XII GAMM Workshop on
Applied an Numerical Linear Algebra
September 02-05 2012, Chateau Liblice (CZ)

Towards a GPU Add-On for the M.E.S.S.

A GPU accelerated inexact Newton iteration for large sparse algebraic Riccati equations

Jens Saak (CSC)
joint work with
Pablo Ezzati (CeCal), Martin Köhler (CSC) and Alfredo Remón (HPCA)

Max Planck Institute for Dynamics of Complex Technical Systems,
Computational Methods in Systems and Control Theory (CSC)

Universitat Jaume I - Spain, High Performance Computing & Architectures (HPCA) group

Universidad de la Repblica - Uruguay, Centro de Cálculo (CeCal)

MAX PLANCK INSTITUTE
FOR DYNAMICS OF COMPLEX
TECHNICAL SYSTEMS
MAGDEBURG

Outline

(Inexact) Newton Methods for AREs

Basic Concepts

[Kleinman '68, Feitzinger/Hylla/Sachs '09]

Consider $\mathfrak{R}(X) := C^T C + A^T XE + E^T XA - E^T XBB^T XE = 0$

Inexact Newton's Iteration for the ARE

$$\mathfrak{R}'|_{X_\ell}(N_\ell) + \mathfrak{R}(X_\ell) = R_\ell, \quad X_{\ell+1} = X_\ell + N_\ell, \quad \ell = 0, 1, \ldots$$

i.e., in every Newton step (approximately) solve a

Lyapunov Equation

$$(A - BB^T X_\ell)^T N_\ell E + E^T N_\ell (A - BB^T X_\ell) = -\mathfrak{R}(X_\ell) + R_\ell.$$

(Inexact) Newton Methods for AREs

Basic Concepts

[Kleinman '68, Feitzinger/Hylla/Sachs '09]

Consider $\mathfrak{R}(X) := C^T C + A^T X E + E^T X A - E^T X B B^T X E = 0$

Inexact Kleinman's Iteration for the ARE

$$\mathfrak{R}'|_{X_\ell}(X_{\ell+1}) - \mathfrak{R}'|_{X_\ell}(X_\ell) + \mathfrak{R}(X_\ell) = R_\ell, \qquad \ell = 0, 1, \ldots$$

i.e., in every Newton step (approximately) solve a

Lyapunov Equation

$$(A - BB^T X_\ell)^T X_{\ell+1} E + E^T X_{\ell+1} (A - BB^T X_\ell) = -C^T C - X_\ell B B^T X_\ell + R_\ell.$$

(Inexact) Newton Methods for AREs

Basic Concepts

[Kleinman '68, Feitzinger/Hylla/Sachs '09]

Consider $$\mathfrak{R}(X) := C^T C + A^T XE + E^T XA - E^T XBB^T XE = 0$$

Inexact Kleinman's Iteration for the ARE

$$\mathfrak{R}'|_{X_\ell}(X_{\ell+1}) - \mathfrak{R}'|_{X_\ell}(X_\ell) + \mathfrak{R}(X_\ell) = R_\ell, \qquad \ell = 0, 1, \ldots$$

i.e., in every Newton step (approximately) solve a

Lyapunov Equation

$$F_\ell^T X_{\ell+1} E + E^T X_{\ell+1} F_\ell = -G_\ell G_\ell^T + R_\ell.$$

(Inexact) Newton Methods for AREs

Convergence Result [Kleinman '68, Lancaster/Rodman '95, Feitzinger/Hylla/Sachs '09]

Theorem

Let Assumption 1 hold,

$$0 \leq R_\ell \leq C^T C \quad \text{and} \quad 0 \leq R_\ell \leq N_\ell B B^T N_\ell.$$

Then the iterates defined by

$$F_\ell^T X_{\ell+1} + X_{\ell+1} F_\ell = -G_\ell G_\ell^T + R_\ell,$$

converge to the unique symmetric matrix X_∞, such that

- $\mathfrak{R}(X_\infty) = 0$
- *and $A - BB^T X_\infty$ is stable.*

Furthermore the convergence is quadratic and monotone with

$$0 \leq X_\infty \leq \cdots \leq X_{k+1} \leq X_k \leq \cdots \leq X_1.$$

(Inexact) Newton Methods for AREs

Convergence Result (Remarks)

Weaker Condition [Hylla '10]

Replacing

$$R_\ell \leq C^T C$$

by

$$R_\ell \leq C^T C + X_\ell B B^T X_\ell$$

keeps the iteration well defined.

(Inexact) Newton Methods for AREs

Convergence Result (Remarks)

Weaker Condition [Hylla '10]

Replacing

$$R_\ell \leq C^T C$$

by

$$R_\ell \leq C^T C + X_\ell B B^T X_\ell$$

keeps the iteration well defined.

Large Scale Difficulty

None of the conditions

- $R_\ell \leq C^T C$,
- $R_\ell \leq C^T C + X_\ell B B^T X_\ell$,
- $0 \leq R_\ell \leq N_\ell B B^T N_\ell$,

can be tested efficiently in large scale applications.

(Inexact) Newton Methods for AREs

Convergence Result (Remarks)

Weaker Condition [Hylla '10]

Replacing

$$R_\ell \leq C^T C$$

by

$$R_\ell \leq C^T C + X_\ell B B^T X_\ell$$

keeps the iteration well defined.

Large Scale Difficulty

None of the conditions

- $R_\ell \leq C^T C$,
- $R_\ell \leq C^T C + X_\ell B B^T X_\ell$,
- $0 \leq R_\ell \leq N_\ell B B^T N_\ell$,

can be tested efficiently in large scale applications.

Crucial for the proof. $\Rightarrow$ Lyapunov solver needs to ensure this.

Solving Large Lyapunov Equations

(G-)LRCF-ADI

e.g.,[Benner/Li/Penzl '08, S. '09]

Consider $FXE^T + EXF^T = -GG^T \qquad E, F \in \mathbb{R}^{n\times n}, G \in \mathbb{R}^{n\times p}$

Task Find $Z \in \mathbb{R}^{n,n_Z}$, such that $n_Z \ll n$ and $X \approx ZZ^T$.

Algorithm

$$V_1 = (F + p_1 E)^{-1} G, \qquad Z_1 = \sqrt{-2\,\mathrm{Re}\,(p_1)}V_1,$$

$$V_i = \left[I - (p_i + \overline{p_{i-1}})(F + p_i E)^{-1}\right] EV_{i-1}, \quad Z_i = [Z_{i-1} \quad \sqrt{-2\,\mathrm{Re}\,(p_i)}V_i].$$

For certain shift parameters $\{p_1, \ldots, p_J\} \subset \mathbb{C}_{<0}$.

Stop if

- $\|V_i V_i^H\|$ is small, or
- $\|FZ_i Z_i^T E^T + EZ_i Z_i^T F^T + GG^T\|$ is small.

Solving Large Lyapunov Equations

(G-)LRCF-ADI

e.g.,[Benner/Li/Penzl '08, S. '09]

Consider $FXE^T + EXF^T = -GG^T \qquad E, F \in \mathbb{R}^{n\times n}, G \in \mathbb{R}^{n\times p}$

Task Find $Z \in \mathbb{R}^{n,n_Z}$, such that $n_Z \ll n$ and $X \approx ZZ^T$.

Algorithm

$$V_1 = (F + p_1 E)^{-1} G, \qquad Z_1 = \sqrt{-2\,\mathrm{Re}\,(p_1)}\,V_1,$$

$$V_i = \left[I - (p_i + \overline{p_{i-1}})(F + p_i E)^{-1}\right] EV_{i-1}, \qquad Z_i = \left[Z_{i-1} \quad \sqrt{-2\,\mathrm{Re}\,(p_i)}\,V_i\right].$$

For certain shift parameters $\{p_1, \ldots, p_J\} \subset \mathbb{C}_{<0}$.

Stop if

- $\|V_i V_i^H\|$ is small, or
- $\|FZ_iZ_i^H E^T + EZ_iZ_i^H F^T + GG^T\|$ is small.

$Z_0 = [\,] \Rightarrow X_0 = 0$

$\Rightarrow R_0 = GG^T \geq 0$

[Hylla '10]: Then $\forall i: R_i \geq 0$,

and $\exists i_0 \quad \forall i \geq i_0 \quad R_i \leq C^T C$.

(Inexact) Newton Methods for AREs

Low-Rank Newton-ADI (LRCF-NM) for AREs [Kleinman '68, Feitzinger/Hylla/Sachs '09]

Consider $\mathfrak{R}(X) := C^TC + A^TXE + E^TXA - E^TXBB^TXE = 0$

Inexact Kleinman's Iteration for the ARE

$$\mathfrak{R}'|_{X_\ell}(X_{\ell+1}) - \mathfrak{R}'|_{X_\ell}(X_\ell) + \mathfrak{R}(X_\ell) = R_\ell, \qquad \ell = 0, 1, \dots$$

i.e., in every Newton step (approximately) solve a

Lyapunov Equation

$$(A - BB^TX_\ell)^TX_{\ell+1}E + E^TX_{\ell+1}(A - BB^TX_\ell) = -C^TC - X_\ell BB^TX_\ell + R_\ell.$$

LRCF-NM for the ARE

Low-Rank Newton-ADI (LRCF-NM) for AREs

Factored Newton-Kleinman Iteration [Benner/Li/Penzl '99/'08]

$F_\ell^T = A - BB^T X_\ell E =: A - BK_\ell$ is "sparse + low rank"

$G_\ell = [C^T \; K_\ell^T]$ is low rank factor

Find low rank factor $Z_\ell \in \mathbb{R}^{n,n_Z}$, where $n_Z \ll n$ and $X_\ell = Z_\ell Z_\ell^T$.

LRCF-NM for the ARE

Low-Rank Newton-ADI (LRCF-NM) for AREs

Factored Newton-Kleinman Iteration [Benner/Li/Penzl '99/'08]

$$F_\ell^T = A - BB^T X_\ell E =: A - BK_\ell \quad \text{is "sparse + low rank"}$$
$$G_\ell = [C^T \ K_\ell^T] \quad \text{is low rank factor}$$

- apply LRCF-ADI in every Newton step
- exploit structure of F_ℓ using Sherman-Morrison-Woodbury formula

$$(A - BK_\ell + p_k^{(\ell)}E)^{-1} =$$
$$(I_n + (A + p_k^{(\ell)}E)^{-1}B(I_m - K_\ell(A + p_k^{(\ell)}E)^{-1}B)^{-1}K_\ell)(A + p_k^{(\ell)}E)^{-1}$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

Main Problem:

How can we ensure quadratic convergence without checking

$$0 \leq R_\ell \leq C^T C \quad \text{and} \quad 0 \leq R_\ell \leq N_\ell B B^T N_\ell?$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

Main Problem:

How can we ensure quadratic convergence without checking

$$0 \leq R_\ell \leq C^T C \quad \text{and} \quad 0 \leq R_\ell \leq N_\ell B B^T N_\ell?$$

Due to the quadratic nature of $\mathfrak{R}(.)$ we have

$$\mathfrak{R}(Y) = \mathfrak{R}(X) + \mathfrak{R}'|_X(Y - X) + \frac{1}{2}\mathfrak{R}''|_X(Y - X, Y - X).$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

Main Problem:

How can we ensure quadratic convergence without checking

$$0 \leq R_\ell \leq C^T C \quad \text{and} \quad 0 \leq R_\ell \leq N_\ell B B^T N_\ell?$$

Due to the quadratic nature of $\mathfrak{R}(.)$ we have

$$\mathfrak{R}(Y) = \mathfrak{R}(X) + \mathfrak{R}'|_X(Y - X) + \frac{1}{2}\mathfrak{R}''|_X(Y - X, Y - X).$$

Recall the Inexact Kleinman step:

$$R_\ell = \mathfrak{R}'|_{X_\ell}(X_{\ell+1}) - \mathfrak{R}'|_{X_\ell}(X_\ell) + \mathfrak{R}(X_\ell)$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

Main Problem:

How can we ensure quadratic convergence without checking

$$0 \leq R_\ell \leq C^T C \quad \text{and} \quad 0 \leq R_\ell \leq N_\ell B B^T N_\ell?$$

Due to the quadratic nature of $\mathfrak{R}(.)$ we have

$$\mathfrak{R}(Y) = \mathfrak{R}(X) + \mathfrak{R}'|_X(Y - X) + \frac{1}{2}\mathfrak{R}''|_X(Y - X, Y - X).$$

Recall the Inexact Kleinman step:

$$R_\ell = \mathfrak{R}'|_{X_\ell}(X_{\ell+1}) - \mathfrak{R}'|_{X_\ell}(X_\ell) + \mathfrak{R}(X_\ell) = \mathfrak{R}'|_{X_\ell}(X_{\ell+1} - X_\ell) + \mathfrak{R}(X_\ell),$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

Main Problem:

How can we ensure quadratic convergence without checking

$$0 \leq R_\ell \leq C^T C \quad \text{and} \quad 0 \leq R_\ell \leq N_\ell B B^T N_\ell?$$

Due to the quadratic nature of $\mathfrak{R}(.)$ we have

$$\mathfrak{R}(Y) = \mathfrak{R}(X) + \mathfrak{R}'|_X(Y-X) + \frac{1}{2}\mathfrak{R}''|_X(Y-X, Y-X).$$

Recall the Inexact Kleinman step:

$$R_\ell = \mathfrak{R}'|_{X_\ell}(X_{\ell+1}) - \mathfrak{R}'|_{X_\ell}(X_\ell) + \mathfrak{R}(X_\ell) = \mathfrak{R}'|_{X_\ell}(X_{\ell+1} - X_\ell) + \mathfrak{R}(X_\ell),$$

and thus

$$\mathfrak{R}(X_{\ell+1}) = R_\ell + \frac{1}{2}\mathfrak{R}''|_{X_\ell}(X_{\ell+1} - X_\ell, X_{\ell+1} - X_\ell).$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

Main Problem:

How can we ensure quadratic convergence without checking

$$0 \leq R_\ell \leq C^T C \quad \text{and} \quad 0 \leq R_\ell \leq N_\ell B B^T N_\ell?$$

Due to the quadratic nature of $\mathfrak{R}(.)$ we have

$$\mathfrak{R}(Y) = \mathfrak{R}(X) + \mathfrak{R}'|_X(Y-X) + \frac{1}{2}\mathfrak{R}''|_X(Y-X, Y-X).$$

Recall the Inexact Kleinman step:

$$R_\ell = \mathfrak{R}'|_{X_\ell}(X_{\ell+1}) - \mathfrak{R}'|_{X_\ell}(X_\ell) + \mathfrak{R}(X_\ell) = \mathfrak{R}'|_{X_\ell}(X_{\ell+1} - X_\ell) + \mathfrak{R}(X_\ell),$$

and thus

$$\mathfrak{R}(X_{\ell+1}) = R_\ell + \frac{1}{2}\mathfrak{R}''|_{X_\ell}(X_{\ell+1} - X_\ell, X_{\ell+1} - X_\ell).$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

Main Problem:

How can we ensure quadratic convergence without checking

$$0 \leq R_\ell \leq C^T C \quad \text{and} \quad 0 \leq R_\ell \leq N_\ell B B^T N_\ell?$$

Due to the quadratic nature of $\mathfrak{R}(.)$ we have

$$\mathfrak{R}(Y) = \mathfrak{R}(X) + \mathfrak{R}'|_X(Y - X) + \frac{1}{2}\mathfrak{R}''|_X(Y - X, Y - X).$$

Recall the Inexact Kleinman step:

$$R_\ell = \mathfrak{R}'|_{X_\ell}(X_{\ell+1}) - \mathfrak{R}'|_{X_\ell}(X_\ell) + \mathfrak{R}(X_\ell) = \mathfrak{R}'|_{X_\ell}(X_{\ell+1} - X_\ell) + \mathfrak{R}(X_\ell),$$

and thus

$$\mathfrak{R}(X_{\ell+1}) = R_\ell + \frac{1}{2}\mathfrak{R}''|_{X_\ell}(N_\ell, N_\ell).$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

Main Problem:

How can we ensure quadratic convergence without checking

$$0 \leq R_\ell \leq C^T C \quad \text{and} \quad 0 \leq R_\ell \leq N_\ell B B^T N_\ell?$$

Due to the quadratic nature of $\mathfrak{R}(.)$ we have

$$\mathfrak{R}(Y) = \mathfrak{R}(X) + \mathfrak{R}'|_X(Y - X) + \frac{1}{2}\mathfrak{R}''|_X(Y - X, Y - X).$$

Recall the Inexact Kleinman step:

$$R_\ell = \mathfrak{R}'|_{X_\ell}(X_{\ell+1}) - \mathfrak{R}'|_{X_\ell}(X_\ell) + \mathfrak{R}(X_\ell) = \mathfrak{R}'|_{X_\ell}(X_{\ell+1} - X_\ell) + \mathfrak{R}(X_\ell),$$

and thus

$$\mathfrak{R}(X_{\ell+1}) = R_\ell + \frac{1}{2} N_\ell B B^T N_\ell$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

New Question

How can we exploit $\mathfrak{R}(X_{\ell+1}) = R_\ell + \frac{1}{2} N_\ell B B^T N_\ell$ to control the ADI accuracy?

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

New Question

How can we exploit $\mathfrak{R}(X_{\ell+1}) = R_\ell + \frac{1}{2} N_\ell B B^T N_\ell$ to control the ADI accuracy?

Riccati residual

inner Lyapunov residual

$$N_\ell B B^T N_\ell = (X_{\ell+1} - X_\ell) B B^T (X_{\ell+1} - X_\ell)$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

New Question

How can we exploit $\mathfrak{R}(X_{\ell+1}) = R_\ell + \frac{1}{2} N_\ell B B^T N_\ell$ to control the ADI accuracy?

Riccati residual

inner Lyapunov residual

$$\begin{aligned} N_\ell B B^T N_\ell &= (X_{\ell+1} - X_\ell) B B^T (X_{\ell+1} - X_\ell) \\ &= X_{\ell+1} B B^T X_{\ell+1} + X_\ell B B^T X_\ell - X_\ell B B^T X_{\ell+1} - X_{\ell+1} B B^T X_\ell \end{aligned}$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

New Question

How can we exploit $\mathfrak{R}(X_{\ell+1}) = R_\ell + \frac{1}{2}N_\ell BB^T N_\ell$ to control the ADI accuracy?

Riccati residual — $\mathfrak{R}(X_{\ell+1})$

inner Lyapunov residual — R_ℓ

$$\begin{aligned} N_\ell BB^T N_\ell &= (X_{\ell+1} - X_\ell)BB^T(X_{\ell+1} - X_\ell) \\ &= X_{\ell+1}BB^T X_{\ell+1} + X_\ell BB^T X_\ell - X_\ell BB^T X_{\ell+1} - X_{\ell+1}BB^T X_\ell \\ &= K_{\ell+1}^T K_{\ell+1} + K_\ell^T K_\ell - K_{\ell+1}^T K_\ell - K_\ell^T K_{\ell+1} \end{aligned}$$

Inexact LRCF-NM for the ARE

Accuracy control for the (G-)LRCF-ADI

New Question

How can we exploit $\mathfrak{R}(X_{\ell+1}) = R_\ell + \frac{1}{2} N_\ell B B^T N_\ell$ to control the ADI accuracy?

Riccati residual

inner Lyapunov residual

$$\begin{aligned} N_\ell B B^T N_\ell &= (X_{\ell+1} - X_\ell) B B^T (X_{\ell+1} - X_\ell) \\ &= X_{\ell+1} B B^T X_{\ell+1} + X_\ell B B^T X_\ell - X_\ell B B^T X_{\ell+1} - X_{\ell+1} B B^T X_\ell \\ &= K_{\ell+1}^T K_{\ell+1} + K_\ell^T K_\ell - K_{\ell+1}^T K_\ell - K_\ell^T K_{\ell+1} \end{aligned}$$

- $\|R_\ell\|_2$ is stopping criterion in the LRCF-ADI.
- $\|\frac{1}{2} N_\ell B B^T N_\ell\|_2$ can be approximated via eigensolver due to symmetry.

Inexact LRCF-NM for the ARE

Implementation

$K_{\ell+1}$ can be accumulated during LRCF-ADI.

Recall $Z_{i+1} = [Z_i \;\; V_i]$ in LRCF-ADI.

$\Rightarrow K_{\ell+1}^{(i+1)} = B^T Z_{i+1} Z_{i+1}^T E = B^T Z_i Z_i^T E + B^T V_i V_i^T E = K_{\ell+1}^{(i)} + B^T V_i V_i^T E$

Inexact LRCF-NM for the ARE

Implementation

$K_{\ell+1}$ can be accumulated during LRCF-ADI.

Recall $Z_{i+1} = [Z_i \;\; V_i]$ in LRCF-ADI.

$\Rightarrow K_{\ell+1}^{(i+1)} = B^T Z_{i+1} Z_{i+1}^T E = B^T Z_i Z_i^T E + B^T V_i V_i^T E = K_{\ell+1}^{(i)} + B^T V_i V_i^T E$

We can force quadratic convergence via:

$$\begin{aligned} \|\mathfrak{R}(X_{\ell+1})\|_2 &\leq \|R_\ell\|_2 + \frac{1}{2}\|K_{\ell+1}^T K_{\ell+1} + K_\ell^T K_\ell - K_{\ell+1}^T K_\ell - K_\ell^T K_{\ell+1}\|_2 \\ &\leq \varepsilon_\ell := \alpha \mathfrak{R}(X_\ell)^2 \end{aligned}$$

M.E.S.S.

Features (Basic)

$$\begin{bmatrix} M & E \\ S & S \end{bmatrix}^{\mathrm{C}}$$

http://www.mpi-magdeburg.mpg.de/mess

M.E.S.S. Includes:

- Solvers for large sparse **Lyapunov** and **Riccati** equations,
- routines for $\mathcal{H}_2$-optimal and balanced truncation **model reduction**,
- **linear quadratic regulator** feedback computation,
- interfaces to: BLAS, LAPACK, SLICOT, SuiteSparse (partial), . . .

M.E.S.S.

Features (shared memory parallel)

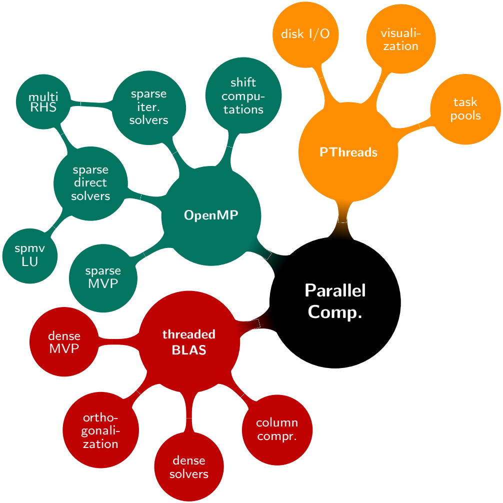

M.E.S.S. Includes:

- **automatic conversion** of sparse matrix formats,
- **uniform API** for sparse and dense linear algebra,
- plain/gzip/bzip2 file I/O,
- basic visualization (GNUPlot, TikZ, X11),
- **Matrix Market** data exchange **format**.

Cu.M.E.S.S.

Features (upcoming)

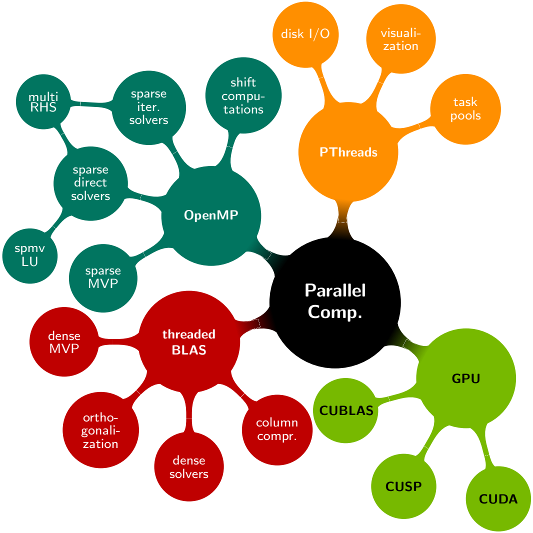

Cu.M.E.S.S. goals:

- offload **expensive computations** to **GPU** (shifted linear system solves),
- perform **less expensive** tasks **on host CPU in parallel** (residual evaluations, updates of solution factors),
- **minimizes communication** overhead,
- employ **GPU tailored** storage **formats** (ELLR-T, ICRS),
- **mixed precision** codes in the implementation.

Preliminary Results

Hardware and Example

Processor	#cores	Freq. (GHz)	L2 cache (MB)	Mem. (GB)
Intel®Xeon® QuadCore	4	2.83	8	24
NVIDIA®GEFORCE®GTX 580	512	1.5	-	3

Selective cooling of steel profiles

- boundary control of a 2d linearized heat equation ($n_{in} = 6, \; n_{out} = 7$)
- spatial FEM semi-discretization (n=20209)
- available in Oberwolfach Benchmark Collection
  `http://www.imtek.de/simulation/benchmark/`

Preliminary Results

Test Setup

- Simple BiCG iteration
  - block version planned using ELLR-T matrix-matrix-kernel from Almeria,
  - recycling investigation planned with K. Ahuja
- No preconditioning
  - e.g., SPAI available on GPUs
- CRS matrix storage (slow on GPUs)
  - ELLR-T data structures ready. Waiting for MVP-kernel update (CUDA 3 to CUDA 4)
  - ICRS comparison planned to reduce memory consumption.

Timings use maximum iterations limited to Newton: 1, ADI: 40, BiCG: 100

Preliminary Results

Test Setup

- Simple BiCG iteration
  - block version planned using ELLR-T matrix-matrix-kernel from Almeria,
  - recycling investigation planned with K. Ahuja
- No preconditioning
  - e.g., SPAI available on GPUs
- CRS matrix storage (slow on GPUs)
  - ELLR-T data structures ready. Waiting for MVP-kernel update (CUDA 3 to CUDA 4)
  - ICRS comparison planned to reduce memory consumption.

Timings use maximum iterations limited to Newton: 1, ADI: 40, BiCG: 100

Preliminary Results

Test Setup

- Simple BiCG iteration
  - block version planned using ELLR-T matrix-matrix-kernel from Almeria,
  - recycling investigation planned with K. Ahuja
- No preconditioning
  - e.g., SPAI available on GPUs
- CRS matrix storage (slow on GPUs)
  - ELLR-T data structures ready. Waiting for MVP-kernel update (CUDA 3 to CUDA 4)
  - ICRS comparison planned to reduce memory consumption.

Timings use maximum iterations limited to Newton: 1, ADI: 40, BiCG: 100

Preliminary Results

M.E.S.S.

Stage	Time (in sec.)	% of total time
BiCG	124.96	93.1 %
ADI parameters	1.28	0.9 %
ADI others	6.20	4.6 %
Newton others	1.95	1.4 %
Total time	134.40	

Plain M.E.S.S. on host CPU only

Computations by A. Remón and P. Ezzati

Preliminary Results

M.E.S.S. + GPU

Stage	Time (in sec.)	% of total time
BiCG	100.84	91.6 %
ADI parameters	1.28	0.9 %
ADI others	5.90	5.3 %
Newton others	2.02	1.8 %
Total time	110.04	

M.E.S.S. on host CPU + Cu.M.E.S.S. on GPU

Computations by A. Remón and P. Ezzati

Preliminary Results

M.E.S.S. + GPU

Stage	Time (in sec)	% of total time
BiCG	100.84	91.6 %
ADI parameters	1.28	0.9 %
ADI others	5.90	5.3 %
Newton others	2.02	1.8 %
Total time	110.04	

-23%

-20%

M.E.S.S. on host CPU + Cu.M.E.S.S. on GPU

Computations by A. Remón and P. Ezzati

Preliminary Results

M.E.S.S. + GPU

Stage	Time (in se[illegible])	% of total time
BiCG	100.84	91.6 %
ADI parameters	1.28	0.9 %
ADI others	5.90	5.3 %
Newton others	2.02	1.8 %
Total time	110.04	

-23%

-20%

M.E.S.S. on host CPU + Cu.M.E.S.S. on GPU

Thank you for the attention!

Computations by A. Remon and P. Ezzati