

Monitoring Errors in Integration Workflows

Rafael Z. Frantz

UNIJUÍ University

Department of Technology

Rua do Comércio, 3000, Ijuí, 98700-000, RS, Brazil

Email: rzfrantz@unijui.edu.br

Rafael Corchuelo

University of Seville

ETSI Informática

Avda. Reina Mercedes, s/n. Seville 41012. Spain

Email: corchu@us.es

Carlos R. Rivero

University of Seville

ETSI Informática

Avda. Reina Mercedes, s/n. Sevilla 41012. Spain

Email: carlosrivero@us.es

Carlos Molina-Jiménez

Newcastle University

School of Computing Science

Newcastle upon Tyne, NE1 7RU, United Kingdom

Email: carlos.molina@ncl.ac.uk

Abstract—Enterprise Application Integration (EAI) is a field of Software Engineering. Its focus is on helping software engineers integrate existing applications at a sensible costs, so that they can support new business processes or optimise existing ones. EAI solutions are distributed in nature, which makes them inherently prone to failures. In this paper, we report on a proposal to address error detection in EAI solutions. The main contribution is that it runs in linear time, it deals with both choreographies and orchestrations, and that it is independent from the execution model used.

Keywords: Distributed systems; Enterprise Application Integration; Fault-tolerance; Error detection algorithms.

I. INTRODUCTION

Companies are relying heavily on computer-based applications to run their businesses processes. Such processes must evolve and adapt as companies evolve and adapt to varying contextual conditions. Common problems include that the applications were not designed to facilitate integrating them with others, i.e., they do not provide a business level API, and that they were implemented using a variety of technologies that do not inter-operate easily. The goal of Enterprise Application Integration (EAI) is to help reduce the costs of EAI solutions to facilitate the implementation and evolution of business processes.

Figure §1 sketches two sample EAI solutions that involve four applications and three integration processes. Note that a solution is only a logical means to organise a set of processes: different solutions can share the same processes, and a solution can contain another solution. The processes interact with the applications using the facilities they provide, e.g., an API in the best case, a user interface, a file, a database or other kinds of resources. They help implement message-based workflows to keep a number of applications' data in synchrony or to build new functionality on top of them. Processes use ports to communicate with each other or with applications over communication channels. Ports encapsulate reading from or writing to a resource, which helps abstract away from the details of the communication mechanism, which may range

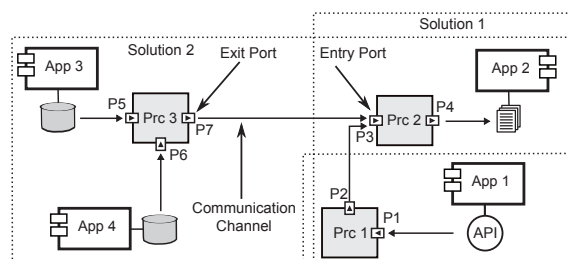


Figure 1. Sample EAI solutions.

from an RPC-based protocol over HTTP to a document-based protocol implemented on a database.

The Service Oriented Architecture initiative has gained importance within the field of EAI, since it provides appropriate technologies to wrap applications (so that they provide business APIs) and to implement message workflows. Centralised workflows, aka orchestrations, rely on a single process that helps co-ordinate a workflow of messages amongst a number of other processes and applications; contrarily, decentralised workflows, aka choreographies, do not rely on such a central co-ordinator. The tools used to implement workflows include conventional systems [1, 8], others based on BPEL, and others like BizTalk [5] or Camel [10].

EAI solutions are distributed in nature, since they involve several applications and processes that may easily fail to communicate with each other [8], which argues for real-world EAI solutions to be fault-tolerant. There seems to be a general consensus that the provisioning fault-tolerance includes the following stages: event reporting, error monitoring, error diagnosing, and error recovery. Event reporting happens when processes report that they have read or written a message by means of a port; the goal of error monitoring is to analyse traces of events to find invalid correlations, i.e., anomalous sets of messages that have been processed together; such correlations must later be diagnosed to find the cause of the anomalies, and appropriate actions to recover from the error

must be taken in the error recovery stage.

Orchestration workflows rely on an external mechanism that analyses inbound messages, correlates them, and starts a new instance of the orchestration whenever a correlation is found. The typical execution model is referred to as process-based since a thread must be allocated to run a process on a given correlation; contrarily, the task-based execution model relies on a pool of threads that are allocated to the tasks. Simply put, in the process-based model threads remain allocated to a process even if that process is waiting for the answer to a request to another process; contrarily, in the task-based model, no thread shall be idle as long as a task in a process is ready for execution.

In this paper, we report on a proposal to build an error monitor for EAI solutions. The key contribution is that it works with both orchestrations and choreographies, and that it is independent from the execution model used. In Section §II, we report on other proposals in the literature; in Section §III, we present an overview of our proposal; in Section §IV, we delve into our proposal to detect errors; in Section §V, we analyse our proposal both from a theoretical and a practical point of view; finally, we present our conclusions and future work in Section §VI; appendix §A provides a few ancillary proofs that support our theoretical analysis.

II. RELATED WORK

Error detection is relatively easy in orchestration systems because either correlations are found prior to starting an orchestration process and everything happens within the boundaries of this process. Contrarily, in choreographies, a correlation may typically involve several processes that run in total asynchrony, and there is not a single point of control; furthermore, EAI solutions may overlap since it is common that processes are reused across several business processes. This makes it more difficult to endow choreographies with fault-tolerance capabilities.

The research on fault tolerance that has been conducted by the workflow community is closely related to our work. Chiu and others [4] presented an abstract model for workflows with embedded fault-tolerance capabilities; it set the foundations for other proposals in this field. Hagen and Alonso [8] presented a proposal that builds on the two-phase commit protocol, and it is suitable for orchestrations in which the execution model is process-based. Alonso and others [1] provided additional details on the minimum requirements to deal with fault tolerance in orchestrated systems. Liu and others [12] discussed how to deal with fault tolerance in settings in which recovery actions are difficult or infeasible to implement; the authors also assume the existence of a centralised workflow engine, i.e., they also focus on orchestrations. Li and others [11] reported on a theoretical solution that is based on using Petri nets; they see processes as if they were controllers, and report on detecting some classes of errors by means of linear parity checks; the key is that they focus on systems in which a fault can involve an arbitrarily large number of correlated messages, which are consumed and produced by distributed

processes, but are assume that they are choreographed by a central processor. An architecture for fault-tolerant workflows, based on finite state machines that recognise valid sequences of messages was discussed in [6]; this proposal is suitable for both orchestrated and choreographed processes; however it is aimed at process-based executions.

The study of fault tolerance in the context of choreographies has been paid less attention in the literature. Chen and others [3] presented a proposal that deviates from the previous ones in that their results can be applied to both orchestrations and choreographies. They assume that the system under consideration is organised into three logical layers (front-end, application server, and database server), plus an orthogonal layer (the logging system). Since they can deal with choreographies, they need to analyse message traces to detect errors. They assume that each message has a unique identifier that allows to trace it throughout the execution flow; unfortunately, they cannot deal with EAI solutions in which messages are split or aggregated, since this would require to find correlations amongst messages, which is not supported at all. Due to this limitation, it can easily deal with both process- and task-based execution models. Yan and Dague [15], Yan and others [14] suggested to re-use the body of knowledge about error detection in industrial discrete event systems, in error detection in web services applications; they discussed runtime error detection of orchestrated web services; a salient feature of this proposal is that, similarly to [13], the authors assume that failure events are not observable; the granularity of execution in this approach is at process level. Baresi and others [2] discussed some preliminary ideas for building an error monitor that can be used for both orchestrated and choreographed processes. No implementation or evaluation was provided.

Our analysis of the literature reveals most authors in the EAI field focus on orchestrations and the process-based execution model; choreographies and the task-based execution model have been paid little attention so far. Another conclusion is that the distinction amongst the stages required to provision fault tolerance is often blurred. The reason is that many proposals focus on error recovery since error detection or error diagnosing is quite a trivial task. In many proposals, the presence of an error can be derived from a single event. For instance, the conventional try-catch mechanism involves the notification of a single event to be caught by the exception mechanisms [7]. However, there is a large class of applications in which the presence of an error can only be deduced from the analysis of traces of events that are related to each other, e.g., by order, parent-child relationships, propagation, or causations. Error detection in these cases is a challenging problem, in particular, when the number of events is large.

III. OVERVIEW OF OUR PROPOSAL

Our proposal builds on a monitor to which each port must report events, and a set of rules that help determine if correlations are valid or not, cf. Figure §2. A monitor is composed of three modules called Registrar, Event Handler, and




Figure 2. Structure of the monitor.

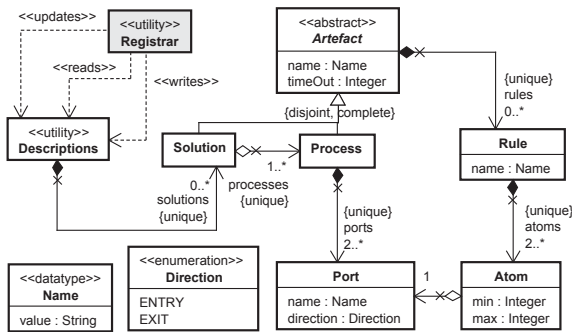


Figure 3. Model of the Registrar module.

Error Detector, three databases called Descriptions Database, Graphs Database, and History Database, and a queue called Graphs Queue.

The Registrar module is responsible for maintaining the Descriptions Database up to date. This database provides the other modules a description of the solutions, processes, ports, and rules the monitor handles. Figure §3 presents the abstract model of this module. (Note that we use term ‘artefact’ to refer to both solutions and processes.)

The Event Handler uses the events reported by ports to update the Graphs Database and the Graphs Queue. Figure §4 presents an abstract model of this module. 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 OK to mean that no problem was detected, RF to mean that there was a reading failure, or WF to mean that there was a writing failure.

The Graphs Database stores an entry per artefact in the Descriptions Database; such entries contain a graph that the Event Handler builds incrementally, as it receives events. Figure §5 shows a sample Graphs Database for the EAI solution in Figure §1. For instance, let us focus on bindings b_6 and b_4 : the former is

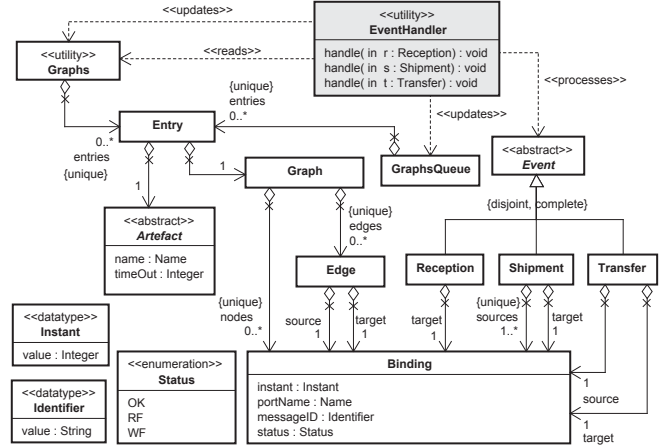


Figure 4. Model of the Event Handler module.

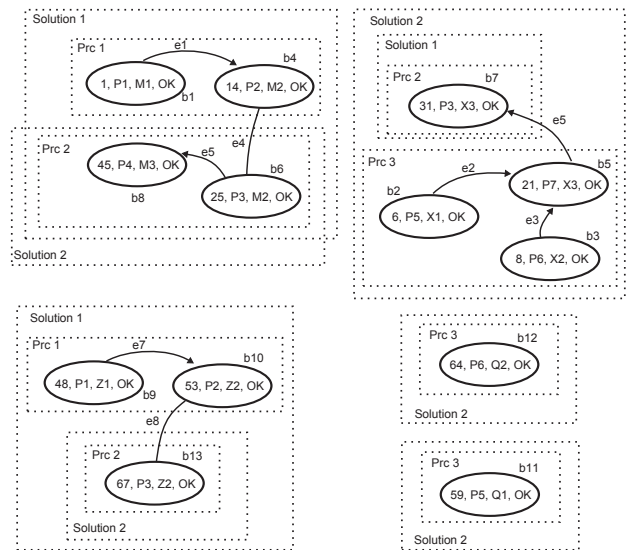


Figure 5. Sample Graphs Database.

involved in process Prc_2 and both solutions, and it denotes that port P_3 dealt with message M_2 at instant 25, and that the result was successful; the later is involved in process Prc_1 and $Solution_1$ only, and it indicates that port P_2 dealt with message M_2 at instant 14, and that the result was successful; furthermore, the edge between them both indicates that binding b_6 originates from binding b_4 .

The Graphs Queue is used to refer to the entries in the Graphs Database that have changed since the database was analysed for the last time. This helps minimise the work performed by the Error Detector, whose abstract model is presented in Figure §6. Note that it is relatively easy to find correlations in a graph like the one in Figure §5 since this task amounts to finding the connected components of the graph [9]. Contrarily, verifying them depends completely on the semantics of the EAI solutions involved. This is why we assign each artefact an upper bound to the total amount of time it is expected to take to produce a valid correlation, i.e., a time out, and a set of rules of the following form, cf. Figure §3:

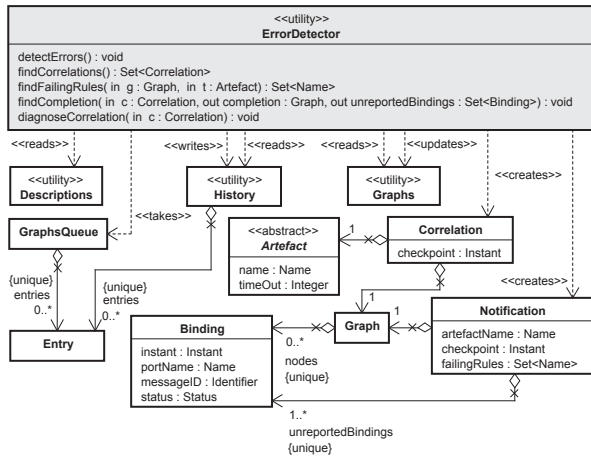


Figure 6. Model of the Error Detector module.

$$P_1[m_1..n_1], \dots, P_p[m_p..n_p] \rightarrow Q_1[r_1..s_1], \dots, Q_q[r_q..s_q],$$

where P_i and Q_j are port names and m_i, n_i, r_i, s_i denote the minimum and maximum number of messages a correlation allows in each port so that it can be considered valid. For instance, a rule like $P_5[1..1], P_6[1..1] \rightarrow P_4[1..10]$ regarding $Solution_2$ in Figure §1 means that it is a requirement for a correlation to be considered valid that it has one message at port P_5 , one message at port P_6 and then 1–10 messages at port P_4 .

The correlations found by the Error Detector module are removed from the Graphs Database and stored in the History Database. This helps us complete them if new messages are reported later.

IV. DETECTING ERRORS

Due to space limitations, we do not provide any additional details on the Registrar or the Event Handler modules. Instead, we focus on the Error Detector, which is the central module. Its algorithm is as follows:

```

1: to detectErrors() do
2:   repeat
3:      $s = \text{findCorrelations}()$ 
4:     for each correlation  $c$  in  $s$  do
5:        $\text{verifyCorrelation}(c)$ 
6:     end for
7:   end repeat
8: end

```

It runs continuously; in each iteration, it first finds a set of correlations and then verifies them sequentially. In the following subsections, we delve into the algorithms to find correlations and to verify them.

A. Finding Correlations

The algorithm to find correlations is as follows:

```

1: to findCorrelations(): Set<Correlation> do
2:   take entry  $f$  from the Graphs Queue
3:    $\text{checkpoint} = \text{getTime}()$ 
4:    $s = \text{find connected components of } f.\text{graph}$ 
5:    $\text{result} = \emptyset$ 

```

```

6:   for each graph  $g$  in  $s$  do
7:      $c = \text{new Correlation}(\text{artefact} = f.\text{artefact}, \text{graph} = g,$ 
8:                            $\text{checkpoint} = \text{checkpoint})$ 
9:     add  $c$  to  $\text{result}$ 
10:  end for
11: end

```

This algorithm starts by taking an entry f from the Graphs Queue at line §2; if there is not an entry available, then we assume that the algorithm blocks here until an entry is available. Note that the core of the algorithm is line §4, in which we find the connected components of the graphs that corresponds to the entry we have taken from the Graphs Queue. We rely on the well-known Hopcroft and Tarjan's algorithm to find connected components since it runs in linear time and has no additional space requirements [9]. In the loop at lines §6–§10, we simply use the components we have found to create the corresponding Correlation objects.

B. Verifying Correlations

A correlation can be diagnosed as on-going, valid or invalid. A correlation is on-going if its deadline has not expired yet. Bear in mind that correlations are analysed within the context of an artefact, which must have an associated time out and set of rules. The deadline for a correlation is defined as the time of its earliest binding plus this time out. This provides a time frame within which all of the messages involved in the correlation are expected to be reported. A correlation is valid if all of the messages it involves were read or written by the expected deadline, there was no reading or writing failure, and all of the rules involved are passed; otherwise, it is considered invalid and a notification must be generated so that it can be diagnosed and appropriate recovery actions can be executed later.

The algorithm to verify a correlation is as follows:

```

1: to verifyCorrelation(in  $c$ : Correlation) do
2:    $\text{findCompletion}(c, \text{out } \text{completedGraph}, \text{out } \text{unnotifiedBindings})$ 
3:    $\text{status} = \text{every binding } b \text{ in } \text{completedGraph}.\text{nodes} \text{ has status OK?}$ 
4:    $\text{earliestInstant} = \text{minimum of } \text{completedGraph}.\text{nodes}.\text{instant}$ 
5:    $\text{latestInstant} = \text{maximum of } \text{completedGraph}.\text{nodes}.\text{instant}$ 
6:    $\text{deadline} = \text{earliestInstant} + c.\text{artefact}.\text{timeOut}$ 
7:    $\text{notPassedRules} = \text{checkRules}(\text{completedGraph}, c.\text{artefact})$ 
8:    $\text{isValid} = \text{deadline} \leq c.\text{checkpoint}$  and  $\text{latestInstant} \leq \text{deadline}$  and
9:              $\text{status} == \text{true}$  and  $\text{notPassedRules} == \emptyset$ 
10:   $\text{isInvalid} = (\text{deadline} < \text{latestInstant})$  or
11:               $(\text{deadline} \leq c.\text{checkpoint}$  and  $(\text{not } \text{status}$  or  $\text{notPassedRules} \neq \emptyset))$ 
12:  if  $\text{isValid}$  then
13:     $f = \text{find the entry for } c.\text{artefact}$  in the Graphs Database
14:    remove  $c.\text{graph}$  from  $f.\text{graph}$ 
15:     $g = \text{new Entry}(\text{artefact} = c.\text{artefact}, \text{graph} = c.\text{graph}, \text{isValid} = \text{true})$ 
16:    add  $g$  to History database
17:  elsif  $\text{isInvalid}$  then
18:     $f = \text{find the entry for } c.\text{artefact}$  in the Graphs Database
19:    remove  $c.\text{graph}$  from  $f.\text{graph}$ 
20:     $g = \text{new Entry}(\text{artefact} = c.\text{artefact}, \text{graph} = \text{completedGraph},$ 
21:                    $\text{isValid} = \text{false})$ 
22:    add  $g$  to History database
23:     $n = \text{new Notification}(\text{artefactName} = c.\text{artefact}.\text{name},$ 

```



```

24:         graph = completedGraph,
25:         unnotifiedBindings = unnotifiedBindings,
26:         checkpoint = c.checkpoint,
27:         notPassedRules = notPassedRules)
28:     send n to the notification port of the monitor
29:   elsif
30:     - Nothing to do, since c is an on-going correlation
31:   end if
32: end

```

The algorithm gets a Correlation c as input; the first thing it has to do is to complete it with the help of the History Database. Note that correlations that are not on-going are removed from the Graphs Database; due to the asynchronous nature of EAI solutions, that implies that after a correlation is verified, additional correlated messages may be reported. This is the reason why before verifying a correlation, it must be completed using the History Database. Algorithm `findCompletion`, which is explained later, performs this task; given a correlation c , it returns a graph that includes $c.graph$ and additional nodes and edges found in the History Database, as well the subset of bindings in the completed correlation that have not been notified, yet. In lines §3–§10 it calculates the status of the correlation, its deadline, the set of rules that are not passed, and determines if the correlation is valid, invalid or on-going. (Note that a correlation is on-going when it is neither valid nor invalid.) If correlation c is found to be valid, we then locate the entry that corresponds to its associated artefact in the Graphs Database, remove the correlation from it, and create a new entry in the History database (lines §13–§16). If it is found to be invalid, the process is similar, but a Notification object is created and sent to the notification port of the monitor so that the correlation can be diagnosed and the appropriate recovery actions can be executed. If it is an on-going correlation, then we just have to wait.

C. Completing Correlations

The algorithm to complete a correlation is as follows:

```

1: to findCompletion(in c: Correlation,
2:   out completedGraph: Graph,
3:   out unnotifiedBindings: Set(Binding) ) do
4:   completedGraph = new Graph(nodes = shallow copy of c.graph.nodes,
5:     edges = shallow copy of c.graph.edges)
6:   unnotifiedBindings = shallow copy of c.graph.nodes
7:   s = find all of the entries for c.artefact in the History database
8:   for each entry f in s do
9:     intersection = c.graph.nodes ∩ f.graph.nodes
10:    if intersection ≠ ∅ then
11:      merge f.graph into completedGraph
12:      if not f.isValid then
13:        remove intersection from unnotifiedBindings
14:      end if
15:    end if
16:  end for
17: end

```

This algorithm takes a correlation c as input and returns a graph that is a completed version of $c.graph$ and a set of bindings that have not been notified so far. It first creates

an initial completed graph at line §4 from a shallow copy of the nodes and the edges of the graph of correlation c . A shallow copy is made because otherwise line §11 would modify the original graph in correlation c . Line §6 also makes a shallow copy of all bindings from correlation c into the set of unnotified bindings, i.e., we initially assume that all of them have been notified. At line §7, the algorithm finds all entries for the artefact associated with correlation c and stores them in variable s . The loop at lines §8–§16 iterates over all of the entries in s ; it discovers if there are common bindings between correlation c and entry f . This is done at line §9 by calculating the intersection amongst the bindings of c and the bindings of f . If the intersection returns a non-empty set, it means that the bindings of f can complete the bindings of c ; in this case, the graph associated with entry f must be merged into the resulting completed graph at line §11. Line §13 removes the bindings that were detected to be already in the graph of entry f from the set of unnotified bindings, leaving only new bindings that were not reported yet. Note that this is done only if graph f represents an invalid graph; otherwise all bindings are new.

D. Checking Rules

The algorithm to check rules is as follows:

```

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

```

This algorithm takes a graph that represents a correlation and an artefact as input; it returns the subset of rules associated with the artefact that the correlation does not pass. The loop at lines §3–§10 iterates over the rules and the internal loop at lines §4–§9 checks every atom. The algorithm is simple since we just need to count the number of bindings that involve the port referenced in the atom; if this figure is not within the margins that the atom establishes, then it is added to the result of the algorithm since that rule is not passed.

V. ANALYSIS OF THE PROPOSAL

In this section, we analyse our proposal from a both a theoretical and a practical point of view. We first prove that it behaves linearly, i.e., it is computationally tractable.

Theorem 1: Algorithm `detectErrors()` takes $O(b + ch)$ time to process an entry in the Graphs Queue, where b denotes the average number of bindings that have been reported since the last checkpoint, c denotes the average number of correlations found at each checkpoint, and h the average number of entries for an artefact in the History Database.

Proof: According to Theorems §2 and §4 in Appendix §A, lines §3 and §5 of the algorithm run in $O(b)$

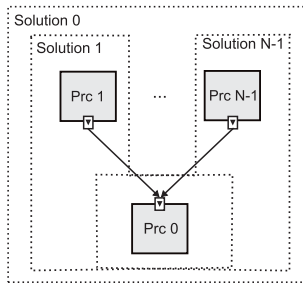


Figure 7. Experimental system.




Figure 8. Experimental results.

and $O(h)$ time, respectively. Assume that `findCorrelations()` returns c correlations in average. Therefore, the loop at lines §4–§6 iterates c times in average. As a conclusion, algorithm `detectErrors()` takes $O(b + ch)$ time to process an entry in the Graphs Queue. ■

Note that there must be an upper bound to the average number of bindings reported between checkpoints; such bound is unknown since it depends on the hardware used, but it exists as long as the Descriptions Database does not change, i.e., no new artefacts are added to the system, and the message production rate is not monotonically increasing. The previous assumptions seem sensible since a typical company does not introduce new artefacts day after day and they cannot grow their hardware continuously. In turn, this implies that there must be an upper bound to the number of correlations that can be found in real-world scenarios. Contrarily, h increases monotonically as time goes by. This implies that after some time, the complexity of algorithm `detectErrors()` is dominated by h , i.e., the algorithm behaves linearly in the average number of entries per artefact in the History Database.

To prove that our algorithm makes sense in practice, we have also carried out a series of experiments. Due to space limitations, we report on one of the worst-case scenarios only, cf. Figure §7. It consists of $N - 1$ processes that report to a single process denoted as `Prc0`. Note that there are N solutions and that process `Prc0` belongs to each of them; consequently, every time a message is sent to this process, $N + 2$ artefacts are involved (`Prc0`, the process that sent the message, and the solutions).

The experiments were run on a machine that was equipped with an Intel Pentium D processors that ran at 3.4Ghz, had 2 GB of RAM memory, Windows Server 2003 (32-bit edition), and JRE 1.6. Each experiment consisted of executing the previous system for 24 hours with a fixed number of processes and a fixed maximum history size; these parameters changed from experiment to experiment. We set the fault rate at 10%, i.e., one out of every 10 messages was not delivered successfully, was intentionally replicated, or contained erroneous data with equal probability. Note that other parameters equal, the fault rate does not have an impact on the efficiency of our algorithm, since checking a correlation takes $O(ra)$ time, where r denotes the number of rules associated with an artefact and a the average number of atoms in these rules, cf. Theorem §5 in Appendix §A. The time to transmit messages was less than a millisecond since we considered small-sized messages that were sent across a high-speed local area network. The message production rate was set to a message per second to simulate a continuously-loaded system.

Figure §8 shows our results; the abscissa reports on the number of processes in the system, which varied from $P = 3$ to $P = 10$, and the ordinate reports on the average time the algorithm to detect errors took; each line represents the results we gathered when we varied the history size from $H = 100$ to $H = 500$ in steps of 100. The conclusion is that adding new processes to the system has obviously an impact on the time to detect errors, since these processes result in additional bindings that need to be processed by our algorithm. The impact is however linear, with a slight slope. In average, the impact of adding a new process to the system is 0.022 ± 0.017 seconds. Increasing the maximum size of the History Database by 100 entries also has an impact of 0.161 ± 0.022 seconds in average.

VI. CONCLUSIONS AND FUTURE WORK

We have presented a proposal to detect errors in the context of EAI solutions. It is novel in that it is not bound to orchestrations or choreographies, neither to a process- nor a task-based execution model; it is totally independent. We have also proven that its time complexity is $O(b + ch)$, where b denotes the average number of bindings that have been reported by means of events since the last checkpoint, c denotes the average number of correlations found at each checkpoint, and h the average number of entries for an artefact in the History Database. Recall that the only purpose of this database is to complete correlations that are found in the Graphs Database, just in case a message is processed by a port after the deadline for the corresponding correlation expires. In practice, it makes sense to remove old information from the database periodically; this puts an upper bound to the size of the History Database, which, in turn, puts an upper bound to the total time the algorithm may take to detect errors. The experimental results prove that not only is the proposal computationally tractable, but also efficient. Future work includes exploring how the proposal may benefit from multi-threading and reducing the amount of work required to analyse the database; note that the same

correlation may be analysed several times in cases in which solutions overlap.

APPENDIX

Theorem 2: Algorithm `findCorrelations()` terminates in $O(b)$ time, where b is the number of bindings that have been reported by means of events since the last checkpoint for a given artefact.

Proof: This algorithm takes an entry from the `Graphs Queue` whenever it is available. Note that this may take an arbitrary time since it depends on the events being handled, which, in turn, depends on the system being monitored. Therefore, the time complexity refers to the time the algorithm takes to process an entry once it is taken from the `Graphs Queue`. The instruction at line §3 runs in $O(1)$ time; contrarily, the instruction at line §4 has to find the connected components of the graph associated with an entry, which is accomplished in $O(\max\{b, r\})$ time [9], where b denotes the number of nodes in the graph being analysed (bindings), and r the number of edges amongst them. In our scenario, it is expected that $b \approx r$ since edges are added to a graph when a shipment event is handled (n source bindings, one target binding, then n edges), or when a transfer event is handled (one source binding, one target binding, then one edge). Thus, without loss of generality, we can assume that line §4 runs on $O(b)$ time. The loop at lines §6–§10 iterates over each connected component to create new correlations, i.e., it runs in $O(c)$ time, where c denotes the average number of connected components found. As a conclusion, `findCorrelations()` terminates in $O(b + c)$ time. Note that b is usually expected to dominate c , since the total number of bindings reported in between checkpoints is proportional to the number of correlations c , i.e., $b = kc$ for an unknown k . Thus, we can conclude that `findCorrelations()` actually runs in $O(b)$ time. ■

Theorem 3: Algorithm `verifyCorrelation(c)` terminates in $O(h)$ time, where h denotes the number of entries in the `History Database` that involve $c.artefact$.

Proof: This algorithm is dominated by the calls to algorithms `findCompletion` and `checkRules` at lines §2 and §7, respectively. The rest of the lines may be assumed to execute in $O(1)$ time since they involve iterating over a completion of a correlation, or manipulating their associated graphs, which is expected to involve a relatively small number of bindings. According to Theorems §2 and §4, lines §2 and §7 are expected to run in $O(h + ra)$ time, where h denotes the number of entries in the `History Database` that involve $c.artefact$, r is the number of rules associated with this artefact, and a is the average number of atoms in these rules. Note that this is expected to be dominated by $O(h)$ as time goes by and the `History Database` grows. Therefore, `verifyCorrelation(c)` terminates in $O(h)$ time. ■

Theorem 4: Algorithm `findCompletion(c, out cg, out ub)` terminates in $O(h)$ time, where h denotes the number of entries for artefact $c.artefact$ in the `History Database`.

Proof: The time complexity of this algorithm is dominated by the loop at lines §8–§16. It iterates over all of the

entries associated with $c.artefact$ in the `History Database`. Let h denote the number of such entries. We can safely assume that the set operations within this loop can be implemented in $O(1)$ time, since they all operate on correlations, which are expected to involve a relatively small number of bindings. `findCompletion(c, out cg, out ub)` thus terminates in $O(h)$ time. ■

Theorem 5: Algorithm `checkRules(g, t)` terminates in $O(ra)$ time, where r denotes the number of rules associated with artefact t and a the average number of atoms in these rules.

Proof: The proof is straightforward since the loop at lines §3–§10 iterates a total of r times, where r denotes the number of rules associates with artefact t , and the loop at lines §4–§9 iterates a times in average, where a denotes the average number of atoms per rule. The algorithm then runs in $O(ra)$ time. ■

REFERENCES

- [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] L. Baresi, S. Guinea, R. Kazhamiakin, and M. Pistore. An integrated approach for the run-time monitoring of BPEL orchestrations. In *ServiceWave*, pages 1–12, 2008
- [3] M. Y. Chen, A. Accardi, E. Kiciman, D. A. Patterson, A. Fox, and E. A. Brewer. Path-based failure and evolution management. In *Networked Systems Design and Implementation*, pages 309–322, 2004
- [4] D. K. W. Chiu, Q. Li, and K. Karlapalem. A meta modeling approach to workflow management systems supporting exception handling. *Inf. Syst.*, 24(2):159–184, 1999
- [5] G. Dunphy and A. Metwally. *Pro BizTalk 2006*. Apress, 2006
- [6] V. Ermagan, I. Krüger, and M. Menarini. A fault tolerance approach for enterprise applications. In *IEEE SCC*, pages 63–72, 2008
- [7] J. B. Goodenough. Exception handling: Issues and a proposed notation. *Commun. ACM*, 18(12):683–696, 1975
- [8] C. Hagen and G. Alonso. Exception handling in workflow management systems. *IEEE Trans. Software Eng.*, 26(10):943–958, 2000
- [9] J. E. Hopcroft and R. E. Tarjan. Efficient algorithms for graph manipulation [h] (algorithm 447). *Commun. ACM*, 16(6):372–378, 1973
- [10] C. Ibsen and J. Anstey. *Camel in Action*. Manning Publications, 2010
- [11] L. Li, C. N. Hadjicostis, and R. S. Sreenivas. Designs of bisimilar petri net controllers with fault tolerance capabilities. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 38(1):207–217, 2008
- [12] C. Liu, M. E. Orłowska, X. Lin, and X. Zhou. Improving backward recovery in workflow systems. In *Database Systems for Advanced Applications*, pages 276–286, 2001
- [13] M. Sampath, R. Sengupta, and S. Lafortune. Failure diagnosis using discrete-event models. *IEEE Trans. on Control Syst. Technol.*, 4(2):105–124, 1996
- [14] Y. Yan, M.-O. Cordier, Y. Pencolé, and A. Grastien. Monitoring Web service networks in a model-based approach. In *International Conference on Web Services*, pages 192–203, 2005
- [15] Y. Yan and P. Dague. Modeling and diagnosing Orchestrated Web service processes. In *International Conference on Web Services*, pages 51–59, 2007