Collegium Helveticum

Digital Societies

Friday, December 8, 2017, 10:10–18:00

After Digital? Emerging Computing Paradigms

Workshop

In Cooperation with Università della Svizzera italiana (USI) and École polytechnique fédérale de Lausanne (EPFL)

The Workshop is part of the focal topic «Digital Societies» of the fellow period 2016–2020 at Collegium Helveticum.

Università della Svizzera italiana

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

ETH zürich

Universität Zürich UZH

hdk

Science and Society have been transformed by computing.

Yet, the existing computing paradigm is becoming insufficient and ineffective as we are reaching the end of the emblematic Moore's law. The energy demands of CMOS-based computers are soaring. Current and foreseeable demands on computing capacity, storage and communications are reaching the limitations of physical laws in terms of reliability and economics. At the same time, the vast volume, veracity, velocity and variety of data are challenging classical methods of thinking and the very foundations of computing.We are in need for alternatives and, indeed, new and powerful technologies such as Quantum, Bio-inspired, and Crowd computing, have begun to emerge. For the first time in 80 years, our generation is presented with a unique opportunity to restructure and reinvent computing as an enabling technology and an intellectual endeavor. Beyond advances that require harnessing these emerging computing paradigms, we have the opportunity to make substantial contributions based on lessons learned. In the past, computing developed across distinct branches of mathematics, software and hardware. Later efforts, aiming for integration, largely failed due to disciplinary and structural boundaries. We believe that future developments should traverse disciplinary boundaries to solve the pressing problems of our society by developing the mathematical, algorithmic and software innovations required to enable this transformation by taking advantage of emerging modes of computing technologies.

Program

Friday, December 8, 2017

10:10–10:15
Introduction
Petros Koumoutsakos

10:15–11:00
Talk 1
Cambrian Explosion of Computing in the Post-Moore Era
Satoshi Matsuoka

11:00–11:45
Talk 2
SpiNNaker – biologically-inspired massively-parallel computing
Steve Furber

11:45–13:15
Lunch Break

13:15–14:00
Talk 3
Quantum computers - models, applications, platforms
Frank Wilhelm-Mauch

14:00–14:45
Talk 4
Overcoming the Power Wall by Exploiting Application Inexactness
Sven Leyffer

14:45–15:00
Coffee Break

15:00–15:45
Talk 5
Crowdsourcing as a Tool for Research and Public Engagement
Edith Law

15:45–16:30
Talk 6
Scale-dependent reduced precision for weather and climate modelling
Tim Palmer

16:30–17:00
Break

17:00–18:00
Round Table Discussion

18:00
Apéro

Speakers

Steve Furber
University of Manchester, UK

Edith Law
University of Waterloo, Canada

Sven Leyffer
Argonne National Laboratory, USA

Satoshi Matsuoka
Tokyo Tech and Advanced Institute of Science and Technology, Japan

Tim Palmer
Oxford University, UK

Frank Wilhelm-Mauch
Saarland University, Germany

Steve Furber
University of Manchester, UK

SpiNNaker – biologically-inspired massively-parallel computing

The SpiNNaker (Spiking Neural Network Architecture) project has delivered a machine incorporating half-a-million ARM processor cores designed primarily for brain-modelling applications, but it can also be used to explore how biological inspiration may advance the capabilities of artificial neural networks and machine learning in the future.

Edith Law
University of Waterloo, Canada

Crowdsourcing as a Tool for Research and Public Engagement

Science is increasingly data-intensive; yet, many research tasks involving the collection, annotation and analysis of data are too complex to be fully automated. The idea of research-oriented crowdsourcing is to engage online workers or volunteers to contribute and process data towards an academic inquiry. In this talk, I will discuss the variety of research questions that arise from the process of designing crowdsourcing systems for research and public engagement, and illustrate our approach through a description of several ongoing projects on CrowdCurio, our experimentation platform for crowdsourcing.

Abstracts

Sven Leyffer
Argonne National Laboratory, USA

Overcoming the Power Wall by Exploiting Application Inexactness

Energy and power consumption are major limitations to continued scaling of computing systems. Inexactness, where the quality of the solution can be traded for energy savings, has been proposed as an approach to overcoming those limitations. In the past, however, inexactness necessitated the need for highly customized or specialized hardware. The current evolution of commercial off-the-shelf (COTS) processors facilitates the use of lower-precision arithmetic in ways that reduce energy consumption. We study these new opportunities, using the example of an inexact Newton algorithm for solving nonlinear equations. We also describe a set of techniques that use reduced precision to improve the quality of the computed result by reinvesting the energy saved by reduced precision.

Satoshi Matsuoka
Tokyo Tech and Advanced Institute of Science and Technology, Japan

Cambrian Explosion of Computing in the Post-Moore Era

The so-called “Moore’s Law”, by which the performance of the processors will increase exponentially by factor of 4 every 3 years or so, is slated to be ending in 10–15 year timeframe due to the lithography of VLSIs reaching its limits around that time, and combined with other physical factors. Although there are multitudes of work trying to address this imminent and ultimate threat, most of the work are point solutions to limited scope of problems domains, and moreover, lack the holistic view on (1) how to advance computing on some conventional algorithms such as PDE solvers that has no apparent acceleration techniques, and (2) create a total system architectural stack of post-Moore systems in a holistic fashion, from devices and hardware, abstracted by system software and programming models and languages, and optimized according to the device characteristics with new algorithms and applications that exploit them. Such systems will have multitudes of varieties according to the matching characteristics of applications to the underlying architecture, leading to what can be metaphorically described as Cambrian Explosion of computing systems.

For example, the promising new parameter in place of the transistor count is the perceived increase in the capacity and bandwidth of storage, driven by device, architectural, as well as packaging innovations: DRAM-alternative Non-Volatile Memory (NVM) devices, 3-D memory and logic stacking evolving from VIAs to direct silicone stacking, as well as next generation terabit optics and networks. Here, exploiting the memory and bandwidth capacities will instead be the acceleration methodology. However, such shift in compute-vs-data tradeoffs would not exactly be return to the old vector days, since other physical factors such as latency will not change when spatial communication is involved in X-Y directions, as well as other parameters, and system software, programming, algorithms, and applications will have to cope with the new parametric changes. Such will be the case for all means of acceleration, such as neuromorphic computing, quantum computing, etc., and we need to advance research, especially in algorithms and computing systems, to address such changes as a whole.

Tim Palmer
Oxford University, UK

Scale-dependent reduced precision for weather and climate modelling

The development of stochastic parametrisation for weather and climate models has opened the possibility of working with significantly reduced numerical precision in weather and climate modelling, thereby reducing data transport inside computers. Results using the spectral IFS show the ability to run globally at single precision without degradation with 40% saving in run time. Using emulators it is shown possible to run at almost half precision for large parts of the IFS dynamical core. Work with FPGAs shows the ability to code reduced precision in hardware. I will conclude with some speculative remarks about the possible use of neural nets run at half precision and trained on output from parametrisations.

Frank Wilhelm-Mauch
Saarland University, Germany

Quantum computers – models, applications, platforms

Quantum computing uses the unique properties of quantum physics to speed up some computational tasks. A core principle for this acceleration is their ability to operate fully parallel operations on a single instance of quantum hardware. There are three currently pursued to approaches to realize this promise: Circuit-based quantum advantage for problems, universal fault-tolerant quantum computing, and quantum annealing. These are suited for different types of problems and have different requirements on hardware, reflected in a different status of implementation.

There is a wealth of implementation candidates for quantum computing from atomic and solid-state physics - currently led by trapped ions and superconducting circuits with a few interesting runners-up such as semiconductor platforms. The current understanding of these platforms and related platforms allows for speculations about the future development of quantum computing including a range of milestones that could be reached in the next decade.

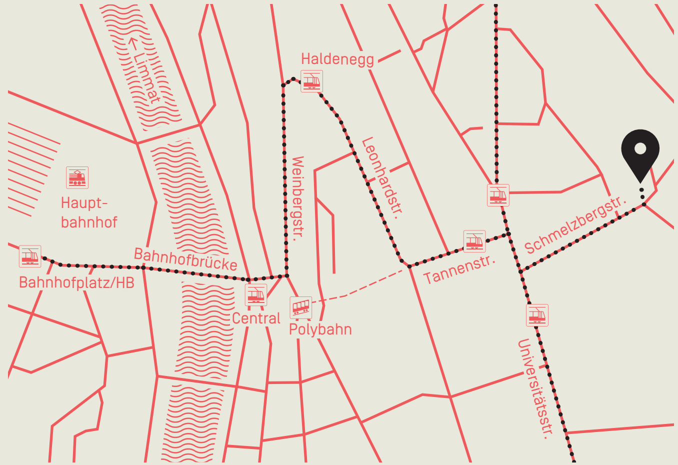

Venue

Collegium Helveticum
Semper-Sternwarte
Schmelzbergstrasse 25
8006 Zürich

Orgnisers

Petros Koumoutsakos, ETH Zürich,
Fellow am Collegium Helveticum
Olaf Schenk, USI, Lugano
Jan S Hesthaven, EPFL, Lausanne

Registration

Since the number of participants is limited, registration is necessary and will be processed in the order of arrival.

Please register at lewiss@ethz.ch (Susanne Lewis).
Deadline is December 4, 2017.

Laboratorium für Transdisziplinarität
Schmelzbergstrasse 25 • CH-8092 Zürich • www.collegium.ethz.ch