From *SIAM News*, Volume 40, Number 8, October 2007

Computing, Business, and Applied Mathematics: The Next Challenges

By William R. Pulleyblank

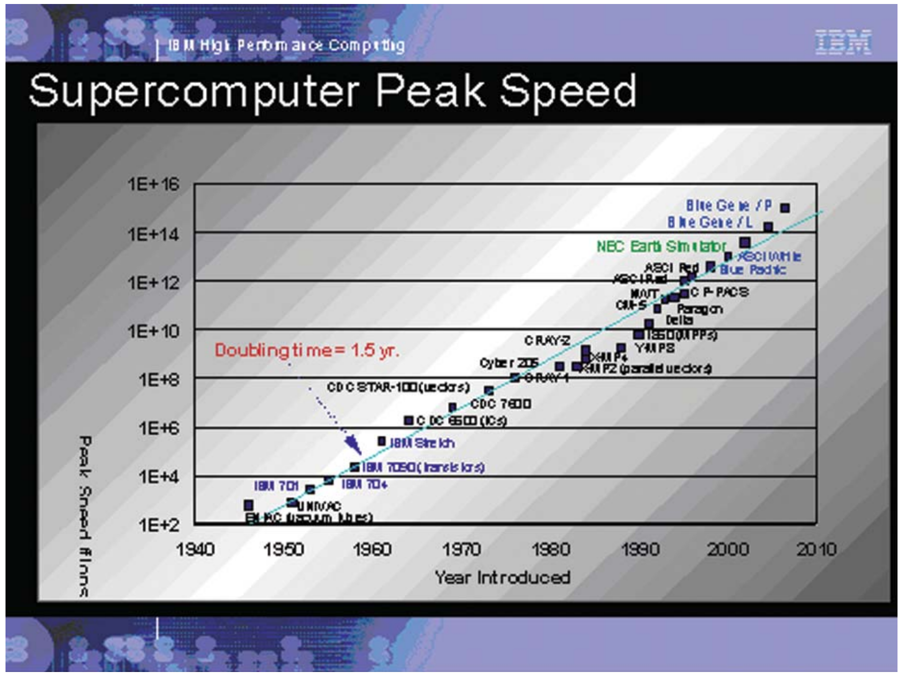

It has been six decades since ENIAC, the first large-scale programmable digital computer, was put into operation. During that period we have seen steady exponential growth in the power of computers and similar growth in the range and variety of their applications. The accompanying chart shows the consistent improvement in the performance of the world's most powerful computers since 1946. Each line on the vertical scale represents an increase of two orders of magnitude in the peak performance of these computers, measured in floating-point operations per second.

The slope of the regression line represents the increase in performance commonly referred to as "Moore's law"—for Gordon Moore's empirical observation in 1965 that the density of transistors on computer chips was doubling every 1.5 years. Subsequent improvements in computer performance and price/performance ratios have followed this same improvement trajectory.

Some of the most rapid growth we are now seeing in applications is in the fields of biology and medicine. IBM's BlueGene/L supercomputer, installed at the Lawrence Livermore National Laboratory, was officially rated the world's fastest supercomputer in November 2004, with a sustained performance of 280.6 trillion (10^{12}) floating-point operations per second on the Linpack benchmark. (See www.top500.org.) The principal applications on the second largest BlueGene/L system, installed at IBM's T.J. Watson Research Center, are in the area of protein science—in particular, simulating the interactions of multiple proteins at an atomic level.

Simulation at this level requires both advanced algorithmic design and huge amounts of computation. If we implemented a molecular dynamics simulation of protein folding in an "obvious" way, the number of calculations required would grow as the square of the number of atoms. On a hypothetical supercomputer approximately four times more powerful than the Law-rence Livermore BlueGene/L, it would take one year of uninterrupted computing to simulate one millisecond of folding time. With careful attention to the application design, however, it becomes possible to run the simulation in time that grows as $n \log(n)$, where n is the number of atoms in the protein and the surrounding liquid. This reduces the computation time significantly and makes the simulation practical [1].

Problems that arise in business are also creating huge demands on computational resources and on the mathematical methods used to solve them. This is in part a result of the strong movement toward globalization. Today's international companies need to deal in an integrated fashion with a multitude of external forces and very large and complex problems as they try to optimize their profitability.

Another significant factor in today's business world is the increasing domination of national economies by the services component. It is forecast that by 2010 the services component will be larger than the more traditional agricultural and industrial components combined. Solution of the resulting range of new problems will probably require new mathematical methods.

We identify five basic challenges:

1. *We must deal with massive amounts of noisy data, both repository and streaming.*

Unprecedented quantities of data are being captured and must be incorporated into decision-making processes. Part of this is traditional business data that is being captured and stored in repositories; another part consists of new forms of data, such as Internet transactions and user "click" histories. But in terms of quantity, both types are rapidly being swamped by machine-generated data, such as data obtained by sensor networks, video surveillance, and medical imaging.

Some machine-generated data is stored in conventional databases. However, we are also seeing a desire to process streaming data, such as stock market ticker data, in near real time so that it can be integrated into decision-making. By its nature, this data will be noisy and will contain errors; there may not be enough time or resources to remove the errors before the data must be used.

2. *Businesses of all types need effective methods for reducing risk and managing uncertainty.*

Many decisions must be made in the presence of uncertainty and are affected by external factors that are often difficult, or impossible, to predict. Examples range from production planning in the face of uncertain demand, to setting prices for products without knowing the competition's strategies, to planning res-ponses to massive natural disasters like hurricanes and floods.

3. *We need new methods that can utilize distributed computation and data.*

Data is generated and stored in a variety of forms in a variety of locations. For many applications on modern parallel computers, runtimes are dominated by the time required to move the data rather than by the actual computation time. The emergence of grid computing has taken this to another scale. The decomposition of computations becomes a challenge, with pieces of a computation to be performed "close" to the relevant data, after which the results are combined.

A benefit of a distributed setting is a much higher level of reliability—if some part of a network or enterprise disappeared, temporarily or permanently, it might still be possible to continue the necessary calculations on the remaining parts.

4. *We need to develop methods for solving operations problems, in addition to those of strategy and planning.*

Historically, most methods in operations research and mathematical programming were developed for planning problems. Computer systems have been used for more than two decades to build optimal crew schedules for airlines, for example; only recently have they been used to manage recovery after disruptions to a schedule. In part this is because significant amounts of computational power are required if resolution of these problems is to be completed within the permitted time. In addition, it is often unclear how these problems should be modeled mathematically so that the results will yield acceptable solutions to the real problems.

5. *Businesses are becoming dynamic networks of components, often created for the duration of an event, rather than the monolithic structures of the last century.*

Such businesses are forced to focus on the interfaces between systems, where the output from one process becomes the input for another. The need to combine processes, often at completely different scales, arises in many areas.

We would like, for example, to combine atomic-level models of proteins to form cell models, which could then be combined to model a human's physiology. An example in the business world is the complex coordination of a variety of functions—insurance, police, disaster recovery, emergency supplies, transportation, and communication—in the event of a natural disaster.

In summary, many of the most fundamental advances in applied mathematics result as new problems or opportunities arise and attract the attention of the mathematicians, statisticians, computer scientists, and engineers who can solve these problems. We may be entering such a period. The extraordinary growth in computational capability of our computer systems provides an opportunity. The rapidly evolving needs in business optimization provide a new set of challenges that complement the continuing challenges presented in more traditional areas.

References

[1] B.G. Fitch, A. Rayshubskiy, M. Eleftheriou, T.C. Ward, M. Giampapa, M.C. Pitman, and R.S. Germain, *Molecular dynamics—Blue matter: Approaching the limits of concurrency for classical molecular dynamics*, SC '06: Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, ACM Press, 2006.

William Pulleyblank is vice-president of the Center for Business Optimization, IBM Global Services.