

# On the impact of quantum computing technology on future developments in high-performance scientific computing

Matthias Möller, Numerical Analysis,  
Delft Institute of Applied Mathematics, TU Delft




M. Möller and C. Vuik  
Ethics and Information Technology  
December 2017, Volume 19, Issue 4, pp. 253-269  
DOI: <https://link.springer.com/article/10.1007/s10676-017-9438-0>


# Outline

- High-performance scientific computing
  - Past, present and possible future trends
  - FEM making friends with exotic hardware
- Q-accelerated scientific computing
  - Potential use case examples
  - Challenges and open problems


# History of computing



1971: Intel 4004 first commercial microprocessor



1976: Cray-1 first vector processor



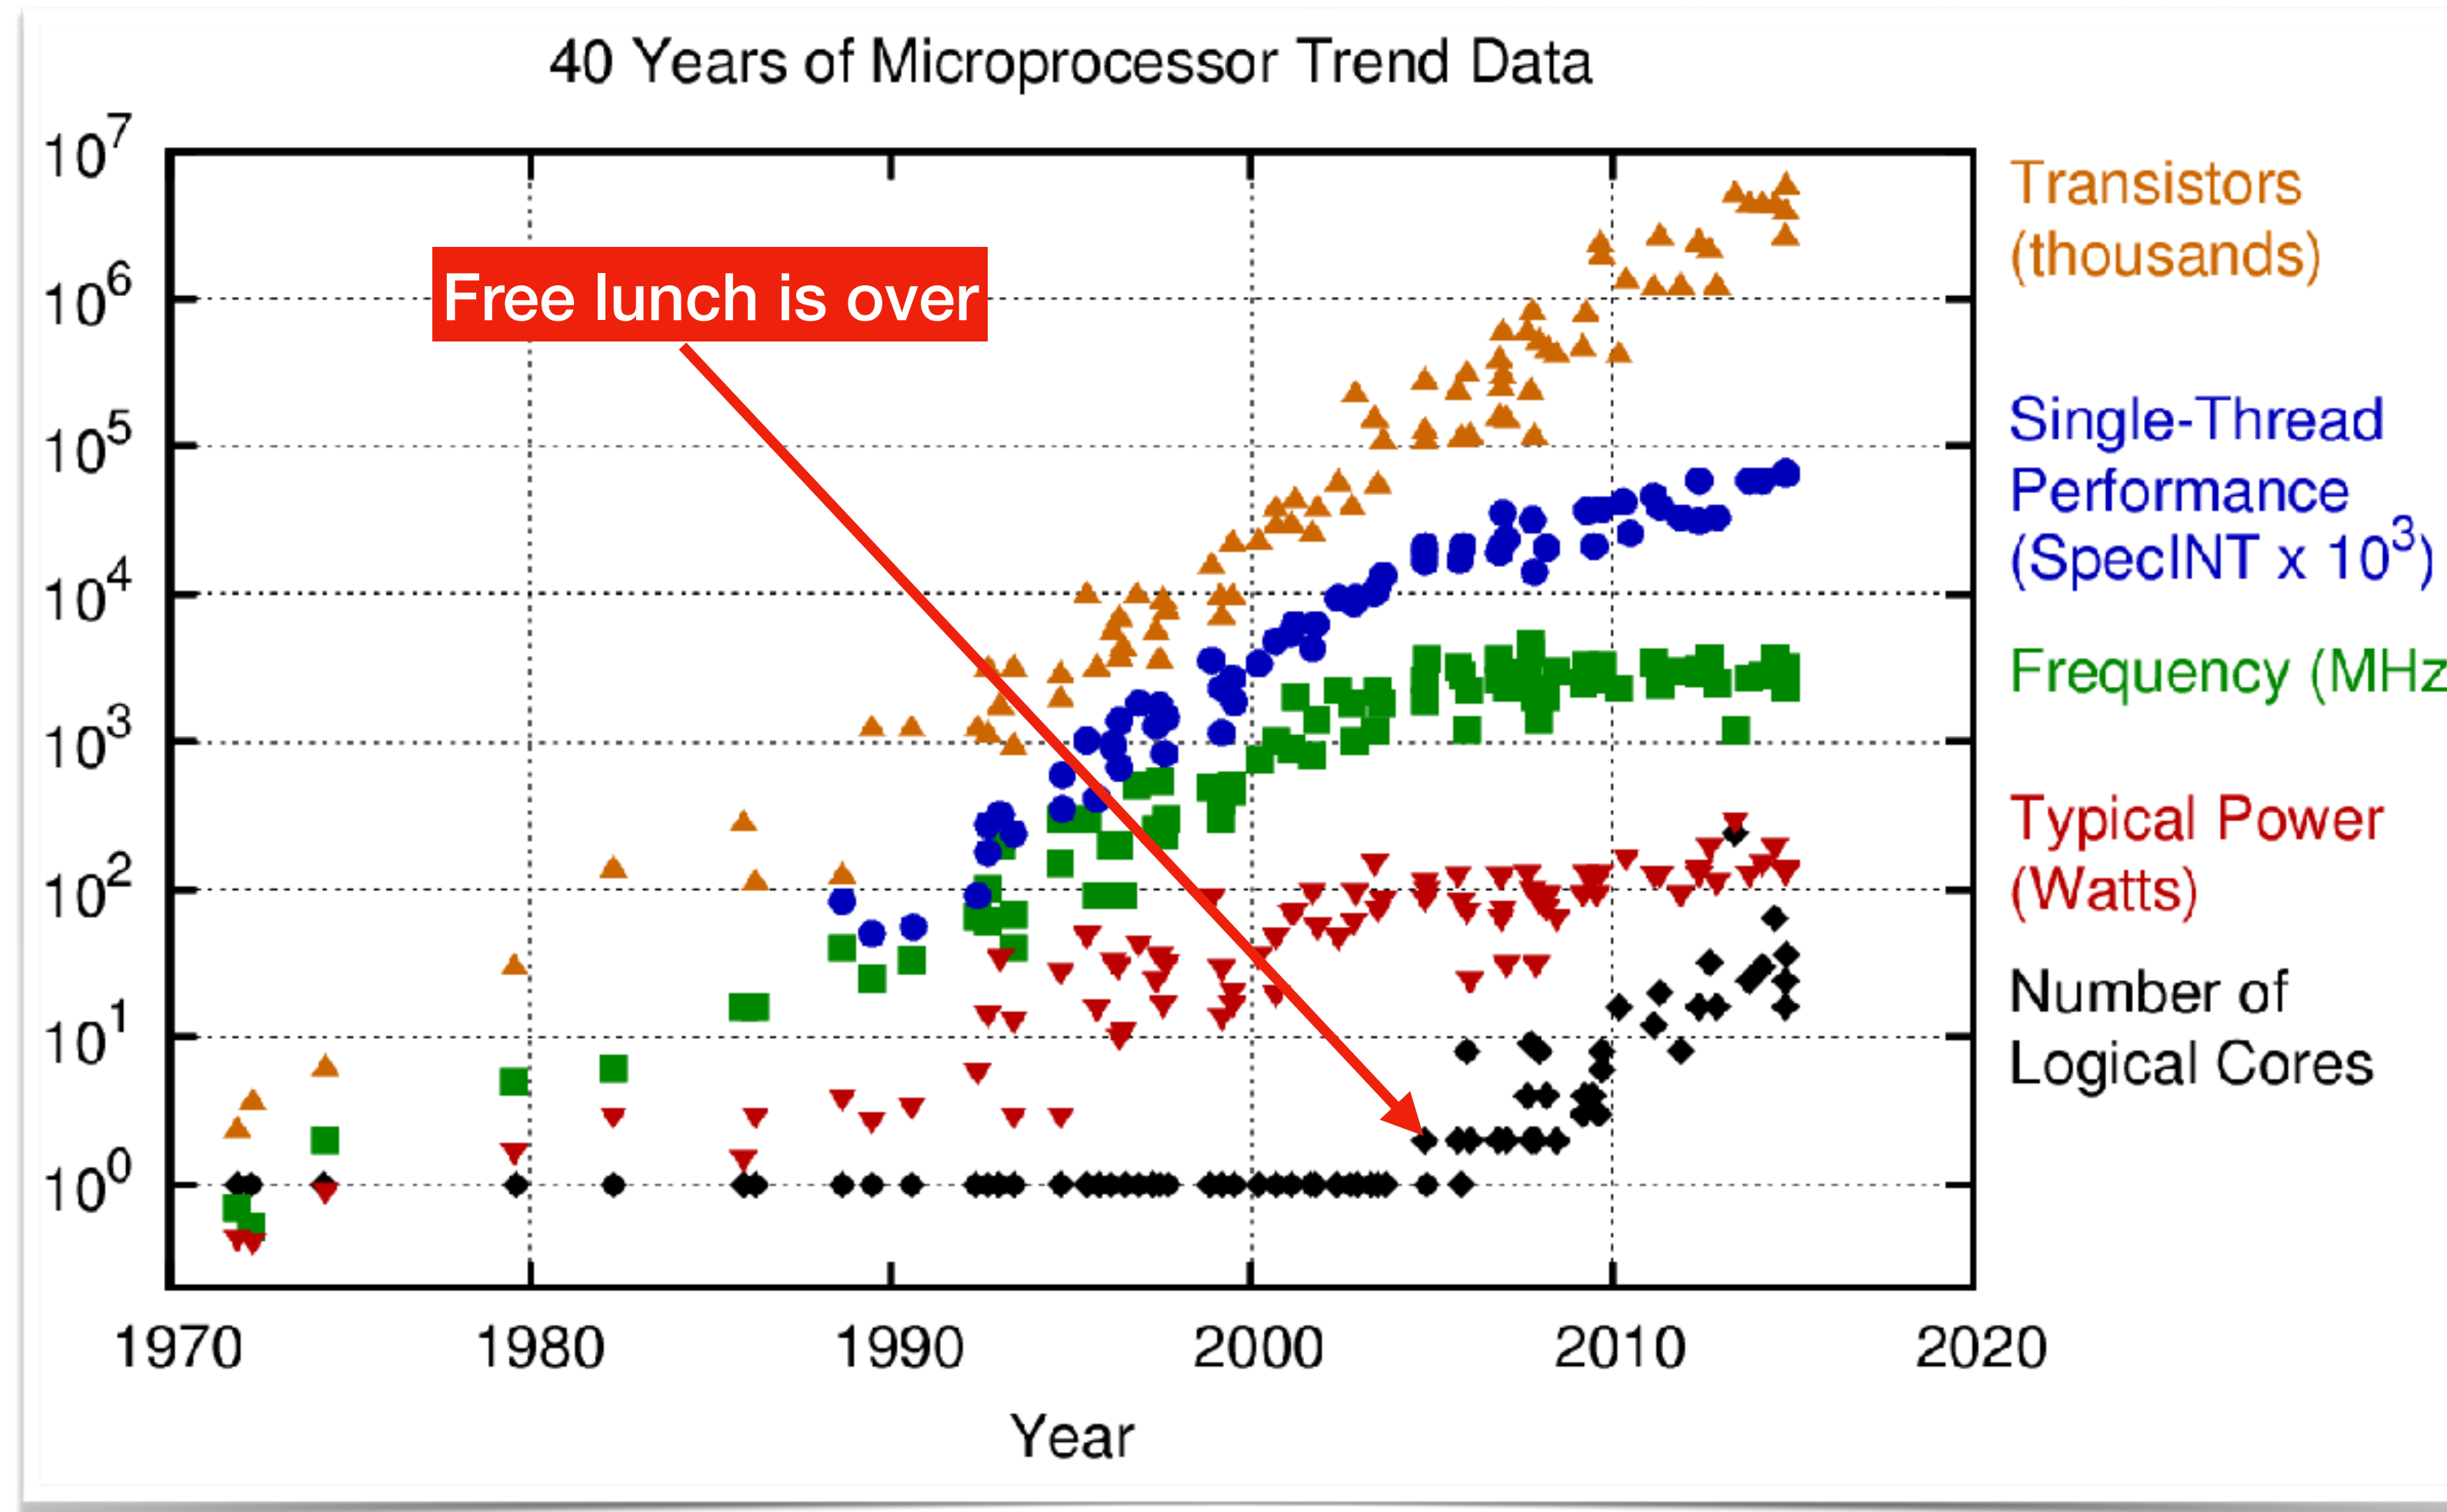
2001: IBM Power4 first multi-core CPU



2017: Intel Xeon Platinum 28 cores and AVX-512

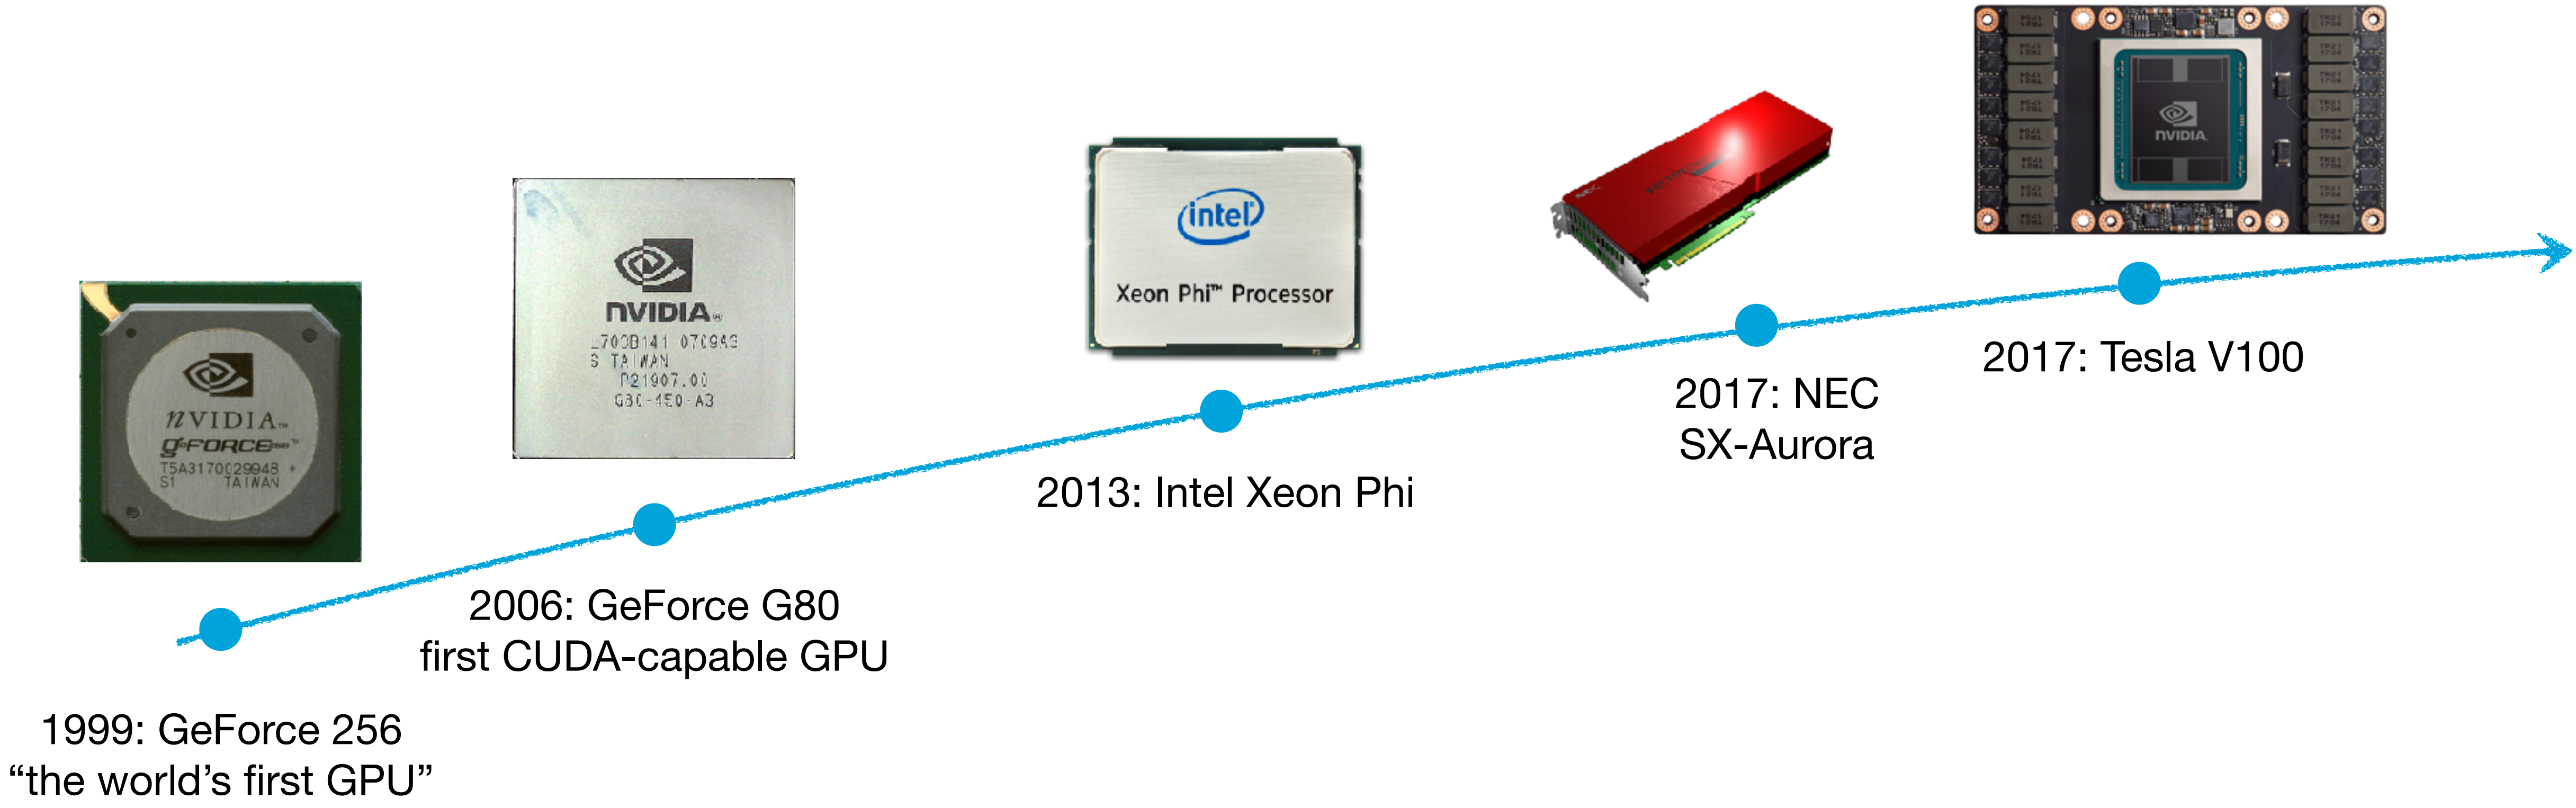
- **Redesign of algorithms** is mandatory to exploit compute capabilities of modern hardware (multi-threading, vectorisation)
- “Algorithm follows hardware”

# History of computing






# History of accelerated-computing

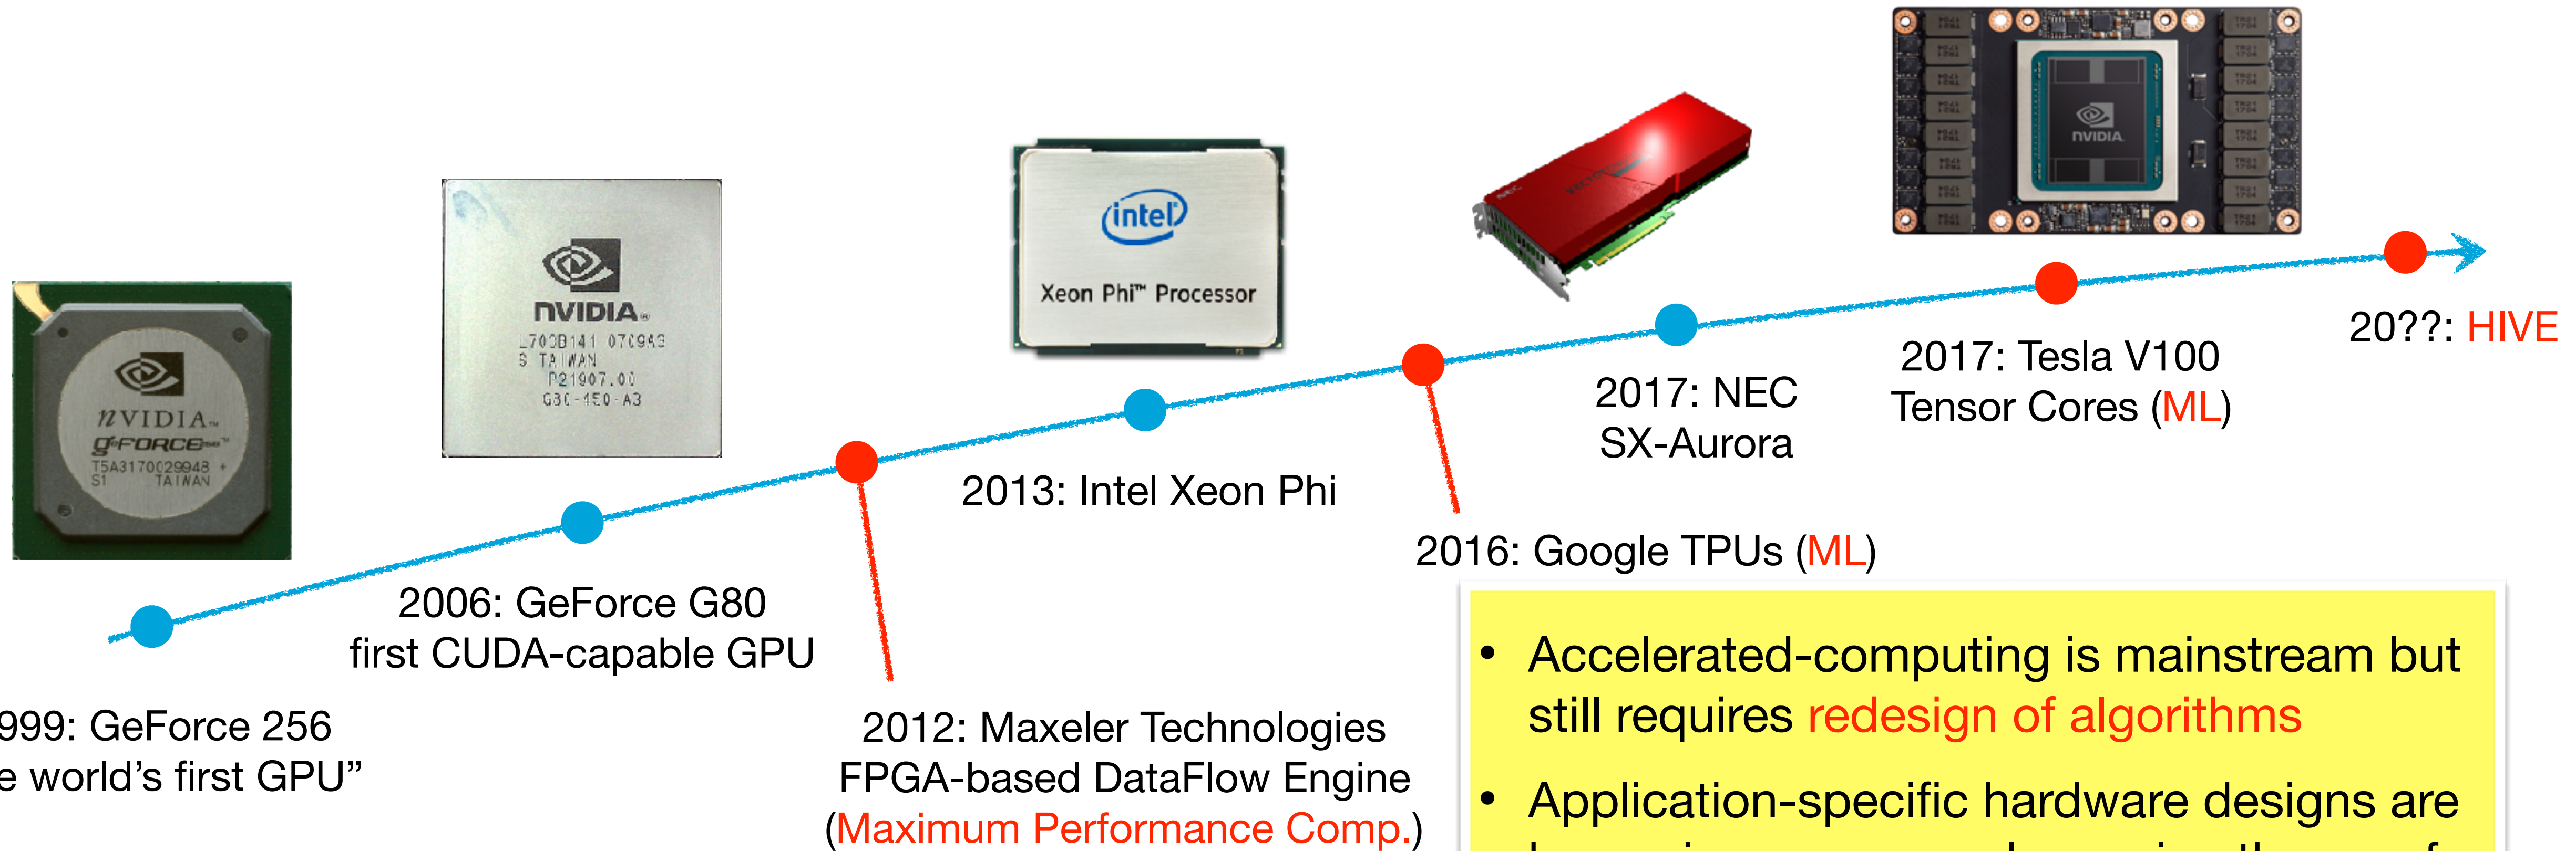


# History of accelerated-computing





# History of accelerated-computing



- Accelerated-computing is mainstream but still requires **redesign of algorithms**
- Application-specific hardware designs are becoming more popular paving the way for **special-purpose functional acceleration**
- "Hardware follows application"



# Future trends in scientific computing

- Heterogeneous (most probably cloud-based) HPC clusters
  - multi-core/socket CPU-based host nodes
  - general-purpose many-core accelerators (GPUs, MICs, VPs)
  - special-purpose functional accelerators (TPUs, FP{G,A}As, **QC?**)

## Maxeler Real Time Risk available on Amazon EC2 F1 Instances

April 20, 2017

Maxeler is proud to announce its [Real Time Risk product](#) on the [Amazon Elastic Compute Cloud \(Amazon EC2\) F1 Instances](#), starting with a Counterparty Risk instance for pre-trade Credit Value Adjustment (CVA) on Amazon Web Service (AWS). With [General Availability of the Amazon](#)

Build and train machine learning models on our new Google Cloud TPUs



*J. Low Power Electron. Appl.* 2016, 6(3), 14; doi:10.3390/jlpea6030014

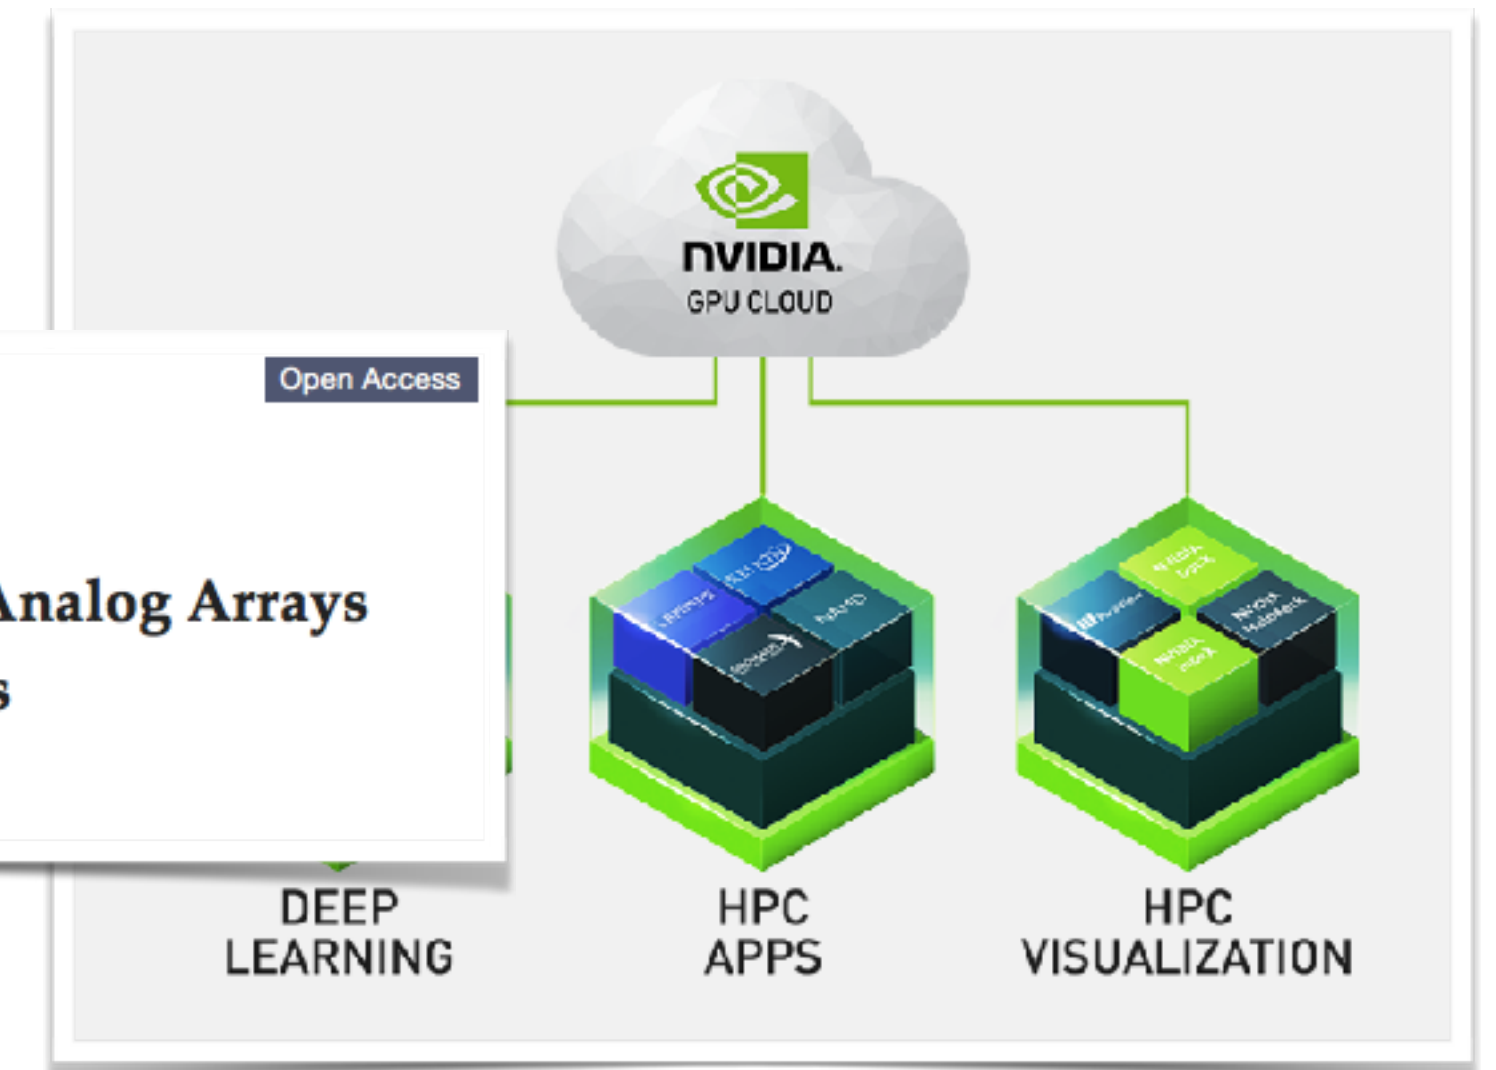
Article

## Remote System Setup Using Large-Scale Field Programmable Analog Arrays (FPAA) to Enabling Wide Accessibility of Configurable Devices

Jennifer Hasler \*, Sahil Shah, Sihwan Kim, Ishan Kumal Lal and Michelle Collins

## IBM makes 20 qubit quantum computing machine available as a cloud service

Posted Nov 10, 2017 by [Ron Miller \(@ron\\_miller\)](#)






# Future trends in scientific computing

- Expected paradigm shifts (“*since the free lunch is over*”)
  - closer co-design of compute hardware and applications
  - more fine-grained application-specific solution algorithms
  - new concepts like in-memory-computing, computing in space
  - rediscovery of ‘old’ approaches
    - *mixed-precision*: Wilkinson ’63, Strzodka et al. ’08, NVIDIA ’16

## **An Analog Accelerator for Linear Algebra**

**Yipeng Huang**, Ning Guo, Mingoo Seok,  
Yannis Tsividis, Simha Sethumadhavan



# **The Finite Element Method**

**From textbook version to HPC implementation**


# FEM in a nutshell

- Poisson eq:  $-\mathbf{u}_{xx}-\mathbf{u}_{yy}=\mathbf{f}$  in  $\Omega$  s.t.  $\mathbf{u}=0$  on  $\Gamma$
- Discretisation by the finite element method:
  - *Basis expansion* of the solution

$$\mathbf{u}_h(x,y)=\mathbf{u}_1\mathbf{B}_1(x,y)+\dots+\mathbf{u}_n\mathbf{B}_n(x,y)$$

Unknown  
coefficients

User-definable  
basis functions





# FEM in a nutshell

- Poisson eq:  $-\mathbf{u}_{xx}-\mathbf{u}_{yy}=\mathbf{f}$  in  $\Omega$  s.t.  $\mathbf{u}=0$  on  $\Gamma$
- Discretisation by the finite element method:
  - *Basis expansion of the solution*

$$\mathbf{u}_h(x,y)=\mathbf{u}_1B_1(x,y)+\dots+\mathbf{u}_nB_n(x,y)$$


- *Assembly of stiffness matrix  $\mathbf{A}_h$*

$$\mathbf{a}_{ij}=\int_{\Omega} B_{i,x}B_{j,x} + B_{i,y}B_{j,y} d\Omega$$

and right-hand side vector  $\mathbf{f}_h$

$$\mathbf{f}_i=\int_{\Omega} B_i \mathbf{f}(x,y) d\Omega$$

Q-accelerated integration (Heinrich 2002)?



# FEM in a nutshell

- Poisson eq:  $-\mathbf{u}_{xx}-\mathbf{u}_{yy}=\mathbf{f}$  in  $\Omega$  s.t.  $\mathbf{u}=0$  on  $\Gamma$
- Discretisation by the finite element method:

- *Basis expansion of the solution*

$$\mathbf{u}_h(x,y)=\mathbf{u}_1B_1(x,y)+\dots+\mathbf{u}_nB_n(x,y)$$


- *Assembly of stiffness matrix  $\mathbf{A}_h$*

$$\mathbf{a}_{ij}=\int_{\Omega} B_{i,x}B_{j,x} + B_{i,y}B_{j,y} d\Omega$$

*and right-hand side vector  $\mathbf{f}_h$*


$$\mathbf{f}_i=\int_{\Omega} B_i \mathbf{f}(x,y) d\Omega$$

- *Solution of linear system  $\mathbf{A}_h\mathbf{u}_h=\mathbf{f}_h$*

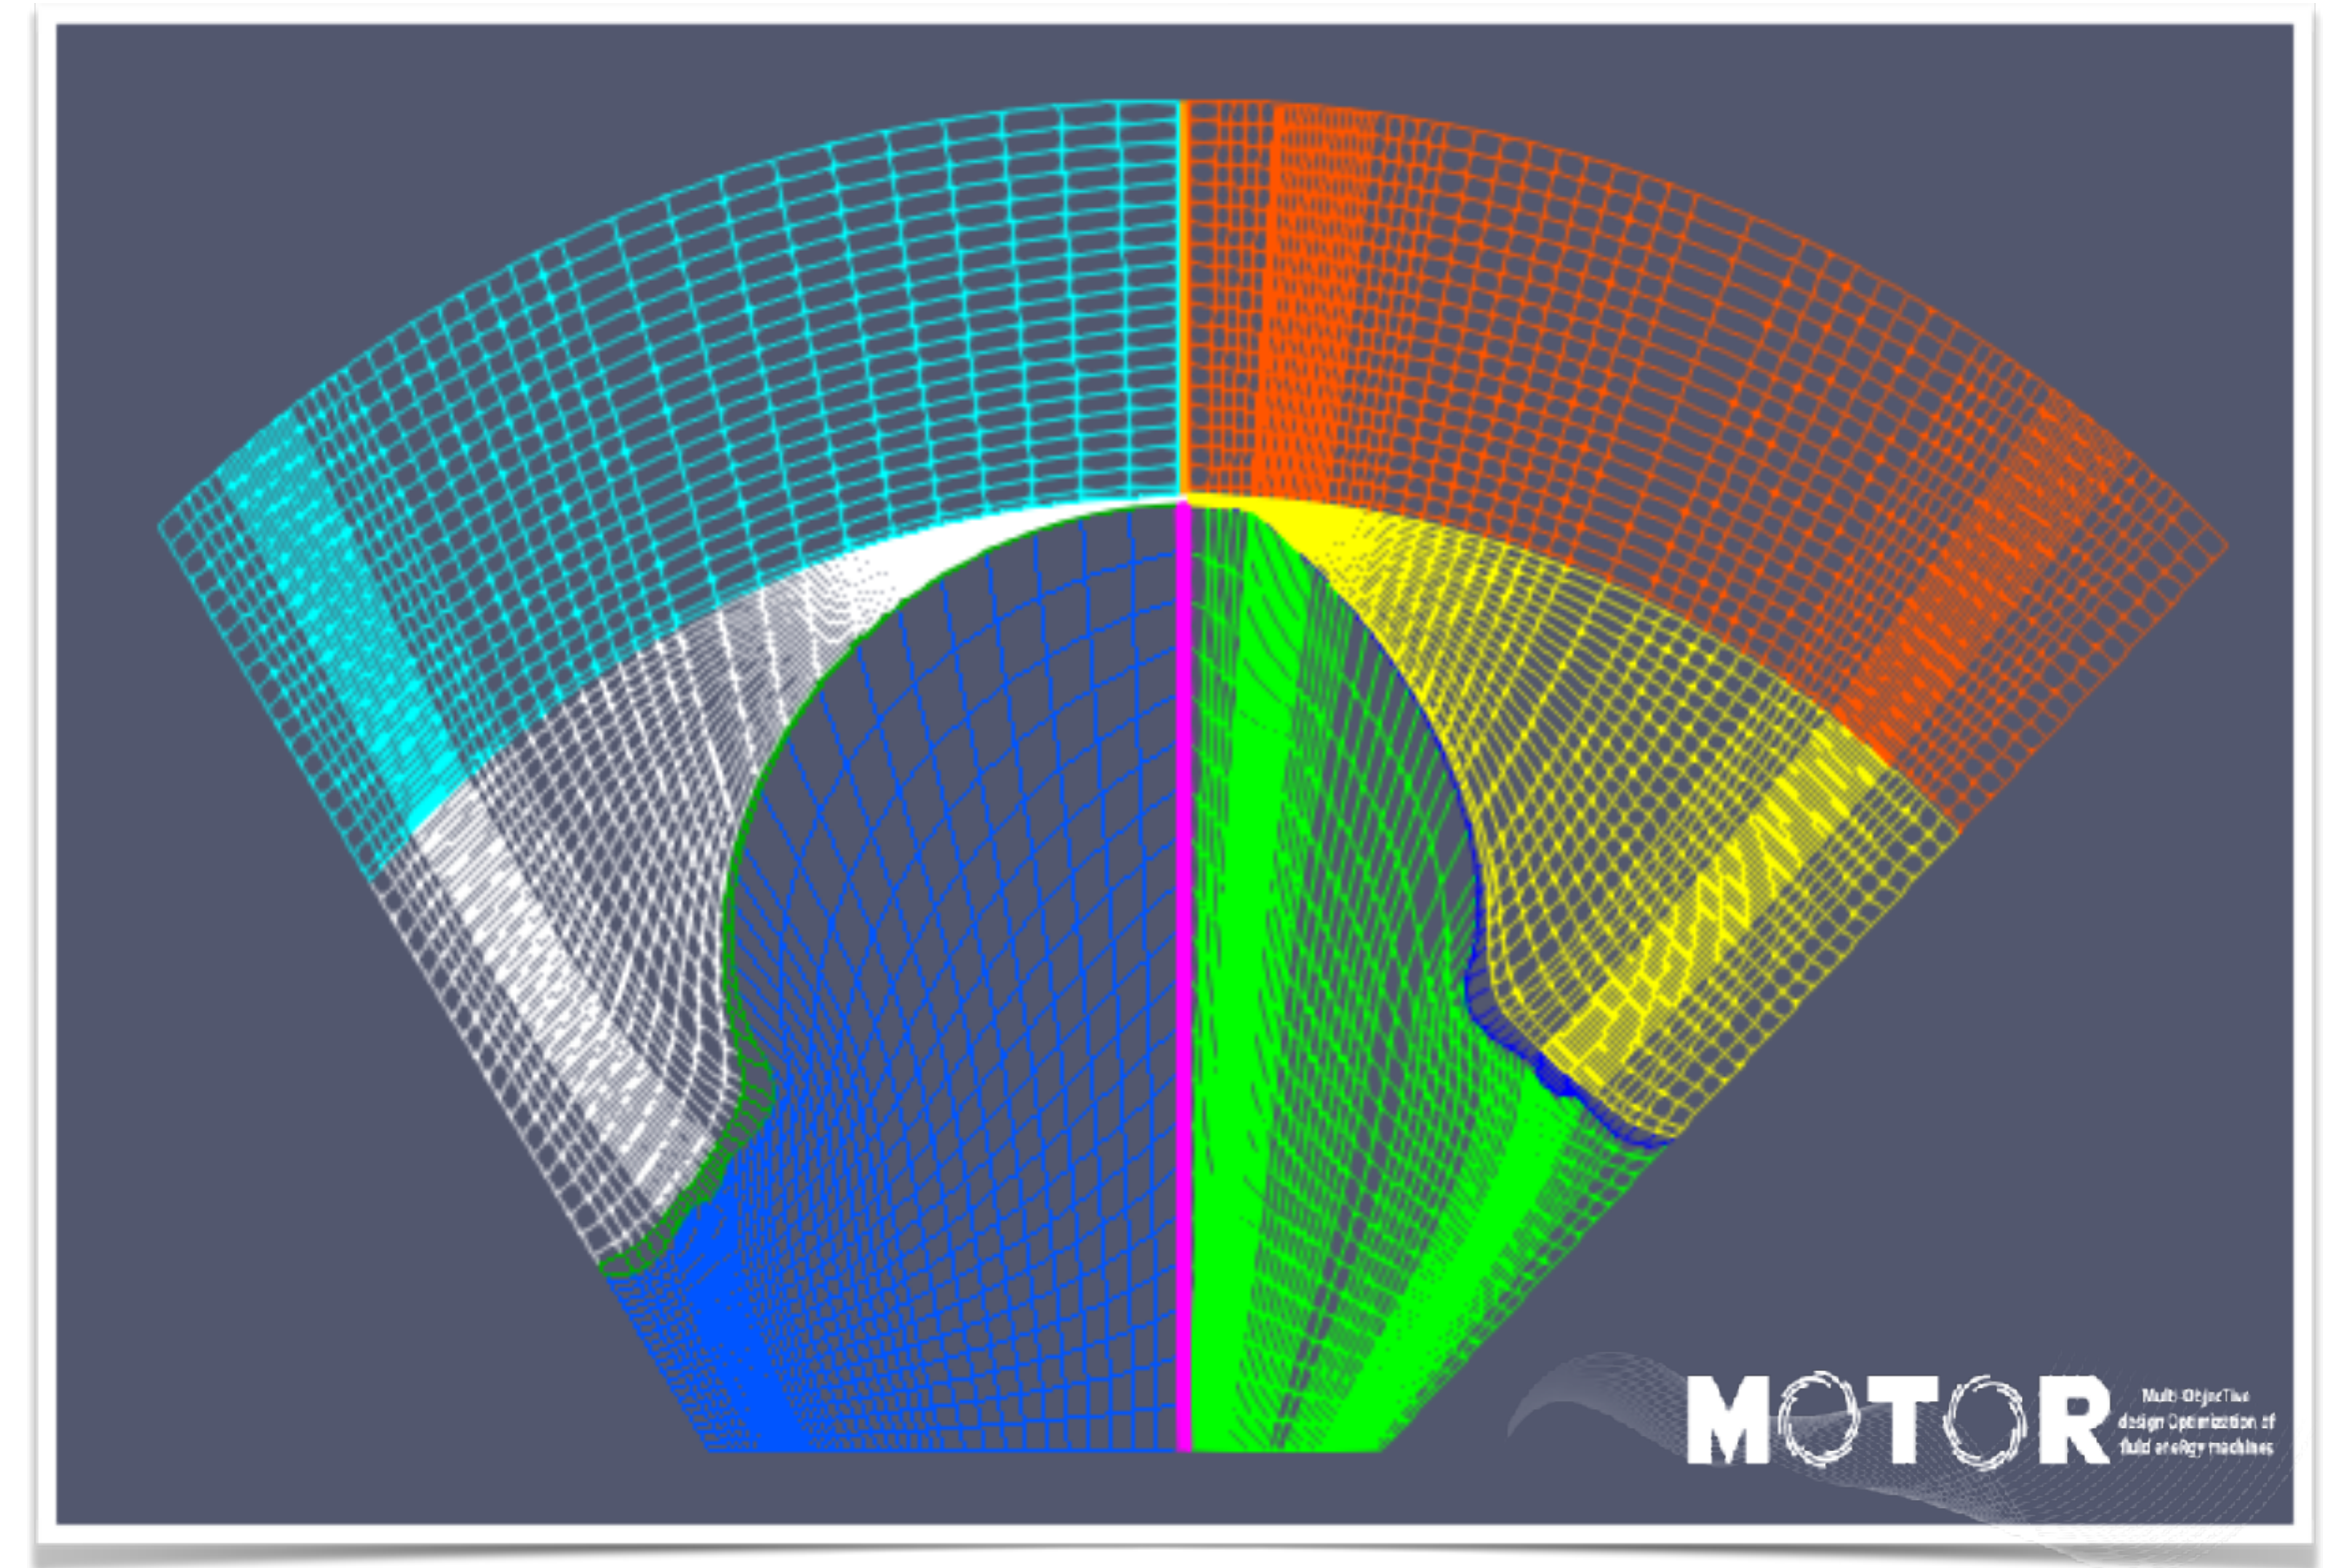




# Beyond textbook FEM



HPC is a challenge

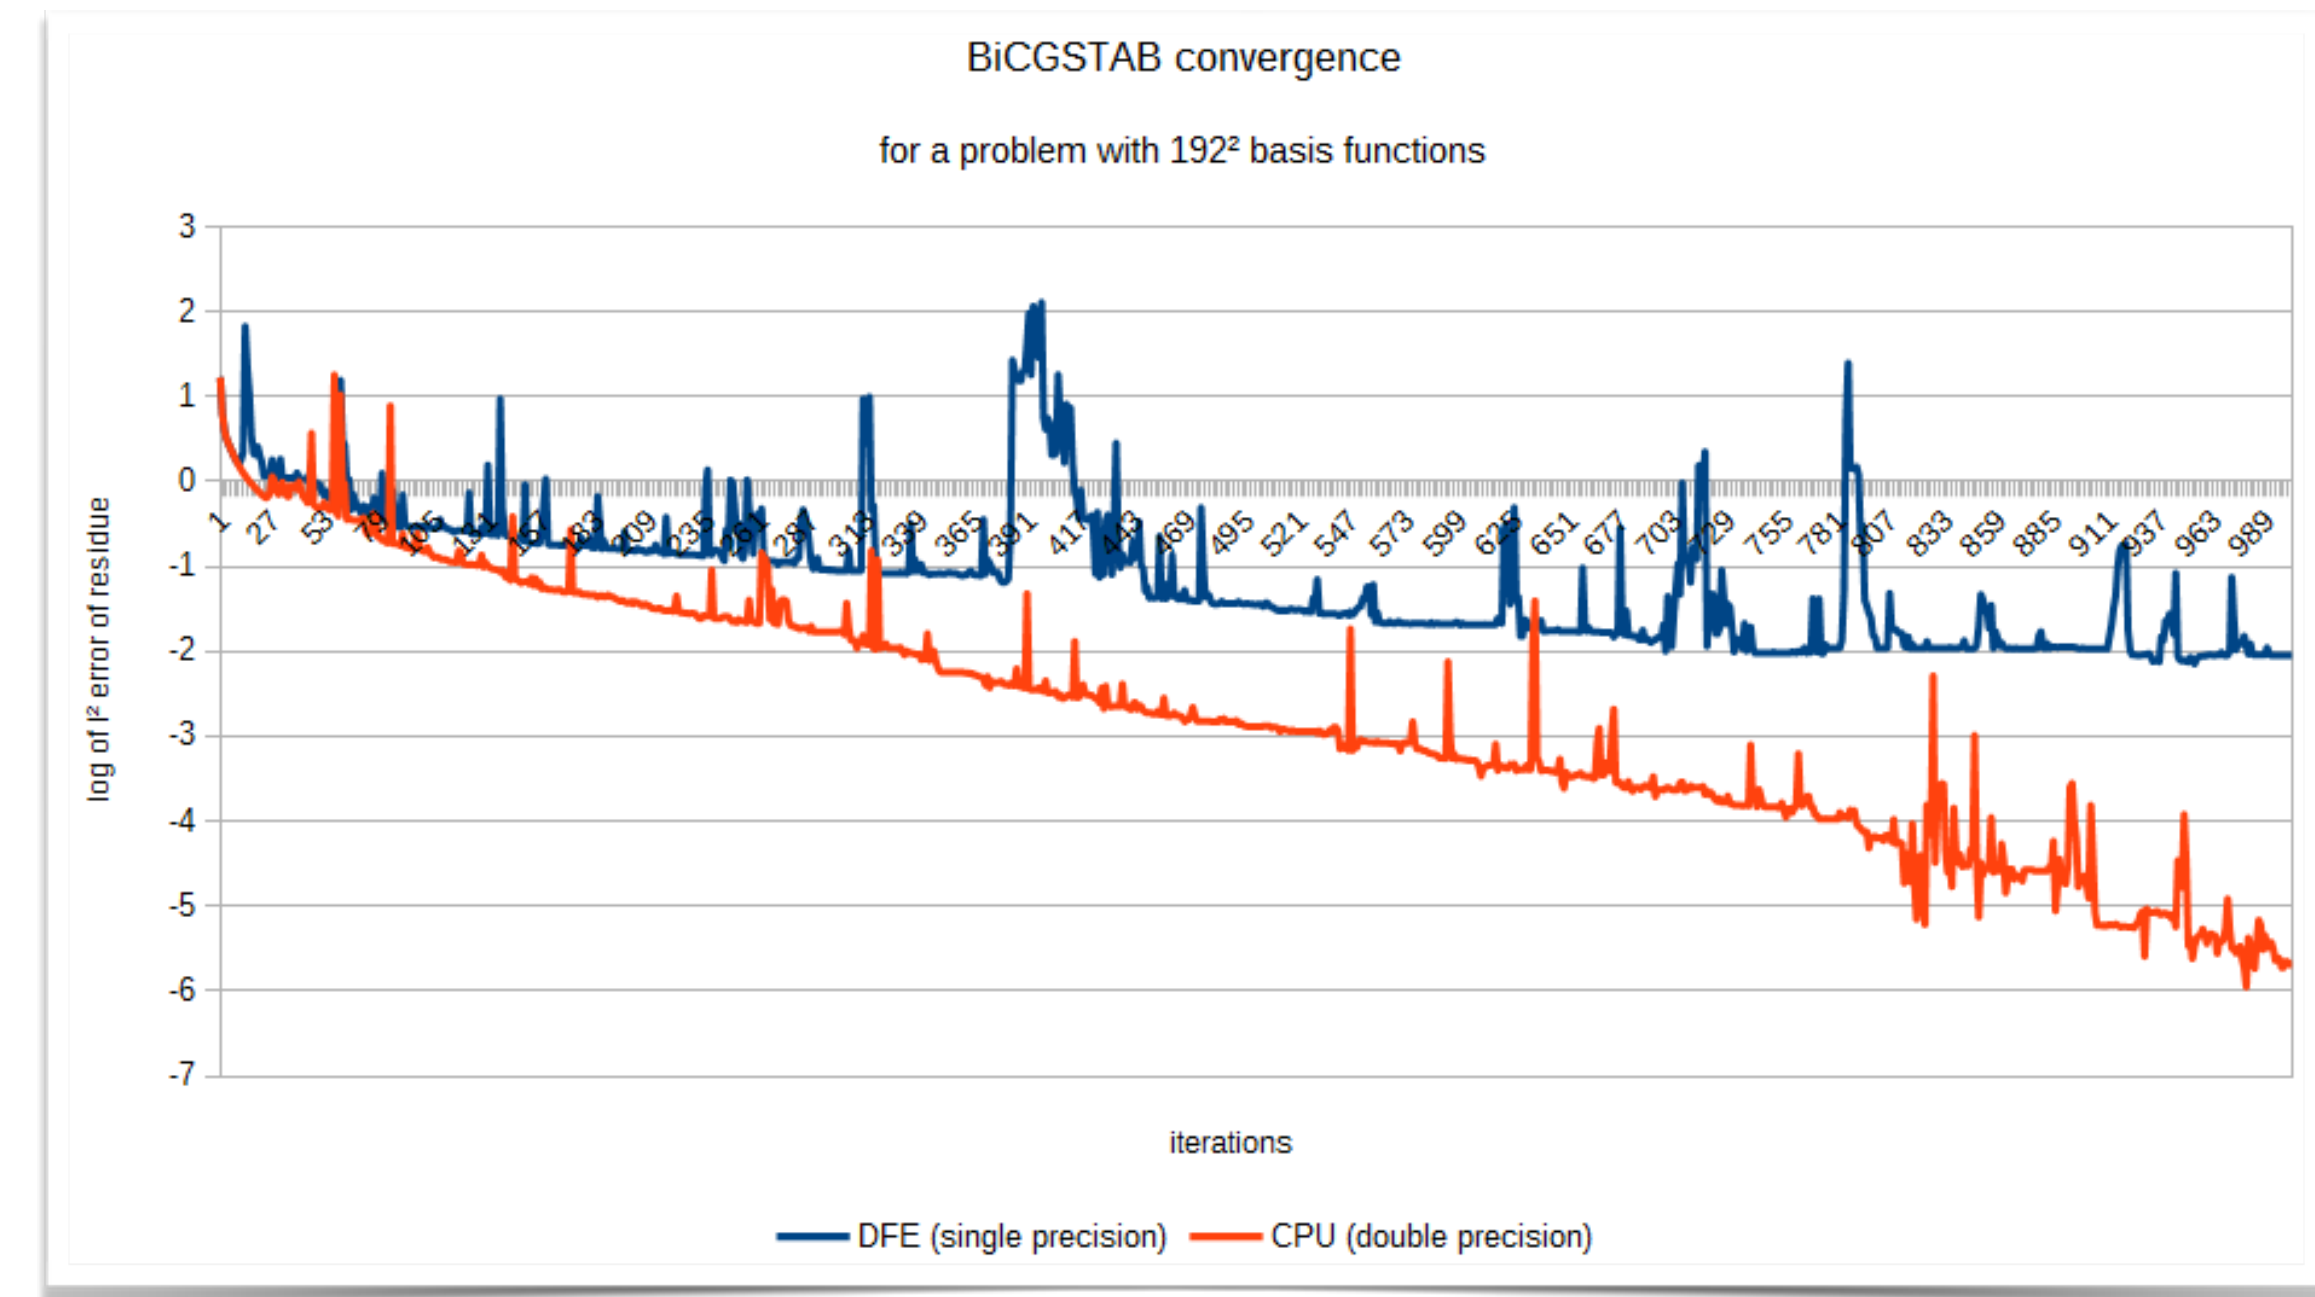
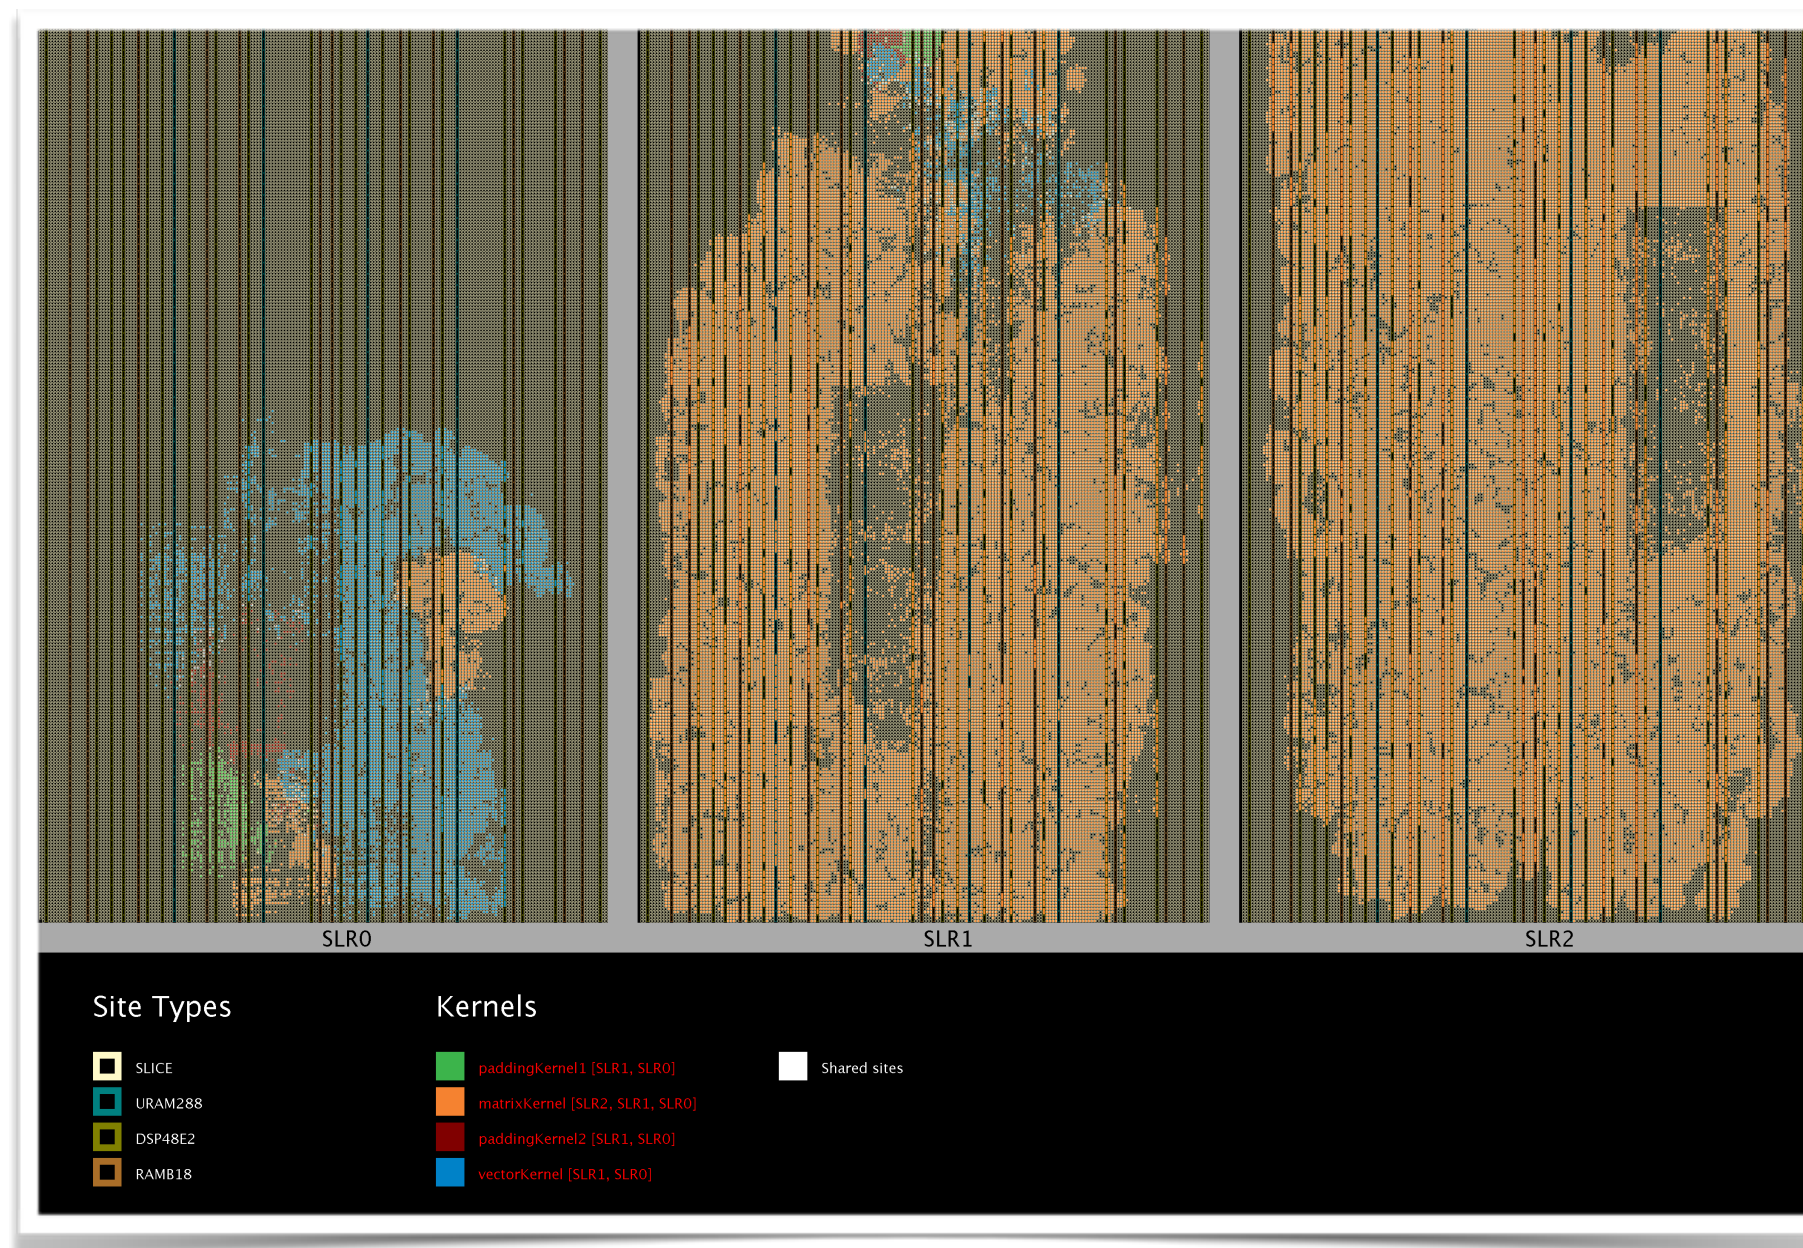


HPC-suitable by design



# An exotic FEM solver

- Matrix-free solver for quadratic B-Spline basis functions; on-the-fly calculation of matrix coefficients
- MAX4/MAX5 implementations



50 s on CPU @2.0GHz (DP)

10 s on MAX4 @20MHz (SP) <- 5x faster

25-30x faster (projected) on MAX5 with  
48bit custom data format for reals


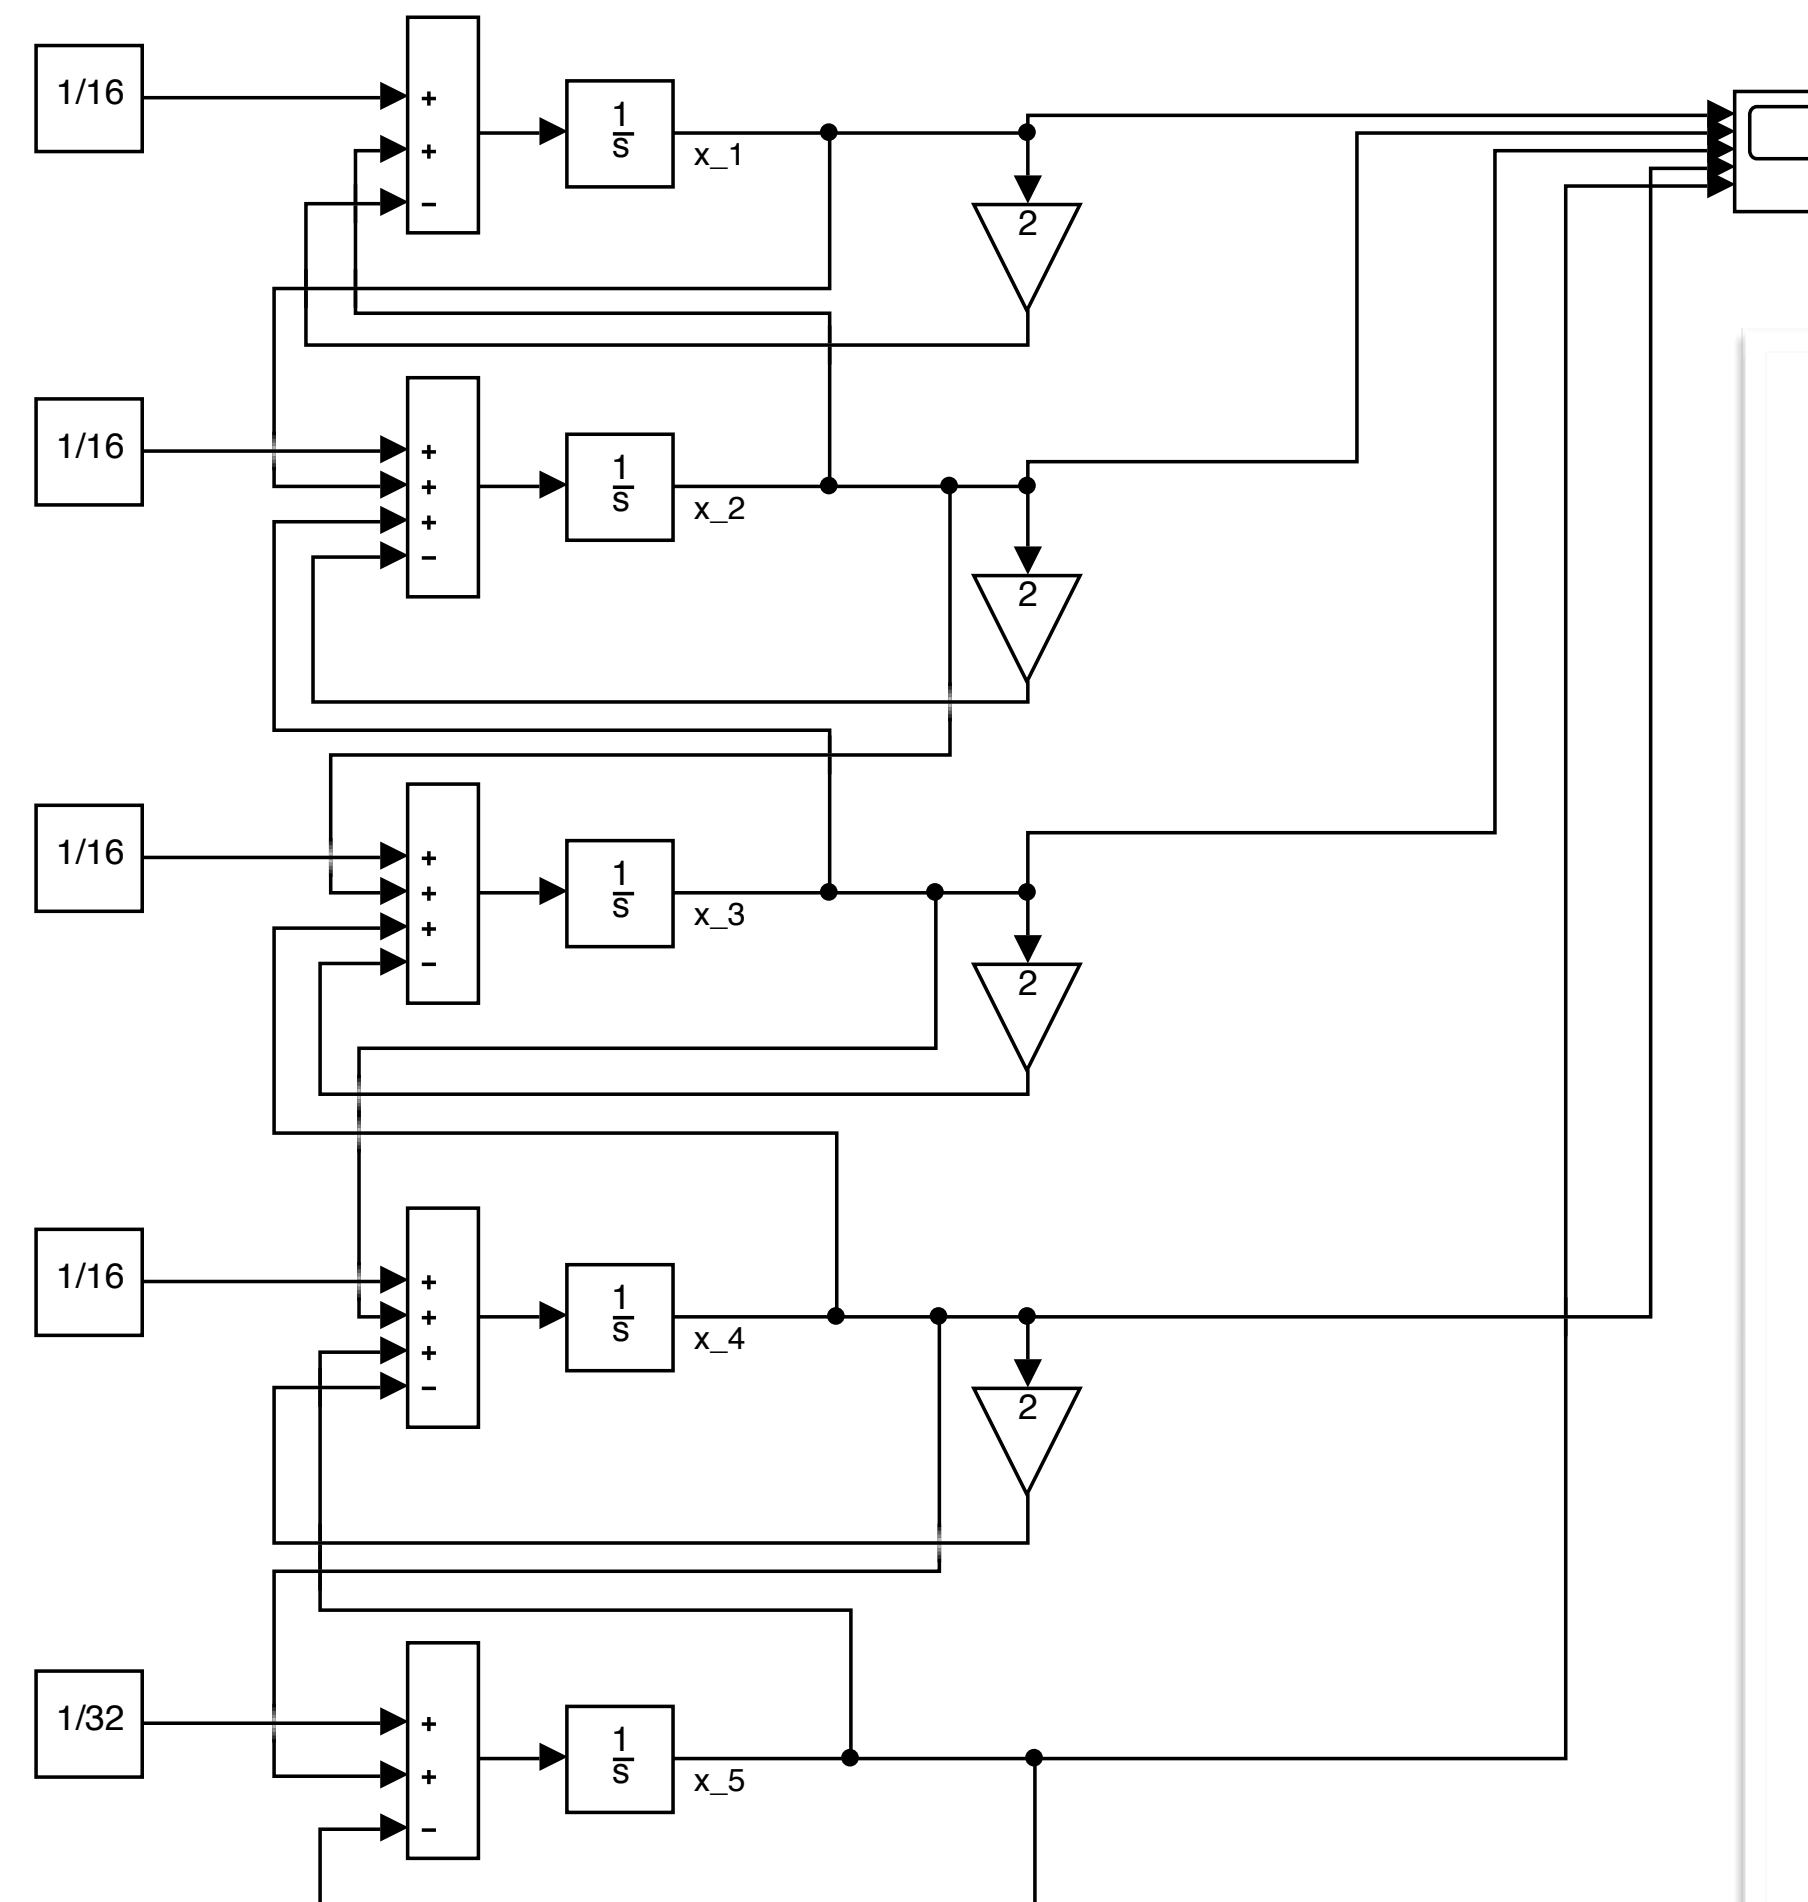


# Another exotic FEM solver

- Solution to the linear system  $\mathbf{A}_h \mathbf{u}_h = \mathbf{f}_h$  can be interpreted as the steady-state limit of the initial value problem

$$\begin{aligned} d\mathbf{u}_h(t)/dt &= \mathbf{f}_h - \mathbf{A}_h \mathbf{u}_h(t), \\ \mathbf{u}_h(0) &= \mathbf{u}_0 \end{aligned}$$

- Acceleration potential using (virtual) analog computing



# Another exotic FEM solver

- Solution to the linear system  $\mathbf{A}_h \mathbf{u}_h = \mathbf{f}_h$  can be interpreted as the steady-state limit of the initial value problem

$$\begin{aligned} d\mathbf{u}_h(t)/dt &= \mathbf{f}_h - \mathbf{A}_h \mathbf{u}_h(t), \\ \mathbf{u}_h(0) &= \mathbf{u}_0 \end{aligned}$$

- Acceleration potential using (virtual) analog computing

- Next steps:
  - $(\mathbf{A}_h, \mathbf{f}_h)/s - ts$  value-time scaling to reduce the dynamic range
  - “Assembly” of  $\mathbf{A}_h, \mathbf{f}_h$  via steady state continuous process
  - “Reduction” of  $\mathbf{u}_h$  into scalar output quantity  $J = \mathbf{u}_h^T \mathbf{M} \mathbf{u}_h$
  - Prototype FPGA-implementation
- Inspiration for Q-FEM solver(?)




# Quantum Computers

The next step in accelerator technologies?

# Simulation-based forecasting

Weather forecast modeling

Timestep 5–10 minutes  
Grid spacing 10–20 km



Variables at the surface:

Temperature  
Humidity  
Pressure  
Moisture fluxes  
Heat fluxes  
Radiation fluxes

Variables in the atmospheric column:

Wind vectors  
Humidity  
Clouds  
Temperature  
Height  
Precipitation  
Aerosols


- PDE-based mathematical model:  
*physical conservation laws*
- Discretisation in space:  
*5 million grid points*  
*x 100 vertical levels*  
*x 10 prognostic variables*

---

**= 5 billion unknowns in space**
- Discretisation in time:  
**864 time steps for 72h-forecast**



# Simulation-based forecasting





- Quantity of interest (QoI):  
*Time evolution of field variables*
- ECMWF at Reading, UK today:  
*simulation of single 10-day forecast in one hour with 10,000 processors*
- Future challenges:  
*increase model resolution by factor 2,000 (1km, 200 levels, 100 variables) and time-step size to improve forecast accuracy requires 20 million processors!*
- Suitable for Q-acceleration:  
*no, due to large size of output data*



# Simulation-based engineering



- QoI: key performance functionals  
*“find  $J_1(U)$ ,  $J_2(U)$ ,... such that  $U$  solves the discretised problem formulation”*
- Example: *cruise L/D (higher is better)*  
*B747-200 (1969): 15.3*  
*B777-200 (1994): 19.3*
- Future challenges:  
*increase model resolution and complexity;*  
*turbulence modelling; CAD integration;*  
*virtual twins for lifetime analysis*
- Suitable for Q-acceleration:  
*maybe, since it fits into QLSA setup but*  
*the problem sizes might be too large*






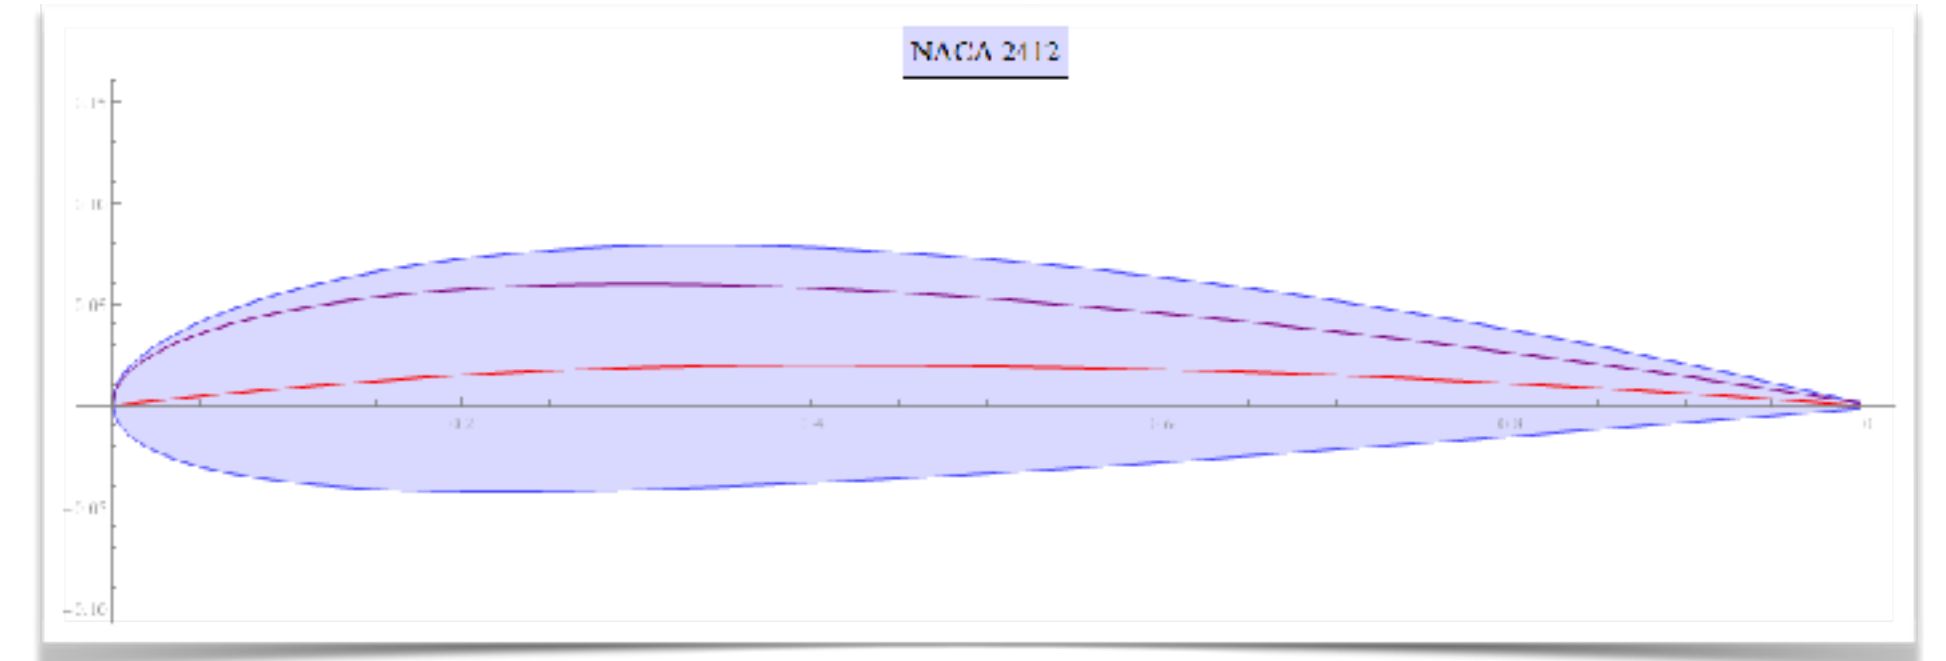
# Simulation-based optimisation

- Input: *design parameters*  $p_1, p_2, \dots$  controlling the shape of the wing
- QoI: *Pareto front of optimal design parameters*  $\mathbf{p}^*$  such that the solutions  $U(\mathbf{p}^*)$  to the discretised PDE problems maximise 'all' key performance functionals  $f_1, f_2, \dots$
- Solution approaches:  $\mathbf{p}^k \rightarrow \mathbf{p}^{k+1}$ 
  - *Gradient-free methods*: evolution-inspired algorithms for generating populations of parameters
  - *Gradient-based methods*: choose parameters 'in the direction' that brings  $U(\mathbf{p}^{k+1})$  closer to an optimal state





# Simulation-based optimisation

- Future challenges:  
*“curse of dimensionality”*; efficient exploration of entire design space
- Suitable for Q-acceleration:  
*maybe, since efficient quantum algorithms for gradient estimation and quantum-based optimisation exist, however, mainly for discrete optimisation problems*





# Uncertainty quantification



- Ensemble forecasting: *Monte-Carlo analysis accounting for uncertainties in initial conditions, mathematical model, data etcetera*
- QoI: *range of possible scenarios (in terms of target functionals) and their likelihood of occurrence*
- Suitable for Q-acceleration: *maybe, since QC requires multiple simulation runs anyway but the problem sizes might be too large*

# Quantum linear solvers

- QLSA: HHL '08, Ambainis '10, Clader et al. '13, Childs et al. '15, Kerenidis '17
- Input:
  - $N \times N$   $s$ -sparse Hermitian matrix  $\mathbf{A}$  with condition number  $\kappa = \|\mathbf{A}\| \|\mathbf{A}\|^{-1}$  and eigenvalues  $1/\kappa \leq |\lambda_i| \leq 1$ , and a right-hand side unit vector  $\mathbf{f}$
- Output:
  - Scalar QoI  $J(\mathbf{u}) = \mathbf{u}^T \mathbf{M} \mathbf{u}$  ( $\mathbf{M}$  is a matrix) such that  $\mathbf{u}$  solves  $\mathbf{A}\mathbf{u} = \mathbf{f}$
- Complexity:
  - Best classical algorithm  $O(\kappa^{1/2} N)$
  - Quantum algorithms  $O(\kappa \log^3 \kappa \log N) \rightarrow O(\kappa^2 \log N)$  ← exponential speed-up

# Q-FEM solver for Poisson's equation

- $\mathbf{A}_h$  is s.p.d.,  $s$ -sparse, well-conditioned ✓
- $\mathbf{f}_h$  is unit vector (after dynamic range scaling) ✓
- Matrix entries in row are accessible in time  $O(s)$  ✓
- Efficient 'preparation' of right-hand side quantum state (✓)
- Can we compute  $\|\mathbf{u}_h\|^2 = \mathbf{u}_h^T \mathbf{u}_h$  where  $\mathbf{u}_h$  solves Poisson's equation?



# Q-FEM solver for Poisson's equation

PHYSICAL REVIEW A **93**, 032324 (2016)

## Quantum algorithms and the finite element method

Ashley Montanaro and Sam Pallister\*

*School of Mathematics, University of Bristol, Bristol BS8 1TW, United Kingdom*

(Received 5 January 2016; published 17 March 2016)

- “when one compares quantum and classical algorithms for the FEM fairly by considering every aspect of the problem – including the complexity of producing an accurate approximation of the desired classical output – an apparent **exponential quantum advantage can sometimes disappear**”
- “there are still two types of problem where quantum algorithms for the FEM could achieve a significant advantage over classical algorithms: those where the **solution has large higher-order derivatives**, and those where the **spatial dimension is large**”

# Q-accelerated scientific computing

- **Challenges**

- Sufficient #qubits for realistic problem sizes
- Q-system hw/sw infrastructure, error correction

- **Open problems**

- Real numbers (IEEE-754, continuous encoding, custom formats)
- Divide-and-conquer/domain decomposition (no-cloning principle!)
- Validation of results (using classical supercomputers?)
- Reproducibility of simulations (but that's a problem in HPC as well)



**Let's get ready for Q-accelerated scientific computing and exchange ideas both ways!**

**Thank you for your attention!**