

Supercomputer **VERSUS** Supercluster



Say you are performing biomolecular investigations that are extremely compute intensive. You have a finite amount of money and time. You could get (1) a supercomputer (fast custom CPUs and high-speed interconnect facilitating parallelization of a single computation) or (2) a big cluster (lots of cheap commodity CPUs with slower commodity networking). What approach will deliver the most science per resources consumed? Professor Vijay Pande overheard the two types of systems arguing the pros and cons and recorded it word-for-word here ...

Supercomputer: Supercomputers have reached a point where they are very powerful, are readily available via supercomputer centers, and can do it all. In terms of physics-based simulation, models have become sufficiently accurate that we can make useful, quantitative predictions.

Supercluster: Yes, no argument about the utility of simulations and the need for large-scale computer resources, but traditional supercomputers are expensive and the network can cost as much as the processors. For those of us with limited budgets, buying lots of computers to build a “super-cluster” with twice the computing power and a cheaper network seems more cost effective.

Supercomputer: Well, ok, but what would you do with all those processors without a fast network? Without the fast network, the processors are useless because they can’t work together. For example, in molecular simulation, one can only simulate about a nanosecond (10^{-9} second), maybe 10 ns, in a day on a single processor. That’s not going to get you far. Using a tightly coupled (i.e., fast network, supercomputer)

machine, one can use thousands of processors’ cores to get orders of magnitude longer simulations.

Supercluster: Yes, you can, but not very efficiently. Even these fast networks aren’t truly instantaneous so there can be a heavy overhead cost associated with processors communicating. For example, if one needs to

Supercomputer: “Supercomputers have reached a point where they are very powerful, are readily available via supercomputer centers, and can do it all.”

simulate a 10,000 atom system, a 10,000-core supercomputer is likely not going to be useful, since breaking the problem up that small won’t scale. In other words, the processors will spend too much time communicating and not enough time calculating.

Supercomputer: True, but scalability is a classic problem in computer science and the solution is a *faster* network, not a slower one. I don’t see your point.

Supercluster: Well, there is another way for many types of simulations interested in kinetic or thermodynamic properties, such as simulations on the molecular scale. Instead of using the processors to work together to

simulate a single long trajectory (i.e., a single simulation of a protein going through its dynamics), one could run many shorter simulations in parallel. Since molecular processes are inherently stochastic, one would be able to use multiple, *independent* simulations to get a performance boost *without* a fast network.

Supercomputer: Yes, I've heard of these tricks. The problem is that they have a very limited regime of applicability. These methods assume single exponential

Supercluster: True, for very, very large systems (which, for example, could not even be run on single processors due to memory or other constraints), scaling works reasonably well on multi-core boxes present in superclusters or on multiple cores in supercomputers, so I suggest a combination of methods: scale as far as one can go well (i.e., linearly, where doubling the number of processors doubles the speed), and then use additional methods on top of that. The MSM approach, for example, just seeks to use trajectories as efficiently as possible,

Supercluster: "For those of us with limited budgets, buying lots of computers to build a "super-cluster" with twice the computing power and a cheaper network seems more cost effective."

kinetics, i.e., that the probability of an event occurring after t nanoseconds of simulation time looks like $p(t) = k \exp^{-kt}$. This distribution is sharply peaked at short times, so one could see events even at times much shorter than the average time $\langle t \rangle = 1/k$. The problem here is that most complex systems have multiple states so this simple two-state approximation would break down.

Supercluster: A simple method like that would break down. However, over the last five years or so, several groups have been working on a much more sophisticated method, called Markov State Models (MSMs). Here, one combines many (relatively short) simulation trajectories with Bayesian statistics to build a kinetic model of the process of interest. One does not have to make approximations regarding single exponential kinetics of the overall system, assume reaction coordinates for the system, *a priori* identify initial and final states, etc. Recent advances in adaptive methods where one builds an MSM by gathering some simulation information, then adaptively decides where to run new simulations in order to optimize some property (such as minimizing the uncertainty in some variable of interest) have shown that this approach can be *more* efficient than a few long runs, *even* if there were no cost for traditional parallelization. This is because MSMs are much more efficient at skipping over traps and other places where simulations simply "wait" for some rare stochastic event.

Supercomputer: OK, but what about large systems that would take a lot of computing power to generate even short trajectories? Even if you use cool methods to combine short trajectories, you are bound to be stuck if the trajectories are too short. And this can be a problem for large systems, where generating even short trajectories can be a challenge on a single CPU core or even a multi-core CPU.

but there are other possible synergies as well. The generation of those trajectories still becomes an interesting challenge for the future.

Supercomputer: Isn't there another alternative? I think you're forgetting about GPUs. While traditional processors (i.e., CPUs) have not been getting faster (just with more CPU cores packed on a single chip), GPUs have been getting more and more powerful, in part due to their unique architecture with lots of floating point units for scientific calculations. Indeed, GPUs remind me of supercomputers of the past, like the old Cray vector supercomputers, which could do certain types of heavy floating point calculations quickly, with some coding effort to take advantage of this unique hardware. Indeed, this opens the door to smaller *personal* supercomputers like a GPU-accelerated under-your-desk minicluster. That can pack a lot of cycles into a small space and you don't have to share it with anyone. It's cheap and you can easily replace it with the next generation.

Supercluster: I hate to agree with you, but it's true: Both the supercomputer and megacluster have problems associated with unleveraged acquisition of capital equipment, real estate, cooling, and power. Maintenance, obsolescence, and other "big science" issues could be avoided by equipping each scientist with a smaller, personalized resource. Scientists can afford even small GPU clusters which can be quite powerful. Of course, the next question will be how one wants to use all of those GPUs in parallel! □

SEND US IDEAS Got your own opinions on this topic? Or have another topic you'd like to write about for these pages? Send us your thoughts on the Feedback page of our Web site: <http://biomedicalcomputationreview.org/feedback.html>.