Concurrency Theory

Winter 2025/26

Lecture 15: True Concurrency Semantics of Petri Nets: Distributed Runs

Thomas Noll, Peter Thiemann
Programming Languages Group
University of Freiburg

https://proglang.github.io/teaching/25ws/ct.html

Thomas Noll, Peter Thiemann

Winter 2025/26

Nets

Definition (Petri net)

A Petri net N is a triple (P, T, F) where:

- P is a finite set of places,
- T is a finite set of transitions with $P \cap T = \emptyset$, and
- $F \subseteq (P \times T) \cup (T \times P)$ are the arcs.[a]

Places and transitions are generically called nodes.

[a] F is also called the flow relation.

Example

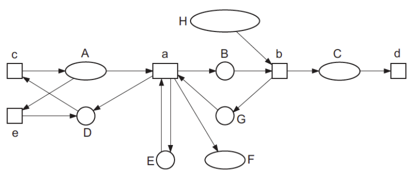

$$P = \{A, B, C, \ldots\}$$
$$T = \{a, b, c, \ldots\}$$
$$F = \{(A, a), (a, B), (B, b), \ldots\}$$

Markings

Definition (Marking)

- A marking M of a net $N = (P, T, F)$ is a mapping $M : P \to \mathbb{N}$.
- For net $N = (P, T, F)$ and marking M_0, the tuple (P, T, F, M_0) is called an elementary system net with initial marking M_0.

Intuition:

- A marking can be seen as a multiset of places.
- It defines a distribution of tokens across places.
- Tokens are depicted as black dots.

$M(p) = 3$

Remark: In generic (= non-elementary) system nets, several types (colours) of tokens can be distinguished.

Transition Occurrence

Definition (Enabling and occurrence of a transition)

Let (P, T, F, M) be an elementary system net.

- Marking M enables a transition t if $M(p) \geq 1$ for each place $p \in {}^\bullet t$.
- Transition t can occur in marking M if t is enabled in M.
- Its occurrence or firing leads to marking M', and defined for place $p \in P$ by:

$$M'(p) := M(p) - F(p,t) + F(t,p)$$

  where we represent F by its characteristic function.

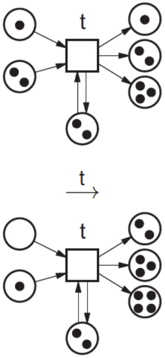

Reachable Markings

Definition (Step sequence)

Let (P, T, F, M_0) be an elementary system net.

- A sequence of transitions $\sigma = t_1\, t_2 \ldots t_n \in T^*$ is a step sequence if there exist markings $M_1, \ldots, M_n$ such that

$$M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \cdots \xrightarrow{t_{n-1}} M_{n-1} \xrightarrow{t_n} M_n.$$

- Marking M_n is then reached by the occurrence of σ, denoted $M_0 \xrightarrow{\sigma} M_n$.
- M is a reachable marking if there exists a step sequence σ such that $M_0 \xrightarrow{\sigma} M$.

Example

In the previous example,

$$\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \xrightarrow{cbedeab} \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$$

Marking Graph

Definition (Marking graph)

The marking graph of a net N has as nodes the reachable markings of N and as edges the corresponding steps of N.[a]

[a] Since firing an (enabled) transition in a marking yields a unique successor marking, marking graphs are a deterministic LTS.

Ex

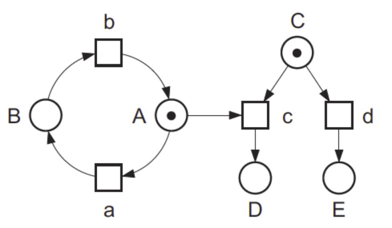

A sample elementary system net

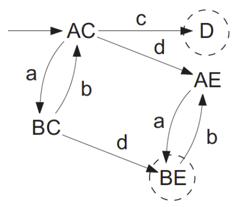

. . . and its marking graph

Interleaving semantics

Interleaving vs. True Concurrency I

The interleaving thesis

The total order assumption is a reasonable abstraction, adequate for practical purposes and leading to nice mathematics.

The true concurrency thesis

The total order assumption does not correspond to physical reality and leads to awkward representations of simple phenomena.

Motivation

Example 15.1 (cf. Example 14.20)

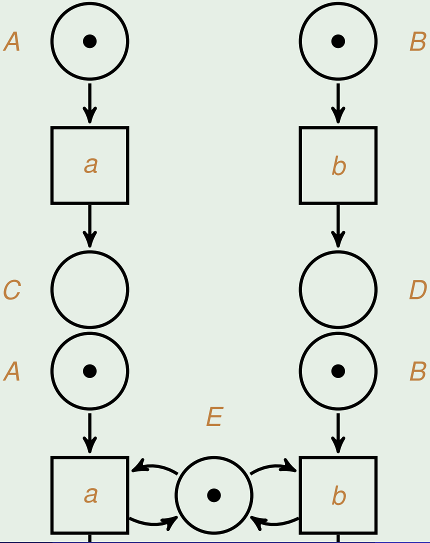

Interleaving vs. True Concurrency II

- In interleaving semantics, a system composed of n independent components

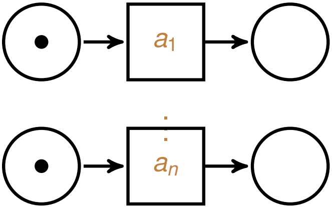

  has $n!$ different sequential runs.
- Its marking graph has 2^n states.
- In true concurrency semantics, it has only one (non-sequential) execution (“distributed run”).

Roadmap:

Net	$\mapsto$	{sequential runs}	$\widehat{=}$	marking graph (as seen)
Net	$\mapsto$	{distributed runs} (today)	$\widehat{=}$	branching process (next lecture)

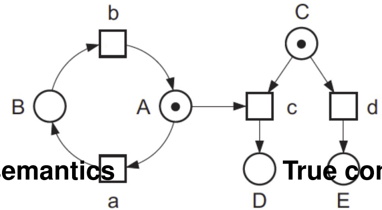

Interleaving semantics

Sequential runs (step sequences):

$$AC \xrightarrow{a} BC \xrightarrow{b} AC \xrightarrow{c} D$$
$$AC \xrightarrow{d} AE \xrightarrow{a} BE \xrightarrow{b}$$
$$AE \xrightarrow{a} \cdots$$

Marking graph (LTS):

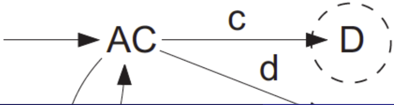

True concurrency semantics

Distributed runs (causal nets):

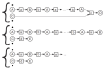

(Max.) Distributed process
(occurrence net):

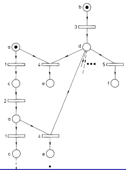

Nets with Infinite Node Sets

Definition 15.2 (Elementary system net – extending Definition 14.6)

An elementary system net is a tuple $N = (P, T, F, M_0)$ where

- P is a countable set of places,
- T is a countable set of transitions with $P \cap T = \emptyset$,
- $F \subseteq (P \times T) \cup (T \times P)$ are the arcs such that, for every transition $t \in T$, ${}^\bullet t$ and $t^\bullet$ are non-empty and finite, and
- $M_0 : P \to \mathbb{N}$ is the initial marking.

Places and transitions are generically called nodes.

Example 15.3 (A distributed run)

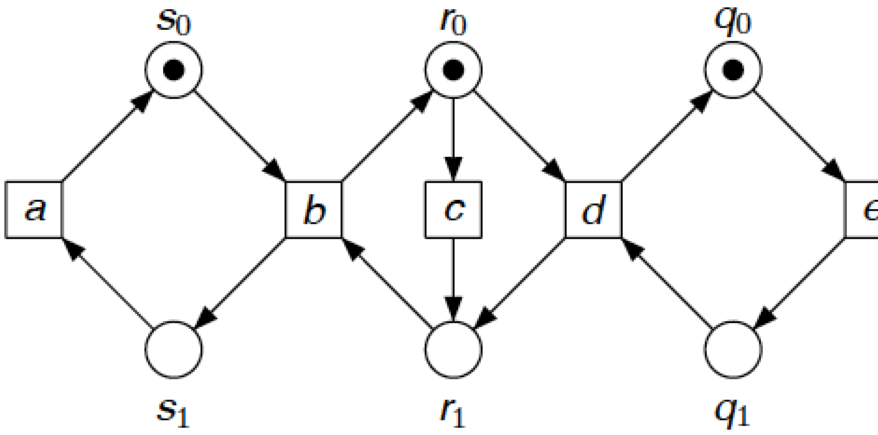

Example 15.3 (continued)

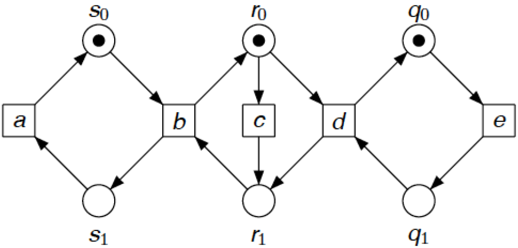

Example 15.3 (continued)

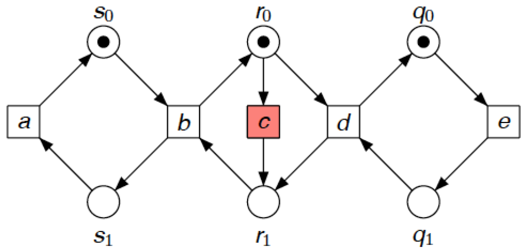

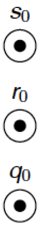

Example 15.3 (continued)

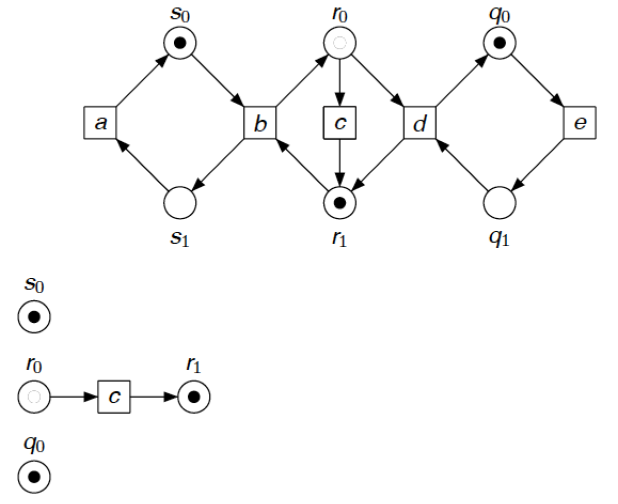

Example 15.3 (continued)

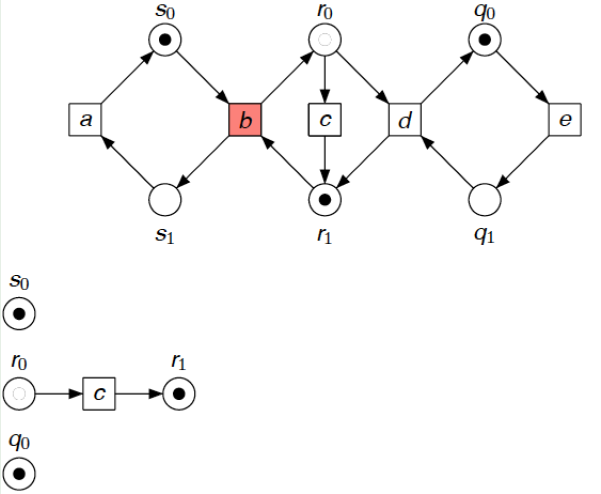

Example 15.3 (continued)

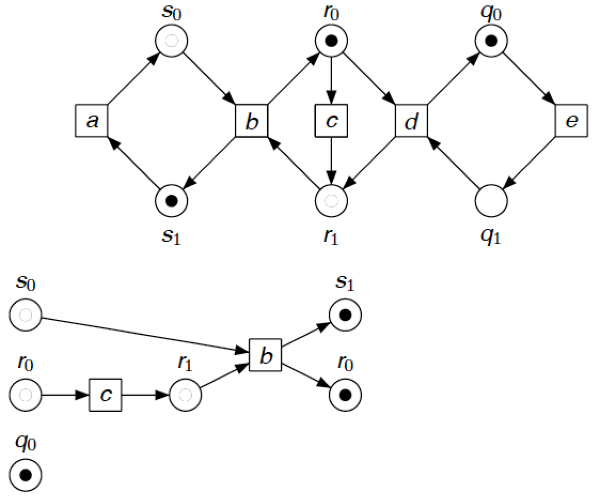

Example 15.3 (continued)

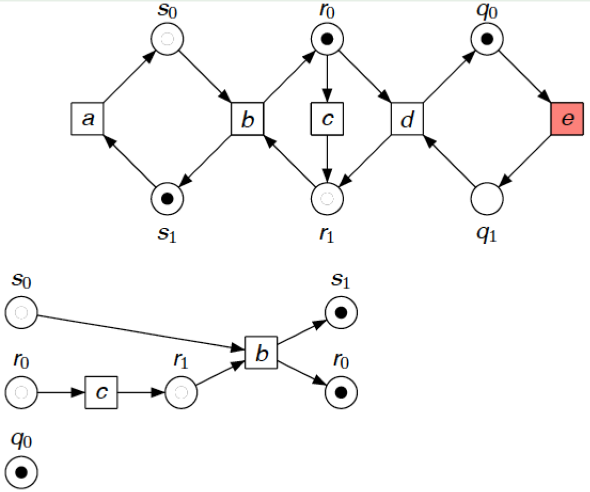

Example 15.3 (continued)

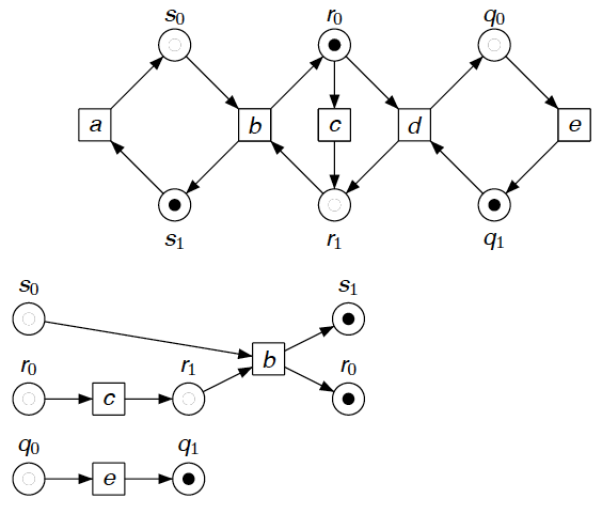

Example 15.3 (continued)

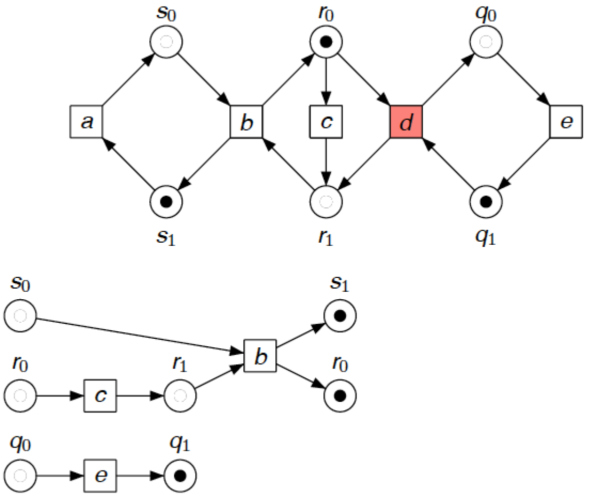

The True Concurrency Semantics of Petri Nets X

Example 15.3 (continued)

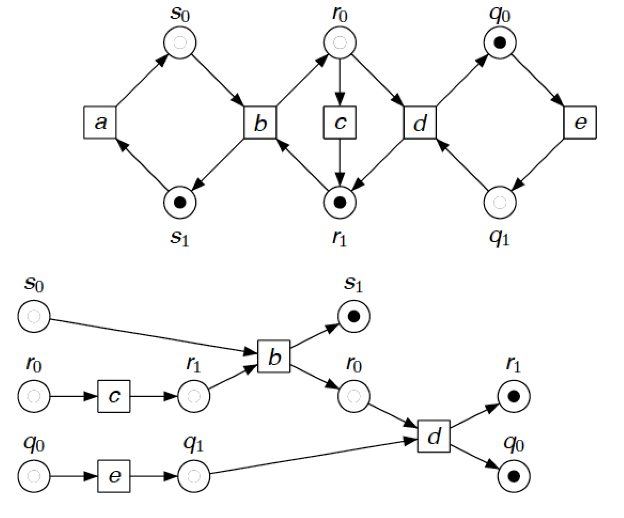

Example 15.3 (continued)

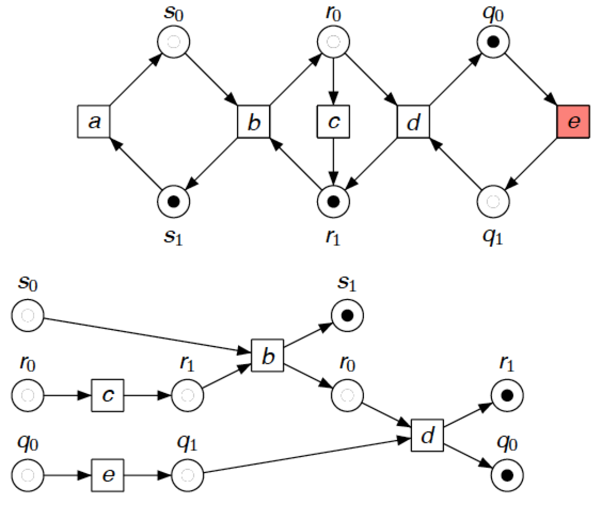

The True Concurrency Semantics of Petri Nets XII

Example 15.3 (continued)

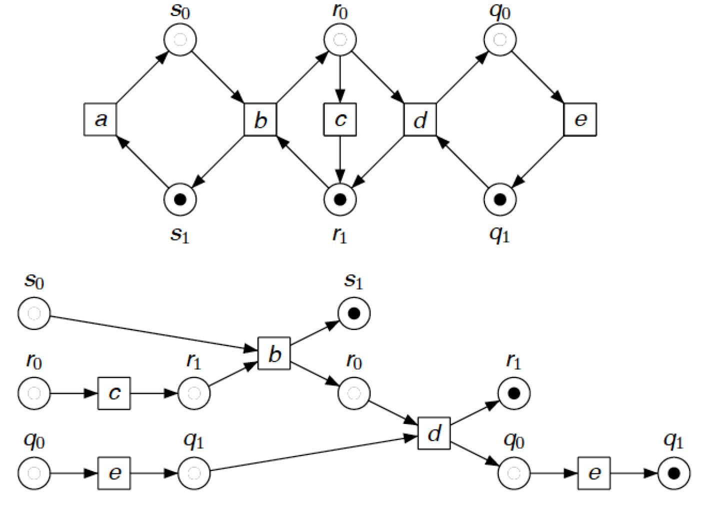

Example 15.3 (continued)

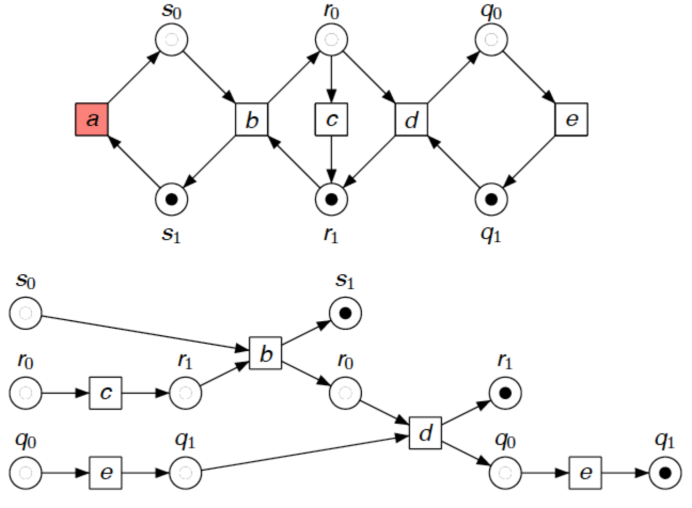

Example 15.3 (continued)

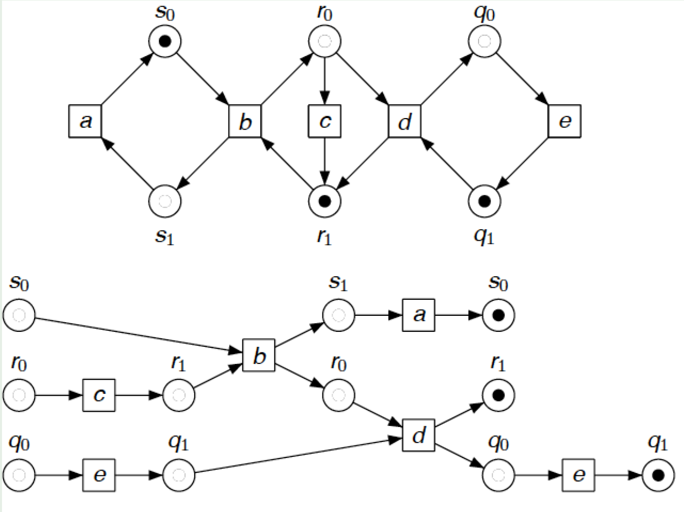

Example 15.3 (continued)

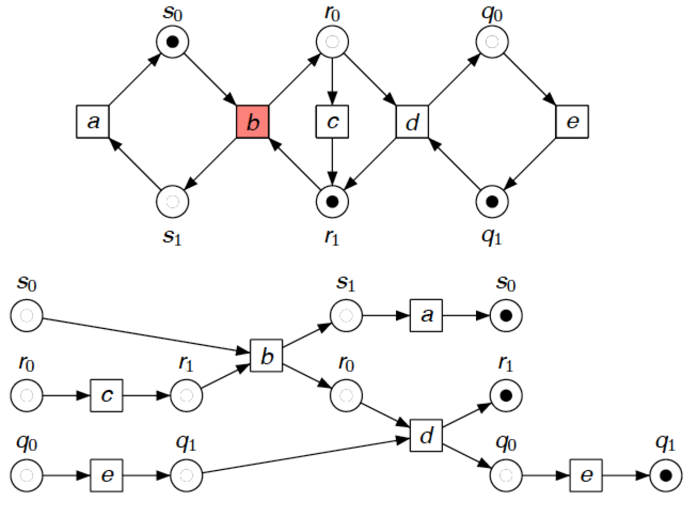

Example 15.3 (continued)

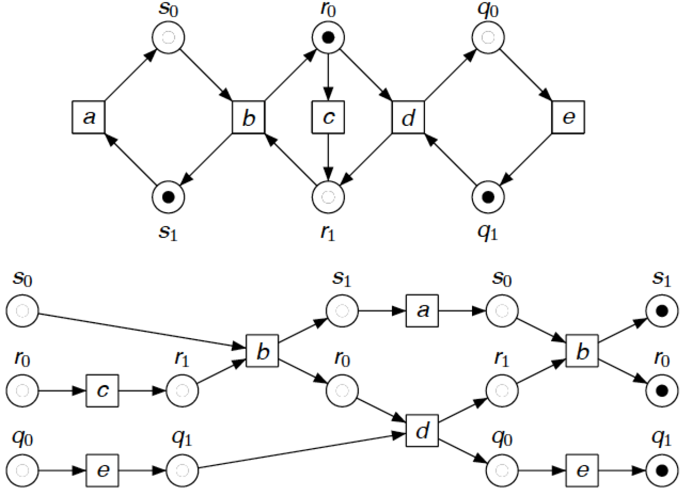

Actions

A distributed run of a net is a partial order represented as a net whose basic building blocks are simple nets denoted as actions.[1]

Definition 15.4 (Action)

An action is a labelled net $A = (Q, \{v\}, G)$ with ${}^\bullet v \cap v^\bullet = \emptyset$ and ${}^\bullet v \cup v^\bullet = Q$.

Actions represent transition occurrences of elementary system nets. If A represents transition t, then the elements of Q are labelled with in- and output places of t, and v is labelled with t.

Example 15.5

(action) 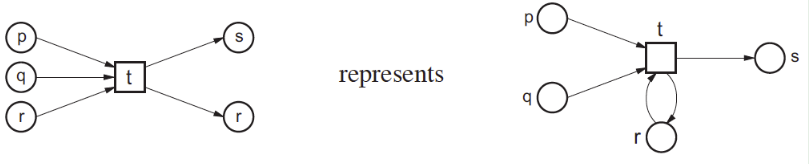

(transition)

Example 15.6

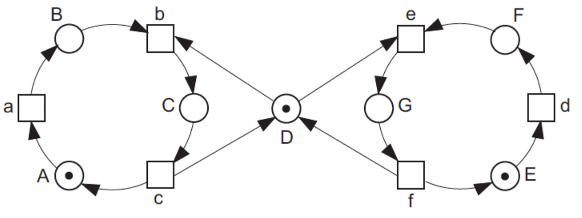

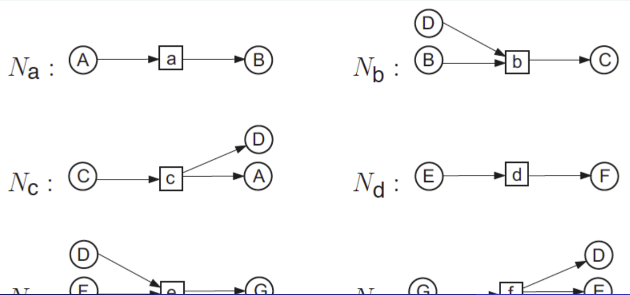

Actions Represent (Repeated) Transition Occurrences

Example 15.7

Original net

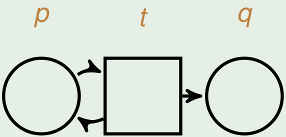

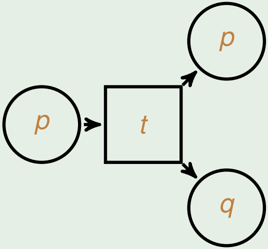

Action of t

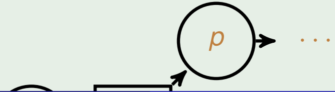

Causal Nets Informally

A causal net constitutes the basis of a “distributed” run.

It is a (possibly infinite) elementary system net with the following properties:

(1) It has no place branches: at most one arc ends or starts in a place.
(2) It is acyclic, i.e., no sequence of arcs forms a loop.
(3) Each sequence of arcs (flows) has a unique first element.
(4) The initial marking contains all places without incoming arcs.

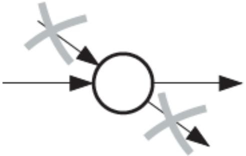

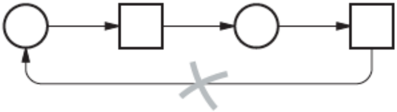

Causal Nets Formally

A causal net is a (possibly infinite) net with the following properties:

(1) It has no place branches: at most one arc ends or starts in a place.
(2) It is acyclic, i.e., no sequence of arcs forms a loop.
(3) Each sequence of arcs (flows) has a unique first element.
(4) The initial marking contains all places without incoming arcs.

Definition 15.8 (Causal net)

A (possibly infinite) net $K = (Q, V, G, M_0)$ is called a causal net if

(1) for each $q \in Q$, $|^\bullet q| \leq 1$ and $|q^\bullet| \leq 1$,
(2) the transitive closure (called causal order) G^+ of G is irreflexive,
(3) for each node $x \in Q \cup V$, the set $\{y \mid (y, x) \in G^+\}$ is finite (i.e., G is well-founded), and
(4) M_0 equals the minimal set of places in K under G^+, i.e.,

$$M_0 = {}^\circ K := \{q \in Q \mid {}^\bullet q = \emptyset\}.$$

Note: The “runs” of the nets in Examples 15.2 and 15.7 (with appropriate

Properties of Causal Nets

In a causal net places can be marked at most once.

Lemma 15.9

Let $K = (Q, V, G, M_0)$ be a causal net. Then every step sequence

$$M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \ldots\ldots \xrightarrow{t_k} M_k$$

of K satisfies $M_j \cap t_k^\bullet = \emptyset$ for all $j \in \{0, \ldots, k-1\}$.

Proof (by contraposition).

- Suppose that $p \in M_j \cap t_k^\bullet$ for some $p \in V$ and some $0 \leq j < k$.
- This is impossible for $j = 0$ as, by Definition 15.8, no place in M_0 has incoming arcs, and thus $M_0 \cap t^\bullet = \emptyset$ for each $t \in T$.
- Hence, $j > 0$. Given that $p \in M_j$ and $p \notin M_0$, it follows $p \in t_i^\bullet$ for some $0 < i \leq j$. (Before reaching M_j, some transition must have placed a token in p.)
- Thus $t_i, t_k \in {}^\bullet p$ with $t_i \neq t_k$, as G is well-founded.

Boundedness of Causal Nets

Lemma

Let $K = (Q, V, G, M_0)$ be a causal net. Then every step sequence

$$M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \ldots\ldots \xrightarrow{t_k} M_k$$

of K satisfies $M_j \cap t_k^\bullet = \emptyset$ for all $j \in \{0, \ldots, k-1\}$.

Theorem 15.10 (Boundedness of causal nets)

Every causal net is one-bounded, i.e., in every reachable marking every place will hold at most one token.

Proof.

Follows directly from the fact that the initial marking M_0 is one-bounded, and by Lemma 15.9. □

Completeness of Causal Nets I

Lemma 15.11 (Absence of superfluous places and transitions)

Let $K = (Q, V, G, M_0)$ be a causal net. Then there exists a (possibly infinite) step sequence $M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \cdots\cdots \xrightarrow{t_k} M_k \xrightarrow{t_{k+1}} \cdots$ of K such that $Q = \bigcup_{k \geq 0} M_k$ and $V = \{t_k \mid k > 0\}$.

A causal net thus contains no superfluous places and transitions, as every place is marked and every transition is fired in the above step sequence.

Completeness of Causal Nets II

Lemma (Absence of superfluous places and transitions)

Let $K = (Q, V, G, M_0)$ be a causal net. Then there exists a (possibly infinite) step sequence $M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} \cdots\cdots \xrightarrow{t_k} M_k \xrightarrow{t_{k+1}} \cdots$ of K such that $Q = \bigcup_{k \geq 0} M_k$ and $V = \{t_k \mid k > 0\}$.

Proof.

Let $depth : Q \cup V \to \mathbb{N}$ give a node's (maximal) "distance" from the initial marking:

$$depth(p) := \max\left\{\frac{d}{2} \mid d \in \mathbb{N}, \exists p_0 \in M_0 : (p_0, p) \in G^d\right\}$$
$$depth(t) := \max\{depth(p) \mid p \in {}^\bullet t\}$$

where $G^0 := \mathrm{id}_Q$ and $G^{d+1} := G \circ G^d$.

Now construct the step sequence as follows:

$$M_0 \quad \underbrace{\longrightarrow \cdots \longrightarrow}_{\text{all transitions of depth } 0} \quad M_i \quad \underbrace{\longrightarrow \cdots \longrightarrow}_{\text{all transitions of depth } 1} \quad M_j \longrightarrow \cdots$$

Example 15.12 (illustrating Lemma 15.11)

Causal net with depth information:

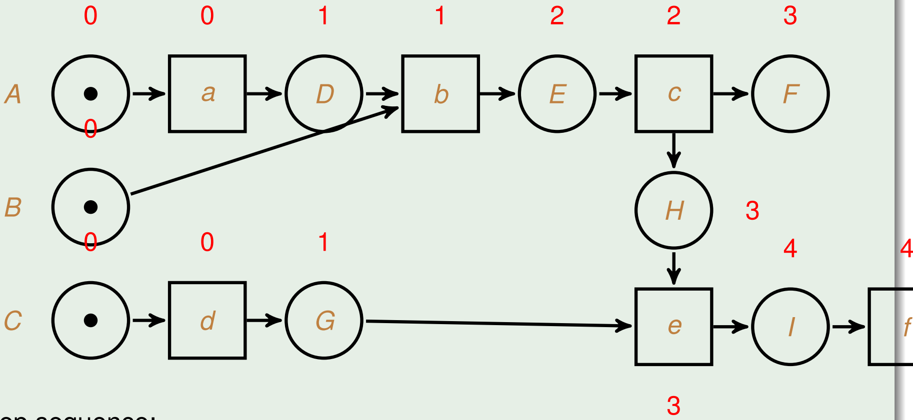

Step sequence:

Outset and End of Causal Nets

Definition 15.13 (Outset and end of a causal net)

The outset and end of causal net $K = (Q, V, G, M_0)$ are defined by the places without an incoming or outgoing arc, respectively:

$$^\circ K := \{q \in Q \mid {}^\bullet q = \emptyset\} \qquad K^\circ := \{q \in Q \mid q^\bullet = \emptyset\}.$$

What Is a Distributed Run?

Definition 15.14 (Distributed run)

A distributed run of a one-bounded elementary system net $N = (P, T, F, M_0)$ is

(1) a labelled causal net $K_N = (Q, V, G, M)$

(2) in which each transition $t \in V$ (with ${}^\bullet t$ and $t^\bullet$) is an action of N.

A distributed run K_N of N is complete if the marking $M = {}^\circ K_N$ represents the initial marking M_0 of N and the marking K_N° does not enable any transition in N.

If N is clear from the context, we just write K for K_N.

A Distributed Run for Mutual Exclusion

Example 15.15 (cf. Example 15.6)

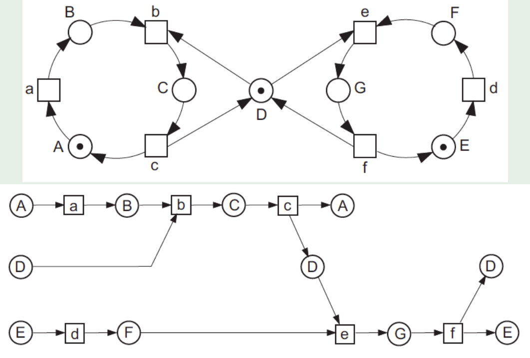

Mutual exclusion algorithm (left) and an (incomplete) distributed run (right)

- Actions N_a, N_b, N_c and N_d causally precede N_e. They form a chain.
- N_a and N_d are not linked by actions; they are causally independent.

Example 15.16

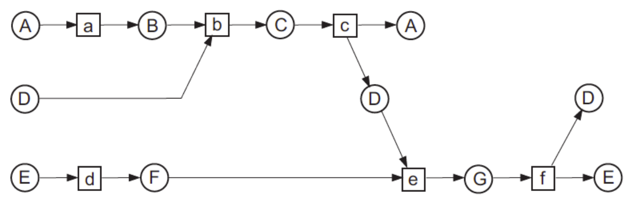

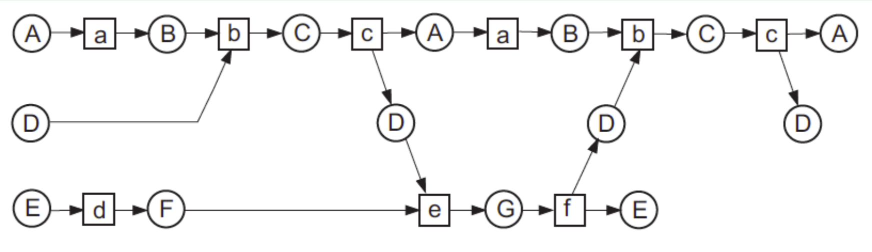

More Distributed Runs

Example 15.17

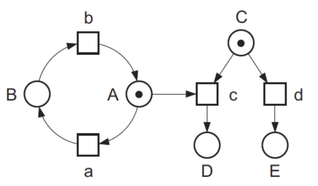

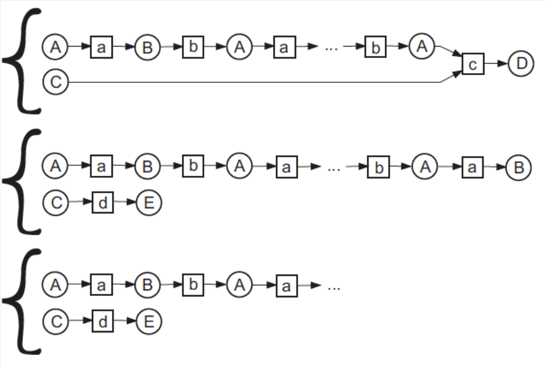

Two finite distributed runs (first complete, second incomplete)

Causality Revisited

In contrast to sequential runs, distributed runs show the causal order of actions.

Example 15.18 (cf. Example 14.20)

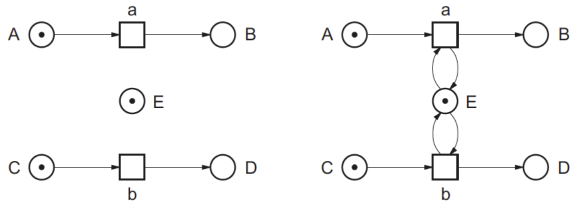

- Both nets have identical sequential runs (*a* occurs before *b*, or vice versa).
- But the left net only has the left distributed run below, the right net both ones:

Composition of Distributed Runs

Definition 15.19 (Composition of distributed runs)

For $i \in \{1, 2\}$, let $K_i = (Q_i, V_i, G_i, M_i)$ be causal nets, labelled with ℓ_i. Let $(Q_1 \cup V_1) \cap (Q_2 \cup V_2) = K_1^\circ = {}^\circ K_2$ and for each $p \in K_1^\circ$, let $\ell_1(p) = \ell_2(p)$.

The composition of K_1 and K_2, denoted $K_1 \bullet K_2$, is the causal net $(Q_1 \cup Q_2, V_1 \cup V_2, G_1 \cup G_2, M_1)$ labelled with ℓ where $\ell(x) = \ell_i(x)$ if $x \in K_i$.

Intuition: The composition $K_1 \bullet K_2$ is formed by identifying the end K_1° of K_1 with the outset ${}^\circ K_2$ of K_2. To do this, K_1° and ${}^\circ K_2$ must represent the same marking.

Example 15.20 (Composable runs)