Lattice Field Theory for High Energy Physics

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

The long term objective of high energy physicists is to identify the fundamental building blocks of matter and to determine the forces among them that lead to the physical world we observe. Major progress has been made towards this very ambitious goal through the development of the standard model of high energy physics, which encompasses our current understanding of the fundamental interactions of subatomic physics. It consists of two quantum field theories: the Weinberg-Salam theory of electromagnetic and weak interactions, and QCD, the theory of the strong interactions. The standard model has been enormously successful in explaining a wealth of data produced in accelerator and cosmic ray experiments over the past thirty-five years; however, our knowledge of it is incomplete because it has been difficult to extract many of the most interesting predictions of QCD, those that depend on the strong coupling regime of the theory, and therefore require non-perturbative calculations. At present the only means of carrying out non-perturbative QCD calculations from first principles and with controlled errors is through large scale numerical simulations within the framework of lattice gauge theory. These simulations are needed to obtain a quantitative understanding of the physical phenomena controlled by the strong interactions, to determine a number of the fundamental parameters of the standard model, and to make precise tests of the standard model. Despite the many successes of the standard model, it is believed by high energy physicists that to understand physics at the shortest distances, or highest energies, a more general theory will be required. The standard model is expected to be a limiting case of this more general theory. A central objective of the experimental program in high energy physics, and of lattice gauge theory simulations, is to determine the range of validity of the standard model, and to search for physical phenomena that will require new theoretical ideas for their understanding. Thus, lattice gauge theory simulations play an important role in efforts to obtain a deeper understanding of the fundamental laws of physics.

Numerical studies of QCD and of theories that have been proposed for generalizations of the standard model support major portions of the world-wide experimental effort in high energy physics. The United States alone spends approximately $800 million per year on experimental high energy physics, and the European Organization for Nuclear Research (CERN) spends over $1 billion per year on the operation of the Large Hadron Collider (LHC) and its detectors. Japan also has a substantial experimental high energy physics program, and China has a growing one. It is no coincidence that all of these countries provide substantial support for lattice gauge theory studies. Furthermore, numerical studies of QCD provide critical support for nuclear physics experiments. The United States spends approximately $500 million for its experimental program in this area. Thus, although lattice gauge theory calculations require large computational resources, they support large and expensive experimental programs that address fundamental problems in physics.

QCD and beyond-the-standard-model (BSM) theories are formulated in the four-dimensional space-time continuum; however, in order to carry out numerical calculations one must reformulate these theories on a four-dimensional grid or lattice. Physical results are obtained by performing calculations for several small values of the lattice spacing, and extrapolating results to the continuum (zero lattice-spacing) limit. Clearly, the precision of calculations increase as the lattice spacings at which they are performed decrease, but the required computer resources increase rapidly with decreasing lattice spacings as well.

Lattice QCD and BSM calculations proceed in two steps. In the first, one uses importance sampling techniques to generate gauge configurations, which are representative samples from the Feynman path integrals that define the theories. These configurations are saved, and in the second step they are used to calculate a wide variety of physical quantities. (A set of configurations with the same lattice spacing, lattice dimensions and quark masses is referred to as an ensemble. A typical ensemble has of order 1,000 configurations.) The generation of gauge configurations presents a capability challenge, and is the rate limiting step in the work. Each configuration evolves from the one before it, so the generation of an ensemble must be performed in a single stream, or a small number of them. In order to make progress on state-of-the-art ensembles at a reasonable rate, the most capable available supercomputers are required. Access to the new generation of supercomputers expected to become available over the next few years would have a major impact on the field, enabling very significant improvements in the precision of ongoing calculations, and providing the ability to address important new problems that are not feasible with present computers. Access to the exascale computers envisioned for the more distant future would have an even greater impact on this field.

The measurement of physical quantities on stored gauge configuration ensembles presents a capacity challenge. The total number of floating point operations used in a measurement campaign is typically larger than the number needed to generate the configurations. However, configurations can be analyzed independently, and therefore in parallel, either on the same machine or on different ones. A large campaign may involve ten to fifteen ensembles, each containing of order a thousand configurations. So, there are workflow challenges that must be met, as well as a need for significant archival storage for the gauge configurations. The largest configurations currently being worked with require approximately 250 gigabytes of storage each. Since ensembles are shared among groups, and are used for a wide range of applications, stable, easily accessible archival storage is of great importance. Ensembles are typically useful for five to ten years.

Lattice gauge theory calculations are well suited for massively parallel computers. The bulk of the computer resources for both configuration generation and measurements go into the inversion of a large, sparse matrix, the lattice Dirac operator, which can be effected by Krylov space techniques, multigrid or a number of other approaches. The Department of Energy has supported the development of code and algorithms for a number of years through its SciDAC Program. This support has enabled the development of community codes, and their optimization for specific machines. For example, optimization for computers with GPU and Xeon Phi accelerators is in progress. Lattice gauge theory codes typically achieve outstanding performance on parallel computers.

In 2013, USQCD, which consists of nearly all of the high energy and nuclear physicists engaged in the numerical study of lattice field theories, wrote four white papers setting out the scientific opportunities in this area, and estimating the computing resources needed to make significant progress in the next few years. These white papers can be found at the URL `http://www.usqcd.org/collaboration.html#2013White papers`. More recently, prospects for lattice field theory in high energy physics were examined as part of the Snowmass study of priorities in high energy physics as a whole. The report can be found at the URL `http://arxiv.org/pdf/1310.6087.pdf`. The upshot of this report and of the white papers is that there are many calculations of importance to the experimental programs in high energy and nuclear physics that can be carried out with computers that sustain petaflop/s, and even more with ones that sustain exaflop/s. Figure 1 below provides an example of progress that has been made in recent years with the current generation of supercomputers. In both cases, the results are of a precision of a fraction of 1%.

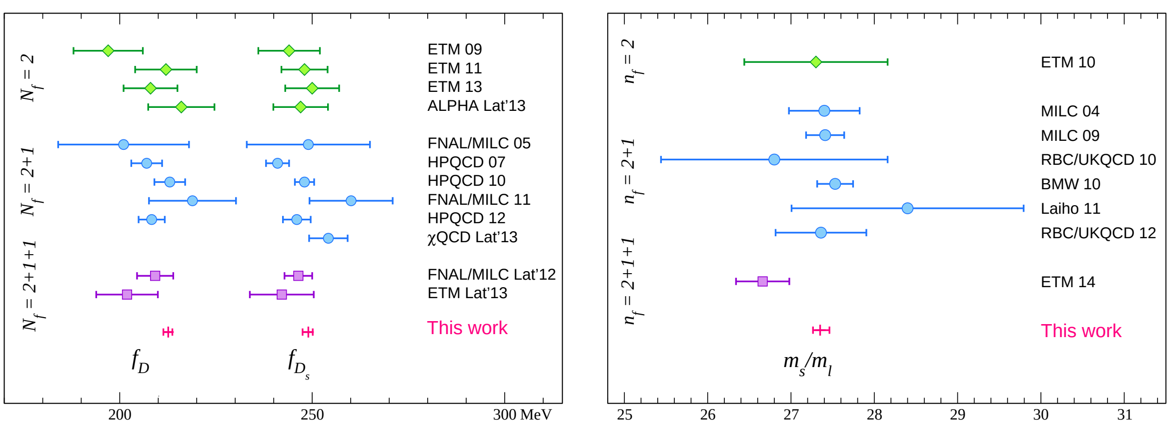

Figure 1: Comparison of recent results for the leptonic decay constants f_D and f_{D_s}, which play an important role in test of the standard model [left panel], and of the ratio of the masses of the strange and light quarks, which are fundamental parameters of the standard model [right panel]. Results are grouped by the number of flavors of sea quarks from top to bottom: nf=2 (green diamonds), nf = 2+1 (blue circles), and nf=2+1+1 (purple squares). Within each grouping, the results are in chronological order. The latest results appear at the bottom of each figure and are labeled *This Work*. These figures are from the Fermilab Lattice and MILC Collaborations, Phys. Rev. D**90**, 074509 (2014) [arXiv:1407.3772].