

**07191 Abstracts Collection**  
**Event Processing**  
— **Dagstuhl Seminar** —

Opher Etzion<sup>1</sup>, Mani Chandy<sup>2</sup> and Rainer von Ammon,<sup>3</sup>

<sup>1</sup> IBM Haifa, IL

opher@il.ibm.com

<sup>2</sup> Cal Tech Pasadena US

kmchandy@gmail.com

<sup>3</sup> CITT & Univ. of Regensburg, DE

**Abstract.** From 06.05. to 11.05.2007, the Dagstuhl Seminar 07191 “Event Processing” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

**Keywords.** Event Processing, Real-time Information Systems, Reactive systems, Proactive systems, Active Technologies

## 07191 Summary – Event Processing

This is the summary of the Dagstuhl Seminar on Event Processing.

*Keywords:* Event Processing

*Joint work of:* Chandy, Mani; Etzion, Opher; von Ammon, Rainer; Niblett, Peter

*Extended Abstract:* <http://drops.dagstuhl.de/opus/volltexte/2007/1148>

## Consistent Streaming Through Time: A Vision for Event Stream Processing

*Roger Barga (Microsoft Research - Redmond, USA)*

Event processing will play an increasingly important role in constructing enterprise applications that can immediately react to business critical events.

Various technologies have been proposed in recent years, such as event processing, data streams and asynchronous messaging (e.g. pub/sub). We believe these technologies share a common processing model and differ only in target workload, including query language features and consistency requirements. We argue that integrating these technologies is the next step in a natural progression. In this paper, we present an overview and discuss the foundations of CEDR, an event streaming system that embraces a temporal stream model to unify and further enrich query language features, handle imperfections in event delivery, define correctness guarantees, and define operator semantics. We describe specific contributions made so far and outline next steps in developing the CEDR system.

*Keywords:* Event stream processing; temporal stream model; model imperfections; correctness guarantees; operator semantics

*Joint work of:* Barga, Roger; Jonathan Goldstein; Mohamed Ali; Mingsheng Hong

*Full Paper:*

<http://www-db.cs.wisc.edu/cidr/cidr2007/index.html>

*See also:* In the Proceedings of CIDR '07, Conference on Innovative Database Research

## Software Engineering Issues

*Mikael Berndtsson (University of Skövde, S)*

Tools and methodologies for designing ECA applications have been requested since the early 1990s. With the recent big interest in Event Processing, similar requests for tools and methodologies have resurfaced.

The talk presents various approaches to providing software engineering guidelines. The presented approaches have appeared within the active database community. Thus, they cannot directly be carried across to the event processing community. Instead they need to be adapted and complemented with approaches from other (sub)areas that are part of the Event Processing community.

A first step, is to collect guidelines for when to use events/rules and when not to use events/rules. Today, these guidelines are scattered in research articles and in teaching material.

*Keywords:* ECA rules, active databases, software engineering, tools, methodologies

## BiCEP - Benchmarking Complex Event Processing Systems

*Pedro Bizarro (University of Coimbra, P)*

BiCEP is a new project being started at the University of Coimbra to benchmark Complex Event Processing systems (CEP). Although BiCEP is still in the early stages, we list here some of the design considerations that will drive our future work and some of the metrics we plan to include in the benchmark.

*Keywords:* Complex Event Processing, benchmark, synthetic benchmark, events, response time, throughput, scalability, adaptivity, query processing

*Extended Abstract:* <http://drops.dagstuhl.de/opus/volltexte/2007/1143>

## Twelve Theses on Reactive Rules for the Web

*François Bry (Universität München, D)*

Reactivity, the ability to detect and react to events, is an essential functionality in many information systems. In particular, Web systems such as online marketplaces, adaptive (e.g., recommender) systems, and Web services, react to events such as Web page updates or data posted to a server.

This article investigates issues of relevance in designing high-level programming languages dedicated to reactivity on the Web. It presents twelve theses on features desirable for a language of reactive rules tuned to programming Web and Semantic Web applications.

*Keywords:* Event-Condition-Action Rules, Web

*Joint work of:* Bry, François; Eckert, Michael

*Extended Abstract:* <http://drops.dagstuhl.de/opus/volltexte/2007/1144>

*Full Paper:*

[http://www.pms.ifi.lmu.de/publikationen/PMS-FB/PMS-FB-2006-4/reactive\\_theses.pdf](http://www.pms.ifi.lmu.de/publikationen/PMS-FB/PMS-FB-2006-4/reactive_theses.pdf)

*See also:* In: Proceedings of Workshop

## Infrastructure for Smart Cities: The Killer Application for Event-Based Computing

*Alejandro P. Buchmann (TU Darmstadt, D)*

Infrastructures for smart cities are considered a potential killer app for event-based computing. Event services are a crucial part of the infrastructure.

The complexity of the event services is compounded by the richness of the events, the number of (mobile) sensors and devices, heterogeneity, requirements for seamless integration, unstable communication and interference, quality of service requirements, the need for context awareness and device orchestration and self-X properties.

*Keywords:* Events, event-based computing

*Extended Abstract:* <http://drops.dagstuhl.de/opus/volltexte/2007/1145>

## Ubiquitous Nature of Event-Driven Approaches: A Retrospective View

*Sharma Chakravarthy (University of Texas at Arlington, USA)*

This paper retrospectively analyzes the progress of event-based capability and their applicability in various domains. Although research on event-based approaches started in a humble manner with the intention of introducing triggers in database management systems for monitoring application state and to automate applications by reducing/eliminating user intervention, currently it has become a force to reckon with as it finds use in many diverse domains. This is primarily due to the fact that a large number of real-world applications are indeed event-driven and hence the paradigm is apposite.

In this paper, we briefly overview the development of the ECA (or event-condition-action) paradigm. We briefly discuss the evolution of the ECA paradigm (or active capability) in relational and Object-oriented systems. We then describe several diverse applications where the ECA paradigm has been used effectively. The applications range from customized monitoring of web pages to specification and enforcement of access control policies using RBAC (role-based access control). The multitude of applications clearly demonstrate the ubiquitous nature of event-based approaches to problems that were not envisioned as the ones where the active capability would be applicable.

Finally, we indicate some future trends that can benefit from the ECA paradigm.

*Keywords:* Complex event processing, stream processing, semantics, event-driven applications

*Joint work of:* Chakravarthy, Sharma; Adaikkalavan, Raman

*Full Paper:* <http://drops.dagstuhl.de/opus/volltexte/2007/1150>

## What is event processing?

*Mani Chandy (CalTech - Pasadena, USA)*

Discussion of what is and is not event processing. This is a strawman proposal for the workshop.

*Keywords:* Event processing, models, abstractions, disciplines

## Business Performance Observation Model

*Henry Chang (IBM TJ Watson Research Center, USA)*

Business in a globally integrated enterprise are hard pressed to make strategic and tactical decisions based on latest business performance data that are continuously generated from scattered sources and organizations. Because of the silo and un-integrated operational data, there are big gaps between what is needed and what is available for near real-time operational performance observation. The business event processing is a new and structured paradigm shift for handling business operational data to build an event-driven performance observation platform for business executives separate from the traditional approach of batch-oriented data replication and integration.

The business event processing approach provides flexibility, near real-time performance, event-to-metrics aggregation and colleration, business situation detection, and businesss decision support.

The observation model governs how events are collected, information extracted, and data visibility supported. The scalability challenges are met by carefully analyzing the declarative models and transform into exeuction plan of event processing that utilizes the latest technology in the event stream world.

*Keywords:* Business event processing, active metrics network, business performance management, business observation

*Full Paper:*

<http://www.research.ibm.com/journal/sj/453/chowdhary.html>

*See also:* Dr. Henry Chang (hychang@us.ibm.com) is a senior technical staff member and a research manager in Business Informatics department at the IBM T.J. Watson Research Center. He leads the research effort in business performance monitoring and management framework with technical impacts to IBM Websphere BPM suits and IBM internal supply chain visibility initiatives. His recent research interests include business event processing optimization, continuous process improvement, and event-based business collaboration. He received an IBM Innovate Award for his work on B2B collaboration solutions. Before joining IBM at the Thomas J. Watson Research Center, he received Ph.D. and MS. in Computer Sciences from U. Wisconsin-Madison at 1987 and a B.S. in Electrical Engineering from National Taiwan University in 1979, respectively. He is a long time member of ACM and IEEE.

## Model Driven Development for Business Performance Management

*Henry Chang (IBM TJ Watson Research Center, USA)*

Business process integration and monitoring provides an invaluable means for an enterprise to adapt to changing conditions.

However, developing such applications using traditional methods is challenging because of the intrinsic complexity of integrating large-scale business processes and existing applications. Model Driven Development (MDD) is an approach to developing applications from domain-specific models to platform-sensitive models that bridges the gap between business processes and information technology. We describe the MDD framework and methodology used to create the IBM Business Performance Management (BPM) solution. We describe how we apply model-driven techniques to BPM and present a scenario from a pilot project in which these techniques were applied. Technical details on models and transformation are presented. Our framework uses and extends the IBM business observation metamodel and introduces a data warehouse metamodel and other platform-specific and transformational models. We discuss our lessons learned and present the general guidelines for using MDD to develop enterprise-scale applications.

*Keywords:* Model driven development, business performance management

*Joint work of:* P. Chowdhary; Henry; K. Bhaskaran; etc.

*Full Paper:*

<http://www.research.ibm.com/journal/sj/453/chowdhary.html>

## **Evaluation and Optimization of Composite Event Queries - A Relational Algebra Perspective**

*Michael Eckert (Universität München, D)*

Queries for composite events typically involve the four complementary dimensions of event data, event composition, relationships between events (esp. temporal and causal), and accumulating events over time windows for negation and aggregation. We consider a datalog-like rule language for expressing such composite event queries and show that their evaluation can be understood as a problem of incrementally evaluating relational algebra expressions. We then show how temporal relationships between events can be utilized to make the evaluation of joins more efficient by avoiding evaluation of certain subexpressions and by making storage of some intermediate results unnecessary.

*Keywords:* Composite event queries, incremental evaluation, optimization

*Joint work of:* Eckert, Michael; Bry, Francois

## **Semantics of event processing systems - introduction**

*Opher Etzion (IBM - Haifa, IL)*

This presentation provides some thoughts about semantics of event processing systems.

*Keywords:* Event processing, semantics

## On the Efficient Monitoring of Events

*Avigdor Gal (Technion - Haifa, IL)*

We shall present a dual formalism for the efficient monitoring of events, given constraints, set by both event generators (servers) and event consumers (clients). We shall then briefly present three specific monitoring problems and their efficient solutions:

1. Monitoring under politeness constraints.
2. Satisfying User Profiles.
3. The proxy dilemma.

*Keywords:* Event monitoring, online data delivery, pull solutions

*Joint work of:* Roitman Haggai; Gal, Avigdor; Louiqa Raschid; J. Eckstein; S. Reiner

*Full Paper:*

<http://iew3.technion.ac.il/~avigal/dqd.pdf>

*See also:* J. Eckstein, A. Gal, and S. Reiner. Monitoring an information source under a politeness constraint. *INFORMS Journal on Computing*, 2007. forthcoming.

## A Model for Reasoning with Uncertain Rules in Event Composition Systems

*Avigdor Gal (Technion - Haifa, IL)*

In recent years, there has been an increased need for the use of active systems - systems required to act automatically based on events, or changes in the environment. Such systems span many areas, from active databases, through applications that drive the core business processes of today's enterprises. However, in many cases, the events to which the system must respond are not generated by monitoring tools, but must be inferred from other events based on complex temporal predicates. In addition, in many practical applications, such inference is inherently uncertain. In this paper, we introduce a formal framework for knowledge representation and reasoning enabling such event inference. Based on probability theory, we define the representation of the associated uncertainty. In addition, we formally define the probability space, and show how the relevant probabilities can be calculated by dynamically constructing a Bayesian network. To the best of our knowledge, this is the first work that enables taking such uncertainty into account in the context of active systems. Therefore, our contribution is twofold: We formally define a probabilistic representational model for event composition, and show how to apply this model to the quantification of the occurrence probability of events. This results in a framework enabling any active system to handle such uncertainty.

*Keywords:* Composite events, uncertainty

*Joint work of:* Wasserkrug, Segev ; Gal, Avigdor; Etzion, Opher

*Full Paper:*

<http://iew3.technion.ac.il/~avigal/UncertainEventsUAI.pdf>

*See also:* Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI-2005). Edinburgh, Scotland, July 26-29, 2005

## Event Technologies in the Oracle Database

*Dieter Gawlick (Oracle Corp. - Redwood Shores, USA)*

This presentation was extended and refined for a tutorial at Sigmod 2007 with the title Event Processing Using Database Technology. Slides 78 and later of the Sigmod presentation map to the Dagstuhl presentation. When reading the Sigmod presentation please use the presentation mode - some slide builds are important. I like to point especially to slides 92 and 155. Slide 92 (How Are Events and Messages Created?) discusses the controversy about the term event. I think it explains why there so divers views and why the view message = event has only limited application. Slide 155 (Challenges for Development Teams) intends to remind the research community that very innovative work is done by development teams; that researchers should become more aware of it and avoid competing with development teams as much as possible. A significant part of the contributions from the research community at the Dagstuhl seminar did not take this warning into account.

*Joint work of:* Gawlick, Dieter; Chandy, Mani; Mishra, Shailendra

## An algebra for event composition

*Annika Hinze (University of Waikato, NZ)*

Event-based systems are used in a wide variety of applications such as digital libraries, stock tickers, traffic control, or facility management. Composite events have been introduced to capture richer situations. Composite events in different systems seem to follow a common semantics. However, the evaluation semantics of the considered events may differ heavily according to the approach chosen.

We propose a parameterized event algebra that describes the detailed semantics of composite events. We define the event operators that form composite events and we introduce parameters for event instance selection and event instance consumption. These parameters serve as a support for handling duplicates in both primitive and composite events.

The event algebra provides a foundation for

- a (meta) language for events



- adapting to changing systems/settings
- meta-system (similar meta-search)

Our approach flexibly supports the integration of event-based applications and changing event sources without the need to redefine user profiles.

*Keywords:* Events, composition, semantics, adaptation, language

*Joint work of:* Hinze, Annika; Voisard, Agnes

## Sensor Event Processing on Grid

*Eui-Nam Huh (Kyung Hee University, ROK)*

Wireless sensor networks are increasingly being deployed in many important applications. For sharing huge amount of sensor data efficiently with diverse users, an information dissemination mechanism is very necessary and important component. In this paper, we have proposed an efficient architecture integrated with sensor network and Grid technology. To disseminate the sensed data to users geographically distributed, an experimental method using Data Grid on pub/sub (publish/subscription) is designed for a u-Healthcare application and its performance is evaluated for various predicate cases.

*Keywords:* Event processing, sensor networks, Grid, Dissemination

*Extended Abstract:* <http://drops.dagstuhl.de/opus/volltexte/2007/1146>

## The OASIS WS-Topics standard

*Peter Niblett (IBM - Winchester, GB)*

WS-Topics defines a mechanism to organize and categorize items of interest for subscription known as "topics". These are used in conjunction with the notification mechanisms defined in WS-BaseNotification. WS-Topics defines three topic expression dialects that can be used as subscription expressions in subscribe request messages and other parts of the WS-Notification system. It further specifies an XML model for describing metadata associated with topics.

These slides give a brief introduction to the WS-Topics specification.

*Keywords:* Publish/Subscribe Topics Web Services

## **View from the field on SOA, EDA, BAM, BSM**

*Marc Peters (IBM Deutschland - Köln, D)*

View from the field on SOA, EDA, BAM, BSM including some thoughts about the comparison between SOA and EDA including similarities and differences and the positioning of active dashboards vs. traditional dashboards in BAM with some remarks on BSM and the rising importance and demand for BSM and BAM in the market. Conclusion on possible good entry points.

*Keywords:* SOA, EDA, BAM, BSM

## **Denial of Information Attacks in Event Processing**

*Calton Pu (Georgia Institute of Technology, USA)*

It is a common assumption in event processing that the events are "clean", i.e., they come from well-behaved and trustworthy sources. Some researchers have studied uncertain event streams [Mok et al, RTCSAŠ06], but few have considered malicious event sources. In the real world, event sources from open environments (e.g., large scale sensor networks and Internet) can be influenced by adversaries injecting misleading or noise data. This has happened to all media that have become valuable in open environments. Spammers have been active with email spam, web spam, blog spam, spam over VoIP, and fake profiles in social networks. We call this automated injection of false or noise fabricated events "Denial of Information" (DOI) attacks. The automated nature of DOI attacks makes it inexpensive and easy to implement.

DOI attacks introduce some fundamental research challenges. For example, consider a set of audio sensors for detection of activities through sound. If half of the sensors report sound and the other half are silent, it is difficult to decide whether the silent one are reporting real phenomena (with the sound purposefully produced by the adversary) or the noisy ones are reporting real phenomena (assuming the silent ones have been incapacitated). This problem is often called Deceptive Information Detection. Furthermore, DOI attacks are different from typical information security problems (e.g., multi-level security) is the "arms race" between DOI attacks and defenses. This can be illustrated with the co-evolution of spam messages and automated email filters employed by spam victims. The evolution of email spam (with randomized camouflaged content and image spam) is a good example. This problem is usually called Adversarial Learning.

Robust event processing of the future must be able to tolerate and resist DOI attacks, by introducing techniques and tools that can counter DOI attacks in areas such as Deceptive Information Detection and Adversarial Learning. We discuss some progress made in specific applications such as email to illustrate the difficulty of these challenges and some promising approaches.

*Keywords:* Denial of information, information security, adversarial learning, deceptive information detection

*Full Paper:*

<http://www-static.cc.gatech.edu/projects/doi/>

*Extended Abstract:* <http://drops.dagstuhl.de/opus/volltexte/2007/1147>

## Customizable Parallell Scientific Stream Processing

*Tore Risch (Uppsala University, S)*

The talk presents work done at Uppsala University on processing high volume scientific events from the application area of space physics. The events are produced by distributed digital receivers of signals from space. The event streams are collected in a central processing facility where computations are made based on combining events from one or several receivers. The often expensive computations are specified in continuous queries. Clusters must be used for achieving acceptable throughput. For performance it is particularly important that users can specify parallelizations of computations and such parallelization often depends on what kind of computations are made in the continuous queries. To enable customized parallelizaion, our systems provide user extensions for parallel processing of high volume continuous queries. The talk presents the user definable parallel event stream processing facilities shown to improve stream throughput significantly on a Linux cluster. Furthermore, the architecture is presented of our new parallel stream processor, SuperComputer Stream Query Processor, SCCQ, for heterogeneous computing environments and applied on our application environment consisting of Linux clusters, a BlueGene supercomputer, and other computers. The talk is based on publications in the VLDB conference and an ICDE workshop.

*Keywords:* Scientic event streams, high volument event streams, parallell event processing

*Joint work of:* Risch, Tore; Ivanova, Milena ; Zeitler, Erik

## Event Recognition and Processing with Embedded Devices

*Kirsten Terfloth (FU Berlin, D)*

With the emergence of wireless sensor networks, the challenges of event recognition and processing have been partially shifted into the embedded domain.

While new processing capabilities on small devices allow for physically close event monitoring and fast filtering, new challenges due to sparse resources or medium contention when relying on wireless communication arise.

Within this talk, a short introduction into an experimental setup featuring event-detection on a construction fence is given. The goal has been to distinguish a person climbing over the fence from other events, which it may be exposed to, with the help of a wireless sensor network. The results obtained will be presented thoroughly. Regarding those, we discuss in which situations distributed event recognition and processing is to be preferred over a conventional server-centered deployment. Therefore, the costs such as communication, hardware and deployment related costs implying an architectural decision are examined.

*Keywords:* Event processing, embedded devices

*Joint work of:* Terfloth, Kirsten; Hahn, Katharina; Voisard, Agnès

*Full Paper:* <http://drops.dagstuhl.de/opus/volltexte/2007/1141>

## **An XML Framework for Integrating Continuous Queries, Composite Event Detection, and Database Condition Monitoring for Multiple Data Streams**

*Susan Urban (ASU - Tempe, USA)*

With advancements in technology over the last ten years, data management issues have evolved from a stored persistent form to also include streaming data generated from sensors and other software monitoring tools.

Furthermore, distributed, event-based systems are becoming more prevalent, with a need to develop applications that can dynamically respond to information extracted from data streams. This research is investigating the integration of stream processing and event processing techniques, with expressive filtering capabilities that include queries over persistent databases to provide application context to the filtering process. Distributed Event Processing Agents (DEPAs) continuously filter events from multiple data streams of different formats that provide XML views. Composite events for data streams are expressed using CXQ, a language that extends XQuery with temporal, composite event language features, including operators for expressing sequence, disjunction, conjunction, repetition, aggregation, and time windows for events. Continuous queries and composite event filters are integrated with techniques for materialized view maintenance and incremental evaluation in condition monitoring to provide efficient ways of enhancing stream filters with database queries. The filtering and event detection load is distributed among multiple DEPAs, with CXQ expressions decomposed to allocate subcomponents of the expression to DEPAs that efficiently communicate in the global detection of composite events. A unique aspect of our research is that it extends XQuery with temporal, composite event features to

combine techniques for continuous queries in stream processing, incremental evaluation in condition monitoring, and detection and filtering of composite events, creating an expressive environment for the extraction of meaningful events from multiple data streams with XML views.

Our approach to the integration of stream processing and event processing is essential to the support of applications such as those in the medical field for health monitoring, the financial domain for executive dashboards, supply chain and other B2B applications for monitoring consumer activity, or for autonomic behavior within computer systems and embedded systems. With XML becoming a standard for data representation, our research will provide extensions to XQuery that will provide a uniform framework of time-based composite event detection to streams of data and events in different formats.

On a broader scale, our research will enhance the analysis of data and event streams through the use of database context filters and through the correlation of multiple streams, thus providing a way of extracting more meaningful, application-oriented events from streams of data in the support of dynamic, data-driven applications.

*Keywords:* XML, continuous queries, composite event detection, database context filters

*Joint work of:* Urban, Susan; Dietrich, Suzanne; Chen, Yi

*Extended Abstract:* <http://drops.dagstuhl.de/opus/volltexte/2007/1142>

## **Filtering Features for a Composite Event Definition Language**

*Susan Urban (ASU - Tempe, USA)*

This research has enhanced a distributed, rule-based application integration environment with a composite event definition language (CEDL) and detection system. CEDL builds on existing composite event operators and selection modes, adding features to support the filtering of primitive and composite events. The filtering features includes basic parameter filtering on primitive and composite events, aggregate and quantifier filters on cumulative event parameters, and time filters for defining the lifetime of the composite event detection process. CEDL is supported by a composite event detection system that implements the filtering capabilities. This research contributes to the expression of more application-oriented events through the aggregation and correlation of distributed events.

This paper appeared in IEEE International Conference on Software Applications on the Internet (SAINT) 2006, Phoenix, AZ, January 2006.

*Keywords:* Parameter filtering, composite events

## **A Composite Event Definition Language and Detection System for the Integration Rules Environment**

*Susan Urban (ASU - Tempe, USA)*

This is I. Biswas's MS thesis that describes the full details on the design and implementation of CEDL.

## **Complex Events in SQL—a critique of the newly proposed ANSI standards**

*Carlo Zaniolo (Univ. California - Los Angeles, USA)*

Powerful SQL:2003 extensions [1] for querying complex events were recently proposed by an ANSI working group in which DBMS vendors and data stream management systems (DSMS) startup companies joined forces. This presentation will provide an introduction and in-depth analysis of the proposed specs—which should be discussed at the workshop since they (i) will impact research directions and commercial systems, and (ii) raise issues on semantics, expressive power, and query optimization that should be discussed in this experts' forum. The proposed ANSI specs are largely based on SQL-TS [2], which proposed (1) a sequence pattern language for time series, based on regular expressions for describing composite events, and (2) query optimization techniques based on a generalization of the Knuth, Morris & Pratt string-search algorithm. The newly proposed ANSI specs [1] add several extensions—including constructs to support disjunctive patterns and others that entail finer control on the answers returned by queries. Through various examples, this presentation will illustrate the main constructs proposed in [2], their power in expressing complex event queries, and the semantic and optimization issues raised by these constructs.

[1] Zemke, Andrew Witkowski, Mitch Cherniak, Latha Colby:

Pattern matching in sequences of rows—ISO/IEC JTC1/SC3 ANSI NCITS H2-2006-nnn, Change Proposal, March 2, 2007.

[2] Reza Sadri, Carlo Zaniolo, Amir M. Zarkesh, Jafar Adibi:

Expressing and optimizing sequence queries in database systems, ACM Transactions on database systems, Vol 29, No. 2, June 2004, pages 282-318.

*Keywords:* Sequence queries, complex events, time series

*Joint work of:* Zaniolo, Carlo; Thakkar, Hetal—UCLA

*See also:* [1] Zemke, Andrew Witkowski, Mitch Cherniak, Latha Colby: Pattern matching in sequences of rows—ISO/IEC JTC1/SC3 ANSI NCITS H2-2006-nnn, Change Proposal, March 2, 2007. [2] Reza Sadri, Carlo Zaniolo, Amir M. Zarkesh, Jafar Adibi: Expressing and optimizing sequence queries in database systems, ACM Transactions on database systems, Vol 29, No. 2, June 2004, pages 282-318.