

IoT Security and Privacy: Remote Attestation

CPS and IoT Security

Alessandro Brighente

Master Degree in Cybersecurity



UNIVERSITÀ
DEGLI STUDI
DI PADOVA



SPRITZ
SECURITY & PRIVACY
RESEARCH GROUP



- IoT devices have their own software and this can be compromised by malicious entities
- It is not easy to detect attacks in on-field devices, as Stuxnet showed
- We need solutions to attest the legitimacy of (software and hardware) components by distantly looking at them and their behavior
- We need to account for the constrained resources of these devices



- Attestation is the activity of making a claim to an appraiser about the properties of a target by supplying evidence which supports that claim
- An attester is a party performing the attestation activity
- An attestation protocol is a cryptographic protocol involving a target an attester, an appraiser and possibly other principals serving as trust proxies
- Target: supply evidence that will be considered authoritative by the appraiser while respecting privacy goals of its target



- We denote as remote attestation a protocol whereby a challenge (Chal) verifies the internal state of a device called a prover (Prov)
- This protocol is performed *remotely*, i.e., over the Internet
- Goal: an honest Prov should create an authentication token that convinces Chal that the former is some well-defined expected state
- If Prov has been compromised by an adversary, the authentication token must reflect this



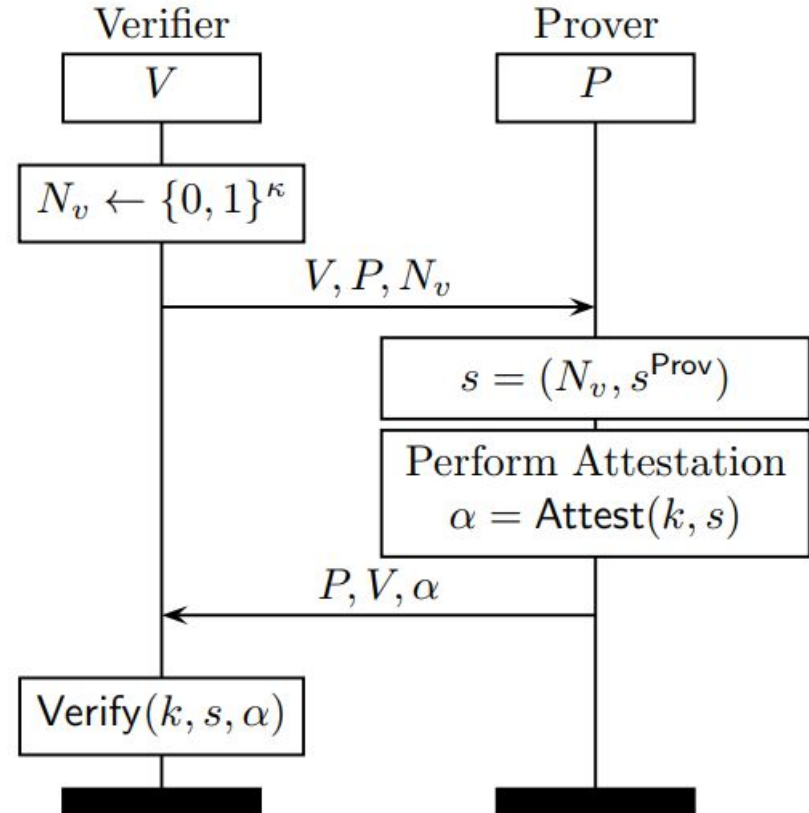
An attestation protocol P is comprised by the following components

- $\text{Setup}()$ - a probabilistic algorithm that, given a security parameter 1^k outputs a long-term key k ;
- $\text{Attest}(k, s)$ - a deterministic algorithm that, given a key k and a device state s , outputs an attestation token a
- $\text{Verify}(k, s, a)$ - a deterministic algorithm that, given a key k , a device state s , and an attestation token a , outputs 1 iff a corresponds to s , i.e., iff $\text{Attest}(k, s) = a$, and outputs 0 otherwise

Attestation Protocol



- The verifier challenges the prover with a fresh nonce (uniformly random and from a large pool)
- The prover attests its state with the key and creates a token
- The verifier receives the token and decides whether to accept it





Let us define the following game

Game 1 ($\text{Att-Forgery}_{\text{Chal,Prov}}(k)$): Chal running P interacts with Prov as follows

1. Chal runs $k \leftarrow \text{Setup}(1^k)$ and outputs s^{Chal} to Prov
2. Prov is given oracles access to *Attest*, i.e., adaptively submit q device states and receive the corresponding token
3. Eventually Prov outputs a ; the game outputs 1 iff $\text{Verify}(k,s,a) = 1$, i.e., iff a corresponds to $s = (s^{\text{Chal}}, s^{\text{Prov}})$



- An honest node has no problem winning the previous game
- If instead Prover has been compromised, its s^{Prov} has changed and must attempt to simulate the operation of Attest
- We can define the following security notion for an attestation protocol

Definition: Att-Forgery Security. An attestation protocol $P = (\text{Setup}, \text{Attest}, \text{Verify})$ is Att-Forgery-secure if there exists a negligible function negl such that for any probabilistic polynomial time prover Prover and sufficiently large k it holds $\Pr[\text{Att-Forgery}_{\text{Chal Prover}}(k) = 1] \leq \text{negl}(k)$



- The central goal of attestation is to verify Prov's state
- Successful execution however does not guarantee that the entire Prov's system can be trusted or that it cannot be compromised after attestation completion
- We assume Prov to be a low end embedded device with a single thread of execution, limited storage capacity, and general complexity
- Although valid for any device, Att-Forgery-security is stronger if the cost of the device is smaller than that of a TPM



- Prov has the following characteristics
 - *Single memory space* (no separation from kernel and user memory)
 - *Single thread of execution* with exception of interrupts (no direct memory access)
 - Ability to *disable interrupts* and force a region of code to execute automatically
 - Availability of *Read Only Memory* (ROM)
 - Ability to *securely cleanup* (erase) memory upon device reset
 - Hardware-based control mechanism to *prevent unauthorized access* to certain memory location



- We make no assumptions about Chal
- A malicious Chal may perform a DoS by forcing Prov to take part in the RA protocol at will
- Malicious Chal does not learn any new information about an honest Prov by performing RA, since Chal must already know the desired state of Prov in order to verify the attestation token
- *We assume that Chal is honest*
- **Note:** Challenger is the term from crypto, verifier is the term from RA. We can use them interchangeably



- We assume the adversary can compromise Prover at any time
- Once it is compromised, the adversary has control over the *prover* device
- There however needs to be *a key that the adversary cannot access* to although being in control of the prover
- We assume that the adversary cannot modify the hardware components of the compromised device
- We also assume that there is a way to protect Attest against side channel attacks



- Attest needs to have specific security properties
- We assure there exists a secure algorithm to compute a based on the prover's state s and a prover-specific key k (e.g., via HMAC)
- Attest must satisfy the following security properties
 - Only Attest can compute a valid token a
 - a accurately captures s , i.e., $\text{Attest}(k,s) = \text{Attest}(k,s')$ with negligible probability
- Two ways to attacker remote attestation
 - Attack 1: The adversary simulates Attest and correctly computes a
 - Attack 2: returned a does not reflect s , i.e., escape detection



- The key k is the only secret held by Prov, and access to k allows the adversary to simulate Attest (i.e., type 1 attack)
- Exclusive access: attest must have exclusive access to k .
- No leaks: Attest leaks no function of k other than a , i.e., after Attest completes, the entire state of Prov is statistically independent from k
- Immutability: Attest code is immutable. This means that it needs to be executed in-place from immutable memory
- Uninterruptibility: the attacker has no means to interrupt the execution of attest



- We derive a set of features that are both necessary and sufficient for remote attestation that achieve the five security properties
- **Exclusive access to k** : in our system model, the best solution is to add a small hardware-based check that monitors the address byte and the program counter and enforces that k is only accessible when PC is within attest
- **No leaks**: we need a way to erase all intermediate values that depend on k , except the attestation token a , when they are no longer needed



- **Immutability:** to ensure Attest to be immutable, we place it in ROM and execute it in place. We consider it as an inexpensive way to enforce immutability.
- **Uninterruptibility:** on a platform with a single thread of execution, the adversary can still regain control after invoking Attest by scheduling an interrupt. Both Attest and instructions to enable and disable interrupts should be *atomic*



- **Invocation from start:** we must enforce exclusive invocation of Attest from its very first instruction. To this aim, custom hardware is needed to enforce the logic: *if the program counter is an address within the Attest code, other than the first instruction address, then the previous instruction must also be within Attest*
- This prevents the adversary to jump in the middle of attest, there is no way to enforce this without OS support
- We hence need to monitor the Attest region and reset the device if detecting illegal behavior



- Remote attestation can be performed in several ways, with different requirements in terms of device capabilities, equipment, and security guarantees
- At a high level, we can distinguish between software-based attestation and root of trust-based attestation



- In software-based attestation, typically we use timing information to allow the verifier to assess the correctness of the firmware running on the prover
- These approaches generally require strict timing requirements on the network, which might not always be feasible in generic IoT
- They also generally consider a one-hop communication between verifier and prover, which makes it hard to be realized in large IoT networks



- A Naive approach for verifying the prover's memory content would be to challenge the prover in computing a MAC of the memory content with a verifier-provided key
- The verifier knows the memory content of the legitimate device
- However, an attacker could easily cheat on this
- Indeed, the attacker could save the original memory content and move it to an empty portion of the memory or to an external device that could be accessible when needing to compute the MAC-based proof



- The embedded device has a memory-content verification procedure that can be remotely activated by the verifier
- The procedure uses a pseudorandom memory traversal
- In particular, the challenge is used as seed for a pseudorandom number generator that generates a list of memory addresses to be checked
- The adversary has no mean to know in advance the portion of the memory that will be checked




- In case the verification procedure is heading towards a portion of the memory that has been altered, the attacker needs to divert it to the a memory location where the correct copy is stored
- This however causes an increase in the time needed to compute the verification
- Thus the verifier will either see a non valid authentication token or a suspiciously long time needed to compute the authentication token



- These schemes leverage a root of trust residing inside the prover
- This component is assumed to be trusted and is the endpoint of the attestation protocol
- It usually comprises a combination of hardware and software
- The value to be attested is stored inside the root of trust
- In IoT devices, the root of trust is realized using hardware with minimal security capabilities, such as code and memory isolation
- Four strategies: interactive RA, Interactive Self-RA, Non-interactive RA, and non-interactive self-RA

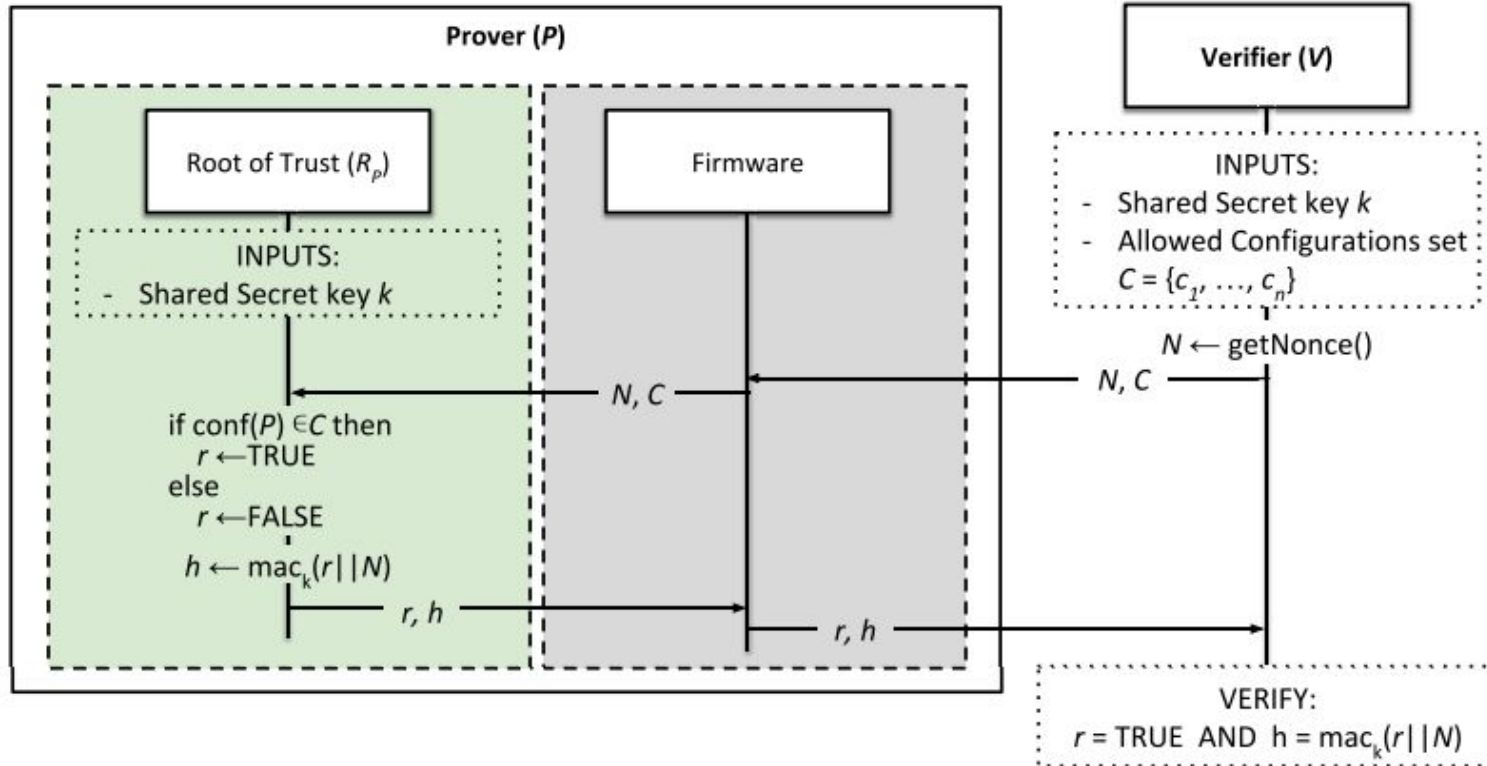


- It consists of an interactive protocol between prover P and verifier V
- V sends a challenge N to P 's root of trust R_p , which responds with a proof of the device's configuration, e.g., $c = \text{hash}(\text{conf}(P)) || N$
- The proof h is either signal (public key crypto) or tagged via MAC (symmetric key)
- V verifies the integrity of P by verifying the authenticity of h





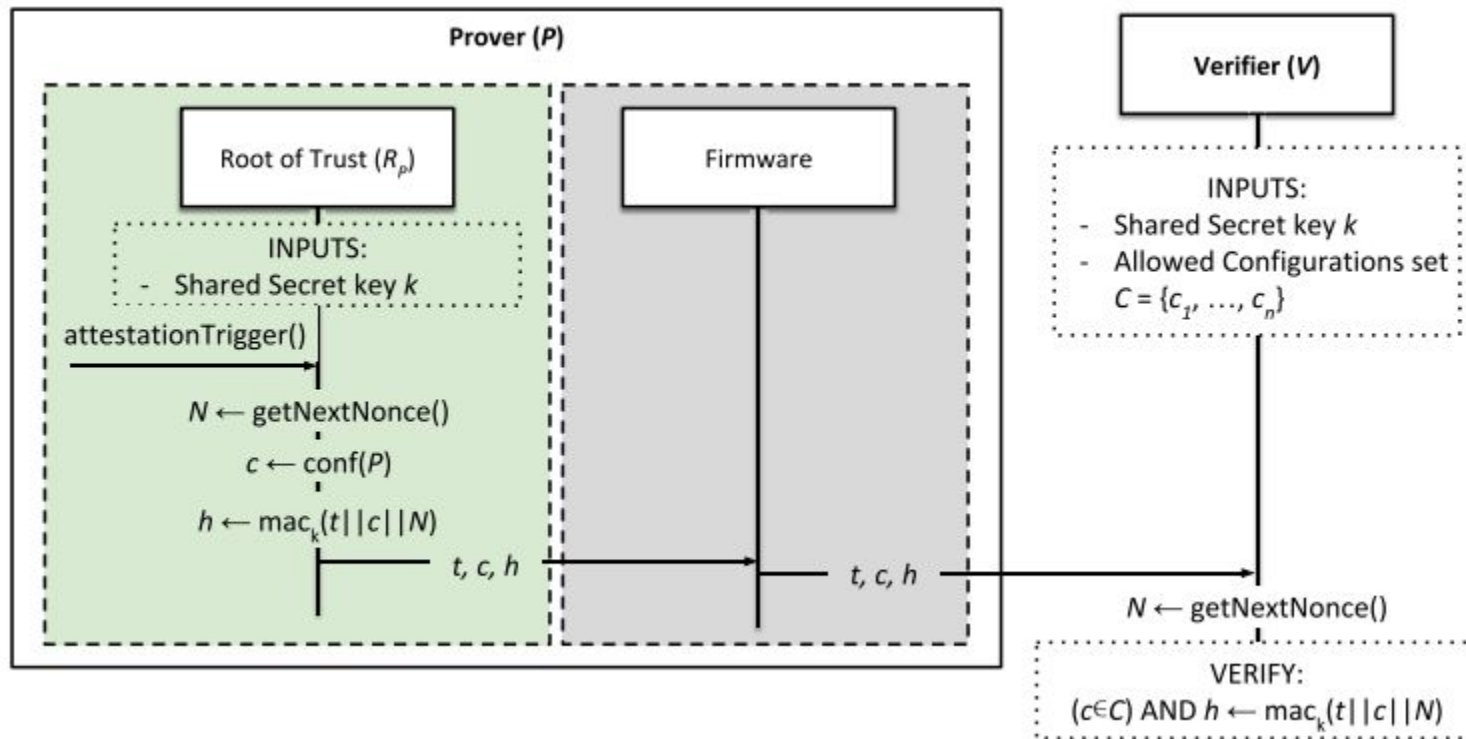
- Leveraging the capabilities of Trusted Execution Environments, it is possible to verify c at P 's side given the list of potential allowed configurations securely stored and accessible by R_p
- After receiving N from V and computing c , R_p produces a signed/MACed token h authenticating a binary result r (true or false)
- This is then delivered to V as a customized token



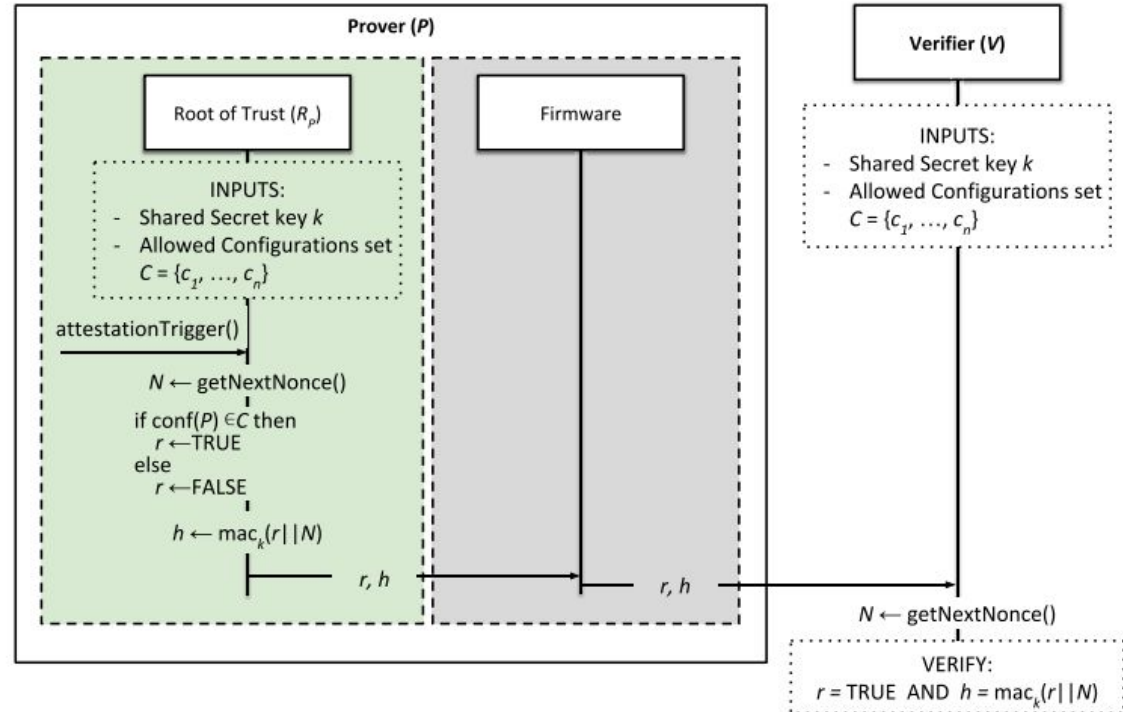


- The prover P autonomously decides the time at which attestation should happen and locally generates a pseudo-random nonce N
- Remove the need for V to start the process, but require additional hardware, e.g., secure source of time
- Examples of additional needs include Real Time Clocks, Attestation Trigger Circuit, or Reliable Read-Only Clocks

Non-Interactive RA



- Variation of the previous one, where P's TEE know the set of possible configurations and can therefore perform self attestation





- In large networks, it might be convenient to assess the status of multiple nodes instead of performing single attestations
- Collective remote attestation has been introduced to limit the time needed to perform attestation of multiple interconnected devices
- We consider a large network of low-end devices which are heterogeneous in terms of software and hardware configurations
- These are provers, and of course we have the verifier
- We need an additional device, i.e., the aggregator



- Verifiable Remote Attestation for Simple Embedded Devices (VRASED)
- A hybrid hardware/software solution to provide formal verification
- Security of hardware while minimizing cost thanks to software
- First formally verified remote attestation scheme



- VRASED is composed of a HW module (HW-Mod) and a SW implementation (SW-Att) of Prv's behavior according to the RA protocol
- HW-Mod enforces access control to K (Prv's unique secret key) in addition to secure and atomic execution of SW-Att
- SW-Att is responsible for computing an attestation report



- Computer-aided formal verification involves three basic steps:
 - The system of interest must be described via a formal model (e.g., finite state machine)
 - Properties that the model should satisfy must be formally specified
 - The system model must be checked against formally specified properties to guarantee that the system retains such properties
- Checking can be achieved via either *theorem proving* or *model checking*



- In model checking, properties are specified as formulae using temporal logic and system models are represented as FSMs
- A system is represented by a triple (S, S_0, T) where S is a finite set of states, $S_0 \subseteq S$ is the set of possible initial states, and $T \subseteq S \times S$ is the transition relation set (set of states that can be reached in a single step from each state)
- Thanks to temporal logic we can represent expected system behavior over time
- We use NuSMV



- In NuSMV, properties are specified in Linear Temporal Logic, particularly useful for sequential systems
- Use of propositional connectives such as conjunction (\wedge), disjunction (\vee), negation (\neg), and implication (\rightarrow).
- LTL includes also temporal connectives enabling sequential reasoning
- We define a list of useful temporal connectives

- We define a list of useful temporal connectives
 - $\mathbf{X}\phi$ – ne \mathbf{X} t ϕ : holds if ϕ is true at the next system state.
 - $\mathbf{F}\phi$ – Future ϕ : holds if there exists a future state where ϕ is true.
 - $\mathbf{G}\phi$ – Globally ϕ : holds if for all future states ϕ is true.
 - $\phi \mathbf{U} \psi$ – ϕ Until ψ : holds if there is a future state where ψ holds and ϕ holds for all states prior to that.
- NuSMV works by checking LTL specifications against the system FSM for all reachable states in such FSM



- We consider an adversary A that can control the entire software state, code, and data of Prv
- A can modify any writable memory and read any memory that is not explicitly protected by access control rules. A can hence read anything that is not protected by HW-Mod
- A can also relocate malware from one memory segment to another to avoid being detected
- A may also have full control over direct memory access controller on Prv



- We focus on attestation functionality of Prv
- The verification approach relies on the following axioms
- **A1 - Program counter:** the PC always contains the address of the instruction being executed in a given cycle
- **A2 - Memory address:** whenever memory is read or written, a data-address signal (D_{addr}) contains the address of the corresponding memory location. R_{en} and W_{en} bits must be set for read and write access, respectively



- **A3 - DMA:** whenever a DMA controller attempts to access main system memory, a DMA-address signal DMA_{addr} reflects the address of the memory location being accessed and a DMA_{en} bit must be set. DMA cannot access memory when DMA_{en} is off
- **A4 - MCU reset:** at the end of a successful reset routine, all registers (including PC) are set to zero before resuming normal software execution flow. Since resets are handled by MCU hardware, no way yo modify them
- **A5 - Interrupts:** when interrupt, the corresponding irq signal is set



- The [MSP430](#) design supports all the five axioms
- Sw-Att uses HACL* HMAC-SHA256 function implemented and verified in [F*](#)
- We assume that the standard compiler can be trusted to semantically preserve its expected behavior
- **A6 - Callee-Saves-Register:** any register touched in a function is cleaned by default when the function returns
- **A7 - Semantic preservation:** functional correctness of HMAC semantically preserved when converted in C

VRASED System Architecture

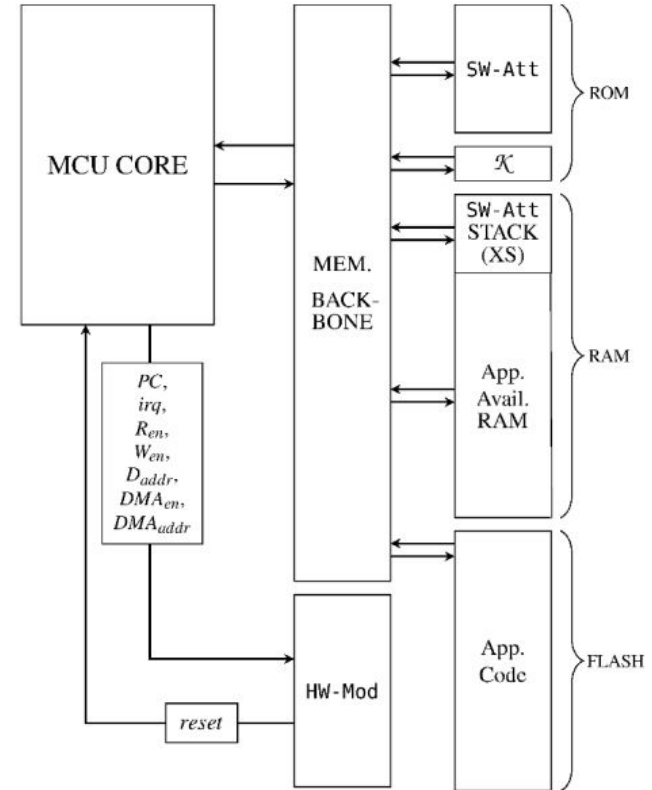


SPRITZ
SECURITY & PRIVACY
RESEARCH GROUP



UNIVERSITÀ
DEGLI STUDI
DI PADOVA

- Implemented by adding HW-Mod to the MCU architecture
- Memory layout is extended to include ROM housing SW-Att code and K
- Access control and SW-Att atomicity are enforced by HW-Mod
- HW-Mod takes 7 input signals from the MCU core



Notation	Description
PC	Current Program Counter value
R_{en}	Signal that indicates if the MCU is reading from memory (1-bit)
W_{en}	Signal that indicates if the MCU is writing to memory (1-bit)
D_{addr}	Address for an MCU memory access
DMA_{en}	Signal that indicates if DMA is currently enabled (1-bit)
DMA_{addr}	Memory address being accessed by DMA, if any
irq	Signal that indicates if an interrupt is occurring (1-bit)
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: $MR = [MAC_{addr}, MAC_{addr} + MAC_{size} - 1]$. The same region is also used to pass the attestation challenge as input to SW-Att
AR	(Attested Region) Memory region to be attested. Can be fixed/predefined or specified in an authenticated request from $\mathcal{V}rf$: $AR = [AR_{min}, AR_{max}]$
$reset$	A 1-bit signal that reboots the MCU when set to logic 1
A1, A2, ..., A7	Verification axioms (outlined in section 3.1)
P1, P2, ..., P7	Properties required for secure RA (outlined in section 3.2)



- We use the following notation
- AR_{\min} and AR_{\max} : first and last physical addresses of the memory region to be attested
- CR_{\min} and CR_{\max} : physical addresses of the first and last instructions of SW-Att in ROM
- K_{\min} and K_{\max} : first and last physical addresses of the ROM region where K is stored
- XS_{\min} and XS_{\max} : first and last physical addresses of the RAM region reserved for SW-Att computation



- We use the following notation
- MAC_{addr} : fixed address that stores the result of SW-Att computation (HMAC)
- MAC_{size} : size of HMAC result
- $[A, B]$ denotes that contiguous memory region between A and B
- Therefore, $C \in [A, B] \Leftrightarrow (C \leq B \wedge C \geq A)$
- $PC \in CR$ Holds when PC is within CR_{min} and CR_{max}



- Each FSM output changes in time as a function of both the current state and current input values
- Each FSM has as inputs a subset of the following signals and wires
$$\{PC, irq, R_{en}, W_{en}, D_{addr}, DMA_{en}, DMA_{addr}\}$$
- Each FSM has only one output, *reset*, that indicates whether any security property was violated. Implicit representation:
 - *reset* is 1 whenever FSM transitions to reset state
 - *reset* remains 1 until a transition leaving the reset state is triggered
 - *reset* is 0 in all other states



- RA soundness corresponds to computing an integrity ensuring function over memory at time t
- We use an HMAC computed on memory AR with a one-time key derived from K and $Chal$
- Since SW-Att is not instantaneous, RA soundness must endure that attested memory does not change during the computation of the HMAC (temporal consistency)
- In other words, the result of SW-Att call must reflect the entire state of the attested memory at the time when SW-Att is called



- In other words, the result of SW-Att call must reflect the entire state of the attested memory at the time when SW-Att is called

Definition 1. *End-to-end definition for soundness of RA computation*

$$G : \{ PC = CR_{min} \wedge AR = M \wedge MR = Chal \wedge [(\neg reset) U (PC = CR_{max})] \rightarrow \\ F : [PC = CR_{max} \wedge MR = HMAC(KDF(\mathcal{K}, Chal), M)] \}$$

where M is any AR value and KDF is a secure key derivation function.



Definition 2.

2.1 RA Security Game (RA-game):

Assumptions:

- *SW-Att* is immutable, and \mathcal{K} is not known to \mathcal{A}
 - l is the security parameter and $|\mathcal{K}| = |\text{Chal}| = |\text{MR}| = l$
 - $\text{AR}(t)$ denotes the content in *AR* at time t
 - \mathcal{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:** \mathcal{A} is given oracle access to *SW-Att*.
 2. **Challenge:** A random challenge $\text{Chal} \leftarrow \{0,1\}^l$ is generated and given to \mathcal{A} . \mathcal{A} continues to have oracle access to *SW-Att*.
 3. **Response:** Eventually, \mathcal{A} responds with a pair (M, σ) , where σ is either forged by \mathcal{A} , or the result of calling *SW-Att* at some arbitrary time t .
 4. \mathcal{A} wins if and only if $\sigma = \text{HMAC}(\text{KDF}(\mathcal{K}, \text{Chal}), M)$ and $M \neq \text{AR}(t)$.
-

2.2 RA Security Definition:

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



- To minimize required hardware feature, hybrid RA approaches implement integrity ensuring functions (e.g., HMAC) in software
- Derive a new unique context-specific key (key) from the master key K via HMAC-based key derivation function on Chal
- Call HACL*'s HMAC using key as the HMAC key
- Sw-Att resides in ROM. guaranteeing software immutability
- Moreover, HW-Mod enforces that no other software running on Prv can access memory allocated by SW-Att



```
1 void HACL_HMAC_SHA2_256_hmac_entry() {
2     uint8_t key[64] = {0};
3     memcpy(key, (uint8_t*) KEY_ADDR, 64);
4     hacl_hmac((uint8_t*) key, (uint8_t*) key, (uint32_t) 64, (uint8_t*)
5             CHALL_ADDR, (uint32_t) 32);
6     hacl_hmac((uint8_t*) MAC_ADDR, (uint8_t*) key, (uint32_t) 32, (uint8_t*)
7             ATTEST_DATA_ADDR, (uint32_t) ATTEST_SIZE);
8     return();
9 }
```



```
1 void HACL_HMAC_SHA2_256_hmac_entry() {
2     uint8_t key[64] = {0};
3     memcpy(key, (uint8_t*) KEY_ADDR, 64);
4     hacl_hmac((uint8_t*) key, (uint8_t*) key, (uint32_t) 64, (uint8_t*)
5             CHALL_ADDR, (uint32_t) 32);
6     hacl_hmac((uint8_t*) MAC_ADDR, (uint8_t*) key, (uint32_t) 32, (uint8_t*)
7             ATTEST_DATA_ADDR, (uint32_t) ATTEST_SIZE);
8     return();
9 }
```

Memory range to be attested

Definition 3. *SW-Att functional correctness*

$$G : \{ PC = CR_{min} \wedge MR = Chal \wedge [(\neg reset \wedge \neg irq \wedge CR = SW-Att \wedge KR = \mathcal{K} \wedge AR = M) \cup PC = CR_{max}] \\ \rightarrow F : [PC = CR_{max} \wedge MR = HMAC(KDF(\mathcal{K}, Chal), M)] \}$$

where M is any arbitrary value for AR .

Natural language: if the memory counter is always preserved until the PC gets to CM_{max} , finally MR will contain a valid response

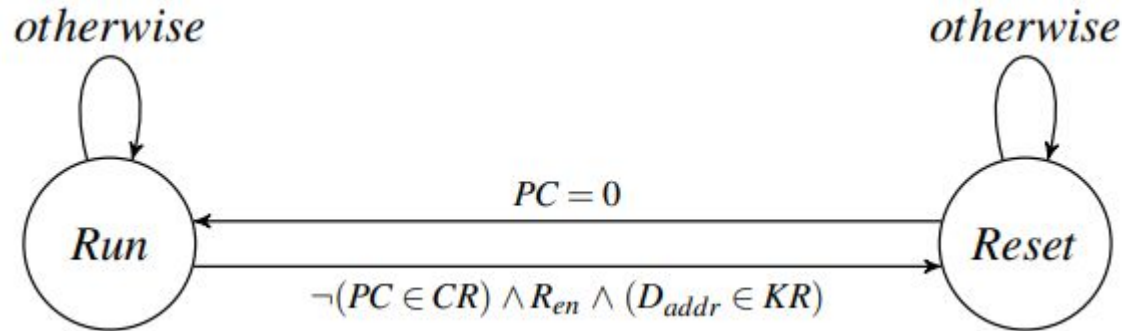


- If malware manages to read K from ROM, it can reply to Vrf with a forged result
- MW-Mod access control (AC) sub-module enforces that K can only be accessed by SW-Att

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

- The system must transition to the Reset state whenever code from outside CR tries to read from D_{addr} within the key space

- We design a FSM from the LTL specification with two states: *run* and *reset*
- Outputs $s_{\text{reset}} = 1$ when the AC submodule transitions to *reset*, implying a hard reset on the MCU





- We define two LTL specifications for atomic execution and controlled invocation

$$\mathbf{G} : \{ [\neg reset \wedge (PC \in CR) \wedge \neg(\mathbf{X}(PC) \in CR)] \rightarrow [PC = CR_{max} \vee \mathbf{X}(reset)] \}$$

$$\mathbf{G} : \{ [\neg reset \wedge \neg(PC \in CR) \wedge (\mathbf{X}(PC) \in CR)] \rightarrow [\mathbf{X}(PC) = CR_{min} \vee \mathbf{X}(reset)] \}$$

$$\mathbf{G} : \{ irq \wedge (PC \in CR) \rightarrow reset \}$$



- Enforce that the only way for SW-Att execution to terminate is through its last instruction $PC = CR_{max}$
- Specified by checking current and next PC values using neXt operator
- If current PC value is within SW-Att region and next PC value is out of SW-Att region, then either current PC value is the address of the last instruction in SW-Att(CR_{max}), or reset is triggered in the next cycle

$$\mathbf{G} : \{ [\neg reset \wedge (PC \in CR) \wedge \neg(\mathbf{X}(PC) \in CR)] \rightarrow [PC = CR_{max} \vee \mathbf{X}(reset)] \}$$



- Enforce that the only way for PC to enter SW-Att region is through the very first instruction CR_{\min}
- This and the previous invariant (LTL specification) imply that it is impossible to jump into the middle of SW-Att or leave SW-Att before reaching the last instruction

$$\mathbf{G} : \{ [\neg reset \wedge \neg(PC \in CR) \wedge (\mathbf{X}(PC) \in CR)] \rightarrow [\mathbf{X}(PC) = CR_{\min} \vee \mathbf{X}(reset)] \}$$

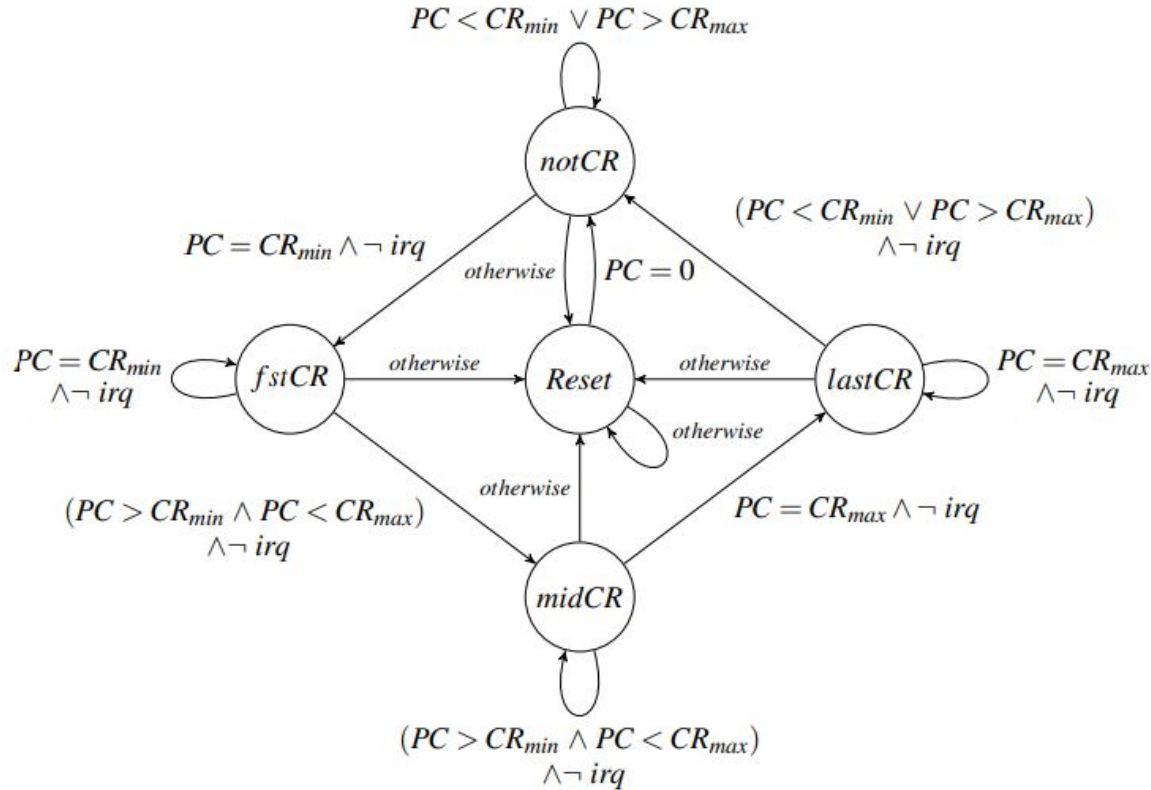


- Atomicity is verified by the following LTL specification
- Although atomicity could be violated by interrupts, this LTL specification prevents an interrupt to happen while SW-Att is executing
- Therefore, if interrupts are not disabled by software running on Prv before calling SW-Att, any interrupt that could violate SW-Att atomicity will necessarily cause an MCU reset

$$\mathbf{G} : \{ irq \wedge (PC \in CR) \rightarrow reset \}$$



- The FSM has 5 states
 - Two basic states notCR and midCR represent movements when PC points to an address 1) outside CR, and 2) within CR, respectively
 - Two fstCR and lstCR represent states when the PC points to the first and last instructions of SW-Att, respectively
 - A reset state
- Transition to reset state whenever 1) any sequence of values for PC does not obey the aforementioned conditions, 2) irq is logical 1 while executing SW-Att



Next



SPRITZ
SECURITY & PRIVACY
RESEARCH GROUP



UNIVERSITÀ
DEGLI STUDI
DI PADOVA

- VRASED
- TOCTOU
- CASU
- GAROTA



- Let us consider an IoT distributed system where nodes communicate in an asynchronous manner via a publish/subscribe pattern
- The verifier performs attestation in two steps: initialization at time T_0 and attestation at time T_1
- During initialization time, the verifier initiates the attestation procedure with one or more services (publishers)
- The publisher then performs the local attestation and publishes the result together with the data it produced



- Every subscriber service receives the publish the data and also perform the attestation
- At attestation time, verifier sends an attestation request to one or more subscribers, which act as prover for the whole network
- Subscribers report an attestation result that includes the result of all the previous services that were directly or indirectly involved in triggering a given event to which the subscriber was registered

An Example of Collective RA



Verifier

Publisher

```
Time  $T_0$   $R \leftarrow \{0,1\}^n$ ;  
 $\sigma_{\text{Vrf}} \leftarrow \text{sig}(\text{SK}_{\text{Vrf}}; P \parallel R)$ ;  $\text{Ch} = \{P, R, \sigma_{\text{Vrf}}\}$   $\rightarrow$   
if ( $\text{vrfsig}(\text{PK}_{\text{Vrf}}; P \parallel R, \sigma_{\text{Vrf}})$ ) then  
  Begin  
    IF ( $\text{timestamp}_p$  is null) then  
       $\text{timestamp}_p[P] = 0$ ;  
     $\text{ServID} \leftarrow P$ ;  
     $\text{GHV}_{\text{prev}} \leftarrow 0$ ;  
     $\text{Input}_p \leftarrow \text{read}()$ ;  
     $\text{Output}_p \leftarrow \text{exec}(\text{ServID}, \text{Input})$ ;  
    ② attest()  
      Begin  
         $\text{LHV}_p \leftarrow \text{checksum}(P)$ ;  
         $\text{timestamp}_p[P] \leftarrow \text{timestamp}_p[P] + 1$ ;  
         $\tau \leftarrow \text{ServID} \parallel \text{timestamp}_p \parallel \text{LHV}_p \parallel \text{Output}_p \parallel \text{Input}_p \parallel \text{GHV}_{\text{prev}}$ ;  
         $\text{GHV}_p \leftarrow \text{Enc}(\text{PK}_{\text{Vrf}}; \tau)$ ;  
      End  
    ③ publish()  
      Begin  
         $\text{msg}_p \leftarrow \text{Output}_p \parallel \text{GHV}_p \parallel \text{timestamp}_p$ ;  
         $\sigma_p \leftarrow \text{sig}(\text{SK}_p; \text{msg}_p)$ ;  
      End  
    Else  
      Reject  $\text{Ch}$ ;  
    End.
```

An Example of Collective RA



Verifier

Publisher

Subscriber

$data = \{msg_p, \mu_p\}$

```
→ IF (vrfsig(PKp; msgp, σp)) then
  Begin
  IF (timestamps is null) then
    timestamps[S] = 0;
  ServID ← S;
  Outputp || GHVp || timestampp ← msgp;
  FOR (i=0; i < length(timestampp); i++) {
    timestamps[i] = max(timestampp[i], timestamps[i]);
  }
  Inputs ← Outputp;
  GHVprev ← GHVp;
  Outputs ← exec (ServID, Inputs);
```

```
④ attest()
  Begin
  LHVs ← checksum(S);
  timestamps[S] ← timestamps[S] + 1;
  τ ← ServID || timestamp || LHVs || Outputs || Inputs || GHVprev;
  GHVs ← Enc(PKVer; τ);
  End
  Else
  Reject data;
  End.
```

