Fully Composable and Adequate Verified Compilation with Direct Refinements between Open Modules

Ling Zhang [1] **Yuting Wang** [1] Jinhua Wu [1] Jérémie Koenig [2] Zhong Shao [2]

[1]Shanghai Jiao Tong University, China

[2]Yale University, USA

POPL, January 2024, London

Verified Compilation

CompCert: the State-of-the-Art

- Verified compilation of a subset of C into assembly in Coq
- Many Applications: CertiKOS, VST, Critical control system, etc.

Compiler Correctness = Refinement of Semantics

- $[\![M_1]\!] \preccurlyeq [\![M_2]\!]$ denotes the semantics of M_1 refines that of M_2

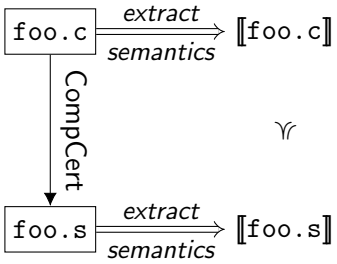

Verified Compilation

CompCert: the State-of-the-Art

- Verified compilation of a subset of C into assembly in Coq
- Many Applications: CertiKOS, VST, Critical control system, etc.

Compiler Correctness = Refinement of Semantics

- $[\![M_1]\!] \preccurlyeq [\![M_2]\!]$ denotes the semantics of M_1 refines that of M_2

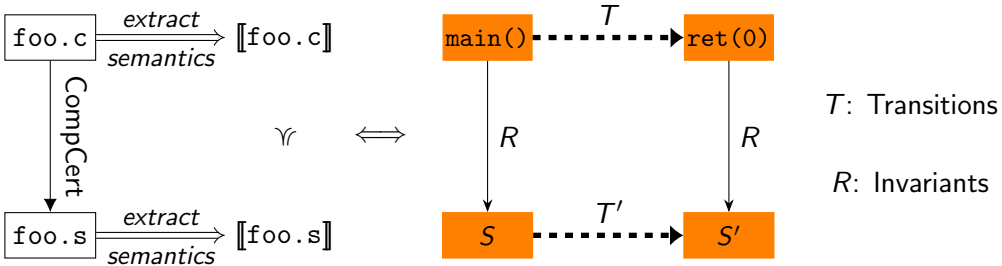

Verified Compilation of Open Modules

Intuition: Get a refinement directly relating semantics of C and assembly modules.

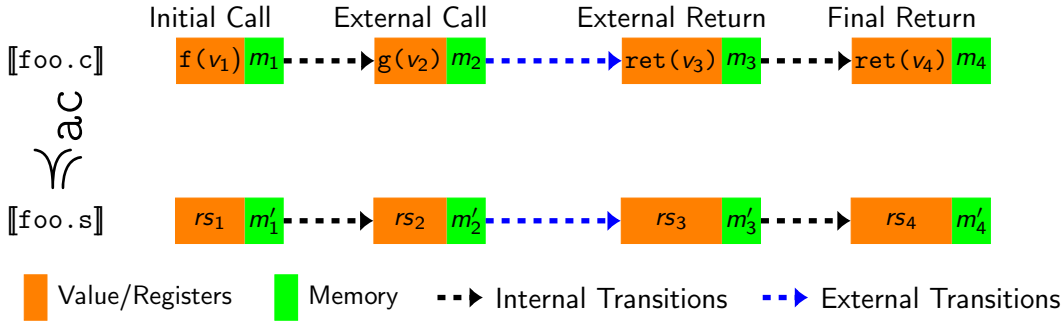

Verified Compilation of Open Modules

Intuition: Get a refinement directly relating semantics of C and assembly modules.

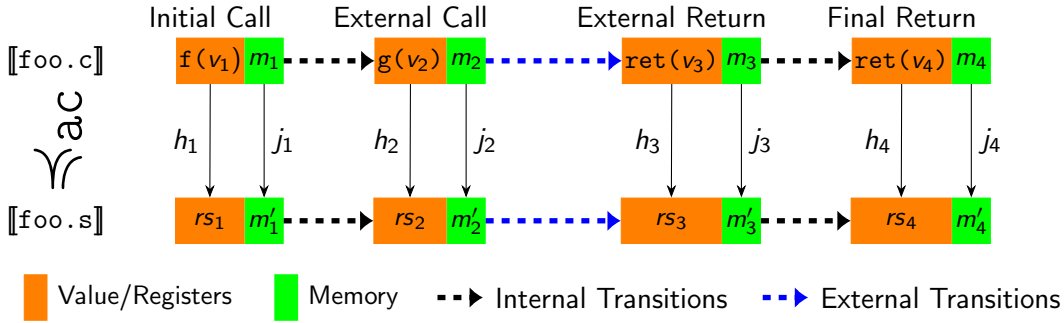

Problem: Indirect Refinements

Observation: No existing work on CompCert produces direct refinement

Compositional CompCert [*POPL'15*]: $[\![\texttt{foo.s}]\!] \quad \preccurlyeq_{\mathcal{C}} \quad [\![\texttt{foo.c}]\!]$

CompCertM [*POPL'20*]: $[\![\texttt{foo.s}]\!] \preccurlyeq_1 + \preccurlyeq_2 + \ldots + \preccurlyeq_n [\![\texttt{foo.c}]\!]$

CompCertO [*PLDI'21*]: $[\![\texttt{foo.s}]\!] \quad \preccurlyeq_1 \cdot \preccurlyeq_2 \cdot \ldots \cdot \preccurlyeq_n \quad [\![\texttt{foo.c}]\!]$

Problem: Indirect Refinements

Observation: No existing work on CompCert produces direct refinement

Compositional CompCert [*POPL'15*]: $[\![\texttt{foo.s}]\!]$ $\preccurlyeq_C$ $[\![\texttt{foo.c}]\!]$

Only with C interfaces

CompCertM [*POPL'20*]: $[\![\texttt{foo.s}]\!] \preccurlyeq_1 + \preccurlyeq_2 + \ldots + \preccurlyeq_n [\![\texttt{foo.c}]\!]$

CompCertO [*PLDI'21*]: $[\![\texttt{foo.s}]\!] \preccurlyeq_1 \cdot \preccurlyeq_2 \cdot \ldots \cdot \preccurlyeq_n [\![\texttt{foo.c}]\!]$

Problem: Indirect Refinements

Observation: No existing work on CompCert produces direct refinement

Compositional CompCert [*POPL'15*]: $[\![\texttt{foo.s}]\!]$ $\preccurlyeq_C$ $[\![\texttt{foo.c}]\!]$

Only with C interfaces

CompCertM [*POPL'20*]: $[\![\texttt{foo.s}]\!]$ $\preccurlyeq_1 + \preccurlyeq_2 + \ldots + \preccurlyeq_n$ $[\![\texttt{foo.c}]\!]$

Union of Refinements

CompCertO [*PLDI'21*]: $[\![\texttt{foo.s}]\!]$ $\preccurlyeq_1 \cdot \preccurlyeq_2 \cdot \ldots \cdot \preccurlyeq_n$ $[\![\texttt{foo.c}]\!]$

Problem: Indirect Refinements

Observation: No existing work on CompCert produces direct refinement

Compositional CompCert [*POPL'15*]: $[\![\texttt{foo.s}]\!]$ $\preccurlyeq_C$ $[\![\texttt{foo.c}]\!]$

Only with C interfaces

CompCertM [*POPL'20*]: $[\![\texttt{foo.s}]\!]$ $\preccurlyeq_1 + \preccurlyeq_2 + \ldots + \preccurlyeq_n$ $[\![\texttt{foo.c}]\!]$

Union of Refinements

CompCertO [*PLDI'21*]: $[\![\texttt{foo.s}]\!]$ $\preccurlyeq_1 \cdot \preccurlyeq_2 \cdot \ldots \cdot \preccurlyeq_n$ $[\![\texttt{foo.c}]\!]$

Concatenation of Refinements

Main Challenge

Challenge: Vertical composition of direct refinements is difficult

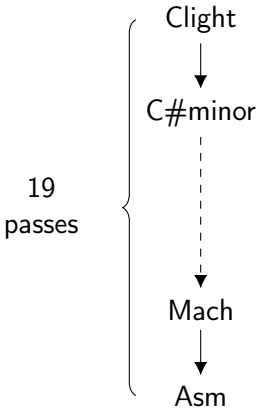

Main Challenge

Challenge: Vertical composition of direct refinements is difficult

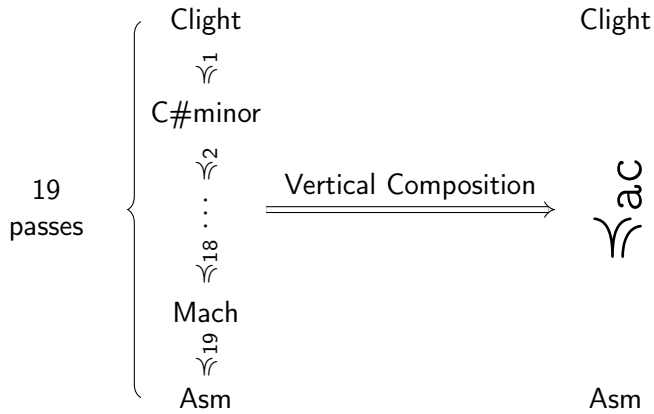

Main Challenge

Challenge: Vertical composition of direct refinements is difficult

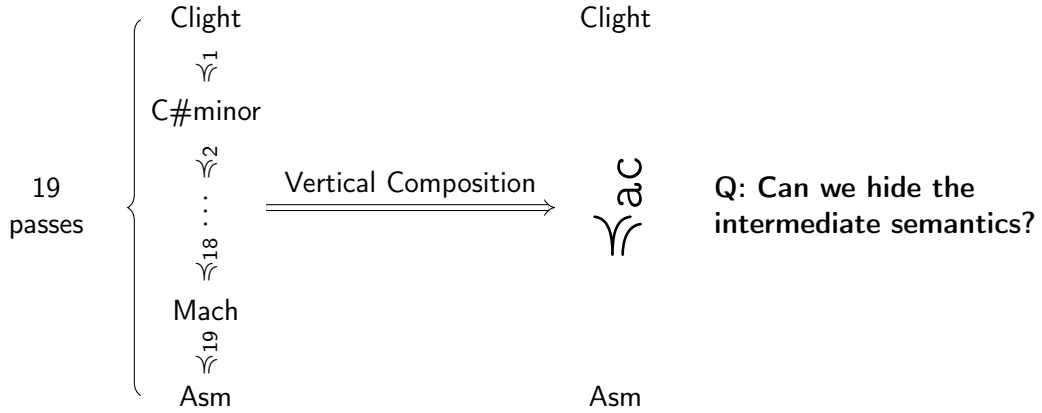

Q: Can we hide the intermediate semantics?

Our Contributions

Approach to Direct Refinements Supporting:

- Vertical and horizontal composition
- Equivalence of semantics and syntactic linking (i.e., Adequacy)
- Heterogeneous modules with mutual calls

Applications:

- CompCert's full compilation chain
- Extension to user-level verification

Notice: We focus on imperative programs with global memory and pointers.

A Running Example

Heterogeneous Modules with Callbacks and Pointer Passing:

- A client written in C;
- An encryption server written in X86 assembly.

```c
/* Client.c */
int result;
void encrypt(int i, void(*p)(int*));

static void rcd(int *r) {
  result = *r;
}
// Entry point
int request(int i) {
  encrypt(i,rcd);
  return result;
}
```

```
/* Server.s */
key: .long 42
encrypt:
  ... // Alloc 24-bytes frame
  // RSP[8] = key XOR i
  mov key RAX
  xor RAX RDI
  mov RDI 8(RSP)
  // call p(RSP + 8)
  lea 8(RSP) RDI
  call RSI
  ...
```

A Running Example

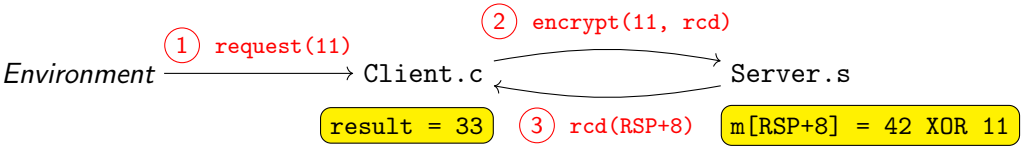


```c
/* Client.c */
int result;
void encrypt(int i, void(*p)(int*));

static void rcd(int *r) {
  result = *r;
}
// Entry point
int request(int i) {
  encrypt(i,rcd);
  return result;
}
```

```asm
/* Server.s */
key: .long 42
encrypt:
  ... // Alloc 24-bytes frame
  // RSP[8] = key XOR i
  mov key RAX
  xor RAX RDI
  mov RDI 8(RSP)
  // call p(RSP + 8)
  lea 8(RSP) RDI
  call RSI
  ...
```

Verification Steps

(1) Prove CompCert has the direct refinement $\preccurlyeq_{\mathtt{ac}}$;

(2) Prove $[\![\mathtt{Server.s}]\!] \preccurlyeq_{\mathtt{ac}} L_{\mathtt{S}}$;

(3) Exploit the compositionality and adequacy of $\preccurlyeq_{\mathtt{ac}}$.

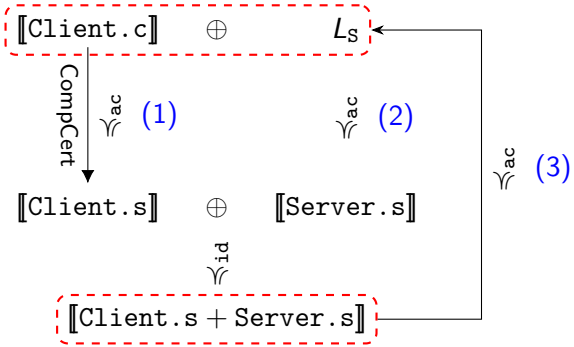

Key Ideas

1. Direct Refinements for Adequacy and Horizontal Composition
2. Transitive Kripke Memory Relation for Vertical Composition

Direct Refinement for CompCert

Direct refinement $\preccurlyeq_{\mathsf{ac}}$ as forward simulation with

- Invariant for source and target program states;

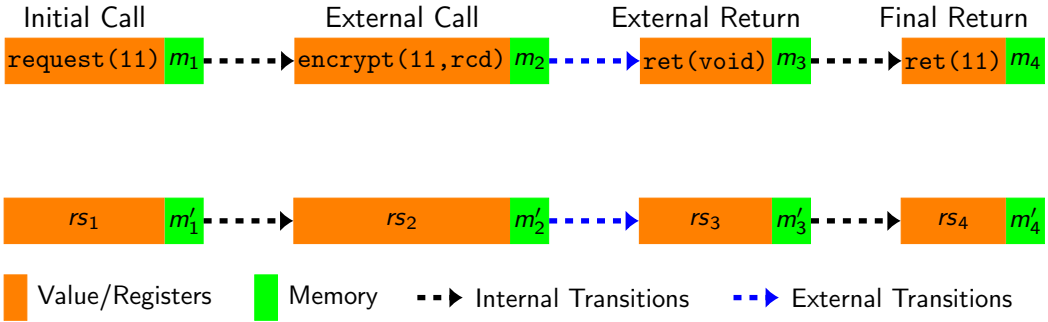

$$[\![\texttt{Client.s}]\!] \preccurlyeq_{\mathsf{ac}} [\![\texttt{Client.c}]\!]$$

Direct Refinement for CompCert

Direct refinement $\preccurlyeq_{\mathsf{ac}}$ as forward simulation with

- Invariant for source and target program states;

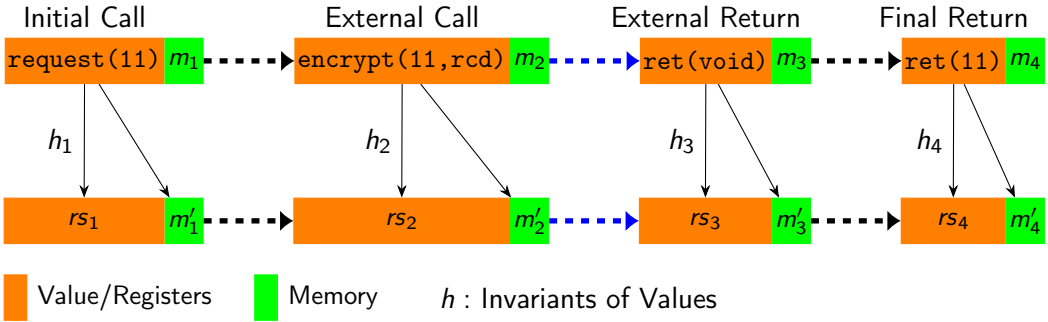

$$[\![\texttt{Client.s}]\!] \preccurlyeq_{\mathsf{ac}} [\![\texttt{Client.c}]\!]$$

Direct Refinement for CompCert

Direct refinement $\preccurlyeq_{\mathsf{ac}}$ as forward simulation with

- Invariant for source and target program states;

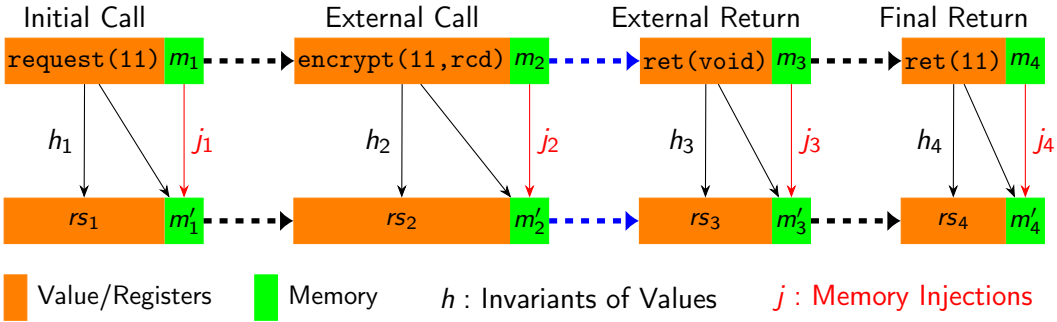

$$[\![\texttt{Client.s}]\!] \preccurlyeq_{\mathsf{ac}} [\![\texttt{Client.c}]\!]$$

Direct Refinement for CompCert

Direct refinement $\preccurlyeq_{\mathsf{ac}}$ as forward simulation with

- Invariant for source and target program states;
- Protection for program states across external calls.

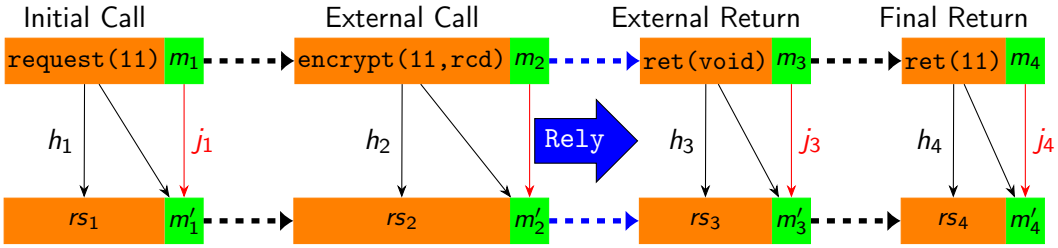

$$[\![\texttt{Client.s}]\!] \preccurlyeq_{\mathsf{ac}} [\![\texttt{Client.c}]\!]$$

Adequacy of Direct Refinements

Adequacy trivially holds as invariants directly relates C and assembly states:

- Invariants formalize the CompCert C calling convention;
- Source function arguments are mapped directly to registers and the stack.

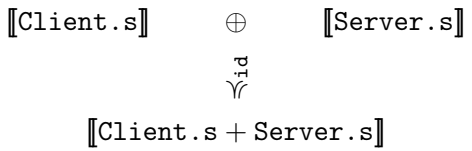

Horizontal Composition of Direct Refinements

Direct protection of private states against external calls :

- Callee-saved registers and stack pointer must be restored upon returning.
- Private stack memory (e.g., spilled registers) must not be modified

Rely-guarantee reasoning

Horizontal Composition of Direct Refinements

Direct protection of private states against external calls :

- Callee-saved registers and stack pointer must be restored upon returning.
- Private stack memory (e.g., spilled registers) must not be modified

Rely-guarantee reasoning

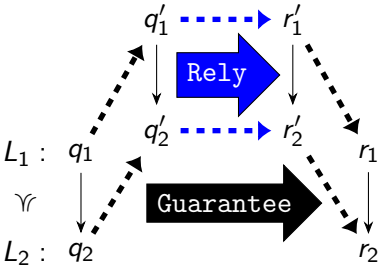

Horizontal Composition of Direct Refinements

Direct protection of private states against external calls :

- Callee-saved registers and stack pointer must be restored upon returning.
- Private stack memory (e.g., spilled registers) must not be modified

Rely-guarantee reasoning

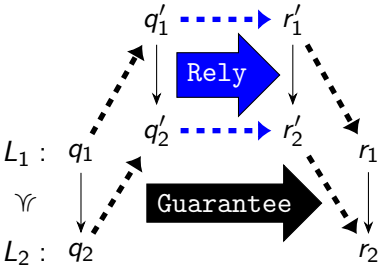

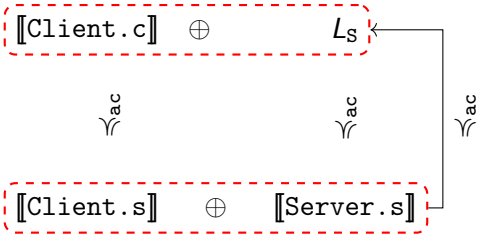

($\preccurlyeq_{ac}$ has symmetric rely and guarantee conditions)

Key Ideas

1. Direct Refinements for Adequacy and Horizontal Composition
2. Transitive Kripke Memory Relation for Vertical Composition

Recall the Challenge

Challenge: Vertical composition of refinements

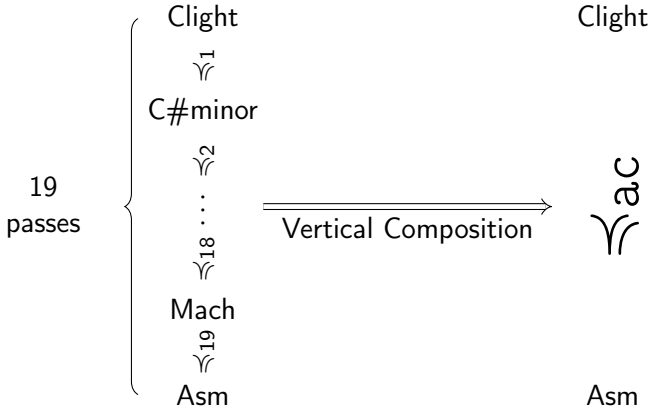

Recall the Challenge

Challenge: Vertical composition of rely-guarantee conditions

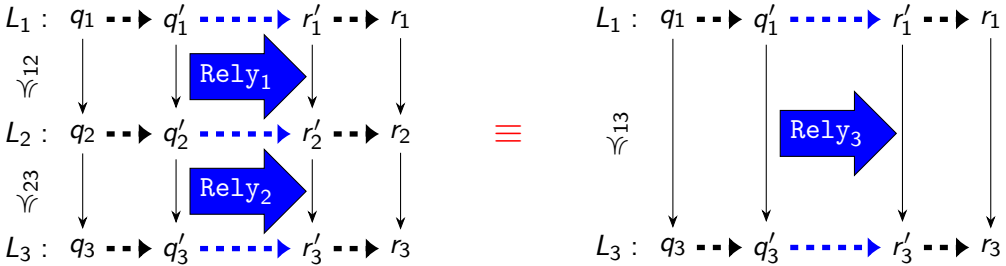

A Kripke Relation with Memory Protection

Kripke relation `injp` for protection:

- At an external call, infer private memory from the injection;
- No modification to private memory allowed during the call.

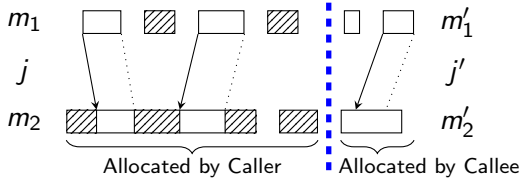

Private memory are the shaded areas, including

- Source caller's memory NOT in the domain of j
- Target caller's memory NOT in the image of j

Example of Memory Protection by `injp`

Before the server calls back `rcd`:

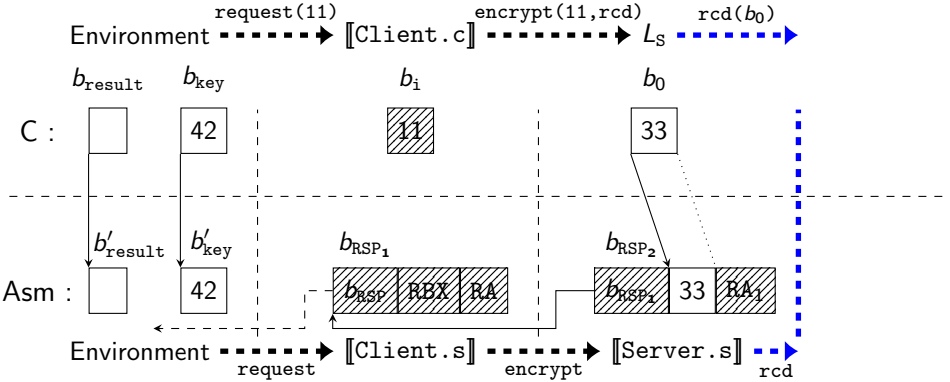

Protected Memory: b_i, $b_{\mathtt{RSP}_1}$, and part of $b_{\mathtt{RSP}_2}$

Example of Memory Protection by `injp`

During the server calls back `rcd`:

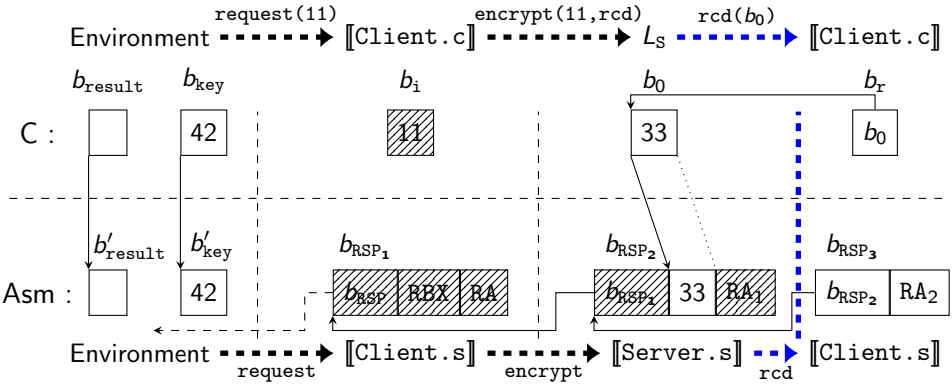

Protected Memory: b_i, b_{RSP_1}, and part of b_{RSP_2}

Example of Memory Protection by `injp`

After the server calls back `rcd`:

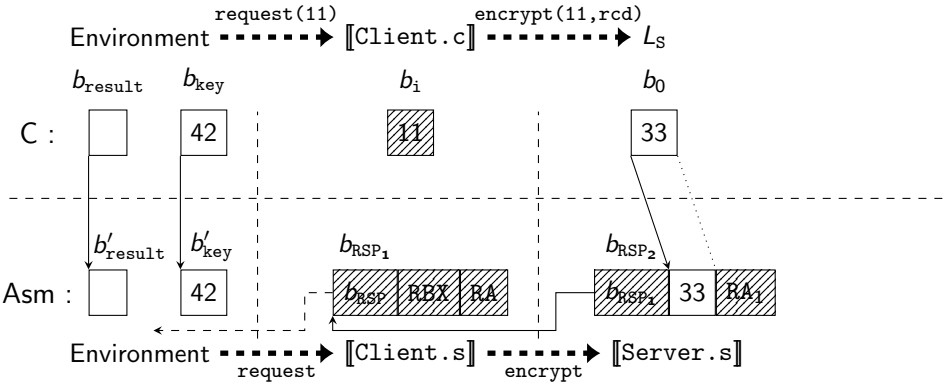

Protected Memory: b_i, b_{RSP_1}, and part of b_{RSP_2}

Vertical Composition of Direct Refinements

Observations:

- `injp` is uniform: its protection works for all passes;
- `injp` is transitive: $\texttt{injp} \cdot \texttt{injp} \equiv \texttt{injp}$.

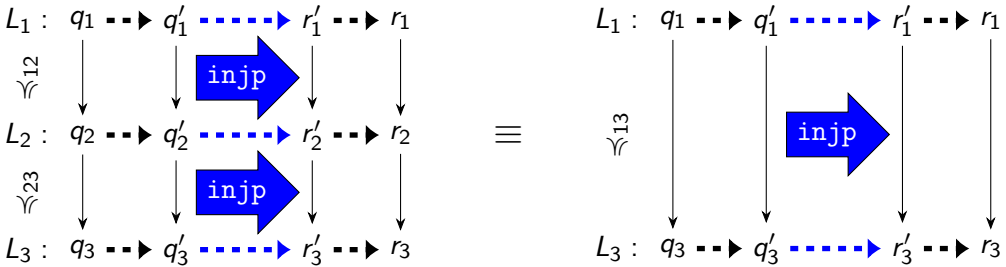

Transitivity of `injp`

Key to prove $\texttt{injp} \cdot \texttt{injp} \equiv \texttt{injp}$:

- Construct an interpolating state when the external call returns.

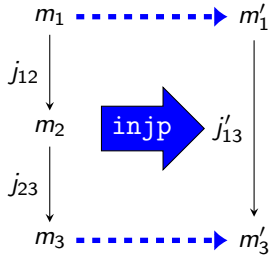

Transitivity of `injp`

Key to prove $\texttt{injp} \cdot \texttt{injp} \equiv \texttt{injp}$:

- Construct an interpolating state when the external call returns.

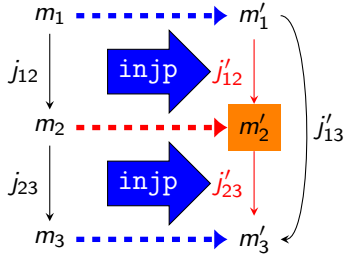

Transitivity of `injp`

Key to prove $\texttt{injp} \cdot \texttt{injp} \equiv \texttt{injp}$:

- Construct an interpolating state when the external call returns.

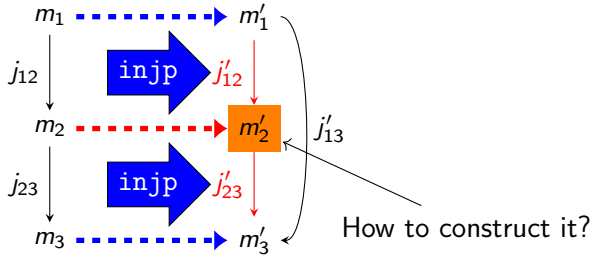

Construction of the Interpolating State

Protection after Composition $\geqslant$ Protection before Composition

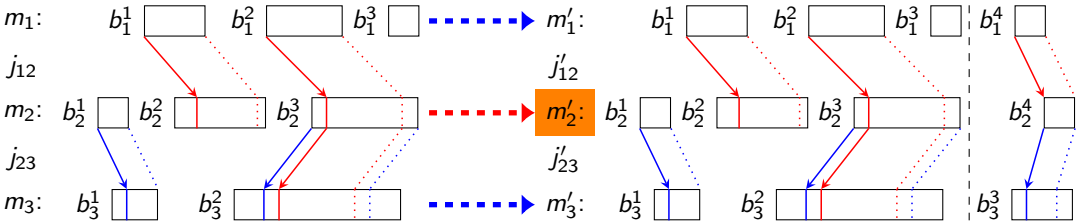

- Public memory of m_2' = (Image of j_{12}) $\cap$ (Domain of j_{23});
- m_2' is unchanged from m_2 except for its public memory is projected from m_1'.

Construction of the Interpolating State

Protection after Composition $\geqslant$ Protection before Composition

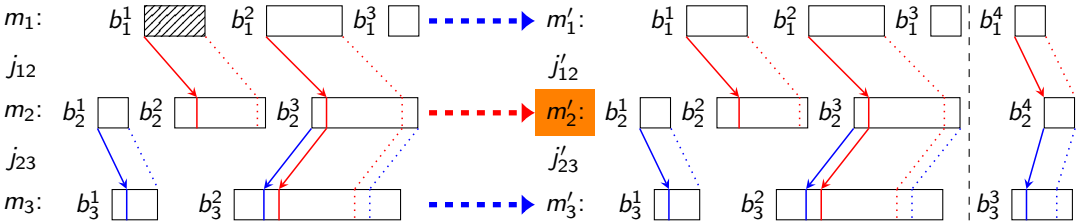

- Public memory of m_2' = (Image of j_{12}) $\cap$ (Domain of j_{23});
- m_2' is unchanged from m_2 except for its public memory is projected from m_1'.

Construction of the Interpolating State

Protection after Composition $\geqslant$ Protection before Composition

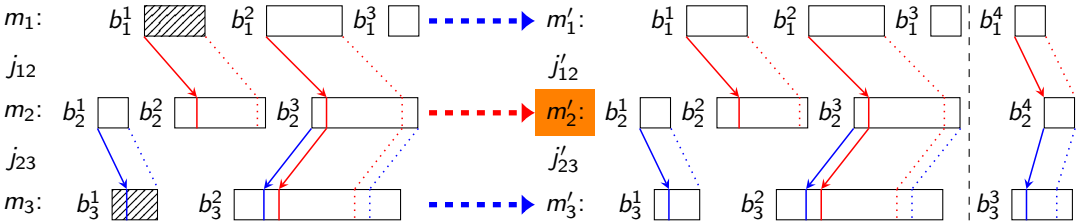

- Public memory of m_2' = (Image of j_{12}) $\cap$ (Domain of j_{23});
- m_2' is unchanged from m_2 except for its public memory is projected from m_1'.

Construction of the Interpolating State

Protection after Composition $\geqslant$ Protection before Composition

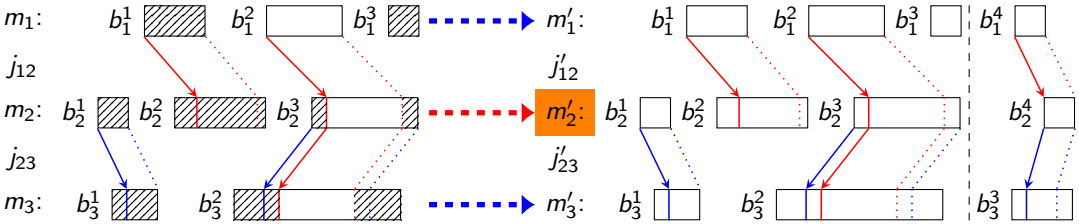

- Public memory of m_2' = (Image of j_{12}) $\cap$ (Domain of j_{23});
- m_2' is unchanged from m_2 except for its public memory is projected from m_1'.

Construction of the Interpolating State

Protection after Composition $\geqslant$ Protection before Composition

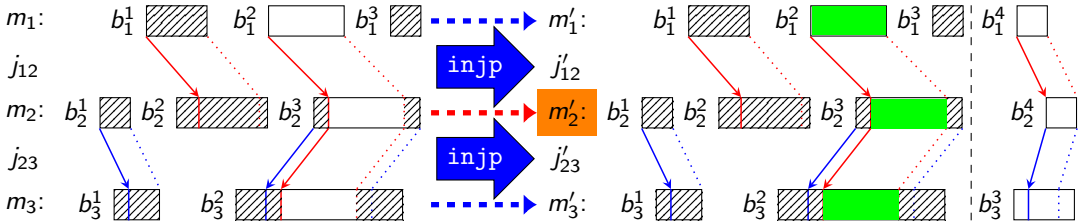

- Public memory of m_2' = (Image of j_{12}) $\cap$ (Domain of j_{23});
- m_2' is unchanged from m_2 except for its public memory is projected from m_1'.

CompCert(O) with Direct Refinement

Passes	Rely $\twoheadrightarrow$ Guarantee
Self-Sim	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
SimplLocals	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{inj}}$
Cminorgen	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{inj}}$
Selection	$\mathtt{wt} \cdot \mathtt{c_{ext}} \twoheadrightarrow \mathtt{wt} \cdot \mathtt{c_{ext}}$
RTLgen	$\mathtt{c_{ext}} \twoheadrightarrow \mathtt{c_{ext}}$
Self-Sim	$\mathtt{c_{inj}} \twoheadrightarrow \mathtt{c_{inj}}$
Tailcall	$\mathtt{c_{ext}} \twoheadrightarrow \mathtt{c_{ext}}$
Inlining	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{inj}}$
Self-Sim	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{injp}}$
Constprop	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
CSE	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
Deadcode	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
Unusedglob	$\mathtt{c_{inj}} \twoheadrightarrow \mathtt{c_{inj}}$
Allocation	$\mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{CL} \twoheadrightarrow \mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{CL}$
Tunneling	$\mathtt{ltl_{ext}} \twoheadrightarrow \mathtt{ltl_{ext}}$
Stacking	$\mathtt{ltl_{injp}} \cdot \mathtt{LM} \twoheadrightarrow \mathtt{LM} \cdot \mathtt{mach_{inj}}$
Asmgen	$\mathtt{mach_{ext}} \cdot \mathtt{MA} \twoheadrightarrow \mathtt{mach_{ext}} \cdot \mathtt{MA}$
Self-Sim	$\mathtt{asm_{inj}} \cdot \mathtt{asm_{injp}} \twoheadrightarrow \mathtt{asm_{inj}} \cdot \mathtt{asm_{injp}}$

Significant Passes

$$\mathtt{c_{injp}} \cdot \mathtt{c_{injp}} \equiv \mathtt{c_{injp}}$$

$$\mathtt{c_{ext}} \cdot \mathtt{c_{injp}} \equiv \mathtt{c_{injp}}$$

$$\cdots$$

CompCert(O) with Direct Refinement

Passes	Rely ↠ Guarantee
Self-Sim	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
SimplLocals	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{inj}}$
Cminorgen	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{inj}}$
Selection	$\mathtt{wt} \cdot \mathtt{c_{ext}} \twoheadrightarrow \mathtt{wt} \cdot \mathtt{c_{ext}}$
RTLgen	$\mathtt{c_{ext}} \twoheadrightarrow \mathtt{c_{ext}}$
Self-Sim	$\mathtt{c_{inj}} \twoheadrightarrow \mathtt{c_{inj}}$
Tailcall	$\mathtt{c_{ext}} \twoheadrightarrow \mathtt{c_{ext}}$
Inlining	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{inj}}$
Self-Sim	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{injp}}$
Constprop	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
CSE	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
Deadcode	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
Unusedglob	$\mathtt{c_{inj}} \twoheadrightarrow \mathtt{c_{inj}}$
Allocation	$\mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{CL} \twoheadrightarrow \mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{CL}$
Tunneling	$\mathtt{ltl_{ext}} \twoheadrightarrow \mathtt{ltl_{ext}}$
Stacking	$\mathtt{ltl_{injp}} \cdot \mathtt{LM} \twoheadrightarrow \mathtt{LM} \cdot \mathtt{mach_{inj}}$
Asmgen	$\mathtt{mach_{ext}} \cdot \mathtt{MA} \twoheadrightarrow \mathtt{mach_{ext}} \cdot \mathtt{MA}$
Self-Sim	$\mathtt{asm_{inj}} \cdot \mathtt{asm_{injp}} \twoheadrightarrow \mathtt{asm_{inj}} \cdot \mathtt{asm_{injp}}$

Significant Passes

$\preccurlyeq_{\mathsf{ac}}$:

$$\mathtt{ro} \cdot \mathtt{c_{injp}} \cdot \mathtt{c_{injp}} \cdot \mathtt{c_{injp}} \cdot \mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{c_{ext}} \cdot \mathtt{c_{inj}} \cdot \mathtt{c_{ext}} \cdot \mathtt{c_{injp}} \cdot \mathtt{c_{injp}} \cdot \mathtt{ro} \cdot \mathtt{c_{injp}} \cdot \mathtt{ro} \cdot \mathtt{c_{injp}} \cdot \mathtt{ro} \cdot \mathtt{c_{injp}} \cdot \mathtt{c_{inj}} \cdot \mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{CL} \cdot \mathtt{ltl_{ext}} \cdot \mathtt{ltl_{injp}} \cdot \mathtt{LM} \cdot \mathtt{mach_{ext}} \cdot \mathtt{MA} \cdot \mathtt{asm_{inj}} \cdot \mathtt{asm_{injp}}$$

$$\twoheadrightarrow$$

$$\mathtt{ro} \cdot \mathtt{c_{injp}} \cdot \mathtt{c_{inj}} \cdot \mathtt{c_{inj}} \cdot \mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{c_{ext}} \cdot \mathtt{c_{inj}} \cdot \mathtt{c_{ext}} \cdot \mathtt{c_{inj}} \cdot \mathtt{c_{injp}} \cdot \mathtt{ro} \cdot \mathtt{c_{injp}} \cdot \mathtt{ro} \cdot \mathtt{c_{injp}} \cdot \mathtt{ro} \cdot \mathtt{c_{injp}} \cdot \mathtt{c_{inj}} \cdot \mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{CL} \cdot \mathtt{ltl_{ext}} \cdot \mathtt{LM} \cdot \mathtt{mach_{inj}} \cdot \mathtt{mach_{ext}} \cdot \mathtt{MA} \cdot \mathtt{asm_{inj}} \cdot \mathtt{asm_{injp}}$$

CompCert(O) with Direct Refinement

Passes	Rely $\twoheadrightarrow$ Guarantee
Self-Sim	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
SimplLocals	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{inj}}$
Cminorgen	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{inj}}$
Selection	$\mathtt{wt} \cdot \mathtt{c_{ext}} \twoheadrightarrow \mathtt{wt} \cdot \mathtt{c_{ext}}$
RTLgen	$\mathtt{c_{ext}} \twoheadrightarrow \mathtt{c_{ext}}$
Self-Sim	$\mathtt{c_{inj}} \twoheadrightarrow \mathtt{c_{inj}}$
Tailcall	$\mathtt{c_{ext}} \twoheadrightarrow \mathtt{c_{ext}}$
Inlining	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{inj}}$
Self-Sim	$\mathtt{c_{injp}} \twoheadrightarrow \mathtt{c_{injp}}$
Constprop	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
CSE	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
Deadcode	$\mathtt{ro} \cdot \mathtt{c_{injp}} \twoheadrightarrow \mathtt{ro} \cdot \mathtt{c_{injp}}$
Unusedglob	$\mathtt{c_{inj}} \twoheadrightarrow \mathtt{c_{inj}}$
Allocation	$\mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{CL} \twoheadrightarrow \mathtt{wt} \cdot \mathtt{c_{ext}} \cdot \mathtt{CL}$
Tunneling	$\mathtt{ltl_{ext}} \twoheadrightarrow \mathtt{ltl_{ext}}$
Stacking	$\mathtt{ltl_{injp}} \cdot \mathtt{LM} \twoheadrightarrow \mathtt{LM} \cdot \mathtt{mach_{inj}}$
Asmgen	$\mathtt{mach_{ext}} \cdot \mathtt{MA} \twoheadrightarrow \mathtt{mach_{ext}} \cdot \mathtt{MA}$
Self-Sim	$\mathtt{asm_{inj}} \cdot \mathtt{asm_{injp}} \twoheadrightarrow \mathtt{asm_{inj}} \cdot \mathtt{asm_{injp}}$

Significant Passes

$\preccurlyeq_{\mathtt{ac}}$:

$$\mathtt{ro} \cdot \mathtt{wt} \cdot \mathtt{CA_{injp}} \cdot \mathtt{asm_{injp}}$$

$$\twoheadrightarrow$$

$$\mathtt{ro} \cdot \mathtt{wt} \cdot \mathtt{CA_{injp}} \cdot \mathtt{asm_{injp}}$$

Conclusion

Direct refinements of realistic verified compilers are feasible:

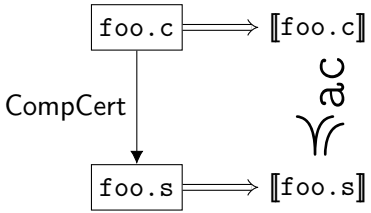

Discovery: Transitivity of Kripke Relation with Memory Protection

Ongoing/Future work:

- Reduce to the original CompCert
- Connect with Program Verification
- Verified Compilation of Safe/Unsafe Rust

https://doi.org/10.5281/zenodo.10036618