Department of Computer Science
University of Houston

SEMINAR SPRING 2011

WHEN: FRIDAY, MAY 13, 2011
WHERE: PGH 563
TIME: 11:00 AM

SPEAKER: Dr. Torsten Hoefler, University of Illinois at Urbana-Champaign

Host: Dr. Edgar Gabriel

Title: Characterizing the Influence of System Noise on Large-Scale Parallel Applications

Abstract:
System noise is increasingly a concern as HPC systems continue to grow in scale. Good operating systems can minimize noise, however, Some sources of asynchronous slowdowns, such as recoverable hardware error remain. Existing studies with artificial noise models provide only limited insight into application behavior under the influence of noise.
This paper presents an in-depth analysis of the impact of system noise on large-scale parallel application performance in realistic settings.

Our analytical model shows the particular circumstances under which noise is propagated or absorbed. The model shows that not only collective operations but also point-to-point communications influence the application's sensitivity to noise. We present a simulation toolchain that injects noise delays from traces gathered on four common large-scale architectures into a LogGPS simulation and allows new insights into the scaling of applications in noisy environments. Our simulation framework enables large-scale simulations up to 8 million processes with more than 1 million events per second. We investigate collective operations in noisy settings with up to 1 million processes and three applications (Sweep3D, AMG, and POP) with up to 32.000 processes. We show that the scale at which noise becomes a bottleneck is system-specific and depends on the structure of the noise. Simulations with different network speeds show that a 10x faster network does not improve application scalability because noise becomes a bottleneck at scale. We quantify this noise bottleneck and conclude that our tools can be utilized to tune the noise signatures of a specific system for minimal noise propagation. For example, our simulations verify the long-standing conjecture that co-scheduling prevents significant application slowdown.

Bio:
Torsten Hoefler is leading the performance modeling and simulation efforts of the Blue Waters project at NCSA/UIUC. He is also an active member of the MPI Forum where he leads the collective operations and topology working group. His main interests lie in performance modeling and simulation, parallel programming languages and runtimes, and large-scale networking. He received his Ph.D. from Indiana University in 2008. He also holds an appointment as Adjunct Assistant Professor of Computer Science at UIUC.