

Overview of CCSM Infrastructure – Implications for ACCESS

Anthony Craig

National Center for Atmospheric Research, Boulder, Colorado, USA

The National Center for Atmospheric Research (NCAR) Community Climate System Model (CCSM) is a global climate model consisting of four primary components (atmosphere, ocean, land, and sea-ice) that are coupled together to simulate historical, present, and future climate. While the scientific focus of the model has not fundamentally changed over its ten-year history; the models, infrastructure, and organisation have evolved dramatically. CCSM3 (Collins et al. 2005) was publicly released in 2004 with significant improvements in resolution, complexity, and quality compared to the first version of CSM1 of the mid-1990s. Over the last decade, the CCSM project has evolved from an NCAR based model to a large community project, and many aspects of the infrastructure have evolved concurrently with the project.

In 2001, the CCSM Software Engineering Working Group (SEWG) was formed to bring together community interests associated with the technical aspects of CCSM. At about the same time, a core group of software engineers at NCAR were reorganised to specifically support CCSM model and process infrastructure (CSEG). The goals of CSEG are to support scientific and technical model development, lead system coupling and coupler development, oversee process infrastructure, put in place more robust testing and validation standards, improve model performance and portability, manage and develop common source code and scripts, and take leadership of hardware and software issues. This reorganisation led to several changes in CCSM process and many improvements in the technical infrastructure of CCSM.

CCSM infrastructure process now includes source code control (Subversion), change review, new web pages and wikis to facilitate communication and lightweight project tracking, bug tracking (Bugzilla), bulletin boards (vBulletin), and the implementation of a database (MySQL) to track model runs. Every effort has been made to keep process lightweight. CCSM documentation is much improved with documentation (both hard and soft) provided in support of CCSM process and policies, plans and goals, portability and performance, tags and test results, coding standards, and user guides.

Within the CCSM model infrastructure, significant improvements have been made in the coupler, in shared source code, in the use of external infrastructure packages, in build and run scripts, and in automated testing and validation. These efforts have resulted in less redundant code, greater portability, reduced maintenance, improved extensibility and increased robustness. One metric that highlights the success of these efforts is CCSM portability. In the 1990s, CSM1 ran on a very limited number of platforms and porting was extremely difficult. This has improved dramatically in the last 5 years, and CCSM has successfully run on many diverse systems including IBM Power 3/4/5, NEC Earth Simulator, Cray X1E, Cray XT3, and Linux clusters. Because of this capability, CCSM was able to contribute 9100 simulated years and 9.6 Tbytes of T85/1deg coupled model data to the IPCC AR4 while continuing to have access to hardware resources to further model development.

CCSM is a complex model. Model infrastructure and process are constantly reviewed and always evolving, and there are many challenges ahead. However, changes in the last few years to both CCSM process and model infrastructure have resulted in significant improvements in key metrics like quality, capability, and productivity.

Collins, W. D., C. M. Bitz, M. L. Blackmon, G. B. Bonan, C. S. Bretherton, J. A. Carton, P. Chang, S. C. Doney, J. J. Hack, T. B. Henderson, J. T. Kiehl, W. G. Large, D. S. McKenna, B. D. Santer, and R. D. Smith, 2006. The Community Climate System Model Version 3 (CCSM3). *Journal of Climate*, 19, 2122-2143.