

Modeling and Abstraction of Memory Management in a Hypervisor

Pauline Bolignano ^{1,2} Thomas Jensen ¹ Vincent Siles ²

¹Inria, Rennes, France

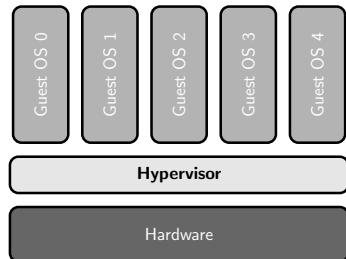
²Prove & Run, Paris, France

April 6th, 2016



Hypervisors

- ▶ Allow to run several OSES on the same platform.
- ▶ Allow to run legacy OS.
- ▶ Are security critical components.
- ▶ Bugs are severe.
 - ▶ Xen bug in Page Tables Management [1].



[1] <https://github.com/QubesOS/qubes-secpack/blob/master/QSBs/qsb-022-2015.txt>

Targeted Properties

Isolation:

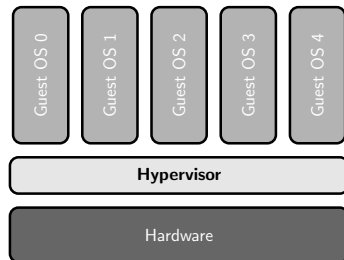
- ▶ Integrity.
- ▶ Confidentiality.

On:

- ▶ Memory.
- ▶ Registers.

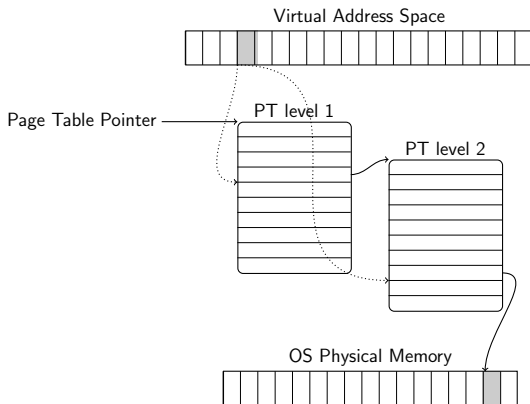
Out of the model:

- ▶ Side Channels.



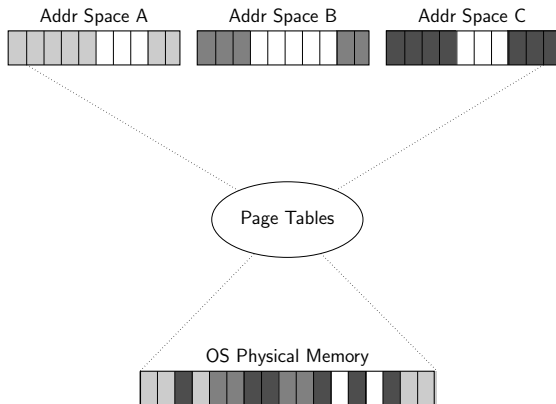
Memory Management for an OS

- ▶ The MMU translates virtual addresses to physical addresses with the Page Tables.

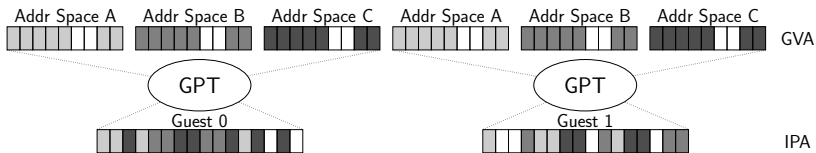


Memory Management for an OS

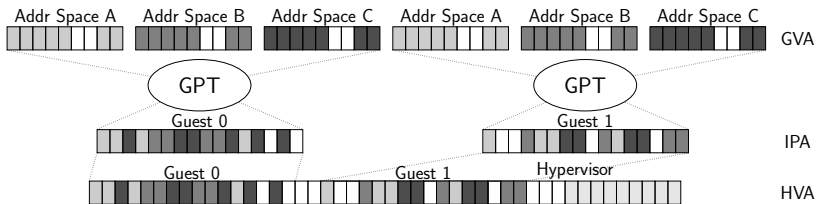
- ▶ The MMU translates virtual addresses to physical addresses with the Page Tables.
- ▶ Each Page Table defines an address space.



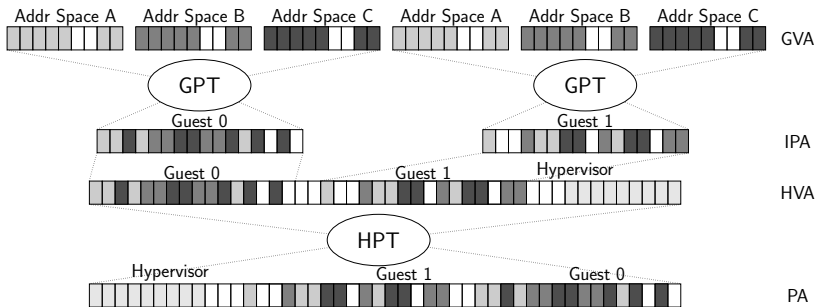
Memory Management for a Hypervisor



Memory Management for a Hypervisor

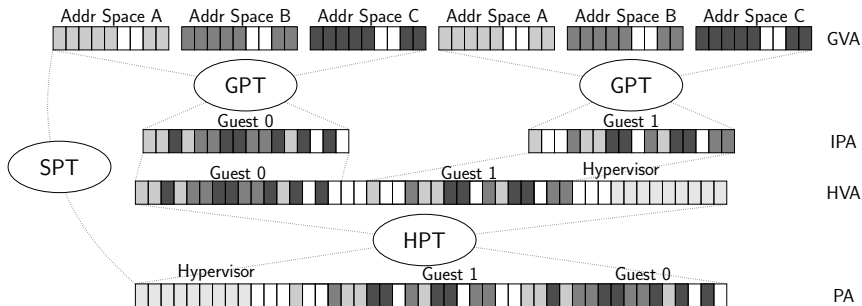


Memory Management for a Hypervisor



Memory Management for a Hypervisor

Shadow Page Tables combine Guest Page Tables and Hypervisor Page Tables.



Context

Prove & Run:

- ▶ Develops Tools:
 - ▶ Language (Smart).
 - ▶ Prover.
- ▶ Develops and proves highly secure systems:
 - ▶ Micro-kernels.
 - ▶ Hypervisors.

Our Work

Which hypervisor?

- ▶ Developed by SecT team (TU Berlin).
- ▶ Para-virtualized.
- ▶ With Shadow Page Tables.
- ▶ With static configuration of memory access rights.

Which tools?

- ▶ Prove & Run Language and Prover.

Outline

Problem Overview

Hypervisors

Shadow Page Tables

Context

Proving Properties on a Hypervisor

Methodology

Abstract State Transition System

Concrete State Transition System

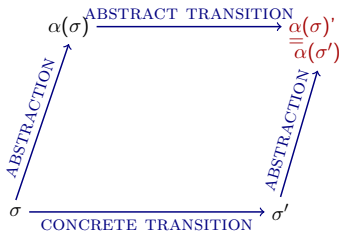
Implementation

Smart

State of the Proof

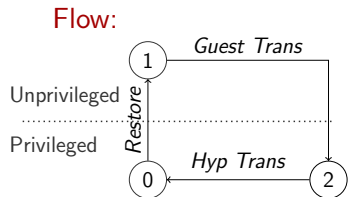
Conclusions

Proof by Abstraction



- ▶ Isolation on the **abstract** system implies isolation on the **concrete** system.
- ▶ The abstract state must be **precise** enough to **express our properties** on it.

Correspondence between Concrete and Abstract Transitions



Hypervisor Transitions:

Concrete	Abstract
Switch Page Fault with MMU Page Fault without MMU Flush Activate MMU Deactivate MMU	Memory Management
Schedule	Schedule
Inject SWI Inject UND Inject ABT Access Privileged Registers	Modify Registers
Handle IRQ Passthrough IRQ Fetch IRQ	Nop

Abstract State

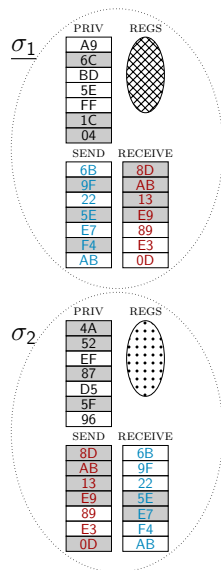
Abstract State :

$$\sigma_\alpha = \langle curr, \sigma_1, \dots, \sigma_n \rangle$$

Where, for all i :

$$\sigma_i = \langle regs, priv, \langle s_1, \dots, s_n \rangle, \langle r_1, \dots, r_n \rangle \rangle$$

- ▶ $priv$: private segment of i .
- ▶ $\sigma_i.s_k$: send segment from i to k .
- ▶ $\sigma_i.r_k$: receive segment from k to i ,
in particular $\sigma_i.r_k = \sigma_k.s_j$.



Properties on the Abstract Model

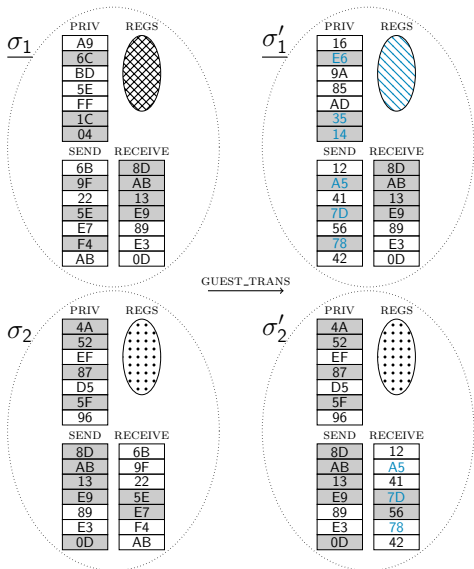
Integrity Theorem: Let i and j be two guest indexes such that $i \neq j$. Consider a transition where j is the running guest. If

$$\langle j, \sigma_1, \dots, \sigma_i, \dots, \sigma_n \rangle \rightarrow \langle j', \sigma'_1, \dots, \sigma'_i, \dots, \sigma'_n \rangle$$

then

- ▶ $\sigma'_i.\text{priv} = \sigma_i.\text{priv}$
- ▶ $\forall k, \sigma'_i.s_k = \sigma_i.s_k$
- ▶ $\forall k \neq j, \sigma'_i.r_k = \sigma_i.r_k$

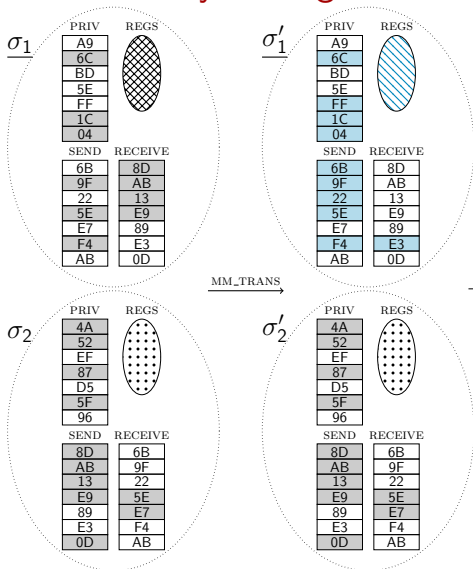
Abstract Guest Transition



$$\sigma'_i = \text{run}(\sigma_i)$$

$$\frac{\forall k \neq i, \sigma'_k = \sigma_k[r_i \leftarrow^{VAL} \sigma'_i.s_k]}{\langle i, \sigma_1, \dots, \sigma_n \rangle \rightarrow \langle i, \sigma'_1, \dots, \sigma'_n \rangle}$$

Abstract Memory Management Transition



$$\frac{\text{decode}(\sigma_i.\text{abs_regs}) = \text{mm}(\sigma'_i)}{\langle i, \sigma_1, \dots, \sigma_i, \dots, \sigma_n \rangle \rightarrow \langle i, \sigma_1, \dots, \sigma'_i, \dots, \sigma_n \rangle}$$

Concrete State

Concrete State:

$$\sigma = \langle \sigma_{\text{HW}}, \sigma_{\text{HYP}}, \text{exception} \rangle$$

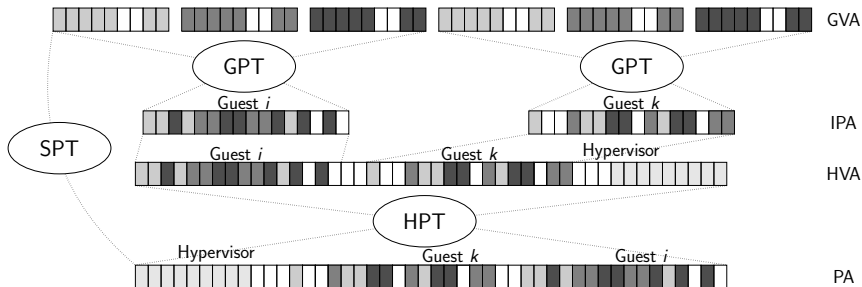
Where:

- ▶ $\sigma_{\text{HW}} = \langle \text{mem}, \text{base}, \text{level}, \text{regs}_{\text{gp}}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gic}} \rangle$
- ▶ $\sigma_{\text{HYP}} = \langle \text{curr}, \sigma_{\text{int}}, \langle \sigma_{G1}, \dots, \sigma_{Gn} \rangle \rangle$
- ▶ $\sigma_{Gi} = \langle \text{vbase}, \text{vmode}, \text{vbnk}, \text{vregs}_{\text{gp}}, \text{vregs}_{\text{mmu}}, \text{vregs}_{\text{gic}} \rangle$

Page Fault Transition:

$$\begin{aligned}
 \text{decode}(\text{abt}, \sigma_{\text{HW}}) &= \text{pf}(\text{gva}) & \sigma_{G_i}.v\text{regs}_{\text{mmu}}.pg &= \text{enabled} \\
 \text{hpt}(\sigma_{G_i}.v\text{base}) &= (\text{pbase}, -) & \{(gva, (ipa, r_0))\} &\in \Gamma_{pt(\text{mem}, \text{pbase})} \\
 \text{hpt}(ipa) &= (pa, -) & \exists r_1 \geq r_0, \text{allowed}(pa, i, r_1) & \\
 \Gamma_{pt(\text{mem}', \text{base})} &= \Gamma_{pt(\text{mem}, \text{base})} \cup \{(gva, (pa, r_0))\} & & \\
 \sigma'_{\text{HYP}} &= \sigma_{\text{HYP}}[\sigma_{\text{int}} \leftarrow \text{alloc}(\sigma_{\text{int}}, \text{mem}, \text{gva})] & &
 \end{aligned}$$

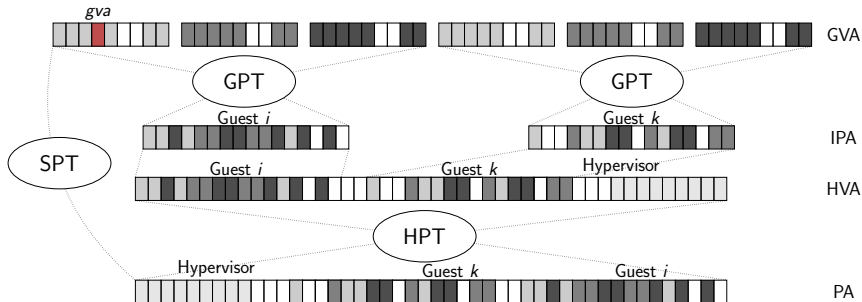
$$\left\langle \begin{array}{l} \langle \text{mem}, \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle \rightarrow \left\langle \begin{array}{l} \langle \text{mem}', \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma'_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle$$



Page Fault Transition:

$$\begin{aligned}
 & \text{decode}(abt, \sigma_{HW}) = pf(gva) \quad \sigma_{Gi}.vregs_{mmu}.pg = \text{enabled} \\
 & \text{hpt}(\sigma_{Gi}.vbase) = (pbase, -) \quad \{(gva, (ipa, r_0))\} \in \Gamma_{pt(mem, pbase)} \\
 & \text{hpt}(ipa) = (pa, -) \quad \exists r_1 \geq r_0, \text{allowed}(pa, i, r_1) \\
 & \Gamma_{pt(mem', base)} = \Gamma_{pt(mem, base)} \cup \{(gva, (pa, r_0))\} \\
 & \sigma'_{HYP} = \sigma_{HYP}[\sigma_{int} \leftarrow \text{alloc}(\sigma_{int}, mem, gva)]
 \end{aligned}$$

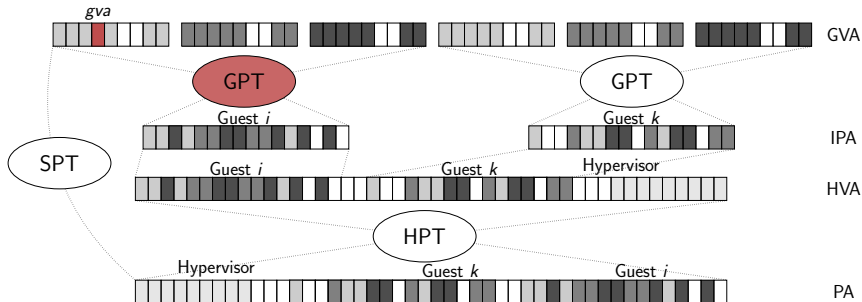
$$\left\langle \begin{array}{l} \langle mem, base, regs_{mmu}, regs_{gp}, pl1, regs_{gic} \rangle \\ \sigma_{HYP} \\ abt \end{array} \right\rangle \rightarrow \left\langle \begin{array}{l} \langle mem', base, regs_{mmu}, regs_{gp}, pl1, regs_{gic} \rangle \\ \sigma'_{HYP} \\ abt \end{array} \right\rangle$$



Page Fault Transition:

$$\begin{aligned}
 & \text{decode}(\text{abt}, \sigma_{\text{HW}}) = \text{pf}(gva) && \sigma_{G_i}.v\text{regs}_{\text{mmu}}.pg = \text{enabled} \\
 & \text{hpt}(\sigma_{G_i}.v\text{base}) = (\text{pbase}, -) && \{(gva, (\text{ipa}, r_0))\} \in \Gamma_{pt(\text{mem}, \text{pbase})} \\
 & \text{hpt}(\text{ipa}) = (\text{pa}, -) && \exists r_1 \geq r_0, \text{allowed}(\text{pa}, i, r_1) \\
 & \Gamma_{pt(\text{mem}', \text{base})} = \Gamma_{pt(\text{mem}, \text{base})} \cup \{(gva, (\text{pa}, r_0))\} \\
 & \sigma'_{\text{HYP}} = \sigma_{\text{HYP}}[\sigma_{\text{int}} \leftarrow \text{alloc}(\sigma_{\text{int}}, \text{mem}, gva)]
 \end{aligned}$$

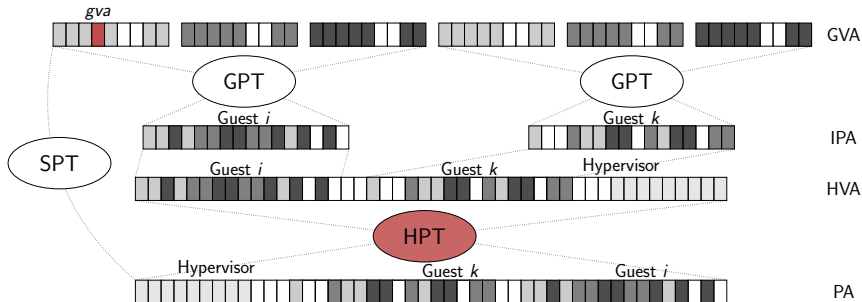
$$\left\langle \begin{array}{l} \langle \text{mem}, \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle \rightarrow \left\langle \begin{array}{l} \langle \text{mem}', \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma'_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle$$



Page Fault Transition:

$$\begin{aligned}
 & \text{decode}(\text{abt}, \sigma_{\text{HW}}) = \text{pf}(\text{gva}) \quad \sigma_{G_i}.v\text{regs}_{\text{mmu}}.pg = \text{enabled} \\
 & \text{hpt}(\sigma_{G_i}.v\text{base}) = (\text{pbase}, -) \quad \{(gva, (ipa, r_0))\} \in \Gamma_{pt(\text{mem}, \text{pbase})} \\
 & \text{hpt}(ipa) = (pa, -) \quad \exists r_1 \geq r_0, \text{allowed}(pa, i, r_1) \\
 & \Gamma_{pt(\text{mem}', \text{base})} = \Gamma_{pt(\text{mem}, \text{base})} \cup \{(gva, (pa, r_0))\} \\
 & \sigma'_{\text{HYP}} = \sigma_{\text{HYP}}[\sigma_{\text{int}} \leftarrow \text{alloc}(\sigma_{\text{int}}, \text{mem}, \text{gva})]
 \end{aligned}$$

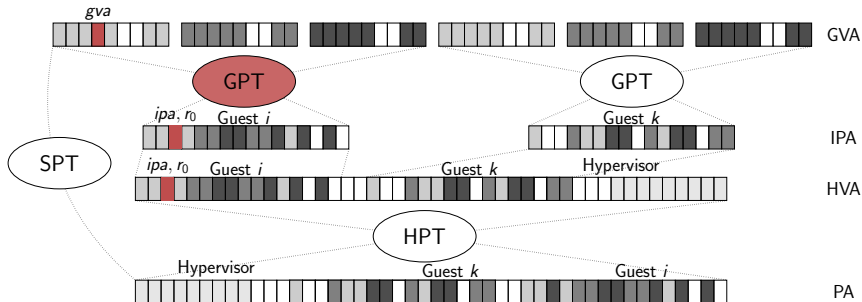
$$\left\langle \begin{array}{l} \langle \text{mem}, \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle \rightarrow \left\langle \begin{array}{l} \langle \text{mem}', \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma'_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle$$



Page Fault Transition:

$$\begin{aligned}
 \text{decode}(\text{abt}, \sigma_{\text{HW}}) &= \text{pf}(gva) & \sigma_{G_i}.v\text{regs}_{\text{mmu}}.pg &= \text{enabled} \\
 \text{hpt}(\sigma_{G_i}.v\text{base}) &= (\text{pbase}, -) & \{(gva, (ipa, r_0))\} &\in \Gamma_{\text{pt}(\text{mem}, \text{pbase})} \\
 \text{hpt}(ipa) &= (pa, -) & \exists r_1 \geq r_0, \text{allowed}(pa, i, r_1) & \\
 \Gamma_{\text{pt}(\text{mem}', \text{base})} &= \Gamma_{\text{pt}(\text{mem}, \text{base})} \cup \{(gva, (pa, r_0))\} & & \\
 \sigma'_{\text{HYP}} &= \sigma_{\text{HYP}}[\sigma_{\text{int}} \leftarrow \text{alloc}(\sigma_{\text{int}}, \text{mem}, gva)] & &
 \end{aligned}$$

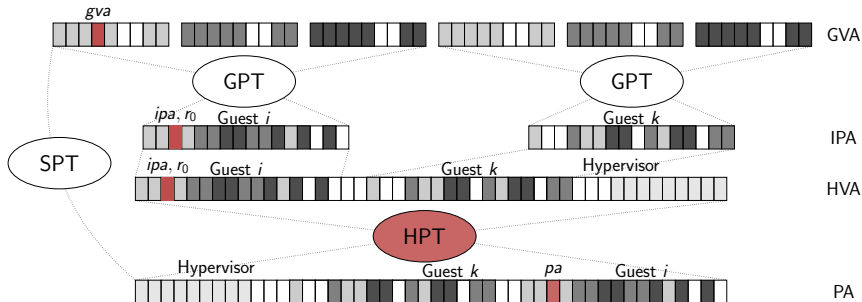
$$\left\langle \begin{array}{l} \langle \text{mem}, \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle \rightarrow \left\langle \begin{array}{l} \langle \text{mem}', \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma'_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle$$



Page Fault Transition:

$$\begin{aligned}
 \text{decode}(\text{abt}, \sigma_{\text{HW}}) &= \text{pf}(gva) & \sigma_{G_i}.v\text{regs}_{\text{mmu}}.pg &= \text{enabled} \\
 \text{hpt}(\sigma_{G_i}.v\text{base}) &= (\text{pbase}, -) & \{(gva, (ipa, r_0))\} &\in \Gamma_{pt(\text{mem}, \text{pbase})} \\
 \text{hpt}(ipa) &= (pa, -) & \exists r_1 \geq r_0, \text{allowed}(pa, i, r_1) & \\
 \Gamma_{pt(\text{mem}', \text{base})} &= \Gamma_{pt(\text{mem}, \text{base})} \cup \{(gva, (pa, r_0))\} & & \\
 \sigma'_{\text{HYP}} &= \sigma_{\text{HYP}}[\sigma_{\text{int}} \leftarrow \text{alloc}(\sigma_{\text{int}}, \text{mem}, gva)] & &
 \end{aligned}$$

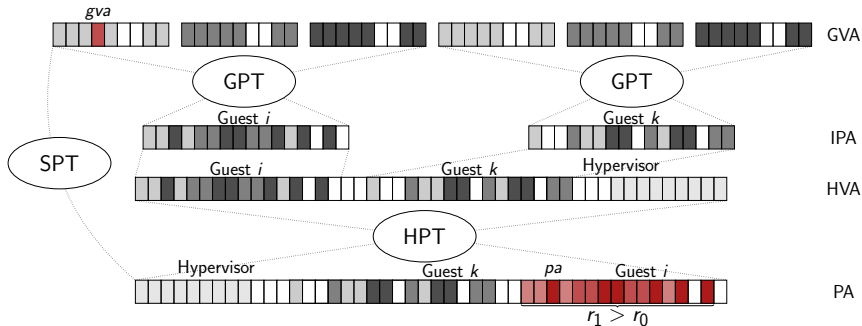
$$\left\langle \begin{array}{l} \langle \text{mem}, \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle \rightarrow \left\langle \begin{array}{l} \langle \text{mem}', \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma'_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle$$



Page Fault Transition:

$$\begin{aligned}
 \text{decode}(\text{abt}, \sigma_{\text{HW}}) &= \text{pf}(gva) & \sigma_{G_i}.v\text{regs}_{\text{mmu}}.pg &= \text{enabled} \\
 \text{hpt}(\sigma_{G_i}.v\text{base}) &= (\text{pbase}, -) & \{(gva, (\text{ipa}, r_0))\} &\in \Gamma_{pt(\text{mem}, \text{pbase})} \\
 \text{hpt}(\text{ipa}) &= (\text{pa}, -) & \exists r_1 \geq r_0, \text{allowed}(\text{pa}, i, r_1) & \\
 \Gamma_{pt(\text{mem}', \text{base})} &= \Gamma_{pt(\text{mem}, \text{base})} \cup \{(gva, (\text{pa}, r_0))\} & & \\
 \sigma'_{\text{HYP}} &= \sigma_{\text{HYP}}[\sigma_{\text{int}} \leftarrow \text{alloc}(\sigma_{\text{int}}, \text{mem}, gva)] & &
 \end{aligned}$$

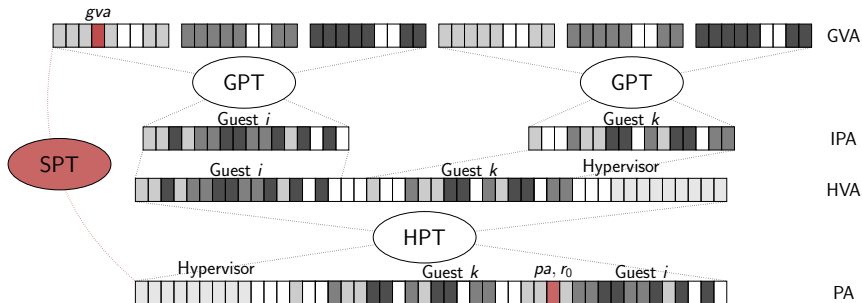
$$\left\langle \begin{array}{l} \langle \text{mem}, \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle \rightarrow \left\langle \begin{array}{l} \langle \text{mem}', \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma'_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle$$



Page Fault Transition:

$$\begin{aligned}
 \text{decode}(\text{abt}, \sigma_{\text{HW}}) &= \text{pf}(\text{gva}) & \sigma_{G_i}.v\text{regs}_{\text{mmu}}.pg &= \text{enabled} \\
 \text{hpt}(\sigma_{G_i}.v\text{base}) &= (\text{pbase}, -) & \{(gva, (ipa, r_0))\} &\in \Gamma_{pt(\text{mem}, \text{pbase})} \\
 \text{hpt}(ipa) &= (\text{pa}, -) & \exists r_1 \geq r_0, \text{allowed}(\text{pa}, i, r_1) & \\
 \Gamma_{pt(\text{mem}', \text{base})} &= \Gamma_{pt(\text{mem}, \text{base})} \cup \{(gva, (\text{pa}, r_0))\} \\
 \sigma'_{\text{HYP}} &= \sigma_{\text{HYP}}[\sigma_{\text{int}} \leftarrow \text{alloc}(\sigma_{\text{int}}, \text{mem}, \text{gva})]
 \end{aligned}$$

$$\left\langle \begin{array}{l} \langle \text{mem}, \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle \rightarrow \left\langle \begin{array}{l} \langle \text{mem}', \text{base}, \text{regs}_{\text{mmu}}, \text{regs}_{\text{gp}}, \text{pl1}, \text{regs}_{\text{gic}} \rangle \\ \sigma'_{\text{HYP}} \\ \text{abt} \end{array} \right\rangle$$



Invariants on SPT

Property ensuring isolation:

if:

- ▶ $base \in \mathcal{B}_{\text{SPT}}(\sigma_{\text{int}}, i)$
- ▶ $(pa, r_0) \in \text{Im}(pt(mem, base))$

then:

- ▶ $\exists r_1 \geq r_0 \wedge \text{allowed}(pa, i, r_1)$

Invariants on SPT

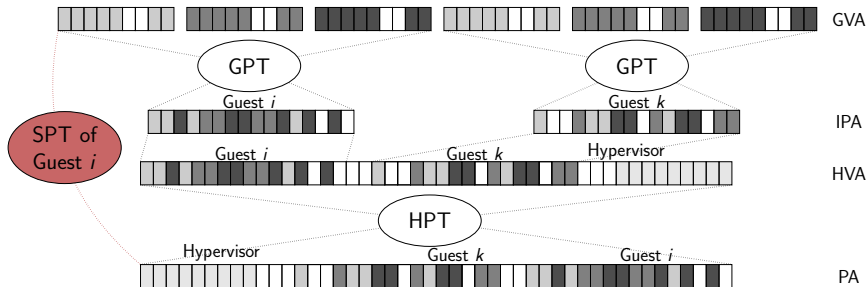
Property ensuring isolation:

if:

- ▶ $base \in \mathcal{B}_{SPT}(\sigma_{int}, i)$
- ▶ $(pa, r_0) \in Im(pt(mem, base))$

then:

- ▶ $\exists r_1 \geq r_0 \wedge allowed(pa, i, r_1)$



Invariants on SPT

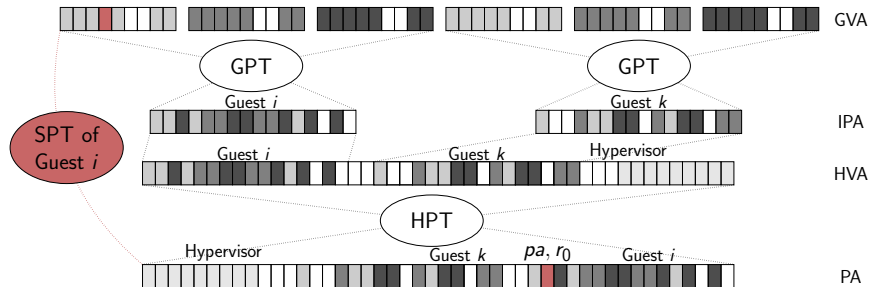
Property ensuring isolation:

if:

- ▶ $base \in \mathcal{B}_{SPT}(\sigma_{int}, i)$
- ▶ $(pa, r_0) \in Im(pt(mem, base))$

then:

- ▶ $\exists r_1 \geq r_0 \wedge allowed(pa, i, r_1)$



Invariants on SPT

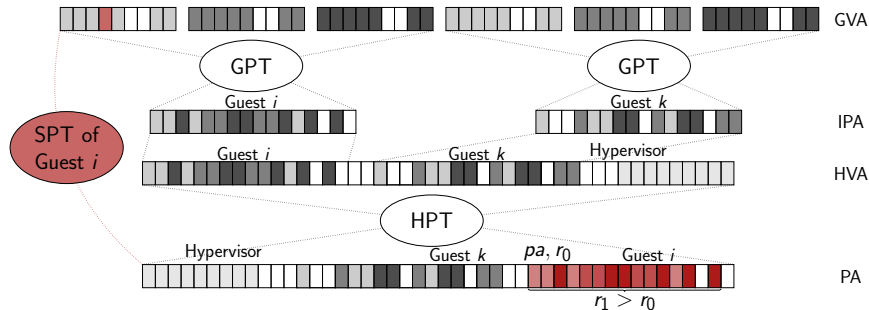
Property ensuring isolation:

if:

- ▶ $base \in \mathcal{B}_{SPT}(\sigma_{int}, i)$
- ▶ $(pa, r_0) \in Im(pt(mem, base))$

then:

- ▶ $\exists r_1 \geq r_0 \wedge allowed(pa, i, r_1)$



Outline

Problem Overview

Hypervisors

Shadow Page Tables

Context

Proving Properties on a Hypervisor

Methodology

Abstract State Transition System

Concrete State Transition System

Implementation

Smart

State of the Proof

Conclusions

Characteristics of Prove & Run Tools

Smart:

- ▶ is a **functional** language.
- ▶ is **pure**.
- ▶ has an **imperative** syntax.
- ▶ allows to write both code and properties.

The Prover:

- ▶ is interactive.
- ▶ provides automatic resolution for simple goals.

Integrity Expressed in Smart

```

public lemma transition_integrity(state st, oracle o,
  idx i)
program
  {{ state nst, seg priv, segstore send, segstore rec,
  segment npriv, segstore nsend, segstore nrec, idx curr }}
  {
    state@curr_guest(st, curr+);
    i != curr => // [i] is not the current guest
    valid(st) =>
    ? transition(st, o, nst+) =>

    ? glo_get_segments(st, i, priv+, send+, rec+) =>
    ? glo_get_segments(nst, i, npriv+, nsend+, nrec+) &&
    npriv = priv &&
    nsend = send &&
    differ_only_on_curr(rec, nrec, curr);
  }

```

Integrity expressed in Smart

"If

- ▶ *curr* is the index of the running guest.
- ▶ *i* is a guest index such that $i \neq curr$.
- ▶ $\langle curr, \sigma_1, \dots, \sigma_i, \dots, \sigma_n \rangle \rightarrow \langle curr', \sigma'_1, \dots, \sigma'_i, \dots, \sigma'_n \rangle$ "

... in Smart :

```
state@curr_guest(st, curr+);
i != curr => // [i] is not the current guest
valid(st) =>
? transition(st, o, nst+) =>
```

Integrity expressed in Smart

"then

- ▶ $\sigma'_i.\text{priv} = \sigma_i.\text{priv}$
- ▶ $\forall k, \sigma'_i.s_k = \sigma_i.s_k$
- ▶ $\forall k \neq \text{curr}, \sigma'_i.r_k = \sigma_i.r_k$ "

... in Smart :

```
? glo_get_segments(st, i, priv+, send+, rec+) =>
? glo_get_segments(nst, i, npriv+, nsend+, nrec+) &&
npriv = priv &&
nsend = send &&
differ_only_on_curr(rec, nrec, curr);
```

Implementation Details

Measure: a *hint* means an interaction with the prover.

State of the proof:

- ▶ Properties of the **concrete** level: **5756 hints**.
 - ▶ **3413** of which for the proof of preservation of invariants over the map operation.
- ▶ Proof of isolation at the abstract level: **336 hints** for **confidentiality** and **235 hints** for **integrity**.
- ▶ Proof of commutation of the **Guest transition**: **1779 hints**.
- ▶ Proof of commutation of the **Page Fault transition**: **1885 hints**.

Conclusions

Recent advances of Formal Methods in complex systems:

- ▶ Sel4, PROSPER, Verisoft XT project.

SPT management is:

- ▶ Complex.
- ▶ Error prone.

The work presented:

- ▶ Provides a **formalization** of a hypervisor, in particular of its **SPT management**.
- ▶ Studies the **properties** needed on SPT to ensure **isolation**.
- ▶ Aims at **simplifying further proofs** by linking the hypervisor model to an abstract model without SPT.

Thank you for your attention

