Partitioned Adaptive Parallel Integrators for Coupled Stiff Systems

Philipp Birken, Peter Meisrimel, Azahar Monge

Lund University, Sweden

Integrating the Integrators for Nonlinear Evolution Equations: from Analysis to Numerical Methods, High-Performance-Computing and Applications
Banff, December 3rd, 2018

Outline

1. 3M: Multiphysics, Multicore, Multirate
2. Thermal FSI: Neumann-Neumann waveform iterations
3. Asynchronous waveform iterations

Outline

Example: Vegetation feedback

- Complex feedbacks between precipitation, temperature, vegetation
- Climate models incorporate PDE models for air- and waterflow
- Add differential equations for clouds, CO2 sources,...
- Vegetation growth based on PDE as well
- Coupling of different models implemented in different codes

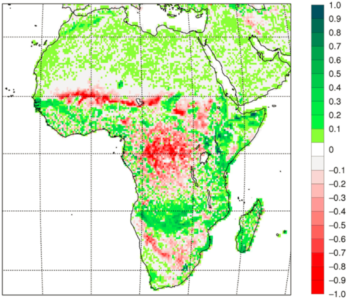

Figure: Output of RCA-GUESS: Leaf area index change

Example: Steel forging

Inductive heating

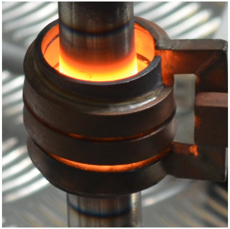

Thermo-mechanical forming

Local air-cooling

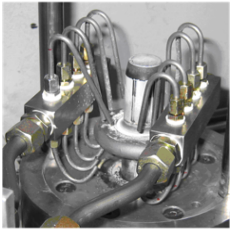

- Entire process chain is characterized by the exchange of energy
- Crystalline structure of steel influenced by heat

Thermal interaction between air and steel needs to be modelled

Applications of thermal FSI

- Engines (rocket, car,...)
- Heating of reentry vehicles in space flight
- Turbine blade cooling
- Thermal anti-icing of airplanes
- Generally cooling systems

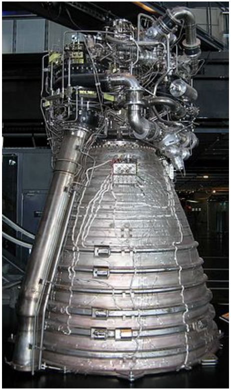

Figure: Vulcain engine for Ariane 5; CC-by-sa 3.0, Pline, Wikimedia Commons

Goal: Solve this fast without too much pain

What do we want?

- Time adaptivity
- High order in time
- Reuse of existing codes (partitioned approach)
- Parallel execution of coupled codes
- Loadbalancing
- Different time steps in different models
- Fast solvers for equation systems
- Should be usable and robust for large class of models

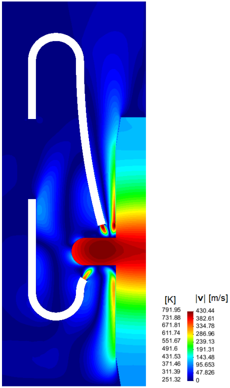

Figure:
Navier-Stokes (FV)
+ nonl. heat (FE)

For steel forging problem: Dirichlet-Neumann works well, but is not parallel, Monge, B., Comp. Mech. 18

Influence of Software

Good coders write great code, great coders use good coders code

- Physics and mathematical models is not all
- Software influences design of numerical methods
- Option 1: Full access to code and willingness to edit it
- Option 2: Only access to code of certain functions, allowed to call some specific ones
- Option 3: No access, allowed to call specific interface functions
- Option 4: No access, Allowed only to call main function

From ODE world: Waveform Iteration

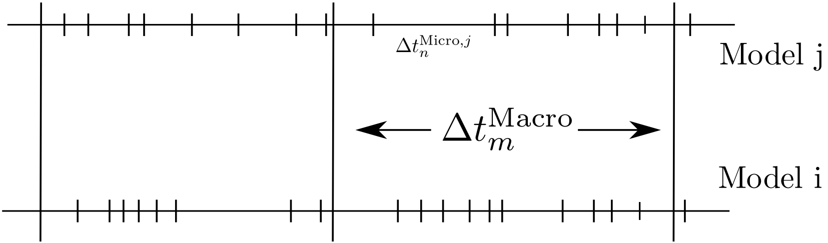

- Synchronize models at macro steps
- In between, subsolvers run with their own time step
- Need to use information from other model at microsteps
- Is black box, but we require possibility to repeat a macrostep

Outline

Model Problem: Coupled heat equations

- Coupled PDEs: Nonoverlapping Domain Decomposition

$$
\begin{aligned}
\alpha_m \frac{\partial u_m(\mathbf{x},t)}{\partial t} - \nabla \cdot (\lambda_m \nabla u_m(\mathbf{x},t)) = 0, \\
t \in [t_0, t_f], \quad \mathbf{x} \in \Omega_m \subset \mathbb{R}^d, \ m = 1,2 \\
u_m(\mathbf{x},t) = 0, \quad \mathbf{x} \in \partial\Omega_m \setminus \Gamma \\
u_1(\mathbf{x},t) = u_2(\mathbf{x},t), \quad \mathbf{x} \in \Gamma \\
\lambda_2 \frac{\partial u_2(\mathbf{x},t)}{\partial \mathbf{n}_2} = -\lambda_1 \frac{\partial u_1(\mathbf{x},t)}{\partial \mathbf{n}_1}, \quad \mathbf{x} \in \Gamma \\
u_m(\mathbf{x},0) = g_m(\mathbf{x}) \quad \mathbf{x} \in \Omega_m
\end{aligned}
$$

- In the weak sense equivalent to

$$
\begin{aligned}
\alpha(x) \frac{\partial u_m(\mathbf{x},t)}{\partial t} - \nabla \cdot (\lambda(x) \nabla u_m(\mathbf{x},t)) = 0 \qquad & \mathbf{x} \in \Omega \\
u(\mathbf{x},t) = 0, \quad \mathbf{x} \in \partial\Omega, \ u(\mathbf{x},0) = g_m(\mathbf{x}), \qquad & \mathbf{x} \in \partial\Omega
\end{aligned}
$$

Neumann-Neumann multirate waveform method

Neumann-Neumann method

Given initial guess for values at boundary, iterate

1. Solve Dirichlet problems in both domains with that data
2. Solve Neumann problems in both domains with derivatives from above
3. Get new Dirichlet data, do relaxation step

- Can be done in parallel!
- Gander, Kwok, Mandal, ETNA '16: Fully continuous version with all material parameters one
- Prove optimal relaxation parameter of $1/4$
- With this, exact solution at interface after one step!!

Neumann-Neumann waveform relaxation (NNWR)

1. *Dirichlet* :
$$\begin{cases} \alpha_m \frac{\partial u_m^{k+1}(\mathbf{x},t)}{\partial t} - \nabla \cdot (\lambda_m \nabla u_m^{k+1}(\mathbf{x},t)) = 0, & \mathbf{x} \in \Omega_m, \\ u_m^{k+1}(\mathbf{x},t) = 0, & \mathbf{x} \in \partial\Omega_m \setminus \Gamma, \\ u_m^{k+1}(\mathbf{x},t) = g^k(\mathbf{x},t), & \mathbf{x} \in \Gamma, \\ u_m^{k+1}(\mathbf{x},0) = u_1^0(\mathbf{x}), & \mathbf{x} \in \Omega_m. \end{cases}$$

2. *Neumann* :
$$\begin{cases} \alpha_m \frac{\partial \psi_m^{k+1}(\mathbf{x},t)}{\partial t} - \nabla \cdot (\lambda_m \nabla \psi_m^{k+1}(\mathbf{x},t)) = 0, & \mathbf{x} \in \Omega_m, \\ \psi_m^{k+1}(\mathbf{x},t) = 0, & \mathbf{x} \in \partial\Omega_m \setminus \Gamma, \\ \lambda_m \frac{\partial \psi_m^{k+1}(\mathbf{x},t)}{\partial \mathbf{n}_1} = \lambda_1 \frac{\partial u_1^{k+1}(\mathbf{x},t)}{\partial \mathbf{n}_1} + \lambda_2 \frac{\partial u_2^{k+1}(\mathbf{x},t)}{\partial \mathbf{n}_2}, & \mathbf{x} \in \Gamma, \\ \psi_m^{k+1}(\mathbf{x},0) = 0, & \mathbf{x} \in \Omega_m. \end{cases}$$

3. *Relaxation* : $g^{k+1}(\mathbf{x},t) = g^k(\mathbf{x},t) - \Theta(\psi_1^{k+1}(\mathbf{x},t) + \psi_2^{k+1}(\mathbf{x},t)), \quad \mathbf{x} \in \Gamma.$

Going discrete: Multirate

- *Nonmatching time grids* at the interface, linear interpolation
- *Time discretization*: Implicit Euler and SDIRK2 - Order is acchieved

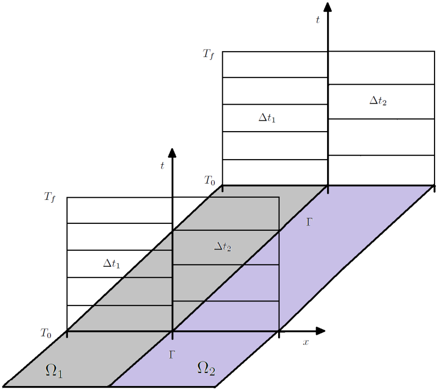

Multirate 1D solution using NNWR algorithm

Iteration 1

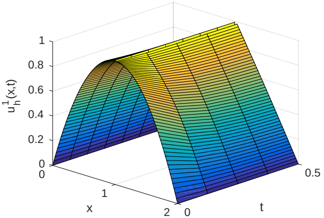

Iteration 2

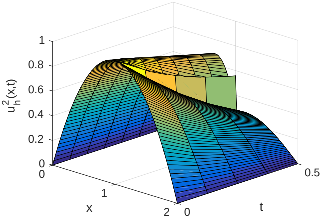

Iteration 3

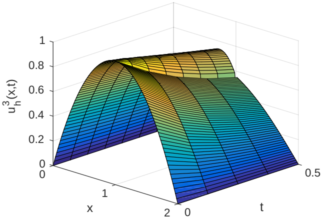

Iteration 20

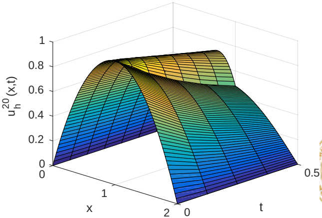

Relaxation parameter

Q: How to choose the relaxation parameter Θ?

- Write iteration in terms of interface unknowns at final time, $\mathbf{u}_\Gamma(T_f)$:

$$\mathbf{u}_\Gamma^{k+1,T_f} = \Sigma(\Theta)\mathbf{u}_\Gamma^{k,T_f} + \Psi$$

- Do only one time step only
- Find Θ that minimizes the spectral radius of $\Sigma(\Theta)$
- $\Theta_{opt} = \rho(2 + \mathbf{S}^{(1)^{-1}}\mathbf{S}^{(2)} + \mathbf{S}^{(2)^{-1}}\mathbf{S}^{(1)})$
- Give exact formula for model discretization
- *Space discretization*: equidistant FE/FE in 1D
- *Matching space grid* at the interface, unknowns on interface
- *Time integration*: nonmultirate Implicit Euler or SDIRK2

One-dimensional analysis, several pages later

Use eigendecomposition of tridiagonal Toeplitz matrices $\mathbf{M}/\Delta t + \mathbf{A}$ to get

$$\Theta_{opt} = \left(2 + \frac{(6\Delta x(\alpha_2 \Delta x^2 + 3\lambda_2 \Delta t) - (\alpha_2 \Delta x^2 - 6\lambda_2 \Delta t)^2 s_2)}{(6\Delta x(\alpha_1 \Delta x^2 + 3\lambda_1 \Delta t) - (\alpha_1 \Delta x^2 - 6\lambda_1 \Delta t)^2 s_1)} + \frac{(6\Delta x(\alpha_1 \Delta x^2 + 3\lambda_1 \Delta t) - (\alpha_1 \Delta x^2 - 6\lambda_1 \Delta t)^2 s_1)}{(6\Delta x(\alpha_2 \Delta x^2 + 3\lambda_2 \Delta t) - (\alpha_2 \Delta x^2 - 6\lambda_2 \Delta t)^2 s_2)}\right)^{-1}.$$

with

$$s_m = \sum_{i=1}^{N} \frac{3\Delta x^2 \sin^2(i\pi\Delta x)}{2\alpha_m \Delta x^2 + 6\lambda_m \Delta t + (\alpha_m \Delta x^2 - 6\lambda_m \Delta t)\cos(i\pi\Delta x)}.$$

Asymptotics: 1D FEM/FEM

- Temporal limit of Θ_{opt}: $\Theta_{\{c\to 0\}} = \frac{\alpha_1\alpha_2}{(\alpha_1+\alpha_2)^2}$.
- Spatial limit of Θ_{opt}: $\Theta_{\{c\to \infty\}} = \frac{\lambda_1\lambda_2}{(\lambda_1+\lambda_2)^2}$.

Convergence Rates, observed in 2D

Air-Steel

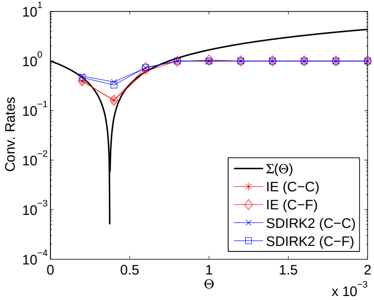

Air-Water

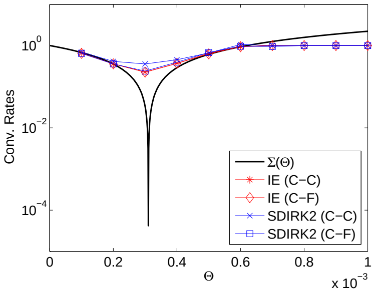

Simplifications in analysis don't affect location of Θ_{opt}.

And now once more, with time adaptivity!

- Add two step size controllers on the Dirichlet problems
- Problem: Θ_{opt} depends on α_m, λ_m, $m = 1, 2$, Δx and Δt
- Initial iteration: $\Theta = (\Theta_{\{c \to 0\}} + \Theta_{\{c \to \infty\}})/2$
- Then take average of time steps on each domain

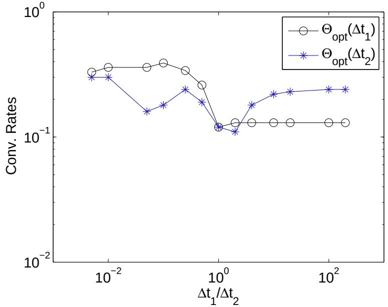

- If $\overline{\Delta t_1} < \overline{\Delta t_2}$, $\Theta_{k+1} = \Theta_{opt}(\overline{\Delta t_2})$; else, $\Theta_{k+1} = \Theta_{opt}(\overline{\Delta t_1})$

Numerical results: Air-Steel

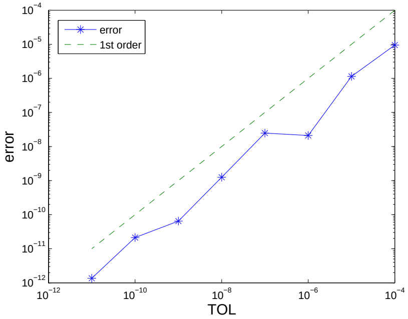

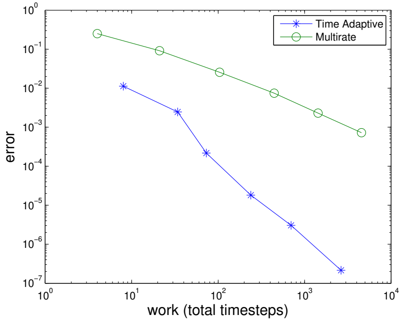

- Multirate uses respective minimal timestep from adaptive method
- See Monge, B.: A time adaptive Neumann-Neumann waveform relaxation method for thermal fluid-structure interaction, DD25, submitted

Where are we?

- Time adaptivity
- High order in time
- Parallel execution of coupled codes
- Loadbalancing to be done
- Different time steps in different models
- Fast solvers for equation systems
- Is being implemented in open source coupling software framework PreCICE (with B. Rueth, B. Ueckermann, M. Mehl)
- Robustness problematic: Relaxation parameter sensitive. Now look at Dirichlet-Neumann again with waveform relaxation with pipeline implementation
- Specific to thermal FSI

More at: Monge, B., *A multirate Neumann-Neumann waveform relaxation method for heterogeneous coupled heat equations*, SISC, submitted, arXiv:1805.04336

Outline

Discretely asynchronous

- Consider one macro step and solution operators ϕ_x, ϕ_y
- Use adaptive numerical method for both problems
- Gauß-Seidel waveform iteration not parallel
- Jacobi waveform iteration parallel, but in fact double the work
- Use continuous interpolations $\mathbf{x}^{[t_n,t_{n+1}],k}(t)$, $\mathbf{y}^{[t_n,t_{n+1}],k}(t)$ in time.
- Update interpolations after every microstep

$$\mathbf{x}^{m+1} = \mathbf{x}^m + \Phi_x(\mathbf{x}^m, \mathbf{y}^{[t_m,t_{m+1}],k}(t))$$

$$\mathbf{y}^{m+1} = \mathbf{y}^m + \Phi_y(\mathbf{y}^m, \mathbf{x}^{[t_m,t_{m+1}],k}(t))$$

- Is multiphysics, multirate, parallel and adaptive
- See also Frommer, Szyld 2000

Preliminary results using MPI Window, Put, Get

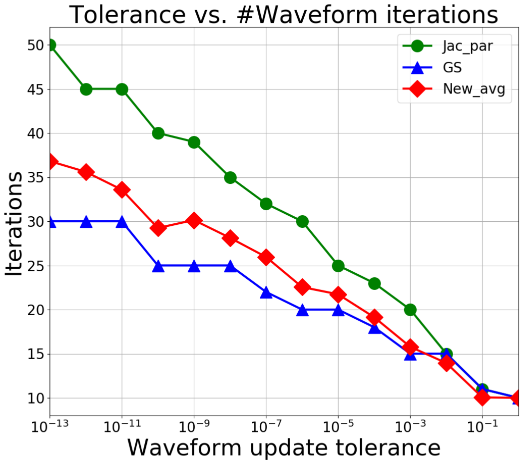

Figure: Averaged iterations for 1000 runs for 2-component linear system

Where are we with this?

- Time adaptivity
- High order in time
- Parallel execution of coupled codes
- Loadbalancing to be done
- Different time steps in different models
- Fast solvers for equation systems: Not clear
- Robustness unclear