TORSTEN HOEFLER & DAVID KEYES

The PASC16 Technical Papers Selection

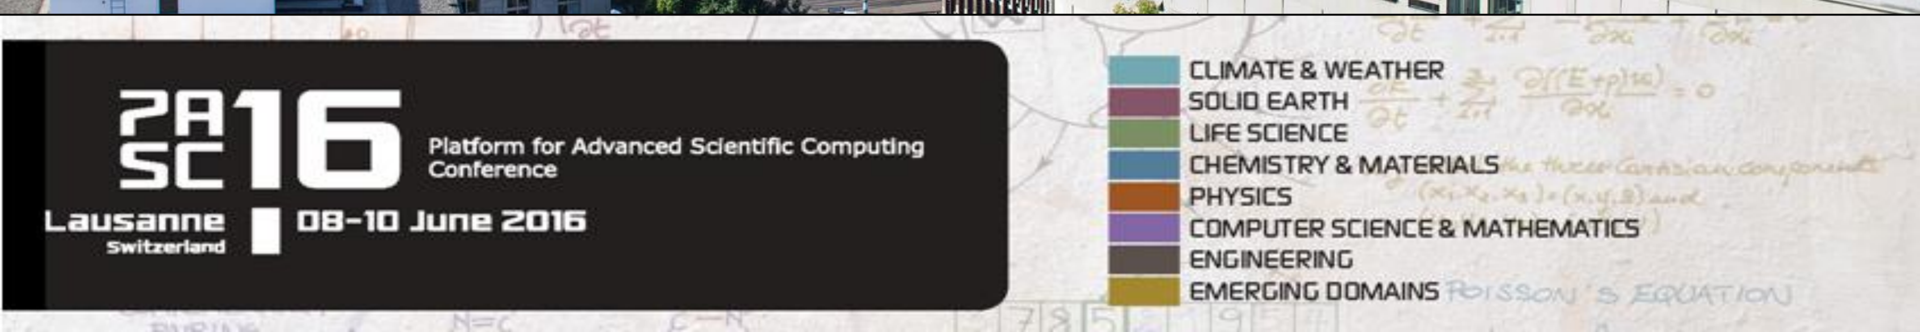

Bridging communities

Journals

acm sighpc

PASC16
Lausanne Switzerland | 08-10 June 2016

Conferences

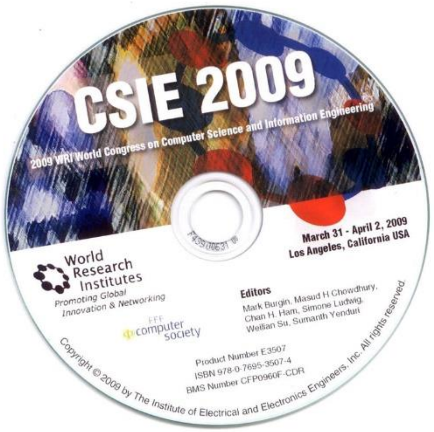

- Considered mature publications
- Thorough revision process
- Expert reviewers for each submission
- Long process (~1 year)
- No dissemination component
- …

- Top-class in computer science
- Very quick turn-around (4-6 months)
- Streamlined review process
- Dissemination at conferences
- Pre-selected committee
- Rebuttals are a waste of time
- …

Bridging communities

Journals

acm sighpc

PASC16
Lausanne Switzerland
08-10 June 2016

Conferences

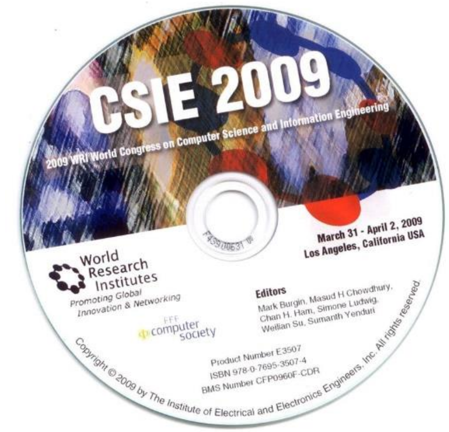

Papers Co-Chairs


Proceedings Chair

Area Editors:
- Michael Wehner
- Omar Ghattas
- George Biros
- Ioannis Xenarios
- Mark von Schilfgaarde
- George Lake
- Jeroen Tromp

The PASC process: four pillars

No pre-selected committee
- Area editors pick reviewers
- More appropriate reviewers
- More work for the chairs

Short revision process
- Two-week revision
- Similar to journals (no rebuttal)
- Pressure on authors

acm sighpc

PASC16
Lausanne Switzerland | 08-10 June 2016

Fully double-blind
- Blind to reviewers and chairs
- Reduces bias significantly
- COI management harder

Suggested Expert Reviews
- Round-1 reviewers asked
- Improved expertise in round 2
- Potential inconsistencies

Submissions overview:

- **44 submissions (stage 1)**
  - Authors: US: 53, CH: 43, UK: 9, SA, JP, FR: 6, Others: 26
- **Most in Math & CS, others reasonably balanced**

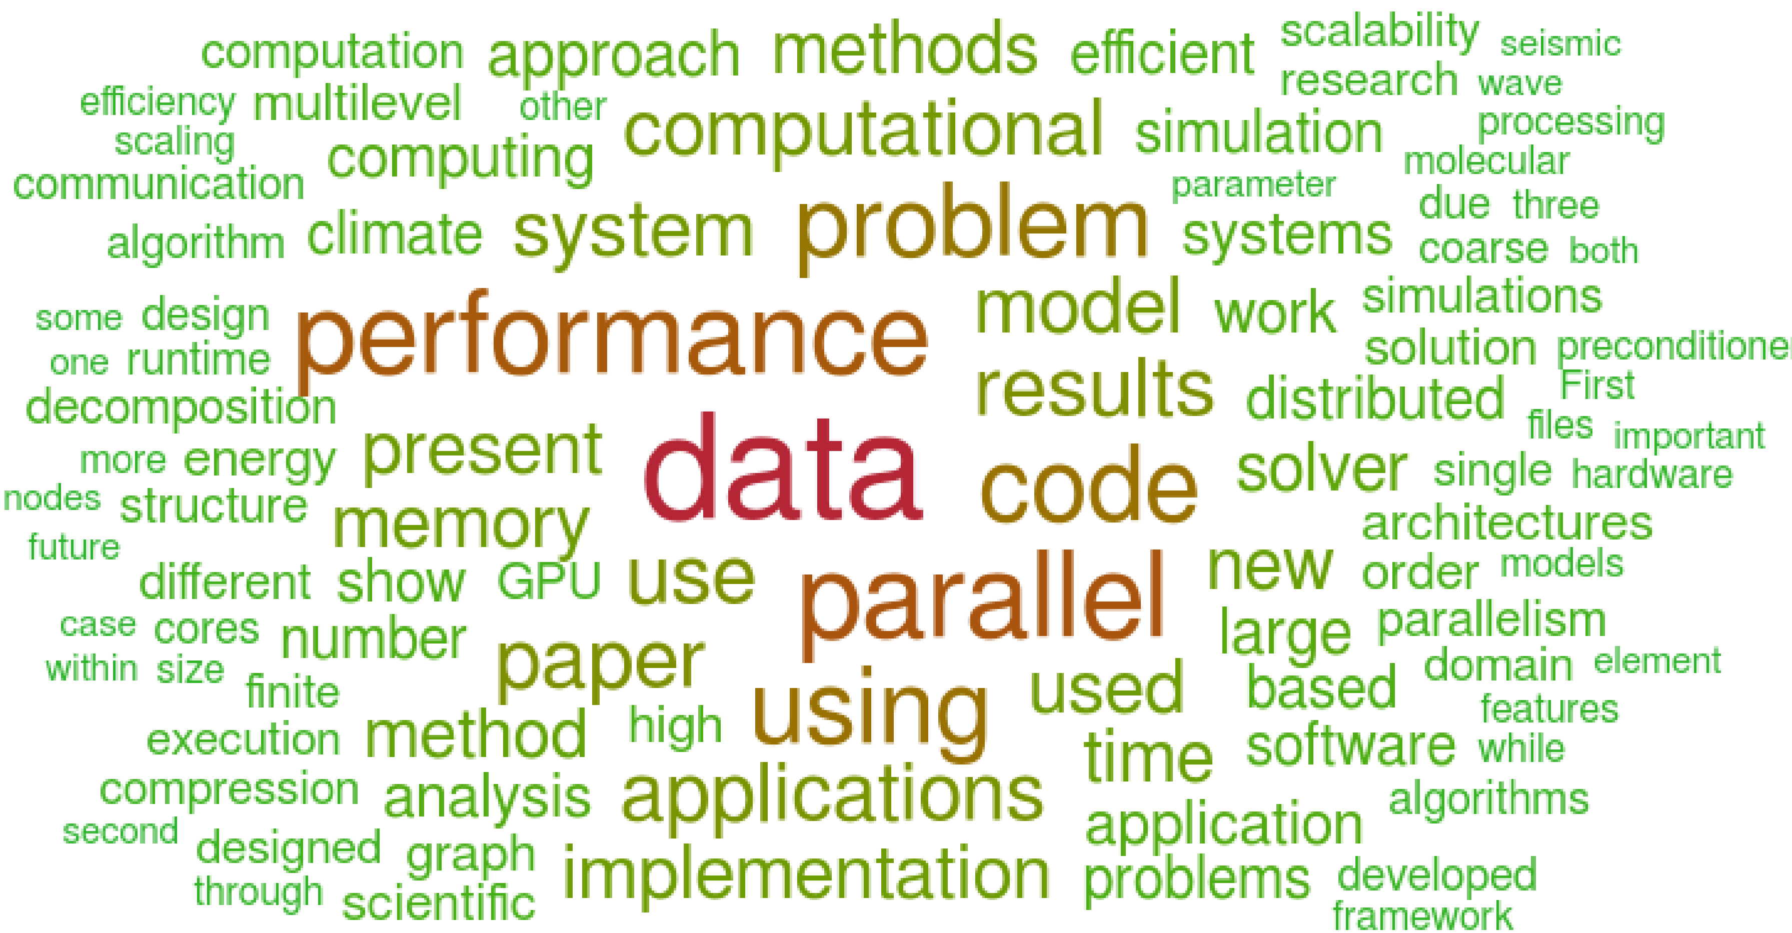

Reviews and acceptance

- **182 reviews total (two stages, ~60k words total)**
- **12 accepted papers (each paper discussed in physical meeting)**
  - Authors: CH: 12, US: 8, Fr, JP: 6, Others: 16

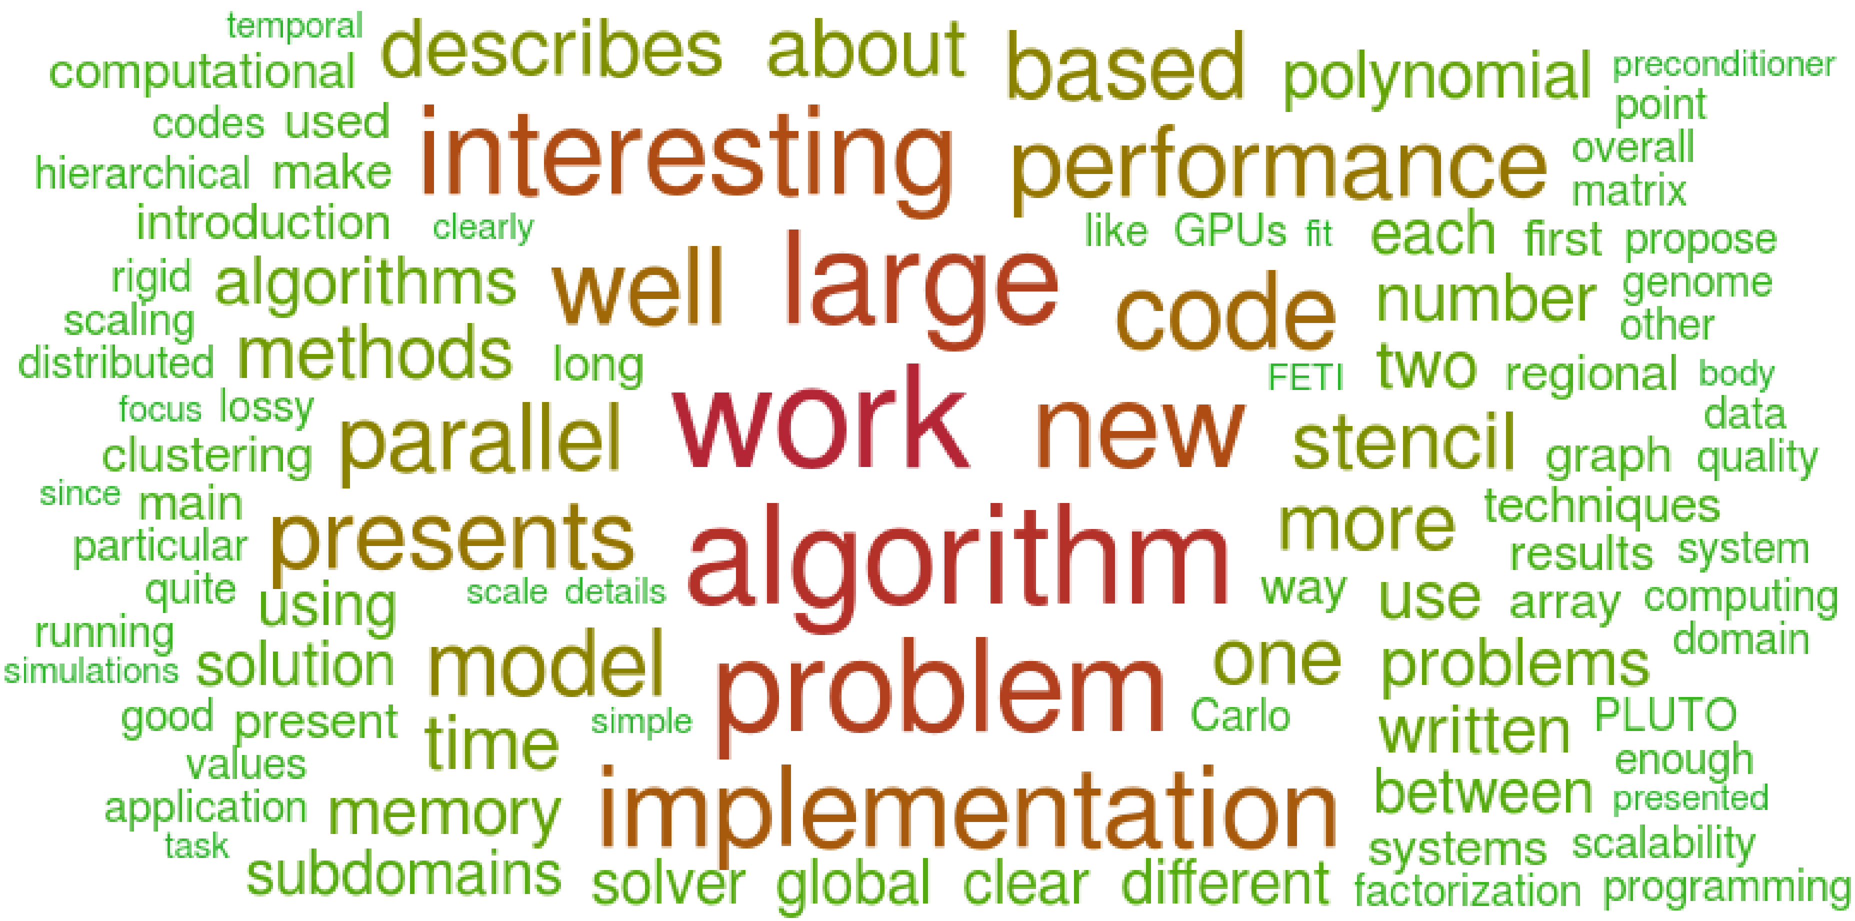

Selection purely based on scientific excellence

- **Carefully briefed expert reviews**
  - 23 papers were invited to stage 2
    Were asked to mark differences made in revision
  - Full review after revision (+ recommended experts)
- **Face-to-face meeting in Lausanne (1 day)**
- **Discussed each paper, asked questions**
  - What did I learn while reading the paper? (quality)
  - How many people would attend the talk? (relevance)
  - Would I recommend my colleagues to read it? (presentation)
- **Committee discussion:**
  - Needs session for software frameworks that may have little novelty but huge impact → should be implemented for PASC17 (cf. State of Practice)
- **Mantra: never go against an expert**
  - It was not necessary but could be tough

Impact of expert reviewers

- **Expert reviewers were suggested by reviewers in stage 1**
- **Invited in stage 2 (short review time)**
  - Nearly all agreed (some very enthusiastically)
- **All 23 stage 2 submissions received expert reviews**
  - 2 were accepted due to expert reviews
  - 2 were rejected due to expert reviews
  - 19 did not change (decision reinforced)
- **Most expert reviews were longer than average**
  - Some nearly as long as the paper …

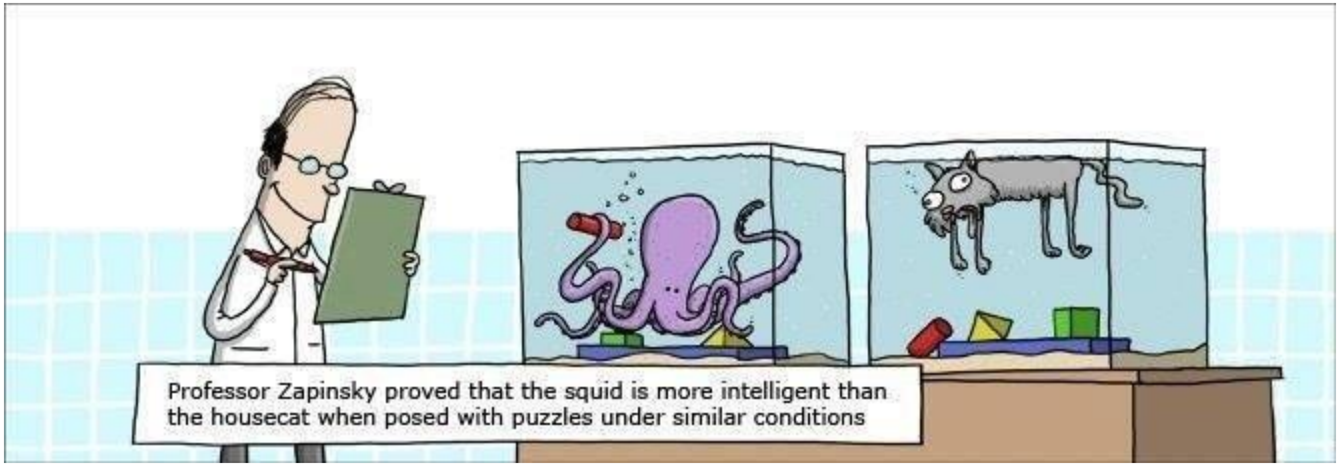

Side note: performance reporting

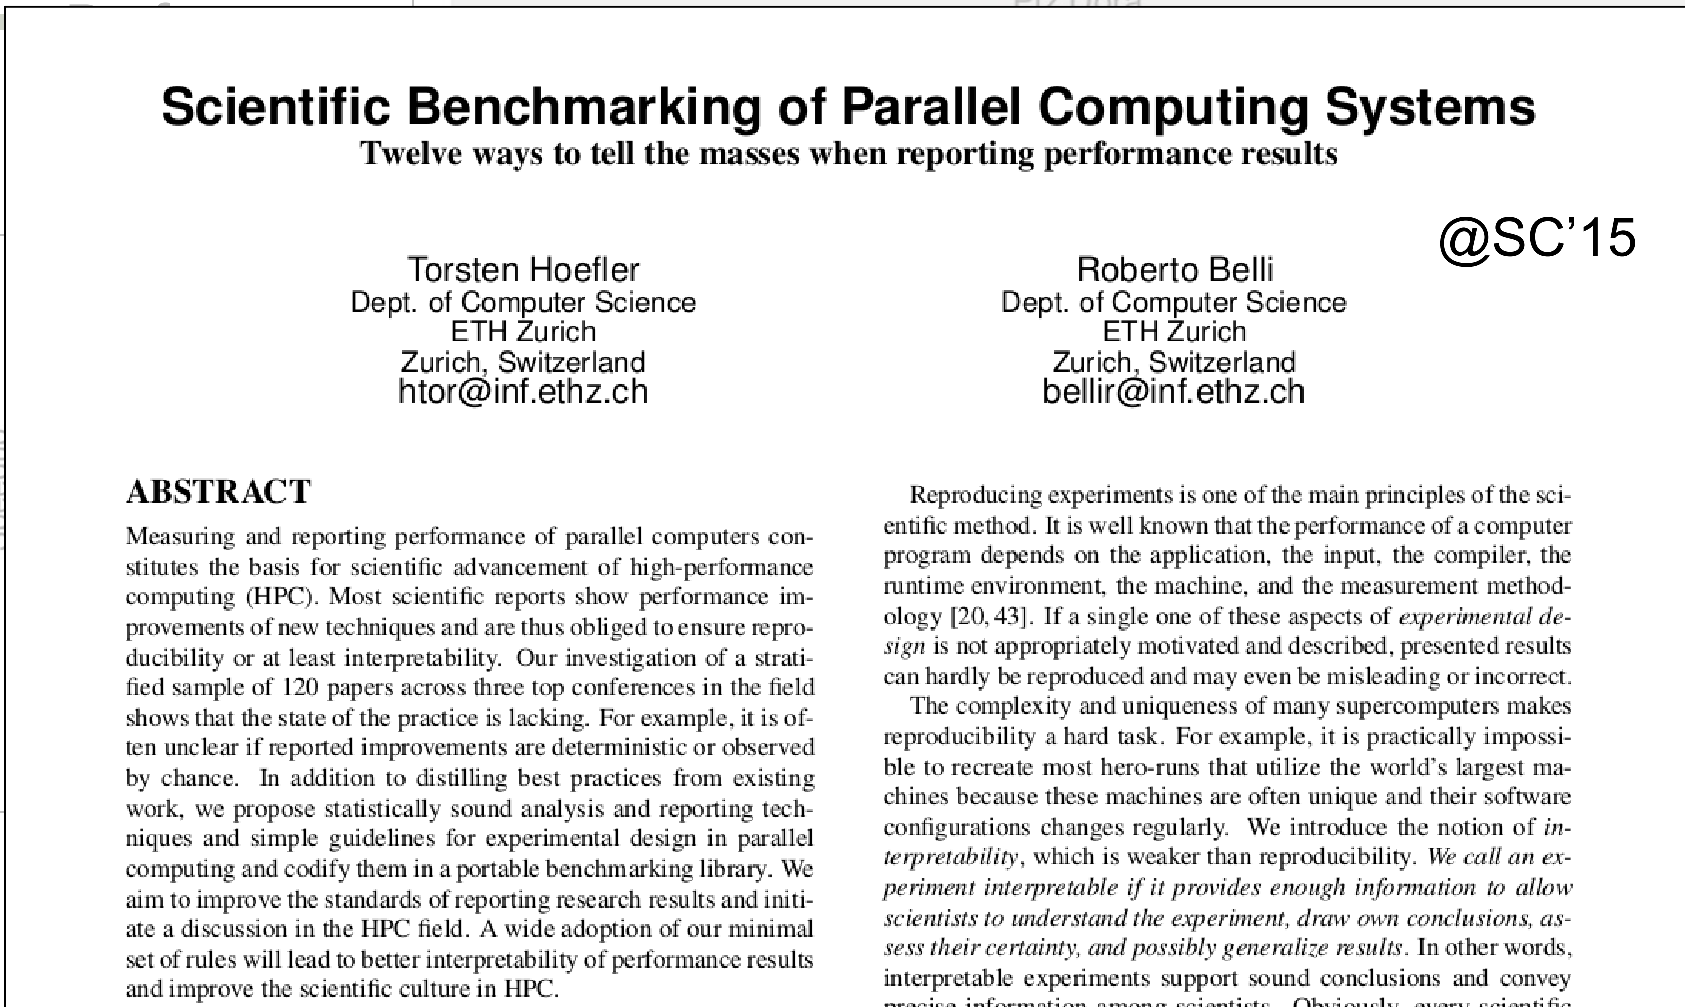

Scientific Benchmarking of Parallel Computing Systems

Twelve ways to tell the masses when reporting performance results

@SC'15

Torsten Hoefler
Dept. of Computer Science
ETH Zurich
Zurich, Switzerland
htor@inf.ethz.ch

Roberto Belli
Dept. of Computer Science
ETH Zurich
Zurich, Switzerland
bellir@inf.ethz.ch

ABSTRACT

Measuring and reporting performance of parallel computers constitutes the basis for scientific advancement of high-performance computing (HPC). Most scientific reports show performance improvements of new techniques and are thus obliged to ensure reproducibility or at least interpretability. Our investigation of a stratified sample of 120 papers across three top conferences in the field shows that the state of the practice is lacking. For example, it is often unclear if reported improvements are deterministic or observed by chance. In addition to distilling best practices from existing work, we propose statistically sound analysis and reporting techniques and simple guidelines for experimental design in parallel computing and codify them in a portable benchmarking library. We aim to improve the standards of reporting research results and initiate a discussion in the HPC field. A wide adoption of our minimal set of rules will lead to better interpretability of performance results and improve the scientific culture in HPC.

Reproducing experiments is one of the main principles of the scientific method. It is well known that the performance of a computer program depends on the application, the input, the compiler, the runtime environment, the machine, and the measurement methodology [20, 43]. If a single one of these aspects of *experimental design* is not appropriately motivated and described, presented results can hardly be reproduced and may even be misleading or incorrect.

The complexity and uniqueness of many supercomputers makes reproducibility a hard task. For example, it is practically impossible to recreate most hero-runs that utilize the world's largest machines because these machines are often unique and their software configurations changes regularly. We introduce the notion of *interpretability*, which is weaker than reproducibility. *We call an experiment interpretable if it provides enough information to allow scientists to understand the experiment, draw own conclusions, assess their certainty, and possibly generalize results*. In other words, interpretable experiments support sound conclusions and convey precise information among scientists. Obviously, every scientific

Side note: overall process

- **The timeframe was way to tight**
  - Issues with ACM sponsorship, will be extended by 2x for PASC17
- **Face-to-face meeting**
  - Very efficient, should be kept
- **Engineering/Software/Experience track**
  - Special session on software systems (potentially high impact)
- **Conflict handling**
  - Can be improved by allowing authors to specify conflicts
- **Chair load**
  - Biggest concern as number of submissions grows

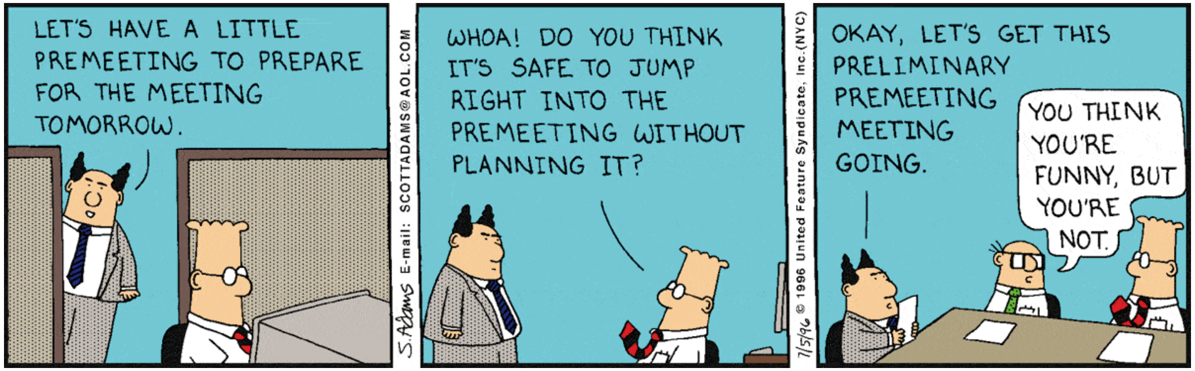

http://dl.acm.org/citation.cfm?id=2929908

Do you have feedback?

Mark Adams	UC Berkeley
Kadir Akbudak	Bilkent University
Srinivas Aluru	Georgia Instit…
William Anderson	NASA
Peter Bastian	University o…
Ugo Becciani	INAF - Cata…
Mauro Bianco	ETH Zurich
Xiao Bo	Baidu Inc.
Ebru Bozdag	University o…
Jed Brown	University o…
Tan Bui-Thanh	University of Texas at Austin
Po Chen	University o…
Diego Darriba	Heidelberg
Bronis de Supinski	Lawrence L…
Sebastian Deorowicz	Silesian Un…
Petros Drineas	Rensselaer Polytechnic Inst…
Stephane Ethier	Princeton P…
Kurt Ferreira	Sandia Nati…
Juan Gómez Luna	Universidad…
Longfei Gao	KAUST
Xin Gao	KAUST
Kai Germaschewski	University o…
Amir Gholami	University o…
Mike Giles	Oxford Univ…
Dominik Goeddeke	University o…
Jorge Gonzalez	University o…
Steven Gottlieb	Indiana University
Ivan Graham	University o…
Damien Gratadour	Paris Observatory
Bill Gropp	University of Illinois Urbana-Champaign
Jean-Luc Guermond	Texas A&M University
Georg Hager	University of Erlangen-Nuremberg
Michael Heroux	Sandia National Laboratories
Des Higham	University of Strathclyde
Jayanth Jagalur Mohan	Massachusetts Institute of Technology
Ajay Jasra	National University of Singapore
Frank Jenko	UC Los Angeles
Jan Christian Kässens	Kiel University
Chandrika Kamath	Lawrence Livermore National Laboratory
Marat Khairoutdinov	Stony Brook University
Scott Klasky	Oak Ridge National Laboratory
Axel Klawonn	University of Koeln

Matthew Knepley	Rice University
Hari Krishnan	Lawrence Berkeley National Laboratory
	University of Tennessee
Dominique Lavenier	IRISA
	Oak Ridge National Laboratory
	Baidu Inc.
	University of Toronto
	Georgia Institute of Technology
	KAUST
	University of Colorado at Boulder
	Sandia National Laboratories
	Lawrence Berkeley National Laboratory
	KAUST
	University of Utah
	UC Merced
	Sandia National Laboratories
	Lawrence Berkeley National Laboratory
	University of Washington
	University of Frankfurt
	University of Manchester
	Technische Universitaet Bergakademie Freiberg
	University of Paris VI
	University of Texas at Austin
	Università della Svizzera italiana
	New York University
	KAUST
	NVIDIA Inc.
	ETH Zurich
Eric Sonnendrucker	Max Planck Institute for Plasma Physics
Joachim Stadel	University of Zurich
Georg Stadler	New York University
Ion Stoica	UC Berkeley
Carl Tape	University of Alaska Fairbanks
Philippe Theirry	Intel Inc.
George Turkiyyah	KAUST
Brian Van Straalen	Lawrence Berkeley National Laboratory
Tobias Weinzierl	Durham University
Olof Widlund	New York University
Lars Wienbrandt	University of Kiel
Stefan Wild	Argonne National Laboratory
Ulrike Yang	Lawrence Livermore National Laboratory
Rio Yokota	Tokyo Institute of Technology