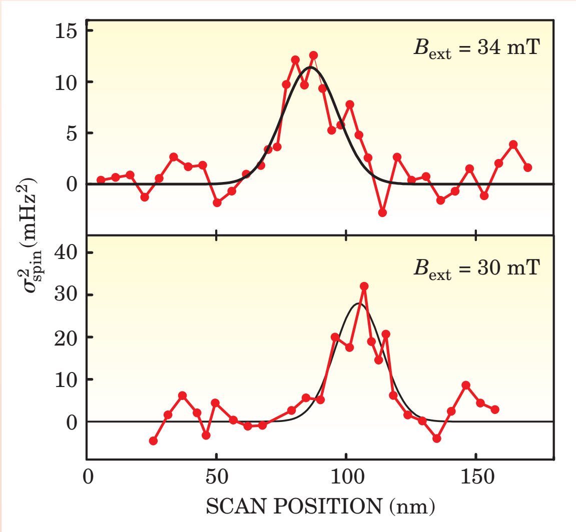

Figure 2. A peak in the spin signal reveals the position of the spin along the scan axis. Reducing the strength of the static magnetic field from 34 mT (top scan) to 30 mT (bottom scan), displaces the resonant slice and the spin's apparent position. (Adapted from ref. 2.)

netic gradient and m_{eff}, T, and τ are, respectively, the cantilever's effective mass, operating temperature, and damping time.

Since the early 1990s, various MRFM groups have steadily improved cantilever performance. Fabricated from single-crystal silicon, the IBM cantilever is about 85 μm long, 100 nm thick, and weighs 90 picograms. The tip, a lump of samarium cobalt, delivers a gradient of 2 G/nm. Experiments are done within a dilution refrigerator at about 1 K.

How the cantilever is operated is as important as its properties. When it scans a sample, the IBM cantilever is set vibrating by a feedback loop at its mechanical resonant frequency of 5.5 kHz. Any spin that lies in the path of the resonant slice will flip up and down in time with the cantilever. If, after each flip, the spin and field are left parallel, the repulsive force between the two will quicken the cantilever's vibration. If the spin and field are left antiparallel, the attractive force will slow the vibration.

That change—a jump or drop in vibration frequency when the cantilever nears a spin—is undetectable amid the noise caused by far stronger electrostatic and van der Waals forces that tug on the tip. To make the signal stand out, the IBM team turns off the oscillating magnetic field every 64 cycles for exactly half the cantilever's vibration cycle at precisely the point when the cantilever is at full swing. Deprived of resonant energy from the oscillating field, the spin misses a beat of its up-down-up-down rhythm. When the field is turned back on, an initially parallel spin will end up antiparallel and vice versa.

That syncopation introduces a weak, square-wave signal of frequency shift into the stream of data from the interferometer. Careful data analysis teases out the signal, but only after considerable signal averaging. To detect a single electron spin, the IBM team operated its cantilever for 13 hours. Figure 2 shows the result: a 5-standard-deviation signal that localizes the spin to 25 nm.

As a test sample, the IBM team created a dilute system of single electron spins by irradiating a block of glass with a weak beam of gamma rays. The gammas displace electrons and leave sparsely spaced dangling bonds called E' centers.

Not knowing where the spins were, Rugar and his team not only detected a spin but found it too. To make sure, they reduced the static magnetic field, thereby displacing the resonant slice and, with it, the separation needed between tip and spin. The result, shown in figure 2, was a shift in the apparent position of the spin.

The quantum nature of MRFM brings with it challenges and opportunities. If a spin relaxes too quickly when flipped, it can't push or pull the cantilever for long enough to produce a measurable effect. In previous experiments, Rugar had observed an apparent acceleration of spin relaxation when the tip neared the sample surface. Uncorrected, the effect would have prevented the detection of a single spin. Fortunately, a quantum mechanical analysis by several groups revealed the culprit, the cantilever's high-order vibrational modes, and a solution, modifying the tip.

Although biomolecules were the original inspiration for developing MRFM, other groups are working to apply it to other systems, notably the buried, irregular interfaces found in polymer systems. There, the absence of order means that a single small set of numbers, derived from diffraction, can't fully describe the material's structure. An atom-scale image would be ideal.

But to reach either goal—polymers or proteins—requires detecting protons and the nuclei of carbon, oxygen, and so on. Having improved the sensitivity of MRFM by 12 orders of magnitude to see one electron spin, Rugar is optimistic about obtaining the next three orders of magnitude to see a single proton spin.

Charles Day

References

1. J. A. Sidles, *Appl. Phys. Lett.* **58**, 2854 (1991).
2. D. Rugar, R. Budakian, H. J. Mamin, B. W. Choi, *Nature* **430**, 329 (2004).

Removing Nodes in a Network Can Protect Against Cascading Failures

Can a mathematical idealization of network behavior be adapted to real systems?

In typical networks, some nodes are more connected than others. On the Web, think Google; on airline route maps, think O'Hare International. Though intricate and tangled, the topological structure of many such networks exhibits scale-free properties, because added nodes connect preferentially to others that are already well connected. The distribution in the number of links scales as a power law—most nodes are connected by links to only a few other nodes, but a small number are connected to thousands or millions.

Four years ago, Albert-László Barabási and colleagues at the University of Notre Dame found that networks organized that way are extremely robust against random failure. But that error tolerance comes at a price: These networks are also extremely vulnerable to failures or targeted attacks on the most highly connected nodes.[1]

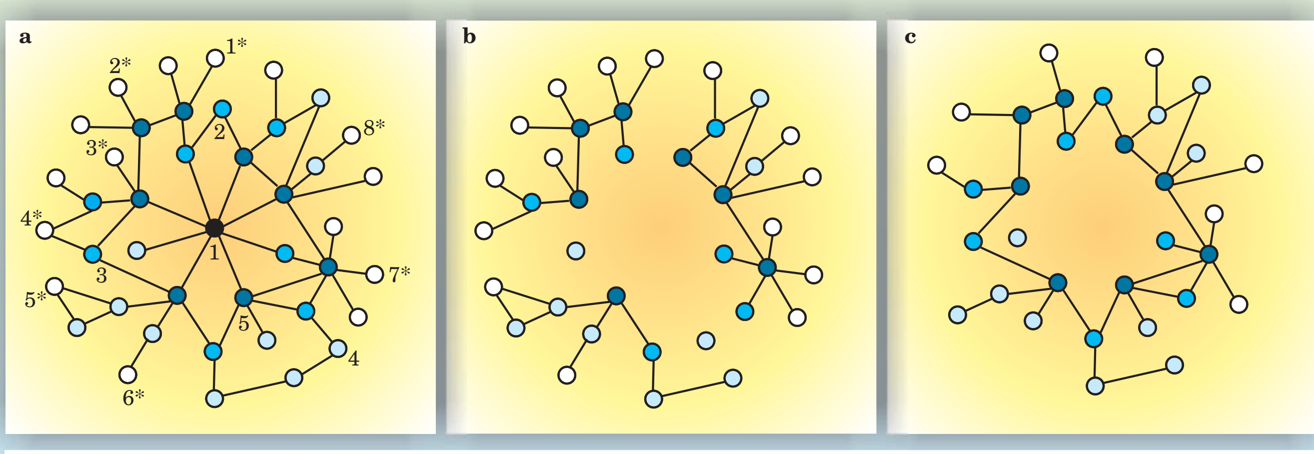

When a network of 42 nodes, panel **(a)**, suffers a breakdown in one of its key nodes (1), the load redistribution may overload others and cascade through the network—killing nodes 2, 3, 4, and 5—and break the network into clusters or islands **(b)**. As a precautionary act to protect the system after node 1 has failed, removing peripheral nodes 1*, 2* . . . 8* **(c)** preserves 31 connected nodes in the network. (Courtesy of Adilson Motter.)

If a network is dynamic and carries a physical quantity—such as e-mails or airplanes—specific failures can be catastrophic. Removing even a single, carefully chosen node may trigger a large-scale cascade that overloads or kills subsequent nodes in the network. One devastating cascade occurred on 14 August 2003, when an instability in the North American power grid prompted a generator in Ohio to fail. That event spawned new instabilities in the grid, shutting down an aggregate 62 000 MW load to 50 million people in the northeastern US and parts of Canada.

Building firebreaks

Using an idealized model of a scale-free network, mathematician Adilson Motter from the Max Planck Institute for the Physics of Complex Systems in Dresden, Germany, has devised a strategy to mitigate the extent of the damage early in the process before the cascade can fracture the network.[2] The model's basis is simple: Each node in the grid serves as a source, destination, and typically a conduit of some quantity—a packet of information or a signal, perhaps—exchanged between every pair of nodes and transmitted along a path that connects them. Each node also is assigned a finite capacity for what it can handle, the "load" being the amount of information or signal that passes through it per unit time. The scale-free nature of the model guarantees a heterogeneous distribution of loads; some nodes are positioned to handle more load than others. Removing one (or more) of the highly connected nodes forces the load to redistribute among the others that remain and can trigger a cascade, thus overloading other nodes down the line. When overloaded, a node is removed from the network.

Motter's prescription to halt the cascade's size and spread can be likened to what foresters do when they deliberately set up firebreaks to control the spread of burning forest. If certain nodes trigger the problem, then intentionally removing others may cure it. But which ones? Networks are more intricately arranged than forests.

The goal is to identify the nodes whose removal would lessen the load on the most central, connected ones. Shutting down the least connected nodes should do the job best; that is, preserve a large connected component of the network. Specifically, an operator chooses those nodes by calculating at each stage in the cascade which ones have the largest disparity between the load they generate and the load they channel. The removal of those peripheral nodes alleviates the load in the whole network—they're no longer expecting or sending out a packet—with little penalty to connectivity because the nodes had channeled the least traffic. Other parts of the system would then pick up the slack.

Made more resilient, the network can then handle the loss of key nodes as shown in the figure above. Alternatively, removing well-chosen links between nodes could also redistribute load to effectively preserve the network's connectivity. In simulations using a network of 5000 nodes, Motter showed that the largest connected component remaining after a cascade began was 10 times greater if his prescription was followed than if the cascade was left untouched.

Although the model was designed to simulate a simple idealized case of network breakup, one can reasonably ask whether it may be adapted to realistic cases.

Real systems

Networks do not have to be scale-free to undergo cascading failure. Power grids are structured networks, for instance. What makes Motter's model potentially relevant are the heterogeneous loads he sets up in a dynamic system. In a real grid, generators produce power that transmission lines distribute to load centers formed by aggregates of consumers. Most nodes are either generators or consumers, but not both. In the event of a cascade, shutting down selected consumer nodes and permitting local blackouts on the ends of the grid can sometimes make sense.

MIT's George Verghese wonders if Motter's model is too idealized. "You don't have anything like a packet going from each node to every other node in power systems," he says. And, he adds, there are real capacities on transmission lines (or links) with potential-driven flows through them that can lead to counterintuitive redistributions of power as the system responds to changes and follows Kirchhoff's circuit laws. Moreover, blackouts occur not because nodes are overloaded specifically, but because generators and transformers are hardwired to save themselves in response to a drastic change. Motter's model does account for both global and local effects, and the strategy of playing with topology to influence stability is one that interests power-grid designers. But Verghese questions whether, once the rules are rewritten to reflect more closely the constraints of power flow, the prescription wouldn't change.

The model may be more directly relevant for the internet, where information flow is bidirectional. Worms or virus attacks can certainly infect a communications network, causing router congestion that leads to denial-of-service errors; these behave in a way analogous to overloading, though no nodes are really killed.

The extent to which some nodes are more vulnerable than others is controversial. Many network theorists argue that the internet is scale-free on the level of autonomous domains (university.edu domains, or huge service providers like AOL or MCI). But AT&T network researcher Walter Willinger cautions that this scaling perspective is a red herring. "Technologically, routers in the core of the network just cannot accommodate large numbers of high-bandwidth connections." The core networks of large service providers really consist of a large number of high capacity routers, for instance, that may be widely spread around the world. "It's hard to imagine taking out a node . . . by effectively attacking all the core routers as a group" he says.

Motter's method of cascade control does not depend on any particular network structure or details of the flow. The task of catering to specific requirements still remains. In fact, he is now adapting his model to the much different system of social networks, where each node—a person—can influence just a few others or up to tens of thousands (see PHYSICS TODAY, September 1998, page 17). Imagine that a rumor starts somehow and you'd like to shut down its spread. The idea would be to prevent a transition between a phase in which few are persuaded and one in which many are, without fracturing the social network. The algorithm for such an action might be one that politicians and their spin doctors find useful.

Mark Wilson

References

1. R. Albert, H. Jeong, A.-L. Barabási, *Nature* **406**, 378 (2000).
2. A. E. Motter, *Phys. Rev. Lett.* (in press).

New Experiments Highlight Universal Behavior in Copper Oxide Superconductors

Similar patterns of spin fluctuations and charge ordering in different compounds suggests that these features are intrinsic to the cuprates.

Since 1986, when lanthanum barium copper oxide (LBCO) was found to superconduct above 30 K, many other such high-critical-temperature (high-T_c) superconductors have been found, all of them copper oxides having layered structures. The superconductivity appears to arise from the planes of copper and oxygen atoms common to these compounds.

The undoped parents of all high-T_c superconductors have one valence electron per copper atom. The ground state is an antiferromagnetic insulator, in which the spin of each electron is aligned opposite that of its neighbors. Once the compound is sufficiently doped with holes or electrons, it goes superconducting below T_c, and the electrons move in coherent pairs throughout the sample. Just how the material makes the transition between these very disparate states has been the subject of intense study.

Of particular interest is the nature of any spin or charge ordering, either in the fully superconducting state or in the underdoped state. When the material is underdoped but above T_c, it's in the so called pseudogap region, where the material is in some intermediate state between antiferromagnetism and superconductivity. Do the charges, for example, arrange themselves in some regular pattern? Is that pattern stripelike, or more like a checkerboard? Probing the charge or spin structure might elucidate the underlying order and perhaps answer whether magnetic order competes with superconductivity or is essential to it.

The latest revelations about spin or charge order have come largely from complementary types of experiments. Scanning tunneling microscopy (STM) measures the static density of electrons, thereby giving a satisfying visual image of how the charge is ordered. But it can only see the surface. By contrast, neutron scattering is sensitive to the spins and not the charge. It's a probe of the dynamical spin excitations throughout the bulk of a sample.

Ideally one would like to apply both methods to the same material. Such attempts have been largely stymied, however, because some materials lend themselves more readily to one type of measurement than to another. It's hard to tell whether the data reflect individual or universal properties.

STM has primarily probed the surface charge order in crystals of bismuth strontium calcium copper oxide ($Bi_2Sr_2CaCu_2O_{8+\delta}$, or BSCCO for short) because those crystals cleave so cleanly. Neutron scattering studies have been done on yttrium barium copper oxide ($YBa_2Cu_3O_{6+x}$, or YBCO for short) and on lanthanum strontium copper oxide ($La_{2-x}Sr_xCuO_4$, or LSCO), which one can grow in sufficiently large crystals. Neutron scattering has been hard to do on BSCCO because of the small size of available single crystals, and STM studies of YBCO have been thwarted by the lack of a clean surface.

Broadened perspectives

A spate of experiments reported in the past six months has now enriched the picture regarding charge and spin order. Some of the experiments were done on materials that had not previously been available, and some used new or improved techniques. As a result, patterns previously seen in one member of the copper oxide family are showing up in other compounds, or under different conditions.

In the case of STM studies, evidence has been accumulating in the past few years for a checkerboard-like pattern of electronic modulation in BSCCO.[1–3] The most recent such evidence comes from an experiment done earlier this year[4] that found charge order in the pseudogap region by probing an underdoped sample above T_c. Last month, as seen on the cover of this issue, another experiment revealed[5] static checkerboard charge order in a very different cuprate, known as sodium-doped calcium copper oxychloride ($Ca_{2-x}Na_xCuO_2Cl_2$, or Na-CCOC). In both the new STM studies, the charge order was found under conditions that suppress superconductivity.

As for neutron scattering experiments, recent work has shown, contrary to naive expectations, that the energy dependence of the magnetic excitations in superconducting YBCO is similar to that of a member of the lanthanum copper oxide family ($La_{1.875}Ba_{0.125}CuO_4$) that is not superconducting but does have a stripelike arrangement of spins and charges.[6,7] Other experiments[8,9] concur on the behavior of the spin excitations in copper oxide compounds, although researchers disagree on the finer details.

The STM and neutron scattering experiments have broadened our knowledge of high-T_c materials, but it's not clear how their separate findings are related to one other. Only