MULTIMEDIA COMMUNICATIONS TECHNICAL COMMITTEE IEEE COMMUNICATIONS SOCIETY

http://www.comsoc.org/~mmc



Vol. 6, No. 4, April 2011



CONTENTS

Message from E-Letter Director	3
Special Issue on Emerging Techniques in 3D: 3D Data Fusion, Motion Trackin	ng in
Multi-View Video, 3DTV Archives and 3D Content Protection	4
SPECIAL ISSUE ON DECISION AND GAME THEORY FOR SECURITY	6
Decision and Game Theory for Security (GameSec 2010)	6
Tansu Alpcan, Technical University of Berlin, Germany, alpcan@sec.t-labs.tu-	
berlin.de	6
Adversarial Control in a Delay Tolerant Network	8
Eitan Altman ¹ , Tamer Başar ² , and Veeraruna Kavitha ¹	8
¹ INRIA, France, ² University of Illinois, Urbana-Champaign, USA	8
{Eitan.Altman, Kavitha.Voleti_Veeraruna}@sophia.inria.fr, basar1@illinois.edu	ı 8
ISPs and Ad Networks Against Botnet Ad Fraud	11
Nevena Vratonjic, Mohammad Hossein Manshaei and Jean-Pierre Hubaux	11
School of Computer and Communication Sciences, EPFL, Switzerland	11
{ nevena.vratonjic, hossein.manshaei, jean-pierre.hubaux}@epfl.ch	11
When Do Firms Invest in Privacy-Preserving Technology?	15
Murat Kantarcioglu, Alain Bensoussan, and SingRu(Celine) Hoe	15
University of Texas at Dallas, USA	15
{muratk, alain.bensoussan}@utdallas.edu, hoceline02@yahoo.com	15
Access Control: Allocating Resources to Selfish Agents	18
Farzad Salim, Jason Reid, Uwe Dulleck, Ed Dawson	18
Queensland University of Technology, Brisbane, Australia	18
{f.salim, jf.reid, uwe.dulleck, e.dawson}@qut.edu.au	18
SPECIAL ISSUE ON TECHNOLOGY ADVANCES	22
Distributed Computing for Multimedia Applications	22
Guest Editor: Jian Tan, IBM T.J. Watson Research, USA, tanji@us.ibm.com	22
Flexible MapReduce Scheduling	24
Joel Wolf, Deepak Rajan, Kirsten Hildrum, Rohit Khandekar, Vibhore Kumar, Su	ıjay
Parekh, Kun-Lung Wu and Andrey Balmin, IBM Research	24
{jlwolf, drajan, hildrum, rohitk, vibhorek, sujay, klwu, abalmin}@us.ibm.com	24
MapReduce Workload Management & Multimedia Applications	29
Jordà Polo, David Carrera, Yolanda Becerra, Jordi Torres,	29
Eduard Ayguadé (UPC/BSC) and Malgorzata Steinder (IBM Research)	29

{jorda.polo,david.carrera,yolanda.becerra, jordi.torres, eduard.ayguade}@bsc.es	29
steinder@us.ibm.com	. 29
Understanding Network Communication Performance in Virtualized Cloud	. 32
Guohui Wang, T. S. Eugene Ng, Rice University, USA	. 32
{gswang, eugeneng}@cs.rice.edu	. 32
Information Security in Distributed Computing	. 35
Winnie Cheng (IBM Research), Carlos R. P. dos Santos (UFRGS) and Nikos	
Anerousis (IBM Research)	. 35
{wcheng, nikos}@us.ibm.com, crpsantos@inf.ufrgs.br	. 35
Parallel Lasso for Large-Scale Video Concept Detection	. 39
Bo Geng ¹ , Yangxi Li ¹ , Dacheng Tao ² , Meng Wang ³ , Zheng-Jun Zha ³ and Chao Xu ¹	. 39
¹ Peking University, China, ² University of Technology Sydney, Australia, ³ National	
University of Singapore, Singapore	. 39
E-Letter Editorial Board	. 43
MMTC Officers	43

Message from E-Letter Director

Dear MMTC fellow members,

Thank you for your support to MMTC and welcome to the April issue of E-Letter!

In 2011, E-letter has resumed to be a monthly publication. Including the April issue, 2011 E-letter has so far brought together a total of 47 research articles from many world renowned researchers in the multimedia research community. The January to April E-letter features a collection of 8 special issues and technology advances columns, including

- Human-centric multimedia communications
- Mobile Internet Television
- Video-Aware Wireless Networks
- Interactive and Smart Multimedia Services
- Multimedia Distributed Networks
- Multimedia over Heterogeneous Networks
- Decision and Game Theory for Security
- Distributed Computing for Multimedia Applications

Thanks to E-letter director, editors as well as chairs and vice-chairs of SIs for their dedication and efforts in putting together these special issues. I also want to thank Dr. Jianwei Huang and Dr. Haohong Wang and other MMTC officers for their strong support to E-letter.

Looking forward, in 2011 E-letter board members will continue to work together with special interest groups of MMTC to bring to Eletter readers the latest technical advances in multimedia signal processing and communications. Please follow closely our Eletter at the beginning of every month! I would like to draw your attention to the call for paper of Special Issue on Emerging Techniques in 3D: 3D Data Fusion, Motion Tracking in Multi-View Video, 3D TV Archives and 3D Content Protection. This special issue is fully supported by the 3DRPC IG of MMTC. This is a great opportunity to share your research findings with other researchers in the community. I encourage all MMTC members to submit your paper to this special issue.

Special thanks to guest editor Dr. Tansu Alpcan (Technical University of Berlin, Germany) for organizing a fantastic special issue on *Decision and Game Theory for Security*, which features a collection of four excellent papers. Please find more details about this special issue from Dr. Alpcan's editorial on Page 6.

In the Technology Advances Column, Dr. Jian Tan from IBM T.J. Watson Research has put together an excellent issue on *Distributed Computing for Multimedia Applications*. This issue includes five research articles from topnotch researchers from worldwide. Please check it out starting from Dr. Tan's editorial on Page 22.

As always, I want to thank the editors and authors for their great efforts. I hope you enjoy reading this issue of E-letter!

Kai Yang

Co-Director IEEE ComSoc MMTC E-letter Bell Labs, Alcatel-Lucent

Call for Papers IEEE Signal Processing Society IEEE Journal of Selected Topics in Signal Processing

Special Issue on Emerging Techniques in 3D: 3D Data Fusion, Motion Tracking in Multi-View Video, 3DTV Archives and 3D Content Protection

As a result of increasing consumer demand for 3D content, content creation associated with this new modality has increased significantly in conjunction with some recent standardization activities on this data type. Hence, a scientific revisit is required particularly to some challenging problems associated with the conventional video, considering the fact that multi-view video has many promising solutions to such problems. Moreover, any 3D representation also produces its own requirements to be dealt with. This special issue is an effort to compile and review the current advances in 3D multimedia and multimodal information analysis and processing. Particularly, it deals with emerging 3D techniques for 3D data integration and object analysis based on multi-views, as well as 3D content protection.

We would like to invite authors to submit their recent and original research results as well as experience reports. In the following, a non-exclusive list of related topics is suggested:

- Creation of 3D Content
 - o 3D multi-view and multimodal data fusion
 - Calibration methods for 3D multi-camera system
- Tracking, Registering and Processing of 3D Content
 - o 3D multi-view image/video processing
 - o 3D registration of multi-view data
 - Motion tracking in stereo and multi-camera systems
- Archiving of 3D Content
 - o Creation, compression in 3D digital archives
 - Indexing and retrieval of 3D content
- Security Issues for 3D Content
 - o Methods for 3D content protection
 - o 3D digital watermarking, fingerprinting and related security solutions
- 3D Objective Quality Measures
 - o Artifacts Characterization
 - Full-reference and partial-reference measures
- Multimedia systems and applications using emerging techniques in 3D

Prospective authors should visit <u>http://www.signalprocessingsociety.org/publications/periodicals/jstsp/</u> for information on paper submission. Manuscripts should be submitted using the Manuscript Central system at <u>http://mc.manuscriptcentral.com/jstsp-ieee</u>. Manuscripts will be peer reviewed according to the standard IEEE process.

Manuscript submission due:	July 10, 2011
First review completed:	October 20, 2011
Revised manuscript due:	November 20, 2011
Second review completed:	February 1, 2012
Final manuscript due:	March 1, 2012

Lead guest editor:

A. Aydın Alatan, Dept.of Electrical & Electronics Engineering, M.E.T.U., Turkey, alatan@eee.metu.edu.tr

Guest editors:

Joern Ostermann, Institut für Informationsver., Leibniz Universität Hannover, Germany, ostermann@tnt.unihannover.de

Levent Onural, Dept. of Electronics Engineering, Bilkent University, Ankara, Turkey, onural@ee.bilkent.edu.tr

Ghassan AlRegib, School of Electrical & Computer Eng., Georgia Institute of Technology, Atlanta, USA, <u>alregib@gatech.edu</u>

Stefano Mattoccia, Faculty of Engineering of the University of Bologna, Italy, <u>smatt@ieee.org</u>

Chunrong Yuan, Cognitive Neuroscience, University of Tuebingen, Germany, <u>chunrong.yuan@uni-tuebingen.de</u>

SPECIAL ISSUE ON DECISION AND GAME THEORY FOR SECURITY

Decision and Game Theory for Security (GameSec 2010)

Tansu Alpcan, Technical University of Berlin, Germany, alpcan@sec.t-labs.tu-berlin.de

Securing complex systems and managing associated risks become more and more important as these systems and networks play an increasingly important role in modern life. At the same time, ubiquitous communication, data, and computing, create novel research challenges in terms of security and risk management. approaches to address Analytical these challenges through quantitative models are in early stages of research, and there are a number of future directions. Decision, control, and game theories provide excellent mathematical and analytical foundations to support novel quantitative approaches for media processing and communications security.

GameSec 2010, the inaugural Conference on Decision and Game Theory for Security has taken place on the campus of Technical University Berlin, Germany, on November 22-23, 2010. GameSec conference brings together researchers who aim to establish a theoretical foundation for making resource allocation decisions that balance available capabilities and perceived security risks in a principled manner. The conference focuses analytical models based information, on game, communication, optimization, decision, and control theories that are applied to diverse security topics. At the same time, the connection between theoretical models and real world security problems are emphasized to establish the important feedback loop between theory and practice.

Focusing on a deeper theoretical understanding of the underlying incentive and resource allocation issues in security, this special issue aims to bring a different perspective to media processing and communications security.

The article "Adversarial Control in a Delay Tolerant Network" considers a multi-criteria control problem that arises in delay tolerant networks with two adversarial controllers: source and jammer and two types of information structures: closed and open loop. The equilibrium solution is characterized under energy constraints and shown to be similar under both open-loop and closed-loop structures. In the article "ISPs and Ad Networks Against Botnet Ad Fraud", a game-theoretic analysis of the behavior and interactions of the ISPs and ad networks facing botnet ad fraud is provided. It is shown that cooperation could emerge under certain conditions that mostly depend on factors such as the number of infected devices, the aggregate power with which bots divert revenue from the ad networks, and the efficiency of the botnet detection system. In addition, the outcome of cooperation is analysed for various situations.

The article "When Do Firms Invest in Privacy-Preserving Technology?" analyses how the relationship between a customer's valuation of his private information and a customer's profitability to a firm affects a firm's investment decision in privacy-preserving technologies and formulates guidelines for entrepreneurs' adoption decision.

An authorisation mechanism as a game between a *selfish employee* and the *benevolent employer* who is the sole authority in making authorisation decisions, is formulated in the article "Access Control: Allocating Resources to Selfish Agents". The adopted game theoretic framework allows to reason about the factors that may affect users' likelihood to misuse a permission at the time of an access decision. Hence, game theory provides a useful perspective in authorisation theory: the notion of the user as a self-interested player who selects among a range of possible actions depending on their pay-offs.



Tansu Alpcan received the B.S. degree in electrical engineering from Bogazici University, Istanbul, Turkey in 1998. He received the M.S.

and Ph.D. degrees in electrical and computer engineering from University of Illinois at Urbana-Champaign (UIUC) in 2001 and 2006, respectively. His research involves applications of distributed decision making, game theory, and control to various security and resource allocation problems in complex and networked systems. He is recipient of multiple research and best paper awards from UIUC and IEEE. He has played a role in organization of several workshops and conferences such as IEEE Infocom, GameComm, and GameSec as TPC member, associate editor, co-chair, chair, and steering board member. He is the (co-)author of more than 100 journal and conference articles, an edited volume, as well as the book "Network Security: A Decision and Game Theoretic Approach" published by Cambridge University Press in 2011. He has worked as a senior research scientist in Deutsche Telekom Laboratories, Berlin, Germany, between 2006 and 2009. Tansu Alpcan is currently assistant professor (Juniorprofessur) in Technical University Berlin while continuing his affiliation with Deutsche Telekom Laboratories.

7/43

Adversarial Control in a Delay Tolerant Network

Eitan Altman¹, Tamer Başar², and Veeraruna Kavitha¹ ¹INRIA, France, ² University of Illinois, Urbana-Champaign, USA. {Eitan.Altman, Kavitha.Voleti_Veeraruna}@sophia.inria.fr, basar1@illinois.edu

1. Introduction and the Model

Consider *n* relay mobile nodes, a source and a destination, assumed to be static. Whenever a relay mobile meets the source, the source may forward a packet to it. A mobile that receives a copy of the packet from the source can forward the same only if it meets the destination (two hop routing). The source meets each relay node according to a Poisson process with a parameter λ . Each relay node meets the destination according to a Poisson process with parameter γ . The source maximizes the probability that a packet arrives successfully at a given destination by time ρ . A second transmitter, tries to jam the transmission, and hence minimizes this probability. The jammer is located close to the source. Let X_t, u_t, w_t denote, respectively, the fraction of mobiles with the message, the source's control and the jammer's control, where u_t is the probability to transmit at time t if at that time the source meets a relay, w_t is the probability of jamming at time t. If jamming and transmission occur simultaneously, then the transmitted packet is lost. Let $x_t = E[X_t]$ be the expected value of X_{t} and it is generated by

$$x_t = u_t (1 - w_t) \lambda(n - x_t), \quad x_0 = x \quad (1)$$

During the incremental time interval [t, t+dt), the number of copies of the packet in the network is X_t . Then, the number of packets the destination receives during [t, t+dt), is a Poisson random variable with $\gamma X_t dt$ as parameter. So the probability of not receiving any copy of the packet during $[0, \rho]$, conditioned on $\{X_t\}$ is

$$\exp\left(-\int_0^{\rho}\gamma X_t dt\right).$$

Let T denote the first instant a copy reaches the destination. Then the failure probability is

$$P(T > \rho) = E \left[\exp\left(-\int_{0}^{\rho} \gamma X_{t} dt\right) \right].$$

Instead of minimizing $P(T > \rho)$, the failure probability, we will minimize its upper bound, obtained using Jensen's inequality:

$$\exp\left(-E\left[\int_{0}^{\rho}\gamma X_{t}dt\right]\right).$$

Minimizing the latter is equivalent to maximizing:

$$J(u,w) \coloneqq \int_{0}^{p} \gamma x_{t} dt \cdot$$

We assume that the jammer wants to minimize this quantity and the source wants to maximize it. Let Π_c, Π_j represent the set of policies for the source and the jammer, respectively. We say that $u^* \in \Pi_c$ and $w^* \in \Pi_j$ are saddle-point (SP) policies for the game (J, Π_c, Π_j) if for every $u \in \Pi_c$ and $u \in \Pi_j$ we have

$$J(u, w^*) \le J(u^*, w^*) \le J(u^*, w)$$
.

Quantity $J(u^*, w^*)$ is called the value of the game. A policy is said to be *open loop* if it does not depend on the state of the system. It is said to be *Markov* (or a *feedback* policy) if its action at time t depends upon t as well as the state x_t . A *pure policy* is one for which the actions at all times are deterministic, i.e., either 0 or 1. Considering soft energy constraints, the source (jammer) maximizes L_u (minimizes L_w), where

$$L_{u} \coloneqq J(u, w) - \mu \int_{0}^{\rho} \mathbf{u}_{t} dt$$

and $L_{w} \coloneqq J(u, w) + \theta \int_{0}^{\rho} \mathbf{w}_{t} dt$,

which thus results in a multi-criteria game. Let

$$\mathbf{L}(\mathbf{x},\mathbf{u},\mathbf{w}) = \mathbf{J}(\mathbf{x},\mathbf{u},\mathbf{w}) - \mu \int_{0}^{\rho} \mathbf{u}_{t} dt + \theta \int_{0}^{\rho} \mathbf{w}_{t} dt.$$

Let G_{zs} be the zero-sum game (ZSG) in which the source maximizes L(x, u, w) while the jammer minimizes it. Note that L(x, u, w) is the result of either adding an extra term to L_u or subtracting a term from L_w and that the addition or the subtraction of these additional terms have

not changed the Nash equilibrium (NE) of the multi-criteria game (this argument is not valid if the control policies depend on the state, that is if they are for example feedback policies). Then clearly if (u^*, w^*) is an open-loop NE for $(L_u, -L_w)$ where both players are maximizers, it is also an open-loop NE for (L, -L), and hence an open-loop SP of L (i.e., the game G_{zs}). Likewise, any open-loop NE of $(L_u, -L_w)$. Below, we first consider a game with

$$\mathbf{L}(\mathbf{x},\mathbf{u},\mathbf{w}) = \int_{0}^{\rho} (\gamma \mathbf{x}_{t} + \mathbf{r}(\mathbf{u}_{t},\mathbf{w}_{t})) dt,$$

where $r(u, w) = -\mu u + \theta w$ and x_t is the solution of equation (1).

2. Static Game

We begin with u and w that are constants in time, in which case (1) has the unique solution:

$$x_t = n + (x_0 - n) \exp(-\lambda \kappa t)$$

where $\kappa := u(1 - w)$. Static NE is a SP of G_{zs} and has the following properties:

Theorem 1 i) If $x_0 < n$, the game has a SP.

ii) If $\gamma(n-x_0)\rho\lambda < \mu$, the game has (0,0) as the unique SP.

iii) The game cannot have a SP with w = 1.

iv) If the SP is in the open square $(0,1)\times(0,1)$, then it is unique. \diamond

3. NE of Open-Loop and Closed-Loop Dynamic Games

We first consider the open-loop case. Here, every NE is also a SP of G_z . Hence, we have a single Hamiltonian:

 $H = -\mu u + \theta w + \gamma x + pu(1-w)\lambda(n-x)$ which we will be maximizing over $u \in [0,1]$ and minimizing over $w \in [0,1]$. The co-state variable *p* satisfies the co-state equation:

$$\dot{p} = -\frac{\partial H}{\partial x} = pu(1-w) - \gamma$$

and x satisfies the original state equation (1). The open-loop SP solution is captured below:

Theorem 2 i) Let $\theta < \mu$. There exists a t_s such that $u^* = w^* = 0$ for $t > t_s$ and for $t < t_s$,

 $u^* = \theta/m(t)$ and $w^* = 1 - \mu/m(t)$, where $m(t) = p(t)\lambda(n - \xi(t))$, with p and ξ solving the coupled differential equations:

$$\dot{\xi} = \frac{\theta\mu}{p^2\lambda(n-\xi)}, \ \xi(0) = x_0, \ \dot{p} = \frac{\theta\mu}{p\lambda^2(n-\xi)} - \gamma,$$

$$p(t_s) = \gamma(\rho - t_s) \ and \ t_s \ solves \ m(t_s) = \mu.$$

$$ii) When \ \theta \ge \mu, \ there \ exists \ an \ additional threshold \ t_{\overline{s}} \ with \ t_s \le t_{\overline{s}} \ such \ that \ the \ source policy \ changes \ to \ u^* = 1 \ when \ t \in [t_s, t_{\overline{s}}].$$

$$Policy \ u^* \ for \ t \notin [t_s, t_{\overline{s}}] \ and \ w^* \ for \ all \ t \ are \ as \ in (i). \ \diamond$$

We next consider the closed-loop (feedback) case. Here we have to stay with the non-cooperative game framework, and seek for NE. Let V^{μ} and V^{ψ} be the value functions. The associated HJB equations are:

$$\frac{\partial \mathbf{V}^{u}}{\partial t} + \max_{\mathbf{u} \in [0,1]} \left[\frac{\partial \mathbf{V}^{u}}{\partial \mathbf{x}} \mathbf{u}(1 - \mathbf{w}^{*})\lambda(\mathbf{n} - \mathbf{x}) + \gamma \mathbf{x} - \mu \mathbf{u} \right] = 0$$
$$\frac{\partial \mathbf{V}^{w}}{\partial t} + \min_{\mathbf{w} \in [0,1]} \left[\frac{\partial \mathbf{V}^{w}}{\partial \mathbf{x}} \mathbf{u}^{*}(1 - \mathbf{w})\lambda(\mathbf{n} - \mathbf{x}) + \gamma \mathbf{x} + \theta \mathbf{w} \right] = 0.$$

with $V^{u}(\rho, x) \equiv V^{w}(\rho, x) \equiv 0$ (boundary conditions) and where (u^{*}, w^{*}) is a NE. The corresponding NE is the argmax and argmin of these equations and has the following structure:

Theorem 3 *i*) For any NE, $w^* < 1$ for all t. *ii*) If $\gamma \lambda$ (n-x₀) $\rho < \mu$, then $u^* \equiv w^* \equiv 0$. *iii*) Let $t_c(x) \coloneqq \rho \gamma \lambda (n-x) - \mu / \gamma \lambda (n-x)$. If $\mu - \lambda \theta (\rho - t_c(x_0)) > \theta$, the NE exists with the optimal state trajectory given by:

$$\dot{\mathbf{x}}(\mathbf{t}) = \frac{\mu \,\theta \,(\mathbf{n} - \mathbf{x}) \mathbf{1}_{\{\mathbf{t} \le \mathbf{t}_{c}(\mathbf{x})\}}}{(\rho - \mathbf{t}) \,(\gamma^{2} \lambda \,(\rho - \mathbf{t}) \,(\mathbf{n} - \mathbf{x})^{2} - \mu \,\theta)}$$

and the optimal controls are given by,

$$\mathbf{u}^{*}(\mathbf{t}) = \frac{\theta \mathbf{1}_{\{\mathbf{t} \leq \mathbf{t}_{c}(\mathbf{x}_{t})\}}}{\lambda(\mathbf{n} - \mathbf{x}_{t}) \left(\gamma(\rho - \mathbf{t}) - \frac{\mu\theta}{\gamma\lambda(n - x_{t})^{2}}\right)}$$
$$\mathbf{w}^{*}(\mathbf{t}) = \left(1 - \frac{\mu}{\gamma(\rho - \mathbf{t}) \lambda(\mathbf{n} - \mathbf{x}_{t})}\right) \mathbf{1}_{\{\mathbf{t} \leq \mathbf{t}_{c}(\mathbf{x}_{t})\}}.$$

iv) When θ is larger, the optimal policy has two switch time thresholds as in Theorem 2.iii. \Diamond

4. Conclusions

We have considered a multi-criteria control problem that arises in DTNs with two adversarial

controllers: source and jammer and two types of information structures: closed and open loop, and in the latter case also with restriction to static policies. The structure of the equilibrium is similar under both open-loop and closed-loop structures. When the jammer has a tighter constraint on its energy than the source, the policies have two switch times. After the first switch time, the jammer switches off and the source transmits at maximum probability and after the second switch time, the source also switches off. When the source has a tighter energy constraint, there exists only one switch time after which both are switched off.

Acknowledgement The first and last authors were supported by the Indo-French Centre for the Promotion of Advanced Research (IFCPAR), project 4000-IT-1. The second author's research was supported by a grant from AFOSR. Collaboration between the first and second authors was facilitated by an INRIA-UIUC exchange program.

Reference

[1] E. Altman, T. Başar and V. Kavitha, "Adversarial Control in a Delay Tolerant Network," GameSec 2010 (Conference on Decision and Game Theory for Security, 22-23 November 2010, Berlin, Germany).



Eitan Altman is a senior researcher at INRIA Sophia-Antipolis, and member of the SFR Research Institute on Science and Technology of Numerical Cultures and Society at Avignon University. He specializes on optimization, control and games for network applications.



Tamer Başar (S'71-M'73-SM'79-F'83) is a Swanlund Chair holder and CAS Professor of Electrical and Computer Engineering at the University of Illinois, Urbana-Champaign, IL, USA. He is also Research Professor with the Coordinated Science Laboratory and the Information Trust Institute. His research interests center on fundamental problems in systems, control, communications, networks, and static and dynamic games. See his web site <u>https://netfiles.uiuc.edu/basar1/www/</u> for more details.



Veeraruna Kavitha is currently a post doctoral researcher with MAESTRO, INRIA, Sophia Antipolis, France and LIA, University of Avignon, France. Her research interests span communication theory, wireless networks, signal processing, game theory, optimal control and optimization. Her website is http://www-sop.inria.fr/members/Kavitha.Voleti_Veeraruna/

ISPs and Ad Networks Against Botnet Ad Fraud

Nevena Vratonjic, Mohammad Hossein Manshaei and Jean-Pierre Hubaux School of Computer and Communication Sciences, EPFL, Switzerland { nevena.vratonjic, hossein.manshaei, jean-pierre.hubaux}@epfl.ch

1. Botnet Ad Fraud

Today, botnets (collections of software agents running autonomously and automatically, typically on compromised end users' PCs) are a very popular tool for perpetrating distributed attacks on the Internet. Botnets are a serious threat for a number of entities: end users, enterprises with online businesses, websites, Internet Service Providers (ISPs), advertisers and ad networks (ANs). Consequently, thwarting botnets would benefit everyone and would reduce the level of online crime on the Internet. However, the problem of botnets in general cannot be solved exclusively by users (lack of know-how), ISPs (too expensive to fight botnets alone), ad networks, advertisers, websites and enterprises (lack of tools and resources).

Recent initiatives propose that ISPs should perform the detection of botnets and remediation of the infected devices [1, 2]. Indeed, it is the ISPs that are in the best position to thwart the botnets. Yet, the revenues of ISPs are not (directly) affected by the botnets and ISPs would probably welcome some external funding in the efforts to fight botnets (e.g., governmentsponsored programs in Australia [2]). In the case governments are unwilling to fund these initiatives, ISPs need to find a way to make them, at the very least, cost neutral if not cost positive.

Over the last decade, online advertising has become a major component of the Web, leading to annual revenues expressed in tens of billions of US Dollars (e.g., 22.4 billion in the US in 2009 [3]). The business model of a fast growing number of online services is based on online advertising and much of the Internet activity depends on that source of revenue. Unsurprisingly, people started abusing the advertising system in various ways. Lately, it is becoming more and more popular to use botnets for ad fraud [4, 5], which creates a loss of ad revenue for advertisers, associated websites and ANs and security threats for end users (e.g., fraudulent ads that lead to phishing attacks). Therefore, ANs have economic incentives to fight botnets.

However, ANs are not in the best position to thwart botnets themselves and thus ANs might be willing to subsidize the ISPs to achieve that goal. We investigate whether ad fraud botnets are a strong enough reason for ISPs and ANs to cooperate. Such cooperation would help ISPs deploy detection and remediation mechanisms and would be a first step towards fighting all botnets.

2. System Model

We consider a system consisting of an *online advertising system*, a number of *bots* that attempt to exploit the online advertising system and an *ISP*, as depicted in Figure 1.



Figure 1. System Model: Online Advertising System, ISP and bots exploiting the advertising system.

The most prevalent model of serving online ads to end users is depicted in Figure 1. To have their ads appear with the appropriate web content, Advertisers (AV) subscribe with an ad network (AN) whose role is to automatically embed ads into web pages. Ad networks have contracts with Websites (WS) that want to host advertisements. When a User (U) visits a website (step 1) that hosts ads, while downloading the content of the web page (step 2), the user's browser will be directed to communicate with one of the Ad Servers (AS) belonging to the AN (step 3). The AS chooses and serves (step 4) the most appropriate ads to the user, such that users' interests are matched and the potential revenue is maximized.

We consider the types of ad fraud: (i) in which malware causes infected devices to display altered ads, which have been the most prominent lately [4, 5] and (ii) in which subverted users' routers modify ad traffic on-the-fly between a web server and a user [6]. When users click on the altered ads, the clicks generate revenue for

the bot master instead of the AN. Thus, the bots divert a part of the ad revenue from the AN.

3. Countermeasures

One possible approach for ANs to protect their revenue is to improve the security of the online advertising systems, thus making it more difficult for an adversary to successfully exploit those systems. For example, ad fraud can be reduced if webpages and ads are served over HTTPS instead of HTTP. The cost of implementing HTTPS at a web server includes the cost of obtaining a valid X.509 authentication certificate. Usually, website owners are not willing to bear this cost. Thus, if an AN wants the secure protocol to be deployed, it should cover the costs itself [7]. The AN may decide to selectively secure only the websites that generate sufficient ad revenue that compensates the costs.

Another possible approach for ANs to protect their revenue is to cooperate with ISPs and eliminate the major cause of the revenue loss, namely botnets. They can do so by funding the existing initiatives for ISPs to detect and remove botnets, since ISPs are in a privileged position to fight botnets. As removing botnets would benefit ANs, they have economic incentives to subsidize ISPs to thwart botnets.

4. Game-Theoretic Model

We introduce a static game to analyze the interaction between an ISP and an AN. In our model, the ISP can choose between the two actions: Abstain (A) and Cooperate (C). The AN can choose one of the following four actions: Abstain (A), Cooperate (C), Secure and Cooperate (S+C), and Secure (S). The Abstain models the ISP that is not willing to fight botnets and the AN not willing to perform any countermeasures. When cooperating, the ISP first detects the bots and then remediates infected devices, for which it receives a reward from the cooperative AN. The AN can secure the websites by choosing the action S. Finally, the AN can choose to simultaneously secure some of the websites and cooperate with the ISP to remediate some of the infected devices. The analytical results of our game-theoretic analysis are presented in details in [8].

In order to understand the implications of the analytical results in reality, we simulate the game using the real data on the ad traffic generated on the 1000 most popular websites obtained from *Compete.com*. We compute numerically the payoffs of the static game and identify the

resulting equilibria. We represent the outcomes of the game for 10^4 bots in the system in Figure 2 and Figure 3. Figure 2 shows the number of secured websites depending on the fraction λ of the revenue that bots divert from the AN. When the AN cooperates with the ISP, the fraction of remediated devices depending on the level of threat λ is shown in Figure 3. We consider three scenarios, for three different efficiencies P_D of the detection system employed by the ISP.



Figure 2. Number of the most popular websites to be secured



Figure 3. Fraction of infected devices remediated by the ISP

The obtained results illustrate that: (i) For a very low level of threat ($\lambda < 2 \cdot 10^{-4}$) no countermeasures will be taken (no websites are secured and no infected devices are remediated); (ii) When the fraction λ of the diverted revenue increases, the AN secures a number of websites that depends on λ ; (iii) Securing websites is not sufficient for an even higher level of threat (e.g., $\lambda > 10^{-4}$ for $P_D=0.1$), thus the AN will in addition cooperate with the ISP to remediate infected devices.

4. Conclusions

The game-theoretic analysis of the behavior and interactions of the ISPs and ANs facing botnet ad fraud shows that cooperation could emerge under

http://www.comsoc.org/~mmc/

certain conditions that mostly depend on: (i) the number of infected devices (ii) the aggregate power with which bots divert revenue from the AN and (iii) the efficiency of the botnet detection system. The cooperation is a win-win situation where: (i) users benefit from the ISP's help in maintaining the security of users' devices; (ii) the AN protects its ad revenue as the botnet ad fraud is reduced; (iii) it is at least cost neutral, if not cost positive for the ISP to fight botnets. Cooperation between the AN and the ISP would help to reduce the level of online crime and improve the Web security in general.

References

[1] Jason Livingood, N. Mody, Michael O'Reirdan and Comcast Communications. Recommendations for the Remediation of Bots in ISP Networks. IETF, 2009.

[2] Jason Livingood, N. Mody, Michael O'Reirdan and Comcast Communications. ISP Voluntary Code of Practice for Industry Self-regulation in the Area of esecurity. Internet Industry Code of Practice, 2009.

[3] Internet Advertising Revenue Report. Interactive Advertising Bureau, 2009.

[4] Click Forensics Discovers Click Fraud Surge from New Sophisticated Bahama Botnet. http://www.clickforensics.com/newsroom/pressreleases/144-bahama-botnet.html

[5] Viral Web Infection Siphons Ad Dollars from Google.

http://www.theregister.co.uk/2009/05/14/viral_web_in fection

[6] Nevena Vratonjic, Julien Freudiger and Jean-Pierre Hubaux. Integrity of the Web Content: The Case of Online Advertising. In Usenix CollSec '10: Workshop on Collaborative Methods for Security and Privacy, Washington, DC, USA, 2010. ACM.

[7] Nevena Vratonjic, Maxim Raya, Jean-Pierre Hubaux and David C. Parkes. Security Games in Online Advertising: Can Ads Help Secure the Web? In WEIS '10: Workshop on Economics of Information Security, Cambridge, MA, USA, 2010.

[8] Nevena Vratonjic, Hossein Manshaei, Maxim Raya, and Jean-Pierre Hubaux. ISPs and Ad Networks Against Botnet Ad Fraud. In GameSec' 10: Proceedings of the First Conference on Decision and Game Theory for Security, Berlin, Germany, 2010.



Nevena Vratonjic is a PhD student in the Laboratory for Computer Communications and Applications (LCA) at EPFL. Prior to pursuing a PhD degree, she completed one year of the Master Program in Computer Science at EPFL in 2007. She obtained a MSc degree in Communication Systems and a BSc degree in Electrical Engineering from University of Belgrade, Serbia in 2006. Her research interests are in the area of network and system security. Her current research work focuses on online advertising frauds and security of online ad serving systems. For more details, see http://people.epfl.ch/nevena.vratonjic



Mohammad Hossein Manshaei earned his B.Sc. degree in electrical engineering and his M.Sc. degree in communication engineering from the Isfahan University of Technology (IUT). Iran, in 1997 and 2000, respectively. He earned another M.Sc. degree in computer science and his Ph.D. in computer science and distributed systems from the University of Nice Sophia-Antipolis, France, in 2002 and 2005, respectively. He completed his thesis work at INRIA Sophia-Antipolis. He currently works as a senior researcher and lecturer at the Laboratory for Computer Communications and Applications (LCA) in EPFL. His research interests include wireless networking, security and privacy, social networks, cognitive radios, and game theory. For more details, see

http://people.epfl.ch/manshaei

13/43



Jean-Pierre Hubaux has been a faculty member at EPFL since 1990. His current research activity is focused on privacy preservation mechanisms in pervasive communications. He has authored and co-authored more than 150 publications in networking, network security and privacy. In 2008, he completed a graduate textbook entitled "Security and Cooperation in Wireless Networks", with Levente Buttyan. He held visiting positions at the IBM T.J. Watson Research Center and at UC Berkeley. He is a Fellow of both ACM and IEEE. For more details, see <u>http://people.epfl.ch/jeanpierre.hubaux</u>

When Do Firms Invest in Privacy-Preserving Technology?

Murat Kantarcioglu, Alain Bensoussan, and SingRu(Celine) Hoe University of Texas at Dallas, USA {muratk, alain.bensoussan}@utdallas.edu, hoceline02@yahoo.com

1. Introduction

Privacy is a central concern in the information age. In many cases, individuals consider privacy issues when deciding whether to use a firm's services. A California Health Foundation Survey [5] shows that one of the barriers to consumer acceptance of digitized personal health records is the fear that their private information may not be adequately protected. At the same time, different customers could demand different levels of privacy protection. For example, in medical research domain, HIV positive patients may be more sensitive about their privacy. On the other hand, a person who is HIV positive could be more valuable to medical researchers working on a vaccine for HIV. Given the customers' demand for privacy protection, entrepreneurs need to decide whether to invest in privacy preserving technologies or not. We analyze how the relationship between a customer's valuation of his private information and a customer's profitability to a firm affects a firm's investment decision and formulate guidelines for entrepreneurs' adoption decision.

2. The Model

We view firms' project valuation processes as a Stackelberg leader-follower game under complete information with customers taking the role of the follower.

For each individual customer, we define his utility function as:

$$U(x_i, D) = (2D - 1) \left(\alpha_P - \left(a \times \frac{|x_i - \mu_X|}{\sigma_X} + b \right) \right)$$

where x_i represents customer *i*'s valuation of his private information, $D=\{0,1\}$, μ_X and σ_X represent the mean and the standard deviation of X (a random variable related to customers' valuation of private information) respectively, both a>0 and b>0 (a basic level of privacy related to values of private information common to the population) are constants, and α_P is the privacy-protection level associated with the privacy-preserving technology P. Indeed, the terms included in the last parenthesis measures customer *i*'s fair level of privacy. This specification is motivated by [2] in which the authors show that an individual in a group would demand a higher value of private information if his trait deviates from that of the average population segment in the group, and the further his trait is away from that of average population segment, the higher the value of private information is demanded.

Given the firm's adoption of the technology P, an individual customer chooses his utility maximization strategy. Obviously, rational customers would choose D=1 if $\alpha_P \ge a \times \frac{|x_i - \mu_X|}{\sigma_X} + b$

and D=0 otherwise. That is, a customer's optimal strategy in response to his utility maximization solution relies solely on his value of private information and privacy protection that the firm's privacy-preservation technology can provide. We define such a rule as a customer's decision function given:

$$D(x_i) = 1_{\substack{\alpha_P \ge a \times \frac{|x_i - \mu_X|}{\sigma_X} + b.}}$$

Once customers' optimal decisions have been solved, the firm integrates customers' decisions into their valuation. The expected project value is the expected value that the firm can receive from customers under the joint distribution function of two random variables, customers' valuation of private information (*X*) and customers' profitability to a firm (Y), denoted by $F_{XY}(x, y)$. We have the formula given:

$$V(x, y, D(x)) = \int_{X \times Y} y D(x) dF_{XY}(x, y).$$

From the formula, a clearly defined joint distribution function is required for valuation. We propose to employ copula functions for the bivariate joint distribution function. This allows us to study project values through a richer class of joint distribution functions fitting into different market segments. It also allows us to investigate the impact of dependence structure on project value.

In return for the project value, the firm needs to pay a fixed investment *K*. Clearly, the market mechanism makes it the case that the firm will undertake the investment if and only if $V(x, y, D(x)) \ge K$.

3. Valuation Results from Utilizing Copula

If *X* and *Y* are independent, then:

 The firm's investment tendency in privacypreserving technologies increases with increases in the mean of customers' profitability to a firm.
 The smaller the "adjusted" weight *a* and the

basic level of privacy b, the greater the probability that firms will make the investment.

(3) Firms would likely invest in privacypreserving technologies requiring significant investment costs if government takes intervention.

The normal distribution is often used to describe any variable that tends to cluster around the mean, and this may well be the case in our two variables. If $F_{XY}(x,y)$ is a binormal distribution, we find that the correlation structure does not affect valuation, and the investment rule prediction is exactly the same as that proposed in the independent case.

We then consider the exponential distribution extensively used due to its mathematical simplicity. If X and Y are from exponential distributions with Gaussian correlation structure, then:

(1) The mean and the standard deviation of customers' valuation of private information are irrelevant to the valuation.

(2) The correlation ρ (the Pearson correlation) affects the valuation where the firm's expected revenue is negatively related to ρ .

(3) The necessity of government intervention is positively related to ρ .

We next exploit the case where two random variables are from Pareto distributions. The Pareto distribution shows rather well in describing the allocation of wealth among individuals. In some cases, customers' profits to the firm may be well correlated to individual wealth, for example the usage of some banking services. If *X* and *Y* are from pareto distributions with Gaussian correlation structure, then:

(1) The impact of ρ is similar to the exponential distribution with Gaussian dependence structure with an additional property that given ρ , the positivity results in better valuation.

(2) Firms have stronger tendency to invest in privacy-preserving technologies if the distribution of customers' valuation of private information is less volatile, even linking to a smaller mean.

By Pareto principle, which says that 20% of the

population controls 80% of the wealth, we may expect that the mean of total wealth is not large. However, the extremely wealthy person may be very profitable to the firm, bringing huge project values. Therefore, the second observation may explain the existence of mythical Swiss numbered bank account where an account is only identified with a number to provide increased anonymity to its user.

4. Conclusions

The mean of customers' profitability to a firm is always positively related to the firm's revenues. However, the impact of the volatility of customers' profitability to a firm is inconclusive. In addition, the impact of the mean and the volatility of customers' valuation of private information varies with two underlying marginals. The dependence structure under Gaussian copulas exhibits different effect for different underlying univariate marginals.

Three cases, which are the independent case and the binormal and the biexponential under Gaussian dependence structure, suggests that the government intervention may be required to have firms invest in privacy-preserving technologies. This seems to support the observed phenomenon that in general we do not see that firms are motivated to invest in privacy-preserving technologies requiring significant costs.

For appropriate valuation and investment decision making, we suggest that firms should be cautious about estimating underlying univariate distributions and dependence structures. If distribution validation is not empirically possible, firms should proceed with distributions and dependence structures practically justifiable.

References

[1] A. Acquisiti and J. Grossklags. Losses, Gains, and Hyperbolic Discounting: An Experimental Approach to Information Security Attitudes and Behavior. Proc. 2nd Int'l Workshop Economics and Info. Security, 2003.

[2] B.A. Huberman, E. Adar, and L.R. Fine. Valuating Privacy. IEEE Security & Privacy, 3(5):22–25, Sep.-Oct. 2005.

[3] M. Kantarcioglu, A. Bensoussan, and C. Hoe.
When Do Firms Invest in Privacy-Preserving Technology. GameSec 2010 - Volume 6442, edited by Tansu Alpcan, Levente Buttyan, and John Baras, 2010.
[4] J. Kleinberg, C.H. Papadimitriou, and P. Raghavan, "On the Value of Private Information," Proc. 8th Conf. Theoretical Aspects of Rationality and Knowledge (TARK-2001), J. van Benthem, ed., Morgan

Kaufmann, 2001, pp. 249-257.

[5] "Consumers and Health Information Technology: A National Survey," California Health Foundation, April 2010.



Murat Kantarcioglu is an Assistant Professor in the Computer Science Department and Director of the UTD Data Security and Privacy Lab at the University of Texas at Dallas. He holds a B.S. in Computer Engineering from Middle East Technical University, and M.S. and Ph.D degrees in Computer Science from Purdue University. He is a recipient of NSF Career award and Purdue CERIAS Diamond Award for Academic excellence.

Dr. Kantarcioglu's research focuses on creating technologies that can efficiently extract useful information from any data without sacrificing privacy or security. His research has been supported by grants from NSF, AFOSR, ONR, NSA, and NIH. He has published over 60 peer reviewed papers. Some of his research work has been covered by the media outlets such as Boston Globe, ABC News etc. and has received two best paper awards.



Alain Bensoussan is Ashbel Smith Professor and the Director of ICDRiA at the University of Texas at Dallas. He is also Chair Professor of Risk and Decision Analysis at the Hong Kong Polytechnic University and World Class University Distinguished Professor at Ajou University. He is Professor Emeritus at the University of Paris Dauphine and has an extensive research background in stochastic control, probability and stochastic processes.

Professor Bensoussan served as President of National Institute for Research in Computer Science and Control (INRIA) from 1984 to 1996; President of the French Space Agency (CNES) from 1996 to 2003; and Chairman of the European Space Agency (ESA) Council from 1999 to 2002. He is a member of the French Academy of Sciences, French Academy of Technology, Academia Europae, and International Academy of Astronautics. His distinctions include IEEE Fellow, SIAM Fellow, Von Humboldt award, and the NASA public service medal. Professor Bensoussan is a decorated Officer of Legion d'Honneur, Commandeur Ordre National du Merite and Officer Bundes Verdienst Kreuz.



Celine Hoe is a post-doc researcher at the University of Texas at Dallas. She holds a master's degree from George Washington University and a Ph.D. degree from the University of Texas at Arlington.

Access Control: Allocating Resources to Selfish Agents

Farzad Salim, Jason Reid, Uwe Dulleck, Ed Dawson Queensland University of Technology, Brisbane, Australia {f.salim, jf.reid, uwe.dulleck, e.dawson}@qut.edu.au

1. Background and Motivations

The ultimate goal of an authorisation system is to allocate each user the level of access they need to complete their job - no more and no less. This proves to be challenging in an organisational setting because on one hand employees need enough access to perform their tasks, while on the other hand more access will bring about an increasing risk of misuse - either *intentionally*, where an employee uses the access for personal benefit, or *unintentionally* through carelessness, losing the information or being socially engineered to give access to an adversary [2].

The approaches to access control have traditionally assumed a command and control model and focused on technological means (i.e., policy models, policy enforcement and usage monitoring tools) to ensure employees cannot misuse their assigned privileges. At the core of these approaches is the assumption that a security policy can be constructed a priori and maintained correctly (by a security officer). However, this proves to be rather challenging in a real organisation for three main reasons [3]. First, various tasks are usually performed by different individuals, each of whom has special information concerning his particular sphere of activity. Usually this information is not available to other individuals (security officers) to correctly determine what resources users may need to complete their jobs. Even if such information is available, the dynamism and uncertainty that exists in today's organisations makes maintaining the correctness of the policy specifically challenging. Second, there may be divergence of preferences between the action an employee considers optimal and what action is optimal for the organisation. This divergence arises because when making a decision agents also seek to maximize their own self-interest; this becomes particularly problematic when employees draw personal benefit from misusing the organisation's resources. Finally, considering the widespread use of digital resources and their non-rivalrous property, full monitoring and verification of employees' resource usage has become impossible or prohibitively costly.

So far, little attention has been directed towards addressing the above mentioned complexities. We are interested in context-based approaches to access control that explicitly take into account, and reason about users' *incentives to misuse*' a privilege and the *expected consequences* of such misuse while making an authorisation decision. Through this, it is not only the sensitivity of the resources that is taken into account, but also the users' propensity to misuse, which may be influenced by adjusting factors such as the accuracy of misuse detection mechanisms, introducing random monitoring/audit and employing punishment (or reward) for detected misuses (for abiding by the policy).

With the goal of developing a more dynamic authorisation model, we have adopted a game theoretic framework to reason about the factors that may affect users' likelihood to misuse a permission at the time of an access decision. Game theory provides a useful but previously ignored perspective in authorisation theory: the notion of the user as a self-interested player who selects among a range of possible actions depending on their pay-offs.

2. Game Theoretic Analysis of Authorisation

We formulate an authorisation mechanism as a game between a *selfish employee*¹ and the *benevolent employer* who is the sole authority in making authorisation decisions [1]. The game starts with a request from the employee for access to a resource. Along with the request, the employee indicates the anticipated outcome of such action for the employer, denoted as *proposal* (**p**). Given this request, the employer shall decide whether to *authorise* or *deny* the access to the resource. On the other hand, the strategy space of the employee consists of either *attack*.

¹We assume the utility function of selfish employees (as opposed to malicious ones) is not inversely proportional to the organisation's utility function - they are considered to be profit maximising agents with no concern about the organisation's payoffs.

The game centers around a resource valuable to both players. An employee may use the resource to either make a personal profit (i.e., attack) or perform a job that actions the proposal (\mathbb{P}) for the benefit of the employer. The employer, hence, is concerned about the expected cost of the attack ($\mathbb{P}_{\mathbb{P}}$) (e.g., the cost of a potential privacy breach).

The employer is also susceptible to opportunity cost: the benefit forgone by denying a request. The quantification of opportunity is determined by the proposal (p) made by the employee to access the resource. Through such a formulation, distinct from the existing authorisation models, denying an access request as well as authorising it may incur a cost for the authorisation system.

Now consider the employee's cost factors. An employee may incur cost through fines, denoted by (c_f) (i.e., given they attack). However, usually a fine is not certain - it is only applicable if the employer can detect the attack, which is a function of the accuracy of detection techniques and the ability to enforce the fine. For now the ability to detect and enforce the fine is combined and referred to as the probability of being fined, denoted by $\Psi \subseteq [0,1]$ which is assumed to be common knowledge. For example, when an employee is out-sourced from another country there may be less chance of enforcement of the fine in comparison to a circumstance where the employee is local. In addition, in order to attack, the employee is assumed to incur a preparation cost, denoted by $(\mathbf{c}_{\mathbf{r}})$. This cost abstracts the effort the employee must expend in order to acquire access to the resource to use it for personal benefit. For instance, if the resource is commercially valuable, finding a buyer requires time and effort. In other cases, the employee may need to prove to the employer that the proposal amount is attainable by him and this could require training courses and faking trustworthiness.

Sometimes the employee is given a personal benefit for the opportunity they realize. This is represented as a rate of return, $\epsilon \in [0,1]$ on the proposed opportunity, p. We regard the predictions of the employee in terms of the actual achievement of p to be always correct if the access is granted. On the other hand, the actual personal profit for the employee from an attack is a portion $\alpha \in [0,1]$ of the cost of the resource (ϵ_r) , if the access is given. Note that

this may not always be the case as sometimes a very costly resource for the employer has a very low value for a selfish employee or vice versa.

Given the above game setting, the payoff for the players is as follows. The authorisation problem is, given complete information about the payoffs, when should the employer authorise the access?



Assuming that the employer always takes the pure strategy deny, then the employee's best response is not attack. However, this is not an equilibrium as the pure strategy of not attack by the employee motivates the rational employer to change his strategy to *authorise* whenever p > 0. By switching to authorise, the employee only attacks when:

$$\epsilon \cdot p < (1 - \psi) \alpha \cdot c_r - \psi \cdot c_f - c_t$$

This finding is interesting and rather counter intuitive in the context of authorisation. It suggests that in making an authorisation decision the authorisation mechanism may only need to focus on the employee's payoff instead of its own. This is contrary to the existing approaches to authorisation, where there exists a policy, assumed to incorporate access rules which result from a tradeoff analysis between some implicit contextual factors, for all current and future requests. Here instead, the decision factors are explicit and abstract enough to adapt to the required application. For example. the probability of misuse detection can depend on the existing monitoring techniques, audit, accuracy of forensic techniques or physical security employed.

The application of such an authorisation model can provide incentives to users such that they don't misuse the permissions, rather than simply providing authorise/deny responses. Through such an interpretation, given an authorisation request, a game theoretic authorisation model may also attempt to meet the above inequality, through taking either or a combination of *deterrence* or *appeasement* incentive policies.

The former aims to increase the cost of attacking for the user, so that the above inequality is met. This may be achieved through either increasing any or a combination of ψ , c_{f^*} , c_t or by reducing α . On the other hand the appeasement policy attempts to increase the benefit for not attacking, by align-ing users' utility function with the organisation through increasing $\boldsymbol{\varepsilon}$.

4. Conclusions

The problem of authorisation is at its core analogous to the principal and agent problem studied in the field economics. The adoption of a game theoretic framework provides a new paradigm of thinking for designing dynamic authorisation models. The models that emerge from this approach attempt to construct incentive mechanisms such that users observe their best alternative as the action the organisation deems optimal. This is where traditional access control solutions to user compliance primarily differ from game theoretic solutions - the former approaches must rely on often-unrealiable policy enforcement mechanisms to force users' compliance whereas the later pragmatically engages their self-interest.

References

[1] Farzad Salim, Jason Reid, Uwe Dulleck, and Ed Dawson. Towards a game theoretic approach to authorisation. In *Decision and Game Theory for Security (GameSec)*, volume 6442 of *Lecture Notes in Computer Science*, pages 208–219, Springer, 2010.

[2] Farzad Salim, Jason Reid, and Ed Dawson. Towards authorisation models for secure information sharing: A survey and research agenda. *The ISC International Journal of Information Security* (*ISeCure*), 2:67–85, 2010.

[3] Sara Sinclair and Sean W. Smith. What's wrong with access control in the real world? IEEE Security & Privacy, 8(4):74–77, 2010



Farzad Salim is a PhD candidate at the Information Security Institute – Queensland University of Technology. Prior to this he was a Research Fellow at the School of Computer

Science and Software Engineering, University of Wollongong, where he also obtained a Master's degree in Computer Science. His research interests include access control, privacy and especially the application of ideas and techniques from the field of economics to information security.



Jason Reid holds a PhD in the area of trusted computing and distributed systems security. Since 1999, he has worked for the Information Security Institute – Queensland University of Technology where he holds the position of Senior Research Fellow. His research interests and areas of expertise include access control, trusted systems, trusted computing and smart card security as well as the privacy implications of these technologies. He has extensive research and consulting experience in such fields as authentication and identity management and general network security.



Uwe Dulleck is Professor of Economics at Queensland University of Technology. He published several papers on Game Theory and Information Economics in leading Economics journals. His research on Expert Services (Credence Goods) was covered by the *The Economist* in an Economics Focus. He is one of the Editors of Economic Analysis and Policy.



Ed Dawson is a Professor Emeritus in the In formation Security Institute (ISI) at Queensland University of Technolgy (QUT). Prior to this he was the Research Director of the ISI and has been a member of the academic staff at QUT since 1974. Professor Dawson has extensive research experience in many aspects of cryptology and its applications. He has published over 250 referred research papers, supervised 25 PhD students to completion and received numerous research grants. He is on the editorial board of several journals and has chaired several conferences on various aspects of information security.

SPECIAL ISSUE ON TECHNOLOGY ADVANCES

Distributed Computing for Multimedia Applications

Guest Editor: Jian Tan, IBM T.J. Watson Research, USA, tanji@us.ibm.com

Multimedia applications witness challenges and opportunities with new technologies and frameworks emerging in distributed computing paradigms recently. In this special issue, we invite five groups of top-notch researchers to share with us their interesting and inspiring work on a broad range of topics, including efficient scheduling algorithms and resource management tools for MapReduce, performance in virtualized cloud, information security for distributed computing and large-scale video concept detection using parallel computation.

MapReduce provides a framework for processing huge datasets by executing parallel computations on a large cluster of computers. Due to its simplicity and flexibility, this model becomes highly popular recently, e.g., for multimedia applications. Scheduling possibly a large number of running jobs with different constraints and managing varying workloads in such a shared environment is quite challenging. In this issue, we have two articles investigating new ways to improve the MapReduce performance on top of the framework Hadoop. The article "Flexible MapReduce Scheduling" presents a newly implemented scheduler called FAIR, which is based on a malleable packing scheme for assigning tasks and can be tailored to a suitable metric out of a number of candidates for better performance, differing from existing ones that usually focus on specific metrics. The article "MapReduce Workload Management & Multimedia Applications" contains an enjoyable introduce to MapReduce as well as some applications related to multimedia services. In addition, it presents another scheduler for management of accelerated MapReduce workloads in heterogeneous clusters using highlevel performance goals and more fine-grained resource management.

Cloud computing emerges as a promising new paradigm for computing and storage. To encourage migrating applications from traditional platforms to cloud, one key issue that must be addressed is the system performance. For enterprise users, this is especially important because of the concern regarding the quality of service after moving to a virtualized cloud. For example, the unstable network can degrade the performance of many data intensive applications, e.g., for multimedia applications such as video processing and media hosting. The article "Understanding Network Communication Performance in Virtualized Cloud" presents extensive measurements on Amazon EC2 cloud and discusses the engineering implications. The observations from this report can help system designers to improve the stability of virtualization infrastructure.

Given the proliferation of web applications, the power of these new distributed computing platforms gives rise to a critical issue on information security. Specifically, how to help developers to avoid bugs or mistakes that may unintentionally leak sensitive information out of the system generates lots of research activities. To this end, the article "Information Security in Computing" Distributed introduces а programming methodology that constructs a controlled environment for the information flow. The appealing characteristic of this technique is that it can capture and enforce good data protection during the whole development. With improving security, the demands for these information services can be further enhanced.

With the ease of provisioning multiple virtual machine instances on cloud, designing efficient algorithms suitable for parallel distributed computation draws increased attention. The article "Parallel Lasso for Large-Scale Video Concept Detection" proposes an algorithm to reduce time and space complexities for large scale video concept detection. This approach is based on parallel incomplete Cholosky factorization for preprocess and parallel primal-dual interior point method for parameter optimization. These types of parallel algorithms are particularly suitable for a cloud environment.

Research in multimedia computing and networking has seen both challenges and opportunities with emerging technologies and paradigms. We would like to thank all the authors for their contributions and hope these articles can stimulate further research interests in this area.



Jian Tan received the B.S.E. and Ph.D. degrees in electrical engineering from University of Science and Technology of China in 2002, and Columbia University, New York in 2008, respectively. After that, he spent one year working as a postdoctoral researcher at the Networking and Communications Research Lab in the Ohio State University. Then, he joined IBM T.J. Watson Research Center, New York, in the System Analysis and Optimization Group. He was with Lucent Bell Laboratories, NJ, during the summers of 2005 and 2006, and with Microsoft Research, Cambridge, UK, in the winter of 2007. His research interests include mathematical foundations as well as related applications for information networks and complex service systems. His current research focuses on cloud computing, MapReduce Hadoop and massive storage systems.

Dr. Tan was awarded the Eliahu Jury Award for his Ph.D. thesis work from Columbia University. He received the Best Student Paper Award at the 20th International Teletraffic Congress and the Best Paper Award at the 3rd IEEE International conference on Distributed Computing in Sensor Systems. He serves and has served as TPC members of various conferences, such as INFOCOM and WASA.

Flexible MapReduce Scheduling

Joel Wolf, Deepak Rajan, Kirsten Hildrum, Rohit Khandekar, Vibhore Kumar, Sujay Parekh, Kun-Lung Wu and Andrey Balmin, IBM Research {jlwolf, drajan, hildrum, rohitk, vibhorek, sujay, klwu, abalmin}@us.ibm.com

1. Introduction

Google's MapReduce [1] and its open source implementation Hadoop [2] have become highly popular in recent years. There are many reasons for this success: MapReduce is simple to use. It is automatically parallelizable, naturally scalable, and can be implemented on large clusters of commodity hardware. Important built-in features include fault tolerance, communications and scheduling. In addition, open source implementations are available.

We focus in this paper on the problem of scheduling MapReduce work. The original scheduler implemented in Hadoop was First In, First Out (FIFO). But FIFO has well-known fairness problems, which arise because a large job can starve subsequently arriving small jobs. The Hadoop Fair Scheduler (FAIR) [3] was motivated by this issue. In particular, it ensures that each job get at least some minimum number of slots, the basic unit of resource in a MapReduce cluster. And it distributes any remaining slots (the so-called *slack*) in an equitable manner. But while FAIR demonstrates its performance in terms of standard scheduling metrics such as average response time, it makes no actual attempt to optimize any such metrics. (Common examples of scheduling metrics [4] include average response time, makespan, average or maximum stretch, lateness, tardiness and SLA costs. In computer science applications the first four and the last are typically most important. Stretch is actually the truest measure of fairness.) It is worth noting that MapReduce schedules designed to optimize one metric will generally be quite different from those designed to optimize another.

2. FLEX

In this short paper we describe a different MapReduce Scheduler known as FLEX. The goal is to optimize any of a wide variety of standard scheduling theory metrics while ensuring the same minimum job slot guarantees as in FAIR. Thus FLEX will also be fair in the sense of [3], avoiding job starvation just as FAIR does. Further, by an appropriate choice of scheduling metrics, starvation can be avoided

even when no minimums are given. The metrics can be chosen by a system administrator on a cluster-wide basis, or even, perhaps, by individual users on a job-by-job basis. And these metrics can be chosen from a menu that includes all the previously mentioned alternatives. The metrics can be weighted or not, and one can minimize either the sum (or average, equivalent from the perspective of optimization) of all the individual job metrics, or the maximum of all of them. Because of these combinatorial alternatives, FLEX can optimize any of 16 separate metrics, and is thus, as the name implies, quite flexible.

We have implemented FLEX as an extension of FAIR, basically only overriding its calculation of minimum slot allocations to each job. Thus we make use of FAIR's extensive infrastructure while computing more intelligent schedules. In fact, FLEX can be regarded as an add-on module that works synergistically with the Hadoop FAIR scheduler.

Below we give a high level FLEX overview, and a brief indication of its performance. See [5] for further details.

3. FLEX Overview

FLEX works identically for either the Map or the Reduce phases, so we describe it in a phaseagnostic manner. The FLEX algorithm builds a so-called *malleable* schedule [6], which basically means that the number of slots allocated to each job may vary throughout the schedule. See Figure 1, on which the x-axis represents the slots and the y-axis represents time. (In a more standard description of malleable scheduling the x-axis would represent processors rather than slots.) There are 4 jobs of various colors. Note how their slot allocations change over time. (*Moldable* schedules are special cases in which the number of slots do not vary by job.)

FLEX depends on two key ideas, described for the sake of exposition in reverse order of their actual execution.



Figure 1. A Malleable Schedule

First, given any particular "priority" ordering of the jobs we devise a *Malleable Packing Scheme* (MPS) to compute a high quality schedule. In fact, it turns out that for any of our possible metrics, there will exist a priority ordering of the jobs for which the MPS schedule will actually be optimal. (The priority ordering is strictly an input to the packing scheme. There is no guarantee, for example, that the jobs will complete in that order in the MPS schedule.)

Second, finding a priority ordering that produces an optimal MPS solution is, unfortunately, difficult. So we content ourselves instead with finding a "high quality" ordering for any of our possible metrics. We do this in one of two ways: The first approach is to employ an essentially generic algorithm to solve a so-called Resource Allocation Problem (RAP) [7]. The solution to this RAP will actually be optimal in the context of moldable scheduling [6], assuming positive minima for each job. In terms of malleable schedules the solution will not be optimal. The second approach, which we will not describe further, involves schemes which are tailored to the selected metric. (For the important case of average response time we have designed 4 such algorithms.)

Combining all of this into a single algorithm, FLEX creates a generic ordering and at most a modest number of specific orderings, feeding each of them to MPS. The final output is the best solution found. Next we provide summary descriptions of MPS and the generic RAP schemes.

Malleable Packing Scheme: Given a priority ordering, the scheme proceeds iteratively. At any iteration a current list of remaining jobs is maintained, ordered by priority. To start, time is initialized to 0. The current list is initialized to be

all of the jobs, and one job is removed as each iteration completes. The number of slots allocated to a given job may vary from interval to interval, thus producing a malleable schedule.

Any given iteration of the algorithm involves the following steps: First, the scheme allocates the minimum number of slots to each job in the list. This is feasible, since the minima have been normalized, necessary, during if precomputation step. After allocating these minima, some slack may remain. The idea is to allocate the remaining allowable slots to the jobs in priority order. The first several may get their full allocations, and those jobs will be allocated their maximum number of slots. But ultimately all slots may get allocated in this manner, leaving at most one job with a ``partial" remaining allocation of slots, and all jobs having lower priority with only their original, minimum number of slots. Given this set of job allocations, one of the jobs will complete first. (Ties may be adjudicated in priority order.) Now this job is removed from the list, and the necessary bookkeeping is performed to compute the remaining work for those jobs remaining in the list. We repeat until all jobs have been removed from the list.

Generic RAP Schemes: The good news is that regardless of the specific metric chosen, finding the optimal moldable schedule can be formulated as an RAP. In the first category (minisum), we wish to minimize the sum of each metric. In the second category (*minimax*), we wish to minimize the maximum of each metric. Such RAPs are called *separable*. Minisum separable RAPs can be solved exactly by one of several means, the key being whether or not the functions involved are convex. (Response time and lateness are convex, and the other metrics are not.) Convex minisum separable RAPs can be solved by a very fast scheme, because for such problems greedy algorithms are exact. Non-convex minsum separable RAPs can be solved by a slower dynamic programming algorithm. All of the minimax separable RAPs can be transformed into convex minisum problems and solved by the faster algorithm. See [5] or [7] for further details.

4. FLEX Performance

Figures 2 and 3 show FLEX performance compared with FAIR and FIFO for 8 minisum metrics and 8 minimax metrics, respectively. In both cases we have overlaid the average and worst case ratios (in 100 experiments) for each

metric and scheme, and normalized by the optimal solution found by brute force. Note how well FLEX does relative to the other schemes, and also relative to optimal, particularly for key metrics. See [5] for additional experiments.



Figure 2. Minisum Metrics



Figure 3. Minimax Metrics

References

[1] J. Dean, S and Ghemawat. Simplified Data Processing on Large Clusters, ACM Transactions on Computer Systems 51(1), pages 107-113,2008, New York, NY, USA, 2009.

[2] Hadoop. http://hadoop.apache.org.

[3] M. Zaharia, D. Borthakur, J. Sarma, K. Elmeleegy, S. Shenker and I. Stoica. Job Scheduling for Multi-User Mapreduce Clusters. Technical Report EECS-2009-55, UC Berkeley, 2009.

[4] M. Pinedo. Scheduling Theory, Algorithms and Systems. Prentice Hall, 1995.

[5] J. Wolf, D. Rajan, K. Hildrum, R. Khandekar, V. Kumar, S. Parekh, K.-L. Wu and A. Balmin, FLEX: A Slot Allocation Scheduling Optimizer for MapReduce Workloads. Proceedings, Middleware, Bangalore, India, 2010.

[6] J. Leung. Handbook of Scheduling. Chapman and Hall / CRC, 2004. [7] T. Ibaraki and N. Katoh. Resource Allocation Problems. MIT Press, 1988.



Joel Wolf received his Ph.D. from Brown University and his Sc.B. from MIT, both in mathematics. He is currently a research staff member at the IBM T.J. Watson Research with Center. interests in mathematical optimization. He has won 2 Outstanding Innovation and 1 Outstanding Technical Achievement awards at IBM, was made a Master Inventor at IBM, and is a Fellow of the IEEE. He has also been an Assistant Professor of Mathematics at Harvard University, as well as a Distinguished Member of Technical Staff and manager at Bell Laboratories.



Deepak Rajan received his B.Tech in M.E. from IIT, India, and his M.S. and Ph.D. from UC Berkeley in Industrial Engineering and Operations Research, His thesis was titled "Designing capacitated survivable networks: Polyhedral analysis and algorithms". His research interests are mainly in the areas of scheduling and network problems, with underpinnings in integer theoretical programming, approximation algorithms, and computational optimization. Recently, he has been working on a variety of scheduling problems arising from distributed computing systems.



Kirsten Hildrum received her Ph.D. from UC Berkeley and her undergraduate degrees in Math and Computer Science from the University of Washington in Seattle. Since joining IBM's T.J. Watson Research Center, she has worked on several parts of IBM's InfoSphere Streams. She is particularly interested in theoretical versions of problems that arise in practice.



Rohit Khandekar received his B.Tech. and Ph.D. degrees in Computer Science from Indian Institute of Technology, Delhi, India. He is currently a Research Staff Member at IBM T. J. Watson research center. His research interests include theoretical and practical aspects of algorithms, combinatorial optimization and mathematical programming.



Vibhore Kumar is a Research Staff Member at IBM's T. J. Watson Research Center. He received his B.Tech in Computer Science & Engineering from IT-BHU, India and his M.S. and Ph.D. in Computer Science from Georgia Institute of Technology, USA. His research interests include large-scale data analysis systems, stream processing systems, autonomic computing and distributed systems. Currently, his research is focused on building systems that enable large scale data analysis.



Sujay Parekh is an Advisory Software Engineer at IBM T.J. Watson Research Center. He received his Ph.D. and M.S. in Computer Science from the University of Washington and his B.S. in Computer Science from Cornell University. He is a co-author of the book "Feedback Control of Computing Systems" and has published over 40 journal articles and conference publications. His research interests include distributed systems, large scale analytics, and resource and performance management.



Kun-Lung Wu is the manager of the Data-Intensive Systems and Analytics Group at IBM T. J. Watson Research Center. He is also the development manager of the InfoSphere Language and Developer Tools Team, Information Management, IBM Software Group. The combined research and development team has been engaged since 2009 in the research and product development of the IBM InfoSphere Streams - a large-scale distributed stream processing middleware. In addition to product development, the team also conducts a wide range of research issues in data-intensive systems and analytics -- including programming languages, compiler and various optimization techniques for distributed stream processing: advanced analytic algorithms and applications for stream applications; job management, scheduling, resource management and system optimization for large-scale distributed systems.



Andrey Balmin joined IBM Almaden Research Center as a Research Staff Member,

after receiving his Ph.D. degree in computer science from UC San Diego. His main area of expertise is search and query processing of semistructured and graph-structured data. He received an Outstanding Technical Achievement award at IBM for development of the cost based optimizer of DB2 pureXML.

MapReduce Workload Management & Multimedia Applications

Jordà Polo, David Carrera, Yolanda Becerra, Jordi Torres, Eduard Ayguadé (UPC/BSC) and Malgorzata Steinder (IBM Research) {jorda.polo,david.carrera,yolanda.becerra, jordi.torres, eduard.ayguade}@bsc.es steinder@us.ibm.com

1. MapReduce

MapReduce [1] is a framework originally designed by Google to exploit large clusters to perform parallel computations. It is based on an implicit parallel programming model that provides an easy and convenient way to express certain kinds of distributed computations, particularly those that process large data sets. The framework is composed of an execution runtime and a distributed file system [7] that helps with the distributed file system [7] that across nodes. The runtime and the distributed file system also provide fault tolerance and reliability, which are crucial in such a large-scale environment.

The MapReduce runtime consists of a single master process and a large number of slave processes. When a MapReduce application (or "job") is submitted to the runtime, it is split into a large number of Map and Reduce tasks, which are executed by the slave nodes. Map tasks take the input and generate a set of key-value pairs which are then merged and processed by the Reduce tasks to generate the final output. The runtime is responsible for dispatching all of these tasks to slave nodes and ensuring their completion.

Since MapReduce was first introduced, there have been a number of different implementations. Hadoop is a state of the art MapReduce runtime [8] and its supporting distributed file system [9]. It is widely used in production clusters, and has become the de-facto standard, open-source implementation. But in addition to large clusters, MapReduce has also been ported to other platforms, such as GPUs [12] and SMPs [11].

Furthermore, while MapReduce was originally used primarily for batch data processing, it is now being used in shared, multi-user environments in which submitted jobs may have completely different priorities. At the same time, MapReduce is being adopted for a wider range of applications thanks to the flexibility of the model. In addition to data analytics, it is now also used as a platform for scientific simulations as well as for some multimedia-oriented jobs, such as large-scale image processing [4] and conversion [6], facial similarity and recognition across large datasets, and video rendering [5].

Other multimedia applications such as streaming and real-time still remain an open issue in terms of workload and cluster management.

2. Workload & Resource Management

The widespread adoption of MapReduce for different kinds of applications with different needs is proving to be an interesting challenge, and it is also making scheduling, which is responsible for selecting tasks for execution across multiple jobs, even more relevant. Task selection and slave node assignment govern a job's performance.

There have been several proposals to manage workload heterogeneity in MapReduce clusters. The work in [2] describes a multi-job scheduler that is able to meet high-level performance goals. Such high-level goals can be leveraged to coschedule various jobs, and prioritize them according to their desired completion objectives and progress rate (e.g. a long-running large-scale image processing job can be run at the same time as other jobs that require a more immediate response).

Another interesting approach for multimedia applications would be to exploit the capabilities of next generation data centers composed of hybrid hardware. Some applications such as video rendering are an obvious target for specific purpose hardware like accelerators and GPUs. Other kinds of accelerable applications have already been successfully deployed in such clusters with MapReduce [10] using our scheduler as presented in [3].

Finally, there are also ongoing efforts to further improve the resource management of MapReduce, which is becoming more important as the heterogeneity of both, workloads and clusters, grows. We expect that a more finegrained resource management [13] will allow

running jobs for which MapReduce wasn't originally designed, such as multimedia applications, more efficiently.

3. Conclusions

MapReduce is a model to execute massively parallel computations on large clusters of computers. It is becoming a platform of choice for an increasing number of applications, including multimedia, due to its simplicity and flexibility. Managing the growing number of applications is challenging, but possible thanks to certain improvements made on top of frameworks such as Hadoop. Features like highlevel performance goals and a more fine-grained resource management not only ease workload management, but at the same time provide the basis for further increasing the number of available MapReduce applications.

Acknowledgements

This work is partially supported by the Ministry of Science and Technology of Spain and the European Union (FEDER funds) under contract TIN2007-60625, and by IBM through the 2010 IBM Faculty Award program.

References

[1] Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. OSDI '04: Sixth Symposium on Operating System Design and Implementation, pages 137-150, San Francisco, CA, USA.

[2] Jordà Polo, David Carrera, Yolanda Becerra, Jordi Torres, Eduard Ayguadé, Malgorzata Steinder and Ian Whalley. Performance-Driven Task Co-Scheduling for MapReduce Environments. 2010 12th IEEE/IFIP Network Operations and Management Symposium, pages 373-380, Osaka, Japan.

[3] Jordà Polo, David Carrera, Yolanda Becerra, Vicenc Beltran, Jordi Torres, Eduard Ayguadé. Performance Management of Accelerated MapReduce Workloads in Heterogeneous Clusters. 2010 39th International Conference on Parallel Processing, pp. 653-662, San Diego, CA, USA.

[4] Shimin Chen and Steven W. Schlosser. MapReduce Meets Wider Varieties of Applications. Technical Report IRP-TR-08-05, Intel Research Pittsburgh, 2008.

[5] Thomas Sandholm and Kevin Lai. MapReduce optimization using regulated dynamic prioritization. 11th International joint conference on Measurement and modeling of computer systems. 2009, NY, USA, 299-310.

[6] New York Times. Large scale image conversions.

http://open.blogs.nytimes.com/2007/11/01/self-service-prorated-super-computing-fun/

[7] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. The Google File System. 19th ACM Symposium on Operating Systems Principles, Lake George, NY, USA, 2003.

[8] Apache Software Foundation. Hadoop MapReduce Tutorial.

http://hadoop.apache.org/mapreduce/

[9] K. Shvachko, H. Huang, S. Radia, R. Chansler. The Hadoop Distributed File System. 6th IEEE International Workshop on Storage Network Architecture and Parallel I/Os, Lake Tahoe, NV, USA, 2010.

[10] Y. Becerra, V. Beltran, D. Carrera, M. Gonzalez, J. Torres, E. Ayguadé. Speeding up distributed mapreduce applications using hardware accelerators. ICPP '09: Proceedings of the 2009 International Conference on Parallel Processing, IEEE Computer Society, Washington, DC, USA, 2009, pp. 42–49.

[11] M. Gschwind, P. Hofstee, B. Flachs, M. Hopkins, Y. Watanabe, T. Yamazaki, A novel SIMD architecture for the Cell heterogeneous chip-multiprocessor, 2005.

[12] Catanzaro, B., Sundaram N., & Keutzer K. A Map Reduce Framework for Programming Graphics Processors. Third Workshop on Software Tools for MultiCore Systems (STMCS), 2008.

[13] Adaptive MapReduce Scheduler.

http://issues.apache.org/jira/browse/MAPREDU CE-1380



Jordà Polo is a PhD student at the department of Computer Architecture, Technical University of Catalonia (UPC). He got his engineer's degree from the same university in 2009 and is currently working at Barcelona Supercomputing Center (BSC) as a research student in the Autonomic Systems and e-Business Platforms group.



Dr. **Yolanda Becerra** holds a PhD in Computer Science since 2006 from the UPC. In 2007 she joined to the Autonomic Systems and eBusiness platforms research line. In this research group she is conducting research about application profiling and performance estimation, energyefficient computing, parallel and distributed systems, and about non-centralized performance management for heterogeneous execution environments.



Dr. David Carrera received the MS degree at the Technical University of Catalonia (UPC) in 2002 and his Ph.D. from the same university in 2008. He has been involved in several EU and industrial research projects. He has authored a number of papers in international journals and conferences. His research interests are focused on the integrated performance management of virtualized data centers that host applications with differentiated QoS objectives.



Prof. Jordi Torres has a Masters degree in

Computer Science from the Technical University of Catalonia (UPC, 1988) and also holds a Ph.D. from the same. Currently he is a full professor in the Computer Architecture Department at UPC and is a manager for the Autonomic Systems and eBusiness Platforms research line in Barcelona Supercomputing Centre (BSC). He has worked in a number of EU and industrial research and development projects. Hi has more than 80 publications and was involved in several conferences in the area.



Prof. Eduard Ayguade received the Engineering degree in Telecommunications in 1986 and the Ph.D. degree in Computer Science in 1989, both from the Universitat Politècnica de Catalunya (UPC), Spain. He is full professor of the Computer Architecture Department at UPC. He is currently associate director for research on Computer Sciences the at Barcelona Supercomputing Center (BSC). His research interests cover the areas of multicore architectures, and programming models and compilers for high-performance architectures.



Dr. **Malgorzata Steinder** is an IBM research staff member working on topics pertaining to the management of computation and data in largescale distributed environments. She manages the Middleware and Virtualization Management group. She holds a PhD in computer science from University of Delaware and MSc in computer science from AGH University of Science and Technology, Krakow, Poland.

Understanding Network Communication Performance in Virtualized Cloud

Guohui Wang, T. S. Eugene Ng, Rice University, USA {gswang, eugeneng}@cs.rice.edu

1. Introduction

Cloud service allows enterprise class and individual users to acquire computing resources from large scale data centers of service providers. Users can rent machine instances with different capabilities as needed and pay at a certain per machine hour billing rate. Despite concerns about security and privacy, cloud service attracts much attention from both users and service providers. Recently, many companies, such as Amazon, Google and Microsoft, have launched their cloud service businesses. Users have also started to run a large variety of applications on cloud, such as content delivery, media hosting, web hosting and streaming processing.

Most cloud service providers use machine virtualization to provide flexible and costeffective resource sharing among users. For example, both Amazon EC2 [1] and GoGrid [5] use Xen virtualization [3] to support multiple virtual machine instances on a single physical server. Virtual machine instances normally share physical processors and I/O interfaces with other instances. It is expected that virtualization can impact the computation and communication performance of cloud services. However, very few studies have been performed to understand the network performance of these large scale virtualized environments.

We study the network communication performance in virtualized cloud environments. The goal is to understand the impact of virtualization on network performance in cloud and its implications to cloud applications. We perform a large scale measurement study on the network performance of Amazon EC2 cloud. We measure the processor sharing, TCP/UDP throughput and end-to-end delays among Amazon EC2 virtual machines and observe that even though the data center network is lightly utilized, virtualization can still cause significant throughput instability and abnormal delay variations. We discuss the implications of unstable network on various cloud applications, especially on multimedia applications that requires stable and predictable network.

2. Network performance of Amazon EC2 Cloud

Our measurement is based on small instances and high CPU medium instances on Amazon EC2 cloud. We perform a spatial experiment to study the network performance of 750 small instance pairs and 150 medium instance pairs at different network locations covering 177 subnets in Amazon EC2 us-east clouds. We also perform a temporal experiment to measure 6 small instance pairs and 3 medium instance pairs continuously over one week. In this section, we high light some of our findings from the measurement study. Please refer to our full version paper [6] for more details.

Processor sharing: we use a simple CPUTest program to test the processor sharing property of EC2 instances. This program consists of a loop that runs for 1 million times. In each iteration, the program simply gets the current time by calling gettimeofday() and saves the timestamp into a pre-allocated array in memory. When the loop finishes, the program dumps all the saved timestamps to the disk. Normally, if the program is executed continuously, all loop iterations should take a similar amount of time. However, virtual machine scheduling can cause some iterations to take much longer than the others. If the instance is scheduled off from the physical processor, we would observe a gap in the timestamp trace.



Figure 1. CPUTest timestamp trace

Figure 1 shows the CPUTest timestamp trace on different environments. When the CPUTest program is run on a non-virtualized machine or a medium instance, the timestamp traces produced indicate the CPUTest program achieves a steady execution rate with no significant interruption.

However, the timestamp trace of the small instance shows very obvious scheduling effects.



Figure 2. Fine grain TCP/UDP throughput trace for small instances

When the CPUTest program is run on a small instance, periodically there is a big timestamp gap between two adjacent loop iterations. The gaps are on the order of tens of milliseconds. We believe the on/off execution pattern is caused by virtual machine scheduling. The large timestamp gaps represent the periods when the instance running the CPUTest program is scheduled off from the physical processor.

From the timestamp trace of CPUTest program, we can estimate the CPU share of our virtual machine instances by compute the fraction of on period. Our statistics shows that small instances always get 40-50% of CPU share, while medium instances mostly get the full CPU share. This means Amazon EC2 performs very restricted control on the processor sharing of small instances.

TCP/UDP throughput: Figure 2 demonstrates the fine-grain TCP and UDP throughput of a typical small instance pair in 1-second transmission. We consistently observe the same transmission pattern on all the small instances. To make the results clearly visible, we only pick one small instance pair and plot the throughput in 1 second period. We observe the drastically unstable TCP throughput switching between full link rate at near 1 Gb/s and close to 0 Gb/s. The quiet periods last for tens of milliseconds. Considering the processor sharing behavior observed in our CPUTest experiments, we believe that the quiet periods are caused by the processor sharing among small instances. During these quiet periods, either the TCP sender instance or the receiver instance is scheduled off from the physical processor, therefore no packet



metrics

We observe a similar unstable UDP throughput on small instances. The difference between UDP and TCP transfers is that, in many cases, after a low throughput period, there is a period where the receiver receives UDP traffic at a high burst rate (even higher than the network's full link rate). We believe the reason is, during the low UDP throughput periods, the receiver is scheduled off from the processor, but the sender instance is scheduled on. All the UDP traffic sent to the receiver will be buffered in the Xen driver domain. When the receiver is scheduled on later. all the buffered data will be copied from driver domain memory to the receiver's memory. Since the data is copied from memory to memory, the receiver can get them at a much higher rate than the full link rate.

End-to-end delay: we measure the packet round trip delay (RTT) of 750 small instance pairs and 150 medium instance pairs using 5000 ping probes. For each instance pair, we compute the minimum, median, average, maximum RTTs and the RTT standard deviation from the probes. Figure 3 shows the cumulative distribution of these RTT statistical metrics for small and medium instances (note that the x-axis is in log scale).

From this graph, we can see that the delays among these instances are not stable. The minimum delays are smaller than 0.2 ms for most of the small instance pairs. However, on 55% of the small instance pairs, the maximum RTTs are higher than 20 ms. The standard deviation of RTTs is an order of magnitude larger than the minimum delay and the maximum RTTs are 100 times larger than the propagation delays. The delays of medium

instances are more stable than the small instances. But we still observe that, for 20% medium instance pairs, the maximum RTTs are larger than 10ms. Considering the Amazon EC2 cloud is based on a large cluster of computers in a data center, these large delay variations are abnormal.

3. Research issues

The measurement results from our study identify a few research questions that need to be addressed in cloud. For cloud service providers, the fundamental issue is how to balance the resource sharing and performance isolation in cloud. System designers need to improve the virtualization infrastructure for more stable and predictable network performance in cloud.

For cloud users, the unstable network can obviously degrade the performance of many data intensive applications. More importantly, many multimedia applications such as video processing and media hosting require stable and predictable network performance. The unstable throughput and abnormal delay variations raise significant challenges to run multimedia applications in cloud. The question is how we can adjust these applications when migrate them to cloud. Users may need to synchronize the sender and receivers to reduce the impact of virtual machine scheduling on unstable throughput.

Virtualization also makes it hard to infer the network congestion and bandwidth properties from end-to-end measurement. The abnormal variations in network performance measurements could also be detrimental to adaptive applications and protocols (e.g. TCP vegas [4], channel-adaptive video streaming [2]) that conduct network performance measurements for self-tuning. Given the observations from our measurement study, many adaptive applications need to be reconsidered to achieve optimal performance in virtualized cloud environments. Inferring and diagnosing virtual network is also an interesting problem that needs to be studied with future research.

References

 Amazon EC2, http://aws.amazon.com/ec2/.
 B. Girod, M. Kalman, N. Liang, R. Zhang, Advances in channel-adaptive video streaming, in proceedings of ICIP'02, Dec, 2002. [3] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, Xen and the art of virtualization, In *Proceedings of SOSP'03*, Oct. 2003.

[4] L. S. Brakmo and L. Peterson, TCP Vegas: End-to-end congestion avoidance on a global internet, *IEEE Journal of Selected Areas in Communication*, vol. 13, no. 8, Oct. 1995.

[5] Gogrid,. http://www.gogrid.com/.

[6] G. Wang, T. S. Eugene Ng, "The impact of virtualization on network performance of Amazon EC2 data center", in proceedings of IEEE INFOCOM'10, Mar. 2010



Guohui Wang is a final year Ph.D. student of Computer Science at Rice University. He got a MS from Chinese Academy of Sciences and a BS from University of Science and Technology of China. Guohui's research interests are on data center networking and cloud computing.



T. S. Eugene Ng is an Associate Professor of Computer Science at Rice University. He is a recipient of a NSF CAREER Award (2005) and an Alfred P. Sloan Fellowship (2009). He received a Ph.D. in Computer Science from Carnegie Mellon University in 2003. His research interest lies in developing new network models, network architectures, and holistic networked systems that enable a robust and manageable network infrastructure.

Information Security in Distributed Computing

Winnie Cheng (IBM Research), Carlos R. P. dos Santos (UFRGS) and Nikos Anerousis (IBM Research) {wcheng, nikos}@us.ibm.com, crpsantos@inf.ufrgs.br

1. Challenges

Advances in multimedia technologies have made it easy to collect and integrate information from diverse data sources and in disparate media formats. This has been a crucial step forward in enhancing intelligence in distributed computing. It opens up opportunities for the next generation in computing where data analytics and crossdomain knowledge representation and understanding will play a leading role.

We have already seen the proliferation of web applications harnessing the power of social media and crowd-sourcing to tackle many practical problems. Tools such as mashups allow seamless integration of human computer interfaces with deep analytics engines. For example, Yahoo!Pipes [1] can create web-based applications that combine web feeds, web pages, user input with computing services such as generation of the geographical distribution of users groups.

The power of these new distributed computing platforms lie in their flexible programming interfaces that allow developers to connect their services with other developers to form larger and more sophisticated web applications. However, if not managed properly, this flexibility comes at a great cost. Data security concerns are bound to arise as user information is gathered and proprietary data models drive the algorithms behind these services. Unintended information leak can violate regulatory compliance, damage company's reputations and lead to costly lawsuits.

Every few days, we read about information leakage and data breaches in the newspapers. Rudder was touted as a convenient web-based financial planning tool where users can link their bank and credit card accounts all in one place [2]. However, in May 2009, it was reported that their software had inadvertently shown users each other's bank account information. In Canada, Ryerson University violated the Freedom of Information and Protection of Privacy Act with a software glitch that listed student names, ID numbers, and grades on its web site and has been publicly criticized for their negligence [3].

WellPoint, a Fortune 500 company and one of the largest health insurer, exposed 130,000 of its customers' protected health information and personal records online [4]. The list goes on and on. There are numerous other documented cases [5] where software errors have jeopardized users' identity (e.g., Virginia Bureau of Insurance, Comcast, Automatic Data Processing, University of Virginia), personal financial data Citigroup/ABN (e.g., Amro Mortage. CompuCredit, City of Riverside California), and patient health records (e.g., Ohio State University Medical Center, Georgetown University Hospital).

While research on information security has focused intensively on detecting and preventing data breaches due to sophisticated malware and hackers with malicious intents, an interesting and important observation is that many of these incidents are actually caused by *simple* software bugs that *inadvertently* expose sensitive data. This calls for a re-examination and shift of focus on how we should safeguard data in distributed computing.

The research questions are then: If these are simple bugs, why can't developers catch them? How can we equip developers and practitioners with the right tools and methodologies to construct safer programs?

2. A Different Approach to Information Security

Software complexity is on the rise with many software projects exceeding millions of lines of codes and getting them all correctly is very unlikely. Furthermore, the information security problem is exacerbated by this shift in the distributed computing paradigm to mashup style applications. Data passes through many intermediate software modules and gets transformed in a variety of ways. It is difficult to assess the sensitivity of data at various points in the computation flow and specify what types of data can be disclosed and at which point.

For over a decade, computing systems have relied on access control for data protection.

Access control allows one to specify who can access what data. For example, one can create a private file directory and specify that only a specific user can read and modify the files within it while having a public file directory that can be accessed by any user of the system. Access control techniques focus on relationships between users and static data. Dynamic content are common in web applications where user input directs the program execution to interact with different services and retrieve situationdependent types of data. One example is a web application that interacts with Google Maps to display the locations of shops as well as addresses of family relatives. A user will want to have different disclosure policies for these two modes of operation in the web application. Access control is inherently poor at capturing these types of constraints.

Information flow control is a class of data security models that overcomes these limitations with access control. Information flow control has been of interest in military systems [6] since the 1970s, it classifies data into different categories (e.g., "secret", "top-secret") and defines conditions on when information can flow from one category to another. The latter makes it possible to effectively and efficiently capture transient and dynamic relationships between different computing modules. In recent years, this research area is receiving renewed interests and researchers have been extending this concept to the design of programming languages and operating systems. Distributed computing can also benefit tremendously by adopting the ideas behind information flow control as it calls for a new programming model that must coordinate the computations and data accesses across many nodes in the network. Information flow control allows one to specify flow constraints without enumerating all the possible data-computation combinations in the system, making it a very practical approach to data security in an environment where modules can be composed in a variety of manner to form web applications.

Information flow control also shifts the focus to disclosure policies rather than access policies. That is, it emphasizes how data is to be used rather than what data can be read. For example, an administrative assistant in a medical clinic needs to know who the patients are in order to schedule appointments. If security is provided by access control, nothing prevents the administrator from retrieving this data and then leaking the information. Information flow control allows the access while preventing the administrator from using the system to disclose the information inappropriately, e.g., sending all patient files in email.

3. Programming in Distributed Computing

While mashup development platforms such as Yahoo!Pipe allow application developers to compose and share useful web services, when it comes to the data accessed and changed between these services, they rely on access control techniques tied to role-based access controls in databases and user login controls. Our research team has been working on a new secure mashup development platform, *Maestro* [8], which demonstrates how we can apply a programming model based on information flow control [9] to distributed computing.

In Maestro, developers can label sensitive data with different tags to express confidentiality concerns. During the runtime of the mashup application, Maestro keeps track of these tags as one module interacts with another. For example, in Figure 1, a module that makes use of geolocation information (tagged *IPLocation*) and network interface details (tagged *SystemDesc*) will have tag *[Geo, SystemDesc]* in the data it generates. In this way, the system retains the sensitivity information of data that transit through modules.

Developers can then specify modules that have privilege to declassify information (i.e., lower its sensitivity). For example, a module that produces overall network utilization information without specific IP addresses and server identifying information may be granted the authority to remove the *IPLocation* and *SystemDesc* tags. This module can then display this result to the public since Maestro prevents tagged data from leaking out of the system.



Figure 1. Defining data flows in Maestro

4. Conclusions

Distributed computing is changing as data becomes more integrated and web developers can tap into other developers' modules to build large-scale applications. This has implications on how we design next generation multimedia web applications. Data security will be a core concern and should be designed into these systems rather than as an afterthought. Programming methodologies will also have to evolve to better suit the needs of these applications. Information flow control is a promising direction to capturing and enforcing the types of data protection policies common in this environment.

References

[1] Yahoo!Pipes, http://pipes.yahoo.com/pipes

[2] Josh Lowensohn. *Rudder steers personal finance to your in-box.* CNET News:

http://news.cnet.com/8301-17939 109-10040526-2.html, September 2008.

[3] Rafael Ruffolo. Ryerson privacy breach highlights immature IT, analyst says. IT

World Canada: http://www.itworldcanada.com/a/Leadership/cbe e770a-74c7-

44b4-aa95-9dccf1bc037b.html, February 2009.

[4] Financial Week. *Data breech at WellPoint puts* 130,000 *customers at risk.* <u>http://www.financialweek.com/apps/pbcs.dll/arti</u> <u>cle</u>? AID=/20080410/REG/756233065/1036, April 2008.

[5] Privacy Rights Clearinghouse. A Chronology of Data Breeches.

http://www.privacyrights.org/ar/ChronDataBreac hes.htm.

[6] U.S. Department of Defense Computer Security Center. *Trusted Computer System*

Evaluation Criteria (The Orange Book). DoD 5200.28-STD, December 1985.

[7] C. dos Santos, R. Bezerra, L. Bertholdo, L. Granville, W. Cheng, N. Anerousis, *A Data Confidentiality Architecture for Developing Management Mashups*, IEEE Interational Symposium on Integrated Network Management, May 2011.

[8] Winnie Cheng. MIT PhD Dissertation, Information Flow for Secure Distributed Applications, MIT-CSAIL-TR-2009-040, August 2009.



Winnie Cheng is a Research Scientist at the IBM T.J. Watson Research Center. She is part of the Services Research laboratory and has published papers on distributed computing and systems security. Dr. Cheng has over 10 years experience in the IT industry worldwide and has worked for companies such as Microsoft, HP and Nvidia. She has a PhD from MIT and a MS from Stanford.



Carlos R. P. dos Santos is a PhD candidate in Computer Science at the Institute of Informatics of the Federal University of Rio Grande do Sul (UFRGS), where he also hold an M.Sc. (2008) degree in Computer Science. In 2005, Carlos received degree in Telematics from the Federal Center of Technological Education of Ceará (CEFET-CE). Currently he is an intern at the IBM T.J. Watson Research Center. His main topics of interest include network management, Web services-based management, P2P-based systems, and Web 2.0/3.0 technologies (specially Mashups).



Nikos Anerousis is the head of the Service Engineering Group at the IBM T.J. Watson Research center. In this role, he and his team are developing new technologies to manage and

transform service organizations with primary emphasis on services quality. Prior to IBM, Dr. Anerousis was the CTO of Voicemate, a venture funded startup that pioneered voice authoring technology for the financial markets. He started his career at AT&T Labs-Research working on network and service management. He also served as adjunct faculty at the department of Electrical Engineering at Columbia University. Dr. Anerousis received the Ph.D., M.Phil. and M.S. degrees in electrical and computer engineering from Columbia University, and a B.S. degree in electrical engineering and computer science from the National Technical University of Athens.

Parallel Lasso for Large-Scale Video Concept Detection

Bo Geng¹, Yangxi Li¹, Dacheng Tao², Meng Wang³, Zheng-Jun Zha³ and Chao Xu¹ ¹Peking University, China, ² University of Technology Sydney, Australia, ³ National University of Singapore, Singapore

{gengbo, liyangxi, xuchao}@cis.pku.edu.cn; dacheng.tao@uts.edu.au eric.mengwang@gmail.com; zhazj@comp.nus.edu.sg

1. Introduction

Mots video concept detectors are built upon the kernel based machine learning techniques, e.g., Support Vector Machines [1], Regularized Least Squares [2] and Logistic Regression [3], just to name a few . However, in order to build robust detectors, the learning process suffers from the scalability issues from the following perspectives [4, 5]:

- (1) These data driven based learning approaches require a huge amount of data to ensure their effectiveness and generalization in practice, e.g., millions or even billions. Conventional learning algorithms are usually infeasible to efficiently and effectively optimize model parameters under such problem scales.
- (2) It has been widely acknowledged that single visual feature is not robust enough to detect the high-level semantics in video footages. Usually a large amount of features extracted from various sources are combined to boost the detector's performance and end up with the *curse of dimensionality*.
- (3) Kernel machines have proven to be effective in many situations, whereas their storage and computational complexities are usually at the super-quadratic order to the size of the training set.

In this paper, to tackle the aforementioned three challenging problems, we propose a new parallel *lasso* (the l_1 regularized least squares) learning algorithm (Plasso), and apply it to the TRECVID datasets for large-scale video concept detection. The Parallel Incomplete Cholosky Factorization (PICF) is utilized to approximate the covariance statistics. This utilization significantly reduces the time and space complexities. Afterward, the parallel primal-dual interior point optimization algorithm together with the Sherman-Morrison-Woodbury formula is applied to learning the model parameters iteratively in a parallel fashion. Besides the regression model designed in the linear space, we also obtain the kernel version of Plasso, and consider the sample reweighting strategy to improve the detection performance and to accommodate the sample distribution bias problem. Theoretically, we show that both space

and time complexities are reduced linearly with the increment of the number of processors, and quadratically with the dimension of the matrix produced by PICF. Beyond the video concept annotation, our algorithm can be directly applied to any applications that are fit for *lasso*. In addition, it is very direct to obtain the kernel version of P*lasso*.

Experiments on the TRECVID video concept detection task suggest that *Plasso* obtained time and space savings for training effective detectors with little communication overhead.

2. Learning Algorithm

The learning problem of *lasso* is trying to find an optimal model parameter w that is able to minimize the least square cost function, balanced with the l_1 norm regularization. The solution to this problem is hard to be parallelized due to the non-smooth objective function brought by the l_1 norm regularization, where conventional gradient descent based algorithms cannot be directly applied.

To parallelize the learning phase, we adopt the primal-dual interior point method to optimize the model parameters iteratively, which is comprised of solving the Newton system for the update direction Δw and Δu , and applying the line search method to estimate the update step. By distributing the data and the optimization variables to different processors, the algorithm can be parallelized by distributing the computing at each processor.

However, it's inefficient to perform the direct parallelization, because the huge covariance matrix is involved to solve the Newton system brought by the primal-dual interior point method. Inspired by the Parallel SVM [6], we firstly apply the Parallel Incomplete Cholosky Factorization to decompose the covariance matrix as $HH^T \approx XX^T$, and subsequently adopt the Sherman-Morrison-Woodbury formula to compute the time consuming inverse operation



Figure. 1 The learning flowchart for PLasso. The data feature is distributed to each machine according to their feature index, and PICF is applied to the data matrix for preprocess. PIPM together with SMW is applied to optimize the model parameters iteratively, until the convergence condition is met.

$$(D + HH^{T})^{-1} z$$

= $D^{-1}z - D^{-1}H (I + H^{T}D^{-1}H)^{-1} H^{T}D^{-1}z$
= $D^{-1}z - D^{-1}HG^{-1}H^{T}D^{-1}z$

For a dataset with *n* samples in a *d*-dimensional space, compared with *lasso*, P*lasso* significantly reduces complexities from the original $O(d^3)$

for computational time and $O(d^2)$ for storage

space to $O(h^2 d / m)$ and O(h d / m), respectively,

if the system has m processors and the reduced dimension h is much smaller than the original dimension d.

5. Conclusions

This paper tackles the scalability problem by the parallel distributed computation perspective, and presents an algorithm Parallel *lasso* (*Plasso*) by using Parallel Incomplete Cholosky Factorization for preprocess and Parallel primal-dual Interior Point Method for parameter optimization. Experimental results on TRECVID 2005 and 2007 datasets showed that *Plasso* can effectively reduce time and space complexities linearly with respect to the number of processors, whereas

more efficiency gain can be achieved for the larger scale datasets. Details of Plasso are given in [7].

References

[1] V. N. Vapnik, Statistical Learning Theory. Wiley-Interscience, 1998.

[2] M. Wang, Y. Song, and X.-S. Hua, "Concept representation based video indexing," ACM SIGIR, pp. 654–655, 2009.

[3] R. Yan and A. G. Hauptmann, "Probabilistic latent query analysis for combining multiple retrieval sources," ACM SIGIR, pp. 324–331, 2006.

[4] C. G. M. Snoek, M. Worring, J. C. van Gemert, J.-M. Geusebroek, and A. W. M. Smeulders, "The challenge problem for automated detection of 101 semantic concepts in multimedia," ACM Multimedia, pp. 421–430, 2006.

[5] M. Naphade, J. R. Smith, J. Tesic, S.-F. Chang, W. Hsu, L. Kennedy, A. Hauptmann, and J. Curtis, "Large-scale concept ontology for multimedia," IEEE Multimedia, vol. 13, pp. 86–91, July 2006.

[6] E. Y. Chang, K. Zhu, H. Wang, H. Bai, J. Li, and Z. Qiu, "PSVM: Parallelizing support vector machines on distributed computers," NIPS, 2007.
[7] B. Genget al., "Parallel Lasso," under review.



Bo Geng received the B.Sci. degree from the Fudan University in 2007. Currently, he is a Ph.D candidate with the Key Laboratory of Machine Perception (Ministry of Education) in the Peking University. Previously, he was a research intern with the Internet Media group in the Microsoft Research Asia, and a research assistant with the Department of Computing in the Hong Kong Polytechnic University. His research interests lie primarily in machine learning, multimedia search, information retrieval and computer vision. He is a student member of IEEE.



Yangxi Li received the B.S. and M.S. degrees from Xi'an Jiaotong University and Lanzhou University, respectively. Currently, he is a Ph.D candidate with the Key Laboratory of Machine Perception (Ministry of Education) in the Peking University. Previously, he was a research intern with the Media Computing group in the Microsoft Research Asia. His research interests lie primarily in multimedia search, information retrieval and computer vision.



Dacheng Tao (M'07) is Professor of Computer

Science with Centre for Quantum Computation & Information Systems and Faculty of Eng. & Info. Tech. in University of Technology, Sydney. He mainly applies statistics and mathematics for data analysis problems in data mining, computer vision, machine learning, multimedia, and video surveillance. He has authored and co-authored 100+ scientific articles at top venues including IEEE TPAMI, TKDE, TIP, NIPS, ICDM, CVPR, AISTATS, ECCV; ACM TKDD, and KDD, with the best theory/algorithm paper runner up award in IEEE ICDM'07.



Meng Wang is a research staff member in National University of Singapore. He received the B.E. degree and Ph.D. degree in the Special Class for the Gifted Young and signal and information processing from the University of Science and Technology of China (USTC), Hefei, China, respectively. He previously worked as an associate researcher at Microsoft Research Asia, and then a core member in a startup in Bay area. His current research interests include multimedia content analysis, search, mining. recommendation, and large-scale computing. He has authored 6 book chapters and over 70 journal and conference papers in these areas, including TMM, TCSVT, TOMCCAP, ACM MM, WWW, SIGIR, ICDM, etc. He serves as an associate editor of Information Sciences, an editorial board member of International Journal of Multimedia Intelligence and Security, a guest editor of ACM/Springer Multimedia Systems journal for the special issue on "Interactive Multimedia Computing", Elsevier Journal of Visual Communication and Image Representation for the special issue on "Large-Scale Image and Video Search: Challenges, Technologies, and Trends", and Springer Multimedia Tools and Applications for the special issue on "Social Media Mining and Search", a Technical Committee Member and reviewer for various international conferences and journals. He received the best paper awards continuously in the 17th and 18th ACM International Conference on Multimedia and the best paper award in the Multimedia 16th International Modeling Conference. He is a member of IEEE and ACM.



Zheng-Jun Zha is a Research Stuff Member in School of Computing, National University of Singapore (NUS). He received the B.E. degree in automation and Ph.D. degree in pattern recognition and intelligent system from the University of Science and Technology of China (USTC), Hefei, China, in 2004 and 2009 respectively. His current research interests include multimedia content analysis, search, mining, recommendation, and large-scale computing. He received Microsoft Fellowship in 2007 and President Scholarship of Chinese Academy of Science in 2009. He is a member of the IEEE and the ACM.



Chao Xu received the B.E. degree from Tsinghua University in 1988, the M.S. degree from University of Science and Technology of China in 1991 and the Ph.D degree from Institute of Electronics, Chinese Academy of Sciences in 1997. Between 1991 and 1994 he was employed as an assistant professor by University of Science and Technology of China. Since 1997 Dr. Xu has been with School of EECS at Peking University where he is currently a Professor. His research interests are in image and video coding, processing and understanding. He has authored or co-authored more than 80 publications and 5 patents in these fields.

E-Letter Editorial Board

DIRECTOR

CO-DIRECTOR

Chonggang Wang InterDigital Communications USA

EDITOR

Mischa Dohler CTTC Spain

Kyungtae Kim NEC Laboratories America USA

Jian Tan IBM T. J. Watson USA

Xiaoqing Zhu Cisco USA Kai Yang Bell Labs, Alcatel-Lucent USA

Takahiro Hara Osaka University Japan

Vijay Subramanian Hamilton Institute Ireland

Weiyi Zhang North Dakota State University USA

MMTC Officers

CHAIR

Haohong Wang TCL Corporation USA

VICE CHAIRS

Madjid Merabti Liverpool John Moores University UK Bin Wei AT&T Labs Research USA

Jianwei Huang The Chinese University of Hong Kong China