Computational Scale Bridging using PETSc on ~~458K~~ 786K Cores

Oliver Rheinbach (TU Bergakademie Freiberg) and Axel Klawonn (U of Cologne);
Martin Lanser (U of Cologne)

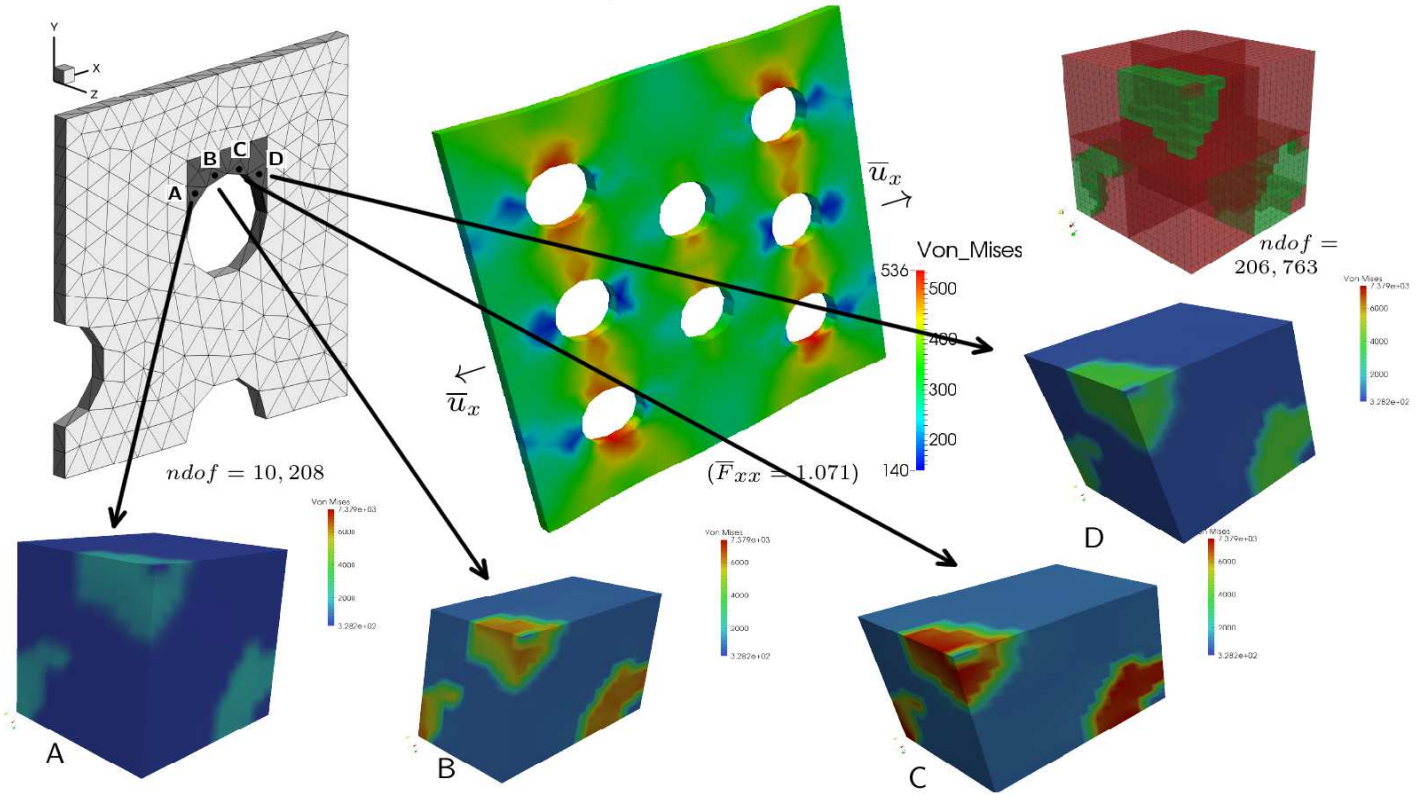

DFG-Schwerpunktprogram SPP 1648 - Software for Exascale Computing

EXASTEEL - Bridging Scales for Multiphase Steels

- FE^2: highly concurrent multiscale method (reduction $10^3 \ldots 10^6$)
- Hybrid Domain Decomposition/Multigrid
- Modelling of advanced multiphase steels,
- Deep drawing of metal sheets
- Software environment based on **PETSc**, FEAP, BoomerAMG

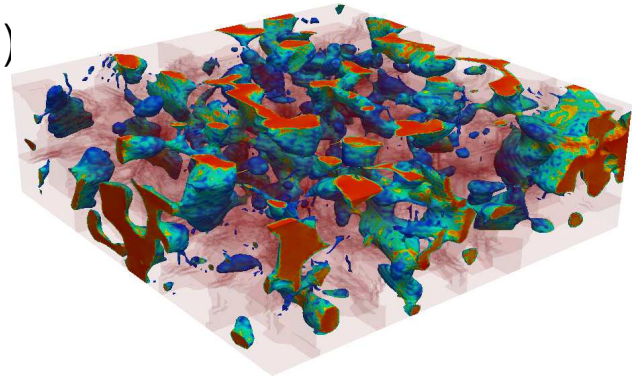

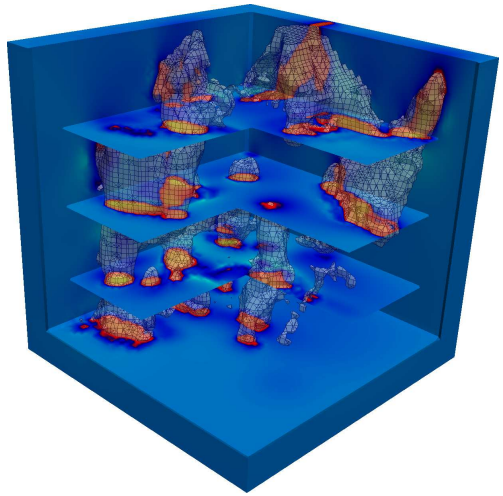

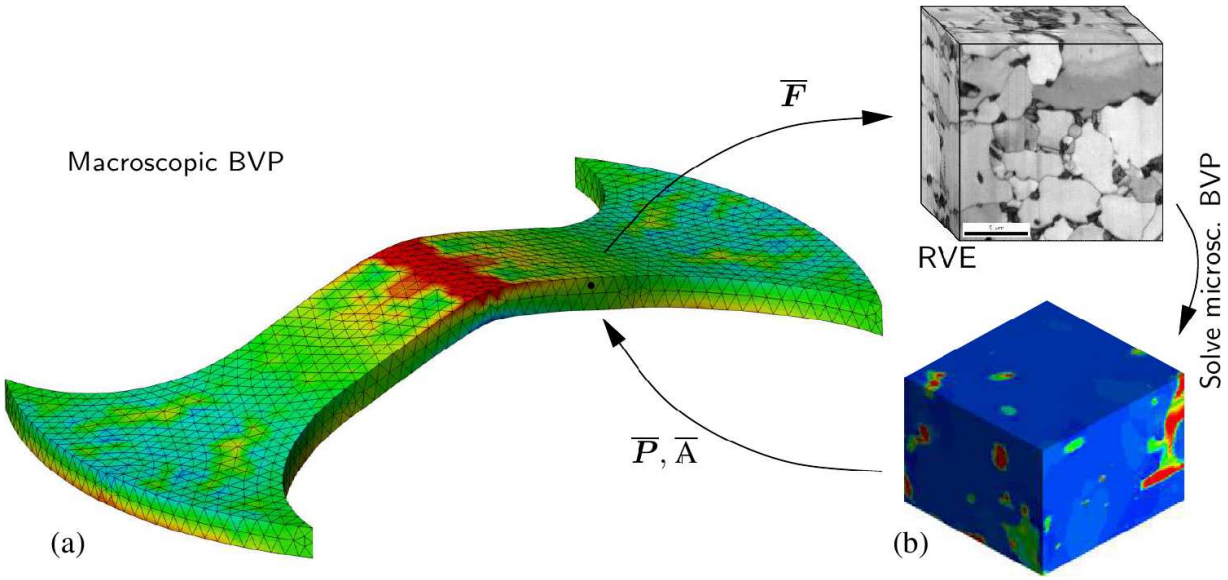

PIs: D. Balzani, A. Klawonn, O. Rheinbach, J. Schröder, G. Wellein

My oldest PETSc output I could still find

```
FetiDP example program
Built Feb  8 2005 (23:32:19). Machine using 8 byte arithmetics (linux).
Started on 2 processors, 4096 domains: Processors will hold 2048 domains each.
Creating Mat_Feti... 1-2048 2049-4096 ...Creating finished. Setting use_Q false.
Loading Mat_Feti...
Tolerances now: 1e-10 1e-16 1e+05 500
Calculating RHS...(Coarse Problem Size:24066)...RHS finished. Starting iterations...
  0 KSP Residual norm 2.206897289671e-03 % max 1 min 1 max/min 1
  1 KSP Residual norm 3.106042932023e-04 % max 1.25444 min 1.25444 max/min 1
  2 KSP Residual norm 9.086829705003e-05 % max 1.39086 min 1.17765 max/min 1.18104
[...]
 16 KSP Residual norm 6.644544424066e-13 % max 2.94886 min 1.00449 max/min 2.93568
 17 KSP Residual norm 1.585043365976e-13 % max 2.95118 min 1.00388 max/min 2.93978
ConvergedReason (2 is rtol): 2 ...Iterations: 17
```

One of the most recent PETSc outputs

```
Warning N<nx*ny: 16<1024 * 1024. Setting N = nx * ny
Started on 1048576 processors, 1048576 domains: Processors will hold 1 domains each.
dofs_per_node: 2 Creating Mat_Feti... 1-1 -coarse_part: 1.000000 Inexact. Global:./FetiDP_ Local:./F
[...]
 RHS...Starting iterations...
  0 KSP Residual norm 5.780049774075e-01 % max 1.000000000000e+00 min 1.000000000000e+00 max/min 1.0
  1 KSP Residual norm 2.435888256695e-01 % max 2.262906444502e+00 min 2.262906444502e+00 max/min 1.0
  2 KSP Residual norm 1.461838416609e-01 % max 2.694505665485e+00 min 1.556463078557e+00 max/min 1.7
[...]
 43 KSP Residual norm 7.830756595320e-11 % max 1.544657844221e+01 min 7.071218296542e-01 max/min 2.1
 44 KSP Residual norm 5.582699597676e-11 % max 1.545478453116e+01 min 7.040560438391e-01 max/min 2.1
Finished iterating...
ConvergedReason (2 is rtol): 2 ...Iterations: 44
```

Short Introduction: Two-Scale Modeling with the FE^2-Method

- Characteristic lengths: $L/d \approx 10^4$–10^6.
- Brute force FE discretization not feasible
- Scale-bridging procedure is essential.
- FE^2: FE-discretization of both scales, reduces problem size by factor 10^3–10^6.
- Algorithmic procedure for FE^2-scheme:
  1. Apply boundary conditions
  2. Solve microscopic BVP
  3. Return volumetric averages

Miehe, Schröder, Schotte 1999; Schröder 2000; Feyel 1999; Smit, Brekelmans, Meijer 1998; Kouznetsova, Brekelmans, Baaijens 2001

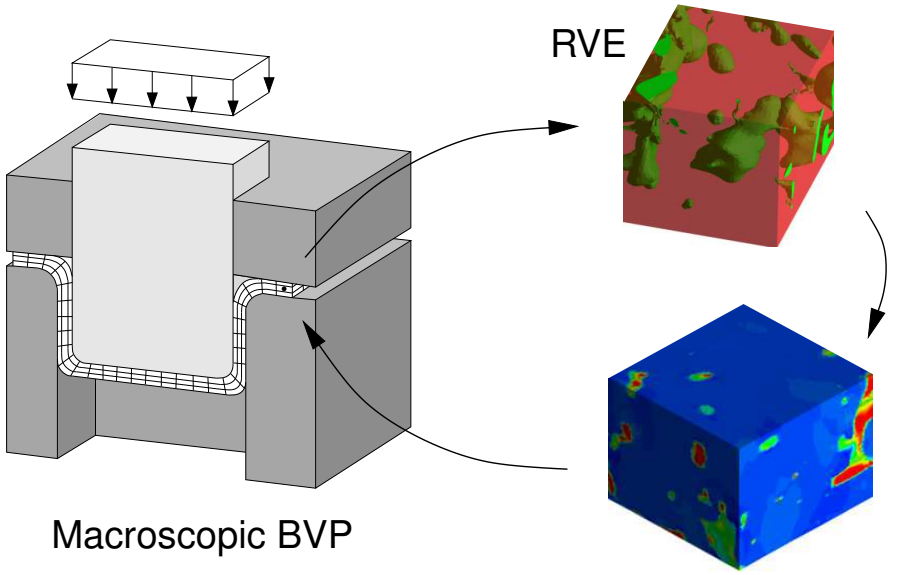

- Improvement of material models:
  - temperature-dependent FE^2-scheme
  - finite strain phase transformations

Remaining orders of magnitude are resolved by highly parallel solver algorithms and performance engineering.

FE^2 Computational Scale Bridging

(a) Macroscopic boundary value problem (BVP).
(b) Microscopic BVP (one in each macroscopic Gauss point).
The results on the microscale replace a phenomenological constitutive law on the macroscale:

$$\overline{P} = \frac{1}{V} \int_{\mathcal{B}} P \, dV$$

$$\overline{A} = \frac{\partial \overline{P}}{\partial \overline{F}} = \frac{\partial}{\partial \overline{F}} (\frac{1}{V} \int_{\mathcal{B}} P \, dV)$$

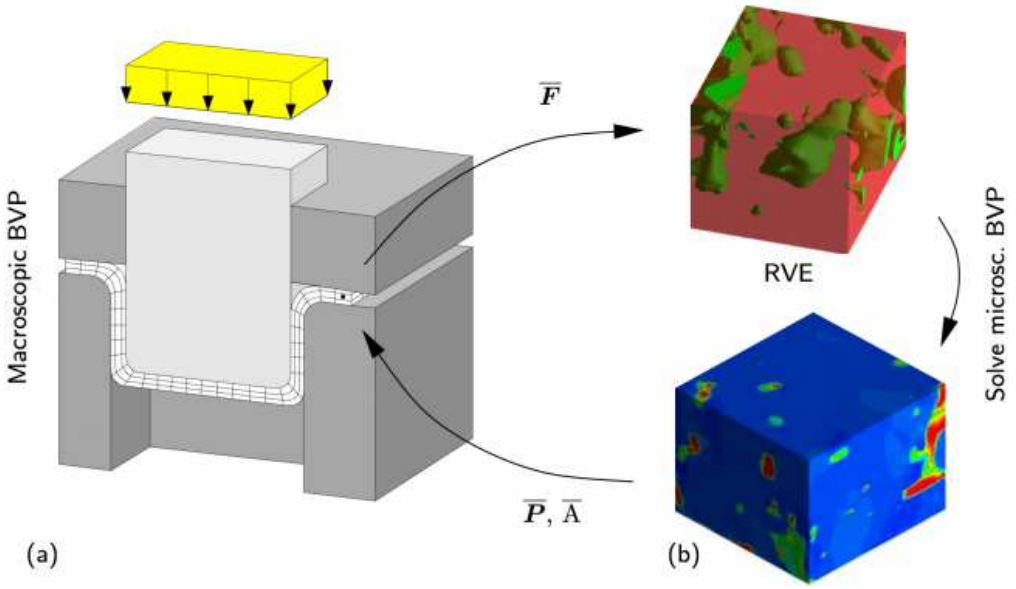

Classical FETI-DP-Methods (Farhat, Lesoinne, LeTallec, Pierson, Rixen)

Decompose into nonoverlapping subdomains $\Omega^{(i)}$. Assemble in the primal variables using partial assembly operator R_Π and enforce continuity of the dual variables by jump operator B and Lagrange multipliers.
Defining

$$K = \begin{bmatrix} K^{(1)} & & \\ & \ddots & \\ & & K^{(N)} \end{bmatrix}$$

we then obtain the FETI-DP master system

$$\begin{bmatrix} R_\Pi^T K R_\Pi & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \tilde{u} \\ \lambda \end{bmatrix} = \begin{bmatrix} R_\Pi^T f \\ 0 \end{bmatrix}$$

using the partial assembly operator R_Π^T.
After elimination of $\tilde{u}$ solve

$$F\lambda = d$$

iteratively.

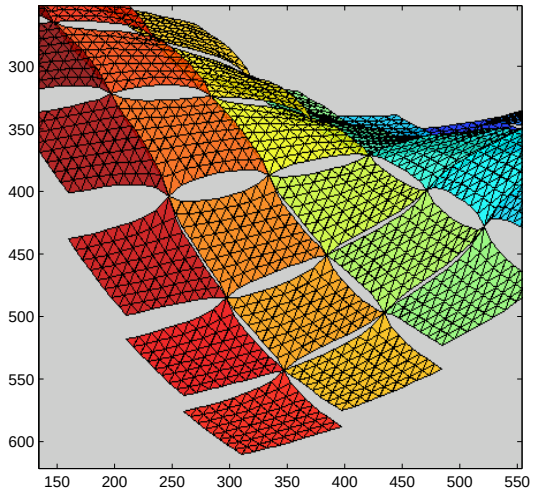

Remark: Reformulation which allows use of AMG for the coarse problem was denoted irFETI-DP (Klawonn, Rheinbach 2007).

FETI and FETI-DP Domain Decomposition Methods

- **F**inite **E**lement **T**earing and **I**nterconnecting (FETI) methods were introduced by Farhat/Roux (1991), FETI-DP by Farhat/Lesoinne/LeTallec/Pierson/Rixen (2001) .

  Further algorithmic and theoretical work by Brenner, Dostal, Dryja, Feyel, Farhat, Gosselet, Jarosova, Kim, Klawonn, Kozubek, Langer, Lesoinne, Le Tallec, Li, Mandel, Pavarino, Pechstein, Pierson, Rey, R., Rixen, Scheichl, Steinbach, Tezaur, Toselli, Tu, Vasseur, Widlund, Zampini ...
- FETI-DP for elasticity with large coefficient jumps aligned with the subdomain boundary: Klawonn/Widlund CPAM 2006 (Theory), Klawonn/R. CMAME (2007) (Numerical results)
- FETI and FETI-DP methods for large coefficient jumps not aligned with the interface: Pechstein/Scheichl (2008, 2011, FETI for diffusion, theory+numerics), (Dryja/Sarkis, 2010, FETI-DP for diffusion, theory) Pechstein/Sarkis/Scheichl 2011, FETI-DP, theory+numerics), Klawonn/R., 2007, Spillane, Jolivet, Nataf (2013).
- FETI and FETI-DP (and related, e.g., BDDC) for (almost) incompressible elasticity: Jing Li (2002), FETI for Stokes; Li/Widlund (2005), BDDC, a.i.e., Dohrmann/Widlund (2009), Overlapping Schwarz for a.i.e.; Pavarino/Widlund/Zampini (2010), BDDC, spectral elements for a.i.e, Badia (2014), Zampini (2014,15)

FE^{2}TI

Repeat until convergence:

Microscopic calculations:

1. Apply boundary conditions $x = \overline{F}X$ on $\partial\mathcal{B}$.
2. Solve microscopic nonlinear boundary value problem **using (ir)FETI-DP** or related methods.
3. Compute macroscopic stresses $\overline{P} = \frac{1}{V}\int_{\mathcal{B}} P\,dV$.
4. Compute macroscopic tangent moduli $\overline{A} = \frac{\partial}{\partial \overline{F}}(\frac{1}{V}\int_{\mathcal{B}} P\,dV)$.

Macroscopic calculations:

5. Set up tangent matrix and rhs of linearized macroscopic BVP using $\overline{P}$ and $\overline{A}$.
6. Solve linearized macroscopic boundary value problem.
7. Update macroscopic deformation gradient $\overline{F}$.

Similar methods are around, e.g., as HMM (Hierarchical Multiscale Method) or, e.g., presented at DARPATech (2007) by the SIPS Team (fatigue of aluminum).

Consistent Tangent

$\overline{P}$ is a simple volumetric average, but $\overline{A}$ is more involved.

$$\overline{A} = \frac{\partial}{\partial \overline{F}}\left(\frac{1}{V}\int_{\mathcal{B}} P(F)\, dV\right) = \frac{1}{V}\int_{\mathcal{B}} \frac{\partial P(F)}{\partial \overline{F}}\, dV = \frac{1}{V}\int_{\mathcal{B}} \frac{\partial P(F)}{\partial F} : \frac{\partial F}{\partial \overline{F}}\, dV$$

Splitting $F = \overline{F} + \widetilde{F}$ into the constant part $\overline{F}$ and a fluctuating part $\widetilde{F}$ leads to

$$\overline{A} = \frac{1}{V}\int_{\mathcal{B}} \frac{\partial P(F)}{\partial F} dV + \frac{1}{V}\int_{\mathcal{B}} \frac{\partial P(F)}{\partial F} : \frac{\partial \widetilde{F}}{\partial \overline{F}} dV = \frac{1}{V}\int_{\mathcal{B}} A dV + \frac{1}{V}\int_{\mathcal{B}} A : \frac{\partial \widetilde{F}}{\partial \overline{F}} dV$$

$\Rightarrow$ Volumetric average over the tangent moduli of all finite elements in the RVE.

$\Rightarrow$ Reformulation using the balance of momentum finally leads to $-\frac{1}{V} L^T K^{-1} L$, where K is the stiffness matrix of the microscopic problem.**

$\Rightarrow$ In the final microscopic Newton step, 9 additional linear solves have to be performed. (4 add. solves in 2D)

See **Hackl, Schröder (Springer 2014)** for details.

Micro-Macro Analysis for Thermomechanics

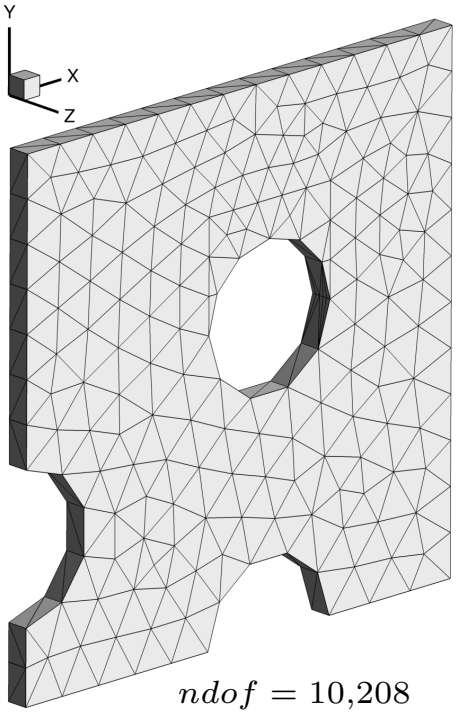

$ndof = 10{,}208$

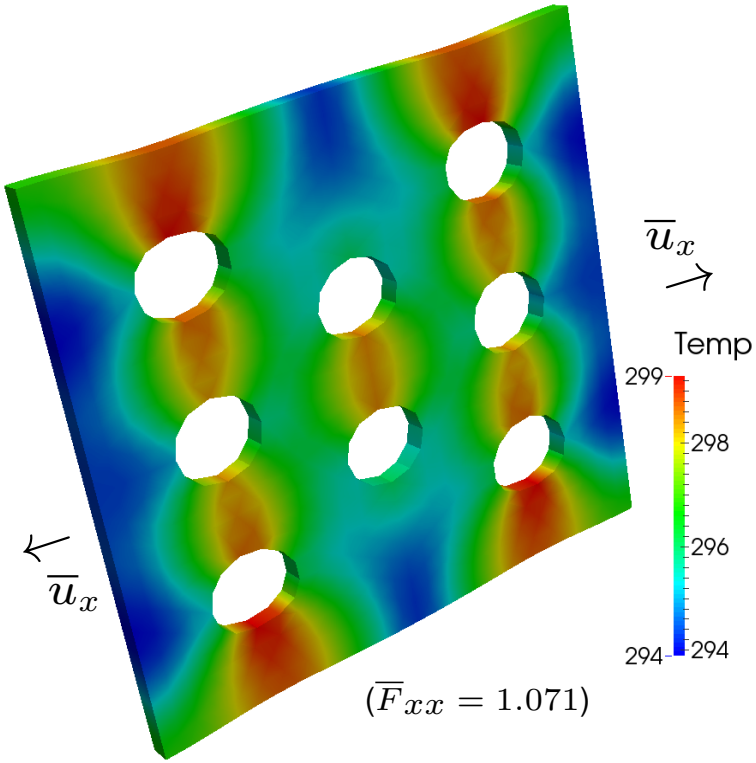

$(\overline{F}_{xx} = 1.071)$

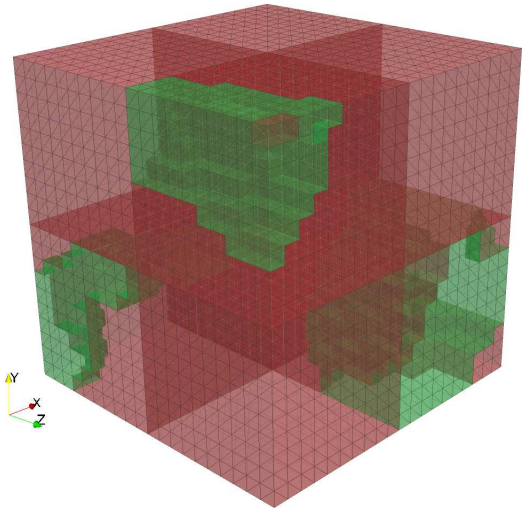

$ndof = 206{,}763$

Material parameters

$E = 206{,}000\,\text{MPa}$
$\nu = 0.3$
$y_0 = 280.0\,\text{MPa}$
$y_\infty = 600.0\,\text{MPa}$
$k = 55.0\,\frac{\text{N K}}{\text{s}}$
$c = 0.46 \times 10^9\,\frac{\text{mm}^2}{\text{s}^2 K}$
$\alpha_t = 1 \cdot 10^{-5}\,\frac{1}{\text{K}}$

Macroscopic problem

(For $1/4^{th}$ geometry)

- Total strains in X-dir.- 7.1%
- Total process time: 10s
- Time increment: 0.001s
- Room temperature, θ_0 = 293K
- Linear hardening, h=1250 MPa

Yield strength data

(in MPa)	ferrite	martensite
y_0	260.0	1000.0
y_∞	580.0	2750.0

Temperature dependency:

$$y_{0/\infty} = \langle \omega(\theta - \theta_0) + y_{0/\infty} - \tilde{y}_{0/\infty} \rangle + \tilde{y}_{0/\infty}$$

with $\omega_f = -0.295$ and $\omega_m = -0.586$

Thermoplasticity model, see SIMO & MIEHE [1992] and publications by SIMO, MIEHE, WRIGGERS, ORTIZ, GEERS, ...

Micro-Macro Analysis using FETI-DP

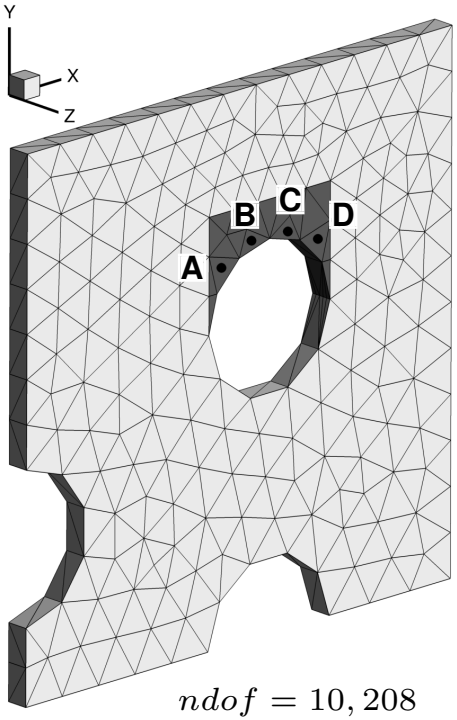

$ndof = 10,208$

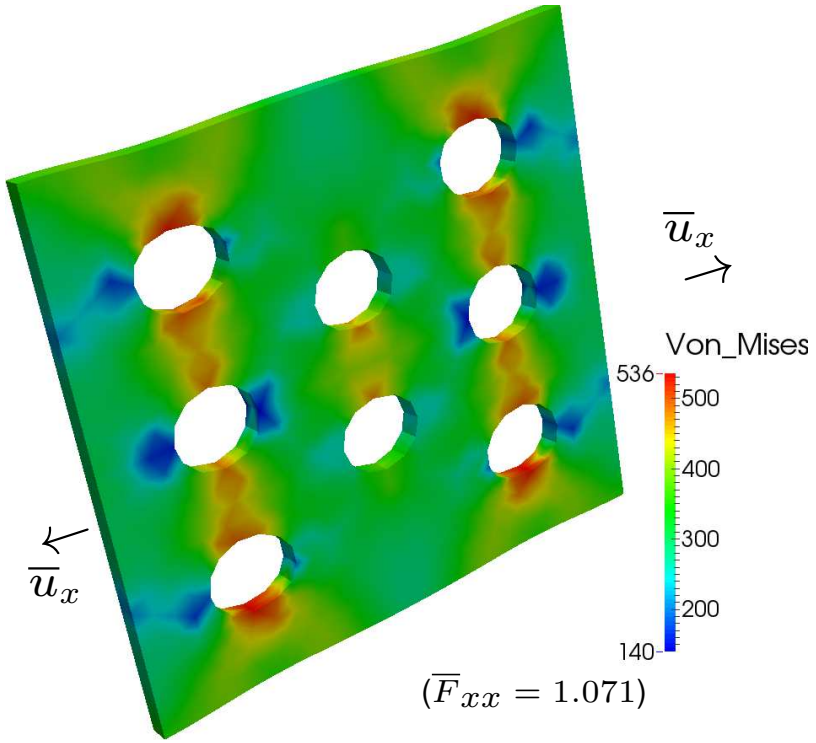

$(\overline{F}_{xx} = 1.071)$

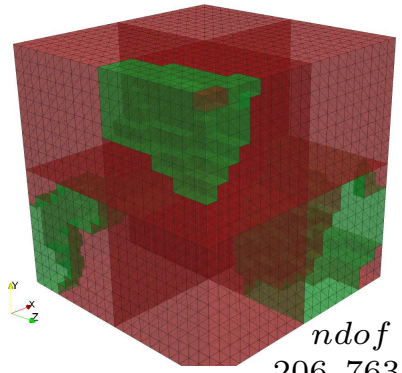

$ndof = 206,763$

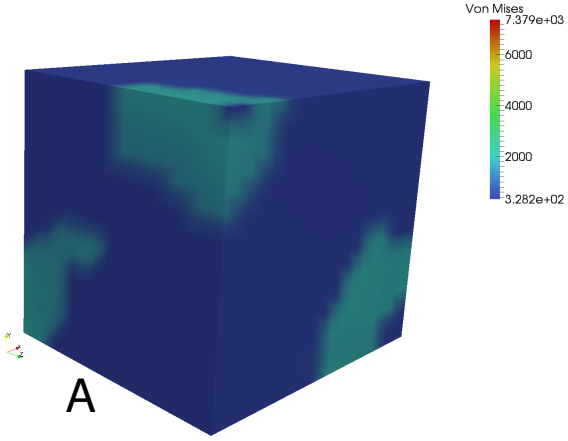

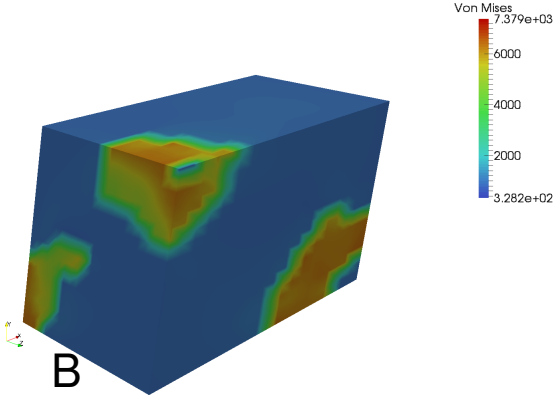

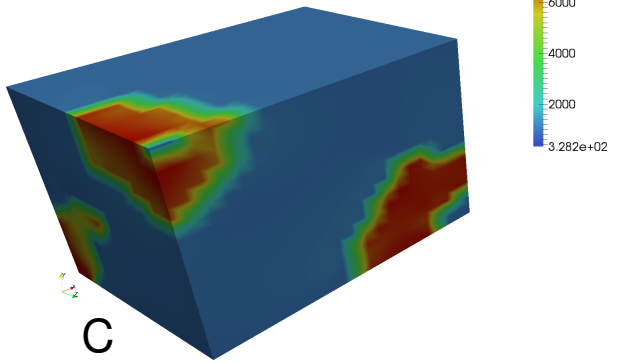

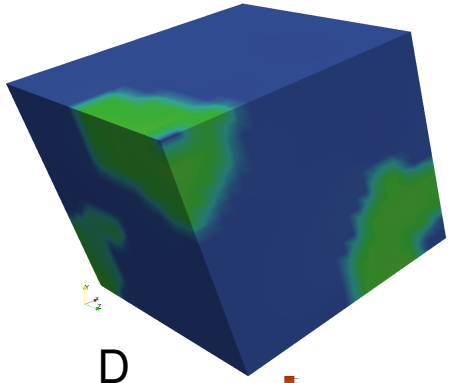

(D. Balzani, D. Brands, A. Gandhi, S. Holger, A. Klawonn, M. Lanser, O. Rheinbach, J. Schröder, and G. Wellein)

Hybrid MPI/OpenMP Parallelization on SandyBridge

Simple hybrid parallelization approach:

- OpenMP parallel finite element assembly into a MPI parallel **PETSc** matrix
- Pardiso shared memory sparse direct solver; developer Olaf Schenk now has joined the consortium.

Threads	Time	Efficiency
1	133.7s	**100%**
2	84.0s	**80%**
3	74.6s	**60%**
4	57.7s	**58%**
5	54.5s	*49%*
6	53.5s	*42%*
7	51.9s	*37%*
8	46.3s	*36%*
9	45.2s	*33%*
10	44.2s	*30%*

Runtime for one to ten threads per MPI process (Setup 1, i.e., one MPI process on each of the four nodes, Neo-Hookean hyperelasticity problem in 2D)

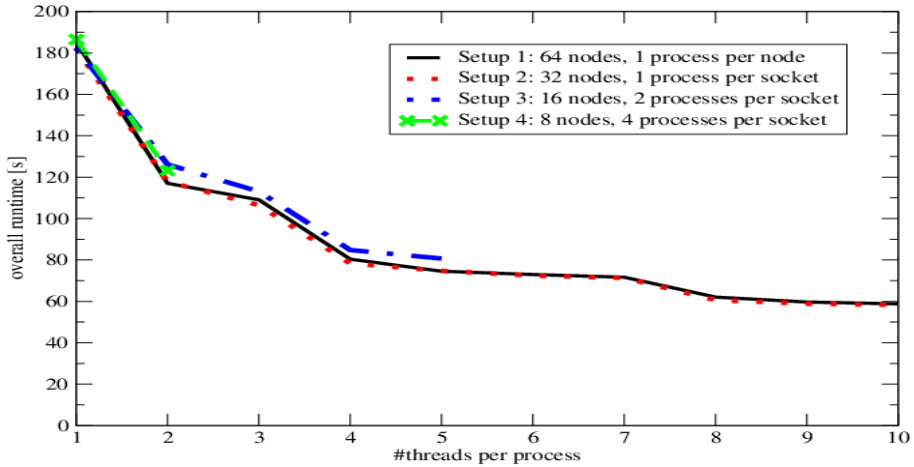

Socket scaling (hyperelasticity in 2D) 64 (lower) MPI ranks and one to 10 threads for each rank.

Klawonn, Lanser, Rheinbach, Stengel, and Wellein, "Hybrid MPI/OpenMP parallelization in FETI-DP methods". Springer Lect. Notes (accepted).
TU Freiberg Preprint 2015-02.

Collaboration with Olaf Schenk (Lugano) on Pardiso has started.

Fully Coupled FE^2 Computational Scale Bridging in 3D on JUQUEEN (JSC Jülich)

bg_size	MPI-ranks	#RVEs	Time to Solution	
512	8 192	16	184.86s	100.0%
1 024	16 384	32	185.09s	99.9%
2 048	32 768	64	185.61s	99.6%
4 096	65 536	128	185.72s	99.5%
8 192	131 072	256	186.43s	99.2%
16 384	262 144	512	186.61s	99.1%
24 576	393 216	768	187.32s	98.7%
28 672	458 752	896	187.65s	98.5%

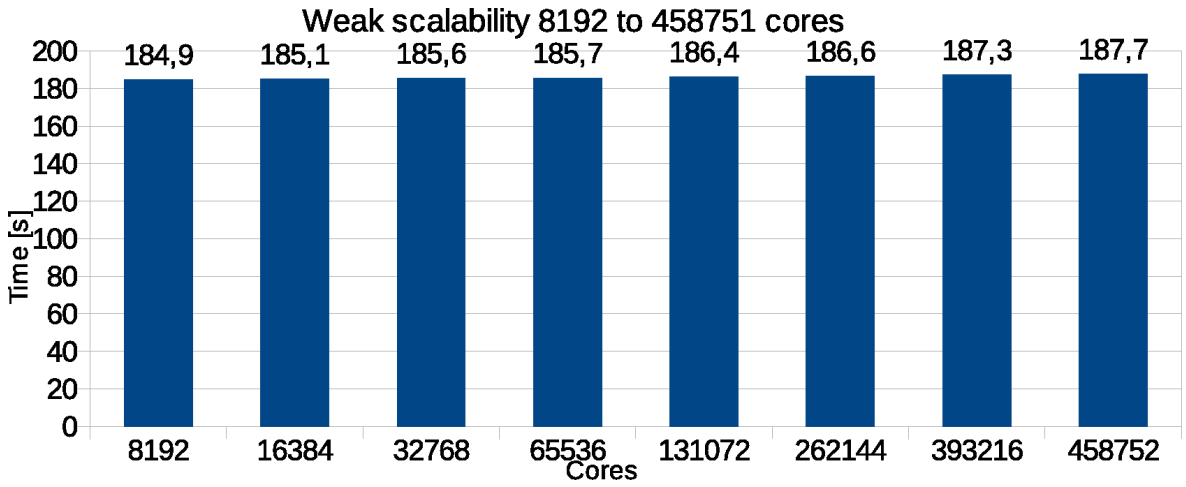

FE^{2}TI: FE^2 in 3D using FETI-DP on each RVE; heterogeneous hyperelasticity using Q1 finite elements macro, P2 finite elements micro; UMFPACK; 1.6M d.o.f. on each RVE; 512 subdomains for each RVE.

Scaling starts at 512 BG/Q nodes, i.e., the smallest reasonable macroscopic problem size (16 Gauß points).

Increasing the RVE sizes

FE^2 in 2D (Increasing RVE sizes)					
bg_size	**MPI-ranks**	**#RVEs**	**RVE-size**	**Total Dof =RVE-size $\times$ #RVEs**	**Time to Solution**
28 672	458 752	1 792	5 126 402	9 186 512 384	161.8s
28 672	458 752	1 792	7 380 482	13 225 823 744	248.2s
28 672	458 752	1 792	13 117 442	23 506 456 064	483.7s
28 672	458 752	1 792	20 492 802	36 723 101 184	817.1s

Increasing the RVE sizes; heterogeneous hyperelasticity;
P1 finite elements macro, P2 finite elements micro.

Using hardware threads by pure MPI (overcommit) on JUQUEEN

FE^2 in 3D (1x, 2x, 4x MPI overcommit)					
bg_size	**ranks per node**	**MPI-ranks**	**#RVEs**	**Time to Solution**	
28 672	16	458 752	896	215.4s	(100%)
28 672	32	917 504	1 792	266.5s	(81%)
28 672	64	1 835 008	3 584	522.1s	(41%)

"Weak scaling efficiency" using 16 / 32 / 64 MPI-ranks per node
(16 BG/Q cores with 64 hardware threads per node). FE^2 in 3D using FETI-DP on each RVE.

Fully Coupled FE^2 Computational Scale Bridging in 3D on Mira

FE^2 in 3D (Weak scaling)					
Cores	**MPI ranks**	**#RVEs**	**Total dof**	**Time to Solution**	**Eff.**
8 192	16 384	16	200M	914.34s	100.0%
16 384	32 768	32	401M	932.96s	98.0%
32 768	65 536	64	801M	932.48s	98.1%
65 536	131 072	128	1.6B	929.35s	98.4%
131 072	262 144	256	3.2B	935.26s	97.8%
262 144	524 288	512	6.4B	937.78s	97.5%
393 216	786 432	768	9.6B	941.25s	97.2%
524 288	1 048 576	1 024	12.8B	948.91s	96.4%
786 432	1 572 864	1 536	19.3B	943.81s	96.9%

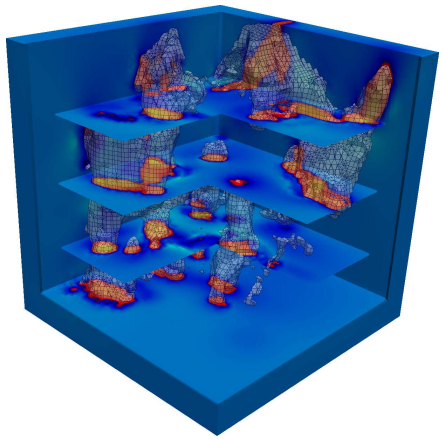

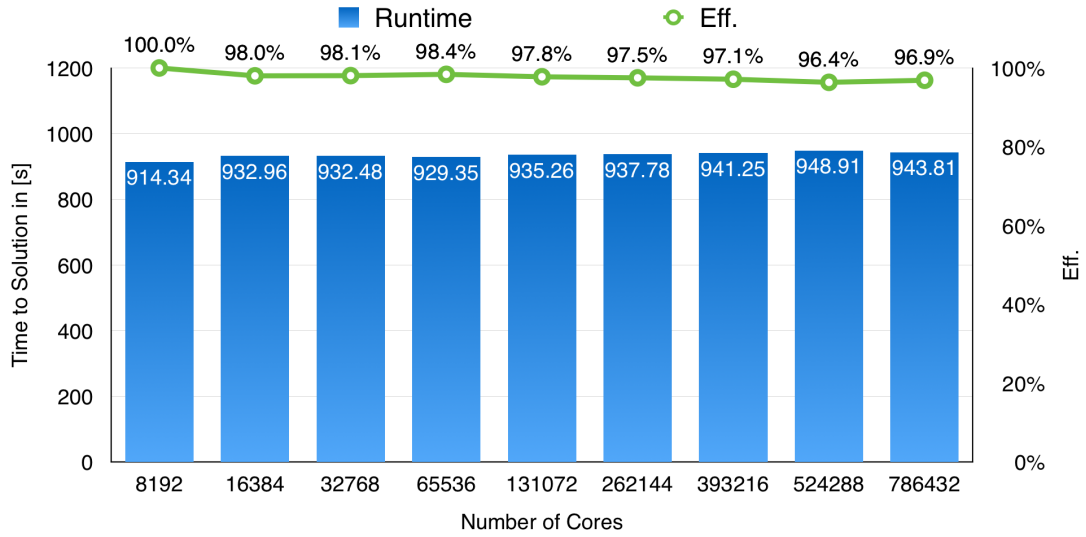

FE^{2}TI: FE^2 in 3D using FETI-DP on each RVE; heterogeneous hyperelasticity; MUMPS;
Q1 finite elements macro, P2 finite elements micro; 12.5M dof per RVE;
4 096 subdomains for each RVE; 1 024 MPI ranks and 512 cores per RVE.

FETI-DP for the Outer Loop

Classic:
Linearization first

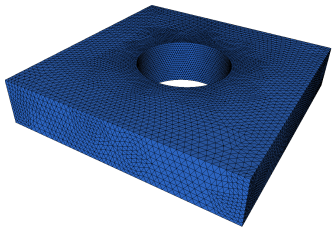

New:
Decomposition first

Decomposition

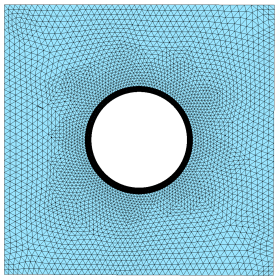

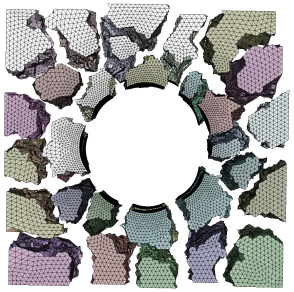

Linearization

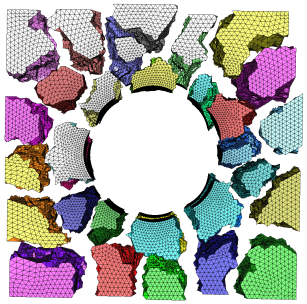

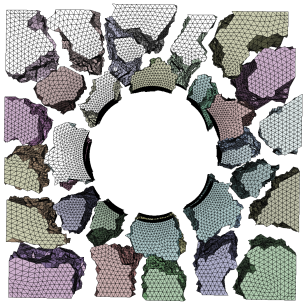

Higher Concurrency, Localization of Work

- Decompose the nonlinear problem into concurrent subproblems - **before linearization**.
- Combine parallel domain decomposition (FETI-DP) and parallel multigrid methods.

Nonlinear FETI-DP and Nonlinear BDDC: Klawonn, Lanser, Rheinbach (2012, 2013, 2014)
Related Methods: **ASPIN** (Cai, Keyes 2002; Cai, Keyes, and Marcinkowski 2002; Hwang, Cai 2005, 2007; Groß, Krause 2010,13);
Nonlinear Neumann-Neumann (Bordeu, Boucard, Gosselet 2009); **Nonlinear FETI-1** (Pebrel, Rey, Gosselet 2008);

Nonlinear FETI-DP Methods

Can we:

- Increase local work?
- Decrease need for synchronization and increase latency tolerance?

We consider the minimization of a nonlinear energy functional $J : V^h \to \boldsymbol{R}$,

$$\min_{\bar{u} \in V^h} J(\bar{u}), \tag{1}$$

where the following locality assumption has to hold:

Assumption 1. *There exist local energy functionals $J_i : W_i \to \boldsymbol{R}$, $i = 1, \ldots, N$, such that for $\bar{u} \in V^h$, the global energy functional can be represented as a sum of local energies*

$$J(\bar{u}) = \sum_{i=1}^{N} J_i(u_i),$$

where $u_i := R_i \bar{u}$.

Assumption 1 is, e.g., satisfied for our nonlinear model problems (p-Laplace, hyperelasticity).

Nonlinear FETI-DP Methods

FETI-DP methods are based on the **linear** master system

$$\begin{aligned} \widetilde{K}\tilde{u} + B^T\lambda &= \tilde{f} \\ B\tilde{u} &= 0. \end{aligned}$$

Nonlinear FETI-DP methods (for nonlin. problems) are based on the **nonlinear** master system

$$\boxed{\begin{aligned} \widetilde{K}(\tilde{u}) + B^T\lambda - \tilde{f} &= 0 \\ B\tilde{u} &= 0. \end{aligned}}$$

Assumption. *Let U be an open neighborhood of $\tilde{u}^*$. The function $\widetilde{K}$ is continuously differentiable in U. The derivative $D\widetilde{K}(\tilde{u}^*)$ of $\widetilde{K}$ in $\tilde{u}^*$ is a regular matrix.*

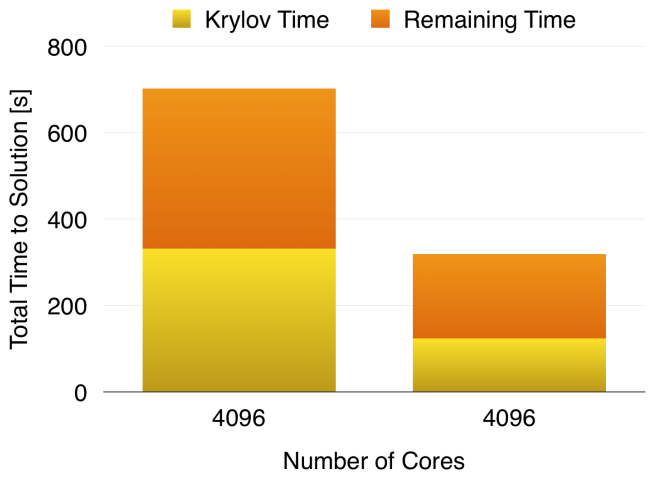

New Approach - An Inexact-Reduced-Nonlinear-FETI-DP (irNL-1)

Based on the nonlinear master system

$$\begin{array}{rcl} \widetilde{K}(\tilde{u}) + B^T\lambda - \tilde{f} &=& 0 \\ B\tilde{u} &=& 0. \end{array}$$

The Newton linearization with respect to $(\tilde{u}, \lambda)$ results in the linear system

$$\begin{pmatrix} D\widetilde{K}(\tilde{u}) & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} \delta\tilde{u} \\ \delta\lambda \end{pmatrix} = \begin{pmatrix} \widetilde{K}(\tilde{u}) + B^T\lambda - \tilde{f} \\ B\tilde{u} \end{pmatrix}.$$

where $\delta\tilde{u} = (\delta u_B^T, \delta\tilde{u}_\Pi^T)^T$, or in more details

$$\begin{bmatrix} (D\widetilde{K})_{BB} & (D\widetilde{K})_{\Pi B}^T & B_B^T \\ (D\widetilde{K})_{\Pi B} & (D\widetilde{K})_{\Pi\Pi} & 0 \\ B_B & 0 & 0 \end{bmatrix} \begin{bmatrix} \delta u_B \\ \delta\tilde{u}_\Pi \\ \delta\lambda \end{bmatrix} = \begin{bmatrix} (\widetilde{K}(\tilde{u}))_B + B_B^T\lambda - f_B \\ (\widetilde{K}(\tilde{u}))_\Pi - \tilde{f}_\Pi \\ B\tilde{u} \end{bmatrix}.$$

(Exact) Nonlinear FETI-DP methods are considered in
Klawonn, Lanser, Rheinbach, SIAM J. Sci. Comput., April 2014.

Reduced System and Preconditioner (irNL-1) (Phase 2)

Klawonn, Rheinbach, Lanser (DD22 2013, GAMM 2014)

Elimination of δu_B yields

$$\begin{bmatrix} \widetilde{S}_{\Pi\Pi} & -(D\widetilde{K})_{\Pi B}(D\widetilde{K})_{BB}^{-1}B_B^T \\ -B_B(D\widetilde{K})_{BB}^{-1}(D\widetilde{K})_{\Pi B}^T & -B_B(D\widetilde{K})_{BB}^{-1}B_B^T \end{bmatrix} \begin{bmatrix} \delta\tilde{u}_\Pi \\ \delta\lambda \end{bmatrix} = \text{r.h.s} \qquad (2)$$

Elimination of $\delta\tilde{u}_\Pi$ yields **Nonlinear-FETI-DP-1 (NL-1)**: $F_{NL1}\delta\lambda = d_{NL1}$.

We instead solve (2) using the block-triangular preconditioner

$$\hat{\mathcal{B}}_{r,L}^{-1} := \begin{bmatrix} \hat{S}_{\Pi\Pi}^{-1} & 0 \\ -M^{-1}B_B(DK)_{BB}^{-1}(D\widetilde{K})_{\Pi B}^T\hat{S}_{\Pi\Pi}^{-1} & -M^{-1} \end{bmatrix}$$

as in **Klawonn, Rheinbach (IJNME 2007, ZAMM 2010)** for linear elasticity.

We use BoomerAMG (**Henson, Meier-Yang**) for $\hat{S}_{\Pi\Pi}^{-1}$ and the Dirichlet preconditioner M^{-1}.

Choosing initial values for irNL-1 (Phase 1)

A possible choice of an initial value $\tilde{u}^{(0)}$ can be obtained from solving the nonlinear problem

$$\widetilde{K}(\tilde{u}^{(0)}) = \tilde{f} - B^T \lambda^{(0)} \tag{3}$$

by some Newton type iteration for some given initial value $\lambda^{(0)}$. We solve local nonlinear subdomain problems which are only coupled in the primal unknowns.

- Only one-time communication over the interface to calculate $B^T \lambda^{(0)}$!
- Only a few Krylov iterations needed for the inexact solution with $S_{\Pi\Pi}$ in each Newton step using GMRES and BoomerAMG.
- Local (independent) problems associated with $(D\widetilde{K})_{BB}$ solved in parallel per Newton step.

Note that a second **Nonlinear-FETI-DP** method is possible. This method is characterized by **eliminate before linearize**. Note that **nonlinear BDDC** methods are also possible.

Both, **Nonlinear FETI-DP-2** and **Nonlinear BDDC**, are also considered in **Klawonn, Lanser, and Rheinbach (SIAM J. Sci. Comput., April 2014)**.

Weak Scaling on the RVE - irNL-1/AMG on the Vulcan BG/Q

Nonlinear Hyperelasticity in 2D with Stiff Inclusions

#Cores (=N)	d.o.f.	Σ Krylov It.	Total Krylov Time	Runtime	eff.
32	1 642 242	93	23.2s	**250.3s**	100%
128	6 561 282	107	27.2s	**256.0s**	98%
512	26 229 762	109	27.9s	**257.7s**	97%
2 048	104 888 322	109	28.4s	**258.7s**	97%
8 192	419 491 842	107	28.7s	**261.0s**	96%
32 768	1 677 844 482	105	29.3s	**266.6s**	94%
131 072	6 711 132 162	102	26.7s	**278.9s**	90%

Vulcan BlueGene/Q, Lawrence Livermore National Laboratory; 2D Neo-Hooke; stiff circular inclusions in soft material; piecewise quadratic fem. irNonlinear FETI-DP

PETSc (Argonne National Laboratory): **Balay, Brown, Buschelman, Gropp, Kaushik, Knepley, Curfman McInnes, Smith and Zhang**

BoomerAMG (Lawrence Livermore National Laboratory): **Henson and Meier-Yang**

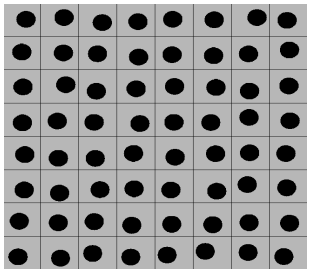

Domain decomposition into 64 subdomains; one circular inclusion for each subdomain

Weak Scaling on the RVE - irNL-1/AMG on the Mira BG/Q

Nonlinear Hyperelasticity in 2D with Stiff Inclusions

Mira BlueGene/Q Supercomputer (ANL)					
#Cores	**Problem Size**	**Phase 1 Time / Newton**	**Phase 2 Time / Newton / ΣGMRES**	**Total Time**	**Parallel Efficiency**
16	1.3M	158.7s / 4	205.3s / 3 / 83	**364.0s**	100%
64	5M	159.5s / 4	220.9s / 3 / 109	**380.4s**	96%
256	21M	160.1s / 4	238.9s / 3 / 135	**399.0s**	91%
1 024	84M	160.3s / 4	245.2s / 3 / 136	**405.5s**	90%
4 096	334M	182.0s / 4	246.5s / 3 / 110	**428.4s**	85%
16 384	1 337M	137.3s / **3**	249.0s / 3 / 110	**433.3s**	84%
65 536	5 348M	145.3s / **3**	180.3s / **2** / 85	**325.6s**	112%
262 144	21 393M	144.9s / **3**	177.5s / **2** / 83	**322.4s**	113%
524 288	42 786M	177.6s / **3**	200.2s / **2** / 82	**377.8s**	96%

Hybrid nonlinear FETI-DP-1/multigrid algorithm on the Mira supercomputer at Argonne National Laboratory; 2D heterogeneous hyperelastic material; Scalability $16 \rightarrow 524K$ for each Newton step: 67% (Phase 1), 68% (Phase 2); GMRES rtol 1e-10; Newton rtol 1e-8;

Weak Scaling: irNL-1/AMG on the Mira BG/Q

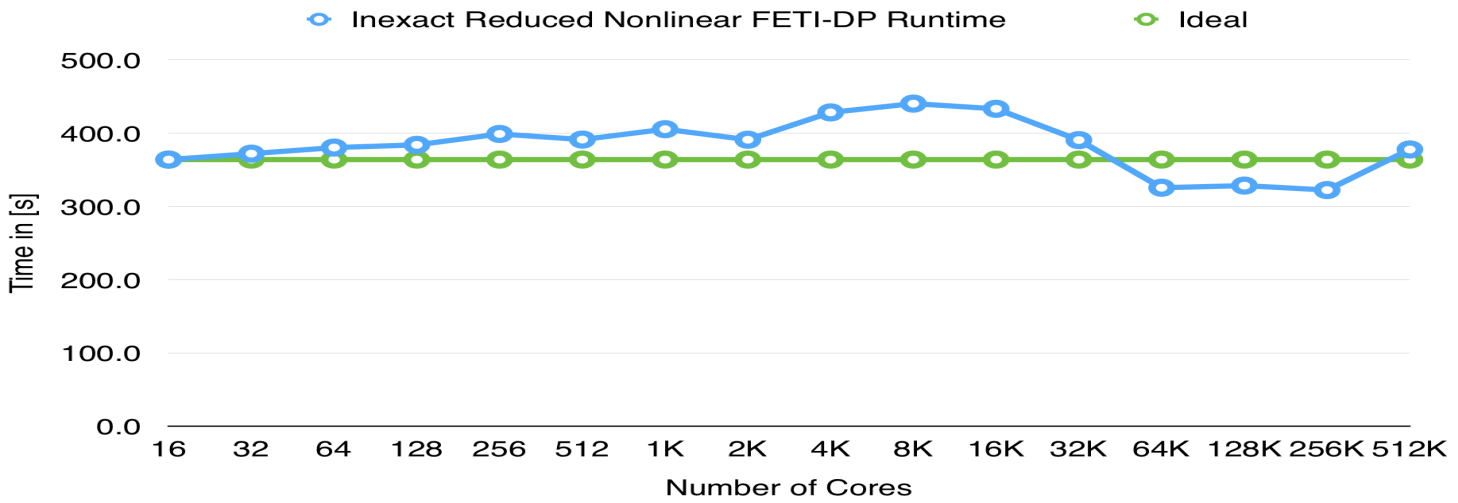

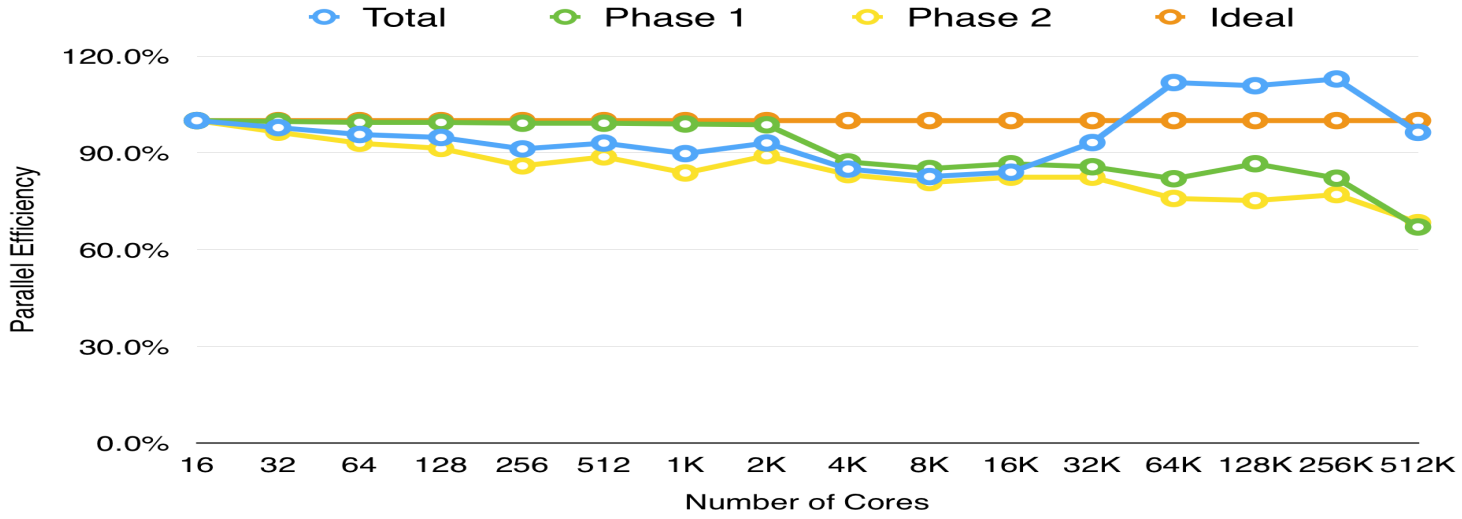

Strong Scaling on the RVE - irNL-1/AMG on the Vulcan BG/Q

Nonlinear Hyperelasticity in 2D with Stiff Inclusions

Vulcan BlueGene/Q (LLNL)						
Cores	**Subdomains**	**Problem Size**	**Total Execution Time**	**Actual Speedup**	**Ideal Speedup**	**Parallel Effic.**
1 024	131 072	419 471 361	3 365.1s	1.0	1	100%
2 048	131 072	419 471 361	1 726.4s	1.9	2	97%
4 096	131 072	419 471 361	868.0s	3.9	4	97%
8 192	131 072	419 471 361	453.5s	7.4	8	93%
16 384	131 072	419 471 361	231.4s	14.6	16	91%
32 768	131 072	419 471 361	119.8s	28.1	32	88%
65 536	131 072	419 471 361	64.3s	51.6	64	81%
131 072	131 072	419 471 361	41.7s	80.6	128	63%

Hybrid nonlinear FETI-DP/multigrid (irNL-1 with BoomerAMG) algorithm on the Vulcan supercomputer at Lawrence Livermore National Laboratory; hyperelastic material;

irNL-1/AMG on the Vulcan BG/Q

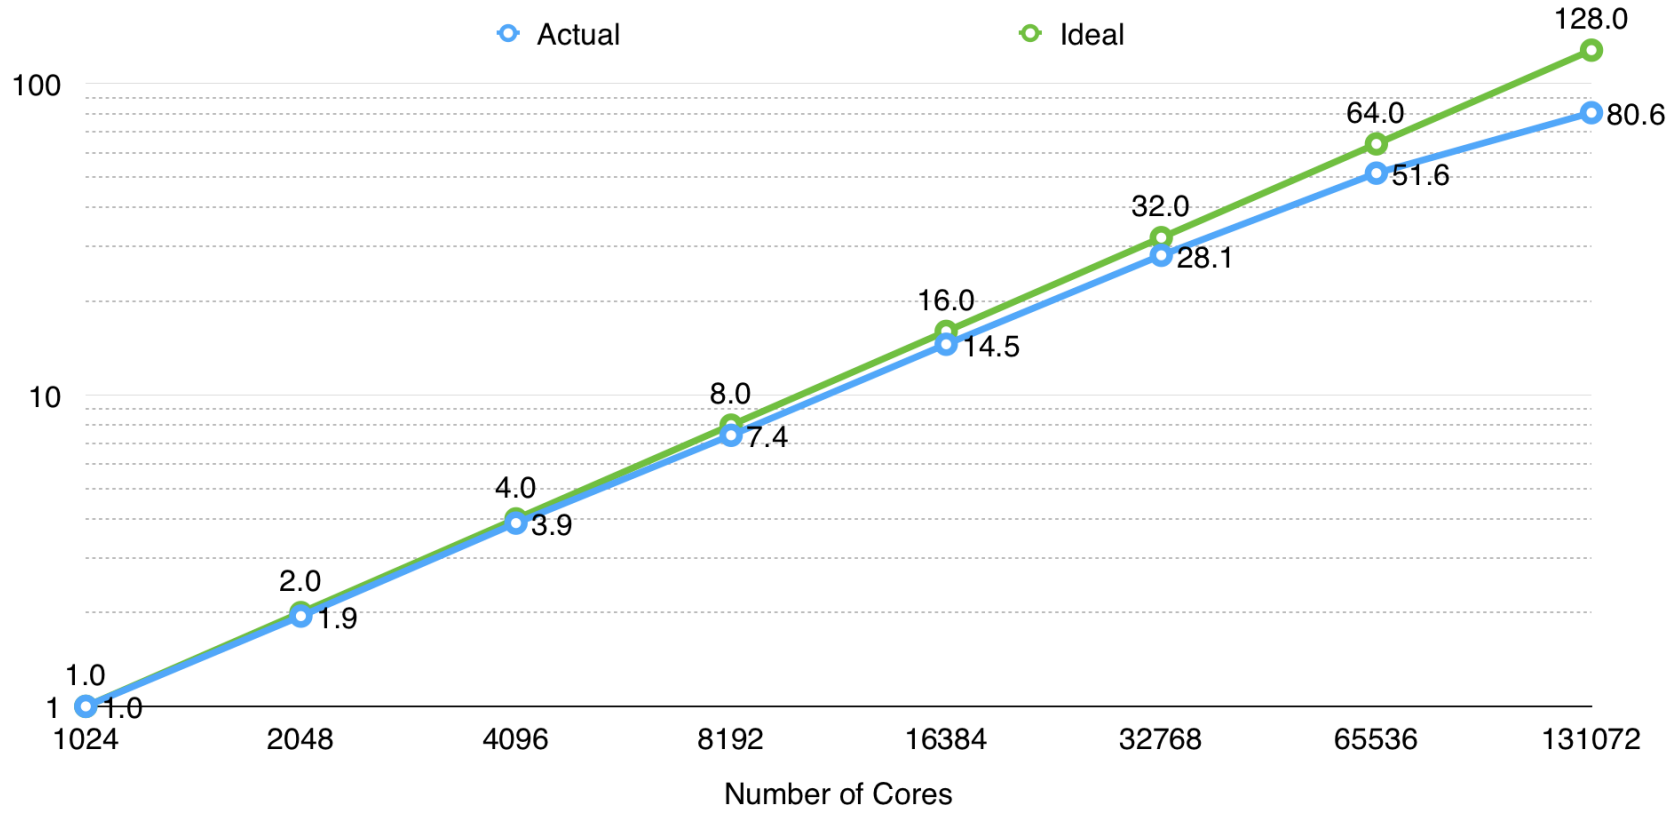

Strong Scaling on Vulcan: Visualization of the speedup.

Conclusions

- Nonlinear-FETI-DP localizes work and reduces communication
- Scalability for FE^2 with Newton-Krylov-FETI-DP on each micscroscopic RVE for up to 486K cores on the JUQUEEN BG/Q (FZ Jülich) and for up to 786K cores on the Mira BG/Q (ANL)
- Scalability for irNL-1/AMG for up to 131K cores on the Vulcan BG/Q (LLNL) and for up to 524K cores on the Mira BG/Q (ANL)

Support is gratefully acknowledged by Deutsche Forschungsgemeinschaft (DFG) within the priority program SPP 1648 Software for Exascale Computing.

The authors gratefully acknowledge the Gauss Centre for Supercomputing (GCS) for providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS share of the supercomputer **JUQUEEN** at Jülich Supercomputing Centre (JSC).

This research used resources, i.e., **Mira**, of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.

The authors gratefully acknowledge the use of **Vulcan** at Lawrence Livermore National Laboratory.

PETSc, happy birthday!