A Vision of Computing in 10+ Years

Steven R. Brandt
sbrandt@cct.lsu.edu
Center for Computation and Technology &
Department of Computer Science
Louisiana State University
Baton Rouge, Louisiana 70803, USA

Frank Löffler
knarf@cct.lsu.edu
Center for Computation and Technology
Louisiana State University
Baton Rouge, Louisiana 70803, USA

Once upon a time, single scientists could essentially know everything that was known in science, could understand the concepts of their time. Their expertise could stretch from physics, chemistry and biology to medicine, art, and even fields of study we would not consider science today. This picture has drastically changed with the increasing amount of knowledge in each of these fields, separating scientists in different fields, often even linguistically.

In addition, especially in experimental fields, performing science became so complex that single teams or even small collaborations of teams are no longer sufficient. Prime examples of well-known teams are the LHC, or the Ligo/Virgo collaboration, with specialization happening even within a single, well-defined field of science. This transition will only accelerate during the next 10+ years.

One of the main drivers of this transition is an increase in problem complexity beyond what a single individual or team can handle. This is increasingly the case in scientific computing, and here we include both two "extremes": high-performance computing (HPC), and single (or few) core applications that are used by a much wider scientific audience. HPC is leading in the increase in complexity due to hardware changes, as it always incorporates new hardware developments first, but small-scale computing will be eventually be affected also. Moreover, software paradigms change. Nobody can answer with certainty that computers in 10 years will look like, or which programming languages or concepts will be used. It is not even certain what the future of current concepts (OpenCL, CUDA, OpenACC, ...) will look like in five years; the lifetime of a cluster.

However, hardware and implementation complexity are not the only drivers of this transition. The increase in possibilities offered by modern computing infrastructure brings with it an increase of model complexity, to the extent that experts in different science fields work on different parts of a single computational model and its implementation, often without being experts in collaborative software development, or formal training in software development. Despite the obvious shortcomings of such an approach, successful examples of such teams are not difficult to find, which underlines the importance of efforts understanding the reasons. We expect that essentially every scientist in the very near future will need to have a certain amount of skills also outside of their base discipline, specifically within domain science, computing, algorithms, and hardware, as well as collaborative software development. That burden needs to be minimized, both by reducing the amount of interdisciplinary knowledge a scientist needs to acquire and in the effort needed to acquire it.

We do not have an answer, and we believe that no general answer exists. However, some approaches have potential to be helpful:

1) We find a way to separate the specification of (a) science simulation and computational problems (b) numerical methods and (c) low-level optimizations

Doing this would create a library of physics problems which can be implemented by independent groups of computer scientists and optimized by specialists in hardware. It would enable better partnerships with industry and academia, and give vendors better targets for new hardware and software designs.

2) We standardize on the basic infrastructure of computing, i.e. parameter files, compilation and configuration tools, performing high performance I/O.

To the extent that this can be achieved, it will lower the barrier to sharing and understanding codes and reduction of effort.

3) We find a better way to recognize and reward the people making the biggest contributions to science. It will be essential to find a way to distribute rewards, including recognition, within teams that not only span hundreds or thousands of people, but also science fields, continents and cultural barriers. Some teams try to do this using author lists that include hundreds of scientists, but this in the end only shows the gross inability of the current reward system to deal with modern science. Unfortunately, these rewards are tightly coupled to the carriers of the scientists through narrow-minded ways of comparing the "scientific value" of, for example, faculty candidates or promotion within departments. The earlier these problems can at least begun to be attacked, even just by spreading discussions about them, the higher will be the pay-off by keeping talented individuals within academia that currently just don't quite

fit the reward structure.

4) We need to find a way to make the ever-increasing complexity of science better understandable. While here we mainly aim to be able to effectively communicate between experts of different fields, it should also have a beneficial effect on the image of science by the general population.

Currently, attempting to read a scientific paper that is even slightly outside one's area of specialization is nearly impossible. This need not be the case, and we hypothesize that in ten years we find a better way to integrate our corpus of scientific knowledge.

Scientific terms, methods, and notations could all be hyperlinked to standardized explanatory texts with technology no more advanced than the hyperlink. Done correctly, all scientific knowledge could be crafted into a single tree making it possible for one to start with any given scientific publication and to systematically learn all the concepts needed to comprehend the paper.

To summarize, we expect scientific computing in 10+ years to look quite different from today. There will be more and more specialization within teams, with team sizes increasing. In done right, we will also see better, more appropriate scientific reward systems, as well as more emphasis on better "documentation of science", even for scientists.