Obtaining Provably Secure Services from Formally Verified Remote Attestation



Gene Tsudik¹

Joint work with: Ivan De Oliveira Nunes¹, Karim Eldefrawy², Norrathep Rattanavipanon¹

University of California Irvine¹, SRI International²

In this talk, I'm skipping:

- Talk Outline
- Background on IoT/CPS devices
- Detailed motivation for securing devices

IoT (In) Security

Webcam Maker Takes FTC's Heat

for Internet-of-Things Security Homeland Security warns of 'BrickerBot' malware that destroys unsecured internetconnected devices

Reminiscent of the Mirai botnet that brought down large swathes of the US internet last year, this new malware ngs devices and renders them useless.

Worms

By Richard Adhikari

Sep 5, 2013 3:56 PM PT

Stuxnet worm heralds new era of global

cyberwar

Home > Security

FEATURE

Attack aimed at Iran nuclear plant and at US base show spread of cyber weapo

The Mirai botnet explained: How teen scammers and CCTV cameras almost brought down the internet

1319 GMT (06:19 PDT) | Topic: Security

Mirai took advantage of insecure IoT devices in a simple but clever way. It scanned big blocks of the internet for open Telnet ports, then attempted to log in default passwords. In this way, it was able to amass a botnet army.

Low-end IoT/CPS Devices (amoebas of the computing world)



- Designed for: Low Cost, Low Energy, Small Size, High Scale
- Memory: Program (≈32kB) and Data (≈2-16 kB)
- Single core CPU (8-16MHz; 8 or 16 bits)
- Simple Communication Interfaces for IO (a few kbps)
- Examples: TI MSP-430, AVR ATMega32 (Arduino)

Attack/Compromise Detection vs. Prevention

- Prevention is hard & expensive:
 - Simple devices can not perform fancy crypto, run anti-malware, verify certificates, etc.
- Detection is the next best thing:
 - Goal: Remotely measure internal state of device and detect anomalous/compromised states

Remote Attestation (RA)

- A general approach for detecting malware presence on devices
- Two-party interaction between:
 - **Verifier**: trusted entity
 - **Prover**: potentially infected and untrusted **remote** IoT device
- Goal: measure current internal state of **prover**

RA Interaction



(4) Verify response

Adversarial Model [DAC'15]

- A. Remote malware adversary
 - Exploits various vulnerabilities to inject malware from afar
 - Exploits scale/popularity, probably not narrowly targeted
- B. Local communication adversary
 - Eavesdrops on, and manipulates, communication channel(s)
 - Note: A can lead to B...
- C. Physical adversary (up close and personal)
 - Non-Invasive: mounts hardware side-channel attacks
 - Invasive: (1) read-only, (2) hw-modifying

RA Techniques

Hardware-based

- Effective, but...
- Dedicated hardware support (e.g., a TPM)
- Expensive & overkill for lower-end devices

Software-based

- Relies on precise timing measurement and no real-time accomplices
- Unrealistic assumptions for remote prover except for peripheral/legacy devices

• Hybrid

- SW/HW co-design
- Minimal hardware impact
- <u>Best fit for resource constrained IoT devices?</u>
- **Examples**: SMART, TrustLite, TyTaN, SeED, HYDRA, ERASMUS, SMARM

Why bother with formally verified RA?

- FV promises higher confidence and concrete security guarantees (towards provable security for concrete implementations)
- Current RA techniques do not offer concrete assurances and rigor stemming from FV to guarantee security of designs and their implementations
- Since existing techniques are not systematically designed from abstract models, soundness and security are hard to argue formally
- Subtle issues are easy to miss (indeed they have been!)
- Verification of hybrid (HW/SW) designs is both important and challenging

Overview of VRASED: A Formally Verified RA Architecture

Verification Approach



- 1) Define end-to-end (general) secure RA property
- 2) Break it down into multiple sub-properties
- Prove that sub-properties together imply end-to-end RA security
- 4) Implement VRASED HW/SW design
- Prove that each HW/SW module satisfies each subproperty

Based on (1-5), VRASED implementation satisfies secure RA property

Notation

Notation	Description
PC	Current Program Counter value (16-bits)
Ren	Signal that indicates if the MCU is reading from memory (1-bit)
Wen	Signal that indicates if the MCU is writing to memory (1-bit)
D_{addr}	Address for an MCU memory access (16-bits)
DMAen	Signal that indicates if DMA is currently enabled (1-bit)
DMA_{addr}	Memory address being accessed by DMA, if any (16-bits)
CR	(Code ROM) Memory region where SW-Att is stored: $CR = [CR_{min}, CR_{max}]$
KR	(\mathcal{K} ROM) Memory region where \mathcal{K} is stored: $KR = [K_{min}, K_{max}]$
XS	(eXclusive Stack) secure RAM region reserved for SW-Att computations: $XS = [XS_{min}, XS_{max}]$
MR	(MAC RAM) RAM region in which SW-Att computation result is written: $MR = [MAC_{addr}, MAC_{addr} + MAC_{size} - 1]$
reset	A 1-bit signal that reboots the MCU when set to logic 1

RA Security

Definition 2. 2.1 RA Security Game (RA-game): Assumptions:

- SW-Att is immutable, and K is not known to A

- l is the security parameter and $|\mathcal{K}| = |Chal| = |MR| = l$

- AR(t) denotes the content in AR at time t

- A can modify AR and MR at will; however, it loses its ability to modify them while SW-Att is running

RA-game:

- 1) Setup: A can make poly(l) oracle calls to SW-Att, for arbitrary values of AR and $MR \neq Chal$.
- 2) Challenge: at time t, A is presented with challenge Chal.
- 3) Response: A responds with a pair M, σ and wins if and only if $M \neq AR(t)$ and $\sigma = HMAC(KDF(\mathcal{K}, Chal), M)$.

2.2 RA Security Definition:

An RA protocol is considered secure if there is no ppt A capable of winning the game defined in 2.1 with $P_r[A, RA\text{-game}] > negl(l)$

Intuition behind sub-properties



HW Implementation



Subset of Linear Temporal Logic (LTL):

- $\mathbf{X}\phi \mathbf{ne}\mathbf{X}\mathbf{t}\phi$: holds if ϕ is true at the next system state.
- Fφ <u>F</u>uture φ: holds if there exists a future state where φ is true.
- $\mathbf{G}\phi \mathbf{\underline{G}}$ lobally ϕ : holds if for all future states ϕ is true.
- $\phi \mathbf{U} \psi \phi \underline{\mathbf{U}}$ ntil ψ : holds if there is a future state where ψ holds and ϕ holds for all states prior to that.
- φ W ψ φ Weak Until ψ: holds if φ holds for all states prior to the state when ψ holds. However, the state when ψ holds is not guaranteed to happen. In that case φ will remain true forever.

Example 1: Sub-module Verification

Sub-property: Key Access Control

 $\mathbf{G}: \{\neg (PC \in CR) \land R_{en} \land (D_{addr} \in KR) \rightarrow reset \}$



Example 2: Sub-module Verification

Sub-property: Atomicity + Controlled Invocation

$$\begin{array}{l} \mathbf{G}: \left\{ \begin{array}{c} PC < CR_{min} \lor PC > CR_{max} \\ \neg reset \land (PC \in CR) \land \neg (\mathbf{X}(PC) \in CR) \rightarrow \\ PC = CR_{max} \lor \mathbf{X}(reset) \right\} \\ \mathbf{G}: \left\{ \begin{array}{c} PC = CR_{min} \land \neg irq \\ \neg reset \land \neg (PC \in CR) \land (\mathbf{X}(PC) \in CR) \rightarrow \\ \mathbf{X}(PC) = CR_{min} \lor \mathbf{X}(reset) \right\} \end{array} \right. \\ \mathbf{F} = CR_{min} \land \neg irq \\ \mathbf{F} = CR_{min} \land \neg irq \\ (PC = CR_{min} \land \neg irq \\ \land \neg irq \\ (PC = CR_{max} \land \neg irq \\ \land \neg irq \\ (PC = CR_{max} \land \neg irq \\ \land \neg irq \\ (PC = CR_{max} \land \neg irq \\ \land \neg irq \\ (PC = CR_{max} \land \neg irq \\ \land \neg irq \\ (PC = CR_{max} \land \neg irq \\ \land \neg irq \\ (PC = CR_{max} \land \neg irq \\ (P$$

 $\wedge \neg ira$

SW Implementation

- Most of SW is due to HMAC
- Use verified HMAC (SHA2-256-based) implementation from HACL*
- HACL*: a cryptographic library written and verified using F^*
 - Functional correctness (according to the primitive's spec)
 - o Memory Safety
 - o Secret Independence
- Low* (subset of F^*) can be automatically translated to ${f C}$

What else can we do with VRASED?



What is RA actually good for?

RA alone is not enough!

What to do when malware is remotely detected on Prover?

- Physically re-flash Prover? Inconvenient...
- Remotely update Prover software?
 - * How to ensure that software is indeed updated and starts correctly? Malware can always lie.
- Maybe reset Prover? Erase its memory?
 - * Same issues

Extending VRASED

- PoU: Proof of software update
- PoE: Proof of memory erasure
- PoR: Proof of system-wide reset

PURE: Architecture for <u>Proofs of Update</u>, <u>Reset and Erasure</u>

<u>Main feature:</u> proof of subsequent malware-free state on Prover

Overview of PURE Approach

- For each security service:
 - o State generic protocol definition
 - o State security definition
 - o Extend VRASED to obtain that service
 - Prove that construction is secure according to security definition as long as VRASED is secure
 Using reductions from VRASED security game
- Side-goal: Minimize mods to VRASED
- Start with PoR, then PoU, and PoE

PoR: Formal Definition

DEFINITION 2 (SYNTAX). POR is a tuple (Request, Reset, Verify) of algorithms:

- Request^{Vrf→Prv}(): algorithm initiated by Vrf at time t to request a proof that Prv has performed a reset at some point t' in time, where it must hold that t' > t. As a part of Request, Vrf sends a challenge to Prv.
- Reset^{Prv→Vrf} (Chal) : algorithm executed by Prv to perform a reset and use Chal to provide an unforgeable proof H that a reset has happened.
- Verify^{Vrf} (H, Chal) : algorithm executed by Vrf upon receiving H. It outputs 1 if H is a valid proof in response to Request. Otherwise, it outputs 0.

PoR: Security Definition

Definition 3.

3.1 PoR Security Game (PoR-game): Challenger plays the following game with Adv:

- Adv is given full control over Prv software state and oracle access to Reset calls.
- (2) At time t, Adv is presented with Chal.
- (3) Adv wins iff, after t, Adv can produce H_{Adv} , such that Verify $(H_{Adv}, Chal) = 1$, without causing Prv to reset.

3.2 PoR Security Definition:

A **PoR** scheme is considered secure if for all PPT adversaries Adv, there exists a negligible function negl such that:

 $Pr[\mathcal{A}dv, \mathcal{P}OR\text{-}game] \leq negl(I)$

PoR: SW Implementation

- Extend VRASED SW to support PoR functionality
- New SW is called "TCB"

```
void Hacl_HMAC_SHA2_256_hmac_entry(uint8_t operation) {
 2
         uint8_t key[64] = \{0\};
 3
4
5
         memcpy(key, (uint8_t*) KEY_ADDR, 64);
                                                                                     PoR code (PoR.C): Compute HMAC
         if (operation == RESET) {
                                                                                      on challenge and reset. Reset is
             // fst(PoR.C) instruction:
                                                                                     enforced by VRASED HW.
             hacl_hmac((uint8_t*) RST_RESULT_ADDR, (uint8_t*) key, (uint32_t) 64,
                    (uint8_t*) CHALL_ADDR, (uint32_t) 32);
 789
             return;
             // lst(PoR.C) instruction.
          }else{___
10
             //Unmodified VRASED code for regular attestation
11
             hacl_hmac((uint8_t*) key, (uint8_t*) key, (uint32_t) 64, (uint8_t*)
                    CHALL_ADDR, (uint32_t) 32);
                                                                                      Unmodified VRASED SW
12
             hacl_hmac((uint8_t*) MAC_ADDR, (uint8_t*) key, (uint32_t) 32, (
                    uint8_t*) ATTEST_DATA_ADDR, (uint32_t) ATTEST_SIZE);
13
14
         return();
15
```

PoR: HW Implementation

• Add one new HW sub-module satisfying the following:

 $G:\{[PC = fst(PoR.C) \land F:(PC = lst(TCB))] \rightarrow [\neg(PC = lst(TCB)) \quad U \quad reset]\}$

- Reads as: <u>"After PoR code is invoked (when PC = fst(PoR.C))</u>, PC does not reach the last TCB instruction before a system reset is triggered."
- Since VRASED triggers a reset whenever PC leaves the TCB from any instruction other than lst(TCB)...
- ...it must reset or stay inside TCB forever.
- But, it cannot stay inside TCB forever since HMAC is proven to terminate
- Therefore, it must reset!

Verified HW sub-module

LTL Specification: $G:\{[PC = fst(PoR.C) \land F:(PC = lst(TCB))] \rightarrow [\neg(PC = lst(TCB)) \cup reset]\}$



Recall <u>Controlled Invocation</u>:

- Normally VRASED only allows exiting TCB from the last instruction; otherwise <u>reset!</u>
- This FSM will disallow that, if the TCB call is for a PoR, i.e., if PC = fst(PoR.C)
- Closes the only exit door => the only option is <u>reset</u>!



PoR: Construction (more formally)

CONSTRUCTION 1 (PURE POR). Suppose PoR.C is the program memory region inside the TCB storing PoR software binary: $PoR.C \in TCB$. Construction of PURE PoR is defined as follows:

- Request $V^{rf \rightarrow \mathcal{P}rv}$ (): Vrf generates a random challenge Chal \leftarrow \${0, 1}^l and sends it to $\mathcal{P}rv$.
- Reset $\mathcal{P}^{\mathrm{rv} \to \mathcal{V}_{\mathrm{rf}}}(C \mathrm{hal})$:
 - (1) Use VRASED's HMAC to compute HMAC(K, Chal), and write the result at RST_RESULT_ADDR, where RST_RESULT_ADDR is in a persistent storage (flash) memory region.
 - (2) Enforce the following LTL invariant:

 $\mathbf{G}: \{PC = fst(PoR.C) \to [(PC \in TCB \land \neg DMA_{en}) \ \mathbf{W} \ reset]\}$ (1)

• Verify \mathcal{V}^{rf} (H, Chal): \mathcal{V}^{rf} returns 1 if and only if H = $HMAC(\mathcal{K}, Chal)$.

PoR Proof

- Reduction from VRASED RA security to PoR security
- Intuition:
 - If there is an Adv that wins PoR-game without calling PoR code, same Adv can be used to win VRASED RA-game.
 - o Therefore, Adv <u>must</u> call PoR code.
 - However, PURE LTL specification enforces that whenever PoR code is invoked, a system-wide reset must eventually happen, before PoR result becomes accessible.

Proofs of Software Update

• Verifier wants to install new software (SW) on Prover:

1 - Verifier sends SW to Prover, along with memory region (MEM) where to install it, and a challenge.

- 2 Untrusted (non-RA) code is responsible for installing SW in MEM.
- 3 Prover runs Attestation on MEM and replies with the result.

4 – If result is valid for MEM == SW, Verifier is assured that SW was successfully installed in MEM on Prover.



PoU: Formal Definition

DEFINITION 4 (SYNTAX). A POU scheme consists of a tuple (Request, Install, Verify) fulfilling the following:

- Request ^{Vrf→Prv} (S): algorithm initiated by Vrf to request software S to be installed on Prv at a memory range UR, where |UR| = |S|. Prv also receives a challenge Chal as a part of the request.
- Install^{Prv→Vrf} (Chal, S): algorithm executed by Prv to update software such that UR = S. Upon successful update, it outputs a proof H that the update has completed.
- Verify ^{Vrf} (H, Chal, S): algorithm executed by Vrf upon receiving H. It outputs 1 if H is a valid proof in response to Request. Otherwise, it outputs 0.

PoU: Security Definition

DEFINITION 5. 5.1 PoU Security Game (PoU-game):

Notation:

- UR is the memory region on Prv that should be updated, and UR(t) denotes the content of UR at time t.

Challenger plays the following game with Adv:

- (1) Adv is given full control over Prv software state and oracle access to Install.
- (2) Adv is presented with Chal and S, and continues to have oracle access to Install.
- (3) Eventually at time t, Adv queries Install and outputs H.
- (4) Adv wins if and only if Verify(H, Chal, S) = 1 and $UR(t) \neq S$.

5.2 **PoU** Security Definition:

A **PoU** scheme is secure if, for all Probabilistic Polynomial-Time (PPT) adversaries Adv, there exists a negligible function negl such that:

 $Pr[\mathcal{A}dv, PoU\text{-}game] \leq negl(I)$

PoU Construction & Verification

Construction:

- Use untrusted software to perform a software update
- Then call VRASED to compute a measurement on updated software

Verification:

- No modification to VRASED HW/SW => no verification effort for actual implementation
- Only need reduction from VRASED RA to PoU

PoU Construction (more formally)

CONSTRUCTION 2 (PURE POU). Since Prv has a VRASED-compliant architecture, where memory range AR = UR, the construction is defined as follows:

- Request ^{Vrf→Prv} (S): Vrf outputs S and a random challenge Chal ← \${0, 1}^l to Prv.
- Install^{Prv→Vrf} (Chal, S): Prv performs three steps:
 - (1) Unprivileged software receives S and Chal and writes them into UR and MR, respectively.
 - (2) Call VRASED SW-Att's remote attestation function to compute H = HMAC(KDF(K, Chal), AR) = HMAC(KDF(K, Chal), UR) and write H into MR.
 - (3) Control returns to unprivileged software which is responsible for sending the content of MR to Vrf.
- Verify^{Vrf} (H, Chal, S): Vrf returns 1 if and only if H = HMAC(KDF(K, Chal), S).

Proofs of Memory Erasure

- Special case of PoU
 - Can be viewed as an update to "all zeros": {000...0}
- To erase a region n Prover's memory:
 - Verifier sends erasure request to Prover along with the memory region (MEM) to erase and a challenge.
 - 2. Untrusted (non-RA) code writes "zeros" to MEM.
 - 3. Prover runs Attestation on MEM and replies with the result
 - 4. If Prover's result is valid, i.e., $\underline{H} = \underline{HMAC(K, Challenge||000...0)}$, Verifier knows that MEM was successfully erased.

Proof of Memory Erasure (PoE)

DEFINITION 4 (SYNTAX). A PoU scheme consists of a tuple (Request, Install, Verify) fulfilling the following:

- Request ^{Vrf→Prv} (S): algorithm initiated by Vrf to request software S to be installed on Prv at a memory range UR, where |UR| = |S|. Prv also receives a challenge Chal as a part of the request.
- Install^{Prv→Vrf} (Chal, S): algorithm executed by Prv to update software such that UR = S. Upon successful update, it outputs a proof H that the update has completed.
- Verify ^{Vrf} (H, Chal, S): algorithm executed by Vrf upon receiving H. It outputs 1 if H is a valid proof in response to Request. Otherwise, it outputs 0.

Basically, PoE is a special case of PoU where S = {0}*.

PoE construction and verification follow those of PoU.

Implementation

- PURE was instantiated on Open Cores
 OpenMSP430 Verilog Design
- Synthesized on Basys3 FPGA



See paper for details \rightarrow <u>https://github.com/sprout-uci/vrased</u>

Evaluation (vis-a-vis VRASED)

Architecture	Software (in bytes)			Hardware	
menneeture	ROM Size	Helper Code Size	Max Runtime Mem. Usage	LUT	Reg
OpenMSP430 [35]	-	-	-	1842	684
VRASED [9]	4500	112	2332	1964	721
PURE	4550	138	2336	1968	724

< 1% HW/SW overhead

Table 2: PURE (additional) hardware and software cost



Serially Composing Proofs of Update, Reset, and Erasure

• Composition:

1 - Proof of Update on Prover's program memory to make sure that the proper software was installed

2 - Proof of Erasure to make sure that <u>nothing remains in data</u> <u>memory</u> (e.g., because malware could be hiding there)
3 - After (1) and (2), Proof of Reset to make sure that <u>newly installed</u> software initializes correctly

- Altogether these proofs assure that Prover is moved to a valid, malware-free ("PURE") state.
- Or do they??? 🙂

Conclusion



Next step

Proofs of Remote Execution (PoX):

- ✓ Cryptographically binds:
 - Instructions executed
 - Any output produced by this execution
 - Temporally consistent attestation of such instructions
- ✓ Can be used to build sensors (actuators?) that "cannot lie" even under full software compromise
- ✓ Could yield a provably secure approach to the TOCTOU problem in RA

Next step

Definition 7. Necessary Sub-Properties for Secure Proofs of Execution in LTL.	
Ephemeral Immutability:	
$\mathbf{G}: \; \{[W_{en} \land (D_{addr} \in ER)] \lor [DMA_{en} \land (DMA_{addr} \in ER)] \rightarrow \neg EXEC\}$	(3)
Ephemeral Atomicity:	
$\mathbf{G}: \{ (PC \in ER) \land \neg (\mathbf{X}(PC) \in ER) \to PC = ER_{max} \lor \neg \mathbf{X}(EXEC) \}$	(4)
$\mathbf{G}: \{\neg (PC \in ER) \land (\mathbf{X}(PC) \in ER) \rightarrow \mathbf{X}(PC) = ER_{min} \lor \neg \mathbf{X}(EXEC)\}$	(5)
$\mathbf{G}:\;\{(PC\in ER)\wedge irq ightarrow egree XEC\}$	(6)
Output Protection:	
$\mathbf{G}: \ \{ [\neg (PC \in ER) \land (W_{en} \land D_{addr} \in OR)] \lor (DMA_{en} \land DMA_{addr} \in OR) \lor (PC \in ER \land DMA_{en}) \rightarrow \neg EXEC \}$	(7)
Executable/Output (ER/OR) Boundaries & Challenge Temporal Consistency:	
$\mathbf{G}: \{ ER_{min} > ER_{max} \lor OR_{min} > OR_{max} \to \neg EXEC \}$	(8)
$\mathbf{G}: \{ ER_{min} \leq CR_{max} \lor ER_{max} > CR_{max} \rightarrow \neg EXEC \}$	(9)
$\mathbf{G}: \; \{[W_{en} \land (D_{addr} \in METADATA)] \lor [DMA_{en} \land (DMA_{addr} \in METADATA)] \rightarrow \neg EXEC\}$	(10)
Remark: Note that $Chal_{mem} \in METADATA$.	
Response Protection:	
$\mathbf{G}: \ \{\neg EXEC \land \mathbf{X}(EXEC) \rightarrow \mathbf{X}(PC = ER_{min})\}$	(11)
$\mathbf{G}: \{reset ightarrow eg EXEC\}$	(12)

Secure PoX requires (verification of) several other properties and proving secure composition

Next step

And an architecture of its own (on top of RA):



The end

Info & Pointers:

VRASED: A Verified Hardware/Software Co-Design for Remote Attestation USENIX Security Symposium, 2019 Implementation, etc: <u>https://github.com/sprout-uci/vrased</u>

PURE: Using Verified Remote Attestation to Obtain Proofs of Update, Reset and Erasure in Low-End Embedded Systems International Conference On Computer Aided Design (ICCAD), 2019.

A Verified Architecture for Proofs of Execution on Remote Devices under Full Software Compromise Available at : <u>https://arxiv.org/abs/1908.02444</u>

Advancing remote attestation via computer-aided formal verification of designs and synthesis of executables ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSEC), 2019.

