

How to Navigate Next Generation Programming Models for Next Generation Computer Architecture

Thursday, February 25, 2010 9:40 AM (40 minutes)

Something strange has been happening in the slowly evolving, placid world of high performance computing. Software and hardware vendors have been introducing new programming models at a breakneck pace. At first blush, the proliferation of parallel programming models might seem confusing to software developers, but is it really surprising? In fact, programming models have been rapidly evolving for the better part of two decades, thanks in no small part to the boom in web-based application development frameworks and tools. The fundamental forces driving this are twofold: the first is the importance of domain-specific specialization and optimizations productively use modern hardware infrastructure and the second a base of software developers that is better able to adapt to these programming models and use “the right tool for the job”. The key question that remains which must be answered by vendors of these tools is: What tool is right for you?

Intel provides a broad set of programming tools and programming models that is a microcosm of the broader diversity available in the software ecosystem. I will discuss how these tools and programming models relate and interoperate with each other in a way that developers can use to navigate their choices. I will pay particular attention to our work in adding data parallelism in C++ via Intel’s Ct technology in ways that eliminate the traditional modularity tax associated with C++ frameworks. I will also show how this work has been applied in particle physics workloads.

Primary author: Dr GHULOUM, Anwar (Intel Corporation)

Presenter: Dr GHULOUM, Anwar (Intel Corporation)

Session Classification: Thursday, 25 February - Plenary Session

Track Classification: Computing Technology for Physics Research