

Markovian Behavioral Equivalences: Their Spectrum, Some Known Results, and One Open Problem

Marco Bernardo

University of Urbino – Italy



© September 2010

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 $t_0 < t_1 < \dots < t_n < t_{n+1}$, and states $s_0, s_1, \dots, s_n, s_{n+1} \in S$:

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

- 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 $\pi()$** 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 Q** .
 - ⊙ The sojourn time in any state is exponentially distributed.
 - ⊙ Given $\pi(0)$, the transient solution $\pi(t)$ is obtained by solving:

$$\pi(t) \cdot Q = \frac{d\pi(t)}{dt}$$

- ⊙ The stationary solution $\pi = \lim_{t \rightarrow \infty} \pi(t)$ is obtained (if any) by solving:

$$\begin{aligned} \pi \cdot Q &= \mathbf{0} \\ \sum_{s \in S} \pi[s] &= 1 \end{aligned}$$

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

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

- Every CTMC (time-aware model) has an **embedded DTMC** (time-abstract model):
 - ⊙ State transitions are described by a **probability matrix P** .
 - ⊙ P is obtained from 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 $\pi(0)$, the transient solution $\pi(n)$ is computed as follows:

$$\pi(n) = \pi(0) \cdot P^n$$

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

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

- 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?

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].
 - ◉ EMPA_{gr} [Bernardo, Bravetti, Gorrieri].
 - ◉ $S\pi$ [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: $\langle a, \lambda \rangle. _-$
 - ◉ The choice among several actions is probabilistic.
 - ◉ TIPP, PEPA, MPA, EMPA_{gr}, S π , 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, $S\pi$.

- Asymmetric synchronizations:

- ⊙ Passive actions of the form $\langle a, *_w \rangle$ whose duration is unspecified.
- ⊙ An exponentially timed action can synchronize only with a passive action, thus determining the duration of the synchronization.
- ⊙ PEPA, EMPA_{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).
- $Name_v$: set of visible action names.
- $Name = Name_v \cup \{\tau\}$: set of all action names.
- $Rate = \mathbb{R}_{>0} \cup \{*_w \mid w \in \mathbb{R}_{>0}\}$: set of action rates.
- $Act_M = Name \times Rate$: 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}_M :

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

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

- The duration of $\langle a, \lambda \rangle$ is the exponentially distributed random variable Exp_λ , where $\Pr\{Exp_\lambda \leq t\} = 1 - e^{-\lambda \cdot t}$ and $E\{Exp_\lambda\} = 1 / \lambda$.
- 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 $\langle a, *_w \rangle$ 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 $(\lambda_1, \dots, \lambda_h)$ 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:

$$\min(\text{Exp}_{\lambda_1}, \dots, \text{Exp}_{\lambda_h}) = \text{Exp}_{\lambda_1 + \dots + \lambda_h}$$

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

- $P_1 + P_2$ 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, $S = Name_V$ implies P_1 and P_2 fully synchronized).

- $\underline{0}$ is a terminated process and hence cannot execute any action.
- $\langle a, \tilde{\lambda} \rangle . P$ 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}_M : set of closed and guarded process terms (fully defined, **finitely branching**).

- **Running example** (MPC syntax):
 - ⊙ Producer-consumer system: producer, finite 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).
 - ⊙ Assumption 1: the buffer has only two positions.
 - ⊙ Assumption 2: identical items, hence the discipline is not important.

- ⊙ The only observable activities are deposits and withdrawals.
- ⊙ Names of visible actions: *deposit* and *withdraw*.
- ⊙ Structure-independent process algebraic description:

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

- ⊙ 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 ($\langle a, \lambda \rangle . \underline{0} + \langle a, \lambda \rangle . \underline{0}$ vs. $\langle a, \lambda \rangle . \underline{0}$).
- Every $P \in \mathbb{P}_M$ is mapped to a **labeled multitransition system** $\llbracket P \rrbracket_M$:
 - ⊙ 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 $P \in \mathbb{P}_{M,pc}$ is mapped to a CTMC (**performance closure** if no passive trans.):
 - ⊙ Dropping action names from all transitions of $\llbracket P \rrbracket_M$.
 - ⊙ Collapsing all the transitions between any two states of $\llbracket P \rrbracket_M$ 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** (\cdot $+$), **static operators** (\parallel $/$ \backslash $[]$).
- The **multitransition relation** $\longrightarrow_{M,P}$ of $\llbracket P \rrbracket_M$ is contained in the smallest multiset of elements of $\mathbb{P}_M \times Act_M \times \mathbb{P}_M$ that:
 - Satisfy the operational semantic rules.
 - Keep track of all the possible ways of deriving each transition.
- No rule for $\underline{0}$: $\llbracket \underline{0} \rrbracket_M$ has a single state and no transitions.

- Operational semantic rules for action prefix:

$$\boxed{\begin{array}{l} \langle a, \lambda \rangle . P \xrightarrow{a, \lambda}_{\text{M}} P \\ \langle a, *_w \rangle . P \xrightarrow{a, *_w}_{\text{M}} P \end{array}}$$

- Operational semantic rules for alternative composition:

$$\boxed{\begin{array}{c} \frac{P_1 \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P'}{P_1 + P_2 \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P'} \qquad \frac{P_2 \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P'}{P_1 + P_2 \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P'} \end{array}}$$

- Operational semantic rule for recursion:

$$\boxed{\begin{array}{c} \frac{P\{\text{rec } X : P \hookrightarrow X\} \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P'}{\text{rec } X : P \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P'} \qquad \left(\frac{B \triangleq P \quad P \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P'}{B \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P'} \right) \end{array}}$$

- 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:

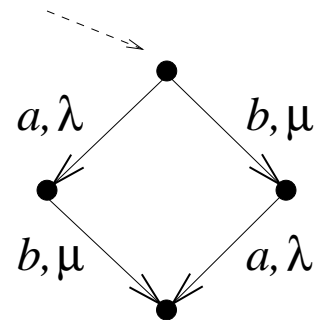
$\frac{P_1 \xrightarrow{a, \tilde{\lambda}}_M P'_1 \quad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{a, \tilde{\lambda}}_M P'_1 \parallel_S P_2}$	$\frac{P_2 \xrightarrow{a, \tilde{\lambda}}_M P'_2 \quad a \notin S}{P_1 \parallel_S P_2 \xrightarrow{a, \tilde{\lambda}}_M P_1 \parallel_S P'_2}$
--	--

- The following process terms represent structurally different systems:

$$\begin{aligned} & \langle a, \lambda \rangle . \underline{0} \parallel_{\emptyset} \langle b, \mu \rangle . \underline{0} \\ & \langle a, \lambda \rangle . \langle b, \mu \rangle . \underline{0} + \langle b, \mu \rangle . \langle a, \lambda \rangle . \underline{0} \end{aligned}$$

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 :

$$weight(P, a) = \sum \{ w \in \mathbb{R}_{>0} \mid \exists P' \in \mathbb{P}_M. P \xrightarrow{a, *w}_M P' \}$$

- 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:

$$\boxed{
\begin{array}{c}
\frac{P_1 \xrightarrow{a, \lambda}_M P'_1 \quad P_2 \xrightarrow{a, *w}_M P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, \lambda \cdot \frac{w}{\text{weight}(P_2, a)}}_M P'_1 \parallel_S P'_2} \\
\\
\frac{P_1 \xrightarrow{a, *w}_M P'_1 \quad P_2 \xrightarrow{a, \lambda}_M P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, \lambda \cdot \frac{w}{\text{weight}(P_1, a)}}_M P'_1 \parallel_S P'_2}
\end{array}
}$$

- Operational semantic rule for reactive-reactive synchronization:

$$\boxed{
\frac{P_1 \xrightarrow{a, *w_1}_M P'_1 \quad P_2 \xrightarrow{a, *w_2}_M P'_2 \quad a \in S}{P_1 \parallel_S P_2 \xrightarrow{a, *_{\text{norm}}(w_1, w_2, a, P_1, P_2)}}_M P'_1 \parallel_S P'_2
}$$

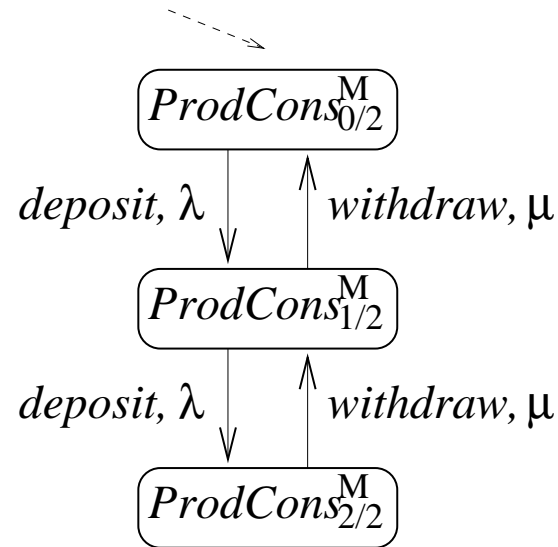
- Operational semantic rules for hiding, restriction, relabeling:

$$\begin{array}{c}
\frac{P \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P' \quad a \in H}{P / H \xrightarrow{\tau, \tilde{\lambda}}_{\text{M}} P' / H} \qquad \frac{P \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P' \quad a \notin H}{P / H \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P' / H} \\
\\
\frac{P \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P' \quad a \notin L}{P \setminus L \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P' \setminus L} \\
\\
\frac{P \xrightarrow{a, \tilde{\lambda}}_{\text{M}} P'}{P[\varphi] \xrightarrow{\varphi(a), \tilde{\lambda}}_{\text{M}} P'[\varphi]}
\end{array}$$

- $\llbracket P \rrbracket_{\text{M}}$ is **finite state** if no recursive definition in P contains static ops.

- **Running example** (MPC semantics):

- ◉ Labeled multitransition system $\llbracket \text{ProdCons}_{0/2}^M \rrbracket_M$ with explicit states:



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

Part II:
The Markovian Spectrum

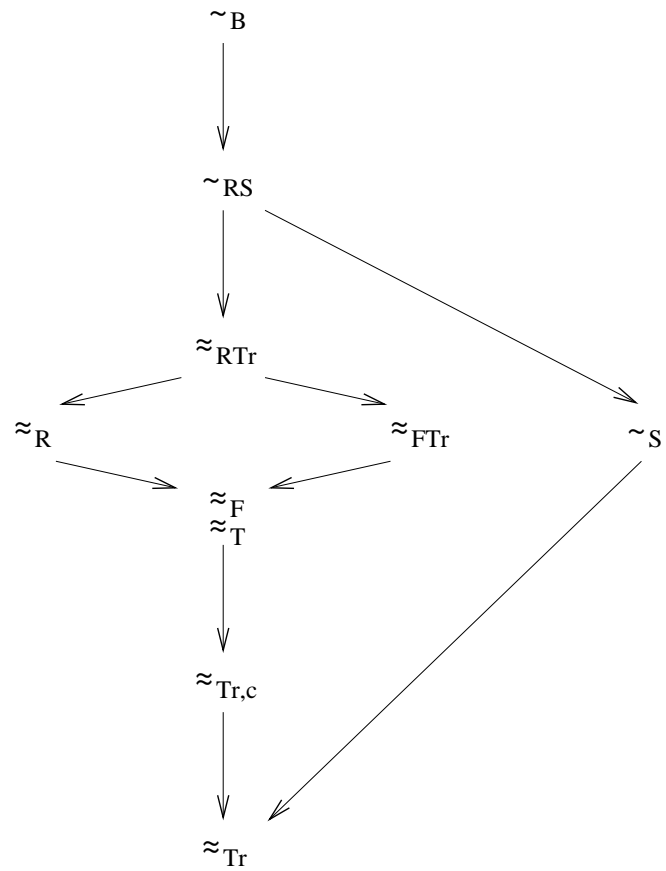
Relating Classical Behavioral Equivalences

- Most studied approaches to the definition of behavioral equivalences:
 - ◉ **Trace approach**: two process terms are equivalent if they are able to execute the same sequences of actions (\approx_{Tr}).
 - ◉ Abstraction from branching points leads to *deadlock insensitivity*.
 - ◉ **Testing approach**: two process terms are equivalent if no difference can be discovered when interacting with them by means of tests and comparing their reactions (\approx_{T}).
 - ◉ Checking whether *every test* may/must be passed.
 - ◉ **Bisimulation approach**: two process terms are equivalent if they are able to mimic each other's behavior after each action execution (\sim_{B}).
 - ◉ Faithful account of branching points leads to *overdiscrimination*.

- Variants of bisimulation equivalence:
 - ◉ **Simulation equivalence**: it is the intersection of two preorders, each of which considers the capability of stepwise behavior mimicking in one single direction (\sim_s).
 - ◉ **Ready-simulation equivalence**: same as simulation equivalence, with in addition the fact that each of the two preorders checks for the equality of the sets of actions that are stepwise enabled (\sim_{RS}).
- Less discriminating than bisimulation equivalence.
- The distinctions they make can be traced provided that reasonable operators are included in the process language.

- Deadlock-sensitive variants of trace equivalence:
 - ◉ **Completed-trace equivalence**: it compares process terms also with respect to traces that lead to deadlock ($\approx_{\text{Tr},c}$).
 - ◉ **Failure equivalence**: it takes into account the set of visible actions that can be refused after executing a trace (\approx_{F}).
 - ◉ **Ready equivalence**: it takes into account the set of visible actions that can be performed after executing a trace (\approx_{R}).
 - ◉ **Failure-trace equivalence**: it takes into account the sets of visible actions that can be refused at each step of a trace (\approx_{FTr}).
 - ◉ **Ready-trace equivalence**: it takes into account the sets of visible actions that can be performed at each step of a trace (\approx_{RTTr}).

- Linear-time/branching-time spectrum (for processes without invisible actions):



Relating Markovian Behavioral Equivalences

- Markovian linear-time/branching-time spectrum:

$$\sim_{\text{MB}} = \sim_{\text{MRS}} = \sim_{\text{MS}} \subset$$

$$\sim_{\text{MRTr}} = \sim_{\text{MFTr}} \subset$$

$$\sim_{\text{MR}} = \sim_{\text{MF}} = \sim_{\text{MT}} \subset$$

$$\sim_{\text{MCTr}} = \sim_{\text{MTr}}$$

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

Part III:
Properties of Markovian Behavioral Equivalences

Comparison Criteria

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.

Markovian Bisimulation Equivalence

- 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 $P \in \mathbb{P}_M$ executes actions of name $a \in \text{Name}$ and level $l \in \{0, -1\}$ (0 exp. timed, -1 passive) that lead to destination $D \subseteq \mathbb{P}_M$:

$$rate_e(P, a, l, D) = \begin{cases} \sum \{ \lambda \in \mathbb{R}_{>0} \mid \exists P' \in D. P \xrightarrow{a, \lambda}_M P' \} & \text{if } l = 0 \\ \sum \{ w \in \mathbb{R}_{>0} \mid \exists P' \in D. P \xrightarrow{a, *w}_M P' \} & \text{if } l = -1 \end{cases}$$

- Overall exit rate of P w.r.t. a at level l : $rate_o(P, a, l) = rate_e(P, a, l, \mathbb{P}_M)$.
- Total exit rate of P at level l : $rate_t(P, l) = \sum_{a \in \text{Name}} rate_o(P, a, l)$.
- $1 / rate_t(P, 0)$ is the average sojourn time of P when $P \in \mathbb{P}_{M,pc}$.

- 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 $P \in \mathbb{P}_M$ executes actions of name $a \in Name$ and level $l \in \{0, -1\}$ that lead to destination $D \subseteq \mathbb{P}_M$:

$$prob_e(P, a, l, D) = \begin{cases} rate_e(P, a, l, D) / rate_t(P, l) & \text{if } l = 0 \\ rate_e(P, a, l, D) / rate_o(P, a, l) & \text{if } l = -1 \end{cases}$$

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

$$\text{rate}_e(P_1, a, l, D) = \text{rate}_e(P_2, a, l, D)$$

- Markovian bisimulation equivalence \sim_{MB} is the union of all the Markovian bisimulations.
- *A consequence of the coinductive nature of \sim_{MB} is that the derivatives of two equivalent terms are still equivalent.*

- **Running example** (\sim_{MB}):

- ◉ Concurrent implementation with two independent one-pos. buffers:

$$PC_{\text{conc},2}^{\text{M}} \triangleq \text{Prod}^{\text{M}} \parallel_{\{\text{deposit}\}} (\text{Buff}^{\text{M}} \parallel_{\emptyset} \text{Buff}^{\text{M}}) \parallel_{\{\text{withdraw}\}} \text{Cons}^{\text{M}}$$

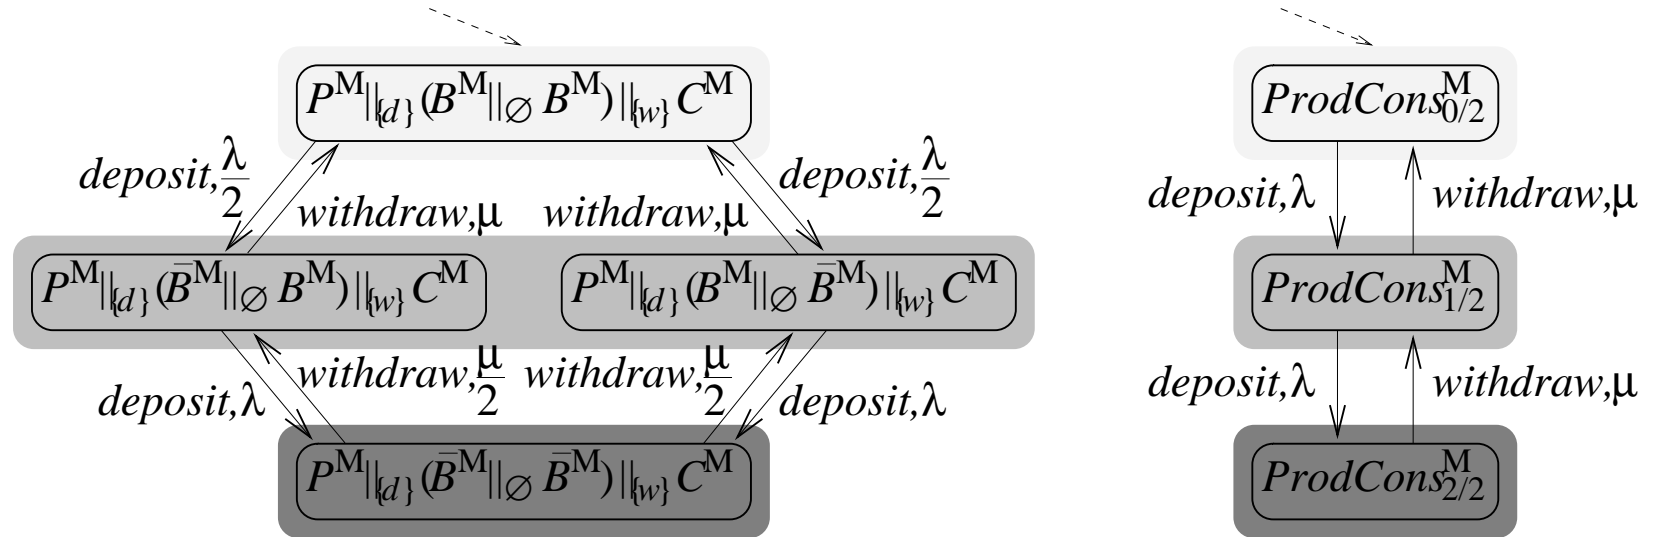
$$\text{Prod}^{\text{M}} \triangleq \langle \text{deposit}, \lambda \rangle . \text{Prod}^{\text{M}}$$

$$\text{Buff}^{\text{M}} \triangleq \langle \text{deposit}, *_1 \rangle . \langle \text{withdraw}, *_1 \rangle . \text{Buff}^{\text{M}}$$

$$\text{Cons}^{\text{M}} \triangleq \langle \text{withdraw}, \mu \rangle . \text{Cons}^{\text{M}}$$

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

- ◉ Markovian bisimulation proving $PC_{\text{conc},2}^M \sim_{\text{MB}} \text{ProdCons}_{0/2}^M$, 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 .

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

$$\boxed{\text{rate}_o(P_1, a, l) = \text{rate}_o(P_2, a, l)}$$

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

$$\boxed{\text{rate}_e(P_1, a, l, D) = \text{rate}_e(P_2, a, l, D)}$$

- Focus on important pairs of process terms that form a bisimulation.
- In order for $P_1 \sim_{\text{MB}} P_2$, it is sufficient to find a Markovian bisimulation up to \sim_{MB} that contains (P_1, P_2) .

- \sim_{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}_{M} is a **separate Markovian bisimulation** iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all action names $a \in \text{Name}$ and levels $l \in \{0, -1\}$:

$$\text{rate}_o(P_1, a, l) = \text{rate}_o(P_2, a, l)$$

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

$$\text{prob}_e(P_1, a, l, D) = \text{prob}_e(P_2, a, l, D)$$

- Separate Markovian bisimulation equivalence $\sim_{\text{MB},s}$ is the union of all the separate Markovian bisimulations.
- For all $P_1, P_2 \in \mathbb{P}_{\text{M}}$:

$$P_1 \sim_{\text{MB},s} P_2 \iff P_1 \sim_{\text{MB}} P_2$$

- \sim_{MB} is a **congruence** with respect to all the dynamic and static operators as well as recursion.
- Let $P_1, P_2 \in \mathbb{P}_M$. Whenever $P_1 \sim_{\text{MB}} P_2$, then:

$$\begin{array}{c}
\langle a, \tilde{\lambda} \rangle . P_1 \sim_{\text{MB}} \langle a, \tilde{\lambda} \rangle . P_2 \\
P_1 + P \sim_{\text{MB}} P_2 + P \qquad P + P_1 \sim_{\text{MB}} P + P_2 \\
P_1 \parallel_S P \sim_{\text{MB}} P_2 \parallel_S P \qquad P \parallel_S P_1 \sim_{\text{MB}} P \parallel_S P_2 \\
P_1 / H \sim_{\text{MB}} P_2 / H \\
P_1 \setminus L \sim_{\text{MB}} P_2 \setminus L \\
P_1[\varphi] \sim_{\text{MB}} P_2[\varphi]
\end{array}$$

- Recursion: extend \sim_{MB} to open process terms by replacing all variables freely occurring outside rec binders with every closed process term.
- Let $P_1, P_2 \in \mathcal{PL}_M$ be guarded process terms containing free occurrences of $k \in \mathbb{N}$ process variables $X_1, \dots, X_k \in \text{Var}$ at most.
- We define $P_1 \sim_{\text{MB}} P_2$ iff:

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

for all $Q_1, \dots, Q_k \in \mathbb{P}_M$:

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

$$\text{rec } X : P_1 \sim_{\text{MB}} \text{rec } X : P_2$$

- \sim_{MB} has a **sound and complete axiomatization** over the set $\mathbb{P}_{\text{M,nrec}}$ of nonrecursive process terms of \mathbb{P}_{M} .
- Basic laws (commutativity, associativity, and neutral element of $+$):

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

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

$$\begin{array}{ll}
 (\mathcal{X}_{\text{MB},4}) & \langle a, \lambda_1 \rangle . P + \langle a, \lambda_2 \rangle . P = \langle a, \lambda_1 + \lambda_2 \rangle . P \\
 (\mathcal{X}_{\text{MB},5}) & \langle a, *_{w_1} \rangle . P + \langle a, *_{w_2} \rangle . P = \langle a, *_{w_1 + w_2} \rangle . P
 \end{array}$$

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

$$\begin{aligned}
(\mathcal{X}_{\text{MB},6}) \quad & \sum_{i \in I} \langle a_i, \tilde{\lambda}_i \rangle . P_i \parallel_S \sum_{j \in J} \langle b_j, \tilde{\mu}_j \rangle . Q_j = \\
& \sum_{k \in I, a_k \notin S} \langle a_k, \tilde{\lambda}_k \rangle . \left(P_k \parallel_S \sum_{j \in J} \langle b_j, \tilde{\mu}_j \rangle . Q_j \right) + \\
& \sum_{h \in J, b_h \notin S} \langle b_h, \tilde{\mu}_h \rangle . \left(\sum_{i \in I} \langle a_i, \tilde{\lambda}_i \rangle . P_i \parallel_S Q_h \right) + \\
& \sum_{k \in I, a_k \in S, \tilde{\lambda}_k \in \mathbb{R}_{>0}} \sum_{h \in J, b_h = a_k, \tilde{\mu}_h = *w_h} \langle a_k, \tilde{\lambda}_k \cdot \frac{w_h}{\text{weight}(Q, b_h)} \rangle . (P_k \parallel_S Q_h) + \\
& \sum_{h \in J, b_h \in S, \tilde{\mu}_h \in \mathbb{R}_{>0}} \sum_{k \in I, a_k = b_h, \tilde{\lambda}_k = *v_k} \langle b_h, \tilde{\mu}_h \cdot \frac{v_k}{\text{weight}(P, a_k)} \rangle . (P_k \parallel_S Q_h) + \\
& \sum_{k \in I, a_k \in S, \tilde{\lambda}_k = *v_k} \sum_{h \in J, b_h = a_k, \tilde{\mu}_h = *w_h} \langle a_k, *_{\text{norm}(v_k, w_h, a_k, P, Q)} \rangle . (P_k \parallel_S Q_h) \\
(\mathcal{X}_{\text{MB},7}) \quad & \sum_{i \in I} \langle a_i, \tilde{\lambda}_i \rangle . P_i \parallel_S \underline{0} = \sum_{k \in I, a_k \notin S} \langle a_k, \tilde{\lambda}_k \rangle . P_k \\
(\mathcal{X}_{\text{MB},8}) \quad & \underline{0} \parallel_S \sum_{j \in J} \langle b_j, \tilde{\mu}_j \rangle . Q_j = \sum_{h \in J, b_h \notin S} \langle b_h, \tilde{\mu}_h \rangle . Q_h \\
(\mathcal{X}_{\text{MB},9}) \quad & \underline{0} \parallel_S \underline{0} = \underline{0}
\end{aligned}$$

- Distribution laws (for unary static operators):

$(\mathcal{X}_{\text{MB},10})$	$\underline{0} / H = \underline{0}$	
$(\mathcal{X}_{\text{MB},11})$	$(\langle a, \tilde{\lambda} \rangle . P) / H = \langle \tau, \tilde{\lambda} \rangle . (P / H)$	if $a \in H$
$(\mathcal{X}_{\text{MB},12})$	$(\langle a, \tilde{\lambda} \rangle . P) / H = \langle a, \tilde{\lambda} \rangle . (P / H)$	if $a \notin H$
$(\mathcal{X}_{\text{MB},13})$	$(P_1 + P_2) / H = P_1 / H + P_2 / H$	
$(\mathcal{X}_{\text{MB},14})$	$\underline{0} \setminus L = \underline{0}$	
$(\mathcal{X}_{\text{MB},15})$	$(\langle a, \tilde{\lambda} \rangle . P) \setminus L = \underline{0}$	if $a \in L$
$(\mathcal{X}_{\text{MB},16})$	$(\langle a, \tilde{\lambda} \rangle . P) \setminus L = \langle a, \tilde{\lambda} \rangle . (P \setminus L)$	if $a \notin L$
$(\mathcal{X}_{\text{MB},17})$	$(P_1 + P_2) \setminus L = P_1 \setminus L + P_2 \setminus L$	
$(\mathcal{X}_{\text{MB},18})$	$\underline{0}[\varphi] = \underline{0}$	
$(\mathcal{X}_{\text{MB},19})$	$(\langle a, \tilde{\lambda} \rangle . P)[\varphi] = \langle \varphi(a), \tilde{\lambda} \rangle . (P[\varphi])$	
$(\mathcal{X}_{\text{MB},20})$	$(P_1 + P_2)[\varphi] = P_1[\varphi] + P_2[\varphi]$	

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

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

- \sim_{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}_{MB} ($a \in \text{Name}, \lambda, w \in \mathbb{R}_{>0}$):

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

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

$P \models_{\text{MB}} \text{true}$	
$P \models_{\text{MB}} \neg\phi$	if $P \not\models_{\text{MB}} \phi$
$P \models_{\text{MB}} \phi_1 \wedge \phi_2$	if $P \models_{\text{MB}} \phi_1$ and $P \models_{\text{MB}} \phi_2$
$P \models_{\text{MB}} \langle a \rangle_{\lambda} \phi$	if $\text{rate}_e(P, a, 0, \text{sat}(\phi)) \geq \lambda$
$P \models_{\text{MB}} \langle a \rangle_{*w} \phi$	if $\text{rate}_e(P, a, -1, \text{sat}(\phi)) \geq w$

where:

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

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

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

- \sim_{MB} is **decidable in polynomial time** over the set $\mathbb{P}_{\text{M,fin}}$ of finite-state process terms of \mathbb{P}_{M} : **Paige-Tarjan partition refinement algorithm**.
- Based on the fact that \sim_{MB} can be characterized as the limit of a sequence of successively finer equivalence relations:

$$\sim_{\text{MB}} = \bigcap_{i \in \mathbb{N}} \sim_{\text{MB},i}$$

- $\sim_{\text{MB},0} = \mathbb{P}_{\text{M}} \times \mathbb{P}_{\text{M}}$ hence it induces the trivial partition $\{\mathbb{P}_{\text{M}}\}$.
- Whenever $P_1 \sim_{\text{MB},i} P_2$, $i \in \mathbb{N}_{\geq 1}$, then for all $a \in \text{Name}$, $l \in \{0, -1\}$, and $D \in \mathbb{P}_{\text{M}} / \sim_{\text{MB},i-1}$:

$$\text{rate}_e(P_1, a, l, D) = \text{rate}_e(P_2, a, l, D)$$

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

- Steps of the algorithm for checking whether $P_1 \sim_{\text{MB}} P_2$:
 1. Build an initial partition with a single class including all the states of $\llbracket P_1 \rrbracket_M$ and $\llbracket P_2 \rrbracket_M$.
 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 $a \in \text{Name}_{P_1, P_2}$ 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 $\llbracket P_1 \rrbracket_M$ and $\llbracket P_2 \rrbracket_M$ 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_{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 $O' \in \mathcal{O}$:

$$\sum \{ \lambda \in \mathbb{R}_{>0} \mid \exists s' \in O'. s_1 \xrightarrow{\lambda} s' \} = \sum \{ \lambda \in \mathbb{R}_{>0} \mid \exists s' \in O'. s_2 \xrightarrow{\lambda} s' \}$$

- 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 $\mathbb{P}_{\text{M,pc}}$ are guaranteed to possess the same performance characteristics.

Markovian Testing Equivalence

- 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 $P \in \mathbb{P}_M$ 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.
- $\mathcal{C}_f(P)$: 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 $c \in \mathcal{C}_f(P)$ is the sequence of action names labeling the transitions of c :

$$trace_c(c) = \begin{cases} \varepsilon & \text{if } |c| = 0 \\ a \circ trace_c(c') & \text{if } c \equiv P \xrightarrow{a, \tilde{\lambda}}_M c' \end{cases}$$

- We denote by $trace(c)$ the visible part of $trace_c(c)$, i.e., the subsequence of $trace_c(c)$ obtained by removing all the occurrences of τ .

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

$$prob(c) = \begin{cases} 1 & \text{if } |c| = 0 \\ \frac{\lambda}{rate_t(P,0)} \cdot prob(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_M c' \end{cases}$$

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

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

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

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

$$time_a(c) = \begin{cases} \varepsilon & \text{if } |c| = 0 \\ \frac{1}{rate_t(P,0)} \circ time_a(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_M c' \end{cases}$$

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

$$C_{\leq \theta} = \{ c \in C \mid |c| \leq |\theta| \wedge \forall i = 1, \dots, |c|. time_a(c)[i] \leq \theta[i] \}$$

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

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

$$time_d(c) = \begin{cases} \varepsilon & \text{if } |c| = 0 \\ Exp_{rate_t(P,0)} \circ time_d(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_M c' \end{cases}$$

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

$$prob_d(C, \theta) = \sum_{c \in C}^{|\theta|} prob(c) \cdot \prod_{i=1}^{|c|} \Pr\{time_d(c)[i] \leq \theta[i]\}$$

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

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

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

$$\begin{aligned} & \langle g, \gamma \rangle . \langle a, \lambda \rangle . \langle b, \mu \rangle . \underline{0} + \langle g, \gamma \rangle . \langle a, \mu \rangle . \langle d, \lambda \rangle . \underline{0} \\ & \langle g, \gamma \rangle . \langle a, \lambda \rangle . \langle d, \mu \rangle . \underline{0} + \langle g, \gamma \rangle . \langle a, \mu \rangle . \langle b, \lambda \rangle . \underline{0} \end{aligned}$$

- Maximal computations of the first term:

$$\begin{aligned} C_{1,1} &\equiv \cdot \xrightarrow[g, \gamma]{\text{pink}}_M \cdot \xrightarrow[a, \lambda]{\text{pink}}_M \cdot \xrightarrow[b, \mu]{\text{pink}}_M \cdot \\ C_{1,2} &\equiv \cdot \xrightarrow[g, \gamma]{\text{cyan}}_M \cdot \xrightarrow[a, \mu]{\text{cyan}}_M \cdot \xrightarrow[d, \lambda]{\text{cyan}}_M \cdot \end{aligned}$$

- Maximal computations of the second term:

$$\begin{aligned} C_{2,1} &\equiv \cdot \xrightarrow[g, \gamma]{\text{cyan}}_M \cdot \xrightarrow[a, \lambda]{\text{cyan}}_M \cdot \xrightarrow[d, \mu]{\text{cyan}}_M \cdot \\ C_{2,2} &\equiv \cdot \xrightarrow[g, \gamma]{\text{pink}}_M \cdot \xrightarrow[a, \mu]{\text{pink}}_M \cdot \xrightarrow[b, \lambda]{\text{pink}}_M \cdot \end{aligned}$$

- 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}_R of **reactive tests** ($a \in Name_V, w \in \mathbb{R}_{>0}$):

$$\begin{array}{l} T ::= s \mid T' \\ T' ::= \langle a, *_{\mathbf{w}} \rangle . T \mid T' + T' \end{array}$$

- 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 $s + T$: two-level syntax for tests.

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

$$P \parallel_{Name_v} T \in \mathbb{P}_{M,pc}$$

- 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 $P \in \mathbb{P}_{M,pc}$ and $T \in \mathbb{T}_R$:
 - ⊙ A *configuration* is a state of $\llbracket P \parallel_{Name_v} T \rrbracket_M$.
 - ⊙ A *test-driven computation* is a computation of $\llbracket P \parallel_{Name_v} T \rrbracket_M$.
 - ⊙ 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 $P \parallel_{Name_v} T$.

- 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 $\langle \tau, \lambda \rangle . \underline{0}$ and $\langle \tau, \mu \rangle . \underline{0}$ 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 $\mathcal{SC}_{\leq \theta}^l(P, T)$ for $0 \leq l \leq |\theta|$.
- $\mathcal{SC}_{\leq \theta}^{|\theta|}(P, T)$ 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 ($\lambda > \mu \Rightarrow \frac{1}{\lambda} < \frac{1}{\mu}$):

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

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

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

- 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.

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

$$\begin{array}{l} \text{trace}_c(c_k) = \text{trace}_c(c_h) \\ \text{time}_a(c_k) = \text{time}_a(c_h) \end{array}$$

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

$$\text{rate}_o(P_k^i, a, 0) = \text{rate}_o(P_h^i, a, 0)$$

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 $\mu \neq \gamma$):

$$\begin{array}{l} \langle a, \lambda_1 \rangle . \langle b, \mu \rangle . \underline{0} + \langle a, \lambda_2 \rangle . \langle c, \gamma \rangle . \underline{0} \\ \langle a, \lambda'_1 \rangle . \langle b, \mu \rangle . \underline{0} + \langle a, \lambda'_2 \rangle . \langle c, \gamma \rangle . \underline{0} \end{array}$$

- \sim_{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 $\mathbb{T}_{\text{R,lib}}$ of tests with the following more liberal syntax:

$$T ::= s \mid \langle a, *_w \rangle . T \mid T + T$$

- In this setting, a successful configuration is a configuration whose test projection includes s as top-level summand.
- For all $P_1, P_2 \in \mathbb{P}_{\text{M,pc}}$:

$$P_1 \sim_{\text{MT,lib}} P_2 \iff P_1 \sim_{\text{MT}} P_2$$

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

$$\begin{aligned} T &::= s \mid T' \\ T' &::= \langle a, *_w \rangle.T \mid \langle \tau, \lambda \rangle.T \mid T' + T' \end{aligned}$$

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

$$P_1 \sim_{\mathbf{MT},\tau} P_2 \iff P_1 \sim_{\mathbf{MT}} P_2$$

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

$$\text{prob}_d(SC^{|\theta|}(P_1, T), \theta) = \text{prob}_d(SC^{|\theta|}(P_2, T), \theta)$$

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

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

- \sim_{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 $(\text{Name}_v \times 2^{\text{Name}_v})^*$ is an **extended trace** iff either ξ is the empty sequence ε or:

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

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

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

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

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

- $c \in \mathcal{C}_f(P)$ is **compatible** with $\xi \in \mathcal{ET}$ iff:

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

- $\mathcal{CC}(P, \xi)$: multiset of computations in $\mathcal{C}_f(P)$ 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 $c \in \mathcal{CC}(P, \xi)$:

$$prob_{\xi}(c) = \begin{cases} 1 & \text{if } |c| = 0 \\ \frac{\lambda}{rate_o(P, \mathcal{E} \cup \{\tau\}, 0)} \cdot prob_{\xi'}(c') & \text{if } c \equiv P \xrightarrow{a, \lambda}_M c' \\ & \text{with } \xi \equiv (a, \mathcal{E}) \circ \xi' \\ \frac{\lambda}{rate_o(P, \mathcal{E} \cup \{\tau\}, 0)} \cdot prob_{\xi}(c') & \text{if } c \equiv P \xrightarrow{\tau, \lambda}_M c' \\ & \text{with } \xi \equiv (a, \mathcal{E}) \circ \xi' \\ \frac{\lambda}{rate_o(P, \tau, 0)} \cdot prob_{\xi}(c') & \text{if } c \equiv P \xrightarrow{\tau, \lambda}_M c' \wedge \xi \equiv \varepsilon \end{cases}$$

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

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

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

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

$$time_{a,\xi}(c) = \begin{cases} \varepsilon & \text{if } |c| = 0 \\ \frac{1}{rate_o(P, \mathcal{E} \cup \{\tau\}, 0)} \circ time_{a,\xi'}(c') & \text{if } c \equiv P \xrightarrow{a,\lambda}_M c' \\ & \text{with } \xi \equiv (a, \mathcal{E}) \circ \xi' \\ \frac{1}{rate_o(P, \mathcal{E} \cup \{\tau\}, 0)} \circ time_{a,\xi}(c') & \text{if } c \equiv P \xrightarrow{\tau,\lambda}_M c' \\ & \text{with } \xi \equiv (a, \mathcal{E}) \circ \xi' \\ \frac{1}{rate_o(P, \tau, 0)} \circ time_{a,\xi}(c') & \text{if } c \equiv P \xrightarrow{\tau,\lambda}_M c' \wedge \xi \equiv \varepsilon \end{cases}$$

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

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

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

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

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

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

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

- 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 $\mathbb{T}_{R,c}$ of canonical reactive tests ($a \in \mathcal{E}, \mathcal{E} \subseteq \text{Name}_V$ finite):

$$T ::= s \mid \langle a, *_1 \rangle . T + \sum_{b \in \mathcal{E} - \{a\}} \langle b, *_1 \rangle . \langle z, *_1 \rangle . S$$

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

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

- **Running example** (\sim_{MT}):

- ◉ Concurrent implementation with two independent one-pos. buffers:

$$PC_{\text{conc},2}^{\text{M}} \triangleq \text{Prod}^{\text{M}} \parallel_{\{\text{deposit}\}} (\text{Buff}^{\text{M}} \parallel_{\emptyset} \text{Buff}^{\text{M}}) \parallel_{\{\text{withdraw}\}} \text{Cons}^{\text{M}}$$

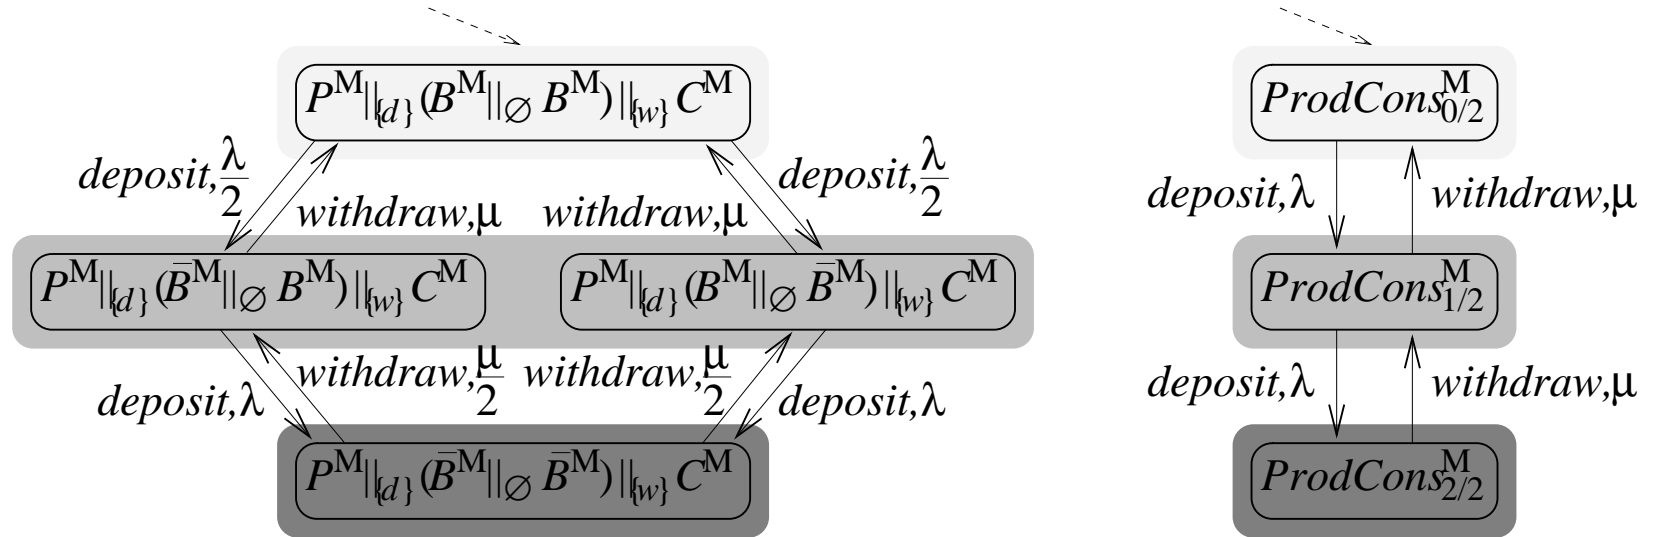
$$\text{Prod}^{\text{M}} \triangleq \langle \text{deposit}, \lambda \rangle . \text{Prod}^{\text{M}}$$

$$\text{Buff}^{\text{M}} \triangleq \langle \text{deposit}, *_1 \rangle . \langle \text{withdraw}, *_1 \rangle . \text{Buff}^{\text{M}}$$

$$\text{Cons}^{\text{M}} \triangleq \langle \text{withdraw}, \mu \rangle . \text{Cons}^{\text{M}}$$

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

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

- \sim_{MT} is a **congruence** over $\mathbb{P}_{\text{M,pc}}$ with respect to all the dynamic and static operators as well as recursion.
- Let $P_1, P_2 \in \mathbb{P}_{\text{M,pc}}$. Whenever $P_1 \sim_{\text{MT}} P_2$, then:

$$\begin{array}{c}
\langle a, \lambda \rangle . P_1 \sim_{\text{MT}} \langle a, \lambda \rangle . P_2 \\
P_1 + P \sim_{\text{MT}} P_2 + P \qquad P + P_1 \sim_{\text{MT}} P + P_2 \\
P_1 \parallel_S P \sim_{\text{MT}} P_2 \parallel_S P \qquad P \parallel_S P_1 \sim_{\text{MT}} P \parallel_S P_2 \\
P_1 / H \sim_{\text{MT}} P_2 / H \\
P_1 \setminus L \sim_{\text{MT}} P_2 \setminus L \\
P_1[\varphi] \sim_{\text{MT}} P_2[\varphi]
\end{array}$$

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

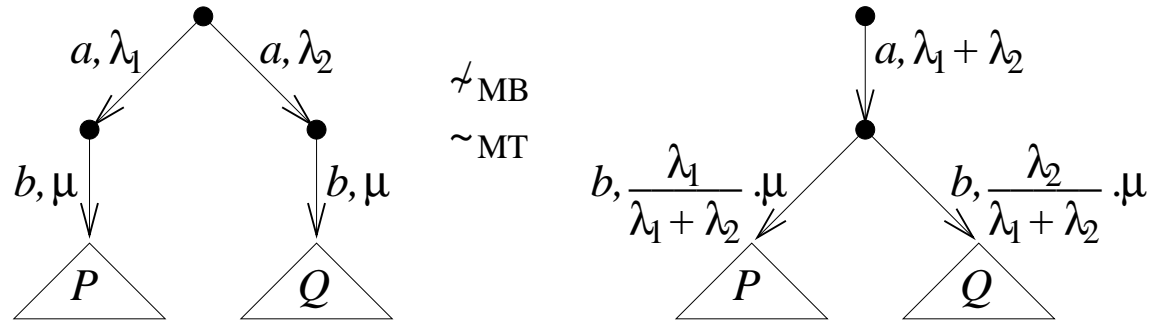
- Recursion: extend \sim_{MT} to open process terms by replacing all variables freely occurring outside rec binders with every closed process term.
- Let $P_1, P_2 \in \mathcal{PL}_M$ be guarded process terms containing free occurrences of $k \in \mathbb{N}$ process variables $X_1, \dots, X_k \in \text{Var}$ at most.
- We define $P_1 \sim_{\text{MT}} P_2$ iff there exist $Q_1, \dots, Q_k \in \mathbb{P}_M$ 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 $\mathbb{P}_{M, \text{pc}}$ and for each such group of process terms $Q_1, \dots, Q_k \in \mathbb{P}_M$:

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

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

$$\text{rec } X : P_1 \sim_{\text{MT}} \text{rec } X : P_2$$

- \sim_{MT} has a **sound and complete axiomatization** over the set $\mathbb{P}_{\text{M,pc,nrec}}$ of nonrecursive process terms of $\mathbb{P}_{\text{M,pc}}$.
- The axioms for \sim_{MB} are sound but not complete for \sim_{MT} ($P \not\sim_{\text{MB}} Q$):



- *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 $\{ \langle b, \mu \rangle \}$ after each of the two a -branches).*

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

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

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

$$(\mathcal{X}_{\text{MT},4}) \quad \sum_{i \in I} \langle a, \lambda_i \rangle \cdot \sum_{j \in J_i} \langle b_{i,j}, \mu_{i,j} \rangle \cdot P_{i,j} = \\
 \langle a, \sum_{k \in I} \lambda_k \rangle \cdot \sum_{i \in I} \sum_{j \in J_i} \langle b_{i,j}, \frac{\lambda_i}{\sum_{k \in I} \lambda_k} \cdot \mu_{i,j} \rangle \cdot P_{i,j}$$

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 $b \in \text{Name}$:

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

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

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

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

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

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

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

- \sim_{MT} has a **modal logic characterization** over $\mathbb{P}_{\text{M,pc}}$ based on a variant of the **Hennessey-Milner logic**.
- Negation is not included and conjunction is replaced by disjunction (decreased discriminating power with respect to \sim_{MB}).
- Syntax of the modal language \mathcal{ML}_{MT} ($a \in \text{Name}_{\text{v}}$):

$$\begin{array}{l} \phi ::= \text{true} \mid \phi' \\ \phi' ::= \langle a \rangle \phi \mid \phi' \vee \phi' \end{array}$$

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):

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

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

$$\text{init}(\text{true}) = \emptyset \quad \text{init}(\phi_1 \vee \phi_2) = \text{init}(\phi_1) \cup \text{init}(\phi_2) \quad \text{init}(\langle a \rangle \phi) = \{a\}$$

- 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}_{MT} over $\mathbb{P}_{\text{M,pc}}$:

$$\llbracket \phi \rrbracket_{\text{MT}}^{|\theta|}(P, \theta) = \begin{cases} 0 & \text{if } |\theta| = 0 \wedge \phi \not\equiv \text{true} \\ & \text{or } |\theta| > 0 \wedge \text{rate}_o(P, \text{init}(\phi) \cup \{\tau\}, 0) = 0 \\ 1 & \text{if } |\theta| = 0 \wedge \phi \equiv \text{true} \end{cases}$$

otherwise:

$$\begin{aligned}
\llbracket \text{true} \rrbracket_{\text{MT}}^{|t \circ \theta|}(P, t \circ \theta) &= \begin{cases} \sum_{P \xrightarrow{\tau, \lambda}_{\text{M}} P'} \frac{\lambda}{\text{rate}_{\text{o}}(P, \tau, 0)} \cdot \llbracket \text{true} \rrbracket_{\text{MT}}^{|\theta|}(P', \theta) & \text{if } \frac{1}{\text{rate}_{\text{o}}(P, \tau, 0)} \leq t \\ 0 & \text{if } \frac{1}{\text{rate}_{\text{o}}(P, \tau, 0)} > t \end{cases} \\
\llbracket \langle a \rangle \phi \rrbracket_{\text{MT}}^{|t \circ \theta|}(P, t \circ \theta) &= \begin{cases} \sum_{P \xrightarrow{a, \lambda}_{\text{M}} P'} \frac{\lambda}{\text{rate}_{\text{o}}(P, \{a, \tau\}, 0)} \cdot \llbracket \phi \rrbracket_{\text{MT}}^{|\theta|}(P', \theta) + \\ \sum_{P \xrightarrow{\tau, \lambda}_{\text{M}} P'} \frac{\lambda}{\text{rate}_{\text{o}}(P, \{a, \tau\}, 0)} \cdot \llbracket \langle a \rangle \phi \rrbracket_{\text{MT}}^{|\theta|}(P', \theta) & \text{if } \frac{1}{\text{rate}_{\text{o}}(P, \{a, \tau\}, 0)} \leq t \\ 0 & \text{if } \frac{1}{\text{rate}_{\text{o}}(P, \{a, \tau\}, 0)} > t \end{cases}
\end{aligned}$$

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

where:

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

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

with p_j representing the conditional probability with which P performs actions whose name is in $\text{init}(\phi_j)$ 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 $P_1, P_2 \in \mathbb{P}_{\text{M,pc}}$:

$$P_1 \sim_{\text{MT}} P_2 \iff \forall \phi \in \mathcal{ML}_{\text{MT}}. \forall \theta \in (\mathbb{R}_{>0})^*. \llbracket \phi \rrbracket_{\text{MT}}^{|\theta|}(P_1, \theta) = \llbracket \phi \rrbracket_{\text{MT}}^{|\theta|}(P_2, \theta)$$

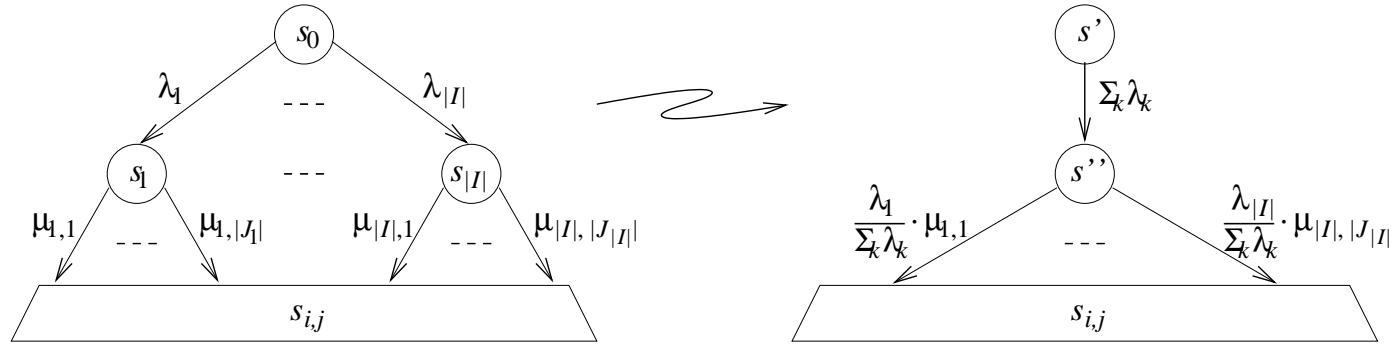
- \sim_{MT} is **decidable in polynomial time** over the set $\mathbb{P}_{\text{M,pc,fin}}$ of finite-state process terms of $\mathbb{P}_{\text{M,pc}}$.
- The reason is that:
 - ◉ \sim_{MT} coincides with the Markovian version of ready equivalence.
 - ◉ Probabilistic ready equivalence can be decided 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 $P_1 \sim_{\text{MT}} P_2$:
 1. Transform $\llbracket P_1 \rrbracket_{\text{M}}$ and $\llbracket P_2 \rrbracket_{\text{M}}$ 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 $\llbracket P_1 \rrbracket_{\text{M}}$ and $\llbracket P_2 \rrbracket_{\text{M}}$ 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 $\llbracket P_1 \rrbracket_{\text{M}}$ and $\llbracket P_2 \rrbracket_{\text{M}}$ 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 $\llbracket P_1 \rrbracket_M$ and $\llbracket P_2 \rrbracket_M$ after reading ς does not belong to the vector space generated by V , then:
 - i. For each $a \in \text{NameReal}_{P_1, P_2}$, add $\varsigma \circ a$ to the queue.
 - ii. Add the state probability vector to V .

4. Build a three-valued state vector \mathbf{u} whose generic element is:
 - a. 0 if it corresponds to a nonaccepting state.
 - b. 1 if it corresponds to an accepting state of $\llbracket P_1 \rrbracket_M$.
 - c. -1 if it corresponds to an accepting state of $\llbracket P_2 \rrbracket_M$.
 5. For each $\mathbf{v} \in V$, check whether $\mathbf{v} \cdot \mathbf{u}^T = 0$.
 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_{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_{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 $\mathbb{P}_{\text{M,pc}}$ are guaranteed to possess the same performance characteristics.

Markovian Trace Equivalence

- 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.

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

$$trace(c) = \alpha$$

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

- Consider subsets of $\mathcal{CC}_{\leq \theta}(P, \alpha)$ including all the trace-compatible computations of the same length.
- They are $\mathcal{CC}_{\leq \theta}^l(P, \alpha)$ for $0 \leq l \leq |\theta|$.
- $\mathcal{CC}_{\leq \theta}^{|\theta|}(P, \alpha)$ 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 ($\lambda > \mu \Rightarrow \frac{1}{\lambda} < \frac{1}{\mu}$):

$$\text{prob}(\mathcal{CC}_{\leq \frac{1}{\lambda}}^1(<\tau, \lambda>.\underline{0}, \varepsilon)) = 1 \neq 0 = \text{prob}(\mathcal{CC}_{\leq \frac{1}{\lambda}}^1(<\tau, \mu>.\underline{0}, \varepsilon))$$
- $P_1 \in \mathbb{P}_{\text{M,pc}}$ is Markovian trace equivalent to $P_2 \in \mathbb{P}_{\text{M,pc}}$, written $P_1 \sim_{\text{MTr}} P_2$, iff for all traces $\alpha \in (\text{Name}_v)^*$ and sequences $\theta \in (\mathbb{R}_{>0})^*$ of average amounts of time:

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

- **Running example** (\sim_{MTr}):

- ◉ Concurrent implementation with two independent one-pos. buffers:

$$PC_{\text{conc},2}^{\text{M}} \triangleq \text{Prod}^{\text{M}} \parallel_{\{\text{deposit}\}} (\text{Buff}^{\text{M}} \parallel_{\emptyset} \text{Buff}^{\text{M}}) \parallel_{\{\text{withdraw}\}} \text{Cons}^{\text{M}}$$

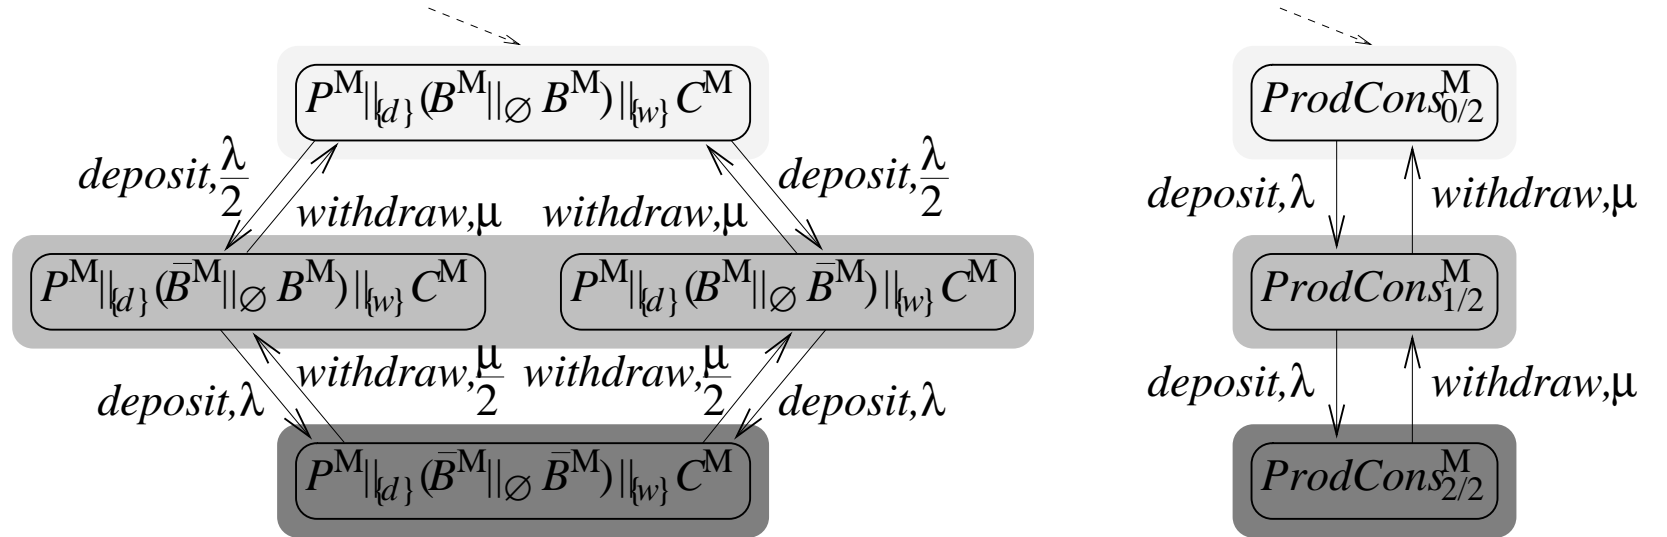
$$\text{Prod}^{\text{M}} \triangleq \langle \text{deposit}, \lambda \rangle . \text{Prod}^{\text{M}}$$

$$\text{Buff}^{\text{M}} \triangleq \langle \text{deposit}, *_1 \rangle . \langle \text{withdraw}, *_1 \rangle . \text{Buff}^{\text{M}}$$

$$\text{Cons}^{\text{M}} \triangleq \langle \text{withdraw}, \mu \rangle . \text{Cons}^{\text{M}}$$

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

- 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 traces to be considered are those coinciding with such prefixes.
- ⊙ Any two computations of $ProdCons_{0/2}^M$ and $PC_{conc,2}^M$ 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.

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

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

$$\begin{aligned} \text{trace}_c(c_k) &= \text{trace}_c(c_h) \\ \text{time}_a(c_k) &= \text{time}_a(c_h) \end{aligned}$$

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

$$\text{rate}_t(P_k^i, 0) = \text{rate}_t(P_h^i, 0)$$

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 $\mu \neq \gamma$):

$$\begin{aligned} &\langle a, \lambda_1 \rangle . \langle b, \mu \rangle . \underline{0} + \langle a, \lambda_2 \rangle . \langle c, \gamma \rangle . \underline{0} \\ &\langle a, \lambda'_1 \rangle . \langle b, \mu \rangle . \underline{0} + \langle a, \lambda'_2 \rangle . \langle c, \gamma \rangle . \underline{0} \end{aligned}$$

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

$$\text{prob}_d(\text{CC}^{|\theta|}(P_1, \alpha), \theta) = \text{prob}_d(\text{CC}^{|\theta|}(P_2, \alpha), \theta)$$

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

$$P_1 \sim_{\text{MTr,d}} P_2 \iff P_1 \sim_{\text{MTr}} P_2$$

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

$$\begin{array}{l} \langle a, \lambda \rangle . P_1 \sim_{\text{MTr}} \langle a, \lambda \rangle . P_2 \\ P_1 + P \sim_{\text{MTr}} P_2 + P \qquad P + P_1 \sim_{\text{MTr}} P + P_2 \end{array}$$

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

$$\begin{array}{l} \langle a, \lambda_1 \rangle . \langle b, \mu \rangle . \underline{0} + \langle a, \lambda_2 \rangle . \langle c, \mu \rangle . \underline{0} \\ \langle a, \lambda_1 + \lambda_2 \rangle . (\langle b, \frac{\lambda_1}{\lambda_1 + \lambda_2} \cdot \mu \rangle . \underline{0} + \langle c, \frac{\lambda_2}{\lambda_1 + \lambda_2} \cdot \mu \rangle . \underline{0}) \end{array}$$

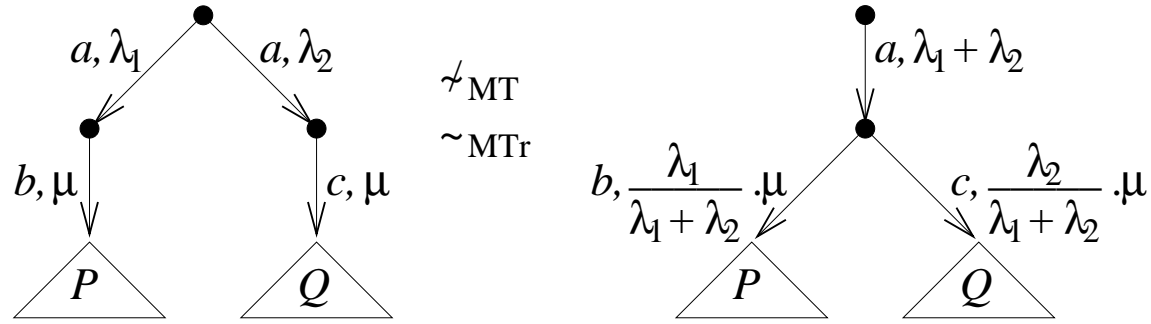
are distinguished in the following context:

$$- \parallel_{\{a,b,c\}} \langle a, * \rangle . \langle b, * \rangle . \underline{0}$$

by the following trace:

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

- \sim_{MTr} has a **sound and complete axiomatization** over the set $\mathbb{P}_{\text{M,pc,dyn}}$ of process terms of $\mathbb{P}_{\text{M,pc}}$ comprising only dynamic operators.
- The axioms for \sim_{MT} are sound but not complete for \sim_{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_{MT}):

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

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

$$\begin{aligned}
 (\mathcal{X}_{\text{MTr},4}) \quad \sum_{i \in I} \langle a, \lambda_i \rangle \cdot \sum_{j \in J_i} \langle b_{i,j}, \mu_{i,j} \rangle \cdot P_{i,j} &= \\
 \langle a, \sum_{k \in I} \lambda_k \rangle \cdot \sum_{i \in I} \sum_{j \in J_i} \langle b_{i,j}, \frac{\lambda_i}{\sum_{k \in I} \lambda_k} \cdot \mu_{i,j} \rangle \cdot P_{i,j}
 \end{aligned}$$

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 $i_1, i_2 \in I$:

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

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

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

- \sim_{MTr} has a **modal logic characterization** over $\mathbb{P}_{\text{M,pc}}$ 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_{MT}).
- Syntax of the modal language $\mathcal{ML}_{\text{MTr}}$ ($a \in \text{Name}_{\text{V}}$):

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

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

$$\llbracket \phi \rrbracket_{\text{MTr}}^{|\theta|}(P, \theta) = \begin{cases} 0 & \text{if } |\theta| = 0 \wedge \phi \not\equiv \text{true} \\ & \text{or } |\theta| > 0 \wedge \text{rate}_t(P, 0) = 0 \\ 1 & \text{if } |\theta| = 0 \wedge \phi \equiv \text{true} \end{cases}$$

otherwise:

$$\llbracket \text{true} \rrbracket_{\text{MTr}}^{|t \circ \theta|}(P, t \circ \theta) = \begin{cases} \sum_{P \xrightarrow{\tau, \lambda}_{\text{M}} P'} \frac{\lambda}{\text{rate}_t(P, 0)} \cdot \llbracket \text{true} \rrbracket_{\text{MTr}}^{|\theta|}(P', \theta) & \text{if } \frac{1}{\text{rate}_t(P, 0)} \leq t \\ 0 & \text{if } \frac{1}{\text{rate}_t(P, 0)} > t \end{cases}$$

$$\llbracket \langle a \rangle \phi \rrbracket_{\text{MTr}}^{|t \circ \theta|}(P, t \circ \theta) = \begin{cases} \sum_{P \xrightarrow{a, \lambda}_{\text{M}} P'} \frac{\lambda}{\text{rate}_t(P, 0)} \cdot \llbracket \phi \rrbracket_{\text{MTr}}^{|\theta|}(P', \theta) + \\ \sum_{P \xrightarrow{\tau, \lambda}_{\text{M}} P'} \frac{\lambda}{\text{rate}_t(P, 0)} \cdot \llbracket \langle a \rangle \phi \rrbracket_{\text{MTr}}^{|\theta|}(P', \theta) & \text{if } \frac{1}{\text{rate}_t(P, 0)} \leq t \\ 0 & \text{if } \frac{1}{\text{rate}_t(P, 0)} > t \end{cases}$$

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

$$P_1 \sim_{\text{MTr}} P_2 \iff \forall \phi \in \mathcal{ML}_{\text{MTr}}. \forall \theta \in (\mathbb{R}_{>0})^*. \llbracket \phi \rrbracket_{\text{MTr}}^{|\theta|}(P_1, \theta) = \llbracket \phi \rrbracket_{\text{MTr}}^{|\theta|}(P_2, \theta)$$

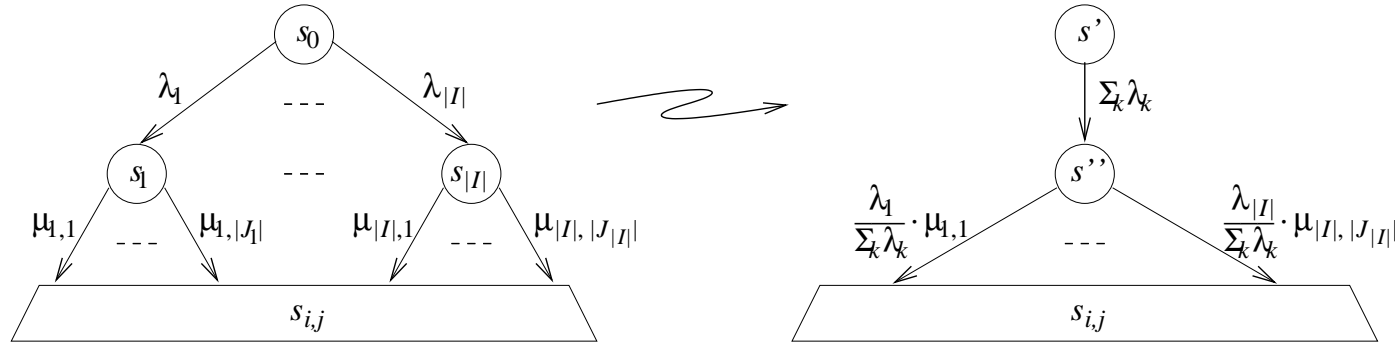
- \sim_{MTr} is **decidable in polynomial time** over the set $\mathbb{P}_{\text{M,pc,fin}}$ of finite-state process terms of $\mathbb{P}_{\text{M,pc}}$.
- Reworking of Tzeng algorithm for probabilistic language equivalence.
- Given two process terms, their name-labeled CTMCs are Markovian trace equivalent 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 $P_1 \sim_{\text{MTr}} P_2$:
 1. Transform $\llbracket P_1 \rrbracket_{\text{M}}$ and $\llbracket P_2 \rrbracket_{\text{M}}$ 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 $\llbracket P_1 \rrbracket_{\text{M}}$ and $\llbracket P_2 \rrbracket_{\text{M}}$ 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 $\llbracket P_1 \rrbracket_M$ and $\llbracket P_2 \rrbracket_M$ after reading ς does not belong to the vector space generated by V , then:
 - i. For each $a \in \text{NameReal}_{P_1, P_2}$, add $\varsigma \circ a$ to the queue.
 - ii. Add the state probability vector to V .

4. Build a three-valued state vector \mathbf{u} whose generic element is:
 - a. 0 if it corresponds to a nonaccepting state.
 - b. 1 if it corresponds to an accepting state of $\llbracket P_1 \rrbracket_M$.
 - c. -1 if it corresponds to an accepting state of $\llbracket P_2 \rrbracket_M$.
 5. For each $\mathbf{v} \in V$, check whether $\mathbf{v} \cdot \mathbf{u}^T = 0$.
 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_{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_{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 $\mathbb{P}_{\text{M,pc}}$ are guaranteed to possess the same performance characteristics.

Summary of Known Results

- Comparing Markovian behavioral equivalences based on given criteria:

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

- 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.

Part IV:
What Remains To Do?

Open Problems

- Markovian behavioral equivalence inducing the coarsest exact nontrivial CTMC-level aggregation?
- Minimization algorithms for \sim_{MT} and \sim_{MT_r} (and T-lumping)?
- Uniform definitions for nondeterministic, probabilistic, and Markovian processes? [*Bernardo - De Nicola - Loreti, 2010*]
- Approximated versions of \sim_{MB} , \sim_{MT} , and \sim_{MT_r} that relax the comparison on exit rates or execution probabilities? [*Aldini, 2010*]
- Weaker versions of \sim_{MB} , \sim_{MT} , and \sim_{MT_r} that abstract from internal exponentially timed actions while preserving nontrivial exactness?

Abstracting from Internal Actions

- When comparing nondeterministic processes, internal actions can be abstracted away via *weak behavioral equivalences*: $a.\tau.b.\underline{0} \approx a.b.\underline{0}$.
- Abstraction not always possible when comparing Markovian processes.
- Immediate internal actions: invisible and take no time [Her,Ret,MT,AB].
- An exponentially timed internal action is invisible **but takes time**.
- $\langle a, \lambda \rangle . \langle \tau, \gamma \rangle . \langle b, \mu \rangle . \underline{0}$ is not equivalent to $\langle a, \lambda \rangle . \langle b, \mu \rangle . \underline{0}$ because a nonzero delay can be observed between a and b in the first case.
- However $\langle a, \lambda \rangle . \langle \tau, \gamma_1 \rangle . \langle \tau, \gamma_2 \rangle . \langle b, \mu \rangle . \underline{0} \approx \langle a, \lambda \rangle . \langle \tau, \gamma \rangle . \langle b, \mu \rangle . \underline{0}$ if the average duration of the sequence of the two τ -actions on the left is equal to the average duration of the τ -action on the right: $\frac{1}{\gamma_1} + \frac{1}{\gamma_2} = \frac{1}{\gamma}$ or equivalently $\gamma = \frac{\gamma_1 \cdot \gamma_2}{\gamma_1 + \gamma_2}$.

- *To what extent can we abstract from exp. timed internal actions?
Only from sequences (of at least two) or also from branches?*
- *How to define a Markovian behavioral equivalence abstracting from
sequences/branches of exponentially timed internal actions?*
- *Will it be a congruence with respect to typical operators?*
- *Will it have a sound and complete axiomatization?*
- *Will it induce an exact CTMC-level aggregation?*
- *Any tradeoff among these properties?*
- *Conduct the study in a process algebraic framework.*

Definition of Weak Markovian Bisimilarity

- Basic idea: weaken the distinguishing power of \sim_{MB} by viewing every sequence of exponentially timed τ -actions as a single exponentially timed τ -action with the same average duration as the sequence.
- $\mathbb{P}_{\text{M},\text{s}}$: set of stable process terms, which can perform no exponentially timed τ -action.
- $\mathbb{P}_{\text{M},\text{u}}$: set of unstable process terms, which can perform at least one exponentially timed τ -action.
- $\mathbb{P}_{\text{M},\text{fu}}$: set of fully unstable process terms, which can perform only exponentially timed τ -actions (most natural candidates for abstraction).
- A computation c having the form $P_1 \xrightarrow{\tau, \lambda_1}_{\text{M}} P_2 \xrightarrow{\tau, \lambda_2}_{\text{M}} \dots \xrightarrow{\tau, \lambda_n}_{\text{M}} P_{n+1}$ is **reducible** iff $P_i \in \mathbb{P}_{\text{M},\text{fu}}$ for all $i = 1, \dots, n$.

- Length-abstracting measure of a reducible computation c :

$$\boxed{proptime(c) = \left(\prod_{i=1}^n \frac{\lambda_i}{rate(P_i, \tau, \mathbb{P}_M)} \right) \cdot \left(\sum_{i=1}^n \frac{1}{rate(P_i, \tau, \mathbb{P}_M)} \right)}$$

- The first factor is the product of the execution probabilities of the transitions of c .
- The second factor is the sum of the average sojourn times of the states traversed by c .
- Finite-length reducible computations are enough to distinguish between fully unstable process terms that must be told apart ($\lambda_1 \neq \lambda_2$ and $a \in Name_V$):

$$\langle \tau, \lambda_1 \rangle . \langle a, \lambda \rangle . P \quad \text{vs.} \quad \langle \tau, \lambda_2 \rangle . \langle a, \lambda \rangle . P$$

$$\text{rec } X : \langle \tau, \lambda_1 \rangle . X \quad \text{vs.} \quad \text{rec } X : \langle \tau, \lambda_2 \rangle . X$$

- The weak variant of \sim_{MB} should work like \sim_{MB} over $\mathbb{P}_{\text{M},\text{nfu}}$ and abstract from the length of reducible computations while preserving their execution probability and average duration over $\mathbb{P}_{\text{M},\text{fu}}$.
- Need to lift measure *probtme* from a single reducible computation to a multiset of reducible computations with the same origin and destination:

$$pbtm(P, D) = \{ \sum_{c \in \text{redcomp}(P, D, t)} \text{probtme}(c) \mid t \in \mathbb{R}_{>0} \}$$

where $\text{redcomp}(P, D, t)$ is the multiset of reducible computations from $P \in \mathbb{P}_{\text{M}}$ to $D \subseteq \mathbb{P}_{\text{M}}$ whose average duration is $t \in \mathbb{R}_{>0}$.

- Measures must be summed up (otherwise $\langle \tau, \lambda_1 \rangle . \underline{0} + \langle \tau, \lambda_2 \rangle . \underline{0}$ not equivalent to $\langle \tau, \lambda_1 + \lambda_2 \rangle . \underline{0}$) only in case of equal average durations.

- An equivalence relation \mathcal{B} over \mathbb{P}_M is a **weak Markovian bisimulation** iff, whenever $(P_1, P_2) \in \mathcal{B}$, then for all equivalence classes $D \in \mathbb{P}_M/\mathcal{B}$:
 - If $P_1, P_2 \in \mathbb{P}_{M,\text{nfu}}$, for all action names $a \in \text{Name}$:

$$\text{rate}(P_1, a, D) = \text{rate}(P_2, a, D)$$
 - If $P_1, P_2 \in \mathbb{P}_{M,\text{fu}}$, when $D \subseteq \mathbb{P}_{M,\text{nfu}}$:

$$\text{pbtm}(P_1, D) = \text{pbtm}(P_2, D)$$
- Weak Markovian bisimilarity, denoted \approx_{MB} , is the union of all the weak Markovian bisimulations.

- **Example 1** – Consider the two process terms:

$$\bar{P}_1 \equiv \langle \tau, \mu \rangle . \langle \tau, \gamma \rangle . Q \quad (\equiv \langle \tau, \gamma \rangle . \langle \tau, \mu \rangle . Q)$$

$$\bar{P}_2 \equiv \langle \tau, \frac{\mu \cdot \gamma}{\mu + \gamma} \rangle . Q$$

with $Q \in \mathbb{P}_{M, \text{nfu}}$.

- Then $\bar{P}_1 \approx_{\text{MB}} \bar{P}_2$ because:

$$\begin{aligned} pbtm(\bar{P}_1, [Q]_{\approx_{\text{MB}}}) &= \{ (1 \cdot 1) \cdot (\frac{1}{\mu} + \frac{1}{\gamma}) \} = \\ &= \{ 1 \cdot \frac{\mu + \gamma}{\mu \cdot \gamma} \} = pbtm(\bar{P}_2, [Q]_{\approx_{\text{MB}}}) \end{aligned}$$

- In general, for $l \in \mathbb{N}_{>0}$ we have:

$$\langle \tau, \mu \rangle . \langle \tau, \gamma_1 \rangle . \dots . \langle \tau, \gamma_l \rangle . Q \approx_{\text{MB}} \langle \tau, \left(\frac{1}{\mu} + \frac{1}{\gamma_1} + \dots + \frac{1}{\gamma_l} \right)^{-1} \rangle . Q$$

- **Example 2** – Consider the two process terms:

$$\bar{P}_3 \equiv \langle \tau, \mu \rangle . (\langle \tau, \gamma_1 \rangle . Q_1 + \langle \tau, \gamma_2 \rangle . Q_2)$$

$$\bar{P}_4 \equiv \langle \tau, \frac{\gamma_1}{\gamma_1 + \gamma_2} \cdot \left(\frac{1}{\mu} + \frac{1}{\gamma_1 + \gamma_2} \right)^{-1} \rangle . Q_1 + \langle \tau, \frac{\gamma_2}{\gamma_1 + \gamma_2} \cdot \left(\frac{1}{\mu} + \frac{1}{\gamma_1 + \gamma_2} \right)^{-1} \rangle . Q_2$$

with $Q_1, Q_2 \in \mathbb{P}_{M, \text{nfu}}$ and $Q_1 \not\approx_{\text{MB}} Q_2$.

- Then $\bar{P}_3 \approx_{\text{MB}} \bar{P}_4$ because:

$$pbtm(\bar{P}_3, [Q_1]_{\approx_{\text{MB}}}) = \left\{ \left\langle \frac{\gamma_1}{\gamma_1 + \gamma_2} \cdot \left(\frac{1}{\mu} + \frac{1}{\gamma_1 + \gamma_2} \right)^{-1} \right\rangle \right\} = pbtm(\bar{P}_4, [Q_1]_{\approx_{\text{MB}}})$$

$$pbtm(\bar{P}_3, [Q_2]_{\approx_{\text{MB}}}) = \left\{ \left\langle \frac{\gamma_2}{\gamma_1 + \gamma_2} \cdot \left(\frac{1}{\mu} + \frac{1}{\gamma_1 + \gamma_2} \right)^{-1} \right\rangle \right\} = pbtm(\bar{P}_4, [Q_2]_{\approx_{\text{MB}}})$$

- In general, for $n \in \mathbb{N}_{>0}$ we have:

$$\begin{aligned} \langle \tau, \mu \rangle . (\langle \tau, \gamma_1 \rangle . Q_1 + \dots + \langle \tau, \gamma_n \rangle . Q_n) &\approx_{\text{MB}} \langle \tau, \frac{\gamma_1}{\gamma_1 + \dots + \gamma_n} \cdot \left(\frac{1}{\mu} + \frac{1}{\gamma_1 + \dots + \gamma_n} \right)^{-1} \rangle . Q_1 + \\ &\dots + \\ &\langle \tau, \frac{\gamma_n}{\gamma_1 + \dots + \gamma_n} \cdot \left(\frac{1}{\mu} + \frac{1}{\gamma_1 + \dots + \gamma_n} \right)^{-1} \rangle . Q_n \end{aligned}$$

- **Example 3** – Consider the two process terms:

$$\bar{P}_5 \equiv \langle \tau, \mu_1 \rangle . \langle \tau, \gamma \rangle . Q_1 + \langle \tau, \mu_2 \rangle . \langle \tau, \gamma \rangle . Q_2$$

$$\bar{P}_6 \equiv \langle \tau, \frac{\mu_1}{\mu_1 + \mu_2} \cdot \left(\frac{1}{\mu_1 + \mu_2} + \frac{1}{\gamma} \right)^{-1} \rangle . Q_1 + \langle \tau, \frac{\mu_2}{\mu_1 + \mu_2} \cdot \left(\frac{1}{\mu_1 + \mu_2} + \frac{1}{\gamma} \right)^{-1} \rangle . Q_2$$

with $Q_1, Q_2 \in \mathbb{P}_{M, \text{nfu}}$ and $Q_1 \not\approx_{\text{MB}} Q_2$.

- Then $\bar{P}_5 \approx_{\text{MB}} \bar{P}_6$ because:

$$pbtm(\bar{P}_5, [Q_1]_{\approx_{\text{MB}}}) = \left\{ \frac{\mu_1}{\mu_1 + \mu_2} \cdot \left(\frac{1}{\mu_1 + \mu_2} + \frac{1}{\gamma} \right) \right\} = pbtm(\bar{P}_6, [Q_1]_{\approx_{\text{MB}}})$$

$$pbtm(\bar{P}_5, [Q_2]_{\approx_{\text{MB}}}) = \left\{ \frac{\mu_2}{\mu_1 + \mu_2} \cdot \left(\frac{1}{\mu_1 + \mu_2} + \frac{1}{\gamma} \right) \right\} = pbtm(\bar{P}_6, [Q_2]_{\approx_{\text{MB}}})$$

- In general, for $n \in \mathbb{N}_{>0}$ we have:

$$\begin{aligned} \langle \tau, \mu_1 \rangle . \langle \tau, \gamma \rangle . Q_1 + \dots + \langle \tau, \mu_n \rangle . \langle \tau, \gamma \rangle . Q_n &\approx_{\text{MB}} \langle \tau, \frac{\mu_1}{\mu_1 + \dots + \mu_n} \cdot \left(\frac{1}{\mu_1 + \dots + \mu_n} + \frac{1}{\gamma} \right)^{-1} \rangle . Q_1 + \\ &\dots + \\ &\langle \tau, \frac{\mu_n}{\mu_1 + \dots + \mu_n} \cdot \left(\frac{1}{\mu_1 + \dots + \mu_n} + \frac{1}{\gamma} \right)^{-1} \rangle . Q_n \end{aligned}$$

- **Example 4** – None of the variants of \bar{P}_5 related to actions $\langle \tau, \gamma \rangle$ leads to a reduction.

- If we consider:

$$\bar{P}_7 \equiv \langle \tau, \mu_1 \rangle . \langle \tau, \gamma_1 \rangle . Q_1 + \langle \tau, \mu_2 \rangle . \langle \tau, \gamma_2 \rangle . Q_2$$

$$\bar{P}_8 \equiv \langle \tau, \frac{\mu_1}{\mu_1 + \mu_2} \cdot \left(\frac{1}{\mu_1 + \mu_2} + \frac{1}{\gamma_1} \right)^{-1} \rangle . Q_1 + \langle \tau, \frac{\mu_2}{\mu_1 + \mu_2} \cdot \left(\frac{1}{\mu_1 + \mu_2} + \frac{1}{\gamma_2} \right)^{-1} \rangle . Q_2$$

with $\gamma_1 \neq \gamma_2$, then $\bar{P}_7 \not\approx_{\text{MB}} \bar{P}_8$.

- If we consider:

$$\bar{P}_9 \equiv \langle \tau, \mu_1 \rangle . \langle \tau, \gamma \rangle . Q_1 + \langle \tau, \mu_2 \rangle . Q_2$$

$$\bar{P}_{10} \equiv \langle \tau, \frac{\mu_1}{\mu_1 + \mu_2} \cdot \left(\frac{1}{\mu_1 + \mu_2} + \frac{1}{\gamma} \right)^{-1} \rangle . Q_1 + \langle \tau, \mu_2 \rangle . Q_2$$

then $\bar{P}_9 \not\approx_{\text{MB}} \bar{P}_{10}$.

- Let:

- $I \neq \emptyset$ be a finite index set.
- $J_i \neq \emptyset$ be a finite index set for all $i \in I$.
- $P_{i,j} \in \mathbb{P}_M$ for all $i \in I$ and $j \in J_i$.

Whenever $\sum_{j \in J_1} \gamma_{i_1,j} = \sum_{j \in J_2} \gamma_{i_2,j}$ for all $i_1, i_2 \in I$, then:

$$\sum_{i \in I} \langle \tau, \mu_i \rangle \cdot \sum_{j \in J_i} \langle \tau, \gamma_{i,j} \rangle \cdot P_{i,j} \approx_{\text{MB}} \sum_{i \in I} \sum_{j \in J_i} \langle \tau, \frac{\mu_i}{\sum_{k \in I} \mu_k} \cdot \frac{\gamma_{i,j}}{\sum_{h \in J_i} \gamma_{i,h}} \cdot \left(\frac{1}{\sum_{k \in I} \mu_k} + \frac{1}{\sum_{h \in J_i} \gamma_{i,h}} \right)^{-1} \rangle \cdot P_{i,j}$$

Congruence Property

- \approx_{MB} is a congruence with respect to action prefix and hiding.
- Let $P_1, P_2 \in \mathbb{P}_M$. Whenever $P_1 \approx_{\text{MB}} P_2$, then:
 - $\langle a, \lambda \rangle.P_1 \approx_{\text{MB}} \langle a, \lambda \rangle.P_2$ for all $\langle a, \lambda \rangle \in \text{Act}_M$.
 - $P_1/H \approx_{\text{MB}} P_2/H$ for all $H \subseteq \text{Name}_v$.
- Not a congruence with respect to alternative and parallel composition due to fully unstable process terms:

$$\langle \tau, \mu \rangle. \langle \tau, \gamma \rangle. \underline{0} \approx_{\text{MB}} \langle \tau, \frac{\mu \cdot \gamma}{\mu + \gamma} \rangle. \underline{0}$$

but:

$$\langle \tau, \mu \rangle. \langle \tau, \gamma \rangle. \underline{0} + \langle a, \lambda \rangle. \underline{0} \not\approx_{\text{MB}} \langle \tau, \frac{\mu \cdot \gamma}{\mu + \gamma} \rangle. \underline{0} + \langle a, \lambda \rangle. \underline{0}$$

$$\langle \tau, \mu \rangle. \langle \tau, \gamma \rangle. \underline{0} \parallel_{\emptyset} \langle a, \lambda \rangle. \underline{0} \not\approx_{\text{MB}} \langle \tau, \frac{\mu \cdot \gamma}{\mu + \gamma} \rangle. \underline{0} \parallel_{\emptyset} \langle a, \lambda \rangle. \underline{0}$$

both for $a \neq \tau$ and for $a = \tau$.

- In order to avoid congruence violations for alternative composition, apply the exit rate equality check also to fully unstable process terms (but consider the equivalence classes with respect to \approx_{MB}).
- We say that P_1 is weakly Markovian bisimulation congruent to P_2 , written $P_1 \simeq_{\text{MB}} P_2$, iff for all action names $a \in \text{Name}$ and equivalence classes $D \in \mathbb{P}_{\text{M}}/\approx_{\text{MB}}$:

$$\text{rate}(P_1, a, D) = \text{rate}(P_2, a, D)$$

- $\sim_{\text{MB}} \subset \simeq_{\text{MB}} \subset \approx_{\text{MB}}$, with \simeq_{MB} and \approx_{MB} coinciding over $\mathbb{P}_{\text{M}, \text{nfu}}$.
- $\langle a, \lambda \rangle.P_1 \simeq_{\text{MB}} \langle a, \lambda \rangle.P_2$ iff $P_1 \approx_{\text{MB}} P_2$.

- Let $P_1, P_2 \in \mathbb{P}_M$. Whenever $P_1 \simeq_{\text{MB}} P_2$, then:
 - $\langle a, \lambda \rangle.P_1 \simeq_{\text{MB}} \langle a, \lambda \rangle.P_2$ for all $\langle a, \lambda \rangle \in \text{Act}_M$.
 - $P_1 + P \simeq_{\text{MB}} P_2 + P$ and $P + P_1 \simeq_{\text{MB}} P + P_2$ for all $P \in \mathbb{P}_M$.
 - $P_1/H \simeq_{\text{MB}} P_2/H$ for all $H \subseteq \text{Name}_v$.
- \simeq_{MB} is the coarsest congruence contained in \approx_{MB} over the set $\mathbb{P}_{M,\text{seq}}$ of process terms of \mathbb{P}_M that do not contain any occurrence of the parallel composition operator.
- Let $P_1, P_2 \in \mathbb{P}_{M,\text{seq}}$. Then $P_1 \simeq_{\text{MB}} P_2$ iff $P_1 + P \approx_{\text{MB}} P_2 + P$ for all $P \in \mathbb{P}_{M,\text{seq}}$.

Sound and Complete Axiomatization

- \simeq_{MB} has a sound and complete axiomatization over the set $\mathbb{P}_{\text{M,seq,nrec}}$ of nonrecursive process terms of $\mathbb{P}_{\text{M,seq}}$.
- Set of basic axioms (the first four coincide with those of \sim_{MB}):

$$\begin{aligned}
 (\mathcal{X}_{\text{MB},1}) \quad & P_1 + P_2 = P_2 + P_1 \\
 (\mathcal{X}_{\text{MB},2}) \quad & (P_1 + P_2) + P_3 = P_1 + (P_2 + P_3) \\
 (\mathcal{X}_{\text{MB},3}) \quad & P + \underline{0} = P \\
 (\mathcal{X}_{\text{MB},4}) \quad & \langle a, \lambda_1 \rangle . P + \langle a, \lambda_2 \rangle . P = \langle a, \lambda_1 + \lambda_2 \rangle . P \\
 (\mathcal{X}_{\text{MB},5}) \quad & \langle a, \lambda \rangle . \sum_{i \in I} \langle \tau, \mu_i \rangle . \sum_{j \in J_i} \langle \tau, \gamma_{i,j} \rangle . P_{i,j} = \\
 & \langle a, \lambda \rangle . \sum_{i \in I} \sum_{j \in J_i} \langle \tau, \frac{\mu_i}{\mu} \cdot \frac{\gamma_{i,j}}{\gamma} \cdot \left(\frac{1}{\mu} + \frac{1}{\gamma} \right)^{-1} \rangle . P_{i,j}
 \end{aligned}$$

if: $I \neq \emptyset$ is a finite index set
 $J_i \neq \emptyset$ is a finite index set for all $i \in I$
 $\mu = \sum_{i \in I} \mu_i$
 $\gamma = \sum_{j \in J_i} \gamma_{i,j}$ for all $i \in I$

- For proving completeness, we cannot resort to normal form saturation as this would alter the quantitative behavior.
- Let $P_1, P_2 \in \mathbb{P}_{\text{M,seq,nrec}}$. If $P_1 \approx_{\text{MB}} P_2$ but $P_1 \not\approx_{\text{MB}} P_2$, then at least one between P_1 and P_2 (both of which must be fully unstable) is of the form:

$$\sum_{i \in I} \langle \tau, \mu_i \rangle \cdot \sum_{j \in J_i} \langle \tau, \gamma_{i,j} \rangle \cdot P_{i,j}$$

where $I \neq \emptyset$ is a finite index set, $J_i \neq \emptyset$ is a finite index set for all $i \in I$, and one of the following two properties holds:

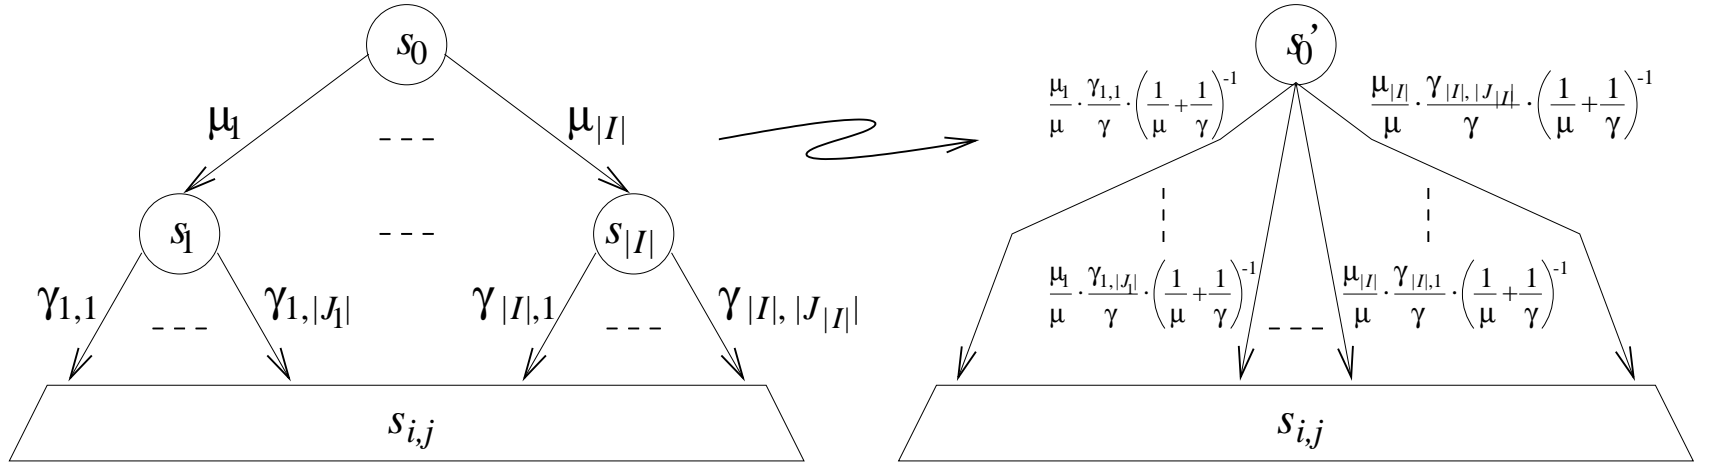
- $\sum_{j \in J_{i_1}} \langle \tau, \gamma_{i_1,j} \rangle \cdot P_{i_1,j} \approx_{\text{MB}} \sum_{j \in J_{i_2}} \langle \tau, \gamma_{i_2,j} \rangle \cdot P_{i_2,j}$ for all $i_1, i_2 \in I$.
- $\sum_{j \in J_{i_1}} \gamma_{i_1,j} = \sum_{j \in J_{i_2}} \gamma_{i_2,j}$ for all $i_1, i_2 \in I$.

- Let $P_1, P_2 \in \mathbb{P}_{\text{M,seq,nrec}}$. Then $\mathcal{X}_{\text{MB},1..5} \vdash P_1 = P_2 \iff P_1 \simeq_{\text{MB}} P_2$.

Exactness of CTMC-Level Aggregation

- \approx_{MB} and \simeq_{MB} allow every sequence of exponentially timed τ -actions to be considered equivalent to a single exponentially timed τ -action having the same average duration.
- This amounts to approximating a hypoexponentially or Erlang distributed random variable with an exponentially distributed random variable having the same expected value.
- This can be exploited to assess more quickly properties expressed in terms of the mean time to certain events.
- *Is there any other performance property that is preserved?*

- Since \sim_{MB} is consistent with ordinary lumpability and the only new axiom is $\mathcal{X}_{\text{MB},5}$, we can concentrate on this axiom.
- The induced CTMC-level aggregation, called **W-lumpability**, eliminates $|I|$ states and $|I|$ transitions by merging the first $1 + |I|$ states into a single one:



- W-lumpability is exact at steady state, i.e., the stationary probability of being in a macrostate of a CTMC obtained via W-lumpability is the sum of the stationary probabilities of being in one of the constituent microstates of the CTMC from which the reduced one has been derived.
- Unlike ordinary lumpability and T-lumpability, properties expressed in terms of transient state probabilities may not be preserved.
- Reconsider $\bar{P}_1 \equiv \langle \tau, \mu \rangle \cdot \langle \tau, \gamma \rangle \cdot Q$ and $\bar{P}_2 \equiv \langle \tau, \frac{\mu \cdot \gamma}{\mu + \gamma} \rangle \cdot Q$.
- The probability of being in the first state of $[\bar{P}_2]_M$ at time $t \in \mathbb{R}_{>0}$ is $1 - (1 - e^{-\frac{\mu \cdot \gamma}{\mu + \gamma} \cdot t}) = e^{-\frac{\mu \cdot \gamma}{\mu + \gamma} \cdot t}$, which reduces to $e^{-\frac{\mu}{2} \cdot t}$ when $\mu = \gamma$.
- The sum of the probabilities of being in one of the first two states of $[\bar{P}_1]_M$ at the same time instant is $\frac{\gamma}{\gamma - \mu} \cdot e^{-\mu \cdot t} - \frac{\mu}{\gamma - \mu} \cdot e^{-\gamma \cdot t}$ for $\mu \neq \gamma$ or $(1 + \mu \cdot t) \cdot e^{-\mu \cdot t}$ for $\mu = \gamma$.

Generalization to Concurrent Processes

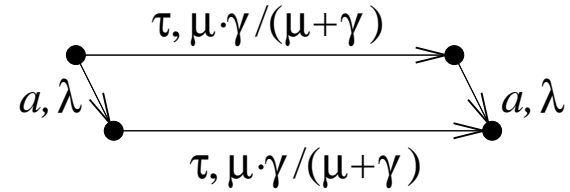
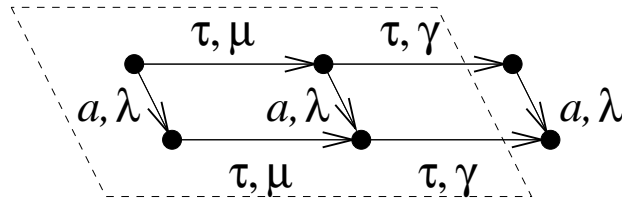
- Not being a congruence with respect to parallel composition significantly reduces the usefulness of \simeq_{MB} for compositional state space reduction.
- It is possible to modify the definition of the equivalence so that it becomes a congruence with respect to parallel composition ...
- ...but exactness will hold at steady state only for a certain class of processes.
- Revise the notion of reducible computation by admitting the traversal of unstable states (that are not fully unstable) satisfying certain conditions.
- Local computations may traverse fully unstable local states that are part of global states that are not fully unstable.

- The following two process terms:

$$\langle \tau, \mu \rangle . \langle \tau, \gamma \rangle . \underline{0} \parallel_{\emptyset} \langle a, \lambda \rangle . \underline{0}$$

$$\langle \tau, \frac{\mu \cdot \gamma}{\mu + \gamma} \rangle . \underline{0} \parallel_{\emptyset} \langle a, \lambda \rangle . \underline{0}$$

should be considered equivalent and should give rise to the following CTMC-level aggregation:



- Trees of sequences of exponentially timed τ -actions may be replicated.
- Take into account at once all the corresponding trees of computations and pinpoint their initial and final states.

- Let $P \in \mathbb{P}_M$, $m, n \in \mathbb{N}_{>0}$, and $P_1, P_2, \dots, P_m \in \mathbb{P}_M$ reachable from P and different from each other.
- Let C_1, C_2, \dots, C_m be m sets each containing n distinct finite-length computations all starting from P_1, P_2, \dots, P_m , respectively, with:
 - Each computation traversing different states except at most the last state and one of the preceding states.
 - Computations in different sets being disjoint.
- Assume that those computations can be partitioned into n groups each consisting of m computations from all the m sets, such that all the computations in the same group have the same length and are labeled with the same sequence of exponentially timed τ -actions.
- $$C_k = \{c_{k,i} \mid c_{k,i} \equiv P_{k,i,1} \xrightarrow{\tau, \lambda_{i,1}}_M P_{k,i,2} \xrightarrow{\tau, \lambda_{i,2}}_M \dots \xrightarrow{\tau, \lambda_{i,l_i}}_M P_{k,i,l_i+1}, 1 \leq i \leq n\}$$

- The family of computations $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ is generally reducible iff either $m = 1$ and for all $i = 1, \dots, n$:
 - $P_{1,i,j} \in \mathbb{P}_{M, \text{fu}}$ for all $j = 1, \dots, l_i$.
 - $P_{1,i,l_i+1} \in \mathbb{P}_{M, \text{nfu}}$ or $P_{1,i,l_i+1} \equiv P_{1,i,j}$ for some $j = 1, \dots, l_i$.
- or $m \geq 1$, not all the states $P_{k,i,j}$, $1 \leq j \leq l_i$, are fully unstable when $m = 1$, and for all $i = 1, \dots, n$:
 - For all $k = 1, \dots, m$, $j = 1, \dots, l_i$, and $\langle a, \lambda \rangle \in \text{Act}_M$:
 1. Whenever $P_{k,i,j} \xrightarrow{a, \lambda}_M P'$ with $P' \not\equiv P_{k,i,j+1}$, then:
 - a) either $P' \equiv P_{k',i,j}$ for some $k' = 1, \dots, m$;
 - b) or $P' \equiv P_{k,i',j'}$ with $a = \tau$ and $\lambda = \lambda_{i',j'-1}$ for some $i' = 1, \dots, n$ such that $i' \neq i$ and some $j' = 2, \dots, l_{i'}+1$.

2. For all $k' = 1, \dots, m$, it holds that $P_{k,i,j} \xrightarrow{a,\lambda}_M P_{k',i,j}$ iff $P_{k,i,j'} \xrightarrow{a,\lambda}_M P_{k',i,j'}$ for all $j' = 1, \dots, l_i$.
3. For all $i' = 1, \dots, n$ such that $i' \neq i$ and $j' = 2, \dots, l_{i'+1}$, it holds that $P_{k,i,j} \xrightarrow{a,\lambda}_M P_{k,i',j'}$ iff $P_{k',i,j} \xrightarrow{a,\lambda}_M P_{k',i',j'}$ for all $k' = 1, \dots, m$.

– One of the following holds:

- $\widetilde{4}$. There is no $\lambda_{i,l_i+1} \in \mathbb{R}_{>0}$ such that $P_{k,i,l_i+1} \xrightarrow{\tau,\lambda_{i,l_i+1}}_M P_{k,i,l_i+2}$ for all $k = 1, \dots, m$.
- $\overline{4}$. If there exists $\lambda_{i,l_i+1} \in \mathbb{R}_{>0}$ such that $P_{k,i,l_i+1} \xrightarrow{\tau,\lambda_{i,l_i+1}}_M P_{k,i,l_i+2}$ for all $k = 1, \dots, m$, then at least one of conditions 1, 2, and 3 above is not satisfied by P_{k',i,l_i+1} for some $k' = 1, \dots, m$.
- $\widehat{4}$. $P_{k,i,l_i+1} \equiv P_{k,i,j}$ for all $k = 1, \dots, m$ and some $j = 1, \dots, l_i$.

- $source(\mathcal{C}) = \{P_k \mid 1 \leq k \leq m\}$: set of initial states of \mathcal{C} .
- $target(\mathcal{C}) = \{P_{k,i,l_i+1} \mid 1 \leq k \leq m, 1 \leq i \leq n\}$: set of final states of \mathcal{C} .
- In order to avoid interferences among computations in C_1, C_2, \dots, C_m and computations across C_1, C_2, \dots, C_m , we redefine:

$$proptime(c_{k,i}) = \left(\prod_{j=1}^{l_i} \frac{\lambda_{i,j}}{rate(P_{k,i,j}, \tau, \mathcal{P}'_k)} \right) \cdot \left(\sum_{i=1}^{l_i} \frac{1}{rate(P_{k,i,j}, \tau, \mathcal{P}'_k)} \right)$$

where $\mathcal{P}'_k = \{P_{k,i',j'} \mid 1 \leq i' \leq n, 2 \leq j' \leq l_{i'+1}\}$.

- We redefine $reducomp(P_k, D, t)$ as the multiset of computations identical to those in C_k that go from P_k to D and have average duration t .
- If \mathcal{C} is generally reducible, then $proptime(c_{k,i}) = proptime(c_{k'},i)$ and $pbtm(P_k, target(\mathcal{C})) = pbtm(P_{k'}, target(\mathcal{C}))$.

- An equivalence relation \mathcal{B} over \mathbb{P}_M is a **g-weak Markovian bisimulation** iff, whenever $(P_1, P_2) \in \mathcal{B}$, then:
 - For all visible action names $a \in Name_v$ and equivalence classes $D \in \mathbb{P}_M/\mathcal{B}$:

$$rate(P_1, a, D) = rate(P_2, a, D)$$
 - If P_1 is not an initial state of any g-reducible family of computations, then P_2 is not an initial state of any g-reducible family of computations either, and for all equivalence classes $D \in \mathbb{P}_M/\mathcal{B}$:

$$rate(P_1, \tau, D) = rate(P_2, \tau, D)$$
 - If P_1 is an initial state of some g-reducible family of computations, then P_2 is an initial state of some g-reducible family of computations too, and for all g-reducible families of computations \mathcal{C}_1 with $P_1 \in source(\mathcal{C}_1)$ there exists a g-reducible family of computations \mathcal{C}_2 with $P_2 \in source(\mathcal{C}_2)$ such that for all equivalence classes $D \in \mathbb{P}_M/\mathcal{B}$:

$$pbtm(P_1, D \cap target(\mathcal{C}_1)) = pbtm(P_2, D \cap target(\mathcal{C}_2))$$
- G-weak Markovian bisimilarity, denoted $\approx_{MB,g}$, is the union of all the g-weak Markovian bisimulations.

- All the examples that we have seen before for \approx_{MB} are valid for $\approx_{\text{MB,g}}$, because a tree of computations reducible in the sense of the original definition forms a g-reducible family of computations.
- Unlike \approx_{MB} , it turns out that $\approx_{\text{MB,g}}$ is a congruence with respect to parallel composition too.
- Let $P_1, P_2 \in \mathbb{P}_M$. Whenever $P_1 \approx_{\text{MB,g}} P_2$, then:
 - $\langle a, \lambda \rangle.P_1 \approx_{\text{MB,g}} \langle a, \lambda \rangle.P_2$ for all $\langle a, \lambda \rangle \in \text{Act}_M$.
 - $P_1 \parallel_S P \approx_{\text{MB,g}} P_2 \parallel_S P$ and $P \parallel_S P_1 \approx_{\text{MB,g}} P \parallel_S P_2$ for all $S \subseteq \text{Name}_v$ and $P \in \mathbb{P}_M$.
 - $P_1/H \approx_{\text{MB,g}} P_2/H$ for all $H \subseteq \text{Name}_v$.
- Like \approx_{MB} , we have that $\approx_{\text{MB,g}}$ is not a congruence with respect to alternative composition either.

- We say that P_1 is g-weakly Markovian bisimulation congruent to P_2 , written $P_1 \simeq_{\text{MB,g}} P_2$, iff for all action names $a \in \text{Name}$ and equivalence classes $D \in \mathbb{P}_M / \approx_{\text{MB,g}}$:

$$\text{rate}(P_1, a, D) = \text{rate}(P_2, a, D)$$

- Let $P_1, P_2 \in \mathbb{P}_M$. Whenever $P_1 \simeq_{\text{MB,g}} P_2$, then:
 - $\langle a, \lambda \rangle . P_1 \simeq_{\text{MB,g}} \langle a, \lambda \rangle . P_2$ for all $\langle a, \lambda \rangle \in \text{Act}_M$.
 - $P_1 + P \simeq_{\text{MB,g}} P_2 + P$ and $P + P_1 \simeq_{\text{MB,g}} P + P_2$ for all $P \in \mathbb{P}_M$.
 - $P_1 \parallel_S P \simeq_{\text{MB,g}} P_2 \parallel_S P$ and $P \parallel_S P_1 \simeq_{\text{MB,g}} P \parallel_S P_2$ for all $S \subseteq \text{Name}_v$ and $P \in \mathbb{P}_M$.
 - $P_1 / H \simeq_{\text{MB,g}} P_2 / H$ for all $H \subseteq \text{Name}_v$.
- Let $P_1, P_2 \in \mathbb{P}_M$. Then $P_1 \simeq_{\text{MB,g}} P_2$ iff $P_1 + P \approx_{\text{MB,g}} P_2 + P$ for all $P \in \mathbb{P}_M$.

- The CTMC-level aggregation induced by $\approx_{\text{MB},g}$ and $\simeq_{\text{MB},g}$ is exact at steady-state only for those process terms with a restricted use of synchronization.
- This limitation stems from insensitivity conditions for GSMPs (with GSMPs coming into play due to the reduction of sequences of exponentially timed τ -transitions) and emphasizes a tradeoff between achieving compositionality over concurrent processes and preserving exactness at steady state.
- GW-lumpability is exact at steady state over each process term $P \in \mathbb{P}_M$ such that, for all g-reducible families of computations \mathcal{C} in $\llbracket P \rrbracket_M$ with size $m \geq 2$ or size $m = 1$ and not all the traversed states being fully unstable, no transition in $\llbracket P \rrbracket_M$ arising from action synchronization has an element of $\text{source}(\mathcal{C})$ as its target state.

- **Example 5** – Consider the two process terms with synchronization:

$$\text{rec } X : \langle \tau, \mu \rangle . \langle \tau, \gamma \rangle . \langle d, \delta \rangle . X \parallel_{\{d\}} \text{rec } Y : \langle a, \lambda \rangle . \langle d, *_w \rangle . Y$$

$$\text{rec } X : \langle \tau, \frac{\mu \cdot \gamma}{\mu + \gamma} \rangle . \langle d, \delta \rangle . X \parallel_{\{d\}} \text{rec } Y : \langle a, \lambda \rangle . \langle d, *_w \rangle . Y$$

- Resulting CTMC-level aggregation:



- Not exact due to the following steady-state probabilities ($\mu = \gamma = \lambda = \delta = 1$):

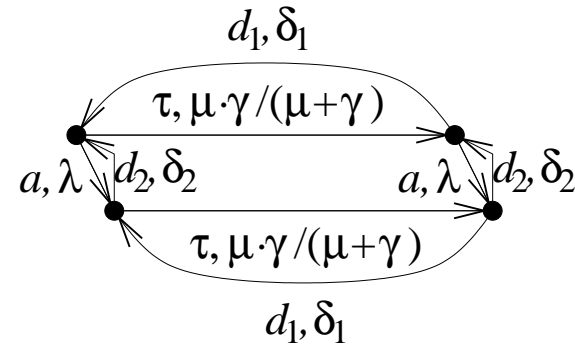
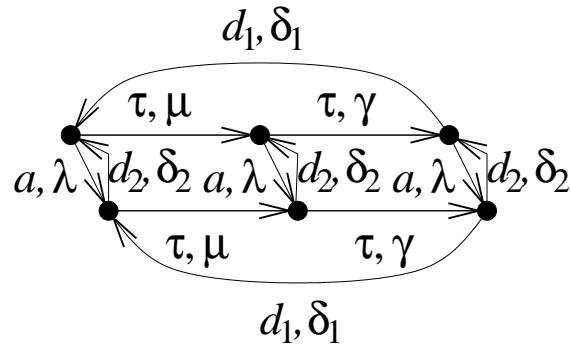
2/13	1/13	1/13	2/10	1/10
2/13	3/13	4/13	4/10	3/10

- **Example 6** – Consider the two process terms without synchronization:

$$\text{rec } X : \langle \tau, \mu \rangle . \langle \tau, \gamma \rangle . \langle d_1, \delta_1 \rangle . X \parallel_{\emptyset} \text{rec } Y : \langle a, \lambda \rangle . \langle d_2, \delta_2 \rangle . Y$$

$$\text{rec } X : \langle \tau, \frac{\mu \cdot \gamma}{\mu + \gamma} \rangle . \langle d_1, \delta_1 \rangle . X \parallel_{\emptyset} \text{rec } Y : \langle a, \lambda \rangle . \langle d_2, \delta_2 \rangle . Y$$

- Resulting CTMC-level aggregation:



- Exact due to the following steady-state probabilities ($\mu = \gamma = \lambda = \delta_1 = \delta_2 = 1$):

$$\begin{array}{ccc} 1/6 & 1/6 & 1/6 \\ 1/6 & 1/6 & 1/6 \end{array}$$

$$\begin{array}{cc} 2/6 & 1/6 \\ 2/6 & 1/6 \end{array}$$

Related and Future Work

- Problem originally addressed in [Hil1996] through a relation called weak isomorphism, from which we have taken the idea of preserving the average duration of internal action sequences.
- Congruence and steady-state exactness of weak isomorphism have been investigated, **but no axiomatization is known** (too strong, no branches).
- Different approach proposed in [Bra2002], where a variant of Markovian bisimilarity is defined that checks for exit rate equality with respect to all equivalence classes apart from the one including the process terms under examination.
- Congruence and axiomatization results have been provided, **but nothing is said about exactness**.
- **Axiomatization of $\simeq_{\text{MB},g}$, modal logic characterization, algorithms.**