Ocean and Sea Ice Modeling in the High-Mountain Desert of New Mexico: A Brief History of COSIM

This article highlights the people, events, and activities that led to the establishment of the Climate, Ocean, and Sea-Ice Modeling (COSIM) Project at Los Alamos National Laboratory (LANL). It does not cover the entire history of the project.

Before COSIM: Building a relationship with NCAR

The seed for COSIM was planted in 1979, when Bob Malone (the author of this article) accepted a position in the Theoretical Division, charged with building a climate modeling capability at LANL. To learn about the state of climate modeling, Malone arranged to spend the summer of 1979 at the National Center for Atmospheric Research (NCAR) in Boulder, Colorado. There he became involved in a very small project to develop a global atmospheric model based on spectral-transform methods. The project was led by Eric Pitcher and V. Ramanathan; Pitcher was a long-term visitor at NCAR from the University of Miami, and Ramanathan was a scientist at NCAR. Ramanathan had developed a new cloud and radiative-transfer scheme that was being installed in the model framework. Malone returned to NCAR for several months during each of the next four summers to work with Pitcher and others on documenting, validating, and applying the model. This model was adopted by NCAR in 1981 as the first version of the NCAR Community Climate Model (CCM0).

This visit was also very helpful in establishing working relationships with Maurice Blackmon, Warren Washington, Bert Semtner, Bob Chervin, and Steve Schneider. In the future, each of them (especially Warren Washington) would influence in some way the direction of climate modeling at Los Alamos. Thus began a long-term collaboration between NCAR and Los Alamos that came to fruition under COSIM.

Prospects for growth

From the beginning, Malone was supported by the Laboratory's Institutional Supporting Research and Development (ISRD) program. Prospects for increased funding from ISRD were nil, so if the program were to grow beyond one person, it would be necessary to obtain external funding, hopefully from DOE. In 1981, Malone submitted a proposal to the DOE, requesting $50K (which in those days was half the cost of a Technical Staff Member or the full cost of a TEC programmer) for work on CCM0. It was rejected, and he was informed that DOE policy was not to support work on global atmospheric models at national laboratories, but to limit such funding to institutions like NCAR. In order to expand beyond the 1-FTE level, Malone took a half-time position as assistant to the Associate Director for Physics and Mathematics to free up half of his ISRD funds; this allowed him to hire a full-time programmer to work with him.

Nuclear Winter

The first real opportunity for growth came in 1983, when the "Nuclear Winter" hypothesis was elevated by Carl Sagan and colleagues to a national scientific issue. Projects were begun immediately at NCAR and Lawrence Livermore National Laboratory (LLNL) to model the global effects of massive injections of smoke into the Earth's atmosphere using multi-dimensional atmospheric models. This was important because the seminal publication on Nuclear Winter by Turco, Toon, Ackerman, Pollack, and Sagan (TTAPS) was based entirely on a physically detailed one-dimensional (vertical) model. Consequently, the TTAPS model could not respond to inhomogeneities in the smoke distribution or at the surface boundary (e.g., ocean versus land). LLNL used a two-dimensional model (ZAM2) and NCAR used a modified version of the three-dimensional CCM0. At a federally sponsored symposium on Nuclear Winter in October 1983, both groups presented preliminary results on the radiative effects of imposed, static smoke distributions on the land surface temperature.

In August 1983, Malone was put in charge of organizing a Nuclear Winter modeling project at Los Alamos. Steve Schneider, who was leading the Nuclear Winter modeling effort at NCAR, encouraged him to undertake the project. In October 1983, an additional FTE of ISRD funding was provided. This allowed Malone to return to full-time research and to bring Larry Auer into the project half time. Auer modified the radiative-transfer scheme in CCM0 to include the absorption of sunlight by smoke and later added refinements involving dust. In November, the Defense Nuclear Agency (which had been designated as the lead federal agency for Nuclear Winter research) committed to provide an additional FTE of funding. It arrived in July 1984, and Gary Glatzmaier was hired immediately. As a graduate student and postdoc, Glatzmaier had developed a 3D model of the solar dynamo using spectral-transform methods, so he was able to quickly become familiar with CCM0.

With Glatzmaier on board, the project really took off. By February 1985, the Los Alamos model had a comprehensive set of physical processes including smoke injection as a function of location, altitude, and time; horizontal and vertical transport and mixing of smoke; interactions of smoke with sunlight; the resultant heating of the atmosphere; and the heat-induced modification of the atmospheric structure, dynamics, and precipitation. The Los Alamos model had moved incrementally ahead of the NCAR model by including a fully interactive scheme for precipitation scavenging of smoke, which was validated against measured lifetimes of passive aerosols. This work was well received at a National Academy of Science symposium on Nuclear Winter in March 1985. The Los Alamos team received a LANL Distinguished Performance Award in 1985, and was especially proud when Warren Washington chose the Los Alamos modeling work to include in his 1986 book on climate modeling[†].

CHAMMP: A new opportunity

This success brought recognition and credibility to the Los Alamos project and laid a foundation that was important in creating a role for Los Alamos at the beginning of the Department of Energy's Computer Hardware, Advanced Mathematics, and Model Physics (CHAMMP) program in 1990. CHAMMP was a new computational initiative to bring massively parallel computing to climate modeling. The establishment of CHAMMP opened the door for Los Alamos to build on its reputation and to obtain stable, long-term funding from DOE for the first time. David Bader from Pacific Northwest National Laboratory was appointed by Ari Patrinos to run the program. Bader then asked Malone to coordinate the Model Development component of CHAMMP, while Bader ran the

[†] An Introduction to Three-Dimensional Climate Modeling, Warren M. Washington and Claire L. Parkinson, University Science Books, 1986, pp. 283-288.

Science Team component. The latter encompassed all university and some DOE laboratory research grants. The former was more programmatic, designed to actually carry out revisions of existing component (atmosphere, ocean, sea ice, land surface) models at our "client" institutions, NCAR and GFDL.

Why do ocean modeling at Los Alamos?

Malone's experience working on Nuclear Winter had convinced him that it would be unprofitable to pursue work at Los Alamos on atmospheric modeling. Not only was NCAR a dominant force, but also there were on the order of 20 atmospheric general circulation models in the international community. The situation was quite different in ocean modeling. There were only a handful of ocean GCMs, and most of them were derivatives of the original Cartesian-coordinate ocean model developed by Kirk Bryan in the late 1960's at the NOAA Geophysical Fluid Dynamics Laboratory (GFDL) in Princeton, NJ. Thus, a small project had a better chance to have a significant impact in ocean modeling.

There is an important difference between atmospheric and ocean modeling. Atmospheric models require a large set of people having the collective expertise to deal with resolved-scale fluid dynamics and parameterizations of subgrid-scale processes such as fluid turbulence, cloud formation and precipitation, cloud optical properties, radiative transfer (solar and infrared), aerosols, convective mixing and the planetary boundary layer, and interactions of the atmosphere with land and ocean surfaces. Water exists in vapor, liquid and ice phases in the atmosphere, and interacts with aerosols, which act as cloud condensation nuclei. Because there are many interactions that take place among these parameterized processes, it is important that the experts work closely to ensure consistent treatments. This makes it desirable (but not essential) to have these experts in-house, which is clearly only possible at a large facility dedicated to climate modeling.

By comparison, ocean modeling is simple: the basic ingredient is fluid dynamics. The ice phase does occur, but only at the surface, so it can be treated as a separate entity: sea ice. The main subgrid-scale process in the ocean that must be parameterized is turbulence. Los Alamos has tremendous resources in fluid dynamics and turbulence to apply to ocean modeling, and a small team can accomplish this. It is noteworthy that the first six people to join the ocean modeling project were members of the Fluid Dynamics group T-3.

Moreover, massively parallel computers seemed to be particularly well suited to global ocean models. Mesoscale eddies in the ocean are about ten times smaller (20-200 km) than analogous eddies in the atmosphere. This means that ten times as many grid-points are required in both horizontal dimensions and, because the time step must be reduced by a factor of ten also, a factor of a thousand more computation is required. These requirements can be well met with massively parallel computers.

These considerations led Malone to decide that Los Alamos would focus on ocean model development. He then approached Semtner and Chervin about obtaining a copy of their ocean GCM as a starting point, which they generously provided.

POP: The foundation for COSIM

Rick Smith and John Dukowicz soon joined the effort to develop a massively parallel ocean model for the Connection Machine, which culminated in the Parallel Ocean

Program (POP). With PhDs in nuclear physics and aerophysics, respectively, they were immune to the “this is the way we’ve always done it” mindset. Smith rewrote the entire code in Connection Machine Fortran, an early variant of Fortran90. They completely reformulated the barotropic mode in the model and replaced algorithms with new ones that would parallelize efficiently. This involved several important and revolutionary changes. First, the streamfunction formulation of the barotropic mode was replaced by a surface pressure formulation. The most important effect of this change was that the boundary conditions became local, making it possible to parallelize the code with good data locality (essential for distributed memory computers) for the first time. This made it possible to tackle very large, highly resolved problems. In addition, the barotropic operator became better behaved near coastlines and near steep topographic gradients, thus improving accuracy. Second, the rigid-lid surface boundary condition approximation was eliminated in favor of a physically realistic variable surface displacement, and the resulting parabolic equation for the surface elevation was discretized implicitly. This, together with the introduction of a preconditioned conjugate-gradient solver, greatly improved the computational efficiency of the barotropic mode and of the code as a whole. What was remarkable was that these changes simultaneously improved physical realism as well as computational efficiency. This new code was dubbed the “Parallel Ocean Program” (POP) by Smith during a visit to GFDL, as a play on GFDL’s “Modular Ocean Model” (MOM). In 1992, the POP team received a Los Alamos Distinguished Performance Award; and in 1994, the team won the Computerworld-Smithsonian Award in the Science category. The first public release of POP took place in 1994.

These changes made it possible to push on to higher-resolution simulations. In collaboration with Bert Semtner, POP was used to extend his pioneering work with Bob Chervin on high-resolution ocean modeling to still higher resolutions on the Connection Machine at Los Alamos. As computers became more powerful, COSIM team members and collaborators refined the grid to 0.1° (~10 km), first in the North Atlantic Ocean and later in the full global ocean.

SKYHI: Attempting a partnership with GFDL

The CHAMMP program was conceived as having two main “clients”, NCAR and GFDL. The idea was to select models developed at these institutions and to establish collaborative teams to carry out the parallelization. Discussions between GFDL and CHAMMP management led to the decision to parallelize the GFDL SKYHI atmospheric model to run on the CM-200 at LANL. SKYHI had been developed by Jerry Mahlman, who by this time had risen to the position of GFDL Director. This required a new hire at LANL with expertise in CM Fortran and large-code development. Phil Jones had come to Los Alamos as a postdoc to work on parallelizing Glatzmaier’s dynamo code to run on LANL’s new CM-200. Fortunately, Jones was looking for a staff member position at just the right time. Jones did an outstanding job on SKYHI, in collaboration with several GFDL scientists. The team published one paper on the parallelization process and two papers on scientific applications of SKYHI. For whatever reasons, this technological and scientific success failed to arouse sufficient enthusiasm within GFDL for parallel computing to draw GFDL into a long-term collaboration with CHAMMP comparable to the collaboration with NCAR. After SKYHI, Jones became the main developer of POP.

At last, a real oceanographer!

In 1993, Mat Maltrud became the first member of the team to hold a PhD in oceanography. He had grown up in Los Alamos and had been an undergraduate student employee in T-3. He moved back to take a postdoctoral position in Group T-3 working with Frank Harlow on turbulence modeling. Rick Smith and John Dukowicz were already working on ocean modeling in T-3. Naturally, Maltrud became interested in their work, and eventually joined COSIM. As POP progressed, it became clear that a good sea-ice model would be a valuable companion to POP. Maltrud had begun to work with Semtner's 1976 thermodynamics-only sea-ice model, when in 1994, Elizabeth Hunke, our second postdoc, joined COSIM to work on sea-ice modeling. This allowed Maltrud to move to the POP team to work on model verification and validation. He and Rick Smith were the first to use POP at 0.1°.

CICE: A radically different sea ice model

Hunke teamed up with Dukowicz, and they repeated the POP experience by thoroughly rethinking traditional approaches to sea-ice modeling. They developed a completely new Elastic-Viscous-Plastic (EVP) rheology and dynamics scheme that addressed the most pressing shortcomings of sea-ice models. Traditionally, sea-ice stresses are modeled by a visco-plastic rheology that resembles that of a highly nonlinear viscous fluid whose effective viscosity becomes singular at low strain rates. To make such a model tractable, the practice was to hold the effective viscosity constant at each time step and to regularize the singularity by bounding the viscosity at some very high value. This still made the model difficult and expensive to solve and the viscous linearization introduced a nonphysical response lag to rapidly changing forcing. Physically, the singularity at very low strain rates does not exist because sea ice responds elastically in this regime, in which it resembles a rigid body. Hunke and Dukowicz introduced a simplified elastic regularization that mimicked the physical response but was primarily intended to improve the numerical behavior. Instead of solving a difficult implicit problem by iteration, the solution was obtained explicitly by sub-cycling with the effective viscosity being allowed to vary on the much shorter sub-cycling time scale. This not only improved the efficiency and parallelizability but also made the model much more responsive to rapid forcing changes. At first, the entrenched traditionalists refused to accept this new approach but repeated comparisons of EVP with standard models demonstrated that EVP was better both physically and computationally. Now EVP has been adopted by almost every major climate-modeling center in the world. Hunke went on to create a full sea-ice model, CICE, by implementing and parallelizing Semtner's 1976 sea ice thermodynamics package. The first public release of EVP took place in 1997 and the full CICE model followed in 1998.

Meeting CHAMMP's objectives

Two early objectives were established for CHAMMP:

1. To have, by 1992, publishable results having high scientific impact, obtained with a new model run on a new massively parallel computer.

2. To develop components for a new state-of-the-science atmosphere-ocean GCM for climate modeling that can be run on a variety of architectures, particularly distributed-memory massively parallel systems.

The ocean-modeling project at Los Alamos satisfied both objectives. The first was met by a paper by Smith, Dukowicz, and Malone, published in *Physica D* in 1992, entitled "Parallel Ocean General Circulation Modeling", in which the breakthrough methods described above were presented with results on the CM-200 at Los Alamos. The second objective was met in 1994 at a small CHAMMP meeting in Santa Fe. At that meeting, Warren Washington discussed his desire to create a new Parallel Climate Model (PCM), based on the model development work accomplished under CHAMMP. For the ocean component, he chose POP. For the atmospheric component of PCM, he adopted a new version of the Community Climate (atmospheric) Model that had been parallelized by Argonne and Oak Ridge National Laboratories under the CHAMMP Model Development Program. At the time of this meeting, Hunke had been at LANL for only two months, so EVP and CICE were far in the future. The sea ice model in PCM was based initially on the 1976 Semtner scheme. However, when the Hunke-Dukowicz EVP dynamics had been demonstrated to be superior to older models, it was quickly adopted in PCM.

Partly due to Washington's success with POP, the ocean-modeling group at NCAR decided to replace the first-generation NCAR ocean model, a modified version of the GFDL Modular Ocean Model (MOM), with POP. Later, the NCAR-based Community Climate System Model (CCSM) project also adopted POP and a sea-ice model based on CICE. At this point, NCAR had two climate models that by different paths had come to have the same component models: a parallelized CCM2, POP, a sea ice model almost identical to CICE, and the community land-surface model. PCM was soon phased out and all effort devoted to CCSM. CCSM is available to anyone worldwide over the web, and it serves the university community in the US. This (probably) makes it the most widely used climate model in the world. The members of COSIM are proud of their contributions to CCSM and look forward to continuing this very successful partnership.

The formation of COSIM

Bader was very supportive of the Los Alamos ocean and sea-ice modeling effort. Due to the success of POP and CICE, Bader encouraged Malone to form a center for climate modeling at Los Alamos, but this proved to be too ambitious. Therefore, the Climate, Ocean, and Sea Ice Modeling (COSIM) Project was formed instead in 1995. The name and acronym were carefully chosen, nonetheless, so that the acronym could also stand for *Center for Ocean and Sea-Ice Modeling*, if the opportunity were to arise to form a center!

Thus, COSIM became an umbrella that encompassed a limited set of activities focused on model development, integration of components, and model applications, all on massively parallel computers. This was COSIM's mandate from DOE, and because the financial resources were constrained, it was important to maintain that rather narrow focus.

Because COSIM is not a formal organization, members of COSIM belong to one of several "host groups". In recent years, there have been three host groups:

- T-3, the Fluid Dynamics Group in the Theoretical Division;

- CCS-2, the Continuum Dynamics Group in the Computer and Computational Sciences Division;
- EES-2, the Atmospheric and Climate Dynamics Group in the Earth and Environmental Sciences Division.

Each member of COSIM was encouraged to consider which of these groups would be most compatible with his/her intrinsic interests. People with a physical science background in meteorology or oceanography might choose EES-2. Those interested in computational fluid dynamics might choose either T-3 or CCS-2. This organizational diversity and freedom of choice has worked well and continues to the present day. COSIM members perceive the advantages of being in host groups where their talents are marketable and fellow group members "speak the same language".

The growth of COSIM

The preceding discussion focuses solely on the people and activities that led to the formation of COSIM. Many people have joined COSIM since it was founded more than a decade ago, and a number of old-timers have retired. The roster of past and present members of COSIM appears in the table below.

Joined	**Name**	**Group**	**Interests**	**Status**
1990	Bob Malone	CCS-2	Project manager, POP	2004(r)
1991	Rick Smith	T-3	POP, turbulence	2005(d)
1991	John Dukowicz	T-3	CFD, model development	2001(r*)
1992	John Baumgardner	T-3	Geodesic grids & methods	2004(r)
1992	Phil Jones	T-3	POP, now project manager	TSM
1993	Mat Maltrud	T-3	Ocean modeling	PD/TSM
1994	Elizabeth Hunke	T-3	Sea ice modeling, CICE	PD/TSM
1995	*COSIM founded*			
1995	Sumner Dean	CCS-2	MICOM/HYCOM	TSM
1995	Balu Nadiga	CCS-2	Ocean modeling	PD/TSM
1996	Matthew Hecht	CCS-2	Ocean modeling	PD/TSM
1998	Bill Lipscomb	T-3	CICE, glacial modeling	PD/TSM
1999	Scott Elliott	CCS-2	Ocean biogeochemistry	TSM
1999	Rainer Bleck	EES-2	MICOM/HYCOM	2005(r)
2000	Synte Peacock	T-3	Application and analysis	PD/2003(t)†
2001	Shao-ping Chu	EES-2	Ocean biogeochemistry	PD/TSM
2001	Beth Wingate	CCS-2	Applied Math/GFD	TSM
2001	John Davis	CCN-12	Software engineering	2004(r)
2002	JoAnne Lysne	CCS-2	Oceanography, analysis	PD/2005(t)
2005	Mark Petersen	CCS-2	Alpha model in POP	PD/TSM
2006	Todd Ringler	T-3	HYPOP, alternate grids	TSM
2006	Wilbert Weijer	CCS-2	HYPOP, implicit modeling	TSM

The starting point for "COSIM membership" is taken to be coincident with the start of CHAMMP: 1990. Only the current (or last) group affiliation is listed. PD in the Status column indicates that the person joined COSIM as a postdoc; a TSM indicates that the person was/is a Technical Staff Member. A year in the Status column indicates that the person left the lab for the following reasons: (r) retirement, (d) disability leave, (t) terminated employment.

*Dukowicz has continued to work on a 40% annualized basis as a retired Associate Fellow.

†Peacock accepted a faculty position at the University of Chicago.

At its peak, COSIM had a dozen full-time scientists and one or two post-docs of various backgrounds working on POP, CICE, and Rainer Bleck's MICOM and HYCOM ocean models. COSIM is now a few people smaller but still very active. Its funding primarily comes from DOE's Climate Change Prediction Program (CCPP) and Scientific Discovery through Advanced Computing (SciDAC) program with lesser and variable amounts from NASA and the Navy.

The intentional focus of this article on COSIM should not be interpreted as implying that no other atmospheric or climate-related research has gone on at the laboratory. For many years before COSIM, the laboratory was a major partner in a global-scale atmospheric sampling program associated with detection of nuclear-weapons tests. The measurements made improved our understanding of atmospheric turbulence and transport properties. Dating back to the early 1970s, the laboratory has had a strong program in regional atmospheric modeling. In recent years, climate-related research at Los Alamos has grown far beyond COSIM. Unfortunately, it is not possible to delve into these topics herein.

Special recognition

Many people not in COSIM have made valuable contributions to COSIM over the years. Four stand out.

Early in CHAMMP, Len Margolin recognized the value of isopycnic coordinates in ocean models and worked to bring such a model into the Los Alamos project. There were two such models in existence at the time: one developed in Germany by Josef Oberhuber, called OPYC; and one developed by Rainer Bleck at the University of Miami, called the Miami Isopycnic Coordinate Ocean Model (MICOM). Len first approached Oberhuber about collaborating on transforming OPYC into a parallel code to run on distributed memory parallel computers. However, it was soon realized that the fundamental structure of the model did not lend itself easily to parallelization. Next Len approached Rainer Bleck, who was a CHAMMP Science-Team grantee, and invited him to visit Los Alamos. By the mid-1990's, Bleck had become a regular summer visitor at LANL, where he developed a long-lasting partnership with Sumner Dean. Eventually, it was possible to bring Rainer onboard full time, a rare opportunity in light of Rainer's stature in the international ocean modeling-community.

Darryl Holm has been closely associated with COSIM for many years. One of Darryl's long-term objectives has been to develop a turbulence model that has been validated over a wide range of parameter values, suitable for use in an ocean model. Beth Wingate and Mark Petersen have implemented and tested Darryl's alpha-model for turbulence in POP. In comparisons of high-resolution (0.2°) POP-only simulations with coarser (0.4°) POP plus alpha-model simulations, the alpha-model generates a distribution and amplitude of turbulent eddy kinetic energy remarkably close to the higher resolution. The lower-resolution runs require a factor of 8-10 less computer time.

Chick Keller has been exceptionally supportive of the climate modeling effort since its inception. As group leader of the Geoanalysis Group (ESS-5) of the Earth and Space Sciences Division, he took particular interest in the Nuclear Winter research that was being performed in his group. Later as the Director of the Institute for Geophysics and Planetary Physics (IGPP), he was supportive of the COSIM project.

While he was Director of the LANL Advanced Computing Laboratory, Andy White took a great interest in the CHAMMP program. He recognized the climate-modeling problem as a grand challenge problem that would tax the limits of the largest supercomputer. He worked with Bader to provide CHAMMP scientists access to the big CM-200 in the ACL. In later acquisitions, CHAMMP would typically be a 10-20% co-owner.

Accomplishments and long-range plans

The numerous accomplishments during the first decade of COSIM are described in a Los Alamos National Laboratory report entitled “A Decade of Progress in the Los Alamos Climate, Ocean and Sea Ice Modeling (COSIM) Project” (LA-UR-01-1886, May 2001). A five-year plan for COSIM was presented in a companion document “Long-range Plan for COSIM” (LA-UR-01-1887, May 2001). A major program review was held in January 2002, and another such review is planned for August 2007.