Markovian Behavioral Equivalences: A Comparative Survey

Marco Bernardo

University of Urbino – Italy

Table of Contents

Part I:
Introduction

Performance-Oriented Notations

- Building performance-aware models of computing systems:
  - Predicting the satisfiability of QoS requirements.
  - Choosing among alternative designs based on their expected QoS.
- *Theory*:
  - Queueing networks (1950's).
  - Stochastic Petri nets (1980's).
  - Stochastic process calculi (1990's).
- *Practice*:
  - System/software performance engineering approaches.
  - Object-oriented modeling languages (UML profiles).
  - Architectural description languages (ÆMILIA).
  - Formal modeling languages (MODEST).
  - Coordination languages (STOKLAIM).

- Performance-oriented notations usually produce behavioral models.
- These models can be uniformly expressed as state transition graphs.
- Representation of the current state:
  - ⊙ Current number of customers in each service center.
  - ⊙ Current Petri net marking.
  - ⊙ Current process term.
- Cause of the state change associated with a transition:
  - ⊙ Execution of a certain activity.
  - ⊙ Occurrence of a certain event.
- Adoption of an interleaving view of concurrency in which independent activities can be executed in any order but not simultaneously.

Behavioral Equivalences

- Behavioral models are equivalent whenever they represent systems that behave the same.
- Need for the introduction of behavioral equivalences.
- Useful for theoretical and applicative purposes:
  - ⊙ Comparing models that are syntactically different on the basis of the behavior they exhibit.
  - ⊙ Relating models of the same system at different abstraction levels (top-down modeling).
  - ⊙ Manipulating models in a way that preserves certain properties (state space reduction before analysis).

- Most studied approaches developed in a purely functional framework:
  - Bisimilarity: two models are equivalent if they are able to *mimic* each other's behavior *stepwise*.
  - Testing: two models are equivalent if an *external observer* cannot distinguish between them by interacting with them by means of *tests* and comparing their reactions.
  - Trace: two models are equivalent if they are able to perform the same *sequences* of activities.
- How to extend behavioral equivalences to performance-aware models?
- It is necessary to take into account *quantitative aspects* related to system evolution over time (event probabilities, activity durations, costs/gains, ...).

Markovian Framework

- A Markov chain is a discrete-state stochastic process $\{RV(t) \mid t \in \mathbb{R}_{\geq 0}\}$ such that for all $n \in \mathbb{N}$, time instants $t_0 < t_1 < \ldots < t_n < t_{n+1}$, and states $s_0, s_1, \ldots, s_n, s_{n+1} \in S$:

$$\boxed{\begin{aligned}&\Pr\{RV(t_{n+1}) = s_{n+1} \mid RV(t_0) = s_0 \wedge RV(t_1) = s_1 \wedge \ldots \wedge RV(t_n) = s_n\}\\ &\quad = \Pr\{RV(t_{n+1}) = s_{n+1} \mid RV(t_n) = s_n\}\end{aligned}}$$

- The past history is completely summarized by the current state.

- Equivalently, the stochastic process has no memory of the past.

- Time homogeneity: probabilities independent of state change times.

- The solution of a Markov chain is its state probability distribution $\boldsymbol{\pi}()$ at an arbitrary time instant (CTMC vs. DTMC).

- Representation and solution of a continuous-time Markov chain (CTMC):
  - State transitions are described by a rate matrix $\boldsymbol{Q}$.
  - The sojourn time in any state is exponentially distributed.
  - Given $\boldsymbol{\pi}(0)$, the transient solution $\boldsymbol{\pi}(t)$ is obtained by solving:

$$\boxed{\boldsymbol{\pi}(t) \cdot \boldsymbol{Q} \;=\; \frac{\mathrm{d}\boldsymbol{\pi}(t)}{\mathrm{d}t}}$$

  - The stationary solution $\boldsymbol{\pi} = \lim\limits_{t \to \infty} \boldsymbol{\pi}(t)$ is obtained (if any) by solving:

$$\boxed{\begin{array}{rcl} \boldsymbol{\pi} \cdot \boldsymbol{Q} & = & \mathbf{0} \\ \sum\limits_{s \in S} \pi[s] & = & 1 \end{array}}$$

- Exponentially distributed random variables are the only continuous random variables satisfying the memoryless property:

$$\boxed{\Pr\{RV \leq v + v' \mid RV > v'\} \;=\; \Pr\{RV \leq v\}}$$

- Every CTMC (time-aware model) has an embedded DTMC (time-abstract model):
  - State transitions are described by a probability matrix $\boldsymbol{P}$.
  - $\boldsymbol{P}$ is obtained from $\boldsymbol{Q}$ by dividing the rate of each transition by the sum of the rates of the transitions that depart from the source state.
  - The sojourn time in any state is geometrically distributed.
  - Given $\boldsymbol{\pi}(0)$, the transient solution $\boldsymbol{\pi}(n)$ is computed as follows:

$$\boxed{\boldsymbol{\pi}(n) \; = \; \boldsymbol{\pi}(0) \cdot \boldsymbol{P}^n}$$

  - The stationary solution $\boldsymbol{\pi} = \lim_{n \to \infty} \boldsymbol{\pi}(n)$ is obtained (if any) by solving:

$$\boxed{\begin{aligned} \boldsymbol{\pi} \; &= \; \boldsymbol{\pi} \cdot \boldsymbol{P} \\ \sum_{s \in S} \pi[s] \; &= \; 1 \end{aligned}}$$

- A CTMC is a state transition graph in which every transition is labeled with a positive real number expressing the rate at which the state change takes place.
- Rates subsume both time information and probability information:
  - The sojourn time in a state is exponentially distributed with rate given by the sum of the rates of the outgoing transitions.
  - The probability of executing a transition is proportional to its rate.
- A CTMC can thus be viewed as a state transition graph in which:
  - Every state has an exponentially distributed random variable associated with it that expresses the sojourn time.
  - Every transition has a positive real number not greater than 1 associated with it that expresses the execution probability.

Markovian Behavioral Equivalences

- Focus on exponential distributions for activity durations.
- Their memoryless property results in a simpler mathematical treatment:
  - ◁ Compliance with the *interleaving view* of concurrency.
  - ◁ Easy calculation of *state sojourn times* and *transition probabilities*.

  without sacrificing expressiveness:
  - ▷ Adequate for modeling the timing of many *real-life phenomena* like arrival processes, failure events, and chemical reactions.
  - ▷ Most appropriate stochastic approximation in the case in which only the *average duration* of an activity is known.
  - ▷ Proper combinations (phase-type distributions) approximate most of *general distributions* arbitrarily closely.
- How to define Markovian behavioral equivalences?

- Comparison criteria for Markovian behavioral equivalences:

  1. Discriminating power:
     which of them is finer/coarser than the others?
  2. Congruence:
     do they support compositional reasoning?
  3. Sound and complete axiomatization:
     what are their fundamental equational laws?
  4. Modal logic characterization:
     what behavioral properties do they preserve?
  5. Complexity (of their verification algorithms):
     can they be checked for efficiently?
  6. Exactness (of their induced CTMC-level aggregations):
     do they make sense from a performance viewpoint?

- Congruence enables the equivalence-based compositional reduction of models obtained as the combination of submodels.
- Axioms can be used as rewriting rules that syntactically manipulate models in a way that is consistent with the equivalence.
- The modal logic characterization provides diagnostic information in the form of distinguishing formulas that explain model inequivalence.
- *Exactness*: the probability of being in a macrostate of an aggregated CTMC is the sum of the probabilities of being in one of the constituent microstates of the original CTMC (transient/stationary).
- Exactness guarantees the preservation of performance characteristics when going from the original CTMC to the aggregated one induced by the equivalence.

Running Example: Producer-Consumer System

- General description:
  - Three components: producer, finite-capacity buffer, consumer.
  - The producer deposits items into the buffer at rate $\lambda \in \mathbb{R}_{>0}$ as long as the buffer capacity is not exceeded.
  - Stored items are then withdrawn by the consumer at rate $\mu \in \mathbb{R}_{>0}$ according to some predefined discipline (like FIFO or LIFO).
- Specific scenario:
  - The buffer has only two positions.
  - Items are identical, hence the discipline is not important.

Part II:
Markovian Process Algebra

Process Algebraic Markovian Modeling

- Behavioral equivalences abstract from the specific kind of model but ...
- ... are better investigated and understood in a process algebraic setting.
- Action-based modeling relying on a set of behavioral operators.
- Performance-oriented process calculi with CTMC semantics:
  - TIPP [Götz, Herzog, Rettelbach].
  - PEPA [Hillston].
  - MPA [Buchholz].
  - $\mathrm{EMPA}_{\mathrm{gr}}$ [Bernardo, Bravetti, Gorrieri].
  - $\mathrm{S}\pi$ [Priami].
  - IMC [Hermanns].
  - PIOA [Stark, Cleaveland, Smolka].

- Markovian process calculi differ for the *action representation*.
- Durational actions (integrated time):
  - An action is executed while time passes.
  - Single action prefix operator comprising the name a of the action and the rate $\lambda \in \mathbb{R}_{>0}$ of the exponentially distributed random variable quantifying the duration of the action: $<a, \lambda>.\_$
  - The choice among several actions is probabilistic.
  - TIPP, PEPA, MPA, $\mathrm{EMPA}_{\mathrm{gr}}$, S$\pi$, PIOA.
- Action names separated from time (orthogonal time):
  - An action is instantaneously executed after some time has elapsed.
  - Two action prefix operators: $(\lambda).\_$ and $a.\_$
  - The choice among several actions is nondeterministic.
  - IMC.

- Markovian process calculi also differ for the discipline adopted for *action synchronization*.
- In the orthogonal time case, action synchronization is governed as in nondeterministic process calculi.
- In the integrated time case, action synchronization can be handled in different ways.
- The more natural choice for deciding the duration of the synchronization of two exponentially timed actions would be to take the maximum of their durations.
- *The maximum of two exponentially distributed random variables is not exponentially distributed* (phase-type: IMC).

- Symmetric synchronizations:
  - The synchronization of two exponentially timed actions is assumed to be exponentially timed.
  - Its rate is defined through an associative and commutative operator applied to the two original rates (multiplication, min, max).
  - TIPP, PEPA, MPA, Sπ.
- Asymmetric synchronizations:
  - Passive actions of the form $<a, *_w>$ whose duration is unspecified.
  - An exponentially timed action can synchronize only with a passive action, thus determining the duration of the synchronization.
  - PEPA, $\mathrm{EMPA}_{\mathrm{gr}}$, PIOA.
- *Bounded capacity assumption*: the rate of an action should not increase when synchronizing that action with other actions.

Markovian Process Calculus: Syntax

- Basic design choices: durational actions (more natural modeling style) and asymmetric synchronizations (exp. timed action can synch. only with passive actions).

- $Name_{\mathrm{v}}$: set of visible action names.

- $Name = Name_{\mathrm{v}} \cup \{\tau\}$: set of all action names.

- $Rate = \mathbb{R}_{>0} \cup \{*_w \mid w \in \mathbb{R}_{>0}\}$: set of action rates.

- $Act_{\mathrm{M}} = Name \times Rate$: set of exponentially timed and passive actions.

- $Relab = \{\varphi : Name \rightarrow Name \mid \varphi^{-1}(\tau) = \{\tau\}\}$: set of visibility-preserv. relabeling functions.

- Var: set of process variables ($Const$: set of process constants).

- Process term syntax for process language $\mathcal{PL}_{\mathrm{M}}$:

$P ::= \underline{0}$	inactive process	
$\mid\ <a, \lambda>.P$	exp. timed action prefix	$(a \in \mathit{Name}, \lambda \in \mathbb{R}_{>0})$
$\mid\ <a, *_w>.P$	passive action prefix	$(a \in \mathit{Name}, w \in \mathbb{R}_{>0})$
$\mid\ P + P$	alternative composition	
$\mid\ P \parallel_S P$	parallel composition	$(S \subseteq \mathit{Name}_{\mathrm{v}})$
$\mid\ P \,/\, H$	hiding	$(H \subseteq \mathit{Name}_{\mathrm{v}})$
$\mid\ P \setminus L$	restriction	$(L \subseteq \mathit{Name}_{\mathrm{v}})$
$\mid\ P[\varphi]$	relabeling	$(\varphi \in \mathit{Relab})$
$\mid\ X$	process variable	$(X \in \mathit{Var})$
$\mid\ \mathrm{rec}\, X : P$	recursion	$(X \in \mathit{Var})$

(process constants are defined by means of equations of the form $B \triangleq P$).

- The duration of $<a, \lambda>$ is the exponentially distributed random variable Exp_λ, where $\Pr\{Exp_\lambda \le t\} = 1 - \mathrm{e}^{-\lambda \cdot t}$ and $\mathrm{E}\{Exp_\lambda\} = 1 / \lambda$.

- The choice among exp. timed actions is generative (prob. over arbitrary names) and is solved by applying the race policy (exec. prob. proportional to action rates).

- The duration of $<a, *_w>$ is unspecified (synch. with exponentially timed action).

- The choice among passive actions is reactive (prob. restricted to same name):
  - Probabilistic for passive actions with the same name and solved by applying the preselection policy (exec. prob. proportional to action weights).
  - Nondeterministic for passive actions with different names.

- The choice between an exponentially timed action and a passive action is nondeterministic.

- Applying the race policy to the exponentially timed actions $(\lambda_1, \ldots, \lambda_h)$ enabled by a process term means executing the *fastest* of those actions.
- The sojourn time associated with that term is thus the *minimum* of the random variables quantifying the durations of those actions.
- The sojourn time is exponentially distributed because:

$$\boxed{\min(Exp_{\lambda_1}, \ldots, Exp_{\lambda_h}) \ = \ Exp_{\lambda_1 + \ldots + \lambda_h}}$$

- The average sojourn time is therefore given by $1 \, / \, (\lambda_1 + \ldots + \lambda_h)$.
- The execution probability of exponentially timed action with rate λ_i is $\lambda_i \, / \, (\lambda_1 + \ldots + \lambda_h)$.

- $P_1 + P_2$ behaves as P_1 or P_2 depending on which executes first.
- The choice among several enabled actions is solved by applying either the race policy or the preselection policy.
- The choice is internal if the enabled actions are all invisible, otherwise the choice can be influenced by the external environment.
- $P_1 \parallel_S P_2$ behaves as P_1 in parallel with P_2 under synchronization set S.
- Actions whose name does not belong to S are executed autonomously by P_1 and by P_2 (order determined by race/preselection policy).
- Synchronization is forced between any action enabled by P_1 and any action enabled by P_2 that have the same name belonging to S, in which case the resulting action has the same name as the two original actions ($S = \emptyset$ implies P_1 and P_2 fully independent, $S = Name_{\mathrm{v}}$ implies P_1 and P_2 fully synchronized).

- $\underline{0}$ is a terminated process and hence cannot execute any action.
- $<a, \tilde{\lambda}>.P$ can perform action a at rate $\tilde{\lambda}$ and then behaves as P (action-based sequential composition).
- $P \,/\, H$ behaves as P but every action belonging to H is turned into τ (abstraction mechanism; can be used for preventing a process from communicating).
- $P \setminus L$ behaves as P but every action belonging to L is forbidden (same effect as $P \parallel_L \underline{0}$).
- $P[\varphi]$ behaves as P but every action is renamed according to function φ (redundance avoidance; encoding of the previous two operators if φ is non-visib.-pres./partial).
- Operator precedence: unary operators $> + > \parallel$.
- Operator associativity: $+$ and $\parallel$ are left associative.

- $\text{rec}\,X : P$ behaves as P with every free occurrence of process variable X being replaced by $\text{rec}\,X : P$.
- A process variable is said to occur *free* in a process term if it is not in the scope of a rec binder for that variable, otherwise it is said to be *bound* in that process term.
- A process term is said to be *closed* if all of its occurrences of process variables are bound, otherwise it is said to be *open*.
- A process term is said to be *guarded* iff all of its occurrences of process variables are in the scope of action prefix operators.
- $\mathbb{P}_{\mathrm{M}}$: set of closed and guarded process terms (fully defined, finitely branching).

- **Running example** (MPC syntax):
  - Conventions: action names are verbs composed of lower-case letters, process constant names are nouns starting with an upper-case letter.
  - The only observable activities are deposits and withdrawals.
  - Names of visible actions: *deposit* and *withdraw*.
  - Structure-independent process algebraic description:

$$
\begin{aligned}
ProdCons^{\mathrm{M}}_{0/2} &\triangleq <deposit, \lambda>.ProdCons^{\mathrm{M}}_{1/2} \\
ProdCons^{\mathrm{M}}_{1/2} &\triangleq <deposit, \lambda>.ProdCons^{\mathrm{M}}_{2/2} + \\
& <withdraw, \mu>.ProdCons^{\mathrm{M}}_{0/2} \\
ProdCons^{\mathrm{M}}_{2/2} &\triangleq <withdraw, \mu>.ProdCons^{\mathrm{M}}_{1/2}
\end{aligned}
$$

  - Specification to which every correct implementation should conform.

Markovian Process Calculus: Semantics

- State transition graph expressing all computations and branching points and accounting for transition multiplicity ($<a, \lambda>.\underline{0} + <a, \lambda>.\underline{0}$ vs. $<a, \lambda>.\underline{0}$).

- Every $P \in \mathbb{P}_{\mathrm{M}}$ is mapped to a labeled multitransition system $[\![P]\!]_{\mathrm{M}}$:
  - Each state corresponds to a process term into which P can evolve.
  - The initial state corresponds to P.
  - Each transition from a source state to a target state is labeled with the action that determines the corresponding state change.

- Every $P \in \mathbb{P}_{\mathrm{M,pc}}$ is mapped to a CTMC (performance closure if no passive trans.):
  - Dropping action names from all transitions of $[\![P]\!]_{\mathrm{M}}$.
  - Collapsing all the transitions between any two states of $[\![P]\!]_{\mathrm{M}}$ into a single transition by summing up the rates of the original transitions.

- Derivation of one single transition at a time by applying suitable operational semantic rules to the source state of the transition.

- Rules defined by induction on the syntactical structure of process terms.

- Basic rules for action prefix, inductive rules for all the other operators.

- Different formats: dynamic operators (. +), static operators (∥ / \ [ ]).

- The multitransition relation $\xrightarrow{\quad}_{\mathrm{M},P}$ of $[\![P]\!]_{\mathrm{M}}$ is contained in the smallest multiset of elements of $\mathbb{P}_{\mathrm{M}} \times Act_{\mathrm{M}} \times \mathbb{P}_{\mathrm{M}}$ that:
  - Satisfy the operational semantic rules.
  - Keep track of all the possible ways of deriving each transition.

- No rule for $\underline{0}$: $[\![\underline{0}]\!]_{\mathrm{M}}$ has a single state and no transitions.

- Operational semantic rules for action prefix:

$$<a, \lambda>.P \xrightarrow{a,\lambda}_{\mathrm{M}} P$$

$$<a, *_w>.P \xrightarrow{a,*_w}_{\mathrm{M}} P$$

- Operational semantic rules for alternative composition:

$$\frac{P_1 \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P'}{P_1 + P_2 \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P'} \qquad \frac{P_2 \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P'}{P_1 + P_2 \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P'}$$

- Operational semantic rule for recursion:

$$\frac{P\{\mathrm{rec}\, X : P \hookrightarrow X\} \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P'}{\mathrm{rec}\, X : P \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P'} \qquad \left(\frac{B \triangleq P \qquad P \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P'}{B \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P'} \right)$$

- Classical interleaving semantics for parallel composition:
  - ⊙ *Due to the memoryless property of the exponential distribution, the execution of an exponentially timed action can be thought of as being started in the last state in which the action is enabled.*
  - ⊙ *Due to the infinite support of the exponential distribution, the probability of simultaneous termination of two concurrent exponentially timed actions is zero.*
- Operational semantic rules for parallel execution:

$$\boxed{\frac{P_1 \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P_1' \qquad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P_1' \parallel_S P_2} \qquad \frac{P_2 \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P_2' \qquad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P_1 \parallel_S P_2'}}$$

- The following process terms represent structurally different systems:

$$<a, \lambda>.\underline{0} \parallel_{\emptyset} <b, \mu>.\underline{0}$$

$$<a, \lambda>.<b, \mu>.\underline{0} + <b, \mu>.<a, \lambda>.\underline{0}$$

  but they are indistinguishable by an external observer.

- Black-box semantics given by the same labeled multitransition system:

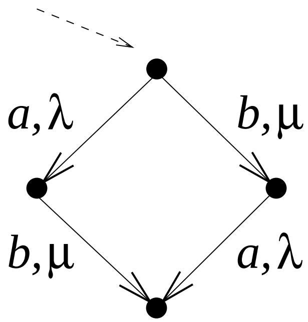

- Interleave concurrent exponentially timed actions without the need of adjusting their rates inside transition labels.

- Synchronization admitted among several actions with the same name, provided that at most one of them is exponentially timed.

- Generative-reactive or reactive-reactive synchronizations.

- The rate of the synchronization of an exponentially timed action with a passive action is given by the rate of the former multiplied by the execution probability of the latter (complies with the bounded capacity assumption).

- Weight of a process term P with respect to passive actions of name a:

$$weight(P, a) \;=\; \sum \{\!| \, w \in \mathbb{R}_{>0} \mid \exists P' \in \mathbb{P}_{\mathrm{M}} .\, P \xrightarrow{a, *_w}_{\mathrm{M}} P' \, |\!\}$$

- Normalizing function for reactive-reactive synchronizations:

$$norm(w_1, w_2, a, P_1, P_2) \;=\; \frac{w_1}{weight(P_1, a)} \cdot \frac{w_2}{weight(P_2, a)} \cdot (weight(P_1, a) + weight(P_2, a))$$

- Operational semantic rules for generative-reactive synchronization:

$$\frac{P_1 \xrightarrow{a,\lambda}_{\mathrm{M}} P'_1 \qquad P_2 \xrightarrow{a,*_w}_{\mathrm{M}} P'_2 \qquad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a,\lambda\cdot\frac{w}{weight(P_2,a)}}_{\mathrm{M}} P'_1 \parallel_S P'_2}$$

$$\frac{P_1 \xrightarrow{a,*_w}_{\mathrm{M}} P'_1 \qquad P_2 \xrightarrow{a,\lambda}_{\mathrm{M}} P'_2 \qquad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a,\lambda\cdot\frac{w}{weight(P_1,a)}}_{\mathrm{M}} P'_1 \parallel_S P'_2}$$

- Operational semantic rule for reactive-reactive synchronization:

$$\frac{P_1 \xrightarrow{a,*_{w_1}}_{\mathrm{M}} P'_1 \qquad P_2 \xrightarrow{a,*_{w_2}}_{\mathrm{M}} P'_2 \qquad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a,*_{norm(w_1,w_2,a,P_1,P_2)}}_{\mathrm{M}} P'_1 \parallel_S P'_2}$$

- Operational semantic rules for hiding, restriction, relabeling:

$$\frac{P \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P' \qquad a \in H}{P \,/\, H \xrightarrow{\tau,\tilde{\lambda}}_{\mathrm{M}} P' \,/\, H} \qquad \frac{P \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P' \qquad a \notin H}{P \,/\, H \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P' \,/\, H}$$

$$\frac{P \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P' \qquad a \notin L}{P \setminus L \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P' \setminus L}$$

$$\frac{P \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} P'}{P[\varphi] \xrightarrow{\varphi(a),\tilde{\lambda}}_{\mathrm{M}} P'[\varphi]}$$

- $[\![P]\!]_{\mathrm{M}}$ is finite state if no recursive definition in P contains static ops.

- **Running example** (MPC semantics):
  - Labeled multitransition system $[\![\mathit{ProdCons}^{\mathrm{M}}_{0/2}]\!]_{\mathrm{M}}$ with explicit states:

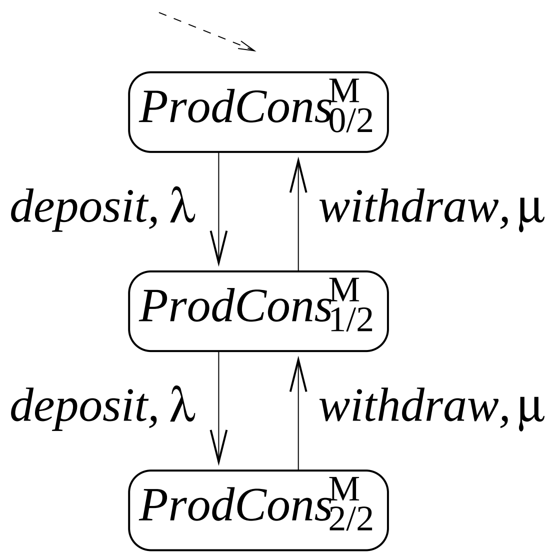

  - Obtained by mechanically applying the operational semantic rules for process constant, alternative composition, and action prefix.

Part III:
Markovian Bisimulation Equivalence

Equivalence Definition

- Two process terms are equivalent if they are able to mimic each other's functional and performance behavior stepwise.
- Whenever a process term can perform actions with a certain name that reach a certain set of terms at a certain speed, then any process term equivalent to the given one has to be able to respond with actions with the same name that reach an equivalent set of terms at the same speed.
- Comparison of process term exit rates rather than individual transitions (different from bisimulation equivalence for nondeterministic processes).
- High sensitivity to the branching structure of process terms.

- The exit rate of a process term is the rate at which the process term can execute actions of a given name that lead to a given set of terms (sum of the rates of those actions due to the race policy).

- Exit rate at which $P \in \mathbb{P}_{\mathrm{M}}$ executes actions of name $a \in Name$ and level $l \in \{0, -1\}$ (0 exp. timed, -1 passive) that lead to destination $D \subseteq \mathbb{P}_{\mathrm{M}}$:

$$\boxed{rate_{\mathrm{e}}(P, a, l, D) \;=\; \begin{cases} \sum \{\!| \lambda \in \mathbb{R}_{>0} \mid \exists P' \in D.\, P \xrightarrow{a,\lambda}_{\mathrm{M}} P' |\!\} & \text{if } l = 0 \\ \sum \{\!| w \in \mathbb{R}_{>0} \mid \exists P' \in D.\, P \xrightarrow{a,*_w}_{\mathrm{M}} P' |\!\} & \text{if } l = -1 \end{cases}}$$

- Overall exit rate of P w.r.t. a at level l: $rate_{\mathrm{o}}(P, a, l) = rate_{\mathrm{e}}(P, a, l, \mathbb{P}_{\mathrm{M}})$.

- Total exit rate of P at level l: $rate_{\mathrm{t}}(P, l) \;=\; \sum_{a \in Name} rate_{\mathrm{o}}(P, a, l)$.

- $1 \,/\, rate_{\mathrm{t}}(P, 0)$ is the average sojourn time of P when $P \in \mathbb{P}_{\mathrm{M,pc}}$.

- The exit probability of a process term is the probability with which the process term can execute actions of a given name that lead to a given set of terms.

- Generative probability for exponentially timed actions (arbitrary names).

- Reactive probability for passive actions (restriction to same name).

- Exit probability with which $P \in \mathbb{P}_{\mathrm{M}}$ executes actions of name $a \in \mathit{Name}$ and level $l \in \{0, -1\}$ that lead to destination $D \subseteq \mathbb{P}_{\mathrm{M}}$:

$$\boxed{\mathit{prob}_{\mathrm{e}}(P, a, l, D) \; = \; \begin{cases} \mathit{rate}_{\mathrm{e}}(P, a, l, D) \, / \, \mathit{rate}_{\mathrm{t}}(P, l) & \text{if } l = 0 \\ \mathit{rate}_{\mathrm{e}}(P, a, l, D) \, / \, \mathit{rate}_{\mathrm{o}}(P, a, l) & \text{if } l = -1 \end{cases}}$$

- An equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M}}$ is a Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \mathit{Name}$, levels $l \in \{0, -1\}$, and equivalence classes $D \in \mathbb{P}_{\mathrm{M}}/\mathcal{B}$:

$$\boxed{\mathit{rate}_{\mathrm{e}}(P_1, a, l, D) \; = \; \mathit{rate}_{\mathrm{e}}(P_2, a, l, D)}$$

- Markovian bisimulation equivalence $\sim_{\mathrm{MB}}$ is the union of all the Markovian bisimulations.

- *A consequence of the coinductive nature of* $\sim_{\mathrm{MB}}$ *is that the derivatives of two equivalent terms are still equivalent.*

- $\sim_{\mathrm{MB}}$ is strictly finer than classical bisimilarity ($a \neq b$ and $\lambda \neq \mu$):

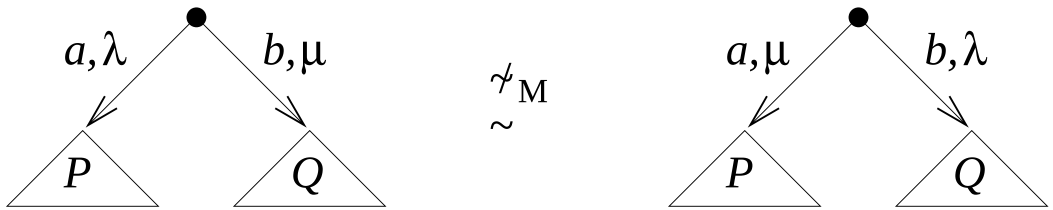

- $\sim_{\mathrm{MB}}$ is strictly finer than probabilistic bisimilarity:

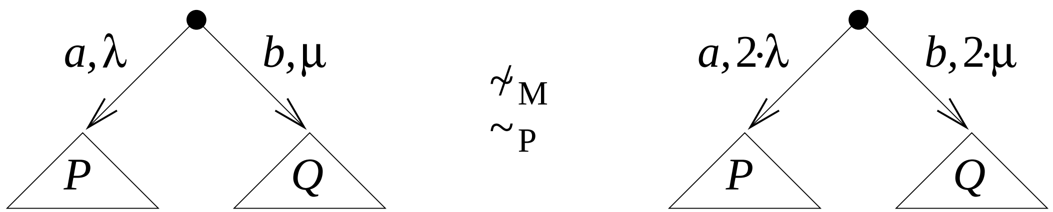

- **Running example** ($\sim_{\rm MB}$):
  - Concurrent implementation with two independent one-pos. buffers:

$$\begin{array}{rcl} PC^{\rm M}_{\rm conc,2} & \triangleq & Prod^{\rm M} \parallel_{\{deposit\}} (Buff^{\rm M} \parallel_{\emptyset} Buff^{\rm M}) \parallel_{\{withdraw\}} Cons^{\rm M} \\ Prod^{\rm M} & \triangleq & <deposit, \lambda>.Prod^{\rm M} \\ Buff^{\rm M} & \triangleq & <deposit, *_1>.<withdraw, *_1>.Buff^{\rm M} \\ Cons^{\rm M} & \triangleq & <withdraw, \mu>.Cons^{\rm M} \end{array}$$

  - All the actions occurring in the buffer are passive, consistent with the fact that the buffer is a passive entity.
  - Is $PC^{\rm M}_{\rm conc,2}$ a correct implementation of $ProdCons^{\rm M}_{0/2}$?
  - Yes, because it turns out that $PC^{\rm M}_{\rm conc,2} \sim_{\rm MB} ProdCons^{\rm M}_{0/2}$.
  - Proved by finding a suitable Markovian bisimulation.

- Markovian bisimulation proving $PC^{\mathrm{M}}_{\mathrm{conc},2} \sim_{\mathrm{MB}} \mathit{ProdCons}^{\mathrm{M}}_{0/2}$, with states of the same color belonging to the same equivalence class:

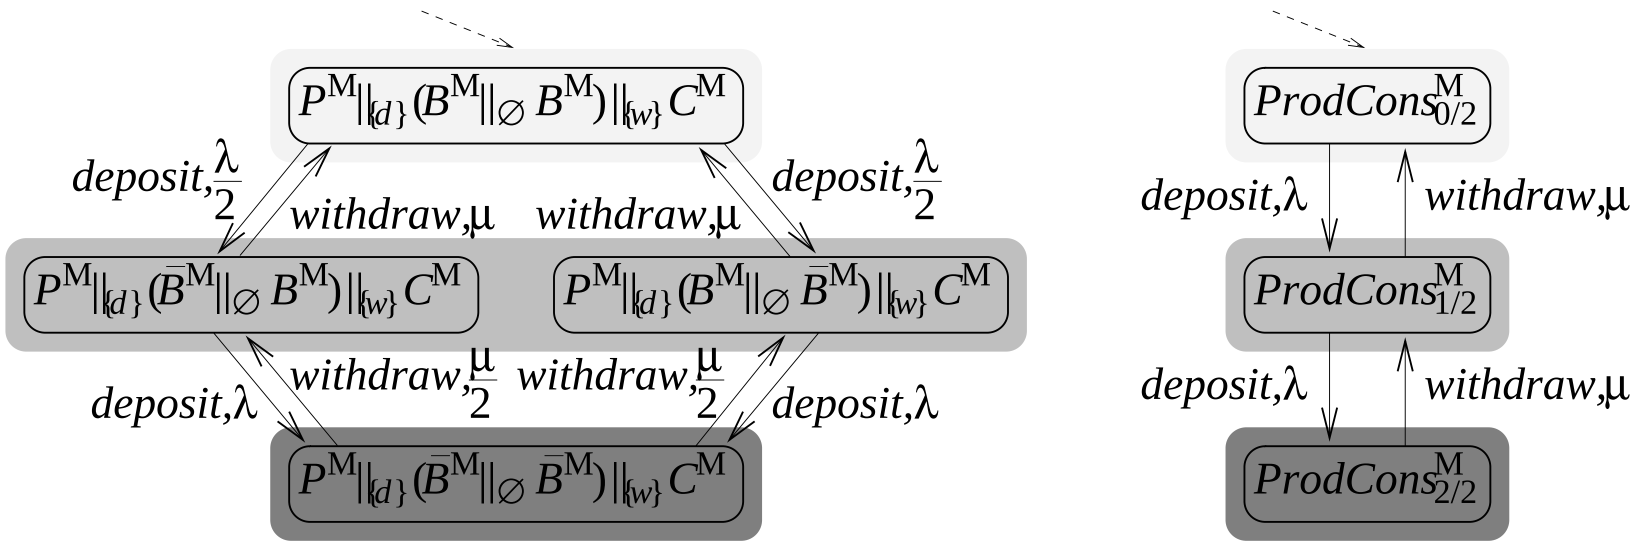

- The initial state on the left-hand side has both outgoing transitions labeled with $\lambda/2$, not λ.

- The bottom state on the left-hand side has both outgoing transitions labeled with $\mu/2$, not μ.

Conditions and Characterizations

- In order for $P_1 \sim_{\mathrm{MB}} P_2$, it is necessary that for all $a \in \mathit{Name}$ and $l \in \{0, -1\}$:

$$\boxed{\mathit{rate}_{\mathrm{o}}(P_1, a, l) \; = \; \mathit{rate}_{\mathrm{o}}(P_2, a, l)}$$

- A binary relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M}}$ is a Markovian bisimulation up to $\sim_{\mathrm{MB}}$ iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \mathit{Name}$, levels $l \in \{0, -1\}$, and equivalence classes $D \in \mathbb{P}_{\mathrm{M}} / (\mathcal{B} \cup \mathcal{B}^{-1} \cup \sim_{\mathrm{MB}})^{+}$:

$$\boxed{\mathit{rate}_{\mathrm{e}}(P_1, a, l, D) \; = \; \mathit{rate}_{\mathrm{e}}(P_2, a, l, D)}$$

- Focus on important pairs of process terms that form a bisimulation.

- In order for $P_1 \sim_{\mathrm{MB}} P_2$, it is sufficient to find a Markovian bisimulation up to $\sim_{\mathrm{MB}}$ that contains (P_1, P_2).

- $\sim_{\mathrm{MB}}$ has an alternative characterization in which time and probability are kept separate (instead of being both subsumed by rates).

- An equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M}}$ is a separate Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \mathit{Name}$ and levels $l \in \{0, -1\}$:

$$\boxed{\mathit{rate}_{\mathrm{o}}(P_1, a, l) \;=\; \mathit{rate}_{\mathrm{o}}(P_2, a, l)}$$

  and for all equivalence classes $D \in \mathbb{P}_{\mathrm{M}}/\mathcal{B}$:

$$\boxed{\mathit{prob}_{\mathrm{e}}(P_1, a, l, D) \;=\; \mathit{prob}_{\mathrm{e}}(P_2, a, l, D)}$$

- Separate Markovian bisimulation equivalence $\sim_{\mathrm{MB,s}}$ is the union of all the separate Markovian bisimulations.

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M}}$:

$$\boxed{P_1 \sim_{\mathrm{MB,s}} P_2 \iff P_1 \sim_{\mathrm{MB}} P_2}$$

Equivalence Properties

- $\sim_{\mathrm{MB}}$ is a congruence with respect to all the dynamic and static operators as well as recursion.

- Let $P_1, P_2 \in \mathbb{P}_{\mathrm{M}}$. Whenever $P_1 \sim_{\mathrm{MB}} P_2$, then:

$$\boxed{\begin{array}{ccc}
\multicolumn{3}{c}{<a, \tilde{\lambda}>.P_1 \ \sim_{\mathrm{MB}} \ <a, \tilde{\lambda}>.P_2} \\
P_1 + P \ \sim_{\mathrm{MB}} \ P_2 + P & & P + P_1 \ \sim_{\mathrm{MB}} \ P + P_2 \\
P_1 \parallel_S P \ \sim_{\mathrm{MB}} \ P_2 \parallel_S P & & P \parallel_S P_1 \ \sim_{\mathrm{MB}} \ P \parallel_S P_2 \\
\multicolumn{3}{c}{P_1 \,/\, H \ \sim_{\mathrm{MB}} \ P_2 \,/\, H} \\
\multicolumn{3}{c}{P_1 \setminus L \ \sim_{\mathrm{MB}} \ P_2 \setminus L} \\
\multicolumn{3}{c}{P_1[\varphi] \ \sim_{\mathrm{MB}} \ P_2[\varphi]}
\end{array}}$$

- Recursion: extend $\sim_{\mathrm{MB}}$ to open process terms by replacing all variables freely occurring outside rec binders with every closed process term.

- Let $P_1, P_2 \in \mathcal{PL}_{\mathrm{M}}$ be guarded process terms containing free occurrences of $k \in \mathbb{N}$ process variables $X_1, \ldots, X_k \in \mathit{Var}$ at most.

- We define $P_1 \sim_{\mathrm{MB}} P_2$ iff:

$$\boxed{P_1\{Q_i \hookrightarrow X_i \mid 1 \leq i \leq k\} \ \sim_{\mathrm{MB}} \ P_2\{Q_i \hookrightarrow X_i \mid 1 \leq i \leq k\}}$$

  for all $Q_1, \ldots, Q_k \in \mathbb{P}_{\mathrm{M}}$:

- Whenever $P_1 \sim_{\mathrm{MB}} P_2$, then:

$$\boxed{\mathrm{rec}\, X : P_1 \ \sim_{\mathrm{MB}} \ \mathrm{rec}\, X : P_2}$$

- $\sim_{\rm MB}$ has a sound and complete axiomatization over the set $\mathbb{P}_{\rm M,nrec}$ of nonrecursive process terms of $\mathbb{P}_{\rm M}$.

- Basic laws (commutativity, associativity, and neutral element of +):

$$\boxed{\begin{array}{lrcl} (\mathcal{X}_{\rm MB,1}) & P_1 + P_2 & = & P_2 + P_1 \\ (\mathcal{X}_{\rm MB,2}) & (P_1 + P_2) + P_3 & = & P_1 + (P_2 + P_3) \\ (\mathcal{X}_{\rm MB,3}) & P + \underline{0} & = & P \end{array}}$$

- Characterizing laws (race policy and preselection policy, instead of + idempotency):

$$\boxed{\begin{array}{lrcl} (\mathcal{X}_{\rm MB,4}) & <a, \lambda_1>.P + <a, \lambda_2>.P & = & <a, \lambda_1 + \lambda_2>.P \\ (\mathcal{X}_{\rm MB,5}) & <a, *_{w_1}>.P + <a, *_{w_2}>.P & = & <a, *_{w_1 + w_2}>.P \end{array}}$$

- Expansion law (interl. view of conc. supported by mem. prop.; I, J nonempty and finite):

$$(\mathcal{X}_{\mathrm{MB},6}) \quad \sum_{i \in I} {<} a_i, \tilde{\lambda}_i {>} . P_i \parallel_S \sum_{j \in J} {<} b_j, \tilde{\mu}_j {>} . Q_j \quad = $$

$$\sum_{k \in I, a_k \notin S} {<} a_k, \tilde{\lambda}_k {>} . \left(P_k \parallel_S \sum_{j \in J} {<} b_j, \tilde{\mu}_j {>} . Q_j \right) \; +$$

$$\sum_{h \in J, b_h \notin S} {<} b_h, \tilde{\mu}_h {>} . \left(\sum_{i \in I} {<} a_i, \tilde{\lambda}_i {>} . P_i \parallel_S Q_h \right) \; +$$

$$\sum_{k \in I, a_k \in S, \tilde{\lambda}_k \in \mathbb{R}_{>0}} \; \sum_{h \in J, b_h = a_k, \tilde{\mu}_h = *_{w_h}} {<} a_k, \tilde{\lambda}_k \cdot \frac{w_h}{weight(Q, b_h)} {>} . (P_k \parallel_S Q_h) \; +$$

$$\sum_{h \in J, b_h \in S, \tilde{\mu}_h \in \mathbb{R}_{>0}} \; \sum_{k \in I, a_k = b_h, \tilde{\lambda}_k = *_{v_k}} {<} b_h, \tilde{\mu}_h \cdot \frac{v_k}{weight(P, a_k)} {>} . (P_k \parallel_S Q_h) \; +$$

$$\sum_{k \in I, a_k \in S, \tilde{\lambda}_k = *_{v_k}} \; \sum_{h \in J, b_h = a_k, \tilde{\mu}_h = *_{w_h}} {<} a_k, *_{norm(v_k, w_h, a_k, P, Q)} {>} . (P_k \parallel_S Q_h)$$

$$(\mathcal{X}_{\mathrm{MB},7}) \quad \sum_{i \in I} {<} a_i, \tilde{\lambda}_i {>} . P_i \parallel_S \underline{0} \quad = \quad \sum_{k \in I, a_k \notin S} {<} a_k, \tilde{\lambda}_k {>} . P_k$$

$$(\mathcal{X}_{\mathrm{MB},8}) \quad \underline{0} \parallel_S \sum_{j \in J} {<} b_j, \tilde{\mu}_j {>} . Q_j \quad = \quad \sum_{h \in J, b_h \notin S} {<} b_h, \tilde{\mu}_h {>} . Q_h$$

$$(\mathcal{X}_{\mathrm{MB},9}) \quad \underline{0} \parallel_S \underline{0} \quad = \quad \underline{0}$$

- Distribution laws (for unary static operators):

$$
\begin{array}{|llcll|}
\hline
(\mathcal{X}_{\mathrm{MB},10}) & \underline{0} \,/\, H & = & \underline{0} & \\
(\mathcal{X}_{\mathrm{MB},11}) & (<a, \tilde{\lambda}>.P) \,/\, H & = & <\tau, \tilde{\lambda}>.(P \,/\, H) & \text{if } a \in H \\
(\mathcal{X}_{\mathrm{MB},12}) & (<a, \tilde{\lambda}>.P) \,/\, H & = & <a, \tilde{\lambda}>.(P \,/\, H) & \text{if } a \notin H \\
(\mathcal{X}_{\mathrm{MB},13}) & (P_1 + P_2) \,/\, H & = & P_1 \,/\, H + P_2 \,/\, H & \\
 & & & & \\
(\mathcal{X}_{\mathrm{MB},14}) & \underline{0} \setminus L & = & \underline{0} & \\
(\mathcal{X}_{\mathrm{MB},15}) & (<a, \tilde{\lambda}>.P) \setminus L & = & \underline{0} & \text{if } a \in L \\
(\mathcal{X}_{\mathrm{MB},16}) & (<a, \tilde{\lambda}>.P) \setminus L & = & <a, \tilde{\lambda}>.(P \setminus L) & \text{if } a \notin L \\
(\mathcal{X}_{\mathrm{MB},17}) & (P_1 + P_2) \setminus L & = & P_1 \setminus L + P_2 \setminus L & \\
 & & & & \\
(\mathcal{X}_{\mathrm{MB},18}) & \underline{0}[\varphi] & = & \underline{0} & \\
(\mathcal{X}_{\mathrm{MB},19}) & (<a, \tilde{\lambda}>.P)[\varphi] & = & <\varphi(a), \tilde{\lambda}>.(P[\varphi]) & \\
(\mathcal{X}_{\mathrm{MB},20}) & (P_1 + P_2)[\varphi] & = & P_1[\varphi] + P_2[\varphi] & \\
\hline
\end{array}
$$

- $DED(\mathcal{X}_{\mathrm{MB}})$: deduction system based on all the previous axioms plus:
  - ⊙ Reflexivity: $\mathcal{X}_{\mathrm{MB}} \vdash P = P$.
  - ⊙ Symmetry: $\mathcal{X}_{\mathrm{MB}} \vdash P_1 = P_2 \implies \mathcal{X}_{\mathrm{MB}} \vdash P_2 = P_1$.
  - ⊙ Transitivity: $\mathcal{X}_{\mathrm{MB}} \vdash P_1 = P_2 \wedge \mathcal{X}_{\mathrm{MB}} \vdash P_2 = P_3 \implies \mathcal{X}_{\mathrm{MB}} \vdash P_1 = P_3$.
  - ⊙ Substitutivity: $\mathcal{X}_{\mathrm{MB}} \vdash P_1 = P_2 \implies \mathcal{X}_{\mathrm{MB}} \vdash <a, \tilde{\lambda}>.P_1 = <a, \tilde{\lambda}>.P_2 \wedge \dots$
- The deduction system $DED(\mathcal{X}_{\mathrm{MB}})$ is sound and complete for $\sim_{\mathrm{MB}}$ over $\mathbb{P}_{\mathrm{M,nrec}}$; i.e., for all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,nrec}}$:

$$\boxed{\mathcal{X}_{\mathrm{MB}} \vdash P_1 = P_2 \iff P_1 \sim_{\mathrm{MB}} P_2}$$

- $\sim_{\mathrm{MB}}$ has a modal logic characterization based on a variant of the Hennessy-Milner logic.

- Basic truth values and propositional connectives, plus modal operators expressing how to behave after executing actions with certain names.

- Diamond operator decorated with a lower bound on the rate/weight with which exponentially timed/passive actions with the given name should be executed (consistent with capturing step-by-step behavior mimicking).

- Syntax of the modal language $\mathcal{ML}_{\mathrm{MB}}$ ($a \in \mathit{Name}$, $\lambda, w \in \mathbb{R}_{>0}$):

ϕ	$::=$	true	basic truth value
	$\mid$	$\neg\phi$	negation
	$\mid$	$\phi \wedge \phi$	conjunction
	$\mid$	$\langle a \rangle_{\lambda} \phi$	exponentially timed possibility
	$\mid$	$\langle a \rangle_{*_w} \phi$	passive possibility

- Interpretation of $\mathcal{ML}_{\mathrm{MB}}$ over $\mathbb{P}_{\mathrm{M}}$:

$$\boxed{\begin{array}{lllll} P & \models_{\mathrm{MB}} & \mathrm{true} & & \\ P & \models_{\mathrm{MB}} & \neg\phi & & \text{if } P \not\models_{\mathrm{MB}} \phi \\ P & \models_{\mathrm{MB}} & \phi_1 \wedge \phi_2 & & \text{if } P \models_{\mathrm{MB}} \phi_1 \text{ and } P \models_{\mathrm{MB}} \phi_2 \\ P & \models_{\mathrm{MB}} & \langle a \rangle_\lambda \phi & & \text{if } \mathit{rate}_{\mathrm{e}}(P, a, 0, \mathit{sat}(\phi)) \geq \lambda \\ P & \models_{\mathrm{MB}} & \langle a \rangle_{*w} \phi & & \text{if } \mathit{rate}_{\mathrm{e}}(P, a, -1, \mathit{sat}(\phi)) \geq w \end{array}}$$

where:

$$\mathit{sat}(\phi) \;=\; \{P' \in \mathbb{P}_{\mathrm{M}} \mid P' \models_{\mathrm{MB}} \phi\}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M}}$:

$$\boxed{P_1 \sim_{\mathrm{MB}} P_2 \iff (\forall \phi \in \mathcal{ML}_{\mathrm{MB}}.\, P_1 \models_{\mathrm{MB}} \phi \iff P_2 \models_{\mathrm{MB}} \phi)}$$

- $\sim_{\mathrm{MB}}$ is decidable in polynomial time over the set $\mathbb{P}_{\mathrm{M,fin}}$ of finite-state process terms of $\mathbb{P}_{\mathrm{M}}$: Paige-Tarjan partition refinement algorithm.

- Based on the fact that $\sim_{\mathrm{MB}}$ can be characterized as the limit of a sequence of successively finer equivalence relations:

$$\boxed{\sim_{\mathrm{MB}} \; = \; \bigcap_{i \in \mathbb{N}} \sim_{\mathrm{MB},i}}$$

- $\sim_{\mathrm{MB},0} \; = \; \mathbb{P}_{\mathrm{M}} \times \mathbb{P}_{\mathrm{M}}$ hence it induces the trivial partition $\{\mathbb{P}_{\mathrm{M}}\}$.

- Whenever $P_1 \sim_{\mathrm{MB},i} P_2$, $i \in \mathbb{N}_{\geq 1}$, then for all $a \in \mathit{Name}$, $l \in \{0, -1\}$, and $D \in \mathbb{P}_{\mathrm{M}} / \sim_{\mathrm{MB},i-1}$:

$$\mathit{rate}_{\mathrm{e}}(P_1, a, l, D) \; = \; \mathit{rate}_{\mathrm{e}}(P_2, a, l, D)$$

- $\sim_{\mathrm{MB},1}$ refines $\{\mathbb{P}_{\mathrm{M}}\}$ by creating an equivalence class for each set of process terms that satisfy the necessary condition for $\sim_{\mathrm{MB}}$.

- Steps of the algorithm for checking whether $P_1 \sim_{\mathrm{MB}} P_2$:
  1. Build an initial partition with a single class including all the states of $[\![P_1]\!]_{\mathrm{M}}$ and $[\![P_2]\!]_{\mathrm{M}}$.
  2. Initialize a list of splitters with the above class as its only element.
  3. While the list of splitters is not empty, select a splitter and remove it from the list after refining the current partition for each $a \in \mathit{Name}_{P_1,P_2}$ and $l \in \{0, -1\}$:
     a. Split each class of the current partition by comparing the exit rates of its states when performing actions of name a and level l that lead to the selected splitter.
     b. For each class that has been split, insert into the list of splitters all the resulting subclasses except for the largest one.
  4. Return yes/no depending on whether the initial states of $[\![P_1]\!]_{\mathrm{M}}$ and $[\![P_2]\!]_{\mathrm{M}}$ belong to the same class of the final partition or not.
- The time complexity is $O(m \cdot \log n)$ if a *splay tree* is used for representing the subclasses arising from the splitting of a class (they can be more than two).

- $\sim_{\mathrm{MB}}$ induces an exact aggregation known as ordinary lumping.

- A partition $\mathcal{O}$ of the state space of a CTMC is an ordinary lumping iff, whenever $s_1, s_2 \in O$ for some $O \in \mathcal{O}$, then for all $O' \in \mathcal{O}$:

$$\boxed{\begin{array}{l} \sum\{\!|\, \lambda \in \mathbb{R}_{>0} \mid \exists s' \in O'.\, s_1 \xrightarrow{\lambda} s' \,|\!\} \ = \\ \qquad\qquad\qquad \sum\{\!|\, \lambda \in \mathbb{R}_{>0} \mid \exists s' \in O'.\, s_2 \xrightarrow{\lambda} s' \,|\!\} \end{array}}$$

- The probability of being in a macrostate of an ordinarily lumped CTMC is the sum of the probabilities of being in one of its constituent microstates of the original CTMC.

- Two Markovian bisimilar process terms in $\mathbb{P}_{\mathrm{M,pc}}$ are guaranteed to possess the same performance characteristics.

Part IV:
Variants of Markovian Bisimilarity

Markovian Bisimilarity and Rewards

- Specific performance measures may distinguish between ordinarily lumpable states by ascribing them a different meaning.
- Make $\sim_{\mathrm{MB}}$ sensitive to performance measures.
- Instant-of-time measures vs. interval-of-time (or cumulative) measures, which refer to stationary or transient behavior.
- Specification of these performance measures for CTMC-based models through reward structures:
  - ⊙ The yield reward $yr \in \mathbb{R}$ associated with a state expresses the rate at which a gain/loss is accumulated while sojourning in that state.
  - ⊙ The bonus reward $br \in \mathbb{R}$ associated with a transition expresses the instantaneous gain/loss implied by the execution of that transition.

- Value of an instant-of-time performance measure expressed via rewards:

$$\boxed{\sum_{s \in S} yr(s) \cdot \pi[s] + \sum_{s \xrightarrow{\lambda} s'} br(s, \lambda, s') \cdot \phi(s, \lambda, s')}$$

  where:

  - $yr(s)$ is the yield reward associated with state s.
  - $\pi[s]$ is the probability of being in state s at the considered instant of time.
  - $br(s, \lambda, s')$ is the bonus reward associated with transition $s \xrightarrow{\lambda} s'$.
  - $\phi(s, \lambda, s')$ is the frequency of transition $s \xrightarrow{\lambda} s'$ at the considered instant of time: $\phi(s, \lambda, s') = \pi[s] \cdot \lambda$.

- *Ascribing a different meaning to ordinarily lumpable states amounts to giving different rewards to such states or their outgoing transitions.*
- How to specify rewards at the process algebraic level?
- Bonus rewards can naturally be associated with actions.
- Yield rewards are problematic, as in process calculi the concept of state is implicit.
- Process calculi are action-based, hence the idea of associating yield rewards with actions too. What if a term enables several actions?
- *Additivity assumption*: the yield reward of a state corresponding to a process term is given by the sum of the yield rewards associated with the actions enabled by that term.

- $\mathrm{MPC_r}$: Markovian process calculus with rewards.
- New action syntax $(yr,\, br \in \mathbb{R})$:

$$\boxed{\begin{array}{c} <a, \lambda, yr, br> \\ <a, *_w, *, *> \end{array}}$$

- $\mathbb{P}_{\mathrm{M,r}}$: set of closed and guarded process terms.
- New semantic rule for action prefix:

$$\boxed{<a, \tilde{\lambda}, \tilde{yr}, \tilde{br}>.P \xrightarrow{a, \tilde{\lambda}, \tilde{yr}, \tilde{br}}_{\mathrm{M,r}} P}$$

- The other semantic rules are modified accordingly.

- In particular, here are the semantic rules for synchronization:

$$\frac{P_1 \xrightarrow{a,\lambda,yr,br}_{\mathrm{M,r}} P_1' \qquad P_2 \xrightarrow{a,*_w,*,*}_{\mathrm{M,r}} P_2' \qquad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a,\lambda\cdot\frac{w}{weight(P_2,a)},yr\cdot\frac{w}{weight(P_2,a)},br}_{\mathrm{M,r}} P_1' \parallel_S P_2'}$$

$$\frac{P_1 \xrightarrow{a,*_w,*,*}_{\mathrm{M,r}} P_1' \qquad P_2 \xrightarrow{a,\lambda,yr,br}_{\mathrm{M,r}} P_2' \qquad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a,\lambda\cdot\frac{w}{weight(P_1,a)},yr\cdot\frac{w}{weight(P_1,a)},br}_{\mathrm{M,r}} P_1' \parallel_S P_2'}$$

$$\frac{P_1 \xrightarrow{a,*_{w_1},*,*}_{\mathrm{M,r}} P_1' \qquad P_2 \xrightarrow{a,*_{w_2},*,*}_{\mathrm{M,r}} P_2' \qquad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a,*_{norm(w_1,w_2,a,P_1,P_2)},*,*}_{\mathrm{M,r}} P_1' \parallel_S P_2'}$$

- Yield rewards normalized in the same way as rates.

- Exit reward with which $P \in \mathbb{P}_{\mathrm{M,r}}$ executes actions of name $a \in \mathit{Name}$ and level $l \in \{0, -1\}$ that lead to destination $D \subseteq \mathbb{P}_{\mathrm{M,r}}$:

$$\boxed{\begin{array}{rcl} \mathit{reward}_{\mathrm{e}}(P, a, 0, D) & = & \sum \{\!| \, yr + \lambda \cdot br \in \mathbb{R} \mid \exists P' \in D. \, P \xrightarrow{a, \lambda, yr, br}_{\mathrm{M,r}} P' \, |\!\} \\ \mathit{reward}_{\mathrm{e}}(P, a, -1, D) & = & 0 \end{array}}$$

- An equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M,r}}$ is a reward Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \mathit{Name}$, levels $l \in \{0, -1\}$, and equivalence classes $D \in \mathbb{P}_{\mathrm{M,r}} / \mathcal{B}$:

$$\boxed{\begin{array}{rcl} \mathit{rate}_{\mathrm{e}}(P_1, a, l, D) & = & \mathit{rate}_{\mathrm{e}}(P_2, a, l, D) \\ \mathit{reward}_{\mathrm{e}}(P_1, a, l, D) & = & \mathit{reward}_{\mathrm{e}}(P_2, a, l, D) \end{array}}$$

- Reward Markovian bisimulation equivalence $\sim_{\mathrm{MB,r}}$ is the union of all the reward Markovian bisimulations.

- $\sim_{\mathrm{MB,r}}$ enjoys the same properties as $\sim_{\mathrm{MB}}$.

- Axioms characterizing $\sim_{\mathrm{MB,r}}$:

$$\boxed{\begin{array}{rcl} <a, \lambda_1, yr_1, br_1>.P + <a, \lambda_2, yr_2, br_2>.P & = & \\ \multicolumn{3}{r}{<a, \lambda_1 + \lambda_2, yr_1 + yr_2, \frac{\lambda_1}{\lambda_1+\lambda_2} \cdot br_1 + \frac{\lambda_2}{\lambda_1+\lambda_2} \cdot br_2>.P} \\ <a, *_{w_1}, *, *>.P + <a, *_{w_2}, *, *>.P & = & <a, *_{w_1+w_2}, *, *>.P \end{array}}$$

- Yield rewards summed up in the same way as rates (additivity assumption).

- Weighted sum of bonus rewards by considering execution probabilities.

- Equivalent characterizing axioms in yield-normal-form:

$$
\begin{aligned}
<a, \lambda, yr, br>.P &= <a, \lambda, yr + \lambda \cdot br, 0>.P \\
<a, \lambda_1, yr_1, 0>.P + <a, \lambda_2, yr_2, 0>.P &= <a, \lambda_1 + \lambda_2, yr_1 + yr_2, 0>.P \\
<a, *_{w_1}, *, *>.P + <a, *_{w_2}, *, *>.P &= <a, *_{w_1 + w_2}, *, *>.P
\end{aligned}
$$

- Equivalent characterizing axioms in bonus-normal-form:

$$
\begin{aligned}
<a, \lambda, yr, br>.P &= <a, \lambda, 0, br + \tfrac{yr}{\lambda}>.P \\
<a, \lambda_1, 0, br_1>.P + <a, \lambda_2, 0, br_2>.P &= \\
&\quad <a, \lambda_1 + \lambda_2, 0, \tfrac{\lambda_1}{\lambda_1+\lambda_2} \cdot br_1 + \tfrac{\lambda_2}{\lambda_1+\lambda_2} \cdot br_2>.P \\
<a, *_{w_1}, *, *>.P + <a, *_{w_2}, *, *>.P &= <a, *_{w_1 + w_2}, *, *>.P
\end{aligned}
$$

Markovian Bisimilarity and Nondeterminism

- Nondeterminism is a useful abstraction whenever not all the details of a model are known in the early design stages.

- Combine nondeterministic process calculi and CTMCs.

- Separate exponential delays from interacting actions (orthogonal time).

- Markovian branchings and nondeterministic branchings.

- Interprocess communication implemented through the synchronization of visible interacting actions.

- Interacting actions take no time.

- *Maximal progress: τ-actions take precedence over time passing.*

- $\mathrm{MPC_i}$: interactive Markovian process calculus.
- New syntax for prefixing:

$$\boxed{\begin{array}{c} a \,.\, P \\[1ex] (\lambda) \,.\, P \end{array}}$$

- $\mathbb{P}_{\mathrm{M,i}}$: set of closed and guarded process terms.
- Two transition relations are necessary: one for actions, one for delays.
- New semantic rules for prefixing:

$$\boxed{\begin{array}{c} a \,.\, P \xrightarrow{a}_{\mathrm{M,i,a}} P \\[1ex] (\lambda) \,.\, P \xrightarrow{\lambda}_{\mathrm{M,i,d}} P \end{array}}$$

- The other semantic rules are modified accordingly.

- In particular, here are the semantic rules for parallel composition:

$$\frac{P_1 \xrightarrow{a}_{\mathrm{M,i,a}} P_1' \qquad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{a}_{\mathrm{M,i,a}} P_1' \parallel_S P_2} \qquad \frac{P_2 \xrightarrow{a}_{\mathrm{M,i,a}} P_2' \qquad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{a}_{\mathrm{M,i,a}} P_1 \parallel_S P_2'}$$

$$\frac{P_1 \xrightarrow{a}_{\mathrm{M,i,a}} P_1' \qquad P_2 \xrightarrow{a}_{\mathrm{M,i,a}} P_2' \qquad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a}_{\mathrm{M,i,a}} P_1' \parallel_S P_2'}$$

$$\frac{P_1 \xrightarrow{\lambda}_{\mathrm{M,i,d}} P_1'}{P_1 \parallel_S P_2 \xrightarrow{\lambda}_{\mathrm{M,i,d}} P_1' \parallel_S P_2} \qquad \frac{P_2 \xrightarrow{\lambda}_{\mathrm{M,i,d}} P_2'}{P_1 \parallel_S P_2 \xrightarrow{\lambda}_{\mathrm{M,i,d}} P_1 \parallel_S P_2'}$$

- CMTC derivation by superposing source and destination state of each action transition, if there are no states with several outgoing action transitions or a non-τ-action transition alternative to delay transitions (i.e., only in the absence of nondeterminism).

- Exit rate of $P \in \mathbb{P}_{\mathrm{M,i}}$ with destination $D \subseteq \mathbb{P}_{\mathrm{M,i}}$:

$$\boxed{\mathit{rate}_{\mathrm{e,d}}(P, D) \; = \; \sum \{\!| \, \lambda \in \mathbb{R}_{>0} \mid \exists P' \in D.\, P \xrightarrow{\lambda}_{\mathrm{M,i,d}} P' \, |\!\}}$$

- An equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M,i}}$ is an interactive Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then:
  - For all action names $a \in \mathit{Name}$, $P_1 \xrightarrow{a}_{\mathrm{M,i,a}} P_1'$ implies $P_2 \xrightarrow{a}_{\mathrm{M,i,a}} P_2'$ for some P_2' with $(P_1', P_2') \in \mathcal{B}$.
  - For all equivalence classes $D \in \mathbb{P}_{\mathrm{M,i}}/\mathcal{B}$, $P_1 \stackrel{\tau}{\not\longrightarrow}_{\mathrm{M,i,a}}$ implies $P_2 \stackrel{\tau}{\not\longrightarrow}_{\mathrm{M,i,a}}$ with:

$$\boxed{\mathit{rate}_{\mathrm{e,d}}(P_1, D) \; = \; \mathit{rate}_{\mathrm{e,d}}(P_2, D)}$$

- Interactive Markovian bisimulation equivalence $\sim_{\mathrm{MB,i}}$ is the union of all the interactive Markovian bisimulations.

- $\sim_{\mathrm{MB,i}}$ enjoys properties similar to those of $\sim_{\mathrm{MB}}$.
- Axioms characterizing $\sim_{\mathrm{MB,i}}$:

$$\boxed{\begin{array}{rcl} a\,.\,P + a\,.\,P & = & a\,.\,P \\ (\lambda_1)\,.\,P + (\lambda_2)\,.\,P & = & (\lambda_1 + \lambda_2)\,.\,P \\ \tau\,.\,P + (\lambda)\,.\,Q & = & \tau\,.\,P \end{array}}$$

- Idempotency of $+$ like in nondeterministic process calculi.
- Race policy like in CTMCs.
- Maximal progress too.

- τ-actions should be ignored when playing the bisimulation game.
- They are invisible and take no time.
- Need for a weak variant of $\sim_{\mathrm{MB,i}}$.
- After any nonpreemptable exponential delay, skip all the states that can evolve via a finite sequence of τ-transitions to a given class.
- Internal backward closure of $D \subseteq \mathbb{P}_{\mathrm{M,i}}$:

$$\boxed{D_\tau \; = \; \{P' \in \mathbb{P}_{\mathrm{M,i}} \mid \exists P \in D.\, P' \xRightarrow{\tau^*}_{\mathrm{M,i,a}} P\}}$$

where $Q \xRightarrow{a_1 \ldots a_n}_{\mathrm{M,i,a}} Q'$ iff:

  - either $n = 0$ and $Q \equiv Q'$, meaning that Q stays idle;
  - or $n \in \mathbb{N}_{\geq 1}$ and there exist $Q_0, Q_1, \ldots, Q_n \in \mathbb{P}_{\mathrm{M,i}}$ such that:
    - $Q \equiv Q_0$;
    - $Q_{i-1} \xrightarrow{a_i}_{\mathrm{M,i,a}} Q_i$ for all $1 \leq i \leq n$;
    - $Q_n \equiv Q'$.

- An equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M,i}}$ is a weak interactive Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then:
  - For all visible action names $a \in Name_{\mathrm{v}}$, $P_1 \xrightarrow{a}_{\mathrm{M,i,a}} P_1'$ implies $P_2 \xRightarrow{\tau^* a \tau^*}_{\mathrm{M,i,a}} P_2'$ for some P_2' with $(P_1', P_2') \in \mathcal{B}$.
  - $P_1 \xrightarrow{\tau}_{\mathrm{M,i,a}} P_1'$ implies $P_2 \xRightarrow{\tau^*}_{\mathrm{M,i,a}} P_2'$ for some P_2' with $(P_1', P_2') \in \mathcal{B}$.
  - For all equivalence classes $D \in \mathbb{P}_{\mathrm{M,i}}/\mathcal{B}$, $P_1 \xRightarrow{\tau^*}_{\mathrm{M,i,a}} P_1' \stackrel{\tau}{\not\longrightarrow}_{\mathrm{M,i,a}}$ implies $P_2 \xRightarrow{\tau^*}_{\mathrm{M,i,a}} P_2' \stackrel{\tau}{\not\longrightarrow}_{\mathrm{M,i,a}}$ for some P_2' with:

$$\boxed{rate_{\mathrm{e,d}}(P_1', D_\tau) \; = \; rate_{\mathrm{e,d}}(P_2', D_\tau)}$$

- Weak interactive Markovian bisimulation equivalence $\approx_{\mathrm{MB,i}}$ is the union of all the weak interactive Markovian bisimulations.
- The first part of the third clause distinguishes τ-divergent processes from non-τ-divergent processes that cannot execute any visible action (otherwise congruence w.r.t. par. comp. would be lost in a timed setting with maximal progress).

- $\approx_{\mathrm{MB,i}}$ is strictly coarser than $\sim_{\mathrm{MB,i}}$ but it is not a congruence with respect to alternative composition.

- Initial τ-actions need a different treatment (as for classical weak bisimilarity).

- $P_1 \in \mathbb{P}_{\mathrm{M,i}}$ is weakly interactive Markovian bisimulation congruent to $P_2 \in \mathbb{P}_{\mathrm{M,i}}$, written $P_1 \simeq_{\mathrm{MB,i}} P_2$, iff:

  - For all action names $a \in \mathit{Name}$, $P_1 \xrightarrow{a}_{\mathrm{M,i,a}} P_1'$ implies $P_2 \xRightarrow{\tau^* a \tau^*}_{\mathrm{M,i,a}} P_2'$ for some P_2' with $P_1' \approx_{\mathrm{MB,i}} P_2'$.

  - For all action names $a \in \mathit{Name}$, $P_2 \xrightarrow{a}_{\mathrm{M,i,a}} P_2'$ implies $P_1 \xRightarrow{\tau^* a \tau^*}_{\mathrm{M,i,a}} P_1'$ for some P_1' with $P_1' \approx_{\mathrm{MB,i}} P_2'$.

  - $P_1 \not\xrightarrow{\tau}_{\mathrm{M,i,a}}$ iff $P_2 \not\xrightarrow{\tau}_{\mathrm{M,i,a}}$.

  - For all equivalence classes $D \in \mathbb{P}_{\mathrm{M,i}} / \approx_{\mathrm{MB,i}}$, $P_1 \not\xrightarrow{\tau}_{\mathrm{M,i,a}}$ implies:

$$\boxed{\mathit{rate}_{\mathrm{e,d}}(P_1, D) \; = \; \mathit{rate}_{\mathrm{e,d}}(P_2, D)}$$

- $\sim_{\rm MB,i} \subset \simeq_{\rm MB,i} \subset \approx_{\rm MB,i}$ with $\simeq_{\rm MB,i}$ having the same properties as $\sim_{\rm MB,i}$.

- Additional axioms characterizing $\simeq_{\rm MB,i}$:

$$\boxed{\begin{array}{rcl} a\,.\,\tau\,.\,P & = & a\,.\,P \\ P + \tau\,.\,P & = & \tau\,.\,P \\ a\,.\,(P + \tau\,.\,Q) + \tau\,.\,Q & = & a\,.\,(P + \tau\,.\,Q) \\ (\lambda)\,.\,\tau\,.\,P & = & (\lambda)\,.\,P \end{array}}$$

- τ-laws witnessing the capability of $\simeq_{\rm MB,i}$ of abstracting from τ-actions that are not initial.

Markovian Bisimilarity and Immediate Actions

- Combinations of exponential distributions approximate many general distributions arbitrarily closely, still some useful durations cannot be represented in the integrated time case, specially zero durations.
- Performance abstraction mechanism for integrated time.
- Useful for handling activities that are several orders of magnitude faster than those important for certain performance measures.
- Necessary to manage situations with which no timing can be associated like choices among logical events (e.g., the reception of a message vs. its loss).
- Zero durations implemented through immediate actions à la GSPN.
- Markovian branchings and prioritized/probabilistic branchings.
- *Preemption: immediate τ-actions take precedence over all the lower priority actions.*

- MPC$_x$: Markovian process calculus extended with immediate actions.

- New action syntax ($l \in \mathbb{N}_{>0}$, $l' \in \mathbb{N}$):

$$\boxed{\begin{array}{c} <a, \lambda>.P \\ <a, \infty_{l,w}>.P \\ <a, *_{l',w}>.P \end{array}}$$

- $\mathbb{P}_{M,x}$: set of closed and guarded process terms.

- *Preselection policy*: each of the highest priority enabled immediate actions is given an execution probability proportional to its weight.

- *Priority constraints* to control process priority interrelation (congruence):
  - ⊙ An exponentially timed action can synchronize only with a passive action with priority constraint $l' = 0$.
  - ⊙ An immediate action with priority level l can synchronize only with a passive action with priority constraint $l' = l$.

- Additional semantic rules for immediate actions:

$$<a, \infty_{l,w}>.P \xrightarrow{a,\infty_{l,w}}_{\mathrm{M,x}} P$$

$$\frac{P_1 \xrightarrow{a,\infty_{l,w}}_{\mathrm{M,x}} P_1' \qquad P_2 \xrightarrow{a,*_{l,v}}_{\mathrm{M,x}} P_2' \qquad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a,\infty_{l,w \cdot \frac{v}{weight(P_2,a,l)}}}_{\mathrm{M,x}} P_1' \parallel_S P_2'}$$

$$\frac{P_1 \xrightarrow{a,*_{l,v}}_{\mathrm{M,x}} P_1' \qquad P_2 \xrightarrow{a,\infty_{l,w}}_{\mathrm{M,x}} P_2' \qquad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a,\infty_{l,w \cdot \frac{v}{weight(P_1,a,l)}}}_{\mathrm{M,x}} P_1' \parallel_S P_2'}$$

- The other semantic rules are modified by taking into account priority constraints associated with passive actions.

- $\mathbb{P}_{\mathrm{M,x,pc}}$: set of performance closed process terms of $\mathbb{P}_{\mathrm{M,x}}$.

- CTMC derivation for performance closed process terms by eliminating all *vanishing states* (those with outgoing immediate transitions, hence zero sojourn time):

  ⊙ Make as many copies of every transition entering a vanishing state as there are highest priority immediate transitions departing from the vanishing state (e.g., λ entering $w_1, \ldots, w_n$).

  ⊙ Connect each copy to the destination state of one of the highest priority immediate transitions leaving the vanishing state.

  ⊙ Assign to each copy the rate of the original incoming transition multiplied by the exec. probability of the highest priority immediate transition corresponding to the copy ($\lambda \cdot w_1 / \sum_i w_i, \ldots, \lambda \cdot w_n / \sum_i w_i$).

- Exit rate at which $P \in \mathbb{P}_{\mathrm{M,x}}$ executes actions of name $a \in \mathit{Name}$ and level $l \in \mathbb{Z}$ that lead to destination $D \subseteq \mathbb{P}_{\mathrm{M,x}}$:

$$\boxed{\mathit{rate}_{\mathrm{e}}(P, a, l, D) \; = \; \begin{cases} \sum\{\!|\, \lambda \in \mathbb{R}_{>0} \mid \exists P' \in D.\, P \xrightarrow{a,\lambda}_{\mathrm{M,x}} P' \,|\!\} & \text{if } l = 0 \\ \sum\{\!|\, w \in \mathbb{R}_{>0} \mid \exists P' \in D.\, P \xrightarrow{a,\infty_{l,w}}_{\mathrm{M,x}} P' \,|\!\} & \text{if } l > 0 \\ \sum\{\!|\, w \in \mathbb{R}_{>0} \mid \exists P' \in D.\, P \xrightarrow{a,*_{-l-1,w}}_{\mathrm{M,x}} P' \,|\!\} & \text{if } l < 0 \end{cases}}$$

- $\mathit{pri}^{\tau}_{\infty}(P)$: priority level of the highest priority immediate τ-action enabled by $P \in \mathbb{P}_{\mathrm{M,x}}$.

- $\mathit{pri}^{\tau}_{\infty}(P) = 0$ if P does not enable any immediate τ-action.

- $\mathit{no\text{-}pre}(l, P)$ if no action of level $l \in \mathbb{Z}$ can be preempted in P:

$$\boxed{\mathit{no\text{-}pre}(l, P) \iff l \geq \mathit{pri}^{\tau}_{\infty}(P) \vee -l-1 \geq \mathit{pri}^{\tau}_{\infty}(P)}$$

- The exit rate comparison should be conducted only when no preemption can be exercised.

- An equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M,x}}$ is an extended Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \mathit{Name}$, levels $l \in \mathbb{Z}$ such that $\mathit{no\text{-}pre}(l, P_1)$ and $\mathit{no\text{-}pre}(l, P_2)$, and equivalence classes $D \in \mathbb{P}_{\mathrm{M,x}}/\mathcal{B}$:

$$\boxed{\mathit{rate}_{\mathrm{e}}(P_1, a, l, D) \ = \ \mathit{rate}_{\mathrm{e}}(P_2, a, l, D)}$$

- Extended Markovian bisimulation equivalence $\sim_{\mathrm{MB,x}}$ is the union of all the extended Markovian bisimulations.

- $\sim_{\mathrm{MB,x}}$ enjoys the same properties as $\sim_{\mathrm{MB}}$.

- Axioms characterizing $\sim_{\mathrm{MB,x}}$:

$$
\boxed{
\begin{array}{rcll}
<a,\lambda_1>.P + <a,\lambda_2>.P & = & <a,\lambda_1+\lambda_2>.P & \\
<a,\infty_{l,w_1}>.P + <a,\infty_{l,w_2}>.P & = & <a,\infty_{l,w_1+w_2}>.P & \\
<a,*_{l,w_1}>.P + <a,*_{l,w_2}>.P & = & <a,*_{l,w_1+w_2}>.P & \\
& & & \\
<\tau,\infty_{l,w}>.P + <a,\lambda>.Q & = & <\tau,\infty_{l,w}>.P & \\
<\tau,\infty_{l,w}>.P + <a,\infty_{l',w'}>.Q & = & <\tau,\infty_{l,w}>.P & \text{if } l > l' \\
<\tau,\infty_{l,w}>.P + <a,*_{l',w'}>.Q & = & <\tau,\infty_{l,w}>.P & \text{if } l > l'
\end{array}
}
$$

- The first three axioms encode race policy and preselection policy.

- The last three axioms encode preemption exercised by immediate τ-actions over lower priority actions.

- Immediate τ-actions should be ignored in the bisimulation game.
- They are invisible and take no time.
- Need for a weak variant of $\sim_{\mathrm{MB,x}}$.
- After any nonpreemptable action, skip all the states that can evolve via a finite sequence of immediate τ-transitions to a given class.
- Harder than weakening $\sim_{\mathrm{MB,i}}$:
  - Need to keep track of quantitative information associated with the actions to be abstracted away.
  - Need to take into account the degree of observability of the classes of process terms to be reached.

- Process term $P \in \mathbb{P}_{\mathrm{M,x}}$ is l-unobservable, $l \in \mathbb{N}_{>0}$, iff $pri^{\tau}_{\infty}(P) = l$ and P does not enable any immediate visible action with priority level $l' \geq l$, nor any passive action with priority constraint $l' \geq l$.

- A computation c of length $n \in \mathbb{N}_{>0}$:
$$P_1 \xrightarrow{\tau,\infty_{l_1,w_1}}_{\mathrm{M,x}} P_2 \xrightarrow{\tau,\infty_{l_2,w_2}}_{\mathrm{M,x}} \dots \xrightarrow{\tau,\infty_{l_n,w_n}}_{\mathrm{M,x}} P_{n+1}$$
is unobservable iff for all $i = 1, \dots, n$ process term P_i is l_i-unobservable.

- The probability of executing the unobservable computation c is:
$$\boxed{prob(c) \; = \; \prod_{i=1}^{n} \frac{w_i}{rate_{\mathrm{o}}(P_i, \tau, l_i)}}$$

- Weak exit rate at which $P \in \mathbb{P}_{\mathrm{M,x}}$ executes actions of name $a \in \mathit{Name}$ and level $l \in \mathbb{Z}$ that lead to destination $D \subseteq \mathbb{P}_{\mathrm{M,x}}$:

$$\boxed{\mathit{rate}_{\mathrm{e,w}}(P, a, l, D) \; = \sum_{P' \in D_{\mathrm{w}}} \mathit{rate}_{\mathrm{e}}(P, a, l, \{P'\}) \cdot \mathit{prob}_{\mathrm{w}}(P', D)}$$

  where:

  - D_{w} is the weak backward closure of D:

$$\boxed{D_{\mathrm{w}} \; = \; D \cup \{Q \in \mathbb{P}_{\mathrm{M,x}} - D \mid Q \text{ can reach } D \text{ via unobs. computations}\}}$$

  - $\mathit{prob}_{\mathrm{w}}(P', D)$ is the sum of the probabilities of all the unobservable computations from a process term in D_{w} to D:

$$\boxed{\mathit{prob}_{\mathrm{w}}(P', D) \; = \; \begin{cases} 1 & \text{if } P' \in D \\ \sum\{\!|\, \mathit{prob}(c) \mid c \text{ unobs. computation from } P' \text{ to } D \,|\!\} & \\ & \text{if } P' \in D_{\mathrm{w}} - D \end{cases}}$$

- The weak exit rate comparison should be conducted only with respect to certain classes of process terms.
- An observable process term is a term that enables a visible action that cannot be preempted by any enabled immediate τ-action.
- An initially unobservable process term is a term in which all the enabled visible actions are preempted by some enabled immediate τ-action, but at least one of the computations starting at this term with one of the higher priority enabled immediate τ-actions reaches an observable term.
- A fully unobservable process term is a term in which all the enabled visible actions are preempted by some enabled immediate τ-action, and all the computations starting at this term with one of the higher priority enabled immediate τ-actions are unobservable.
- $\mathbb{P}_{\mathrm{M,x,fu}}$: set of fully unobservable process terms of $\mathbb{P}_{\mathrm{M,x}}$.

- The weak exit rate comparison with respect to observable and fully unobservable classes must obviously be performed.
- The comparison should be made with respect to all fully unobservable classes together, in order to maximize the abstraction power despite the quantitative information attached to immediate τ-actions.
- The comparison with respect to initially unobservable classes should be skipped, otherwise process terms like:

$$<a, \lambda>.<\tau, \infty_{l_1, w_1}>.<b, \mu>.\underline{0}$$

$$<a, \lambda>.<\tau, \infty_{l_2, w_2}>.<b, \mu>.\underline{0}$$

$$<a, \lambda>.<b, \mu>.\underline{0}$$

could not be considered equivalent to each other.

- An equivalence relation $\mathcal{B}$ over $\mathbb{P}_{\mathrm{M,x}}$ is a weak extended Markovian bisimulation iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \mathit{Name}$ and levels $l \in \mathbb{Z}$ such that $\mathit{no\text{-}pre}(l, P_1)$ and $\mathit{no\text{-}pre}(l, P_2)$:

$$\boxed{\begin{array}{rcll} \mathit{rate}_{\mathrm{e,w}}(P_1, a, l, D) & = & \mathit{rate}_{\mathrm{e,w}}(P_2, a, l, D) & \forall\, D \in \mathbb{P}_{\mathrm{M,x}}/\mathcal{B} \text{ obs.} \\ \mathit{rate}_{\mathrm{e,w}}(P_1, a, l, \mathbb{P}_{\mathrm{M,x,fu}}) & = & \mathit{rate}_{\mathrm{e,w}}(P_2, a, l, \mathbb{P}_{\mathrm{M,x,fu}}) & \end{array}}$$

- Weak extended Markovian bisimulation equivalence $\approx_{\mathrm{MB,x}}$ is the union of all the weak extended Markovian bisimulations.

- $\approx_{\mathrm{MB,x}}$ enjoys the same properties as $\sim_{\mathrm{MB,x}}$ except for congruence with respect to parallel composition.

- Need to restrict to a well-prioritized subset of $\mathbb{P}_{\mathrm{M,x,nd}}$, the set of nondivergent process terms of $\mathbb{P}_{\mathrm{M,x}}$.

- A fully unobservable process term like $<\tau, \infty_{l,w}>.\underline{0}$ allows concurrent exp. timed actions to be executed, while the equivalent divergent process term $\mathrm{rec}\, X : <\tau, \infty_{l,w}>.X$ prevents time from passing.

- State observability and preemption schemes for two equivalent terms may change differently when composing each of them in parallel with some term, thus exposing parts of their behavior not compared before.

- A set of terms of $\mathbb{P}_{\mathrm{M,x}}$ is well prioritized if, taken two arbitrary terms P_1 and P_2 in the set, any immediate/passive transition of each of $[\![P_1]\!]_{\mathrm{M,x}}$ and $[\![P_2]\!]_{\mathrm{M,x}}$ has priority level/constraint less than the priority level of any highest priority immediate τ-transition departing from an unobservable state of the other one.

- Additional axioms characterizing $\approx_{\rm MB,x}$:

$$<a, \lambda>. \sum_{i \in I} <\tau, \infty_{l,w_i}>.P_i \;=\; \sum_{i \in I} <a, \lambda \cdot w_i / \Sigma_{k \in I}\, w_k>.P_i$$

$$<a, \infty_{l',w'}>. \sum_{i \in I} <\tau, \infty_{l,w_i}>.P_i \;=\; \sum_{i \in I} <a, \infty_{l', w' \cdot w_i / \Sigma_{k \in I}\, w_k}>.P_i$$

$$<a, *_{l',w'}>. \sum_{i \in I} <\tau, \infty_{l,w_i}>.P_i \;=\; \sum_{i \in I} <a, *_{l', w' \cdot w_i / \Sigma_{k \in I}\, w_k}>.P_i$$

- τ, ∞-laws showing the ability of $\approx_{\rm MB,x}$ of abstracting from immediate τ-actions and of encoding the procedure for removing vanishing states.

- No abstraction from initial immediate τ-actions, hence $\approx_{\rm MB,x}$ does not incur the congruence problem with respect to alternative composition found in $\approx_{\rm MB,i}$ (a consequence of the way the weak exit rate is defined).

- **Running example** ($\approx_{\mathrm{MB,x}}$):
  - Pipeline implementation with two communicating one-pos. buffers:

$$\begin{array}{rcl}
PC^{\mathrm{M}}_{\mathrm{pipe},2} & \triangleq & Prod^{\mathrm{M}} \parallel_{\{deposit\}} (LBuff^{\mathrm{M}} \parallel_{\{pass\}} RBuff^{\mathrm{M}}) \,/\, \{pass\} \parallel_{\{withdraw\}} Cons^{\mathrm{M}} \\
Prod^{\mathrm{M}} & \triangleq & <deposit, \lambda>.Prod^{\mathrm{M}} \\
LBuff^{\mathrm{M}} & \triangleq & <deposit, *_{0,1}>.<pass, \infty_{1,1}>.LBuff^{\mathrm{M}} \\
RBuff^{\mathrm{M}} & \triangleq & <pass, *_{1,1}>.<withdraw, *_{0,1}>.RBuff^{\mathrm{M}} \\
Cons^{\mathrm{M}} & \triangleq & <withdraw, \mu>.Cons^{\mathrm{M}}
\end{array}$$

  - Action *pass* is immediate in the left buffer, as it has been assumed that its execution takes a negligible amount of time compared to deposits and withdrawals.
  - Is $PC^{\mathrm{M}}_{\mathrm{pipe},2}$ a correct implementation of $ProdCons^{\mathrm{M}}_{0/2}$?
  - Yes, because it turns out that $PC^{\mathrm{M}}_{\mathrm{pipe},2} \approx_{\mathrm{MB,x}} ProdCons^{\mathrm{M}}_{0/2}$.
  - Proved by finding a suitable weak extended Markovian bisimulation.

- Weak extended Markovian bisimulation for $PC^{\mathrm{M}}_{\mathrm{pipe},2} \approx_{\mathrm{MB,x}} ProdCons^{\mathrm{M}}_{0/2}$, with states of the same color belonging to the same equivalence class:

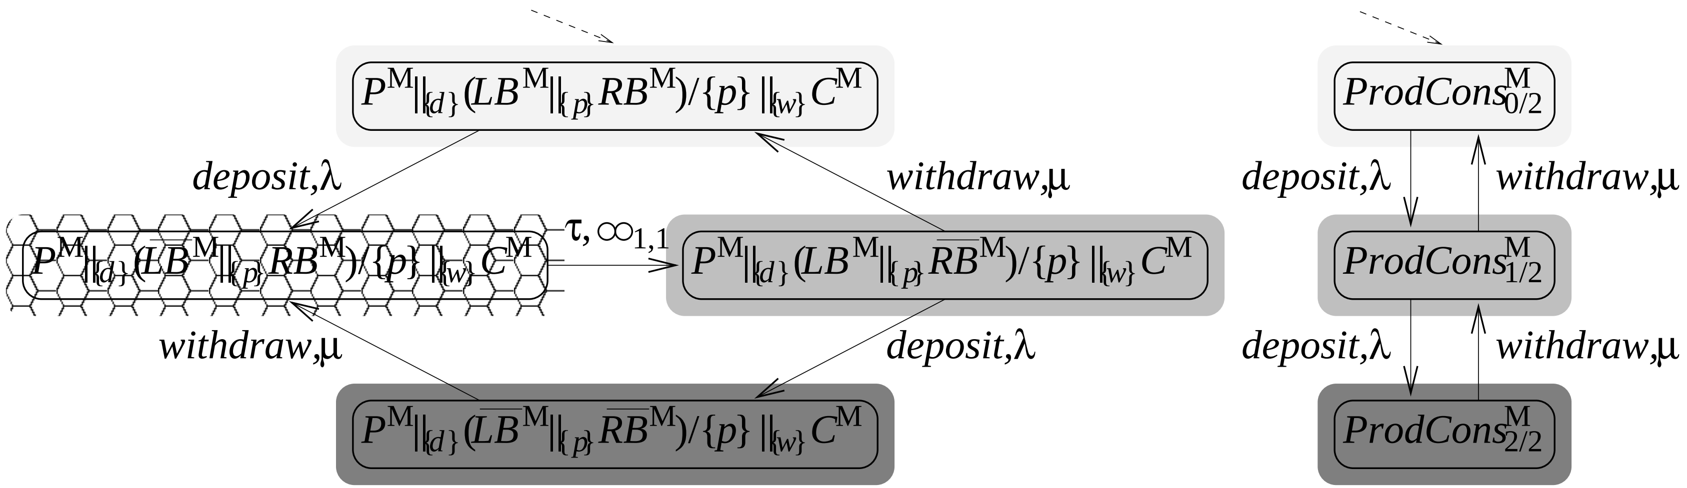

- All the states are observable except for the leftmost state, which is initially unobservable.

Part V:
Markovian Testing Equivalence

Equivalence Definition

- Two process terms are equivalent if an external observer cannot distinguish between them, with the only way for the observer to infer information about their functional and performance behavior being to interact with them by means of tests and compare their reactions.

- Was the test passed?
  If so, with which probability?
  And how long did it take to pass the test?

- Tests formalized as process terms.

- Interaction formalized as parallel composition of process term and test with synchronization enforced on any visible action name.

- Comparison of process term probabilities of performing successful test-driven computations within arbitrary time upper bounds.

- A computation of a process term $P \in \mathbb{P}_{\mathrm{M}}$ is a sequence of transitions that can be executed starting from P.

- The *length* of a computation is given by the number of its transitions.

- $\mathcal{C}_{\mathrm{f}}(P)$: multiset of finite-length computations of P.

- Two distinct computations are *independent* of each other iff neither is a proper prefix of the other one.

- Focus on finite multisets of independent, finite-length computations.

- Attributes of a finite-length computation:
  - ⊙ *Trace.*
  - ⊙ *Probability.*
  - ⊙ *Duration.*

- Given a set of sequences, we use:
  - Operator $\_ \circ \_$ for sequence concatenation.
  - Operator $|\_|$ for sequence length.
- The concrete trace associated with the execution of $c \in \mathcal{C}_{\mathrm{f}}(P)$ is the sequence of action names labeling the transitions of c:

$$\boxed{\mathit{trace}_{\mathrm{c}}(c) \;=\; \begin{cases} \varepsilon & \text{if } |c| = 0 \\ a \circ \mathit{trace}_{\mathrm{c}}(c') & \text{if } c \equiv P \xrightarrow{a,\tilde{\lambda}}_{\mathrm{M}} c' \end{cases}}$$

- We denote by $\mathit{trace}(c)$ the visible part of $\mathit{trace}_{\mathrm{c}}(c)$, i.e., the subsequence of $\mathit{trace}_{\mathrm{c}}(c)$ obtained by removing all the occurrences of τ.

- For the quantitative attributes, we assume $P \in \mathbb{P}_{\mathrm{M,pc}}$.
- The probability of executing $c \in \mathcal{C}_{\mathrm{f}}(P)$ is the product of the execution probabilities of the transitions of c:

$$\boxed{\mathit{prob}(c) \; = \; \begin{cases} 1 & \text{if } |c| = 0 \\ \frac{\lambda}{\mathit{rate}_{\mathrm{t}}(P,0)} \cdot \mathit{prob}(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_{\mathrm{M}} c' \end{cases}}$$

- Probability of executing a computation in $C \subseteq \mathcal{C}_{\mathrm{f}}(P)$:

$$\boxed{\mathit{prob}(C) \; = \; \sum_{c \in C} \mathit{prob}(c)}$$

assuming that C is finite and all of its computations are independent.

- The stepwise average duration of $c \in \mathcal{C}_{\mathrm{f}}(P)$ is the sequence of average sojourn times in the states traversed by c:

$$\boxed{\mathit{time}_{\mathrm{a}}(c) \;=\; \begin{cases} \varepsilon & \text{if } |c| = 0 \\ \frac{1}{\mathit{rate}_{\mathrm{t}}(P,0)} \circ \mathit{time}_{\mathrm{a}}(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_{\mathrm{M}} c' \end{cases}}$$

- Multiset of computations in $C \subseteq \mathcal{C}_{\mathrm{f}}(P)$ whose stepwise average duration is not greater than $\theta \in (\mathbb{R}_{>0})^*$:

$$\boxed{C_{\leq\theta} \;=\; \{\!| \, c \in C \mid |c| \leq |\theta| \wedge \forall i = 1, \ldots, |c|. \, \mathit{time}_{\mathrm{a}}(c)[i] \leq \theta[i] \, |\!\}}$$

- C^l: multiset of computations in $C \subseteq \mathcal{C}_{\mathrm{f}}(P)$ having length $l \in \mathbb{N}$.

- The stepwise duration of $c \in \mathcal{C}_{\mathrm{f}}(P)$ is the sequence of random variables quantifying the sojourn times in the states traversed by c:

$$\boxed{\mathit{time}_{\mathrm{d}}(c) \;=\; \begin{cases} \varepsilon & \text{if } |c| = 0 \\ \mathit{Exp}_{\mathit{rate}_{\mathrm{t}}(P,0)} \circ \mathit{time}_{\mathrm{d}}(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_{\mathrm{M}} c' \end{cases}}$$

- Probability distribution of executing a computation in $C \subseteq \mathcal{C}_{\mathrm{f}}(P)$ within a sequence $\theta \in (\mathbb{R}_{>0})^*$ of time units:

$$\boxed{\mathit{prob}_{\mathrm{d}}(C, \theta) \;=\; \sum_{c \in C}^{|c| \leq |\theta|} \mathit{prob}(c) \cdot \prod_{i=1}^{|c|} \Pr\{\mathit{time}_{\mathrm{d}}(c)[i] \leq \theta[i]\}}$$

  assuming that C is finite and all of its computations are independent.

- Factor $\Pr\{\mathit{time}_{\mathrm{d}}(c)[i] \leq \theta[i]\} = 1 - \mathrm{e}^{-\theta[i]/\mathit{time}_{\mathrm{a}}(c)[i]}$ stems from the cumulative distribution function of the exponentially distributed random variable $\mathit{time}_{\mathrm{d}}(c)[i]$ (whose expected value is $\mathit{time}_{\mathrm{a}}(c)[i]$).

- Why not summing up sojourn times? (standard duration instead of stepwise one)
- Consider process terms ($\lambda \neq \mu$, $b \neq d$, identical nonmaximal computations):

$$<g,\gamma>.<a,\lambda>.<b,\mu>.\underline{0} + <g,\gamma>.<a,\mu>.<d,\lambda>.\underline{0}$$
$$<g,\gamma>.<a,\lambda>.<d,\mu>.\underline{0} + <g,\gamma>.<a,\mu>.<b,\lambda>.\underline{0}$$

- Maximal computations of the first term:

$$c_{1,1} \equiv \cdot \xrightarrow{g,\gamma}_{\mathrm{M}} \cdot \xrightarrow{a,\lambda}_{\mathrm{M}} \cdot \xrightarrow{b,\mu}_{\mathrm{M}} \cdot$$
$$c_{1,2} \equiv \cdot \xrightarrow{g,\gamma}_{\mathrm{M}} \cdot \xrightarrow{a,\mu}_{\mathrm{M}} \cdot \xrightarrow{d,\lambda}_{\mathrm{M}} \cdot$$

- Maximal computations of the second term:

$$c_{2,1} \equiv \cdot \xrightarrow{g,\gamma}_{\mathrm{M}} \cdot \xrightarrow{a,\lambda}_{\mathrm{M}} \cdot \xrightarrow{d,\mu}_{\mathrm{M}} \cdot$$
$$c_{2,2} \equiv \cdot \xrightarrow{g,\gamma}_{\mathrm{M}} \cdot \xrightarrow{a,\mu}_{\mathrm{M}} \cdot \xrightarrow{b,\lambda}_{\mathrm{M}} \cdot$$

- Same sum of average sojourn times $\frac{1}{2\cdot\gamma} + \frac{1}{\lambda} + \frac{1}{\mu}$ and $\frac{1}{2\cdot\gamma} + \frac{1}{\mu} + \frac{1}{\lambda}$ but ...
- *... an external observer would be able to distinguish between the two terms by taking note of the instants at which the actions are performed.*

- Comparing probabilities of passing a test within a time upper bound.
- Syntax of the set $\mathbb{T}_{\mathrm{R}}$ of reactive tests $(a \in \mathit{Name}_{\mathrm{v}},\ w \in \mathbb{R}_{>0})$:

$$\boxed{\begin{array}{rcl} T & ::= & \mathrm{s} \mid T' \\ T' & ::= & <a, *_w>.T \mid T' + T' \end{array}}$$

- Asymmetric action synchronization: only passive actions within tests.
- Performance closure: passive τ-actions not admitted within tests.
- Presence of a time upper bound: recursion not necessary within tests.
- Denoting test passing: zeroary success operator s (success action may interfere).
- Avoiding ambiguous tests like $\mathrm{s} + T$: two-level syntax for tests.

- Interaction system of $P \in \mathbb{P}_{\mathrm{M,pc}}$ and $T \in \mathbb{T}_{\mathrm{R}}$:

$$\boxed{P \,\|_{Name_{\mathrm{v}}}\, T \;\in\; \mathbb{P}_{\mathrm{M,pc}}}$$

- In any of its states, P generates the proposal of an action to be executed by means of a race among the exponentially timed actions enabled in that state.

- If the name of the proposed action is τ, then P advances by itself.

- Otherwise T:

  ⊙ either reacts by participating in the interaction with P through a passive action having the same name;

  ⊙ or blocks the interaction if it has no passive actions with the proposed name.

- Let $P \in \mathbb{P}_{\mathrm{M,pc}}$ and $T \in \mathbb{T}_{\mathrm{R}}$:
  - A *configuration* is a state of $[\![P \parallel_{Name_{\mathrm{v}}} T]\!]_{\mathrm{M}}$.
  - A *test-driven computation* is a computation of $[\![P \parallel_{Name_{\mathrm{v}}} T]\!]_{\mathrm{M}}$.
  - A configuration is formed by *process projection* and *test projection*.
  - A configuration is *successful* iff its test projection is s.
  - A test-driven computation is *successful* iff it traverses a successful configuration.
  - $\mathcal{SC}(P, T)$: multiset of successful computations of $P \parallel_{Name_{\mathrm{v}}} T$.

- If P has no exponentially timed τ-actions:
  - All the computations in $\mathcal{SC}(P, T)$ have a finite length due to the restrictions imposed on the test syntax.
  - All the computations in $\mathcal{SC}(P, T)$ are independent of each other because of their maximality.
  - The multiset $\mathcal{SC}(P, T)$ is finite because both P and T are finitely branching.
- Same considerations for $\mathcal{SC}_{\leq \theta}(P, T)$.
- If there are exponentially timed τ-actions:
  - Are the computations in $\mathcal{SC}_{\leq \theta}(P, T)$ independent of each other?
  - How to distinguish among process terms having only exponentially timed τ-actions, like $<\tau, \lambda>.\underline{0}$ and $<\tau, \mu>.\underline{0}$ with $\lambda > \mu$?

- *Consider subsets of* $\mathcal{SC}_{\leq\theta}(P,T)$ *including all the successful test-driven computations of the same length.*
- They are $\mathcal{SC}^{l}_{\leq\theta}(P,T)$ for $0 \leq l \leq |\theta|$.
- $\mathcal{SC}^{|\theta|}_{\leq\theta}(P,T)$ is enough as shorter successful test-driven computations can be taken into account when imposing prefixes of θ as time upper bounds.
- Process terms having only exponentially timed τ-actions are compared after giving them the possibility of executing the same number of τ-actions ($\lambda > \mu \Rightarrow \frac{1}{\lambda} < \frac{1}{\mu}$):

$$prob(\mathcal{SC}^{1}_{\leq\frac{1}{\lambda}}(<\tau,\lambda>.\underline{0},\mathrm{s})) = 1 \neq 0 = prob(\mathcal{SC}^{1}_{\leq\frac{1}{\lambda}}(<\tau,\mu>.\underline{0},\mathrm{s}))$$

- $P_1 \in \mathbb{P}_{\mathrm{M,pc}}$ is Markovian testing equivalent to $P_2 \in \mathbb{P}_{\mathrm{M,pc}}$, written $P_1 \sim_{\mathrm{MT}} P_2$, iff for all reactive tests $T \in \mathbb{T}_{\mathrm{R}}$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of average amounts of time:

$$\boxed{prob(\mathcal{SC}^{|\theta|}_{\leq \theta}(P_1, T)) \; = \; prob(\mathcal{SC}^{|\theta|}_{\leq \theta}(P_2, T))}$$

- Not defined as the intersection of may- and must-equivalence as the possibility and the necessity of passing a test are qualitative concepts, hence they are not sufficient (probability > 0, probability $= 1$).

- Not defined as the kernel of a Markovian testing preorder as such a preorder would have boiled down to an equivalence relation.

- The presence of time upper bounds makes it possible to decide whether a test is passed or not even if the process term under test can execute infinitely many exponentially timed τ-actions.

- $\sim_{\mathrm{MT}}$ is strictly finer than classical testing equivalence ($a \neq b$ and $\lambda \neq \mu$):

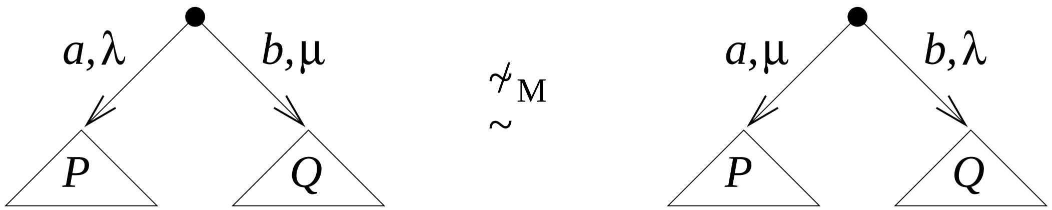

- $\sim_{\mathrm{MT}}$ is strictly finer than probabilistic testing equivalence:

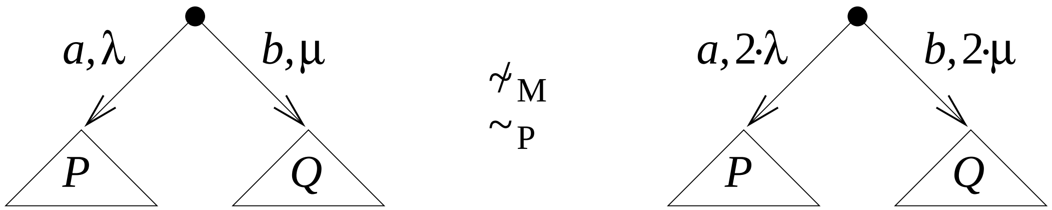

- $\sim_{\mathrm{MT}}$ is strictly coarser than $\sim_{\mathrm{MB}}$ ($P \not\sim_{\mathrm{MB}} Q$):

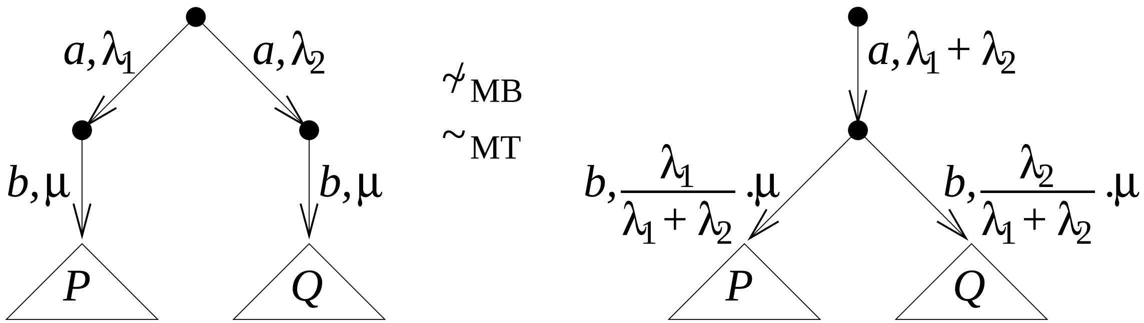

Conditions and Characterizations

- In order for $P_1 \sim_{\mathrm{MT}} P_2$, it is necessary that for all $c_k \in \mathcal{C}_{\mathrm{f}}(P_k)$, $k \in \{1, 2\}$, there exists $c_h \in \mathcal{C}_{\mathrm{f}}(P_h)$, $h \in \{1, 2\} - \{k\}$, such that:

$$\boxed{\begin{array}{rcl} \mathit{trace}_{\mathrm{c}}(c_k) & = & \mathit{trace}_{\mathrm{c}}(c_h) \\ \mathit{time}_{\mathrm{a}}(c_k) & = & \mathit{time}_{\mathrm{a}}(c_h) \end{array}}$$

  and for all $a \in \mathit{Name}$ and $i \in \{0, \ldots, |c_k|\}$:

$$\boxed{\mathit{rate}_{\mathrm{o}}(P_k^i, a, 0) \; = \; \mathit{rate}_{\mathrm{o}}(P_h^i, a, 0)}$$

  with P_k^i (resp. P_h^i) being the i-th state traversed by c_k (resp. c_h).

- Process terms satisfying the necessary condition that are not Markovian testing equivalent ($\lambda_1 + \lambda_2 = \lambda'_1 + \lambda'_2$ with $\lambda_1 \neq \lambda'_1$, $\lambda_2 \neq \lambda'_2$ and $b \neq c$ or $\mu \neq \gamma$):

$$<a, \lambda_1>.<b, \mu>.\underline{0} + <a, \lambda_2>.<c, \gamma>.\underline{0}$$

$$<a, \lambda'_1>.<b, \mu>.\underline{0} + <a, \lambda'_2>.<c, \gamma>.\underline{0}$$

- $\sim_{\mathrm{MT}}$ has three alternative characterizations, each providing further justifications for the way in which the equivalence has been defined.

- The first characterization establishes that the discriminating power does not change if we consider a set $\mathbb{T}_{\mathrm{R,lib}}$ of tests with the following more liberal syntax:

$$\boxed{T \ ::= \ \mathrm{s} \mid <a, *_w>.T \mid T + T}$$

- In this setting, a successful configuration is a configuration whose test projection includes s as top-level summand.

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MT,lib}} P_2 \iff P_1 \sim_{\mathrm{MT}} P_2}$$

- The second characterization establishes that the discriminating power does not change if we consider a set $\mathbb{T}_{\mathrm{R},\tau}$ of tests capable of moving autonomously by executing exponentially timed τ-actions:

$$\boxed{\begin{array}{rcl} T & ::= & \mathrm{s} \mid T' \\ T' & ::= & <a, *_w>.T \mid <\tau, \lambda>.T \mid T' + T' \end{array}}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MT},\tau} P_2 \iff P_1 \sim_{\mathrm{MT}} P_2}$$

- The third characterization establishes that the discriminating power does not change if we consider the probability distribution of passing tests within arbitrary sequences of amounts of time.

- Considering the (more accurate) stepwise durations of test-driven computations leads to the same equivalence as considering the (easier to work with) stepwise average durations.

- $P_1 \in \mathbb{P}_{\mathrm{M,pc}}$ is Markovian distribution-testing equivalent to $P_2 \in \mathbb{P}_{\mathrm{M,pc}}$, written $P_1 \sim_{\mathrm{MT,d}} P_2$, iff for all reactive tests $T \in \mathbb{T}_{\mathrm{R}}$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of amounts of time:

$$\boxed{prob_{\mathrm{d}}(\mathcal{SC}^{|\theta|}(P_1, T), \theta) \; = \; prob_{\mathrm{d}}(\mathcal{SC}^{|\theta|}(P_2, T), \theta)}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MT,d}} P_2 \iff P_1 \sim_{\mathrm{MT}} P_2}$$

- $\sim_{\mathrm{MT}}$ has another alternative characterization that *fully abstracts* from comparing process term behavior in response to tests.

- Based on traces that are extended at each step with the set of visible action names permitted by the environment at that step.

- An element ξ of $(\mathit{Name}_{\mathrm{v}} \times 2^{\mathit{Name}_{\mathrm{v}}})^*$ is an extended trace iff either ξ is the empty sequence ε or:

$$\boxed{\xi \;\equiv\; (a_1, \mathcal{E}_1) \circ (a_2, \mathcal{E}_2) \circ \ldots \circ (a_n, \mathcal{E}_n)}$$

  for some $n \in \mathbb{N}_{>0}$ with $a_i \in \mathcal{E}_i$ and $\mathcal{E}_i$ finite for each $i = 1, \ldots, n$.

- $\mathcal{ET}$: set of extended traces.

- Trace associated with $\xi \in \mathcal{ET}$:

$$\boxed{trace_{\mathrm{et}}(\xi) \; = \; \begin{cases} \varepsilon & \text{if } |\xi| = 0 \\ a \circ trace_{\mathrm{et}}(\xi') & \text{if } \xi \equiv (a, \mathcal{E}) \circ \xi' \end{cases}}$$

- $c \in \mathcal{C}_{\mathrm{f}}(P)$ is compatible with $\xi \in \mathcal{ET}$ iff:

$$\boxed{trace(c) \; = \; trace_{\mathrm{et}}(\xi)}$$

- $\mathcal{CC}(P, \xi)$: multiset of computations in $\mathcal{C}_{\mathrm{f}}(P)$ compatible with ξ.

- The probability and the duration of any computation of $\mathcal{CC}(P, \xi)$ have to be calculated by considering only the action names permitted at each step by ξ.

- Probability w.r.t. ξ of executing $c \in \mathcal{CC}(P, \xi)$:

$$\boxed{prob_{\xi}(c) \; = \; \begin{cases} 1 & \text{if } |c| = 0 \\ \frac{\lambda}{rate_{\mathrm{o}}(P, \mathcal{E} \cup \{\tau\}, 0)} \cdot prob_{\xi'}(c') & \text{if } c \equiv P \xrightarrow{a, \lambda}_{\mathrm{M}} c' \\ & \quad \text{with } \xi \equiv (a, \mathcal{E}) \circ \xi' \\ \frac{\lambda}{rate_{\mathrm{o}}(P, \mathcal{E} \cup \{\tau\}, 0)} \cdot prob_{\xi}(c') & \text{if } c \equiv P \xrightarrow{\tau, \lambda}_{\mathrm{M}} c' \\ & \quad \text{with } \xi \equiv (a, \mathcal{E}) \circ \xi' \\ \frac{\lambda}{rate_{\mathrm{o}}(P, \tau, 0)} \cdot prob_{\xi}(c') & \text{if } c \equiv P \xrightarrow{\tau, \lambda}_{\mathrm{M}} c' \wedge \xi \equiv \varepsilon \end{cases}}$$

- Probability w.r.t. ξ of executing a computation in $C \subseteq \mathcal{CC}(P, \xi)$:

$$\boxed{prob_{\xi}(C) \; = \; \sum_{c \in C} prob_{\xi}(c)}$$

assuming that C is finite and all of its computations are independent.

- Stepwise average duration w.r.t. ξ of $c \in \mathcal{CC}(P, \xi)$:

$$\boxed{time_{\mathrm{a},\xi}(c) \;=\; \begin{cases} \varepsilon & \text{if } |c| = 0 \\ \frac{1}{rate_{\mathrm{o}}(P, \mathcal{E} \cup \{\tau\}, 0)} \circ time_{\mathrm{a},\xi'}(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_{\mathrm{M}} c' \\ & \quad \text{with } \xi \equiv (a, \mathcal{E}) \circ \xi' \\ \frac{1}{rate_{\mathrm{o}}(P, \mathcal{E} \cup \{\tau\}, 0)} \circ time_{\mathrm{a},\xi}(c') & \text{if } c \equiv P \xrightarrow{\tau,\lambda}_{\mathrm{M}} c' \\ & \quad \text{with } \xi \equiv (a, \mathcal{E}) \circ \xi' \\ \frac{1}{rate_{\mathrm{o}}(P, \tau, 0)} \circ time_{\mathrm{a},\xi}(c') & \text{if } c \equiv P \xrightarrow{\tau,\lambda}_{\mathrm{M}} c' \wedge \xi \equiv \varepsilon \end{cases}}$$

- Multiset of computations in $C \subseteq \mathcal{CC}(P, \xi)$ whose stepwise average duration w.r.t. ξ is not greater than $\theta \in (\mathbb{R}_{>0})^*$:

$$\boxed{C_{\leq \theta, \xi} \;=\; \{\!| \, c \in C \mid |c| \leq |\theta| \wedge \forall i = 1, \dots, |c|. \, time_{\mathrm{a},\xi}(c)[i] \leq \theta[i] \, |\!\}}$$

- C^l: multiset of computations in $C \subseteq \mathcal{CC}(P, \xi)$ having length $l \in \mathbb{N}$.

- Consider $\mathcal{CC}^{|\theta|}_{\leq\theta,\xi}(P, \xi)$ in order to ensure independence.
- $P_1 \in \mathbb{P}_{\mathrm{M,pc}}$ is Markovian extended-trace equivalent to $P_2 \in \mathbb{P}_{\mathrm{M,pc}}$, written $P_1 \sim_{\mathrm{MTr,e}} P_2$, iff for all extended traces $\xi \in \mathcal{ET}$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of average amounts of time:

$$\boxed{prob_\xi(\mathcal{CC}^{|\theta|}_{\leq\theta,\xi}(P_1, \xi)) \; = \; prob_\xi(\mathcal{CC}^{|\theta|}_{\leq\theta,\xi}(P_2, \xi))}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MTr,e}} P_2 \iff P_1 \sim_{\mathrm{MT}} P_2}$$

- Extended traces identify a set of reactive tests necessary and sufficient in order to establish whether two terms are Markovian testing equivalent.
- Each canonical reactive test admits a main computation leading to success, whose intermediate states can have additional computations each leading to failure in one step.
- Failure is represented through a visible action name z that can occur within tests but not within process terms under test.
- Syntax of the set $\mathbb{T}_{\mathrm{R,c}}$ of canonical reactive tests ($a \in \mathcal{E}, \mathcal{E} \subseteq Name_{\mathrm{v}}$ finite):

$$\boxed{T \ ::= \ \mathrm{s} \mid <a, *_1>.T + \sum_{b \in \mathcal{E} - \{a\}} <b, *_1>.<\mathrm{z}, *_1>.\mathrm{s}}$$

- $P_1 \sim_{\mathrm{MT}} P_2$ iff for all $T \in \mathbb{T}_{\mathrm{R,c}}$ and $\theta \in (\mathbb{R}_{>0})^*$:

$$\boxed{prob(\mathcal{SC}^{|\theta|}_{\leq \theta}(P_1, T)) \ = \ prob(\mathcal{SC}^{|\theta|}_{\leq \theta}(P_2, T))}$$

- **Running example** ($\sim_{\rm MT}$):
  - Concurrent implementation with two independent one-pos. buffers:

$$\begin{aligned}
PC^{\rm M}_{\rm conc,2} &\triangleq Prod^{\rm M} \parallel_{\{deposit\}} (Buff^{\rm M} \parallel_{\emptyset} Buff^{\rm M}) \parallel_{\{withdraw\}} Cons^{\rm M} \\
Prod^{\rm M} &\triangleq <deposit, \lambda>.Prod^{\rm M} \\
Buff^{\rm M} &\triangleq <deposit, *_1>.<withdraw, *_1>.Buff^{\rm M} \\
Cons^{\rm M} &\triangleq <withdraw, \mu>.Cons^{\rm M}
\end{aligned}$$

  - All the actions occurring in the buffer are passive, consistent with the fact that the buffer is a passive entity.
  - Is $PC^{\rm M}_{\rm conc,2}$ a correct implementation of $ProdCons^{\rm M}_{0/2}$?
  - It turns out that $PC^{\rm M}_{\rm conc,2} \sim_{\rm MT} ProdCons^{\rm M}_{0/2}$.
  - Proved by exploiting the fully abstract alternative characterization.

- Here are the underlying labeled multitransition systems:

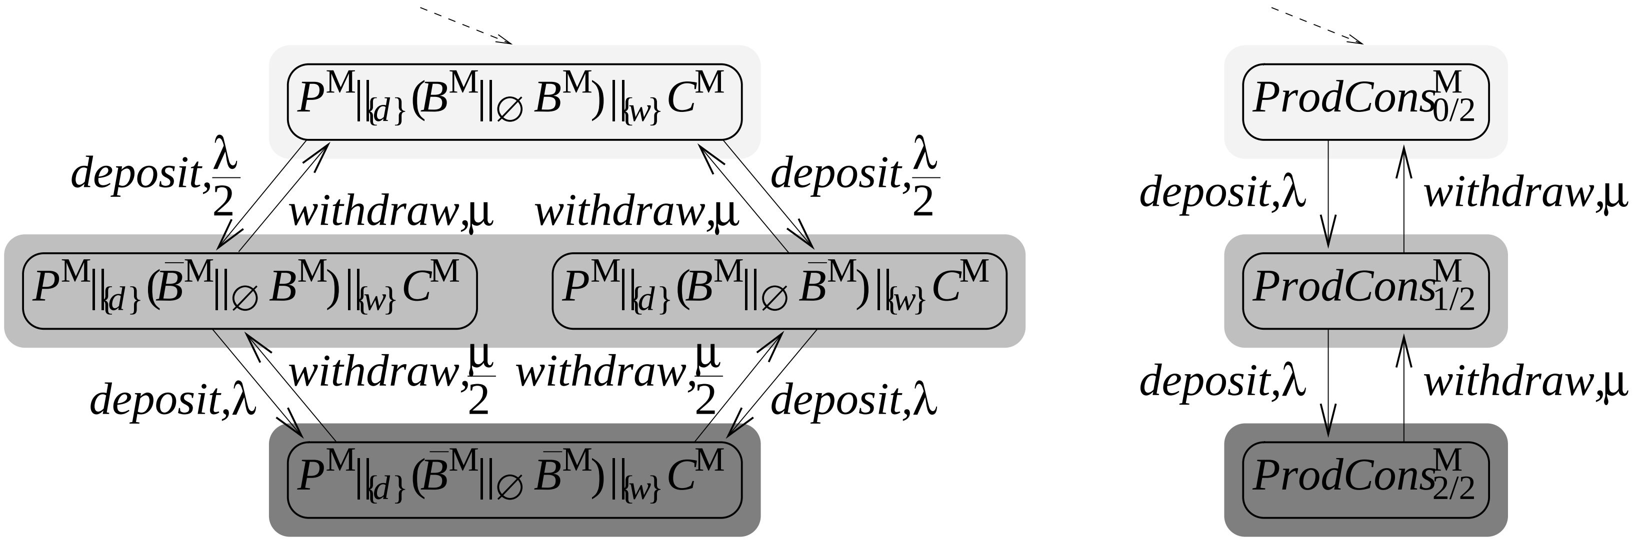

- The initial state on the left-hand side has both outgoing transitions labeled with $\lambda/2$, not λ.

- The bottom state on the left-hand side has both outgoing transitions labeled with $\mu/2$, not μ.

- The only sequences of visible actions that the two systems are able to perform are the prefixes of the strings complying with:

$$(deposit \circ (deposit \circ withdraw)^* \circ withdraw)^*$$

- The only significant extended traces to be considered are those whose associated traces coincide with such prefixes.

- Their nonempty finite sets of visible actions permitted at the various steps necessarily contain at least one between *deposit* and *withdraw*.

- Any two computations of $ProdCons^{\mathrm{M}}_{0/2}$ and $PC^{\mathrm{M}}_{\mathrm{conc},2}$ compatible with such a ξ traverse states that pairwise enable sets of actions with the same names and total rates.

- Therefore the stepwise average durations with respect to ξ of the considered computations are identical.

- $\odot$ Four basic cases for the execution probabilities with respect to ξ of $\mathcal{CC}(PC^{\mathrm{M}}_{\mathrm{conc},2}, \xi)$ and $\mathcal{CC}(ProdCons^{\mathrm{M}}_{0/2}, \xi)$:
  - $*$ If $\xi \equiv (deposit, \mathcal{E})$, then for both sets of computations the execution probability is 1.
  - $*$ If $\xi \equiv (deposit, \mathcal{E}_1) \circ (withdraw, \mathcal{E}_2)$, then for both sets of computations the execution probability is 1 if $\mathcal{E}_2$ does not contain *deposit*, $\frac{\mu}{\lambda+\mu}$ otherwise.
  - $*$ If $\xi \equiv (deposit, \mathcal{E}_1) \circ (deposit, \mathcal{E}_2)$, then for both sets of computations the execution probability is 1 if $\mathcal{E}_2$ does not contain *withdraw*, $\frac{\lambda}{\lambda+\mu}$ otherwise.
  - $*$ If $\xi \equiv (deposit, \mathcal{E}_1) \circ (deposit, \mathcal{E}_2) \circ (withdraw, \mathcal{E}_3)$, then for both sets of computations the execution probability is 1 if $\mathcal{E}_2$ does not contain *withdraw*, $\frac{\lambda}{\lambda+\mu}$ otherwise.

Equivalence Properties

- $\sim_{\rm MT}$ is a congruence over $\mathbb{P}_{\rm M,pc}$ with respect to all the dynamic and static operators as well as recursion.

- Let $P_1, P_2 \in \mathbb{P}_{\rm M,pc}$. Whenever $P_1 \sim_{\rm MT} P_2$, then:

$$\boxed{\begin{array}{c}
<\!a, \lambda\!>.P_1 \;\sim_{\rm MT}\; <\!a, \lambda\!>.P_2 \\[4pt]
P_1 + P \;\sim_{\rm MT}\; P_2 + P \qquad\qquad P + P_1 \;\sim_{\rm MT}\; P + P_2 \\[4pt]
P_1 \,\|_S\, P \;\sim_{\rm MT}\; P_2 \,\|_S\, P \qquad\qquad P \,\|_S\, P_1 \;\sim_{\rm MT}\; P \,\|_S\, P_2 \\[4pt]
P_1 \,/\, H \;\sim_{\rm MT}\; P_2 \,/\, H \\[4pt]
P_1 \setminus L \;\sim_{\rm MT}\; P_2 \setminus L \\[4pt]
P_1\,[\varphi] \;\sim_{\rm MT}\; P_2[\varphi]
\end{array}}$$

provided that $P \in \mathbb{P}_{\rm M,pc}$ for the alternative composition operator and $P_1 \,\|_S\, P, P_2 \,\|_S\, P \in \mathbb{P}_{\rm M,pc}$ for the parallel composition operator.

- Recursion: extend $\sim_{\mathrm{MT}}$ to open process terms by replacing all variables freely occurring outside rec binders with every closed process term.

- Let $P_1, P_2 \in \mathcal{PL}_{\mathrm{M}}$ be guarded process terms containing free occurrences of $k \in \mathbb{N}$ process variables $X_1, \ldots, X_k \in \mathit{Var}$ at most.

- We define $P_1 \sim_{\mathrm{MT}} P_2$ iff there exist $Q_1, \ldots, Q_k \in \mathbb{P}_{\mathrm{M}}$ such that both $P_1\{Q_i \hookrightarrow X_i \mid 1 \leq i \leq k\}$ and $P_2\{Q_i \hookrightarrow X_i \mid 1 \leq i \leq k\}$ belong to $\mathbb{P}_{\mathrm{M,pc}}$ and for each such group of process terms $Q_1, \ldots, Q_k \in \mathbb{P}_{\mathrm{M}}$:

$$\boxed{P_1\{Q_i \hookrightarrow X_i \mid 1 \leq i \leq k\} \ \sim_{\mathrm{MT}} \ P_2\{Q_i \hookrightarrow X_i \mid 1 \leq i \leq k\}}$$

- Whenever $P_1 \sim_{\mathrm{MT}} P_2$, then:

$$\boxed{\mathrm{rec}\, X : P_1 \ \sim_{\mathrm{MT}} \ \mathrm{rec}\, X : P_2}$$

- $\sim_{\mathrm{MT}}$ has a sound and complete axiomatization over the set $\mathbb{P}_{\mathrm{M,pc,nrec}}$ of nonrecursive process terms of $\mathbb{P}_{\mathrm{M,pc}}$.

- The axioms for $\sim_{\mathrm{MB}}$ are sound but not complete for $\sim_{\mathrm{MT}}$ $(P \not\sim_{\mathrm{MB}} Q)$:

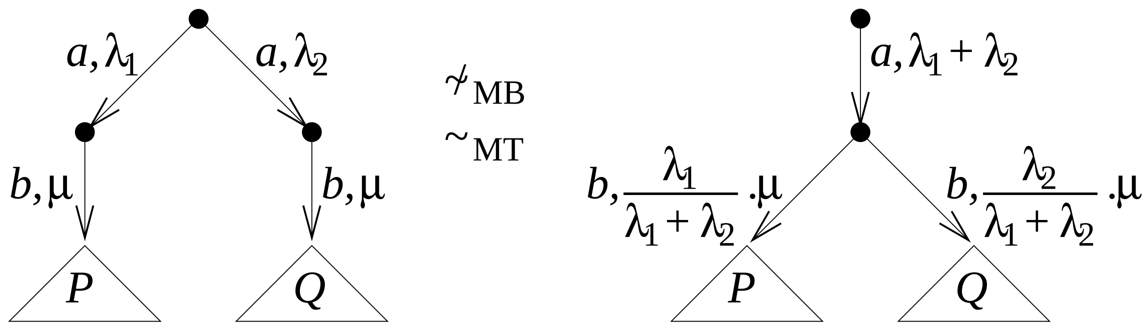

- *Possibility of deferring choices related to branches starting with the same action name (see the two a-branches on the left-hand side) that are immediately followed by sets of actions having the same names and total rates (see* $\{<b, \mu>\}$ *after each of the two a-branches).*

- Basic laws (identical to those for $\sim_{\mathrm{MB}}$):

$$
\begin{array}{llcl}
(\mathcal{X}_{\mathrm{MT},1}) & P_1 + P_2 & = & P_2 + P_1 \\
(\mathcal{X}_{\mathrm{MT},2}) & (P_1 + P_2) + P_3 & = & P_1 + (P_2 + P_3) \\
(\mathcal{X}_{\mathrm{MT},3}) & P + \underline{0} & = & P
\end{array}
$$

- Characterizing law (subsumes $\sim_{\mathrm{MB}}$ axiom for race policy):

$$
(\mathcal{X}_{\mathrm{MT},4}) \quad \sum_{i \in I} {<}a, \lambda_i{>}. \sum_{j \in J_i} {<}b_{i,j}, \mu_{i,j}{>}.P_{i,j} \;=\; {<}a, \sum_{k \in I} \lambda_k{>}. \sum_{i \in I} \sum_{j \in J_i} {<}b_{i,j}, \frac{\lambda_i}{\sum_{k \in I} \lambda_k} \cdot \mu_{i,j}{>}.P_{i,j}
$$

if: I is a finite index set with $|I| \geq 2$;
for all $i \in I$, index set J_i is finite and its summation is $\underline{0}$ if $J_i = \emptyset$;
for all $i_1, i_2 \in I$ and $b \in \mathit{Name}$:

$$
\sum_{j \in J_{i_1}} \{\!| \mu_{i_1,j} \mid b_{i_1,j} = b |\!\} \;=\; \sum_{j \in J_{i_2}} \{\!| \mu_{i_2,j} \mid b_{i_2,j} = b |\!\}
$$

- Expansion law (identical to that for $\sim_{\mathrm{MB}}$):

$$
\begin{array}{lrcl}
(\mathcal{X}_{\mathrm{MT},5}) & \sum\limits_{i \in I} {<} a_i, \tilde{\lambda}_i {>} . P_i \parallel_S \sum\limits_{j \in J} {<} b_j, \tilde{\mu}_j {>} . Q_j & = & \\
& & & \sum\limits_{k \in I, a_k \notin S} {<} a_k, \tilde{\lambda}_k {>} . \left(P_k \parallel_S \sum\limits_{j \in J} {<} b_j, \tilde{\mu}_j {>} . Q_j \right) \; + \\
& & & \sum\limits_{h \in J, b_h \notin S} {<} b_h, \tilde{\mu}_h {>} . \left(\sum\limits_{i \in I} {<} a_i, \tilde{\lambda}_i {>} . P_i \parallel_S Q_h \right) \; + \\
& & & \sum\limits_{k \in I, a_k \in S, \tilde{\lambda}_k \in \mathbb{R}_{>0}} \; \sum\limits_{h \in J, b_h = a_k, \tilde{\mu}_h = *_{w_h}} {<} a_k, \tilde{\lambda}_k \cdot \frac{w_h}{weight(Q, b_h)} {>} . (P_k \parallel_S Q_h) \; + \\
& & & \sum\limits_{h \in J, b_h \in S, \tilde{\mu}_h \in \mathbb{R}_{>0}} \; \sum\limits_{k \in I, a_k = b_h, \tilde{\lambda}_k = *_{v_k}} {<} b_h, \tilde{\mu}_h \cdot \frac{v_k}{weight(P, a_k)} {>} . (P_k \parallel_S Q_h) \; + \\
& & & \sum\limits_{k \in I, a_k \in S, \tilde{\lambda}_k = *_{v_k}} \; \sum\limits_{h \in J, b_h = a_k, \tilde{\mu}_h = *_{w_h}} {<} a_k, *_{norm(v_k, w_h, a_k, P, Q)} {>} . (P_k \parallel_S Q_h) \\
(\mathcal{X}_{\mathrm{MT},6}) & \sum\limits_{i \in I} {<} a_i, \tilde{\lambda}_i {>} . P_i \parallel_S \underline{0} & = & \sum\limits_{k \in I, a_k \notin S} {<} a_k, \tilde{\lambda}_k {>} . P_k \\
(\mathcal{X}_{\mathrm{MT},7}) & \underline{0} \parallel_S \sum\limits_{j \in J} {<} b_j, \tilde{\mu}_j {>} . Q_j & = & \sum\limits_{h \in J, b_h \notin S} {<} b_h, \tilde{\mu}_h {>} . Q_h \\
(\mathcal{X}_{\mathrm{MT},8}) & \underline{0} \parallel_S \underline{0} & = & \underline{0}
\end{array}
$$

- Distribution laws (identical to those for $\sim_{\rm MB}$):

$$\begin{array}{lrcll}
(\mathcal{X}_{\rm MT,9}) & \underline{0} \,/\, H & = & \underline{0} & \\
(\mathcal{X}_{\rm MT,10}) & (<a, \tilde{\lambda}>.P) \,/\, H & = & <\tau, \tilde{\lambda}>.(P \,/\, H) & \text{if } a \in H \\
(\mathcal{X}_{\rm MT,11}) & (<a, \tilde{\lambda}>.P) \,/\, H & = & <a, \tilde{\lambda}>.(P \,/\, H) & \text{if } a \notin H \\
(\mathcal{X}_{\rm MT,12}) & (P_1 + P_2) \,/\, H & = & P_1 \,/\, H + P_2 \,/\, H & \\
 & & & & \\
(\mathcal{X}_{\rm MT,13}) & \underline{0} \setminus L & = & \underline{0} & \\
(\mathcal{X}_{\rm MT,14}) & (<a, \tilde{\lambda}>.P) \setminus L & = & \underline{0} & \text{if } a \in L \\
(\mathcal{X}_{\rm MT,15}) & (<a, \tilde{\lambda}>.P) \setminus L & = & <a, \tilde{\lambda}>.(P \setminus L) & \text{if } a \notin L \\
(\mathcal{X}_{\rm MT,16}) & (P_1 + P_2) \setminus L & = & P_1 \setminus L + P_2 \setminus L & \\
 & & & & \\
(\mathcal{X}_{\rm MT,17}) & \underline{0}[\varphi] & = & \underline{0} & \\
(\mathcal{X}_{\rm MT,18}) & (<a, \tilde{\lambda}>.P)[\varphi] & = & <\varphi(a), \tilde{\lambda}>.(P[\varphi]) & \\
(\mathcal{X}_{\rm MT,19}) & (P_1 + P_2)[\varphi] & = & P_1[\varphi] + P_2[\varphi] &
\end{array}$$

- $DED(\mathcal{X}_{\mathrm{MT}})$: deduction system based on all the previous axioms plus:
  - Reflexivity: $\mathcal{X}_{\mathrm{MT}} \vdash P = P$.
  - Symmetry: $\mathcal{X}_{\mathrm{MT}} \vdash P_1 = P_2 \implies \mathcal{X}_{\mathrm{MT}} \vdash P_2 = P_1$.
  - Transitivity: $\mathcal{X}_{\mathrm{MT}} \vdash P_1 = P_2 \wedge \mathcal{X}_{\mathrm{MT}} \vdash P_2 = P_3 \implies \mathcal{X}_{\mathrm{MT}} \vdash P_1 = P_3$.
  - Substitutivity: $\mathcal{X}_{\mathrm{MT}} \vdash P_1 = P_2 \Longrightarrow \mathcal{X}_{\mathrm{MT}} \vdash <a, \lambda>.P_1 = <a, \lambda>.P_2 \wedge \dots$

- The deduction system $DED(\mathcal{X}_{\mathrm{MT}})$ is sound and complete for $\sim_{\mathrm{MT}}$ over $\mathbb{P}_{\mathrm{M,pc,nrec}}$; i.e., for all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc,nrec}}$:

$$\boxed{\mathcal{X}_{\mathrm{MT}} \vdash P_1 = P_2 \iff P_1 \sim_{\mathrm{MT}} P_2}$$

- $\sim_{\rm MT}$ has a modal logic characterization over $\mathbb{P}_{\rm M,pc}$ based on a variant of the Hennessy-Milner logic.

- Negation is not included and conjunction is replaced by disjunction (decreased discriminating power with respect to $\sim_{\rm MB}$).

- Syntax of the modal language $\mathcal{ML}_{\rm MT}$ ($a \in \mathit{Name}_{\rm v}$):

$$\boxed{\begin{array}{rcl} \phi & ::= & \mathrm{true} \mid \phi' \\ \phi' & ::= & \langle a \rangle \phi \mid \phi' \vee \phi' \end{array}}$$

where each formula of the form $\phi_1 \vee \phi_2$ satisfies the following constraint (consistent with the name-deterministic nature of canonical reactive tests):

$$\mathit{init}(\phi_1) \cap \mathit{init}(\phi_2) \; = \; \emptyset$$

with $\mathit{init}(\phi)$ being defined as follows:

$$\mathit{init}(\mathrm{true}) = \emptyset \qquad \mathit{init}(\phi_1 \vee \phi_2) = \mathit{init}(\phi_1) \cup \mathit{init}(\phi_2) \qquad \mathit{init}(\langle a \rangle \phi) = \{a\}$$

- No quantitative decorations in the syntax because the focus is on entire computations rather than on step-by-step behavior mimicking, but ...

- ... replacement of the boolean satisfaction relation with a quantitative interpretation function measuring the probability with which a process term satisfies a formula quickly enough on average.

- Interpretation of $\mathcal{ML}_{\mathrm{MT}}$ over $\mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{[\![\phi]\!]_{\mathrm{MT}}^{|\theta|}(P, \theta) \; = \; \begin{cases} 0 & \text{if } |\theta| = 0 \wedge \phi \not\equiv \mathrm{true} \\ & \quad \text{or } |\theta| > 0 \wedge \mathit{rate}_{\mathrm{o}}(P, \mathit{init}(\phi) \cup \{\tau\}, 0) = 0 \\ 1 & \text{if } |\theta| = 0 \wedge \phi \equiv \mathrm{true} \end{cases}}$$

otherwise:

$$[\![\mathrm{true}]\!]_{\mathrm{MT}}^{|t\circ\theta|}(P, t\circ\theta) = \begin{cases} \sum\limits_{P \xrightarrow{\tau,\lambda}_{\mathrm{M}} P'} \frac{\lambda}{rate_{\mathrm{o}}(P,\tau,0)} \cdot [\![\mathrm{true}]\!]_{\mathrm{MT}}^{|\theta|}(P',\theta) & \text{if } \frac{1}{rate_{\mathrm{o}}(P,\tau,0)} \leq t \\ 0 & \text{if } \frac{1}{rate_{\mathrm{o}}(P,\tau,0)} > t \end{cases}$$

$$[\![\langle a\rangle\phi]\!]_{\mathrm{MT}}^{|t\circ\theta|}(P, t\circ\theta) = \begin{cases} \sum\limits_{P \xrightarrow{a,\lambda}_{\mathrm{M}} P'} \frac{\lambda}{rate_{\mathrm{o}}(P,\{a,\tau\},0)} \cdot [\![\phi]\!]_{\mathrm{MT}}^{|\theta|}(P',\theta) + \\ \quad \sum\limits_{P \xrightarrow{\tau,\lambda}_{\mathrm{M}} P'} \frac{\lambda}{rate_{\mathrm{o}}(P,\{a,\tau\},0)} \cdot [\![\langle a\rangle\phi]\!]_{\mathrm{MT}}^{|\theta|}(P',\theta) \\ & \text{if } \frac{1}{rate_{\mathrm{o}}(P,\{a,\tau\},0)} \leq t \\ 0 & \text{if } \frac{1}{rate_{\mathrm{o}}(P,\{a,\tau\},0)} > t \end{cases}$$

$$
\begin{aligned}
[\![\phi_1 \vee \phi_2]\!]_{\mathrm{MT}}^{|t \circ \theta|}(P, t \circ \theta) \quad = \quad & p_1 \cdot [\![\phi_1]\!]_{\mathrm{MT}}^{|t_1 \circ \theta|}(P_{\textit{no-init-}\tau}, t_1 \circ \theta) + \\
& p_2 \cdot [\![\phi_2]\!]_{\mathrm{MT}}^{|t_2 \circ \theta|}(P_{\textit{no-init-}\tau}, t_2 \circ \theta) + \\
& \sum_{P \xrightarrow{\tau, \lambda}_{\mathrm{M}} P'} \frac{\lambda}{\mathit{rate}_{\mathrm{o}}(P, \mathit{init}(\phi_1 \vee \phi_2) \cup \{\tau\}, 0)} \cdot [\![\phi_1 \vee \phi_2]\!]_{\mathrm{MT}}^{|\theta|}(P', \theta)
\end{aligned}
$$

where:

- $P_{\textit{no-init-}\tau}$ is P without computations starting with a τ-transition.
- For $j \in \{1, 2\}$:

$$
\begin{aligned}
p_j \quad & = \quad \frac{\mathit{rate}_{\mathrm{o}}(P, \mathit{init}(\phi_j), 0)}{\mathit{rate}_{\mathrm{o}}(P, \mathit{init}(\phi_1 \vee \phi_2) \cup \{\tau\}, 0)} \\
t_j \quad & = \quad t + \left(\frac{1}{\mathit{rate}_{\mathrm{o}}(P, \mathit{init}(\phi_j), 0)} - \frac{1}{\mathit{rate}_{\mathrm{o}}(P, \mathit{init}(\phi_1 \vee \phi_2) \cup \{\tau\}, 0)}\right)
\end{aligned}
$$

with p_j representing the conditional probability with which P performs actions whose name is in $\mathit{init}(\phi_j)$ and t_j representing the extra average time granted to P for satisfying ϕ_j.

- The constraint on disjunctions guarantees that their subformulas exercise independent computations of P (correct probability calculation).

- In the absence of p_1 and p_2, the fact that $\phi_1 \vee \phi_2$ offers a set of initial actions at least as large as the ones offered by ϕ_1 alone and by ϕ_2 alone may lead to an overestimate of the probability of satisfying $\phi_1 \vee \phi_2$.

- Considering t instead of t_j in the satisfaction of ϕ_j in isolation may lead to an underestimate of the probability of satisfying $\phi_1 \vee \phi_2$ within the given time upper bound, as P may satisfy $\phi_1 \vee \phi_2$ within $t \circ \theta$ even if P satisfies neither ϕ_1 nor ϕ_2 taken in isolation within $t \circ \theta$.

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MT}} P_2 \iff \forall \phi \in \mathcal{ML}_{\mathrm{MT}}.\, \forall \theta \in (\mathbb{R}_{>0})^*.\, [\![\phi]\!]^{|\theta|}_{\mathrm{MT}}(P_1, \theta) = [\![\phi]\!]^{|\theta|}_{\mathrm{MT}}(P_2, \theta)}$$

- $\sim_{\mathrm{MT}}$ is decidable in polynomial time over the set $\mathbb{P}_{\mathrm{M,pc,fin}}$ of finite-state process terms of $\mathbb{P}_{\mathrm{M,pc}}$.
- The reason is that:
  - $\sim_{\mathrm{MT}}$ coincides with the Markovian version of ready equivalence.
  - Probabilistic ready equivalence can be decided in polynomial time through a suitable reworking of Tzeng algorithm for probabilistic language equivalence.
- Given two process terms, their name-labeled CTMCs are Markovian ready equivalent iff the corresponding embedded name-labeled DTMCs are probabilistic ready equivalent.
- Markovian ready equivalence and probabilistic ready equivalence coincide on corresponding models if the total exit rate of each state of a name-labeled CTMC is encoded inside the names of all transitions departing from that state in the associated embedded DTMC.

- Steps of the algorithm for checking whether $P_1 \sim_{\mathrm{MT}} P_2$:
  1. Transform $[\![P_1]\!]_{\mathrm{M}}$ and $[\![P_2]\!]_{\mathrm{M}}$ into their corresponding embedded discrete-time versions:
     a. Divide the rate of each transition by the total exit rate of its source state.
     b. Augment the name of each transition with the total exit rate of its source state.
  2. Compute the relation $\mathcal{R}$ that equates any two states of the discrete-time versions of $[\![P_1]\!]_{\mathrm{M}}$ and $[\![P_2]\!]_{\mathrm{M}}$ whenever the two sets of augmented action names labeling the transitions departing from the two states coincide.
  3. For each equivalence class R induced by $\mathcal{R}$, consider R as the set of accepting states and check whether the discrete-time versions of $[\![P_1]\!]_{\mathrm{M}}$ and $[\![P_2]\!]_{\mathrm{M}}$ are probabilistic language equivalent.
  4. Return yes/no depending on whether all the checks performed in the previous step have been successful or not.

- Tzeng algorithm for probabilistic language equivalence visits in breadth-first order the tree containing a node for each possible string and studies the linear independence of the state probability vectors associated with a finite subset of the tree nodes.

- Refinement of each iteration of step 3:

  1. Create an empty set V of state probability vectors.
  2. Create a queue whose only element is the empty string ε.
  3. While the queue is not empty:
     a. Remove the first element from the queue, say string ς.
     b. If the state probability vector of the discrete-time versions of $[\![P_1]\!]_{\mathrm{M}}$ and $[\![P_2]\!]_{\mathrm{M}}$ after reading ς does not belong to the vector space generated by V, then:
        i. For each $a \in \mathit{NameReal}_{P_1,P_2}$, add $\varsigma \circ a$ to the queue.
        ii. Add the state probability vector to V.

4. Build a three-valued state vector $\boldsymbol{u}$ whose generic element is:
   a. 0 if it corresponds to a nonaccepting state.
   b. 1 if it corresponds to an accepting state of $[\![P_1]\!]_{\mathrm{M}}$.
   c. -1 if it corresponds to an accepting state of $[\![P_2]\!]_{\mathrm{M}}$.
5. For each $\boldsymbol{v} \in V$, check whether $\boldsymbol{v} \cdot \boldsymbol{u}^{\mathrm{T}} = 0$.
6. Return yes/no depending on whether all the checks performed in the previous step have been successful or not.

- The time complexity of the overall algorithm is $O(n^5)$.

- $\sim_{\mathrm{MT}}$ induces an exact aggregation called T-lumping.

- T-lumping is strictly coarser than ordinary lumping and graphically definable as follows (name-abstracting axiom schema characterizing $\sim_{\mathrm{MT}}$):

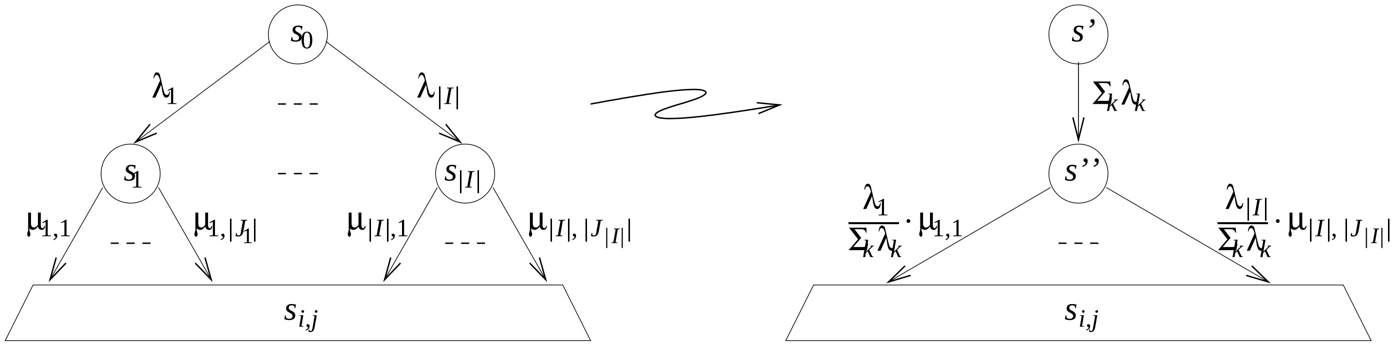

where for all $i_1, i_2 \in I$:

$$\sum_{j \in J_{i_1}} \mu_{i_1,j} = \sum_{j \in J_{i_2}} \mu_{i_2,j}$$

- Exact aggregation not previously known in the CTMC field, but entirely characterizable in a process algebraic framework like ordinary lumping.

- Two Markovian testing equivalent process terms in $\mathbb{P}_{\mathrm{M,pc}}$ are guaranteed to possess the same performance characteristics.

Part VI:
Markovian Trace Equivalence

Equivalence Definition

- Two process terms are equivalent if they can perform computations with the same functional and performance characteristics.

- Test passing replaced by trace acceptance (linear, unconstrained environment).

- Was the trace accepted?
  If so, with which probability?
  And how long did it take to accept the trace?

- Comparison of process term probabilities of performing trace-compatible computations within arbitrary time upper bounds.

- Branching points in process term behavior are all overridden.

- A computation of a process term $P \in \mathbb{P}_{\mathrm{M}}$ is a sequence of transitions that can be executed starting from P.
- The *length* of a computation is given by the number of its transitions.
- $\mathcal{C}_{\mathrm{f}}(P)$: multiset of finite-length computations of P.
- Two distinct computations are *independent* of each other iff neither is a proper prefix of the other one.
- Focus on finite multisets of independent, finite-length computations.
- Attributes of a finite-length computation:
  - ⊙ *Trace.*
  - ⊙ *Probability.*
  - ⊙ *Duration.*

- Given a set of sequences, we use:
  - Operator $\_ \circ \_$ for sequence concatenation.
  - Operator $|\_|$ for sequence length.
- The concrete trace associated with the execution of $c \in \mathcal{C}_{\mathrm{f}}(P)$ is the sequence of action names labeling the transitions of c:

$$\boxed{\mathit{trace}_{\mathrm{c}}(c) \; = \; \begin{cases} \varepsilon & \text{if } |c| = 0 \\ a \circ \mathit{trace}_{\mathrm{c}}(c') & \text{if } c \equiv P \xrightarrow{a, \tilde{\lambda}}_{\mathrm{M}} c' \end{cases}}$$

- We denote by $\mathit{trace}(c)$ the visible part of $\mathit{trace}_{\mathrm{c}}(c)$, i.e., the subsequence of $\mathit{trace}_{\mathrm{c}}(c)$ obtained by removing all the occurrences of τ.

- For the quantitative attributes, we assume $P \in \mathbb{P}_{\mathrm{M,pc}}$.

- The probability of executing $c \in \mathcal{C}_{\mathrm{f}}(P)$ is the product of the execution probabilities of the transitions of c:

$$\boxed{prob(c) \; = \; \begin{cases} 1 & \text{if } |c| = 0 \\ \frac{\lambda}{rate_{\mathrm{t}}(P,0)} \cdot prob(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_{\mathrm{M}} c' \end{cases}}$$

- Probability of executing a computation in $C \subseteq \mathcal{C}_{\mathrm{f}}(P)$:

$$\boxed{prob(C) \; = \; \sum_{c \in C} prob(c)}$$

  assuming that C is finite and all of its computations are independent.

- The stepwise average duration of $c \in \mathcal{C}_{\mathrm{f}}(P)$ is the sequence of average sojourn times in the states traversed by c:

$$\boxed{\mathit{time}_{\mathrm{a}}(c) \;=\; \begin{cases} \varepsilon & \text{if } |c| = 0 \\ \frac{1}{\mathit{rate}_{\mathrm{t}}(P,0)} \circ \mathit{time}_{\mathrm{a}}(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_{\mathrm{M}} c' \end{cases}}$$

- Multiset of computations in $C \subseteq \mathcal{C}_{\mathrm{f}}(P)$ whose stepwise average duration is not greater than $\theta \in (\mathbb{R}_{>0})^*$:

$$\boxed{C_{\leq \theta} \;=\; \{\!| \, c \in C \mid |c| \leq |\theta| \wedge \forall i = 1, \dots, |c|.\, \mathit{time}_{\mathrm{a}}(c)[i] \leq \theta[i] \, |\!\}}$$

- C^l: multiset of computations in $C \subseteq \mathcal{C}_{\mathrm{f}}(P)$ having length $l \in \mathbb{N}$.

- The stepwise duration of $c \in \mathcal{C}_{\mathrm{f}}(P)$ is the sequence of random variables quantifying the sojourn times in the states traversed by c:

$$\boxed{\mathit{time}_{\mathrm{d}}(c) \;=\; \begin{cases} \varepsilon & \text{if } |c| = 0 \\ \mathit{Exp}_{\mathit{rate}_{\mathrm{t}}(P,0)} \circ \mathit{time}_{\mathrm{d}}(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_{\mathrm{M}} c' \end{cases}}$$

- Probability distribution of executing a computation in $C \subseteq \mathcal{C}_{\mathrm{f}}(P)$ within a sequence $\theta \in (\mathbb{R}_{>0})^{*}$ of time units:

$$\boxed{\mathit{prob}_{\mathrm{d}}(C, \theta) \;=\; \sum_{c \in C}^{|c| \le |\theta|} \mathit{prob}(c) \cdot \prod_{i=1}^{|c|} \Pr\{\mathit{time}_{\mathrm{d}}(c)[i] \le \theta[i]\}}$$

  assuming that C is finite and all of its computations are independent.

- Factor $\Pr\{\mathit{time}_{\mathrm{d}}(c)[i] \le \theta[i]\} = 1 - \mathrm{e}^{-\theta[i] / \mathit{time}_{\mathrm{a}}(c)[i]}$ stems from the cumulative distribution function of the exponentially distributed random variable $\mathit{time}_{\mathrm{d}}(c)[i]$ (whose expected value is $\mathit{time}_{\mathrm{a}}(c)[i]$).

- Why not summing up sojourn times? (standard duration instead of stepwise one)
- Consider process terms ($\lambda \neq \mu$, $b \neq d$, identical nonmaximal computations):

$$<g,\gamma>.<a,\lambda>.<b,\mu>.\underline{0} + <g,\gamma>.<a,\mu>.<d,\lambda>.\underline{0}$$
$$<g,\gamma>.<a,\lambda>.<d,\mu>.\underline{0} + <g,\gamma>.<a,\mu>.<b,\lambda>.\underline{0}$$

- Maximal computations of the first term:

$$c_{1,1} \quad \equiv \quad \cdot \xrightarrow{g,\gamma}_{\mathrm{M}} \cdot \xrightarrow{a,\lambda}_{\mathrm{M}} \cdot \xrightarrow{b,\mu}_{\mathrm{M}} \cdot$$
$$c_{1,2} \quad \equiv \quad \cdot \xrightarrow{g,\gamma}_{\mathrm{M}} \cdot \xrightarrow{a,\mu}_{\mathrm{M}} \cdot \xrightarrow{d,\lambda}_{\mathrm{M}} \cdot$$

- Maximal computations of the second term:

$$c_{2,1} \quad \equiv \quad \cdot \xrightarrow{g,\gamma}_{\mathrm{M}} \cdot \xrightarrow{a,\lambda}_{\mathrm{M}} \cdot \xrightarrow{d,\mu}_{\mathrm{M}} \cdot$$
$$c_{2,2} \quad \equiv \quad \cdot \xrightarrow{g,\gamma}_{\mathrm{M}} \cdot \xrightarrow{a,\mu}_{\mathrm{M}} \cdot \xrightarrow{b,\lambda}_{\mathrm{M}} \cdot$$

- Same sum of average sojourn times $\frac{1}{2\cdot\gamma} + \frac{1}{\lambda} + \frac{1}{\mu}$ and $\frac{1}{2\cdot\gamma} + \frac{1}{\mu} + \frac{1}{\lambda}$ but ...
- *... an external observer would be able to distinguish between the two terms by taking note of the instants at which the actions are performed.*

- Comparing probabilities of accepting a trace within a time upper bound.
- $c \in \mathcal{C}_{\mathrm{f}}(P)$ is compatible with $\alpha \in (Name_{\mathrm{v}})^*$ iff:

$$\boxed{trace(c) \; = \; \alpha}$$

- $\mathcal{CC}(P, \alpha)$: multiset of computations in $\mathcal{C}_{\mathrm{f}}(P)$ compatible with α.
- If P has no exponentially timed τ-actions:
  - $\odot$ All the computations in $\mathcal{CC}(P, \alpha)$ are independent.
  - $\odot$ The multiset $\mathcal{CC}(P, \alpha)$ is finite.
- Same properties for $\mathcal{CC}_{\leq \theta}(P, \alpha)$.
- If there are exponentially timed τ-actions:
  - $\odot$ Are the computations in $\mathcal{CC}_{\leq \theta}(P, \alpha)$ independent of each other?
  - $\odot$ How to distinguish among process terms having only exponentially timed τ-actions, like $<\tau, \lambda>.\underline{0}$ and $<\tau, \mu>.\underline{0}$ with $\lambda > \mu$?

- *Consider subsets of* $\mathcal{CC}_{\leq\theta}(P, \alpha)$ *including all the trace-compatible computations of the same length.*

- They are $\mathcal{CC}^{l}_{\leq\theta}(P, \alpha)$ for $0 \leq l \leq |\theta|$.

- $\mathcal{CC}^{|\theta|}_{\leq\theta}(P, \alpha)$ is enough as shorter trace-compatible computations can be taken into account when imposing prefixes of θ as time upper bounds.

- Process terms having only exponentially timed τ-actions are compared after giving them the possibility of executing the same number of τ-actions ($\lambda > \mu \Rightarrow \frac{1}{\lambda} < \frac{1}{\mu}$):

$$\mathit{prob}(\mathcal{CC}^{1}_{\leq\frac{1}{\lambda}}(<\tau, \lambda>.\underline{0}, \varepsilon)) \;=\; 1 \;\neq\; 0 \;=\; \mathit{prob}(\mathcal{CC}^{1}_{\leq\frac{1}{\lambda}}(<\tau, \mu>.\underline{0}, \varepsilon))$$

- $P_1 \in \mathbb{P}_{\mathrm{M,pc}}$ is Markovian trace equivalent to $P_2 \in \mathbb{P}_{\mathrm{M,pc}}$, written $P_1 \sim_{\mathrm{MTr}} P_2$, iff for all traces $\alpha \in (\mathit{Name}_{\mathrm{v}})^*$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of average amounts of time:

$$\boxed{\mathit{prob}(\mathcal{CC}^{|\theta|}_{\leq\theta}(P_1, \alpha)) \;=\; \mathit{prob}(\mathcal{CC}^{|\theta|}_{\leq\theta}(P_2, \alpha))}$$

- $\sim_{\mathrm{MTr}}$ is strictly finer than classical trace equivalence ($a \neq b$ and $\lambda \neq \mu$):

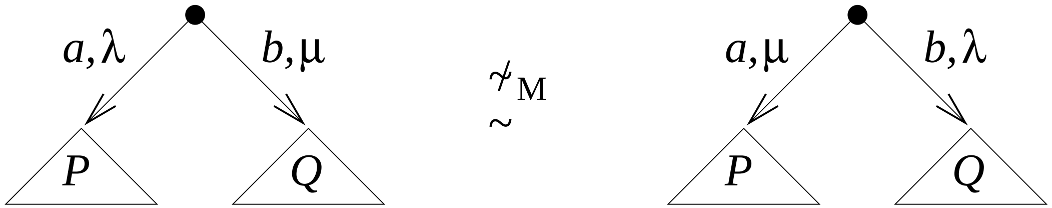

- $\sim_{\mathrm{MTr}}$ is strictly finer than probabilistic trace equivalence:

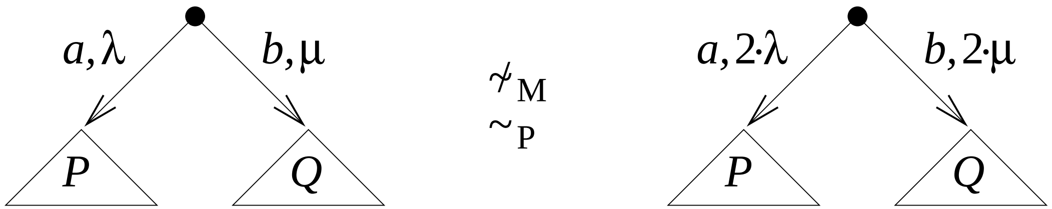

- $\sim_{\mathrm{MTr}}$ is strictly coarser than $\sim_{\mathrm{MT}}$ ($b \neq c$):

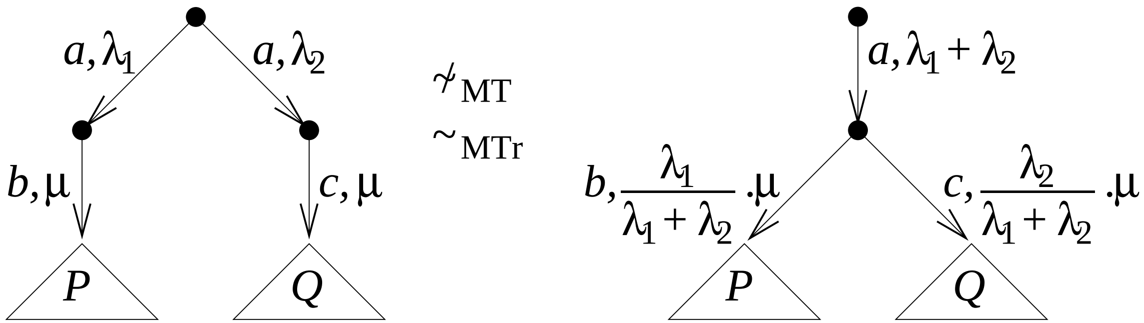

- **Running example** ($\sim_{\rm MTr}$):
  - Concurrent implementation with two independent one-pos. buffers:

$$\begin{array}{rcl} PC^{\rm M}_{\rm conc,2} & \triangleq & Prod^{\rm M} \parallel_{\{deposit\}} (Buff^{\rm M} \parallel_{\emptyset} Buff^{\rm M}) \parallel_{\{withdraw\}} Cons^{\rm M} \\ Prod^{\rm M} & \triangleq & <deposit, \lambda>.Prod^{\rm M} \\ Buff^{\rm M} & \triangleq & <deposit, *_1>.<withdraw, *_1>.Buff^{\rm M} \\ Cons^{\rm M} & \triangleq & <withdraw, \mu>.Cons^{\rm M} \end{array}$$

  - All the actions occurring in the buffer are passive, consistent with the fact that the buffer is a passive entity.
  - Is $PC^{\rm M}_{\rm conc,2}$ a correct implementation of $ProdCons^{\rm M}_{0/2}$?
  - It turns out that $PC^{\rm M}_{\rm conc,2} \sim_{\rm MTr} ProdCons^{\rm M}_{0/2}$.

- Here are the underlying labeled multitransition systems:

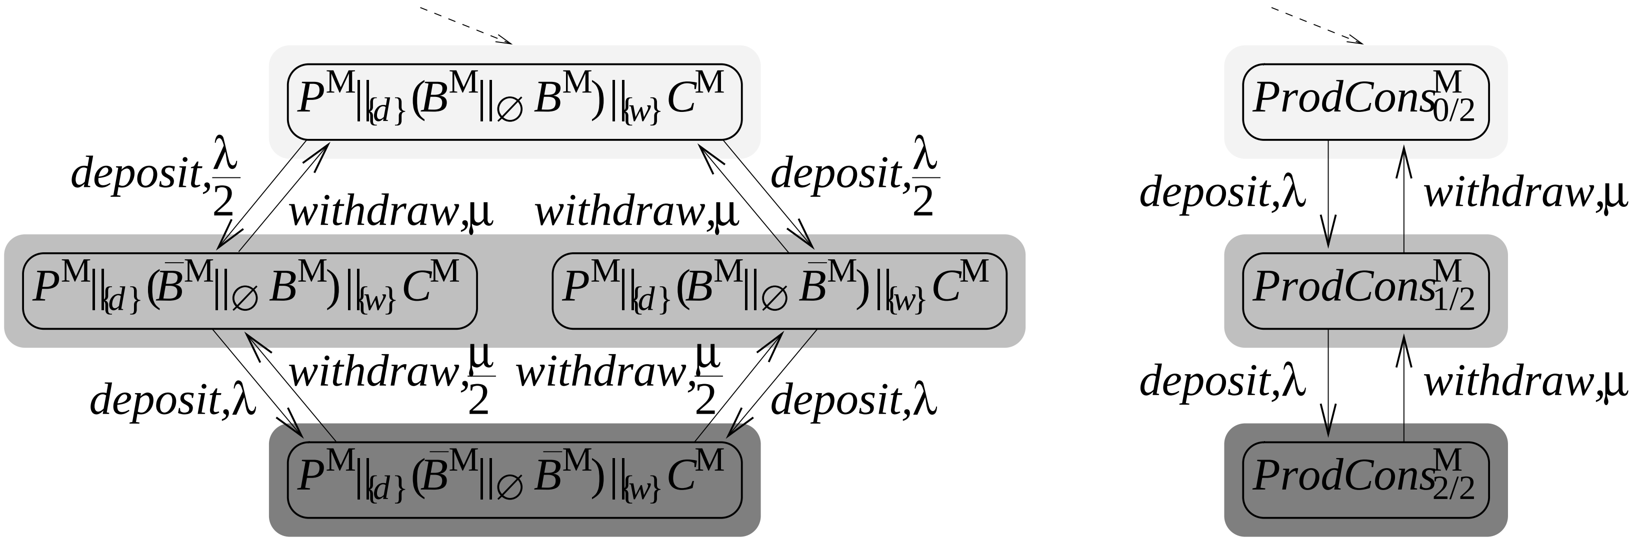

- The initial state on the left-hand side has both outgoing transitions labeled with $\lambda/2$, not λ.

- The bottom state on the left-hand side has both outgoing transitions labeled with $\mu/2$, not μ.

- The only sequences of visible actions that the two systems are able to perform are the prefixes of the strings complying with:

$$(\mathit{deposit} \circ (\mathit{deposit} \circ \mathit{withdraw})^* \circ \mathit{withdraw})^*$$

- The only significant traces to be considered are those coinciding with such prefixes.
- Any two computations of $\mathit{ProdCons}^{\mathrm{M}}_{0/2}$ and $\mathit{PC}^{\mathrm{M}}_{\mathrm{conc},2}$ compatible with such an α traverse states that pairwise have the same average sojourn time.
- Therefore the stepwise average durations of the considered computations are identical.

- $\odot$ Four basic cases for the execution probabilities of $\mathcal{CC}(PC^{\mathrm{M}}_{\mathrm{conc},2}, \alpha)$ and $\mathcal{CC}(ProdCons^{\mathrm{M}}_{0/2}, \alpha)$:
  - $*$ If $\alpha \equiv deposit$, then for both sets of computations the execution probability is 1.
  - $*$ If $\alpha \equiv deposit \circ withdraw$, then for both sets of computations the execution probability is $\frac{\mu}{\lambda+\mu}$.
  - $*$ If $\alpha \equiv deposit \circ deposit$, then for both sets of computations the execution probability is $\frac{\lambda}{\lambda+\mu}$.
  - $*$ If $\alpha \equiv deposit \circ deposit \circ withdraw$, then for both sets of computations the execution probability is $\frac{\lambda}{\lambda+\mu}$.

Conditions and Characterizations

- In order for $P_1 \sim_{\mathrm{MTr}} P_2$, it is necessary that for all $c_k \in \mathcal{C}_{\mathrm{f}}(P_k)$, $k \in \{1,2\}$, there exists $c_h \in \mathcal{C}_{\mathrm{f}}(P_h)$, $h \in \{1,2\} - \{k\}$, such that:

$$\boxed{\begin{array}{rcl} trace_{\mathrm{c}}(c_k) & = & trace_{\mathrm{c}}(c_h) \\ time_{\mathrm{a}}(c_k) & = & time_{\mathrm{a}}(c_h) \end{array}}$$

  and for all $i \in \{0, \ldots, |c_k|\}$:

$$\boxed{rate_{\mathrm{t}}(P_k^i, 0) \;=\; rate_{\mathrm{t}}(P_h^i, 0)}$$

  with P_k^i (resp. P_h^i) being the i-th state traversed by c_k (resp. c_h).

- Process terms satisfying the necessary condition that are not Markovian trace equivalent ($\lambda_1 + \lambda_2 = \lambda_1' + \lambda_2'$ with $\lambda_1 \neq \lambda_1', \lambda_2 \neq \lambda_2'$ and $b \neq c$ or $\mu \neq \gamma$):

$$<a, \lambda_1>.<b, \mu>.\underline{0} + <a, \lambda_2>.<c, \gamma>.\underline{0}$$

$$<a, \lambda_1'>.<b, \mu>.\underline{0} + <a, \lambda_2'>.<c, \gamma>.\underline{0}$$

- $\sim_{\mathrm{MTr}}$ has an alternative characterization showing that its discriminating power does not change if we consider the probability distribution of accepting traces within arbitrary sequences of amounts of time.

- Considering the (more accurate) stepwise durations of trace-compatible computations leads to the same equivalence as considering the (easier to work with) stepwise average durations.

- $P_1 \in \mathbb{P}_{\mathrm{M,pc}}$ is Markovian distribution-trace equivalent to $P_2 \in \mathbb{P}_{\mathrm{M,pc}}$, written $P_1 \sim_{\mathrm{MTr,d}} P_2$, iff for all traces $\alpha \in (Name_{\mathrm{v}})^*$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of amounts of time:

$$\boxed{prob_{\mathrm{d}}(\mathcal{CC}^{|\theta|}(P_1, \alpha), \theta) \; = \; prob_{\mathrm{d}}(\mathcal{CC}^{|\theta|}(P_2, \alpha), \theta)}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MTr,d}} P_2 \iff P_1 \sim_{\mathrm{MTr}} P_2}$$

Equivalence Properties

- $\sim_{\rm MTr}$ is a congruence over $\mathbb{P}_{\rm M,pc}$ w.r.t. all the dynamic operators.
- Let $P_1, P_2 \in \mathbb{P}_{\rm M,pc}$. Whenever $P_1 \sim_{\rm MTr} P_2$, then:

$$\boxed{\begin{array}{c} <a, \lambda>.P_1 \;\sim_{\rm MTr}\; <a, \lambda>.P_2 \\ P_1 + P \;\sim_{\rm MTr}\; P_2 + P \qquad\qquad P + P_1 \;\sim_{\rm MTr}\; P + P_2 \end{array}}$$

- *Not a congruence with respect to parallel composition.*
- For instance, the Markovian trace equivalent process terms $(b \neq c)$:

$$<a, \lambda_1>.<b, \mu>.\underline{0} + <a, \lambda_2>.<c, \mu>.\underline{0}$$

$$<a, \lambda_1 + \lambda_2>.(<b, \tfrac{\lambda_1}{\lambda_1+\lambda_2} \cdot \mu>.\underline{0} + <c, \tfrac{\lambda_2}{\lambda_1+\lambda_2} \cdot \mu>.\underline{0})$$

  are distinguished in the following context:

$$\_ \, \|_{\{a,b,c\}} \; <a, *_1>.<b, *_1>.\underline{0}$$

  by the following trace:

$$\alpha \;\equiv\; a \circ b$$

- $\sim_{\mathrm{MTr}}$ has a sound and complete axiomatization over the set $\mathbb{P}_{\mathrm{M,pc,dyn}}$ of process terms of $\mathbb{P}_{\mathrm{M,pc}}$ comprising only dynamic operators.

- The axioms for $\sim_{\mathrm{MT}}$ are sound but not complete for $\sim_{\mathrm{MTr}}$ $(b \neq c)$:

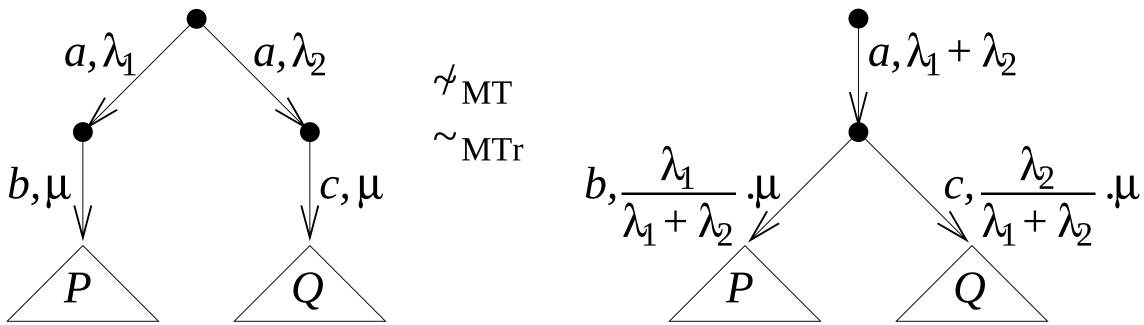

- *Possibility of deferring choices related to branches starting with actions having the same name (a) that are immediately followed by process terms having the same total exit rate (μ).*

- Names and total rates of the initial actions of such derivative terms can be different in the various branches.

- Basic laws (identical to those for $\sim_{\rm MT}$):

$$\begin{array}{|lrcl|}
\hline
(\mathcal{X}_{\rm MTr,1}) & P_1 + P_2 & = & P_2 + P_1 \\
(\mathcal{X}_{\rm MTr,2}) & (P_1 + P_2) + P_3 & = & P_1 + (P_2 + P_3) \\
(\mathcal{X}_{\rm MTr,3}) & P + \underline{0} & = & P \\
\hline
\end{array}$$

- Characterizing law (subsumes $\sim_{\rm MT}$ characterizing law):

$$(\mathcal{X}_{\rm MTr,4}) \quad \sum_{i \in I} {<}a, \lambda_i{>}. \sum_{j \in J_i} {<}b_{i,j}, \mu_{i,j}{>}.P_{i,j} \;=\; {<}a, \sum_{k \in I} \lambda_k{>}. \sum_{i \in I} \sum_{j \in J_i} {<}b_{i,j}, \frac{\lambda_i}{\Sigma_{k \in I} \lambda_k} \cdot \mu_{i,j}{>}.P_{i,j}$$

if: I is a finite index set with $|I| \geq 2$;
for all $i \in I$, index set J_i is finite and its summation is $\underline{0}$ if $J_i = \emptyset$;
for all $i_1, i_2 \in I$:

$$\sum_{j \in J_{i_1}} \mu_{i_1,j} \;=\; \sum_{j \in J_{i_2}} \mu_{i_2,j}$$

- $DED(\mathcal{X}_{\mathrm{MTr}})$: deduction system based on all the previous axioms plus:
  - ⊙ Reflexivity: $\mathcal{X}_{\mathrm{MTr}} \vdash P = P$.
  - ⊙ Symmetry: $\mathcal{X}_{\mathrm{MTr}} \vdash P_1 = P_2 \implies \mathcal{X}_{\mathrm{MTr}} \vdash P_2 = P_1$.
  - ⊙ Transitivity: $\mathcal{X}_{\mathrm{MTr}} \vdash P_1 = P_2 \wedge \mathcal{X}_{\mathrm{MTr}} \vdash P_2 = P_3 \implies \mathcal{X}_{\mathrm{MTr}} \vdash P_1 = P_3$.
  - ⊙ Substitutivity: $\mathcal{X}_{\mathrm{MTr}} \vdash P_1 = P_2 \implies \mathcal{X}_{\mathrm{MTr}} \vdash {<}a, \lambda{>}.P_1 = {<}a, \lambda{>}.P_2 \wedge \dots$

- The deduction system $DED(\mathcal{X}_{\mathrm{MTr}})$ is sound and complete for $\sim_{\mathrm{MTr}}$ over $\mathbb{P}_{\mathrm{M,pc,dyn}}$; i.e., for all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc,dyn}}$:

$$\boxed{\mathcal{X}_{\mathrm{MTr}} \vdash P_1 = P_2 \iff P_1 \sim_{\mathrm{MTr}} P_2}$$

- $\sim_{\rm MTr}$ has a modal logic characterization over $\mathbb{P}_{\rm M,pc}$ based on a variant of the Hennessy-Milner logic.

- Neither negation nor any binary connective is included, only diamond (decreased discriminating power with respect to $\sim_{\rm MT}$).

- Syntax of the modal language $\mathcal{ML}_{\rm MTr}$ ($a \in Name_{\rm v}$):

$$\boxed{\phi \ ::= \ \text{true} \mid \langle a \rangle \phi}$$

- No quantitative decorations in the syntax because the focus is on entire computations rather than on step-by-step behavior mimicking, but ...

- ...replacement of the boolean satisfaction relation with a quantitative interpretation function measuring the probability with which a process term satisfies a formula quickly enough on average.

- Interpretation of $\mathcal{ML}_{\mathrm{MTr}}$ over $\mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{[\![\phi]\!]^{|\theta|}_{\mathrm{MTr}}(P,\theta) \;=\; \begin{cases} 0 & \text{if } |\theta| = 0 \wedge \phi \not\equiv \text{true} \\ & \quad\text{or } |\theta| > 0 \wedge \mathit{rate}_{\mathrm{t}}(P,0) = 0 \\ 1 & \text{if } |\theta| = 0 \wedge \phi \equiv \text{true} \end{cases}}$$

otherwise:

$$[\![\text{true}]\!]^{|t\circ\theta|}_{\mathrm{MTr}}(P, t\circ\theta) \;=\; \begin{cases} \sum\limits_{P \xrightarrow{\tau,\lambda}_{\mathrm{M}} P'} \frac{\lambda}{\mathit{rate}_{\mathrm{t}}(P,0)} \cdot [\![\text{true}]\!]^{|\theta|}_{\mathrm{MTr}}(P',\theta) & \text{if } \frac{1}{\mathit{rate}_{\mathrm{t}}(P,0)} \leq t \\ 0 & \text{if } \frac{1}{\mathit{rate}_{\mathrm{t}}(P,0)} > t \end{cases}$$

$$[\![\langle a\rangle\phi]\!]^{|t\circ\theta|}_{\mathrm{MTr}}(P, t\circ\theta) \quad = \quad \begin{cases} \displaystyle\sum_{P \xrightarrow{a,\lambda}_{\mathrm{M}} P'} \frac{\lambda}{rate_{\mathrm{t}}(P,0)} \cdot [\![\phi]\!]^{|\theta|}_{\mathrm{MTr}}(P',\theta) + \\ \displaystyle\qquad \sum_{P \xrightarrow{\tau,\lambda}_{\mathrm{M}} P'} \frac{\lambda}{rate_{\mathrm{t}}(P,0)} \cdot [\![\langle a\rangle\phi]\!]^{|\theta|}_{\mathrm{MTr}}(P',\theta) & \text{if } \frac{1}{rate_{\mathrm{t}}(P,0)} \leq t \\ 0 & \text{if } \frac{1}{rate_{\mathrm{t}}(P,0)} > t \end{cases}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MTr}} P_2 \iff \forall \phi \in \mathcal{ML}_{\mathrm{MTr}}.\, \forall \theta \in (\mathbb{R}_{>0})^*.\, [\![\phi]\!]^{|\theta|}_{\mathrm{MTr}}(P_1,\theta) = [\![\phi]\!]^{|\theta|}_{\mathrm{MTr}}(P_2,\theta)}$$

- $\sim_{\mathrm{MTr}}$ is decidable in polynomial time over the set $\mathbb{P}_{\mathrm{M,pc,fin}}$ of finite-state process terms of $\mathbb{P}_{\mathrm{M,pc}}$.
- Reworking of Tzeng algorithm for probabilistic language equivalence.
- Given two process terms, their name-labeled CTMCs are Markovian trace equivalent iff the corresponding embedded name-labeled DTMCs are probabilistic trace equivalent.
- Probabilistic trace equivalence is decidable in polynomial time through the algorithm for probabilistic language equivalence.
- Markovian trace equivalence and probabilistic trace equivalence coincide on corresponding models if the total exit rate of each state of a name-labeled CTMC is encoded inside the names of all transitions departing from that state in the associated embedded DTMC.

- Steps of the algorithm for checking whether $P_1 \sim_{\mathrm{MTr}} P_2$:
  1. Transform $[\![P_1]\!]_{\mathrm{M}}$ and $[\![P_2]\!]_{\mathrm{M}}$ into their corresponding embedded discrete-time versions:
     a. Divide the rate of each transition by the total exit rate of its source state.
     b. Augment the name of each transition with the total exit rate of its source state.
  2. Check whether the discrete-time versions of $[\![P_1]\!]_{\mathrm{M}}$ and $[\![P_2]\!]_{\mathrm{M}}$ are probabilistic language equivalent when all of their states are considered as accepting states.
  3. Return yes/no depending on whether the check performed in the previous step has been successful or not.

- Tzeng algorithm for probabilistic language equivalence visits in breadth-first order the tree containing a node for each possible string and studies the linear independence of the state probability vectors associated with a finite subset of the tree nodes.
- Refinement of step 2:
  1. Create an empty set V of state probability vectors.
  2. Create a queue whose only element is the empty string ε.
  3. While the queue is not empty:
     a. Remove the first element from the queue, say string ς.
     b. If the state probability vector of the discrete-time versions of $[\![P_1]\!]_{\mathrm{M}}$ and $[\![P_2]\!]_{\mathrm{M}}$ after reading ς does not belong to the vector space generated by V, then:
        i. For each $a \in \mathit{NameReal}_{P_1,P_2}$, add $\varsigma \circ a$ to the queue.
        ii. Add the state probability vector to V.

4. Build a three-valued state vector $\boldsymbol{u}$ whose generic element is:
   a. 0 if it corresponds to a nonaccepting state.
   b. 1 if it corresponds to an accepting state of $[\![P_1]\!]_{\mathrm{M}}$.
   c. -1 if it corresponds to an accepting state of $[\![P_2]\!]_{\mathrm{M}}$.
5. For each $\boldsymbol{v} \in V$, check whether $\boldsymbol{v} \cdot \boldsymbol{u}^{\mathrm{T}} = 0$.
6. Return yes/no depending on whether all the checks performed in the previous step have been successful or not.

- The time complexity of the overall algorithm is $O(n^4)$.

- $\sim_{\mathrm{MTr}}$ induces an exact aggregation called T-lumping.

- T-lumping is strictly coarser than ordinary lumping and graphically definable as follows (name-abstracting axiom schema characterizing $\sim_{\mathrm{MTr}}$):

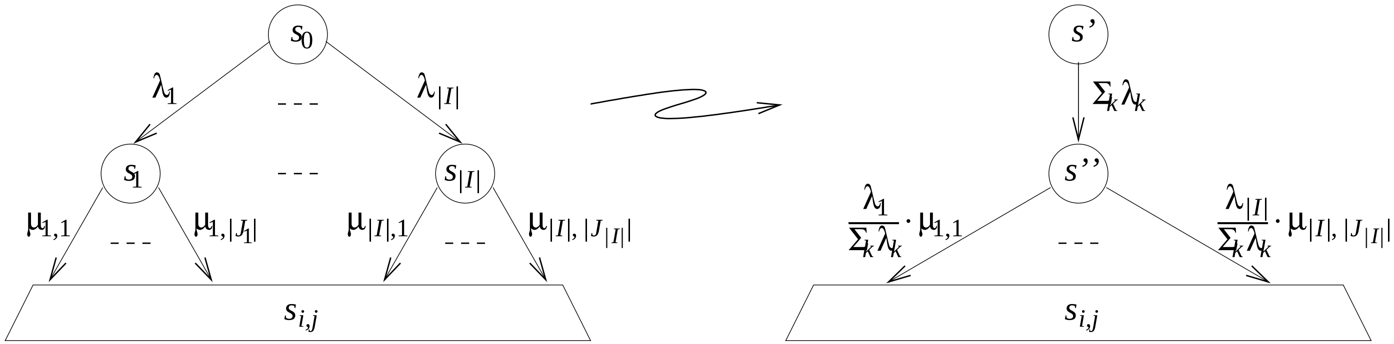

where for all $i_1, i_2 \in I$:

$$\sum_{j \in J_{i_1}} \mu_{i_1,j} = \sum_{j \in J_{i_2}} \mu_{i_2,j}$$

- Exact aggregation not previously known in the CTMC field, but entirely characterizable in a process algebraic framework like ordinary lumping.

- Two Markovian trace equivalent process terms in $\mathbb{P}_{\mathrm{M,pc}}$ are guaranteed to possess the same performance characteristics.

Part VII:
The Markovian Spectrum

Markovian Versions of Similarity Variants

- Simulation equivalence is the intersection of two simulation preorders each considering stepwise behavior mimicking in one single direction (less discriminating than bisimulation equivalence in the nondeterministic setting).

- For *Markovian simulation equivalence* and *Markovian ready-simulation equivalence*, we need to take rates into account as well.

- Separate time and probability information subsumed by rates and view the destination of a transition as a probability distribution.

- Hence lift the simulation relation from states to distributions on states, precisely the next-state distributions encoded by $prob_{\mathrm{e}}(P, a, 0, .)$.

- Use *weight functions* for relating pairs of distributions in a way that takes into account the simulation relation on states and preserves the probability mass associated with each state by the original distributions.

- A distribution on a countable set S is a function $d : S \rightarrow \mathbb{R}_{[0,1]}$ such that:

$$\boxed{\sum_{s \in S} d(s) \leq 1}$$

- A distribution d on S is called a probability distribution iff:

$$\boxed{d(\perp) \triangleq 1 - \sum_{s \in S} d(s) = 0}$$

- $Distr(S)$: set of distributions on S.

- A function $\omega : (S_1 \cup \{\perp\}) \times (S_2 \cup \{\perp\}) \to \mathbb{R}_{[0,1]}$ is a weight function for $d_1 \in \mathit{Distr}(S_1)$ and $d_2 \in \mathit{Distr}(S_2)$ with respect to $\mathcal{R} \subseteq S_1 \times S_2$ iff for all $s_1 \in S_1 \cup \{\perp\}$ and $s_2 \in S_2 \cup \{\perp\}$:

$$\boxed{\begin{array}{rcl} \omega(s_1, s_2) > 0 & \Longrightarrow & (s_1, s_2) \in \mathcal{R} \vee s_1 = \perp \\ d_1(s_1) & = & \displaystyle\sum_{s_2 \in S_2 \cup \{\perp\}} \omega(s_1, s_2) \\ d_2(s_2) & = & \displaystyle\sum_{s_1 \in S_1 \cup \{\perp\}} \omega(s_1, s_2) \end{array}}$$

- We write $d_1 \sqsubseteq_{\mathcal{R}} d_2$ in that case.

- A relation $\mathcal{S}$ over $\mathbb{P}_{\mathrm{M,pc}}$ is a Markovian simulation iff, whenever $(P_1, P_2) \in \mathcal{S}$, then for all action names $a \in \mathit{Name}$:

$$\boxed{\begin{array}{rcl} \mathit{rate}_{\mathrm{o}}(P_1, a, 0) & \leq & \mathit{rate}_{\mathrm{o}}(P_2, a, 0) \\ \mathit{prob}_{\mathrm{e}}(P_1, a, 0, .) & \sqsubseteq_{\mathcal{S}} & \mathit{prob}_{\mathrm{e}}(P_2, a, 0, .) \end{array}}$$

- Markovian simulation preorder $\preceq_{\mathrm{MS}}$ is the largest Markovian simulation.

- Markovian simulation equivalence is the kernel of Markovian simulation preorder:

$$\boxed{\sim_{\mathrm{MS}} \; = \; \preceq_{\mathrm{MS}} \cap \preceq_{\mathrm{MS}}^{-1}}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MS}} P_2 \iff P_1 \sim_{\mathrm{MB}} P_2}$$

- A relation $\mathcal{S}$ over $\mathbb{P}_{\mathrm{M,pc}}$ is a Markovian ready simulation iff, whenever $(P_1, P_2) \in \mathcal{S}$, then for all action names $a \in \mathit{Name}$:

$$\boxed{\begin{array}{rcl} \mathit{rate}_{\mathrm{o}}(P_1, a, 0) & \leq & \mathit{rate}_{\mathrm{o}}(P_2, a, 0) \\ \mathit{prob}_{\mathrm{e}}(P_1, a, 0, .) & \sqsubseteq_{\mathcal{S}} & \mathit{prob}_{\mathrm{e}}(P_2, a, 0, .) \\ \mathit{rate}_{\mathrm{o}}(P_1, a, 0) = 0 & \Longrightarrow & \mathit{rate}_{\mathrm{o}}(P_2, a, 0) = 0 \end{array}}$$

- Markovian ready-simulation preorder $\preceq_{\mathrm{MRS}}$ is the largest Markovian ready simulation.

- Markovian ready-simulation equivalence is the kernel of Markovian ready-simulation preorder:

$$\boxed{\sim_{\mathrm{MRS}} \; = \; \preceq_{\mathrm{MRS}} \cap \preceq_{\mathrm{MRS}}^{-1}}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MRS}} P_2 \iff P_1 \sim_{\mathrm{MB}} P_2}$$

Markovian Versions of Trace Equivalence Variants

- *Completed trace*: trace ending up in a deadlock state.
- *Failure set*: set of names of visible actions that cannot be executed in a certain state.
- *Ready set*: set of names of all the visible actions that must be executable in a certain state.
- *Failure trace*: trace extended at each step with a failure set.
- *Ready trace*: trace extended at each step with a ready set.
- Variants more discriminating in the nondeterministic setting:
  - Completed traces for gaining deadlock sensitivity.
  - Failures for reasoning about safety.
  - Readies for reasoning about liveness.

- $c \in \mathcal{C}_{\mathrm{f}}(P)$ is a maximal computation compatible with $\alpha \in (\mathit{Name}_{\mathrm{v}})^*$ iff $c \in \mathcal{CC}(P, \alpha)$ and the last state reached by c has no outgoing transitions.

- $\mathcal{MCC}(P, \alpha)$: multiset of maximal computations in $\mathcal{C}_{\mathrm{f}}(P)$ compatible with α.

- $P_1 \in \mathbb{P}_{\mathrm{M,pc}}$ is Markovian completed-trace equivalent to $P_2 \in \mathbb{P}_{\mathrm{M,pc}}$, written $P_1 \sim_{\mathrm{MCTr}} P_2$, iff for all traces $\alpha \in (\mathit{Name}_{\mathrm{v}})^*$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of average amounts of time:

$$\boxed{\begin{array}{rcl} \mathit{prob}(\mathcal{CC}^{|\theta|}_{\leq \theta}(P_1, \alpha)) & = & \mathit{prob}(\mathcal{CC}^{|\theta|}_{\leq \theta}(P_2, \alpha)) \\ \mathit{prob}(\mathcal{MCC}^{|\theta|}_{\leq \theta}(P_1, \alpha)) & = & \mathit{prob}(\mathcal{MCC}^{|\theta|}_{\leq \theta}(P_2, \alpha)) \end{array}}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MCTr}} P_2 \iff P_1 \sim_{\mathrm{MTr}} P_2}$$

- $c \in \mathcal{C}_{\mathrm{f}}(P)$ is compatible with the failure pair $\beta \equiv (\alpha, F) \in (Name_{\mathrm{v}})^* \times 2^{Name_{\mathrm{v}}}$ iff $c \in \mathcal{CC}(P, \alpha)$ and the last state reached by c cannot perform any visible action whose name belongs to the failure set F.

- $\mathcal{FCC}(P, \beta)$: multiset of computations in $\mathcal{C}_{\mathrm{f}}(P)$ compatible with β.

- $P_1 \in \mathbb{P}_{\mathrm{M,pc}}$ is Markovian failure equivalent to $P_2 \in \mathbb{P}_{\mathrm{M,pc}}$, written $P_1 \sim_{\mathrm{MF}} P_2$, iff for all failure pairs $\beta \in (Name_{\mathrm{v}})^* \times 2^{Name_{\mathrm{v}}}$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of average amounts of time:

$$\boxed{prob(\mathcal{FCC}^{|\theta|}_{\leq\theta}(P_1, \beta)) \; = \; prob(\mathcal{FCC}^{|\theta|}_{\leq\theta}(P_2, \beta))}$$

- For all $P_1, P_2 \in \mathbb{P}_{\mathrm{M,pc}}$:

$$\boxed{P_1 \sim_{\mathrm{MF}} P_2 \iff P_1 \sim_{\mathrm{MT}} P_2}$$

- $c \in \mathcal{C}_{\rm f}(P)$ is compatible with the ready pair $\rho \equiv (\alpha, R) \in (Name_{\rm v})^* \times 2^{Name_{\rm v}}$ iff $c \in \mathcal{CC}(P, \alpha)$ and the set of names of visible actions executable by the last state reached by c coincides with the ready set R.

- $\mathcal{RCC}(P, \rho)$: multiset of computations in $\mathcal{C}_{\rm f}(P)$ compatible with ρ.

- $P_1 \in \mathbb{P}_{\rm M,pc}$ is Markovian ready equivalent to $P_2 \in \mathbb{P}_{\rm M,pc}$, written $P_1 \sim_{\rm MR} P_2$, iff for all ready pairs $\rho \in (Name_{\rm v})^* \times 2^{Name_{\rm v}}$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of average amounts of time:

$$\boxed{prob(\mathcal{RCC}^{|\theta|}_{\leq \theta}(P_1, \rho)) \; = \; prob(\mathcal{RCC}^{|\theta|}_{\leq \theta}(P_2, \rho))}$$

- For all $P_1, P_2 \in \mathbb{P}_{\rm M,pc}$:

$$\boxed{P_1 \sim_{\rm MR} P_2 \iff P_1 \sim_{\rm MT} P_2}$$

- $c \in \mathcal{C}_{\mathrm{f}}(P)$ is compatible with the failure trace $\zeta \in (Name_{\mathrm{v}} \times 2^{Name_{\mathrm{v}}})^*$ iff c is compatible with the trace projection of ζ and each state traversed by c cannot perform any visible action whose name belongs to the corresponding failure set in the failure projection of ζ.

- $\mathcal{FTCC}(P, \zeta)$: multiset of computations in $\mathcal{C}_{\mathrm{f}}(P)$ compatible with ζ.

- $P_1 \in \mathbb{P}_{\mathrm{M,pc}}$ is Markovian failure-trace equivalent to $P_2 \in \mathbb{P}_{\mathrm{M,pc}}$, written $P_1 \sim_{\mathrm{MFTr}} P_2$, iff for all failure traces $\zeta \in (Name_{\mathrm{v}} \times 2^{Name_{\mathrm{v}}})^*$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of average amounts of time:

$$\boxed{prob(\mathcal{FTCC}^{|\theta|}_{\leq \theta}(P_1, \zeta)) \; = \; prob(\mathcal{FTCC}^{|\theta|}_{\leq \theta}(P_2, \zeta))}$$

- $c \in \mathcal{C}_{\rm f}(P)$ is compatible with the ready trace $\eta \in (Name_{\rm v} \times 2^{Name_{\rm v}})^*$ iff c is compatible with the trace projection of η and the set of names of visible actions executable by each state traversed by c coincides with the corresponding ready set in the ready projection of η.

- $\mathcal{RTCC}(P, \eta)$: multiset of computations in $\mathcal{C}_{\rm f}(P)$ compatible with η.

- $P_1 \in \mathbb{P}_{\rm M,pc}$ is Markovian ready-trace equivalent to $P_2 \in \mathbb{P}_{\rm M,pc}$, written $P_1 \sim_{\rm MRTr} P_2$, iff for all ready traces $\eta \in (Name_{\rm v} \times 2^{Name_{\rm v}})^*$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of average amounts of time:

$$\boxed{prob(\mathcal{RTCC}^{|\theta|}_{\leq \theta}(P_1, \eta)) \; = \; prob(\mathcal{RTCC}^{|\theta|}_{\leq \theta}(P_2, \eta))}$$

- For all $P_1, P_2 \in \mathbb{P}_{\rm M,pc}$:

$$\boxed{P_1 \sim_{\rm MRTr} P_2 \iff P_1 \sim_{\rm MFTr} P_2}$$

Relating Markovian Behavioral Equivalences

- Markovian linear-time/branching-time spectrum:

$$\sim_{\mathrm{MB}} = \sim_{\mathrm{MRS}} = \sim_{\mathrm{MS}} \subset$$
$$\sim_{\mathrm{MRTr}} = \sim_{\mathrm{MFTr}} \subset$$
$$\sim_{\mathrm{MR}} = \sim_{\mathrm{MF}} = \sim_{\mathrm{MT}} \subset$$
$$\sim_{\mathrm{MCTr}} = \sim_{\mathrm{MTr}}$$

- Similar to the probabilistic spectrum.
- More linear than the nondeterministic spectrum.
- The considered processes do not exhibit nondeterministic behavior.

Part VIII:
Conclusion

Summary of Results

- Comparing Markovian behavioral equivalences based on given criteria:

	congruence property	sound & complete axiomatization	modal logic characteriz.	verification complexity	exact aggreg.
$\sim_{\rm MB}$	$\surd$	$\surd$	$\surd$	$O(m \cdot \log n)$	$\surd$
$\sim_{\rm MT}$	$\surd$	$\surd$	$\surd$	$O(n^5)$	$\surd$
$\sim_{\rm MTr}$	*dynamic*	*dynamic*	$\surd$	$O(n^4)$	$\surd$

- Not only intuitively appropriate from the functional viewpoint, but also meaningful and useful from the performance standpoint.

- Aggregating the state space of a model by exploiting symmetries or reducing the state space of a model before analysis takes place without altering the performance properties to be assessed.

Open Problems

- *Markovian behavioral equivalence inducing the coarsest exact nontrivial CTMC-level aggregation?*
- *Minimization algorithms for* $\sim_{\mathrm{MT}}$ *and* $\sim_{\mathrm{MTr}}$ *(and T-lumping)?*
- *Weaker versions of* $\sim_{\mathrm{MB}}$, $\sim_{\mathrm{MT}}$, *and* $\sim_{\mathrm{MTr}}$ *that abstract from internal exponentially timed actions while preserving nontrivial exactness?*
- *Approximated versions of* $\sim_{\mathrm{MB}}$, $\sim_{\mathrm{MT}}$, *and* $\sim_{\mathrm{MTr}}$ *that relax the comparison on exit rates or execution probabilities?*
- *Uniform definitions for nondeterministic, probabilistic, and Markovian processes?*

References

[1] M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis, *"Modelling with Generalized Stochastic Petri Nets"*, John Wiley & Sons, 1995.

[2] J.C.M. Baeten and W.P. Weijland, *"Process Algebra"*, Cambridge University Press, 1990.

[3] C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen, *"Automated Performance and Dependability Evaluation Using Model Checking"*, in *Performance Evaluation of Complex Systems: Techniques and Tools*, Springer, LNCS 2459:261–289, 2002.

[4] C. Baier, J.-P. Katoen, H. Hermanns, and V. Wolf, *"Comparative Branching-Time Semantics for Markov Chains"*, in Information and Computation 200:149–214, 2005.

[5] S. Balsamo, M. Bernardo, and M. Simeoni, *"Performance Evaluation at the Software Architecture Level"*, in *Formal Methods for Software Architectures*, Springer, LNCS 2804:207–258, 2003.

[6] J.A. Bergstra, A. Ponse, and S.A. Smolka (eds.), *"Handbook of Process Algebra"*, Elsevier, 2001.

[7] M. Bernardo, *"Non-Bisimulation-Based Markovian Behavioral Equivalences"*, in Journal of Logic and Algebraic Programming 72:3–49, 2007.

[8] M. Bernardo, *"Towards State Space Reduction Based on T-Lumpability-Consistent Relations"*, in Proc. of the *5th European Performance Engineering Workshop (EPEW 2008)*, Springer, LNCS 5261:64–78, Palma de Mallorca (Spain), 2008.

[9] M. Bernardo, *"Uniform Logical Characterizations of Testing Equivalences for Nondeterministic, Probabilistic and Markovian Processes"*, in Proc. of the *7th Int. Workshop on Quantitative Aspects of Programming Languages (QAPL 2009)*, Elsevier, ENTCS 253(3):3–23, York (UK), 2009.

[10] M. Bernardo, *"Markovian Testing Equivalence and Exponentially Timed Internal Actions"*, in Proc. of the *1st Int. Workshop on Quantitative Formal Methods (QFM 2009)*, EPTCS 13:13–25, Eindhoven (The Netherlands), 2009.

[11] M. Bernardo and A. Aldini, *"Weak Markovian Bisimilarity: Abstracting from Prioritized/Weighted Internal Immediate Actions"*, in Proc. of the 10th Italian Conf. on Theoretical Computer Science (ICTCS 2007), World Scientific, pp. 39–56, Rome (Italy), 2007.

[12] M. Bernardo and S. Botta, *"A Survey of Modal Logics Characterizing Behavioral Equivalences for Nondeterministic and Stochastic Systems"*, in Mathematical Structures in Computer Science 18:29–55, 2008.

[13] M. Bernardo and M. Bravetti, *"Performance Measure Sensitive Congruences for Markovian Process Algebras"*, in Theoretical Computer Science 290:117–160, 2003.

[14] M. Bernardo and R. Cleaveland, *"A Theory of Testing for Markovian Processes"*, in Proc. of the *11th Int. Conf. on Concurrency Theory (CONCUR 2000)*, Springer, LNCS 1877:305–319, State College (PA), 2000.

[15] M. Bernardo, L. Donatiello, and P. Ciancarini, *"Stochastic Process Algebra: From an Algebraic Formalism to an Architectural Description Language"*, in *Performance Evaluation of Complex Systems: Techniques and Tools*, Springer, LNCS 2459:236–260, 2002.

[16] M. Bernardo and R. Gorrieri, *"A Tutorial on EMPA: A Theory of Concurrent Processes with Nondeterminism, Priorities, Probabilities and Time"*, in Theoretical Computer Science 202:1–54, 1998.

[17] H. Bohnenkamp, P.R. D'Argenio, H. Hermanns, and J.-P. Katoen, "MODEST*: A Compositional Modeling Formalism for Hard and Softly Timed Systems"*, in IEEE Trans. on Software Engineering 32:812–830, 2006.

[18] M. Bravetti, *"Revisiting Interactive Markov Chains"*, in Proc. of the *3rd Int. Workshop on Models for Time-Critical Systems (MTCS 2002)*, Elsevier, ENTCS 68(5):1–20, Brno (Czech Republic), 2002.

[19] M. Bravetti, M. Bernardo, and R. Gorrieri, *"A Note on the Congruence Proof for Recursion in Markovian Bisimulation Equivalence"*, in Proc. of the *6th Int. Workshop on Process Algebra and Performance Modelling (PAPM 1998)*, pp. 153–164, Nice (France), 1998.

[20] S.D. Brookes, C.A.R. Hoare, and A.W. Roscoe, *"A Theory of Communicating Sequential Processes"*, in Journal of the ACM 31:560–599, 1984.

[21] P. Buchholz, *"Exact and Ordinary Lumpability in Finite Markov Chains"*, in Journal of Applied Probability 31:59–75, 1994.

[22] P. Buchholz, *"Markovian Process Algebra: Composition and Equivalence"*, in Proc. of the *2nd Int. Workshop on Process Algebra and Performance Modelling (PAPM 1994)*, Technical Report 27-4, pp. 11–30, Erlangen (Germany), 1994.

[23] I. Christoff, *"Testing Equivalences and Fully Abstract Models for Probabilistic Processes"*, in Proc. of the *1st Int. Conf. on Concurrency Theory (CONCUR 1990)*, Springer, LNCS 458:126–140, Amsterdam (The Netherlands), 1990.

[24] G. Clark, S. Gilmore, and J. Hillston, *"Specifying Performance Measures for PEPA"*, in Proc. of the *5th AMAST Int. Workshop on Formal Methods for Real Time and Probabilistic Systems (ARTS 1999)*, Springer, LNCS 1601:211–227, Bamberg (Germany), 1999.

[25] R. Cleaveland, Z. Dayar, S.A. Smolka, and S. Yuen, *"Testing Preorders for Probabilistic Processes"*, in Information and Computation 154:93–148, 1999.

[26] T.M. Cover and J.A. Thomas, *"Elements of Information Theory"*, John Wiley & Sons, 1991.

[27] R. De Nicola and M. Hennessy, *"Testing Equivalences for Processes"*, in Theoretical Computer Science 34:83–133, 1983.

[28] R. De Nicola, D. Latella, and M. Massink, *"Formal Modeling and Quantitative Analysis of* KLAIM-*Based Mobile Systems"*, in Proc. of the *20th ACM Symp. on Applied Computing (SAC 2005)*, ACM Press, pp. 428–435, Santa Fe (NM), 2005.

[29] S. Derisavi, H. Hermanns, and W.H. Sanders, *"Optimal State-Space Lumping in Markov Chains"*, in Information Processing Letters 87:309–315, 2003.

[30] D. Ferrari, *"Considerations on the Insularity of Performance Evaluation"*, in IEEE Trans. on Software Engineering 12:678–683, 1986.

[31] S. Gilmore, J. Hillston, and M. Ribaudo, *"An Efficient Algorithm for Aggregating PEPA Models"*, in IEEE Trans. on Software Engineering 27:449–464, 2001.

[32] R.J. van Glabbeek, *"The Linear Time – Branching Time Spectrum I"*, in [6], pp. 3–99, 2001.

[33] R.J. van Glabbeek, S.A. Smolka, and B. Steffen, *"Reactive, Generative and Stratified Models of Probabilistic Processes"*, in Information and Computation 121:59–80, 1995.

[34] N. Götz, U. Herzog, and M. Rettelbach, *"Multiprocessor and Distributed System Design: The Integration of Functional Specification and Performance Analysis Using Stochastic Process Algebras"*, in Proc. of the *16th Int. Symp. on Computer Performance Modelling, Measurement and Evaluation (PERFORMANCE 1993)*, Springer, LNCS 729:121–146, Rome (Italy), 1993.

[35] B.R. Haverkort and K.S. Trivedi, *"Specification Techniques for Markov Reward Models"*, in Discrete Event Dynamic Systems: Theory and Applications 3:219–247, 1993.

[36] M. Hennessy, *"Algebraic Theory of Processes"*, MIT Press, 1988.

[37] M. Hennessy and R. Milner, *"Algebraic Laws for Nondeterminism and Concurrency"*, in Journal of the ACM 32:137–162, 1985.

[38] H. Hermanns, *"Interactive Markov Chains"*, Springer, LNCS 2428, 2002.

[39] H. Hermanns and M. Rettelbach, *"Syntax, Semantics, Equivalences, and Axioms for MTIPP"*, in Proc. of the *2nd Int. Workshop on Process Algebra and Performance Modelling (PAPM 1994)*, Technical Report 27-4, pp. 71–87, Erlangen (Germany), 1994.

[40] J. Hillston, *"The Nature of Synchronisation"*, in Proc. of the *2nd Int. Workshop on Process Algebra and Performance Modelling (PAPM 1994)*, Technical Report 27-4, pp. 51–70, Erlangen (Germany), 1994.

[41] J. Hillston, *"A Compositional Approach to Performance Modelling"*, Cambridge University Press, 1996.

[42] C.A.R. Hoare, *"Communicating Sequential Processes"*, Prentice Hall, 1985.

[43] R.A. Howard, *"Dynamic Probabilistic Systems"*, John Wiley & Sons, 1971.

[44] D.T. Huynh and L. Tian, *"On Some Equivalence Relations for Probabilistic Processes"*, in Fundamenta Informaticae 17:211–234, 1992.

[45] B. Jonsson and K.G. Larsen, *"Specification and Refinement of Probabilistic Processes"*, in Proc. of the *6th IEEE Symp. on Logic in Computer Science (LICS 1991)*, IEEE-CS Press, pp. 266–277, Amsterdam (The Netherlands), 1991.

[46] C.-C. Jou and S.A. Smolka, *"Equivalences, Congruences, and Complete Axiomatizations for Probabilistic Processes"*, in Proc. of the *1st Int. Conf. on Concurrency Theory (CONCUR 1990)*, Springer, LNCS 458:367–383, Amsterdam (The Netherlands), 1990.

[47] P.C. Kanellakis and S.A. Smolka, *"CCS Expressions, Finite State Processes, and Three Problems of Equivalence"*, in Information and Computation 86:43–68, 1990.

[48] L. Kleinrock, *"Queueing Systems"*, John Wiley & Sons, 1975.

[49] M.Z. Kwiatkowska and G.J. Norman, *"A Testing Equivalence for Reactive Probabilistic Processes"*, in Proc. of the *2nd Int. Workshop on Expressiveness in Concurrency (EXPRESS 1998)*, Elsevier, ENTCS 16(2):114–132, Nice (France), 1998.

[50] K.G. Larsen and A. Skou, *"Bisimulation Through Probabilistic Testing"*, in Information and Computation 94:1–28, 1991.

[51] E.D. Lazowska, J. Zahorjan, G. Scott Graham, and K.C. Sevcik, *"Quantitative System Performance: Computer System Analysis Using Queueing Network Models"*, Prentice Hall, 1984.

[52] R. Milner, *"Communication and Concurrency"*, Prentice Hall, 1989.

[53] M.F. Neuts, *"Matrix-Geometric Solutions in Stochastic Models – An Algorithmic Approach"*, John Hopkins University Press, 1981.

[54] R. Paige and R.E. Tarjan, *"Three Partition Refinement Algorithms"*, in SIAM Journal on Computing 16:973–989, 1987.

[55] D. Park, *"Concurrency and Automata on Infinite Sequences"*, in Proc. of the *5th GI Conf. on Theoretical Computer Science*, Springer, LNCS 104:167–183, Karlsruhe (Germany), 1981.

[56] G.D. Plotkin, *"A Structural Approach to Operational Semantics"*, Tech. Rep. DAIMI-FN-19, Aarhus University (Denmark), 1981 (published in Journal of Logic and Algebraic Programming 60/61:17–139, 2004).

[57] C. Priami, *"Stochastic π-Calculus"*, in Computer Journal 38:578–589, 1995.

[58] M. Rettelbach, *"Probabilistic Branching in Markovian Process Algebras"*, in Computer Journal 38:590–599, 1995.

[59] W.H. Sanders and J.F. Meyer, *"A Unified Approach for Specifying Measures of Performance, Dependability, and Performability"*, in Dependable Computing and Fault Tolerant Systems 4:215–237, 1991.

[60] C. Smith, *"Performance Engineering of Software Systems"*, Addison-Wesley, 1990.

[61] J. Sproston and S. Donatelli, *"Backward Bisimulation in Markov Chain Model Checking"*, in IEEE Trans. on Software Engineering 32:531–546, 2006.

[62] E.W. Stark, R. Cleaveland, and S.A. Smolka, *"A Process-Algebraic Language for Probabilistic I/O Automata"*, in Proc. of the *14th Int. Conf. on Concurrency Theory (CONCUR 2003)*, Springer, LNCS 2761:189–203, Marseille (France), 2003.

[63] W.J. Stewart, *"Introduction to the Numerical Solution of Markov Chains"*, Princeton University Press, 1994.

[64] W.-G. Tzeng, *"A Polynomial-Time Algorithm for the Equivalence of Probabilistic Automata"*, in SIAM Journal on Computing 21:216–227, 1992.

[65] V. Wolf, C. Baier, and M. Majster-Cederbaum, *"Trace Machines for Observing Continuous-Time Markov Chains"*, in Proc. of the *3rd Int. Workshop on Quantitative Aspects of Programming Languages (QAPL 2005)*, Elsevier, ENTCS 153(2):259–277, Edinburgh (UK), 2005.

[66] C.M. Woodside, *"From Annotated Software Designs (UML SPT/MARTE) to Model Formalisms"*, in *Formal Methods for Performance Evaluation*, Springer, LNCS 4486:429–467, 2007.

[67] C.M. Woodside, J.E. Neilson, D.C. Petriu, and S. Majumdar, *"The Stochastic Rendezvous Network Model for Performance of Synchronous Client-Server-Like Distributed Software"*, in IEEE Trans. on Computers 44:20–34, 1995.