or possibly that one or both doors were open at the time.

The lack of evidence for movement from the hallway into the OR does not conclusively rule out any such motion. While the testing appears to indicate that particulate moves in the opposite direction, it also appears to show that events generating significant amounts of particulate matter are occurring within the OR and not in the hallway. If possible, a future test of the environment might include intentional generation

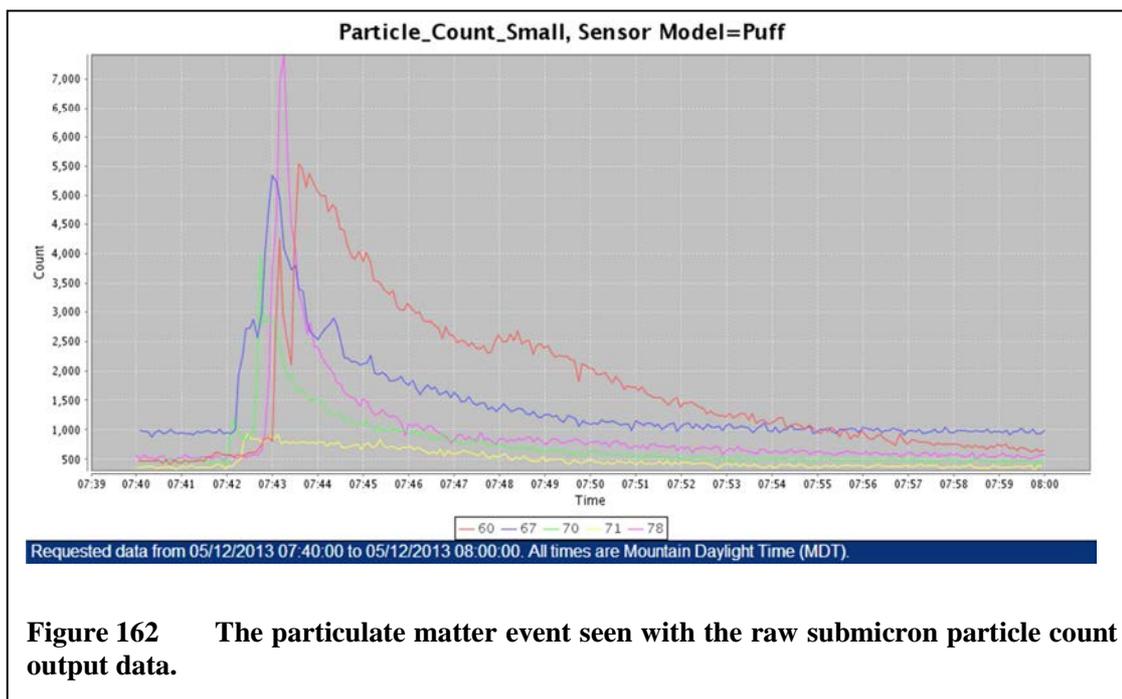


Figure 162 The particulate matter event seen with the raw submicron particle count output data.

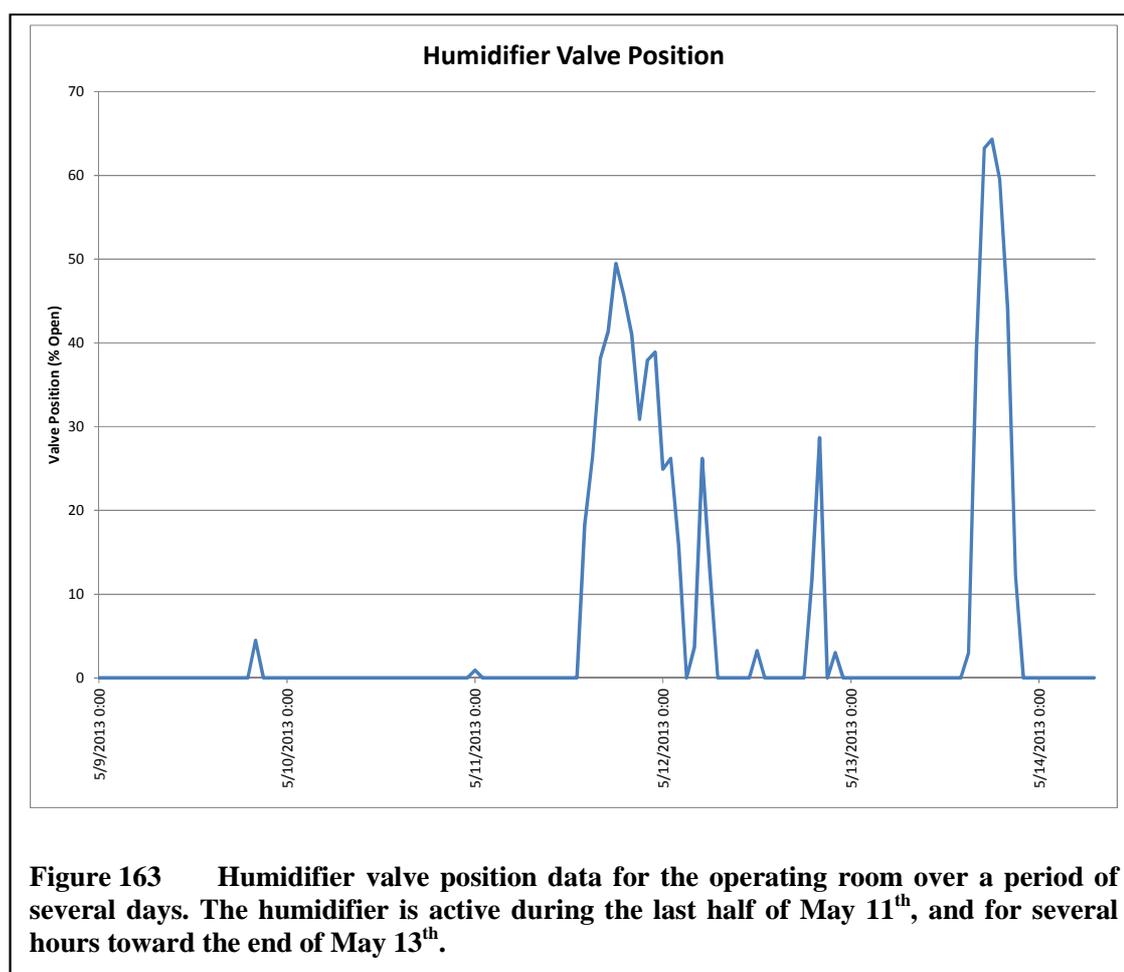
of particulate matter in the hallway to test for flow into the Operating room. Unfortunately, due to the sensitive nature of the environment, this was not possible during the test period.

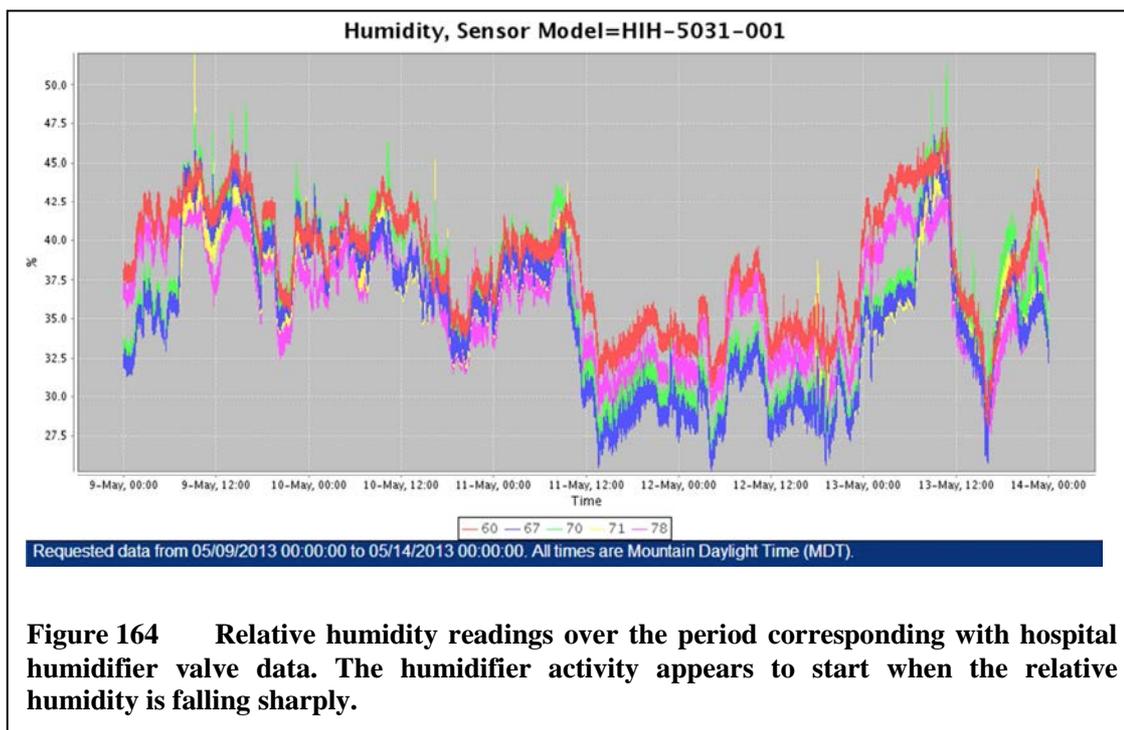
Humidifier Impact on Particulate Matter Levels

The air supply into the OR and hospital in general is controlled for relative humidity levels with a set point at 40% RH based on a feedback sensor in the air return. As described above, steam is injected into the air supply by a humidifier to maintain the humidity level. The humidifier might conceivably impact humidity levels in two ways. The first method is the direct creation of particulate matter by injecting steam into the air that contains some substance or ionic content. This substance may act as seed particulate that could nucleate to create measurable particulate matter. The second method involves the nucleation of the injected steam onto existing very small particles, growing them to a measurable size.

To test this theory, the humidifier valve position data were supplied by the hospital staff. These data are very coarse, and only provides a reading of the humidifier valve position once per hour. This does give some indication of the time periods during which the humidifier was active, but is not very useful for short-term correlation. Figure 163 shows the humidifier valve position data for five consecutive days. The humidifier appears to be normally inactive, but does have two periods of significant use during the afternoons and evenings of May 11th and May 13th.

The humidity data for the OR can be seen in Figure 164, which shows that the relative humidity levels are falling sharply to low levels when the humidifier first turns



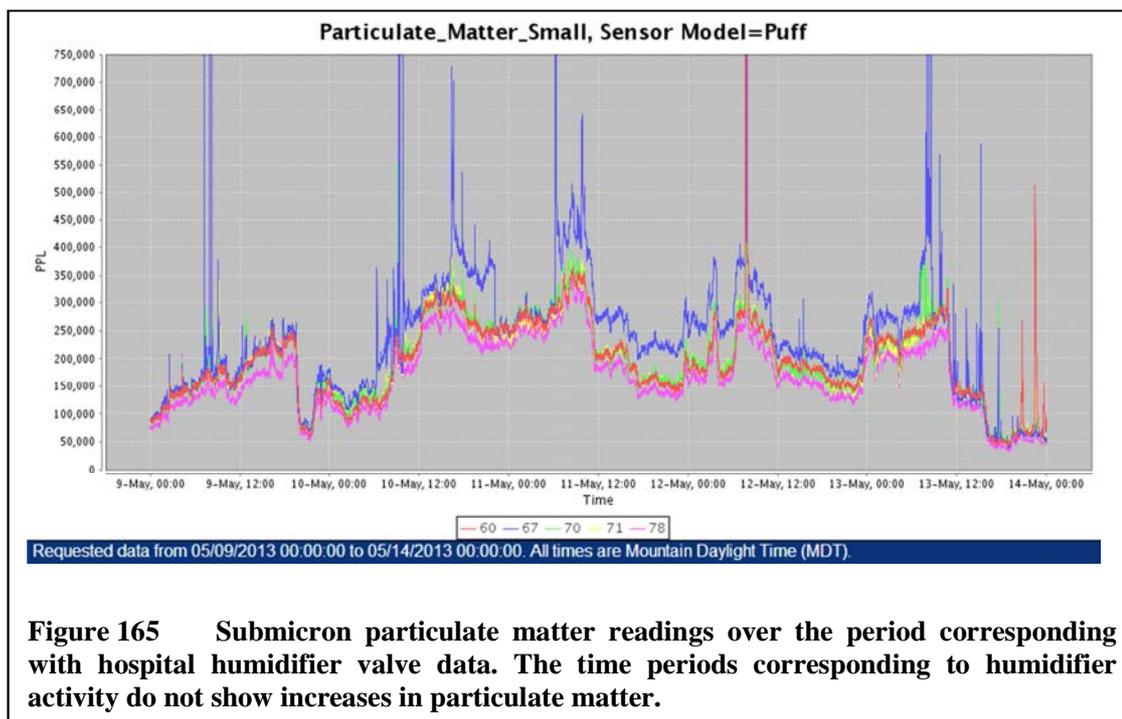


on. This is true for both the May 11th and May 13th humidifier events. As would be expected, the humidifier is turning on when the relative humidity drops significantly, triggering the feedback sensor in the return air path.

Figure 165 shows the submicron particulate matter concentrations in the OR during the five day period corresponding to the humidifier valve position data. The periods when the humidifier is active do not show any marked increase in particulate matter. If any particulate matter is being generated either directly or indirectly by the humidifier, it is not at a level that significantly impacts the particulate readings in the OR.

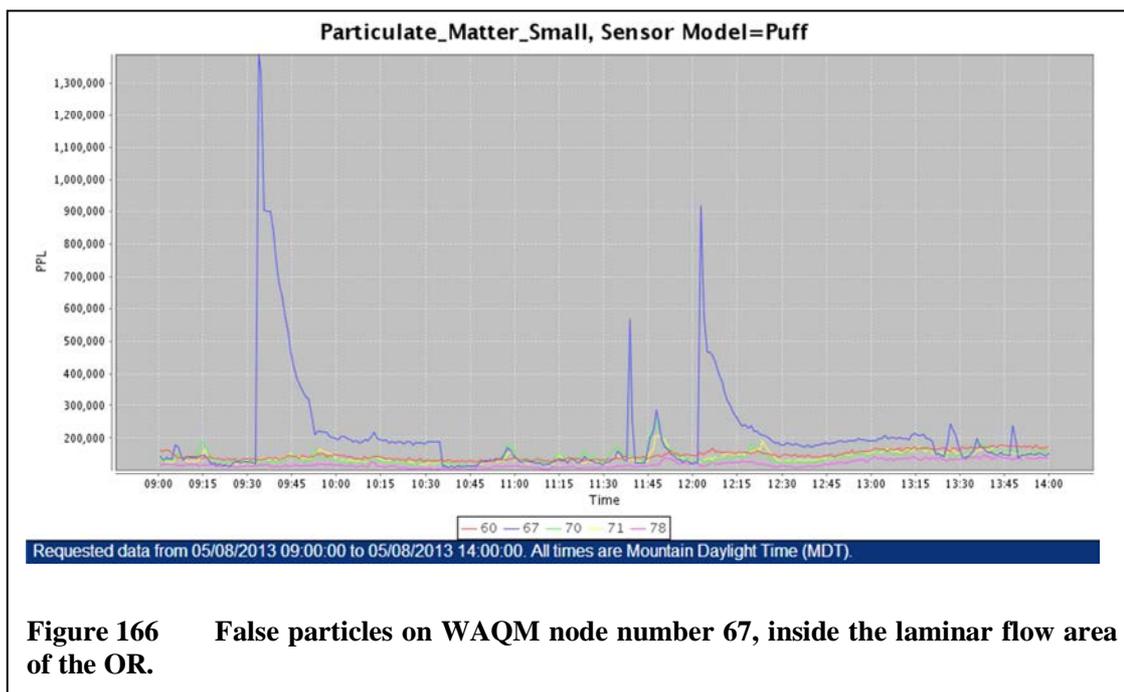
False Particle Peaks

During the course of the time in the OR, it was noticed that one unit in particular would peak to a very high concentration and then fall to an elevated level from time to time. The unit would then remain at the elevated level for a significant duration of time



before falling back to a level that matched the other units in the room, often staying high for multiple hours at a time. The unit in question was WAQM node 67, which was mounted to the gas and power delivery arm inside the laminar flow area. Figure 166 shows an example of this behavior, with Unit 67 peaking twice over the course of five hours. Note that in each case the other sensors in the room do not register a significant shift in concentration.

The problem was very difficult to diagnose in real time, since the units were not accessible inside the OR once they were installed. At first it was thought that there was either an electrical interference issue, or that possibly the OR lighting contained an oscillating component that caused the particulate levels to rise. The issue with both of these possible sources was the initial peaking followed by elevated levels did not match any signature seen before by such noise sources. It was discovered that a change to the laser supply voltage several months prior to the OR deployment was causing the light



output to oscillate at times. While this impacted the data collected for the OR, the characteristic shape of the false particle episodes was fairly easy to recognize, and could be qualified against the other sensors in the room.

Summary of Operating Room Testing

This experiment was entered into without having much idea of what data might be found. Some significant changes might be made to equipment and procedures for better correlation with actual particulate matter sources in future experiments. While the experiment with the hospital operating room was not performed in the most optimal way, several conclusions may be drawn from the data.

First, there is a relatively high level of airborne submicron particulate in the OR, which appears to be due mostly to particulate entering the OR from the outdoors through the ventilation system. While the experiment did not include a unit set up to measure the air at the ventilation system's inlet, based on a unit operating concurrently at another site,

it would appear that the filtering mechanisms at the hospital aren't removing much of the ambient air particulate matter before it enters the building. Since the particulate is submicron, it is very likely not microbial in nature, so likely poses minimal risk to patient infection.

Second, some events in the OR during surgeries appear to generate significant amounts of submicron particulate matter. The experiment setup was not able to correlate this directly with any particular activity, but it might originate with the use of electro-surgical apparatus.

Third, some events in the OR appear to generate small peaks in particulate matter greater than 1 μm . The experiment setup was not able to correlate this directly with any particular activity. Some of these events appear to bracket a procedure, and could be due to movement in and out of the OR, or the movement and/or activation of materials or equipment that had been idle.

Fourth, there was no clear evidence of airborne movement of particulate matter from the hallway or substerile room into the OR. This could be simply due to the fact that most high-concentration events were occurring in the OR itself, rather than in the hallway. A contrived high-concentration event in the hallway might provide more conclusive results.

Fifth, the humidifier does not appear to be producing any of the submicron particulate in the OR.

Improvements might have been made to the quality of data if there were better correlation between actual events and sensor data. In this case, the direct observation of

the operating room during active surgeries was not possible, so the correlations were not available. As a result, the causes of particulate matter generation can only be surmised, not proven. Another significant factor missing from the experiment was the inability to monitor air entering the ventilation supply from the outdoors. This would have been helpful in determining whether or not the source of particulate matter was outdoors or within the hospital itself.

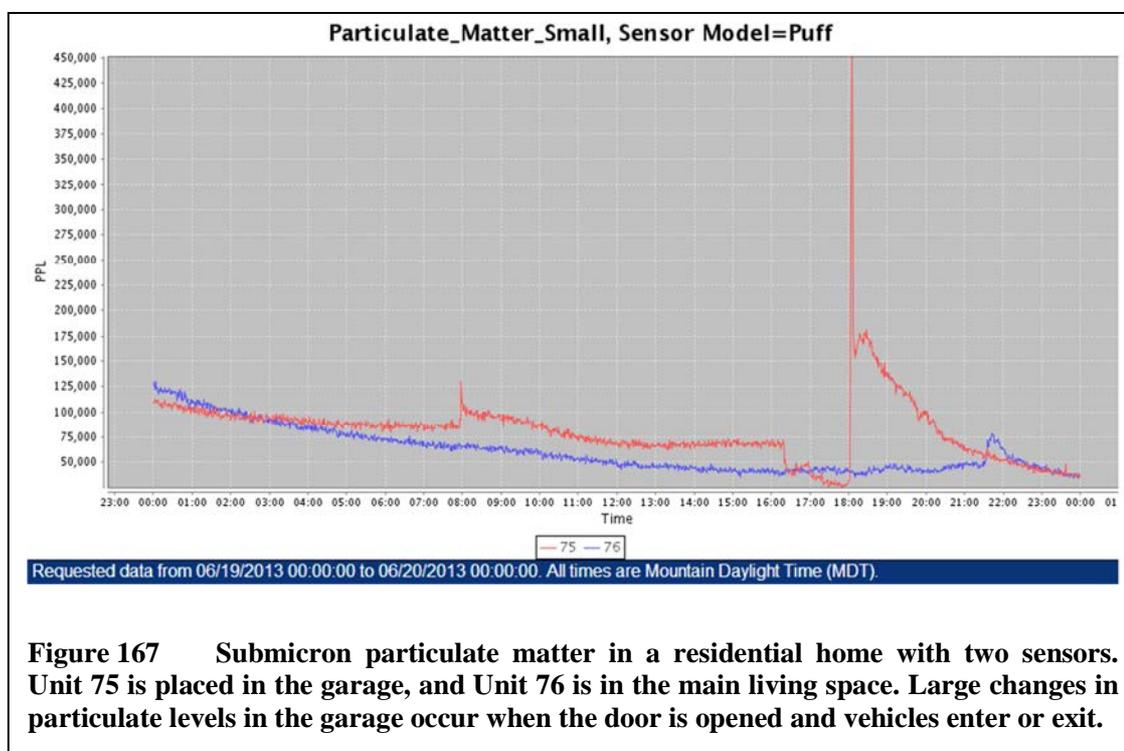
Field Environment Testing in Residential Homes

The air quality within residential homes is of great interest due to its potential impact on the health of the occupants, especially since most people in the U.S. spend 90% or more of their time indoors [21]. In general a person might have access to data regarding the outdoor air quality in the city where he or she resides, but not much quantitative data regarding the indoor air quality experienced daily. For example, a person may not understand when it is best to open windows for fresh air. The levels of particulate matter, carbon dioxide, relative humidity, and temperature can vary greatly indoors depending on many fixed and changing, environmental and activity factors. Taking advantage of data such as those provided by the WAQM system can provide insight into the indoor air quality of a building, which can help one formulate actions that may improve living conditions.

Understanding the interaction between the environments in a home's automobile garage and main living spaces can demonstrate a potential for some health hazards. Testing was performed in modern home with a pair of WAQM sensors, one located in the main living space of the home and the other in the garage. The next several figures show data taken over the course of a day. This was a work and school day for the occupants,

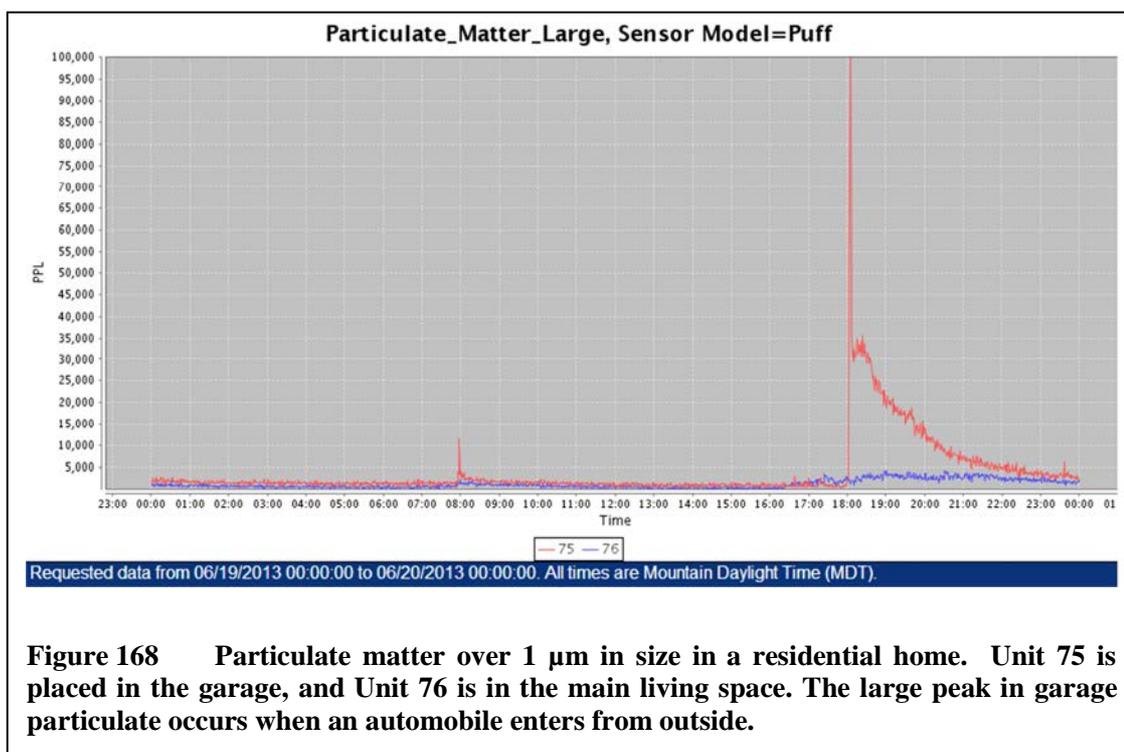
who left home in the morning and returned late in the afternoon. The home sat empty for most of the day.

Figure 167 shows the submicron particulate matter over the course of a day in the home and garage. Note that the particulate matter concentration tends to change drastically in the garage when the door is opened. The owner of the home left for work at approximately 08:00, and returned at approximately 16:20. The garage door was opened again at 18:00. The large changes in particulate concentration in the garage may be due to the sudden exposure of the garage to the outdoor environment when the garage door opens, or could be due to particulate generated by an automobile entering or leaving the garage. Negative shifts are most certainly due to the former, but the cause for positive shifts is more difficult to determine without having an outdoor reference available. The



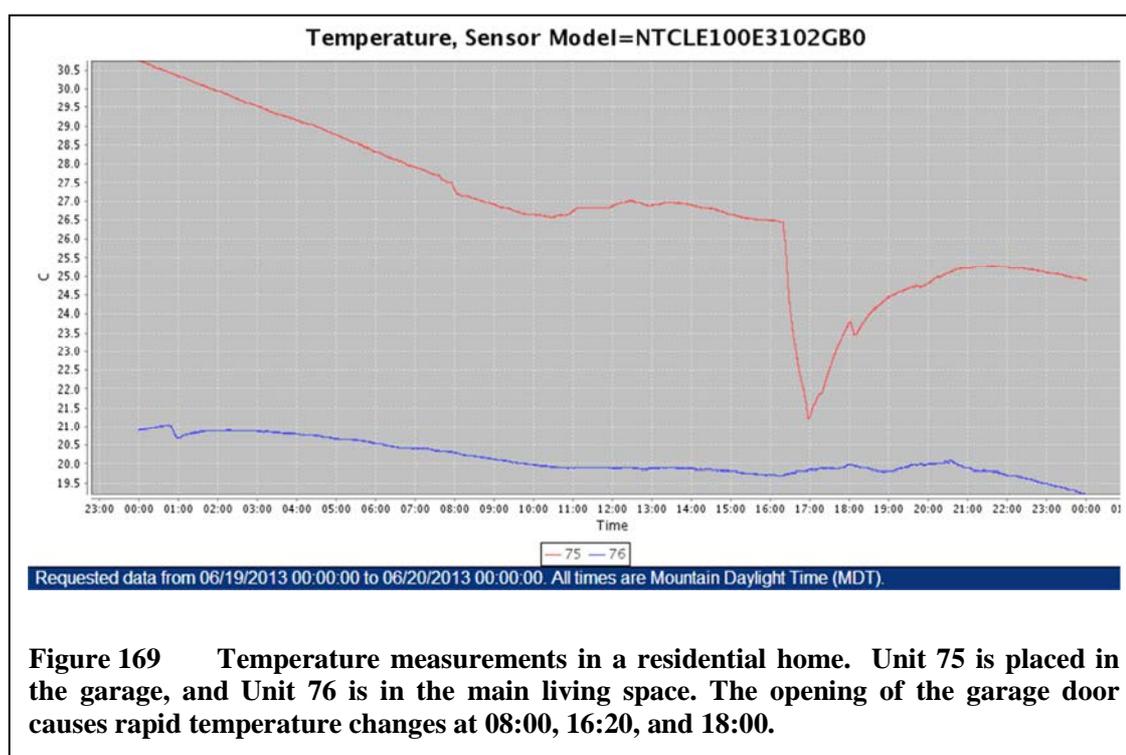
peak in particulate in the main living space at 21:30 may be due to cooking or some other occupant-caused activity.

The measurements of particulate matter over 1 μm in size show a similar profile, but with peaks in particulate only appearing for certain garage door openings. Figure 168 shows the large particulate matter data from the same time period. The small peak at 08:00 in the garage corresponds to the shift seen in submicron particulate when the garage is opened in the morning. This peak in particulate matter is likely due to either exhaust or dust disturbed by an automobile leaving. Even more interesting is the large peak in particulate seen at 18:00 when an automobile returns. Even though another automobile returned previously at 16:20, it did not cause such a large spike. It is very unlikely that this much large particulate was present in the ambient outdoor air, and was likely either exhaust or dust brought in by the automobile.



The temperature profile measured for the same period is shown in Figure 169. The temperature in the garage tells a more complete story about the opening and closing of the garage door. The internal temperature of the garage appears to have been elevated from the previous day's temperature, and dropped rapidly any time the garage door was opened. The door was opened briefly at 08:00 when the automobiles left the garage. The door was left open from 16:20 to 17:00 when the first automobile returned. The door was opened briefly again at 18:00 when the second automobile returned. Temperatures inside the home are relatively stable, with air conditioning controlling the temperature over the course of the day.

Carbon dioxide levels in the home and garage for the same period are shown in Figure 170. The levels inside the home are a good indication of human presence. Note that levels in the home tend to increase during the night, and fall off during the day while



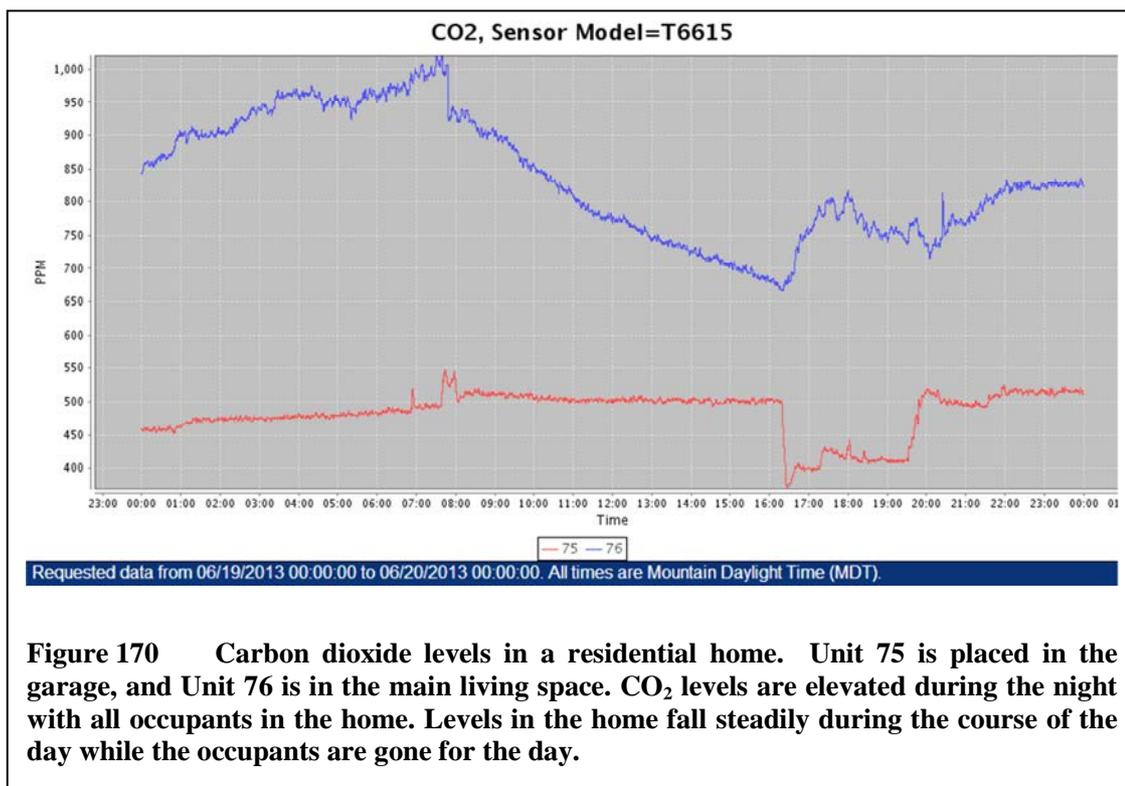
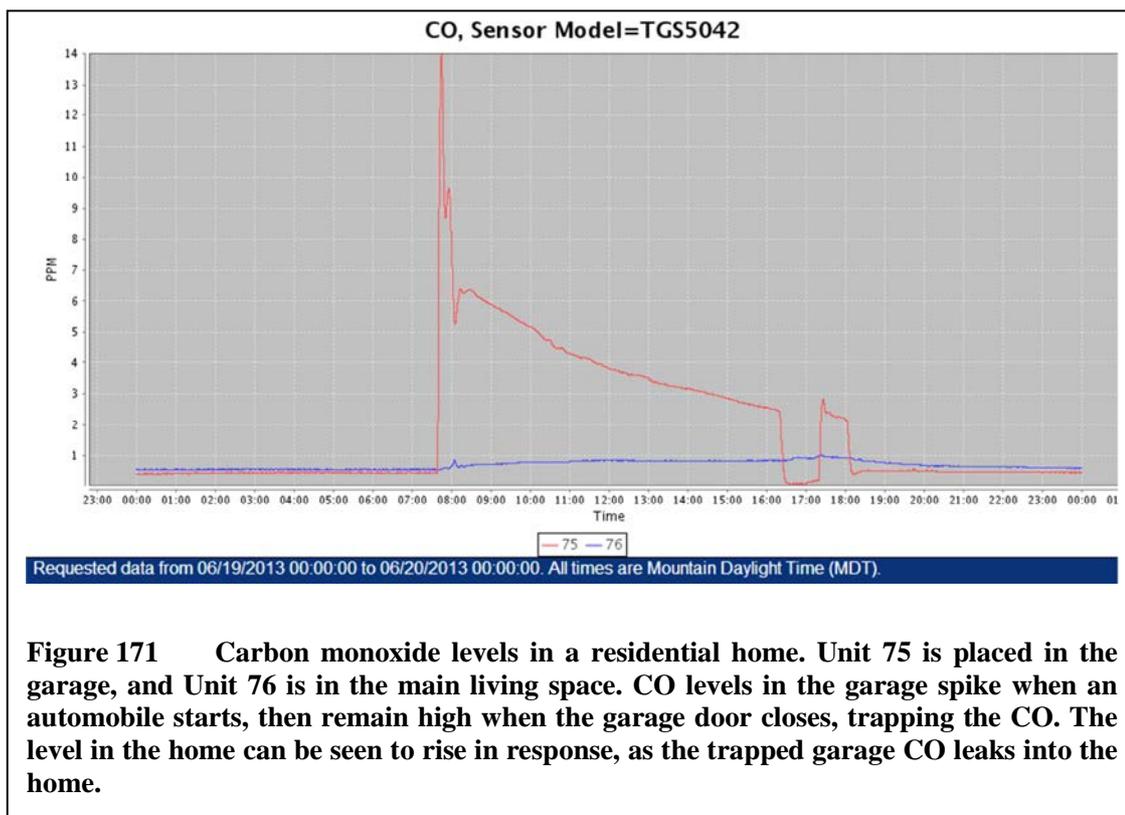


Figure 170 Carbon dioxide levels in a residential home. Unit 75 is placed in the garage, and Unit 76 is in the main living space. CO₂ levels are elevated during the night with all occupants in the home. Levels in the home fall steadily during the course of the day while the occupants are gone for the day.

the home's occupants are away. Levels in the garage are less clear, with openings sometime causing a rise in CO₂ levels and sometimes a fall. This may be dependent on the amount of automobile exhaust trapped in the garage at the time the door is opened. The rise in CO₂ in the garage at 19:30 does not appear to correspond to any of the garage door openings, and may be due to the presence of one or more of the home's occupants in the garage at this time.

The carbon monoxide levels in the garage and home show an interesting correlation between the two spaces. The CO data for the same period are shown in Figure 171. The CO from automobile exhaust is captured in the garage in the morning, and slowly diffuses out during the day. While there does not appear to be any direct generation of CO inside the home itself, there does appear to be leakage from the garage into the home during the course of the day. While the absolute levels are not very high,



the trapping of CO in the garage does cause elevated levels in the garage and the home. The high levels in the garage could be hazardous to people working or pets living there, especially since the repeated activity of the automobile leaving the garage would likely cause repeated exposure to any occupants.

CHAPTER TEN: FUTURE WORK

Given the broad set of problems encountered in the development of the WAQM system, there is an equally broad set of areas for improvement and enhancement. As the WAQM system evolves over time, some of these areas may be considered and developed. This section breaks down improvements into an organization matching the layout of the previous chapters. Improvements are discussed for the WSN system, the OPC hardware, the simulation environment, the detection and sizing algorithms, the particle calibration system, and the event detection framework.

Wireless Sensor Network System Improvements

There are a number of improvements that might be made to the WSN system as it evolves. Several issues impact the performance of the system, and might be critical to address in future developments. Three of these issues directly impact the delivery and integrity of the data stream, including the mesh network software, the buffer space in the system, and misbehaving tasks.

Mesh Network Software

The mesh network software stack currently used in the WAQM system is provided by a commercial vendor. This software removes much of the low-level complexity of forming and running a mesh network. This greatly simplifies the effort in building the mesh networking solution, but the details of the low level operation are

hidden from the development team. For most situations, this is acceptable. However, when the network begins to see congestion, such as when a modest number of nodes join the network, things can go wrong.

In testing of the WAQM system, it was initially found that communication using relatively inefficient text-based messaging would limit the number of nodes on the network to approximately 12 units that were generating on the order of 10 messages per second each. Moving to a binary message format and aggregating multiple measurements into each packet pushed this up to approximately 20 units. Beyond this limit, the network would become unstable, typically seeing a large increase in dropped packets and a significant drop in throughput. Adding back off mechanisms to the network layer helped the system recover from congestion, but overall throughput was as much as an order of magnitude less than the advertised 250 kbits/second of the ZigBee radio modules.

Taking on the task of developing the low-level mesh networking software could help increase overall system throughput. This is difficult to quantify, since the current commercial stack reveals very little of what is occurring when the network starts to see congestion. There is risk in taking on this development, as at the end the performance may be no better than that of the commercial stack. Regardless of the outcome, better control of the lower layers should remove at least some redundancy between the WAQM network layer and the ZigBee software stack, with the potential to improve throughput significantly.

Flash Buffering

In the current system, when the mesh network sees congestion, the buffering mechanisms in the WAQM sensor node start to fill. These buffers are relatively small due

to the limited amount of memory in the microcontroller, and are only able to hold a small number of packets before overflowing. When this occurs, the measurements are lost from the data stream and can only be recovered by manual retrieval from the SD memory card in the system.

One method of addressing this data loss would be to buffer data through the flash-based SD memory card. Since the SD card is typically written to by default in parallel with the mesh network, a system could be implemented where messages were first written to the SD card and then read from the card to be sent into the mesh network. This flash buffer would hold a very large amount of data in the case that the network connection was lost, allowing the data to be sent once the network connection was restored. The flash buffer could be permanently in line with the network data stream, or could be implemented just as a backup store. The power, timing, and complexity of various different solutions should be evaluated before proceeding with a design.

Implementing the flash buffer would potentially introduce a host of other issues to the mesh network. When the network sees congestion in the current system, data loss is one of the mechanisms allowing the network to recover. If all the data were stored to be sent later, the network may never fully recover once the sensor nodes start to build up a backlog of data. The aggregate network bandwidth must be weighed against the advisability of sending all data generated by the sensor nodes. This might be alleviated by a data rarefaction algorithm that only sends a sparse representation of the data stream once congestion occurs. Much research exists on mechanisms to handle these issues, and extensive study should be undertaken to find solutions that apply.

Preemptive Multitasking

The cooperative multitasking model employed by the WAQM sensor node has functioned relatively well in most cases. The simple task loop is easy to manage, and as long as the tasks behave by taking relatively small amounts of processing at any given time there are no issues. One place where the system falls apart is with respect to the SD memory card. The firmware for the file system has been leveraged from the microprocessor vendor's set of example code, and is quite complex and difficult to modify. Multiple issues causing the code to lock up have been addressed, but a fundamental issue with blocking remains. When data are passed to the file system, it will attempt to write it to the SD card and will block until this occurs. Depending on many different factors impacting flash writes, this could take as long as 100 milliseconds, an extremely long time in terms of the system. When the file system blocks, it is not typically doing useful activity, but is simply waiting for the flash write to complete. When this happens, measurement timing may be displaced, or data may be lost. This is particularly bad for the OPC, as the sampling buffers will quickly overflow without being processed in a timely manner leading to holes in the sampling stream.

This type of issue is not addressable in the cooperative multitasking system without rewriting the entire file system, a task not seen as particularly valuable or rewarding in terms of the development of the WAQM system. One method of addressing this would be to move to a preemptive multitasking model. Having multiple threads executing in parallel would allow code that needs to block to exist alongside code that is time critical. Future designs of the HSIL system are looking to use a preemptive, multithreaded operating system to overcome these types of issues.

Particle Counter Hardware Areas for Improvement

A large part of the work for the development of the WAQM OPC has been focused on the particle counter hardware. In particular the design and testing of the AFE, the flow channel, and the light path were iterative processes that evolved the design over time. Further improvements are certainly possible, and worth pursuing to create better versions of the OPC. Several candidates for improvement to the hardware are listed.

Improvements to the AFE

The AFE circuitry has evolved to a point where decent results are achieved with a fair amount of external noise immunity. The most serious remaining issue with the AFE is the total RMS noise. This noise currently consumes three least significant bits of the 10-bit ADC values, which represents a significant portion of the system's accuracy in detecting and sizing particles. The high gain of the small channel exacerbates the problem, with the largest noise source being the input referred voltage noise of the operational amplifier used for the transimpedance amplifier. Finding an operational amplifier with a lower input referred voltage noise would be a direct way to help address the issue. Other circuit changes may also yield improved total RMS noise results, allowing the lowering of the detection threshold and the more accurate determination of size based on pulse height.

Laser Source

The laser source is one of the most critical components of the OPC design. The power output, beam waist characteristics, and stability of this component are key parameters to producing repeatable measurements with sensitivity to sub-micron

particles. The work for this dissertation has been performed using laser modules purchased from third parties, where the laser diode is pre-mounted in a metal housing with power-control circuitry. In targeting low-cost modules for the prototypes, the choices for defining these parameters are limited. In particular, requiring an adjustable focus on the module precludes the use of many small laser modules that are preset to a collimated beam for use in laser pointers.

One direction to investigate in particular would be the integration of discrete components into the OPC housing rather than using a pre-assembled module. The discrete components might include a laser diode, a plastic lens, a heat sink diode mount, and laser diode driver circuitry. The laser diode should be an industry standard package, such as the standard 5.6 mm housing, and should include a feedback photodiode to measure power output. The plastic lens may be acrylic with or without antireflection coating. The heat sink diode mount is necessary to manage the temperature of the laser diode, replacing the metal sink of the laser module housing. The diode driver circuitry could be customized as necessary.

With these discrete components, one could better control some of the key parameters of the laser source. For example, control of the driver circuitry design could allow variable power operation that helps to preserve battery life or extend the working life of the laser diode. By controlling the mounting of the laser diode on the heat sink, one could better control the beam waist measurements, and one could also expose the heat sink to the exhaust path of the airflow channel to help cool the diode for longer working life. Direct control of the lens could lead to more optimal total length of the laser source and more efficient capture of laser light as it diverges from the diode. Also,

splitting the module into discrete components might allow direct sourcing of the components from a wider set of vendors, potentially helping to reduce costs and achieve better assurance of supply.

These cost and performance gains would need to be measured against the added complexity of sourcing more discrete components and the increased complexity of assembly for the product. In the end, these additional costs might be better spent on a more expensive custom device from a laser module vendor, but further study will be required to reach this conclusion.

Airflow Improvements

The flow channel of the WAQM OPC has been developed as an evolutionary process over several generations of designs. The original design was very angular, with emphasis placed on blocking ambient light at the cost of a smooth transition from ambient air to the sensing area. This has since evolved into a much improved airflow with features that allow better retention of particulate matter in the air without sacrificing too much in the way of ambient light immunity. Though largely a mechanical and aerodynamic exercise, further study and experimentation in this area is warranted to improve the performance of the OPC.

Light Management Improvements

While unwanted incident and reflected laser light is controlled fairly well by the countermeasures implemented in the WAQM OPC, there are still issues with unwanted ambient light. This is particularly true for devices used outdoors, where ambient light levels can be very high. Currently, units deployed outdoors must be shielded from

ambient light to prevent measurement issues when the ambient light reaches the photodiode. Further work on the flow channel at both the inlet and outlet are required to allow a unit to operate in outdoor ambient light levels without additional shielding.

To accomplish this, one might add light blocking features into the flow channel. This may require lengthening the flow channel to further remove the photodiode from the source of ambient light, and to allow other features such as baffles or turns to be added to the channel. Great care must be exercised in doing this, as the added features may also lead to a reduction in airflow, or a loss of larger particles due to impaction on the features.

Simulation Areas for Improvement

The main purpose for the creation of the simulation flow was to provide insight into the development of the WAQM OPC. Fairly early in the development cycle it was realized that a more analytical method was needed to explain the results being seen in testing. As such, the goal has been to provide insight into what is happening with the OPC under development, rather than to provide a general simulation tool. This targeted mind set has perhaps constrained the tool where further expansion could allow it to be employed for a more diverse set of problems. In this section, a brief set of possible improvements is explored as guidance to future development of the simulation tool.

Front-End Improvements

The tool as it exists is a set of MATLAB scripts, and requires significant understanding of the script implementation details to operate effectively. At a high level, this constrains operation of the simulation tool to a select few who might spend the effort

to understand the scripts to a level required to achieve successful operation. To open up the simulation tool to a broader audience, a better front-end might be developed.

The first challenge in developing a better front-end will be in selecting which simulations to better automate. The scripts as they exist provide a great deal of flexibility in generating different scenarios and output plots. A set of scenarios would need to be chosen to target towards front-end development. Along these lines, something that produced an operational simulation with varying sensor positions might be most useful. A graphical user interface could be produced that showed a representation of the sensor arrangements, and then output simulation results for the arrangement.

Laser Beam Intensity Assumptions

One of the largest areas of discrepancy between the simulations and the WAQM OPC implementation is in the modeling of the laser. The simulation assumes that the laser is a Gaussian beam with a circular profile. The actual shape of the laser beam in the WAQM OPC is much more complex. To start, the typical edge-emitting laser diode produces a beam of light that is highly elliptical in shape, diverging at a much greater rate normal to the diode junction than in the plane of the junction. Once the beam passes through the collimating lens, the intensity profile remains very elliptical in shape. Since most of the simulation work used particle positions in the center of the beam relative to the x-axis of the system coordinate framework, and the employed polarization kept the minor axis of the ellipse aligned with the x-axis, the shape of the beam was deemed to be not particularly important. This could change if a different polarization were desired, or if simulation of the particle movement through the beam were performed.

There are also other issues with the beam that likely cause deviation from a simple Gaussian profile. The module used in the WAQM OPC is of relatively low quality, and does not appear to be well suited for a focal point so near to the laser diode. The seating of the laser diode within the module is not ideal, and allows a great deal of play prior to the diode being glued in place. This results in a variety of beam orientations and shapes. When the lens is positioned for a 35 mm focal point, some portion of the beam often intersects the brass casing around the lens rather than the lens itself. This causes beam artifacts that lower the efficiency of the module and may change the profile of the beam.

An example laser beam profile from a WAQM OPC laser module is shown in Figure 172, taken with a high-resolution Newport Laser Beam Profiler. The profile is somewhat elliptical in shape, with a number of partial ring-like structures especially on the long axis. This non-uniformity may be due to lens imperfections or beam self-interference. This type of fine structure varied greatly between laser modules, and no attempt was made to model this for the simulations.

The simulation framework also assumes that the incident light at the particle is a plane wave of constant irradiance and uniform direction for an infinite width. This of course is not true, more so for positions nearer to the beam waist where the irradiance may change drastically over a very small distance. A study by Tsai and Pogorzelski [90] explores the impact of a finite beam on scattering by spherical objects, showing that the scattering intensity decreases as the particle diameter approaches the beam diameter. The results also suggest that there will be minimal impact on scattering intensity if the beam diameter is in the range of 10 times the diameter of the particle. A more accurate

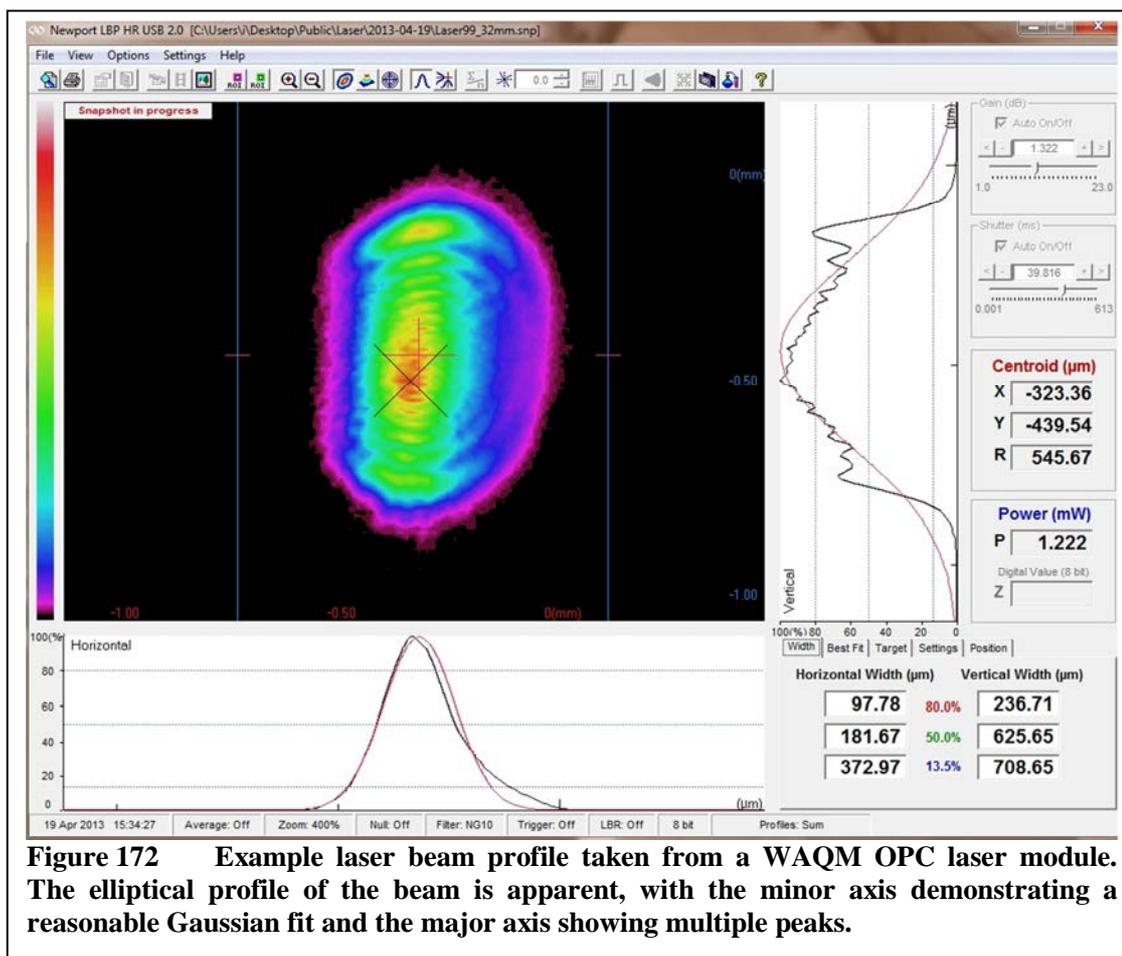


Figure 172 Example laser beam profile taken from a WAQM OPC laser module. The elliptical profile of the beam is apparent, with the minor axis demonstrating a reasonable Gaussian fit and the major axis showing multiple peaks.

simulation could be produced by taking the finite beam width into account, but at the cost of increased simulation complexity.

Particle Shape Assumptions

The only particle shape currently supported by the simulation is a sphere. Extensive research has been done into the modeling of scattering from other shapes such as cylinders [66], ellipsoids [91], and other arbitrary shapes [92], which might be incorporated into the simulation framework. For the purposes of the WAQM OPC development, it was thought that the spherical assumption was a good enough approximation for the expected use. In particular, the main vehicle for calibration of the

WAQM OPC is the PSL sphere, which is best approximated by a spherical simulation. Further study might be done on specific particulate types that the WAQM OPC may measure in the future, which could lead to a need to model other shapes.

Simulation of X-Axis Movement

The simulations in their current form are used to find the peak scattering intensities produced by a particle as it moves through the air flow channel and intersects the laser. This is a two-dimensional process in assuming that the peak intensity will come from the intersection of the particle with the $x=0$ point in the flow channel, where the laser is most intense.

To expand upon this, there may be some use in incorporating a position-dependent model of the particle velocities in the channel. This would allow the extraction of the scattering intensity over time could occur as the particle passes through the beam. This could then be analyzed through a circuit simulator to build actual waveforms as would be seen from the output of the OPC amplifier circuit. Incorporation of the circuit model may lead to a better understanding and balancing of the broader interaction between the air flow, laser, sensor positions, and circuit design.

Detection and Sizing Areas for Improvement

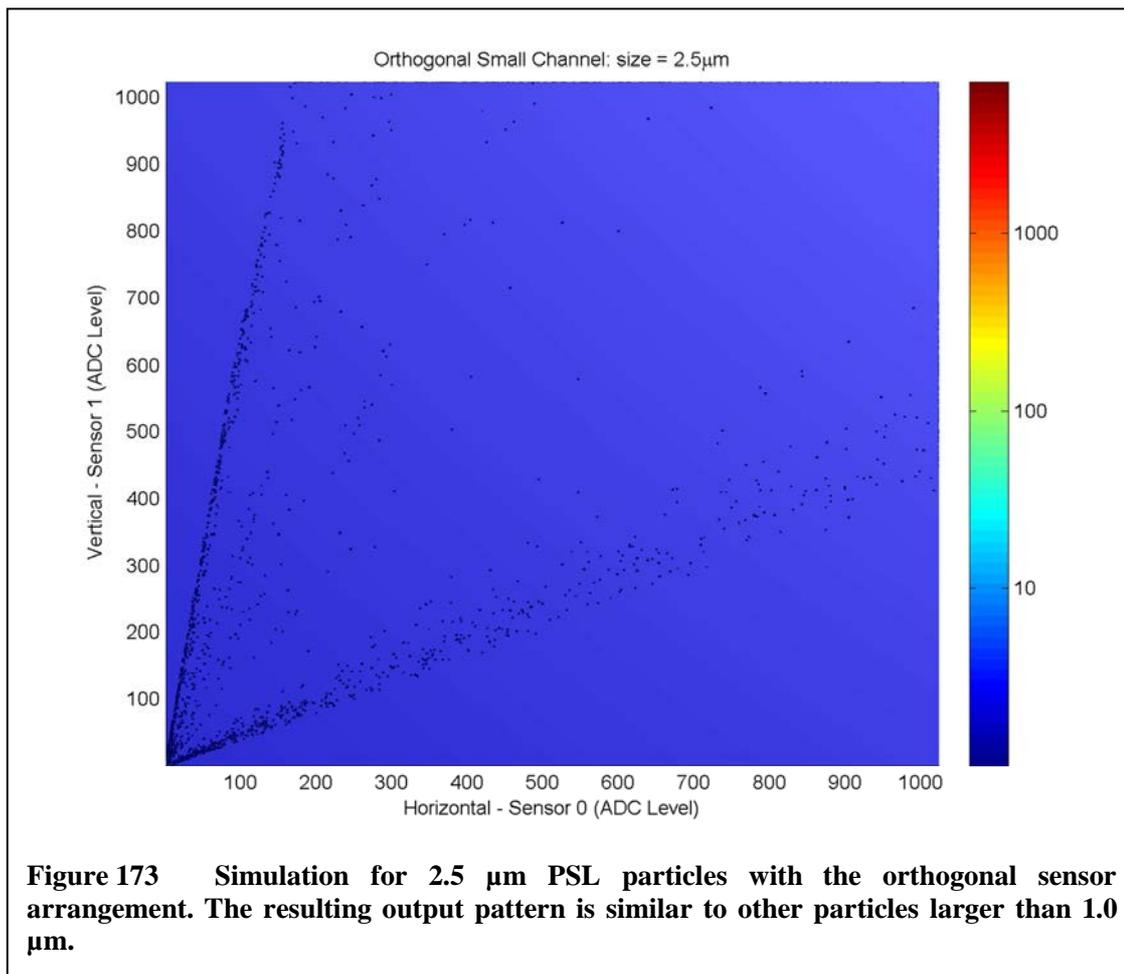
As shown, performing detection and sizing with the WAQM OPC is a challenging endeavor, with room for improvement in many areas. The creation of the sensing volume by passing a focused laser through a constrained airflow leads to uncertainty in generating accurate particle counts for a given size. The following is a brief discussion of some of the areas for improvement that may be undertaken as future work.

Sizing Larger Particles

As previously mentioned, the WAQM OPC has set its large particle size threshold at 1.0 μm . One significant reason for using this size as a threshold was in the inability to produce an aerosol from particles of sizes much greater than 2.0 μm with the development equipment available. There will however be some specific interest in setting thresholds at larger particle sizes. In particular, if classification meeting the $\text{PM}_{2.5}$ and PM_{10} designations were desired, a sizing threshold at 2.5 μm would be needed. This would categorize particles under 2.5 μm diameter as $\text{PM}_{2.5}$ and particles over 2.5 μm diameter as PM_{10} .

Further improvements in the particle generation system will be needed to accomplish this. To accurately characterize the threshold, the reliable generation of particles larger than the desired threshold is required. The potential issues with the particle generation system and the creation of an aerosol from particles of size greater than 2.0 μm will be discussed further in the chapter on calibration.

Once an appropriate particle generation system is available, moving the large particle threshold to 2.5 μm should be relatively straightforward. Particles of sizes in this range and larger tend to scatter light similarly to particles in the 1.0 μm range. Figure 173 shows the predicted scattering intensities for 2.5 μm from the horizontal and vertical sensors from the simulation framework. As with the scattering from 1.0 μm and 1.6 μm PSL particles shown in Figure 71 and Figure 93, the 2.5 μm particles produce strong groups of pulses with high and low horizontal to vertical intensity ratios. The method used in differentiation of the 2.5 μm and larger particles from those of smaller diameter will depend on the sensor arrangement. For the single sensor arrangement, the intensity



threshold can be moved up to reject a larger proportion of the particles falling under the new size threshold. For the orthogonal sensor arrangement, the pulse duration threshold would need to be increased to raise the sizing threshold. In both cases, further study should be done with the appropriate particle generation systems to ensure that the sizing algorithms will produce acceptable results.

The current implementation of the WAQM OPC does not attempt to limit the counting of particles larger than a certain size threshold. To create an accurate PM₁₀ concentration, the particle sizes greater than 10 μ m would need to be removed from the air flow physically, or removed from the count algorithmically. To physically remove the

particles, a method for separating these large particles from the air flow would need to be added to the WAQM OPC, either with filtering, impaction, or some other means. The addition of a filter would likely reduce the air flow through the device, add cost to the bill-of-materials, and increase the size of the OPC. Impaction is a potential candidate that could remove the bulk of large particles by simply adjusting the OPC inlet shape. Further work would need to be done to model and test this type of large particle removal. Algorithmic removal of the large particles is also possible, though pollution into the desired size range would always present a problem due to the laser beam intensity variation across the channel. As with physical removal, further study would be required to implement this type of solution.

Further Testing of the Orthogonal Sensor Arrangement

In comparison with the single sensor arrangement, the orthogonal sensor arrangement is relatively new. The particulate matter sizing and detection methods based on the orthogonal design have mostly been tested in the lab with a very small number of prototype devices. The bulk of the WAQM OPC systems have been built with the single sensor arrangement, and the system demonstrations shown in later chapters have been performed with these devices.

The orthogonal system design is at the point where it is ready to be produced in higher volumes and deployed in new systems. With additional units available, more work may be done to compare results between units and evaluate the effectiveness of the detection and sizing algorithms in dealing with unit-to-unit variations. Since the orthogonal sensor arrangement does not greatly add to the cost or size of the OPC, future devices should include the necessary sensor hardware. If the large channel is left as an

option for the horizontal sensor, the single sensor arrangement can still be supported as a fallback in the case issues are found that limit the effectiveness of the orthogonal arrangement.

Other Algorithms and Metrics

As part of the development of the WAQM OPC and the various sensor arrangements, some other methods for particle detection and sizing were examined. While not included in the current WAQM OPC design, these methods may merit further study and may be appropriate for different combinations of sensors and processing hardware.

A relatively simple metric that might prove valuable in particle sizing would be the use of pulse integration. For this method, the sum of intensities across the duration of the particle would be collected and used as a metric in sizing. There is some possibility that the pulse integration value would represent intensity and duration well enough that the two metrics currently used could be replaced by a single value. This was examined with a moderate amount of experimentation during the course of the WAQM OPC development, and it was deemed that there would not be enough data in this metric to differentiate between long duration, low intensity pulses caused by large particles moving through a part of the laser far from the beam waist and short duration, high intensity pulses moving through a part of the laser close to the waist. Even though it does not currently show promise as a replacement metric, further study could show some value in adding it as an additional metric to help in size differentiation or large particle rejection.

A more advanced method for recognizing false particles, beyond those mentioned previously, has been explored that employs a wavelet transform. Early on in the

development of the WAQM OPC, it was found that the analog front-end design was susceptible to AC line noise when powered from a DC wall transformer. As discussed above in the chapter on OPC design, further investigation showed that a fair amount of this noise could be attributed to silicon controlled rectifiers operating near the OPC. This type of noise produced pulses on the small channel of the OPC that could be distinguished from actual particle caused pulses based on the pulse shape. To detect these noise generated pulses, an algorithm was developed using a Haar transform to determine whether or not the pulse shape matched a profile that should be rejected. Once the particle was detected using the normal means mentioned above, a predetermined number of ADC values for a portion of the pulse starting just prior to the detected start of the pulse would be processed. The wavelet transform provided a convenient method for providing metrics on the shape of the pulse temporally and spectrally. In this case, the features in the pulse indicating noise could be highlighted by the transform.

This method showed some promise in the detection of false particles when simulated, but was rendered unnecessary by hardware changes to the analog front-end of the OPC that electronically rejected these same pulses. Even if employed in the WAQM OPC, some severe challenges were still to be overcome. Not the least of these would be the processing load to implement the algorithm on a per-particle basis. The requirements to continuously store a length of the pulse ADC levels, and the processing of the transform, would likely have severely limited the number of particles that could be processed in a given time. However, there may be some interest in pursuing this method, especially if the system could afford dedicated hardware in the form of an FPGA or higher performance processor to implement the algorithm.

Particle Counter Calibration System Improvements

As the OPC calibration system was not the main focus of the project, but was still critical to its successful implementation, there are many areas where improvement may be had. This includes known deficiencies as well as areas where further study will be required. Several of the key issues where improvement may be found are examined.

Large Particle Aerosol Generation

The most apparent known issue with the OPC calibration system is in the inability to generate particles in the 2-5 μm range. This range is critical for differentiation between $\text{PM}_{2.5}$ and PM_{10} particles, which may be important to some future studies. As mentioned, this may be due to multiple factors including low air pressure and low PSL solution concentration. This issue should be investigated and corrected prior to the start of any study requiring this level of differentiation.

Monodisperse Aerosol

Improvement of the quality of the aerosol produced by the OPC calibration system is one of the most impactful areas where further study and experimentation may raise the quality of the WAQM OPC design. Having a truly monodisperse aerosol with which to work would allow better correlation with simulation results and help in understanding the system response to a known particle size. The best aerosol that the system can currently produce contains a broad mix of particle sizes, even when starting with a known calibrated PSL solution. There is always a small particle component that is present in the aerosol alongside the PSL particles, presumably from the PSL solution surfactant or other substances in the system. The PSL particles themselves also display a

fairly broad size signature, possibly due to incomplete drying prior to entering the test chamber. Working with polydisperse aerosols introduces a fair amount of guesswork in identifying how particles of a particular size impact the OPC response.

A first simple solution would be to test adding more diffusion dryers into the system. This would have to be carefully monitored to make sure that the additional path length in the system between the nebulizer and the test chamber did not consume too much of the particulate matter in the aerosol. Additional drying could significantly reduce the amount of moisture in the aerosol. While a single dryer might reduce the relative humidity of the aerosol from 60% at the inlet to 20% at the outlet, a second dryer might reduce this to 8% [83]. Reducing the moisture content of the aerosol should help to produce a size distribution that is closer to that of the calibrated PSL particles.

Other options are much more expensive. For the best results, adding a differential mobility analyzer (DMA) to the system would allow the production of an aerosol that was very close to monodisperse. Unfortunately these devices typically cost more than US\$50,000, which was well outside the budget for the development of the WAQM OPC. These devices also may have issues in dealing with large, supermicron particles that would be of interest to the development of future WAQM OPC designs.

Other devices such as condensation monodisperse aerosol generators (CMAGs) can grow monodisperse aerosols from stearic acid or carnauba wax, using atomized salt solution as the initial seed particles. These types of devices are also expensive and may be messy, producing particulate films that must be regularly cleaned from the OPC calibration system. These factors would need to be considered prior to adding such a device to the OPC calibration system.

Monitoring Equipment

Along with the ability to generate a good monodisperse aerosol, it is important to have the ability to accurately monitor the aerosol composition. For the current state of the OPC calibration system this is less important than the generation piece. This is due to the including of the PMS Lasair II 110 into the system, which is a relatively high quality OPC. An additional device that may be of significant interest in characterizing the aerosol produced by the OPC calibration system would be an aerodynamic particle sizer (APS), which uses pulse height and particle velocity to estimate the size of a particle. These devices can typically provide a much higher resolution view of the spectrum of particle sizes in an aerosol, which would be very valuable in monitoring generated size distributions. As with most quality particulate matter measurement devices, these are typically quite expensive and difficult to acquire with limited funding.

Cold and Hot Testing

The development of Cold and Hot testing for the calibration of the WAQM OPC provided a low-cost method for creating aerosols with two distinct size distributions. This allows the calibration of the sensitivity and absolute particle count for the small channel of the OPC. The process as used in the development of the OPC depended on the use of tap water, a decidedly uncontrolled substance. While fairly consistent results were achieved in the laboratory environment where the OPC calibration system was deployed, there is no guarantee that future results will hold. As such, it would be preferable if the source water solution could be created from known quantities, both to better control the amount of particulate matter generated and to assure consistent future results. A possible route to test would be the use of a salt or other soluble compound that would not be

harmful to the OPC calibration system or its operators. Further study and testing will be required to produce a reliable solution.

Aerosol Mixing

In general the OPC calibration system appears to demonstrate relatively consistent mixing of the aerosol as it enters the chamber. A small fan blowing air across the aerosol inlet to the test chamber helps ensure that the particulate matter is dispersed throughout the volume. However, occasional anomalous results may indicate that certain situations or device arrangements might lead to insufficient mixing. Further study and experimentation will be required to prove that a problem exists, and to test possible solutions.

A secondary concern is the inefficient use of the testing volume. The current chamber is taller than either dimension of its footprint, but only the bottom of the chamber is used to hold devices. A rack system might be introduced to the test chamber that allowed multiple levels of devices to be tested simultaneously. However, pursuing this path must include a fair amount of testing to make sure that the introduction of a vertical component in the device distribution does not skew the calibration results. For example, effects such as gravitational settling might impact higher levels in the stack more than those on the bottom.

Event Detection Future Work

Several features originally defined for the event detection framework were not completed prior to the writing of this dissertation. Adding these pieces to the framework will make it easier to interact with from the user's standpoint, and will set it up for the

addition of learning algorithms in future work. This additional work includes the implementation of WSN-Level Event Triggers, server acceptance of input user labeling of events, and the dynamic programmability of event triggers from a user visualization device or server.

Each of these has been described in the event detection chapter. The first feature is really just an addition to the framework that will make event detection more useable at the WSN-level. The second and third features are important to the future addition of a learning algorithm to the WAQM system. The ability for the server to dynamically change triggers will allow a learning algorithm at the server to tune the triggers to provide an improved set of event notifications. The learning algorithm could then work with user input labeling of the events to tune the triggers.

Analytics

While the collection of data with a set of sensor nodes is a first step towards understanding an environment, a higher level of machine intelligence can be applied in examining the data and looking for trends post-collection and in real-time. Higher-level analytics of the data sets coming from the WAQM sensor nodes was mostly performed in a manual fashion for this research. Future efforts might employ data analytics in an attempt to quantify trends and recognize large-scale or long-term events not easily discerned by human observation.

Much of this work may need to be implemented at the server level of the WAQM system due to the processing requirements, though the sensor and coordinator nodes may provide a first level of pre-processing. Taking advantage of the vast quantities of data available from these types of sensor systems will become more common in the future as

industry and academia focus on solving world-scale problems through sensing and data mining. The WAQM system and its subsequent generations should play a part in this, leading the way with novel methods for air quality monitoring.

CHAPTER ELEVEN: CONCLUSIONS

The Environmental Protection Agency and other organizations have made great strides over the last several decades toward providing regulation, information, and awareness about air quality. Even so, there is a gap between the air quality experienced outdoors where such institutions monitor air quality and the air quality experienced by the average person. Little information is available about the air in homes and offices where the average citizens of the industrialized world spend much of their time. A move toward making low-cost air quality monitoring available can help bridge this gap.

The development of the WAQM system and its optical particle counter show that such low-cost systems can be constructed and that decent results can be produced from these systems. While sizing remains a challenge, and the amount of air flow reduces the total number of particles available for counting, results very close to much more expensive commercial OPCs can be produced. This allows for much more cost effective systems, and is conceivably the only financially feasible way to produce a multi-sensor system for most budget constrained programs. This type of sensor system can open the door to many applications where previously particulate matter measurements were infeasible.

The application benefits of this work are numerous, and many of the individual pieces involved will provide benefits for future programs. The novel solution for rejecting power line coupled noise removes this type of pollution from the data that

would otherwise plague a wall transformer powered system. The simulation system will provide a platform for evaluation future generations of optical particle counter systems. The dual-sensor prototypes will provide advances in particle detection and sizing for new WAQM system designs that are already being developed. The calibration and test system innovations, especially the cold and hot testing methodology, provide a low-cost method for calibrating detection thresholds without expensive particle generation materials.

Many challenges were encountered in the development of the WAQM system, and much future work remains to be done. The completion of the unfinished features of the event detection framework is at the top of the list. Work on the optical particle counter analog front end may yield results in sensitivity that have not yet been achievable with the current state of the design. Algorithmic work on particle detection and sizing may allow a new level of accuracy in matching more expensive commercial devices. The author looks forward to seeing the next generation of WAQM system developers carry on the work.

In field testing the WAQM system, several interesting environments have been explored. The work with a hospital in monitoring an operating room provided valuable insight to the medical team evaluating bacterial infection sources during operations. The experimentation in the aircraft cabin simulator demonstrated how the particulate matter sensor could be used to track substance movement in commercial airliners. Ad-hoc testing in the residences of the author and others involved in the program showed some interesting results relative to particulate matter filtering and carbon monoxide movement. Future experimentation is also in the works, with proposals for using the WAQM system and its progeny in air quality monitoring environments from working mines to outdoor

inversion event monitoring to long-term mounting aboard aircraft. The future holds a great deal of promise for using the work from this dissertation in many research and commercial applications.

REFERENCES

- [1] United States Environmental Protection Agency, “Integrated Science Assessment for Particulate Matter,” Washington, D.C., 2009.
- [2] G. B. Diette, M. C. McCormack, N. N. Hansel, P. N. Breyse, and E. C. Matsui, “Environmental Issues in Managing Asthma,” *Respir. Care*, vol. 53, no. 5, pp. 602–617, May 2008.
- [3] W. W. Nazaroff, “Indoor Particle Dynamics,” *Indoor Air*, vol. 14, no. 7, pp. 175–183, 2004.
- [4] C. A. Pope and D. Dockery, “Health Effects of Fine Particulate Air Pollution: Lines that Connect,” *J. Air Waste Manage. Assoc.*, vol. 56, pp. 709–742, 2006.
- [5] C. A. Pope, “Epidemiology of Fine Particulate Air Pollution and Human Health: Biologic Mechanisms and Who’s at Risk?,” *Environ. Health Perspect.*, vol. 108, no. 4, pp. 713–723, Aug. 2000.
- [6] R. D. Brook, B. Franklin, W. Cascio, Y. Hong, G. Howard, M. Lipsett, R. Luepker, M. Mittleman, J. Samet, S. C. Smith, and I. Tager, “Air Pollution and Cardiovascular Disease: A Statement for Healthcare Professionals from the Expert Panel on Population and Prevention Science of the American Heart Association,” *Circulation*, vol. 109, no. 21, pp. 2655–2671, Jun. 2004.
- [7] A. H. Mokdad, J. S. Marks, D. F. Stroup, and J. L. Gerberding, “Actual Causes of Death in the United States, 2000,” *JAMA*, vol. 291, no. 10, pp. 1238–1245, Mar. 2004.
- [8] California Air Resource Board, “Estimate of Premature Deaths Associated with Fine Particle Pollution (PM_{2.5}) in California Using a U.S. Environmental Protection Agency Methodology,” 2010.
- [9] World Health Organization, “WHO Air quality guidelines for particulate matter, ozone, nitrogen dioxide and sulfur dioxide. Online,” 2005.
- [10] L. Wallace, “Correlations of Personal Exposure to Particles with Outdoor Air Measurements: A Review of Recent Studies,” *Aerosol Sci. Technol.*, vol. 32, no. 1, pp. 15–25, 2000.

- [11] C. A. Pope, M. Ezzati, and D. W. Dockery, "Fine-Particulate Air Pollution and Life Expectancy in the United States," *N. Engl. J. Med.*, vol. 360, no. 4, pp. 376–86, Jan. 2009.
- [12] D. Liu, "Particle Deposition onto Enclosure Surfaces," in *Developments in Surface Contamination and Cleaning: Particle Deposition, Control and Removal*, R. Kohli and K. L. Mittal, Eds. Amsterdam: Elsevier, 2010.
- [13] D. Cooper, "Particulate Contamination and Microelectronics Manufacturing: An Introduction," *Aerosol Sci. Technol.*, vol. 5, no. 3, pp. 287–299, 1986.
- [14] W. W. Nazaroff, L. G. Salmon, and G. R. Cass, "Concentration and Fate of Airborne Particles in Museums," *Environ. Sci. Technol.*, vol. 24, no. 1, pp. 66–77, 1990.
- [15] S. Scaltriti, S. Cencetti, S. Rovesti, I. Marchesi, A. Bargellini, and P. Borella, "Risk Factors for Particulate and Microbial Contamination of Air in Operating Theatres," *J. Hosp. Infect.*, vol. 66, no. 4, pp. 320–326, Aug. 2007.
- [16] D. a Grantz, J. H. B. Garner, and D. W. Johnson, "Ecological effects of particulate matter.," *Environ. Int.*, vol. 29, no. 2–3, pp. 213–39, Jun. 2003.
- [17] American Conference of Governmental Industrial Hygienists, *2010 TLVs and BEIs: Based on the Documentation of the Threshold Limit Values for Chemical Substances and Physical Agents & Biological Exposure Indices*. Cincinnati, Ohio: ACGIH Signature Publications, 2010.
- [18] Occupational Safety and Health Administration, "Federal Register 29 CFR 1910, OSHA Safety and Health Standards, General Industry." Washington, D.C., 1976.
- [19] Department of Health and Human Services, Centers for Disease Control and Prevention, and National Institute for Occupational Safety and Health, "NIOSH Pocket Guide to Chemical Hazards." NIOSH Publications, Cincinnati, Ohio, 2007.
- [20] M. H. Sherman, "ASHRAE's First Residential Ventilation Standard," in *Proceedings of Buildings IX Conference*, 2004, pp. 1–13.
- [21] United States Environmental Protection Agency, "Buildings and their Impact on the Environment: A Statistical Summary," 2009.
- [22] A. Jones, "Indoor Air Quality and Health," *Atmos. Environ.*, vol. 33, no. 28, pp. 4535–4564, 1999.
- [23] T. Hussein, T. Glytsos, J. Ondracek, P. Dohanyosova, V. Zdimal, K. Hameri, M. Lazaridis, J. Smolik, and M. Kulmala, "Particle Size Characterization and

- Emission Rates During Indoor Activities in a House,” *Atmos. Environ.*, vol. 40, no. 23, pp. 4285–4307, Jul. 2006.
- [24] S. Kim and E. Paulos, “inAir: Measuring and Visualizing Indoor Air Quality,” in *Proceedings of the 11th International Conference on Ubiquitous Computing*, 2009, pp. 81–84.
- [25] A. R. Ferro, R. J. Kopperud, and L. M. Hildemann, “Elevated personal exposure to particulate matter from human activities in a residence.,” *J. Expo. Anal. Environ. Epidemiol.*, vol. 14 Suppl 1, pp. S34–40, Jan. 2004.
- [26] J. Richmond-Bryant, A. D. Eisner, L. a. Brixey, and R. W. Wiener, “Short-Term Dispersion of Indoor Aerosols: Can It Be Assumed the Room is Well Mixed?,” *Build. Environ.*, vol. 41, no. 2, pp. 156–163, Feb. 2006.
- [27] L. D. Sabin, E. Behrentz, A. M. Winer, S. Jeong, D. R. Fitz, D. V Pankratz, S. D. Colome, and S. a Fruin, “Characterizing the Range of Children’s Air Pollutant Exposure During School Bus Commutes,” *J. Expo. Anal. Environ. Epidemiol.*, vol. 15, no. 5, pp. 377–387, Sep. 2005.
- [28] J. Kiepert, S. M. Loo, D. Klein, and M. Pook, “Wireless Sensor Networks for Aircraft Cabin Environmental Sensing,” in *AIAA-2011-5108 41st International Conference on Environmental Systems*, 2011, pp. 1–8.
- [29] National Research Council (US). Committee on Air Quality in Passenger Cabins of Commercial Aircraft, *The Airliner Cabin Environment and the Health of Passengers and Crew*. National Academies Press, 2002.
- [30] S. Osman, M. T. La Duc, A. Dekas, D. Newcombe, and K. Venkateswaran, “Microbial Burden and Diversity of Commercial Airline Cabin Air During Short and Long Durations of Travel,” *ISME J.*, vol. 2, no. 5, pp. 482–497, May 2008.
- [31] P. La Roche and M. Milne, “Effects of Window Size and Thermal Mass on Building Comfort Using an Intelligent Ventilation Controller,” *Sol. Energy*, vol. 77, no. 4, pp. 421–434, Oct. 2004.
- [32] United States Environmental Protection Agency, “Federal Register Part II 40 CFR Part 50, National Ambient Air Quality Standards for Particulate Matter,” vol. 71, no. 200, pp. 61144–61233, 2006.
- [33] E. G. Patashnick, H.; Rupprecht, “Continuous PM-10 Measurements Using the Tapered Element Oscillating Microbalance,” *J. Air Waste Manag. Assoc.*, vol. 41, no. 8, pp. 1079–1083, 1991.
- [34] P. A. Baron, P. Kulkarni, and K. Willeke, *Aerosol Measurement: Principles, Techniques, and Applications*, 3rd ed. John Wiley & Sons, 2011.

- [35] C. O’Konski, “Means for Circulating a Particle Carrying Stream Through an Illumination Zone,” US 27327531956.
- [36] X. Wang, J. Agarwal, and G. Chancellor, “Size Segregated Aerosol Mass Concentration Measurement with Inlet Conditioners and Multiple Detectors,” 8047055 B22011.
- [37] S. A. Schleusener, “Laser Small-Particle Detector and Method,” US 3,406,2891968.
- [38] R. G. Knollenberg, “Particle Measurement Utilizing Orthogonally Polarized Components of a Laser Beam,” US 4,636,0751987.
- [39] Dylos Corporation, “DC1100 Air Quality Monitor.” [Online]. Available: <http://www.dylosproducts.com/ornodcairqum.html>. [Accessed: 21-Mar-2012].
- [40] Sharp Microelectronics, “GP2Y1010AU0F Compact Optical Dust Sensor Datasheet,” 2006. [Online]. Available: <http://www.sharpsme.com/optoelectronics/sensors/air-sensors/GP2Y1010AU0F>.
- [41] J. P. Black, A. Elium, R. M. White, M. G. Apte, L. A. Gundel, and R. Cambie, “MEMS-Enabled Miniaturized Particulate Matter Monitor Employing 1.6 GHz Aluminum Nitride Thin-Film Bulk Acoustic Wave Resonator (FBAR) and Thermophoretic Precipitator,” in *IEEE Ultrasonics Symposium*, 2007, vol. 4, no. c, pp. 476–479.
- [42] A. Hajjam, J. Wilson, and A. Rahafrouz, “Self-Sustained Micromechanical Resonant Particulate Microbalance/Counters,” in *IEEE 24th International Conference on Micro Electro Mechanical Systems (MEMS)*, 2011, pp. 629–632.
- [43] X. Wang and J. Agarwal, “Size Segregated Aerosol Mass Concentration Measurement Device,” US 20090039249A12008.
- [44] T. Nakajima, “Particle Counting Method,” US007733487B22010.
- [45] A. Trampe, T. Nierhoff, D. Semleit, and H. Fissan, “Digital Signal Processing for Optical Particle Counters in Order to Improve the Signal-to-Noise Ratio,” *Aerosol Sci. Technol.*, vol. 25, pp. 399–410, 1996.
- [46] J. Luo, A. Trampe, and H. Fissan, “A New Particle Counter Using Nonimaging Optics and Real- Time Correlation Filter for Particle Detection,” *Aerosol Sci. Technol.*, vol. 30, pp. 545–555, 1999.
- [47] J.-J. Kim, “Wireless Monitoring of Indoor Air Quality by a Sensor Network,” *Indoor Built Environ.*, vol. 19, no. 1, pp. 145–150, Mar. 2010.

- [48] S. Choi, N. Kim, H. Cha, and R. Ha, "Micro Sensor Node for Air Pollutant Monitoring: Hardware and Software Issues.," *Sensors (Basel)*, vol. 9, no. 10, pp. 7970–87, Jan. 2009.
- [49] B. Dong, B. Andrews, K. P. Lam, M. Höynck, R. Zhang, Y.-S. Chiou, and D. Benitez, "An information technology enabled sustainability test-bed (ITEST) for occupancy detection through an environmental sensing network," *Energy Build.*, vol. 42, no. 7, pp. 1038–1046, Jul. 2010.
- [50] G. P. Sharp, J. P. Schwartz, D. J. Buzzeo, E. M. Desrochers, and D. L. Farrington, "Air Quality Monitoring Systems and Methods," US007302313B22007.
- [51] A. Wood, G. Virone, and T. Doan, "ALARM-NET : Wireless Sensor Networks for Assisted-Living and Residential Monitoring," *Univ. Virginia ...*, 2006.
- [52] O. A. Postolache, J. M. D. Pereira, and P. M. B. S. Girão, "Smart sensors network for air quality monitoring applications," *Instrum. Meas. IEEE Trans.*, vol. 58, no. 9, pp. 3253–3262, 2009.
- [53] Atmel Corporation, "32-bit Microcontroller AT32UC3A3/A4 Datasheet." [Online]. Available: <http://www.atmel.com/Images/doc32072.pdf>. [Accessed: 21-Mar-2012].
- [54] J. Yick, B. Mukherjee, and D. Ghosal, "Wireless Sensor Network Survey," *Comput. Networks*, vol. 52, no. 12, pp. 2292–2330, Aug. 2008.
- [55] Raspberry Pi Foundation, "Raspberry Pi Website." [Online]. Available: <http://www.raspberrypi.org/>. [Accessed: 03-Jan-2014].
- [56] Texas Instruments, "TLV1117 Datasheet." [Online]. Available: <http://www.ti.com/product/tlv1117>.
- [57] OSRAM Opto Semiconductors GmbH, "BPW34 Datasheet." [Online]. Available: <http://catalog.osram-os.com/>.
- [58] J. Graeme, *Photodiode Amplifiers: Op Amp Solutions*. McGraw Hill, 1996.
- [59] P. Hobbs, *Building Electro-Optical Systems: Making It All Work*, 2nd ed. Hoboken, New Jersey: John Wiley & Sons, 2011.
- [60] S. Mak and D. Reed, "TWACSTM, A New Viable Two-Way Automatic Communication System for Distribution Networks. Part I: Outbound Communication," *Power Appar. Syst. IEEE ...*, no. 8, pp. 2941–2949, 1982.
- [61] Apinex, "Apinex Laser Module Datasheet." [Online]. Available: <http://www.apinex.com/ret2/03008B.html>.

- [62] Coherent Incorporated, “Coherent PowerMax-USB Datasheet.” [Online]. Available: <http://www.coherent.com/products/?1714/PowerMax-USB-Sensors>.
- [63] Newport Corporation, “Newport Laser Beam Profiler LPB HR Datasheet.” [Online]. Available: <http://www.newport.com/>.
- [64] G. Mie, “Contributions to the Optics of Turbid Media, Particularly of Colloidal Metal Solutions,” *Contrib. to Opt. Turbid media, Part. Colloid. Met. Solut. Transl. into ENGLISH from Ann. Phys.(Leipzig)*, vol. 25, no. 3, pp. 377–445, 1976.
- [65] A. Naqwi and F. Durst, “Focusing of Diode Laser Beams: A Simple Mathematical Model,” *Appl. Opt.*, vol. 29, no. 12, 1990.
- [66] C. F. Bohren and D. R. Huffman, *Absorption and scattering of light by small particles*. New York: John Wiley & Sons, 1983.
- [67] C. Mätzler, “MATLAB Functions for Mie Scattering and Absorption,” Bern, 2002.
- [68] W. J. Wiscombe, “Improved Mie scattering algorithms,” *Appl. Opt.*, vol. 19, no. 9, pp. 1505–1509, 1980.
- [69] E. Hecht, *Optics*. Reading, Massachusetts: Addison-Wesley, 1998.
- [70] A. Van Oosterom and J. Strackee, “The Solid Angle of a Plane Triangle,” *Biomed. Eng. IEEE ...*, no. 2, pp. 125–126, 1983.
- [71] X. Ma, J. Q. Lu, R. S. Brock, K. M. Jacobs, P. Yang, and X.-H. Hu, “Determination of Complex Refractive Index of Polystyrene Microspheres from 370 to 1610 nm,” *Phys. Med. Biol.*, vol. 48, no. 24, pp. 4165–72, Dec. 2003.
- [72] W. C. Hinds, *Aerosol Technology: Properties, Behavior, and Measurement of Airborne Particles*. New York: Wiley, 2012.
- [73] G. S. Kent, G. K. Yue, U. O. Farrukh, and a Deepak, “Modeling atmospheric aerosol backscatter at CO₂ laser wavelengths. 1: Aerosol properties, modeling techniques, and associated problems.,” *Appl. Opt.*, vol. 22, no. 11, p. 1655, Jun. 1983.
- [74] M. Ebert, S. Weinbruch, P. Hoffmann, and H. M. Ortner, “The chemical composition and complex refractive index of rural and urban influenced aerosols determined by individual particle analysis,” *Atmos. Environ.*, vol. 38, no. 38, pp. 6531–6545, Dec. 2004.
- [75] OSRAM Opto Semiconductors GmbH, “BPW34 Silicon PIN Photodiode.” [Online]. Available: <http://catalog.osram-os.com>. [Accessed: 18-Oct-2012].

- [76] TSI Incorporated, "TSI Aerotrak 9303 Datasheet." [Online]. Available: <http://www.tsi.com/aerotrak-handheld-particle-counter-9303/>.
- [77] Particle Measuring Systems, "PMS Lasair II 110 Datasheet." [Online]. Available: <http://www.pmeasuring.com/particleCounter/air/mobile/L2110>.
- [78] R. L. Unger, "Compact, Low Cost Particle Sensor," US008009290B22011.
- [79] K. Whitey, "The physical characteristics of sulfur aerosols," *Atmos. Environ.*, vol. 41, pp. 25–49, 2007.
- [80] International Standards Organisation., *ISO 21501-4. determination of particle size distribution : single particle light interaction methods : light scattering airborne particle counter for clean spaces. part 4 part 4*. Geneva, Switzerland: International Standards Organisation, 2007.
- [81] V. H. Grassian, P. T. O'shaughnessy, A. Adamcakova-Dodd, J. M. Pettibone, and P. S. Thorne, "Inhalation Exposure Study of Titanium Dioxide Nanoparticles with a Primary Particle Size of 2 to 5 nm.," *Environ. Health Perspect.*, vol. 115, no. 3, pp. 397–402, Mar. 2007.
- [82] BGI Incorporated, "BGI Collison Nebulizer Datasheet." [Online]. Available: <http://www.bgiusa.com/agc/collison.htm>.
- [83] TSI Incorporated, "TSI 3062 Diffusion Dryer Datasheet." [Online]. Available: http://www.tsi.com/Products/_Product-Accessories/Diffusion-Dryer-3062.aspx.
- [84] S. V Hering and M. R. Stolzenburg, "A Method for Particle Size Amplification by Water Condensation in a Laminar , Thermally Diffusive Flow," *Aerosol Sci. Technol.*, vol. 39, pp. 428–436, 2005.
- [85] Thermo Scientific, "Thermo Scientific 3000/4000 Series Particle Counter Standards." [Online]. Available: <http://www.thermoscientific.com/en/product/3000-4000-series-particle-counter-standards-1.html>. [Accessed: 02-Jan-2014].
- [86] R. O. Darouiche, "Treatment of Infections Associated with Surgical Implants," *N. Engl. J. Med.*, vol. 350, no. 14, pp. 1422–9, Apr. 2004.
- [87] J. Whitehouse and N. Friedman, "The Impact of Surgical-Site Infections Following Orthopedic Surgery at a Community Hospital and a University Hospital: Adverse Quality of Life, Excess Length of Stay, and Extra Cost," *Infect. Control Hosp. Epidemiol.*, vol. 23, no. 4, pp. 183–189, 2002.
- [88] "Raspberry Pi Home Page." [Online]. Available: <http://www.raspberrypi.org/>.

- [89] M. Pereira, R. Vilain, T. Leivas, and A. Tribess, “Measurement of the Concentration and Size of Aerosol Particles and Identification of the Sources in Orthopedic Surgeries,” *HVAC&R Res.*, no. February 2013, pp. 37–41, 2012.
- [90] W. Tsai and R. Pogorzelski, “Eigenfunction Solution of the Scattering of Beam Radiation Fields by Spherical Objects,” *JOSA*, vol. 65, no. 12, pp. 1457–1463, 1975.
- [91] S. Borrmann, B. Luo, and M. Mishchenko, “Application of the T-Matrix Method To the Measurement of Aspherical (Ellipsoidal) Particles With Forward Scattering Optical Particle Counters,” *J. Aerosol Sci.*, vol. 31, no. 7, pp. 789–799, Jul. 2000.
- [92] E. Purcell and C. Pennypacker, “Scattering and Absorption of Light By Nonspherical Dielectric Grains,” *Astrophys.*, 1973.