

Hardware Information Flow Tracking

WEI HU, Northwestern Polytechnical University, China

ARMAITI ARDESHIRICHAM, University of California, San Diego, USA

RYAN KASTNER, University of California, San Diego, USA

Information flow tracking (IFT) is a fundamental computer security technique used to understand how information moves through a computing system. Hardware IFT techniques specifically target security vulnerabilities related to the design, verification, testing, manufacturing, and deployment of hardware circuits. Hardware IFT can detect unintentional design flaws, malicious circuit modifications, timing side channels, access control violations, and other insecure hardware behaviors. This article surveys the area of hardware IFT. We start with a discussion on the basics of IFT, whose foundations were introduced by Denning in the 1970s. Building upon this, we develop a taxonomy for hardware IFT. We use this to classify and differentiate hardware IFT tools and techniques. Finally, we discuss the challenges yet to be resolved. The survey shows that hardware IFT provides a powerful technique for identifying hardware security vulnerabilities as well as verifying and enforcing hardware security properties.

CCS Concepts: • **Security and privacy** → **Logic and verification**; **Information flow control**; **Formal security models**.

Additional Key Words and Phrases: Hardware security, information flow security, information flow tracking, security verification, formal method, survey

ACM Reference Format:

Wei Hu, Armaiti Ardeshiricham, and Ryan Kastner. 2020. Hardware Information Flow Tracking. *ACM Comput. Surv.* 0, 0, Article 00 (April 2020), 38 pages. <https://doi.org/10.0000/0000000.0000000>

1 INTRODUCTION

A core tenet of computer security is to maintain the confidentiality and integrity of the information being computed upon. Confidentiality ensures that information is only disclosed to authorized entities. Integrity maintains the accuracy and consistency of information – verifying that it is correct, complete, and not modified in any unauthorized manner. Understanding how information flows throughout a computing system is crucial to determine if that system operates securely with respect to confidentiality and integrity.

Information flow tracking (IFT) is a security technique that models how information propagates as a system computes. It labels data objects with a tag to denote security classes, which are assigned different meanings depending on the type of security property under analysis. IFT updates the tags as the data is computed upon, and verifies information flow properties by observing the state of the tags.

Information flow tracking can reason about the confidentiality and integrity properties. IFT verifies confidentiality by determining if secret information can ever be disclosed to an unclassified location. An example confidentiality property states that information related to a secret key should not leak outside a restricted memory space. IFT can also verify data integrity properties by not allowing untrusted entities to operate on trusted information. An example integrity property states that a user level program can never modify the system configuration data.

Authors' addresses: Wei Hu, Northwestern Polytechnical University, 127 Youyi Xi Lu, Xi'an, Shaanxi, 710072, China, weihu@nwpu.edu.cn; Armaiti Ardeshiricham, University of California, San Diego, 9500 Gilman Dr., La Jolla, CA, 92093, USA, aardeshi@ucsd.edu; Ryan Kastner, University of California, San Diego, 9500 Gilman Dr., La Jolla, CA, 92093, USA, kastner@ucsd.edu.

2020. Manuscript submitted to ACM

IFT has been used to enforce security properties across the computing system stack including operating systems [35, 78, 160], programming languages [126], distributed systems [104, 161], and cloud computing [6]. It can detect and prevent a wide range of software security issues including buffer overflow, memory corruptions, SQL injection, formatted string attack, cross-site scripting attacks, and malware [1, 36, 101, 108, 137]. IFT has been applied to verify the functional correctness and security of complex software systems like the *SeL4* secure OS kernel [74, 103]. More recently, hardware IFT techniques have been used to detect and assess security threats related to hardware design vulnerabilities [65, 135], insecure debug modes [15, 60, 123], timing channels [2, 4, 29, 64, 162], and hardware Trojans [44, 53, 107, 152]. A survey on security verification of hardware/software systems cited information flow analysis as the most commonly used technique for modeling and reasoning about security [26].

This article introduces a hardware IFT taxonomy, which is used to survey, classify, and compare hardware IFT techniques. We start by discussing the initial information flow model put forth by Denning [30, 31]. We identify the distinguishing characteristics of Denning’s information flow model including the *security property* enforced, *level of abstraction*, *operator precision*, and *verification technique*. We use this to create a taxonomy to help compare and contrast the hardware IFT techniques. Figure 1 shows a taxonomy which consists of the following elements:




Fig. 1. A taxonomy for hardware information flow tracking. We start by describing the characteristics of an information flow model and define basic terminology. After that, we expand on different ways to classify the IFT techniques. The precision of class combining operator as well as the type of security property under consideration play an important role in defined the characteristics of the IFT technique. Another distinguishing characteristic is the abstraction level of the technique’s information flow model. Finally, the IFT techniques enable different forms of verification, which provides another method of differentiation.

- **Operator Precision:** How does the class combining operator update security classes?
 - **Conservative:** IFT operator uses a least upper bound.
 - **Precise:** IFT operator considers the effect of data values.
 - **Tradeoffs:** IFT operator performs tradeoffs between precision and computational complexity.
- **Security Property:** What types of security properties are supported?
 - **Confidentiality:** IFT technique prevents leakage of sensitive information.
 - **Integrity:** IFT technique prohibits overwriting of trusted data by an untrusted entities.
 - **Isolation:** IFT technique prevents communication between two entities of different trust.

- **Timing channel:** IFT technique captures information flows through a timing channel.
- **Hardware Trojan:** IFT technique detects malicious design modifications.
- **Level of Abstraction:** What is the level of abstraction of the information flow model?
 - **System:** IFT technique considers system level flows.
 - **Algorithmic:** IFT technique is deployed during high-level synthesis.
 - **Architecture:** IFT technique is deployed at the instruction set architecture (ISA) level.
 - **RTL:** IFT technique targets register transfer level (RTL) design.
 - **Gate:** IFT technique considers a gate-level netlist.
 - **Circuit:** IFT technique targets analog and mixed-signal hardware designs.
- **Verification Technique:** Which verification techniques are supported?
 - **Simulation:** IFT technique uses simulation tools to determine information flows.
 - **Formal verification:** IFT technique employs formal methods to verify security properties.
 - **Emulation:** IFT technique allows for hardware emulation of information flow behaviors.
 - **Virtual prototyping:** IFT technique creates a software version of hardware to measure information flows.
 - **Runtime:** IFT technique can dynamically track information flows during runtime.

We use this taxonomy to classify and describe different hardware IFT techniques. This helps us discuss the relative advantages, weaknesses, and drawbacks of these techniques. We attempt to draw a clear picture of the past and future of hardware IFT research with a focus on its capability for developing new hardware security tools.

The remainder of this article is organized as follows. In Section 2, we discuss the core concepts of information flow, covering the fundamental elements in Denning’s information flow model [31], the classical noninterference model [42], the different types of flow relations as well as basics of covert and side channels. Section 3 formalizes the concept of class combining operator precision and discusses IFT precision and computational complexity tradeoffs. In Section 4, we describe the different types of security properties that can be modeled and enforced under the information flow framework. Sections 5 reviews existing hardware IFT techniques from the perspective of the level of abstraction where they are deployed while Section 6 elaborates on the various mechanisms for performing information flow security verification. We discuss potential research vectors and challenges in Section 7 and conclude the article in Section 8.

2 INFORMATION FLOW

Information flow tracking (or information flow control) verifies that no unauthorized flow of information is possible. IFT techniques work by labeling storage objects with a security class, tracking those labels as their data is computed upon, and enforcing information flow policies to understand, verify, and/or control the movement of information in the system. Simple examples of policies include: 1) confidential information should never leak to an unclassified location and 2) untrusted information should never overwrite a trusted memory component.

Denning developed the earliest models for *information flow* [30, 31] that enables one to define the allowable flows of information between storage objects \mathcal{N} using security classes \mathcal{SC} and flow relations \rightarrow . Processes \mathcal{P} operate on those storage objects. And a class combining operator \oplus provides a relationship between security classes. We use Denning’s information flow model to discuss and compare different hardware IFT techniques.

An information flow model \mathcal{FM} is defined by $\mathcal{FM} = \langle \mathcal{N}, \mathcal{P}, \mathcal{SC}, \oplus, \rightarrow \rangle$ where:

- $\mathcal{N} = \{a, b, \dots\}$ is a set of *storage objects*. Elements of \mathcal{N} can vary depending on the level of abstraction under consideration, e.g., files, segments, registers, and flip-flops.

- $\mathcal{P} = \{p, q, \dots\}$ is a set of *processes* where information flows through, e.g., functions, arithmetic operations and Boolean gates.
- $SC = \{A, B, \dots\}$ is a set of *security classes* corresponding the security classification of data objects, e.g., secret and unclassified. The security class associated with the information stored in an object a can either be bound statically or updated dynamically with the propagation of information. These are also commonly called *labels* or *tags*.
- The *class combining operator* \oplus takes a pair of security classes and calculates the resulting security class, i.e., $SC \times SC \mapsto SC$.
- A *flow relation* \rightarrow is defined on a pair of security classes. It is a “can flow” relation on SC , i.e., $\rightarrow \subseteq SC \times SC$. Information flows from class A to class B whenever information associated with A could affect the value of data object associated with B . In other words, there exists a situation where changing the value of object associated with A leads to a change in the value of object associated with B . We write $A \rightarrow B$ if and only if information in class A is allowed to flow to class B .

We attempt to adopt this classic notation as much as possible throughout this article. Our goal is to classify the work in hardware IFT using this notation so that it can provide a common language to compare and contrast the different techniques and methodologies. In the following subsections, we elaborate on the elements of this information flow model. Section 2.1 describes *storage objects* and *processes* from a hardware viewpoint. Section 2.2 discusses relevant information related to the *flow relations* and *security classes*. Section 2.3 introduces the idea of *noninterference*. Section 2.4 categorizes and illustrates the different types of *information flow relations*. Section 2.5 covers the basics of *covert* and *side channels*.

2.1 Storage Objects and Processes

An *object* is a container of information and a *process* describes computations performed upon those objects. Storage objects and processes vary depending on the security policies and level of abstraction. The initial work by Denning [30, 31] largely considered the security of software processes or programs. Much of it focused on security at the operating system level, and thus objects were things like files, directories, and users. Hardware storage objects are constructs defined in hardware description languages (HDLs) and their corresponding logical representations and physical implementations.

In hardware design, storage objects are defined using HDL types that can take or store values, e.g., input and output ports, internal wires and signals, registers, flip-flops, and memory blocks. It is often convenient to consider objects as any stateful elements in the hardware, i.e., to ignore wires and signals and focus on registers, flip-flops, and memory blocks. These objects are operated upon by processes \mathcal{P} that can be defined as register transfer operations, finite state machine (FSM) actions, combinational functions, and other common hardware models of computation.

Unfortunately, “process” is an overloaded term since *process* is a defined keyword in VHDL which denotes a set of sequential statements. The security notion of a process (\mathcal{P}) covers a much broader scope; it refers to any arithmetic, logical or analog hardware component that operates on input data objects and produces a resulting object. Examples of hardware process \mathcal{P} include arithmetic, logical, and circuit-level operations.

2.2 Security Classes and Flow Relations

Information flow policies define allowable relationships between data objects. To do this, IFT associates every object with a security class, i.e., gives objects a *security label* or *security tag*. Objects define the functional state of the system; their security labels determine the system security state.

The flow relation operator \rightarrow defines an allowable flow of information on two security classes. $SC = \{A, B\}$, $A \rightarrow B$ states that information from security class A is allowed to flow to security class B . We can also denote the lack of an information flow with the operator \nrightarrow , e.g., $B \nrightarrow A$ states that information in security class B should never flow to security class A .

Security practitioners often use a *lattice* to describe flow relations. For example, Denning defined information flow policies as a finite lattice [31]. A lattice is in the form of $\mathcal{L} = \{\mathcal{E}, \sqsubseteq\}$, where \mathcal{E} is the set of elements and \sqsubseteq is a partial order relation defined on \mathcal{E} . The tuple $\{\mathcal{E}, \sqsubseteq\}$ constitutes a *lattice* if there exists a *least upper bound* element and a *greatest lower bound* element for any $A, B \in \mathcal{E}$ in the element set. The class combining operator \oplus is also a least upper bound operator in that for all $A, B, C \in \mathcal{E}$:

$$\begin{aligned} A \rightarrow A \oplus B \text{ and } B \rightarrow A \oplus B \\ A \rightarrow C \text{ and } B \rightarrow C \implies A \oplus B \rightarrow C \end{aligned} \quad (1)$$

The greatest lower bound operator \odot is defined as $A \odot B = \oplus LSet(A, B)$, where $LSet(A, B) = \{C \mid C \rightarrow A \text{ and } C \rightarrow B\}$ is the set of security classes from which information can flow to both A and B .

We usually consider lattices with finite number of elements. Let $\mathcal{E} = \{A_1, A_2, \dots, A_n\}$. The least upper bound (also called *maximum element*, denoted as *high*) and the greatest lower bound (also called *minimum element*, denoted as *low*) on the lattice are defined as follows:

$$\begin{aligned} \text{high} &= A_1 \oplus A_2 \oplus \dots \oplus A_n \\ \text{low} &= A_1 \odot A_2 \odot \dots \odot A_n \end{aligned} \quad (2)$$

An information flow policy/property can be modeled using a *security lattice* $\mathcal{L} = \{SC, \sqsubseteq\}$, where SC is a set of security classes and \sqsubseteq is the partial order defined on SC . Given any two security classes $A, B \in SC$, we say A is lower than B (or B is higher than A) when $A \sqsubseteq B$.




Fig. 2. Example security lattice structures. The lattice defines allowable information flows. Any operations that lead to an unallowable flow (e.g., allowing Untrusted information to flow to a Trusted data object or Confidential information to flow to an Unclassified data object) result in a security violation. Information flow is allowed to flow upward along the arrow in the security lattice; downward flows of information are not permitted.

Figure 2 provides four examples of lattices commonly used in modeling security. Figure 2 (a) shows the high and low security lattice that we mentioned earlier. Here, the security class set has two elements $SC = \{\text{high}, \text{low}\}$. The partial order (reflected by the arrow) defines allowable flows of information among different security classes, e.g., $\text{low} \sqsubseteq \text{high}$ indicating that low information is allowed to flow to a high data object.

Figure 2 (a) to (c) are linear security lattices, where any two security classes are ordered or comparable. There are also more complex lattice structures with partial ordering. Figure 2 (d) gives an example of such non-linear security lattice, which contains non-comparable security classes, e.g., security classes A and B as well as A and BC. Information flow between security classes without partial order defined by the lattice should be prohibited.

In practice, most policies use a two-element lattice like in Fig. 2 (a). This allows the modeling of security properties related to confidentiality and integrity as $\text{low} \sqsubseteq \text{high}$. In cases where multi-level security needs to be enforced, often times we employ conservative label operators (e.g., the least upper bound or greatest lower bound operators) to determine the security class that the output of a process should take. More precise security class updating rules expand the two-level technique to account for the influence of the data values on security class operation [54, 55].

2.3 Noninterference

Noninterference is a strict multi-level security model proposed by Goguen and Meseguer [42]. It creates a model of information flow by modeling all system inputs, outputs, and state as either high or low. It states that any changes in high inputs shall not result in changes in low outputs. That is, low objects can learn nothing about high data. Another way of stating this is that a computer system produces the same low outputs for any set of low inputs regardless of the high inputs. That is, it is impossible to learn any information about the high values by changing the low inputs. Equivalently, the computer system responds exactly the same to any input sequence of low values regardless of the high values.

Noninterference provides a notion of strong, end-to-end security properties [128]. That is, it tracks the release and subsequent transmissions of the data to ensure that they adhere to the desired security properties. It extends the notion of access control [32, 80, 154], which strictly limits accesses to objects in the system. These lists enable restrictions on who can access an object; yet once it is released, its propagation is not restricted, and this information can be subsequently transmitted in an insecure manner.

Definitions of noninterference often rely on semantic models of program execution. Early work in this space includes Cohen’s work using the notion of strong dependency [22, 23] and McLean’s definition of noninterference using trace semantics [100]. Volpano *et al.* formalize Denning’s lattice approach as a type system viewable as a form of noninterference [148]. Many follow-on works extend the expressiveness of the model and extend the typing system to handle concurrency, nondeterminism, synchronization and exceptions; for more information, see Sabelfield and Myers’ excellent survey on language-based information flow security [126].

Much of this early work on noninterference was related to accessing user data and isolating programs running on shared computers. Thus, it largely dealt with files, resources, and users as objects and took a very “software” or operating system centric view. Regardless, the general notion of noninterference translates well to hardware. Perhaps the most straightforward example is the SecVerilog project [162], which employs the notion of noninterference using semantic models from Verilog. SecVerilog adds security labels to the Verilog language and extends previous security typing languages to hardware design. The challenge is accounting for all the eccentricities of the hardware models of computation, and in particular of vast use of concurrency. Regardless, the general idea that low objects should not be able to infer any information about the high objects still holds and can be translated to hardware.

A major challenge is that the number of objects in hardware is typically much larger than in software. Hardware objects could be comprised of the entire memory space, registers, and other low-level state – any stateful storage element could be considered. Of particular note is that objects can vary by the level of abstraction. For example, objects at the operating systems are files and users. ISA-level objects are visible programmer state, e.g., registers, memory, and

some program state (PC, status and configuration registers). Objects at the RTL would start to include intermediate registered values and state related to control flow. Objects at the gate level would likely denote every flip-flop which would include all registers (programmer visible or not) and the state of different controllers (i.e., finite state machine bits). Enforcing noninterference across all of these objects can be challenging. Yet, regardless of the number of objects, the notion of noninterference remains true – any information related to high objects should not flow to low objects.

2.4 Types of Information Flow Relations

While flow relations provide a way to specify security policies related to whether or not information is allowed to be transferred between objects, they do not strictly define the method to determine if there is a flow of information. There are many different ways in which information can be transmitted. Understanding the different types of information flow relations and how they manifest themselves in hardware is an important first step for precisely measuring all flows of information and further tightly enforcing the security policies.

Information can flow through functional (e.g., Boolean gates, arithmetic operators, transactions) and physical channels (e.g., power, electromagnetic, and thermal). A major difference between functional and physical flows lies in how information is encoded and transmitted. Functional flows encode information by using specified behavior of the hardware. The simplest example is directly writing data from one object to another. Physical information flows use some notion related to the physical process of performing the computation. This could be the time that an operation takes where examples include the time to receive data through a cache or compute some operation whose execution time depends on the data involved in the computation (e.g., modular exponentiation [77]). Additionally, any circuit switching activity results in power consumption, electromagnetic radiation, and thermal emission. Each of these can be used to infer what is being computed, and thus creates a flow of information.

In this article, we primarily focus on functional information flows, which traditionally have been categorized as *explicit* and *implicit*. In the following, we illustrate the difference between explicit and implicit flows. We also discuss a special type of implicit flow called *timing flow*, which usually leads to leakage of sensitive information.

2.4.1 Explicit Flow. *Explicit flow* is associated with the direct movement of data. Thus, it is also called *data flow*. As shown in Fig. 3, information flows explicitly from the source operand(s) to the destination operand (from A and B to Sum in this example) when an expression $sum := A + B$ is evaluated.



Fig. 3. A simple example of explicit flow. The data from A and B explicitly flow to Sum when the add operation is executed.

2.4.2 Implicit Flow. *Implicit flow* is a more subtle type of information flow that is caused by context dependent execution such as conditional operations. Here, information can flow between two objects even when there is no direct data assignment. Consider the example shown in Fig. 4 (a), which implements a simple if/else control structure.

It is clear that both A and B flow explicitly to Mux through the statements. The more challenging question is whether there is an information flow between Sel and Mux . There is obviously no direct explicit flow. However, it is possible to infer the value of Sel depending on what is written into Mux , i.e., whether Mux is assigned the value from A or B . Thus, there is an information flow, or more precisely an implicit information flow between Sel and Mux . Information




Fig. 4. A simple example of implicit flow. (a) The if/else hardware control structure causes a flow of information from *Sel* to *Mux* even though there is no direct assignment of the *Sel* data to *Mux*. As the design is synthesized to RTL (b) and gate level (c), the implicit flow becomes explicit as the multiplexer in (b) and the gates in (c) now denote an explicit flow between *Sel* and *Mux*.

flows implicitly from the conditional variable *Sel* to *Mux*, where how *Mux* will be updated depends on the value of *Sel*. It is important to understand that an implicit flow can occur even if there is no explicit data assignment.

These implicit flows start to become more like explicit flows as the level of abstraction becomes lower. For example, in Fig. 4 (b) and (c), the if/else statement from Fig. 4 (a) is synthesized into a multiplexer, which is a three input (*A*, *B*, and *Sel*) one output (*Mux*) function block. As synthesis progresses, the multiplexer is further decomposed into gates which again implements the same function. Thus, the implicit flow from Fig. 4 (a) becomes an explicit flow after RTL (Fig. 4 (b)) and logic synthesis (Fig. 4 (c)). As RTL conditional branch statements are synthesized to multiplexers and even smaller primitive Boolean gates, the information flow from *Sel* to *Mux* appears no different from those from *A* or *B* to *Mux*. In other words, all functional information flows are explicit at the gate level. At the same time, it becomes harder to decouple the different types of flows as the level of abstraction is lowered. For example, assume a security policy needed to differentiate between explicit and implicit flows. This would be more challenging to distinguish at the gate level.

2.4.3 Timing Flow. A *timing flow* is a special type of implicit flow where information is flows through timing-related design behavior. Figure 5 shows a simple example of cache timing flow. The cache signal *hit* input affects the time at which the *valid* signal is asserted. It is not a functional flow because *valid* will eventually be asserted regardless of the value of *hit*. Yet, by observing the time at which *valid* is asserted, an attacker can infer the value of *hit*. Thus, *hit* flows to *valid* through a timing channel.




Fig. 5. A cache timing flow. The signal *hit* flows to the signal *valid* through a timing channel, i.e., an attacker can determine information about *hit* by observing when *valid* is asserted.

A more subtle case is that an information flow measurement can be a mix of functional and timing. Consider a RSA implementation through repeated square-and-multiply used in Kocher’s timing attack [77]. There is functional flow from the key to ciphertext since the value of key directly affects the encryption result. However, there is also a timing flow to *cipher* in this implementation since the individual key bits determine the amount of time needed to encrypt a given message. It is possible to distinguish between timing flows from functional flows using hardware IFT [2, 113].

2.5 Covert and Side Channels

A *covert channel* is a transfer of information using an unintended source. The Orange Book classifies two types of covert channels - storage and timing channels [99]. The *storage covert channel* transmits information to a “storage location” (e.g., disk sectors) from parties in different security levels that are not allowed to communicate. The *timing covert channel* transmits information through the usage of system resources, which are observable by another process.

Lampson defined a covert channel as a channel “not intended for information transfer at all, such as the service program’s effect on system load” [79]. Sabelfeld and Myers followed this notion and characterized covert channels as ways to transfer information within a computing system using mechanisms not intended to function as information channels. They classified covert channels to several categories, including implicit flows, termination channels, timing channels, probabilistic channels, resource exhaustion channels, and power channels [126]. Covert channels can be a security threat depending on what attackers can observe about the computing system, e.g., execution time, power consumption, resource usage and probability distribution of observable data.

By comparison, a side channel is a leakage of information through non-functional (often physical) characteristics, e.g., execution time, power consumption, EM radiation, and acoustic emissions. Kocher pioneered in developing side channel attacks to recover the cryptographic key from runtime and power trace measurements using statistical analysis [76, 77]. *Meltdown* [88], *Spectre* [75], and *Foreshadow* [13, 153] leverage architectural level security flaws to load kernel information into cache and then use timing channel attack to retrieve protected information. These side channels are caused by access to shared resources. There also can be a security threat in shared computing environments such as the cloud server [155].

An interesting consequence of modeling information flow at the hardware level, is that it provides cycle accurate timing information enabling timing channels (or lack thereof) to be modeled and formally verified using IFT [2, 4]. Section 4 describes this in more detail.

Covert and side channels often exploit similar phenomenon, e.g., timing, power consumption, and thermal radiation. A major difference between a covert and side channel is intent. A covert channel often times involves one party attempting to transfer information to another; a side channel is usually unintentional, but it often exploited as a vulnerability.

3 OPERATOR PRECISION

IFT techniques can use different class combining operators (or label propagation policies) to calculate the security class of the output upon input security labels. The type of operator employed can affect the precision of an IFT technique in measuring information flows and further the computational complexity to do the measurement. This section categorizes IFT techniques by operator precision and also discusses precision and complexity tradeoffs.

3.1 Precision of IFT

Precision is an attribute that reflects whether an IFT technique can accurately measure all (and only) existing information flows in hardware designs. Let O be an n -input function denoted as $O = f(I_1, I_2, \dots, I_k, \dots, I_n)$. A precise IFT technique indicates a flow from I_k to O if and only if the value of input I_k ($1 \leq k \leq n$) has an effect on that of O , i.e., a change in the value of I_k will result in a change in O .

$$I_k \rightarrow O \Leftrightarrow f(I_1, I_2, \dots, I_k, \dots, I_n) \neq f(I_1, I_2, \dots, I'_k, \dots, I_n) \quad (3)$$

Precise IFT techniques can accurately model the information flow behaviors of hardware designs while imprecise ones can either indicate non-existent information flows (i.e., false positives) or miss actual information flows (i.e., false negatives). The false positives in information flow measurement correspond to information flow behaviors that would not actually happen. Most imprecise IFT techniques tend to introduce false positives, which are safe but may lead to conservative design decisions, i.e., a perfectly secure design may be verified as insecure. However, it should be noted that false negatives can arise due to incomplete analysis, e.g., when an information flow security property cannot be verified within reasonable amount of time and resource.

3.2 Imprecise IFT

Early hardware IFT techniques all tend to employ the conservative least upper bound operator in Denning's information flow framework to calculate the output security class and determine the flow of information [12, 137]. These techniques assume that information will always propagate from the inputs to outputs driven by them. In other words, information flows are independent of the functionality of the component and input conditions; the output will take a high label as long as any input is high. Let $\mathcal{L}(I_1), \mathcal{L}(I_2), \dots, \mathcal{L}(I_n)$ be the security labels of the inputs I_1, I_2, \dots, I_n respectively and $\mathcal{L}(O)$ be the security label of output O . The conservative label propagation policy can be formalized as Equation (4), where the \oplus symbol represents the least upper bound operator. Imprecise IFT techniques employing this label propagation policy has the computational complexity of $O(2^n)$.

$$\mathcal{L}(O) = \mathcal{L}(I_1) \oplus \mathcal{L}(I_2) \oplus \dots \oplus \mathcal{L}(I_n) \quad (4)$$

The label propagation policy shown in Equation (4) can be conservative for certain hardware components according to Equation (3). Consider the two-to-one multiplexer (MUX-2), whether an input could propagate to the output depends on the value of the select line.

Although imprecise IFT techniques can lead to conservative verification results, they usually allow a quick profile of potential information flow security vulnerabilities. This can be useful for identifying security violations that occur under rare conditions, e.g., a hardware Trojan that leaks information only when triggered [53]. However, the verification performance benefits that come at the cost of loss in precision may later be counteracted by additional efforts needed to exclude false alarms, e.g., a disabled debug port that had access to critical memory locations [52].

3.3 Precise IFT

Precise IFT techniques take into account both the functionality of a hardware component and the input values when determining flow relations. The label propagation policy can be formalized as a function of both the inputs and their labels as shown in Equation (5). Thus, the complexity of precise IFT techniques is $O(2^{2n})$.

$$\mathcal{L}(O) = f(I_1, I_2, \dots, I_n, \mathcal{L}(I_1), \mathcal{L}(I_2), \dots, \mathcal{L}(I_n)) \quad (5)$$

Gate level information flow tracking (GLIFT) [144] is the first well-established precise IFT technique. Figure 6 shows the difference in class combining operation with/without considering data values.

Figure 6 (b) and (c) show some partial flow relation and class combining operations without/with consideration of data values respectively. The entries in red text reveal the differences between these two rule sets. In Fig. 6 (b), the least upper bound operator is used for class combining and thus high information will always flow from B to O . However, data value can sometimes determine the output and prevent the flow of information from the other input. Take the




Fig. 6. Examples of flow relation and class combining operation. (a) Two-input AND gate. (b) Flow relation and class combining operation without considering data values. (c) Flow relation and class combining operation with consideration of data values.

second and third rows in Fig. 6 (c) for example. The (low, 0) input in A will dominate the output to be (low, 0). In such cases, the high information in input B does not have any effect on the output and thus cannot flow to O .

Oberg and Hu *et al.* have formalized precise IFT logic for Boolean gates under the notion of GLIFT [57, 111]. In a successive work, they performed a formal analysis on the complexity of the precise hardware IFT logic generation problem and revealed the impreciseness that can arise when tracking information flows in a constructive manner [58].

Ardehiricham *et al.* proposed RTLIFT [3], a precise hardware IFT technique at the RTL. The major challenges are handling the complex conditional branch structures such as if-else, case and loop statements. However, RTLIFT allows more effective handling of timing flows modeled as conditional register updates controlled by tainted values [2].

Although RTLIFT has observed a magnitude of improvement over GLIFT in terms of verification performance and scalability, precise IFT can inevitably be expensive for large hardware designs since it is inherently an NP-complete problem [58]. Thus, it is desirable to derive practical IFT methods that can balance between the performance benefits and precision to achieve verification goals in acceptable amount of time and resource.

3.4 Precision and Complexity Tradeoffs

Precision and complexity are two contradictory factors for evaluating hardware IFT techniques [52]. We have to tradeoff one for the other like we have done for area and performance in logic synthesis.

Switching circuit theories provide insightful explanations about the impreciseness of hardware IFT techniques. Oberg and Hu *et al.* have identified the root cause of false positives in IFT analysis as static logic hazards [57, 111]. They used a MUX-2 example to demonstrate how false positive will arise when there is either static-0 or static-1 logic hazard.

Hu *et al.* associate the source of impreciseness with internal variable correlations caused by reconvergent fanouts [58]. They also formally proved that precise hardware IFT logic generation is an NP-complete problem. To overcome the exponential scale complexity of the precise IFT logic generation problem, constructive as well as heuristic methods were proposed to create imprecise tracking logic in polynomial time at the cost of loss in precision.

In [52], Hu *et al.* investigated the effect of logic synthesis optimizations on precision of IFT. This work reveals that the false positives of hardware IFT hide behind unreachable satisfiability don't care conditions in digital circuits. It provides an efficient approach to derive different version of hardware IFT logic with variable precision and complexity for Boolean gates. The observation is that precise IFT logic more accurately measures information flows because it considers the value of variables in label propagation. We would move towards the less precise by gradually ignoring the inputs to relax precision. Figure 7 illustrate this process using the two-input AND (AND-2) gate as an example, where A , B and O are the inputs and output of AND-2; A_t , B_t and O_t are their security labels, respectively.




Fig. 7. Deriving simplified hardware IFT logic for AND-2.

At the top level, we have the most precise hardware IFT logic for AND-2. The second level shows three simplified versions of hardware IFT logic derived by setting one of the inputs to don't-care. By eliminating an additional input, we reach the least precise (but most simplified) hardware IFT logic for AND-2 at the bottom.

In [52], both graph and QBF-SAT based search techniques are proposed to gradually introduce false positives into hardware IFT logic to accelerate security verification and pinpoint security vulnerabilities. Becker *et al.* took a step further to quantitatively introduce a certain percentage of false positives into IFT logic with the assistance of QBF solvers, allowing a finer-grained tradeoff between precision and complexity [7].

4 SECURITY PROPERTIES

gInformation flow properties are expressed as flow relations over objects. The property specifies how information stored in hardware objects can or can not flow to other objects. For example, a confidentiality property is specified by stating certain flows cannot occur (e.g., *secret* \rightarrow *unclassified*). An integrity property is also specified by restricting flows (e.g., *untrusted* \rightarrow *trusted*).

gInformation flow tracking properties are examples of *hyperproperties* since they express expected behaviors that require verification across multiple traces [21]. A *trace* is an execution of the hardware, e.g., a set of waveforms denoting the values of the hardware registers over time. An individual trace can be determined to satisfy or violate a *trace property*. Functional hardware verification techniques use trace properties, which are defined over sets of traces, e.g., a register should be set to '0' upon reset. While trace properties can encode functional properties, IFT properties require comparing the hardware's behavior with respect to more than one trace. Hyperproperties are defined on sets of sets of traces (also called systems) and thus more complex than trace properties.

To illustrate the relationship between hyperproperties and hardware IFT, consider an *integrity* property, which indicates that the behavior of trusted objects should remain intact with respect to values of untrusted objects. This policy compares the hardware behavior over multiple traces with differing values for untrusted data. Consequently, the integrity policy cannot be represented as a trace property. However, integrity can be represented as a hyperproperty that considers multiple traces which are identical except for the value of the untrusted data. To preserve integrity, the system is expected to act identically (in terms of the trusted objects) across all the traces assuming they start from equivalent initial states. Using formal methods, this hyperproperty can be verified for all possible values of untrusted data components to provide a formal proof of integrity in all possible executions. gOr one could verify it in simulation or emulation by insuring that the security label never reach an undesired state. This likely will not provide total coverage, but typically scales better than formal verification techniques.

gHardware IFT tools model movement of information through the hardware using labels that store security relevant information. The labels indicate whether objects are trusted or untrusted, confidential or public, contains timing variation or not, etc. The property defines the initial conditions on the labels and a set of conditions for how the labels can or can not change during execution. For an integrity property, IFT uses labels to indicate if objects are trusted or not.

Objects which carry untrusted values are initialized to have a *high* label while other objects have a *low* label by default. To ensure integrity, the property states that the system compute in a manner that maintains the *low* labels for the objects which are assumed to be trusted throughout execution. Maintaining a *low* label in this scenario demonstrates that untrusted values have not influenced the trusted components.

In this section, we show how different security policies can be modeled from the perspective of information flow and provide example hardware IFT properties. We describe the properties as an assertion-based IFT property language [49]. Assertion-based verification is commonly used for functional hardware verification [40], and thus should be familiar with hardware designers. Additionally, commercial IFT hardware tools have adopted assertion-based verification for IFT properties [97, 98]. Assertion-based IFT verification works by initially setting labels in the system (through *assume* statements) and verifying that these labels do or do not propagate to some other object (through *assert* statements). We consider an information flow model where relevant hardware components are extended with a security label, and a two level lattice $low \sqsubseteq high$ shown in Fig. 2 (a). We use the System Verilog Assertion (SVA) language to specify these properties.

4.1 Confidentiality

Confidentiality properties ensure that any information related to sensitive objects (labeled as high) never flows to an unclassified object (using a low label). For instance, in a cryptographic core the secret key should not flow to a point that is publicly observable. Confidentiality properties are written to protect secret assets in a given design. This is done by marking the secret assets or inputs with a high security label and monitoring the label of public ports and storage units. To preserve confidentiality, all public ports and storage units should maintain a low label throughout execution. Examples of confidentiality properties for various hardware designs such as crypto cores and arithmetic units are shown in the first four rows of Table 1. Here we wish to verify that sensitive/secret objects do not affect publicly viewable objects. To do this, we set the sensitive objects labels as high using the *assume* statement, and *assert* that the publicly viewable objects remain low.

Table 1. Examples of hardware IFT properties related to confidentiality, integrity, and isolation.

Benchmark	Synopsis	Formal Representation
SoC Arbiter	(Confidentiality) Acknowledgement signal is not driven from sensitive requests	<code>assume (req[i] == high); assert (ack[j] == low)</code>
Scheduler	(Confidentiality) Grant signal is not driven from sensitive modes	<code>assume (mode[i] == high); assert (grant[j] == low)</code>
Crypto Core	(Confidentiality) Ready signal is not driven from secret inputs	<code>assume (key == high & plain_text == high); assert (ready == low)</code> <code>assume (operand1 == high & operand2 == high);</code>
Floating Point Unit	(Confidentiality) Ready signal is not driven from the inputs	<code>assert (ready == low)</code> <code>assume (user_inp == high);</code>
Crypto Core	(Integrity) Key register is not modified by public inputs	<code>assert (key == low)</code>
Debug Unit	(Integrity) Debug flag is not modified by public inputs	<code>assume (user_inp == high);</code> <code>assert (debug_en == low)</code>
Processor	(Integrity) PC, private memory and control flow conditions are not modified by public inputs	<code>assume (user_inp == high);</code> <code>assert (PC == low & private_mem == low & cond == low)</code>
Access Control	(Integrity) Unauthorized users cannot access protected units	<code>assume (user[i] == high);</code> <code>assert (asset[j] == low)</code>
SoC	(Isolation) Accesses to different cores on an SoC are isolated	<code>assume (req[i] == high); assert (ack[j] == low)</code> <code>assume (req[j] == high); assert (ack[i] == low)</code>
Memory	(Isolation) Memory locations are isolated	<code>assume (mem[i] == high); assert (mem[j] == low)</code> <code>assume (mem[j] == high); assert (mem[i] == low)</code>

4.2 Integrity

Integrity is the dual of confidentiality - here we mark untrusted hardware resources with a high label and verify that they do not affect critical components with low labels. For example, in a processor the program counter (PC) should not be overwritten by data from an unprotected network.

Integrity properties can be specified for any design where certain memory locations, registers, or flags should be protected against unauthorized access. This is modeled by marking public access such as user or network input with a high label and constraining the sensitive variables to maintain a low security label. Table 1 shows four integrity properties written for crypto cores, debug units, processors, and access control units.

4.3 Isolation

Isolation can also be enforced as an information flow security property. Isolation states that there should never be information exchange between two components with different trust levels. For example, in SoC designs, trusted IP cores sitting in the secure world with low labels should be separated from those which are untrusted and are in the insecure domain with high labels. It should be noted that isolation is a two-way property as shown in the examples of Table 1.

4.4 Timing Channel

Information flow models can be used to capture timing side-channels in hardware designs. These properties assess whether sensitive information can be retrieved by measuring the computation time. To precisely capture timing flows, the information flow model needs to distinguish between logical and timing flows. Take as example a floating point division unit expected to run in constant time independent of the value of the operands. In this case, logical flow exists from the data inputs (i.e., the divider and dividend) to the data outputs (i.e., quotient and remainder) since the outputs are computed from the inputs. However, whether or not there is timing flow from the inputs to the outputs (i.e., if the arithmetic unit runs in constant time or not) depends on the implementation of the floating point unit.

Table 2 summarizes the properties used in [2, 4] to verify timing side channel in different hardware designs. Here, we assume an IFT model where the flow relation tracks timing-based information flows (denoted by *time*), e.g. as described in Clepsydra [2]. To verify timing leakage, the assertions are written over these “*time*” labels.

Table 2. Summary of properties used for detecting timing side channels.

Benchmark	Synopsis	Formal Representation
Sequential Divider	Result is ready in constant time	assume (dividend == high & divisor == high); assert (quotient == <i>time</i> low & remainder == <i>time</i> low)
Sequential Multiplier	Result is ready in constant time	assume (dividend == high & divisor == high); assert (quotient == <i>time</i> low & remainder == <i>time</i> low)
Cache	Data is available in constant time	assume (index[i] == high); assert (data[j] == <i>time</i> low)
SoC Arbiter	Requests are granted in constant time	assume (req[i] == high); assert (ack[j] == <i>time</i> low)
Thread Scheduler	Scheduling is done in constant time	assume (thread_active[i] == high); assert (thread_grant[j] == <i>time</i> low)
AES Cipher	Cipher text is ready in constant time	assume (key == high & plain_text == high); assert (cipher_text == <i>time</i> low)
RSA Cipher	Cipher text is ready in constant time	assume (key == high & plain_text == high); assert (cipher_text == <i>time</i> low)

4.5 Hardware Trojan

Information flow tracking can be used to detect certain types of hardware Trojans that leak sensitive information by inserting malicious information channels in the design. For example, Trust-HUB benchmarks [127] include examples of Trojans added to crypto cores using a small circuitry that transfers the secret key to a public output under certain conditions. This class of hardware Trojans can be detected using a property that observes information flow from

sensitive data (e.g., the secret key) to the public ports. Table 3 summarizes the properties for detecting information leakage by gate level information flow tracking in the Trust-HUB [127] benchmarks as used by Hu et. al. [53].

Table 3. Summary of properties used for detecting hardware Trojans

Benchmark	Synopsis	Formal Representation
AES-T100	Key does not flow to Antena	assume (key == high); assert (Antena == low)
AES-T400	Key does not flow to shift register	assume (key == high); assert (TSC-SHIFTReg == low)
AES-T1100	Key does not flow to capacitance	assume (key == high); assert (capacitance == low)
RSA-T200	Key does not flow to count	assume (key == high); assert (count == low)

The Common Weakness Enumeration (CWE) is a community-developed list of software and hardware weakness types [96]. The hardware CWE list started in 2019 and includes a large number of enumerations that can be modeled using IFT properties. This includes most of the CWE categories including CWE-1195: Manufacturing and Life Cycle Management Concerns, CWE-1196: Security Flow Issues, CWE-1197: Integration Issues, CSE-1198: Privilege Separation and Access Control Issues, CWE-1199: General Circuit and Logic Concerns, CWE-1201: Core and Compute Issues, CWE-1202: Memory and Storage Issues, CWE-1205 Security Primitives and Cryptography Issues, CWE-1207: Debug and Test Problems, and CWE-1208: Cross-Cutting Problems. Many of these weaknesses can be modeled as IFT properties discussed in this section. The IFT based security verification tool *Radix* [97, 98] claims support of over 80% of the hardware CWEs.

5 LEVEL OF ABSTRACTION

Hardware designs are translated through many levels of abstraction before implementation. Each abstraction level has different capability and accuracy in describing hardware behaviors, which in turn allow verifying of different types of security properties. For example, timing properties require cycle accurate precision which is only possible at the RTL and below. In the following, we categorize Hardware IFT techniques according to the abstraction level where they are developed and deployed, ranging from system, to RTL, and all the way down to Boolean gates and the circuits [3, 9, 39, 54, 67, 119, 129, 132, 137, 144, 157, 162].

Figure 8 shows some hardware design abstractions and the IFT tools at these different levels of abstraction. A question that would arise naturally is when and where to deploy hardware IFT? It is important to choose the right abstraction for the job [85]. At higher levels (e.g., system and algorithmic), the models abstract away significant amount of design details, which can lead to better verification performance and scalability. In addition, we can detect potential security vulnerabilities earlier in the design process and fix them at much lower cost. However, we usually need to make conservative assumptions about the hardware behaviors due to the lack of implementation details at such an early design phase. The architectural level has enough design information for verifying the security related interference between hardware and software as well as firmware. However, it is a non-trivial task to account for explicit flows and timing channels at this level. As we move to RTL, the model is cycle accurate and can capture timing information flows. However, the complex HDL syntax features (i.e., conditional branch statements) renders automatic IFT logic generation a complicated process. At the gate level, the complex syntax structures are simplified to Boolean gates and all information flows explicitly, significantly lowering the effort for IFT logic formalization. However, it is well-known that gate level simulation and verification is typically one or two orders of magnitude slower than doing the same job




Fig. 8. Hardware IFT techniques at different levels of abstraction.

at the RTL [8]. However, gate level and circuit level IFT techniques are still essential complements to higher level IFT methods. Some designs may come in the form of synthesized or layout netlist or we may want to understand if security properties still hold after logic synthesis and technology mapping.

5.1 System Level

A number of approaches explore hardware extension for dynamic information flow tracking (DIFT) to perform system level verification [20, 48, 72, 82, 83, 117, 130, 149, 151]. Some techniques use loosely-coupled processor or dedicated hardware to perform IFT at runtime while others employ IFT to verify system security during design time.

5.1.1 IFT Through Co-processing Unit. Lee *et al.* use CDI (Core Debug Interface) to fetch meta data [82, 83]. The ARM CoreSight ETM (Event Trace Macrocell) trace component is used to collect information for each CPU instruction. This trace component is later updated to CoreSight PTM (Program Trace Macrocell), which provides information only on instructions that modify the PC. Wahab *et al.* use this new debug feature in ARMv7 processor to retrieve details on instructions committed by the CPU to perform DIFT. A DIFT coprocessor is implemented in the program logic of Xilinx Zynq device [150] for tag processing. *ARMHex* [151] is a DIFT solution leveraging the CoreSight PTM debug component. Static analysis and instrumentation are employed to retrieve the missing information resulting from the update. *ARMHex* also considers the security of the DIFT unit by using ARM Trust-Zone to protect it. A successive work develops a more adaptive information flow monitor called *BFBlare*, which supports multi-thread, virtual tag memory, floating point and Linux kernel features [149].

Heo *et al.* propose an application specific instruction set processor (ASIP) architecture for system level programmable dynamic program analysis [48]. The proposed approach is implemented as a hardware engine called program analysis unit (PAU), whose key components are the tag processing core and the main controller. The main controller communicates the execution traces with the host processor through the system bus. The tag processing core is an application specific instruction set architecture for management of tags.

Shin *et al.* proposed the implicit flow tracking unit (*IFTU*), an external hardware module connected to the host processor through system bus. In *IFTU*, for a conditional branch, the taint tag of data used for branch condition checking is first propagated to the program counter. Then, for the instructions after the conditional branch, the taint tag of the program counter is propagated to the destinations where the values are affected by the branch result. *IFTU* can detect more recent advanced attacks that exploited implicit flows at reasonable area costs [130].

5.1.2 IFT for Verifying and Enforcing SoC Security. Kastner *et al.* use IFT to prevent secret information from leaking (for confidentiality), and untrusted data from being used in the making of critical decisions (for integrity) in SoC systems using untrusted IPs [72]. They use a gate level information flow model but use it to verify system level properties.

Hassan *et al.* and Pieper *et al.* propose an approach for SoC security validation through Virtual Prototypes (VP) modeled using SystemC TLM 2.0 [46, 118]. They employ static information flow analysis to verify the *no flow* security property and detect potential security breaches such as data leakage and untrusted access. However, information flow is determined by static path search between the source and sink points, which can lead to conservative verification results.

Piccolboni *et al.* add support of DIFT to loosely coupled accelerators in heterogeneous SoCs, which are vulnerable to attacks [117]. It is inspired by two earlier works [119, 120] on DIFT targeting such accelerators. A standalone shell circuit is added for interacting with the accelerator, e.g., loading, propagating and storing tags.

5.1.3 IFT for Enforcing Embedded System Security. In [102], a bottom-up approach is proposed for enforcing information flow security in embedded systems. The paper implements bit-tight information flow control from the level of Boolean gates. It then exposes this information flow measurement capability to higher abstraction levels in the system stack ranging from architecture and all the way up to secure application to enforce security properties.

Cherupalli *et al.* propose a fine-grained IFT technique for enforcing information flow security for IoT systems [20]. They introduce application specific GLIFT [144] for IoT systems. The paper builds a gate-level symbolic analysis framework that leverages the knowledge of the target application to efficiently identify the potential information flow security vulnerabilities. It designs a software tool that takes as inputs the gate level netlist of the processor, the original binary of a given application and information flow policies of interest. The tool performs symbolic gate level simulation of the application binary on the netlist to determine if there is any information flow security policy violation and identify the offending instructions.

5.2 Algorithmic Level

Some recent works synthesize DIFT enhanced hardware accelerators from high-level specification [64, 119, 129].

ASSURE [64] is a security-constrained high-level synthesis (HLS) framework that automatically synthesizes verifiably secure hardware accelerators from high-level descriptions under user-defined security policies and constraints. *ASSURE* extends the LegUp HLS framework [17] to perform IFT. It generates an additional enforcement FSM to manage public I/O accesses separately in a way that all secret-conditioned branches are balanced by applying additional latency constraints. Other operations that do not affect timing of the public outputs are controlled by the main FSM. As a result, the public outputs have constant latency.

GARUDA [129] is a high-level security policy language and a compiler to compile security policies described in *GARUDA* to Verilog modules. It enables the modular construction and composition of runtime hardware monitors for securing a variety of security policies, including fault isolation, secure control flow, and DIFT via taint checking. The high-level security policies described in *GARUDA* are first converted into a policy intermediate representation (IR) and then compiled to Verilog code using the extraction feature of the Coq theorem prover.

TaintHLS [119] is an HLS methodology to automatically generate DIFT-enabled accelerators from high-level specification. TaintHLS creates a microarchitecture to support baseline operations and a shadow microarchitecture for intrinsic DIFT support in hardware accelerators. The DIFT-enabled accelerators have negligible performance overhead while adding no more than 30% hardware. These accelerators achieve the same value tags and identical number of false positives/negatives as compared to software IFT implementation.

5.3 Architecture Level

5.3.1 Capability Vector Architecture. Architectural level hardware IFT research dates back to as early as Denning’s information flow theoretic framework. In 1978, Saal and Gat from IBM proposed a *capability vector machine* to enforce access and information flow control [125]. The machine used a capability vector to describe information flow policy. For example, an instruction moving data from *A* to *B* required the capability vectors $\langle(A, \text{READ})$ and $(B, \text{WRITE})\rangle$. An instruction would be allowed to execute when its implied capability vector existed otherwise the movement of information should be prohibited.

Bondi and Branstad designed an architecture for fine-grained secure computing [12]. The architecture’s operands are 2-tuples, consisting of a data object and an associated security tag. A dedicated subprocessor and tag memory are added to perform security tag manipulation in parallel to the data path. The coupled subprocessors enforce multi-level security (MLS) access control and information flow control. While such an architecture provides efficient security support, it is infeasible for modern processors with their large instruction set.

5.3.2 DIFT and Related Architectures. Dynamic information flow tracking (*DIFT*) is an architectural mechanism to protect programs against buffer overflow and format string attacks by identifying spurious information flows from untrusted I/O and restricting the usage of untrusted information [137]. Although DIFT employs conservative tracking rules, the performance overheads could be as high as 23% for certain benchmarks. To reduce the overheads, Kannan *et al.* propose to decouple DIFT with a co-processor [70]. The co-processor receives instruction tuples and loads tag information from tag memory/cache to perform security checks. The two processors synchronize on system calls. When a tag check fails, an exception would be generated and the execution results of the instruction and tag update would be rejected.

Runtime Information Flow Engine (*RIFLE*) enforces user-defined information flow security policies [146]. In RIFLE, program binaries are translated from a conventional ISA to an information flow secure ISA. The translated programs execute on hardware with IFT capability and interact with a security-enhanced operating system, which is responsible for enforcing the security policies. The authors also show that language-based information-flow systems are less secure than previously thought while architectural level IFT approaches can be at least as powerful as language-based information-flow systems.

Raksha is an architecture that combines both hardware and software DIFT mechanisms [24]. Hardware manages security tags at low performance overhead for user code, OS code, and data that crosses multiple processes while software performs security checks and handles exceptions. Raksha supports a flexible and programmable mechanism for specifying security policies, allowing multiple concurrently active security policies to deal with a wide range of attacks. It enables security exceptions that run at the same privilege level and address space as the protected program. This allows the integration of the hardware security mechanisms with software analyses, without incurring the performance overhead of switching to the operating system.

FlexiTaint [147] is another hardware accelerator for DIFT with support of flexible and configurable taint propagation rules. It extends the processor pipeline and implements taint-related logic for tag processing. Tags are stored as a packed array in virtual memory; a taint propagation cache tagged by operation type and input-operand taints to improve performance. Reported results showed that FlexiTaint incurred negligible performance overheads even when simultaneously supporting two different taint propagation policies.

In [27, 28], Deng *et al.* implement DIFT using a soft-core processor. The CPU pipeline is modified in order to allow the processing core to forward its execution trace for tag processing. In [25], executable binary is instrumented to extract the required information for DIFT. However, it causes up to 86% performance overhead.

Shioya *et al.* and Joannou *et al.* tried to reduce the memory overhead of DIFT techniques by exploiting tag memory access characteristics, e.g., non-uniformity, locality, and cache behavior. They propose to use multi-level table, cache structures and memory access optimizations to accelerate tag memory access [69, 131]. Another attempt to optimize taint label manipulation is to separate trusted and untrusted data objects into different pages, which prevents the overhead for storing and loading taint labels [92, 93].

Locality-Aware Taint Checker (*LATCH*) is an architecture for optimizing DIFT [145]. The key observation is that that DIFT exhibit strong temporal locality, with typical applications manipulating sensitive data during limited phases of computation. This property allows *LATCH* to invoke precise, computationally intensive tracking logic only during the execution involves sensitive data otherwise perform lightweight, coarse-grained checks to reduce performance overheads. The authors design and implement three different DIFT systems that incorporate the *LATCH* model, i.e., S-*LATCH* for optimizing software-based DIFT on a single core, P-*LATCH* to optimize DIFT monitoring using a separate core and H-*LATCH* that targets hardware-based DIFT.

DataSafe is a software-hardware architecture that provides dynamic instantiations of secure data compartments (SDCs) [19]. The architecture enforces hardware monitoring of the information flows from the compartment using hardware policy tags associated with the data at runtime to prevent leakage of sensitive information. Security tags are derived from security policies associated with sensitive data and passed to the SDC for tag propagation. These tags would be checked when data were processed by unvetted applications.

5.3.3 Execution Lease Architectures. Tiwari *et al.* describe how a class of secure architectures can be constructed, from the gates up, to completely capture all information flows [141–143]. In [142], hardware IFT is applied to an FPGA device to create a micro-processor that implements bit-tight information flow control. The architecture design is a full implementation that is programmable and precise enough to track all flows of information and prevent untrusted data from flowing to critical memory locations such as the PC. Novel techniques are proposed to handle conditional branches, loops, loads and stores to prevent over tainting. A prototype processor is implemented on Altera Stratix II FPGA, where IFT logic adds about 70% resource overhead [143].

The *execution lease* is another IFT enhanced architecture [87, 141]. This lease architecture allows execution contexts to be tightly quarantined and their side effects to be tightly bounded. Information flow security properties such as isolation between trusted and untrusted execution contexts can be verified all the way down to the gate-level using the precise information flow measurement capability of IFT. However, all the leases have a fixed time bound, which does not support performance optimization micro-architecture features such as caches, pipelining, branch prediction or TLBs due to the timing variation they introduce. In addition, the lease architecture adds substantial design and performance overheads resulting from the additional information flow tracking logic.

To overcome the shortcomings of the lease architecture, in [143] a minimal but configurable architectural skeleton is crafted to operate along with a small piece of software, which together composes the minimal essential functionality through which information flow security properties can be verified all the way down to Boolean gates. Such strict core component is then used to create a hardware-software system that allows unbounded operation, inter-process communication, pipelining, I/O with traditional devices and other architectural level performance optimizations. This work also proposes a more scalable verification technique called star-Logic (*-logic), which represents indeterministic values as the abstract value * in order to verify all possible executions arising from unknown values in a single run.

5.3.4 IFT on Multi-core Architectures. Most architectural level IFT techniques discussed above used a dedicated co-processor for tag processing. There are also works that distribute data processing and metadata processing to different cores in a single processor [18, 63, 91, 105, 110, 115, 124]. DIFT on multi-core processor can be a much more challenging task in that data processing and metadata processing are normally decoupled and in different orders. Maintaining metadata coherence is a fundamental problem for ensuring the correctness of DIFT in a multi-core environment. To address possible inconsistency, Santos *et al.* proposed *METACE* that included architectural enhancement in the memory management unit and leveraged cache coherence hardware protocol to enforce metadata coherence [91]. In addition, there is also performance penalty due to inter-process communication [105].

ShadowReplica [63] is another dynamic data flow tracking technique on multi-core processor. It uses a shadow thread and spare CPU cores to decouple execution and data flow tracking. The two processes communicate through a shared data structure. A major difference from previous approaches lies in that it introduces an off-line application analysis phase that utilizes both static and dynamic analysis methodologies to generate optimized code for decoupling execution and implementing DFT. This additional phase also minimizes the amount of information that is to be communicated between the two threads.

SMT-based IFT (*SIFT*) is a hardware IFT technique on SMT processors [115]. Taint propagation and policy checking are performed by a separate thread executed in spare context of an SMT processor. The instructions for the checking thread are generated in hardware using self-contained off-the-critical path logic at the commit stage of the pipeline. Instruction generation can be completed in one additional cycle at commit time. Experimental results using SPEC CPU 2006 benchmarks showed 4.5% area overhead, 20% overhead in performance and 23% additional power consumption, which is significantly lower than several software IFT implementations.

Another benefit of performing IFT on multi-core processors is the possibility of parallelizing DIFT in order to reduce the high performance overhead of sequential implementations. The major challenge lies in serial dependencies in DIFT, i.e., label propagation usually depends on resulting labels of prior operations. To accelerate DIFT, different optimizations are proposed to resolve and simplify such dependencies, e.g., preserving only operation that affect the taint status of critical points [110] and tracking the information flow only through unary operations [18, 124]. Afterwards, the parallelized DIFT instances are distributed to multiple cores for processing.

5.3.5 IFT on RISC-V Architectures. The *SAFE* project developed one of the earliest verified information flow secure architectures based on the RISC instruction set. A simple stack-and-pointer machine with “hard-wired” dynamic IFC was designed to provide IFT support to higher-level software. The instructions for implementing the IFT mechanism are formalized and verified using Coq [5]. More recent projects such as *SSITH* and *AISS* both aim to develop secure RISC-V architectures, where IFT is employed for verifying and enforcing security guarantees [109, 133].

Programmable Unit for Metadata Processing (*PUMP*) is an enhanced RISC processor architecture with ISA-level extension [34]. *PUMP* changes the pipeline to incorporate programmable tag propagation rule checking. *PUMP* performs

single-cycle common-case computation on metadata with support of the tag propagation rule cache. On every instruction, the tags of the inputs are used to determine if the operation is allowed, and if so to determine the tags for the results. A tag propagation rule cache miss will trigger the policy miss handler to determine if the instruction is allowed.

HDFI is a hardware-assisted solution for data flow isolation by extending memory words with single bit tags and the RISC-V ISA to perform tag checking [134]. The integrity level of a memory unit is determined by the last write operation to the unit. At memory read, a program would check if the tag is allowed by the security policy. HDFI enforces the Biba integrity and Bell-LaPadula confidentiality security models with <2% performance overhead.

HyperFlow [39] is another extended RISC-V architecture with security features for enforcing information flow control. The processor is designed and verified using a type enforced hardware design language *ChiselFlow* [65]. It supports rich configurable security policies described by complex lattice structures and controlled security class downgrading to allow inter-process communication and system calls. This new architecture also prevents timing channels by enforcing the security policy that the latency of an instruction should not depend on operand value or memory access from another process with a higher security type. A later work extends this architecture by also incorporating language-level IFC for securing the system architecture of secure autonomous vehicles [90].

Several other works build hardware IFT capability into RISC-V processor to provide hardware-assisted security [116, 132]. In [159], ISA based on IFC is designed to mitigate the gap between IFC hardware and software. This ISA is used to prove strong timing-sensitive security conditions about software, leveraging the IFC capabilities and guarantees provided by hardware.

5.3.6 IFT for Speculative Architecture. Speculative Taint Tracking (*STT*) [157] is a framework that executes and selectively forwards the results of speculative access instructions to younger instructions, as long as those younger instructions cannot form a covert channel. *STT* tracks the flow of results from access instructions in a manner similar to *DIFT*, until those results reach an instruction, or sequence of instructions, that may form a covert channel. It automatically “untaints” the result once the instruction that produces it becomes non-speculative to improve performance. The framework enforces a novel form of non-interference, with respect to all speculatively accessed data.

5.4 RTL

A large body of research work aims to develop hardware IFT techniques at the RTL or extend existing HDLs with IFT capability [2–4, 10, 11, 29, 37, 66–68, 84, 86, 162].

5.4.1 State Machine Based IFT Language. *Caisson* [86] and *Sapper* [84] are HDLs enhanced with IFT capability for automated generation of circuits that enforce IFT related security properties. Both languages model hardware designs using state machines and prevent illegal flows from a high state to a low state. In *Caisson*, security policies are enforced through static type checking. However, this technique can cause large area overheads since static property checking requires that resource being hard partitioned or even replicated. *Sapper* improves *Caisson* by taking a hybrid approach to label checking. It uses static analysis to determine a set of dynamic checks to be inserted into the hardware design to enable dynamical tracking. This enables resource re-use and thus lowers the area overheads. In addition, *Sapper* also made an attempt to enforce timing-sensitive noninterference by accounting for information flow related to when events happen.

5.4.2 PCH and Coq Based IFT. Jin and Makris *et al.* integrated IFT into the proof carrying hardware (*PCH*) framework [10, 11, 66–68]. Their techniques define Verilog-to-Coq conversion rules [11] to automate Verilog design to Coq formal

logic translation. The resulting Coq semantic circuit model facilitates tracking secrecy labels and proving information flow security properties. To track the flow of information, the Coq circuit model is extended with a sensitivity label so that the Coq circuit primitives return a `value*sensitivity` pair, where `sensitivity` can be either `secure` or `normal`. The enhanced circuit model is then formally verified under the Coq proof environment to check if the hardware design adheres to desired security properties. A security violation would be identified when a `normal` signal takes a `secure` tag.

The PCH framework provides a method for protecting the hardware design and supply chain. However, it employs a conservative sensitivity label propagation policy, i.e., always allowing secure labels to propagate independent of the operation performed or input values. While this reduces the complexity of the formal Coq circuit model and allows fast profiling of security properties, conservative IFT usually leads to false positives (as we discuss in Section 3). Qin *et al.* made an attempt to eliminate such false positives by incorporating the precise security label propagation policy from GLIFT [144] into the PCH framework [122, 123]. This allows GLIFT logic to be represented in Coq syntax and verified using a new proof tool, but it incurs large overheads in the complexity of the Coq circuit model and proof time.

5.4.3 SecVerilog and SecChisel. *SecVerilog* is a language-based approach for enforcing information flow security [162]. *SecVerilog* augments Verilog with labels and performs static type checking to determine if the design meets the desired security properties. It enforces timing-sensitive noninterference and can be used to design secure architectures that are free of timing channels. The research group later improves the precision of *SecVerilog* by using finer-grained labels to reason about the security of individual array memory elements and bits in data packets. The improved security-type based HDL can be used to verify security properties such as noninterference and uncover hardware vulnerabilities [38]. *SecVerilogLC* supports mutable dependent types to allow secure sharing of hardware resources. A major challenge addressed is implicit downgrading [162] caused by unsynchronized update of data value and security label. *SecVerilogLC* proposes a new type system and introduces an explicit notion of cycle-by-cycle transitions into the syntax, semantics, and type system of the new language to solve the implicit downgrading problem [37].

ChiselFlow is another type enforced HDL with timing labels for developing secure hardware architectures [39]. It has been used to create cryptographic accelerators with hardware-level information flow control and formally verify the security of these implementations [65]. A successive work extends the Chisel hardware design language with security features to create a design-time security verification framework, which employs information flow analysis to verify hardware security properties in order to capture information leaks caused by hardware security flaws and Trojans [29].

5.4.4 RTLIFT. *RTLIFT* precisely captures all the logical information flows, including timing flows, leveraging the cycle-accurate timing information at the RTL [3]. The major challenge lies in precisely accounting for the implicit flows resulting from the complex RTL syntax such as nested conditional branches. This is resolved by converting the branch statements to multiplexer network so that we only need to instantiate IFT logic for the multiplexers. *RTLIFT* allows flexible precision and complexity tradeoffs by instantiating different versions of tracking logic for logical and arithmetical primitives. It achieves over 5X speedup in verification performance over GLIFT [144] and thus has better scalability.

Clepsydra [2] provides a new IFT model for distinguishing timing flows from functional ones and formally verifying timing channel security. A timing information flow occurs when a register is conditionally updated under the control of a tainted value. In addition to the taint label in IFT, *Clepsydra* introduces a timing label to track the propagation of timing variation. The *Clepsydra* tool can formally prove either the existence of a timing leakage or constant execution time.

The *VeriSketch* [4] tool takes a step further and automatically synthesizes hardware designs that satisfy information flow security properties, using the sketch technique for synthesizing designs that meets specified functional constraints.

5.4.5 Control and Data Flow Graph Based RTL IFT. High-level information flow tracking (*HLIFT*) was developed for detecting hardware Trojans that leak secret information through exclusive output [152]. This method extracts IFT features from coarse grained control and data flow graph (CDFG) of hardware Trojan benchmarks. Trojan detection is performed by matching the statement level CDFG of given RTL with the extracted feature library.

5.5 Gate Level

Researchers also aimed to understand the flow of information through Boolean gates. This can be a complement to higher level IFT methods when the hardware design comes in the form of gate netlist, e.g., an IP core. In addition, the complex syntax structures such as conditional branch statements are all flattened to primitive gates and thus, hardware IFT model can be created with significantly lower effort at this level of abstraction.

5.5.1 Gate Level Information Flow Tracking. Gate level information flow tracking (*GLIFT*) is a well-established IFT technique [144]. A key insight is that all information flows, whether explicit or implicit, all appear in a unified form at the gate level and have good mathematical representations. *GLIFT* employs fine-grained taint label and label propagation policy to precisely account for each bit of information flow. It precisely measures the actual flows of information by analyzing the influence of inputs on the output and greatly eliminates the false positives of conservative IFT techniques. As an example, *GLIFT* precisely indicates that a low *reset* clears the high label of a counter while conservative IFT methods cannot.

Hu *et al.* built the theoretical fundamentals of *GLIFT* [57, 111]. They derived formal representations of *GLIFT* logic for arbitrary primitive gates. A successive work formally proved that precise *GLIFT* logic generation is an NP-complete problem. It then presented various algorithms for creating *GLIFT* logic for large digital circuits. The authors also identified the root cause of the differences in precision of *GLIFT* logic created using different methods as static logic hazards or variable correlation resulting from reconvergent fanouts. Such differences motivated IFT precision and complexity tradeoff research [7, 52] discussed in Section 3.4.

Several works attempt to reduce the complexity of *GLIFT* logic by employing optimized label encoding technique [56, 59, 71]. The key observation is that the value of a tainted signal can be ignored in taint propagation, i.e., *GLIFT* logic will produce identical output taint label when a tainted bit takes the value of either 0 or 1. Thus, we can combine the tainted states and reduce the possible signal states to three, namely untainted 0, untainted 1 and tainted *. The don't care condition in tainted signals provides flexibility for circuit optimization. Further, when untainted 0, untainted 1 and tainted * are encoded as 00, 11 and 01 (or 10) respectively, the *GLIFT* logic can serve as hardware redundancy if we set all inputs to be untainted. This allows both information flow security and fault tolerance properties to be enforced with the same *GLIFT* logic [59, 71].

5.5.2 Multi-level, Multi-valued and Multi-flow *GLIFT*. In [54, 55], Hu *et al.* expand *GLIFT* for multilevel security lattices for enforcing multi-level security (MLS). They generalize the *GLIFT* method to arbitrary linear and non-linear security lattices by defining unified yet precise class combining operators. Another work expands *GLIFT* for multi-valued logic systems [138]. It derives taint propagation rules for additional design states such as *unknown* and *high impedance* in order to support multi-valued IFT simulation. Multi-flow IFT [139, 140] is another recent extension to *GLIFT*. It aims to understand the effect of multiple inputs on a signal, e.g., how many key or random number bits are affecting a target bit

at the same time. This technique enables modeling simultaneous information flow behaviors and proving quantitative information flow security properties.

5.5.3 GLIFT Enhanced with Timing Label. Oberg *et al.* demonstrated how GLIFT could be used to enforce timing information flow security [112–114]. In [112], GLIFT is employed for testing timing information flow security in shared bus architectures such as I²C and USB. GLIFT shows that there can be unintended interference between trusted and untrusted devices through timing related design behaviors. They use GLIFT to test for information flow isolation properties between trusted and untrusted IP components, which share computing resources such as the WISHBONE system bus and cryptography core [114]. A later work [113] formalizes the notion of timing channel and provides criteria to separate timing flows from functional ones. A *Timing-only flow* is defined as the case in which a set of inputs affects only the timing of the output. As an example, a change in the key of the RSA cipher affects when the encryption done signal is asserted. The paper then formally proves that GLIFTed state machines could detect timing-only flows.

5.5.4 Gate Level IFT Through Structural Checking. Another work by Le *et al.* tracked data flow in gate level design netlist through structural checking [81]. The proposed method assigns “asset” (or “property”) to signals in hardware designs and then tracks the propagation of “asset” via data flow. It models “asset” related to cryptographic core, debug interface, timing channel and hardware Trojan. By observing the resulting “asset” of signals, the designer can understand important security properties such as integrity and confidentiality. However, structural checking is basically conservative connectivity analysis and thus cannot precisely determine if and when “asset” could propagate like GLIFT.

5.6 Circuit Level

Bidmeshki *et al.* made a first attempt to understand the flow of information in analog/mixed-signal designs [9]. This work defines information flow policies for various analog components such as MOSFETs, bipolar transistors, capacitors, inductors, resistors and diodes and extends the existing PCHIP based IFT method [67, 68] established for digital designs to analog/mixed-signal circuits. While providing information flow control at the level of transistors is desirable due to recent security attacks that target analog hardware, IFT in analog/mixed-signal hardware designs can be a challenging problem. How to encode taint information using continuous analog signals and precisely determine their influence on the output still needs to be investigated.

VeriCoq [11] is recently extended to track sensitive signals in the mixed-signal domain [44]. It detects charge-domain Trojans, a generalization of the A2 Trojan [156] and RowHammer [73] attacks. It proposes an abstracted model of charge-domain leakage structures tainted information flowing to user controllable or accessible flip-flops. Data flow graph based path search is employed for taint propagation along analog components such as MOSFET, Diode and Resistor. A charge-domain Trojan is identified if there is a path between the taint source and target.

Apart from a hierarchical view of hardware IFT methods at different levels of abstraction, a horizontal view helps understand the developing trend in this realm. We mark a few milestones in the development of hardware IFT research as shown in Fig. 9. Early hardware IFT methods focus on the architectural level coupled by numerous works in language based IFT techniques [126]. The strong connection between instructions in software programs and functional units in computer architecture enables a natural extension of the IFT spectrum to the hardware domain. Gate-level IFT methods allow finer grained measurement and classification of information flow. More recent IFT developments primarily target different HDLs, e.g., Verilog and HLS C, for better scalability and verification performance. In addition, IFT techniques have also been extended to combat analog hardware security exploits.




Fig. 9. Milestones in the development of hardware IFT.

6 VERIFICATION TECHNIQUE

IFT techniques can operate in a static or dynamic manner. Static hardware IFT techniques [3, 4, 9, 20, 67, 162] check if the design adheres to desired security properties through *simulation*, *formal methods*, *emulation* or *virtual prototyping* during design time. The IFT model will be removed when verification completes. Dynamic hardware IFT [27, 28, 70, 83, 130, 134, 137, 151] is a *runtime* mechanism that augments the original hardware design with tracking logic that monitors information flow behaviors and prevents harmful flows of information. This typically comes at a cost, e.g., additional resource usage and performance overheads.

6.1 Simulation

Simulation is a common method of verifying the behaviors of hardware designs. Functional simulation involves feeding a testbench to a circuit design and observing if the results meet the desired specification. Hardware security simulation plays a similar role except the key output is not functional correctness rather if the circuit has violated any security policies. Figure 10 illustrates the difference between functional and IFT based security simulation for identifying information leakage.




Fig. 10. The difference between functional and IFT based security simulation.

In both the functional and security simulation waveforms, there are switching activities in the output signal *Antena* (as spelled in the Trust-HUB AES-T1700 benchmark). However, functional simulation cannot reveal the connection between such switching activities and the key. IFT based security simulation associates security labels (i.e., *key_t* and *Antena_t*) with the signals, reflecting their security attributes. Here, we label the key as high, i.e., *key_t* = 128'hFFFFFFFF_FFFFFFFF_FFFFFFFF_FFFFFFFF. From Fig. 10 (b), the security label *Antena_t* can be logical 1 in the dotted red boxes, indicating that *Antena* can carry high information. In other words, *key* leaks to *Antena* during these time slots. The security label *Antena_t* accurately indicates when and where leakage happens even if there might be no switching activity (i.e., leaking a logical 0) at all in the *Antena* signal.

Some hardware IFT techniques describe their information flow model using standard HDL (e.g., GLIFT [144] and RTLIFT [3]). In this case, it is possible to perform information flow security simulation using standard or open source EDA simulation tools such as *Mentor Graphics ModelSim* or *QuestaSim*, *Synopsys VCS*, and *Icarus Verilog*.

Tortuga Logic and *Cadence* have released a commercial hardware security simulation tool, namely the *Radix-S* [98] hardware root of trust security verification framework, which automates the information flow model annotation process for IFT based hardware security simulation using standard HDL and EDA simulation environment.

6.2 Formal Verification

Verification tools can be used for formal (or semi-formal) verification of security properties on information flow models through equivalence checking, SAT solving, theorem proving, or type checking. The benefits of formal approaches are that the hardware is guaranteed to be secure with respect to the properties proven. However, a major downside is that formal methods inherently do not scale due to the exponential size of design state space. Certain properties may not be proved using hundreds of Giga bytes of memory after running for days, e.g., a hardware Trojan triggered under a specific input pattern or sequence. In addition, the verification outcomes heavily rely on the quality and completeness of the security properties specified.

6.2.1 Equivalence Checking: Equivalence checking is the most frequently used formal verification technique for verifying functional correctness. To check if tainted information propagates from a source signal to a destination point, we can set the taint source to logical 0 and 1 respectively and create two instances of the circuit, i.e., miter circuits. We then check for equivalence at the destination point to see if the taint source has an effect on the destination. This provides an approach to secure path and X-propagation verification [14–16, 43, 45, 94]. These verification tools can be used to prove security properties such as confidentiality and integrity and identify insecure design paths that could lead to information flow security property violations.

6.2.2 SAT Solving: SAT solvers such as *minSAT* and *zchaff* are effective tools for proving functional properties. Several open source EDA tools such as *ABC* and *Yosys* integrate these SAT solvers and HDL to CNF (conjunctive normal form) formula converters in order to prove design properties on digital circuits. Therefore, they can parse information flow models written in HDL, e.g., IFT models generated by GLIFT [144] and RTLIFT [3], and prove information flow security properties as illustrated in Section 4. Such security properties can be easily converted to proof constraints and incorporated by the solver. The following shows an example of security property mapping for *Yosys* where we constrain the 32-bit *key* to high while *clk* to low and prove that the *ready* output should always be low.

A major difference from simulation is that a counter example will be given when the solver captures a security property violation, indicating exactly when such violation could possibly occur. This provides insightful hints for hardware designers to pinpoint security vulnerabilities. While SAT tools can be powerful for verifying simple security properties

```

1: set key := high
2: assert ready == low
3: sat -prove ready_t 1'b0 -set key_t 32'hFFFFFFF -set clk_t 1'b0

```

on combinational circuits, they have some inherent limitations as well. They can only verify qualitative security properties on CNF formula. SMT solvers such as *Yices* can complete such limitation and prove quantitative security properties on QBF (Quantified Boolean Formula) [7, 52]. In addition, SAT solving on sequential digital circuits can be a challenging task. SAT tools cannot automatically trace beyond register boundaries and thus usually restrict the proof to a minimal combinational region. Unrolling the design by a number of time frames is a frequently used technique for proving properties on sequential designs. However, this technique can lead to huge memory and verification performance overheads.

Commercial formal proof tools can be more effective in verifying security properties and search for security violations. *Mentor Graphics Questa Formal* takes a semi-formal approach to IFT verification. It takes as inputs the initialization vectors and proof-time constraints, simulates the design to a specified proof radius in number of clock cycles and also performs state space search to identify all reachable design states. As shown in Tables 1 to 3, security properties are specified using property specification language (e.g., SVA) assertions, which are familiar to hardware designers.

With the security constraints specified in initialization and directive files, the SVA property asserts that at the rising edge of the *ready* signal, it should always be low. Existing work has demonstrated how the *Mentor Graphics Questa Formal* tool can be used to prove information flow security properties and identify hard-to-detect security vulnerabilities such as timing channels, insecure debug ports and hardware Trojans [53, 123].

6.2.3 Theorem Proving: A number of hardware IFT methods verify security properties through theorem proving. Information flow security properties are specified as security theorems while security constraints as pre-conditions. Formal tools such as *Coq* can be used to verify these properties [10, 11, 66–68, 122, 123]. There are several HDL annotation tools that convert RTL hardware designs or gate level netlists to formal Coq circuits augmented with security labels [11]. While the formal IFT model construction has largely been automated, the proof process still requires significant amount of human interaction. New security theorems need to be derived as tainted information reaches new signals. The designer needs to decide the signals to be checked in the next time frame.

6.2.4 Type Checking: *Type system* is a set of rules that assign a property called *type* to each computed value, examines the flow of these values, and attempts to ensure or prove that no type error can occur. It identifies a type error when a certain type of value is used in an operation where not allowed, e.g., when an untrusted variable is used as the jump target in a secure execution context. Type system can be employed for proving security properties by associating variables with security-related types, such as Confidential and Unclassified. Access or information flow control can be enforced through either static or runtime type checking.

The SecVerilog project [121] releases an open source tool for verifying timing information flow security in Verilog designs. The tool extends Verilog with annotations that support reasoning about information flows including timing flows through compile time type checking. SecVerilog associates variables with security types (e.g., {H} and {L}) and labels to specify the information flow policy to be enforced. Verification is done by checking if information flows cause violation of type rules.

6.3 Emulation

With the evolution of IC verification platforms, leading EDA companies have also developed tools to verify information flow security properties of hardware designs on FPGA emulation servers. Hardware emulation typically can achieve 10X+ verification performance than software simulation [95]. Another advantage with hardware emulation is hardware-software security co-verification, which is not yet possible for formal verification tools. Emulation can be a promising approach that leverages the information flow measurement capability of the underlying hardware to perform security co-verification and detect software attacks that exploit hardware security vulnerabilities.

The *Radix-M* [97] hardware security platform from *Tortuga Logic* and *Cadence* is one such emulation tool that performs firmware security validation on full system-on-chip (SoC) designs. The tool deploys hardware IFT models on an FPGA emulation server to run security verifications. Such verification technology can also be used to verify the security of firmware and even software, providing a promising solution to system level security.

6.4 Virtual Prototyping

Virtual prototyping involves creating abstract models/libraries for the different hardware components, which allows verifying the system software before the hardware design is complete. Hassan *et al.* proposed an approach for SoC security validation through Virtual Prototypes (VP) modeled using SystemC TLM 2.0 [46, 118]. They employ static information flow analysis to verify the *no flow* security property and detect potential security breaches such as data leakage and untrusted access. High-level abstractive information flow models will lead to better verification performance and scalability. However, it can only employ conservative label propagation policy to profile the possible information flow violations due to the lack of design implementation details at this stage.

6.5 Runtime

Dynamic IFT techniques deploy dedicated hardware to monitor and control information flows during runtime. These dynamic IFT techniques usually target the RTL level or higher, as it becomes costly to perform dynamic IFT at lower levels of abstractions. Additional label memory and IFT logic needs to be integrated into the circuit. At the architectural level, this typically involves a tag propagation unit or co-processor that sits along with the main processor to determine the taint labels of instruction outputs. The labels will then get involved in successive operations or be written back to update the label memory. These techniques usually need to load all the parameters of instructions to the tag propagation unit or co-processor in order to calculate the output taint label. Such parameters can be retrieved either from the execution path of the processor or core debug interface.

Dynamic IFT techniques perform information flow security policy checking on the execution trace, which is typically a small subset of the entire state space. This prevents the state space explosion problem of formal methods and allows the methods to scale to large circuits. However, the additional taint propagation and checking logic typically increases resource usage and power consumption. Even worse, such logic and its synchronization with the original execution context will inevitably cause considerable performance overheads.

As a comparison, formal verification based IFT techniques can completely verify IPs and small subsystems but does not scale to complex system level designs due to the state space explosion problem. Simulation and emulation based IFT approaches have better scalability. However, they cannot cover the entire design state space. Virtual Prototyping allows a quick but conservative profiling of possible information flow security violations. Runtime IFT mechanisms can monitor realistic information flow behaviors but incurs large area and performance penalties.

Table 4 summarizes the different hardware IFT techniques in terms of their verification mechanisms. We note that some techniques can be used with different verification techniques, e.g., GLIFT can be used for simulation, emulation, formal verification, and during runtime. Other techniques target a specific area, e.g., VeriCoq [11, 67] and SecVerilog [121, 162] solely use formal verification techniques.

Table 4. Summary of hardware IFT techniques in terms of verification mechanisms used.

Simulation	Caiison[86], Sapper[84], GLIFT[144], GLIFT-SoC[72], GLIFT Timing[112–114], Multi-level GLIFT[54, 55], Multi-value GLIFT[138], Multi-flow IFT[140] Radix-S[98], RTLIFT[3], Clepsydra[2]
Formal Verification	ASSURE[64], GARUDA[129], TaintHLS[119], ChiselFlow[65], Zagieboylo[159], PCH-Coq[10, 66–68], VeriCoq[11], Qin[122, 123], SecVerilog[121, 162] SecVerilogLC[37], Ferraiuolo[38], SecChisel[29], RTLIFT[3], Clepsydra[2], VeriSketch[4], HLIFT[152], Hu[52], Becker[7], GLIFT Timing[112–114] Multi-level GLIFT[54, 55], Multi-value GLIFT[138], Multi-flow IFT[140], Le[81], VeriCoq-analog[9], Guo[44], Nicholas[109], Siddiqui[133], SCRIPT[106]
Emulation	GLIFT[144], RTLIFT[3], Clepsydra[2], Radix-M[97]
Virtual Prototyping	IFT-VP[46], IFT-VP+[118]
Runtime	CDI-IFT[82, 83], PTM-IFT [150], ARMHEX[151], BFBlare[149], PAU[48], IFTU[130], IFT-Accelerator [117, 120], GLIFT-Embedded[102], GLIFT-IoT[20] TaintHLS[119], Saal[125], Bondi[12], DJFT[137], Kannan[70], RIFLE[146], Raksha[24], FlexiTaint[147], Deng[27, 28], HAFIX[25], Shioya[131], Joannou[69] PIFT[92, 93], LATCH[145], DataSafe[19], GLIFT-Lease[87, 141–143], METACE[91], Nagarajan[105], ShadowReplica[63], SIFT[115], Nightingale[110] Chen[18], Ruwase[124], SAFE[5], Nicholas[109], Siddiqui[133], PUMP[34], HDFI[134], HyperFlow[39], Liu[90], Palmicro[116], Shirley[132] Zagieboylo[159], STI[157], Caiison[86], Sapper[84], GLIFT[144], GLIFT-TMR[59, 71]

7 FUTURE RESEARCH DIRECTIONS AND CHALLENGES

7.1 Automated Security Property Generation

Security properties are crucial for hardware IFT verification. High-quality security properties allow verification tools to quickly pinpoint potential security flaws. Unfortunately, specifying security properties that cover a wide range of vulnerabilities is all too often a challenging and daunting process. Crafting hardware security properties requires expertise in formal verification, expertise in security, and knowledge of the design under verification. The job is difficult and time consuming, and critical properties are easily missed. For example, participants of Hack@DAC’18 were able to write properties covering only 17 of the 31 bugs, which “indicates that the main challenge with using formal verification tools is identifying and expressing security properties that the tools are capable of capturing and checking” [33]. Automated generation of high-quality security properties is an important problem, especially for developing property driven hardware security solutions [49, 50]. Two interesting directions in this regard are: 1) developing parameterized security properties and 2) higher level property languages that build upon the existing hardware IFT models, e.g., an access control policy language. We discuss both in the following.

7.2 Parameterized Security Properties

Verification IP (VIP) plays an important role in the hardware design process by providing a set of functional inputs (testbenches, assertions, etc.) for functional verification of a hardware design [47, 158, 164]. The general goal is to determine whether your IP core correctly interfaces with memory protocols (DDR, QSPI, SD), PCIe, USB, ethernet and other communication protocols, and AXI, WISHBONE, and other on-chip interfaces. A similar approach could be done to develop *security VIP* that focuses on security verification rather than functional verification. One potential direction is to develop security property templates based upon information flow metrics [51] and leverage functional property templates from existing assertion based verification approaches [41]. Ideally, the designer would minimally change property parameters based upon their desired security constraints and their hardware design, and the tool would automatically provide a series relevant properties. Transys provides evidence that this approach is possible [163], but we feel there is much room to generalize and enhance this process.

7.3 Hardware Access Control Policy Language

Modern hardware is comprised of tens to hundreds of different IP cores – processors, custom accelerators, memories, peripherals, timers, and other system resources – that are interconnected through a network of on-chip busses and routers. These IP cores, busses, and memories each have their own security policies which govern how they can behave and how other IP cores can access their data, control registers, and status signals. Since access control is a prevalent and important aspect of hardware security, a property specification language specifically targeted towards access control could easily allow security experts to specify a wide range of on-chip access control properties such that they can be compiled to properties that can be used by existing hardware IFT tools. The properties could utilize lattices to denote allowable accesses. Additionally, the properties should be able to specify temporal rules, e.g., a memory location cannot be read by certain processes after sensitive information has been written there. Also, they should allow for easy specification of common policies like compartmentalization, access control lists, secure handoff, Chinese wall, and redaction. Huffmire et al. defined a language for policies related to off-chip memory controllers [62]. This could be used as a baseline for a more general access control language by modifying it for on-chip resources and provide a clear mapping to hardware security verification properties models.

7.4 Novel Information Flow Models

Hardware IFT opens the door for unique flow models. For example, timing flows become apparent at the RTL [113] and Clepsydra developed a model that updates labels based only on whether there is a timing channel. Note that GLIFT does not do this as it encodes functional and timing models simultaneously (i.e., it can only say a flow occurs, and cannot differentiate if that flow is due to a timing channel or a functional channel). Thus, Clepsydra provided a more expressive IFT model for tracking timing flows. We see opportunities to use IFT for other side channels like power side channel [106], faults, and electromagnetic radiation. We believe this would require more quantitative flow models as opposed to the vast majority of the existing models with take a qualitative approach (there is a flow violation or not). For these side channels, it is often important to know how much information is flowing in addition to whether there was a flow or not. There is some initial work to develop tools for measuring quantitative information flow properties such as leakage and randomness [49, 140]. We believe there is more to explore in these directions. Additionally, IFT models that can answer questions related to dependability and resilience properties such as availability, reliability, safety and fault tolerance is also an interesting research vector.

7.5 Automated Debugging

Hardware IFT verification often uncovers a security violation. The verification tools may provide a witness, i.e., an example trace describing an execution where the property is violated, but nothing more than that. The onus is on the hardware designer to modify the design and eliminate the flaw. This is a challenging and mostly manual process. Any debugging tools that aid this process provide immense value. Error localization is a fundamental part of this process which identifies the parts of the hardware that need to be modified in order to fix the flaw. If this can be done accurately, one could imagine additional tools to help the verification engineer modify the design to fix the flaws. While fully automated debugging is a huge challenge, there seems to be ample opportunities and value for partially automating the hardware security debugging process.

7.6 Firmware IFT

Firmware is an attractive attack surface for hackers since critical system credentials (e.g., encryption keys and low level access credentials) are stored in or managed by firmware. In addition, firmware provides direct access to the hardware and can be exploited to root the system, uncover confidential system security parameters, and perform other nefarious activities. Firmware requires tracking of information across hardware and software. This is challenging. Firmware IFT verification tools must formalize the flow models for the computational constructs at this level of abstraction. This requires research in combined hardware/software flow models. Subramanyan et al. [136] provide some great initial progress in this space. But there is much more to consider including scalability, mixed hardware/software IFT modeling, and property modeling.

7.7 Analog/Mixed-signal Hardware IFT

Another potential research direction is to move to even lower levels of abstraction and develop IFT techniques for analog or mix-signal hardware designs. Recent work has made a first attempt to develop IFT techniques that allow co-verification of digital and analog hardware to help uncover the cross-domain security vulnerabilities [9]. A fundamental problem of performing IFT in analog hardware is label encoding. This is fundamentally different from the digital world since analog signals are continuous signals. Binary labels are no longer eligible for encoding and distinguishing the complex label information. Consequently, we need to define policies for multi-level and/or multi-bit taint label propagation.

7.8 Relationship between IFT and Traditional Function Hardware Verification and Test

Security testing and verification closely relate to several classic problems in switching circuit theory, including controllability, observability, X-propagation, and fault propagation [89]. Taint propagation and X-propagation can be viewed as the same problem [61]. As an example, *Mentor Graphics* and *Cadence* have both integrated security path and X-propagation Apps into their formal verification tool chain. An interesting research direction is leveraging hardware IFT models to understand these classic testing problems and attempts to formalize them under a unified framework. A recent work has shown how hardware IFT provides a unified model for verifying both taint- and X-propagation properties [61]. On the other hand, existing controllability, observability, X-propagation and fault propagation analysis tools also provides an approach to performing IFT. A benefit with these tools is that these tools all work on the original circuit model, which has significantly lower complexity than IFT models.

8 CONCLUSION

This article performs a comprehensive review of existing work in the hardware IFT spectrum. Our aim is to draw a clear picture of the development of this technique, bring forward the potential research directions, and discuss the challenges to be resolved. Our intention is to introduce the hardware IFT technique to a wider range of audiences in the hardware security, testing and verification communities and show how this technique can be employed to verify if a hardware design adheres to both security and reliability properties.

REFERENCES

- [1] S. Ao and Guo Shuangzhou. 2011. A enhancement technology about system security based on dynamic information flow tracking. In *2011 2nd International Conference on Artificial Intelligence, Management Science and Electronic Commerce (AIMSEC)*. 6108–6111.

- [2] A. Ardeshiricham, W. Hu, and R. Kastner. 2017. Clepsydra: Modeling timing flows in hardware designs. In *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 147–154.
- [3] A. Ardeshiricham, W. Hu, J. Marxen, and R. Kastner. 2017. Register transfer level information flow tracking for provably secure hardware design. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2017*. 1691–1696.
- [4] Armaiti Ardeshiricham, Yoshiki Takashima, Sicun Gao, and Ryan Kastner. 2019. VeriSketch: Synthesizing Secure Hardware Designs with Timing-Sensitive Information Flow Properties. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (London, United Kingdom) (CCS '19)*. ACM, New York, NY, USA, 1623–1638.
- [5] 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 *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14)* (San Diego, California, USA). ACM, New York, NY, USA, 165–178. <https://doi.org/10.1145/2535838.2535839>
- [6] Jean Bacon, David Eysers, Thomas FJ-M Pasquier, Jatinder Singh, Ioannis Papagiannis, and Peter Pietzuch. 2014. Information flow control for secure cloud computing. *IEEE Transactions on Network and Service Management* 11, 1 (2014), 76–89.
- [7] A. Becker, W. Hu, Y. Tai, P. Brisk, R. Kastner, and P. Ienne. 2017. Arbitrary precision and complexity tradeoffs for gate-level information flow tracking. In *2017 54th ACM/EDAC/IEEE Design Automation Conference (DAC)*. 1–6.
- [8] Lionel Bening and Harry Foster. 2001. *Principles of Verifiable RTL Design* (2nd ed.). Kluwer Academic Publishers, USA.
- [9] M. Bidmeshki, A. Antonopoulos, and Y. Makris. 2017. Information flow tracking in analog/mixed-signal designs through proof-carrying hardware IP. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2017*. 1703–1708.
- [10] M. Bidmeshki, X. Guo, R. G. Dutta, Y. Jin, and Y. Makris. 2017. Data Secrecy Protection Through Information Flow Tracking in Proof-Carrying Hardware IP—Part II: Framework Automation. *IEEE Transactions on Information Forensics and Security* 12, 10 (Oct 2017), 2430–2443.
- [11] M. Bidmeshki and Y. Makris. 2015. Toward automatic proof generation for information flow policies in third-party hardware IP. In *2015 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*. 163–168.
- [12] J. O. Bondi and M. A. Branstad. 1989. Architectural support of fine-grained secure computing. In *Proceedings of the Fifth Annual Computer Security Applications Conference*. 121–130.
- [13] Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F. Wenisch, Yuval Yarom, and Raoul Strackx. 2018. Foreshadow: Extracting the Keys to the Intel SGX Kingdom with Transient Out-of-Order Execution. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, Baltimore, MD, 991–1008. <https://www.usenix.org/conference/usenixsecurity18/presentation/bulck>
- [14] G. Cabodi, P. Camurati, S. F. Finocchiaro, C. Loiacono, F. Savarese, and D. Vendraminetto. 2016. Secure Path Verification. In *2016 1st IEEE International Verification and Security Workshop (IVSW)*. 1–6.
- [15] G. Cabodi, P. Camurati, S. F. Finocchiaro, F. Savarese, and D. Vendraminetto. 2017. Embedded Systems Secure Path Verification at the Hardware/Software Interface. *IEEE Design Test* 34, 5 (Oct 2017), 38–46.
- [16] Cadence. 2016. JasperGold Security Path Verification App. https://www.cadence.com/content/cadence-www/global/en_US/home/tools/system-design-and-verification/formal-and-static-verification/jasper-gold-verification-platform/security-path-verification-app.html.
- [17] Andrew Canis, Jongsok Choi, Blair Fort, Bain Syrowik, Ruo Long Lian, Yu Ting Chen, Hsuan Hsiao, Jeffrey Goeders, Stephen Brown, and Jason Anderson. 2016. *LegUp High-Level Synthesis*. Springer International Publishing, Cham, 175–190.
- [18] Shimin Chen, Michael Kozuch, Theodoros Strigkos, Babak Falsafi, Phillip B. Gibbons, Todd C. Mowry, Vijaya Ramachandran, Olatunji Ruwase, Michael Ryan, and Evangelos Vlachos. 2008. Flexible Hardware Acceleration for Instruction-Grain Program Monitoring. In *Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA '08)*. IEEE Computer Society, USA, 377–388.
- [19] Yu-Yuan Chen, Pramod A. Jamkhedkar, and Ruby B. Lee. 2012. A Software-hardware Architecture for Self-protecting Data. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security (Raleigh, North Carolina, USA) (CCS '12)*. ACM, New York, NY, USA, 14–27.
- [20] Hari Cherupalli, Henry Duwe, Weidong Ye, Rakesh Kumar, and John Sartori. 2017. Software-based Gate-level Information Flow Security for IoT Systems. In *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture (Cambridge, Massachusetts) (MICRO-50 '17)*. ACM, New York, NY, USA, 328–340.
- [21] Michael R Clarkson and Fred B Schneider. 2010. Hyperproperties. *Journal of Computer Security* 18, 6 (2010), 1157–1210.
- [22] Ellis Cohen. 1977. Information transmission in computational systems. In *ACM SIGOPS Operating Systems Review*, Vol. 11. ACM, 133–139.
- [23] Ellis S Cohen. 1978. Information transmission in sequential programs. *Foundations of Secure Computation* (1978), 297–335.
- [24] Michael Dalton, Hari Kannan, and Christos Kozyrakis. 2007. Raksha: A Flexible Information Flow Architecture for Software Security. *SIGARCH Comput. Archit. News* 35, 2 (June 2007), 482–493.
- [25] L. Davi, M. Hanreich, D. Paul, A. Sadeghi, P. Koeberl, D. Sullivan, O. Arias, and Y. Jin. 2015. HAFIX: Hardware-Assisted Flow Integrity eXtension. In *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*. 1–6.
- [26] Onur Demir, Wenjie Xiong, Faisal Zaghoul, and Jakub Szefer. 2016. Survey of Approaches for Security Verification of Hardware/Software Systems. Cryptology ePrint Archive, Report 2016/846. <https://eprint.iacr.org/2016/846>.
- [27] D. Y. Deng, D. Lo, G. Malysa, S. Schneider, and G. E. Suh. 2010. Flexible and Efficient Instruction-Grained Run-Time Monitoring Using On-Chip Reconfigurable Fabric. In *2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture*. 137–148.
- [28] D. Y. Deng and G. E. Suh. 2012. High-performance parallel accelerator for flexible and efficient run-time monitoring. In *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2012)*. 1–12.

- [29] Shuwen Deng, Dođuhan Gümüőođlu, Wenjie Xiong, Sercan Sari, Y. Serhan Gener, Corine Lu, Onur Demir, and Jakub Szefer. 2019. SecChisel Framework for Security Verification of Secure Processor Architectures. In *Proceedings of the 8th International Workshop on Hardware and Architectural Support for Security and Privacy* (Phoenix, AZ, USA) (HASP '19). ACM, New York, NY, USA, Article 7, 8 pages.
- [30] Dorothy Elizabeth Robling Denning. 1975. *Secure Information Flow in Computer Systems*. Ph.D. Dissertation. West Lafayette, IN, USA. AAI7600514.
- [31] Dorothy Elizabeth Robling Denning. 1976. A lattice model of secure information flow. *Commun. ACM* 19, 5 (1976), 236–243.
- [32] Jack B Dennis and Earl C Van Horn. 1966. Programming semantics for multiprogrammed computations. *Commun. ACM* 9, 3 (1966), 143–155.
- [33] Ghada Dessouky, David Gens, Patrick Haney, Garrett Persyn, Arun Kanuparthi, Hareesh Khattri, Jason M Fung, Ahmad-Reza Sadeghi, and Jeyavijayan Rajendran. 2019. Hardfails: Insights into software-exploitable hardware bugs. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*. 213–230.
- [34] Udit Dhawan, Catalin Hritcu, Raphael Rubin, Nikos Vasilakis, Silviu Chiricescu, Jonathan M. Smith, Thomas F. Knight, Benjamin C. Pierce, and Andre DeHon. 2015. Architectural Support for Software-Defined Metadata Processing. *SIGPLAN Not.* 50, 4 (March 2015), 487–502.
- [35] Petros Efstathopoulos, Maxwell Krohn, Steve VanDeBogart, Cliff Frey, David Ziegler, Eddie Kohler, David Mazieres, Frans Kaashoek, and Robert Morris. 2005. Labels and event processes in the Asbestos operating system. *ACM SIGOPS Operating Systems Review* 39, 5 (2005), 17–30.
- [36] Manuel Egele, Christopher Kruegel, Engin Kirda, Heng Yin, and Dawn Song. 2007. Dynamic spyware analysis. (2007).
- [37] A. Ferraiuolo, Weizhe Hua, A. C. Myers, and G. E. Suh. 2017. Secure information flow verification with mutable dependent types. In *2017 54th ACM/EDAC/IEEE Design Automation Conference (DAC)*. 1–6.
- [38] 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. *SIGARCH Comput. Archit. News* 45, 1 (April 2017), 555–568.
- [39] Andrew Ferraiuolo, Mark Zhao, Andrew C. Myers, and G. Edward Suh. 2018. HyperFlow: A Processor Architecture for Nonmalleable, Timing-Safe Information Flow Security. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (Toronto, Canada) (CCS '18)*. ACM, New York, NY, USA, 1583–1600.
- [40] Harry Foster. 2009. *Applied assertion-based verification: An industry perspective*. Now Publishers Inc.
- [41] Harry Foster, Adam Krolnik, and David Lacey. 2005. *Assertion-Based Design*. Springer US, Boston, MA. 61–102 pages.
- [42] Joseph A. Goguen and Jose Meseguer. 1982. Security Policies and Security Models. In *IEEE Symposium on Security & Privacy*. 11–20.
- [43] Mentor Graphics. 2016. Questa secure check - exhaustive verification of secure paths to critical hardware storage. <https://www.mentor.com/products/fv/questa-secure-check>.
- [44] X. Guo, H. Zhu, Y. Jin, and X. Zhang. 2019. When Capacitors Attack: Formal Method Driven Design and Detection of Charge-Domain Trojans. In *2019 Design, Automation Test in Europe Conference Exhibition (DATE)*. 1727–1732.
- [45] Ziyad Hanna. 2013. Jasper case study on formally verifying secure on-chip datapaths. <http://www.deepchip.com/items/0524-03.html>.
- [46] Muhammad Hassan, Vladimir Herdt, Hoang M. Le, Daniel Groőe, and Rolf Drechsler. 2017. Early SoC Security Validation by VP-Based Static Information Flow Analysis. In *Proceedings of the 36th International Conference on Computer-Aided Design (Irvine, California) (ICCAD '17)*. IEEE Press, 400–407.
- [47] Hannibal Height. 2013. *A practical guide to adopting the universal verification methodology (UVM)*. Lulu. com.
- [48] Ingoo Heo, Minsu Kim, Yongje Lee, Changho Choi, Jinyong Lee, Brent Byunghoon Kang, and Yunheung Paek. 2015. Implementing an Application-Specific Instruction-Set Processor for System-Level Dynamic Program Analysis Engines. *ACM Trans. Des. Autom. Electron. Syst.* 20, 4, Article 53 (Sept. 2015), 32 pages.
- [49] W. Hu, A. Althoff, A. Ardeshiricham, and R. Kastner. 2016. Towards Property Driven Hardware Security. In *2016 17th International Workshop on Microprocessor and SOC Test and Verification (MTV)*. 51–56.
- [50] Wei Hu, Armaiti Ardeshiricham, Mustafa S Gobulukoglu, Xinmu Wang, and Ryan Kastner. 2018. Property specific information flow analysis for hardware security verification. In *2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 1–8.
- [51] Wei Hu, Armaiti Ardeshiricham, and Ryan Kastner. 2017. Identifying and measuring security critical path for uncovering circuit vulnerabilities. In *2017 18th International Workshop on Microprocessor and SOC Test and Verification (MTV)*. IEEE, 62–67.
- [52] Wei Hu, A. Becker, A. Ardeshiricham, Yu Tai, P. lenne, D. Mu, and R. Kastner. 2016. Imprecise security: Quality and complexity tradeoffs for hardware information flow tracking. In *2016 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 1–8.
- [53] W. Hu, B. Mao, J. Oberg, and R. Kastner. 2016. Detecting Hardware Trojans with Gate-Level Information-Flow Tracking. *Computer* 49, 8 (Aug 2016), 44–52.
- [54] Wei Hu, Dejun Mu, Jason Oberg, Baolei Mao, Mohit Tiwari, Timothy Sherwood, and Ryan Kastner. 2014. Gate-Level Information Flow Tracking for Security Lattices. *ACM Trans. Des. Autom. Electron. Syst.* 20, 1, Article 2 (Nov. 2014), 25 pages.
- [55] W. Hu, J. Oberg, J. Barrientos, D. Mu, and R. Kastner. 2013. Expanding Gate Level Information Flow Tracking for Multilevel Security. *IEEE Embedded Systems Letters* 5, 2 (June 2013), 25–28.
- [56] W. Hu, J. Oberg, A. Irturk, M. Tiwari, T. Sherwood, D. Mu, , and R. Kastner. 2011. An Improved Encoding Technique for Gate Level Information Flow Tracking. In *International Workshop on Logic & Synthesis*.
- [57] W. Hu, J. Oberg, A. Irturk, M. Tiwari, T. Sherwood, D. Mu, and R. Kastner. 2011. Theoretical Fundamentals of Gate Level Information Flow Tracking. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 30, 8 (Aug 2011), 1128–1140.
- [58] W. Hu, J. Oberg, A. Irturk, M. Tiwari, T. Sherwood, D. Mu, and R. Kastner. 2012. On the Complexity of Generating Gate Level Information Flow Tracking Logic. *IEEE Transactions on Information Forensics and Security* 7, 3 (June 2012), 1067–1080.

- [59] Wei Hu, Jason Oberg, Dejun Mu, and Ryan Kastner. 2012. Simultaneous Information Flow Security and Circuit Redundancy in Boolean Gates. In *Proceedings of the International Conference on Computer-Aided Design* (San Jose, California) (ICCAD '12). ACM, New York, NY, USA, 585–590.
- [60] W. Hu, X. Wang, and D. Mu. 2018. Security Path Verification Through Joint Information Flow Analysis. In *2018 IEEE Asia Pacific Conference on Circuits and Systems (APCCAS)*. 415–418.
- [61] Wei Hu, Linjuan Wu, Jing Tan, Yu Tai, and Jiliang Zhang. 2020. A Unified Formal Model for Proving Security and Reliability Properties. In *Asian Test Symposium (ATS)*.
- [62] Ted Huffmire, Shreyas Prasad, Tim Sherwood, and Ryan Kastner. 2006. Policy-Driven Memory Protection for Reconfigurable Hardware. In *European Conference on Research in Computer Security* (Hamburg, Germany) (ESORICS'06). Springer-Verlag, Berlin, Heidelberg, 461–478. https://doi.org/10.1007/11863908_28
- [63] Kangkook Jee, Vasileios P. Kemerlis, Angelos D. Keromytis, and Georgios Portokalidis. 2013. ShadowReplica: Efficient Parallelization of Dynamic Data Flow Tracking. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security* (Berlin, Germany) (CCS '13). ACM, New York, NY, USA, 235–246.
- [64] Zhenghong Jiang, Steve Dai, G. Edward Suh, and Zhiru Zhang. 2018. High-Level Synthesis with Timing-Sensitive Information Flow Enforcement. In *Proceedings of the International Conference on Computer-Aided Design* (San Diego, California) (ICCAD '18). ACM, New York, NY, USA, Article 88, 8 pages.
- [65] Zhenghong Jiang, Hanchen Jin, G. Edward Suh, and Zhiru Zhang. 2019. Designing Secure Cryptographic Accelerators with Information Flow Enforcement: A Case Study on AES. In *Proceedings of the 56th Annual Design Automation Conference 2019* (Las Vegas, NV, USA) (DAC '19). ACM, New York, NY, USA, Article 59, 6 pages.
- [66] Y. Jin, X. Guo, R. G. Dutta, M. Bidmeshki, and Y. Makris. 2017. Data Secrecy Protection Through Information Flow Tracking in Proof-Carrying Hardware IP—Part I: Framework Fundamentals. *IEEE Transactions on Information Forensics and Security* 12, 10 (Oct 2017), 2416–2429.
- [67] Y. Jin and Y. Makris. 2012. Proof carrying-based information flow tracking for data secrecy protection and hardware trust. In *2012 IEEE 30th VLSI Test Symposium (VTS)*. 252–257.
- [68] Y. Jin, B. Yang, and Y. Makris. 2013. Cycle-accurate information assurance by proof-carrying based signal sensitivity tracing. In *2013 IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*. 99–106.
- [69] A. Joannou, J. Woodruff, R. Kovacsics, S. W. Moore, A. Bradbury, H. Xia, R. N. M. Watson, D. Chisnall, M. Roe, B. Davis, E. Napierala, J. Baldwin, K. Gudka, P. G. Neumann, A. Mazzinghi, A. Richardson, S. Son, and A. T. Marketos. 2017. Efficient Tagged Memory. In *2017 IEEE International Conference on Computer Design (ICCD)*. 641–648.
- [70] H. Kannan, M. Dalton, and C. Kozyrakis. 2009. Decoupling Dynamic Information Flow Tracking with a dedicated coprocessor. In *2009 IEEE/IFIP International Conference on Dependable Systems Networks*. 105–114.
- [71] Ryan Kastner. 2013. Circuit Primitives for Monitoring Information Flow and Enabling Redundancy. In *Proceedings of the 8th International Conference on Hardware and Software: Verification and Testing* (Haifa, Israel) (HVC'12). Springer-Verlag, Berlin, Heidelberg, 6–6.
- [72] R. Kastner, J. Oberg, W. Hu, and A. Irturk. 2011. Enforcing Information Flow Guarantees in Reconfigurable Systems with Mix-trusted IP. In *International Conference on Engineering of Reconfigurable Systems and Algorithms*.
- [73] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. 2014. Flipping Bits in Memory without Accessing Them: An Experimental Study of DRAM Disturbance Errors. *SIGARCH Comput. Archit. News* 42, 3 (June 2014), 361–372.
- [74] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. 2009. seL4: Formal Verification of an OS Kernel. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles* (Big Sky, Montana, USA) (SOSP '09). ACM, New York, NY, USA, 207–220. <https://doi.org/10.1145/1629575.1629596>
- [75] Paul Kocher, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. 2018. Spectre Attacks: Exploiting Speculative Execution. *ArXiv e-prints* (Jan. 2018). arXiv:1801.01203
- [76] Paul Kocher, Joshua Jaffe, and Benjamin Jun. 1999. Differential Power Analysis. In *Advances in Cryptology — CRYPTO' 99*, Michael Wiener (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 388–397.
- [77] Paul C. Kocher. 1996. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Advances in Cryptology — CRYPTO '96*, Neal Koblitz (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 104–113.
- [78] Maxwell Krohn, Alexander Yip, Micah Brodsky, Natan Cliffer, M Frans Kaashoek, Eddie Kohler, and Robert Morris. 2007. Information flow control for standard OS abstractions. *ACM SIGOPS Operating Systems Review* 41, 6 (2007), 321–334.
- [79] Butler W. Lampson. 1973. A Note on the Confinement Problem. *Commun. ACM* 16, 10 (Oct. 1973), 613–615.
- [80] Butler W. Lampson. 1974. Protection. *ACM SIGOPS Operating Systems Review* 8, 1 (1974), 18–24.
- [81] T. Le, J. Di, M. Tehraniipoor, and L. Wang. 2016. Tracking data flow at gate-level through structural. In *2016 International Great Lakes Symposium on VLSI (GLSVLSI)*. 185–189.
- [82] J. Lee, I. Heo, Y. Lee, and Y. Paek. 2015. Efficient dynamic information flow tracking on a processor with core debug interface. In *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*. 1–6.
- [83] Jinyong Lee, Ingoo Heo, Yongje Lee, and Yunheung Paek. 2016. Efficient Security Monitoring with the Core Debug Interface in an Embedded Processor. *ACM Trans. Des. Autom. Electron. Syst.* 22, 1, Article 8 (May 2016), 29 pages.

- [84] Xun Li, Vineeth Kashyap, Jason K. Oberg, Mohit Tiwari, Vasanth Ram Rajarathinam, Ryan Kastner, Timothy Sherwood, Ben Hardekopf, and Frederic T. Chong. 2013. Position Paper: Sapper – a Language for Provable Hardware Policy Enforcement. In *Proceedings of the Eighth ACM SIGPLAN Workshop on Programming Languages and Analysis for Security* (Seattle, Washington, USA) (*PLAS '13*). ACM, New York, NY, USA, 39–44.
- [85] Xun Li, Mohit Tiwari, Ben Hardekopf, Timothy Sherwood, and Frederic T. Chong. 2010. Secure Information Flow Analysis for Hardware Design: Using the Right Abstraction for the Job. In *Proceedings of the 5th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security* (Toronto, Canada) (*PLAS '10*). ACM, New York, NY, USA, Article 8, 7 pages.
- [86] 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 *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation* (San Jose, California, USA) (*PLDI '11*). ACM, New York, NY, USA, 109–120.
- [87] X. Li, M. Tiwari, T. Sherwood, and F. T. Chong. 2010. Function flattening for lease-based, information-leak-free systems. In *ASAP 2010 - 21st IEEE International Conference on Application-specific Systems, Architectures and Processors*. 349–352.
- [88] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. 2018. Meltdown. *ArXiv e-prints* (Jan. 2018). arXiv:1801.01207
- [89] P. Lisherness and K. Cheng. 2009. An instrumented observability coverage method for system validation. In *2009 IEEE International High Level Design Validation and Test Workshop*. 88–93.
- [90] Jed Liu, Joe Corbett-Davies, Andrew Ferraiuolo, Alexander Ivanov, Mulong Luo, G. Edward Suh, Andrew C. Myers, and Mark Campbell. 2018. Secure Autonomous Cyber-Physical Systems Through Verifiable Information Flow Control. In *Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and Privacy* (Toronto, Canada) (*CPS-SPC '18*). ACM, New York, NY, USA, 48–59.
- [91] Juan Carlos Martinez Santos and Yunsi Fei. 2013. Micro-Architectural Support for Metadata Coherence in Multi-Core Dynamic Information Flow Tracking. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy* (Tel-Aviv, Israel) (*HASP '13*). ACM, New York, NY, USA, Article 6, 8 pages.
- [92] Juan Carlos Martinez Santos, Yunsi Fei, and Zhijie Jerry Shi. 2009. PIFT: Efficient Dynamic Information Flow Tracking Using Secure Page Allocation. In *Proceedings of the 4th Workshop on Embedded Systems Security* (Grenoble, France) (*WESS '09*). ACM, New York, NY, USA, Article 6, 8 pages.
- [93] Juan Carlos Martinez Santos, Yunsi Fei, and Zhijie Jerry Shi. 2012. Static Secure Page Allocation for Light-Weight Dynamic Information Flow Tracking. In *Proceedings of the 2012 International Conference on Compilers, Architectures and Synthesis for Embedded Systems* (Tampere, Finland) (*CASES '12*). ACM, New York, NY, USA, 27–36.
- [94] Jamil Mazzawi and Ziyad Hanna. 2013. Formal Analysis of Security Data Paths in RTL Design. In *Hardware and Software: Verification and Testing*, Armin Biere, Amir Nahir, and Tanja Vos (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 7–7.
- [95] Mentor Graphics. 2015. From Simulation to Emulation – A Fully Reusable UVM Framework. <https://www.mentor.com/products/fv/resources/overview/from-simulation-to-emulation-a-fully-reusable-uvvm-framework-0def891c-ab7a-453d-b079-2c99f584650e>
- [96] MITRE. 2020. Common Weakness Enumeration (CWE). <https://cwe.mitre.org/>.
- [97] Tortuga Logic. 2019. Radix-M hardware security platform for firmware security validation. <https://www.tortugalogic.com/radix-m/>.
- [98] Tortuga Logic. 2019. Radix-S Hardware Root of Trust security verification framework. <https://www.tortugalogic.com/radix-s/>.
- [99] US Department of Defense. 1985. *Department of Defense Trusted Computer System Evaluation Criteria*. Palgrave Macmillan UK, London, 1–129.
- [100] John McLean. 1992. Proving noninterference and functional correctness using traces. *Journal of Computer security* 1, 1 (1992), 37–57.
- [101] Andreas Moser, Christopher Kruegel, and Engin Kirda. 2007. Exploring multiple execution paths for malware analysis. In *2007 IEEE Symposium on Security and Privacy (SP'07)*. IEEE, 231–245.
- [102] D. Mu, W. Hu, B. Mao, and B. Ma. 2014. A bottom-up approach to verifiable embedded system information flow security. *IET Information Security* 8, 1 (Jan 2014), 12–17.
- [103] T. Murray, D. Matchuk, M. Brassil, P. Gammie, T. Bourke, S. Seefried, C. Lewis, X. Gao, and G. Klein. 2013. seL4: From General Purpose to a Proof of Information Flow Enforcement. In *2013 IEEE Symposium on Security and Privacy*. 415–429.
- [104] Andrew C Myers and Barbara Liskov. 1997. A decentralized model for information flow control. *ACM SIGOPS Operating Systems Review* 31, 5 (1997), 129–142.
- [105] Vijay Nagarajan, Ho-Seop Kim, Youfeng Wu, and Rajiv Gupta. 2008. Dynamic Information Flow Tracking on Multicores. In *Workshop on Interaction between Compilers and Computer Architectures*.
- [106] Adib Nahiyani, Jungmin Park, Miao He, Yousef Iskander, Farimah Farahmandi, Domenic Forte, and Mark Tehranipoor. 2020. SCRIPT: A CAD Framework for Power Side-Channel Vulnerability Assessment Using Information Flow Tracking and Pattern Generation. *ACM Trans. Des. Autom. Electron. Syst.* 25, 3, Article 26 (May 2020), 27 pages.
- [107] A. Nahiyani, M. Sadi, R. Vittal, G. Contreras, D. Forte, and M. Tehranipoor. 2017. Hardware trojan detection through information flow security verification. In *2017 IEEE International Test Conference (ITC)*. 1–10.
- [108] James Newsome and Dawn Xiaodong Song. 2005. Dynamic Taint Analysis for Automatic Detection, Analysis, and Signature-Generation of Exploits on Commodity Software.. In *NDSS*, Vol. 5. Citeseer, 3–4.
- [109] G. S. Nicholas, Y. Gui, and F. Saqib. 2020. A Survey and Analysis on SoC Platform Security in ARM, Intel and RISC-V Architecture. In *IEEE International Midwest Symposium on Circuits and Systems (MWSCAS)*. 718–721.
- [110] Edmund B. Nightingale, Daniel Peek, Peter M. Chen, and Jason Flinn. 2008. Parallelizing Security Checks on Commodity Hardware. In *Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems* (Seattle, WA, USA) (*ASPLOS XIII*).

- ACM, New York, NY, USA, 308–318.
- [111] J. Oberg, W. Hu, A. Irturk, M. Tiwari, T. Sherwood, and R. Kastner. 2010. Theoretical analysis of gate level information flow tracking. In *Design Automation Conference*. 244–247.
 - [112] Jason Oberg, Wei Hu, Ali Irturk, Mohit Tiwari, Timothy Sherwood, and Ryan Kastner. 2011. Information Flow Isolation in I2C and USB. In *Proceedings of the 48th Design Automation Conference (San Diego, California) (DAC '11)*. ACM, New York, NY, USA, 254–259.
 - [113] J. Oberg, S. Meiklejohn, T. Sherwood, and R. Kastner. 2014. Leveraging Gate-Level Properties to Identify Hardware Timing Channels. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 33, 9 (Sept 2014), 1288–1301.
 - [114] J. Oberg, T. Sherwood, and R. Kastner. 2013. Eliminating Timing Information Flows in a Mix-Trusted System-on-Chip. *IEEE Design Test* 30, 2 (April 2013), 55–62.
 - [115] M. Ozsoy, D. Ponomarev, N. Abu-Ghazaleh, and T. Suri. 2014. SIFT: Low-Complexity Energy-Efficient Information Flow Tracking on SMT Processors. *IEEE Trans. Comput.* 63, 2 (Feb 2014), 484–496.
 - [116] C. Palmiero, G. Di Guglielmo, L. Lavagno, and L. P. Carloni. 2018. Design and Implementation of a Dynamic Information Flow Tracking Architecture to Secure a RISC-V Core for IoT Applications. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*. 1–7.
 - [117] L. Piccolboni, G. Di Guglielmo, and L. P. Carloni. 2018. PAGURUS: Low-Overhead Dynamic Information Flow Tracking on Loosely Coupled Accelerators. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 37, 11 (Nov 2018), 2685–2696.
 - [118] P. Pieper, V. Herdt, D. Große, and R. Drechsler. 2020. Dynamic Information Flow Tracking for Embedded Binaries using SystemC-based Virtual Prototypes. In *ACM/IEEE Design Automation Conference (DAC)*. 1–6. <https://doi.org/10.1109/DAC18072.2020.9218494>
 - [119] C. Pilato, K. Wu, S. Garg, R. Karri, and F. Regazzoni. 2019. TaintHLS: High-Level Synthesis for Dynamic Information Flow Tracking. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 38, 5 (May 2019), 798–808.
 - [120] Joël Porquet and Simha Sethumadhavan. 2013. WHISK: An Uncore Architecture for Dynamic Information Flow Tracking in Heterogeneous Embedded SoCs. In *Proceedings of the Ninth IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (Montreal, Quebec, Canada) (CODES+ISSS '13)*. IEEE Press, Article 4, 9 pages.
 - [121] SecVerilog Project. 2015. SecVerilog: Verilog + Information Flow V1.0. <http://www.cs.cornell.edu/projects/secverilog/>.
 - [122] Maoyuan Qin, Wei Hu, Dejun Mu, and Yu Tai. 2018. Property Based Formal Security Verification for Hardware Trojan Detection. In *International Verification and Security Workshop (IVSW '18)*. IEEE, Costa Brava, Spain, 62–67.
 - [123] Maoyuan Qin, Wei Hu, Xinmu Wang, Dejun Mu, and Baolei Mao. 2019. Theorem proof based gate level information flow tracking for hardware security verification. *Computers & Security* 85 (2019), 225 – 239.
 - [124] Olatunji Ruwase, Phillip B. Gibbons, Todd C. Mowry, Vijaya Ramachandran, Shimin Chen, Michael Kozuch, and Michael Ryan. 2008. Parallelizing Dynamic Information Flow Tracking. In *Proceedings of the Twentieth Annual Symposium on Parallelism in Algorithms and Architectures (Munich, Germany) (SPAA '08)*. ACM, New York, NY, USA, 35–45.
 - [125] Harry J. Saal and Israel Gat. 1978. A Hardware Architecture for Controlling Information Flow. In *Proceedings of the 5th Annual Symposium on Computer Architecture (ISCA '78)*. ACM, New York, NY, USA, 73–77.
 - [126] A. Sabelfeld and A. C. Myers. 2003. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications* 21, 1 (Jan 2003), 5–19.
 - [127] H. Salmani, M. Tehranipoor, and R. Karri. 2013. On design vulnerability analysis and trust benchmarks development. In *2013 IEEE 31st International Conference on Computer Design (ICCD)*. 471–474.
 - [128] Jerome H Saltzer, David P Reed, and David D Clark. 1984. End-to-end arguments in system design. *Technology* 100 (1984), 0661.
 - [129] S. Sefton, T. Siddiqui, N. S. Amour, G. Stewart, and A. K. Kodi. 2018. GARUDA: Designing Energy-Efficient Hardware Monitors From High-Level Policies for Secure Information Flow. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 37, 11 (Nov 2018), 2509–2518.
 - [130] Jangseop Shin, Hongce Zhang, Jinyong Lee, Ingo Heo, Yu-Yuan Chen, Ruby Lee, and Yunheung Paek. 2016. A hardware-based technique for efficient implicit information flow tracking. In *2016 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 1–7.
 - [131] R. Shioya, D. Kim, K. Horio, M. Goshima, and S. Sakai. 2009. Low-Overhead Architecture for Security Tag. In *2009 15th IEEE Pacific Rim International Symposium on Dependable Computing*. 135–142.
 - [132] G. Shirley and F. Saqib. 2019. Information Flow Tracking in RISC-V. In *2019 IEEE 16th International Conference on Smart Cities: Improving Quality of Life Using ICT IoT and AI (HONET-ICT)*. 199–200.
 - [133] A. S. Siddiqui, G. Shirley, S. Bendre, G. Bhagwat, J. Plusquellic, and F. Saqib. 2019. Secure Design Flow of FPGA Based RISC-V Implementation. In *2019 IEEE 4th International Verification and Security Workshop (IVSW)*. 37–42.
 - [134] C. Song, H. Moon, M. Alam, I. Yun, B. Lee, T. Kim, W. Lee, and Y. Paek. 2016. HDFI: Hardware-Assisted Data-Flow Isolation. In *2016 IEEE Symposium on Security and Privacy (SP)*. 1–17.
 - [135] Pramod Subramanyan and Divya Arora. 2014. Formal Verification of Taint-Propagation Security Properties in a Commercial SoC Design. In *Proceedings of the Conference on Design, Automation & Test in Europe (Dresden, Germany) (DATE '14)*. European Design and Automation Association, Leuven, BEL, Article 313, 2 pages.
 - [136] P. Subramanyan, S. Malik, H. Khattri, A. Maiti, and J. Fung. 2016. Verifying information flow properties of firmware using symbolic execution. In *2016 Design, Automation Test in Europe Conference Exhibition (DATE)*. 337–342.
 - [137] G. Edward Suh, Jae W. Lee, David Zhang, and Srinivas Devadas. 2004. Secure Program Execution via Dynamic Information Flow Tracking. In *Proceedings of the 11th International Conference on Architectural Support for Programming Languages and Operating Systems (Boston, MA, USA)*

- (ASPLOS XI). ACM, New York, NY, USA, 85–96.
- [138] Y. Tai, W. Hu, L. Guo, B. Mao, and D. Mu. 2017. Gate Level Information Flow analysis for multi-valued logic system. In *2017 2nd International Conference on Image, Vision and Computing (ICIVC)*. 1102–1106.
 - [139] Y. Tai, W. Hu, D. Mu, B. Mao, and L. Zhang. 2018. Towards Quantified Data Analysis of Information Flow Tracking for Secure System Design. *IEEE Access* 6 (2018), 1822–1831.
 - [140] Y. Tai, W. Hu, L. Zhang, D. Mu, and R. Kastner. 2021. A multi-flow information flow tracking approach for proving quantitative hardware security properties. *Tsinghua Science and Technology* 26, 1 (2021), 62–71.
 - [141] M. Tiwari, X. Li, H. M. G. Wassel, F. T. Chong, and T. Sherwood. 2009. Execution leases: A hardware-supported mechanism for enforcing strong non-interference. In *2009 42nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 493–504.
 - [142] M. Tiwari, X. Li, H. M. G. Wassel, B. Mazloom, S. Mysore, F. T. Chong, and T. Sherwood. 2010. Gate-Level Information-Flow Tracking for Secure Architectures. *IEEE Micro* 30, 1 (Jan 2010), 92–100.
 - [143] M. Tiwari, J. K. Oberg, X. Li, J. Valamehr, T. Levin, B. Hardekopf, R. Kastner, F. T. Chong, and T. Sherwood. 2011. Crafting a usable microkernel, processor, and I/O system with strict and provable information flow security. In *2011 38th Annual International Symposium on Computer Architecture (ISCA)*. 189–199.
 - [144] Mohit Tiwari, Hassan M.G. Wassel, Bitu Mazloom, Shashidhar Mysore, Frederic T. Chong, and Timothy Sherwood. 2009. Complete Information Flow Tracking from the Gates Up. In *the 14th International Conference on Architectural Support for Programming Languages and Operating Systems* (Washington, DC, USA). 109–120.
 - [145] Daniel Townley, Khaled N. Khasawneh, Dmitry Ponomarev, Nael Abu-Ghazaleh, and Lei Yu. 2019. LATCH: A Locality-Aware Taint Checker. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (Columbus, OH, USA) (MICRO '52)*. ACM, New York, NY, USA, 969–982.
 - [146] N. Vachharajani, M. J. Bridges, J. Chang, R. Rangan, G. Ottoni, J. A. Blome, G. A. Reis, M. Vachharajani, and D. I. August. 2004. RIFLE: An Architectural Framework for User-Centric Information-Flow Security. In *37th International Symposium on Microarchitecture (MICRO-37'04)*. 243–254.
 - [147] G. Venkataramani, I. Doudalis, Y. Solihin, and M. Prvulovic. 2008. FlexiTaint: A programmable accelerator for dynamic taint propagation. In *2008 IEEE 14th International Symposium on High Performance Computer Architecture*. 173–184.
 - [148] Dennis Volpano, Cynthia Irvine, and Geoffrey Smith. 1996. A sound type system for secure flow analysis. *Journal of computer security* 4, 2-3 (1996), 167–187.
 - [149] M. A. Wahab, P. Cotret, M. N. Allah, G. Hiet, A. K. Biswas, V. LapÔtre, and G. Gogniat. 2018. A small and adaptive coprocessor for information flow tracking in ARM SoCs. In *2018 International Conference on ReConfigurable Computing and FPGAs (ReConFig)*. 1–8.
 - [150] Muhammad Abdul Wahab, Pascal Cotret, Mounir Nasr Allah, Guillaume Hiet, Vianney Lapotre, and Guy Gogniat. 2016. Towards a hardware-assisted information flow tracking ecosystem for ARM processors. In *26th International Conference on Field Programmable Logic and Applications, FPL 2016, Lausanne, Switzerland, August 29 - September 2, 2016*. 1–2.
 - [151] M. A. Wahab, P. Cotret, M. N. Allah, G. Hiet, V. LapÔtre, and G. Gogniat. 2017. ARMHEX: A hardware extension for DIFT on ARM-based SoCs. In *2017 27th International Conference on Field Programmable Logic and Applications (FPL)*. 1–7.
 - [152] C. Wang, Y. Cai, and Q. Zhou. 2018. HLIFT: A high-level information flow tracking method for detecting hardware Trojans. In *2018 23rd Asia and South Pacific Design Automation Conference (ASP-DAC)*. 727–732.
 - [153] Ofir Weisse, Jo Van Bulck, Marina Minkin, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Raoul Strackx, Thomas F. Wenisch, and Yuval Yarom. 2018. Foreshadow-NG: Breaking the Virtual Memory Abstraction with Transient Out-of-Order Execution. <https://foreshadowattack.eu/foreshadow-NG.pdf>
 - [154] William Wulf, Ellis Cohen, William Corwin, Anita Jones, Roy Levin, Charles Pierson, and Fred Pollack. 1974. Hydra: The kernel of a multiprocessor operating system. *Commun. ACM* 17, 6 (1974), 337–345.
 - [155] C. Yang, W. Liu, Y. Wang, Q. Tong, and L. Li. 2018. Interference-based VM Migration to Mitigate Cache-based Side-channel Attacks in Cloud. In *2018 IEEE 4th International Conference on Computer and Communications (ICCC)*. 1188–1192.
 - [156] K. Yang, M. Hicks, Q. Dong, T. Austin, and D. Sylvester. 2016. A2: Analog Malicious Hardware. In *2016 IEEE Symposium on Security and Privacy (SP)*. 18–37.
 - [157] Jiyong Yu, Mengjia Yan, Artem Khyzha, Adam Morrison, Josep Torrellas, and Christopher W. Fletcher. 2019. Speculative Taint Tracking (STT): A Comprehensive Protection for Speculatively Accessed Data. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (Columbus, OH, USA) (MICRO '52)*. ACM, New York, NY, USA, 954–968.
 - [158] Young-Nam Yun, Jae-Beom Kim, Nam-Do Kim, and Byeong Min. 2011. Beyond UVM for practical SoC verification. In *2011 International SoC Design Conference*. IEEE, 158–162.
 - [159] D. Zagieboylo, G. E. Suh, and A. C. Myers. 2019. Using Information Flow to Design an ISA that Controls Timing Channels. In *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)*. 272–27215.
 - [160] Nickolai Zeldovich, Silas Boyd-Wickizer, Eddie Kohler, and David Mazières. 2011. Making information flow explicit in HiStar. *Commun. ACM* 54, 11 (2011), 93–101.
 - [161] Nickolai Zeldovich, Silas Boyd-Wickizer, and David Mazières. 2008. Securing Distributed Systems with Information Flow Control. In *NSDI*, Vol. 8. 293–308.

- [162] Danfeng Zhang, Yao Wang, G. Edward Suh, and Andrew C. Myers. 2015. A Hardware Design Language for Timing-Sensitive Information-Flow Security. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems (Istanbul, Turkey) (ASPLOS '15)*. ACM, New York, NY, USA, 503–516.
- [163] R. Zhang and C. Sturton. 2020. Transys: Leveraging Common Security Properties Across Hardware Designs. In *IEEE Symposium on Security and Privacy (SP)*. 1713–1727.
- [164] Hu Zhaohui, Arnaud Pierres, Hu Shiqing, Chen Fang, Philippe Royannez, Eng Pek See, and Yean Ling Hoon. 2012. Practical and efficient SOC verification flow by reusing IP testcase and testbench. In *2012 International SoC Design Conference (ISOCC)*. IEEE, 175–178.