4-15-2019

Simulations of Ocean Circulation Under an Ice-Shelf: Problem Feasibility Study Using Non-Hydrostatic Unified Model of the Ocean (NUMO)

Chithkala Dhulipati
Boise State University

Michal Kopera
Boise State University

Simulations of Ocean Circulation Under an Ice-Shelf: Problem Feasibility Study Using Non-Hydrostatic Unified Model of the Ocean (NUMO)

Abstract

Antarctic ice-cover extends from the ice-cap on the continent (ice-sheet) into the surrounding ocean forming floating ice-shelves. The interaction of ice-shelves and ocean underneath (in the so-called ice-shelf cavity) controls the rate at which the ice is exported from the ice-sheet to the ocean, which directly contributes to sea-level rise. The ocean circulation underneath Antarctic ice-shelves is poorly understood and is one of the focal points of worldwide climate modeling efforts.

In this work, we are conducting a feasibility study of deploying a new ocean circulation model NUMO to an ice-shelf cavity circulation problem. We are testing the performance of the model on a simplified cavity geometry, with a particular focus to the area close to the grounding line. The feasibility study involves measuring the time it takes the model to produce a 2-week simulation given the spatial resolution used, the order of the numerical scheme, and the number of processors of Boise State's R2 cluster used for the simulation. The results of this study will be used to improve the performance of the NUMO model.

Simulations of ocean circulation under an ice-shelf

Chithkala Dhulipati, Dr. Michal Kopera, Department of Mathematics, Boise State University

BOISE STATE UNIVERSITY

Problem feasibility study using Non-hydrostatic Unified Model of the Ocean (NUMO)

INTRODUCTION

Antarctic ice-cover extends from the ice-cap on the continent into the surrounding ocean forming floating ice-shelves. The interaction of ice-shelves and ocean in the ice-shelf cavity controls the rate at which the ice is moved from the ice-sheet to the ocean. This directly contributes to sea-level rise. While of utmost importance for climate science, this phenomenon is not very well represented in modern climate ocean models.

AIM

In this work, we conducted a feasibility study of deploying a new ocean circulation model NUMO to an ice-shelf cavity circulation problem. We tested the performance of this model on a simplified two-dimensional cavity geometry. The study involves measuring the time it takes the model to produce a two-week simulation given the spatial resolution used, the order of the numerical scheme, and the number of processors used for the simulation. These results will ultimately be used to improve the performance of the NUMO model.

SIMULATION SET-UP

The simulation used 8192 elements to represent the sub-ice shelf cavity domain. We have used 4th order method, which created 25 points within each 2D element, for a total number of 204800 points. This resulted in a resolution varying from 0.2 m near the grounding line (left) and 3 m near the open ocean (right). To run the model on a parallel computer, we have used up to 140 cores on Boise State's R2 High Performance Cluster. The figure below shows the salinity field on a small fraction of the domain closest to the grounding line with. One square in the background represents 1 m x 1 m area. Total cavity length is 10 km, and the depth at the ocean interface is 100 m.

SCALABILITY ANALYSIS

We measured the wall clock time (the time taken for the model to produce a simulation of 60 s of flow evolution) of each model run for a variable number of processors on R2, up to 140 cores. The figure below shows the wall clock time as a function of processor count plotted against ideal time (wall time on 1 processor divided by the number of processors). As the number of processors increases, the wall time decreases at almost ideal rate, with some indication of flattening out near the upper limit of processor count.
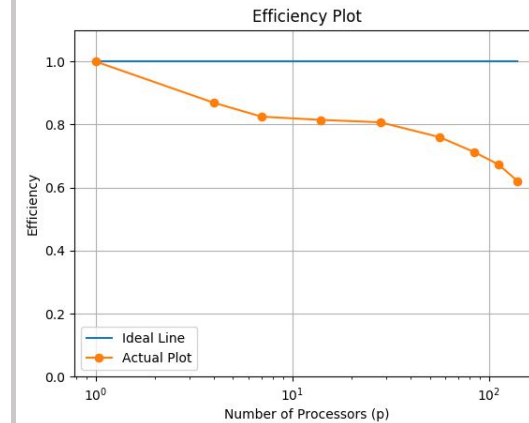
SPEED UP PLOT

$$S_p = \frac{t(1)}{t(p)}$$

The timing result above can be presented in terms of simulation speed-up (S_p), which is the ratio of wall time on one processor ($t(1)$) to the wall time on p processors ($t(p)$) We see that we achieve speed-up close to the ideal line for most of the range of the number of processor used.
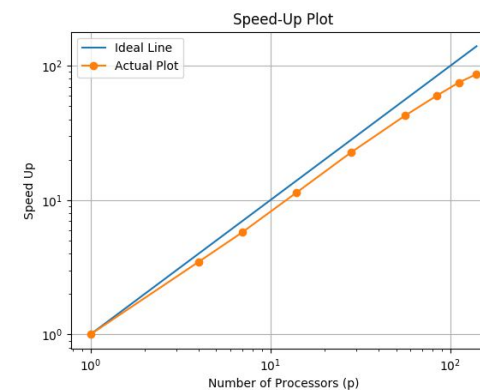
EFFICIENCY ANALYSIS

The efficiency plot allows us to see how effectively we use increasing number of processors for given problem. The efficiency is defined as a ratio of speed-up to the processor number. We see that for this problem we achieve efficiency >70% for almost the entire range of p.

CONCLUSION

We found that the simulation runs most efficiently on 140 processors, but the speed up and efficiency plots suggest that using more cores will yield improvement in time-to-solution. With this configuration, we are able to simulate 0.24 model days in one wall-clock day, that is, it takes the model 1 day of computer time to produce 6 hours of simulation.

In the future, we will test the time-to-solution on more R2 cores and increase the simulation resolution, to see how the speed-up and efficiency plots change for the new configuration. Also, we will look for ways to decrease the time-to-solution