

A Proposal to Detect Errors in Enterprise Application Integration Solutions*

Rafael Z. Frantz¹, Rafael Corchuelo², and Carlos Molina-Jiménez³

¹ UNIJUÍ University, Department of Technology
Rua do Comércio, 3000, Ijuí, 98700-000, RS, Brazil
rzfrantz@unijui.edu.br

² Universidad de Sevilla, ETSI Informática
Avda. de la Reina Mercedes, s/n, Sevilla 41012, Spain
corchu@us.es

³ Newcastle University, School of Computing Science
Newcastle upon Tyne, NE1 7RU, United Kingdom
carlos.molina@ncl.ac.uk

Abstract. Enterprise Application Integration (EAI) solutions comprise a set of specific-purpose processes that implement exogenous message workflows. The goal is to keep a number of applications' data in synchrony or to develop new functionality on top of them. Such solutions are prone to errors because they are highly distributed and involve applications that were not usually designed with integration concerns in mind. This has motivated many authors to work on provisioning EAI solutions with fault-tolerance capabilities. In this article we analyse EAI solutions from two orthogonal perspectives: viewpoint (orchestration versus choreography) and execution model (process- versus task-based model). A review of the literature shows that current proposals are bound to a specific viewpoint or execution model or have important limitations. In this article, we report on an error monitor that can be used to provision EAI solutions with fault-tolerance capabilities. Our theoretical analysis proves that it is computationally tractable, and our experimental results prove that it is efficient enough to be used in situations in which the workload is very high.

Key words: Enterprise Application Integration, Fault-Tolerance, Error Monitoring, Error Detection.

1 Introduction

The computer infrastructure of a typical today's enterprise can be seen as a software ecosystem that involves several complementary applications purchased

* Complementary material for this article is available at the following Internet site: <http://www.tdg-seville.info/rzfrantz/JSS-2010>. This material includes an implementation of the system, source code, documentation, and screen casts.

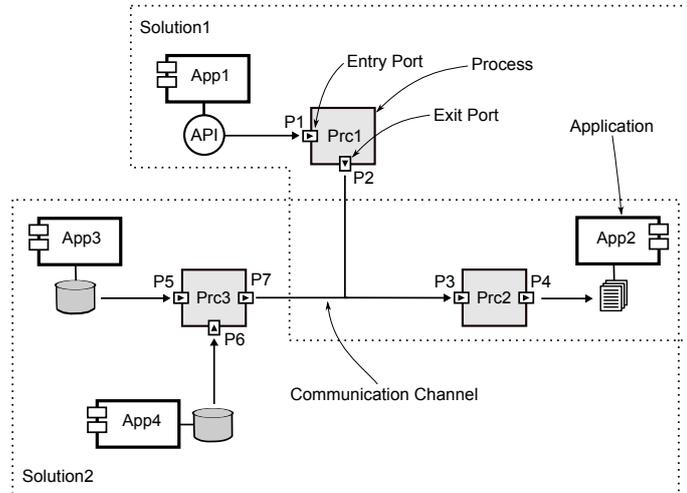


Figure 1. Sample integration solutions.

from different providers or built at home [26]. A recurrent challenge is to make these applications inter-operate with each other to keep their data synchronised or to create a new piece of functionality [14, 15]. This problem is known as Enterprise Application Integration (EAI).

A typical EAI solution consists of one or more processes that interact with each other and with the existing applications by means of ports that read/write messages from/to communication channels. Roughly speaking, they implement an exogenous workflow in which messages are read from a subset of applications, routed through the processes, which may transform them, and the results are written to another subset of applications. Ports abstract away from the details of a specific communication mechanism, which may range from an RPC-based protocol over HTTP to a document-based protocol implemented on a database management system. Processes build on tasks to filter, enrich, split, aggregate, route or transform messages, to name a few [14, 15]. It is common that solutions share processes, which makes them overlap or even include others. Figure §1 shows two solutions that we shall use throughout the article to illustrate our proposal. In this example, there are two overlapping solutions, namely: *Solution1*, which integrates applications App1 and App2 by means of processes Prc1 and Prc2, and *Solution2*, which integrates applications App3, App4 and App2 by means of processes Prc3 and Prc2. Note that process Prc2 is shared by both solutions.

EAI solutions can be characterised from several perspectives. In this article, we focus on viewpoint and execution model.

By viewpoint, we refer to whether a solution is specified an orchestration or a choreography. It is an orchestration if there is a single process that co-ordinates every exchange of messages [29]. This process plays the role of a centralised point

of control that can consequently have an accurate global view of the current state of execution. This is a valuable piece of information that can be used, for instance, to identify the applications and processes that are involved in an error. Contrarily, an EAI solution is a choreography if there is not a centralised point of control, i.e., applications and processes interact in a peer-to-peer fashion; this consequently implies that no process can have an accurate global view of the execution [29]. Currently, WS-BPEL [28] and WS-CDL [31] are the de facto standards to specify orchestrations and choreographies, respectively.

Regarding the execution model, it is worth mentioning that EAI solutions must be deployed to a run time system that must provide an execution engine. Depending on the granularity of execution, we distinguish between the process- and the task-based execution models. In the process-based model, the engine controls process instances as a whole, i.e., there is no means that it can interact with the internal tasks; contrarily, in the task-based model, the engine may control both process instances and their internal tasks. The implication is that the process-based execution model requires a mechanism to gather messages, correlate them, and decide when a new process instance can be started; furthermore, each process instance requires a thread to be allocated exclusively, which may have a negative impact on performance in cases in which a process instance sends a request and has to wait for a long time before it gets the answer (for instance, think of a request that requires the intervention of a person or a configuration that assigns low priorities to request that come from integration solutions [13]). Since the task-based execution model deals with the tasks inside a process instance, it can allocate threads more efficiently; in other words, no thread shall be idle as long as there is a task ready to be executed, independently from the process to which this task belongs.

EAI solutions are inherently distributed; they are thus vulnerable to a variety of errors due to which they may behave abnormally. Errors are due to faults, which can be either permanent, e.g., due a software defect, or transient, e.g., due to a resource that is temporarily unavailable. Errors that are not dealt with properly are perceived as failures by end users [2, 6]. Fault-tolerance proposals aim to help keep systems delivering their functionality in spite of faults. Typically, they can be modelled as a pipeline that goes through the following stages: event reporting, error monitoring, error diagnosing, and error recovering. The event reporting stage deals with reporting whether a port was able to handle a message or not. In the error monitoring stage, events are stored and analysed to find correlations that shall later be checked for validity. When an error is detected, a notification is created and sent to the error diagnosing stage, whose aim is to identify the cause of the error, the messages and the parties involved. The error recovering stage attempts to execute recovery actions to help the system compensate for the existence of faults and the occurrence of errors.

In the literature there are many proposals to deal with fault tolerance [1, 3, 4, 7–10, 13, 21–25, 33–35]. We have analysed them from the viewpoint and execution model perspectives. Our conclusion is that most of them aim at the orchestration viewpoint and the process-based execution model, since they focus

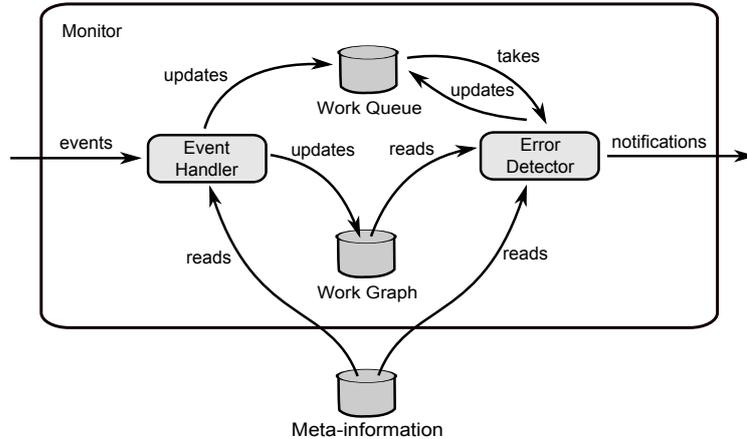


Figure 2. Abstract view of the monitor.

on Web Services, WS-BPEL, and traditional workflow systems. There are a few exceptions that take the choreography viewpoint and/or the task-based execution model into account, but they have important limitations that make them difficult to use in a general context. The main limitation of [7] is that processes can have only one input and messages cannot be split or aggregated inside the workflow; these are common tasks in EAI solutions. In proposals [3, 9], although the authors suggest that it can be applied to the choreography viewpoint, no implementation or evaluation is provided, which makes them difficult to assess and apply in practice. Finally, the limitation of proposal [21] is that it is theoretical and transforming it into a practical tool that deals with distribution problems in a software ecosystem does not seem straightforward.

In this article, we focus on the error monitoring stage. Our proposal is sketched in Figure 2. It relies on a so-called **Meta-Information** database that stores meta-information about the solutions being monitored. (Note that this database is external to the monitor since it is intended to be shared with other stages of the fault-tolerance pipeline.) The monitor itself is composed of two sub-systems and two databases, namely: **Event Handler**, **Error Detector**, **Work Graph** and **Work Queue**. The **Event Handler** uses inbound events to build a graph structure that is stored in the **Work Graph** database; this graph keeps track of the messages processes exchange and their relationships. The **Error Detector** is responsible for analysing this graph to find and verify correlations. The **Work Queue** database is used as an intermediate buffer that allows the **Event Handler** and the **Error Detector** to work in total asynchrony. Every time the **Event Handler** processes an event, it stores a piece of information in the **Work Queue** database; this information instructs the **Error Detector** to analyse the **Work Graph** database at a specific point in time in order to find the correlation in which a specific message is involved. To validate correlations, the **Error Detector** builds on both built-in and user-defined rules; the former allows to detect communications or deadline

errors; the latter allows to detect structural errors that depend on the semantics of a given process or solution, i.e., correlations that lack messages or have more messages than expected. The only assumption we make is that the clock resolution of the monitor is enough to distinguish between every two messages that are read or written in a row; in other words, we can distinguish between multiple events that involve the same message at the same port. Note that this is not a shortcoming in practice since current clock resolutions are in the order of nanoseconds, whereas reading or writing to a port usually consumes much more time.

Our main contribution is that our proposal is not bound with a particular viewpoint and/or execution model; neither imposes it any practical limitations. We have analysed our proposal from a theoretical point of view, and have proved that the algorithms on which it relies are computationally tractable; furthermore, we have carried out a series of experiments that prove that it performs quite well under heavy workloads.

The rest of the article is structured as follows: Section §2, discusses the related work; the **Meta-Information** database shared by all of the stages of the fault-tolerance pipeline is presented in Section §3; Section §4, reports on the **Event Handler**; Section §5, reports on the **Error Detector**; the time complexity analysis of our algorithms is presented in Section §6; the experiments are introduced in Section §7; and, finally, we draw our conclusions in Section §8.

2 Related Work

The literature distinguishes between static and run-time fault tolerance proposals. As discussed in [32], static analysis is concerned with the detection of logical errors in the specification of a system; for instance, race condition errors that emerge from overlooking data dependencies. In this proposal, the authors use a directed acyclic graph to specify the synchronisation constraints imposed on the execution of individual tasks of business processes; they then convert the graph into a coloured Petri net and analyse it to detect potential synchronisation errors. In contrast, we assume that the EAI solutions with which we deal are logically consistent; we then focus on errors that emerge during their execution due, for instance, to delayed or missing messages. The work on static analysis is orthogonal to ours, which is the reason why we do not report on these proposals.

We have realised that the distinction between the error detection and the error recovery stages is blurred in many proposals. This is common in cases in which the complexity of the error detection algorithm is trivial or abstracted away from the user's view by means of another mechanism. In our proposal, we make a clear distinction between error detection and error recovery.

In some cases, the presence of an error can be determined from the analysis of a single event, e.g., the widely-used try-catch mechanism falls within this category [11]. These cases fall outside the scope of our research. We are interested in cases in which it is necessary to process large traces of events that are related to each other. Traces like these are prone to propagate faulty data.

Representative examples include discrete event-based systems, e.g., workflow systems [13], automatic controllers used in manufacturing processes [21], and computer-based controllers used to automatically operate aircrafts, train traffic crossing and medical devices [20]. Error detection is a challenging problem in this context, in particular, when the number of events notified to the monitor is large, e.g., in the order of hundreds of events per second.

To place our proposal in context, it is worth clarifying that our solution builds on a centralised error monitor that can receive events from several EAI solutions concurrently. We are aware that some authors have pointed out that centralised error detectors are problematic regarding scalability [27]. Alternatively, they favour decentralised error detectors which are composed of two or more local error detectors that gather and process their own local information and collaborate to detect global errors. A similar idea to handle scalability is also suggested in [17–19]. Our reservations about decentralised error detectors is the complexity involved in the co-ordination and synchronisation of their activities, and the difficulty to deal with open systems in which processes and solutions are not pre-defined, but may be created and destroyed as time goes by; we leave this alternative out of our discussion on the basis that centralised error detectors seem to satisfy our requirements well even in situations in which the workload is high, cf. Section §7.

A distinctive feature of our research is that we do not only suggest an error detector, but we also provide a rule-based language to express the expected behaviours of processes and solutions. The rules provide flexibility to our error detection mechanism. Finally, it is worth clarifying that the proposals that provide a rule-based language consider that the messages or events in the log are already correlated, presumably because they aim only at orchestrated EAI solutions. Since our proposal aims to be independent from the viewpoint and the execution model, our monitor accounts for the arrival of uncorrelated messages and then finds their correlations. A salient feature of our proposal is that it emphasises the relevance of evaluating the error detection algorithm, from several dimensions. This is an issue frequently overlooked in current literature.

There are several proposals in the literature on fault-tolerance aiming to help keep systems delivering their functionality in spite of errors. Our conclusion is that most of them aim at the orchestration viewpoint and the process-based execution model, since they focus on Web Services, WS-BPEL, and traditional workflow systems. The few proposals that cover choreography and the task-based execution model have important limitations, which prevent them from being used in a general context. For example, in [7] processes can have only one input and messages cannot be split or aggregated inside the workflow; in [3, 9], no implementation or evaluation is provided, which makes it difficult to assess and apply; in practice [21] addresses only a single class of errors. Our proposal tackles the problem of endowing EAI solutions, independently from the viewpoint and execution model, with a fault-tolerance mechanism, that does not impose in practice any limitation. Our failure semantics includes communication, structural and deadline errors.

Proposal	Orchestration Choreography Process-based Task-based			
Zeng and others [35]	+	-	+	-
Liu and others [24]	+	-	+	-
Erradi and others [10]	+	-	+	-
Liu and others [23]	+	-	+	-
Chiu and others [8]	+	-	+	-
Hagen and Alonso [13]	+	-	+	-
Liu and others [25]	+	-	+	-
Alonso and others [1]	+	-	+	-
Chen and others [7]	+	+†	+	+†
Ermagan and others [9]	+	+†	+	-
Baresi and others [3]	+	+†	+	-
Li and others [22]	+	-	+	-
Borrego and others [4]	+	-	+	-
Li and others [21]	-	+†	-	+†
Yan and Dague [33]	+	-	+	-
Yan and others [34]	+	-	+	-

† With important limitations.

Table 1. Summary of the related work.

In the following, we report on more details about the proposals we have summarised in Table §1.

Fault tolerance in business processes and web service composition has been addressed by several authors, but mainly with a focus on error recovery. For instance, in [35] the authors realise that service compositions are vulnerable to failures and produce different exceptions; they then show how so-called ECA rules can be used to handle them at run time. The focus of this proposal is on error recovery. Error detection is mentioned in passing only, presumably, because their context does not require sophisticated error detection algorithms. A similar discussion on exception handling for WS-BPEL can be found in [24], which is complementary to [35]. In [10], the authors propose a policy-driven middleware solution to handle exceptions in web service compositions. With this article, we share the view that communication operations are the most error-prone. Relevant to us is also the failure semantics, which covers the communication failure types addressed in our proposal. An algorithm for the execution of WS-BPEL processes with relaxed transactions is presented in [23]. In [24] the authors enhance their approach with a rule-based language for defining specific-purpose error recovery

rules. It is worth emphasising that papers [10, 23, 24, 35] focus on orchestrated applications.

The research conducted by the workflow community on fault-tolerance is also related to our proposal. An abstract model for workflows with fault-tolerance features is discussed in [8]; this article also provides a good survey on fault-tolerance approaches. An algorithm for handling faults automatically in applications integrated by means of workflow management systems, which is amenable to compensation actions or to the two-phase commit protocol, is suggested in [13]. Our reservation about this idea is that it does not isolate fault-tolerance mechanisms from business logic. In addition, it is suitable only for orchestrated applications that are based on the process-based execution model. Error recovery in workflow applications in which compensation actions are difficult or infeasible to implement is discussed in [25]. The authors assume that there is a centralised workflow engine, which suggests that they focus on orchestrated solutions. In [1], the authors discuss error recovery in orchestrated workflow applications that build on the process-based execution model. An architecture to implement fault-tolerance based on ad-hoc workflows is discussed in [7]; runtime error detection is central in this proposal. The architecture relies on a web server as a front-end, an application server, a database server, and a logging system. Like in our proposal, error detection is achieved by means of the analysis of message traces. The reservation we have about this approach is that processes can have only one input and messages cannot be split or aggregated inside the workflow; however, at least in theory, it is suitable for both orchestrated and choreographed processes; it supports both the process- and the task-based execution model.

An architecture for fault-tolerant workflows, based on finite state machines that recognise valid sequence of messages, is discussed in [9]. Error recovery actions are triggered when a message is found to be invalid, or the execution time of the state machine goes beyond a pre-defined deadline. This proposal is suitable for both orchestrated and choreographed processes; however it aims at process-based executions. In [3], the authors discuss some preliminary ideas for building an error monitor that can be used for both orchestrated and choreographed processes. No implementation or evaluation is provided, which makes it difficult to assess and apply in practice.

An approach for runtime detection of errors that emerge from potential corrupted data and faulty web services in WS-BPEL orchestrated processes is discussed in [22]. The central idea is to regard WS-BPEL processes as discrete event systems and map them onto coloured Petri nets. The Petri net model is then integrated into additional WS-BPEL processes that are triggered when exceptions are thrown. Although the authors believe that this approach can be extended to handle choreographed processes, they do not discuss how this can be done. In [4], the authors study runtime error detection in business processes specified in BPMN and discuss a framework that includes a diagnosis layer that runs in parallel and independently from the main BPMN process. The authors consider that a BPMN process is composed of several activities that might deliver incorrect output, for instance, due to incorrect inputs provided by humans. The

correct functionality of each activity is specified by means of compliance rules that are later mapped onto constraint satisfaction problems; errors are detected by means of a constraint solver. The diagnosis layer is conceptually similar to the monitor in our proposal. This proposal seems to focus on functional errors (also called business errors); in contrast, we focus on non-functional errors related to the computer infrastructure. A limitation of this proposal is that, in its current state, it can handle only a single instance of a BPMN process.

Error detection has also been intensively studied by designers of controllers that automate the operation of manufacturing plants, which can also be viewed as discrete event systems. These controllers need to process events that arrive from different sensors. Consequently, to determine whether the plant is operating correctly, the controller needs to perform intensive event correlation. In [21], the authors explain how one can detect errors in manufacturing plants that can be specified and controlled by means of Petri nets. They focus on place faults (for instance, tokens that are not removed after firing a transition) caused by sensor failures and bit flips. To provide the controller with error detection capabilities they embed it into a larger controller that provides additional redundant places, tokens and transitions that are used to specify correctness invariants. Errors are detected by linear parity checks at run time. The similarity between this proposal and ours lies in the use of event correlation for error detection. This proposal aims at choreographed applications. Our criticism is that it addresses only a single class of errors, namely, place errors. Neither is it clear how these ideas can be ported to business processes that are not modelled as Petri nets and involve a large amount of events and process instances executing simultaneously.

A core feature of our proposal is that we assume that all of the events produced by an EAI solution are notified (in other words, observable to the error detector). Several authors have studied error detection in discrete event systems that are considered to be in failure when they do not produce one or more expected events, cf. [30]. The challenge here is to analyse the observable events to infer what unobservable event or events drove the system into an abnormal state. This kind of systems is outside the scope of our current work. In [33, 34], the authors suggest to re-use the body of knowledge about error detection in industrial discrete event systems. They discuss runtime error detection of orchestrated web services. A salient feature of this proposal is that, similarly to [30], the authors assume that failure events are not observable. The granularity of execution in this approach is at the process level. Our criticism against this proposal is that it focuses only on orchestrated web services and the process-based execution model.

3 The Meta-Information database

In this section, we report on the meta-data we use to represent the information our proposal requires about the artefacts it monitors. This meta-data are stored in the *Meta-Information* database. We do not provide many details on how this database is managed since this is a typical information system with

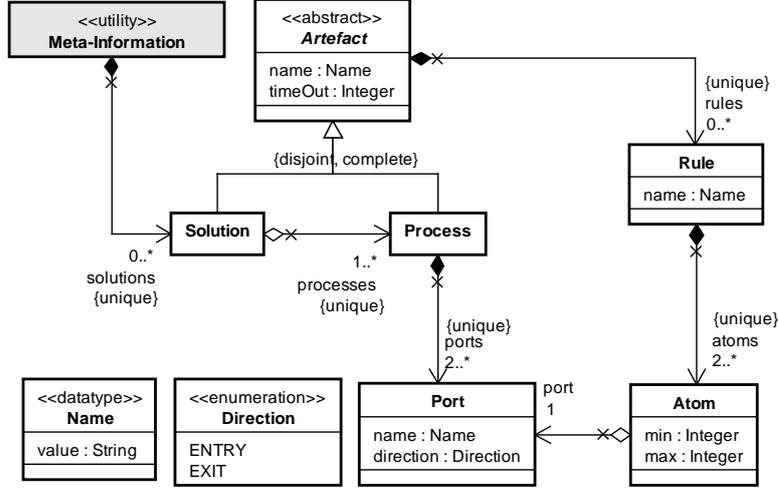


Figure 3. Model of the Meta-Information database.

$$\text{Rule} = \underbrace{P_1 [n_1..m_1] \& P_2 [n_2..m_2] \dots P_k [n_k..m_k]}_{\text{Atoms regarding entry ports}} \longrightarrow \underbrace{P_{k+1} [n_{i+1}..m_{j+1}] \& P_{k+2} [n_{i+2}..m_{j+2}] \dots P_{k+i} [n_{i+q}..m_{i+r}]}_{\text{Atoms regarding exit ports}}$$

Figure 4. Textual syntax for rules.

an interface in which administrators can register, unregister, or list information about the solutions to be monitored. Instead, we focus on the meta-data it stores, whose model is presented in Figure §3, and provide a few hints on our implementation.

Note that we use term artefact to refer to both solutions and processes. A solution must be composed of at least one process, and every process must have at least two ports with different directions (`Direction::ENTRY` or `Direction::EXIT`). Every artefact, port, or rule has a name that identifies it uniquely. In addition, artefacts have a time out, which denotes the maximum time they may require to process a set of correlated messages, and a set of rules, which help verify the correlations in which they are involved.

Rules are very important in our proposal. Figure §4 presents the syntax we use to write them textually. They are composed of two groups of atoms that are separated by means of an arrow. The left hand side refers to entry ports and the right hand side to exit ports. Each atom is of the following form: $P[\text{min}..\text{max}]$, where P refers to a port name, and min and max are natural numbers that represent the minimum and the maximum number of messages that are allowed at port P in a given correlation ($\text{min} \leq \text{max}$). For the sake of brevity, we use the common syntactic sugar depicted in Figure §5.

- a) $P[+] = P[1 \dots \text{Integer.MAX_VALUE}]$
- b) $P[*] = P[0 \dots \text{Integer.MAX_VALUE}]$
- c) $P[?] = P[0..1]$
- d) $P[n] = P[n..n]$

Figure 5. Syntactic sugar for rules.

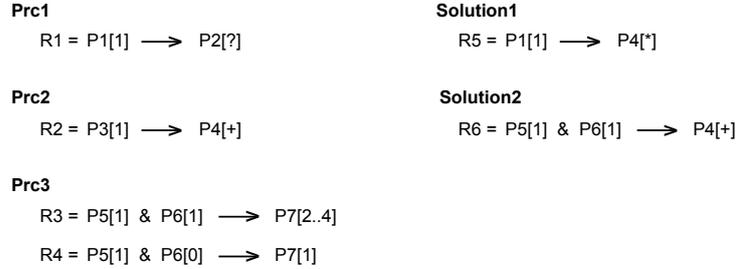


Figure 6. Sample rules.

Figure §6 presents a few rules for the artefacts in the sample solutions we introduced in Figure §1. For instance, Rule R1 is associated with process Prc1 and involves entry port P1 and exit port P2; it states that it validates a given correlation as long as there is one message at port P1 and zero or one correlated message at port P2. Similarly, Rule R6 is associated with Solution2; it states that it validates a given correlation as long as there is one message at port P5, one correlated message at port P6, and one or more correlated messages at port P4.

Regarding our implementation, the **Meta-Information** database is a simple set of data, and none of our queries involve joining information from other databases. In our prototype we use a hash function to index the entries of the **Meta-Information** database, and we have implemented maps from port names onto processes, processes onto solutions, and so on. The space required by this design is proportional to the number of artefacts, ports, and rules. As a conclusion, it is possible to retrieve information from the **Meta-Information** database in $O(1)$ time. The **Work Queue** database relies on a Brodal’s priority queue [5], which allows to insert entries or retrieve the next to be analysed in $O(1)$ time, whereas removing an entry takes $O(\log n)$ time, where n denotes the size of the queue.

4 The Event Handler

Figure §7 depicts the model we have devised for the **Event Handler**. Roughly speaking, it handles events that inform about a port reading or writing a message, either successfully or unsuccessfully. The events are used to build a graph that is maintained incrementally in the **Work Graph** database. This graph records information about the messages being exchanged in an EAI solution and their parent-child relationships, i.e., which messages originate from which ones. In

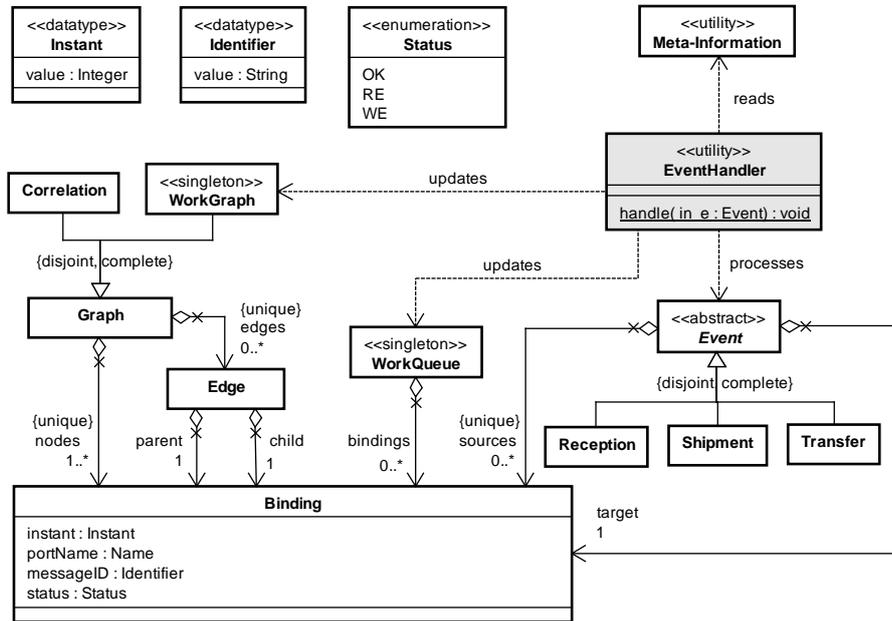


Figure 7. Model of the Event Handler.

addition, it updates the Work Queue database in order to schedule the activation of Error Detector appropriately.

An event can be of type Reception, which happens at ports that read data from an application (either successfully or unsuccessfully) and other ports that fail to read data at all, Shipment, which occurs when a port writes information (either successfully or unsuccessfully), and Transfer, which happens when a port succeeds to read data that was written previously by another port. Every event has a target binding and zero, one, or more source bindings. We use this term to refer to the data involved in an event, namely: the instant when the event happened, the name of the port, the identifier of the message read or written, and a status, which can be either Status::OK to mean that no problem was detected, Status::RE to mean that there was a reading error, or Status::WE to mean that there was a writing error. Recall that the only assumption we make is that the clock resolution of the monitor is enough to distinguish between every two messages that are read or written in a row. Thus, from now on, we assume that no confusion regarding the same message being read from or written to the same port may happen.

The Event Handler provides only one method that handles all types of events. The algorithm for this method is presented in Figure §8. It gets an event *e* as input and proceeds as follows: it first finds the process to which the event refers and the solutions to which this process belongs. Then, it computes the minimum

```

1: to handle(in  $e$ : Event) do
2:    $p$  = find the process to which a port called  $e.target.portName$  belongs
3:     in the Meta-Information database
4:    $s$  = find all solutions to which  $p$  belongs
5:     in the Meta-Information database
6:    $notBefore$  =  $e.target.instant$  + maximum time out of artefacts in  $s \cup \{p\}$ 
7:    $g$  = WorkArea.getInstance()
8:    $q$  = WorkQueue.getInstance()
9:   add  $e.target$  to  $g.nodes$ 
10:  add  $e.target$  with priority  $notBefore$  to  $q$ 
11:  for each binding  $b$  in  $e.sources$  do
12:     $r$  = new Edge(parent =  $b$ , child =  $e.target$ )
13:    add  $r$  to  $g.edges$ 
14:  end for
15: end

```

Figure 8. Algorithm to handle events.

time at which the Error Detector should analyse the target binding, which is the instant when the message in the target binding was read or written plus the maximum time out involved; this is a safe deadline that guarantees that every artefact should have enough time to process the corresponding correlation. Note that the time we calculate is just a hint that must be interpreted as “the Error Detector should not analyse that binding before this time”; obviously, the sooner the binding is analysed after this time has passed, the better, but it is not a real-time requirement. The algorithm then fetches both the Work Graph and the Work Queue instances and adds the target binding to them both; then, it iterates over the source bindings, if any, and adds them to the Work Graph together with an edge to link them to the target binding.

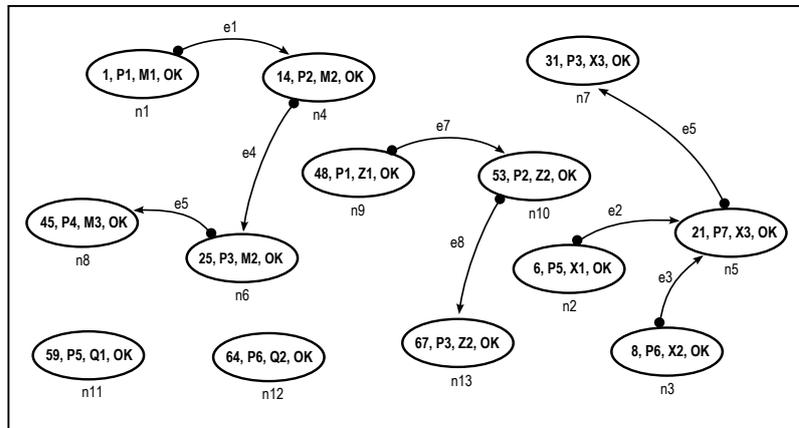
Figure §9 illustrates a graph that results from executing the previous algorithm on a series of bindings regarding the sample system in Figure §1. Ellipses denote bindings and arrows denote edges that connect parent bindings to their corresponding child bindings. For instance, $n1$ is the parent binding of $n4$, and the latter is the parent of $n6$, which is in turn the parent of $n8$. Inside each binding we represent the instant, the port name, the message id, and the status, respectively. A snapshot of the Work Queue is also presented in this figure.

5 The Error Detector

The model of the Error Detector is presented in Figure §10, and the algorithm to detect errors is presented in Figure §11.

The Error Detector executes a never-ending loop in which it fetches entries from the Work Queue, finds the correlations in which the corresponding bind-

Work Graph



Work Queue

n1	n2	n3	n4	n5	n6	...
101	107	109	115	122	126	

Figure 9. Sample work graph.

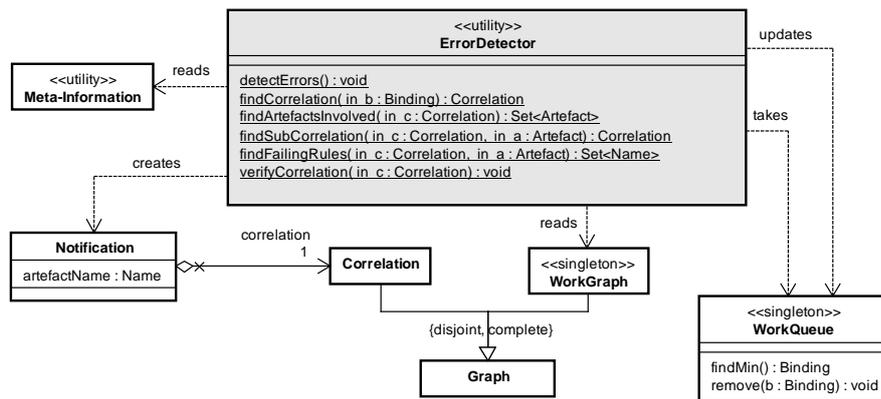


Figure 10. Model of Error Detector.

ings are involved, and then verifies them. Recall that each binding in the Work Queue is scheduled to be processed not before a given time; this implies that this algorithm may need to block for some time in cases in which the earliest binding is scheduled to be analysed after the current moment. After a correlation is verified, its bindings are removed from the Work Queue, since analysing them would not result in new correlations.

```

1: to detectErrors() do
2:    $q = \text{WorkQueue.getInstance}()$ 
3:   repeat
4:      $b = \text{fetch minimum of } q \text{ (waiting if necessary)}$ 
5:      $c = \text{findCorrelation}(b)$ 
6:      $\text{verifyCorrelation}(c)$ 
7:     for each binding  $b$  in  $c.\text{nodes}$  do
8:       remove  $b$  from  $q$ 
9:     end for
10:  end repeat
11: end

```

Figure 11. Algorithm to detect errors.

In the following subsections, we report on the sub-algorithms on which the Error Detector relies. In Section §5.1, we present our algorithm to find the correlation in which a binding is involved; we then report on a number of ancillary sub-algorithms, namely: an algorithm to find the artefacts that are involved in a given correlation, cf. Section §5.2, an algorithm to find the sub-correlation that corresponds to a given artefact, cf. Section §5.3; and an algorithm to find the subset of sub-rules according to which a correlation is invalid, cf. Section §5.4; the previous algorithms are used to implement the algorithm to verify a correlation, which we present in Section §5.5.

5.1 Finding Correlations

A correlation is represented as a graph that has a single connected component that represents a subset of messages that are correlated to each other, cf. Figure §7 and [16].

The Error Detector provides a method called `findCorrelation` whose algorithm is presented in Figure §12. It gets a binding as input and calculates the correlation in which it is involved. The main loop navigates from the binding that is passed as a parameter to all of the bindings that are reachable from it, either directly or transitively, and to all of the bindings from which it can be reached, either directly or transitively. In other words, it implements a breadth-first search to calculate the expansion of a node in a graph [12].

For instance, if algorithm `findCorrelation` is invoked on bindings `n1`, `n4`, `n6`, or `n8` in Figure §9, it then would return the correlation in Figure §13.

5.2 Finding the Artefacts Involved in a Correlation

A correlation may involve several artefacts. This situation is very common since typical EAI solutions involve several processes, some of which may be

```

1: to findCorrelation(in b: Binding): Correlation do
2:   result = new Correlation(nodes =  $\emptyset$ , edges =  $\emptyset$ )
3:   g = WorkArea.getInstance()
4:   q =  $\emptyset$ 
5:   add b to q
6:   whilst q <>  $\emptyset$  do
7:     d = take a binding from q
8:     add d to result.nodes
9:     cs = (children of d in g) \ result.nodes
10:    add cs to q
11:    for each c in cs do
12:      h = new Edge(parent = d, child = c)
13:      add h to result.edges
14:    end for
15:    ps = (parents of d in g) \ result.nodes
16:    add ps to q
17:    for each p in ps do
18:      h = new Edge(parent = p, child = d)
19:      add h to result.edges
20:    end for
21:  end whilst
22:  return result
23: end

```

Figure 12. Algorithm to find correlation.

shared. This implies that an event reported from a port results in a binding that actually involves several artefacts.

The Error Detector provides a method called `findArtefactsInvolved` whose algorithm is presented in Figure §14. It gets a correlation as input and returns a set of artefacts. The main loop of this method iterates over the set of bindings in the correlation that is passed as a parameter. In each iteration, it firsts finds the processes that own the ports from which the corresponding events were reported, and the solutions in which they are involved. The loop simply adds all of these artefacts to the result, and then returns the whole collection.

For instance, if we invoke method `findArtefactsInvolved` on the correlation in Figure §13, the following artefacts involved would be returned: `Solution1`, `Solution2`, `Prc1`, and `Prc2`, cf. Figure§16.

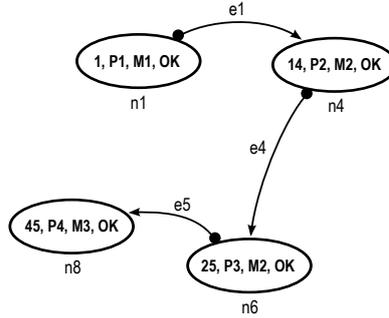


Figure 13. Sample correlation.

```

1: to findArtefactsInvolved(in c: Correlation): Set(Artefact) do
2:   result =  $\emptyset$ 
3:   for each binding b in c.nodes do
4:     p = find the process to which a port called b.portName belongs in
5:       the Meta-Information database
6:     s = find the solutions to which a process p belongs in
7:       the Meta-Information database
8:     add  $s \cup \{p\}$  to result
9:   end for
10:  return result
11: end
  
```

Figure 14. Algorithm to find artefacts involved in a correlation.

5.3 Finding Sub-Correlations

By sub-correlation, we refer to a subset of a correlation in which the ports involved belong to a given artefact. Note that there is not any structural differences between correlations and sub-correlations: they both are represented as graphs. In the sequel, we write (sub)correlation wherever we wish to emphasise that there is a single artefact from which all of the events represented in a correlation were reported.

The Error Detector provides a method called `findSubCorrelation` whose algorithm is presented in Figure §15. It gets a correlation and an artefact as input and returns a (sub)correlation. The algorithm is straightforward: it first finds which ports belong to the artefact, and then finds all of the bindings in the correlation whose ports are in the previous set; to create the resulting sub-correlation we just need to find the whole collection of edges that connect the previous bindings.

For instance, Figure §16 shows all of the sub-correlations we can find in the correlation in Figure §13.

- 1: to findSubCorrelation(in c : Correlation, in a : Artefact): Correlation do
- 2: $p =$ find all ports of artefact a in the Meta-Information database
- 3: $b =$ find all bindings b in $c.nodes$ such that binding $b.portName$
- 4: belongs to $p.portName$ in the Meta-Information database
- 5: $e =$ find all edges x such that $\{x.source, x.target\} \subseteq b$
- 6: return new Correlation(nodes = b , edges = e)
- 7: end

Figure 15. Algorithm to find sub-correlations.

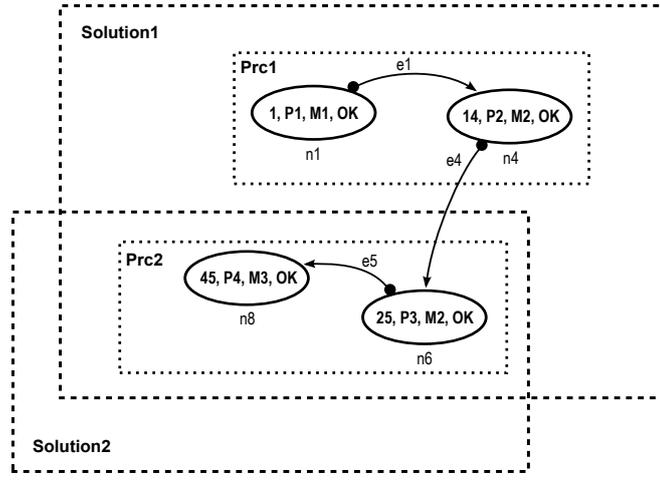


Figure 16. Artefacts involved in a correlation.

5.4 Finding Failing Rules

A part of the verification of a (sub)correlation relies on a set of user-defined rules. The Error Detector provides a method called `findFailingRules` that takes a (sub)correlation and an artefact as input and returns a set of names that denote the rules associated with the artefact according to which the correlation can be considered invalid.

The algorithm to the `findFailingRules` method is presented in Figure §17. It iterates through the collection of rules that are associated with a given artefact, and then through their atoms. The algorithm basically counts the number of bindings in the given (sub)correlation that correspond to events that were reported from the ports to which each atom refers; if this count is not within the limits that the atom specifies, then the corresponding rule does not validate the correlation and can thus be added to the result of the algorithm.

Figures §18 and §19 illustrate the two situations in which a correlation can be considered invalid due to a failing rule. The left side of the figures represent the (sub)correlation being analysed, whereas the right side represent the times

```

1: to findFailingRules(in c: Correlation, in a: Artefact): Set(Name) do
2:   result =  $\emptyset$ 
3:   for each rule r in a.rules do
4:     for each atom t in r.atoms do
5:       n = count bindings b in c.nodes such that b.portName == t.port.name
6:       if n <= t.min or n >= t.max then
7:         add r.name to result
8:       end if
9:     end for
10:  end for
11:  return result
12: end

```

Figure 17. Algorithm to find failing rules.

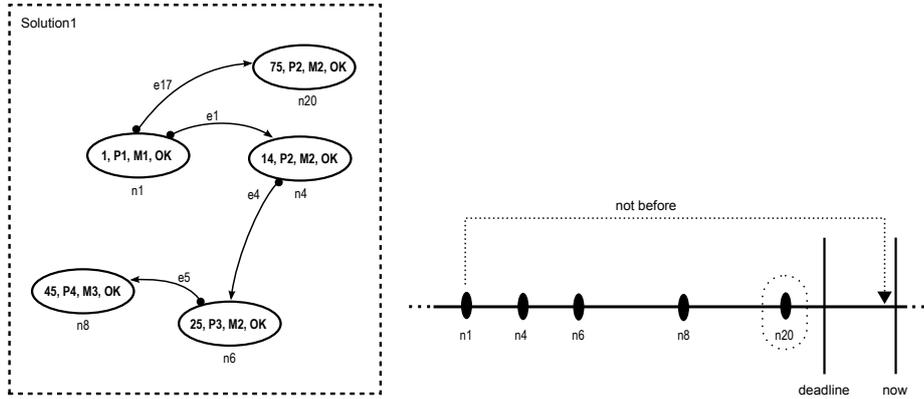


Figure 18. Correlation which causes a rule to fail due to excess of bindings.

at which events were reported; *now* represents the instant at which the analysis is performed, and *deadline* the maximum time at which a correlation is expected to be produced (see more on this below); *not before* represents the time not before which the initial binding in a correlation must be analysed. Consider, for instance, rule R1 for artefact *Solution1*, which was introduced in Figure §6; it states that it is expected zero or one correlated bindings at port P2 for each binding at port P1. Note that in the (sub)correlation in Figure §18 there are two correlated bindings at port P2, namely n4 and n20, which causes rule R1 to fail due to excess of bindings. Contrarily, in Figure §19 the (sub)correlation contains less bindings than expected.

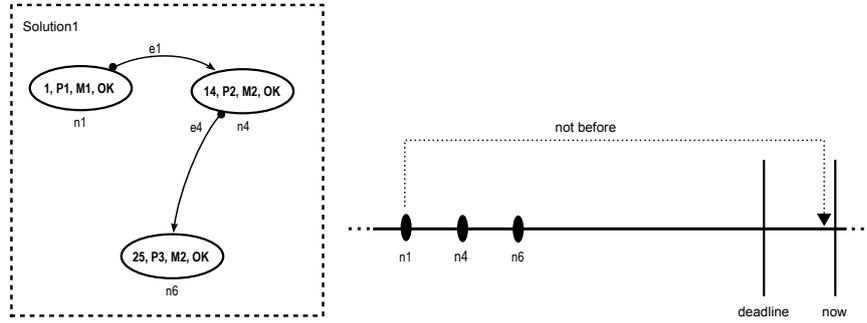


Figure 19. Sub-correlation which causes a rule to fail due to lack of bindings.

```

1: to verifyCorrelation(in c: Correlation) do
2:   artefacts = findArtefactsInvolved(c)
3:   for each artefact a in artefacts do
4:     s = findSubCorrelation(c,a)
5:     status = for all binding b in s, b.status == Status :: OK
6:     earliestInstant = minimum of s.nodes.instant
7:     latestInstant = maximum of s.nodes.instant
8:     deadline = earliestInstant + a.timeOut
9:     isValid = (latestInstant <= deadline) and
10:      ( status == true) and
11:      findFailingRules(s,a) == ∅
12:   if not isValid then
13:     n = new Notification(artefactName = a.name, correlation = c)
14:     send n to the next fault-tolerance stage
15:   end if
16: end for
17: end

```

Figure 20. Algorithm to verify correlations.

5.5 Verifying Correlations

Verifying (sub)correlations is the central task of the Error Detector since this is the task that actually detects errors. A (sub)correlation can be either valid or invalid. It is valid if all of its bindings have status `Status::OK`, no rule fails to validate it, and all of the messages to which it refers were read or written within a given deadline; otherwise, it is invalid. The deadline refers to the time when the first message in a correlation was read or written plus the time out of the corresponding artefact. Recall that each artefact is associated with a time

Notation Meaning

s	Maximum number of solutions to which a process can belong.
u	Maximum number of source bindings in an event.
b	Maximum number of bindings in a correlation.
c	Maximum number of child bindings of a given binding.
p	Maximum number of parent bindings of a given binding.
r	Maximum number of rules associated with an artefact.
t	Maximum number of atoms in a rule.
a	Maximum number of artefacts involved in a correlation.
n	Number of entries in the Work Queue.

Table 2. Notation used in our complexity analysis.

out that represents the maximum time it is expected to produce a correlation, cf. Figure §3.

The Error Detector provides a method called `verifyCorrelation` to perform this task. The algorithm to this method is presented in Figure §20. It takes a correlation as input, but does not return any results. The algorithm first calculates the artefacts involved in the correlation and iterates through them to find their corresponding (sub)correlations; each (sub)correlation is checked for validity according to the definition in the previous paragraph. (Sub)correlations that are found invalid are transformed into notifications that are sent to the following fault-tolerance stage so that they can be diagnosed.

6 Complexity Analysis

In this section we analyse our proposal and characterise its complexity. It can deal with an arbitrary number of processes and solutions, but we assume that there is a sensible upper bound; this does not amount to loss of generality since the number of artefacts of which a company’s software ecosystem is composed must be necessarily finite. Table §2 summarises the notation we use in this section.

The analysis proves that our proposal is computationally tractable since handling events runs in $O(1)$ time and detecting errors runs in $O(\log n)$ time for a given ecosystem. This makes the proposal appealing from a theoretical point of view since it is logarithmic on the size of the Work Queue database, which is expected not to be monotonically increasing or decreasing, but to grow and shrink as time goes by. Our experiments support this conjecture, cf. Section §7.

6.1 Handling Events

We first report on the complexity of the algorithm to handle events, and prove that it is computationally tractable because it runs in constant time for a given ecosystem.

Theorem 1. *Algorithm handle in Figure §8 terminates in $O(s + u)$ time.*

Proof. Handling an event involves finding a process by means of a port name and the solutions to which this process belongs. Finding this information can be accomplished in $O(1)$ time in our implementation, cf. Section §3 (lines §2 and §4). The computation of the maximum time out can be easily accomplished in $O(s)$ time (line §6). Getting the Work Graph and the Work Queue instances can be accomplished in $O(1)$ time (lines §7 and §8). Lines §9 and §10 add the target binding to the Work Graph and to the Work Queue databases, respectively, which can also be accomplished in $O(1)$ time. The loop in lines §11–§14 iterates u times at most. Lines §12 and §13 can be implemented in $O(1)$ time since they just create an object and add it to a set. As a conclusion, Algorithm handle terminates in $O(s + u)$ time.

Corollary 1. *In a given ecosystem, there must be an upper bound to $s + u$ because the number of solutions in a company's ecosystem must be finite, which in turn implies that there must also be an upper bound to the number of source bindings in an event. As a conclusion algorithm handle terminates in $O(1)$ time in a given ecosystem.*

6.2 Detecting Errors

We now analyse the complexity of the algorithm to detect errors. Note that this algorithm does not terminate, since it is a never ending loop. In this case, the complexity refers to the complexity of an iteration of this loop.

Theorem 2. *Every iteration of Algorithm detectErrors in Figure §11 terminates in $O(b(1 + c + p + a + \log n) + a r t)$ time.*

Proof. In each iteration, the algorithm first fetches an entry from the Work Queue database, which is accomplished in $O(1)$ time in our implementation, cf. Section §3 (Line §4). According to Theorems §3 and §4 below, lines §5 and §6 run in $O(b(c + p))$ and $O(b + a(b + r t))$ time, respectively. The loop in lines §7–§9 iterates b times at most; in each iteration, Line §8 removes a binding from the Work Queue database, which is accomplished in $O(\log n)$ time, cf. Section §3. As a conclusion, each iteration of Algorithm detectErrors terminates in $O(1 + b(c + p) + b + a(b + r t) + b \log n) = O(b(1 + c + p + a + \log n) + a r t)$ time.

Corollary 2. *In a given ecosystem, b , c , p , a , r , and t are constants because there is an upper limit to the number of artefacts a company runs. As a conclusion, each iteration of Algorithm detectErrors terminates in $O(\log n)$ time in a given ecosystem.*

Next, we analyse the complexity of Algorithm `findCorrelation`.

Theorem 3. *Algorithm `findCorrelation` in Figure §12 terminates in $O(b(c+p))$ time.*

Proof. The loop in lines §6–§21 iterates b times at most. In each iteration, the algorithm executes two inner loops. The first one, iterates c times at most and the operations inside terminate in $O(1)$ time since they just involve creating objects and adding them to a set. Similarly, the second one iterates p times at most and the operations inside also terminate in $O(1)$ time. As a conclusion, Algorithm `findCorrelation` terminates in $O(b(c+p))$ time.

Now, we analyse the complexity of Algorithm `verifyCorrelation`.

Theorem 4. *Algorithm `verifyCorrelation` in Figure §20 terminates in $O(b + a(b + r t))$ time.*

Proof. The algorithm first calculates the artefacts involved in a given correlation, which terminates in $O(b)$ time according to Theorem §5 below (Line §2). It then executes a loop that iterates a maximum of a times (Lines §3–§16). In each iteration, it first needs to find a number of sub-correlations, which terminates in $O(b)$ time according to Theorem §6 (Line §4); then, it calculates the status, the earliest instant, and the latest instant, which requires iterating a maximum of b times (Lines §5–§7); calculating the deadline can be accomplished in $O(1)$ time (Line §8), but determining if a sub-correlation is valid involves executing algorithm `findFailingRules`, which runs in $O(r t)$ time according to Theorem §7 below. Lines §12–§15 run in $O(1)$ time since they just require creating an object and sending it to another stage. As a conclusion, Algorithm `verifyCorrelation` terminates in $O(b + a(b + r t))$ time.

In the following theorems, we analyse the complexity of the three sub-algorithms on which `verifyCorrelation` relies.

Theorem 5. *Algorithm `findArtefactsInvolved` in Figure §14 terminates in $O(b)$ time.*

Proof. The loop at lines §3–§9 iterates b times at most. Lines §4–§8 can be implemented in $O(1)$ time, cf. Section §3. As conclusion, finding the artefacts involved in a given correlation terminates in $O(b)$ time.

Theorem 6. *Algorithm `findSubCorrelation` in Figure §15 terminates in $O(b)$ time.*

Proof. The algorithm first finds the ports that belong to a given artefact, which can be accomplished in $O(1)$ time, cf. Section §3 (Line §2). It then finds the bindings in a correlation whose port belongs to the previous set, which requires iterating through the bindings in the correlation (Line §2); this requires $O(b)$ time in the worst case. Finally, the algorithm finds the edges that connect the previous bindings, which also requires $O(b)$ time (Line §5), and creates an object in $O(1)$ time (Line §6). As a conclusion, Algorithm `findSubCorrelation` terminates in $O(1 + 2 b) = O(b)$ time.

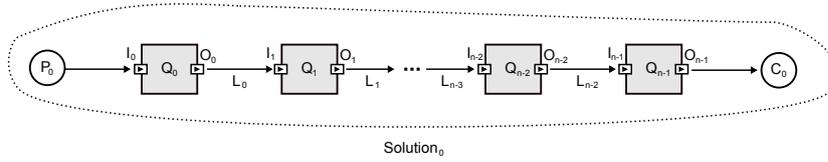


Figure 21. Pipeline pattern.

Theorem 7. *Algorithm findFailingRules in Figure §17 terminates in $O(r t)$ time.*

Proof. The loop at lines §3–§10 iterates a maximum of r times, and the inner loop at lines §4–§9 iterates t times at most. As conclusion, Algorithm findFailingRules runs in $O(r t)$ time.

7 Experiments

We have conducted a series of experiments to evaluate our proposal in the laboratory. We implemented them on top of a discrete event simulation layer that allowed us to run the experiments in simulated time. In the following sections, we first provide additional details on the patterns with which we have experimented, and then on the experimentation parameters; later, we draw our conclusions about the experimental results.

7.1 Experimentation Patterns

We set up four well-known patterns that lay at the core of many real-world EAI solutions, namely: pipeline, dispatcher, merger, and request-reply [14, 15]. In the sequel, we use term producer to refer to the process or application that produces the messages that are fed into a pattern; similarly, we use term consumer to refer to the process or application that consumes the messages the pattern produces. Furthermore, we use variable n to denote the total number of processes involved in each pattern.

In the pipeline pattern, messages flow from a producer to a consumer in sequence: the messages a process produces are consumed by the next process in the pipeline, cf. Figure §21. There is a single producer and a single consumer, i.e., changes to n have an impact on the number of intermediate processes only. In other words, $2n$ events are reported to the monitor per message the producer feeds into the pattern.

In the dispatcher pattern, there is a process that routes the messages it produces to only a specific consumer, cf. Figure §22. Note that changes to n have an effect on the number of consumers, not on the number of producers, which is one. That is, 4 events are reported to the monitor per message the producer feeds into the pattern.

The goal of the merger pattern is to gather messages from several producers and route them to a unique consumer, cf. Figure §23. Changes to n have an

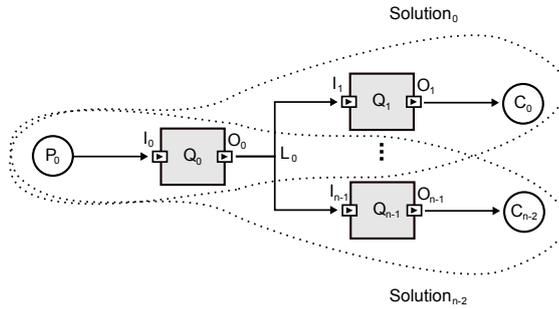


Figure 22. Dispatcher pattern.

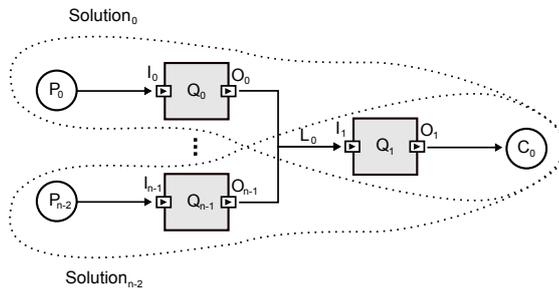


Figure 23. Merger pattern.

impact on the number of producers, not on the number of consumers, which is one. Note that 4 events are reported to the monitor per message a producer feeds into the pattern.

In the request-reply pattern, there are a number of processes that require a service from another process, cf. Figure §24. Changes to n have an effect on the number of client processes that send requests to the single server process. Every message fed into the pattern results in 6 events that are reported to the monitor.

7.2 Experimentation Parameters and Variables

Each experiment consisted of running an instance of a pattern with a fixed number of processes (n) and a fixed mean message production rate (t); we varied n in the range 3..15 processes, and t in the range 100..1000 milliseconds, with increments of 100 milliseconds. In total, we ran 520 experiments to draw our conclusions

We sampled the production rate from a negative exponential with parameter t . Similarly, we sampled both the time to transmit messages and the time each process took to produce a correlation from a negative exponential with parameter 250 milliseconds; the time out of every artefact was set to 5 minutes. We carried out additional experiments with other values for these parameters and found

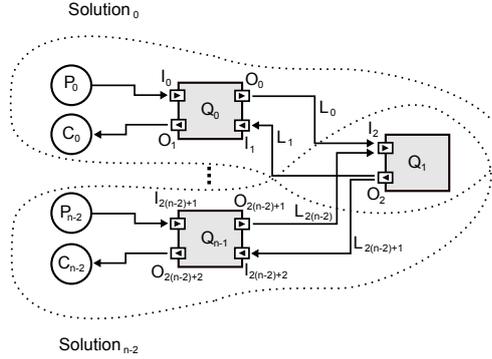


Figure 24. Request-Reply pattern.

that they did not have an impact on the conclusions. Each experiment was run for a duration of 24 hours.

In each experiment, we measured the following variables: the time to handle an event (*THE*), the size of the Work Queue (*QS*), the time each binding spent in the Work Queue after it was scheduled to be analysed (*TSQ*), and the time the `detectError` algorithm took to perform each iteration of its main loop (*TDE*). (In the sequel, we report on the averaged values of these variables.)

We ran the experiments on a cluster of virtual machines that were equipped with four-core Intel Xeon processors running at 3.00 GHz, 16 GB of RAM, Windows Server 2008 64-bit, and the 1.6.0 version of the Java runtime Environment.

7.3 Experimentation Results

Figures §25, §26, §27, and §28 present the results we have gathered regarding variables *THE*, *QS*, *TSQ*, and *TDE* for each pattern.

The time to handle events (*THE*) remains low in all of the experiments, and the changes to n or t do not seem to have an impact on it. According to Theorem §1, the time to handle an event depends on the maximum number of solutions to which a process can belong and on the maximum number of source bindings in an event, which are fixed constants for a given ecosystem. In Corollary §1, we argued that there is an upper bound to these figures, and we concluded that the time to handle events might be considered $O(1)$ in practice. The experiments corroborate this idea, since *THE* seems to be totally independent from n or t in practice.

The size of the Work Queue (*QS*) is important insofar as it has an impact on the memory footprint and the time required to complete an iteration of Algorithm `detectErrors`. It depends on the number of events reported in each experiment. In most cases, it behaves linearly in the number of processes and the slope depends on the mean message production rate (it decreases as this parameter increases). The only exception is the dispatcher pattern, in which *QS* seems to be a constant that depends on the mean message production rate only. The reason for this

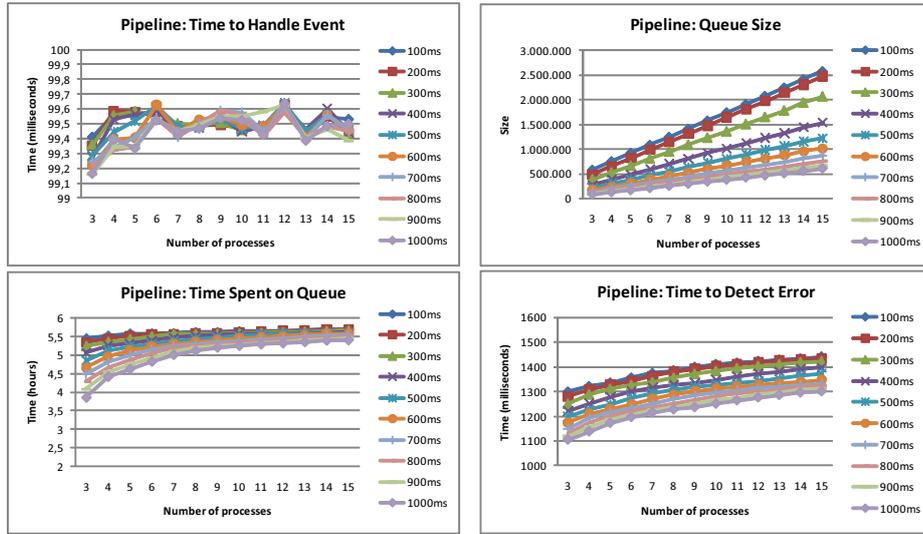


Figure 25. Results of the experiments for the pipeline pattern.

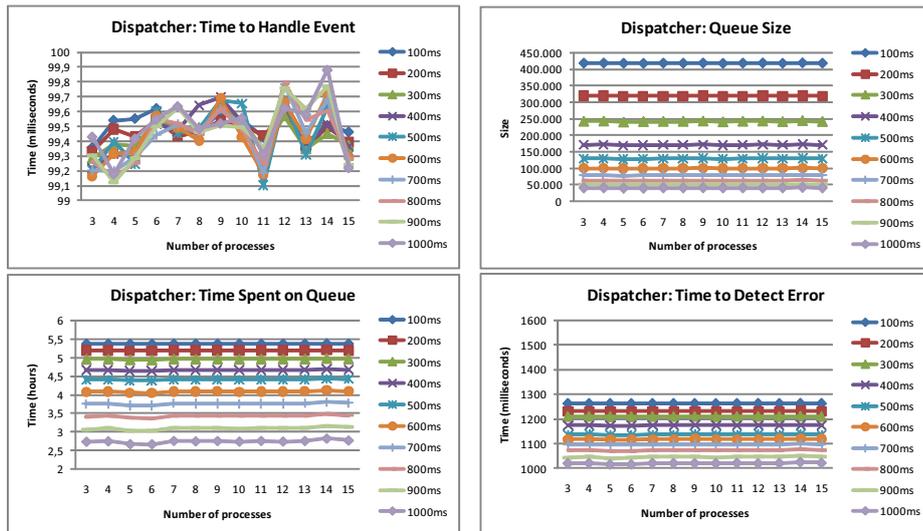


Figure 26. Results of the experiments for the dispatcher pattern.

behaviour is that changes to n do not have an impact on the number of events that are reported. Note that every new process in the pipeline pattern contributes with 2 additional events, new processes in the merger pattern contribute with 4 additional events, and new processes in the request-reply pattern contribute with 6 additional events. Contrarily, adding a new process to the dispatcher pattern

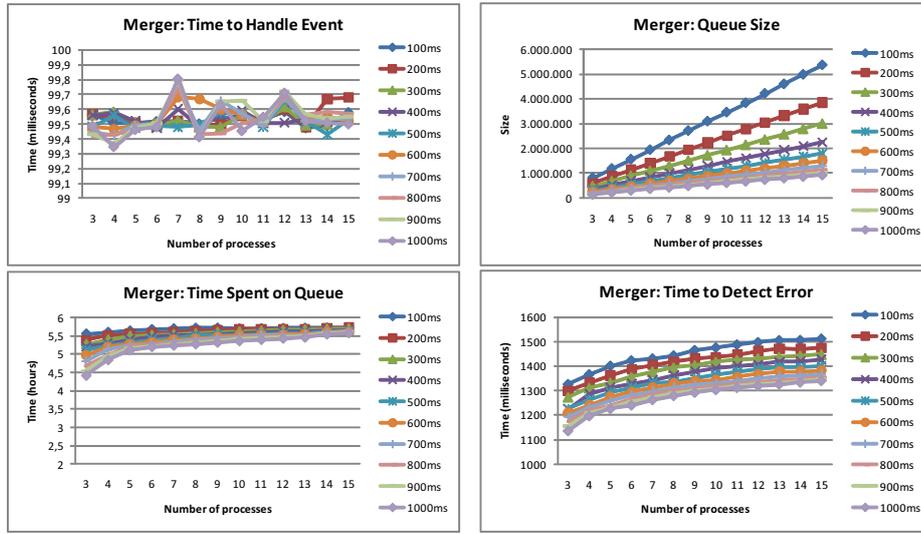


Figure 27. Results of the experiments for the merger pattern.

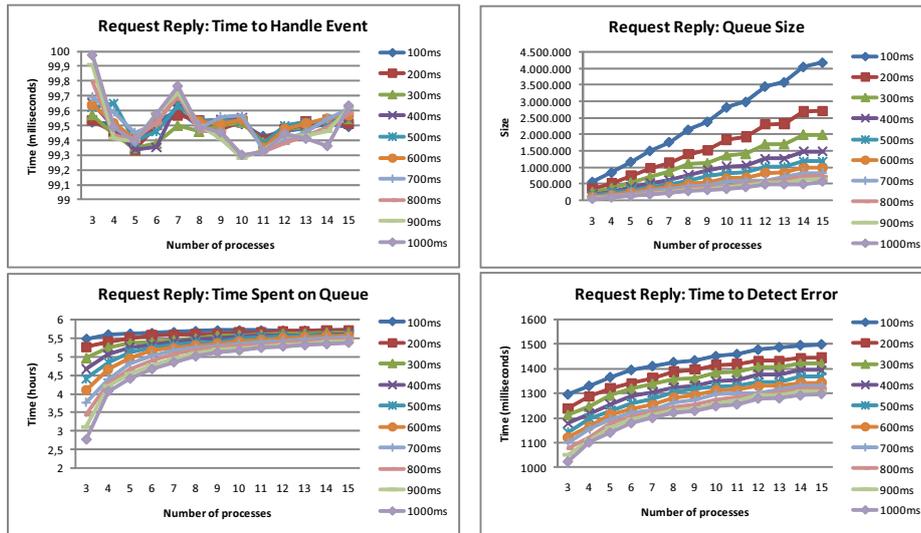


Figure 28. Results of the experiments for the request-reply pattern.

does not contribute with new events. The reason is that adding a new consumer implies that there is a new process that competes for the messages the producer feeds into the pattern; in other words, the events are reported from different sources, but the total number of events remains constant.

Variable TSQ is very important for our conclusions. Recall that this variable measures the time a binding spends in the **Work Queue** after it is scheduled to be analysed by the **Error Detector**. Spending too much time in this queue is not problematic insofar the goal is to detect errors so that they can be diagnosed and recovered. There are no real-time constraints involved. However, generally speaking, the less time, the better since this implies that errors shall be detected, diagnosed, and recovered sooner. Our experiments prove that TSQ seems to behave logarithmically in the number of processes in all cases, except for the dispatcher pattern, in which it behaves constantly. This implies that a change to the number of processes does not usually have a significant impact on the time bindings spend in the **Work Queue**. Neither does the mean message production rate seem to have a negative impact on this variable.

Regarding the time to detect errors (TDE), we proved that it behaves logarithmically in the size of the **Work Queue**, cf. Theorem §2. This theoretical result was promising as long as the size of the **Work Queue** was not monotonically increasing, as we conjectured in Section §6. Our results support this conjecture since variable TDE ranges in the order of seconds in all of our experiments.

8 Conclusions

In this article we have addressed the problem of detecting errors in EAI solutions. We have reported on a monitor that receives events with information about the messages that are read or written at each port. These events are used to build a graph that keeps record of the messages exchanged within the EAI solutions being monitored. We have reported on a series of algorithms to analyse this graph and find invalid correlations. Our failure semantic includes communication, structural and deadline errors.

In our analysis of the related work we have found out that the majority of the proposals in the literature focus on EAI solutions that are specified as orchestrations and build on the process-based execution model. There are a few that take choreographies or take the task-based execution model into account, but they have important limitations that hinder their applicability in practice.

We have analysed our proposal from a theoretical point of view and our conclusion is that it is computationally tractable. We have also carried out a series of experiments that proof they can be used in settings in which the workload is very high (15 processes and 100 milliseconds as the mean message production rate).

Acknowledgements.

The first author conducted part of this work at the Newcastle University, UK as visiting member of staff. His work is partially funded by the Evangelischer Entwicklungsdienst e.V. (EED). The first and the second authors are partially funded by the European Commission (FEDER), the Spanish and the Andalusian R&D&I programmes (grants TIN2007-64119, P07-TIC-2602, P08-TIC-

4100, TIN2008-04718-E, TIN2010-21744, TIN2010-09809-E, TIN2010-10811-E, and TIN2010-09988-E). The third author is partially funded by UK EPSRC Platform Grant Number EP/D037743/1.

Bibliography

- [1] G. Alonso, C. Hagen, D. Divyakant, A. E. Abbadi, and C. Mohan. Enhancing the fault tolerance of workflow management systems. *IEEE Concurrency*, 8(3):74–81, 2000
- [2] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. on Depend. and Secure Comp.*, 1(1):11–33, Jan-Mar 2004
- [3] L. Baresi, S. Guinea, R. Kazhamiakin, and M. Pistore. An Integrated Approach for the Run-Time Monitoring of BPEL Orchestrations. In *Towards a Service-Based Internet*, volume 5377 of *LNCS*, pages 1–12. Springer, 2008
- [4] D. Borrego, R. M. Gasca, M. T. Gómez-López, and L. Parody. Contract-based diagnosis for business process instances using business compliance rules. In *Proc. 21th Int'l Workshop on Principles of Diagnosis (DX'10)*, 2010
- [5] G. Brodal. Worst-case efficient priority queues. In *7th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 52–58, 1996
- [6] R. Campbell and B. Randell. Error recovery in asynchronous systems. *IEEE Trans. Soft. Eng.*, 12(8):811–826, 1986
- [7] M. Chen, A. Accardi, E. Kiciman, J. Lloyd, D. Patterson, A. Fox, and E. Brewer. Path-based failure and evolution management. In *Proc. Int'l Symp. Netw. Syst. Des. and Impl.*, page 23, 2004
- [8] D. Chiu, Q. Li, and K. Karlapalem. A meta modeling approach to workflow management systems supporting exception handling. *Inf. Syst.*, 24(2):159–184, 1999
- [9] V. Ermagan, I. Kruger, and M. Menarini. A fault tolerance approach for enterprise applications. In *Proc. IEEE Int'l Conf. Serv. Comput.*, volume 2, pages 63–72, 2008
- [10] A. Erradi, P. Maheshwari, and V. Tasic. Recovery policies for enhancing web services reliability. In *Proc. Int'l Conf. on Web Serv. (ICWS'06)*, pages 189–196, Sept 2006
- [11] J. Goodenough. Exception handling: Issues and proposed notation. *Communications of the ACM*, 18(12):683–696, December 1975
- [12] J. Gross and J. Yellen. *Handbook of Graph Theory*. CRC Press, 2003
- [13] C. Hagen and G. Alonso. Exception handling in workflow management systems. *IEEE Trans. Softw. Eng.*, 26(10):943–958, 2000
- [14] G. Hohpe. Enterprise Application Integration. In *Proc. 9th Conf. on Pattern Language of Programms*, page #14, 2002
- [15] G. Hohpe and B. Woolf. *Enterprise Integration Patterns - Designing, Building, and Deploying Messaging Solutions*. Addison-Wesley, 2003
- [16] J. E. Hopcroft and R. E. Tarjan. Efficient algorithms for graph manipulation. *Communications of the ACM*, 16(6):372–378, 1973
- [17] G. Khanna, I. Laguna, F. A. Arshad, and S. Bagchi. Stateful detection in high throughput distributed systems. In *Proc. 26th IEEE Int'l Symp.*

- on *Reliable Distrib. Syst. (SRDS 2009)*, pages 275–287. IEEE Computer Society, 2007
- [18] G. Khanna, P. Varadharajan, and S. Bagchi. Automated online monitoring of distributed applications through external monitors. *IEEE Trans. on Depend. and Secure Comp.*, 3(2):115–129, Apr–Jun 2009
- [19] I. Laguna, F. A. Arshad, D. M. Grothe, and S. Bagchi. How to keep your head above water while detecting errors. In *Proc. 10th Int’l Middleware Conf. 2009*, volume 5896, pages 205–225. Springer, 2009
- [20] N. Levenson. Software safety in embedded computer systems. *Communications of the ACM*, 34(2):34–46, February 1991
- [21] L. Li, C. Hadjicostis, and R. Sreenivas. Designs of bisimilar petri net controllers with fault tolerance capabilities. *IEEE Trans. Syst. Man Cybern. Part A: Syst. Humans*, 38(1):207–217, 2008
- [22] Y. Li, T. Melliti, and P. Dague. A colored petri nets model for diagnosing data faults of BPEL services. In *Proc. 20th Int’l Workshop on Principles of Diagnosis (DX’09)*, pages 267–274, 2009
- [23] A. Liu, L. Huang, Q. Li, and M. Xiao. Fault-tolerant orchestration of transactional web services. In *Proc. Int’l Conf. Web Inf. Syst. Eng.*, pages 90–101, 2006
- [24] A. Liu, Q. Li, L. Huang, and M. Xiao. A declarative approach to enhancing the reliability of BPEL processes. In *Proc. IEEE Int’l Conf. Web Serv.*, pages 272–279, 2007
- [25] C. Liu, M. Orlowska, X. Lin, and X. Zhou. Improving backward recovery in workflow systems. In *Proc. 7th Int’l Conf. Database Syst. Adv. Appl.*, page 276, 2001
- [26] D. Messerschmitt and C. Szyperski. *Software Ecosystem: Understanding an Indispensable Technology and Industry*. MIT Press, 2003
- [27] M. S. Mouchaweh. Decentralized fault detection and isolation of manufacturing systems. In *Proc. 21th Int’l Workshop on Principles of Diagnosis (DX’10)*, 2010
- [28] OASIS. Web Services Business Process Execution Language Version 2.0 Specification, 2007. Available at <http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-OS.html>
- [29] C. Peltz. Web services orchestration and choreography. *Computer*, 36(10):46–52, 2003
- [30] M. Sampath, R. Sengupta, and S. Lafortune. Failure diagnosis using discrete-event models. *IEEE Trans. on Control Syst. Technol.*, 4(2):105–124, march 1996
- [31] W3C. Web Services Choreography Description Language Version 1.0 Specification, 2005. Available at <http://www.w3.org/TR/ws-cdl-10/>
- [32] Q. Wu, C. Pu, and A. Sahai. DAG synchronization constraint language for business process. In *IEEE Int’l Conf. on Commerce and Enterprise Computing (CEC’09)*. IEEE Computer Society, 2006
- [33] Y. Yan and P. Dague. Modeling and diagnosing orchestrated web service processes. In *Proc. IEEE Int’l Conf. on Web Serv. (ICWS 2007)*, pages 51–59, July 9–13, Salt Lake City, Utah, 2007. IEEE Computer Society

- [34] Y. Yan, Y. Pencole, M.-O. Cordier, and A. Grastien. Monitoring web service networks in a model-based approach. In *Proc. Third European Conf. on Web Serv. (ECOWS'05)*, Nov 14-16, Vajxo Sweden, 2005. IEEE Computer Society
- [35] L. Zeng, H. Lei, J.-J. Jeng, J.-Y. Chung, and B. Benatallah. Policy-driven exception-management for composite web services. In *Proc. IEEE Int'l Conf. E-Commer. Technol.*, pages 355–363, 2005