

Survey of Approaches for Security Verification of Hardware/Software Systems

FERHAT ERATA*, Yale University, USA

SHUWEN DENG*, Yale University, USA

FAISAL ZAGHLOUL*, Yale University, USA

WENJIE XIONG*, Yale University, USA

ONUR DEMIR[†], Yeditepe Üniversitesi, Turkey

JAKUB SZEFER*, Yale University, USA

This paper surveys the landscape of security verification approaches and techniques for computer systems at different levels: from a software-application level all the way to the physical hardware level. Different existing projects are compared, based on the tools used and security aspects being examined. Since many systems require both hardware and software components to work together to provide the system's promised security protections, it is not sufficient to verify just the software levels or just the hardware levels in a mutually exclusive fashion. This survey highlights common sets of system levels that are verified by the different existing projects and presents to the readers the state of the art in hardware and software system security verification. Few approaches come close to providing full-system verification, and there is still much room for improvement. In this survey, readers will gain insights into existing approaches in formal modeling and security verification of computer systems, and gain insights for future research directions.

CCS Concepts: • **Security and privacy** → **Formal security models**; *Security in hardware*.

Additional Key Words and Phrases: Formal methods, theorem provers, model checkers, security verification, processor architectures

1 INTRODUCTION

News articles and opinion pieces by top security researchers constantly remind us that as computing becomes more pervasive, security vulnerabilities are more likely to translate into real-world disasters [74]. Computing systems today are very complex systems, and if the design of the hardware, software, or the way the hardware and software interact are not perfect, then there may be security vulnerabilities that attackers can exploit.

To help find these potential vulnerabilities and prove the designed system is trustworthy, formal methods can be used. For instance, in the development of eXecute Only Memory (XOM) [59], with formal verification, a possible replay attack was identified and then the designed was improved and proved to be secure. Since the security of the systems depends on the correctness of the protections that both the hardware and software components provide, there is the need to verify the security of both the hardware and the software.

Many “secure architectures” such as XOM have been designed to provide enhanced security features in hardware, however, only a few of them come with formal verification. In academia, these include [18, 28, 44, 47, 51, 59, 83, 84, 92]. For these secure architecture they mostly do not come with any formal specifications or proofs for security. In industry, only a small number of designs from hardware vendors provide some hardware features for security, e.g. ARM TrustZone

*This work was supported in part by the National Science Foundation (NSF) grants 1419869 and 1524680; and Semiconductor Research Corporation (SRC) contract 2015-TS-2633. Shuwen Deng was supported through Google PhD Fellowship.

[†]The author's work is supported by TUBITAK grant 2219.




Fig. 1. Hardware and software levels found in a typical computer systems are shown on the left. The projects and levels which they consider are shown on the right. The verification projects surveyed in this work focus typically on one or more of the levels. Broadly, projects either focus on multiple software levels, or hardware levels; some projects span both software and hardware levels, but not all the levels. Note, some projects “skip” certain levels, as indicated by the breaks in the arrows on the right-hand side of the figure.

[88], Intel SGX [63], and most recently AMD Memory Encryption [1]. These designs also all rely on the assumption that the hardware is correct – the industry designs we know about do not have any publicly available formal security specifications nor proofs. With formal verification, the designers could prove the system design and implementations are secure and trustworthy.

To help promote more security verification of computer systems, this survey aims to show readers about current approaches to security verification of computer systems. In this survey, we compare different projects that consider both the hardware and the software levels of a system, and which use formal methods to verify security properties of such systems. We show the state of the art in security verification and lower the barrier to entry into this field for interested researchers.

1.1 Software and Hardware System Levels Considered in Verification Process

A computer system is typically composed of multiple hardware and software levels, as shown in Figure 1. The typical software levels are: Application, Operating System (OS), and Hypervisor. These levels cover typical software running on a commodity computing system. The typical hardware levels in a computer system are: ISA (Instruction Set Architecture), Microarchitecture, RTL (Register Transfer Level), Gate, and Physical.

Traditionally, upper levels depend on the lower levels for functionality and security. E.g., a guest OS relies on the Hypervisor to provide isolation from other malicious guests, if the more privileged Hypervisor has a security vulnerability, the OS can not make any guarantees about security. At the hardware level, for example, ISA is not secure if the microarchitecture that implements it has a bug; and a microarchitecture realized using a flawed RTL that implements it is likewise not secure, and so forth. The relationship is not strictly linear in that upper level always depends on all lower levels. Some of the secure architectures have introduced hardware that allow higher software levels to be protected from intermediate software levels. For example, in Bastion [18] applications are able to communicate with the Hypervisor while bypassing the OS; or in HyperWall [84] a virtual machine does not need to rely on hypervisor for isolation as the hardware provides some of the basic memory management functionality. Thus, the security verification approach needs to consider which levels are important for security verification.




Fig. 2. General procedure for security verification.

The software and hardware levels needed for ensuring security of the system constitute the trusted computing base (TCB) which contains all the software and hardware that need to be trusted (although may not be trustworthy). Thus, the TCB should be verified for security to make it truly trustworthy. Effectively, TCB consists of different components in different levels, and security verification tools and methods should include all the levels in when checking the security of the system. Figure 1 on the right side shows the various surveyed projects and the different system levels that their security verification covers. Because different projects consider different levels, it may be difficult to select the right approach (or mix of approaches) for security verification that one may desire. Especially, some of the works skip certain levels, which may not be needed for their verification, for example Moat [81] verifies applications with respect to ISA, and assumes hardware fully protects the applications from OS or Hypervisor, so OS and Hypervisor levels are skipped. By studying each group of projects, this survey aims to show the state of the art in security verification of hardware/software and allow researchers and practitioners to understand how to better approach security verification of their designs. We also show limitations of current approaches, and suggest new or worthwhile research directions to help achieve computer system verification methodologies that can check all the levels in TCBs of today's and tomorrow's systems, not just a subset of them.

2 GENERAL TOOLS AND MECHANISMS

This section presents a background on the different tools, mechanisms, and approaches typically needed to check security guarantees. The general flow of the security verification process is shown at a high level in Figure 2. The starting point is the actual system, either an already existing system or a design of some new system whose security properties need to be verified. From the actual system, or design, a representation of the system needs to be obtained in the verification tools, (a) in Figure 2. In parallel, the security properties of the system need to be specified, (b) in Figure 2. The security properties are closely tied to the system's assumed threat model. The security properties can be specified separately or together within the representation of the system, in which case (a) and (b) would be done together. The final step is the actual verification process which takes the system representation and security properties as input, and returns whether the verification passed or failed, (c) in Figure 2. If the verification fails, the design needs to be updated and re-evaluated, (d) in Figure 2.

2.1 System Representation

In order to check if a system complies with some properties, we need a representation of the system that accurately expresses the behavior of the system. Ideally, the actual system description can be used, such as the source hardware description language (HDL) code for hardware components, or a programming language source code for software components. Otherwise, a model in the

verification tool is needed. One reason a model may be needed is that the way a system is described in HDL or programming language may not be compatible with the verification tool that is being used, or the way the system is described is too complex for the verification process to handle.

Hardware components can be described with Hardware Description Languages (HDLs). The most popular HDLs are Verilog [85] and VHDL [60]. Recently, a new tier of HDLs is emerging that feature more high-level abstractions and reusability than circuit-level HDLs. This new generation of HDLs is called Hardware Generation Languages (HGLs) [62], including Chisel [5], BlueSpec [68], and Genesis2 [77]. There are some attempts to link hardware and circuit system representation methods with security-property specification towards security verification, which we will present in Section 3.

Software components can be described by their high-level implementation in programming languages such as C, C++, or Java. There are also ongoing efforts to integrate design processes and system specification with security-property specification and include verification information inside programming languages as annotations. Examples include TAL [22] and Dafny [52].

Alternatively, some tools require a model of the system based on its original description. For example, this survey later discusses VeriCoq which is a tool that can be used to translate (annotated) Verilog code to code understood by Coq verification tools. If an automated method of creating a model is not available, then model has to be created manually by engineers. However, when creating models manually, proving the correspondence between the model and the actual system is an open research problem.

Formal verification is done with respect to a system representation, as described above. Most projects assume a trusted compiler or tool chain such that the system realization indeed matches the system representation, and does not contain extra hidden, or unwanted, functionality that may compromise the security of the system. For example, after verifying the C code of an application, there is still a concern that the compiler may not generate the correct machine code from the C code. A malevolent compiler might insert malicious code into the binary, as demonstrated in [86], where a virus-infected compiler is able to inject back-doors into applications during compilation. A number of projects include “trusted” compilers that are guaranteed not to inject behavior that was not specified. One example of such a compiler is CompCert [53] which is a certified compiler that generates binaries from Coq code. All surveyed work assumes trusted toolchains.

2.2 Security Properties Representation

Depending on the verification mechanism, either *deductive mechanism* or *algorithmic mechanism* in Section 2.3 and 2.4, the security properties are represented in the corresponding verification tools. In *deductive mechanisms*, the security properties can be represented in terms of logical formula. A logical formula serves as a limitation on the states the system is allowed throughout its execution. Some specialized forms of logic are used to express the relations between the states of the system. In *algorithmic mechanisms* security properties can be expressed as invariants within a system, and their validity is checked against all possible execution paths.

2.3 Formal Verification via Deductive Mechanisms

When using deductive mechanisms, verification is achieved by deducing properties from a system representation. Theorem provers fall in this category. The key element in deductive mechanisms is a proof. Deductive mechanisms use formal proofs to verify that a system complies with some given properties.

Theorem provers (also referred to as proof assistants) aid the verification process by providing frameworks for creating a mathematical model of the system, for specifying the security properties, and for formally proving whether the model complies with the properties or not. Theorem provers

are generally composed of a language (such as Coq), and an environment for describing the proofs (such as CoqIDE). There is a number of proof assistants used actively in academia and industry such as: Coq [10], Isabelle/HOL [69], PVS (Prototype Verification System) [70], ACL2 (A Computational Logic for Applicative Common Lisp) [46], and Twelf (LF) [39]. Theorem proving typically requires a lot of effort and time to complete, and learning the required tools is seen as one of the difficult aspects of verification using theorem provers. In the following paragraphs we will introduce different theorem provers and give examples on their usage in functional and security verification.

2.4 Formal Verification via Algorithmic Mechanisms

Algorithmic mechanisms typically use an algorithmic search, which is performed over a system's representation and its states, rather than using deduction. Model checkers, SMT (Satisfiability Modulo Theories) Solvers, and Symbolic Execution fall in this category.

According to [6] “model checking is an automated technique that, given a finite-state model of a system and a formal property, systematically checks whether this property holds for a given state in that model.” The security property that is being verified has to be defined using a logical form. After the model and the property definition, the model checker can be run to see if the given security property is valid in the system model. The checks can be done either for each transition or each state using invariants, pre- and post-conditions. The execution time of the model checker is determined by the invariants and the complexity of the model. The output can be positive (property satisfied), negative (property violated), or the execution runs indefinitely. There are a number of model checking tools: SPIN [43], Murφ [27], SMV [64], CBMC [49]. For further details about model-checker design, we point the readers to an early survey by Clarke et al. [20]. Model checking has a well-known *state explosion problem*, which is the exponential growth rates of states. This may lead to memory insufficiency or extremely long run times. For fairly complex systems, model checking needs to use more abstraction to simplify the model. However, as the level of abstraction gets higher, we run the risk of missing some important details of the system design.

Satisfiability Modulo Theories (SMT) solvers are used to solve satisfiability problems expressed in first-order logic regarding some logical theory. For verifying systems, first they have to be transferred into formulas that SMT solvers can work with. The validity of the property is then checked by SMT solvers [9]. Note that SMT solver execution time can vary from a few seconds to hours depending on the size of the problem. There are many verification tools using SMT solvers. For instance, Z3 [25] is an SMT solver used by an intermediate verification language, Boogie [94], and Dafny [52] is a programming language and verifier for functional correctness that uses Boogie as its target language. Vale [14], a language for expressing and verifying high-performance assembly code, uses Dafny/Z3 as its verifier backend.

Unlike model checking, which requires a model of the system, symbolic execution [7] deals directly with the program semantics assuming symbolic values for inputs. It thus arrives at expressions in terms of those symbols and constraints them with the possible outcomes of each conditional branch. Finally, program expressions can be evaluated by solving the constraints (e.g. by an SMT-Solver [25]). In this way, all possible execution states can be evaluated simultaneously, at great cost of storage and slow execution during analysis. There are many symbolic execution engines targeting different software levels: while KLEE [16] symbolically executes programs in the LLVM Intermediate Representation (IR) [50], Angr [78] executes lifted binary programs in VEX IR [67], and JPF [41] executes Java byte code.

3 SECURITY VERIFICATION FOCUSING ON THE HARDWARE LEVELS OF A SYSTEM

In this section, we present projects which focus on the security verification of the hardware levels. As listed in Section 2, the formal verification tools have their own languages with formal semantics

for describing system specification, security properties or invariants, and for doing the verification and proofs. However, these languages differ significantly from common HDLs that are used in hardware design.

To bridge the gap, researchers either manually or automatically convert system designs in HDL to system representations in formal verification tools' languages; or, conversely, generate HDL specification from the system representation written in a verification tool's language. With the former approach, hardware designer can design the system in HDL as usual, and main new effort is in generating models in the verification tools' language and describing the security specification. With the latter approach, hardware designers need to learn the language used by the verification tools and develop hardware, as well as the security specification and proofs with that language. The tools then generate HDLs from the system representation, so the system can be synthesized normally with existing tool chain.

The following classes of projects are discussed next: manually modeling systems in verification tools' languages, automatically converting system designs in HDLs to models in verification tool's languages, adding verification features to existing HDLs, or generating HDLs from system model written in verification tool's language.

3.1 Approaches Requiring Manually Modeling Systems in Verification Tools' Languages

Since the languages used in verification tools are usually very different from HDL. For security-critical modules, sometimes it makes sense to manually model the system and verify the model with security specifications. Depending on the size of the system, this process is time-consuming. Moreover, this process does not guarantee that the model faithfully represents the real system.

If a proof assistant is used, the system is modeled as a set of definitions, and the security properties are formalized as theorems. Then the proof is developed manually and checked by the proof assistant. If a model checker is used, the model of the system is built in the form of a finite state machine (FSM). The security properties are represented by a set of invariants. The model checker can automatically search for all possible states, and check if the invariants always hold. If so, the security properties are said to be proved. Usually, the model is simplified to avoid the state explosion problem.

3.1.1 Micro-Policies. A recent work on Programmable Unit for Metadata Processing (PUMP) [26] added programmable metadata processing unit alongside with the data computation. PUMP allows programmers to create policies and rules that enforce IFT mechanisms by manipulating the metadata tags in each instruction. Metadata processing can thus support many safety and security policies. However, given a high-level specification, it is nontrivial to design metadata processing rules. Whether the metadata processing rules in PUMP comply with a high-level security properties needs to be proved. Micro-Policies [4, 24] presented an approach for formalizing and verifying the IFT policies.

Micro-Policies verification process is shown in Figure 3. To design a set of metadata rules, first, an abstract machine specification with a set of instructions and information flow policies is defined, showing the security properties the machine should have. Then, programmers design the metadata rules (concrete machine), where the information flow policy is implemented into the PUMP hardware. To reason whether the concrete machine reflects the abstract machine specification, an intermediate layer *symbolic machine* is added manually, as shown by arrows labeled (1) in Figure 3. The Micro-Policies prove the equivalence by backward refinement, which means if there is a state transition in low-level machine, there exists a corresponding transition in high-level machine. They use Coq to formally prove whether the concrete machine backward refines the symbolic




Fig. 3. Micro-Policies verification process.

machine (2), and whether the symbolic machine backwards refines the abstract machine (3). If the backward refinement verification in both (4) and (5) passes, then the concrete machine has the security properties of the abstract machines.

The work shows the proof of a variety of security policies, including noninterference, sealing, compartmentalization, control flow integrity, and memory safety. The whole verification cost about 17.7k lines of code. To apply this methodology to other architectures, abstract, concrete, and symbolic machines need to be specified by the designer manually for each architecture. The refinement proofs depend on the system and also need to be re-done. Currently there is no programmatic way to generate these from the HDL code. Reusability of this approach is low.




Fig. 4. Processor cache security verification process using $\text{Mur}\varphi$ model checker.

3.1.2 Cache Verification. Processor caches are integral part of any modern processor. They are small, but fast memory components that are used to provide quick access to frequently accessed data. Through a fixed algorithm, the cache logic decides which data to keep in the cache and which data to send back to memory if a new request comes and there is not sufficient space in the cache. Memory access timing changes depending on whether a request “hits” or “misses” in the cache. Based on this timing difference between hits and misses, researchers have presented numerous side-channel attacks, e.g., [61], that are able to compromise data confidentiality and potentially leak out cryptographic keys.

In [97] researchers create side-channel leakage models based on the non-interference property between an attacker and a victim process that are using same processor cache. First, they model the cache architecture in $\text{Mur}\varphi$ as an FSM with states representing which process is currently using the cache line and transitions between the states based on cache operations (e.g. attacker cache hit, victim cache miss, etc.), (1) in Figure 4. By modeling the cache operation and transitions, the authors were able to obtain probabilities for how different operations of the victim (e.g. cache hit, cache miss, etc.) are observed by the attacker. Zhang and Lee used $\text{Mur}\varphi$ to enumerate all possible states and transitions, and count the number of interferences between attacker and victim for the different state transitions, (2) in Figure 4. Based on this data, mutual information [21] is then used to quantitatively analyze the interference between the two processes, and reveal side-channel vulnerabilities, (3) in Figure 4.

Authors of [97] applied their work to six cache architectures and revealed that most cache architectures do not satisfy the non-interference property, thus fail the verification. To apply this method to other designs, designers need to manually create the $\text{Mur}\varphi$ system representation from the cache architecture description, as there is currently no automated way to extract these models from the system representation (e.g. from HDL code). Reusability of this approach is low.




Fig. 5. Process for verifying XOM architecture with $\text{Mur}\varphi$ model checker. Two sets of states and state transitions are show, corresponding to the “actual” and “ideal” worlds that XOM verification process compares during verification.

3.1.3 XOM. The eXecute Only Memory [59] is a hardware design with embedded cryptographic functionality and access control. By adding new hardware and new instructions, XOM is able to protect user data from a malicious operating system. On-chip data is isolated using hardware tags which label the identity of the owner of the data, while off-chip data is protected by encryption and hashing. In [58], XOM was formally specified and then verified in $\text{Mur}\varphi$.

A model of XOM and its adversary is build in $\text{Mur}\varphi$ as shown in Figure 5. The model of the XOM hardware contains arrays representing the registers, cache, and memory, including data and tags; the possible values and states the hardware is modeled as system states, (1) in Figure 5. The effects of each operation of the processor are represented as state transitions. To model the effect of the adversary, two identical sets of states are used, dubbed the “actual world” and the “ideal world”. In the actual world, the adversary is modeled by a set of primitive actions she can perform as state transitions. The ideal world does not include the effect of the adversary. The actual world states and ideal world states are concatenated, and thus updated together during model checking. With the model and state transition function, $\text{Mur}\varphi$ is able to exhaustively search for all possible combinations of these actions. Invariants are defined according to the security properties to be verified, (2) in Figure 5: to prove the adversary cannot read user data, the model checker verifies that the on-chip user data is tagged with user’s XOM ID and off-chip user data is always encrypted and hashed with the user’s key. To prove that the adversary is not able to write the user data without halting the system, the model checker compares the state of the ideal world against the state actual world, and thus, knows whether the adversary will succeed.

The authors of XOM, during verification found a replay attack and fixed it. Moreover, it was shown that if the operating system does not behave maliciously the liveness of the system is guaranteed. To apply this method to other designs, designers need to manually create the $\text{Mur}\varphi$ system representation for the architecture. Especially, invariants about any tags need to be specified. Again, there is currently no automated way to extract system models from the HDL system representation. Reusability of this approach is low.

3.1.4 Formal Foundation of Enclave Secure Remote Execution. Enclave is the special kind of CPU that is able to maintain a protected memory region and make advantage of this to do operation or isolation of sensitive code and data. The formal foundation for Secure Remote Execution (SRE)




Fig. 6. Process for verifying formal models of trusted hardware platform like Intel SGX and Sanctum

of Enclaves[82] provides a framework and methodology to reason about the security guarantees provided by enclave platforms.

As is shown in Figure 6, first, secure remote execution of enclave is defined. It is used to provide formal security features to let users remotely outsource the enclave's execution including attestation and secure operations of data. Three key security properties that entail SRE are integrity, confidentiality and secure measurement. Secure measurement allows the user to verify that the platform is running unmodified enclave programs. The secure measurement property states that any two enclaves with the same measurement must also have the same semantics: they must produce equivalent execution traces for equivalent input sequences. Secondly, a Trusted Abstract Platform (TAP) is introduced to specify trusted primitives of enclaves' behavior. Along with that a parameterized attacker model is defined to verify TAP's confidentiality, etc. Proof is provided that secure remote execution holds for TAP. In the final step, the ideal TAP is refined and the refined platform is shown to have equivalent functionality and security compared with some practical enclaves like Intel SGX and MIT Sanctum. Consequently, these practical trusted hardware platforms are verified to hold SRE.

All the hardware platform models including TAP, Intel SFX and MIT Sanctum are constructed by BoogiePL and verified by Z3 SMT solver. This formal foundation is proved to be able to efficiently and effectively verify SRE of enclaves.

3.2 Approaches Requiring Using Automatic Conversion from HDL to System Models in Verification Tools' Languages

To lower the verification efforts, there are attempts to develop tools that automatically convert designs in HDL to systems models that are used in verification tools. However, the security specification and proofs still need to be done as verification effort. For certain security properties, it is possible to automatically generate security specifications and proofs.

3.2.1 VeriCoq. VeriCoq is a tool that provides mechanisms to transform Verilog code into code with PCHIP (Proof-Carrying Hardware Intellectual Property), which makes it possible to verify the security of the design written in Verilog [12]. Original VeriCoq supports an essential subset of Verilog, but requires the design to be flattened and have no nested modules. The newer VeriCoq-IFT [11] has same constraints, but adds ability to verify information flow properties automatically. The information flow proofs need "initial sensitivity list" as input and labeling the variables in the design. Given this input, VeriCoq-IFT automatically creates theorems and proofs for guaranteeing the information flow property.

The verification process is shown in Figure 7. First, the input is the Verilog code, which is then converted into Coq by VeriCoq, (1) in Figure 7. Based on the security properties requested, designers create the theorems to be verified, (2) in Figure 7. With the design represented in Coq, alongside




Fig. 7. VeriCoq and VeriCoq-IFT verification process. Dashed portion on left-hand side show the VeriCoq-IFT process that can automatically generate the theorems and the proofs.

with the theorems, developers come up with formal proofs showing that the code has the desired security properties, (3) and (4) in Figure 7. They can then verify that the design aligns with the defined security properties by using Coq, (5) and (6) in Figure 7, and either the design passes, meaning that it conforms the properties, or it fails to pass, (7) in Figure 7.

The advantage of VeriCoq is the automated conversion of Verilog code into Coq. VeriCoq-IFT also adds ability to automatically generate the theorems and proofs for information flow. To apply this method to other designs, security properties need to be specified, and the theorems and proofs developed manually they are not focusing on information flow. Reusability of this approach is medium.

3.2.2 Formal-HDL. *Formal-HDL* [45] is a hardware description language in Coq proof assistant. In [38], a tool is developed to automatically convert design in VHDL to *Formal-HDL*. Different from VeriCoq, which only allow a flattened hierarchical design (a one-level design), *Formal-HDL* supports instantiation of modules within other modules.




Fig. 8. Formal-HDL verification process, the lightly shaded portions are not yet done and are presumably future work of the authors.

The advantage of Formal-HDL is the automated conversion of VHDL code into Coq. Also, it supports instantiation of modules within other modules. However, currently no actual security verification is done using the Coq model. To apply this method to other designs, VHDL can be automatically translated to Coq, but all security verification work has to be done manually. Reusability of this approach is low (as the current work [45] does not do any actual proofs about security, just produces Coq model).

3.2.3 RTLIFT. Information Flow Tracking (IFT) has been widely used to enforce security properties, such as confidentiality, integrity, and non-interference [73]. To precisely reason about the security properties in a hardware design, RTLIFT [3] tracks information flow at the RTL level. The RTLIFT software generates extra IFT logic in the circuit and then, with standard functional verification tools, it can evaluate the security property (information flow) of the hardware design, i.e. make sure no *High* data flows to *Low* outputs. After the verification, the extra IFT logic is removed from the design – so the verification does not introduce overhead into the final circuit design.




Fig. 9. Flow of the RTLIFT verification process.

To generate IFT logic automatically, flow tracking libraries were developed for Verilog for each basic module such as a multiplexer, decoder, etc. For each basic module two types of information flow tracking were considered, along with an associated library: a precise library and a conservative library. The precise library propagates the security tags of signals in such a way as to minimize the number of false positives, while the conservative library gives smaller tracking logic with simple OR expression, but may generate more false positives. As shown in Figure 9 (1), Verilog code, along with the flow tracking libraries and specification of the desired precision (precise or conservative), is used by RTLIFT to replace each basic module in the original Verilog code with its corresponding module from the library that allows for tracking of flows through that basic module. To deal with implicit information flows in the hardware design code, conditional statements are treated as explicit multiplexers where security tags of multiplexer control signals propagate to the multiplexer output. Next, as shown in Figure 9 (2), The generated circuit with IFT features is then feed to simulation and verification tools (Quetsa Formal Verification tool in this case) to analyze whether unwanted information flows exist. If the circuit passes the verification, then the extra IFT checking logic that was added can be removed from the hardware design, while the design maintains its security properties.

Compared to tracking the IFT at the gate level [87], RTLIFT has more information about the high-level circuit, and thus, can propagate tags faster and more precisely. An RSA core, an AES core and a bus architecture were verified using this method and hardware Trojans in the designs, which leaked secret key to output, were detected. To apply this method to other designs, the information flow libraries can be re-used. Reusability of this approach is medium-high (slightly more reusable than others so far as design specific theorems or libraries need not be developed if only *High* to *Low* IFT is considered).

3.3 Adding Security Verification Features in HDL

Another approach for security verification is to add security verification features into an HDL, e.g. by either introducing a new HDL language or introducing new syntax into existing language. Caisson [56], Sapper [55], and SecVerilog [96] take this approach and introduce information flow tracking (IFT) features into an HDL language. System designers can use new syntax to specify the information flow tags and policies in their designs. If the verification passes, then designers know their designs do not have any information flow violations.




Fig. 10. Static information flow verification process of Caisson.

3.3.1 Caisson. Caisson [56] is a hardware description language for static information flow verification at the design time. The verification process is shown in Figure 10. First, the original Caisson code is written containing labeling of security tags, especially data ports in a hardware module are assigned with security labels (i.e. *Low* to *High*), (1) in Figure 10. With the design written in Caisson, and with the security labels on each register and wires, the information flow can be checked at the design time by running the Caisson compiler, (3) in Figure 10. During compilation, it is checked whether the information flow strictly follows the policy that data labeled *High* should not end up in a port labeled *Low*. If so, there will be no information flow from *High* to *Low* during the system runtime. The checking is done based on the typed Caisson language and type checking rules in the Caisson tools. Caisson can generate standard Verilog code as the output as well, with no labels, and can the code be synthesized using existing tools, (4) in Figure 10, i.e. labels are removed and have no impact on final design or performance. In the paper, authors use manual proofs to formally prove that Caisson enforces timing-sensitive non-interference in designed hardware, (2) in Figure 10.

Using Caisson, the authors were able to create the first provably information-flow secure processor that contains a time-multiplexed pipeline and a partitioned cache [56]. To apply this method to other designs, the designer needs to augment his or her Verilog code with the security labels. Reusability of this approach is medium.




Fig. 11. Dynamic information flow verification process in Sapper.

3.3.2 Sapper. Sapper [55] is a hardware description language that is based on a synthesizable subset of Verilog. Sapper compiler automatically ensures non-interference in the generated hardware logic, and is able to generate Verilog code with added dynamic information flow tags. Figure 11 shows the verification flow.

First, the Sapper code is written, which includes labeling of security tags and in particular input and output ports in a hardware modules are assigned with security labels (i.e. *Low* to *High*), (1) in Figure 11. Example of how security labels indicating an IFT policy are inserted into the code in Sapper language is shown in Figure 12. Sapper's policy is that the hardware logic should ensure that data flow to any output port never allows *High* data to reach a *Low* port. Especially, in the presence of an active attacker (e.g. a malicious software in the system), who has full control over all *Low* input ports, the non-interference enforced by the policy can protect all the data tagged with *High*. Sapper statically analyzes the hardware logic and automatically inserts dynamic IFT

	Sapper	Verilog
check	<pre> reg [7:0] a: %*\`hl{L}*; reg [7:0] b,c; a <= b & c; </pre>	<pre> reg [7:0] a,b,c; reg a_tag ,b_tag ,c_tag ; if (a_tag >= (b_tag c_tag)) a <= b & c; </pre>
track	<pre> reg [7:0] a,b,c; a <= b & c; </pre>	<pre> reg [7:0] a,b,c; reg a_tag ,b_tag ,c_tag ; a <= b & c; a_tag <= (b_tag c_tag) </pre>

Fig. 12. Example Sapper code and generated Verilog code, modeled after [55], with a security label highlighted.

logic and generates Verilog code with extra logic for the dynamic information flow tracking, (3) in Figure 11. In the paper, authors use pen-and-paper proofs to formally prove that Sapper enforces non-interference of the generated system, (2) in Figure 10.

Static analysis enables the system to cover explicit, implicit, and timing-based information flows. With the inserted IFT logic, the synthesized hardware can track and check security policy at runtime, and any policy violations will be detected. Authors designed a processor simulated the hardware with ModelSim [40]. A micro-kernel and a compiler were also implemented, and processes in different security levels could run on the processor. To apply this method to other designs, the designer needs to write his or her Sapper code with the security labels. Reusability of this approach is medium.

3.3.3 SecVerilog. SecVerilog [96] is a well-typed language and built on top of Verilog to include information flow annotations. It is first proposed in [95] to mitigate the timing channel in program execution. The language semantics can be used to analyze and formally prove the security of the system.




Fig. 13. Static information flow verification process of SecVerilog.

SecVerilog enables static checking of hardware information flows and uses an SMT checker to verify non-interference between modules with different security levels. First, the designers define a security policy, for example, the design has two security levels: *Low* and *High*. The policy may be such that the adversary, who has access to all information at or below the *Low* security level, and can measure the clock cycles of hardware operations, never has access to any data labeled *High*. Also, during the implementation of the design in SecVerilog, each variable has to be labeled with its corresponding security label, (1) in Figure 13. Example of SecVerilog labeling the code is given in Figure 14. Using these labels, SecVerilog generate models in Z3 for verification, (2) in Figure 13. Then, in Z3, the information flow is checked, and report is given whether the design passes or fails the verification, (4) in Figure 13. On the other hand, SecVerilog generate design in Verilog, (5) in Figure 13. In the paper, authors use pen-and-paper proofs to prove SecVerilog enforces timing-sensitive noninterference, (3) in Figure 10.

SecVerilog allows sharing of resources within a module. Static labeling does not solve all the problems of information flow, especially if resources are shared. In the case of shared resources,

```

1  reg [18:0] %*\h1{\{L\}\}* tag0[256], tag1[256];
2  reg [18:0] %*\h1{\{H\}\}* tag2[256], tag3[256];
3  wire [7:0] %*\h1{\{L\}\}* index;
4  //Par(0)=Par(1)=L Par(2)=Par(3)=H
5  wire [1:0] %*\h1{\{ Par(way)\}\}* way;
6  wire [18:0] %*\h1{\{ Par(way)\}\}* tag_in;
7  wire %*\h1{\{ Par(way)\}\}* write_enable;
8
9  always @ (posedge clock) begin
10    if (write_enable) begin
11      case (way)
12        0: begin tag0[index] = tag_in; end
13        1: begin tag1[index] = tag_in; end
14        2: begin tag2[index] = tag_in; end
15        3: begin tag3[index] = tag_in; end
16      endcase
17    end
18  end

```

Fig. 14. Example of the split cache in SecVerilog, modeled after [96], with the security labels highlighted.

the labels might change during runtime. SecVerilog use dependent types to handle runtime label changes. A design of split cache is shown in Figure 14 as an example. Type changes are detected and updated dynamically during the runtime, e.g. *Par(way)* in Figure 14. The dependent types can be determined by type-valued functions: For a variable v , the type of the variable can be determined dynamically during runtime by a function, e.g. *Par(v)*, line 4 Figure 14.

SecVerilogLC [30] extends SecVerilog [96] to allow more sufficient hardware resources sharing for different security levels. For dependent labels, the information flow control type system along with syntax and semantics supports signals to propagate on clock edges explicitly. In order to test next clock cycles' label, there are also related syntax supported. The type system permits registers to update values with labels securely and statically. Furthermore, SecVerilogLC explicitly divides sequential and combinational variables to do corresponding security checks. Following the changes illustrated above, it is also able to avoid implicit downgrading by explicit implementing it.

SecVerilogBL [31] [32] also extends SecVerilog [96], to support packed data structures, and downgrading mechanism. It provides an improved type system to cover the extensions. The first new feature allows complex data structures such as arrays, network packets to be tagged with finer granularity. That allows tagging of individual elements within arrays or packed data structures. The second feature supports modifying the security tag of an element dynamically.

A secure MIPS processor and caches were designed in SecVerilog [96] and SecVerilogLC [30]. Dynamic labeling makes the shared ports of the cache possible. SecVerilogBL is used to verify a secure architecture based on ARM TrustZone which provides isolated memory regions for providing confidentiality and integrity [31] [32]. SecVerilog also provides timing-sensitive non-interference, which is proved in the paper [96]. To apply this method to other designs, the designer needs to write his or her Verilog code with the annotations. Reusability of this approach is medium.

In [95], a well-typed language is proposed to mitigate the timing channel in program execution. Each command in the program is extended with security labels for confidentiality and integrity, and a new command "mitigate" is introduced to bound the execution time of another command. The language made some assumptions on the properties of the underlying hardware. The language semantics can be used to analyze and formally prove the security of the system. Meanwhile, a secure hardware architecture satisfying the properties required by the language is designed, explicitly formalized and experimentally shown to have only moderate overhead.

3.4 Generate HDL from System Model in Verification Tools

Another approach is to develop a new domain specific language, model and verify the system using this domain specific language (and associated tools), and then generate HDL. The hardware designers need to learn and use the new language, but the tools will then automatically generate HDLs, so there is one-to-one relationship between the code used for verification and the final HDL code.

3.4.1 ReWire. ReWire [72] is a functional programming language and compiler that translates high-level designs into HDL description of the hardware. It is a subset of Haskell, which produces a suitable foundation for writing formal specifications. ReWire enables modular, high-level, semantics-directed hardware circuit designs.




Fig. 15. ReWire verification flow

In ReWire, combinational logic is represented by pure, non-recursive first-order functions. The sequential logic in one clock domain is represented by a structure called “reactive resumption monad.” This structure uses tail recursive type and functions as a continuation to map an input to a “new” sequential logic. Monad is the method to produce new type with function of specific computation by incorporating fundamental data type values in a robust way. More information on monads can be found in [57]. As shown in Figure 15 (1), to generate synthesizable VHDL code from ReWire code, first the ReWire abstract syntax tree (AST) is produced by parsing Haskell concrete syntax. (2) A preHDL is generated by compiling the monadic operations and loop flattening. (3) By replacing the loop structures with VHDL processes, preHDL can be converted to VHDL.

In a sample dual-core processor with shared register, one core is designated as *High* core and the other is *Low*. In ReWire, to verify the separation between the two cores as theorems, a formal proof is written in Haskell, as in Figure 15 (4). The verification precludes storage channels, timing channels and control flow channels. In the proof, a “harness security” function enables precise control of information flow.

With ReWire, a single-core processor and a secure dual-core processor based on the single-core processor were designed and synthesized, showing that ReWire compiler can produce VHDL implementation from the high-level specification and that it supports modular design. To apply this method to other designs, the designer needs to write his or her ReWire code and specify the security properties. Formal proofs have to be done manually but the ReWire AST and preHDL steps are done automatically, so VHDL will be automatically generated from the ReWire code. Reusability of this approach is medium.

3.5 Comparison of Verification Focusing on the Hardware Levels of a System

Comparison of Caisson, Sapper, and SecVerilog:

Caisson and SecVerilog only do compile-time IFT checking. Sapper does both compile-time checking and adds run-time IFT checking into the design, at the cost of extra hardware and decreased performance. Caisson and Sapper do not allow sharing of resources, while SecVerilog adds dynamic labels which allow one module to work on both *High* and *Low* data. The static IFT

checking in SecVerilog makes sure that there is no possible combination of operations or inputs that would leak the *High* data to *Low* data – the cost is that without run-time IFT, the design has to be very conservative and consider worst-case scenarios.

Unlike Caisson [56], data with different security labels can share resources in Sapper, e.g. registers, resulting in a lower overhead.

Table 1. Comparisons of Caisson, SecVerilog and Sapper

	Caisson	SecVerilog	Sapper
Type of Information Flow Checking	Compile-Time	Compile-Time	Run-Time
Type of Label	Static Label	Dynamic Label	Dynamic Label
Prototype Circuit	Time Mux pipeline; Split Cache, Split Ports	Time Mux pipeline; Split Cache, Shared Ports	Time Mux pipeline; Split Cache, Shared Ports

A comparison is shown in the Table 1. Caisson and SecVerilog use compile-time information flow control, while Sapper uses run-time information flow control. Compile-time information flow tracking is done at the design time, and does not introduce any run-time overhead for the system. Also, designers can fix information leakage at design time, and thus do not need to worry about the effects of security violations at runtime, since no violations will appear then. Compared to Caisson [56], dynamic labeling in SecVerilog and run-time IFT in Sapper make resources sharing possible, thereby reducing area and timing overheads.

Table 2. Comparison of verification tools focusing on the hardware levels of a system.

	Verif. to HDL Code Relationship	Reusability	Source Code Public
Micro-Policies	manual	low	https://github.com/micro-policies
Cache Verif.	manual	low	no
XOM	manual	low	no
VeriCoq	programmatic	medium	no
Formal-HDL	programmatic	low	no
RTLIFT	programmatic	medium-high	no
Caisson	programmatic	medium	https://github.com/vineethk/Caisson
Sapper	programmatic	medium	no
SecVerilog	programmatic	medium	http://www.cs.cornell.edu/projects/secverilog/
ReWire	programmatic	medium	http://mu-chaco.github.io/ReWire/

3.6 Commercial Tools

Even though they are relatively new, there is a number of commercial security verification tools for hardware design. These tools are quite similar to works in Section 3.2.

Mentor Graphics Questa Secure Check. The application is part of Mentor Graphics Questa package. It receives RTL data and a spec for secure storage and paths. The spec is defined in TCL language. Secure Check then finds ports/black box inputs and generates properties for integrity and confidentiality. Black box inputs are generated in a way that it assures that no information flows outside its designated path. The application then verifies these properties. The output of the application is an exhaustive proof of integrity and confidentiality of the design and/or counterexamples showing how your spec can be violated [36].

Cadence JasperGold Security Path Verification (SPV) App. Similar to Questa Secure Check, SPV App takes RTL data and path specs. The user defines illegal sources and destinations of the data. SPV App proves that the defined secure data maintains confidentiality and integrity during operation and even after a hardware fault occurs. Verification is performed exhaustively using Jasper's path sensitization technology. Path sensitization technology utilizes the path cover property in which there is a source signal and a destination signal. By proving path cover property, the signal at the source of the path is tainted. The app formally verifies if it is possible to cover a tainted signal at the destination. When the property is covered, a waveform displays how data can propagate from source to destination. The property can also be determined to be unreachable, which means that it is not possible for data to propagate from source to destination. Verification can also be tuned by the user by creating black box modules where data can enter or not. This will simplify the process of verification to scale well [17].

4 SECURITY VERIFICATION FOCUSING ON SOFTWARE LEVELS OF A SYSTEM

The second class of projects that our survey deals with focuses on verifying security properties of software, while considering the ISA or a machine model of the hardware. Here, we investigate how the security of software is verified in the literature with a hardware model, e.g., some memory model, register files, and other components of the hardware that constitute the environment on which the code will run. Software security verification work that does not consider any hardware in the verification process is outside the scope of this survey. For software-only security verification, we refer the reader to the following surveys [8, 35, 71, 79, 90].

The surveyed projects fall in two categories. First, verification with respect to ISA is where the verification process involves generating assembly code that is considered correct and embodies the program with desired security properties. Typically, assembly code has one-to-one correspondence to the ISA thus the verification process ties the software to the hardware ISA level. Second, verification with respect to a machine model is where the verification process involves a model of the target machine, such as the memory, registers, etc. The machine model is typically very simplified, but it considers key hardware features in the verification process nevertheless.

4.1 Verification with respect to ISA

4.1.1 seL4. [48] was the first operating system microkernel that was formally verified for functional correctness. The aim of the seL4 verification effort is to provide a system free of programming errors that introduce vulnerabilities that may cause failures or facilitate attacks. seL4 is a software-only work and assumes that the underlying hardware, the compiler, and the low-level device driver code are provided free of errors. It uses capability-based security model [54] for access control to enable formal reasoning about object accessibility. seL4's implementation is formally proven correct against its specification, has been proved to enforce strong security properties.

The verification process is shown in Figure 16. In Figure 16, the system prototype is coded in Haskell according to a high-level specification (1). The specification includes a detailed functional and behavioral description of the system (2). Isabelle/HOL theorem prover generates an Executable Specification out of the Haskell code (3). This process is critical since it will directly impact the correctness of the system, any misrepresentations can render the verification ineffective. This specification contains all implementation details and data structures that the low-level implementation must have. The last layer is the actual C implementation of seL4 (4). These three layers used in the formal verification are: abstract specification, executable specification, and C implementation, (5) in Figure 16. The total effort for SeL4 was 11 person years with 14k lines in Haskell/C and 33k lines in Isabelle. The total size of the proof is 200k including generated proofs.




Fig. 16. The seL4 verification process.

4.1.2 CertiKOS and Deep Specifications. [37] presents a design technique based on modern computer system architectures (such as OSes) where each system consists of abstraction levels such as kernels, hypervisors, device drivers, network protocols. Each hides the implementation through a definition of an interface. Deep Specifications is based on the verification of abstraction layers that define interfaces to other layers hiding the implementation details. In mCertiKOS, each layer represents an abstraction, and its behavior is defined in a specification as shown in Figure 17. These specifications are called deep specifications and any two implementations that have the same deep specification must have contextually equal behavior regardless of the implementation method. Hence, mCertiKOS relies on deep specification of layers rather than their specific implementations and as long as an implementation of an abstraction layer can be proven to be equivalent to its deep specification, it can be used without violating the general correctness of the system. An error-free and functionally-correct implementation of the whole system relies on implementing the abstraction layers correctly. Unlike seL4 [48], where the whole system is verified at once, mCertiKOS can be verified layer by layer or as a whole.




Fig. 17. CertiKOS verification process.

mCertiKOS uses two core languages for high-level and assembly-level code to describe the behavior of the system: ClightX and LAsm. ClightX is based on CompCert Clight language [13], a formally verified optimizing compiler for a large subset of the C99 programming language (known as Clight), and LAsm is an assembly language customized for CertiKOS development. These languages can be used to implement abstraction layers. The layer interfaces and Deep Specifications are described using Coq. mCertiKOS uses the CompCertX compiler for both languages. CompCertX is a specialized version of CompCert compiler that works with the mCertiKOS memory and machine

model. If implementations M1 and M2 implement the same DeepSpec, they should have contextually equivalent behavior. The whole CertiKOS took 11.5 person months to finish.

4.1.3 Verve and Ironclad Apps. Verve [93] is an operating system that is verified to guarantee memory and type safety. Verve's architecture consists of two levels. The first level is called the "nucleus" that implements the core functionality needed to access memory and hardware. On top of nucleus, there is a kernel level which supports functionality such as preemptive threads. The applications run on top of these two levels.

Verve uses two strategies to verify the nucleus and kernel. The nucleus is written in Boogie programming language and verified by Boogie. The code of nucleus is manually annotated with assertions that include preconditions, post-conditions and loop invariants. Some of the code is written in assembly and the assembly instructions from the nucleus code are also converted into Boogie so that they can have the annotations. The kernel ensures type safety using Typed Assembly Language (TAL) [65] and a TAL-checker [19]. The kernel is written in safe C# and the code is then compiled to TAL by a special compiler. TAL-checker is used to verify that the assembly code does not violate the primitive abstractions of the language.

Verve is still an experimental OS which lacks some modern features such as exception handling and multiprocessor support, and it assumes the hardware is trusted. However, it supports type safety in the whole OS including the applications. It demonstrates that using automated techniques, high level code (such as safe C#) can be verified for type safety in assembly level using type-safe assembly languages (such as TAL). The specification and proof cost 5494 lines of Boogie code, while the system implementation uses 1377 instructions, resulting a 4 \times annotation ratio.

Ironclad Apps [42] focuses on the execution of remote applications in a secure and a functionally-verified manner. Ironclad uses Verve as the operating system. The verification process covers the code that is executed remotely, the remote OS, libraries, and drivers. Therefore, Ironclad Apps can be regarded as a multi-level verification system which assumes that the hardware is secure. However, the BIOS, and peripheral devices can be malicious. Ironclad Apps eliminates data leaks and software based vulnerabilities. However, it is not designed for hardware-based attacks (side-channels, etc.) nor denial-of-service attacks.




Fig. 18. Ironclad verification process.

The verification process requires an implementation in a high-level language and a high-level specification of the application code which are written in Dafny [52]. The spec and the code are handled in parallel. The code is compiled to output assembly code in the BoogieX86 assembly language (note that the verifier Boogie and the assembly language BoogieX86 are different) [93], (5), (6), and (7) in Figure 18. Meanwhile, the high-level spec is translated into a low-level spec by a spec-translator tool, (1), (2), and (3) in Figure 18. The low level spec and assembly code are then verified together to see if they are functionally equivalent and free of software vulnerabilities. If

the verifier fails, the process has to be restarted with fixed code and spec. When verification passes successfully, the assembler and linker tools convert the assembly code into machine code, (8) in Figure 18.

4.1.4 Komodo. Even though hardware based security mechanisms provide powerful solutions, they are somehow slower to adapt new changes and to provide defenses for new security threats. For example, Intel SGX has not been improved to provide defenses against “controlled-channel” attacks that leaks information using the ability of the OS to observe page faults in enclaves. Komodo [29] provides a different approach to Intel SGX-like architectures by moving management structure of enclaves to a privileged software monitor. Komodo is the first formally-verified, software-based implementation of an SGX-like enclave isolation mechanism [29]. Its design decouples enclave hardware primitives from security-critical but formally verified software, enabling independent evolution of the two. It employs noninterference to prove high-level guarantees of confidentiality and integrity.

The specification of Komodo including its monitor code is then formally proved that it protects the confidentiality and integrity of enclave code and data from the other software (including OS and hypervisor) running on the same machine. The proof establishes that enclave state and out-of-enclave state does not interfere with each other. As in SGX, Komodo does not prove that user code inside enclave cannot leak information.

The implementation uses the Vale programming language [14], which consists of assembly language instructions together with annotations, such as preconditions, postconditions, and loop invariants, that describe the behavior of the instructions. The Vale generates an abstract-syntax-tree (AST) representation of the instructions and proof about the behavior of the instructions in Dafny Language [52]. Dafny uses Z3 to verify the proofs generated by Vale. A trusted assembly printer turns the instruction ASTs into GNU assembly format. A prototype of Komodo has been implemented in ARM TrustZone, since it is capable of providing its basic hardware requirements. The hardware specification covers a subset of the ARMv7 architecture.

4.2 Verification with respect to a Machine Model

4.2.1 SecVisor. SecVisor [34, 76] is a hypervisor designed to provide execution and code integrity. It guarantees that code can execute in kernel mode only if the code is approved by user, and the code can only be modified by SecVisor. SecVisor leverages hardware memory protections and kernel privilege level to achieve execution and code integrity. The design assumes that the attacker has control of everything except the CPU, the memory controller, and the system memory. The small codebase makes the formal verification of SecVisor possible.

A model in $\text{Mur}\varphi$ is developed to verify the system. The model consists of three parts: a hardware model, a SecVisor model, and an attacker model. Since the security of SecVisor is based on the hardware memory protections, it is crucial to specify the hardware model and the page table in the SecVisor model correctly. The hardware model includes physical memory, CPU mode bits, program counter, and a Device Exclusion Vector (DEV) that controls DMA permissions. The initialization, CPU mode transitions, and page-table synchronization in SecVisor are modeled in $\text{Mur}\varphi$. To deal with the state space explosion problem, the authors simplify the model conservatively to avoid false negative. So when $\text{Mur}\varphi$ returns with success, the SecVisor is proved to satisfy all the security properties.

To model the attacker, an actual model where the attacker behavior is modeled and an ideal model without the attacker model, akin to the models used in the verification of XOM [58]. In the actual model, the attacker can write to any memory pages with the permission bits set and can update page tables. The execution integrity is the equality between the actual model and ideal

model where the attacker behavior is not modeled. The code integrity means that the attacker cannot modify the approved code. The execution integrity and code integrity invariants are checked by the $\text{Mur}\varphi$ model checker. The whole model cost 500 lines in $\text{Mur}\varphi$, and takes 343.97 s to finish the model checking for models with 4-page table entries.

4.2.2 MinVisor. MinVisor [23] is a simple hypervisor, which protects its own memory from malicious guests. This work was presented as a follow-up work on SecVisor, but using theorem proving approach. The goal of the project is to fully verify the MinVisor at the assembly level using ACL2. A series of detailed and accurate models of the AMD64 instruction set architecture (ISA), including the memory model, registers, and state transitions, were developed. Several theorems, such as the one where isolation of model specific registers and MinVisor memory are guaranteed against guest modifications, are proved to show the security properties of MinVisor.

4.2.3 AAMP7G. The AAMP7G microprocessor [91] provides “Intrinsic partitioning”, where each partition has exclusive time slices of CPU execution, and exclusive memory space. The time and space partitioning is achieved by its “separation kernel” in microcode. To verify separation kernel, a formal security specification abstractly describing the separation kernel, and a microcode-level functional design model closely corresponding to the implementation are build in ACL2. The entire AAMP7G model is about 3000 lines of definitions. National Security Agency evaluation team conduct a code-to-spec review to validate the microcode-level model. It is then proved that an abstract model enforces the security specification, and the microcode-level corresponds to the abstract model. The strict partition is formally verified in ACL2. Furthermore, a formal model of the instruction set is build, which can be used for analysis of user programs.

4.2.4 Verification of Noninterference at ISA Level. Fox [33], who improved upon work of Myreen, et al. [66], presented a framework for decompilation of machine (assembly) code into statements that can be processed by the HOL4 interactive theorem prover. One of the contributions of [33] is design of a domain specific language, L3, to describe the properties of an ISA. L3 can be converted to statements which can be processed by HOL4. Another of their contributions is definition of numerous instruction behavior of ISAs in L3. Later, Schwarz, et al. [75], derive noninterference properties of ARM and MIPS ISAs using the ISA definitions from [33]. Their framework determines automatically which system components (e.g. program counter or status registers) are accessible at given privilege level, based on the ISA definition. Noninterference is proved by checking how different components (e.g. status registers used by a given instruction) affect state or any return value of an instruction. For the verification, user has to manually label certain components as “low”, such as program counter is low. Then the tools check all possible instructions from the ISA to determine which components can affect the “low” component, and these components are themselves re-labeled as “low.” At the end, the tools output which components should be considered as “low”, given the initial specification.

4.2.5 XMHF. XMHF [89] is an extensible and modular hypervisor. The focus of verification is to preserve the fundamental hypervisor security property of memory integrity (i.e., ensuring that the hypervisor’s memory is not modified by software running at a lower privilege level). To verify the memory integrity, security invariants are inserted into the C code as assertions. However, the full functional correctness is not verified. 5208 lines of the C code is verified automatically by CBMC model checker [49], while the remaining 422 lines of C and 388 lines of assembly are manually audited.

4.3 Tools Automatically Converting Software to System Models in verification tools

Tools are also developed to convert the system implementation and automatically insert assertions for verification. Many architectures provide security features like isolated memory region, e.g. ARM TrustZone, Intel SGX, and AMD memory encryption. In the following, there are two examples that verify the security of application with the security feature provided by hardware.

4.3.1 MOAT. MOAT [81] proposed to find vulnerabilities in enclave user programs that run on Intel SGX architecture. To protect sensitive data and code from disclosure or modification by infrastructure attackers (e.g. malicious OS) or other malicious programs, Intel developed Software Guard Extensions(SGX) [2]. Intel SGX makes such protection possible by providing an isolated memory region called *enclave*. The hardware primitives provided by SGX enforce that only the code inside the enclave can access data within the enclave. However, it cannot protect an enclave user program from leaking sensitive information from within if the software running in the enclave is not programmed properly, thus the need for verification.




Fig. 19. MOAT and SIR verification.

The input of MOAT is the x86+SGX assembly code of an enclave user program, (1) in Figure 19, alongside with annotations indicating the location of secret data. The usage of assembly code as input to the verification process eliminates the need for a trusted compiler. MOAT then translates the assembly code to BAP (Binary Analysis Platform) assembly, which is a simple, RISC-like instruction set [15], as shown by (2) in Figure 19. MOAT uses BAP assembly for precise modeling of x86 and SGX instructions in Boogie verifier.

Inside MOAT, BAP assembly and the secret annotations are converted to code with assertions that Boogie can process. Two kinds of adversaries are considered: active adversaries who can write to any locations in non-enclave memory and passive adversaries who can read any location in non-enclave memory. To model the adversary, MOAT introduced a *havocing adversary*, “who symbolically modifies all the non-enclave memory after every instruction of the enclave code, and is able to observe all non-enclave memory.” To show the effect of the adversary, a havoc instruction (`havoc mem=epc`) is added before every BAP instruction, (2) in Figure 19. To reason about confidentiality, ghost variables (C_x) are also added. If C_x is true then the data x in registers or memory is dependent on a secret. Based on the value of the C_x , one can judge whether there is secret data leaking to non-enclave memory. E.g. line 4 of (2) in Figure 19, asserts data in `%eax` can be written to `mem[%esp]` only if `%esp` does not depend on any secret (no control flow), and if `%esp` is in non-enclave memory ($\neg enc(epb)$) then `%eax` must not depend on a secret. This way, assumptions and assertions about the ghost variable are added, see (2) in Figure 19. The system’s security assertions are verified by Boogie verifier (3). If the assertions do not always hold, then

there is violation to confidentiality, and the verifier returns the violating piece of code, otherwise the design passes; see (4) in Figure 19.

MOAT provides a methodology to prove the security properties of software developed for the Intel SGX architecture. Several applications such as One-time password (OTP) service, query processing over encrypted database are verified as an example. The query processing enclave code, consisting 575 instructions, needs 9 policy annotations and takes 55 sec to proof. It is also the first work to create formal a model of Intel’s new SGX instructions.

4.3.2 Verification of user programs in SIR. Another similar work [80] considers the applications in containers that provide isolation, referred to as Secure Isolated Regions (SIR), such as SGX. This approach decomposes an application to user code (U) that implement the functionality of the application and a small runtime library (L) that provides a narrow interface between U and the untrusted platform outside SIR. The focus of the work is to prove the confidentiality of the U running in SIR by verifying that U satisfies the “WCFI-RW”, which is restriction properties on reads and writes.

This work first uses compiler to generate machine code of U with runtime checks to guarantee WCFI-RW. As compiler is not trusted, it can further optimize the runtime checks for runtime performance. Then the assembly generated from compiler is taken as input to the verification. Since the verifier doesn’t differentiate if it’s generated from the original user code or belongs to the runtime checks, compiler can be untrusted. As (1) in Figure 19, to model the x86 and SGX assembly code, BAP assembly [15] is used. So the U is treated as a set of procedures. There is also a havoc adversary introduced which controls the host OS, hypervisor, network, storage and other datacenter nodes. (2), A static verifier generates proof obligations for each instruction in the procedure by inserting assertions, and (3) these static assertions will be discharged automatically by an SMT solver, Boogie verifier. Three large MapReduce examples are verified and evaluated. The overhead of the runtime checks is 15% on average, and the static verification takes less than 20s.

Different from MOAT where annotations from programmer are needed for fine-grained information flow tracking in the application memory, this work requires U to perform communication with outside SIR only through a narrowed constrained interface provided by L, and everything in U’s memory is considered confidential. This is thus a modular and scalable approach.

5 ANALYSIS OF EXISTING WORK

Tables 3 and 4 present a summary of the main projects reviewed in Sections 3 and 4. In these tables, we compare the existing works in terms of their verification methods, the levels they consider in a system, and the security aspects being verified.

5.1 Verification Methods

Most of the projects use general purpose verification tools, as shown in the *Tool* column of Tables 3 and 4. The current general-purpose tools used in security verification are not compatible with conventional hardware or software languages, such as C or Verilog, and verification is performed as an additional step after design and implementation. Security specifications are described as formulas in theorem provers like Coq, or as invariants in model checkers like Mur φ , as shown in the *Specification* column of Tables 3 and 4. A model (system representation) separate from the system implementation (actual system) is built, e.g. Micro-Policies [24], Cache verification [97], XOM [59], SecVisor [34]. Designers have to make sure their model accurately mirrors the system implementation, otherwise the result of verification might be not correct.

Some projects take the approach of designing new domain-specific languages that allow making verification an integral part of the design and implementation process. In these projects, tools

are developed to transform the system description in the new domain-specific language into another form that is amenable to use with verification tools, e.g. VeriCoq, Dafny, or TAL-compiler. For example, in Dafny, the code has annotations for pre- and post-conditions, invariants, and ghost variables. With use of annotations and through automatic transformation SMT solvers can check if the invariants always hold. Meanwhile, other projects embed security-related tags into a conventional language, and facilitate describing the security-properties to be verified. These projects tend to develop custom tools, as shown in Tables 3 and 4 to make sure the generated design has the desired security properties, such as Sapper [55], Caisson [56], and SecVerilog [96].

Designers should decide which approach to take for their development cycle. On the one hand, they can develop their system in a “traditional” language. This allows for quick development of the functional design with tools familiar to engineers, but does incur the effort of having to also separately write their design in a representation that their preferred security verifier understands. On the other hand, they can implement their system in a verification-friendly language. This has higher initial effort, but may pay off in long term with less effort due to not having to write the representation second time for verification. The drawback is that the verification-friendly language may not support all the aspects the designer desires to verify.

5.2 Verification Aspects

Confidentiality and integrity are the two main security properties often sought in a system. The verification aspect of a system often covers these properties, but can be formulated in a more generic form (e.g. non-interference) or a more specific form (e.g. memory integrity). The formulation of these properties depends on the levels that the system spans, and on the tools used. The analysis of information flow provides a useful basis for proving these security properties of a system. Monitoring information flow requires data labeling, declassification, and information flow rules specific to the system. We observe that many hardware projects use the analysis of information flow for proving information flow policies, non-interference, and confidentiality and integrity, as seen in Table 3. Software projects, as illustrated in Table 4, have a wider variety of verification aspects, which try to verify confidentiality and/or integrity, but only within the selected levels. Designers generally try to provide partial integrity or confidentiality for a system. For example, SecVisor [34] verifies execution and code integrity which is a subsection of the whole memory.

Designers should decide which security aspects to prove in their design, that sit well with both implementation and verification of the system at the same time.

5.3 Verification Levels

The Trusted Computing Base (TCB) often encompasses multiple levels of the system from hardware to software. However, as can be seen in Tables 3 and 4, verification projects are typically focused on the hardware levels, or focused on the software levels. Bringing the hardware and software levels together is difficult, however, needed. For example, enhancing the security of software levels by using support in hardware levels is becoming a more viable approach, especially for remote computing. Hardware based TCBs are emerging quite rapidly, such as ARM TrustZone [88], Intel SGX [63], and AMD memory encryption[1]. The working of this hardware with software needs to be verified for security, and requires spanning many system levels.

Designers should consider expanding their approaches to include more levels into security verification, to allow truly full-system verification.

6 CONCLUSION

Formal verification research has been mostly focused on the functional correctness of the hardware or software systems. Security verification of software-only is also well studied. Hardware security

verification, however, is an emerging research area which is necessitated by the fact that modern systems require both software and hardware for their correct and secure operation. Especially with introduction of security-focused hardware, such as Intel SGX. Trusting remote software and hardware is more critical now than before, as it handles users' ever-increasing sensitive information. Any vulnerabilities in these computing systems can be exploited by attackers. Thus, the whole system, including both the hardware and software parts, should be considered in the security verification.

Security verification is a branch of formal verification where the correctness properties are extended to include security properties, e.g. confidentiality and integrity. The process requires a formal specification of the security properties, an accurate representation of the implementation, and some verification mechanisms, e.g. theorem proving and model checking, to prove that the implementation complies with the needed security properties.

In this survey, we focused on the security verification projects that involve at least some hardware and software levels. Since security properties are provided by multiple levels in the system, only verifying some particular level or levels cannot guarantee whole system's security. With the improvement of verification tools and methods, as presented in this survey, there is a trend to include more and more system levels in verification, but not yet all levels. We provide an insight into the tools and mechanisms used for security verification, and compare projects based on security verification of hardware and software levels they consider.

There are many open research topics in the security verification of hardware and software systems. The most critical, however, is the need for full-system security verification, which spans more levels than can be done through today's existing projects.

Table 3. Summary of projects that focus on hardware verification. These projects were detailed in Section 3.

Name	System representation	Tool	Custom Tool	Levels					Verification Method	Verification Aspect		
				Physical	Gate	RTL	ISIA	Arch	OS	App		
Micro-Policies [24]	Coq Language	Coq IDE	none				✓				Theorem Prover	Non-interference, sealing, etc.
Cache Verification [97, 98]	Murφ Language	Murφ	none					✓			Model Checking	Confidentiality, Integrity
XOM [59]	Murφ Language	Murφ	none					✓	✓		Model Checking	Confidentiality, Integrity
VeriCoq [12]	Verilog	Coq IDE	VeriCoq				✓	✓			Theorem Prover	Information Flow, etc.
Formal-HDL [38]	VHDL	Coq IDE	VHDL converter				✓	✓			Theorem Prover	
RTLIFT [73]	Verilog	none	RTLIFT Tool					✓			Model Checking	Information Flow
Caisson [56]	Cassion Language	none	Cassion Tool				✓	✓			Pen-and-paper Proof	Non-interference
Sapper [55]	Sapper Language	none	Sapper Tool				✓	✓			Pen-and-paper Proof	Non-interference
SecVerilog [96]	SecVerilog Language	none	SecVerilog Tool				✓	✓			Pen-and-paper Proof	Non-interference
ReWire [72]	ReWire	Haskell	ReWire compiler				✓	✓			Theorem Prover	Non-interference, etc.

Table 4. Summary of projects that focus on software verification with respect to the ISA or machine model. These projects were detailed in Section 4.

Name	System representation	Tool	Levels				Verification Method	Verification Aspect	Verification Effort
			Hypervisor	ISA	OS	Arch			
SeL4 [48]	C, Haskell	Isabelle/HOL		✓			Theorem Prover	Functional Correctness, Capability-based Security	200k lines of Isabelle to verify 8700 lines of C code in 22 person-years
CetrikOS [37]	Clight, LAsm	Coq IDE		✓	✓	✓	Theorem Prover	Functional Correctness, Non-interference	11.5 person-months
Verve [93]	TAL, C#, Boogie Lang.	Boogie/Z3 verifier, TAL checker		✓		✓	SMT solver	Type & Memory Safety	5.5k lines of Boogie in 9 person-months
Ironclad Apps [42]	Dafny	Boogie/Z3 verifier, Custom Compiler	✓			✓	SMT solver	Functional Correctness, Memory Safety	36k lines spec and proof in 3 person-years
Komodo [29]	Vale, Dafny	Boogie/Z3 verifier, Custom Translator			✓	✓	SMT solver	Non-interference	23K spec and proof in 2 person-years
XMHF [89]	C, assertions	CBMC			✓	✓	Model Checking	Memory Integrity	5208 lines of C code
SecVisor [34]	Murφ Lang.	Murφ			✓		✓	Model Checking	Execution and code Integrity
MinVisor [23]	ACL2 Lang.	ACL2			✓	✓	✓	Theorem Prover	Code Integrity
AAMP7G [91]	ACL2 Lang.	ACL2			✓	✓	✓	Theorem Prover	Non-interference
ISA [75]	L3	HOL4			✓	✓	Theorem Prover	Non-interference	N/A
MOAT [81]	Assembly	Boogie/Z3 verifier, BAP	✓			✓	SMT solver	Confidentiality	a few policy annotations
Verification of SIR [80]	Assembly	Boogie/Z3 verifier, BAP	✓			✓	SMT solver	Confidentiality	less than 20s

REFERENCES

- [1] AMD. 2016. AMD Memory Encryption. http://amd-dev.wpengine.netdna-cdn.com/wordpress/media/2013/12/AMD_Memory_Encryption_Whitepaper_v7-Public.pdf, accessed May 2016.
- [2] Ittai Anat, Shay Gueron, Simon Johnson, and Vincent Scarlata. 2013. Innovative technology for CPU based attestation and sealing. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy*.
- [3] Armaithi Ardeshiricham, Wei Hu, Joshua Marxen, and Ryan Kastner. 2017. Register transfer level information flow tracking for provably secure hardware design. In *2017 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 1691–1696.
- [4] Arthur Azevedo de Amorim, Nathan Collins, André DeHon, Delphine Demange, Cătălin Hrițcu, David Pichardie, Benjamin C Pierce, Randy Pollack, and Andrew Tolmach. 2014. A verified information-flow architecture. In *ACM SIGPLAN Notices*, Vol. 49. ACM, 165–178.
- [5] Jonathan Bachrach, Huy Vo, Brian Richards, Yunsup Lee, Andrew Waterman, Rimas Avižienis, John Wawrzynek, and Krste Asanović. 2012. Chisel: constructing hardware in a scala embedded language. In *Proceedings of the 49th Annual Design Automation Conference*. ACM, 1216–1225.
- [6] Christel Baier, Joost-Pieter Katoen, et al. 2008. *Principles of model checking*. MIT press Cambridge.
- [7] Roberto Baldoni, Emilio Coppa, Daniele Cono D'elia, Camil Demetrescu, and Irene Finocchi. 2018. A survey of symbolic execution techniques. *ACM Computing Surveys (CSUR)* 51, 3 (2018), 1–39.
- [8] Béatrice Bérard, Michel Bidot, Alain Finkel, François Laroussinie, Antoine Petit, Laure Petrucci, and Philippe Schnoebelen. 2013. *Systems and software verification: model-checking techniques and tools*. Springer Science & Business Media.
- [9] Josh Berdine and Nikolaj Bjørner. 2014. Computing all implied equalities via SMT-based partition refinement. In *Automated Reasoning*. Springer, 168–183.
- [10] Yves Bertot and Pierre Castéran. 2013. *Interactive theorem proving and program development: Coq'Art: the calculus of inductive constructions*. Springer Science & Business Media.
- [11] M. M. Bidmeshki and Y. Makris. 2015. Toward automatic proof generation for information flow policies in third-party hardware IP. In *International Symposium on Hardware Oriented Security and Trust (HOST)*. 163–168.
- [12] Mohammad-Mahdi Bidmeshki and Yiorgos Makris. 2015. VeriCoq: A Verilog-to-Coq converter for proof-carrying hardware automation. In *International Symposium on Circuits and Systems (ISCAS)*. IEEE, 29–32.
- [13] Sandrine Blazy and Xavier Leroy. 2009. Mechanized semantics for the Clight subset of the C language. *Journal of Automated Reasoning* 43, 3 (2009), 263–288.
- [14] Barry Bond, Chris Hawblitzel, Manos Kapritsos, K Rustan M Leino, Jacob R Lorch, Bryan Parno, Ashay Rane, Srinath Setty, and Laure Thompson. 2017. Vale: Verifying high-performance cryptographic assembly code. In *26th USENIX Security Symposium (USENIX Security 17)*. 917–934.
- [15] David Brumley, Ivan Jager, Thanassis Avgerinos, and Edward J. Schwartz. 2011. *BAP: A Binary Analysis Platform*. Springer Berlin Heidelberg, Berlin, Heidelberg, 463–469.
- [16] Cristian Cadar, Daniel Dunbar, Dawson R Engler, et al. 2008. Klee: unassisted and automatic generation of high-coverage tests for complex systems programs.. In *OSDI*, Vol. 8. 209–224.
- [17] CADENCE. 2016. JasperGold Security Path Verification App. http://www.cadence.com/products/fv/jaspergold_security/pages/default.aspx.
- [18] David Champagne and Ruby B Lee. 2010. Scalable architectural support for trusted software. In *International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 1–12.
- [19] Juan Chen, Chris Hawblitzel, Frances Perry, Mike Emmi, Jeremy Condit, Derrick Coetzee, and Polyvios Pratikaki. 2008. Type-preserving compilation for large-scale optimizing object-oriented compilers. In *ACM SIGPLAN Notices*, Vol. 43. ACM, 183–192.
- [20] Edmund Clarke, Orna Grumberg, and D Long. 1993. Verification tools for finite-state concurrent systems. In *A decade of concurrency reflections and perspectives*. Springer, 124–175.
- [21] Thomas M Cover and Joy A Thomas. 2012. *Elements of information theory*. John Wiley & Sons.
- [22] Karl Crary. 2003. *Toward a foundational typed assembly language*. ACM.
- [23] Mike Dahlin, Ryan Johnson, Robert Bellarmine Krug, Michael McCoyd, and William Young. 2011. Toward the verification of a simple hypervisor. *arXiv preprint arXiv:1110.4672* (2011).
- [24] Arthur Azevedo De Amorim, Maxime Dénes, Nick Giannarakis, Catalin Hritcu, Benjamin C Pierce, Antal Spector-Zabusky, and Andrew Tolmach. 2015. Micro-policies: Formally verified, tag-based security monitors. In *IEEE Symposium on Security and Privacy (SP)*. IEEE, 813–830.
- [25] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An efficient SMT solver. In *Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340.

- [26] Udit Dhawan, Nikos Vasilakis, Raphael Rubin, Silviu Chiricescu, Jonathan M Smith, Thomas F Knight Jr, Benjamin C Pierce, and André DeHon. 2014. Pump: a programmable unit for metadata processing. In *Proceedings of the Third Workshop on Hardware and Architectural Support for Security and Privacy*.
- [27] David L Dill. 1996. The Mur ϕ verification system. In *Computer Aided Verification*. Springer, 390–393.
- [28] Jeffrey S Dwoskin and Ruby B Lee. 2007. Hardware-rooted trust for secure key management and transient trust. In *Proceedings of the 14th ACM conference on Computer and communications security*. ACM, 389–400.
- [29] Andrew Ferraiuolo, Andrew Baumann, Chris Hawblitzel, and Bryan Parno. 2017. Komodo: Using verification to disentangle secure-enclave hardware from software. In *Proceedings of the 26th Symposium on Operating Systems Principles*. ACM, 287–305.
- [30] Andrew Ferraiuolo, Weizhe Hua, Andrew C Myers, and G Edward Suh. 2017. Secure information flow verification with mutable dependent types. In *Design Automation Conference (DAC), 2017 54th ACM/EDAC/IEEE*. IEEE, 1–6.
- [31] Andrew Ferraiuolo, Rui Xu, Danfeng Zhang, Andrew C. Myers, and G. Edward Suh. 2017. *Lightweight Verification of Secure Hardware Isolation Through Static Information Flow Analysis*. Technical Report. Technical report, Cornell University.
- [32] Andrew Ferraiuolo, Rui Xu, Danfeng Zhang, Andrew C Myers, and G Edward Suh. 2017. Verification of a Practical Hardware Security Architecture Through Static Information Flow Analysis. In *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 555–568.
- [33] Anthony Fox. 2015. Improved tool support for machine-code decompilation in HOL4. In *International Conference on Interactive Theorem Proving*. Springer, 187–202.
- [34] Jason Franklin, Arvind Seshadri, Ning Qu, Sagar Chaki, and Anupam Datta. 2008. *Attacking, repairing, and verifying SecVisor: A retrospective on the security of a hypervisor*. Technical Report. Technical Report CMU-CyLab-08-008, Carnegie Mellon University.
- [35] Iván García-Ferreira, Carlos Laorden, Igor Santos, and Pablo García Bringas. 2014. A survey on static analysis and model checking. In *International Joint Conference SOCO*. 443.
- [36] Mentor Graphics. 2016. Mentor Graphics Questa Secure Check. <https://www.mentor.com/products/fv/questa-secure-check>
- [37] Ronghui Gu, Jérémie Koenig, Tahina Ramananandro, Zhong Shao, Xiongnan Newman Wu, Shu-Chun Weng, Haozhong Zhang, and Yu Guo. 2015. Deep specifications and certified abstraction layers. In *ACM SIGPLAN Notices*, Vol. 50. ACM, 595–608.
- [38] Xiaolong Guo, Raj Gautam Dutta, Prabhat Mishra, and Yier Jin. 2016. Automatic RTL-to-Formal Code Converter for IP Security Formal Verification. In *Microprocessor and SOC Test and Verification (MTV), 2016 17th International Workshop on*. IEEE, 35–38.
- [39] Robert Harper, Furio Honsell, and Gordon Plotkin. 1993. A framework for defining logics. *Journal of the ACM (JACM)* 40, 1 (1993), 143–184.
- [40] Uwe Hatnik and Sven Altmann. 2004. Using ModelSim, Matlab/Simulink and NS for simulation of distributed systems. In *International Conference on Parallel Computing in Electrical Engineering*. IEEE, 114–119.
- [41] Klaus Havelund and Thomas Pressburger. 2000. Model checking java programs using java pathfinder. *International Journal on Software Tools for Technology Transfer* 2, 4 (2000), 366–381.
- [42] Chris Hawblitzel, Jon Howell, Jacob R Lorch, Arjun Narayan, Bryan Parno, Danfeng Zhang, and Brian Zill. 2014. Ironclad Apps: End-to-end security via automated full-system verification. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 165–181.
- [43] Gerard J Holzmann. 1997. The model checker SPIN. *IEEE Transactions on software engineering* 23, 5 (1997), 279.
- [44] Seongwook Jin, Jeongseob Ahn, Sanghoon Cha, and Jaehyuk Huh. 2011. Architectural support for secure virtualization under a vulnerable hypervisor. In *Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, 272–283.
- [45] Yier Jin and Yiorgos Makris. 2013. A proof-carrying based framework for trusted microprocessor IP. In *Computer-Aided Design (ICCAD), 2013 IEEE/ACM International Conference on*. IEEE, 824–829.
- [46] Matt Kaufmann and J Strother Moore. 2008. An ACL2 tutorial. In *Theorem Proving in Higher Order Logics*. Springer, 17–21.
- [47] Taeho Kgil, Laura Falk, and Trevor Mudge. 2005. ChipLock: support for secure microarchitectures. *ACM SIGARCH Computer Architecture News* 33, 1 (2005), 134–143.
- [48] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammadika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, et al. 2009. seL4: Formal verification of an OS kernel. In *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*. ACM, 207–220.
- [49] Daniel Kroening and Michael Tautschnig. 2014. CBMC-C bounded model checker. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 389–391.
- [50] Chris Lattner. 2008. LLVM and Clang: Next generation compiler technology. In *The BSD conference*, Vol. 5.

- [51] Ruby B Lee, Peter Kwan, John P McGregor, Jeffrey Dwoskin, and Zhenghong Wang. 2005. Architecture for protecting critical secrets in microprocessors. In *Computer Architecture, 2005. ISCA'05. Proceedings. 32nd International Symposium on*. IEEE, 2–13.
- [52] K Rustan M Leino. 2010. Dafny: An automatic program verifier for functional correctness. In *Logic for Programming, Artificial Intelligence, and Reasoning*. Springer, 348–370.
- [53] Xavier Leroy. 2012. The CompCert C verified compiler. *Documentation and user's manual. INRIA Paris-Rocquencourt* (2012).
- [54] Henry M Levy. 2014. *Capability-based computer systems*. Digital Press.
- [55] Xun Li, Vineeth Kashyap, Jason K Oberg, Mohit Tiwari, Vasanth Ram Rajarathinam, Ryan Kastner, Timothy Sherwood, Ben Hardekopf, and Frederic T Chong. 2014. Sapper: A language for hardware-level security policy enforcement. In *ACM SIGARCH Computer Architecture News*, Vol. 42. ACM, 97–112.
- [56] Xun Li, Mohit Tiwari, Jason K Oberg, Vineeth Kashyap, Frederic T Chong, Timothy Sherwood, and Ben Hardekopf. 2011. Caisson: a hardware description language for secure information flow. In *ACM SIGPLAN Notices*, Vol. 46. ACM, 109–120.
- [57] Sheng Liang. 1998. Modular monadic semantics and compilation. (1998).
- [58] David Lie, John Mitchell, Chandramohan A Thekkath, and Mark Horowitz. 2003. Specifying and verifying hardware for tamper-resistant software. In *Symposium on Security and Privacy*. IEEE, 166–177.
- [59] David Lie, Chandramohan Thekkath, Mark Mitchell, Patrick Lincoln, Dan Boneh, John Mitchell, and Mark Horowitz. 2000. Architectural support for copy and tamper resistant software. *ACM SIGPLAN Notices* 35, 11 (2000), 168–177.
- [60] Roger Lippett, Carl F Schaefer, and Cary Ussery. 2012. *VHDL: Hardware description and design*. Springer Science & Business Media.
- [61] Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B Lee. 2015. Last-level cache side-channel attacks are practical. In *IEEE Symposium on Security and Privacy*. 605–622.
- [62] Derek Lockhart and Christopher Batten. 2014. Hardware Generation Languages as a Foundation for Credible, Reproducible, and Productive Research Methodologies. In *Workshop on Reproducible Research Methodologies (REPRODUCE)*.
- [63] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R Savagaonkar. 2013. Innovative instructions and software model for isolated execution. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy*. ACM.
- [64] Kenneth L McMillan. 1993. The SMV system. In *Symbolic Model Checking*. Springer, 61–85.
- [65] Greg Morrisett, David Walker, Karl Crary, and Neal Glew. 1999. From System F to typed assembly language. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 21, 3 (1999), 527–568.
- [66] Magnus O Myreen and Michael JC Gordon. 2007. Hoare logic for realistically modelled machine code. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 568–582.
- [67] Nicholas Nethercote and Julian Seward. 2007. Valgrind: a framework for heavyweight dynamic binary instrumentation. *ACM Sigplan notices* 42, 6 (2007), 89–100.
- [68] Rishiyur Nikhil. 2004. Bluespec System Verilog: efficient, correct RTL from high level specifications. In *Formal Methods and Models for Co-Design, 2004. MEMOCODE'04. Proceedings. Second ACM and IEEE International Conference on*. IEEE, 69–70.
- [69] Tobias Nipkow, Lawrence C Paulson, and Markus Wenzel. 2002. *Isabelle/HOL: a proof assistant for higher-order logic*. Vol. 2283. Springer Science & Business Media.
- [70] Sam Owre, John M Rushby, and Natarajan Shankar. 1992. PVS: A prototype verification system. In *Automated Deduction—CADE-11*. Springer, 748–752.
- [71] Christine Paulin-Mohring. 2011. Introduction to the Coq proof-assistant for practical software verification. In *Tools for Practical Software Verification*. Springer, 45–95.
- [72] Adam Procter, William L Harrison, Ian Graves, Michela Becchi, and Gerard Allwein. 2015. Semantics driven hardware design, implementation, and verification with ReWire. In *ACM SIGPLAN Notices*, Vol. 50. ACM, 13.
- [73] Andrei Sabelfeld and Andrew C Myers. 2003. Language-based information-flow security. *Selected Areas in Communications, IEEE Journal on* 21, 1 (2003), 5–19.
- [74] Bruce Schneier. 2016. The Internet of Things Will Turn Large-Scale Hacks into Real World Disaster. <https://motherboard.vice.com/read/the-internet-of-things-will-cause-the-first-ever-large-scale-internet-disaster>, accessed July 25, 2016..
- [75] Oliver Schwarz and Mads Dam. 2016. Automatic Derivation of Platform Noninterference Properties. In *International Conference on Software Engineering and Formal Methods*. Springer, 27–44.
- [76] Arvind Seshadri, Mark Luk, Ning Qu, and Adrian Perrig. 2007. SecVisor: A tiny hypervisor to provide lifetime kernel code integrity for commodity OSes. *ACM SIGOPS Operating Systems Review* 41, 6 (2007), 335–350.
- [77] Ofer Shacham, Megan Wachs, Andrew Danowitz, Sameh Galal, John Brunhaver, Wajahat Qadeer, Sabarish Sankaranarayanan, Artem Vassiliev, Stephen Richardson, and Mark Horowitz. 2012. Avoiding game over: Bringing design to

the next level. In *Proceedings of the 49th Annual Design Automation Conference*. ACM, 623–629.

[78] Yan Shoshitaishvili, Ruoyu Wang, Christopher Salls, Nick Stephens, Mario Polino, Andrew Dutcher, John Grosen, Siji Feng, Christophe Hauser, Christopher Kruegel, et al. 2016. Sok: (state of) the art of war: Offensive techniques in binary analysis. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 138–157.

[79] Vijay D Silva, Daniel Kroening, and Georg Weissenbacher. 2008. A survey of automated techniques for formal software verification. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on* 27, 7 (2008), 1165–1178.

[80] Rohit Sinha, Manuel Costa, Akash Lal, Nuno P Lopes, Sriram Rajamani, Sanjit A Seshia, and Kapil Vaswani. 2016. A design and verification methodology for secure isolated regions. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, 665–681.

[81] Rohit Sinha, Sriram Rajamani, Sanjit Seshia, and Kapil Vaswani. 2015. Moat: Verifying Confidentiality of Enclave Programs. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15)*. 1169–1184.

[82] Pramod Subramanyan, Rohit Sinha, Ilya Lebedev, Srinivas Devadas, and Sanjit Seshia. [n. d.]. A Formal Foundation for Secure Remote Execution of Enclaves. ([n. d.]).

[83] G Edward Suh, Dwayne Clarke, Blaise Gassend, Marten Van Dijk, and Srinivas Devadas. 2003. AEGIS: architecture for tamper-evident and tamper-resistant processing. In *Proceedings of the 17th annual international conference on Supercomputing*. ACM, 160–171.

[84] Jakub Szefer and Ruby B Lee. 2012. Architectural support for hypervisor-secure virtualization. In *ACM SIGPLAN Notices*, Vol. 47. ACM, 437–450.

[85] Donald Thomas and Philip Moorby. 2008. *The Verilog® Hardware Description Language*. Springer Science & Business Media.

[86] Ken Thompson. 1984. Reflections on trusting trust. *Commun. ACM* 27, 8 (1984), 761–763.

[87] Mohit Tiwari, Hassan M.G. Wassel, Bita Mazloom, Shashidhar Mysore, Frederic T. Chong, and Timothy Sherwood. 2009. Complete Information Flow Tracking from the Gates Up. In *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XIV)*. 109–120.

[88] ARM Trustzone. 2016. *TrustZone Information Page*. Technical Report. <http://www.arm.com/products/processors/technologies/trustzone/>

[89] Amit Vasudevan, Jonathan M Mccune, James Newsome, and Cylab Carnegie Mellon. 2012. Design and Implementation of an eXtensible and Modular Hypervisor Framework. (2012).

[90] Dolores R Wallace and Roger U Fujii. 1989. Software verification and validation: an overview. *Ieee Software* 6, 3 (1989), 10.

[91] Matthew M Wilding, David A Greve, Raymond J Richards, and David S Hardin. 2010. Formal verification of partition management for the AAMP7G microprocessor. In *Design and Verification of Microprocessor Systems for High-Assurance Applications*. Springer, 175–191.

[92] Jonathan Woodruff, Robert NM Watson, David Chisnall, Simon W Moore, Jonathan Anderson, Brooks Davis, Ben Laurie, Peter G Neumann, Robert Norton, and Michael Roe. 2014. The CHERI capability model: Revisiting RISC in an age of risk. In *Proceeding of the 41st annual international symposium on Computer architecuture*. IEEE Press, 457–468.

[93] Jean Yang and Chris Hawblitzel. 2010. Safe to the last instruction: automated verification of a type-safe operating system. In *ACM Sigplan Notices*, Vol. 45. ACM, 99–110.

[94] Jean Yang and Chris Hawblitzel. 2011. Safe to the last instruction: automated verification of a type-safe operating system. *Commun. ACM* 54, 12 (2011), 123–131.

[95] Danfeng Zhang, Aslan Askarov, and Andrew C Myers. 2012. Language-based control and mitigation of timing channels. *ACM SIGPLAN Notices* 47, 6 (2012), 99–110.

[96] Danfeng Zhang, Yao Wang, G Edward Suh, and Andrew C Myers. 2015. A hardware design language for timing-sensitive information-flow security. In *ACM SIGARCH Computer Architecture News*, Vol. 43. ACM, 503–516.

[97] Tianwei Zhang and Ruby B Lee. 2014. New models of cache architectures characterizing information leakage from cache side channels. In *Proceedings of the 30th Annual Computer Security Applications Conference*. ACM, 96–105.

[98] Tianwei Zhang and Ruby B Lee. 2014. *Secure cache modeling for measuring side-channel leakage*. Technical Report. Tech. Report, <http://palms.ee.princeton.edu/node/>.

Received December 2021