

Scale+Fidelity+Speed+Integration=Necessary+Possible in Big Network Simulations

Tuesday, September 11, 2012 11:26 AM (12 minutes)

We are in violent agreement with the vision and outlook of the COMBINE workshop. It is clear that effects such as feedback and rapidity of change are becoming so pronounced in global computer networks that modeling and simulation of big networks is both a timely need as well as a formidable challenge.

We are convinced that simulation is the third leg of network science, just as it has been accepted so in physical sciences recently. However, capture of scale, fidelity, integrated capture of complexity and similar challenges underlie the realization of simulation as a first-class scientific principle. In relation to this view, our team's expertise and interests have been in taking network modeling and simulation tools to the extreme along these important dimensions *simultaneously* – largest scale (multi-million node scenarios), highest fidelity (virtual machines, packet-level models, packet-fluid hybrids), and realistic/active behaviors (combined with other man-made phenomena, such as electric, vehicular, and other grids).

Our team's research has been driven by the understanding that (a) network effects at scale are very difficult to capture sans simulations, (b) feedback effects dominate to such a degree that fidelity in simulations is difficult to be abstracted, (c) the effects of interest when viewed in the overall system-level, at nation- or global-scale, are best captured by large-scale simulations and analyses alone.

Scale: We have dealt with the scale aspect in the past, by advancing techniques to sustain large-scale parallel simulations of network models on some of the largest supercomputing installations (MASCOTS'03). We executed multi-million node simulations (using pdns) on a Pittsburg Supercomputing Center machine (the largest of that time) as part of Georgia Tech team (with Fujimoto and Riley) on a DARPA NMS project in 2002-03. These runs still stand today as some of the largest packet-level network simulations to date.

Complexity: Our multi-million node computer network worm propagation simulations reported in ACSAC'04 still represent some of the largest, high-fidelity cyber security simulations with respect to scale. We are convinced that the effects of scale are best captured by simulations of this size, to uncover unforeseen effects or emergent phenomena.

Fidelity: We are currently uncovering the issues, and devising solutions, to increasing the model fidelity to the extreme, using virtual machines (VMs) as surrogates for either end-hosts or intermediate routers or both. While VMs have been employed in network simulations in recent times, a new challenge remains to be solved, namely, the incorporation of a first-class notion of virtual time into the VMs. Native VM schedulers are a gross mismatch to VM-based network simulations. We recently showed that using native schedulers (which are typically optimized for throughput) in fact can give wrong answers from simulations. With funding from Army Research Laboratory and others, we are developing new virtual time-ordered schedulers (PADS'11, MASCOTS'12) that are indispensable for large-scale high-fidelity network simulations in the future when executed on many-core host platforms of supercomputing-scale configurations.

Related: Our parallel discrete event simulator, musik, is being developed to suit the vision of very large-scale simulations in the future, and has now matured to the point of being the only simulation engine in the world to be tested on the largest extant supercomputing platforms (Jaguar, 216K+ cores), with multiple discrete event applications exercised as proofs of concept (epidemiology, radio signal, vehicular transportation, etc.)

We are highly interested in interacting with the network modeling and designing experts at the workshop and sharing ideas towards meeting the grand vision painted in the call, namely, a community analogous to that of the climate simulation community, for network science.

Primary author: Prof. PERUMALLA, Kalyan (Oak Ridge National Laboratory)

Presenter: Prof. PERUMALLA, Kalyan (Oak Ridge National Laboratory)