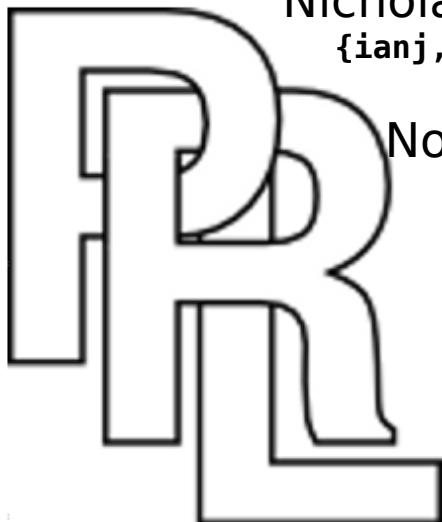


Optimizing

Abstract Abstract Machines



J. Ian Johnson,

Nicholas Labich, David Van Horn

{ianj, labichn, dvanhorn}@ccs.neu.edu

Northeastern University

Boston, MA, USA

Matt Might

matt@might.net

University of Utah

Salt Lake City, UT, USA

Abstract
abstract machines?

abc

Abstracting Abstract Machines

David Van Horn *
Northeastern University
dvanhorn@ccs.neu.edu

Matthew Might
University of Utah
might@cs.utah.edu

Abstract

We describe a derivational approach to abstract interpretation that yields novel and transparently sound static analyses when applied to well-established abstract machines. To demonstrate the technique and support our claim, we transform the CEK machine of Felleisen and Friedman, a lazy variant of Krivine's machine, and the stack-inspecting CM machine of Clements and Felleisen into abstract interpretations of themselves. The resulting analyses bound temporal ordering of program events; predict return-flow and stack-inspection behavior; and approximate the flow and evaluation of by-need parameters. For all of these machines, we find that a series of well-known concrete machine refactorings, plus a technique we call store-allocated continuations, leads to machines that abstract into static analyses simply by bounding their stores. We demonstrate that the techniques scale uniformly to allow static analysis of realistic language features, including tail calls, conditionals, side effects, exceptions, first-class continuations, and even garbage collection.

Categories and Subject Descriptors F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages—Program analysis, Operational semantics; F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Lambda calculus and related systems

General Terms Languages, Theory

Keywords abstract machines, abstract interpretation

1. Introduction

Abstract machines such as the CEK machine and Krivine's machine are first-order state transition systems that represent the core of a real language implementation. Store-allocated program analyses, on the other hand, is concerned with safety and maintaining intensional properties of such a machine as it runs a program. It seems natural then to want to systematically derive analyses from machines to approximate the core of realistic run-time systems.

Our goal is to develop a technique that enables direct abstract interpretations of abstract machines by methods for transforming a given machine description into another that computes its finite approximation.

* Supported by the National Science Foundation under grant 0937060 to the Computing Research Association for the CIFellow Project.

Permission to make digital or hard copies of all or part of this work for personal classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICFP '10, September 27–29, 2010, Baltimore, Maryland, USA.
Copyright © 2010 ACM 978-1-60558-794-3/10/09 . \$10.00.

We demonstrate that the technique of refactoring a machine with **store-allocated continuations** allows a direct structural abstraction¹ by bounding the machine's store. Thus, we are able to convert semantic techniques used to model language features into static analysis techniques for reasoning about the behavior of those very same features. By abstracting well-known machines, our technique delivers static analyzers that can reason about by-need evaluation, higher-order functions, tail calls, side effects, stack structure, exceptions and first-class continuations.

The basic idea behind store-allocated continuations is not new. SML/NJ has allocated continuations in the heap for well over a decade [28]. At first glance, modeling the program stack in an abstract machine with store-allocated continuations would not seem to provide any real benefit. Indeed, for the purpose of defining the meaning of a program, there is no benefit, because the meaning of the program does not depend on the stack-implementation strategy. Yet, a closer inspection finds that store-allocating continuations eliminate recursion from the definition of the state-space of the machine. With no recursive structure in the state-space, an abstract machine becomes eligible for conversion into an abstract interpreter through a simple structural abstraction.

To demonstrate the applicability of the approach, we derive abstract interpreters of:

- a call-by-value λ -calculus with state and control based on the CESK machine of Felleisen and Friedman [13],
- a call-by-need λ -calculus based on a tail-recursive, lazy variant of Krivine's machine derived by Ager, Danvy and Midgaard [1], and
- a call-by-value λ -calculus with stack inspection based on the CM machine of Clements and Felleisen [3];

and use abstract garbage collection to improve precision [25].

Overview

In Section 2, we begin with the CEK machine and attempt a structural abstract interpretation, but find ourselves blocked by two recursive structures in the machine: environments and continuations. We make three refactorings to:

1. store-allocate bindings,
2. store-allocate continuations, and
3. time-stamp machine states;

resulting in the CESK, CESK*, and time-stamped CESK* machines, respectively. The time-stamps encode the history (context) of the machine's execution and facilitate context-sensitive abstractions. We then demonstrate that the time-stamped machine abstracts directly into a parameterized, sound and computable static analysis.

¹ A structural abstraction distributes component-, point-, and member-wise.

s?

abc

x 1000

s?

Abstracting Abstract Machines

David Van Horn *
Northeastern University
dvanhorn@ccs.neu.edu

Matthew Might
University of Utah
might@cs.utah.edu

Abstract

We describe a derivational approach to abstract interpretation that yields novel and transparently sound static analyses when applied to well-established abstract machines. To demonstrate the technique and support our claim, we transform the CEK machine of Felleisen and Friedman, a lazy variant of Krivine's machine, and the stack-inspecting CM machine of Clements and Felleisen into abstract interpretations of themselves. The resulting analyses bound temporal ordering of program events; predict return-flow and stack-inspection behavior; and approximate the flow and evaluation of by-need parameters. For all of these machines, we find that a series of well-known techniques we call store-allocated continuations allows a direct structural abstraction¹ by bounding the machine's store. Thus, we are able to convert semantic techniques used to model language features into static analysis techniques for reasoning about the behavior of those very same features. By abstracting well-known machines, our technique delivers static analyzers that can reason about by-need evaluation, higher-order functions, tail calls, side effects, stack structure, exceptions and first-class continuations.

The basic idea behind store-allocated continuations is not new. SML/NJ has allocated continuations in the heap for well over a decade [1]. However, the meaning of a continuation in a heap-based system is not clear. In particular, it is not clear how to implement a stack-inspecting continuation in such a system. We demonstrate that the meaning of a continuation in a heap-based system is clear if we abstract away from the state-space of the machine. Instead, we implement continuations as abstract objects in a separate state-space. By doing so, we are able to reason about them directly.

Categories and Subject Terms: Semantics, Operational Semantics, Formal Languages, Related Systems

General Terms: Languages, Abstract Interpretation

Keywords: abstract interpretation

1. Introduction

Abstract machines such as the CEK machine and Krivine's machine are first-order state transition systems that represent the core of a real language implementation. Semantics-based program analysis, on the other hand, is concerned with safety and maintaining intensional properties of such a machine as it runs a program. It seems natural then to want to systematically derive analyses from machines to approximate the core of realistic run-time systems.

Our goal is to develop a technique that enables direct abstract interpretations of abstract machines by methods for transforming a given machine description into another that computes its finite approximation.

* Supported by the National Science Foundation under grant 0937060 to the Computing Research Association for the CIFellow Project.

Permission to make digital or hard copies of all or part of this work for personal classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICFP '10, September 27–29, 2010, Baltimore, Maryland, USA
Copyright © 2010 ACM 978-1-60558-794-3/10/09 . \$10.00.

We demonstrate that the technique of refactoring a machine with store-allocated continuations allows a direct structural abstraction¹ by bounding the machine's store. Thus, we are able to convert semantic techniques used to model language features into static analysis techniques for reasoning about the behavior of those very same features. By abstracting well-known machines, our technique delivers static analyzers that can reason about by-need evaluation, higher-order functions, tail calls, side effects, stack structure, exceptions and first-class continuations.

The basic idea behind store-allocated continuations is not new. SML/NJ has allocated continuations in the heap for well over a decade [1]. However, the meaning of a continuation in a heap-based system is not clear. In particular, it is not clear how to implement a stack-inspecting continuation in such a system. We demonstrate that the meaning of a continuation in a heap-based system is clear if we abstract away from the state-space of the machine. Instead, we implement continuations as abstract objects in a separate state-space. By doing so, we are able to reason about them directly.

control based on the [13], recursive, lazy variant of Krivine's machine derived by Ager, Danvy and Midgaard [11], and

• a call-by-value λ -calculus with stack inspection based on the CM machine of Clements and Felleisen [3];

and use abstract garbage collection to improve precision [25].

Overview

In Section 2, we begin with the CEK machine and attempt a structural abstract interpretation, but find ourselves blocked by two recursive structures in the machine: environments and continuations. We make three refactorings to:

1. store-allocate bindings,
2. store-allocate continuations, and
3. time-stamp machine states;

resulting in the CESK, CESK*, and time-stamped CESK* machines, respectively. The time-stamps encode the history (context) of the machine's execution and facilitate context-sensitive abstractions. We then demonstrate that the time-stamped machine abstracts directly into a parameterized, sound and computable static analysis.

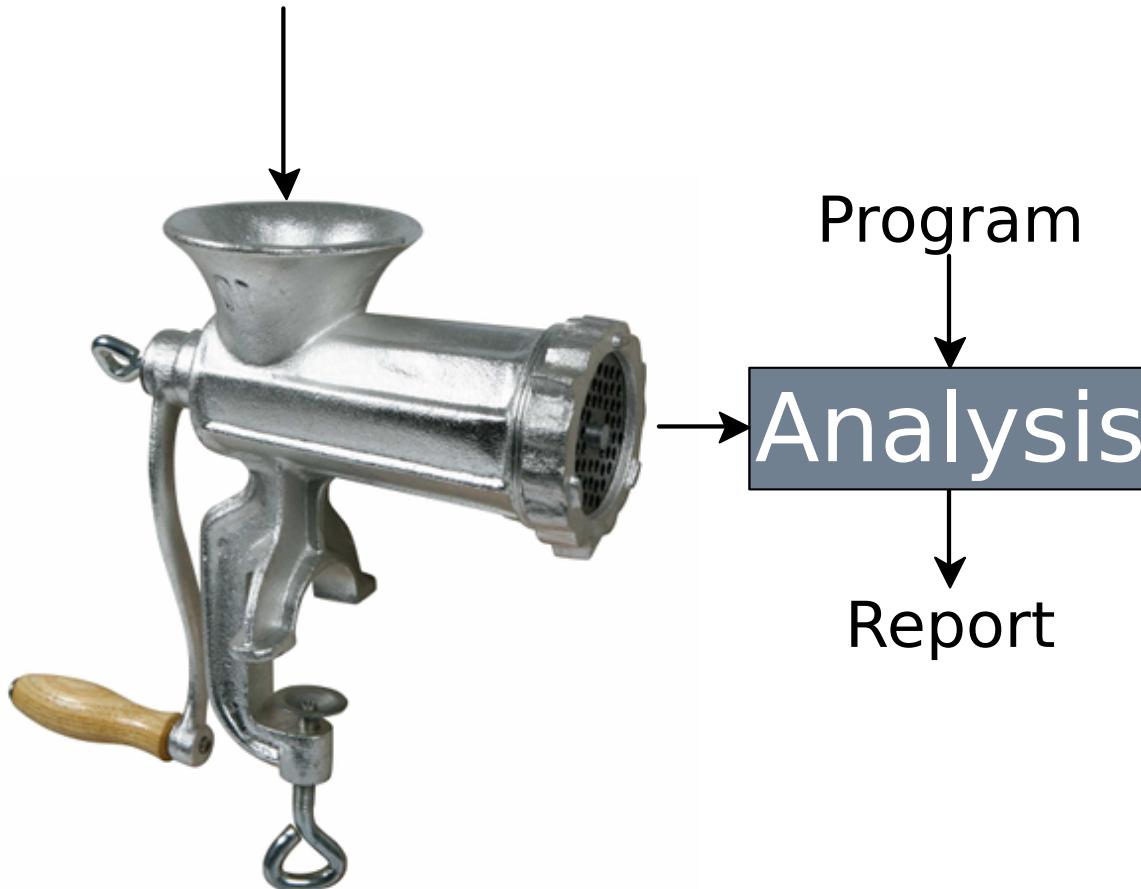
¹ A structural abstraction distributes component-, point-, and member-wise.

Semantics



→Analysis

Semantics



Program

Analysis

Report

Semantics



Program

Analysis

Maybe

Goal: Directly implementable math

Competitive performance,*

minimal effort,

simple proofs

same ballpark, nosebleed seats

Goal: Directly implementable math

Competitive performance,*

Slogan:

Engineering tricks
as semantics refactorings

same ballpark, nosebleed seats

An ideal static analysis has

Performance

An ideal static analysis has

Soundness

Maintainability

Precision

Performance

An ideal static analysis has

[Van Horn and Might ICFP 2010]
(AAM)

Soundness
Maintainability
Precision
Performance

An ideal static analysis has

[Van Horn and Might ICFP 2010]
(AAM)

Soundness
Maintainability
Precision
Performance

[Johnson, et al. ICFP 2013]
(OAAM)

OAAM outline

OAAM outline

- Store-allocate values
- Frontier-based semantics
- Lazy non-determinism
- Abstract compilation
- Locally log-based store-deltas
- Store-counting
- Mutable frontier and store

OAAM outline

Decrease state space

- Store-allocate values
- Frontier-based semantics
- Lazy non-determinism
- Abstract compilation
- Locally log-based store-deltas
- Store-counting
- Mutable frontier and store

OAAM outline

Decrease state space

- Store-allocate values
- Frontier-based semantics
- Lazy non-determinism
- Abstract compilation
- Locally log-based store-deltas
- Store-counting
- Mutable frontier and store

Explore faster

OAAM outline

Decrease state space

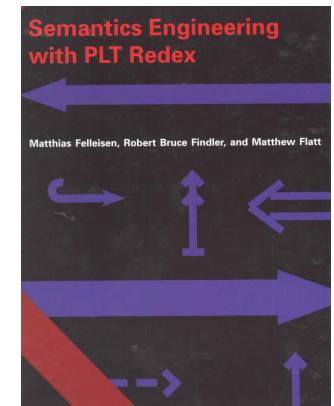
- Store-allocate values
 - Frontier-based semantics
1. Lazy non-determinism
 2. Abstract compilation
 3. Locally log-based store-deltas
 - Store-counting
 - Mutable frontier and store

Explore faster

First: recap of AAM

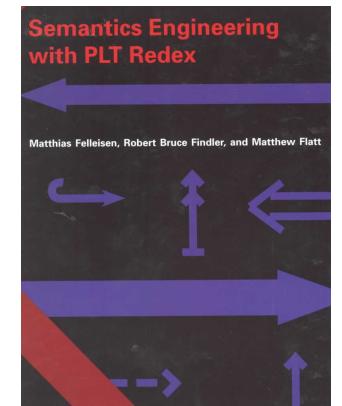
Start with CESK machine

Control
↓
 $\langle e, p, \sigma, K \rangle$



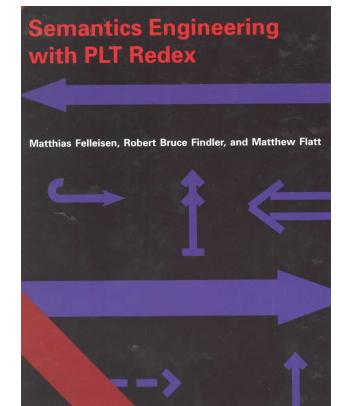
Start with CESK machine

$\langle e, p, \sigma, k \rangle$
↑
Environment



Start with CESK machine

Store
↓
 $\langle e, p, \sigma, k \rangle$

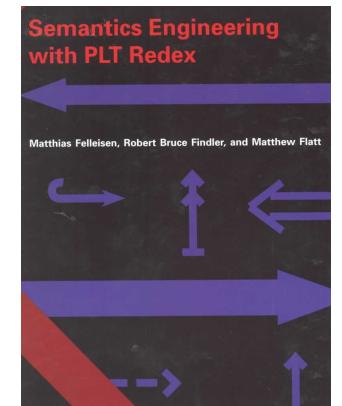


Start with CESK machine

$\langle e, p, \sigma, K \rangle$

↑

Kontinuation



$$\begin{array}{lcl}
 \rho \in \text{Env} & = & \text{Var} \rightarrow \text{Addr} \\
 \kappa \in \text{Kont} & ::= & [] \mid \varphi : \kappa \\
 \sigma \in \text{Store} & = & \text{Addr} \rightarrow (\text{Value} \times \text{Env}) \\
 \langle x, \rho, \sigma, \kappa \rangle & \mapsto & \langle v, \rho', \sigma, \kappa \rangle
 \end{array}$$

if $(v, \rho') = \sigma(\rho(x))$

$$\langle (e_0 e_1), \rho, \sigma, \kappa \rangle \mapsto \langle e_0, \rho, \sigma, \text{ar}(e_1, \rho) : \kappa \rangle$$

$$\langle v, \rho, \sigma, \text{ar}(e, \rho) : \kappa \rangle \mapsto \langle e, \rho, \sigma, \text{fn}(v, \rho) : \kappa \rangle$$

$$\langle v, \rho, \sigma, \text{fn}(\lambda x. e, \rho') : \kappa \rangle \mapsto \langle e, \rho'', \sigma', \kappa \rangle$$

where $\rho'' = \rho' [x \mapsto a]$

$$\sigma' = \sigma[a \mapsto (v, \rho)]$$

a fresh

$$\begin{aligned}
 \rho \in \text{Env} &= \text{Var} \rightarrow \text{Addr} \\
 \kappa \in \text{Kont} &::= [] \mid \varphi : \textcolor{blue}{a} \\
 \sigma \in \text{Store} &= \text{Addr} \rightarrow (\text{Value} \times \text{Env} + \text{Kont}) \\
 \langle x, \rho, \sigma, \kappa \rangle &\mapsto \langle v, \rho', \sigma, \kappa \rangle
 \end{aligned}$$

if $(v, \rho') = \sigma(\rho(x))$

$$\langle (e_0 e_1), \rho, \sigma, \kappa \rangle \mapsto \langle e_0, \rho, \textcolor{blue}{\sigma'}, \text{ar}(e_1, \rho) : \textcolor{blue}{a} \rangle \quad \sigma' = \sigma[a \mapsto \kappa]$$

$$\langle v, \rho, \sigma, \text{ar}(e, \rho) : \textcolor{blue}{b} \rangle \mapsto \langle e, \rho, \sigma, \text{fn}(v, \rho) : \textcolor{blue}{b} \rangle$$

$$\langle v, \rho, \sigma, \text{fn}(\lambda x. e, \rho') : \textcolor{blue}{b} \rangle \mapsto \langle e, \rho'', \sigma', \kappa \rangle \quad \kappa = \sigma(b)$$

where $\rho'' = \rho'[x \mapsto a]$

$$\sigma' = \sigma[a \mapsto (v, \rho)]$$

a fresh

$$\begin{aligned}
 \rho \in \text{Env} &= \text{Var} \rightarrow \text{Addr} \\
 \kappa \in \text{Kont} &::= [] \mid \varphi : a \\
 \sigma \in \text{Store} &= \text{Addr} \rightarrow \wp(\text{Value} \times \text{Env} + \text{Kont}) \\
 \langle x, \rho, \sigma, \kappa \rangle &\mapsto \langle v, \rho', \sigma, \kappa \rangle
 \end{aligned}$$

if $(v, \rho') \in \sigma(\rho(x))$

$$\langle (e_0 e_1), \rho, \sigma, \kappa \rangle \mapsto \langle e_0, \rho, \sigma', \text{ar}(e_1, \rho) : a \rangle \sigma' = \sigma \sqcup [a \mapsto \{\kappa\}]$$

$$\langle v, \rho, \sigma, \text{ar}(e, \rho) : b \rangle \mapsto \langle e, \rho, \sigma, \text{fn}(v, \rho) : b \rangle$$

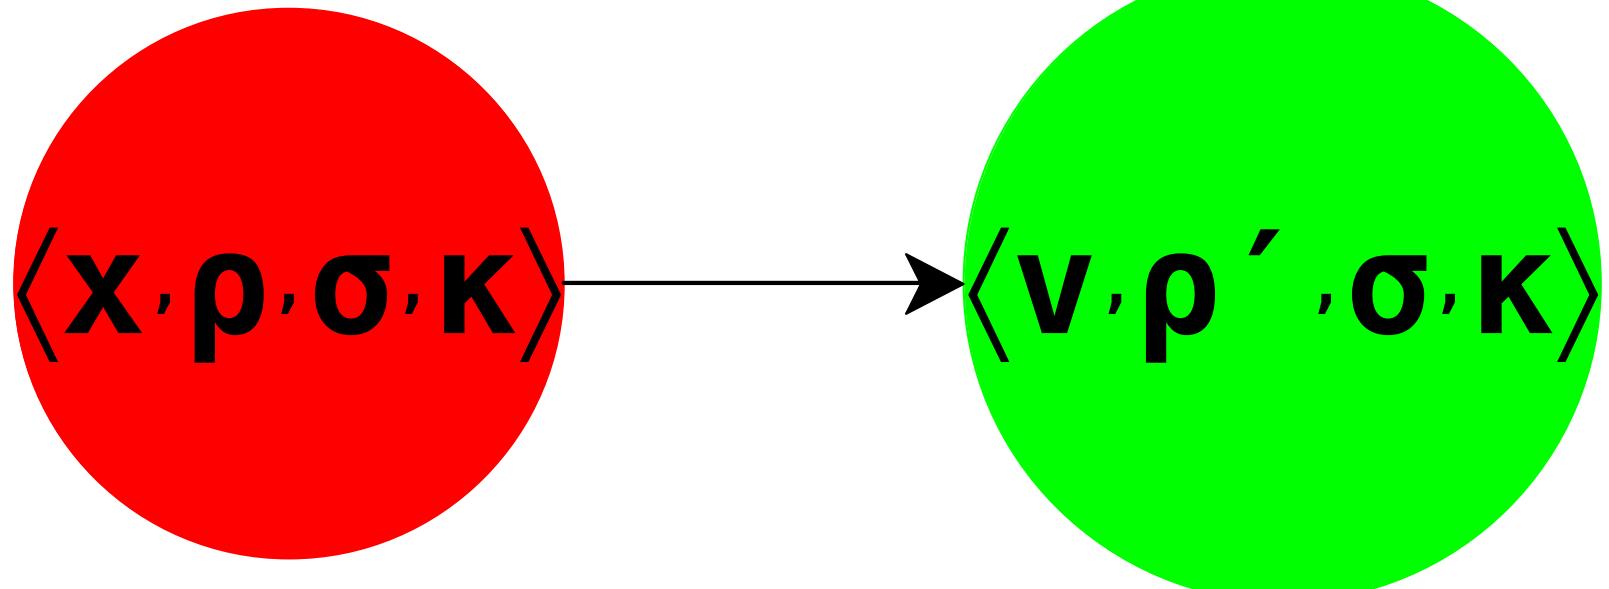
$$\langle v, \rho, \sigma, \text{fn}(\lambda x. e, \rho') : b \rangle \mapsto \langle e, \rho'', \sigma', \kappa \in \sigma(b) \rangle$$

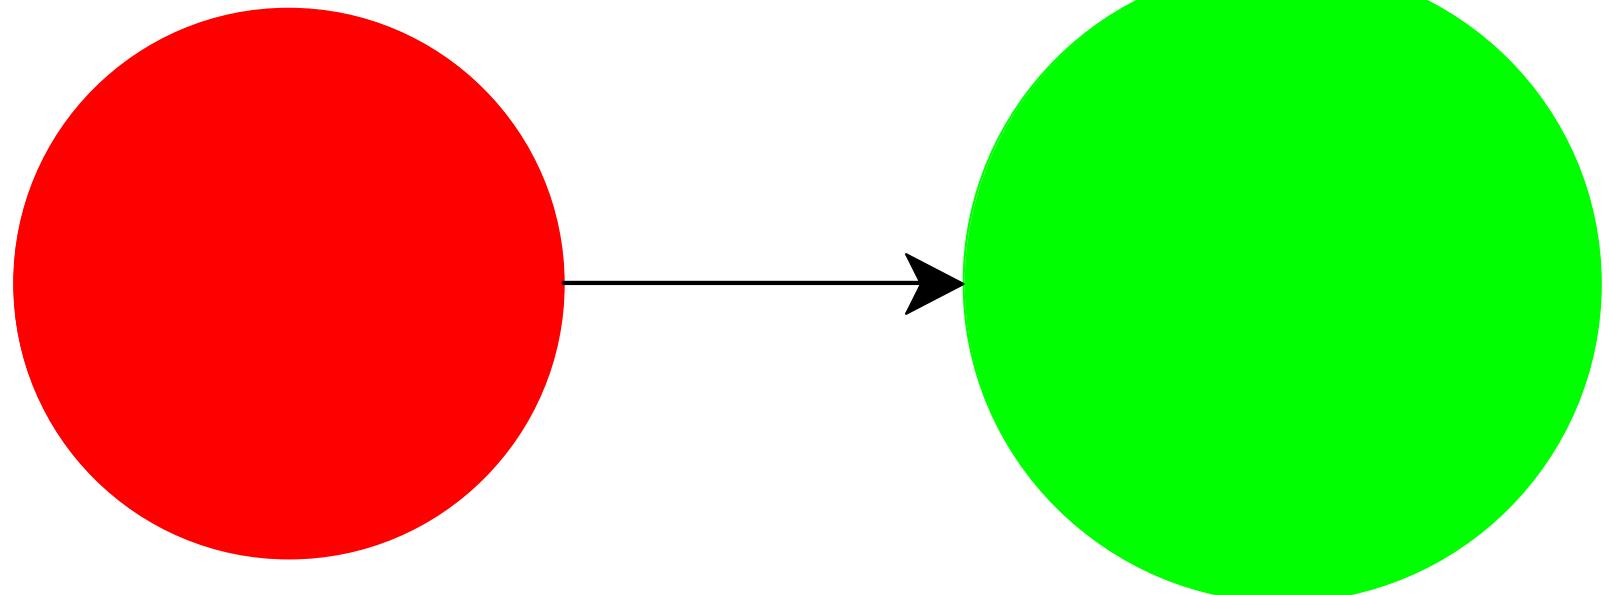
where $\rho'' = \rho' [x \mapsto a]$

$$\sigma' = \sigma \sqcup [a \mapsto \{(v, \rho)\}]$$

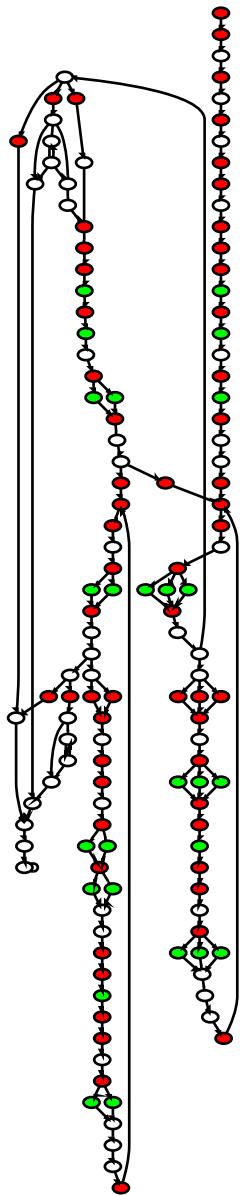
a = alloc(ς)

$$\langle x, \rho, \sigma, \kappa \rangle \rightarrow \langle v, \rho', \sigma, \kappa \rangle$$

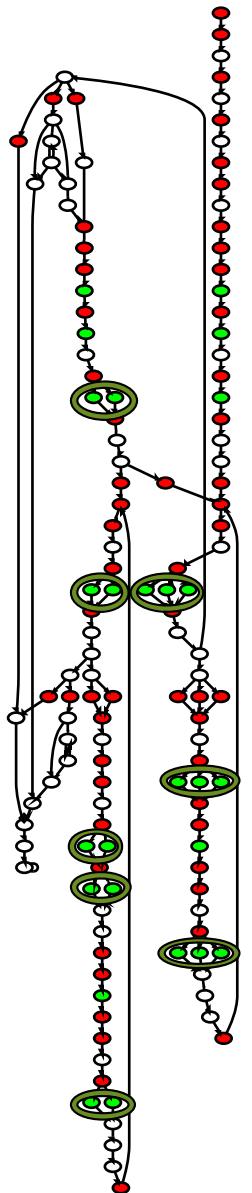




Finite reduction relation = graph

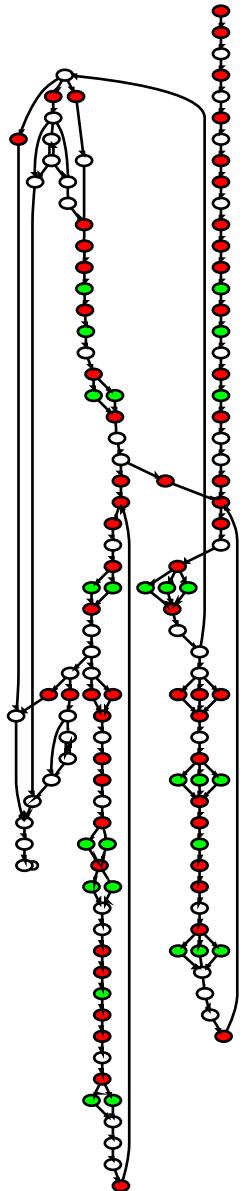


Finite reduction relation = graph

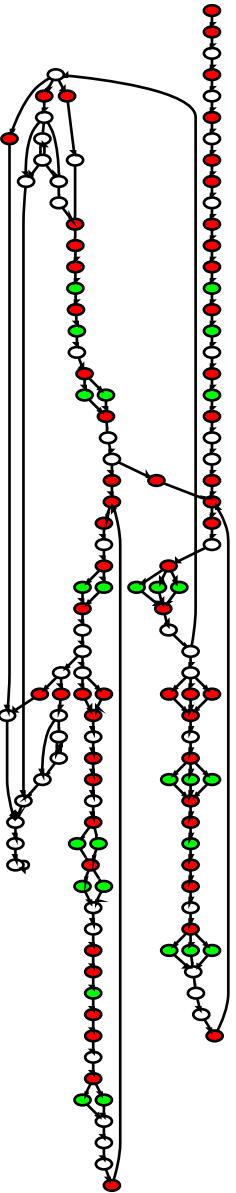


1. Lazy non-determinism

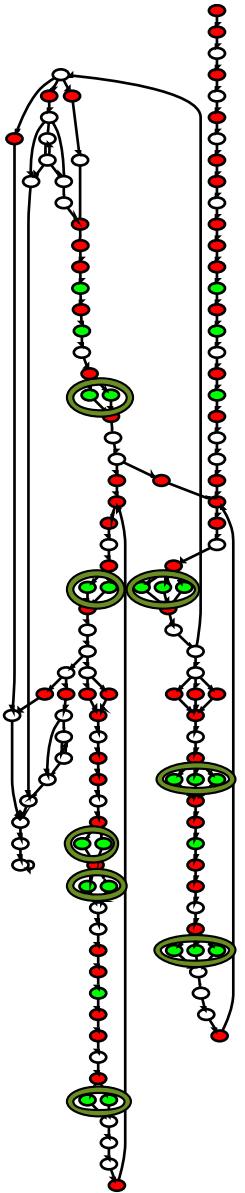
Finite reduction relation = graph



1. Lazy non-determinism
2. Abstract compilation



1. Lazy non-determinism
2. Abstract compilation
3. Store deltas



1. Lazy non-determinism
2. Abstract compilation
3. Store deltas

$$\begin{aligned}
 \rho \in \text{Env} &= \text{Var} \rightarrow \text{Addr} \\
 \kappa \in \text{Kont} &::= [] \mid \varphi : a \\
 \sigma \in \text{Store} &= \text{Addr} \rightarrow \wp(\text{Value} \times \text{Env} + \text{Kont}) \\
 \langle x, \rho, \sigma, \kappa \rangle &\mapsto \langle v, \rho', \sigma, \kappa \rangle
 \end{aligned}$$

if $(v, \rho') \in \sigma(\rho(x))$

$$\langle (e_0 e_1), \rho, \sigma, \kappa \rangle \mapsto \langle e_0, \rho, \sigma', \text{ar}(e_1, \rho) : a \rangle \quad \sigma' = \sigma \cup [a \mapsto \{\kappa\}]$$

$$\langle v, \rho, \sigma, \text{ar}(e, \rho) : b \rangle \mapsto \langle e, \rho, \sigma, \text{fn}(v, \rho) : b \rangle$$

$$\langle v, \rho, \sigma, \text{fn}(\lambda x. e, \rho') : b \rangle \mapsto \langle e, \rho'', \sigma', \kappa \rangle \quad \kappa \in \sigma(b)$$

where $\rho'' = \rho' [x \mapsto a]$

$$\sigma' = \sigma \cup [a \mapsto \{(v, \rho)\}]$$

a = alloc(ζ)

$$\begin{aligned}
 \rho \in \text{Env} &= \text{Var} \rightarrow \text{Addr} \\
 \kappa \in \text{Kont} &::= [] \mid \varphi : a \\
 \sigma \in \text{Store} &= \text{Addr} \rightarrow \wp(\text{Value} \times \text{Env} + \text{Kont}) \\
 \langle x, \rho, \sigma, \kappa \rangle &\mapsto \langle \text{addr}(\rho(x)), \sigma, \kappa \rangle
 \end{aligned}$$

$$\langle (e_0 e_1), \rho, \sigma, \kappa \rangle \mapsto \langle e_0, \rho, \sigma', \text{ar}(e_1, \rho) : a \rangle \quad \sigma' = \sigma \cup [a \mapsto \{\kappa\}]$$

$$\langle v, \rho, \sigma, \text{ar}(e, \rho) : b \rangle \mapsto \langle e, \rho, \sigma, \text{fn}(v, \rho) : b \rangle$$

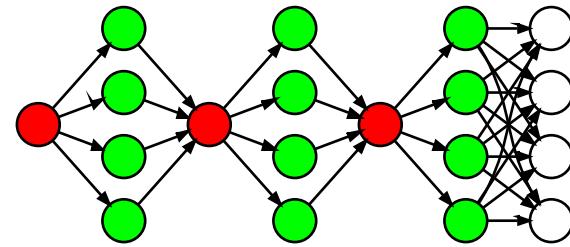
$$\langle v, \rho, \sigma, \text{fn}(\lambda x. e, \rho') : b \rangle \mapsto \langle e, \rho'', \sigma', \kappa \rangle \quad \kappa \in \sigma(b)$$

where $\rho'' = \rho' [x \mapsto a]$

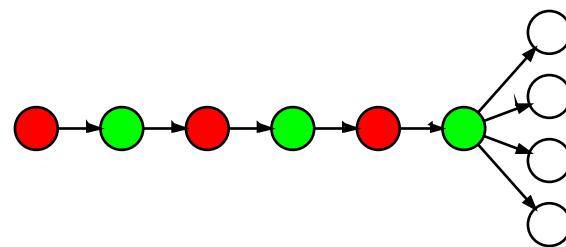
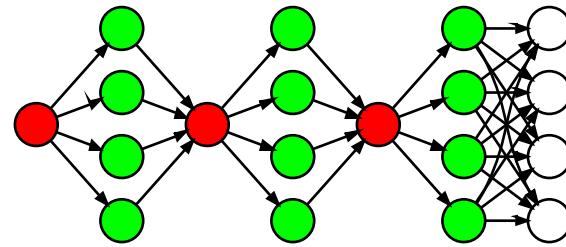
$$\sigma' = \sigma \cup [a \mapsto \text{force}(v)]$$

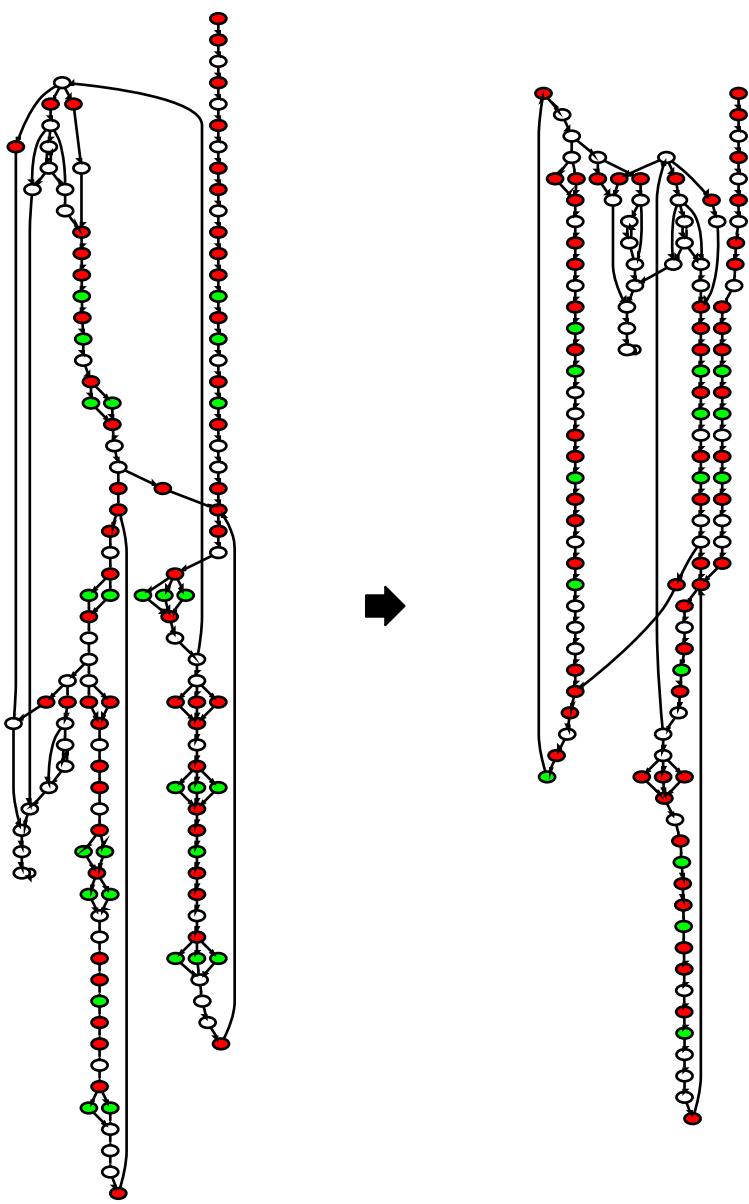
$a = \text{alloc}(\zeta)$

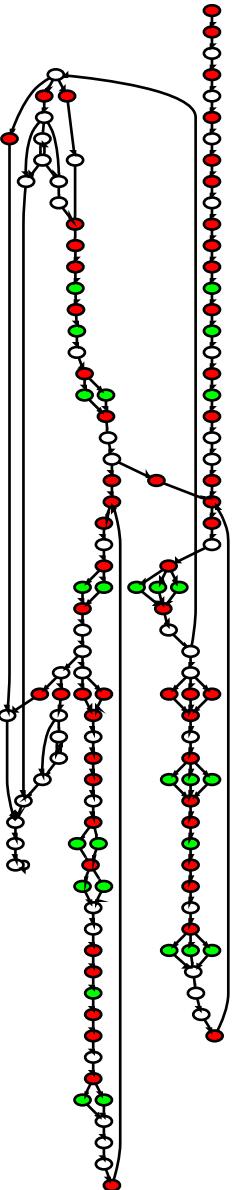
(\mathbf{f} \mathbf{x} \mathbf{y})



$(f \ x \ y)$







1. Lazy non-determinism
2. Abstract compilation
3. Store deltas

$\rho \in \text{Env} = \text{Var} \rightarrow \text{Addr}$ $\kappa \in \text{Kont} ::= [] \mid \varphi : a$ $\sigma \in \text{Store} = \text{Addr} \rightarrow \wp(\text{Value} \times \text{Env} + \text{Kont})$ $\langle x, \rho, \sigma, \kappa \rangle \mapsto \langle v, \rho', \sigma, \kappa \rangle$

if $(v, \rho') \in \sigma(\rho(x))$

 $\langle (e_0 e_1), \rho, \sigma, \kappa \rangle \mapsto \langle e_0, \rho, \sigma', \text{ar}(e_1, \rho) : a \rangle \quad \sigma' = \sigma \cup [a \mapsto \{\kappa\}]$ $\langle v, \rho, \sigma, \text{ar}(e, \rho) : b \rangle \mapsto \langle e, \rho, \sigma, \text{fn}(v, \rho) : b \rangle$ $\langle v, \rho, \sigma, \text{fn}(\lambda x. e, \rho') : b \rangle \mapsto \langle e, \rho'', \sigma', \kappa \rangle \quad \kappa \in \sigma(b)$

where $\rho'' = \rho' [x \mapsto a]$

 $\sigma' = \sigma \cup [a \mapsto \{(v, \rho)\}]$

a = alloc(ζ)

```
(match e ((var x) ((continue v σ κ) : v ∈ (get σ (get ρ x)))) ((
```

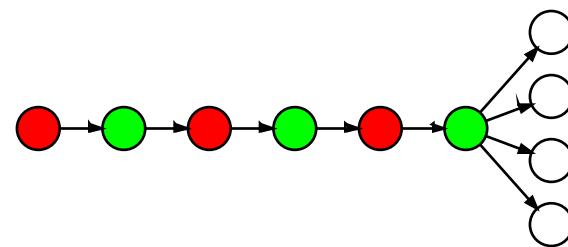
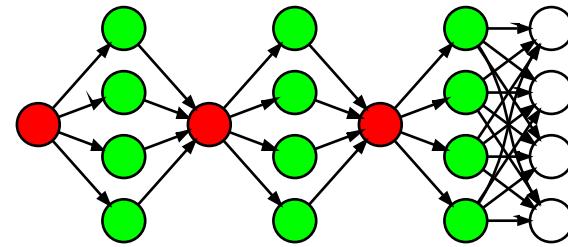

Compile away dispatch overhead

⟨e,ρ,σ,κ⟩

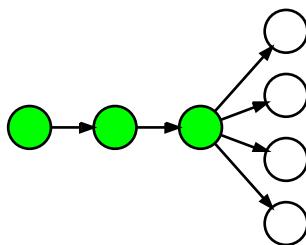
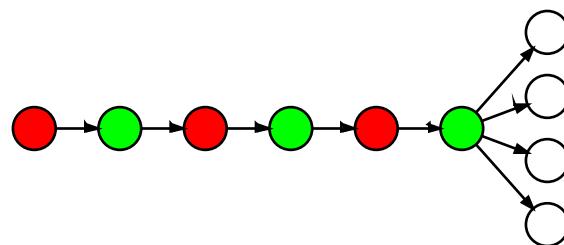
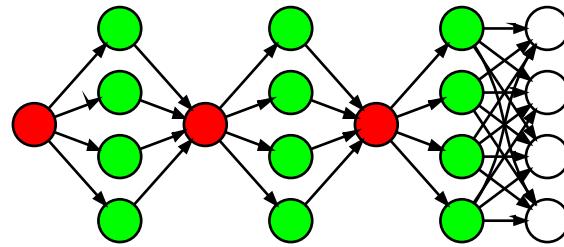
[[e]](ρ, σ, κ)

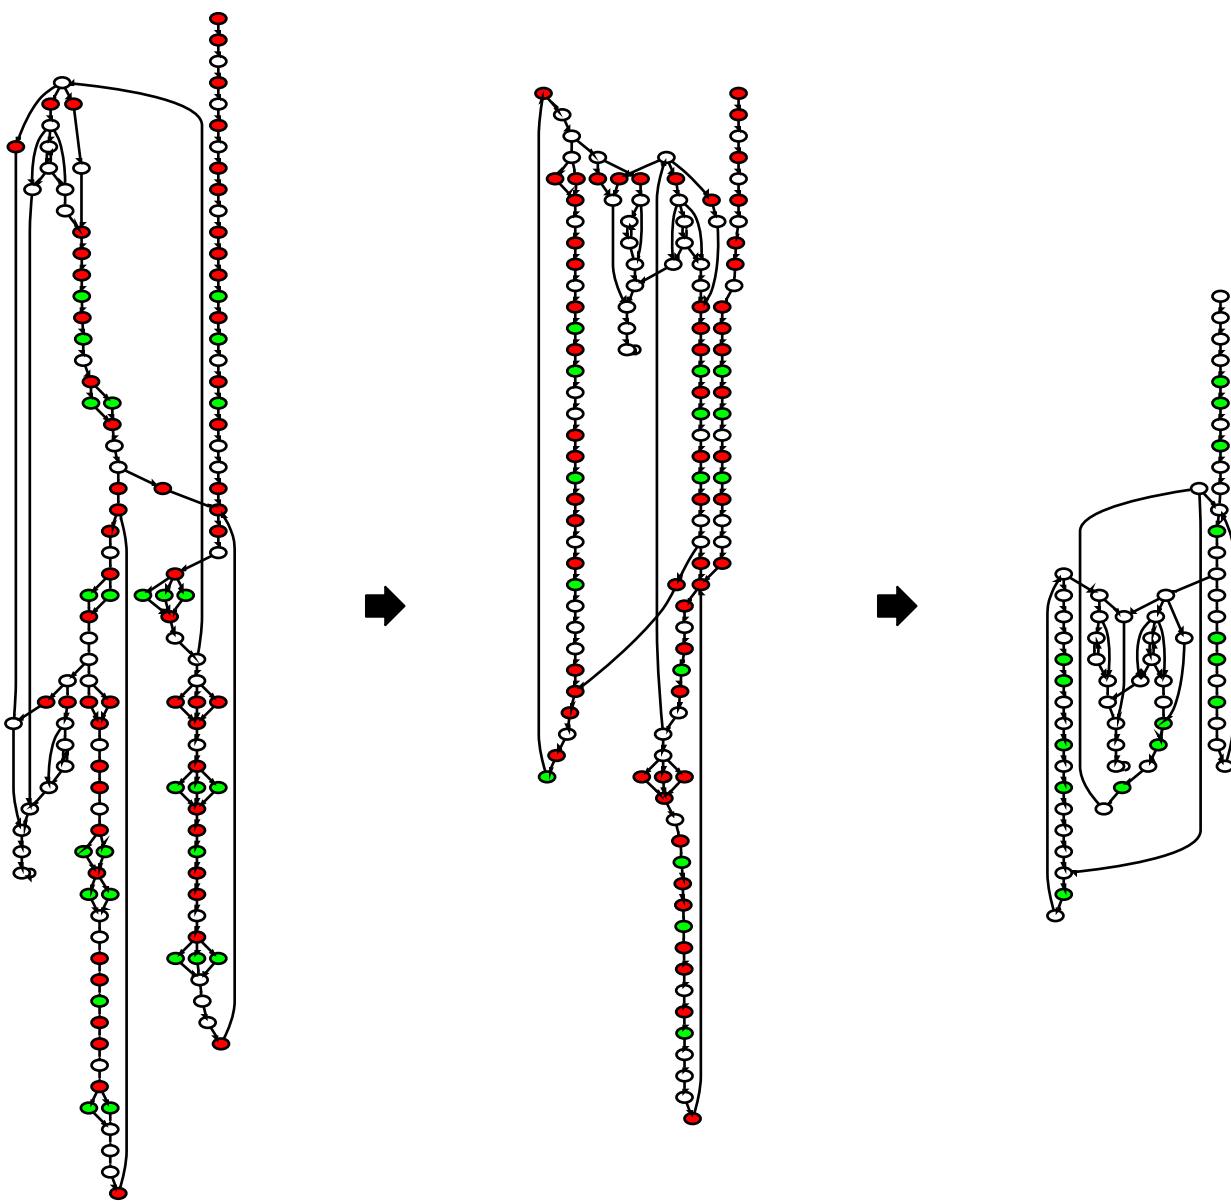
$\llbracket e \rrbracket(\rho, \sigma, \kappa)$ $\langle e, \rho, \sigma, \kappa \rangle \mapsto \text{RHS}$  $\llbracket e \rrbracket = \lambda \rho, \sigma, \kappa. \text{ RHS}$

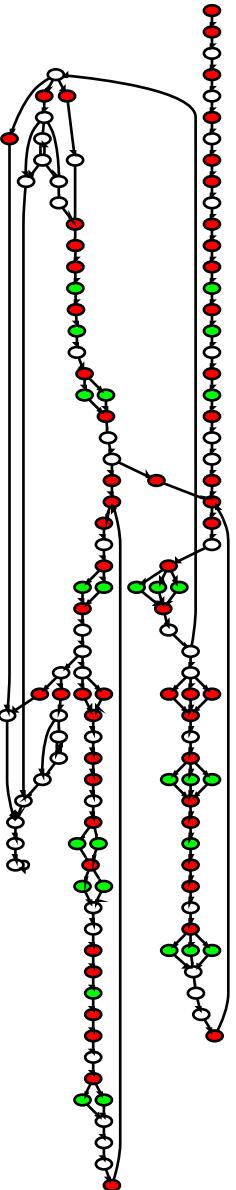
$(f \ x \ y)$



$(f \ x \ y)$







1. Lazy non-determinism
2. Abstract compilation
3. Store deltas

Store widening / Global store

↪ implemented as **step** : **State** → $\wp(\text{State})$

Lift and iterate to fixed point in $\wp(\text{State}) \rightarrow \wp(\text{State})$

$$\{\langle e_0, \rho_0, \sigma_0, K_0 \rangle, \quad \langle e'_0, \rho'_0, \sigma'_0, K'_0 \rangle,$$

:

:

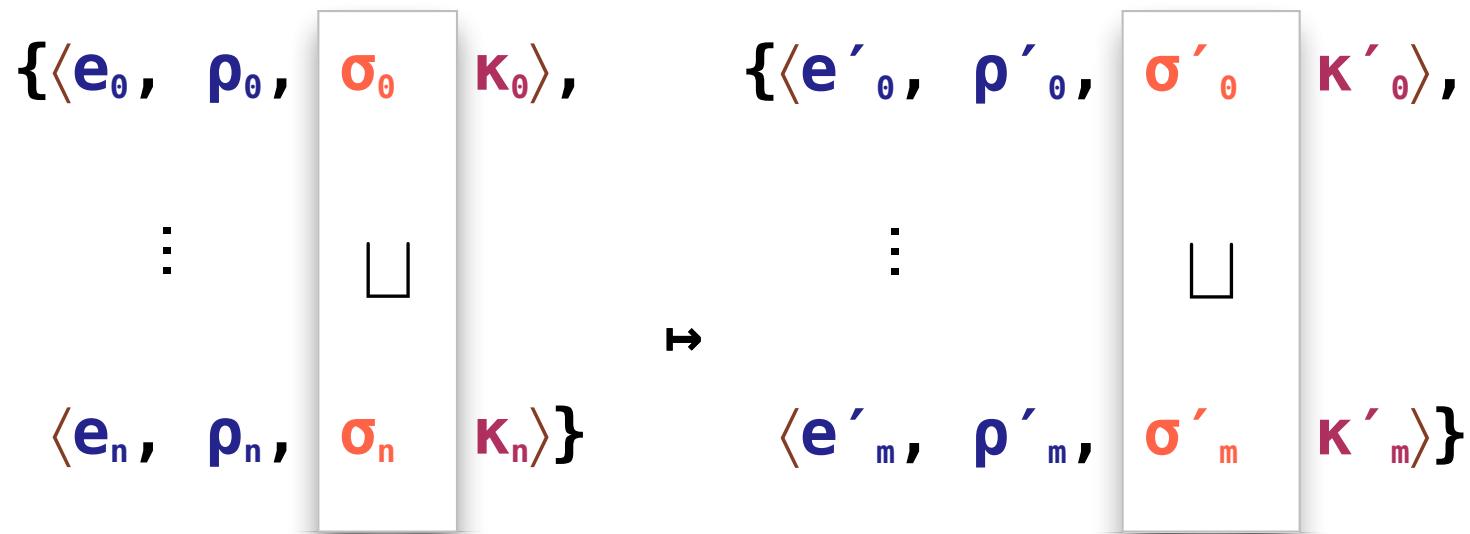
↪

$$\langle e_n, \rho_n, \sigma_n, K_n \rangle\} \quad \langle e'_m, \rho'_m, \sigma'_m, K'_m \rangle\}$$

Store widening / Global store

↪ implemented as **step** : **State** → $\wp(\text{State})$

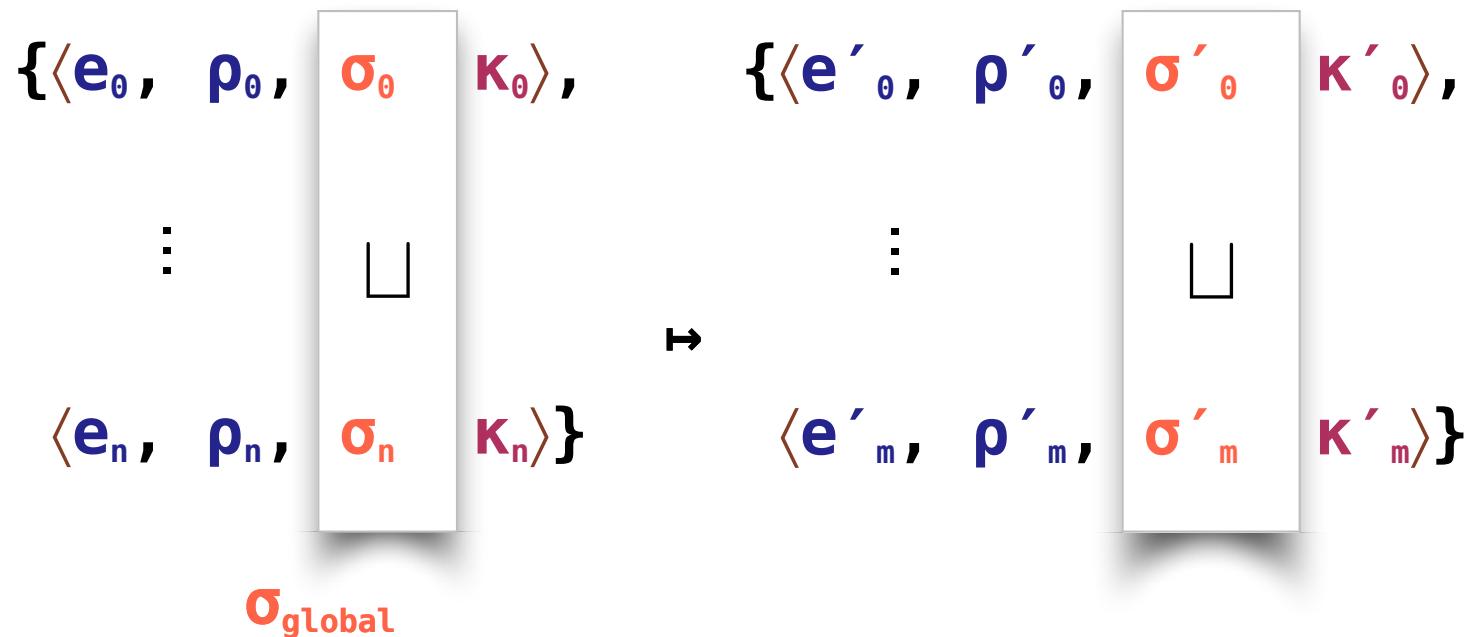
Lift and iterate to fixed point in $\wp(\text{State}) \rightarrow \wp(\text{State})$



Store widening / Global store

↪ implemented as **step** : **State** → $\wp(\text{State})$

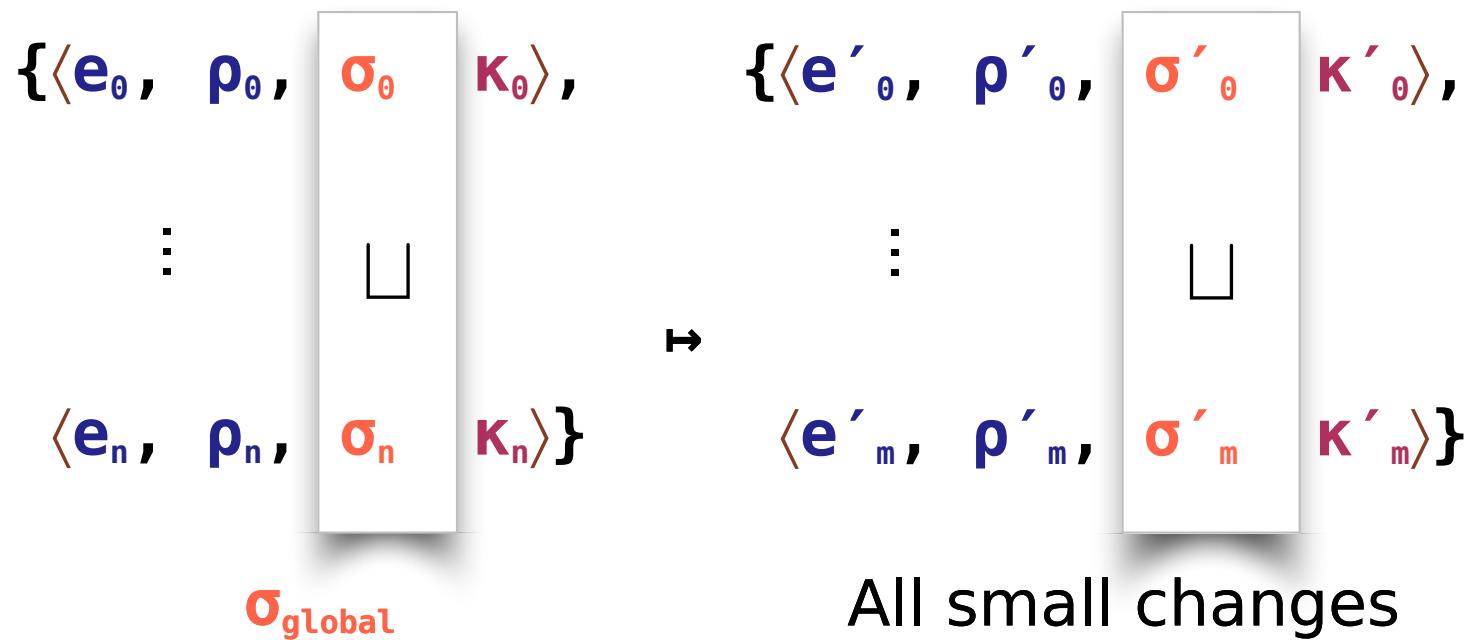
Lift and iterate to fixed point in $\wp(\text{State}) \rightarrow \wp(\text{State})$



Store widening / Global store

↪ implemented as **step** : **State** → $\wp(\text{State})$

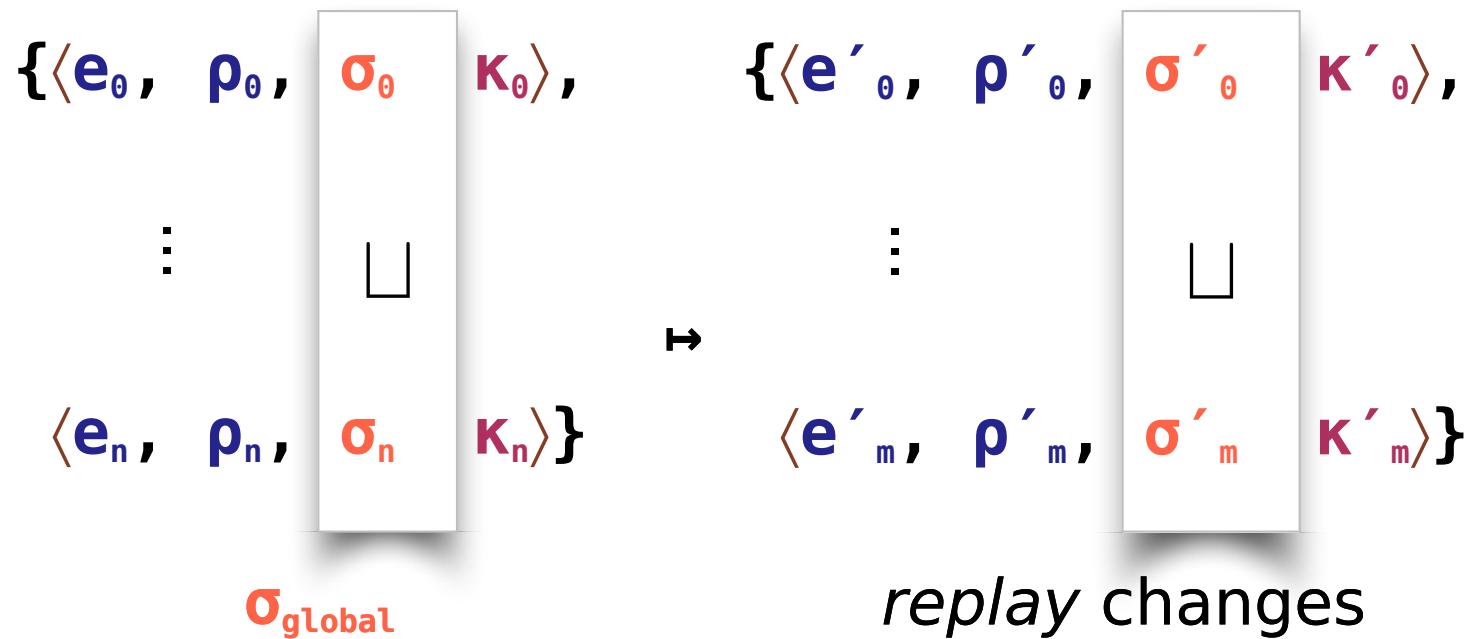
Lift and iterate to fixed point in $\wp(\text{State}) \rightarrow \wp(\text{State})$



Store widening / Global store

↪ implemented as **step** : **State** → $\wp(\text{State})$

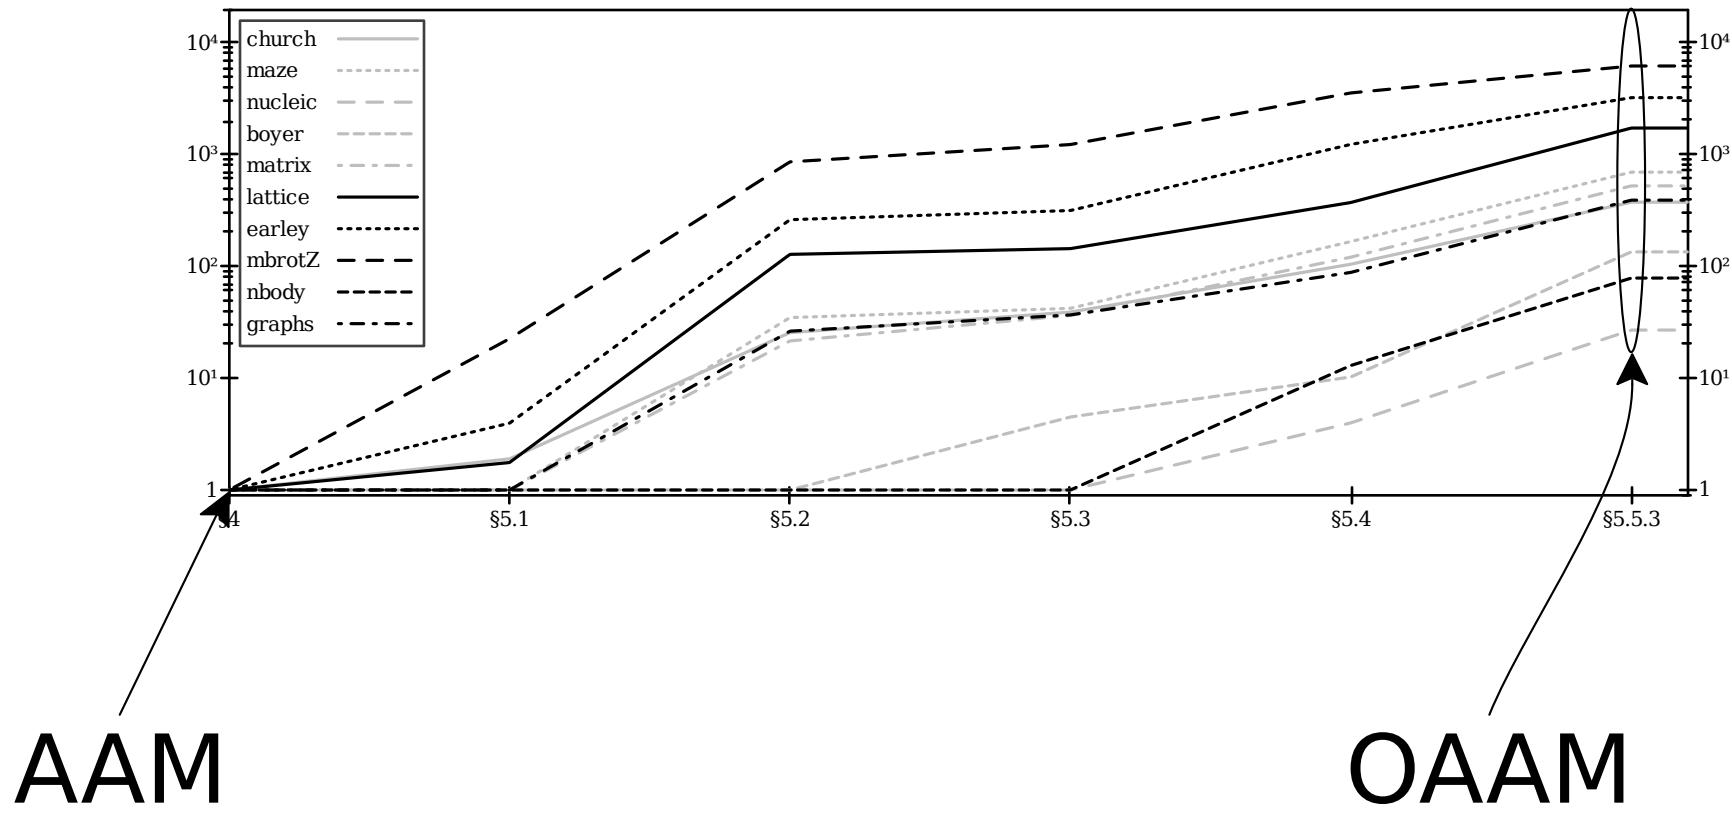
Lift and iterate to fixed point in $\wp(\text{State}) \rightarrow \wp(\text{State})$



What does all this buy us?

100000% improvement

Factor speed-up over naive vs. paper section



AAM

OAAM

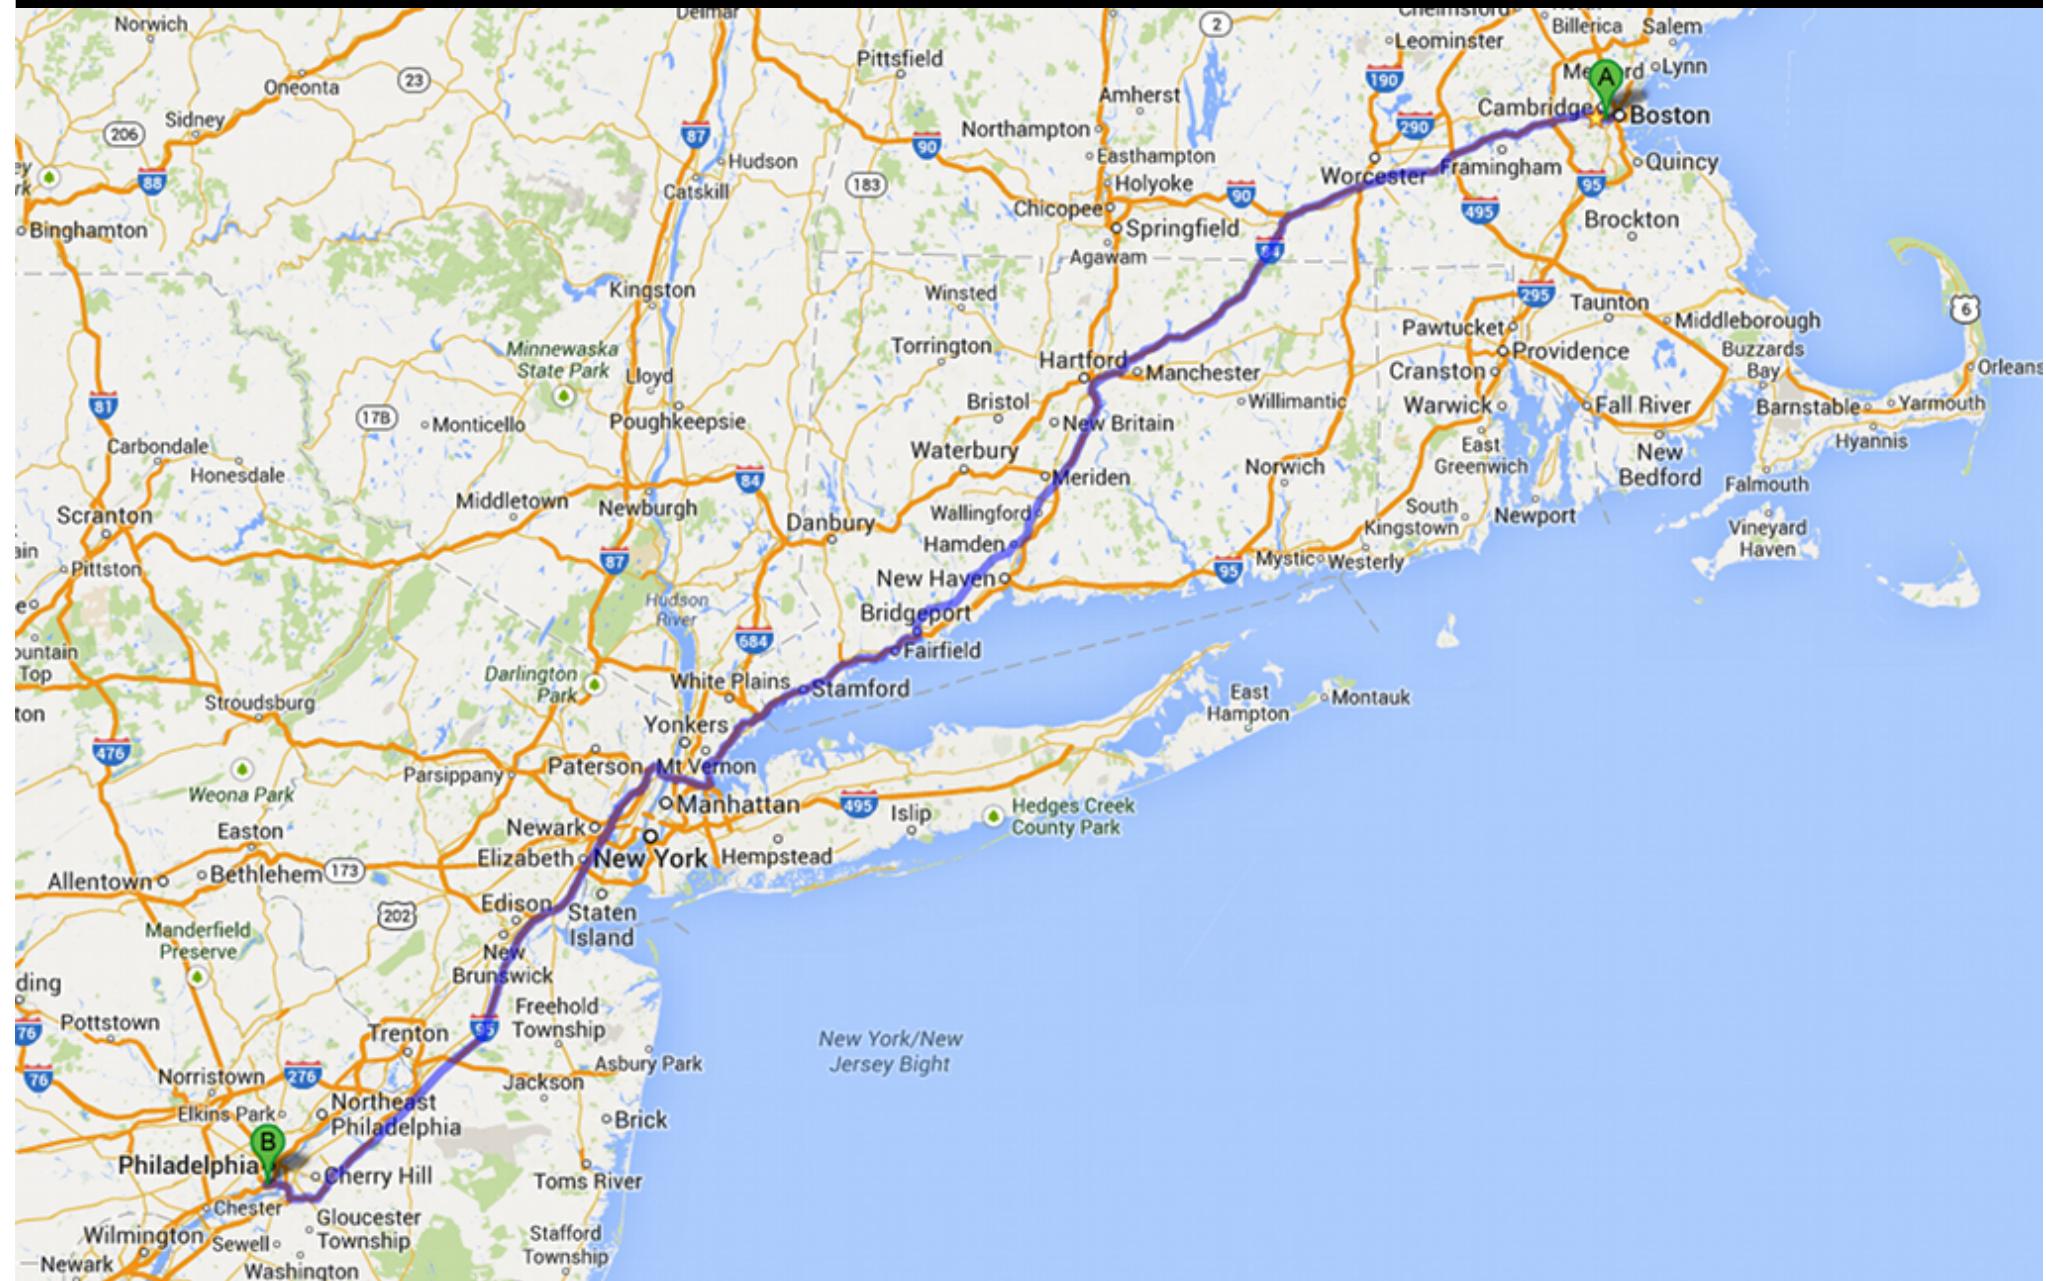
No change in evaluated precision

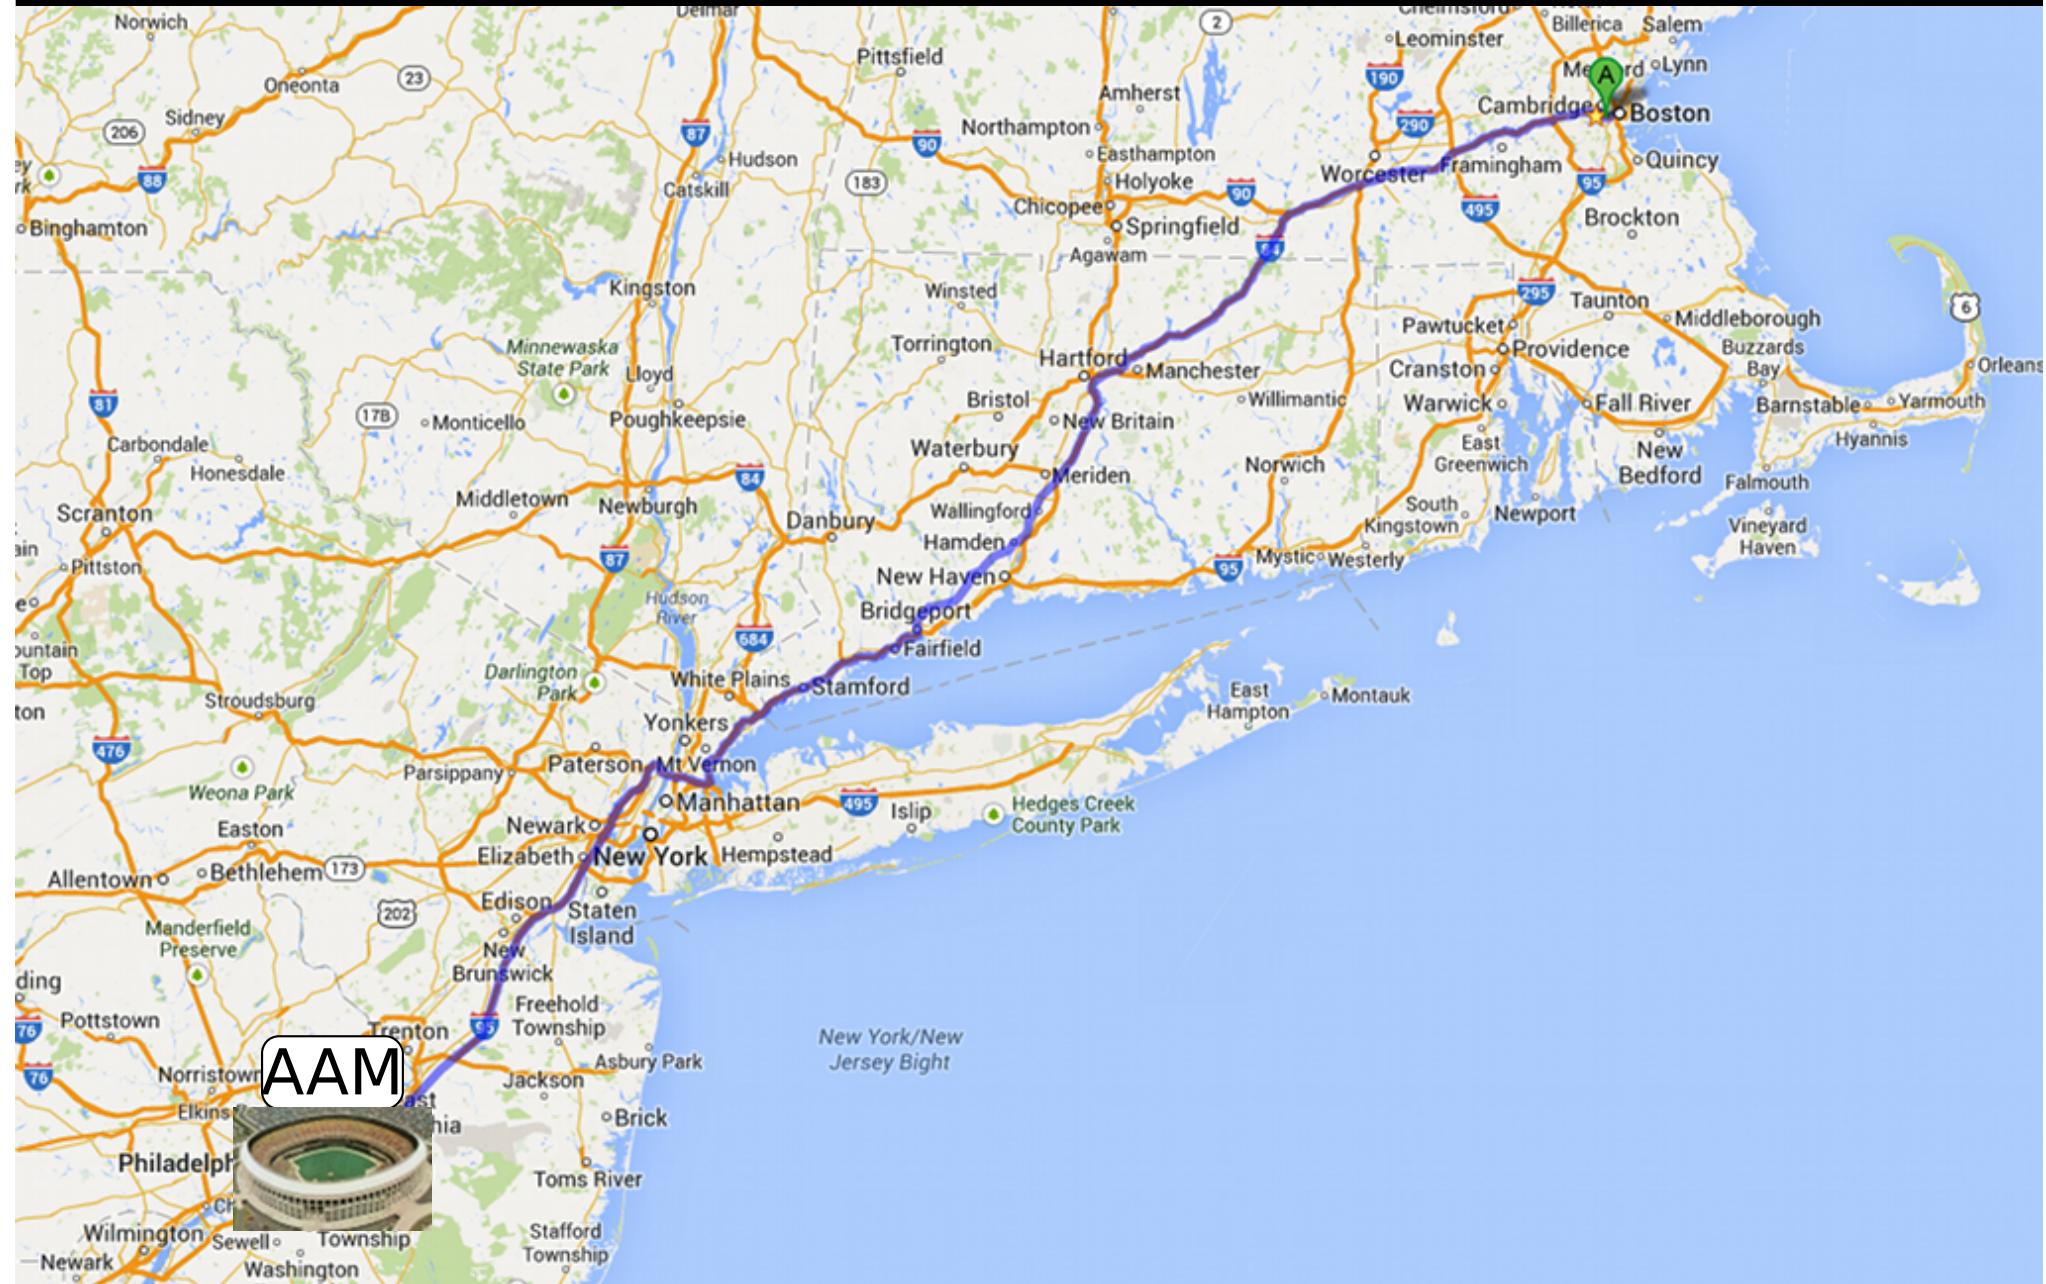


OAAM



Hand-optimized





The evolution of AAM to OAAM

Semantics



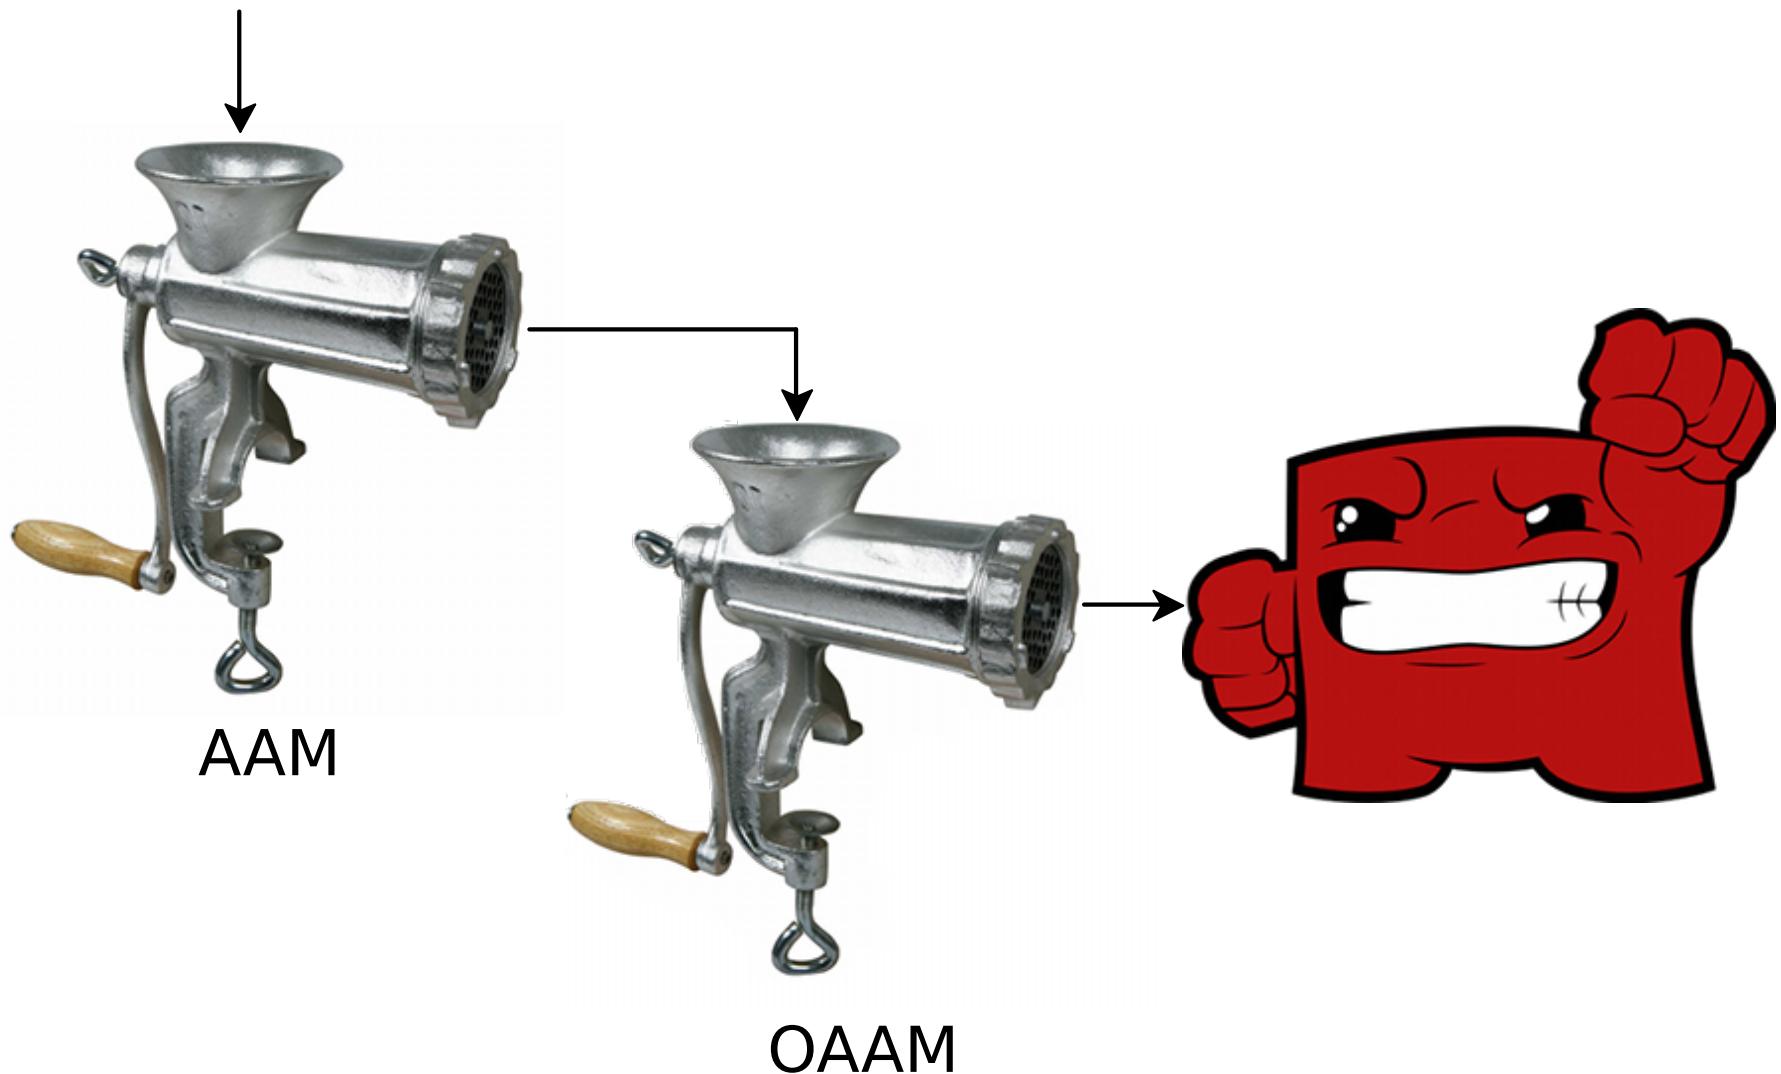
AAM



OAAM

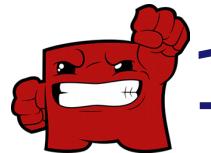
The evolution of AAM to OAAM

Semantics



Before we part

Simple, systematic implementation 



1000x speedup

"A simple matter of engineering?"

Build your tricks into the semantics

Thank you