

## In this issue:

- Welcome Message
- Research Highlights
- New Member
- New Projects
- Awards
- Publications

**Welcome Message from the Executive Committee**

Welcome to EcoCloud's first electronic newsletter!

We had an amazing kickoff in May and we are proud to say that EcoCloud has been moving forward at a strong pace in the past six months. We already took off on excellent grounds with over 1.5M CHF of external annual funding on cloud research, and funds to build a prototype datacenter observatory (in collaboration with ETHZ and Lugano). We are happy to report that since the kickoff our center has grown in a number of ways. We have added new members, our members have won a number of prestigious awards, and we have a few exciting research highlights to report. We also would like to announce a number of projects funded through Swiss-NSF and our affiliates. Indeed, 2011 has been a fruitful year! We look forward to seeing you in 2012 at our annual retreat to be scheduled for the beginning of the summer.

**Research Highlights****3D Single-Chip Cloud Systems with Active Cooling**

Cloud computing requires server-level performance at minimal energy. While 3D-integrated processor and memory chips promise high-bandwidth and low latency access to memory with significant power reductions, emerging 3D designs suffer from drastic increases in temperature and thermal gradients. Prof. Atienza in collaboration with Oracle Labs, IBM and Boston University is exploring suitable thermal-aware 3D architectures with inter-tier liquid cooling techniques. In a recent paper in IEEE Micro special issue on Big Chips, they demonstrate for the first time that integrating predictive cooling control with chip-level thermal-management – such as job scheduling and voltage-frequency scaling – enables 3D designs to achieve an overall minimal energy while mitigating thermal gradients.

**Clearing the Clouds**

The emergence of global-scale online services has galvanized scale-out software, characterized by splitting vast datasets and massive computation across many independent servers. In a paper appearing in ASPLOS 2012, Profs. Ailamaki and Falsafi and their teams identify the inefficiencies in modern server processors and memory system when running emerging scale-out workloads (e.g., analytics, data serving, debugging as a service, video streaming and web) and advocate server chip architectures and hardware mechanisms that maximize silicon efficiency for these workloads. For more information see, [Clearing the Clouds: A Study of Emerging Workloads on Modern Hardware](#) by Ferdman et al., available as an EPFL Tech. Report.

## Online Updates in Data Warehouses

Data warehouses have been traditionally optimized for read-only query performance, allowing only for offline updates at night, essentially trading off data freshness for performance. The need for 24x7 operation in global markets and the rise of online businesses make concurrent online updates increasingly indispensable. Unfortunately, state-of-the-art approaches to update warehouses can dramatically slow down query performance. In a paper in SIGMOD 2011, Prof. Ailamaki's team in collaboration with Intel Research Pittsburgh successfully introduce the use of solid-state drives to enable efficient online updates in data warehouses. For more information, see [MASM: Efficient Online Updates in Data Warehouses](#) by Athanassoulis et al., in SIGMOD 2011.

## Resource Allocation in Virtualized Environments

Effective resource management of virtualized environments in the clouds is a challenging task. State-of-the-art managements systems either rely on analytical models or testing various resource allocation strategies by running real experiments. Both approaches face a significant overhead once the workload changes. Prof. Kostic's team and his collaborators at Rutgers propose DejaVu – a framework that accelerates resource allocation in virtualized environments. Preliminary evaluation of DejaVu to be presented at ASPLOS 2012 indicates an order of magnitude faster adaptation time to workload changes for representative network services on the Amazon EC2.

## Toward Dark Silicon in Servers

In a recent paper in IEEE Micro special issue on Big Chips, July 2011, Profs. Ailamaki and Falsafi in collaboration with researchers at Northwestern University project that server chips will not scale beyond a few tens to low hundreds of cores, and an increasing fraction of the chip in future technologies will be “dark silicon” that one cannot afford to power up. Specialized on-chip architectures can leverage the underutilized die area to overcome the initial power barrier, delivering significantly higher performance for the same bandwidth and power envelope. For more information see, [Toward Dark Silicon in Servers](#) by Hardavellas et al., in IEEE Micro, July of 2011.

## New Member

### Arjen Lenstra

We welcome Arjen Lenstra, a world-renowned expert in security and cryptography to EcoCloud. Arjen's research focuses on computational and implementation aspects and design of cryptologic methods and the impact of mass market electronic devices such as PlayStation 3 game consoles on security assessments of cryptographic methods. Arjen Lenstra was closely involved in the development of the first polynomial-time algorithm for the factorization of polynomials with rational coefficients, internet computing, the (special) number field sieve, XTR, cryptographic key length recommendations, VSH, rogue CA construction, and numerous integer factorization records.

## New Projects

### Energy Monitoring in Datacenters

Prof. Atienza in collaboration with Credit Suisse is exploring energy/temperature monitoring technologies to instrument datacenters and enable synergistic management of server utilization and cooling. For more information the project [website](#).

### Topology-Aware Transaction Processing

Prof. Ailamaki in collaboration with Oracle Research is investigating technologies to automate tuning and configuration of transactional workloads to adapt to a wide spectrum of processor and memory hierarchy characteristics in commodity server hardware.

## Trustworthy Cloud Storage

A Swiss-wide team of researchers, led by Prof. Ailamaki, received CHF 1.5M to research and innovate technologies for improved performance predictability, security and verifiability of cloud storage services.

## Vertically-Integrated Server Architecture

Transistor energy efficiency is reaching its limits in modern silicon technologies and circuits requiring efficiency at the chip design level to enable further scalability. Prof. Falsafi in collaboration with HP Labs and Microsoft Research are pursuing server designs for dark silicon where only a fraction of a server chip, specialized for a particular task, is powered up at a time. Much is in embedded and portable digital platforms, specialization in servers for popular services would lead to orders of magnitude improvement in energy efficiency. For more information visit the project [website](#).

## Awards

### Conference Awards

Dan Alistarh, Seth Gilbert, Rachid Guerraoui, and Corentin Travers received the ICDCN 2011 **Best Paper Award** for the paper titled "Generating Fast Indulgent Algorithms". ICDCN is emerging as a leading forum for research in distributed computing and networking.

Vitaly Chipounov, Volodymyr Kuznetsov and George Candea received the ASPLOS 2011 **Best Paper Award** for the paper titled "S2E: A Platform for In-Vivo Multi-Path Analysis of Software Systems". ASPLOS is one of the flagship computer systems conferences and the award is a first for EPFL.

Nitin Gupta, Lucja Kot, Sudip Roy, Gabriel Bender, Johannes Gehrke, Christoph Koch received the SIGMOD 2011 **Best Paper Award** for the paper titled "Entangled queries: enabling declarative data-driven coordination".

Ippokratis Pandis, Pinar Tozun, Miguel Branco, Dimitris Karampinas, Danica Porobic, Ryan Johnson, Anastasia Ailamaki received the SIGMOD 2011 **Best Demo Award** for their system demonstration titled "A Data-oriented Transaction Execution Engine and Supporting Tools". SIGMOD is one of the flagship conferences in databases and data management systems and both awards are a first for EPFL.

### Faculty Awards

Prof. Atienza is the recipient of an Oracle Outstanding Research Award in 2011 for contributions to thermal control in Oracle's enterprise servers.

Profs. Candea, Koch and Seeger all received the prestigious ERC Starting Grant in 2011. The ERC Starting Grant "aims to provide adequate support to researchers that demonstrate the potential to perform world-class research" with awards up to 2 millions euros. Together with Prof. Kostic, a 2010 recipient, EcoCloud now boasts four ERC winners with EPFL ranking among the top with the highest number of ERC awardees in all of Europe for a single campus.

### PhD Fellowships

Manos Athanassoulis (advised by Prof. Ailamaki) and Radu Banabic (advised by Prof. Candea and Prof. Guerraoui) each received an IBM Ph.D. Fellowship for the academic year 2011-2012. This IBM Ph.D. Fellowship is an intensely competitive worldwide program honoring exceptional students interested in solving problems that are important to IBM.

Stefan Bucur received a Google European Doctoral Fellowship in 2011. This program was created to support outstanding Doctoral students conducting exceptional research in Computer Science or a closely related discipline. Stefan is one of 16 recipients of this prestigious fellowship this year.

## Publications

### Cloud Programming

- **An efficient decision procedure for imperative tree data structures**,  
T. Wies, M. Muñiz, and V. Kuncak, CADE, 2011.
- **Constraints as control**,  
A. S. Köksal, V. Kuncak, and P. Suter, POPL, 2012.
- **Data Races vs. Data Race Bugs: Telling the Difference with Portend**,  
B. Kasikci, C. Zamfir and G. Candea, ASPLOS, March 2012.
- **Interactive Synthesis of Code Snippets**,  
T. Gvero, V. Kuncak and R. Piskac, CAV, 2011.
- **Laws of Order: Expensive Synchronization in Concurrent Algorithms Cannot be Eliminated**,  
H. Attiya, R. Guerraoui, D. Hendler, P. Kuznetsov and M. Michael, M. Vechev, POPL, 2011.
- **Performance Profiling of Virtual Machines**,  
J. Du, N. Sehrawat and W. Zwaenepoel, VEE, 2011.
- **S2E: A Platform for In-Vivo Multi-Path Analysis of Software Systems**,  
V. Chipounov, V. Kuznetsov and G. Candea, ASPLOS, 2011.
- **Satisfiability Modulo Recursive Programs**,  
P. Suter, A. S. Köksal and V. Kuncak, SAS, 2011.
- **Scala to the Power of Z3: Integrating SMT and Programming**,  
A. S. Köksal, P. Suter and V. Kuncak, CADE, 2011.
- **Sets with Cardinality Constraints in Satisfiability Modulo Theories**,  
P. Suter, R. Steiger and V. Kuncak, VMCAI, 2011.
- **Striking a New Balance Between Program Instrumentation and Debugging Time**,  
O. Cramer, R. Bianchini and W. Zwaenepoel, EuroSys, 2011.
- **The S2E Platform: Design, Implementation, and Applications**,  
V. Chipounov, V. Kuznetsov and G. Candea, ACM TOCS, 2012.
- **Towards Complete Reasoning about Axiomatic Specifications**,  
S. Jacobs and V. Kuncak, VMCAI, 2011.
- **Trustworthy Numerical Computation in Scala**,  
E. Darulova and V. Kuncak, OOPSLA, 2011.

### Cooling

- **Attaining Single-Chip, High-Performance Computing Through 3D Systems with Active Cooling**,  
A. K. Coskun, D. Atienza, M. Sabry, J. Meng, IEEE Micro, 2011.
- **Characterization and prediction of two-phase flow regimes in miniature tubes**,  
E. Rahim, R. Revellin, J. Thome, and A. Bar-Cohen, International Journal of Multiphase Flow, 2011.
- **Efficiency Improvements of a Thermal Power Plant by Making Use of the Waste Heat of Large Datacenters using Two-Phase On-Chip Cooling**,  
N. Lamaison, J. B. Marcinichen, and J. R. Thome, In Proceedings of the World Engineer's Convention, 2011.
- **Energy-Efficient Multi-Objective Thermal Control for Liquid-Cooled 3D Stacked Architectures**,  
M. M. Sabry, A. K. Coskun, D. Atienza, T. Simunic, T. Brunschwiler, IEEE T-CAD, 2011.
- **On-Chip Micro-Evaporation: Experimental Evaluation of Liquid Pumping and Vapor Compression Cooling Systems**,  
J. B. Marcinichen and J. R. Thome, Macro and Nano Flows Conference, 2011.
- **Reasons to Use Two-phase Refrigerant Cooling**,  
J. B. Marcinichen, J. Olivier, and J. R. Thome, ElectronicsCooling, 2011.
- **Thermal-Aware System-Level Modeling and Management for Multi-Processor Systems-on-Chip**,  
F. Zanini, D. Atienza Alonso, L. Benini and G. De Micheli, ISCAS, 2011.
- **Towards Thermally-Aware Design of 3D MPSoCs with Inter-Tier Cooling**,  
M. Sabry, A. Sridhar, D. Atienza Alonso, Y. Temiz, Y. Leblebici, S. Szczukiewicz, N. Borhani, J. R. Thome, T. Brunschwiler, and B. Michel, DATE, 2011.

## Data Management

- **A Data-oriented Transaction Execution Engine and Supporting Tools,**  
I. Pandis, P. Tozun, M. Branco, D. Karampinas, D. Porobic, R. Johnson, A. Ailamaki, SIGMOD, 2011.
- **Agile Views in a Dynamic Data Management System,**  
O. A. Kennedy, Y. Ahmad and C. Koch, CIDR, 2011.
- **CoPhy: A Scalable, Portable, and Interactive Index Advisor for Large Workloads,**  
D. Dash, N. Polyzotis and A. Ailamaki, VLDB, 2011.
- **Coordination through querying in the Youtopia system,**  
N. Gupta, L. Kot, G. Bender, S. Roy, J. Gehrke, and C. Koch, SIGMOD, 2011.
- **Entangled queries: enabling declarative data-driven coordination,**  
N. Gupta, L. Kot, S. Roy, G. Bender, J. Gehrke, and C. Koch, SIGMOD, 2011.
- **Here are my Data Files. Here are my Queries. Where are my Results?,**  
S. Idreos, I. Alagiannis, R. Johnson and A. Ailamaki, CIDR, 2011.
- **MaSM: Efficient Online Updates in Data Warehouses,**  
M. Athanassoulis, S. Chen, A. Ailamaki, P. Gibbons and R. I. Stoica, SIGMOD, 2011.
- **PLP: Page Latch-free Shared-everything OLTP,**  
I. Pandis, P. Tözün, F. R. Johnson and A. Ailamaki, VLDB, 2011.
- **Predictable Performance and High Query Concurrency for Data Analytics,**  
G. Candea, N. Polyzotis, R. Vingralek, VLDB, 2011.
- **Predicting cost amortization for query services,**  
V. Kantere, D. Dash, G. Gratsias, A. Ailamaki, SIGMOD, 2011.
- **Probabilistic Databases,**  
D. Suciu, D. Olteanu, C. Re, and C. Koch, Morgan & Claypool Publishers, 2011.
- **Speeding Up Range Queries For Brain Simulations,**  
F. Tauheed, T. Heinis, F. Schürmann, H. Markram, and A. Ailamaki, ICDE, 2012.

## Data-centric Chip Architecture

- **A Mapping Flow for Dynamically Reconfigurable Multi-Core System-on-Chip Design,**  
V. Rana, D. Atienza, D. Sciuto, I. Beretta, IEEE T-CAD, 2011.
- **Cuckoo Directory: Efficient and Scalable CMP Coherence,**  
M. Ferdman, P. Lotfi-Kamran, K. Balet and B. Falsafi, HPCA, 2011.
- **Design Methods and Tools for 3D Integration,**  
G. De Micheli, V. Pavlidis, D. Atienza Alonso and Y. Leblebici, VLSI, 2011.
- **Proactive Instruction Fetch,**  
M. Ferdman, C. Kaynak and B. Falsafi, MICRO, 2011.
- **Toward Dark Silicon in Servers,**  
N. Hardavellas, M. Ferdman, A. Ailamaki and B. Falsafi, IEEE Micro, 2011.

## Energy-proportional Systems

- **Clearing the Clouds: A Study of Emerging Workloads on Modern Hardware,**  
M. Ferdman, A. Adileh, O. Kocberber, S. Volos, M. Alisafae, D. Jevdjic, C. Kaynak, A.D. Popescu, A. Ailamaki and B. Falsafi, ASPLOS, 2012.
- **DejaVu: Accelerating Resource Allocation in Virtualized Environments,**  
N. Vasic, D. Novakovic, S. Miucin, D. Kostic and R. Bianchini, ASPLOS, 2012.
- **Efficiency Optimizations for Implementations of Deadlock Immunity,**  
H. Jula, S. Andrica and G. Candea, RV, 2011.
- **Identifying and Using Energy-Critical Paths,**  
N. Vasic, D. Novakovic, S. Shekhar, P. Bhurat, M. Canini and D. Kostic, CoNEXT, 2011.
- **Insomnia in the Access (or How to Curb Access Network Related Energy Consumption),**  
E. Goma, M. Canini, A. Lopez, N. Laoutaris and D. Kostic, P. Rodriguez, R. Stanojevic, and P. Yague, SIGCOMM, 2011.

## Intelligence

- **Fast Convergent Algorithms for Expectation Propagation Approximate Bayesian Inference,**  
M. Seeger and H. Nickisch, Artificial Intelligence and Statistics 14, 2011.
- **Large Scale Bayesian Inference and Experimental Design for Sparse Linear Models,**  
M. Seeger and H. Nickisch, SIAM Journal on Imaging Sciences, 2011.

## Robust Systems & Networks

- **Communix: A Framework for Collaborative Deadlock Immunity,**  
H. Jula, P. Tözün and G. Candea, DSN, 2011.
- **Efficient Testing of Recovery Code Using Fault Injection,**  
P. Marinescu and G. Candea, ACM TOCS, 2011.
- **Finding Almost-Invariants in Distributed Systems,**  
M. Yabandeh, A. Abhishek, M. Canini and D. Kostic, SRDS, 2011.
- **Generating Fast Indulgent Algorithms,**  
D. Alistarh, S. Gilbert, R. Guerraoui and C. Travers, ICDCN, 2011.
- **Parallel symbolic execution for automated real-world software testing,**  
S. Bucur, V. Ureche, C. Zamfir and G. Candea, EuroSys, 2011.
- **Sahara: Guiding the Debugging of Failed Software Upgrades,**  
R. Bachwani, O. Crameri, R. Bianchini, D. Kostic and W. Zwaenepoel, ICSM, 2011.
- **Shifting Network Tomography Toward A Practical Goal,**  
D. Ghita, C. Karakus, A. Argyraki and P. Thiran, CoNEXT, 2011.
- **Timing Analysis of Leader-Based and Decentralized Byzantine Consensus Algorithms,**  
F. Borran, M. Hutle and A. Schiper, LADC, 2011.
- **Toward Online Testing of Federated and Heterogeneous Distributed Systems,**  
M. Canini, V. Jovanovic, D. Venzano, B. Spasojevic, O. Crameri and D. Kostic, USENIX, 2011.
- **Tuning Paxos for high-throughput with batching and pipelining,**  
N. Santos and A. Schiper, ICDCN, 2012.
- **WaRR: A Tool for High-Fidelity Web Application Record and Replay,**  
S. Andrica and G. Candea, DSN, 2011.