



PARALLELISM AND CSE

David J. Kuck

TRENDS IN MANY AREAS OF COMPUTATION ARE MOVING TOWARD MULTIFUNCTIONAL APPLICATIONS. IN SCIENCE AND ENGINEERING, FOR EXAMPLE, THIS MIGHT TAKE THE FORM OF MULTI PHYSICS APPLICATIONS, WHEREBY FUNCTIONALITIES (SUCH as structures and fluids applications) integrate into a single new application. It might also take the form of a computational chemistry application that embeds a spreadsheet format and allows users to specify a set of runs via a high-level graphical user interface. Even in nontechnical areas, combining applications is the order of the day—witness databases embedded in enterprise and Internet applications.

Developing multifunctional applications requires reusing existing applications. This has led many forward-looking independent software vendors (ISVs) to team up to produce joint products. The products that result provide users with more benefits but at the cost of requiring more capable computer systems.

Continuing system performance improvement is necessary in areas ranging from CSE to ERP (enterprise resource planning), and parallelism can provide the increasing speed and memory sizes needed. Parallelism in computing is as old as the first two decades of electronic computers, and it is as new as four-processor parallel workstations or 4,000-processor massively parallel supercomputers. Most ISVs are now motivated to develop parallel applications, but new efforts are often haunted by the broken schedules and superlinear resource demands of many past parallelism projects.

While the Department of Energy's ASCI project, for example, can still afford to put multiyear efforts into new application development, most development projects cannot. Most of the software used today, even in large enterprises, is produced by ISVs—not users or enterprises. Because the 20,000-odd ISVs are mostly small, discipline-focused companies, practical approaches to parallel software development are crucial.

A few simple principles must be followed today to develop successful parallel applications:

1. *do* use parallel software engineering technology with maximum power,
2. *do* evolve sequential applications toward parallelism to preserve sequential functionality,
3. *don't* use hardware with insufficient bandwidths or excessive latencies for chosen applications, and
4. *don't* include desired functionality that requires high communication if it blocks scalability.

The good news is that at the desktop and server levels, practical SMP (shared memory parallel) software-engineering techniques are now available, and increasing numbers of ISVs are taking advantage of them via principles one and two. However, to exploit larger systems, the two don'ts must be mitigated before scalable parallelism is practically useful. Today, to reach very-high performance goals, it is often necessary to develop from scratch, focus on low communication, and drop any functionality that blocks scalability.

The main reason for building DMP (distributed memory parallel) or NUMA (nonuniform memory access) SMP systems is that the physics of current technology, over time, has prevented arbitrarily large SMP system construction. Barring major changes in computer technology, something beyond uniform memory SMP systems will be needed indefinitely, and so we must address the two don'ts.

What can we learn from the dos that might suggest solutions for the don'ts? Both dos are related to progress in directive-based parallel programming over the past 20 years, culminating in the past year with the new OpenMP programming model and directives for Fortran, C, and C++. A powerful SMP programming model and effective parallel software engineering tools are now available, making hand-threaded assembly-language-level thinking and programming unnecessary for SMP developers. Proof of this claim comes from the increasing number of ISVs adopting OpenMP.

NUMA systems (SMP with widely varying main-memory access times) and clustered SMP systems (message passing between SMP nodes) currently address the

These essays

Editors-in-Chief George Cybenko and Lewis Holmes have asked a number of leading scientists and engineers to comment on the effect computation is having on their work. These essays will appear in this and subsequent issues of *Computing in Science & Engineering*. — Eds.

long-term requirements for arbitrarily scalable parallelism. Because SMP programming is basic to both architectures, can the SMP programming dos be naturally extended to mitigate principle four in some way?

Today, message passing for DMP systems is difficult, but it is the most successful DMP technique. High-performance Fortran has not lived up to its promise—mostly due to difficulties with irregular data and program dynamics. Presumably it might have been more successful if it had made programming less labor-intensive. However, success with message passing comes from serious thinking about domain decomposition and the assembly-language-level expression of ideas—both very labor-intensive activities. In other words, the issue is not really programming languages or even human effort—it is the underlying mismatches between algorithms and DMP architectures that trap unwary users and foil performance.

When an application needs more parallelism than current SMP systems provide, the parallelism might have arisen for one of three reasons:

- the natural parallelism used in the SMP system is simply greater than the system can handle,
- nested parallelism exists (a second level is found above the SMP level already in use), or
- multifunctional computations occur that decompose at a very high level, with distinct functionalities running on distinct SMP systems.

An approach to scalable parallelism in these cases is to iteratively refine applications through SMP-style programming. First, you develop the best SMP applications possible and then study their structure and performance, locating ideal points for decomposing them into NUMA/cluster parallelism and introducing messaging or nested parallelism directives at these crucial points. Note that finding and expressing such parallelism

(which OpenMP allows with directives) is only half the problem. The remaining issue is having a system that executes the program for each data set in a low-communication, load-balanced manner.

This requires new tools and extends the dos above. It does not allow for weak systems, of course—increasingly fast memory systems must remove principle three. Supposing that hardware advances and user awareness *do* control number three, the goal of mitigating number four might be reached by new software research that extends the directive-based programming model while preserving sequential functionality.

After some 30 years of increasingly intense R&D, we now have effective SMP software-engineering methods. We can expect continuing progress toward scalable practical parallelism that will make most parallelism relatively easy to deal with. However, there might always be some applications that yield only to totally original thinking and message passing for high-level decomposition. These open questions for the future naturally lead us to CSE's education and research agenda.

A number of university CSE programs exist today, but with great diversity. Some are looking for ways to focus and integrate themselves in the larger CSE community. One way of accelerating progress would be for them to share software and educational tools. Specifically, they could use a common software base to focus on improving real applications relative to realism, multifunctionality, performance, numerical quality, and usability. Effective study and development could be carried out by students in various departments—from the viewpoints of science and engineering models, compilers, performance tuning, numerical methods, GUIs, and so forth. The architectural focus would range from parallel workstations to distributed supercomputing.

A lasting contribution to performance

improvement would be in parallelization. Source code containing OpenMP directives for SMP parallelism could be provided that meet the dos above, and would provide an easy learning experience for new graduate students relative to the application's functionality and its improvement potential. This effort would begin in fields having workhorse public-domain codes and then spread to partially revealed source code for ISV applications. Software ISVs could provide powerful tools. (For information about university courses and applications areas already moving in this direction, visit www.kai.com/kpts/applications.) This would lead to a cross-university fabric of applications, tools, and development activities in areas with widely used public-domain codes.

The benefits would include the development of better application software, algorithms, parallel software engineering tools, and performance information. MS theses of all kinds would quickly arise, and eventually PhD students would write theses in these areas. This might even lead to consensus grand-challenge questions that the community could work on jointly, replacing after-the-fact attempts to claim grand-challenge status for any high-performance project. Eventually, this would benefit the entire CSE educational community as well as CSE as a field.

David J. Kuck is a chairman of the board at Kuck and Associates Inc. (KAI) and a professor emeritus of computer science and electrical and computer engineering at the University of Illinois at Urbana-Champaign. He received his BSEE from the University of Michigan, and his MS and PhD in engineering from Northwestern University. He is a fellow of the American Association for the Advancement of Science, ACM, and IEEE, and a member of the National Academy of Engineering. Contact him at Kuck and Associates Inc., 1906 Fox Dr., Champaign, IL, 61820; kuck@kai.com

MOLECULAR DYNAMICS SIMULATION

D.C. Rapaport

THE PROPERTIES OF MATTER ARE DETERMINED BY THE COLLECTIVE BEHAVIOR OF THE COMPONENT ATOMS.

THEORY—IN SOME CASES, WORK SPANNING WELL OVER A CENTURY—HAS TAUGHT US A GREAT DEAL IN FIELDS SUCH AS MATER-

ial science, polymers, fluids, and surfaces. But there are often serious simplifications involved, such as replacing the discrete atoms with a continuous medium. More detailed exploration at the atomistic level requires solving various incarnations of the N -body problem, a task we can only perform numerically. Such techniques, in which the particles are classical (that is, nonquantum) representations of the atoms or molecules, are known as molecular dynamics (MD) simulation.

MD requires a description of the molecules and the forces that act between them; a well-known example is the Lennard-Jones potential, in which spherical particles repel one another at close range but otherwise attract. The MD simulation itself amounts to numerically integrating the equations of motion for systems of between a few hundred and a few million particles over many thousand (or more) timesteps. The paths the particles follow during the computation represent actual molecular trajectories.

A variety of issues must be addressed when developing an MD simulation. The integration technique used, such as the leapfrog method, must provide an accurate solution; an important accuracy test is the ability to conserve energy in a mechanically isolated system. The particles are typically confined to a box-shaped region; periodic boundaries—which treat the box as just one member of a spatially repeating array—effectively confine the particles without the perturbing effects of real hard walls. The initial state must be specified. Property measurements must be implemented, and their accuracy and statistical reliability ensured as in experiment. An isolated system will normally reach an equilibrium state, perhaps resembling a crystalline or amorphous solid, or a liquid; adequate time must be allowed to reach this state before making any measurements.

Molecules more complex than simple spherical particles are readily modeled. Multiple interaction sites can represent rigid molecules of various kinds. Flexible molecules, both short oligomer and long polymer chains, can be constructed by permanently linking particles with suitable springs. To simulate complex biopolymers such as proteins, molecules with a partially rigid skeleton (to eliminate fast internal vibrations that would lead to stiff equations) can be built with geometric constraints. The interactions themselves need not be based on simple pair potentials; we can incorporate orientational dependence and three-body effects. Weak quantum effects can be expressed in terms of corrections to measured properties, if needed. A combination of MD and density functional methods provides the means for modeling semiconductors and liquid metals when we have to take the valence electron behavior into account. MD is not limited to equilibrium studies; because it is capable of handling open systems, such as a fluid mechanically sheared by a sliding wall, it can address problems beyond the capabilities of current theory. It can also explore environments unobtainable experimentally, such as extreme pressure conditions.

Because an MD simulation produces a complete trajectory history, it can measure any conceivable property. These include thermodynamic features such as the equation of state (temperature is proportional to kinetic energy and pressure follows from the virial), dynamical properties such as viscosity, elastic constants, and structural characteristics of which the radial distribution function is the simplest. Measurements such as these all correspond directly with what we can obtain in the laboratory, using, for example, calorimetry or X-ray scattering. Other measurements have no laboratory counterparts but can nevertheless be used to study the underlying behavior in greater detail; examples include three-body correlations and trajectory properties.

Computational efficiency is important because interesting problems tend to push the envelope in terms of system size and/or the number of integration timesteps. While the number of particle pairs is proportional to N^2 , the force computation can be reduced to order N —in the case of short-range forces—thanks to suitable bookkeeping techniques. This is based on the fact that each particle interacts with others in a small neighborhood. Alternative techniques have been developed for long-range forces based on Ewald summation and cell-multipole methods.

Despite these efforts, MD is limited in

Further reading

1. M.P. Allen and D.J. Tildesley, eds., *Computer Simulation in Chemical Physics*, Kluwer Academic, Dordrecht, The Netherlands, 1993.
2. K. Binder and G. Ciccotti, eds., *Monte Carlo and Molecular Dynamics of Condensed Matter Systems*, Italian Physical Soc., Bologna, Italy, 1996.
3. D.C. Rapaport, *The Art of Molecular Dynamics Simulation*, Cambridge University Press, Cambridge, UK, 1995.




Figure 1. 3D convection patterns (top-down view).

what it can achieve: an effect whose description embraces a wide range of time-scales, such as protein folding or glass formation, remains a computational challenge. Ideally (and this is a subject of some interest), one would like to merge various computational approaches, so that different parts of the system are represented at varying levels of detail. In this way, the computational effort could be better focused.

MD is the most general of the atomistic simulation techniques available. Among the alternatives, Monte Carlo is based on statistical sampling and used in equilibrium studies, while lattice-based models such as cellular automata are very limited in their descriptive powers. MD's richness also has a negative side, because the wealth of detail implies heavy computation; fortunately, increasing processor power alleviates this problem's severity to some extent. MD has been implemented on a variety of advanced computer architectures, including vector and parallel supercomputers, and even special-purpose machines. Large-scale parallelism is beneficial only within limits, however, because a certain minimal number of particles per processing node are needed; this implies that massive parallelism is unlikely to provide a complete solution to the timescale problem.

MD applications are ubiquitous. Work on complex fluids includes molecular liquids of many kinds, the structure and dynamics of glasses, water and ionic solutions, liquid crystals, and thin films. Studies of solids include molecular crystals, structural transformations, fracture, radiation damage, mechanical properties, friction, shock waves, and epitaxial growth. Biochemical simulations include protein dynamics and the docking processes between molecules important in drug design. Other more fundamental uses aim to test theory. MD methodology is also crucial for modeling granular flow, where the particles correspond to grains and the interactions include a dissipative component.

Visualization plays an important role in situations where the behavior is complex and time dependent. It might well be the static or animated visual imagery derived from the results of an MD simulation that best summarizes system behavior. Two figures from my current work illustrate this point.

Figure 1 demonstrates the use of MD in hydrodynamic instability studies—in this case, the convection patterns that form when a thin fluid layer is heated from below (the Rayleigh-Bénard problem). The figure shows a top-down view of a 3D system containing four million particles, with color distinguishing hot fluid upflow from cold downflow. The formation of a set of thermal convection cells is clearly apparent. The surprise is that the same familiar self-organized flow occurring in macroscopic fluids should also appear in this microscopically small system whose layer thickness is a mere 100 Å.

Figure 2 is from a study of supramolecular self-assembly in which simple structural elements combine spontaneously to form closed polyhedra reminiscent of the icosahedral protein shells that appear in many viruses. The robustness of this process in nature suggests that a simplified model based on rigid bodies with suitable interactions might be used in simulating assembly.




Figure 2. Growth of 60-faced polyhedra using an MD approach.

What does the future hold? MD simulation covers length scales ranging from the atomistic to entire microstructures. It has proved capable of studying a broad range of phenomena associated with both simple and complex molecules. It is free of many of the simplifying assumptions that tend to dominate theory and other modeling techniques. So, after making the reasonable extrapolation that computer power will continue to grow at its present rate, I have little doubt that MD is destined to play an ever-increasing role in both science and engineering.

D.C. Rapaport is a professor in the Physics Department at Bar-Ilan University, Israel. Rapaport received his PhD in theoretical physics from the University of London. His technical interests include simulational and computational physics and statistical mechanics. Contact him at the Physics Department, Bar-Ilan University, Ramat-Gan, Israel 52900; rapaport@bluegum.ph.biu.ac.il.

THE IMPORTANCE OF COMPUTER SIMULATIONS OF GRANULAR FLOW

Hans J. Herrmann

GRANULAR MATERIALS ARE VERY COMMON IN OUR DAILY LIFE BUT ALSO IMPORTANT FOR TECHNOLOGICAL APPLICATIONS AND IN GEOPHYSICS. UNTIL THE LAST CENTURY, THEY HAD BEEN STUDIED BY SUCH PHYSICISTS AS MICHAEL FARADAY, James Maxwell, and Osborne Reynolds. However, because of the difficulties in formulating a continuum theory for their deformation, they were abandoned by physicists for many decades and only rediscovered about 10 years ago. In the meantime (that is, since the 1950s), two engineering communities (soil mechanics and chemical engineering) succeeded in formulating continuum equations of motion that work rather well in the dense¹ and dilute² limits of granular systems. But the foundation of these theories is, in view of modern statistical mechanics, still quite open and has recently become a very active research field. Most important, these theories are not applicable in most of the realistic situations in which the granular material changes its density spontaneously and in which different types of grains induce interesting segregation patterns.

Many rather astonishing phenomena occur when granular materials such as sand or powders move.³ Examples are the “Brazil nut” segregation, heap formation under vibration,⁴ density waves emitted from outlets, and $1/f$ noise in the power spectra of local forces. All these effects originate in the ability of granular materials to form a hybrid state between a fluid and a solid: When the material’s density exceeds a certain value, the *critical dilatancy*,⁵ the material is resistant to shear, like solids, while below this density it will “fluidify.” This fluidified state can be rather complex, especially in the presence of density fluctuations and density gradients.

Particularly suited to study this fluidization is an experiment where sand is put on a loudspeaker or on a vibrating table. Under gravity, the sand jumps up and down; although kinetic energy is strongly dissipated, collisions among the grains reduce the sand’s density, thereby allowing it to flow (fluidization). Under certain circumstances, flow between the top and bottom can occur in the form of convection cells, as researchers have observed experimentally in the case of inhomogeneities in the vibration’s amplitude. More striking is that sand spontaneously can form heaps, as Faraday first described in 1831.⁴ Within these heaps, convection occurs that might even be the motor for the heap formation: Inside the heap the sand rises, pops out at the top, and slides down on the surface. Usually these heaps have complicated shapes that change in time, and sometimes you can also see ripples and other regular structures on their surface. If you put particles of different sizes but equal density on the vibrating plate, the larger particles tend to rise. After some time, you’ll see a segregation into regions with larger particles and regions with smaller particles. When the plate’s vibration also has a horizontal component, the material will flow in one direction, a technique often used in powder transport.


Researchers have made various attempts to formalize and quantify the complicated rheology of granular media. They’ve proposed continuum equations of motion and a kinetic theory,² thermodynamic formulations, cellular automata,⁶ and

random-walk approach.⁷ But these theories haven’t been able to satisfactorily explain phenomena, such as size segregation or density fluctuations, that are part of many of the effects we mentioned previously. This is because incorporating static friction, local rotations, and other relevant microscopic mechanisms into these theories is very difficult.

Therefore, to better understand the rheological effects of granular media, computer simulations are very important.⁸ For over a decade, researchers have used discrete-element methods that, instead of treating the granular material as a continuum, treat it as an assemblage of particles interacting through their contacts. Peter Cundall introduced this technique to study the motion of rock masses.⁹ Since then it has been applied to statistical micromechanics, constitutive behavior of granular soils, creep of soils, analysis of rock-support interaction, and other applications of soil mechanics. This technique has also been applied to model-size segregation, outflow from a hopper, shear flow, and flow down an inclined chute.

The discrete-element methods described above are very similar to what physicists have been calling, for decades, *molecular dynamics*.¹⁰ In the case of frictional hard-core interactions that we encounter in granular matter, different variants of molecular dynamics are used, depending on the application. If the particles move rapidly, collisions are considered instantaneous (with infinite forces), and a lookup table for the velocities can describe a collision’s outcome. These techniques are called *event-driven* and are particularly fast on workstations (but difficult to parallelize). Figure 1 shows the application of event-driven techniques to illustrate the surface waves that occur when a shallow bed of grains vibrates vertically. The picture, which was generated by Stefan Luding,¹¹ was produced in 2D with fixed walls. (Similar simulations in 3D were

Figure 1. Surface waves in a 2D vibrated box. Colors indicate a particle's kinetic energy (blue = low, red = high).¹¹



performed at the University of Texas, Austin,¹² also where Paul Umbanhowar, Francisco Melo, and Harry Swinney conducted their spectacular experiments on surface waves,¹³ which triggered the simulation of Figure 1 and other similar work.) Another important application for event-driven simulations is cluster formations in intergalactic dust.

For dense systems and for certain applications, such as the propagation of sound in granular packings, using elastic repulsion potentials between the particles can avoid infinite collisions in a finite time interval. These soft-core molecular dynamics have been applied successfully to describe the clogging of silos and pipes, segregation effects under shear or under vibration, and the pressure distribution below sand heaps.

Although the above methods have been pushed to simulate more than 10^9 particles during several microseconds' real time, current workstations can typically simulate only approximately 10^4 particles, using one day of CPU time to simulate realistic experimental times. For many applications, in particular in 3D, this number is not large enough because collective phenomena occur on very large scales. Imagine a silo filled with wheat grains that has a 10-cm-diameter outlet that clogs!

To simulate a much larger number of particles than with a discrete-element method, new approaches have been proposed. One is *Discrete Simulation Monte Carlo*, which treats collision events stochastically. Instead of calculating each particle's trajectory and waiting for the particles to collide, DSMC divides the system into boxes, out of which it randomly chooses collision partners. These methods allow an increase of the number of particles by a factor of 100.¹⁴

Other new algorithms similar to cellular automata place the particles on a lattice, simplifying the calculation of the coordinates. Such methods have been very successful in reproducing experimental details such as the logarithmic tail at the bottom of the sand heap.¹⁵

The importance of the technological applications of granular media in silos or conveyor belts, in ball mills or under bulldozers, demands more efficient simulation techniques to reproduce entire industrial processes. Although several commercial programs for the simulation of granular material exist, their use is not nearly as common as in fluid or solid mechanics. This is mainly due to historical reasons and to limitations in the number of simulated particles. The advancement of computer hardware and the rapid advance in simulation algorithms ensure that particle simulations will become much more commonplace in power technology. Other areas will certainly profit from this progress in computational granular materials, such as geosciences and astrophysics. One beautiful example is the simulation of the transport of particles under the action of wind, and the subsequent formation of dunes.

References

1. I. Vardoulakis and J. Sulem, *Bifurcation Analysis in Geomechanics*, Chapman & Hall, London, 1995.
2. P.K. Haff, "Grain Flow as a Fluid-Mechanical Phenomenon," *J. Fluid Mechanics*, Vol. 134, 1983, pp. 401-430.
3. H.M. Jaeger and S.R. Nagel, "Physics of the Granular State," *Science*, Vol. 255, No. 3, Mar. 1992, pp. 1523-1531.
4. M. Faraday, "On a Peculiar Class of Acoustical Figures; and on Certain Forms Assumed by Groups of Particles upon Vibrating Elastic Surfaces," *Philosophical Trans. Royal Soc. London*, Vol. 52, 1831, p. 299.
5. O. Reynolds, "On the Dilatancy of Media Composed of Rigid Particles in Contact," *The London, Edinburgh and Dublin Philosophical Magazine and J. Science*, Vol. 20, No. 127, 1885, p. 469.
6. G. Peng and H.J. Herrmann, "Density Waves of Granular Flow in a Pipe Using Lattice-Gas Automata," *Physical Rev. E*, Vol. 49, No. 3, Mar. 1994, pp. R1796-R1799.
7. H. Caram and D.C. Hong, "Random-Walk Approach to Granular Flows," *Physical Rev. Letters*, Vol. 67, No. 7, Aug. 1991, pp. 828-831.
8. O.R. Walton, "Numerical Simulation of Inelastic, Frictional Particle-Particle Interactions," *Particulate Two Phase Flow*, M.C. Roco, ed., Butterworth-Heinemann, Boston, 1992, pp. 884-911.
9. P.A. Cundall, *A Computer Model for Rock-Mass Behavior Using Interactive Graphics for the Input and Output of Geometrical Data*, Tech. Report AD/A-001 602, Nat'l Technical Information Service, Springfield, Va., 1974.
10. M.P. Allen and D.J. Tildesley, *Computer Simulation of Liquids*, Oxford Univ. Press, Oxford, UK, 1987.
11. S. Luding et al., "Simulations of Pattern Formation in Vibrated Granular Media," *Europhysics Letters*, Vol. 36, No. 4, Nov. 1996, pp. 247-252.
12. C. Bizon et al., "Patterns in 3D Vertically Oscillated Granular Layers: Simulation and Experiment," *Physical Rev. Letters*, Vol. 80, No. 1, 1998, pp. 57-60.
13. F. Melo, P.B. Umbanhowar, and H.L. Swinney, "Transition to Parametric Wave Patterns in a Vertically Oscillated Granular Layer," *Physical Rev. Letters*, Vol. 72, 1994, pp. 172-175.
14. M. Müller, S. Luding, and H.J. Herrmann, "Simulations of Vibrated Granular Media in 2D and 3D," *Friction, Arching, and Contact Dynamics*, World Scientific, Singapore, 1997, pp. 335-340.
15. J.J. Alonso and H.J. Herrmann, "On the Shape of the Tail of a Sand Pile," *Physical Rev. Letters*, Vol. 76, 1996, pp. 4911-4914.

Hans J. Herrmann is the director of the Institute for Computer Applications 1 at the University of Stuttgart. His research interests include granular media, fracture, irreversible growth, cellular automata, biophysics, traffic models, computers, and geophysics. He received his doctorate in physics from the University of Cologne. He is a Guggenheim Fellow and a corresponding member of the Brazilian Academy of Science. Contact him at the Inst. for Computer Applications 1, Univ. of Stuttgart, Pfaffenwaldring 27, D-70569 Stuttgart, Germany; hans@ica1.uni-stuttgart.de; <http://www.ica1.uni-stuttgart.de/home/hans/WWW/personal.html>.

ECONOPHYSICS: CAN PHYSICISTS CONTRIBUTE TO THE SCIENCE OF ECONOMICS?

H. Eugene Stanley

THE FUNDAMENTAL PRINCIPLES GOVERNING THE COMPLEX SYSTEM CALLED ECONOMICS ARE NOT COMPLETELY UNCOVERED. THIS OBSERVATION SEEMS TO BE ALMOST GENERALLY ACCEPTED; FOR EXAMPLE, THE 23 AUGUST 1997 ISSUE OF *THE ECONOMIST* FEATURED THE COVER ARTICLE, “THE PUZZLING FAILURE OF ECONOMICS.”

Then how can computational physicists contribute to the search for solutions to the puzzles posed by modern economics that economists themselves cannot solve? An approach—not very commonly used in economics—is to begin empirically, with real data that you can analyze in some detail, but without prior models. In economics, a great deal of real data is available. If you, moreover, have at your disposal the tools of computational physics and the computing power to carry out any number of approaches, this abundance of data is a great advantage. Thus, for physicists, studying the economy means studying a wealth of data on a well-defined complex system. Indeed, physicists in increasing numbers are finding problems posed by economics sufficiently challenging to engage their attention, independent of any personal profit that might be made. Various terms have been applied to this new interdisciplinary subfield of physics. Some French physicists prefer the term *phynance*, while others prefer other terms. In an analogy with the terms *biophysics*, *geophysics*, and *astrophysics*, in 1994 or 1995 I introduced the term *econophysics* to attempt to legitimize why physics graduate students should be allowed to work on problems originating in economics.

If we physicists have any prior bias, it might be the lesson learned years ago when many of us worked on critical phenomena: Everything depends on everything else. A careful analysis of any system involves studying the propagation of correlations from one unit of the system to the next. We learned that these correlations propagate both directly and indirectly. At one time, it was imagined that “scale-free” phenomena are relevant to only a fairly narrow slice of physical phenomena. However, the range of systems that apparently display power-law and hence scale-invariant correlations has increased dramatically in recent years. Such systems range from base-pair correlations in noncoding DNA, lung inflation, and interbeat intervals of the human heart, to complex systems involving large numbers of interacting subunits that display “free will,” such as animal behavior¹ and even human behavior.² In particular, economic time series—for example, stock market indices or currency exchange rates—depend on the evolution of a large number of strongly interacting systems far from equilibrium, and belong to the class of complex evolving systems. Thus, the statistical properties of economic time series have attracted the interests of many physicists. Space limitations mo-

tivate me to focus mainly on the Boston group’s results; the work of other research groups is described elsewhere.

Economic time series: correlations or the lack thereof

The recent availability of “high-frequency” data lets us study economic time series on a wide range of time scales varying from seconds up to years. Consequently, researchers have applied a large number of methods known from statistical physics to characterize the time evolution of stock prices and foreign exchange rates.

Much of our recent work is based on analysis of the S&P 500 index, an index of the New York Stock Exchange (NYSE) that consists of the 500 largest companies in the US. It is a market-value weighted index (stock price times the number of shares outstanding), with each stock’s weight in the index proportionate to its market value.³⁻⁵ The S&P 500 index is one of the most widely used benchmarks of US equity performance. Data typically cover a long period, such as 13 years (from January 1984 to December 1996), with a recording frequency of one minute or shorter. The total number of data points in this 13-year period exceeds one million, three orders of magnitude greater than Benoit Mandelbrot’s classic analysis of cotton price fluctuations.⁶

The S&P 500 index $Z(t)$ from 1984 to 1996 tends to drift constantly upward on a semilog graph—except during crashes, such as in October 1987 and May 1990. We analyze the difference of the logarithm of the index values $G(t) \equiv \log_e Z(t + \Delta t) - \log_e Z(t)$, where Δt is the time lag. We count only the number of minutes during the stock market’s opening hours and remove the nights, weekends, and holidays from the data set. That is, the market’s closing and next opening is continuous.

The distributions of the increments of economic time series, both in stock market indices and foreign currency exchange rates, turn out to be nearly symmetric and have

very fat tails (strong leptokurtic wings). Index increments as a function of time show exponentially decaying correlations that are at noise level after a few minutes. This makes these increments fundamentally different from many well-studied examples of complex dynamical systems in physics. One such example is turbulent flow, which commonly displays power-law correlations on long time scales.^{7,8}

The situation is different for the volatility, which is calculated, for example, averaging market fluctuations over a suitable time interval. The volatility has long time persistence—much larger than the correlation time for price changes.^{4,5} Quantifying the volatility's dynamics is important. Volatility is the key input of virtually all option pricing models, including the classic Black and Scholes model, which is based on estimates of the asset's volatility over the option's remaining life.

Specifically, using both traditional power-spectrum methods as well as a new method—*detrended fluctuation analysis* (DFA)⁹—Yanhui Liu and his coworkers^{4,5} detect long-range volatility correlations embedded in a nonstationary time series. This new method avoids the spurious detection of apparent long-range correlations that are an artifact of nonstationarities.

Both methods show the existence of two distinct regions of power-law behavior for the autocorrelation function of volatility, with the exponents $\alpha_1 = 0.66$ and $\alpha_2 = 0.95$ for t less than or greater than a characteristic time scale t_x on the order of one day. It is as if the information used in a single day to make trading decisions differs from the long-term information used.

To test whether this correlation is a spurious artifact of the distribution function, which might have long tails, Liu and his colleagues shuffle each point of the volatility time series. The random shuffling keeps the volatility distribution unchanged but totally kills any correlations in the time series. DFA analysis of this randomly shuffled data does not show any correlations and

gives the exponent $\alpha = 0.50$. This tells us that the long-range correlations are due to the economic system's dynamics and not simply due to the fat-tailed distribution, because the distribution does not change when the data are shuffled.

Histograms of price changes

Although the correlations in the price change $G(t)$ are not particularly novel, the histograms of price changes are. Because economic systems consist of a large number of interacting units, it is plausible that they might be amenable to scaling analysis. Mandelbrot in 1963 demonstrated that the histogram of fluctuations in cotton prices obeys a scaling distribution, the *Lévy distribution*.⁶ A recent study determined that the high-frequency fluctuations in the S&P 500 index also exhibit scaling behavior.³ Analyzing almost one million records at one-minute intervals over six years of trading, Rosario N. Mantegna and I determined that fluctuations on a one-minute time interval were reflected in 10-minute, 100-minute, and 1,000-minute intervals.³ The distribution of index returns fits a Lévy distribution with a sharp drop-off in the tails. These scaling properties mean that viewing stock market returns at one-minute intervals provides insight on the behavior at 1,000-minute intervals.

Thus, the Lévy part of the S&P 500 distribution agrees with Mandelbrot's 1963 cotton-price results, but the tail truncation does not (presumably because the tail statistics in the low-frequency results are not above the noise level). Recently, Parameswaran Gopikrishnan and his colleagues asked whether this discrepancy could be because the S&P is an average over many firms.¹⁰ To this end, they analyze a database documenting every trade in the three major US stock markets—the NYSE, the American Stock Exchange (AMEX), and the National Association of Securities Dealers Automated Quotation (NASDAQ)—for the entire two-year period of January 1994 to December 1995. They thereby extract a

sample of approximately 40 million data points, which is much larger than the one million data points analyzed by Liu and his colleagues^{4,5} and the approximately 1,000 data points studied by Mandelbrot. Gopikrishnan and his colleagues find, remarkably, an asymptotic power-law behavior, with an exponent $\alpha \approx 3$ for the cumulative distribution (see Figure 1) that is well outside the Lévy regime ($0 < \alpha < 2$).¹⁰

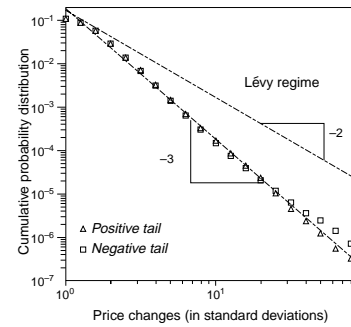
In summary, previous proposals for the histogram of index changes have included a Gaussian distribution, a Lévy distribution,⁶ and a truncated Lévy distribution, where the tails become “approximately exponential.”³ The inverse cubic result differs from all three proposals. Unlike the Gaussian distribution and the truncated Lévy distribution, it has diverging higher moments, and unlike the Gaussian distribution and the Lévy distribution, it is not a stable distribution.

Economic organizations

Economics is of course much broader than just analyzing economic time series. Many physicists imagine they can add new ideas on how to analyze a time series, but what about general questions in social science, which concerns itself with the organization of individuals—each with free will? Taking the same “empirical” approach, the Boston group has also studied a range of data on economic organizations—viewing the data through the special eyeglasses of critical phenomena (imagining that “everything depends on everything else”). Specifically, in collaboration with a card-carrying economist, Michael A. Salinger, we studied the possibility that all the companies in a given economy might interact, more or less, like an Edwards-Anderson spin glass.¹¹ In that spin glass, each spin interacts with every other spin—but not with the same coupling and not even with the same sign.

For example, a 10% decrease in the sales of a given business firm will have repercussions in the economy. Some of the repercussions will be favorable—firm B, which

Figure 1. A log-log plot of the cumulative probability distribution $P(g)$ of the normalized price increments where g is calculated in units of a standard deviation. The lines are power-law fits to the data over the range from two to 100 standard deviations. The regression lines yield $\alpha = 2.93$ and $\alpha = 3.02$ for the positive and negative tails. (Figure courtesy of Parameswaran Gopikrishnan, Luis A.N. Amaral, and Martin Meyer.)



competes with A, might experience an increase in market share. Others will be negative—service industries that provide personal services for firm A employees might experience a drop-off in sales because employee salaries will surely decline. Almost any economic change has positive and negative correlations. Can we view the economy as a complicated spin glass?

To approach this interesting bit of statistical “poetry” and make sense of it, Michael H.R. Stanley and Michael A. Salinger first located and secured a database—called Compustat—that lists the annual size of every US firm. With this information, they and their colleagues calculated histograms of how firm sizes change from one year to the next.¹¹ They then made 15 histograms for each of 15 bins of firm sizes. The largest firms have very narrow growth-rate distributions—plausible because the percentage of size change from year to year for the largest firms cannot be that great. On the other hand, a tiny firm or a garage-based start-up can radically increase (or decrease) in size from year to year. Thus, these 15 histograms have widths that depend on the firm size. When this width is plotted on the y -axis of log-log paper as a function of firm size on the x -axis, the data are approximately linear over eight orders of magnitude, from the tiniest firms in the database to the largest. The width scales with the firm size to an exponent β , with $\beta \approx 1/6$.¹¹ We can therefore normalize the growth rate and show that all the data collapse on a single curve—demonstrating the scaling of this measure of firm size.

Why does this data collapse occur? Researchers are working on that. Sergey V. Buldyrev models this firm structure as an approximate Cayley tree, in which each subunit of a firm reacts to its directives from above with a certain probability distribution.¹² More recently, Luis A.N. Amaral and his colleagues have proposed a microscopic model that reproduces both the exponent and the distribution function.¹³

Hideki Takayasu and Kenji Okuyama extended the empirical results to a wide range of countries and developed still another model.¹⁴

It is not impossible to imagine some very general principles of complex organizations are at work here, because similar empirical laws appear to hold for data on a range of systems that at first sight might not seem to be so closely related. For example, instead of studying the growth rate of firms, you can study the growth rates of countries by analyzing the ratio of a country’s GDP (gross domestic product) in one year compared to its value in the previous year. The histograms of country GDP sizes appear to behave the same way as the histograms of firm sizes¹⁵ (see Figure 2), even with the same exponent value $\beta \approx 1/6$. Very recently, Vasiliki Plerou and her colleagues analyzed in the same way a database comprising research budgets of 719 US universities and found similar qualitative results, but a slightly larger exponent value, $\beta \approx 1/4$.

Instead of a firm’s size at time t (or the size of a GDP or a university budget), you might analyze the population $N_i(t)$ of a species i in successive years. Such data exist for a 30-year period for every species sighted in North America. Very recently, Timothy H. Keitt and I have analyzed this database using the same sort of techniques used to describe long-term data sets on economics and finance.² We find statistical properties that are remarkably similar, and consistent with the idea that “every bird species interacts with every other bird species,” just as the economic analysis supports the notion that “every economic entity interacts with every other economic entity.”

These empirical results are not without interest, because they cast doubt on models of economic systems—and bird populations—that partition the entire data set into strongly interacting and weakly interacting subsets and then oversimplify or ignore the interactions in the weakly interacting subset.

What can we say so far, other than just that apparently a number of natural questions in economics can be investigated quantitatively, using empirical analysis methods not unlike those used in the study of critical phenomena? And that the quantitative behavior of these complex economic systems—comprising many animate subunits—is not unlike that found in interacting systems comprising many inanimate subunits. Can we understand why methods developed in, say, critical phenomena to quantify systems comprising inanimate subunits should apparently apply to complex systems comprising animate subunits? Indeed, the conceptual framework of critical phenomena is increasingly finding application in other fields, ranging from chemistry and biology to econophysics and even liquid water. Why is this? One possible answer concerns the way in which correlations spread throughout a system comprising subunits in which “everything depends on everything else.”

The paradox is simply stated: our intuition suggests that the correlation $C(r)$ between subunits separated by a distance r should decay exponentially with r —for the same reason the value of money stored in your mattress decays exponentially with time (each year it loses a constant fraction of its worth). Thus, we might expect that $C(r) \sim e^{-r/\xi}$, where ξ , the correlation length, is the characteristic length scale above which the correlation function is negligibly small. Experiments and calculations on mathematical models confirm that correlations usually do decay exponentially. But, if the system is at its critical point, the rapid exponential decay magically turns into a long-range power-law decay: magically $\xi \rightarrow \infty$.

So then how can correlations actually propagate an infinite distance, without re-

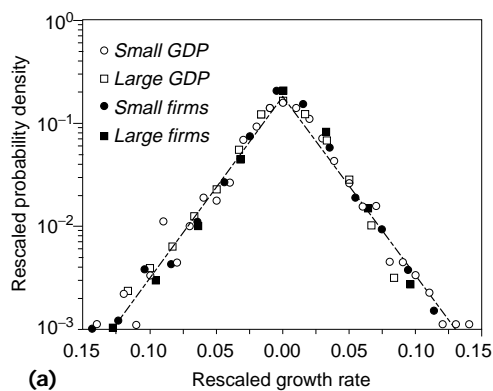
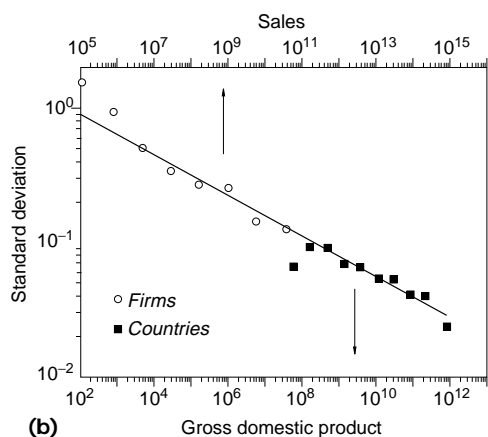


Figure 2. A test of the similarity of the results for the growth of firms and countries: (a) The conditional probability density of annual growth rates for countries and firms. All rescaled data collapse onto a single curve, showing that the distributions have indeed the same functional form. (b) Standard deviation of the distribution of annual growth rates. The standard deviations decay with size with the same exponent for both firms and countries. The size is measured in sales for the firms and in GDP for the countries. The firm data include all 4,000 publicly-traded manufacturing firms from the 19-year period 1974–1993,¹¹ while the GDP data include 152 countries for the 43-year period 1950–1992.¹⁵ (Courtesy of Youngki Lee, Luis A. N. Amaral, David Canning, and Martin Meyer.)



Acknowledgments

I thank those under whose tutelage I learned what little I understand of this subject. These include the students, postdocs, and faculty visitors to our research group with whom I have enjoyed the pleasure of scientific collaboration, especially those whose research provided the basis of this short report: Luis A.N. Amaral, Sergey V. Buldyrev, David Canning, Pierre Cizeau, Parameswaran Gopikrishnan, Shlomo Havlin, Timothy H. Keitt, Youngki Lee, Heiko Leschhorn, Yanhui Liu, Philipp Maass, Rosario N. Mantegna, Martin Meyer, Chung-Kang Peng, Vasiliki Plerou, Jeffrey Sachs, Michael A. Salinger, and Michael H.R. Stanley. I also thank Luis A.N. Amaral and Parameswaran Gopikrishnan for helpful comments on the early drafts of this manuscript.

References

1. T. H. Keitt and H.E. Stanley, "Scaling in the Dynamics of North American Breeding-Bird Populations," *Nature*, Vol. 393, 1998, pp. 257–259.
2. H.A. Makse et al., "Modeling Urban Growth Patterns with Correlated Percolation," *Physical Rev. E*, Vol. 58, 1998, pp. 7054–7062.
3. R.N. Mantegna and H.E. Stanley, "Scaling Behaviour in the Dynamics of an Economic Index," *Nature*, Vol. 376, 1995, pp. 46–49.
4. Y. Liu et al., "Correlations in Economic Time Series," *Physica A*, Vol. 245, 1997, pp. 437–440.
5. P. Cizeau et al., "Volatility Distribution in the S&P500 Stock Index," *Physica A*, Vol. 245, 1997, pp. 441–445.
6. B.B. Mandelbrot, "The Variation of Certain Speculative Prices," *J. Business*, Vol. 36, 1963, pp. 394–419.
7. S. Ghashghaie et al., "Turbulent Cascades in Foreign Exchange Markets," *Nature*, Vol.

381, 1996, pp. 767–770.

8. R.N. Mantegna and H.E. Stanley, "Turbulence and Financial Markets," *Nature*, Vol. 383, 1996, pp. 587–588.
9. C.-K. Peng et al., "Mosaic Organization of DNA Nucleotides," *Physical Rev. E*, Vol. 49, 1994, pp. 1685–1689.
10. P. Gopikrishnan et al., "Inverse Cubic Law for the Probability Distribution of Stock Price Variations," *European Physical J. B*, Vol. 3, No. 2, May 1998, pp. 139–140.
11. M.H.R. Stanley et al., "Scaling Behavior in the Growth of Companies," *Nature*, Vol. 379, 1996, pp. 804–806.
12. S.V. Buldyrev et al., "Scaling Behavior in Economics: II. Modeling of Company Growth," *Journal de Physique I*, Vol. 7, No. 4, Apr. 1997, pp. 635–650.
13. L.A.N. Amaral et al., "Power Law Scaling for a System of Interacting Units with Complex Internal Structure," *Physical Rev. Letters*, Vol. 80, 1998, pp. 1385–1388.
14. H. Takayasu and K. Okuyama, "Country Dependence on Company Size Distributions and a Numerical Model Based on Competition and Cooperation," *Fractals*, Vol. 6, 1998, pp. 67–79.
15. Y. Lee et al., "Universal Features in the Growth Dynamics of Complex Organizations," *Physical Rev. Letters*, Vol. 81, 1998, pp. 3275–3278.

quiring a series of amplification stations all along the way? We can understand such "infinite-range propagation" as arising from the huge multiplicity of interaction paths that connect two spins. The correlation between two spins along each of the interaction paths that connect them decreases exponentially with the path's length. On the other hand, the number of such interaction paths increases exponentially, with a characteristic length that is temperature-independent, depending primarily on the lattice dimension. This exponential increase is multiplied by a "gently decaying" power law that is negligible except for one special circumstance—the critical point. Right at the critical point, the gently decaying power-law correction factor in the number of interaction paths, normally negligible, emerges as the victor in this standoff between the two warring exponential effects. So, two spins are well-correlated even at an arbitrarily large separation.

Will the power laws found empirically to describe complex economic systems ever be understood in analogous terms? We will see—as the flip saying goes, "stay tuned."

H. Eugene Stanley is a professor of physics at Boston University, the director of the university's Center for Polymer Studies, and a professor of physiology at the Boston University School of Medicine. He received his BS in physics from Wesleyan University and his PhD in physics from Harvard, and honorary doctorates from Bar-Ilan University and the University of Budapest. He is a fellow of the American Physical Society and of the American Association for the Advancement of Science. He was awarded the 1997 Richtmyer Prize and the 1998 Turnbull Prize. Contact him at the Center for Polymer Studies, Boston Univ., Boston, MA 02215; hes@argento.bu.edu; <http://polymer.bu.edu/hes.html>.

WHY SEX—ARE MEN USEFUL FOR ANYTHING?

Dietrich Stauffer

SEX HAS BEEN QUITE IMPORTANT IN 1998, IF FOR NO OTHER REASON THAN TO REDUCE UNEMPLOYMENT FOR THE NEWS MEDIA NEAR WASHINGTON, DC. BUT THE FIRST SMALL ANIMALS OR PLANTS COULD HARDLY HAVE PREDICTED THIS EFFECT HUNDREDS of million years ago. Why did they deviate from the traditional bacterial behavior of replicating by splitting? Is there an evolutionary advantage of sexual over asexual reproduction?

A special “Evolution of Sex” section in *Science* magazine¹ addressed the traditional arguments but failed to review recent computer simulations of this problem, starting with Rosemary Redfield.² Any mutation that gives an advantage after thousands of generations but kills the population during the first few generations is not evolutionary realistic. Nature does not tunnel through a large energy barrier, and it did not learn multicanonical Ising simulations, nor has it invented cluster-flip algorithms. Thus each small mutation should give an advantage, or at least not too much of a disadvantage, to individuals with this mutation.

Redfield’s elegant computer model assumed that every mutation reduces slightly the survival probability (on average a child gets half of the maternal mutations and half of the paternal mutations).² If the male and female mutation rate is the same, then the child’s stationary survival probability is the same with both sexual and asexual reproduction. Nevertheless, males give birth much more seldomly than they drink beer while watching soccer, so Nature would be much better off without them. (The female tradition of eating the no-longer-required male³ has yet to be adopted widely.) Things get even worse if the male mutation rate is assumed as much higher than the female one. In this case, sexual reproduction gives much lower offspring survival rates than if females reproduced asexually for millions of years.

Obviously, Mother Nature did not read Redfield’s article, because sex is quite widespread among life on Earth. A Brazilian group found out why: Real mutations can be divided into common recessive and rare dominant mutations.^{4,5} Thus the child of a father with eight and a mother with four bad mutations will not have on average six bad mutations to reduce its survival chances. Instead, if among the many genes of a species, one of the father’s differs from the corresponding one of the mother’s, then it adversely affects the child only if the mutation is dominant. Recessive mutations affect the child only if both parents had them. Eighty or 90 percent of real mutations are recessive, perhaps because the individuals with more dangerous dominant mutations died long ago. As soon as the Brazilian group took this aspect into account in the Redfield model,^{4,5} the advantages of sex became very clear through improved survival rates, even if the male mutation rate was six times higher than the female mutation rate.

This distinction between recessive and dominant mutations is not a DNA repair mechanism, nor does it remove bad mutations by death. It is merely a cover-up—the hereditary disease that bad recessive mutations cause is stored in the child’s genes. However, it does not affect the child’s health as long as one of the two sets of genes (from both the mother and father) is still unmutated. It seems we men play

the role of back-up diskettes (for the main disk represented by the female genome)—normally useless, but helpful in case of accidental loss of information.

A more realistic simulation has to take into account the various stages in an individual’s life—this means biological aging. Most simulations use Penna’s bit-string model,^{6–8} one of the parent magazines⁹ of this new magazine reviewed its asexual version. The model was also applied to sexual reproduction.^{10,11} With sex included, the Fortran program became much longer (200 lines) but confirmed that sexual reproduction provides higher survival probabilities than asexual reproduction.^{4,5} Again, rare dominant mutations were distinguished from more common recessive ones. In line with the tradition of aging theories, all mutations were assumed detrimental. Positive mutations occur so seldomly that it took billions of years of evolution to produce Nature’s masterpiece, the German Herr Professor.

Unfortunately, this justification of male existence was questioned¹² by the suggestion that Nature should have widely adopted the meiotic parthenogenesis compromise. Here, no males are needed. Females have two sets of genes in all their cells, and they pass on to their offspring a random combination of them, just as sexual production produces a random combination of male and female genes for the child—in short, a genetic algorithm with only one sex. The information is still stored twice, and an error in one of the two sets of genes can still be covered up by the same recessive-versus-dominant-mutation trick. But now no males take food away from the females without getting pregnant. Simulations showed about the same survival chances with meiotic parthenogenesis as with sexual reproduction. So why didn’t Nature follow this simple compromise?

The answer came quite late.^{13,14} Sex gives more variety than meiotic parthenogenesis, and this greater variety aids survival after catastrophic environmental changes. Imagine another big meteor like the one that killed

the dinosaurs arrived not only in the movies but actually to Earth. It would drastically reduce surface temperature because of the increased dust and smoke levels. If all individuals alive before the crash adjusted to the same temperature, the species might die out. If, thanks to the greater variety produced by the sexual combination of different genes, different individuals of the same species adjusted to different temperatures, some minority should survive. It was quite difficult for J.S. Sa Martins and Suzana Moss de Oliveira⁹ to confirm this plausible speculation through explicit Penna-model simulation of a sudden catastrophe and compare meiotic parthenogenesis with sexual reproduction. But even sex did not help the dinosaurs survive the catastrophic climate change.

Why do women live longer than men? Cebrat¹⁵ explained this effect through dominant versus recessive mutations. Female mammals have two X chromosomes, and males combine one X with one Y chromosome. So a bad mutation in the X chromosome is usually recessive for females but is always dominant for males. Computer simulations¹⁶ confirmed this idea and detailed a male mortality twice as high as that of females, except for the highest ages where equal rights prevail. Figure 1 shows how similar this is to human reality.

Biologists have debated for years about why menopause exists. Why didn't evolution increase the population by giving women the same possibility as men, who can father children in their 80s? Or why doesn't Nature kill women after menopause, similar to the rapid death of salmon after reproduction? One reason^{4,5} why genetics cannot kill women and leave men alive after 50 years exists—which chromosome should store this genetically programmed suicide if males and females share most of the genes? (This argument only explains how menopause is possible, not why evolution prefers it.) Human civilization is




Figure 1. Sexual Penna model with nine chromosome pairs plus XX or XY. Mortality, defined as the negative time derivative of the logarithm of the number of survivors, versus age for men (diamonds, squares) and women (+, x) in West Germany around 1986 (squares and x) compared with computer simulations (diamonds and +).^{15,16}

also hardly a reason for menopause. Analogous effects are observed in most mammals, provided they are protected against hunger and predators.¹⁷ Presumably the best explanation comes from a recent Penna-type computer simulation,¹⁸ where the offspring's need for some period of child care¹⁹ together with an increasing reproduction risk that rises with age means a danger arises from a greater possibility that childbirth will kill the mother as she ages. This combination causes the maximum age of reproduction to self-organize to some intermediate age.

Besides death or the deep impact of Armageddon meteors, studies have simulated other activities more successfully, such as female sexual infidelity. Solange Martins and Thadeu Penna followed bird patterns to present simulations detailing which females should conduct affairs with older males.²⁰ This work might be a consolation for men who don't want to be regarded by these computer simulations as merely an insurance against catastrophic loss of information or change of climate. Perhaps the war between the sexes could end with Suzana Moss de Oliveira's peace proposal—males are disgusting, yet necessary.²¹ ❏

Acknowledgments

I thank, in temporal order, Naeem Jan, Thadeu Penna, Suzana Moss de Oliveira,

Paulo Murilo de Oliveira, Rita Zorzenon dos Santos, Stanislaw Cebrat, and Johannes Schneider for collaboration in my aging efforts, and Lewis Holmes for a critical reading of the first draft.

References

1. "The Evolution of Sex," special section with various authors, *Science*, Vol. 281, No. 5385, Sept. 1998, pp. 1979–2008.
2. R.J. Redfield, "Male Mutation Rates and the Cost of Sex for Females," *Nature*, Vol. 369, No. 6483, May 1994, pp. 145–147.
3. N. Vandewalle, "The Effect of Sexual Cannibalism on the Evolution of Large Populations," *Physica A*, Vol. 245, No. 1–2, Oct. 1997, pp. 113–123.
4. D. Stauffer et al., "Monte Carlo Simulation of Sexual Reproduction," *Physica A*, Vol. 231, No. 4, Oct. 1996, pp. 504–514.
5. S. Moss de Oliveira, "Ageing with Sexual and Asexual Reproduction: Monte Carlo Simulation of Mutation Accumulation," *Braz. J. Physics*, Vol. 26, No. 3, Sept. 1996, pp. 626–630.
6. T.J.P. Penna, "A Bit-String Model for Biological Aging," *J. Statistical Physics*, Vol. 78, No. 5/6, Mar. 1995, pp. 1629–1633.
7. T.J.P. Penna, A. Racco, and M.A. de Menezes, "Getting Older with Computers," *Comp. Physics Comm.*, to be published in 1998.
8. S. Moss de Oliveira, "A Small Review of the Penna Model for Biological Ageing," *Physica A*, Vol. 257, No. 1–4, Aug. 1998, pp. 465–469.
9. D. Stauffer, "Getting Older—Monte Carlo Simulations of Biological Ageing," *Computers in Physics*, Vol. 10, No. 4, July/Aug. 1996, pp. 341–348.

10. A.T. Bernardes, "Mutational Meltdown in Large Sexual Populations," *J. Physique I*, Vol. 5, No. 12, Nov. 1995, pp. 1501–1515.
11. A.T. Bernardes, "Strategies for Reproduction and Ageing," *Ann. Physik*, Vol. 5, No. 6, 1996, pp. 539–550.
12. A.T. Bernardes, "Can Males Contribute to the Genetic Improvement of a Species?" *J. Statistical Physics*, Vol. 86, No. 1/2, Jan. 1997, pp. 431–439.
13. S. Dasgupta, "Genetic Crossing vs. Cloning by Computer Simulation," *Int'l J. Modern Physics C*, Vol. 8, No. 3, June 1997, pp. 605–608.
14. J.S. Sa Martins and S. Moss de Oliveira, "Why Sex—Monte Carlo Simulations of Survival after Catastrophes," *Int'l J. Modern Physics C*, Vol. 9, No. 3, May 1998, pp. 421–432.
15. S. Cebrat, "Penna Model from the Perspective of One Geneticist," *Physica A*, Vol. 258, No. 3–4, Sept. 1998, pp. 493–498.
16. J. Schneider, S. Cebrat, and D. Stauffer, "Why Do Women Live Longer Than Men? A Monte Carlo Simulation of Penna Type Models with X and Y Chromosomes," *Int'l J. Modern Physics C*, Vol. 9, No. 5, July 1998, pp. 721–726.
17. S.N. Austad, "Postreproductive Survival," *Between Zeus and the Salmon*, K.W. Wachter and C.E. Finch, eds., National Academy Press, Washington, D.C., 1997, p. 161.
18. S. Moss de Oliveira, A.T. Bernardes, and J.S. Sa Martins, "Self-Organisation of Female Menopause in Populations with Child-Care and Reproductive Risks," to be published in *European Physical J. B*, Vol 10, Jan. 1999.
19. J. Thoms et al., "Senescence from Reproduction," *J. Physique I*, Vol. 5, No. 8, Aug. 1995, pp. 935–940.
20. S.G.F. Martins and T.J.P. Penna, "Computer Simulation of Sexual Selection in Age-Structured Populations," *Int'l J. Modern Physics C*, Vol. 9, No. 3, May 1998, pp. 491–496.
21. S. Moss de Oliveira, P.M.C. de Oliveira, and D. Stauffer, *Evolution, Money, War, and Computers*, Teubner, Stuttgart-Leipzig, Germany, to be published in 1999.

Dietrich Stauffer is a professor at the Institute for Theoretical Physics, Cologne University, Germany. Stauffer received a PhD from Munich Technical University. His technical interests include Monte Carlo simulations of statistical physics and of interdisciplinary models. Contact him at the Inst. for Theoretical Physics, Cologne Univ., 50923 Köln, Germany; stauffer@thp.Uni-Koeln.de.