PL and Crypto Case studies in their combination



Michael Hicks, Professor of CS @ UMD





Cryptography

- Original focus: data security
- Expanded to computation security
 - Homomorphic encryption
 - Secure Multi-party computation
 - Authenticated data structures
 - Zero-knowledge proofs
 - Oblivious RAM

•

Programming Languages

- Originally just for writing programs
- Expanded to help ensure programs are
 - Correct
 - Efficient
 - Secure

••••

- Maintainable
- In a variety of domains

Crypto Reasoning

 Crypto has developed several techniques for reasoning about security in the presence of a (bounded) adversary

- Reasoning is often done per application (a.k.a. construction)
- Use a simple model of computation to make reasoning easier
 - Turing machines, boolean circuits, etc.

PL Reasoning

- PL focuses on per-language reasoning
 - Type system that ensures all programs P have property Q
 - Program analysis that ensures programs P have no flaws of class B

 PL mathematical semantics highly developed for expressing realistic and accurate models of behavior

Idea: Bring Them Together

- Start with a crypto idea
- Broaden and strengthen it by embedding it in a programming language
 - Prove security of all programs in the language, not just a particular one
 - Optimize the performance of *all* programs
 - Taking advantage of finer-grained, PL-style semantic reasoning

This talk: Two Examples

- Privacy-preserving outsourced computation
 - Generalize use of Oblivious RAM to strengthen security and (greatly) optimize performance
- Integrity-preserving outsourced data management
 - Generalize idea underlying Merkle trees to more optimized data structures

<u>Memory-Trace Oblivious</u> <u>Program Execution</u>

0

Joint work with Chang Liu, Austin Harris, Martin Maas, Mohit Tiwari, Elaine Shi

Outsourcing computation



Security goal

 Privacy: Bob cannot learn about the input or output of the computation Bob is **honest but curious** – will snoop but won't corrupt

Challenge: Physical access

• The cloud provider has physical access to the machine running the computation





Solution(?): Encrypted data/code

 Use a secure co-processor (SC) to encrypt code and data



Access pattern is a side channel

- It turns out that the address trace alone can leak private data
 - For example, prior work has shown that the control flow graph can be inferred from it
- Two sources of leakage
 - The **pattern** of addresses
 - The **total number** of memory events



Access pattern is a side channel

Program

• a[x]:=s

- Input: x=1, s=2
- Memory bus read(x, |) read(s, 2)write(a+1, 2)Information Leakage!

Purple variables are secret Red values are encrypted



Variables accessed leak info





Instructions fetched leak info



Oblivious RAM (ORAM)

- A primitive to hide RAM access patterns
 - Due to Goldreich and Ostrovsky (STOC'87)
 - Reads cannot be distinguished from writes
 - The memory address cannot be distinguished
- Practical implementations now exist
 E.g., Path ORAM (CCS'13) by Stefanov et al.



Architecture with ORAM



ORAM drawbacks

Performance

- ORAM accesses are O(poly-log N) where N is the size of the memory bank
 - Normal accesses O(I)
- N can be large for many cloud computations

Security

- Does not protect leaks via the length of the address trace
 - "Padding out" to a fixed length is impractical





PL to the rescue!

- Use program analysis (a type system) to
 - Optimize the use of ORAM based on the program's semantics
 - Use only where necessary, not by default
 - Reason about the influence of secret data on the length of the address trace



Results

- Formal notion of memory-trace oblivious (MTO) execution
 - Like noninterference, but applied to a physical attacker with access to the memory trace
- Formal language and type system
 - Proved that type-correct programs are MTO
- Prototype compiler
 - Performs transformations and type inference to establish MTO on secure programs

Hardware implementation

 Demonstrate significant speedups compared to naïve ORAM allocation

Memory Trace Oblivious Program Execution, *CSF 2013*, Winner of 2013 NSA Best Scientific Cybersecurity Paper Competition GhostRider: A Hardware-Software System for Memory Trace Oblivious Computation, *ASPLOS* 2015 (Best paper).

Assume: Program not secret

- Then memory access trace *may* reveal no addition information when ...
- Certain data stored in DRAM
 - With encryption, or not
 - Eliminates ORAM overhead entirely
- Data stored in multiple ORAM banks
 - Reduces impact of logarithmic factor



Examples

- Program I a[3]:=0
- a[] can be stored in *normal* DRAM since it contains no secret data
- Program 2 for i=1 to 10 a[i]=s
- a[] stored in encrypted DRAM since access pattern is predictable (i is nonsecret)

- Program 3
 b[s]=1
 c[s]=2
- b[] and c[] must be in ORAM, since s is secret, but can be stored in separate ORAM banks



Hybrid Architecture





Formal language

 Formalize simple imperative source language and its semantics

Variables	x,y,z	\in	Vars
Numbers	n	\in	Nat
ORAM bank	0	\in	ORAMbanks
Expressions	e	::=	$x \mid e \; op \; e \mid x[e] \mid n$
Statements	\boldsymbol{s}	::=	skip $ x := e x[e] := e $
			$if(e, S, S) \mid while \ (e, S)$
Programs	S	::=	$p:s \mid S; S$
Locations	p	::=	$n \mid o$

Big Step Semantics

- $\langle M,S \rangle \Downarrow_t M'$
- *M* is a memory
- S is a program
- t is a trace event
- *M*' is the resulting memory

Big Step Semantics

- Semantics has two novelties
 - Variables in DRAM, and in ORAM banks
 - Execution steps may produce events t
 - Read(x,n) read variable x, has value n
 - Write(x,n) write n to variable x
 - *o* read or write to ORAM bank o
 - Fetch(p) fetch instruction at address p

Memory Trace Obliviousness

- MTO: Given a program S and two indistinguishable memories M1 and M2
 - Running S under MI produces the same trace as running S under M2
- Def. of memory indistinguishability
 - All variables allocated in the same banks
 - **DRAM variables** have **same values** in both
 - ORAM variables may have different values

Type System

- Extends standard security type system with a trace effect T
 - Static abstraction of possible run-time trace
 - Trace effects have notion of equivalence ~
 - Only traces whose lengths we can statically reason about are potentially equivalent
- Judgment Γ , $I \vdash S$; T states
 - Program S is type correct having trace effect T when variables have types given by Γ and in context I (the "pc label")

Type Rule for If

- if (e) then SI else S2 produces trace T under environment Γ only if
 - SI produces trace TI
 - S2 produces trace T2
 - e mentions secret variables, or the context does, implies TI ~ T2, in which case T = TI

• Else T = TI + T2 ("TI or T2")

• (Plus some other details I'm skipping)

Type Rule for If

 $\begin{array}{c} \Gamma \vdash e : \operatorname{Nat} l; T \quad \Gamma, l \sqcup l_0 \vdash S_i; T_i \ (i = 1, 2) \\ l \sqcup l_0 \neq \mathrm{L} \Rightarrow T_1 \sim_L T_2 \quad T' = select(T_1, T_2) \\ \hline \Gamma, l_0 \vdash \operatorname{if}(e, S_1, S_2); T@T' \end{array}$

- S2 produces trace T2
- e mentions secret variables, or the context does, implies TI ~ T2, in which case T = TI
 - Else T = T I + T2 ("T I or T2")
- (Plus some other details I'm skipping)

Type Rule for Loops

- while (e) do S produces trace
 loop(T1,T2) under environment Γ only if
 - e produces trace TI
 - S produces trace T2
 - **e mentions no secret variables**, nor does the context
- Hence: No secret variables in loop guards
 No loops in secure conditionals

Type Rule for Loops

T-While $\frac{\Gamma \vdash e : \text{Nat } l; T_1 \quad \Gamma, l_0 \vdash S; T_2 \quad l \sqcup l_0 \sqsubseteq L}{\Gamma, l_0 \vdash p: \text{while}(e, S); \text{Loop}(p, T_1, T_2)}$

- e produces trace I I
- S produces trace T2
- e mentions no secret variables, nor does the context
- Hence: No secret variables in loop guards
 No loops in secure conditionals



Controlling leaks

while (i < H) do S

Can be rewritten to be
 while (i < N) do
 if (i < H) then S else equiv(S)

Where

- H is secret, but N is a public constant
- equiv(S) is an inert code sequence that produces the same sequence as S



Security

- Theorem: If Γ, I ⊢ S;T then S is memorytrace oblivious
- Proof by standard techniques (induction on derivations)



If rule, for assembly

$$\begin{split} I &= \iota_{1}; I_{t}; \iota_{2}; I_{f} \quad |I_{t}| = n_{1} - 2 \quad |I_{f}| + 1 = n_{2} \\ \iota_{1} &= \mathbf{br} \ r_{1} \ rop \ r_{2} \hookrightarrow n_{1} \qquad \iota_{2} = \mathbf{jmp} \ n_{2} \\ \ell' &= \ell \sqcup \Upsilon(r_{1}) \sqcup \Upsilon(r_{2}) \\ \ell' \vdash I_{t} : \langle \Upsilon, Sym \rangle \to \langle \Upsilon', Sym' \rangle; T_{1} \\ \ell' \vdash I_{f} : \langle \Upsilon, Sym \rangle \to \langle \Upsilon', Sym' \rangle; T_{2} \\ \ell' &= \mathbf{H} \Rightarrow \begin{cases} T_{1}@\mathbf{F} \equiv T_{2} \land \\ \ell = \mathbf{L} \Rightarrow \vdash_{const} \ Sym \land \\ \forall r. \ \Upsilon'(r) = \mathbf{L} \Rightarrow \vdash_{safe} \ Sym'(r) \end{cases} \\ \mathbf{T} \text{-IF} \\ \hline T = ite(\ell' = \mathbf{H}, \ \mathbf{F}@T_{1}@\mathbf{F}, \ \mathbf{F}@((T_{1}@\mathbf{F}) + T_{2})) \\ \ell \vdash I : \langle \Upsilon, Sym \rangle \to \langle \Upsilon', Sym' \rangle; T \end{split}$$

While rule, for assembly

$$\begin{split} I &= I_c; \iota_1; I_b; \iota_2 \\ |I_b| &= n_1 - 2 \qquad |I_c| + n_1 = 1 - n_2 \\ \iota_1 &= \mathbf{br} \ r_1 \ rop \ r_2 \hookrightarrow n_1 \qquad \iota_2 = \mathbf{jmp} \ n_2 \\ \ell \sqcup \Upsilon(r_1) \sqcup \Upsilon(r_2) \sqsubseteq \mathbf{L} \\ \ell \vdash I_c : \langle \Upsilon, Sym \rangle \to \langle \Upsilon', Sym' \rangle; T_1 \\ \ell \vdash I_b : \langle \Upsilon', Sym' \rangle \to \langle \Upsilon, Sym \rangle; T_2 \\ \end{split}$$

$$\mathsf{T-LOOP} \underbrace{ \begin{array}{c} I \\ \ell \vdash I_b : \langle \Upsilon', Sym' \rangle \to \langle \Upsilon, Sym \rangle; T_2 \\ \ell \vdash I : \langle \Upsilon, Sym \rangle \to \langle \Upsilon', Sym' \rangle; \mathbf{loop}(T_1, T_2) \end{array}}_{} \end{split}$$



FPGA-based evaluation

Non-Secure Baseline	All data in DRAM
One ORAM	All data in one ORAM
GhostRider	Optimized version utilizing the hybrid memory model

- We compare One ORAM and GhostRider to see the improvement from the compiler
 - measure the overhead of One ORAM and GhostRider over the non-secure baseline

Result: up to 10x faster than one ORAM



Little overhead over nonsecure baseline for some programs For programs whose memory trace patterns **heavily depend on the input**, speedup is small

Opening the Black Box

- Ghostrider treats ORAM as a primitive
 - Part of the trusted computing base (TCB)
 - Full flexibility costs some performance

Supports random access

• We are now working on a language in which ORAM can be implemented

• Less to trust

- Permits efficient oblivious data structures
- Key feature: Proper use of randomness
- http://www.cs.umd.edu/~mwh/papers/darais17obliv.html



Summary

- Goal: Outsourced computation that is private despite physical snooping
 And gets good performance
- Solution: Memory-trace Oblivious program execution enabled by Program analysis and Oblivious RAM
 - HW/SW co-design
 - Designed/validated using formal semantics of PL: type systems and operational semantics





• (switch decks)

Other PL/crypto connections

- Techniques (or PLs) for analyzing/verifying implementations of crypto
 - E.g., side-channel freeness
- Verified crypto-style proofs of security
 - Adopt PL ideas of contextual equivalent and full abstraction
 - Using PL-style automation (theorem proving)
- Real/ideal paradigm of security
 - Matches specification/implementation in PL

Further reading

- The Synergy between Programming Languages and Cryptography
 - <u>http://www.pl-enthusiast.net/2014/12/17/synergy-programming-languages-cryptography/</u>
- Formal Reasoning in PL and Crypto
 - <u>http://www.pl-enthusiast.net/2014/12/23/formal-</u> reasoning-pl-crypto/
- What is PL research and how is it useful?
 - <u>http://www.pl-enthusiast.net/2015/05/27/what-is-pl-research-and-how-is-it-useful/</u>



Broader Message



- Crypto is cool, PL is cool
- Each has something to offer the other:
 - Result is better than both!

http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=14492

- Other related work from our group
 - PL-optimized secure multiparty computation
 - ASPLOS'15, IEEE S&P'14 x 2, PLAS'13, PLAS'12
 - and works in progress

