

Mining Lifecycle Event Logs for Enhancing Service-based Applications

Franco Maria Nardini¹, Gabriele Tolomei^{1,3}, Philipp Leitner², Fabrizio Silvestri¹, Schahram Dustdar²

¹ISTI-CNR
Pisa, Italy

{firstname.lastname}@isti.cnr.it

²Distributed Systems Group
Vienna University of Technology, Vienna, Austria

{lastname}@infosys.tuwien.ac.at

³Ca' Foscari University
Venice, Italy

{lastname}@dsi.unive.it

Abstract—*Service-Oriented Architectures (SOAs), and traditional enterprise systems in general, record a variety of events (e.g., messages being sent and received between service components) to proper log files, i.e., event logs. These files constitute a huge and valuable source of knowledge that may be extracted through data mining techniques. To this end, process mining is increasingly gaining interest across the SOA community. The goal of process mining is to build models without a priori knowledge, i.e., to discover structured process models derived from specific patterns that are present in actual traces of service executions recorded in event logs. However, in this work we focus on detecting frequent sequential patterns, thus considering process mining as a specific instance of the more general sequential pattern mining problem. Furthermore, we apply two sequential pattern mining algorithms to a real event log provided by the Vienna Runtime Environment for Service-oriented Computing, i.e., VRESCO. The obtained results show that we are able to find services that are frequently invoked together within the same sequence. Such knowledge could be useful at design-time, when service-based application developers could be provided with service recommendation tools that are able to predict and thus to suggest next services that should be included in the current service composition.*

I. INTRODUCTION

The vast majority of nowadays software-based systems, ranging from the simplest, i.e., small-scale, to the most complex, i.e., large-scale, record massive amounts of data in the form of *logs*. Such logs could either refer to the functioning of the system as well as keep trace of any possible software or human interaction with the system itself. For this reason, logs represent a valuable source of hidden knowledge that can be exploited in order to enhance the overall performances of any software-based system.

Well-known examples of systems that have started trying to improve their performances by analyzing event logs are surely Web Search Engines (SEs). Roughly, SEs are increasingly exploiting past user behaviors recorded in *query logs* in order to better understand people search intents, thus, for providing users with better search experiences. Indeed, by accurately recognizing and predicting actual user information needs, SEs are now able to offer more sophisticated functionalities (e.g., *query suggestion*) as well as better relevant result sets in response to a specific query (e.g., *query diversification*).

Moreover, there are plenty of modern enterprise software systems that need to operate in highly dynamic and distributed environments in a standardized way. Such systems implement their business logic according to the *Service-oriented Architecture* (SOA) principles, thus, assembling their business processes as the composition and orchestration of autonomous, protocol-independent, and distributed logic units, i.e., *software services*.

Service-based systems and applications (SBAs) require proper *run-time* environments where their composing services can be searched, bound, invoked, monitored and managed. Therefore, SBA's run-time support might keep track of what is going on during the whole application lifecycle by roughly recording all such events to log files, i.e., *service event logs*.

Analysis of such service event logs could reveal interesting *patterns*, which in turn might be exploited for improving the overall performances of SOA's run-time frameworks as well as supporting SBA designers during the whole application lifecycle.

The main contribution of this work concerns the application of data mining techniques to a real-life service event log collected by the VRESCO SOA run-time framework. Our aim is to analyze the historical events stored on VRESCO in order to discover software services that are frequently invoked and composed together, i.e., *process mining*.

Although traditional process mining refers to a set of techniques and methodologies whose aim is to distill a structured process description from a set of actual traces of executions recorded in event logs, here we treat it as an instance of the *sequential pattern mining problem*.

The remaining of the paper is structured as follows. Section II describes the information collected by SOA lifecycle event logs, in particular focusing on the VRESCO run-time framework. In Section III, we propose how VRESCO event log may be analyzed for approaching our research challenge. Therefore, Section IV shows the experiments we conduct on a real VRESCO log data set. In Section V, we describe past work that somehow concerns with service event log analysis. Finally, Section VI summarizes the contributions we provide in this work together with any further idea that could be better investigated as future work.

II. SOA LIFECYCLE EVENT LOGS

In the following, we will discuss SOA lifecycle event logs as they are used in this paper. We then present the lifecycle events emitted by the VRESCO SOA runtime environment as a concrete example used in Section III and Section IV of this paper.

A. Lifecycle Events

Events and complex event processing [11] (CEP) are frequently used tools to document and track the lifecycle of applications in various domains. For instance, in the business domain the idea of business activity monitoring [9] (BAM) uses events to monitor business process performance. Analogously, technical implementations of business processes on top of SOAs (service compositions) are often monitored using CEP. To this end, many service composition engines can be configured to track their current state in event logs. For instance, the Apache ODE WS-BPEL engine triggers a rich model of execution events¹. Similarly, service compositions implemented using Windows Workflow Foundation can use the .NET tracking service² to persist event logs. However, tracking system state via event logs in SOA is not confined to composition engines. For instance, The Vienna Runtime Environment for Service-Oriented Computing (VRESCO) [4] uses events to track not only service compositions, but all entities and interactions in a SOA (services, users, compositions, metadata and interactions).

In its most general form, an event log \mathcal{E} consists of a sequence of n recorded events, i.e., $\mathcal{E} = \langle e_1, e_2, \dots, e_n \rangle$. Each event $e_i \in \mathcal{E}$ usually contains at least an unique identifier, an event timestamp, the publisher of the event (e.g., the BPEL engine), the subject of the event (e.g., the composition instance that triggered the event), and the event type. Depending on the concrete event type, more detailed information is available. This type-specific information cannot be describe generally, i.e., it is different from event type to event type as well as from system to system. In the following we describe the event types triggered by the VRESCO system as an example of the possibilities provided by event logs.

B. SOA Event Log: VRESCO

VRESCO is an experimental runtime environment developed at Vienna University of Technology. VRESCO is being developed under an open source license, and can be accessed via the project Web page³. The project aims at solving some of the research problems identified in [15], e.g., dynamic selection of services based on Quality-of-Service (QoS), dynamic rebinding and service composition, service metadata and event-based services computing.

¹<http://ode.apache.org/ode-execution-events.html>

²[http://msdn.microsoft.com/en-us/library/ms735887\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/ms735887(v=vs.85).aspx)

³<http://www.infosys.tuwien.ac.at/prototypes/VRESCO/>

In the following we focus on the latter aspect. The foundations of event-based service-oriented computing have been discussed in [3], [5]. In a nutshell, the goals of this earlier work were to track what is going on in a service-based application by constantly triggering events and using CEP to construct meaningful information from those events.

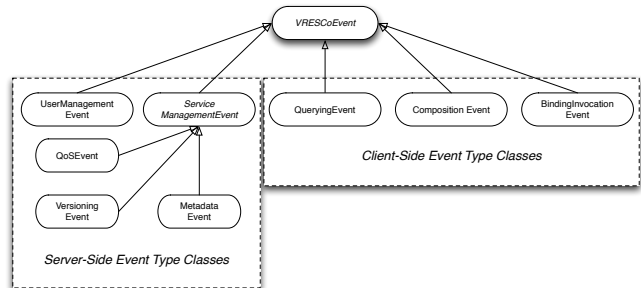


Figure 1. VRESCO Event Type Classes

In a VRESCO system, events of various types are triggered. A simplified taxonomy of event type classes is depicted in Figure 1. As can be seen, events are triggered when services are queried, bound and invoked. Additionally, events indicate if the data or metadata about services changes (e.g., the QoS is changed, new operations are available). Each of the concrete event type classes (those with non-italic name) in turn contains a number of concrete event types that can be triggered. For full details on all events refer to [3].

Events in VRESCO can be triggered either on client- or server-side. While events concerning metadata are triggered by the VRESCO server, all querying and invocation events are triggered by clients and only processed throughout the VRESCO event engine. These client-triggered events are listed in more detail in Table I. In the table, we provide the condition that triggers each event along with the event type and event type class of the event. All of these events provide the basic information discussed above (sequence number, timestamp, ...). In addition, events generally provide some type-specific additional information, which we also summarize in the table. For reasons of brevity, we have omitted composition events. Composition events in VRESCO are of comparable expressiveness as the events triggered by Apache ODE.

The VRESCO event engine stores triggered events in an event log. Therein, events are serialized as XML and can be accessed and analyzed via a RESTful service interface. In Listing 1 we provide an example event serialized to XML. Evidently, the event reflects a service invocation with a very simple input message (`<order>`) as payload.

III. LIFECYCLE EVENT LOG MINING

In this work, we are interested in exploring how event log collected by the VRESCO framework, i.e., *service event*

Table I
CLIENT-TRIGGERED VRESCO EVENTS

Event Type Class	Event Type	Event Condition	Additional Event Information
<i>BindingInvocationEvent</i>	ServiceInvokedEvent	Specific service is invoked	Message sent to service Invoking user
	ServiceInvocationFailedEvent	Service invocation failed	Message sent to service Triggered fault
	ProxyRebindingEvent	Service proxy is (re-)bound to a new service	Old service New service
<i>QueryingEvent</i>	RegistryQueriedEvent	Registry is queried	Query string
	ServiceFoundEvent	Specific service is found by a query	Query string Query results
	NoServiceFoundEvent	No services are found by a query	Query string

```

1 <ServiceInvokedEvent xmlns="http://www.vitalab.tuwien.
  ac.at/vresco/usertypes">
2   <Priority>0</Priority>
3   <Publisher>guest</Publisher>
4   <PublisherGroup>GuestGroup</PublisherGroup>
5   <UserName>a007b09b-8c23-4fac-af30-0142a61f3795</
  UserName>
6   <SeqNum>74006756-64f1-40cb-858e-565d4bc6a94c:24</
  SeqNum>
7   <Timestamp>2010-11-09T09:58:49</Timestamp>
8   <CurrentRevisionId>180</CurrentRevisionId>
9   <CurrentRevisionWsdL>
10    http://localhost:60000/AssemblyAtomicServices/
    IAssemblingPlanningService?wsdl
11  </CurrentRevisionWsdL>
12  <FeatureName>GetPartFeature</FeatureName>
13  <InvocationInfo>
14    <service_input>
15      <order><part1>text</part1></order>
16    </service_input>
17  </InvocationInfo>
18 </ServiceInvokedEvent>

```

Listing 1. Serialized Invocation Event

log, could be harnessed for better supporting service-based applications during their whole lifecycle.

Roughly, analysis of VRESCO event log data is finalized to the discovery of *sequences* of services that are frequently invoked together, thus, to detect *processes*, or part of those, as result of the compositions of highly co-invoked services.

This knowledge could be useful for improving the overall performances of the VRESCO run-time framework as well as supporting SBA designers during the whole application lifecycle. As an example, service-based application developers could be provided with *service recommendation tools* that are able to predict and, thus, to suggest next services that should be included in the current service composition at *design-time*.

Due to the huge amount of data collected in the VRESCO event log, data mining techniques represent a suitable approach for addressing our research challenge. In the following, we describe how processes may be mined from the VRESCO event log, namely how sequences of co-invoked services that frequently appear in the event log may be discovered and exploited.

A. Process Mining

According to van der Aalst *et al.* [19], the term *process mining*, also referred to as *workflow mining*, describes a set of techniques and methodologies whose aim is to distill a structured process description from a set of actual traces of executions recorded in event logs.

In our vision, a service log might be viewed as a database consisting of sequences of events that change with time, i.e., a *time-series database* [8]. Such kind of database records the valid time of each data set. For example, in a time-series database that records service invocation transactions, each transaction includes the unique identifier of the invoked service as well as an extra time-stamp attribute indicating when the event happened [24].

Several kinds of patterns can be extracted from various types of time-series data. In this work, we are interested in finding *sequences* of services that are frequently invoked together in a specific order, i.e., *sequential patterns* [2].

Each process instance recorded on event logs might be expressed as an *unrolled* trace of invoked services. Thus, let $S = \{s_1, s_2, \dots, s_m\}$ be a set of services and let $S^j \subseteq S$ be an itemset of services invoked at the same time (or within a small time window), i.e., $S^j = \{s_{j_1}, s_{j_2}, \dots, s_{j_k}\}$. Therefore, a process $p_j = \langle S^j_1, S^j_2, \dots, S^j_{|p_j|} \rangle$ has a unique identifier and represents a sequence of service itemsets, chronologically-ordered according to their time-stamps. Globally, a process database is a sequence database $P = \{p_1, p_2, \dots, p_n\}$ of executed and recorded processes.

Roughly, a sequential pattern is a sequence of service itemsets that occur *frequently* in P according to a specific order, i.e., appear as subsequence of a large percentage of sequences of P . More formally, a sequence $p' = \langle S'_1, S'_2, \dots, S'_u \rangle$ is a *subsequence* of $p'' = \langle S''_1, S''_2, \dots, S''_v \rangle$, i.e., $p' \preceq p''$, if there exists integers $1 \leq i_1 < \dots < i_u \leq v$ such that $\forall_{1 \leq j \leq u} S'_j \subseteq S''_{i_j}$.

Then we may define the support $supp(p')$ of a sequence of service itemsets p' as the proportion of processes in the database P that contains p' as its subsequence, that is:

$$supp(p') = \frac{|\{p_j \mid p' \preceq p_j\}|}{|P|}$$

Therefore, *sequential pattern mining* is the process of extracting certain sequential patterns whose support exceed a predefined minimal support threshold min_supp .

To this end, as a first approach we apply one of the more efficient sequence pattern mining algorithm to the VRESCO event log, namely *PrefixSpan* [16].

Furthermore, we extend our first approach in order to exploit the temporal information associated with each service invocation in a different way. Indeed, in traditional sequential pattern mining, event time-stamps are only used for establishing the chronological order between service invocations, i.e., for simply stating that service s_i is invoked *before* service s_j . However, observing that s_i and s_j are invoked really closed to each other, e.g., within 5 seconds, rather than noting that s_i and s_j are farther away from each other, e.g. 5 minutes, could lead to different conclusions. Thus, we apply another sequential pattern mining algorithm, which is able to deal with this issue, i.e., *MiSTA* [7].

Finally, in Section IV, we describe the different results we obtain on the VRESCO event log when using the two approaches described above.

IV. EXPERIMENTATION

In order to test our claim about finding frequent sequential patterns inside service event logs, we use a real-life log of events collected by the VRESCO runtime framework.

This event log consists of 89 transactions. Each transaction is in turn composed of several events according to the ones described in Table I (e.g., *ServiceInvokedEvent*, *ServiceInvocationFailedEvent*, etc.).

For the sake of our purposes, namely for discovering sequences of services that frequently are invoked together, we only consider the list of *ServiceInvokedEvent* for each transaction.

Firstly, we run the *PrefixSpan* algorithm [16] on the VRESCO data set. We use several thresholds on the *minimum support* of the sequential patterns to be extracted, ranging from 20% to 100%. However, the maximum support for frequent sequences found in our data set is 66%.

The following figures show the distribution of the lengths of extracted sequence patterns, i.e., the number of invoked services that are inside a frequent sequence. In particular, Figure 2, Figure 3, and Figure 4 show the results obtained by using a minimum support threshold of 25%, 50%, and 66%, respectively.

As one would expect, independently from the minimum support chosen, 2-length sequences are the most popular since for each k -length frequent sequences ($k > 2$) any i -length subsequence, i.e., $2 \leq i < k$, is also frequent by definition. Moreover, as the minimum support threshold increases, the maximum length of frequent sequences decreases as well from 17 to 11, and, also, most popular frequent sequences result to be globally shorter.

Finally, on average frequent sequences are composed of 5.86, 4.60, and 4.07 services for minimum support values of 25%, 50%, and 66%, respectively.




Figure 2. *PrefixSpan* [$min_supp = 25\%$]: Pattern length distribution.

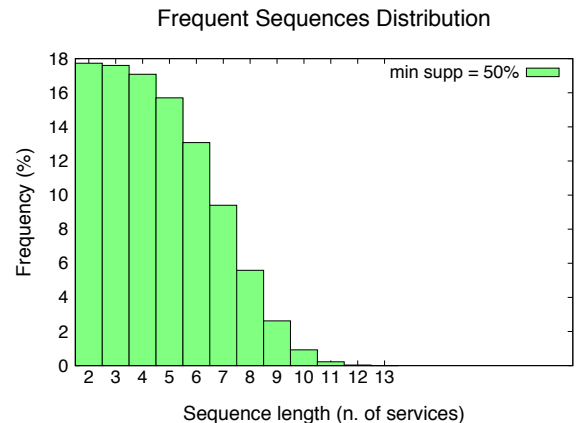


Figure 3. *PrefixSpan* [$min_supp = 50\%$]: Pattern length distribution.

As a second step of our experimental phase, we also run the *MiSTA* algorithm [7] on the VRESCO event log. This algorithm differs from classical sequential pattern mining algorithms like *PrefixSpan* because it also takes care of the time gaps between consecutive items in a sequence. In other words, *MiSTA* extracts frequent sequences by considering not only the *minimum support* but also another parameter, i.e., τ , which basically represents a threshold on the time-gap between pairs of consecutive items.

In our experiments, we use several combinations of such two parameters, namely min_supp and τ . In the following, we describe the distribution of frequent sequence size obtained by using two values for the minimum support, i.e.,




Figure 4. PrefixSpan [$min_supp = 66\%$]: Pattern length distribution.

20% and 32%, and three values for the time threshold τ , i.e., 5, 60, and 300 seconds, which are shown in Figure 5, 6, 7, 8, 9, 10, respectively.

The total number of frequent sequences found when $min_supp = 20\%$ is 33, both for $\tau = 5$ and $\tau = 60$, whereas it reaches the value of 35 when $\tau = 300$. Moreover, when $min_supp = 32\%$ the total amount of frequent sequences is 20, both for $\tau = 5$ and $\tau = 60$, whereas it reaches the value of 27 when $\tau = 300$.

Finally, the maximum size of frequent sequences is 7, whereas the average ranges from a minimum of 2.80 to a maximum of 3.07. These results show that, by taking into account the time gap between service invocations using *MiSTA* instead of *PrefixSpan*, we are able to detect less and shorter frequent sequences on average.

We argue that *MiSTA* could provide better results if we previously analyze how time gaps are distributed across consecutive service invocations, and we leave this as a possible future work.

V. RELATED WORK

In this paper, we present a use case for event log mining in service-based systems. This idea bears some resemblance to the established idea of business activity management [9] (BAM). BAM considers the event-driven governance of business processes, and is, hence, mostly a term from the business domain. Technically, BAM is enabled by monitoring runtime monitoring of services and their interactions within company SOAs. To this end, event-based monitoring approaches [22], [21], [6] produce a steady stream of low-level lifecycle events, similarly to the lifecycle events discussed in Section II-B and to the event logs produced by VRESCO. These low-level events need to be aggregated so that real business information can be gained from them. Existing techniques to do this include SLA

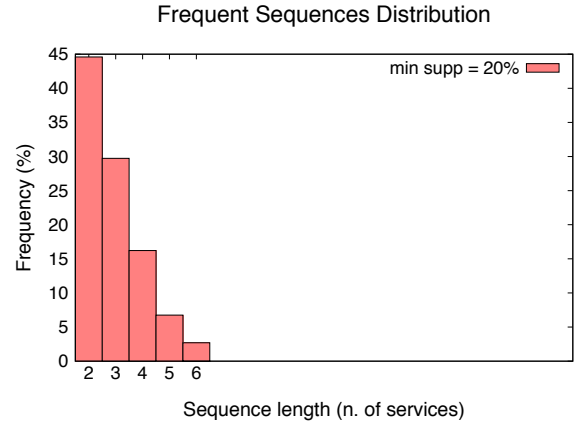


Figure 5. MiSTA [$min_supp = 20\%$, $\tau = 5$ sec.]: Pattern length distribution.




Figure 6. MiSTA [$min_supp = 20\%$, $\tau = 60$ sec.]: Pattern length distribution.

aggregation [18] or event-based SLA monitoring [17], [13]. Related to the ideas of BAM is research work by [14], which considers event-based monitoring of business compliance. Our research, specifically mining for invocation sequences that lead to failed service invocations, is complementary to BAM. While BAM is mostly concerned with discovering failures, our research can be used to identify or predict them in advance.

In literature, events emitted by service-based applications have found various other uses. For instance, [12] uses eventing information (along with various other data sources) to generate visualizations of the past behavior and quality of Web services, mostly to ease management and selection of services for other uses. This is related to our work, which identifies services which have in the past been used together, mostly to suggest suitable services for new uses. [23],




Figure 7. MiSTA [$min_supp = 20\%$, $\tau = 300$ sec.]: Pattern length distribution.

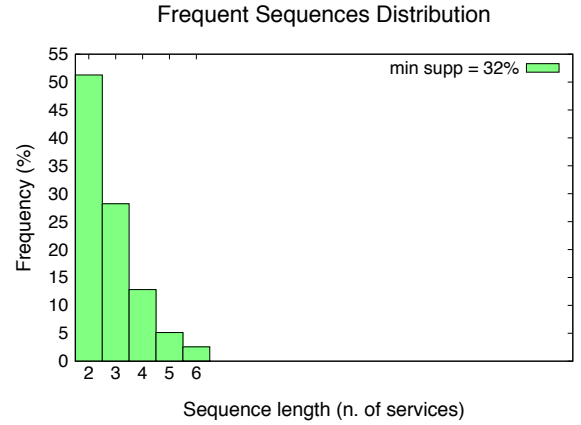


Figure 9. MiSTA [$min_supp = 32\%$, $\tau = 60$ sec.]: Pattern length distribution.




Figure 8. MiSTA [$min_supp = 32\%$, $\tau = 5$ sec.]: Pattern length distribution.




Figure 10. MiSTA [$min_supp = 32\%$, $\tau = 300$ sec.]: Pattern length distribution.

[10] have used SOA events to generate predictions of SLA violations. This is done by training machine learning models from the collected event data, and using runtime event information as input to those models.

Finally, our research is related to the idea of process mining [20], [1], [19]. Generally speaking, process mining considers discovering structures (mostly business processes) from traces of earlier executions of information systems. This also includes making implicit processes (that people are subconsciously following) explicit, so that they can be optimized using the techniques of business process reengineering. Our work is different to process mining in the sense that we do not suggest new processes from event logs. Instead, we rather give recommendations which combinations of services have in the past been used together (indicating that it might make sense to use them in combination).

VI. CONCLUSIONS AND FUTURE WORK

Service-Oriented Architectures (SOAs), and traditional enterprise systems in general, record a variety of events (e.g., messages being sent and received between service components) to proper log files, i.e., *event logs*. These files constitute a huge and valuable source of knowledge that may be extracted through data mining techniques.

In this work, we focus on *process mining* as a specific instance of the more general *sequential pattern mining* problem. Basically, our aim is to detect *frequent sequential patterns* that might be present in actual traces of service executions recorded in event logs. To this end, we apply two sequential pattern mining algorithms to a real event log provided by the Vienna Runtime Environment for Service-oriented Computing, i.e., VRESco.

The obtained results show that we are able to find services that are frequently invoked together within the same sequence. Such knowledge could be useful at *design-time*, when service-based application developers could be provided with *service recommendation tools* that are able to predict and thus to suggest next services that should be included in the current service composition.

Finally, as a future work, we aim at discovering possible sequences of invoked services, which frequently lead to failures or unexpected behaviors. This knowledge could be exploited either for preventing SBA designers to deploy possible faulty service compositions as well as for devising novel run-time adaptation mechanisms in response to undesired events.

ACKNOWLEDGEMENT

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement 215483 (S-Cube).

REFERENCES

- [1] W. M. P. V. D. Aalst, B. F. V. Dongen, C. Günther, A. Rozinat, H. M. W. Verbeek, and A. J. M. M. Weijters, "ProM : The Process Mining Toolkit," in *Industrial Engineering*, vol. 489, Technische Universiteit Eindhoven Eindhoven, The Netherlands. CEUR-WS.org, 2009, pp. 1–4. [Online]. Available: http://prom.win.tue.nl/research/wiki/_media/publications/prom2009bpmdemo.pdf
- [2] R. Agrawal and R. Srikant, "Mining sequential patterns," in *ICDE '95*. IEEE, 1995, pp. 3–14.
- [3] Anton Michlmayr and Florian Rosenberg and Philipp Leitner and Schahram Dustdar, "Advanced Event Processing and Notifications in Service Runtime Environments," in *Proceedings of the 2nd International Conference on Distributed Event-Based Systems (DEBS'08)*, 2008.
- [4] —, "End-to-End Support for QoS-Aware Service Selection, Binding, and Mediation in VRESCO," *IEEE Transactions on Services Computing*, vol. 3, pp. 193–205, July 2010.
- [5] Anton Michlmayr and Philipp Leitner and Florian Rosenberg and Schahram Dustdar, "Publish/Subscribe in the VRESCO SOA Runtime," in *Proceedings of the 2nd International Conference on Distributed Event-Based Systems (DEBS'08)*, 2008, invited demo paper.
- [6] L. Baresi, S. Guinea, M. Pistore, and M. Trainotti, "Dynamo + Astro: An Integrated Approach for BPEL Monitoring," in *Proceedings of the 2009 IEEE International Conference on Web Services (ICWS'09)*. Washington, DC, USA: IEEE Computer Society, 2009, pp. 230–237.
- [7] F. Giannotti, M. Nanni, D. Pedreschi, and F. Pinelli, "Mining sequences with temporal annotations," in *Proceedings of the 2006 ACM Symposium on Applied Computing*, ser. SAC '06. New York, NY, USA: ACM, 2006, pp. 593–597. [Online]. Available: <http://doi.acm.org/10.1145/1141277.1141413>
- [8] J. Han and M. Kamber, *Data Mining: Concepts and Techniques*. Morgan Kaufmann, 2000.
- [9] H. Kochar, "Business activity monitoring and business intelligence," 2005. [Online]. Available: <http://www.ebizq.net/topics/bam/features/6596.html>
- [10] P. Leitner, B. Wetzstein, F. Rosenberg, A. Michlmayr, S. Dustdar, and F. Leymann, "Runtime Prediction of Service Level Agreement Violations for Composite Services," in *Proceedings of the 3rd Workshop on Non-Functional Properties and SLA Management in Service-Oriented Computing (NFPSLAM-SOC'09)*, 2009, pp. 176–186.
- [11] D. C. Luckham, *The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 2001.
- [12] A. Michlmayr, F. Rosenberg, P. Leitner, and S. Dustdar, "Selective Service Provenance in the VRESCO Runtime," *International Journal of Web Services Research*, vol. 7, no. 2, pp. 65–86, 2008. [Online]. Available: <http://www.scopus.com/inward/record.url?eid=2-s2.0-77954090267&partnerID=40&md5=10b3d8328dbd86f8a2f623dd1d6cf886>
- [13] Michlmayr, Anton and Rosenberg, Florian and Leitner, Philipp and Dustdar, Schahram, "Comprehensive QoS Monitoring of Web Services and Event-Based SLA Violation Detection," in *Proceedings of the 4th International Workshop on Middleware for Service Oriented Computing (MWSOC'09)*. New York, NY, USA: ACM, 2009, pp. 1–6.
- [14] E. Mulo, U. Zdun, and S. Dustdar, "Monitoring Web Service Event Trails for Business Compliance," in *IEEE International Conference on Service-Oriented Computing and Applications (SOCA)*, 2009, pp. 1–8.
- [15] M. P. Papazoglou, P. Traverso, S. Dustdar, and F. Leymann, "Service-Oriented Computing: State of the Art and Research Challenges," *IEEE Computer*, vol. 40, no. 11, pp. 38–45, November 2007.
- [16] J. Pei, J. Han, B. Mortazavi-Asl, and H. Pinto, "Prefixspan: Mining sequential patterns efficiently by prefix-projected pattern growth," in *ICDE '01*. IEEE, 2001.
- [17] A. Sahai, V. Machiraju, M. Sayal, A. P. A. van Moorsel, and F. Casati, "Automated SLA Monitoring for Web Services," in *Proceedings of the 13th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM)*, 2002.
- [18] T. Unger, F. Leymann, F. Leymann, and T. Scheibler, "Aggregation of Service Level Agreements in the Context of Business Processes," in *Proceedings of the Twelfth IEEE Enterprise Distributed Object Conference (EDOC 2008)*, vol. 0, Universität Stuttgart, Fakultät Informatik, Elektrotechnik und Informationstechnik, Germany. IEEE Computer Society, 2008, pp. 43–52. [Online]. Available: http://www2.informatik.uni-stuttgart.de/cgi-bin/NCSTR/_view.pl?id=INPROC-2008-93&eng=0

- [19] W. van der Aalst, T. Weijters, and L. Maruster, "Workflow Mining: Discovering Process Models from Event Logs," *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 9, pp. 1128–1142, Sep. 2004. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1316839>
- [20] W. van der Aalst, "Process mining: a research agenda," *Computers in Industry*, vol. 53, no. 3, pp. 231–244, Apr. 2004. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0166361503001945>
- [21] B. Wetzstein, S. Strauch, and F. Leymann, "Measuring Performance Metrics of WS-BPEL Service Compositions," in *Proceedings of the Fifth International Conference on Networking and Services (ICNS'09)*. IEEE Computer Society, Apr. 2009.
- [22] L. Zeng, H. Lei, and H. Chang, "Monitoring the QoS for Web Services," in *Proceedings of the 5th International Conference on Service-Oriented Computing (ICSOC'07)*. Berlin, Heidelberg: Springer-Verlag, 2007, pp. 132–144.
- [23] L. Zeng, C. Lingenfelder, H. Lei, and H. Chang, "Event-Driven Quality of Service Prediction," in *Proceedings of the 6th International Conference on Service-Oriented Computing (ICSOC'08)*. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 147–161.
- [24] Q. Zhao and S. S. Bhowmick, "Sequential pattern matching: A survey," 2003.