

MindModeling@Home ... and Anywhere Else You Have Idle Processors

Jack Harris (jack.harris@us.af.mil)

Kevin A. Gluck (kevin.gluck@us.af.mil)
Air Force Research Laboratory, 6030 S. Kent St.
Mesa, AZ, 85212 USA

Thomas Mielke (thomas.mielke@mesa.afmc.af.mil)

L-3 Communications at AFRL, 6030 S. Kent St.
Mesa, AZ, 85212 USA

L. Richard Moore, Jr. (larry.moore@mesa.afmc.af.mil)

Lockheed Martin at AFRL, 6030 S. Kent St.
Mesa, AZ, 85212 USA

Keywords: MindModeling; Computational Cognitive Process Modeling; Volunteer Computing; Distributed Computing

Introduction

As we, the ICCM community, continue to expand the scope of our cognitive modeling ambitions, we increasingly face computational requirements that are an impediment to progress. Computational complexity grows quickly with increases in the granularity of models, the fidelity of the models' operating environment, and the time scales across which these models interact. Additional processing demands are encountered when studying the breadth of a cognitive model's performance capabilities such as through observing the model's sustained fitness while varying the environment or conducting sensitivity analyses of interactions between internal model parameters in a controlled experiments. Such computational demands are not unique to the cognitive modeling community. Other scientific fields (bioinformatics, meteorology, physics, etc.) have already pioneered a variety of platforms and methodologies for dealing with similarly computationally complex problems. We will achieve faster progress toward the broader scientific objectives of cognitive modeling and the specific goals of particular research projects if we pay attention to the lessons learned and capabilities developed in other computational sciences.

Volunteer Computing

An exciting methodological development of the past decade has been the rise of volunteer computing as a means of acquiring access to large numbers of computer processors distributed across the internet. Volunteer computing represents a huge and increasingly powerful computational resource due to the continuous growth rate of end-user processing capability around the world. The first volunteer computing project was SETI@Home. It was established in 1999 for the purpose of demonstrating the utility of "distributed grid computing" by providing a mechanism for analysis of massive amounts of observational radio

telescope data. The scientific Search for Extra-Terrestrial Intelligence rapidly caught and held the public imagination, and SETI@Home remains the longest running and one of the most popular volunteer computing projects in the world. This actually is an impressive feat, given that the volunteer computing concept has caught on in an assortment of other scientific communities and there are now approximately three dozen volunteer computing projects available to those interested in donating their idle processor time to scientific pursuits. Most of them, including SETI@Home, run on a software architecture called the Berkeley Open Infrastructure for Network Computing (BOINC). Some of the other large BOINC-based scientific volunteer computing projects include: Climateprediction.net (sensitivity analyses of models that predict Earth's climate up to 2080), MilkyWay@Home (investigating optimization methods for Internet-based computing and developing 3-dimensional models of the Milky Way galaxy), and Einstein@Home (searching for pulsars in data from the LIGO gravitational wave detector). In total, as of April 14, 2009 (the submission date for ICCM 2009), BOINC-based volunteer computing projects include 291,956 active volunteers, offering 531,174 computer hosts. That level of volunteerism is producing an average computational throughput, across all projects, of 1,368 TeraFLOPS.

The largest existing volunteer computing project does not run on the BOINC platform. It is called Folding@Home and is dedicated to understanding protein folding. Folding@Home currently, as of the submission date, has more than 4 million volunteered computers and is producing a throughput of 4,782 TeraFLOPS. By comparison, the world's current fastest centrally-managed High Performance Computing system, at the United States Department of Energy's Oak Ridge National Laboratory, has a peak processing capacity of 1,640 TeraFLOPS, so clearly there is a lot of computational power and potential scientific return available through distributed, volunteer grid computing. It will benefit computational cognitive scientists to begin taking advantage of this platform. Happily, there are now

two cognitive science-related volunteer computing projects. One, called Artificial Intelligence System, is an AI project hoping to achieve large scale artificial intelligence by reverse engineering the brain. The other is MindModeling@Home.

MindModeling@Home

Launched in March of 2007, and still (intentionally) in Beta status, MindModeling@Home focuses on utilizing computational cognitive process modeling to better understand the human mind and to improve on the scientific foundations that explain the mechanisms and processes that enable and moderate human performance and learning. It accomplishes this goal using the BOINC software augmented with in-house development of web-based user interfaces and community portals. Together these tools attempt to bridge the gap between the complex engineering challenges of large scale computing and the domain specific requirements of the cognitive science research community.

Since its inauguration two years ago, the system has completed over 50 jobs — most of which involve millions of model runs — which substantially contributed to various research efforts both within and external to our organization. Most of these jobs were completed exclusively by volunteers donating computing time from over 3000 machines (typically 200 to 900 at any given moment in time). However, MindModeling@Home is not limited to volunteers, as we have also achieved integration with local resources as well as several high performance computer clusters.

There are many lessons learned and remaining challenges associated with using heterogeneous computational resources beyond those faced when attempting to use homogenous computing clusters. Some challenges include how to schedule work with virtually no consistent expectation of availability, how to gauge progress and report status to customers, how to ensure that models are written properly and behave appropriately on Linux, Windows and Mac OS X, how to recover when those resources fail, and how to ensure the level of data integrity required for scientific publications. These are non-trivial engineering efforts, and it behooves the computational cognitive science community to leverage existing work in this space.

MindModeling@Home currently supports Common Lisp-based cognitive models. Work has begun to support models written in other languages such as Java, Python and Scheme. The hope of this expansion is to open the door for different types of cognitive research to be supported by this framework; thereby making the computational resource pool available to a broader cross-section of the cognitive modeling community. Other future work includes the exploration of special purpose processors such as the Graphics Processing Units (GPUs) which have enormous computational ability but do not support general process calculations currently used in the models of MindModeling@Home. Also, an effort is underway to better parallelize model component pieces. The ability to

parallelize not only experiment parameterization configurations but also the cognitive model and environment itself provides the ability to support modeling efforts at a depth not possible in single processing environments.

Our attempts to address cognitive modeling's growing computational demands are not limited to acquiring computational resources. The problem is also being tackled through research in exploration and search algorithms. MindModeling.org already incorporates an experimental "adaptive mesh refinement" (AMR) algorithm to intelligently prune and interpolate parameter spaces, and as a result of some very recent research efforts (Best et al., submitted), a smoothing algorithm has been added to help reduce resampling requirements. Integrating intelligent algorithms in the context of large scale computing has proven to be surprisingly challenging, but we see these challenges as opportunities for active involvement by the broader cognitive modeling and computational sciences communities.

A long range goal of MindModeling.org is to abstract away the challenges of using large-scale resources (and using large-scale resources well) for cognitive modelers. Unlike most volunteer computing projects, in which a single research project/team/center drives all of the computational demands, we would prefer to turn MindModeling@Home into a cognitive modeling community resource, enabling cognitive modelers all over the world to harness the power of the system by submitting their own cognitive model batch runs. We believe the distributed power of computational resources available to a distributed MindModeling community will facilitate new advances in computational cognitive process modeling otherwise not possible.

Acknowledgments

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. This research was sponsored by the Air Force Research Laboratory's Warfighter Readiness Research Division.

References

- Best, B. J., Furjanic, C., Gerhart, N., Fincham, J., Gluck, K. A., Gunzelmann, G., & Krusmark, M. A. (submitted). Adaptive mesh refinement for adaptive exploration of cognitive architectures and cognitive models. Under consideration for presentation at the 9th *International Conference on Cognitive Modeling*.
- Harris, J. (2008a). MindModeling@Home: A large-scale computational cognitive modeling infrastructure. In *Proceedings of the 6th Annual Conference on Systems Engineering Research* (<http://cser.lboro.ac.uk/CSER08/>).
- Harris, J. (2008b). Maximizing the Utility of MindModeling@Home Resources. In *Proceedings of The Eleventh World Conference on Integrated Design and Process Technology*.