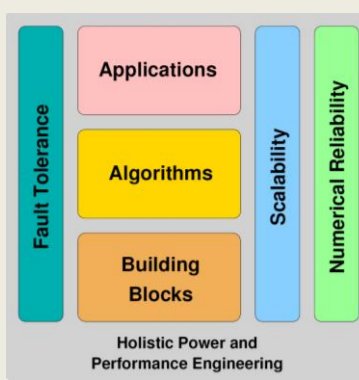


ESSEX-II: Equipping Sparse Solvers for Exascale



Gerhard Wellein
University of Erlangen
Department of Computer Science

Achim Basermann
German Aerospace
Center



Holger Fehske
University of Greifswald
Institute for Physics



Georg Hager
Erlangen Regional
Computing Center



Bruno Lang
University of Wuppertal
Applied Computer Science



Resilience, Performance Engineering,
Parallelization, Optimization

Scalable Preconditioners,
Eigenvalue & Linear Solvers

Quantum Physics Algorithms &
Applications

Performance Engineering, Tools,
Parallelization, Optimization

Efficient Direct and Iterative
Eigensolvers

Tetsuya Sakurai
University of Tsukuba
Department of Computer Science



Kengo Nakajima
University of Tokyo
Information Technology Center



Collaborations

J. Dongarra
Blocked sparse
kernels
ICL
OR

M. Hoemmen
Node performance,
FT
Sandia
National
Laboratories

F. Mueller
Skeptical
programming
COZ

SPPEXA impact

EXASTEEL
PE for kernels,
Pardiso

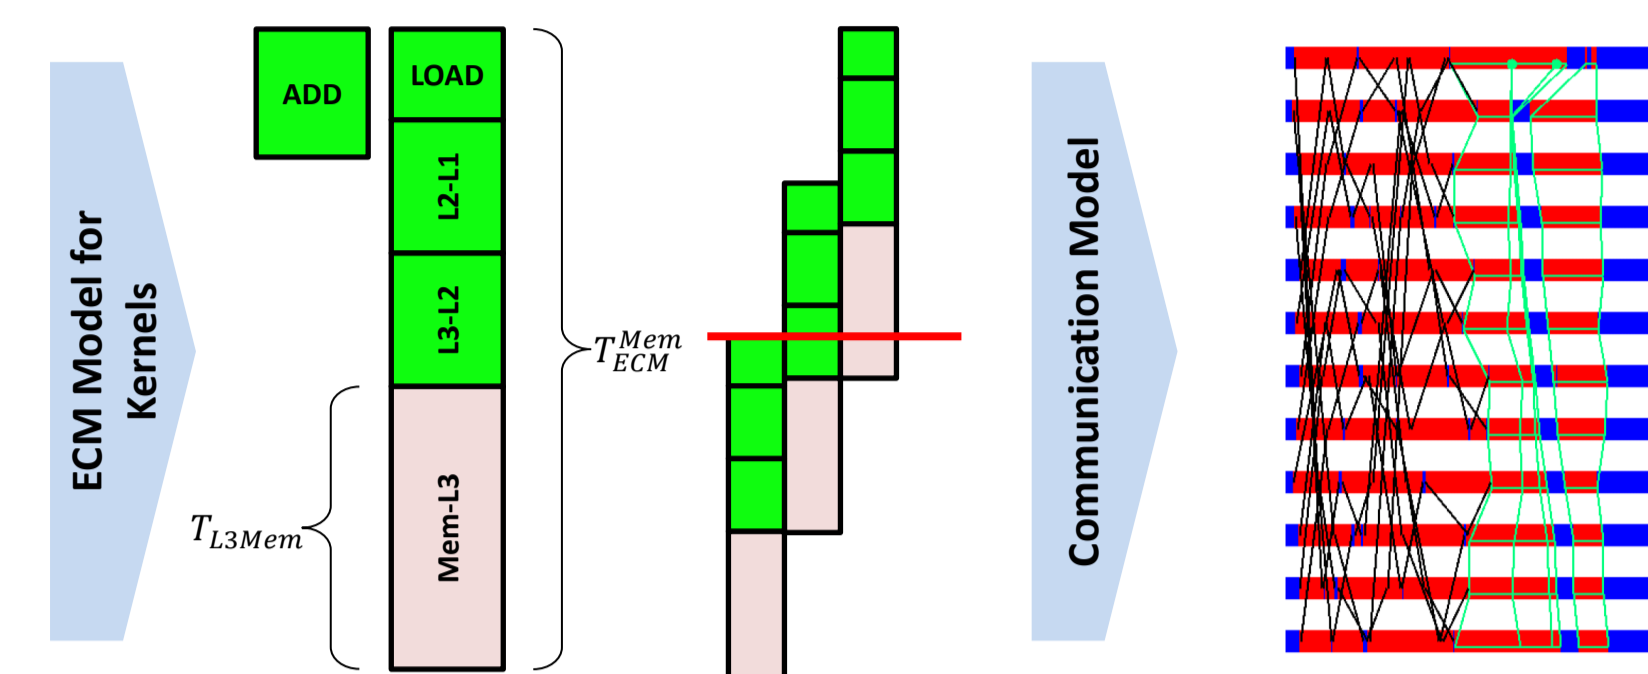
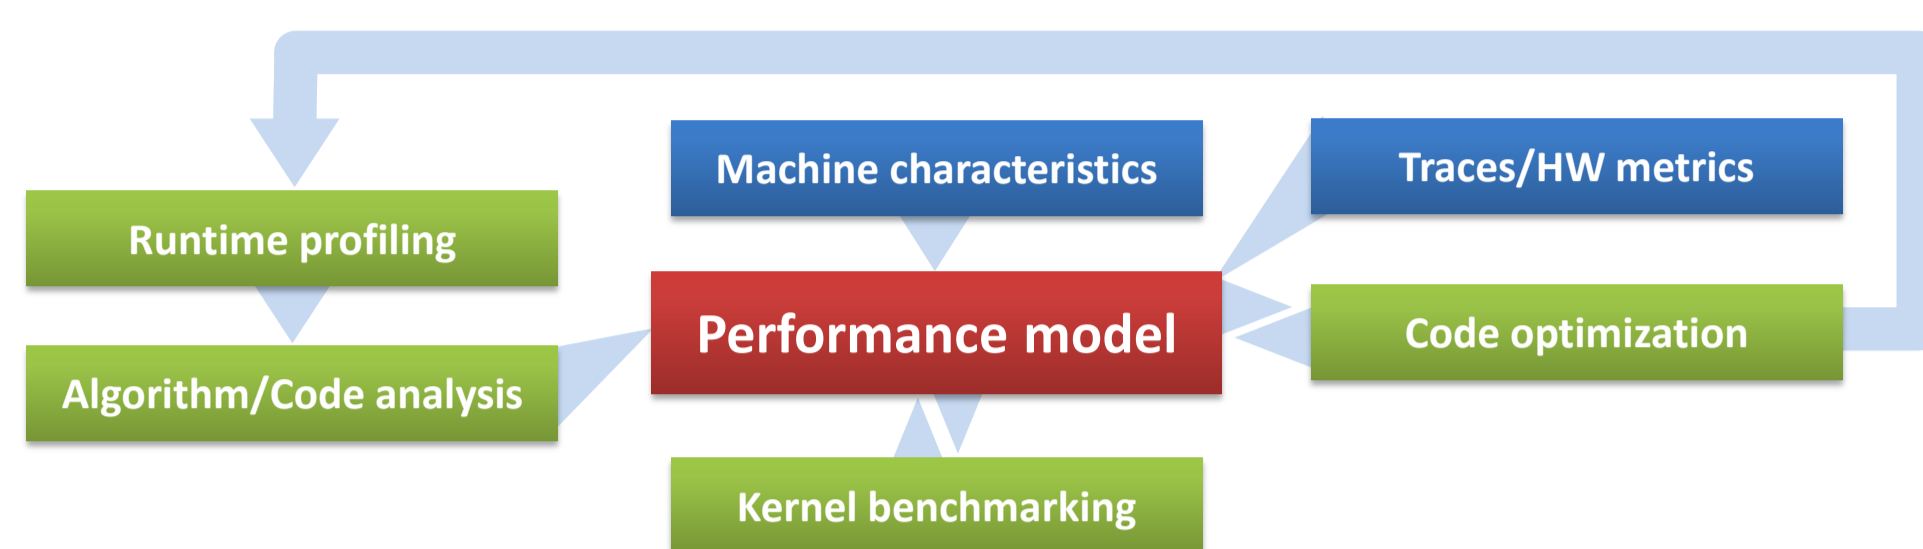
EXA-DUNE
Sparse matrix
formats

**EXASTENCILS,
TERRA-Neo**
PE for kernels

Advanced Eigensolvers

Scalable Preconditioners, AMG

Holistic Performance and Energy Engineering



```

[V] := [v]_{k,R-1}
[W] := [w]_{k,R-1}
[V] := [v]_{k,R-1}
for m = 1 to M/2 do
  Initialization steps and computation of \eta_{m,1}
  swap([V], [W])
  [W] := 2\eta(H - \eta)V - [W] &
  \eta_{m+1} := ([W] \cdot [V]) &
  \eta_{m+1} := ([W] \cdot [V])
end for
(1.88/R+0.35) Bytes/Flop

for r = 0 to R-1 do
  [v] := [v]_{k,R-1}
  Initialization steps and computation of \eta_{r,1}
  for m = 1 to M/2 do
    swap([v], [w])
    [w] := [w] - \eta[v]
    [w] := [w] - \eta[v]
    [w] := [w] - 2\eta[v]
    \eta_{m+1} := ([w] \cdot [v]) &
    \eta_{m+1} := ([w] \cdot [v])
  end for
  3.4 Bytes/Flop
end for

for r = 0 to R-1 do
  [v] := [v]_{k,R-1}
  Initialization steps and computation of \eta_{r,1}
  for m = 1 to M/2 do
    swap([v], [w])
    [w] := 2\eta(H - \eta)v - [w] &
    \eta_{m+1} := ([w] \cdot [v]) &
    \eta_{m+1} := ([w] \cdot [v])
  end for
  2.2 Bytes/Flop
end for
    
```

From Physics to Performance

Physical Problem: microwave drive, LC circuit, vibrating capacitor, $C(\hat{x})$

Efficient DoF numbering: QSE, Hilbert space reduction

Problem-specific: Reordering Techniques

Generic: Reordering Techniques

Concurrency: Basic Ops, Algorithms, Communication Patterns

GHOST-II

Insight into Constituents Performance

```

# parallel Kahan sum
for (i=0; i<N; ++i) {
  double prod = a[i]*b[i];
  double y = prod-c;
  double t = sum+y;
  c = (t-sum)-y;
  sum = t;
}
    
```

Augmented Kernel Modeling: Model, Validation

Error-Correcting Reductions: $\|w_0\| = \|Aq_j\| \leq \|A\|_2 \|q_j\|_2$, $\|w_0\| \leq \|A\|_2 \leq \|A\|_F$, $|h_{i,j}| \leq \|A\|_F$

Skeptical Programming: Resiliency & Reliability Thrust

PE support: Performance/Energy Thrust

Efficient Building Blocks: ESSR, Fundamental algorithms

Harnessing HPC for Quantum Systems

Conservative Quantum Systems: $i\hbar \frac{\partial}{\partial t} \Psi = H\Psi$

Dissipative Quantum Systems: $i\hbar \frac{\partial}{\partial t} \rho = [H, \rho] + D$

$A = H$, $A = A^+$

$A = A(H, D)$, $A \neq A^+$

General Physics Topics: Electronic Structure, Correlations, Quantum Coherence, Condensed Matter ↔ Light-Matter Coupling ↔ Environment

Relevant Physical Quantities: Density of States, Spectral Functions, Density Matrix & Wave Function Dynamics, Topological Invariants, Entanglement Measures

Number of Required Eigenstates: Few (extremal) - Many (interior) - All (approximation)

Polynomial Techniques: Adaptive: Search Space & Degree Control

Example: ChebFD for inner Eigenvalue bulks

Example: ChebTP/KPM for dynamic & static properties

Idea: High-order polynomial filters

Propagation $|\psi(t + \Delta t)\rangle = e^{-iH\Delta t/\hbar} |\psi(t)\rangle$

Expansion $e^{-iH\Delta t/\hbar} |\psi(t)\rangle \sim \sum_{n=1}^{N_{ch}} c_n |v_n\rangle$

No synchronization

Chebyshev recurrence $|v_{n+1}\rangle = 2H|v_n\rangle - |v_{n-1}\rangle$

Functional parallelism $|\psi(t + \Delta t)\rangle = e^{-iH\Delta t/\hbar} |\psi(t)\rangle$, $\langle \psi(t) | Op | \psi(t) \rangle$

Top. Ins., $N=10^9$, polynomial degree = 7839, search space = 256

From Algorithms to Parallel Computing

Linear Eigenvalue Problem: $Ax = \lambda x$, $A = A^+$, $Ax = \lambda Bx$, $A \neq A^+$, $P^{-1}Ax = P^{-1}b$

physical properties: dynamic, static

Scalability: Reordering & Partitioning, Asynchronous Preconditioning, Communication Avoiding/Hiding

Applications: Matrix Generation, FEAST, Krylov, JaDa, KPM, ChebTP, ChebFD

Automatic testing & validation: Trifinos, Thread & Affinity Mgmt, Block vector kernels, Interoperability, GHOST

Numerical Reliability: High-Precision Reductions, Skeptical Programming, Mixed/Adapted Precision

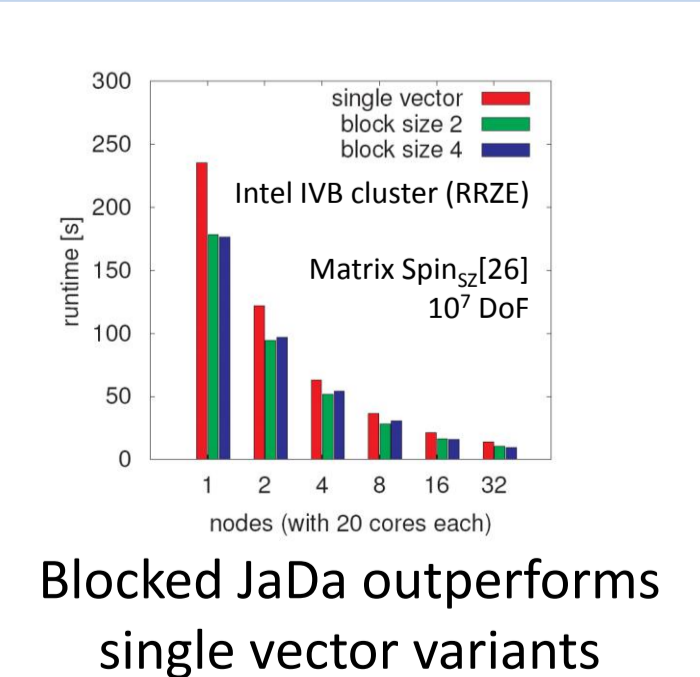
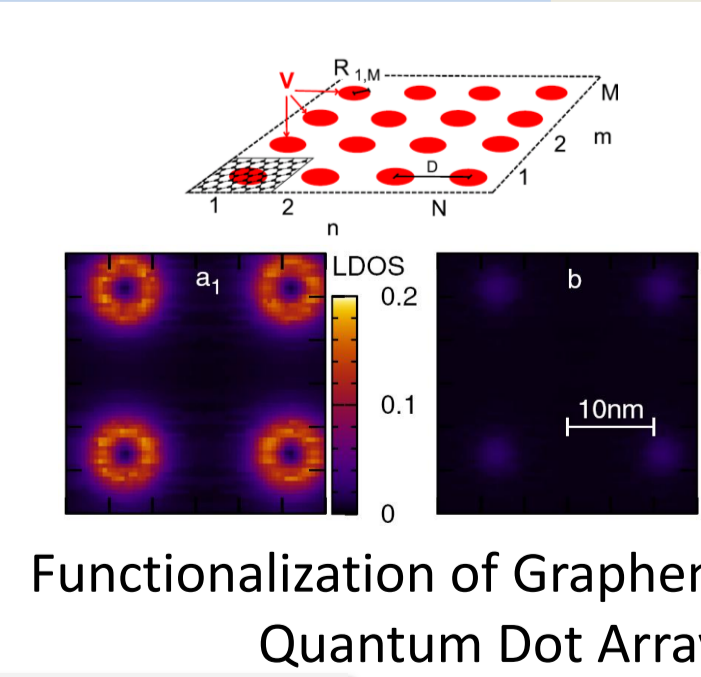
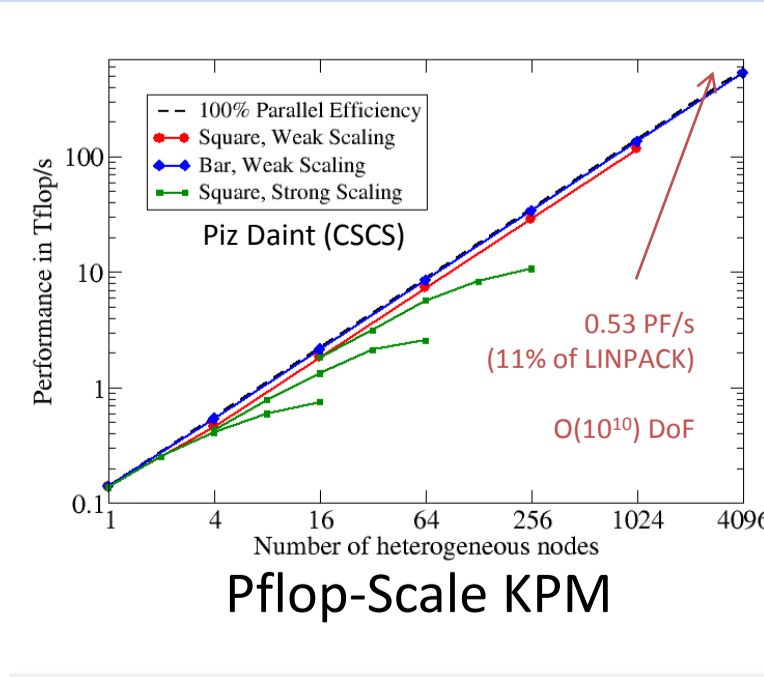
Software Infrastructure: MPI-X, Tools, Topology & Affinity Awareness

z-Pares: Generalization, $A(\lambda)x = 0$, Highly Parallel, $Ax = B$

Nonlinear Eigenvalue Problem: GHOST

MAGMA, ViennaCL, EXA-DUNE, SELL-C-σ performance-portable sparse matrix format

GHOST, Intel Xeon Phi, CPU, Nvidia GPU, Fully Heterogeneous General Hybrid Optimized Sparse Toolkit



KPM, Compute disjoint bulks of eigenvalues, Loop over Gauss points z_j , Solve k systems $(z_j I - A)u_j = v_j$, CARP-CG, Basic linear preconditioners and iterative linear solvers, KPM estimators for FEAST

Starting Point: ESSEX-I

Silent Data Corruption Detection

NSSM for Nonlinear Problems

AMG Preconditioner

Advanced Performance Models

2016

2017

2018

High-Precision Reductions

KPM/ChebTP & Block JaDa Release

Advanced Block JaDa Preconditioning

ChebFD/Krylov Release

Application Benchmark Suite

Final Library Release



tiny.cc/GHOST