



The quest for scalability: Some tools and approaches
Brian Sumner, Principal Engineer
Performance Engineering
SGI

Coding an application to try to take advantage of concurrent execution is getting to be routine. It's also fairly routine to see some reduction in elapsed time when running these codes on multiple processors. However, it is still far from routine to achieve expected performance and scaling on large numbers of processors. In fact, it's almost always a challenge.

The parallel programmer will face numerous obstacles on the path to scalability (including the programmer himself/herself). In this presentation, I intend to discuss tools and approaches to detecting these obstacles, and techniques for limiting their impact.

While I primarily assume the runtime environment provided by large cache coherent NUMA systems, much of the discussion should be applicable to other parallel environments as well.