

## PERSPECTIVES

to ~20 attoliters (1 aL =  $10^{-18}$  l), and near-field scanning optical microscopy (NSOM, panel D) (10), which allows attoliter detection volumes. NSOM is best suited to cell-surface studies where scanning is important.

Single-molecule detection in reduced detection volumes can also be achieved with stimulated emission depletion (STED, panel E) (11, 12). With this method, detection volumes have been reduced to 0.67 attoliters (11). The detection volume can be placed anywhere in solution, including inside a cell. However, use of STED with more than one color of fluorophore is difficult, and the STED beam is of high intensity.

Another possibility is the use of nanofluidic channels to restrict the movement of molecules to a width and depth smaller than

the confocal detection volume (panel F) (13). Molecules would not be near metal surfaces, and diffusion times would be longer because diffusion would be effectively one-dimensional. Finally, if super-resolution lenses can be made from materials with a negative index of refraction, they may in the future be used for single-molecule detection in small volumes (14).

All these methods have strengths in different areas and are suitable for different applications. The strength of the zero-mode waveguide of Levene *et al.* (6) is that it allows parallel detection of many single molecules while providing the smallest detection volumes of any method reported to date, enabling single-molecule analysis of much weaker interactions than previously possible.

## References

1. H. P. Lu, L. Xun, X. S. Xie, *Science* **282**, 1877 (1998).
2. K. Adachi *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 7243 (2000).
3. A. A. Deniz *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 5179 (2000).
4. B. Schuler, E. A. Lipman, W. A. Eaton, *Nature* **419**, 743 (2002).
5. X. W. Zhuang *et al.*, *Science* **288**, 2048 (2000).
6. M. J. Levene *et al.*, *Science* **299**, 682 (2003).
7. C. Zander, J. Enderlein, R. A. Keller, *Single Molecule Detection in Solution: Methods and Applications* (Wiley-VCH, Berlin, 2002).
8. R. X. Bian, R. C. Dunn, X. S. Xie, P. T. Leung, *Phys. Rev. Lett.* **75**, 4772 (1995).
9. T. E. Starr, N. L. Thompson, *Biophys. J.* **80**, 1575 (2001).
10. F. de Lange *et al.*, *J. Cell Sci.* **114**, 4153 (2001).
11. T. A. Klar *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 8206 (2000).
12. M. Dyba, S. W. Hell, *Phys. Rev. Lett.* **88**, 163901 (2002).
13. M. Foquet *et al.*, *Anal. Chem.* **74**, 1415 (2002).
14. J. B. Pendry, *Phys. Rev. Lett.* **85**, 3966 (2000).

## COMPUTER SCIENCE

# Rough Times Ahead

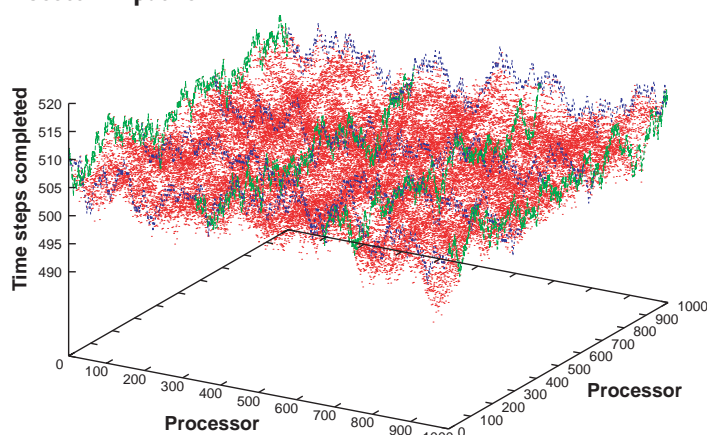
Scott Kirkpatrick

High-performance computing has moved from being a problem of optimizing the architecture of an individual supercomputer to one of optimizing the organization of large numbers of ordinary computers operating in parallel (1). This development has been made possible by rapid progress in microprocessor, memory, and storage components.

On page 677 of this issue, Korniss *et al.* (2) point out a key weakness of this approach. Simulated time, the common parameter linking the many loosely coupled elements of a distributed simulation, can get rough. Just as the surfaces of crystalline materials grown by depositing individual atoms roughen (3), so does the temporal surface of a computation spread continuously as the modeled system evolves.

This roughening becomes a problem when measurements must be continuously extracted from the complex simulation. The solution proposed by Korniss *et al.* takes advantage of a property found in the Internet: its "small world" nature of having occasional links between points that would otherwise be thought to be far apart. The authors show that intermittent synchronization over random distances can suppress temporal roughening.

The author is in the Department of Computer Science and Engineering, Hebrew University, Jerusalem, 91904 Israel. E-mail: kirk@cs.huji.ac.il



**Out of sync.** This simulated temporal surface was created by allowing 1000 x 1000 processing elements to proceed in equal time steps under local synchronization.

The best evidence of the changing nature of high-performance computing comes from the regularly updated "TOP500" survey of the world's 500 biggest computing complexes (4). By 2002, clusters of conventional computers had come to make up 93% of the listed machines, with old-fashioned, vector-based supercomputers constituting the remaining 7%. Ten years ago, when the TOP500 surveys began, 27% of the listed systems used a conventional architecture, while 66% were specialized vector machines. The remaining 7% were variations on the massively parallel Connection Machine (5), a now-extinct species.

For some computations, which are vast in extent but simple in organization, a new computing resource with lower costs is being tapped. "Embarrassingly parallel" computations are increasingly computed at

the "edge of the network." A computation is distributed to many cooperating workstations, which perform their pieces of the work with little or no coordination as background tasks. They return the completed result to a central host, which checks the results, assigns new tasks, and compiles the partial results. This approach has been used, for example, for factoring extremely large numbers and for folding a protein. But the best known example is the SETI@home project (6), in which over 4 million users have analyzed data from a radio telescope, seeking evidence of narrowband transmissions at likely communications frequencies. The work is done at the lowest possible priority—it is a screen saver.

The SETI@home project estimates that the total cpu power of the thousands of users active on any given day is 13 teraflops, greater than the combined power of the world's three biggest supercomputers, as listed in the 2001 TOP500 survey (4). However, the coordination involved in distributing and collecting all the work for the SETI@home project has consumed a significant fraction of the communications bandwidth to the outside world from the University of California's Berkeley campus.

Efforts are under way to make approaches like this applicable to a much wider class of tasks. The GRID consortium (7) of open-source developers is creating an infrastructure that permits jobs to be submitted anywhere within a "power grid" of reasonably reliable computer servers, under control mechanisms that will permit costs and resources to be allocated fairly. In the GRID

vision, not only weather and weapons simulations but also business computing tasks, such as accounting and supply-chain logistics, can be accommodated in the “grid.” Peer-to-peer file- and music-sharing schemes (8) have demonstrated that enormous storage reserves also exist at the “edge of the network,” and that these can be tapped to meet a shared interest.

Synchronizing all this computation to allow extracting the state of the computation at a particular time is a well-known problem. SETI@home emphasizes that a participant whose screen saver announces that aliens have left a message should under no circumstances call the press. The real information content of the data from that part of the sky and time period will not be completely assessed until all relevant work units are returned, which may take weeks. SETI@home is willing to wait weeks to assemble all the data because interesting events are very rare. In other cases, such as weather prediction, such delays are unacceptable, and synchronization becomes crucial.

Korniss *et al.* consider a simple model of synchronization in which no processing element proceeds to the next time step until it has checked that it is in sync with its neighbors. To expose the worst case, they consider a one-dimensional array of processing elements, each

with only two neighbors. The temporal dispersion that emerges is enormous.

One might question whether simulations of two-dimensional (2D) (such as, the temperatures in all European cities) or three-dimensional phenomena (such as, the weather) exhibit such a dramatic temporal roughening. The figure exhibits the result of Korniss *et al.*'s local synchronization model for a 2D array of  $1000 \times 1000$  processing elements. Asynchronously, in arbitrary order, each processing element has attempted to take a time step ahead, first checking with its four neighbors. On average, each processor has made 1600 attempts. Roughly one-third of them are successful, and already the time steps completed are scattered over a range of at least 30 time steps. Data over this temporal range must be stored for subsequent combination with data from distant points in the simulation.

Korniss *et al.* show that the roughening satisfies equations developed to explain crystal growth. These equations predict that the temporal width continues to grow over time until it is limited by the size of the simulation. The challenge it poses will thus continue to grow as problem sizes increase. The physical analogy to roughening also implies that long-range interactions will eliminate or drastically reduce the interface width.

There are trade-offs between computing and communication costs that can be made in exploiting this insight. One can eliminate temporal roughening by requiring each processing element to check for synchronization with a randomly chosen (and probably distant) other element before starting some small fraction of its time steps. Alternatively, one can ask the same small fraction of the processing elements to synchronize with a single distant partner before every time step. Both of these approaches result in essentially the same smoothing. Which solution is preferred will depend on the cost of the fewer links to distant sites required versus the imbalance in the workload of the elements on which falls the extra demand of participating in the small world far away from them.

#### References

1. C. G. Bell, J. Gray, *Commun. ACM* **45**, 91 (2002).
2. G. Korniss, M. A. Novotny, H. Guclu, Z. Toroczka, P. A. Rikvold, *Science* **299**, 677 (2003).
3. M. Kardar, G. Parisi, Y.-C. Chang, *Phys. Rev. Lett.* **56**, 889 (1986).
4. www.top500.org
5. D. Hillis, *The Connection Machine* (MIT Press, Cambridge, 1985).
6. setiathome.ssl.berkeley.edu
7. www.gridforum.org
8. A. Oram, Ed., *Peer-to-Peer: Harnessing the Power of Disruptive Technologies* (O'Reilly, Sebastopol, CA, 2001).

## ATMOSPHERIC SCIENCE

# Deciphering the Energetics of Lightning

Philip E. Krider

On page 694 of this issue, Dwyer *et al.* (1) report the first observations of energetic radiation—x-rays, gamma rays, and/or relativistic electrons—during rocket-triggered lightning. Using an instrument designed to operate in the highly disturbed environment near a lightning strike (see the figure), the authors detected intense bursts of radiation (much more than 10 keV) just before 84% of the return strokes that they were able to trigger. The radiation began while dart leaders were propagating downward—before they contacted ground and initiated return strokes, intense pulses of current that begin at the ground and propagate upward at close to the speed of light (2).

Their sensor, a NaI(Tl) scintillation counter operating in conjunction with a

control detector, was placed within 25 m of the rocket launcher at a triggering site in central Florida, the International Center for Lightning Research and Testing (ICLRT). In all cases, energetic radiation began when the dart leader was close to the ground and before it initiated an upward-propagating return stroke. The radiation appeared to cease at, or a few microseconds after, the onset of the stroke current. Although the observed bursts of radiation were short and started at most 160  $\mu$ s before the onset of the current, the total energy deposited in the detector was large, typically many tens of MeV per stroke.

The dart leader/return stroke sequences in triggered lightning are similar to the return strokes that come after the first stroke in natural cloud-to-ground lightning (2). Because Dwyer *et al.* detected high-energy ra-

diation during 31 out of 37 triggered strokes, they suggest that such radiation is probably present in all cloud-to-ground flashes. This hypothesis is consistent with a recent report of energetic radiation during the final stages of the stepped-leaders that initiated three natural flashes to a mountain in New Mexico (3). The fact that energetic radiation has been observed in Florida near sea level (1 atm pressure) adds credence to other observations and models of energetic radiation at higher altitudes, where such radiation can be more easily produced.

The search for radiation from thunderstorms and lightning has a long history, dat-



**Lightning strikes.** A rocket-triggered lightning flash striking on or near a short rod at the ICLRT.

The author is in the Institute of Atmospheric Physics, University of Arizona, Tucson, AZ 85721, USA. E-mail: krider@atmo.arizona.edu