

REAL-TIME OCEAN MODELING SYSTEMS

The Naval Research Laboratory has developed the world's first eddy-resolving global ocean nowcast and forecast system. It uses satellite observations and an ocean model running on a high-performance-computing platform to enhance real-time knowledge of the marine environment in which naval submarines and ships must operate.

The first operational weather prediction occurred in May 1955 as a joint United States Air Force, Navy, and Weather Bureau project. In principle, numerical ocean modeling is similar to atmospheric modeling, but global operational oceanography has lagged far behind (see the “Atmospheric versus Oceanic Prediction” sidebar). This has been a concern for the US Naval Research Laboratory, because, as a nation protected from adversaries and linked to partners by the world's great oceans, it is fundamental that the US understand its surrounding marine environment. Consequently, for the past decade, the NRL has been working on the problem of eddy-resolving global ocean modeling and prediction. Furthermore, it has developed the world's first global ocean nowcast and forecast system using the Department of Defense's High Performance Computing Modernization Program (HPCMP) com-

puting resources. It has been running in real time at the Naval Oceanographic Office (NAVO) since October 2000. Here, we describe the computational requirements of numerical ocean modeling and how the NRL system operates.

Computational requirements

As far back as 1989, the President's Office of Science and Technology recognized global ocean modeling and prediction as a Grand Challenge problem, defined as requiring a computer system capable of sustaining at least one trillion floating-point adds or multiplies per second. We are solving the problem on today's systems—capable of only a fraction of this performance—by taking a multifaceted approach to cost minimization.

One facet is using the NRL Layered Ocean Model (NLOM),¹ specifically designed for eddy-resolving global ocean prediction. It is tens of times faster than other ocean models in computer time per model year for a given horizontal resolution and model domain. NLOM's performance is due to a range of design decisions, the most important of which is the use of isopycnal (density-tracking) layers in vertical—rather than fixed-depth—cells.

Density is the natural vertical coordinate system for the stratified ocean, and it lets seven NLOM layers replace the 50 or more fixed levels that would be needed at 1/16-degree (or 7

1521-9615/02/\$17.00 © 2002 IEEE

ALAN J. WALLCRAFT, HARLEY E. HURLBURT, E. JOSEPH METZGER, ROBERT C. RHODES, AND JAY F. SHRIVER

Naval Research Laboratory

OLE MARTIN SMEDSTAD

Planning Systems Inc.

km mid-latitude) resolution. NLOM's semi-implicit time scheme allows a longer time step by making it independent of all gravity waves. However, it requires solving a 2D Helmholtz's equation for each gravity mode at each time step. We can solve internal modes with 5 to 10 red-black successive over-relaxation (SOR) sweeps, but efficient solution of the single external gravity mode requires direct solution using the Capacitance Matrix Technique. CMT involves solving a dense system of linear equations across all coastline points. This is a huge matrix for global regions ($90,000 \times 90,000$ elements at $1/32$ -degree resolution). However, it does not change with time, so we can invert it once at the start of the simulation, leaving a simple matrix-vector product to be performed at each time step.

The CMT is an example of trading more memory for less CPU time. Because memory is plentiful on distributed memory machines, this tradeoff is now viable even for high-resolution global models. NLOM uses 2D domain decomposition for scalability, based on the tiled data parallel programming style.² This is sufficiently general that it can use any one of several parallel programming approaches, including autotasking Fortran, OpenMP Fortran, Co-Array Fortran, Cray's SHMEM one-sided communication library, the MPI message-passing library, or both OpenMP and MPI.

The NA824 benchmark consists of a typical NLOM simulation of three model days on a $1/32$ -degree five-layer Atlantic Subtropical Gyre region (grid size $2,048 \times 1,344 \times 5$). Like most heavily used benchmarks, this is for a problem smaller than those now typically run. The NA824 speedup from 28 to 56 processors is similar to the 112 to 224 speedup for the $1/64$ -degree Atlantic model, which is four times larger.

Figure 1 summarizes the NA824 performance results for the best parallel programming approach on a wide range of machines. The sustained Mflops estimate is based on the number of floating-point operations reported by a hardware trace of a single-processor Origin 2000 run (without combined multiply-add operations)—that is, only useful flops (adds, multiplies, divides). A constant Mflops rate for all processor counts would indicate perfect scalability.

We tested on seven machines, and the four slowest are several years old. Three of these have excellent scalability from 14 to 112 processors, due to the low latency of one-sided (direct to memory) communications. The IBM SP is the fastest of the four but doesn't scale as well, primarily because it

Atmospheric versus Oceanic Prediction

Operational oceanography has lagged far behind atmospheric modeling because of two major complications. First, oceanic space and time scales are much different than those of the atmosphere. Ocean eddies are typically 100 km in diameter, which makes them 20 to 30 times smaller than comparable atmospheric highs and lows. As a result, approximately four orders of magnitude more computer time and three orders of magnitude more computer memory are required.

Second, unlike the meteorological radiosonde network that provides initial conditions from the surface to near the top of the atmosphere, there are few observations below the ocean surface at the synoptic time scale. Thus, effective oceanic data assimilative techniques are limited to surface satellite observations, which were not available until the 1990s. One advantage ocean prediction enjoys is that forecast skill for many ocean features, including ocean eddies and the meandering of ocean currents and fronts, is longer than the 10 to 14 day limit for atmospheric pressure systems.

Glossary of Terms

| | |
|-------|--|
| CMT | Capacitance Matrix Technique |
| HPCMP | High Performance Computing Modernization Program |
| NAVO | Naval Oceanographic Office |
| NLOM | NRL Layered Ocean Model |
| NRL | Naval Research Laboratory |
| RMS | Root-mean-square |
| SSH | Sea surface height |
| SST | Sea surface temperature |
| WSC | War-fighting Support Center |

must use two-sided MPI communications.

IBM and SGI recently introduced the three fastest machines, which are significantly faster than the previous generation of machines. The two new IBM systems use the same micro-processor but with different memory systems, numbers of processors per node, and switches between nodes. All three machines initially get superlinear scalability, typically due to better cache performance as the problem size per processor decreases. Cache effects appear to be much more significant on this latest generation of systems. Eventually, MPI's high latency leads to a fall off in scalability on the IBM SP systems. The particular Origin 3800 we tested was too small to measure large processor count effects.

Another facet of our efficiency drive is the use of an inexpensive data assimilation scheme²

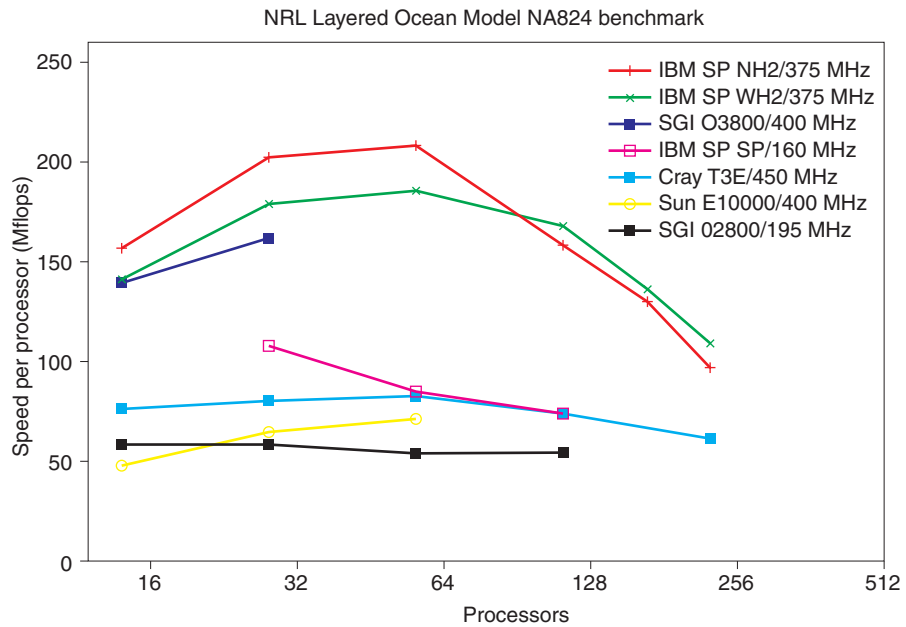


Figure 1. Performance of the NRL Layered Ocean Model NA824 benchmark on seven machines: the IBM NightHawk 2 and WinterHawk 2 (using MPI), the SGI Origin 3800 (using SHMEM), the IBM SP (using MPI), the Cray T3E (using SHMEM), the Sun E10000 (using OpenMP), and the SGI Origin 2800 (using SHMEM).

backed by a statistical technique for relating surface satellite data to subsurface fields.³ The statistics are from an atmospherically forced 20-year interannual simulation of the same ocean model, an application that requires a model with high simulation skill.

The NLOM system's focus on minimizing the computational cost is necessary if we are to provide near-global eddy-resolving capability on existing computers, but it comes at the price of relatively low vertical resolution and the exclusion of the Arctic (above 65 degrees North) and all coastal regions (shallower than 200 m). NRL is working on a second-generation global system without these limitations, but deployment is not scheduled until 2006 because of its much higher computational cost.

In October 2000, we achieved the major goal of our Fiscal Year 1998–2000 HPC Challenge—transitioning the world's first eddy-resolving nearly global (excluding the Arctic) ocean prediction system to NAVO for operational testing and evaluation. NAVO made this NLOM-based system an operational Navy product in September 2001.

The system consists of the 1/16-degree seven-layer, thermodynamic, finite-depth version of the NLOM for the global ocean (72 degrees S to 65 degrees N) and includes a mixed layer and sea surface temperature (SST).⁴ It was spun-up to real time using high-frequency wind

and thermal forcing from the Fleet Numerical Meteorology and Oceanography Center's Navy Operational Global Atmospheric Prediction System (FNMOC's NOGAPS). It assimilates SST plus real-time satellite altimeter data from three satellites using NAVO's Altimeter Data Fusion Center. It runs in real time on 216 Cray T3E or IBM WinterHawk 2 processors, with daily updates and a 30-day forecast performed every Wednesday. It provides a real-time view of the ocean down to the 50 km to 200 km scale of ocean eddies and the meandering of ocean currents and fronts. See www7320.nrlssc.navy.mil/global_nlom for real-time and archived results that include many zoom regions, nowcasts and forecasts of sea surface height (SSH), upper ocean currents, and SSTs, plus forecast verification statistics, the amount of altimeter data that went into each nowcast from each satellite, and nowcast comparisons with unassimilated data. Greater than 30-day forecast skill is routinely obtained globally.

Research at NRL and elsewhere suggests that we must resolve the oceanic eddy space scale, with about 3.5-km (mid-latitude) resolution needed globally for optimal performance.⁵ So far, HPC systems are only fast enough to allow about half this resolution under the strict real-time constraints of operational nowcasting and forecasting. However, our Fiscal Year 2001–2003




Figure 2. Sea-surface-height analysis (nowcast) in the Gulf Stream region from the real-time 1/16-degree global NRL Layered Ocean Model for (a) 4 June 2001 and (b) 11 June 2001. Superimposed on each is an independent Gulf Stream north-wall frontal analysis determined from satellite IR imagery (white lines) by the Naval Oceanographic Office for the same days. The color palette was chosen to emphasize the location of the Gulf Stream and associated eddies.

HPC Challenge project aims to double the resolution of the near-global NLOM system, with transition of a 1/32-degree version for operational testing and evaluation planned for 2003.

SSH nowcast comparisons with frontal analyses

The NRL has developed evaluation software and has been monitoring the performance of the 1/16° global NLOM system to establish the baseline metrics for this first-generation operational system. One evaluation monitors the system's ability to nowcast the positions of major fronts and eddies on the global scale. The Warfighting Support Center (WSC) at NAVO relies on satellite infrared (IR) SST data to locate fronts and eddies for the global ocean and re-

lease frontal analysis products to the fleet. The NLOM system lets the WSC analysis use daily nowcasts and animations of SSH to improve the quality of frontal analysis products.

This is particularly significant because SSH is a better indicator of subsurface frontal location than SST. Specifically, NLOM provides a daily map of the ocean mesoscale SSH field, which can help the WSC interpret cloud-filled IR images. In addition, by using animations of the NLOM SSH field, analysts can better track front and eddy movements to help analyze the space and time continuity of the ocean mesoscale in areas where frontal analysis is required.

Figure 2 shows the NLOM SSH nowcast for the Gulf Stream, overlaid with a portion of the daily WSC frontal analysis product for 4 June 2001 (Figure 2a) and 11 June 2001 (Figure 2b).

We must be careful when interpreting these results because of the uncertain quality of the WSC frontal analysis. Due to clouds, the analyst might not have seen the front for several days and might be depending on a previous (older) frontal analysis or analyzing for a shallow shingle of SST off the Gulf Stream. In the case for June, the WSC had good quality images and the frontal position has high confidence along most of the Gulf Stream path. However, the WSC analysis did not use the NLOM system output and is thus a completely independent analysis.

For both 4 and 11 June 2001, the NLOM Gulf Stream pathway agreed with the WSC frontal analysis. Both agreed on the position and shape of the meanders at 67 degrees West and 53 degrees West, features that changed shape between 4 and 11 June. During this period, the meanders deepen, as seen in both the model and the frontal analyses. The model also captures a warm ring-shedding event near 62 degrees West that the WSC frontal analysis verified. On 4 June, both products indicate a deep ridge near 60 degrees West that bends back toward the West, and on 11 June, both indicate that a warm eddy has shed from the Gulf Stream. There is also good agreement on a trough at 47 degrees West, which indicates that the model has a realistic pathway from Cape Hatteras to the Grand Banks when compared with the WSC frontal analysis.

Of course, the real value of the NLOM SSH frontal analysis is during times when the IR data is cloudy. The WSC plans to use the NLOM product as a first guess for their frontal-analysis product. The errors in the WSC analyses are mainly due to IR imagery that is difficult or impossible to interpret because of clouds. NLOM provides a useful guidance product for the WSC analysts to improve the frontal analysis in areas where it is cloudy, or during the warmer months when the fronts lose their surface SST signature.

SSH forecasts

NLOM's ability to forecast SSH and the positions of major fronts and eddies represents a new Naval product that can be used for future operational planning and to help users gauge the product's quality (by comparing forecasts with the analysis for that same day when it becomes available). The future positions of major ocean fronts will give the war-fighter some guidance on how changes in the ocean environment could affect future missions. An accurate SSH forecast would let the Navy predict changes in locations

of mesoscale features (fronts and eddies) that affect the 3D temperature and salinity field by using the predicted NLOM SSH and SST to derive synthetic profiles from the Modular Ocean Data Assimilation System.⁶ Changes in the 3D ocean environment would change acoustic conditions in the deep water and could be used to plan for future deployment of fleet assets.

Figure 3 shows an SSH example of a nowcast for 3 January 2001 (Figure 3a); the 15-day forecast (Figure 3b) and verifying analysis (Figure 3c) for 18 January 2001; and the 30-day forecast (Figure 3d) and verifying analysis (Figure 3e) for 2 February 2001 for the Kuroshio Extension region. The NLOM system performs a 30-day forecast once a week and calculates the verification statistics by comparing the model's daily prediction to the corresponding validating analysis for that same day. The Kuroshio is in "meander mode" near 140 degrees East and, in this case, the model predicts a widening of the meander after 30 days, which is also seen in the verifying analysis. The model also correctly predicts the deepening of the trough meander near 147 degrees East and the growth of the ridge just to the east at 149 degrees East, as well as the movement of the warm and cold eddies near 150 degrees East.

Figure 4 shows mean 30-day forecast verification statistics for the 19 monthly forecasts made each week from 20 December 2000 to 16 May 2001. It provides statistics for SSH root-mean-square (RMS) error and SSH anomaly correlation (a standard similar to that used in atmospheric forecast verification) versus forecast length for the entire global domain (Figures 4a and 4b), the Gulf Stream region (Figure 4c and 4d), and the Kuroshio region (Figure 4e and 4f). We calculated these statistics by comparing the model forecast SSH field versus the analysis SSH from NLOM valid on the day of the forecast. For both RMS error and anomaly correlation on the global domain, the model is beating persistence (that is, a forecast of no changes) and climatology (for RMS error) for more than 30 days. Anomaly correlation values greater than 0.6 indicate useful forecast skill.⁷

Verification statistics for the Gulf Stream and Kuroshio areas indicate the model's skill in regions dominated by flow instabilities and high SSH variability. In the Gulf Stream area, NLOM shows forecast skill only for approximately 18 days. The Kuroshio results are better than the Gulf Stream and are more similar to the global case, with skill out beyond 30 days.

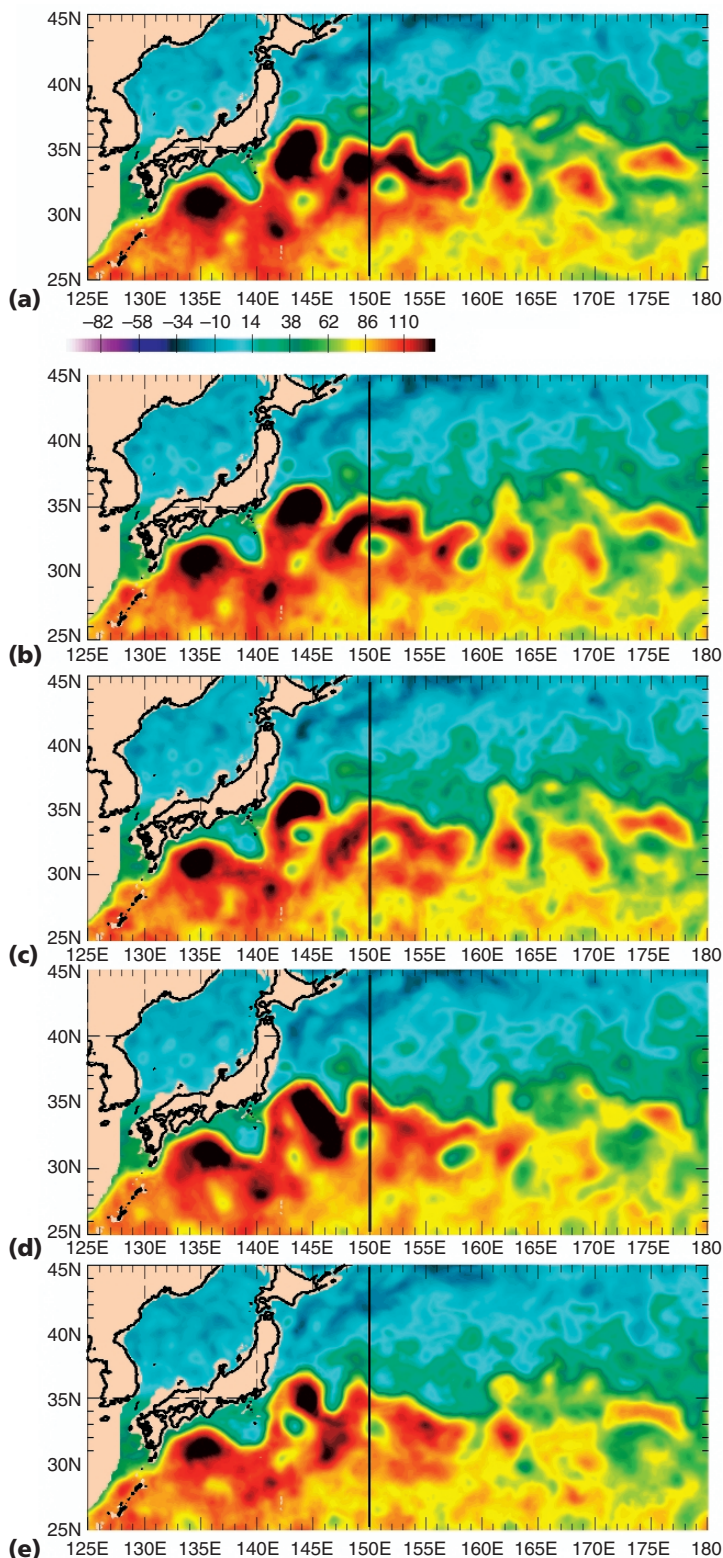
Although not shown here, NLOM SST pre-


Figure 3. Sea surface height (cm) for the Kuroshio region from the 1/16-degree global NRL Layered Ocean Model running in forecast mode for a 30-day forecast. (a) We initiated the model forecast on 3 January 2001 and ran it for 30 days. We obtained (b) forecast and (c) validating analysis for 18 January 2001, as well as (d) forecast and (e) validating analysis for 2 February 2001. The model forecast was forced by the Fleet Numerical Meteorology and Oceanography Center's Navy Operational Global Atmospheric Prediction System surface winds and heat fluxes for four days and then gradually switched to climatological forcing for the remainder of the 30-day forecast.

diction metrics are also calculated for 30-day forecasts and show that NLOM, on average, beats both persistence and climatology for the global domain with the help of relaxation to climatologically corrected persistence. Real-time SST prediction metrics are similar to those shown for SSH. In addition, nowcast SST comparisons with unassimilated buoy data are updated daily and are available on NRL's NLOM Web site (www.7320.nrlssc.navy.mil/global_nlom).

The difference in model skill between the Kuroshio and Gulf Stream is due to the resolution of the NLOM system and the dynamics of the two strong western boundary currents. The results shown here are consistent with other NRL research,⁵ which indicates that at least 1/32-degree resolution is required for proper simulation of the Gulf Stream. The forecast statistics indicate that the Kuroshio dynamics are well represented by the model. For most of the ocean, 1/16-degree resolution is enough to allow skillful forecasts of SSH. Even though the current system does not show 30-day skill for the Gulf Stream, the results indicate that the system has skill to at least 18 days. NRL is currently working on a 1/32-degree global system in research and development, which should significantly improve the Gulf Stream results.

In today's mood of military uncertainty and heightened awareness, it is imperative that the US Navy operate with improved efficiency and navigation. In an effort to fully exploit the marine environment in which naval submarines and ships must operate, HPC is making a significant contribution to the success of the Navy's mission through a cost-effective system of global oceanic modeling with high-resolution,



real-time nowcasts and forecasts. Together, HPCMP and NRL are working continually to apply state-of-the-art computational capability to the system in a concerted effort to provide the most accurate and timely information available. 

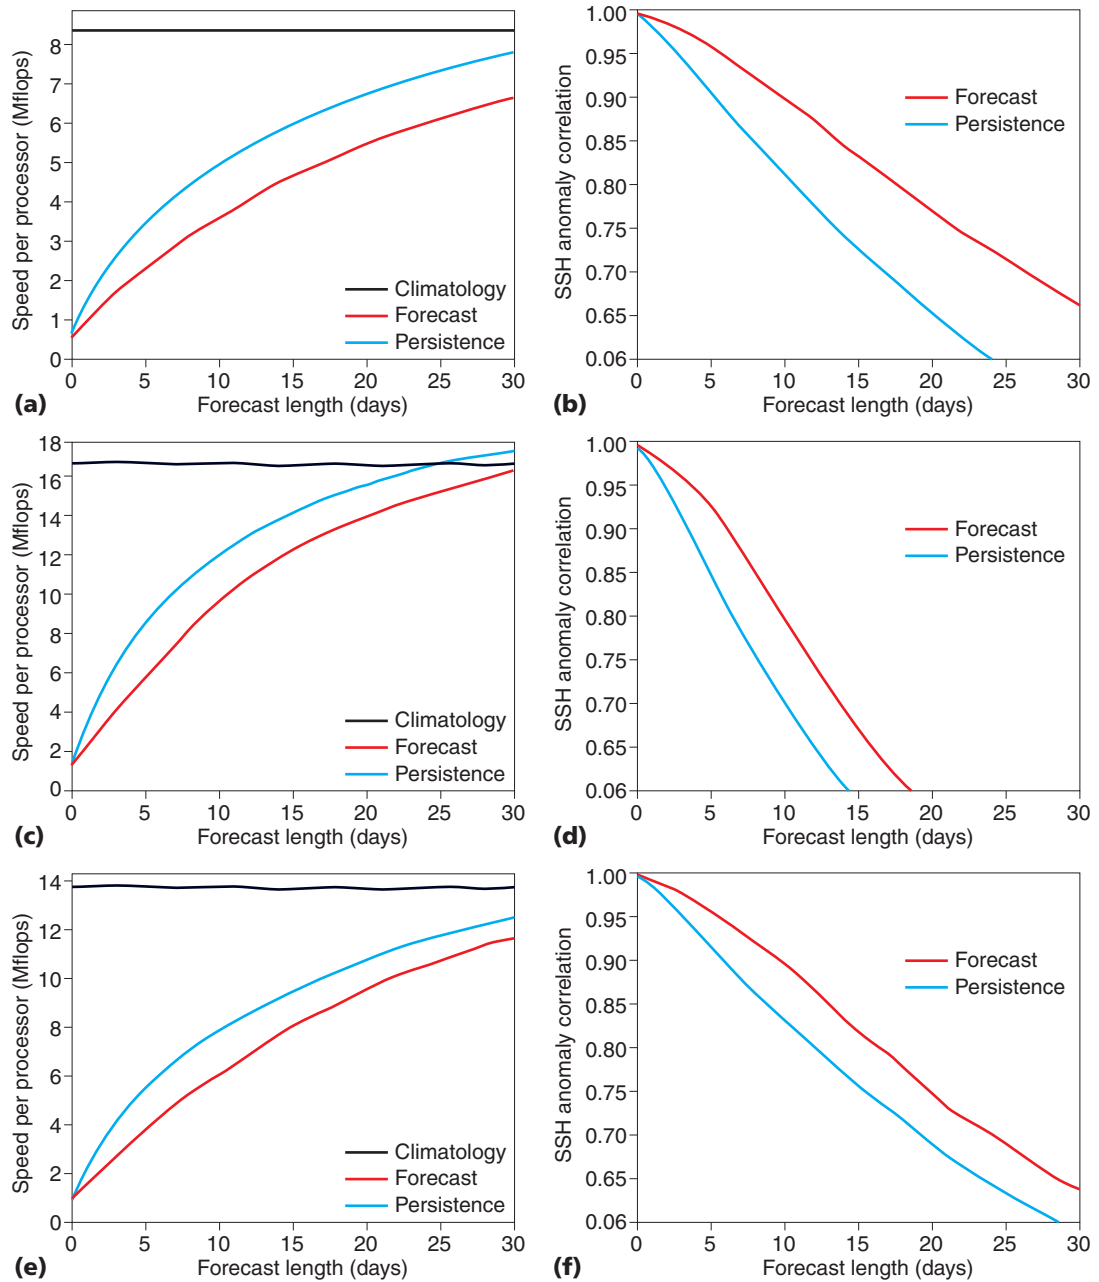


Figure 4. Mean sea-surface-height forecast verification statistics for 19 weekly 30-day forecasts from 20 December 2000 to 16 May 2001 for the 1/16-degree global NRL Layered Ocean Model. Left column shows mean SSH RMS error (cm) and the right column shows mean anomaly correlation versus forecast length (days) for NLOM forecast (red curve), persistence forecast (blue curve), and climatology forecast (black curve). The top row (a, b) is for the global domain, the middle row (c, d) is the Gulf Stream region, and the bottom row (e, f) is the Kuroshio region.

Acknowledgments

This work is a culmination of many projects over many years. Individuals who have made significant contributions include A. Birol Kara (Sverdrup Technology) and Jean-Francois Cayula (Planning Systems). This is a contribution to the 6.2 Basin-Scale Ocean Prediction System project funded

by the Office of Naval Research under program element 62435N and to the 6.4 projects Large Scale Ocean Modeling and Ocean Data Assimilation funded by the Space and Naval Warfare Systems Command under program element 63207N. The numerical model was run on the Cray T3E and IBM SP at the Naval Oceanographic

Office, Stennis Space Center, Mississippi, and on the Cray T3E at the US Army Engineer Research and Development Center, Vicksburg, Mississippi. Both are part of the Defense Department's High Performance Computing Modernization Program.

References

1. A.J. Wallcraft and D.R. Moore, "The NRL Layered Ocean Model," *Parallel Computing*, vol. 23, 1997, pp. 2227–2242.
2. O.M. Smedstad et al., "A Real-Time 1/16° Global Ocean Nowcast/Forecast System," submitted to *J. Marine Systems*, 2001 (available from Smedstad; smedstad@nrlssc.navy.mil).
3. H.E. Hurlburt, D.N. Fox, and E.J. Metzger, "Statistical Inference of Weakly Correlated Subthermocline Fields from Satellite Altimeter Data," *J. Geophysical Research*, vol. 95, 1990, pp. 11375–11409.
4. A.J. Wallcraft et al., "The NRL Layered Global Ocean Model (NLOM) with an Embedded Mixed Layer Sub-Model: Formulations and Model Simulations," submitted to *J. Atmospheric Oceanic Technology*, 2001 (available from Wallcraft; wallcraft@ajax.nrlssc.navy.mil).
5. H.E. Hurlburt and P.J. Hogan, "Impact of 1/8° to 1/64° Resolution on Gulf Stream Model-Data Comparisons in Basin-Scale Subtropical Atlantic Ocean Models," *Dynamics Atmosphere and Ocean*, vol. 32, 2000, pp. 283–329.
6. D.N. Fox et al., "The Modular Ocean Data Assimilation System (MODAS)," to be published, *J. Atmospheric and Oceanic Technology*, 2002.
7. A.H. Murphy and E.S. Epstein, "Skill Scores and Correlation Coefficients in Model Verification," *Monthly Weather Rev.*, vol. 117, 1989, pp. 572–581.

Alan J. Wallcraft is a computer scientist for the Ocean Dynamics and Prediction Branch of the Naval Research Laboratory. His research interests include the numerical solution of finite difference partial differential equations as well as ocean model design, development, and maintenance. He received his BSc in mathematics and computer science from the University of Essex, England, and his PhD in numerical analysis from Imperial College, University of London. Contact him at wallcraft@nrlssc.navy.mil.

Harley E. Hurlburt is the senior scientist for Ocean Modeling and Prediction at the Naval Research Laboratory, where he leads a research team on eddy-resolving global ocean modeling and prediction. He is also an adjunct faculty member at Florida State University and the University of Southern Mississippi. He received his BS in physics from Union College, NY, and his MS and a PhD in meteorology from Florida State University. He is a member of the US and International Steering Teams for the Global Ocean Data Assimilation Experiment (GO-DAE), a multinational project designed to help justify a permanent global ocean observing system by demonstrating real-time global ocean products that are widely available. Contact him at hurlburt@nrlssc.navy.mil.

E. Joseph Metzger works at the Naval Research Laboratory in the Large-Scale Modeling Group, which uses

both global and basin-scale versions of the NRL Layered Ocean Model (NLOM). His research has focused on the dynamical interactions between the Pacific Ocean and the South China Sea, with an emphasis on how the Kuroshio intrudes and sheds eddies in this semi-enclosed sea. Other work includes SST forecasting studies using NLOM. He received his BA in physical geography from Ohio State University and his MS in meteorology from the University of Wisconsin-Madison. He is a member of the American Geophysical Union and Sigma Xi. Contact him at metzger@nrlssc.navy.mil.

Robert C. Rhodes is the head of the Ocean Monitoring and Prediction System Section at the Naval Research Laboratory. He has been working on the development and transition to Naval operations of real-time ocean prediction systems on the global and regional scale. His research interests include ocean modeling, ocean data assimilation, and development of ocean nowcast and forecast systems. He received his BS in atmospheric science from the State University of New York at Albany and his MS in meteorology from Texas A&M University. He is a member of the American Geophysical Union and Sigma Xi. Contact him at rhodes@nrlssc.navy.mil.

Jay F. Shriver is an oceanographer in the Ocean Dynamics and Prediction Branch at Naval Research Laboratory. He has published in the areas of El Niño, decadal variability, and global- and basin-scale ocean circulation. His current research interests include the global thermohaline circulation, nonsteric sea surface height variability, and nonlinear dynamics. He received his BS in meteorology and mathematics from the State University of New York College at Brockport and his MS and PhD in meteorology from the Florida State University. Contact him at shriver@nrlssc.navy.mil.

Ole Martin Smedstad is a contractor for the Naval Research Laboratory at the Stennis Space Center and works for Planning Systems, Inc. He has published more than 20 papers in the area of open boundary conditions in numerical models, ice modeling, and data assimilation in ocean models. His experience with data assimilation methods range from simple techniques such as OI to advanced methods such as variational (adjoint) assimilation techniques. He received a BS in applied mathematics and an MS in physical oceanography from the University of Oslo, Norway. He received a PhD in geophysical fluid dynamics from Florida State University. Contact him at smedstad@nrlssc.navy.mil.

For more information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.