NATIONAL AND KAPODISTRIAN UNIVERSITY OF ATHENS
DEPARTMENT OF INFORMATICS & TELECOMMUNICATIONS

ABSTRACTS OF DOCTORAL DISSERTATIONS

Athens 2013

Volume 8

NATIONAL AND KAPODISTRIAN UNIVERSITY OF ATHENS
DEPARTMENT OF INFORMATICS & TELECOMMUNICATIONS

ABSTRACTS OF DOCTORAL DISSERTATIONS

ISSN: 1791-7948

Volume 8

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications
Panepistimiopolis, 15784 Athens, Greece

PREFACE

This volume contains the extended abstracts of the Doctoral Dissertations conducted in the Department of Informatics and Telecommunications, National and Kapodistrian University of Athens and completed in the time period 12/2011 to 12/2012.

The goal of this volume is to demonstrate the breadth and quality of the original research performed by our Ph.D. students and to facilitate the dissemination of their innovative research results. We are happy to present the 8th collection of this kind and expect this initiative to continue in the years to come. The submission of an extended abstract in English is required by all graduating Ph.D. students of the Department.

We would like to thank our Ph.D. graduates who contributed to this volume and hope that this has been a positive experience for them. Finally, we would like to thank PhD candidate Nikos Bogdos for his help in putting together this publication. The painting in the cover is called "*Απάγκιο*" by artist *Διονύσης Καμπόλης*.

The Committee of Research and Development

A. Eleftheriadis
M. Koubarakis
E. Manolakos (Chair)
T. Theoharis

Athens, June 2013

Table of Contents

Advanced Location Prediction Techniques in Mobile Computing

Theodoros Anagnostopoulos*

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications,
thanag@di.uoa.gr

Abstract— Context-awareness is viewed as one of the most important aspects in the emerging pervasive computing paradigm. Mobile context-aware applications are required to sense and react to changing environment conditions. Such applications, usually, need to recognize, classify and predict context in order to act efficiently, beforehand, for the benefit of the user. Firstly, we propose an efficient spatial context classifier and a short-term predictor for the future location of a mobile user in cellular networks. Secondly, we propose a novel adaptive mobility prediction algorithm, which deals with location context representation and trajectory prediction of moving users. Thirdly, we propose a short-memory adaptive location predictor that realizes mobility prediction in the absence of extensive historical mobility information. Fourthly, we assume the existence of a pattern base and try to compare the movement pattern of a certain user with stored information in order to predict future locations. Our findings, compared with other schemes, are very promising for the location prediction problem and the adoption of proactive context-aware applications and services.

1. Introduction

In order to render mobile context-aware applications intelligent enough to support users everywhere and anytime, information on the present *context* of the users has to be captured and processed accordingly. A well-known definition of context is the following: "*context is any information that can be used to characterize the situation of an entity. An entity is a person, place or object that is considered relevant to the integration between a user and an application, including the user and the application themselves"* [1]. Context refers to the current values of specific ingredients that represent the activity of an entity, situation and environmental state e.g., location, time, walking, attendance of a meeting, driving a car, traveling. One of the more intuitive capabilities of the mobile applications is their *proactivity*. The prediction of the user's mobility behavior enables a new class of location-aware applications to be developed along with the improvement of the existing location-based services [2]. Even at the network level, mobility prediction assists in critical operations like handoff management, resource allocation, and quality of service provisioning.

Two classes of location (path) prediction schemes can be found in the current, mobile computing, literature. The first class includes schemes based on extensive historical data of the user movement. Such a scheme can be characterized as stateful. In a stateful scheme the prediction process relies on the *matching* of established (historical) movement patterns with the user movement experienced up to moment of prediction in order to estimate the future location of the user. Pattern-based and data mining approaches as well as machine learning techniques (e.g., learning automata) can be classified in this category. Contrary to the stateful scheme, a stateless model does not take into account extensive historical movement information for the prediction process. Instead, it uses of *a short sliding window* of historical movement information. Such scheme applies statistical techniques (e.g., extrapolation) on the recent movement information (window) in order to predict the future user location. The stateless scheme does not assume regularity in the user movement, as opposed to the stateful scheme, but proceeds with predictions based only on short-term spatiotemporal knowledge. Moreover, the prediction process comes along with adaptation techniques for certain parameters of the statistical techniques to fully cover the potential randomness of the user movement. One could also define hybrid schemes based on the stateful and stateless mechanisms that are invoked and collectively taken into account for joint decisions (e.g., weighted decisions).

---

*Dissertation Advisor: Stathes Hadjiefthymiades, Assistant Professor

2. Dissertation Summary

Firstly, we propose an efficient spatial context classifier and a short-term predictor for the future location of a mobile user in cellular networks [3], [4] and [5]. The process of location / context prediction is based on pattern classification in the sense that the user's mobility patterns are collected, processed, classified and used for predicting future mobility behavior. Typically, Machine Learning (ML) is primarily concerned with the design and development of algorithms for pattern classification and prediction. Such algorithms train a system (classifier) to classify observations in order to predict (predictor) unknown situations based on a history of patterns. Hence, if a mobile application is capable of tracing, learning and classifying the mobility patterns of a user, (i.e., the observed trajectories) then, it can predict his/her future mobility behavior. In this approach, we adopt pattern classification in order to predict the future location of a mobile user based on the spatial context i.e.,

- the current position and direction of the user,
- the history of the trajectories followed by the user, and
- the information on the user's surroundings (neighbouring network cells).

Each observed trajectory (pattern) is assigned to a specific class (a symbolic location) from a fixed set of classes. Our objective is to predict the symbolic location of a trajectory based on a collected set of popular trajectories. We propose a Location Predictor (LP), which is based on a Trajectory Classifier (TC), in order to address the location prediction problem. We design, implement and evaluate different versions of the proposed LP. Each variant exploits differently the derived knowledge on the mobility behavior of the user. The proposed, classifier-based, LP relies on short-term historical data of the user movement. Prediction is accomplished through the matching of established (historical) movement patterns with the short-term history of the user movement (spatiotemporal knowledge) seen up to the moment of prediction. In our case, a user mobility pattern is a series of symbolic locations. The size of the mobility pattern is directly linked to prediction accuracy. Limiting the size of the user's short-term history leads to loss of information, which results in false predictions and, on the other hand, redundant or excess historical knowledge results to noisy information and biased predictions. The challenge is to introduce an approach based on (a) short term spatio–temporal knowledge on user mobility behavior, and, (b) low computational cost on the decision predictor based on short-term historical movement data. The outcome of the proposed LP is provided through efficient and low complexity ML algorithms. Our findings show a clear advantage on the adoption of ML algorithms for location prediction as we obtain increased prediction accuracy (probability of correctness at a reasonable computing cost. In addition, we perform a comparative analysis of the proposed LP with popular prediction algorithms reported in the literature.

Secondly, we propose a novel adaptive mobility prediction algorithm, which deals with location context representation and trajectory prediction of moving users [6], [7]. In this context ML provides algorithms for learning a system to:

- clusters the user movements,
- identifies changes in the user movements
- adapts its knowledge structure to such changes, and,
- predicts the future user location.

In addition, ML provides solutions suitable for location prediction. Context-aware applications have a set of pivotal requirements e.g., flexibility and adaptation. These requirements would strongly benefit if the learning and prediction processes could be performed in real time. We argue that the most appropriate solutions for location prediction are offline and online clustering and classification. Offline clustering is performed through the Offline *k*Means algorithm. Online clustering is accomplished through the Online *k*Means and Adaptive Resonance Theory (ART). Offline learners typically perform complete model building, which can be very costly, if the amount of samples rises. On the other hand, online learning algorithms are able to (i) detect changes in patterns and (ii) update only parts of the model. The later provides adaptation of the model. Both forms of algorithms extract a subset of patterns, i.e., clusters in a knowledge base, from an initial dataset i.e., a database of user trajectories.

Online learning is more suited for the task of prediction through classification than offline learning. That is because, in real life, user movement data often needs to be processed in an online manner, each time after a new portion of the data arrives. This is caused by the fact that such data is organized in the form of a data stream (e.g., a sequence of time-stamped visited locations).

Classification involves the matching of an unseen pattern with existing clusters in the knowledge base. We rely on a

Hausdorff-like distance for matching unseen trajectories to clusters. Such metric applies to convex patterns and is considered ideal for user trajectories. Hence, location prediction boils down to location classification w.r.t. Hausdorff-like distance. In this approach, we introduce two training methods for training our algorithm: (i) the "nearly" *zero-knowledge* and (ii) the *supervised* method. In the first method, our algorithm is incrementally trained starting with a little knowledge on the user mobility behavior. In the second method, sets of known trajectories are fed to our algorithm. In addition, we introduce two learning methods for the proposed algorithm regarding the success of location prediction: (i) the *non-reinforcement learning* (nRL) and (ii) the *reinforcement learning* (RL). In the nRL method, a misclassified trajectory is no further used in the model-training phase. Hence, the algorithm is no longer aware of unsuccessful predictions. In the RL method, a misclassified trajectory is introduced into the knowledge base, thus, updating appropriately the model.

We evaluate the performance of our algorithm against real users' movement. We establish important metrics for the performance assessment process w.r.t. (i) low system-requirements, i.e., storage capacity, and (ii) the required effort for model building, i.e., processing power. We assess the prediction accuracy of our algorithm, i.e., the precision of location predictions, along with the required size of the derived knowledge base. This size indicates the portion of the produced clusters out of the volume of the training patterns. Surely, we need to keep storage capacity as low as possible while maintaining good prediction accuracy. We also study the capability of our algorithm in adapting the derived model to unseen patterns. *Adaptivity* indicates the capability of the proposed algorithm to detect and update appropriately a specific part of the trained model. The algorithm should rapidly detect changes in the behavior of the mobile user and adapt accordingly through model updates. This is, however, achieved often at the expense of the prediction accuracy.

Thirdly, we propose a short-memory adaptive location predictor that realizes mobility prediction in the absence of extensive historical mobility information [8]. All the reported drawbacks led us to introduce a stateless scheme for movement prediction without considering user movement profiles and pattern-based techniques. The challenge is to introduce an approach based on

- short term spatiotemporal knowledge on user mobility behavior,
- detection and quick adaptation to changes of the user mobility behavior, and,
- very low requirements on storing and manipulating short-term historical movement data.

We present an approach for mobility prediction exploiting only the current spatiotemporal knowledge on the user's mobility. Our aim is to address the problem of mobility prediction in the absence of large historical mobility information for the considered user(s). The proposed mobility prediction approach can be applied whenever the mobility behavior of a user is not known a-priori. For instance, consider tourists that visit a town or a freshman in a University campus. However, due to the absence of past mobility patterns, mobility prediction has also to regularly adapt to unseen user movements. The proposed scheme can, quite accurately, predict the traveling trajectory. This is achieved by using and dynamically adjusting only local current spatiotemporal knowledge on the user's mobility behavior having very short relevant historical information. The main objective of our approach is to circumvent the difficulties that arise in predicting the user's future location when extensive knowledge on the history of user's traveling patterns is not available or the user behaves quite randomly. Another challenging point is the determination of the capacity of the relevant historical information used for mobility prediction. The capacity of the relevant historical information varies among mobile users since they exhibit different movement behaviors. In addition, more interestingly, the capacity of the relevant historical information is time-varying for a single user itself. This is because the mobile user can abruptly change his/her movement behavior, which has to be reflected in the capacity of the relevant historical information in order for the mobility prediction system to adjust to such unseen, new changes.

In this approach, we discuss the design and implementation of an *adaptive*, *short-memory* Location Predictor (LP). The LP does not require an extensive knowledge base of past user movements. Instead, the LP estimates the future location of the mobile user based on short-term movement information. That is, the LP exploits only local spatiotemporal knowledge of the user movement trying to predict future locations near in time. Evidently, this requires that the LP dynamically adjusts its model for future decisions on location estimations. Since there are no stored historical patterns or other representative knowledge on the user's mobility behavior, the LP has also to be efficient even in the case of random (spontaneous) movement changes, e.g., sudden turns. On the other hand, the LP should reach a stable condition (model) whenever the user does not change his/her mobility behavior for long. Therefore, the LP should be able to rapidly detect changes in mobility behavior and adapt its model to the *current* behavior, thus, providing for fast adaptation of the model.

The criteria that we establish for the performance assessment of the adaptive LP take into account the system requirements (storage capacity) and the computational effort for prediction and fast adaptation. Besides the prediction accuracy, i.e., the precision of location predictions, we are interested in the size / capacity of the information that LP has to process for making predictions. We should stress that in our case, the consumed memory is extremely low if compared with other adaptive techniques for location prediction. Lastly, our objective is to assess the adaptive behavior of the LP that is its capability to rapidly detect changes in the movement of the mobile user and react accordingly.

Fourthly, we assume the existence of a pattern base and try to compare the movement pattern of a certain user with stored information in order to predict future locations [9]. Movement traces are always quite noisy and, therefore, undermine the performance of mobile computing mechanisms. A location-aware application requires online motion pattern learning and adaptation. On-line learning is more suited for the task of prediction. That is because, in real life, user movement data often needs to be processed in an on-line manner, each time new location information becomes available. However, most location-aware applications require training on the movement space of a moving object (supervised learning). Once the training is completed, the application is ready for prediction, and no additional modification is permitted. This case is acceptable provided that the motion patterns of the moving object correspond to well-defined boundaries. Nevertheless, in many realistic situations, the full spectrum of mobility patterns is not available during the training phase. The training process restricts the capability of the system to adapt its knowledge base to unseen patterns. The proposed system delays decision making (i) for trajectory classification and knowledge base adaptation, and for (ii) unseen pattern insertion in the knowledge base. Such delay makes the system 'confident' to decision making for classification, prediction, learning and adaptation. This system does not require a training process, learns, and revises the knowledge base in an on-line manner. Potential applications, as those mentioned above, which require efficient location prediction could adopt such delay tolerant decision making mechanism. Such functionality might enhance the prediction score for such location-aware applications by being more intelligent in decision making for location prediction and on-line motion pattern learning.

Our system manages the entire life-cycle of an on-line motion pattern learning, trajectory (sequential) classification, and adaptive movement prediction. It works in a relatively stateless way and tries to keep the underlying knowledge base K with the lowest possible total spatial variance if the stored patterns. The proposed system demonstrates some temporal tolerance (i.e., patience - delay) in classifying a trajectory to an existing (stored) pattern or abandon the similarity matching process and treat the assessed trajectory as completely new. It is worth noting that the proposed system is not forced to proceed with prediction once not being able to classify an unseen pattern. The system delays decision making for being more certain on prediction and learning decisions. We also introduce a trajectory similarity metric, which provides a holistic view of how close a stored pattern and an observed trajectory are. Should the classifier conclude, spatial variance reduction process decides (based on spatial statistics) on the exact pattern to be kept in the underlying knowledge base, K. If the new (assessed) trajectory demonstrates smaller spatial variance than the best matching pattern then the latter is substituted by the former. Otherwise K remains unmodified. In this approach we propose,

- a closed-loop system for sequential trajectory classification and k-step ahead location prediction;
- a dynamic programming approach (optimal stopping rules) for sequential trajectory classification;
- a trajectory similarity metric which supports sequential classification;
- an incremental and adaptive motion pattern learning scheme. The system does not require training phase and incrementally refines the knowledge base;
- a total spatial variance reduction process for the stored patterns in K, thus, avoiding classification decisions based on noisy patterns.

3. Results and Discussion

In the third approach, we adopt a local regression model based on kernel weighted functions in order to determine the future user location through extrapolation. The LP exploits a fuzzy controller in order to decide on the appropriate size of the memory of the local regression model that minimizes the prediction error. This means that, the fuzzy controller adjusts the memory size of the regression model based on the current user mobility behavior. Any detection of change on the mobility behavior is treated through a fuzzy control signal that adjusts the current memory size (history window) of the regression process. The overall model of the proposed LP is illustrated in Figure 1. Specifically, the regression model accumulates the last $m \geq 1$ locations and constructs a statistic regression function f at time instance t. Then, the future

location x, which is predicted for the next $l > 1$ time instances (i.e., at time index $t + l$), is an extrapolation point based on the mean value of the user velocity and direction of the last m locations. The prediction error e at time $t + l$ is calculated whenever the predicted location x is not the actual location at time $t + l$. Specifically the prediction error e is the spatial distance between the predicted location and the actual location at time $t + l$. The fuzzy controller is fed with the prediction error e and adjusts the length of the m last locations in the LP.

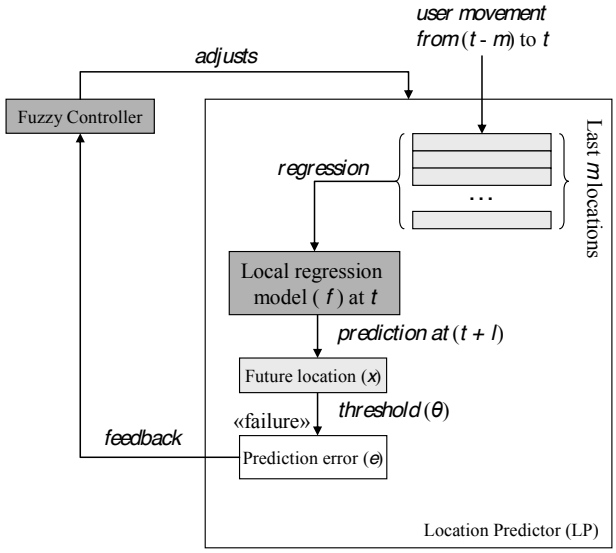

Figure 1. The adaptive short-memory Location Predictor (LP) model and fuzzy controller

3.1 Local Linear Regression Model

We now elaborate on the local regression component of the LP as illustrated in Figure 1. The estimating local regression function (model) $f(X)$ over a domain $X \subset R^p$ fits a function f separately at each target point x_0. This is achieved by using only those observations (points) close to x_0 to fit the model in a way that the resulting estimated function $\hat{f}$ is *smooth* (has derivatives of all orders) in the domain R^p. Local linear regression becomes less useful for $p > 2$. In our case $p = 1$ as discussed below. This localization is based on a weighted function or kernel $K_\lambda(x_0, x)$. The kernel assigns weights to x_i based on its distance from x_0. The typical smoothing parameter λ indicates the width of the neighborhood. A larger λ implies lower variance. In such methods the model is structured by a set of observations and requires no training, i.e., all the work is performed at the location estimation time. The locally weighted regression solves a separate weighted least squares problem at each target point x_0 for a set of m observations $(x_i, f(x_i) = y_i)$. It should be noted that although we fit an entire linear model to the data in the region of m observations, we only use it to evaluate the fit at x_0. The extrapolated $\hat{f}(x_0)$ is generated by the m previous observations $f(x_i)$, $i = 1, \ldots, m$. The kernel function K_λ might be, for instance, the Epanechnikov quadratic kernel, defined as:

$$K_\lambda(x_0, x_i) = D\left(\frac{|x - x_0|}{\lambda}\right), D(u) = \begin{cases} \frac{3}{4}\left(1 - u^2\right), |u| \le 1 \\ 0, \text{elsewhere} \end{cases}$$

According to such function $K_\lambda(x_0, x_i)$, as we "move" the point x_0 from left to right, neighboring points x_i enter the history window initially with weight zero, and then their contribution slowly increases. The regression model is used for interpolation among the location information of the user trajectory.

3.2 Location and Trajectory Information Model

We adopt a location and trajectory information model, in which the location of the user at time t is represented by a 2D vector $(x(t), y(t))$ of longitude $x(t)$ and latitude $y(t)$. The exact position information can be easily estimated. In our

case we use position information from GPS receivers embedded in the mobile user's handheld terminal. Based on the discussed representation, the history of the user movement up to time index n is represented as a sequence –trajectory- of n points $k(t) = (x(t), y(t))$, $t = 1 \ldots n$, that is $\mathbf{k} = [k(t)] = [(x(1), y(1)), \ldots, (x(n), y(n))]$. The sequence $\mathbf{k}$ for $t = 1 \ldots n$ is the user trajectory for the last n points starting from the point at time index $t = 1$. Each point $k(t)$ is sampled with frequency q, e.g., in our traces, the GPS locations are samples with $q = 1$Hz. A window $\mathbf{w}$ is a sub-sequence of the m first points from $\mathbf{k}$, that is $\mathbf{w} = [k(t)]$, $t = 1, \ldots, m$, $m < n$. The estimation-prediction of the future position of a mobile user is based on the window $\mathbf{w}$ of m points. In other words, the predictor is fed with a trajectory of m points in order to predict future points. The discussed predictor is referred to as an $m+l$ predictor model [12]. That is the predictor estimates the future point indexed as $m+l$, i.e., predicts the point $k(m+l) = (x(m+l), y(m+l))$, with $l > 0$ from a trajectory (history window) of m points (see also Figure 2). For instance, we set $l = 1$ for the prediction of the future point at the next time instance. The window of length m, the number of the l time indices for predicting the $k(l)$ point and the frequency of sampling q are the basic parameters in our LP, as illustrated in Figure 2.

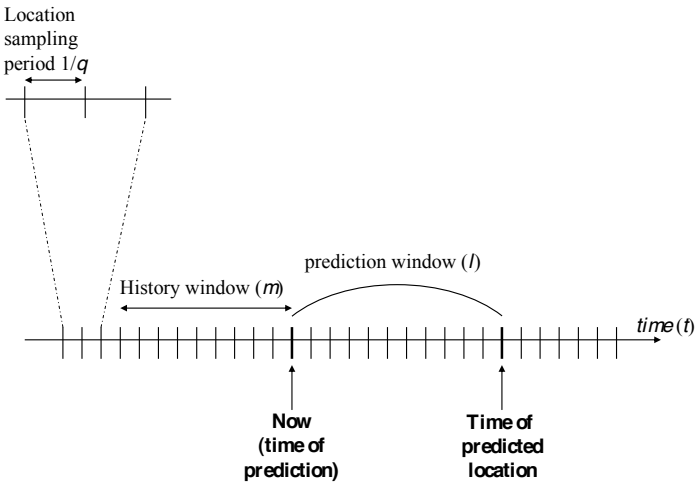

Figure 2. The basic parameters of the $m+l$ prediction.

3.3 Location Predictor Design

We discuss the adoption of the local linear regression in location prediction. We can distinguish the following cases for the user movement for a very small duration $\Delta t > 0$:

- Case-1: the user moves more steadily in the longitude $x(t)$ direction than in the latitude direction. This means that, it may hold true that the difference $|x(t+\Delta t) - x(t)| << |y(t+\Delta t) - y(t)|$ in a Δt.
- Case-2: the user moves more steadily in the latitude $y(t)$ direction than in the longitude direction, that is, $|y(t+\Delta t) - y(t)| << |x(t+\Delta t) - x(t)|$, and,
- Case-3: the user moves in such a way that $|x(t+\Delta t) - x(t)| = |y(t+\Delta t) - y(t)|$ in Δt.

Consider that for a time frame m, we calculate the mean differences for both longitude and latitude directions. That is, consider the vectors:

$$\Delta x(m) = \frac{1}{m-1}\sum_{i=2}^{m}\left(x(i) - x(i-1)\right) = \frac{1}{m-1}\left(x(m) - x(1)\right)$$
$$\Delta y(m) = \frac{1}{m-1}\sum_{i=2}^{m}\left(y(i) - y(i-1)\right) = \frac{1}{m-1}\left(y(m) - y(1)\right) \quad (1)$$

The predictor can estimate the point after l time instances when the following cases holds true:

- If $|\Delta x(m)| < |\Delta y(m)|$ then $k(m+l) = (x(m+l), y(m+l))$ with the extrapolation $x(m+l) = x(m) + l \cdot \Delta x(m)$ and the local regression estimate $y(m+l) = \hat{f}\left(x(m+l)\right)$. In this case, the $y(m+l)$ is generated by the window w and the local linear estimation of $\hat{f}(x(m+l))$ at $x(m+l)$ having local regression in the 1-dimension domain ($p = 1$);

- If $|\Delta x(m)| > |\Delta y(m)|$ then $k(m+l) = (x(m+l), y(m+l))$ with the extrapolation $y(m+l) = y(m) + l \cdot \Delta y(m)$ and the local regression estimate $x(m+l) = \hat{f}(y(m+l))$. In this case, the $x(m+l)$ is generated by the window w and the local linear estimation of $\hat{f}(y(m+l))$ at $y(m+l)$.

It is worth noting to study the case where $|\Delta x(m)| = |\Delta y(m)|$. Specifically, the time window m used for the calculation of $|\Delta x(m)|$ or $|\Delta y(m)|$ may contain smaller spatial transitions not captured by the absolute values of relocation in the x or y axis. A small time window m renders the scheme more sensitive to such (small) transitions while a large time window may 'absorb' movements and derive predictions that are not representative of the actual movement (not necessarily wrong predictions). Surely, due to the above shortcoming, if $|\Delta x(m)| = |\Delta y(m)| = 0$, the predictor fails to distinguish between the totally static case and a recurring (cycle) path. Hence, in this case, the future point $k(m+l)$ is $(x(m+l), y(m+l))$ with the extrapolation either $y(m+l) = y(m) + l \cdot \Delta y(m)$ (and $x(m+l) = \hat{f}(y(m+l))$) or $x(m+l) = x(m) + l \cdot \Delta x(m)$ (and $y(m+l) = \hat{f}(x(m+l))$).

In order to judge whether the predictor is successful in predicting the future point $k(m+l) = (x(m+l), y(m+l))$, we introduce an error threshold $\theta > 0$. Such threshold is the minimum distance of the actual point $z(m+l)$ at time instance $m+l$ and the predicted point $k(m+l)$. The $z(m+l)$ point is the actual (observed) location of the mobile user at time instance $m+l$. Hence, a prediction is considered successful if $\|z(m+l) - k(m+l)\| \leq \theta$. Otherwise, the prediction fails.

If we use the $m+l$ predictor for N times (for large N) then we can estimate the *probability of a successful* $m+l$ *prediction* $P(m,l)$. $P(m,l)$ is the fraction of the correctly predicted positions, i.e., those predicted points $k(m+l)$ that fall within a circle of radius θ, out of N predictor invocations, i.e.,

$$P(m,l) = \frac{\left|\left\{k(m+l) : \|z(m+l) - k(m+l)\| \leq \theta\right\}\right|}{N}$$

where $z(m+l)$ is the actual (observed) location of the mobile user at time instance $m+l$ and $|U|$ is the cardinality of the set U.

3.3 Fuzzy Control for Adaptive Location Prediction

Fuzzy control is a direct nonlinear mapping between inputs (e.g., location prediction error, change of the prediction error) and outputs (e.g., change in the memory size m). Fuzzy control provides formal techniques to represent and implement human experts' heuristic knowledge for controlling a system. Such techniques rely on certain *if-then* rules instead of relying on mathematical modeling of the system. Fuzzy logic controllers are used extensively, especially, in the form of the proportional, integral and derivative (PID), PD, and PI controllers. In this paper we consider a fuzzy, proportional-integral controller PIm for adjusting the window length m in location prediction. The control signal is based on a set of fuzzy rules that generate the required output.

Figure 2 depicts the high level structure of the proposed closed-loop, real-time system for adjusting the history window length m. To achieve the maximum value of the probability of successful predictions $P(m(t), l)$ at time t, which is denoted by the reference model $a(t) = 1$ at time t (the upper bound in the probability of successful predictions), the PIm controller computes the required value of m, denoted as $u(t)$, based on the error $e(t)$, change of error $\Delta e(t)$ and the past decision of the controller on changing m, $\Delta m(t)$. Specifically, the error $e(t)$ at time instance t is:

$$e(t) = a(t) - P(m(t), l) \quad (2)$$

In location prediction, we aim at maximizing the $P(m(t), l)$ at any given t that is, $a(t) = 1$ at each prediction. The error change at time t is:

$$\Delta e(t) = e(t) - e(t-1) \quad (3)$$

and the change of the controller's past decision on m is:

$$\Delta m(t) = m(t) - m(t-1) \quad (4)$$

The exact value of u which is the window length that minimizes the error cannot be directly estimated. Instead, the PIm adjusts $m(t)$ in order to minimize $e(t)$ for each t. In addition, it should be noted that the m value that minimizes $e(t)$ (denoted by m^*) is not constant since the mobile user can spontaneously change its moving behavior. In such case, the controller has to adapt to such change and adjust m so as to minimize $e(t)$ until the next change of the user's behavior.

We define the required change in the value of m, $\Delta u(t) \in [-1, 1]$, which denotes the signed portion of changing the value of $m(t)$ -increase, decrease or leave constant- and results to the value of $m(t+1)$ at the next time step, t+1. In other words we obtain the following adaptation rule:

$$m(t+1) = m(t) + \mu(\Delta u(t) m(t)) \tag{5}$$

with $\mu > 0$ an adaptation (learning) rate. The proposed fuzzy controller calculates the value of the control signal $\Delta u(t) = u(t) - u(t-1)$. That is, the decision on the next value of m (that minimizes the error) is based on the past decision on m along with the observed error induced by the LP. The change in m, Δm, and the change in error, Δe, are the basic components in our controller which, along with the value of error, determine the direction (positive, zero, negative) of Δu in the decision on the next value of m, as will be described subsequently.

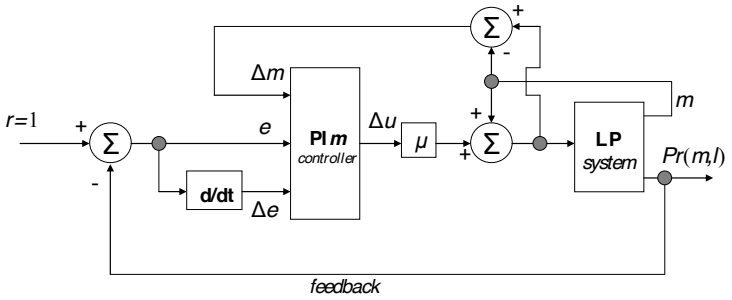

Figure 2. The control flow and feedback of the fuzzy controller PI m to the location predictor LP

The recursive equation of a discrete PID controller is:

$$u(t) = G_P e(t) + G_I \sum_{\tau=0}^{t} e(\tau) + G_D \Delta e(t)$$

where G_P, G_I, G_D are the corresponding constant proportional, integral and derivative gains. The definition of the difference between two consecutive control outputs $\Delta u(t)$ is:

$$\Delta u(t) = u(t) - u(t-1) = G_I e(t) + G_P \Delta e(t) + G_D \Delta^2 e(t) \tag{6}$$

where $\Delta^2 e(t) = \Delta e(t) - \Delta e(t-1)$. In our case, the PI$m$ has three inputs –$e(t)$, $\Delta e(t)$, and $\Delta m(t)$- and one output –the current change in m, Δu(t). Our PIm controller has the following form:

$$u(t) = G_P e(t) + G_I \sum_{\tau=0}^{t} e(\tau) + G_D(\Delta e(t) + \Delta m(t)) \tag{7}$$

The iterative form of $u(t)$ is obtained by taking the derivative of both sides of the Equation (8) having also $\Delta m(t) = m(t) - m(t-1)$, thus, in our case,

$$\Delta u(t) = G_P \Delta e(t) + G_I e(t) + G_D(\Delta^2 e(t) + \Delta^2 m(t)) \tag{8}$$

and then, $u(t) = u(t-1) + G_I e(t) + G_P \Delta e(t) + G_D(\Delta^2 e(t) + \Delta^2 m(t))$ with $\Delta^2 m(t) = \Delta m(t) - \Delta m(t-1)$. Based on Equation (8), the inputs of our controller (e, Δe and Δm) and the output ($\Delta u(t)$) turn the controller into a PI controller. Figure 2 depicts the architecture of the proposed controller.

3.5 Fuzzy Controller

We describe the basic fuzzy control system for inferring the required control signal based on fuzzy inference rules. A fuzzy controller is a fuzzy logic system with n inputs and k outputs. In our case, the PIm controller is a Multiple Input Single Output (MISO) fuzzy logic system with $n = 3$ and $k = 1$, such that the input at time t is $\mathbf{p}(t) = [e(t), \Delta e(t), \Delta m(t)]$ and output $\mathbf{q}(t) = [\Delta u(t)]$. The fuzzy system consists of the processes:

- fuzzification,

- fuzzy inference process, and
- defuzzification.

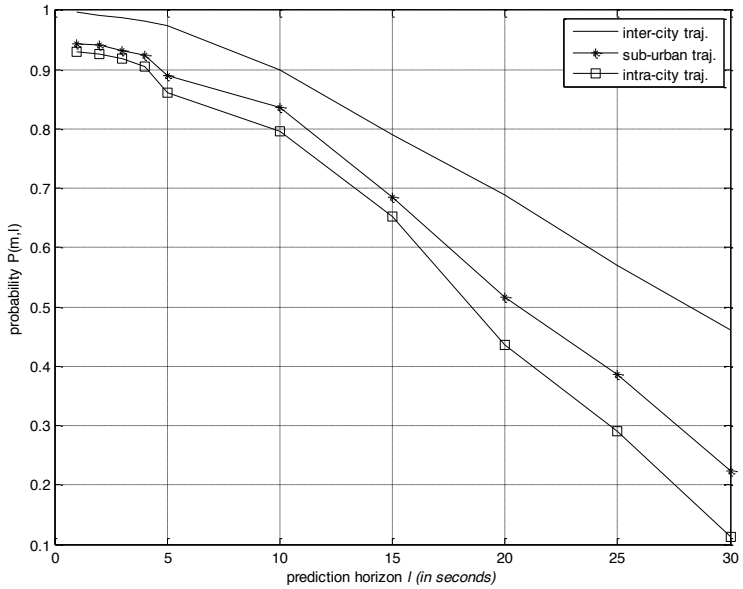

Figure 3. Probability of successful prediction for m=10 s and different l values

3.6 Performance of the Fuzzy Controlled Location Predictor

We assess the performance of the proposed LP and the capability of the proposed PIm controller to control the LP targeting to adaptively minimize the prediction error. At first, we examine the performance of the LP by experimenting with real GPS traces. We examine the probability of correct predictions $P(m, l)$ and determine the best value of the window length m^* that minimizes the prediction error. The value m^* cannot be determined beforehand (e.g., analytically) due to the inherent randomness in the mobility behavior of humans. For this reason, we require that the proposed controller adjusts the value of m in order to converge to m^*. Once the user changes its mobility behavior, the PIm has to re-adjust m. With the aim of evaluating the PIm controller, we have determined the best values for m for which the LP assumes minimum prediction error. Hence, we examine whether the PIm controller converges to such values. Moreover, we examine the adaptive behavior of the controller. That is, its capability to detect changes in the mobility behavior of the user and to re-adjust its decisions to new m^* values.

3.7 Experimental Movement Trajectories - Traces

We examine the behavior of the proposed LP with real movement traces of mobile users in German, Italy, France, and Denmark. In those traces the mobile user moves

- among different locations in a city (termed intra-city trajectory),
- among different suburban areas of a city (termed suburban trajectory), and,
- among different cities in highways (inter-city trajectory).

3.8 Performance Assessment of the Fuzzy Controlled Location Predictor

At first, we examine the probability of successful predictions derived from LP by experimenting with various values of m and l without introducing the PIm controller. Figure 3 depicts the probability of successful prediction $P(10, l)$ for various values of l for $N = 500$ predictions. Such size of m is not the most suitable for making predictions with minimum

prediction errors. As illustrated in Figure 3, for a low value of l (short-term prediction), $P(10, l)$ assumes high values. For instance, for $l = 5 < m$, $P(10, l)$ assumes high values for any type of trajectories ($r = 0.5$). In addition, for $r = 1$, that is $m = l = 10$, the probability of prediction assumes values 0.8 and 0.9 for intra-city and intercity trajectories, respectively. When the value of l increases the prediction error also increases, for constant m. For $r = 2$, $P(10, 20)$ assumes values between 0.45 to 0.7 for all types of trajectories.

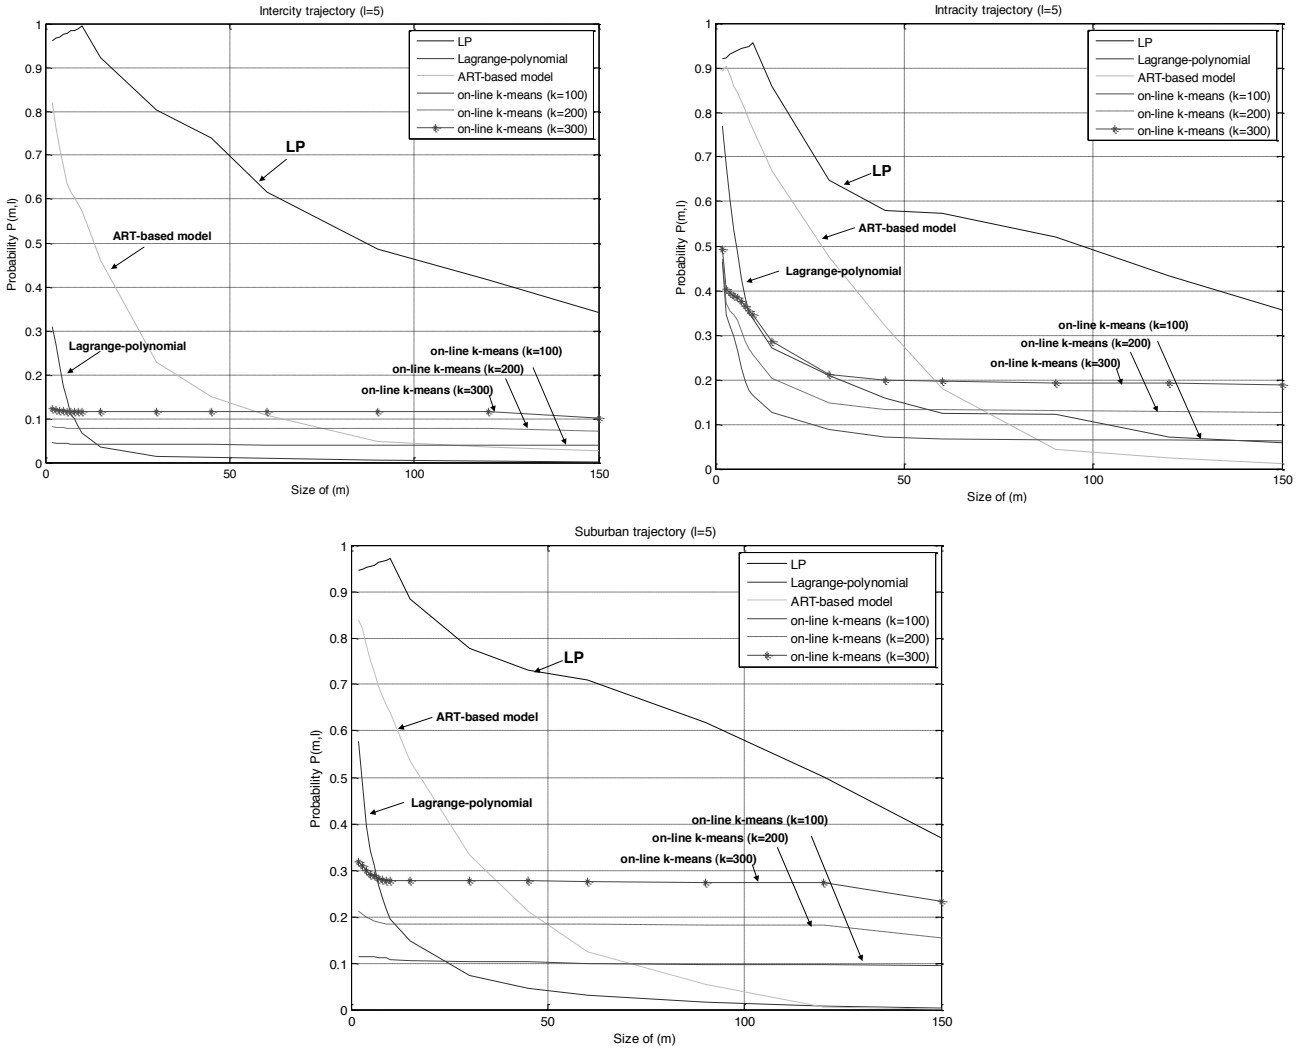

Figure 4. The $P(m, l)$ vs. size of m for the different trajectory types

3.9 Comparison with other Models

We compare the behavior of our stateless model with other stateful and stateless models w.r.t. the probability of successful prediction, space requirements and time complexity. At first, we compare our model with the (stateless) Lagrange polynomial method for location information extrapolation (prediction). In addition, we experiment with two stateful models, the Adaptive Resonance Theory (ART) model enhanced with reinforcement learning and the on-line k-means algorithm. Such models are based on mobility patterns and use pattern classification for a given trajectory in order to predict the future user location. It is worth noting that all considered models adapt their spatiotemporal knowledge base to the mobility behavior of the user. The main criterion for comparison is then how accurately and efficiently, in terms of space and time, each adaptive model predicts the future user location.

3.10 Comparative Performance Assessment

At first, we examine the probability of successful prediction $P(m, l)$ obtained by the stateless and stateful models for the three trajectory types . We experiment with the 'm+l' prediction with l = 5 seconds, vigilance h and threshold value θ set to 10 meters, and k = 100, 200 and 300 in the on-line k-means algorithm. Figure 4 depicts the $P(m, l)$ (vs. the size m) achieved by the algorithms for the intercity, intra-city and sub-urban trajectories. We observe that the LP algorithm achieves better $P(m, l)$ values than the other models in all types of trajectories; specifically, Figure 4 depicts the m^* value for LP, in which $P(m^*, l)$ is maximum (m^* = 10 for l = 5). On the other hand the ART-based model achieves good prediction accuracy, especially, for values of $m < 25$ for the three trajectory types. Specifically, the ART-based model demonstrates good prediction accuracy for the intracity trajectory (only 10% lower than that of LP) due to the fact that the mobile user 'repeats' some mobility patterns moving within the city. However, for $m > 50$ the ART-based model assumes very low $P(m, l)$ and for $m > 100$ it shows the worst performance. Furthermore the Lagrange polynomial model obtains quite good prediction accuracy when m is low. Specifically, it obtains $P(m, l) = 0.78$ with $m = 5$ for the intercity trajectory. For $m > 10$ the Lagrange polynomial is not suitable for predictions. Finally all the on-line k-means models (k = 100, 200 and 300), for every m and in all kinds of trajectories, achieve low performance in terms of prediction accuracy. This is attributed to the fact that such model is not able to increase the predefined number of clusters. Instead it can only readjust them. It is worth noting that the on-line k-means model is independent of the value of m, except for the case of the intra-city trajectory, in which $m < 5$. But, even in this case the performance in prediction accuracy is low.

4. Conclusions

In the first approach we proposed efficient LP schemes based on ML algorithms for trajectory classification. Specifically, the proposed spatial context classifier and a short-term predictor for predicting the location of a mobile user in cellular networks exploits (i) the current position and direction of the user, (ii) history of the trajectories of the user, and, (iii) surrounding location information. We design, implement and evaluate different variants of the proposed LP. Each variant exploits differently the derived knowledge on the mobility behavior of the user (namely the macro and the micro LP). We define the parameters of the short-term LP and introduce certain metrics for evaluating the ability of correct predictions and the efficiency in the prediction process. Moreover, we compare our LP (all the corresponding variants) with popular predictors discussed in the relevant literature. Such predictors are also based on ML algorithms. Simulations with synthetic and real-world mobility data shown that the proposed short-term MiD predictor achieves high prediction efficiency and accuracy, thus, delivering LPs suitable for advanced context-aware applications.

In the second approach we presented how ML can be applied to the engineering of mobile context-aware applications for location prediction. Specifically, we proposed an adaptive ML algorithm for location prediction using ART (a special Neural Network Local Model). We introduce two learning methods: one with non-reinforcement learning and one with reinforcement learning. Furthermore, we deal with two training methods for each learning method: in the supervised method the model uses training data in order to make classification and in the zero-knowledge method the model incrementally learns from unsuccessful predictions. We evaluated our models (versions of the proposed algorithm) with different spatial and temporal parameters. We examine the knowledge bases storage cost (i.e., emerged clusters) and the precision measures (prediction accuracy). Our findings indicate that the C-RLnT suits better to context-aware systems. The advantage of C-RLnT is that (1) it does not require pre-existing knowledge in the user movement behavior in order to predict future movements, (2) it adapts its on-line knowledge base to unseen patterns and (3) it does not consumes much memory to store the emerged clusters. For this reason, C-RLnT is quite useful in context-aware applications where no prior knowledge about the user context is available. Furthermore, through experiments, we decide on which vigilance value achieves the appropriate precision w.r.t. memory limitations and prediction error.

In the third approach we study the proactivity feature of mobile, location-dependent applications and present an approach for mobility prediction exploiting only recent user movement knowledge. We propose a short-memory adaptive LP to address the problem of mobility prediction in the absence of extensive historical information. The location prediction of the proposed LP is obtained by a local linear regression model, while the adaptive capability is achieved through a fuzzy-driven PIm controller. Such controller produces control signals for estimating the best size for the mobility history window attempting to minimize the location prediction error. We experiment with real GPS traces and examine the predictability and adaptability behavior of our LP.

The LP dynamically controls the size of the historical mobility behavior, by the produced control signals, for intra-city,

sub-urban and inter-city trajectories. LP stabilizes to small m^* when dealing with sudden movement changes. This leads to minimizing the probability of inducing noise into the local regression model. In addition, we can conclude that the control signals are more aggressive in the case of intra-city trajectory than in intercity trajectory since the former movement integrates abrupt movement changes while the latter does not. Finally, we show that LP achieves fast adaptation to sudden changes experienced when the user changes mobility behavior (i.e., transition between different trajectory types).

In the forth approach, we propose a sequential trajectory classification and spatial variance reduction system for noise resilient movement prediction of moving objects. The system deals with noisy motion patterns due to random deviations from previously seen patterns, which negatively impacts the accuracy of the prediction result. The system relies on stochastic dynamic programming which relaxes the classification task so that slightly different patterns can be treated as equivalent. Moreover, the system adopts SVRP for keeping K concise and with minimum spatial variance. We provided a comparative assessment of the model with on-line classifiers and a variant adopting the odds algorithm. The proposed model achieves high prediction scores along with efficient data storage of motion patterns.

References

[1] A. Dey, Understanding and using context, *Personal and Ubiquitous Computing*, 5(1), pp. 4-7, 2001.

[2] J. Hightower, G. Borriello, *Location Systems for Ubiquitous Computing*, IEEE Computer, 34(8), August, 2001.

[3] T. Anagnostopoulos, C. Anagnostopoulos, S. Hadjiefthymiades, *Efficient Location Prediction in Mobile Cellular Networks*, International Journal of Wireless Information Networks, *Springer,* Springer Verlag, December, 2011 [doi:10.1007/s10776-011-0166-9].

[4] T. Anagnostopoulos, Christos Anagnostopoulos, Stathes Hadjiefthymiades, M. Kyriakakos, A. Kalousis, *Predicting the location of mobile users: a machine learning approach,* ACM International Conference on Pervasive Services (ICPS09), Imperial College, London, UK., July, 2009.

[5] T. Anagnostopoulos, C. Anagnostopoulos, S. Hadjiefthymiades, A. Kalousis, M. Kyriakakos , *Path Prediction through Data Mining*, IEEE Conference on Pervasive Services, ICPS07, Istanbul, Turkey, July, 2007.

[6] T. Anagnostopoulos, C. Anagnostopoulos, S. Hadjiefthymiades, *An Adaptive Machine Learning Algorithm for Location Prediction*, International Journal of Wireless Information Networks (2011) 18:88-99, Springer, April, 2011.

[7] T. Anagnostopoulos, C. Anagnostopoulos, S. Hadjiefthymiades, *An Online Adaptive Model for Location Prediction*, (3rd Conf.) ICST Autonomics 2009, Limassol, Cyprus, September, 2009 .

[8] T. Anagnostopoulos, C. Anagnostopoulos, S. Hadjiefthymiades, *An Adaptive Location Prediction Model based on Fuzzy Control*, Computer Communications 34 (2011) 816-834, Elsevier, Elsevier, January, 2011.

[9] C. Anagnostopoulos, S. Hadjiefthymiades , T. Anagnostopoulos, *Intelligent Trajectory Classification for Improved Movement Prediction*, IEEE Transactions on Mobile Computing, November, 2011., [revision].

Extraction of Semantic Information from Events

Panagiota Antonakaki *

Department of Informatics and Telecomunications
National and Kapodistrian University of Athens
ganton@iit.demokritos.gr

Abstract. Behavior understanding from events can be considered as a typical classification problem using predefined classes. In our research presented in this thesis, the main idea is to deal with this classification problem using motion analysis for behavior recognition. We describe our research work from its first steps to the final accomplishments. Novel methods are introduced including a methodology whereby frame information is coded in graphs (called Optical Flow Proximity Graphs - OFPGs), using only optical flow to form the feature vector. A symbolic method including two levels of graph representation is also proposed dealing with the same recognition problem. OFPGs are also proposed for video indexing. Furthermore, a bottom-up approach for anomaly detection using a multi camera system is proposed. In the framework of this system, we present the use of one class continuous Hidden Markov Models (cHMMs) for the task of human behavior recognition. An approximation algorithm, called Observation Log Probability Approximation (OLPA), is proposed to overcome numerical stability problems in the calculation of probability of emission for very long observations.

Keywords: video processing, behavior recognition, event detection, anomaly detection

1 Introduction

One of the most important goals of visual surveillance systems is to track objects and further analyze their activities in order to recognize behaviors and even detect anomalies. To this end we try to detect objects in a scene, track them, recognize their action and interactions in order to understand and describe their behaviors. As a behavior we determine a performed motion by a human subject (object of interest). This motion is captured by low-cost cameras (web-cameras) placed in the room where the object of interest is moving, from a distance in order to observe the whole room.

This thesis is concerned with two main subjects. The first subject is based on methods used during our research in order to classify observed motion to predefined behaviors. The part of the thesis focusing on this subject includes all

* Dissertation Advisor: Sergios Theodoridis, Professor - Dr Stavros Perantonis, Researcher

our research steps for behavior recognition and the corresponding experimental results per approach. Before the completion of this part, a novel model applied on raw data of the video for behavior recognition using graph representation is proposed. The behaviors recognized are also used for video indexing. The second subject deals with the classification of behaviors as normal or abnormal. In the part considering this subject we also propose a novel approach of anomaly detection based on short term behavior and trajectory classification using one class Support Vector Machines and an alternation of continuous Hidden Markov Models also used as one class classifiers.

2 Related Work

Human motion analysis is receiving increasing attention from computer vision researchers [8]. For behavior recognition in video streams, many features based on motion, trajectory or shape of the object are used in different methods. For example, Ribeiro et al. in [9], deal with the problem of feature selection and they present features based on the motion and the trajectory of the object, like velocity, speed, optical flow, etc.

Most approaches for 2D interpretation of human body structure focus on motion estimation of the joints of body segments between consecutive frames. Leo et al. in [10] attempt to classify actions at an archaeological site. They present a system that uses binary patches and an unsupervised clustering algorithm to detect human body postures. A discrete HMM is used to classify the sequences of poses into a set of four different actions. In [11], the used feature is called *Star Skeletonization* and is a distance of the extremities of the silhouette from the blob's center of gravity. In [12] the vertical and horizontal projections of the blob are used to recognize people posture. A simple nearest neighbour classifier, that compares the current projections with manually tuned models, is adopted. Roh et al. in [13] base their action recognition task on curvature scale space templates of a player's silhouette.

Recently, several researchers have dealt with the problem of anomaly detection, which is the process of behavior classification as normal or abnormal. A variety of methods, ranging from fully supervised [14, 15], to semi-supervised [16] and unsupervised systems [17–20], have been proposed in the existing literature. It should be noted, however, that most of the existing approaches do not use multi-camera information, except for Zhou et al. in [21], where multiple video streams are combined via a coupled Hidden Markov Model.

Tarassenko et al. in [22] are proponents of the idea that learning normality alone is all that is required for the detection of abnormality.

3 Human Behavior Recognition and Video Indexing

In this thesis we present all steps made during our research work for one person's behavior recognition, including two similar approaches.

3.1 First Approaches for Short Term Behaviors ([3])

Our first approaches propose a hierarchical method separated into steps, each step dealing with a distinct sub-problem. The main idea is to recognize short term behavior and detect events in order to use this information for long term behavior recognition. As short term behaviors we define actions observed under specific time and space constraints and as events we define the change between two different successive short term behaviors. Long term behaviors are described as scenarios which involves different short term behaviors and events.

Because the low performance of the first approaches, the fact that the number of features used in the feature vector increases the computational cost, and based on the assumption that optical flow by itself can be descriptive enough, we also propose a different approach for both one person and multiple persons behavior recognition, using features calculated directly from raw data without tracking being necessary ([23], [26], [**?**]).

3.2 Final Approach for Behavior Recognition and Video Indexing

For the behavior recognition problem we propose two methods based on graph representation. The first method includes graphs (OFPGs) extracted by each frame containing motion neighbourhood information from the whole frame (WFGA). The second method is a symbolic approach including two levels of graph representation (SGA).

In Whole Frame Motion Representation (WFGA), we use the optical flow vectors calculated from the whole frame using a pixel-by-pixel analysis. In the training step we calculate the optical flow vectors from the whole frame, for each frame. Then, the corresponding OFPG per frame is extracted. Thus, all observed behaviors in a frame are represented by the same graph. The graphs extracted from each frame are merged using the U operator into the graphs which represent the frame's assigned classes.

After that we describe a frame as a feature vector. The vector contains the similarities of our frame OFPG to all the class graphs.

Then, an one-class Support Vector Machine (SVM) SVM_c per class c is trained (see Figure 1).

In the testing step, a similar procedure is used. Each testing instance (frame) is represented by one graph. We create the similarity-based feature vector. Each trained SVM_c model decides whether that feature vector (and the corresponding frame) belongs to class c (see Figure 2)..

In the Symbolic Craph-Based Approach (SGA) the frame is segmented into equal areas, each of which is represented by its OFPG (see Figure 3). The resulting graphs of this level are used as symbols forming an index of graphs to hold the different symbols observed in the frame. At the second level, we use this index in order to map each segment of the frame with a specific symbol.

In the second level of graphs, we use the index of graphs to create a symbolic representation of the original frame in the video sequence (see Figure 4 for an overview of the process). We create and update an index, mapping symbols to

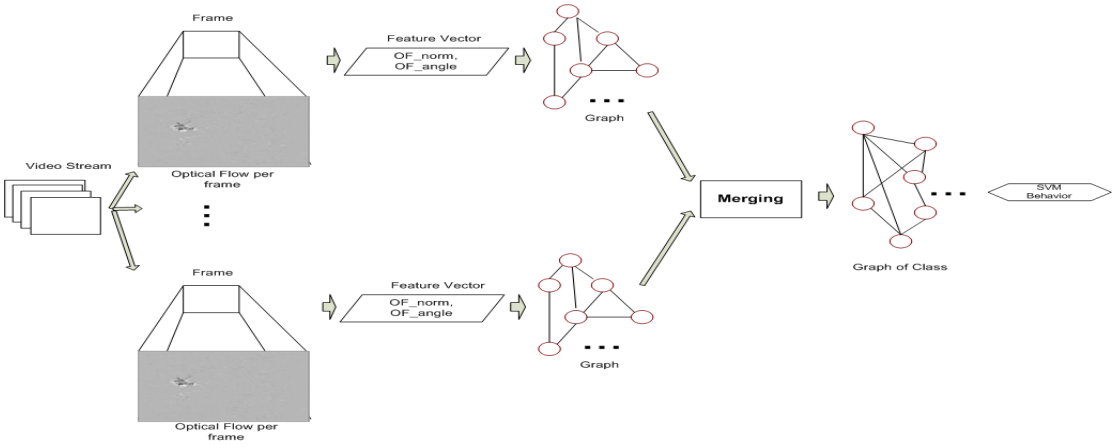

Fig. 1. The training step of WFGA.

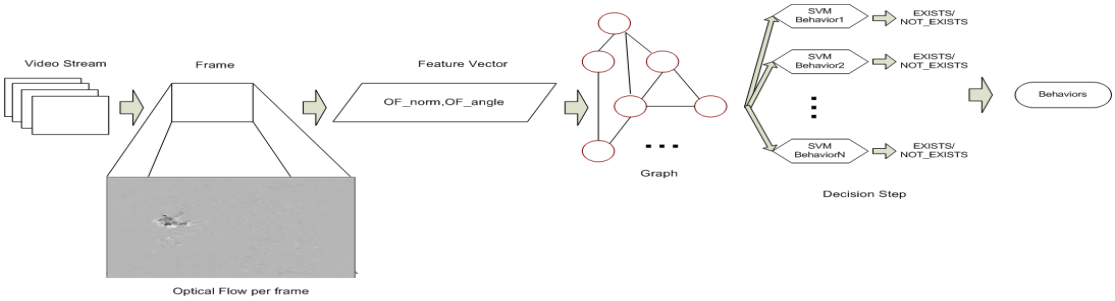

Fig. 2. The testing step of the WFGA framework.

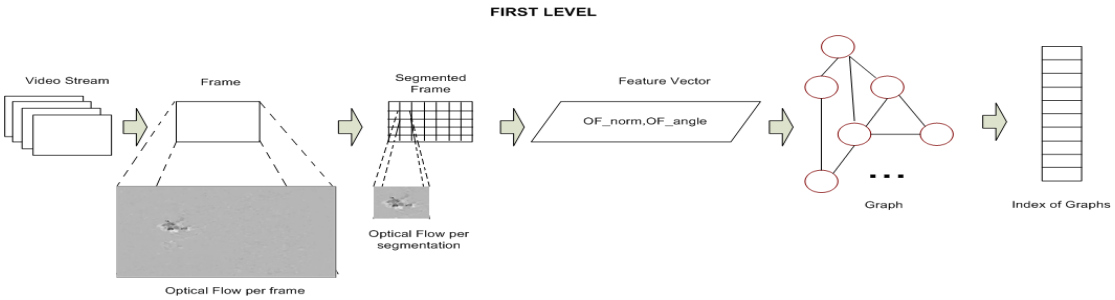

Fig. 3. Extraction of OFPGs from each segment.

subgraphs. These subgraphs are OFPGs of image segments. Each subgraph is assigned by the index a symbol.

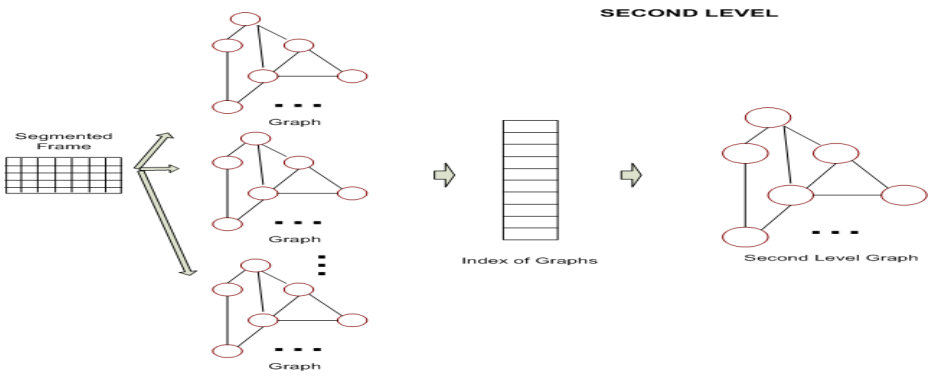

Fig. 4. Extraction of second level graphs.

Once the second level graph of each known class is generated in the training step, the classification step follows. For the classification, each testing instance is represented by a second level graph. The similarity feature vectors are extracted and the SVM classifiers give a binary answer on whether a particular behavior is observed in the frame or not. See Figure 5 for the overview of the testing step in the SGA case. Similarly to WFGA, for the video indexing procedure all behaviors observed in the video are used to tag this video.

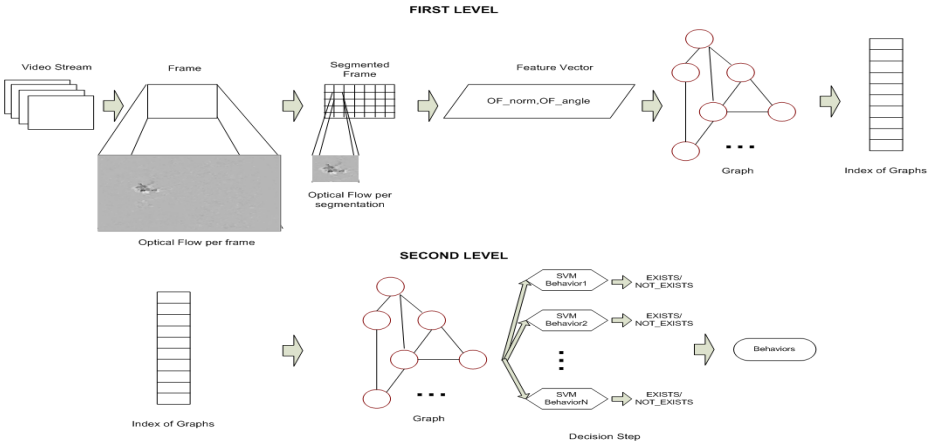

Fig. 5. The testing step of SGA.

4 Anomaly Detection

One of the goals of Smart Surveillance is anomaly detection. Anomaly detection refers to detecting patterns in a given data set that do not conform to an established normal behavior. The patterns are called anomalies and often translate to critical and actionable information in video surveillance calling for human attention if necessary. The second subject of this thesis deals with this problem proposing a novel method for anomaly detection.

The proposed methodology is based on the fusion of data that we collect from several cameras with overlapping fields of view ([**?**]). We perform classification using two different one-class classifiers, a Support Vector Machine (for motion classification) and a continuous Hidden Markov Model (for trajectory classification). A trajectory is the path that a moving object of interest follows through space as a function of time. In our experiments the trajectory of a person includes points (x,y) on the common views of multiple cameras. The final decision on the behavior is made by taking into account outputs from both classifiers.

In motion representation and analysis, our methodology uses information obtained by preprocessing, namely the object's bounding box, the object's blob and sequential positions.

The short term activity is represented by a 7-dimensional feature vector, as follows:

$$f = (v(t), \widehat{v_T}(t), R_T(t), F(t), \Delta F(t), \max(\Delta H(t)), \max(\Delta SD(t))) \quad (1)$$

The corresponding features calculated are: speed, algebraic mean speed, algebraic mean bounding box difference, mean optical flow, mean optical flow difference, max entropy difference and maz standard deviation difference.

The decision whether a short-term behavior is normal or not can be taken by employing a one-class SVM as proposed by Scholkopf [6].

Our second information source for evaluating behavior is the trajectory. The problem of discriminating between normal/abnormal trajectories concerns the definition of a measure that would give sufficiently different values for the two classes. The variable length of the trajectories poses additional difficulties. Long, normal trajectories would have cHMM generation probability values comparable to small values of short, abnormal trajectories, so the observation's length factor needs to be removed.

We proved that for a normal observation sequence (O_{normal}) and for an abnormal one ($O_{abnormal}$) the following condition must hold:

$$\frac{\log P(O_{abnormal}|\lambda)}{length(O_{abnormal})} \ll \frac{\log P(O_{normal}|\lambda)}{length(O_{normal})} \quad (2)$$

then we will be able to use it as a classification measure.

Long normal sequences give small values of cHMM probabilities, due to successive multiplications, making the logarithm of those probabilities to be too high to let the 1 to be damaging. Assuming that the approximation $\frac{\log P(O|\lambda)}{length(O)}$ with $\frac{\lfloor \log P(O|\lambda) \rfloor}{length(O)}$ is acceptable, it can be inserted to Forward Backward algorithm.

According to the above approximations, we can express the algorithm as follows.

1. Initialization:
   $\tilde{\alpha}_1(i) \simeq \lfloor \log \pi_i \rfloor + \lfloor \log b_i(O_1) \rfloor$
2. Induction:
   $\tilde{\alpha}_t(i) \simeq max_j(\tilde{\alpha}_{t-1}(j) + \lfloor \log a_{ij} \rfloor) + \lfloor \log b_j(O_t) \rfloor)$
3. Termination:
   $\tilde{P}(O|\lambda) \simeq \max_i \tilde{\alpha}_t(i)$

This Observation Log Probability Approximation (OLPA) algorithm helps us overcome the problem of consecutive multiplications, by making it possible to use a sum of integers.

5 Experiments

5.1 Behavior Recognition and Video Indexing

For our experiments in behavior recognition and video indexing, we used three different sets of data. In the first data set, simple behaviors were captured in our lab[1](for Semveillance project). The second data set is the commonly used data set from the PETS04 workshop [4]. These latter video sequences have already been used by the CAVIAR project. The third data set is the dataset collected in *The Weizmann Institute of Science*, used by the authors in [5], for comparison reasons.

Our evaluation of the final approach started with the PETS04 data set to determine whether WFGA is better than SGA. On our corpus, the WFGA failed to run in some cases, because of memory complexity.

WFGA

Behavior	Precision	Recall	F-measure
browser	0.2093	0.3459	0.3273
walker	0.9423	0.9491	0.9456
fighters	0.1263	0.9461	0.2223
meeters	0.2934	0.9810	0.4294

SGA

	Precision	Recall	F-measure
browser	0.8065	0.7366	0.7377
walker	0.9918	0.8480	0.9129
fighters	0.5608	0.8766	0.6437
meeters	0.6685	0.8537	0.7448

a

SGA

Behavior	Precision	Recall	F-measure
walk	0.7165	0.9437	0.8138
run	0.3367	0.9376	0.4952
abrupt	0.4003	0.9457	0.5616
active	0.7520	0.9479	0.8381
inactive	0.9853	0.9478	0.9662

b

Behavior	Precision	Recall	F-measure
walk	0.9687	0.9464	0.9573
run	0.7343	0.9473	0.8271
abrupt	0.7782	0.9468	0.8533
active	0.9899	0.9468	0.9678
inactive	0.9999	0.9442	0.9712

c

Fig. 6. (a). Experimental results on PETS04 dataset with noise removal. (b) Precision, Recall and F-measure of SGA on our dataset with noise removal and (c) without noise removal.

[1] The dataset can be made available on demand to the author.

The results of our experiments using the PETS04 data for both proposed approaches are shown in Figure 6. The results of the SGA are higher than those of the WFGA. This indicates that SGA has both lower computational cost and better experimental results.

The application of the SGA method on our dataset yielded promising results as well, especially if one considers that we only use optical flow to detect behavior. The results from experiments using noise removal are shown in Figure 6.

The results of our experiments using the Weizmann data set follow the leave-one-out method, i.e. for every video sequence we remove the entire sequence from the database while other actions of the same person remain. The results are shown in the Figure 7. The results presented in those tables include again low precision values but higher recall values. This is, again, due to the fact that the classes are unbalanced. In Figure 7 we included the results after noise removal and sampling our data set. The latter results are in most cases lower than the results without sampling, due to the fact that video sequences in this data set are small enough (few frames per video) and important information was excluded due to sampling. The same is valid for noise removal.

The experimental results for video indexing using the PETS04 data are shown in Figure 7. These results imply that, when a behavior observed in the video is not included in the training set — and thus not modeled — the system can make the decision that the specific behavior does not belong to any observed and modeled behavior. The specificity values are high enough (~ 0.83), allowing us to conclude that in video indexing the success of the approach follows the success of the behavior recognition case.

Behavior	Precision	Recall	F-measure
walk	0.3188	0.9460	0.4768
run	0.1780	0.9413	0.2994
skip	0.2149	0.9428	0.3499
jack	0.7452	0.9467	0.8339
jump	0.2040	0.9437	0.3354
pjump	0.2041	0.9449	0.3357
side	0.4636	0.9447	0.6219
wave1	0.2556	0.9462	0.4024
wave2	0.2190	0.9460	0.3557
bend	0.5634	0.9456	0.7058

a

Behavior	WFGA Specificity	SGA Specificity
browser	0.3521	0.8444
walker	0.1829	0.8613
fighters	0.5257	0.8157
meeters	0.2605	0.8077

b

Fig. 7. (a). Precision, Recall and F-measure of SGA on Weizmann dataset with noise removal. (b) Experimental results for video indexing.

5.2 Anomaly Detection

As a scene for our experiments we have used our lab, where we installed three cameras, and there we tried to simulate some common scenarios[2]. The experiments measure the performance of two variations of our process, namely the offline and the real-time process.

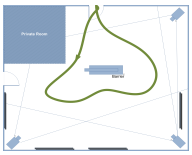

Fig. 8. Example of normal trajectory in the scene.

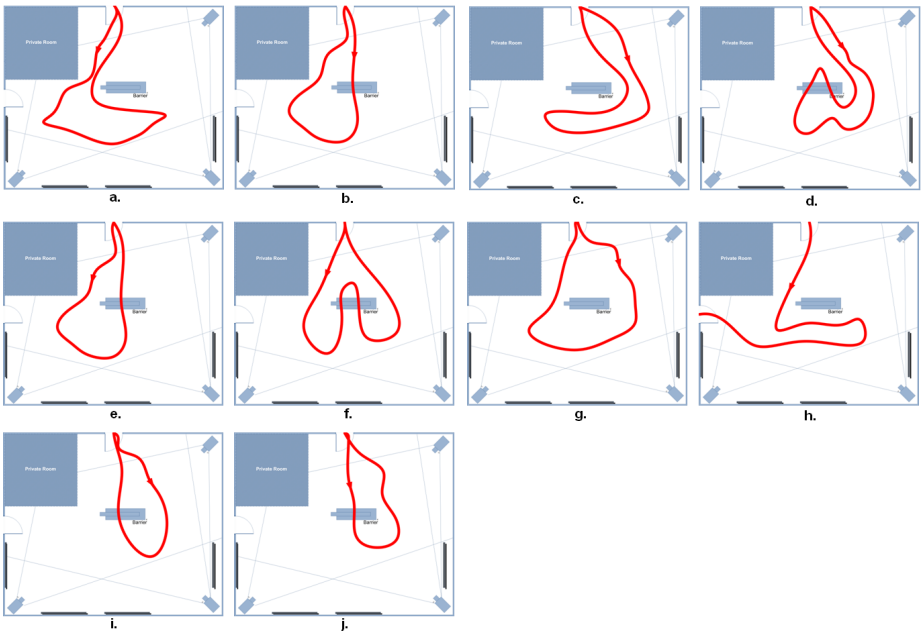

Fig. 9. Examples of abnormal behaviors in the scene.

We have performed a 10-fold cross validation method to test the effectiveness of our system using the offline approach. The videos with normal behaviors

[2] The custom corpus used within our experiments can be made available to any interested party, via e-mail correspondence.

illustrate a person entering the room, buying a ticket, browsing and looking around for several minutes and exiting the room using a preset path (Figure 8). The abnormal behaviors consist of running, abrupt motion or unexpected trajectory (Figure 9).

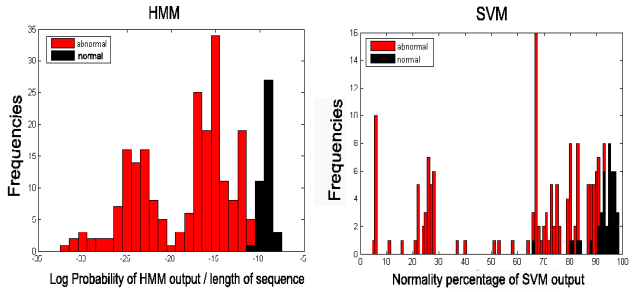

Fig. 10. a. Percentage normality in normal and abnormal behaviors for Support Vector Machine, b. output of continuous Hidden Markov Model for normal and abnormal behaviors. Black color is for normal behaviors and red for abnormal behaviors

For SVM-based classification we set the threshold to be the following function of the mean and the standard deviation of the distribution of the number of allowed abnormal frames within a normal sequence:

$$threshold_{SVM} = mean(Hsvm_{normal}) - 2.5 \cdot std(Hsvm_{normal}) \quad (3)$$

For HMM outputs the minimum value of the distribution of normalized log-probabilities of the normal instances was considered to be the threshold value that separates normal trajectories from the abnormal ones:

$$threshold_{HMM} = \min(Hhmm_{normal}), \quad (4)$$

where Hsvm is the histogram of SVM's outputs and Hhmm is the histogram with HMM's outputs.

Precision and recall have been calculated for the offline and the real-time experiments. For each approach we give the performance for both the SVM and HMM classifier models separately, as well as for the whole system in Figure 11.

6 Conclusions

This thesis was concentrated on high level processes of behavior recognition problem extracting semantic information from events taking place in a video. The first part of this thesis includes preprocessing steps such as background subtraction and tracking in order for the proposed high level approach that we propose to accomblish behavior recognition. In the second part of this thesis a

3-cameras

Offline

	SVM		HMM		Overall	
	Precision	Recall	Precision	Recall	Precision	Recall
Normal	0.9048	0.9286	1	0.9762	1	0.9286
Abnormal	0.7674	0.7071	0.95	1	0.88	1

Real-time

	SVM		HMM		Overall	
	Precision	Recall	Precision	Recall	Precision	Recall
Normal	0.9875	0.9228	0.9960	0.9770	0.9960	0.9105
Abnormal	0.2419	0.6788	0.8478	0.9704	0.8478	0.9375

one camera

Offline

	SVM		HMM		Overall	
	Precision	Recall	Precision	Recall	Precision	Recall
Normal	0.9788	0.8375	1	0.95	1	0.8
Abnormal	0.6708	0.9464	0.913	1	0.7366	1

Real-time

	SVM		HMM		Overall	
	Precision	Recall	Precision	Recall	Precision	Recall
Normal	0.9945	0.9148	0.9953	0.9597	0.9975	0.8525
Abnormal	0.2696	0.8569	0.7544	0.9637	0.5042	0.9861

Fig. 11. Precision and Recall for 3-camera and one camera System on our dataset. The column "Overall" indicates the performance of the combined decision.

bottom-up approach for human recognition understanding is presented, using a multi camera system for anomaly detection.

For behavior recognition adn video indexing we have presented a unified system, which needs no preprocessing steps and a priori knowledge about the surveillance room or how many people are being observed. We have proposed two innovative approaches for classification — and consequently semantic annotation — and indexing of videos, using graph-based representations and analysis methods, with promising experimental results.

Our experimental results demonstrated the good performance of the proposed approaches in the task of recognizing human behaviors in a somewhat noisy environment, with different scenarios of action and participation of different actors. The experiments were implemented using two different datasets with good performance, implying the robustness of the method.

For anomaly detection we have presented a set of theoretical and practical tools for the domain of behaviour recognition, which have been integrated within a unified, automatic, bottom-up system based on the use of multiple cameras performing human behaviour recognition in an indoor environment, without a uniform background.

References

1. Hu, W. and Tan, T. and Wang, L. and Maybank, S.: A Survey on Visual Surveillance of Object Motion and Behaviors. IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews. 34, 3, 334–352 (2004)
2. Zhou, H. and Kimber, D. and com Inc, A. and Seattle, WA: Unusual Event Detection via Multi-Camera Video Mining. 18th International Conference on Pattern Recognition. 3 (2006)
3. Efros, A.A. and Berg, A.C. and Mori, G. and Malik, J.: Recognizing Action at a Distance. IEEE International Conference on Computer Vision. 2, 726–733 (2003)
4. Fisher, R.B.: The PETS04 Surveillance Ground-Truth Data Sets. Proc. 6th IEEE International Workshop on Performance Evaluation of Tracking and Surveillance. 1–5 (2004)

5. Lena Gorelick and Moshe Blank and Eli Shechtman and Michal Irani and Ronen Basri: Actions as Space-Time Shapes. Transactions on Pattern Analysis and Machine Intelligence (2007)
6. Schölkopf, B. and Platt, J.C. and Shawe-Taylor, J. and Smola, A.J. and Williamson, R.C.: Estimating the Support of a High-Dimensional Distribution. Neural Computation. 13, 7, 1443–1471 (2001)
7. Manevitz, L.M. and Yousef, M.: One-Class SVMs for Document Classification. The Journal of Machine Learning Research. 2, 154 (2002)
8. Moeslund, T.B. and Hilton, A. and Kruger, V.: A Survey of Advances in Vision-Based Human Motion Capture and Analysis. Computer Vision and Image Understanding. 104, 90–126 (2006)
9. Ribeiro, P.C. and Santos-Victor, J. and Lisboa, P.: Human Activity Recognition from Video: Modeling, Feature Selection and Classification Architecture. Proceedings of International Workshop on Human Activity Recognition and Modelling. 61–78 (2005)
10. Leo, M. and D'Orazio, T. and Gnoni, I. and Spagnolo, P. and Distante, A.: Complex Human Activity Recognition for Monitoring Wide Outdoor Environments. Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on. 4 (2004)
11. Fujiyoshi, H. and Lipton, A.J. and Kanade, T.: Real-Time Human Motion Analysis by Image Skeletonization. IEICE Transactions on Information and Systems E Series D. 87, 1, 113–120 (2004)
12. Haritaoglu, I. and Harwood, D. and Davis, L.S.: Ghost: A Human Body Part Labeling System using Silhouettes. International Conference on Pattern Recognition. 14, 77–82 (1998)
13. Roh, M. and Christmas, B. and Kittler, J. and Lee, S.: Robust Player Gesture Spotting and Recognition in Low-Resolution Sports Video. LECTURE NOTES IN COMPUTER SCIENCE. 3954, 347 (2006)
14. Dee, H. and Hogg, D.: Detecting Inexplicable Behaviour. British Machine Vision Conference. 447, 486 (2004)
15. Duong, T. and Bui, H. and Phung, D. and Venkatesh, S.: Activity Recognition and Abnormality Detection with the Switching Hidden Semi-Markov Model. IEEE Computer Society Conference on Computer Vision and Pattern Recognition. 1, 838 (2005)
16. Zhang, D. and Gatica-Perez, D. and Bengio, S. and McCowan, I.: Semi-Supervised Adapted HMMs for Unusual Event Detection. IEEE Computer Society Conference on Computer Vision and Pattern Recognition. 1, 661 (2005)
17. Jiang, F. and Wu, Y. and Katsaggelos, A.K.: Abnormal Event Detection from Surveillance Video by Dynamic Hierarchical Clustering. Proceedings IEEE International Conference on Image Processing. 5, 145–148 (2007)
18. Lee, C.K. and Ho, M.F. and Wen, W.S. and Huang, C.L. and Hsin-Chu, T.: Abnormal Event Detection in Video using N-cut Clustering. International Conference on Intelligent Information Hiding and Multimedia Signal Processing (IIH-MSP). (2006)
19. Mahajan, D. and Kwatra, N. and Jain, S. and Kalra, P. and Banerjee, S.: A Framework for Activity Recognition and Detection of Unusual Activities. Indian Conference on Computer Vision, Graphics and Image Processing. 15–21 (2004)
20. Xiang, T. and Gong, S.: Incremental and Adaptive Abnormal Behaviour Detection. Computer Vision and Image Understanding. (2008)

21. Zhou, H. and Kimber, D. and com Inc, A. and Seattle, WA: Unusual Event Detection via Multi-Camera Video Mining. 18th International Conference on Pattern Recognition. 3 (2006)
22. Tarassenko, L. and Nairac, A. and Townsend, N. and Buxton, I. and Cowley, P.: Novelty Detection for the Identification of Abnormalities. International Journal of Systems Science. 31, 11, 1427–1439 (2000)
23. Dimitrios I. Kosmopoulos, Panagiota Antonakaki, Konstandinos Valasoulis, Dimitrios Katsoulas: Monitoring Human Behavior in an Assistive Environment using Multiple Views. Proceedings of the 1st international conference on PErvasive Technologies Related to Assistive Environments (PETRA 2008). 32, (2008)
24. Dimitrios I. Kosmopoulos, Panagiota Antonakaki, Konstandinos Valasoulis, Anastasios L. Kesidis, Stavros J. Perantonis: Human Behavior Classification Using Multiple Views. Proceedings of the 5th Hellenic conference on Artificial Intelligence: Theories, Models and Applications (SETN 2008). 123–134, (2008)
25. Dimitrios I. Kosmopoulos, Anastasios Kesidis, Panagiota Antonakaki, Konstandinos Valasoulis, Stavros J. Perantonis: SemVeillance System: Tracking and Behavior Recognition Under Occlusions. European Conference on Artificial Intelligence (ECAI 2008). (2008)
26. Panagiota Antonakaki, Dimitrios I. Kosmopoulos, Stavros J. Perantonis: Detecting Abnormal Human Behaviour using Multiple Cameras. Signal Processing. 89(9), 1723–1738 (2009)

A Web-based Case-based Learning Environment – Use in the Didactics of Informatics

Maria Boubouka*

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications
mboub@di.uoa.gr

Abstract. The present dissertation examines the didactic utilization of cases, specifically through digital learning environments. The review of existing digital case-based learning environments indicated the need for the development of an environment that enables the creation of learning activities of different types based on one or more than one cases. Towards this end, the design of the learning environment CASTLE (CASes for Teaching and LEarning) is proposed. Innovative characteristics of CASTLE concern: the distinction of cases and activities that are handled as different entities, the generic design that can be specified for teaching different subject matters and the support of interaction between users. A prototype of CASTLE for teaching Programming has been created and is described, along with a proposed template of Programming cases which is supported by experimental evidence. Finally, the evaluation of CASTLE prototype and an investigation for the specification of CASTLE in Didactics are presented.

Keywords. Cases, Problems, Learning Activities, Educational scenarios, Didactics of Programming

1 Introduction

One of the most salient goals of modern education is the development of problem solving skills [1-2]. Cases which are context-rich problem descriptions are essential ingredients of learning environments aiming at enabling students to deal with problems. Context information included in cases helps students understand the meaningfulness of dealing with the problems and therefore increases their motivation for engagement in the problem solving process [3-4].

In bibliography there is neither a unique definition for the term case nor a consensus for case content or structure [5-7]. This is due to the fact that cases are being used in a great variety of learning activities, some of which demand only the problem description and others necessitate that this description be accompanied by one or more

* Dissertation Advisor: Maria Grigoriadou, Emeritus Professor

solutions [8]. Jonassen [9] characterizes cases as "building blocks of problem-based learning environments" and summarizes seven different ways of case utilization in teaching: 1) cases as problems to solve, 2) cases as worked examples, 3) case studies, 4) cases as analogues [10], 5) cases as prior experiences [11-13], 6) cases as multiple perspectives [7], 7) cases as simulations. The aforementioned applications of cases differ in terms of: a) the structuredness of the problems (well structured or ill structured) [1], b) the number of different cases that are being involved (one or more than one case) and c) how the case is being used i.e. as an example to be studied or as a problem to be solved. Within the wide range of different learning activities that can be designed based on cases there have been included activities adequate for both novices and experts on a subject matter. Moreover, the fact being that case authoring is a rather demanding task, the need to maximize case utilization through the creation of more than one learning activities based on one case is emphasized.

In the last decades a large number of digital case-based learning environments have been developed, covering a variety of learning subject matters. The extended review performed, and analytically described in the dissertation, focused on the structure and content of cases as well as on the types of activities based on the cases in the examined learning environments.

Considerable differences have been noticed in the length of the cases included in the reviewed environments. Cases in the form of story problems in learning environments for Mathematics [14] or Physics [15] are rather short (a couple of paragraphs), while cases describing design problems in Architecture [16] or Computer Engineering [17,18] are quite long (multiple pages, collections of files connected by analytical narratives of the problem situation and the solution process). Moreover, besides the problem description, a case may contain or not the description of its solution process given by an expert.

As far as the case-based learning activities are concerned, each learning environment usually contains one single type of activity. In particular, there are environments where students are asked to solve the case problems [19]; others where students can study cases already solved by experts [16] and others where students have to participate in group discussions based on one case [20]. In another group of environments students working on a case are engaged in a sequence of activities, which is always the same for all cases included in the environment [21]. Finally, only a small number of environments support activities asking for comparisons between cases [22].

Furthermore, few learning environments allow case authoring [23]. In the majority of the environments the case collection is created by their developers and no authoring functionalities are available for the final users (instructors or learners). Finally, none of the reviewed case-based environments have social software functionalities that would allow the exchange of comments on the learning material between users.

To sum up, the review of existing case-based digital learning environments has indicated that the creation of distinct learning activities based on one case remains an open research issue. The development of such a learning environment is of critical importance as there are case-based activities adequate for both novices and advanced learners in a subject matter. Moreover, case authoring is an exacting task. Thus, a learning environment that supports case authoring and case reuse in distinct case-

based activities may consist a powerful tool for the instructor by enabling him/her to create learning activities adapted to the educational needs of his/her students.

Towards this end the present dissertation proposes the design of the learning environment CASTLE (CASes for Teaching and LEarning), which aims to support teaching and learning through cases. The design of CASTLE is generic and can be adapted to the specific needs of a learning subject matter through the specification of the design of cases and case-based learning activities. In the dissertation experimental evidence is presented concerning: the development of a prototype of CASTLE for teaching Programming, the evaluation of this prototype by novice and expert Programming teachers who used the environment as authors, and an investigation for the specification of CASTLE in Didactics.

2 The design principles of CASTLE

The main requirement for the design of CASTLE was to support the authoring of cases as well as the authoring and elaboration of different types of learning activities based on cases. Within the different types of activities, the same case may play a role ranging from being an already solved problem offered to be studied as an example, to being presented as a novel problem to be solved. Other types of activities may involve two or more cases. Such an activity asks students to compare the solutions of two isomorphic case problems (different story, same procedure)[24] in order to identify structural similarities, or to study a solved case and engage in solving an isomorphic one. This fundamental requirement is being met in CASTLE by handling cases and case-based learning activities as distinct entities which are interrelated between them.

Three different user roles are being supported in CASTLE: teacher, learner and administrator. In CASTLE a teacher can: author cases and case-based activities; organize learning activities in collections; create sequences of learning activities in order to form educational scenarios; assign these scenarios to learners and follow their elaboration. Learners engage in the elaboration of educational scenarios individually or in groups. Learners are also allowed to author cases themselves. Administrators in CASTLE are responsible for the management of user accounts, subject matters and learning material in the form of cases, activities and scenarios.

Additionally, CASTLE supports the interaction between the groups of teachers and learners thus supporting the formation of communities. In other words, CASTLE has social software functionalities [25]. For example, a case author, after publicizing a case in CASTLE, can receive both verbal comments on the case and an arithmetical grade from 1 to 5 by other users, teachers and learners. The received comments and feedback can facilitate the optimization of the learning material as authors have the possibility to commit improvements. Every case is owned by its author in the sense that only its author can make changes on a case and create an updated version of it. However, no ownership is required on a case in order to create a learning activity based on it. This means that there are interactions between users during the creation of the learning material.

Furthermore, data from user interaction is collected in CASTLE while students elaborate educational scenarios. This data is used as an input to the open learner model [26, 27] of CASTLE and enables the provision of personalized feedback to the learners.

The generic design of CASTLE presented here needs further specification in order to meet the special demands of each learning subject matter. This procedure takes place through the selection of case content and structure as well as the types of case-based learning activities adequate for a given subject matter. In the following paragraph the specification process leading to the prototype of CASTLE for Programming teaching and learning is described.

3 Specification process of CASTLE in Programming .

3.1 Specification of case structure

In order to specify the case structure for the subject matter of Programming, a review of the relevant literature was conducted [28]. Two different Programming *case study* structures have been found in bibliography: Structure A proposed by Linn and Clancy [29] and Structure B proposed by Spooner and Skolnick [30] (see Table 1). These structures have both similarities and differences when compared. For example, both structures propose that a case should contain the problem description (*Programming problem statement* in Structure A, *Motivation*, *Background* in Structure B), an explanation for the problem solution (*Solution process description* in Structure A, *Algorithm development* and *New Programming concepts* in Structure B), the code of the program produced to solve the problem (*Code listing* in Structure A, *Solution program* in Structure B) and questions on the solution (*Study* & *Test questions* in Structure A, *Discussion* & *Further study* in Structure B).

Table 1. Programming Case study structures

Structure A	Structure B
Programming Problem statement	Motivation
Solution process description	Background
Code listing	Algorithm development
	New Programming concepts
	Solution program
Study questions	Discussion
Test questions	Further study

In CASTLE, *Cases* and case-based *Activities* are considered different entities and are handled independently. In the *Case study* structures described above, there are parts that correspond to the *Case* entity of CASTLE and parts that correspond to the *Activi-*

ties (e.g *Study & Test questions* in Structure A and *Discussion & Further study* in Structure B). The comparison between the two structures shows that differences reside in the *Case* part of the *Case studies* and, more specifically, in the parts referring to the explanation of the problem solution. In particular, Structure A proposes an explanation of the critical decisions reached by an expert programmer while writing the solution program. On the contrary, Structure B suggests a gradual introduction of the solution, starting with the description of the algorithm development and proceeding with the new programming concepts required in the program code.
In order to select the structure of Cases in the CASTLE prototype for Programming teaching, the two structures have been compared in two empirical studies.

First empirical study.

The first empirical study aimed to compare the two structures in terms of their efficiency and their acceptance by the learners.

In particular, the research questions were:

- does the structure of a programming case affect the ability of learners to develop programs for resolving similar problems?
- does the structure of a programming case affect the time required to study it?
- what are the opinions of the learners about the structure and the usefulness of cases?

Participants.

102 first-year students participated in the study. They had enrolled to the course "Introduction to Informatics and Telecommunications" at the Department of Informatics and Telecommunications of the National and Kapodistrian University of Athens. The students formed two groups of 51 members each: Group 1 and Group 2.

Procedure.

The empirical study consisted of the following phases:

- Phase A, *Pre-test* (30 min): the students of both groups worked on a programming problem and were asked to develop a program in order to solve it.
- Phase B, *Case study* (60 min): the students of Group 1 worked on a Case study structured according to Structure A and the students of Group 2 worked on a Case study structured according to Structure B. The case problem for both groups was the same.
- Phase C, *Post-test* (30 min): the students of both groups worked on a problem isomorphic to the problem of the pre-test phase. As in phase A, their task was to develop a program.
- Phase D, *Filling the questionnaire* (15 min). All students were asked to fill in a questionnaire on their opinions on the Case study assigned to them in phase B.

Students were also asked to write down in their response sheet the actual time they had spent on the completion of phases A, B and C.

Results.

The programs developed by the students during the pre-test and post-test phases have been evaluated independently according to criteria set by two experienced Programming teachers who assigned a five-point scale grade (0-4) for each criterion. The total grade of each program has been calculated as the mean value of the distinct criteria grades. Disagreements between evaluators were resolved through discussion.

Pre-test performance between the two groups has been compared using an independent samples t-test. No significant difference was found in the t-test performed (t(100)= 0,107, p=0,915) (mean pre-test scores Group1=2,5620, Group2=2,5375).

Subsequently, repeated General Linear Model (GLM) measures analysis of variance on students' performances in pre- and post-test, with Case (pre-test vs. post-test) as the within-subjects factor, and Case structure (Group1 vs. Group2) as the between-subjects factor has been conducted. The corresponding multivariate tests revealed that the Case factor contributes to the improvement of the post–test performance of both groups'. The "Case x Case structure" interaction yielded no significant effects (mean post-test scores Group1=2,995, Group2=2,743). The results of the repeated measures analysis are presented in Table 2.

Table 2. Multivariate tests for time factor and for the interaction Case x Case structure

Factor	F(1)	*p*
Case	11.68	0.001
Case x Case structure	1.48	0.226

As can be seen in Table 2, no evidence was found supporting the assumption that Case structure may affect the ability of students to develop programs that resolve similar (to the case) problems.

Data concerning the other two research questions has also been analyzed. In specific, this data concerned: a) the time students spent studying the Case in Phase B, b) the students' opinions on the structure and usefulness of Case. The data analysis showed that: a) there was a significant difference in the time students of the two groups spent studying the Case, with students of the Group 1 spending less time than those of Group 2, b) there was a significant difference of opinions of students of the two groups about the redundancy of the information contained in the Cases, with students of the Group 2 claiming more often that some parts of the Case they worked on should be omitted.

Second empirical study.

The second empirical study focused on the case authoring process aiming at the investigation of the difficulties pre-service Computer Science teachers face while acting as case authors. In this framework, the proposed case structures (Structure A and Structure B) have also been compared. Significant differences have been noticed in the frequency pre-service teachers select Structure A and Structure B templates, with Structure A being preferred by the majority of authors. Moreover, significantly less

problems in the case structure have been observed to cases structured according to Structure A than those structured according to Structure B.

Case template in CASTLE Programming prototype.

Evidence from both empirical studies described above has been considered in order to design the Case template for CASTLE Programming prototype. Namely, the selected template followed the Structure A outline containing the parts: *Problem* where the case problem is described, *Solution* where the solution program code is listed and *Explanation* where the most important decisions in the development of the solution program are commented.

3.2 Specification of case-based activities

Next, the specification of CASTLE Programming prototype proceeded with the selection of the case-based activities types. Ideas for the types of activities were found in the Didactics of Programming bibliography, including the Programming case studies structure proposals described above as well as digital learning environments specialized in Programming instruction.

Some of the selected types of activities engage learners to respond to questions after studying an entire case and others demonstrate only some of the case parts and require the completion of the case by the learners. For example, in *perturbation* activities, the entire case template (including the parts *problem*, *solution* and *explanation*) is presented to learners. Learners have to study the proposed solution and explanation thoroughly and solve a new problem themselves. The new problem is created by the alteration of the case problem conditions. On the contrary, in *explanation* activities only the *problem* and *solution* parts of the case are being presented to the learners who have to complete the explanation part themselves. In another set of activities, two or more cases are involved. For example, in *comparison* activities students have to study two entire cases and find similarities and differences in their solutions.

3.3 CASTLE Programming prototype outline

CASTLE layout.

The main screen of CASTLE is divided into three areas: the concept tree (Figure 1, A), the information frame (Figure 1, B) and the presentation tabs (Figure 1, C). Cases appear as leaves in the concept tree, appended from the concept they refer to. Once a user clicks on the name of a given case, the case unfolds into a new presentation tab. Simultaneously, a list of all activities based on this case appears in the information frame. A user may select an activity from the list in order to view it in a new presentation tab.

Case management.

Case authoring in CASTLE is a two-step activity. In the first step the author fills in the case name, the problem description, the problem constraints (if any), the main concept and (optionally) the group of cases containing isomorphic problems. In the second step the problem solution (code listing) together with the explanation of critical decisions are completed. This enables the submission of multiple solutions to a given problem description. Every solution is divided into parts which are labeled according to their function. Explanation is listed below the solution, following the same labels.

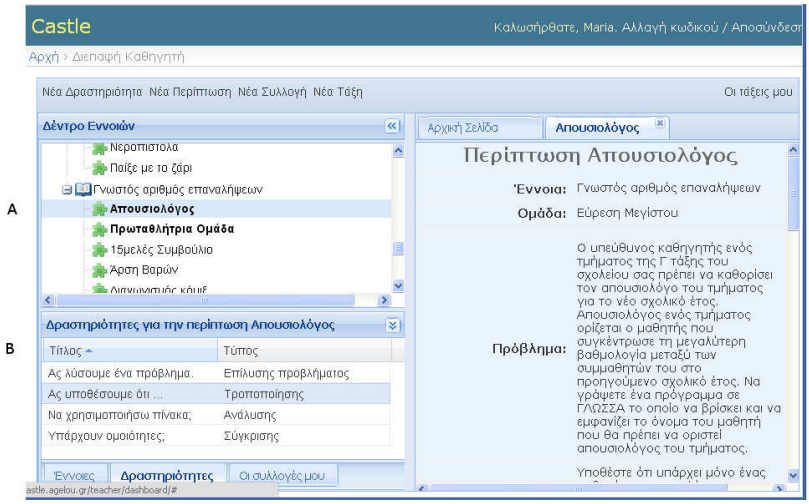

Fig. 1 CASTLE Snapshot where A. the concept tree, B. the information frame & C. the presentation tabs

Activity management.

CASTLE facilitates the case activity authoring by appropriately presenting and hiding parts of a case according to the selected activity type. An activity author should complete the following steps in order to create an activity:

1. Name the activity
2. Select the activity type out of the list of activities included in CASTLE
3. Select the case (or the cases) the activity is based on
4. Write the activity question
5. Provide information about: the estimated difficulty level (1 to 5) and the required elaboration time (in minutes).

An activity author may base his/her activities on cases previously authored by him/herself, but also on cases authored by others. If the adequate case is not already included in CASTLE, the author should first create the case and then proceed with the creation of the activity.

Comment exchange functionalities.

CASTLE users may exchange comments on the cases and activities included in CASTLE. This commenting mechanism provides authors with feedback on the learning material they have created thus enabling them to make improvements. In specific, a case author may receive both verbal comments and a numerical evaluation in a five-point scale (1 to 5). Activities authors, in addition to the verbal and numerical evaluation, receive feedback concerning the estimation of other users on the activity difficulty level (1 to 5) and the required elaboration time.

Other functionalities.

CASTLE supports the creation of *educational scenarios*, which are sequences of learning activities. In order to help teachers create their scenarios, CASTLE provides them the opportunity to organize sets of activities in *collections*. Moreover, a teacher is able to monitor and support the elaboration of scenarios by his/her students who work with CASTLE.

3.4 CASTLE Programming prototype evaluation

The prototype of CASTLE for Programming has been used and evaluated by a group of 66 pre-service programming teachers and a group of 17 expert programming teachers. Pre-service teachers were four-year students at the Department of Informatics and Telecommunications of the National and Kapodistrian University of Athens who had enrolled to the course "Introduction to Informatics and Telecommunications". Expert teachers were in-service teachers of Informatics with many years of teaching experience in secondary education. Teachers had to work with CASTLE prototype in order to create cases and case-based learning activities. Next, they were asked to submit comments on cases and case-based learning activities authored by their peers. At the same time, they had to fill in an evaluation questionnaire containing both Likert-scale and open-ended questions. The analysis of the teachers' answers in the questionnaires revealed that both groups of teachers had expressed positive opinions on CASTLE. In specific, teachers agreed that CASTLE manages to handle successfully cases and activities as distinct entities, supporting users adequately in finding activities based on a given case. Moreover, teachers of both groups found the comment exchange functionalities of CASTLE really useful, recognizing how important it is not only to receive but also to submit comments themselves. Few significant differences on the opinions of the two groups of students have been observed. For example, pre-service teachers responded that CASTLE is more adequate for beginner learners in Programming than it is for more advanced learners. Expert teachers do not share this opinion as they consider CASTLE equally useful for learners of all levels. Suggestions for the next versions of CASTLE have also been collected from the evaluation of the prototype. Evaluators indicated that help information should be upgraded. Finally, interesting ideas for layout improvements have been reported and will be taken into consideration.

4 Specification process of CASTLE in Didactics.

The specification of CASTLE in Didactics has also been investigated in the framework of this dissertation. Many examples of case utilization in teacher education have been reported in the last years. An empirical study was conducted to examine whether the way a case is being studied (through an online learning environment or in class) affects the ability of teachers to design learning activities [31]. No significant differences have been reported. The examination of how case-based activities of Didactics can be implemented through a digital learning environment provided evidence that CASTLE design is adequate for this subject matter as well.

5 Conclusions

The review of relevant bibliography indicated the need and usefulness of a learning environment that enables the authoring of learning activities of different types based on one case, as this remains an open issue in the domain of digital case-based learning environments. There is a large variety of learning activities that may be created on a given case, including activities both for novices and for advanced learners. Additionally, case authoring is a demanding task. These facts stress the importance of case reuse in different types of activities.

In the framework of the present dissertation the architectural design of the web-based case-based learning environment CASTLE (CASes for Teaching and LEarning) is presented. CASTLE handles cases and case-based learning activities as distinct entities. An author may create his/her own cases in CASTLE or find interesting cases created by others and use them to create different types learning activities. Some types of learning activities in CASTLE may refer to two or more cases (e.g. comparison activities). Additionally, CASTLE enables the creation and elaboration of educational scenarios that are formed as sequences of case-based learning activities. Finally, comment exchange on cases, learning activities and educational scenarios is supported, enabling authors to constantly improve their learning material.

The specification of the generic design of CASTLE in two subject matters has been also examined. Programming and Didactics have been selected as representative subject matters that contain mainly well-structured and ill-structured problems respectively. The investigation showed that CASTLE may be specified in both subject matters. In Programming, the investigation went further and a prototype of CASTLE for Programming teaching has been developed. The Programming case template used in CASTLE prototype has been selected after the experimental comparison of two templates found in bibliography and is proposed. Finally, CASTLE prototype has been evaluated by pre-service and expert teachers. Both groups of teachers expressed overall satisfaction with CASTLE, found that it succeeds to support users to find the learning activities based on a given case and recognize the importance of receiving and submitting comments on the learning material.

References

1. Jonassen, D. H.: Instructional design model for well-structured and ill-structured problem-solving learning outcomes. Educational Technology Research and Development, vol 45(1), 65-95 (1997)
2. Jonassen, D. H.: What is problem solving. In Jonassen, D. H. (Ed.) Learning to solve problems: instructional design guide, pp 1-18. NJ: Lawrence Erlbaum Associates (2003b)
3. Brown, J.S., Collins, A., Duguid, P. : Situated cognition and the culture of learning. Educational Researcher, 18, 32-42 (1989)
4. Vosniadou, S. How children learn. Educational Practices Series, The International Academy of Education (IAE) and the International Bureau of Education (UNESCO) (2001).
5. Herreid, C. F.: Start with a story: The case study method of teaching college science. NSTA Press: Arlington, VA (2006)
6. Kolodner, J.: Case-Based Reasoning. San Mateo, CA: Morgan Kaufmann Publishers Inc. (1993)
7. Spiro, R. J., Coulson, R. L., Feltovich, P. J., Anderson, D. K.: Cognitive flexibility theory: Advanced knowledge acquisition in ill-structured domains. Technical Report No. 441. Champaign, IL: University of Illinois, Center for the Study of Reading (1988)
8. Golich, V.L., Boyer, M., Franko, P., Lamy, S.: The ABC of Case Teaching, Pew Case Studies in International Affairs. Institute for the Study of Diplomacy, Georgetown University, Washington D. C. (2000)
9. Jonassen, D.H.: Research Issues in Problem Solving. In 11th International Conference on Education Research. (2010)
10. Gentner, D., Loewenstein, J., Thompson, L.: Analogical encoding: Facilitating Knowledge transfer and integration. In Forbus, K. D., Gentner, D., Reiger, T. (Eds.) Proceedings of the 26th Annual Conference of the Cognitive Science Society, pp 452–457.Cognitive Science Society, Austin, TX (2004)
11. Aamodt, A., Plaza, E.: Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. AI Communications. IOS Press, vol7 (1), 39-59 (1994)
12. Kolodner, J.L.: An introduction to case-based reasoning. Artificial Intelligence Review, vol 6(1), 3-34 (1992)
13. Maher, M. L., Balachandran, M. B., Zhang, D.: Case-Based Reasoning in Design. Lawrence Erlbaum, Hillsdale, NJ. Lawrence Erlbaum Associates(1995)
14. Derry, S. J., the TiPS Research Group: Development and Assessment of Tutorials in Problem Solving (TiPS): A Remedial Mathematics Tutor. Final Report to the Office of Naval Research (N00014-93-1-0310), Wisconsin Center for Education Research. University of-Wisconsin-Madison, Madison, WI (2001)
15. Jonassen, D. H.: Designing research-based instruction for story problems. Educational Psychology Review, vol 15(3), 267-296 (2003a)
16. Zimring, C.M., Bafna, S., Do, E.: Structuring cases in a case based design aid. Vanegas J., Chinowsky. P. (eds), pp. 308-313 (1996)
17. Carroll, J.M., Rosson, M.B.: Case studies as minimalist information. IEEE Transactions on Professional Communication, vol49(4), 297-311 (2006)
18. Xiao, L., Carroll, J., Rosson, M.: Support Case-Based Authentic Learning Activities: A Collaborative Case Commenting Tool and a Collaborative Case Builder. Paper presented at the Human-Computer Interaction. HCI Applications and Services, pp. 371-380. Beijing, China (2007)

19. Riedel, J., Fitzgerald, G., Leven, F., Toenshoff, B.: The Design of Computerized Practice Fields for Problem Solving and Contextualized Transfer. Journal of Educational Multimedia and Hypermedia, 12(4), 377-398. Norfolk, VA: AACE (2003).
20. Rocha, A.R., Moreira, G., Campos, F., Rabelo, L.: CardioCaseDiscussion: a cooperative learning environment for patients' cases discussion in Cardiology. In Barker, P., Rebelsky, S. (Eds.) Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications 2002, pp. 1371-1373. Chesapeake, VA: AACE (2002)
21. Choi, I., Lee, S. Jung, J.: Designing Multimedia Case-Based Instruction accommodating Students' Diverse Learning Styles. Jurnal of Educational and Educational Multimedia and Hypermedia. Vol17(1), 5-25 (2008)
22. Leclet, D., Joiron, C.: Design of a Distance Learning Environment - DIACOM : An Interactive Forum Based on Collaborative Learning for Continuing Medical Eduaction. In Barker, P., Rebelsky, S. (Eds.) Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications 2002, pp 876-881. Chesapeake, VA: AACE (2002)
23. Rudy, M., Jaksch, S.: Building Science by Design: Developing a Building Information System and Teaching Architecture with Analytical Case Studies. ED-MEDIA 2004: World Conference on Educational Multimedia, Hypermedia & Telecommunications, Association for the Advancement of Computing in Education, pp. 4815-4820. Norfolk (VA) (2004)
24. Reed, S. K.: A structure-mapping model for word problems. Journal of Experimental Psychology: Learning, Memory, and Cognition, 13, 124–139 (1987)
25. Shirky, C.: A Group is Its Own Worst Enemy. Networks, Economics, and Culture mailing list, Retrieved 21 2012 from:. http://www.shirky.com/writings/group_enemy.html (2003)
26. Bull, S.: Supporting Learning with Open Learner Models. In Proceedings of 4th Hellenic Conference in Information and Communication Technologies in Education, pp. 47-61. Athens (2004)
27. Bull, S., Kay, J. : A Framework for Designing and Analysing Open Learner Modelling. Proceedings of Workshop on Learner Modelling for Reflection. In International Conference on Artificial Intelligence in Education, pp. 81-90 (2005)
28. Boubouka, M., Verginis, I., Grigoriadou, M. & Zafiri, I. (2009) "Towards the enhancement of the learning process with different types of case based activities" in Proceedings of the 8th IEEE International Conference on Advanced Learning Technologies (ICALT2009), Riga, Latvia, 15-17 July 2009, pp 652-654.
29. Linn, M., Clancy, M.: The case for case studies of programming problems. Communications of the ACM, vol 35(3), 121-132 (1992)
30. Spooner, D. & Skolnick, M.: Science and Engineering Case Studies in Introductory Computing Course for Non-Majors, pp. 154-158. SIGCSE CA (1997)
31. Boubouka M., Verginis I., Grigoriadou M.. Supporting the Implementation of Case Activities using e-learning. In Spector M. J., Ifenthaler, D., Isaías, P., Kinshuk, & Demetrios, G. Sampson (Eds.), *Learning and Instruction in the Digital Age Technologies*, Springer, 2009, ISBN 978-1-4419-1550-4 (2009).

Distributed RDF Query Processing and Reasoning in Peer-to-Peer Networks

Zoi Kaoudi*

Dept. of Informatics and Telecommunications
National and Kapodistrian University of Athens, Greece

Abstract. With the vast amount of available RDF data sources on the Web increasing rapidly, there is an urgent need for RDF data management and RDFS reasoning. In this thesis, we focus on distributed RDF data management in peer-to-peer (P2P) networks. More specifically, we present results that advance the state-of-the-art in the research area of distributed RDF query processing and reasoning in P2P networks. We fully design and implement a P2P system, called Atlas, for the distributed query processing and reasoning of RDF and RDFS data. Atlas is built on top of distributed hash tables (DHTs), a commonly-used case of P2P networks. Initially, we study RDFS reasoning algorithms on top of DHTs. We design and develop distributed forward and backward chaining algorithms, as well as an algorithm which works in a bottom-up fashion using the magic sets transformation technique. We study theoretically the correctness of our reasoning algorithms and prove that they are sound and complete. We also provide a comparative study of our algorithms both analytically and experimentally. In the experimental part of our study, we obtain measurements in the realistic large-scale distributed environment of PlanetLab as well as in the more controlled environment of a local cluster. Moreover, we propose algorithms for SPARQL query processing and optimization over RDF(S) databases stored on top of distributed hash tables. We fully implement and evaluate a DHT-based optimizer. The goal of the optimizer is to minimize the time for answering a query as well as the bandwidth consumed during the query evaluation. The optimization algorithms use selectivity estimates to determine the chosen query plan. Our algorithms and techniques have been extensively evaluated in a local cluster.

1 Introduction

More than just a vision nowadays, the Semantic Web has begun to be realized by the publication of large datasets according to the principles of the Linked Data initiative[1]. The Linked Data initiative aims at connecting data sources on the Web and exposing real life data using semantic technologies offering a new way of data integration and interoperability. The result of this effort is a Web

* Dissertation Advisor: Manolis Koubarakis, Professor
[1] `http://linkeddata.org/`

of Data, where URIs identify real life things, dereferencing URIs returns RDF information about those things, and this RDF information contains related URIs which are links to other resources enabling further exploration. The Linked Data community have established a set of best practices for collaboratively publishing and interlinking structured data on the Web. There are numerous sources that expose their data on the Web in the form of Linked Data ranging from community-driven efforts to governmental bodies or scientific groups. DBpedia[2], BBC music information [18], open government data[3] are only a few examples of the constantly increasing Linked Data cloud[4].

With the vast amount of available RDF data sources on the Web increasing rapidly, there is an urgent need for RDF data management. RDF storage, query processing and reasoning have been at the center of attention during the last years in the Semantic Web community and more recently in other research fields as well. Many systems have been developed for storing and querying RDF data. The first attempts were centralized approaches, such as Jena [32], Sesame [4] and RSSDB [1]. However, managing the avalanche of available RDF data has become a challenge for such RDF stores. This has necessitated the careful performance evaluation of existing RDF stores on appropriately designed benchmarks and very big data sets and the development of novel implementations based on efficient indexing techniques and relational-style statistics-based query optimization [22, 31]. Performance results published very recently indicate that state-of-the-art systems like RDF-3X [22] can execute complex join queries on RDF data sets containing close to a billion triples in a few seconds.

Although some existing RDF stores have excellent performance, they can be overwhelmed by user requests when used in wide-area network applications such as content-sharing, Web/Grid service registries, distributed digital libraries and social networks such as the ones discussed in [16, 27] More generally, since centralized RDF stores are lacking the reliability properties typically associated with large distributed systems, (e.g., fault-tolerance, load balancing, availability), researchers have also studied parallel and distributed solutions for RDF and RDFS query processing and reasoning. These include solutions based on peer-to-peer (P2P) systems, distributed computing platforms built on powerful clusters and, more recently, cloud computing platforms using the MapReduce framework [6].

In this thesis we concentrate on RDF query processing and reasoning using P2P networks. The results of this thesis have been published in major international conferences [14, 15] of the Semantic Web community, workshops [12, 17] and journals [13, 16].

[2] `http://dbpedia.org`
[3] `http://www.data.gov/`, `http://data.gov.uk/`
[4] `http://www4.wiwiss.fu-berlin.de/lodcloud/state/`

2 Dissertation Summary

In a setting where several heterogeneous sources of data are geographically distributed, P2P systems enable the aggregation and integration of these data sources in an efficient way. P2P networks have gained much attention in the last ten years, given all the good features they can provide to Internet-scale applications. There have been several proposals of P2P architectures and amongst them distributed hash tables (DHTs) [2] are the most prominent class. DHTs allow for *full distribution*, *high-performance*, *scalability*, *resilience to failures*, *robustness* and *adaptivity* in applications such as distributed digital libraries and others we mentioned above.

An RDF repository built on top of a DHT simplifies the integration of data from many distributed heterogeneous data sources compared to other distributed approaches. Additionally, DHTs can be used to ensure *efficient* query answering from all these heterogeneous sources. DHTs have been proposed for the storage and querying of RDF data at Internet scale by several works such as [5, 11, 19]. This thesis also focuses on DHTs as the P2P architecture of choice. Thus, the algorithms of this thesis can run on commodity machines deployed all over the world, as it is the case with many other P2P applications. This is in contrast to other distributed approaches that rely on distributed platforms built on powerful clusters such as [7, 9, 23] and cloud computing platforms using MapReduce [20, 28], which typically demand high-end, locally deployed infrastructures whose cost can be very high in many cases.

When designing a DHT-based system for RDF data management, there are several challenges that have to be faced. The first one is how to distribute the data among the nodes of the network. DHTs utilize an efficient protocol for indexing data items in the network and thus, an indexing scheme for RDF data that conforms with this protocol may be adopted. Another issue that has to be faced is whether RDF data and RDFS ontologies should be handled uniformly or RDFS ontologies should be globally known by all nodes.

The adopted storage scheme will help us deal with the second challenge: answering SPARQL queries efficiently. A key aspect here is to design efficient query processing algorithms which are able to combine RDF data distributed across different nodes of the network to answer user queries. Another important need in Semantic Web applications is modeling application knowledge and reasoning about this knowledge. Therefore, in the context of RDF, we have to deal not only with a huge amount of distributed data, but also with a set of RDFS ontologies that give meaning to this data. Naturally, SPARQL queries need to be answered in a way that take into account both RDF data and RDFS ontologies, as well as the RDFS entailment rules given in [10].

Another challenging issue for RDF data management in a DHT environment is query optimization. RDF query optimization techniques in a DHT-based system have to be carefully considered given that data are distributed across all nodes of the network. Although query optimization has been extensively studied in the database area and is widely used in modern DBMSs, RDF query optimization has been addressed only recently even in centralized environments [21,22,26].

The contributions of this thesis are summarized in the following paragraphs.

In this thesis we fully design and implement a DHT-based system for the distributed query processing and reasoning of RDF and RDFS data. The indexing scheme we deploy in our system is the triple indexing algorithm originally presented in [5] where each RDF triple is indexed in the DHT three times. An important aspect of our indexing scheme is that data and schema information is handled uniformly. Although other distributed approaches such as [28, 30] assume that each node keeps all RDF schema information, we adopt a more generic approach where no global knowledge about the schema is required. In this way, our system can also handle scenarios with very big ontologies where other systems such as the above might not scale.

With respect to RDFS reasoning, our contribution is the design and development of distributed *forward* and *backward* chaining algorithms on top of a DHT. The forward chaining (FC) approach has minimal requirements during query answering, but needs a significant amount of storage for all the inferred data. In contrast, the backward chaining (BC) approach has minimal storage requirements, at the cost of an increase in query response time. There is a time-space trade-off between these two approaches, and only by knowing the query and update workload of an application, we can determine which approach would suit it better. This trade-off has never been studied in detail in a *distributed Web-scale* scenario and this is a challenge we undertake. Our backward chaining algorithm is the first distributed top-down algorithm proposed for RDFS reasoning in a decentralized environment in general. Current forward chaining approaches in various distributed architectures demonstrate a big rate of redundant information occurred from the inferred RDF triples [28, 30]. Our forward chaining algorithm (FC*) is the first one that deals with an important case of generating redundant RDF information. In addition, we present an algorithm (MS) which works in a bottom-up fashion using the magic sets transformation technique [3], a technique that has not been studied in the literature for distributed RDFS reasoning. We study theoretically the correctness of our reasoning algorithms and prove that they are sound and complete. We also provide a comparative study of our algorithms both analytically and experimentally. In the experimental part of our study, we obtain measurements in the realistic large-scale distributed environment of PlanetLab as well as in the more controlled environment of a local cluster.

We propose a query processing algorithm adapted to use a query graph model to represent SPARQL queries in order to avoid the computation of Cartesian products. In addition, as URIs and literals may consist of long strings that are transferred in the network and processed locally at the nodes, we show how to benefit from a mapping dictionary to further enhance the efficiency of the query processing algorithm. Although mapping dictionaries are by now standard in centralized RDF stores, our work is the first that discusses how to implement one in a DHT environment. Our experiments conducted in both PlanetLab and a local cluster showcase the importance of having a distributed mapping dictionary in our system.

In the context of query optimization, we fully implement and evaluate a DHT-based optimizer. The goal of the optimizer is to minimize the time for answering a query as well as the bandwidth consumed during the query evaluation. We propose three greedy optimization algorithms for this purpose: two static, namely NA and SNA, and one dynamic, namely DA. The static query optimization is completely executed before the query evaluation begins, while the dynamic query optimization take places during the query evaluation creating query plans incrementally. These algorithms use selectivity estimates to determine the chosen query plan. We propose methods for estimating the selectivity of RDF queries utilizing techniques from relational databases. We discuss which statistics should be kept at each network node and use histograms for summarizing data distributions. We demonstrate that it is sufficient for a node to create and maintain local statistics, i.e., statistics for its locally stored data. These local statistics are in fact global statistics needed by the optimization algorithms and can be obtained by other nodes by sending low cost messages. This is a very good property of the indexing scheme we adopt from [5] that has not been pointed out in the literature before.

Using the above techniques, we have implemented a P2P system, called Atlas, for distributed query processing and reasoning of RDF and RDFS data. Atlas is publicly available as open source under the LGPL license[5]. Although our proposed algorithms and techniques have been implemented in Atlas using the Bamboo DHT [24], they are DHT-agnostic; they can be implemented on top of any DHT.

3 Results and Discussion

In this section, we present a brief experimental evaluation of our reasoning algorithms and optimization techniques. All algorithms have been implemented as an extension to our prototype system Atlas. In the latest version of Atlas, we have adopted SQLite as the local database of each peer since the Berkeley DB included in the Bamboo implementation was inefficient. In our algorithms, we have also utilized the dictionary encoding implemented in Atlas, where URIs and literals are mapped to integer identifiers. For our experiments, we used as a testbed both the PlanetLab network as well as a local shared cluster (`http://www.grid.tuc.gr/`). Although we have extensively tested our techniques on both testbeds, here we present results only from the cluster where we achieve much better performance. The cluster consists of 41 computing nodes, each one being a server blade machine with two processors at 2.6GHz and 4GB memory. We used 30 of these machines where we run up to 4 peers per machine, i.e., 120 peers in total.

For our evaluation, we use the Lehigh University benchmark (LUBM) [8] that provides synthetic RDF datasets of arbitrary sizes and 14 SPARQL queries. LUBM benchmark consists of a university domain ontology modeling an academic setting and is widely used for testing RDF stores. Each dataset can be defined by the number of universities generated. For example, the dataset LUBM-1

[5] `http://atlas.di.uoa.gr`

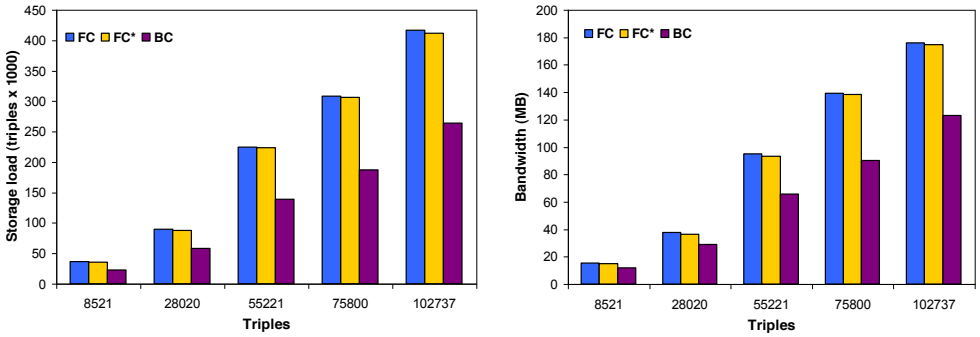

(a) Message storage load

(b) Bandwidth

Fig. 1. Storing LUBM-1

involves one university, while the dataset LUBM-10 involves 10 universities. The more universities are generated the more triples are produced. Since answering a query in the LUBM benchmark might involve RDFS reasoning, All measurements are averaged over 10 runs using the geometric mean which is more resilient to outliers.

3.1 Comparing the reasoning algorithms

In this section, we compare the performance of the forward chaining algorithm (FC) with the backward chaining algorithm (BC) when storing RDF(S) data in the network. Alongside the results of the forward chaining algorithm we have presented in [15], we also present results from the forward chaining algorithm which generates less redundant information (FC*) and the algorithm that uses the magic sets transformation (MS).

For this set of experiments, we stored an increasing number of triples from the LUBM-1 dataset which consists of 102, 737 triples. We compare the behaviour of BC, FC and FC*. Figure 1 shows results regarding the store message load and the bandwidth consumption. The x-axis shows the number of triples initially inserted in the network. The results of this experiment are similar regarding the comparison of BC with FC and FC*. However, in this case, FC and FC* produce almost the same number of messages resulting in similar bandwidth consumption and cause almost the same load in the network, with FC* performance slightly better. The difference between FC and FC* in this case is not so evident because of the nature of the LUBM schema. LUBM schema includes only small RDFS class hierarchies whose depth is at most 2. Therefore, the number of redundant triples generated from FC is very small.

In the next experiment, we use the LUBM dataset whose schema contains several independent class hierarchies. We created a network of 156 nodes in the cluster and stored the complete LUBM-20 dataset consisting of 2, 782, 435 triples. We want to retrieve the instances of the following classes of the LUBM

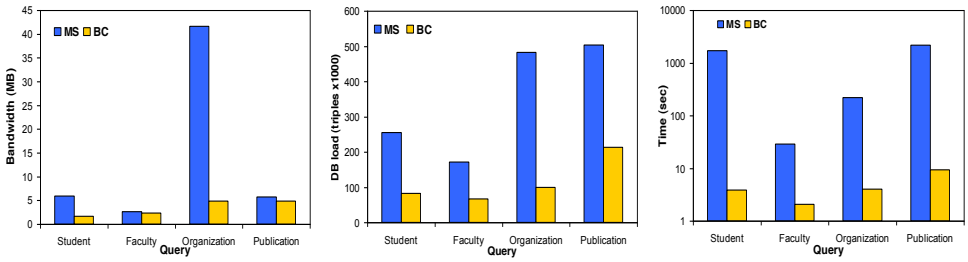

(a) Bandwidth consumption (b) DB load (c) Time

Fig. 2. MS vs. BC for LUBM-20

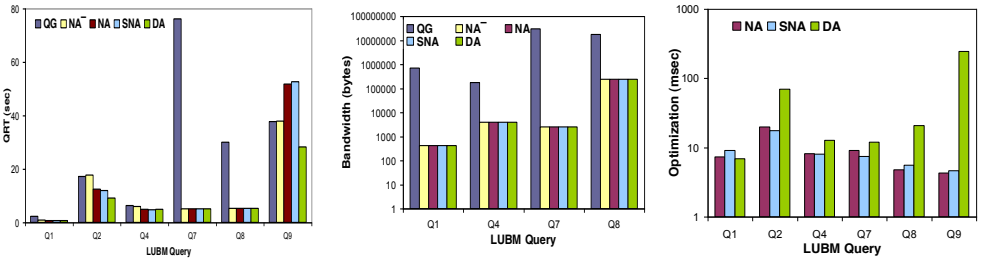

(a) Total query response time (b) Total bandwidth usage (c) Optimization time

Fig. 3. Query optimization performance for LUBM-50

schema: `ub:Student`, `ub:Faculty`, `ub:Organization`, `ub:Publication`, where `ub` is the appropriate namespace. For MS, we sent a request with the predicate `m_type` and argument the class name, while for BC we run the query that asks for all instances of the respective class. Figure 2 shows the bandwidth consumption, the total DB load and the time required for each algorithm to terminate. In this experiment as well, we observe that BC outperforms MS. Since we deal with a bigger dataset the advantage of using BC is more evident, as shown in Figure 2(c).

3.2 Comparing the optimization algorithms

In the following, QG denotes that the query graph is used to avoid Cartesian products but no other optimization is utilized. The naive algorithm using the bound-is-easier heuristic is denoted by NA^-, while the naive and semi-naive algorithm using the analytical estimation is denoted by NA and SNA, respectively. Finally, DA denotes the dynamic optimization algorithm. Details about these algorithms can be found in [14]. In this section, we compare and evaluate the optimization algorithms. For this set of experiments, we store all the inferred triples of the LUBM-50 dataset ($9,437,221$ triples) in a network of 120 peers. Then, using each optimization algorithm, we run the queries.

In all graphs of Fig. 3, the x-axis shows the LUBM queries while the y-axis depicts the metric of interest. Figure 3(a) shows the query response time (QRT)

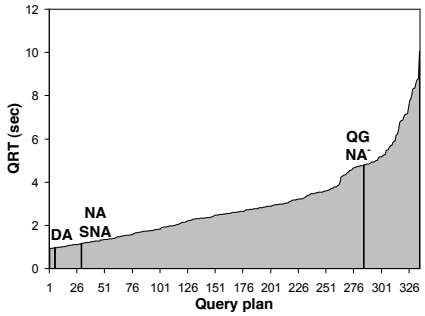

(a) LUBM Q2 query plan space

LUBM Query	Min QRT	Max QRT	DA QRT
Q2	0.91 s	10.06 s	0.96 s
Q4	2.54 s	17.39 s	2.89 s
Q7	1.53 s	16.62 s	1.55 s
Q8	2.88 s	10.20 s	3.06 s
Q9	3.30 s	64.06 s	4.41 s

(b) QRT of LUBM query plans

Fig. 4. Exploring the query plan space for LUBM-10

for the different LUBM queries. QRT is the total time required to answer a query and it also includes the time required by the query optimizer for determining a query plan (optimization time). Figure 3(b) shows the total bandwidth consumed during query evaluation.

Queries Q2 and Q9 consist of 6 triple patterns having only their predicates bound. In both queries, there exists a join among the last three triple patterns (in the order given by the benchmark) and the combination of all three triple patterns is the one that yields a small result set. DA finds a query plan that combines these three triple patterns earlier than the other algorithms. This results in producing smaller intermediate result sets, as it is also shown by the bandwidth consumption in Fig. 3(b), and thus results in better QRT. Although NA and SNA perform close to DA for query Q2, they fail to choose a good query plan for Q9 affecting both the QRT and the bandwidth consumption. At this point, we should note that QG and NA^- depend on the initial order of a query's triple patterns. For this reason, both algorithms choose a relatively good query plan for query Q9 since the order in which its triple patterns are given by the benchmark is a good one. Q4 is a star-shape query with all its triple patterns sharing the same subject variable, while only the first two triple patterns have two bound components. Therefore, since these two triple patterns are the more selective ones, all optimization algorithms choose the same query plan and perform identically in terms of both QRT and bandwidth. The same holds for query Q7 where QRT is significantly reduced when using either optimization technique compared to QG. Q8 is a query similar to Q7.

In Fig. 3(c), we show the total optimization time in msec on a logarithmic scale. For QG and NA^- the optimization overhead is negligible and is not shown in the graph. The optimization time contains the time for retrieving the required statistics from the network, the time for the selectivity estimation and the time spent by the optimization algorithm. As expected, DA spends more time than the other optimization algorithms since it runs at each query processing step. However, the optimization time is still one order of magnitude smaller than the time required by the query evaluation process. Therefore, although DA requires more time than the other optimization algorithms, the system manages to per-

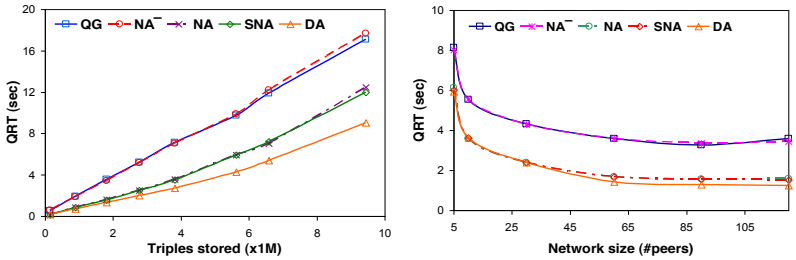

(a) Varying dataset size for query Q2 (b) Varying network size for query Q2

Fig. 5. Varying dataset and network size

form efficiently for all queries when DA is used. We observe similar results for the bandwidth consumed by the query optimizer. NA and SNA consume $\sim 2KB$ while DA consumes $\sim 7KB$, still one order of magnitude less than the bandwidth spent during query evaluation. We omit this graph due to space limitations.

3.3 Effectiveness of query optimization

In this section, we explore the query plan space of the LUBM queries to show how effective the optimization algorithms are. The size of the query plan space of a query consisting of N triple patterns is $N!$. Since query plans that involve Cartesian products are very inefficient to evaluate in a distributed environment, we consider only triple pattern permutations which do not produce any Cartesian product. In this experiment, we store the LUBM-10 dataset in a network of 120 peers and run all possible query plans for several LUBM queries. In Fig. 4(a), we depict the QRT of all possible query plans for query Q2 in ascending order. The query plan space of Q2 consists of 335 query plans which do not involve any Cartesian product. In this figure, we highlight the position of the query plans chosen by the different optimization algorithms. We observe that DA chooses one of the best query plans, while NA and SNA perform worse choosing the 27th best query plan. NA^- performs poorly choosing one the worst query plans. Similar results are observed for the other queries as well. In Fig. 4(b), we list the QRT for all queries of the best and the worst query plan together with the QRT when using DA. We observe that the QRT when using DA is very close to the QRT of the optimal query plan for all queries. Note that without the query plans that involve Cartesian products, the difference between the min and the max QRT of all queries is not very large.

3.4 Varying the dataset and network size

In these sets of experiments, we study the performance of our system when varying the number of triples stored in the network and the number of peers. We show results only for Q2 which involves a join among three triple patterns.

Figure 5(a) shows the behavior of our system using each optimization algorithm as the dataset stored in the network grows. In a network of 120 peers, we stored datasets from LUBM-1 to LUBM-50. Every time we measured the QRT of query Q2 using each optimization algorithm. As expected, QRT increases as the number of triples stored in the network grows. This is caused by two factors.

Firstly, the local database of each peer grows and as a result local query processing becomes more time-consuming. Secondly, the result set of query Q2 varies as the dataset changes. For example, for LUBM-1 the result set is empty, while for LUBM-50 the result set contains 130 answers. This results in transferring larger intermediate result sets through the network which also affects the QRT of the query. Besides, this experiment brings forth an interesting conclusion regarding the optimization techniques. While query plans chosen by NA, SNA and DA perform similarly up to approximately $1.8M$ triples stored (i.e., LUBM-10), we observe that for bigger datasets the query plan chosen by DA outperforms the others. This shows that the system becomes more scalable with respect to the number of triples stored in the network when using DA. Similar results are observed for Q9, while for the rest queries all optimization algorithms choose the same query plan independently of the dataset size.

In the next set of experiments, we start networks of 5, 10, 30, 60, 90 and 120 peers and store the LUBM-10 dataset. We then run the queries using all optimization techniques. In Fig. 5(b), we show the QRT for Q2 as the network size increases. We observe that QRT improves significantly as the network size grows up to 60, while it remains almost the same for bigger network sizes. The decrease in the QRT for small networks is caused by the fact that the more peers join the network the less triples are stored in each peer's database and thus local processing load is reduced. The same result was observed in other queries where QRT either improved or remained unaffected as the number of peers increased.

3.5 Discussion

We have also experimented with different datasets using the SP^{2}B benchmark [25] as well as a real world dataset of the US Congress vote results presented in [29]. The results were similar to the ones observed using LUBM. For all datasets, DA consistently chooses a query plan close to the optimal regardless of the query type or dataset stored and without posing a significant overhead neither to the total time for answering the query nor to the bandwidth consumed. On the contrary, the static optimization methods are dependent on the type of the query and the dataset, which make them unsuitable in various cases (as shown earlier for query Q9). In addition, we have also tested indexing all possible combinations of the triples' components, as proposed in [19]. In this case, we have used histograms at each peer for combinations of triples' components as well. However, we did not observe any difference in the choice of the query plan and thus, showed results only with the triple indexing algorithm. This results from the nature of the LUBM queries which mostly involve bound predicates and object-classes for which we kept an exact distribution in both cases.

4 Conclusions

We presented a system that is able to support full-fledged management of RDF data in a large-scale decentralized environment. In particular, we fully designed

and implemented a P2P system, called Atlas, for the distributed query processing and reasoning of RDF and RDFS data. One of the main focus of this thesis was to enable Atlas to support RDFS reasoning. Two well-known reasoning techniques we studied are *forward chaining* and *backward chaining*. Our work was the first that investigated the trade-off of the two algorithms in detail in a *distributed Web-scale* scenario. Our forward chaining algorithm was the first one that dealt with an important case of generation of redundant RDF information, while our backward chaining algorithm was the first distributed backward chaining algorithm proposed for RDFS reasoning in a decentralized environment in general. Moreover, a magic sets transformation technique for distributed RDFS reasoning has not been studied in the literature before.

Finally, this thesis addressed the problem of query optimization over RDF data stored on top of DHTs and fully implemented and evaluated a DHT-based optimizer. Our optimization algorithms ranged from static to dynamic ones and focused on minimizing the size of the intermediate results computed during query evaluation. In this way, we achieved to decrease the time required for answering a query and the bandwidth consumed during the query evaluation. Our algorithms utilized selectivity estimates to determine a query plan that minimizes the size of the intermediate results. We also proposed methods for estimating the selectivity of RDF queries utilizing techniques from relational databases. We defined which statistics are required at each node of the network for the computation of the selectivity estimates and used histograms for summarizing the distribution of these statistics.

References

1. S. Alexaki, V. Christophides, G. Karvounarakis, and D. Plexousakis. On Storing Voluminous RDF Descriptions: The case of Web Portal Catalogs. In *Proceedings of the 4rth International Workshop on the Web and Databases (WebDB 2001, co-located with SIGMOD 2001)*, Santa Barbara, California, USA, 2001.
2. H. Balakrishnan, M. F. Kaashoek, D. R. Karger, R. Morris, and I. Stoica. Looking up Data in P2P Systems. *Communications of the ACM*, 46(2):43–48, 2003.
3. C. Beeri and R. Ramakrishnan. On the Power of Magic. In *PODS '87: Proceedings of the sixth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 269–284, New York, NY, USA, 1987. ACM.
4. J. Broekstra, A. Kampman, and F. van Harmelen. Sesame: A Generic Architecture for Storing and Querying RDF and RDF Schema. In *Proceedings of the 1st International Semantic Web Conference (ISWC 2002)*, Sardinia, Italy, 2002.
5. M. Cai and M. Frank. RDFPeers: A Scalable Distributed RDF Repository based on A Structured Peer-to-Peer Network. In *Proceedings of the 13th World Wide Web Conference (WWW 2004)*, New York, USA, 2004.
6. J. Dean and S. Ghemawat. Mapreduce: Simplified Data Processing on Large Clusters. In *Proceedings of the USENIX Symposium on Operating Systems Design & Implementation (OSDI)*, pages 137–147, 2004.
7. O. Erling and I. Mikhailov. Towards Web Scale RDF. In *Proceeindgs of the 4th International Workshop on Scalable Semantic Web knowledge Base Systems (SSWS2008)*, Karlsruhe, Germany, October 2008.
8. Y. Guo, Z. Pan, and J. Heflin. LUBM: A Benchmark for OWL Knowledge Base Systems. *Journal of Web Semantics*, 3(2-3):158–182, 2005.
9. A. Harth, J. Umbrich, A. Hogan, and S. Decker. YARS2: A Federated Repository for Querying Graph Structured Data from the Web. In *Proceedings of the 6th International Semantic Web Conference (ISWC/ASWC 2007)*, pages 211–224, Busan, South Korea, 2007.
10. P. Hayes. RDF Semantics. W3C Recommendation, February 2004. `http://www.w3.org/TR/rdf-mt/`.

11. F. Heine, M. Hovestadt, and O. Kao. Processing Complex RDF Queries over P2P Networks. In *Proceedings of Workshop on Information Retrieval in Peer-to-Peer-Networks (P2PIR 2005)*, Bremen, Germany, 2005.
12. Z. Kaoudi, M. Koubarakis, K. Kyzirakos, M. Magiridou, I. Miliaraki, and A. Papadakis-Pesaresi. Publishing, Discovering and Updating Semantic Grid Resources using DHTs. In *CoreGRID Workshop on Grid Programming Model, Grid and P2P Systems Architecture, Grid Systems, Tools and Environments*, Heraklion, Crete, Greece, June 2007.
13. Z. Kaoudi, M. Koubarakis, K. Kyzirakos, I. Miliaraki, M. Magiridou, and A. Papadakis-Pesaresi. Atlas: Storing, Updating and Querying RDF(S) Data on Top of DHTs. *Journal of Web Semantics*, 2010.
14. Z. Kaoudi, K. Kyzirakos, and M. Koubarakis. SPARQL Query Optimization on Top of DHTs. In *Proceedings of the 9th International Conference on The Semantic Web (ISWC 2010)*, Shanghai, China, 2010.
15. Z. Kaoudi, I. Miliaraki, and M. Koubarakis. RDFS Reasoning and Query Answering on Top of DHTs. In *Proceedings of the 7th International Conference on The Semantic Web (ISWC 2008)*, Karlsruhe, Germany, 2008.
16. Z. Kaoudi, I. Miliaraki, M. Magiridou, E. Liarou, S. Idreos, and M. Koubarakis. Semantic Grid Resource Discovery in Atlas. *Knowledge and Data Management in Grids*, 2006. Talia Domenico and Bilas Angelos and Dikaiakos Marios D. (editors), Springer.
17. Z. Kaoudi, I. Miliaraki, M. Magiridou, A. Papadakis-Pesaresi, and M. Koubarakis. Storing and Querying RDF Data in Atlas. 3rd European Semantic Web Conference, 11 - 14 June 2006, Budva, Montenegro, Demo paper, 2006.
18. G. Kobilarov, T. Scott, Y. Raimond, S. Oliver, C. Sizemore, M. Smethurst, C. Bizer, and R. Lee. Media Meets Semantic Web - How the BBC Uses DBpedia and Linked Data to Make Connections. In *Proceedings of the 6th European Semantic Web Conference (ESWC)*, Heraklion, Crete, Greece, 2009.
19. E. Liarou, S. Idreos, and M. Koubarakis. Evaluating Conjunctive Triple Pattern Queries over Large Structured Overlay Networks. In *Proceedings of 5th the International Semantic Web Conference (ISWC 2006)*, Athens, GA, USA, November 2006.
20. P. Mika and G. Tummarello. Web Semantics in the Clouds. *IEEE Intelligent Systems*, 23(5):82–87, 2008.
21. T. Neumann and G. Weikum. RDF-3X: a RISC-style engine for RDF. In *Proceedings of 34th International Conference on Very Large Data Bases (VLDB 2008)*, volume 1, pages 647–659, Auckland, New Zealand, 2008.
22. T. Neumann and G. Weikum. Scalable Join Processing on Very Large RDF Graphs. In *Proceedings of the 35th SIGMOD international conference on Management of data*, pages 627–640, Providence, Rhode Island, USA, 2009.
23. E. Oren, S. Kotoulas, G. Anadiotis, R. Siebes, A. ten Teije, and F. van Harmelen. MaRVIN: A platform for large-scale analysis of Semantic Web data. In *Proceedings of Web Science Conference*, 2009.
24. S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz. Handling Churn in a DHT. In *USENIX Annual Technical Conference*, 2004.
25. M. Schmidt, T. Hornung, G. Lausen, and C. Pinkel. SP^{2}Bench: A SPARQL Performance Benchmark. In *Proceedings of the 29th International Conference on Data Engineering (ICDE2009)*, Shanghai, China, 2009.
26. M. Stocker, A. Seaborne, A. Bernstein, C. Kiefer, and D. Reynolds. SPARQL Basic Graph Pattern Optimization using Selectivity Estimation. In *Proceedings of the 17th International World Wide Web Conference (WWW 2008)*, Beijing, China, 2008.
27. C. Tryfonopoulos, S. Idreos, and M. Koubarakis. LibraRing: An Architecture for Distributed Digital Libraries Based on DHTs. In *ECDL*, pages 25–36, 2005.
28. J. Urbani, S. Kotoulas, E. Oren, and F. van Harmelen. Scalable Distributed Reasoning using MapReduce. In *Proceedings of the 8th International Semantic Web Conference (ISWC2009)*, October 2009.
29. M.-E. Vidal, E. Ruckhaus, T. Lampo, A. Martínez, J. Sierra, and A. Polleres. Efficiently Joining Group Patterns in SPARQL Queries. In *Proceedings of the 7th Extended Semantic Web Conference (ESWC2010)*, pages 228–242, Heraklion, Greece, 2010.
30. J. Weaver and J. Hendler. Parallel Materialization of the Finite RDFS Closure for Hundreds of Millions of Triples. In *8th International Semantic Web Conference (ISWC2009)*, October 2009.
31. C. Weiss, P. Karras, and A. Bernstein. Hexastore: Sextuple Indexing for Semantic Web Data Management. In *Proceedings of 34th International Conference on Very Large Data Bases (VLDB 2008)*, volume 1, pages 647–659, Auckland, New Zealand, 2008.
32. K. Wilkinson, C. Sayers, H. A. Kuno, and D. Raynolds. Efficient RDF Storage and Retrieval in Jena2. In *Proceedings of the 1st International Workshop on Semantic Web and Databases (SWDB 2003, co-located with VLDB 2003)*, Berlin, Germany, 2003.

Nonlinear signal processing applied to telecommunications

Diamantis Kotoulas *

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications
kotoulas@ote.gr

Abstract. Nonlinear behaviour appears in almost all digital communication systems, such as satellite systems, telephone channels, mobile cellular communications, wireless LAN devices, radio and TV channels, digital magnetic systems, etc. Linear approximations that do not take into account this type of behaviour, may lead to system performance degradation as well as loss of information. Therefore, appropriate models should be developed that tackle nonlinear system characteristics. Another important issue in studying both linear and nonlinear systems is that of the order (memory length) of the associated subsystems. It is critical, because knowing the exact subsystem orders may lead to accurate system identification and channel equalization. The primary objective of this dissertation is the solution of the structure determination problem in system identification. More specifically, the following two main objectives are pursued:

- To develop methods and algorithms for the order determination of both linear and nonlinear systems.
- To develop methods and algorithms for system identification and channel equalization.

1 Introduction

Equalization or deconvolution is essentially a signal processing procedure to restore a set of source signals which were distorted by an unknown linear or nonlinear system, whereas system identification is a signal processing procedure to identify and estimate the unknown linear or nonlinear system. The two problems prove important in a variety of areas of telecommunication applications. In this dissertation, we study blind equalization and identification methods, not only for linear systems but also for nonlinear systems modelled by finite Volterra series [1]. The approach we take towards blind identification/equalization is the following. We develop a method that identifies the orders (memory lengths) of the discrete subsystems that comprise the total system (either linear or nonlinear). This method is implemented by a computationally efficient algorithm. The algorithm detects the different subsystem orders as well as the number of

* Dissertation Advisor: Nicholas Kalouptsidis, Professor

subsystems that attain the same order. This is done both for Single Input Multiple Output (SIMO) and Multiple Input Multiple Output (MIMO) Linear Time Invariant (LTI) Finite Impulse Response (FIR) systems as well as for nonlinear SIMO Volterra systems [2, 3] .

Once the orders have been determined, subsystems are clustered in groups of the same memory length. For linear systems, Blind Source Separation techniques (BSS) are used to identify the system kernels. For nonlinear Volterra systems equalization is performed [3]. The complexity of the algorithm that is introduced depends on the number of subsystems that comprise the total system, the number of output channels, the degree of nonlinearity, the number of output samples and finally the statistical properties of the input. Our effort is concentrated on developing an algorithm that minimizes complexity with respect to all the above parameters.

This summary is organized as follows; First, in section 2 the model with the underlying assumptions, the data structures used and the order determination algorithm are presented. Section 3 provides the kernel identification algorithm for LTI FIR MIMO systems and the equalization algorithm for SIMO Volterra systems. In Section 4, conclusions and future work are presented.

2 Order Determination

Given an LTI FIR system or a nonlinear system described by a finite Volterra series, our objective is to establish that

(a) The number of discrete inputs can be computed.

(b) The orders (memory lengths) of the subsystems that comprise the total system are determined.

(c) Computations can be carried out by an efficient algorithm. Specifically, if $L_1, L_2, \cdots, L_P$ are the different subsystem orders, they may be grouped into r distinct numbers J_1, J_2, $\cdots$, J_r such that $J_1 < J_2 < \cdots < J_r$. Then for all i, $1 \leq i \leq r$, the number m_i of subsystems that have order J_i can be computed by the proposed algorithm.

2.1 LTI FIR MIMO Systems

We shall be concerned with FIR MIMO systems of the form:

$$\mathbf{x}(k) = \sum_{i=1}^{P} \sum_{j=0}^{L_i} \mathbf{h}_i(j) s_i(k-j) \tag{1}$$

The system has P inputs and M outputs. Thus the output signal $\mathbf{x}(k)$ is an $M \times 1$ dimensional vector. The input sequences consist of the signals $s_1(k)$, $s_2(k)$, $\cdots$, $s_P(k)$. The orders of the P subsystems are given by the integers $L_1, L_2, \cdots, L_P$. For each $1 \leq i \leq P$, and $0 \leq j \leq L_i$, $\mathbf{h}_i(j)$ is the corresponding $M \times 1$ kernel tap.

Memory lengths of the P subsystems are given by the integers $L_1, L_2, \cdots, L_P$. In the general case, some of the above integers may be equal to each other. Therefore, the r distint integer values appearing in the set $L_1, L_2, \cdots, L_P$ that denote the different subsystem orders, will in the following be denoted by J_1, J_2, $\cdots$, J_r and we shall assume, without loss of generality, that $J_1 < J_2 < \cdots < J_r$. In addition, for all i, $1 \leq i \leq r$, m_i will denote the number of subsystems attaining the same order J_i.

2.2 SIMO Volterra Systems

The nonlinear systems under consideration are assumed to be SIMO, discrete-time, time-invariant, causal and of finite memory. Furthermore, we assume that any small changes to the system's input $s(n)$ result in small changes in the system's output. Any such system can be approximated over a uniformly bounded set of input signals by a truncated Volterra series expansion of finite order P.

The output $y(n)$ of a real valued SIMO Volterra system, as stated in [1], can be described as:

$$x(n) = \mathbf{h}_0 + \sum_{p=1}^{P} \sum_{m_i=0}^{N-1} \cdots \sum_{m_p=0}^{N-1} \mathbf{h}_p(m_1, \cdots, m_p) s(n-m_1) \cdots s(n-m_p) + \eta(n) \quad (2)$$

As stated in [1], for a narrowband communication system input-output relationship is described by the equation:

$$x(n) = \mathbf{h}_0 + \sum_{p=1}^{P} \sum_{m_i=0}^{N-1} \cdots \sum_{m_p=0}^{N-1} \mathbf{h}_p(m_1, \cdots, m_{2k+1}) s(n-m_1) \cdots s(n-m_{k+1}) s^*(n-m_{k+2}) \cdots s^*(n-m_{2k+1}) + \eta(n) \quad (3)$$

2.3 Model Equivalence and Assumptions

We may cast a nonlinear system described by a finite Volterra series as a linear MIMO system. It should be noticed that there is no physical equivalence of the two types of systems since the inputs of the casted Volterra system are products of the original input. However, the casting is useful for algebraic manipulations. Due to the casting, any of the systems under consideration can be described as:

$$\mathbf{x}(k) = \sum_{i=1}^{B} \sum_{j=0}^{L_i} \mathbf{h}_i(j) s_i(k-j) \quad (4)$$

or equivalently

$$\mathbf{x}(k) = [\mathbf{H}(z)]\mathbf{s}(k) \tag{5}$$

where $[\mathbf{H}(z)]$ is the system transfer function and $\mathbf{s}(k) = [s_1(k) \cdots s_P(k)]^t$.

If L_q, $1 \leq q \leq P$, denotes the maximum of $L_1, L_2, \cdots, L_P$, then the channel polynomial matrix $\mathbf{H}(z)$ can be written as

$$\mathbf{H}(z) = \sum_{i=0}^{L_q} \mathbf{H}(i) z^{-i} \tag{6}$$

with

$$\mathbf{H}(i) = [\mathbf{h}_1(i)\mathbf{h}_2(i) \cdots \mathbf{h}_P(i)] \tag{7}$$

Again, the r distint integer values that denote the different subsystem orders, will in the following be denoted by J_1, J_2, $\cdots$, J_r and we shall assume, without loss of generality, that $J_1 < J_2 < \cdots < J_r$. In addition, for all i, $1 \leq i \leq r$, m_i will denote the number of subsystems attaining the same order J_i.

For LTI FIR MIMO systems the following assumptions are made:
A1) The input sequences $s_1(k), s_2(k), \cdots, s_P(k)$ are stationary Independent Identically Distributed (I.I.D.) zero mean signals of finite variance that are mutually independent with each other.
A2) An upper bound L of the subsystems' orders is known.
A3) The number of inputs P is strictly less than the number of outputs M. Furthermore, the channel polynomial matrix $\mathbf{H}$(z) is irreducible and column reduced.

For SIMO Volterra systems the following assumptions are made:
B1) The input sequence $s(n)$ is zero mean, i.i.d., with values in a finite alphabet of at least $P + 1$ complex numbers. Examples include PSK or QAM signals.
B2) The system transfer matrix $\mathbf{G}$(z) is irreducible. This guarantees that there are no common zeros in the (FIR) transfer functions of every pair of the subchannels involved. It is a common assumption in all methods based on Second Order Statistics.
B3) The memory of the linear kernel is strictly greater than the memory of any nonlinear term. Furthermore, the first element of the zero-th tap of this kernel is equal to unity.
B4) An upper bound L of B is known.

2.4 Definition of Data Structures

Assuming that $k, Q \in \mathcal{N}$and $Q > k$, we collect successive output vectors in the matrix

$$\mathbf{T}(k) = [\mathbf{x}(k) \quad \mathbf{x}(k+1) \quad \cdots \mathbf{x}(Q)] \tag{8}$$

Likewise, we combine successive input values in row vector form as follows

$$\begin{aligned}\mathbf{s}_1(k) &= [s_1(k) \quad s_1(k+1) \quad \cdots s_1(Q)] \\ &\vdots \quad \vdots \\ \mathbf{s}_B(k) &= [s_B(k) \quad s_P(k+1) \quad \cdots s_B(Q)]\end{aligned} \tag{9}$$

Using a window of length w, we form the following array of output samples:

$$\mathbf{X}_{k,w} = \begin{pmatrix} \mathbf{T}(k) \\ \vdots \\ \mathbf{T}(k-w+1) \end{pmatrix} \tag{10}$$

$\mathbf{X}_{k,w}$ is the data matrix defined by stacking w consecutive such observations, starting with $\mathbf{T}(k)$ and going back to $\mathbf{T}(k-w+1)$. Similarly, for all $j, 1 \leq j \leq B$, we have

$$\mathbf{S}^j_{k,w} = \begin{pmatrix} \mathbf{s}_j(k) \\ \vdots \\ \mathbf{s}_j(k-w+1) \end{pmatrix} \tag{11}$$

Next, we consider the row spaces

$$\mathcal{X}_{k,w} = \mathcal{R}\left(\mathbf{X}_{k,w}\right) \tag{12}$$

$$\mathcal{S}^j_{k,w} = \mathcal{R}\left(\mathbf{S}^j_{k,w}\right) \tag{13}$$

Finally, for any $l \in \mathcal{N}$ we define the vector space

$$\begin{aligned}\dot{\mathcal{S}}_{k,l} = &\left(\mathcal{S}^1_{k-1,L_1+w} \oplus \cdots \oplus \mathcal{S}^B_{k-1,L_B+w}\right) \\ &\bigcup \left(\mathcal{S}^1_{k+l+w,L_1+w} \oplus \cdots \oplus \mathcal{S}^B_{k+l+w,L_B+w}\right)\end{aligned} \tag{14}$$

If we think of l as a smoothing window, we see that $\dot{\mathcal{S}}_{k,l}$ is a subspace constructed from both past and future data observations. Based on the above, we state the following theorem that establishes the isomorphic relationship between input and output subspaces.

Theorem 1: For all $w \geq w_0$, $\mathcal{X}_{k,w} \doteq \mathcal{S}^1_{k,L_1+w} \oplus \cdots \oplus \mathcal{S}^B_{k,L_B+w}$.

Then, we define the projection error matrix and the kernel-input product matrices that are necessary to the development of the algorithm. The projection error matrix $\mathbf{E}_{k,l}$ is defined as

$$\mathbf{E}_{k,l} = \begin{pmatrix} \mathbf{T}(k+l) - \mathbf{T}(k+l)_{|\dot{\mathcal{S}}_{k,l}} \\ \mathbf{T}(k+l-1) - \mathbf{T}(k+l-1)_{|\dot{\mathcal{S}}_{k,l}} \\ \vdots \\ \mathbf{T}(k) - \mathbf{T}(k)_{|\dot{\mathcal{S}}_{k,l}} \end{pmatrix}$$

$\mathbf{E}_{k,l}$ is a $(M(l+1))\times(Q-k+1)$ matrix. Each block entry $\mathbf{T}(k+m)-\mathbf{T}(k+m)_{|\dot{\mathcal{S}}_{k,l}}$ of $\mathbf{E}_{k,l}$ is formed by the error resulting from the projection of $\mathbf{T}(k+m)$ on the space $\dot{\mathcal{S}}_{k,l}$ generated by past and future input values. In the following analysis,we shall fix k, without loss of generality since the carried out analysis is valid for any choice of it. For all $i, 1 \leq i \leq r$, let $\mathbf{h}_{i_1}, \mathbf{h}_{i_2}, \cdots, \mathbf{h}_{i_{m_i}}$, be the m_i subsystems that attain order J_i. Denote by $s_{i_1}, s_{i_2}, \cdots, s_{i_{m_i}}$ the inputs to the above subsystems.

Then, for a given subsystem $\mathbf{h}_{i_s}$, for all $l, J_1 \leq l \leq L$ and $t, 1 \leq t \leq m_i$, we define the matrix:

$$\mathbf{D}_{l,J_i}(\mathbf{h}_{i_t}) = \begin{pmatrix} \mathbf{h}_{i_t}(0) & \cdots & \mathbf{0}_{M\times 1} \\ \vdots & \ddots & \vdots \\ \mathbf{h}_{i_t}(J_i) & \vdots & \mathbf{h}_{i_t}(0) \\ \vdots & \ddots & \vdots \\ \mathbf{0}_{M\times 1} & \cdots & \mathbf{h}_{i_t}(J_i) \end{pmatrix},$$

Next, for all subsystems attaining the same order J_i, we form the matrix:

$$\mathbf{G}_{l,J_i} = \left(\mathbf{D}_{l,J_i}(\mathbf{h}_{i_1}) \cdots \mathbf{D}_{l,J_i}(\mathbf{h}_{i_{m_i}})\right),$$

Collect input values of the above subsystems, to form the matrix:

$$\tilde{\mathbf{S}}_{l,J_i} = \begin{pmatrix} \mathbf{s}_{i_1}(k+l-J_i) \\ \vdots \\ \mathbf{s}_{i_1}(k) \\ \vdots \\ \mathbf{s}_{i_{m_i}}(k+l-J_i) \\ \vdots \\ \mathbf{s}_{i_{m_i}}(k) \end{pmatrix}$$

Finally, for all $i, 1 \leq i \leq r$ and for all l, $J_i \leq l \leq L$ we define the matrices

$$\mathbf{\Psi}_{l,J_i} = \mathbf{G}_{l,J_i}\tilde{\mathbf{S}}_{l,J_i} \tag{15}$$

Having defined the projection error matrix and the kernel-input product matrices we proceed to:

Theorem 2

(i) For any $l : l < J_1$, $\mathbf{E}_l = 0$

(ii) For $l : J_1 \leq l < J_2$,

$$\mathbf{E}_l = \mathbf{\Psi}_{l,J_1} \tag{16}$$

(iii) For $n : 3 \leq n \leq r$ and $l : J_{n-1} \leq l < J_n$,

$$\mathbf{E}_l - \sum_{m=1}^{n-2} \mathbf{\Psi}_{l,J_m} = \mathbf{\Psi}_{l,J_{n-1}} \tag{17}$$

(iv) For $l : J_r \leq l \leq L$,

$$\mathbf{E}_l - \sum_{m=1}^{r-1} \mathbf{\Psi}_{l,J_m} = \mathbf{\Psi}_{l,J_r} \tag{18}$$

Next we establish the ranks of the projection error matrices $\mathbf{E}_l$.
Theorem 3: The following statements are true:
(i) For any $l : l < J_1$, $rank(\mathbf{E}_l) = 0$
(ii) For $l : J_1 \leq l < J_2$,

$$rank(\mathbf{E}_l) = (l - J_1 + 1)m_1 \tag{19}$$

(iii) For $n : 3 \leq n \leq r$ and $l : J_{n-1} \leq l < J_n$,

$$rank(\mathbf{E}_l) = \sum_{i=1}^{n-1}(l - J_i + 1)m_i \tag{20}$$

(iv) For $l : J_r \leq l \leq L$,

$$rank(\mathbf{E}_l) = \sum_{i=1}^{r}(l - J_i + 1)m_i \tag{21}$$

Theorem 3 indicates how to compute the system's different orders, as well as the number of subsystems that attain it in a straightforward manner. Indeed, starting with $l = 0$ we compute $rank(\mathbf{E}_0)$ and we increase l by one until $rank(\mathbf{E}_l) > 0$. Equation (19) suggests that this value of l equals the smaller of the orders J_1. Moreover, it also gives the number m_1 of the subsystems that attain it. Having determined m_1, we increase l in steps of one. As long as $l < J_2$, $rank(\mathbf{E}_l)$ remains a multiple of m_1. When this stops to hold (20) suggests that $l = J_2$. At this point having computed m_1, J_1 we use (20) to determine the number m_2 of the subsystems that attain the order J_2. We continue increasing l. As long as $l < J_3$, $rank(\mathbf{E}_l)$ increases by $m_1 + m_2$ each time l increases by one. Again, when this stops to hold, (20) suggests that $l = J_3$. At this point, having computed J_1, m_1, J_2, m_2 we use (20) to determine the number m_3 of the subsystems that have order equal to J_3. We keep increasing l by one until we reach L and proceed in the same way, using (20) and (21) to determine J_i, m_i for all $i : 1 \leq i \leq r$. When $l = L$, we compute the number P of input signals as:

$$P = \sum_{i=1}^{r} m_i \tag{22}$$

3 LTI FIR MIMO Systems Kernel Identification and Volterra Systems Equalization

Given an LTI FIR system or a nonlinear system described by a finite Volterra series the objective of the algorithm is

(a) To identify the kernels of the system in case it is LTI FIR MIMO.

(b) To equalize input symbols in case of SIMO Volterra systems.

The model as well as the assumptions used have already been defined in the previous chapter. An extra assumption is required for SIMO Volterra systems. Specifically, the memory of the linear kernel is strictly greater than the memory of any nonlinear term, while the first element of the zero-th tap of this kernel is assumed to be equal to unity. Data structures required, have already been defined in the previous section. In this section we state and prove a number of lemmas and theorems that establish the validity of the algorithm. In addition, the steps of the algorithm are presented.

Assuming k is fixed, we shall denote $\mathbf{E}_{k,l}$ by $\mathbf{E}_l$, to simplify notation. We state theorem 4 that is essential to the development of the proposed algorithm:

Theorem 4

For all $i, \quad 1 \leq i \leq r$ and each $l, \quad J_i \leq l \leq L$, the elements of each one of the matrices $\mathbf{\Psi}_{l,J_i}$ can be computed from the elements of the matrices $\mathbf{E}_{J_1}, \mathbf{E}_{J_2}, \cdots, \mathbf{E}_{J_i}$.

3.1 Algorithm Description

A batch algorithm that performs system identification/equalization is developed. For LTI FIR MIMO systems the algorithm uses the Joint Diagonilization Principle through the JADE algorithm as described in [9]. The outline of the algorithm is the following.

***A*. Order Estimation**

- Step A_1: At the receiver, group the collected symbols in blocks of Q symbols per block.
- Step A_2: Using the first block, for all $l, J_1 \leq l \leq L$ compute the matrices $\mathbf{E}_l$. Store the matrices.
- Step A_3: Compuite the orders $J_1, J_2, \cdots J_r$ of the different subsystems that comprise the overall system following the algorithm established in [12]. Store $J_1, J_2, \cdots J_r$.
- Step A_4: For FIR LTI MIMO systems execute part B of the algorithm, for SIMO Volterra systems execute part C.

***B*. LTI FIR MIMO Systems Identification**

- Step 1: Use matrices $\mathbf{E}_l$ computed in part **A**, to compute the matrices $\mathbf{\Psi}_{l,J_i}$, according to theorem 1.
- Step 2: Set $l = J_1$. Using the equation

$$\mathbf{E}_{J_1} = \mathbf{\Psi}_{J_1,J_1} \tag{23}$$

  and the JADE algorithm identify the kernels corresponding to the subsystems of order J_1 as columns of the matrix $\mathbf{G}_{J_1,J_1}$.
- Step 3: Set $l = J_2$. Using the equation

$$\mathbf{E}_{J_2} - \mathbf{\Psi}_{J_2,J_1} = \mathbf{\Psi}_{J_2,J_2} \tag{24}$$

and the JADE algorithm identify the kernels that correspond to subsystems of order J_2 as columns of the matrix $\mathbf{G}_{J_2,J_2}$.

$\vdots$

– Step $r+1$: Set $l = J_r$. Use equation

$$\mathbf{E}_{J_r} - \sum_{i=1}^{r-1} \mathbf{\Psi}_{J_r,J_i} = \mathbf{\Psi}_{J_r,J_r} \tag{25}$$

and the JADE algorithm to identify the kernels corresponding to subsystems of order J_r as columns of the matrix $\mathbf{G}_{J_r,J_r}$.

– **End of Kernel Identification**

C. **Volterra Systems Equalization**

– Step 1: Based on theorem 3, compute the matrices $\mathbf{\Psi}_{J_r,J_m}$, for all $m, 1 \leq m \leq r-1$.

– Step 2: Set $l = J_r$ and compute the matrix

$$\mathbf{E}_{J_r} - \sum_{m=1}^{r-1} \mathbf{\Psi}_{J_r,J_m} = \mathbf{\Psi}_{J_r,J_r} \tag{26}$$

– Step 3: Use the first row of the computed matrix $\mathbf{\Psi}_{J_r,J_r}$ to equalize the input vector. Due to assumption A3 and the definition of $\mathbf{\Psi}_{J_r,J_r}$, the first row of $\mathbf{\Psi}_{J_r,J_r}$ equals the vector $[s(k) \quad s(k+1) \cdots \quad s(k+Q-1)]$.

– Step 4: Use the next block of Q received symbols, to compute for all $i, 1 \leq i \leq r$ the matrices $\mathbf{E}_{J_i}$. Store the matrices.

– Step 5:Repeat steps $1-4$ to equalize next block of Q symbols. Do so, until all received symbols have been processed.

– Step 6:**End of Equalization**

For Volterra systems, after equalization has been performed, the vector of the linear kernel taps can be evaluated as the first column of $\mathbf{\Psi}_{J_r,J_r}$ divided by $s(k)$. Furthermore, in order to equalize input symbols using equation(26) we only need to know the values of the first row of each of the matrices $\mathbf{\Psi}_{J_r,J_m}$.

4 Conclusions

4.1 Order Determination

The performance of the proposed algorithm in the presence of noise relies heavily on the determination of the effective rank of a matrix perturbed by noise, is the most sensitive. In [5] various criteria are given that estimate the effective rank of a perturbed matrix using its singular values. All but one out of these, use threshold values that either do not appear to be based on any explicit analytical expressions but are selected on an ad hoc basis, or lower and upper bounds for them can be derived analytically assuming known noise statistics.

The only criterion among those presented in [5] that does not use either an ad hoc threshold value or a threshold with upper and lower bounds based on known noise statistics, determines the numerical rank t of a perturbed matrix $\mathbf{B}$ from its singular values $\beta_1 \geq \beta_2 \geq \cdots \geq \beta_r$ as the index t for which $\beta_t >> \beta_{t+1}$. It is straightforward to see that it is equivalent to the criterion used by the proposed method.

Simulation experiments showed the following:

- For MIMO systems, the validity of the method was established even for low SNR values. On the other hand, the rank pattern suggested by Theorem 3 is harder to detect in the presence of noise.
- For SIMO systems the proposed method outperforms the algorithm given in [4] in all cases at least by 6dB, while MDL achieves the 90% success rate at about the same SNR levels with the proposed method. The new method outperforms the J-LSS algorithm by at least 14dB, while J-LSS performed poorly in the case of true microwave channels.
- Denoising improved dramatically the algorithm's performance for MIMO systems. For SIMO systems, denoising had a mixed effect. However, we should mention the performance boost that was achieved when denoising was applied at the J-LSS method.
- In all our simulation experiments w was chosen equal to L the upper bound of system's orders. This is because, as noticed in [8], for a fixed data length, large values of w imply fewer columns in the algorithm's data matrices that correspond to smaller sample size when projections on subspaces are evaluated. Therefore, we tried to keep w as small as possible.
- The algorithm gives analogous results for order detection of Volterra Systems. It is important to notice that using 1500 symbols the algorithm achieves high rates of successful order detection even for low SNR values such as 10dB. This is due to the special structure of the Volterra systems, that is due to linearization of the system, where the new inputs were redefined as subproducts of delayed input values.
- The computational complexity of the proposed algorithm is $O(M^3Q^2(L+1)^4+M^3QL^3(L+1)+M^3Q^2L^3(L+1)+M^3Q(L+1)^3(L-J_1+1))$ for MIMO systems. This reduces to $O(M^3Q^2(L+1)^4+M^3QL^3(L+1)+M^3Q^2L^3(L+1))$ for SIMO systems. In the above M is the number of outputs and Q is the number of output data samples used. This suggests that the oversampling rate or the number of diverse sensors M used at the receiver should be kept small to avoid heavy computations. In addition, the less output samples used and the tighter the estimate of the system orders' upper bound L is, the less expensive computations become.

4.2 LTI FIR MIMO Kernel Identification and SIMO Volterra Systems Equalization

For LTI FIR MIMO systems the algorithm allows kernel identification in steps, based on the JADE algorithm. Kernels are identified in groups, depending on the

memory length of the subsystems they belong to. Performance improvement as compared with the algorithms given in [9, 10] emanates from the fact that kernel identification is performed in steps, starting with the subsystems attaining the lower memory and progressing to the subsystems with the higher memory. The computational complexity of the proposed method is $O(Q^4 + (ML\ Q \log_2(Q)))$, where Q is the size of symbols block, This is because, even though the method is SOS based, the involvement of the JADE algorithm increases the complexity, as it requires the computation of fourth-order cumulants.

For Volterra systems, the performance of the algorithm is quite satisfactory, as compared with the ones given in [6, 11]. The overall algorithm complexity is therefore, $(\frac{N}{Q})O(ML\ Q \log_2(Q))$.

For both LTI FIR MIMO systems and SIMO Volterra systems, selecting the appropriate size Q of symbol blocks affects the correct order determination and by that the matrices E_l based on which identification/equalization are performed. Selection of Q depends on SNR as well as the complexity of the system examined, that is the number of subsystems it comprises of and their orders. Extensive simulation experiments, involving both LTI FIR MIMO and Volterra systems triggered by different types of PSK input, suggested that $Q \geq 1500$.

4.3 Future Work

As it has already been mentioned the algorithm's complexity is $O(M^3Q^2(L+1)^4 + M^3QL^3(L+1) + M^3Q^2L^3(L+1) + M^3Q(L+1)^3(L-J_1+1))$ for the order identification part of MIMO systems and $O(M^3Q^2(L+1)^4 + M^3QL^3(L+1) + M^3Q^2L^3(L+1))$ for SIMO systems. Moreover, we pay $O(Q^4 + (ML\ Q \log_2(Q)))$ operations for identification or $(\frac{N}{Q})O(ML\ Q \log_2(Q))$ for Volterra systems equalization. It is obvious that there should be some future work to improve computational complexity. This can be achieved either by ausing a recursive implementation scheme or by introducing other type of BSS techniques, with less complexity than the JADE algorithm.

References

1. V. J. Mathews, G.L. Sicuranza, Polynomial Signal Processing, Wiley Series in Telecommunications and Signal Processing, John Wiley, 2000.
2. D. Kotoulas, P. Koukoulas, N. Kalouptsidis, "Subspace Projection Based Blind Channel Order Estimation of MIMO Systems", *IEEE Transactions on Signal Processing*, vol. SP-54, no. 4, pp. 1351-1363, April 2006.
3. D. Kotoulas, P. Koukoulas, G.-O. Glentis, "Subspace based blind identification of LTI FIR MIMO systems and equalization of finite memory SIMO Volterra systems", Signal Processing (Elsevier), vol. 91, no. 8, pp. 1941-1950, Aug. 2011.
4. A. P. Liavas, P. A. Regalia and J.-P. Delmas "Blind channel approximation: effective channel order determination," *IEEE Trans. Signal Processing,* vol. 47, pp. 3336-3344, Dec. 1999.
5. K. Konstantinides and K. Yao, "Statistical analysis of effective singular values in matrix rank determination," *IEEE Trans. on Acoust., Speech, Signal Processing*, vol. 36, pp. 757-763, May 1988.

6. G.B. Giannakis, E. Serpedin, Linear Multichannel Blind Equalizers of Nonlinear FIR Volterra Channels, IEEE Transactions on Signal Processing, vol. 45, no. 1, pp. 67-81, January 1997. vol. 49, pp. 2042-2049, Sept. 2001.
7. L. Tong and Q. Zhao, "Joint order detection and blind channel estimation by least squares smoothing," *IEEE Trans. Signal Processing,* vol. 47, pp. 2345-2355, Sept. 1999.
8. D. Kotoulas, P. Koukoulas, N. Kalouptsidis, "Subspace Projection Based Blind Channel Order Estimation of MIMO Systems," IEEE Transactions on Signal Processing, vol. SP-54, no. 4, pp. 1351-1363, April 2006. Sept. 2003, pp. 409-412.
9. J. Liang and Z. Ding, "Blind MIMO system identification based on cumulant subspace decomposition," *IEEE Trans. Signal Processing,* vol. 51, pp. 1457-1468, June 2003.
10. S. An, J. H. Manton, Y. Hua, "A Sequential Subspace Method for Blind Identification of General FIR MIMO Channels," *IEEE Transactions on Signal Processing*, vol. 53, pp. 3906-3910, October 2005.
11. R. Lopez-Valcarce and S. Dasgupta, "Blind Equalization of Nonlinear Channels From Second-Order Statistics," *IEEE Trans. Signal Processing*, vol. 49, pp. 3084-3097, Sept. 2001.

The impact of economic and social environment on the diffusion of Information and Communications Technologies

Vagia Kyriakidou*

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications
bkiriak@di.uoa.gr

Abstract. In this PhD thesis, the impact of economic and social environment on the diffusion of Information and Communications Technologies (ICTs) is studied. Initially, broadband diffusion process in Europe is examined and suggested an approach to measure the digital convergence of countries participating in the analysis. Through a non-parametric analysis between broadband penetration and socio-economic parameters, it is shown that the impact of the influential parameters depends on the level of broadband penetration. Furthermore, a causal analysis between the economic status and the level of broadband penetration is performed in many countries. Then, a new metric, the "Utilization of Communications Network Potential" (UCNP) versus public practices aimed at enhancing new technologies is proposed. The impact of UCNP is examined in a series of socio-economic parameters. Next, a new index is introduced i.e. the "Maturity Level of ICT" (IMLI). The calculation of this index allows the ranking of countries on the basis of their economic status, highlighting best practices. Finally, a theoretical framework for developing business plans operating optical networks is developed, along with an analysis of population restrictions regarding Greece.

Keywords: Broadband, Influential Factors, Econometrical Analysis, Service Diffusion, Decision Making

1 Introduction

The development of new technologies and, particularly those which are known as Information and Communications Technologies (ICT), raises a number of entrepreneur and political considerations. In this dissertation, the impact of economic and social environment on the diffusion of ICTs is studied. The analysis is based on statistics from international data bases, econometric methods and strong mathematical tools. For assessing the reliability of results statistical measures and criteria from information theory were used. In the following sections critical questions regarding the aforementioned impact are attempted to be answered.

* Dissertation Advisor: Thomas Sphicopoulos, Professor

Apart from the continuously increase of ICT diffusion, ICT inequality among countries, in terms of ICT investments, PC skills, Internet skills and the availability of telecommunications networks is still evident. This inequality can be described as a digital gap (or digital divide) between and within countries. Broadband services are considered as key factor of ICTs, and therefore the rate of digital convergence within European borders, in terms of broadband penetration across a number of 28 European countries is studied [1]. Digital convergence is used under the meaning of the process of homogenization of broadband diffusion, based on the status of penetration rate of corresponding countries. In addition, it is especially concerned with the evaluation of the contribution of each country to the process of broadband diffusion, as a whole. The methodology used to measure the digital divide is based on appropriate mathematical or econometrical approach and corresponding assumptions regarding the proxies used to estimate it. The digital gap is assumed to be driven and reflected by the differences in broadband penetration among European countries. An adequate index for the estimation of digital convergence is developed which in turn can be used to provide future forecasts on the convergence.

Despite of the technical characteristics, i.e. broadband penetration rate, there is a group of other important factors which play a crucial role in the process of broadband growth. The analysis of the impact of these factors on broadband adoption could lead to useful outputs and it constitutes the objective of the following analysis [2]. Towards this direction a number of socio-economic factors are assumed to influence broadband diffusion, considered responsible for the different levels of adoption among countries. Therefore, the effect of a wide range of social, economic and political factors over the broadband diffusion process is analyzed, following a non-parametric approach and comparing the results with these of the parametric. Based on criteria from information theory the link function between the level of penetration and the rest variables is derived, providing highly accurate results, during the different stages of the diffusion process, which are defined by the inflection point. The evaluation of the methodology was performed over countries from the wider European area, revealing that the influential parameters vary, depending on the stage of the diffusion cycle.

Under the present economic crisis, much discussion is made for income increase. New technologies, especially broadband, can yield comparative advantages in countries. Based on the assumption that income is related with broadband penetration, a Granger causality test is applied in 83 countries worldwide [3]. The causality case is tested separately in developed and developing countries, as far as in different level of broadband diffusion. The analysis aimed to examine the relationship (one-way or bidirectional) between Gross National Income per capita and broadband penetration level.

Technological adoption, with an emphasis on ICT, is considered as a decisive factor for the overall development of each country. For this, the European Commission (EC) has launched a number of policy frameworks, while recently EC decided to set some additional targets, in order to facilitate a wider adoption of information services and maximize economical and societal benefits. In line with this, the effect of the driving factors that accelerate the uptake of public e-services, together with the impact of technological adoption on the socio-economic status is studied [4]. A new parame-

ter is introduced, the Utilization of Communications Network Potential (UCNP), which echoes the Information Society (IS) maturity level. An analysis, focusing on monitoring the progress of public and European Commission (EC) actions is additionally presented in order to assess the evolution of the IS maturity level in the European area. The impact of two main public depended indexes, i.e. structural and benchmarking indicators, on the UCNP maturity level is evaluated, together with the influence of the latter over socio-economic parameters.

The use of ICT turned out to be a key factor in the process of the wider development of a country. It is therefore very useful to estimate ICT evolution by the means of an appropriate metric. Based on statistical data from 159 countries, the ICT Maturity Level Index (IMLI) is proposed and estimated by using Structural Equation Modelling (SEM) [5]. This index is a metric measuring the information society in a country and consists of three sub-indices which are access, use and skills. It is an improvement of the ICT development index, proposed by the ITU in 2009. The analysis divides the countries into two groups, the developed and the developing, due to major disparities in their statistical data. The criterion used to define the groups was the income, as expressed by the Gross National Income per capita. The impact of a number of influential parameters on the ICT maturity level is evaluated and it becomes obvious that there is a substantial difference in their impact between developed and developing countries. Finally, a procedure that allows the ranking of the countries, based on IMLI, is presented.

Additionally to the above, a pool of business solutions for Metropolitan Area Networks (MANs) development, especially for the Greek case, is proposed based on detail analysis and benchmarking of international practices [6]. Finally, one of the most widely used power laws applied in demographics, the Zipf's law, is tested over urban cities in Greece. Apart from the examination of Zipf's law validation over population, the distribution of population density as far as population differentiations in the last decades are studied [7].

2 Digital Divide Gap Convergence in Europe

Despite the targeted actions initiated by the EC and the member states, there are still large differences in the corresponding broadband penetration rates [7]. However, there is still room for improvement that could lead to a higher broadband adoption. The dataset used in this analysis was extracted from Eurostat and the term penetration describes the number of connections over population (i.e. over the 28 European countries considered) [9]. More specifically, the dataset consists of 28 quarter terms, from December 2001 up to December 2009. The development of the proposed methodology is presented in the following paragraphs, in terms of the assumptions made and the mathematical framework used.

For each quarter term, *t*, and for each one of the considered countries, *i*, the proportion of its penetration over the mean penetration rate is calculated. The results are plotted against time and depict the relative penetration, *RP*, of each considered country and for each period of time. The formulation for calculating the *RP* is described by Eq.1.

$$RP_{(i,t)} = \frac{\frac{x_i}{p_i} - \frac{x_T}{p_T}}{\frac{x_T}{p_T}} \quad (1)$$

where x_i is the cumulative diffusion rate of each country at time t, and x_T is the total penetration of the considered countries, in the same time period, in terms of subscribers. In addition, p_i is the population of each country and p_T is the total population across all of the considered countries. In addition to the above calculations, the width of $RP_{(t)}$, ($\Delta(RP_t)=RP_{max,t}-RP_{min,t}$), in each quarter term is calculated.

Further analysis includes the forecasting of the future values of $\Delta(RP_t)$, which describes the future process of digital convergence. The corresponding results for both the estimated and the forecasted values of $\Delta(RP_t)$ are plotted against time, in order to provide an estimation regarding the time of full convergence. Forecasting is based on appropriate mathematical functions, originated from two representative function families, the exponential and the polynomial.

The calculated $RP_{(i,t)}$ values for all the European countries considered tend to zero, which means that countries tend to equally contribute to the total broadband diffusion in Europe. The statistical accuracy of the two functions used to provide forecast results is based on the R-squared (R^2), the Residual Sum of Squares (RSS), the Mean Square Error (MSE) and the Mean Absolute Error (MAE).

Both functions provide quite acceptable results and therefore can be expected to provide reliable forecasts as well. However, in all cases, the estimated values of statistical measures referred to exponential approach provided more accurate results.

A visual representation of the above findings is illustrated in the graphs of Fig. 1. According to these graphs, it is obvious that both functions provide accurate fittings of the actual observed values.

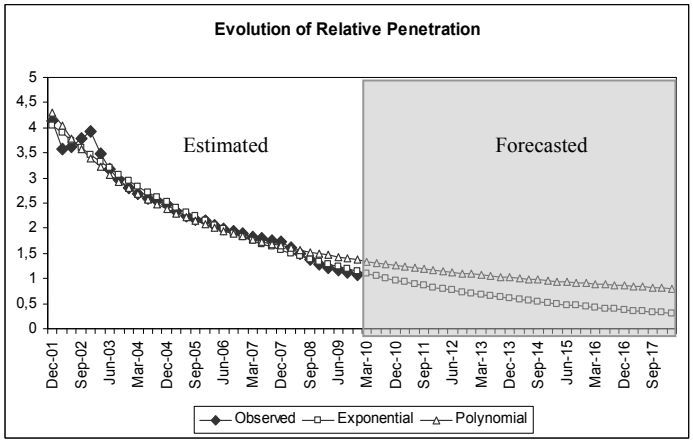

Fig. 1 Evolution of relative penetration over time. The white portion refers to the estimated values and the grey to the forecasted.

The plots of approximate findings that forecast the digital divide convergence are presented in Fig. 1. The polynomial function seems to be more pessimistic, as the corresponding plot tends to keep an almost constant distance from zero, whereas the

exponential function tends to be closer to the X axis. However, both forecasts imply that European countries will come closer, in terms of broadband diffusion. Thus, based on the exponential model, the homogenization is expected to be met after the 67th quarter, which coincides not before the first quarter of 2018. Of course, this result reflects the current dynamics of the system, based on the level of influence of the factor that affect the process, which leads to huge differentiations in broadband penetration rates among countries.

3 Driving Factors During the Different Stages of Broadband Diffusion: A Non-Parametric Approach and the relationship with income

The contribution of this work is, firstly, to use a non-parametric regression in order to identify the a priori unknown nature of the relation between broadband diffusion (dependent variable) and socio-economic factors (independent variables). Secondly is to assume that the influence of socio-economic factors over broadband diffusion differs during the two main stages of the diffusion process, i.e. before and after inflection point (IP).

The first step corresponds to the determination of the relevant social, economic and other parameters and data collection for each considered country, from Eurostat's database [9]. Then the estimation of the IP of each country based on the non-symmetric Gompertz model according to the following procedure is conducted: Based on the existing diffusion data, for each country, the Gompertz model was trained, using non- linear squares (NLS) [10]. This procedure concluded with the estimation of the diffusion parameters, i.e. the diffusion rate and the saturation level, for each country. Following this, the inflection point is calculated at the 37,79% of the estimated saturation level, as derives from the formulation of the model. Next, the data divided into two groups, based on the estimated inflection point of the diffusion process, for each country, according to the procedure described in the previous step. Diffusion data that have lower values than the inflection point are collected into the first group, while the rest are collected into the second. In addition, there is a third group which consists of the whole dataset which is also analyzed, mainly for comparison reasons. After that, the Sliced Inverse Regression (SIR) algorithm applied to all three groups [11]. This will result into the identification of the most influential parameters and the exclusion of the less relevant ones, leading to dimension reduction with no loss of information. Subsequently, a Kernel smoothing technique was applied, in order to reveal the behavior of the dependent against the independent variables. This will serve as a driver for the construction of the pool of the candidate link functions. Finally, Akaike Information Criterion (AIC) [12] and Bayesian Information Criterion (BIC) [13] were estimated, in order to conclude to the link function that best describes the modeled system, in terms of appropriate statistical measures.

The dataset, that the analysis was based on, consists of cross-sectional, time-series data, over 26 European countries [9]. More specifically, semiannual data were used, starting from the second semester of year 2001 up to December 2009, for all countries. Before the analysis was performed after a dataset transformation, based on the

generalized differences [14]. More specifically, the corresponding parameters are: Gross Domestic Product (in €), E-government online availability – Percentage of online availability of 20 basic public services, Tertiary graduates in science and technology per 1 000 of population aged 20-29 years, Persons employed using computers connected to the Internet, Inequality of income distribution – Income quintile share ratio, Proportion of population aged 25-49, School expectancy - Expected years of education over a lifetime, Individuals' level of Computer skills – Percentage of the total number of individuals aged 16 to 74, Individuals' level of Internet skills – Percentage of the total number of individuals aged 16 to 74, Price of telecommunications – Local calls (10 minutes), Communications expenditure as a percentage of GDP and Population Density.

In order to conclude to the link function that better describes the data, the AIC and BIC criteria for the linear and the exponential functions are calculated. According to results, it seems that the exponential model has the lowest values, in both criteria (AIC and BIC). Therefore, it is considered as the most appropriate function to describe the broadband diffusion process in the majority of the cases. As mentioned above, a parametric analysis is also conducted in order to compare the results with those provided by non-parametric method. The comparison was based on appropriate statistical measures such as the coefficient of multiple determinations (R2), the adjusted coefficient of multiple determinations (Ra2), the Residual Sum of Squares (RSS) and the standard error of the estimates (s.e.). Comparison of the results shows that all the statistical measures associated with the non-parametric approach are better, in all cases, than these of the parametric. However, it must be noted that the parametric method provided quite acceptable outputs despite the fact that the non-parametric approach seems to fit the data more accurately.

Based on the assumption that the different stages of diffusion are influenced by different factors, the data used for the analysis were split into two segments, before and after the inflection point of the diffusion process, for each considered country. The analysis revealed that, as initially assumed, the factors that influence diffusion are partially different, during the different stages of the process.

According to the results, it is proved that governmental actions such as E-government online availability and communication expenditures can enhance the process of broadband diffusion. The level of communications expenditures could exert either positive or negative influence depending on the level of cumulative communications investments made in the past. Moreover, the proportion of persons employed using computers connected to the Internet is a driving factor for the proliferation of broadband diffusion for all three groups of considered datasets. It seems that penetration in the private sector is directly and positively related to the broadband diffusion over households. Additionally, population density can be considered as a supporting parameter in the process, as urbanization facilitates the adoption of broadband services.

In the dataset describing the process before the inflection point, Internet skills are added to the influential factors with a positive sign regarding broadband adoption. Thus, policies targeting to motivate people gaining such kind of skills could accelerate broadband diffusion. The results indicate that governmental intervention regard-

less the level of penetration rate can promote broadband diffusion. It is also concluded that the most suitable regression model can be different based on the maturity level of broadband diffusion.

In addition to the above analysis, a more specific relationship is also examined, i.e. the relationship between broadband diffusion and Gross National Income (GNI) per capita. The aim of this analysis is to provide detailed results regarding the interaction of the two aforementioned variables. To do this, a Granger causality test is applied in a number of countries which where divided into developed and developing [15]. The criterion used for the separation was GNI per capita, based on the classification used by World Bank [16]. The main assumption of this analysis lies on the hypothesis that the relationship between income and broadband diffusion differs depending on the level of the later variable. Therefore, each dataset, i.e. 42 developed and 41 developing countries, examined separately regarding the level of broadband diffusion by using inflection point (IP) in order to split data in two groups, i.e. before and after IP. The estimation of IP of each country based on the non-symmetric Gompertz model [10]. For comparison reason Granger causality test was also applied to the whole dataset, in order to indentify differences in results.

On the one hand, based on the results derived from developed countries, it seems that after IP, broadband penetration causes positive changes in income, while before IP, income will precede positive changes in broadband diffusion. Though, it seems that countries can have economic benefits if they enhance broadband diffusion. On the contrary, countries with low broadband penetration, i.e. before IP, could focus their efforts to provide economic incentives in order to boost broadband demand. It is worth mentioned that results for all the developed countries are similar with those before IP.

On the other hand, the results referring to developing countries reveal that one way relationship is existed between considered variables. More specifically, after IP broadband diffusion can be considered as an accelerator for income increase, while before IP the opposite relationship is forced. As far as the results for all the developing countries, a bidirectional relationship is estimated. According to this, it seems that there is a feedback relationship between GNI per capita and broadband penetration.

4 Utilization of Communications Network Potential: Public practices and effects

In the context of the present work a new metric, the Utilization of Communications Network Potential (UCNP), which describes the maturity level of what is called “Information Society- IS”, is proposed. UCNP is a latent (unobserved) parameter which reflects the output of public and EC applied practices on the use and adoption of new technologies. Public sector indicators are considered as the most appropriate indexes connecting governmental initiatives, policy decisions, EC practices etc., with the enhancement of the UCNP. “Structural indicators” and “Benchmarking indicators: Public services- eGovernment” are the most coherent indexes reflecting the results of public interventions and public e-services, respectively.

The analysis performed in this work is twofold and is based on the preceding considerations. On the one hand, the relationship between UCNP maturity and public sector indicators is studied, under the assumption that these variables interact with each other. On the other hand, the impact of UCNP maturity over the main socio-economic parameters is examined. From the first part of the analysis, useful outputs can be derived, regarding the importance of the two considered key constructs according to public practices. According to the results, policy suggestions and rethinking proposals could arise and new opportunities can be explored by business organizations. The analysis was performed based on a well-known multidimensional methodology, Structural Equation Modeling (SEM), which is employed for the development of the proper model and the estimation of the determinants [17-18].

The analysis is conducted in the European area, among a number of countries for years 2007 and 2009, while dataset were extracted from Eurostat [9]. Year 2007 is the time when the EC stated a more focused strategy regarding the development of UCNP and year 2009 corresponds to the most recent available data. Comparison and analysis of the results between these two years are expected to reveal the progress made, due to the introduced strategies.

The model was validated by a number of goodness- of- fit statistics [19], such as the comparative fit index (CFI), the Tucker-Lewis index (TLI) , and the Root Mean Square Error of Approximation (RMSEA). The results are presented in Fig. 2, where values in parentheses refer to year 2007.

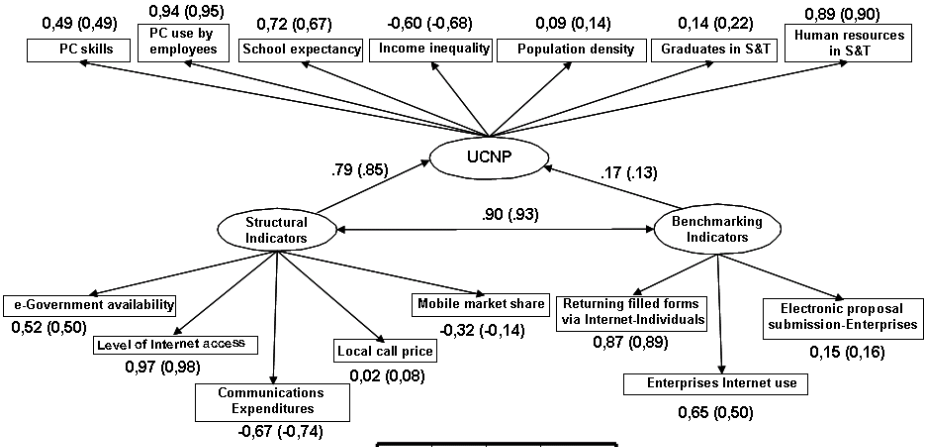

	CFI	TLI	RMSEA
2007	0.944	0.932	0.067
2009	0.921	0.903	0.074

Fig. 2 Estimated standardized regression weights among variables and correlation results; (i)between UCNP and the two main indicators and (ii)between the two main idicators.

According to these results, both indicators have an important influence on UCNP maturity, since they are both positively related with the latent variable of the considered model. More specifically, structural indicators have a greater impact on the maturity level than the benchmarking ones. For this reason, public and EC practices should continue to aim at further expansion of the Internet access. Given the general

economic crisis, alternative access technologies, such as mobile which requires lower initial investments, should be considered and included into the developing business plans. Moreover, Communications expenditures are the second dominant factor in this group of indicators. Thus, as the utilization maturity grows, the necessity for corresponding investments will be minimized. This suggests that, despite the fact that Europe considers telecommunications growth to be very important there are still restrictions in terms of the availability of required infrastructures and investments.

In addition, it seems that mobile telephony operators tend to merge in order to achieve a larger market share. However, market concentration has a negative impact on UCNP maturity, as it is commonly related with higher prices and lower quality of services. Though, the development of UCNP maturity seems to boost competition in the mobile market, restraining the market share of the leading operator in this sector. Moreover, it is worth noting that Price does not seem to influence utilization maturity, indicating an inelastic behavior.

On the other hand, e-services show an increase, in terms of their impact on UCNP maturity. Therefore, decisions makers should rethink the content of the offered e-services, as they have to meet the needs of citizens and business organizations.

Additionally, income inequality has an inverse relationship with UCNP maturity level. Thus, the growth of the last-mentioned variable reflects the minimization of the technological exclusion derived from economical reasons. Moreover, educational variables, such as graduates in S&T and School expectancy could be enhanced by the utilization maturity level. Finally, factors describing the working environment, such as the required qualifications –Computer Skills, Computer Use etc- and the market share of Science and Technology sector in terms of the percentage of the labor force, are heavily influenced by UCNP maturity level.

5 Assessment of Information and Communications Technology maturity level

This work aims to express the level of ICT development in each country, through an adequate conceptual model. To do this, a latent variable is defined, the ICT maturity level, in the context of Structural Equation Modeling (SEM) [17-18]. The ICT maturity level is determined by three elements (sub-indices), “Access”, “Use” and “Skills”, which are also latent variables and which in turn are described by observed factors. These factors are indicators which derive by the countries’ statistical data. Since the estimation of ICT evolution by the means of an appropriate metric is important, the ICT Maturity Level Index (IMLI) is introduced, to improve the ICT Development Index (IDI) proposed by the ITU [20]. Structural equation modeling allows for the precise estimation of the sub-indices and indicators’ weights, which are different from those set by ITU, especially when applied to developing countries. Calculation of the IMLI was based on statistics from 159 countries [16]. Due to their huge differences in the corresponding economic wealth, the considered countries were grouped according to their income level, as expressed by the GNI per capita, in order to avoid bias in estimations. The common approach to do this is the clustering of countries into devel-

oped and developing. This separation, proposed in the present work, proved very important according to the results.

It could be probably expected that the impact of sub-indices and their indicators over the ICT maturity level would change slowly over time. However, within short periods of time, significant changes are observed on indicators statistics describing "Access", "Use" and "Skills" led to the examination of their dynamic nature. Therefore, the analysis was performed for two consecutive years, revealing marginal changes regarding the sub-indices but significant for their indicators. Thus, a different ranking of the countries with respect to IMLI, from year to year, was apparently obtained. In this way, the time behavior of the model was highlighted.

The first hypothesis of the proposed model is that there is a link between GNI per capita and IMLI. The second hypothesis is that IMLI is related with "Access", "Use" and "Skills" respectively. According to the above hypotheses, the new entity that evaluates the ICT Maturity Level (IML), is described by the following Eq.2:

$$IML = w_{Access} * Access + w_{Use} * Use + w_{Skills} * Skills \tag{2}$$

where the considered sub-indices ("Access", "Use" and "Skills") participate with a corresponding weight (w_{Access}, w_{Use} and w_{Skills}). The sub-indices are described by a number of observed indicators. Following the above consideration the estimation of IMLI is based on Eq.3:

$$\begin{aligned} IMLI = {} & w_{Access} * (w_{A1} * A1 + w_{A2} * A2 + w_{A3} * A3 + w_{A4} * A4 + w_{A5} * A5) + \\ & w_{Use} * (w_{U1} * U1 + w_{U2} * U2 + w_{U3} * U3) + w_{Skills} * (w_{S1} * S1 + w_{S2} * S2 + w_{S3} * S3) \end{aligned} \tag{3}$$

In the above expression of IMLI, Ai, Ui and Si are the corresponding observed variables, for each country, while w_{Access}, w_{Use}, w_{Skills}, $w_{A,i}$, $w_{U,i}$ and $w_{S,i}$ are the estimated weights resulting from SEM approach. The implementation of the proposed model was performed by using Ordinary Least Squares (OLS).

Following the methodology described above, the corresponding weights for the sub-indices w_{Access}, w_{Use} and w_{Skills}, as well for the indicators w_{Ai}, i=1,2,..5, w_{Ui}, i=1,2,3, and w_{Si}, i=1,2,3, can be estimated using the SEM approach. Based on these weights, the index IMLI can be calculated for each country, according to Eq.3, assuming that within the same cluster they do not vary significantly among the countries. The above weights were calculated using two different datasets, one formed by the data of the 75 developed countries and the other formed by the data of the 84 developing countries.

According to the results, income level has a significant effect on ICT development, which is more intense for countries with low incomes. Furthermore, it is important to notice that the estimated weights of the three sub-indices of the model have quite different values between the developed and the developing countries. "Access" turned out to be the dominant component regarding IML for both groups of countries. These findings indicate that access availability has a substantial influence on ICT development. More specifically, access speed in developed countries is the most important indicator affecting IML. On the contrary, in developing countries the impact of telephone lines on IML reflects the lack of the required infrastructures.

The "Use" element is significantly more important for developed countries than for developing. In particular, fixed broadband connections are the dominant indicator for further ICT growth. Not surprisingly, the importance of "Skills" is lower in developed countries, indicating the sufficient educational level they have reached. On the other hand, in developing countries "Skills" exert great influence on IML. Hence, it is proved that educational level is directly related to the ICT development. Therefore targeted actions should take place, especially in developing countries, where even nowadays literacy rates remain in extremely low levels.

Europe is the leading area, in terms of ICT development, as the great majority of the countries belonging to the first 20 places in world ranking are European. This conclusion is in line with ITU [20], indicating that ICT uptake in Europe is faster than in other areas, probably due to common regulatory framework and orchestrated policies implemented in all countries. There is an observable improvement each year with respect to ICT development, especially for the developed countries. Further, best practices could be analyzed and adapted to other countries' increasing ICT diffusion.

6 Conclusions

The contribution of this dissertation lies in the results of the above studies. Initially, an index for the assessment of digital divide convergence in terms of broadband penetration is proposed which in turn it can be used as a reliable tool in order to forecast the progress made in the aforementioned convergence process. Next, the evaluation of a number o socio-economic factors affecting broadband penetration in relation to the diffusion level took place. It turned out that before inflection point, e-Government on line service availability is the dominant factor for further broadband diffusion. On the contrary, after inflection point, working environment, and specifically, persons employed using computers connecting to the Internet, fosters broadband adoption. After that, the relation between GNP per capita and broadband diffusion is examined by applying a causality test. It is proved that the above relation depends on the level of broadband penetration. More specifically, before inflection point GNI per capita can cause positive change in broadband penetration, while after inflection point broadband diffusion can cause an increase in income.

Subsequently a new metric, the "Utilization of Communications Network potential –UCNP), is determined to evaluate the effectiveness of public practices aiming to enhance ICTs. According to results, it seems that UCNP is affected more by structural indicators then benchmarking indicators. Finally, ICT maturity level is estimated based on "ICT Maturity Level Index- IMLI" in accordance to "ICT Development Index-IDI", proposed by ITU. The estimation is made with precision and according to results, the influence of sub-indices "Access", "Use" and "Skills" and their corresponding variables, differs between developed and developing countries.

7 References

1. Kyriakidou, V., Michalakelis, C., Sphicopoulos, T.: Digital Divide Gap Convergence in Europe. Technology in Society. 33 (3-4), 265-270 (2011)
2. Kyriakidou, V., Michalakelis, C., Sphicopoulos, T.: Study of broadband penetration influential factors, with non-predetermined functional form. Technological Forecasting and Social Change (accepted July 2012)
3. Kyriakidou, V., Michalakelis, C., Sphicopoulos, T.: Broadband penetration as an economic growth accelerator. (under review, 2012)
4. Kyriakidou, V., Michalakelis, C., Sphicopoulos, T.: Utilization of Communications Network Potential: Public practices and effects. Government Information Quarterly. 29 (2), 182-191 (2012)
5. Kyriakidou, V., Michalakelis, C., Sphicopoulos, T.: Assessment of Information and Communications Technology maturity level. Telecommunications Policy (accepted August 2012)
6. Kyriakidou, V., Katsianis, D., Orfanos, I., Chipouras, A., Varoutas, D.: Business Modeling and Financial Analysis for Metropolitan Area Networks: Evidence from Greece. Telematics & Informatics. 28 (2), 112-124 (2011)
7. Kyriakidou, V., Michalakelis, C., Varoutas, D.: Applying Zipf's power law over population density and growth as network deployment indicator. Journal of Service Science and Management. 2 (6), 132-140 (2011)
8. European Commission. Information society: digital agenda for Europe, http://ec.europa.eu/information_society/digitalagenda/index_en.htm; 2010
9. European Commission, Eurostat. Available: http://ec.europa.eu/eurostat/
10. Gompertz, , B.: On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. Philosophical Transactions of the Royal Society of London. 115, 513-585 (1825)
11. Naik, P., Hagerty, M., Tsai, C-L.: A new dimension reduction approach for data-rich marketing environments: sliced inverse regression. J. Mark. Res. 37, 113-24 (2000)
12. Akaike, H.: Information theory and an extension of the maximum likelihood principle. In: 2nd International Symposium on Information Theory, pp. 267-81. Akademia Kiado, Budapest (1973)
13. Schwartz, G.: Estimating the dimension of a model. Annals of Statistics. 6, 461-464 (1978)
14. Hanke J. H., Wichern, D. W.: Business Forecasting. Pearson International Edition, New Jersey (2009)
15. Granger, C.W.J.: Investigating causal relation by econometric and cross-sectional method. Econometrica. 37, 424–438 (1969)
16. World Bank: http://data.worldbank.org/
17. Wright, S.: Correlation and causation. Journal of Agricultural Research. 10(7), 557-585 (1921)
18. Wright, S.: The method of path coefficients. Annual Mathematical Statistics. 5, 161-215 (1934)
19. Bagozzi, R. P., Yi, Y.: On the evaluation of structural equation models. Academy of Marketing Science. 16, 74-94 (1988)
20. ITU: Measuring the Information Society - The ICT Development Index. (2009)

Web based Adaptive Educational Games - Exploitation in Computer Science Education

Konstantinos Maragos *

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications

kmaragos@di.uoa.gr

Abstract. In the context of this dissertation, issues related to design and development of educational games that support the teaching of programming are studied. Specifically, the dissertation is placed in the context of supporting the learning process by utilizing the Web and especially Web based, Adaptive Educational Games. A Web based Adaptive and Multiplayer Educational Game TALENT (Teaching ALgorithms EnvironmeNT) was designed and implemented, which motivates and supports students in learning programming.

Keywords: Programming, Mini-Languages, Adaptive Systems, Educational Computer Games, Computer Science Education

1 Introduction

Much of the existing research, in the computing education literature at least, focuses on new and interesting ways to teach programming [1]. One way is to focus on game based learning, which can be defined as the use of a computer game based approach to deliver, support, and enhance teaching, learning, assessment and evaluation [2]. Computer games are well suited to use within an educational environment because they build on theories of motivation, constructivism, situated learning, cognitive apprenticeship, problem-based learning and learning by doing [3]. Furthermore, educational computer games improve intrinsic motivation through challenge, curiosity, control and fantasy [4].

On the other hand educators have to get more attention on the fact that today high school students have been brought up in a technologically rich environment. They belong to the generation called “digital natives” or “net generation” [5]. These students were "born digital" and they have been heavily influenced by the latest highly interactive and individual technologies such as iPods, iPhones, iPads, Wii, Xbox, Playstation game consoles, as well as Wi-Fi Internet access and graphic rich multiplayer Internet gaming. In addition, as result of the Internet, students learn much more collaboratively today than in previous generations. Thus, there is an important need

*Dissertation advisor: Maria Grigoriadou, Emeritus Professor

for educators to embrace and adopt approaches to teaching and learning that are better suited to the learning styles that the younger generation of learners now adopt and provide a more stimulating and engaging learning environment.

According to [6], the most common method of introducing students to programming is the gradual presentation of programming structures for a general purpose computer language and also, the presentation of increasing difficulty problems, based on these structures. This approach is inadequate for students, as it provides them a series of obstacles [7]. This is augmented by the fact that students today seem to be motivated by graphically rich learning environments that have multimedia characteristics [8].

An alternative and more effective approach for introducing students to programming is to provide them with mini environments, based on micro worlds and mini languages [1], [9]. According to this approach, students learn to program by instructing a virtual entity (e.g. turtle, robot) in a virtual micro world. This virtual entity can be controlled by a small set of instructions and basic programming structures like. conditions and loops. There are several examples of games that provide interaction through a mini environment and a mini language like Robocode (2000), Gun Tactyx (2004), Ceebot (2005), Robomind (2005), Sheep (2006), Rapunsel (2007), LightBot (2008) and Marvin's Arena. The main issues of these educational games are the type and the story of the game, the programming language they are using, the method of programming, the capabilities of the execution of the program, the ability to support many users at the same time that are able to communicate with each other and the exploitation of the web technologies.

In addition, all the above games they do not take into account the student model which provides an additional resource of the learning process that is believed to enhance the learning [10], [11]. The student model is a representation of understanding, difficulties and misconceptions of the student [12] enabling the system to adapt to student needs and current learning requirements [13].

In the context of supporting programming education the design of a web based multiplayer adaptive educational game is proposed, referred to as TALENT (Teaching ALgorithms EnvironmeNT) [14], [15], [16], [17], which (i) is based on a story-context and the activities are grouped according to learning goals, (ii) uses specially designed mini language that supports all programming structures: sequence, selection and repetition, (iii) utilizes visual programming techniques to create programs with no limitation on the number of program's instructions, (iv) allows the programming of variables corresponding to the properties of objects transforming the virtual world from static to dynamic, (v) allows a step by step execution of program's instructions, (vi) creates and utilizes the model of student providing adaptation, (vii) exploits the potential of the internet, does not require any installation, supports the simultaneous presence and interaction of many users and the synchronous communication between them as well and (vii) includes the design of an authoring tool.

2 The TALENT Environment

TALENT is designed as an adventure multiplayer web based educational game for teaching programming. Adventure games offer powerful opportunities for learning and development of problem-solving abilities [18]. Adventure games are also suitable for science concepts that may be hard to visualize or manipulate with concrete materials [19] and programming is a kind of such a concept. In addition, adventure games, which are consisted of tasks, provide a clear progression overall in the game [20].

2.1 Programming in the game

In the educational game TALENT, when a student is engaged in a programming activity, they have to use the mini language to instruct a robot vehicle to do the necessary actions in order to succeed in mission in the microworld. More complex missions require more complicated programming. For this reason the student in TALENT is able to construct more advanced programs by using programming structures (e.g. selection, repetition). The student is able to construct the programs in the environment with the support of a prototype online programming tool (fig. 1) which combines a program editor and an interpreter running on the web. The online program editor permits students to drag instructions, represented by buttons, from the instruction toolbox and drop them on programming area. The student's program is formed by this drag and drop of several instructions and programming structures.

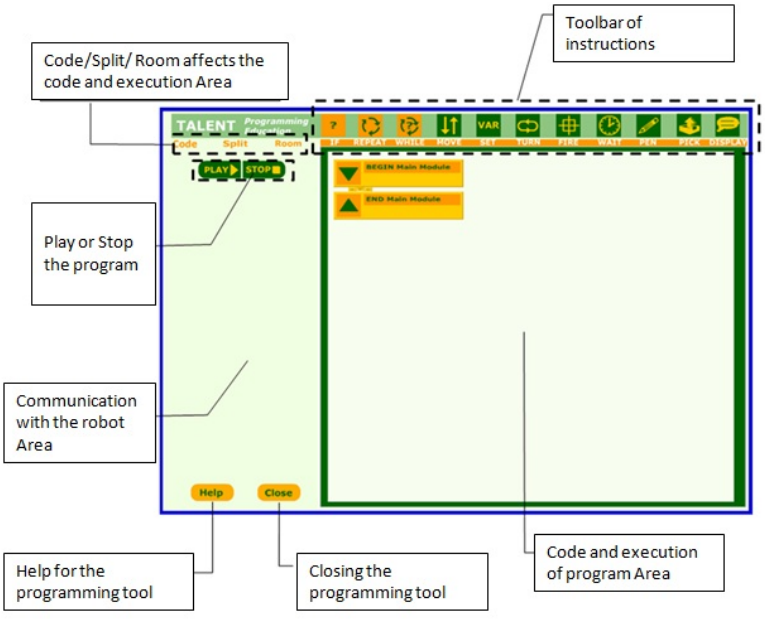

Fig. 1. The programming tool of the Game

When the program is ready, the student can go on with execution using the embedded interpreter. On program execution, the environment maintains both the micro world and the student's program, visible on the screen. The program executes one instruction at a time, while the interpreter highlights programming constructs in the source code as they are being executed and the effect is simultaneously shown in the micro world. This makes the connection between a mini language instruction or construct and its effect on the micro world obvious. The programming tool has no limitation about the number of the instructions, a drawback found in other games.

2.2 Student Modeling and Adaptation

The first step to create an adaptive educational game environment is to create several different activities, indicate for each the learning goals involved and then group these activities into activities groups [21]. Activity is the basic unit of the game structure and indicates a programming task to be performed. Following this methodology we divided the whole game world into three tiers Map, Place and Mission. The Map tier is the general map of the game world that is presented to the student when the game starts and outlines the places that a student can visit. A Place is the part of the game which outlines the missions that are proposed to the student. Finally, Mission corresponds to a programming activity presented to the student. Each programming activity has its own learning goals. There exist many learning goals concerning programming concepts like the use of variables, selection and repetition. Finally, a set of programming activities with the same learning goal is grouped together forming the Activity Groups. The number of the different activity groups is equal to the number of the different learning goals.
The first time the student enters the game, he has to register. During registration the student gives his own personal characteristics (age, gender etc,), chooses a nickname, which must be unique, and builds the presence of his own avatar so to be distinguished by the other students. These data records are saved in the game database and are associated with the student model.

The TALENT game gathers information about learner's navigation, the use of tools, the progress in programming activities and the achievement of learning goals. In addition, data that concerns the sequence of places and programming activity selection, the total time spent on a place or a programming activity, the total number of times that the student asked for help in a programming activity, the visits to a programming activity or a place, the number of times he communicated with the other students in the game and the progress on learning goals are saved in a separate database record according to student's nickname. The system uses all of these data records to build the student model and to adapt on students abilities with curriculum sequencing and navigation support. This means that system proposes the next mission to the student and also helps the student to navigate to this mission. Student model is visible by the student providing student with power over his learning [22]. Furthermore, one can compare the information in his model with that produced with the

mean data of all the other students that is argued to motivate student to try to succeed more in the activities [23].

3 The Empirical studies

For the evaluation of the Educational Game TALENT, two empirical studies were carried out, the first with university students and the second with high school students. In the following subsections we are going to present the procedures and the results of the two studies.

3.1 The 1st empirical study - Use of TALENT by University Students

The first empirical study was conducted during the lesson "Informatics and Education" of the Informatics and Telecommunications Department of the University of Athens. The participants of the research were 122 four year graduate university students of the lesson. The students were already taught concepts related to ICT in Education, educational software, educational games, educational programs, simulations, programming languages, educational resources on the web, educational case studies, distance learning, the role of the teacher and evaluation of educational environments.

The evaluation of TALENT was given as one of the written exercises of the lesson. Students were asked to use the programming environment of the game and then to rate the system's usability, to write down advantages and disadvantages, and finally to propose didactic scenarios to be embedded in the game. For the rating of the system usability the Computer System Usability Questionnaire (CSUQ) [24] was used. The questionnaire includes 19 questions about system usability in a 7-point Likert scale [25] (7=strongly agree, 1=strongly disagree).

Furthermore, in order to collect the students' opinions, comments and suggestions about the operation of system components a second questionnaire with open type questions was used. Indicative questions of the second questionnaire were "what do you think about the usefulness of TALENT in the process of teaching elementary programming concepts", "summarize the advantages and disadvantages of the use of TALENT", and "what was your experience from the design of didactical scenarios with TALENT's mini language".

The research lasted about a month and the procedure of the research consisted of the following three phases:

Phase 1: Prior the evaluation, the students participated in a four hour class lesson. The lesson was divided in one hour lecture and three hours in the computer laboratory. During the one hour lecture several subjects about educational computer games, design issues and results from game based learning research was discussed. During the three hour lesson in computer laboratory several paradigms of educational computer games especially for computer programming were presented, as well as the syntax and the educational tool of the mini language of the TALENT.

Phase 2: After the four hours lesson the university students had to deal with an exercise which was the use of the game mini language in given didactic scenarios about

the three programming structures, sequence, selection, repetition. Specifically, students had to program some activities using the mini language tool of the game.

Phase 3: After the use of the programming tool of the game, students had to answer the two questionnaires, the Computer System Usability Questionnaire and the questionnaire with open type questions about the advantages and disadvantages of the system as an educational tool. Finally they had to propose a programming activity that could be embedded in the game.

3.2 The 2nd Empirical study - Use of TALENT in Classroom settings

The second empirical research was conducted in the curriculum of the lesson "Informatics Applications" in a high school of Athens, Greece. Overall, 65 students (n=65), 33 females and 32 males participated in the research. Students were at the age of 15-16 and were already familiar with programming, as they were taught programming languages from the junior high school. The research was carried out in the computer laboratory of the high school and the duration of the research was 8 weeks. At the beginning, all the students participated in a pre-test. After the pre-test, students were randomly assigned to one of two groups, the control group and the experimental group. There was no significant difference found on the t-test performed between the two groups according to the pre-test performance ($t=-0.255$, $df=63$, 2-tailed $p=0.799$).

The concept of "if-then-else" and "while loops" in informatics teaching was used as the experimental content. At the end, all the students participated in a post test. The procedure of the 8 weeks research consisted of the following five phases:

Phase 1: Introduction to TALENT game environment (1st and 2nd week) – lasted 4 hours. During the first phase the components of the game was presented to students (2 hours) and all the students were engaged in programming tasks about the concept of variables (2 hours).

Phase 2: All the students completed the Computer System Usability Questionnaire (CSUQ), like the university students, about the system usability (3rd week) – lasted 1 hour.

Phase 3: Pre-test (3rd week) - lasted 1 hour. All the students took the pre-test consisted of questions and exercises about "if-then-else" and "while loops".

Phase 4: Working in groups (4th-5th-6th-7th week) - lasted 8 hours. Experimental group studied the concepts of research using only the TALENT environment. The students used the game and mini language tool of the game to program the game activities. An example of such a programming activity for the concept of selection control structure is shown in figure 3. The role of the teacher was limited on the support of system operations, when needed. On the other hand, the control group participated in a traditional lesson where the teacher carried out lectures about the main concepts of the research. During the traditional lesson, the teacher made questions in order to address the prior knowledge, explained the concepts, answered queries and solved exercises.

Phase 5: Post-test (8th week) - lasted 1 hour. All the students from the experimental and control group participated to a final written test.

The research tools were the questionnaire, the log files created by the students' actions in the game and finally the two tests, pre-test and post-test. The questionnaire was based on Computer System Usability Questionnaire (CSUQ) extended with three open questions in order to extract students' opinions about the advantages and disadvantages of the system. The log files used in the research were automatically created by the environment from students' actions. Data in the log files consists of a) the username of the student in the game, b) the student's action (e.g. entering in programming activity, clicking on the help button, sending message), c) the precise time (date and time) the action was done. Data in the log files are used by the system to construct the student model at runtime.

Finally, the first written test (pre-test) was consisted of two programming activities about the research concepts. The goal of the pre-test was to extract students' prior knowledge. The second written test (post-test) was carried out with the completion of the research. It was consisted of five closed type questions to control further knowledge of the concepts and two programming activities. The programming activities were the same as in the first written test.

4 Results

4.1 Results from the 1st Empirical study

The results of the 1st study, according to the score of the questionnaires and the students' answers, can be divided in three categories: system usability, the usage of the system as an educational tool for teaching programming and the capability of designing and embedding new programming activities. The results for these categories were:

- **The usability of the system**: According to the score of Computer System Usability Questionnaire, the rating of the system usability is high. The final mean of the score in each question and for all the system components is 5.18 for a maximum of 7 (SD 1.2). Making the reduction in [0, 100] scale it arises a total score of 74/100 for the whole system usability. In spite of the high total scoring, there were two questions with a low score. These questions where question 9 that "system gives error messages that clearly tell me how to fix problems" where mean score was 3.78 (SD 1.8) and question 11: "the information (such as online help, on-screen messages, and other documentation) provided with system is clear" where mean score was 4.08 (SD 1.69). The low rating in these two questions shows that students need more guidance and better advises from the environment. In order to support these needs some improvements were made in online help and a system tutorial about the use and the capabilities of the mini-language programming tool was developed.
- **The usage of the system as an educational tool**: According to the university students' answers the system is pleasant, is attracting the interest, increases motivation, has a nice interface, is characterized by simplicity, functionality and ease of learning. Furthermore, one of its major advantages is that is web based and it is not depended on particular architecture or installations. Indicative students' comments

were: “Very amusing and satisfying”, “interesting activities with high level of creativity that it will attract curiosity and interest of students”, “Simplicity, ease of understanding and learning”, “with fantasy and desire you can make nice exercises for children”, “very nice interface, drag and drop f instructions is very useful”, “it is very positive that works in web browser”, “it is important that it does not need installation but only an internet connection, that simplify things”.
As a tool for supporting the teaching of programming the answers of university students in the research stated that it is very easy for someone to learn programming with the mini language, drag and drop of instructions makes programming easier, it eliminates syntax errors and is very useful for novice programmers. Indicative students’ comments were: “system simple use helps and encourage students to learn programming through gaming”, “very clever idea that differs from today traditional and boring method of teaching programming”, “an innovative method to introduce children in programming”, “drag and drop is a very intuitive method of inserting instructions and making a program eliminating syntax errors”, “tracing of the program in addition to execution visualization helps student to learn programming easier”.

- **Designing and embedding new programming activities**: In their work students designed several didactic scenarios that need the use of programming with system mini-language tool. Students designed prototype didactic scenarios that could be embedded in the knowledge base of the system. At the end they wrote down their thoughts and proposals about this task. According to students’ answers the design was very pleasant and developed their fantasy, instructions of mini language are sufficient for designing new didactic scenarios and proposed an authoring tool that could help in embodiment of the new activities in the game. Sample students’ answers were: “Scenarios design was very pleasant, I was feeling like designing a game”, “I was impressed! The design of didactic scenarios improved my fantasy limit”, “mini language has all the instructions I needed to design the scenario”, “an authoring tool could be helpful to try out the proposed scenario”.

4.2 Results from the 2nd Empirical study

The results of the 2nd empirical study are based on the scores of the questionnaires, log files and students’ grades in pre-test and post test and they can be divided in three categories: system usability, the utilization of the game as presented by the students’ models and the usefulness of the system as an educational tool for teaching programming to high school students. The results for these categories were:

- **The usability of the system**: According to the score of Computer System Usability Questionnaire, the rating of the system usability is high, like in the first research. The final mean of the score in each question for all the system components is 5.01 for a maximum of 7 (SD 1.6). Making the reduction in [0, 100] scale it arises a total score of 72/100 for the whole system usability. The interventions in the help component of the system made after the first research seemed to have positive effect on students’ evaluation of the system usability. More specific, question 9:

"system gives error messages that clearly tell me how to fix problems" the mean score was 4.85 (SD 1.3) showing an increase from the first research where mean was 3.78 (SD 1.8). Furthermore, on question 11: "the information (such as online help, on-screen messages, and other documentation) provided with system is clear" the mean score in the second research is 4.86 (SD 1.18) showing again an increase from 4.08 (SD 1.69). The demonstration of the game by the teacher played of course an important role in increasing these ratings. In spite of that increase the help component of the system seems to demand further development in order to provide more support to the students. From the students' answers in open type questions of the questionnaire was concluded that the game environment is pleasant, easy to use, arouses interest and supports motivation. From the observation of students' behavior it was clear that students were highly motivated and were trying to succeed in programming activities the same time they had fun. Sample comments were: "I'm satisfied with all", "it is comprehensible and funny", "It was an enjoyable and creative activity", "very easy handling, very nice in general" and "It could be better if there were more activities".

- **Students' models data and system utilization**: As already stated, the student model consists of data about the knowledge level of the student, the time spent in programming activities and also the time spent in other components of the system like help, chatting and wandering in the game world. The data for every action of a student are saved in log files automatically and permit the system to construct the individual learner model at runtime. From the log files of the second research the mean time of the overall system usage for every student in the experimental group calculated equal to 339.18 minutes. During this time students used the programming tool of the game for 235.16 minutes, searched for help for 40.7 minutes, chat with other students for 36.17 minutes and 17.13 minutes were spent on wandering in the game world.

  From the measures it is concluded that the 70% of the total time was spent on programming activities. The percentages make clear that the environment is doing well in arousing the interest and motivating the student to be engaged in programming activities. Furthermore from the total 235.16 minutes participating in programming activities, a percentage of 54% of the time was spent on selection control structures (if then else) and the rest 46% of the time was spent on repetition structures (while). Students of the experimental group worked on 7 programming activities, 4 for selection structure (57% of total activities) and 3 for the repetition one (43% of total activities). From this observation concludes a good analogy to the difficulty of the activities. Finally, as the level of knowledge is concerned, the mean of knowledge level of the students of the experimental group was70%. For the computation of the mean was used the mean of successful activities for every research concept, that was 74% in selection structure (2.94 successful to 4 total activities) and 75% in repetition structure (2 successful to 3 total activities). The result is in agreement with the mean grade of the experimental group in the post-test, which is presented below.
- **The usefulness of the system as an educational tool for teaching programming to high school students:** As stated above, during the 3rd week all of the students

participated in first written test (pre-test) and after that the students were randomly assigned to control group and experimental group. According to the statistical analysis from the t-test in pre-test grades there was no significant difference between the mean grades of the control group (mean=12.13, SD=8.27) and experimental group (mean=11.61, SD=8.11). On the other hand the grades in the second test (post-test) for the students in experimental group was significant higher (mean=68.61, SD=17.22) than those of the control group (mean=48.63, SD=21.18). The results of pre-test and post-test for the two groups are presented in table 1.

Table 1. Mean and SD for the two groups in pre-test and post-test

Students' Group - Method of Teaching – number of Students	**Performance Pre-test Grade**	**Performance Post-test Grade**
Control Group - Traditional Teaching (n=32)	12,13 (8,27)	48,63 (21,18)
Experimental Group - Working with TALENT (n=33)	11,61 (8,11)	68,61 (17,22)
Total (n=65)	11,87 (8,19)	58.62 (19,2)

From the results it is concluded that students in both groups, control and experimental, improved their performance by increasing the average rating from 11.87 (SD=8.19) in pre-test to average 58.62 (SD=19.2) in post-test. As implied by comparing the results of the two groups, students in experimental group had a significantly higher performance than students in control group. This indicates that there is some evidence that the use of TALENT contributes positively in learning programming concepts than the traditional method of teaching.

5 Conclusions

In the context of this thesis, we presented TALENT a web based adaptive educational multiplayer game for supporting the teaching of programming. TALENT is a combination of an adventure and a role play game and it utilizes web technologies in order to motivate the student to engage in programming activities. An innovative web programming tool which exploits visual programming techniques and embed an interpreter which executes student's program step by step. TALENT adapts to student abilities by providing adaptation in curriculum sequencing and adaptive navigation support. It also includes an authoring tool which can be used by an author to add or alter

educational content in knowledge base of the system or used by the teacher to monitor students' performance.

Two empirical studies have been conducted. The results from these studies are positive and indicate acceptance and satisfaction of participating students for the usability of the environment. Furthermore, the environment achieves to attract the interest of students and motivate them to keep on their programming efforts. Finally, the performance of students in an experimental group as opposed to these of a control group shows some evidence that teaching with TALENT has a positive contribution to the learning process in contrast to the traditional way of teaching.

Our future plans include the development of appropriate modules that would increase the collaboration between the students, the analysis of students' communication and also the exploitation of different adaptive methods and techniques in order to provide more personalized learning.

References

1. T. Jenkins, "On the difficulty of learning to program", *In Proceedings of 3rd Annual LTSN_ICS Conference,* The Higher Education Academy, 2002, p.p. 53-58.
2. T.M. Connolly, and M.H. Stansfield, From eLearning to games-based eLearning: Using interactive technologies in teaching an IS course, *International Journal of Information Technology Management,* vol. 26, no. 4, 2007, p.p. 188-208.
3. T.M. Connolly, E. McLellan, M.H. Stansfield, J. Ramsay, and J. Sutherland, "Applying computer games concepts to teaching database analysis and design", *Proceedings of the International Conference on Computer Games, AI, Design and Education*, Reading, 2004.
4. M. R. Lepper, and T.W. Malone, "Intrinsic motivation and instructional effectiveness in computer-based education", In R. E. Snow & M. J. Farr (Eds.). *Aptitude, learning, and instruction: Vol. III. Cognitive and affective process analyses,* Hillsdale, NJ: Erlbaum, 1988, p.p.255-286.
5. S. Bennett, K. Maton, and L. Kervin, The digital natives debate: A critical review of the evidence, *British Journal of Educational Technology* vol.39, no.5, 2008, p.p.775–786.
6. P. Brusilovsky, E. Calabrese, J. Hvorecky, A. Kouchnirenko, and P. Miller, Mini-languages: A way to learn programming principles, *Education and Information Technologies*, vol. 2, no.1, 1997, p.p. 65-83.
7. R. Reichert, "Theory of computation as a vehicle for teaching fundamental concepts of computer science", Dissertation No. 15035, ETH Zürich, 2003.
8. M. Guzdial, and E. Soloway, Teaching the Nintendo generation to program, *Communications of the ACM*, vol.45, no.4, 2002, p.p.17–21.
9. C. Kelleher, and R. Pausch, (2005). Lowering the barriers to programming: A taxonomy of programming environments and languages for novice programmers, A*CM Computing Surveys*, vol.37, no.2, 2005, p.p. 83-137.
10. J. Kay, "Learner Know Thyself: Student Models to Give Learner Control and Responsibility", *Proceedings of International Conference on Computers in Education, Association for the Advancement of Computing in Education (AACE)*, Z. Halim, T. Ottomann & Z. Razak (eds), 1997, p.p. 17-24.
11. K. Maragos, and M. Grigoriadou, "Towards the design of Intelligent Educational Gaming systems", *Proceedings of Workshop on Educational Games as Intelligent learning envi-*

ronments, Artificial Intelligence in Education, University of Amsterdam, Amsterdam, 2005.
12. S. Bull, "Supporting Learning with Open Learner Models", *4th Hellenic Conference with International Participation: Information and Communication Technologies in Education*, Athens, 2004.
13. Y. Cui, and S. Bull, Context and Learner Modelling for the Mobile. *Foreign Language Learner System*, vol.33, no.2, 2005, p.p. 353-367.
14. K. Maragos, and M. Grigoriadou, "Teaching Computer Science Concepts with Educational Computer Games", *Proceedings of EUTIC 2007 International Colloquium on Challenges and Uses of ICT Media and Information diffusion: towards an open society*, Laboratory of New Technologies in Communication, Education and the Mass Media of the University of Athens, Athens-Greece, 2007.
15. K. Maragos, and M. Grigoriadou, "Designing an Educational Onine Multiplayer Game for learning Programming", *Proceedings of Informatics Education Europe II*, ACM, Thessaloniki-Greece, 2007.
16. K. Maragos, and M. Grigoriadou, "Learner Modeling and Adapted Interaction in Educational Games", *Proceedings of 80Days' 1st International Open Workshop on Intelligent Personalization and Adaptation in Digital Educational Games*, University of Graz, Graz-Austria, 2009.
17. K. Maragos, and M. Grigoriadou, Exploiting Talent as a Tool for Teaching and Learning, *The International Journal of Learning*, vol.18, no.1, 2011, pp.431-440.
18. A. McFarlane, ed., *Information Technology and Authentic Learning: Realising the potential of computers in the primary classroom*. London: Routledge, 1997.
19. A. Mitchell, and C. Savill-Smith, The use of computer and video games for learning: A review of the literature, 2004; http://www.lsda.org.uk/files/PDF/1529.pdf [Retrieved 10/9/2005]
20. A. McFarlane, A. Sparrowhawk, and Y. Heald, Report on the Educational Use of Games, 2002; http://www.teem.org.uk/publications [Retrieved 10/9/2005
21. R.M. Carro, A.M. Breda, G. Castillo, and A.L. Bajuelos, A Methodology for Developing Adaptive Educational-Game Environments, In P. De Bra, P. Brusilovsky, and R. Conejo(Eds.): *AH2002, LNCS 2347*, Springer- Verlag, Berlin Heidelber, 2002, p.p. 90-99..
22. A. Kerly, and S. Bull, Open Learner Models: Opinions of School Education Professionals, in K. Koedinger, R. Luckin & J. Greer (eds), *Artificial Intelligence in Education*, IOS Press, Amsterdam, 2007.
23. S. Bull, and M. McKay, "An Open Learner Model for Children and Teachers: Inspecting Knowledge Level of Individuals and Peers", *Intelligent Tutoring Systems: 7th International Conference*, Springer-Verlag, Berlin Heidelberg, 2004, p.p. 646-655.
24. J.R. Lewis, IBM Computer Usability Satisfaction Questionnaires: Psychometric Evaluation and Instructions for Use, *International Journal of Human-Computer Interaction*, vol.7:, no.1, 1995, p.p. 57-78.
25. R. Likert, A Technique for the Measurement of Attitudes, *Archives of Psychology* vol.140, 1932, p.p. 1–55.

E-Negotiations for trading Commodities and Services: Predictive Strategies

Marisa Masvoula[1]

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications
marisa@di.uoa.gr

Abstract. The current thesis takes into account the advances in the field of electronic bi-lateral negotiations, adopting state-of-the-art protocols and strategies that characterize the behavior of each party. The research objective is the application of predictive strategies that have their basis on the estimation of the counterpart's next offer, and give the agent the advantage to establish agreements that are more beneficial. An issue that is contemplated is that of risk when employing predictive skills. To enable the adoption of different risk attitudes, a new strategy is proposed, developed and assessed. This thesis also examines the AI-based models that are incorporated in the strategic core of negotiation software agents to estimate the counterpart's next offer. The main problem of the majority of related applications is their inability to capture the dynamics of turbulent negotiation environments, and to provide accurate estimations also in cases where the data distributions change. For this reason, the utilization of models trained with data that are acquired from the current negotiation discourse is considered of vital importance. This thesis sheds light to the application of models that are retrained in each negotiation round and to models that adapt their structure in time. More specifically, application of neural networks that adapt their structure on the basis of a genetic algorithm, and application of an evolving connectionist structure that does one-pass learning of data (eMLP) are examined. Numerous experiments that result from simulations of different negotiation environments and justify the proposed solutions are provided. Finally, future research issues that relate to the domain of application of the proposed strategy, as well as to the utilization of other learning models that estimate the counterpart's next offer, are suggested.

Keywords: electronic negotiations, negotiating agents, neural networks, genetic algorithms, predictive strategies

1 Introduction

There is a grand variety of problems drawn from everyday life where negotiations are evident. A typical list would contain economic transactions, distribution of services, management of business processes, labor negotiations, political and juridical disputes etc. Such scenarios do not always relate to conflict as might be assumed. Negotiation is the key decision-making approach that is used to reach consensus whenever a person, organization or another entity cannot achieve its goals unilaterally. It can be defined as an iterative communication and distributed decision-making process, where participants are searching for an agreement. Negotiation is thus a mechanism that can be used for allocating and sharing resources. The term 'resource' is used in the broadest possible sense and may involve commodities, services, time, money etc. Yet it is not guaranteed that an agreement always exists or that it will be established.

During the last decades, scientists belonging in various scientific areas such as anthropology, psychology and sociology, law, political science, economics, mathematics and computer science have made efforts to model and study negotiation interactions. These efforts have resulted to different methodologies, architectures and approaches.

Among many significant contributions is the transfer of negotiation encounters in electronic settings. Electronic platforms have been designed to facilitate the conduct of negotiations. Furthermore, computer scientists have contributed to the development of software components that either assist negotiators in various stages of the negotiation process (Negotiation Support Systems, NSSs), or in some cases are capable of undertaking stages or even the whole negotiation process. This thesis focuses on the latter category, where negotiation software agents (NSAs) are used for the representation of market stakeholders.

An agent can be viewed as an encapsulated computer system that is situated in an environment and is capable of flexible, autonomous action in order to meet its design objectives. Negotiation Software

[1] Dissertation Advisor: Drakoulis Martakos, Associate Professor

Agents are good representatives of human negotiators. The specific rules of the interaction, which constitute the negotiation protocol, are predefined in negotiations between autonomous agents. Based on the negotiation protocol, agents need to plan their specific actions, their strategy, in order to meet their objectives. The action planning is usually not disclosed to the other participants and takes place before the actual conduct of negotiation (at a pre-negotiation phase). Yet, it is possible that during the interaction an agent reassesses the negotiation problem and adapts his strategy to the responses of his counterpart. State of the art negotiating agents use learning techniques in order to increase their profits or fulfill their objectives. Learning techniques usually assist agents to select optimal or suboptimal strategies and better model their counterparts.

In order to facilitate comprehension of the domain, in this thesis we provide a categorization of strategies that are enhanced with learning techniques and are used by state of the art automated negotiators. In this respect, we devise agents to those who use explorative, repetitive and predictive strategies, either at the planning phase or during discourse [1]. Explorative strategies imply the search for new solutions and are based on trial and error learning processes, such as Q-learning and Genetic algorithms. Repetitive strategies are based on knowledge reuse, and specific knowledge is acquired by repeated execution of actions. Case-based reasoning is one such technique. Finally, in predictive strategies learning is introduced in the form of predictive decision making, where estimations of factors that influence strategy selection or update serve as input to the agents' decision making.

In this thesis we focus on the third category and particularly on agents who update their strategy based on estimations of their counterpart's future responses. The main objective is to induce more satisfying outcomes, by employing the predictive skill. Previous offers and domain-specific information are used as input to the forecaster in order to calculate the next offer of the opponent. This encourages more sophisticated decision making, irrespective of the type of e-market component (negotiation support system or negotiating software agent). Negotiation support systems, enriched with tools that undertake single-lag predictions are very powerful. Carbonneau, Kersten, and Vahidov [2] depict the development of a neural network predictive model in order to facilitate "What-if" analysis and generate optimal offers in each negotiation round. The support tool simulates the possible response to the offer the user is contemplating and assesses offers and counter-offers based on the users' utility function. This capability allows optimization of potential offers, with the purpose of finding the most beneficial for the current situation. A similar negotiation support tool is applied by Lee and Ou-Yang [3] in a supplier selection auction market, where the demander benefits from the suppliers' forecasts, by selecting the most appropriate alternative in each round. As far as autonomous negotiating agents with predictive abilities are concerned, most models extend the formal decision functions presented in [4] [5]. These functions are used to generate the next offer during a discourse, and are based on a series of tactics dependent on time, resources and opponents' behavior. Learning agents manage to control the negotiation outcome by foreseeing the next move of their opponents, as is depicted by the following papers. Papaioannou, Roussaki, and Anagnostou ([6], [7]) developed an agent who applies the predictive mechanism only at the pre-final step of the process, in order to increase the likelihood of achieving an agreement, and to produce an outcome of maximal utility. An older work concerning single-lag predictions in agents' strategy can be found in [8]. Trading scenarios via an internet platform are facilitated with the use of SmartAgent, enhanced with predictive decision making.

2 Dissertation Summary

In related work, the estimation of the counterparts' next move leads to increased utility when agreements are established. However, in some situations agents tend to prolong the negotiation discourse and this increases the risk that negotiation breaks off, as the counterpart may decide to terminate the process.

The first issue we investigate is how such strategies affect the establishment of negotiating agreements. In this vein, we propose a negotiation strategy that introduces a risk related parameter, mainly linked to the prolongation of the negotiation discourse [9] [10]. This parameter allows the specification of different attitudes towards risk, where risk measures an agent's willingness to stay in negotiation in order to use the predictive mechanism more extensively and heighten his gains. For example a negotiator with a risk-seeking behavior would decide to exhaust negotiation time in order to increase his profits, while a risk-averse agent would act more conservatively, and would not risk prolonging negotiation time.

We have considered a number of negotiation environments and we have measured the average increase in utility that incurs to the predictive agent compared to the non-learning one. An agent with a highly risk-seeking attitude achieves on average 12.05% increase in utility, while a predictive agent with a more conservative behavior (risk-averse) achieves 0.94% increase in utility. However, the trade-off of the highly increased utility is the decrease of the number of agreements, which is due to prolongation of the negotiation time. In the case of the highly risk-seeking agent the number of agreements is decreased by 20.78% compared to the non-learning case. To address this issue, we propose a reformed strategy that appropriately sets the risk-related parameter in each negotiation discourse, based on the estimation of the counterpart's deadline. We have conducted a number of experiments with the reformed strategy which resulted to 12.017% average increase in utility, and 0.61% average decrease of agreements. As it is observed, the average increase approximates the one attained by the risk-seeking agents and the average decrease approximates the one attained by the risk-averse agents.

Another issue that is studied in this thesis is the type of learning mechanism that is used by predictive negotiators who estimate their counterpart's future offers. When it comes to forecasting the partners' future offers, techniques can be summarized into those based on statistical approaches (particularly non-linear regression) [11] [12], mathematical models [13] [14], and connectionist approaches, particularly some special types of neural networks, Multi Layer Perceptrons (MLPs) and Radial Basis Function Networks (RBFN) [2] [3] [6] [7] [8] [15] [16]. From the above methods we argue that neural networks are best applicable for the purpose of forecasting the counterpart's future offers. Experiments have shown that mathematical models give poorer results when compared to non-linear regression models [12], and non-linear regression models are very restrictive, since they require specific assumptions regarding the function form of the opponent's strategy. On the other hand neural networks are applicable in the general case, without assuming implicit knowledge of the function that maps input to output data. This is particularly desirable for negotiation forecasting situations where data relations are not known.

In current research approaches neural networks are trained at a pre-negotiation phase with data extracted from past negotiations, and are then used in the current discourse to provide estimations of the counterpart's future offers. However, the accuracy of the forecasting tool depends heavily on data acquired from previous interactions. We investigate how the forecasting accuracy is affected when data distributions change, and propose building and training neural networks with data extracted from the current interaction. We term agents that exploit data from the discourse session-long learning, and prove that a small neural network with few training examples is capable of capturing the negotiation dynamics [17]. In this thesis we introduce two types of session-long learning agents: Static session-long learning agents (SSLAs), who use a neural network with a static structure during the negotiation process, and adaptive session-long learning agents (ASLAs), who use a neural network which evolves its structure and input features based on a genetic algorithm.

To illustrate the superiority of SSLAs compared to agents that employ pre-trained networks (Pre-Trained Agents, PTAs) in cases where data distributions change, we have conducted a number of experiments considering single-issued negotiations and we have computed the absolute error yielded in each decision making step. SSLAs are proved more accurate, as the mean of the errors is reduced by 92.67% compared to the PTAs. However, they do not yield satisfying results in negotiations with short deadline. This is due to the small size of the training data set compared to the number of parameters of the neural network that need to be learned. The incorporation of an optimization technique for the selection of the networks' architecture to address this issue has led to the introduction of ASLA. ASLA evolves its structure and input features with the use of a genetic algorithm in each negotiation round. We empirically prove the superiority of ASLAs compared to SSLAs as far as accuracy is concerned. More specifically, we have conducted a number of experiments and we have computed the mean, the standard deviation and the maximum values of the absolute errors at the end of each negotiation. ASLAs are proved more accurate as the average mean of the absolute errors is reduced by 38.34%, the average of the maximum values is

reduced by 44.75% and the average of the standard deviation is reduced by 38.03%. ASLA is a smoother predictive model as it proves more accurate with decreased standard deviation and maximum error values. However it is not as fast as SSLA and has higher storage requirements, which makes it difficult to apply in real situations.

We also study the use of another adaptive structure, eMLP, which is a simple evolving connectionist system that engages in one-pass, lifelong learning. eMLP adapts its structure with each new training pattern, and is much faster than ASLA as it conducts one-pass learning [18]. However agents enhanced with eMLP structures are less accurate than agents enhanced with MLP structures.

The idea of static session-long learning agents is extended to support multi-issued negotiations. Forecasting is again conducted with the use of Multilayer Perceptrons (MLPs) and the training set is extracted during the negotiation session. Two cases are examined: one where separate MLPs are used to estimate each negotiable attribute and one where a single MLP is used to estimate the counterpart's response. It is shown that simple MLPs with one hidden layer are adequate for forecasting the counterpart's offer vectors, and are tested in order to find the appropriate number of nodes on input and hidden layer. The network that yields the lowest error, incurs to the predictive agent 10.78% average absolute increase of his individual utility (gain), which is close to the increase in utility incurred by the highly risk-seeking agent enhanced with the ideal forecasting tool.

3 The Proposed Predictive Strategy

In this section we provide an analytical description of the predictive strategy and discuss how it affects the negotiators that adopt it. If the two agents do not employ any learning technique, and each applies the default strategy, as illustrated in [4], an agreement will be established at a point which we term the "Meeting Point" (MP). The proposed predictive strategy is based on the assumption that in each decision making step, if the negotiators decide to send a counter-offer, the risk-averse agent will adopt a conservative behavior and generate an offer that will be accepted by his counterpart, while the risk-seeking agent will provoke gradual concessions of his counterpart, so as to increase his individual gain. The main objective of the proposed predictive strategy is to prolong negotiation beyond MP and increase the incurred utility of the agent, by taking into account the two extreme attitudes. In this respect, the agent pre-defines a reference point that is related to his willingness to prolong the discourse. He adopts the risk seeking behavior until that reference point and the risk-averse behavior from the reference point until expiration of his deadline or termination of the discourse. This behavior is illustrated in Figure 1.

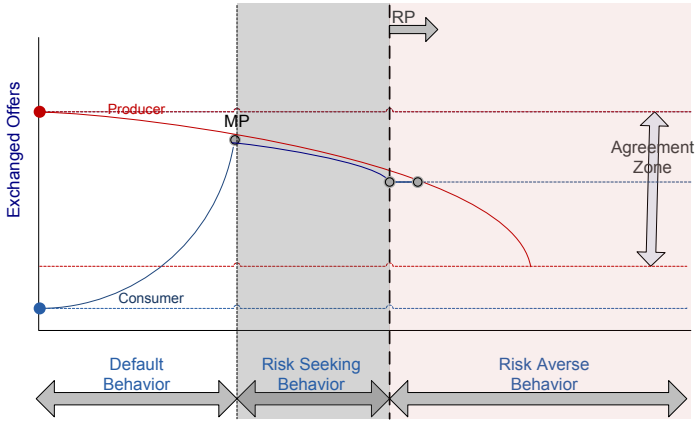

Fig. 1. The Proposed Predictive Strategy

It is essential to distinguish the different usages of the predictive skill. When agents employ a predictive strategy as the one described by Oprea [8], they run into the danger of prolonging the negotiation process, and increase the risk the other party exhausts his deadline and walks out of the process. Prolongation is due to the fact that the counterpart tends to respond with a counter-offer to the predictive agent, which is the case of a risk-seeking attitude. On the other hand, when agents employ a predictive strategy at the pre-final step of the process, as the one discussed in [6] [7], they manage to increase the number of successful negotiations if the counterpart has equal or higher deadline than the predictive agent. In this second case, the predictive agent sends an offer that is likely to be accepted by the counterpart, which is the case of a risk-averse behavior.

The proposed strategy combines the virtues of the two strategies with the introduction of a parameter noted risk portion "RP". It is a predictive strategy that allows agents pre-specify in percentage terms how much they are willing to prolong the negotiation process in order to achieve a more satisfying outcome compared to the outcome they would achieve in the non-learning case.

In each time step t agent α estimates the next offer of his counterpart, $\hat{X}_{b\to a}^{t+1} = \left(\hat{x}_{1(b\to a)}^{t+1}, \hat{x}_{2(b\to a)}^{t+1}, \ldots, \hat{x}_{n(b\to a)}^{t+1}\right)^T$. The proposed decision rule makes use of the default strategy (S^{α}) of the predictive agent to generate offers until the detection of a "meeting point" (MP) with the "opponent". MP is a point which would result an established agreement if the agent was guided solely by his default strategy. When such point is detected, and according to the agent's attitude towards risk, agent risks staying in the negotiation in order to maximize the utility of the final agreement. In this respect two extreme attitudes can be generated: risk-seeking and risk-averse. The risk-seeking agent is willing to spend all the remaining time until expiration of his deadline engaging in an adaptive behavior to turn the estimations of his counterpart's responses to profit. This risk-seeking behavior is based on the decision rule discussed in [8] and is extended to support multiple issues. More specifically:

Risk-seeking Behavior:

For each issue i

If issue value is increasing with time

$x_{i(a\to b)}^{t} = \hat{x}_{i(b\to a)}^{t+1} - \varepsilon$

Else

$x_{i(a\to b)}^{t} = \hat{x}_{i(b\to a)}^{t+1} + \varepsilon$

End For

Generate Offer $X_{a\to b}^{t} = \left(x_{1(a\to b)}^{t}, x_{2(a\to b)}^{t}, \ldots, x_{n(a\to b)}^{t}\right)^T$

where ε is a domain dependent parameter.

On the other hand risk-averse agents follow a more conservative behavior when they detect an MP. They use the decision rule discussed in ([6], [7]) and thereafter do not make any further concessions and insist on sending their previous offer, waiting for the opponent to establish an agreement.

Risk-Averse Behavior:

When MP is detected:

For each issue i

$x_{i(a\to b)}^{t} = \hat{x}_{i(b\to a)}^{t+1}$

End For

If t>MP:

For each issue i

$x_{i(a\to b)}^{t} = x_{i(a\to b)}^{t-2}$

End For

Generate Offer $X_{a\to b}^{t} = \left(x_{1(a\to b)}^{t}, x_{2(a\to b)}^{t}, \ldots, x_{n(a\to b)}^{t}\right)^T$

Fusions of the two extreme attitudes have led to the specification of risk portions (RPs) which characterize the predictive agent's behavior after the detection of MP. RP_{α} determines the percentage of the distance between MP and deadline $T_{\max}^{a}$ that agent α is willing to adopt the risk-seeking behavior. After RP_{α} is consumed agent adopts the risk-averse behavior. For a predictive agent who is not willing to take any risks RP_{α} is set to 0%, while for an agent who is willing to risk until expiration of his deadline RP_{α} is set to 100%. The decision making rule repeated in each step is thus formulated as follows:

If $U^{a}\left(\hat{X}_{b\to a}^{t+1}\right) > U^{a}\left(X_{a\to b}^{t}\right)_{default}$ (*detection of MP*)

  If **RP_{α}** is not consumed

    Generate Offer adopting Risk-Seeking Behavior

  Else

    Generate Offer adopting Risk-Averse Behavior

Else

  Generate Offer $\left(X_{a\to b}^{t}\right)_{default}$

where $\left(X_{a\to b}^{t}\right)_{default}$ is the offer generated by agent α at time t based on his default strategy, and U^{a} is his utility function that permits the comparison of different offer vectors.

In the following example we consider negotiations conducted between an electricity provider and a consumer agent over the service terms of an electricity trade. The negotiable object is characterized by four attributes representing the number of Kwh, the Price per Kwh (measured in euro cents), the Penalty term (percentage of the sum which will be returned to the consumer in case of dissatisfaction), and the duration of service provision (measured in hours). The deadline of the consumer is set to 150 rounds and that of the producer is set to 152 rounds. We assume that the two agents have opposing interests; the consumer will start from a low price and a low number of Kwh, which he will increase in each round, while he will start from a high percentage of returns and high service duration which he will decrease in each round. At the same time the producer will initially request high price per Kwh, and high number of Kwh which he will lower in each round, and low penalty and duration of service provisioning which he will increase in each round. In the example we assume that the two agents have set the same reservation (min and max) values as demonstrated in Table 1.

Table 1. The reservation values of the negotiable attributes

Attribute	Min Value	Max Value	Consumer Role	Provider Role
Duration (in hours)	10	30	Decreasing	Increasing
Kwh	20	200	Increasing	Decreasing
Price (euro cents)	10	100	Increasing	Decreasing
Penalty (% returns)	5	80	Decreasing	Increasing

Figure 2(a) illustrates the negotiation discourse of the two agents when they do not employ any learning techniques. The offered values of each negotiable attribute in each round are depicted with blue for the consumer and with red for the producer.

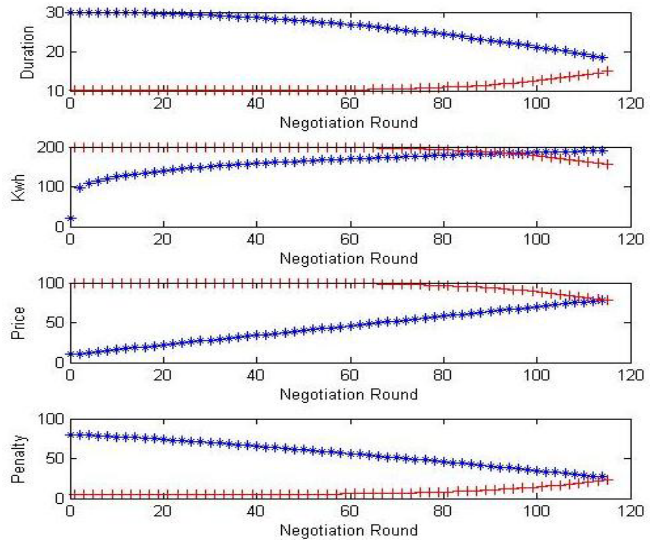

Fig. 2 (a). Negotiation between two non learning agents

In the non-learning case (Figure 2(a)), negotiation terminates at round 116, where the consumer agent decides to accept his counterpart's offer. The final offer vector (Duration, Kwh, Price, Penalty) is (14.95, 155.36, 77.68, 23.59) and the utility incurred to the consumer agent is 0.248.
Figure 2(b) illustrates a discourse where the consumer agent applies the proposed strategy with RP=0%.

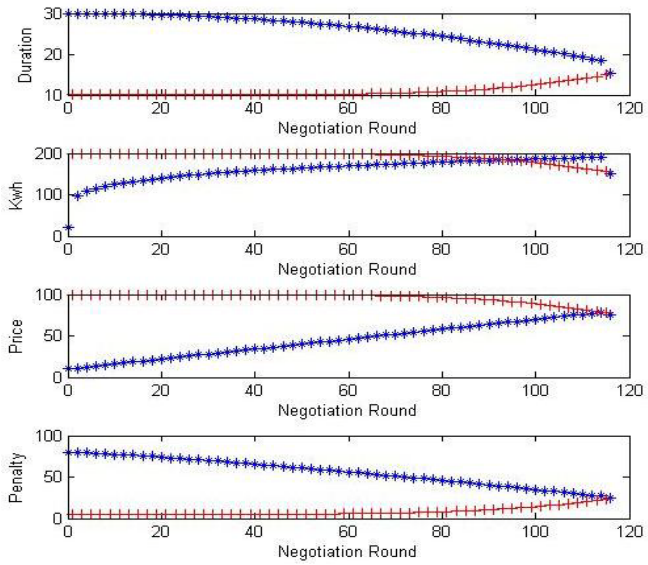

Fig. 2 (b). Negotiation with a consumer agent employing the proposed strategy with RP=0%

At round 116, the consumer agent detects the meeting point (MP) and initiates the predictive behavior. Since RP is set to 0% (it is consumed upon detection of MP), the agent generates the risk-aversive behavior and sends the predicted offer to his counterpart. Negotiation terminates at round 117, where the provider decides to accept the consumer's offer. The final offer vector (Duration, Kwh, Price, Penalty) is (15.40, 151.34, 75.67, 25.27) and the utility incurred to the consumer agent is 0.2703. Note that in case RP=0%, the predictive strategy incurs 2.23% absolute increase in utility, without much prolongation of the negotiation discourse. The maximum prolongation of the discourse when the risk aversive behavior is applied is by 1 round (the counterpart will accept the predictive agent's proposal in the next round).

In continuance, in Figure 2(c) we illustrate a discourse with a consumer agent who uses the proposed strategy with RP=50%. At round 116, the consumer detects the MP and adopts the risk-seeking behavior, until round 134 where the RP is consumed. The offer sent at round 134 is based on the risk-averse behavior and the offer vector (Duration, Kwh, Price, Penalty) = (20.66, 101.39, 51.30, 45.42) is accepted by the provider at round 135. The utility incurred to the consumer in this case is 0.5403, thus the absolute increase compared to the non-learning case is 29.23%.

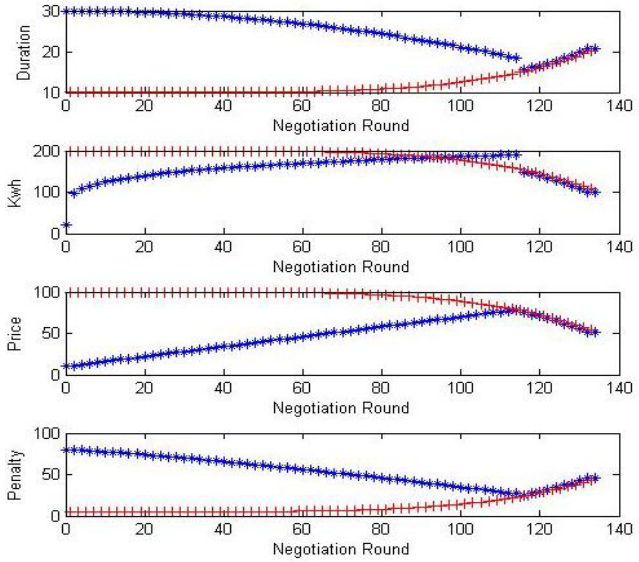

Fig. 2 (c). Negotiation with a consumer agent employing the proposed strategy with RP=50%

4 Results and Discussion

Since the objective is to increase the utility that incurs to the predictive agent, focus is set on studying the change of the agent's utility, as well as the change of the number of agreements with variable RP values. For this reason a number of experiments are conducted assuming negotiations between a predictive agent who makes very accurate estimations and a non-learning counterpart employing many different types of time-dependent behaviors. In the experiments conducted, the strategy of the counterpart is known to the predictive agent, who simply applies the expected values of the counteroffers to the proposed decision rule.

The strategy is tested with consumer agents assumed to have perfect predicting skills and providers following TD strategies. The experimental workbench issues nine different scenarios with respect to deadline, and overlap of agreement zones of the two negotiators ($\{T_{\max}^{Con} = T_{\max}^{\Pr}, T_{\max}^{Con} < T_{\max}^{\Pr}, T_{\max}^{Con} > T_{\max}^{\Pr}\} \times \{\Phi=0, \Phi=0.33, \Phi=0.66\}$), where $T_{\max}^{a} \in [50:100:350]$, $a=\{Con,Pr\}$, and Φ is a parameter which indicates overlap of the agreement zones [4]. In each scenario a variety of concession curves, defined by parameter $\beta = \{0.2, 0.5, 0.8, 1, 3, 5, 7\}$, is considered in order to build the default strategies of the opposing agents. For each of the 2,352 generated negotiation environments different RPs are set to the predictive agent (consumer) ($RP_{Con} \in [0:5:100]$), leading to an overall of 49,392 experiments. The objective is to measure the gain of consumer agent with respect to the RP parameter, and highlight the value of forecasting counterpart's next offer in multi-issued negotiations. From a total of 2,352 negotiation environments, average utilities of negotiations conducted between non-learning agents (AvgUtil_NL) and between predictive and non-learning agents employing different RP values ($AvgUtil\_L_{(RP)}$) are computed. Since utilities are specified in range [0,1], the average absolute increase in utility incurred to the agent who employs the proposed strategy is computed as follows:

$$AvgAbsInc_{(RP)} = (AvgUtil\_L_{(RP)}) - AvgUtil\_NL) * 100\%$$

Figure 3 depicts the average absolute increase for each RP value.

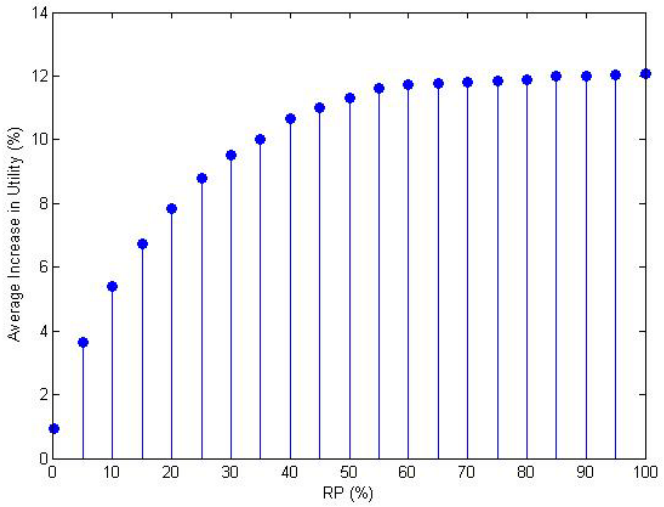

Fig. 3. (%) Average gain in Utility with respect to RP

As it is shown in Figure 3 an agent with RP 0% incurs an average increase of 0.94% in utility, while an agent with RP 100% incurs an average increase of 12.05% in utility. Additionally, the concavity of the curve shows that the marginal increase of the utility decreases as RP grows. For different values of RP>60%, the increase in utility does not change significantly.

However employment of the predictive strategy increases the time (number of negotiation rounds) of the discourse, as the learning agent continues the negotiation after detection of the meeting point (MP). The time consumed, TCons, in each negotiation round is normalized thus: $TCons = \frac{NegotiationRounds}{\min(T_{\max}^{Con}, T_{\max}^{\Pr})}$. For the total of 2,352 negotiations between non-learning agents the average consumed time (AvgCons_NL) is computed. For each RP value the average consumed time $AvgCons\_L_{(RP)}$ is also computed. Finally the average absolute increase of the time consumed by the agent who employs the proposed strategy is derived as follows:

$$AvgAbsCons_{(RP)} = (AvgCons\_L_{(RP)} - AvgCons\_NL) * 100\%$$

Figure 4 illustrates the increase of negotiation time with respect to the risk portion.

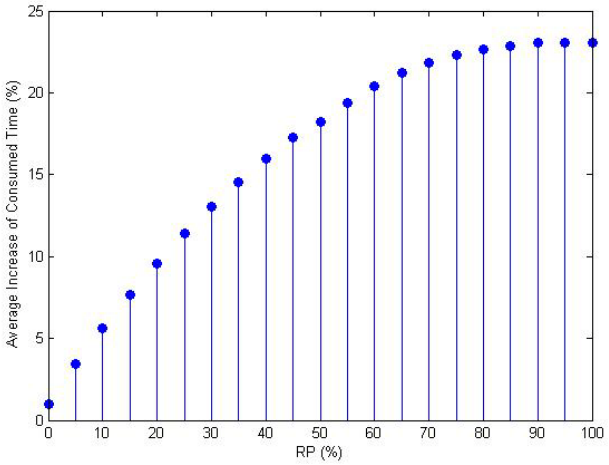

Fig. 4. (%) Average Increase of Negotiation Time with respect to RP

As it is shown in Figure 4 the average increase of negotiation time is 0.96% when RP is set to 0% and 23.07% when RP is set to 100%.
Prolonging the negotiation procedure is the main cause of decrease of negotiating agreements, as it likely that the counterpart reaches his deadline and terminates the process.
From a total of 2,352 negotiations conducted between non-learning agents the number of agreements (Num_NL) is computed. The number of agreements of negotiations conducted between predictive and non-learning agents for each RP value (Num_$L_{(RP)}$) is also computed. The average relative decrease of the number of established agreements by the agent who employs the proposed strategy is derived as follows:

$$AvgDecAgreement_{(RP)} = \frac{Num\_NL - Num\_L_{(RP)}}{Num\_NL} * 100\%$$

Figure 5 illustrates how the number of agreements is affected by the adoption of the predictive strategy.

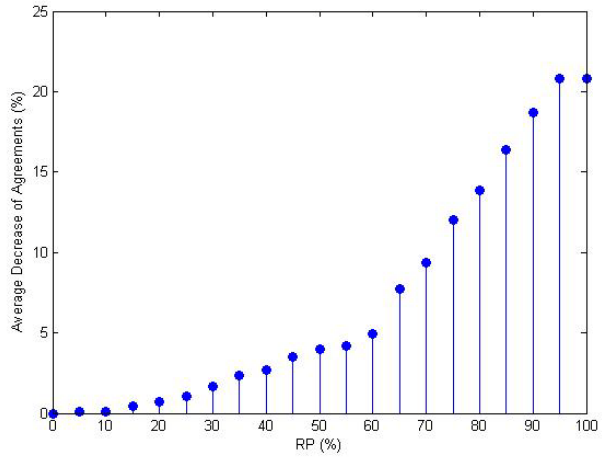

Fig. 5. (%) Average Decrease of the number of agreements with respect to RP

More specifically, when RP is set to 0% the number of agreements is not decreased, while when RP is set to 100%, the number of agreements shows an average decrease of 20.78%.
It is empirically proved that as RP value increases the agent may increase his individual utility. For each negotiation environment experiments were conducted for the 21 predefined RP values. However, some of the selected RPs did not necessarily satisfy the rationality condition, and this led to negotiation breakdowns. As it is observed, failing to establish an agreement increases with increasing RP value, as the counterpart reaches his deadline before the predictive agent adopts the risk-averse behavior. Setting the appropriate RP values is crucial for a predictive agent who wishes to attain the maximum possible utility gain and avoid negotiation failures.
Appropriate selections can be made if estimations of the opponent's deadline are also available. In this case, RP can be set just before the expiration of the opponent's deadline. The predictive agent can make

use of the risk-seeking behavior until that point and then employ the risk-averse behavior, which will result in acceptance from the counterpart, before the expiration of the latter's deadline.

In order to set appropriate RPs, if an estimation of the counterpart's deadline $\hat{T}^b_{\max}$ is available, predictive agent a should distinguish the two cases: If $\hat{T}^b_{\max}$ falls in agent's a turn, then he should adopt the risk-averse behavior at round $\hat{T}^b_{\max}$ - 2. On the contrary, if it falls in agent's b turn, then agent should adopt the risk aversive behavior at round $\hat{T}^b_{\max}$ - 1. In the first case RP can be set upon detection of MP as follows:

$$RP = \frac{(\hat{T}^b_{\max} - 2) - MP}{T^a_{\max} - MP} * 100\% \quad (1)$$

While in the second case RP can be set upon detection of MP as follows:

$$RP = \frac{(\hat{T}^b_{\max} - 1) - MP}{T^a_{\max} - MP} * 100\% \quad (2)$$

Equations (1) and (2) apply if opponent's deadline is shorter than the player's deadline. Contrarily, the predictive agent can set RP to 100%. In this respect the predictive agent can attain the maximal possible increase of his individual utility by adopting the risk seeking behavior after detection of MP, and the risk-averse behavior only in the final round. Estimating the counterpart's deadline is beyond the scope of this thesis. However, to assess the extension with the appropriate RP values, we have assumed knowledge of the opponent's deadline by the predictive agent. The proposed strategy with the appropriate RP values when applied to the same experimental settings, yields very satisfying results as the average absolute increase in utility is 12.017%, the average absolute increase of the consumed time is 22.53%, and the average relative decrease of the number of agreements is reduced to 0.61%.

5 Conclusion

Predictive decision making is characteristic to current state of the art socio-technical systems that guide negotiation processes under electronic settings. From semi-automated negotiation support systems, to fully automated platforms where all processes are undertaken by software agents, the back end participants are particularly benefitted by the use of models of computational intelligence. Such models provide estimations about the behavior of the negotiator's counterparts and allow users or agents acting on their behalf, to adapt their strategy and evaluate risks and dynamics of current negotiation.

This thesis contributes to the field of negotiation with the proposal of a predictive strategy that incorporates different attitudes towards risk, as well as to the field of application of neural networks in negotiations with the introduction of session-long learning agents.

We present the foundations related to the negotiation domain, terminologies and classifications, research methodologies and software platforms, as well as description of negotiation protocols and strategies that constitute state of the art negotiating agents. We categorize "smart" strategies to explorative, repetitive and predictive, and particularly focus on the issue of forecasting the counterpart's next offer in order to increase the individual gain. As different attitudes towards risk may emerge, we present a predictive strategy that incorporates a risk-related parameter and allows the predictive agent to adapt his consequent offers with respect to the estimations of his counterpart's responses. We demonstrate the increase of the individual gain that incurs to the agent and highlight the need to appropriately set the risk related parameter in order to avoid negotiation breakdowns.

The skill of forecasting the counterpart's future offers is further investigated and selection of Multi-layer Perceptrons (MLPs) is preferred to other learning models. A brief overview and comparison of the forecasting tools employed by negotiators is provided, and the selection of MLPs is justified. Current systems which base their learning models on data acquired from previous interactions or from synthetic data provide satisfying results in static negotiation environments (where data distributions do not change). However, when data distributions change, the systems no longer provide accurate estimations. A new perspective to the issue is introduced, by highlighting the need to learn during the negotiation session. Such an approach is viable in open, dynamic negotiation environments. A number of experiments are conducted to support this argument and disclose the inability of initially pre-trained networks to capture the dynamics of changing distributions. "Session-long learning" agents, static (SSLAs) or adaptive (ASLAs), trained with the data of the current negotiation thread, prove capable of capturing the negotiation dynamics. From the experiments conducted it is empirically proved that ASLAs provide the most promising results (forecasts yield the smallest error), however the combination of neural networks with genetic algorithms require a lot of time and resources which is sometimes restricting in negotiation domains. This result makes SSLAs a good selection for the problem of forecasting the counterpart's future offers.

This thesis also outlines a number of issues that can be considered for future research. The first relates to the domain of application of the predictive strategy. The decision rule of the risk-seeking agent, as well as the decision rule discussed in [8], cannot safely be used in the case where counterparts adopt pure Behavior Dependent strategies, as the latter would imitate the 'smart' behavior of the predictive agents and would push back from the agreement as well. For this reason the predictive agents presented in this thesis were only tested with non-learning counterparts who followed Time Dependent strategies. However, in hybrid strategies, where linear combinations of time and behavior dependent tactics are considered, success of the proposed strategy is expected to depend on the weight of the time dependent tactic. In order to broaden the applicability of the proposed strategy, estimation of behavior dependency could also be enhanced to the decision rule. Another issue left for future research is the investigation of predictive strategies in co-operative environments, where the objective is to maximize the joint rather than the individual utility. Moving to the realm of the employed forecasting tool, ASLAs have proved more efficient than SSLAs. However the trade-off is the increased computational resources and time of convergence. An issue left for future research is therefore to test other adaptive and more efficient structures with the predictive agents [19] [20]. Examples of such structures are the Evolving Fuzzy Neural Networks (EFuNNs) and DENFIS, which are Evolving Connectionist Systems (ECoS) that continuously, evolve their structure and functionality to capture the dynamics of turbulent settings.

References

1. Masvoula, M., Kanellis, P., Martakos, D. (2010). *A Review of Learning Methods Enhanced in Strategies of Negotiating Agents*. Proceedings of 12th International Conference on Enterprise Information Systems (June 8-12,2010,Madeira). ICEIS pp. 212-219.
2. Carbonneau, R, Kersten, GE, and Vahidov, R (2008). Predicting opponent's moves in electronic negotiations using neural networks. *Expert Systems with Applications: An International Journal,* 34 (2): 1266-1273.
3. Lee, CC, and Ou-Yang, C (2009). A neural networks approach for forecasting the supplier's bid prices in supplier selection negotiation process. *Expert Systems with Applications,* 36(2): 2961-2970.
4. Faratin, P, Sierra, C, and Jennings, NR (1998). Negotiation Decision Functions for Autonomous Agents. *Int. Journal of Robotics and Autonomous Systems*. 24 (3 - 4): 159-182.
5. Masvoula, M., Kontolemakis, G., Kanellis, P., and Martakos, D. (2005). *Design and Development of an Anthropocentric Negotiation Model*. Proceedings of the Seventh IEEE International Conference on E-

Commerce Technology (July 19-22, 2005, Munich). CEC. IEEE Computer Society, Washington, DC, pp. 383-386.
6. Papaioannou, IV, Roussaki, IG, and Anagnostou, ME (2006a). *Comparing the Performance of MLP and RBF Neural Networks Employed by Negotiating Intelligent Agents*. Proceedings of the IEEE/WIC/ACM Int. Conf. on intelligent Agent Technology, pp. 602-612
7. Papaioannou, IV, Roussaki, IG, and Anagnostou, ME (2006b). *Towards Successful Automated Negotiations based on Neural Networks*. Proceedings of the 5th IEEE/ACIS Int. Conf. on Computer and Information Science, pp. 464-472
8. Oprea, M (2003). *The Use of Adaptive Negotiation by a Shopping Agent in Agent-Mediated Electronic Commerce*. Proceedings of the 3rd Int. Central and Eastern European Conf. on Multi-Agent Systems. Heidelberg, Berlin: Springer Verlag , pp. 594-605.
9. Masvoula, M., Halatsis, C., Martakos, D. (2012). Predictive Automated Negotiators with different Attitudes Towards Risk. *Engineering Intelligent Systems Journal*, 20(1):89-105
10. Masvoula, M., Halatsis, C., Martakos, D. (2011). *Predictive Automated Negotiators Employing Risk-Seeking and Risk-Averse Strategies*. Proceedings of 12th EANN/7th AIAI Joint Conferences (September 15-18, 2011, Corfu, Greece). EANN/AIAI (1) 2011: 325-334.
11. Hou, C (2004). *Predicting Agents Tactics in Automated Negotiation*. Proceedings of the IEEE/WIC/ACM Int. Conf. on Intelligent Agent Technology, pp. 127-133.
12. Brzostowski, J, and Kowalczyk, R (2006a). *Adaptive Negotiation with On-Line Prediction of Opponent Behaviour in Agent-Based Negotiations*. Proceedings of the IEEE/WIC/ACM int. Conf. on intelligent Agent Technology, pp. 263-269.
13. Brzostowski, J, and Kowalczyk, R (2005). *Modelling Partner's Behaviour in Agent Negotiation*. AI 2005. Lecture Notes in Computer Science 3809. Berlin, Heidelberg: Springer-Verlag, pp. 653-663.
14. Brzostowski, J, and Kowalczyk, R (2006b). *Predicting Partner's Behaviour in Agent Negotiation*. Proceedings of the 5th int. joint conf. AAMAS. New York USA: ACM, pp. 355-361.
15. Papaioannou, I, Roussaki, I, and Anagnostou, M (2008). *Detecting Unsuccessful Automated Negotiation Threads When Opponents Employ Hybrid Strategies*. Proceedings of the 4th int. Conf. on intelligent Computing. Heidelberg, Berlin: Springer-Verlag, pp. 27-39.
16. Roussaki, I, Papaioannou, I, and Anagnostou, M (2007). *Building Automated Negotiation Strategies Enhanced by MLP and GR Neural Networks for Opponent Agent Behaviour Prognosis*. (IWANN) 2007. Heidelberg, Berlin: Springer-Verlag, pp. 152-161
17. Masvoula, M., Halatsis, C., Martakos, D. (2012). Forecasting negotiation counterpart's offers with the use of neural networks: A focus on session long learning agents. *Neurocomputing and Applications*. (Submitted for Publication)
18. Masvoula, M., Halatsis, C., Martakos, D. (2011). *Integrating Negotiating Agents with Evolving Connectionist Structures to Facilitate the Prediction of Counterpart's Responses*. Proceedings of the International Conference on Neural Networks (ICNN 2011) (July 27-29, 2011, Paris, France).
19. Masvoula, M., Kanellis, P., Martakos, D. (2010). *Evolving Structures for Predictive Decision Making in Negotiations*. Proceedings of 12th International Conference on Enterprise Information Systems (June 8-12,2010,Madeira). ICEIS pp. 391-394.
20. Masvoula, M., Kanellis, P., and Martakos, D. (2009). *Integrating agents with connectionist systems to extract negotiation routines*. Proceedings of the 11th International Conference on Enterprise Information Systems (6 – 10, May 2009 Milan, Italy). ICEIS (2): 251-256.

Theoretical and Experimental Investigation of Quantum Dot Passively Mode Locked Lasers for Telecomm and Biomedical Applications

Charis Mesaritakis*
National and Kapodistrian University of Athens,
Department of Informatics and Telecommunications,
Panepistimiopolis, Ilisia, 15784, Athens, Greece
cmesar@di.uoa.gr

Abstract. This thesis is focused on the experimental and theoretical study of novel quantum dot passively mode locked lasers. The motivation behind this endeavour was the fundamental need, of many scientific areas, for high power, ultra-short and time stabilized optical pulses, generated directly from highly integrated laser devices. Such devices could allow the study of ultra-fast chemical and physical effects, or alternatively could provide telecomm engineers and medical practitioners with hand-held lasers suitable for a variety of applications. The first step consists of a detailed mathematical investigation of the mode-locking mechanism so as to depict clearly the fundamental limitations associated with ultra-fast high power generation. Through numerical modelling a novel semiconductor material (quantum-dots) is successfully introduced, whereas the benefits of such an approach are presented. The major part of this thesis consists of a detailed experimental investigation of various quantum-dot based mode-locked semiconductor structures. Through this investigation various new effects were demonstrated, that allowed a deeper insight to new regimes of operation. These regimes that namely consist of multi-wavelength emission and mode locking, pulse-width narrowing with increasing current injection, Q-switching stabilization and feedback-free chaotic operation, unlock new areas of applications, present new design guidelines for highly integrated semiconductor lasers. Finally, these regimes originate from the unique properties of the quantum-dot material and were firstly observed and investigated during this thesis. Consequently, this thesis played a major role in the future development of this scientific field and led to more than 15 publications in peer-reviewed journals 4 conference papers and a book chapter.

1 Introduction

The introduction of laser systems during the 1960s allowed the generation of highly spectrally and spatially coherent electromagnetic radiation with immense impact into traditional scientific fields, whereas it allowed a radical expansion of newly emerging

* Dissertation advisor: Dimitris Syvridis, Professor

fields like bio-imaging/surgery, ultra-fast physics and high bitrate optical communications. Optical pulse generation was a highly sought achievement that could further expand the application areas of lasers. Traditional laser structures hinder such attempts due to their very large size and the complicated biasing schemes. On the other hand the semiconductor counterparts offered excellent integrating capabilities but their operating performance were significant inferior. In order to address this issue many novel ideas have been put into practice elevating some of the fundamental limitations of semiconductor pulsed lasers. During the last years novel material like quantum dots (QD) have been introduced offering further improvement compared to conventional semiconductor approaches. In detail QD materials offer three-dimensional confinement of carriers that in turn allow ultra-fast recovery time, low-linewidth enhancement factor, low thermal sensitivity and high saturation gain [1].

In this thesis the major effort was focused on exploiting the inherent advantages offered by the QD materials, aiming through device design of new laser structures and through identification of new regimes of operation, to further expand the capabilities of QD-based mode locked lasers. In detail a detailed numerical model that takes into consideration the unique band-structure of the QD material was developed, whereas by incorporating time-delayed differential rate equations the full behavior of a passive mode locking QD laser was modeled. Through modeling all the basic parameters that affect pulse quality and power were identified, thus the experimental and design efforts were focused on these aspects.

The experimental investigation allowed new regimes of operation that consisted of: Simultaneous dual wavelength mode locking from the two basic QD energy bands (ground-excited state) allowing the creating of two independent train of pulses from different wavebands [2]. Using the same concept dual wavelength mode-locking has been achieved from the ground state allowing tunable repetition rate of pulses [3]. Moreover, the ground state splitting effect allowed the identification of a new regime of operation that provided shorter pulses with increasing injection current in the presence of dual ground state emission [4-5]. This effect is very significant for the design of QD-lasers due to the fact that it allows simultaneous increase of the average power and reduction of the pulse width, a trend that cannot be achieved by conventional devices. The low linewidth enhancement factor of QD materials provided stimulation for a detailed experimental identification of the feedback tolerance of QD mode-locked lasers from both wavebands [6-8]. This study provided useful information about the isolation requirements of each energy-state, whereas it confirmed the strong dependence of pulse stability to the precise length of the external cavity. The significant tolerance of QD-lasers to optical-instabilities induced by feedback setup provided another urge to study the stability of the lasers under asymmetric pumping conditions. In this case optical instabilities and chaotic emission was achieved only by asymmetric current injection in the absence of complicated feedback setups [9]. Although these findings revealed an unstable operating regime it can allow the expansion of QD-lasers to new areas of applications like chaotic cryptography. Moreover, through controlled feedback, Q-switching elimination and wavelength switching was achieved. This process can remove the amplitude modulation that imposes serious restrictions to certain medical modalities [10]. Finally, by designing and evaluating QD-lasers with different structural parameters

like, cavity length and width, number of QD layers, manufacturing guidelines was extracted and the parameters that allow multi-wavelength emission were investigated [11] whereas novel tapered structures were designed and constructed that allowed record level peak power [12].

2 Experimental Setup

In order to perform the above-mentioned experiments different experimental setups were employed. Nonetheless, in all cases some basic components were used. The different laser modules were mounted on a custom made 3-axis manipulator equipped with a vacuum pump and a thermoelectric cooler in a closed control loop, that provided thermal and mechanical stabilization. Optical power from the device was collected through a tapered single mode optical fiber with high numerical aperture also positioned through an independent 3-axis manipulator. Basic measurements have been performed with the help of a RF-spectrum analyzer with electrical bandwidth of 30GHz. The optical spectrum was evaluated through an optical spectrum analyzer with minimum resolution of 0.1nm, optical power measurements were performed through a large area InGaAs photodiode, while the temporal properties of the pulses were evaluated with the help of a background free optical autocorrelator based on second harmonic generation with a minimum resolution of 10fs. In fig. 1 the most elaborate experimental setup is presented that was used in order to study the feedback tolerance of each waveband. It consists of the laser module, a polarization controller aiming to optimize laser-external cavity coupling, a fiber coupler (90/10), where the 10% branch fed an optical isolator and all the measuring components (RF-Optical analyzer and autocorrelator), while the 90% of the coupler was fed to the external cavity. This cavity consisted of a free space section with various neutral density filters in order to tune the amplitude of the back-reflected optical power and a variable delay line, whose aim was to fine tune the length of the cavity. All the other characterization experiments were performed with the same experimental setup but with the absence of the external cavity module.

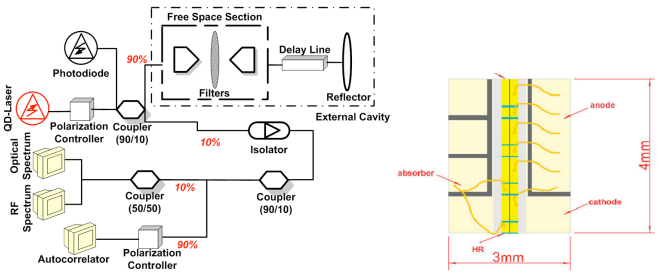

Fig.1. Fundamental experimental setup used to investigate feedback tolerance of each waveband of QD-mode locked lasers. Basic schematic of the laser structure.

The devices tested in this thesis consist of Fabry Perot laser structures that were grown by molecular beam epitaxy on a GaAs (100) substrate and contain self-

assembled InAs/InGaAs QD Layers into 440 nm GaAs waveguide surrounded with $Al_{0.35}Ga_{0.65}As$ claddings. The back and front facets of the devices were high reflective coated (>99%)/low reflective coated (~10%), respectively, for the operating wavelength (1260 nm). Mode locking was achieved though reverse biasing part of the laser. The gain/absorption ratio is set to 85/15 for all devices under test. The length of the devices used in this paper varies from 2 mm to 4 mm, and the number of QD layers was 5, 10, and 15.

3 Results and discussion

A. Pulse Width Narrowing In the Presence of dual Ground-State Emission.

In this experiment, experimental results related to the existence of dual waveband emission from the GS from a multi-section QD mode locked laser are presented. In particular, we demonstrate that under specific bias conditions, GS splitting (GSS) occurs and in turn enables the narrowing of the generated pulses with increasing injection current. This observation leads to pulses of simultaneous increased peak and average power with respect to the usual device operation, which is a desired feature for biomedical applications.
The device under test is a multi-sectional QD laser, where the absorber is split into to two sections (0.3 mm each of them) and reverse biased, while the rest are forward biased. The laser structure contains five self-assembled InAs/InGaAs QD layers. The laser chip has a total length of 4 mm and contains stripes with 6μm width. In fig. 2 corresponding optical spectra for different bias conditions are presented that demonstrate the GS splitting effect.

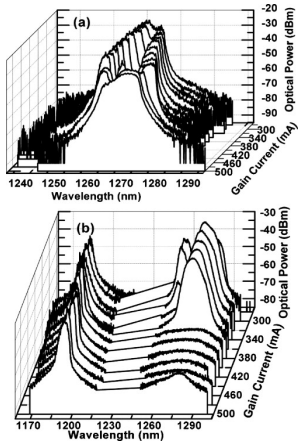

Fig.2. (a) Optical spectra in the absence of reverse voltage at the absorbing section and for various gain currents. (b) Optical spectra for V_{abs}=-7V.

This effect can be directly attributed to mode competition effects through the homogenous and inhomogeneous broadening. In particular, the splitting effect occurs in the GS when the injection current is high enough to enhance spectral hole burning

(SHB) at the center of the lasing spectrum. Carriers from these modes are reallocated through the inhomogeneous broadening to the outer modes, which initiate stimulated emission. The increase in the spectral separation of the two peaks with increasing current can be attributed to the fact that with increased modal gain a larger number of central modes are subject to SHB.
Pulse width evolution with the injection current is shown in Fig. 3. Although the pulse width is expected to increase monotonically with the current, in our case the opposite behavior is observed when double emission from the GS occurs.

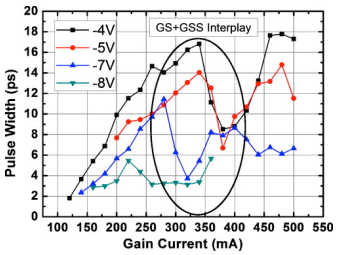

Fig.3. (a) Pulse width versus current injection for different reverse voltage at the absorbing section.

This temporal behavior of the emitted pulses was attributed to the fact that as injection current is increased thus more modes participate in the mode-locking mechanism. Although this effect theoretically can enhance mode locking in practice the larger number of modes tend to lock at different relative phases and consequently competition effects can destabilize mode locking. In our case the spectral splitting effect reduces the number of modes and thus reduces the competition effects. In particular this pulse-width reduction can lead to a significant increase in peak power.

B. Dual Ground-State Mode-Locking

Dual-wavelength mode-locked laser sources have been traditionally developed and investigated with other solid-state materials, most notably in Ti:Sapphire, fiber and most recently, ceramic lasers Research in this area has been motivated by the variety of applications for dual- and multiple-wavelength ultra-short pulses, such as time-domain spectroscopy, nonlinear optical frequency conversion and wavelength division multiplexing. In this context, the compactness, lower cost and direct electrical pumping associated with semiconductor lasers are very attractive features for reducing the footprint and complexity of the aforementioned applications, with the potential to also open up new avenues in ultrafast optical processing and optical interconnects.
By using the same device as described above but employing different bias conditions we tried to achieve simultaneous mode locking from the two independent sub-bands of the GS under splitting conditions. This regime of operation appeared when the reverse voltage at the absorbing section was kept below 3.5 V and the gain current was swept from the laser threshold up to 300 mA. The optical spectrum exhibited again two discrete GS sub-bands, while the RF spectrum had two discrete peaks centered around 10 GHz with separation in the order of 30 MHz. The two RF peaks

correspond to two independent groups of modes that were phase locked. The slightly different repetition rate ($\Delta f \approx 30$ MHz) occurs due to the refractive index variation at each sub-band, which results to a slightly different effective cavity length. The absence of any side peaks in the RF spectrum apart from the two main peaks observed, imply that there is no amplitude modulation in the generated pulses (Q-switching) (fig.4).

Such monolithic lasers emitting mode locked pulses of two different GS sub-bands could be exploited for the realization of a low-cost terahertz pulsed source through beating in a photoconductor. Taking into account the spectral separation of the two sub-bands (2-14 nm), terahertz pulsed radiation in the (350 GHz-2.6 THz) band could be generated. Due to the different repetition rates of the optical pulses, the resulting terahertz radiation will be modulated by the difference $\Delta F_{rep} = F_{GS2}-F_{GS1}$, where F_{GS} are the repetition rates of the fundamental pulses.

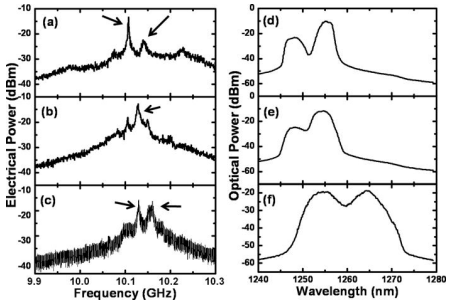

Fig.4. (a)-(c)RF spectra for increasing gain current and reverse voltage equal to -3.5V. (b) Corresponding optical spectra.

C. Dual Ground-Excited State Mode-Locking

In this case the structure used in the laser device was grown by molecular beam epitaxy on a GaAs substrate. The active region incorporated 5 layers of InAs QDs. A two-section QD laser diode was fabricated with a ridge waveguide 6μm wide, a total length of 2mm, while the saturable absorber was 300μm long, and was located near the back facet. The front and back facets were anti-reflection (~3%) and high-reflection (~95%) coated, respectively.

The dual-wavelength mode locking regime was obtained for current levels in the gain section between 330 and 430mA, and values of reverse bias between 6 and 10V in the saturable absorber section. A map depicting the different mode-locking regimes is represented in Fig. 5. It is noteworthy to point out that mode-locking involving solely the ES transition is observed both before and after the dual-wavelength mode-locking regime. Although these results present again simultaneous from transitions from different energy bands available in the QD materials it is worth mentioning that in this case mode-locking is not achieved though two discrete ground-sub-bands but from a different available transition (ES) thus the spectral separation of the two trains of pulses exceeds 80nm.

The exploitation of this mode-locked regime could enable a range of applications extending from time-domain spectroscopy, through to optical interconnects, wavelength-division multiplexing and ultrafast optical processing.

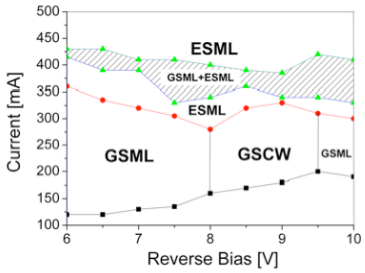

Fig.5. Mapping of the different operational regimes of a QD-passively mode locked laser versus current and reverse bias. The striped region corresponds to dual GS/ES mode locking.

D. Feedback Tolerance of QD Mode Locked Lasers and Chaotic Emission

The unique band-structure and carrier confinement of QD materials have triggered many attempts to study the line-width enhancement factor of such monolithic devices alongside their tolerance to feedback induce optical instabilities. During this thesis the first evaluation of pulse stability from both bands in the presence of optical feedback has been achieved. The experimental setup employed is described previously (fig.1). Two cases have been considered the first consists of single GS mode-locking while feedback effects were evaluated through calculation of the full-width at half maximum (FWHM) of the RF peak corresponding to the fundamental repetition rate. In fig. 6 the FWHM of the RF peak versus feedback strength is presented. It can be concluded that for feedback strength exceeding -34dB a radical increase in the FWHM is observed that corresponds to mode-locking deterioration, while for feedback larger than -25dB a complete coherence collapse is observed. On the other hand, if only ES mode-locking is present (higher gain current and reverse voltage) the system is more resilient to feedback instabilities and the increase of feedback strength even enhances the quality of mode locking by reducing the FWHM of the RF peak and consequently allow a more efficient phase locking of the longitudinal modes. Moreover as it has been shown numerically in the past for quantum well devices the exact external cavity length can play a dominant role in mode-locking stability and even further enhance mode locking if the external cavity round-trip is an integral integer of the intra-cavity round-trip.

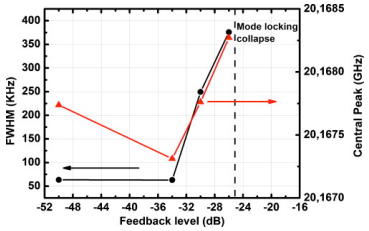

Fig.6. FWHM of the RF peak and the repetition rate for single GS mode locking.

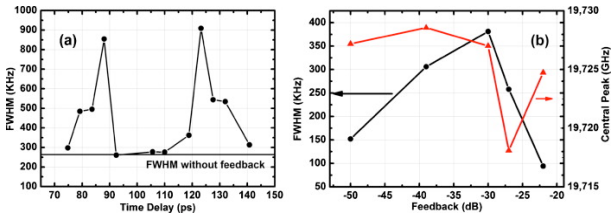

Fig.7. (a) FWHM for single ES mode locking versus cavity delay time. (b) FWHM and repetition rate versus feedback strength.

The main goal of the aforementioned approaches was to stabilize mode locking and investigate the isolation requirements for each waveband. On the other hand there are newly emerging applications that encourage coherence collapse and high spectral incoherence (multi GHz electrical bandwidth). Such applications like chaotic cryptography usually use controlled feedback related schemes that are quite complicated especially when high level of integration is required. In our case an alternative approach has been utilized in order to destabilize the laser. A two section monolithic laser was used where a different current density was employed at each section. The investigation revealed two regimes of operation. The first consisted of stable self-pulsations with low repetition rate. In this case the small section acted as an incomplete slow saturable absorber. The second regime manifest if the current was increased at both sections. In this case a significant increase in low frequency noise was observed with a 3dB bandwidth exhibiting 5GHz, whereas mathematical analysis of the collected experimental time-series revealed high complexity time-series with a correlation dimension exceeding 7. This operation regime originates from the different gain profiles of the two sections and the dynamic feedback induced due to the intra-cavity variations. Furthermore, this regime of operation allows the expansion of QD lasers into chaotic applications without the use of complicated feedback setups.

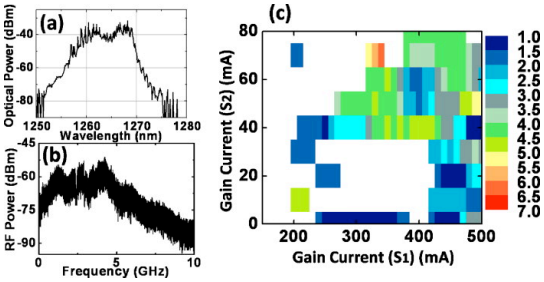

Fig.8. (a) Optical spectrum under coherence collapse (b) RF spectrum for the same regime (d) correlation dimension versus gain current of both sections exhibiting high complexity regimes.

E. Q-Switching Elimination and Wavelength Switching Through Controlled Feedback
In this part of the experimental investigation optical feedback was used in order to trigger wavelength switching. In particular through variation of parameters like feedback strength and external cavity length, wavelength switching under stable mode locked operation was achieved. One of the key benefits of this technique is the variation of the range that each regime of operation manifests (single GS/ES or dual wavelength emission), which in general is directly related only to the fabrication process. Furthermore through application of controlled level of optical feedback we were able to demonstrate stabilization of the ES mode locking and strong reduction of the pulse amplitude modulation effects (Q-switching), whereas significant reduction in device's current threshold was also recorded. In particular in fig. the RF spectrum for various level of optical feedback is presented. It can be clearly seen that by increasing the feedback strength the strong side-bands that imply amplitude modulation of the emitted pulses are strongly suppressed. The physical mechanism behind this effect is based on the fact that excess photons form the external cavity enhance stimulated emission and drive the laser beyond threshold, where the Q-switching effects tend to be dominant (fig.9).

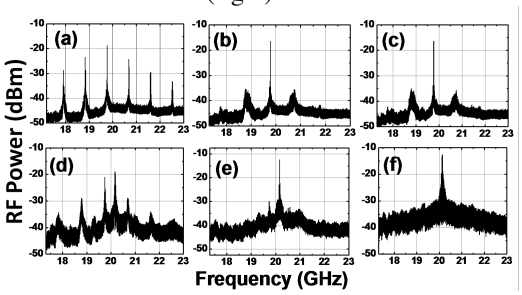

Fig. 9 RF spectra for the ES mode-locking of a 4mm QD device for various levels of controlled optical feedback.

F. Manufacturing Guidelines and Record Level Power Through Tapered Section QD-Mode Locked Lasers.

In order to exploit the advantages of QD-based pulsed sources, many efforts have focused on cavity optimization. These studies have focused on the effect of different fabrication related parameters like the number of QD layers and cavity length but were related to continuous wave (CW) operation. On the other hand, in the case of pulsed operation, most studies focused on the impact of the absorber/gain ratio to the mode locking performance. During this thesis we focus on the experimental identification of the impact of the number of QD layers on the quality of mode locking and the impact of excited state (ES) emission on the ground state (GS) mode locking was investigated for the first time.
In fig. 10 by comparing devices in terms of the number of QD layers independently of their length, a common behavior for all available lengths can be observed. The increase of number of layers from 5 to 10 induces an enhancement of the mode-locking performance in terms of pulse width, alongside an increase in the available range of bias conditions where mode locking is achievable. In particular, the dark grey regimes in the contour plots, which correspond to pulse widths smaller than 5ps,

are non-existent or difficult to achieve in the 5 QD layers (–7 V to –8 V and 100 mA to 200 mA), while the 10 layers devices exhibit a large range of bias conditions where such narrow pulses can be observed (–4 V to –8 V and 200 mA to 450 mA). It is of paramount importance to note that in the 5 QD layers case, the limited bias range that allows stable mode locking coincides with the early onset of ES lasing (high gain current). The increase of the QD Layers number from 10 to 15 does not result in a significant improvement of the pulse width, whereas in many cases, a degradation of system's performance can be observed (4 mm long devices).

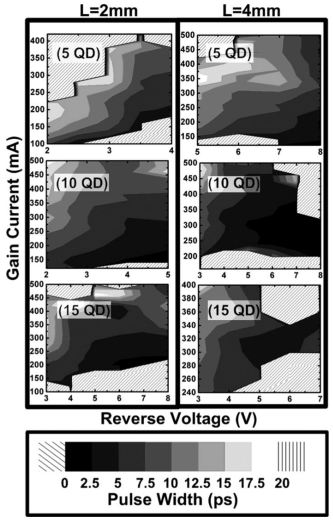

Fig.10. Contour plot of the pulse width versus gain current and reverse voltage for different laser structures.

In order to explain the experimental observed behavior, a standard mathematical formalism was employed that has been used in the past to evaluate the mode-locking performance of different QD laser structures. This approach treats mode locking from each state (GS-ES) as a simple net-gain modulation phasor with a time constant matching the characteristic round-trip of the cavity. Based on this approach, the formula predicted that a monotonical degradation of mode-locking quality should be observed by increasing the QD layers, due to the fact that differential gain increases dramatically for 5 QD layers, whereas remains almost constant for 10 and 15 QD layers devices. Although, these results confirm the experimental behavior of 10 to 15 QD layers, surprisingly in the case of 5 to 10 QD layers, the results are contradictory.
This discrepancy can be attributed to the fact that ES emission is present in the 5 QD layers devices due to early saturation of the GS and due to the absence of any additional frequency selection mechanism in the laser cavity that could suppress lasing from either band.
Finally by employing a tapered structure instead of the conventional rectangular structure gain saturation effects were suppressed due to increased waveguide width. Consequently, a significant gain elevation is observed that allowed sub-pico second pulses due to reduced self-phase modulation effects (fig.11) and a peak power of

7.5W directly from a monolithically integrated tapered mode-locked laser. Last but not least after the end of this thesis the tapered designs that were proposed were utilized for both laser and semiconductor optical amplifiers and the final peak power achieved after two stage amplifications reached the record level of 34W.

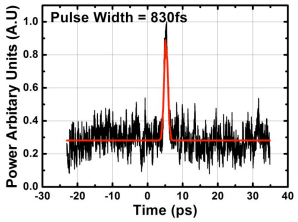

Fig.11. Autocorrelation trace for a two section tapered passively mode locked QD laser.

5 Conclusions

The results reported in this thesis have provided a deeper insight in the mechanism of passive mode locking in quantum dot lasers, whereas it revealed through experimental investigation new regimes of operations. These regimes consist of ground state splitting and the corresponding ability to dual mode lock the two tunable sub-bands, whereas by exploiting the same splitting effect a significant pulse width reduction was observed with simultaneous increase in the gain current and the emitted average power. Moreover, by exploiting the multi-wavelength capabilities of quantum dot devices, novel simultaneous dual wavelength mode-locking from the two main wavebands was achieved directly from a monolithic device.

Furthermore, an experimental evaluation of the feedback tolerance of both main energy-bands under pulsed operation was also performed for the first time, allowing the extraction of isolation requirements for pulses of each wavelength. This effort confirmed the low linewidth enhancement factor observed at the excited state alongside the strong dependence of pulse stability with the exact roundtrip of the external cavity. In the same context an alternative approach utilized in order to destabilize the laser operation was the utilization of inhomogeneous pumping in a two-section quantum dot Fabry Perot laser. This approach allowed coherence collapse and the generation of high complexity time series directly through the device without feedback related complicated setups. The beneficial role of feedback was also demonstrated by exploiting these dynamics in order to suppress Q-switching effects and allow wavelength switching. Finally, throughout this thesis several devices were designed and characterized. Through this process the impact of structural parameters like device length, width and number of quantum dot lasers on the quality of the pulsed operations were extracted, allowing clear guidelines for future devices. Last but not least through cavity optimization record level pulse widths and peak power were achieved that consist of 7.5W peak power directly from a tapered gain guided passively mode locked laser, and 34W through a two stage amplification process that employed tapered laser and semiconductor optical amplifier.

References

1. E. U. Rafailov, M. A. Cataluna, and W. Sibbett, “Quantum Dot Mode Locked Lasers”, Nat. Photonics 1, 395 (2007).
2. M. A. Cataluna, D. I. Nikitichev, S. Mikroulis, H. Simos, C. Simos, C. Mesaritakis, D. Syvridis, I. Krestnikov, D. Livshits, and E. U. Rafailov, “Dual-wavelength mode-locked quantum-dot laser, via ground and excited state transitions: experimental and theoretical investigation”, OSA Optics Express, vol. 18, pages 12832-12838, (2010)
3. C. Mesaritakis, C. Simos, H. Simos, I. Krestnikov, and D. Syvridis. “Dual ground-state pulse generation from a passively mode-locked InAs/InGaAs quantum dot laser” Appl. Phys. Lett. 99, 141109 (2011)
4. C. Mesaritakis, D. Syvridis, “Quantum Dot Devices” Springer ISBN: 978-1-4614-3569-3
5. C. Mesaritakis, C. Simos, H. Simos, S. Mikroulis, I. Krestnikov, and D. Syvridis, “Pulse Width Narrowing due to Dual Ground State Emission in Quantum Dot Mode Locked Lasers” Appl. Phys. Lett. 96, 211110 (2010)
6. C. Mesaritakis, C. Simos, H. Simos, S. Mikroulis, I. Krestnikov, E. Roditi, D. Syvridis “Effect of feedback to the Ground and Excited State of a Quantum dot passively mode locked Laser” AIP: Applied Physics Letters Vol. 97 August (2010).
7. H. Simos, C. Simos, C. Mesaritakis, D. Syvridis, “Two Section Quantum Dot Mode-Locked Lasers under Optical Feedback: Pulse Broadening and Harmonic Operation” IEEE Journal of Quantum Electronics (2012).
8. H. Simos, M. Rossetti, C. Simos, C. Mesaritakis, T. Xu, P. Bardella, I. Montrosset, D. Syvridis, “Numerical analysis of passively mode-locked quantum-dot lasers with absorber section at the low-reflectivity output facet” IEEE Journal of Quantum Electronics, Vol.49 No. 1, pp. 3-10, (2013).
9. C. Mesaritakis, A. Argyris, C. Simos, H. Simos, A. Kapsalis, D. Syvridis “Chaotic emission and tunable self-sustained pulsations in a two-section Fabry–Perot quantum dot laser” AIP: Applied Physics Letters Vol. 98, 051104 (2011).
10. C. Mesaritakis, C. Simos, H. Simos, I. Krestnikov, D. Syvridis “External Optical Feedback-Induced Wavelength Selection and Q-switching Elimination in an InAs/InGaAs Passively Mode Locked Quantum Dot Laser” Journal of Optical Society of America - B Vol. 29, No. 5, pp. 1071-1077 (2012).
11. C. Mesaritakis, C. Simos, H. Simos, A. Kapsalis, E. Roditi, D. Syvridis, I. Krestnikov, “Effect of the Number of Quantum Dot Layers and Dual State Emission on the Performance of InAs/InGaAs Passively Mode-Locked Lasers”, AIP Applied Physics Letters Vol.101, 25 pp. 251115 (2012).
12. Y. Ding, R. Aviles-Espinosa, M. A. Cataluna, D. Nikitichev, M. Ruiz, M. Tran, Y. Robert, A. Kapsalis, H. Simos, C. Mesaritakis, T. Xu, P. Bardella, M. Rossetti, I. Krestnikov, D. Livshits, Ivo Montrosset, D. Syvridis, M. Krakowski, P. Loza-Alvarez, and E. Rafailov, “High peak power picosecond optical pulse generation using a quantum-dot external-cavity passively mode-locked laser and a quantum-dot optical amplifier” accepted for publication OSA Optics Express (2012).

Bayesian signal processing techniques for hyperspectral image unmixing

Themelis Konstantinos*

Department of Informatics and Telecommunications
National and Kapodistrian University of Athens
TYPA Buildings, University Campus, 15784 Athens, Greece
themelis@space.noa.gr

Abstract. In this dissertation the problem of semi-supervised spectral unmixing is studied. We describe a framework of Bayesian techniques that take into account the special characteristics of hyperspectral data and exploit the prior knowledge of the model's constraints. The assumption of sparsity on the unmixing process is introduced, which leads to the development of efficient schemes for spectral unmixing. Besides, a novel method for Bayesian inference is proposed, termed BI-ICE, which is computationally efficient, and can be considered as a first-moments approximation to variational inference methods. Experimental results conducted both on simulated and real hyperspectral data are presented that illustrate the robust estimation performance of the proposed sparsity-inducing methods.

Keywords: Hyperspectral imagery, Spectral unmixing, Sparse linear regression, Bayesian analysis, Variational inference.

1 Introduction

Over the past decades hyperspectral image analysis has emerged as one of the fastest growing technologies in the field of remote sensing. Hyperspectral imagery refers to the process of remotely obtaining data about an object (in our case a geographical area), with the use of hyperspectral sensors. Hyperspectral sensors have the ability to sample the electromagnetic spectrum in hundreds of continuous spectral bands. As a consequence, each pixel of a hyperspectral image is represented by a vector, where each coefficient is a measurement of reflectance at a respective wavelength.

The immense growth of hyperspectral imaging applications has also been accompanied by the development of numerous techniques to effectively process the overwhelming amount of collected data. Sophisticated algorithms have been proposed in the literature that either ameliorate the shortcomings of hyperspectral data, or focus on the exploitation of their informational content. Among all possible signal processing applications that exploit hyperspectral data, in this dissertation we are more interested in the thematic area of spectral unmixing.

The process of hyperspectral unmixing is described by two major steps: (a) the endmember extraction step, and (b) the inversion process. In the endmember

* Dissertation Advisor: Prof. Sergios Theodoridis.

extraction step the spectral signatures of the endmembers contributing to the hyperspectral image are determined. Popular endmember extraction algorithms include the pixel purity index (PPI), [1], the N-FINDR algorithm, [2], and the vertex component analysis (VCA) method, [3]. The inversion process determines the abundances corresponding to the estimated endmembers obtained in the previous step. The abundances should satisfy two constraints, in order to remain physically meaningful; they should be non-negative and sum to one. Under these constraints, spectral unmixing is formulated as a convex optimization problem, which can be addressed using iterative methods, e.g., the fully constrained least squares method, [4], or numerical optimization methods, e.g., [5]. Bayesian methods have also been proposed for the problem, e.g., the Gibbs sampling scheme applied to the hierarchical Bayesian model of [6]. Semi-supervised unmixing, [6, 7], which is considered in this dissertation, assumes that the endmembers' spectral signatures are available. The objective of semi-supervised unmixing is (a) to determine how many and which endmembers are present in the mixed pixel under study and (b) to estimate their corresponding abundances.

2 Work contributions

The scientific contributions that appear in the present work exhibit a Bayesian treatment for the problem of supervised spectral unmixing. A major advantage of the Bayesian approach is that it provides a flexible framework to represent the probabilistic mechanisms of data generation and our prior information about it. To this end, parametric probabilistic models are considered, and appropriate prior distributions are employed to capture the uncertainties of the model parameters.

In the context of hyperspectral signal processing, we were greatly influenced by the wide applicability of Bayesian methods. Our first contribution is a Bayesian maximum a posteriori (MAP) estimator, published in [8], which is specifically designed to account for the convex constraints of the abundance estimation problem. Utilizing the simplicity of the Gaussian distribution and the symmetry of constraints, closed form expressions are derived for the modes of the posterior distribution of the abundances. Following this path an efficient estimator is proposed, which has almost similar estimation performance, but is computationally more efficient than quadratic programming methods that address the same constrained estimation problem.

The applicability of our proposed schemes has also been of concern. A representative example is the case study on the unmixing of real hyperspectral data collected from the OMEGA spectrometer, [9], aboard the Mars Express mission. The objective of the OMEGA spectrometer is to collect information which will help to determine the mineral composition of the Mars surface. This is an interesting application for spectral unmixing. Three methods are considered and compared in the unmixing process: (a) the ENVI-SVD method, which is commercially available through the popular remote sensing software ENVI, (b) the MAP estimator of the preceding paragraph, and (c) a quadratic programming method, [5], which is an iterative Newton method. The results of this comparison

are available in [10, 11], where it is seen that the MAP estimator can provide reliable results and outperform existing methods.

In the sequel, a significant amount of efforts was invested to leverage sparse signal processing techniques for spectral unmixing. The primary assumption is that only a small number of endmembers will be mixed in a single pixel, and hence, the abundance estimation problem will inherently have a sparse solution. To the best of our knowledge, we were among the first to introduce the notion of sparsity to the problem of spectral unmixing. This consideration is exhibited in the publications [7, 12, 13]. In both publications a variant of the least absolute shrinkage and selection operator (lasso), [14], was selected to impose sparsity. In [7], the adaptively weighted lasso, [15], is utilized and special manipulation is provided for the constraints of the problem. To force the nonnegativity constraint, the optimization problem is solved using a modified LARS algorithm, which retrieves only nonnegative solutions. Moreover, the additivity constraint is included in the quadratic loss function of least squares, by means of an appropriate extra linear equation.

Another contribution of this thesis is the development of a Bayesian hierarchical model analogous to the adaptive lasso, [15]. In the proposed Bayesian setup, independent Laplace priors are employed by the model to correspond to the weighted ℓ_1 norm penalization of the adaptive lasso. Besides, the nonnegativity constraint is incorporated to the model by a truncation operation on the prior distributions. A novel method for Bayesian inference is then developed, termed as Bayesian inference iterative conditional expectations algorithm (BI-ICE). BI-ICE appears to be a first-moments approximation to variational approximation methods, and is summarized in Section 3. The proposed Bayesian approach is analytically described in [13].

Finally, a lot of effort has also been invested into the research and development of a sparse reconstruction algorithm, in the framework of Bayesian compressive sensing. The previous Bayesian set-up has been adopted to promote sparse solutions to an underdetermined system of linear equations. To perform Bayesian inference, a recently proposed sub-optimal, type-II maximum likelihood algorithm was adjusted to fit the needs of the present framework. The resulting incremental-type algorithm has superior performance when compared to other Bayesian compressive sensing methods, as illustrated in Chapter 4 of the dissertation.

3 Semi-supervised hyperspectral unmixing via a novel Bayesian approach

An interesting perspective of the semi-supervised spectral unmixing problem arises when the latent sparsity of the abundance vector is taken into account. A reasonable assumption is that only a small number of endmembers are mixed in a single pixel, and hence, the solution to the endmember determination and abundance estimation problem is inherently sparse. This lays the ground for the utilization of sparse signal representation techniques, e.g., [16–18], in semi-supervised unmixing. A number of such semi-supervised unmixing techniques has been recently proposed in [7, 19, 20], based on the concept of ℓ_1 norm penalization

to enhance sparsity. These methods assume that the spectral signatures of many different materials are available, in the form of a spectral library. Since only a small number of the available materials' spectra are expected to be present in the hyperspectral image, the abundance vector is expected to be sparse.

Let $\mathbf{y}$ be a $M \times 1$ hyperspectral image pixel vector, where M is the number of spectral bands. Also let $\boldsymbol{\Phi} = [\boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \ldots, \boldsymbol{\phi}_N]$ stand for the $M \times N$ signature matrix of the problem, with $M > N$, where the $M \times 1$ dimensional vector $\boldsymbol{\phi}_i$ represents the spectral signature (i.e., the reflectance values in all spectral bands) of the ith endmember and N is the total number of distinct endmembers. Finally, let $\mathbf{w} = [w_1, w_2, \ldots, w_N]^T$ be the $N \times 1$ abundance vector associated with $\mathbf{y}$, where w_i denotes the abundance fraction of $\boldsymbol{\phi}_i$ in $\mathbf{y}$.

The linear mixture model (LMM) is adopted, that is, the previous quantities are assumed to be interrelated as follows

$$\mathbf{y} = \boldsymbol{\Phi}\mathbf{w} + \mathbf{n}. \tag{1}$$

The additive noise $\mathbf{n}$ is assumed to be a zero-mean Gaussian distributed random vector, with independent and identically distributed (i.i.d.) elements, i.e., $\mathbf{n}|\beta \sim \mathcal{N}(\mathbf{n}|\mathbf{0}, \beta^{-1}\mathbf{I}_M)$, where β denotes the inverse of the noise variance (precision). Due to the nature of the problem, the abundance vector is usually assumed to satisfy the following two constraints

$$w_i \geq 0,\ i = 1, 2, \ldots, N,\ \text{and}\ \sum_{i=1}^{N} w_i = 1, \tag{2}$$

namely, a non-negativity constraint and a sum-to-one (additivity) constraint. Based on this formulation, a semi-supervised hyperspectral unmixing technique is introduced, where the endmember matrix $\boldsymbol{\Phi}$ is assumed to be known a priori. As mentioned before, each column of $\boldsymbol{\Phi}$ contains the spectral signature of a single material, and its elements are non-negative, since they represent reflectance values. The mixing matrix $\boldsymbol{\Phi}$ can either stem from a spectral library or it can be determined using an endmember extraction technique, e.g., [3]. However, the actual number of endmembers that compose a single pixel's spectrum, denoted as ξ, is unknown and may vary from pixel to pixel. Sparsity is introduced when $\xi \ll N$, that is by assuming that only few of the available endmembers are present in a single pixel. This is a reasonable assumption, that is in line with intuition, since it is likely for a pixel to comprise only a few different materials from a library of several available materials. Summarizing, in semi-supervised unmixing, we are interested in estimating the abundance vector $\mathbf{w}$ for each image pixel, which is non-negative and sparse, with ξ out of its N entries being non-zero.

3.1 Hierarchical Bayesian model

Considering the observation model defined in (1) and the Gaussian property of the additive noise, the likelihood function of $\mathbf{y}$ can be expressed as follows

$$p(\mathbf{y}|\mathbf{w}, \beta) = \mathcal{N}(\mathbf{y}|\boldsymbol{\Phi}\mathbf{w}, \beta^{-1}\mathbf{I}_M) = (2\pi)^{-\frac{M}{2}}\, \beta^{\frac{M}{2}} \exp\left[-\frac{\beta}{2}\|\mathbf{y} - \boldsymbol{\Phi}\mathbf{w}\|_2^2\right]. \tag{3}$$

Accounting for the non-negativity property of $\mathbf{w}$, and assuming that all w_i's are i.i.d., a normal distribution truncated on the non-negative orthant $\mathbf{R}_+^N$ of the N-dimensional Euclidean space $\mathcal{R}^N$ is assigned to $\mathbf{w}$, i.e.,

$$\begin{aligned} p(\mathbf{w}|\boldsymbol{\gamma},\beta) &= \prod_{i=1}^{N}\left[\mathcal{N}(w_i|0,\frac{\gamma_i}{\beta})\,\mathbf{I}_{\mathbf{R}_+^1}(w_i)\right] \\ &= 2^N\,(2\pi)^{-\frac{N}{2}}\,\beta^{\frac{N}{2}}\,|\boldsymbol{\Lambda}|^{\frac{1}{2}}\exp\left[-\frac{\beta}{2}\mathbf{w}^T\boldsymbol{\Lambda}\mathbf{w}\right]\mathbf{I}_{\mathbf{R}_+^N}(\mathbf{w}) \end{aligned} \tag{4}$$

where $\mathbf{R}_+^1$ is the set of non-negative real numbers, $\mathbf{I}_{\mathbf{R}_+^N}(\cdot)$ is the indicator function[1] for $\mathbf{R}_+^N$, $\boldsymbol{\gamma} = [\gamma_1, \gamma_2, \ldots, \gamma_N]^T$ is a $N \times 1$ vector of hyperparameters, $\gamma_i \geq 0$, and $\boldsymbol{\Lambda}^{-1} = \mathrm{diag}(\boldsymbol{\gamma})$.

For β, a conjugate Gamma prior with respect to the Gaussian likelihood of (3) is selected, expressed as

$$p(\beta|\kappa,\theta) = \Gamma(\beta|\kappa,\theta) = \frac{\theta^\kappa}{\Gamma(\kappa)}\beta^{\kappa-1}\exp\left[-\theta\beta\right], \tag{5}$$

where $\beta \geq 0$, and $\kappa \geq 0$, $\theta \geq 0$ are the distribution parameters.

We extend the model of [21, 22], by assigning an independent Gamma distribution to every γ_i, each parameterized by a distinct hyperparameter λ_i, i.e.,

$$p(\gamma_i|\lambda_i) = \Gamma(\gamma_i|1,\frac{\lambda_i}{2}) = \frac{\lambda_i}{2}\exp\left[-\frac{\lambda_i}{2}\gamma_i\right],\ i = 1, 2, \ldots, N, \tag{6}$$

Then, the combination of the hierarchical priors given in (4) and (6) leads to a sparsity-promoting, non-negative (truncated) Laplace distribution for $\mathbf{w}$ (this formulation gives rise to the so-called Bayesian lasso [22]). This distribution can be obtained by marginalizing the hyperparameter vector $\boldsymbol{\gamma}$ from the model, i.e.,

$$p(\mathbf{w}|\boldsymbol{\lambda},\beta) = \int p(\mathbf{w}|\boldsymbol{\gamma},\beta)p(\boldsymbol{\gamma}|\boldsymbol{\lambda})d\boldsymbol{\gamma} = \beta^{\frac{N}{2}}\,|\boldsymbol{\Psi}|^{\frac{1}{2}}\exp\left[-\sqrt{\beta}\sum_{i=1}^{N}\sqrt{\lambda_i}\,|w_i|\right]\mathbf{I}_{\mathbf{R}_+^N}(\mathbf{w}), \tag{7}$$

where $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \ldots, \lambda_N]^T$ and $\boldsymbol{\Psi} = \mathrm{diag}(\boldsymbol{\lambda})$. The motivation to use a hyperparameter vector $\boldsymbol{\lambda}$ instead of a single λ for all γ_i's as in [22, 21], is to form a hierarchical Bayesian analogue to the adaptive lasso, proposed in [15]. Indeed, it can be shown, that the maximum a posteriori (MAP) estimator of $\mathbf{w}$, which is distributed according to (7), is the solution to the following optimization problem,

$$\tilde{\mathbf{w}} = \arg\min_{\mathbf{w}}\left\{\frac{\beta}{2}\left\|\mathbf{y}-\boldsymbol{\Phi}\mathbf{w}\right\|_2^2 + \sum_{i=1}^{N}\alpha_i|w_i|\right\},\ \text{s.t. } \mathbf{w} \in \mathbf{R}_+^N, \tag{8}$$

which, excluding the non-negativity constraint, coincides with the definition of the adaptive lasso, [15].

[1] $\mathbf{I}_{\mathbf{R}_+^N}(\mathbf{x}) = 1(0), \text{if } \mathbf{x} \in \mathbf{R}_+^N (\mathbf{x} \notin \mathbf{R}_+^N)$.

It is obvious from (7) that the quality of the endmember selection procedure depends on the tuning parameter vector $\boldsymbol{\lambda}$. We choose to infer the hyperparameter vector $\boldsymbol{\lambda}$ from the data, by assuming a Gamma hyperprior for each element of $\boldsymbol{\lambda}$,

$$p(\lambda_i|r,\delta) = \Gamma(\lambda_i|r,\delta) = \frac{\delta^r}{\Gamma(r)}{\lambda_i}^{r-1}\exp\left[-\delta\lambda_i\right], \quad i = 1,2,\ldots,N \tag{9}$$

where r and δ are hyperparameters, with $r \geq 0$ and $\delta \geq 0$. Both Gamma priors of β, in (5), and λ_i, in (9), are flexible enough to express prior information, by properly tuning their hyperparameters. In this paper, the hyperparameters $\kappa, \theta, r, \delta$ are set to zero as in [6, 21], forming non-informative (Jeffreys') priors, although other values can, in principle, be selected.

3.2 Bayesian inference

As it is common in Bayesian inference, the estimation procedure is based on the computation of the joint posterior distribution of the parameters. For the model presented in Section 3.1, this posterior is

$$p\left(\mathbf{w},\beta,\boldsymbol{\gamma},\boldsymbol{\lambda}|\mathbf{y}\right) = \frac{p\left(\mathbf{y}|\mathbf{w},\beta\right)p\left(\mathbf{w}|\beta,\boldsymbol{\gamma}\right)p\left(\boldsymbol{\gamma}|\boldsymbol{\lambda}\right)p\left(\boldsymbol{\lambda}\right)p\left(\beta\right)}{p(\mathbf{y})}, \tag{10}$$

which is intractable, because $p(\mathbf{y})$ cannot be computed analytically. To overcome this obstacle, a Markovian Gibbs sampling strategy can be followed, in which the conditional posterior distributions of the associated parameters are utilized.

Posterior conditional distributions In the following, analytical expressions are derived for the posterior conditional distributions of the model parameters $\mathbf{w}$, $\boldsymbol{\gamma}$, $\boldsymbol{\lambda}$ and β. Starting with $\mathbf{w}$, it can be easily shown that its posterior conditional density is the multivariate Gaussian, truncated in $\mathbf{R}_+^N$,

$$p(\mathbf{w}|\mathbf{y},\boldsymbol{\gamma},\boldsymbol{\lambda},\beta) = \mathcal{N}(\mathbf{w}|\boldsymbol{\mu},\boldsymbol{\Sigma})\mathbf{I}_{\mathbf{R}_+^N} = \mathcal{N}_{\mathbf{R}_+^N}(\mathbf{w}|\boldsymbol{\mu},\boldsymbol{\Sigma}) \tag{11}$$

where $\boldsymbol{\Sigma}$ and $\boldsymbol{\mu}$ are respectively expressed as follows,

$$\boldsymbol{\Sigma} = \beta^{-1}\left[\boldsymbol{\Phi}^T\boldsymbol{\Phi} + \boldsymbol{\Lambda}\right]^{-1}, \qquad \boldsymbol{\mu} = \beta\boldsymbol{\Sigma}\boldsymbol{\Phi}^T\mathbf{y}. \tag{12}$$

The posterior conditional for the precision parameter β, is easily shown to be a Gamma distribution, i.e.,

$$p(\beta|\mathbf{y},\mathbf{w},\boldsymbol{\gamma},\boldsymbol{\lambda}) = \Gamma\left(\beta|\frac{M+N}{2} + \kappa, \frac{1}{2}\|\mathbf{y}-\boldsymbol{\Phi}\mathbf{w}\|_2^2 + \theta + \frac{1}{2}\mathbf{w}^T\boldsymbol{\Lambda}\mathbf{w}\right) \tag{13}$$

Straightforward computations yield that the conditional distribution of γ_i given $\mathbf{y}, w_i, \lambda_i, \beta$ is expressed as

$$p(\gamma_i|\mathbf{y},w_i,\lambda_i,\beta) = \left(\frac{\lambda_i}{2\pi}\right)^{\frac{1}{2}}{\gamma_i}^{-\frac{1}{2}}\exp\left[-\frac{\beta w_i^2}{2\gamma_i} - \frac{\lambda_i}{2}\gamma_i + \sqrt{\beta\lambda_i}\,|w_i|\right], i = 1,2,\ldots,N \tag{14}$$

Finally, the conditional posterior of λ_i given $\mathbf{y}, w_i, \gamma_i, \beta$ also turns out to be a Gamma distribution,

$$p(\lambda_i|\mathbf{y}, w_i, \gamma_i, \beta) = \mathit{\Gamma}\left(\lambda_i|1+r, \frac{\gamma_i}{2}+\delta\right),\ i = 1, 2, \ldots, N. \tag{15}$$

In the sequel, we propose a deterministic approximation of the Gibbs sampler, where the randomly generated samples of the Gibbs sampler are replaced by the *means* of the corresponding conditional distributions, (11), (13), (14) and (15), as explained in Section 3.2. Thus, a novel scheme iterating among the conditional means of $\mathbf{w}$, β, γ_i and λ_i arises, which will be termed *Bayesian inference iterative conditional expectations* (BI-ICE) algorithm. It should be emphasized that by following this approach, we depart from the statistical framework implied by the Gibbs sampler and we end up with a new deterministic algorithm for estimating the parameters of the proposed hierarchical model.

The proposed BI-ICE algorithm As mentioned previously, BI-ICE needs the conditional expectations of the model parameters. These are computed analytically as described below:

Expectation of $p(\mathbf{w}|\mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \beta)$ $\mathbf{w}$: As shown in (11), $p(\mathbf{w}|\mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \beta)$ is a multivariate Gaussian distribution, truncated in $\mathbf{R}_+^N$. In the one-dimensional case, the expectation of the truncated Gaussian distribution in $\mathbf{R}_+^1$ can be computed as

$$x \sim \mathcal{N}_{\mathbf{R}_+^1}(x|\mu^*, \sigma^*) \Rightarrow \mathrm{E}[x] = \mu^* + \frac{\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\frac{\mu^{*2}}{\sigma^{*2}}\right)}{1-\frac{1}{2}\mathrm{erfc}\left(\frac{\mu^*}{\sqrt{2}\sigma^*}\right)}\sigma^*, \tag{16}$$

where $\mathrm{erfc}(\cdot)$ is the complementary error function. Unfortunately, to the best of our knowledge, there is no analogous closed-form expression for the N-dimensional case. However, as shown in [23], the distribution of the ith element of $\mathbf{w}$ conditioned on the remaining elements $\mathbf{w}_{\neg i} = [w_1, \ldots, w_{i-1}, w_{i+1}, \ldots, w_N]^T$, can be expressed as

$$w_i|\mathbf{w}_{\neg i} \sim \mathcal{N}_{\mathbf{R}_+^1}(w_i|\mu_i^*, \sigma_{ii}^*) \tag{17}$$

$$\mu_i^* = \mu_i + \boldsymbol{\sigma}_{\neg i}^T \boldsymbol{\Sigma}_{\neg i \neg i}^{-1}(\mathbf{w}_{\neg i} - \boldsymbol{\mu}_{\neg i}) \tag{18}$$

$$\sigma_{ii}^* = \sigma_{ii} - \boldsymbol{\sigma}_{\neg i}^T \boldsymbol{\Sigma}_{\neg i \neg i}^{-1} \boldsymbol{\sigma}_{\neg i}, \tag{19}$$

where matrix $\boldsymbol{\Sigma}_{\neg i \neg i}$ is formed by removing the ith row and the ith column from $\boldsymbol{\Sigma}$, the $(N-1)\times 1$ vector $\boldsymbol{\sigma}_{\neg i}$ is the ith column of $\boldsymbol{\Sigma}$ after removing its ith element σ_{ii} and μ_i is the ith element of $\boldsymbol{\mu}$. Based on this result, an iterative procedure is proposed in order to compute the mean of the posterior $p(\mathbf{w}|\mathbf{y}, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \beta)$. Specifically, the j-th iteration, $j = 1, 2, \ldots$, of this procedure is

described as follows[2]

$$\begin{aligned} &1.\ w_1^{(j)} = \mathrm{E}[p(w_1|w_2^{(j-1)}, w_3^{(j-1)}, \ldots, w_N^{(j-1)})] \\ &2.\ w_2^{(j)} = \mathrm{E}[p(w_2|w_1^{(j)}, w_3^{(j-1)}, \ldots, w_N^{(j-1)})] \\ &\qquad\vdots \\ &N.\ w_N^{(j)} = \mathrm{E}[p(w_N|w_1^{(j)}, w_2^{(j)}, \ldots, w_{N-1}^{(j)})] \end{aligned} \tag{20}$$

This procedure is repeated iteratively until convergence. Experimental results have shown that the iterative scheme in (20) converges to the mean of $\mathbf{w} \sim \mathcal{N}_{\mathbf{R}_+^N}(\mathbf{w}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ after a few iterations.

Expectation of $p(\beta|\mathbf{y}, \mathbf{w}, \boldsymbol{\gamma}, \boldsymbol{\lambda})$: The mean value of the Gamma distribution of (13) is given by

$$\mathrm{E}\left[p(\beta|\mathbf{y}, \mathbf{w}, \boldsymbol{\gamma}, \boldsymbol{\lambda})\right] = \frac{\frac{M+N}{2} + \kappa}{\frac{1}{2}\|\mathbf{y} - \boldsymbol{\Phi}\mathbf{w}\|_2^2 + \theta + \frac{1}{2}\mathbf{w}^T\boldsymbol{\Lambda}\mathbf{w}} \tag{21}$$

Expectation of $p(\gamma_i|\mathbf{y}, w_i, \lambda_i, \beta)$: It can be shown that the expectation of (14) is expressed as

$$\mathrm{E}\left[p(\gamma_i|\mathbf{y}, w_i, \lambda_i, \beta)\right] = \left(\frac{2\lambda_i}{\pi}\right)^{\frac{1}{2}} \left(\frac{\beta w_i^2}{\lambda_i}\right)^{\frac{3}{4}} \exp\left[\sqrt{\beta\lambda_i}\,|w_i|\right] K_{3/2}\left(\sqrt{\beta\lambda_i}\,|w_i|\right), \tag{22}$$

where $K_\nu(\cdot)$ stands for the modified Bessel function of second kind of order ν.

Expectation of $p(\lambda_i|\mathbf{y}, w_i, \gamma_i, \beta)$: Finally, the mean value of the Gamma distribution (15) is

$$\mathrm{E}\left[\lambda_i|\mathbf{y}, w_i, \gamma_i, \beta)\right] = \frac{1+r}{\frac{1}{2}\gamma_i + \delta}. \tag{23}$$

The basic steps of the proposed BI-ICE algorithm are summarized in Table 1. Regarding the updating of parameter $\mathbf{w}^{(t)}$, an auxiliary variable $\mathbf{v}$ has been utilized in Table 1. This is initialized with $\boldsymbol{\mu}^{(t)}$ (the value of $\boldsymbol{\mu}$ at iteration t) and is updated by performing a *single* iteration of the scheme described in (20). The resulting value of $\mathbf{v}$ is assigned to $\mathbf{w}^{(t)}$. The rationale behind this choice is that for $\boldsymbol{\Sigma}$ diagonal (which happens when the columns of $\boldsymbol{\Phi}$ are orthogonal), it easily follows from (18), (19) that the w_i's in (20) are uncorrelated. Thus, a single iteration is sufficient to obtain the mean of $\mathcal{N}_{\mathbf{R}_+^N}(\mathbf{w}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Although, this is not valid when $\boldsymbol{\Sigma}$ is not diagonal, experimental results have evidenced that the estimation of the mean of $\mathcal{N}_{\mathbf{R}_+^N}(\mathbf{w}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ resulting after the execution of a single iteration of the scheme in (20) is also sufficient in the framework of the BI-ICE algorithm.

A basic advantage of the proposed Bayesian approach, which is the Bayesian analogue to the adaptive lasso, is that all parameters are naturally estimated

[2] In the following, for notational simplicity, the expectation $E_{p(x|y)}[x]$ of a random variable x with conditional distribution $p(x|y)$ is denoted as $E[p(x|y)]$.

Table 1. The BI-ICE algorithm

Input $\boldsymbol{\Phi}$, $\mathbf{y}$, $\kappa, \theta, r, \delta$
Initialize $\boldsymbol{\gamma}^{(0)} = \boldsymbol{\lambda}^{(0)} = \mathbf{1}, \beta^{(0)} = 0.01 \|\mathbf{y}\|_2$
for $t = 1, 2, \ldots$ **do**
- Compute $\mathbf{w}^{(t)}$ as follows
  Compute $\boldsymbol{\Sigma}^{(t)}, \boldsymbol{\mu}^{(t)}$ using (12)
  Set $\mathbf{v}^{(0)} = \boldsymbol{\mu}^{(t)}$
  Compute $v_1^{(1)} = \mathrm{E}\left[p(v_1|v_2^{(0)}, \ldots, v_N^{(0)})\right]$,
  using (18), (19), and (16)
  Compute $v_2^{(1)} = \mathrm{E}\left[p(v_2|v_1^{(1)}, v_3^{(0)}, \ldots, v_N^{(0)})\right]$,
  using (18), (19), and (16)
  $\vdots$
  Compute $v_N^{(1)} = \mathrm{E}\left[p(v_N|v_1^{(1)}, v_2^{(1)}, \ldots, v_{N-1}^{(1)})\right]$,
  using (18), (19), and (16)
  Set $\mathbf{w}^{(t)} = \mathbf{v}^{(1)}$
- Compute $\beta^{(t)} = \mathrm{E}\left[p(\beta|\mathbf{y}, \mathbf{w}^{(t)}, \boldsymbol{\gamma}^{(t-1)}, \boldsymbol{\lambda}^{(t-1)}\right]$, using (21)
- Compute $\gamma_i^{(t)} = \mathrm{E}\left[p(\gamma_i|\mathbf{y}, w_i^{(t)}, \lambda_i^{(t-1)}, \beta^{(t)})\right]$,
  $i = 1, 2, \ldots, N$, using (22)
- Compute $\lambda_i^{(t)} = \mathrm{E}\left[p(\lambda_i|\mathbf{y}, w_i^{(t)}, \gamma_i^{(t)}, \beta^{(t)})\right]$,
  $i = 1, 2, \ldots, N$, using (23)
endfor

from the data. In contrast, deterministic algorithms for solving the lasso, e.g. [15], face the problem of fine-tuning specific parameters (corresponding to $\boldsymbol{\lambda}$ of our model), that control the sparsity of the solution. As shown in the simulations presented in the next section, the BI-ICE algorithm converges very fast, and retains the sparsity of the solution. It has been further observed that by initializing $\mathbf{w}$ with $\boldsymbol{\mu}$, a single cycle is sufficient for the inner sampler to converge. The computational complexity of the proposed method can be further reduced by avoiding the explicit computation of the matrices $\boldsymbol{\Sigma}_{\neg i \neg i}^{-1}$ in (18), (19).

4 Experimental results and discussion

The performance of the BI-ICE algorithm is illustrated by unmixing a synthetic hyperspectral image, using data from the USGS spectral library. Specifically, 30 endmembers were selected from the library, to construct a 453×220 endmember matrix, having condition number 36.182×10^6. The number of disparate endmembers composing a single pixel varied between one (pure pixel) and five, whereas the abundances were generated according to a Dirichlet distribution, so as to satisfy the positivity and sum-to-one constraints. The observations were corrupted by Gaussian noise, whose variance was determined by the SNR level. First, the fast convergence and sparse estimations of $\mathbf{w}$ exhibited by the new algorithm are depicted in Fig. 1a. In this experiment, a pixel with three non-

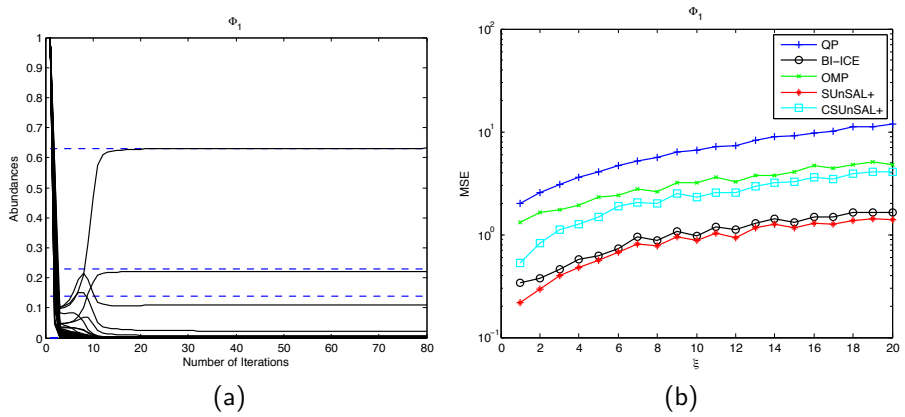

Fig. 1. (a) Estimation of the entries of the sparse vector $\mathbf{w}$, as BI-ICE progresses. The algorithm is applied to simulated data, generated using a highly correlated matrix of spectral data. White noise is added (SNR = 25 dB). Dashed lines: true values. Solid lines: estimated values. (b) MSE as a function of the level of sparsity obtained by different unmixing methods when applied to simulated data with white additive noise (SNR = 20 dB)

zero abundances (0.1397, 0.2305, 0.6298) is considered, and white noise is added to the model, such that the SNR is equal to 25dB. The curves in Fig. 1a are the average of 50 noise realizations. We observe that less than 15 iterations are sufficient for the BI-ICE algorithm to converge to the correct sparse solution of $\mathbf{w}$. That is, it determines correctly the abundance fractions of the endmembers present in the pixel, while all remaining abundance fractions converge to zero.

Next, the BI-ICE algorithm was compared to: (a) the least squares (LS) algorithm, (b) a quadratic programming (QP) technique, which enforces the constraints, but does not specifically exploit the problem's sparsity, [5], (c) the orthogonal matching pursuit (OMP) algorithm, [17], which is a widely used, greedy, sparsity promoting algorithm, (d) the sparse unmixing by variable splitting and augmented Lagrangian (SUnSAL) algorithm, [20, 24], and (e) the constrained version of SUnSAL, CSUnSAL, which solves the constrained version of the lasso operator, (see also [24] for details). In our experiments, the parameters used for SUnSAL are $\mu = 1$, and $\lambda = 1$, while for CSUnSAL we used $\mu = 1$, $\lambda = 10^{-3}$, and $\delta = 10^{-6}$, see also [20]. Based on the following metric, $\text{MSE} = \text{E}\left[\frac{\|\mathbf{w}-\tilde{\mathbf{w}}\|_2^2}{\|\mathbf{w}\|_2^2}\right]$, where $\mathbf{w}$ and $\tilde{\mathbf{w}}$ are the true and the estimated abundance vectors respectively, the corresponding MSE curves for different sparsity levels ranging from 1 (pure pixel) to 20 are shown in Fig. 1b. Due to poor results, the MSE curve of the LS algorithm is not shown in the figure. It can be seen that the proposed algorithm outperforms the OMP, QP, and CSUnSAL algorithms and has similar performance to the SUnSAL algorithm.

In comparison to BI-ICE, the adaptive methods SUnSAL and CSUnSAL are of lower computational complexity. However, it should be pointed out that the

comparable performance, in terms of MSE, of the alternating direction algorithms SUnSAL and CSUnSAL with BI-ICE comes at the additional expense of manually fine-tuning nontrivial parameters, such as the sparsity promoting parameter λ, (see [24]). Thus, an advantage of the proposed BI-ICE algorithm over SUnSAL and CSUnSAL algorithms is that all unknown parameters are directly inferred from the data. Besides that, BI-ICE bears interesting by-products such as: (a) estimates of all model parameters; a useful parameter in many applications is the noise variance, (b) estimates for the variances of the estimated parameters, which may serve as confidence intervals, and (c) approximate posterior distributions for the estimated parameters. In contrast, all other algorithms considered are iterative algorithms that return point estimates of the parameters of interest.

5 Conclusions

In this dissertation we have presented a general framework for semi-supervised Bayesian learning, comprised of soft-constraint Bayesian estimation methods, sparsity-promoting hierarchical Bayesian models and novel methods for Bayesian inference. Under the assumption of the linear mixing model, spectral unmixing has been formulated as a standard linear regression problem, where the parameters of interest are the abundance fractions of the endmembers spectra in each pixel of a hyperspectral scene. A Bayesian MAP estimator has been suitably adjusted to address the abundance estimation problem, taking into account the constraints of nonnegativity and full additivity. To the best of our knowledge, we were among the first to introduce sparsity in the context of spectral unmixing. In order to induce sparsity, our research efforts have focused around the lasso and its variants. First, we have successfully applied the adaptively weighted lasso to the unmixing problem. A fully Bayesian treatment for the unmixing problem was further developed, wherein a suitably selected hierarchical Bayesian model and an efficient method to perform Bayesian inference are proposed. Finally, we attempt to demonstrate the potential of the hierarchical Bayesian model by considering a sparse image reconstruction problem. To this end, a fast, suboptimal, type-II maximum likelihood algorithm has been developed.

References

1. J. W. Boardman, Automating spectral unmixing of AVIRIS data using convex geometry concepts, in: Proc. Summaries 4th Annu. JPL Airborne Geosci. Workshop, Vol. 1, Washington, DC, 1993, pp. 11–14.
2. M. E. Winter, N-FINDR: an algorithm for fast autonomous spectral end-member determination in hyperspectral data, in: Proc. SPIE Imaging Spectrometry V, Vol. 3753, 1999, pp. 266 –275.
3. J. M. P. Nascimento, J. M. Bioucas-Dias, Vertex component analysis: A fast algorithm to unmix hyperspectral data, IEEE Trans. Geosci. Remote Sensing 4 (2005) 898–910.
4. D. C. Heinz, C. I. Chang, Fully constrained least squares linear spectral mixture analysis method for material quantification in hyperspectral imagery, IEEE Trans. Geosci. Remote Sensing 39 (2001) 529–545.

5. T. F. Coleman, Y. Li, A reflective Newton method for minimizing a quadratic function subject to bounds on some of the variables, SIAM Journal on Optimization 6 (1996) 1040–1058.
6. N. Dobigeon, J.-Y. Tourneret, C.-I. Chang, Semi-supervised linear spectral unmixing using a hierarchical Bayesian model for hyperspectral imagery, IEEE Trans. Signal Process. 56 (7) (2008) 2684–2695.
7. K. E. Themelis, A. A. Rontogiannis, K. D. Koutroumbas, Semi-supervised hyperspectral unmixing via the weighted Lasso, in: Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'10), Dallas, Texas, 2010.
8. K. E. Themelis, A. A. Rontogiannis, A soft constrained MAP estimator for supervised hyperspectral signal unmixing, in: Proc. of European Signal Processing Conference, (EUSIPCO), 2008.
9. J. P. Bibring, A. Soufflot, M. Berthe', Y. Langevin, B. Gondet, P. Drossart, M. Bouye', M. Combes, P. Puget, G. Bellucci, et al., OMEGA: Observatoire pour la Mine'ralogie, l'Eau, les Glaces et l'Activit, Vol. SP-1240, ESA, 2004, pp. 37–50.
10. K. E. Themelis, A. A. Rontogiannis, O. Sykioti, K. D. Koutroumbas, I. Daglis, F. Schmidt, On the unmixing of MEx/OMEGA hyperspectral data, in: Proc. European Planetary Science Congress (EPSC), 2010.
11. K. E. Themelis, F. Schmidt, O. Sykioti, A. A. Rontogiannis, K. D. Koutroumbas, I. A. Daglis, On the unmixing of MEx/OMEGA hyperspectral data, Planetary and Space Science, Elsevier (2011) –Accepted.
12. K. E. Themelis, A. A. Rontogiannis, K. D. Koutroumbas, Sparse semi-supervised hyperspectral unmixing using a novel iterative bayesian inference algorithm, in: 19th European Signal Processing Conference (EUSIPCO), 2011, pp. 1165–1169.
13. K. E. Themelis, A. A. Rontogiannis, K. D. Koutroumbas, A novel hierarchical Bayesian approach for sparse semi-supervised hyperspectral unmixing, IEEE Trans. Signal Process. 60 (2) (2012) 585–599.
14. R. Tibshirani, Regression shrinkage and selection via the Lasso, Journal of the Royal Statistical Society 58 (1) (1996) 267–288.
15. H. Zou, The adaptive Lasso and its oracle properties, Journal of the American Statistical Association 101 (2006) 1418–1429.
16. B. Efron, T. Hastie, I. Johnstone, R. Tibshirani, Least angle regression, Annals of Statistics 32 (2002) 407–499.
17. J. Tropp, A. Gilbert, Signal recovery from random measurements via orthogonal matching pursuit, IEEE Trans. Inf. Theory 53 (2007) 4655 –4666.
18. S. Ji, Y. Xue, L. Carin, Bayesian compressive sensing, IEEE Trans. Signal Process. 56 (6) (2008) 2346–2356.
19. M.-D. Iordache, J. Bioucas-Dias, A. Plaza, Unmixing sparse hyperspectral mixtures, in: Geoscience and Remote Sensing Symposium, 2009 IEEE International, IGARSS 2009, Vol. 4, Cape Town, 2009, pp. 85 –88.
20. J. Bioucas-Dias, M. Figueiredo, Alternating direction algorithms for constrained sparse regression: Application to hyperspectral unmixing, in: Proc. IEEE International Workshop on Hyperspectral Image and Signal Processing: evolution in Remote Sensing (WHISPERS'10), Reykjavik, Iceland, 2010.
21. S. D. Babacan, R. Molina, A. Katsaggelos, Bayesian compressive sensing using Laplace priors, IEEE Trans. Image Process. 19 (1) (2010) 53 –63.
22. T. Park, C. G., The Bayesian Lasso, Journal of the American Statistical Association 103 (482) (2008) 681–686.
23. G. Rodriguez-Yam, R. Davis, L. Scharf, A Bayesian model and Gibbs sampler for hyperspectral imaging, in: Proc. IEEE Sensor Array and Multichannel Signal Processing Workshop, 2002, pp. 105–109.
24. M.-D. Iordache, J. Bioucas-Dias, A. Plaza, Sparse unmixing of hyperspectral data, Geoscience and Remote Sensing, IEEE Transactions on 49 (6) (2011) 2014 –2039.

Fault Detection Methodology for Caches in Reliable Modern VLSI Microprocessors based on Instruction Set Architectures

Georgios A. Theodorou*

National and Kapodistrian University of Athens
Department of Informatics & Telecommunications
gthe@di.uoa.gr

Abstract. This PhD thesis introduces a low cost Software-Based Self-Test (SBST) fault detection methodology for small embedded cache memories. The methodology leverages existing powerful mechanisms of modern ISAs by utilizing instructions that we call in this PhD thesis Direct Cache Access (DCA) instructions. Moreover, our methodology exploits the native performance monitoring hardware and the trap handling mechanisms which are available in modern microprocessors. By effectively combining these features of modern microprocessors the proposed methodology applies March test operations with lower cost (code size, execution time, system performance overhead) when compared with other proposed solutions in the literature for fault detection in caches through SBST. Finally, a multithreaded optimization of the proposed methodology is also presented. The proposed methodology was applied to three processor benchmarks: a) OpenRISC 1200 b) LEON3 and c) OpenSPARC T1. Experimental results both for the test code size and test execution time of several March tests demonstrate the significant improvements in terms of test time (86% for instruction L1 cache, 87% for the data L1 cache, about 40% for D-TLB and about 82% for I-TLB) and test code size (83% for instruction L1 cache, 86% for the data L1 cache, 3% for D-TLB and 35% for I-TLB) when the methodology is applied to the same benchmarks (LEON3 for L1 caches and OpenSPARC T1 for TLBs) and such DCA instructions are exploited compared to SBST solutions that don't utilize these types of instructions.

Keywords: Microprocessor testing, Reliability, Software-based Self-test, On-line testing, Memory testing, March tests

1 Introduction

During the past 30 years the semiconductor industry has been characterized by a steady path of constantly shrinking transistor geometries and increasing chip size. However, this technology achievement leads to new reliability challenges for modern systems that have not been considered in the past. Such reliability threats are either

* Dissertation Advisor: Antonis Paschalis, Professor

latent hardware defects that have not been detected by manufacturing tests or hardware defects that may occur during system operation by the increased soft error rate or by aging degradation effects. On-line testing schemes aim to detect such faults both in logic and memory modules of modern chips during their lifetime. Nowadays, in modern processors the relative chip area occupation of cache devices is up to 90%. Thus, high quality cache memory on-line testing in modern processors is essential.

Small memories, such as L1 caches and Translation Lookaside Buffers (TLBs) are not usually equipped with Memory Built-In Self-Test (MBIST) hardware or lack a programmable MBIST scheme because of its impact on chip area and performance. Software-Based Self-Test (SBST) is a flexible and low-cost solution for on-line March test application and error detection in such small memories. Although, L1 caches and TLBs are small components, their reliable operation is crucial for the system performance due to the large penalties caused when L1 cache or TLB misses occur.

This PhD thesis introduces a low cost fault detection methodology for small embedded cache memories that is based on modern Instruction Set Architectures and is applied with SBST routines. The proposed methodology applies March tests through software to detect both storage faults [1] when applied to caches that comprise SRAM memories only, e.g. L1 caches, and comparison faults [2] when applied to caches that apart from SRAM memories comprise CAM memories too, e.g. TLBs. The proposed methodology can be applied to all three cache associativity organizations: direct mapped, set-associative and full-associative and it does not depend on the cache write policy. The methodology leverages existing powerful mechanisms of modern ISAs by utilizing instructions that we call in this PhD thesis Direct Cache Access (DCA) instructions. Moreover, our methodology exploits the native performance monitoring hardware and the trap handling mechanisms which are available in modern microprocessors. By effectively combining these features of modern microprocessors the proposed methodology applies March write and read operations with lower cost (code size, execution time, system performance overhead) when compared with other proposed solutions in the literature for fault detection in caches through SBST. Moreover, the proposed methodology applies March compare operations when needed (for CAM arrays) and verifies the test result with a compact response to comply with periodic on-line testing needs. Finally, a multithreaded optimization of the proposed methodology that targets multithreaded, multicore architectures is also presented in this thesis. The proposed multithreaded optimization exploits the thread level parallelism of multithreaded, multicore architectures and the low level multiple sub-bank organization of modern cache designs to speedup March tests while preserving the March test quality. Hence, in case of multithreaded, multicore architectures that can adopt the proposed multithreaded optimization, test time is further lowered and such SBST routines can be effectively executed periodically during the system's lifetime with an acceptable performance overhead.

The proposed methodology was applied to three processor benchmarks: a) OpenRISC 1200 b) LEON3 and c) OpenSPARC T1. In detail, the methodology was applied to the L1 caches of all three processor benchmarks and to the TLBs of OpenSPARC T1 processor. The multithreaded optimization was demonstrated on the multithreaded, multicore OpenSPARC T1. Experimental results both for the test code size and test execution time of several March tests demonstrate the effectiveness of

the proposed methodology, its high adaptability and the significant improvements in terms of test time and test code size when compared with other proposed solutions in the literature for fault detection in caches through SBST that do not exploit DCA instructions [9] [13].

Finally, a test evaluation framework was implemented in this thesis for several on-line periodic test scenarios in order to evaluate the system performance overhead of the proposed methodology. Simulation results indicate that the proposed March test implementation through SBST slightly influences the system's performance, even in intensive test scenarios with high test frequency requirements.

1.1 Related Work

Several SBST approaches have been proposed [3] – [13] to apply March tests to L1 caches but none for TLBs. Besides, none of the proposed SBST approaches in the literature exploits special instructions that can be considered as DCA instructions. In [3], the first systematic approach for transforming March B test algorithm for in-system testing of Intel 860 processor cache is proposed. In [4], a March test transformation methodology is proposed for in-system testing of both data and tag arrays of data and instruction cache memories with various organizations by taking advantage of features such as enable/disable or freeze. In [5], the authors propose a methodology to transform March tests for in-system testing targeting only the data cache memory tag without providing implementation details. The transformation methodology targets direct mapped caches and is applied on March B and March X tests. In [6], the methodology of [4] is enhanced to exploit microprocessor's performance monitoring hardware and on-line hardware detectors to improve test observability. In [7], a SBST approach is proposed to develop March-based self-test programs for the data array of both instruction and data caches. Experimental results for traditional memory faults are provided for a MIPS R3000 processor model for different March tests. In [8], March tests are translated in order to effectively test the data and tag part of set-associative caches with Least-Recently Used (LRU) replacement and both write-back and write-through policies. In [9], a software-based methodology is presented for testing memory arrays and logic modules of a direct-mapped data cache. Experimental results are provided for an ARM-compatible processor both for the cache arrays and the logic modules. In [10], a hybrid SBST approach is proposed to test data and instruction cache controllers by combining instruction-based pattern generation and an I-IP module insertion for observability. Experimental results for the cache controllers of OpenRISC 1200 processor are provided. In [11], a test program generation approach is proposed to generate suitable programs for testing the replacement logic in set-associative caches. Experimental results of a cache that implements the LRU policy are provided. In [12], the capabilities and limitations of CPU-based at-speed memory testing are presented based on test routine examples for an ATMEL RISC microcontroller. Such SBST routines can be also adopted for testing cache arrays. Recently in [13], a methodology to exploit the ISA to translate generic March tests into SBST programs for set-

associative instruction cache memories is presented and was applied on the instruction cache of the LEON3 microprocessor.

2 Cache arrays testability challenges

There are three basic cache organizations: direct mapped, set-associative and fully-associative. In processor design, data and instruction L1 caches are usually organized as direct mapped and set-associative caches whereas data and instruction TLBs are always organized as fully associative caches.

A typical L1 cache organization comprises of at least two SRAM memory arrays (or two SRAM arrays per way in set-associative organizations) - the data array and the tag array - whereas a fully-associative TLB organization comprises of one SRAM array - the data array - and one CAM array - the tag array. CAM is a special type of memory that compares all the stored data in parallel with incoming data and is utilized in the tag part of fully associative caches to speed up the tag comparison. Further down, those arrays will be denoted as DL1-Data, DL1-Tag, IL1-Data and IL1-Tag for the data and instruction L1 cache whereas for TLBs those arrays will be denoted as DTLB-Data, DTLB-Tag, ITLB-Data and ITLB-Tag for the data and instruction TLB, respectively.

All these cache arrays (either for L1 caches or TLBs) are implicitly accessed because they are not directly visible to the assembly programmer through the ISA. Hence, applying test patterns and observing the test responses through a software test routine is challenging. The challenges of accessing and thus testing those implicitly accessed cache arrays are summarized in Table 1 and thoroughly described in the thesis manuscript.

Table 1: Cache arrays testability challenges

Testability Challenges	***Cache arrays***							
	DL1 Data	***DL1 Tag***	***IL1 Data***	***IL1 Tag***	***DTLB Data***	***DTLB Tag***	***ITLB Data***	***ITLB Tag***
Direct access from generic ISA	×	×	×	×	×	×	×	×
Indirect March write	✓	✓	×	✓	×	×	×	×
Indirect March read	✓	×	×	×	×	×	×	×
Data Backgrounds	✓	×	×	×	×	×	×	×
Ascending Address Order	✓	✓	✓	✓	×	×	×	×
Descending Address Order	✓	✓	×	×	×	×	×	×
March Compare operation	-	-	-	-	-	×	-	×

The proposed methodology overcomes all these testability challenges for all cache arrays of both L1 caches and TLBs and optimizes the SBST routines in terms of test time and test code size by exploiting DCA instructions.

3 DCA instructions in modern ISAs

So far, previous SBST approaches cannot successfully overcome all the above mentioned challenges by using generic instructions to access cache arrays both for write and read operations. Fortunately, modern ISAs include dedicated instructions for debug-diagnostic and performance purposes that provide direct controllability and observability of cache arrays. These instructions are extremely suitable for cache and TLB SBST; we use the term Direct Cache Access (DCA) instructions to refer to them.

In order to gain direct access to the cache arrays for all three cache organizations (direct mapped, set-associative and fully-associative) and implement a software-based March test, an ideal DCA instruction needs to contain the following fields:

Fields for selecting cache/TLB content:

- Way Selection (WS) field.
- Set Selection (SS) field.
- Line Word Selection (LWS) field.

Field for selecting internal cache array:

- Data/Tag Array Selection (AS) field.

Field for selecting March operation:

- Write/Read/Compare operation (WRC) field

Field for addressing register/memory for fetching DBs:

- From/To data Address (A) field

An ideal DCA instruction that contains all these fields gains direct access to any cache array, hence it can apply March operations through ISA to these arrays in a very effective way and overcome all the above described testability challenges. DCAs that access direct mapped caches should contain only fields *SS* and *LWS* while DCAs for accessing set associative L1 caches should contain all three *WS*, *SS* and *LWS* fields. DCAs that access fully associative caches and TLBs should contain only *WS* and *LWS* fields (fully associative caches contain 1 set with many ways) respectively. When the cache organization imposes a uniform cache line (e.g. TLBs) *LWS* field is not required. Finally, if the cache organization does not comprise a CAM memory (e.g. L1 caches), the March operation selection field can be renamed to *WR* field (only write/read operations, no compare).

In Figure 1 and Figure 2, ideal DCA instructions and the way that every field is utilized to access a 2-way set associative L1 cache and a TLB, are presented, respectively.

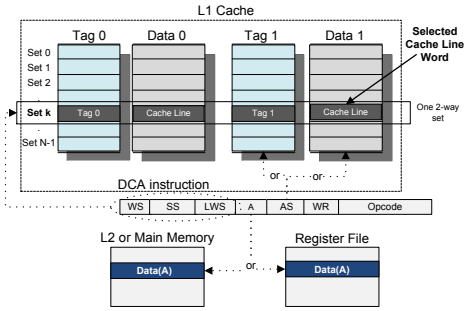

Figure 1: **Ideal DCA instruction for 2-way set associative L1 cache**

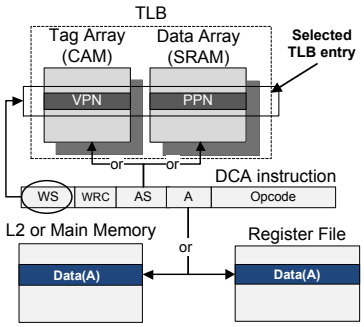

Figure 2: **Ideal DCA instruction for TLB**

In detail, in L1 set-associative caches the data or tag array (selected by the *AS* field) is accessed both for write and read operation (selected by the *WR* field). The selection of a word inside a cache line is controlled in three steps. First, the *SS* field selects the set, then, the *WS* field selects the cache way and finally the *LWS* field selects the word inside the cache line. Furthermore, all DBs can be composed by initializing either a general purpose register or a memory location that can be accessed by the *A* field. In TLBs, the data or tag array of the TLBs can be accessed by controlling the *AS* field. The March operation can be selected with the *WRC* field to write, read or compare a selected TLB entry. Compare operation is valid only for tag array. TLBs are fully associative arrays; hence the *WS* field is needed in the ideal DCA instruction to gain access to every TLB entry. Note that, such an instruction has no limitation to access a cache either in ascending or in descending order.

In practice, such an ideal DCA instruction does not exist in ISAs but it can be indirectly implemented by combining a set of existing DCA instructions that together totally cover all fields of the ideal instruction. Representative examples of such special purpose instructions, which can be considered as DCAs, are present in RISC architectures, such as MIPS, ARM and SPARC architectures and in CISC architectures such as x86 architectures. A detailed presentation of these existing DCA instructions is included in the thesis manuscript.

4 SBST methodology for small caches

The methodology targets all cache arrays for both data and instructions (either L1 caches or TLBs) and is suitable for all three cache organizations with any write policy. An SBST technique that targets L1 caches cannot be cache resident, since the actual L1 cache is under test. However, this is not a limitation in case of on-line testing, since the test routines can be stored and executed from either L2 cache or the chip's main memory that is available at test time. Moreover, when the SBST methodology targets TLBs, the SBST routine should be placed in a non-pageable memory location that is not cached to the instruction TLB since the actual TLB arrays are under test.

The proposed SBST methodology implements low cost SBST March tests that target cache arrays by taking advantage of existing debug-diagnostic instructions in modern ISAs, as described in Section 3. These instructions must cover in total the fields of the ideal DCA instruction to overcome the testability challenges of the cache arrays. The main feature of the methodology is low cost March write/read implementation due to high controllability and observability of DCA instructions. The proposed SBST methodology is briefly summarized in Figure 3 for L1 caches and TLBs respectively and is thoroughly presented in the thesis manuscript. The main features of the proposed methodology are:

- Low cost March writes due to the high controllability of DCA write instructions.
- Low cost March reads due to the high observability of DCA read instructions.
- Low cost March comparisons due to the special features of DCA compare instructions.
- Low cost March reads for tag arrays by exploiting performance monitoring hardware, if DCA read instructions do not exist in the ISA.
- Low cost March comparisons for TLB tag arrays by exploiting the trap handler mechanism, if DCA compare instructions do not exist in the ISA.
- Test response compaction to comply with on-line testing requirements.

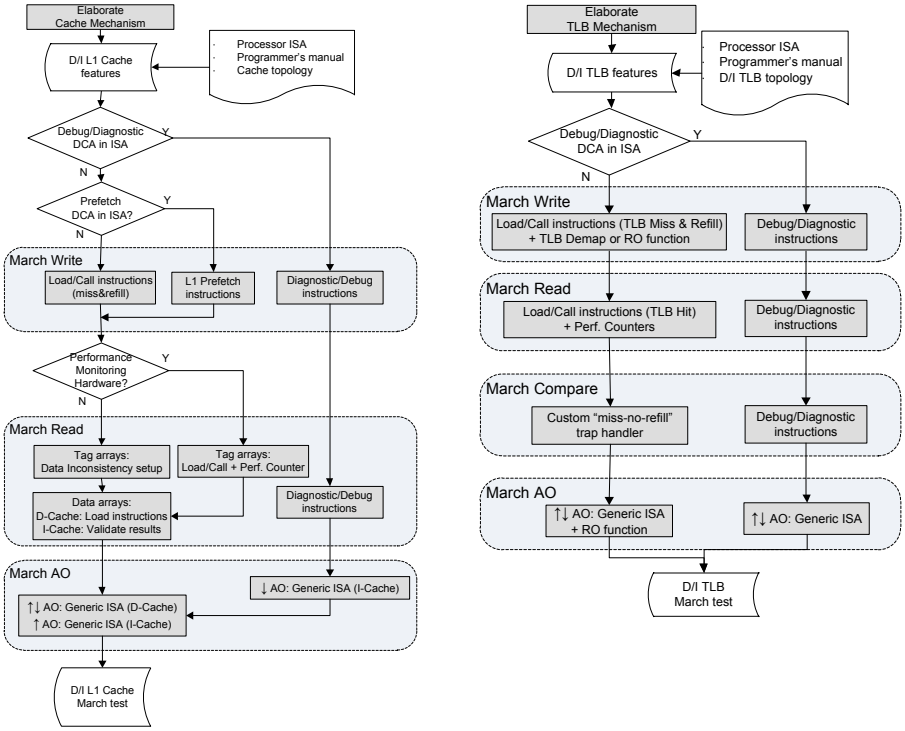

Figure 3: SBST Methodology for a) L1 caches b) TLBs

5 Multithreaded optimization for multi-bank L1 caches

We propose a multithreaded optimization of the SBST methodology that elaborates the low level multiple bank organization of modern cache designs to exploit the thread level parallelism of modern multithreaded multicore processors and speedup March test execution time.

In cache designs, SRAM arrays are partitioned into sub-banks to save power or to tweak the cache dimensions or even multiple banks to minimize internal wiring delay. Such cache designs that consist of one bank but multiple sub-banks are called Uniform Cache Architectures (UCA) whereas these that that consist of both multiple bank and sub-banks are called Static/Dynamic Non-Uniform Cache Architectures (S-NUCA/ D-NUCA) based on the way that the mapping of data into cache banks is achieved [12].

In UCA and S-NUCA architectures, the mapping of data into cache banks and sub-banks is predetermined based on the block index of the architecture and thus can reside in only one bank of the cache. Several cache modeling tools (e.g. Cacti) enable fast exploration of the cache design space by automatically choosing the optimal bank

and sub-bank count, size, and orientation of UCA and S-NUCA architectures. Each sub-bank has a separate memory array, decoder, write drivers, and sense amplifier. The S-NUCA cache consists of multiple banks. Every bank is organized as a UCA with multiple sub-banks.

Each cache sub-bank in UCA and S-NUCA multibank architectures can be considered as a separate SRAM array with a distinct functional fault set, since no coupling faults can occur between memory cells of different cache banks and sub-banks. Therefore, an on-line test strategy that considers every sub-bank as an independent memory array can be developed since the memory mapping is known to test engineers.

The proposed optimization considers exclusive fault sets for every L1 cache array sub-bank and proposes a fine-tuned clustering of the applied March tests in smaller subroutines based on the information provided for the physical implementation of a multibank cache (sub-bank address mapping, address scrambling etc.). Every cluster targets a sub-bank of the cache and the logical union of all clusters ends up to the initial March test for the whole cache array under test. These March test clusters target separate cache array sub-banks which can be executed concurrently without diminishing the March test effectiveness. Such a clustering is well suited to multithreaded processors, where concurrent execution can be achieved by assigning test clusters to different threads and hence SBST March test execution time is divided by the factor of available number of test threads. The test threads are dynamically scheduled software threads among the other executed processes. These test threads should be isolated during on-line testing in order to prevent the rest of the software processes to corrupt the March test effectiveness. In case that the March tests clusters outnumber the available threads, more than one March test clusters are assigned to every thread. Finally, the proposed multithreaded optimization is suitable both for simultaneous multithreading and interleaved multithreaded architectures, since it is independent of the way that the threads are issued and it is compatible with any resource allocation policies (e.g. physical register file size, register windows, register renaming e.t.c).

6 Case studies: LEON3, OpenRISC 1200, OpenSPARC T1

We have applied the proposed SBST methodology to three processor benchmarks: a) LEON3, b) OpenRISC 1200 and c) OpenSPARC T1. We have also applied the multithreaded optimization to the L1 caches of OpenSparc T1.
In order to evaluate the effectiveness of the self-test routines we have used RAMSES memory fault simulator [14].

The first benchmark is LEON3, a publicly available processor designed by Aeroflex Gaisler that implements a SPARC V8 compliant architecture. The SPARC V8 ISA, that LEON3 implements, includes privileged store/load instructions, denoted as alternate load/store (*lda*/*sta* instructions). These instructions can directly access cache arrays for diagnostic purposes by specifying alternate space identifiers (ASIs) that are defined by the SPARC architecture for both write and read access at

supervisor level. These instructions have been used as DCA instructions for March write/read operations to apply and read the test patterns in SBST routines.

The second benchmark is OpenRISC 1200, a publicly available processor core. OpenRISC 1200 lacks debug-diagnostic instructions in its ISA to access the cache arrays. However, it includes a cache prefetch mechanism for both data and instruction L1 caches and maps prefetch operations to valid instructions. These instructions have been utilized as DCA instructions for March write operations. For March read operations, generic load and call instructions have been used. The observability of the March tests has been improved by exploiting the performance counters.

The third benchmark is OpenSPARC T1 that includes both data and instruction L1 caches and fully functional TLBs. Therefore, the SBST methodology has been fully applied to both the L1 cache and TLB arrays. OpenSPARC T1 implements a SPARC V9 compliant ISA and includes privileged store/load instructions, denoted as alternate load/store (*ldxa/stxa* instructions). We have exploited these alternate load/store instructions for March write/read operations to directly access all cache arrays for both L1 caches and TLBs for March write and March read operations at low cost by utilizing the appropriate ASI at the hypervisor level. SPARC V9 ISA does not implement a debug/diagnostic compare instruction for implementing the March compare operation. Therefore, a custom "miss-no-refill" trap handler has been utilized to implement March compare operations to test DTLB-Tag and ITLB-Tag CAM arrays for comparison faults. We have also applied the multithread optimization to the March tests that target the four SRAM arrays of both data and instruction L1 caches of a SPARC core.

We have applied a set of March tests with different test complexities to both data and instruction caches. Solid data backgrounds (all-zero/all-ones) have been used to all tests. The test quality is verified with RAMSES for all the applied March tests and 100% fault coverage has been verified for all cache arrays.

Detailed experimental results and comparisons both for the test code size and test execution time of these March tests are provided in the thesis manuscript and demonstrate the effectiveness of the proposed methodology, its high adaptability and the significant improvements in terms of test time (86% for instruction L1 cache, 87% for the data L1 cache, about 40% for D-TLB and about 82% for I-TLB) and test code size (83% for instruction L1 cache, 86% for the data L1 cache, 3% for D-TLB and 35% for I-TLB) when the methodology is applied to the same benchmarks (LEON3 for L1 caches and OpenSPARC T1 for TLBs) and such DCA instructions are exploited compared to SBST solutions that don't utilize these types of instructions. Moreover, experimental results show a speedup of more than 1.7 (for two threads) and more than 3.7 (for four threads) in test time when the proposed multithreaded optimization is applied to the L1 caches of OpenSPARC T1.

7 Performance overhead evaluation

In this section, we present the evaluation framework that was utilized to estimate the performance overhead of the proposed SBST routines. We have implemented

several on-line periodic testing scenarios and we will present detailed statistics of the performance overhead introduced in a typical workload under these test scenarios.

We have utilized a SunFire T2000 server running a set of multithreaded programs -the PARSEC benchmark suite- over Solaris 10 to evaluate the performance overhead of the deployed SBST routines. Our server is powered by a quad-core UltraSPARC T1 processor. OpenSPARC T1 processor is the free version of UltraSPARC T1 that is utilized in SunFire T2000 servers. Hence, the SBST routines that were developed above for OpenSPARC T1 processor can be directly compiled to our server to evaluate their performance overhead.

We have selected the optimized 2-thread March C- SBST routine that targets the data L1 cache (both DL1-Tag and DL1-Data) to be utilized as our self-test routine in the evaluation framework. The test time of this test routine has been measured about 1.2sec in our system. Any other SBST routine (or a set of March tests) could have been selected. Since we do not have access on the hypervisor level on the UltraSPARC T1 processor of a native system, we have slightly altered the SBST test routines in order to comply with Solaris OS limitations on executing hyper privileged instructions.

The modified SBST routines have the same memory footprint and test execution time with the OpernSPARC's T1 implemented routines, thus the modified self-test routine is sufficient for studying the performance impact of the proposed on-line self-tests.

Here after, we will utilize the terms of Test Period (TP) and Test Latency (TL) as described below:

- Test period (TP) and is the amount of time from the beginning of a self-test on a core to the beginning of the next self-test on the same core.
- Test latency (TL) is the duration of an on-line self-test.

We have applied several on-line testing scenarios with a fixed TL (1.2sec) and several short TPs (< 1min) that are suitable for detecting early-life failures on two different framework configurations, a *1core/4thread* setup and a *4core/16thread* setup as described below in detail.

- ***1core/4threads setup, TL=1.2sec, TP=2, 15 and 60 sec***

Firstly, Solaris capability of creating virtual processor sets has been exploited to isolate a single SPARC core from OpenSPARC T1 and both PARSEC workload applications and SBST routine have been set -by utilizing Light Weight Process (LWP) binding- to be executed in this SPARC core.

Note that we have selected to isolate an idle core that does not execute any other OS process in order to evaluate the real performance overhead due to the SBST routine's periodic execution only. The PARSEC suite has been configured to be executed by four threads (the maximum number of threads in the core). These four threads share the same L1 cache. Hence, both workload and March test application access the same cache SRAM arrays. Afterwards, several TPs have been selected to represent different test scenarios. For example a demanding testing scenario may require intensive test period (e.g. TP=2sec) while a more relaxed test scenario may require less intensive test periods (e.g. TP=60sec).

- ***4cores/16threads setup, TL=1.2sec, TP=10, 30 and 60 sec***

In this configuration we have utilized all four available SPARC cores of our server. Hence, apart from the PARSEC applications and the periodic execution of the SBST routine, the OS processes were also executed in the background. The PARSEC suite has been configured to be executed by all the available sixteen threads of our system (four threads per SPARC core). A script has been composed to call the SBST routine for every SPARC core in a round-robin way in every TP. Moreover, in each SBST routine execution, the script forms a virtual processor set of the 4 threads that belong to the core under test to ensure that the SBST routine will be executed only by the selected core under test. This is a critical requirement to guarantee the test quality by preventing the test patterns to be stored in L1 cache of other cores, apart from the core under test. Several TPs have been selected to represent different test scenarios, in a similar way that was presented for the 1core/4threads setup. The selected TPs were longer in these case studies due to the need of executing the same SBST routine four times (one for the L1 cache of each SPARC core).

The system in both configuration setups was configured to execute all the 13 multithreaded programs of the PARSEC suite. All PARSEC programs have been compiled to utilize the pthreads parallelization model and the native dataset was utilized in all simulations. After estimating the workload execution time in both configuration setups without test, the PARSEC programs were executed several times while the SBST routine was scheduled to be executed in the background at a fixed TP in every test scenario for both configurations.

Detailed statistics for the performance overhead for each test scenario are provided in the thesis manuscript and even in the more demanding test scenario of a quad core processor does not exceed the 11% of the performance of the system without test. Considering that this performance penalty refers to a demanding on-line periodic self-test scenario that applies March C- algorithm to both D-Tag and D-Data SRAM arrays for all four L1 caches of the quad core system every 10sec, such performance degradation can be affordable when a strict test scenario is required. In contrary, in a relaxed test scenario (e.g. not more than a test per minute), the performance overhead is lower. For example, when the SBST routine is periodically executed every minute in a single core system (e.g. a single core system can be an embedded processor), the performance overhead is less than 1%, thus negligible. Thus, the proposed SBST methodology can periodically apply March tests to L1 caches effectively during the system's lifetime with acceptable performance overhead in workload execution.

Conclusions

We have presented a low cost SBST fault detection methodology for small embedded cache memories that leverages existing powerful mechanisms of modern ISAs by utilizing existing DCA instructions, exploits the native performance monitoring hardware and the trap handling mechanisms which are available in modern microprocessors. Furthermore, a multithreaded optimization of the proposed methodology has been presented. The methodology has been applied to three processor benchmarks: a) OpenRISC 1200 b) LEON3 and c) OpenSPARC T1.

Experimental results for several March test implementations demonstrate the significant improvements in terms of test time and test code size when the methodology is applied to the same benchmarks (LEON3 for L1 caches and OpenSPARC T1 for TLBs) and such DCA instructions are exploited compared to SBST solutions that don't utilize these types of instructions. Finally, a performance evaluation framework has been presented to demonstrate that the implemented SBST routines have acceptable performance overhead for periodic on-line testing.

References

1. S. Hamdioui et al, "March SS: A Test for All Static Simple RAM Faults", in Proc. of IEEE Int'l Workshop Memory Technology, Design and Testing (MTDT), 2002
2. K.J. Lin, C.W. Wu, "Testing content-addressable memories using functional fault models and march-like algorithms", IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, Vol.19, no.5, pp.577-588, May 2000
3. A.J. van de Goor, T.J.W Verhallen, "Functional testing of current microprocessors", in Proc. of International Test Conference (ITC), 1992, pp.684-695.
4. J. Sosnowski, "In-system testing of cache memories", in Proc. of International Test Conference (ITC), 1995, pp.384-393.
5. S.M. Al-Harbi et al. "A Methodology for Transforming Memory Tests for In-System Testing of Direct Mapped Cache Tags", in Proc. of VLSI Test Symposium (VTS), 1998.
6. J. Sosnowski, "Improving Software Based Self-Testing for Cache Memories," in Proc. of Design and Test Workshop (DTW), 2007
7. M. Tuna et al. "Software-Based Self-Test Strategies for Memory Caches of RISC Processor Cores", in Proc. of IEEE Latin-American Test Workshop, 2007, pp.124-130.
8. S. Alpe et al., "Applying March Tests to K-Way Set-Associative Cache Memories", in Proc. of the 13th European Test Symposium (ETS), 2008, pp.77-83.
9. Y.C. Lin, Y.Y. Tsai, K.J. Lee, C.W. Yen, C.H. Chen, "A Software-Based Test Methodology for Direct Mapped Data Cache", in Proc. of Asian Test Symposium, 2008
10. W.J. Perez et al., "A Hybrid Approach to the Test of Cache Memory Controllers Embedded in SoCs", in Proc. of IEEE International On-Line Testing Symposium (IOLTS), 2008, pp.143-148.
11. W.J. Perez et al., "On the Generation of Functional Test Programs for the Cache Replacement Logic", in Proc. of IEEE Asian Test Symposium (ATS), 2009, pp.418-423.
12. A.J. van de Goor et al., "Memory testing with a RISC microcontroller", in Proc. of Design, Automation & Test in Europe Conference & Exhibition (DATE), 2010
13. S. Di Carlo, P. Prinetto, A. Savino, "Software-Based Self-Test of Set-Associative Cache Memories," IEEE Transactions on Computers, Vol.60, no.7, pp.1030-1044, July 2011
14. C.F. Wu et al, "RAMSES: a fast memory fault simulator", in Proc. of the International Symposium on Defect and Fault Tolerance in VLSI Systems, 1999, pp.165-173

Resource management in *IaaS*-Clouds

Konstantinos Tsakalozos*

National and Kapodistrian University of Athens
Department of Infromatics and Telecommunications
k.tsakalozos@di.uoa.gr

Abstract. In this thesis, we present our work on four fundamental problems current *IaaS* clouds face: *a)* How consumers can assist the cloud middleware in resource management while the internal structure of the cloud is never revealed. *b)* The *VM*-to-*PM* placement in large non-homogeneous physical infrastructures. *c)* Load balancing though live *VM* migration under time constraints while not depleting the cloud resources. *d)* The maximization of the financial profit when using the cloud. The proposed solutions on these problems show significant improvements in the efficiency of *IaaS* clouds.

1 Introduction

In our everyday interaction with computing systems, we regularly use services that are classified as cloud enabled. Most of these cloud services are made available through the *Internet*, while others are restricted for use within organizations. The large variety and versatility of cloud services may discourage us from finding common properties among them, yet, they are all hosted by the "misty" cloud.

Concealing services inside clouds allows us to concentrate on how consumers are expected to use them. Consumers are kept agnostic to the details of *how* services are provided. In effect, users focus on what really matters to them, consuming a service. Similarly the cloud service providers focus only on aspects of their domain that are largely transparent to the end consumers. This separation of concern introduces a very elegant and convenient way for building services and interacting with them. It essentially sets the interface between service providers and service consumers.

Users contacting an *IaaS* cloud often request multiple *VM*s that collectively form an entire virtual infrastructure. Each *VM* is coupled with specific service level agreements (SLAs) regarding its availability and performance. Often these SLAs involve absolute values when they refer to space shared resources (e.g., RAM) and average values for time shared resources (e.g., network bandwidth). The reason for this is that due to the multi-tenancy, time-shared resources are offered to multiple users simultaneously and they are often over-committed in order to increase hardware efficiency. Most commercial cloud providers charge *VM* access depending on system characteristics and the SLAs offered.

* Dissertation Advisor: Alex Delis, Professor.

Internally, *IaaS* clouds produce *VM*s hosted on physical nodes. The dedicated software that undertakes the task of serving a user request is called *Cloud Middleware* [1–4]. The cloud middleware decides which physical machine (*PM*) is going to host the requested *VM*. All *PM*s are equipped with software capable of virtualizing hardware, that is made available to the customers. The cloud middleware also configures the network so that the provided *VM*s operate in isolation. The end result is that the user has a service that provides access to processing power upon request.

The cloud middleware monitors the overall performance of the cloud and may help facilitate the sharing of load across physical nodes. Hot nodes are cooled down by migrating *VM*s away from them. *VM* migration essentially changes the source of resources consumed by the *VM*. Migrations also serve administrative tasks that require certain parts of the physical infrastructure to be shutdown for maintenance reasons.

In this thesis, address four fundamental problems current *IaaS* clouds face: First, we show *how* consumers may assist the cloud middleware in its decision on where to place the *VM*s requested by the user. The challenge in this is that the internal structure of the cloud should never be revealed to the consumer. Our approach is based on user-provided hints that describe an ideal placement of the *VM*s to *PM*s. The cloud middleware tries to reach a *VM* placement that resembles the ideal. The second problem we discuss is how the *VM*-to-*PM* placement decision can take into account the heterogeneity of large physical infrastructures. The third problem is how load balancing though *VM* migration can take place under time constraints. We address this issue by combining a special purpose file-system with a high level resource allocation policy that oversees the progress of simultaneous *VM* migration tasks. Our goal is to migrate *VM*s under time-constraints while not completely depleting the cloud resources. The final problem we discuss is how to maximize the financial profit of a cloud whose actual performance may change overtime due to events the consumer is not aware of (e.g., sharing cloud resources with other consumers).

2 Related Work

The work in this thesis targets four distinct areas of *IaaS* clouds.

2.1 *VM* Placement Based On Constraints

The *VM*s-to-*PM*s placement policy in [5] exploits the tendency of *VM*s to have certain properties in common. [6] reformulates the problem as a multi-unit combinatorial auction. In [7], placement constraints are treated as separate dimensions in a multi-dimensional *Knapsack* problem. User and administrative preferences expressed through constraints are also employed in [8–11]. Often, as the number of constraints increases more resources are needed to solve the constraint satisfaction problem.

2.2 Workloads and System Consolidation

The allocation of resources in dynamic distributed environments [12] where load and resource availability change over time requires adaptive policies. In [13, 14], such resource sharing policies are proposed for the execution of jobs on the Grid.

Sandpiper [15] detects and monitors performance bottlenecks in a cluster hosting *VM*s. Two approaches are evaluated in the decision making mechanism that produces the *VM* migration actions: the first, termed *black-box*, remains fully OS-and-applications agnostic while the second, termed *gray-box*, exploits statistics originating from both the OS and the application-layer. Modeling the *VM* load is deemed important in the *black-box* approach used in IaaS clouds. [16–18] classify *VM* workloads and develop metrics to model the encountered workloads in an effort to reduce *VM* migration costs.

2.3 System Migration

Virtualization used in *IaaS*-clouds offers a simple, yet, powerful solution to load balancing. "Hot" *PM*s can be cooled down by having *VM*s moved elsewhere. Live migration [19] could potentially reduce the downtime of migrating *VM*s to the scale of milliseconds. Current *VM* Monitors (*VMM*s) [20–22] place a hard requirement for live migration as both the source and target *PM*s must share a storage layer. To address this requirement, small and medium-size clouds resort to the use of distributed file systems (*DFS*s). Red Hat's Global File System (*GFS*) [23] allows for several nodes to combine their storage resources. Each client accessing *GFS* must first acquire appropriate locks by contacting a set of master nodes. *GlusterFS* [24] is a user-space *DFS* in which every node can act both as a storage repository and/or a client. *Ceph* [25] uses separation between metadata and storage nodes and although scales well, its efficiency is hampered by saturated remote components such as switches.

An alternative proposal to the scalability issues of *DFS*s is the use of incremental transfer of virtual systems to different locations [26]. In [27, 28], live *VM* migration in WANs is advocated as a way to relieve overloaded physical nodes. In such environments, migration must handle challenges emanating from rerouting, significant latencies and low bandwidth rates. A workload-driven migration of virtual storage is discussed in [29].

2.4 The Finance of Highly Scalable, Elastic Infrastructures

Often in commodity selling markets, user demand reduces as resource price increases. This kind of user behavior is assumed to study the conditions that maximize the resource provider's profit [30]. Markets with more than one type of commodity (such as CPU, network, storage) are well-suited for use in multi-user computing clusters. Here, policies set the price of all sold resources considering their correlations [31]. Policies that take into account only high-level performance metrics, such as the application response time [32], seem more appealing to the end users since less intervention is required on their part.

By combining clouds with high resource availability [33] and a massively scalable programming framework [34], a user can process large amounts of data. *Dynamic Hadoop* [35] and Condor [36] can harvest the resources of large cloud installations and offer virtual elastic infrastructures on demand. These frameworks add a software layer that efficiently manages administrative concerns of scaling virtual infrastructures.

The work in this thesis addresses problems in the aforementioned areas and greatly improves the efficiency and features of *IaaS* clouds. The instantiation of complex virtual infrastructures is assisted by enhancements in the interaction between the cloud provider and customers [11, 37]. The latter are able to make use of cloud internal properties, such as high availability features, that are normally known only to the cloud administration. Through *VM* deployment hints the customers offer advice to the cloud's resource management, while the physical substrate remains concealed. Taking into account the provided hints, the cloud middleware can make intelligent resource allocation decisions not only during the initial instantiation of *VM*s, but also throughout the *VM*'s life-cycle. Well-informed resource appointment is displayed when the need for balancing load or maintenance tasks rises. Our work allows for the re-organization of virtual infrastructures, through *VM* migrations, under time-constraints [38]. By exploiting this feature the cloud middleware can schedule migrations in periods of low demand, thus customers can benefit from peak performance and reduced SLA failure rates. Furthermore, part of the enhanced interaction between providers and consumers is the cost efficiency of elastic infrastructures. We show how the elasticity of virtual infrastructures can be tuned to match the budget of their owners. As a result, we maximize the profit entailed in consuming cloud resources [39].

Since the success of cloud installations is based on economies of scale, we have put significant effort in promoting scalability. Our approach in matching the user needs with the available resources is computationally demanding. Yet, its design takes advantage of the distributed nature of the cloud and keeps up with the growth of the physical infrastructure [40]. In similar spirit, the *VM* migration facilities we propose, offer live migration in a true share nothing cloud architecture; thus we impose no limits to the size of the physical substrate. Overall, the scalability, our work offers, combined with the profit maximization methods we propose, can greatly benefit huge-data applications setup on the cloud.

3 Hint-based Execution of Workloads in Clouds with Nefeli

In our thesis, we present the design, implementation, and evaluation of a *cloud gateway*, *Nefeli*. *Nefeli* performs intelligent placement of *VM*s onto physical nodes by exploiting user-provided deployment *hints*. Hints realize placement preferences based on knowledge only the cloud consumer has regarding the intended

usage of the requested *VM*s. By modeling workloads as patterns of data flows, computations, control/synchronization points and necessary network connections, users can identify favorable *VM* layouts. These layouts translate to deployment hints. Such hints articulate *1)* resource consumption patterns among *VM*s, *2) VM*s that may become a performance bottleneck and *3)* portions of the requested virtual infrastructure that can be assisted by the existence of special hardware support. For instance, the fact that two *VM*s in a virtual infrastructure will hold mirrors of a database is only known to the cloud consumer. This information should be communicated to the cloud as a deployment hint so that the respective *VM*s will not be deployed on the same host. We refer to *VM* layout patterns as *task-flows* to distinguish them from the traditional workflow concept. Specifically, task-flows illustrate "ideal" deployments of *VM*s described by the cloud consumers using deployment hints. *Nefeli* exploits these hints so as to (re)deploy *VM*s in the cloud and achieve efficient task-flow execution. However, hints must not reveal any cloud internal properties to the consumers. Although hints may offer a desired *VM*-deployment for consumer workloads, *Nefeli* may ultimately elect to ignore part or all of them based on the available physical resources. In addition to hints, *Nefeli* also takes into account high-level *VM* placement policies, set by the cloud administration, whose objectives may entail energy efficiency and load balancing.

The main contribution of our approach is that we present a complete solution in extracting and exploiting the knowledge cloud consumers posses regarding the operational aspects of their virtual infrastructures. Our approach is compatible with the cloud abstractions that dictate users are kept agnostic of the physical infrastructure properties at all times. Furthermore, our approach is able to adapt to dynamic environments where both task-flows and user preferences change over time. *Nefeli* produces suitable *VM* to physical node mappings in response to signals coming from the infrastructures (both physical and virtual) or any other external notification mechanism. The produced mappings are applied through appropriate *VM* placement calls to an underlying cloud middleware.

We have created a detailed prototype and experimented with both simulated and real cloud environments. We compare *Nefeli VM*-placement against *a)* random placement, *b)* a placement that evenly distributes *VM*s among physical nodes, *c)* a policy that minimizes the number of physical nodes used and thus reduces the power footprint of the cloud, and *d)* the match making policy used by the open–source cloud middleware *OpenNebula v.1.2.0*. Our approach consistently displays significant performance improvements when compared to the aforementioned policies. In video transcoding, *Nefeli* achieves 17% reduced processing times compared to the *VM* placement decided by *OpenNebula*. In scientific task-flows and for a variety of simulated clouds, *Nefeli* demonstrates significantly higher throughput rates compared to other *VM* placement policies. Noteworthy savings in terms of power consumption are reported as well. We also present the performance overheads involved in the operation of *Nefeli* as the cloud infrastructure scales out.

4 Handling Non-homgenious Clouds

Infrastructure-as-a-Service cloud providers often face the following challenge: they must offer uniform access (resource provision) over a non-uniform hardware infrastructure. Non-homogeneous infrastructures may be the product of hardware upgrades, where old resources are left operational alongside new ones, or federated environments, where several parties are willing to share hardware resources with diverse characteristics.

In our work, we focus on the problem of instantiating entire virtual infrastructures in large non-homogeneous IaaS clouds. We introduce a service implementing a two-phase mechanism. In the first phase, we synthesize infrastructures out of existing promising physical machines (*PM*s). These dynamically-formed physical infrastructures, termed *cohorts*, host the user-requested *VM*s. In the second phase, we determine the final *VM*-to-*PM* mapping considering all low-level constraints arising from the particular user requests and special characteristics of the most promising selected *cohorts*. Compared to other constraint-based *VM* scheduling systems [8–10], the novelty of our approach mainly lies in the first phase. During this phase, besides resource availability, we also take into account properties such as migration capabilities, network bandwidth connectivity, and user-provided deployment hints. This helps prune out many possible *cohorts* within the cloud, and thus reduces the time required to produce a deployment plan in the second phase. We express both the selection of hosting nodes and the production of *VM*-to-*PM* mappings as constraint satisfaction problems and we use cloud-resources to solve these problems. We harvest these resources with the help of an elastic virtual infrastructure that can be dynamically enhanced with additional nodes depending on the needs of cloud users and the quality of *VM*-to-*PM* mappings cloud administrators want to provide. Our evaluation shows that this approach *1)* scales effectively for hundreds of *PM*s, *2)* reduces plan production time by up to a factor of 9, and *3)* improves plan quality by up to a factor of 4, when compared to a single-phase *VM* placement approach.

5 Time Constrained Live *VM* Migration in Share-Nothing *IaaS*-Clouds

Live migration requires that both the source and target *PM*s (used for hosting the *VM*) have access to the virtual disks of the *VM*. There are three alternatives in achieving this:

Single storage node: having a single common storage node that retains all virtual disk images is often impractical. This directly impacts the cloud scalability. Furthermore, such "single-node" solutions tend to become a performance bottleneck as well as a point of failure.

Distributed file systems: to tackle inefficiencies of a common storage node, clouds may employ distributed file systems such as *GFS* [23] and *GlusterFS* [24]. The latter allow for several *PM*s to collectively form a persistence layer, and thus exchange *VM*s and share load. The key objective of

distributed file systems is to present the *same* view of *all* files to *all PM*s at *all* times. Even though such file systems employ data replication and caching mechanisms they are often unable to comply with the requirements of large cloud installations regarding a *POSIX* API and low I/O latency.

Synchronization of virtual disks: as *VM*s could benefit from the exploitation of local resources including low-latency access to data and availability of RAID arrays, migration of *virtual disks on-demand* is an attractive choice to help re-arrange *VM*s [26, 28]. This approach is very different from that of distributed file systems as it attempts to place *VM*s on *PM*s on-the-fly implementing a scheduling policy of choice. On-demand synchronization entails the *VM* disk images that are being migrated, a target *PM*, and how the actual migration happens [27]. As there are no restrictions regarding the sizes (often multiple Gigabytes) of *VM*s to move, there are always significant overheads that have to be weighted against benefits to be reaped in the long term.

The feasibility of the on-demand synchronization has been established in prior efforts [26–28] and we use it as the foundation of our proposal. In our approach, we employ on-demand *virtual disk synchronization* to accommodate large numbers of migration tasks -involved in operations such as load balancing- across large cloud installations. The novelty of our approach is in the policy appointing resources to migrations. Our objective is to complete each migration within a designated time frame while not depleting the resources of the cloud. In effect, we attempt to achieve real-time load balancing and reduce the performance penalty that migrations inflict. In our approach, each migration task is paired with a *time-constraint* regarding its completion. By complying with their designated time-constraints, *VM*s do not display degraded performance due to migration in periods of high utilization. Our approach employs resource management features that help effectively synchronize virtual disk images across *PM*s. These low-level features predominantly deal with the consumption of *PM*-resources and are realized in the context of our *MigrateFS* file system. Instances of *MigrateFS* communicate over the network for moving virtual disk images between any two *PM*s. During the shipment, we are allowed to set: *a)* the disk bandwidth available to the *VM* internal processes accessing the virtual disks under migration, and *b)* the network bandwidth used for the purposes of migration. In this way, we are able to yield safe estimates on the exact time the migration will finish.

Our approach is weaved around a coordinating *Migrations Scheduler* and a distributed *network of Brokers*. The *scheduler* prioritizes migration tasks while the *Brokers* monitor the progress and appoint resources to each migration. The end-effect is that migrations finish within specified time-limits and "hot" physical network links are not further stressed by virtual disk shipments. The latter is achieved through constraints set by *MigrateFS* on the network bandwidth. Our evaluation, based on both a *MigrateFS* prototype and simulation of large infrastructures, shows up to 24% less stress on saturated *PM*s during migration at the expense of minimal administration effort.

6 Flexible Use of Cloud Resources through Profit Maximization and Price Discrimination

Cloud computing introduces a number of challenges regarding both performance and financial issues. On the one hand, consumers of cloud services try to minimize the execution time of their submitted tasks without exceeding a given budget and on the other, cloud providers are keen on maximizing their financial gain while keeping their customers satisfied. With this work we focus on virtual infrastructures following the master-worker paradigm hosted in a cloud. In this type of infrastructure, there is a master node that dispatches jobs to worker nodes and increasing performance is a matter of adding extra worker nodes. This ease of expansion has been put into use by a number of programming frameworks, such as MapReduce[34, 35], and resource allocation management tools, including Condor[36] and TORQUE[41].

Much of the previous work on master-worker architectures targets scalability bottlenecks. Such bottlenecks may develop due to two factors: first, each infrastructure has its own hardware limitations. Second, the data processing algorithms implemented by the jobs cannot always be efficiently parallelized and therefore are not suited for this type of distributed environment. Research on the aforementioned bottleneck factors has resulted in solutions that can efficiently harvest the hardware resources of infrastructures of any size. Experience in Grids has shown that the combination of a resource allocation tool with a programming library for distributed programming (e.g., MPI) can fully utilize small to medium computer clusters [36, 41]. For larger computing infrastructures, frameworks such as MapReduce are shown to display outstanding scalability[34]. Yet, purchasing and maintaining such large physical infrastructures involves a high investment risk. *Infrastructure as a Service* (*IaaS*) clouds have greatly reduced the investment risk of owning an infrastructure, but introduce a new performance scaling factor: the user's financial capacity to rent virtual resources. This additional factor complicates the deployment and management of cloud architectures.

The value for money spent in a multi-tenant, elastic cloud renders resource sharing policies a key player in both cloud performance and user satisfaction. These policies must take into account the fact that cloud providers need strategies to evaluate and reduce possible financial risks while maximizing their profit. In addition, resource sharing policies must take into consideration the user's budget. Such policies often employ either auctioning [42, 43] or attempt to estimate user demand for resources [31]. In this way, the consumer needs are quantified based on their willingness to pay for the resources available. Microeconomic-based resource sharing policies thus allow for the distributed computing infrastructure to reach an equilibrium where the quality of service provided reflects the money spent.

In our thesis, we propose a virtual-machine provision policy based on marginal cost and revenue functions. Each cloud customer announces her budget as a function of the execution time of the tasks she submits. Knowledge of this function, combined with the machine-hour cost, allows for educated decisions regarding

the amount of virtual resources allocated per customer in the context of an *IaaS-Cloud*. The main contribution of our approach is that we provide an answer to the question of exactly how many virtual machines (*VM*s) a consumer should request from a cloud within a budget. In light of scalable, master-worker-based, virtual infrastructures and the seemingly endless resources of a physical cloud, specifying the exact amount of resources needed must be based on a) the consumer's budget and b) the performance bottlenecks which are known only at runtime. We propose a mechanism that continuously monitors user application performance and either "removes" or "adds" *VM*s in response to the observed performance fluctuations serving the needs of autonomic systems [44]. Our approach automatically adjusts to the ever-changing equilibrium point caused by dynamic workloads and thus ensures that resources are shared proportionally to money spent by the users. In addition, our approach is applicable to a wide range of computational environments since it does not enforce the use of a specific job submission tool. We allow each user to select any virtual infrastructure equipped with the tools of her preference.

7 Conclusions

Cloud computing changes our perspective of the services we provide and consume. Cloud services are meant to be scalable, elastic and support a well defined business model. Infrastructure as a Service (*IaaS*) clouds include services that offer infrastructures on-demand. *IaaS* clouds build their success on virtualization so as to over-commit hardware resources and to take advantage of economies of scale.

Often *IaaS* clouds are seen as the evolution of the computing clusters. In both clouds and clusters the goal is the same: the provision of distributed resources. What has changed with the advent of the clouds is the resource provisioning mechanisms and abstractions. Users contact an *IaaS* cloud in search of processing power offered in the form of virtual machines (*VM*s). *VM*s have different specifications that also define their performance. Cloud service consumers have to choose the *VM* that will satisfy their need for computational power. Resource allocation policies map the *VM*s to the physical nodes.

Resource management within the physical infrastructure needs to be revisited as there are several properties that differentiate *IaaS* clouds from other distributed infrastructures. The physical infrastructure is never revealed to the consumers of the services. Current cloud abstractions do not allow cloud consumers to assist in resource management. This allows clouds to transparently support multi-tenancy and offer service scaling. Furthermore, cloud internal tasks such as migration of virtual resources are entirely concealed from the customers. In effect, users are not aware of the fact that the resources/system they are using are virtual, thus they are not aware of any on-going *VM* migration operations.

In the future, we will attempt to have some of the ideas described in this thesis incorporated into commercial and open-source cloud software. This will require us to examen the effectiveness and robustness of available constraint satisfaction

solvers and realize standard interfaces. Access to a variety of cloud installations will reveal several opportunities for improvement in both our implementation and our approach to certain cloud aspects. We will populate the set of deployment hints and we will attempt to form cohorts in more effective ways. Large cloud installations will allow us to deploy our approach regarding time-constrained migrations and evaluate several resource management policies. Finally, we will have the chance to evaluate alternative profit maximization models and test them in *PaaS* clouds.

References

1. OpenStack: http://www.openstack.org (May 2012)
2. OpenNebula: http://www.opennebula.org (May 2012)
3. Nurmi, D., Wolski, R., Grzegorczyk, C., Obertelli, G., Soman, S., Youseff, L., Zagorodnov, D.: The Eucalyptus Open-Source Cloud-Computing System. In: 9th IEEE/ACM Int. Symposium on Cluster Computing and the Grid (CCGRID), Shanghai, China (May 2009) 124–131
4. VMware: VMware hypervisor. http://www.vmware.com/ (2011)
5. Sindelar, M., Sitaraman, R.K., Shenoy, P.: Sharing-Aware Algorithms for Virtual Machine Colocation. In: Proceedings of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures, San Jose, California, USA (June 2011)
6. Breitgand, D., Epstein, A.: SLA-aware Placement of Multi-Virtual Machine Elastic Services in Compute Clouds. In: IFIP/IEEE International Symposium on Integrated Network Managemen, Dublin, Ireland (May 2011)
7. Singh, A., Korupolu, M., Mohapatra, D.: Server-Storage Virtualization: Integration and Load Balancing in Data Centers. In: Proc. of the 2008 ACM/IEEE conference on Supercomputing (SC'08), Austin, TX (2008) 53:1–53:12
8. Hermenier, F., Lorca, X., Menaud, J., Muller, G., Lawall, J.: Entropy: a Consolidation Manager for Clusters. In: Proc. of the 2009 ACM SIGPLAN/SIGOPS Int'l Conf. on Virtual Execution Environments, Washington, DC (March 2009)
9. Hyser, C., McKee, B., Gardner, R., Watson, B.J.: Autonomic Virtual Machine Placement in the Data Center. http://www.hpl.hp.com/techreports/2007/HPL-2007-189.html (2008)
10. Wang, X., Lan, D., Wang, G., Fang, X., Ye, M., Chen, Y., Q.B. Wang, Q.: Appliance-Based Autonomic Provisioning Framework for Virtualized Outsourcing Data Center. In: Proc. of the 4th Int. Conf. on Autonomic Computing, Washington, DC (2007) 29
11. Tsakalozos, K., Roussopoulos, M., Floros, V., Delis, A.: Nefeli: Hint-based Execution of Workloads in Clouds . In: Proc. of the 30th IEEE Int. Conf. on Distributed Computing Sytems (ICDCS'10), Genoa, Italy (June 2010)
12. Keahey, K., Tsugawa, M., Matsunaga, A., Fortes, J.: Sky Computing. IEEE Internet Computing **13** (September 2009)
13. Lee, K., Paton, N., Sakellariou, R., Deelman, E., Fernandes, A., Mehta, G.: Adaptive Workflow Processing and Execution in Pegasus. In: Proc. of 3rd IEEE Int. Conf. on Grid and Pervasive Computing Workshops, Kunming, PR China (2008) 99–106
14. Lee, K., Paton, N., Sakellariou, R., Fernandes, A.: Utility Driven Adaptive Workflow Execution. In: Proc. of 9th IEEE/ACM Int. Symposium on Cluster Computing and the Grid, Shanghai, PR China (2009) 220–227

15. Wood, T., Shenoy, P., Venkataramani, A., Yousif, M.: Black-box and Gray-box Strategies for Virtual Machine Migration. In: Proc of the 4th USENIX Symposium on Networked Systems Design and Implementation, Cambridge, MA (2007)
16. Khanna, G., Beaty, K., Kar, G., Kochut, A.: Application Performance Management in Virtualized Server Environments. In: Proc of the 10th IEEE/IFIP Network Operations and Management Symposium, Vancouver, Canada (April 2006)
17. Bobroff, N., Kochut, A., Beaty, K.: Dynamic Placement of Virtual Machines for Managing SLA Violations. In: Proc of the 10th IFIP/IEEE International Symposium on Integrated Network Management, Munich, Germany (May 2007)
18. Weng, C., Li, M., Wang, Z., Lu, X.: Automatic Performance Tuning for the Virtualized Cluster System. In: Proc. of the 29th IEEE International Conference on Distributed Computing Systems, Montreal, Canada (June 2009)
19. Clark, C., Fraser, K., Hand, S., Hansen, J.G., Jul, E., Limpach, C., Pratt, I., Warfield, A.: Live Migration of Virtual Machines. In: Proc. of the 2nd Symposium on Networked Systems Design & Implementation - Volume 2. NSDI'05, Boston, MA, USENIX Association (May 2005)
20. Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T.L., Ho, A., Neugebauer, R., Pratt, I., Warfield, A.: Xen and the Art of Virtualization. In: SOSP, New York, NY (2003) 164–177
21. Kernel Based Virtual Machine: http://www.linux-kvm.org (May 2012)
22. VMware: VMware ESXi hypervisor. http://www.vmware.com/products/vsphere/esxi-and-esx/index.html (May 2012)
23. Red Hat: Global File System. http://www.redhat.com/gfs/ (May 2012)
24. Red Hat: GlusterFS. http://www.gluster.org/ (May 2012)
25. Weil, S.A., Brandt, S.A., Miller, E.L., Long, D.D.E., Maltzahn, C.: Ceph: A Scalable, High-Performance Distributed File System. In: OSDI'06, Seattle, WA (Nov. 2006) 307–320
26. Luo, Y., Zhang, B., Wang, X., Wang, Z., Sun, Y., Chen, H.: Live and Incremental Whole-System Migration of Virtual Machines Using Block-Bitmap. In: Cluster Computing, 2008 IEEE International Conference on, Tsukuba, Japan (Sept. 2008)
27. Bradford, R., Kotsovinos, E., Feldmann, A., Schiberg, H.: Live Wide-Area Migration of Virtual Machines Including Local Persistent State. In: In VEE 07: Proc. of the 3rd Int. Conf. on Virtual Execution Environments, San Diego, CA (June 2007)
28. Wood, T., Ramakrishnan, K.K., Shenoy, P., van der Merwe, J.: CloudNet: Dynamic Pooling of Cloud Resources by Live WAN Migration of Virtual Machines. SIGPLAN Not. (March 2011) 121–132
29. Zheng, J., Ng, T.S.E., Sripanidkulchai, K.: Workload-Aware Live Storage Migration for Clouds. In: Proc. of the 7th ACM Int. Conf. on Virtual Execution Environments. (VEE '11), Newport Beach, CA (2011)
30. Marbukh, V., Mills, K.: Demand Pricing & Resource Allocation in Market-Based Compute Grids: A Model and Initial Results. In: ICN '08: Proceedings of the Seventh International Conference on Networking, Cancun, Mexico, IEEE Computer Society (Apr. 2008) 752–757
31. Subramoniam, K., Maheswaran, M., Toulouse, M.: Towards a Micro-Economic Model for Resource Allocation. In: In IEEE Canadian Conference on Electrical and Computer Engineering, IEEE Press (2002) 782–785
32. Volker, C.E., Hamscher, V., Yahyapour, R.: Economic Scheduling in Grid Computing. In: JSSPP '02: Revised Papers from the 8th International Workshop on Job Scheduling Strategies for Parallel Processing, London, UK, Springer (2002) 128–152

33. Amazon: Elastic Cloud. http://aws.amazon.com/ec2/ (2010)
34. Dean, J., Ghemawat, S.: MapReduce: Simplified Data Processing on Large Clusters. In: OSDI'04: Sixth Symposium on Operating System Design and Implementation, San Francisco, CA (Dec. 2004)
35. HP Labs: Dynamic Hadoop Clusters. http://wiki.smartfrog.org/ /wiki/display/sf/Dynamic+Hadoop+Clusters (Feb. 2010)
36. University of Wisconsin Madison: Condor. http://www.cs.wisc.edu/condor/ (Nov. 2010)
37. Tsakalozos, K., Roussopoulos, M., Delis, A.: Hint-based Execution of Workloads in Clouds with Nefeli. In IEEE Transactions on Parallel and Distributed Systems, Under Minor Revision (2012)
38. Tsakalozos, K., Roussopoulos, M., Delis, A.: Time Constrained Live VM Migration in Share-Nothing IaaS-Clouds VM Placement in non-Homogeneous IaaS-Clouds. Under Submission (2012)
39. Tsakalozos, K., Kllapi, H., Sitaridi, E., Roussopoulos, M., Paparas, D., Delis, A.: Flexible Use of Cloud Resources through Profit Maximization and Price Discrimination. In: Proc. of the 27th IEEE Int. Conf. on Data Engineering (ICDE'11), Hannover, Germany (Apr. 2011)
40. Tsakalozos, K., Roussopoulos, M., Delis, A.: VM Placement in non-Homogeneous IaaS-Clouds. In: Proc. of the 9th Int. Conf. on Service Oriented Computing, Paphos, Cyprus (Dec. 2011)
41. http://www.clusterresources.com/pages/products/torque-resource manager.php: TORQUE Resource Manager (2010)
42. Grosu, D., Das, A.: Auctioning resources in Grids: model and protocols: Research Articles. Concurrent Computation : Practice and Experience **18**(15) (2006) 1909–1927
43. Chen, C., Maheswaran, M., Toulouse, M.: Supporting Co-allocation in an Auctioning-based Resource Allocator for Grid Systems. In: Proc. of the International Parallel and Distributed Processing Symposium, Fort Lauderdale, Florida, USA (Apr. 2002) 89–96
44. Kephart, J.O., Chess, D.M.: The Vision of Autonomic Computing. IEEE Computer **36**(1) (2003) 41–50

Modeling Acoustic Rendition of Documents' Typography using Expressive Speech Synthesis for Sighted and Blind Users

Dimitrios Tsonos*

National and Kapodistrian University of Athens
Department of Informatics and Telecommunications
dtsonos@di.uoa.gr

Abstract. The accessibility to printed and electronic documents (books, newspapers, web content) by the print disabled, as well as the moving users and the elderly, is based on the possibility to convert them (in real time) into acoustic and/or haptic modality. Besides its content, a printed or electronic text document contains a number of presentation visual elements that apply design glyphs or typographic elements, such as font (type, size and color) and font style (bold and italics). Regardless the important progress achieved in Text-to-Speech systems, they do not support the efficient sonification of the semantics and cognitive aspects of the Visual Presentation Elements in Documents (VPED). Essentially all this additional metadata vanishes during document's processing towards its acoustic or haptic rendition.

This dissertation deals with the sonification of the VPED metadata during their transformation to speech. The approach includes: a) the automatic extraction of VPED induced emotional states to the reader and b) their acoustic rendition using expressive emotional synthetic speech.

Focusing on the development of a system for the automatic extraction of the VPED induced emotional states and the appropriate document annotation, a novel architecture is proposed for the multimodal universal accessibility of documents, regardless of their natural language, content and culture. A quantitative model is developed for the sonification of the typographic alterations by: i) the mathematical formulation of the induced reader's emotional state, based on the dimensional nature of the emotions ("Pleasure", "Arousal" and "Dominance"), and ii) their mapping into prosodic alterations of the expressive synthetic speech.

The prosodic model was evaluated, using psychoacoustic experiments, whether the listeners can acoustically recognize the typographic alterations. The results were positive even in the case of listeners without any previous training. Furthermore, the evaluation of the developed model by sighted and blind students of primary education shows enhancement of their performance during the didactic process.

Keywords: Human Computer-Interaction, Universal Accessibility, Design-for-All, Emotions, Expressive Speech Synthesis

* Dissertation Advisor: Georgios Kouroupetroglou, Assoc. Professor

1 Introduction

A document is the "medium" in which a "message" (information) is communicated [1]. The term "signal" is introduced as, "the writing device that emphasize aspects of a text's content or structure without adding to the content of the text" [2]. It attempts to pre-announce or emphasize content and/or reveal content relationship [3] [4]. The title, heading, typographic cues are considered as signals. Also, "input enhancement" is an operation whereby the saliency of linguistic features is augmented through e.g. textual enhancement for visual input (i.e. bold) and phonological manipulations for aural input (i.e. oral repetition) [5]. Tsonos and Kouroupetroglou [6] categorize the signals, focusing on visual presentation of documents, into three layers: logical, layout and typographic. All the devices [2], either mentioned as signals or layers:

1. share the goal for directing the reader's attention during reading.
2. facilitate specific cognitive process occurring during reading.
3. ultimate comprehension of text information.
4. may influence memory on text.
5. direct selective access between and within texts.

People with print disabilities (as well as moving and/or elderly) require printed or electronic documents in alternative formats, such as Braille, audio, large print or electronic text. Text-to-Speech (TtS) is a common software technology that converts in real-time any electronic text into speech [7]. Most of the current Text-to-Speech systems do not include effective provision of the semantics and the cognitive aspects of the visual (e.g. font style and size) and non-visual signals (e.g. emphasis).

The acoustic rendition of documents' visual information is a complex procedure. In order to accomplish this task, it is used methodologies such as specific sounds (earcons) before and/or after the text that is presented in a specific typographic manner or the use of the prosodic variations in speech synthesis. This method can be considered as "direct acoustic rendition" of the typographic elements because of the direct link/mapping between typography and prosody. Recently, there was an effort towards Document-to-Audio (DtA) synthesis [34] [35]. It essentially constitutes the next generation of the Text-to-Speech systems, supporting the extraction of the semantics of document metadata [8] and the efficient acoustic representation of text formatting [9-12] through modeling the parameters of the synthesized speech signal [13] [14], considered as "indirect acoustic rendition".

W3C introduces Aural Cascaded Style Sheets [15] recommendations towards the rendition of web documents typographic elements using synthetic speech. Kallinen et al. [16] introduces the term "auditive boldfacing". They examine the effects of auditive "boldfacing", i.e., occasional lowering of the voice by two semitones, on the memory performance of audio business news differing in terms of valence and arousal in 28 subjects. "Audio boldfacing" enhanced the immediate memory performance but decreased the longer term performance. Truillet et al. [17] present an experimental study towards evaluating "sound fonts". They study the acoustic rendition of salient words using speech synthesis, obtained either presented by addition of a verbal description of the typographic attribute ("in bold" with a decrease of 15% of the current

pitch is said before the salient word) or pronounced with an increase of 13% of the default pitch. Through various psycho-acoustic manipulations (pitch, volume and speed variations of synthetic speech), Argyropoulos et al. [18] examine their effectiveness for the understanding of specific information (typographic attributes - bold and italic) by thirty sighted and thirty blind participants. A preliminary study of auditory rendition of typographical and/or punctuation information using expressive speech synthesis is presented in [14]. The aim is to increase the expressiveness of the already existing TTS system of France Telecom using prosodic rules. Four prosodic elements are proposed for use: pitch, rate, volume and break. The acoustic rendition of LATEX documents was proposed by Raman [19].

The above studies simply propose rules for the implementation of the acoustic rendition of specific typographic elements. There is not a systematic approach towards the acoustic rendition of typographic elements. So, it is difficult to implement an automated system, which supports this functionality. The present dissertation aims, not only to propose a set of rules, but also a methodology that through its implementation we are able to create a system that it can automatically produce and apply the typography to acoustic mapping rules utilizing emotions and/or emotional states.

Thus, this work:

1. presents a novel integrated XML-based system, content, language and domain independent, for the real-time production, presentation and navigation of multimodal accessible documents by conforming to international guidelines and standards.
2. proposes a modular, content-free and language independent methodology, for the acoustic rendition of document's typographic elements, using the dimensional approach of emotions (this assumption is based on the universal character of emotions across languages and cultures).
3. introduces a set of rules towards the automated reader's emotional state response on document's typographic elements, based on the dimensional description of emotions.
4. introduces specific rules for the acoustic rendition of font style and size alterations, combining reader's emotional state variations and expressive speech synthesis.
5. evaluates the derived acoustic mapping rules, through listeners' preferences and the level of recognition (psychoacoustic experiments).

2 Multimodal Accessibility of Documents

The discussion on accessibility of documents imposes questions on how the documents can be accessed either in visual, acoustic or haptic modality. Some studies are trying to create accessible documents in the acoustic modality using speech synthesis or by combining earcons, auditory icons and 3D sounds for the auditory [20]. In this section a novel integrated XML-based system is presented, for the real-time production, presentation and navigation of multimodal accessible documents by conforming to international guidelines and standards (such as W3C's and ANSI/NISO). This approach includes a unified methodology for the multimodal rendering of text formatting, text structure, text layout and non-textual elements.

The overall proposed architecture consists of two main parts (Fig. 1): the User Interface and the Document Explorer. The later is responsible for document analysis during the *production process* and the exploration in the documents (during *navigation task*). User Interface is responsible for the multimodal interaction with the documents. It collects the user preferences and the device requests, as well as the navigation commands, and executes the *presentation task*. These parts are implemented by two different modules following the Client – Server model. The Document Explorer can be hosted on a powerful server machine due to the resource demanding tasks that performs. In contradiction, the User Interface can be hosted on any common computer (e.g. a personal computer, a PDA, mobile smart phone). This kind of implementation fulfills the Web Content Accessibility Guidelines, in order to be device and software independent. The implementation and the communication between the modules are XML-based. The proposed architecture also can be slightly altered and can be used e.g. in classroom for sighted and blind students [27] [28] [29].

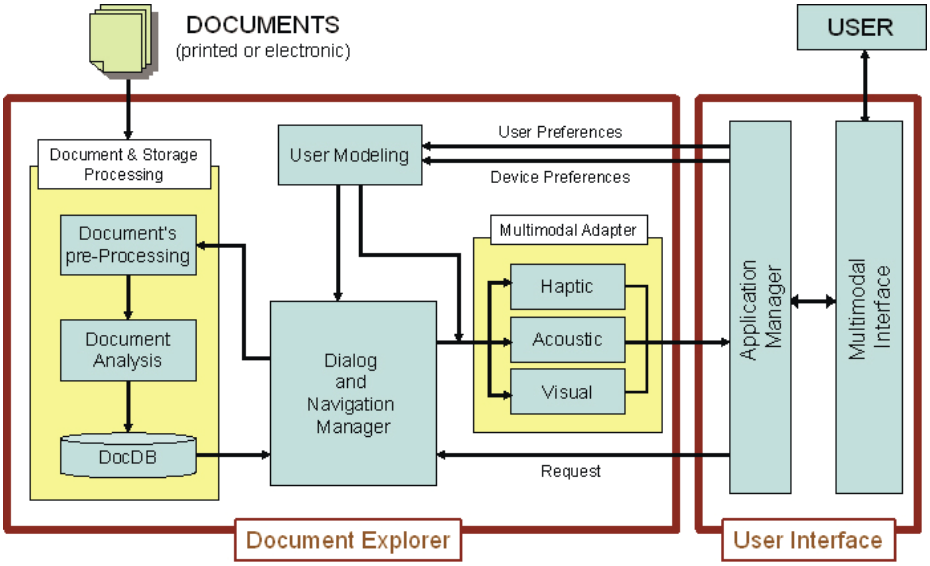

Fig. 1. The XML-based system for multimodal accessibility of documents

3 Acoustic Rendition of Documents' Typography using Expressive Speech Synthesis

While reading a document, time passes and typographic elements alter, thus reader's emotions and emotional state vary. It is a similar syllogism to one used in the expressive speech synthesis model by Schröder [23] who uses the dimensional approach of emotions. It is worthy to mention that the ordinary alterations of typographic elements are not extreme e.g. font size from 12pt to 14pt. or plain text to italized. Thus the dimensional approach of emotions, either in typography or speech, is con-

sidered as an appropriate approach for the mapping of the typographic elements to prosodic variations. The dimensional approach of emotions in speech is based on the continuous nature of the emotion representation and variation, so we are able to utilize this model in order to represent the analogous behavior of typographic elements in documents to the acoustic modality.

3.1 Visual and Acoustic Modality Analogy

The primary goal of using presentation elements and specifically typographic elements in text documents, is to distinguish parts of the text and to create a well formed presentation of the content in order e.g. to augment the reading performance, attract the reader, render semantics through the visual channel. Authors use typography in a specific way, e.g. in scientific documents there are "strict" typographic rules. But in newspapers and magazines, the page-art design department and not the authors has the primary responsibility on applying typography. The notion "typographic profile" of an emotion is introduced, in a similar way that "prosodic profile" is presented, in [24]. The way typographic parameters are used (here are referred as "typographic elements"), constitute the typographic profile of the document. The profile defines the space within each emotion is located (and vice-versa). In Figure 2 are presented the resemblances of prosodic and typographic profile respectively.

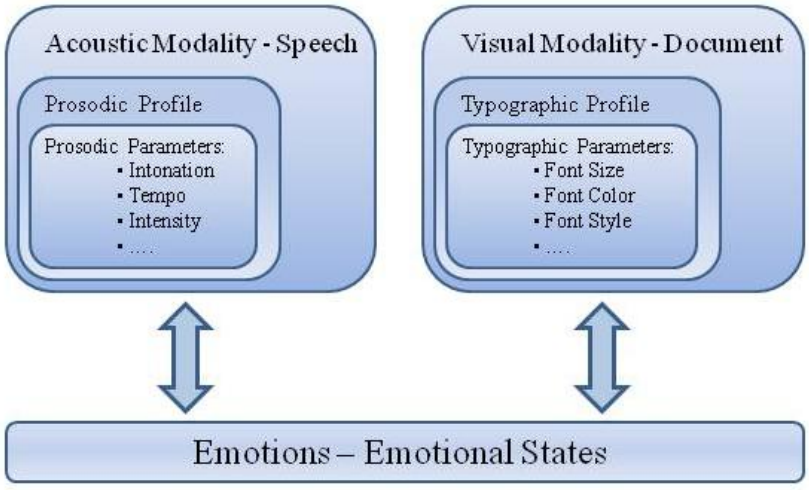

Fig. 2. The prosodic and typographic profiles.

Typography-to-emotions and emotions-to-speech approaches, can map the typographic alterations into speech prosodic variations and we are able to observe which speech parameters are influenced and how they vary, while changing the font style/size attributes. The combination of these models provides a model on how the prosody varies according to typographic alterations. Thus, human emotions constitute the medium for the two channels / modalities to interact - communicate.

3.2 Modeling Reader's Emotional State Response

In [6] are described the results of an experimental study towards modeling the reader's emotional state variations derived by the typographic elements in documents. Based on the experimental results, it is proposed a set of mapping rules on how typographic elements, such as typesetting (plain, italics, bold and bold-italics), and font elements (font/background color combinations, font type and font size) affect the reader's emotional state. The study was based on the dimensional approach of the emotions (namely "Pleasure", "Arousal" and "Dominance") and the experimental procedure was designed and implemented according to IAPS experimental guidelines, using the Self Assessment Manikin test [21]. It is clear that emotional states have similar affection and variations across language and culture. This assumption derives from the comparison of our experimental results with those from similar studies for other languages and cultures.

In the present dissertation is introduced a model for the quantitative representation of the results. It is used the polynomial regression for the case of font size variations (continuous values) and simple percentage variations for the case of font type and typesetting (discrete values). Based on this modeling and the need of implementing the rules for the automated reader's emotional state extraction [22] [33], a generic mathematical description is proposed on how the variations of the typographic elements can be mapped into emotional state values.

Some major findings of this work are: Three polynomial equations describe the way the emotional dimensions are affected by the font size (Table 1). While font size increases up to 26px, "Pleasure" dimension increases respectively. From this value and higher there is a decrease. In contradiction, "Arousal" and "Dominance" decreases up to 15px and 18px respectively. Beyond these values there is an increase in both dimensions. The results denoted that only "Pleasure" is affected by the font type. "Times New Roman" are considered more pleasant than Arial (23% increase of "Pleasure") and only "Pleasure" is affected by typesetting elements.

Table 1. Constants of the modeling of PAD relationship on text size using equation (1) with their standard errors.

	Pleasure		Arousal		Dominance	
		SE		SE		SE
Intercept	-2.45287	0.30771	3.58459	0.73588	2.85307	0.77138
B_1	0.24942	0.03299	-0.59264	0.12133	-0.41265	0.12828
B_2	-0.00523	0.00077	0.02866	0.00632	0.01631	0.0067
B_3	0	0	-0.000427111	0.0001029	-0.000199863	0.0001093
Adjusted R-Square	**0.90687**		**0.78103**		**0.84705**	

3.3 Typography to Expressive Speech Synthesis Modeling

Following the proposed methodology and utilizing the model proposed by Tsonos & Kouroupetroglou [6], we can calculate the emotional state variations from the neutral state derived by the "typographic baseline". The typographic baseline is defined as "the most frequent value of the typographic element that is occurred in the whole document". A statistical survey [25] [32] was conducted on printed school book and newspapers. The more frequent font sizes that consist the baseline of the documents are 10pt, 12pt and 14pt (throughout all books and newspapers). Also it is observed that the increment of font size (occurred e.g. in title, headline) has values of 12pt, 14pt, 16pt, 18pt, 20pt, 22pt and decrement (e.g. footnote) at 8pt, 9pt, 10pt, 11pt, 12pt, 13pt depending the font size baseline.

The mapping of typographic alterations into prosodic variations can be accomplished combining the results derived from the reader's emotional state response model and the expressive speech synthesis model proposed by Schröder [23] and its implementation, the XSLT file used in OpenMary System [26]. The prosodic variations are very small due to little variations of the emotional states derived from the alterations of the font size from the baseline. These variations are not distinct to the human listener. So it is needed to be "augmented" in order to be more comprehensible. Thus, it is proposed the normalization of the values using linear quantization of these variations.

The levels of quantization are defined calculating the difference between minimum and maximum values for each prosodic element and dividing it by 6 levels of increment or decrement. The number of levels is considered as the optimum and derives from observations while trying not to have duplicate cases for each prosodic element variations. Greater and lower values, is observed many duplicate cases.

The value associated for each level of increment/decrement for each prosodic element is:

- 12% for pitch (corresponds approximately to 2 semitones of increment/decrement).
- 10% - 14wpm for speech rate (6 levels of increment correspond to maximum rate of 224 wpm not exciding the upper limit of 225 wpm).
- 15% for intensity corresponding to more than >1dB increment/decrement.
- After and before each prosodic variation there is a speech pause of 150 ms duration.

The model proposes multiple cases for each typographic alteration (3 cases for each font size and 2 cases for each font style alteration from the baseline). Using many prosodic variations in speech and implementing the same typographic element in different ways, the listener can be confused and annoyed. Also, in the present study, we want to create a model that can be applied to any document, examining only the typographic alterations from the baseline and not absolute values. Thus, a psychoacoustic experimental procedure was implemented for the selection of the optimum acoustic rendition of typographic alterations. The results are [30]:

1. plain to bold alteration affects pitch (-12%) and rate (-10%),

2. plain to italics, affects pitch (+36%) and rate (+30%),
3. plain to bold-italics, affects pitch (-36%) and rate (-30%),
4. increasing font size (+2pt) should increase pitch (+12%), decrease rate (-10%) and volume (-15%),
5. decreasing font size (-1pt) should increase pitch (+12%), rate (+10%) and volume (+30%).

4 Prosodic Model Evaluation

In order to evaluate the prosodic model, another experimental procedure was conducted [30]. It is examined whether the listeners are able to recognize the typographic alterations (from plain to bold, italics, bold-italics, increment and decrement of font size) by their corresponding prosodic variations.

Eleven male and eight female participated in the experimental procedure (total 19 participants) - aging from 20 to 32 year-old (Mean Age = 27.8, SD = 3.2). They were undergraduate or postgraduate students from our department; their native language is Greek, naive to any previous experiment. Totally, 20 stimuli were used in Greek language, 12 for the font style elements and 8 for the font size. Each stimulus contained a sentence that is acoustically expressed in a typographic element or a phrase of the sentence consisting from 1, 2 or 3 words. The sentences were converted into speech using the Document-to-Audio platform [12].

There was *no training session* for the participants to be familiarized with prosodic variations due to typographic (font style and size) alterations. They participated each one in two sessions: evaluation of the acoustic rendition of a) font style elements and b) font size elements respectively. The stimuli were displayed simultaneously with the sentence without the typographic elements. The evaluators were asked which typographic element they believe is associated with the sentence or phrase (font style: bold, italics, bold-italics or none / font size: size increment, size decrement or none)

The results denote that participants easily recognized the font style and size variations. In details:

- Italics were recognized successfully.
- Bold-italics failed to be recognized by the evaluators and were confused with the bold (almost all cases were assessed by the majority of the participants as bold). This is due to the similar prosodic variations for the both cases of bold and bold-italics (decrement of pitch and rate).
- Font size variations are also recognized successfully. In one case the correct answers were marginally lower 36.8%. We estimate that probably the false answer is due to stimulus implementation (e.g. the quality of speech synthesis for the specific sentence).

The model was also evaluated through another experimental procedure [31] by 10 sighted and 10 blind students of primary education (Mean Age=11 years and 6 months). The visual stimuli (MS PowerPoint slides) were presented using an interactive whiteboard. The content of the slides were converted into speech using

DEMOSTHeNES TtS applying the proposed prosodic model. The results show enhancement of their performance during the didactic process.

5 Conclusions

To ensure that people with print disability are able equally to participate in society, it is crucial to develop more effective ways for the accessibility of both printed and electronic documents. In this study it is presented an integrated architecture on how a document is structured. Based on international standards and guidelines, it is proposed a novel holistic XML-based system for the real time production, presentation and navigation of multimodal accessible documents.

Also it is presented a novel approach for the acoustic rendition of typographic alterations in documents targeting to a methodology that is modular, content-free and language independent. Based on reader's emotional state variations on documents typographic alterations (from plain text to bold, italics and bold-italics / several font sizes) and using expressive speech synthesis we are able to acoustically render these alterations in order to be clearly distinguished by the human listener. This description can be further extended to other modalities, to visual and/or haptic, with prior knowledge of their mathematical modeling.

Utilizing reader's emotional state model proposed by Tsonos and Kouroupetroglou [6], the difference of the percentage emotional state variations (Pleasure, Arousal and Dominance) is calculated for a set of typographic alterations (from plain text to bold, italics and bold-italics / font size alterations from specific font size baselines, that was studied in the statistical survey presented in [25] [32]). Utilizing expressive speech synthesis [23] we are able to map these emotional state variations into prosodic (pitch, rate and volume) ones. Using this "indirect" manner, we map the typographic alterations into prosodic variations. The resulting variations are below the level of human listener's perception and performance. Thus, in order to enhance the perception, the normalization is proposed through linear quantization of the prosodic variations and assign to each level a minimum perceived value for each prosodic element. Because the resulting acoustic model sets multiple ways for the acoustic rendition of font style and size alterations, an experimental procedure was conducted in order to select the optimum implementation.

In general, it is observed that the recognition results are promising. Some major findings are: a) plain to bold alteration affects pitch (-12%) and rate (-10%), b) plain to italics affects pitch and rate (increasing them by 36% and 30% respectively), c) increasing font size (+2pt) should increase pitch (+12%), decrease rate (-10%) and volume (-15%), d) decreasing font size (-1pt) should increase pitch (+12%), rate (+10%) and volume (+30%). The results denoted the lack of representing bold-italics (but it has a similar behavior with the case of bold, decrement of both pitch and rate). Also we should have in mind that the participants were *not trained* to the acoustic variations and the acoustic representation of typographic elements. So it would be very interesting to examine, in a future study, the proposed methodology in an experimental procedure with trained participants to these prosodic variations. Another pro-

posed future study is to change the minimum perceived levels for pitch, rate and volume in order to eliminate the pitfalls of the current study (e.g. the misunderstanding of bold-italics and bold during their acoustic rendition) or the use of non-linear quantization instead of the proposed linear.

Also it would be very interesting the study of acoustic rendition of typographic elements in a similar methodology using discrete emotions (than the dimensional approach). This would extend the use of current methodology in systems than can process reader's discrete emotions derived by the typographic elements/alterations and the T-t-S systems that only support discrete emotions approach. This study is not limited to the acoustic rendition of typographic elements but it can be extended for the acoustic rendition of emotions that derive from the content itself. Finally, the proposed methodology is targeting to other modalities (multimodal approach). It sets the guidelines towards future work for the haptic modality or for a multimodal system due to its modular architecture.

6 References

1. McLuhan, M., Fiore, Q.: The Medium is the Message. Berkeley, USA, Gingko Press (2005)
2. Lorch, R.F.: Text-Signaling Devices and Their Effects on Reading and Memory Processes. Educational Psychology Review, 1, 209-234 (1989)
3. Spyridakis, J.H.: Signaling effects: A review of the research—Part I. Journal of Technical Writing and Communication. 19, 227-240 (1989)
4. Lemarie, J., Eyrolle, H., Cellier, J. M.: Visual signals in text comprehension: How to restore them when oralizing a text via a speech synthesis?. Computers in Human Behavior, 22, 1096-1115 (2006)
5. Han, Z.H., Park, E.S., Combs, C.: Textual Enhancement of Input: Issues and Possibilities. Applied Linguistics, 29, 597-618 (2008)
6. Tsonos, D., Kouroupetroglou, G.: Modeling reader's emotional state response on document's typographic elements. Advances in Human Computer Interaction, 2011, 1-18, (2011)
7. Fellbaum, K., Kouroupetroglou, G.: Principles of Electronic Speech Processing with Applications for People with Disabilities, Journal Technology and Disability, 20(2), 55-85 (2008)
8. Fourli-Kartsouni, F., Slavakis, K., Kouroupetroglou, G., Theodoridis, S.: A Bayesian Network Approach to Semantic Labelling of Text Formatting in XML Corpora of Documents. LNCS, vol. 4556, pp. 299-308 (2007)
9. Xydas, G., Kouroupetroglou, G.: Augmented Auditory Representation of e-Texts for Text-to-Speech Systems, LNAI, vol. 2166, pp. 134-141 (2001)
10. Xydas, G., Kouroupetroglou, G.: Text-to-Speech Scripting Interface for Appropriate Vocalisation of e-Texts. In: EUROSPEECH 2001, pp. 2247–2250 (2001)
11. Xydas, G., Spiliotopoulos, D., Kouroupetroglou, G.: Modelling Emphatic Events from Non-Speech Aware Documents in Speech Based User Interfaces. In: 10th International Conference on Human-Computer Interaction, pp. 806–810 (2003)
12. Xydas, G., Argyropoulos, V., Karakosta, Th., Kouroupetroglou, G.: An Experimental Approach in Recognizing Synthesized Auditory Components in a Non-Visual Interaction with Documents, In: 11th Int. Conference on Human-Computer Interaction (2005)

13. Gussenhoven, C.: Intonation and interpretation: phonetics and phonology. In: Speech Prosody 2002, pp. 47–57, URL http://aune.lpl.univ-aix.fr/sp2002/pdf/gussenhoven.pdf (2002)
14. Launay, J., Segalen, L., Kanellos, I., Moudenc, T., Otesteanu, C., David, A., Fang, G., Jin, J.: Speech expressiveness: Modeling and implementing the expressive impact of typographic and punctuation marks for textual inputs. In: 3rd International Conference on Information and Communication Technologies: From Theory to Applications, pp. 1-6 (2008)
15. ACSS. Aural Style Sheets: http://www.w3.org/TR/CSS2/aural.html
16. Kallinen, K., Laarni, J., Ravaja, N., Saari, T.: Auditive "Boldfacing", Emotional Characteristics of Message, and Individual Differences in Memory Acquisition of Computer Mediated Business News. In: 11th Conference on Human-Computer Interaction (2005)
17. Truillet, P., Oriola, B., Nespoulous, J.L., Vigoroux, N.: Effect of Sound Fonts in an Aural Presentation. In: 6th ERCIM Workshop, UI4ALL, pp. 135-144 (2000)
18. Argyropoulos, V., Sideridis, G., Kouroupetroglou, G., Xydas G.: Auditory Discriminations of Typographic Attributes of Documents by Students with Blidness, British Journal of Visual Impairment, 27(3), pp. 183-203 (2009)
19. Raman, T.V.: An Audio view of (LA)TEX Documents. In: Annual Meeting. TUGboat, 13(3), pp. 65-70 (1992)
20. Kouroupetroglou G., Tsonos, D.: Multimodal Accessibility of Documents. Chapter in the book: Shane Pinder (ed.) Advances in Human-Computer Interaction, I-Tech Education and Publishing, Vienna, pp. 451-470 (2008)
21. Lang, P. J., Bradley, M., Culthbert, B.: International Affective Picture System (IAPS): Instruction Manual and Affective Ratings, Technical Report A-6, The Center for Research in Psychophysiology, University of Florida, U.S.A. (2005)
22. Kouroupetroglou, G., Tsonos, D., Vlahos, E.: DocEmoX: A System for the Typography-Derived Emotional Annotation of Documents. LNCS, vol. 5616, pp. 550-558 (2009)
23. Schröder M.: Expressing degree of activation in synthetic speech. IEEE Transactions on Audio, Speech and Language Processing, 14(4), pp. 1128-1136 (2006)
24. Tatham, M., Morton, K.: Expression in Speech: Analysis and Synthesis Oxford linguistics, Oxford University Press (2006)
25. Tsonos, D., Ikospentaki, K., Kouroupetrolgou, G.: Towards Modeling of Readers' Emotional State Response for the Automated Annotation of Documents. In: IEEE World Congress on Computational Intelligence, pp. 3252-3259 (2008)
26. Schröder , M., Trouvain, J.: The German Text-to-Speech Synthesis System MARY: A Tool for Research, Development and Teaching. International Journal of Speech Technology, 6, pp. 365-377 (2003)
27. Tsonos, D., Kaccori, H., Kouroupetroglou, G.: A Design-for-All Approach Towards Multimodal Accessibility of Mathematics. In: 10th International Conference of the Association for the Advancement of Assistive Technology in Europe (2009)
28. Tsonos D., Kouroupetroglou, G.: Accessibility of Board and Presentations in the classroom: a Design-for-All Approach. In International Conference on Telehealth and Assistive Technologies, pp. 13-18 (2008)
29. Ikospentaki, K., Vosniadou, S., Tsonos, D., Kouroupetroglou, G.: HOMER: A Design for the Development of Acoustical-Haptic Representations of Document Meta-Data for Use by Persons with Vision Loss. In: 2nd European Cognitive Science, vol. 25, pp. 912 (2007)
30. Tsonos, D., Kouroupetroglou, G.: Prosodic Mapping of Typographic Alterations based on the Dimensional Theory of Emotions. IEEE Transactions on Affective Computing, under review (2012)

31. Ikospentaki, K., Tsonos, D., Vosniadou, S., Kouroupetroglou, G.: Sonification of Board Typographic Elements: Does it Enchance the Traditional Teaching Approach of Sighted and Blind Students?. Journal of Visual Impairment and Blindness, in preparation (2012)
32. Katsoulis, F., Tsonos, D., Kouroupetroglou, G.: Analysis of Text Signaling Devices. International Journal on Document Analysis and Recognition, in preparation (2012)
33. Tsonos, D., Kouroupetroglou, G.: A Methodology for the Extraction of Reader's Emotional State Triggered from Text Typography. Chapter in the book: Paula Fritzsche (ed.) Tools in Artificial Intelligence, In-Tech Publishing, Vienna, pp. 439-454 (2008)
34. Tsonos, D., Xydas, G., Kouroupetroglou, G.: Auditory Accessibility of Metadata in Books: A Design for All Approach. LNCS, vol. 4556, pp. 436-445 (2007)
35. Tsonos, D., Xydas, G., Kouroupetroglou, G.: A Methodology for Reader's Emotional State Extraction to Augment Expressions in Speech Synthesis. In: 19th IEEE Int. Conf. on Tools with Artificial Intelligence, vol. II, pp. 218-225 (2007)