Approximate Simulations for Task-Structured Probabilistic I/O Automata

Sayan Mitra and Nancy Lynch

CSAIL, MIT

Implementation

- Implementation: a fundamental notion in concurrency theory
- "traces" or observable behavior, e.g. sequence of events, timing of events, probabilities
- **A** *implements* ***B*** if *traces(**A**)* $\subseteq$ *traces(**B**)*
- **A** *is equivalent to* ***B*** if they implement each other, i.e., *traces(**A**)* = *traces(**B**)*

Fragility of equality

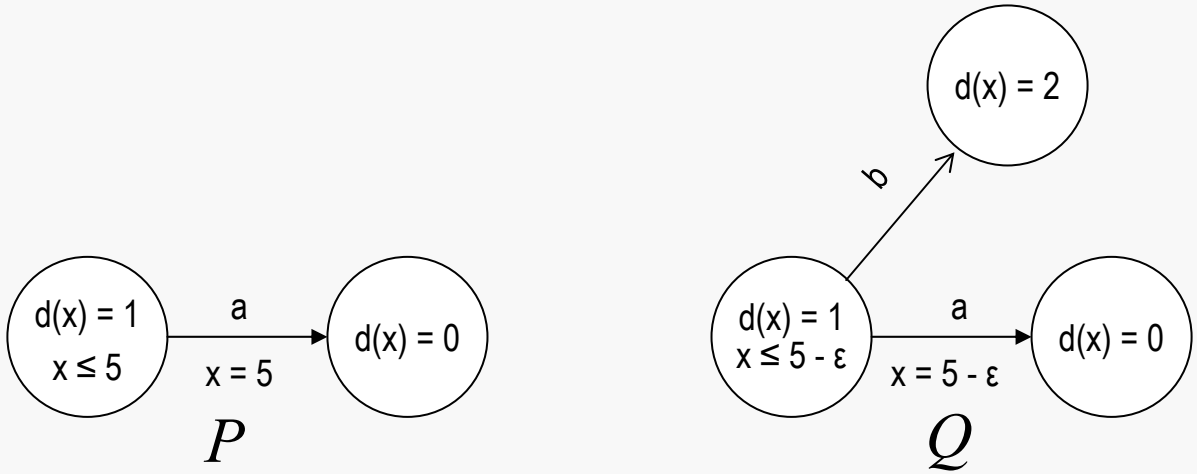

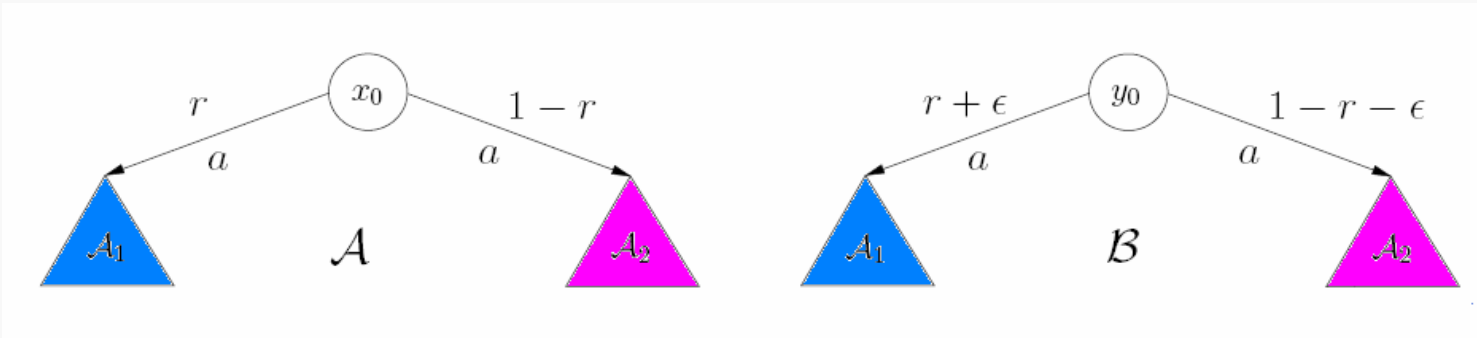

Unequal, but similar

- A metric **d** on the space **T** of traces of **A** (and **B**)
- (**T,d**) is a metric space
- **A** *approximately implements* ***B*** if the one-sided Hausdorff distance from traces(**A**) to traces(**B**) is small.

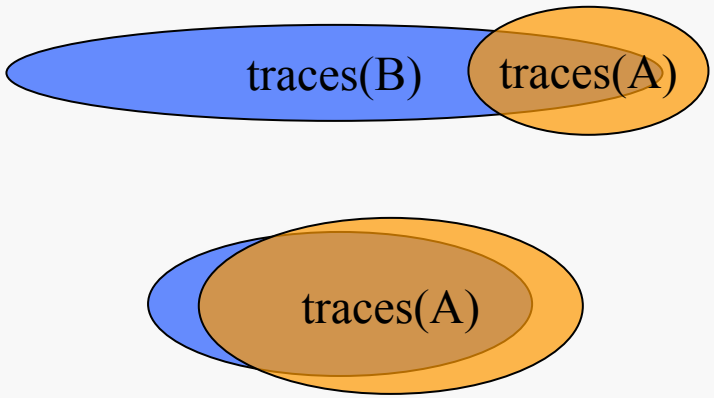

Previously

- Metric-based approximate simulations and bisimultions
  - PIOA **[Jou and Smolka '90]**
  - Labelled Markov Processes **[Desharnais, et. al. '04] [Breugel, Mislove '03]**
  - Hybrid Systems **[Girard, Julius, Pappas '05]**
  - GSMP **[Gupta, Jagadeesan, Panagaden '04]**
  - Linear stochastic hybrid automata **[Julius '06]**

- This talk: approximate implementations & simulations for task-PIOAs
  - Probabilistic and nondeterministic choices
  - Simulation relation on distributions of executions (not states)
  - Metric on trace distributions

    state space / space of external actions need not be metric spaces.

Outline

- [x] Background
- [ ] Task PIOA vocabulary
- [ ] Definitions: metrics and simulations
- [ ] Soundness (sketch)
- [ ] Discussions
  - Generalization
  - Applications
  - Future directions

Task PIOA

$A = (Q,v,A,D,R)$ **[Segala'96] [Canetti, et. al. 2006]**

- Countable set of states Q
- Initial distribution on states v
- Countable set of actions $A = I \cup O \cup H$
  - If $I = \emptyset$ then A is *closed*
  - $O \cup H$ set of *locally controlled* actions
- Set of (q,a,μ) transitions D
- An equivalence R relation on locally controlled actions
  - Each equivalence class of R is a *task*

- Input enabled
- For any action there is at most one transition
- For any task T, there is at most one enabled action in T

Nondeterministic choice over tasks.

Task PIOA Vocabulary

- *Execution fragment* $\alpha = q_0 a_1 q_1 a_2 \ldots$
- α is an *execution* if q_0 in supp(v)
- *trace*(α) is obtained by deleting all q's and the a's in H.
  - *trace* is a measurable function

- *Task scheduler* σ is a sequence of tasks $T_1 T_2 T_3 \ldots$
- **apply(**μ,σ) = probability distribution over fragments

- *tdist*(μ) is the image measure of μ under the *trace* function
- *tdists*(A) = {*tdist*(*apply*(v, σ)): σ is a task scheduler for A}

Previously in PIOA: Exact implementations

In **[Canetti, et. al. 2006]** the *tdists(A) ⊆ tdists(B) exact* implementation relation used for verifying an Oblivious Transfer protocol.

- $R \subseteq Disc(Execs^*(A)) \times Disc(Execs^*(B))$ is a simulation relation if:
  - $\mu_1 R \mu_2$ implies $\text{tdist}(\mu_1) = \text{tdist}(\mu_2)$
  - $v_1 R v_2$
  - If $\mu_1 R \mu_2$, then $\mu_1' \mathcal{E}(R) \mu_2'$.

$\mathcal{E}(R)$ is defined using lifting and flattening

Approximate implementations

- Uniform metric on trace distributions

$$\mathbf{d}_u(\mu_1, \mu_2) = \sup_{C \in F_{Traces}} |\mu_1(C) - \mu_2(C)|$$

- A δ-*implements* B if for every μ_1 there is a μ_2 with $\mathbf{d}_u(\mu_1,\mu_2) \leq \delta$

$$\phi(Disc(Exec^*(A)) \times Disc(Exec^*(B))) \rightarrow \Re_{\geq 0} \cup \{\infty\}$$

is (ε,δ)-approximate simulation function if:

$$\text{Start} : \phi(\nu_1, \nu_2) \leq \varepsilon$$

$$\text{Step} : \phi(\mu_1, \mu_2) \leq \varepsilon \text{ implies } \hat{\phi}(\mu_1', \mu_2') \leq \varepsilon$$

$$\text{Trace} : \phi(\mu_1, \mu_2) \leq \varepsilon \text{ implies } \mathbf{d_u}(\text{tdist}(\mu_1), \text{tdist}(\mu_2)) \leq \delta$$

Phi and Phi Hat

Given $\phi(X,Y) \to \Re_{\geq 0} \cup \{\infty\}$

$$\widehat{\phi}(x_1,y_1) = \min_{\substack{\Gamma \in D(X \times Y) \\ x_1 = \mathbf{E}[\Gamma_x] \\ y_1 = \mathbf{E}[\Gamma_y]}} \left[\max_{(x,y) \in \mathrm{supp}(\Gamma)} \phi(x,y) \right]$$

$\hat{\phi}(x_1,y_1) \leq \varepsilon \Leftrightarrow \exists$ a witnessing joint distribution $\psi \in D(X \times Y)$

such that $\max_{x,y \in \mathrm{supp}(\psi)} \phi(x,y) \leq \varepsilon$

$x_1 = \sum_{x,y} \psi(x,y)x$ and $y_1 = \sum_{x,y} \psi(x,y)y$

Expansion

$\phi(x_1, y_1) \leq \varepsilon$

One witnessing joint distribution is the delta at x_1, y_1

$\hat{\phi}(x_1, y_1) \leq \varepsilon$

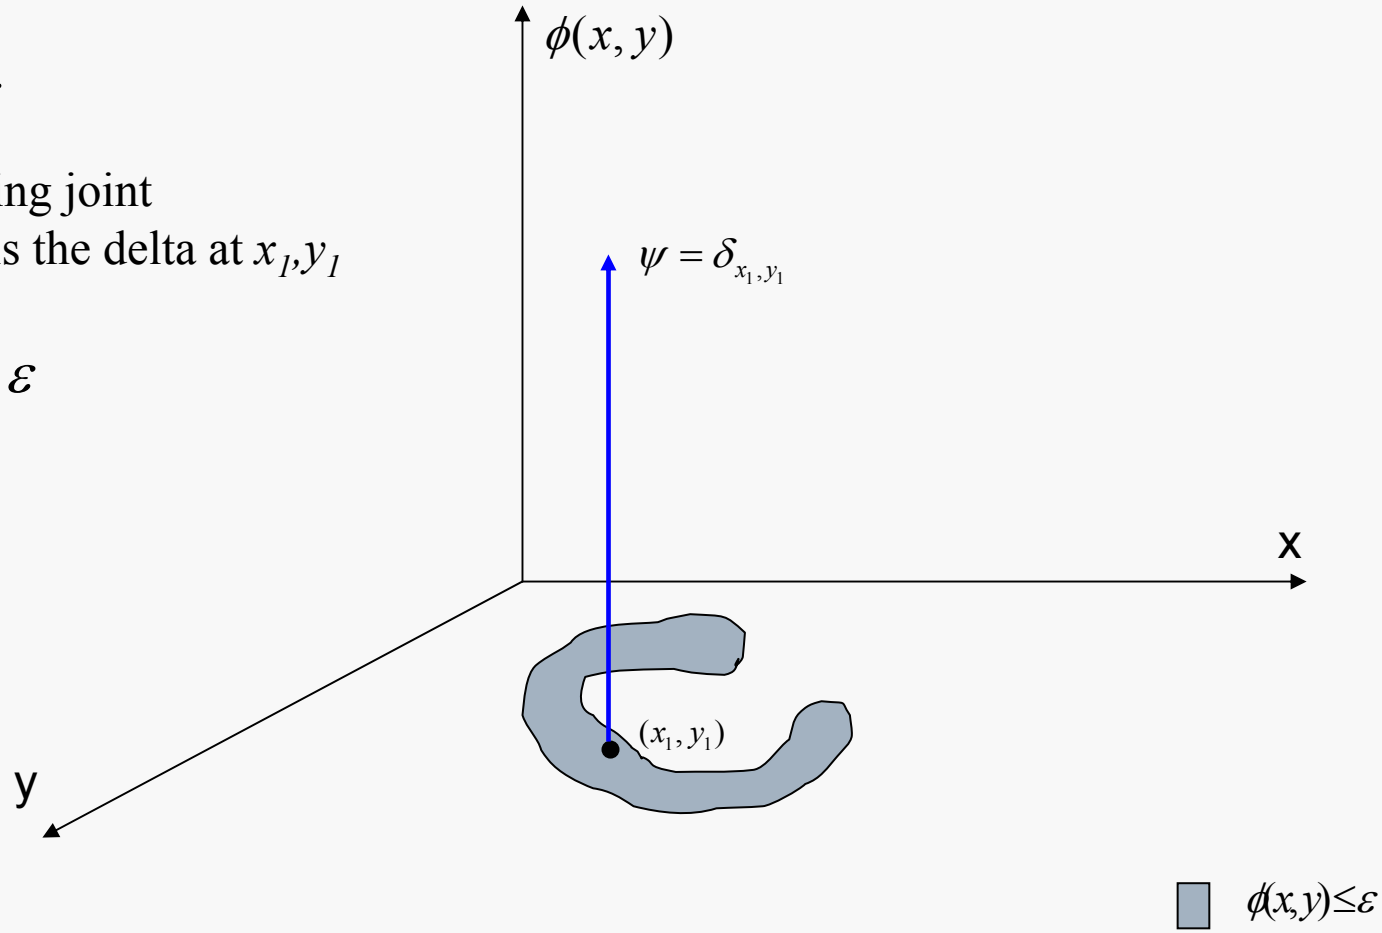

Expansion

$\phi(x,y) > \varepsilon$

$\hat{\phi}(x,y) \leq \varepsilon$

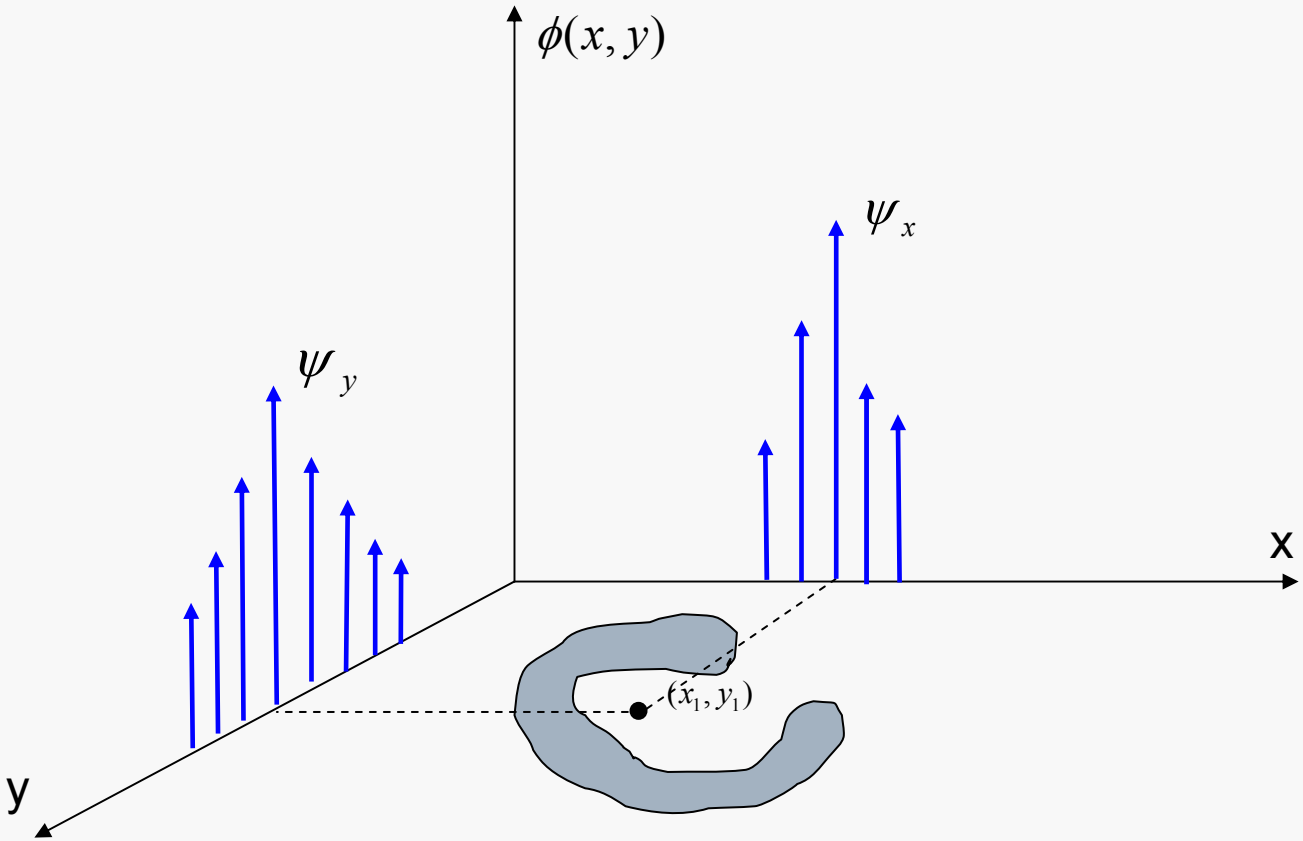

$\phi(x,y) \leq \varepsilon$

Approximate simulation

ϕ is an (ε,δ)-*approximate simulation* function from A to B if:

1. $\phi(\nu_1, \nu_2) \leq \varepsilon$

2. There exists a function $c: R_1^* \times R_1 \rightarrow R_2^*$ such that for any task T of A and any schedule σ of A if μ_1 is *consistent* σ and μ_2 is *consistent* with *full(c)*(σ) then $\phi(\mu_1, \mu_2) \leq \varepsilon$ implies $\hat{\phi}(\text{apply}(\mu_1, T), \text{apply}(\mu_2, c(\sigma, T))) \leq \varepsilon$

3. $\phi(\mu_1, \mu_2) \leq \varepsilon$ implies $\mathbf{d_u}(\text{tdist}(\mu_1), \text{tdist}(\mu_2)) \leq \delta$

Lemma 1: ϕ (ε,0)-*approximate simulation* then R := $\{(\mu_1, \mu_2) \mid \phi(\mu 1, \mu 2) \leq \varepsilon\}$ is an exact simulation relation.

Key Lemmas

- **Lemma 2:** $\widehat{\phi}(\mu_1, \mu_2) \leq \varepsilon$ with witness ψ.

  To prove that $\widehat{\phi}(f_1(\mu_1), f_2(\mu_2)) \leq \varepsilon$ it suffices to show that

  $\forall \rho_1, \rho_2 \in \text{supp}(\psi), \widehat{\phi}(f_1(\rho_1), f_2(\rho_2)) \leq \varepsilon.$

Key Lemmas

- **Lemma 3:** $\bar{\phi}(\mu_1, \mu_2) \leq \varepsilon$ implies $\mathbf{d_u}(\text{tdist}(\mu_1), \text{tdist}(\mu_2)) \leq \delta$
- **Lemma 4:** If $\mu = \text{Lt}_{n\to\infty}\ \mu_n$ then $\text{tdist}(\mu) = \text{Lt}_{n\to\infty}\ \text{tdist}(\mu_n)$.
- **Lemma 5:** If $\mu_{1i} \to \mu_1$ and $\mu_{2i} \to \mu_2$ then $\mathbf{d}_u(\mu_{1i},\mu_{2i}) \to \mathbf{d}_u(\mu_1,\mu_2)$.

Soundness

- **Theorem:** If there exists an (ε, δ)-*approximate simulation* function from *A* to *B*, then *A δ-implements B*.
- Construct chain of distributions for *A* applying one task at a time. Construct the corresponding chain for B.
- Induction on the length of the chain
  - Base case from start condition
  - Induction step from Lemma 2
- Use Lemmas 4 & 5 for n→∞

Probabilistic Safety

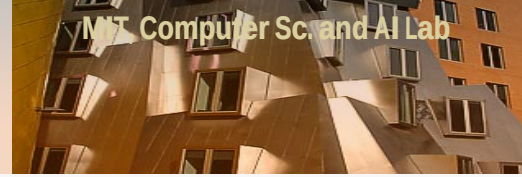

- X be a random variable on $(\mathbf{T}, F_{\mathbf{T}})$.
- If A_1 is δ-equivalent to B and for every trace distribution μ_2 of B, $\mu_2[X=x] = p$ then $\mu_1[X=x] \leq p + \delta$.

- X_u: $\mathbf{T} \rightarrow \{0,1\}$ defined as $X_u(\beta) := 1$ if some unsafe action U occurs in β.
- If B is safe with probability p then A is safe with probability at least $p + \delta$.

Task-PIOAs

- An *environment* E for A is a task-PIOA such that $E||A$ is closed
- *External behavior* of A is a function mapping each environment E of A to the set of trace distributions of $E||A$
- A_1 *δ-implements* B if for every environment E, for every trace distribution μ_1 in $extbeh_A(E)$ there is a trace distribution μ_2 in $extbeh_A(E)$.
- Suppose for every environment E, there exists a (ε_E, δ)-*approximate simulation* function from $A_1||E$ to $B||E$, then A_1 *δ-implements* B.

Future directions

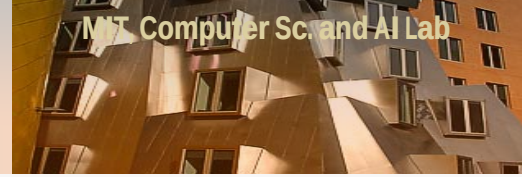

- Applications: randomized consensus protocols,
- Approximate implementations and simulation relations for task-PIOAs with continuous state spaces.
- Simulations as functions of distributions over states (as opposed to distributions over fragments).
- Explore the possibility of automating simulation proofs by solving optimization problems.

  (See my thesis)

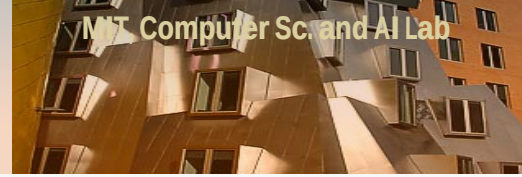

Thank you!

Key Lemmas

Lemma 2: $\widehat{\phi}(\mu_1, \mu_2) \leq \varepsilon$ with witness ψ.

$f_i : \mathrm{Disc}(X_i) \to \mathrm{Disc}(X_i)$ are distributive functions.

If $\forall \rho_1, \rho_2 \in \mathrm{supp}(\psi), \widehat{\phi}(f_1(\rho_1), f_2(\rho_2)) \leq \varepsilon$ then $\widehat{\phi}(f_1(\mu_1), f_2(\mu_2)) \leq \varepsilon$.

For each $\rho_1, \rho_2 \in \mathrm{supp}(\psi)$, let $\psi_{\rho_1,\rho}$ be the witnessing joint for $\widehat{\phi}(f_1(\rho_1), f_2(\rho_2)) \leq \varepsilon$.

Define $\psi' := \sum_{\rho_1,\rho_2 \in \mathrm{supp}(\psi)} \psi(\rho_1, \rho_2)\psi_{\rho_1,\rho_2}$

Show : $f_i(\mu_i) = \sum_{\eta_1,\eta_2} \psi'(\eta_1, \eta_2)\eta_i$ and $\eta_1, \eta_2 \in \mathrm{supp}(\psi')$ implies $\phi(\eta_1, \eta_2) \leq \varepsilon$