PETRI NETS

Petri nets

A Petri net is a tuple $(\mathcal{P}, \mathcal{E}, \mathbf{D}^-, \mathbf{D}^+, \mathbf{x}_{init})$ where:

- $\mathcal{P}$ set of places, drawn as circles
- $\mathcal{E}$ set of transitions drawn as rectangles
- $\mathbf{D}^- : \mathcal{P} \times \mathcal{E} \to \mathbb{N}$ input arc cardinalities
- $\mathbf{D}^+ : \mathcal{P} \times \mathcal{E} \to \mathbb{N}$ output arc cardinalities
- $\mathbf{x}_{init} \in \mathbb{N}^{|\mathcal{P}|}$ initial state, or marking

with $\mathcal{P} \cap \mathcal{E} = \emptyset$

Condition for transition α to be enabled in state $\mathbf{i} \in \mathbb{N}^{|\mathcal{P}|}$: $\alpha \in \mathcal{E}(\mathbf{i}) \;\Leftrightarrow\; \forall p \in \mathcal{P},\ \mathbf{D}^-_{p,\alpha} \leq \mathbf{i}_p$

A transition α enabled in state $\mathbf{i}$ can fire: $\mathbf{i} \xrightarrow{\alpha} \mathbf{j} \;\Leftrightarrow\; \forall p \in \mathcal{P},\ \mathbf{j}_p = \mathbf{i}_p - \mathbf{D}^-_{p,\alpha} + \mathbf{D}^+_{p,\alpha}$

Thus, the next-state function $\mathcal{N}$ satisfies $\mathbf{j} \in \mathcal{N}(\mathbf{i}) \;\Leftrightarrow\; \exists \alpha \in \mathcal{E},\ \mathbf{i} \xrightarrow{\alpha} \mathbf{j}$

The state space, or reachability set, $\mathcal{X}_{reach}$ is defined as usual

Graphical representation of a Petri net

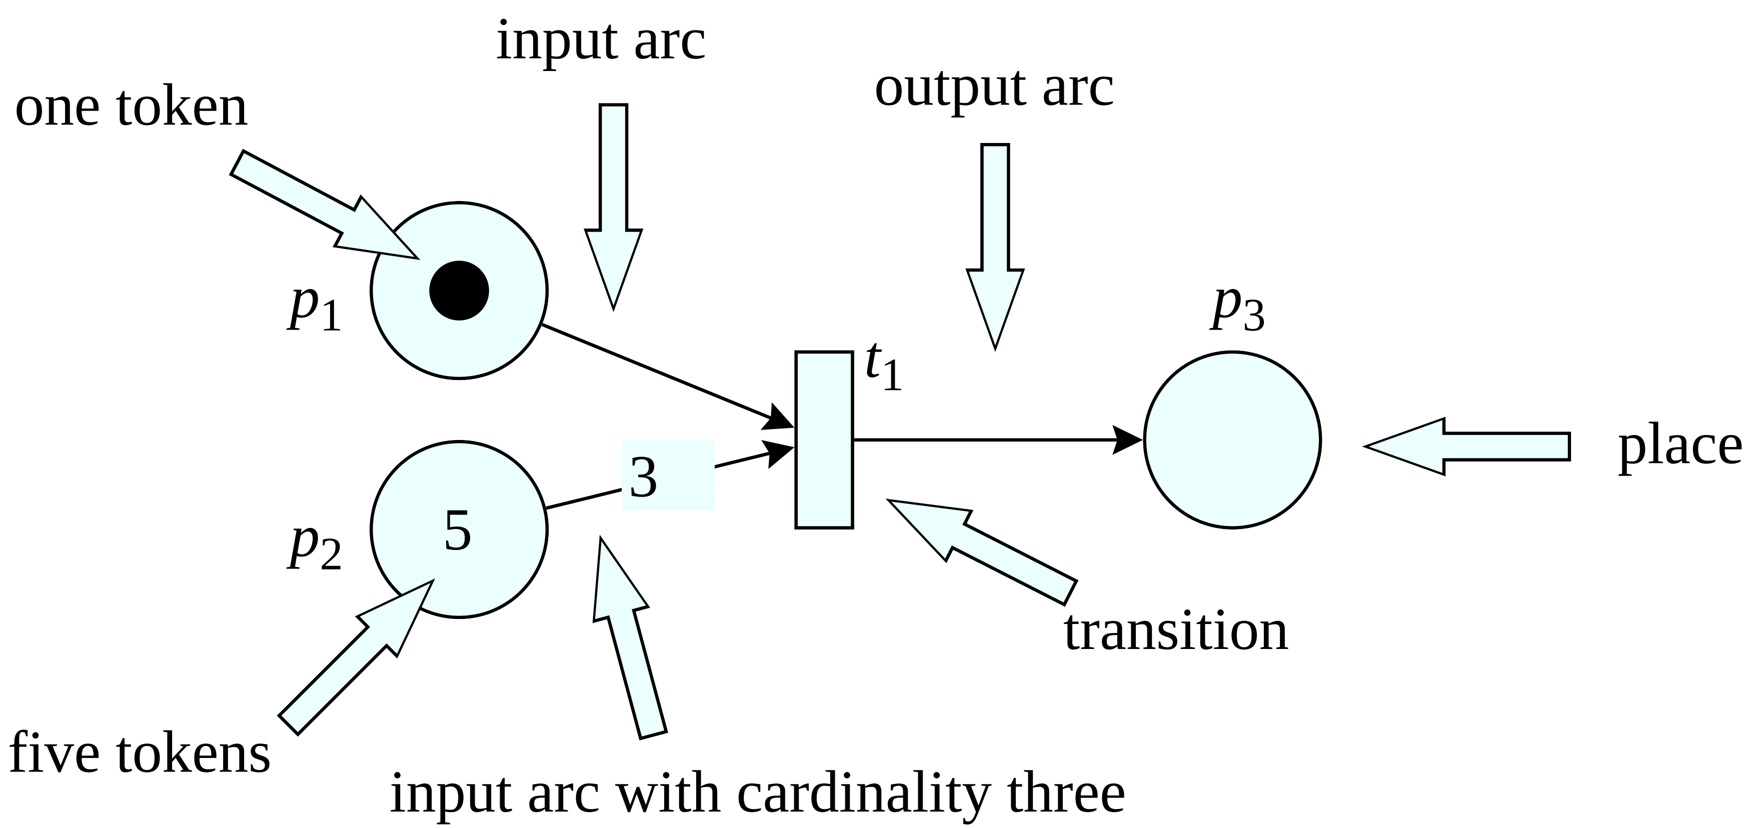

Dynamic evolution of a Petri net

Enabling rule

$e \in \mathcal{E}(\mathbf{i})$ iff each input arc contains at least as many tokens as the cardinality of the input arc:

$$\forall p \in \mathcal{P},\ \mathbf{D}^-_{p,e} \le \mathbf{i}_p \quad \text{or, in vector form} \quad \mathbf{D}^-_{\bullet,e} \le \mathbf{i}$$

Firing rule

If $\mathbf{i} \xrightarrow{e} \mathbf{j}$, we obtain $\mathbf{j}$ by removing tokens from input places and adding tokens to output places:

$$\forall p \in \mathcal{P},\ \mathbf{j}_p = \mathbf{i}_p - \mathbf{D}^-_{p,e} + \mathbf{D}^+_{p,e} \quad \text{or, in vector form} \quad \mathbf{j} = \mathbf{i} - \mathbf{D}^-_{\bullet,e} + \mathbf{D}^+_{\bullet,e} = \mathbf{i} + \mathbf{D}_{\bullet,e}$$

where $\mathbf{D} = \mathbf{D}^+ - \mathbf{D}^-$ is the incidence matrix

For example, if t_1 fires:

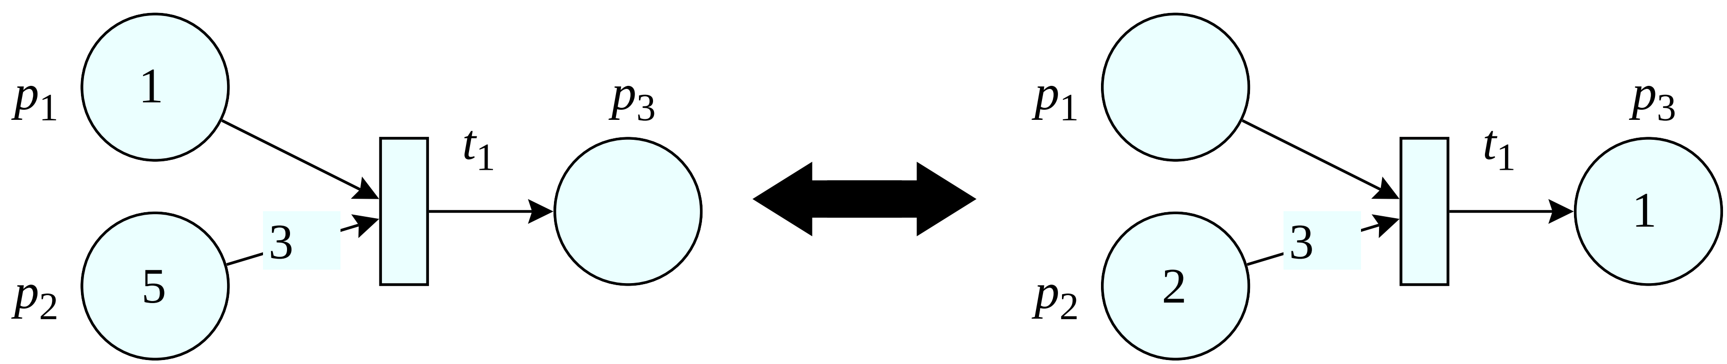

Reachability set and reachability graph

- $\mathcal{X}_{reach}$ is the reachability set: $\mathcal{X}_{reach} = \{\mathbf{i} : \exists \sigma \in \mathcal{E}^*, \mathbf{x}_{init} \xrightarrow{\sigma} \mathbf{i}\}$
- $(\mathcal{X}_{reach}, \mathcal{A})$ is the reachability graph: $\mathcal{A} = \{(\mathbf{i} \xrightarrow{\alpha} \mathbf{j}) : \exists \mathbf{i}, \mathbf{j} \in \mathcal{X}_{reach} \wedge \exists \alpha \in \mathcal{E}(\mathbf{i}) \wedge \mathbf{i} \xrightarrow{\alpha} \mathbf{j}\}$

If the PN is not simple, there can be multiple arcs between the same pair of nodes:

$$(\mathbf{i} \xrightarrow{\alpha} \mathbf{j}) \in \mathcal{A} \wedge (\mathbf{i} \xrightarrow{\beta} \mathbf{j}) \in \mathcal{A} \wedge \alpha \neq \beta$$

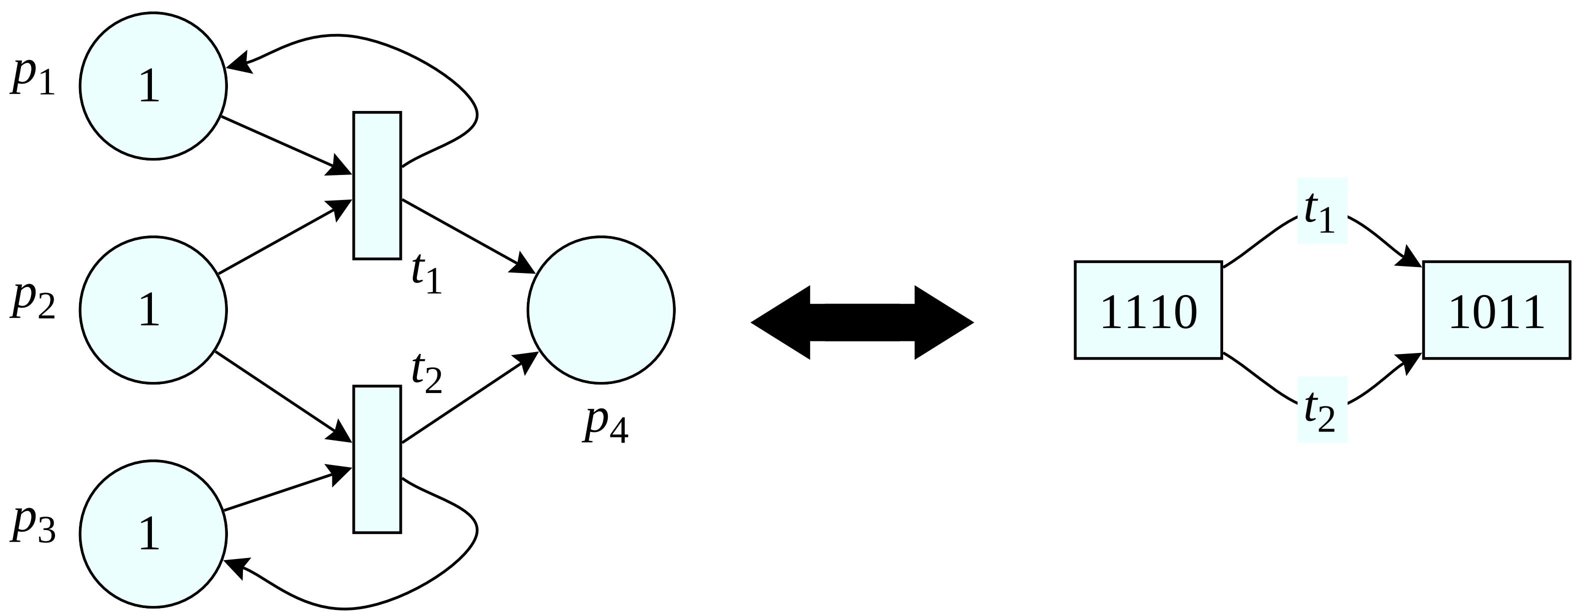

Petri nets and their state space $\mathcal{X}_{reach}$: finite case

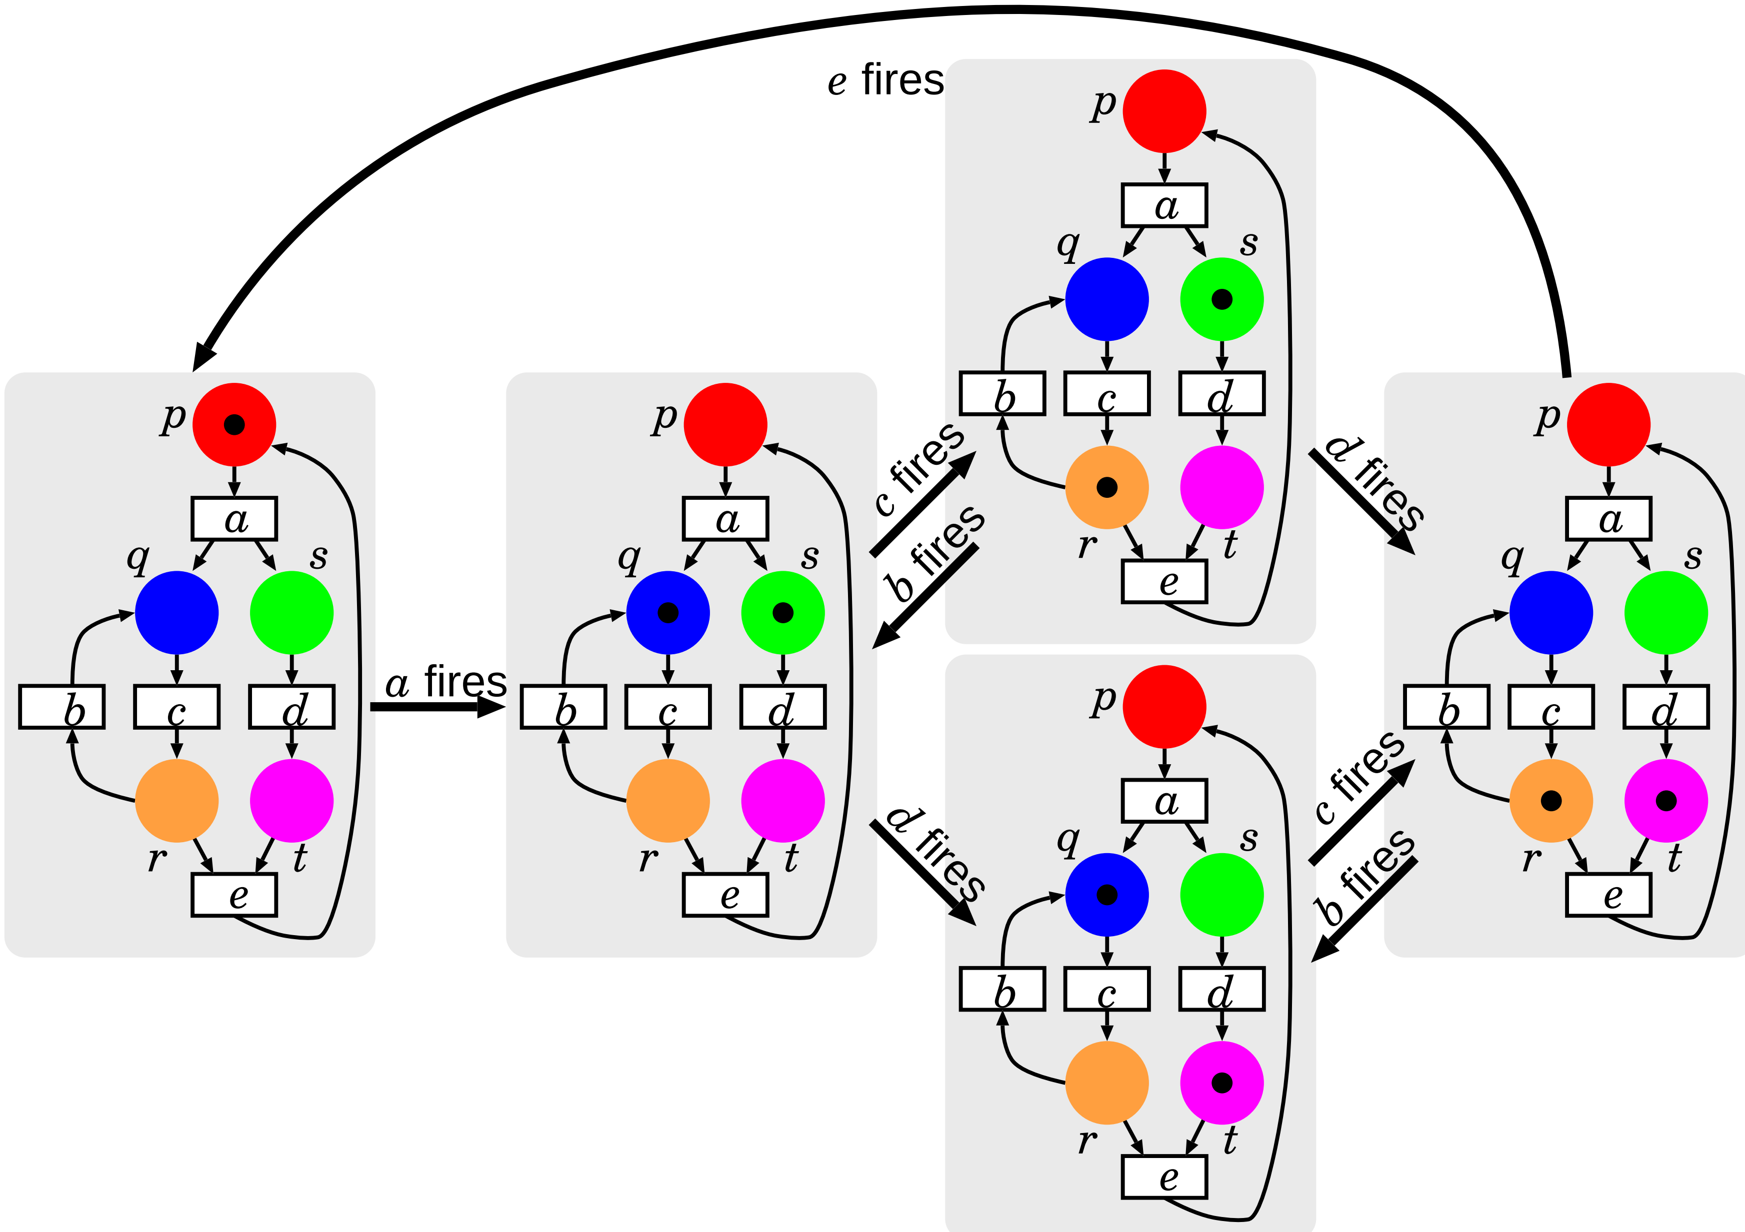

If the initial state is $\mathbf{x}_{init} = (N, 0, 0, 0, 0)$, $\mathcal{X}_{reach}$ contains $\dfrac{(N+1)(N+2)(2N+3)}{6}$ states

Petri nets and their state space $\mathcal{X}_{reach}$: infinite case

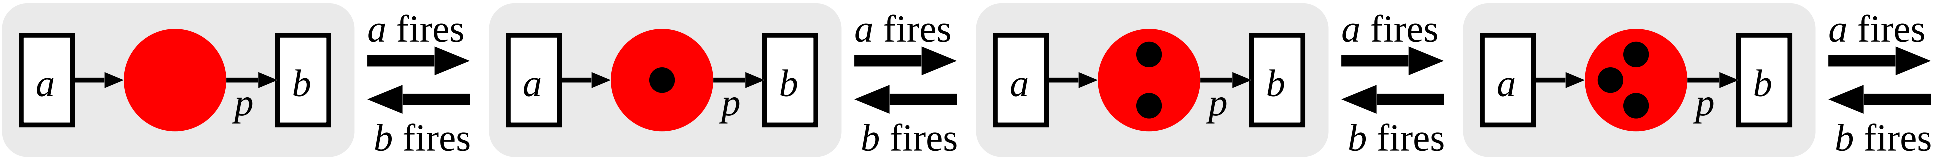

$\mathcal{X}_{reach}$ contains an infinite number of states regardless of the initial state $\mathbf{x}_{init} = (N)$

Example of Petri net

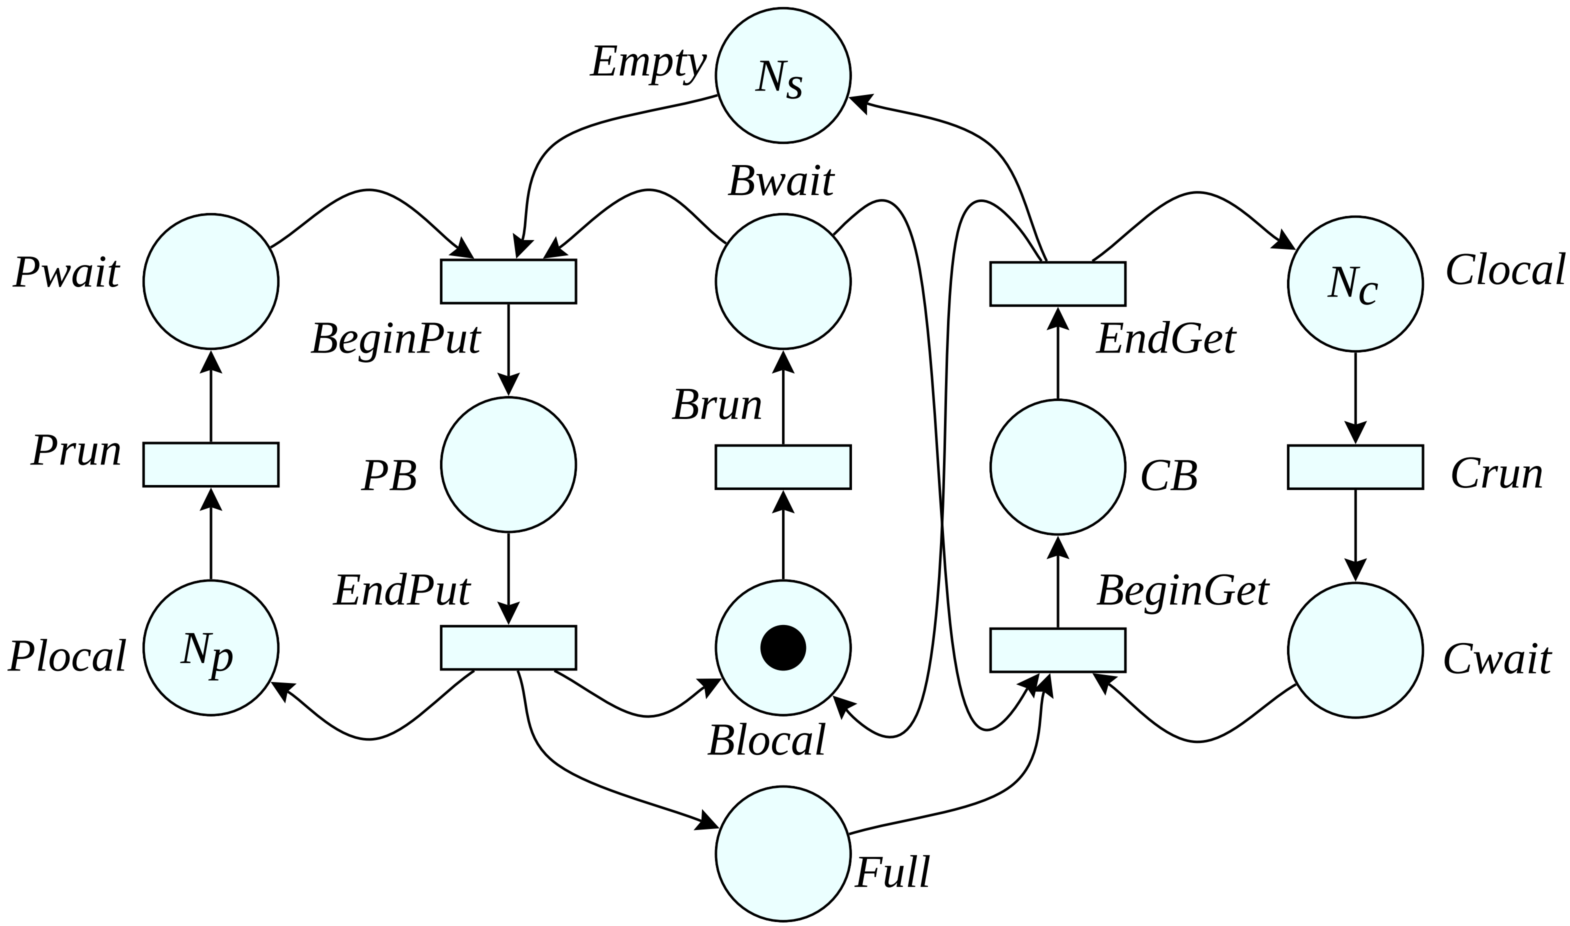

Events $Prun$, $Crun$, and $Brun$ can be changed into a more detailed subnet if more information is available about the actual production, consumption, and buffering processes

By eliminating place $Empty$, we model a system with an infinite number of slots: producers are never blocked because no empty slot is available, but consumers are still blocked if no slot is full

What behaviors can Petri nets model?

The producer-consumer Petri net exhibits:

- Splitting and synchronization (fork and join)

  $BeginPut$ can fire only when there is (at least) one token in $Pwait$, $Bwait$, and $Empty$

  When $EndPut$ fires, the token in PB “splits” back into three tokens

- Parallelism (concurrent activities)

  $Prun$ and $Brun$ are independently enabled by tokens in $Plocal$ and $Blocal$, respectively

- Contention (conflicting activities)

  When both $BeginPut$ and $BeginGet$ are enabled, the firing of one disables the other

- Sequentialization (causality)

  The firing of $EndPut$ enables $Brun$, the two activities happen in sequence

- Mutual exclusion

  There can never be a token in PB and CB at the same time

- Confusion

  $Prun$ and $BeginGet$ are not in conflict

  But the firing of $Prun$ may enable $BeginPut$, which is in conflict with $BeginGet$

Behavioral properties (hold for a given initial state)

Given a particular Petri net with a particular initial state $\mathbf{x}_{init}$, we can ask:

- Reachability: Is state $\mathbf{i}$ reachable from $\mathbf{x}_{init}$?

$$\mathbf{i} \in \mathcal{X}_{reach}$$

  Coverability: Given $\mathbf{i}$, is there a $\mathbf{j} \in \mathcal{X}_{reach}$ with at least as many tokens as $\mathbf{i}$ in each place?

$$\exists \mathbf{j} \in \mathcal{X}_{reach},\ \mathbf{j} \geq \mathbf{i}$$

  E.g., search for "bad" states

  - we don't want a token in PB and a token in CB

- Boundedness: Given a place p, is the number of tokens in p always bounded by b?

$$\forall \mathbf{i} \in \mathcal{X}_{reach},\ \mathbf{i}_p \leq b$$

  An unbounded net may reveal an error either in the system or in the model

  - all places are bounded
  - if we remove $Empty$, the system has an infinite number of slots, $Full$ is unbounded

Behavioral properties (cont.)

- Liveness: Given an event e, can it be eventually fired from any reachable state?

$$\forall \mathbf{i} \in \mathcal{X}_{reach},\ \exists \sigma \in \mathcal{E}^*,\ \exists \mathbf{j} \in \mathcal{X}_{reach},\ \mathbf{i} \xrightarrow{\sigma} \mathbf{j} \wedge e \in \mathcal{E}(\mathbf{j})$$

  A non-live event might indicate a deadlock or livelock

  - all the events are live

- Reversibility: Can the initial state $\mathbf{x}_{init}$ be reached from any reachable state?

$$\forall \mathbf{i} \in \mathcal{X}_{reach},\ \exists \sigma \in \mathcal{E}^*,\ \mathbf{i} \xrightarrow{\sigma} \mathbf{x}_{init}$$

  A non reversible net might indicate multistable behavior or the inability to reset the system

  - the Petri net is reversible for any initial state where N_p, N_s, and N_c are positive

- Persistence: Once event e is enabled, can it be disabled only by its own firing?

$$\forall \mathbf{i} \in \mathcal{X}_{reach},\ \forall f \in \mathcal{E}(\mathbf{i}),\ e \neq f,\ e \in \mathcal{E}(\mathbf{i}) \wedge \mathbf{i} \xrightarrow{f} \mathbf{j} \ \Rightarrow\ e \in \mathcal{E}(\mathbf{j})$$

  If events are persistent, timing does not affect certain logical properties

  - $BeginPut$ and $BeginGet$ disable each other, all other events are persistent

Behavioral properties (cont.)

- Synchronic distance: What is the maximum difference of the number of times e and f can fire in a firing sequence starting in any reachable state?

$$synch\_dist_{e,f} = \max_{\sigma \in \mathcal{E}^*} \{0, |\sigma|_e - |\sigma|_f : \exists \mathbf{i}, \mathbf{j} \in \mathcal{X}_{reach},\ \mathbf{i} \xrightarrow{\sigma} \mathbf{j}\}$$

  ($|\sigma|_e$ is the number of occurrences of e in σ)
  A finite synchronic distance implies that the events are "somewhat synchronized"

  - $synch\_dist_{Prun,BeginPut} = N_p$ $\quad$ $synch\_dist_{Prun,Brun} = +\infty$
    $synch\_dist_{Prun,Crun} = N_p + N_s$ $\quad$ $synch\_dist_{BeginPut,EndPut} = 1$

- Fairness: What is the maximum number of times e can fire without f firing in between?

$$fairness_{e,f} = \max_{\sigma \in \mathcal{E}^*} \{|\sigma|_e : \exists \mathbf{i}, \mathbf{j} \in \mathcal{X}_{reach},\ \mathbf{i} \xrightarrow{\sigma} \mathbf{j} \wedge |\sigma|_f = 0\}$$

  An infinite value might indicate the possibility of starvation

  - $fairness_{Prun,BeginPut} = N_p$ $\quad$ $fairness_{Prun,Brun} = N_p + 1$
    $fairness_{Prun,Crun} = N_p + N_s + N_c$ $\quad$ $fairness_{BeginPut,EndPut} = 1$

Structural properties (hold for any initial state)

- Structural Boundedness: Given a place p, is p bounded for each finite initial state?

$$\forall \mathbf{x}_{init} \in \mathbb{N}^{|\mathcal{P}|},\ \exists b \in \mathbb{N},\ \forall \mathbf{i} \in \mathcal{X}_{reach},\ \mathbf{i}_p \leq b$$

- Conservativeness: Is there a vector $\mathbf{x}$ of non-negative weights such that the weighted sum of tokens in the places is constant?

$$\exists \mathbf{x} \in \mathbb{N}^{|\mathcal{P}|},\ \mathbf{x} \neq \mathbf{0},\ \exists c \in \mathbb{N},\ \forall \mathbf{i} \in \mathcal{X}_{reach},\ \sum_{p \in \mathcal{P}} \mathbf{x}_p \mathbf{i}_p = c = \sum_{p \in \mathcal{P}} \mathbf{x}_p \mathbf{x}_{init\,p}$$

(the value of c is a behavioral property, the value of $\mathbf{x}$ is a structural property)

A place p is structurally bounded if $\mathbf{x}_p > 0$

- $\mathbf{i}_{Plocal} + \mathbf{i}_{Pwait} + \mathbf{i}_{PB} = N_p \qquad \mathbf{i}_{Empty} + \mathbf{i}_{Full} + \mathbf{i}_{PB} + \mathbf{i}_{CB} = N_s$

Structural properties (cont.)

- Repetitiveness: Is there an initial state $\mathbf{x}_{init}$ with an infinite firing sequence σ fireable in it?

$$\exists \mathbf{x}_{init} \in \mathbb{N}^{|\mathcal{P}|},\ \exists \sigma \in \mathcal{E}^{\omega},\ \mathbf{x}_{init} \xrightarrow{\sigma}$$

If an infinite sequence is fireable, then at least a subset of events is "self supporting" (they generate, as a whole, at least as many tokens as they consume)

- (Partial) Consistency: Is there a vector $\mathbf{y}$ of (non-negative) positive firing counts, such that the firing of any sequence σ where event e appears $\mathbf{y}_e$ times returns to the same state?

$$\exists \mathbf{y} \in \mathbb{N}^{|\mathcal{E}|},\ \mathbf{y} \neq \mathbf{0},\ \forall \mathbf{i} \in \mathbb{N}^{|\mathcal{P}|},\ \forall \sigma \in \mathcal{E}^{*},\ (\forall e \in \mathcal{E},\ |\sigma|_e = \mathbf{y}_e\ \Rightarrow\ (\mathbf{i} \xrightarrow{\sigma}\ \Rightarrow\ \mathbf{i} \xrightarrow{\sigma} \mathbf{i}))$$

  - $\mathbf{y}_{Prun} = 1$ $\quad \mathbf{y}_{BeginPut} = 1$ $\quad \mathbf{y}_{EndPut} = 1$ $\quad \mathbf{y}_{Brun} = 2$
    $\mathbf{y}_{Crun} = 1$ $\quad \mathbf{y}_{BeginGet} = 1$ $\quad \mathbf{y}_{EndGet} = 1$

Analysis of Petri nets

To answer questions about Petri net properties, two main methods exist:

- Coverability tree and reachability graph
  - Reachability: compute the set of reachable states
  - Coverbaility: if the set is infinite, compute a finite overapproximation of the reachable states
  - Studies behavioral properties
  - Usually expensive: the reachability set $\mathcal{X}_{reach}$ grows combinatorially

- Incidence matrix
  - Computes relations among places and events, enforced by the structure (arcs) of the Petri net
  - Studies structural properties
  - Usually fast: the incidence matrix $\mathbf{D}$ has size $|\mathcal{P}| \times |\mathcal{E}|$

Reachability and coverability

If the state space is bounded, we can build the reachability graph

If $\mathbf{x}_{init} = (N, 0, 0, 0, 0)$, the total number of states is $(N+1)(N+2)(2N+3)/6$

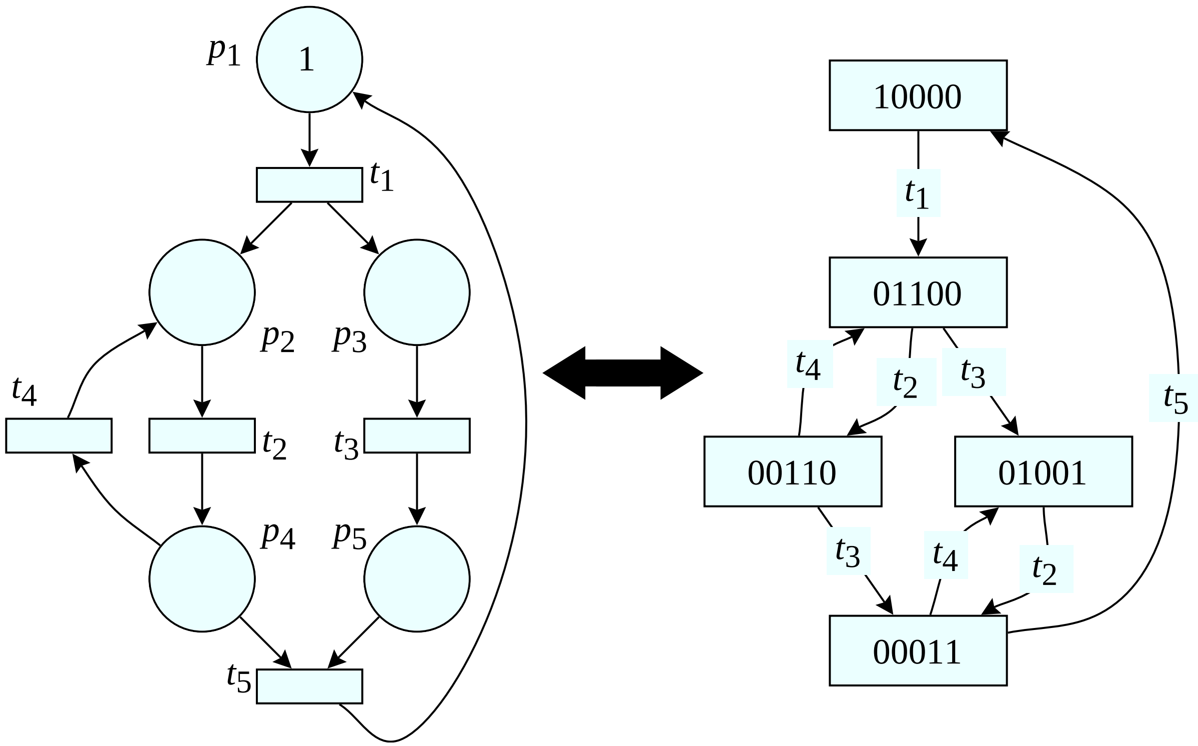

If the state space may be unbounded, we can generate the coverability tree or the coverability graph

We write $\mathbf{i}_p = \omega$ to mean that the number of tokens in place p may grow without bounds:

- $\mathbf{D}^-_{p,e} < \omega$: an input arc from p is always satisfied
- $\omega + \mathbf{D}_{p,e} = \omega$: after the firing of an event e, the number of tokens in p remains unbounded

Coverability graph with an infinite reachability set

The reachability graph is not completely described by the coverability graph

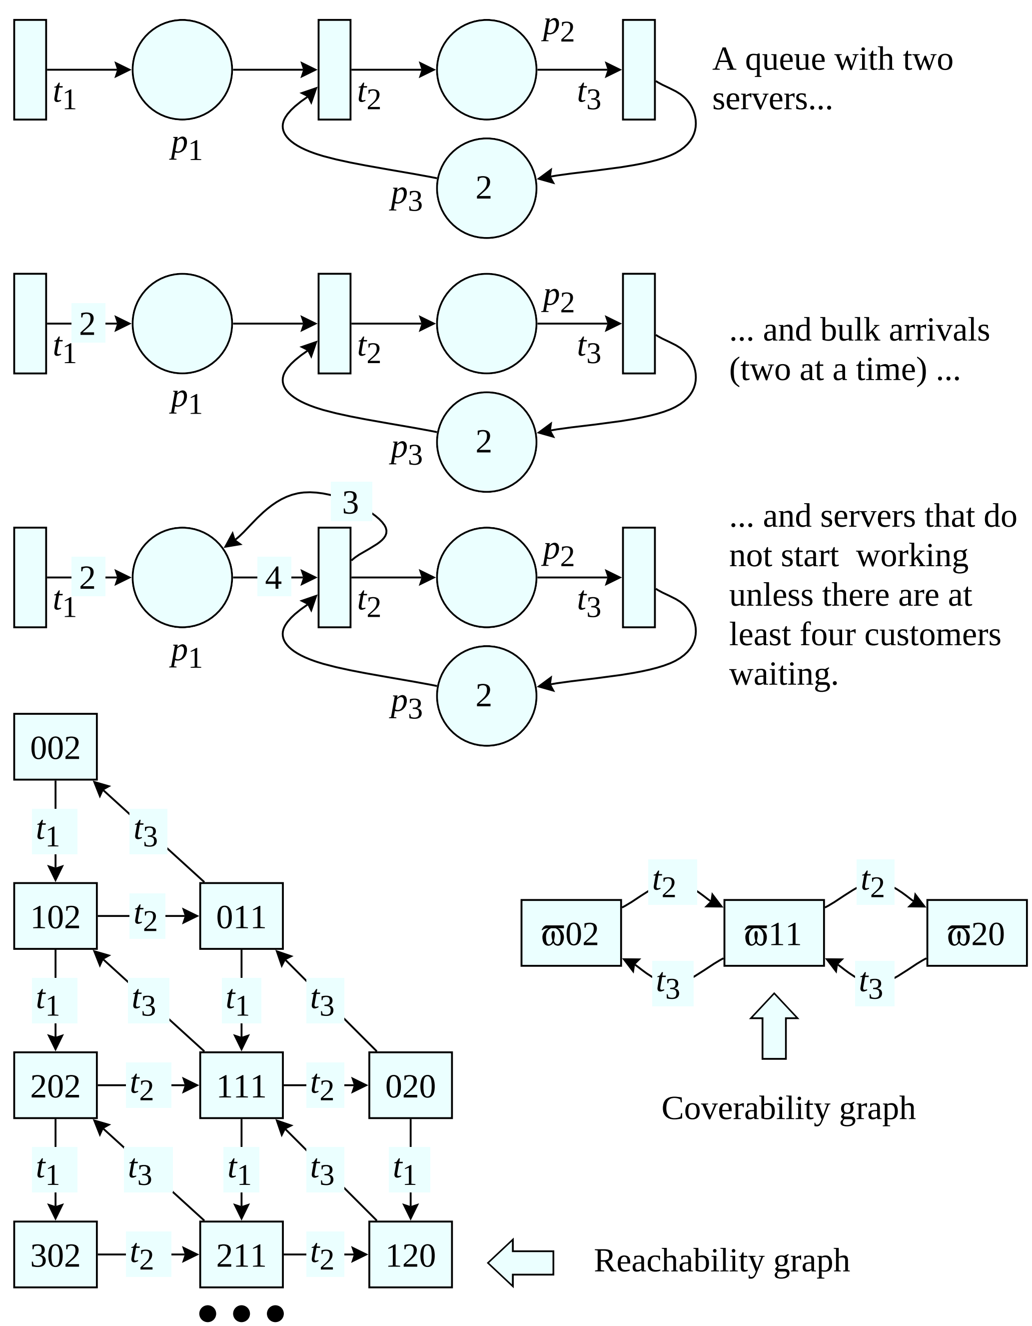

Incidence matrix: t-invariants

$\mathbf{i}$ is reachable from $\mathbf{x}_{init}$ iff there is a sequence $\sigma \in \mathcal{E}^*$ such that $\mathbf{x}_{init} \xrightarrow{\sigma} \mathbf{i}$

If σ consist of a single event e, the net effect on $\mathbf{x}_{init}$ when firing σ is $\boldsymbol{\delta} = \mathbf{D}_{\bullet,e} = \mathbf{D} \cdot \mathbf{1}(e)$ where $\mathbf{1}(e)$ is the column vector of size $|\mathcal{E}|$ having a 1 in position e and 0 elsewhere:

$$\mathbf{i} = \mathbf{x}_{init} + \boldsymbol{\delta} = \mathbf{x}_{init} + \mathbf{D} \cdot \mathbf{1}(e)$$

In general, the net effect of firing the sequence σ is $\boldsymbol{\delta} = \mathbf{D} \cdot (|\sigma|_1, \ldots |\sigma|_{|\mathcal{E}|})^T$:

$$\mathbf{i} = \mathbf{x}_{init} + \boldsymbol{\delta} = \mathbf{x}_{init} + \mathbf{D} \cdot (|\sigma|_1, |\sigma|_2, \ldots |\sigma|_{|\mathcal{E}|})^T$$

where $(|\sigma|_1, |\sigma|_2, \ldots |\sigma|_{|\mathcal{E}|})$ is the firing count vector corresponding to the firing sequence σ

A necessary condition for $\mathbf{i}$ to be reachable is then that the equation

$$\boldsymbol{\delta} = \mathbf{D} \cdot \mathbf{y}$$

has a non-negative integer solution $\mathbf{y} \in \mathbb{N}^{|\mathcal{E}|}$

The condition is not sufficient: we have no way to ensure that σ can be fired starting from $\mathbf{x}_{init}$ ($\mathbf{y}$ does not even tell us the order in which the $|\sigma|$ events should be fired)

Incidence matrix: t-invariants (cont.)

A non-negative, non-zero integer solution $\mathbf{y} \in \mathbb{N}^{|\mathcal{E}|}$ to

$$\mathbf{D} \cdot \mathbf{y} = \mathbf{0}$$

is called a t-semiflow (or t-invariant)

Any sequence σ satisfying $\forall e \in \mathcal{E},\ |\sigma|_e = \mathbf{y}_e$, if fireable, will not change the state:

$$\forall \mathbf{i} \in \mathbb{N}^{|\mathcal{P}|},\ \forall \sigma \in \mathcal{E}^*,\ (\forall e \in \mathcal{E},\ |\sigma|_e = \mathbf{y}_e \ \Rightarrow\ (\mathbf{i} \xrightarrow{\sigma}\ \Rightarrow\ \mathbf{i} \xrightarrow{\sigma} \mathbf{i}))$$

If $\mathbf{y}$ exists, the net is consistent: for a large enough $\mathbf{x}_{init}$, there are infinite firing sequences

- at least one producer, one consumer, one buffer, and one slot

Incidence matrix: p-invariants

The incidence matrix $\mathbf{D}$ can give information about place properties as well

Place p is structurally bounded if $\exists \mathbf{x} \in \mathbb{N}^{|\mathcal{P}|},\ \mathbf{x}_p > 0,\ \ \mathbf{x}^T \cdot \mathbf{D} \leq \mathbf{0}$

Proof:

$$\forall \mathbf{i} \in \mathcal{X}_{reach},\ \exists \mathbf{y} \in \mathbb{N}^{|\mathcal{E}|},\ \mathbf{i} = \mathbf{x}_{init} + \mathbf{D} \cdot \mathbf{y}$$

($\mathbf{y}$ is the firing count vector required to reach $\mathbf{i}$ from $\mathbf{x}_{init}$)

Then

$$\mathbf{x}^T \cdot \mathbf{i} = \mathbf{x}^T \cdot \mathbf{x}_{init} + \underbrace{\mathbf{x}^T \cdot \mathbf{D}}_{\leq \mathbf{0}} \cdot \mathbf{y}$$

which implies

$$\mathbf{x}^T \cdot \mathbf{i} \leq \mathbf{x}^T \cdot \mathbf{x}_{init}$$

and

$$\mathbf{i}_p \leq (\mathbf{x}^T \cdot \mathbf{x}_{init}) / \mathbf{x}_p$$

Incidence matrix: p-invariants (cont.)

A non-negative, non-zero integer solution $\mathbf{x} \in \mathbb{N}^{|\mathcal{P}|}$ to

$$\mathbf{x}^T \cdot \mathbf{D} = \mathbf{0}$$

is called a p-semiflow (or p-invariant)

If there exists a p-semiflow $\mathbf{x}$ such that $\mathbf{x}_p \neq 0$, then place p is bounded regardless of the initial state

If $\mathbf{x}$ and $\mathbf{x}'$ are p-semiflows, any nonzero integer linear combination $c \cdot \mathbf{x} + c' \cdot \mathbf{x}'$ is a p-semiflow

Let the support of a p-semiflow $\mathbf{x}$ be $Supp(\mathbf{x}) = \{p \in \mathcal{P} : \mathbf{x}_p > 0\}$, then a $\mathbf{x}$ is minimal if

- it has minimal support, i.e., no other p-semiflow $\mathbf{x}'$ exists with $Supp(\mathbf{x}') \subset Supp(\mathbf{x})$ and
- it is scaled back, i.e., there is no integer $c > 1$ such that $\mathbf{x}/c$ is a p-semiflow

As the set of p-semiflows is infinite, we seek the generator set containing only the minimal p-semiflows

Methods to compute the generator set have been extensively studied

In the worst case, the generator set contains an exponential number of p-semiflows

Modeling power of Petri nets

Petri nets cannot perform zero testing on an unbounded place

Consider the mutual exclusion problem among N processes (readers or writers)

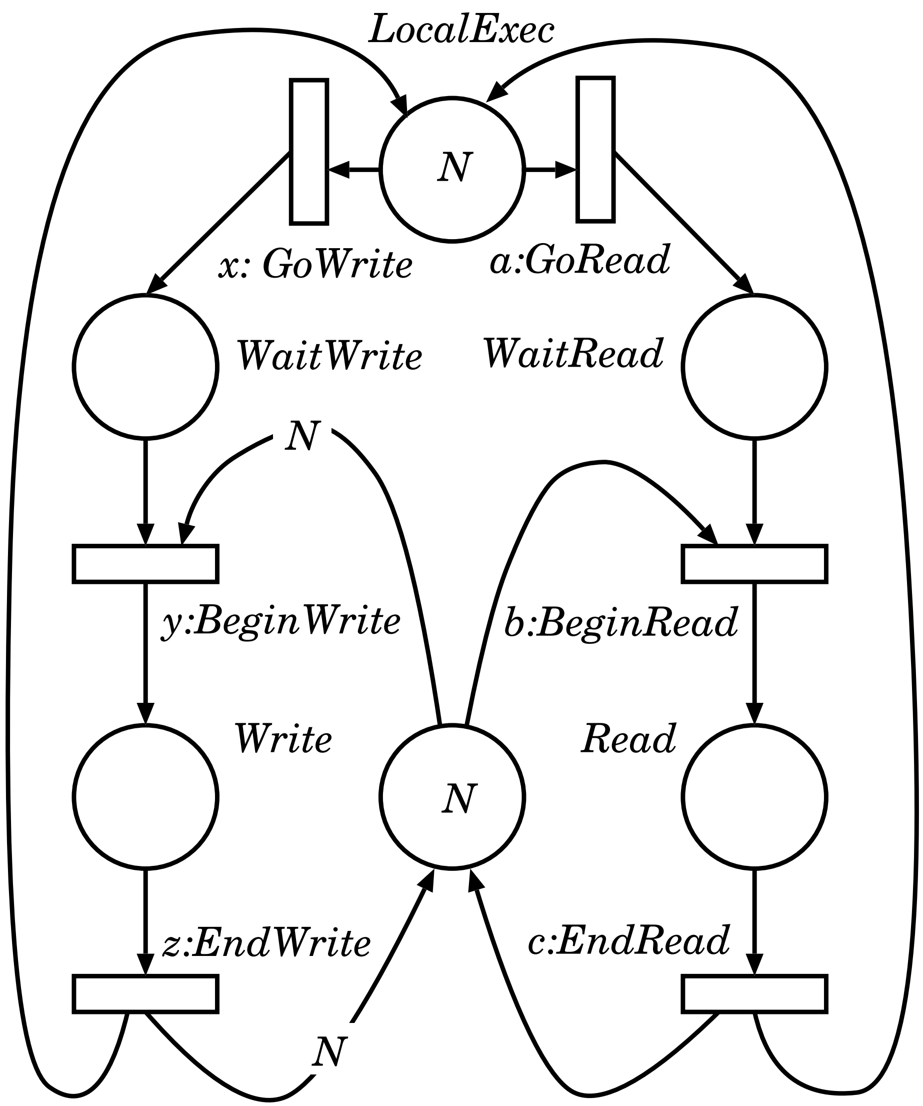

Modeling power of Petri nets (cont.)

Petri nets cannot model this system if the number N of processes is unbounded [Agerwala]

The proof uses the idea of Petri net languages and is a type of pumping lemma

- six actions: $GoRead$, $BeginRead$, and $EndRead$, $GoWrite$, $BeginWrite$, $EndWrite$
- associate the symbols $\{a, b, c, x, y, z\}$ to these six actions
- if multiple events are used to model the same action, label all of them with the same symbol
- any other event has an empty label

There is a firing sequence σ containing the infinite subsequence $bcyzbbccyz \cdots b^i c^i yz \cdots$

Consider $(\mathbf{s}^{[1]}, \mathbf{s}^{[2]}, \ldots, \mathbf{s}^{[i]}, \ldots)$, where $\mathbf{s}^{[i]}$ is the state reached after firing $bcyzbbccyz \cdots b^i$

We can prove by induction on the number of places, i.e., on the size of the state vector, that there are $i < j$ such that $(\mathbf{s}^{[1]}, \mathbf{s}^{[2]}, \ldots, \mathbf{s}^{[i]}, \ldots, \mathbf{s}^{[j]}, \ldots)$ satisfies $\mathbf{s}^{[i]} \leq \mathbf{s}^{[j]}$

Thus, we can reach $\mathbf{s}^{[j]}$ by firing $bcyzbbccyz \cdots b^j$, then fire $c^i yz$ in $\mathbf{s}^{[j]}$, because of monotonicity:

$$\mathbf{s}^{[i]} \stackrel{c^i yz}{\longrightarrow} \wedge \ \mathbf{s}^{[i]} \leq \mathbf{s}^{[j]} \ \Rightarrow \ \mathbf{s}^{[j]} \stackrel{c^i yz}{\longrightarrow}$$

But this means that we allow a writer in the critical section while $j - i > 0$ readers are still in it

Extending the modeling power of Petri nets

Solution: extend the modeling power of Petri nets by adding new constructs

- Inhibitor arc: disable an event if a place is not empty (opposite logic to that of an input arc)

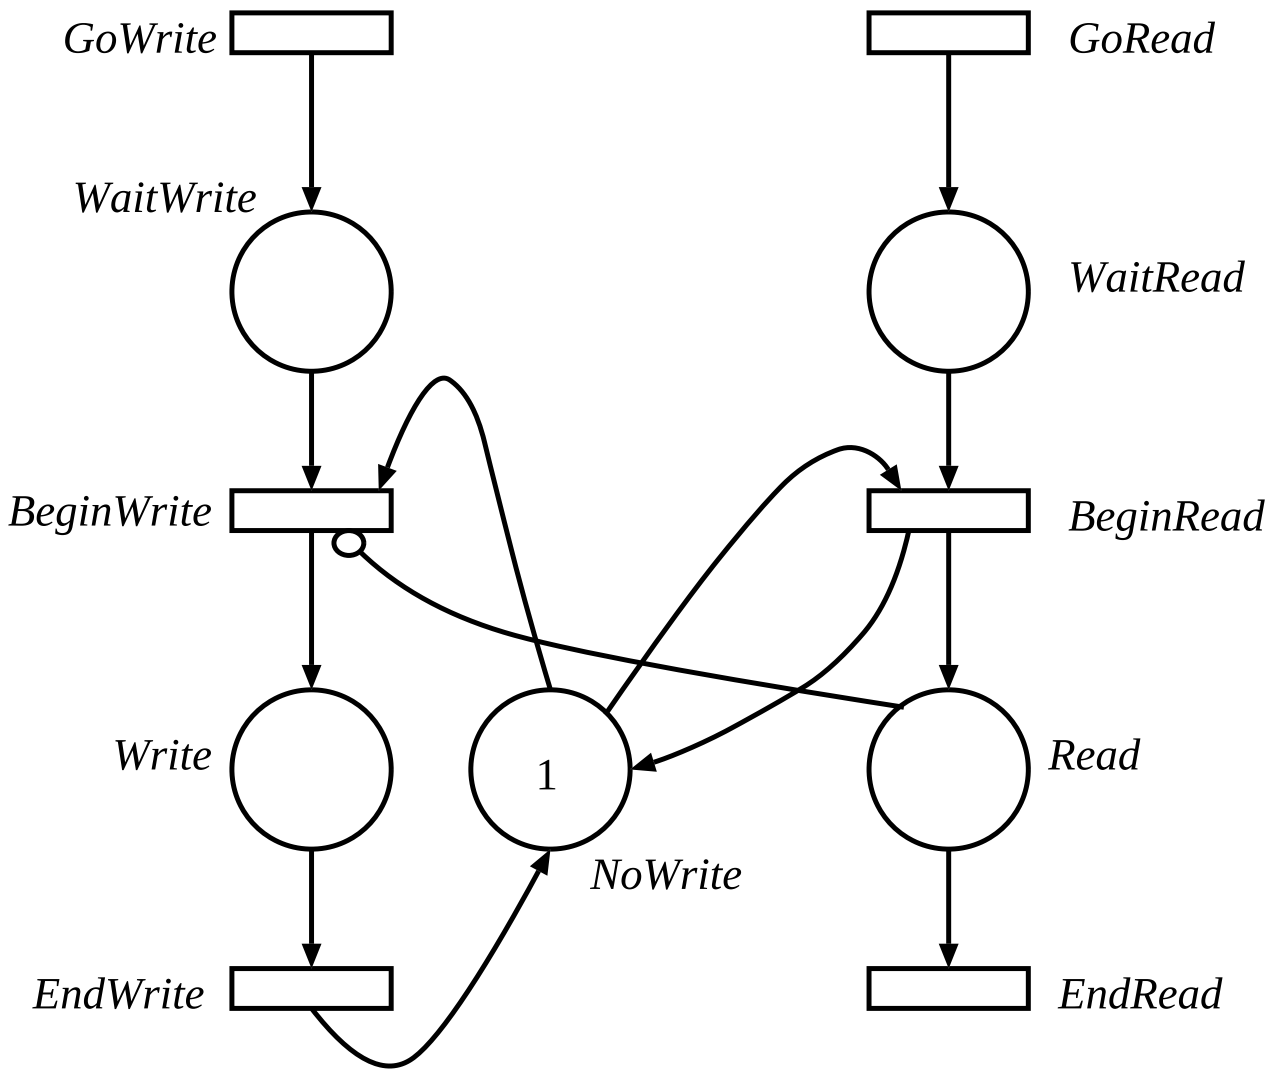

Extending the modeling power of Petri nets (cont.)

- Priority: disable an event if another enabled event has priority over it and is enabled

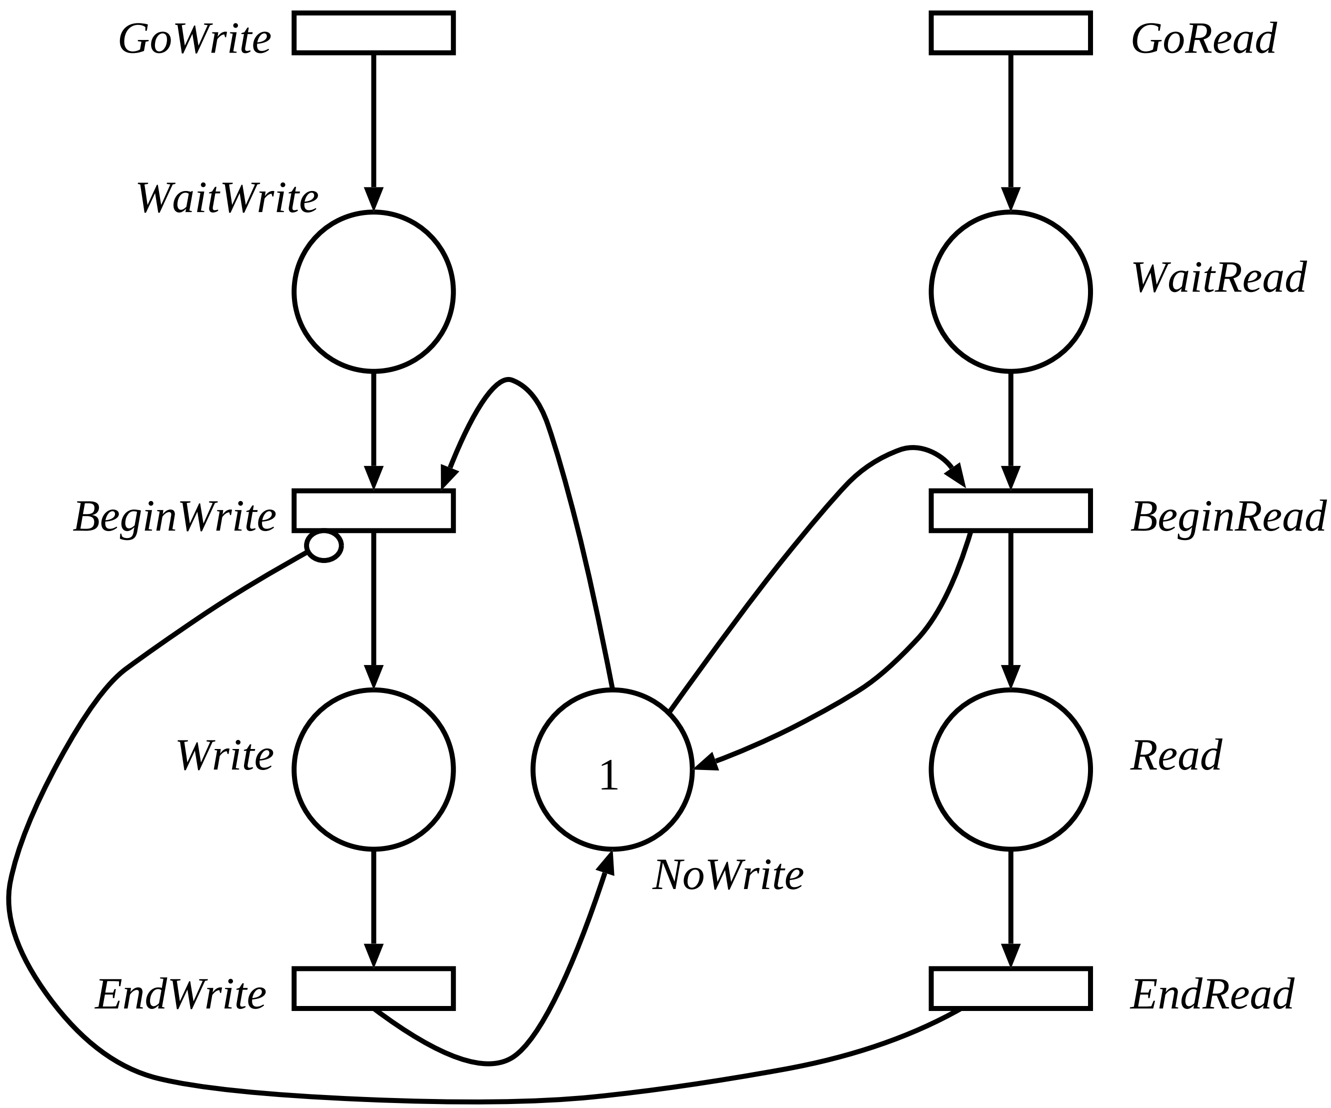

Often defined as a mapping $prio : \mathcal{E} \to \mathbb{N}$ — α has priority over β if $prio(\alpha) > prio(\beta)$

We prefer a partial order relation $\succ \subseteq \mathcal{E} \times \mathcal{E}$ — α has priority over β if $\alpha \succ \beta$

Extending the modeling power of Petri nets (cont.)

- Self-modifying nets: the cardinality of input and output arcs is a function of the state

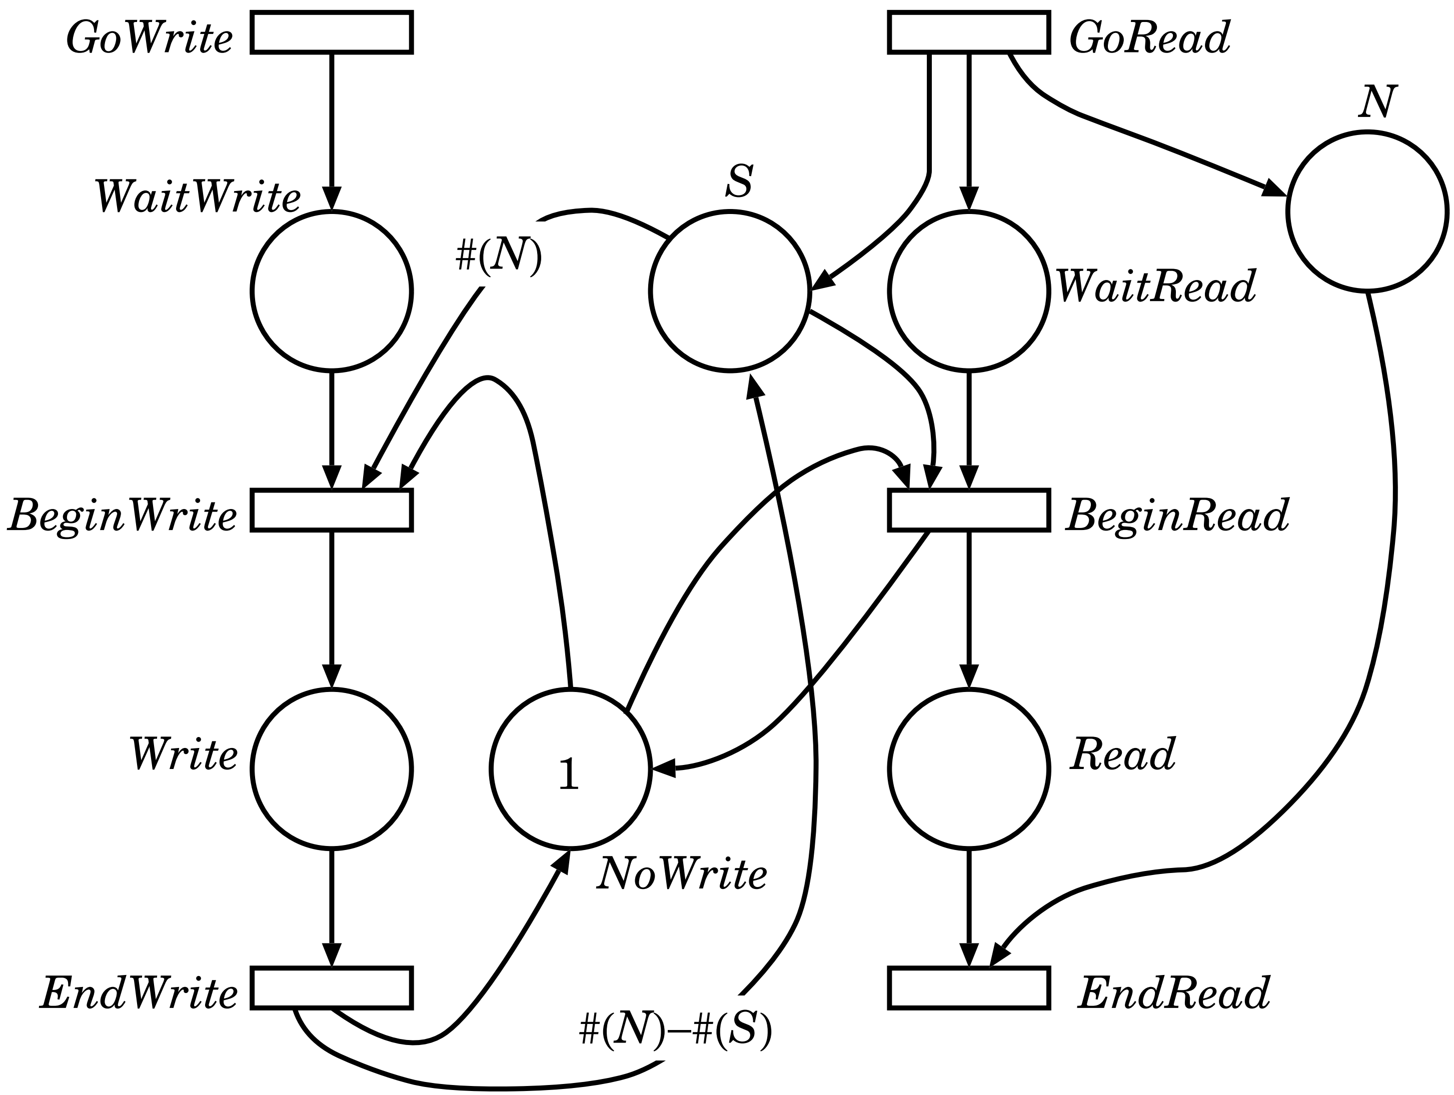

We can also call them Petri nets with marking-dependent arc cardinalities

Extended Petri nets

Any one of these extensions is enough to achieve full Turing-equivalence

- $\succ \subset \mathcal{E} \times \mathcal{E}$ a pre-selection priority (partial order) relation
  if $f \succ e$, event e is disabled whenever event f is enabled

- $\mathbf{D}^\circ : \mathcal{P} \times \mathcal{E} \to \mathbb{N} \cup \{\infty\}$ inhibitor arc cardinalities
  event e is disabled in state $\mathbf{i} \in \mathbb{N}^{|\mathcal{P}|}$ if $\mathbf{D}^\circ_{p,e} \leq \mathbf{i}_p$

- $G : \mathcal{E} \times \mathbb{N}^{|\mathcal{P}|} \to \{0, 1\}$ event guards
  event e is disabled in state $\mathbf{i} \in \mathbb{N}^{|\mathcal{P}|}$ if $G_e(\mathbf{i}) = 0$

Self-modifying nets may achieve Turing-equivalence, depending on the restrictions on arc cardinalities

Self-modifying Petri nets with inhibitor arcs

A self-modifying Petri net with inhibitor arcs is a tuple $(\mathcal{P}, \mathcal{E}, \mathbf{D}^-, \mathbf{D}^+, \mathbf{D}^\circ, \mathbf{i}^{init})$ where:

- $\mathcal{P}$ and $\mathcal{E}$ — places and transitions
- $\mathbf{D}^-, \mathbf{D}^+ : \mathcal{P} \times \mathcal{E} \times \mathbb{N}^{|\mathcal{P}|} \rightarrow \mathbb{N}$ — marking-dependent input, output arc cardinalities
- $\mathbf{D}^\circ : \mathcal{P} \times \mathcal{E} \times \mathbb{N}^{|\mathcal{P}|} \rightarrow \mathbb{N} \cup \{\infty\}$ — marking-dependent inhibitor arc cardinalities
- $\mathbf{i}^{init} : \mathbb{N}^{|\mathcal{P}|}$ — initial marking

Transition α is enabled in marking $\mathbf{i} \in \mathbb{N}^{|\mathcal{P}|}$ iff $\forall p \in \mathcal{P},\ \mathbf{D}^-_{p,\alpha}(\mathbf{i}) \leq i_p \ \wedge \ \mathbf{D}^\circ_{p,\alpha}(\mathbf{i}) > i_p$

If α is enabled in $\mathbf{i}$, it can $fire$ and lead to marking $\mathbf{j}$ $\forall p \in \mathcal{P},\ \ j_p = i_p - \mathbf{D}^-_{p,\alpha}(\mathbf{i}) + \mathbf{D}^+_{p,\alpha}(\mathbf{i})$

The effect of α is deterministic, so we can write $\mathbf{i} \xrightarrow{\alpha} \mathbf{j}$ or use the general notation $\mathbf{j} \in \mathcal{N}_\alpha(\mathbf{i})$

$$\widehat{\mathcal{X}} \equiv \mathbb{N}^{|\mathcal{P}|} \qquad \mathcal{X}_{init} \equiv \{\mathbf{i}^{init}\} \qquad \mathcal{N} \equiv \bigcup_{\alpha \in \mathcal{E}} \mathcal{N}_\alpha$$

Self-modifying Petri nets with inhibitor arcs, guards, priorities

A self-modifying Petri net with inhibitor arcs, guards, and priorities is a tuple

$$(\mathcal{P}, \mathcal{E}, \mathbf{D}^-, \mathbf{D}^+, \mathbf{D}^\circ, G, \succ, \mathbf{x}_{init})$$

- $\mathcal{P}$ and $\mathcal{E}$ — places and events
- $\mathbf{D}^-, \mathbf{D}^+ : \mathcal{E} \times \mathcal{P} \times \mathbb{N}^{|\mathcal{P}|} \rightarrow \mathbb{N}$ — state-dependent input, output arc cardinalities
- $\mathbf{D}^\circ : \mathcal{E} \times \mathcal{P} \times \mathbb{N}^{|\mathcal{P}|} \rightarrow \mathbb{N} \cup \{\infty\}$ — state-dependent inhibitor arc cardinalities
- $G : \mathcal{E} \times \mathbb{N}^{|\mathcal{P}|} \rightarrow \{true, false\}$ — state-dependent guards
- $\succ \ \subset \mathcal{E} \times \mathcal{E}$ — acyclic (preselection) priority relation
- $\mathbf{x}_{init} : \mathbb{N}^{|\mathcal{P}|}$ — initial state

Event α is enabled in a state $\mathbf{i} \in \mathbb{N}^{|\mathcal{P}|}$, written $\alpha \in \mathcal{E}(\mathbf{i})$, iff

$$\forall p \in \mathcal{P},\ \mathbf{D}^-_{\alpha,p}(\mathbf{i}) \leq i_p \ \wedge\ \mathbf{D}^\circ_{\alpha,p}(\mathbf{i}) > i_p \ \wedge G_\alpha(\mathbf{i}) \ \wedge\ \forall \beta \in \mathcal{E},\ \beta \succ \alpha \ \Rightarrow\ \beta \notin \mathcal{E}(\mathbf{i})$$

If $\mathbf{i} \xrightarrow{\alpha} \mathbf{j}$, the new state $\mathbf{j}$ satisfies $\forall p \in \mathcal{P},\ j_p = i_p - \mathbf{D}^-_{\alpha,p}(\mathbf{i}) + \mathbf{D}^+_{\alpha,p}(\mathbf{i})$ (deterministic effect)