

Cornell Virtual Workshop Capabilities Brief

The Cornell University Center for Advanced Computing (CAC) is a leader in the development and deployment of Web-based training programs designed to: (1) enhance the computational science skills of researchers, (2) broaden the participation of underrepresented groups in the sciences and engineering, and (3) accelerate the adoption of new and emerging technologies. The *Cornell Virtual Workshop*SM learning platform was launched in 1995. Since then, over 87,000 unique visitors have accessed Cornell Virtual Workshop training modules on high-performance computing (HPC) topics ranging from parallel computing to visualization.

Besides HPC topics, CAC develops Virtual Workshops for the humanities, social sciences, sciences or engineering on any topics of interest. The client provides access to a Subject Matter Expert (SME) and CAC does the rest – from instructional design to the deployment of the Virtual Workshop on CAC servers. Usage data may be auto-generated for inclusion in funding agency reports.

A Virtual Workshop is made up of multiple, logically-related modules, comprised of full text discussion, short audio or video clips, graphical simulations, examples, exercises, and quizzes. Virtual Workshop modules are always available on the Web as a 24x7 option for users who want to study a topic on demand and at their own pace.

The Cornell Virtual Workshop learning platform is a proven training technology. Cornell has received numerous grants from the National Science Foundation (NSF), the Department of Defense (DOD), and private industry to develop Virtual Workshops on a wide variety of topics. For example, under an NSF grant, CAC developed the *Ranger Virtual Workshop* to train educators and students on how to effectively use the Ranger supercomputer [1]. CAC was selected by Dynamic Research Corporation to develop the *User Productivity Enhancement, Technology Transfer, and Training (PETTT) Virtual Workshop* as an online training resource for the users of DOD Supercomputer Resource Centers (DSRC).

Today, CAC is developing and deploying online training for XSEDE, the NSF's advanced cyberinfrastructure and services program. Twenty-seven Cornell Virtual Workshop modules are currently available through the XSEDE User Portal [2]. Most recently, CAC staff worked with a team of educators and technologist to convert Jim Demmel's *Applications of Parallel Computers* course into the Cornell Virtual Workshop format in order to broaden access to the popular Computer Science course taught at Berkeley [3]. Due to the success of this project, the conversion of other Computer Science courses are planned for the future in order to further enhance XSEDE educational offerings and broaden participation.

CAC is also developing the *Stampede Virtual Workshop* [4] and recently released two new technology modules: *Many Integrated Core (MIC)* (see Figure 1) and *Vectorization in Modern CPUs and the New Intel Xeon Phi*.

Cornell Virtual Workshops are available at all times to the entire scientific community – researchers, HPC practitioners, students, and educators – and they require no travel budget. By leveraging the content and learning methodologies developed over a decade of Virtual Workshop design, Cornell can bring new instructional content online quickly to accelerate new technology adoption.

[1] Mehringer, S., Woody, N., Dolgert, A., Lantz, S. & Stanzione, D. (2011). Maximizing Computational Learning for Faculty and Student Scientists: The Ranger Virtual Workshop. *TeraGrid Conference Proceedings*. Retrieved from: <http://www.cac.cornell.edu/about/pubs/RangerVirtualWorkshop.pdf>.

[2] XSEDE User Portal: On Demand Training (n.d.). Retrieved from <https://portal.xsede.org/web/xup/online-training>.

[3] Cornell Virtual Workshop: Applications of Parallel Computers (2013). Retrieved from <http://www.cac.cornell.edu/VW/apc/>.

[4] Stampede Virtual Workshop: TACC User Portal (n.d.). Retrieved from <https://portal.tacc.utexas.edu/stampede-virtual-workshop>.

Home	Topics	Reference	Glossary	Help	Notebook
-------------	---------------	------------------	-----------------	-------------	-----------------

Welcome shm7

Cornell Virtual Workshop

<p>MIC</p> <hr/> <p>Introduction</p> <p>Goals</p> <p>Prerequisites</p> <hr/> <p>MIC Architecture</p> <ul style="list-style-type: none"> * Processor * Cache and Memory * I/O and Networking * File Systems <p>Programming Paradigms</p> <ul style="list-style-type: none"> * Distributed Memory * Shared Memory * Hybrid * Offload <p>Running Code on the MIC</p> <ul style="list-style-type: none"> * Native Execution (M) <ul style="list-style-type: none"> - Exercise: Interactive - Exercise: OpenMP * Symmetric Execution <ul style="list-style-type: none"> - Exercise: MPI - Exercise: Hybrid * Offload <p>Performance Considerations</p> <ul style="list-style-type: none"> * Vectorization * Threading & Affinity * Memory Access * Tuning * Balancing <hr/> <p>Quiz</p>	<h3>MIC Architecture</h3> <p>In this section, we provide an overview of the MIC architecture and describe several distinguishing features that are likely to be relevant to its performance when given different kinds of workloads.</p> <p>As mentioned in the introduction, MIC (pronounced 'mike') is an acronym for "Many Integrated Core". This is the architecture used by the Xeon Phi coprocessors in Stampede. In essence, each MIC coprocessor is a fully functional Linux-based system running on a PCIe card featuring a specialized CPU that achieves very high floating point throughput and low power consumption through lots of simple cores derived from the x86 architecture. These cores support a modified x86 instruction set, contain wide vector units, and feature hardware threading (SMT) with four threads per core.</p> <p>Surrounding the processor is a limited amount of very fast GDDR5 memory. Each coprocessor in Stampede has 8GB of system memory. Communication with the host machine occurs via the PCIe bus. Viewed from afar, the Xeon Phi coprocessor behaves like an entirely separate compute node that happens to be physically contained within a host node:</p> <div style="text-align: center;"> <p style="text-align: center;">Systems overview of a host and MIC coprocessor</p> </div> <p>In the next few sections we will describe more about the CPU architecture to understand how it is optimized for producing high floating point throughput, as well as its memory and I/O subsystems.</p>
--	--

Figure 1 – Stampede Virtual Workshop screenshot. The Many Integrated Cores (MIC) module features interactive exercises and quizzes focused on programming paradigms and running code on the MIC.