Information Technology

Ereigniskorrelation in heterogenen Umgebungen

Event Correlation in Heterogeneous Environments

Dipl.-Inf. Björn Schilling: IPVS, Universitätsstr., 38, D-70569 Stuttgart, Germany
Tel: +49-711-7816-226, Fax: +49-711-7816-424, E-Mail: bjoern.schilling@ipvs.uni-stuttgart.de
Dipl.-Inf. Björn Schilling received his diploma in software engineering from the University of Stuttgart in 2008. Since then, he is working as a research assistant in the Distributed Systems group. His Ph.D. research topic deals with complex event processing in distributed systems.

Dr. Udo Pletat: IBM Deutschland Research & Development GmbH, Schönaicher Str., 220, D-71034 Böblingen, Germany
Tel: +49-7031-16-2146, Fax: , E-Mail: pletat@de.ibm.com
Dr. Udo Pletat received his Ph. D. in Computer Science from University of Stuttgart in 1984. In 1986 he has been a vising scientist at the Technical University of Denmark and in 1988 at the Computer Science department of Stanford University. Since 1986 Udo Pletat has been working for IBM Germany, starting at the Scientific Center Heidelberg/Stuttgart and since 1992 at the IBM Development lab in Böblingen. He has lead various projects in areas like artificial intelligence/knowledge representation, embedded systems development tools, and manufacturing execution systems. Since 1999 he is a certified Project Management Professional (by Project Management Institute) and since 2008 he is also an IBM Senior Certified IT Specialist. His current interests include semantic information models for industrial applications, complex event processing and the interfaces between service-oriented and event-driven software architectures. Udo Pletat has published 20+ scientific articles at conferences and in journals, he is a member of Gesellschaft für Informatik (GI).

Prof. Dr. Kurt Rothermel: IPVS, Universitätsstr., 38, D-70569 Stuttgart, Germany
Tel: +49-711-7816-434, Fax: +49-711-7816-424, E-Mail: kurt.rothermel@ipvs.uni-stuttgart.de
Kurt Rothermel received his doctoral degree in Computer Science from University of Stuttgart in 1985. From 1986 to 1987 he was "Post-Doctoral Fellow" at IBM Almaden Research Center in San Jose, U.S.A. and then joined IBM's European Networking Center in Heidelberg. Since 1990 he is a Professor for Computer Science at University of Stuttgart. He is Director of Institute of Parallel and Distributed Systems and head of Centre of Excellence Nexus, which is funded by German Science Foundation (DFG) and is conducting research in the area of mobile context-aware systems. His research interest is in the field on distributed and mobile systems, where he contributed more than 170 publications.

Keywords: C.2.1 Network Architecture and Design: Heterogeneous Systems and Networks; C.2.4 Distributed Systems: Distributed Applications, Distributed Complex Event Processing

MS-ID: Heft: / ()	bjoern.schilling@ipvs.uni-stuttgart.de	7. Juli 2009

Published in it - Information Technology 5/2009, pages 270-275, 2009.
© 2009 Oldenbourg Wissenschaftsverlag GmbH, München
http://www.oldenbourg-wissenschaftsverlag.de/olb/de/1.c.1495462.de

Abstract

The condition and location of a business item are of central interest in many business applications such as supply chain management, manufacturing, or ensuring safety and security for people and goods. Recent advances in sensor technology allow to transmit condition and location information about goods, materials and people to enterprise software systems in real-time. In this context, complex event processing is an emerging software technology for detecting business-relevant situations in streams of events and for providing these detected situations to various business processes. While currently complex event processing systems are mostly deployed within a single business domain at a limited scale, the cooperative nature of business applications gives reason that complex event processing will soon address multiple business domains and involve an increasingly large number of business events. In order to ensure interoperability as well as efficient utilization of processing and network capability, we motivate the need for heterogeneous correlation technology in the context of business applications.

In this article we give an overview of the project *Distributed heterogeneous event processing* (DHEP) involving the IBM Böblingen lab and the Universität Stuttgart. In particular, we highlight how business applications can benefit from using event correlation technology in heterogeneous environments. The key aspects of the project address the deployment of collections of event correlation rules to a network of heterogeneous event correlation engines. We give an overview of challenges and possible solutions for the dynamic configuration of such environments and present our architecture which supports network-wide cooperation between different correlation engines.

Zusammenfassung

In vielen Geschäftsanwendungen, wie sie beispielsweise in der Logistik, zum Steuern von Produktionsanlagen oder im Bereich der Sicherheitsüberwachung von Personen und Gebäuden eingesetzt werden, sind Zustand und Position von Objekten von großem Interesse. Jüngste Fortschritte im Bereich der Sensortechnik machen das Wissen über Zustand und Position in Echtzeit verfügbar. Komplexe Ereignisverarbeitung (Complex Event Processing) bietet in diesem Bereich eine Möglichkeit, die entstehenden großen Datenmengen zu verarbeiten und bestimmte Situationen in den Ereignisströmen zu erkennen. Derzeit ist die Verarbeitung typischerweise noch zentral und somit auf einzelne Bereiche beschränkt. Allerdings entstehen durch die zunehmenden Kooperationen von Geschäftspartnern neue Anforderungen für die komplexe Ereignisverarbeitung. Um eine Interoperabilität und effiziente Ereignisverarbeitung auch in solchen geschäftsübergreifenden Szenarien gewährleisten zu können, schlagen wir eine heterogene Ereignisverarbeitung im Kontext von Geschäftsanwendungen vor.

Dieser Artikel gibt eine Übersicht über das Projekt *Distributed Heterogeneous Event Processing* (DHEP), das derzeit vom IBM Forschungs- und Entwicklungszentrum Böblingen gemeinsam mit der Universität Stuttgart durchgeführt wird. Es wird aufgezeigt, wie Geschäftsanwendungen von einer Ereigniskorrelation in heterogenen Umgebungen profitieren können. Im Hauptaspekt beschäftigt sich das Projekt mit der Konfiguration und Verteilung von Korrelationsregeln in einem Netzwerk von heterogenen Korrelationsknoten. Wir geben eine Übersicht über die sich ergebenden Anforderungen und schlagen mögliche Lösungen für eine dynamische Konfiguration vor. Desweiteren präsentieren wir unsere Architektur, die eine netzwerkweite Kooperation von verschiedenen Korrelationsmaschinen ermöglicht.

1 Introduction

With the emerging establishment of event-driven solutions, *complex event processing* (CEP) has become increasingly important in the context of a wide range of business applications such as supply chain management, manufacturing, or ensuring safety and security. CEP allows applications to asynchronously react to the changing conditions of possibly many business items by describing relevant business situations as correlation over many simple and complex events. Each event corresponds either to a change of a business item or the occurrence of a relevant business situation.

CEP provides high flexibility in writing and reconfiguring business applications, e.g. by decoupling low level information related to technical content and high level information related to business content. Consider an organization that wants to monitor access to a restricted area. In this context, sensor readings that provide position information about movements within the area represent technical data, while the business application works on process-relevant events like the entry of a person into a dangerous area. Though substituting the sensor technology might require new correlation rules to detect such a situation, applications that work on this situation will not require any change at all.

Most business applications require CEP systems to satisfy a couple of core requirements: i) the language used by a CEP system needs to be expressive enough to describe all relevant situations; ii) the processing has to be fast and effective to cope with real-time requirements of many applications and possibly high event rates.

While initial research in correlation systems (cf. [1, 3]) lack expressiveness (e.g. by allowing only the correlation of event types), there has been a development towards highly expressive and efficient correlation technology [4] satisfying both requirements.

However, the design of these correlation systems has been inherently centralized, thus limiting scalability as the number of event sources increases in the system. The dramatic increase in the complexity of business situations and the increase of devices that contribute to them (e.g. millions or billions of sensor devices available) seems to confirm that indeed more and more event sources need to be accounted for in the future. This growth is expected to be amplified by the industrial need of interoperability. Business partners connect their systems in order to optimize and manage their collaboration.

Recent research efforts towards distributed correlation systems (cf. [7, 5, 6, 10, 12, 13]) improve scalability by distributing correlation rules over multiple servers. However, available systems are far less efficient and expressive than their centralized counterparts. Furthermore, these systems typically consider homogeneous systems, which means that in order to detect a situation, all nodes are required to use the very same language in describing the rules as well as the same engine used to detect the situation. Interoperability between heterogeneous applications is only possible, if the same technology is used. This, however, forms an obstacle if business situations involve multiple domains like a logistic chain does.

In this article we address the need to cope with *heterogeneity* in distributed event correlation systems in order to i) reuse expressive correlation and efficient technology optimized for processing speed, ii) increase scalability by distributing correlation tasks over various correlation engines, and iii) allow migration of correlation tasks between heterogeneous engines and domains to increase interoperability and availability of correlation results. In particular, we present a framework that copes with such requirements and is currently developed within the DHEP project, a collaboration between IBM and the Institute of Parallel and Distributed Systems (IPVS) of the Universität Stuttgart.

2 Dealing with Heterogeneity

So far there is a large gap between event correlation systems adverted in academia and the ones used in business contexts. We argue that dealing with heterogeneity in distributed event correlation is crucial to close this gap and identify research challenges in coping with heterogeneity.

2.1 The Gap in Supporting Business Requirements

Recent research efforts have addressed the distributed organization of the event correlation process in order to ensure properties like scalability, efficiency and reliability. A distributed setting allows improvements on a wide range of performance characteristics, such as balancing the load by distributing correlation tasks over several nodes, reducing the overall network usage by appropriate placement of correlation tasks (possibly close to the data sources), or increasing the availability by performing correlation tasks at multiple nodes concurrently. Hence, one fundamental requirement is the ability to dynamically migrate rules between different network nodes, for instance if new application requirements are introduced, the network characteristics have changed, new consumers of events are introduced into the system, or failures occur in the system (cf. [5, 10]).

Although distributed event processing has recently evolved to a major research topic, it has hardly been deployed in business applications because it leaves a gap in supporting requirements imposed by business applications.

A characteristic of current approaches is that they consider homogeneous engines for their migration techniques (cf. [10, 12, 13]), whereas in industry various correlation engines are deployed to perform different tasks. Therefore, migration techniques require universally comprehensible rules and the collaboration of heterogeneous engines, which are possibly situated in different domains and communicate via different network protocols. Also, in order to ease the migration process, current approaches try to keep the resource utilization on the nodes low, which typically leads to less expressive languages, while business applications need high expressiveness.

Moreover, adaptation to node resource restrictions has been dealt with by splitting and distributing queries in research, while a promising approach is missing that adapts to node resources not just by adapting the complexity of the query but by the deployment and configuration of correlation engines itself.

Finally, distribution techniques do not seem to be mature enough to be deployed in business applications. They take into account only static node resources, but ignore important aspects for business application, such as dynamic load, additional restrictions for the processing of rules, security aspects and the proximity to context data sources that are necessary for rule processing.

Without these issues being addressed, distributed event processing is unlikely to be applied in business applications.

2.2 Closing the Gap

Closing the gap between academic distributed solutions and the needs of business applications will most likely change the way correlation technology is deployed.

A first step to maintain the power of mature correlation technology is to integrate it into a distributed correlation network without modifying it itself [7]. In addition to connecting input and output of individual powerful correlation nodes and deciding which correlation tasks to perform at which node, the configuration and deployment of each correlation machine becomes a main issue.

As an essential part of this process, the distribution of correlation rules may depend very much on the individual characteristics of the correlation machine. For example, a machine with scarce resources may not be capable of performing a costly correlation task but it can nonetheless be charged with low level correlation tasks to increase the capacity and efficiency of the correlation system as a whole when high-performance nodes reach their capacity bounds. Additional to node resources, in heterogeneous settings, distributed event processing has to take into account available network resources and the desired degree of reliability of rule processing.

In the context of business applications, it can be observed that several levels of data processing are involved. Figure 1 depicts this on the basis of the logistic chain example. Here, the arrows highlight the event flow, with their size indicating the events' complexity, and two levels of data processing represent different complexities of event correlation. At the technical level, processing of mainly simple, but frequent data correlations need to be performed and at the application level, correlation tends to become increasingly complex. The use of identical technology on both levels may lead either to a large overhead at the technical level, thus reducing the throughput, or too little expressiveness at the application level. For example, a powerful but resource consuming correlation engine is intended to process complex rules at the application level, but it is likely to be an overkill at the technical level. Therefore, using heterogeneous technology at different levels can increase the system's overall efficiency.

Finally, there is a trend to spread business processes over multiple domains. For instance, in a global logistic chain several logistic partners have to cooperate (e.g. a parcel service, an airport or a harbor). As a result, interoperability between different interacting domains becomes important in many modern applications. The dashed arrows in Figure 1 indicate such a inter-domain event flow. Heterogeneous event correlation can meet this challenge by enabling efficient communication between the various correlation engines of the participating business parties.

2.3 Current State

So far, work on heterogeneity in complex event processing has mainly addressed interoperability between either heterogeneous nodes (cf. [6, 12]) or networks (cf. [8, 11, 9]). Typically, the former approaches assume homogeneous event correlation technology while the latter do not allow rule migration due to language disparity. A combination of both has not been tackled so far. For instance, the authors in [8] describe a system that allows communication across different domains and networks, each handling its own correlation system. However, the configuration is left to the owner of each domain and rule migration is not possible.

A first step towards exploiting heterogeneous correlation technology to increase the system's overall performance is presented by the authors of [2]. They combine different event processing engines to increase the system's overall performance, namely a classical correlation engine at the application level as well as a fast stream processing engine at the technical level. However, the system has to be set up manually, and the topic is not investigated further.

Ontology-based infrastructures have received more attention recently. For example, CREAM [9] is a middleware based on ontologies that supports event handling from heterogeneous data in heterogeneous environments. It includes a layered rule processing structure that is capable of transferring rules defined by the user in a basic common language into various internal rule representations, dependent on the used optimization criteria. However, adaption by migrating rules is also not dealt with in this approach.

2.4 Challenges

We classify the challenges of *Heterogeneous Event Correlation* in three different levels:

the interoperability level, the communication level and the node level.

To enable *interoperability* between heterogeneous systems, language translation is inevitable. To support direct translation from one correlation language to another, however, introduces a large overhead since every correlation node needs to be aware of all possible counterparts. Furthermore, the translation from one correlation syntax to another can become fairly complex since translation was not the scope in their design.

Interoperability becomes even more complex when security issues have to be considered. The potential of connecting different correlation systems causes a problem with interfering interests of participating domains. On the one hand, a cooperation between the systems is desired, on the other hand, the domain owners do not want to give more information than needed to their partners, especially since correlation rules may contain confidential business process information. Therefore it is important to be able to regulate which information is private and thus has to be processed locally, and which information is allowed to be processed on other domains.

On the *communication* level, we have to consider data proximity. Business event processing often uses context information that has to be applied to the events. The location of each event's context information should be considered when rules have to be deployed and processed.

However, the major research challenge on the placement of rules is coming along with heterogeneous *nodes*. Even in a system that supports interoperability, we still may not be able to process all rules at their optimal position. The reasons for that are twofold. First, a node may not be capable of processing some rules due to engine or node restrictions. Second, nodes may already be fully loaded and are not capable of processing any other rule. In both cases, another, maybe suboptimal placement has to be found.

The availability and reliability of heterogeneous systems are also affected by nodes. Their characteristics, like location, domain or network connectivity may have a major impact on whether a node is used to ensure availability or not. Mechanisms to achieve a reliable event correlation might also have to be investigated within the presence of heterogeneous systems.

3 DHEP

In the DHEP (Distributed Heterogeneous Event Processing) project [15], we are addressing the challenges introduced in Section 2.4 by providing event correlation over heterogeneous environments. In the following we give an overview of the main concepts (cf. Figure 2) of our approach.

3.1 Correlation Language

Supporting interoperability of heterogeneous systems and migration of correlation rules between them requires some translation mechanism. Since direct translation between all used languages becomes far too expensive, we define a meta-language to which all correlation rules that need to be migrated are transferred. The defined meta language is expressive enough to support even sophisticated correlation languages like AMiT (IBM Active Middleware Technology^{TM}) [4]. Hence, our system can easily be extended to support a wide range of correlation technology.

In addition, the meta-language supports the formulation of restrictions as well as the definition of sophisticated models representing context information that might be used during the event processing, e.g. to enrich events with additional data.

3.2 Core Components

DHEP provides a framework that enables heterogeneous event correlation by supporting communication between different correlation engines across the network. Therefore, the engines are not part of the framework itself but located on top of the architecture. The framework, as it is depicted in Figure 2, is deployed on each participating correlation node within the network. Hence, we are not only able to interoperate between heterogeneous engines across the network, but also with heterogeneous engines on one single node.

The rule management component handles rules that need to be distributed and deployed. Basically, the component stores the currently deployed rules and their states. This allows migrations without having to interact with the internal state of the underlying engine. This also enables DHEP to integrate several correlation engines on the same node. Rule management has several sub-components attached to it: i) the translation of correlation rules is performed inside a wrapper component which acts as an interface between the framework and the engine; ii) the configuration component uses knowledge about the network and the defined rule set to find nodes where rules should be deployed (cf. Section 3.3).

Event and rule dissemination is handled by the communication component. Here, it is important to support nodes and networks using heterogeneous communication paradigms, since business applications typically use a wide diversity of communication such as queuing mechanisms, direct synchronous communication or an enterprise service bus. Thus, we provide a generic interface offering multiple communication paradigms. Initially we provide three ways of interaction in our network: classical direct communication via sockets, reliable communication via message queues and a topic based publish/subscribe communication using the Java Message Service™(JMS) [14]. The communication component is the basis of the DHEP framework since every task, either functional or organizational, has to be processed by it. It enables the distribution of the correlation process and therefore the scalability of the event processing system.

3.3 Network Configuration

The flexibility in placing correlation rules raises also the need for good placement strategies that allow to detect situations efficiently as well as to ensure the availability of correlation results in the presence of failures.

The placement of correlation rules heavily influences the network's efficiency. Depending on the desired system behavior, placement algorithms typically try to optimize specific criteria like end-to-end latency, load and rule reusability.

However, when heterogeneity comes into play, the placement problem becomes even more complex. As described in Section 2.4, additional factors like node attributes or available engines may have an impact on the placement of correlation rules. In our project we will face this challenge by providing rule placement strategies and network configuration mechanisms that consider these effects, e.g. by using network and node monitoring.

4 Conclusions & Future Work

In this article, we have motivated the need for heterogeneous event processing so that business applications can benefit from distributed event correlation technology. We have presented the DHEP framework which addresses the challenges introduced by heterogeneity. The framework is currently being implemented and tested in the context of IBM business applications. We are investigating methods to optimize placement strategies in the heterogeneous context and how reliable event processing can be achieved in such a context. Additionally, we will examine further the inter-domain event processing problem.

Acknowledgment

The DHEP project is funded by IBM Deutschland as part of the IBM Technology Partnership Center (ITPC). We thank Marco Völz, Gerald Koch and Boris Koldehofe for their constructive feedback during the development process of this work.

Literatur

[1] S. Chakravarthy and D. Mishra: Snoop: an expressive event specification language for active databases. In *Data Knowl. Eng.*. Elsevier Science, 1994.

[2] S. Chakravarthy and R. Adaikkalavan: Events and streams: harnessing and unleashing their synergy!. In *Proc. of the second international conference on Distributed event-based systems (DEBS '08)*. ACM, 2008.

[3] S. Gatziu and K.R. Dittrich: SAMOS: an Active Object-Oriented Database System. In *Data Engineering, 1992, Vol. 15, pp. 23-26*. IEEE, 1992.

[4] A. Adi and O. Etzion: Amit - the situation manager. In *The VLDB Journal 13, 2 (May. 2004), pp. 177-203*. Springer, 2004.

[5] P.R. Pietzuch et al.: A Framework for Event Composition in Distributed Systems. In *Proc. of the 4th Int. Conf. on Middleware (MW'03), pp. 62-82*. Springer, 2003.

[6] G. Jakobson: The Technology and Practice of Integrated Multi-Agent Event Correlation Systems. In *Proc. of the Int. Conf. on Integration of Knowledge Intensive Multi-Agent Systems, pp. 568-573*. IEEE, 2003.

[7] A. Biger et al.: Stratified implementation of event processing network. *Fast Abstract on the Int. Conf. on Distributed event-based systems (DEBS '08)*, 2008.

[8] P. Arozarena et al.: Madeira: A peer-to-peer approach to network management. In *Wireless World Research Forum, Vol. 16*, 2006.

[9] M. Cilia et al.: CREAM: An Infrastructure for Distributed Heterogeneous Event-Based Applications. In *Proc. of the confederated Int. Conf. CoopIS/DOA/ODBASE, pp. 482-502*. Springer, 2003.

[10] W. Tai et al.: Distributed fault correlation scheme using a semantic publish/subscribe system. In *Proc. of Network Operations and Management Symposium (NOMS 2008), pp. 835-838*. IEEE, 2008.

[11] E. Yoneki: Event Broker Grids with Filtering, Aggregation, and Correlation for Wireless Sensor Data. In *Proc. of the OTM Workshops '05, pp. 304-313*. Springer, 2005.

[12] Y. Liu: The architecture of an event correlation service for adaptive middleware-based applications. In *J. Syst. Softw., Vol. 81, pp. 2134-2145*. Elsevier Science, 2008.

[13] G. Li and H.-A. Jacobsen: Composite Subscriptions in Content-Based Publish/Subscribe Systems. In *Proc. of Middleware 2005, pp. 251-271*. Springer, 2005.

[14] M. Hapner et al.: Java Message Service. Sun Microsystems Inc., 2002.

[15] Distributed Heterogeneous Event Processing: http://itpc.informatik.uni-stuttgart.de/dhep.html

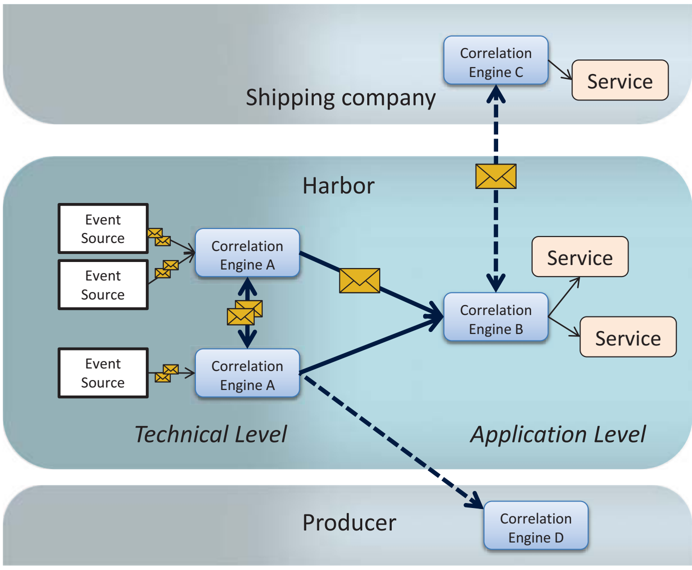

Abbildung 1: Event processing in a logistic chain

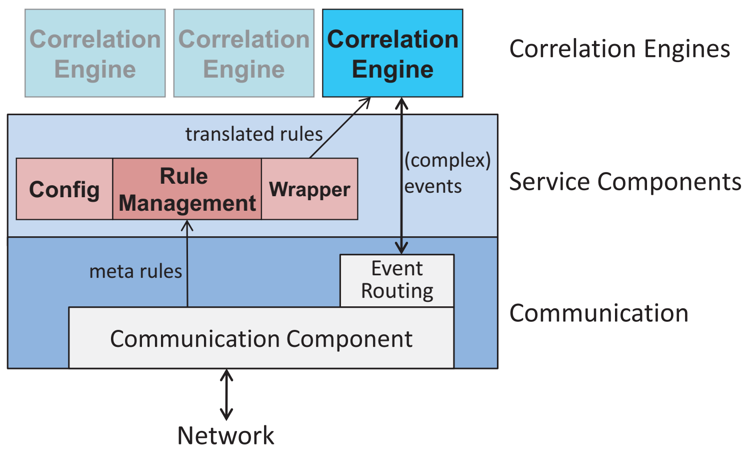

Abbildung 2: The DHEP Framework