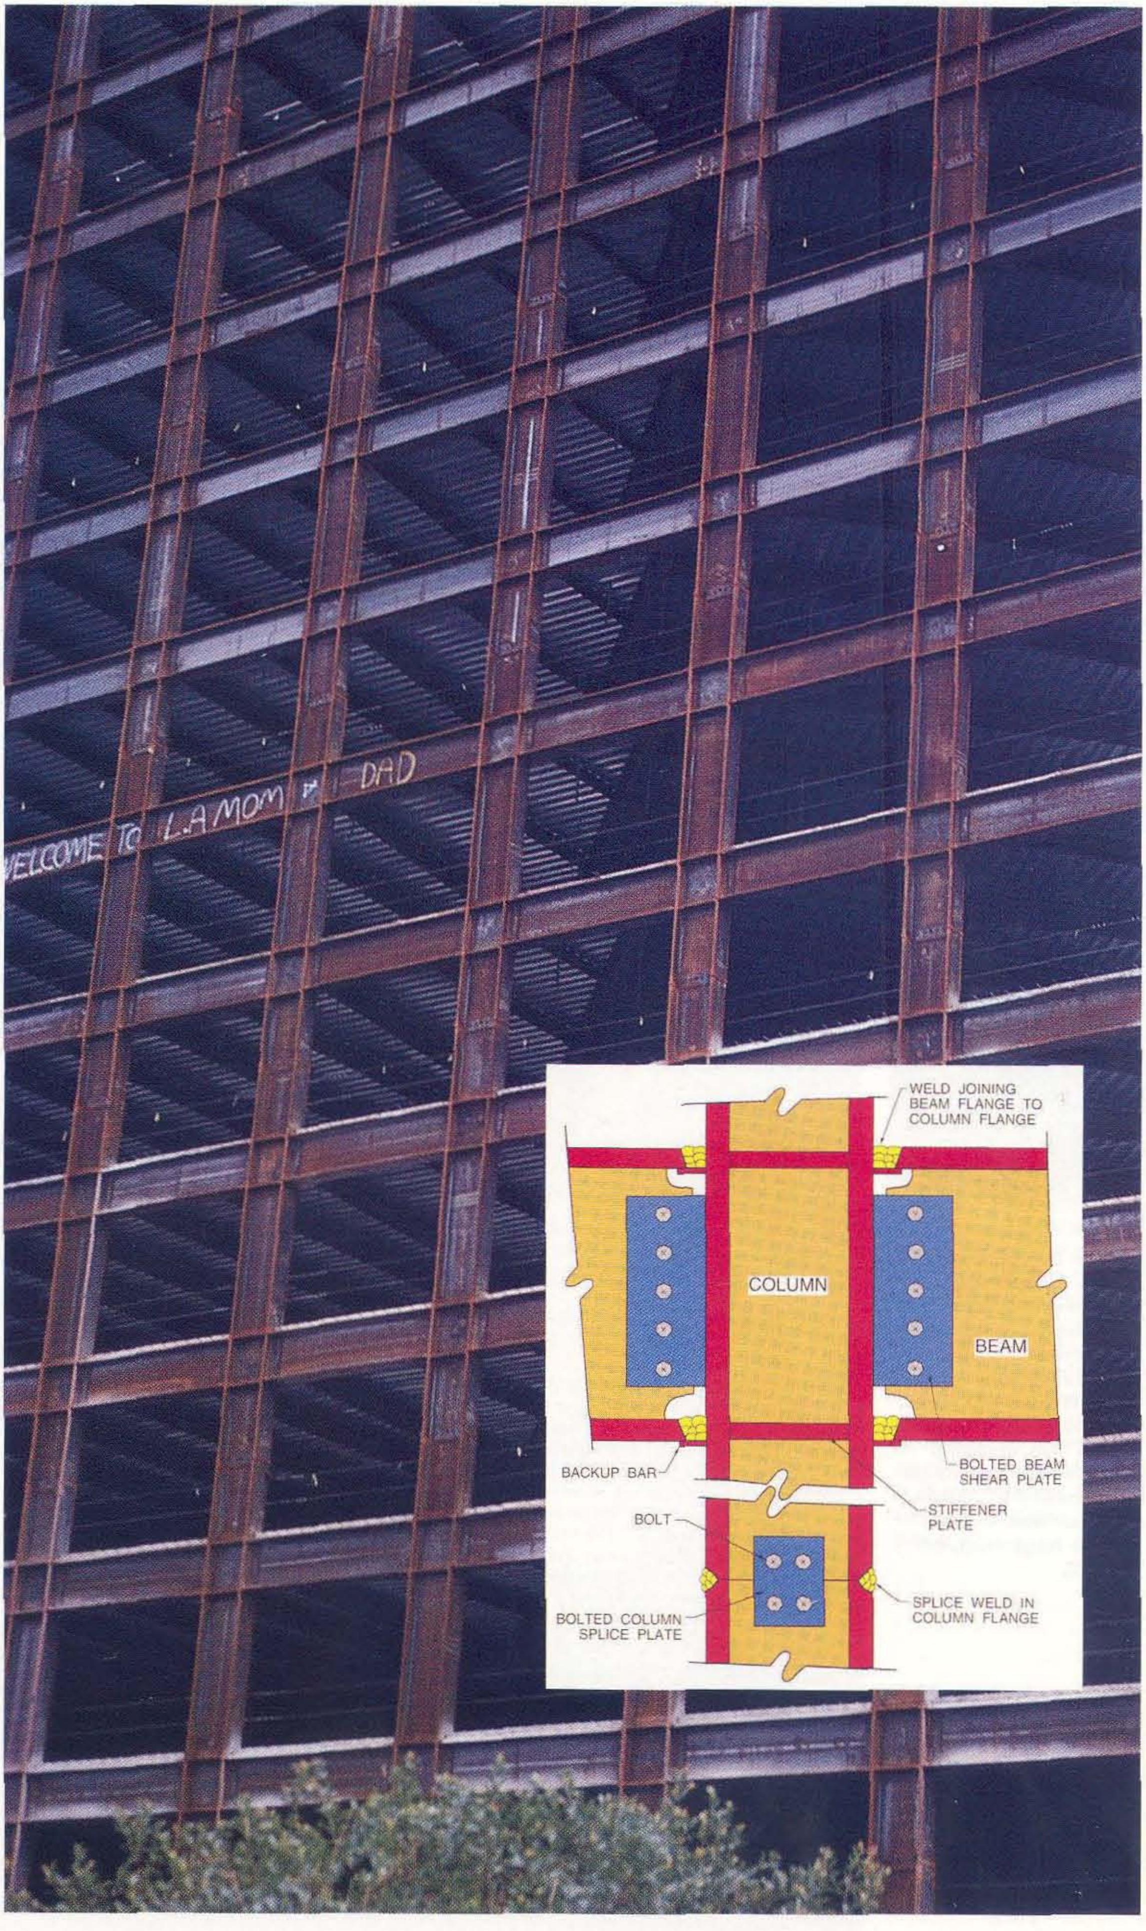
ELCOME TO L.A MOM
DAD
WELD JOINING BEAM FLANGE TO COLUMN FLANGE
COLUMN
BEAM
BACKUP BAR
BOLTED BEAM SHEAR PLATE
STIFFENER PLATE
BOLT
SPLICE WELD IN COLUMN FLANGE
BOLTED COLUMN SPLICE PLATE

Tall Buildings, Bad Welds, Large Earthquakes—Big Problems

by John F. Hall

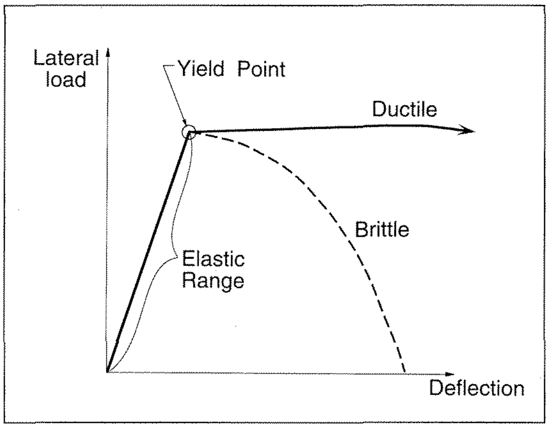

Opposite: The steel frame gives this building its strength and stiffness. (The building is actually in Pasadena, incidentally.) The inset shows the frame's construction in detail. The columns and beams—collectively known as members—are I-shaped in cross section. The web (orange) lies in the plane of the page. The perpendicular flanges (red) resist bending. The beam's web is bolted to the column (blue plates with gray bolts), and the beam's flanges are welded (yellow) to the column's flanges. The backup bars form troughs that contain the molten weld material.
Above: The larger the sideways force applied to a building, the farther it moves, as shown in this load-deflection curve. If the force is less than the yield point, the building will spring back elastically. At larger loads, something gives—either ductilely, in which case the members yield but don't break, or brittlely, in which case they crack.

I'll start by being up front with you: last night I spent most of the evening trying to glue my glasses back together, so you might say I'm not too hot on metal structures at the moment. Having gotten that off my chest, let me begin by giving you a brief introduction to earthquake engineering. In a steel-frame building, the frame supports not only the weight of the building—a vertical load—but also withstands lateral loads from winds and earthquakes. These lateral loads cause the frame members to bend, and the engineering term for the action that causes bending is "moment." Hence these frames are called *moment frames*, or *moment-resisting frames*. The frame consists of vertical columns and horizontal beams, and in order to transfer the bending moments between these members, we need to have very strong connections—usually made with welds.

Now, if you apply a lateral force to a building, it will displace sideways in response. Engineers plot this behavior in a load-deflection curve, such as the one above. In the curve's elastic range, from zero load up to the elastic limit, or yield point, you can apply a load on and off and the building always springs back to its original position—it behaves elastically. At loads above the yield point, the building no longer behaves elastically. The postelastic behavior can be ductile, which means that the members deform—they stretch like chewing gum—but maintain the strength of the building. Or, like my glasses, the behavior can be brittle—as the deflection increases, there's a loss of strength as something snaps. Whenever possible, it's best to design structures to have enough strength to carry their loads in the elastic range to avoid the damage associated with yielding. (For example, airplanes are designed to behave elastically while airborne.)

Wind is one lateral load to be considered when designing a building. The wind exerts a sideways pressure on the building, and engineers understand this force pretty well. They treat wind as a constant pressure, and even though the pressure is significant, it's possible and economical to design the building to withstand it in the elastic range. This is fortunate, because if a windstorm came up strong enough to make the building yield, the steady pressure would actually push it over.

An earthquake, like the wind, causes a building to deflect sideways. But unlike the wind, an earthquake is a back-and-forth action. It reminded the ancient Japanese of how a landed fish wiggles, so in their legends, a giant catfish causes earthquakes. This giant catfish can make the ground move pretty violently, and so earthquake loads are larger than wind loads—in fact, it's not economically possible to design a building to respond elastically to a strong earthquake. That means the building is going to yield. How can we get away with that? How can we be sure that the building won't collapse when it yields in a strong earthquake? The answer has to do with the back-and-forth nature of the ground motion. Say the ground moves to the left, causing the building to start to yield to the right. Then, before the building has time to collapse, the ground moves back to the right and gets under the building again, and so on. You can actually try this at home—walk up behind somebody, give him a shove, and before he falls on his face,

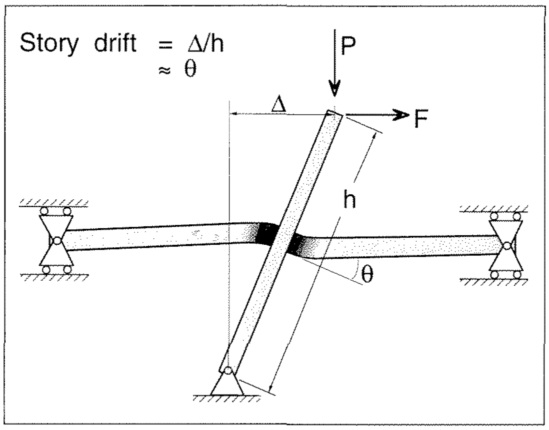

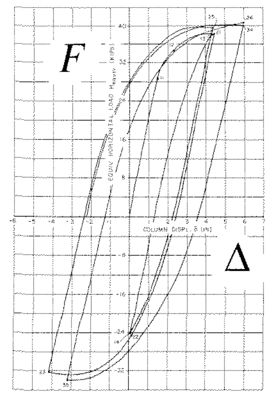

Right: Since shaking an entire building on demand is impractial, engineers use a mock-up of a single beam-to-column joint, plus the adjoining halves of the members surrounding it, as a proxy. The column's base is fixed to a pivot and the beam ends are on sliders, closely reproducing the conditions this subassmebly would feel in a building during a quake. Two loads are applied to the top of the column—a vertical load, P, which represents the building's weight, and the back-and-forth horizontal earthquake force, F. The story drift is determined by dividing the resulting deflection, Δ, by the story height, h. The columns are stronger than the beams, so once the elastic limit is exceeded, the beam kinks where it joins the column; this kink angle (θ) is approximately the story drift.

Far right: A typical force-deflection curve from such an experiment. As in the idealized curve on the previous page, the force (F) is plotted vertically and deflection (Δ) horizontally. But here the force is appiled back and forth, over and over again.

run around to the front and push him back.

This explanation's not quite good enough for engineers, without some calculations to verify that it's possible. So back in the 1960s and 1970s, engineers invented computerized methods to calculate the responses of buildings to earthquakes. These mathematical models were pretty simple, and assumed that the buildings would behave in a ductile manner. The engineers used the ground-motion records that were available at the time, and were thought to be representative of strong ground shaking, for the inputs. This led to two conclusions.

For one, if the building has to yield, it's much better to have the yielding occur in the beams than in the columns. So the engineers started making the columns stronger than the beams. The yielding then showed up as kinks—like in a wire that's been bent too hard—at the ends of the beams where the bending moments are highest. This was good, because the columns held and the building stayed up. The computer programs could also predict the amount of yielding in the structure. I'll quantify that for our purposes by something called "story drift," which is the sideways movement in a story divided by its height from ceiling to ceiling.

This led to the second finding—the engineers calculated that a reasonable story drift for the earthquakes they were using was about 1.5 percent, or a lateral deflection of two inches per 10-foot story. (A building begins to yield at about 0.4 percent, so most of this story drift actually occurs in the yield range.) So they then had to determine whether the actual materials used in a building—the steel beams and columns—could take this kind of drift without losing strength after yielding. In other words, did the members have sufficient ductility?

The only way to determine something like that is in the laboratory, and the easiest method is to build a small piece of the building and apply forces to it to reproduce what it would feel if it were a part of the building during a strong earthquake. Then we measure the story deflection, and the story drift is determined by dividing that number by the story height.

Above is an actual force-deflection curve from such a setup, taken from a report written back in the early 1970s. The curve's bending toward the horizontal is due to the yielding. You can see that the assemblage yields in first one direction, then the other, but you don't see much degradation in strength as the cycles continue. That's very good. That's ductility—the strength is being maintained as the material yields. And if we convert the deflections from this test into story drifts, we get about 4 percent, which is greater than the needed 1.5 percent. So things looked pretty good—the engineers considered their designs to be validated, and the building code was written accordingly. It's important to note that the code is essentially a life-safety document, whose goal is to preserve lives by avoiding building collapses. The code is not intended to prevent damage to buildings.

Now, in the Northridge earthquake, the engineers got a terrible shock of their own—the welded connections in many steel buildings fractured. The fact that the welds failed means

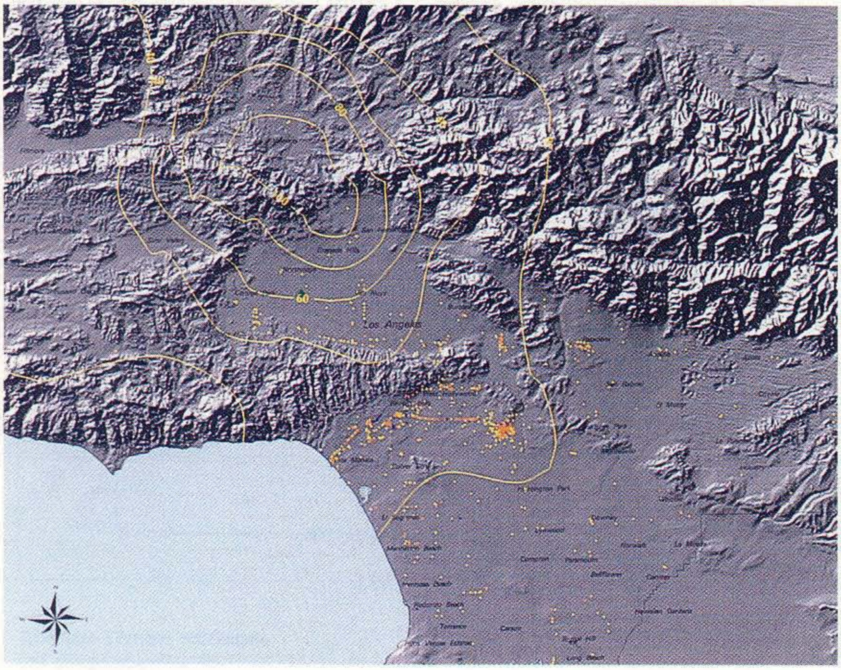

The contours on this map of the L.A. area show peak ground velocities during the Northridge earthquake in centimeters per second. (The green triangle marks the quake's epicenter.) The dots show the locations of steel-frame buildings, as gleaned from the county assessor's records. Red dots are high-rises (six stories or taller), yellow dots are one- to five-story structures, and blue dots are buildings whose height was not recorded. Map prepared by the California Office of Emergency Services.

that these buildings are not as ductile as we thought—they're more on the brittle side. (Remember that ductility is the foundation of our design philosophy.) Furthermore, many welds failed well within their elastic range. Because they never reached yield, the designed strength of those members was never achieved. Now, one optimistic point of view says that since the code is a life-safety document, and since Northridge was a pretty good shake and none of the steel buildings fell down, the code was a success. Sure, we had some damage, but the code really doesn't try to prevent damage. This view is actually still held by some engineers, but you can make a couple of points against it.

Many welds failed well within their elastic range. Because they never reached yield, the designed strength of those members was never achieved.

First, the buildings really didn't get shaken all that hard. In the map above, the dots represent steel buildings, and the contours are the peak ground velocities in the Northridge earthquake. (Peak ground velocity is probably the best single parameter for gauging the damage potential of an earthquake, because even a large acceleration, if applied for a short duration, may not be sufficient to get the building to move.) The map shows that the most damaging ground motions occurred in the Santa Susana Mountains to the north, where there are very few steel buildings—or other buildings, for that matter. So most of the steel buildings got only moderate shaking.

Which leads to the second point: the way in which the code represents an earthquake is deficient. We soon realized that, even for this moderate earthquake, the ground motions and attendant high ground velocities to the north of the epicenter were larger than anticipated by the building code. The records that the engineers used to validate their design procedures back in the 1960s and 1970s didn't show any such velocities. It can be seen in retrospect that California simply wasn't densely instrumented enough back then to catch them. Most of the earthquakes the engineers used, such as the 1940 El Centro (magnitude 6.9) and the 1952 Kern County (7.5), occurred in rural areas where there weren't many strong-motion sensors. The 1971 San Fernando earthquake (6.7), which shook urban Los Angeles, did in fact register a ground velocity of 113 centimeters (about four feet) per second at nearby Pacoima Dam. But this sensor was atop a steep ridge, which was blamed for the strong motions, and so this velocity was discounted as being inapplicable to what a building in the flatlands might feel.

In summary, then, the building code is supposed to be written for larger earthquakes than Northridge, yet the code didn't anticipate the ground motions felt even in this moderate quake. Furthermore, the welds failed in buildings that didn't get the strongest shaking that Northridge had to offer. What does this tell us about what's going to happen in larger earthquakes? I'll come back to that, but first let's take a closer look at what did happen in the Northridge quake.

Most of the steel buildings that were shaken in the Northridge earthquake look fine from the outside. (Remember that no steel buildings collapsed, although other, weaker structures did.) But if you go inside, and uncover some of the beam-to-column connections (which is a lot of

In some cases, more than 50 percent of the welded connections are broken; in a few buildings, nearly every connection has given way.

Right: Exposing the beam-to-column connections so engineers can inspect them means cutting through drywall, stripping away insulation, and sometimes dealing with asbestos.

Below: This beam's lower flange is completely severed where it joins the column. The earthquake also sheared off some of the original connecting bolts—these are replacements. Practically all of the welded joints in this building had something similar happen.

work, by the way), you'll see things like the photo at left. The flanges, which carry most of the bending moment, are cracked clear through at the welds. The cracks sometimes extend into the web of the beam or column, and, very occasionally, the member is torn in two. We know that this problem exists in about 100 or so buildings. In some cases, more than 50 percent of the welded connections are broken; in a few buildings, nearly every connection has given way. And there are perhaps another 200 suspect buildings that we haven't really looked at yet.

Why did this happen? Remember that we confine the building's yielding to the beams, causing them to kink at their ends, which is exactly where the welds are. So the welds were highly stressed, and they didn't hold up. Why not? There are at least four reasons. First of all, quality control, to put it bluntly, is often not very good as these buildings are built. There simply aren't enough building inspectors for the volume of construction, and some contractors just aren't well-educated in the importance of following the code—they either don't have the specs on hand at the job site, or they don't follow them. And buildings aren't like airplanes, which provide a good reading really quickly the first time a test pilot takes one up. A badly built building can stand for quite a while before its weaknesses are revealed in an earthquake. So the welds that fractured probably had lots of small defects to begin with. Second, the material used for the welds is not very fracture-resistant. No one was expecting brittle fracture to be a problem, so why pay more for fracture-resistant material when the need is not apparent? Third, there was little or no heat treating done during the welding, which means that the welds cooled very fast, and that tends to embrittle them. The more slowly a weld cools—if you put an electrically heated blanket on it, for example—the more ductile it will be. And finally, the backup bar—which helps retain the molten material as the weld cools—often didn't fuse completely with the column. That gap between the bar and the column often became the notch where the crack started.

One might reasonably ask why the laboratory tests didn't pick this kind of thing up. There are multiple reasons here, too. For one, the tests were generally done at small scales—say, one-third scale—and at slow loading rates, because there wasn't enough money to buy the large equipment and fast actuators necessary to give full-sized connection specimens the shaking they would really feel in an actual earthquake. Also, the quality control on the laboratory welds that the researchers made was probably a lot better than it is at the construction site. These factors worked together to make the test results better than, and not a fair indication of, what might happen in the field. However, if you go back through the old laboratory reports, you do find a fair number of premature fractures caused by the weld-fracture problem, even in those small-scale specimens. The researchers, when asked about this after Northridge, said, "Well, it's all in the reports," and the engineers replied, "We don't have time to read your reports. Why didn't you yell and scream about it?" And so it goes. It's human nature.

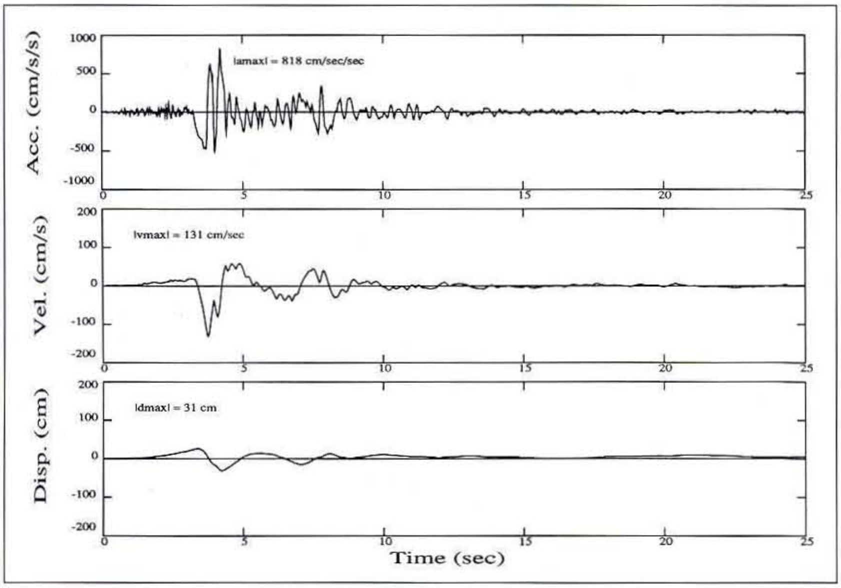

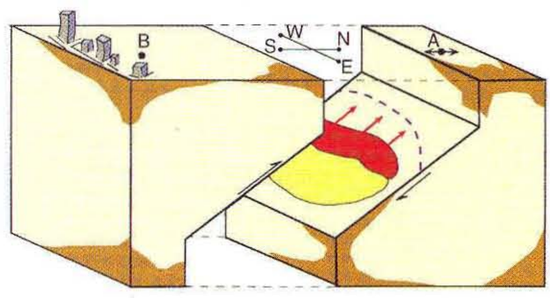

Above: These horizontal displacements, velocities, and accelerations (bottom, middle, and top traces, respectively) were recorded near the Olive View Hospital in Sylmar during the Northridge quake. |dmax| stands for peak displacement, |vmax| is peak velocity, and |amax| is peak acceleration.

Right: The slip-pulse mechanism tends to focus an earthquake's energy. In the Northridge quake, a south-dipping thrust fault (a fault where one side overrides the other instead of slipping by sideways) ruptured at its base. The slip pulse propagated upward and to the north. At the instant of the sketch, the slip pulse is rupturing the red region and is moving up-fault (red arrows). The yellow region has finished slipping. The slip pulse feeds energy into the shear wave traveling ahead (dashed purple line), which will eventually reach the surface near point A. Thus the region to the north experienced more damaging ground motions than did the built-up area to the south, or even the epicenter (point B).

The Federal Emergency Management Agency is now funding a research program to try to find a solution. Phase I, which I was involved in, is just wrapping up, and Phase II is about to start. The first thing the task force did was investigate the scale effect by testing more nearly full-sized connections in the higher-capacity rigs that are now available. And although the task force improved quality control—they used better weld material, ground off the backup bar, and did heat treatments—the cracks appeared, so it seems that our fundamental design was bad. So we're now trying to reduce the stress the welds must carry by welding cover plates over the joints. The cover plates strengthen the connection of the beam to the column, forcing the yielding out into the beam where there's no weld to break. This method has had some successes, although there are still problems that we hope Phase II will solve. I might add that the solution, when one is found, is liable to be pretty expensive.

In the meantime, some of the buildings damaged in the Northridge earthquake still sit vacant, waiting for a solution to emerge. Others have been torn down. But the majority of building owners can't afford to let their real estate sit idle indefinitely, and are fixing their buildings one way or another. In the absence of a definitive solution, the city of Los Angeles has issued its own guidelines for building rehabilitation, essentially saying, "If you take these suggested actions we'll approve your plans expeditiously now, so that you can put your building back in use, but we may require you to do more things later on."

Now let's turn to the ground-motion side of the equation. Above is a record of the ground motion felt in Sylmar during the Northridge earthquake, in a region of strong shaking to the northeast of the epicenter. It shows pretty high accelerations, which are a concern, but I want to focus on the rapid displacement—a roughly 60-centimeter (about two feet) peak-to-trough pulse that happened in less than a second. That kind of motion has a very high damage potential, and it simply wasn't present in the old ground-motion records that the engineers used when they were validating the design procedures.

These large, rapid displacements are what seismologists call "near-source directivity effects"—a very important idea that I want to discuss in some depth. Over the last decade, Professor of Engineering Seismology Tom Heaton (PhD '78) and his colleagues at the U.S. Geological Survey (USGS) and Caltech have discovered some very interesting things about how a rupture proceeds on a fault—namely that, at any given instant, only a small part of the fault is involved in the slip. The slip actually takes place in a pulse that propagates along the fault, as shown above, and the amount of slip within this pulse is quite large. Now, the fault's slip produces shear waves that travel out in all directions. Since the slip pulse travels at a slightly lower speed than the shear wave (a fact also discovered by Heaton, et al.), each successive bit of fault slip contributes more energy to the part of the shear wave being sent out ahead of the rupture, building the wave up to a very large amplitude. So, in general, the largest ground motions are going to be observed in areas toward which the fault is rupturing.

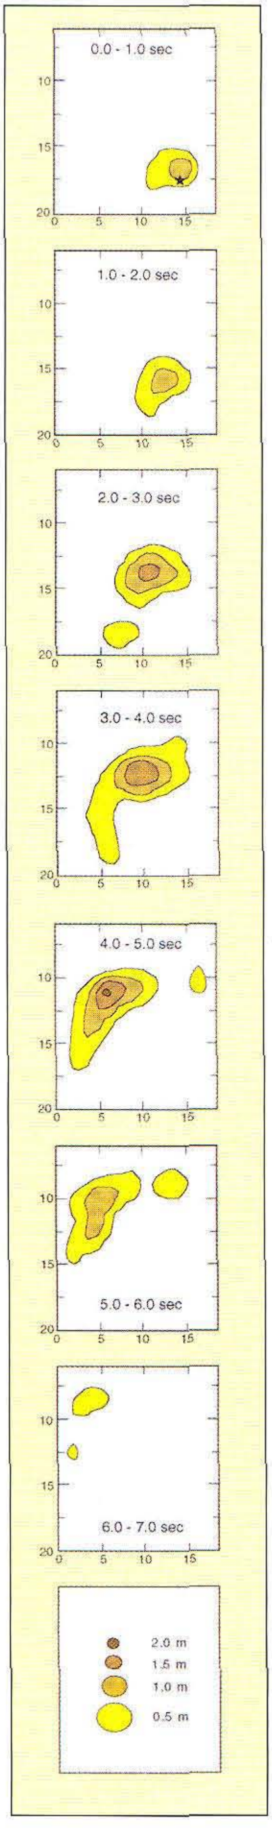

Left, top to bottom: The Northridge quake in one-second intervals, as seen from a vantage point above and perpendicular to the fault plane. The axes are marked in kilometers. The star in the first panel plots the earthquake's hypocenter, or point of origin. From there, the slip pulse travels northward and toward the surface. The darker the color, the larger the slip in meters during that interval, as shown in the bottom panel.

Right: The colored zones are L.A.'s main thrust faults. The sawtooth lines mark the faults' upper edges; those that reach the surface have black teeth. The black blobs represent earthquakes this century (labeled with their year and magnitude). The figures in SHADOWED TYPE show the size earthquake that could happen if an entire fault broke at once, and the recurrance interval in years for that quake. Abbreviated fault names: SSF= Santa Susana, MCF= Malibu Coast, HF= Hollywood, RF= Raymond, C-SF= Clamshell-Sawpit, SJF= San José, COMP= Compton.

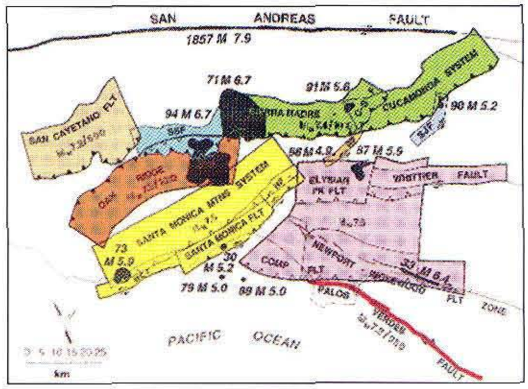

But the fault slip is deep underground—how do the seismologists know what's going on down there? They solved what's called an inversion problem. They took strong ground-motion data and geodetic data—surveyor's measurements of surface displacements caused by the earthquake —and back-calculated what must have happened down there in order to give the observed motions up here. Heaton and Dave Wald (PhD '93) of the USGS developed a lot of the methodology used in those calculations, and also generated the set of images at left, which show the Northridge earthquake from start to finish at one-second intervals. The slip pulse's passage along the fault is clearly visible.

As I said, the Northridge earthquake was only a magnitude 6.7, yet it created stronger ground motions than are represented in the code. But we have even larger earthquakes in California. The San Andreas and the Hayward faults, which are capable of generating large earthquakes, pass close to some of our major cities, which means that we can have very strong near-source effects within our metropolitan areas. This is of real concern. What about Los Angeles?

You may be surprised to learn that in the 1920s, the seismic threat to L.A. was quite a lively topic. Robert Hill, a well-known geologist at the time, wrote a book on the subject. He was so proud of his conclusion that he put it on the cover: "This book completely refutes the prediction… that Los Angeles is about to be destroyed by earthquakes. It proves that this area is not only free from the probability of severe seismic disturbances, but has the least to fear from Acts of God of any city under the American flag." I won't talk about the fires and floods we've had of late, but I can say something about earthquakes in the Los Angeles region.

Measurements by many people, among them Ken Hudnut (a Caltech postdoc now at the USGS) and Andrea Donnellan (MS '88, PhD '92) of Caltech's Jet Propulsion Laboratory, have documented a north-south compression of the Los Angeles region by about one centimeter per year, which is thought to arise from the bend in the San Andreas fault to the city's north. Last January, eight geologists associated with the Southern California Earthquake Center, including Jim Dolan (a Caltech postdoc now at USC), and Caltech Professor of Geology Kerry Sieh as lead authors, published a paper that assumed that this compression is accommodated by the system of thrust faults shown in the map above, and calculated how these faults could plausibly release the accumulated pressure, based on their known slip rates and other data. Now we don't know whether this stress is relieved in a few large earthquakes, or a lot of smaller ones, or some mix in between, but this compression by itself is enough to give us one magnitude-7.3 shaking about every 150 years. In the last 200 years, we've only had two magnitude 6.7s, Northridge and the San Fernando earthquake of 1971, so this seems to indicate that there are going to be some large earthquakes sooner or later, and that one such quake might be overdue.

What might this quake do to L.A.'s steel buildings? For the sake of discussion, I'm going to consider a magnitude-7.0 earthquake on the

The hypothetical magnitude-7 earthquake on the Elysian Park fault (right) starts 15 kilometers below the surface and has a peak slip of four meters, as shown by the colored contours. In the grid of observing stations at the ground's surface, the letter indicates north-south location, and the number is east-west position. The ground motions predicted at grid point C5 (below) are plotted to the same scale as the Sylmar ground motions on page 5.

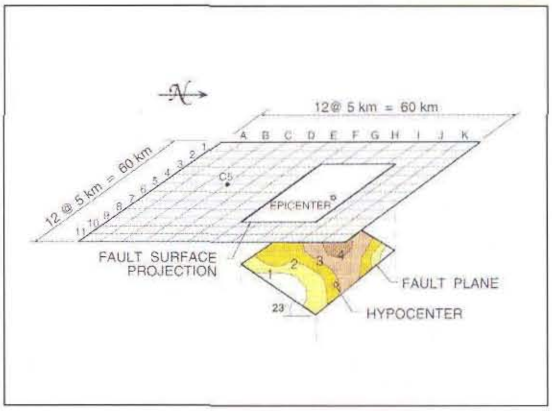

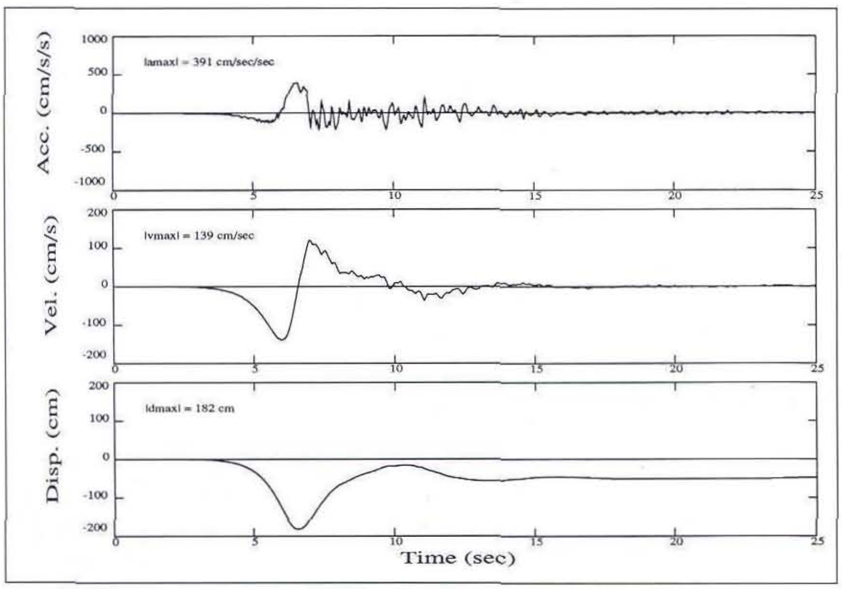

Right: As an earthquake kicks the ground out from under a building (A through C), the lower stories are dragged along while inertia briefly keeps the upper stories at rest. Then by the time the upper stories respond to the initial outward motion, the ground is bringing the lower ones back in, and the two parts of the building are moving in different directions at once (D and E).

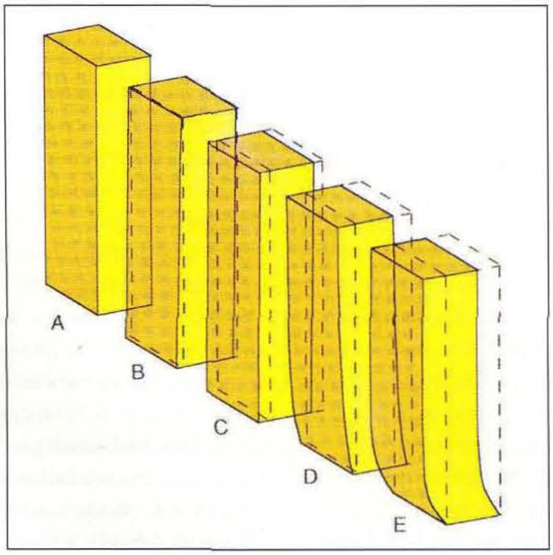

Elysian Park thrust ramp (part of the purple region on the map), which dips to the north and passes directly beneath downtown Los Angeles. Considering that our recent 6.7 was on a blind thrust fault, and that the magnitude-7.5 Kern County quake of 1952 occurred in a similar tectonic setting, this seems pretty plausible.

How do we know how the ground is going to move in a future earthquake, like our hypothetical magnitude 7.0? Well, the seismologists come through again. I've mentioned the inverse problem; this is the forward problem. Through their inverse studies, seismologists have developed a pretty good idea of how ground rupture takes place, so they can impose a reasonable fault-rupture scenario on a mathematical model of a chunk of the earth. From this they can compute the ground motion anywhere, including on the surface. For this hypothetical magnitude 7.0, which Tom and Dave ran for me, the most damaging ground motions occur to the south, in the area toward which the rupture is propagating. In this region, say at location C5, the peak acceleration isn't so big, because we're some distance from the fault. But look at the peak displacement—182 centimeters is about six feet, and this fault doesn't even break the surface! And the accompanying velocity is 139 centimeters per second—about four and a half feet per second—which is a pretty good leap for a piece of solid ground. Needless to say, this is very worrisome.

Let's consider how a building could be affected by this leap, which is actually a double leap—out and back. In other words, the moving ground carries the base of the building out with it and then brings it back. The outward movement gets the building going forward at a high velocity; then the ground doubles back (and the lower stories with it), putting the building under enormous stress. Even if the building can arrest its forward motion, it's liable to experience severe deformations in the lower part of its structure. If the welds are popping on top of this, it's going to have a very hard time stopping, greatly increasing the likelihood of collapse.

Now it's time for some engineering analyses. I fed the ground motions—the Sylmar one from the Northridge earthquake and the C5 one from the simulated magnitude 7.0—into a computer model of how a steel-frame building behaves when shaken. This model is a more sophisticated descendant of the ones that the engineers were using back in the 1970s. One improvement is that this program is able to approximately represent weld fracture. But weld fracture is only one of the ways in which a building can lose strength and stiffness. Another way is that, when a beam

The 20-story building before the C5 ground motion hits. The displacement pulse will be toward the left.

At t=6 seconds, the ground is approaching its maximum horizontal displacement of 182 centimeters.

At t=7 seconds, the ground is returning to its original position, causing the building to "crack the whip."

This flexure creates a ripple of breaking welds that travels up the building.

By t=16 seconds, the building is hopelessly overbalanced and on its way to oblivion.

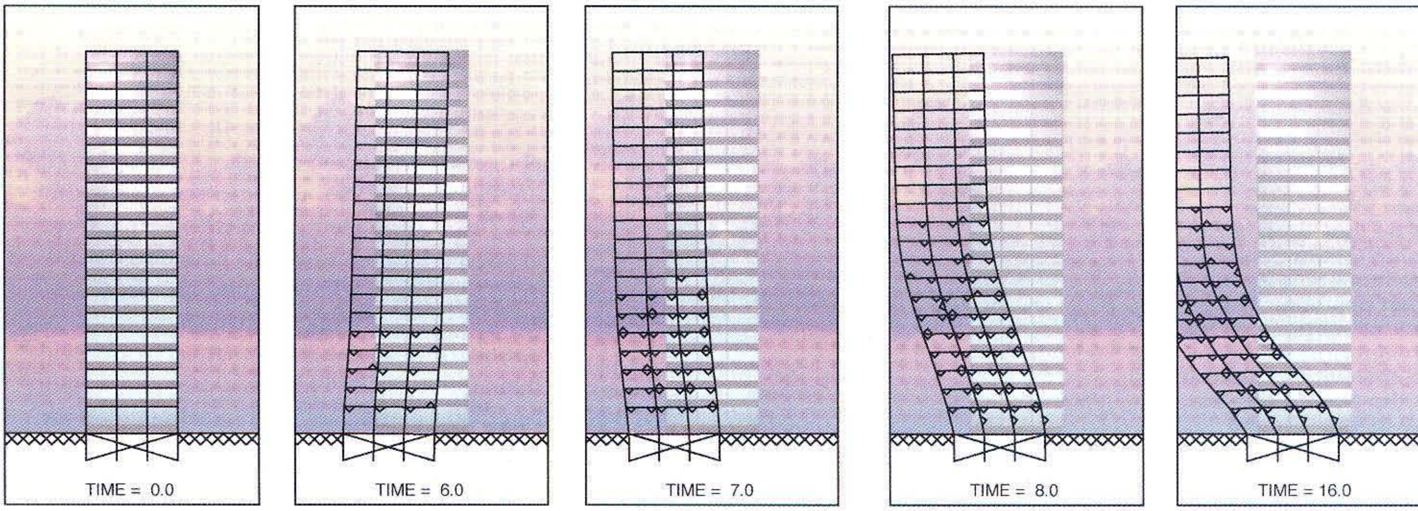

	Sylmar	C5
6-story	3.0	12.4
20-story	2.0	*

Above: Peak story drifts (shown as percentages) calculated for a six-story and a twenty-story steel frame building subjected to the Sylmar and C5 ground motions. The asterisk indicates a collapse predicted by the computer.

kinks, the flange that's in compression can buckle. Flange buckling can be a very significant type of deterioration, but it's extremely hard to model and my program isn't smart enough to do it. And the base plates, which secure columns to their foundations, can fail; concrete slabs can crack; beams can buckle in torsion; the list goes on and on... So sometimes the program computes very large story drifts, and I'd have to think that if it had included more deterioration mechanisms, the building would have collapsed. We should interpret these large story drifts as actual collapses, even though the output doesn't explicitly say so. The table at left shows the peak story drifts computed for a six-story and a 20-story structure subjected to our two ground motions.

The Sylmar numbers are pretty good news. Story drifts of 2 and 3 percent are not unreasonable, especially considering the ground motion's strength and the weld-fracture problem. So even if we'd had more steel-framed buildings hit with near-source directivity effects as measured in Sylmar, we probably shouldn't have seen any collapses. However, the Sylmar record doesn't represent the Northridge earthquake's strongest motion—it's just one of the strongest ones that happened to get recorded. The most damaging ground motions occurred in the mountains north of the San Fernando Valley, and might have caused problems had there been buildings up there to feel them. This is now being studied.

The C5 ground motion is another story. The six-story building has a 12 percent story drift, which is one of those numbers that we have to interpret as a collapse, and the 20-story building collapses outright. The sequences of images across the top of this page are from a computer-animated movie that Wayne Waller of Caltech's Media Integration Lab made from the data generated by my 20-story building model and the C5 ground motion. All of the displacements in these graphics have been amplified by a factor of five for clarity, and the little triangles denote fractured welds. The sequence ends with the building clearly headed for collapse. (Convergence problems in the computer code prevent the model from following the building all the way down.)

So—now that I've shown you these things, here's the big question: Are our steel buildings, which we thought were our most earthquake-resistant type of structure, liable to collapse? We've seen that they're going to behave brittlely during earthquakes, not ductilely as we expected. Also, we can get near-source ground motions from large earthquakes that are considerably stronger than the building code provides for. Furthermore, large earthquakes have duration effects that are not anticipated properly. A magnitude 7.5 can give you 30 seconds of strong shaking, instead of the seven or eight seconds felt in the cases I've shown here, and deterioration is a function of duration. So I think that when we consider these things, we have to admit the possibility that some of our steel buildings will collapse. In Japan, where they build stronger buildings with much better quality control than we do here, they had some problems in the Kobe quake. I've heard from a reliable source that about 30 low-rise modern steel buildings collapsed, although I haven't been able to confirm that.

What about the *real* high-rises? I only looked at a 20-story building; what about the skyscrapers? It turns out that they are actually probably safer, for various reasons. They're relatively stronger than the mid-rise and shorter buildings, because they're designed to carry larger loads—higher wind loading on their bigger surface areas, and, of course, their own heavier weights. Also, skyscrapers like to vibrate back and forth very, very slowly—their natural resonant frequencies are quite low—and only a very large earthquake would have enough low-frequency motion to really grab hold of them and make them move. However, the geologists aren't ruling out such an earthquake, and our experience with Northridge tells us that we have to assume that the welds in these buildings are deficient. So that's something that deserves more study.

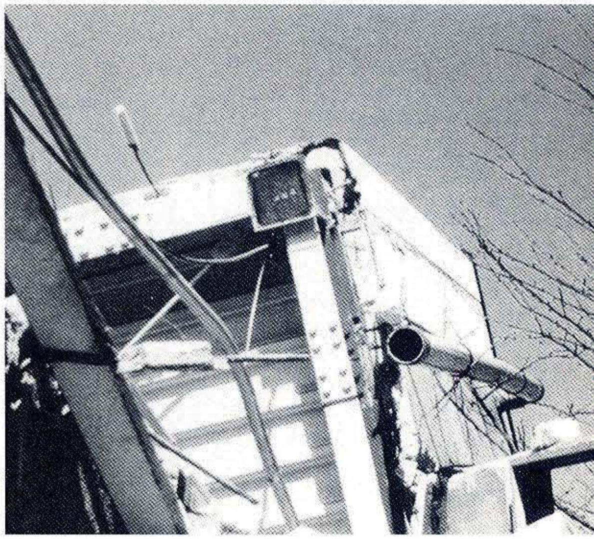

Steel columns in Japanese buildings are not I-shaped but square in cross section. In the Kobe quake, some columns snapped (below), toppling buildings (right). In this picture of the underside of an upper story, you can see the hollow square of the column that used to support the corner of the building.

By now, if you work in a steel building, you're probably starting to wonder about your chances. Life is full of risks, and there are ways to quantify them. (I think it's something we should do more of.) Let's be blunt—what are the chances of getting killed by a steel building if you work in one? Here's how to figure it out. First, you ask a seismologist what the probability of a large earthquake is, and what the probability is that your building will be in the near-source region, and you multiply those numbers together. Then you ask an engineering researcher what the probability is that your building will collapse. I don't know what answer you'll get, but it may be a fairly modest percentage—not every building is going to collapse. Multiply again, and then you multiply that figure by the fraction of your time that you actually spend in the building. If you work there eight hours a day, five days a week, then you only spend about 23 percent of your time there. (This has been a saving grace for many earthquakes—they hit any hour of the day, any day of the week with equal probability, so the odds are good that you won't be in the building when the time comes.) You can reduce your calculated risk still further because most buildings don't pancake when they fail. Usually, only a few floors collapse—we saw that a lot in Kobe. So you want to also consider the odds that you're going to be on one of those floors. If you work all of that out, you may find a number you can live with, especially if you compare it to some other numbers—the probability of being hit and killed by a drunk driver, for example. It's important to keep these things in perspective.

But there's more to an earthquake's toll than lives lost—there's property damage. The Northridge quake cost us about $20 billion at last count; direct property damage from the Kobe quake is currently about $100 billion. An Elysian Park earthquake under downtown Los Angeles would easily cost as much as Kobe. Can our economy take a $100-billion hit? When people were coming up with the building code's philosophy 30 years ago, we weren't having many earthquakes. Therefore it seemed reasonable to design minimal buildings that were just strong enough to avoid collapse (or so they thought), and it wasn't economical to worry about damage control. Today we have a much better idea of the earthquake threat, and things look more ominous. I'd be willing to bet that if it were possible to do a proper economic analysis, it would now make much more sense to design stronger buildings to limit damage. And, of course, stronger buildings would also save more lives.

For many years now, new buildings on the Caltech campus have been designed by increasing the code forces by 50 percent. This is just smart business practice: we sit on top of the Raymond fault; the Sierra Madre fault is just a few miles away; we're self-insured. I think that such designs will become more common as more people, including the code writers and the government, realize the benefits of damage control. The Moore Laboratory of Engineering, currently under construction, is a very strong building with reinforced concrete walls. That's a good design choice for earthquake country, but what's particularly relevant to our discussion of steel frames is the penthouse. We used bolted flange connections there, even though welding is cheaper, as we just weren't comfortable with the defect potential of the welds. Bolted connections,

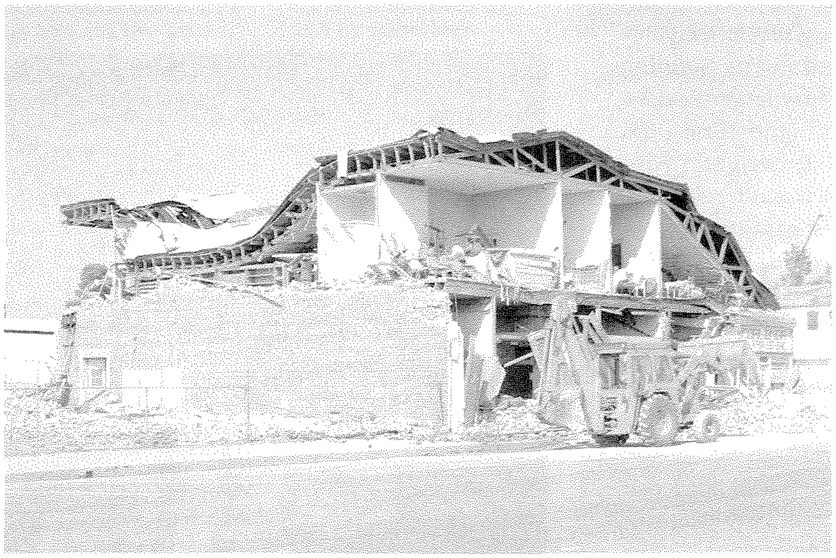

Above: This unreinforced masonry building in downtown Coalinga collapsed in the 1983 earthquake. Many California cities still have large stocks of such buildings, and no retrofit programs.

however, should behave like perfect, defect-free welded ones.

Now, finally, in an effort to make you feel a little better about steel buildings, and to again put things in perspective, let me remind you that there are a lot worse things out there. Unreinforced masonry—seen in buildings predating the 1933 Long Beach earthquake—is one, as has been demonstrated many times, such as in the 1983 Coalinga quake. Several cities, including Long Beach and the city of Los Angeles itself, have tried to address this problem by requiring the owners of such buildings to do nominal retrofits, such as tying the masonry walls to the floors so that the walls don't pull away and come crashing down. (This is the simplest thing you can do to get obvious benefits. It will avert collapse in medium-sized earthquakes, but it probably won't be enough in large ones. You're reinforcing the weakest point, which means that the failure is just transferred to the next weakest point. This is a general problem with retrofits.) Many other cities haven't done anything yet. Unreinforced masonry buildings remain a real problem, much worse than the steel-building situation.

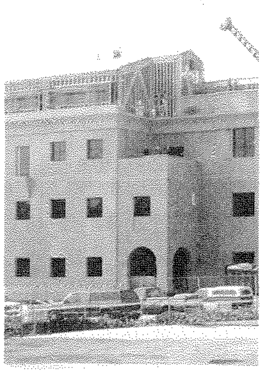

Above: The Gordon and Betty Moore Laboratory of Engineering, currently under construction, has a steel-frame penthouse with fully bolted connections.

Reinforced-concrete-frame structures built before the early 1970s are also very hazardous during earthquakes. They're very brittle, and the things that seem to go first are the columns, which are bad parts of your building to have fail. (I know an engineer who uses the term "ductilely challenged" to refer to this type of construction.) No cities have yet taken action to address their inventories of these nonductile concrete buildings. Two- and three-story wood-frame apartment buildings with an open first story given over to parking are another problem, as we've seen in Northridge; in general, the damage to multistory, multifamily wood-frame buildings was greater than expected. Again, most of these buildings were built before modern codes. (Even something as seemingly minor as using a smaller-diameter nail than the code calls for can make a significant difference in a structure's strength.) The most infamous example, of course, is the Northridge Meadows Apartments, whose collapse killed 16 people. Even some types of modern structures, namely precast concrete parking garages, are known to be collapse hazards—we lost seven of them in the Northridge earthquake.

It has become traditional, in the months following a damaging earthquake in California, for the governor to call on a blue-ribbon panel to investigate the structural failures caused by that quake. The panel eventually issues a report summarizing the engineering lessons learned, and recommending modifications in the building codes and other precautions that—if implemented—should significantly reduce damage in subsequent earthquakes. A glance at the titles of these reports gives us an unintended insight into California's earthquake problem. After the 1989 Loma Prieta quake, the Board of Inquiry viewed the situation as "Competing Against Time." The Seismic Safety Commission, in its recent report on the Northridge earthquake, sees the need for "Turning Loss To Gain," although someone has said that, following the lead of Loma Prieta's Board of Inquiry, a better title would have been "We Lost." Certainly, if we don't pay serious attention to our earthquake threat, we'll be "Picking Up the Pieces" in a future report. □

Associate Professor of Civil Engineering John Hall was the team leader for the Earthquake Engineering Institute's reconnaissance of the Northridge earthquake, and participated in the Seismic Safety Commission's study of that quake. (He was the secretary to the Board of Inquiry into the Loma Prieta earthquake.) He is also a member of Caltrans' Seismic Advisory Board and the White House Office of Science and Technology Policy's National Earthquake Strategy Working Group. His research combines computer simulations, laboratory models, and field testing, and focuses on the nonlinear response of structures, especially high-rise buildings and concrete dams, to earthquakes. Hall's degrees in civil engineering are a BS from West Virginia University in 1972, an MS from the University of Illinois in 1973, and a PhD (with a minor in seismology) from UC Berkeley in 1980; he also has several years' worth of "real-world" experience in a structural design office. This article is adapted from a recent Watson lecture.

Bodies of Steel on Legs of Rubber

Below: A three-story base-isolated building gets bent out of shape by the C5 ground motion in these stills from another Media Integration Lab movie.

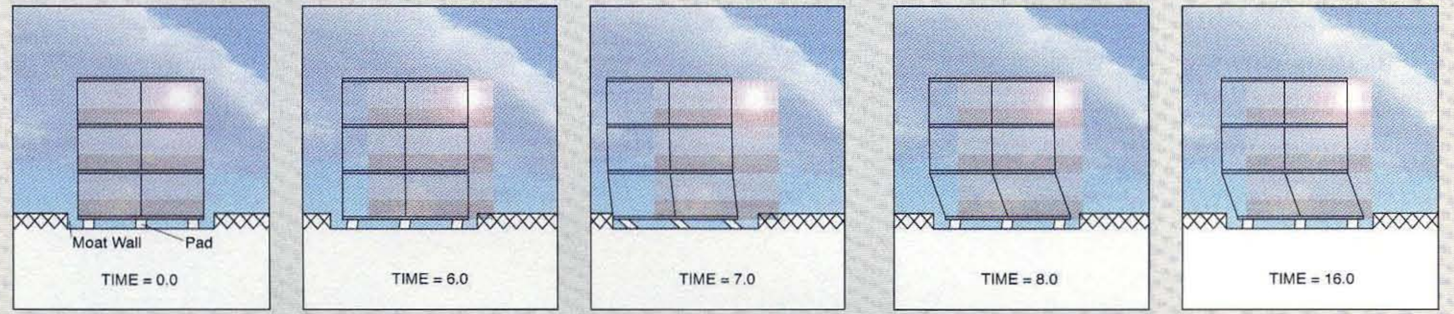

	Sylmar	C5
16-inch	5.0!	*!
20-inch	1.7	19.8!
24-inch	1.1	10.4

Above: The peak story drifts calculated for a three-story base-isolated building with a 16-, 20-, or 24-inch-wide moat, when subjected to the Sylmar and C5 ground motions. An exclamation point indicates that the building hit the moat wall, and an asterisk indicates a collapse predicted by the computer.

It's a common myth that many buildings in Southern California are on rollers. Not so, but we do have about half a dozen base-isolated buildings, which are built on rubber pads, and we're building more. It's a similar idea to the rollers—put something soft between the ground and the building to try to reduce the ground motion that travels up into the building. This is expensive, so it's only been used so far for critical structures, such as hospitals and emergency operations centers, that need to remain functional after earthquakes. How would near-source ground motions from a large quake affect such buildings?

The designs for base-isolated structures are generally more sophisticated than for fixed-base buildings, and the engineers do usually take some account of the near-source directivity effect—it's the controlling issue, in fact. Consequently, a major design goal is to keep the building's displacements reasonable, so that the structure does not move too far on the pads. Otherwise, the building's weight would squash the pad sideways, and the structure would drop down. So as an added precaution, the engineers often put stops—usually low concrete walls—around the building to act as a barrier. This is just so everyone can sleep better at night, because the building isn't supposed to actually hit them. If that ever happened, it would damage the structure and probably wreck the contents—the building wouldn't exactly be functional any more. The zone of free movement between the building and the stops is called the moat; the moat's width, and ensuring that the pads remain stable within this width, is the critical design issue.

I have another computer program—it's rather crude, but it models a lot of the yielding behavior and other nonlinear features that are important for this problem—with which I've analyzed the response of a three-story base-isolated building to the Sylmar and C5 ground motions. I considered three cases: a 16-inch-wide moat, which is typical of the buildings we've already built; a 24-inch moat, a better design that's typical of several buildings now going up close to major faults; and an intermediate 20-inch moat. The results, as seen in the table at left, aren't encouraging. The building collided with the stops in three of the six trials, and collapsed once. There are also some very high story drifts, which again should be interpreted as collapses.

There are only two cases that might appear satisfactory—the two better-designed isolation systems in the Sylmar ground motion. But even there, we're getting story drifts that tell us that the building yields. This is not good, because in order to ensure that the building and its contents will still be in working order after the shaking stops, the engineer usually makes the promise that the building is going to behave elastically. But that's not true even in our best results—there is some structural damage. Across the top of the page are some stills from a movie we made of the 20-inch moat for the C5 ground motion. The displacements and the moat width are amplified by two, in order to see them better. Note how much the building yields after it hits the wall.

So the near-source ground motions being used in the design of base-isolated buildings could be too small, and the resulting buildings may not, in fact, be "earthquake-proof." □—*JH*