

 [Print this Page for Your Records](#)[Close Window](#)

Program#/Poster#: 102.27/AAA7

Title: The Merge framework: software and hardware tools to accelerate computational neuroscience

Location: San Diego Convention Center: Halls B-H

Presentation Start/End Time: Saturday, Nov 03, 2007, 3:00 PM - 4:00 PM

Authors: ***M. D. LINDERMAN**^{1,2}, J. COLLINS², H. WANG², T. H. MENG¹;
¹Elec Eng, Stanford Univ., Stanford, CA; ²Intel Corp., Santa Clara, CA

Driven by high-throughput experimental techniques and aggressive real-time requirements, such as in neural prosthetics, computational workloads in neuroscience research are growing almost without bound. The response, “computation is cheap”, is no answer. Only computers are really cheap, software is very expensive. Developing high performance software packages for neuroscience research is a costly, time-consuming process, requiring individuals skilled in both the research area and software engineering. As a result of the difficulty and cost, researchers often avoid computational tools or are forced onto inferior, but simpler, computational platforms. The Merge framework, a computational model and compiler, reduces the difficulty of software development by insulating the researcher from the implementation details of state-of-the-art heterogeneous parallel computing platforms. The Merge framework provides (1) a parallel-aware high-level library abstraction offering a structured interface between application and library developers and (2) a compiler and runtime which dynamically selects the best available library function for a task for a given input and machine configuration. The user assembles their program in a MATLAB-like manner from a rich and extensible library set (including algorithms like PCA, FFT, K-Means clustering, etc.). When specialized implementations exist for a given library function, such as for a GPGPU or FPGA, these will be invoked, offering potential to-the-metal performance for a wide range of computational architectures, all transparent to the user. The Merge framework has been prototyped on (1) a heterogeneous platform consisting of an Intel Core 2 Duo IA32 CPU and 8-core 32-thread Intel X3000 GPU, and (2) a homogeneous 32-processor Unisys SMP system with Intel Xeon processors. We implemented a set of common data analysis and simulation benchmarks, including a spike classification algorithm used in neural prosthetic systems, using the Merge library abstraction and enhanced the library set with X3000 specific implementations, achieving speedups of 2.1x - 22.2x using the X3000 GPU and 5.3x - 17.1x using the 32-way system relative to straight C on the IA32 CPU alone. All platforms use the same algorithm description with the runtime and compiler automatically selecting and invoking the appropriate libraries.

Disclosures: **M.D. Linderman**, Intel Corporation, A. Employment (full or part-time); **J. Collins**, Intel Corporation, A. Employment (full or part-time); **H. Wang**, Intel Corporation, A. Employment (full or part-time); **T.H. Meng**, None.

Support: FCRP C2S2

[Authors]. [Abstract Title]. Program No. XXX.XX. 2007 Neuroscience Meeting Planner. San Diego, CA: Society for Neuroscience, 2007. Online.

2007 Copyright by the Society for Neuroscience all rights reserved. Permission to republish any abstract or part of any abstract in any form must be obtained in writing by SfN office prior to publication.