

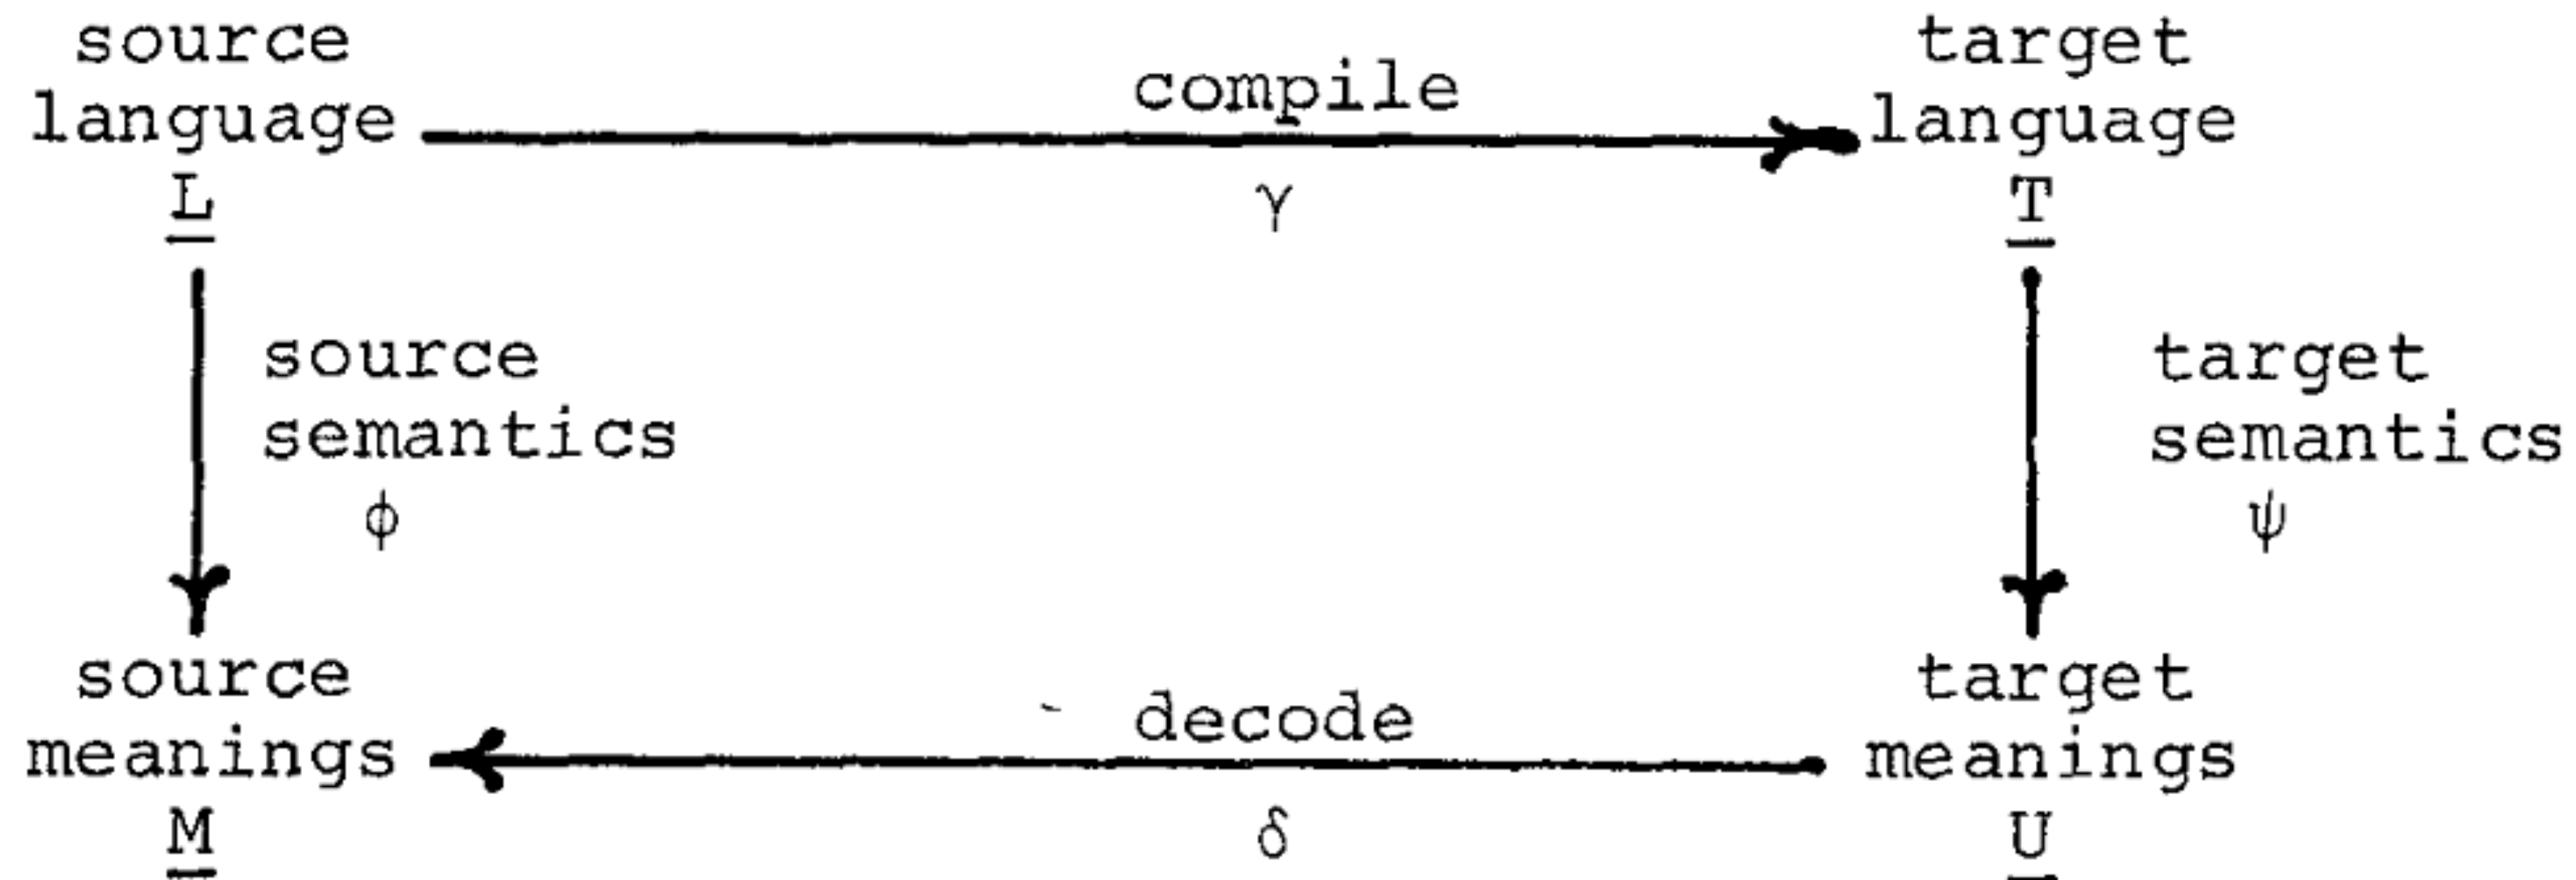
DimSum: A Decentralized Approach to Multi-language Semantics and Verification

Michael Sammler, Simon Spies,
Youngju Song, Emanuele D'Oswaldo,
Robbert Krebbers, Deepak Garg, Derek Dreyer

POPL 2023 Distinguished Paper

Presented by Priya Srikumar, at Cornell University PLDG, Spring 2023

Program verification needs compiler correctness!



F. Lockwood Morris. 1973. Advice on structuring compilers and proving them correct.

In Proceedings of the 1st annual ACM SIGACT-SIGPLAN symposium on Principles of programming languages (POPL '73).

Association for Computing Machinery, New York, NY, USA, 144–152. <https://doi.org/10.1145/512927.512941>

...but compiling whole programs isn't realistic

```
int src = 1;
int dst;

asm ("mov %1, %0\n\t"
     "add $1, %0"
     : "=r" (dst)
     : "r" (src));

printf("%d\n", dst);
```

Verifying multi-language programs is hard (!)

- Languages have different interfaces/data representations
- Prior work limits the structure of multi-language programs
 - Fixed source language
 - Fixed set of languages
 - Fixed memory model
 - Fixed interoperation mechanisms

DimSum employs a *decentralized* approach

- Framework to specify/verify multi-language programs
- Notion of compiler correctness via module refinement
- Language-agnostic combinators to link/translate modules
- Users define labeled transition systems over language events
- Users define embeddings between each pair of languages

Case study: **Asm** and **Rec**

Trace events

- **Jump** and **Syscall**
- **Call** and (function) **Return**
- ! for an outgoing event
 - **Syscall!** (...) is an outgoing syscall
- ? for an incoming event
 - **Return?** (...) is an incoming function return

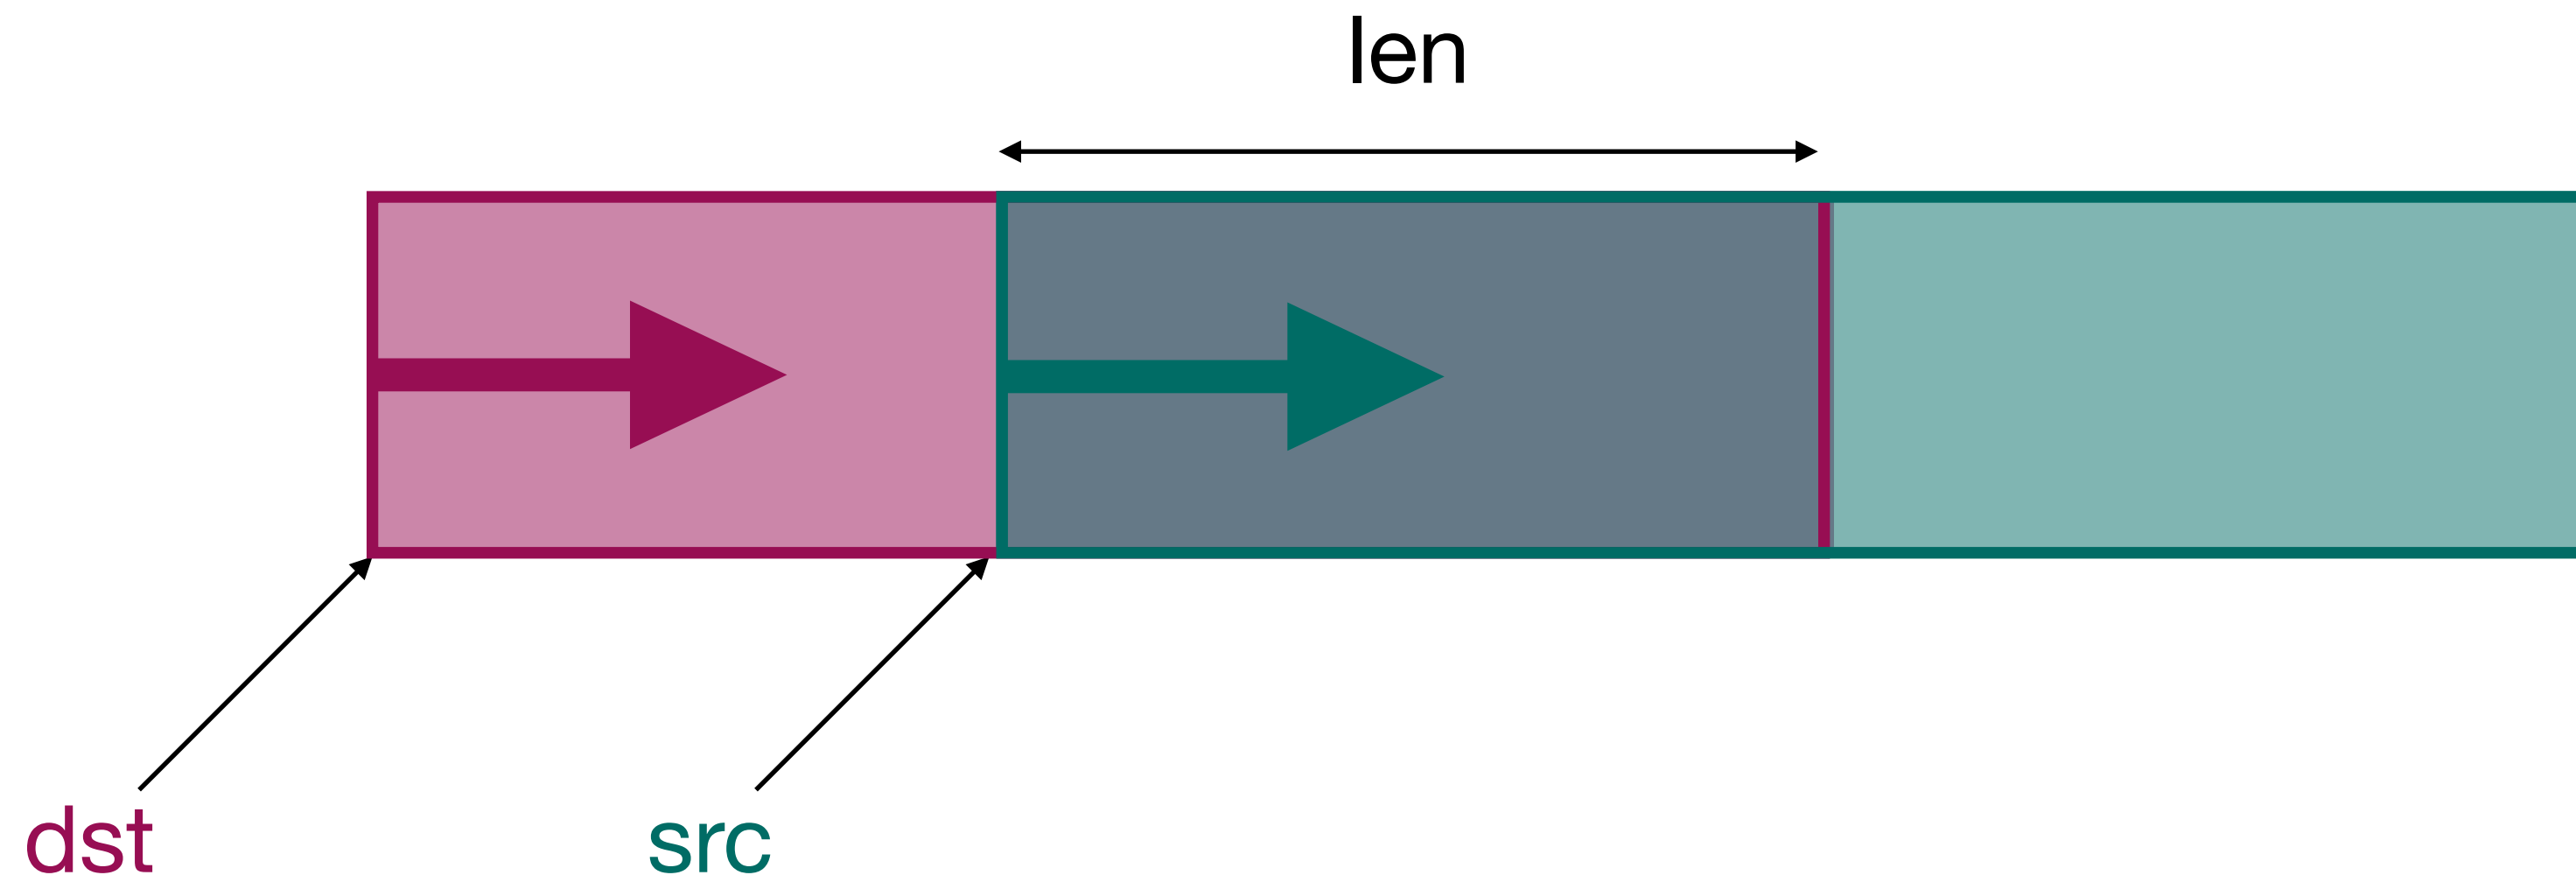
main

```
fn main() {  
    local x[3];  
    x[0] ← 1;  
    x[1] ← 2;  
    // x ↦ [1, 2, 0]  
    memmove(x + 1, x + 0, 2);  
    // x ↦ [1, 1, 2]  
    print(x[1]); print(x[2])  
    // print 1, then 2  
}
```

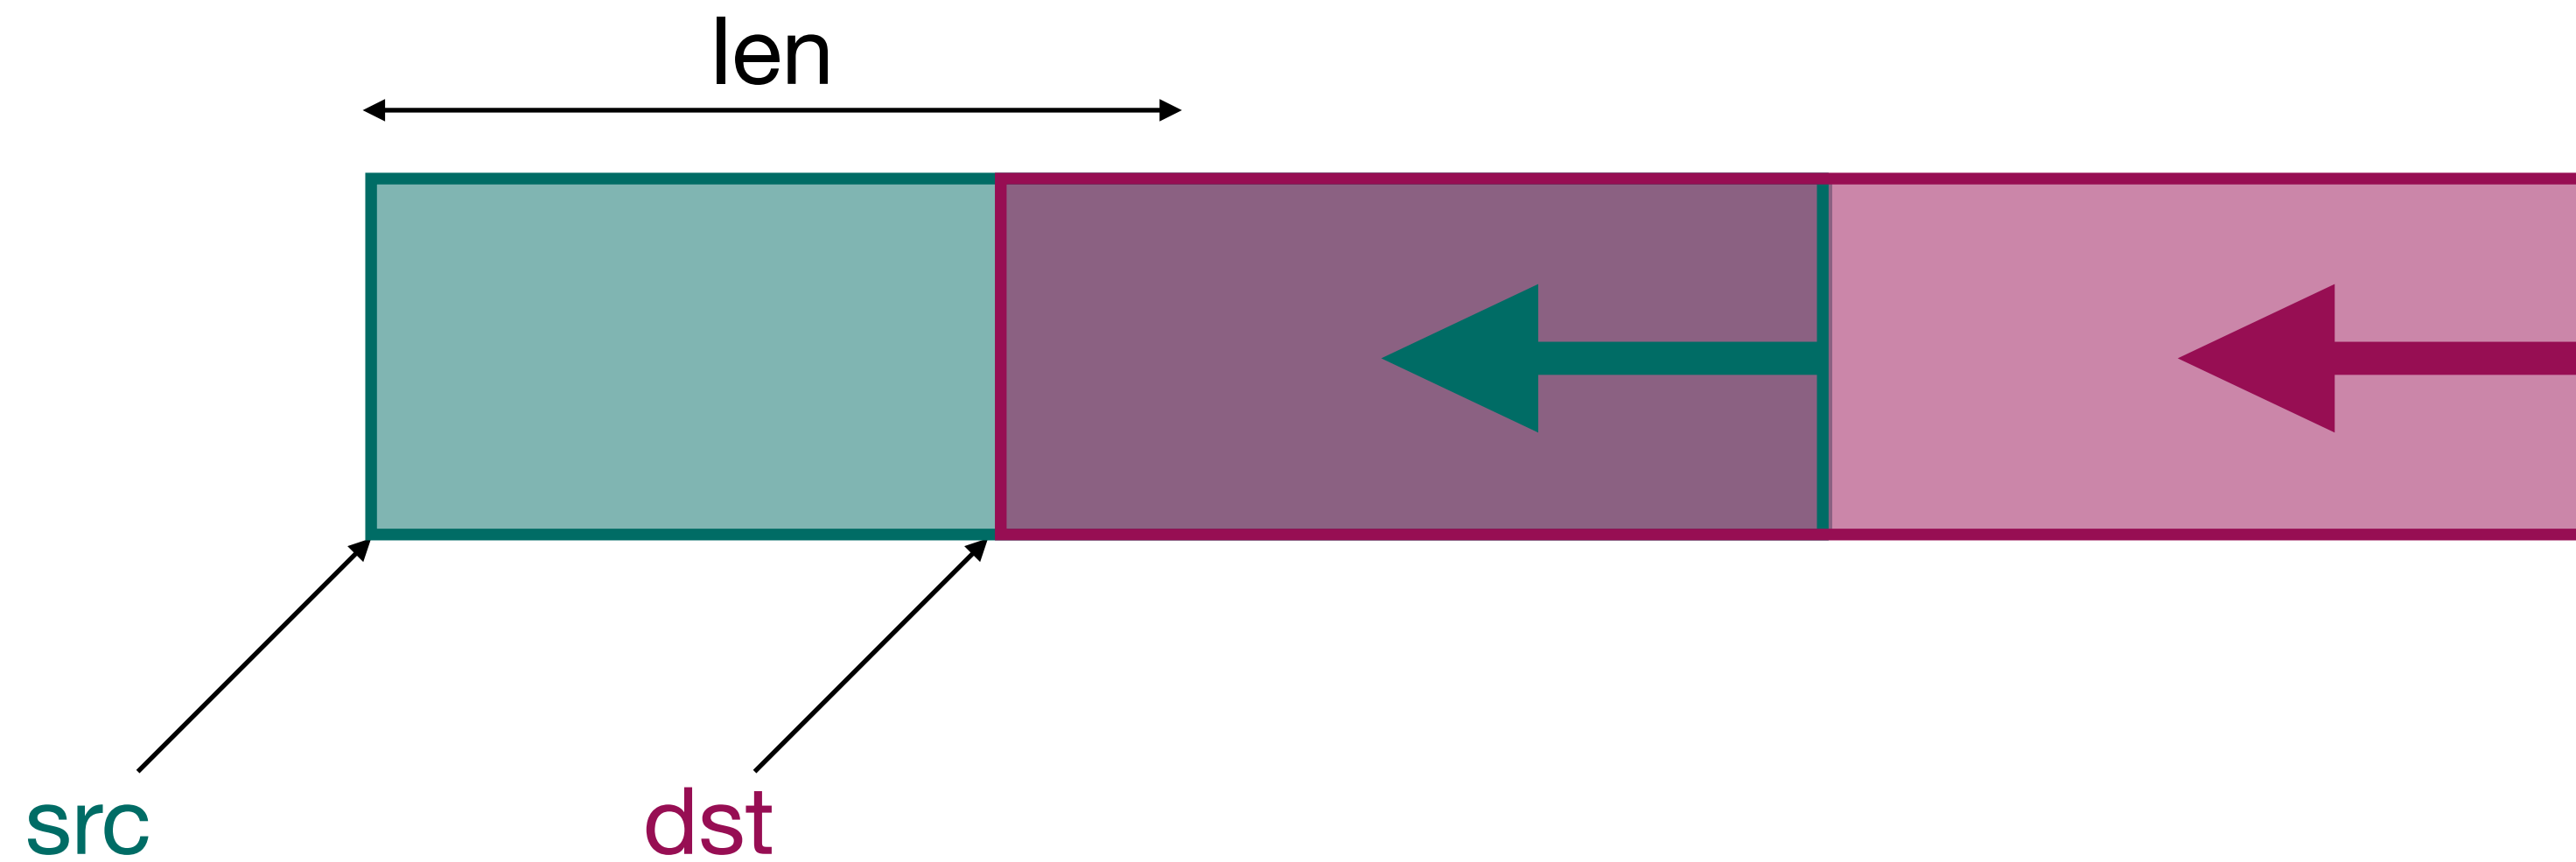
main

```
fn main() {  
    local x[3];  
    x[0] ← 1;  
    x[1] ← 2;  
    // x ↦ [1, 2, 0]  
    memmove(x + 1, x + 0, 2);  
    // x ↦ [1, 1, 2]  
    print(x[1]); print(x[2])  
    // print 1, then 2  
}
```


Correctness of `memmove(dst, src, len)`



Correctness of `memmove(dst, src, len)`



Correctness of `memmove(dst, src, len)`



memmove

```
fn memmove(d, s, n) {  
  if locle(d, s)  
  then  
    memcpy(d, s, n, 1)  
  else  
    memcpy(d + n - 1, s + n - 1, n, -1)  
}
```

memmove

```
fn memmove(d, s, n) {  
  if locle(d, s)  
  then  
    memcpy(d, s, n, 1)  
  else  
    memcpy(d + n - 1, s + n - 1, n, -1)  
}
```

locle

```
locle : sle x0, x0, x1;  
      ret
```

locle

```
locle : sle x0, x0, x1;  
      ret
```


Can't be implemented in **Rec!**

Asm and Rec

Memory Model

Addresses

Values



Block ID + offset



main

```
fn main() {  
    local x[3];  
    x[0] ← 1;  
    x[1] ← 2;  
    // x ↦ [1, 2, 0]  
    memmove(x + 1, x + 0, 2);  
    // x ↦ [1, 1, 2]  
    print(x[1]); print(x[2])  
    // print 1, then 2  
}
```

main

```
fn main() {  
    local x[3];  
    x[0] ← 1;  
    x[1] ← 2;  
    // x ↦ [1, 2, 0]  
    memmove(x + 1, x + 0, 2);  
    // x ↦ [1, 1, 2]  
    print(x[1]); print(x[2])  
    // print 1, then 2  
}
```

print

```
print : mov x8, PRINT;  
      syscall; ret
```

`print`

```
print : mov x8, PRINT;  
      syscall; ret
```

Can't be implemented in `Rec`!

main

```
fn main() {  
    local x[3];  
    x[0] ← 1;  
    x[1] ← 2;  
    // x ↦ [1, 2, 0]  
    memmove(x + 1, x + 0, 2);  
    // x ↦ [1, 1, 2]  
    print(x[1]); print(x[2])  
    // print 1, then 2  
}
```

Our compiled program!

`onetwo` \triangleq \downarrow `main` U_a \downarrow `memmove` U_a `locle` U_a `print`

\downarrow `prog` := compilation from **Rec** to **Asm**

`prog1` U_a `prog2` := *syntactic linking operator* for **Asm**

Our proof goal

$$\text{onetwo} \text{ a} \leq \text{S} \text{ onetwo}_{\text{spec}}$$

1. Each step of the **Asm** program **onetwo** corresponds to zero or more steps of the **Asm**-level specification **onetwo**_{spec}.
2. **onetwo** and **onetwo**_{spec} yield the same external behavior.

From *Libraries* to *Modules*

$[\mathbf{AsmLib}]_a$ is a *semantic module* for the *syntactic library* \mathbf{AsmLib} .

A module $M \in \mathit{Module}(E)$ is a *labeled transition system* emitting events from a language-specific set of events E .

An excerpt of module semantics

ASM-INCOMING

$$\mathbf{r(pc) = a} \quad \mathbf{a \in |A|}$$

$$\mathbf{(Wait, A)} \xrightarrow{\mathbf{Jump?(r,m)}}_{\mathbf{a}} \{ \mathbf{(Run(r, m), A)} \}$$

Defining our proof goal

$$\begin{array}{ccc} \text{onetwo} & \stackrel{a}{\sim}_s & \text{onetwo}_{\text{spec}} \\ & \triangleq & \\ \llbracket \text{onetwo} \rrbracket_a & \sim & \llbracket \text{onetwo}_{\text{spec}} \rrbracket_s \end{array}$$

$M_1 \preceq M_2$ is the *simulation* between modules.

Decomposing our program

$[\text{onetwo}]_a \preceq [\downarrow \text{main} \cup_a \downarrow \text{memmove} \cup_a \text{locle} \cup_a \text{print}]_a$

By definition of `onetwo`.

Decomposing into linked **Asm** modules

$$\llbracket \downarrow \text{main} \cup_a \downarrow \text{memmove} \cup_a \text{locle} \cup_a \text{print} \rrbracket_a \preceq \\ \llbracket \downarrow \text{main} \rrbracket_a \oplus_a \llbracket \downarrow \text{memmove} \rrbracket_a \oplus_a \llbracket \text{locle} \rrbracket_a \oplus_a \llbracket \text{print} \rrbracket_a$$

$M_1 \oplus_a M_2$ is a *semantic linking operator* for **Asm**.

It synchronizes the **Jump** events of M_1 and M_2 .

Semantically linking **Asm** modules

$$M_1 \oplus_a M_2$$

ASM-LINK-JUMP

$$\frac{(d' = L \wedge \mathbf{r(pc)} \in \mathbf{d}_1) \vee (d' = R \wedge \mathbf{r(pc)} \in \mathbf{d}_2) \vee (d' = E \wedge \mathbf{r(pc)} \notin \mathbf{d}_1 \cup \mathbf{d}_2) \quad d \neq d'}{(d, \text{None}, \mathbf{Jump}(\mathbf{r}, \mathbf{m})) \rightsquigarrow_{\mathbf{d}_1, \mathbf{d}_2} (d', \text{None}, \mathbf{Jump}(\mathbf{r}, \mathbf{m}))}$$

Decomposing into linked **Asm** modules

$$\llbracket \downarrow \mathbf{main} \cup_a \downarrow \mathbf{memmove} \cup_a \mathbf{locle} \cup_a \mathbf{print} \rrbracket_a \preceq \\ \llbracket \downarrow \mathbf{main} \rrbracket_a \oplus_a \llbracket \downarrow \mathbf{memmove} \rrbracket_a \oplus_a \llbracket \mathbf{locle} \rrbracket_a \oplus_a \llbracket \mathbf{print} \rrbracket_a$$

$M_1 \oplus_a M_2$ is a *semantic linking operator* for **Asm**.

It synchronizes the **Jump** events of M_1 and M_2 .

Translating **Asm** modules back into **Rec** modules

$$\llbracket \downarrow \text{main} \rrbracket_a \oplus_a \llbracket \downarrow \text{memmove} \rrbracket_a \oplus_a \llbracket \text{locle} \rrbracket_a \oplus_a \llbracket \text{print} \rrbracket_a \leq \\ \llbracket \llbracket \text{main} \rrbracket_r \rrbracket_{r \Rightarrow a} \oplus_a \llbracket \llbracket \text{memmove} \rrbracket_r \rrbracket_{r \Rightarrow a} \oplus_a \llbracket \text{locle} \rrbracket_a \oplus_a \llbracket \text{print} \rrbracket_a$$

$\llbracket \cdot \rrbracket_{r \Rightarrow a}$ is a *semantic wrapper*.

It embeds **Rec** modules into **Asm** and translates between **Rec** and **Asm** events.


Translating **Asm** modules back into **Rec** modules

$$\begin{aligned} & \llbracket \downarrow \text{main} \rrbracket_a \oplus_a \llbracket \downarrow \text{memmove} \rrbracket_a \oplus_a \llbracket \text{locle} \rrbracket_a \oplus_a \llbracket \text{print} \rrbracket_a \leq \\ & \llbracket \llbracket \text{main} \rrbracket_r \rrbracket_{r \Rightarrow a} \oplus_a \boxed{\llbracket \llbracket \text{memmove} \rrbracket_r \rrbracket_{r \Rightarrow a}} \oplus_a \llbracket \text{locle} \rrbracket_a \oplus_a \llbracket \text{print} \rrbracket_a \end{aligned}$$

$\llbracket \cdot \rrbracket_{r \Rightarrow a}$ is a *semantic wrapper*.

It embeds **Rec** modules into **Asm** and translates between **Rec** and **Asm** events.

The semantic wrapper $\llbracket \cdot \rrbracket_{r \Rightarrow a}$ in use



Translating **Asm** modules back into **Rec** modules

$$\begin{aligned} & \llbracket \downarrow \text{main} \rrbracket_a \oplus_a \llbracket \downarrow \text{memmove} \rrbracket_a \oplus_a \llbracket \text{locle} \rrbracket_a \oplus_a \llbracket \text{print} \rrbracket_a \leq \\ & \llbracket \llbracket \text{main} \rrbracket_r \rrbracket_{r \Rightarrow a} \oplus_a \llbracket \llbracket \text{memmove} \rrbracket_r \rrbracket_{r \Rightarrow a} \oplus_a \llbracket \text{locle} \rrbracket_a \oplus_a \llbracket \text{print} \rrbracket_a \end{aligned}$$

$\llbracket \cdot \rrbracket_{r \Rightarrow a}$ is a *semantic wrapper*.

It embeds **Rec** modules into **Asm** and translates between **Rec** and **Asm** events.

Abstracting **Asm** implementations

$$\begin{aligned} & \llbracket \mathbf{main} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{memmove} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{locle} \rrbracket_a \oplus_a \llbracket \mathbf{print} \rrbracket_a \preceq \\ & \llbracket \mathbf{main} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{memmove} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{locle}_{\text{spec}} \rrbracket_s \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{print} \rrbracket_a \end{aligned}$$

locle can have a **Rec**-level specification because it has the same interaction behavior as a **Rec** program (i.e., no syscalls).

Abstracting **Asm** implementations

$$\begin{aligned} & \llbracket \mathbf{main} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{memmove} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{loacle}_{\text{spec}} \rrbracket_s \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{print} \rrbracket_a \leq \\ & \llbracket \mathbf{main} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{memmove} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{loacle}_{\text{spec}} \rrbracket_s \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{print}_{\text{spec}} \rrbracket_s \end{aligned}$$

print makes a syscall, so it must have an **Asm**-level specification.

Moving away from **Asm**

$$\begin{aligned} & \llbracket \mathbf{main} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{memmove} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{loclc} \rrbracket_{\text{spec}} \Big|_s \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{print} \rrbracket_{\text{spec}} \Big|_s \leq \\ & \llbracket \mathbf{main} \rrbracket_r \oplus_r \llbracket \mathbf{memmove} \rrbracket_r \oplus_r \llbracket \mathbf{loclc} \rrbracket_{\text{spec}} \Big|_s \Big|_{r \Rightarrow a} \oplus_a \llbracket \mathbf{print} \rrbracket_{\text{spec}} \Big|_s \end{aligned}$$

Semantic linking distributes over semantic wrappers.

Semantic linking over semantic wrappers

REC-TO-ASM-LINK

$$\llbracket M_1 \rrbracket_{r \Rightarrow a} \oplus_a \llbracket M_2 \rrbracket_{r \Rightarrow a} \leq \llbracket M_1 \oplus_r M_2 \rrbracket_{r \Rightarrow a}$$

Moving away from *Asm*

$$\begin{aligned} & \llbracket \text{main} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \text{memmove} \rrbracket_r \Big|_{r \Rightarrow a} \oplus_a \llbracket \text{locle}_{\text{spec}} \rrbracket_s \Big|_{r \Rightarrow a} \oplus_a \llbracket \text{print}_{\text{spec}} \rrbracket_s \leq \\ & \llbracket \text{main} \rrbracket_r \oplus_r \llbracket \text{memmove} \rrbracket_r \oplus_r \llbracket \text{locle}_{\text{spec}} \rrbracket_s \Big|_{r \Rightarrow a} \oplus_a \llbracket \text{print}_{\text{spec}} \rrbracket_s \end{aligned}$$

Semantic linking distributes over semantic wrappers.

Shifting from semantic to syntactic linking

$$\begin{aligned} & \left[\llbracket \mathbf{main} \rrbracket_r \oplus_r \llbracket \mathbf{memmove} \rrbracket_r \oplus_r \llbracket \mathbf{locle}_{\text{spec}} \rrbracket_s \right]_{r \Rightarrow a} \oplus_a \llbracket \mathbf{print}_{\text{spec}} \rrbracket_s \preceq \\ & \left[\llbracket \mathbf{main} \ U_r \ \mathbf{memmove} \rrbracket_r \oplus_r \llbracket \mathbf{locle}_{\text{spec}} \rrbracket_s \right]_{r \Rightarrow a} \oplus_a \llbracket \mathbf{print}_{\text{spec}} \rrbracket_s \end{aligned}$$

Syntactic linking and semantic linking are equivalent in **Rec**,
so we can link **main** and **memmove** semantically with U_r .

Rec-level reasoning

$$\left[\llbracket \text{main} \cup_r \text{memmove} \rrbracket_r \oplus_r \llbracket \text{locale} \rrbracket_{\text{spec}} \right]_{r \Rightarrow a} \oplus_a \llbracket \text{print} \rrbracket_{\text{spec}} \preceq \left[\llbracket \text{main} \rrbracket_{\text{spec}} \right]_{r \Rightarrow a} \oplus_a \llbracket \text{print} \rrbracket_{\text{spec}}$$

$\llbracket \text{main} \cup_r \text{memmove} \rrbracket_r \oplus_r \llbracket \text{locale} \rrbracket_{\text{spec}} \preceq \llbracket \text{main} \rrbracket_{\text{spec}}$,
so we can use $\llbracket \text{main} \rrbracket_{\text{spec}}$ instead.

locl_espec

Pointers within the same block



Pointers in different blocks



Rec-level reasoning

$$\left[\llbracket \text{main} \cup_r \text{memmove} \rrbracket_r \oplus_r \llbracket \text{locale} \rrbracket_{\text{spec}} \right]_{r \Rightarrow a} \oplus_a \llbracket \text{print} \rrbracket_{\text{spec}} \preceq \left[\llbracket \text{main} \rrbracket_{\text{spec}} \right]_{r \Rightarrow a} \oplus_a \llbracket \text{print} \rrbracket_{\text{spec}}$$

$\llbracket \text{main} \cup_r \text{memmove} \rrbracket_r \oplus_r \llbracket \text{locale} \rrbracket_{\text{spec}} \preceq \llbracket \text{main} \rrbracket_{\text{spec}}$,
so we can use $\llbracket \text{main} \rrbracket_{\text{spec}}$ instead.

Asm-level reasoning

$$\llbracket \text{main}_{\text{spec}} \rrbracket_s \Big|_{r \Leftarrow a} \bigoplus_a \llbracket \text{print}_{\text{spec}} \rrbracket_s \preceq \llbracket \text{onetwo}_{\text{spec}} \rrbracket_s$$

We only need to make sure that the calls in **main** and the jumps in **print** line up to prove the simulation relation.

Finished proof!

$$\llbracket \text{onetwo} \rrbracket_a \preceq \llbracket \text{onetwo}_{\text{spec}} \rrbracket_s$$

The program's module refines its specification's module.

Compiler correctness in DimSum

COMPILER-CORRECT

$\downarrow R$ *defined*

$$\llbracket \downarrow R \rrbracket_a \leq \llbracket \llbracket R \rrbracket_r \rrbracket_{r \rightleftharpoons a}$$

Syntactic translation via the compiler refines
semantic translation via the wrapper.