

COMMENT

Why Experimentalists Mistrust Computer Modeling—and Why They Rely on It

I designed my first simulation more than 30 years ago when I was a graduate student. The goal was to model the trajectories of metastable mercury atoms as they moved through inhomogeneous electric and magnetic fields in an atomic-beam machine. The simulation included the effects of hyperfine structure, the velocity and angular distributions of the atomic beam, and the effect of the residual background gas. It required more than 240 hours of computer time on what was the largest and fastest machine on the campus at that time: a room-sized Control Data 1604.

Since that time, I have worked with my students on simulating the effect of the Earth tides on the magma chamber beneath Yellowstone National Park; on modeling the dynamics of a seismic zone in southern California; and on studying the effect of multipath reflections on measurements made using signals from the Global Positioning System satellites. My most recent simulation was designed to investigate the question, "What is the best way of using data from a primary frequency standard?" and resulted in a new algorithm called AF1 whose formal description is now in press. Simulations are particularly powerful in this case, because they make it possible to observe the performance of a statistical estimation procedure using input data whose noise spectrum is known exactly.

The simulations played a central role in each of these studies. When we were wise enough to use them early in the project, they allowed us to study many different approaches and experimental configurations and to choose the one with the most promise. They provided insight into how the experiment should be performed and what accuracy we might

expect to achieve. At other times—when we chose to measure first and simulate afterwards—the results of the simulation either amplified our self-esteem or, much more frequently, provided insight into why our experiment was not working and what we should have done in the first place.

Much has changed since the time of my first simulation, but the relationship between simulating an experiment and performing the measurements is about the same as it was when I started. Apart from a mastery of the mechanical issues of programming, writing an effective simulation requires a deep understanding of the physics of the problem. The researcher needs to identify the principal effects and know how important each one is and also to determine which constants or parameters govern a process and how well they are known. This understanding can only come in a field that has reached a certain level of maturity.

We can be reasonably confident about our simulation of the performance of a primary frequency standard based on conventional atomic-beam techniques, for example, because we have built a number of standards of this type and have used them for many years. On the other hand, I would have much less confidence in a simulation that was used to predict the performance of a frequency standard based on trapped ions, or one that was used to predict that ion traps might be useful in factoring large numbers. These applications are only in a preliminary stage of development at the present time, and the simulations may have either com-

pletely ignored or inadequately modeled what will turn out to be the principal limitation to using trapped ions to achieve these goals.

Unfortunately, preliminary simulations of an experiment in a young field often err in presenting too optimistic a picture. It is certainly possible in principle that a neglected effect will turn out to improve matters when the measurements are actually performed, but this is not often the case in my experience. Nevertheless, I am not suggesting that we should abandon these projects because the preliminary simulations may be flawed or too optimistic. Quite the opposite—it is important to pursue these problems experimentally until the technique has been fully explored, and its fundamental capabilities and limitations are well understood. But I am suggesting that although simulations may help in the design of these kinds of experiments, it is too soon to take as conclusive their predictions of the accuracy of a frequency standard based on an ion trap or of its usefulness in addressing real-world problems in number theory.

One of the great dangers of simulating the results of an experiment before it is completed is that the result of the simulation biases our expectations and affects our judgment in subtle (and sometimes not so subtle) ways. This is an obvious dilemma. Because of technical, economic, or political limitations, it may be impossible to carry out a particular experiment, and some experiments are so complicated or expensive that simulations are absolutely necessary during the design phase of the work. Even when this is not the case, the information provided by a simulation may more than offset the intellectual "phase locking" that it may cause.

Finally, simulations are useful from a pedagogic point of view. They encourage students to think about their work from a global point of view rather than always worrying about ground loops, vacuum leaks, and syntax errors. This is a lesson that I still find personally helpful as well.

By **Judah Levine**, a fellow of JILA and physicist in the Time and Frequency Division of the National Institute of Science and Technology (NIST) in Boulder, CO. JILA is a joint institute between the University of Colorado and NIST. E-mail: jlevine@utcnist.colorado.edu