An Application Framework for the Rapid Deployment of Ocean Models in Support of Emergency Services: Application to the MH370 Search

Uwe Rosebrock, Peter R. Oke, and Gary Carroll

CSIRO Oceans and Atmosphere Flagship, Hobart, TAS, Australia
{uwe.rosebrock,peter.oke,gary.carroll}@csiro.au

Abstract. Ocean models are beneficial to many different applications, including industry, public-good, and defence. Many applications use high-resolution models to produce detailed maps of the ocean circulation. High-resolution models are historically time-consuming to configure – often taking weeks to months to properly prepare. A system for automatically configuring and executing a model to predict the past, present, or future state of the ocean – named TRIKE – has been developed. TRIKE includes a sophisticated user interface allowing a user to easily set a model in minutes, control data management and execute the model run on a supercomputer almost instantly. TRIKE makes it feasible to configure and execute high-resolution regional models anywhere in the world at a moments notice. TRIKE was used to support the recent search for MH370 that is thought to have crashed somewhere in the South Indian Ocean.

Keywords: integrated environmental modeling, software control framework, connectivity, large data management, model management.

1 Introduction

Ocean models are routinely applied to predict the past, present, and future state of the ocean to better understand ocean dynamics and to support a range of industrial, public-good and defense applications. Since the advent of the Global Ocean Data Assimilation Experiment (GODAE; www.godae.org), a global initiative to develop operational ocean forecasting capabilities, several global and basin-scale ocean forecast systems have been developed with resolution of up to 10 km [5]. This is beneficial to many applications, but delivers forecasts at insufficient resolution for many situations that require resolution of up to 1 km. For these applications requiring high-resolution, it is common to configure and run regional models, where the domain extends only a few hundred kilometres or less. Historically, regional ocean models are time-consuming to configure, requiring an expert modeller to configure a grid – carefully setting the spatial extent of the model domain and the model resolution. The expert continues the process by acquiring the data to support the running of the ocean model. This generally involves the gathering of the data from multiple sources to set the bathymetry, coastline, initial conditions as well as forcing fields. The modeller then reformats the various datasets to meet the requirements of the specific ocean

R. Denzer et al. (Eds.): ISESS 2015, IFIP AICT 448, pp. 235–241, 2015.

model and configures the ocean model dynamics for the chosen region – carefully considering the dominant physical processes and requirements to properly represent the important characteristics of the regional circulation. The modeller selects a suite of model parameters, numerical schemes, and boundary conditions that are appropriate for the chosen application. In the past, this process would take several weeks for an experienced modeller, and several months for a less-experienced modeller, making real-time applications inconceivable. Moreover, the idea of a non-expert user (e.g. a marine biogeochemist) was merely a pipe dream.

However, the development of a system called TRIKE (or 'wheels for modellers'), by the CSIRO has made the configuration and execution of a regional ocean model simple. With a sophisticated user interface, a user can quickly and easily set up a model domain, select the desired input data sources, and execute the model. TRIKE allows a user to run a model simulation for any historical period, or for a forecast – predicting the ocean circulation for several days into the future. This paper documents the salient features of TRIKE, and describes a recent application that supported the Australian Maritime Safety Authority in the search for the ill-fated flight MH370.

2 Trike

The TRIKE user interface (Fig. 1) allows a user to use a graphical interface to set the model domain, the model resolution, the input data sources, and other key specifications of the model run (e.g., what time-period to simulate, and when to execute the

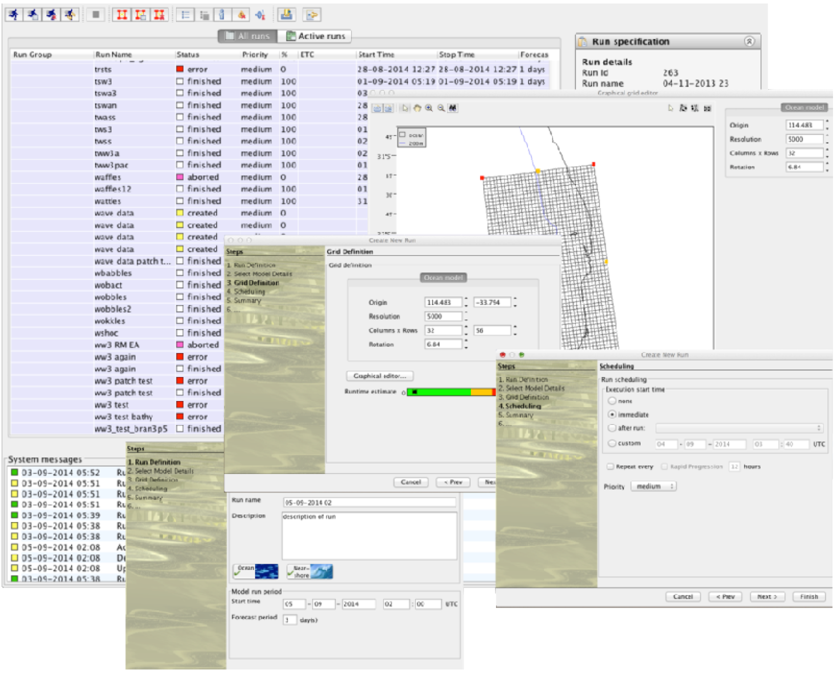

Fig. 1. The TRIKE user interface and wizard

model). The back-end of TRIKE then configures a state-of-the-art ocean model [4] for a specific area. Ocean models solve mathematical equations on a computer to simulate the time-evolution of the ocean circulation; representing impacts of surface forcing (e.g. wind) and tides [2]. A large amount of data is needed for every application. TRIKE automates this process of configuring the model and massively reduces the setup time required to less than a minute (historically weeks to months). TRIKE consists of several software components, including an extensive data management and interrogation component. The software components manage the user interactions and apply extensive heuristics to create the scientific modelling domain.

2.1 Defining the Model Domains

TRIKE provides a simple-to-use application that uses the notion of a "run wizard" to select model characteristics such as the model type and the temporal and geographical extent. The original aim in developing TRIKE was to provide a facility to allow the running of complex models with minimal human effort. This was extended to also allow users with minimal understanding of the actual modelling process access to initiate, configure and execute model runs independently. The result is a "desktop-client" leading a user through the process of setting up a model run with simple steps. (Depending on the permissions of users more sophisticated options are made available.) The complexity of the various parameter ranges from output directives to creating sophisticated interpolations of a domains bathymetry. In addition TRIKE provides the facility of high level scheduling and repetitions (e.g. run process every 12 hours). The scheduling extends from simple run-later to creating dependencies between runs (e.g. process x starts after process y has finished).

2.2 Data Preparation

The nature of numerical models dictates very comprehensive input data requirements. TRIKE is capable of maintaining an extensive database of forcing data and automates the process of data preparation for a particular model; configuring a model process and selects, extracts, prepares and assembles the appropriate input data. To be able to utilise different data sources and to meet the requirements of an ocean model, a matching from provided to the required data is necessary. As an example of a data provider; the Australian Bureau of Meteorology performs many weather forecasts each day. Internationally, many weather centers perform equivalent forecasts to the Bureau, generating equivalent data, but disseminating data in different formats (e.g., with different variable names, units etc). TRIKE matches the data sources it manages to what is required by a model by using metadata vocabularies and runs processes to convert data so the receiving model can ingest the data correctly. The regional data used can be available through online services or stored locally. The ingestion of data can also be localised or provided through online services. The provisioning of the data is a very important aspect of the model preparation process, particularly for data sets updated on a regular basis, the TRIKE sub-system is capable of monitoring and responding to changes in the data availability.

2.3 Running a Model

Once the modeling process has been prepared it can then be run on a remote high-performance computer, on a local compute server, or the workspace can be made available as an archive to be run on user infrastructure outside TRIKE. The progress of the computational process can be monitored through the application's user interface as the modelling process is executed (see Fig. 1). The execution may utilise high per-formance computing job queues or simply execute a process remotely. Either is mon-itored and can be aborted. Deployment into cloud-based compute-resources has also been implemented as a prototype.

TRIKE consists of several components that function independently of each other and communicate with each other via dedicated protocols. Fig. 2 shows an overview of the management component Run Coordination System (RCS), the Data Management Framework (DMF) and the interface to the individual modelling components. Fig. 3 shows a generalised workflow of creating a modeling process. Once the process has successfully finished, post-processing can consist of further analysis and/or packaging. After a process is complete, results are transferred to a server for further analysis. Scientific users of TRIKE can receive bundled model run workspaces that allow the user to fine-tune the model configuration and rerun the process.

Originally TRIKE was developed for the Royal Australian Navy (RAN), under the Bluelink partnership [6]. Today the RAN uses a dedicated customised system routinely for tactical briefings prior to operations. TRIKE also has enormous potential benefits for public applications including search and rescue at sea, modelling of effluent dispersal following oil spills, environmental prediction of ocean currents to inform industry of risks during operations (e.g. oil and gas exploration), and storm surge prediction during tropical cyclones.

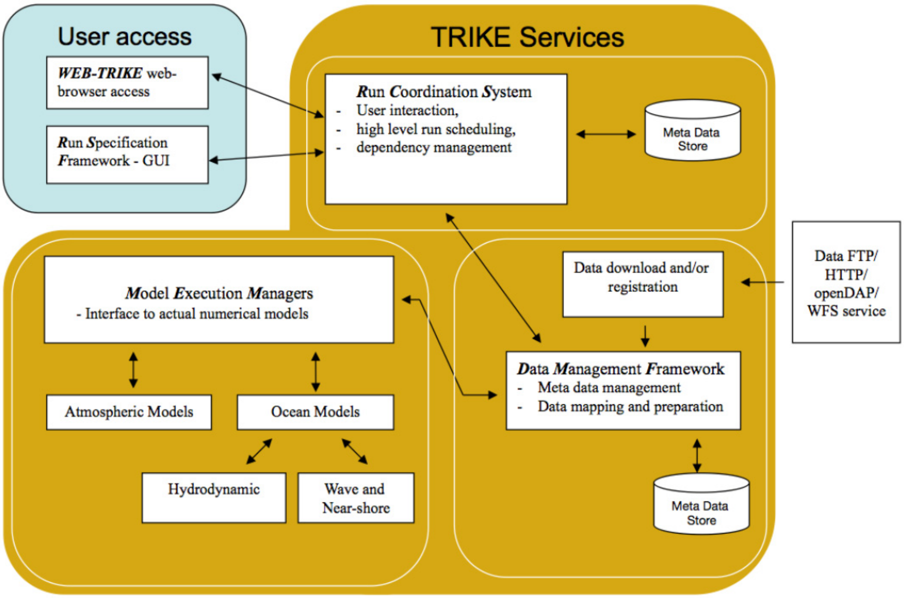

Fig. 2. An overview of the TRIKE framework

As a research tool it can provide easy access to determine for example biological connectivity, path of invasive species [1], marine debris in general [3], abandoned fishing equipment [7] or support more general coastal inundation modelling. TRIKE currently supports several ocean models and wave models as well as a regional atmospheric model. It has the capacity for further integration of models or any other data driven processes as well as creating dependencies between modelling process (nesting of one model within the data space of another, parameter estimation).

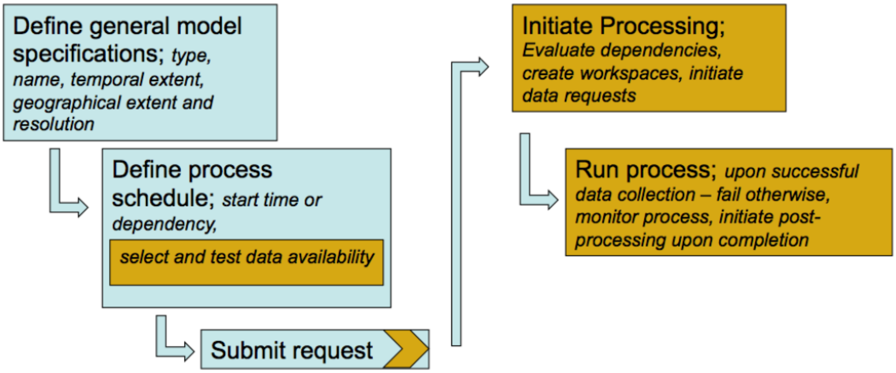

Fig. 3. The general workflow to initiate a TRIKE process, the background colours correspond to the location and correspondence of components in Figure 2

2.4 Current and Future Work

A web-based user interface has been developed under the MARVL (http://www.marvl.org.au) project to simplify maintenance and expand the user uptake. It forms the basis for several other projects within CSIRO such as eReefs (http://www.ereefs.org.au). Further developments under the 'Australian Wave Energy Atlas' are underway to incorporate the advances provided by accessing the local graphical processing unit (GPU) from within a web-browser to present and edit complex geometries defining the modelling domain. The intention for future development is to also move some of the processes, like interpolations currently performed on the server, into the GPU to improve the user experience.

3 Application to MH370

Malaysian Airlines Flight 370 disappeared on 7 March 2014 at 17:20 UTC. A multinational search effort subsequently began, initially focusing on the Gulf of Thailand and the South China Sea, before shifting to the South Indian Ocean. Using a range of different techniques, several possible "splash-points" were identified – but all with relatively high-levels of uncertainty. Many groups around the world began trying to predict the possible fate of surface debris that may have resulted from the crash.

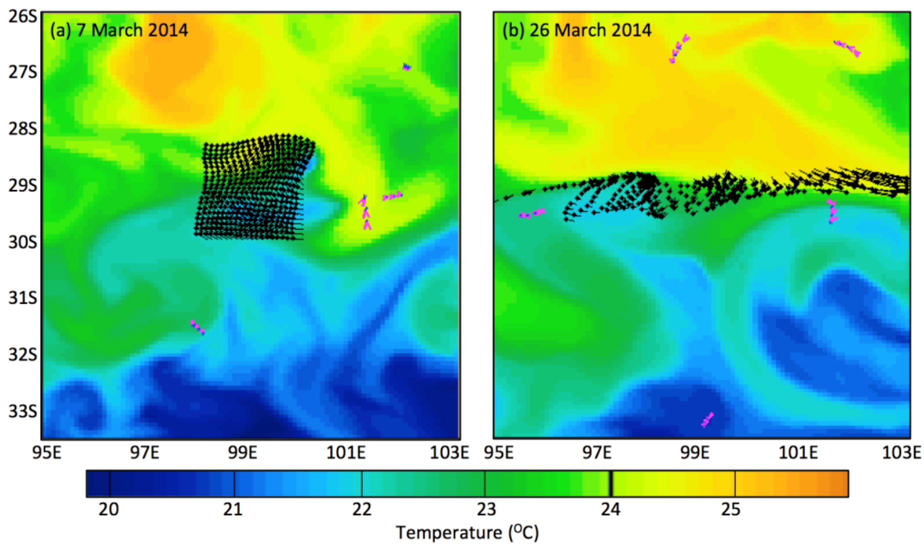

Fig. 4. Examples of the modelled ocean temperature (colour) and modelled particles (black cells) tracking the possible trajectories of surface debris from a possible crash site of the MH370 plane for 7 March 2014 and 26 march 2014. The pink arrowheads show the location and velocity of surface drifters.

Many different estimates of surface velocities were used to support the search, including near-real-time maps of sea-level and surface velocities from satellite altimetry (e.g. oceancurrent.imos.org.au), estimates of the past, present, and future ocean circulation from operational ocean forecasts and regional ocean forecasts generated by TRIKE. An example of the type of analysis performed using 2-km resolution TRIKE forecasts (limited by computation resources) are shown in Fig. 4, showing the locations of a cluster of hypothetical particles (intended to represent debris from a possible crash site) at the time that MH370 is thought to have crashed, and locations of the particles 19 days later. In this example, there was a front separating relatively warm water (yellow-orange in Fig. 4) and colder water (green-blue in Fig. 4). For this simulation, the model predicts that if the plane crashed into the ocean along this front, then any resulting surface debris would likely spread along the front. Information like this was disseminated to decision-makers who were coordinating the search effort for MH370.

The data produced by TRIKE that formed the basis of the analysis outlined as well as the imagery shown in Fig. 4, can easily be incorporated as a post-processing step of a modeling process managed in TRIKE.

4 Conclusion

As environmental computational modelling is becoming a more and more standard tool used in science and environmental services, it appears crucial that a supporting infrastructure is also developed. TRIKE has demonstrated to be capable of supporting numerous numerical models as well as the versatility to meet a broad set of demands. The current state of TRIKE is the result of many years of development. The Royal Australian Navy has made use of a derived version of TRIKE called ROAM (Relocatable Ocean and Atmosphere Modelling framework) for some time. The lessons learnt from that deployment and the re-investment into the concept have saved innumerous hours of science time and promises to lower the hurdle of access to numerical modeling further through the current developments.

Although the search of the MH370 plane did not result in a positive detection of the remains of the aircraft, the applications allowed the developers of TRIKE to showcase its versatility. It will undoubtedly be used for future applications such as support for search and rescue, oil spill response and as a research tool.

References

1. Brinkman, D.: Could ocean currents be responsible for the west to east spread of aquatic invasive species in Maritime Canadian waters? Marine Pollution Bulletin 85, 235–243 (2014)
2. England, M.H., Oke, P.: Ocean modeling and prediction. In: Leslie, L., Peng, P., Shao, Y. (eds.) Environmental modelling and prediction, pp. 125–171. Springer (2001)
3. Hardesty, B.D., Wilcox, C.: Understanding the types, sources and at-sea distribution of marine debris in Australian Waters. Final report to the Department of Sustainability, Environment, Water, Health, Population and Communities (2011)
4. Herzfeld, M.: Improving stability of regional numerical ocean models. Ocean Dynamics 59(1), 21–46 (2009)
5. Hurlburt, H.E., Brassington, G.B., Drillet, Y., Kamachi, M., Benkiran, M., Bourdalle-Badie, R., Chassignet, E.P., Jacobbs, G.A., Le Galloudec, O., Lellouche, J.-M., Metzger, E.J., Oke, P.R., Pugh, T.F., Schiller, A., Smedstad, O.M., Tranchant, B., Tsujino, H., Usui, N., Walcraft, A.J.: High-resolution Global and Basin-Scale Ocean Analyses and Forecasts. Oceanography 22(3), 80–97 (2009)
6. Schiller, A., Meyers, G., Smith, N.: Taming Australia's last frontier. Bull. Am. Met Soc., 436–440 (April 2009)
7. Wilcox, C., Hardesty, B.D., Sharples, R., Griffin, D.A., Lawson, T.J., Gunn, R.: Ghostnet impacts on globally threatened turtles, a spatial risk analysis for northern Australia. Conservation Letters 1(2012), 1–8 (2012)