



This project is co-financed by the European Social Fund
and the state budget of the Czech Republic.



INVESTMENTS IN EDUCATION DEVELOPMENT

Editorial

This is the first edition of Poster Presentations 2013 from the Faculty of Informatics, at Masaryk University, Brno. The purpose of this booklet is to disseminate work done in this faculty. This booklet and the posters is an output of a poster design course at FI, supported by the Platform for Industrial Cooperation (CZ.1.07/2.4.00/17.0041).

The posters presented on the following pages are the work of a group of students and staff members. The students are enrolled for Masters and Doctoral degrees, and the staff members are involved with industry networks in the IT field.

Guidance was given through a series of lectures on poster layout and choosing content, using hundreds of poster examples from academic conferences all over the world. Proper selection of descriptive, scientific titles followed, and students started writing their own abstracts. The authors then made a mock-up by hand on full-sized sheets, to free up their creativity. Abstracts, titles and short CVs were checked and proofread.

The students then spent time on the final design and content of their posters. The wide variety of layouts, innovative approaches and high quality of content is a result of the dedication and hard work of the poster authors themselves. A competition was staged to evaluate the posters and reward the authors of the two best posters in terms of layout and content.

What you see in this book is an example of what can be achieved by a group of creative young people who are passionate about their fields of work and study. For most of them it is their first attempt at producing an academic poster, and as such is commendable.

We believe you will find this book stimulating, and that it will provide food for thought. Feel free to contact the authors if you need more information on their research or the progress of their work!

Sincere thanks go to my doctoral student Lucia Tokárová for doing an expert job on the graphic design and layout for this book.

Prof M Weideman

Visiting Erasmus Mundus Scholar

Cape Peninsula University of Technology, Cape Town, SA

weidemanm@cput.ac.za

February 2013

Contents

Complex Event Processing in Building Management Systems – a Prototype	6
Adam Kučera	
Design of a system for the analysis of social media content	8
Dalibor Toth, Jaroslav Ráček	
Developing the ICT partners network (project at FI MU)	10
Stanislava Sedláčková	
Distributed Event-driven Monitoring Model for Cloud Datacenters	12
Daniel Tovarňák, Tomáš Pitner	
Graphic design of a mobile UI for primary school administration	14
Lenka Plháková	
Implementation of the internship model in ICT area (project at FI MU)	16
Jana Bartáková	
Indoor User Localization Using Mobile Devices	18
Jonáš Ševčík	
Semantically Partitioned Peer to Peer Complex Event Processing Exploiting Information Loss	20
Filip Nguyen	
Supporting the Process of Learning Mobile Application User Interfaces	22
Lucia Tokárová	

Adam Kučera, Faculty of Informatics, Masaryk University

Introduction

Complex Event Processing (CEP) is set of tools and techniques allowing us to handle events from large systems in real time. CEP is commonly used in network security, algorithmic stock-trading, fraud detection or business process monitoring for correlating, aggregating and analyzing various events.

Modern ("intelligent") buildings are often equipped with BMS – network of interconnected devices that ensures integrated building operation, control and monitoring. Building operators and maintenance staff are easily overwhelmed by large amounts of data and events produced by the BMS. In this way BMS are similar to other fields where CEP is successfully used (network monitoring, credit card fraud detection, algorithmic stock trading etc.).

The goals


Aim of this research is to adapt existing CEP framework (Microsoft StreamInsight) to work with data from intelligent buildings and provide building operators with the easily understandable outputs that help to optimize building operation and detect various types of faults that occurs in a building. The CEP framework was extended to be able to work with specific features of building management data (location, source device, measured/controlled quantity). A prototype of the system was implemented and tested in real operation environment.

Test environment – BMS at Masaryk University


- Integrates monitoring and control of University Campus in Brno – Bohunice (first phase of construction ended in 2007) and other university buildings in Brno.
- Uses BACnet as common communication protocol
- Over 700 native BACnet devices by approximately 10 vendors
- Integrated systems/technologies: HVAC (Heating, Ventilation and Air-conditioning), Fire alarms, Security system, Access control system, CCTV, Power systems, UPS monitoring, university information system (controlling door locks according to lecture room timetables)

Selecting, filtering, joining, aggregating data

Using building management data semantics, we are able to perform selections, filtering and computations (average, summa, maximum, minimum) over the data from the BMS. The system uses time windows for performing continuous computations of the data characteristics. We can utilize hierarchy of passport databases and use masks to gather related data into separate groups.




System architecture overview



Building and technology passport

Passports are spatial databases containing information about facilities and devices that university uses. Each item in the databases is identified with unique code with hierarchic structure.



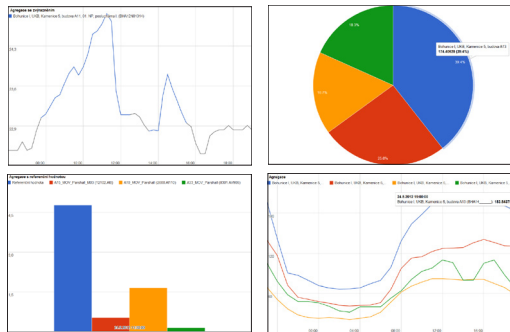
Building management data semantics

Joining data from BMS and from passport allows as to add "machine-readable" information about semantics (meaning) of data points (i.e. inputs, outputs, control variables) in BMS. In addition to the information from technology passport we add information about measured quantity (temperature, pressure, humidity, actual current, voltage...) of each data point.

Attribute	Value	Meaning
Data point	102_AV145	ID of data point in BMS
Location	BHA12P01005	Room 005 in the basement of building 12 at university campus
Device	BAPK	Energy meter
Quantity	ESS0	Electricity consumption

System prototype – Data analysis and visualization

System runs continuous user-defined queries that transform input data from BMS using various methods (selection, filtering, joining, aggregation, grouping). Results of the computations performed by the system can be further analyzed automatically (e.g. Detecting fault on particular device) or stored for later visualization that can be examined by users.



References

- D. Luckham, "The Power of Events", Addison-Wesley, 2002.
- H. Merz, T. Hansermann and C. Hübner, "Building automation", Springer, 2009.
- O. Etzion and R. Nebel, "Event Processing in Action", Manning Publications, 2010.
- A. Kučera, "Complex Event Processing in Building Management Systems", Master thesis, Masaryk University, Faculty of Informatics, 2012.

Complex Event Processing in Building Management Systems – a Prototype

Adam Kučera

Modern (“intelligent”) buildings are often equipped with a Building Management System (BMS) – a network of interconnected devices which ensures integrated building operation, control and monitoring.

The aim of this research is to adapt an existing Complex Event Processing (CEP) framework to process data from intelligent buildings. This will provide building operators with easily understandable outputs that help to optimize building operations and provide early detection of various types of faults.

The research problem is that building operators and maintenance staff are easily overwhelmed by large amounts of data and events produced by the BMS. In this way a BMS is similar to other fields where CEP is successfully used – network monitoring, credit card fraud detection and algorithmic stock trading.

As methodology, a CEP framework was extended to be able to work with specific features of the building management data (location, source device, measured/controlled quantity). A prototype of the system was implemented and tested in a real operation environment on the Campus of the Masaryk University.

The expected outcome of this research is to develop a complete system for fault detection and operation analysis of a BMS, using methods and tools based on CEP.



Adam Kučera is a doctoral student at FI MU. He defended his Master’s thesis on Complex event processing in building management systems in 2012. He now continues his research in the field of the management of building systems infrastructure, fault detection in intelligent buildings and analysis of building management data. He also works at the Facility Management department of the University Campus Bohunice as a programmer since 2008. This gives him an opportunity to apply his results directly in everyday use. He intends on completing his degree in 2016.

Design of a system for the analysis of social media content



Faculty of Informatics, Masaryk University & IBA CZ
Brno, Czech Republic

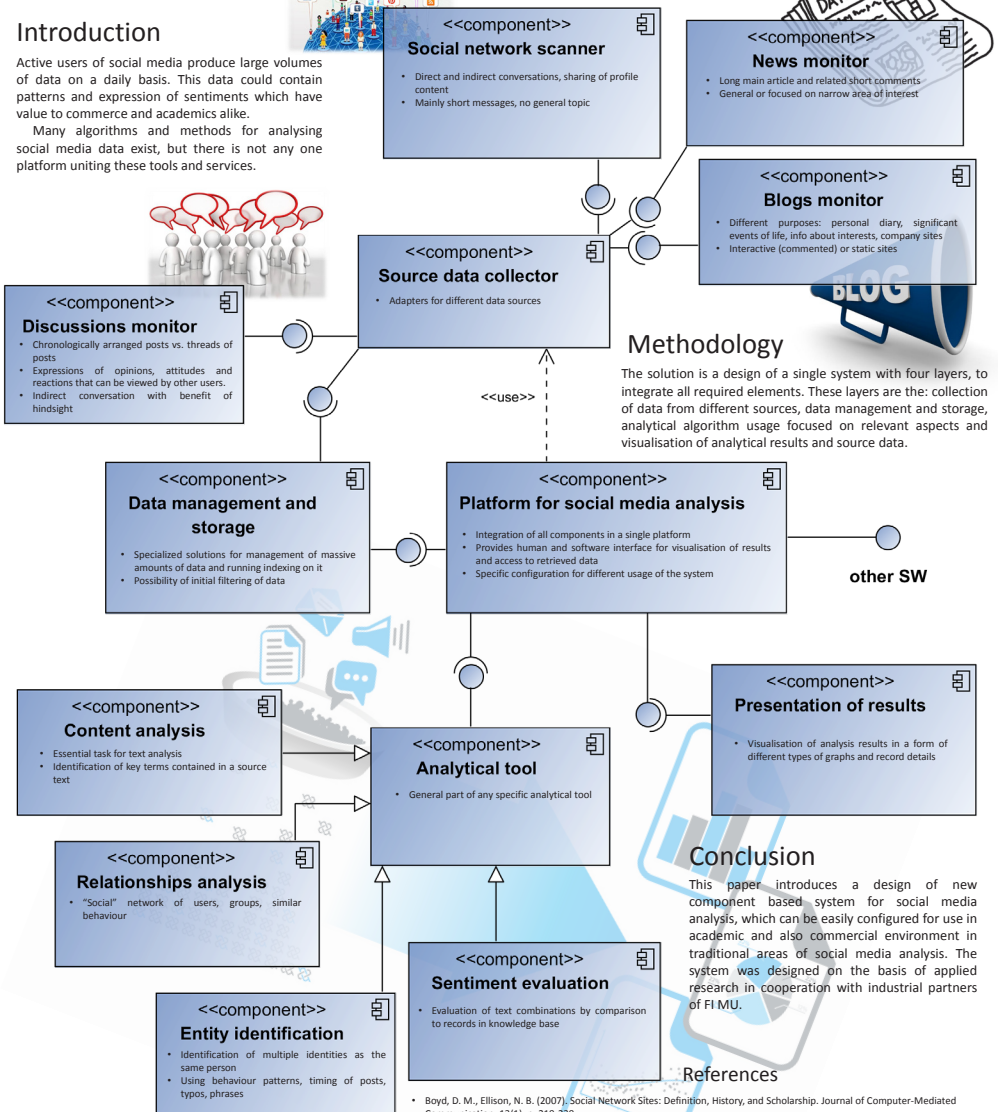
Dalibor Toth
xtoth2@fi.muni.cz

Jaroslav Ráček
Jaroslav.Racek@ibacz.eu

Introduction

Active users of social media produce large volumes of data on a daily basis. This data could contain patterns and expression of sentiments which have value to commerce and academics alike.

Many algorithms and methods for analysing social media data exist, but there is not any one platform uniting these tools and services.



Methodology

The solution is a design of a single system with four layers, to integrate all required elements. These layers are the: collection of data from different sources, data management and storage, analytical algorithm usage focused on relevant aspects and visualisation of analytical results and source data.

Conclusion

This paper introduces a design of new component based system for social media analysis, which can be easily configured for use in academic and also commercial environment in traditional areas of social media analysis. The system was designed on the basis of applied research in cooperation with industrial partners of FI MU.

References

- Boyd, D. M., Ellison, N. B. (2007). Social Network Sites: Definition, History, and Scholarship. *Journal of Computer-Mediated Communication*, 13(1), p. 210-230.
- Pang, B.; Lee, L. (2008). Opinion Mining and Sentiment Analysis. *Foundations and Trends® in Information Retrieval*. 2(2), p. 1-135.
- Ministr, J., Ráček, J. and Toth, D. (2012) "Visualization of the discussion content from the Internet", IDIMT- 2012 ICT Support for Complex Systems – 20th Interdisciplinary Information Management Talks, Jindřichův Hradec, Linz: Trauner Verlag, p. 297-304.

Design of a system for the analysis of social media content

Dalibor Toth, Jaroslav Ráček

Active users of social media produce large volumes of data on a daily basis. This data could contain patterns and expression of sentiments which have value to commerce and academics alike.

The objective of this research is to design a component based platform for processing social media data.

The research problem is that many algorithms and methods exist for analysing social media data, but there is not any one platform uniting these tools and services. The methodology is based on the design of a single system with four layers, to integrate all required elements. These layers are the: collection of data from different sources, data management and storage, analytical algorithm usage focused on relevant aspects and visualisation of analytical results and source data. The final solution will be designed and tuned under laboratory conditions, and tested on accessible social data on the Internet.

The expected results include a versatile social media data analysis tool, which can be personalised for use in many ways in both an academic and a business environment.



Dalibor Toth is a doctoral student at the Faculty of Informatics, Masaryk University (FI MU). He is interested in social media analysis, software engineering and architecture, project management and business processes management. He is a member of the LaSArIS group and participates in teaching of several courses at FI MU (related to his interests). His research is focused mainly in social data analysis, cooperating with industrial partners of FI MU.

Developing the ICT partners network

(project at FI MU)

Q: Best approach?

The faculty currently cooperates with almost 30 companies in area of ICT. Cooperation with industry is one of the key instruments to fulfil FI's vision: "... to become an outstanding research university with significant accomplishments in this area...".

A sustainable approach to developing partners network has to meet some requirements like: a reasonable financial burden, simplicity of implementation of such activities and their attractiveness.

Platform for Industrial Cooperation

Platforma průmyslové spolupráce
CZ.1.07/2.4.00/17.0041

Project partners:


- IBM Česká republika, spol. s r. o.
- Microsoft s. r. o.
- Red Hat Czech s. r. o.
- JIC, zájmové sdružení právnických osob



From idea to realisation

The Platform project addresses those requirements by promoting **four key activities** that are interconnected through the **two-phased internship model**:

- student projects – team or individual • related to curricula or student's interests • projects are lead by faculty supervisors and consulted by company experts
- workshops – orientation on technology and business → student projects and workshops together create a **first phase** of the internship model allowing students to meet experts from companies at the faculty
- internships at companies – correspond to curricula and student projects • with support of the faculty
- → internships are a **second phase** – they offer students opportunities to use their knowledge they have gained by participating in first-phase activities
- own business – support for start-ups and spin-offs • incubation at South Moravian Innovation Centre (JIC) with professional support from the faculty



Q: What is in it for me?

Benefits	FI	S	C
cooperation as a source for new ideas	✓	✓	✓
student as an agent to transfer knowledge	✓	✓	✓
distribution of the "resources load"	✓	✓	✓
start-ups and spin-offs as a future FI partner	✓	✓	✓

FI – Faculty of Informatics, S – student, C – company



This project is co-financed by the European Social Fund and the state budget of the Czech Republic

Contact information

<http://lasaris.fi.muni.cz>

pps@fi.muni.cz

Project lead: Tomáš Pitner

Project manager: Stanislava Sedláčková



Developing the ICT partners network (project at FI MU)

Stanislava Sedláčková

The Platform for Industrial Cooperation project (PPS) builds on existing cooperation between the faculty and companies in industry. As this cooperation grows, it may become disordered and less effective. The main objective of the project is to advocate a systematic approach focused on four parallel activities, creating the platform that enhances cooperation between the academic and application areas in fields of research, development and education.

The four activities – student projects and workshops prepared with experts from companies, internships and the support for start-ups and spin-offs – are interconnected through the so-called two-phased internship model. By participating in these activities, students become familiar with a particular technology on the intermediate level. To enrich their theoretical knowledge with practical experience, students are encouraged to attend internships at selected companies. For those with an entrepreneurial spirit, a programme to gain knowledge and experience in the area of business and innovation in ICT is prepared.

This approach offers intensive involvement of all parties in common activities creating, besides other benefits, an environment that brings better “resource load” distribution, knowledge transfer, new ideas and attractive opportunities for growth and improvement.



Stanislava Sedláčková is a member of the realisation team of the “Platform for Industrial Cooperation” project at the Faculty of Informatics, MU. Her responsibilities include managing particular project activities and communication with external project partners. She believes that cooperation between faculty and commercial companies improves conditions for professional growth of all participants and may generate notable results in applied research and development.

DISTRIBUTED EVENT-DRIVEN MONITORING MODEL FOR CLOUD DATACENTERS



Daniel Tovarňák, Tomáš Pitner


Masaryk University, Faculty of Informatics, Botanická 68a, 60200 Brno, Czech Republic

Introduction

When monitoring distributed infrastructure, huge amounts of monitoring data are typically produced by multiple distributed producers spread across many individual computing nodes. There are reports of monitoring data rates up to 1 MB/s per node [1].

Scope


In order to determine the state and behavior of a particular resource all the relevant data must be collected, processed and evaluated. We are interested in behavior-related monitoring data (e.g. logs) in particular.



Additionally, when monitoring cloud datacenters and all its main layers simultaneously, the monitoring system must also be able to cope with dynamically changing multi-tenant environment [2].

S - sensor, P - producer, PR - processing agent, C - consumer, 1|2|3 - access permissions

Research Goals

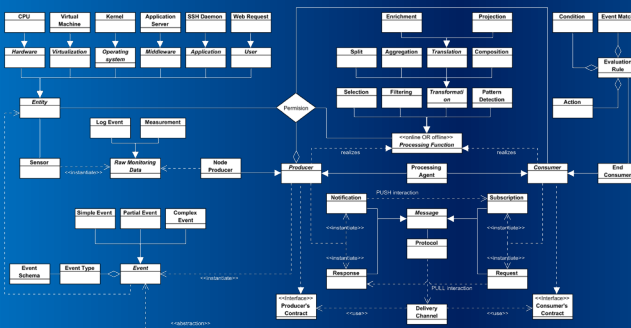


The volume, velocity, and variability of monitoring data produced by modern cloud datacenters multiply and the current approaches are insufficient for online processing of behavior-related monitoring data.

Aims and Objectives

The goal of our research is to design a *Distributed Event-driven Monitoring Architecture* (DEDMA) following a new monitoring model with focus on novel mechanisms and algorithms in the areas of monitoring data generation, production, distributed collection and processing, as well as multi-tenant monitoring.

Distributed Event-driven Monitoring Model



The model is founded on pattern-based publish-subscribe interaction and Complex Event Processing principles [3] with support for multi-tenancy. The monitored entities can belong to any of the main cloud layers.

Methodology

The solution and its respective components and algorithms will be evaluated experimentally and compared to existing approaches. To be able to achieve this, prototype implementation will be developed.

Conclusion and Expected Results

This poster presents the notion of *Distributed Event-driven Monitoring Architecture* (DEDMA) and its model. When compared to existing approaches we expect improvements in the terms of intrusiveness, network overhead, and throughput with respect to the number of producers, consumers, volume, velocity, and variability of monitoring data.

References

- [1] Cretu-Ciocarlie, G. F., Budiu, M., and Goldszmidt, M. Hunting for problems with artemis. *WASL'08, USENIX Association*, 2008
- [2] Etzion, O., and Niblett, P. *Event Processing in Action*, 1st ed. Manning Publications Co., Greenwich, CT, USA, 2010.
- [3] Tovarňák, D., and Pitner, T. Towards Multi-Tenant and Interoperable Monitoring of Virtual Machines in Cloud. *SYNASC'12*, 2013



Distributed Event-driven Monitoring Model for Cloud Datacenters

Daniel Tovarňák, Tomáš Pitner

When monitoring distributed infrastructure, data related to the particular resource are typically produced by multiple distributed producers spread across many individual computing nodes. In order to determine the state and behavior of a resource, all the relevant data must be collected, processed and evaluated without overloading the computing resources and flooding the network. The volume, velocity and variability of monitoring data produced by modern cloud datacentres multiply and there is a need for new approaches since current approaches are insufficient for online monitoring data processing.

The objective of this research is to propose a real-time monitoring architecture utilizing a novel distributed event-driven model.

The methodology will involve experimentation and a comparative study with existing approaches.

Expected outcomes include improvements in the terms of intrusiveness, network overhead, and throughput with respect to the number of producers, consumers, volume, velocity, and variability of monitoring data.



Daniel Tovarňák is a Ph.D. student at Masaryk University (FI). His research interests include distributed infrastructures (e.g. Cloud and Smart Grid) monitoring, and Complex Event Processing. Daniel is also a member of the Lab of Software Architectures and Information Systems and one of the Technology Platform Leaders in Platform for Industrial Cooperation project. He intends on completing his doctoral degree at the end of 2014.

Graphics design of a mobile UI for primary school administration

OBJECTIVE

- The aim of this project is to reduce the load of daily administration on teachers.
- It is focussed on the UI of mobile devices and touch screen interaction.
- Simple design and ease of learning are the main design considerations.

- Administration in the classroom and after class hours takes much time.
- Lost class books have to be rewritten – difficult to retrieve old data.
- Copying of data from one form to the next wastes time.
- The goal is to reduce the administrative load of teachers.
- Currently Web applications are available to reduce the administrative load at schools.

RESEARCH PROBLEM

- Related Web systems and applications are not easy to use.
- Tablet technology is not commonly used in Czech primary and secondary schools yet.
- Technology is quite cheap.
- A specialized GUI should be easy to learn and use.

METHODOLOGY

- Analysis of existing applications for Czech school administration.
- Analyzing tablet UI needs, advantages and disadvantages.
- Meetings with teachers.
- Identifying daily goals and scenarios.
- Testing new GUI with teachers.

EXPECTED RESULTS

- The goal is to create an interactive UI simulator on a tablet device.
- Videos will be used to show interaction, scenarios and the achievement of selected goals.

References:

Cooper, Allan. 2007. *About face 3*, Wiley Publishing, Inc.
Tidwell, Jenifer. 2005. *Designing Interfaces*, O'Reilly

Author: Lenka Plháková



INVESTICE DO ROZVOJE VZDĚLÁVÁNÍ



Graphic design of a mobile UI for primary school administration

Lenka PlhÁková

This project is focused on the Graphical User Interface (GUI) design of an administration system for teachers. A UI is important for teachers to effectively work with IT systems on mobile devices. Some Web-based systems and applications, which are attempting to help teachers with administration do exist, but these systems are not easy to use.

The objective of the project is to produce a mobile (tablet) interface which will enhance the way teachers use IT for administrative purposes. The research problem is that teachers often spend a lot of time on administration instead of teaching or preparing materials for teaching. Data is mostly in paper form and there are problems with lost administrative forms and copying data by hand.

As methodology, teachers will be interviewed as potential customers; a new UI will be designed and finally tested by them. This project will consider one specific educational UI of Czech computer systems. Expected results include an interactive screen design and accompanying video showing selected goals and scenarios.



Lenka PlhÁková is a Masters student at FI MU. She is interested in graphic design and user interfaces. She is currently working in Celebrio software s.r.o. as a graphic designer. Her Master's thesis will be about creating a graphical user interface on mobile devices (especially tablets) for primary school teachers. This work is inspired by the Celebrio system, which is an operating system for seniors (running on computers and tablets).

Implementation of the internship model *in ICT area*

(project at FI MU)

Best concept?

Whilst company resources are significantly limited, the effect of internship (transfer of knowledge and the overall contribution for all parties) needs to be maximized

→ internship is not suitable to draw up like "vacation job" but better approach is to **plan and managed as a project**


→ goals, time limitations, resources has to be set → **student's vision is the key factor**

Platform for Industrial Cooperation

Platforma průmyslové spolupráce
CZ.1.07/2.4.00/17.0041

Project partners:

- IBM Česká republika, spol. s r. o.
- Microsoft s. r. o.
- Red Hat Czech s. r. o.
- JIC, zájmové sdružení právnických osob



Two-phase model

First phase

- students can work on **real projects or theses** in application area with support of teachers from FI
- students advance knowledge by attending **technical workshops, meetings with experts** from business and/or standard curricula lectures
- students have opportunity to meet specialists from companies directly at FI and become familiar with **practical applications of theoretical approaches and modern technologies** used by companies
- participating in activities of the first phase may help students to determine their **own career vision**

Second phase

- **students go to pre-selected companies** where they can use their experience and knowledge gained from first phase of the internship model, and
- their vision can guide them during the process of setting strategy and goals for their internship → **thus they are able to plan and manage it**

Q: What is in it for me?

Benefits	FI	S	C
shaping own future		✓	✓
establishing contacts	✓	✓	✓
praxis related to studies	✓	✓	✓
enrichment of theoretical knowledge by experience	✓	✓	✓
systematic build and extending knowledge	✓	✓	✓
source of potential employees and employers		✓	✓
financial support of PPS		✓	✓

FI – Faculty of Informatics, S – student, C – company



INVESTMENTS IN EDUCATION DEVELOPMENT

This project is co-financed by the European Social Fund and the state budget of the Czech Republic

Contact information

<http://iasaris.fi.muni.cz/pps/internship>
pps@fi.muni.cz

Project lead: Tomáš Píthner

Internship manager: Jana Bartáková



Implementation of the internship model in ICT area (project at FI MU)

Jana Bartáková

The main focus of this project is to introduce and expand the concept of internships at commercial partners of the project. The project aims to connect theory and practice and improve student's preparation for practice. The internship model is popular and provides maximum benefits for students, but the process of internship should operate effectively at the Faculty of Informatics at Masaryk University.

This model has two phases:

1. During the first phase students can get to know the technologies which are used in specific companies involved in the projects. By working on these projects, students will gain an understanding of the standard of education level. This will increase their chances of success in open competition. To assist in this process, several seminars and workshops are being prepared, where the students can develop an understanding of the specific technologies.
2. In the second phase the student is placed and starts working in the specific company. The student cooperates with the company, and in the process the students make progress with his/her diploma thesis.

It is expected that students will learn the relevant skills in the seminars and workshops, and will furthermore experience growth in their careers as they cooperate with the company on the diploma thesis or project.



Jana Bartáková is an Internship programme coordinator and Project administrator for the Platform for Industrial Cooperation at the Faculty of Informatics at MU. Jana has been working on this project since July 2011. She participates on a two-phase internship and cooperates with students who want to work at commercial companies. Her primary focus is the co-ordination of activities with Faculty partners to develop the professional growth of students.

Indoor User Localization Using Mobile Devices

lasaris

pps

Jonáš Ševčík

Introduction

- difficult navigation in complex buildings
- even if equipped with signs
- no direct visibility to GPS satellites
- GPS tracking not possible

Methodology

- localization system based on [1, 2, 3]
- consists of three localization techniques
- techniques merged together for more accurate results

Wi-Fi Localization

- Received Signal Strength fingerprinting
- used to create a database of APs and their RSS
- mapped to location coordinates
- receiver's location estimated from SS maps



Dead Reckoning

- calculating position using previously determined coordinates advanced by known speed and course
- position tracking - calculating steps
- step length estimated using neural network [2]
- course determined by gyrocompass.

Sequential Monte Carlo Filtering

- particles evenly spread in probable location determined by RSS fingerprint database
- particles set to motion by step detection events
- eliminated particles, which hypothetical motion leads through impassable obstacles
- results in improvement of location estimation.

Preliminary Results

- implemented prototype Android application
- results accurate to 2.3 meters

References

- [1] A. Chen et al: "An algorithm for fast, model-free tracking indoors", SIGMOBILE MCCR, vol. 11 no. 3, 2006
- [2] S. Cho: "MEMS Based Pedestrian Navigation System", JN, vol. 59, pp. 135-153, Royal Institute of Navigation, 2006
- [3] M. Arulampalam et al: "A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking", SP, vol. 50 no. 2, 2002

Acknowledgments

Thanks to all the staff members of the Masaryk University Department of Building Passporization for providing maps and consultations. Also, great thanks to Michal Holčík and Adam Bertholdy for implementing localization prototypes.

Indoor User Localization Using Mobile Devices

Jonáš Ševčík

This research addresses techniques suitable for indoor user localization using mobile devices, without the use of GPS technology.

The objective of this research is to design and implement an indoor localization prototype system, with results which are directly usable.

The research problem is that GPS signals are weak or absent in closed spaces, and cannot be used reliably to identify locations inside buildings.

Several technologies are used as part of the experimental methodology to implement the system. These include wi-fi tracking, step detection, dead reckoning, and Sequential Monte Carlo filtering. These technologies were combined and coded in the Java programming language to form the localization system. At this stage the implementation has been done for the Android platform, and could be expanded as future research to include iOS and possibly other platforms.

Preliminary results of the prototype system indicate an accuracy of 2.3 meters.



Jonáš Ševčík is currently part of the PhD program at Masaryk University. His research is focused on mobile platforms and he is developing a pedestrian navigation system for these devices. Besides doing research and working hard for the university, he likes traveling to Asia and learning foreign languages. If everything goes smoothly, he hopes to establish his own company, which will earn him enough money for early retirement.

Semantically Partitioned Peer to Peer Complex Event Processing Exploiting Information Loss



Masaryk University
Faculty of Informatics
Božetěchova
602 00 Brno
Czech Republic

Filip Nguyen
xnnguyen@fi.muni.cz




Abstract

Scaling Complex Event Processing (CEP) applications is inherently problematic. Our solution for scaling CEP applications is fully distributed and aspires to scale CEP to the limits of current hardware. Our solution simplifies resident Event Processing Network abstraction and adds features on the level of CEP that change direction of its usage.

Related Work

Complex Event Processing was introduced by David Luckham. We are mainly concerned with subarea of Luckham's work related to distributed CEP [8] (also studied by [6] and [4]). Motivation of our work stems from our work related to event processing [3]. We have applied our theoretical ideas in concepts introduced in [2] and gave brief introduction to our overall research in [1].

Basic Approach




Very simple query that matches events that happen in the time window of 0.1 second.

```
select name(E), name(E2)
where abs(time(E2) - time(E1))
< 0.01s
&& E2 != E1
```

This query needs all the produced events.

Suppose we know that P1 and P3 produce events at the same time with high probability. Then we can add an engine between them and match events.



Here the second engine between P1 and P3 will match events and is loaded with less events than the former engine. Unfortunately the event produced at 12:20 by P5 will not be matched. This is the trade off situation in our solution.

Event

The definition of an event varies based on context of CEP. However some parameters are always the same. Each event has defined time of creation and producer. Events should be as fine grained as possible - to allow effective CEP. That means thousands, even millions of events per second are desirable. This is not uncommon thing today with advent of social networks, faster networking hardware and computer driven high frequency trading.

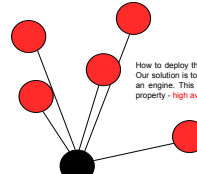
Distributed CEP

There is ongoing research to distribute CEP. Every author makes his own definition of distributed CEP. Usually, it refers to a use of filters on producers or parallelizing existing CEP operators. We see distributed CEP differently. We aim to distribute the processing at semantic level. We do not want to just filter unknown events. We allow users to leverage standard operators and give them framework to easily trade off processing power with matching precision.

Partitioning Algorithm

Partitioning algorithm is used to select O of course, the number of these sets is sets of producers among which the exponential, thus it is necessary to select engine is placed. This is equivalent to such amount that will balance the trade creating channels between peers.

Peer Network



How to deploy the engines dynamically? Our solution is to turn each producer into an engine. This way we gain additional property - high availability.

A node is said to be a peer.

We refer to this model as peer to peer model. The events are distributed throughout the formed network. Some of the events level on dedicated paths, some are broadcasted. This behavior is based on the result of partitioning algorithm.

CEP Based and Monte Carlo Algorithm

There are many possible partitioning algorithms. Beginning with the manual, CEP based partitioning algorithm. This algorithm is dynamic (selects subsets of peers at runtime of matching). Input for this algorithm is CEP rule provided by user. For example:

```
subselect originator(CE1)
where abs(time(CE1) - time(CE2)) < 1 minute
```

The CEP query above will ensure that channel will be created among those peers that generated coarse grained events in window of 1 minute.

Another, fully automated, partitioning algorithm is Monte Carlo partitioning algorithm. This algorithm is based on statistical analysis of events firing and producers that fired the event. This way, the peer network automatically makes itself more effective.

Coarse Grained Event

Event is relatively small, high frequency event, e.g. scanning of barcode in supermarket. Coarse Event is a notion used to describe event that happens much less frequently than Event, e.g. payment by credit card in supermarket.

Results

We believe that fully distributed peer to peer CEP is inevitable solution to high volume event streams. Our implementation of presented concept is called peer CEP (PCEP). The main property of PCEP is semantic scaling. The scaling is not done by brute force or by exploiting specific feature of specific event context, but it is done by exploiting partitioning of peers according to their affiliation to matching rules.

The developed distributed engine is written in Java and thus runs on heterogeneous platforms. In the implementation we leverage distributed algorithms developed in their natural form - not optimized to the state of being obfuscated code. In theoretical point of view, our solution introduces rigorously defined trade off between matching capabilities and throughput of the events. In the future we plan to extend this knowledge by revealing statistical properties of mentioned trade off situation.

There is another result we present - partitioning algorithms.

We believe those algorithms may be extended and generalized to be used in other fields for set partitioning and analysis of data sets. These algorithms join several Distributed Algorithms, Statistics and Complex Event Processing. We theorize, that the partitioning may be done in a distributed fashion.

We also believe in adoption by users. We strive robust architecture. Our solution is Open Source and we plan to apply for Apache Incubation. We believe the science should be done for greater good and sharing the code will improve the implementation. Lastly, our solution is not mutually exclusive with recent research in the area of CEP. It will be possible to use standard CEP engines on the peer nodes and thus augmenting existing tools with PCEP.



References

- [1] Nguyen F., Pilner T. 2012. Information system monitoring and notifications using complex event processing. In Proceedings of the Fifth Balkan Conference in Informatics (BCI'12), ACM, New York, NY, USA, 211-216.
- [2] Kunc P., Nguyen F., Pilner T. 2013. Towards Effective Social Network System Implementation. New Trends in Databases and Information Systems in Intelligent Systems and Computing. Springer Berlin Heidelberg, 327-336.
- [3] Nguyen F., Szabalek J. 2011. NoX service oriented multi-platform notification system. In Computer Science and Information Systems (FedCSIS), Szczecin, Poland, 313-316.
- [4] Wu, E., Dao, Y., Rizvi, S. 2006. High-performance complex event processing over streams. In Proceedings of the 2006 ACM SIGMOD international conference on Management of data - SIGMOD '06.
- [5] Akram, S., Marazakis, M., Bilas, A. 2012. Understanding and improving the cost of scaling distributed event processing. In Proceedings of the 8th ACM International Conference on Distributed Event-Based Systems (DEBS '12), ACM, New York, NY, USA, 290-301.
- [6] Luckham, D. C., Frasca, B. 1998. Complex Event Processing in Distributed Systems. In Stanford University, Vol.26.

Links

GitHub: <https://github.com/nguyenfilip/pep>
 LaSARIS: <http://lasaris.fi.muni.cz/>
 LinkedIn: <http://www.linkedin.com/pub/filip-nguyen/2760/5b4>
 University: <http://www.fi.muni.cz>

Semantically Partitioned Peer to Peer Complex Event Processing Exploiting Information Loss

Filip Nguyen

Scaling event processing is an inherently complex process. The use of distributed architectures in its implementation is a well-known endeavour, but its limits have been identified over time.

The objective of this study is to provide a distributed solution to this problem. The most important drawback of the scaling process is the centralized approach of the event processing. Current distributed architectures use filtering on producers or horizontal scaling of parallelizable operators.

The problem is the fundamental method of matching of related events. If two events are from one event set, the whole set must be considered for the processing. Thus the size of the event set creates a bottleneck.

To overcome this limitation, a semantic partitioning process is proposed, which allows a trade-off between information loss and matching power. The information loss is exploited by partitioning the event set into subsets with least probable correlation among events.

The resulting architecture is completely peer to peer, thus allowing new processing nodes to join the distributed cluster in order to take on an equal load for processing. This will either improve matching capabilities or allow higher throughput.



Filip Nguyen is a JBoss engineer specializing in Java middleware technologies. The main focus of his research is integration technologies and data mining technologies. Filip is an active community member and contributor to several JBoss projects. He is busy with his PhD in applied computer science for the software architecture-oriented laboratory LaSaRIS, and his field is distributed complex event processing. He has published 3 conference papers related to the topic of Java middleware and continues his work toward publishing his on-going results in the area of semantic partitioning for CEP.



Supporting the Process of Learning Mobile Application User Interfaces

Lucia Tokárová, Faculty of Informatics, Masaryk University, Brno, Czech Republic



With the recent expansion of the mobile industry, applications for mobile devices are becoming more complex, empowering people to perform more advanced tasks. However, modern mobile user interfaces introduce several challenges, which affect learnability of mobile applications. For example:



Small screens provide space for displaying high-priority functions and reduce discoverability of advanced features [4].



Gestural UIs are engaging and intuitive for simple tasks but not for advanced operations [4].



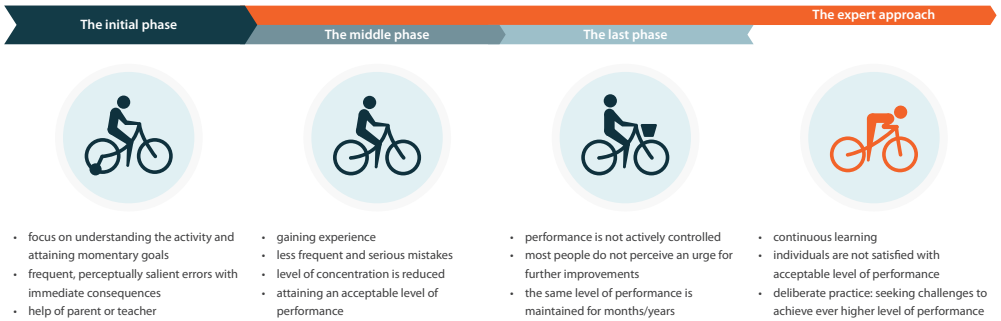
Sessions with mobile applications are short [1, 6] and variable in the context of use [6], which affect users' attention.




Solely visual user interfaces without haptic feedback prevent activation of the muscle memory.

The process of learning

Learning is a long-term process. Individual's needs are changing over time. To achieve the highest levels of expertise, learners should be engaged in deliberate practice to continuously improve their performance. (via [5])



The process of learning mobile applications



Research question: How to continuously support the process of learning mobile application user interfaces?

Objective

The objective of this research project is to investigate how people learn to use mobile applications, and how can this process be supported in different phases so that they quickly perceive the value of the application, accomplish basic tasks, and gradually learn new features in a natural way.

References

- [1] Böhmér, M. et al. 2011. Falling asleep with angry birds, facebook and kindle: a large scale study on mobile application usage. *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI 2011)*, Sweden, Stockholm, 30 August – 2 September 2011:47–56.
- [2] Burnett, M.M. et al. 2011. Gender pluralism in problem-solving software. *Interacting with Computers*, 23(5):450–460.
- [3] Min, A. 2011. 26% of Mobile Application Users are Fickle — or Loyal. <http://www.localytics.com/blog/2011/26percent-of-mobile-app-users-are-either-fickle-or-loyal/> [05 December 2012].
- [4] Norman, D. 2010. Natural user interfaces are not natural. *Interactions*, 17(3):6–10.
- [5] Oulasvirta, A. et al. 2011. What does it mean to be good at using a mobile device? An investigation of three levels of experience and skill. *International Journal of Human-Computer Studies*, 69(3):155–169.
- [6] Oulasvirta, A. et al. 2005. Interaction in 4-Second Bursts: The Fragmented Nature of Attentional Resources in Mobile HCI. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2005)*, USA, Oregon, Portland, 2 – 7 April 2005:910–928.

Methodology & expected results

This study will consist of observation of users' behavior in the process of learning mobile UIs. Patterns in users' behavior will be investigated, in order to create learning profiles of representative user groups. Further profile examination should lead to the design of support mechanisms that will encourage various types of learners in the process of continuous learning of mobile applications.



INVESTMENTS IN EDUCATION DEVELOPMENT

Supporting the Process of Learning Mobile Application User Interfaces

Lucia Tokárová

With the recent expansion of the mobile industry, applications for mobile devices are quickly becoming more complex, empowering people to perform more advanced tasks. However, current mobile user interfaces introduce several challenges, such as direct manipulation, gesture control, solely visual feedback, or limited screen size. These factors affect the learnability of mobile applications.

The primary objective of this research project is to investigate how people learn to use mobile applications and how this process can be supported. This needs to be done so that people can perceive the value of the application, accomplish basic tasks and gradually learn new features in a natural way.

A pilot study will consist of observation of users' behavior in the context of skill acquisition within mobile user interfaces. Patterns in users' behavior and links between behavior and background information obtained by questionnaire will be investigated. It is anticipated that the study will demonstrate the correlation between learning strategies and personal information about users. The results of this study will aid in creating learning profiles of representative user groups. Further examination of these profiles should lead to the design of support mechanisms that will encourage various types of application learners in the process of continuous learning of mobile applications.



Lucia Tokárová is a doctoral student at Masaryk University, Faculty of Informatics. She is currently enrolled in her second year of the Computer Systems and Technologies program and she intends on completing her degree in 2015. Her research is focused on mobile user experience design, particularly the aspects of continuous learning of mobile application user interfaces. Outside of school, she works as a freelance user interface designer, focused on mobile and web applications.