Habanero Extreme Scale Software Research Project

The computer industry is at a major crossroads. Historically, single-core architectures have delivered sustained performance improvements at an exponential rate over multiple generations, and sequential programming languages and compilers have been able to pass these performance increases on to a wide range of applications. However, this is no longer the case due to current hardware trends and power efficiency limits. Instead of building processors with faster clock speeds, all computers--- embedded, mainstream, and high-end --- are being built using chips with an increasing number of processor cores, with little or no increase in clock speed per core. This trend poses a tremendous challenge for software enablement on future extreme scale systems as the number of cores per socket continues to grow, and the cores become more heterogeneous.
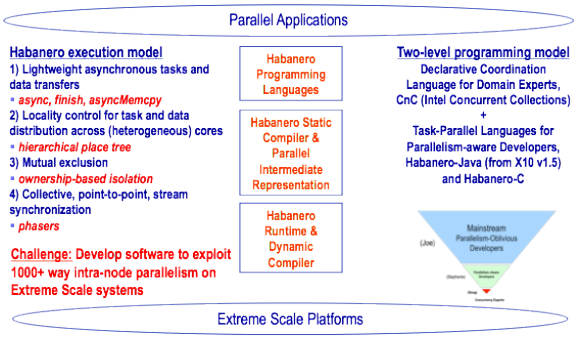
The Habanero project at Rice University was initiated in Fall 2007 to address the multicore software challenge by developing new programming technologies --- languages, compilers, runtime systems, and tools --- that support portable parallel abstractions for future hardware with high productivity and high performance. Our goal is to ensure that future software rewrites are done on software platforms that enable application developers to reuse their investment across multiple generations of homogeneous and heterogeneous extreme scale hardware. We also envision broader impact of this research including: updating the foundations of parallel software in introductory Computer Science courses, building an open source testbed to grow the ecosystem of researchers in the parallel software area, and using our research infrastructure as the basis for building reference implementations of future industry standards.