Modelling Stochastic Nondeterministic Systems:

The challenge of continuous measures

Roberto Segala

University of Verona

Why Stochastic Systems?

- Distributed algorithms
  - Consensus

    … the scheduler is adversarial
  - Leader election
- Security
  - Randomized protocols

    … but users are unpredictable
  - Indistinguishability
- Embedded systems
  - Environment may be stochastic

    … but distributions may be unknown
  - Perturbations may be stochastic

Randomization Difficult

- Intuition often fails
  - Many wrong protocols
- Interplay probability / nondeterminism
  - Independence broken
- Probability gives observational power
  - Language inclusion is branching
- Measurability
  - When can we study probabilities
- Compositional reasoning
  - Need substitutive relations
  - Need projection theorems

Models with Discrete Measures

- Markov Processes
- Markov Decision Processes [Bel57]
  - Probabilistic Automata [Rab63]
  - Reactive model [GSST90]
- Strictly Alternating Automata [Han91]
- Alternating Automata [Var85,PLS02]
- Probabilistic Automata [Seg95]
  - Extend all models above
- Probabilistic Nondeterministic Systems [BA95]
- Concurrent Probabilistic Systems [BK98]
- Probabilistic Reactive Modules [AHJ01]

Probabilistic Automata

$$A = (Q\,,\, q_0,\, E\,,\, H\,,\, D)$$

Transition relation
$D \subseteq Q \times (E \cup H) \times \mathrm{Disc}(Q)$

Internal (hidden) actions

External actions: $E \cap H = \varnothing$

Initial state: $q_0 \in Q$

States

Example: Probabilistic Automata

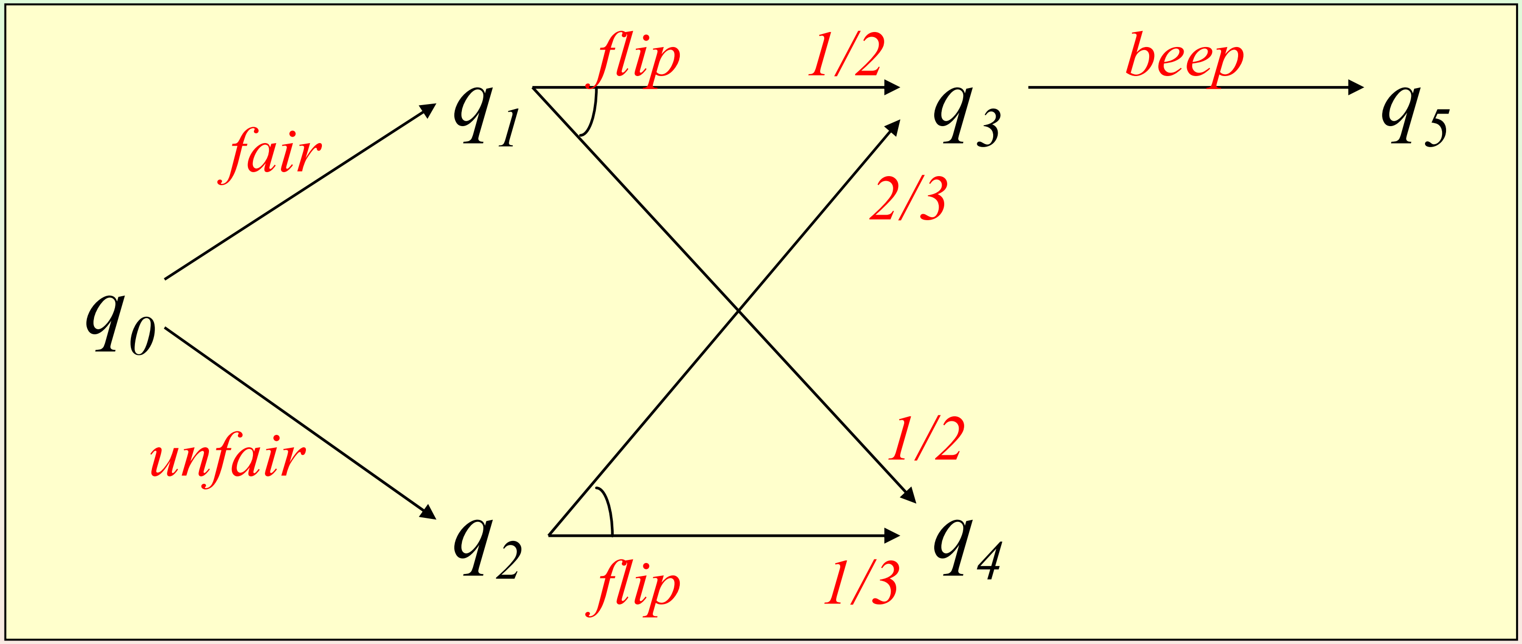

What is the probability of beeping?

Example: Probabilistic Execution

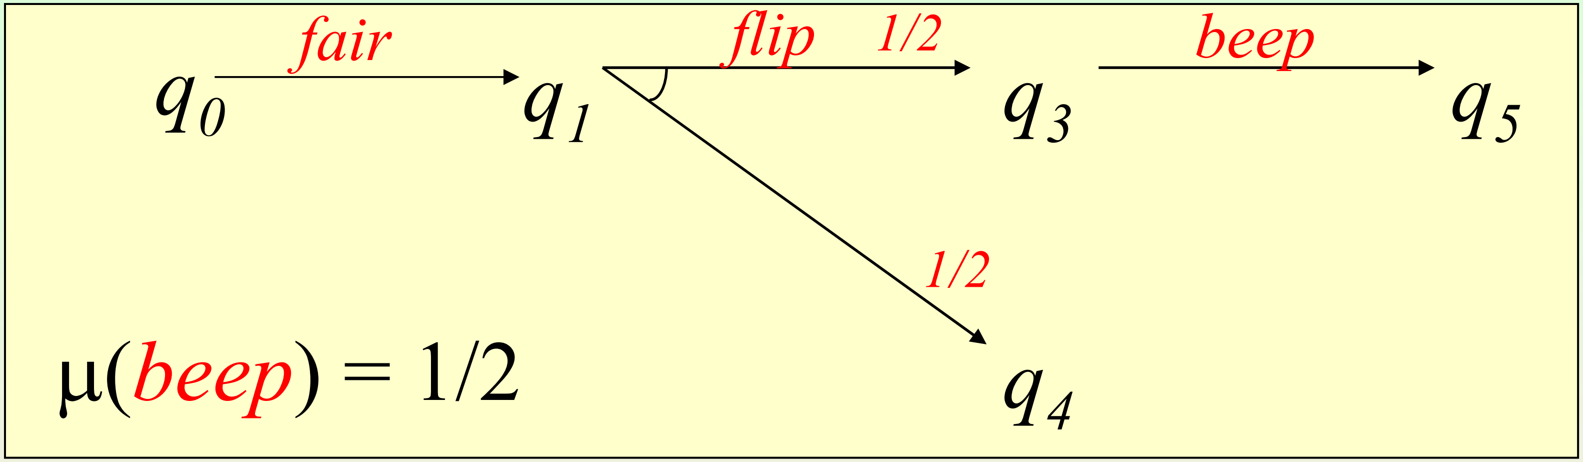

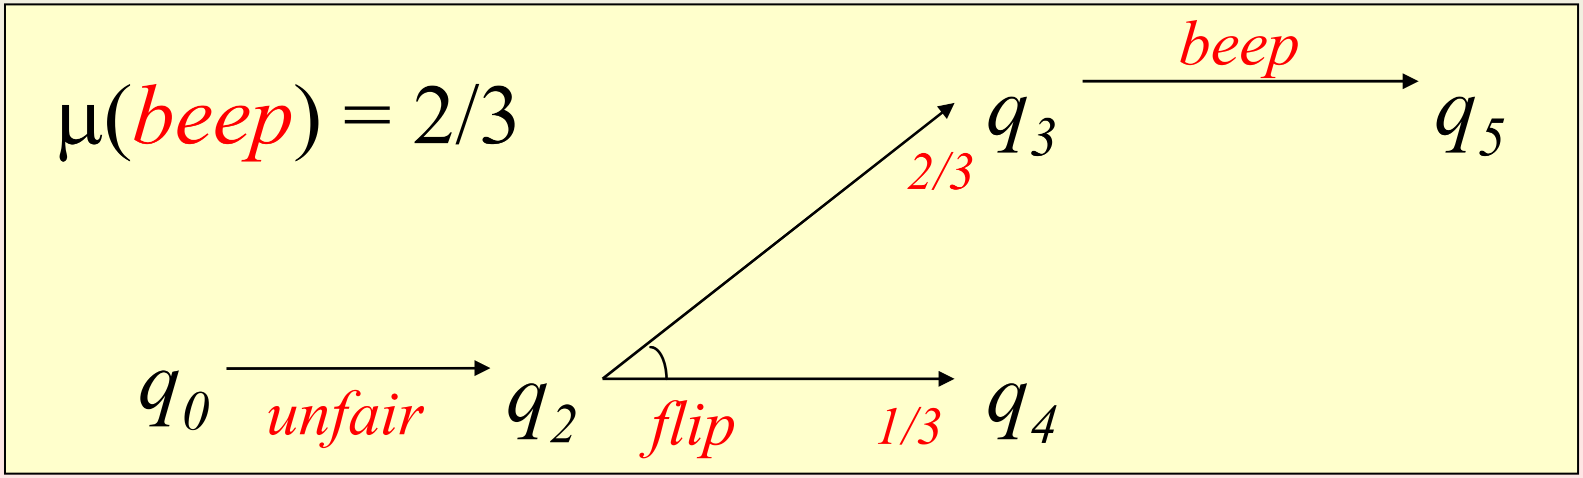

Cones and Measures

- Cone of α
  - Set of executions with prefix α
  - Represent event "α occurs"
- Measure of a cone
  - Product edges of α

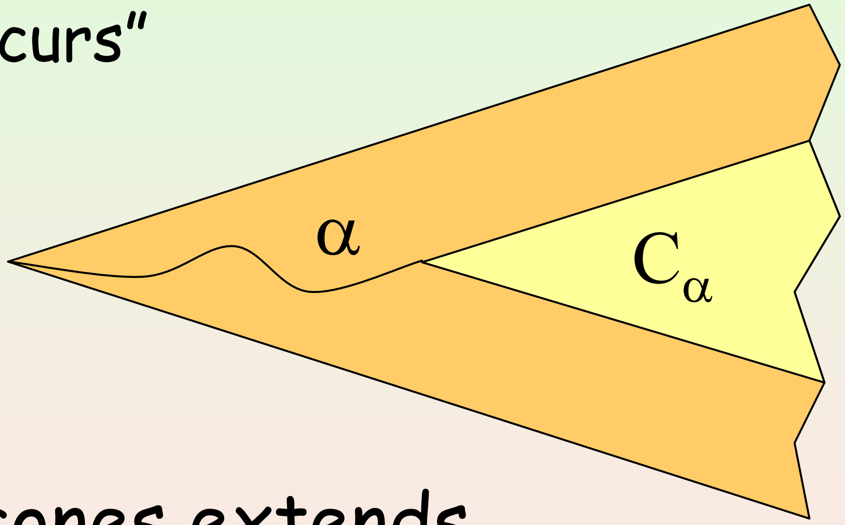

Thm. The measure over cones extends uniquely to a measure over the σ-field generated by cones

Schedulers and Probabilistic Executions

Scheduler σ

$$\sigma : exec^*(A) \to SubDisc(D)$$

$\sigma(\alpha)((q,a,\mu)) > 0$ implies $q = lstate(\alpha)$

Probabilistic execution:

given start state r, measure $\mu_{\sigma,r}$ where

$$\mu_{\sigma,r}(C_r) = 1$$
$$\mu_{\sigma,r}(C_{\alpha a q}) = \mu_{\sigma,r}(C_\alpha)\, \rho$$

$$\rho = \sum_{(s,a,\nu) \in D} \sigma(\alpha)((s,a,\nu))\, \nu(q)$$

Examples of Events

- Eventually action a occurs
  - Union of cones where action a occurs once
- Action a occurs at least n times
  - Union of cones where action a occurs n times
- Action a occurs at most n times
  - Complement of action a occurs at least n+1 times
- Action a occurs exactly n times
  - Intersection of previous two events
- Action a occurs infinitely many times
  - Intersection of action a occurs at least n times for all n
- Execution α occurs and nothing is scheduled after
  - Set consisting of α only
  - C_α intersected complement of cones that extend α

Composition

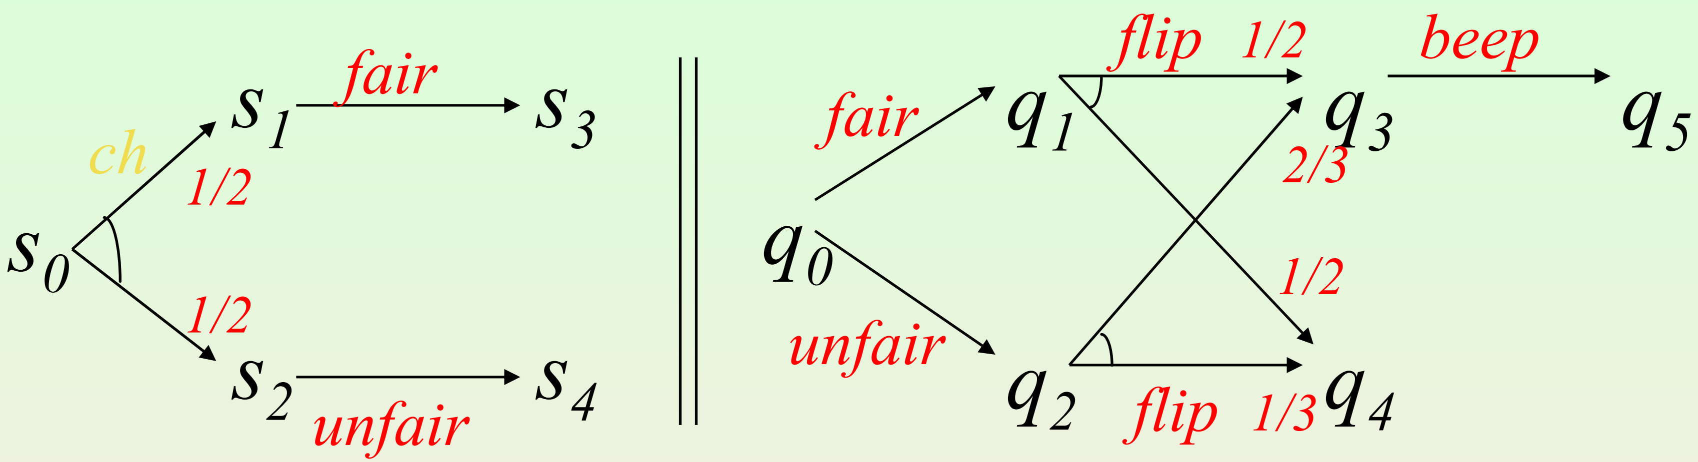

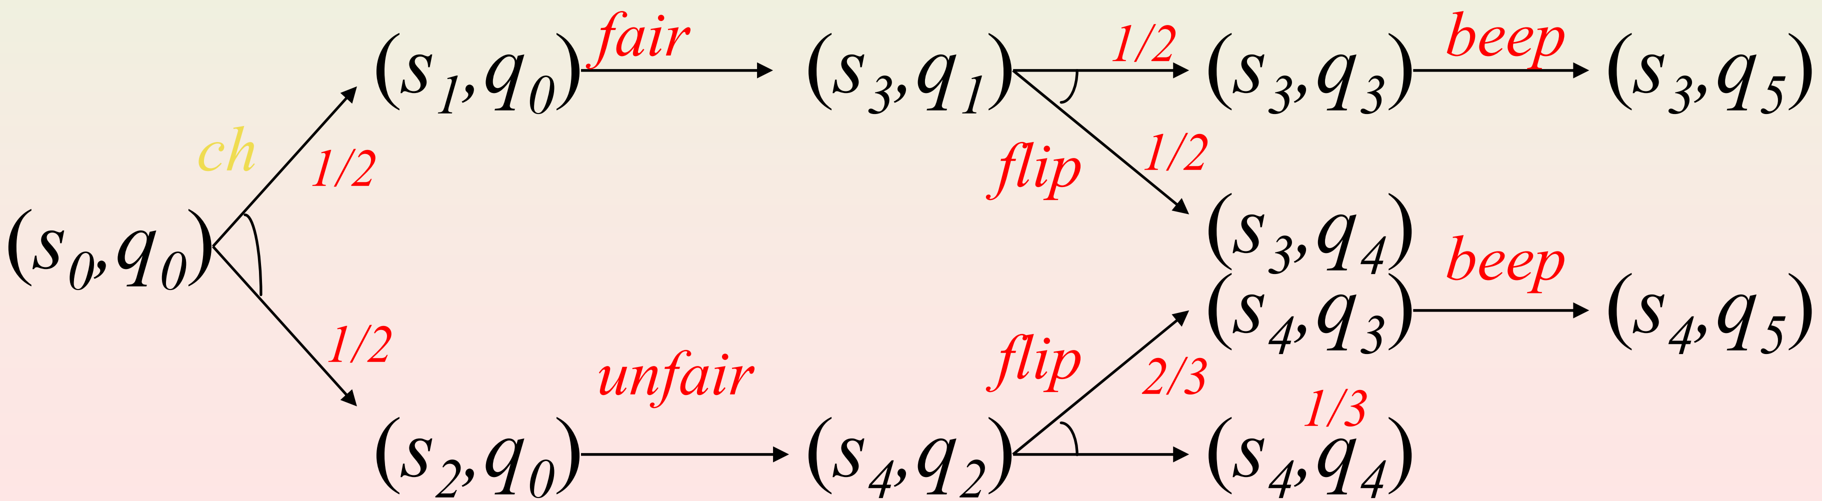

Projections

The projection function is measurable

$\pi(\mu)$: image measure under π of μ

Theorem

If μ is a probabilistic execution of $A_1 \| A_2$

then

$\pi_i(\mu)$ is a probabilistic execution of A_i

Example: Projection

Projection onto right component

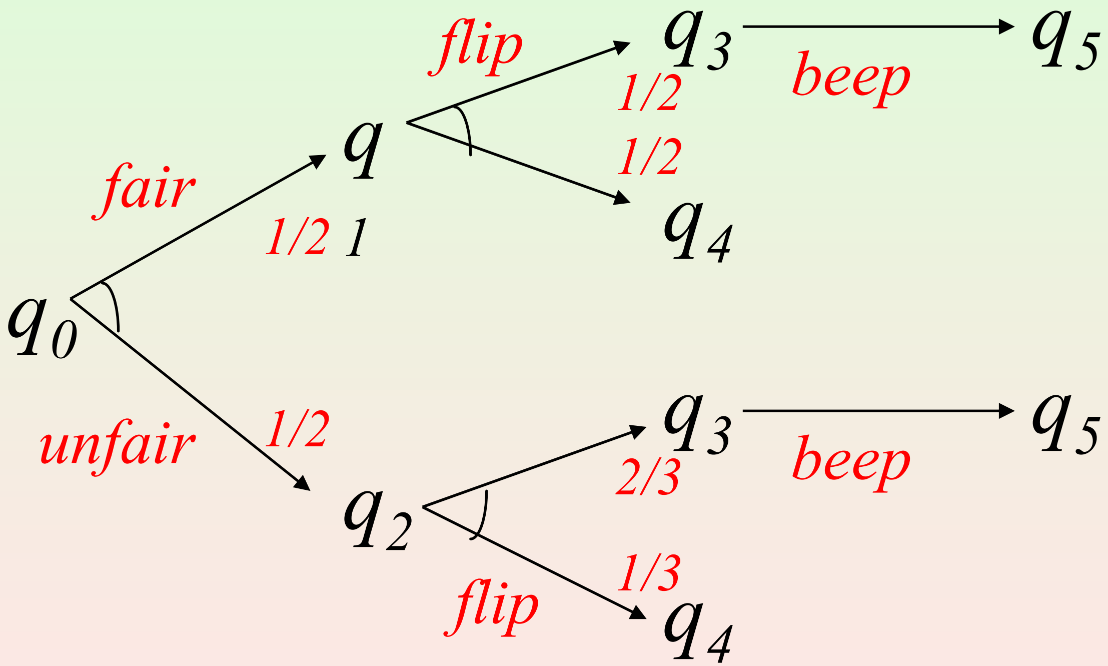

Note that the scheduler is randomized

Use of Projections

Let μ start in a state s of F_0.

μ

$\mu(\pi_2^{-1}(C1 \cap C2)) \geq p$

projection ↙ ↖ inverse image

$\mu \lceil MP$

$\pi_1(\pi_2^{-1}(C1 \cap C2))$ sat. $M1 \cap M2$

$\mu \lceil CF$

$\pi_2(\mu)(C1 \cap C2) \geq p$

$$F_0 \xrightarrow{2}_{p} D$$

Hyerarchical Verification

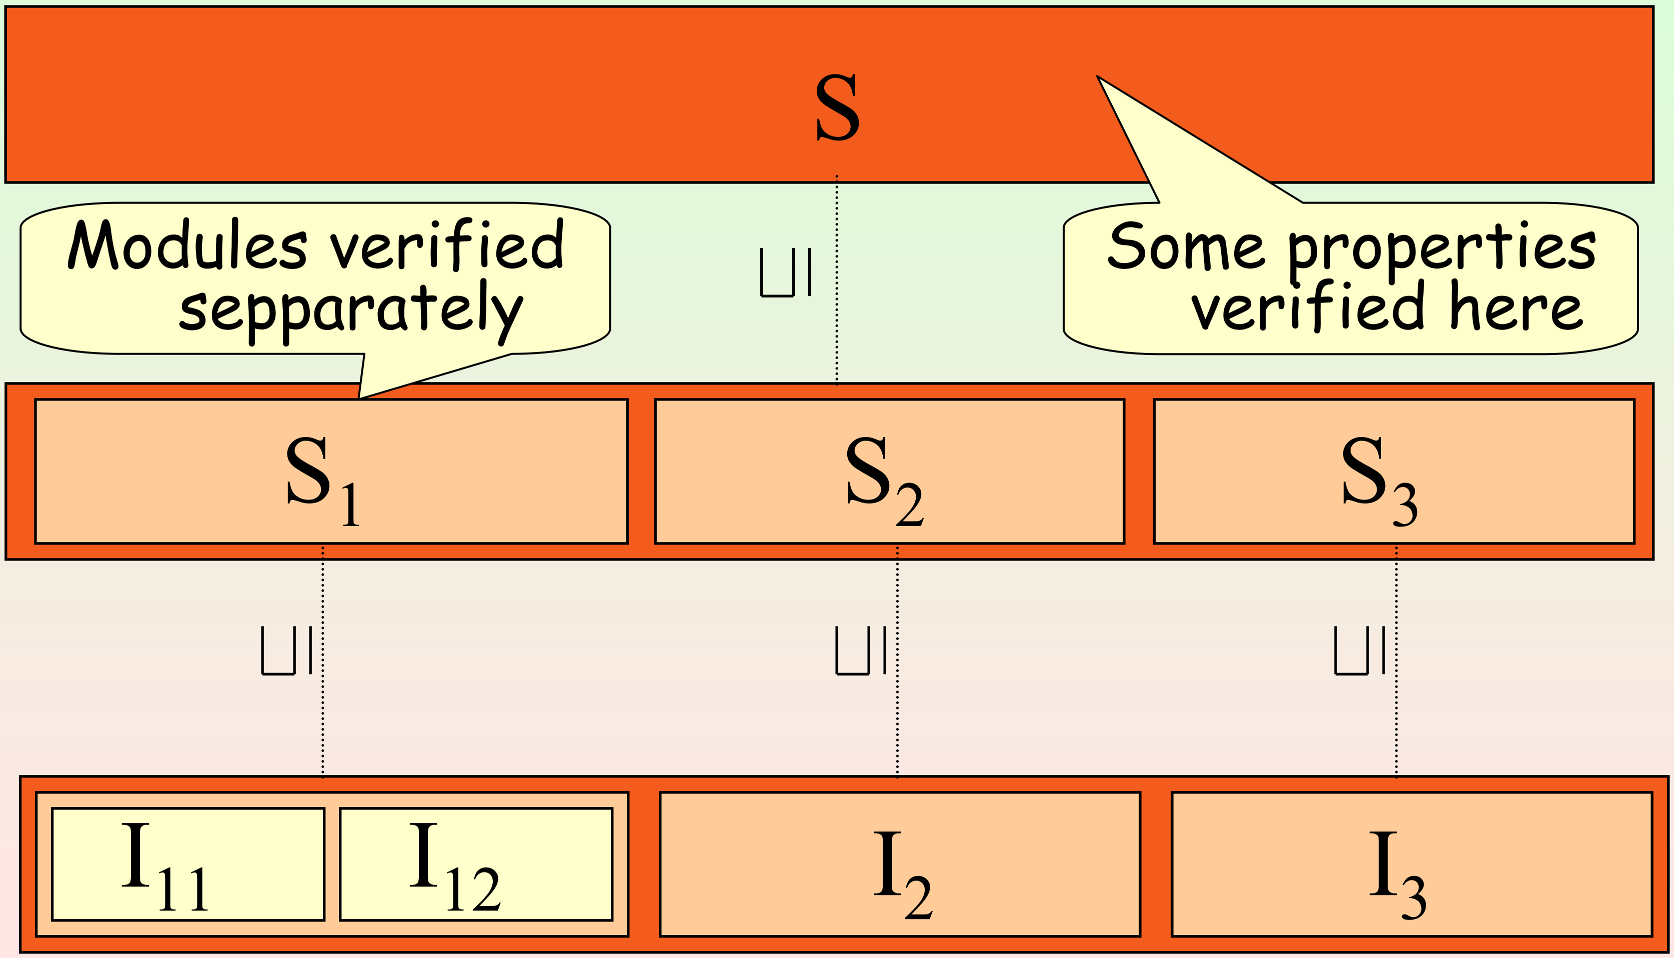

Trace Distributions

The *trace* function is measurable

Trace distribution of μ

tdist(μ) : image measure under *trace* of μ

Trace distribution inclusion preorder

$A_1 \leq_{TD} A_2$ iff $tdists(A_1) \subseteq tdists(A_2)$

Trace Distribution Inclusion is not Compositional

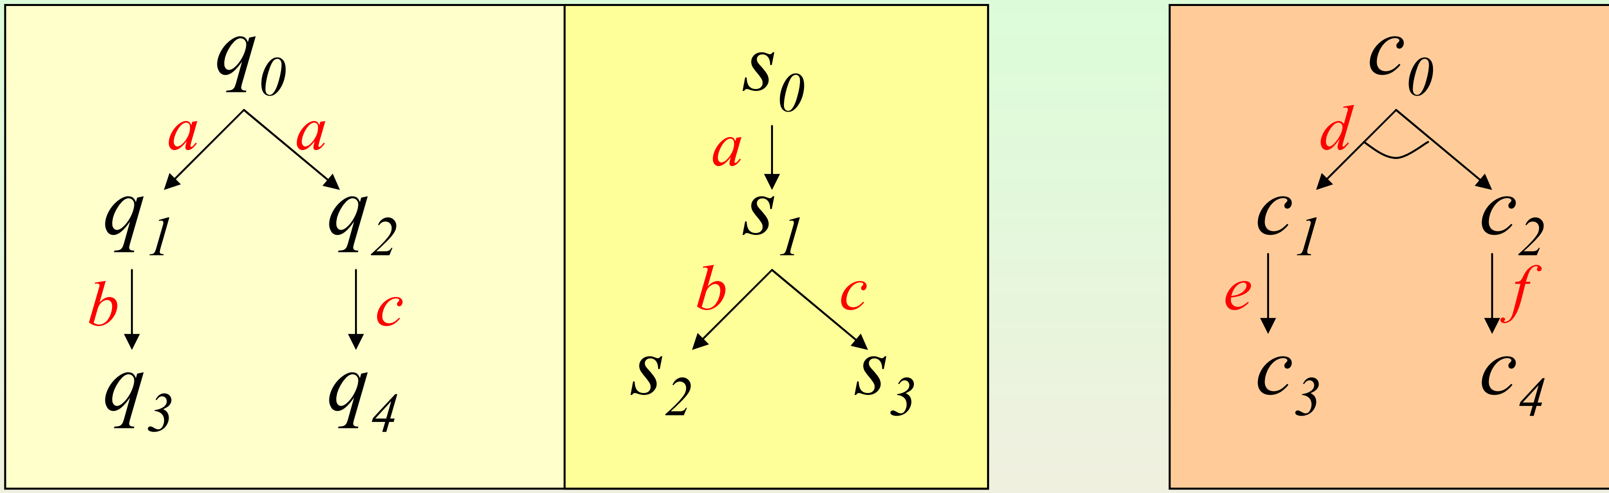

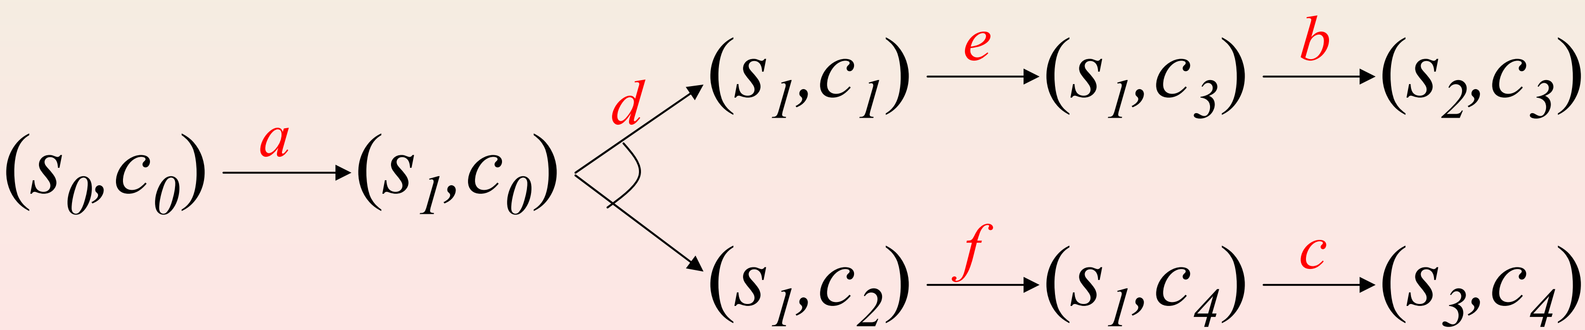

How to Get Compositionality

- Restrict the power of composition
  - Probabilistic reactive modules [AHJ01]
  - Switched probabilistic I/O automata [CLSV04]

- Trace Distribution Precongruence
  - Coarsest precongruence included in preorder
  - Alternative characterizations
    - Principal context [Seg95]
    - Testing [Seg96]
    - Forward simulations [LSV03]

Models with Continuous Measures

- Continuous Time Markov Chains
  - Markov chains with exponential time delays
- Generalized Stochastic Petri Nets [MBC84]
  - Petri Nets extended with exponential delays
- Stochastic Transition Systems [Alf98]
  - GSPNs with nondeterminism
- Labelled Markov Processes [DP97]
  - Markov chains with labels and continuous measures
- Concurrent Timed Probabilistic Graphs [KNNS02]
  - Timed automata with arbitrary measures
- Continuous Markov Decision Processes [BHHK04]
- Stochastic Transition Systems [CSKN05]
  - LMP's with nondeterminism

Problems with Measurability

- Not all sets can be measurable
  - Need to fix a σ-field
- Not all schedulers are nice

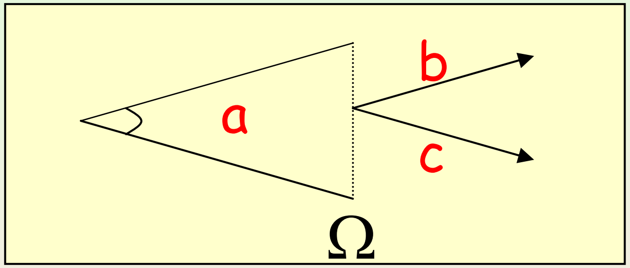

  - Let X be a non-measurable subset of Ω
  - Schedule b from X and c from Ω-X
  - **Solution**: restrict to measurable schedulers [CSKN05]
- Not all automata may be nice
  - Enable b only from X and c only from Ω-X
  - **Solution**: impose measurable transition relations [AW05]

Stochastic Transition Systems

$$A = ((Q,F_Q)\ ,\ q_0,\ (L,F_L)\ ,\ D)$$

Transition relation

$D \subseteq Q \times (E \cup H) \times \mathrm{D}(Q,F_Q)$

Actions with σ-field

Initial state: $q_0 \in Q$

States with σ-field

What Next?

- More general model
  - Stochastic hybrid systems
    - Nondeternimism, continuous evolutions
    - Stochastic differential equations
- Understand relations
  - Substitutivity
  - Aproximate simulation relations (metrics)
    - Exact values do not seem to matter
    - Useful for cryptographyc protocols
- Understand logics
  - CSL, PCTL
- Verification
  - Aproximate reasoning
  - Model checking
- Case studies
  - ???