Modèle probabiliste de systèmes distribués et concurrents Théorèmes limite

Samy Abbes

Projet IRISA DistribCom

IRISA / Université de Rennes 1

14 octobre 2004

Motivation and scientific context

Management of *networked* systems
(e.g., telecommunication networks)

Probabilistic model for:

- performance evaluation
- analysis of observations with *incomplete* information
  - distributed fault diagnosis
  - distributed statistical learning
    (inferring local parameters from local observations)

Requirements

Distributed architecture:

[Benveniste-Haar-Fabre-Jard 03]

occurrences of synchronous and asynchronous events

1. Nodes have *local clocks*
2. Asynchronous events cannot be *chronologically* compared:
   - **no global clock**
   - if a clock is imposed, different interleavings of asynchronous events must be identified
   - events in an execution form a *causal partial order*

Probabilistic extensions of concurrency models

Composition of probabilistic automata ([Segala 95,02])

a *2-steps* semantics

1. synchronisation w.r.t. labelling is *scheduled*
2. random decisions are *private* and do not interfere with synchronisation

- the probability is *not defined up to interleaving*: it does not match our requirements

Probabilistic extensions of concurrency models

Timed approaches (Continuous time Markov chains)

- **Stochastic Petri nets** based on race policies to solve conflicts
  - applications in queuing theory
  - connexions with $(\max, +)$ algebra (Baccelli, Mairesse,...)
- **Stochastic process algebras:** stochastic extensions of languages of timed processes
  - performance evaluation, bisimulation (Hermanns-Herzog-Katoen)

Probabilistic extensions of concurrency models

Timed approaches (Continuous time Markov chains)

- Transform the model, e.g. the Petri net, into a stochastic process $(X_t)_{t \geq 0}$ on a (huge) state space

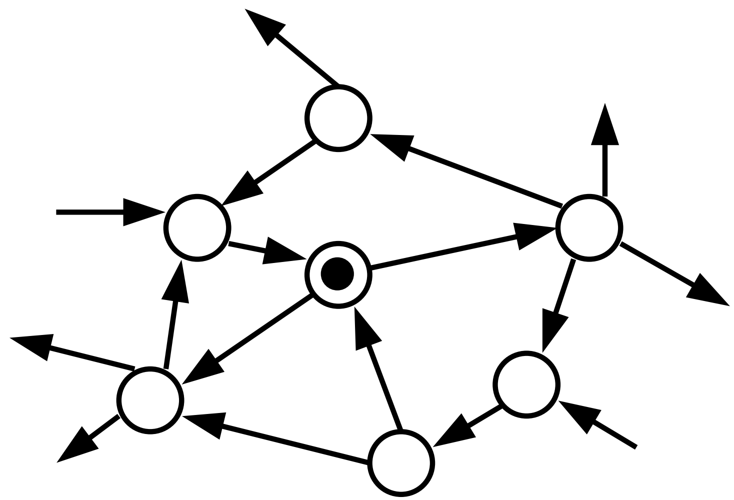

- partial orders are randomised *through the temporisation*: what happens without global clock?

Probabilistic models: a new approach

In the model of **safe Petri nets**

- No global clock $t \geq 0$
- randomize the set Ω of *maximal executions* of the system
- randomize Ω in a *recursive* way

  [Völzer 2001,
  Benveniste-Haar-Fabre 2003,
  Varacca-Winskel-Völzer 2004,
  Abbes 2004]

Summary

Objectives:

- provide a probabilistic framework for *safe Petri nets* without reference to a global clock
- obtain *asymptotic* results and statistical applications

Problematic:

is it true that a probabilistic Petri net is a Markov chain (DTMC) with *several* tokens?

Contributions

- **Continuous domain of configurations**
  - identification of the space of maximal configurations as a *projective limit*
  - *locally finite* unfoldings for extension of probabilities
- **Occurrence nets and event structures**
  - decomposition of true-concurrent processes through *branching cells* (local states)
  - computability
- **Probabilistic model**
  - construction of the *distributed probabilities*
  - stopping operators and the Strong *Markov property*
  - part of a *recurrence* theory
  - Law of large numbers

1. *Background: unfoldings and representations of space* Ω
2. Extension of probabilities
3. Decomposition of true-concurrent processes
4. Distributed product of probabilities
5. Markov nets: the Markov property and the Law of large numbers
6. Computability of local finiteness
7. Conclusion and perspectives

Representation of Ω (1)

Transition system (T.S.):

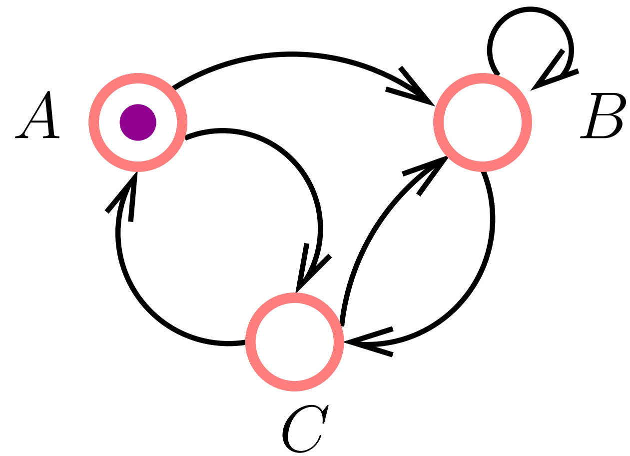

A **partial execution** is a *sequence* of moves of the token

= a *path* in the finite graph of the T.S.

A **maximal execution** $\omega \in \Omega$ is a sequence finite or infinite that cannot be continued

Representation of Ω (1)

covering tree = *transition system* (infinite)

- acyclic
- same dynamics than the T.S.: a path in the T.S. is *lifted* into a unique path in the covering tree

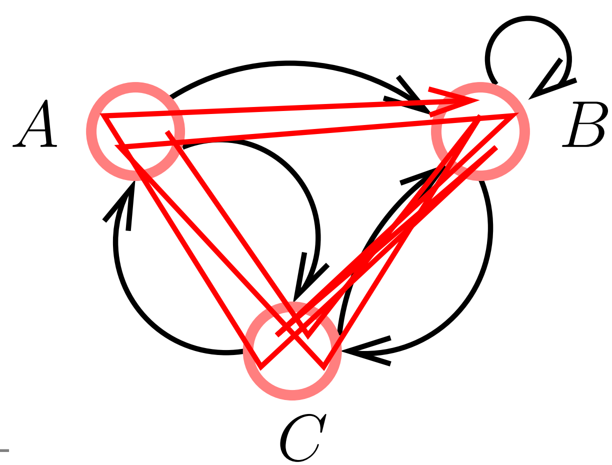

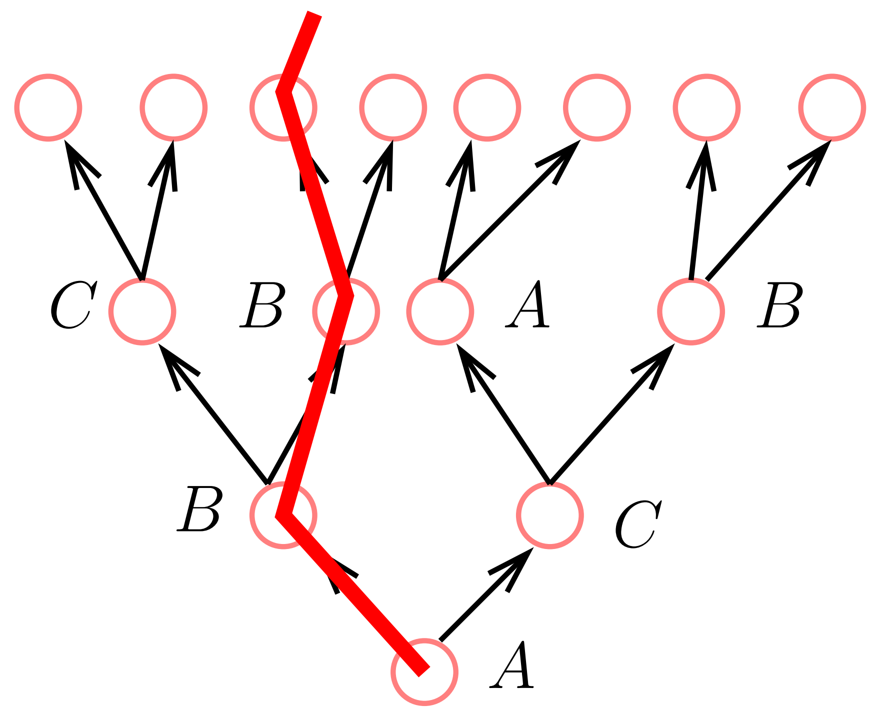

Representation of Ω (1)

covering tree = *transition system* (infinite)

- acyclic
- same dynamics than the T.S.: a path in the T.S. is *lifted* into a unique path in the covering tree

Ω is the **boundary at infinity** of the covering tree

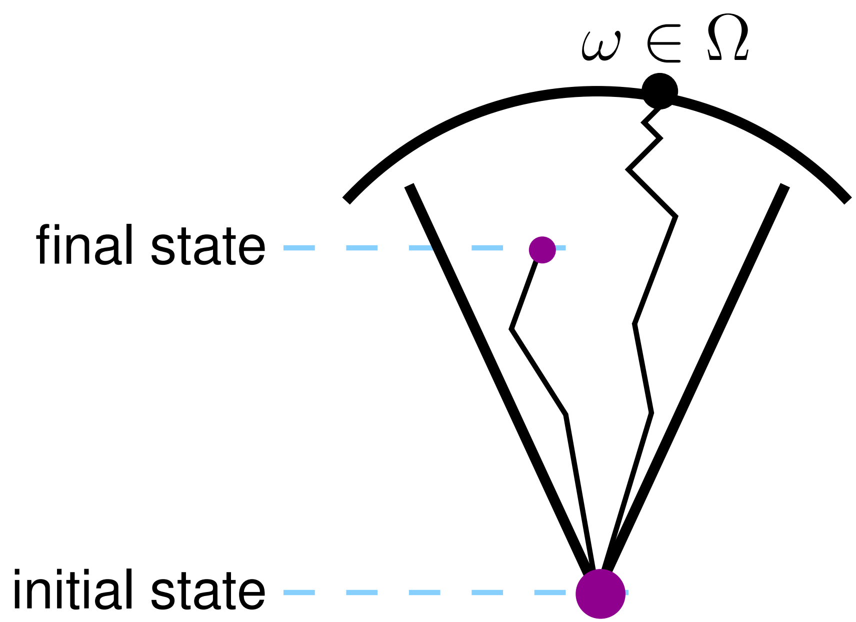

Representation of Ω (2)

Unfolding of a safe Petri net (Winskel 80)

A *firing sequence* of the Petri net is lifted into a *firing sequence* of the unfolding.

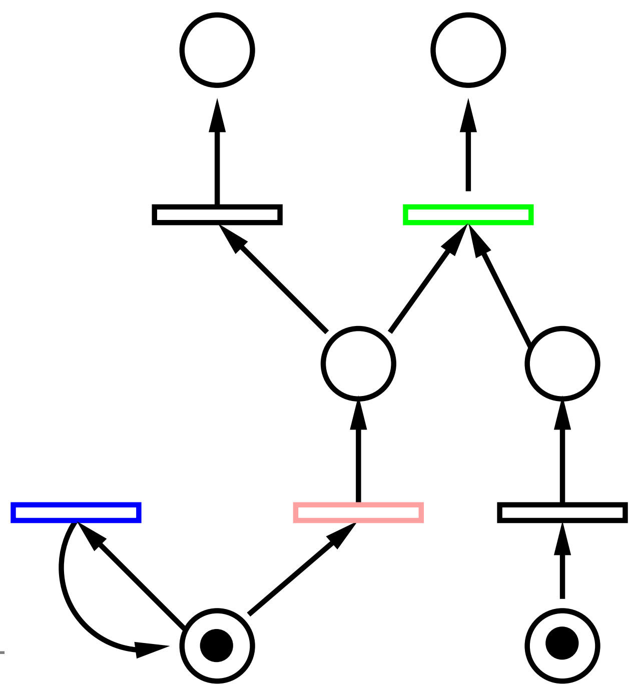

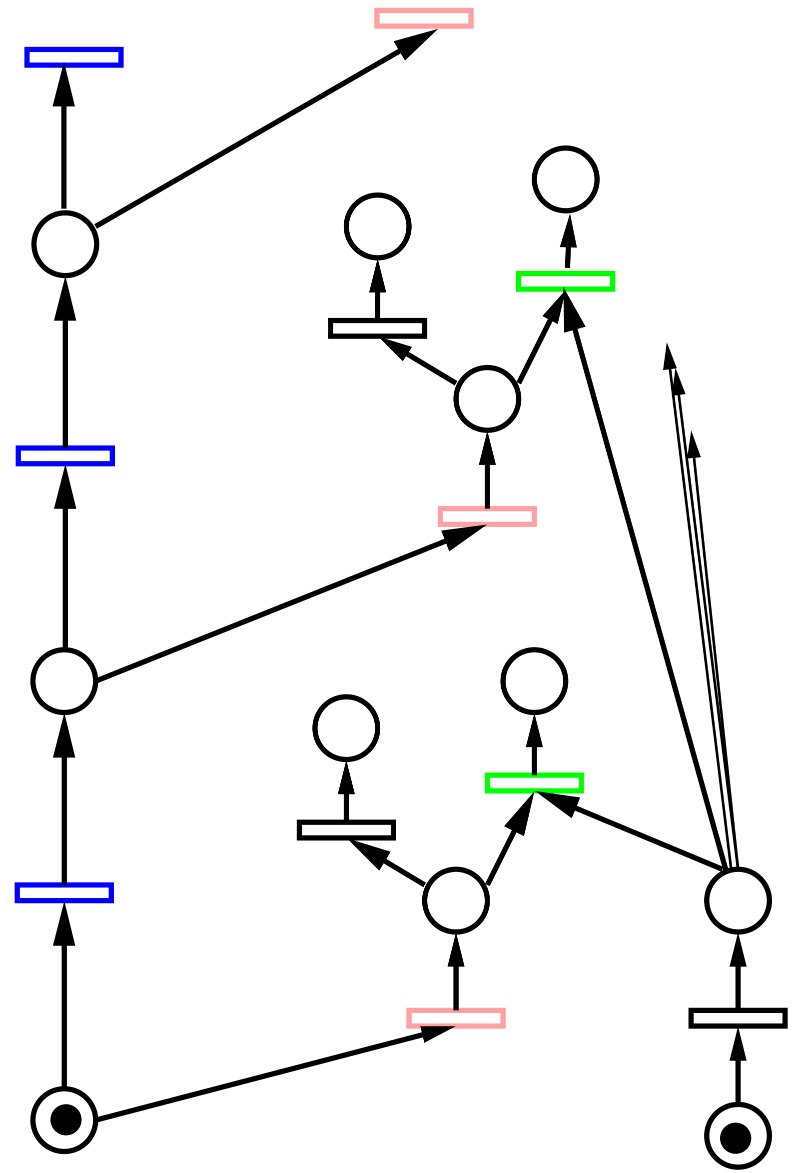

Representation of Ω (2)

Unfolding of a safe Petri net (Winskel 80)

A *firing sequence* of the Petri net is lifted into a *firing sequence* of the unfolding.

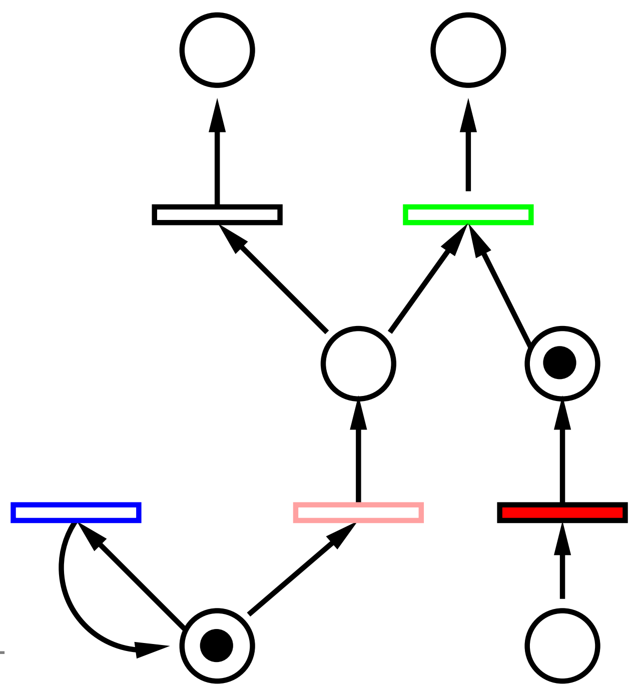

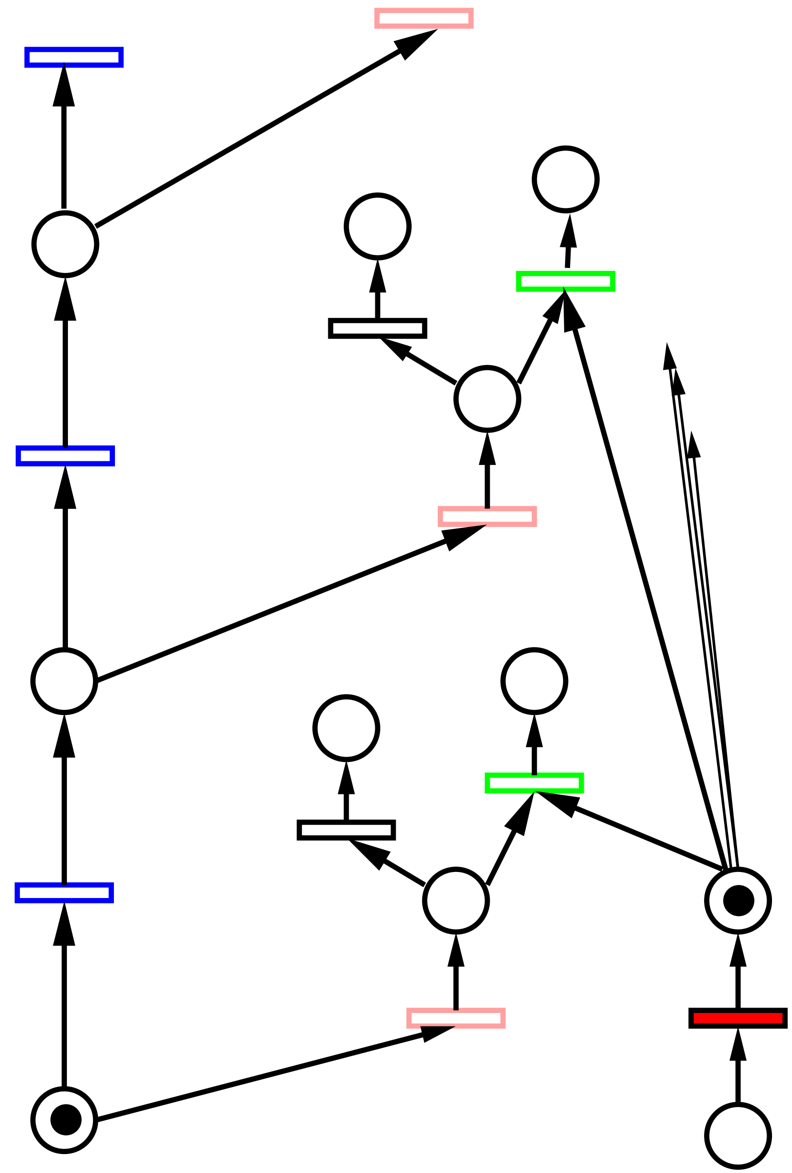

Representation of Ω (2)

Unfolding of a safe Petri net (Winskel 80)

A *firing sequence* of the Petri net is lifted into a *firing sequence* of the unfolding.

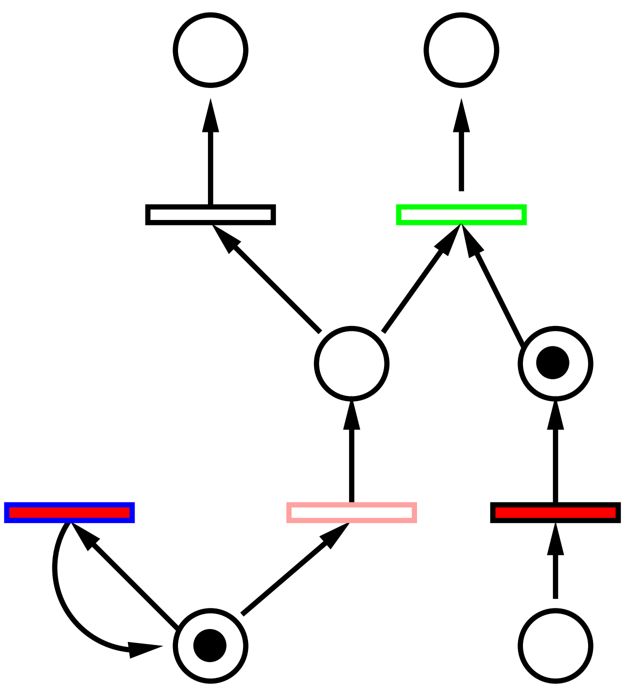

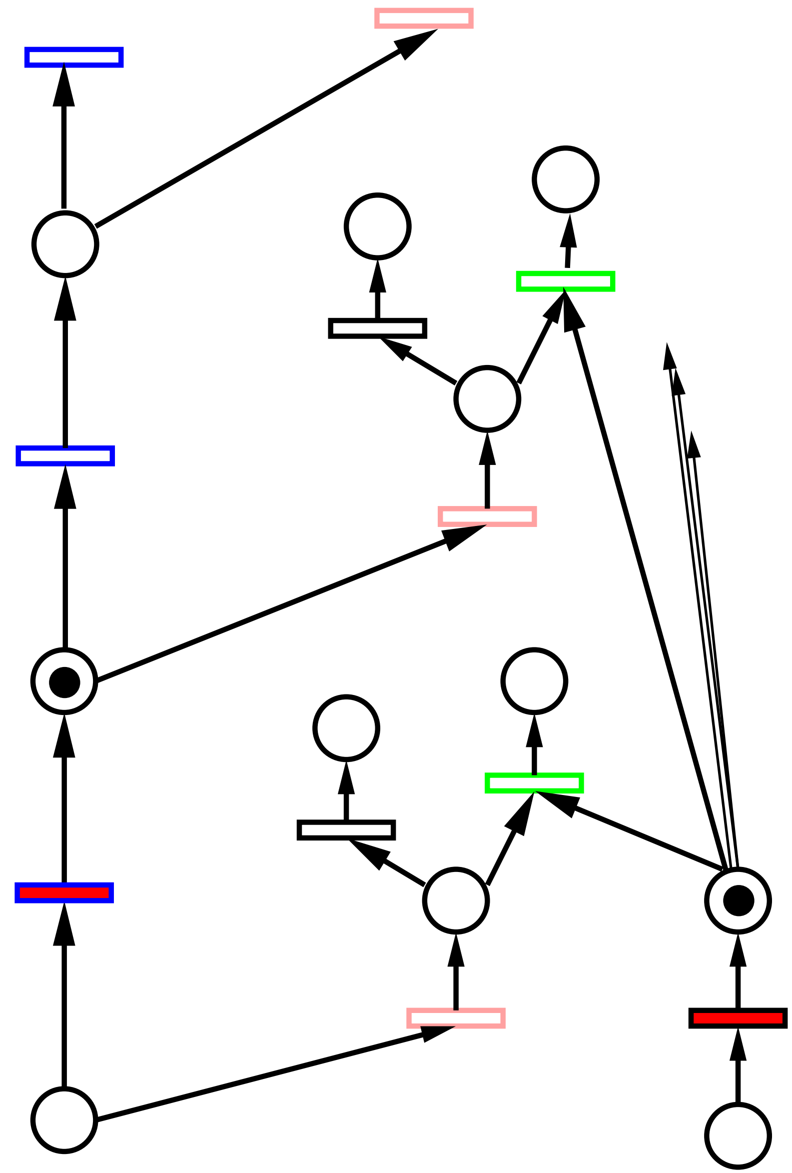

Representation of Ω (2)

Unfolding of a safe Petri net (Winskel 80)

A *firing sequence* of the Petri net is lifted into a *firing sequence* of the unfolding.

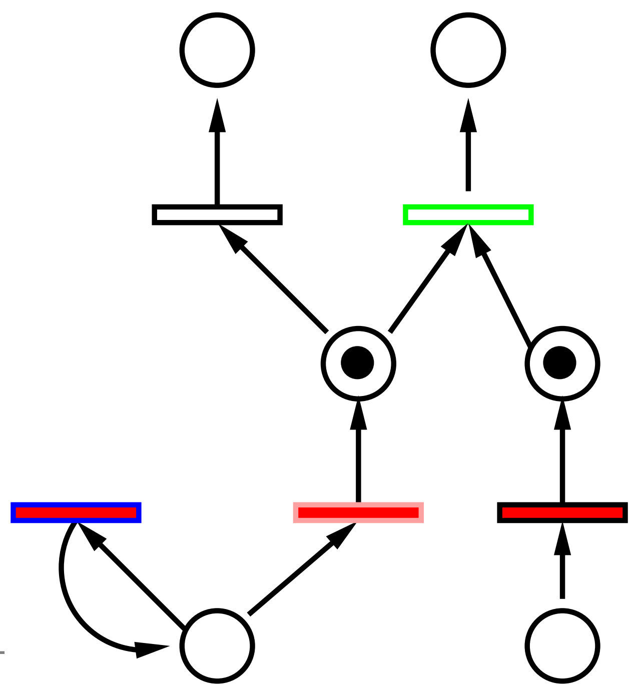

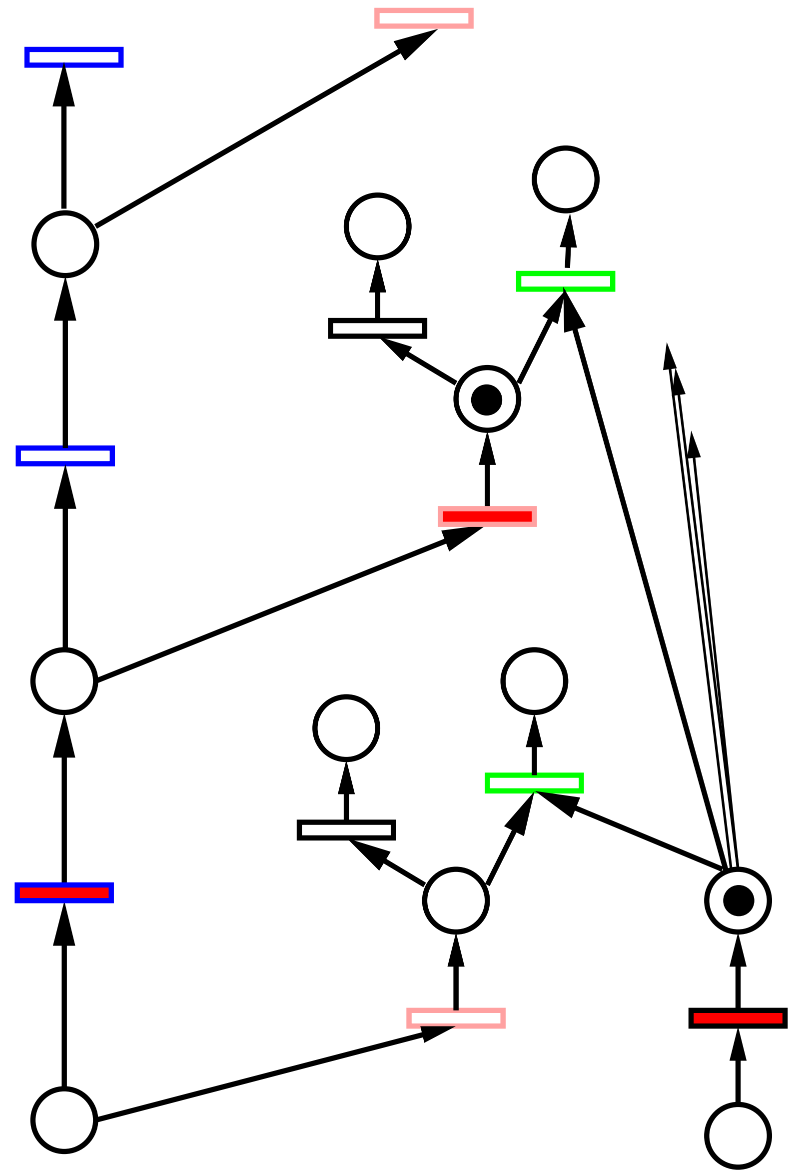

Representation of Ω (2)

Unfolding of a safe Petri net (Winskel 80)

A *firing sequence* of the Petri net is lifted into a *firing sequence* of the unfolding.

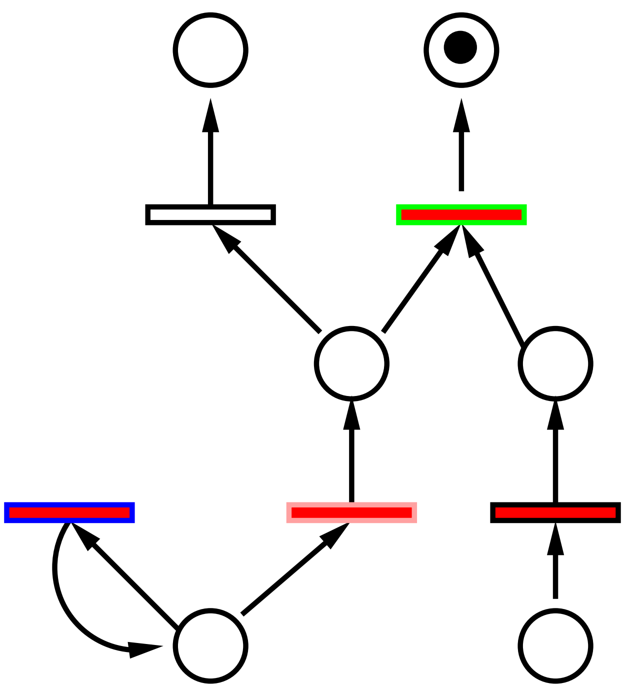

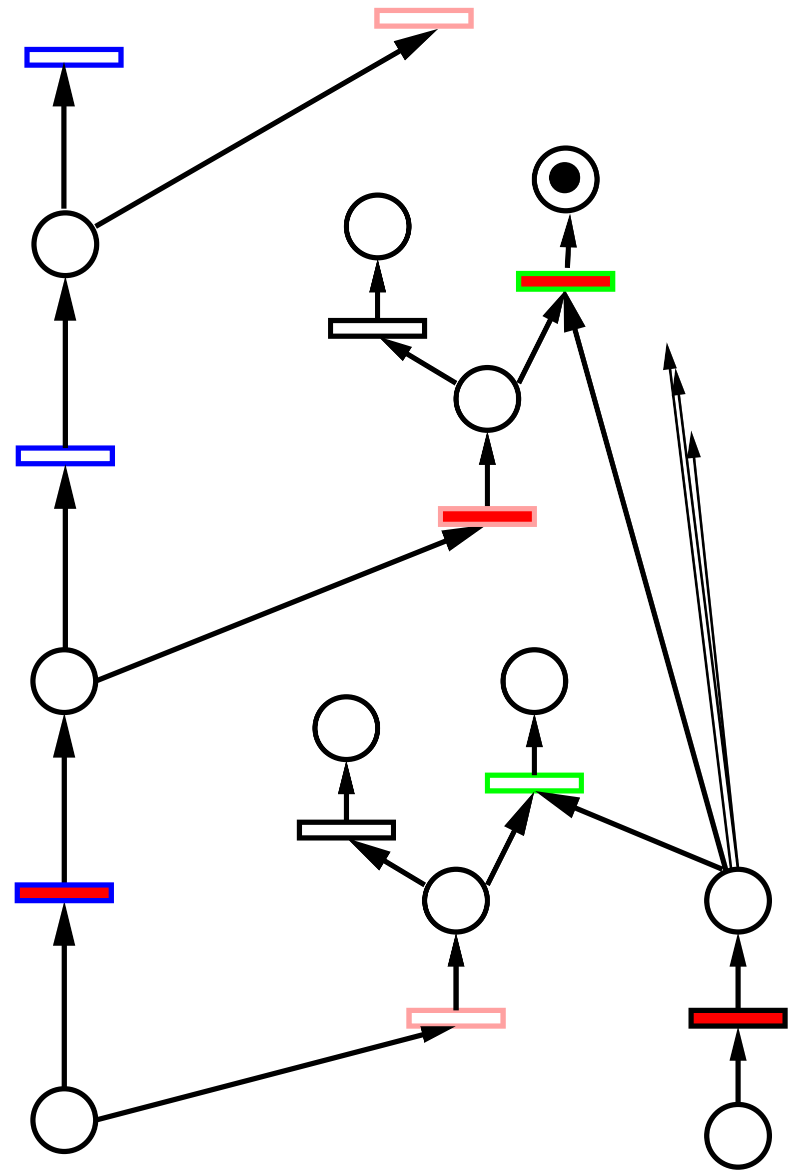

Representation of Ω (2)

Unfolding of a safe Petri net (Winskel 80)

A *trace* of the Petri net is lifted into a *trace* of the unfolding.

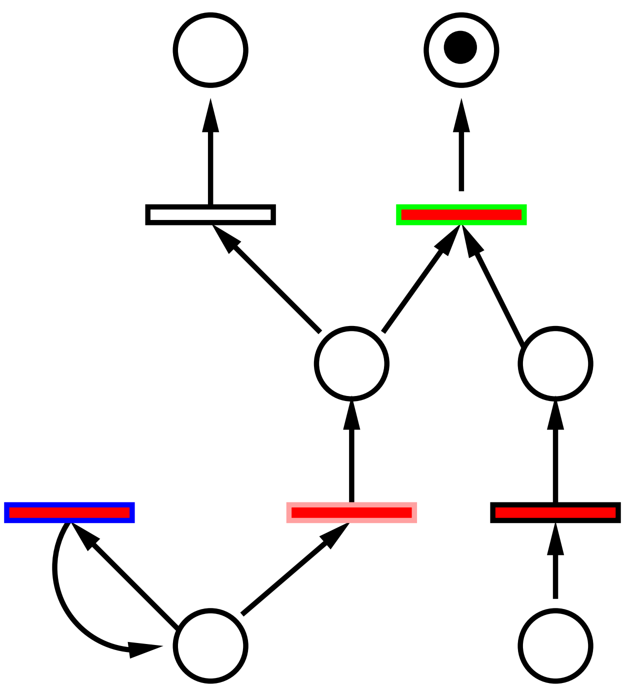

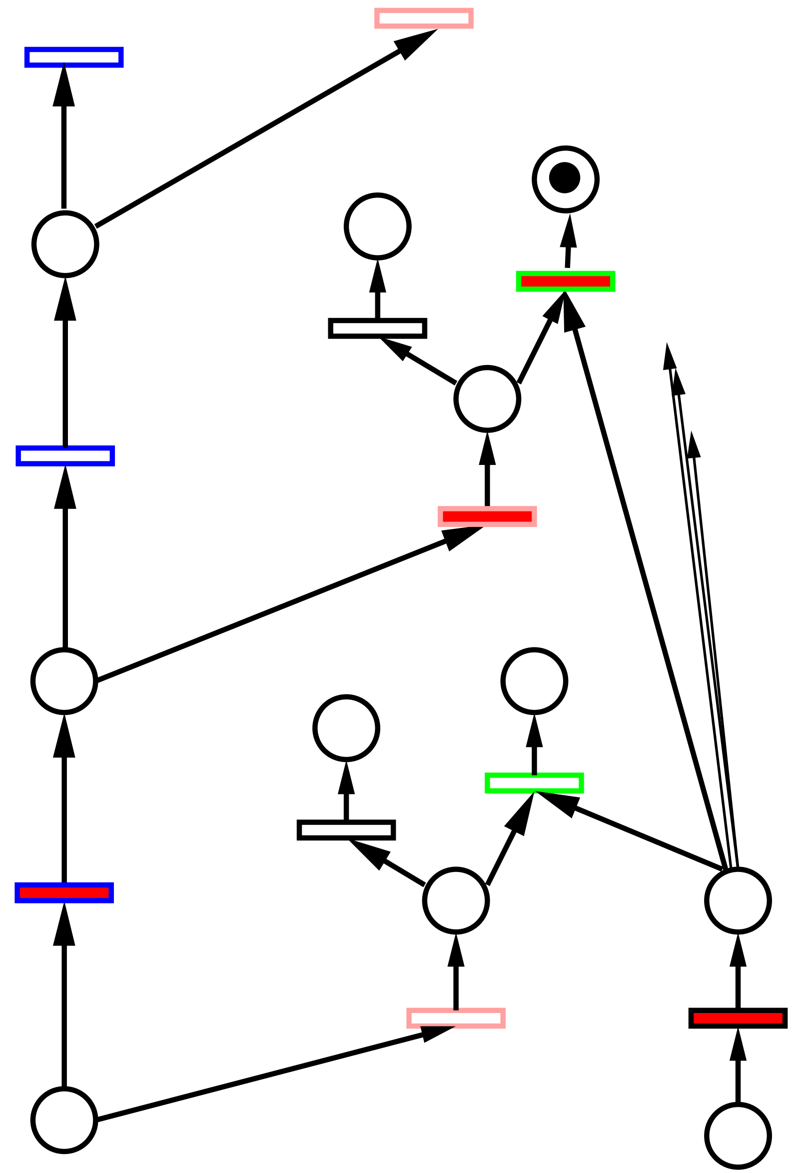

Representation of Ω (2)

- The **unfolding** $\mathcal{U}$ of a safe Petri net $\mathcal{N}$ is a *labelled occurrence net*, labelled by $\mathcal{N}$, and with the same true-concurrent dynamics than $\mathcal{N}$.
- a finite *configuration* of $\mathcal{U}$ leads to a marking of net $\mathcal{N}$

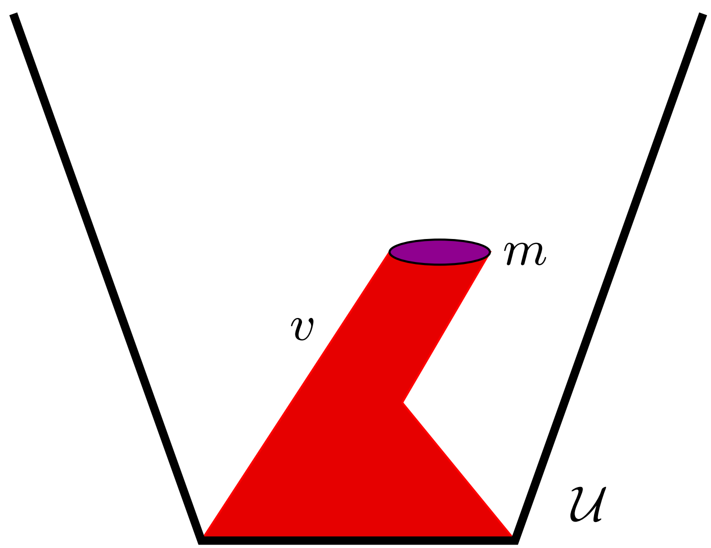

Representation of Ω (2)

- The **unfolding** $\mathcal{U}$ of a safe Petri net $\mathcal{N}$ is a *labelled occurrence net*, labelled by $\mathcal{N}$, and with the same true-concurrent dynamics than $\mathcal{N}$.
- $\Omega = \{$**maximal configurations**$\}$ = boundary at infinity of $\mathcal{U}$

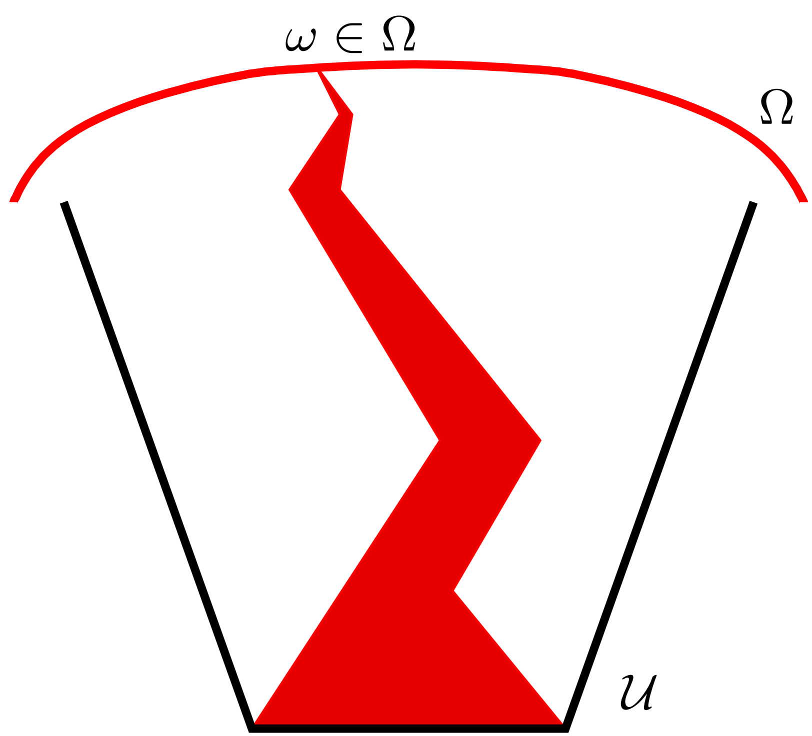

1. Background: unfoldings and representations of space Ω
2. *Extension of probabilities*
3. Decomposition of true-concurrent processes
4. Distributed product of probabilities
5. Markov nets: the Markov property and the Law of large numbers
6. Computability of local finiteness
7. Conclusion and perspectives

Finite approximations of Ω

Case of a transition system.
Ω_n = finite set of executions after n moves

? Existence of $\mathbb{P}$: $\mathbb{P}(X_0, \dots, X_n) = \mathbb{P}_n(X_0, \dots, X_n) \quad \forall n \geq 0$**?**

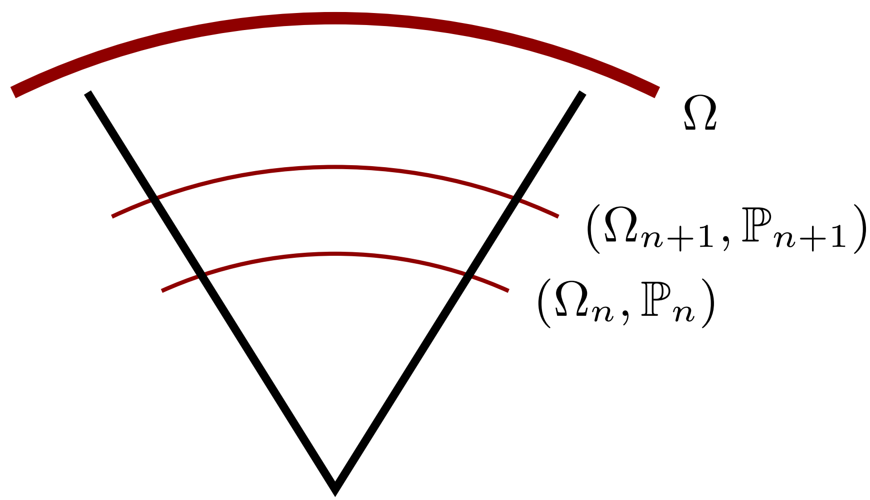

Kolmogorov-Prokhorov *extension* theorem: the extension occurs iff:

$$\mathbb{P}_n\big(X_0, \dots, X_n\big) = \sum_{s \in S} \mathbb{P}_{n+1}\big(X_0, \dots, X_n, s\big)$$

Finite approximations of Ω

Case of a transition system.
Ω_n = finite set of executions after n moves

?**Existence of $\mathbb{P}$:** $\mathbb{P}(X_0, \ldots, X_n) = \mathbb{P}_n(X_0, \ldots, X_n) \quad \forall n \geq 0$?

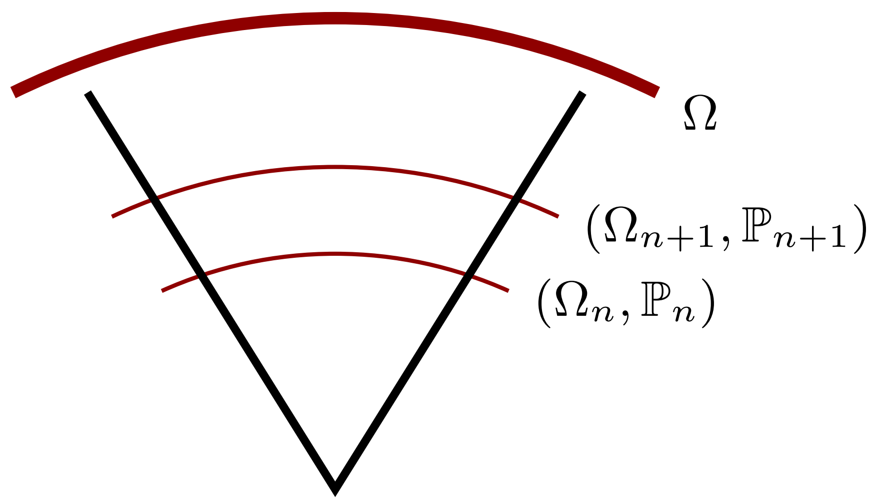

Kolmogorov-Prokhorov *extension* theorem relies on:

$\underbrace{\Omega = \varprojlim_n \Omega_n}_{\text{tree model}}$ + $\underbrace{\text{extension theorem for projective limits of probabilities (Prokhorov, 1930's)}}_{\text{probability theory}}$

Finite approximations of Ω

- For *concurrent* models ($\mathcal{N}$, $\mathcal{U}$, Ω), is there a projective systems of *finite* sets $(\Gamma_n)_n$ such that:

$$\Omega = \varprojlim_n \Gamma_n \; ? \qquad (n \text{ ranges over a countable lattice})$$

Finite approximations of Ω

- For *concurrent* models $(\mathcal{N}, \mathcal{U}, \Omega)$, is there a projective systems of *finite* sets $(\Gamma_n)_n$ such that:

$$\Omega = \varprojlim_n \Gamma_n \, ? \quad (n \text{ ranges over a countable lattice})$$

- **Theorem:** yes, if and only if Ω is *compact* in the Scott topology, and in this case we can take:

$$P \text{ finite prefix of } \mathcal{U}, \quad \Gamma_P = \{\omega \cap P\,,\ \omega \in \Omega\}$$

Ω *is the limit of its traces over finite prefixes of* $\mathcal{U}$

Locally finite unfoldings

A general **compact** case for Ω

- *minimal* conflict on $\mathcal{U}$: $e \#_\mu e'$
- a **stopping prefix** of $\mathcal{U}$ is a prefix $\#_\mu$-closed
- **Property:** for stopping prefix B:

$$\Omega_B = \{\omega \cap B,\ \omega \in \Omega\}$$

the traces of Ω *over stopping prefix* B *coincide with* Ω_B
$\rightarrow$ interesting property from a *computational* point of view

Locally finite unfoldings

A general **compact** case for Ω

- an unfolding $\mathcal{U}$ is *locally finite* if for every node x of $\mathcal{U}$, there is a stopping prefix B s.t.:

$$x \in B, \quad B \text{ is finite}$$

- **Theorem:** if $\mathcal{U}$ is locally finite, then Ω is compact, and:

$$\Omega = \varprojlim_B \Omega_B$$

$\rightarrow$ *restrict the study to nets with locally finite unfoldings*

Examples

- **Locally finite**: unfoldings of
  - transition systems
  - confusion-free and free-choice nets
  - nets with finite unfoldings
  - some other nets
- **Non locally finite**:

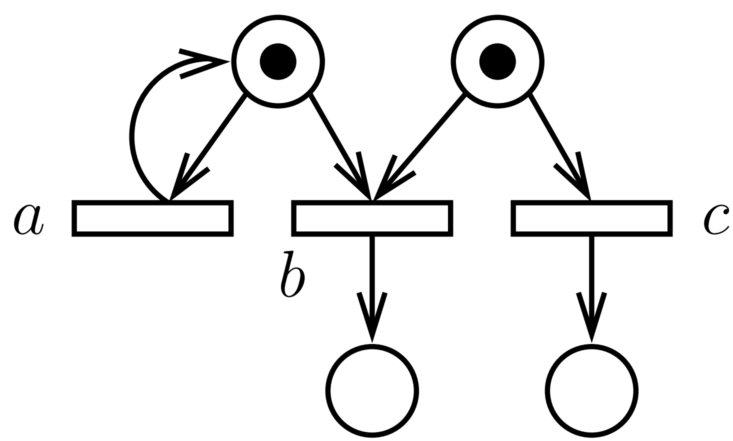

transition c is in competition with *infinitely many* occurrences of transition a

1. Background: unfoldings and representations of space Ω
2. Extension of probabilities
3. *Decomposition of true-concurrent processes*
4. Distributed product of probabilities
5. Markov nets: the Markov property and the Law of large numbers
6. Computability of local finiteness
7. Conclusion and perspectives

Future

$\mathcal{N}$ a finite safe Petri net with unfolding $\mathcal{U}$ from marking M_0

For v a configuration of $\mathcal{U}$

$$\mathcal{U}^v = \textbf{future of } v = \text{unfolding of } \mathcal{N} \text{ from } m(v)$$

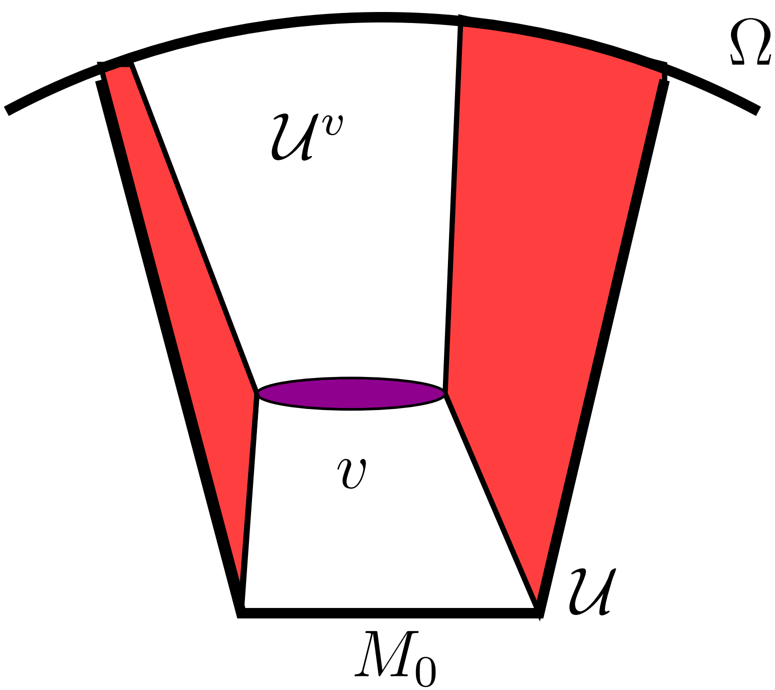

Nodes not in v and
not in conflict with v

Future

$\mathcal{N}$ a finite safe Petri net with unfolding $\mathcal{U}$ from marking M_0

For v a configuration of $\mathcal{U}$

$$\mathcal{U}^v = \textbf{future} \text{ of } v = \text{unfolding of } \mathcal{N} \text{ from } m(v)$$

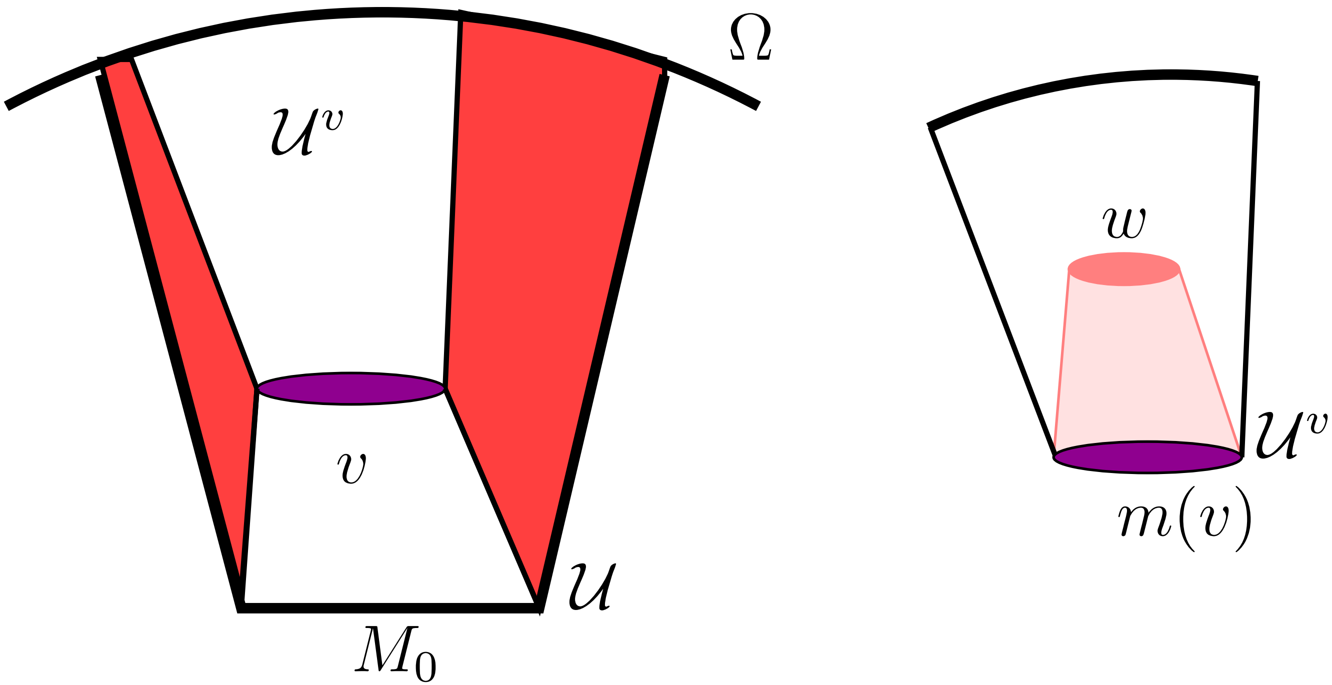

Future

$\mathcal{N}$ a finite safe Petri net with unfolding $\mathcal{U}$ from marking M_0

For v a configuration of $\mathcal{U}$

$$\mathcal{U}^v = \textbf{future} \text{ of } v = \text{unfolding of } \mathcal{N} \text{ from } m(v)$$

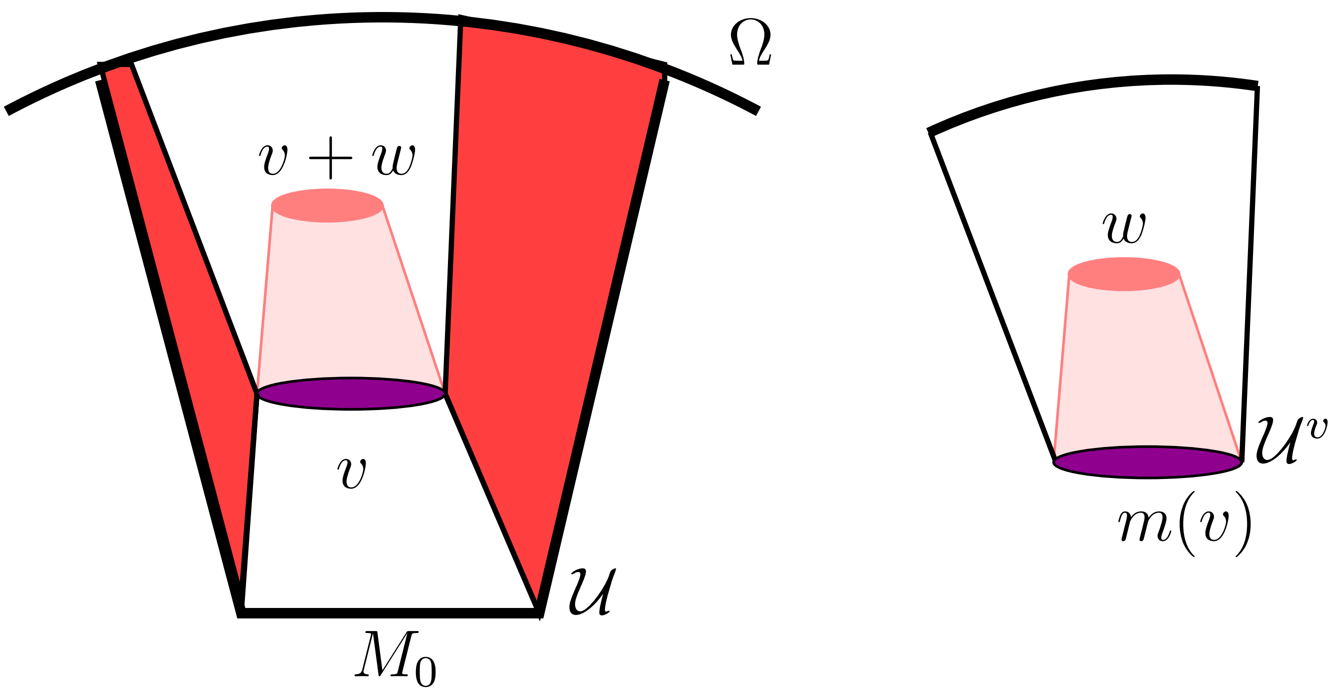

Future

$\mathcal{N}$ a finite safe Petri net with unfolding $\mathcal{U}$ from marking M_0

For v a configuration of $\mathcal{U}$

$$\mathcal{U}^v = \textbf{future of } v = \text{unfolding of } \mathcal{N} \text{ from } m(v)$$

- **associative** composition $v + w$
  defined for w a configuration of $\mathcal{U}^v$
  $\rightarrow$ suggests recursive decompositions of maximal processes

Decompositions of configurations

Safe Petri net $\mathcal{N}$, unfolding $\mathcal{U}$. *Fix* a maximal configuration ω.

Definition: An **initial branching cell** of $\mathcal{U}$ is a minimal $\neq \emptyset$ stopping prefix of $\mathcal{U}$.

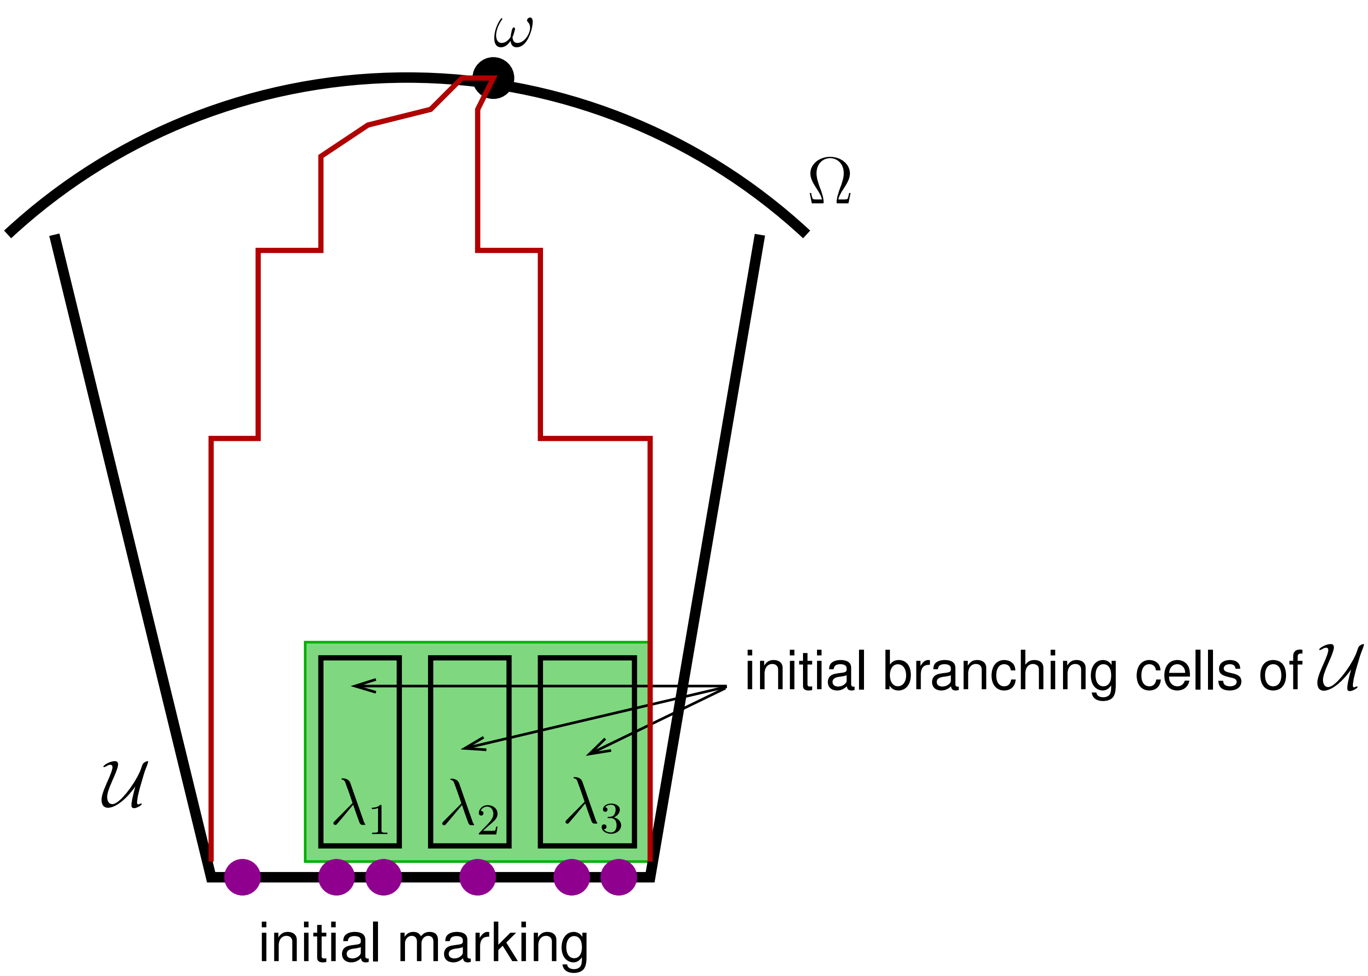

Decompositions of configurations

Safe Petri net $\mathcal{N}$, unfolding $\mathcal{U}$. *Fix* a maximal configuration ω.

Definition: An **initial branching cell** of $\mathcal{U}$ is a minimal $\neq \emptyset$ stopping prefix of $\mathcal{U}$.

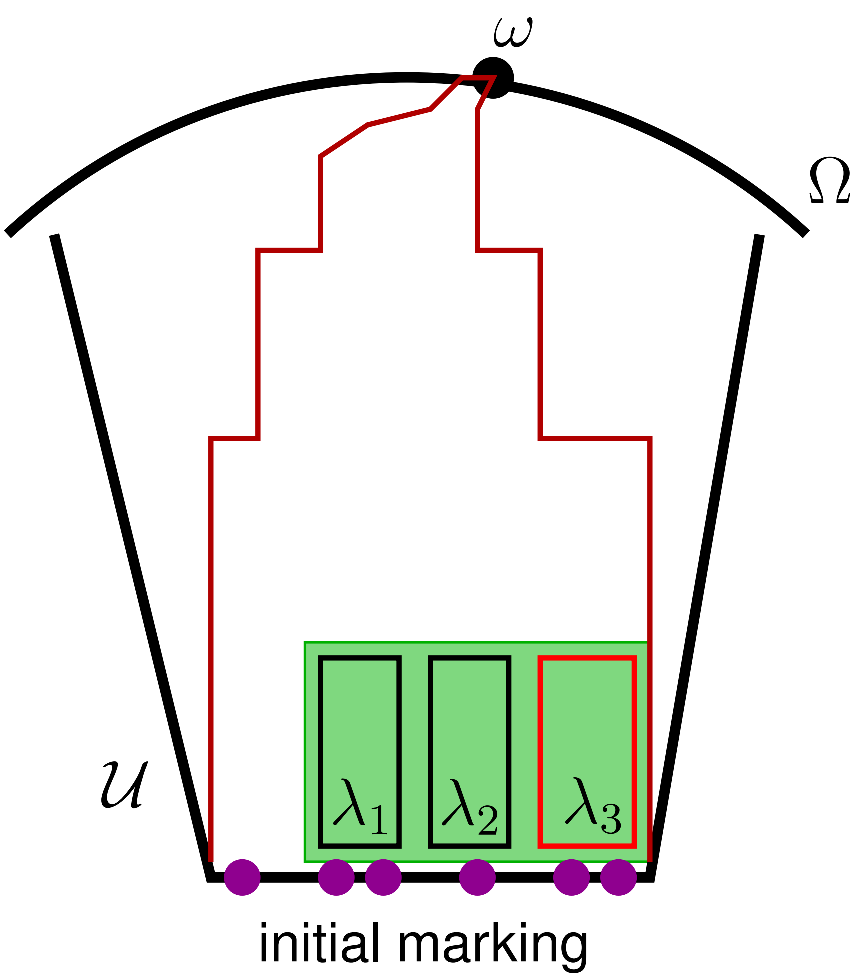

Step 1: select a initial branching cell, say λ_3

Decompositions of configurations

Safe Petri net $\mathcal{N}$, unfolding $\mathcal{U}$. *Fix* a maximal configuration ω.

Definition: An **initial branching cell** of $\mathcal{U}$ is a minimal $\neq \emptyset$ stopping prefix of $\mathcal{U}$.

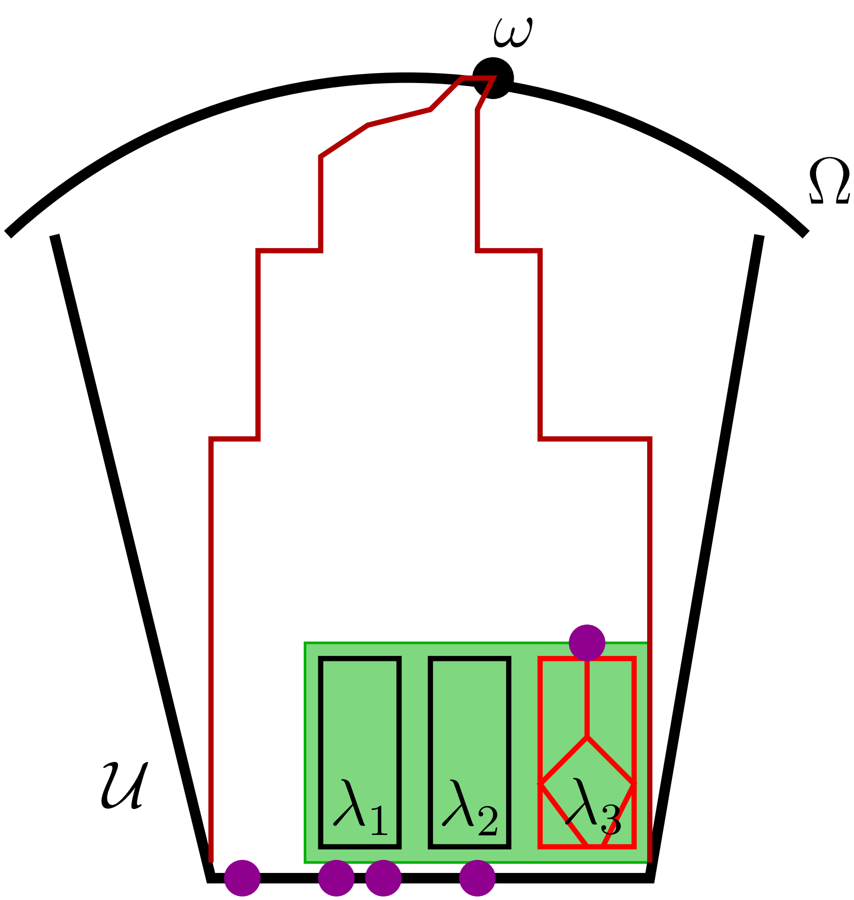

Step 1: select a initial branching cell, say λ_3
Set $v_1 = \omega \cap \lambda_3$.
v_1 is *maximal* in λ_3

Decompositions of configurations

Safe Petri net $\mathcal{N}$, unfolding $\mathcal{U}$. *Fix* a maximal configuration ω.

Definition: An **initial branching cell** of $\mathcal{U}$ is a minimal $\neq \emptyset$ stopping prefix of $\mathcal{U}$.

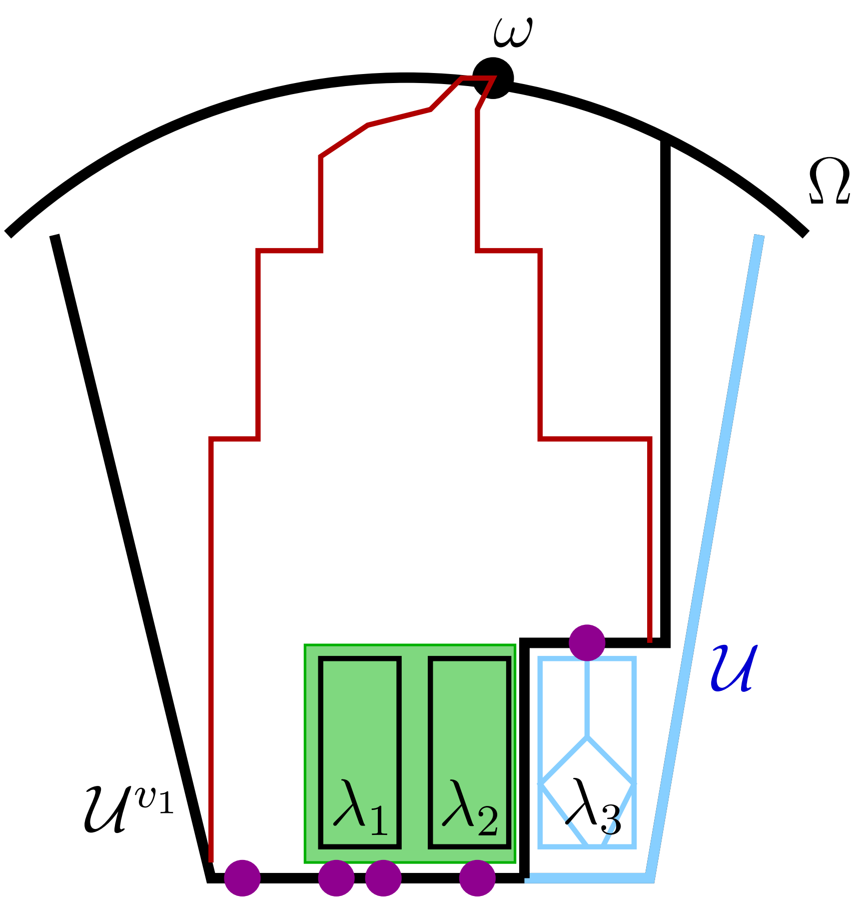

Step 1: select a initial branching cell, say λ_3
Set $v_1 = \omega \cap \lambda_3$.
v_1 is *maximal* in λ_3

Step 2: consider
$\mathcal{U}^{v_1}$ = *future* of v_1

Decompositions of configurations

Safe Petri net $\mathcal{N}$, unfolding $\mathcal{U}$. *Fix* a maximal configuration ω.

Definition: An **initial branching cell** of $\mathcal{U}$ is a minimal $\neq \emptyset$ stopping prefix of $\mathcal{U}$.

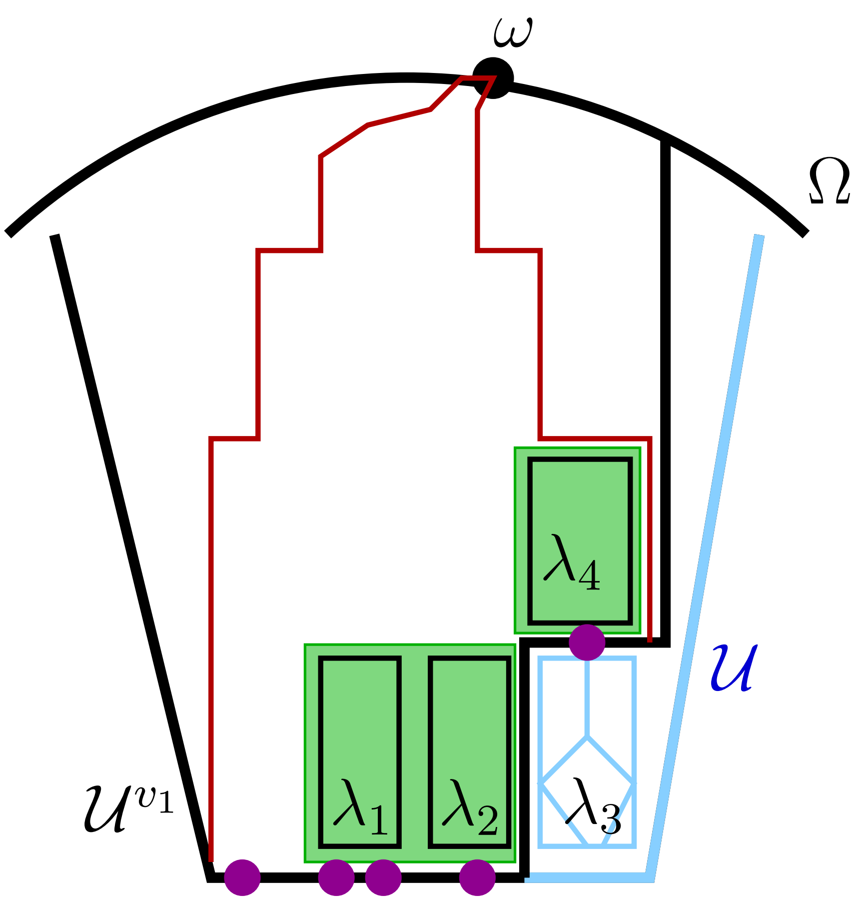

Step 1: select a initial branching cell, say λ_3
Set $v_1 = \omega \cap \lambda_3$.
v_1 is *maximal* in λ_3

Step 2: consider
$\mathcal{U}^{v_1}$ = *future* of v_1
Initial branching cells of $\mathcal{U}^{v_1}$ are: $\lambda_1, \lambda_2, \lambda_4$

Decompositions of configurations

Safe Petri net $\mathcal{N}$, unfolding $\mathcal{U}$. *Fix* a maximal configuration ω.

Definition: An **initial branching cell** of $\mathcal{U}$ is a minimal $\neq \emptyset$ stopping prefix of $\mathcal{U}$.

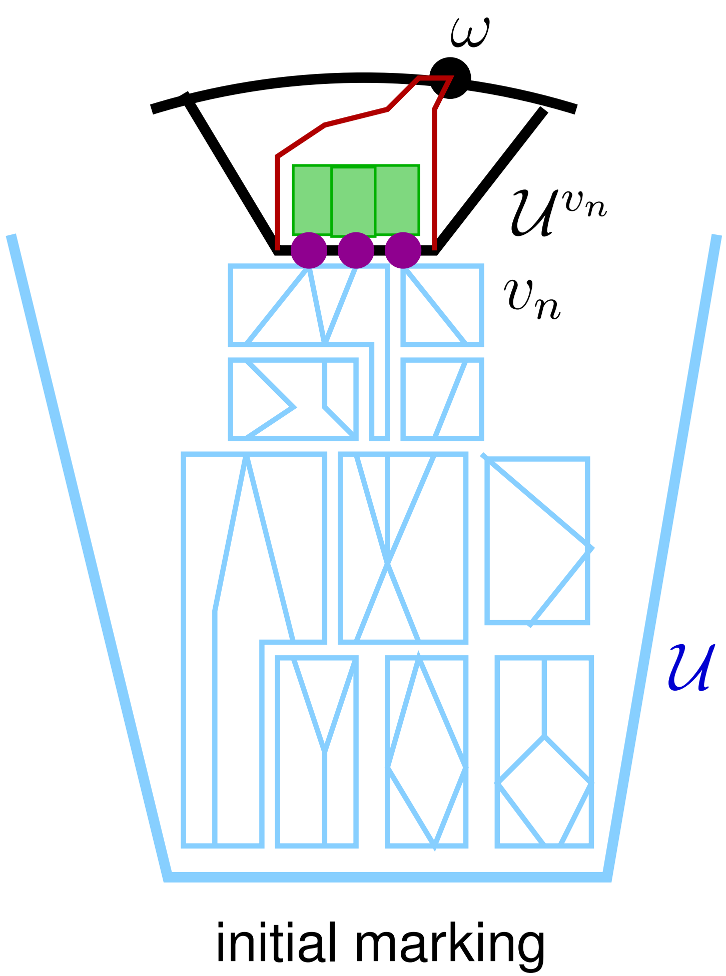

Step 1: select an initial branching cell, say λ_3
Set $v_1 = \omega \cap \lambda_3$.
v_1 is *maximal* in λ_3

Step 2: consider
$\mathcal{U}^{v_1}$ = *future* of v_1
Initial branching cells of $\mathcal{U}^{v_1}$ are: $\lambda_1, \lambda_2, \lambda_4$

Repeat to construct
$v_1, v_2, \ldots$

Branching cells

- **Finite well stopped (w.s)** configurations are all the v_n that can be constructed, for ω ranging over Ω
- **Branching cells** are initial branching cells of $\mathcal{U}^v$, for v finite **w.s**

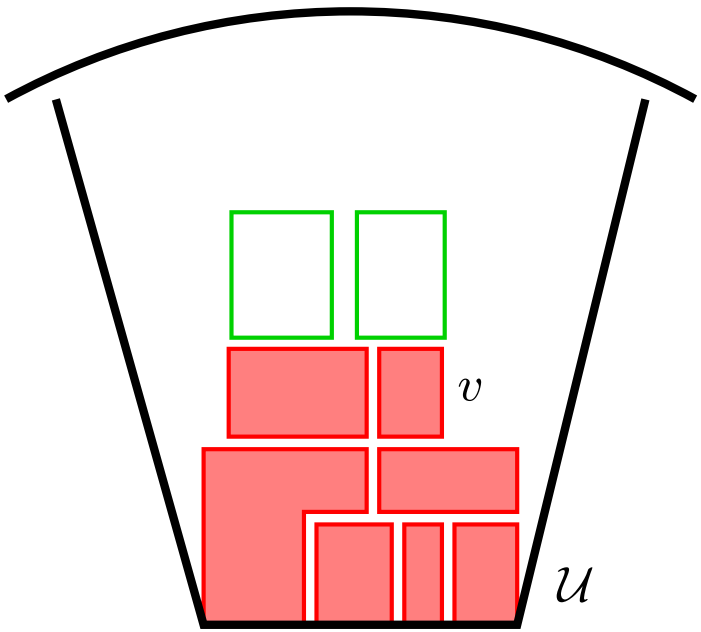

Theorem: a finite **w.s** configuration v admits a *unique* decomposition:

$$v = \bigcup_{\text{finite}} \xi_\lambda, \quad \xi_\lambda \in \Omega_\lambda$$

and branching cells λ are *disjoint*

Branching cells

Properties

- if B is a finite stopping prefix, then every maximal configuration $\omega_B \in \Omega_B$ is well stopped
- **Stability** under concatenation
  if v is well stopped in $\mathcal{U}$, if w is well stopped in $\mathcal{U}^v$, then $v + w$ is well stopped in $\mathcal{U}$

Remark: Well stopped configurations form the smallest class with both properties.

Branching cells

Comments:

- branching cells are *dynamic*, because of *concurrency*
  An event t can belong to different branching cells, according to the *context* of t

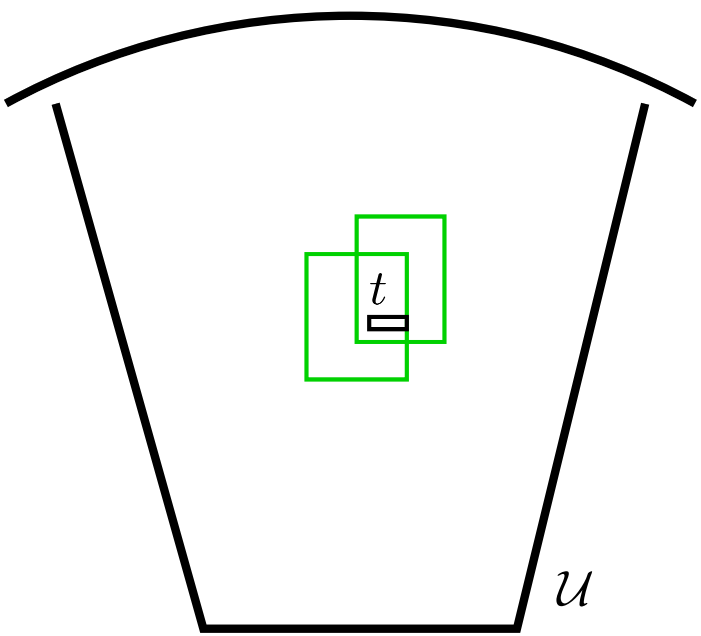

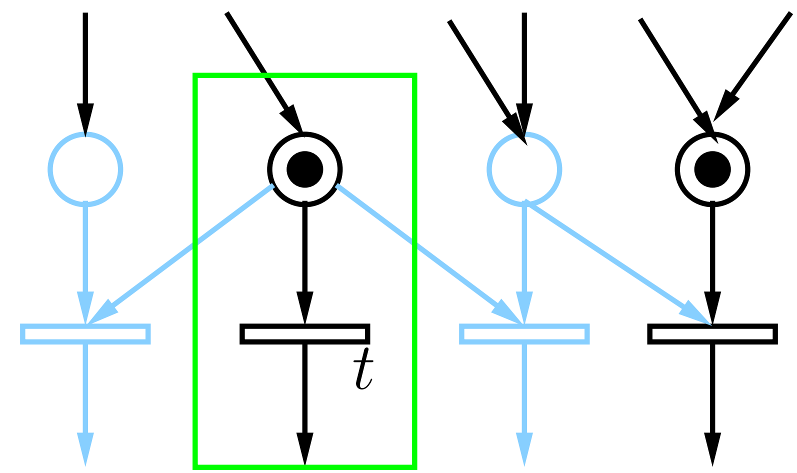

Branching cells

Comments:

- branching cells are *dynamic*, because of *concurrency*
  An event t can belong to different branching cells, according to the *context* of t

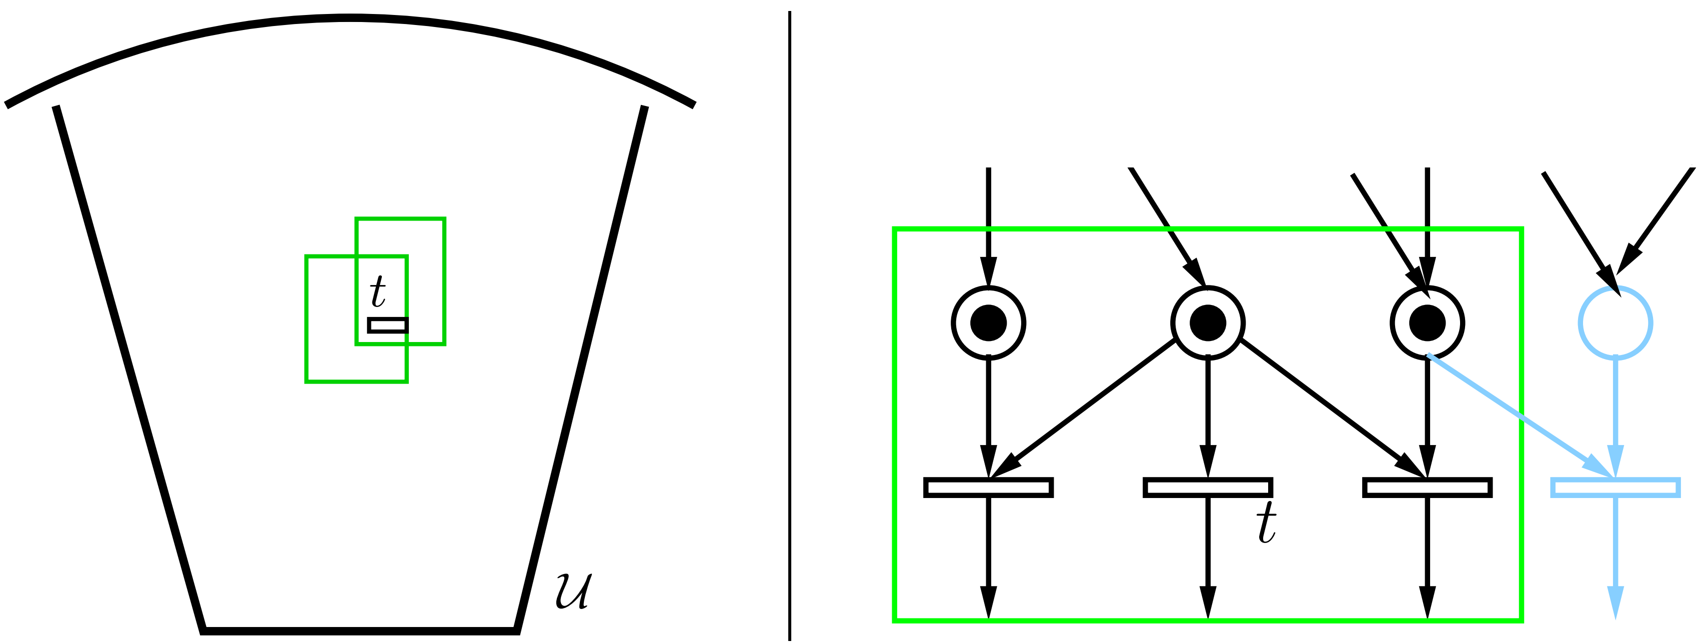

Branching cells

Comments:

- branching cells are *dynamic*, because of *concurrency*
- **Case of a tree:** branching cells do not overlap

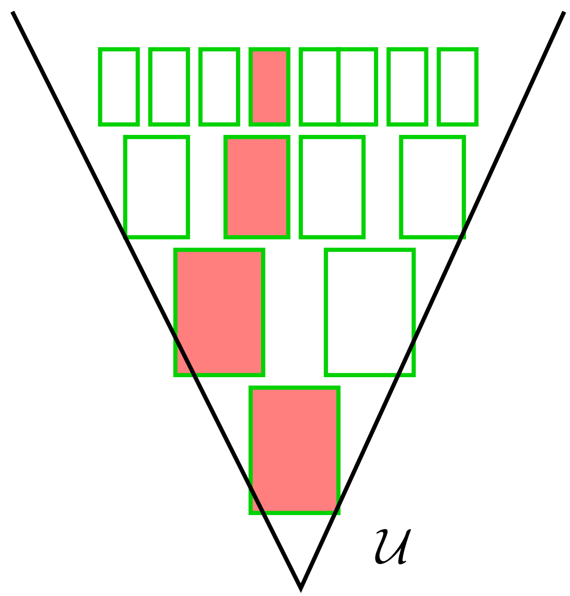

1. Background: unfoldings and representations of space Ω
2. Extension of probabilities
3. Decomposition of true-concurrent processes
4. *Distributed product of probabilities*
5. Markov nets: the Markov property and the Law of large numbers
6. Computability of local finiteness
7. Conclusion and perspectives

Construction of a probability

Data: a *countable* family of *finite* probabilities $(q_\lambda)_\lambda$, λ ranging over the branching cells of $\mathcal{U}$

Define, for configuration $v = \bigcup_\lambda \xi_\lambda$ finite **w.s**

$$p(v) = \prod_\lambda q_\lambda(\xi_\lambda)$$

Construction of a probability

Data: a *countable* family of *finite* probabilities $(q_\lambda)_\lambda$, λ ranging over the branching cells of $\mathcal{U}$

Define, for configuration $v = \bigcup_\lambda \xi_\lambda$ finite **w.s**

$$p(v) = \prod_\lambda q_\lambda(\xi_\lambda)$$

- for every finite stopping prefix B, $\mathbb{P}_B(\omega_B) =_{\text{def}} p(\omega_B)$ is a probability on Ω_B
- $(\mathbb{P}_B)_B$ is a projective system of probabilities on $(\Omega_B)_B$

$$B \subseteq B', \quad \forall \omega_B \in \Omega_B, \quad \mathbb{P}_B(\omega_B) = \sum_{\omega_{B'} \in \Omega_{B'}} \mathbb{P}_{B'}(\omega_{B'})$$

Construction of a probability

Data: a *countable* family of *finite* probabilities $(q_\lambda)_\lambda$, λ ranging over the branching cells of $\mathcal{U}$

Define, for configuration $v = \bigcup_\lambda \xi_\lambda$ finite **w.s**

$$p(v) = \prod_\lambda q_\lambda(\xi_\lambda)$$

Extension theorem: there is a unique probability $\mathbb{P}$ on Ω s.t. for all B:

$$\mathbb{P}(\omega \supseteq \omega_B) = p(\omega_B)$$

and then for every v finite **w.s**: $\mathbb{P}(\omega \supseteq v) = p(v)$

Construction of a probability

$\Omega(v) = \{\omega \in \Omega : \ \omega \supseteq v\}$ is the *shadow* of v

$p(v) = \mathbb{P}\big(\Omega(v)\big)$

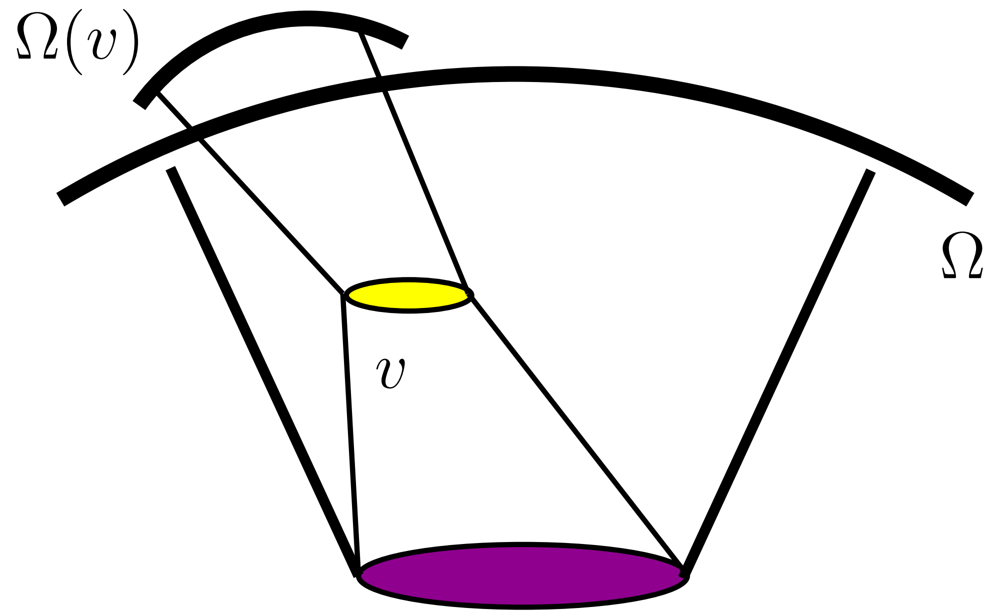

Call $\mathbb{P}$ the **distributed product** of family $(q_\lambda)_{\lambda \subseteq \mathcal{U}}$

$$\mathbb{P} = \bigotimes_{\lambda \subseteq \mathcal{U}}^{\mathsf{d}} q_\lambda$$

Probabilities induced in past and future

$\Omega(v) = \{\omega \in \Omega : \ \omega \supseteq v\}$ is the *shadow* of v

$p(v) = \mathbb{P}\big(\Omega(v)\big)$

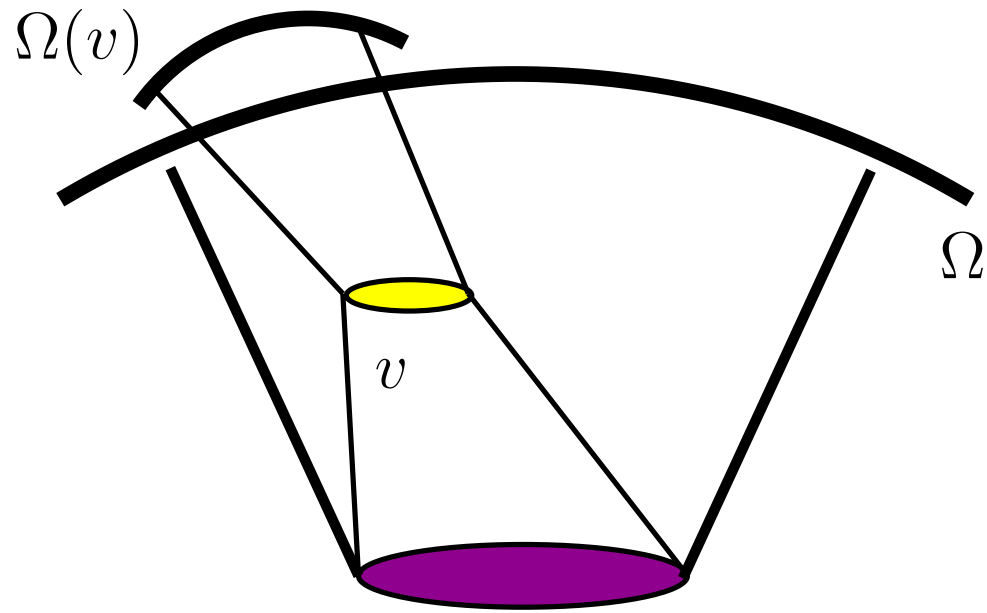

- **Probabilistic future** of configuration v: probability $\mathbb{P}^v$ on the *shadow* $\Omega(v)$

$$A \subseteq \Omega(v), \quad \mathbb{P}^v(A) = \frac{1}{p(v)}\mathbb{P}(A)$$

Probabilities induced in past and future

$\Omega(v) = \{\omega \in \Omega : \ \omega \supseteq v\}$ is the *shadow* of v

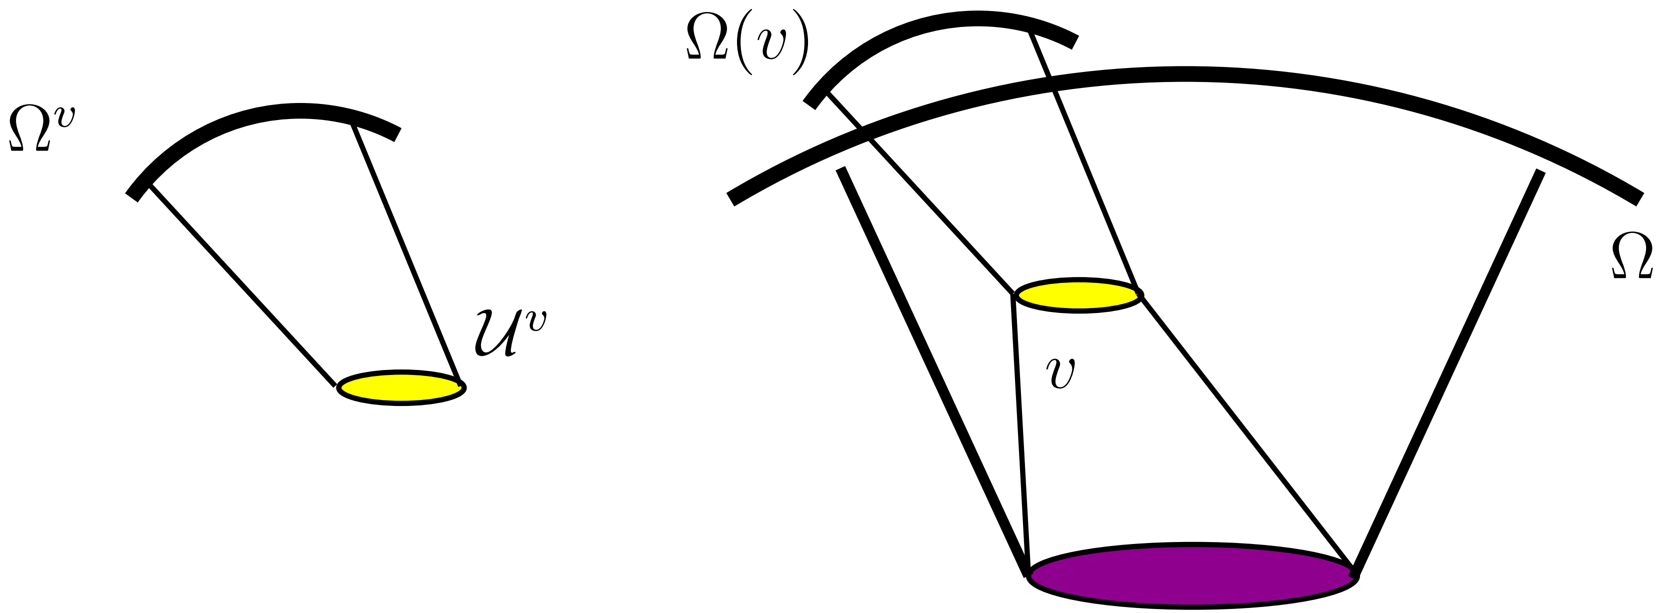

- **Probabilistic future** of configuration v: probability $\mathbb{P}^v$ on the *boundary* at infinity Ω^v of the future $\mathcal{U}^v$

$$A \subseteq \Omega^v, \quad \mathbb{P}^v(A) = \frac{1}{p(v)} \mathbb{P}(A)$$

Probabilities induced in past and future

$\Omega(v) = \{\omega \in \Omega : \ \omega \supseteq v\}$ is the *shadow* of v

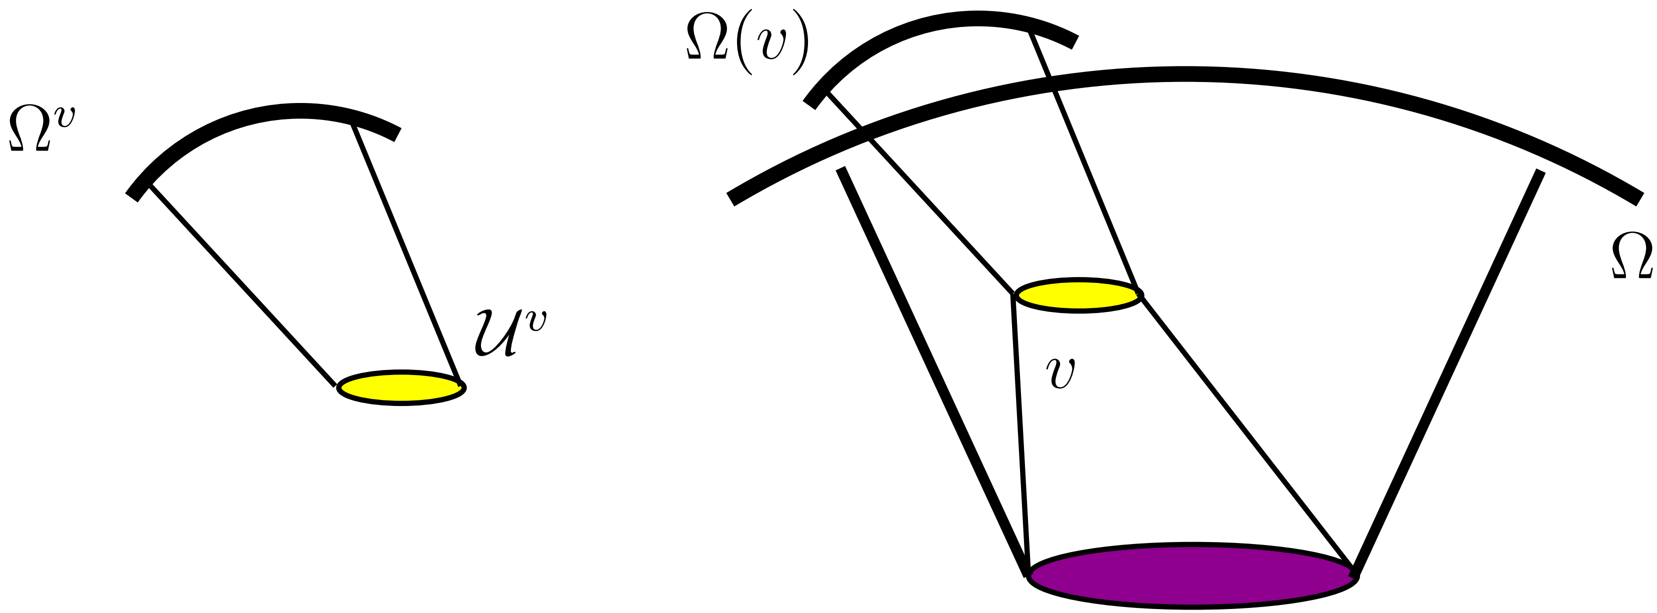

- **Property:** *conservation* of distributed products w.r.t. *future*

$$\mathbb{P} = \bigotimes_{\lambda \subseteq \mathcal{U}}^{\mathrm{d}} q_\lambda \quad \Longrightarrow \quad \mathbb{P}^v = \bigotimes_{\lambda \subseteq \mathcal{U}^v}^{\mathrm{d}} q_\lambda$$

Probabilities induced in past and future

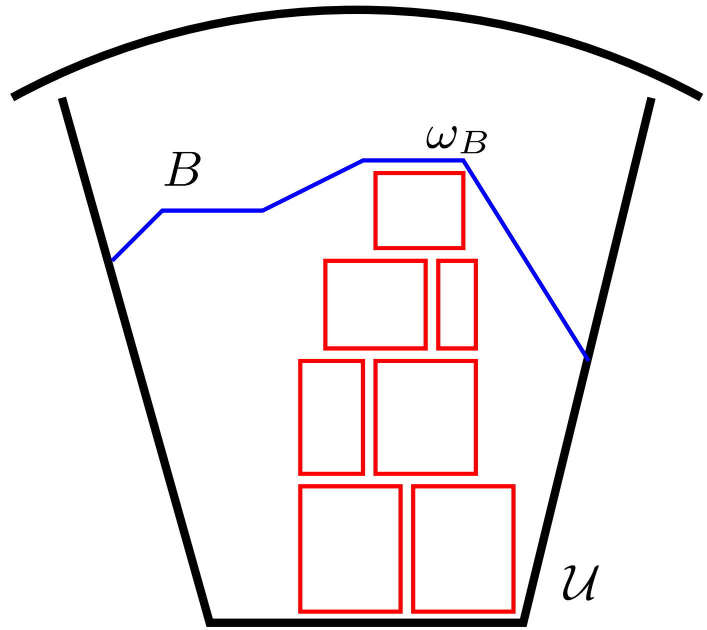

- B **stopping prefix** (*past*) $\rightarrow$ probability $\mathbb{P}_B$ on Ω_B

$$\mathbb{P}_B(\omega_B) = \mathbb{P}\big(\Omega(\omega_B)\big)$$

Probabilities induced in past and future

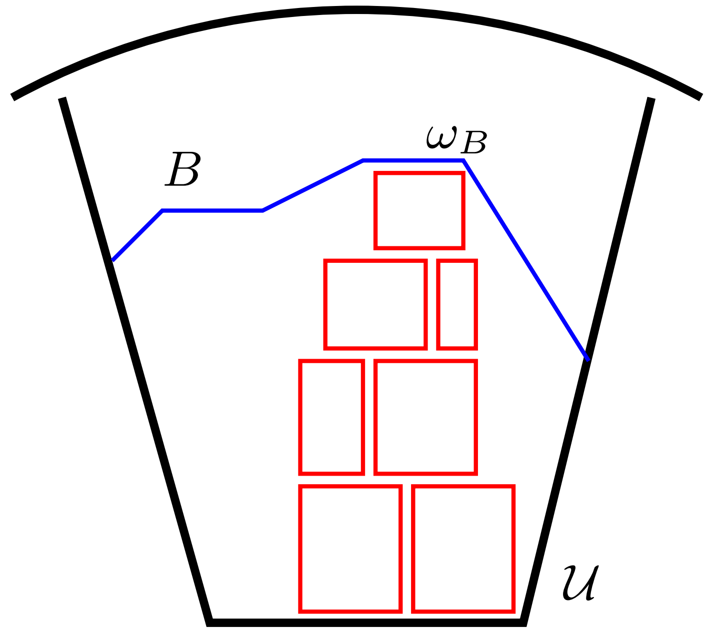

- **Property:** *conservation* of distributed products w.r.t. *past*

$$\mathbb{P} = \bigotimes_{\lambda \subseteq \mathcal{U}}^{\mathsf{d}} q_\lambda \quad \Longrightarrow \quad \mathbb{P}_B = \bigotimes_{\lambda \subseteq B}^{\mathsf{d}} q_\lambda$$

Probability and concurrency

- Let $\mathbb{P}$ be a distributed product $\mathbb{P} = \bigotimes^{\mathrm{d}}_{\lambda \subseteq \mathcal{U}} q_\lambda$. Fix v **w.s**

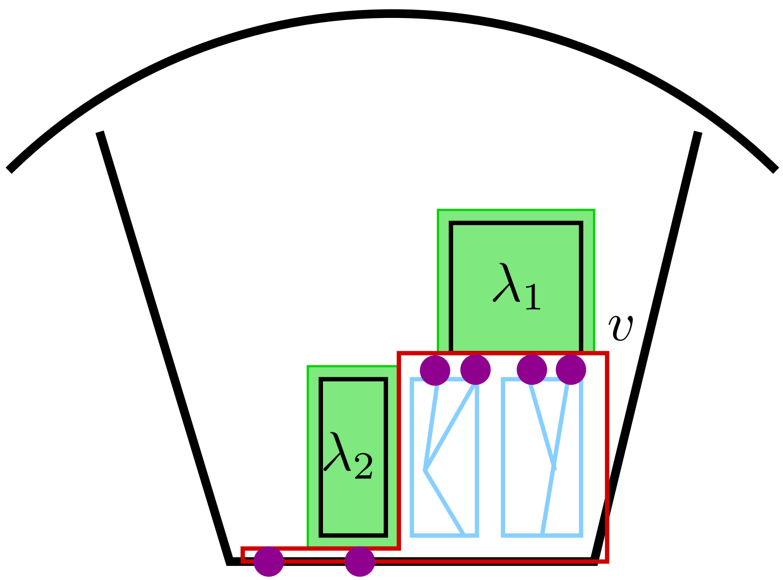

λ_1, λ_2 = branching cells enabled from v

(*minimal* $\neq \emptyset$ *stopping prefixes of* $\mathcal{U}^v$)

Set stopping prefix $B = \lambda_1 \cup \lambda_2$. Then: $\Omega_B = \Omega_{\lambda_1} \times \Omega_{\lambda_2}$ and

$$\mathbb{P}^v_B = \mathbb{P}^v_{\lambda_1} \otimes \mathbb{P}^v_{\lambda_2} = q_{\lambda_1} \otimes q_{\lambda_2}$$

Probability and concurrency

- if $\mathbb{P}$ is a distributed product $\mathbb{P} = \bigotimes_{\lambda \subseteq \mathcal{U}}^{d} q_\lambda$. Fix v **w.s**

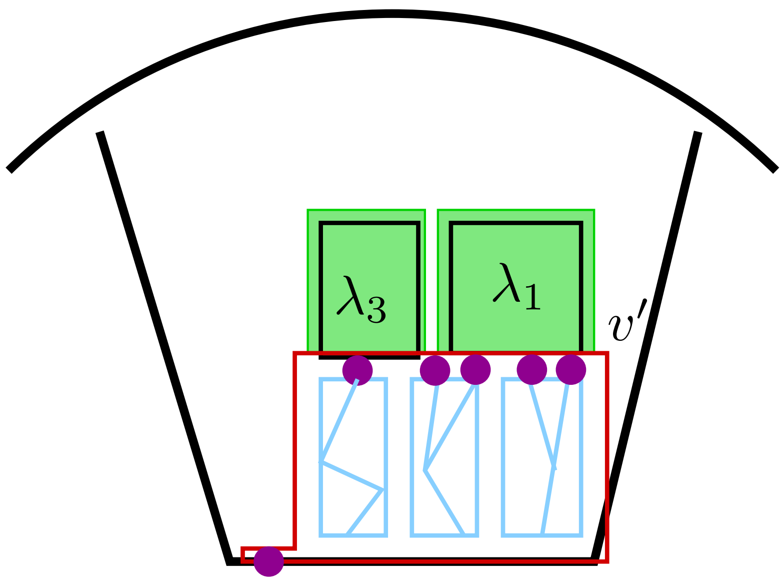

$$v' = v + \xi, \quad \xi \in \Omega_{\lambda_2}$$

Set stopping prefix $B' = \lambda_1 \cup \lambda_3$. Then:

$$\mathbb{P}_{B'}^{v'} = \mathbb{P}_{\lambda_1}^{v'} \otimes \mathbb{P}_{\lambda_3}^{v'} = q_{\lambda_1} \otimes q_{\lambda_3}$$

Probability and concurrency

Theorem: a probability $\mathbb{P}$ is a distributed product *iff*

- for every v finite **w.s**, the *product decomposition* holds:

$$\mathbb{P}^v_B = \mathbb{P}^v_{\lambda_1} \otimes \cdots \otimes \mathbb{P}^v_{\lambda_n}$$

  with $B = B^{\perp}(\mathcal{U}^v) = \lambda_1 \cup \ldots \cup \lambda_n$

- and for λ fixed, $\mathbb{P}^v_\lambda = q_\lambda$ is independent of v

In this case: $\mathbb{P} = \bigotimes^{\mathrm{d}}_{\lambda \subseteq \mathcal{U}} q_\lambda$

$\mathbb{P}$ is a **distributed** probability

Probability and concurrency

- *product form* for a distributed probability $\mathbb{P}$

$$B = \lambda_1 \cup \ldots \cup \lambda_n, \quad \mathbb{P}_B^v = q_{\lambda_1} \otimes \cdots \otimes q_{\lambda_n}$$

  - a horizontal independence due to *concurrency*
  - locality in *space*: **new feature**
  - randomization by *local agents*:
    - dynamic
    - without communication during asynchronous actions

1. Background: unfoldings and representations of space Ω
2. Extension of probabilities
3. Decomposition of true-concurrent processes
4. Distributed product of probabilities
5. *Markov nets: the Markov property and the Law of large numbers*
6. Computability of local finiteness
7. Conclusion and perspectives

Markov nets

- A distributed product is given by a *countable* family of finite probabilities $(q_\lambda)_\lambda$
- there are *finitely many* classes of branching cells as labelled occurrence nets

finite alphabet Σ = { classes of branching cells }

Markov nets

- A distributed product is given by a *countable* family of finite probabilities $(q_\lambda)_\lambda$
- there are *finitely many* classes of branching cells as labelled occurrence nets

finite alphabet Σ = { classes of branching cells }

- a **Markov net** is a pair $(\mathcal{N}, (q_s)_{s\in\Sigma})$, q_s a (finite) probability on Ω_s
- the associated distributed probability:

$$\langle\lambda\rangle = \text{class of } \lambda, \quad \mathbb{P} = \bigotimes_{\lambda\subseteq\mathcal{U}}^{\text{d}} q_{\langle\lambda\rangle}$$

Homogeneity and the Markov property

- **Homogeneity**: Markov net $(\mathcal{N}, (q_s)_\Sigma)$, 2 configurations v, v' finite **w.s** leading to *same* marking m

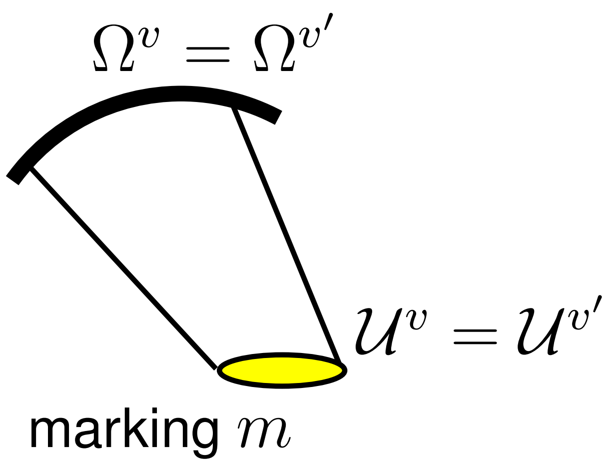

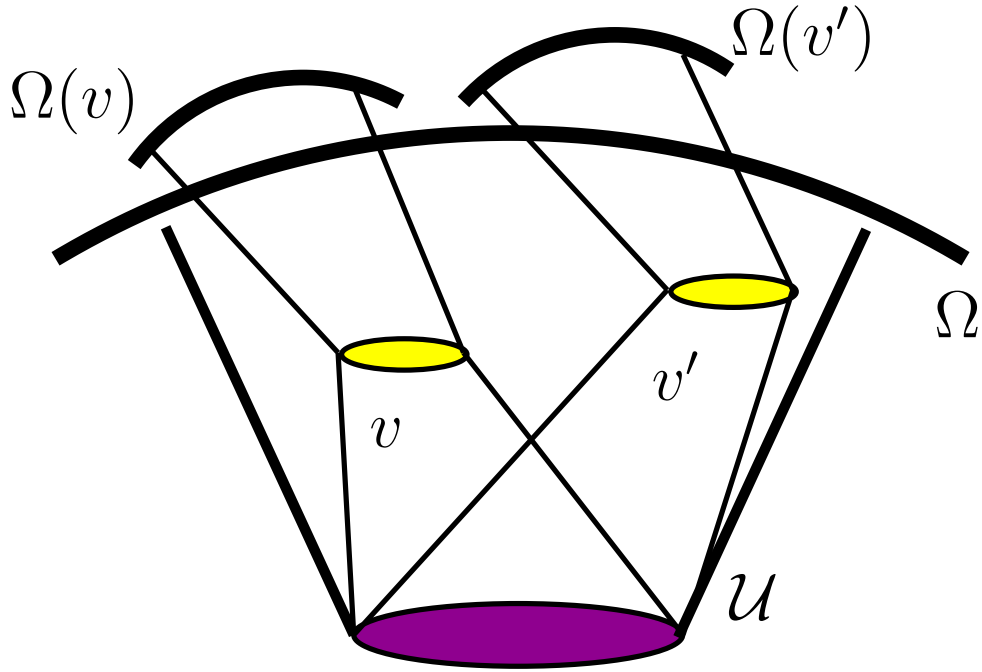

Homogeneity and the Markov property

- **Homogeneity**: Markov net $(\mathcal{N}, (q_s)_\Sigma)$, 2 configurations v, v' finite **w.s** leading to *same* marking m

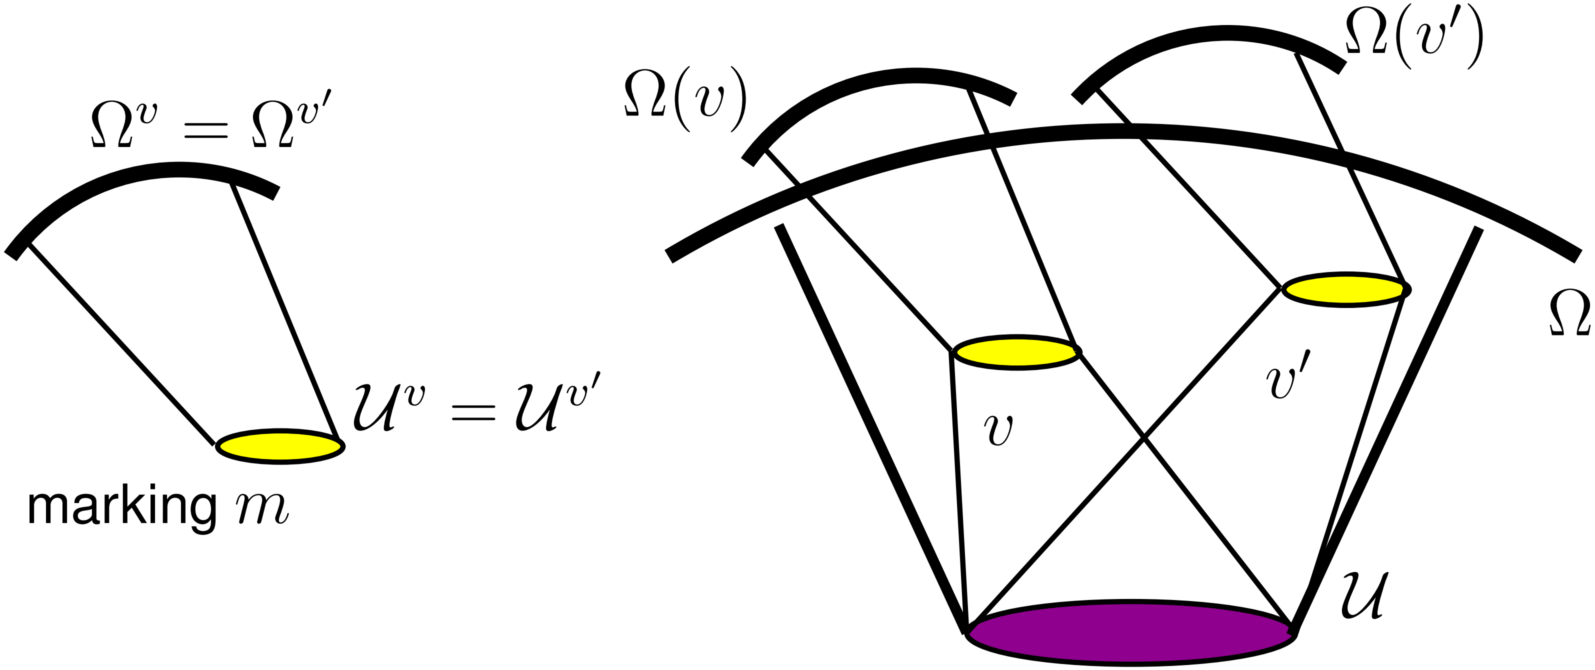

homogeneity $\quad \mathbb{P} = \bigotimes_{\lambda \subseteq \mathcal{U}}^{\mathsf{d}} q_{\langle\lambda\rangle} \quad \Longrightarrow \quad \mathbb{P}^v = \mathbb{P}^{v'} = \bigotimes_{\lambda \subseteq \mathcal{U}^v}^{\mathsf{d}} q_{\langle\lambda\rangle}$

Homogeneity and the Markov property

- **Homogeneity**: Markov net $(\mathcal{N}, (q_s)_\Sigma)$, 2 configurations v, v' finite **w.s** leading to *same* marking m

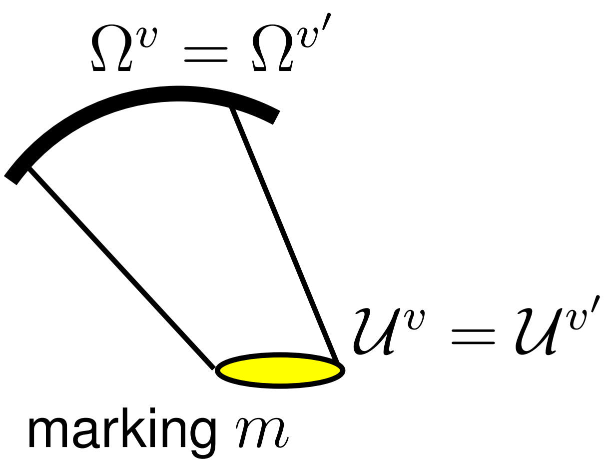

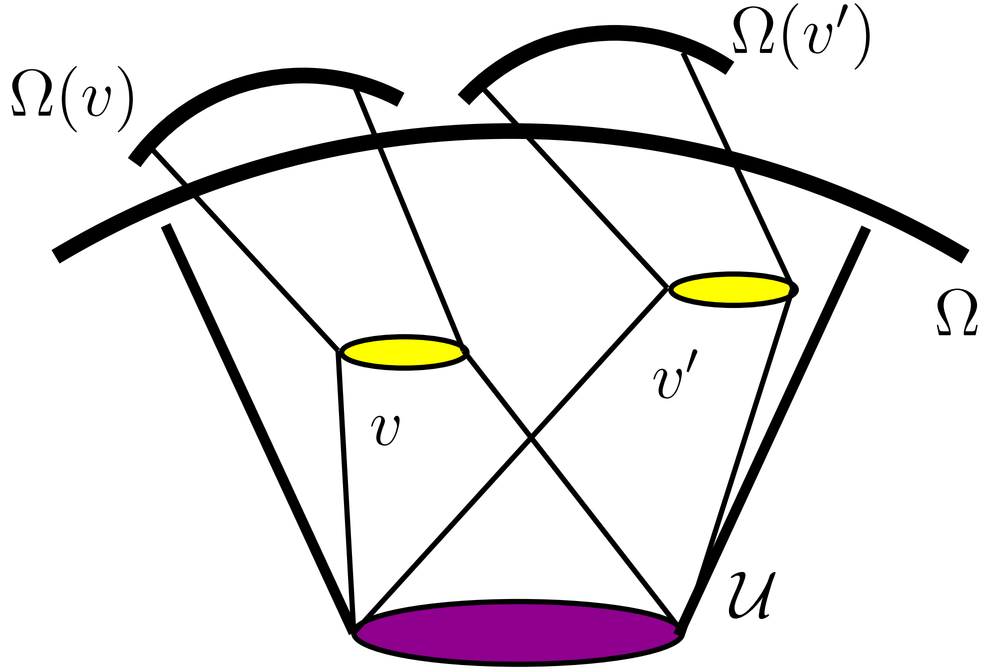

The probabilistic future $\mathbb{P}^v$ only depends on the marking $m(v)$

Homogeneity and the Markov property

- **Stopping operators** generalize *stopping times* for sequential systems
  A stopping operator is a random variable V such that:
  - $V(\omega)$ is a **w.s** configuration, $V(\omega) \subseteq \omega$
  - $\forall \omega, \omega' \in \Omega, \quad \omega' \supseteq V(\omega) \Rightarrow V(\omega') = V(\omega)$
- **Example:** (sequential) first return of the initial state

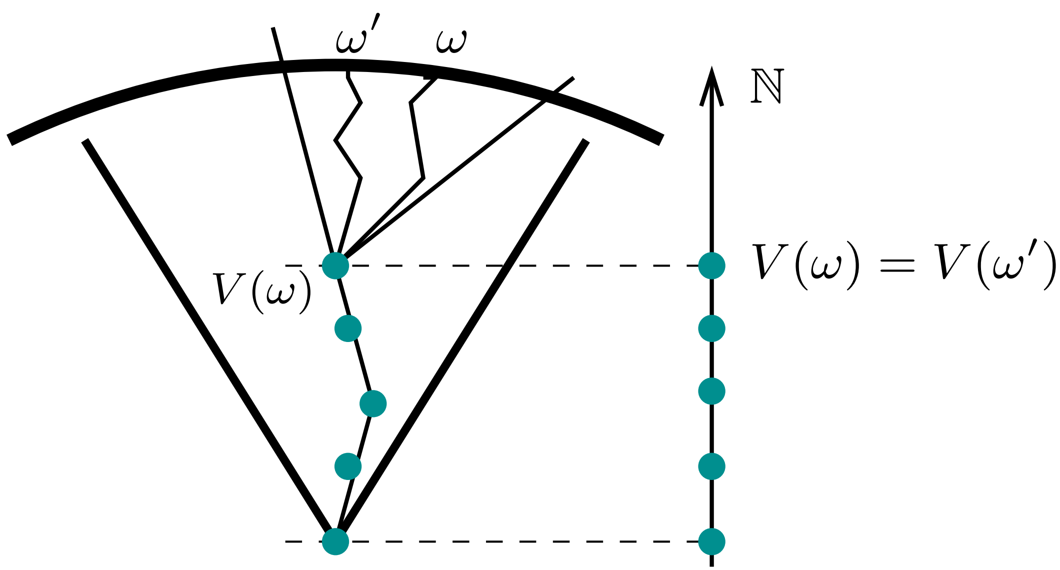

Homogeneity and the Markov property

- Reformulated with stopping operators, the **Strong Markov Property** (from Markov chains) holds for Markov nets
- adapt the (beginning of) recurrence theory of Markov chains $\rightarrow$ *global recurrence* of Markov nets
  - in a **recurrent** Markov net, reachable markings have probability 1 to return infinitely often
- there are results for a *local recurrence* (coincides with global recurrence for Markov chains)

The Law of large Numbers

- **Case of a recurrent Markov chain** $(X_n)_{n\geq 1}$
  with state space S
  $f : S \to \mathbb{R}$ a test function
  For *integer* $n \geq 1$:

  ergodic **sum**: $S_n f = f(X_1) + \cdots + f(X_n)$

  ergodic **mean**: $M_n f = \frac{1}{n} S_n f = \frac{\text{sum of outputs of } f}{\text{time elapsed}}$

  LLN: there is a probability α on S such that:

  $$\lim_{n\to\infty} M_n f = \alpha(f), \quad \mathbb{P}\text{-a.s.} \qquad \alpha(f) = \sum_{s\in S} \alpha(s) f(s)$$

The Law of large Numbers

- **Case of a recurrent Markov chain** $(X_n)_{n\geq 1}$
  with state space S
  $f : S \to \mathbb{R}$ a test function
  For *integer* $n \geq 1$:

  ergodic **sum**: $S_n f = f(X_1) + \cdots + f(X_n)$

  ergodic **mean**: $M_n f = \frac{1}{n} S_n f = \frac{\text{sum of outputs of } f}{\text{time elapsed}}$

- For **concurrent** systems
  - what is the state space? what are the test functions?
  - what is the time elapsed?

The Law of large Numbers

- **Case of a recurrent Markov chain**

f is tested *along* configuration v

$$\langle f, v \rangle = \sum_{x \in v} f(\mathbf{x})$$
$$= f(\mathbf{s_1}) + f(\mathbf{s_2}) + f(\mathbf{s_1})$$

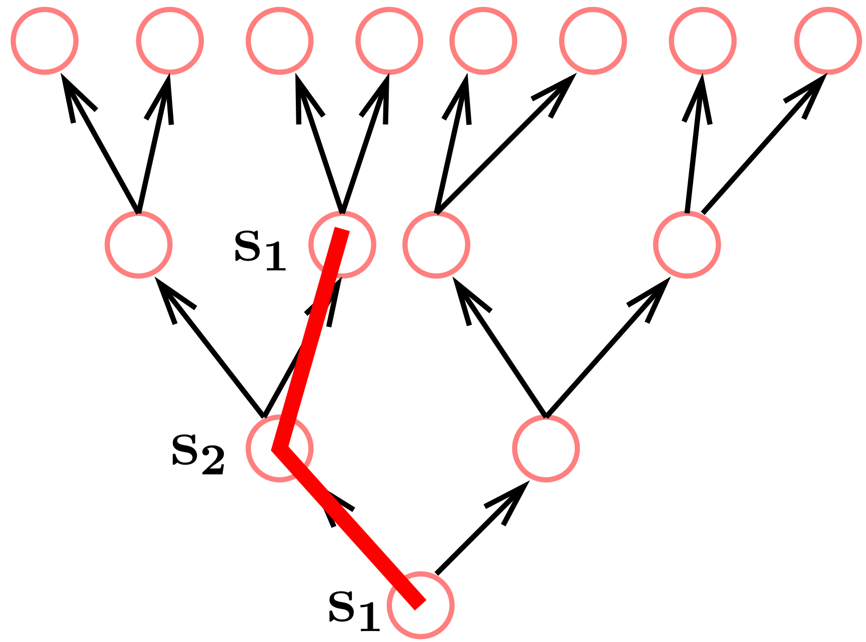

duration of v = $\langle 1, v \rangle$ = 3

Ergodic mean of f *along* v

$$Mf(v) = \frac{\langle f, v \rangle}{\langle 1, v \rangle}$$

The Law of large Numbers

Classes of branching cells act like *local states*

A **distributed function** is a (finite) *family* $f = (f_s)_{s\in\Sigma}$ of real valued functions $f_s : \Omega_s \to \mathbb{R}$

$$\langle f, v\rangle = \sum_\lambda f_{\langle\lambda\rangle}(v \cap \lambda)$$

$\langle \mathbf{1}, v\rangle = 6 =$ duration of v

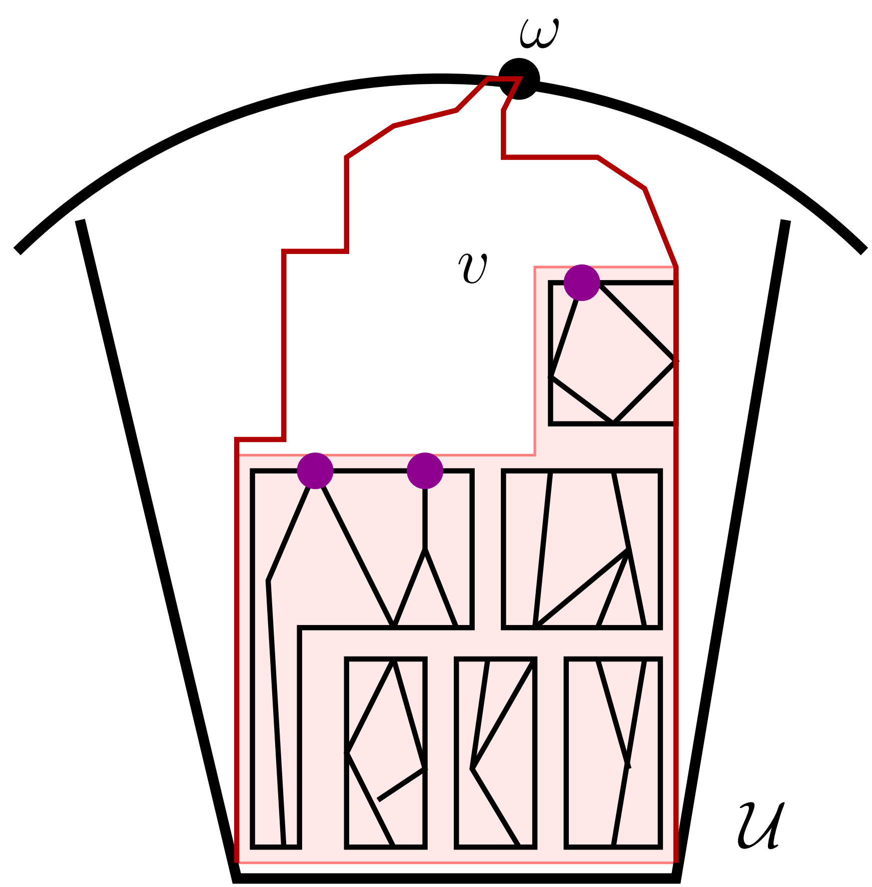

limit of ergodic means: $Mf(v) = \dfrac{\langle f, v\rangle}{\langle \mathbf{1}, v\rangle}, \quad v \to \omega?$

The Law of large Numbers

A sequence of stopping operators $(V_n)_{n\geq 1}$ is **regular** if:

- for all n, $V_n \subseteq V_{n+1}$
- $\bigcup_n V_n(\omega) = \omega$ with probability 1
- there are $K_1, K_2 > 0$ such that for all n:

$$K_1 \leq \frac{\langle \mathbf{1}, V_n \rangle}{n} \leq K_2$$

The Law of large Numbers

A sequence of stopping operators $(V_n)_{n\geq 1}$ is **regular** if:

- for all n, $V_n \subseteq V_{n+1}$
- $\bigcup_n V_n(\omega) = \omega$ with probability 1
- there are $K_1, K_2 > 0$ such that for all n:

$$K_1 \leq \frac{\langle \mathbf{1}, V_n \rangle}{n} \leq K_2$$

Definition: For a distributed function $f = (f_s)_{s\in\Sigma}$, the ergodic means $Mf(\cdot) = \dfrac{\langle f, \cdot \rangle}{\langle \mathbf{1}, \cdot \rangle}$ **converge** to a function $\mu : \Omega \to \mathbb{R}$ if, for *every* regular sequence of stopping operators $(V_n)_{n\geq 1}$:

$$\lim_{n\to\infty} Mf\big(V_n(\omega)\big) = \mu(\omega), \quad \mathbb{P}\text{-a.s.}$$

The Law of large Numbers

The convergence of ergodic means **cannot** hold if net $\mathcal{N}$ is the product of two independent components $\mathcal{N}_1 \cup \mathcal{N}_2$
$\rightarrow$ need for a *synchrony* assumption

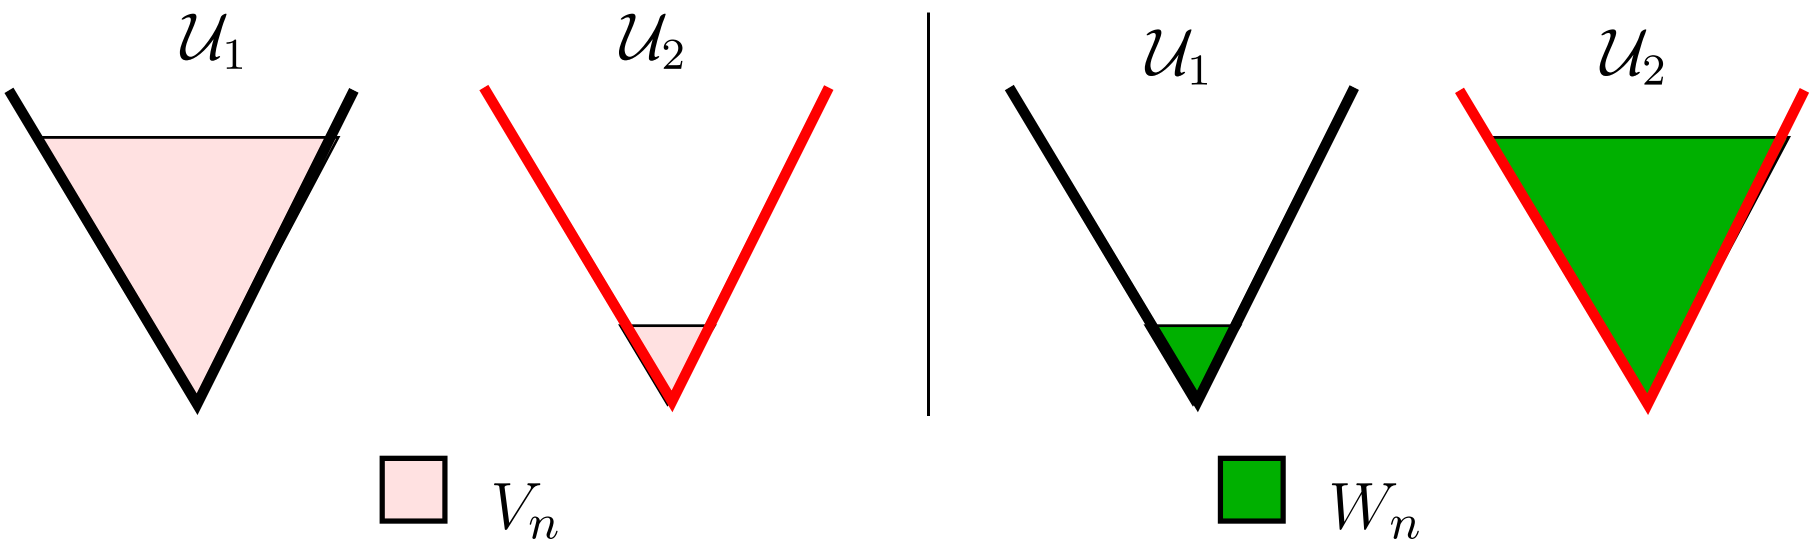

$Mf(V_n)$ and $Mf(W_n)$ have different limits

The Law of large Numbers

The convergence of ergodic means **cannot** hold if net $\mathcal{N}$ is the product of two independent components $\mathcal{N}_1 \cup \mathcal{N}_2$
$\rightarrow$ need for a *synchrony* assumption

- Markov net $\mathcal{N}$ has **integrable concurrency height** if for each partial execution of the system, leading to marking m, and for each place P of m, there is a time of *finite expectation* before the token in place P moves.

The Law of large Numbers

Theorem (LLN) Let $(\mathcal{N}, (q_s)_{s\in\Sigma})$ be a Markov net, **recurrent** and with **integrable concurrency height**.

- For $f = (f_s)_{s\in\Sigma}$ a distributed function, the *ergodic means* $Mf(\cdot)$ *converge* to a function $\mu f : \Omega \to \mathbb{R}$, and μf is *constant* with probability 1.
- There is a (finite) probability α on Σ s.t.:

$$\mu f = \sum_{s\in\Sigma} \alpha(s) q_s(f_s)$$

Comments on the LLN

$$\mu f = \sum_{s \in \Sigma} \alpha(s) q_s(f_s)$$

- If $\mathcal{N}$ is actually a Markov chain, $\alpha(s)$'s are the coefficients from the sequential LLN (*stationary* measure).
- Classes of branching cells $s \in \Sigma$ appear as *local states* of the concurrent system.
- coefficients $\alpha(s)$ is the *asymptotic density* of local state $\alpha(s)$

1. Background: unfoldings and representations of space Ω
2. Extension of probabilities
3. Decomposition of true-concurrent processes
4. Distributed product of probabilities
5. Markov nets: the Markov property and the Law of large numbers
6. *Computability of local finiteness*
7. Conclusion and perspectives

Conjecture and consequences

Conjecture: $\mathcal{U}$ the unfolding of a safe Petri net $\mathcal{N}$.

Assume that for every event e, the set:

$$\{f \in \mathcal{U} : \; f \;\#_\mu\; e\}$$

is *finite*.

Then $\mathcal{U}$ is *locally finite*.

Conjecture and consequences

Consequences

- $\mathcal{U}$ locally finite $\Rightarrow \Omega$ compact **OK**

  With the conjecture:

  - the *converse* holds
  - locally finite constructions are made easy

- Local finiteness is *decidable*,

  branching cells are *computable*

Decidability of local finiteness

$\mathcal{N}$ a safe Petri net with unfolding $\mathcal{U}$

Question: is $\mathcal{U}$ locally finite?

Reductions:

- decide the finiteness of $\{f \in \mathcal{U} \,:\, f \ \#_\mu \ e\}$ for $e \in \mathcal{U}$ (conjecture)
- assume that e is *minimal* in $\mathcal{U}$
- decide the finiteness of $F_{t'}(e) = \{f \in \mathcal{U} \,:\, f \ \#_\mu \ e, \ f$ labelled by $t'\}$ for t' a fixed transition
- assume that $F_{t'}(e)$ contains an event f minimal in $\mathcal{U}$

Decidability of local finiteness

$\mathcal{N}$ a safe Petri net with unfolding $\mathcal{U}$, $\quad \rho : \mathcal{U} \to \mathcal{N}$

- fix e minimal event of $\mathcal{U}$, labelled by transition $\mathbf{t}$
- assume that $F_{\mathbf{t}'}(e) = \{f \in \mathcal{U} \ : \ f \ \#_\mu \ e, \ \rho(f) = \mathbf{t}'\}$ contains an event minimal in $\mathcal{U}$
- **Question:** finiteness of $F_{\mathbf{t}'}$?

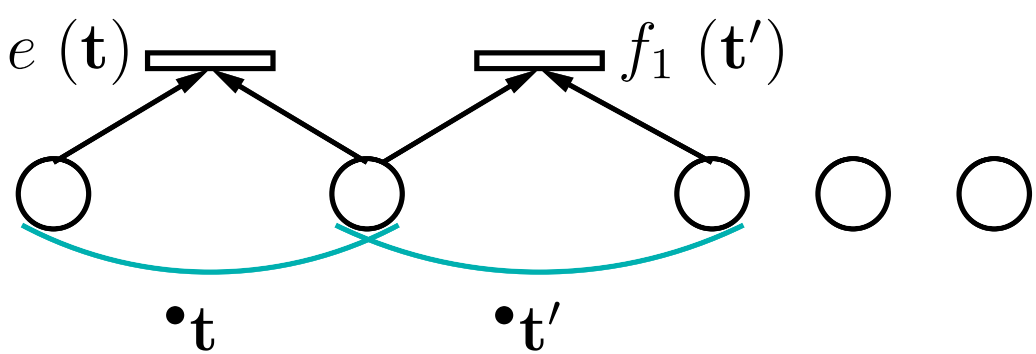

Decidability of local finiteness

$\mathcal{N}$ a safe Petri net with unfolding $\mathcal{U}$, $\quad \rho : \mathcal{U} \to \mathcal{N}$

- fix e minimal event of $\mathcal{U}$, labelled by transition $\mathbf{t}$
- assume that $F_{\mathbf{t}'}(e) = \{f \in \mathcal{U} \ : \ f \ \#_{\mu} \ e, \ \rho(f) = \mathbf{t}'\}$ contains an event minimal in $\mathcal{U}$
- **Question:** finiteness of $F_{\mathbf{t}'}$?

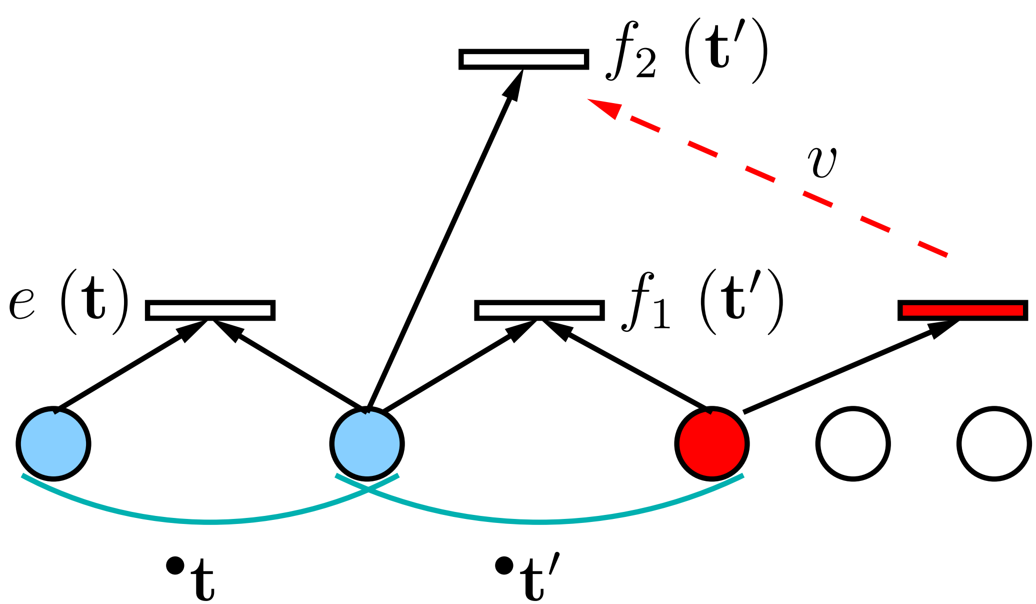

If v is a configuration that enables $f_2 \in F_{\mathbf{t}'}(e)$

- tokens in $^\bullet\mathbf{t}$ have not moved
- at least one token in $^\bullet\mathbf{t}'$ has moved

Decidability of local finiteness

$\mathcal{N}$ a safe Petri net with unfolding $\mathcal{U}$, $\quad \rho : \mathcal{U} \to \mathcal{N}$

- fix e minimal event of $\mathcal{U}$, labelled by transition $\mathbf{t}$
- assume that $F_{\mathbf{t}'}(e) = \{f \in \mathcal{U} \; : \; f \; \#_\mu \; e, \; \rho(f) = \mathbf{t}'\}$ contains an event minimal in $\mathcal{U}$
- **Question:** finiteness of $F_{\mathbf{t}'}$?
  - Draw a finite graph in the submarkings of $M_0 \setminus {}^\bullet\mathbf{t}$, M_0 = initial marking of $\mathcal{N}$
  - $F_{\mathbf{t}'}(e)$ is infinite *if and only if* the graph has a cycle
- **Conclusion:** under the conjecture, local finiteness is decidable

1. Background: unfoldings and representations of space Ω
2. Extension of probabilities
3. Decomposition of true-concurrent processes
4. Distributed product of probabilities
5. Markov nets: the Markov property and the Law of large numbers
6. Computability of local finiteness
7. *Conclusion and perspectives*

Conclusion

- **probabilistic framework for true-concurrency models**
  - *extension* finite Markov chains theory to safe Petri nets
  - construction of a *Markovian* probability,
    from a *finite* number of *local* parameters
  - concurrency matches a probabilistic independence

Conclusion

Contributions

- *Continuous domain of configurations*
  - identification of the space Ω as a *projective limit*
  - *locally finite* unfoldings and extension of probabilities
- *Occurrence nets*
  - decomposition of true-concurrent processes through *branching cells* (local states)
  - computability
- *Probabilistic model*
  - construction of the *distributed product*
  - stopping operators and the Strong *Markov property*
  - part of a *recurrence* theory
  - Law of large numbers

Extensions

Open questions

- above conjecture: topological consequences and decidability of local finiteness
- more about the density coefficients of the LLN: positivity? (potential theory)
- Central Limit Theorem? (Martingales)
- branching cells form a (non prime) event structure?

Extensions

Extensions

- non locally finite nets and *products* of nets
- Distributed *HMM* (Hidden Markov Models)?
  Probabilistic extension of diagnosis algorithms
- Temporisation
  - add *temporisation* after randomization of runs
    $\rightarrow$ performance evaluation
  - Markov nets as a “uniformisation” of stochastic Petri nets?