Z02: Computing, Co-Design, Performance Optimisation

E. Landinez, JSC

U Sinha, JSC

B. Kostrzewa, HISKP, Uni Bonn

E. Suarez, JSC, Uni Bonn

Outline

- State of the Art: MSA
- Connections
- Co-Design
- Work in Progress

Domain Scientists

State of the Art

- Use HPC to solve complex scientific problems
- Combine numerical methods with artificial intelligence (ML/DL)

Domain Scientists

State of the Art

- Use HPC to solve complex scientific problems
- Combine numerical methods with artificial intelligence (ML/DL)

Challenges

- Hardware heterogeneity
- Thread parallelism & Scalability
- Memory management
- Communication
- Performance portability

Domain Scientists

State of the Art

- Use HPC to solve complex scientific problems
- Combine numerical methods with artificial intelligence (ML/DL)

Challenges

- Hardware heterogeneity
- Thread parallelism & Scalability
- Memory management
- Communication
- Performance portability

Z02 supports areas A, B and C:

→ Guidance, documentation, and training for application porting, scaling, and perf. engineering

→ Research Software Engineerig (RSE) and CI/CD infrastructure

→ HPC topics at NuMeriQS lecture weeks and retreats

→ Benchmarking on diverse HPC devices and architectures

⟨NuMeriQS⟩

Increasing hardware heterogeneity

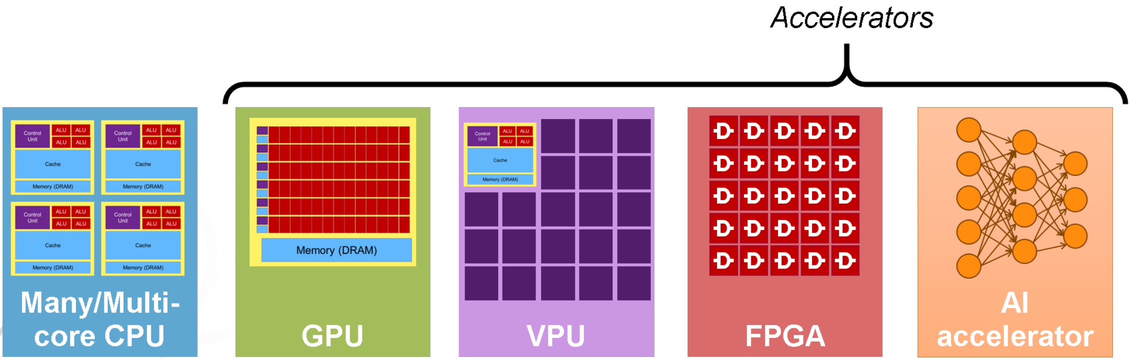

Many/Multi-core CPU	GPU	VPU	FPGA	AI accelerator
10's strong cores	1000's functional units	100's vector arithmetic units	1.000.000's programmable gates	custom ASIC implementations (e.g. TPUs)

Modular Supercomputing Architecture

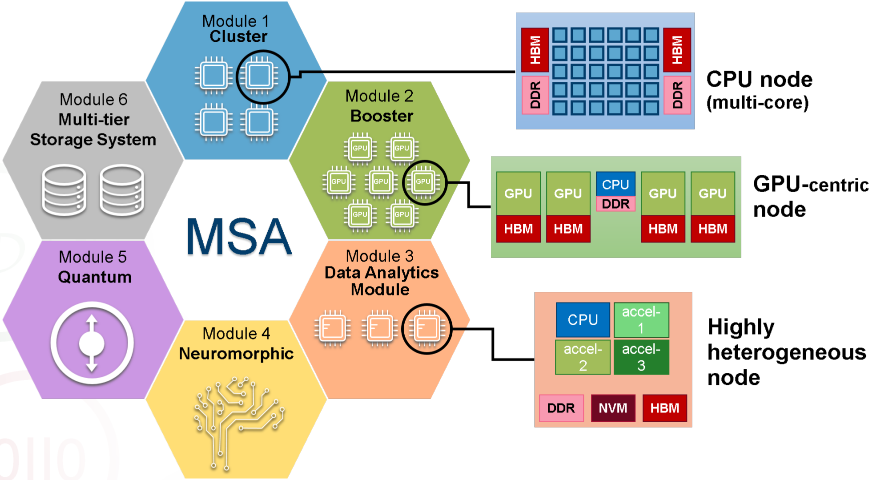

⟨NUMERIQS⟩

Modular Supercomputing Architecture

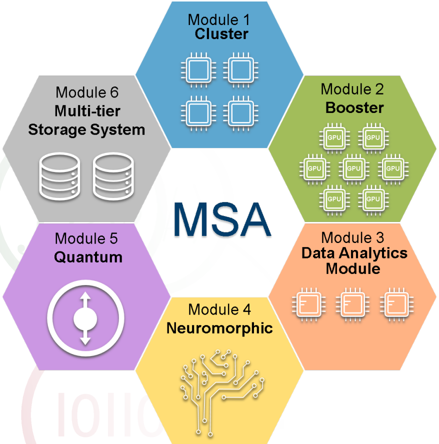

- Heterogeneity at system level
  - ⇒ Scale modules independently
  - ⇒ Gradual integration of disruptive technologies
- Dynamic allocation
  - ⇒ Efficient ressource sharing
- Support application diversity
  - ⇒ Adapt system to user portfolio
  - ⇒ **Co-Design**
- Hardware architecture of:
  - ⇒ **Marvin:** HPC system at UniBonn
  - ⇒ **JUPITER:** 1^{st} Exascale system in Europe

Connections

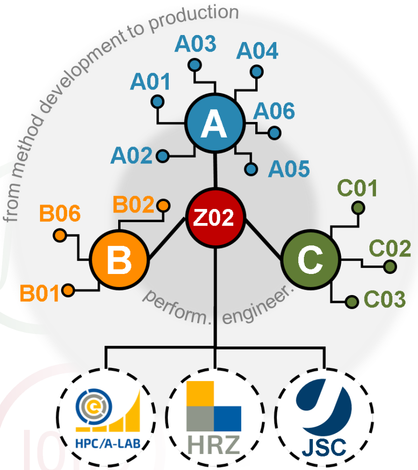

⟨NuMeriQS⟩

Connections , and Activity Areas

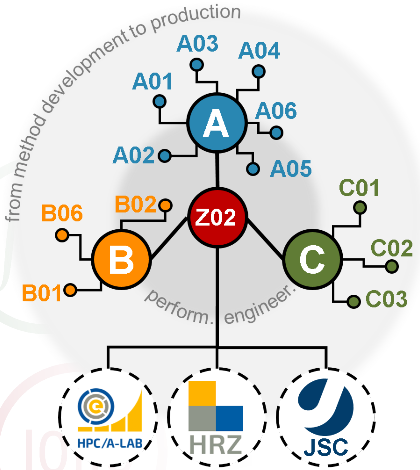

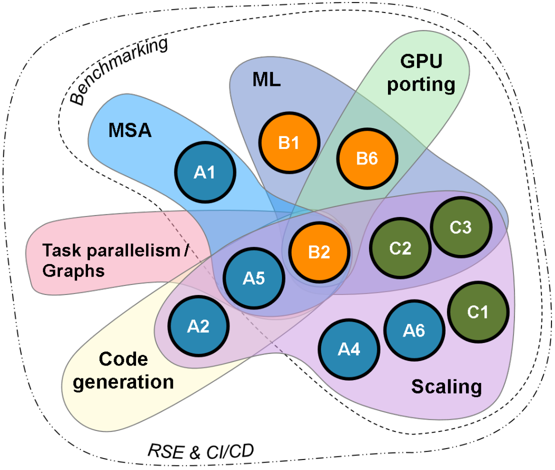

⟨NuMeriQS⟩

Example: Tensor Operations

- Important in projects **A02**, **A06**, **B02**, **C01**, **C02** and **C03**

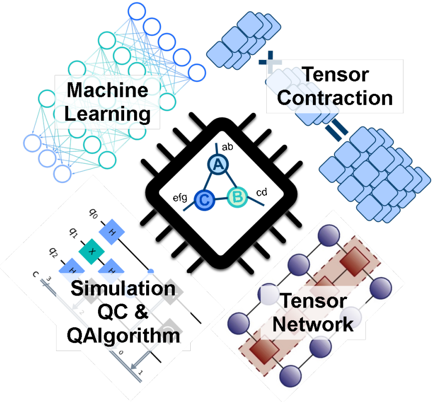

Example: Tensor Operations

- Important in projects **A02**, **A06**, **B02**, **C01**, **C02** and **C03**
- *NuMeriQS* method development → scale-out to larger systems → software bottleneck

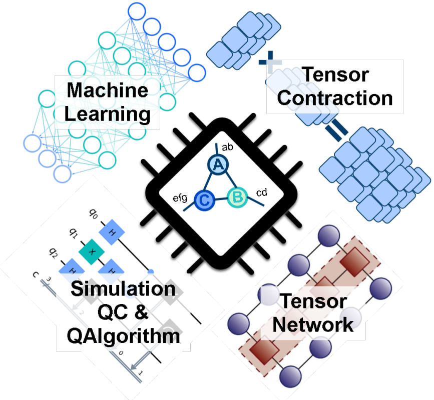

⟨NuMeriQS⟩

Example: Tensor Operations

- Important in projects **A02**, **A06**, **B02**, **C01**, **C02** and **C03**
- *NuMeriQS* method development → scale-out to larger systems → software bottleneck

Current State tensor libraries

P. Bientinesi https://arxiv.org/pdf/2103.13756

- avaialable software (80+)
- no standarization (API):
  - notation, nomenclarutre
  - functionality (app. driven)
- limited performance
  - n-dimensional arrays (GEMMS)
  - overhead: memory access
- limited hardware support

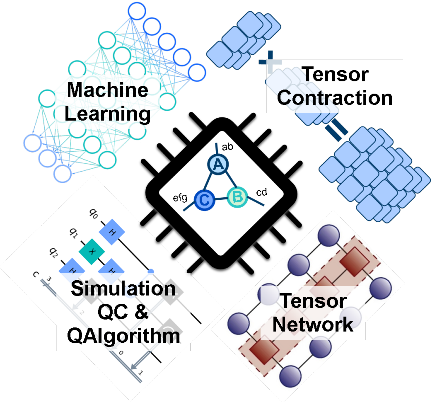

⟨NuMeriQS⟩

Example: Tensor Operations

- Important in projects **A02**, **A06**, **B02**, **C01**, **C02** and **C03**
- *NuMeriQS* method development → scale-out to larger systems → software bottleneck

Current State tensor libraries

P. Bientinesi https://arxiv.org/pdf/2103.13756

- avaialable software (80+)
- no standarization (API):
  - notation, nomenclarutre
  - functionality (app. driven)
- limited performance
  - n-dimensional arrays (GEMMS)
  - overhead: memory access
- limited hardware support

We invest in development now!

- explore existing frameworks
- perform CRC-based case studies
- aggregate cross-project expertise
- integrate HPC know-how from HPC/A-Lab and JSC
- training in best practices

⇒ co-design and prepare for new era

Co-Design

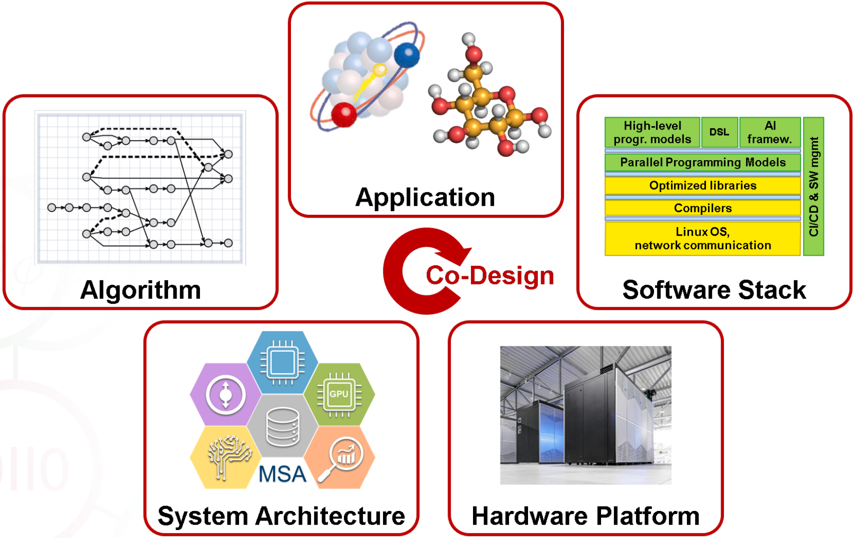

⟨NuMeriQS⟩

Z02: Work in Progress

Research Software Engineering:

- Building templates, examples, and workflows for benchmarking and CI/CD on JSC runner machines.
- **A06: Polynomial Filtering Diagonalization Application**
  - Interface and structural changes to enhance the modularity of the code.
  - Improved build process, usage, and documentation.

HPC: Performance Optimization

- **A02: ORCA**
  - Instrumentation to profile serial and parallel versions.
  - Analysis of coupled cluster and ORCA-AGE.
- **A06:**
  - Profiled CPU/GPU for initial assessment and implemented recommended modifications

⟨NuMeriQS⟩

A06: Polynomial filter diagonalization

Prof. David Luitz, UniBonn

https://scipost.org/SciPostPhys.11.2.021

Leonhard Moske: CPU+GPU

Single node: Arpack-ng/CUDA (C++)

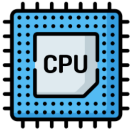

Diagonalization: Arnoldi Iterations

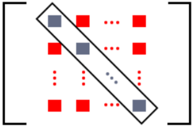

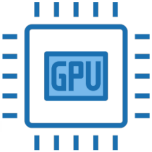

Polynomial transformation

A06: Polynomial filter diagonalization

- **JURECA:** AMD EPYC 7742, 64 cores, 2.25 GHz, NVIDIA A100 GPU

HPC: Performance Optimization

- Reduced blocking calls during device synchronization.
- Reduced memory movements between Host $\leftrightarrow$ Device.
- Reduced CPU calls on C++ objects, using intermediate copies.
- Improved speed and efficiency for small workloads.

GPU: A100- NQubits 20	Baseline	Current
Name	Time (%)	Time (%)
cudaDeviceSynchronize	29.5	0.0
cudaMalloc	28.7	0.1
cudaFree	27.7	0.0
cudaLaunchKernel	7.0	75.9
cudaMemcpy	2.7	24.1
cudaEventCreateWithFlags	2.4	0.0
cudaEventDestroy	2.0	0.0
cudaMemset	0.0	0.0

⟨NuMeriQS⟩

A06: Polynomial filter diagonalization

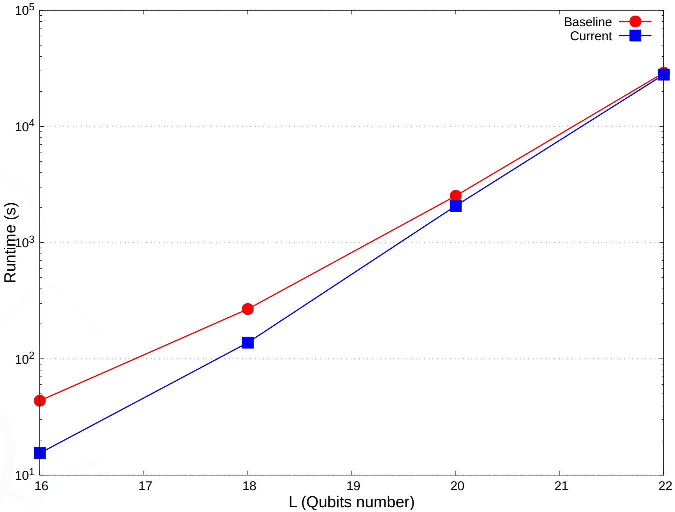

A06: Polynomial filter diagonalization

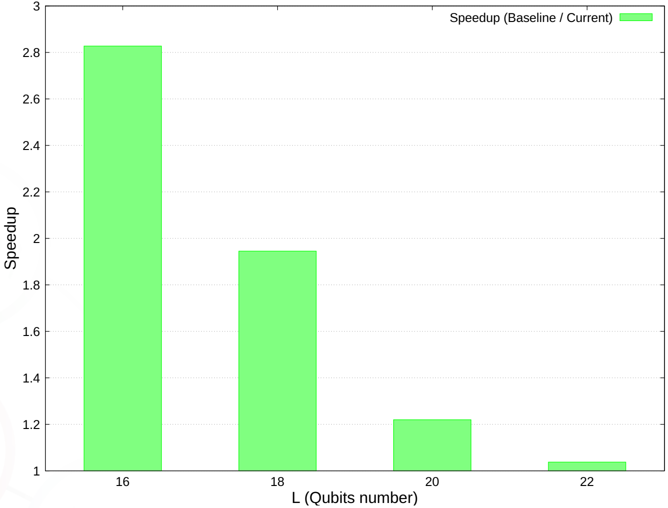

Take Aways

Computing is at the core of the scientific domains in *NuMeriQS*

Take Aways

Computing is at the core of the scientific domains in *NuMeriQS*

- Project **Z02 brings expertise in crucial computer science areas**:
  - Best RSE practices
  - Performance engineering
  - Heterogeneous computing
  - Efficient use of HPC systems

$\Rightarrow$ 1^{st} **funding period:** HPC+RSE support

$\Rightarrow$ 2^{nd} **funding period:** adopt disruptive computing technologies

Take Aways

Computing is at the core of the scientific domains in *NuMeriQS*

- Project **Z02 brings expertise in crucial computer science areas**:
  - Best RSE practices
  - Performance engineering
  - Heterogeneous computing
  - Efficient use of HPC systems

$\Rightarrow$ 1^{st} **funding period:** HPC+RSE support

$\Rightarrow$ 2^{nd} **funding period:** adopt disruptive computing technologies

- **CRC's interdisciplinarity:** ideal ground to share computational methods across domains
  - $\rightarrow$ Cross-fertilization: similar computational methods applied to different science domains
  - $\rightarrow$ Support from HPC experts
  - $\rightarrow$ Co-design feedback for future HPC systems

Take Aways

Computing is at the core of the scientific domains in *NuMeriQS*

- Project **Z02 brings expertise in crucial computer science areas**:
  - Best RSE practices
  - Performance engineering
  - Heterogeneous computing
  - Efficient use of HPC systems

$\Rightarrow$ 1^{st} **funding period:** HPC+RSE support

$\Rightarrow$ 2^{nd} **funding period:** adopt disruptive computing technologies

- **CRC's interdisciplinarity:** ideal ground to share computational methods across domains
  - $\rightarrow$ Cross-fertilization: similar computational methods applied to different science domains
  - $\rightarrow$ Support from HPC experts
  - $\rightarrow$ Co-design feedback for future HPC systems

We aim to bring key scientific codes onto Exascale systems