Review of: A Case for NOW, Anderson, et al. IEEE Micro, February 1995.
Review by: Troy Ronda, September 15, 2005

The question addressed by this paper, in their own words is: "Must we buy a single computer big enough for the biggest task we ever need to run and pay a huge premium for the additional capacity?" The authors argue that resources on networked workstations can provide the additional capacity, at a much lower cost, and include interactive usage. The authors claim that NOW can provide virtual memory and a cache by using the aggregate DRAM across the network, scalable file systems, and parallel processing using the idle cycles on workstations. The paper discusses the best points of MPP (massively parallel processors) and show how they can be implemented in NOW. They balance the needs of interactive users with the needs of large applications on the NOW. I found the argument compelling for several reasons:

- I found the metrics to be applicable to the research. The price-performance metric for supercomputers improves at only 20-30% per year while the workstation improves at 80%. They also found that more than 60% of workstations on the network are available 100% of the time, even during daytime hours.
- Off-the-shelf components means their system will continue to work even after upgrading stations (a very common event, in my experience). They also like to use whatever resources are available on the network. They do this to obtain a "free" supercomputer and they are clear that it is more than parallel processing.
- The authors identified the issues with communications between off-the-shelf components. This is the deal-breaker for the system but they identified the goal and how they will make progress.
- In addition to good metrics and interesting simulation, real "wall clock" timing was obtained. I liked the GATOR timings.
- I had never thought of using NOW for cache purposes or as a large virtual memory. This is interesting; especially with the prevalence of P2P file sharing.

Although the paper is good, I found some parts weak:

- I inferred their cost-performance metric; I would like them to be more explicit.
- They did not show a connection between graduate student computer idle time and industry/home idle time. They also did not discuss slowdown to the network RAM and the co-operative file system if the workstation needs these resources.
- Showing good performance for one application (i.e. GATOR) does not necessarily prove the system will work for all supercomputer applications. It is still a compelling argument because the authors argue that GATOR is representative of a real-world, high performance, and parallel application.
- I would have liked to see some wall clock timings for the time required to migrate a job from one machine to another. Although they argue that the time is less than four seconds, it does not seem that the authors implemented the system or conducted a user survey. Theory and user complaints might be different.

I would like to explore if other authors have solved some of the issues in this paper for Internet distributed systems. For example, I would like to know how the SETI project deals with these issues. In any case, this paper was a good read.