

Technischer Bericht Nr. 2008-2

Unfolding-based Diagnosis of Systems with an Evolving Topology

Paolo Baldan

Thomas Chatain
Barbara König

Stefan Haar

2.7.2008

IMPRESSUM:

Technische Berichte der Abteilung für Informatik und Angewandte
Kognitionswissenschaft, Universität Duisburg-Essen

ISSN 1863-8554

Herausgeber:

Abteilung für Informatik und Angewandte Kognitionswissenschaft
Fakultät für Ingenieurwissenschaften
Universität Duisburg-Essen
Campus Duisburg
47048 Duisburg

<http://duepublico.uni-duisburg-essen.de/informatik/berichte.xml>

Unfolding-based Diagnosis of Systems with an Evolving Topology^{*}

Paolo Baldan¹, Thomas Chatain², Stefan Haar³, and Barbara König⁴

¹ Dipartimento di Matematica Pura e Applicata, Università di Padova, Italy

² LSV, ENS Cachan, CNRS, INRIA, France

³ INRIA, France, and University of Ottawa, Canada

⁴ Abteilung für Informatik und Angewandte Kognitionswissenschaft, Universität
Duisburg-Essen, Germany

Abstract. We propose a framework for model-based diagnosis of systems with mobility and variable topologies, modelled as graph transformation systems. Generally speaking, model-based diagnosis is aimed at constructing explanations of observed faulty behaviours on the basis of a given model of the system. Since the number of possible explanations may be huge we exploit the unfolding as a compact data structure to store them, along the lines of previous work dealing with Petri net models. Given a model of a system and an observation, the explanations can be constructed by unfolding the model constrained by the observation, and then removing incomplete explanations in a pruning phase. The theory is formalised in a general categorical setting: constraining the system by the observation corresponds to taking a product in the chosen category of graph grammars, so that the correctness of the procedure can be proved by using the fact that the unfolding is a right adjoint and thus it preserves products. The theory thus should be easily applicable to a wide class of system models, including graph grammars and Petri nets.

1 Introduction

The *event-oriented model-based diagnosis* problem is a classical topic in discrete event systems [9, 17]. Given an observed alarm stream, the aim is to provide *explanations* in terms of actual system behaviours. Some events of the system are observable (alarms) while others are not. In particular, fault events are usually unobservable; therefore, fault diagnosis is the main motivation of the diagnosis problem. Given a sequence (or partially ordered set) of observable events, the diagnoser has to find all possible behaviours of the model explaining the observation, thus allowing the deduction of invisible causes (faults) of visible events (alarms). The paper [18] provides a survey on fault diagnosis in this direction.

Since the number of possible explanations may be huge, especially in the case of highly concurrent systems, it is advisable to employ space-saving methods. In [18, 12], the global diagnosis is obtained as the fusion of local decisions: this

^{*} Supported by the MIUR Project ART, RNRT project SWAN, INRIA Sabbatical program, the DFG project SANDS and CRUI/DAAD VIGONI.

distributed approach allows one to factor explanations over a set of local observers and diagnoses, rather than centralizing the data storage and handling.

We will build here upon the approach of [6] where diagnoses are stored in the form of unfoldings. The unfolding of a system fully describes its concurrent behaviour in a single branching structure, representing all the possible computation steps and their mutual dependencies, as well as all reachable states; the effectiveness of the approach lies in the use of partially ordered runs, rather than interleavings, to store and handle explanations extracted from the system model.

While [6] and subsequent work in this direction was mainly directed to Petri nets, here we face the diagnosis problem in mobile and variable topologies. This requires the development of a model-based diagnosis approach which applies to other, more expressive, formalisms. Unfoldings of extensions of Petri nets where the topology may change dynamically were studied in [10, 8]. Here we focus on the general and well-established formalism of graph transformation systems.

In order to retain only the behaviour of the system that matches the observations, it is not the model itself that is unfolded, but the product of the model with the observations, which represents the original system constrained by the observation; under suitable observability assumptions, a finite prefix of the unfolding is sufficient. The construction is carried out in a suitably defined category of graph grammars, where such a product can be shown to be the categorical product. A further *pruning* phase is necessary in order to remove incomplete explanations that are only valid for a prefix of the observations.

We show the correctness of this technique, i.e., we show that the runs of the unfolding properly capture all those runs of the model which explain the observation. This non-trivial result is obtained by using the fact that unfolding for graph grammars is a coreflection, hence it preserves limits (and especially products, such as the product of the model and the observation). In order to ensure that the product is really a categorical product, special care has to be taken in the definition of the category.

The rest of the paper is structured as follows. In Section 2 we introduce the category of graph grammars we will deal with and we characterise the corresponding categorical product. Then we show how the unfolding semantics extends to this setting. In Section 3 we review interleaving structures, an event based semantic model which is helpful in proving the correctness of the diagnosis procedure. In Section 4 we formalise the diagnosis problem and we propose a technique for deriving a diagnosis, in form of an unfolding, from a given system model and a corresponding observation. The correctness of such procedure is proved in 5. Finally, in Section 6 we outline some directions of future research.

2 Graph Grammars and Grammar Morphisms

In this section we summarise the basics of graph rewriting in the *single-pushout* (SPO) approach [15]. We introduce a category of graph grammars, whose morphisms are a variation of those in [2] and we characterise the induced categorical product, which turns out to be adequate for expressing the notion of composition

needed in our diagnosis framework. Then we show that the unfolding semantics smoothly extends to this setting, arguing that the unfolding construction can still be characterised categorically as a universal construction. The proof relies on the results in [2]; this motivates our choice of the SPO approach as opposed to the more classical *double-pushout* (DPO) approach, for graph rewriting.

2.1 Graph Grammars and their Morphisms

Given a partial function $f : A \rightharpoonup B$ we write $f(a) \downarrow$ whenever f is defined on $a \in A$ and $f(a) \uparrow$ whenever it is undefined. We will denote by $\text{dom}(f)$ the *domain* of f , i.e., the set $\{a \in A \mid f(a) \downarrow\}$. Let $f, g : A \rightharpoonup B$ be two partial functions. We will write $f \leq g$ when $\text{dom}(f) \subseteq \text{dom}(g)$ and $f(x) = g(x)$ for all $x \in \text{dom}(f)$.

For a set A , we denote by A^* the set of sequences over A . Given $f : A \rightharpoonup B$, the symbol $f^* : A^* \rightarrow B^*$ denotes its extension to sequences defined by $f^*(a_1 \dots a_n) = f(a_1) \dots f(a_n)$, where it is intended that the elements on which f is undefined are “forgotten”. Specifically, $f^*(a_1 \dots a_n) = \varepsilon$ whenever $f(a_i) \uparrow$ for any $i \in \{1, \dots, n\}$. Instead, $f^\perp : A^* \rightharpoonup B^*$ denotes the *strict* extension of f to sequences, satisfying $f^\perp(a_1 \dots a_n) \uparrow$ whenever $f(a_i) \uparrow$ for some $i \in \{1, \dots, n\}$.

A (*hyper*)*graph* G is a tuple (N_G, E_G, c_G) , where N_G is a set of nodes, E_G is a set of edges and $c_G : E_G \rightarrow N_G^*$ is a connection function. Given a graph G we will write $x \in G$ to say that x is a node or edge in G , i.e., $x \in N_G \cup E_G$.

Definition 1 (partial graph morphism). A partial graph morphism $f : G \rightharpoonup H$ is a pair of partial functions $f = \langle f_N : N_G \rightharpoonup N_H, f_E : E_G \rightharpoonup E_H \rangle$ such that:

$$c_H \circ f_E \leq f_N^\perp \circ c_G \quad (*)$$

We denote by **PGraph** the category of hypergraphs and partial graph morphisms. A morphism is called *total* if both components are total, and the corresponding subcategory of **PGraph** is denoted by **Graph**.

Notice that, according to condition (*), if f is defined on an edge then it must be defined on all its adjacent nodes: this ensures that the domain of f is a well-formed graph. The inequality in condition (*) ensures that *any* subgraph of a graph G can be the domain of a partial morphism $f : G \rightharpoonup H$. Instead, the stronger (apparently natural) conditions $c_H \circ f_E = f_N^\perp \circ c_G$ would have imposed f to be defined over an edge whenever it is defined on all its adjacent nodes.

We will work with *typed graphs* [11, 16], which are graphs labelled over a structure that is itself a graph, called the *graph of types*.

Definition 2 (typed graph). Given a graph T , a typed graph G over T is a graph $|G|$, together with a total morphism $t_G : |G| \rightarrow T$. A partial morphism between T -typed graphs $f : G_1 \rightharpoonup G_2$ is a partial graph morphism $f : |G_1| \rightharpoonup |G_2|$ consistent with the typing, i.e., such that $t_{G_1} \geq t_{G_2} \circ f$. A typed graph G is called *injective* if the typing morphism t_G is injective. The category of T -typed graphs and partial typed graph morphisms is denoted by T -**PGraph**.

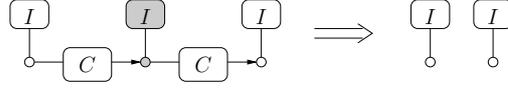


Fig. 1. Dangling edge removal in SPO rewriting.

Definition 3 (graph production, direct derivation). Fixing a graph T of types, a (T -typed graph) production q is an injective partial typed graph morphism $L_q \xrightarrow{r_q} R_q$. It is called consuming if r_q is not total. The typed graphs L_q and R_q are called left-hand side and right-hand side of the production.

Given a typed graph G and a match, i.e., a total injective morphism $g : L_q \rightarrow G$, we say that there is a direct derivation δ from G to H using q (based on g), written $\delta : G \Rightarrow_q H$, if there is a pushout square in $T\text{-PGraph}$ as on the right.

$$\begin{array}{ccc} L_q & \xrightarrow{r_q} & R_q \\ g \downarrow & & \downarrow h \\ G & \xrightarrow{d} & H \end{array}$$

Roughly speaking, the rewriting step removes from G the image of the items of the left-hand side which are not in the domain of r_q , namely $g(L_q - \text{dom}(r_q))$, adding the items of the right-hand side which are not in the image of r_q , namely $R_q - r_q(\text{dom}(r_q))$. The items in the image of $\text{dom}(r_q)$ are “preserved” by the rewriting step (intuitively, they are accessed in a “read-only” manner). Additionally, whenever a node is removed, all the edges incident to such a node are removed as well. For instance, consider production **fail** at the bottom of Fig. 2. Its left-hand side contains a unary edge (i.e., an edge connected to only one node) and its right-hand side is the empty graph. Nodes and edges are represented as circles and boxes, respectively. The application of **fail** to a graph is illustrated in Fig. 1, where the match of the left-hand side is indicated as shaded.

Definition 4 (typed graph grammar). A (T -typed) SPO graph grammar \mathcal{G} is a tuple $\langle T, G_s, P, \pi, A, \lambda \rangle$, where G_s is the (typed) start graph, P is a set of production names, π is a function which associates to each name $q \in P$ a production $\pi(q)$, and $\lambda : P \rightarrow A$ is a labelling over the set A . A graph grammar is consuming if all the productions in the range of π are consuming.

As standard in unfolding approaches, in the paper we will consider *consuming* graph grammars only, where each production deletes some item. Hereafter, when omitted, we will assume that the components of a given graph grammar \mathcal{G} are $\langle T, G_s, P, \pi, A, \lambda \rangle$. Subscripts carry over to the component names.

For a graph grammar \mathcal{G} we denote by $\text{Elem}(\mathcal{G})$ the set $N_T \cup E_T \cup P$. As a convention, for each production name q the corresponding production $\pi(q)$ will be $L_q \xrightarrow{r_q} R_q$. Without loss of generality, we will assume that the injective partial morphism r_q is a partial inclusion (i.e., that $r_q(x) = x$ whenever defined). Moreover we assume that the domain of r_q , which is a subgraph of both $|L_q|$ and $|R_q|$, is the *intersection* of these two graphs, i.e., that $|L_q| \cap |R_q| = \text{dom}(r_q)$, componentwise. Since in this paper we work only with typed notions, we will

usually omit the qualification “typed”, and, sometimes, we will not indicate explicitly the typing morphisms.

In the sequel we will often refer to the runs of a grammar defined as follows.

Definition 5 (runs of a grammar). *Let \mathcal{G} be a graph grammar. Then $Runs(\mathcal{G})$ contains all sequences $r_1 r_2 \dots r_n$ where $r_i \in P$ and $G_s \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \dots \xrightarrow{r_n} G_n$.*

Example. As a running example we will consider the graph grammar \mathcal{M} whose start graph and productions are given in Fig. 2. It models a network with mobility whose nodes are either senders (labelled S), receivers (R) or intermediary nodes (I). Senders may send messages which can then cross connections and should finally arrive at a receiver. However, a connection may be spontaneously corrupted, which causes the corruption of any message which crosses it. The network is variable and of unbounded size as we allow the creation of a new connection between existing intermediary nodes and the creation of a new connection leading from an existing intermediary node to a new intermediary node.

Productions (and the corresponding partial morphisms) are depicted as follows: edges that are deleted or created are drawn with solid lines, whereas edges that are preserved are indicated with dashed lines. Nodes which are preserved are indicated with numbers, whereas newly created nodes are not numbered.

Productions that should be observable (a notion that will be made formal in Section 4) are indicated by bold face letters.

We next define the class of grammars which will focus on in the development.

Definition 6 (semi-weighted SPO graph grammars). *A grammar \mathcal{G} is semi-weighted if (i) the start graph G_s is injective, and (ii) for each $q \in P$, for any x, y in $|R_q| - |L_q|$ if $t_{R_q}(x) = t_{R_q}(y)$ then $x = y$, i.e., the right-hand side graph R_q is injective on the “produced part” $|R_q| - |L_q|$.*

Intuitively, conditions (i) and (ii) ensure that in a semi-weighted grammar each item generated in a computation has a uniquely determined causal history, a fact which is essential for the validity of Theorem 16.

Note that grammar \mathcal{M} of Fig. 2 is not semi-weighted (if we assume the simplest type graph that contains one node and exactly one edge for every edge label). It could easily be converted into a semi-weighted grammar, for instance by creating the start graph (which is not injectively typed) step by step. However, for the sake of simplicity we do not carry out this construction in the paper.

We next introduce a category of graph grammars, which will be used to define products and to characterise the unfolding construction as a coreflection. The choice of arrows of the category is quite subtle: other possible notions of morphisms are conceivable but they do not necessarily provide the right notion of product and the coreflection result. In order to define grammar morphisms we need to introduce the notion of semi-abstract span, which, roughly speaking, provides a categorical generalisation of the notion of multirelation.

Given a category \mathbf{C} , a (concrete) span $f : A \leftrightarrow B$ in \mathbf{C} is a pair of total graph morphisms $f = \langle f^L : X_f \rightarrow A, f^R : X_f \rightarrow B \rangle$, where X_f is called the support.

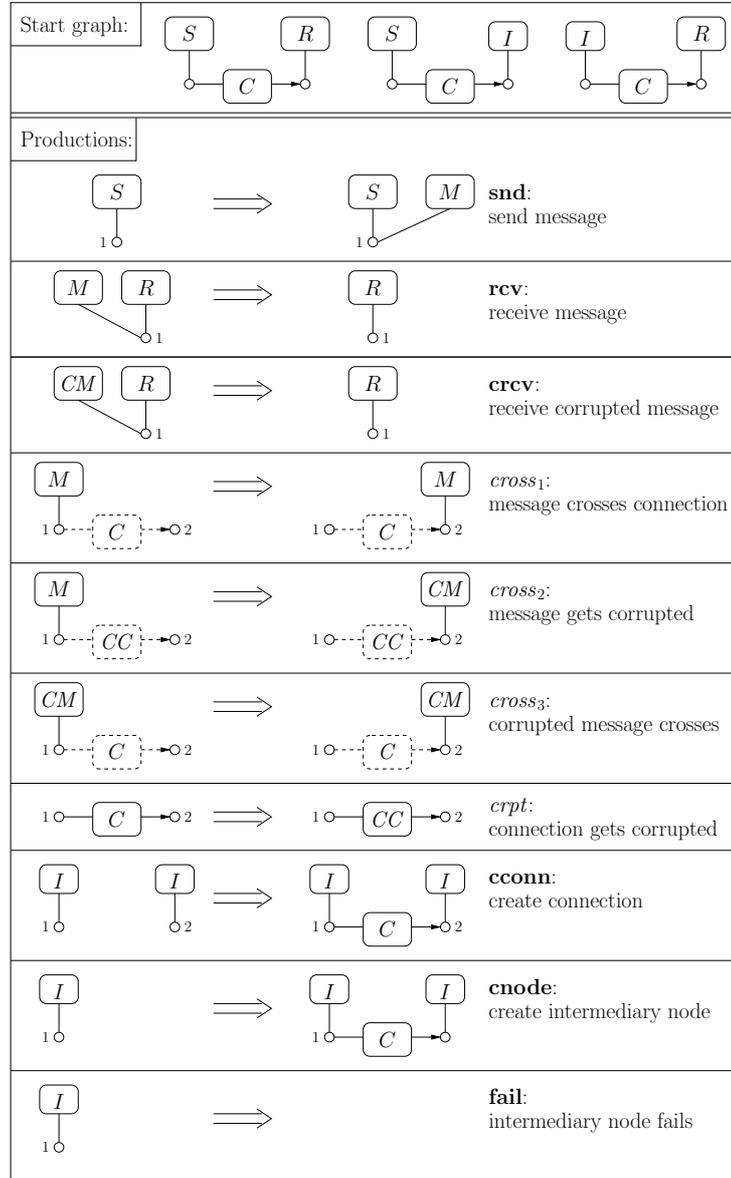


Fig. 2. Example grammar \mathcal{M} : message passing over possibly corrupted connections.

A *semi-abstract span* $[f] : A \leftrightarrow B$ is an equivalence class of spans obtained by considering the support up to isomorphism, i.e., $[f] = \{f' : A \leftrightarrow B \mid \exists k : k \text{ iso } \wedge X_f \rightarrow X_{f'}, f'^L \circ k = f^L \wedge f'^R \circ k = f^R\}$ (see Fig. 3.(a)).

If \mathbf{C} is a category with pullbacks, then semi-abstract spans can be composed as follows: given two semi-abstract spans $[f_1] : A \leftrightarrow B$ and $[f_2] : B \leftrightarrow C$, their composition is the (equivalence class of a) span f constructed as in Fig. 3.(b) (i.e., $f^L = f_1^L \circ y$ and $f^R = f_2^R \circ z$), where the square (1) is a pullback. This allows one to consider a category $\mathbf{Span}(\mathbf{C})$ has the same objects as \mathbf{C} and semi-abstract spans in \mathbf{C} as arrows.

The following definition generalises the notion of image of a set through a multirelation (see, e.g., [7]).

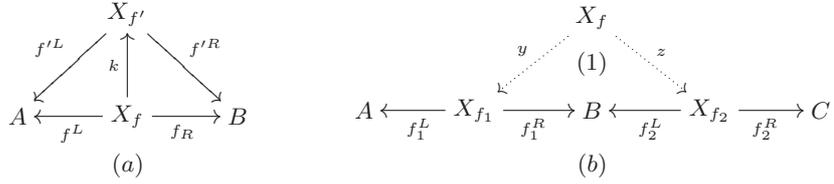
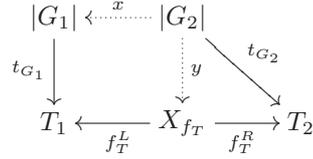


Fig. 3. Equivalence and composition of spans.

Definition 7 (pullback-retyping relation). Let $[f_T] : T_1 \leftrightarrow T_2$ be a semi-abstract span in \mathbf{Graph} , let G_1 be a T_1 -typed graph, and let G_2 be a T_2 -typed graph. Then G_1 and G_2 are related by pullback-retyping (via $[f_T]$) if there exist total morphisms $x : |G_2| \rightarrow |G_1|$ and $y : |G_2| \rightarrow X_{f_T}$ such that the square in the following diagram is a pullback:



In this case we will write $f_T\{x, y\}(G_1, G_2)$, or simply $f_T(G_1, G_2)$ if we are not interested in morphisms x and y .

We are now ready to introduce grammar morphisms. Except for the treatment of the labels, these morphisms coincide with those in [2]. The latter are, in turn, a generalisation of Winskel's morphisms for Petri nets (see [19]), which ensure the existence of products, which can be interpreted as asynchronous compositions, and of some coproducts, modelling nondeterministic choice [20].

Besides the component specifying the (multi)relation between the type graphs, a morphism from \mathcal{G}_1 to \mathcal{G}_2 includes a (partial) mapping between production

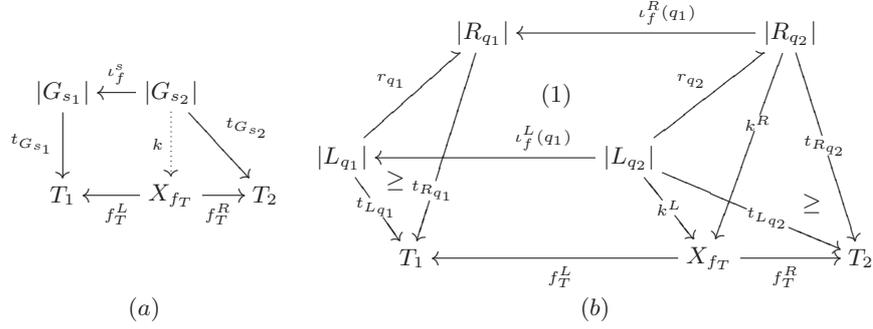


Fig. 4. Diagrams for SPO grammar morphisms.

names. Furthermore a third component explicitly relates the (untyped) graphs underlying corresponding productions of the two grammars, as well as the graphs underlying the start graphs.

Definition 8 (grammar morphism). Let \mathcal{G}_i ($i \in \{1, 2\}$) be graph grammars such that $\Lambda_2 \subseteq \Lambda_1$. A morphism $f : \mathcal{G}_1 \rightarrow \mathcal{G}_2$ is a triple $\langle [f_T], f_P, \iota_f \rangle$ where

- $[f_T] : T_1 \leftrightarrow T_2$ is a semi-abstract span in **Graph**, called the type-span;
- $f_P : P_1 \rightarrow P_2 \cup \{\emptyset\}$ is a total function, where \emptyset is a new production name (not in P_2), with associated production $\emptyset \rightarrow \emptyset$;
- ι_f is a family $\{\iota_f(q_1) \mid q_1 \in P_1\} \cup \{\iota_f^s\}$ of morphisms in **Graph** such that $\iota_f^s : |G_{s2}| \rightarrow |G_{s1}|$ and for each $q_1 \in P_1$, if $f_P(q_1) = q_2$, then $\iota_f(q_1)$ is a pair

$$\langle \iota_f^L(q_1) : |L_{q2}| \rightarrow |L_{q1}|, \iota_f^R(q_1) : |R_{q2}| \rightarrow |R_{q1}| \rangle.$$

such that the following conditions are satisfied:

1. Preservation of the start graph.
There exists a morphism k such that $f_T\{\iota_f^s, k\}(G_{s1}, G_{s2})$, i.e., the diagram in Fig. 4.(a) commutes and the square is a pullback.
2. Preservation of productions.
For each $q_1 \in P_1$, with $q_2 = f_P(q_1)$, there exist morphisms k^L and k^R such that the square (1) in Fig. 4.(b) commutes, and $f_T\{\iota_f^Y(q_1), k^Y\}(Y_{q1}, Y_{q2})$ for $Y \in \{L, R\}$.
3. Preservation of labelling.
For each $q_1 \in P_1$, $f_P(q_1) \neq \emptyset$ iff $\lambda_1(q_1) \in \Lambda_2$ and, in this case, $\lambda_2(f_P(q_1)) = \lambda_1(q_1)$.

For technical convenience, the partial mapping on production names is represented as a total mapping by enriching the target set with a distinguished element \emptyset , representing “undefinedness”. With respect to the morphisms in [2], note that here a morphism from \mathcal{G}_1 to \mathcal{G}_2 might exist only when $\Lambda_2 \subseteq \Lambda_1$ and there are some restrictions on the labelling as expressed by condition 3 above.

Definition 9 (category of graph grammars). We denote by \mathbf{GG} the category where objects are SPO graph grammars and arrows are grammar morphisms. By \mathbf{SGG} we denote the full subcategory of \mathbf{GG} having semi-weighted graph grammars as objects.

The choice of grammar morphisms and, in particular, the conditions on the labelling, lead to a categorical product suited for composing two grammars \mathcal{G}_1 and \mathcal{G}_2 : productions with labels in $\Lambda_1 \cap \Lambda_2$ are forced to be executed in a synchronous way, while the others are executed independently in the two components.

Proposition 10 (product of graph grammars). Let \mathcal{G}_1 and \mathcal{G}_2 be two graph grammars. Their product in $\mathbf{GG} = \mathcal{G}_1 \times \mathcal{G}_2$ is defined as follows. The product object \mathcal{G} has components:

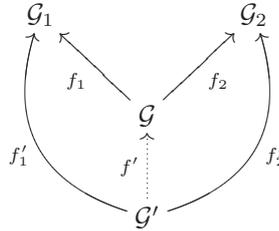
- $T = T_1 \uplus T_2$;
- $G_s = G_{s1} \uplus G_{s2}$, with the obvious typing;
- $P = \{(p_1, p_2) \mid \lambda_1(p_1) = \lambda_2(p_2)\} \cup \{(p_1, \emptyset) \mid \lambda_1(p_1) \notin \Lambda_2\} \cup \{(\emptyset, p_2) \mid \lambda_2(p_2) \notin \Lambda_1\}$;
- $\pi(p_1, p_2) = \pi_1(p_1) \uplus \pi_2(p_2)$, where $\pi_i(\emptyset)$ is the empty rule $\emptyset \rightarrow \emptyset$;
- $\Lambda = \Lambda_1 \cup \Lambda_2$;
- $\lambda(p_1, p_2) = \lambda_i(p_i)$, for any $i \in \{1, 2\}$ such that $p_i \neq \emptyset$;

where, p_1 and p_2 range over P_1 and P_2 , respectively, and disjoint unions are taken componentwise. The projections $f_i : \mathcal{G} \rightarrow \mathcal{G}_i$ are morphisms $f_i = \langle f_{iT}, f_{iP}, \iota^{f_i} \rangle$, where

- $f_{iT} : T_1 \uplus T_2 \leftrightarrow T_i$ has support T_i , the left leg $f_{iT}^L : T_i \rightarrow T_1 \uplus T_2$ is the obvious injection, while the right leg $f_{iT}^R : T_i \rightarrow T_1 \uplus T_2$ is the identity.
- $f_{iP} : P \rightarrow P_i \uplus \{\emptyset\}$ is the obvious projection
- concerning the components of ι_{f_i} , e.g., for the start graph $\iota_{f_i}^s : G_{si} \rightarrow G_{s1} \uplus G_{s2}$ is the obvious inclusion. The situation is analogous for $\iota_{f_i}^R$ and $\iota_{f_i}^L$.

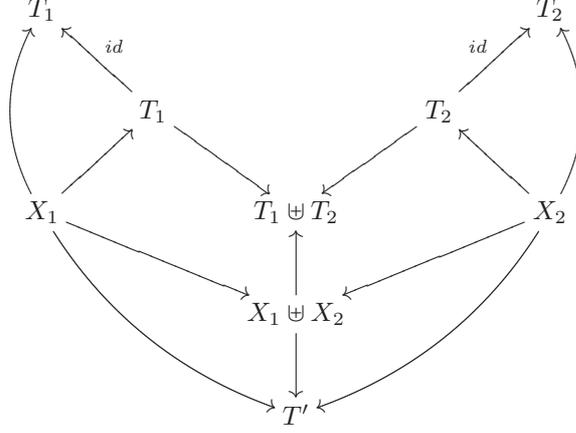
If $\mathcal{G}_1, \mathcal{G}_2$ are both semi-weighted grammars, then \mathcal{G} as defined above is semi-weighted, and it is the product of \mathcal{G}_1 and \mathcal{G}_2 in \mathbf{SGG} .

Proof. Let \mathcal{G}' be another graph grammar with morphisms $f'_i : \mathcal{G}' \rightarrow \mathcal{G}_i$. We have to show that there exists a unique morphism $f' : \mathcal{G}' \rightarrow \mathcal{G}$ with $f'_i \circ f' = f_i$.



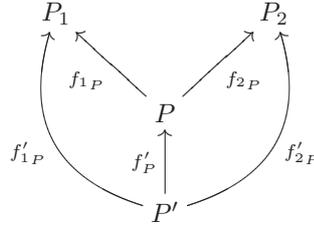
The morphism f' will be defined as follows:

- Assume that $f'_{i_T}: T' \leftrightarrow T_i$ has support X_i . Then for $f'_T: T' \leftrightarrow T_1 \uplus T_2$ take $X_1 \uplus X_2$ as support, where the arrows $X_1 \uplus X_2 \rightarrow T'$ and $X_1 \uplus X_2 \rightarrow T_1 \uplus T_2$ are obtained as mediating morphisms (see diagram below).



The category of graphs with total morphisms is an adhesive category with an initial object (the empty graph), which is hence extensive [14]. This implies that the two squares $X_i, T_i, X_1 \uplus X_2, T_1 \uplus T_2$ must be pullbacks. From this it follows that the composition of the spans f'_T and f_{i_T} equals f'_{i_T} . Furthermore the span f'_T is unique since taking $X_1 \uplus X_2$ as support is the only way to obtain two pullbacks and the arrows of the span are fixed as mediating morphisms.

- Now define $f'_P: P' \rightarrow P \uplus \{\emptyset\}$ as follows: $f'_P(q') = (f'_{1_P}(q'), f'_{2_P}(q'))$ if at least one of the two elements is different from \emptyset and $f'_P(q') = \emptyset$ otherwise. Clearly $f_{i_P} \circ f'_P = f'_{i_P}$. The situation is depicted in the diagram below:



It remains to show that f'_P is unique. Let $f''_P: P' \rightarrow P \uplus \{\emptyset\}$ another mapping with $f_{i_P} \circ f''_P = f'_{i_P}$ and let $q' \in P'$. It holds that $\lambda'(q') \in \Lambda' \supseteq \Lambda = \Lambda_1 \cup \Lambda_2$. Now there are four cases:

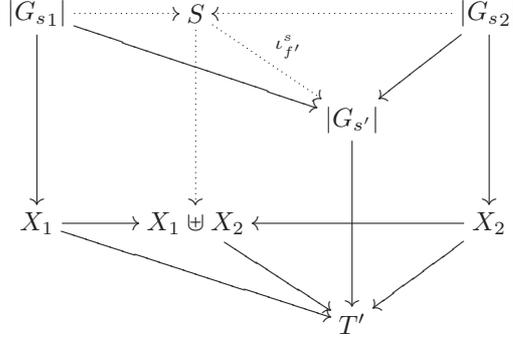
- $\lambda'(q') \in \Lambda_1 \cap \Lambda_2$: in this case $f'_{i_P}(q')$ must both be defined and since the f_{i_P} are simply projections it holds that $f''_P(q') = (f'_{1_P}(q'), f'_{2_P}(q')) = f'_P(q')$.
- $\lambda'(q') \in \Lambda_1 \setminus \Lambda_2$: in this case $f'_{1_P}(q')$ must be defined and $f'_{2_P}(q') = \emptyset$. Again, since the f_{i_P} are projections, it holds that $f''_P(q') = f'_P(q')$.

- $\lambda'(q') \in A_2 \setminus A_2$: analogous to the case above.
- $\lambda'(q') \notin A$: it holds that $f'_{iP}(q') = \emptyset$ for $i \in \{1, 2\}$. Since (\emptyset, \emptyset) is not contained in P , it must hold that $f''_P(q') = \emptyset = f'_P(q')$.

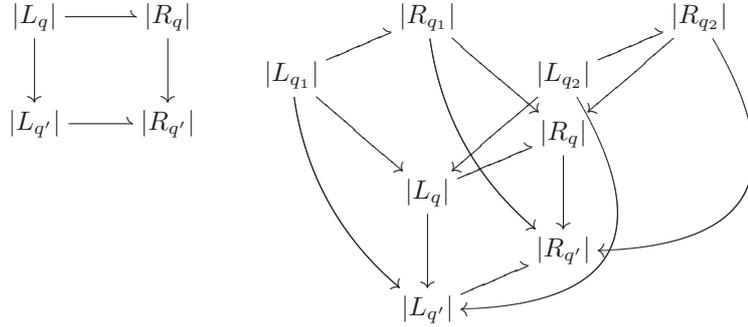
Note that for the argument above it is essential to require that a production cannot be mapped to \emptyset whenever its label is preserved.

- Finally, the components of $\iota_{f'}$ can be defined as follows: in the diagram below take the pullback of $X_1 \uplus X_2 \rightarrow T'$ and $|G_{s'}| \rightarrow T'$ in order to obtain S . The arrows $|G_{s_i}| \rightarrow S$ are obtained as mediating morphisms.

Since the squares $|G_{s_i}|, |G_{s'}|, X_i, T'$ are pullbacks, the squares $|G_{s_i}|, X_i, S, X_1 \uplus X_2$ are also pullbacks, due to pullback splitting. Hence—since we are working in an extensive category— S must be the coproduct of $|G_{s_1}|, |G_{s_2}|$ and is isomorphic to $|G_s|$. That is, we have a morphism $\iota_{f'}^s$, and have shown that it can be obtained from $|G_{s'}|$ via pullback-retyping. The arrow $\iota_{f'}^s$ must be unique since it is the mediating morphism of a coproduct.



Analogously one can define morphisms $\iota_{f'}^L: |L_q| \rightarrow |L_{q'}|, \iota_{f'}^R: |R_q| \rightarrow |R_{q'}|$ for productions $q' \in P', q \in P$ where $f'_P(q') = q$. It remains to show that the square on the left below commutes. Let $q_i = f_{1P}(q)$ and consider the diagram on the right below. It is known that all squares—apart from the bottom square—commute and that $|L_q|$ is the coproduct of $|L_{q_1}|$ and $|L_{q_2}|$. Hence the two coprojections $|L_{q_i}| \rightarrow |L_q|$ are jointly epi (also in the category of partial morphisms) and commutativity follows from a simple diagram-chasing argument.



Note also that whenever \mathcal{G}_1 and \mathcal{G}_2 are semi-weighted, then \mathcal{G} is semi-weighted. And since every arrow in \mathbf{GG} between semi-weighted grammars is also an arrow in \mathbf{SGG} , the same argument applies. \square

2.2 Occurrence Grammars and Unfolding

A grammar \mathcal{G} is *safe* if (i) for all H such that $G_s \Rightarrow^* H$, H is injective, and (ii) for each $q \in P$, the left- and right-hand side graphs L_q and R_q are injective.

In words, in a safe grammar each graph G reachable from the start graph is injectively typed, and thus we can identify it with the corresponding subgraph $t_G(|G|)$ of the type graph. With this identification, a production can only be applied to the subgraph of the type graph which is the image via the typing morphism of its left-hand side. Thus, according to its typing, we can think that a production *produces*, *preserves* or *consumes* items of the type graph, and using a net-like language, we speak of pre-set, context and post-set of a production, correspondingly. Intuitively the type graph T stands for the *places* of a net, whereas the productions P represent the *transitions*.

Definition 11 (pre-set, post-set and context of a production). *Let \mathcal{G} be a graph grammar. For any production $q \in P$ we define its pre-set $\bullet q$, context \underline{q} and post-set q^\bullet as the following subsets of $E_T \cup N_T$:*

$$\bullet q = t_{L_q}(|L_q| - |dom(r_q)|) \quad \underline{q} = t_{L_q}(|dom(r_q)|) \quad q^\bullet = t_{R_q}(|R_q| - r_q(|dom(r_q)|)).$$

Symmetrically, for each item $x \in T$ we define $\bullet x = \{q \in P \mid x \in \bullet q\}$, $x^\bullet = \{q \in P \mid x \in q^\bullet\}$, $\underline{x} = \{q \in P \mid x \in \underline{q}\}$.

Causal dependencies between productions are captured as follows.

Definition 12 (causality). *The causality relation of a grammar \mathcal{G} is the (least) transitive relation $<$ over $Elem(\mathcal{G})$ satisfying, for any node or edge $x \in T$, and for productions $q, q' \in P$,*

1. *if $x \in \bullet q$ then $x < q$;*
2. *if $x \in q^\bullet$ then $q < x$;*
3. *if $q^\bullet \cap \underline{q'} \neq \emptyset$ then $q < q'$.*

As usual \leq is the reflexive closure of $<$. Moreover, for $x \in Elem(\mathcal{G})$ we denote by $\lfloor x \rfloor$ the set of causes of x in P , namely $\{q \in P \mid q \leq x\}$.

As in Petri nets with read arcs, the fact that a production application not only consumes and produces, but also preserves a part of the state, leads to a form of asymmetric conflict between productions; for a thorough discussion of asymmetric event structures see [1].

Definition 13 (asymmetric conflict). *The asymmetric conflict relation of a grammar \mathcal{G} is the binary relation \nearrow over the set of productions, defined by:*

1. *if $\underline{q} \cap \bullet q' \neq \emptyset$ then $q \nearrow q'$;*

2. if $\bullet q \cap \bullet q' \neq \emptyset$ and $q \neq q'$ then $q \nearrow q'$;
3. if $q < q'$ then $q \nearrow q'$.

Intuitively, whenever $q \nearrow q'$, q can never follow q' in a computation. This holds when q preserves something deleted by q' (Condition 1), trivially when q and q' are in conflict (Condition 2) and also when $q < q'$ (Condition 3). Conflicts (in acyclic grammars) are represented by cycles of asymmetric conflict: if $q_1 \nearrow q_2 \nearrow \dots \nearrow q_n \nearrow q_1$ then $\{q_1, \dots, q_n\}$ will never appear in the same computation.

An *occurrence grammar* is an acyclic grammar which represents, in a branching structure, several possible computations beginning from its start graph and using each production at most once. Recall that a relation $R \subseteq X \times X$ is *finitary* if for any $x \in X$, the set $\{y \in X \mid R(y, x)\}$ is finite.

Definition 14 (occurrence grammar). An occurrence grammar is a safe grammar $\mathcal{O} = \langle T, G_s, P, \pi, \Lambda, \lambda \rangle$ such that

1. causality $<$ is irreflexive, its reflexive closure \leq is a partial order, and, for any $q \in P$, the set $[q]$ is finite and asymmetric conflict \nearrow is acyclic on $[q]$;
2. the start graph G_s is the set $\text{Min}(\mathcal{O})$ of minimal elements of $\langle \text{Elem}(\mathcal{O}), \leq \rangle$ (with the graphical structure inherited from T and typed by the inclusion);
3. any item x in T is created by at most one production in P , i.e., $|\bullet x| \leq 1$.

A finite occurrence grammar is deterministic if relation \nearrow^+ , the transitive closure of \nearrow , is irreflexive. We denote by **OGG** the full subcategory of **GG** with occurrence grammars as objects.

Note that the start graph of an occurrence grammar \mathcal{O} is determined by $\text{Min}(\mathcal{O})$. An occurrence grammar is deterministic if it does not contain conflicts (cycles of asymmetric conflict) so that all its productions can be executed in the same computation. In the sequel, the productions of an occurrence grammar will often be called events.

The notion of configuration captures the intuitive idea of (deterministic) computation in an occurrence grammar.

Definition 15 (configuration). Let $\mathcal{O} = \langle T, P, \pi \rangle$ be an occurrence grammar. A configuration is a subset $C \subseteq P$ satisfying the following requirements

1. for any $q \in C$ it holds $[q] \subseteq C$;
2. \nearrow_C , the asymmetric conflict restricted to C , is acyclic and finitary.

It can be shown that, indeed, all the productions in a configuration can be applied in a derivation exactly once in any order compatible with \nearrow .

Since occurrence grammars are particular semi-weighted grammars, there is an inclusion functor $\mathcal{I} : \mathbf{OGG} \rightarrow \mathbf{SGG}$. Such functor has a right adjoint. We remark that this theorem would not hold if we considered **GG** instead of **SGG** due to the failure of uniqueness properties.

Theorem 16. The inclusion functor $\mathcal{I} : \mathbf{OGG} \rightarrow \mathbf{SGG}$ has a right adjoint, the so-called unfolding functor $\mathcal{U} : \mathbf{SGG} \rightarrow \mathbf{OGG}$.

As a consequence of the above result \mathcal{U} , as a right adjoint, preserves all limits and in particular products.

The result is a corollary of [2], which, in turn, is obtained through the explicit definition of the unfolding $\mathcal{U}(\mathcal{G})$. Given a grammar \mathcal{G} the unfolding construction produces an occurrence grammar which fully describes its behaviour recording all the possible graph items which are generated and the occurrences of productions. The unfolding is obtained by starting from the start graph (as type graph), applying productions in any possible way, without deleting items but only generating new ones, and recording such production instances in the type graph. The result is an occurrence grammar $\mathcal{U}(\mathcal{G})$ and a grammar morphism $f : \mathcal{U}(\mathcal{G}) \rightarrow \mathcal{G}$, called the *folding morphism*, which maps each item (instance of production or graph item) of the unfolding to the corresponding item of the original grammar. The construction is not formally defined here. In Section 4 we will show an example of an unfolding.

As an immediate consequence of the fact that the unfolding $\mathcal{U}(\mathcal{G})$ completely captures the behaviour of a grammar \mathcal{G} , we have the following result.

Proposition 17. *For a graph grammar \mathcal{G} it holds that*

$$\lambda^*(Runs(\mathcal{U}(\mathcal{G}))) = \lambda^*(Runs(\mathcal{G})).$$

3 Interleaving Structures

Interleaving structures [3] are a semantic model which captures the behaviour of a system as the collection of its possible runs. They are used as a simpler intermediate model which helps in stating and proving the correctness of the diagnosis procedure.

An interleaving structure is essentially a collection of runs (sequences of events) satisfying suitable closure properties. Given a set E , we will denote by E° the set of sequences over E in which each element of E occurs at most once.

Definition 18 (interleaving structures). *A (labelled) interleaving structure is a tuple $\mathcal{I} = (E, R, \Lambda, \lambda)$ where E is a set of events, $\lambda : E \rightarrow \Lambda$ is a labelling of events and $R \subseteq E^\circ$ is the set of runs, satisfying: (i) R is prefix-closed, (ii) R contains the empty run ε , and (iii) every event $e \in E$ occurs in at least one run.*

The components of an interleaving structure \mathcal{I} will be denoted by E , R , Λ , λ , possibly with subscripts. The category of interleaving structures, as defined below, is adapted from [3] by changing the notion of morphisms in order to take into account the labels. This is needed to obtain a product which expresses a suitable form of synchronised composition.

Definition 19 (interleaving morphisms). *Let \mathcal{I}_i with $i \in \{1, 2\}$ be interleaving structures. An interleaving morphism from \mathcal{I}_1 to \mathcal{I}_2 is a partial function $\theta : E_1 \rightarrow E_2$ on events such that*

1. $\Lambda_2 \subseteq \Lambda_1$;

2. for each $e_1 \in E_1$, $\theta(e_1) \downarrow$ iff $\lambda_1(e_1) \in \Lambda_2$ and, in this case, $\lambda_2(\theta(e_1)) = \lambda_1(e_1)$;
3. for every $r \in R_1$ it holds that $\theta^*(r) \in R_2$.

Morphism θ is called a projection if θ is surjective on runs (as a function from R_1 to R_2). The category of interleaving structures and morphisms is denoted \mathbf{Ilv} .

An occurrence grammar can be easily mapped to an interleaving structure, by simply taking all the runs of the grammar.

Definition 20 (interleaving structures for occurrence grammars). For an occurrence grammar \mathcal{G} we denote by $Ilv(\mathcal{G})$ the interleaving structure which consists of all runs of \mathcal{G} , i.e., $Ilv(\mathcal{G}) = (P, Runs(\mathcal{G}), \Lambda, \lambda)$.

We next characterise the categorical product in \mathbf{Ilv} , which turns out to be, as in \mathbf{GG} , the desired form of synchronised product.

Proposition 21 (product of interleaving structures). Let \mathcal{I}_1 and \mathcal{I}_2 be two interleaving structures. Then the product object $\mathcal{I}_1 \times \mathcal{I}_2$ is the interleaving structure $\mathcal{I} = (E, R, \Lambda, \lambda)$ defined as follows. Let

$$E' = \{(e_1, e_2) \mid e_1 \in E_1, e_2 \in E_2, \lambda_1(e_1) = \lambda_2(e_2)\} \\ \cup \{(e_1, *) \mid e_1 \in E_1, \lambda_1(e_1) \notin \Lambda_2\} \cup \{(*, e_2) \mid e_2 \in E_2, \lambda_2(e_2) \notin \Lambda_1\}$$

and let $\pi_i : E \rightarrow E_i$ be the obvious (partial) projections. Then $R = \{r \in (E')^\odot \mid \pi_1^*(r) \in R_1, \pi_2^*(r) \in R_2\}$, $E = \{e' \in E' \mid e \text{ occurs in some run } r \in R\}$, $\Lambda = \Lambda_1 \cup \Lambda_2$ and λ is defined in the obvious way.

Proof. Let \mathcal{I}' be any interleaving structure and let $\theta_i : \mathcal{I}' \rightarrow \mathcal{I}_i$ be morphisms. Let us define $\theta : \mathcal{I}' \rightarrow \mathcal{I}$ as follows:

$$\theta(e') = \begin{cases} \uparrow & \text{if } \theta_1(e') \uparrow \text{ and } \theta_2(e') \uparrow \\ (\theta_1(e'), \theta_2(e')) & \text{otherwise} \end{cases}$$

where, abusing the notation, $\theta_i(e')$ must be read $*$ when $\theta_i(e')$ is undefined. Trivially, the diagram, seen in the category of sets and partial functions, commutes and θ is uniquely determined.

Hence to conclude we just need to show that θ is a well-defined interleaving morphism. In fact,

1. since there are morphisms $\theta_i : \mathcal{I}' \rightarrow \mathcal{I}_i$, necessarily $\Lambda_i \subseteq \Lambda'$ and thus $\Lambda = \Lambda_1 \cup \Lambda_2 \subseteq \Lambda'$.
2. for each $e' \in E'$, $\theta(e') \downarrow$ iff $\lambda'(e') \in \Lambda$, in fact
 - $\theta_1(e') \uparrow$ or $\theta_2(e') \uparrow$ iff
 - $\lambda'(e') \in \Lambda_1$ or $\lambda'(e') \in \Lambda_2$ iff
 - $\lambda'(e') \in \Lambda_1 \cup \Lambda_2 = \Lambda$.
 and clearly, when defined, $\lambda'(e') = \lambda(\theta(e'))$
3. for any $r' \in R'$, $\theta^*(r') \in R$
 Just notice that, due to commutativity, $\theta_i^*(\pi_i^*(r')) = \theta_i^*(r') \in R_i$. Hence, by construction, $\theta^*(r') \in R$. \square

4 Diagnosis and Pruning

In this section we use the tools introduced so far in order to formalise the diagnosis problem. Then we show how, given a graph grammar model and an observation for such a grammar, the diagnosis can be obtained by first taking the product of the model and the observation, considering its unfolding and finally pruning such unfolding in order to remove incomplete explanations. As already mentioned, typically only a subset of the productions in the system is observable. Hence, for this section, we fix a graph grammar \mathcal{G} with Λ as the set of labels, and a subset $\Lambda' \subseteq \Lambda$ of *observable labels*; an event or production is called *observable* if it has an observable label. In order to keep explanations finite, we will only consider systems that satisfy the following observability assumption (compare [17, 13]): any infinite run must contain an infinite number of observable productions.

In the sequel we will need to consider the runs of a system which have a number of observable events coinciding with the number of events in the observation. For this aim the following definition will be useful.

Definition 22 (*n*-runs of a grammar). *Let \mathcal{G} be a graph grammar. For a given $n \in \mathbb{N}$ we denote by $\text{Runs}^n(\mathcal{G})$ the set of all runs for which the number of observable productions equals n .*

The outcome of the diagnosis procedure is an occurrence grammar which, intuitively, collects all the behaviours of the grammar \mathcal{G} modelling the system, which are able to “explain” the observation.

An observation can be a sequence (in the case of a single observer) or a set of sequences (in the case of multiple distributed observers) of alarms (observable events). Here we consider, more generally, partially ordered sets of observations, which can be conveniently modelled as deterministic occurrence grammars \mathcal{O} .

Definition 23 (observation grammar). *An observation grammar for a given grammar \mathcal{G} , with observable labels Λ' , is a (finite) deterministic occurrence grammar labelled over Λ' .*

Given a sequence of observed events, we can easily construct an observation grammar \mathcal{O} having that sequence as observable behaviour. It will have a production for each event in the sequence (with the corresponding label). Each such production consumes a resource generated by the previous one in the sequence (or an initial resource in the case of the first production). The same construction applies to general partially ordered sets observations.

Example. In the example grammar \mathcal{M} (see Fig. 2), assume that we have the following observation: **snd₂ cconn crcv₂**, i.e., we observe, in sequence, the sending of a message, the creation of a connection and the reception of a corrupted message. These three observations can be represented by a simple grammar \mathcal{O} (see Fig. 5) with three productions, each of which either consumes an initial resource or a resource produced by the previous production. These resources are

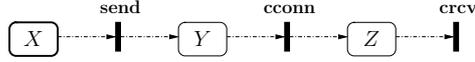


Fig. 5. A graph grammar representing an observation \mathcal{O} .

modeled as 0-ary edges (labelled X , Y , Z). The initial graph is depicted with bold lines, and the left- and right-hand sides of the productions of the occurrence grammar are indicated by a Petri-net-like notation: events are drawn with black rectangles connected to the respective edges by dashed lines.

When unfolding the product of a grammar \mathcal{G} with its observation \mathcal{O} , we obtain a grammar $\mathcal{U} = \mathcal{U}(\mathcal{G} \times \mathcal{O})$ with a morphism $\pi: \mathcal{U} \rightarrow \mathcal{O}$, arising as the image through the unfolding functor of the projection $\mathcal{G} \times \mathcal{O} \rightarrow \mathcal{O}$ (since the unfolding of an occurrence grammar is the grammar itself). Now, as grammar morphisms are simulations, given the morphism $\pi: \mathcal{U} \rightarrow \mathcal{O}$ we know that any computation in \mathcal{U} is mapped to a computation in \mathcal{O} . Say that a computation in \mathcal{U} is a full explanation of \mathcal{O} if it is mapped to a computation of \mathcal{O} including all its productions. As \mathcal{U} can still contain events belonging only to incomplete explanations, the aim of *pruning* is to remove such events.

Definition 24 (pruning). *Let $\pi: \mathcal{U} \rightarrow \mathcal{O}$ be a grammar morphism from an occurrence grammar \mathcal{U} to an observation \mathcal{O} . We define the pruning of π , denoted by $Pr(\pi)$, to be the grammar obtained from \mathcal{U} by removing all events (including their consequences) not belonging to the following set:*

$$\{q \in \mathcal{P}_{\mathcal{U}} \mid \exists C \in Conf(\mathcal{U}): (q \in C \wedge \pi(C) = \mathcal{P}_{\mathcal{O}})\}$$

The next technical lemma shows that applying the pruning operation to a morphism $\pi: \mathcal{U} \rightarrow \mathcal{O}$ no runs in \mathcal{U} which provide a full explanation of \mathcal{O} are lost.

Lemma 25 (interleavings of a pruned grammar). *Let $\pi: \mathcal{U} \rightarrow \mathcal{O}$ be a grammar morphism from an occurrence grammar to an observation with n productions. Then*

$$Runs^n(Pr(\pi)) = \{r \in Runs(\mathcal{U}) \mid \pi^*(r) \in Runs^n(\mathcal{O})\}.$$

Proof. We first show that $Runs^n(Pr(\pi)) \subseteq \{r \in Runs(\mathcal{U}) \mid \pi^*(r) \in Runs^n(\mathcal{O})\}$. Take any run $r \in Runs^n(Pr(\pi))$. Then clearly r is contained in $Runs(\mathcal{U})$ since $Pr(\pi)$ is a “subgrammar” of \mathcal{U} . Furthermore π preserves all observable events of r and hence $\pi^*(r)$ must contain exactly n observables and is hence an element of $Runs^n(\mathcal{O})$.

For the other direction take a run $r \in Runs(\mathcal{U})$ such that $\pi^*(r) \in Runs^n(\mathcal{O})$. Then, by the same argument as above, r must contain exactly n observables. Furthermore take the configuration C that is composed exactly of the events of r . Since C maps to the set of all productions of \mathcal{O} , none of the events of r is removed during the pruning phase. Hence $r \in Runs^n(Pr(\pi))$. \square

Discussing the efficiency of pruning algorithms is outside the scope of the paper; for sequential observations an on-the-fly algorithm is discussed in [6].

As described above, the diagnosis is constructed by first taking the product of \mathcal{G} with the observation (this intuitively represents the system constrained by the observation). This product is then unfolded to get an explicit representation of the possible behaviours explaining the observation. Finally, a pruning phase removes from the resulting occurrence grammar the events belonging (only) to incomplete explanations. This is formalised in the definition below.

Definition 26 (diagnosis grammar). *Let \mathcal{G} be the grammar modelling the system of interest and let \mathcal{O} be an observation. Take the product $\mathcal{G} \times \mathcal{O}$, the right projection $\varphi : \mathcal{G} \times \mathcal{O} \rightarrow \mathcal{O}$ and consider $\pi = \mathcal{U}(\varphi) : \mathcal{U}(\mathcal{G} \times \mathcal{O}) \rightarrow \mathcal{O}$.*

Then the occurrence grammar $Pr(\pi)$ is called the diagnosis grammar of the model and the observation and denoted by $D(\mathcal{G}, \mathcal{O})$.

Note that since the observability assumption holds, it can easily be shown that the diagnosis grammar is finite, whenever the observation is finite.

Example. We can compute the product of grammars \mathcal{M} and \mathcal{O} and unfold it. For reasons of space Fig. 6 shows only a prefix of the unfolding that depicts one possible explanation: here the message is sent (event **a**) and crosses the first connection (*b*). Possibly concurrently a new connection between the two intermediate nodes is created (**c**), which is then also crossed by the message (*d*). Again in a possibly concurrent step the last connection is corrupted (*e*), leading to the corruption of the message (*f*) and its reception by the receiver (**g**). Observable events are indicated by bold face letters.

Several events of the unfolding have been left out due to space constraints, for instance:

- Events belonging to alternative explanations: the corruption of the first connection or the corruption of the newly created middle connection (or the corruption of any non-empty subset of these connections). Alternatively it might also have been the case that the other sender/receiver pair handles the message, while the connection (which is not involved in any way) is created between the two intermediate nodes.
- Events that happen concurrently but are not directly related to the failure, such as the corruption of a connection over which no message is sent.

Furthermore there are events belonging to prefixes of the unfolding that cannot be extended to a full observation. For instance, the unfolding would contain concurrent events representing sending by the right-hand sender and the creation of a new connection leading from right to left instead of left to right. However, this is a false trail since this would never cause the reception of a message by the right-hand receiver. These incomplete explanations are removed from the unfolding in the pruning phase.

Note that—due to the presence of concurrent events—the unfolding is a much more compact representation of everything that might have happened in the system than the set of all possible interleavings of events.

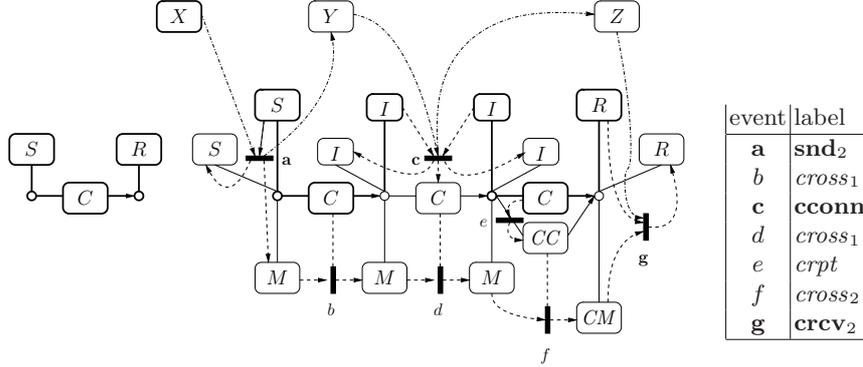


Fig. 6. Running example: prefix of the unfolding of the product.

5 Correctness of the Diagnosis

We now show our main result, stating that the runs of the diagnosis grammar properly capture all those runs of the system model which explain the observation. This is done by exploiting the coreflection result (Theorem 16) and by additionally taking care of the pruning phase (Definition 24). Note that the proofs will not require an explicit construction of the product of occurrence grammars. We first need some technical results.

Lemma 27. *Let $\mathcal{G}_1, \mathcal{G}_2$ be two occurrence grammars. Consider the product of the two grammars and its image through the Ilv functor as shown below. Furthermore consider the product of the interleaving structures $Ilv(\mathcal{G}_1), Ilv(\mathcal{G}_2)$. Then the mediating morphism δ is a projection which is total on events.*

$$\begin{array}{ccc}
 Ilv(\mathcal{G}_1) & \xleftarrow{\delta_1} & Ilv(\mathcal{G}_1) \times Ilv(\mathcal{G}_2) & \xrightarrow{\delta_2} & Ilv(\mathcal{G}_2) \\
 & \searrow \pi_1 & \uparrow \delta & \nearrow \pi_2 & \\
 & & Ilv(\mathcal{G}_1 \times \mathcal{G}_2) & &
 \end{array}$$

Proof. The fact that the mediating morphism δ is a projection can be shown as in [4, Lemma 20].

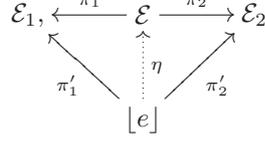
In order to prove that δ is total it suffices to show that, for any event e in $Ilv(\mathcal{G}_1 \times \mathcal{G}_2)$ (and hence in $\mathcal{G}_1 \times \mathcal{G}_2$), either $\pi_1(e)$ or $\pi_2(e)$ are defined. This implies that $\delta(e)$ must be defined, since otherwise either $\delta_1 \circ \delta = \pi_1$ or $\delta_2 \circ \delta = \pi_2$ would not hold.

In order to show this we consider the labelled asymmetric event structures (AES) $\mathcal{E}_1 = Ev(\mathcal{G}_1)$, $\mathcal{E}_2 = Ev(\mathcal{G}_2)$ (where Ev is the functor that associates every occurrence grammar with an AES) and prove an analogous property for AESs.

We know that the functor Ilv factors through Ev , i.e., there is a functor Ilv' mapping event structures to interleaving structures with $Ilv = Ilv' \circ Ev$.

As an immediate consequence of the results in [2], the functor Ev is a coreflection also for the labelled categories, hence it holds that $Ev(\mathcal{G}_1 \times \mathcal{G}_2) = \mathcal{E}_1 \times \mathcal{E}_2 =$:

\mathcal{E} . Now let e be an event of \mathcal{E} such that $\pi_1(e)$ and $\pi_2(e)$ are both undefined. Consider an event structure $[e]$ containing all the causes of e and its morphisms π'_1, π'_2 to \mathcal{E}_1 and \mathcal{E}_2 obtained by restricting π_1 and π_2 (see the diagram below).



There are at least two possibilities for the mediating morphism η : either the event e of $[e]$ is mapped to e in \mathcal{E} or the image is undefined. This is in contradiction to the uniqueness of the mediating morphism. Hence, by contradiction, either $\pi_1(e)$ or $\pi_2(e)$ must be defined.

In order to complete the proof we apply the functor Ilv' to the diagram above. \square

To lighten the notation, hereafter, given an interleaving structure \mathcal{I} we write $\lambda^*(I)$ for $\lambda^*(R_I)$. Recall that, given $f : A_1 \rightarrow A_2$, $f^* : A_1^* \rightarrow A_2^*$ denotes the (non-strict) extension of f to sequences. Then $f^{-1} : \mathcal{P}(A_2^*) \rightarrow \mathcal{P}(A_1^*)$ is its inverse.

Lemma 28. *Let $\mathcal{G}_1, \mathcal{G}_2$ be two occurrence grammars and let $f_i : A_1 \cup A_2 \rightarrow A_i$ ($i \in \{1, 2\}$) be the obvious partial inclusions. Then it holds that*

$$\lambda^*(Ilv(\mathcal{G}_1 \times \mathcal{G}_2)) = f_1^{-1}(\lambda_1^*(Ilv(\mathcal{G}_1))) \cap f_2^{-1}(\lambda_2^*(Ilv(\mathcal{G}_2))).$$

Proof. We set $\mathcal{I}_1 = Ilv(\mathcal{G}_1)$, $\mathcal{I}_2 = Ilv(\mathcal{G}_2)$ and $\mathcal{I} = \mathcal{I}_1 \times \mathcal{I}_2 = Ilv(\mathcal{G}_1) \times Ilv(\mathcal{G}_2)$. Furthermore let $\pi_i : \mathcal{I} \rightarrow \mathcal{I}_i$, $i \in \{1, 2\}$ be the projections. It holds that

$$\begin{aligned}
\lambda^*(\mathcal{I}) &= \lambda^*(R_{\mathcal{I}}) = \{\lambda_{\mathcal{I}}^*(r) \mid r \in (E'_{\mathcal{I}})^{\odot} \wedge \pi_1^*(r) \in R_1 \wedge \pi_2^*(r) \in R_2\} \\
&\stackrel{(*)}{=} \{w \in A^* \mid f_1^*(w) \in \lambda_1^*(R_1), f_2^*(w) \in \lambda_2^*(R_2)\} \\
&= f_1^{-1}(\lambda_1^*(R_1)) \cap f_2^{-1}(\lambda_2^*(R_2)) \\
&= f_1^{-1}(\lambda_1^*(\mathcal{I}_1)) \cap f_2^{-1}(\lambda_2^*(\mathcal{I}_2))
\end{aligned}$$

It remains to show the equality marked (*) which is not obvious.

\subseteq : Let $\lambda_{\mathcal{I}}^*(r)$ be the sequence of labels associated with a run r . Since π_1 is defined exactly on the events having a label contained in A_1 it holds that $\lambda_1^*(\pi_1^*(r)) = f_1^*(\lambda_{\mathcal{I}}^*(r))$, hence $f_1^*(\lambda_{\mathcal{I}}^*(r)) \in \lambda_1^*(R_1)$. Similarly one can show that $f_2^*(\lambda_{\mathcal{I}}^*(r)) \in \lambda_2^*(R_2)$.

\supseteq : Let $w \in A^*$ be a word such that $f_1^*(w) \in \lambda_1^*(R_1)$ and $f_2^*(w) \in \lambda_2^*(R_2)$. We assume that $w = c_1 \dots c_k$, $f_1^*(w) = c_{i_1} \dots c_{i_n}$ and $f_2^*(w) = c_{j_1} \dots c_{j_m}$. Furthermore let $r_1 = f_{i_1} \dots f_{i_n} \in R_1$ be a run with $\lambda_1^*(r_1) = f_1^*(w)$ and $r_2 = g_{j_1} \dots g_{j_m} \in R_2$ be a run with $\lambda_2^*(r_2) = f_2^*(w)$. Now construct a run $r = e_1 \dots e_n$ with

$$e_i = \begin{cases} (f_i, g_i) & \text{if } c_i \in A_1 \cap A_2 \\ (f_i, *) & \text{if } c_i \in A_1 \setminus A_2 \\ (*, g_i) & \text{if } c_i \in A_2 \setminus A_1 \end{cases}$$

It holds that $\pi_1^*(r) = r_1$, $\pi_2^*(r) = r_2$ and obviously $\lambda_1^*(r) = w$. Hence w is contained in the left-hand set.

In the next step we have to show that $\lambda^*(\mathcal{I}) = \lambda^*(\text{Ilv}(\mathcal{G}_1 \times \mathcal{G}_2))$. Due to Lemma 27 there exists a total projection $\delta: \text{Ilv}(\mathcal{G}_1 \times \mathcal{G}_2) \rightarrow \mathcal{I}$, which is therefore surjective on runs. Hence, since δ respects labellings, whenever we map two runs of $\text{Ilv}(\mathcal{G}_1 \times \mathcal{G}_2)$ to one run of \mathcal{I} , they must be associated with the same label sequence. So, if we apply the labelling function to both sides, we obtain the same sets of label sequences. \square

The next proposition shows that considering the product of the original grammar \mathcal{G} and of the observation \mathcal{O} , taking its unfolding and the corresponding labelled runs, we obtain exactly the runs of \mathcal{G} compatible with the observation.

Proposition 29. *Let \mathcal{G} be a grammar and \mathcal{O} an observation, where Λ is the set of labels of \mathcal{G} and $\Lambda' \subseteq \Lambda$ the set of labels of \mathcal{O} . Furthermore let $f: \Lambda \rightarrow \Lambda'$ be the obvious partial inclusion. Then it holds that:*

$$\lambda^*(\text{Ilv}(\mathcal{U}(\mathcal{G} \times \mathcal{O}))) = \lambda^*(\text{Runs}(\mathcal{G})) \cap f^{-1}(\lambda^*(\text{Runs}(\mathcal{O}))).$$

Proof. First, observe that $\lambda^*(\text{Runs}(\mathcal{G})) = \lambda^*(\text{Ilv}(\mathcal{U}(\mathcal{G})))$ and $\lambda^*(\text{Runs}(\mathcal{O})) = \lambda^*(\text{Ilv}(\mathcal{U}(\mathcal{O}))) = \lambda^*(\text{Ilv}(\mathcal{O}))$. Furthermore due to the fact that unfolding is a coreflection, we have $\mathcal{U}(\mathcal{G} \times \mathcal{O}) = \mathcal{U}(\mathcal{G}) \times \mathcal{U}(\mathcal{O}) = (\mathcal{U}(\mathcal{G})) \times (\mathcal{U}(\mathcal{O}))$.

Hence we have to show that

$$\lambda^*(\text{Ilv}((\mathcal{U}(\mathcal{G})) \times (\mathcal{U}(\mathcal{O})))) = \lambda^*(\text{Ilv}(\mathcal{U}(\mathcal{G}))) \cap f^{-1}(\lambda^*(\text{Ilv}(\mathcal{U}(\mathcal{O}))))).$$

This is a corollary of Lemma 28 for $\mathcal{G}_1 = \mathcal{U}(\mathcal{G})$, $\mathcal{G}_2 = \mathcal{U}(\mathcal{O}) = \mathcal{O}$, $\Lambda_1 = \Lambda$, $\Lambda_2 = \Lambda'$. Furthermore f corresponds to f_2 and f_1 is the identity since $\Lambda' \subseteq \Lambda$. \square

We can conclude that the described diagnosis procedure is complete, i.e., given an observation of size n , the runs of the diagnosis grammar $D(\mathcal{G}, \mathcal{O})$ with n observable events are in 1-1 correspondence with those runs of \mathcal{G} that provide a full explanation of the observation. As a preliminary result, on the basis of Proposition 29 one could have shown that the same holds replacing the diagnosis grammar with $\mathcal{U}(\mathcal{G} \times \mathcal{O})$, i.e., the unpruned unfolding. The result below additionally shows that no valid explanation is lost during the pruning phase.

Theorem 30 (correctness of the diagnosis). *With the notation of Proposition 29 it holds that:*

$$\lambda^*(\text{Runs}^n(D(\mathcal{G}, \mathcal{O}))) = \lambda^*(\text{Runs}(\mathcal{G})) \cap f^{-1}(\lambda^*(\text{Runs}^n(\mathcal{O}))).$$

That is, the maximal interleavings of the diagnosis grammar (seen as label sequences) are exactly the runs of the model which explain the full observation.

Proof. By definition

$$\lambda^*(Runs^n(D(\mathcal{G}, \mathcal{O}))) = \lambda^*(Runs^n(Pr(\pi_2))), \quad (1)$$

where $\pi_2: \mathcal{U}(\mathcal{G} \times \mathcal{O}) \rightarrow \mathcal{O}$ is the second projection of the product. By Lemma 25 this is the same as

$$\lambda^*({r \in Runs(\mathcal{U}) \mid \pi^*(r) \in Runs^n(\mathcal{O})}), \quad (2)$$

We will now show that set (2) coincides with

$$\lambda^*(Runs(\mathcal{G})) \cap f^{-1}(\lambda^*(Runs^n(\mathcal{O}))). \quad (3)$$

\subseteq : Let $w \in \{r \in Runs(\mathcal{U}) \mid \pi^*(r) \in Runs^n(\mathcal{O})\}$. There exists a run r of the unfolding with $\pi_2^*(r) \in Runs^n(\mathcal{O})$ and $\lambda^*(r) = w$. Hence it follows from Proposition 29 that $w = \lambda^*(r) \in \lambda^*(Runs(\mathcal{G}))$.

It remains to show that $w \in f^{-1}(\lambda^*(Runs^n(\mathcal{O})))$. It holds $\lambda^*(\pi_2^*(r)) \in \lambda^*(Runs^n(\mathcal{O}))$ is a subsequence of w where all unobservable labels are missing. These labels can be reinserted by f^{-1} and hence w is contained in $f^{-1}(\lambda^*(Runs^n(\mathcal{O})))$. This proves the inclusion (2) \subseteq (3).

\supseteq : Now let $w \in \lambda^*(Runs(\mathcal{G})) \cap f^{-1}(\lambda^*(Runs^n(\mathcal{O})))$. We have to show that $w \in \lambda^*(\pi_2^{-1}(Runs^n(\mathcal{O})))$. Since

$$\begin{aligned} & \lambda^*(Runs(\mathcal{G})) \cap f^{-1}(\lambda^*(Runs^n(\mathcal{O}))) \\ & \subseteq \lambda^*(Runs(\mathcal{G})) \cap f^{-1}(\lambda^*(Runs(\mathcal{O}))) \\ & = \lambda^*(Ilv(\mathcal{U}(\mathcal{G} \times \mathcal{O}))) \quad (\text{Proposition 29}) \end{aligned}$$

there exists a run r of the unfolding of the product with $\lambda^*(r) = w$. Hence $\pi_2^*(r)$ is a run of \mathcal{O} with⁵ $\lambda^*(\pi_2^*(r)) = f^*(\lambda^*(r)) = f^*(w) \in \lambda^*(Runs^n(\mathcal{O}))$ (by choice of w). This implies also that $\pi_2^*(r)$ is contained in $Runs^n(\mathcal{O})$.

Therefore $r \in \{r \in Runs(\mathcal{U}) \mid \pi(r) \in Runs^n(\mathcal{O})\}$ and since $\lambda^*(r) = w$ we obtain $w \in \lambda^*(\pi_2^{-1}(Runs^n(\mathcal{O})))$. □

Observe that, due to the nondeterministic nature of the diagnosis grammar, events which are kept in the pruning phase as they are part of some full explanation of the observation, can also occur in a different configuration. As a consequence, although all inessential events have been removed, the diagnosis grammar can still contain spurious configurations which cannot be extended to full explanations. As an example, consider the graph grammar \mathcal{G} in Fig. 7, given in a Petri-net-like notation. Assume we observe three unordered events **a**, **b**, **c**. Then the unfolding of the product basically corresponds to \mathcal{G} itself. In the pruning phase nothing is removed. However there is a configuration (indicated by the dashed closed line) that cannot be further extended to an explanation.

⁵ Recall that $\lambda^*(\pi_2^*(r)) = f^*(\lambda^*(r))$ holds for every morphism π_2 due to the definition of a morphism.

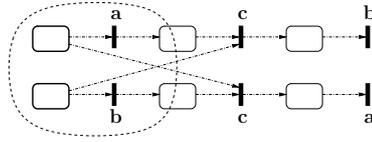


Fig. 7. Spurious runs in a diagnosis grammar.

6 Conclusion

In this paper we formalised event-based diagnosis for systems with variable topologies, modelled as graph transformation systems. In particular we have shown how to exploit the coreflection result for the unfolding of graph grammars in order to show the correctness of a diagnosis procedure generating partially ordered explanations for a given observation.

We are confident that the approach presented in the paper, although developed for transformation systems over hypergraphs, can be generalised to the more abstract setting of adhesive categories. In particular we are currently working on a generalization of the unfolding procedure that works for SPO-rewriting in (suitable variations of) adhesive categories [14]. This would allow one to have a kind of parametric framework which can be used to instantiate the results of this paper to more general rewriting theories.

We are also interested in distributed diagnosis where every observer separately computes possible explanations of local observations that however have to be synchronized. In [3] we already considered distributed unfolding of Petri nets; for *diagnosis* however, the non-trivial interaction of distribution and pruning has to be taken into account. Distribution will require the use of pullbacks of graph morphisms, in addition to products.

References

1. P. Baldan, A. Corradini, and U. Montanari. Contextual Petri nets, asymmetric event structures and processes. *Information and Computation*, 171(1):1–49, 2001.
2. P. Baldan, A. Corradini, U. Montanari, and L. Ribeiro. Unfolding semantics of graph transformation. *Information and Computation*, 205:733–782, 2007.
3. P. Baldan, S. Haar, and B. König. Distributed unfolding of Petri nets. In *Proc. of FOSSACS '06*, pages 126–141. Springer, 2006. LNCS 3921.
4. P. Baldan, S. Haar, and B. König. Distributed unfolding of Petri nets. Technical Report CS-2006-1, Department of Computer Science, University Ca' Foscari of Venice, 2006.
5. Paolo Baldan, Thomas Chatain, Stefan Haar, and Barbara König. Unfolding-based diagnosis of systems with an evolving topology. Technical Report 2008-2, Abteilung für Informatik und Angewandte Kognitionswissenschaft, Universität Duisburg-Essen, 2008.
6. A. Benveniste, E. Fabre, S. Haar, and C. Jard. Diagnosis of asynchronous discrete event systems, a net unfolding approach. *IEEE Transactions on Automatic Control*, 48(5):714–727, 2003.

7. R. Bruni and F. Gadducci. Some algebraic laws for spans (and their connections with multirelations). In W. Kahl, D.L. Parnas, and G. Schmidt, editors, *Proceedings of RelMiS 2001, Relational Methods in Software*, volume 44.3 of *ENTCS*, 2001. Also published as Technical Report, Bericht Nr. 2001-02, Fakultät für Informatik, Universität der Bundeswehr München.
8. R. Bruni and H. C. Melgratti. Non-sequential behaviour of dynamic nets. In *ICATPN*, volume 4024 of *LNCS*, pages 105–124. Springer, 2006.
9. C. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems*. Kluwer Academic Publishers, 1999.
10. T. Chatain and C. Jard. Models for the supervision of web services orchestration with dynamic changes. In *AICT/SAPIR/ELETE*, pages 446–451. IEEE, 2005.
11. A. Corradini, U. Montanari, and F. Rossi. Graph processes. *Fundamenta Informaticae*, 26:241–265, 1996.
12. E. Fabre, A. Benveniste, S. Haar, and C. Jard. Distributed monitoring of concurrent and asynchronous systems. *Journal of Discrete Event Dynamic Systems*, 15(1):33–84, 2005.
13. S. Haar, A. Benveniste, E. Fabre, and C. Jard. Partial order diagnosability of discrete event systems using Petri net unfoldings. In *Proc. 42nd IEEE Conf. on Decision and Control (CDC)*, 2003.
14. S. Lack and P. Sobociński. Adhesive and quasiadhesive categories. *RAIRO – Theoretical Informatics and Applications*, 39(3), 2005.
15. M. Löwe. Algebraic approach to single-pushout graph transformation. *Theoretical Computer Science*, 109:181–224, 1993.
16. M. Löwe, M. Korff, and A. Wagner. An Algebraic Framework for the Transformation of Attributed Graphs. In M.R. Sleep, M.J. Plasmeijer, and M.C. van Eekelen, editors, *Term Graph Rewriting: Theory and Practice*, pages 185–199. Wiley, London, 1993.
17. M. Sampath, R. Sengupta, K. Sinnamohideen, S. Lafortune, and D. Teneketzis. Failure diagnosis using discrete event models. *IEEE Trans. on Systems Technology*, 4(2):105–124, 1996.
18. Y. Wang, T.-S. Yoo, and S. Lafortune. Diagnosis of discrete event systems using decentralized architectures. *Discrete Event Dynamic Systems*, 17(2):233–263, 2007.
19. G. Winskel. Event Structures. In *Petri Nets: Applications and Relationships to Other Models of Concurrency*, volume 255 of *LNCS*, pages 325–392. Springer Verlag, 1987.
20. G. Winskel. Petri nets, algebras, morphisms, and compositionality. *Information and Computation*, 72(3):197–238, 1987.