Prof. David Atienza, EPFL

YINS

ENERGY- AND THERMAL-AWARE DESIGN OF MANY-CORE HETEROGENEOUS DATACENTERS

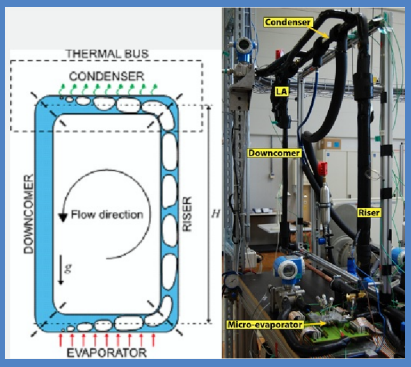

Prof. Luca Benini, ETHZ

Dr. Edouard Bugnion, EPFL

Prof. Babak Falsafi, EPFL

Marcel Ledergerber, Credit Suisse

Dr. Fabrice Roudet, Eaton Manufacturing

Patrick Segu, BrainServe

Prof. Lothar Thiele, ETHZ

Prof. John Thome, EPFL

What it's about…

Developing new technologies and optimization methodologies to develop next-generation energy-efficient datacenters.

Context and project goals

Energy efficiency in datacenters is of strategic importance to Switzerland, as over 75% of the Swiss economy is service-based and depends on information technology (IT), which makes Switzerland one of top spenders of IT per capita. IT is also witnessing a major paradigm shift towards Cloud Computing with datacenters emerging as a key backbone of services at scale, making energy dissipated in datacenters a key concern even with moderate increases in IT electricity demands. In particular, because of Switzerland's decision to abandon nuclear energy by 2034, improvements in energy efficiency are imperative to make up for half of the electricity otherwise furnished by nuclear power. Meanwhile, in Switzerland many IT departments for enterprises, research, and governmental organizations alike are at capacity with regards to their electricity budget. Thus, the only feasible solution to achieve energy sustainability is to drastically increase the power efficiency of data centers.

The design of datacenters today is a really complex process, where important opportunities exist both within and across various server and infrastructure components. Server software, system and silicon technologies as well as infrastructure for cooling and power delivery have historically been designed in isolation with an over provisioning of resources to guarantee a desired quality of service. Unfortunately, due to the diverse nature of workloads and demands on resources, such over provisioning results in prohibitive levels of waste in energy and efficiency. Modern volume server software and hardware is broadly based on designs primarily derived from the desktop market and are ill-suited for serviced-oriented server workloads. Similarly, while most datacenters make use of air-cooling technologies to ensure the correct running of the servers, air-cooling is reaching fundamental physical limits in efficiency with a continued increase in server density requiring innovation in cooling technologies.

In YINS, we propose to develop a radically new thermal-aware design approach for next generation energy-efficient datacenters. This new design approach tightly integrates the cooling infrastructure definition with holistic system-level power, performance and thermal management. This vertically-integrated system-level management paradigm goes beyond hardware and software boundaries by redesigning the entire datacenter to maximize performance given a target power, area and cost budget. Therefore, we propose to develop new server technologies based on Fully Depleted Silicon On Insulator (FDSOI) and specialized server architectures. Finally, these novel architectures interact with new on-chip microfluidic cooling delivery at server-level and passive thermosyphon cooling systems for the rack and room-level, as well as energy recovery strategies for the complete datacenter.

To realize its vision, YINS requires inter-disciplinary research at the boundaries of multiple scientific domains, as well as developing and integrating innovations in critical research areas, namely, computer systems, circuits and semiconductor technologies, and cooling technologies, large-scale simulation, software synthesis and optimization, statistical network modeling and model predictive control theory. Therefore, the YINS consortium includes six world-renown academic partners from EPFL and ETHZ covering key research areas of computer, electrical and mechanical engineering, and three key industrial partners in Switzerland for datacenter design and large-scale IT banking services provisioning (Credit Suisse, Eaton and BrainServe), and three third-party industrial partners (Constellium, Friotherm and Osmoblue).

How it differentiates from similar projects in the field

The vertically-integrated system-level management paradigm that YINS targets for datacenters goes beyond competitors in the field by performing inter-disciplinary research at the boundaries of multiple scientific domains.

It integrates innovations in several research areas, namely, computer engineering and cooling design, large-scale computing system simulation, software generation and optimization, statistical network modeling and model predictive control theory.

Therefore, YINS includes six world-renown academic partners from EPFL and ETHZ covering these key research areas and three key industrial partners in Switzerland for datacenter design and large-scale IT banking services provisioning (Credit Suisse, Eaton and BrainServe).

Quick summary of the project status

The team completed the characterization (energy, performance and temperature) full HW IPs synthetized and manufactured in 28nm FDSOI technology tailored to emulate parts of parallel accelerators for HPC and many-core server systems.

The project developed PowerCool, the first-of-its-kind mathematical model that can be used to explore architectures of micro-scale on-chip microfluidic fuel cell networks for joint cooling and localized power generation and delivery for memories and logic of many-core servers.

A new family of designs was developed for dedicated many-core network interfaces that can offload virtual switching functionality and trade-off low latency vs. high bandwidth according to different data center requirements.

A novel stacked-DRAM cache design, Unison Cache, was introduced, which includes the tag metadata directly into the stacked DRAM to enable scalability to arbitrary stacked-DRAM capacities.

An extensive experimental campaign of server thermal behaviors was performed (uniform and non-uniform heat fluxes by mimicking multi-core microprocessors) applied on a two-phase multi-microchannel evaporator under steady-state and transient regimes, and proving the high thermal (energy) performance and safe operation of such device for challenging applications.

Simulations of an entire thermosyphon cooling loop implementation and preliminary experiments in a new test bench were developed.

The team developed a flow to extract complete system thermal models of a target processor without requiring any hard-to-get information, such as, the detailed processor floorplan or system power traces. They validated it against an Intel Xeon 8-core server processor.

Success stories

Within YINS, IIS-ETHZ members presented a new method to estimate thermal profiles in many-core servers using directly on-line feedback from real sensors that are dynamically calibrated in many-core computing architectures. This approach was validated using the Intel SCC many-core processor and this work was awarded the Best Paper Award at DATE 2014.

The new system developed in YINS to monitor wirelessly energy consumption, temperature and humidity in racks and servers, called Power Monitor System and Management (PMSM), was transferred to the Wispes Sàrl start-up to start commercializing it in a short term (less than a year) in the area of energy-efficient datacenter monitoring, and it is in the process of consolidating a maintenance and extra purchase contract to expand its use in Credit Suisse datacenters in the complete Switzerland.

Presence in the media:

The “Power System Monitoring and Management (PMSM) device and system, developed at ESL-EPFL in cooperation with Credit Suisse, has been cited in the Data Center Dynamics International Magazine.

Main publications

J.B. Marcinichen, S. Szczukiewicz, N. Lamaison, J.R. Thome, Towards Development of a Passive Datacenter Cooling Technology: On-Server Thermosyphon Cooling Loop under Dynamic Workload, IEEE ITHERM 2014, Orlando FL, May 2014.

M.M. Sabry, A. Sridhar, D. Atienza, P. Ruch, B. Michel, Integrated Microfluidic Power Generation and Cooling for Bright Silicon MPSoCs, IEEE/ACM DATE 2014, Dresden, Germany, March 2014.

M. Ferdman, A. Adileh, Y. Onur Koçberber, S. Volos, M. Alisafaee, D. Jevdjic, C. Kaynak, A. D. Popescu, A. Ailamaki, B. Falsafi , A Case for Specialized Processors for Scale-Out Workloads, IEEE Micro 34(3): 31-42 (2014).

S. Novakovic, A. Daglis, E. Bugnion, B. Falsafi, B. Grot, Scale-out NUMA, ASPLOS 2014: 3-18.

Roberto Diversi, Andrea Tilli, Andrea Bartolini, Francesco Beneventi, Luca Benini, Bias-Compensated Least Squares Identification of Distributed Thermal Models for Many-Core Systems-on-Chip, Circuits and Systems I: Regular Papers, IEEE Transactions on , vol.PP, no.99, pp.1,14, doi: 10.1109/TCSI.2014.2312495.

A. Sridhar, M.M. Sabry, D. Atienza, A Semi-Analytical Thermal Modeling Framework for Liquid-Cooled Ics, IEEE Trans. On Computer-Aided Design of Integrated Circuits and Systems (T-CAD), Vol. 33, no. 8, pp. 1145-1158, 2014. doi: 10.1109/TCAD.2014.2323194.

M.M. Sabry, A. Sridhar, D. Atienza, P. Ruch, B. Michel, PowerCool: Simulation of Integrated Microfluidic Power Generation in Bright Silicon MPSoCs, IEEE/ACM ICCAD 2014, San Jose, CA, USA, November 2014.

Stavros Volos, Javier Picorel, Babak Falsafi, Boris Grot, BuMP: Bulk Memory Access Prediction and Streaming, IEEE/ACM MICRO 2014 (MICRO-47), Cambridge, UK, December 2014.

Djordje Jevdjic, Gabriel H. LoH, Cansu Kaynak, Babak Falsafi, Unison Cache: A Scalable and Effective Die-Stacked DRAM Cache, IEEE/ACM MICRO 2014 (MICRO-47), Cambridge, UK, December 2014.

Erfan Azarkhish, Igor Loi, Davide Rossi and Luca Benini, A Modular Shared L2 Memory Design for 3D Integration, IEEE Trans. On Very Large Scale Integration Systems (T-VLSI), Vol. 33, no. 8, pp. 1-21, 2014. doi:10.1109/TVLSI.2014.2340013.

Pietro Mercati, Francesco Paterna, Andrea Bartolini, Luca Benini and Tajana Simunic Rosing, Dynamic Variability Management in Mobile Multicore Processors under Lifetime Constraints, IEEE ICCD 2014, Seoul, Korea, October 2014.

S. Szczukiewicz, N. Lamaison, J.B. Marcinichen, P. J, Beucher, J.R. Thome, Passive Thermosyphon Cooling System for High Heat Flux Servers, ASME InterPACK ICNMM 2015, San Francisco, CA, USA, July 2015 (accepted paper).

Christian Conficoni, Andrea Bartolini, Andrea Tilli, Giampietro Tecchiolli and Luca Benini, Energy-Aware Cooling for Hot-water Cooled Supercomputers, IEEE/ACM DATE 2015, Grenoble, France, March 2015.

Michael Schaffner, Frank K. Gurkaynak, Aljoscha Smolic and Luca Benini, DRAM or no-DRAM? Exploring Linear Solver Architectures for Image Domain Warping in 28nm CMOS, IEEE/ACM DATE 2015, Grenoble, France, March 2015.

Erfan Azarkhish, Igor Loi, Davide Rossi and Luca Benini, High Performance AXI-4.0 Based Interconnect for Extensible Smart Memory Cubes, IEEE/ACM DATE 2015, Grenoble, France, March 2015.

Francesco Conti and Luca Benini, A Ultra-Low-Energy Convolution Engine for Fast Brain-Inspired Vision in Multicore Clusters, IEEE/ACM DATE 2015, Grenoble, France, March 2015.

Alexandros Daglis, Stanko Novakovic, Edouard Bugnion, Babak Falsafi and Boris Grot, Manycore Network Interfaces for In-Memory Rack-Scale Computing, IEEE/ACM ISCA 2015, Portland, Oregon, USA, June 2015.