

ASR : an introduction to
Distributed Systems and Algorithms
(Algorithmes et Systèmes Répartis)
3/3

M. Raynal¹, A. Mostefaoui², E. Fabre³

¹Pr., ASAP team

²MC, ASAP team

³CR INRIA, DistribCom team


Master 2, research in Computer Science, '07


Download the **lecture slides** from my web-page :
http://www.irisa.fr/distribcom/Personal_Pages/fabre/fabre.html


Computing with branching processes


- 1 Summary of last time
- 2 Computing with branching processes
 - Morphisms
 - Combinations of nets
 - Combinations of occurrence nets
 - Application to diagnosis
 - A central theorem
 - Distributed constraint solving
 - Distributed diagnosis
- 3 Open research issues

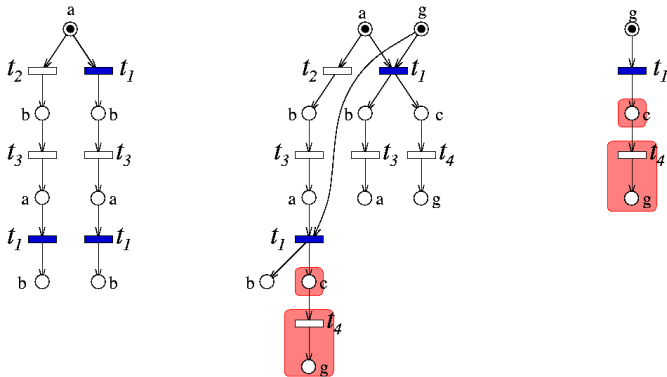
Example:



Example:

Example:




Example:

Example:

Pullback of BP:


- $O_1 \wedge^O O_2$ defined as $\mathcal{U}(O_1 \wedge O_2)$
- computable by a similar algorithm

Example:



Application to diagnosis

- **observation**: let $\Lambda_o \subseteq \Lambda$ be the set of observable labels
- a hidden run κ_h of \mathcal{N} is produced
- one observes only $\lambda(\kappa_h) \cap \Lambda_o$
- these labels are collected either as a sequence or as a partial order Obs (i.e. an occurrence net)



Definition

Diagnosis = the set of runs of \mathcal{N} that explain the observations.

Obtained from $\mathcal{D} = \mathcal{U}_{\mathcal{N}} \times^O Obs$.



- the product synchronizes possible runs of \mathcal{N} with observations Obs
- the hidden run κ_h is present in \mathcal{D}
- not all configurations of \mathcal{D} explain entirely observations Obs they may only explain a prefix of it

A central theorem

Theorem (product is preserved by unfolding)

Let $\mathcal{N} = \times_{i \in I} \mathcal{N}_i$, then $\mathcal{U}_{\mathcal{N}} = \times_{i \in I}^{\circ} \mathcal{U}_{\mathcal{N}_i}$.

- as for languages, trace languages, ...
- in the same way, combinations by pullback are preserved by unfolding


Example:

Projection of BP:

- let $O \subseteq \mathcal{U}_N$ and $N = \times_{i \in I} N_i$
- from $\psi_i : \mathcal{U}_N \rightarrow \mathcal{U}_{N_i}$ we can define the **projection** of O on N_i

$$\Pi_i O \triangleq \psi_i(O) \subseteq \mathcal{U}_{N_i}$$

Example:



Distributed constraint solving


- There exists a general theory of distributed constraint solving.
- We specialize it here to PN, and to 2 components.

Theorem

Let $N = N_1 \wedge N_2$, with N_0 an *interface* automaton between N_1 and N_2 . Let $O = O_1 \wedge^O O_2$ be a BP of N . Then

$$\Pi_1(O) = O_1 \wedge^O \Pi_0(O_2)$$

$$\Pi_2(O) = \Pi_0(O_1) \wedge^O O_2$$



Meaning of these equations


$$\Pi_1(\mathcal{O}) = \mathcal{O}_1 \wedge^{\mathcal{O}} \Pi_0(\mathcal{O}_2)$$

- $\Pi_0(\mathcal{O}_2)$ is a **message** sent from \mathcal{N}_2 to \mathcal{N}_1 about its behaviors on the interface net \mathcal{N}_0 (and symmetrically)
- the computations involve small and **local** BP, \mathcal{O} never appears
- this message-passing principle extends to networks of components

Distributed diagnosis

Assumptions

- $\mathcal{N} = \mathcal{N}_1 \wedge \mathcal{N}_2$, interface net \mathcal{N}_0
- visible labels $\Lambda_{0,i} \subseteq \Lambda_i$ in each \mathcal{N}_i
- a hidden run κ_h is performed by \mathcal{N}
- $Obs_i = \lambda_i(\kappa_h) \cap \Lambda_{0,i}$
represents the observations collected on component \mathcal{N}_i



Problem

- compute the projection on each \mathcal{N}_i of the global explanations $\mathcal{D} = \mathcal{U}_{\mathcal{N}} \times^{\circ} Obs$ where $Obs = Obs_1 \times^{\circ} Obs_2$
- recover runs of each component that contribute to explaining all observations (local views of the global diagnosis)

Formally

$$\begin{aligned} \mathcal{D} &= (\mathcal{U}_{\mathcal{N}_1} \times^{\circ} Obs_1) \wedge^{\circ} (\mathcal{U}_{\mathcal{N}_2} \times^{\circ} Obs_2) \\ \mathcal{D}_1 &= \Pi_1(\mathcal{D}) \\ &= (\mathcal{U}_{\mathcal{N}_1} \times^{\circ} Obs_1) \wedge^{\circ} \Pi_0(\mathcal{U}_{\mathcal{N}_2} \times^{\circ} Obs_2) \\ \mathcal{D} &= \mathcal{D}_1 \wedge^{\circ} \mathcal{D}_2 \end{aligned}$$

Problem

- compute the projection on each \mathcal{N}_i of the global explanations $\mathcal{D} = \mathcal{U}_{\mathcal{N}} \times^{\circ} Obs$ where $Obs = Obs_1 \times^{\circ} Obs_2$
- recover runs of each component that contribute to explaining all observations (local views of the global diagnosis)

Formally

$$\begin{aligned} \mathcal{D} &= (\mathcal{U}_{\mathcal{N}_1} \times^{\circ} Obs_1) \wedge^{\circ} (\mathcal{U}_{\mathcal{N}_2} \times^{\circ} Obs_2) \\ \mathcal{D}_1 &= \Pi_1(\mathcal{D}) \\ &= (\mathcal{U}_{\mathcal{N}_1} \times^{\circ} Obs_1) \wedge^{\circ} \Pi_0(\mathcal{U}_{\mathcal{N}_2} \times^{\circ} Obs_2) \\ \mathcal{D} &= \mathcal{D}_1 \wedge^{\circ} \mathcal{D}_2 \end{aligned}$$

Some open research issues

- 1 compute a finite complete prefix in factorized form, with a distributed algorithm
- 2 distributed reachability test
- 3 distributed deadlock detection
- 4 computations based on other structures than BP
- 5 coupling distributed computations with distributed optimization, to find an optimal diagnosis
- 6 relations to distributed control
- 7 ...