The High-Q Club

Applications scaling to the full JUQUEEN system with 458 752 cores and over 1.8M threads

Dirk Brömmel, Wolfgang Frings, and Brian J. N. Wylie | Jülich Supercomputing Centre | Forschungszentrum Jülich GmbH
Contact: {d.broemmel, w.frings, b.wylie}@fz-juelich.de | **DOI:**10.14529/jsfi180104 | **Website:** www.fz-juelich.de/ias/jsc/high-q-club

- **Showcase of 32 application codes demonstrating extreme scalability from diverse fields and global developer community**
- **Based on effective use of full JUQUEEN IBM Blue Gene/Q system (installed spring 2012, decommissioned spring 2018)**
- **28 racks, 28 672 processors, 458 752 cores, 1 835 008 threads (16 GiB memory/processor, 5D torus network, GPFS filesystem)**
- **Membership supported by 3 Extreme Scaling Workshops, JSC Simulation Laboratories and Cross-Sectional Teams**
- **Application benefits extended beyond BG/Q to other HPC leadership computer systems**
- **Potential insight for expected future exa-scale applications**
- **Reference basis for successor activity**

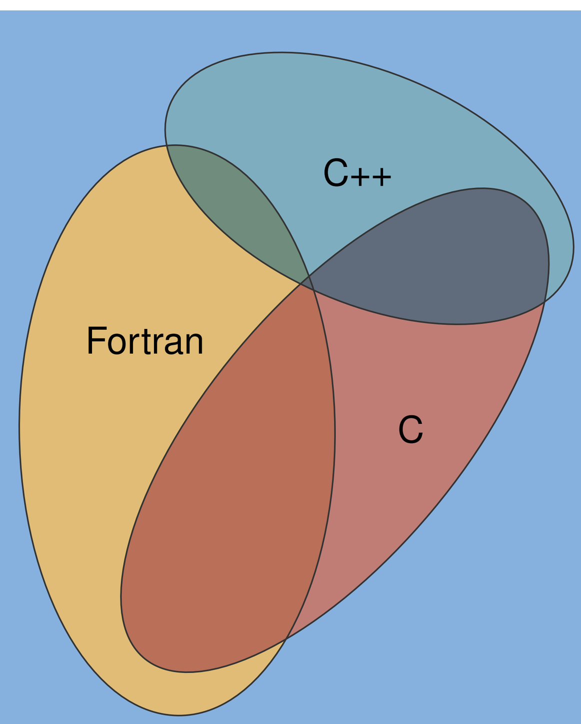

Programming languages used

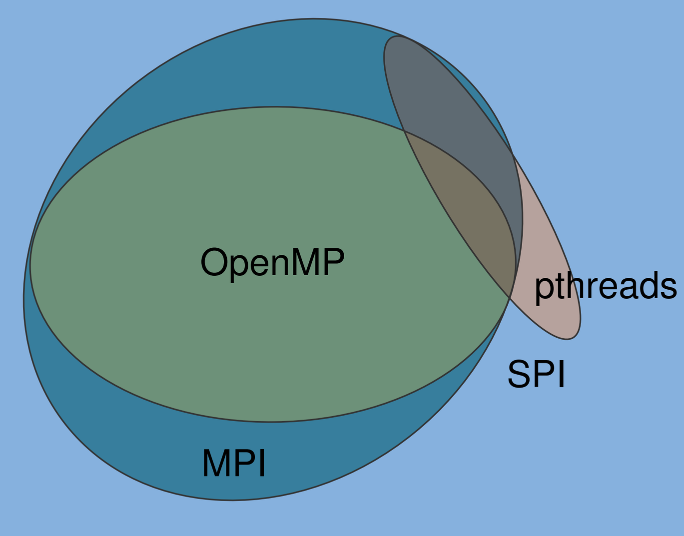

Programming models used

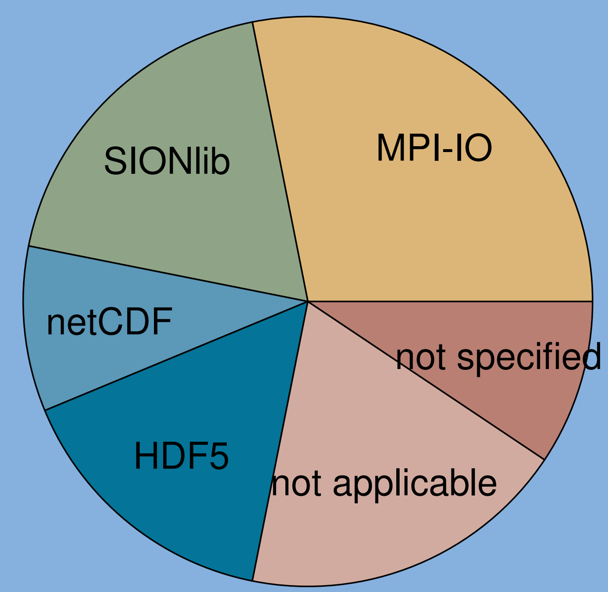

I/O methods used

Lessons

- Compute node memory and file I/O are common constraints
- Need effective parallel I/O solutions for metadata and bandwidth, such as SIONlib
- Standard programming languages and MPI combined with multi-threading suffice
- Optimised libraries developed for efficient mesh partitioning and non-blocking 3D FFTs
- Tuning often required as directed by tools such as Scalasca

Application characteristics

- Fortran, C & C++ languages all used, and often combined
- MPI usage ubiquitous, mostly with OpenMP and/or pthreads mixed-mode combination more memory-efficient
- Additionally used hardware threads benefitted most codes
- SIONlib, MPI-IO & HDF5 equally prevalent for file I/O (often data synthesised or writing disabled as workarounds)

Scaling performance

- Weak scaling is easiest:
  12 codes with scaling efficiency over 80%
- Strong scaling is much more challenging:
  10 codes maintained normalised scaling efficiency over 80% (compute node memory requirements can limit baseline to quarter of machine)
- Cache benefits and machine topology can be important

Strong scaling

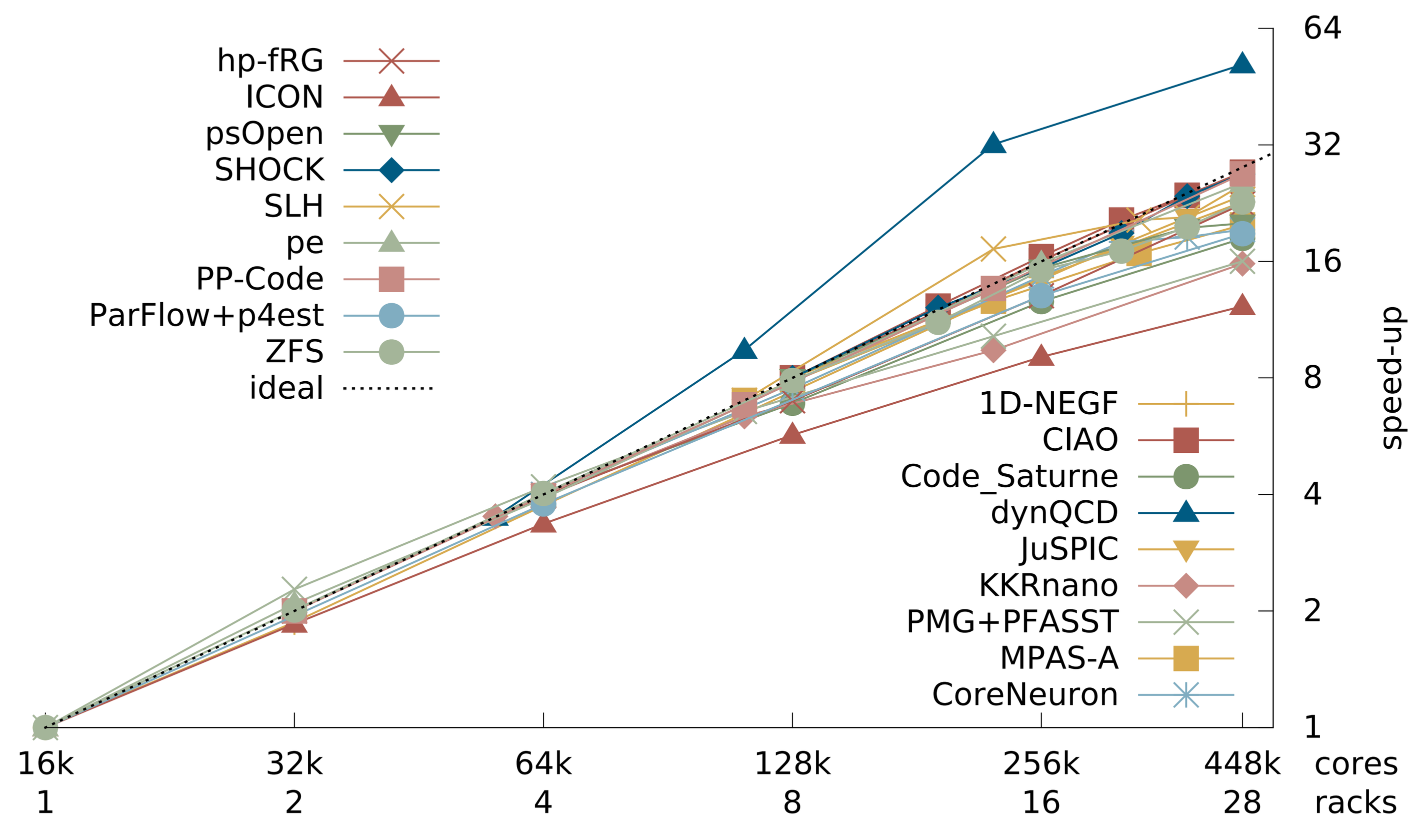

Hardware threads

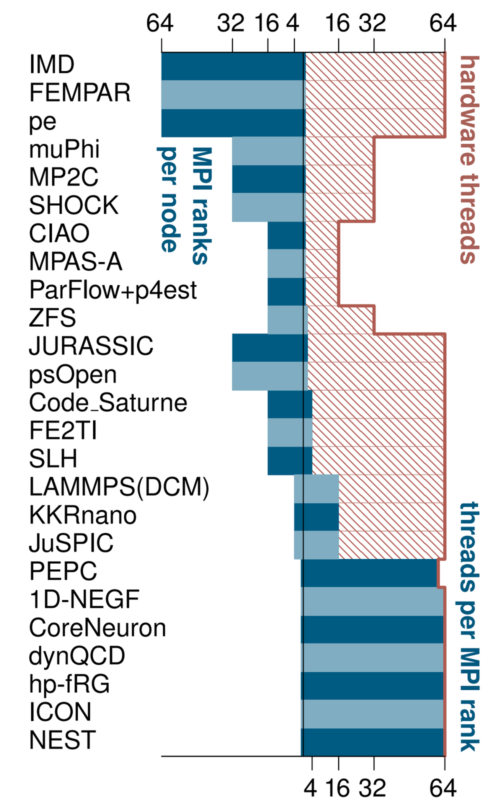

Weak scaling

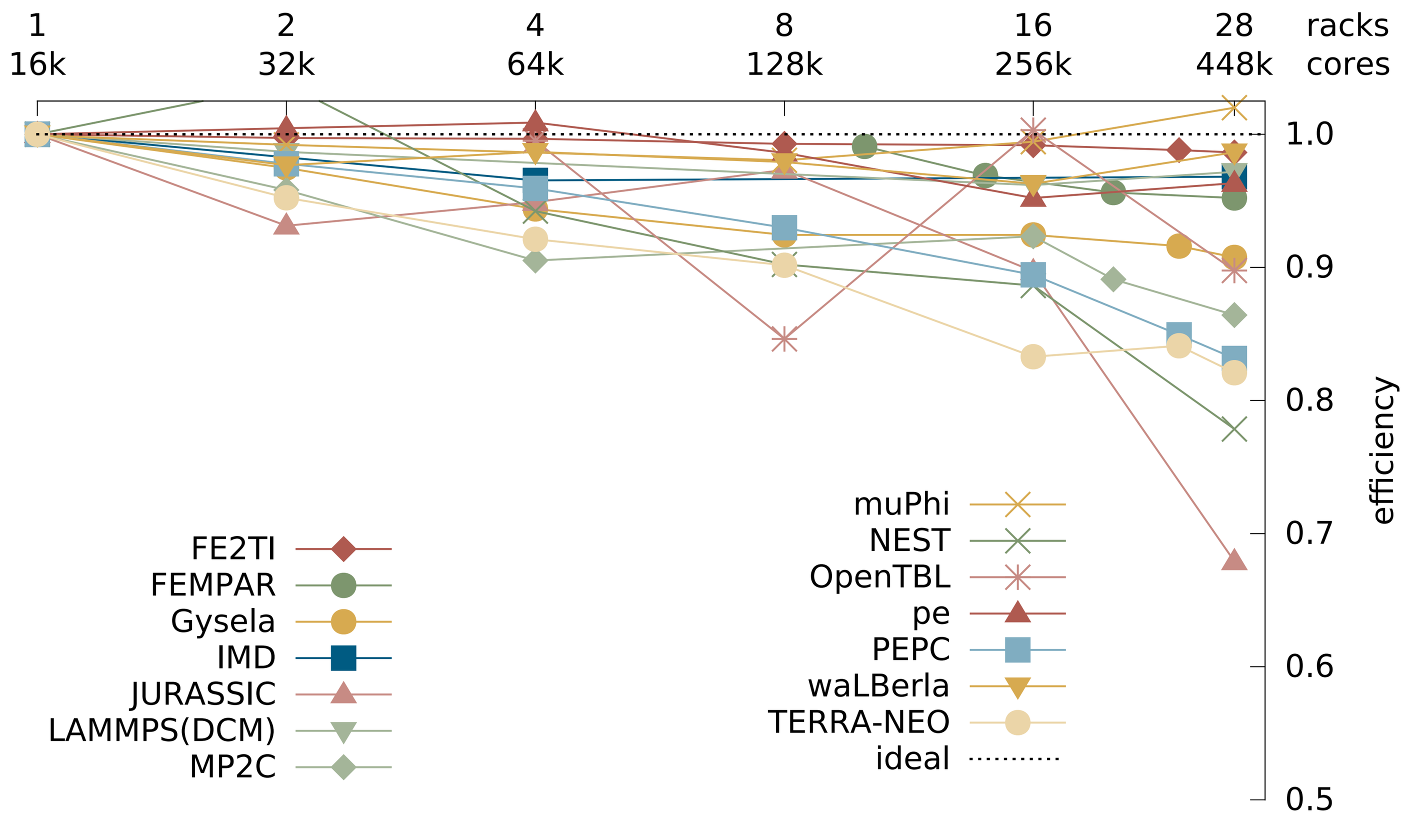

High-Q Club codes

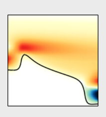
1D-NEGF (JSC SimLab Quantum Materials) 1D Non-Equilibrium Green's Functions framework for transport phenomena


CIAO (RWTH ITV) multiphysics, multiscale NS solver for turbulent reacting flows in complex geometries

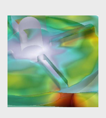
Code_Saturne (EDF & STFC Daresbury Laboratory) multiphysics simulation of the Navier-Stokes equations

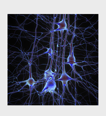
CoreNeuron (EPFL Blue Brain Project) simulation of electrical activity of neuronal networks including morphologically detailed neurons

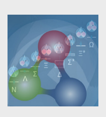
dynQCD (JSC SimLab Nuclear and Particle Physics & Bergische Universität Wuppertal) lattice quantum chromodynamics with dynamical fermions


FE2TI (Universität Köln & TUB Freiberg) scale-bridging incorporating micro-mechanics in macroscopic simulations of multi-phase steels

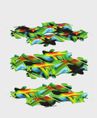
FEMPAR (UPC-CIMNE) massively-parallel finite-element simulation of multi-physics problems governed by PDEs

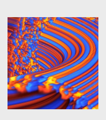
Gysela (CEA-IRFM Cadarache) gyrokinetic semi-Lagrangian code for plasma turbulence simulations

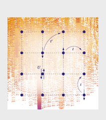
hp-fRG (JSC) hierarchically parallelised functional renormalisation group calculations


ICON (DKRZ & JSC SimLab Climate Science) icosahedral non-hydrostatic atmospheric model


IMD (Ruhr-Universität Bochum & JSC SimLab Molecular Systems) classical molecular dynamics simulations

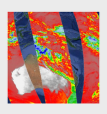
JURASSIC (JSC SimLab Climate Science) solver for infrared radiative transfer in the Earth's atmosphere

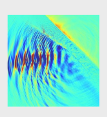
JuSPIC (JSC SimLab Plasma Physics) fully relativistic particle-in-cell code, plasma physics simulations, laser-plasma interaction

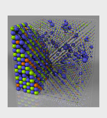
KKRnano (FZJ IAS) Korringa-Kohn-Rostoker Green function code for quantum description of nano-materials in all-electron density-functional calculations

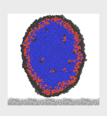
LAMMPS(DCM) (RWTH AICES) molecular dynamics simulation with dynamic cutoff method for arbitrarily large interfacial systems

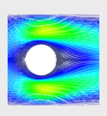
MP2C (JSC SimLab Molecular Systems) massively-parallel multi-particle collision dynamics for soft matter physics and mesoscopic hydrodynamics

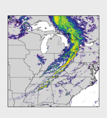
MPAS-A (KIT & NCAR) multi-scale non-hydrostatic atmospheric model for global, convection-resolving climate simulations

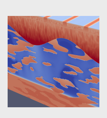
$\mu\varphi$ (muPhi) (Universität Heidelberg) modelling and simulation of water flow and solute transport in porous media, algebraic multi-grid solver

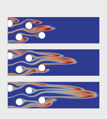
Musubi (Universität Siegen) multi-component Lattice Boltzmann solver for flow simulations

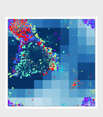
NEST (FZJ INM-6 & IAS-6) large-scale simulations of biological neuronal networks


OpenTBL (Universidad Politécnica de Madrid) direct numerical simulation of turbulent flows

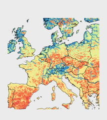
ParFlow+p4est (FZJ IBG-3, Colorado School of Mines, LLNL & Universität Bonn) high resolution parallel simulation of variably saturated flow

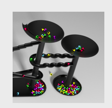
pe (Universität Erlangen-Nürnberg) physics engine framework for simulations of rigid bodies with arbitrary shapes

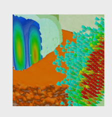
PEPC (JSC SimLab Plasma Physics) tree code for *N*-body simulations, beam-plasma interaction, vortex dynamics, gravitational interaction, MD simulations

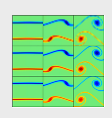
PMG+PFASST (LBNL, Universität Wuppertal, USI & JSC) space-time parallel solver for systems of ODEs with linear stiff terms, e.g. from lines discretisations of PDEs

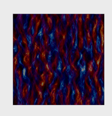
PP-Code (University of Copenhagen) simulations of relativistic and non-relativistic astrophysical plasmas

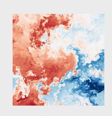
psOpen (RWTH ITV, JARA & CNRS CORIA) highly-resolved direct numerical simulation of fine-scale turbulence

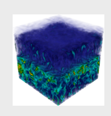
Seven-League Hydro (Heidelberg Institute for Theoretical Studies) all Mach number fluid dynamics in astrophysics

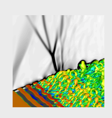
SHOCK (RWTH Shock Wave Laboratory) structured high-order finite-difference kernel for compressible flows

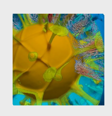
TERRA-NEO (Universität Erlangen-Nürnberg, LMU & TUM) modeling and simulation of earth mantle dynamics

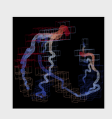
waLBerla (Universität Erlangen-Nürnberg) Lattice-Boltzmann method for the simulation of fluid scenarios

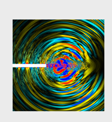
ZFS (RWTH AIA & JARA SimLab Fluids & Solids) Computational fluid dynamics, computational aeroacoustics, conjugate heat transfer, particulate flows

Programming language and model | I/O