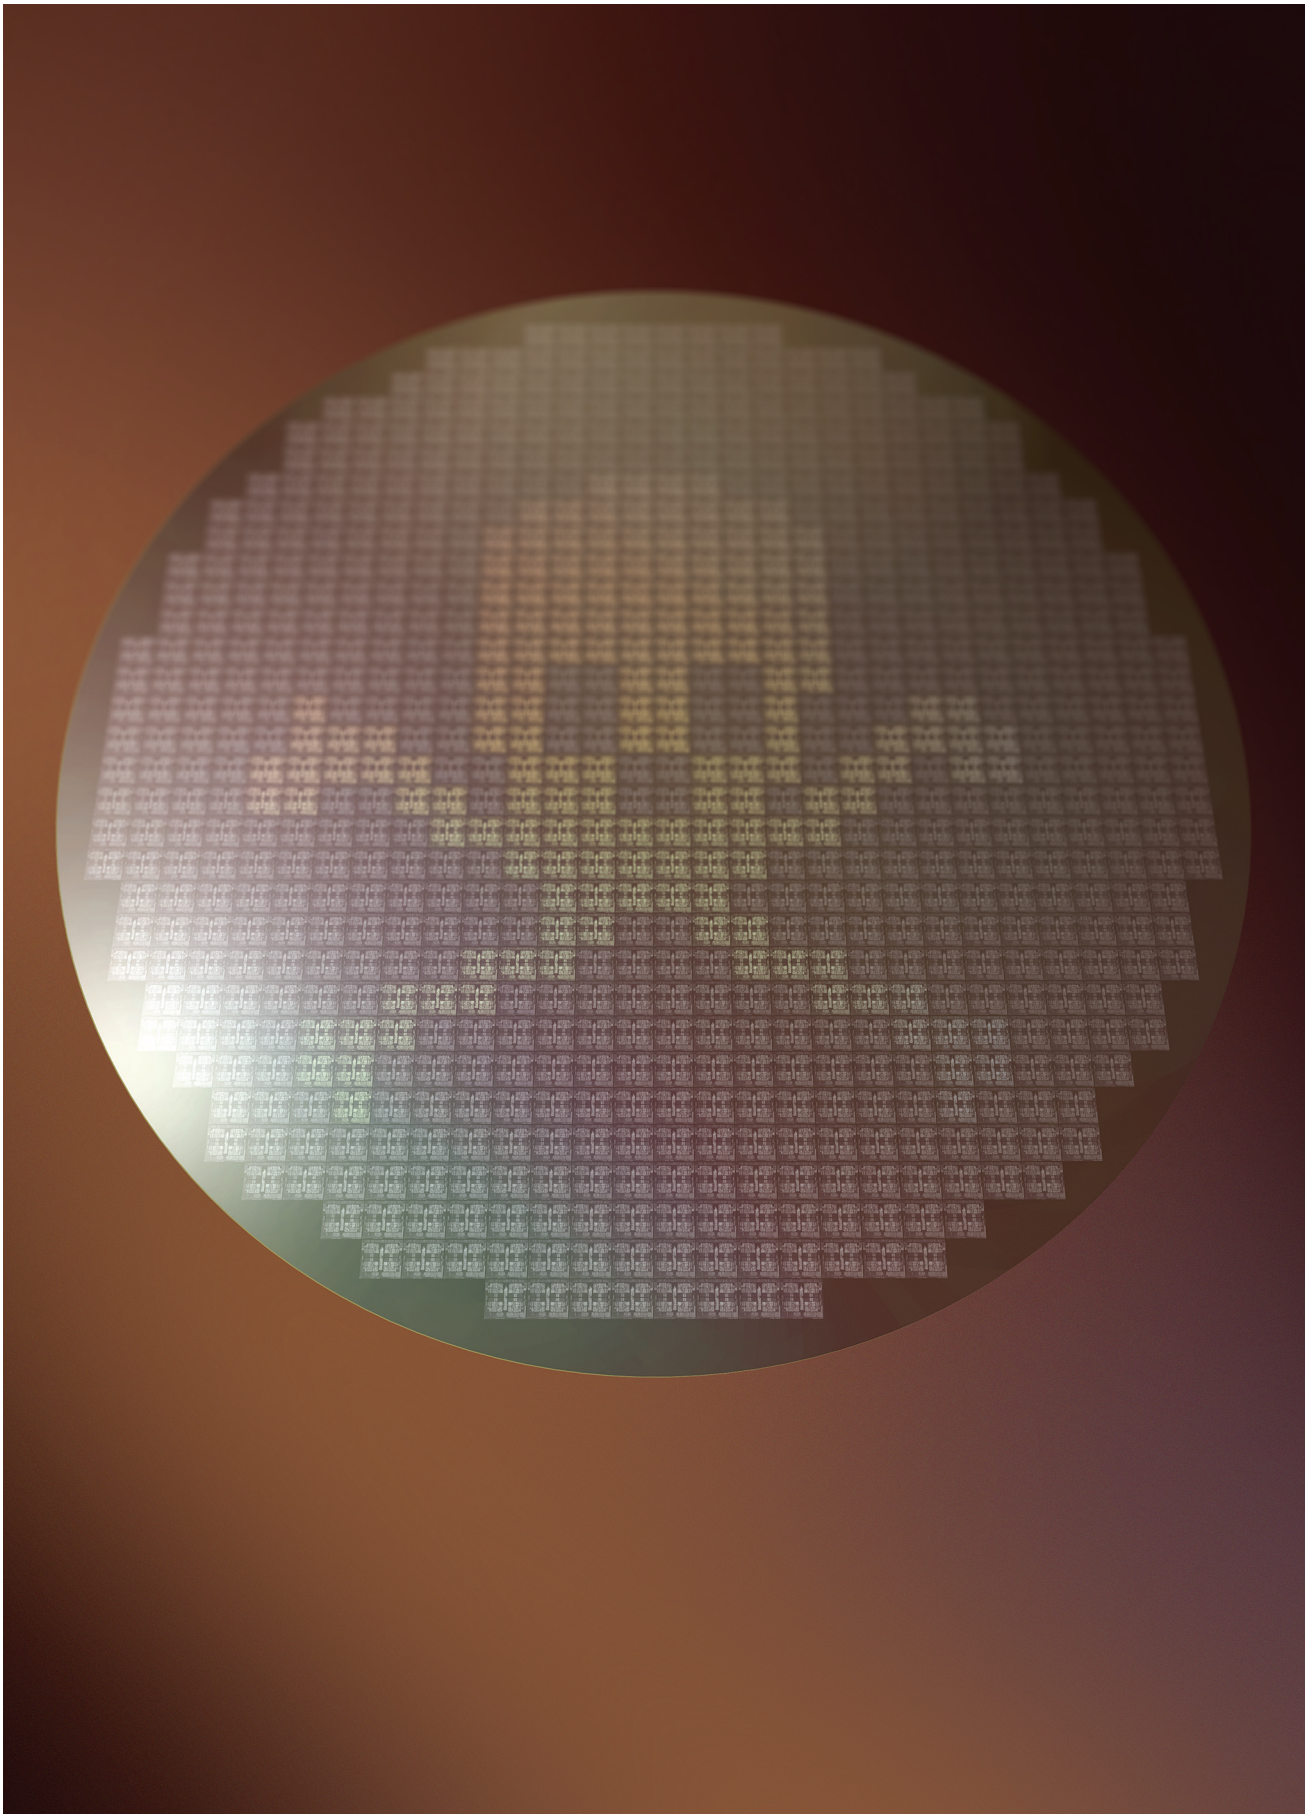

JAMES ARCHER/ANATOMY BLUE

THE HUNT FOR THE KILL SWITCH

ARE CHIP MAKERS BUILDING ELECTRONIC TRAPDOORS IN KEY MILITARY HARDWARE? THE PENTAGON IS MAKING ITS BIGGEST EFFORT YET TO FIND OUT

BY SALLY ADEE

LAST SEPTEMBER, Israeli jets bombed a suspected nuclear installation in northeastern Syria. Among the many mysteries still surrounding that strike was the failure of a Syrian radar—supposedly state-of-the-art—to warn the Syrian military of the incoming assault. It wasn't long before military and technology bloggers concluded that this was an incident of electronic warfare—and not just any kind.

Post after post speculated that the commercial off-the-shelf microprocessors in the Syrian radar might have been purposely fabricated with a hidden "backdoor" inside. By sending a preprogrammed code to those chips, an unknown antagonist had disrupted the chips' function and temporarily blocked the radar.

That same basic scenario is cropping up more frequently lately, and not just in the Middle East, where conspiracy theories abound. According to a U.S. defense contractor who spoke on condition of anonymity, a "European chip maker" recently built into its microprocessors a kill switch that could be accessed remotely. French defense contractors have used the chips in military equipment, the contractor told *IEEE Spectrum*. If in the future the equipment fell into hostile hands, "the French wanted a way to disable that circuit," he said. *Spectrum* could not confirm this account independently, but spirited discussion about it among researchers and another defense contractor last summer at a military research conference reveals a lot about the fever dreams plaguing the U.S. Department of Defense (DOD).

Feeding those dreams is the Pentagon's realization that it no longer controls who manufactures the components that go into its increasingly complex systems. A single plane like the DOD's next generation F-35 Joint Strike Fighter, can contain an "insane number" of chips, says one semiconductor expert familiar with that aircraft's design. Estimates from other sources put the total at several hundred to more than a thousand. And tracing a part back to its source is not always straightforward. The dwindling of domestic chip and electronics manufacturing in the United States, combined with the phenomenal growth of suppliers in countries like China, has only deepened the U.S. military's concern.

Recognizing this enormous vulnerability, the DOD recently launched its most ambitious program yet to verify

RECIPE FOR DISASTER

AT EACH step of the hardware design process, a saboteur could make a particular part of the circuit fail. A typical microprocessor can have up to eight layers, and any layer on a microchip can be targeted. *ILLUSTRATION: EMILY COOPER*

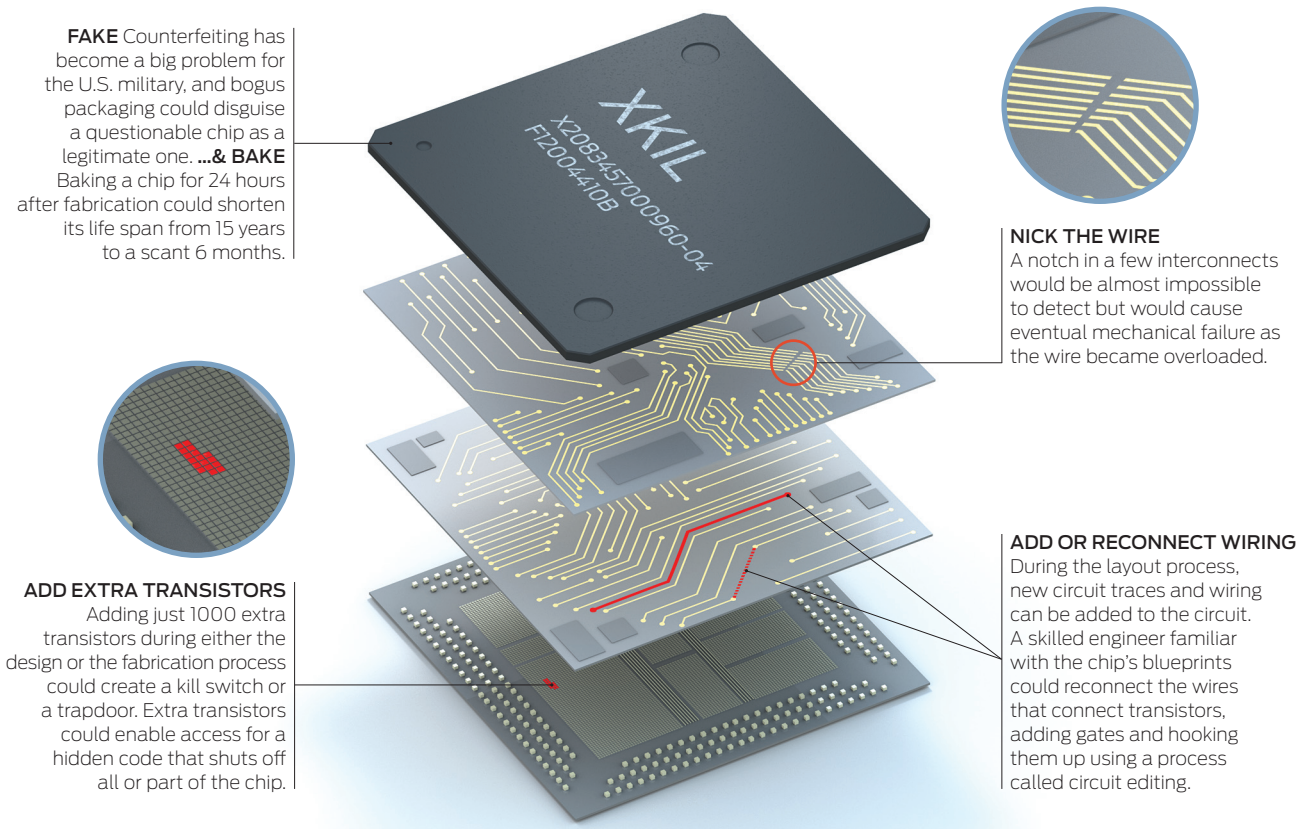

FAKE Counterfeiting has become a big problem for the U.S. military, and bogus packaging could disguise a questionable chip as a legitimate one. **...& BAKE** Baking a chip for 24 hours after fabrication could shorten its life span from 15 years to a scant 6 months.

NICK THE WIRE
A notch in a few interconnects would be almost impossible to detect but would cause eventual mechanical failure as the wire became overloaded.

ADD EXTRA TRANSISTORS
Adding just 1000 extra transistors during either the design or the fabrication process could create a kill switch or a trapdoor. Extra transistors could enable access for a hidden code that shuts off all or part of the chip.

ADD OR RECONNECT WIRING
During the layout process, new circuit traces and wiring can be added to the circuit. A skilled engineer familiar with the chip's blueprints could reconnect the wires that connect transistors, adding gates and hooking them up using a process called circuit editing.

the integrity of the electronics that will underpin future additions to its arsenal. In December, the Defense Advanced Research Projects Agency (DARPA), the Pentagon's R&D wing, released details about a three-year initiative it calls the Trust in Integrated Circuits program. The findings from the program could give the military—and defense contractors who make sensitive microelectronics like the weapons systems for the F-35—a guaranteed method of determining whether their chips have been compromised. In January, the Trust program started its prequalifying rounds by sending to three contractors four identical versions of a chip that contained unspecified malicious circuitry. The teams have until the end of this month to ferret out as many of the devious insertions as they can.

Vetting a chip with a hidden agenda can't be all that tough, right? Wrong. Although commercial chip makers routinely and exhaustively test chips with hundreds of millions of logic gates, they can't afford to inspect everything. So instead they focus on how well the chip performs specific functions. For a microprocessor destined for use in a cellphone, for instance, the chip maker will check to see whether all the phone's various functions work. Any extraneous circuitry that doesn't interfere with the chip's normal functions won't show up in these tests.

"You don't check for the infinite possible things that are not specified," says electrical engineering professor Ruby Lee, a cryptography expert at Princeton. "You could check the obvious possibilities, but can you test for every unspecified function?"

Nor can chip makers afford to test every chip. From a batch of thousands, technicians select a single chip for physical inspection, assuming that the manufacturing process has yielded essentially identical devices. They then laboriously grind away a thin layer of the chip, put the chip into a scanning electron microscope, and then take a picture of it, repeating the process until every layer of the chip has been imaged. Even here, spotting a tiny discrepancy amid a chip's many layers and millions or billions of transistors is a fantastically difficult task, and the chip is destroyed in the process.

But the military can't really work that way. For ICs destined for mission-critical systems, you'd ideally want to test every chip without destroying it.

The upshot is that the Trust program's challenge is enormous. "We can all do with more verification," says Samsung's Victoria Coleman, who helped create the Cyber Trust initiative to secure congressional support for cybersecurity. "My advice to [DARPA director] Tony Tether was 'trust but verify.' That's all you can do."

Semiconductor offshoring dates back to the 1960s, when U.S. chip makers began moving the labor-intensive assembly and testing stages to Singapore, Taiwan, and other countries with educated workforces and relatively inexpensive labor.

Today only Intel and a few other companies still design and manufacture all their own chips in their own fabrication plants. Other chip designers—including LSI Corp. and most recently Sony—have gone "fabless," outsourcing their manufacturing to offshore facilities known as foundries. In doing so, they avoid the huge expense of building a state-of-the-art fab, which in 2007 cost as much as US $2 billion to $4 billion.

Well into the 1970s, the U.S. military's status as one of the largest consumers of integrated circuits gave it some control over the industry's production and manufacturing, so the offshoring trend didn't pose a big problem. The Pentagon could always find a domestic fab and pay a little more to make highly classified and mission-critical chips. The DOD also maintained its own chip-making plant at Fort Meade, near Washington, D.C., until the early 1980s, when costs became prohibitive.

But these days, the U.S. military consumes only about 1 percent of the world's integrated circuits. "Now," says Coleman, "all they can do is buy stuff." Nearly every military system today contains some commercial hardware. It's a pretty sure bet that the National Security Agency doesn't fabricate its encryption chips in China. But no entity, no matter how well funded, can afford to manufacture its own safe version of every chip in every piece of equipment.

The Pentagon is now caught in a bind. It likes the cheap, cutting-edge devices emerging from commercial foundries and the regular leaps in IC performance the commercial sector is known for. But with those improvements comes the potential for sabotage. "The economy is globalized, but defense is not globalized," says Coleman. "How do you reconcile the two?"

In 2004, the Defense Department created the Trusted Foundries Program to try to ensure an unbroken supply of secure microchips for the government. DOD inspectors have now certified certain commercial chip plants, such as IBM's Burlington, Vt., facility, as trusted foundries. These plants are then contracted to supply a set number of chips to the Pentagon each year. But Coleman argues that the program blesses a process, not a product. And, she says, the Defense Department's assumption that onshore assembly is more secure than offshore reveals a blind spot. "Why can't people put something bad into the chips made right here?" she says.

Three years ago, the prestigious Defense Science Board, which advises the DOD on science and technology developments, warned in a report that the continuing shift to overseas chip fabrication would expose the Pentagon's most mission-critical integrated circuits to sabotage. The board was especially alarmed that no existing tests could detect such compromised chips, which led to the formation of the DARPA Trust in IC program.

Where might such an attack originate? U.S. officials invariably mention China and Russia. Kenneth Flamm, a technology expert at the Pentagon during the Clinton administration who is now a professor at the University of Texas at Austin, wouldn't get that specific but did offer some clues. Each year, secure government computer networks weather thousands of attacks over the Internet. "Some of that probing has come from places where a lot of our electronics are being manufactured," Flamm says. "And if you're a responsible defense person, you would be stupid not to look at some of the stuff they're assembling, to see how else they might try to enter the network."

John Randall, a semiconductor expert at Zyvex Corp., in Richardson, Texas, elaborates that any malefactor who can penetrate government security can find out what chips are being ordered by the Defense Department and then target them for sabotage. "If they can access the chip designs and add the modifications," Randall says, "then the chips could be manufactured correctly anywhere and still contain the unwanted circuitry."

So what's the best way to kill a chip? No one agrees on the most likely scenario, and in fact, there seem to be as many potential avenues of attack as there are people working on the problem. But the threats most often mentioned fall into two categories: a kill switch or a backdoor.

A kill switch is any manipulation of the chip's software or hardware that would cause the chip to die outright—to shut off an F-35's missile-launching electronics, for example. A backdoor, by contrast, lets outsiders gain access to the system through code or hardware to disable or enable a specific function. Because this method works without shutting down the whole chip, users remain unaware of the intrusion. An enemy could use it to bypass battlefield radio encryption, for instance.

Depending on the adversary's degree of sophistication, a kill switch might be controlled to go off at a set time, under certain circumstances, or at random. As an example of the latter, Stanford electrical engineering professor Fabian Pease muses, "I'd nick the [chip's] copper wiring." The fault, almost impossible to detect, would make the chip fail early, due to electromigration: as current flowed through the wire, eventually the metal atoms would migrate and form voids, and the wire would break. "If the chip goes into a defense satellite, where it's supposed to work for 15 years but fails after six months, you have a very expensive, inoperative satellite," Pease says.

But other experts counter that such ideas ignore economic realities. "First and foremost, [the foundries] want to make sure their chips work," says Coleman. "If a company develops a reputation for making chips that fail early, that company suffers more than anyone else."

A kill switch built to be triggered at will, as was allegedly incorporated into the European microprocessors, would be more difficult and expensive to pull off, but it's also the more likely threat, says David Adler, a consulting professor of electrical engineering at Stanford, who was previously funded by DARPA to develop chip-testing hardware in an unrelated project.

To create a controlled kill switch, you'd need to add extra logic to a microprocessor, which you could do either during manufacturing or during the chip's design phase. A saboteur could substitute one of the masks used to imprint the pattern of wires and transistors onto the semiconductor wafer, Adler suggests, so that the pattern for just one microchip is different from the rest. "You're printing pictures from a negative," he says. "If you change the mask, you can add extra transistors."

Or the extra circuits could be added to the design itself. Chip circuitry these days tends to be created in software modules, which can come from anywhere, notes Dean Collins, deputy director of DARPA's Microsystems Technology Office and program manager for the Trust in IC initiative. Programmers "browse many sources on the Internet for a component," he says. "They'll find a good one made by somebody in Romania, and they'll put that in their design." Up to two dozen different software tools may be used to design the chip, and the origin of that software is not always clear, he adds. "That creates two dozen entry points for malicious code."

Collins notes that many defense contractors rely heavily on field-programmable gate arrays (FPGAs)—a kind of generic chip that can be customized through software. While a ready-made FPGA can be bought for $500, an application-specific IC, or ASIC, can cost anywhere from $4 million to $50 million. "If you make a mistake on an FPGA, hey, you just reprogram it," says Collins. "That's the good news. The bad news is that if you put the FPGA in a military system, someone else can reprogram it."

Almost all FPGAs are now made at foundries outside the United States, about 80 percent of them in Taiwan. Defense contractors have no good way of guaranteeing that these economical chips haven't been tampered with. Building a kill switch into an FPGA could mean embedding as few as 1000 transistors within its many hundreds of millions. "You could do a lot of very interesting things with those extra transistors," Collins says.

The rogue additions would be nearly impossible to spot. Say those 1000 transistors are programmed to respond to a specific 512-bit sequence of numbers. To discover the code using software testing, you might have to cycle through every possible numerical combination of 512-bit sequences. That's 13.4×10^{153} combinations. (For perspective, the universe has existed for about 4×10^{17} seconds.) And that's just for the 512-bit number—the actual number of bits in the code would almost certainly be unknown. So you'd have to apply the same calculations to all possible 1024-bit numbers, and maybe even 2048-bit numbers, says Tim Holman, a research associate professor of electrical engineering at Vanderbilt University, in Nashville. "There just isn't enough time in the universe."

Those extra transistors could create a kill switch or a backdoor in any chip, not just an FPGA. Holman sketches a possible scenario: suppose those added transistors find their way into a networking chip used in the routers connecting the computers in your home, your workplace, banks, and military bases with the Internet. The chip functions perfectly until it receives that 512-bit sequence, which could be transmitted from anywhere in the world. The sequence prompts the router to hang up. Thinking it was the usual kind of bug, tech support would reset the router, but on restart the chip would again immediately hang up, preventing the router from connecting to the outside world. Meanwhile, the same thing would be happening to similarly configured routers the world over.

The router scenario also illustrates that the nation's security and economic well-being depend on shoring up not just military chips but also commercial chips. An adversary who succeeded in embedding a kill switch in every commercial router could devastate national security without ever targeting the Defense Department directly.

A kill switch or backdoor built into an encryption chip could have even more disastrous consequences. Today encoding and decoding classified messages is done completely by integrated circuit—no more Enigma machine with its levers and wheels. Most advanced encryption schemes rely on the difficulty that computers have in factoring numbers containing hundreds of digits; discovering a 512-bit type of encryption would take some machines up to 149 million years. Encryption that uses the same code or key to encrypt and decrypt information—as is often true—could easily be compromised by a kill switch or a backdoor. No matter what precautions are taken at the programming level to safeguard that key, one extra block of transistors could undo any amount of cryptography, says John East, CEO of Actel Corp., in Mountain View, Calif., which supplies military FPGAs.

"Let's say I can make changes to an insecure FPGA's hardware," says East. "I could easily put a little timer into the circuit. The timer could be programmed with a single command: 'Three weeks after you get your configuration, forget it.' If the FPGA were to forget its configuration information, the entire security mechanism would be disabled."

Alternately, a kill switch might be programmed to simply shut down encryption chips in military radios; instead of scrambling the signals they transmit, the radios would send their messages in the clear, for anybody to pick up. "Just like we figured out how the Enigma machine worked in World War II," says Stanford's Adler, "one of our adversaries could in principle figure out how our electronic Enigma machines work and use that information to decode our classified communications."

Chip alteration can even be done after the device has been manufactured and packaged, provided the design data are available, notes Chad Rue, an engineer with FEI, based in Hillsboro, Ore., which makes specialized equipment for chip editing (albeit for legitimate reasons). FEI's circuit-editing tools have been around for 20 years, Rue says, and yet "chip designers are still surprised when they hear what they can do."

Skilled circuit editing requires electrical engineering know-how, the blueprints of the chip, and a $2 million refrigerator-size piece of equipment called a focused-ion-beam etching machine, or FIB. A FIB shoots a stream of ions at precise areas on the chip, mechanically milling away tiny amounts of material. FIB lab workers refer to the process as microsurgery, with the beam acting like a tiny scalpel. "You can remove material, cut a metal line, and make new connections," says Rue. The process can take from hours to several days. But the results can be astonishing: a knowledgeable technician can edit the chip's design just as easily as if he were taking "an eraser and a pencil to it," says Adler.

Semiconductor companies typically do circuit editing when they're designing and debugging prototypes. Designers can make changes to any level of the chip's wiring, not just the top. "It's not uncommon to dig through eight different layers to get to the intended target," says Rue. The only thing you can't do with a FIB is add extra transistors. "But we can reroute signals to the transistors that are already there," he says. That's significant because chips commonly contain large blocks of unused circuitry, leftovers from previous versions of the design. "They're just along for the ride," Rue says. He thinks it would be possible to use a FIB to rewire a chip to make use of these latent structures. To do so, an adversary would need a tremendous amount of skill with digital circuitry and access to the original design data. Some experts find the idea too impractical to worry about. But an adversary with unlimited funds and time—exactly what the Defense Science Board warned of—could potentially pull it off, Rue says.

In short, the potential for tinkering with an integrated circuit is almost limitless, notes Princeton's Lee. "The hardware design process has many steps," she says. "At each step, you could do something that would make a particular part of the IC fail."

CLEARLY, THE COMPANIES participating in the Trust in IC program have their work cut out for them. As Collins sees it, the result has to be a completely new chip-verification method. He's divided up the Trust participants into teams: one group to create the test chips from scratch; another to come up with malicious insertions; three more groups, which

he calls "performers," to actually hunt for the errant circuits; and a final group to judge the results.

To fabricate the test chips, Collins chose the Information Sciences Institute at the University of Southern California, Los Angeles. He picked MIT's Lincoln Laboratory to engineer whatever sneaky insertions they could devise, and he tapped Johns Hopkins University Applied Physics Laboratory, in Laurel, Md., to come up with a way to compare and assess the performers' results.

The three performers are Raytheon, Luna Innovations, and Xradia. None of the teams would speak on the record, but their specialties offer some clues to their approach. Xradia, in Concord, Calif., builds nondestructive X-ray microscopes used widely in the semiconductor industry, so it may be looking at a new method of inspecting chips based on soft X-ray tomography, Stanford's Pease suggests. Soft X-rays are powerful enough to penetrate the chip but not strong enough to do irreversible damage.

Luna Innovations, in Roanoke, Va., specializes in creating antitamper features for FPGAs. Princeton's Lee suggests that Luna's approach may involve narrowing down the number of possible unspecified functions. "There are ways to determine where such hardware would be inserted," she says. "Where could they gather the most information? Where would they be least likely to be noticed? That is what they're looking for." She compares chip security to a barricaded home. The front door and windows might offer vaultlike protection, but there might be an unknown window in the basement. The Luna researchers, she speculates, may be looking for the on-chip equivalent of the basement window.

Raytheon, of Waltham, Mass., has expertise in hardware and logic testing, says Collins. He believes the company will use a more complex version of a technique called Boolean equivalence checking to analyze what types of inputs will generate certain outputs. Normally, applying specific inputs to a circuit will result in specific, predictable outputs, just as hitting a light switch should always cause the light to turn off. "Now look at that process in reverse," says Collins. Given a certain output (the lights go out), engineers can reconstruct what made it happen (someone hit a switch). Collins says this could help avoid cycling through infinite combinations of inputs to find a single fatal response.

In January, the performers were given a set of four test chips, each containing an unknown (to them) number of malicious insertions. Along with a thorough description of the chips, Collins says, "we told them precisely what the circuits were supposed to be."

Each team's success will be gauged by the number of malicious insertions it can spot. The goal is a 90 percent detection rate, says Collins, with a minimum of false positives. The teams will also have to contend with red herrings: to trip them up, the test set includes fully functioning, uncompromised chips. By the end of this month, the performers will report back to DARPA. After Johns Hopkins has tallied the results, the teams will get a second set of test chips, which they'll have to analyze by the end of the year. Any performer that doesn't pass muster will be cut from the program, while the methods developed by the successful ones will be developed further. By the program's end in 2010, Collins hopes to have a scientifically verifiable method to categorically authenticate a circuit. "There's not going to be a DARPA seal of approval on them," says Collins, but both the Army and the Air Force have already expressed interest in adopting whatever technology emerges.

Meanwhile, other countries appear to be awakening to the chip threat. At a January hearing, a U.S. House Committee on Foreign Affairs addressed Pakistan's ongoing refusal to let the United States help it secure its nuclear arsenal with American technology. Pakistan remains reluctant to allow such intervention, citing fears that the United States would use the opportunity to cripple its weapons with—what else?—a kill switch. ❑

SCAVENGER HUNT

DARPA'S PROGRAM has formed teams to test chip integrity. USC's group creates the chips, which MIT's group then compromises with unknown additions. The "performers," Xradia, Luna Innovations, and Raytheon, are supposed to find the malicious alterations. And the Johns Hopkins group judges the results. The program's three phases get progressively harder, with the number of insertions increasing and the testing time decreasing.

TEST-ARTICLE TEAM
USC Information Sciences Institute creates test chips.

RED TEAM
MIT Lincoln Labs inserts malicious circuits.

PERFORMER
Xradia studies X-ray analysis.

PERFORMER
Luna Innovations studies FPGAs.

PERFORMER
Raytheon studies design process, ASICs, and FPGAs.

METRICS TEAM
Johns Hopkins Applied Physics Lab develops ways to measure success.

TO PROBE FURTHER For a comprehensive look into the failure of the Syrian radar, see "Cyber-Combat's First Shot," Aviation Week & Space Technology, *26 November 2007 by David A. Fulghum, Robert Wall, and Amy Butler. For more on the DARPA program and on the kill-switch debate, go to http://spectrum.ieee.org/may08/chiptrust.*