

Layout Randomization and Nondeterminism

Martín Abadi Jérémy Planul Gordon Plotkin

PrakashFest, Oxford, May 2014

Hoping to work with Prakash: Choquet capacities

Josée, Vineet, Radha, and Prakash, CONCUR '02:
Weak Bisimulation is Sound and Complete for pCTL* .

- Superadditive Choquet capacity:

$$\mu(A \cup B) \leq \mu(A) + \mu(B) \quad (A, B \subseteq \Sigma, \text{ disjoint})$$

- Given measures μ_i , their **lower envelope**:

$$\mu(A) =_{\text{def}} \min_i \mu_i(A)$$

is a superadditive Choquet capacity.

- Can we do ordinary + probabilistic nondeterminism this way?

Hoping to work with Prakash: Choquet capacities

Josée, Vineet, Radha, and Prakash, CONCUR '02:
Weak Bisimulation is Sound and Complete for pCTL* .

- Superadditive Choquet capacity:

$$\mu(A \cup B) \leq \mu(A) + \mu(B) \quad (A, B \subseteq \Sigma, \text{ disjoint})$$


- Given measures μ_i , their **lower envelope**:

$$\mu(A) =_{\text{def}} \min_i \mu_i(A)$$


is a superadditive Choquet capacity.

- Can we do ordinary + probabilistic nondeterminism this way?
- Sadly, this map is many-to-one from convex sets of measures to capacities, even when $|\Sigma| = 3$.

Barycentric coordinates




Cartesian coordinates



Barycentric coordinates

Counterexamples in plenty (Cozman)



Trying to work with Prakash: quantum computation

A Language and Type Theory for Quantum Computing

Prakash Panangaden*
School of Computer Science
McGill University
Montreal, Quebec, Canada

Gordon D. Plotkin†
Department of Computer Science
University of Edinburgh
Edinburgh, Scotland, U. K.

February 28, 2003

Abstract

The abstract is very uncertain.

1 Introduction

Outline the nature of quantum comp as a web of interactions Our philosophy: evolution is fundamental, density ops, intervention operators are fundamental, no relativistic effects, entanglement, expressing unitaries; cite Peres and NC

Computation is dynamics. One prepares a dynamical system in some initial state - implicitly or explicitly containing the “input” - and lets it evolve. The “answer” is contained in the final state. Understanding quantum computation must, at its lowest level, involve understanding quantum evolution.

Normally in quantum mechanics understanding quantum evolution means understanding the generator of time evolution, i.e. the Hamiltonian operator. If H is the Hamiltonian operator for a system then the evolution of states is governed by the unitary operator e^{-iHt} . Typically this is cast as a differential equation, the Schrödinger equation. In typical text books one spends a significant amount of time solving the Schrödinger equation for a variety of Hamiltonians.

In quantum computation textbooks [NC00] one takes various unitary operations on states of collections two-state systems (qubits) as given. These two-state systems, or qubits, are conceived of as separate carriers of information, just like classical bits. Physically they are generally realized as spin-states of spin $\frac{1}{2}$ particles or as polarization states of photons. What is very different about qubits versus bits is that (i) a qubit may exist in a superposed state; i.e. in a linear combination of its two basic states and (ii) states of different qubits may be *entangled* i.e. correlated in nontrivial ways.

Algorithms are developed by combining the primitive unitary operators and occasionally by making measurements. However, the primitive unitaries do not all just act on one qubit at a time. In some computations of the primitive, there is at least one primitive that involves multiple qubits.

Somewhat working with Prakash: Markov processes

Approximating Markov Processes By Averaging

PHILIPPE CHAPUT
McGill University
VINCENT DANOS
University of Edinburgh
PRAKASH PANANGADEN
McGill University
GORDON PLOTKIN
University of Edinburgh

October 14, 2013

Abstract

Normally, one thinks of probabilistic transition systems as taking an initial probability distribution over the state space into a new probability distribution representing the system after a transition. We, however, take a dual view of Markov processes as transformers of bounded measurable functions. This is very much in the same spirit as a “predicate-transformer” view, which is dual to the state-transformer view of transition systems.

We redevelop the theory of labelled Markov processes from this view point, in particular we explore approximation theory. We obtain three main results:

(i) It is possible to define bisimulation on general measure spaces and show that it is an equivalence relation. The logical characterization of bisimulation can be done straightforwardly and generally. (ii) A new and flexible approach to approximation based on averaging can be given. This vastly generalizes and streamlines the idea of using conditional expectations to compute approximations. (iii) We show that there is a minimal process bisimulation-equivalent to a given process, and this minimal process is obtained as the limit of the finite approximants.

Low-Level Attacks and Protection


- Many attack techniques:
 - Buffer overflows
 - Exception overwrites
 - return-to-libc
 - jump-to-libc
 - use-after-free attacks
- Often with knowledge and control of the heap
- Many defenses:
 - Stack canaries
 - Safe exception handling
 - NX (No eXecute) data
 - Layout randomization
- Useful mitigations
- But not necessarily perfect in a precise sense
- Nor all well understood

Layout randomization

Runtime attacks often depend on addresses (e.g. jump to `libc`).

- Let us randomize the addresses!
 - Considered for data at least since the rise of large virtual address spaces (e.g., [Druschel & Peterson, 1992] on fbufs).
 - Now present in Linux (PaX), Windows, Mac OS X, iOS, Android (4.0).

`jump0x4`




A theory of layout randomization

- Abadi and Plotkin did a first formalization, using operational methods.
- Further work [Jagadeesan et al., Abadi et al.] concerned memories containing functions and arrays.
- Here we consider commands that make nondeterministic choices (as a first step towards parallelism).
- We use denotational methods rather than operational ones (for interest, but not only).
- We use random variables for probability rather than distributions, as events are not independent.


Probabilities and nondeterminism

Probabilities:



Not red with
 ≥ 0.9
probability

Nondeterminism:



Never red

- Probabilistic cryptographic schemes
- Differential privacy
- Low-level software protection (Layout randomization)

- Network communications
- Nondeterministic scheduling
- User interaction

Probabilities \longleftrightarrow Nondeterminism: Game theory

- Player 1 chooses a bit (b_1) by flipping a coin
- Player 2 chooses a bit (b_2) nondeterministically
- Player 1 wins if he guesses right:

$$b_1 = b_2$$

- The program:


$$b_1 := 0 + \frac{1}{2} 1; (b_2 := 0 \text{ OR } b_2 := 1)$$

Probabilities \longleftrightarrow Nondeterminism: Game theory

$$b_1 := 0 + \frac{1}{2} 1; (b_2 := 0 \text{ OR } b_2 := 1)$$


What should its semantics be?

- a set of distributions on outcomes



- Player I always wins with probability 1/2.

- a distribution on sets of outcomes




- Player I can win or lose.

Probabilities \leftrightarrow Nondeterminism: Game theory

$$b_1 := 0 + \frac{1}{2} 1; (b_2 := 0 \text{ OR } b_2 := 1)$$

The obvious semantics for sets of distributions on outcomes does not compose as expected:



Probabilities \longleftrightarrow Nondeterminism: Algebra

- Resolving probabilistic choice first:
 - Distribute nondeterministic choice over probabilistic choice:

$$x \cup (y +_p z) = (x \cup y) +_p (x \cup z)$$

- One can prove that for all p, q : $x +_p y = x +_q y$ so the $+_p$ become a second semilattice, distributing over the first one.
 - Free Algebra over a set X : Non-empty finite \cup -closed sets of non-empty finite subsets of X .
- Resolving nondeterministic choice first:
 - Distribute probabilistic choice over nondeterministic choice:

$$x +_p (y \cup z) = (x +_p y) \cup (x +_p z)$$

- Free Algebra over a set X : Non-empty finitely generated convex sets of distributions over X .
- **But gives the "wrong" semantics to our program.**

Roadmap

- Define **high-level** and **low-level** languages, each with access to (the same) private and public locations.
⇒ View the implementations as compilations.
- View attackers as (high- or low-level) “public” contexts: those with no access to private locations.
⇒ Relate commands depending on their public behaviour in public contexts.
- Characterize the contextual relations using semantical simulations (a kind of logical relation)
⇒ prove that high-level contextual relations correspond to low-level contextual relations (full abstraction):

$$\begin{aligned} \llbracket c \rrbracket \preceq \llbracket c' \rrbracket &\iff \forall \text{high-level public } C[\cdot]. \llbracket C[c] \rrbracket \leq_L \llbracket C[c'] \rrbracket \\ \Downarrow & \\ \llbracket c^\downarrow \rrbracket \preceq \llbracket c'^\downarrow \rrbracket &\iff \forall \text{low-level public } C[\cdot]. \llbracket C[c^\downarrow] \rrbracket \leq_L \llbracket C[c'^\downarrow] \rrbracket \end{aligned}$$

High-level language

$$\begin{aligned} e &::= k \mid !l_{\text{loc}} \mid e + e \mid e * e \\ b &::= e \leq e \mid \neg b \mid \text{true} \mid \text{false} \mid b \vee b \mid b \wedge b \\ c &::= l_{\text{loc}} := e \mid \text{if } b \text{ then } c \text{ else } c \mid \text{skip} \mid c; c \mid \\ &\quad c + c \mid \text{while } b \text{ do } c \end{aligned}$$

- Finitely many locations.
- They are divided into **private** (high) and **public** (low).
- store s : locations \rightarrow values
- Semantics $\llbracket c \rrbracket$: stores $\rightarrow \mathcal{H}(\text{stores}_{\perp})$


High-level language

$$\begin{aligned} e &::= k \mid !l_{loc} \mid e + e \mid e * e \\ b &::= e \leq e \mid \neg b \mid \text{true} \mid \text{false} \mid b \vee b \mid b \wedge b \\ c &::= l_{loc} := e \mid \text{if } b \text{ then } c \text{ else } c \mid \text{skip} \mid c; c \mid \\ &\quad c + c \mid \text{while } b \text{ do } c \end{aligned}$$

- Finitely many locations.
- They are divided into **private** (high) and **public** (low).
- store s : locations \rightarrow values
- Semantics $\llbracket c \rrbracket$: stores $\rightarrow \mathcal{H}(\text{stores}_{\perp})$

Refinement relation (high-level)

- For $X, Y \in \mathcal{H}(\text{stores}_{\perp})$, X **refines** Y if the public part of any store in X is the public part of a store in Y .



Formally:

$$X \leq_L Y \equiv \forall s \in \text{stores}. s_L \in X \Rightarrow s_L \in Y$$

- For commands, c **refines** c' , written $c \leq_L c'$ when, for every input, the outputs of c refine those of c' :

$$\forall s \in \text{stores}. \llbracket c \rrbracket(s) \leq_L \llbracket c' \rrbracket(s)$$

Simulation (high-level)


- Simulation relation \preceq on commands such that

Theorem 1


For all high-level commands c, c' , we have:

$$\llbracket c \rrbracket \preceq \llbracket c' \rrbracket \iff \forall \text{ high-level public } C[\cdot]. \llbracket C[c] \rrbracket \leq_L \llbracket C[c'] \rrbracket$$

- Simple without nondeterminism or nontermination:



Simulation (details)



- So begin with $R \subseteq \mathcal{H}(\text{private-stores}_\perp)^2$, subject to conditions
- Then construct $R^+ \subseteq \mathcal{H}(\text{stores}_\perp)^2$ from R
- Then

$$[[c]] \preceq [[c']] \equiv_{\text{def}} \forall X, Y \in \mathcal{H}(\text{stores}_\perp). XR^+Y \Rightarrow [[c]](X)R^+[[c']](Y)$$

Low-level memory model

- **Memories** ($m \in \text{Mem}$) are partial mappings from the set of **memory addresses** $\text{Add} = \{1, \dots, r\}$ to values.
- **Memory layouts** ($w \in \text{Lay}$) are injective mappings from locations to the addresses Add .

\implies stores map to memories via layouts.

$$s \begin{cases} l_1 \mapsto \color{red}\blacksquare \\ l_2 \mapsto \color{yellow}\blacksquare \\ l_3 \mapsto \color{cyan}\blacksquare \end{cases} \qquad w \begin{cases} l_1 \mapsto 7 \\ l_2 \mapsto 28 \\ l_3 \mapsto 14 \end{cases}$$




Layout randomization

- We fix a public layout w_p .
- We randomize the private layout (using a distribution d)
- We rely on one value in our theorems, the minimum probability for a guessed private address to be unallocated:

$$\begin{aligned}\delta &= \min\{P(i \notin \text{ran}(w)) \mid i \in \text{Add} \setminus \text{ran}(w_p)\} \\ &= 1 - |\text{PriLoc}|/|\text{Add} \setminus \text{ran}(w_p)| \quad (d \text{ uniform})\end{aligned}$$

$$s \begin{cases} l_1 \mapsto \text{red} \\ l_2 \mapsto \text{yellow} \\ l_3 \mapsto \text{cyan} \end{cases} \quad w_3 \begin{cases} l_1 \mapsto 7 \\ l_2 \mapsto 28 \\ l_3 \mapsto 14 \end{cases}$$



Low-level language

$$\begin{aligned} e & ::= k \mid l_{\text{nat}} \mid !e \mid e + e \mid e * e \\ b & ::= e \leq e \mid \neg b \mid \text{true} \mid \text{false} \mid b \vee b \mid b \wedge b \\ c & ::= e := e \mid \text{if } b \text{ then } c \text{ else } c \mid \text{skip} \mid c; c \mid \\ & \quad c + c \mid \text{while } b \text{ do } c \end{aligned}$$

- We can compute over l_{nat} , ex: $(x + l_{\text{nat}}) :=_{\text{nat}} (x + l'_{\text{nat}})$.

Low-level language

$$\begin{aligned} e &::= k \mid l_{\text{nat}} \mid !e \mid e + e \mid e * e \\ b &::= e \leq e \mid \neg b \mid \text{true} \mid \text{false} \mid b \vee b \mid b \wedge b \\ c &::= e := e \mid \text{if } b \text{ then } c \text{ else } c \mid \text{skip} \mid c; c \mid \\ &\quad c + c \mid \text{while } b \text{ do } c \end{aligned}$$

- We can compute over l_{nat} , ex: $(x + l_{\text{nat}}) :=_{\text{nat}} (x + l'_{\text{nat}})$.


Low-level semantics

- The final memory depends on the initial one and the layout,
- We cannot prove security using a semantics of the form:

$$\text{Mem} \times \text{Lay} \rightarrow \mathcal{H}(\text{Mem}_{\xi, \perp})$$

- for consider the command:

$$(1 := 1) + (2 := 1) + (3 := 1) + (4 := 1)$$



Low-level semantics

- So we need to resolve probabilities after nondeterminism:
- However

$$\text{Mem} \rightarrow \mathcal{H}(\mathcal{V}(\text{Mem}_{\xi, \perp}))$$

does not work as our events are not independent.

- For example, suppose there is just one private location l_1 , and 1,2 are not public addresses. Then:
 - $1 := 1 ; 2 := 1$ should always give an error.
 - $1 := 1 ; 1 := 2$ should give an error or overwrite l_1

Low-level semantics

- Solution idea: replace distributions

$$d \in \mathcal{V}(\text{Mem}_{\xi, \perp})$$

by **random variables**

$$\zeta : \text{Lay} \rightarrow \text{Mem}_{\xi, \perp}$$

with layouts as the sample space, obtaining:

$$\text{Mem} \rightarrow \mathcal{H}(\text{Lay} \rightarrow \text{Mem}_{\xi, \perp})$$

which compose, using Kleisli structure on \mathcal{H} .

Low-level semantics

- Solution idea: replace distributions

$$d \in \mathcal{V}(\text{Mem}_{\xi, \perp})$$

by **random variables**

$$\zeta : \text{Lay} \rightarrow \text{Mem}_{\xi, \perp}$$

with layouts as the sample space, obtaining:

$$\text{Mem} \rightarrow \mathcal{H}(\text{Lay} \rightarrow \text{Mem}_{\xi, \perp})$$

- Sadly, these do not compose (cf intro) and we instead use:


$$(\text{Lay} \rightarrow \text{Mem}_{\xi, \perp}) \rightarrow \mathcal{H}(\text{Lay} \rightarrow \text{Mem}_{\xi, \perp})$$

which compose, using Kleisli structure on \mathcal{H} .

Low level semantics

- Example:

$$(1:=1) + (2:=1) + (3:=1) + (4:=1)$$



- So, as desired, always get an error with high probability.

Low-level semantics

$$\llbracket c \rrbracket : (\text{Lay} \rightarrow \text{Mem}_{\xi, \perp}) \rightarrow \mathcal{H}(\text{Lay} \rightarrow \text{Mem}_{\xi, \perp})$$

Most cases are straightforward:

$$\llbracket c + c' \rrbracket(\zeta) = \llbracket c \rrbracket(\zeta) \cup \llbracket c' \rrbracket(\zeta)$$

$$\llbracket c; c' \rrbracket = \llbracket c' \rrbracket^\dagger \circ \llbracket c \rrbracket$$

$$\llbracket \text{skip} \rrbracket = \eta$$

$$\llbracket e := e' \rrbracket(\zeta) = \begin{cases} \eta(\lambda w. m[\llbracket e \rrbracket_{\zeta(w)}^w \mapsto \llbracket e' \rrbracket_{\zeta(w)}^w]) & (\text{if } \llbracket e \rrbracket_{\zeta(w)}^w \in \text{dom}(m)) \\ \eta(\lambda w. \xi) & (\text{otherwise}) \end{cases}$$

$$\llbracket \text{if } b \text{ then } c \text{ else } c' \rrbracket = \text{Cond}(\llbracket b \rrbracket, \llbracket c \rrbracket, \llbracket c' \rrbracket)$$

$$\llbracket \text{while } b \text{ do } c \rrbracket = \mu\theta. \text{Cond}(\llbracket b \rrbracket, \theta^\dagger \circ \llbracket c \rrbracket, \eta)$$

Example conditional semantics

$$\llbracket b \rrbracket \left(\begin{array}{l} w_1 \mapsto m_1 \\ w_2 \mapsto m_2 \\ w_3 \mapsto m_3 \end{array} \right) =$$

$$\left[\begin{array}{l} w_1 \mapsto \text{false} \\ w_2 \mapsto \text{true} \\ w_3 \mapsto \text{false} \end{array} \right]$$

$$\llbracket c \rrbracket \left(\begin{array}{l} w_1 \mapsto m_1 \\ w_2 \mapsto m_2 \\ w_3 \mapsto m_3 \end{array} \right) =$$

$$\left[\begin{array}{l} w_1 \mapsto m_{c1} \\ w_2 \mapsto m_{c2} \\ w_3 \mapsto m_{c3} \end{array} \right]$$

$$\left[\begin{array}{l} w_1 \mapsto m'_{c1} \\ w_2 \mapsto m'_{c2} \\ w_3 \mapsto m'_{c3} \end{array} \right]$$

$$\llbracket c' \rrbracket \left(\begin{array}{l} w_1 \mapsto m_1 \\ w_2 \mapsto m_2 \\ w_3 \mapsto m_3 \end{array} \right) =$$

$$\left[\begin{array}{l} w_1 \mapsto m'_{c'1} \\ w_2 \mapsto m_{c'2} \\ w_3 \mapsto m_{c'3} \end{array} \right]$$

$$\left[\begin{array}{l} w_1 \mapsto m'_{c'1} \\ w_2 \mapsto m'_{c'2} \\ w_3 \mapsto m'_{c'3} \end{array} \right]$$

$$\llbracket \text{if } b \text{ then } c \text{ else } c' \rrbracket \left(\begin{array}{l} w_1 \mapsto m_1 \\ w_2 \mapsto m_2 \\ w_3 \mapsto m_3 \end{array} \right) =$$

$$\left[\begin{array}{l} w_1 \mapsto m_{c'1} \\ w_2 \mapsto m_{c2} \\ w_3 \mapsto m_{c'3} \end{array} \right]$$

$$\left[\begin{array}{l} w_1 \mapsto m_{c'1} \\ w_2 \mapsto m'_{c2} \\ w_3 \mapsto m_{c'3} \end{array} \right]$$

$$\left[\begin{array}{l} w_1 \mapsto m'_{c'1} \\ w_2 \mapsto m_{c2} \\ w_3 \mapsto m'_{c'3} \end{array} \right]$$

$$\left[\begin{array}{l} w_1 \mapsto m'_{c'1} \\ w_2 \mapsto m'_{c2} \\ w_3 \mapsto m'_{c'3} \end{array} \right]$$

Refinement relation (low-level)

- **Input** We are interested in the distinguishing capacities of the adversary on inputs that directly correspond to stores. These are called **store projections** and are random variables ($\text{Lay} \rightarrow \text{Mem}_{\perp, \xi}$) of the form $w \mapsto w(s)$ for some store s (but also allow \perp and ξ).
- **Output** For $X, Y \in \mathcal{H}(\text{Lay} \rightarrow \text{Mem}_{\perp, \xi})$, $X \leq_L Y$ iff (roughly) the public parts of the random variables in X are less than those of Y *with probability greater than δ* . More precisely, for all $\zeta \in X$, there exists $\zeta' \in Y$ such that:
 - $\zeta_L \leq \zeta'_L$, or
 - $P(\zeta(w) \in \{\xi, \perp\}) \geq \delta$ and $P(\zeta'(w) = \xi) \geq \delta$.
- **Command refinement** c refines c' , written $c \leq_L c'$, when:

$$\forall s \in \text{stores}. \llbracket c \rrbracket(w \mapsto w(s)) \leq_L \llbracket c' \rrbracket(w \mapsto w(s))$$

Simulation relation

- We mostly consider public contexts containing **safe commands**.
- Safe commands transform store projections into **store projection sets**, which are (downsets of) sets of
 - store projections, or
 - random variables with high probability of error ($> \delta$).
- We project those semantics and define a simulation relation similar to the high-level one.

Theorem 2

For all safe low-level commands c, c' , we have:

$$\llbracket c \rrbracket \preceq \llbracket c' \rrbracket \iff \forall \text{ low-level public } C[\cdot]. \llbracket C[c] \rrbracket \leq_L \llbracket C[c'] \rrbracket$$

Simulation relation - more detail

- Begin with $R \subseteq \mathcal{H}(\text{private-stores}_{\perp, \xi})^2$, subject to conditions
- Then construct $R^+ \subseteq \mathcal{H}(\text{stores}_{\perp, \xi})^2$ from R , much as before
- Then construct $R^\times \subseteq \mathcal{H}(\text{Lay} \rightarrow \text{Mem}_{\xi, \perp})^2$ from R^+ by:

$$\begin{aligned} XR^\times Y &\equiv_{\text{def}} X \text{ and } Y \text{ are store projection sets} \\ &\quad \wedge \chi(X)R^+\chi(Y) \end{aligned}$$

(where χ extracts the stores, bottom, and, possibly, error from a store projection set).

- Then

$$[[c]] \preceq [[c']] \equiv_{\text{def}} \forall X, Y. XR^\times Y \Rightarrow [[c]](X)R^\times [[c']](Y)$$

Compilation

High-level commands c are compiled to low-level commands c^\downarrow by:

$$\begin{aligned} (!l_{\text{loc}})^\downarrow &= !l_{\text{nat}} \\ (l_{\text{loc}} := e)^\downarrow &= l_{\text{nat}} := e^\downarrow \end{aligned}$$

Theorem 3

For all high-level commands c, c' we have: $c \preceq c'$ iff $c^\downarrow \preceq c'^\downarrow$

Theorem 4

For all high-level commands c, c' ,

$$\forall \text{ high-level public } C[\cdot]. \llbracket C[c] \rrbracket \leq_L \llbracket C[c'] \rrbracket$$

holds iff

$$\forall \text{ low-level public } C[\cdot]. \llbracket C[c^\downarrow] \rrbracket \leq_L \llbracket C[c'^\downarrow] \rrbracket$$

does

Naturality of the semantics

Parameterising over sets of layouts:

$$\begin{array}{ccc} \text{Mem}_{\xi \perp}^{W'} & \xrightarrow{\llbracket c \rrbracket_{W'}} & \mathcal{H}(\text{Mem}_{\xi \perp}^{W'}) \\ \downarrow & & \downarrow \\ \text{Mem}_{\xi \perp}^{\ell} & & \mathcal{H}(\text{Mem}_{\xi \perp}^{\ell}) \\ \downarrow & & \downarrow \\ \text{Mem}_{\xi \perp}^{W''} & \xrightarrow{\llbracket c \rrbracket_{W''}} & \mathcal{H}(\text{Mem}_{\xi \perp}^{W''}) \end{array}$$

Suggests working with presheaves.

Summary of this work

- A semantic approach to layout randomization

$$\begin{aligned} \llbracket c \rrbracket \preceq \llbracket c' \rrbracket &\iff \forall \text{ public } C[\cdot]. \llbracket C[c] \rrbracket \leq_L \llbracket C[c'] \rrbracket \\ &\Updownarrow \\ \llbracket c^\downarrow \rrbracket \preceq \llbracket c'^\downarrow \rrbracket &\iff \forall \text{ public } C[\cdot]. \llbracket C[c^\downarrow] \rrbracket \leq_L \llbracket C[c'^\downarrow] \rrbracket \end{aligned}$$

- Nondeterminism (and its interaction with probabilities).
- Concurrency would be an interesting next step (or higher-order)

Summary of this work

- A semantic approach to layout randomization

$$\begin{aligned} \llbracket c \rrbracket \preceq \llbracket c' \rrbracket &\iff \forall \text{ public } C[.]. \llbracket C[c] \rrbracket \leq_L \llbracket C[c'] \rrbracket \\ &\iff \\ \llbracket c^\downarrow \rrbracket \preceq \llbracket c'^\downarrow \rrbracket &\iff \forall \text{ public } C[.]. \llbracket C[c^\downarrow] \rrbracket \leq_L \llbracket C[c'^\downarrow] \rrbracket \end{aligned}$$

- Nondeterminism (and its interaction with probabilities).
- Concurrency would be an interesting next step (or higher-order)
- But there is still something to understand.

What is the corresponding operational semantics?

- To make sense of our semantics, we want operational/denotational consistency, akin to:

$$\zeta' \in \llbracket c \rrbracket(\zeta) \iff \forall w. w \models \langle c, \zeta(w) \rangle \rightarrow^* \zeta'(w)$$

- Once again, the nondeterminism is chosen after the layout.
- We set the nondeterminism with an oracle, then the operational semantics is deterministic:

$$\zeta' \in \llbracket c \rrbracket(\zeta) \iff \exists \pi. \forall w. w \models \langle c, \zeta(w), \pi \rangle \rightarrow^* \zeta'(w)$$