

Floating Channels between Communicating Nets

Ludwik Czaja^{1,2}

¹ Academy of Finance and Business Vistula, Warsaw

² Institute of Informatics, The University of Warsaw, Poland
lczaja@mimuw.edu.pl

Abstract. A network system is given as a set of Petri net-like structures called agents. Each agent has a singled out place interpreted as a communication port with ingoing edges labelled with $send(p_1, \dots, p_n)$ and $receive(q_1, \dots, q_m)$ commands, where p_i, q_j are names of ports of its interlocutors. Every such edge exits a transition emitting a request for send or receive message. A transmission channel between the agent and its interlocutors is established when its port holds a send or receive command, while ports of its interlocutors hold respective (matching) communication commands. This gives rise to communication between the agent and its interlocutors, after which the channel is disrupted: hence *floating channels*. Some behavioural properties of such network system are examined, their decision complexity, deadlock and fairness in their number.

1 Introduction

A system of communicating agents here is a collection of Petri net-like structures [Rei 1985], such that in every net there is a singled out place serving for communication and called a port. Each arrow entering the agent's port is labelled with a *send* or *receive* communication statement with parameters being names of ports the agent sends a message to, or receives from. Firing a transition the arrow outgoes, results in putting the arrow's label in the port. If it is a send (receive) statement and all ports - its parameters - hold matching receive (send) statements, then a communication channel between senders and receivers is set up ("matching" in the sense of "hand-shaking" [Hoa 1978], [Hoa 1985], [OCCAM 1984]). The channel is realized as a special transition, called a *transmission*, with sending and receiving ports as its preset and postset respectively. Firing such transition represents a message transfer between the ports involved. After firing, this transition disappears, thus the channel is disrupted. That is why we say that the channels are floating. Such systems are defined, examples shown and some behavioural properties investigated. If the 1-safe Petri nets are taken as the agents, then complexity of a number of decision problems for systems with floating channels become straightforward conclusions from known results, collected e.g. in [ESP 1998]. Some problems, namely deadlock and two kinds of fairness is analysed in the framework of the proposed model and their set-theoretic characteristics are given.

2 Communicating Agents

Let $A = \{A_{p_1}, A_{p_2}, \dots, A_{p_d}\}$ be a set of agents, each agent A_{p_i} ($i = 1, 2, \dots, d$) equipped with a single communication port p_i , their set $P = \{p_1, p_2, \dots, p_d\}$. A is treated as a distributed system whose agents are capable of intercommunicating through their ports. Suppose that the agents are autonomous, i.e. do not share any of their constituents. Let $!(p_{k_1}, p_{k_2}, \dots)$ and $?(p_{l_1}, p_{l_2}, \dots)$ be a shorter notation of $send(p_{k_1}, p_{k_2}, \dots)$ and $receive(p_{l_1}, p_{l_2}, \dots)$ operations respectively, i.e. sending a message by an agent to ports p_{k_1}, p_{k_2}, \dots and receiving a message from ports p_{l_1}, p_{l_2}, \dots . Here k_1, k_2, \dots and l_1, l_2, \dots are subsequences of the sequence $1, 2, \dots, d$. These *communication operations* may assume a varying number of parameters and are executed in the synchronous, i.e. hand-shaking mode. Let C denote a set of all possible communication operations of all the agents, along with the empty (no communication) operation Θ . Since apart from communication, other computational activity of the agents is inessential, such fragments of their activity are not taken into consideration. That is why we assume that agent A_p with port p is represented as a single place net with a specific firing rule (semantics):

$A_p = (\{p\}, T_p, F_p)$ for $p \in P$ where:

T_p is a set of transitions, i.e. actions inserting send or receive operations in the port p ,

$F_p : T_p \times \{p\} \rightarrow C_p$ is a set of arrows from transitions to place p , each arrow labelled with a send or receive operation the agent A_p can issue, i.e. $C_p \subseteq C - \{\Theta\}$. Suppose no agent can send/receive message to/from itself. That is:

$F_p(t, p)$ is either $!(p_{k_1}, p_{k_2}, \dots)$ or $?(p_{l_1}, p_{l_2}, \dots)$ with $p_{k_i} \neq p \neq p_{l_j}$ ($i = 1, 2, \dots; j = 1, 2, \dots$).

The *local communication state* (for short: a *local state*) of the agent A_p is a function $M_p : \{p\} \rightarrow C_p \cup \{\Theta\}$.

The set of all states of the agent A_p is $\mathbb{S}_p = (C_p \cup \{\Theta\})^{\{p\}}$

Semantics of transition $t \in T_p$ is a relation $[[t]] \subseteq \mathbb{S}_p \times \mathbb{S}_p$ defined by $(M_p, M'_p) \in [[t]]$ iff $M_p(p) = \Theta \wedge M'_p(p) = F_p(t, p)$ (M'_p is the next state following M_p obtained in effect of firing transition t)

Semantics of agent $A_p : [[A_p]] = \bigcup_{t \in T_p} [[t]]$

Fig.1 depicts agent A_p capable of communicating with agents $A_{p_1}, A_{p_2}, A_{p_3}, A_{p_4}, A_{p_5}$ and passing from the state $M_p = \{(p, \Theta)\}$ to the state $M'_p = \{(p, !(p_1, p_4))\}$




Fig. 1. Example of agent A_p and result of firing transition t_3

as a result of firing transition t_3 . This means that A_p issued a request for sending a message to A_{p_1} and A_{p_4} .

The *global communication state* (for short: a *global state*) of the system A is a function

$M : P \rightarrow C$, their set $S = C^P$, thus the local state M_p is a restriction of M to $\{p\}$. The state (global and local) will be treated as a set of pairs of the form $(p, !(p_{k_1}, p_{k_2}, \dots))$ and $(p, ?(p_{l_1}, p_{l_2}, \dots))$ for $p, p_{k_1}, p_{k_2}, p_{l_1}, p_{l_2}, \dots \in P$.

2.1 Transmissions

For $n, m \geq 1$, let a_1, \dots, a_n and b_1, \dots, b_m , pairwise distinct, be ports of agents A_{a_1}, \dots, A_{a_n} and A_{b_1}, \dots, A_{b_m} . Let $a_i :!(b_1, \dots, b_m)$ denote the pair $(a_i, !(b_1, \dots, b_m))$ meaning "agent A_{a_i} sends a message to agents A_{b_1}, \dots, A_{b_m} " and $b_j :?(a_1, \dots, a_n)$ the pair $(b_j, ?(a_1, \dots, a_n))$ meaning "agent A_{b_j} receives a message from agents A_{a_1}, \dots, A_{a_n} ". A *transmission* (matching send and receive operations) is a pair $\mathbf{t} = (\bullet \mathbf{t}, \mathbf{t} \bullet)$ of sets of the form:

$\bullet \mathbf{t} = \{a_1 :!(b_1, \dots, b_m), \dots, a_n :!(b_1, \dots, b_m)\}$ (pre-set of transmission \mathbf{t})

$\mathbf{t} \bullet = \{b_1 :?(a_1, \dots, a_n), \dots, b_m :?(a_1, \dots, a_n)\}$ (post-set of transmission \mathbf{t})

Let $\bullet \mathbf{t} \bullet = \bullet \mathbf{t} \cup \mathbf{t} \bullet$ and $\bullet \mathbf{t} \bullet \downarrow P$ be a projection of $\bullet \mathbf{t} \bullet$ onto the set P , i.e.

$\bullet \mathbf{t} \bullet \downarrow P = \{a_1, \dots, a_n, b_1, \dots, b_m\}$, that is, the set of ports the transmission \mathbf{t} is involved in. Note that $\bullet \mathbf{t} \bullet$ is of the same type as the global state M : both are sets of pairs of the form $(x, !(...))$ or $(x, ?(...))$.

Expressions $a_i :!(b_1, \dots, b_m)$ and $b_j :?(a_1, \dots, a_n)$ denote *matching labelled communication operations*.

Note that a transmission depends on a state: it may come into existence in a certain global state and disappear in another. Such emerging and disappearing during system's activity transmissions are typed in bold letters, to distinguish them from the static transitions of the agents.

Let \mathbf{TR} denote the set of all possible transmissions in the system. If a transmission $\mathbf{t} \in \mathbf{TR}$ exists in a state M (i.e. $\bullet \mathbf{t} \bullet \subseteq M$) then its semantics is a relation $[[\mathbf{t}]] \subseteq \mathbb{S} \times \mathbb{S}$ defined by $(M, M') \in [[\mathbf{t}]]$ iff $M' = M - \bullet \mathbf{t} \bullet \cup \{(x, \Theta) \mid x \in \bullet \mathbf{t} \bullet \downarrow P\}$. This means that M' is M in which all pairs $(x, !(\dots))$ and $(x, ?(\dots))$ belonging to $\bullet \mathbf{t} \bullet$ are replaced with pairs (x, Θ) , i.e. M' is the result of "firing" transmission \mathbf{t} at the state M . This models the transfer of a message from senders to receivers and disruption of the communication channel.

In Fig.2 a collection of 6 ports of agents $A_p, A_{p_1}, A_{p_2}, A_{p_3}, A_{p_4}, A_{p_5}$ are depicted. The global state of this system is

$$M = \begin{array}{|c|c|c|c|c|c|} \hline p & p_1 & p_2 & p_3 & p_4 & p_5 \\ \hline !(p_1, p_4) & ?(p) & \Theta & \Theta & ?(p) & \Theta \\ \hline \end{array}$$

Transmission $\mathbf{t} = (\{(p:!(p_1, p_4))\}, \{(p_1:?(p), (p_4:?(p)))\})$ transforms M

into $M' = \begin{array}{|c|c|c|c|c|c|} \hline p & p_1 & p_2 & p_3 & p_4 & p_5 \\ \hline \Theta & \Theta & \Theta & \Theta & \Theta & \Theta \\ \hline \end{array}$ in effect of sending

simultaneously a message from agent A_p to A_{p_1} and A_{p_4}




Fig. 2. Transmission of a message from agent A_p to A_{p_1} and A_{p_4} through channel \Rightarrow . Here, $\bullet \mathbf{t} \bullet \downarrow P = \{p, p_1, p_4\}$

2.2 Existence of transmissions

A transmission $\mathbf{t} \in \mathbf{TR}$ exists in a global state M of the system A iff $\bullet \mathbf{t} \bullet \subseteq M$. Given a global state M_0 and a transmission \mathbf{t} , the realizability of \mathbf{t} starting computation from M_0 is expressed as: does there exist a state M reachable from M_0 such that $\bullet \mathbf{t} \bullet \subseteq M$? Thus, the existence problem for \mathbf{t} reduces to some versions of state reachability and inclusion problems. Their solution in the

form of yes/no decision and, possibly, their complexity, depends obviously on the formal description of agents. For example, let us assume that agents are described as 1-safe finite Petri nets (places are valued in the set $\{0,1\}$) obtained by replacing labels of arrows entering ports by weights 1. Then, the existence of \mathbf{t} reduces to the problem "For a given marking M_0 and place p , is there a reachable from M_0 marking with a token in p ?", which is known to be the PSPACE-hard (PSPACE - the set of all decision problems solvable by a Turing machine with a polynomial amount of space), see e.g. [ESP 1998]. Indeed, after the replacement of arrow labels, the whole system becomes one disconnected 1-safe Petri net. Denote by m_0 the marking of it, such that each place (port) $x \in \bullet \mathbf{t} \bullet \downarrow P$ holds a token (i.e. $m_0(x) = 1$) iff $M_0(x) \neq \emptyset$. In such system net each $x \in \bullet \mathbf{t} \bullet \downarrow P$ has no outgoing arrow, thus, if a token enters this place at a certain marking reachable from m_0 , then it will stay there indefinitely. Now, decide if there is a marking m reachable from m_0 and satisfying $m(x) = 1$ for all $x \in \bullet \mathbf{t} \bullet \downarrow P$. If yes (and only if), then in the original system (before replacement of the labels of arrows entering ports) there exists the transmission \mathbf{t} , because $\bullet \mathbf{t} \bullet \subseteq M$, where M is m restricted to ports $x \in \bullet \mathbf{t} \bullet \downarrow P$ holding communication operations $!(\dots)$ and $?(\dots)$ instead of tokens.

Note that the assumption on agents' internal (i.e. without communication) activity as specified by Petri nets, corresponds to the concept of self-modifying nets ([B-D 1997], [Val 1978], [Val 1981], [Cza 2013]). Indeed, transmissions are in fact a special kind of transitions appearing and disappearing, so the system changes its structure in the course of its performance.

3 Semantics of the System A and some PSPACE-hard Decision Problems of its Behaviour

Let $T = \bigcup_{p \in P} T_p$ and $F = \bigcup_{p \in P} F_p$ i.e. the set of all transitions and arrows in the system A respectively. The triple $A = (P, T, F)$, denoted also by A , is a *net representation* of the system. Its semantics is the union of semantics of the transitions $t \in T$ and message transmissions $\mathbf{t} \in \mathbf{TR}$: $[[A]] = \bigcup_{\tau \in T \cup \mathbf{TR}} [[\tau]]$. If $(M, M') \in [[A]]$ then M' is the next to M state evoked by a transition $t \in T$ or a transmission $\mathbf{t} \in \mathbf{TR}$. For $\tau \in V = T \cup \mathbf{TR}$ denote $M \xrightarrow{\tau} M'$ iff $(M, M') \in [[\tau]]$. A *run* starting at M_0 is a chain $M_0 \xrightarrow{\tau_1} M_1 \xrightarrow{\tau_2} M_2 \xrightarrow{\tau_3} \dots$, finite or infinite, but if finite $M_0 \xrightarrow{\tau_1} M_1 \xrightarrow{\tau_2} M_2 \xrightarrow{\tau_3} \dots \xrightarrow{\tau_n} M_n$ then none M satisfies $(M_n, M) \in [[A]]$. A finite or infinite word $v = \tau_1 \tau_2 \tau_3 \dots \in V^\omega = V^* \cup V^\infty$ occurring at this run is a *path* starting at M_0 . If finite $v = \tau_1 \tau_2 \tau_3 \dots \tau_n \in V^*$ then $M \xrightarrow{v} M'$ denotes $M \xrightarrow{\tau_1} M_1 \xrightarrow{\tau_2} M_2 \xrightarrow{\tau_3} \dots \xrightarrow{\tau_n} M'$. The set of all finite and infinite runs starting at M is $RUN_*(M)$ and $RUN_\infty(M)$ respectively and $RUN(M) = RUN_*(M) \cup RUN_\infty(M)$. The set of respective paths is $PATH(M) = PATH_*(M) \cup PATH_\infty(M)$, thus $PATH(M) \subseteq V^\omega$.

Assuming, as above, that agents are described by 1-safe Petri nets obtained by replacing labels of arrows entering ports by weights 1, one can simulate behaviour

of the system by a 1-safe net as follows. Let a state M be given. For each transmission $t \in TR$ create a transition $t \notin T$ defined as $t = (\bullet t, t \bullet)$ with $\bullet t = \bullet t \bullet \downarrow P$, $t \bullet = \emptyset$, and make arrows from ports $p \in \bullet t \bullet \downarrow P$ to t . The extended net is a triple $\bar{A} = (P, \bar{T}, \bar{F})$, where $\bar{T} = T \cup$ set of newly created transitions, and $\bar{F} = F \cup$ set of newly created arrows weighted with 1. A marking of \bar{A} is obtained from marking of A by replacing operations $!(\dots)$, $?(\dots)$ with tokens wherever such operations are in some ports and removing \emptyset from remaining ports. Fig.3 depicts a simulation of transmission t from Fig.2 by the newly created transition t and result of its firing. Remember: while t appears and disappears in the course of the system activity, the transition $t \in \bar{T}$ is the ordinary transition of the Petri net \bar{A} simulating system A , thus a unchangeable member of the \bar{A} 's static structure.




Fig. 3. Transition $t = (\{p, p_1, p_4\}, \emptyset)$ with empty post-set simulates behaviour of message transmission t in Fig.2

Some problems concerning behaviour of the system A may be reduced to problems concerning behaviour of the Petri net \bar{A} . To mention a few (suppose runs start from a given marking):

- a. Existence of run with a given message transmission occurrence
- b. Existence of reachable dead marking (no transition can fire at it)
- c. Existence of finite run (equivalent to b)
- d. Existence of infinite run
- e. Existence of run with infinite number of a given message transmission occurrence
- f. Existence of run with never accomplished a given request for communication

All these problems are PSPACE-hard for 1-safe Petri nets ([ESP 1998]) and \bar{A} is such net. Therefore, by virtue of the obviously polynomial simulation procedure described above, the problems for systems specified like A in this paper, are PSPACE-hard provided that internal activity of agents is specified by 1-safe Petri nets.

4 Deadlock and Fairness: Emptiness and Finiteness of Sets of Paths

Out of several concepts and kinds of deadlock and fairness found in diverse models of distributed computing, let us consider those arising from communication and described in terms of the model pursued here.

4.1 Deadlock

System A is deadlock-free at a state M if for each agent requesting for communication there is a finite path starting at M , such that the agent will be permitted to accomplish the request on this path. A deadlock is a negation of this property. For an agent $A_p \in A = (P, T, F)$, with port p and for a state $M \in \mathbb{S}$ define:

$$D_p(M) \stackrel{def}{\iff} \neg[\exists M'. \exists v. (M \xrightarrow{v} M' \wedge \exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet))]$$

where $t \in^* v$ means "transmission $t \in \mathbf{TR}$ occurs on the path v " and $p:M(p) \in \bullet t \bullet$ means "t accomplishes request for communication issued by agent A_p and pending at the state M ".

In words: never agent A_p requesting for communication at the state M will be permitted to accomplish the request by a certain transmission occurring on whichever finite path starting at M .

The system is subject to a deadlock at the state M iff:

$$\exists p. M(p) \neq \emptyset \wedge D_p(M).$$

Proposition 4.1.1 (set-theoretic characterization)

$D_p(M)$ if and only if $PATH(M) \cap V^* t V^* = \emptyset$ for each t satisfying $p:M(p) \in \bullet t \bullet$

Proof

$$\begin{aligned} D_p(M) &\iff \\ \neg \exists M'. \exists v. (M \xrightarrow{v} M' \wedge \exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet)) &\iff \text{(swapping quantifiers)} \\ \neg \exists v. \exists M'. (M \xrightarrow{v} M' \wedge \exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet)) &\iff \text{(De Morgan law)} \\ \forall v. \neg \exists M'. (M \xrightarrow{v} M' \wedge \underbrace{\exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet)}_{\text{no } M' \text{ in this formula}}) &\iff \\ \forall v. \neg((\exists M'. M \xrightarrow{v} M') \wedge \exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet)) &\iff \text{(De Morgan law)} \\ \forall v. (\neg(\exists M'. M \xrightarrow{v} M') \vee \neg \exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet)) &\iff \\ \forall v. ((\exists M'. M \xrightarrow{v} M') \Rightarrow \neg \exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet)) &\iff \\ \forall v. (v \in \{u \mid \exists M'. M \xrightarrow{u} M'\} \Rightarrow \neg \exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet)) &\iff \text{(definition of } PATH(M)) \\ \forall v. (v \in PATH(M) \Rightarrow \neg \exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet)) &\iff \\ \{v \mid v \in PATH(M)\} \subseteq \{v \mid \neg \exists t. (t \in^* v \wedge p:M(p) \in \bullet t \bullet)\} &\iff \end{aligned}$$

$$\begin{aligned} PATH(M) &\subseteq \{v \mid \neg \exists t. (t \overset{*}{\in} v \wedge p:M(p) \in \bullet t \bullet)\} \iff \\ PATH(M) &\subseteq V^* - \{v \mid \exists t. (t \overset{*}{\in} v \wedge p:M(p) \in \bullet t \bullet)\} \end{aligned}$$

Therefore:

$PATH(M) \subseteq V^* - V^*tV^*$ for each t satisfying $p:M(p) \in \bullet t \bullet$ where V^*tV^* is the set of all finite words over V where t occurs. Thus:

if $p:M(p) \in \bullet t \bullet$ then $PATH(M) - (V^* - V^*tV^*) = \emptyset$.

Since $X - (Y - Z) = (X - Y) \cup (X \cap Z)$ for any sets X, Y, Z then

$$PATH(M) - (V^* - V^*tV^*) = (PATH(M) - V^*) \cup (PATH(M) \cap V^*tV^*) = \emptyset$$

(Because $PATH(M) - V^* = \emptyset$). Finally:

$$D_p(M) \text{ iff } \forall t. (p:M(p) \in \bullet t \bullet \Rightarrow PATH(M) \cap V^*tV^* = \emptyset) \quad \square$$

Theorem 4.1.1

A deadlock at a state M occurs iff:

$$\exists p. [M(p) \neq \emptyset \wedge (\forall t. (p:M(p) \in \bullet t \bullet \Rightarrow PATH(M) \cap V^*tV^* = \emptyset))] \quad \square$$

Thus decidability of such deadlocks reduces to deciding whether transmission t does not occur on any path starting from M (provided that there are a finite number of agents, thus also transmissions), which depends on algebraic structure of the set $PATH(M)$.

4.2 Weak fairness

System A is *weakly* fair at a state M if each agent requesting for communication at M will be permitted to accomplish the request on every infinite path starting from M . This is expressed by the formula:

$$\forall p. [M(p) \neq \emptyset \Rightarrow \forall v. (v \in PATH_\infty(M) \Rightarrow \exists t. (t \overset{*}{\in} v \wedge p:M(p) \in \bullet t \bullet))]$$

Theorem 4.2.1

System A is weakly fair at a state M iff

$$\forall p. [M(p) \neq \emptyset \Rightarrow (\forall t. (p:M(p) \in \bullet t \bullet \Rightarrow PATH_\infty(M) - V^*tV^\infty = \emptyset))]$$

Proof

$$\forall v. (v \in PATH_\infty(M) \Rightarrow \exists t. (t \overset{*}{\in} v \wedge p:M(p) \in \bullet t \bullet)) \iff$$

$$\{v \mid v \in PATH_\infty(M)\} \subseteq \{v \mid \exists t. (t \overset{*}{\in} v \wedge p:M(p) \in \bullet t \bullet)\} \iff$$

$PATH_\infty(M) \subseteq V^*tV^\infty$ for each t satisfying $p:M(p) \in \bullet t \bullet$. Thus

$$\forall t. (p:M(p) \in \bullet t \bullet \Rightarrow PATH_\infty(M) - V^*tV^\infty = \emptyset) \quad \square$$

4.3 Strong fairness

System A is *strongly fair* at a state M if each agent requesting for communication at M will be permitted to accomplish the request on every finite path starting at M if all these paths are "sufficiently long", i.e. of the length at least k , for a certain k . So, all these paths may be jointly ("uniformly") bounded in length. This is expressed by the formula:

$$\forall p.[M(p) \neq \emptyset \Rightarrow \exists k.F_p(M, k)] \quad \text{where}$$

$$F_p(M, k) \stackrel{\text{def}}{\iff} \forall v.((v \in \text{PATH}_*(M) \wedge |v| \geq k) \Rightarrow \exists \mathbf{t}.(\mathbf{t} \in v \wedge p:M(p) \in \bullet \mathbf{t} \bullet))$$

Theorem 4.3.1

System A is strongly fair at a state M iff

$$\forall p.[M(p) \neq \emptyset \Rightarrow (\forall \mathbf{t}.(p:M(p) \in \bullet \mathbf{t} \bullet \Rightarrow |\text{PATH}_*(M) - V^* \mathbf{t} V^*| < \infty))]$$

Proof

$$F_p(M, k) \iff$$

$$\{v | v \in \text{PATH}_*(M)\} \cap \{v | |v| \geq k\} \subseteq \{v | \exists \mathbf{t}.(\mathbf{t} \in v \wedge p:M(p) \in \bullet \mathbf{t} \bullet)\} \iff$$

$$\text{PATH}_*(M) \cap V^k V^* \subseteq V^* \mathbf{t} V^* \quad \text{for each } \mathbf{t} \text{ satisfying } p:M(p) \in \bullet \mathbf{t} \bullet$$

where $V^k V^*$ is the set of finite words of the length at least k and $V^* \mathbf{t} V^*$ the set of finite words where \mathbf{t} occurs. Thus:

$$\forall \mathbf{t}.(p:M(p) \in \bullet \mathbf{t} \bullet \Rightarrow \text{PATH}_*(M) \cap V^k V^* - V^* \mathbf{t} V^* = \emptyset).$$

Now, let $H_p(M, k) = \text{PATH}_*(M) \cap V^k V^* - V^* \mathbf{t} V^*$ for each \mathbf{t} satisfying $p:M(p) \in \bullet \mathbf{t} \bullet$.

We show that $|H_p(M, 0)| < \infty$ if and only if $\exists k.H_p(M, k) = \emptyset$

($|X|$ is cardinality of the set X).

(\Rightarrow) Let $\forall k.H_p(M, k) \neq \emptyset$. Then $\lambda(H_p(M, k)) < \lambda(H_p(M, k+1))$ where $\lambda(L)$ is the length of a shortest word in the set L .

Thus $\lim_{k \rightarrow \infty} \lambda(H_p(M, k)) = \infty$, which implies

(since $H_p(M, k) \supset H_p(M, k+1)$) that $H_p(M, 0)$ contains words of arbitrary length, hence $|H_p(M, 0)| = \infty$.

(\Leftarrow) Let $|H_p(M, 0)| = \infty$. Then $H_p(M, 0)$ contains words of arbitrary length, thus, for any k it contains a word w with $|w| \geq k$. Because $w \in V^k V^*$ and $H_p(M, k) = H_p(M, 0) \cap V^k V^*$ we have $w \in H_p(M, k)$ hence $H_p(M, k) \neq \emptyset$. \square

4.4 Equivalence of weak and strong fairness

To demonstrate the equivalence between the two kinds of fairness in the model considered here, let us recall a version of the:

König's Lemma [Kön 1927]:

Let Σ be a set and \mathcal{Y} a tree of the properties:

- the number of sons of every node in \mathcal{Y} is finite;
- for any $k \geq 0$ there is a finite branch b in \mathcal{Y} with $|b| \geq k$ and $b \subseteq \Sigma$.

Then there exists a infinite branch B in \mathcal{Y} with $|B| \subseteq \Sigma$.

Theorem 4.4.1

The weak and strong fairness are equivalent.

Proof

By the Theorem 4.2.1 and 4.3.1 it suffices to demonstrate that $|PATH_*(M) - V^*tV^*| < \infty \iff PATH_\infty(M) - V^*tV^\infty = \emptyset$ for each transmission t such that $p:M(p) \in \bullet t \bullet$ for every port p with $M(p) \neq \emptyset$. Implication " \Rightarrow " is evident, it remains to show " \Leftarrow ". Suppose $|PATH_*(M) - V^*tV^*| = \infty$. Note that the set of paths starting at M is prefix-closed: each prefix of $v \in PATH(M)$ belongs to $PATH(M)$. To each v assign a unique element $node(v)$ in this way that $v_1 \neq v_2 \Rightarrow node(v_1) \neq node(v_2)$ and let $NODE(M) = \{node(v) | v \in PATH_*(M)\}$. This set with internodal relation defined by " $node(u)$ is father of $node(v)$ iff $v = u\tau$ for a certain $\tau \in V$ " is a tree \mathcal{Y} with $node(\varepsilon)$ as the root (ε is the empty path) and finitely many sons of each father. So, every path $v \in PATH_*(M)$ is a branch in \mathcal{Y} . By assumption $|PATH_*(M) - V^*tV^*| = \infty$ there are infinitely many finite paths, thus branches in \mathcal{Y} on which no t exists. Therefore there must be an arbitrarily long branch in the tree. Setting $\Sigma = PATH_*(M)$ and applying the König's Lemma, we come to contradiction. \square

Summing up the results obtained above, the set-theoretic characteristics of the deadlock and fairness at a state M are in the following table:

Deadlock	$PATH(M) \cap V^*tV^* = \emptyset$
Weak fairness	$PATH_\infty(M) - V^*tV^\infty = \emptyset$
Strong fairness	$ PATH_*(M) - V^*tV^* < \infty$

for every transmission t .

5 Counting States

If the agents do not send and receive messages to/from themselves then the total number of (global) states of n -agent system is $(2^n - 1)^n$. Indeed, each agent may issue $\frac{2^n}{2} - 1$ send $!(\dots)$ requests and the same number of receive $?(\dots)$ requests, that is $2^n - 2$ requests for communication. Since the agent may assume \emptyset as its local state, the number of local states it may assume is $2^n - 1$. The set of global states is the Cartesian product of sets of the local states of all agents. Therefore the number of global states is $\underbrace{(2^n - 1) \cdot \dots \cdot (2^n - 1)}_{n \text{ times}} = (2^n - 1)^n$. For instance,

for agents p_1, p_2, p_3 :

the set of local states of $p_1 = \{\emptyset, !(p_2), ?(p_2), !(p_3), ?(p_3), !(p_2, p_3), ?(p_2, p_3)\}$

the set of local states of $p_2 = \{\emptyset, !(p_1), ?(p_1), !(p_3), ?(p_3), !(p_1, p_3), ?(p_1, p_3)\}$

the set of local states of $p_3 = \{\emptyset, !(p_1), ?(p_1), !(p_2), ?(p_2), !(p_1, p_2), ?(p_1, p_2)\}$

Thus, the system of three agents has $7^3 = 343$ global states.

References

- [B-D 1997] Badouel E., Darondeau P., *Stratified Petri Nets*, FCT'97, Lecture Notes in Computer Science vol. 1279 (1979), pp. 117-128
- [Cza 2013] Czaja L., *Self-Modifying Nets for Synchronous, Connection-Oriented, Multicast Communication*, Fundamenta Informaticae, to appear
- [ESP 1998] Esparza J., *Decidability and Complexity of Petri Net Problems - An Introduction*, Lecture Notes in Computer Science, Vol. 1491, 1998, pp. 374-428
- [Hoa 1978] Hoare C.A.R., *Communicating Sequential Processes*, Comm. of the ACM, Vol 21, pp. 666-677, 1978
- [Hoa 1985] Hoare C.A.R., *Communicating Sequential Processes*, Prentice-Hall, 1985
- [Kön 1927] König D., *Über eine Schlussweise aus dem Endlichen in Uendliche*, Acta Litt. Ac. Sci. Hung. Fran. Josep. 3 (1927), pp. 121-130
- [OCCAM 1984] INMOS Limited: *OCCAM Programming Manual*, Prentice-Hall, 1984
- [Rei 1985] Reisig W., *Petri Nets, An Introduction*, EATC Monographs on Theoretical Computer Science, Springer-Verlag, 1985
- [Val 1978] Valk R., *Self-Modifying Nets, a Natural Extension of Petri Nets*, Icalp'78, Lecture Notes in Computer Science vol. 62 (1978), pp. 464-476
- [Val 1981] Valk R., *Generalization of Petri Nets*, MFCS'81 Lecture Notes in Computer Science vol. 118 (1981), pp. 140-155