Restricting schedulers

Restricting the nondeterministic choices arising from the interleaving

S. Giro

Sección de Ciencias de la Computación
Facultad de Astronomía, Matemática y Física
Universidad Nacional de Córdoba

VOSS 2007

Markov decision processes (MDPs) have both non-deterministic and probabilistic choices

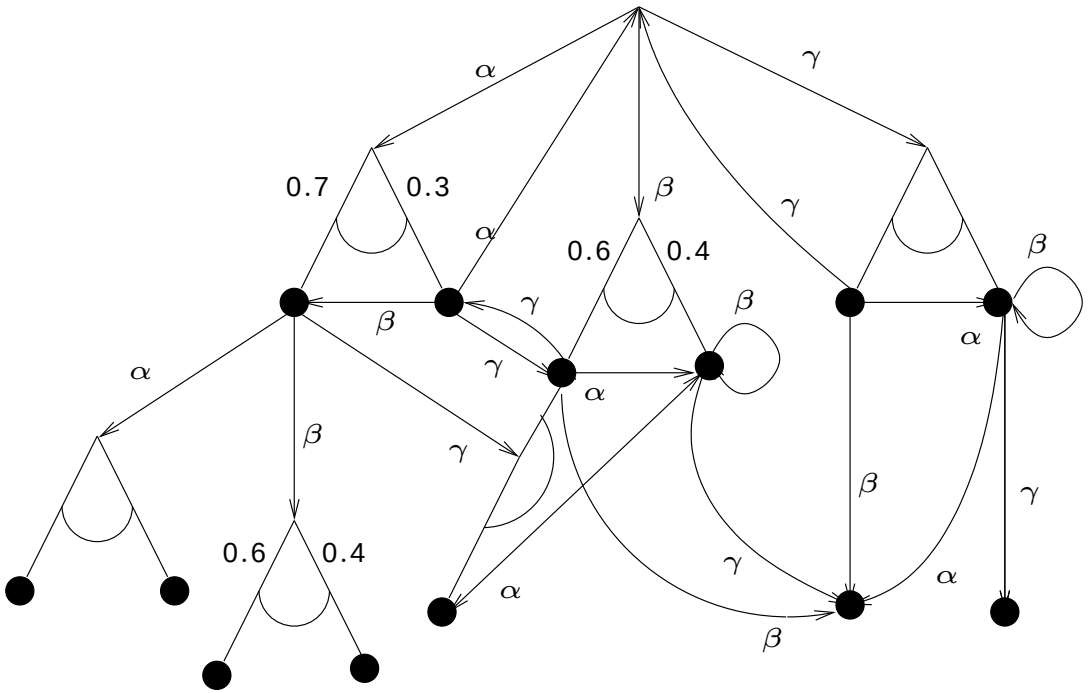

- In our framework, a choice between the different non-deterministic options is performed, and this choice determines a distribution for the next state of the system

Markov decision process need schedulers to yield probabilities

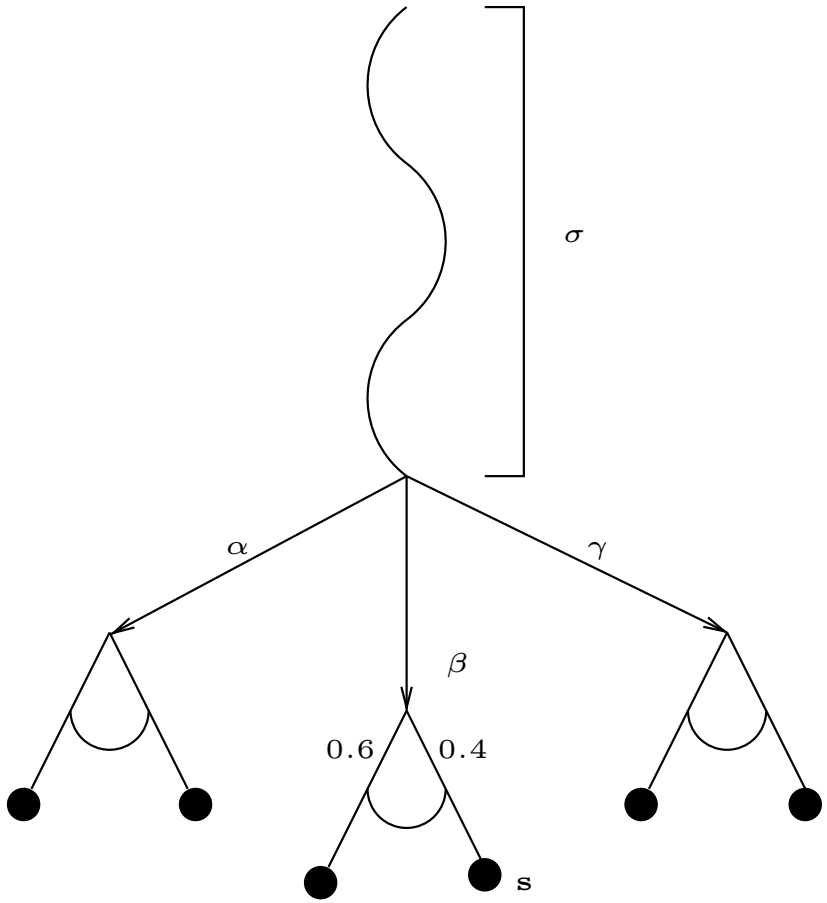

- Schedulers are mathematical entities that resolve non-determinism
- A scheduler map traces of the system to (distributions on) transitions
- By fixing the scheduler, the MDP becomes a Markov chain

Markov decision process need schedulers to yield probabilities

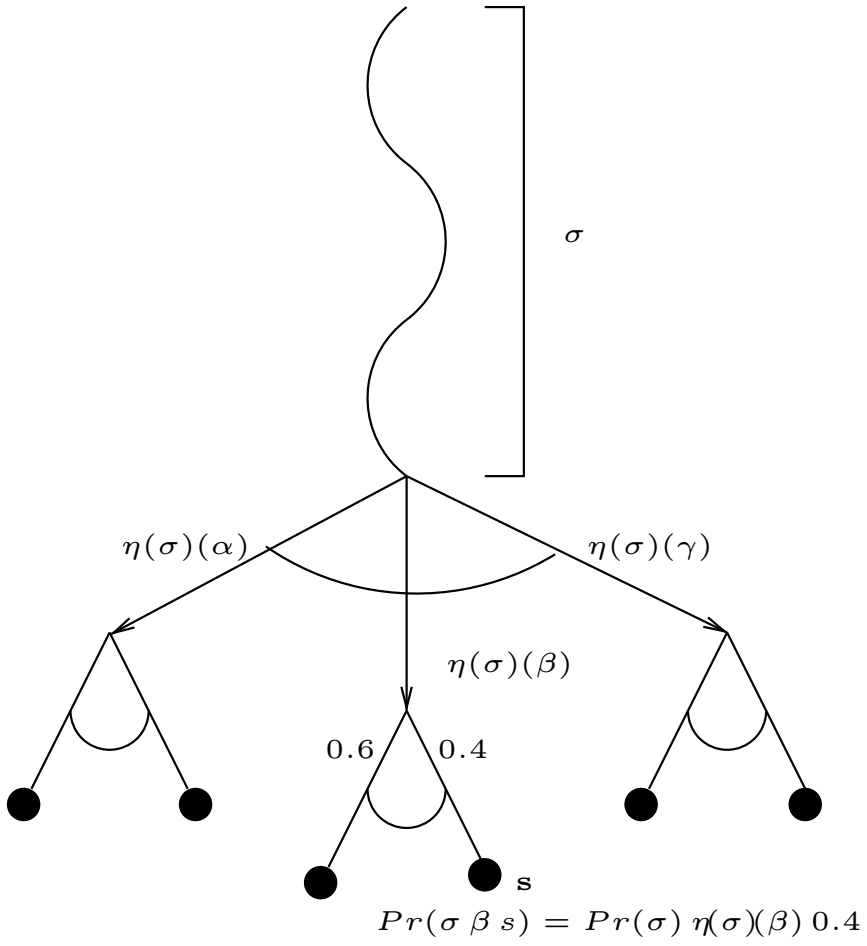

- Schedulers are mathematical entities that resolve non-determinism
- A scheduler map traces of the system to (distributions on) transitions
- By fixing the scheduler, the MDP becomes a Markov chain

Markov decision process need schedulers to yield probabilities

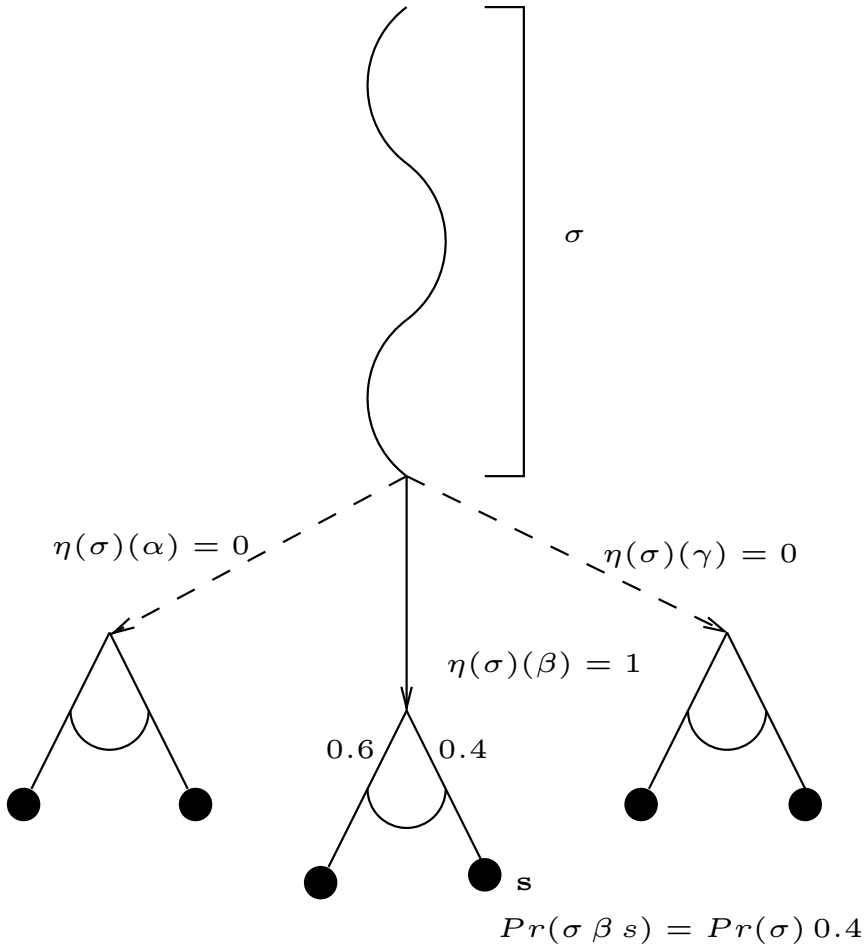

- Schedulers are mathematical entities that resolve non-determinism
- A scheduler map traces of the system to (distributions on) transitions
- By fixing the scheduler, the MDP becomes a Markov chain
- The schedulers that always choose a single transition are called *deterministic*

Schedulers determine a probability value for a given property

The infimum/supremum probability of a property depends on the set of schedulers considered

- A property determines a set of infinite traces in which the property holds
- The probability of a property under a given scheduler is the probability of the traces complying with the property under the scheduler
- We are interested in the maximum/minimum probabilities of a property
- Temporal logics used for MDPs are based on statements related maximum/minimum probabilites (for example, the maximum probability is less than 0.5)

Schedulers may be more powerful than needed

- Usually, the set of all schedulers is considered
- However, they may result too pessimistic
- In this example,

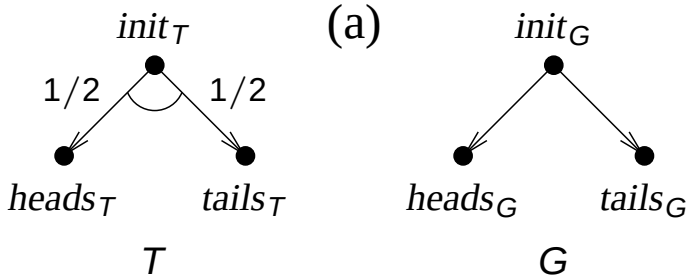

it is assumed that a component may guess the behaviour of the other one.

Distributed schedulers yield better bounds for probabilities

- Distributed schedulers are schedulers that, when choosing among transitions of a given component, the choice is based solely on the local history of that component
- This schedulers can be obtained by composing local schedulers for each component

Note that, in the example, the choice among head or tails in G cannot be based on the result of the transition executed by T, thus yielding a probability of agreement of $1/2$

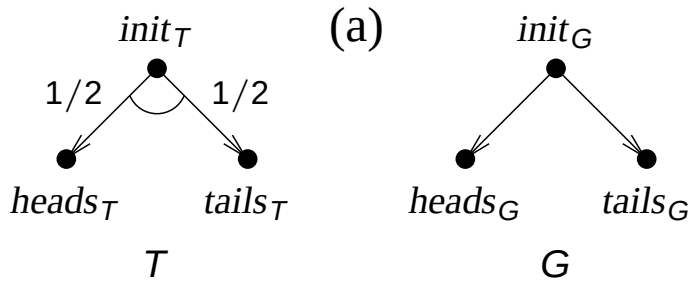

Distributed schedulers have been introduced to allow compositional reasoning

- In the example, the trace distributions of Early and Late are the same. However, the trace distribution shown below is a t.d. of Late $\|$ Coin but not a t.d. of Early $\|$ Coin
- If the schedulers are restricted to be distributed, such a t.d. is not a t.d. of Late $\|$ Coin anymore.

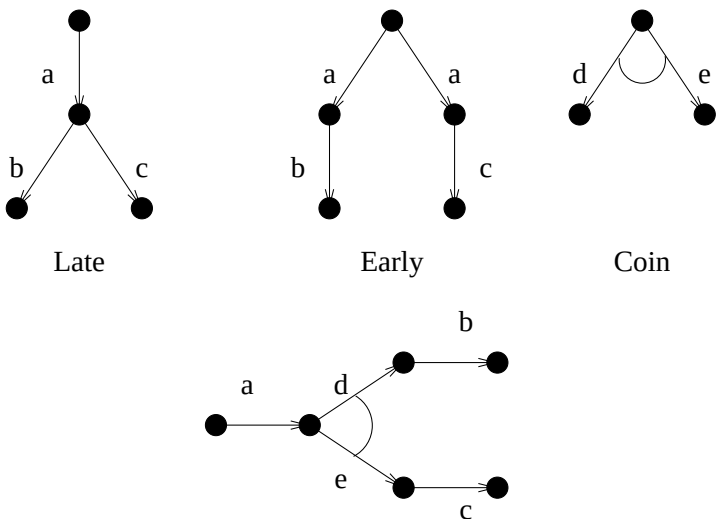

The model checking problem is undecidable when restricted to distributed schedulers

- The model-checking problem is decidable in polynomial time if the schedulers are not restricted
- We showed that the maximum probability of a given property cannot be calculated nor even approximated when the schedulers are restricted to be distributed (Giro/D'Argenio, FORMATS 2007)

So far, there are no nondeterministic choices concerning the interleaving

- In the framework proposed in (de Alfaro/Henzinger/Jhala, CONCUR 2001) the components execute in a completely synchronous fashion.
- In the framework introduced in Ling Cheung's thesis (also Cheung/Lynch/Segala/Vaandrager, TCS) the component that owns a token is the one allowed to execute

If *arbitrary* nondeterministic choices are allowed, compositionality is lost

- In the example, the trace distributions of Product and A ‖ B are the same. However, the trace distribution shown below is a t.d. of A ‖ B ‖ Coin but not a t.d. of Process ‖ Coin
- What kinds of restrictions on the interleaving should we introduce in order to avoid A ‖ B to guess the outcome of Coin?

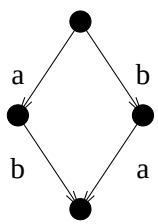

Product

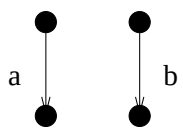

A B

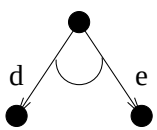

Coin

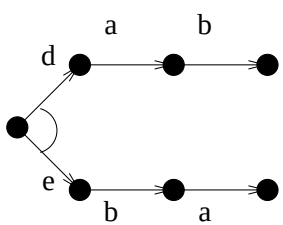

The challenge: to find a set of schedulers in which the interleaving choices are nondeterministic, but still yielding compositionality and/or realistic bounds.