

A Separation Logic for Communicating Virtual Machines

@2nd Iris Workshop

Zongyuan Liu

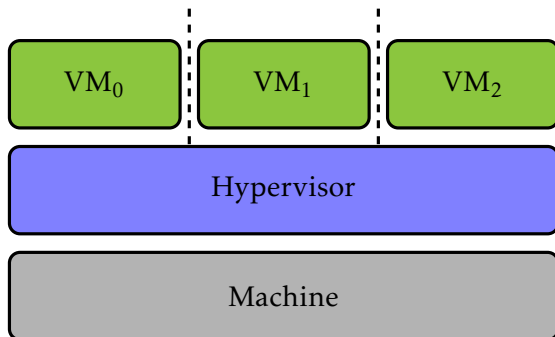
**Joint work with Sergei Stepanenko,
Aslan Askarov, Jean Pichon-Pharabod,
Amin Timany, Lars Birkedal**

Aarhus University

May 2, 2022

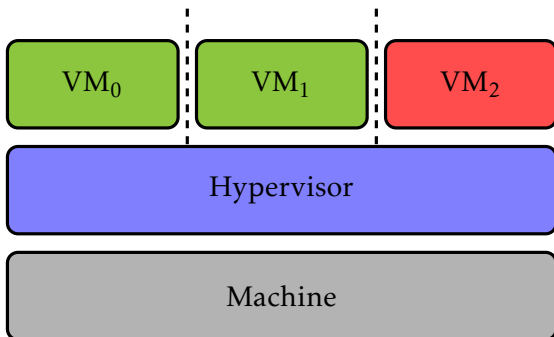
Hypervisor and Virtual Machines

- Allows one host machine to run multiple guest VMs
- Ensures VMs run as if on bare metal, with their own CPUs, registers, memory etc.
- Provides *isolation* between VMs
- Allows controlled communication



Hypervisor and Virtual Machines

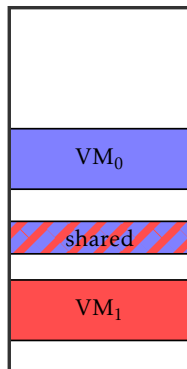
- Allows one host machine to run multiple guest VMs
- Ensures VMs run as if on bare metal, with their own CPUs, registers, memory etc.
- Provides *isolation* between VMs
- Allows controlled communication



Memory Management in Hypervisors

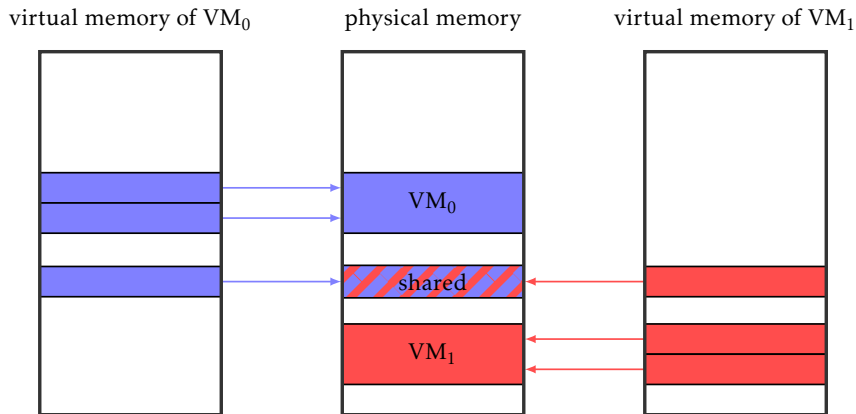
- Controlling memory access of VMs is crucial for isolation

physical memory



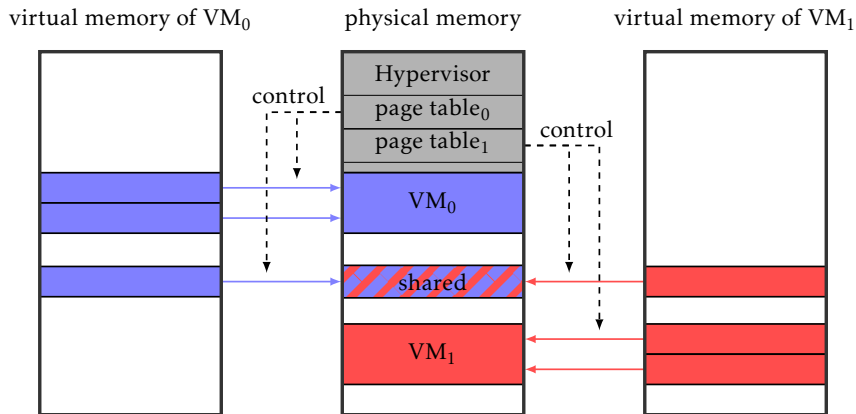
Memory Management in Hypervisors

- Controlling memory access of VMs is crucial for isolation
- Access control is implemented by address translation



Memory Management in Hypervisors

- Controlling memory access of VMs is crucial for isolation
- Access control is implemented by address translation
- Page tables of VMs are managed by the hypervisor



Verifying Communicating VMs

Separation logic nicely captures domain concepts:

Hypervisor	Separation Logic
Communicating VMs	Cooperative threads
Permissions: access, share, ...	Ownership
Sharing memory pages	Transferring ownership
Memory isolation	Separation
Containment of unknown code	Logical relation

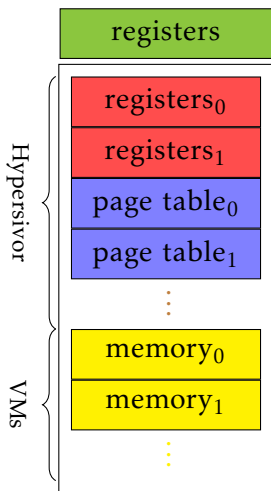
Contributions

- An **operational semantics** for the combination of a machine and a hypervisor
 - With hypercalls (hvc) based on Google's Hafnium hypervisor: run, yield, share, reclaim, etc.
- A **program logic** for modular reasoning about VMs with communication
- A **logical relation** for robust safety property¹
 - Allows us to verify VMs interacting with arbitrary untrusted VMs

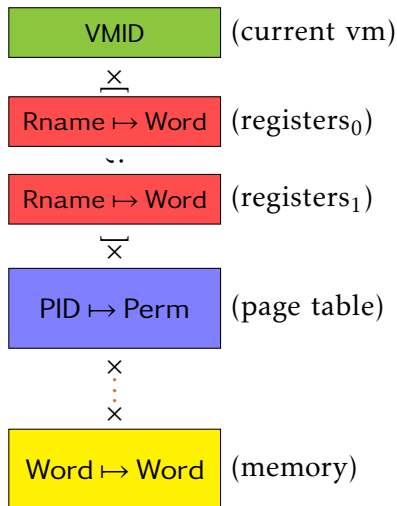
¹[Swasey et al. 2017]

Abstraction of Real Machines

Real machines



Our abstraction



Program Logic for Reasoning about VMs

Resources for machine state

- Registers are indexed by VMID: $\text{PC } @0 \xrightarrow{\text{reg}} a$
- Points-to for page table access: $\text{Pgt } @0 \xrightarrow{\text{acc}} \{p\}$
- Regular points-to for memory cells: $a \xrightarrow{\text{mem}} w$

A variant of weakestpre:

$$\text{WP } m @i \{\Phi\}$$

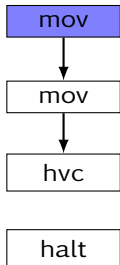
$m \in \text{Mode} \triangleq \text{ExecI} \mid \text{HaIt} \mid \text{PageFault} \mid \dots^2$ $i \in \text{VMID}$

Support for partial evaluation:

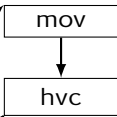
$$\text{WP } m @i \{\Phi\} \dashv\vdash \text{SSWP } m @i \{m', \text{WP } m' @i \{\Phi\}\}$$

²from Cerise

VM₀
(primary)



VM₁
(secondary)



PC @0 $\xrightarrow{\text{reg}}$ $a_{pc0} + 0$

* R0 @0 $\xrightarrow{\text{reg}}$ -

* R1 @0 $\xrightarrow{\text{reg}}$ -

* $a_{pc0} + 0$ $\xrightarrow{\text{mem}}$ mov R0 Run

* $a_{pc0} + 1$ $\xrightarrow{\text{mem}}$ mov R1 1

* $a_{pc0} + 2$ $\xrightarrow{\text{mem}}$ hvc

* $a_{pc0} + 3$ $\xrightarrow{\text{mem}}$ halt

* Pgt @0 $\xrightarrow{\text{acc}}$ {pid(a_{pc0})}

regs

program

pgt

PC @1 $\xrightarrow{\text{reg}}$ $a_{pc1} + 0$

* R0 @1 $\xrightarrow{\text{reg}}$ -

* $a_{pc1} + 0$ $\xrightarrow{\text{mem}}$ mov R0 Yield

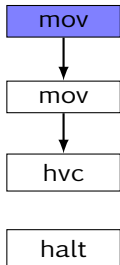
* $a_{pc1} + 1$ $\xrightarrow{\text{mem}}$ hvc

* Pgt @1 $\xrightarrow{\text{acc}}$ {pid(a_{pc1})}

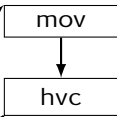
WP ExecI @0 {m,...}

WP ExecI @1 {m,...}

VM₀
(primary)



VM₁
(secondary)



PC @0 $\xrightarrow{\text{reg}}$ $a_{pc0} + 0$

* R0 @0 $\xrightarrow{\text{reg}}$ -

* R1 @0 $\xrightarrow{\text{reg}}$ -

 * $a_{pc0} + 0$ $\xrightarrow{\text{mem}}$ mov R0 Run

* $a_{pc0} + 1$ $\xrightarrow{\text{mem}}$ mov R1 1

* $a_{pc0} + 2$ $\xrightarrow{\text{mem}}$ hvc

* $a_{pc0} + 3$ $\xrightarrow{\text{mem}}$ halt

 * Pgt @0 $\xrightarrow{\text{acc}}$ {pid(a_{pc0})}

PC @1 $\xrightarrow{\text{reg}}$ $a_{pc1} + 0$

* R0 @1 $\xrightarrow{\text{reg}}$ -

 * $a_{pc1} + 0$ $\xrightarrow{\text{mem}}$ mov R0 Yield

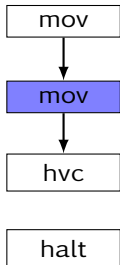
* $a_{pc1} + 1$ $\xrightarrow{\text{mem}}$ hvc

 * Pgt @1 $\xrightarrow{\text{acc}}$ {pid(a_{pc1})}

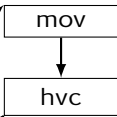
SSWP ExecI @0 {m, WP m @0 {...}}

WP ExecI @1 {m, ...}

VM₀
(primary)



VM₁
(secondary)



PC @0 $\xrightarrow{\text{reg}}$ $a_{pc0} + 1$

* R0 @0 $\xrightarrow{\text{reg}}$ *Run*

* R1 @0 $\xrightarrow{\text{reg}}$ -

* $a_{pc0} + 0$ $\xrightarrow{\text{mem}}$ mov R0 *Run*

* $a_{pc0} + 1$ $\xrightarrow{\text{mem}}$ mov R1 1

* $a_{pc0} + 2$ $\xrightarrow{\text{mem}}$ hvc

* $a_{pc0} + 3$ $\xrightarrow{\text{mem}}$ halt

* Pgt @0 $\xrightarrow{\text{acc}}$ {pid(a_{pc0})}

PC @1 $\xrightarrow{\text{reg}}$ $a_{pc1} + 0$

* R0 @1 $\xrightarrow{\text{reg}}$ -

* $a_{pc1} + 0$ $\xrightarrow{\text{mem}}$ mov R0 *Yield*

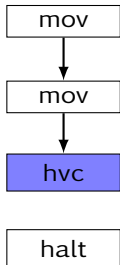
* $a_{pc1} + 1$ $\xrightarrow{\text{mem}}$ hvc

* Pgt @1 $\xrightarrow{\text{acc}}$ {pid(a_{pc1})}

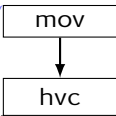
WP ExecI @0 {*m*,...}

WP ExecI @1 {*m*,...}

VM₀
(primary)



VM₁
(secondary)



RO @0, R1 @0

PC @0 $\xrightarrow{\text{reg}}$ $a_{pc0} + 2$

* RO @0 $\xrightarrow{\text{reg}}$ Run

* R1 @0 $\xrightarrow{\text{reg}}$ 1

* $a_{pc0} + 0$ $\xrightarrow{\text{mem}}$ mov R0 Run

* $a_{pc0} + 1$ $\xrightarrow{\text{mem}}$ mov R1 1

* $a_{pc0} + 2$ $\xrightarrow{\text{mem}}$ hvc

* $a_{pc0} + 3$ $\xrightarrow{\text{mem}}$ halt

* Pgt @0 $\xrightarrow{\text{acc}}$ {pid(a_{pc0})}

PC @1 $\xrightarrow{\text{reg}}$ $a_{pc1} + 0$

* RO @1 $\xrightarrow{\text{reg}}$ -

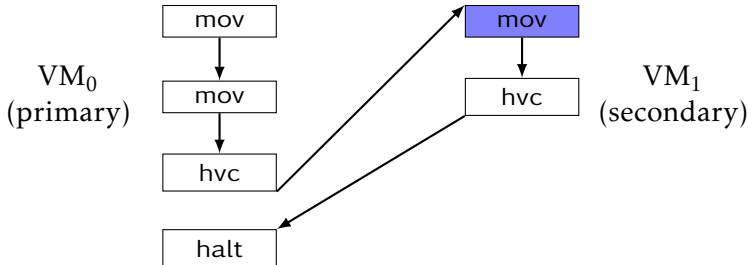
* $a_{pc1} + 0$ $\xrightarrow{\text{mem}}$ mov R0 Yield

* $a_{pc1} + 1$ $\xrightarrow{\text{mem}}$ hvc

* Pgt @1 $\xrightarrow{\text{acc}}$ {pid(a_{pc1})}

WP ExecI @0 {m,...}

WP ExecI @1 {m,...}



PC @0 $\xrightarrow{\text{reg}}$ $a_{pc0} + 3$

 * $a_{pc0} + 0 \xrightarrow{\text{mem}}$ mov R0 Run
 * $a_{pc0} + 1 \xrightarrow{\text{mem}}$ mov R1 1
 * $a_{pc0} + 2 \xrightarrow{\text{mem}}$ hvc
 * $a_{pc0} + 3 \xrightarrow{\text{mem}}$ halt

 * Pgt @0 $\xrightarrow{\text{acc}}$ {pid(a_{pc0})}

WP ExecI @0 { m, \dots }

PC @1 $\xrightarrow{\text{reg}}$ $a_{pc1} + 0$

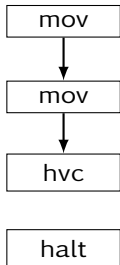
* R0 @1 $\xrightarrow{\text{reg}}$ -
 * R0 @0 $\xrightarrow{\text{reg}}$ Run
 * R1 @0 $\xrightarrow{\text{reg}}$ 1

 * $a_{pc1} + 0 \xrightarrow{\text{mem}}$ mov R0 Yield
 * $a_{pc1} + 1 \xrightarrow{\text{mem}}$ hvc

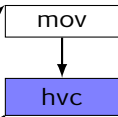
 * Pgt @1 $\xrightarrow{\text{acc}}$ {pid(a_{pc1})}

WP ExecI @1 { m, \dots }

VM₀
(primary)



VM₁
(secondary)



PC @0 $\xrightarrow{\text{reg}}$ $a_{pc0} + 3$

 * $a_{pc0} + 0 \xrightarrow{\text{mem}}$ mov R0 Run
 * $a_{pc0} + 1 \xrightarrow{\text{mem}}$ mov R1 1
 * $a_{pc0} + 2 \xrightarrow{\text{mem}}$ hvc
 * $a_{pc0} + 3 \xrightarrow{\text{mem}}$ halt

 * Pgt @0 $\xrightarrow{\text{acc}}$ {pid(a_{pc0})}

WP ExecI @0 {m,...}

PC @1 $\xrightarrow{\text{reg}}$ $a_{pc1} + 1$

* R0 @1 $\xrightarrow{\text{reg}}$ Yield

* R0 @0 $\xrightarrow{\text{reg}}$ Run

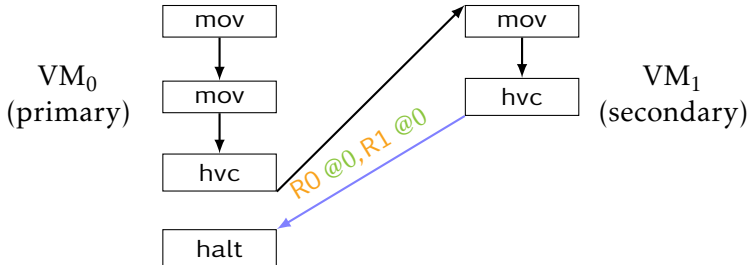
* R1 @0 $\xrightarrow{\text{reg}}$ 1

 * $a_{pc1} + 0 \xrightarrow{\text{mem}}$ mov R0 Yield

* $a_{pc1} + 1 \xrightarrow{\text{mem}}$ hvc

 * Pgt @1 $\xrightarrow{\text{acc}}$ {pid(a_{pc1})}

WP ExecI @1 {m,...}



PC @0 $\xrightarrow{\text{reg}}$ $a_{pc0} + 3$

 * $a_{pc0} + 0 \xrightarrow{\text{mem}} \text{mov R0 Run}$
 * $a_{pc0} + 1 \xrightarrow{\text{mem}} \text{mov R1 1}$
 * $a_{pc0} + 2 \xrightarrow{\text{mem}} \text{hvc}$
 * $a_{pc0} + 3 \xrightarrow{\text{mem}} \text{halt}$

 * Pgt @0 $\xrightarrow{\text{acc}} \{\text{pid}(a_{pc0})\}$

WP ExecI @0 {m,...}

PC @1 $\xrightarrow{\text{reg}}$ $a_{pc1} + 2$

* R0 @1 $\xrightarrow{\text{reg}} \text{Yield}$

* R0 @0 $\xrightarrow{\text{reg}} \text{Yield}$

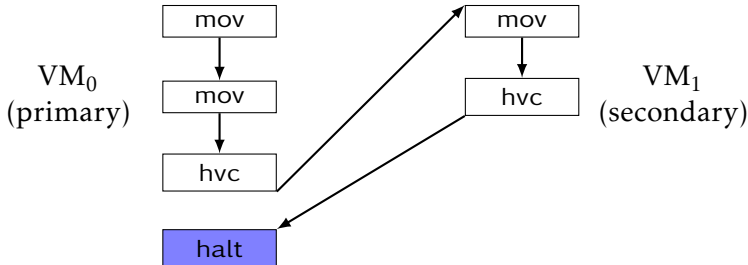
* R1 @0 $\xrightarrow{\text{reg}} 1$

 * $a_{pc1} + 0 \xrightarrow{\text{mem}} \text{mov R0 Yield}$

* $a_{pc1} + 1 \xrightarrow{\text{mem}} \text{hvc}$

 * Pgt @1 $\xrightarrow{\text{acc}} \{\text{pid}(a_{pc1})\}$

WP ExecI @1 {m,...}



PC @0 $\xrightarrow{\text{reg}}$ $a_{pc0} + 3$

* R0 @0 $\xrightarrow{\text{reg}}$ *Yield*

* R1 @0 $\xrightarrow{\text{reg}}$ 1

 * $a_{pc0} + 0 \xrightarrow{\text{mem}}$ mov R0 Run

* $a_{pc0} + 1 \xrightarrow{\text{mem}}$ mov R1 1

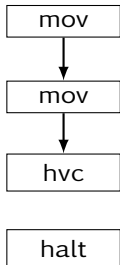
* $a_{pc0} + 2 \xrightarrow{\text{mem}}$ hvc

* $a_{pc0} + 3 \xrightarrow{\text{mem}}$ halt

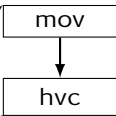
 * Pgt @0 $\xrightarrow{\text{acc}}$ {pid(a_{pc0})}

WP ExecI @0 {m,...}

VM₀
(primary)



VM₁
(secondary)



PC @0 $\xrightarrow{\text{reg}}$ $a_{pc0} + 4$

* R0 @0 $\xrightarrow{\text{reg}}$ *Yield*

* R1 @0 $\xrightarrow{\text{reg}}$ 1

* $a_{pc0} + 0 \xrightarrow{\text{mem}}$ mov R0 Run

* $a_{pc0} + 1 \xrightarrow{\text{mem}}$ mov R1 1

* $a_{pc0} + 2 \xrightarrow{\text{mem}}$ hvc

* $a_{pc0} + 3 \xrightarrow{\text{mem}}$ halt

* Pgt @0 $\xrightarrow{\text{acc}}$ {pid(a_{pc0})}

WP Halt @0 {m,...}

Resumption Conditions - Session-type Like Protocols

Allow us to specify what is needed to resume the execution of a VM, and transfer resources accordingly

$$\begin{array}{l} \text{ResumCond } @1 \ 1/2 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Run} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \\ * \text{ResumCond } @0 \ 1/2 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \end{array} \right\} \end{array} \right\} \\ * \text{ResumCond } @0 \ 1 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \end{array} \right\} \end{array}$$

Resumption Conditions - Session-type Like Protocols

$$\text{ResumCond } @1 \text{ } 1/2 \left\{ \begin{array}{l} R0 @0 \xrightarrow{\text{reg}} \text{Run} \\ * R1 @0 \xrightarrow{\text{reg}} 1 \\ * \text{ResumCond } @0 \text{ } 1/2 \left\{ \begin{array}{l} R0 @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 @0 \xrightarrow{\text{reg}} 1 \end{array} \right\} \end{array} \right\}$$
$$* \text{ResumCond } @0 \text{ } 1 \left\{ \begin{array}{l} R0 @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 @0 \xrightarrow{\text{reg}} 1 \end{array} \right\}$$

$$* R0 @0 \xrightarrow{\text{reg}} \text{Run}$$

$$* R1 @0 \xrightarrow{\text{reg}} 1$$

$$* a_{pc0} + 0 \xrightarrow{\text{mem}} \text{mov R0 Run}$$

$$* a_{pc0} + 1 \xrightarrow{\text{mem}} \text{mov R1 1}$$

$$* a_{pc0} + 2 \xrightarrow{\text{mem}} \text{hvc}$$

$$* a_{pc0} + 3 \xrightarrow{\text{mem}} \text{halt}$$

WP ExecI @0 {m,...}

Resumption Conditions - Session-type Like Protocols

$$\text{ResumCond } @0 \text{ } 1/2 \left\{ \begin{array}{l} \text{R0 } @0 \xrightarrow{\text{reg}} \text{Yield} \\ * \text{R1 } @0 \xrightarrow{\text{reg}} 1 \end{array} \right\}$$

* $a_{pc0} + 0 \xrightarrow{\text{mem}} \text{mov R0 Run}$

* $a_{pc0} + 1 \xrightarrow{\text{mem}} \text{mov R1 1}$

* $a_{pc0} + 2 \xrightarrow{\text{mem}} \text{hvc}$

* $a_{pc0} + 3 \xrightarrow{\text{mem}} \text{halt}$

$$\frac{\text{ResumCondHolds } @i * \text{ResumCond } @i \text{ } 1/2 \{P\}}{\text{ResumCond } @i \text{ } 1 \{P\} * \triangleright P}$$

ResumCondHolds @0 * WP ExecI @0 {m,...}

Resumption Conditions - Session-type Like Protocols

$$\text{ResumCond } @0 \ 1 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \end{array} \right\}$$
$$* R0 \ @0 \xrightarrow{\text{reg}} \text{Yield}$$
$$* R1 \ @0 \xrightarrow{\text{reg}} 1$$
$$* a_{pc0} + 0 \xrightarrow{\text{mem}} \text{mov } R0 \ \text{Run}$$
$$* a_{pc0} + 1 \xrightarrow{\text{mem}} \text{mov } R1 \ 1$$
$$* a_{pc0} + 2 \xrightarrow{\text{mem}} \text{hvc}$$
$$* a_{pc0} + 3 \xrightarrow{\text{mem}} \text{halt}$$

$\frac{\text{ResumCondHolds } @i * \text{ResumCond } @i \ 1/2 \ \{P\}}{\text{ResumCond } @i \ 1 \ \{P\} * \triangleright P}$
--

$$\text{WP ExecI } @0 \ \{m, \dots\}$$

Resumption Conditions - Session-type Like Protocols

$$\text{ResumCond } @1 \text{ } 1/2 \left\{ \begin{array}{l} \text{R0 } @0 \xrightarrow{\text{reg}} \text{Run} \\ * \text{R1 } @0 \xrightarrow{\text{reg}} 1 \\ * \text{ResumCond } @0 \text{ } 1/2 \left\{ \begin{array}{l} \text{R0 } @0 \xrightarrow{\text{reg}} \text{Yield} \\ * \text{R1 } @0 \xrightarrow{\text{reg}} 1 \end{array} \right\} \end{array} \right\}$$

$$\begin{array}{l} * \text{R0 } @1 \xrightarrow{\text{reg}} - \\ * a_{pc0} + 0 \xrightarrow{\text{mem}} \text{mov R0 Yield} \\ * a_{pc0} + 1 \xrightarrow{\text{mem}} \text{hvc} \end{array}$$

$\frac{\text{ResumCondHolds } @i * \text{ResumCond } @i \text{ } 1/2 \{P\}}{\text{ResumCond } @i \text{ } 1 \{P\} * \triangleright P}$
--

$$\text{ResumCondHolds } @1 \text{ } \rightarrow * \text{WP ExecI } @1 \{m, \dots\}$$

Resumption Conditions - Session-type Like Protocols

$$\text{ResumCond } @1 \ 1 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Run} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \\ * \text{ResumCond } @0 \ 1/2 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \end{array} \right\} \end{array} \right\}$$

$$* \triangleright \text{ResumCond } @0 \ 1/2 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \end{array} \right\}$$

$$* R0 \ @0 \xrightarrow{\text{reg}} \text{Run}$$

$$* R1 \ @0 \xrightarrow{\text{reg}} 1$$

$$* R0 \ @1 \xrightarrow{\text{reg}} -$$

$$* a_{pc0} + 0 \xrightarrow{\text{mem}} \text{mov } R0 \ \text{Yield}$$

$$* a_{pc0} + 1 \xrightarrow{\text{mem}} \text{hvc}$$

$\frac{\text{ResumCondHolds } @i \ * \ \text{ResumCond } @i \ 1/2 \ \{P\}}{\text{ResumCond } @i \ 1 \ \{P\} \ * \ \triangleright \ P}$
--

WP ExecI @1 {m,...}

Resumption Conditions - Session-type Like Protocols

$$\text{ResumCond } @1 \ 1 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Run} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \\ * \text{ResumCond } @0 \ 1/2 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \end{array} \right\} \end{array} \right\}$$

$$* \text{ResumCond } @0 \ 1/2 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \end{array} \right\}$$

$$* R0 \ @0 \xrightarrow{\text{reg}} \text{Run}$$

$$* R1 \ @0 \xrightarrow{\text{reg}} 1$$

$$* R0 \ @1 \xrightarrow{\text{reg}} \text{Yield}$$

$$* a_{pc0} + 0 \xrightarrow{\text{mem}} \text{mov } R0 \ \text{Yield}$$

$$* a_{pc0} + 1 \xrightarrow{\text{mem}} \text{hvc}$$

WP ExecI @1 {m,...}

Resumption Conditions - Session-type Like Protocols

$$\begin{aligned} & \text{ResumCond } @1 \ 1 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Run} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \\ * \text{ResumCond } @0 \ 1/2 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \end{array} \right\} \end{array} \right\} \\ & * \text{ResumCond } @0 \ 1/2 \left\{ \begin{array}{l} R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ * R1 \ @0 \xrightarrow{\text{reg}} 1 \end{array} \right\} \\ & * R0 \ @0 \xrightarrow{\text{reg}} \text{Yield} \\ & * R1 \ @0 \xrightarrow{\text{reg}} 1 \\ & * R0 \ @1 \xrightarrow{\text{reg}} \text{Yield} \\ & * a_{pc0} + 0 \xrightarrow{\text{mem}} \text{mov } R0 \ \text{Yield} \\ & * a_{pc0} + 1 \xrightarrow{\text{mem}} \text{hvc} \end{aligned}$$

WP ExecI @1 {m,...}

Resumption Conditions - Session-type Like Protocols

$$\text{ResumCond } @1 \ 1 \left\{ \begin{array}{l} \text{R0 } @0 \xrightarrow{\text{reg}} \text{Run} \\ * \text{R1 } @0 \xrightarrow{\text{reg}} 1 \\ * \text{ResumCond } @0 \ 1/2 \left\{ \begin{array}{l} \text{R0 } @0 \xrightarrow{\text{reg}} \text{Yield} \\ * \text{R1 } @0 \xrightarrow{\text{reg}} 1 \end{array} \right\} \end{array} \right\}$$

$$\begin{array}{l} * \text{R0 } @1 \xrightarrow{\text{reg}} \text{Yield} \\ * a_{pc0} + 0 \xrightarrow{\text{mem}} \text{mov R0 Yield} \\ * a_{pc0} + 1 \xrightarrow{\text{mem}} \text{hvc} \end{array}$$

$\frac{\text{ResumCondHolds } @i * \text{ResumCond } @i \ 1/2 \ \{P\}}{\text{ResumCond } @i \ 1 \ \{P\} * \triangleright P}$
--

$\text{ResumCondHolds } @1 \ \rightarrow * \text{WP ExecI } @1 \ \{m, \dots\}$

What If We Don't Know Code of VM₁?

Robust safety: arbitrary unknown code won't break known VMs:
it can only change memory it has (or can get) access to

$$\frac{? * \text{ResumCond } @1 \text{ }^{1/2} \{?\}}{\text{ResumCondHolds } @1 \text{ } \rightarrow * \text{ WP ExecI } @1 \{m, \dots\}}$$

What If We Don't Know Code of VM₁?

Robust safety: arbitrary unknown code won't break known VMs:
it can only change memory it has (or can get) access to

$$\frac{\text{Owned}(pgt) * \text{ResumCond } @1 \text{ } 1/2 \text{ } \{\text{Protocol}(pgt)\}}{\text{ResumCondHolds } @1 \text{ } -* \text{ WP ExecI } @1 \text{ } \{m, \dots\}}$$

We define a logical relation that enforces **robust safety**

- No assumptions on contents of memory, only on the page tables
- Challenge is to consider all possible interactions, incl. in-flight memory sharing, etc.

What If We Don't Know Code of VM₁?

Robust safety: arbitrary unknown code won't break known VMs:
it can only change memory it has (or can get) access to

FUNDAMENTAL THEOREM

$$\text{Owned}(pgt) * \text{ResumCond } @i \text{ } 1/2 \{ \text{Protocol}(pgt) \}$$

$$\text{ResumCondHolds } @i \text{ } -* \text{ WP ExecI } @i \{ m, \top \}$$

We define a logical relation that enforces **robust safety**

- No assumptions on contents of memory, only on the page tables
- Challenge is to consider all possible interactions, incl. in-flight memory sharing, etc.

Conclusion

Contributions

- An **operational semantics** for the combination of a machine and a hypervisor
- A **program logic** for modular reasoning about VMs with communication
 - Verified the run & yield example with resumption conditions (saved ~200 LOC comparing to invariants)
- A **logical relation** for robust safety property
 - The run & yield example with an unknown VM (~80 more LOC)

Next Steps

- Adding interrupts/exceptions - crucial for scheduling
- Considering multi-core CPUs - concurrency
- Non-interference