Lattice Field Theory Computations

**Steven Gottlieb (Indiana University), Anna Hasenfratz (University of Colorado),
Paul Mackenzie (Fermilab), and Robert Sugar (UCSB)**

Research needs/opportunities

As discussed in the companion paper, Lattice Field Theory for High Energy Physics, the most pressing challenge in high energy physics now is testing the standard model, and searching for evidence of new physics more fundamental than the standard model ("beyond the standard model" physics). This involves determining as accurately as possible the parameters of the standard model (the masses of the quarks, the strengths of the forces, etc.), and searching for evidence of discrepancies in the relationships between them that are predicted by the standard model. Lattice gauge theory calculations are required to extract many of these parameters from experiment. Many of these calculations have been done decently in existing work, and this has contributed greatly to the value the related experiments. However, in many cases the existing experimental data is much more precise than the theoretical calculations needed to interpret it, and will produce much more precise tests of the standard model when the lattice calculations have been improved to the precision of the experiments.

Although the US lattice gauge theory community used nearly two billion core-hours in 2014 and comparable efforts exist in Europe and Japan, much greater precision is still required. Calculations have been mapped out that could profitably use an order of magnitude more cycles than are currently available. Detailed road maps exist of future data sets that will help in producing the precision required by completed experiments. The value of these very expensive experiments will be much enhanced when the required calculations can be completed.

Advanced computing capabilities, facilities, requirements

Lattice field theory calculations make heavy use of almost all types of computing hardware: laptops and desktops, research group clusters, department or campus commodity cluster systems, mid- to large-scale, shared capacity systems such as those provided by XSEDE, and leadership-class capability systems such as Blue Waters and Mira. The only significant kind of computing that is little used in lattice calculations is cloud computing. (Cloud computing systems such as the Amazon Cloud generally do not have fast enough interconnects between nodes to make multi-node calculations practical, and they tend not to be cost-effective compared with purchasing dedicated hardware. A study was done at NERSC some time ago comparing performance and cost-effectiveness of NERSC and cloud services, with NERSC coming out ahead. However, cloud systems are fast evolving and there are now some HPC cloud providers such as Penguin Computing on Demand (POD) that features a QDR inter-node interconnect. Unless volume discounts are very steep, it would still appear that this is not cost-effective compared to XSEDE. Our recent XSEDE allocation was 8.56 million SU and 540 TB of tape storage. This was valued at \$307K. However, POD price is \$0.12/core/hr and \$2.32/GPU/hr. The computer time alone would cost over \$1 million if purchased from Penguin.) Parts of the lattice research program require the highest capability computing, but there are even larger needs for flops from capacity computing which may be of smaller, but still substantial, capability. Roughly speaking, we have similar needs for cycles at all scales: one-node, ten-node, hundred-node, and thousand-node scales. Therefore, access to both leadership-class resources and high-capacity resources is essential. Research group and campus clusters are an important supplement to lattice hardware and some lattice groups make heavy use of them.

Lattice data sets in the US are quite large, of order ten petabytes in aggregate They consist in large part of moderate numbers of very large files. We rely on a few large depositories to store the main parts, and relatively simple tools are used to wrangle the data.

Since lattice field theory calculations depend on access to the largest computing hardware possible, the new GPU- and Intel-based systems are critical to further progress. Major efforts in lattice field theory are going into preparing for them, for example in the NERSC NESAP program and, in the future, through the Summit CAAR program. Our calculations involve intensive multiplication of small matrices, and are based on regular grids. They are therefore highly vectorizable, and so in principle are can be adapted to the coming machines. However, each new architecture requires its own data layout, which is a real challenge. Highly skilled programmers devote a lot of effort into porting our most important codes to new machines. It would be very valuable to have a single code base that has high performance across several architectures.

Bandwidth to memory has long been a limiting factor to the performance of our codes, so improvements to memory are of critical importance for us. Use of advanced computing hardware is enabled by USQCD's extensive base of publicly available codes and libraries. This code is developed, maintained and ported to new hardware by the US lattice gauge theory community, partly supported by the DOE SciDAC program.

Challenges and suggestions

The most important component limiting scientific progress in our computations is the number and productivity of the scientific staff who design and carry out the computational programs of lattice gauge theory. The next most important component is the number and type of cycles available. The precision of our calculations is severely limited by the number of cycles available to lattice gauge theory in the US. We could immediately make good use of an order of magnitude more cycles than we currently have access to. Furthermore, the composition of the cycles is critical. As previously stated, some of our calculations can only be done on the largest leadership-class resources, like Blue Waters. An even larger fraction of cycles can be done well on medium scale resources such as those provided by the XSEDE program. These include some of the analysis done on the large sets of gauge configurations generated on the leadership-class resources, and a diverse set of smaller calculations.

After a long period of stability in computer architecture, we are entering a new era in which the basic node architecture is variable and two or three levels of parallel coding (such as message passing, threading, and vectorization) are required. The NSF may not have the resources to change the development of commercial machines, so a valuable approach would be to understand the important codes and algorithms in order to guide scientists to the best architectures for their problems. It will require the work of computer scientists, programmers and application scientists to do the detailed performance analysis needed to pick the best computers, or mix of computers, for the NSF. However, the potential payoffs of this work are enormous.

Another challenge for the lattice community is dealing with instability in computer centers, for example in mass storage and staffing. Whenever a mass storage system is shut down and users are forced to migrate data, a large amount of of precious scientist effort is required to migrate the data that could be better spent on improving and carrying out the calculations. When resources are moved from one center to another, relationships between scientific users and computing professionals are disrupted and new relationships must be formed with support personnel who must become familiar with new users, codes and work flows. In the past, the benefits to the users of competition and turnover of centers have been outweighed by the costs of unproductive migration to new equipment and establishment of relationships with new personnel.

Other components of the lattice computational program are also very important. USQCD has a SciDAC-supported software program which involves 20–30 people and is critical to the success of our computations. A similarly-sized effort is devoted to algorithmic development. Algorithmic improvements have produced an even greater improvement in the speed of our calculations than the enormous improvement in hardware over the last thirty years. Since our most important limiting factor is the size and quality of scientific staff, training of young scientists is also of crucial importance. The US lattice gauge theory community devotes great effort to these areas and continued effort is always needed.

For lattice gauge theory, the most important new investments that could be made to benefit the lattice community would be greater numbers of cycles on both leadership-class and capacity resources. The most important policy improvement NSF could make would be to create better stability at its computing centers.

Until recently, there was an Office of Cyberinfrastructure (OCI) reporting directly to the Director of the NSF. This seemed like a good system for top-down coordination of NSF's overall strategy. The head of OCI was in a good position to query the heads of directorates as to their computing needs and to present those needs at the highest level. As far as bottom up efforts go, service providers like NERSC and ESnet have a history of regular requirements workshops for their user communities. Perhaps the NSF would benefit from such workshops, or activities could be added to the annual XSEDE meeting. Finally, maybe it is time for an NSF Supercomputer Users Group.