Parallel Simulation of Large PNNI Networks

March 24, 1999

Kalyan Perumalla

We have conducted sample simulation runs of TED models of PNNI networks. Here, we report the run-time performance of the parallel simulation system used for the simulations.

We have used three sample network configurations in the study: a 200 node network, a 600 node network and finally a 1000 node network. The 200 node network was provided by Iraj Sainee of Bellcore. The 600 and 1000 node networks were synthetically generated by Fang Hao of Georgia Tech using the Georgia Tech Internetwork Topology Model generation (GT-ITM) software. Using a few heuristics, we pre-processed all the three networks to automatically partition and generate hierarchical PNNI ATM addresses and other protocol-specific configuration parameters for the network nodes.

The networks were configured to include a single user node attached to each network node. Each user node included a client that initiates SVC connections and a server that accepts SVC connections. Two types of client behavior were used: *uniform* and *localized*. In networks with *uniform* clients, the clients pick a SVC destination uniformly at random from the list of *all* servers. In networks with *localized* clients, the clients choose destinations that are closer to them in the PNNI hierarchy more often than those that are farther in the hierarchy. In each simulation, a total of over 100,000 SVCs were established by all the clients combined. The number of failed SVC connection requests was negligible.

The simulations were performed on a SGI Origin shared-memory multiprocessor system at Georgia Tech, containing R10000 processors and 4 Giga bytes of memory. The speedup numbers are self-relative (i.e. with respect to the performance of the *parallel* simulator on 1 processor).

As Figure 1 demonstrates, we observe a speedup of 6.15 when the 1000-node network is simulated on 8 processors. Figure 4 demonstrates that approximately 1,000,000 SVC connection setups can be simulated in approximately 3 hours of wall-clock time for a large (1000-node) PNNI network, and in less than 0.5 hours for smaller (200-node) networks.

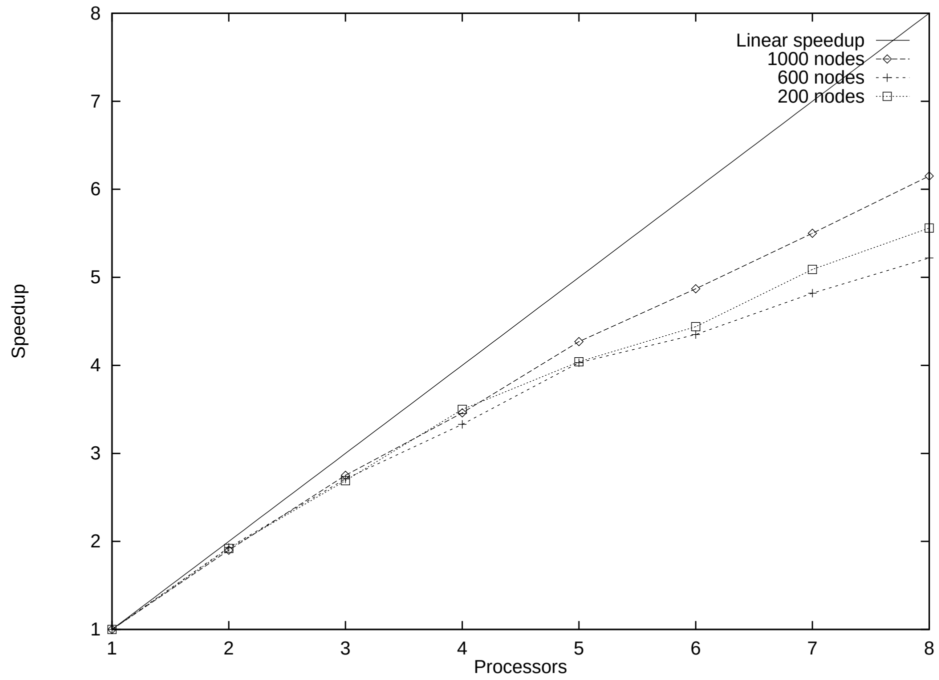

Figure 1: Simulation speedup for network configurations with *localized communication* clients

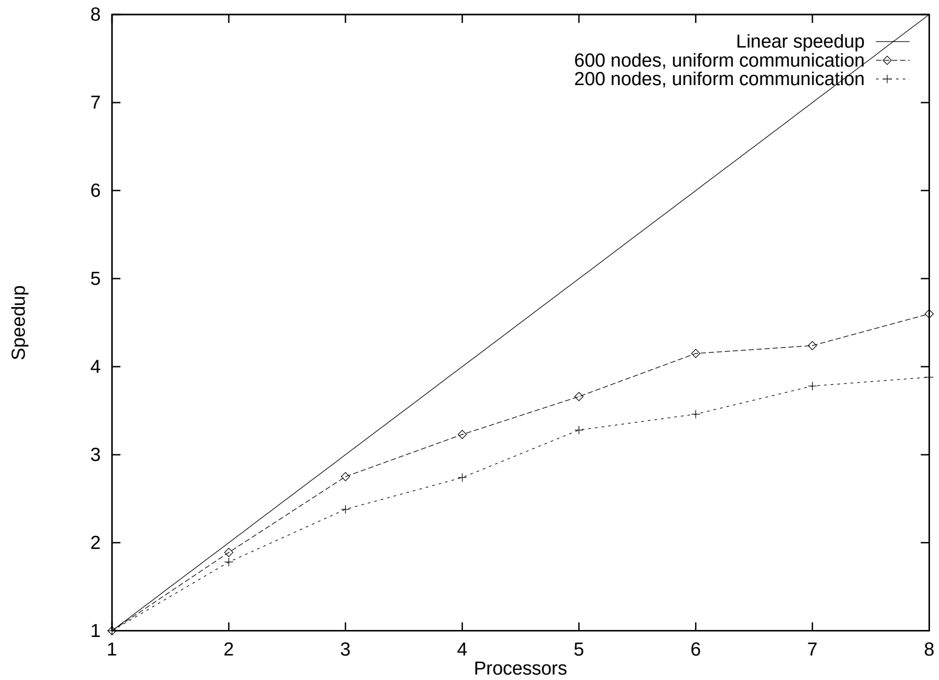

Figure 2: Simulation speedup for network configurations with *uniform communication* clients

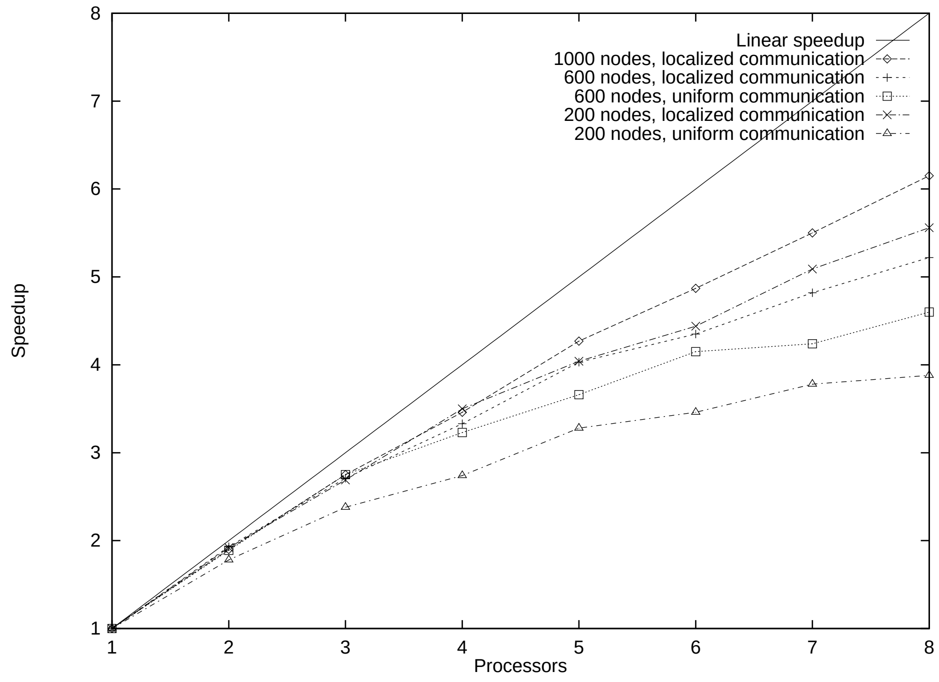

Figure 3: Simulation speedup for all network configurations

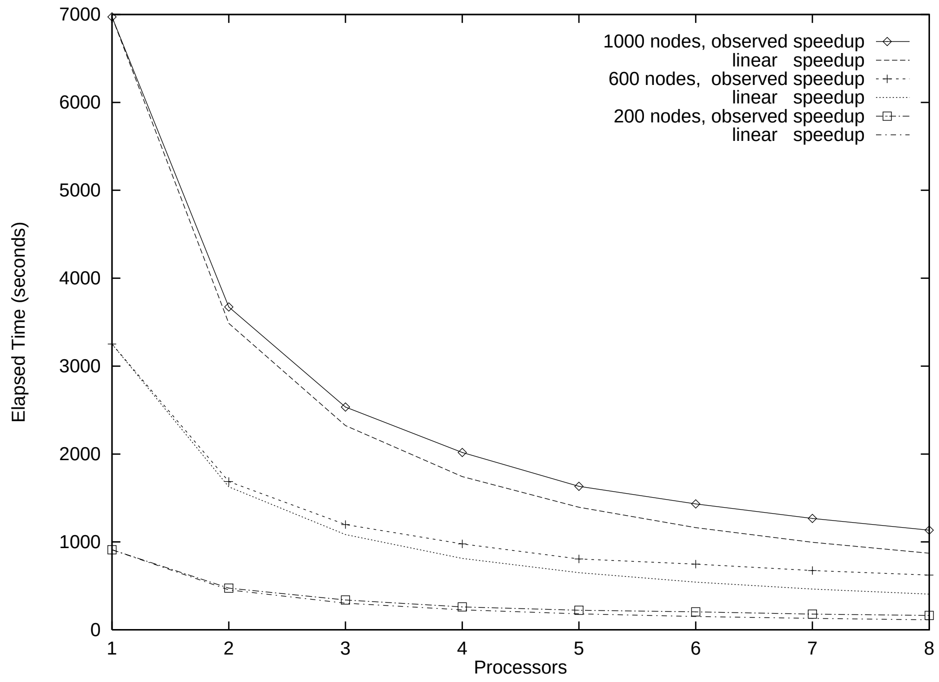

Figure 4: Elapsed time for network configurations with *localized communication* clients

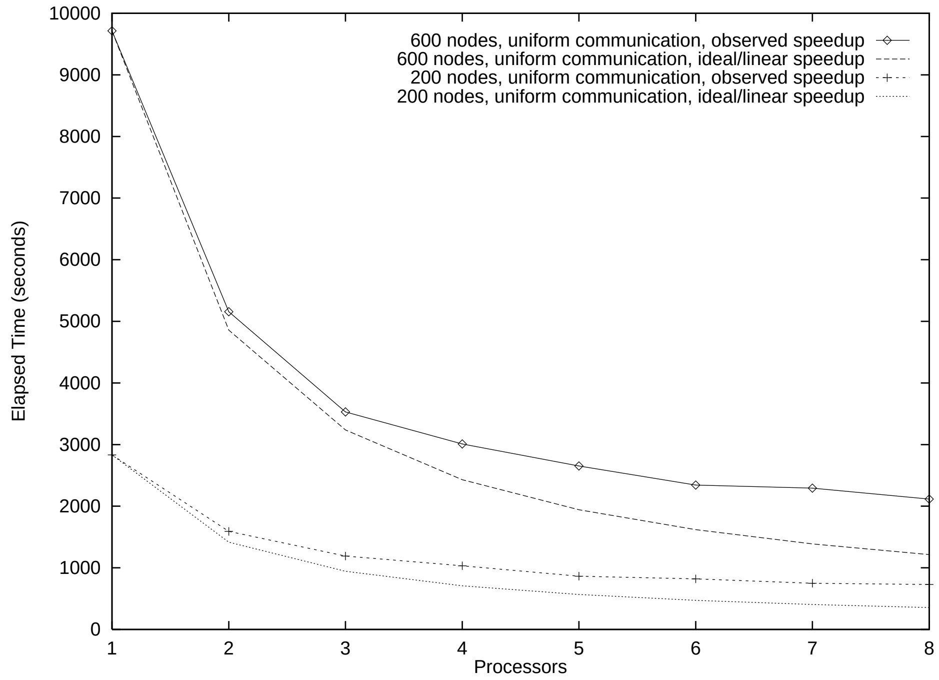

Figure 5: Elapsed time for network configurations with *uniform communication* clients

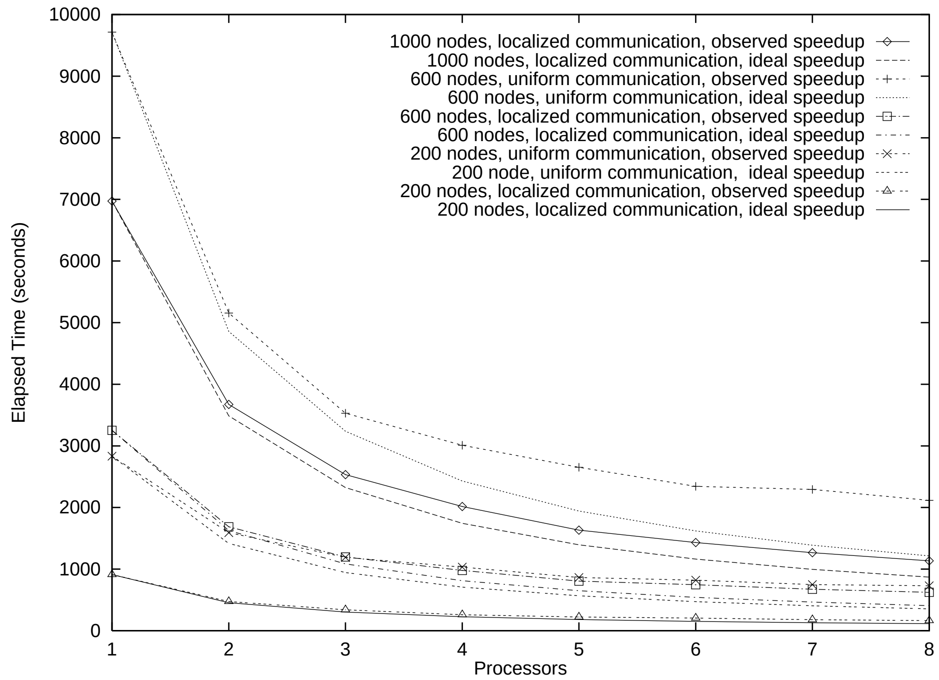

Figure 6: Elapsed time for all network configurations