CONET Newsletter

INFSO-ICT-224053 www.cooperating-objects.eu

Issue nr.9, June 2010

Contents

Editorial

Welcome to the ninth issue of the CONET newsletter. CONET is the EU FP7 network of excellence on Cooperating Objects, merging the fields of embedded systems for robotics and control, pervasive computing and wireless sensor networks. CONET focuses on establishing the field of Cooperating Objects within the research and industrial community, thus strengthening the position of Europe in the research landscape.

This issue presents a guest column from Prof. Panos K. Chrysanthis (University of Pittsburgh) about next generation data stream management systems. In the member profile section, we take the opportunity to present the Mobile Computing Laboratory of the Department of Computer Science, University of Cyprus (UCY). This issue also includes information about two research projects, the SpitFire project from DERI and the InteliSense project from UCY.

If you are interested in obtaining up-to-date information about the CONET project please visit our website at: http://www.cooperating-objects.eu

We hope you will enjoy this issue ■

AQSIOS – Next Generation Data Stream Management System

By Prof. Panos K. Chrysanthis, University of Pittsburgh, USA.

On-line monitoring applications are widely enabled today by Wireless Sensor Networks (WSN), in conjunction with Data Stream Management Systems (DSMS). Environmental monitoring is an example of a monitoring application, which combines both surveillance and detection, as well as scientific exploration and discovery. Individual sensor devices produce simple periodic readings such as temperature, humidity, and sound measurements, which are all combined into streams of data. Depending on the underlying setup of the WSN, these streams could be heterogeneous and have a steady and/or a bursty rate. Environmental monitoring applications compose these individual readings into more meaningful complex events (e.g., fire, flood). At the same time, these applications also perform aggregations and calculate various useful statistics for future reference and analysis. In both cases, the applications' requests can be expressed in terms of Continuous Queries (CQs) running in a DSMS. These CQs have different Quality of Service (QoS) requirements such as response time and Quality of Data (QoD) requirements such as freshness. For example, CQs that detect hazardous conditions require a lower response time than CQs that collect statistics.

A critical challenge for DSMS is the scalable processing of time-varying, transient and possibly unbounded data streams to incrementally provide timely response to many long-running and heterogeneous CQs. The AQSIOS project [1] and prototype DSMS of the Advanced Data Management Technologies Laboratory (ADMT Lab) at the University of Pittsburgh has approached this challenge from three different angles: (1) Multiple Query Optimization, (2) Operator Scheduling and (3) Load Management. Multiple Query Optimization (MQO) aims to generate query execution plans that minimizes the tuple processing delay. Scheduling, on the other hand, aims to select the operator to execute next in a query plan that minimizes queuing delays. whereas load management deals with overloading situations. The

AQSIOS prototype is built upon the STREAM DMSM [2].

Multiple Query Optimization: Aggregate Continuous Queries (ACQs) are one of the most common queries across all applications. For example, an aggregate function, e.g., Count(), could be used to monitor the number of moving objects to detect dangerous congestion beyond a threshold. Another example of aggregate function is Average() which could be used to monitor the average temperature and humidity in the past hour. Given the cost and commonality of ACQs, their shared processing is crucial if scalability is to be achieved. In order to reduce the total number of aggregates performed, AQSIOS optimizer exploits partial aggregation to process input tuples only once and assemble the partial aggregates into final aggregate values. The current state-of-the-art shared MCQ scheme also shares partial aggregation to blindly group all ACQs on the same aggregate function into one execution tree. AQSIOS optimizer, on the other hand, achieves orders of magnitude improvement over the state-of-the-art by selectively grouping ACQs into multiple execution trees.

Operator Scheduling: AQSIOS aimed to identify the most appropriate metric to optimize for a given application and to develop an efficient scheduling policy for it. This goal is complicated by the fact that the scheduling policy must take into account that the CQs are heterogeneous, i.e., they may have different time complexities (the amount of processing required to find if input data represents an event), and different productivity or selectivity i.e., the number of events detected by the CQs.

AQSIOS showed that if the objective is to optimize the response time, then the ``right'' strategy is to schedule CQs according to their *output rate* [3]. Specifically, it presented a new scheduling policy called *Highest Rate* (HR). Under HR, the priority of a query is set to its output rate where the output rate of the query is the ratio between its expected selectivity and its expected cost.

Although scheduling to minimize average response time works well for homogeneous workloads, it is well known that it is not ``fair'' for end-users to use average response time as the metric to optimize when the workload is heterogeneous. A user that issued a complex query expects a higher response time than the user that issued simple query. A metric that captures this phenomenon is *average slowdown*. The slowdown of a job is the ratio of the response time of the job to its ideal processing time. Given that HR does not optimize average slowdown because of the ``probabilistic'' nature of CQs where the selectivity might not equal 1, AQSIOS proposed *Highest Normalized Rate* (HNR) policy, which sets the priority of a query to the ratio of its selectivity over the product of its expected cost and its ideal total processing cost [4].

Improving the average-case performance usually comes at the expense of fairness toward certain classes of queries that might experience *starvation.* AQSIOS eliminates starvation by developing a hybrid scheduling policy that optimizes for both the average and worst cases of slowdown.

In addition to new scheduling policies, AQSIOS considered two special problems that are unique to DSMSs and should be exploited by the query scheduler [4]. The first problem is inherent in the dynamic nature of data streams where the distribution of data may vary significantly over time. Towards solving this problem, it proposed an *adaptive scheduling* mechanism that allows the proposed policies to react quickly to changes in data distribution. The second problem is inherent in the inter-dependency between operators in CQs due to the presence of join or shared operators. Towards this, it developed a scheme to appropriately set the priority of shared operators.

All existing CQ scheduling policies including HR and HNR are oblivious to the different importance levels of different CQs. This means that the critical queries for detecting, for example, fire or flood might be dragged down by the statistics gathering queries. As a first step in handling different ranks of CQ classes, AQSIOS experimented with *Weighted HR* and proposed *Continuous Query Class* (CQC), a new scheduling policy as an alternative. CQC employs two-level scheduling that combines both the Round-Robin and HR schedulers [5]. It provides the lowest response times for classes of critical CQs, while at the same time keeping reasonable response times for the other classes down the (criticality) rank.

Finally, AQSIOS was the first to identify the trade-off between freshness and performance for DSMSs and proposed FAS-MCQ, a *freshness-aware* continuous query scheduling policy. It also proposed a weighted version of FAS-MCQ that supports applications in which queries have different priorities and a parameterized version that is able to balance the trade-off between QoD and QoS according to the application's requirements [6].

Load Management: AQSIOS explored a proactive approach based on admission control of continuous queries and reactive approach based on load shedding. The developed admission control schemes are auction-based and primarily targeted for data stream monitoring/management services, which seek to profit from it [6]. Specifically, when submitting a query, each user also submits a bid

of how much she will commit to paying for that query to run. The admission control auction mechanism then determines which queries to admit, and how much to charge each user. Our goal was to admit queries and set payments in a way that maximizes system revenue while incentivizing users to use the system honestly.

Towards load shedding, AQSIOS first developed a simple yet complete load manager that monitors the system workload at runtime and automatically decides the appropriate amount of random shedding from the input data when the system is overloaded, so as to keep the response time of all outputs to be within a user-specified threshold. AQSIOS also implemented a new control-based load shedder that overcomes the two weaknesses of the state-of-the-art load shedder. It does not require the manual tuning of the parameters used to estimate the system capacity, and is dynamic enough to self-correct the errors in its cost estimation models for complex queries. AQSIOS further proceeded to propose an alternative scheme that achieves equivalent performance as the control-based load shedder, but follows a different design, so that it can be used in a DSMS where the delay estimation model, the heart of the control-based scheme, is not applicable.

In summary, AQSIOS is an on-going project and the investigations in multiple query optimizations, operator scheduling and load management mentioned above currently continues. Recently, AQSIOS expanded its goal to consider data dissemination. Specifically it considered the situation where the results of the continuous queries, which are in the form of individual data streams, are disseminated to the users of hand-held devices (or to activators) over a shared broadcast medium [8]. The goal in this work is to design operator placement algorithms that work in synergy with the broadcast organization so that to minimize the total energy consumption on the hand-held devices. Our approach is based on the observation that the energy costs for receiving data are typically much higher than the energy costs for processing that data, an observation that has also motivated in-network processing in wireless sensor networks.

The AQSIOS project is funded by the National Science Foundation under awards IIS-0534531 and NSF-IIS 0746696 and directed by Professors Panos K. Chrysanthis, Alexandros Labrinidis and Kirk Pruhs. If you are interested in obtaining up-to-date information about AQSIOS visit http://db.cs.pitt.edu/streams ■

Member Profile: UCY

By Prof. George Samaras, University of Cyprus, Cyprus

The University of Cyprus was established in 1989 and admitted its first students in 1992. Postgraduate studies started in 1995. The main objectives of the University are twofold: the promotion of scholarship and education through teaching and research, and the enhancement of the cultural, social and economic development of Cyprus. The University is a public corporate body, governed by the Council and the Senate. The Faculties and Departments are administered by Boards; each Faculty is headed by a Dean and each Department is headed by a chairperson.

Figure 1: View of the University of Cyprus main building

The Department of Computer Science is a research-oriented department active in several areas of Information Society Technologies. The research interests of the Department vary from purely theoretical and foundational topics to applied development where, to the extent possible, this applied research is aimed to be directly useful to the local community of Cyprus. Over the last years, the Department has been continuing and strengthening its many collaborations with other universities and research institutions all over the world, expanding its research into new areas. The general research areas of the Department include: Sensor Networks, Data Management, Mobile computing, Parallel and Distributed Systems and Computation, High-Speed Communication Networks, Internet Technologies, Computer Architecture, Intelligent Systems, eLearning, Multimedia Systems, eGovernment, eHealth and eLearning. Furthermore, the Computer Science Department has participated in numerous national and international (FP4-7) projects. Finally, the University of Cyprus is in very close collaboration with the Cyprus Department of Information Technology and the Ministry of Health for the definition of related IST strategies and their implementation. ■

SPITFIRE: Semantic-Service Provisioning for the Internet of Things using Future Internet Research by Experimentation

By Marcel Karnstedt, Digital Enterprise Research Institute, NUI Galway, Ireland

We are currently witnessing the integration of two long-lost relatives: on the one hand general purpose computing and its networking backbone, the Internet, and on the other hand embedded computing. However, the development of applications exploiting this merged infrastructure is currently exceedingly difficult. SPITFIRE works towards the realisation of a stronger connection between the natural and the digital worlds. The goal of this project is hence to investigate unified concepts, methods, and software infrastructures that allow the efficient development of applications that span and integrate the Internet and the embedded world. Hence, the key metric of success for the SPITFIRE project is the effort required for development of robust, interoperable, and scalable applications in the Internet of Things (IoT).

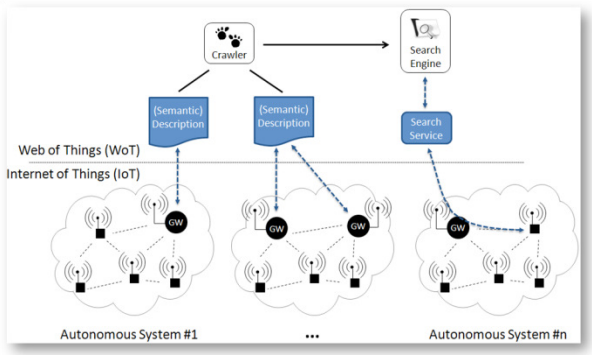

Figure 2: Exemplary use of the SPITFIRE infrastructure

The key objectives of the SPITFIRE project are:

- A minimal set of services to act as a unifying abstraction between the large varieties of applications built and the large variety of available technologies.
- Enabling search, interpretation and transformation of sensor data by giving it explicit semantics.
- Develop the concept of semantic entities as a common abstraction of real-world entities and devise services to process queries on them.

RESEARCH OVERVIEW

By design, the Internet combines two opposing aspects: On the one hand, it works in a highly decentralised manner, allowing easy expansion and replacement; on the other hand, it has become a central unifying force for global information exchange. This paradox yields a constant stream of new problems and solutions, demonstrated, for example, by the range of projects within the FIRE initiative.

In the meantime, embedded computing has also seen a tremendous development, using tiny autonomous devices to build distributed sensing systems. However, limited node capabilities and the tedious task of implementing applications still make it demanding to integrate them into large and powerful distributed networks. Hence, sensor networks and related areas still give rise to a large spectrum of research problems.

Only very recently, serious attempts have been made to integrate these two long-lost relatives. One such activity is the WISEBED project. Building on the experimental facilities developed in WISEBED and consisting of a tightly knit team of experts in semantics, embedded systems, middleware, algorithms, and two cutting-edge SMEs, SPITFIRE endeavours to go one step further towards progress that can be felt by the general public.

WORK PACKAGE STRUCTURE

The project is structured into six RTD Work Packages (WPs). WP one (“Services”) will design, implement, evaluate, and refine the “Service Protocol” (SP). The goal of SP is to provide the same abstraction that the Internet Protocol (IP) provided to create the Internet as a network of networks.

The first goal of the second WP (“Semantics”) is to provide semantic models in order to capture the characteristics of sensors, both static (e.g., type, manufacturer, owner) and dynamic (e.g., real time output). The second goal is to discover parts of the semantic description of sensors semi-automatically. These techniques finally enable querying and inference on real-time streamed sensor data.

Sensor networks possess a unique property separating them from other networks: The individual device has no relevance; the smallest meaningful unit is a sub-network providing certain functionality. The algorithms developed in WP 3 (“Algorithms for Semantic Entities”) will introduce concepts that reflect this situation. The network will then expose these entities via services, so that the outside view of the network installation consists of real-world entities only.

The purpose of the fourth WP ("Experimentation and Applications") is to integrate and evaluate the developed technologies and to provide refinement to the innovative services. A prototypical yet working application, which will be refined in each year to include newly usable technologies, will serve to prove the feasibility of the proposed architecture and technologies.

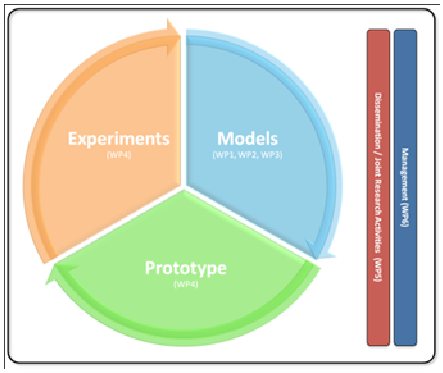

Figure 3: Experimentally-Driven methodology of SPITFIRE

TESTBEDS

The embedded component of the IoT is largely affected by its surrounding real-world environment. Therefore, experimentation on real platforms in realistic environments is a key for successful IoT research. The FIRE initiative and its large-scale experimental facilities provide the unique opportunity to evaluate and drive SPITFIRE research at large scale and in a realistic environment.

WISEBED is the only FIRE facility project that explicitly deals with sensor networks for use in the IoT. It provides a federation of nine Wireless Sensor Network testbeds providing about 2000 nodes, which are physically spread over Europe and federated by a service-based overlay network.

PlanetLab is a testbed for experimentally-driven Future Internet research. As of October 2009, it consists of 1000+ computers hosted at more than 400 different sites worldwide. PlanetLab Europe, the European contribution to PlanetLab, is managed by the project OneLab2.

Finally, iLab.t is an extensive wireless mesh and sensor network infrastructure installed at IBBT office premises. It consists of 400 802.11 wireless mesh network nodes and 200 sensor nodes, supporting a mix of sensor platforms. Its unique features include power control and measurement, and environmental emulation on all nodes.

If you are interested in obtaining up-to-date information about SpitFire visit http://www.spitfire-project.eu ■

InteliSense: Cooperating Objects in the use of eLearning environments

By Panagiotis Germanakos, University of Cyprus, Cyprus

The emotional state of a user at any given moment is susceptible to constant variations due to factors that are either triggered within a specific environment or may as well be extraneous and unrelated to human-computer interactions. As a result, users may experience a wide range of emotions while interacting with eLearning hypermedia content, both positive and negative. In parallel, research on emotions has shown that the association of certain emotions to specific events may have a significant impact on information processing and/or on the consolidation of newly acquired information. The role of emotions is also been discussed specifically in relation to human-computer interaction, with constantly increasing research interest.

Hence, the monitoring and identification of users' emotional state while interacting with hypermedia content could be useful, since it would allow the employment of corresponding personalization techniques in order to reduce the influence of negative emotions. However, identifying human emotions with high accuracy is a complex procedure, whereas altering the emotional state of a user is even more challenging.

Our research is generally focused on web-based learning, information processing, and personalization. Learners' individual differences are represented by a multi-dimensional user model, which also includes the notion of emotional processing. This term refers to a combination of regulatory and affective factors, with anxiety as the core element of this construct. Thus, in terms of identifying learners' emotions, we are mainly interested in the different forms of anxiety (state and trait anxiety, cognitive test anxiety, computer anxiety), and its manifestations as an increase of stress and negative emotional arousal with consequent negative effects on learning performance.

The use of questionnaires in user profiling procedures is a valid form of measurement regarding learners' predispositions towards developing high levels of anxiety in a web-based learning environment, at least as shown by our previous research findings. Nevertheless, the variability of learners' emotional state remains evasive, whereas the notion of a system that dynamically

adapts on individuals' fluctuating emotional state cannot be realized.

Our framework, coined InteliSense, features a real-time method for user- profiling, involving the use of biometric sensors providing measurements of skin conductance, heart rate, and blood volume pulse, in an effort to capture users' state-like emotional characteristic in a more accurate way. The IntelliSense architecture is depicted in Figure 4.

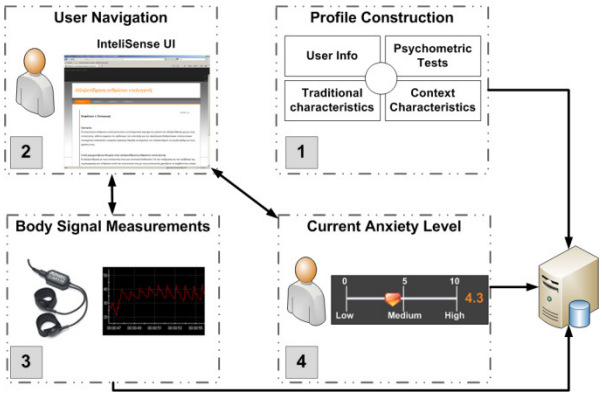

Figure 4: InteliSense Architecture

The design of the study was a mixed between-within subjects design. The number of participants was 10 (6 female and 4 male), and they were all students at the Department of Computer Science of the University of Cyprus. Two identical experiments with the same participants took place in a computer lab, with a 5 day interval, providing a total of 20 samples.

Each participant was asked to log into the system and to firstly fill in the psychometric questionnaires of anxiety and emotional regulation. Thereupon, the sensors were placed on the hand of the participant, who was subsequently asked to navigate in an e-learning course about HCI design.

The participants were aware that an exam would follow, in order for a certain degree of emotional arousal to be evoked; the sensors provided real-time measurements throughout the learning procedure, while there were not any time constraints imposed on learners. Additionally, users were able to report their self-perceived levels of anxiety every five minutes.

The participants, after the completion of the learning course, took an exam on the subject that they were taught, with their scores indicating their academic performance in terms of efficient information processing and consolidation of the newly acquired information.

According to the findings of the first experiment, there is a statistically significant positive correlation between learners' mean heart rate and their levels of trait anxiety (STAI questionnaire): Pearson's r=0.668, p=0.035. However, neither high levels of trait anxiety nor elevated heart rate were correlated with academic performance. There was also a negative correlation between self-reported levels of anxiety and emotional regulation ability, though not reaching statistical significance (r=-0.565, p=0.089). As it concerns skin conductance and blood volume pulse, there was only a marginal variability without any correlations to academic performance, trait anxiety, or emotional regulation.

The second experiment confirmed the relation between heart rate and trait anxiety (r=0.640, p=0.046). Additionally, heart rate was positively correlated with learners' self reported anxiety (r=0.739, p=0.015), possibly due to the fact that users had become more familiar with the environment and the self-report procedure. The negative correlation between self-reported levels of anxiety and emotional regulation was also observed, though still not reaching statistical significance (r=0.610, p=0.61). Skin conductivity and blood volume pulse measurements did not provide any useful data in terms of variability or correlations. Academic performance was also unrelated to any other variable. As expected, learners' scores in the second experiment were higher, due to their previous learning experience (the mean rose from 71% to 85%, paired samples t=3.157, p=0.012). All the other variables were basically unaffected. Thus, it seems that previous experience certainly has a positive effect on information consolidation, but not on anxiety levels that seem to be essentially related to trait-like characteristics.

The aforementioned experiments provided initial results on the use of bio-feedback sensors within the context of educational hypermedia environments. Even though this preliminary research does not allow the formation of definitive conclusions, the main finding is that an e-learning course does not evoke significant levels of emotional arousal that would manifest in the form of physiological responses. On the contrary, pre-existing higher levels of trait anxiety are related to an increase in heart rate, while both seem independent from academic performance; this may be explained by the fact that a certain amount of anxiety is rather beneficial in learning as a motivational factor.

The next step of our research regarding biofeedback and learning will be the design of a larger scale study, with a far more elaborate methodology, involving positive and negative stimuli in an educational environment. This will allow a bet-

ter identification of learners' emotional state in web-based learning environments and the proper manipulation of a learning course for its dynamic reconstruction and adaptation in order to positively affect learners' levels of emotional arousal and valence.

If you are interested in obtaining up-to-date information about IntelliSense visit http://www4.cs.ucy.ac.cy/adaptiveweb/ ■

Announcements

HDMS2010 – 9th Hellenic Data Management Symposium

June 30 - July 3, 2010 Ayia Napa, Cyprus

ⓘ http://www.cs.ucy.ac.cy/~hdms2010/en

Important dates:
Paper submission deadline: April 15, 2010
Notification of acceptance: May 20, 2010
Poster submission deadline: June 2, 2010
Camera ready submission: June 3, 2010

In Bio We Trust 2010

Intl. Workshop on Bio-Inspired Trust Management for Information Systems

December 1-3, 2010, Boston, USA

ⓘ http://inbiowetrust.org/

Important dates:
Paper Submission: 1 September 2010
Notification of Acceptance: 24 September 2010
Final Version Due: 10 October 2010

CAT2010 – European Context Awareness & Trust 2010

4th Workshop on Combining Context with Trust, Security, and Privacy

co-located with CENTRIC 2010

23-24 August 2010, Nice, France

ⓘ http://www.cat10.uni.lu/

Important dates:
Paper submission: April 5, 2010
Notification: 25 April 2010
Registration: 15 May 2010
Final version due: 22 May 2010

RTSS 2010 - The 31st IEEE Real-Time Systems Symposium

30 Nov – 3 Dec, 2010, San Diego, CA, USA

ⓘ http://www.rtss.org/

Important dates:
Paper submissions: 16 May 2010
Notification of acceptance: 8 August 2010
Camera-ready deadline: 12 September 2010

EuroPKI2010 – European Workshop on Public Key Services Applications and Infrastructures

23-24 September 2010, Athens, Greece

ⓘ http://www.ds.unipi.gr/europki10/

Important dates:
Submission deadline: 28 May 2010
Notification to authors: 2 July 2010
Camera-ready version: 18 July 2010

CONET Summer School 2010

Aug 15 – 21, 2010, Schloss Dagstuhl, Germany

ⓘ http://www.cooperating-objects.eu/school/2010/home

Important dates:
Application deadline: 31 May 2010
Admittance notification: 15 June 2010

Special Issue Call for Papers – Query Processing in Sensor Networks

Springer 2010

ⓘ http://www.springer.com

Important dates:
Paper submission deadline: 15 May2010
Notification to authors: 1 August 2010
Online Publication: Beginning of 2011

Workshop on Real-World Wireless Sensor Networks (REALWSN2010)

December 16 – 17, 2010, Colombo, Sri Lanka

ⓘ **http://www.ucsc.cmb.ac.lk/realwsn10/**

Important dates:
Paper Submission Deadline: mid-July, 2010
Acceptance Notification: Sep, 2010
Final Paper Submission Deadline: Oct, 2010

Latest News

- The 2nd CONET Summer School: "Act-Control-Move: Beyond Networked Sensors", will take place in Schloss, Dagstuhl, Germany, on 15th – 21st of August, 2010.