Methods for the specification and verification of business processes

MPB (6 cfu, 295AA)

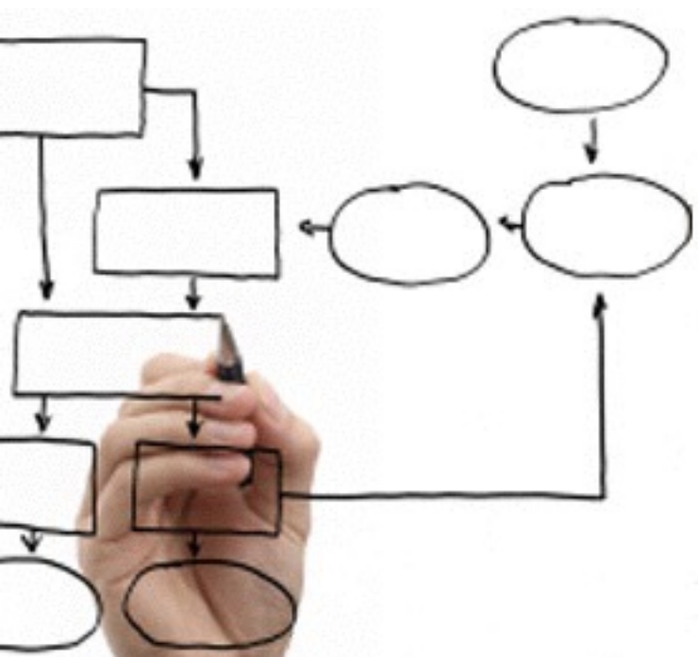

Roberto Bruni

http://www.di.unipi.it/~bruni

09 - Petri nets properties

Object

$$N \vdash \psi$$

We give a formal account of some key properties of Petri nets

Free Choice Nets (book, optional reading)
https://www7.in.tum.de/~esparza/bookfc.html

Petri nets: behavioural properties

Properties of Petri nets

We introduce some of the properties of Petri nets that can play an important role in the verification of business processes

Liveness
Deadlock-freedom
Boundedness
Cyclicity (also Reversibility)

Liveness

A transition t is **live**, if from any reachable marking M another marking M' can be reached where t is enabled

In other words:
at any point in time of the computation, we cannot exclude that t will fire in the future
or, equivalently,
at any point in time of the computation, it is still possible to enable t in the future

A Petri net is **live** if all of its transitions are live

Liveness illustrated

For any reachable marking M

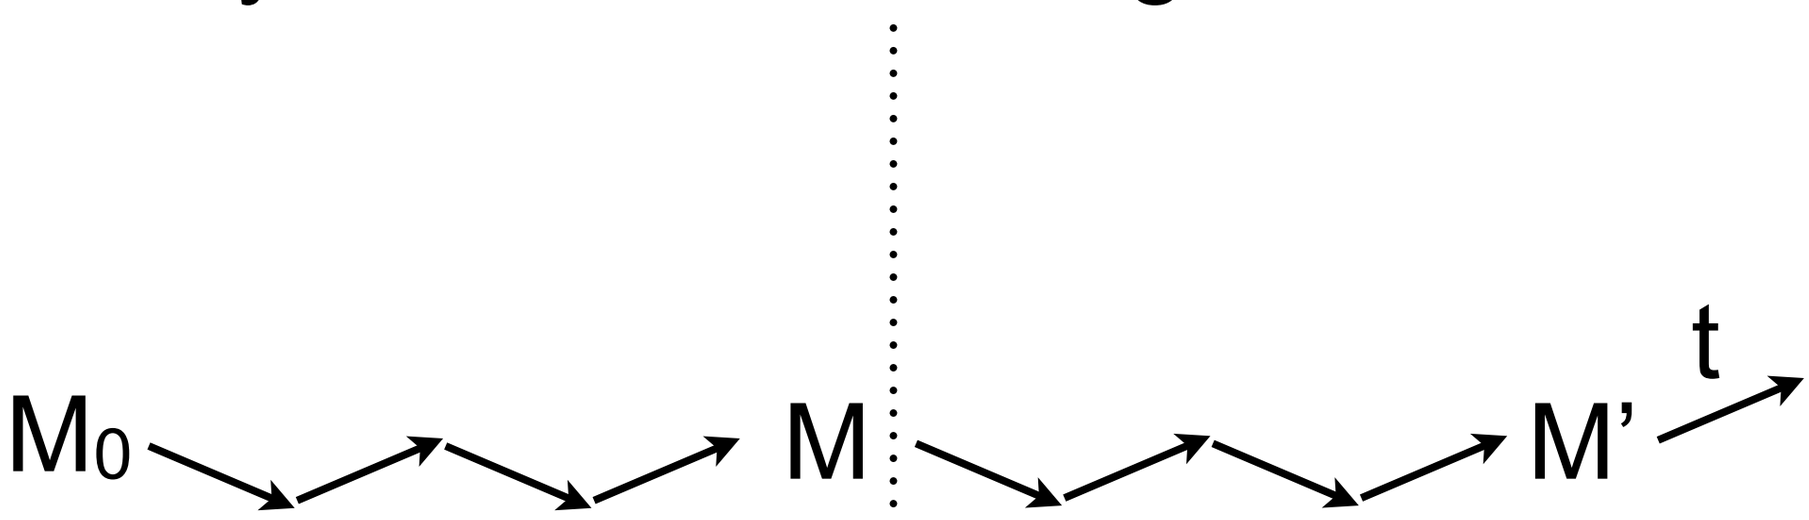

Can we find a way to enable t?

Liveness, formally

$$(P, T, F, M_0)$$

$$\forall t \in T, \quad \forall M \in [\, M_0 \,\rangle, \quad \exists M' \in [\, M \,\rangle, \quad M' \xrightarrow{t}$$

Digression

Order of quantifier is important:

quantification of the same kind can be switched

the order of universal and existential quantification is important

$$\forall n.\ \exists m.\ n < m \quad \not\equiv \quad \exists m.\ \forall n.\ n < m$$

Liveness: pay attention!

Liveness should not be confused with the following property:

starting from the initial marking M_0 it is possible to reach a marking M that enables t

$$\exists M \in [M_0\rangle.\ M \xrightarrow{t}$$

(this property just ensures that t is not "dead" in M_0)

Liveness: example

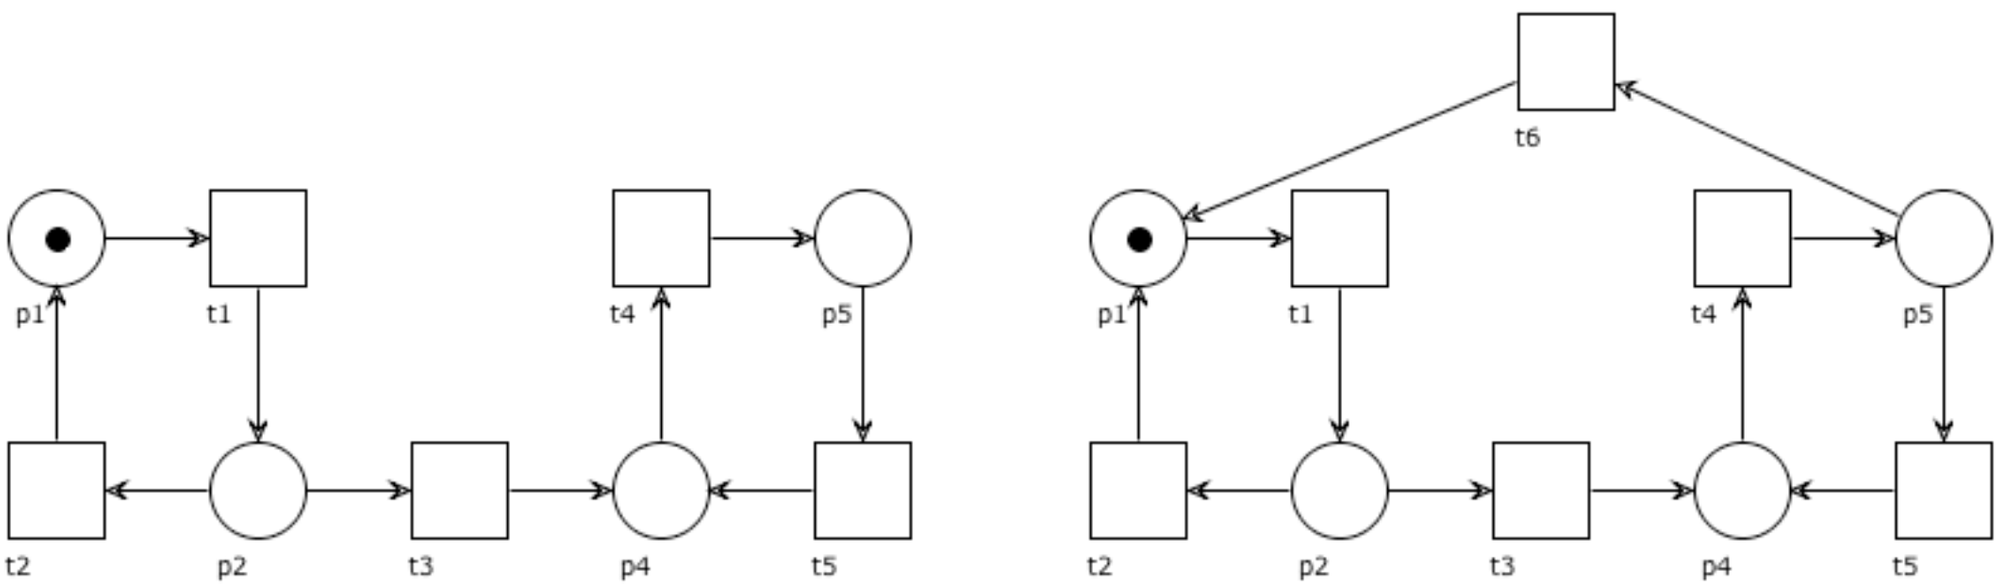

Which transitions are live?
Which are not?
Is the net live?

Marked place

Given a marking M

We say that a place p is **marked** (in M)
if M(p) > 0
(i.e., there is a token in p in the marking M)

We say that p is **unmarked**
if M(p) = 0
(i.e., there is no token in p in the marking M)

Live place, intuitively

A place p is live

if every time it becomes unmarked

there is still the possibility to be marked in the future

(or if it always stays marked)

Live place

Definition: Let (P, T, F, M_0) be a net system.

A place $p \in P$ is **live** if $\forall M \in [\, M_0 \,\rangle . \exists M' \in [\, M \,\rangle . M'(p) > 0$

Place liveness

Definition:

A net system (P, T, F, M_0) is **place-live** if every place $p \in P$ is live

Dead nodes, intuitively

Given a marking M

A transition t is **dead** at M
if t will never be enabled in the future
(i.e., t is not enabled in any marking reachable from M)

A place p is **dead** at M
if p will never be marked in the future
(i.e., p is unmarked in any marking reachable from M)

Dead nodes

Definition: Let (P, T, F) be a net

A transition $t \in T$ is **dead** at M if $\forall M' \in [\, M \,\rangle . \, M' \not\xrightarrow{t}$

A place $p \in P$ is **dead** at M if $\forall M' \in [\, M \,\rangle . \, M'(p) = 0$

Non-live vs Dead

If a transition is dead at some reachable marking M
then it is non-live

If a place is dead at some reachable marking M
then it is non-live

being non-live implies possibly becoming dead
(but not necessarily in the current marking)

Some obvious facts

If a system is not live, it must have a transition dead at some reachable marking

If a system is not place-live, it must have a place dead at some reachable marking

If a place / transition is dead at M, then it remains dead at any marking reachable from M
(the set of dead nodes can only increase during a run)

Every transition in the pre- or post-set of a dead place is also dead

Liveness implies place-liveness

Proposition: Live systems are also place-live

Take any p and any $t \in \bullet p \cup p \bullet$

Let $M \in [\, M_0 \,\rangle$

By liveness: there is $M', M'' \in [\, M \,\rangle$ s.t. $M' \xrightarrow{t} M''$

Then $M'(p) > 0$ or $M''(p) > 0$

Digression

Contrapositive property

$$P \Rightarrow Q \qquad \equiv \qquad (\neg Q) \Rightarrow \neg P$$

Exercise

Draw a net that
is place-live but not live
(if you can)

Deadlock-freedom

A Petri net is **deadlock free**, if every reachable marking enables some transition

In other words, we are guaranteed that at any point in time of the computation, some transition can be fired

Deadlock-freedom illustrated

For any reachable marking M

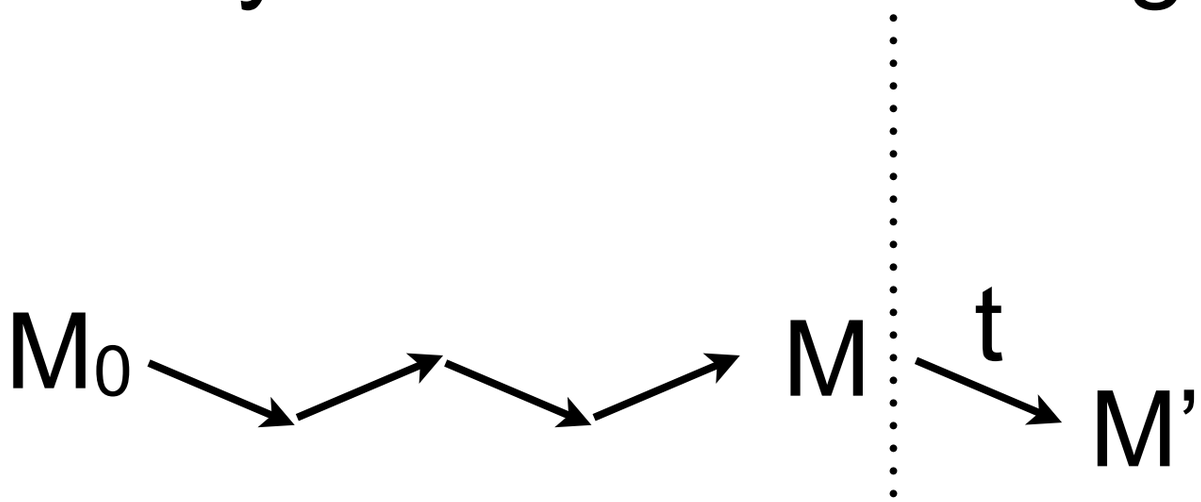

Can we fire some transition?

Deadlock freedom, formally

$$(P, T, F, M_0)$$

$$\forall M \in [\, M_0 \,\rangle, \quad \exists t \in T, \quad M \xrightarrow{t}$$

Deadlock-freedom: example

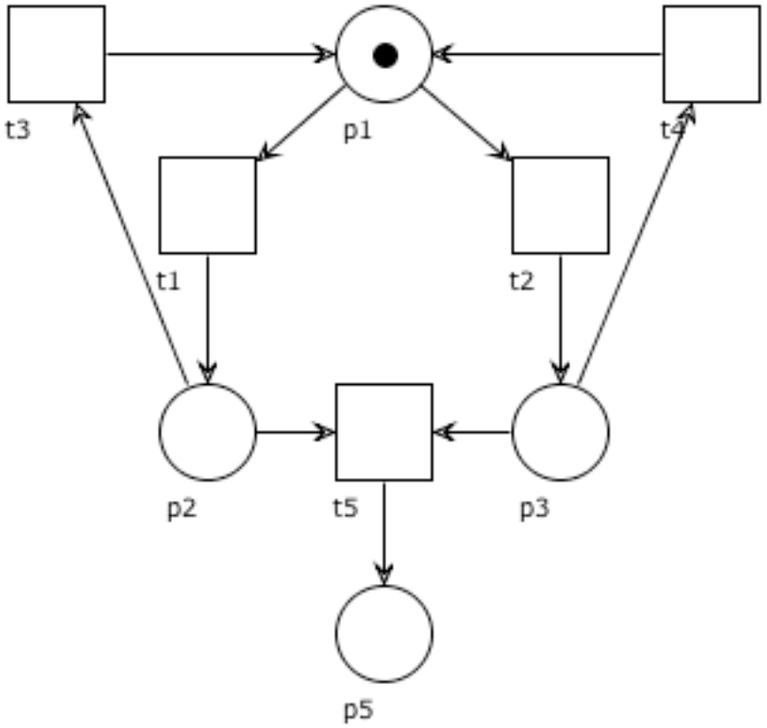

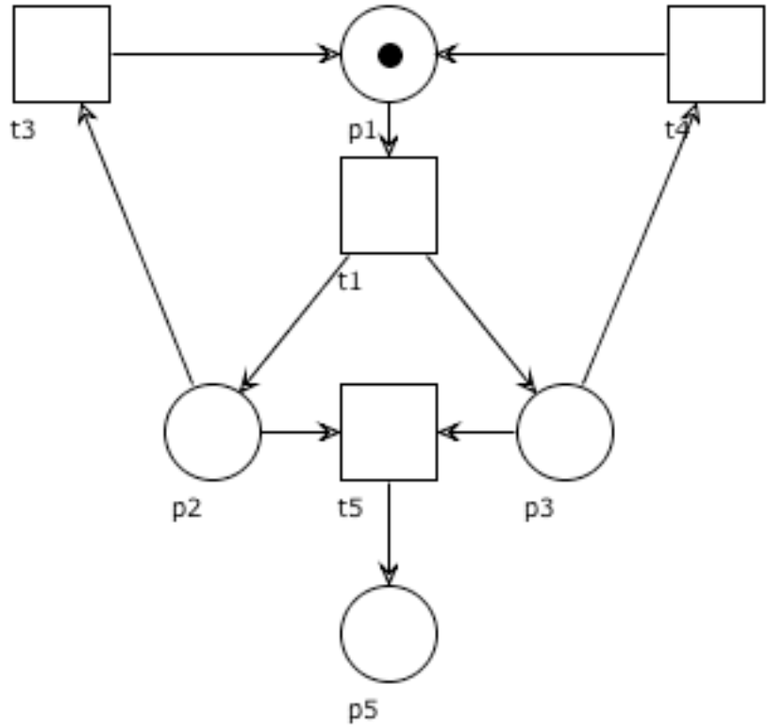

Is the net deadlock-free?

Question time

Does liveness imply deadlock-freedom?
(Can you exhibit a live Petri net that is not deadlock-free?)

Does deadlock-freedom imply liveness?
(Can you exhibit a deadlock-free net that is not live?)

Liveness implies deadlock freedom

Lemma If (P, T, F, M_0) is live, then it is deadlock-free

By contradiction, let $M \in [\, M_0 \,\rangle$, with $M \not\to$

Let $t \in T$ (T cannot be empty).
By liveness, $\exists M' \in [\, M \,\rangle$ with $M' \xrightarrow{t}$.

Since M is dead, $[\, M \,\rangle = \{\, M \,\}$.
Therefore $M = M' \xrightarrow{t}$, which is absurd.

Exercises

Prove each of the following properties
or give some counterexamples

If a system is not place-live, then it is not live

If a system is not live, then it is not place-live

If a system is place-live, then it is deadlock-free

If a system is deadlock-free, then it is place-live

k-Boundedness

Let k be a natural number

A place p is **k-bounded** if no reachable marking has more than k tokens in place p

A net is **k-bounded** if all of its places are k-bounded

In other words, if a net is k-bounded, then k is a capacity constraint that can be imposed over places without any risk of causing “overflow”

Safe nets

A place p is **safe** if it is 1-bounded

A net is **safe** if all of its places are safe

In other words, if the net is safe, then we know that, in any reachable marking, each place contains one token at most

Boundedness

A place p is **bounded** if it is k-bounded for some natural number k

A net is **bounded** if all of its places are bounded

A net is **unbounded** if it is not bounded
(i.e., there is at least one place in which any number of tokens can appear)

Boundedness, formally

$$(P, T, F, M_0)$$

$$\exists k \in \mathbb{N}, \quad \forall M \in [\, M_0 \,\rangle, \quad \forall p \in P, \quad M(p) \leq k$$

Boundedness: example

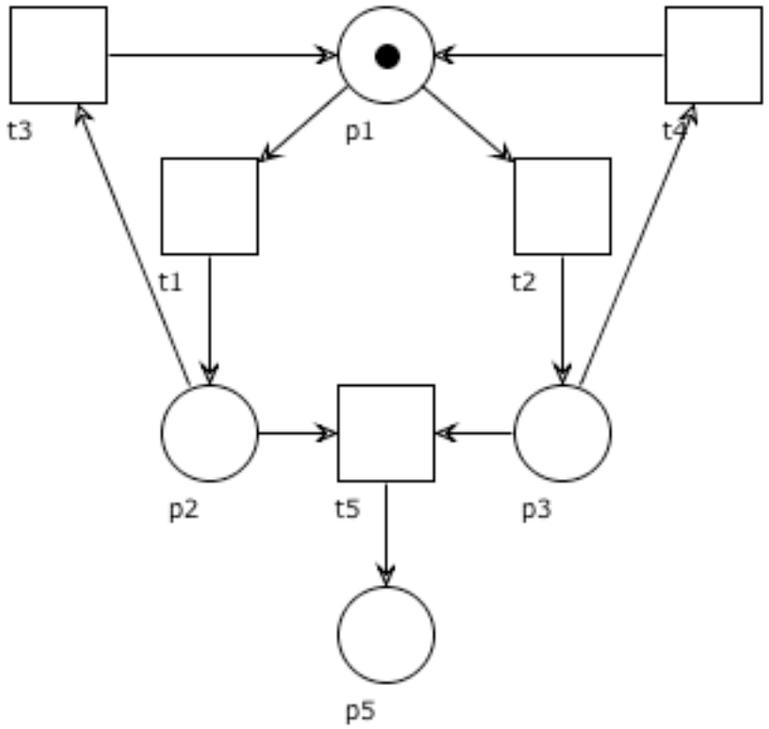

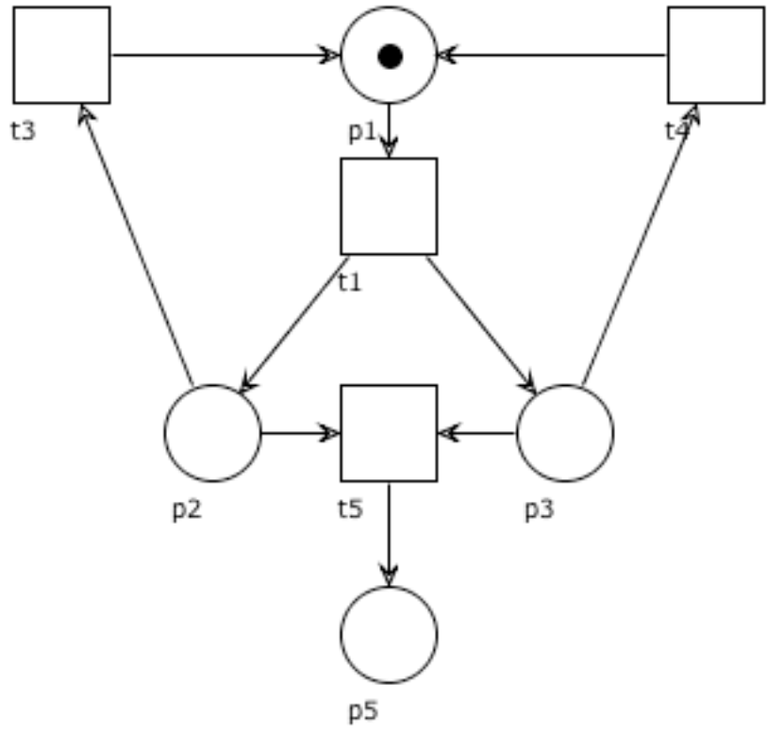

Which places are bounded?
Is the net bounded?
Which places are safe?
Is the net safe?

A puzzle about reachability

Theorem: If a system is...
then its reachability graph is finite

Theorem: A system is...
iff its reachability graph is finite

(fill the dots and the proofs)

A puzzle about boundedness

Theorem: If a system is k-bounded then any reachable marking contains a number of tokens less than or equal to

…

Theorem: If a system is safe then any reachable marking contains a number of tokens less than or equal to

…

(fill the dots and the proofs)

Cyclicity
(aka Reversibility)

A marking M is a **home marking** if it can be reached from every reachable marking

A net is **cyclic** (or **reversible**) if its initial marking is a home marking

Orthogonal properties

Liveness, boundedness and cyclicity are independent of each other

In other words, you can find nets that satisfy any arbitrary combination of the above three properties (and not the others)

Exercises

For each of the following nets, say if they are live, deadlock-free, bounded, safe, cyclic

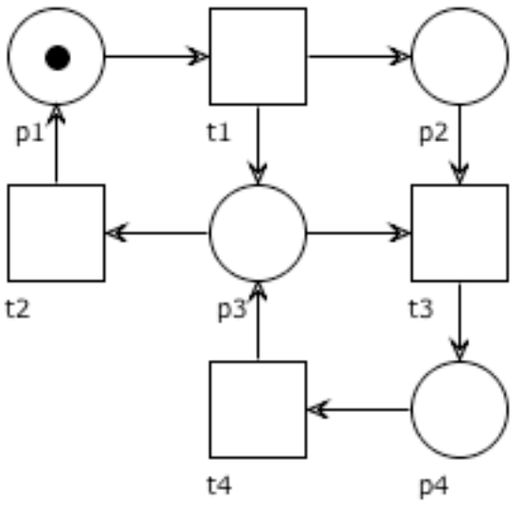

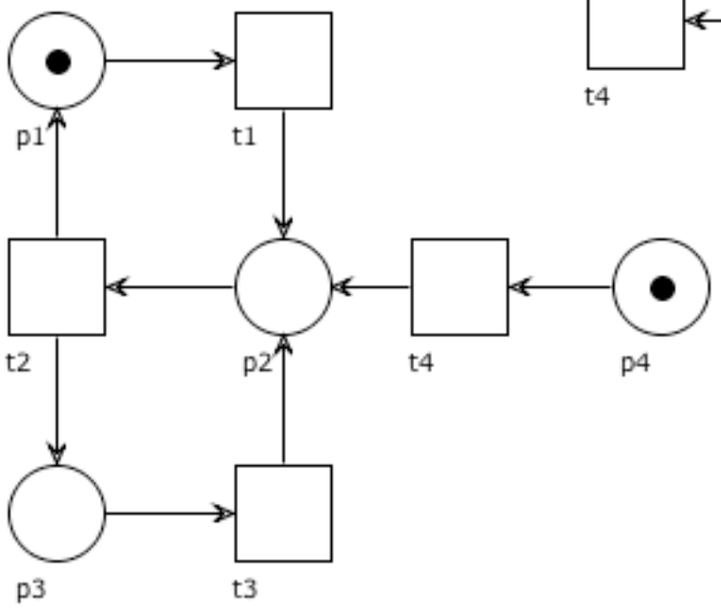

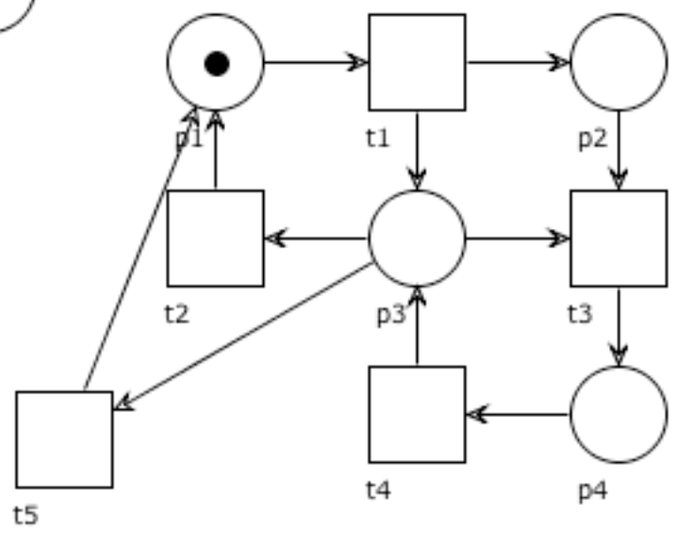

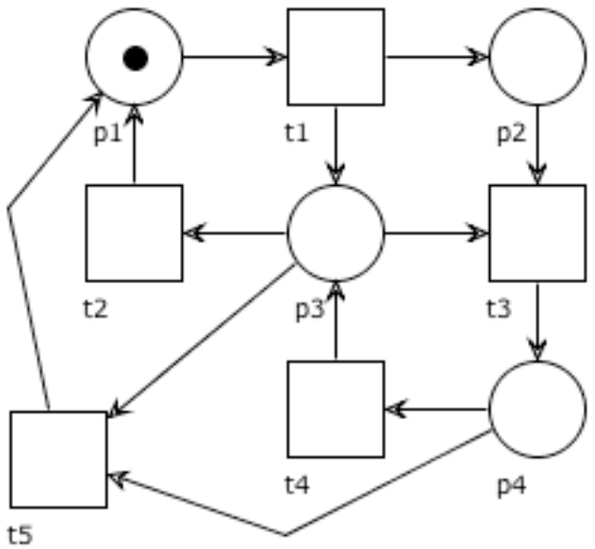

Exercises

For each of the following nets, say if they are live, deadlock-free, bounded, safe, cyclic

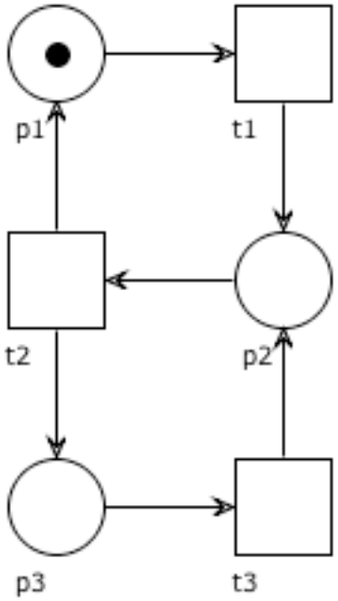

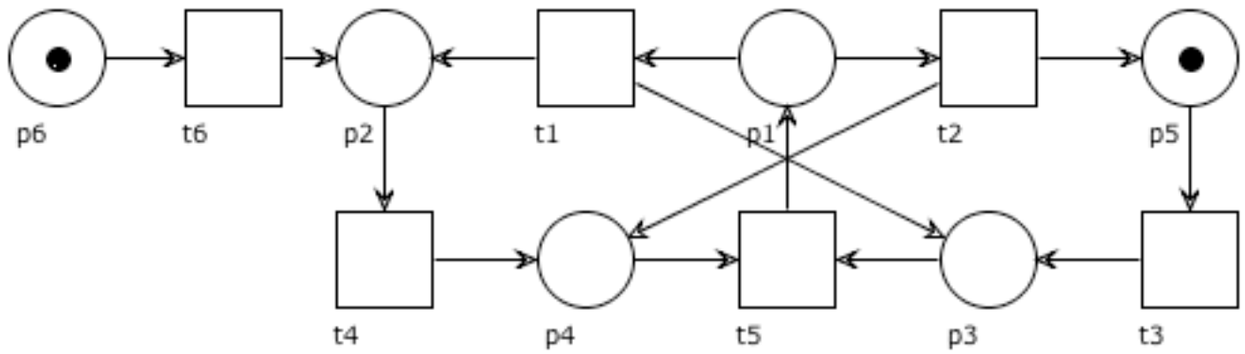

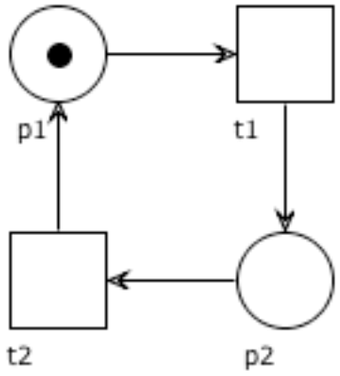

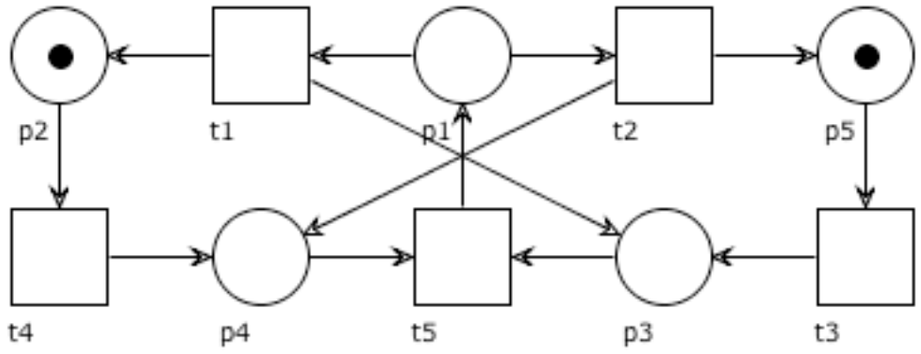

Petri nets:
structural properties

Structural properties

All the properties we have seen so far are
behavioural (or **dynamic**)
(i.e. they depend on the initial marking and firing rules)

It is sometimes interesting to connect them to
structural properties
(i.e. the shape of the graph representing the net)

This way we can give **structural characterization** of behavioural properties for a class of nets
(computationally less expensive to check)

A matter of terminology

To better reflect the above distinction, it is frequent:

to use the term **net system** for denoting a Petri net with a given initial marking
(we study behavioural properties of systems)

to use the term **net** for denoting a Petri net without specifying any initial marking
(we study structural properties of nets)

Paths and circuits

A **path** of a net (P, T, F) is a non-empty sequence $x_1 x_2 ... x_k$ such that

$$(x_i, x_{i+1}) \in F \quad \text{for every } 1 \leq i < k$$

(and we say that it leads from x_1 to x_k)

A path from x to y is called a **circuit** if:
no element occurs more than once in it and $(y, x) \in F$
(since for any x we have $(x, x) \notin F$, hence a circuit involves at least two nodes)

Connectedness

A net (P,T,F) is **weakly connected** iff it does not fall into (two or more) unconnected parts
(i.e. no two subnets (P_1,T_1,F_1) and (P_2,T_2,F_2) with disjoint and non-empty sets of elements can be found that partition (P,T,F))

A weakly connected net is **strongly connected** iff for every arc (x,y) there is a path from y to x

Connectedness, formally

A net (P, T, F) is **weakly connected** if every two nodes x, y satisfy

$$(x, y) \in (F \cup F^{-1})^*$$

(here the * denotes the reflexive and transitive closure of a binary relation)

(i.e. if there is an undirected path from x to y)

It is **strongly connected** if $(x, y) \in F^*$

(here the * denotes the reflexive and transitive closure of a binary relation)

A note

In the following we will consider (implicitly) weakly connected nets only

(if they are not, then we can study each of their subsystems separately)

S-systems / S-nets

A Petri net is called **S-system** if every transition has one input place and one output place
(S comes from *Stellen*, the German word for place)

This way any synchronization is ruled out

The theory of S-systems is very simple

T-systems / T-nets

A Petri net is called **T-system** if every place has one input transition and one output transition

This way all choices/conflicts are ruled out

T-systems have been studied extensively since the early Seventies

Interference of conflicts and synch

Typical situation:

initially t1 and t2 are not in conflict

but when t3 fires they are in conflict
(the firing of t3 is not controllable)

How to rule this situation out?

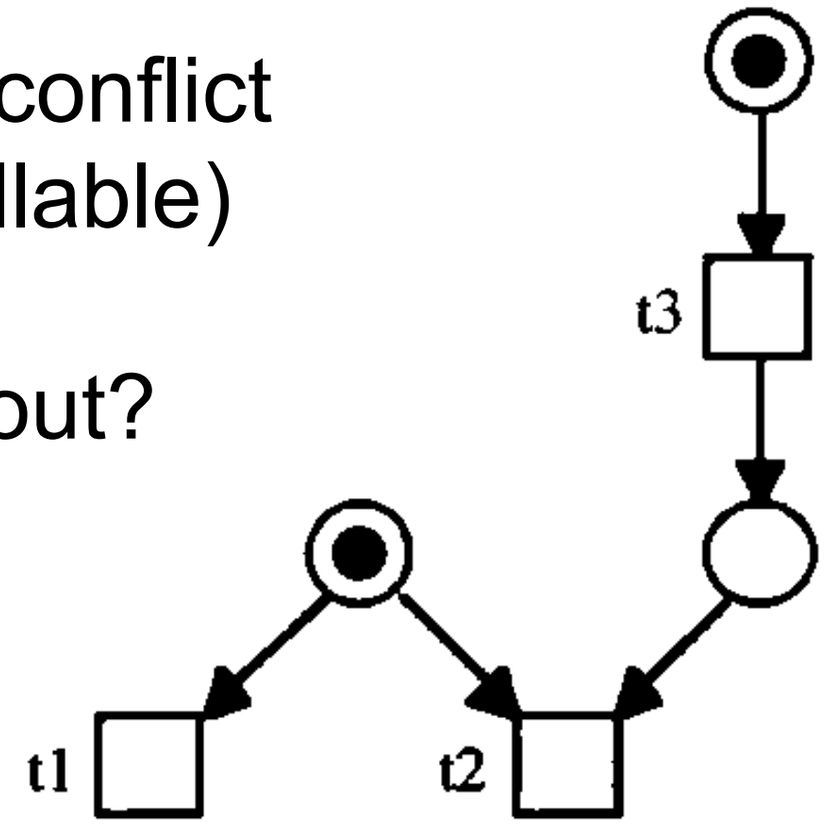

Free-choice nets

The aim is to avoid that a choice between transitions is influenced by the rest of the system

Easiest way:
keep places with more than one output transition apart from transitions with more than one input place

In other words, if (p,t) is an arc, then it means that
t is the only output transition of p (no conflict)
OR
p is the only input place of t (no synch)

Free-choice systems / nets

But we can study a slightly more general class of nets by requiring a weaker constraint

A Petri net is **free-choice** if
whenever there is an arc (p,t), then there is an arc
from any input place of t
to any output transition of p

Question time

Is the net an S-net, a T-net, free-choice?