




Towards ARM for HPC at Jülich Supercomputing Centre

Forschungszentrum Jülich

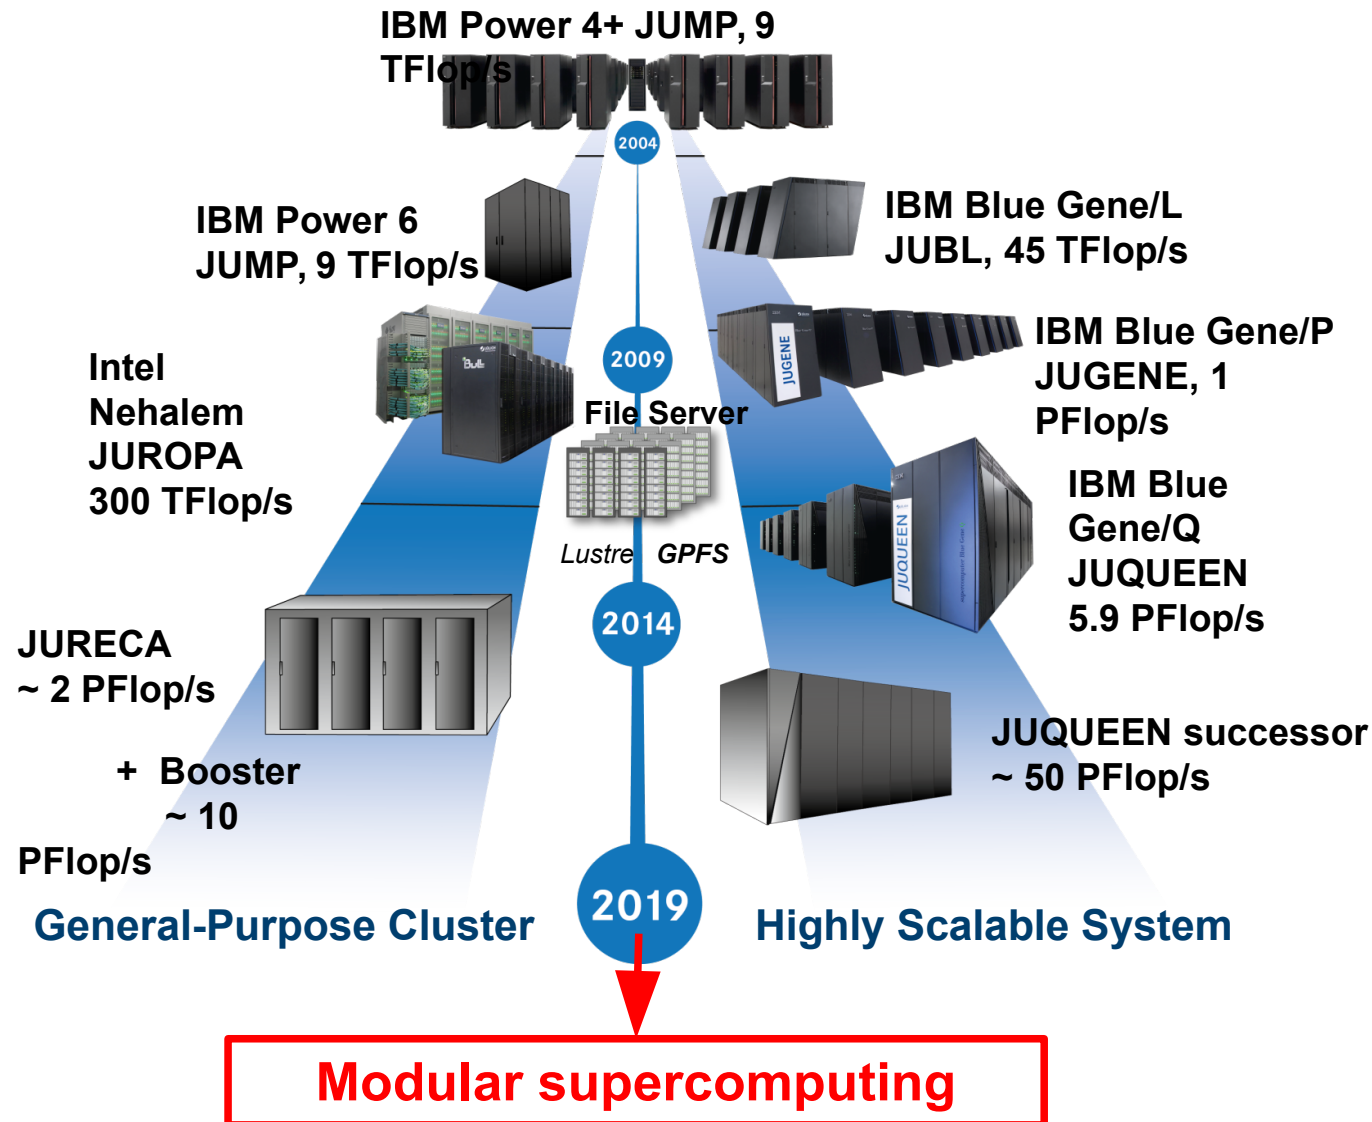


JSC



- One of Europe's largest interdisciplinary research centres; 5,000 employees
- Special expertise in physics, materials science, nanotechnology, neuroscience and medicine, and information technology
- Leader in various European HPC projects, including PRACE

HPC Infrastructure at JSC: Dual track concept



petascale

pre-exascale

exascale

Applications: Structure of Matter

Goal: Explore structure of matter based on ab initio calculations

- Simulation of theory of strong interactions on a lattice: Lattice Quantum Chromodynamics

Numerical challenges


- Perform calculations at
 - Sufficiently large lattice volumes
 - Realistic quark masses
 - Sufficiently fine lattice spacings

Progress limited by computational resources

- Need for exascale performance resources

Application consider here: BQCD

[H. Stüben, Y. Nakamura, 2010]




Applications: Materials Science

Goal: Explore properties of materials using Density Functional Theory (DFT) method

- DFT has become major tool for exploring properties of materials, e.g.
 - Phase-change materials
 - Dilute magnetic semiconductors

Challenge: Need to simulate many atoms

- Complexity of method typically scales $O(N_{\text{atom}}^3)$
 - In certain cases linear scaling is possible
- Aim for $O(10^5 \dots 10^6)$ atoms → exascale challenge



Application considered here: KKRnano

[R. Zeller et al., 2008]

Applications: Brain Research

Goal: Improve the understanding of the human brain by means of models

- Simulations at different organisational levels


Challenge of whole brain simulations

- Extremely complex network
 - Huge network: $O(10^{11})$ neurons
 - High connectivity: Neuron connected to $O(10^4)$ neurons
- Memory capacity limitations

Application considered here: NEST

- Simulation of spiking neuronal network
- Focus on large networks, use of simple point neurons

[M.O. Gewaltig, M. Diesmann, 2007]



[Potjans, Diesmann, 2012]

Performance Characteristics

Performance critical regions

- BQCD + KKRnano: Linear solver
- NEST performance is less dominated by single kernel

Control flow

- Simple and predictable in case of BQCD + KKRnano
- Less regular control flow in case of NEST

Dominating arithmetic operations

- BQCD + KKRnano: Complex FMA

Arithmetic intensity


KKRnano	~4 (typical)
BQCD	0.9
NEST	0

Performance Characteristics (cont.)

Consider hypothetical ARM-v8 server processor


- 8 cores, 128-bit SIMD, 2.5 GHz
- 4 DDR channels, 2.133 MT/s

Roofline model



Performance Observations: STREAM

Memory bandwidth on X-Gene 2



.L26:

```

add    w0, w0, 1
cmp     w0, w19
ld1     {v0.2d}, [x6], 16
st1     {v0.2d}, [x5], 16
bcc     .L26
cmp     w22, w23
beq     .L25
  
```


Observations

- Internal bandwidth does not scale with number of threads
- Maximum memory bandwidth below 10 Byte/cycle

Performance Observations: KKRnano

KKRnano kernel dominated by sparse matrix multiplication

- Sub-structure with matrices of size b
- Micro-benchmark based on OpenBlas



Throughput $B_{fp} \approx 1 \text{ Flop/cycle}$

[P. Baumeister, 2016]

- Compiler fails to map to SIMD instructions
- Single thread performance full application: Haswell vs. X-Gene 1 about 10x

Performance Observations: BQCD

Kernel performance dominated by sparse matrix-vector multiplication

- Stencil-type pattern

Single processor scaling


- Reasonable scalability observed on 32 cores

Performance per core

- About 10% of peak in sub-kernel

Performance limiters


- Lacking exploitation of SIMD instructions
- LLC cache utilization is difficult due to limited size



LQCD: Exploiting SIMD

SIMD support in upcoming community code „grid“

- Strategy: Domain decomposition to exploit SIMD parallelism



[P. Boyle et al., 2015]

- Scales to very large SIMD width → candidate for SVE
- C++ classes hide details to higher levels

Port to NEON completed

- Performance analysis pending
- Part of QPACE 4 project


[N. Meyer, 2016]



Performance Observations: NEST

Comparative evaluation of IPC for spike delivery

- Performance critical region



Observations

- Smaller IPC corresponds to less pipelines
- Oversubscribing core slightly improves IPC

Conclusions and Outlook

Early results for a set of HPC applications

- All scale well on Blue Gene/Q

Need for software efficiently exploiting hardware

- LQCD community code exploiting SIMD intrinsics
- SIMD support in numerical libraries

KKRnano and BQCD could benefit from SVE

- Can be mapped to maximum width

Memory subsystem performance crucial

- Potential for better balance of throughput of arithmetic operations vs. memory bandwidth
- Suitable cache sizes