

A Concurrency-Preserving Translation from Time Petri Nets to Networks of Timed Automata

Sandie Balaguer, Thomas Chatain, Stefan Haar

INRIA & LSV – ENS Cachan, CNRS

MeFoSyLoMa Seminar – May 6, 2011

- 1 Introduction
 - Motivation
 - Timed and concurrent models
- 2 Partial order semantics
 - Timed traces
 - Distributed timed language
- 3 Decomposing a PN in processes
 - S-invariants
 - Decomposition
- 4 Translation from TPN to NTA
 - Adding clocks
 - Know thy neighbour!
- 5 Conclusion

- 1 Introduction
 - Motivation
 - Timed and concurrent models
- 2 Partial order semantics
 - Timed traces
 - Distributed timed language
- 3 Decomposing a PN in processes
 - S-invariants
 - Decomposition
- 4 Translation from TPN to NTA
 - Adding clocks
 - Know thy neighbour!
- 5 Conclusion

Motivation

Concurrency

- Two actions that might be performed in **any order** leading to the **same state** are **concurrent**. Concurrency can be used to improve the analysis of distributed systems.
- The definition of concurrency in **timed systems** is not clear since events are ordered both by their occurrence dates and by causality.

2 formalisms

- Networks of timed automata (NTA)
- Time Petri nets (TPN)

Translation between formalisms

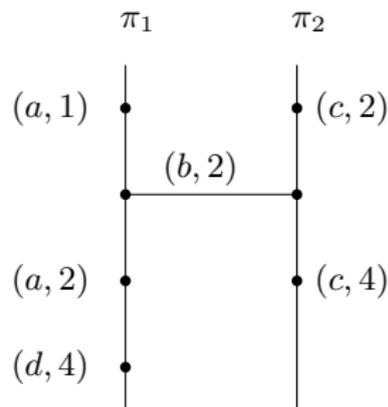
- Theoretical reasons (comparison)
- Practical reasons (verification tools)

Motivation

- Translations from TPN to NTA with preservation of timed words but **loss of concurrency**

Concurrency-preserving translation

- Runs are represented as timed traces \neq timed words. The translation preserves timed traces.
- Some hidden dependencies caused by time are made explicit.

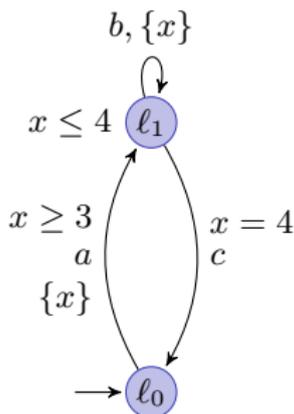


Timed Automata [Alur, Dill, 94]

Definition (Timed Automaton)

A timed automaton is a tuple $\mathcal{A} = (L, \ell_0, C, \Sigma, E, Inv)$ where:

- L is a set of locations,
- $\ell_0 \in L$ is the initial location,
- C is a finite set of clocks,
- Σ is a finite set of actions,
- $E \subseteq L \times \mathcal{B}(C) \times \Sigma \times 2^C \times L$ is a set of edges,
- $Inv : L \rightarrow \mathcal{B}(C)$ assigns invariants to locations.



- A location must be leaved when its **invariant** reaches its limit.
- An edge cannot be taken if its **guard** is not satisfied.

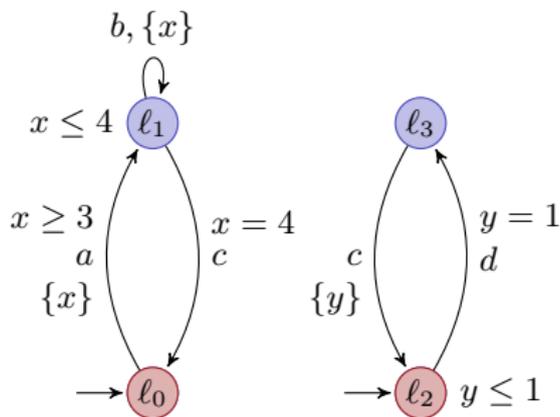
Networks of Timed Automata: $\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n$

Action step: $(\vec{\ell}, v) \xrightarrow{a} (\vec{\ell}', v')$

- If all the automata that share a are ready to perform it.
- Edges labeled by a are taken simultaneously in these automata.

Delay step: $\forall d \in \mathbb{R}_{\geq 0}, (\vec{\ell}, v) \xrightarrow{d} (\vec{\ell}, v + d)$

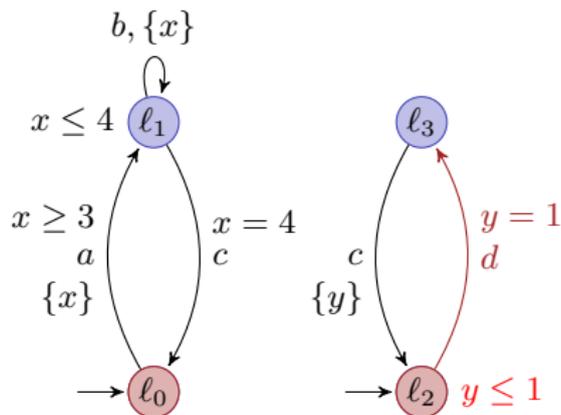
- $v + d$ respects the invariants of the current locations.



$$(\ell_0, \ell_2) \xrightarrow{1} (0, 0)$$

Networks of Timed Automata: $\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n$

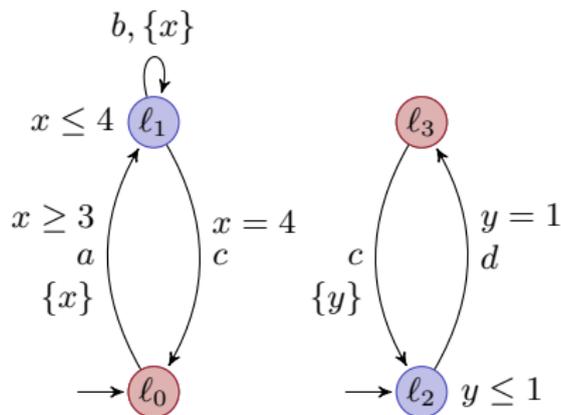
Example run



$$\begin{array}{c} (\ell_0, \ell_2) \\ (0, 0) \end{array} \xrightarrow{1} \begin{array}{c} (\ell_0, \ell_2) \\ (1, \mathbf{1}) \end{array} \xrightarrow{d}$$

Networks of Timed Automata: $\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n$


Example run



$$\begin{array}{c}
 (\ell_0, \ell_2) \\
 (0, 0)
 \end{array}
 \xrightarrow{1}
 \begin{array}{c}
 (\ell_0, \ell_2) \\
 (1, 1)
 \end{array}
 \xrightarrow{d}
 \begin{array}{c}
 (\ell_0, \ell_3) \\
 (1, 1)
 \end{array}
 \xrightarrow{2.5}$$

Networks of Timed Automata: $\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n$


Example run



$$\begin{array}{ccccccc}
 (\ell_0, \ell_2) & \xrightarrow{1} & (\ell_0, \ell_2) & \xrightarrow{d} & (\ell_0, \ell_3) & \xrightarrow{2.5} & (\ell_0, \ell_3) & \xrightarrow{a} \\
 (0, 0) & & (1, 1) & & (1, 1) & & (3.5, 3.5) &
 \end{array}$$

Networks of Timed Automata: $\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n$

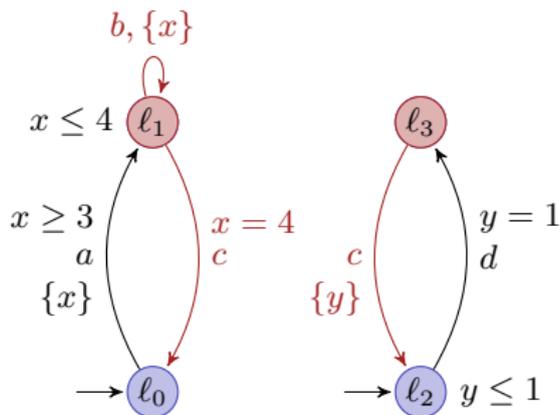
Example run



$$\begin{array}{ccccccc}
 (\ell_0, \ell_2) & \xrightarrow{1} & (\ell_0, \ell_2) & \xrightarrow{d} & (\ell_0, \ell_3) & \xrightarrow{2.5} & (\ell_0, \ell_3) & \xrightarrow{a} & (\ell_1, \ell_3) & \xrightarrow{4} & \\
 (0, 0) & & (1, 1) & & (1, 1) & & (3.5, 3.5) & & (0, 3.5) & &
 \end{array}$$

Networks of Timed Automata: $\mathcal{A}_1 \parallel \dots \parallel \mathcal{A}_n$

Example run

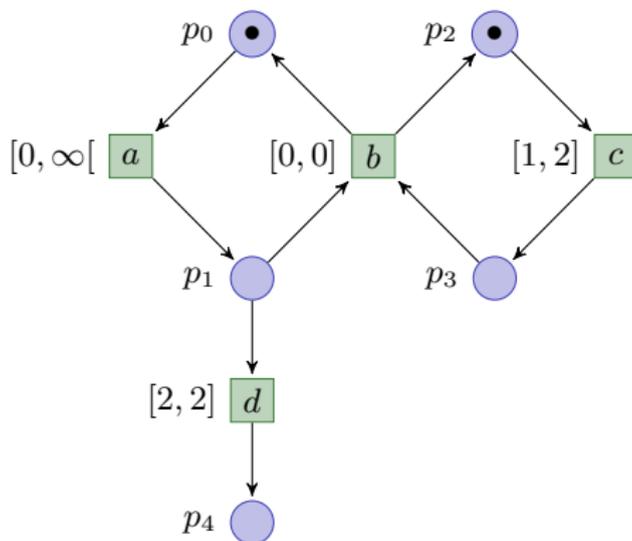


$$\begin{array}{ccccccc}
 (\ell_0, \ell_2) & \xrightarrow{1} & (\ell_0, \ell_2) & \xrightarrow{d} & (\ell_0, \ell_3) & \xrightarrow{2.5} & (\ell_0, \ell_3) & \xrightarrow{a} & (\ell_1, \ell_3) & \xrightarrow{4} & (\ell_1, \ell_3) & \xrightarrow{c} & \dots \\
 (0, 0) & & (1, 1) & & (1, 1) & & (3.5, 3.5) & & (0, 3.5) & & (4, 7.5) & &
 \end{array}$$

Time Petri Nets [Merlin, 74]

(P, T, F, M_0, efd, lfd)

- $efd : T \rightarrow \mathbb{R}$ earliest firing delay
- $lfd : T \rightarrow \mathbb{R} \cup \{\infty\}$ latest firing delay



TPN Semantics

- t is enabled in M : $t \in \text{enabled}(M) \Leftrightarrow \bullet t \subseteq M$
- firing t from M : $M \xrightarrow{t} (M' = M - \bullet t + t \bullet)$
- t' is newly enabled by the firing of t from M : intermediate semantics
 $\uparrow \text{enabled}(t', M, t) = (t' \in \text{enabled}(M')) \wedge (t' \notin \text{enabled}(M - \bullet t))$

Discrete transition: $\forall t \in \text{enabled}(M), (M, \nu) \xrightarrow{t} (M', \nu')$ iff

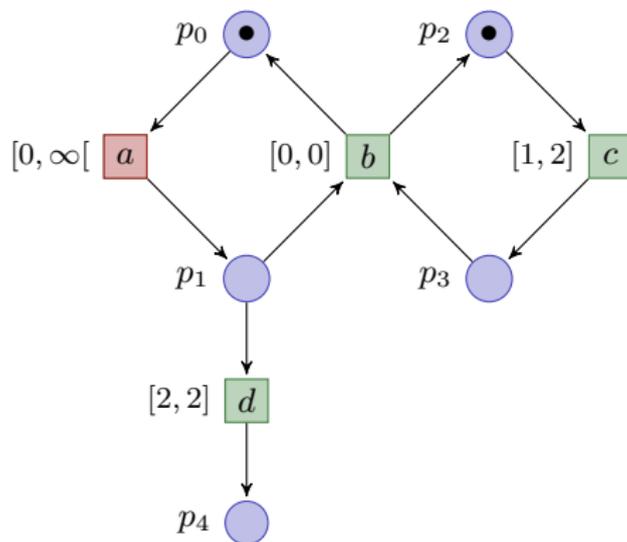
- $\text{efd}(t) \leq \nu(t)$,
- $\forall t' \in T, \nu'(t') = \begin{cases} 0 & \text{if } \uparrow \text{enabled}(t', M, t) \\ \nu(t') & \text{otherwise.} \end{cases}$

Continuous transition: $\forall d \in \mathbb{R}_{\geq 0}, (M, \nu) \xrightarrow{d} (M, \nu + d)$ iff

- $\forall t \in \text{enabled}(M), \nu(t) + d \leq \text{lfd}(t)$ urgency

TPN Semantics

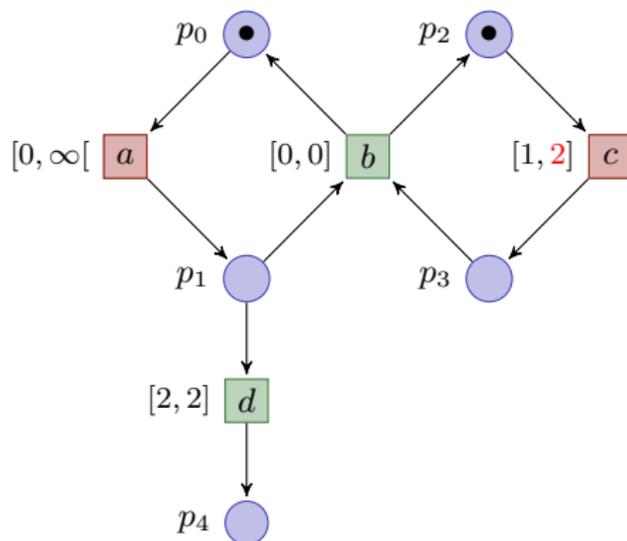
Example run



$$\{p_0, p_2\} \xrightarrow{2} (0, -, 0, -)$$

TPN Semantics

Example run




$$\{p_0, p_2\} \xrightarrow{a} \{p_0, p_2\} \xrightarrow{c}$$

$$(0, -, 0, -) \xrightarrow{a} (2, -, 2, -)$$

TPN Semantics

Example run




$$\{p_0, p_2\} \xrightarrow{2} \{p_0, p_2\} \xrightarrow{c} \{p_0, p_3\} \xrightarrow{10}$$

$$(0, -, 0, -) \xrightarrow{2} (2, -, 2, -) \xrightarrow{c} (2, -, -, -)$$

TPN Semantics

Example run

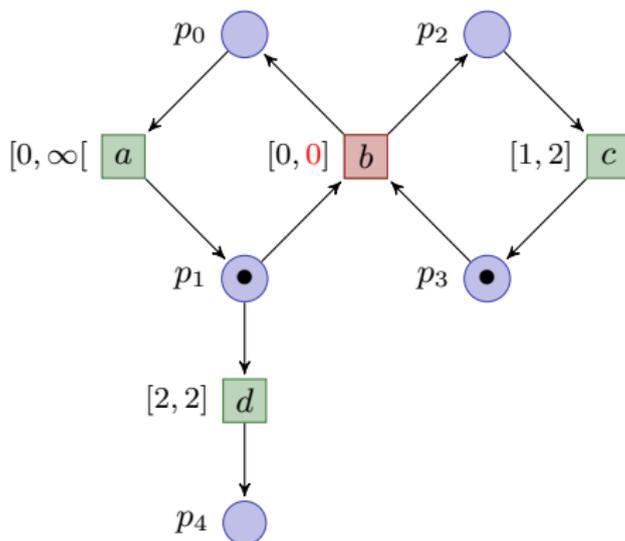


$$\{p_0, p_2\} \xrightarrow{2} \{p_0, p_2\} \xrightarrow{c} \{p_0, p_3\} \xrightarrow{10} \{p_0, p_3\} \xrightarrow{a}$$

$$(0, -, 0, -) \xrightarrow{2} (2, -, 2, -) \xrightarrow{c} (2, -, -, -) \xrightarrow{10} (12, -, -, -)$$

TPN Semantics

Example run



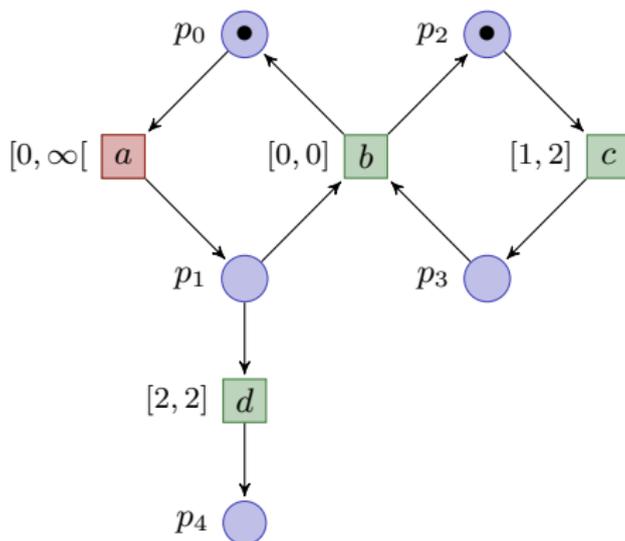
$$\{p_0, p_2\} \xrightarrow{2} \{p_0, p_2\} \xrightarrow{c} \{p_0, p_3\} \xrightarrow{10} \{p_0, p_3\} \xrightarrow{a} \{p_1, p_3\} \xrightarrow{b}$$

$$(0, -, 0, -) \longrightarrow (2, -, 2, -) \longrightarrow (2, -, -, -) \longrightarrow (12, -, -, -) \longrightarrow (-, 0, -, 0)$$

b and d are newly enabled.

TPN Semantics

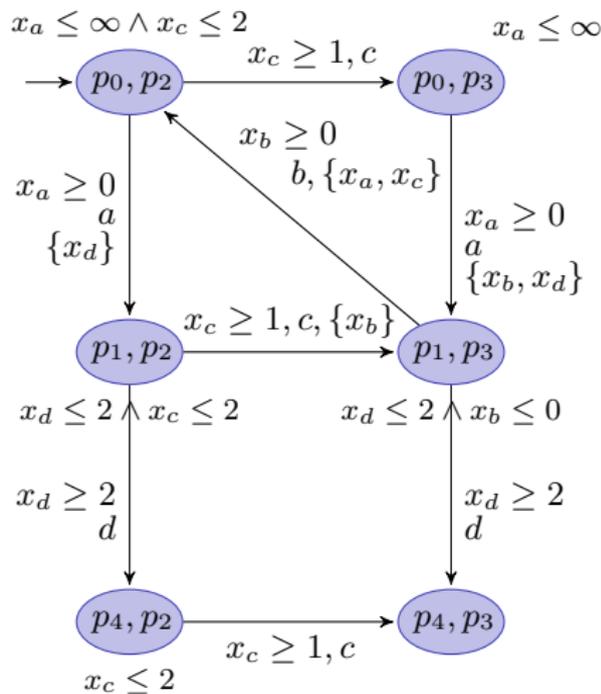
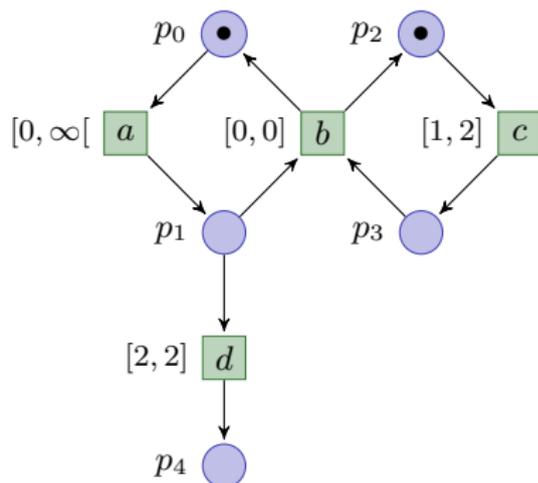
Example run



$$\begin{array}{ccccccc}
 \{p_0, p_2\} & \xrightarrow{2} & \{p_0, p_2\} & \xrightarrow{c} & \{p_0, p_3\} & \xrightarrow{10} & \{p_0, p_3\} & \xrightarrow{a} & \{p_1, p_3\} & \xrightarrow{b} & \{p_0, p_2\} \\
 (0, -, 0, -) & & (2, -, 2, -) & & (2, -, -, -) & & (12, -, -, -) & & (-, 0, -, 0) & & (0, -, 0, -)
 \end{array}$$



TPN Semantics

Can be seen as a TA



TPN Semantics

Can be seen as a TA



- 1 Introduction
 - Motivation
 - Timed and concurrent models
- 2 Partial order semantics
 - Timed traces
 - Distributed timed language
- 3 Decomposing a PN in processes
 - S-invariants
 - Decomposition
- 4 Translation from TPN to NTA
 - Adding clocks
 - Know thy neighbour!
- 5 Conclusion

Partial order semantics for distributed systems

NTA and TPN represent distributed systems

- Composition of several (physical) components
- Notion of **process**
 - In a NTA, each automaton is a process.
 - PNs usually built as products of transition systems

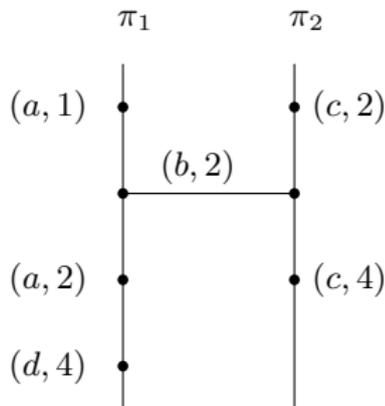
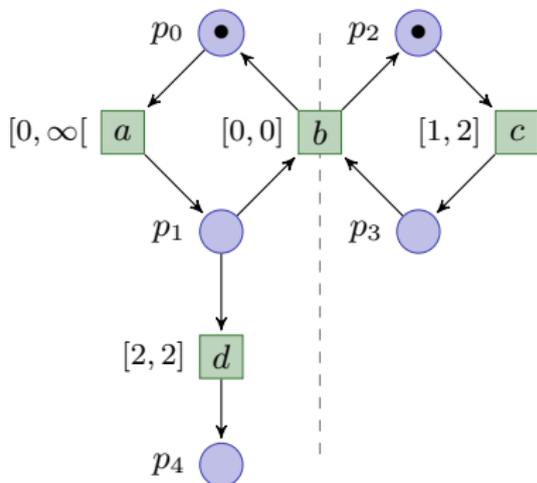
Usual semantics as timed words does not reflect the distribution of actions.

Partial order semantics reflects the distribution of actions.

Timed traces

A **timed trace** over the alphabet Σ , and the set of **processes** $\Pi = (\pi_1, \dots, \pi_n)$ is a tuple $\mathcal{W} = (E, \preceq, \lambda, t, \text{proc})$ where:

- E is a set of events,
- $\preceq \subseteq (E \times E)$ is a **partial order** on E ($\preceq|_{\pi_i}$ is a total order),
- $\lambda : E \rightarrow \Sigma$ is a labeling function,
- $t : E \rightarrow \mathbb{R}_{\geq 0}$ maps each event to a date,
- $\text{proc} : \Sigma \rightarrow 2^{\Pi}$ is a distribution of actions.



Distributed timed language

Definition (Distributed timed language)

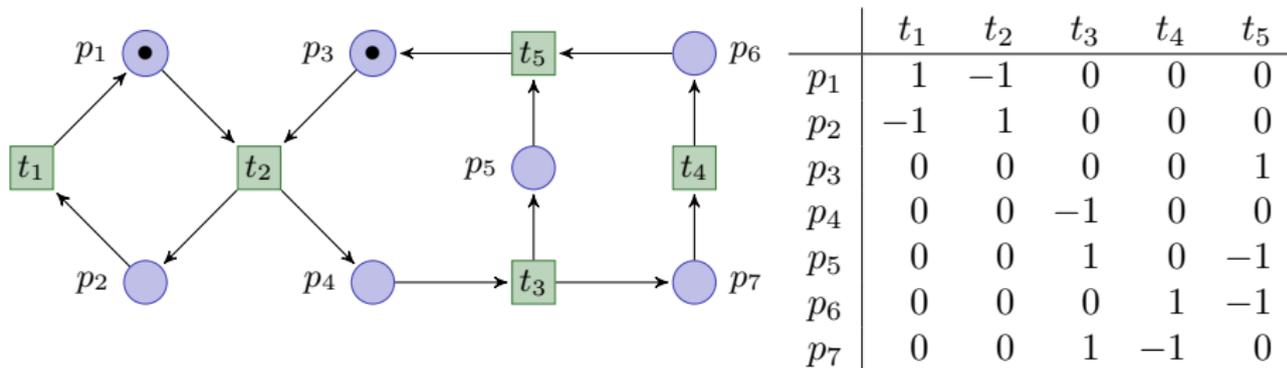
A **distributed timed language** is a set of timed traces.

- A **timed trace** is defined by a **timed word** and a **distribution of actions** ($proc : \Sigma \rightarrow 2^{\Pi}$).
- A **distributed timed language** is defined by a **timed language** and a **distribution of actions**.

- 1 Introduction
 - Motivation
 - Timed and concurrent models
- 2 Partial order semantics
 - Timed traces
 - Distributed timed language
- 3 **Decomposing a PN in processes**
 - **S-invariants**
 - **Decomposition**
- 4 Translation from TPN to NTA
 - Adding clocks
 - Know thy neighbour!
- 5 Conclusion

S-invariants [Lautenbach, 75], [Reisig, 85], [Desel, Esparza, 95]...

$X : P \rightarrow \mathbb{N}$, solution of the equation $X \cdot \mathbf{N} = \mathbf{0}$, where \mathbf{N} is the incidence matrix.



S-invariants [Lautenbach, 75], [Reisig, 85], [Desel, Esparza, 95]...

$X : P \rightarrow \mathbb{N}$, solution of the equation $X \cdot \mathbf{N} = \mathbf{0}$, where \mathbf{N} is the incidence matrix.

We consider S-invariants X s.t. $X : P \rightarrow \{0, 1\}$ (subsets of places).

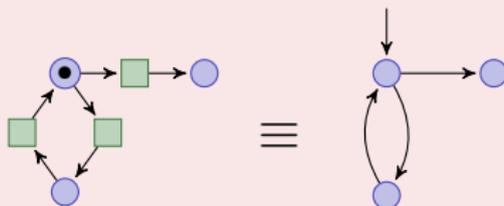
Properties

- X is an S-invariant $\Leftrightarrow \forall t \in T, \sum_{p \in \bullet t} X(p) = \sum_{p \in t \bullet} X(p)$ i.e. $|\bullet t \cap X| = |t \bullet \cap X|$
- X is an S-invariant $\Rightarrow \forall M, X \cdot M = X \cdot M_0$ i.e. $\sum_{p \in X} M(p) = \sum_{p \in X} M_0(p)$

S-invariants as processes

- A net (P, T, F) is an **S-net** if $\forall t \in T, |\bullet t| = |t \bullet| = 1$.

An S-net with one token can be seen as an automaton.



- The subnet (P', T', F') of N is a **P-closed** subnet of N if $T' = \bullet P' \cup P' \bullet$.

Definition

The net $N = (P, T, F)$ is **decomposable** iff there exists a set of P-closed S-nets $N_i = (P_i, T_i, F_i)$ that **covers** N .

[Desel, Esparza, 95] Well-formed free-choice nets are covered by strongly connected P-closed S-nets (S-components).

Decomposition

Proposition

A Petri net (P, T, F) is decomposable in the subnets N_1, \dots, N_n iff there exists a set of S-invariants $\{X_1, \dots, X_n\}$ such that,

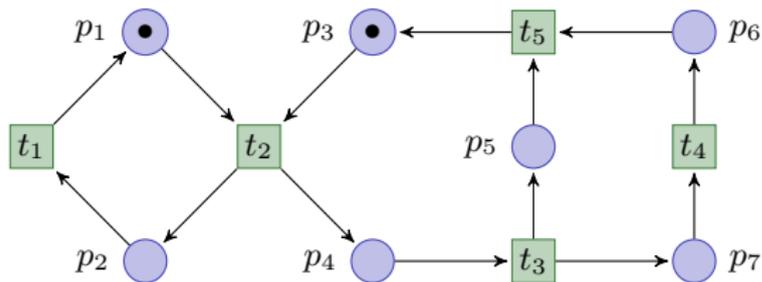
- $\forall i \in [1..n], X_i : P \rightarrow \{0, 1\}$,
 X_i is the characteristic function of P_i over P .
- $\forall i \in [1..n], \forall t \in T, \sum_{p \in \bullet t} X_i(p) = 1$ ($= \sum_{p \in t^\bullet} X_i(p)$),
 N_i is an S-net.
- $\forall p \in P, \sum_i X_i(p) \geq 1$
 Each place is in at least one component.

The processes are the P-closed subnets spanned by the supports of these S-invariants.

Since each place is in at least one subnet and the subnets are P-closed, each transition is also in at least one subnet and the net is covered.

Decomposition

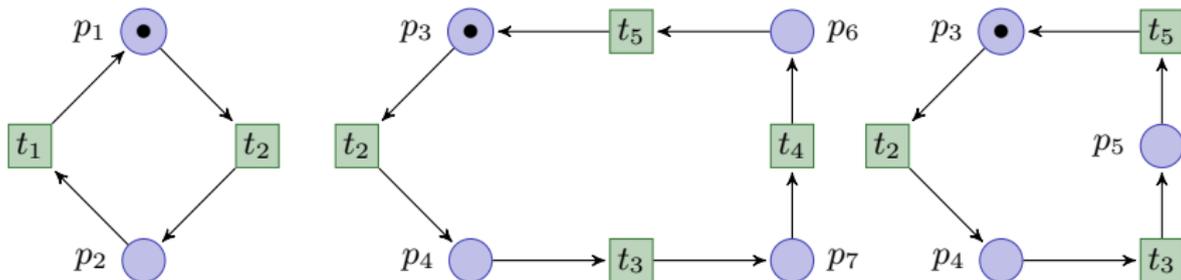
An example



$$X \cdot \mathbf{N} = 0$$

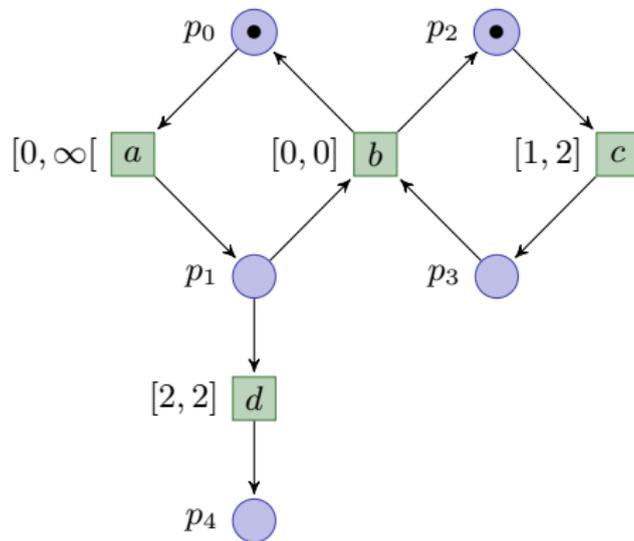
$$X_1 = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 \end{bmatrix}$$

p_1 p_2 p_3 p_4 p_5 p_6 p_7





- 1 Introduction
 - Motivation
 - Timed and concurrent models
- 2 Partial order semantics
 - Timed traces
 - Distributed timed language
- 3 Decomposing a PN in processes
 - S-invariants
 - Decomposition
- 4 Translation from TPN to NTA
 - Adding clocks
 - Know thy neighbour!
- 5 Conclusion

Translation



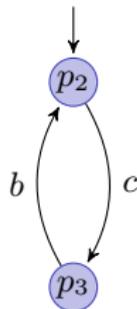
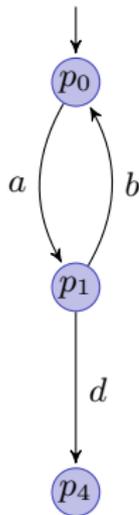
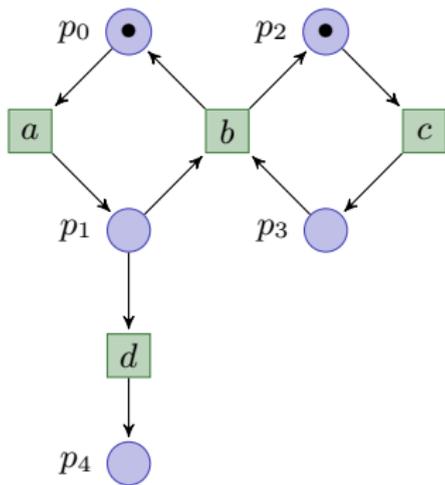
Translation

Decomposing the **untimed** PN.



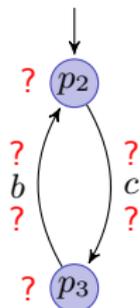
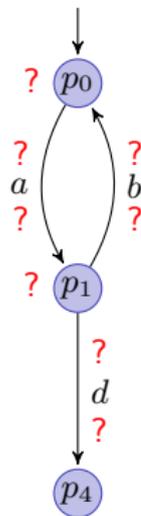
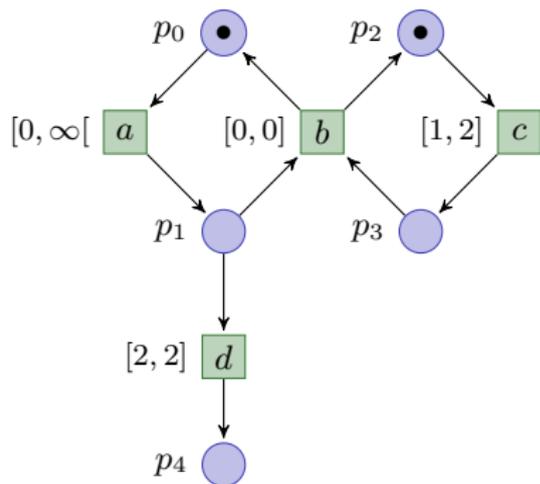
Translation

Translating each subnet into an automaton.



Translation




Adding timing constraints (resets, guards and invariants).



Translation

$$t \text{ enabled} \implies \nu(t) = \min_{\{i | t \in \Sigma_i\}} (\nu(x_i))$$

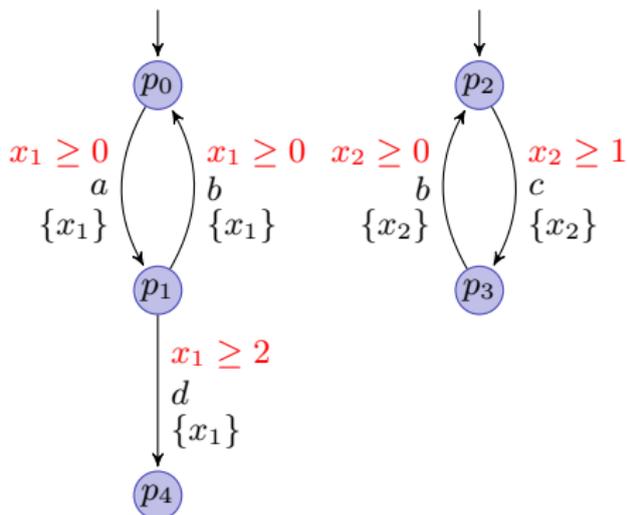

We add one clock to each automaton. The clock is reset on each edge.



Translation

$$t \text{ enabled} \implies \nu(t) = \min_{\{i | t \in \Sigma_i\}} (\nu(x_i))$$



We add guards. $\min_{\{i | t \in \Sigma_i\}} (\nu(x_i)) \geq efd(t) \Leftrightarrow \forall i \text{ s.t. } t \in \Sigma_i, \nu(x_i) \geq efd(t)$



Translation



$$t \text{ enabled} \implies \nu(t) = \min_{\{i|t \in \Sigma_i\}} (\nu(x_i))$$

We add invariants. $Inv_i(p) \equiv \bigwedge_{t \in P} (t \text{ enabled} \implies \nu(t) \leq lfd(t))$



$$\begin{aligned}
 Inv(p_1) &\equiv \overbrace{(p_1 \implies x_1 \leq 2)}^{Inv(d)} \wedge \overbrace{((p_1 \wedge p_3) \implies (\min(x_1, x_2) \leq 0))}^{Inv(b)} \\
 &\equiv (x_1 \leq 2) \wedge (\neg p_3 \vee (x_1 \leq 0 \vee x_2 \leq 0)) \\
 Inv(p_3) &\equiv (p_1 \wedge p_3) \implies (\min(x_1, x_2) \leq 0) \\
 &\equiv (\neg p_1 \vee (x_1 \leq 0 \vee x_2 \leq 0))
 \end{aligned}$$

Translation



$$Inv(p_1) \equiv (x_1 \leq 2) \wedge (\neg p_3 \vee (x_1 \leq 0 \vee x_2 \leq 0))$$

$$Inv(p_3) \equiv (\neg p_1 \vee (x_1 \leq 0 \vee x_2 \leq 0))$$

It is **unavoidable** to share clocks and states.

Properties of the translation

- ① **Timed bisimulation:** (M, ν) denotes a state of the NTA \mathcal{S} and (M, ν) a state of the TPN \mathcal{N} .

$$(M, \nu) \mathcal{R} (M, \nu) \Leftrightarrow \forall t \in \text{enabled}(M), \nu(t) = \min_{\{i | t \in \Sigma_i\}} (\nu(x_i))$$

We show that \mathcal{R} is a **timed bisimulation**.

- ② **Distributed timed language equivalence:**
- **Timed bisimulation** between the TTS of \mathcal{S} and \mathcal{N} .
 - **Bijection** between the processes of \mathcal{S} and those of \mathcal{N} (same distribution of actions up to a renaming of processes).

Size of the resulting NTA

Decomposition: at most $|P|$ processes


- at most $|P|^2$ locations,
- at most $|T| \times |P|$ edges (exactly $\sum_{t \in T} |\{i \mid t \in \Sigma_i\}|$ edges).

Timing information:

- at most $|P|$ clocks,
- $\sum_{t \in T} |\{i \mid t \in \Sigma_i\}|$ guards,
- $\sum_{t \in T} |\{i \mid t \in \Sigma_i\}|$ clock comparisons in the invariants ($Inv(t)$ can be attached to one place).

Know thy neighbour!

Given a TPN \mathcal{N} , in general, there does not exist any NTA \mathcal{S} using the local syntax (clocks and current locations are not shared) such that \mathcal{N} and \mathcal{S} have the same distributed timed language.



$$Inv(p_1) \equiv (x_1 \leq 2) \wedge (\neg p_3 \vee (x_1 \leq 0 \vee x_2 \leq 0))$$

$$Inv(p_3) \equiv (\neg p_1 \vee (x_1 \leq 0 \vee x_2 \leq 0))$$

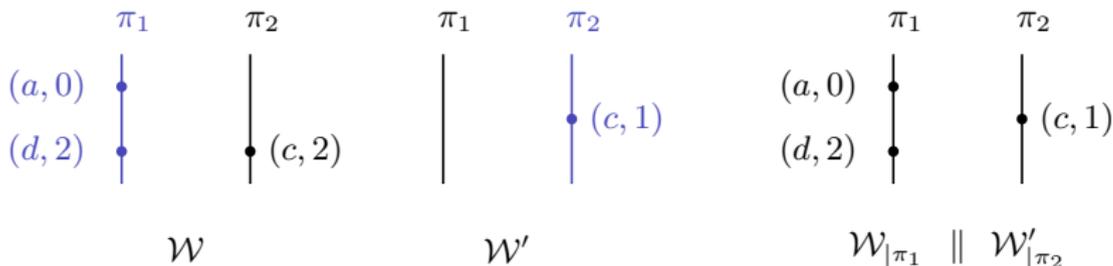
Know thy neighbour!

Given a TPN \mathcal{N} , in general, there does not exist any NTA \mathcal{S} using the local syntax (clocks and current locations are not shared) such that \mathcal{N} and \mathcal{S} have the same distributed timed language.

Lemma

Let \mathcal{S} be a network of n timed automata that do not read the state of the other automata, then for any $\mathcal{W}_1, \dots, \mathcal{W}_n$ admissible timed traces without synchronization and stopping at a same date θ , $\mathcal{W}_1|_{\pi_1} \parallel \dots \parallel \mathcal{W}_n|_{\pi_n}$ is also an admissible timed trace stopping at θ .

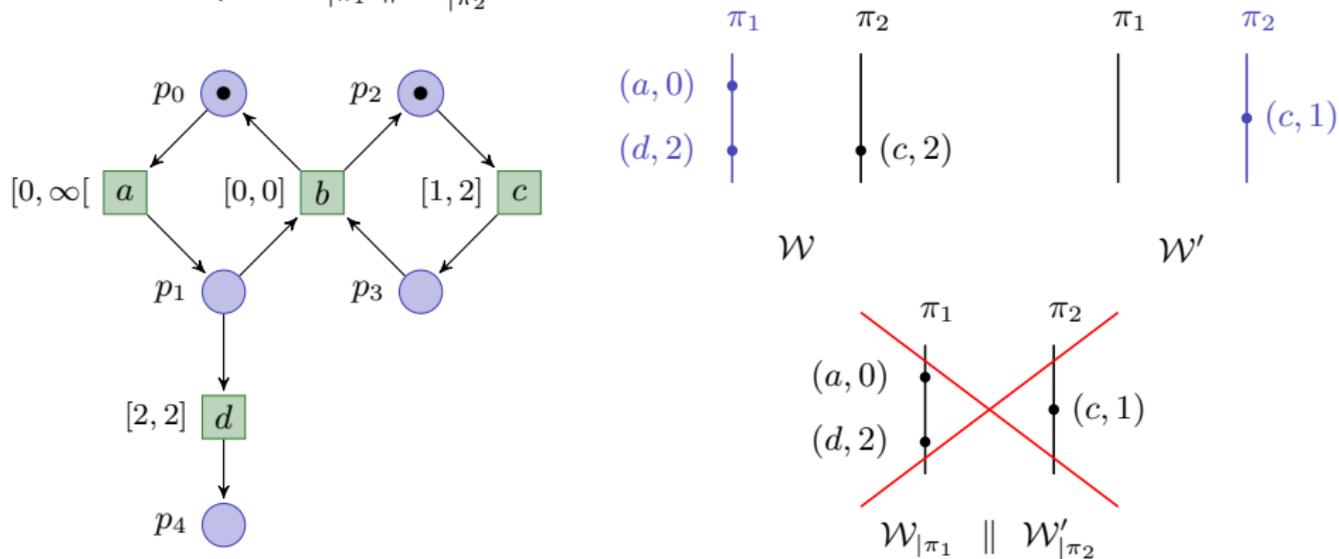
Proof



Know thy neighbour!

Given a TPN \mathcal{N} , in general, there does not exist any NTA \mathcal{S} using the local syntax (clocks and current locations are not shared) such that \mathcal{N} and \mathcal{S} have the same distributed timed language.

Counterexample: $\mathcal{W}_{|\pi_1} \parallel \mathcal{W}'_{|\pi_2}$ should be admissible.



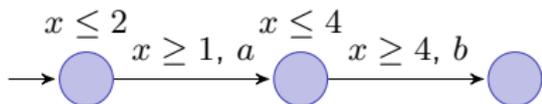
Reverse translation: from NTA to TPN

Sequential semantics: [Bérard, Cassez, Haddad, Lime, Roux, 06] When are Timed Automata weakly timed bisimilar to Time Petri Nets?

But we want to preserve the distributed semantics.

- 1 Translation of each TA in a finite “time S-net” with one token

But finite time S-nets with 1 token are strictly less expressive than TA with 1 clock



\neq

time S-net with one token

Reverse translation: from NTA to TPN

Sequential semantics: [Bérard, Cassez, Haddad, Lime, Roux, 06] When are Timed Automata weakly timed bisimilar to Time Petri Nets?

But we want to preserve the distributed semantics.

- 1 Translation of each TA in a finite “time S-net” with one token

But finite time S-nets with 1 token are strictly less expressive than TA with 1 clock

- 2 Considering the translation into more general nets,



Reverse translation: from NTA to TPN

Sequential semantics: [Bérard, Cassez, Haddad, Lime, Roux, 06] When are Timed Automata weakly timed bisimilar to Time Petri Nets?

But we want to preserve the distributed semantics.

- 1 Translation of each TA in a finite “time S-net” with one token

But finite time S-nets with 1 token are strictly less expressive than TA with 1 clock

- 2 Considering the translation into more general nets,
- 3 Composing the nets.



- 1 Introduction
 - Motivation
 - Timed and concurrent models
- 2 Partial order semantics
 - Timed traces
 - Distributed timed language
- 3 Decomposing a PN in processes
 - S-invariants
 - Decomposition
- 4 Translation from TPN to NTA
 - Adding clocks
 - Know thy neighbour!
- 5 Conclusion

Conclusion

Summary

- **Timed trace** and **distributed timed language**: description of a distributed semantics where concurrency is not erased
- **Translation** from a TPN to a NTA based on the decomposition in processes
 - Correctness w.r.t. the **distributed timed language**
 - Usable in practice (small tests with Uppaal)
 - Readable and close to the modeled system: processes are preserved

Future work

- Identification of TPN with good decompositional properties (no need to share clocks).
- Explore timed concurrency
 - Definition and properties
 - Use in verification tools
 - [Lugiez, Niebert, Zennou, 05] A partial order semantics approach to the clock explosion problem of timed automata
 - [Niebert, Qu, 06] invariants