Viaduct

An Extensible, Optimizing Compiler for Secure Distributed Programs

Coşku Acay*
Cornell University
Ithaca, NY, USA
cacay@cs.cornell.edu

Rolph Recto*
Cornell University
Ithaca, NY, USA
rr729@cornell.edu

Joshua Gancher
Cornell University
Ithaca, NY, USA
jrg358@cornell.edu

Andrew C. Myers
Cornell University
Ithaca, NY, USA
andru@cs.cornell.edu

Elaine Shi
Cornell University
Ithaca, NY, USA
runting@gmail.com

Abstract

Modern distributed systems involve interactions between principals with limited trust, so cryptographic mechanisms are needed to protect confidentiality and integrity. At the same time, most developers lack the training to securely employ cryptography. We present Viaduct, a compiler that transforms high-level programs into secure, efficient distributed realizations. Viaduct's source language allows developers to declaratively specify security policies by annotating their programs with information flow labels. The compiler uses these labels to synthesize distributed programs that use cryptography efficiently while still defending the source-level security policy. The Viaduct approach is general, and can be easily extended with new security mechanisms.

Our implementation of the Viaduct compiler comes with an extensible runtime system that includes plug-in support for multiparty computation, commitments, and zero-knowledge proofs. We have evaluated the system on a set of benchmarks, and the results indicate that our approach is feasible and can use cryptography in efficient, nontrivial ways.

CCS Concepts: • Security and privacy → Information flow control; Cryptography; Domain-specific security and privacy architectures.

Keywords: information flow, multiparty computation, zero knowledge

*Equal contribution.

This work is licensed under a Creative Commons Attribution 4.0 International License.

PLDI '21, June 20–25, 2021, Virtual, Canada

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8391-2/21/06.

<https://doi.org/10.1145/3453483.3454074>

ACM Reference Format:

Coşku Acay, Rolph Recto, Joshua Gancher, Andrew C. Myers, and Elaine Shi. 2021. Viaduct: An Extensible, Optimizing Compiler for Secure Distributed Programs. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '21)*, June 20–25, 2021, Virtual, Canada. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3453483.3454074>

1 Introduction

Modern distributed applications such as federated systems and decentralized blockchains typically involve parties from multiple administrative domains each with its own security policy. Companies might be required by law (such as the European Union's GDPR [21]) to protect user privacy when they process user data or share it with other companies. The lack of full trust among parties makes it difficult to develop such systems, especially when the security requirements necessitate the use of cryptographic mechanisms. Recent efforts from the cryptography community have pushed these mechanisms from theory to practical deployment [5], but a gap remains: they still require too much expertise to use successfully [15, 17, 22].

We introduce Viaduct, a system that makes it easier for non-expert programmers to develop secure distributed programs that employ cryptography. It puts a variety of sophisticated cryptographic mechanisms in the hands of developers, including secure multiparty computation (MPC) protocols, zero-knowledge proofs (ZKP), and commitment schemes. Viaduct's *security-typed* language allows developers to annotate programs with information-flow labels to specify fine-grained security policies regarding the confidentiality and integrity of data and computation. An inference algorithm allows these annotations to be lightweight, and enables Viaduct to reject inherently insecure programs. Viaduct then enforces these policies by compiling high-level source code to secure distributed programs, automatically choosing efficient use of cryptography without sacrificing security. The compiler supports a range of cryptographic protocols whose security guarantees are characterized using information-flow

Figure 1. Architecture of Viaduct.

labels. New protocols can be added to Viaduct by specifying their security properties and by implementing well-defined interfaces.

Although prior efforts have attempted to bridge this gap, most existing work focuses on compiling programs to a fixed set of cryptographic mechanisms. For example, some focus on compiling programs to MPC (e.g., Wysteria [40], OblivM [33], SCALE-MAMBA [3]); others focus on ZKP (e.g., Pinocchio [37], Buffet [46], xjSNARK [31]). To our knowledge, by providing a unified abstraction to both specify security policies of programs and to specify security guarantees of cryptographic mechanisms, Viaduct is the first system to compile secure, distributed programs with an *extensible* suite of cryptography.

We make the following contributions:

- An algorithm to infer minimum consistent security requirements of data storage and computation for programs written in a security-typed language. (§3)
- A technique to compile secure distributed programs, deploying an extensible set of cryptographic protocols while minimizing a customizable notion of cost. (§4)
- An extensible runtime system for running compiled programs. Cryptographic mechanisms are added as plug-ins to the runtime. (§5, §6)
- An evaluation that shows that the Viaduct compiler can synthesize a wide variety of secure and efficient distributed programs, that the compilation technique is scalable, and that the annotation burden of the source language is minimal. (§7)
- An open-source implementation of the Viaduct compiler and runtime system.¹

¹Available at <https://github.com/apl-cornell/viaduct>.

```

1 host alice: {A ∧ B←}
2 host bob  : {B ∧ A←}
3
4 val a1, a2, a3 = input int alice
5 val b1, b2, b3 = input int bob
6 val a = min(a1, a2, a3)
7 val b = min(b1, b2, b3)
8 val b_richer = declassify a < b to {A □ B}
9 output b_richer to alice, bob

```

Figure 2. Implementation of the historical millionaires' problem in Viaduct. Viaduct uses MPC for the comparison $a < b$, but computes the minima locally.

```

1 host alice: {A}
2 host bob  : {B}
3
4 val n: {B ∧ A←} =
5   endorse (input int bob) from {B}
6 var tries: {A □ B} = 5
7 var win: {A □ B} = false
8 while (0 < tries ∧ !win) {
9   val guess =
10    declassify (input int alice) to {A □ B→}
11   val tguess: {A □ B} =
12    endorse guess from {A □ B→}
13   win = declassify (n == tguess) to {A □ B}
14   tries -= 1
15 }
16 output win to alice, bob

```

Figure 3. Guessing game, where Alice attempts to guess Bob's secret number. Viaduct uses zero-knowledge proofs so Alice learns nothing more than whether her guesses are correct. Most labels in this code can be inferred automatically.

2 Overview of Viaduct

Figure 1 gives a high-level overview of Viaduct. Its compiler takes a high-level source program partially annotated with information-flow labels. The compiler infers labels consistent with programmer-supplied annotations to determine security requirements for all program components. Then for each component the compiler selects a protocol that matches these requirements, guiding the selection with a cost model. The output is a secure and efficient distributed program, which hosts execute using the Viaduct runtime system. The Viaduct architecture has a small set of well-defined extension points, allowing developers to add support for new protocols with relative ease.

We give two examples to motivate and describe the Viaduct compilation process.

Historical Millionaires’ Problem. Our first example is a slightly modified version of the “millionaires’ problem” [47]. As in the classic formulation, two individuals, Alice and Bob, want to determine who has more money without revealing how much money they have to the other person. Rather than comparing their current wealth, in our “historical” variant Alice and Bob want to see who was richer at their *poorest*. Figure 2 shows an implementation of the historical millionaires’ problem in Viaduct. The program compares Alice’s lowest wealth with Bob’s, and outputs the answer (`b_richer`) to both Alice and Bob.

Viaduct programs must specify the *hosts* that participate in the program, along with the *authority* that each host has, as shown in lines 1–2. All security policies in Viaduct are represented using *security labels* (in blue), which are defined formally in §2.1. Security labels capture both *confidentiality* and *integrity*. For example, host `alice` is given label $A \wedge B^\leftarrow$. Here, $B^\leftarrow$ is the *integrity component* of B (similarly, $B^\rightarrow$ is the confidentiality component of B). This label means that Alice fully trusts host `alice` (with both confidentiality and integrity), while Bob trusts host `alice` to execute the program correctly, but does not trust the host with his secret data.

All variables and expressions in Viaduct carry a security label, which is derived from the possible flows of information in the program. The variables in lines 4–7 carry the same label as their respective hosts, since they only involve data local to that host. However, the comparison `a < b` involves *both* hosts’ private data, so has the higher security label $A \wedge B$. This label corresponds to data that is secret to and trusted by both principals. Since $A \wedge B$ corresponds to secret data, we require an explicit *declassification* to the label $A \sqcap B$, which describes data that both hosts can see and trust.

During protocol selection (§4), Viaduct chooses cryptographic protocols to securely and efficiently execute our example. The central idea that allows Viaduct to select protocols automatically is that the security guarantees of protocols can also be captured by labels. Neither Alice nor Bob alone has enough authority to be responsible for the comparison, so Viaduct generates the following distributed implementation: Alice and Bob compute their respective minima locally but perform the comparison `a < b` in semi-honest MPC. A semi-honest MPC protocol works here because the authority labels assigned to the hosts indicate that Alice and Bob trust each other’s hosts for integrity. Without that assumption, Viaduct is instead forced to select another protocol such as maliciously secure MPC.

There are typically multiple ways to assign protocols to a given program expression. For example, the computation of Alice’s minimum on line 6 could be securely performed in MPC, but since the computation requires the authority of Alice alone, it is cheaper yet still secure to do the computation locally on Alice’s machine. Using its cost estimator, Viaduct compiles the optimal program described above.

After protocol selection, Viaduct outputs a distributed program which captures the required cryptography to execute the source program. Hosts can execute this distributed program using Viaduct’s runtime system.

Guessing Game. Figure 3 presents a contrasting example. Here, Alice and Bob have security labels A and B respectively, modeling a *malicious* corruption scenario. Since they do not trust each other to execute the program correctly, semi-honest MPC is not applicable. Bob inputs a number n , and Alice has five attempts to guess the number. Since Bob’s input initially has label B , it must first be *endorsed* to the label $B \wedge A^\leftarrow$, raising integrity so that Bob cannot unilaterally modify the value. This endorsement requires a cryptographic mechanism to protect the integrity and secrecy of variable n throughout program execution.

Viaduct synthesizes a program in which Bob commits to n so that its value remains secret to Alice but Bob cannot later lie about the committed value. The statement `n == tguess` is computed by having Bob send a zero-knowledge proof (ZKP) to Alice, so that Alice can trust the outcome but learns no additional information. All other variables are replicated in plaintext across the two hosts.

These examples show that Viaduct is general, as it treats protocols such as MPC and ZKP uniformly.

2.1 Specifying Security Policies

In Viaduct, security policies capture a notion of authority. Policies are represented by *principals*, formulas composed of conjunctions and disjunctions over a set of base principals $\{A, B, C, \dots\}$ and two special principals $\mathbf{0}$ and $\mathbf{1}$. Principal $\mathbf{0}$ represents maximal authority and corresponds to the conjunction of all base principals; principal $\mathbf{1}$ represents minimal authority and corresponds to the disjunction of all base principals. We distinguish authority over *confidentiality* and over *integrity*. The security requirements of information are thus characterized by *labels* consisting of pairs $\langle p_c, p_i \rangle$ of two principals p_c and p_i , for confidentiality and integrity respectively.

A conjunction of principals $p_1 \wedge p_2$ represents combined authority. For confidentiality, this means the principal is allowed to read data that p_1 may read and also data that p_2 may read. For integrity, the conjunction may influence data that p_1 may influence, and also data p_2 may influence. A disjunction $p_1 \vee p_2$ corresponds to common authority, which may read or influence exactly the data that either p_1 and p_2 may individually.

Principals carry a natural partial order based on their authority. We write $p_1 \Rightarrow p_2$ to mean p_1 “acts for”, or is at least as trusted as, p_2 . This relation coincides with logical implication: for example, $p_1 \wedge p_2 \Rightarrow p_1$ and $p_1 \Rightarrow p_1 \vee p_2$.

It is convenient to have syntax that works over both components of labels simultaneously. So, we extend $\mathbf{0}$, $\mathbf{1}$, $\wedge$, $\vee$, and $\Rightarrow$ pointwise, and write one principal to mean that the

two components are the same. For example, the annotation $\{A\}$ denotes the label $\langle A, A \rangle$. To talk about confidentiality and integrity separately, we use projections, writing $\ell^\rightarrow$ for the confidentiality projection of ℓ and $\ell^\leftarrow$ for its integrity. Thus, $\{B \wedge A^\leftarrow\}$ expands to $\langle B, B \wedge A \rangle$, meaning Bob’s sole confidentiality and the combined integrity of Alice and Bob. These projections are defined formally as follows:

$$\langle p_c, p_i \rangle^\rightarrow \triangleq \langle p_c, 1 \rangle \quad \langle p_c, p_i \rangle^\leftarrow \triangleq \langle 1, p_i \rangle.$$

The *reflection* operator [48] swaps the two components:

$$\bar{X}(\langle p_c, p_i \rangle) \triangleq \langle p_i, p_c \rangle$$

Viaduct programs assign labels to hosts to indicate the amount of trust placed in them, but there are also labels on data. The important insight, borrowed from FLAM [4], is that the same set of labels can be used to talk about both authority and information flow. When placed on data, a label takes on an information flow interpretation, specifying the minimum authority required to read and influence that data. As in FLAM, standard operations from information flow literature can be reformulated in terms of authority:

$$\begin{aligned} \ell_1 \sqsubseteq \ell_2 &\iff \ell_2^\rightarrow \Rightarrow \ell_1^\rightarrow \quad \text{and} \quad \ell_1^\leftarrow \Rightarrow \ell_2^\leftarrow \quad (\text{flows to}) \\ \ell_1 \sqcup \ell_2 &\triangleq (\ell_1 \wedge \ell_2)^\rightarrow \wedge (\ell_1 \vee \ell_2)^\leftarrow \quad (\text{join}) \\ \ell_1 \sqcap \ell_2 &\triangleq (\ell_1 \vee \ell_2)^\rightarrow \wedge (\ell_1 \wedge \ell_2)^\leftarrow \quad (\text{meet}) \end{aligned}$$

The flows-to relation $\ell_1 \sqsubseteq \ell_2$ orders information flow policies: it means label ℓ_1 is more permissive about the use of information than ℓ_2 . The join $\ell_1 \sqcup \ell_2$ is more restrictive than both ℓ_1 and ℓ_2 , and the meet $\ell_1 \sqcap \ell_2$ is more permissive than either ℓ_1 or ℓ_2 . The most restrictive label—that of completely secret, untrusted data—is $\mathbf{0}^\rightarrow = \langle \mathbf{0}, \mathbf{1} \rangle$, and the least restrictive (public, trusted data) is $\mathbf{0}^\leftarrow = \langle \mathbf{1}, \mathbf{0} \rangle$.

2.2 Threat Model

Compiled programs run in a distributed setting in which each host executes a single thread concurrently with other hosts. Hosts communicate via message passing over secure, private, asynchronous channels. There is no shared memory that spans multiple hosts. We assume the attacker cannot observe wall-clock timing. Additionally, we are not concerned with availability, so the attacker can halt execution at any time.

In the setting of Viaduct, there is no single notion of an attacker. For example, in the historical millionaires problem, neither Alice nor Bob fully trust the other. To Alice, Bob is a potential attacker; Alice expects her security requirements to be met as long as the behavior of Bob’s (partially trusted) host is accurately described by the label assigned to it ($B \wedge A^\leftarrow$). Conversely, to Bob, Alice is a potential attacker. Hence, we are concerned with security versus all possible attackers.

We model the power of an attacker using a label. The attacker can read the data on a host if the confidentiality of the attacker label is at least as trusted as that of the host, and can change data and code on the host if the integrity of the attacker label is at least as trusted as that of the host. We do

not consider unreasonable attack scenarios in which a host has compromised integrity but still enforces confidentiality.²

For example, in the historical millionaires’ problem, there are five interesting corruption scenarios: no corrupted hosts; alice has corrupted confidentiality; bob has corrupted confidentiality; both have corrupted confidentiality; or both alice and bob are fully corrupted. The full corruption of a single host is not possible because the hosts trust each other, so if the integrity of one is corrupted then the other’s integrity must be corrupted also.

2.3 Label Inference

Viaduct selects a protocol for every piece of data and computation in the program based on their authority requirements, represented as labels. Intuitively, program components must be executed by protocols with enough authority to defend the confidentiality of host inputs and the integrity of host outputs. These authority requirements are captured formally by a type system (§3.1), and Viaduct uses a novel inference algorithm (§3.2) to compute for all program components the minimum-authority labels that still respect the information-flow constraints on the program.

The only required label annotations on Viaduct programs are the authority labels on host declarations and labels on declassify/endorse expressions—all labels on variables can be elided, making annotation burden low. As we show in our evaluation, in practice these required annotations are enough to capture programmer intent: minimally annotated programs compile to the same distributed programs as their fully annotated versions.

2.4 Protocol Selection

After label inference, Viaduct performs *protocol selection*, which assigns a protocol to compute and store each subexpression and variable. Protocols encompass storage and computation performed “in the clear” as well as cryptographic mechanisms such as commitments, MPC and zero-knowledge proofs.

Each protocol P carries an associated authority label $\mathbb{L}(P)$, which approximates the security guarantees the protocol provides. Given a program component with minimum authority requirement ℓ , protocol selection only assigns P to execute that component if $\mathbb{L}(P) \Rightarrow \ell$ —that is, if P meets the authority requirement for the program component.

Intuitively, given a program s and protocol P , we may imagine an *ideal functionality* P^s (in the style of UC [9]) which executes the program fragments of s that are assigned to P . The fragments of s that are assigned to P may depend on the computational abilities of P . For example, if P is a commitment protocol, then P^s is only able to store values

²The semi-honest and malicious threat models common in cryptography correspond to corrupting only hosts’ confidentiality and corrupting both hosts’ confidentiality and integrity respectively.

Protocol	Authority label
Local(h)	$\mathbb{L}(h)$
Replicated(H)	$\prod_{h \in H} \mathbb{L}(h)$
Commitment(h_p, h_v)	$\mathbb{L}(h_p) \wedge \mathbb{L}(h_v)^\leftarrow$
ZKP(h_p, h_v)	$\mathbb{L}(h_p) \wedge \mathbb{L}(h_v)^\leftarrow$
MAL-MPC(H)	$\bigwedge_{h \in H} \mathbb{L}(h)$
SH-MPC(H)	let $I = \bigvee_{h \in H} \mathbb{L}(h)^\leftarrow$ $(\mathbb{X}(I) \vee \bigwedge_{h \in H} \mathbb{L}(h)^\rightarrow) \wedge I$

Figure 4. Example protocols and security labels that represent their authority.

but not perform any computations. If P is an MPC protocol, then P^s can execute computations that can be translated into circuits—the standard interface for MPC implementations.

P^s guarantees that the storage and computation it performs are protected at label $\mathbb{L}(P)$. In particular, the adversary cannot observe storage or computation performed by P^s unless its confidentiality is at least $\mathbb{L}(P)$; dually, the adversary cannot influence storage or computation performed by P^s unless its integrity is at least $\mathbb{L}(P)$.

Examples of protocols and their corresponding authority labels are given in Figure 4. Following the above intuition for the security of functionalities P^s , the authority label of protocols are determined to be the least authority required of the adversary to corrupt the protocol (in confidentiality or integrity). We explain the example protocols below:

Local(h). No cryptography is performed, and data is stored and computations performed on host h in the clear. It provides exactly the authority of h .

Replicated(H). Data and computations are replicated on all hosts in set H , and replicated data is checked for equality when necessary. This protocol provides confidentiality $\bigvee_{h \in H} \mathbb{L}(h)^\rightarrow$ since all hosts hold the plaintext value. It provides integrity $\bigwedge_{h \in H} \mathbb{L}(h)^\leftarrow$ since all hosts must corrupt their local values for the value to be globally corrupted. Together, these labels form the label $\prod_{h \in H} \mathbb{L}(h)$.

Commitment(h_p, h_v). Data is stored on h_p and commitments are placed on h_v . Commitments are computationally inexpensive but usually no computations can be performed with them. Commitments increase integrity without sacrificing confidentiality. Its confidentiality is $\mathbb{L}(h_p)^\rightarrow$ since only h_p holds the plaintext value, while h_v only holds a commitment. Its integrity is $(\mathbb{L}(h_p) \wedge \mathbb{L}(h_v))^\leftarrow$ for the same reason as for replication.

ZKP(h_p, h_v). A zero-knowledge proof protocol where h_p is the prover and h_v is the verifier. The prover computes over its private data and sends the result to the verifier, along with a *proof* that attests the value computed is correct. The proof reveals nothing about the private data except what can be gleaned from the result itself. Zero-knowledge proofs provide the same authority as commitments, for essentially

the same reason: the prover holds all secret information and performs all computation, while the verifier only holds information which allows it to believe in the correctness of the result, but nothing more.

MAL-MPC(H). A corrupt-majority, maliciously secure multiparty computation protocol [8, 10, 24] performed by hosts H . The protocol allows hosts to jointly perform a computation over their private inputs, keeping these inputs secret to the other hosts and revealing only the result. The label $\bigwedge_{h \in H} \mathbb{L}(h)$ reflects that the confidentiality (resp., integrity) of data computed in MPC is compromised only if *all* participating hosts have compromised confidentiality (resp. integrity).

SH-MPC(H). A corrupt-majority, semi-honest secure multiparty computation protocol performed by hosts H . While the combined authority label is complex, its confidentiality and integrity projections are easy to understand. The integrity is equal to $\bigvee_{h \in H} \mathbb{L}(h)^\leftarrow$, since the integrity of the MPC computation may be compromised if *any* host behaves maliciously. The confidentiality is equal to

$$\left(\bigvee_{h \in H} \mathbb{X}(\mathbb{L}(h)^\leftarrow) \right) \vee \left(\bigwedge_{h \in H} \mathbb{L}(h)^\rightarrow \right).$$

The first disjunct captures the fact that confidentiality guarantees are discarded if the integrity of any host is compromised. The second disjunct states that, if all hosts follow the protocol correctly, the adversary can only learn the state of intermediate MPC computations if all hosts have corrupted confidentiality. Overall, this means that in order to compromise confidentiality guarantees of semi-honest MPC, either the integrity of any host or the confidentiality of all hosts must be compromised.

In particular, for the historical millionaires' example, the label of SH-MPC(alice, bob) is $A \wedge B$. This is because hosts alice and bob are both assumed to have the high integrity of $(A \wedge B)^\leftarrow$. If alice and bob only have their own integrity, however, then the label is computed to be $A \vee B$. The protocol only has enough authority to perform computations over data public to both hosts, and neither host trusts the result. Indeed, semi-honest MPC offers little to no benefit if any host has lower integrity than any other.

2.5 Runtime

Viaduct provides a modular runtime system for executing compiled distributed programs, implemented as an interpreter. All hosts run the interpreter with the same compiled program, which then executes each host's portion of the program. During execution, the interpreter calls out to back ends implementing the cryptographic mechanisms used in the program. Back ends translate computations in the source language into their cryptographic realizations. For instance, the back ends for MPC and ZKP in our implementation build a circuit representation of the program as it executes.

Figure 5. Execution of the compiled distributed program for the historical millionaires’ problem using a cleartext back end and an MPC back end. Sends and receives are over protocol–host pairs (P, h) . These messages are processed by the back end for protocol P at host h .

Protocol back ends can send data to and receive data from each other, supporting the composition of protocols. Source-level declassification and endorsement induce this communication. For example, in Figure 2 on line 8, the computation $a < b$ is declassified from label $A \wedge B$ to $A \sqcap B$. This declassification causes the MPC protocol between Alice and Bob to execute its stored circuit for this comparison, and to output the result in cleartext.

Figure 5 shows the execution of the program compiled by Viaduct for the historical millionaires’ problem. The program runs as follows. (1) First, the cleartext back ends on Alice and Bob’s machines receive input locally and compute their respective minima. The back ends send the minima as secret inputs to their respective MPC back ends, which create input gates for these inputs. (2) Next, the MPC back ends on Alice and Bob’s machines each create an operation gate that compares Alice and Bob’s secret inputs. The back ends jointly execute the circuit with the comparison result as output, which they send to their respective cleartext back ends. (3) Finally, the cleartext back ends on Alice and Bob’s machines both receive from their MPC back ends and output the result.

3 Source Language

The syntax for Viaduct’s source language, a simplified version of the *surface* language, is given in Figure 6. The language supports base types such as booleans and integers, along with their usual operators. Surface-level assignables (**val** and **var** declarations) and arrays are uniformly represented as *data types*, a restricted form of objects. Like regular objects, they are created using constructors (**new** declarations) and contain methods. For simplicity, we only include three data types: immutable/mutable cells, which model surface-level assignables, and arrays. Arrays are dynamically sized but statically allocated: the size of an array can depend on values known only at run time, but array references cannot be rebound to different names or stored in arrays.

We distinguish between fully evaluated atomic expressions a , and expressions e that evaluate to values and may have side effects. Methods include **get** and **set** operations

Temporaries	t	$\in \mathbb{T}$
Assignables	x	$\in \mathbb{X}$
Hosts	h	$\in \mathbb{H}$
Labels	ℓ	$\in \mathbb{L}$
Base Types	β	$::= \text{unit} \mid \text{bool} \mid \text{int}$
Data Types	D	$::= \text{Cell}_\beta \mid \text{Array}_\beta$
Values	v	$::= () \mid \text{true} \mid \text{false} \mid i \in \mathbb{Z}$
Unary Operators	op_1	$::= \text{not} \mid - \mid \dots$
Binary Operators	op_2	$::= \wedge \mid \vee \mid + \mid \times \mid = \mid \dots$
Methods	m	$::= \text{get} \mid \text{set} \mid \dots$
Atomic Expr.	a	$::= v \mid t$
Expressions	e	$::=$ $\mid a \mid op_n(a_1, \dots, a_n) \mid x.m(a_1, \dots, a_n)$ $\mid \text{declassify } a \text{ to } \ell \mid \text{endorse } a \text{ from } \ell$ $\mid \text{input}_\beta h \mid \text{output } a \text{ to } h$
Statements	s	$::=$ $\mid \text{let } t = e \text{ in } s \mid \text{new } x = D(a_1, \dots, a_n) \text{ in } s$ $\mid \text{if } a \text{ then } s_1 \text{ else } s_2 \mid b : \text{loop } s \mid \text{break } b$ $\mid s_1; s_2 \mid \text{skip}$

Figure 6. Abstract syntax of Viaduct’s source language

for both mutable cells and arrays (for which they take an index as an extra argument). Input/output expressions allow programs to interact with hosts. The **declassify** expression marks locations where private data is explicitly allowed to flow to public data, while the **endorse** expression marks locations where untrusted data is explicitly allowed to influence trusted data.

Statements consist of let-bindings, assignable declarations, as well as the usual conditionals, loops, and sequential composition. Temporaries bind values while assignables bind instances of data types. We require all intermediate computations to be let-bound by a temporary, enforcing a variant of *A-normal form* [18]. We use the more general loop-until-break statements instead of the more traditional while loops, simplifying the conversion to A-normal form. A break statement (**break** b) includes an identifier b that names the loop it breaks out of.

3.1 Label Checking

Viaduct's type system enforces secure information flow in a standard way. The type system serves two purposes. First, it helps programmers ensure there are no unintended information flows: secrets are not leaked to and data is not corrupted by unauthorized principals. Second, it specifies what labels can be assigned to variables and expressions that the user did not explicitly annotate.

Figure 7 presents label checking rules for expressions and selected statements. Expressions are checked by the judgment $\Gamma; pc \vdash e : \ell$, which means that e has label ℓ under the context on the left. Here, Γ is a finite partial map from temporaries, assignables, or loop names to labels:

Label Contexts $\Gamma ::= \cdot \mid \Gamma, t : \ell \mid \Gamma, x : \ell \mid \Gamma, b : \ell$

The *program counter* label pc is a standard way to prevent implicit flows of information via control flow [43]. The rules for method calls and input/output expressions differ from those in standard security-typed languages in that they also include premises with pc checks. These checks are required because these expressions may induce communication between hosts, and hosts may learn secrets based on which requests they receive. Prior work that targets the distributed setting contains similar checks to control *read channels* [51].

Statement checking rules have the form $\Gamma; pc \vdash s$; they are largely standard [43]. Because we assume attackers cannot observe timing nor analyze traffic, the rule for conditional statements does not require branches to have the same timing behavior or effects (e.g., method calls, input/output).

Nonmalleable Information Flow Control. Information flow type systems typically aim to enforce a compositional security property such as *noninterference* [23]. Noninterference is a strong property but it is too restrictive for practical applications, which usually have a more nuanced policy for secure information flow. Hence, like most languages supporting information flow control (e.g., [6, 36, 39]), Viaduct allows programmers to signify the exceptions to a noninterference policy through *downgrading* expressions.

Downgrading enables information flows that would violate noninterference, so it can be dangerous. This is especially true in the distributed setting, where storage and computation can be performed by hosts that one does not fully trust. Downgrading confidentiality (declassification) allows secret information to be treated as public information—a necessity for many applications, but doing so might allow a corrupted host to control when information is released or what information is released. Downgrading integrity (endorsement) allows untrusted information to be treated as trusted information, but might enable a corrupted host to trick an honest one into accepting mauled secrets.

The property of *nonmalleable information flow control* (NMIFC) [11] prevents both of these abuses of downgrading by combining two properties: *robust declassification* [50]

and *transparent endorsement* [11]. Robust declassification requires that principals to which data is declassified could not have influenced either the decision to declassify or the data itself. Meanwhile, transparent endorsement prevents trusting mauled secrets by ensuring that information can only be endorsed if the providing principal can read it.

The declassification and endorsement rules in Figure 7 enforce NMIFC using the reflection operator $\bar{X}$ (§2.1). The rules prevent the program from downgrading information with *compromised labels* [48], in which confidentiality exceeds integrity. These rules generate authority requirements that prevent the Viaduct compiler from placing data and computation on insufficiently trustworthy hosts. For example, consider a program where a server releases secret information to a client when the client guesses the correct password:

```
host server: {S}, client: {1}
val info: int{S}, pw: int{S}, guess: int{1}
if (declassify (pw == guess) to {1})
  output (declassify info to {1}) to client
```

This program violates robust declassification, because the decision to declassify `info` depends on (low-integrity) `guess`. Without the restrictions on downgrading, Viaduct could compile the program to store the guard `pw == guess` (with label 1) on the *client*. The client could simply claim to the server that its guess is correct! For this program to type-check with NMIFC, endorsement is needed to make the guard high-integrity. A naive programmer might think to endorse the entire guard, but this (nontransparent) endorsement could still be compiled in a way that lets an untrusted host supply its value. The correct solution is to explicitly endorse `guess` before declassifying the comparison; since `guess` is not secret, the endorsement is transparent. The resulting labels correctly force Viaduct to put the comparison on the server.

3.2 Label Inference

Checking secure information flow is not enough; for protocol selection, the compiler also needs the labels of all expressions. We present an algorithm to infer these labels.

As in prior work on inferring information flow labels [36, 39], information flow checking reduces to a system of flows-to ($\sqsubseteq$) constraints over label constants and label variables. Type inference collects these premises from Figure 7, and generates fresh label variables for labels that appear in a premise of a rule but not its conclusion (e.g., pc' in the rule for if statements). The inference algorithm finds a label-variable assignment that satisfies all the constraints, if possible.

The algorithm computes the *minimum-authority* solution, the choice of labels requiring the least amount of confidentiality and integrity for each component. Minimum-authority labels are desirable because higher authority is achieved only through more trust or costly cryptography.

First, we translate the flows-to ($\sqsubseteq$) constraints over *labels*, which appear in rule premises, to acts-for ($\Rightarrow$) constraints

$$\begin{array}{c}
\boxed{\Gamma \vdash a : \ell} \quad \boxed{\Gamma; pc \vdash e : \ell} \quad \frac{}{\Gamma \vdash v : \ell} \quad \frac{\Gamma(t) = \ell_t \quad \ell_t \sqsubseteq \ell}{\Gamma \vdash t : \ell} \quad \frac{\Gamma \vdash a_i : \ell}{\Gamma; pc \vdash op_n(a_1, \dots, a_n) : \ell} \\
\\
\frac{\Gamma(x) = \ell_x \quad pc \sqsubseteq \ell_x \quad \Gamma \vdash a_i : \ell_x \quad \ell_x \sqsubseteq \ell}{\Gamma; pc \vdash x.m(a_1, \dots, a_n) : \ell} \quad \frac{pc \sqsubseteq \ell_t \quad \Gamma \vdash a : \ell_t \quad \ell_f^{\leftarrow} = \ell_t^{\leftarrow} \quad \ell_f^{\rightarrow} \sqsubseteq \ell_t^{\rightarrow} \sqcup \mathbb{X}(\ell_f^{\leftarrow}) \quad \ell_t \sqsubseteq \ell}{\Gamma; pc \vdash \text{declassify } a \text{ to } \ell_t : \ell} \quad \frac{pc \sqsubseteq \ell_t \quad \Gamma \vdash a : \ell_t \quad \ell_f^{\rightarrow} = \ell_t^{\rightarrow} \quad \ell_f^{\leftarrow} \sqsubseteq \ell_t^{\leftarrow} \sqcup \mathbb{X}(\ell_f^{\rightarrow}) \quad \ell_t \sqsubseteq \ell}{\Gamma; pc \vdash \text{endorse } a \text{ from } \ell_t : \ell} \\
\\
\frac{pc \sqsubseteq \mathbb{L}(h) \quad \mathbb{L}(h) \sqsubseteq \ell}{\Gamma; pc \vdash \text{input}_\beta h : \ell} \quad \frac{pc \sqsubseteq \mathbb{L}(h) \quad \Gamma \vdash a : \mathbb{L}(h)}{\Gamma; pc \vdash \text{output } a \text{ to } h : \ell} \\
\\
\boxed{\Gamma; pc \vdash s} \quad \frac{\Gamma; pc \vdash e : \ell \quad pc \sqsubseteq \ell \quad (\Gamma, t : \ell); pc \vdash s}{\Gamma; pc \vdash \text{let } t = e \text{ in } s} \quad \frac{\Gamma \vdash a_i : \ell \quad pc \sqsubseteq \ell \quad (\Gamma, x : \ell); pc \vdash s}{\Gamma; pc \vdash \text{new } x = D(a_1, \dots, a_n) \text{ in } s} \quad \frac{pc \sqsubseteq pc' \quad \Gamma \vdash a : pc' \quad \Gamma; pc' \vdash s_1 \quad \Gamma; pc' \vdash s_2}{\Gamma; pc \vdash \text{if } a \text{ then } s_1 \text{ else } s_2}
\end{array}$$

Figure 7. Selected information flow checking rules for expressions and statements.

$$\begin{aligned}
\ell_1 \sqsubseteq \ell_2 &\rightsquigarrow C(\ell_2) \Rightarrow C(\ell_1), I(\ell_1) \Rightarrow I(\ell_2) \\
\ell_f^{\rightarrow} \sqsubseteq \ell_t^{\rightarrow} \sqcup \mathbb{X}(\ell_f^{\leftarrow}) &\rightsquigarrow I(\ell_f) \wedge C(\ell_t) \Rightarrow C(\ell_f) \\
\ell_f^{\leftarrow} \sqsubseteq \ell_t^{\leftarrow} \sqcup \mathbb{X}(\ell_f^{\rightarrow}) &\rightsquigarrow I(\ell_f) \Rightarrow C(\ell_t) \vee I(\ell_t)
\end{aligned}$$

Figure 8. Translating flows-to constraints over labels to acts-for constraints over label components.

Constraint	Update rule
$L_1 \Rightarrow L_2$	$L_1^{i+1} := L_1^i \wedge L_2^i$
$L_1 \wedge p_2 \Rightarrow L_3$	$L_1^{i+1} := L_1^i \wedge (p_2 \rightarrow L_3^i)$
$L_1 \Rightarrow L_2 \vee L_3$	$L_1^{i+1} := L_1^i \wedge (L_2^i \vee L_3^i)$

Figure 9. Update rules for solving acts-for constraints.

over the underlying *label components* as shown in Figure 8. Here, $C(\ell)$ and $I(\ell)$ are functions that project the confidentiality and integrity components, respectively, of label ℓ . These components are constants p when the label is known, and variables L otherwise.

We then adapt the algorithm of Rehof and Mogensen [41] for iteratively solving semilattice constraints. All principal variables are initialized to **1** and unsatisfied constraints are used to update variables repeatedly, until a fixed point is reached, according to the rules in Figure 9. Constraints of the form $L_1 \Rightarrow L_2$ or $L_1 \Rightarrow L_2 \vee L_3$ are used to perform the corresponding update.

However, the rules in Figure 8 can also generate constraints of the form $L_1 \wedge p_2 \Rightarrow L_3$, arising from the typing rule for robust declassification. The term p_2 is always a constant since Viaduct requires annotations on declassify operations, so the value of L_1 can be updated safely to $p_2 \rightarrow L_3$, which denotes the weakest authority p such that $p \wedge p_2 \Rightarrow L_3$. When a lattice supports the $\rightarrow$ operation, it is a *Heyting*

algebra [42], allowing each update rule to lower the left-hand-side variable to the minimum authority satisfying the constraint. Any free distributive lattice, such as our lattice of principals, is a Heyting algebra. We prove this fact, as well as the fact that iterative analysis always terminates with the minimum-authority solution, in the supplemental technical report [2].

4 Protocol Selection

The protocol selection phase of Viaduct assigns a protocol to each program component. Formally, a *protocol assignment* is a function $\Pi : (\mathbb{T} \cup \mathbb{X}) \rightarrow \mathbb{P}$ from temporaries and assignables to protocols. For a temporary t , $\Pi(t)$ is the protocol that executes the expression associated with t . Similarly, $\Pi(x)$ is the protocol that stores and responds to method calls on the data type instance bound to x .

4.1 Validity of Protocol Assignments

Figure 10 outlines the conditions under which a protocol assignment is valid. The judgement $\Pi \models e : P$ means that expression e can be executed by protocol P under assignment Π . Similarly, the judgement $\Pi \models s$ means that Π is a valid assignment for statement s .

We now describe the rules for validity. The rule for temporaries states that t can only be read by protocol P if $\Pi(t)$, the protocol storing t , can communicate with P , written $\text{comm}(\Pi(t), P)$. Not all pairs of protocols can communicate; the customizable *protocol composer*, discussed further in §5.1, defines the valid set of protocol compositions.

Other rules restrict where certain expressions can be executed. A method call on x must be executed by $\Pi(x)$, the protocol that stores x . Similarly, input/output expressions must be executed locally on the relevant host.

The rules for temporary and assignable declarations ensure that the protocol selected for a temporary or assignable

$$\begin{array}{c}
\boxed{\Pi \models e : P} \qquad \boxed{\Pi \models s} \\
\\
\frac{\text{comm}(\Pi(t), P)}{\Pi \models t : P} \quad \frac{\Pi \models a_i : \Pi(x)}{\Pi \models x.m(a_1, \dots, a_n) : \Pi(x)} \quad \frac{}{\Pi \models \text{input}_\beta h : \text{Local}(h)} \quad \frac{\Pi \models a : \text{Local}(h)}{\Pi \models \text{output } a \text{ to } h : \text{Local}(h)} \\
\\
\frac{\mathbb{L}(\Pi(t)) \Rightarrow \mathbb{L}(t) \quad \Pi \models s}{\Pi \models \text{let } t = e \text{ in } s} \quad \frac{\mathbb{L}(\Pi(x)) \Rightarrow \mathbb{L}(x) \quad \Pi \models a_i : \Pi(x) \quad \Pi \models s}{\Pi \models \text{new } x = D(a_1, \dots, a_n) \text{ in } s} \quad \frac{H = \text{hosts}(\Pi, s_1) \cup \text{hosts}(\Pi, s_2) \quad \forall h \in H. \mathbb{L}(a) \rightarrow \sqsubseteq \mathbb{L}(h) \rightarrow \quad \forall h \in H. \Pi \models a : \text{Local}(h)}{\Pi \models s_1 \quad \Pi \models s_2} \\
\frac{}{\Pi \models \text{if } a \text{ then } s_1 \text{ else } s_2}
\end{array}$$

Figure 10. Selected rules for the validity of a protocol assignment.

$$\begin{array}{c}
\boxed{\Pi(s) : 2^{\mathbb{P}}} \quad \boxed{\text{hosts}(\Pi, s) : 2^{\mathbb{H}}} \\
\\
\Pi(\text{let } t = e \text{ in } s) = \Pi(t) \cup \Pi(s) \\
\Pi(\text{new } x = D(a_1, \dots, a_n) \text{ in } s) = \Pi(x) \cup \Pi(s) \\
\Pi(\text{if } a \text{ then } s_1 \text{ else } s_2) = \Pi(s_1) \cup \Pi(s_2) \\
\Pi(b : \text{loop } s) = \Pi(s) \\
\Pi(\text{break } b) = \Pi(b : \text{loop } s) \\
\\
\text{hosts}(\Pi, s) = \bigcup_{P \in \Pi(s)} \text{hosts}(P)
\end{array}$$

Figure 11. Protocols and hosts involved in the execution of a statement. Here, $\text{hosts}(P)$ is the set of hosts that protocol P runs on, which is specified individually for each protocol.

$$\begin{aligned}
\text{cost}(\Pi, \text{let } t = e \text{ in } s) &= \\
&\quad c_{\text{exec}}(\Pi(t), e) + \sum_{P \in \text{readers}(\Pi, t, s)} c_{\text{comm}}(\Pi(t), P) + \text{cost}(\Pi, s) \\
\text{cost}(\Pi, \text{if } a \text{ then } s_1 \text{ else } s_2) &= \max(\text{cost}(\Pi, s_1), \text{cost}(\Pi, s_2)) \\
\text{cost}(\Pi, b : \text{loop } s) &= W_{\text{loop}} \times \text{cost}(\Pi, s) \\
\text{cost}(\Pi, s_1; s_2) &= \text{cost}(\Pi, s_1) + \text{cost}(\Pi, s_2) \\
\text{cost}(\Pi, s) &= 0 \text{ otherwise}
\end{aligned}$$

Figure 12. Abstract cost model.

has enough authority to securely store it. Formally, the label $\mathbb{L}(\Pi(t))$ of the protocol storing temporary t must act for $(\Rightarrow)$ the minimum required authority label $\mathbb{L}(t)$ computed for t in §3.2 (and similarly for assignables). Labels $\mathbb{L}(\Pi(t))$ are the ones explained in Figure 4.

The rule for conditional statements ensures that all hosts involved in the execution of a conditional statement (Figure 11) can learn which branch is taken. The first premise requires that involved hosts have enough confidentiality to read the value of the conditional guard, while the second premise ensures that the protocol computing the value of the

guard can forward it to all involved hosts. Both premises are trivially satisfied when the guard is a constant expression.

Where necessary, the Viaduct compiler removes these guard visibility constraints by multiplexing [34] conditional statements into straight-line code. This allows, for example, the compilation of conditionals with secret guards that require execution in MPC.

4.2 Cost of Protocol Assignments

There can be many valid protocol assignments that securely realize a source program. To select an optimal assignment, Viaduct attributes a cost to each assignment using an abstract cost model, shown in Figure 12. Developers can instantiate the abstract model by modifying the customizable *cost estimator*, which specifies $c_{\text{exec}}(P, s)$, the cost of executing statement s in protocol P ; $c_{\text{comm}}(P_1, P_2)$, the cost of communicating between P_1 and P_2 ; and the global constant W_{loop} , the number of times a loop is assumed to execute when its iteration count is not statically known.

Our implementation configures c_{exec} to assign a small cost to executing “in the clear” and a large cost to the use of cryptography, so the compiler avoids the use of cryptography except when required for security. We also configure the communication cost c_{comm} to minimize data movement. For example, a frequently accessed public variable would be replicated on two hosts so that each host has a local copy. Placing the variable only on one of the hosts could reduce storage cost but entails frequently sending its value to the other host.

4.3 Computing an Optimal Protocol Assignment

To compute an optimal protocol assignment given a program s , the Viaduct compiler constructs a constrained optimization problem over the following sets of variables:

- *Assignment variables* (α_i). These represent the protocols that execute let-bindings or declarations.
- *Cost variables* (β_i). These represent the cost of executing let-bindings or declarations.

- *Participating host variables* ($\gamma_{i,j}$). These are true if host j is participating in the execution of a statement i .

The compiler generates a set of constraints $\{\phi_1, \dots, \phi_n\}$ over these assignment, cost, and participating host variables, as well as an expression β_s capturing the cost of s as in Figure 12. These constraints are drawn from a grammar consisting of logical connectives, an equality predicate between assignment variables and protocols, and an equality predicate between cost variables and cost expressions. The compiler uses an off-the-shelf solver to find a solution for assignment variables α_i and participating host variables $\gamma_{i,j}$ such that all constraints $\{\phi_1, \dots, \phi_n\}$ are satisfied and β_s is minimized. Given the set of valid protocol assignments VA for s such that $VA = \{\Pi \mid \Pi \models s\}$, this solution for the assignment variables corresponds to a protocol assignment Π_{opt} such that

$$\Pi_{\text{opt}} = \arg \min_{\Pi \in VA} \text{cost}(\Pi, s).$$

Protocol Factory. To construct the optimization problem, the compiler draws the set of available protocols from the customizable *protocol factory*. Developers wishing to add new protocols to Viaduct must extend the protocol factory so that the compiler can generate assignments with these protocols during protocol selection.

The protocol factory defines a function $\text{viable} : \mathbb{T} \cup \mathbb{X} \rightarrow 2^{\mathbb{P}}$ that returns a set of viable protocols that can execute a let-binding or declaration. This allows developers to specify limitations regarding the use of particular protocols. For example, commitment protocols may be unable to compute over commitments. Other protocols may lack support for certain operators.

Example. Consider the following source program to be executed by hosts a and b :

let $t_1 = 1 + 1$ in let $t_2 = t_1 \times 2$ in skip

and the following data from the compiler's label inference phase and extension points:

1. $\text{viable}(t_1) = \{P_1, P_3, P_4\}$, $\text{viable}(t_2) = \{P_1, P_2\}$
2. $\mathbb{L}(P_1) \Rightarrow \mathbb{L}(t_1)$, $\mathbb{L}(P_3) \Rightarrow \mathbb{L}(t_1)$, $\mathbb{L}(P_4) \not\Rightarrow \mathbb{L}(t_1)$
3. $\mathbb{L}(P_1) \Rightarrow \mathbb{L}(t_2)$, $\mathbb{L}(P_2) \Rightarrow \mathbb{L}(t_2)$
4. $\text{hosts}(P_1) = \{a\}$, $\text{hosts}(P_2) = \{b\}$, $\text{hosts}(P_3) = \{a, b\}$
5. $c_{\text{exec}}(P_1, _) = 5$, $c_{\text{exec}}(P_2, _) = 5$, $c_{\text{exec}}(P_3, _) = 3$
6. $c_{\text{comm}}(P_1, P_1) = 0$, $c_{\text{comm}}(P_3, P_2) = 2$
7. $\text{comm}(P_1, P_1)$, $\neg \text{comm}(P_3, P_1)$
8. $\text{comm}(P_3, P_2)$, $\neg \text{comm}(P_1, P_2)$

Then the compiler constructs the problem of minimizing cost $\beta_1 + \beta_2$ while satisfying the following constraints:

$$\begin{aligned} &(\alpha_1 = P_1 \vee \alpha_1 = P_3) \wedge (\alpha_2 = P_1 \vee \alpha_2 = P_2) \\ &\alpha_1 = P_1 \rightarrow (\gamma_{1,a} \wedge \neg \gamma_{1,b} \wedge \beta_1 = 5) \\ &\alpha_1 = P_3 \rightarrow (\gamma_{1,a} \wedge \gamma_{1,b} \wedge \beta_1 = 3) \\ &\alpha_2 = P_1 \rightarrow (\gamma_{2,a} \wedge \neg \gamma_{2,b} \wedge \alpha_1 \neq P_3 \wedge (\alpha_1 = P_1 \rightarrow \beta_2 = 5 + 0)) \\ &\alpha_2 = P_2 \rightarrow (\neg \gamma_{2,a} \wedge \gamma_{2,b} \wedge \alpha_1 \neq P_1 \wedge (\alpha_1 = P_3 \rightarrow \beta_2 = 3 + 2)) \end{aligned}$$

Note that α_1 , β_1 , $\gamma_{1,a}$, and $\gamma_{1,b}$ are variables associated with t_1 while α_2 , β_2 and $\gamma_{2,a}$, $\gamma_{2,b}$ are variables associated with t_2 . The first constraint bounds the possible values of assignment variables α_1 and α_2 and is generated from the viable protocols returned by the protocol factory. Viable protocols that do not meet authority requirements are filtered out, so P_4 is not a possible value for α_1 . The rest of the constraints describe the relationship between protocol assignments, participating hosts, possible protocol compositions, and cost.³ From this optimization problem the compiler then computes the optimal assignment Π_{opt} where $\Pi_{\text{opt}}(t_1) = P_3$ and $\Pi_{\text{opt}}(t_2) = P_2$.

5 Viaduct Runtime

Once it has computed a protocol assignment, the Viaduct compiler outputs a program where every let-binding and assignable declaration is annotated with the protocol that will execute it. This annotated program can be executed by the Viaduct runtime, which consists of an extensible interpreter that interacts with a set of *protocol back ends*, each of which implement a set of protocols. The interface for protocol back ends is straightforward: back ends must implement methods to execute let-bindings and assignable declarations, and methods to communicate with other protocol back ends.

Each host runs a copy of the interpreter with the annotated program as input. For each statement, the interpreter checks whether the host participates in its execution, as defined by $\text{hosts}(\Pi, \cdot)$ —if not, the statement is treated like **skip**. If a host participates in executing a let-binding or a declaration, the interpreter calls the back end for the protocol assigned to the statement. To execute a conditional, the host retrieves the cleartext value of the guard from the protocol back end that stores it, and executes the appropriate branch. The validity rules for protocol assignments ensure the host is allowed to see the cleartext value, and that it is able to retrieve it.

5.1 Protocol Composition

The protocol back end executing a let-binding must send the computed value to back ends executing statements that read the bound temporary. How one back end sends a value to another depends on the protocols involved. For example, a statement executed in $\text{Replicated}(h_1, h_2)$ reading a temporary computed in $\text{SH-MPC}(h_1, h_2)$ corresponds to executing

³Note that participating host variables are unused here, but in general they encode the guard visibility constraint for conditionals.

Sending protocol (s)	Receiving protocol (r)	Communication	Explanation
Local(h_1)	SH-MPC(h_1, h_2)	$(s, h_1) \xrightarrow{in} (r, h_1)$	create input gate for MPC circuit
Local(h_p)	Commitment(h_p, h_v)	$(s, h_p) \xrightarrow{cc} (r, h_p)$	create commitment
Replicated(h_1, h_2)	SH-MPC(h_1, h_2)	$(s, h_1) \xrightarrow{ct} (r, h_1), (s, h_2) \xrightarrow{ct} (r, h_2)$	read replicated data
SH-MPC(h_1, h_2)	Replicated(h_1, h_2)	$(s, h_1) \xrightarrow{ct} (r, h_1), (s, h_2) \xrightarrow{ct} (r, h_2)$	execute circuit and reveal output
Commitment(h_p, h_v)	Local(h_v)	$(s, h_p) \xrightarrow{occ} (r, h_v), (s, h_v) \xrightarrow{ohc} (r, h_v)$	open commitment
ZKP(h_p, h_v)	Local(h_v)	$(s, h_v) \xrightarrow{ct} (r, h_v)$	send result and proof to verifier

Figure 13. Selected examples of protocol composition. The ct port of various protocols stands for cleartext input; the in port of the MPC protocol represents secret input from a host; the cc port of the Commitment protocol represents creating a commitment; the occ and ohc ports of the Local protocol respectively represent receiving the cleartext value of an opened commitment and the commitment itself.

an MPC circuit and revealing the output to the hosts. On the other hand, a temporary computed in Local(h_3) might not meaningfully be read by a statement executed under SH-MPC(h_1, h_2) as it is unclear how the MPC back end should read local data from an unrelated host.

Viaduct uses the customizable *protocol composer* to define the set of source and destination protocols that can communicate. The composer translates communication between two protocols to a set of messages between hosts participating in the protocols. Developers who want to extend Viaduct with support for a new protocol must enumerate the set of allowed compositions for the protocol and ensure that such compositions are secure.

Formally, the protocol composer translates communication between two protocols P_1 and P_2 to a set of messages, each of the form $(P_1, h_1) \xrightarrow{a} (P_2, h_2)$, where the back end for protocol P_1 at host h_1 sends a message to the back end for protocol P_2 at host h_2 along port a . For a pair (P, h) , it must be the case that $h \in \text{hosts}(P)$. The Viaduct runtime handles the delivery of these messages between back ends.

Each protocol provides a set of ports that define how its back end processes input from another protocol back end. The ZKP protocol, for instance, has two ports: a secret input port, and a public input port. The ZKP back end treats data from the secret input port as the secret input of the prover, while it treats data from its public input port as data known to both the prover and verifier.

Recalling the previous example, when SH-MPC(h_1, h_2) sends a value to Replicated(h_1, h_2), the MPC back ends in h_1 and h_2 jointly execute a circuit in an MPC protocol. The MPC back end at h_1 then sends the revealed circuit output to the cleartext back end (which implements the Replicated protocol) at h_1 along its cleartext port. There is a corresponding message between the MPC and cleartext back ends at h_2 . Step (3) in Figure 5, which depicts execution of the historical millionaires' problem, shows this protocol composition in the context of a larger program.

Figure 13 shows a table of selected compositions and the messages that constitute them. The table illustrates our insight that protocol composition is a general abstraction to represent the use of cryptographic mechanisms. The creation of a commitment and its opening; the execution of an MPC circuit and the revealing of its output; a prover sending a zero-knowledge proof to a verifier—all of these are captured by a composition of one protocol with another.

6 Implementation

We implemented the Viaduct compiler in about 20 KLoC of Kotlin code, which includes code for the parser, the label constraint solver, protocol selection, and the runtime system. The code written against the compiler's extension points—the protocol factory, the protocol composer, the cost estimator, and the protocol back ends—runs to about 4 KLoC. Viaduct uses the Z3 SMT solver [14] to solve the optimization problem generated during protocol selection.

The compiler supports the more liberal surface syntax seen in Figure 2 and Figure 3, as well as functions with bounded polymorphism on parameter labels. The compiler specializes functions at each call site, allowing different compiled implementations for the same function.

We implemented four protocol back ends for Viaduct:

Local/Replicated. The cleartext back end executes code in Local and Replicated protocols. It maintains a store for objects that directly represent the temporaries and assignables of the source program. Computations performed by the cleartext back end are executed directly.

SH-MPC. This back end links Viaduct to ABY, a library for two-party semi-honest MPC [16]. It maintains a store of gate objects that represent circuit components executed by ABY. Computations performed by the back end build gate objects that represent the operation performed (e.g., an addition in the source program creates an ADD gate).

The ABY framework supports execution of circuits in three different schemes—arithmetic sharing, boolean sharing, and

Yao’s garbled circuits—as well as conversions between these, allowing for execution of mixed-protocol circuits. Viaduct represents each scheme as a separate protocol, but all three are implemented by a single back end. To generate efficient mixed circuits, we follow Demmler et al. [16] and Ishaq et al. [28] and estimate inputs to the cost estimator by measuring execution time of individual operations under a particular scheme and conversions between schemes. We perform measurements for two settings: low-latency, high-bandwidth (LAN), and high-latency, low-bandwidth (WAN).⁴ Thus the cost estimator has two modes, each of which optimizes compiled programs for a specific network environment.

Commitment. This back end manages commitments, implemented using SHA-256 hashes of data along with a nonce. The back end for the commitment creator maintains a store of cleartext values along with metadata for commitments. The back end for the commitment receiver maintains the set of commitments, as hashes. The commitment back end cannot support computation.

ZKP. This back end links to libsnark [1], a library for zkSNARKs (zero-knowledge Succinct Non-interactive Arguments of Knowledge). This back end maintains a store of circuit gate objects. The prover and verifier both manage cleartext values for the public inputs to the proof, while only the prover manages cleartext values for the secret inputs. To ensure the prover cannot modify secret inputs mid-execution, all secret inputs are “committed” by sending their hash to the verifier. All proofs that use a secret input then include a clause that equates the input to the pre-image of the hash held by the verifier.

The libsnark library requires proving and verifying keys to be generated for each unique circuit before the protocol is executed. The current prototype requires a “dummy” run of the compiled program to generate these keys.

7 Evaluation

To evaluate Viaduct, we address these research questions:

- RQ1: Is Viaduct expressive enough?
- RQ2: Is its compilation performance acceptable?
- RQ3: Does it generate efficient distributed programs?
- RQ4: How much does label inference reduce the annotation burden for programmers?
- RQ5: What is the overhead of the runtime system?

Experiments used Dell OptiPlex 7050 machines with an 8-core Intel Core i7 7th Gen CPU and 16 GB of RAM. Note that for experiments involving time measurements (RQ2, RQ3, RQ5), the numbers reported are over 5 trials and the relative standard error is at most 6% of the sample mean.

⁴Existing work such as Büscher et al. [7] and Ishaq et al. [28] focus on optimizing mixed circuits for ABY specifically, and as such these employ more sophisticated reasoning about cost for ABY circuits. We consider it future work to incorporate such techniques into Viaduct.

RQ1 - Expressiveness. Figure 14 shows the benchmarks used for the experiments and the cryptography synthesized by Viaduct for each benchmark. Several are from prior work, rewritten in the Viaduct source language. Host configurations are either semi-honest, as in Figure 2, where hosts A and B trust each other for integrity; mutually distrusting as in Figure 3; or are “hybrid” configurations where A and B trust each other but host C is trusted by neither.

Our benchmarks show that Viaduct can compile programs whose security demands a variety of cryptographic mechanisms. With hybrid configurations (interval, bet), Viaduct combines MPC and ZKP to implement different components of a single distributed program. Code for selected benchmarks can be found in the supplemental technical report [2].

RQ2 - Scalability of Compilation. The two main phases of the Viaduct compiler are label inference and protocol selection. Our benchmarks indicate that the overhead of label inference is negligible: at most several hundred milliseconds. As seen in Figure 14, the overhead for protocol selection is more significant, but still on the order of several seconds for most benchmarks. The longest running benchmark, k-means, performs most of its computations in MPC. In this case, it may be harder to converge to the optimal solution since the solver generates a large mixed circuit, choosing between the three MPC schemes supported by ABY.

RQ3 - Cost of Compiled Programs. To show that Viaduct can compile efficient distributed programs, we chose a subset of our benchmarks requiring the use of MPC and compared the execution of optimal programs generated by Viaduct—for each benchmark, one optimized for local area networks (LAN) and another for wide area networks (WAN)—with naive protocol assignments that perform all computation in MPC. The naive ABY assignments use either boolean sharing or Yao garbled circuits, since arithmetic sharing can only perform arithmetic operations. We measured executions in a 1 Gbps LAN and simulated WAN (100 Mbps bandwidth and 50 ms latency). We configured ABY to use 32-bit integers and set its security parameter to 128 bits.

Figure 15 summarizes our results. For some benchmarks (HHI score, hist. millionaires, median, two-round bidding), computation can be securely moved from MPC to cleartext protocols, making execution much more efficient. Even for benchmarks that require computations to be almost entirely in MPC (bio. match, k-means), Viaduct chooses efficient mixed circuits that perform much better than the naive assignments entirely in boolean sharing or Yao circuits. Viaduct replicates the result in Büscher et al. [7] (which specifically targets the ABY framework) in choosing a mix of arithmetic and Yao circuits as optimal assignments for the two benchmarks from that paper, with the exception of the k-means benchmark in the WAN setting.

Benchmark	Description	Protocols		LoC	Ann	Selection	
		LAN / WAN				Vars	Time
battleship	model of the board game	RZ / RZ		79	12	1022	1.0
bet	C bets who wins hist. millionaires b/w A & B	CLRY / CLRY		79	7	1022	1.0
biometric match	min distance b/w sample & database (from [7])	ALRY / ALRY		40	8	708	2.0
guessing game	same as in fig. 3	RZ / RZ		16	6	193	0.4
HHI score	compute market concentration index (from [45])	ALRY / LRY		22	3	285	1.1
historical millionaires interval	same as in fig. 2 but with arrays	LRY / LRY		17	3	187	0.7
	A & B compute interval of combined points, C attests point is in interval	RYZ / RYZ		45	9	660	2.8
k-means	cluster secret points from A & B (from [7])	ARY / RY		82	3	1684	7.9
k-means (unrolled)	k-means w/ 3 unrolled iterations	ARY / RY		174	3	3629	29.0
median	compute median of A & B's lists (from [30])	RY / RY		36	6	386	1.0
rock-paper-scissors	A & B commit to moves then reveal	CR / CR		56	6	741	1.0
two-round bidding	A & B bid for a list of items	LRY / LRY		34	4	575	1.7

Figure 14. Benchmark programs. **Protocols** give the protocols used in the compiled program for either the LAN or WAN setting. Legend for protocols used: **A, B, Y**—ABY arithmetic/boolean/Yao sharing; **C**—Commitment; **L**—Local; **R**—Replicated; **Z**—ZKP. **Ann** gives the minimum number of label annotations needed to write the program. **Selection** gives the number of symbolic variables and run time in seconds for protocol selection, averaged across five runs.

Benchmark	Bool			Yao			Opt-LAN			Opt-WAN		
	LAN	WAN	Comm	LAN	WAN	Comm	LAN	WAN	Comm	LAN	WAN	Comm
bio. match	3.6	95.9	56.0	2.8	7.1	52.3	1.0	2.2	3.9	same as Opt-LAN		
HHI score	0.8	9.7	7.0	0.5	1.6	2.7	0.3	1.1	0.5	0.3	0.9	0.6
hist. million.	1.0	90.6	4.8	0.6	1.6	3.1	0.3	0.7	0.005	same as Opt-LAN		
k-means	56.5	696.1	1273.1	44.4	117.4	1051.3	17.7	35.8	180.0	same as Yao		
median	11.5	1098.7	197.1	12.8	35.4	327.8	0.7	31.7	1.0	same as Opt-LAN		
2-R bidding	17.3	184.7	233.0	17.8	184.5	233.0	3.1	155.5	4.7	same as Opt-LAN		

Figure 15. Run time (in seconds) and communication (in MB) of select benchmark programs, averaged across five runs. **Bool** and **Yao** are naive assignments using boolean sharing and Yao sharing respectively to execute MPC computations. **Opt-LAN** and **Opt-WAN** are optimal assignments generated by Viaduct for the LAN and WAN setting respectively. Optimal time and communication for a benchmark and execution setting pair are in **bold**.

Benchmark	LAN		WAN	
	Time	Slowdown	Time	Slowdown
bio. match	0.4	150%	1.5	50%
HHI score	0.3	0%	1.0	10%
hist. million.	0.3	0%	0.7	0%
k-means	1.2	1380%	4.1	770%
median	0.5	40%	31.5	0%
2-R bidding	1.6	90%	154.7	0%

Figure 16. Run time (in seconds) of LAN-optimized benchmarks hand-written to use ABY directly and the slowdown of running the same benchmarks through the Viaduct runtime in LAN and WAN settings.

RQ4 - Annotation Burden of Security Labels. Security-typed languages add some annotation burden when writing programs. In practice, labels on host declarations and

downgrading operations suffice to specify intended security policies in Viaduct programs. To substantiate this claim, we created two versions of each benchmark program. In one, every variable has a label annotation; in the other, “erased” version, all such labels are omitted.

For all benchmarks, Viaduct generates the same compiled program for the fully labeled and the erased versions. Although the inferred labels for the erased programs are not exactly the same as in their manually labeled counterparts, the differences do not affect the protocols chosen.⁵ The **Ann** column in Figure 14 counts label annotations on erased programs. This is the minimum number of annotations needed to write the program: effectively, the number of downgrades

⁵This mostly occurs with data publicly known to hosts (e.g. loop indices, array lengths). Given hosts Alice and Bob, a fully-annotated benchmark might have label $A \sqcap B$ for the data, but Viaduct infers label $(A \wedge B)^{\leftarrow}$ in the erased version.

plus the number of host declarations, each of which need an authority label. The table shows that the annotation burden is low: most benchmarks need only a few label annotations.

RQ5 - Overhead of Runtime System. The Viaduct runtime introduces some overhead compared to using cryptographic libraries like ABY directly. To measure this overhead, we translated Viaduct’s LAN-optimized outputs for the MPC benchmarks in Figure 15 to directly use the ABY framework’s API. We then measured the performance of these hand-written programs in the LAN and WAN settings.⁶

Figure 16 gives running times for the hand-written programs and the overhead of using the Viaduct runtime. For most benchmarks, the Viaduct runtime incurs an overhead of at most 150% in the LAN setting; the overhead is reduced to at most 50% in the WAN setting where network delay is a more significant factor. This overhead is due to the cost of interpretation and dynamic circuit generation, and can be eliminated by moving circuit generation to compile time [7, 33].

The markedly larger overhead of the k-means benchmark is due to Viaduct recomputing intermediate results. The benchmark has 8 outputs; while Viaduct evaluates 8 smaller MPC circuits each with one output, the hand-written version evaluates one larger circuit with 8 outputs, taking advantage of shared intermediate computations. The compiler could, with additional analysis, determine when output gates can be grouped and executed in the same circuit. We leave this to future work.

8 Related Work

Compilation to Cryptographic Protocols. The idea of compiling a high-level program to a cryptographic protocol has been explored in the context of multiparty computation [27] (e.g., Fairplay [34], SCVM [32], OblivM [33], OblivC [49], Wysteria [40], HyCC [7], SCALE-MAMBA [3]), and that of zero-knowledge proofs (e.g., Pinocchio [37], Gepheto [13], Buffet [46], xjSNARK [31]). Earlier work is generally limited to the domain of a particular fixed cryptographic task (e.g., MPC or ZKP); Viaduct’s novelty is synthesizing efficient protocols *across* cryptographic tasks. Like SCVM [32], Viaduct can synthesize “hybrid” programs that perform computations locally, replicated between hosts, or under MPC. This is impossible in the simple two-point label model that many MPC compilers [3, 33] use, which only distinguish between public (low) and secret (high) information. Viaduct also does not fix the number of hosts in a program (unlike [32–34]), nor fix compiling programs only under a semi-honest or malicious setting (unlike [31–33, 37, 40, 46]).

⁶Running LAN-optimized programs in the WAN setting does not skew the results since Figure 15 shows that LAN-optimized programs perform roughly the same as WAN-optimized programs in the WAN setting.

Program Partitioning. Another line of related work [19, 20, 51, 52] describes distributed computations using sequential programs and captures security requirements using information-flow labels. The Jif/split compiler [51, 52] synthesizes simple cryptographic primitives such as cryptographic commitments to satisfy security constraints that would otherwise be impossible without relying on trusted principals. Unlike Viaduct, Jif/split is not extensible to new protocols. Later work [19, 20] proves computational soundness for a similar system under a strong attacker that controls the network and some of the hosts. However, this work does not support replicating computations (only *data* replication is supported), or the other protocols that Viaduct supports.

9 Conclusion

The prototype implementation of the Viaduct compiler compiles high-level, security-typed programs into efficient distributed programs that employ a variety cryptographic mechanisms to ensure security. And the compiler is agnostic to the set of available protocols, making it easily extensible.

Promising avenues for future work remain. The label model could be extended with availability policies [53], guiding selection of fault-tolerant protocols like quorum replication [54] and MPC with guaranteed output delivery [26]. A more full-fledged implementation of Viaduct could support executing code on trusted execution environments like hardware enclaves [25, 29, 35], the use of special-purpose protocols like private set intersection [12, 38] and Oblivious RAM [44], and the incorporation of a more detailed and accurate cost model [28].

Finally, a full correctness proof for the Viaduct compiler would be a significant research achievement, bridging security notions defined by the programming-languages and cryptography communities. One can see Viaduct source programs as *ideal functionalities* and the distributed programs generated by the compiler as *hybrid protocols* using ideal functionalities implemented by cryptographic mechanisms. The conjectured correctness statement for Viaduct is a simulation proof in the Universal Composability (UC) framework [9], relating a Viaduct source program to the distributed implementation generated by the compiler.

Acknowledgments

This work was supported by the National Science Foundation under grant CNS-1704788, and by the IARPA HECTOR program under a subcontract from IBM. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any of these funders.

We thank our shepherd Mike Hicks and our anonymous reviewers for their insightful suggestions. We also thank Alexa van Hattum, Tobias Kappe, Ralph Recto, and Drew Zagieboylo for feedback during the drafting of this paper.

References

- [1] [n.d.]. <https://github.com/scipr-lab/libsnark>.
- [2] Coşku Acay, Rolph Recto, Joshua Gancher, Andrew Myers, and Elaine Shi. 2021. Viaduct: An Extensible, Optimizing Compiler for Secure Distributed Programs (Technical Report). <https://eprint.iacr.org/2021/468>
- [3] Abdelrahman Aly, Daniele Cozzo, Marcel Keller, Emmanuela Orsini, Dragos Rotaru, Peter Scholl, Nigel P. Smart, and Tim Wood. 2019. SCALE–MAMBA v1.6 : Documentation. <https://homes.esat.kuleuven.be/~nsmart/SCALE>
- [4] Owen Arden, Jed Liu, and Andrew C. Myers. 2015. Flow-Limited Authorization. In *28th IEEE Computer Security Foundations Symp. (CSF)*. 569–583. <https://doi.org/10.1109/CSF.2015.42>
- [5] Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas Jakobsen, Mikkel Krøigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, Michael Schwartzbach, and Tomas Toft. 2009. Financial Cryptography and Data Security. Springer-Verlag, Berlin, Heidelberg, Chapter Secure Multiparty Computation Goes Live, 325–343. https://doi.org/10.1007/978-3-642-03549-4_20
- [6] Niklas Broberg, Bart van Delft, and David Sands. 2013. Paragon for Practical Programming with Information-Flow Control. In *11th ASIAN Symposium on Programming Languages and Systems, APLAS 2013*. Springer, 217–232. https://doi.org/10.1007/978-3-319-03542-0_16
- [7] Niklas Büscher, Daniel Demmler, Stefan Katzenbeisser, David Kretzmer, and Thomas Schneider. 2018. HyCC: Compilation of Hybrid Protocols for Practical Secure Computation. In *25th ACM Conf. on Computer and Communications Security (CCS)*. ACM, New York, NY, USA, 847–861. <https://doi.org/10.1145/3243734.3243786>
- [8] Ran Canetti. 2000. Security and Composition of Multiparty Cryptographic Protocols. *Journal of Cryptology* (2000), 143–202. <https://doi.org/10.1007/s001459910006>
- [9] Ran Canetti. 2001. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *42nd Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 136–145. <https://doi.org/10.1109/SFCS.2001.959888>
- [10] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. 2002. Universally composable two-party and multi-party secure computation. In *34th Annual ACM Symposium on Theory of Computing*. ACM, 494–503. <https://doi.org/10.1145/509907.509980>
- [11] Ethan Cecchetti, Andrew C. Myers, and Owen Arden. 2017. Non-malleable Information Flow Control. In *24th ACM Conf. on Computer and Communications Security (CCS)* (Dallas, TX). ACM, 1875–1891. <https://doi.org/10.1145/3133956.3134054>
- [12] Hao Chen, Kim Laine, and Peter Rindal. 2017. Fast Private Set Intersection from Homomorphic Encryption. In *24th ACM Conf. on Computer and Communications Security (CCS)*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). 1243–1255. <https://doi.org/10.1145/3133956.3134061>
- [13] Craig Costello, Cédric Fournet, Jon Howell, Markulf Kohlweiss, Benjamin Kreuter, Michael Naehrig, Bryan Parno, and Samee Zahur. 2015. Geppetto: Versatile verifiable computation. In *IEEE Symp. on Security and Privacy*. IEEE, 253–270. <https://doi.org/10.1109/SP.2015.23>
- [14] Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: an efficient SMT solver. In *Proceedings of the Theory and Practice of Software, 14th Int'l Conf. on Tools and Algorithms for the Construction and Analysis of Systems* (Budapest, Hungary). Springer-Verlag, Berlin, Heidelberg, 337–340. https://doi.org/10.1007/978-3-540-78800-3_24
- [15] Christian Decker and Roger Wattenhofer. 2014. Bitcoin transaction malleability and MtGox. In *19th European Symposium on Research in Computer Security*. Springer, 313–326. https://doi.org/10.1007/978-3-319-11212-1_18
- [16] Daniel Demmler, Thomas Schneider, and Michael Zohner. 2015. ABY - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation. In *Network and Distributed System Security Symp.* The Internet Society. <https://doi.org/10.14722/ndss.2015.23113>
- [17] Manuel Egele, David Brumley, Yanick Fratantonio, and Christopher Kruegel. 2013. An Empirical Study of Cryptographic Misuse in Android Applications. In *ACM Conf. on Computer and Communications Security (CCS)* (Berlin, Germany). 73–84. <http://doi.acm.org/10.1145/2508859.2516693>
- [18] Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. 1993. The Essence of Compiling With Continuations. In *ACM SIGPLAN Conf. on Programming Language Design and Implementation (PLDI)* (Albuquerque, New Mexico, USA) (PLDI '93). 237–247. <https://doi.org/10.1145/155090.155113>
- [19] Cédric Fournet, Guervan le Guernic, and Tamara Rezk. 2009. A Security-Preserving Compiler for Distributed Programs: From Information-Flow Policies to Cryptographic Mechanisms. In *16th ACM Conf. on Computer and Communications Security (CCS)*. 432–441. <https://doi.org/10.1145/1653662.1653715>
- [20] Cédric Fournet and Tamara Rezk. 2008. Cryptographically sound implementations for typed information-flow security. In *35th ACM Symp. on Principles of Programming Languages (POPL)*. 323–335. <https://doi.org/10.1145/1328438.1328478>
- [21] GDPR 2016. General Data Protection Regulation. <https://gdpr-info.eu>
- [22] Martin Georgiev, Subodh Iyengar, Suman Jana, Rishita Anubhai, Dan Boneh, and Vitaly Shmatikov. 2012. The most dangerous code in the world: validating SSL certificates in non-browser software. In *19th ACM Conf. on Computer and Communications Security (CCS)*, Ting Yu, George Danezis, and Virgil D. Gligor (Eds.). ACM, 38–49. <https://doi.org/10.1145/2382196.2382204>
- [23] Joseph A. Goguen and Jose Meseguer. 1982. Security Policies and Security Models. In *IEEE Symp. on Security and Privacy*. 11–20. <https://doi.org/10.1109/SP.1982.10014>
- [24] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game. In *19th Annual ACM Symposium on Theory of Computing*, Alfred V. Aho (Ed.). 218–229. <https://doi.org/10.1145/28395.28420>
- [25] Anitha Gollamudi, Stephen Chong, and Owen Arden. 2019. Information Flow Control for Distributed Trusted Execution Environments. In *32nd IEEE Computer Security Foundations Symp. (CSF)*. IEEE, 304–318. <https://doi.org/10.1109/CSF.2019.00028>
- [26] S. Dov Gordon, Feng-Hao Liu, and Elaine Shi. 2015. Constant-Round MPC with Fairness and Guarantee of Output Delivery. , 63–82 pages. https://doi.org/10.1007/978-3-662-48000-7_4
- [27] Marcella Hastings, Brett Hemenway, Daniel Noble, and Steve Zdancewic. 2019. SoK: General Purpose Compilers for Secure Multi-Party Computation. In *IEEE Symp. on Security and Privacy*. 1220–1237. <https://doi.org/10.1109/SP.2019.00028>
- [28] Muhammad Ishaq, Ana Milanova, and Vassilis Zikas. 2019. Efficient MPC via Program Analysis: A Framework for Efficient Optimal Mixing. In *26th ACM Conf. on Computer and Communications Security (CCS)*. ACM, 1539–1556. <https://doi.org/10.1145/3319535.3339818>
- [29] David Kaplan, Jeremy Powell, and Tom Woller. 2016. AMD memory encryption.
- [30] Florian Kerschbaum. 2011. Automatically Optimizing Secure Computation. In *18th ACM Conf. on Computer and Communications Security (CCS)*. <https://doi.org/10.1145/2046707.2046786>
- [31] Ahmed Kosba, Charalampos Papamanthou, and Elaine Shi. 2018. xJSnark: A Framework for Efficient Verifiable Computation. In *IEEE Symp. on Security and Privacy*. IEEE, 944–961. <https://doi.org/10.1109/SP.2018.00018>
- [32] Chang Liu, Yan Huang, Elaine Shi, Jonathan Katz, and Michael Hicks. 2014. Automating efficient RAM-model secure computation. In *IEEE Symp. on Security and Privacy*. IEEE, 623–638. <https://doi.org/10.1109/SP.2014.46>
- [33] Chang Liu, Xiao Shaun Wang, Kartik Nayak, Yan Huang, and Elaine Shi. 2015. OblivVM: A Programming Framework for Secure Computation. In

- 25th ACM Symp. on Operating System Principles (SOSP). IEEE, 359–376. <https://doi.org/10.1109/SP.2015.29>
- [34] Dahlia Malkhi, Noam Nisan, Benny Pinkas, and Yaron Sella. 2004. Fairplay - A Secure Two-Party Computation System. In *13th Usenix Security Symposium* (San Diego, CA). 287–302. <http://www.usenix.org/publications/library/proceedings/sec04/tech/malkhi.html>
- [35] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday Savagaonkar. 2013. Innovative Instructions and Software Model for Isolated Execution. In *Workshop on Hardware and Architectural Support for Security and Privacy*. <https://doi.org/10.1145/2487726.2488368>
- [36] Andrew C. Myers. 1999. JFlow: Practical Mostly-Static Information Flow Control. In *26th ACM Symp. on Principles of Programming Languages (POPL)*. 228–241. <https://doi.org/10.1145/292540.292561>
- [37] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. 2013. Pinocchio: Nearly practical verifiable computation. In *IEEE Symp. on Security and Privacy*. IEEE, 238–252. <https://doi.org/10.1109/SP.2013.47>
- [38] Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. 2019. SpOT-Light: Lightweight Private Set Intersection from Sparse OT Extension. In *Advances in Cryptology – CRYPTO 2019*. Springer, 401–431. https://doi.org/10.1007/978-3-030-26954-8_13
- [39] François Pottier and Vincent Simonet. 2002. Information flow inference for ML. In *29th ACM Symp. on Principles of Programming Languages (POPL)*. 319–330. <https://doi.org/10.1145/503272.503302>
- [40] Aseem Rastogi, Matthew A. Hammer, and Michael Hicks. 2014. Wysteria: A Programming Language for Generic, Mixed-Mode Multi-party Computations. In *IEEE Symp. on Security and Privacy*. 655–670. <https://doi.org/10.1109/SP.2014.48>
- [41] Jakob Rehof and Torben Æ. Mogensen. 1996. Tractable Constraints in Finite Semilattices. In *3rd International Symposium on Static Analysis (Lecture Notes in Computer Science)*. Springer-Verlag, 285–300. https://doi.org/10.1007/3-540-61739-6_48
- [42] Daniel Edwin Rutherford. 1965. *Introduction to Lattice Theory*. Oliver and Boyd.
- [43] Andrei Sabelfeld and Andrew C. Myers. 2003. Language-Based Information-Flow Security. *IEEE Journal on Selected Areas in Communications* 21, 1 (Jan. 2003), 5–19. <https://doi.org/10.1109/JSAC.2002.806121>
- [44] Emil Stefanov, Marten van Dijk, Elaine Shi, Christopher W. Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. 2013. Path ORAM: an extremely simple oblivious RAM protocol. In *20th ACM Conf. on Computer and Communications Security (CCS)*, Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung (Eds.). 299–310. <https://doi.org/10.1145/2508859.2516660>
- [45] Nikolaj Volgushev, Malte Schwarzkopf, Ben Getchell, Mayank Varia, Andrei Lapets, and Azer Bestavros. 2019. Conclave: secure multi-party computation on big data. In *ACM SIGOPS/EuroSys European Conference on Computer Systems*, George Candea, Robbert van Renesse, and Christof Fetzer (Eds.). 3:1–3:18. <https://doi.org/10.1145/3302424.3303982>
- [46] Riad S. Wahby, Srinath Setty, Zuocheng Ren, Andrew J. Blumberg, and Michael Walfish. 2015. Efficient RAM and Control Flow in Verifiable Outsourced Computation. In *Network and Distributed System Security Symp.* The Internet Society. <https://doi.org/10.14722/ndss.2015.23097>
- [47] Andrew C. Yao. 1982. Protocols for Secure Computations. In *23rd annual IEEE Symposium on Foundations of Computer Science*. 160–164. <https://doi.org/10.1109/SFCS.1982.38>
- [48] Drew Zagieboylo, G. Edward Suh, and Andrew C. Myers. 2019. Using Information Flow to Design an ISA that Controls Timing Channels. In *32nd IEEE Computer Security Foundations Symp. (CSF)*. <https://doi.org/10.1109/CSF.2019.00026>
- [49] Samee Zahur and David Evans. 2015. Obliv-C: A Language for Extensible Data-Oblivious Computation. *IACR Cryptol. ePrint Arch.* (2015). <http://eprint.iacr.org/2015/1153>
- [50] Steve Zdancewic and Andrew C. Myers. 2001. Robust Declassification. In *14th IEEE Computer Security Foundations Workshop (CSFW)* (Cape Breton, Nova Scotia, Canada). 15–23. <https://doi.org/10.1109/CSFW.2001.930133>
- [51] Steve Zdancewic, Lantian Zheng, Nathaniel Nystrom, and Andrew C. Myers. 2002. Secure Program Partitioning. *ACM Trans. on Computer Systems* 20, 3 (Aug. 2002), 283–328. <https://doi.org/10.1145/566340.566343>
- [52] Lantian Zheng, Stephen Chong, Andrew C. Myers, and Steve Zdancewic. 2003. Using Replication and Partitioning to Build Secure Distributed Systems. In *IEEE Symp. on Security and Privacy* (Oakland, California). 236–250. <https://doi.org/10.1109/SECPRI.2003.1199340>
- [53] Lantian Zheng and Andrew C. Myers. 2005. End-to-End Availability Policies and Noninterference. In *18th IEEE Computer Security Foundations Workshop (CSFW)*. 272–286. <https://doi.org/10.1109/CSFW.2005.16>
- [54] Lantian Zheng and Andrew C. Myers. 2014. A Language-Based Approach to Secure Quorum Replication. In *9th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security (PLAS)*. <https://doi.org/10.1145/2637113.2637117>