

What's the Worst That Can Happen?


Bob Colwell

When engineers design systems to work in the real world—and what other kind of product-for-profit is there?—they must repeatedly ask themselves, “What’s the worst that can happen?” They may argue that nominal operating conditions are far from those limits; they may even seek relief from those limits in their design if accommodating worst-case scenarios would be prohibitively expensive or technically infeasible. But they must always know where the limits are.

In reduced form, this principle is usually introduced early in the educational curriculum. When designing a series of logic gates to solve a problem, the designer tallies up all of the delays through each gate, plus the transit time of the signals between the gates, to arrive at an overall worst-case, never-to-exceed time delay through the chain. Making that circuit operate faster than that worst-case delay would require changes to logic, circuits, or layout.

Programmers have equivalent concerns with the worst-case scenarios. Critical sections, semaphores, flags, protected regions—all are intended to maintain correct algorithm execution even in the statistically unlikely case that multiple accesses to shared control variables are simultaneously pending.

Granted, “statistically unlikely” is a relative term here; how the system handles such shared resources is of first-order importance to overall per-



Engineers must always know where the system's limits are.

formance and scalability. But on a case-by-case basis, such collisions are unusual.

We all know Amdahl's law, which essentially says that to speed up a system, you concentrate on making the common case fast, but as you succeed at that, the remaining unspeeded-up part looms proportionately ever larger. Intel Fellow Dave Papworth points out that there's a dual to Amdahl's law, which he calls “Slhadma's” law (get it?): Unlikely corner cases of a design generally don't figure into overall system performance, but if neglected sufficiently (while you're off chasing Amdahl's law), they will eventually become so relatively slow that they drag down performance and require remedies.

NONENGINEERING WORST-CASE SCENARIOS

This question of “what's the worst that can happen?” appears in the nonengineering world in many guises.

Some self-help books intended to bolster a person's confidence in public speaking use this idea. However, I must admit that it's not clear to me how this helps. As a frequent public speaker, it seems to me that the worst that could happen is that the audience would rush the stage and pummel me senseless. Pondering that prospect doesn't improve my confidence.

I've also seen this strategy used in parental how-to books: When little Janey is disassembling an old broken toy, let her do it. What's the worst that can happen? But if little Johnny has a screwdriver in his hand and is looking inquiringly at your new big-screen TV, parental intervention probably is warranted.

When the kids get older, you're advised to keep a wary eye on them. If what they want to do seems like a bad idea, but it won't hurt them and might actually teach them a useful lesson, then the parenting books say to let them do it. But you can't use this approach to teach a child not to play with fire or run across a busy street. The worst that can happen would be unacceptable.

We know the kinds of mischief that little Johnny's teenage brother can get into when he decides to see if he can create self-replicating computer programs on the Internet. Hackers, worms, viruses, and spyware are the effluent of our emerging computer-based economies. The worst that Johnny's big brother can do with a computer is bad, expensive, and sometimes supremely annoying, but so far not terribly dangerous. (There are reasons to think things won't stay this way, but that's a topic for another column.)

HOW MUCH TROUBLE COULD AN EAGLE SCOUT GET INTO?

The worst that can happen in non-computer fields is, well, worse. Much

worse. For instance, in many ways, it's humankind's great good luck that nuclear weaponry is so hard to make. (It might have been even better luck had it been impossible.) Fissionable material isn't widely available in nature, and it's not easily made from anything else. It takes detailed knowledge of chemistry, along with good training (or sheer dumb luck) to avoid getting killed while muddling up the learning curve. Consider the sad fates of Marie Curie and the dozens of people who died from accidents (or ignorance) in the nuclear weapons and nuclear power industries.

A cautionary tale for modern times

Radioactivity is dangerous. Which is why a book titled *The Radioactive Boy Scout* was so immediately intriguing, especially since the subtitle is *The Frightening True Story of a Whiz Kid and His Homemade Nuclear Reactor* (Ken Silverstein, Villard, 2004). This book is a cautionary tale for our modern times on many levels.

What Boy Scout David Hahn tried to do was to build a fission reactor in the potting shed of his parent's home. Hahn had read a book from the 1950s extolling the unlimited future of nuclear energy, particularly breeder reactors, and decided his mission in life was to further that cause.

The saying that a little knowledge is a dangerous thing certainly applies in this case. Hahn had acquired a deep knowledge of chemistry, which he combined with a shallow knowledge of the actual state of affairs in the nuclear industry. For example, Hahn didn't know that the actual history of real breeders was one of failure after failure. Hahn said he purposely steered clear of such "negativity": "If I knew [a technical account of a nuclear initiative] had a critical perspective, I wouldn't even pick it up."

To the extent that this book's account of him is accurate, Hahn sets a new high-water mark for individual hubris and irresponsibility throughout

the narrative. Irresponsible, yes, but bright and determined as well.

Hahn couldn't find a way to obtain enough fissionable material like uranium-235 or plutonium. However, he knew from his chemistry studies that bombarding thorium-232 with neutrons would yield thorium-233, which emits a beta particle and becomes uranium-233, a manmade fissionable element. In other words, Hahn needed thorium and a neutron gun.

Engineers are preconditioned to let the data tell us where our analysis should go.

Getting advice and finding fuel

Through various subterfuges, such as posing as a physics professor in mail correspondence, Hahn managed to get useful advice and coaching from experts in the nuclear industry. Combining that advice with his own burgeoning knowledge of chemistry, he took the americium out of hundreds of smoke detectors and combined it with aluminum in a lead frame to make a neutron gun. He even knew how to test it: Paraffin gives off protons when hit by neutrons, and his Geiger counter could detect those protons.

Looking for fuel, Hahn drove around many mining sites with a Geiger counter on his dashboard. When he couldn't find enough radioactive material to suit his purposes, he bought some ore on the Internet and tried to refine it to U-238. Had he succeeded at that, and used his neutron gun on it, he would probably have produced plutonium, the deadliest substance on Earth, and the story would have ended there.

Instead, he turned to thorium, stole hundreds of Coleman gas mantles (each coated with thorium dioxide), scraped off the thorium, reduced it to

ash, and spent weeks using lithium extracted from batteries to purify it. After further adventures with new sources of radium and beryllium, Hahn put the whole thing together, and his Geiger counter went crazy. Only then did he realize that he: (a) had not provided any means for controlling the reaction, and (b) had forgotten that heat was one of the by-products.

After a few days of increasing radiation, Hahn finally panicked and scattered the entire apparatus until he could figure a way to control it.

Hahn's nuclear activities were eventually detected by sheer accident, when a policeman stopped him for suspicious activity, looked in the car's trunk, and saw what appeared to be an elaborate bomb. Federal agencies dismantled the potting shed, and Hahn joined the Navy, which wouldn't let him anywhere near its nuclear reactors after learning what he'd done for his Boy Scout Eagle project. (The Boy Scouts of America had a vigorous debate about whether Hahn should be allowed to keep his Eagle status, given what his project was really about. Evidently they lack a merit badge for irradiating your neighbors.)

Where were the adults?

Where, you might ask, were his parents during all this? Due to divorce, there were two sets of parents. One set was so thoroughly dysfunctional that they never really asked him what he was doing; the other lacked the imagination to see what trouble he was brewing.

The parents mirrored what all of the other adults in this story did. The high school chemistry and physics teachers, the nuclear industry experts who helped "Professor Hahn" along the way, friends of the family who had the technical background to understand that this kid was inhabiting rarified chemical atmospheres, but simply couldn't envision him doing what he said he was doing—they all failed to imagine the worst-case scenario.

AND NOW FOR SOMETHING REALLY SCARY

Perhaps you and I would find the prospect of an eerie green glow emanating from the neighbor's potting shed somewhat disturbing, especially if you have a Geiger counter that finds it equally interesting. If I knew the neighbor's kid was trying to build a nuclear reactor, I would definitely be concerned.

Not Richard A. Posner. Posner is a federal judge who has written a book titled *Catastrophe: Risk and Response* (Oxford University Press, 2004). If you lived near David Hahn's potting shed, you might well feel justified in labeling it a catastrophe, but Posner is chasing considerably bigger fish: He's trying to compile a list of all events that might cause extinction of human life on Earth. The great flu pandemic of 1918 killed tens of millions of people worldwide—a disaster, yes, but not a catastrophe by Posner's definition.

Posner is literally asking "what's the worst that can happen?" on behalf of our entire species, both the living and untold generations not yet born. He isn't just worried in the sense that we ought to be thinking about how to protect our species. His point is that we as a society aren't well prepared to deal with events that, though rare, have such dire consequences that the default—simply ignoring the issue—might not be the right course of action.

We computer designers and programmers have come to understand very well that if an event has an extremely small likelihood of occurrence, but there are a large number of trials, we *will* see that event transpire. Odds are low that a single fair coin will come up heads each time if flipped 10 times. But gather enough people flipping enough coins, and it becomes nearly certain that somebody's coin will do exactly that. This is what statistics predicts, and this is what we see in labs and the real world.

Posner's list includes pandemics, but he is concerned about a "juiced up" version of an existing pathogen, one

for which nobody on Earth has any antibodies. He argues that we should not be reassured by the fact that this disaster has never yet befallen us because it's only very recently that our knowledge of genetics and other medical fields has reached the necessary level of sophistication to make the threat plausible.

When worst-case design fails, it's usually due to a failure of the designer's imagination.

Asteroids make Posner's list, and one of them features prominently on the front cover of his book. He laboriously defines the sizes and speeds of astral bodies that could pose a threat to all life on Earth by causing fires, tsunamis, enough dust and smoke to block out most sunlight, and acid rain. He considers megavolcano eruptions (such as Yellowstone in the US) as possible threats, but says they're far less frequent than near-misses by asteroids and generally are more localized phenomena.

THE STRANGELET SCENARIO

One of the most interesting of Posner's catastrophes is the "strangelet scenario." When the first atomic bomb was tested, some theorized that it might be hot enough to ignite the atmosphere, and a runaway global conflagration would ensue. Today's equivalent is a concern that when particle accelerator collision energies are high enough, they might produce a shower of quarks that form themselves into an arrangement called a *strangelet*. Strangelets might be contagious—any matter near them might also be assimilated into the strangelet.

Posner says that at least one existing accelerator exceeds the minimum threshold for this scenario, and so far we're all still here. But, as he points

out, it's not clear how reassuring that should be. It could be that the creation of strangelets is a statistical game with very low probability. Try it often enough, and surprise, all your quarks are permanently reassigned.

Other Posnerian catastrophes include runaway nanomachines (K. Eric Drexler's gray goo—*Engines of Creation*, Anchor Books, 1986), genetically modified crops run amok, and artificial intelligence in the form of robots that find us more trouble than we're worth.

Posner considers other problems that, while serious and potentially very deadly, don't seem likely to cause extinction. Among these he includes global warming, and his take on this problem is fascinating, given his day job as a judge. He correctly points out that there are fundamental differences among experts on whether global warming is taking place, whether it is of manmade origin, and whether and to what extent it will be bad for us.

FACING SCIENTIFIC DISCORD

Since Posner is not himself an expert, what should he do in the face of such scientific discord? After considering the viewpoints of many prominent voices in the matter, he tries something novel: He statistically samples the 20 most prominent journals on atmospheric science and meteorology. Based on this sampling, he concludes that, as usual, there are some voices of dissent, but the vast majority of scientists believe that global warming is real and is a product of human activities.

He also considers overpopulation, nuclear winter, exhaustion of natural resources, bioterrorism, and loss of biodiversity, but concludes that none of these is likely to result in extinction.

Posner's aim is to overcome normal human repugnance at these possibilities, to apply statistical methods to them, and to have that analysis guide our actions. He believes the risks of global catastrophe are higher than generally understood, and are growing.

He points out that we don't have a way to value human life for any rational cost-benefit analysis; even raising this issue when any number short of "infinite" is proposed is likely to cut the discussion off before it starts. Yet the outcome of that willful blindness is to leave all of us unprepared for events that could wipe us all out.

Posner posits a fascinating question: When we as a species try to make purposeful decisions about what level of risk we're willing to accept, to what extent should we take into account future generations? He doesn't mean our children and grandchildren—he means over millions of years.

We have no collective practice at thinking in these terms, but Posner is quite right in his insistence on this point. Remember: If, say, the strangelet scenario is of vanishingly low probability, and you run the experiment enough times, the odds of a strangelet

event changes from very low to nearly certain—assuming the model that predicts strangelets can occur is itself correct. For a disaster of this magnitude, it would be wise to err on the side of caution. Posner points out that it isn't just the six billion people on the planet right now that we're risking, it's all of their descendants over potentially vast stretches of time.

When worst-case design fails, it's usually due to a failure of the designer's imagination: It didn't occur to most people that gasket resilience would allow hot gases to blow by a solid rocket booster's O-ring, let alone two of them; we didn't see that something as light as foam could impact a Shuttle wing hard enough to destroy its integrity. Posner has a good, rational imagination, and his instincts that we should be apply-

ing our statistical abilities to these prospects are excellent.

Complacency is the precondition for radioactive pottery sheds or assimilation into the nearest strangelet. We engineers don't always get this right, but we're at least preconditioned to let the data tell us where our analysis should go, and we can handle the requisite math for combining low-probability events with high-consequence effects.

Posner's on the right path. Keep practicing your worst-case forecasting, and try to expand your imagination. When you think you've got it, remember the radioactive Boy Scout and try even harder. ■

Bob Colwell was Intel's chief IA32 architect through the Pentium II, III, and 4 microprocessors. He is now an independent consultant. Contact him at bob.colwell@comcast.net.

IEEE TRANSACTIONS ON DEPENDABLE AND SECURE COMPUTING



Learn more about this new publication and become a subscriber today.

www.computer.org/tdsc

Learn how others are achieving systems and networks design and development that are dependable and secure to the desired degree, without compromising performance.

This new journal provides original results in research, design, and development of dependable, secure computing methodologies, strategies, and systems including:

- Architecture for secure systems
- Intrusion detection and error tolerance
- Firewall and network technologies
- Modeling and prediction
- Emerging technologies

Publishing quarterly

Member rate:

\$31 print issues

\$25 online access

\$40 print and online

Institutional rate: \$275

