The Die is Cast: Hardware Security *is Not* Assured

EDLYN V. LEVINE, PH.D.

In 2011 a fictitious company was created by the U.S. GAO (Government Accountability Office) to gain access to vendors of military-grade integrated circuits (ICs) used in weapons systems. Upon successfully joining online vendor platforms, the GAO requested quotes for bogus part numbers not associated with any authentic electronics components. No fewer than 40 offers returned from vendors in China to supply the bogus chips, and the GAO successfully obtained bogus parts from a handful of these vendors.³ The ramifications of the GAO findings are stark: The assumption of trusted hardware is inappropriate to invoke for cybersecure systems.

Injection of counterfeit electronics into the market is only a subset of vulnerabilities that exist in the global IC supply chain. Other types of attacks include trojans built into the circuitry, piracy of intellectual property, and reverse engineering. Modern ICs are exceptionally

complex devices, consisting of upwards of billions of transistors, miles of micron-scale interconnecting wires, advanced packaging configurations, and multisystem integration into chips sized on the order of a U.S. quarter. These ICs are designed, manufactured, and assembled by an equivalently complicated, globally distributed supply chain. A semiconductor company can have more

From Specs and Sand to Semiconductors: How ICs Are Made

Break open your laptop and you will find on the order of 100 to 1,000 ICs. These range from the CPU to microprocessors to memory. Each of these circuits has crossed the globe multiple times, moving among geographically distributed supply-chain vendors during their evolution from an initial specification to final assembly as a component in the machine sitting in your home or office. IC manufacturing can be broken into three primary stages—design, fabrication, and assembly and testing—each of which presents opportunities for hardware to be altered or assembled systems to be compromised.

SPECIFICATIONS AND DESIGN

Designing a new IC begins once the desired specifications for the chip are established. The specs determine the required performance of a chip for a targeted environment, including function, power, size, and timing. Semiconductor design is typically undertaken by teams of engineers who translate the IC specification into an RTL (register transfer level) description of the circuit in an HDL (hardware description language) such as VHDL (Very High-speed Integrated Circuit HDL) or Verilog. The RTL description is synthesized into a gate-level netlist using the logic gates and components from the desired technology library. The netlist is then converted to the transistor level with a fully placed and routed physical

layout (shown in a GDSII file, the standard format used to represent the layout) using EDA (electronic design automation) software, thereby completing the circuit description.

Design is undertaken by both IDMs (integrated device manufacturers) that own fabrication facilities and fabless semiconductor companies that outsource semiconductor manufacturing. Throughout the design process, engineers incorporate IP (intellectual property) from external vendors. The third-party IP companies develop and license circuit blocks, called IP cores, that are integrated into the overall design of a new chip. IP cores can take the form of synthesizable RTL or of a GDSII representation of the fully placed and routed core design. Leading IP vendors can have their IP cores included in tens of billions of chips manufactured each year.

FABRICATION

Completed GDSII files are sent to a semiconductor fabrication facility, called a foundry, for manufacturing. Foundries are either owned and operated by IDMs or exist as stand-alone fabrication companies contracted by fabless semiconductor companies. GDSII files are converted by the foundry or a third party into mask sets that are used for patterning the physical circuit layout into layers in a silicon wafer during photolithography.

The full fabrication process includes multiple steps of material deposition, etching, and patterning, along with the processes of ion implantation and annealing that fine-tune electrical properties of the integrated elements. Once the transistor level has been fabricated, patterned metal wires are deposited to link transistor elements. The geometrical configuration of these interconnections is optimized for the functional specification of the chip, with complex ICs having upwards of 20 metal layers. A completed fabricated wafer is tested and cut into individual silicon chips (dies) that are shipped for assembly and further testing.

ASSEMBLY, TESTING, AND DISTRIBUTION

The packaging of individual silicon dies creates a protective interface between the die and the external environment. Package integration incorporates the silicon die with package wiring, substrates, heat spreaders, and ground planes, thereby creating the required electrical, mechanical, and thermal environment for the chip to interface properly with an external system. The packaged ICs are tested, binned according to performance, and distributed to electronics assembly plants that incorporate the ICs into end-user products.

than 16,000 suppliers spread around the world.¹⁰ While globalization has drastically reduced industry costs by tapping inexpensive labor markets and economies of scale, it has simultaneously opened many windows of opportunity for attackers to maliciously modify hardware without the knowledge of ODMs (original device manufacturers) or their customers.

The tenet that “trust starts in silicon” underscores hardware as the root of security upon which software protections are implemented. Secure systems cannot be architected on a foundation of compromised hardware. Unlike software, there is no patch update that can fix a malicious hardware insertion short of replacing the device. Securing hardware is a multifaceted problem consisting of shoring up the manufacturing chain, developing robust means to detect malicious insertions, and designing systems to be secure against the inevitability of hardware compromise.

Innovative research efforts spanning DARPA's TRUST (Trusted Integrated Circuits) program to its LADS (Leveraging the Analog Domain for Security) program emphasize the increasing spotlight on hardware security as do high-profile reports ranging from the Defense Science Board to the President's Council of Advisors on Science and Technology. Modern economies and critical systems depend on IC technologies, making the ramifications of hardware attacks increasingly dire.

THE SPECTRUM OF INVASIVE HARDWARE ATTACKS

An invasive hardware attack consists of changing the physical layout of a single IC or assembly of ICs. Specific classes of attacks include hardware trojans that modify the layout of a legitimate IC during design and fabrication, counterfeit attacks that substitute an illegitimate chip for a legitimate one, and assembly attacks that include incorporating additional ICs in the end-user device. (This last type of attack was the subject of a now-famous 2018 *Bloomberg Businessweek* article concerning data-center motherboards.⁸ Even if the events of that article are not verifiable, the attack described represents a realistic threat vector.)

An invasive attack seeks to incorporate a malicious capability in an end-user device. An overt attack has signatures that are potentially detectable by the targeted system once implemented. Examples include kill switches that destroy a system's function, backdoors that enable illegitimate access, and control circuitry that changes a system's behavior. A covert attack seeks to operate undetected for long periods of time, often with the

objective of collecting information to route to the attacker, and may never be detected. The execution of a hardware attack requires knowledge of how ICs are fabricated and how they can be compromised.

Semiconductor manufacturing includes hundreds of steps from specification to distribution, providing many opportunities for invasive attacks (see sidebar). Counterfeit attacks and assembly attacks are conducted during the assembly, distribution, and second-hand supply-chain stages. Insertion of malicious hardware trojans can occur at any stage during IC manufacturing.

Trojans can be categorized according to the fabrication step at which they are inserted, to yield insight into supply-chain risk mitigation. The three classes of trojan insertion are pre-silicon, in-silicon, and post-silicon. Trojans range in their impact on IC performance (function change, backdoors, kill switches, decrease in service lifetime, information leakage), their activation mechanism (always on, internally triggered, externally triggered), their physical location on the chip (I/O, logic, memory, power distribution, clock), and the hardware abstraction level at which they occur.⁴

A pre-silicon attack occurs during the specification and design stages. A trojan can be inserted by changing functional characteristics during specification, such as timing or power consumption, or by modifying features at different hardware-abstraction layers during design, such as RTL (register transfer level), gate level, transistor level, and place-and-route. Every stage of design and every software tool used during design is a potential security vulnerability. The pervasive use of third-party IP cores and standard cell libraries in circuit design affords

increased opportunity for external parties to insert malicious functionality. Computer-aided design tools can be tampered with to create compromised IC designs.⁹ Malicious modifications can even be made during the inclusion of design for test functionality before a design is sent to fabrication.

An in-silicon attack occurs during fabrication. An attack of this type requires both detailed knowledge of and access to manufacturing stages for the targeted device. These attacks can range from editing or exchanging the masks to altering the types or concentrations of chemicals used during fabrication. Changing the fine-tuned electrical properties of IC materials can have serious impacts on the function and lifetime of the device. Altering transistor dopant concentration can impact circuit function,¹ and altered composition or dimension of interconnects can lead to increased electromigration of metal atoms and early circuit failure.

A post-silicon attack is conducted after fabrication is completed. Attacks that can occur at this stage include circuit editing, modified package-level circuitry, untrusted testing that fails to reveal trojans, package counterfeiting, and malicious assembly of trusted ICs on a printed circuit board. Assembly attacks can manifest as the inclusion of unwanted ICs or the use of unshielded connections between trusted ICs and their environment, giving rise to electromagnetic-coupling-mediated information leakage.

DETECTING INVASIVE ATTACKS

Many variants of hardware trojans can be implemented to achieve a range of attacks: from the addition of extra

Secure systems cannot be architected on a foundation of compromised hardware.

transistors creating new logic to the modification of the wire width of the clock distribution network introducing clock skew. Overt kill switches and shortening of service lifetimes to covert backdoors and information leakage also have different activation mechanisms. Some trojans are always on, whereas others require either internal or external triggers for attack payload activation. A universal objective for all trojans, however, is to escape detection throughout manufacturing and deployment until the trojan's attack is executed.

A trojan is designed to be of minimal size and consume minimal resources on a chip, posing a serious challenge to any effort to detect it. Because of the potential impact of hardware attacks, extensive research efforts have led to the development of sophisticated means of detecting trojans, but there is no smoking gun that ensures the trust of an IC. In principle, detection can be accomplished either by activating the trojan and observing its impact on chip performance compared with known performance specifications, or by comparing the questionable design or fabricated chip with the physicality and functionality of a trusted [golden] copy. Methods for detecting pre-silicon attacks differ from those for in- and post-silicon attacks, the latter ranging from nondestructive to destructive.

Detecting trojans in IC designs requires evaluating and ensuring the trust of third-party IP cores, libraries, and electronic design-automation tools. This is not easy. IP cores are challenging to verify for trust since there is no golden version with which to compare. As such, establishing trust in IP cores typically takes the form of searching for unexpected components or signal output

during design performance testing. Internal verification of IP functionality and code coverage analysis is used to identify suspect components and signals.

ATPG (automatic test pattern generation) uses digital signal inputs to sequentially generate output patterns from a simulation of the designed chip. ATPG can detect trojans consisting of modifications to the known functionality of the chip, but it will not be successful finding trojans that have *added* functionality, such as additional logic, to the design. Having no information about the additional logic makes it impossible for ATPG to conduct a directed search of all possible digital signal inputs that could cause trojan activation. Furthermore, a trojan that activates physical side-channel leakage will go undetected with ATPG alone.

Once the chip has been fabricated, a new suite of trojan-detection methods is brought to bear. Sophisticated tools such as scanning electron microscopy and picosecond imaging circuit analysis can be used to do a full teardown of an IC to extract its physical layout for comparison with a trusted design. This is expensive and time consuming, resulting in partial to full destruction of the device under test, and thus is infeasible for widescale testing of chips set to enter the consumer market.

More tractable, less thorough nondestructive physical inspection and electrical testing leverage everything from x-ray imaging to parametric testing of chip behavior. Other testing methods include trojan activation via ATPG on the physical device, as well as side-channel analysis. The latter method investigates the physical characteristics of the device under test, such as timing and power consumption,

to compare with known or golden side-channel behavior. Process variations that naturally occur during the course of fabrication, however, decrease the efficacy of side-channel analysis for trojan detection.

There is as yet no assured way of definitively determining whether or not a chip has been tampered with, despite the large arsenal of testing methods. In many cases the sheer volume of ICs, as well as the lack of access to sophisticated testing equipment, hinders assurance of devices on the market. Testing is typically done by the ODMs or third-party specialists. Testing methods make heavy use of established means used by the microelectronics industry to test for device quality assurance. These techniques, including performance assessment and failure analysis, similarly extend to counterfeit and assembly attacks. Although powerful, these methods are not comprehensive, and increasing emphasis is being placed on adopting either design for security or zero trust in IC manufacturing.

BROADENING THE SPECTRUM:

SEMI-INVASIVE AND NON-INVASIVE ATTACKS

The notoriety of recent microarchitectural attacks such as Spectre and Meltdown clearly shows that the book on hardware security does not end with the supply chain. Latent vulnerabilities of trusted ICs can be taken advantage of using semi-invasive attacks such as fault injections and non-invasive attacks leveraging side channels. If you've ever been warned not to yell in a data center, you are familiar with the faults that can be introduced in disk-head readers by mechanical vibrations.

Analogous fault injection can be introduced by physical coupling or manipulation of ICs. Many examples exist, ranging from corrupted memory isolation induced by disturbance errors injected into DRAM by repeated row hammering,⁶ to violations of trusted execution environments such as Arm TrustZone, to Intel SGX [Software Guard Extensions].⁵

The physical attack plane can also be leveraged for side-channel attacks such as Spectre and Meltdown. Unintended physical or microarchitectural signatures that manifest during the operation of the IC can be leveraged by an attacker to learn information about the circuit that allows the attacker either to compromise secure data or to yield access to secure functions. This was famously first demonstrated with timing attacks.⁷ Increasingly, designing for security seeks to understand and preempt the physical signatures of ICs at the design stage to anticipate or detect side-channel security vulnerabilities that manifest in the post-fabrication stage.

THE FUTURE OF HARDWARE SECURITY

Recognition of the importance of hardware security has shifted focus from traditional software threats to lower levels of the computing hierarchy. Research across hardware security areas from supply chains to side channels has led to a better understanding of hardware threats and increased development of detection and mitigation techniques. Resources such as the TrustHub Trojan database and conferences such as IEEE's HOST (Hardware-oriented Security and Trust) and PAINE (Physical Assurance and Inspection of Electronics) are

signs of this shifting focus toward hardware security.

Despite the increased attention and growing corpus of research, no common standards or tools exist and no definitive solutions have been developed. The spectrum

Related articles

➡ Why Is It Taking So Long to Secure Internet Routing?

Routing security incidents can still slip past deployed security defenses.

Sharon Goldberg

<https://queue.acm.org/detail.cfm?id=2668966>

➡ Kode Vicious

What is a CSO Good For?

Security requires more than an off-the-shelf solution.

<https://queue.acm.org/detail.cfm?id=3357152>

➡ Building Systems to be Shared Securely

Want to securely partition VMs? One option is to put 'em in jail.

Poul-Henning Kamp and Robert Watson

<https://queue.acm.org/detail.cfm?id=1017001>

of invasive to non-invasive vulnerabilities at the physical attack plane makes hardware assurance a daunting if not insurmountable challenge. As with the rest of the cybersecurity community, hardware security benefits from the recognition that a prevention-only approach to assurance leaves systems vulnerable to successful attacks. This is analogous to a home security system solely dependent on an external fence, with no internal alarms, locks, safe rooms, or police response force should an intruder hop the barrier.

As such, focus increasingly

leans toward designing hardware capable of identifying, operating through, mitigating, and recovering from an attack.¹¹ However, the economic benefits of security often remain unclear due to the high cost of security and the prevalence of consumers who are willing to risk security for increased compute capability (or who are ignorant of the vulnerabilities).

The future of hardware security will evolve with hardware. As packaging advances and focus moves to beyond Moore's law technologies, hardware security experts will need to keep ahead of changing security paradigms, including system and process vulnerabilities. Research focused on quantum hacking is emblematic of the translation of principles of security on the physical attack plane for emerging communications and computing technologies.² Perhaps the commercial market will evolve such that the GAO will run a study on compromised quantum technologies in the not-too-distant future.

References

1. Becker, G. T., Regazzoni, F., Paar, C., Burleson, W. P. 2014. Stealthy dopant-level hardware trojans: extended version. *Journal of Cryptographic Engineering* 4(1), 19-31; <https://link.springer.com/article/10.1007/s13389-013-0068-0>.
2. Emerging Technology from the arXiv. 2014. The next battleground in the war against quantum hacking. *MIT Technology Review* (August 20); <https://www.technologyreview.com/2014/08/20/171574/the-next-battleground-in-the-war-against-quantum-hacking/>.
3. GAO. 2012. DoD supply chain: suspect counterfeit electronic parts can be found on Internet purchasing platforms. GAO-12-375; <https://www.gao.gov/products/GAO-12-375>.
4. Karri, R., Rajendran, J., Rosenfeld, K., Tehranipoor, M. 2010. Trustworthy hardware: identifying and classifying hardware trojans. *Computer* 43 (10), 39-46; <https://ieeexplore.ieee.org/document/5604161>.

5. Keegan, R. 2019. Hardware-backed heist: extracting ECDSA keys from Qualcomm's TrustZone. *NCC Group Whitepaper* (April 23); <https://www.nccgroup.com/us/our-research/extracting-ecdsa-keys-from-qualcomms-trustzone/>.
6. Kim, Y., et al. 2014. Flipping bits in memory without accessing them: an experimental study of DRAM disturbance errors. *ACM SIGARCH Computer Architecture News* 42(3), 361-372; <https://dl.acm.org/doi/10.1145/2678373.2665726>.
7. Kocher, P. C. 1996. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: *Advances in Cryptology—CRYPTO '96*, ed. N. Koblitz. Lecture Notes in Computer Science 1109. Berlin, Heidelberg: Springer; https://link.springer.com/chapter/10.1007/3-540-68697-5_9.
8. Robertson, J., Riley, M. 2018. The big hack: how China used a tiny chip to infiltrate U.S. companies. *Bloomberg Businessweek* (Oct. 4); <https://www.bloomberg.com/news/features/2018-10-04/the-big-hack-how-china-used-a-tiny-chip-to-infiltrate-america-s-top-companies>.
9. Roy, J. A., Koushanfar, F., Markov, I. L. 2008. Extended abstract: circuit CAD tools as a security threat. Published in *2008 IEEE International Workshop on Hardware-Oriented Security and Trust*, 65-66; <https://ieeexplore.ieee.org/document/4559052>.
10. Semiconductor Industry Association, Nathan Associates. 2016. Beyond borders: the global semiconductor value chain; <https://www.semiconductors.org/wp-content/uploads/2018/06/SIA->

Beyond-Borders-Report-FINAL-June-7.pdf.

11. Villasenor, J. 2010. The hacker in your hardware. *Scientific American* 303(2), 82-87; <https://www.scientificamerican.com/article/the-hacker-in-your-hardware/>.

Edlyn V. Levine, Ph.D. is Chief Engineer of MITRE Engenuity and a research associate in the Department of Physics at Harvard University. Dr. Levine is internationally recognized for her contributions in information technology as an AFCEA 40-under-40 award winner. She earned her Ph.D. in Applied Physics from Harvard University in 2016 where she was a National Defense Science and Engineering Fellow and a National Science Foundation Graduate Fellow.

Copyright © 2020 held by owner/author. Publication rights licensed to ACM.