

## Comments on Topics 4 and 6 on the NRC List

Klaus Bartschat  
Department of Physics and Astronomy  
Drake University  
Des Moines, IA 50311, USA  
klaus.bartschat@drake.edu

As shown in the title of this document, I believe that these two topics are closely related. We are certainly facing significant challenges already, and there would be serious consequences if the NSF were no longer able to provide state-of-the-art computing resources for its research community. In fact, chances are that the time has already come when NSF computing resources could be classified as insufficient.

Let me start by providing a brief background of my work and the relationship to NSF programs and computing facilities. My group's research is currently funded by the NSF through two programs, namely Theoretical Atomic, Molecular, and Optical Physics (TAMOP) at about \$90,000 per year and Computational Physics (CP) at about \$70,000 annually. These funds are mostly for personnel, travel, and small supply items, but nothing serious for computing. Our computing resources are provided through the XSEDE (formerly TeraGrid) program. I am the PI of one of the larger XSEDE projects with about 10,000,000 service units (SU) per year. [In NSF accounting, 1 SU is worth about \$0.05; hence this is by far the biggest grant of my group – even though my university never sees the actual dollars and hence administrators tend to ignore them.] One of the Co-PIs on this project is Barry Schneider, who was originally responsible for the transition and significant extension of TeraGrid into XSEDE. As you may imagine, we talk a lot about the computing situation in the US and worldwide, and some of my comments will likely reflect his input. It seems that Europe and Asia have clearly recognized the problem and are (re-)directing significant resources toward scientific computing.

Looking at the above description, one problem appears immediately. Namely: Even when the NSF funds research proposals through TAMOP and CP, and we indicate the need for substantial computational resources in the submission, there is no guarantee that we will actually get these resources. That decision comes after another serious, outside, competitive review, through which some, but certainly not all (my impression is that the number is about 30%, but I am not sure) of the requests for supercomputing time get funded. Furthermore, even if we are successful in obtaining resources for a while, a renewal proposal has to be written every year for computing time, and once again there is no guarantee that things will continue. If, indeed, the computing resources were to dry up, the situation could be compared to an experimental project where the personnel to run, for example, a big laser is approved, but equipment funds for the laser are not. Or, if the laser breaks after a year, there are no funds to repair it. In that case, ***one would simply not be able to deliver what one proposed.***

As we all know, resources have become scarce in general, not just at the NSF, but also at other federal agencies – and presumably in industry as well. [Of course, "for-profit organizations" are likely to take for free what they can get for free, i.e., to change the situation and secure support from such organizations may not be possible without

conveying a serious threat that their operations would be affected negatively unless they chip in.] Budgets are often flat at best, and with inflation (which certainly does exist on most salaries, even if the increases are small) the "real" value of the support is going down. Looking at our own recent experience, ***the situation is much worse on the computing front.*** To some extent, XSEDE has become a victim of its own success. As the word spreads how great the program is, more and more people would like to get access. At the same time, we would like to solve bigger and ever more challenging problems. ***So the demand in both the number and the size of the allocations increases!*** Even though new machines come on line every once in a while (e.g. Stampede, Comet), others go away (e.g. Kraken, maybe a little too early?). As a result, our actual award on XSEDE for the present year is less than what it was the year before. Despite the fact that our proposal was rated very high, the Allocations Committee simply had no choice but to cut everybody by at least 30% – obviously we are not the only group who has a worthy project.

Another issue to consider is funding across agencies. I believe that a significant part of the NSF-provided computing resources is going to researchers who are otherwise not supported by the NSF. These include grantees from the DoE, NASA, NIH, etc., as well as researchers from outside the US. In some cases (e.g. the DoE) there is reciprocity, but in other cases there does not seem to be. On the positive side, we have used cycles on NERSC machines and just recently applied for access to Titan at Oak Ridge – basically forced by the above 30% cut in our XSEDE request. Also, collaborations with groups from outside the US can be very beneficial when it comes to cross-fertilization and exchange of ideas, expertise, and software. So it is by no means a trivial problem with an obvious (at least not to me) solution.

One might, of course, say that simply crying for more money is unlikely to work well. Our group is interested in computing and presumably thinks that a few extra million dollars (ultimately a small part of the NSF and certainly the federal budget) will solve all our problems. Realistically speaking, there will be other groups who think exactly the same about the field of their passion. Also, if more resources are allocated, more people will try to use them. [As we all know, independent of the size of somebody's hard drive, it is always 95% full and hence seems too small.] So, once again, it is not obvious how to ***make a case for increased resources without appearing to be completely unrealistic.*** It will not be an easy task.

Then there is the question whether the review/allocation process can be improved sufficiently to free up extra resources or to use those that exist more effectively. My experience with the XSEDE program management has been very positive. As far as I can tell, they try everything possible to help out, for example, in transfer or supplement requests. As mentioned above, XSEDE has become so successful that it is, or will soon be, facing serious challenges to fulfill the likely further increasing expectations.

The strongest case one can presumably make for an increased, ***strategic investment*** (in reality more likely a reallocation of existing resources or a redirection of discretionary funds) towards computing is the little-disputed fact that ***scientific computing has become a major enabler for discoveries in nearly all fields of science and engineering.*** We cannot simply give up, but should continue to be aggressive in getting this message out. Maybe somebody will ultimately listen?