

Compromised By Design? Securing the Defense Electronics Supply Chain

John Villasenor
November 2013

Executive Summary

John Villasenor is a nonresident senior fellow in Governance Studies and the Center for Technology Innovation at Brookings. He is also a professor of electrical engineering and public policy at UCLA.

Electronic “chips” are found everywhere—not just in critical defense systems, but also in the broader infrastructure for power, finance, communications, and transportation. All of these systems function effectively only when the electronic circuits at their heart can be trusted to operate as intended.

Unfortunately, ensuring trust has become much more difficult in recent years. Concern over the growth of counterfeit electronics (parts that have been harvested from discarded systems, relabeled, and sold as new to unsuspecting buyers) has grown in recent years.¹ These parts can fail prematurely, with potentially disastrous consequences. Thanks to recent congressional attention, improved detection methods, and heightened screening requirements for parts destined for defense systems, however, the threat of counterfeits is being actively addressed.

Yet the supply chain is almost completely unprotected against a threat that may turn out to be more significant in the long term: Chips could be *intentionally* compromised during the design process, before they are even manufactured. If placed into the design with sufficient skill, these built-in vulnerabilities would be extremely difficult to detect during testing. And, they could be exploited months or years later to disrupt—or exfiltrate data from—a system containing the compromised chip.

As chips have gotten more complex and design teams have grown larger and more globalized, the opportunities to insert hidden malicious functionality have increased. If the history of cybersecurity has taught us anything, it is that these opportunities will be exploited. The prudent question, therefore, is not “will intentionally compromised hardware will end up in the defense electronics supply chain?” but “how do we maintain security when it inevitably does?” This paper aims to help frame the discussion regarding how best to respond to this important and underappreciated aspect of cybersecurity.

Introduction: The Growth of Chip Complexity

It is almost impossible to overstate the importance of integrated circuits—or “chips”—to modern society. We rely on them to run the Internet, the power grid, the financial markets, laptop computers and mobile phones, food distribution networks, medical equipment, and an essentially endless list of other networks and devices. Defense systems are similarly reliant on chips. To be effective, the military needs the ability to gather, analyze, and move information, and, when necessary, to move troops, supplies, and weapons systems. In today’s military, chips are essential to all of these tasks.

Thanks to continued advances in electronics, chips have become staggeringly complex. In accordance with Moore’s Law,² the number of transistors—the basic electrical elements used for implementing logic functions—that can be built into a single chip has roughly doubled every two years since the 1960s. In the early 1970s, a chip with several thousand transistors was considered large; today, chips with well over a billion transistors are routine. These advances are directly responsible for many of the conveniences we have learned take for granted, including the ability to perform Internet searches and to carry smartphones having more computational power than a roomful of 1960s-era computers.

Yet this complexity, which has so enhanced our capabilities, also complicates trust. In the days when the largest chips had only a few thousand transistors, it would have been nearly impossible for a rogue designer to maliciously compromise circuitry without being detected. Post-manufacturing testing could explore nearly all of the functionality of the chips in that era. In addition, the people performing the testing had often been deeply involved in the design. Just as a longtime resident can know every street and building in a small town, the designers of a small chip could know the role of every circuit, and easily recognize any unauthorized changes.

A cyberattack launched using a chip containing compromised circuits could:

- Exfiltrate data while making the chip appear to function normally
- Corrupt data within the chip
- Stop the chip from functioning

Hardware-based cyberattacks:

- Harder to conduct than software attacks, since far fewer people have the necessary skills and access
- Harder to defend against, since replacing corrupted hardware can be extremely difficult and expensive

Today’s chips have become so complex that no single person can understand every detail of a chip’s design. Even the fastest automated testing methods would take many years to exhaustively test everything that a modern large chip can do. To avoid this obviously impractical outcome, testing is done on a statistical basis. A small fraction³ of all possible inputs are provided to the chip, and the resulting observations are used to infer behavior even for inputs that weren’t

specifically tested. Historically, this approach has worked quite well. The laws of probability ensure that a properly designed suite of tests will be extremely effective at

identifying accidental design flaws. However, intentionally hidden alterations inserted by a skilled designer would be much harder to find. In particular, latent functionality configured to be dormant until triggered months or years later would be invisible to testing protocols built on the assumption that all of the possible behaviors of a chip can be easily explored during the verification process.

Hardware: A Gaping Cybersecurity Exposure

For all the attention paid in recent years to cybersecurity, it remains largely software-focused, both in terms of the techniques employed and the expertise of the people and companies working in the field. This is a blind spot; hardware represents a gaping and exploitable hole in the current approach to cybersecurity. While software cybersecurity remains critically important, a complete cybersecurity strategy now requires consideration of hardware as well.

The varied means of attack illustrate how hardware-level vulnerabilities can be exploited to completely sidestep software-based security countermeasures. For example, a team of university researchers recently demonstrated that carefully chosen alterations in portions of a chip involved in encryption processing could allow an attacker to extract encryption keys.⁴ These compromises did not require the addition of any additional circuitry, but instead involved introducing subtle modifications in the electrical behavior of certain transistors in the chip. The modifications would not be noticed by an unsuspecting observer, and would not be detectable by any of the software running on the chip. But an attacker with specific knowledge of where to look could exploit them to decrypt data coming off the chip.

Another possible attack, demonstrated in a 2011 paper,⁵ involves corrupting the circuitry responsible for governing the data movement within a chip. When activated, the effect would be similar to turning all of the traffic lights in a city to red: Data movement on the chip would simply grind to a halt. Inspection of the software running on the system containing the chip would not provide any substantive insight into the nature of the problem. In some systems, a software-based inspection would not even be capable of identifying the chip as the location of the problem. And even if the chip were suspected, software-based techniques would typically be unable to pinpoint the circuitry within the chip responsible for the problem, and would be unable to rectify it.

Yet another attack involves intentional corruption of data within a chip. The corruption could be designed to act only after receipt of a specific externally delivered trigger. Or, it could be triggered automatically by the arrival of a pre-programmed date, by use of the chip in a defense facility of particular interest, use in close proximity to certain GPS coordinates, or any combination of the above. The presence of the malicious circuitry would be nearly impossible to detect during pre-deployment testing.

The impact of a hardware-based cyberattack could range from relatively modest to catastrophic. In some cases, an attack might cause a relatively minor malfunction that, while undesirable, could initially be handled using existing protocols put in place to

detect and mitigate non-intentionally-caused failures.⁶ However, many attacks would be far more alarming. An attack aimed at enabling exfiltration and/or eavesdropping of defense/intelligence communications networks would pose grave national security concerns. An attack targeting navigation or control systems in defense aircraft, ships, submarines, or land vehicles could put lives at risk.

Attacks targeting systems for industrial control or manufacturing automation could cause extensive physical and economic damage, and even lead to loss of life. In 1999, the failure of a control system—believed to have been accidental—contributed to a gas pipeline rupture that killed three people and injured multiple others.⁷ In early 2003, a computer worm left safety monitoring systems in an Ohio nuclear plant inoperable for several hours.⁸

Unfortunately, the lists of potential attacks and U.S. defense (and non-defense) systems that could be targeted are almost infinitely long. And, as explained above, detecting intentional compromises is much harder than in the past due to the size of today’s chips. But how would a rogue designer insert malicious functionality in the first place? To answer that question, it’s helpful to consider how chips are designed and how the dynamics of the global semiconductor market have changed in recent years.

Building a Chip

While the details of creating a new chip are complex, at the highest level the process involves a small number of basic steps. The first is to specify what the chip will need to do. This includes identifying the functions it will perform, how fast it will need to perform them, and constraints on power, size, and cost. The next step is design, which entails mapping the desired functionality first into a set of logical operations, and then into the corresponding electrical circuits. A description of the completed design is then sent to a semiconductor manufacturing facility for fabrication into an actual, physical chip. The chip is subjected to testing, and if there are no problems identified, it can be shipped to customers and incorporated into a product.⁹




Figure 1: Steps in creating a chip

In the earliest days of the semiconductor industry, a single company would often do the specification, design, manufacturing and testing. Some companies, including IBM, Intel, Samsung, and Texas Instruments still operate in this manner.¹⁰ However, the costs of building manufacturing facilities—sometimes referred to as fabrication facilities, or more commonly, “fabs”—have gone from extremely expensive to stratospheric. Back in the 1980s, a fab could cost over \$200 million in 2013 dollars.¹¹ As technology advances enabled the production of chips with ever-smaller features, semiconductor manufacturing equipment got more advanced—but also more expensive. Fab costs, which are closely

coupled to the purchase prices of manufacturing equipment, have risen accordingly. In late 2012, Samsung broke ground on a new fab in Xian, China that will cost \$2.3 billion¹² initially and involve a total investment of \$7 billion¹³ by the time the project is completed. Gartner, Inc. has projected that “[i]ncreasing costs of manufacturing equipment will drive the average cost of semiconductor fabs to between \$15 billion and \$20 billion by 2020.”¹⁴

The prohibitive costs of building an in-house manufacturing capability spurred the growth of services enabling semiconductor companies without their own fabs to send their designs to an external facility known as a “foundry” for manufacturing. (A “foundry” is the term used for a fab—or, equivalently, a manufacturing facility—that can be accessed on a contract basis by companies that have designed a chip, but need access to external manufacturing capacity in order to build it). This made it possible to start a semiconductor company without incurring the immense capital costs associated with constructing a fab. Both Qualcomm,¹⁵ which was founded in 1985, and Broadcom,¹⁶ which was founded in 1991, were built using this model. Today, the ranks of “fabless” semiconductor companies include not only giants like Qualcomm, Broadcom, AMD,¹⁷ and Nvidia, but also hundreds of smaller companies, some with only a few dozen employees.

Manufacturing: A Recognized Security Concern—But Only Part of the Problem

As the semiconductor industry grew through the 1990s, cost pressures pushed more and more manufacturing of American chip designs to offshore foundries. While this raised some concerns related to intellectual property security among American semiconductor companies focused on the consumer and business markets, it created a particularly acute challenge for those designing chips for sensitive defense and intelligence systems. Classified designs, of course, could not be shipped overseas for manufacturing. And even unclassified designs can contain sensitive information that could raise national security concerns if improperly disclosed.¹⁸

To address this, in 2004, the Department of Defense and the National Security Agency jointly funded a “trusted foundry” at a preexisting IBM semiconductor manufacturing facility in Vermont.¹⁹ The Trusted Foundry Program is administered by NSA’s Trusted Access Program Office²⁰ and aims “to ensure that mission-critical national defense systems have access to leading-edge integrated circuits from secure, domestic sources.”²¹ The program has now grown to over 50 accredited suppliers,²² including over a dozen trusted foundries as well as companies offering design, test, and other services.²³ While the original focus of the program was manufacturing, there are now multiple participating suppliers focused purely on providing trusted design services.

While the Trusted Foundry Program has been vitally important for enabling the secure production of the most sensitive chips, it is used for only a small fraction of the chips in defense systems. When purchasing computers, routers, navigation and communications

equipment and most other electronics hardware, The Department of Defense (DoD) is heavily reliant on the commercial supply chain—and therefore exposed to any associated vulnerabilities.

In addition, from the standpoint of someone wishing to insert malicious functionality, the manufacturing stage—which has been the traditional focus of chip security concerns—is not the weak link in the chain. Designs are provided to manufacturers as descriptions of the shapes and locations of all the silicon and metal structures that must be built into the chip. It is possible, but very expensive and time consuming, to reverse engineer the full functionality of a chip from the information provided to a manufacturer. Attempting to insert malicious functionality by directly modifying the description of on-chip structures would be difficult, and in some (though not all)²⁴ cases would create easily detectable defects. The task facing an attacker is much easier if he or she can get access at an earlier stage of the supply chain, when the design is still being created.

The Chip Design Ecosystem: More Globalized, More Complex

The semiconductor industry has evolved significantly over the last decade. In 2003, according to an analysis by PwC, global semiconductor expenditures for consumption (as opposed to production) were \$166 billion, in non-inflation-adjusted terms.²⁵ In 2003, Japan accounted for about 23% of this total, and each of the Americas, Europe, China, and the rest of the world (ROW) contributed between 18% and 20%.²⁶ By 2012, the global market had grown to slightly under \$300 billion. China had risen to the dominant position, with 52.5% of the total.²⁷ The Americas, Europe, and Japan accounted for only 12.4%, 11.1%, and 7.3% respectively.

Revenues related to semiconductor production have also experienced both growth and offshore migration, including to countries that are not U.S. alliance partners. In China, chip manufacturing (including packaging and testing) revenues increased from \$3.7 billion in 2003 to over \$24 billion in 2012.²⁸ Chip design revenue to companies in China was \$130 million in 2000 and \$540 million in 2003.²⁹ By 2012, it had grown to \$9.87 billion³⁰—an amount, after adjusting for inflation, over 14 times higher than in 2003 and over 56 times higher than in 2000.³¹ India has experienced rapid growth in its semiconductor design market as well. According to a 2011 report from the India Semiconductor Association, semiconductor hardware design revenues in India were over \$1.4 billion in 2010 and expected to grow to over \$2 billion by 2012.³²

Statistics aside, anyone who has spent significant time interacting with semiconductor companies in recent years has seen plenty of anecdotal evidence that design has become a much more globalized—and less American—endeavor. The same cost pressures that pushed semiconductor manufacturing offshore during the 1990s have been acting during the past decade to push design offshore. In the early 2000s, fabless American semiconductor companies would often perform much of the circuit design in house and outsource the manufacturing to a foundry in Asia. Today, American companies do more

and more of their design overseas. In some cases, this is due to outsourcing. Alternatively, American companies will sometimes establish an overseas branch or acquire an overseas company with the specific goal of getting access to foreign design talent. This makes eminent economic sense. Salaries for highly skilled designers in China, India, Eastern Europe, and South Korea are lower than in the United States. And in global economic terms, the inflow of design revenue in these and other places has become an important contributor to economic growth and prosperity.

For highly sensitive American defense and intelligence applications, these changes raise challenges. It is extremely unlikely that the U.S. government would want the design of chips destined for use in such applications outsourced to China. And, to be fair, it is equally unlikely that the Chinese government would be comfortable entrusting the design of similarly sensitive chips to an American company. However, the global nature of the semiconductor design industry can leave governments with less choice than they might like over such matters. Chip design has become so globally interconnected that, for all but the most narrowly tailored applications and systems, there is no longer any economically practical way to avoid complex international supply chains.

Furthermore, it is the complexity of today's chip design ecosystem, as much as its international nature, that complicates the security picture. Just as a city has residential, commercial, and industrial areas, the real estate on a chip is partitioned into different sections, often called "blocks," with different functions. Some blocks of a chip are devoted to memory. Another block might be used to decode a JPEG file into an image that can be displayed on a screen. Often, a chip will contain a programmable block of circuitry that can be instructed using software to do many different things at different times as the processing needs change. In addition, in a chip, just as in a city, plenty of real estate is devoted to facilitating movement. The analog of city streets in a chip is the network of microscopic interconnecting metal lines that allow data to move within and among the various blocks.

A company leading the design of a large chip will typically rely not only on its own engineers, but will also obtain a significant fraction of the design by purchasing blocks created by other companies. Purchased blocks arrive not as physical pieces of silicon, but as files of computer-like code expressing the logic that will eventually be converted into a physical circuit. The lead company assembles all of these design files into a single description of the entire chip. Next, the company performs a set of computer simulations aimed at confirming that the chip will behave as expected. If problems are found, the design is modified and checked again through another round of simulations. When all of the identified problems have been ironed out, the design is sent to a fabrication facility for manufacturing.³³

Why Design Corruption Is a Growing Threat

There are multiple security exposures in this process. One concern lies in the large number of organizations and people involved in the design of a single large chip. In addition to the lead company, there are companies subcontracted to provide design

services or provide pre-made designs. Some subcontracted companies may in turn farm out some of their work through a further level of subcontracting. Pieces of the chip design are stored and exchanged using a myriad of computers and networks, some of which may be insufficiently protected against external intrusions—potentially allowing an attacker to hack in and corrupt a previously healthy portion of the logic. There is also the risk of an insider threat among the dozens or hundreds of engineers with access to the design.

The overwhelming majority of designs will emerge from this process without any intentionally introduced flaws. But there are over 5000 new chips designed each year³⁴ in a globe-spanning ecosystem involving thousands of companies and hundreds of thousands of people. To conclude that chip designs will *never* be intentionally compromised flies in the face of much of what the last twenty years of cybersecurity have taught us. Design corruption is a very real, growing threat for multiple reasons:

1. The laws of statistics guarantee that there are people with the skills, access, and motivation to intentionally compromise a chip design

Globally, hundreds of thousands of people are employed as chip designers.³⁵ The overwhelming majority of them aim to produce the best, most effective designs possible. But some small percentage could be induced to intentionally compromise a design. Even if only 1/10 of 1% of chip designers would consider corrupting a chip for financial gain, competitive advantage, or other reasons, that still corresponds to hundreds of people with exactly the right skills and access. It would defy logic to assume that none of them will ever try.

2. A skilled attacker could compromise a design in a manner minimizing the chance of detection

An unskilled attacker who engaged in wholesale replacement of commonly used functionality within a critical block of a chip would easily be detected during the pre- or post-manufacturing testing. But a skilled attacker with access to detailed information about the targeted chip would be well positioned to identify the locations within the chip where carefully placed malicious circuitry would likely remain undetected. As noted above, chips have become so complex that testing is only partial. Malicious circuitry designed to lay dormant and avoid impeding the normal operation of the chip would be extremely difficult to detect using verification protocols aimed at accidental, as opposed to intentional, design flaws.

3. The threat of attribution is not a sufficiently strong disincentive

One of the most commonly articulated arguments against the risk of intentionally compromised chips holds that the attacker would easily be identified once the flaws were discovered. Thus, the argument goes, the near certainty of getting caught would dissuade a would-be attacker from actually compromising a chip in the first place. This logic is incorrect for several reasons.

First, a skilled attacker could introduce a flaw with plausible deniability (see point 4 below). Second, just as occurs in traditional software cyberattacks, the source of a design attack could be disguised. Consider a subcontracted design company with weak network security, allowing an outsider to break in and introduce malicious circuitry into the design. Years later, the compromise might be identified and traced to the company. The company's executives and designers—many of whom may have moved on to other jobs by that point in time—could be interviewed, and all would express bafflement about how the chip could have been corrupted. The computer records that might help identify a network intrusion as the attack vector might be long gone, and the attacker would be off the hook.

More complex combinations of the above are also possible. For example, an insider intent on corrupting a design could create a software vulnerability in his or her employer's networks, and then introduce the design corruption by accessing those networks from an anonymized, off-site location.

Despite all of the above, a thorough after-the-fact investigation of a malicious design alteration may eventually be able to identify the person who introduced it and prove that it was intentional. All of that, however, would occur downstream, potentially years after the design was corrupted. As long as the attacker believed—perhaps incorrectly—that he or she was unlikely to get caught, the threat of attribution would be a weak disincentive.

4. A skilled attacker could introduce a flaw with plausible deniability

Consider a skilled attacker who builds a back door into the circuitry in a chip, with the intent of exploiting it years later to exfiltrate data or configure the chip to impede its function. What would happen if the back door were discovered and studied, and the designer who inserted it identified and confronted? He or she could respond to accusations of tampering by characterizing the back door as a feature that was introduced to assist in testing early prototypes of the chips. The intent, he or she could claim, was to remove it before high volume manufacturing. But when there are a billion transistors to think about, some things fall through the cracks. Leaving the back door open, the attacker could say, was simply a mistake—albeit one that other people could later attempt to exploit for malicious purposes.

To make matters more complex, honest designers sometimes make mistakes that can increase the vulnerability of a chip to being called into service in a later hardware-based cyberattack. In some cases, disentangling genuine mistakes from intentional changes intended to weaken security could be extremely difficult.

5. An attacker could afford to cast a wide net, knowing that only a tiny fraction of the corrupted chips would end up in systems of interest

In many cases, a single design is used to manufacture millions of identical copies of a chip, only a tiny fraction of which might end up in systems of interest to an attacker. Knowing this, the attacker could cast a wide net by modifying the design to provide an

extremely small back door in every one of these systems, with the full knowledge that most of the back doors would never be exploited. The back door could be designed so that a software- or firmware-based probe could be used later to query the chip and obtain information regarding the system in which it is installed.

Growing Recognition of the Threat

While the globalization of semiconductor design is a more recent trend, U.S. government concerns about the threat of intentionally corrupted chips were initially spurred by the offshore migration of manufacturing that occurred during back in the 1990s and early 2000s. In October 2003, then-Deputy Secretary of Defense Paul Wolfowitz wrote in a widely distributed unclassified memo that the “country needs a defense industrial base that includes leading edge, trusted commercial suppliers for critical integrated circuits used in sensitive defense weapons, intelligence, and communications systems.”³⁶ A few months later, at DoD request, the Defense Science Board convened a task force on “High Performance Microchip Supply.”³⁷ The task force’s report, published in 2005, offered a blunt conclusion:

The Department of Defense and its suppliers face a major integrated circuit supply dilemma that threatens the security and integrity of classified and sensitive circuit design information, the superiority and correct functioning of electronic systems, system reliability, continued supply of long system-life and special technology components.³⁸

The report also stated that “[t]rust cannot be added to integrated circuits after fabrication; electrical testing and reverse engineering cannot be relied upon to detect undesired alterations in military integrated circuits.”³⁹

In 2007, the Defense Advanced Research Projects Agency (DARPA) launched⁴⁰ the Trust in Integrated Circuits⁴¹ program to address the challenges of “ensuring IC hardware works as intended and ensuring the design and fabrication process can be trusted.”⁴² This was followed, in 2011, by DARPA’s Integrity and Reliability of Integrated Circuits⁴³ program and by an Intelligence Advanced Research Projects Activity (IARPA) program in Trusted Integrated Circuits.⁴⁴

These programs have generated important technological advances in the ability to secure the chip supply, and in some cases have addressed not only manufacturing concerns but also design vulnerabilities. But the threat landscape, particularly with regard to chip design, is continuing to evolve, and the scale of the challenge is immense. Even if sufficiently effective technological solutions are found, they will have limited impact in the absence of a broader national strategy and an accompanying set of policy responses.

In 2013, there have been encouraging signs that supply chain integrity is gaining visibility within the broader cybersecurity dialog. In February, President Obama issued an Executive Order⁴⁵ requiring, among other things, that NIST develop a “Cybersecurity Framework” addressing risks to critical infrastructure. While the scope of the Order is

very broad, a discussion draft of the framework, while primarily focused on software, identified the supply chain as one of the areas needing increased attention,⁴⁶ observing that “[s]upply chain risk management, particularly in terms of product and service integrity, is an emerging discipline characterized by diverse perspectives, disparate bodies of knowledge, and fragmented standards and best practices.”⁴⁷

A June 2013 report from Semiconductor Research Corporation and the Computing Community Consortium, “Research Needs for Secure, Trustworthy, and Reliable Semiconductors,” focused specifically on emerging challenges to ensuring hardware integrity.⁴⁸ The report identified an “overarching need . . . for research in ‘Design for Security,’ ” and identified seven specific research challenges requiring attention.⁴⁹ The report concluded that “[n]ow is the time to launch a collaborative program of research with industry and government support in ‘design for security.’ ”⁵⁰ Technology measures are a critical aspect of the solution. But a complete solution will also require reframing the policy discussion to better reflect today’s electronics industry.

Addressing the Threat: Some Guiding Principles

Against this backdrop, here are some principles that should guide the development of a more comprehensive strategy to secure the defense electronics supply chain against intentionally compromised hardware:

- **Many of the most significant security exposures today lie in chip design; focusing on manufacturing alone is no longer sufficient**

The current approach to supply chain security still largely reflects a decade-old view that the most significant vulnerabilities lie in manufacturing. This is no longer the case. As chip complexities have increased, the vulnerabilities have expanded upstream in the supply chain, to include design. Corrupting a chip during manufacturing first requires determining which subset of the millions or billions of transistors to target. To obtain this knowledge, a malicious manufacturer would either need to engage in a lengthy and time consuming reverse engineering effort, or, more likely, to obtain inside information from someone directly involved in the design process. By contrast, design-stage corruption is easier, less expensive, and offers a far broader range of opportunities to an attacker.

- **Outsourcing is only part of the concern**

To the extent that offshore outsourcing can reduce the ability to maintain oversight into design integrity for defense electronics components, it is a concern. But given the complexity of today’s design environment, it would be a mistake to assume that onshore designs are of necessity less vulnerable to corruption than those created offshore. Intentional design corruption requires either 1) the ability to hack into the computer systems used to hold a design as it is being created, or 2) access to an insider willing to insert malicious circuitry. Geography offers little protection against either of these attack vectors.

- **Trust should not be assumed**

Electronics components destined for defense systems are often subjected to an initial screening process aimed at ensuring they are not counterfeit, and that they operate as designed. Once those tests are passed, trust is assumed. However, a skilled attacker could embed latent malicious functionality and triggering it long after the system containing the part was deployed. To address this, systems should be designed to actively assess trust of their components throughout their service lifetimes.

- **Most of the global semiconductor market is outside the U.S.**

A 2009 White House-directed “Cyberspace Policy Review” concluded with respect to supply chain concerns that, somewhat circularly, “[t]he best defense may be to ensure U.S. market leadership through continued innovation that enhances U.S. market leadership and the application of best practices in maintaining diverse, resilient supply chains and infrastructures.”⁵¹

However, the U.S. now accounts for less than 13% of the global market for semiconductor consumption.⁵² Semiconductor production (including design) is also largely based outside the U.S. A practical supply chain security strategy needs to reflect the economic reality that most defense systems, and nearly all non-defense systems, will need to rely on chips that were designed and built, at least in part, outside the U.S.

- **The security of the commercial supply chain needs more attention**

The greatest supply chain security exposure for defense applications comes not from the small fraction of chips designed and manufactured uniquely for defense systems, but from the massive inflow of commercial chips into those systems. While the designers of chips for sensitive defense and intelligence applications have long recognized concerns raised by the hardware supply chain, in the commercial world cybersecurity attention is almost never directed to the supply chain. In order to better protect DoD systems as well as the broader non-defense infrastructure, it is critical to raise the level of attention regarding the commercial chip supply as potential cyberattack vector.

- **Design practices within the semiconductor industry should be changed**

Design practices in the semiconductor industry should be modified to specifically recognize and address the potential for malicious hardware insertion. Companies engaged in chip design should adopt a greater level need-to-know partitioning of information.⁵³ When subcontracting design work to external parties, companies should do diligence not only the design skills of the contractor, but their procedures for ensuring design security.

- **Responses to hardware-based cyberattacks should be formulated in advance**

To the extent that there is not already a well-formulated plan that would guide the response to a significant, hardware-based cyberattack on defense and/or other critical

infrastructure systems, one should be developed. The governmental entity that would be charged with overseeing the response to a hardware attack should be identified preemptively, and procedures should be put in place for reporting an attack and for engaging with the appropriate companies and governmental organizations.⁵⁴

Making Hardware Cybersecurity a Reality

Addressing hardware cybersecurity will require action at multiple levels on the part of various governments around the world. Much can be done to improve the level of government awareness and preparedness regarding hardware-based cybersecurity threats. In the U.S., the government generally and DoD in particular should expand its current software-centered approach to give more attention to the threat from maliciously modified hardware. Currently, awareness of the threat is high within some subsets of the DoD research community (notably DARPA and IARPA), but at the operational level, cybersecurity everywhere—both within and outside of government—is still considered a software-only concern. An indirect measure of this can be found by looking at the companies that provide cybersecurity products aimed at government and corporate customers, the overwhelming majority of which appear to be focused on software. Other important government steps include development of a response plan for addressing significant hardware-based attacks, and procurement incentives aimed at encouraging suppliers to actively mitigate hardware cybersecurity risks.

Companies that make chips have a role to play as well. As noted above, they can change their design practices to reduce the likelihood of chip corruption. They can more carefully vet their design suppliers and more carefully track design changes as a chip is developed. Chipmakers should also consider building countermeasures into chips that can help identify and respond to an attack as it occurs. Importantly, many of these effective countermeasures could be introduced at negligible additional cost.⁵⁵

Companies that provide design tools—i.e., the software products used by hardware designers during the chip design process—can also give increased attention to the potential for malicious design insertion. Design tools can be improved to 1) enable more careful tracking of design changes, 2) enable increased testing for intentionally introduced flaws, and 3) recognize that different blocks within a chip may have different levels of trustworthiness. The coming years will also see significant private-sector opportunities to provide hardware cybersecurity expertise for government and corporate customers. This will require a very different set of technical skills than those used in software cybersecurity environments.

Conclusions

Too often, we wait for catastrophe to spur change. As there has not yet been a string of publicly disclosed examples of defense hardware with malicious design alterations, it is hard to spur interest in investing significant effort to address the inevitability of intentionally compromised hardware. But given the critical role of chips in nearly every

defense system, there are good reasons to be proactive as opposed to purely reactive with respect to hardware cybersecurity.

None of this means that hardware cybersecurity will require the same level of effort and expense that has been directed to software cybersecurity. Software has always been, and will remain, the more significant vulnerability. But the commonly held view that software is the *only* vulnerability is out of step with the reality of how today's systems are designed and built.

At the 2011 Aspen Security Forum, retired Gen. Michael Hayden, who formerly headed both the CIA and NSA, said with respect to compromised hardware, "Frankly, it's not a problem that can be solved . . . This is a condition that you have to manage."⁵⁶ Unfortunately, this dark assessment is accurate. But while the bad news is that the problem can't be completely solved, the good news is that there is enormous opportunity to more proactively manage and mitigate the risks.

Acknowledgements

I would like to thank Chris Daverse of Defined Business Solutions, Peter Singer of Brookings, and Ian Wallace, also of Brookings, for providing comments on an earlier draft.

¹ Recycled components are the largest, but not the only, types of counterfeit parts. See John Villasenor & Mohammad Tehranipoor, *Chop Shop Electronics - Clever Counterfeiters Sell Old Components as New, Threatening Both Military and Commercial Systems*, IEEE SPECTRUM, Sept. 20, 2013, <http://spectrum.ieee.org/semiconductors/processors/the-hidden-dangers-of-chopshop-electronics>.

² Intel co-founder Gordon Moore was one of the first people to observe the exponential growth in chip complexity. In 1965, while he was at Fairchild Semiconductor, he wrote a paper suggesting a doubling every year. See Moore, Gordon E., "Cramming More Components onto Integrated Circuits," Reprinted from ELECTRONICS, Vol. 38, No. 8, pp. 114-17 (Apr. 9, 1965)," PROCEEDINGS OF THE IEEE, Vol. 86, No. 1, pp. 82-85 (Jan. 1998), available at <http://www.cs.utexas.edu/~fussell/courses/cs352h/papers/moore.pdf>. In 1975, he offered a revised prediction of "a doubling every two years, rather than every year." See 1965 – "Moore's Law" Predicts the Future of Integrated Circuits, COMPUTER HISTORY MUSEUM, <http://www.computerhistory.org/semiconductor/timeline/1965-Moore.html> (last visited on Oct. 5, 2013).

³ The number of possible inputs is so large that exhaustive testing is impossible.

⁴ Zeljka Zorz, *Researchers create undetectable layout-level hardware Trojans*, HELP NET SECURITY, Sept. 17, 2013, <http://www.net-security.org/secworld.php?id=15589>.

⁵ Kim, Lok-Won & Villasenor, John D., *A System-On-Chip Bus Architecture for Thwarting Integrated Circuit Trojan Horses*, IEEE TRANSACTIONS ON VLSI SYSTEMS, Vol. 19, No. 10, pp. 1921-1926 (Oct. 2011), available at <http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=5556060>.

⁶ In the ensuing investigation, it would of course be alarming to discover that the failure was intentionally caused. Such a discovery could require an expensive and complex response to rid similar systems of the corrupted parts.

⁷ Ryan Singel, *Industrial Control Systems Killed Once and Will Again, Experts Warn*, WIRED.COM, Apr. 9, 2008, <http://www.wired.com/threatlevel/2008/04/industrial-cont/>.

⁸ Id.

⁹ In the interest of a simple, readable description, this high-level characterization of how a chip is made omits many details and potential variations.

¹⁰ This does not mean that companies such as TI manufacture all of their chips in house. For example, TI has outsourced some manufacturing to UMC. See United Microelectronics Corp. Press Release, UMC

Recognized for Excellence by Texas Instruments, Apr. 27, 2009, <http://www.umc.com/English/news/2009/20090427-1.asp>.

¹¹ See Thomas C. Hayes, *Texas Instruments In Gamble*, N.Y. TIMES, Mar. 29, 1986, <http://www.nytimes.com/1986/03/29/business/texas-instruments-in-gamble.html>, noting that “Texas Instruments’ efforts have included spending \$200 million to build wafer-fabrication plants at its headquarters in the Dallas suburb of Richardson and in Miho, Japan.” While the article does not break down the cost across the two plants, one of the plants must have cost at least \$100 million. Using CPI numbers published by the Social Security Administration at *Average CPI By Quarter And Year*, SOCIAL SECURITY ADMINISTRATION, <http://www.ssa.gov/oact/STATS/avgcpi.html> (last visited on Sept. 1, 2013), \$100 million in 1986 corresponds to about \$210 million in 2013 dollars.

¹² Samsung Press Release, *Samsung Breaks Ground for Memory Manufacturing Complex in China*, Sept. 12, 2012, <http://www.samsung.com/global/business/semiconductor/news-events/press-releases/detail?newsId=12043>.

¹³ Id.

¹⁴ Johnson, B., Freeman, D., Christensen, D. & Wang, S.T., *Market Trends: Rising Costs of Production Limit Availability of Leading-Edge Fabs* [Abstract] (ID: G00238123), GARTNER, INC. (Sept. 1, 2012), available at http://www.gartner.com/DisplayDocument?doc_cd=238123.

¹⁵ *Who We Are: History*, QUALCOMM INC., <http://www.qualcomm.com/who-we-are/history/story> (last visited on October 5, 2013).

¹⁶ *Facts at a Glance*, BROADCOM CORP., <http://www.broadcom.com/docs/company/BroadcomQuickFacts.pdf> (last visited on Oct. 5, 2013).

¹⁷ AMD’s manufacturing arm was spun off in 2009 to form GLOBALFOUNDRIES. See *About GLOBALFOUNDRIES*, GLOBALFOUNDRIES INC., <http://www.globalfoundries.com/about/> (last visited on Oct. 5, 2013).

¹⁸ A separate, related issue that arises for all chips, not only those designed specifically for defense applications, is the intellectual property contained in the chip design. Chip designs are prime targets for economic espionage.

¹⁹ McCormack, Richard, *\$600 Million Over 10 Years For IBM’s ‘Trusted Foundry’ Chip Industry’s Shift Overseas Elicits National Security Agency, Defense Department Response*, MANUFACTURING & TECHNOLOGY NEWS, Vol. 11, No. 3 (Feb. 3, 2004), available at <http://www.manufacturingnews.com/news/04/0203/art1.html>.

²⁰ *Trusted Access Program Office (TAPO)*, NATIONAL SECURITY AGENCY, July 1, 2013, <http://www.nsa.gov/business/programs/tapo.shtml>.

²¹ *DoD Trusted Foundry Program*, ACQUISITION COMMUNITY CONNECTION, Oct. 23, 2011, <https://acc.dau.mil/CommunityBrowser.aspx?id=480636>.

²² Ortiz, Catherine, *DOD Trusted Foundry Program – Ensuring “Trust” for National Security & Defense Systems* [PowerPoint presentation], TRUSTED FOUNDRY PROGRAM, at pg. 17 (June 20, 2013), available at <http://www.ndia.org/Divisions/Divisions/SystemsEngineering/Documents/Past%20Meetings/06-20-12%20Division%20Meeting/04%20-%20Trusted%20Foundry%20Program,%20Ortiz,%20NDIA%20SE%20Talk%2020120620%20cjo2.pdf>.

²³ For a map showing companies and services provided, see id. For a list of trusted foundries, see id. at pg. 18.

²⁴ Given sufficiently detailed knowledge of the design, there is a class of intentional circuit corruptions that could be introduced without creating easily detectable defects. See., e.g., *Researchers create undetectable layout-level hardware Trojans*, supra note 4.

²⁵ See *Continuing to grow: China’s impact on the semiconductor industry 2013 update* [Report] (the “2013 PwC Report”), PRICEWATERHOUSECOOPERS LLP, at pg. 3 (Sept. 2013), available at <http://www.pwc.com/gx/en/technology/chinas-impact-on-semiconductor-industry/assets/china-semicon-2013.pdf>.

²⁶ Id.

²⁷ Id.

²⁸ Id., at pg. 13. The chart separately considers “manufacturing” and “packaging and testing.” In 2003, manufacturing and packaging/testing comprised a total of 44.6% of \$8.3 billion, corresponding to \$3.7 billion. In 2012, manufacturing and packaging/testing comprised a total of 43.3% of \$56.3 billion, corresponding to \$24.4 billion.

²⁹ Id., at pg. 17.

³⁰ Id.

³¹ The PwC figures for chip design revenue to companies in China of \$130 million in 2000 and \$540 million in 2003 are not inflation adjusted. Using CPI numbers published by the Social Security Administration at *Average CPI By Quarter And Year*, SOCIAL SECURITY ADMINISTRATION, <http://www.ssa.gov/oact/STATS/avgcpi.html> (last visited on Sept. 27, 2013), \$130 million in 2000 corresponds to \$174 million in 2012 dollars. \$540 million in 2003 corresponds to \$679 million in 2012 dollars. By contrast, 2012 revenue was \$9.87 billion, over 14 times as high as the inflation-adjusted 2003 amount. The inflation-adjusted ratio of 2012 revenues to 2000 revenues is about 57.

³² *Study on semiconductor design, embedded software and services industry* [Report], INDIA SEMICONDUCTOR ASSOCIATION & ERNST & YOUNG GLOBAL LTD., at pg. 24 (Apr. 2011), available at http://deity.gov.in/hindi/sites/upload_files/dithindi/files/Semiconductor06April11_020511_0.pdf. The report states that total “semiconductor design industry” revenues were \$7.5 billion in 2010. However, that includes \$6.076 billion for embedded software development. “VLSI design” and “board design” were reported as \$944 million and \$476 million respectively; adding these figures gives \$1.42 billion. For the projected 2012 total of \$10.6 billion, the VLSI design and board design components were \$1.33 billion and 6.72 million respectively.

³³ The full design process involves not only ensuring that the design logic is correct, but that it will function as expected when, after a process called “layout,” it is mapped into the structures of the target semiconductor process.

³⁴ Ramamoorthy, Ganesh, *Market Trends: ASIC and ASSP Chip Design Start Trends, Worldwide, 2013* [Abstract] (ID: G00246635), GARTNER, INC. (Aug. 16, 2013), available at <http://www.gartner.com/id=2575215>.

³⁵ Countries with very large numbers of people involved in chip design include the United States, India, and China. In China, for example, the 2013 PwC Report, *supra* note 25 at pg. 17, noting that “The total number of employees in the IC design sector increased by 6% in 2012 to about 112,500 . . .”

³⁶ *Defense Science Board Task Force On High Performance Microchip Supply* [Report], OFFICE OF THE UNDER SECRETARY OF DEFENSE FOR ACQUISITION, TECHNOLOGY AND LOGISTICS, at pg. 87 (February 2005), available at <http://www.acq.osd.mil/dsb/reports/ADA435563.pdf>.

³⁷ Id., at pg. 83.

³⁸ Id., at pg. 5.

³⁹ Id., at pg. 3.

⁴⁰ Sally Adee, *Contracts awarded for DARPA’s Trust in Integrated Circuits program*, IEEE SPECTRUM, Dec. 6, 2007, http://spectrum.ieee.org/tech-talk/semiconductors/devices/contracts_awarded_for_darpas_t.

⁴¹ The program is now referred to as the “Trusted Integrated Circuits program.” See *Microsystems Technology Office: Trusted Integrated Circuits (Trust)*, DARPA, [http://www.darpa.mil/Our_Work/MTO/Programs/Trusted_Integrated_Circuits_\(TRUST\).aspx](http://www.darpa.mil/Our_Work/MTO/Programs/Trusted_Integrated_Circuits_(TRUST).aspx) (last visited on Oct. 5, 2013).

⁴² *TRUSTED in Integrated Circuits Program BAA07-24 Industry Day Notice* (Solicitation No. SN07-24), DARPA, Mar. 16, 2007, <https://www.fbo.gov/index?s=opportunity&mode=form&id=56160bfade0d4086dbd06abb45668b18>.

⁴³ Adam Rawnsley, *Can Darpa Fix the Cybersecurity ‘Problem From Hell?’*, WIRED.COM, Aug. 5, 2011, <http://www.wired.com/dangerroom/2011/08/problem-from-hell/>. DARPA’s description of the program is at *Microsystems Technology Office: Integrity and Reliability of Integrated Circuits (IRIS)*, DARPA, [http://www.darpa.mil/Our_Work/MTO/Programs/Integrity_and_Reliability_of_Integrated_Circuits_\(IRIS\).aspx](http://www.darpa.mil/Our_Work/MTO/Programs/Integrity_and_Reliability_of_Integrated_Circuits_(IRIS).aspx) (last visited on Oct. 5, 2013).

⁴⁴ Michael Cooney, *US intelligence group wants to change the way chips are made*, NETWORK WORLD, Nov. 1, 2011, <http://www.networkworld.com/news/2011/110111-iarpa-chips-252624.html>.

⁴⁵ Exec. Order No. 13,636, 78 Fed. Reg. 11,739 (Feb. 19, 2013), available at <http://www.gpo.gov/fdsys/pkg/FR-2013-02-19/pdf/2013-03915.pdf>.

⁴⁶ *Discussion Draft of the Preliminary Cybersecurity Framework* [Report], NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY (Aug. 28, 2013), at pg. 11, available at <http://www.acuta.org/wcm/acuta/legreg/083013b.pdf>.

⁴⁷ Id., at pg. 13.

⁴⁸ *Research Needs for Secure, Trustworthy, and Reliable Semiconductors* [Report], COMPUTING

COMMUNITY CONSORTIUM & SEMICONDUCTOR RESEARCH CORP. (June 2013), available at <http://www.src.org/emerging-initiative/cybersecurity/>.

⁴⁹ Id., at pg. 16.

⁵⁰ Id., at pg. 19.

⁵¹ *Cyberspace Policy Review: Assuring a Trusted and Resilient Information and Communications Infrastructure* [Report], THE WHITE HOUSE, at pg 34 (May 29, 2009), available at http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_Review_final.pdf. The report document is undated and does not clearly identify specific agencies involved in writing the report at White House direction. However, a May 29, 2009 White House press release accompanying the report's release states that "In February 2009, President Obama directed the National Security Council (NSC) and Homeland Security Council to conduct a 60-day review of the plans, programs, and activities underway throughout government that address our communications and information infrastructure (i.e., 'cyberspace'), in order to develop a strategic framework to ensure that the U.S. government's initiatives in this area are appropriately integrated, resourced, and coordinated" and that "The review team's report to the President contains five main chapters, outlined below, and includes a near-term action plan for U.S. Government activities to strengthen cybersecurity." See The White House Press Release, *Cyberspace Policy Review: Assuring a Trusted and Resilient Information and Communications Infrastructure*, May 29, 2009, <http://www.whitehouse.gov/the-press-office/cybersecurity-event-fact-sheet-and-expected-attendees>.

⁵² According to the 2013 PwC Report, *supra* note 25, at pg. 3, the Americas as a whole accounted for 14% of the global semiconductor consumption in 2011 and 12.4% in 2012. The report does not specify the subset of those percentages attributable specifically to the United States.

⁵³ The recommendation for more effective information partitioning was also discussed in a 2011 Brookings paper. See Villasenor, John D., *Ensuring Hardware Cybersecurity* [Report], THE BROOKINGS INSTITUTION, No. 9 (May 2011), available at <http://www.brookings.edu/research/papers/2011/05/hardware-cybersecurity>.

⁵⁴ Some of the text in this recommendation was previously published in a 2011 Brookings paper. See *id.*

⁵⁵ Ways to actively monitor the chip behavior and identify suspicious activity were discussed in a 2011 Brookings paper (see Villasenor, *Ensuring Hardware Cybersecurity*, *supra* note 53) and, in more detail, in a 2010 Scientific American article (see Villasenor, John, *The Hacker in Your Hardware* [Preview], SCIENTIFIC AMERICAN, Vol. 303, No. 2, pp. 82-87 (Aug. 4, 2010), available at <http://www.scientificamerican.com/article.cfm?id=the-hacker-in-your-hardware>).

⁵⁶ Rawnsley, *supra* note 43.

About the Centers

B | Center for **Technology Innovation** at BROOKINGS

Founded in 2010 and led by Director Darrell West, the Center for Technology Innovation (CTI) in Governance Studies at Brookings focuses on delivering research that impacts public debate and policymaking in the arena of U.S. and global technology innovation. At CTI, our research centers on identifying and analyzing key developments to increase innovation; developing and publicizing best practices to relevant stakeholders; briefing policymakers about actions needed to improve innovation; and enhancing the public and media's understanding of technology innovation.

B | Center for **21ST Century Security and Intelligence** at BROOKINGS

The Center for 21st Century Security and Intelligence (21CSI) in Foreign Policy at Brookings was created to address the key issues shaping security policy over the coming decades. The Center seeks to answer the critical questions emerging in defense, cybersecurity, arms control, and intelligence in an all-encompassing manner, seeking not just to explore important new policy challenges but also how they cross traditional fields and domains. Under the leadership of Peter W. Singer, one of the world's leading experts on changes in warfare, the Center focuses on delivering cutting-edge research, analysis and outreach that shapes public understanding and official decision-making across a broad range of security issues.

The views expressed in this piece are those of the author and should not be attributed to the staff, officers or trustees of the Brookings Institution.