




CONTENTS

Welcome	4
The Center and its Facilities	5
Faculty Members and Labs	7
Industry Affiliates	10
Current Affiliates	13
Future Initiative: UrbanTwin	14
The coolest microchips around, in designer swimwear	16
Using the matrix to help Oracle gear up	19
Beating hackers at bug hunting	21
Midgard - a paradigm shift in memory use	23
ECOGreen and CloudProphet - to see into the future	25
Sustainable technology for SKACH	27
Extendable heterogeneous energy-efficient platform: x-heap	29
Contact	30



WELCOME TO ECOCLOUD



Welcome to EcoCloud, the only academic center of its kind, promoting the development of IT infrastructure and cloud computing technology as key enablers for a sustainable society. On the one hand, this mission requires the effective use of cloud computing to preserve biodiversity, natural resources, and climate in the years to come for different core applications (e.g., digital twins for smart cities, sustainable artificial intelligence technologies, energy-aware scientific computing, etc.). On the other hand, we are targeting the definition of a new circular economy of large IT infrastructures or datacenters to minimize carbon footprint and long-term environmental damage.

Our mission is to provide support in networking the EPFL community with key industrial players in the IT sector to enable environmental sustainability for our digital world, thanks to IT. In so doing, we aim to address the major national and global IT challenges that affect us all. Moreover, our success synergistically builds on strong support for, and a growing appreciation by, private industry.

With our strong emphasis on industry collaboration and technology transfer, the Industry Affiliates Program (IAP) aims to build long-term partnerships founded on research collaborations, large-scale University-Industry partnered research grants, PhD student dissemination activities with industry, internships, fellowships, executive education and more.



THE CENTER AND ITS FACILITIES

On January 1st, 2022, Professor David Atienza took over the direction of the EPFL EcoCloud Centre from predecessor, Professor Babak Falsafi. EcoCloud's scientific mission has been expanded with a strong new focus on fundamental research and education in the domain of sustainable cloud computing.

"Historically, EcoCloud's main focus has been to deliver technologies jointly with top companies in the information technologies (IT) sector to help them optimize the large cloud computing infrastructure of public cloud systems", says Atienza. "We are now focusing on the whole IT ecosystem to develop sustainable multiscale computing from the cloud to the edge", he adds. "Our goal is to rethink the whole ecosystem and how we can provide IT solutions that can make computing more sustainable. In particular, the goal is to optimize the used resources for computing to minimize the environmental and social impact of IT infrastructures and practices. This includes the monitoring of materials, energy, water as well as other rare resources, and the creation of a circular economy for IT infrastructure, considering the impact of electronics on the environment from production to the recycling of cloud computing components."

IT infrastructure as enabler for a sustainable society

"In collaboration with the School of Engineering (STI), the School of Computer and Communication Sciences (IC), the School of Architecture, Civil and Environmental Engineering (ENAC), and the School of Basic Sciences (SB) we have defined multi-disciplinary IT application pillars or directions that are strategic for them", says Atienza.

Four multi-center discussions, and multiple projects kicked off in 2022 in the following research areas: energy-constrained and sustainable deep learning (in collaboration with the Center for Intelligent Systems (CIS) and the Center for Imaging), computational and data storage sustainability for scientific computing (in collaboration with the Space Center and the Energy Center), sustainable smart cities and transportation systems (in partnership with the FUSTIC Association, CIS and CLIMACT Center) and energy-constrained trustworthy systems, including Bitcoin technology (in collaboration with the Center for Digital Trust).

In addition to its multi-center research projects on specific applications, EcoCloud is also working on fundamental technologies to enable sustainable IT infrastructures, such as minimal-energy computing and storage platforms, or approaches to maximize the use of renewable energy in data centers and IT services deployment.



THE CENTER AND ITS FACILITIES

Moreover, EcoCloud will keep developing and strengthening, in this new era of sustainable cloud computing research, its collaboration of many years with IT partners through its Industrial Affiliates Program (IAP), such as Microsoft, HPE, Intel, IBM, Huawei and Meta, who have confirmed their interest in continuing to collaborate with the center on its new research topics.

A new facility for research on sustainable computing

"We are creating an experimental facility dedicated to multi-disciplinary research on sustainable computing at EPFL," says Atienza. In this facility, EcoCloud will provide specialized IT personnel to assist and support the EPFL laboratories in performing tests related to the proposed multi-center IT research projects and cloud infrastructures. "This year, research activities will focus on the agreed projects with the different schools and centers at EPFL, but in the future, we expect to make open calls for anyone at EPFL interested in research related to sustainable computing to be supported by EcoCloud."

An example of this is the new heating plant, inaugurated in September 2022, which has the capacity to heat and cool the Lausanne campus solely by drawing water from Lake Geneva and recovering excess heat from a connected data center.

Best practices for IT infrastructure

The dissemination of best practices for sustainable IT infrastructure is another core mission of EcoCloud. "In cooperation with the Vice-Presidency for Responsible Transformation (VPT), we are going to develop a course about the fundamentals of sustainable computing for EPFL students at the master level, which will be offered by the Section of Electrical Engineering (SEL) and the Section of Computer Science (SIN) for the complete campus," says Atienza. "Continuous education for professionals is also important. We plan to offer training to companies to support and assist them in their digitalization processes and help them understand how to implement the most sustainable IT technologies and processes possible."

"IT is the engine of our digital world. With a compound annual growth rate of more than 16%, cloud computing must embrace a strategy of digital responsibility to support economic progress and societal development without compromising the future of our planet," concludes Atienza.

The key pillars of EcoCloud activity

<< IT infrastructure as enabler for a sustainable society

- Energy-constrained and sustainable deep learning
- Sustainable smart cities and transportation systems
- Computational and data storage sustainability for scientific computing
- Energy-constrained trustworthy systems

<< Sustainable IT infrastructure


- Minimal-energy computing and storage cloud platforms
- Sustainable use of renewable energy in IT infrastructures

<< Dissemination of best practices for IT infrastructure in a sustainable society

<< Preparation of courses and focused programs on sustainable computing, for IT professionals


<< Annual EcoCloud event on sustainable computing trends and forward-looking research

FACULTY MEMBERS AND LABS - IN ALPHABETICAL ORDER



Anastasia Ailamaki
*Data-Intensive Applications
and Systems Laboratory*

Enabling discoveries in scientific domains through automating physical database design, revolutionizing exploration algorithms in very large data repositories



Alexandre Alahil
*Visual Intelligence
for Transportation*

Socially aware AI applying computer vision, deep learning and human-robot interaction to transportation applications




David Atienza
*Embedded
Systems Laboratory*

Efficient machine-learning based resource management in servers and data centers. Low-power design of edge AI and heterogeneous server architectures



Antoine Bosselut
*Natural Language Processing
Lab*

Natural language processing, machine learning, artificial intelligence



Thomas Bourgeat
*Architecture, Language and
Verification Laboratory*

We leverage the power of formal methods and high-level hardware programming languages to ensure the correctness and the security of tomorrow's computers



Maria Brbic
*Machine Learning for
Biomedical Discovery Lab*

Developing machine learning methods that solve real-world data challenges, paving the way for new biomedical discoveries



Edouard Bugnon
*Data Center
Systems Laboratory*

Data center efficiency, infrastructure support in network and data planes for OLDI applications. System security, Trusted Execution Environments in hardware



Andreas Burg
*Telecommunications
Circuits Laboratory*

Design of technology systems, prototypes and demonstrators for the development of robust, reliable and energy efficient systems



Volkan Cevher
*Parallel Systems
Architecture Laboratory*

Robust machine learning and optimization, reinforcement learning, game theory, and deep learning.



Babak Falsafi
*Parallel Systems
Architecture Laboratory*

Computer architecture, vertically integrated data center systems, post-Moore server design



Olga Fink
*Intelligent Maintenance and
Operations Systems Laboratory*

Development of intelligent algorithms for complex infrastructures and industrial systems. Deep learning and hybrid algorithms for intelligent maintenance systems

FACULTY MEMBERS AND LABS



Pascal Frossard
*Signal Processing
Laboratory*

Computer vision, medical imaging,
network machine learning, robust
machine learning, deep learning



Martin Jaggi
*Machine Learning
and Optimization Laboratory*

Machine learning, optimization
algorithms and text understanding,
as well as several application
domains



Colin Jones
*Automatic
Control Laboratory*

Theory and practice of
optimization-based, or model
predictive control with a particular
emphasis on problems arising from
renewable energy challenges



Sanidhya Kashyap
*Robust Scalable
Systems Software Lab*

Robust and high-performance
software for heterogeneous
hardware: concurrency, scheduling,
networks, analytics and fuzzing



Elison Matioli
POWERlab

Microchannel liquid cooling of data
center components, ultra-efficient
purpose-built cooling solutions



Gabriele Manoli
Optics Laboratory

Analysis and conceptualization of
complex urban and environmental
dynamics, to guide the design
of greener and more sustainable
territories



Giovanni de Micheli
*Laboratory of
Integrated Systems*

Modelling of hardware with
dedicated languages, co-design of
software and hardware, system-
level optimization with efficient
performance, energy consumption
and yield



Christophe Moser
*Laboratory of Applied
Photonics Devices*

Non-linear transformation in fiber
optics to simplify machine
learning tasks



Rachid Guerraoui
*Distributed
Computing Lab*

Distributed machine learning with
Byzantine resilience and privacy.
Distributed algorithms for new
technologies: RDMA and NVRAM



Paolo Ienne
*Processor
Architecture Laboratory*

Computer and processor
architecture, FPGAs and
reconfigurable computing,
electronic design automation,
computer arithmetic



Anne-Marie Kermarrec
*Scalable Computing
Systems Laboratory*

Large-scale distributed systems,
failure resilience, performance
and privacy-preservation, frugal
distributed learning systems



Jean-Paul Kneib
*Laboratory
of Astrophysics*

Reliable transport and precise
integration of a flow of 707
Petabytes per year of data from
large arrays of radiotelescopes



Martin Odersky
*Programming Methods
Laboratory*

The design and implementation
of Scala, to achieve a fusion of
object-oriented and functional
programming, compatible with
platforms such as Java and .NET



Mario Paolone
*Distributed Electrical
Systems Laboratory*

Developing smart grid concept
solutions to efficiently deliver
sustainable, economic and secure
electricity supply



Mathias Payer
HexHive Laboratory

Software testing to discover
security bugs. Sanitization for
memory, type, and API violations.
Fuzzing of complex code
to trigger bugs



Clément Pit-Claudel
Systems and Formalisms Lab

Can we leverage compilers,
languages, and proofs to build
more robust, more efficient, and
more trustworthy systems?



Christoph Koch
*Data Analysis Theory and
Applications Laboratory*

Efficient and scalable massively
parallel real-time analytics engines,
complex expressive declarative
and domain-specific languages in
databases



Viktor Kuncak
*Lab for Automated
Reasoning and Analysis*

Precise automated reasoning
techniques: tools, algorithms and
languages, for the construction of
reliable computer systems



Zhengmao Lu
*Energy Transport
Advances Laboratory*

Towards a deeper understanding
of phase change phenomena,
creating sustainable energy and
water technologies by optimizing
interfacial transport.



Demetri Psaltis
Optics Laboratory

Optical systems such as
spatiotemporal nonlinearities in
multimode optical fibers, used as
neuromorphic neural networks



Mirjana Stojilovic
*Parallel Systems
Architecture Laboratory*

Electronic design automation,
reconfigurable computing,
electromagnetic-compatibility and
signal-integrity issues, hardware
security



Carmela Troncoso
Spring Lab

Designing strong, embedded
security and privacy guarantees in
complex systems. Quantification of
the information an adversary can
infer from acquired data



Amir Zamir
*Visual Intelligence and
Learning Lab*

The development of computer
vision models that can function as
part of larger intelligent systems

FACULTY MEMBERS AND LABS

INDUSTRY AFFILIATES

MISSION

The EcoCloud Industry Affiliate Program (IAP) offers companies a unique opportunity to collaborate with EPFL faculty, students and researchers. Affiliates are given unparalleled access to new technologies and ideas as they move from laboratory to marketplace. The ideal platform for communication and discovery between the research and corporate communities, the program catalyzes collaborative research, customizes educational programs and facilitates graduate recruiting.

The EcoCloud IAP was created to enable connections and strengthen collaborations between EcoCloud and industry. While there are numerous benefits to joining the IAP and becoming an Affiliate, the advantages boil down to three key reasons:

<< To gain early awareness of the latest research - Through meetings, visits and online resources, EcoCloud Industry Affiliates get to preview the latest research findings from across our labs before they are published.

<< To explore potential research collaborations and sponsorships - Companies can get much more by sponsoring EcoCloud research directly. Becoming an Affiliate allows you to see how we work and what we do, giving you the insights you need to identify research partnerships.

<< To recruit EcoCloud students - Our students are one of our most valuable assets. They can add substantial value to your company as interns or employees. We post student profiles on the members-only website, and you can meet them at our Annual Event. EcoCloud can also host talks geared towards students, and distribute job and internship announcements.



INDUSTRY AFFILIATES

MEMBER BENEFITS

The EcoCloud IAP is a corporate membership program whereby companies pay an annual membership fee in return for facilitated access to the research programs, the researchers and the graduate students, offering the ability to capitalize on the unique, dynamic trans-disciplinary innovation culture at EPFL. Spanning entrepreneurial startups emerging from our research, to delivering customizable executive education, industry collaboration and technology transfer lie at the heart of the EcoCloud's IAP.

Benefits of becoming an Affiliate include:

<< Annual Event - An annual conference exclusively designed for our existing and prospective Industry Affiliates and research colleagues to showcase the activities of the center. The event's program includes activities such as presentations by EcoCloud faculty and researchers on the latest research results organized around chosen research themes; student poster sessions; lab tours; demonstrations; discussions on grand challenges, applications and technology roadmaps; and opportunities to meet and network with EcoCloud researchers, students and colleagues.

<< Research Monitoring - Throughout the year, EcoCloud will enable Affiliates to remain engaged by means of newsletters, seminars, talks and virtual meetings. When feasible, these events will be broadcast live for our IAP member companies, so participants can join in remotely to hear from different thought leaders at EcoCloud and keep informed about the latest research. Affiliates also have access to the comprehensive information about research outcomes, EcoCloud events, video and publication archives, and other research outputs.

<< Graduate Student Recruiting - EcoCloud will organize events (including the Annual Event), during which Affiliates have access to soon-to-be graduating students to facilitate recruiting. We work with our Affiliates to facilitate recruiting activities throughout the year, including advertising job and internship announcements, hosting talks and seminars, and other student-targeted networking events.

<< Occasional Visits - Affiliates may arrange visits to EcoCloud's affiliated laboratories and the experimental facility created in the new data center of EPFL. Visits enable previews of EcoCloud's research programs and results and demonstrations of emerging technologies. We work with our Affiliates to identify appropriate EcoCloud researchers, ongoing projects and potential opportunities for collaboration.



INDUSTRY AFFILIATES MEMBER BENEFITS

<< Joint Research Projects - Member companies have opportunities to engage with EcoCloud in research projects and collaborations into deployable technology. These include, but are not limited to: the opportunity to contribute and formally participate in EcoCloud research projects, customization of educational programs and the opportunity to develop and sponsor structured research programs.

<< Advertising - Our affiliates have the opportunity to promote their company's brand on EcoCloud's website and reports.

<< Outreach and Executive/Continuing Education - EcoCloud develops and hosts at least two outreach programs per year. Past events included a summer workshop on cooling technologies, a winter school on data-centric systems, and co-hosting the 3D silicon integration conference. Moreover, EcoCloud is eager to develop special executive courses, as well as continuing education courses, to address the specific needs of our affiliates.

<< Technical Advisory Board - EcoCloud Affiliates designate a technical staff to the EcoCloud Technical Advisory Board which meets once a year to discuss grand challenges, research and industrial trends, and EcoCloud research direction. This annual meeting can vary from a one-hour meeting to a full-blown two-day retreat to present research and solicit feedback.

<< Visiting Scholars and Fellows Program - The EcoCloud Visiting Scholar and Fellows Program stimulates and supports our research by engaging promising scholars and practitioners in order to foster exchange. Each year, a number of distinguished academics (Visiting Scholars) and junior faculty and students (Fellows) will be selected on the basis of their qualifications, the quality of their research plans, and the relevance to both EcoCloud's mission and targeted research objectives. EcoCloud's Visiting Scholars and Fellows will work on projects that offer joint collaborative opportunities.

CURRENT ECOCLOUD AFFILIATES IN ALPHABETICAL ORDER



IBM Research Europe



ORACLE

FUTURE INITIATIVE: URBANTWIN

SEEING DOUBLE FOR SUSTAINABILITY



David Atienza
ESL, EPFL



Mario Paolone
DESL, EPFL



François Maréchal
IPESE, EPFL



Patrik Hoffmann
EPFL, EMPA

A consortium of Swiss research institutes have begun working on UrbanTwin, to make an AI driven, ecologically sensitive model of the energy, water and waste systems the town of Aigle to help boost sustainability.

A collaboration of Swiss research institutions plans to make identical twins of another kind, using neural networks instead of DNA, to create a double of a Swiss town. Aigle has been chosen due to its size and because it has an extensive range of water sources. It also has very detailed energy monitoring infrastructure previously developed by the Energy Center of EPFL. Lausanne is a potential partner.

EPFL involvement in UrbanTwin

Labs from the faculties of:

Engineering
Architecture, Civil and Environmental Engineering
Computer and Communication Sciences
Basic Sciences

Centers:

EcoCloud
Energy Center
Climate Impact and Action (CLIMACT)
Intelligent Systems (CIS)

One of ten nationally funded Joint Initiatives of the ETH Board addressing the strategic areas of energy, climate,

and environmental sustainability, UrbanTwin aims to develop and validate a holistic tool to support decision-makers in achieving environmental goals, such as the Energy Strategy 2050 and the vision of climate-adaptive "sponge cities". The tool will be based on a detailed model of critical urban infrastructure, such as energy, water, buildings, and mobility, accurately simulating the evolution of these interlinked infrastructures under various climate scenarios and assessing the effectiveness of climate-change-related actions.

"Urban areas are responsible for 75% of greenhouse gas emissions while rising temperatures significantly impact their liveability. They represent a natural integrator of several systems, including energy, water, buildings, and transport. So, they represent the ideal setting for implementing a coordinated, multi-sectoral response to climate changes leveraging digitalization as a systemic approach," explains David Atienza, Director of EcoCloud and Head of EPFL's Embedded Systems Laboratory (ESL). David Atienza and François Marechal (Head of EPFL's Industrial Process and Energy Systems Engineering Group) are the coordinators of UrbanTwin.

"In UrbanTwin, we want to collect information from multiple sources by using new edge artificial intelligence (AI) platforms and integrate them using cloud computing technologies on a detailed model of critical urban infrastructures, such as energy, water (both clean and waste), buildings, and mobility and their inter-dependencies," continues Atienza.

FUTURE INITIATIVE: URBANTWIN

SEEING DOUBLE FOR SUSTAINABILITY



Participating Swiss Federal Institutes :

EPFL, Lausanne
ETHZ, Zurich
WSL (Institute for Forest, Snow and Landscape Research)
EMPA (Laboratories for Materials Science and Technology)
EAWAG (Institute of Aquatic Science and Technology)

"As a cutting-edge example of what digitalization and AI can offer, this tool will be able to consider underlying socio-economic and environmental factors, while assessing the effectiveness of climate-change-related actions beforehand," adds Atienza. "The goal is to develop a technology that is open and can be applied to other urban areas in any region of Switzerland".

Also, it was key to have a flexible and realistic urban environment, such as Aigle, to use as case study. "By making reference to the Aigle demonstrator, we will develop an advanced modelling and control framework for the day-ahead and intraday control of urban/rural multi-energy systems," explains Mario Paolone, Head of EPFL's Distributed Electrical Systems Laboratory (DESL). "The framework will be capable of integrating the physical constraints of the electrical, mobility, heating/cooling, and water systems, along with the representation of the stochastic nature of the available resources. Based on this framework, a planning tool for integrated energy systems that considers their daily operation will be developed. The aim is to produce planning decisions inherently satisfying daily and intra-day operational needs."

Inter-institutional cooperation is essential to UrbanTwin, with researchers working between multiple laboratories from different institutions. An example of this is the water monitoring system, which will include selecting the best source of freshwater supplies and measuring its quality, as well as modelling the disposal of waste water. This research will include AI technologies that will do detective work, with a system that will track sources of pollution as quickly as possible and send alarms with the origin located and reported.

To this end, Giulio Masinelli, a doctoral student jointly co-supervised by EMPA and EPFL, will work on developing new smart multi-parametric sensing systems to create the digital twin. He has a good appreciation of this project as he is working on a similar approach for advanced manufacturing. "We can generate data by installing sensors on sinks," explains Masinelli, "measuring water quality around the city, the pH level, the salt concentration

and other metrics. We will use machine learning to collect observations, and then make predictions - with physical constraints. These constraints are what make a simulation powerful because it becomes a flexible model with lots of parameters.

"Masses of work goes into applying partial differential equations to the data so that the system can be generalized," continues Masinelli, "without a drop in quality coming from physical constraints and unfamiliar data. The result is a neural network that can generate results in a couple of milliseconds: the resolution of the partial differential equations. Then you can fine-tune the parameters so that they will work with all data. You must not stay too close to one dataset if you want good predictions."

UrbanTwin represents a welcome opportunity for these researchers to collaborate with a range of different teams at a difficult time for Swiss scientists. Participation in Horizon Europe was lost to Swiss researchers since the country broke off negotiations with the EU in 2021, making national funding the only current option. François Maréchal, is hopeful that UrbanTwin can repay the investment of the Swiss government: "If we can improve the way city administrators deal with their resources and raise levels of efficiency, it would be a really big step."

Currently, AI and cloud computing are used in an ever-increasing number of ways in research, as exemplified by the EcoCloud center of EPFL. Atienza and Maréchal are convinced that "Sustainable digital twin technologies will be implemented through UrbanTwin, which will provide a great tool to complement decision-makers in their work, searching through vast stores of data to find anomalies, or recommendations, that would take a person too long to find. UrbanTwin will be an AI system, and a holistic one: we expect unexpected results."

Unexpected results should not come as a big surprise here: it's a twin thing.



THE COOLEST MICROCHIPS AROUND

WITH DESIGNER SWIMWEAR



Elisa Matioli
POWERlab

Having come through a heatwave it is easy to forget that humans are not the only ones that need to shelter from oppressive heat. Animals suffer too, and huge mechanical infrastructures like railways, even major runways. But what about computers?

From the beginning, data centers have faced problems with cooling. In the old days, it was known that Google had a fantastic set-up in their data centers where super computers would be laid out in open rows, in a refrigerated environment. Technicians in California could go and replace velcro-mounted components, while enjoying the cool air for a few refreshing minutes. The real benefit was that the rooms would be kept at a low temperature, to lighten the load on each server's individual cooling system. Nowadays, however, this would be regarded as extraordinarily inefficient: why waste energy refrigerating an entire room, when it is only the PCs that have to be kept cool?

Prof. Elisa Matioli from POWERlab has taken this question a step further. Why should the entire PC be cooled down by having air blown at it, when each individual chip could have its own liquid-cooled system?

As data center demand goes up, so does the cost of cooling

Unless you are unusually well-versed in data center technology, it is unlikely you will be aware of the full extent to which you make use of data centers. If you use Gmail or Outlook.com it is likely that every email you send or receive is stored in a center in the USA; if you have your photos backed up on Apple iCloud or Google Photos, they will be stored in several of the many data centers these companies run all over the world. As more and more people increasingly make use of the Internet-of-Things, from smart cars to smart doorbells to fitness monitors, our data center usage is constantly growing.

At EPFL's EcoCloud Center we have many professors working on ways to decrease the energy consumption of data centers, even as demand grows and global temperatures increase. At the smallest level, Prof. Moser and Prof. Psaltis have been researching light propagation in optical fibers, and using it to perform practical computational tasks, with much lower energy consumption than traditional digital techniques. At the city-wide level, Prof. Paolone has been building smart grids that turn static power networks into self-regulating, highly-efficient intelligent systems.

Prof. Elisa Matioli is working in between these two extremes, at the level of computer components. "Our vision

THE COOLEST MICROCHIPS AROUND

WITH DESIGNER SWIMWEAR

is to contribute to the development of an integrated chip – a single unit for powered electronics where you have loads of devices integrated – smaller and more energy-efficient than anything that can be achieved currently."

Microchips continue to get smaller in size, so scientists around the globe are seeking out alternatives to silicon, as the natural limits of this tried and tested material impose themselves.

Prof. Matioli has identified the best alternative, but there is an inevitable problem:

"Using Gallium Nitride allows us build electronic devices like power units, memory chips or graphic cards, much smaller than can be achieved with silicon. We can deliver better performances in energy conversion with a much smaller volume but, as a consequence of this, we are generating greater amounts of heat over a smaller surface area.

"It is vital that components do not overheat, nor cause neighboring devices to overheat."


Cooling revolutionary chip sets like the above became a main focus of Prof. Matioli's team, and led to some radical solutions. In turn, these solutions opened new possibilities for cooling all kinds of chip sets, including data center infrastructure.

Walking through the lab, it comes across like science fiction.

Just as an astronaut uses a space suit with built-in liquid-cooling, so these microchips are each housed in a liquid-cooled membrane. Cooling with air in the traditional way is fine, but liquid conducts heat faster than air, so the cooling is much more efficient. In the devices being pioneered by POWERlab, microchannels of varying diameter provide a cooling system that is tailored to the needs of each chip, as part of a cooling network designed for the entire machine – the hot spots having been identified in advance. Crucially, it is only the hot spots themselves that are targeted – an ultra-efficient strategy.

Getting mechanical and electronics engineers to work together

This "co-designing" of the ultra-compact microchips and their cooling system makes the approach unique, and



beats at the heart of spin-off Corintis SA, a start-up which has evolved out of POWERlab, and is currently recruiting.

"Corintis is bringing a service to chip manufacturers, providing them with heat maps for their devices. Their experts can optimize microfluidic cooling while the customer is designing their microchips. They can then design the heat sink in a way that is tailor-made for their chipset."

Interdisciplinarity is a key feature to this work: "Very often the departments looking at thermal issues and electronic devices are in different buildings: mechanical engineering and electrical engineering. So you build a chip and then send it to another department to find a way to cool it down. But by this time you have already missed many opportunities!

"In our lab I brought mechanical engineers and electrical engineers to work together," explains Prof. Matioli, "and that is what makes us different."

THE COOLEST MICROCHIPS AROUND WITH DESIGNER SWIMWEAR



Remco van Erp, CEO of Corintis SA

In conversation with Remco van Erp of Corintis

The annual increase in computing power of general-purpose chips has been slowing down in recent years. Many of the biggest tech companies in the world are now designing their own application-specific chips to meet their future needs: Apple designs chips for their phones and laptops, Amazon designs chips for their data center, and YouTube even designs chips for video processing. There is a large amount of heterogeneity. The custom design of chip-sets can greatly benefit from tailor-made cooling solutions to improve energy efficiency and performance, especially where data centers are concerned. Increasingly, companies are coming to us looking for better cooling solutions.

What is the range of expertises at Corintis?

This is a very multidisciplinary problem, requiring expertise ranging all the way from mathematics to plumbing. At Corintis, we have computational and hardware experts working together to achieve our goals. The modelling

aspect is very important, since we want to predict power and temperature distribution, and optimize the microfluidic cooling design, before a chip is even manufactured. It's also a multi-scale problem: on the one hand, we are dealing with channels at the micrometer scale, on the other it is integrated into chips that are several centimeters big. This requires clever innovations in modelling and simulation.

Does EPFL still have a role to play?

We keep strong links with EPFL: our microfabrication experts are working in the clean rooms there, we have four interns from EPFL and other international institutions, and we are applying for research funding in collaboration with POWERlab.



More about
Designer Swimwear



Read more about Corintis

USING THE MATRIX TO HELP META GEAR UP



Martin Jaggi

Machine Learning
and Optimization Laboratory

Just 12-months after it was created, in December 2004, 1-million people were active on Facebook. As of December 2021 it had an average 1.93 billion daily active users. EPFL is in a unique collaboration with its parent company Meta for distributed deep learning research.

For a user base of this size, large-scale automated-systems must be utilized to understand user experience in order to ensure accuracy and success. EPFL's Machine Learning and Optimization Laboratory (MLO), led by Professor Martin Jaggi, is in active collaboration with Meta Platforms, Inc., Facebook's parent company, to solve this unique challenge.

With funding from EPFL's EcoCloud Research Center, MLO collaborates with Meta through internships at the company for MLO researchers and the use by Meta of a pioneering MLO invention: PowerSGD. MLO is helping Meta to analyze and better understand millions of users' experiences while at the same time respecting user privacy. This requires collaborative learning, that is, privacy-preserving analysis of information from many devices for the training of a neural network that gathers, and even predicts, patterns of behavior.

To do this, a key strategy is to divide the study of these patterns over "the edge", using both the user's device,



and others that sit between it and the data center, as a form of distributed training. This requires a fast flow of information and efficient analysis of the data. PowerSGD is an algorithm which compresses model updates in matrix form, allowing a drastic reduction in the communication required for distributed training. When applied to standard deep learning benchmarks, such as image recognition or transformer models for text, the algorithm saves up to 99 % of the communication while retaining good model accuracy.

PowerSGD was used to speed up training of the XLM-R model by up to 2x. XLM-R is a critical Natural Language Processing model powering most of the text understanding services at Meta. Facebook, Instagram, WhatsApp and Workplace all rely on XLM-R for their text understanding needs. Use cases include:

Content Integrity: detecting hate speech, violence, bullying and harassment;

Topic Classification: the classification of topics enabling feed ranking of products like Facebook;

Business Integrity: detecting any policy violation for Ads across all products;

Shops: providing better product understanding and recommendations for shops.

USING THE MATRIX TO HELP META GEAR UP



Tao Lin of the Machine Learning and Organization Laboratory

MLO doctoral researcher Tao Lin explains the PowerSGD process:

There are three aspects to the process. The first is to develop gradient compression algorithms to speed up the training, reducing the time required to prepare this information for its transfer to a centralized hub.

The second is efficient training of the neural network within a data center – it would normally take several weeks to process all the information, but we distribute the training, reducing computation from months to days.

As a third aspect, privacy is a constant factor under consideration. We have to distinguish between knowledge and data. We need to ensure users' privacy by making sure that our learning algorithms can extract knowledge without extracting their data and we can do this through federated learning.

The PowerSGD algorithm has been gaining in reputation over the last few years. The developers of deep learning software PyTorch have included it as part of their software suite (PyTorch 1.10), which is used by Meta, OpenAI, Tesla and similar technology corporations that rely on artificial intelligence.



BEATING HACKERS AT BUG HUNTING



Mathias Payer
HexHive Laboratory

On the 9th of December, 2021 the world of IT security abruptly went into a state of shock. An alarming message was spreading like wildfire:

RCE 0-day exploit found in log4j

For the uninitiated, there is a lot to unpack here. "RCE" stands for remote code execution: similar to when somebody takes control of your computer with TeamViewer to run programs of their choosing. In this context, however, control is exerted without the consent, or even the knowledge of the owner.



A zero-day exploit is a major software vulnerability, previously unknown to the developer. They must act quickly to develop a patch because, by the time the developer learns about it, adversaries could already be exploiting the opening.

The log4j library allows Java software to log (report) certain conditions. A vulnerability in it could allow an adversary to execute arbitrary code in the context of the underlying software in applications around the globe.

Put it all together and you get this: at the time the above headline was published, a system tool used by companies all over the world - in cloud servers, game servers and financial platforms - was already being exploited by hackers, allowing them to take control of servers and data centers.

93% of the world's cloud services affected

According to the Wall Street Journal, "U.S. officials say hundreds of millions of devices are at risk, hackers could use the bug to steal data, install malware or take control."

One estimate stated that the vulnerability affected 93% of enterprise cloud environments. At EPFL, all IT administrators were sent instructions to patch their server software immediately. Even Oracle Corporation, world leaders in information security, had to send out a distress call:

"Due to the severity of this vulnerability and the publication of exploit code on various sites, Oracle strongly recommends that customers apply the updates provided by [our] Security Alert as soon as possible."

It is hard to gauge the full extent of the damage caused, but it is clear that these vulnerabilities have real-world

BEATING HACKERS AT BUG HUNTING

impact: among confirmed victims of the log4j bug are the Belgian Ministry of Defence, the UK's National Health Service and a range of financial trading platforms. So the question begs itself - what are corporations like Oracle doing about it?

As a matter of fact, Oracle had already been working against this kind of vulnerability long before the log4j zero day. The log4j library uses deserialization: a server receives structured data (a form of code and object relationships) for processing. If the checks during deserialization are insufficient, and allow the attacker leeway in how the data is interpreted, it often results in RCE. Identifying the vulnerabilities exposed during the deserialization process had long been a subject of interest to Oracle researchers by 2020, when they reached out to Prof. Mathias Payer of EPFL's HexHive lab:

"We had already covered fuzzing and program analysis, and had worked on cloud security as part of EPFL's EcoCloud Center," explains Prof. Payer, "but we had not approached these deserialization bugs. Then we got to work with Oracle Labs (part of Oracle Inc), who provided funding via a gift. Dr. François Gauthier and Dr. Kostyantyn Vorobyov, two Oracle researchers from Oracle Labs, introduced us to the complex technical issues that they were facing. And then we worked together, and developed a platform for discovering deserialization vulnerabilities.

"People have been attempting to find and exploit vulnerabilities in deserialization code, including Oracle's, for years: either intent on gaining some kind of direct advantage, or to earn money by submitting bug reports. Either way, these are dedicated, manual attacks. In these manual attacks, the analyst thoroughly analyzes the source code of the target and then painstakingly crafts the exploit. What we have developed is a mechanism that automates the process, and allows Oracle to get ahead of the attackers.

Eight moves ahead, like a chess grandmaster

"In addition to this, the bugs that we are finding can be much more complex than the ones that experts are finding manually. Most analysts are trained to search to a depth of two manipulations: an entry and a control vector. Our platform creates an abstract dependency graph for all available classes, and can do a fuzzy search to a depth of up to eight manipulations."

The battle between IT security managers and attackers

is one where the defenders hope to find bugs before the attackers do. However, Payer explains that security managers have one key advantage when it comes to using HexHive's platform: "Although our tool is neutral, i.e., it can be used by both attackers and defenders, developers have full access to and understanding of their own code, which gives them a huge advantage over a hacker when it comes to interpreting the results. They therefore have a very good chance of finding weak points before the attacker."

Negotiations are under way to set up internships for HexHive researchers at Oracle Corporation. "This will be good for Oracle because they will have people who actually developed some of the code on site, which will make it easier to integrate the platform into their pipeline. Another thing I appreciate is that our prototype will remain open source, and bug reports will be published."

So long as information technology is around, the battle between security managers and hackers will rage on. Thanks to their collaboration with HexHive, however, Oracle will be able to keep one step ahead of the aggressor: faster, higher, stronger.

ORACLE



MIDGARD

A PARADIGM SHIFT IN MEMORY USE



Babak Falsafi
Parallel Systems
Architecture Laboratory



Mathias Payer
HexHive Laboratory



David Atienza
ESL, EPFL



Boris Grot
University of Edinburgh



Abhishek Bhattacharjee
Yale University

Our researchers have pioneered an innovative approach to implementing virtual memory in data centers, which will greatly increase server efficiency.

Virtual memory has always been a pillar for memory isolation, protection and security in digital platforms. The use of virtual memory is non-negotiable, even in widely-used hardware accelerators like GPUs, NICs, FPGAs and secure CPU architectures. It is therefore vital that silicon should be used as frugally as possible.

As services host more data in server memory for faster access, the traditional virtual memory technologies that

look up data in server memory and check for protection have emerged as a bottleneck. Modern graph analytics workloads (e.g., on social media) spend over 20% of their time in virtual memory translation and protection checks. Server virtualization for cloud computing, to help increase utilization of infrastructure and return on investment in data centers, dramatically exacerbates this problem by requiring lookups and protection checks across multiple layers of guest (customer) and host (cloud provider) software.

The way in which virtual memory is assigned in these servers is critical because, with such huge quantities of data involved, changes in strategy can have a massive effect on server efficiency and data security.

Memory has become the most precious silicon product in data centers in recent years, as more services are brought online. Virtual memory traditionally divides up the physical storage into fixed size units, for optimal capacity management. This division slows down lookups and protection checks as memory capacity increases, because large regions of memory in application software (gigabytes) is disintegrated into millions of pages

MIDGARD

A PARADIGM SHIFT IN MEMORY USE

(kilobytes). Modern chips (e.g., the recently announced Apple M2) employ thousands of table entries per processor to do lookups and perform protection checks for each memory access.

Namespaces are used to store unique references for data, in structured hierarchies. Removing some of this hierarchy and reducing the number of translations would represent a net gain in efficiency. The authors propose Midgard, which introduces a namespace for data lookup and memory protection checks in the memory system without making any modifications to the application software or the programming interface in modern platforms (e.g.: Linux, Android, macOS/iOS).

With Midgard, data lookups and protection checks are done directly in the Midgard namespace in on-chip memory, and a translation to fixed size pages is only needed for access to physical memory. Unlike traditional virtual memory, whose overhead grows with memory capacity, Midgard future-proofs virtual memory as the overhead of translation and protection check to physical memory decreases, with growing on-chip memory capacity in future products, filtering traffic to physical memory.

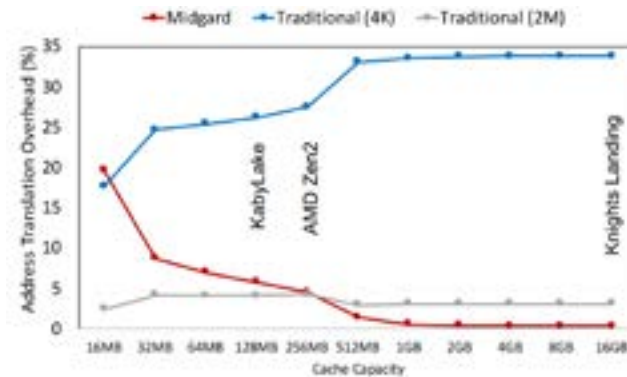
Analytic and empirical results described in the paper show a remarkable performance from Midgard when compared to traditional technology, or even rival new technologies (e.g., the larger fixed size pages used in certain applications). At low capacity the Midgard system was 5% behind standard performance, but with loads of 256 MB aggregate large cache it can match and even outperform traditional systems in terms of virtual memory overheads.

"This initial work is the first of several steps needed to demonstrate a fully functioning system with Midgard," explains Prof. Babak Falsafi, EPFL Professor, and founder of EcoCloud. We focused on a proof-of-concept software-modelled prototype of key architectural components. Future work will address the wide spectrum of topics needed to realize Midgard in real systems."

Global semiconductor manufacturer Intel is now funding this research, bringing together experts from PARSA, HEXHIVE and ESL of EPFL, as from the Computer System Lab at Yale, and the EASE Lab of the University of Edinburgh.


The project is entitled "Virtual Memory for Post-Moore Servers", and is part of Intel's wider research into improving power performance and total cost of ownership for servers in big-scale datacenters.

Midgard etymology: a middle realm between heaven (Asgard) and hell (Helheim)



ECOGREEN AND CLOUDPROPHET

TO SEE INTO THE FUTURE



Luis Costero

formerly of ESL, now Associate Professor at UC Madrid



David Atienza

Embedded Systems Laboratory

Reducing the carbon footprint of data centers means taking control of two main factors: sourcing the energy supply wisely, and using data center resources efficiently. The first of these factors is getting increasingly complicated: energy production was a relatively consistent process when we mainly relied on regular suppliers such as nuclear facilities, coal-burning power stations and hydroelectric dams. However, wind and solar are highly volatile methods of putting electricity into the grid: when the sun sets, solar panels stop producing energy, so there is a marked decrease at night — a volatile, if regular, occurrence. Moreover, solar supply intensity is subject to clouds blocking out or revealing the sun, while wind farm productivity is affected as the wind changes. Many data center managers have an added level of complexity given that their centers themselves have solar panels and other forms of energy production, so that load balancing is complicated and highly variable. When to switch to which source, in order to maximize the use of renewables?

As for using data resources efficiently, if there was a regularity in the use of data center facilities by their customers, it would be an easy game. However, the resources of a data center are used by customers (more often customers of customers) in a way that is not only unpredictable, but lacking in transparency even in real

time: data center staff are not permitted to observe the processes being used by their customers. So how do you allocate resources efficiently, when you are being handed black boxes?

For both of these complex questions, it would appear that a crystal ball is necessary. In the face of this challenge, EcoCloud's Embedded Systems Lab (ESL) has come up with two artificially intelligent crystal balls: ECOGreen and CloudProphet, with a team that includes David Atienza, Luis Costero and Darong Huang.

Prof. David Atienza explains: "ECOGreen is a system that can anticipate needed resources, having characterized current applications, and will attribute a set of jobs to a particular data center, balancing all the options that are available in the energy market.

"There is a template that allows the modelling of a data center, in terms of its network behavior, computer behavior, cooling and other factors. Then there is a model for the energy supply system of each country and even for geo-distributed data centers, so that tasks can be assigned to different data centers throughout the world.

"A manager can then say, 'As far as is possible, we will configure our system to use data centers where a maximum of renewable energy is available. Then we'll set it to migrate tasks to other data centers as the situation develops.'

"In this way we can let them match the high level of their



intel®

ECOGREEN AND CLOUDPROPHET TO SEE INTO THE FUTURE

expectations to the potential performance of the data centers."

This highly developed knowledge-based system is therefore tailored to the needs of each data center, providing suggestions that take into account many different factors: the cost, quality and origin of various potential sources of electrical energy, power loss in battery banks due to aging and charging sequences, quality of service constraints, while monitoring the situation in real time. Lab tests have predicted savings of up to 71% in electricity cost, compared to the state-of-the-art.

CloudProphet, on the other hand, is aimed at dealing with the black boxes that data centers have to accommodate.

"Companies like Google and Amazon provide virtual machines for customers," explains Prof. Atienza, "but these customers do not tell you anything about what they are actually doing, and we are not usually permitted to look inside. Using a mixed black-white box analysis is possible in case the particular company/entity owns the applications, as we have tried with a few industrial partners affiliated with the project. But otherwise the behavior of applications is hard to predict - black boxes."

By identifying application processes within the black box, and basing performance prediction only on hardware counter information, CloudProphet learns to anticipate an application's demands on resources. Neural networks multiply in a balletic form of machine learning, building up a picture of the predicted requirements.

"Data centers do have diagnostic tools for the identification and performance prediction of applications, but they can only claim a success rate of 18%," explains Darong Huang. "We are achieving results that are orders of magnitude above that. We hope that CloudProphet will pave the way to a more intelligent resource management system for modern data centers, thus reducing their carbon footprint."

"For these projects, we have PhD and Post-Docs funded by industry," explains Prof. Atienza, "but they can do the work from here, maintaining systems at a distance. The case of China was more complicated because of



bandwidth issues, so we set up a replica of a Huawei facility here.


"Industry project managers are invited to give advice on the physical and logistical constraints in place, so that we get as close as possible to a real-world application.

"Once we have hard data about the reduction in the carbon data footprint, we can decide whether the next step is to license the software, or even to start a spin-off company."

All this is in the future, and nobody can predict the future. Although with the benefit of machine learning, knowledge-based systems, research collaboration and dialogue with industry, we can get close...



Darong Huang
of ESL



SUSTAINABLE TECHNOLOGY FOR SKACH GET READY FOR EXORBITANT AMOUNTS OF DATA



Jean-Paul Kneib
Laboratory of Astrophysics,
SKACH Scientific Delegate to
the SKAO Council

The Square Kilometre Array Observatory (SKAO) is a radio-astronomy project that will consist of two radio telescope arrays in Australia and South Africa, on which construction began December 5, 2022. The data that come from these telescopes will revolutionize our understanding of the universe and its origins, as well as the laws of fundamental physics. It promises to have a major impact on society, in science and beyond.

SKACH is the organisation that will be delivering Swiss contributions to SKAO, and is led by Prof. Jean-Paul Kneib. During its operation, the SKAO will generate a flow of 707 petabytes of data per year. As part of this amazing project, participating scientists will need to handle these large data streams, preferably in the most energy-efficient way possible. This is where EcoCloud comes in.

Right now, EcoCloud scientists Dr. Denisa Constantinescu and Dr. Miguel Peón Quirós are looking at ways that the mountains of data that the SKAO will need to produce and move can be handled with as little energy expenditure as necessary.

The collaboration, which began in September of this year, was borne out of the EcoCloud Annual Event, organized in May 2022. Researchers from SCITAS, who work on high-performance computing at EPFL, presented their work for SKACH, and a new collaboration kicked off.

"We saw that the type of computation they are doing is similar to what we have been doing for the biomedical domain at a different scale," says Peón Quirós. "And that they have exactly the same problem now, that they are less concerned about speed of calculation than they are about energy consumption."



Because EcoCloud already has this expertise in the biomedical domain, Peón Quirós and Constantinescu believed that they might be able to help with SKACH's big data problem, and are now in the middle of a six-month exploratory process to see if the collaboration could be fruitful.

"Right now, I'm just profiling the server to see exactly what kind of computation is in there, see if maybe we can optimize any of the kernels, and what exactly is the best kind of accelerator that could be used for each one of these kernels," explains Constantinescu, "then we'll explore different kinds of algorithms, different ways of implementing the same thing, looking for the most efficient imaging process. We'll test different hypotheses and see which is going to bring the most benefit in terms of energy consumption.

"It is likely that we will be using heterogenous computing, specialized architectures, for different parts of the whole pipeline." Heterogenous computing involves the bringing together of different types of processors, with a particular task in mind. The huge quantities of data involved in this project mean that there will be a great benefit from processing hardware that is not only diverse, but also made-to-measure.

Then there is the time element: the SKAO project is projected to last for many years, so making the data processing energy efficient has important long-term implications.

"If you want it running for 50 years, think of the prospect of providing enough energy for such a long time," suggests Constantinescu. "If you could get that down to zero energy costs, that would be wonderful. This is pure fantasy, but our goal will be to get as close to that as possible."

Constantinescu and Peón Quirós will continue this initial exploratory phase, which is funded by EcoCloud, through the beginning of 2023. In the meantime, they are applying for funding to continue this promising line of research.

SUSTAINABLE TECHNOLOGY FOR SKACH

GET READY FOR EXORBITANT AMOUNTS OF DATA



Denisa Constantinescu and Miguel Peón Quirós of EcoCloud

Just like the SKAO itself, this will be a long-term project. Depending on what they come up with, it may require not only making it work, but also figuring out how to mass produce this new technology.

"Right now we are making a promise, we are defining research," explains Constantinescu, "and we are hopeful that we will get something interesting. But we don't have anything concrete as yet. We have just an informed intuition with which we are formulating hypotheses for how it should work."

The Square Kilometre Array Observatory is gearing up to be fifty times more sensitive than any telescope currently in existence. But this will only be useful if the resulting data is efficiently accessible to scientists in the long term.

EcoCloud Center researchers relish a challenge, but this one involves exorbitant amounts of data, over decades, with the ever-present demands for low energy: it's huge. In fact, it's astronomical.



EXTENDABLE HETEROGENEOUS ENERGY-EFFICIENT PLATFORM:

x-heep

The advancement of continuous healthcare monitoring depends on the development of new and more efficient ultra-low power wearable platforms as well as new algorithms. However, the teams that develop the algorithms are usually not the same as those that design the platforms, thus, optimization opportunities are often lost in the way. Not less importantly, the availability of widespread high-performance, low-power edge-computing platforms will help to alleviate the reliance on cloud data centers and contribute to a more distributed use of energy compatible with local energy generation or harvesting.

To ensure the continuing development of new algorithms for biosignal processing, EcoCloud aims to develop a prototyping environment that allows algorithm developers to assess the implications of their high-level decisions, without overwhelming them with the low-level platform details. This project is based on the open-source x-heep (eXtendable Heterogeneous Energy-Efficient Platform) developed at the Embedded Systems Laboratory (ESL) of the EPFL. The x-heep project is a fully open-source RISC-V microcontroller, and in fact, it will join the Open Hardware Group to be adopted by the open-source community at large as a reference RISC-V microcontroller.

"x-heep can be configured to target small and tiny platforms as well as extended with accelerators for more performance demanding applications," says Dr. Davide Schiavone, the leader of the project at ESL. "The interesting thing about x-heep is that we provide a simple, customizable MCU composed by an open-source, industrial-grade RISC-V CPU, common peripherals (e.g., UART, SPI, JTAG), and memories, and the users will extend it with their own accelerators in an agile way,



without the need of modifying internally the MCU, but only by instantiating it in their custom design. By doing so, the user inherits an IP capable of booting an RTOS (such as freeRTOS) with the whole SW stack, including drivers and SDK, and focus only on building the special HW that makes the user's application unique."

x-heep targets both FPGA and ASIC implementations. In fact, the support for Xilinx-based FPGAs is available open-source, while its first silicon implementation was just delivered last month. Once the chip (called HEEPocrates) is back at ESL, it will be tested and characterized in real-life settings. The measured results will be used by EcoCloud so that programmers will be able to estimate the energy consumption of their algorithms running directly on a commonly available FPGA.

EcoCloud will exploit the FPGA version of x-heep as baseline and extend it with additional characteristics such as virtual ADCs and time compression to enable easy prototyping and benchmarking of biomedical algorithms. EcoCloud will work to ensure that the platform is as accessible as possible for a wide range of users - particularly those in other domains such as bioengineering and the mathematical development of new biomedical algorithms.



CONTACT

ECOCLOUD



Exploring new paradigms in data center hardware technology, pioneering strategies of cloud data management and innovative security techniques, in collaboration with industrial partners



EcoCloud

EPFL ECOCLOUD
INJ 234 (Bâtiment INJ)
Station 14
CH-1015 Lausanne
Tel: +41 21 693 13 24
contact.ecocloud@epfl.ch
<https://ecocloud.epfl.ch>

Prof. David Atienza

Chief Scientific Officer
EPFL STI IEM ESL
ELG 130 (Bâtiment ELG)
Station 11
CH-1015 Lausanne
Tel: +41 21 693 11 31
david.atienza@epfl.ch

Valérie Locca

Administrative Assistant and IAP Manager
EPFL ECOCLOUD
INJ 234 (Bâtiment INJ)
Station 14
CH-1015 Lausanne
Tel: +41 21 693 13 24
support.ecocloud@epfl.ch

CREDITS

AND RIGHTS




Texts copyright EcoCloud 2022, written by

4	David Atienza
5-6	Leila Ueberschlag
10-13	Babak Falsafi and David Atienza
14-15	John Maxwell and David Atienza
16-18	John Maxwell
19-20	John Maxwell and Tanya Petersen
21-22	John Maxwell and Tanya Petersen
23-24	John Maxwell
25-26	John Maxwell
27-28	Stephanie Parker and John Maxwell
29	Miguel Peón Quirós and Davide Schiavone

Image Credits - All Rights Reserved

1	Peach-adobe, Adobe Stock: 470470740
4	David Atienza, EPFL
5, 11	Eduardo Garland, Adobe Stock: 317397107
7-9	EPFL
10	cherezoff, Adobe Stock: 467673383
12	Issaranow, Adobe Stock: 280047831
14-15	Who is Danny, Adobe Stock: 541155495
	EPFL
16	Elison Matioli, EPFL; Alex Widerski, EPFL
17	Alex Widerski, EPFL
18	Remco van Erp, Corintis
19	Ulia Koltyrina, Adobe Stock: 236642518
	Margin Jaggi, EPFL
20	Alex Widerski, EPFL
21	LunaSec; Matthias Payer, EPFL
	ZinetroN, Adobe Stock: 367623665
23	Babak Falsafi, Mathias Payer, David Atienza, EPFL
	Abhishek Bhattacharjee, Yale
	Boris Grot, University of Edinburgh
	ZinetroN, Adobe Stock: 367623665
24	EPFL, first printed in the Proceedings of ISCA '21
25	Dlgilife, Adobe Stock: 392825932
	David Atienza, EPFL; Luis Costero, UCMadrid
	Ipopba, Adobe Stock: 323829465
26	Darong Huang, EPFL
	Adobe Stock: 480306328
27	Jean-Paul Kneib, EPFL; SKAO
28	John Maxwell, EPFL
29	Christoph Müller, EPFL
30	Ipopba, Adobe Stock: 480306328



EPFL