

The Augmented Block-Cimmino Distributed Method

Improvements to the Conjugate Gradient accelerated Block Cimmino
Method

Mohamed Zenadi¹, Iain S. Duff² Ronan Guivarch¹ Daniel Ruiz¹,
IRIT - ENSEEIHT¹, CERFACS and Rutherford Appleton Laboratory²

Jun 17, 2013

Block Cimmino : The basic facts

- Block row projection method [Elfving (Numer. Math. 1980)]

Partitioning the system $Ax = b$

$$\begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_p \end{pmatrix} x = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_p \end{pmatrix}$$

Block Cimmino : The basic facts

- Block row projection method [Elfving (Numer. Math. 1980)]

Partitioning the system $Ax = b$

$$\begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_p \end{pmatrix} x = \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_p \end{pmatrix}$$

Partitions can be obtained by cutting uniformly the matrix (with no permutation) or using a partitioner (we use PaToH in our examples)

Block Cimmino : The basic facts

- Block row projection method [Elfving (Numer. Math. 1980)]

The Block Cimmino Iteration

$$\delta_i^{(k)} = A_i^+ b_i - P_{\mathcal{R}(A_i^T)} x^{(k)}$$
$$x^{(k+1)} = x^{(k)} + \nu \sum_{i=1}^p \delta_i^{(k)}$$

where:

$$A_i^+ = A_i^T (A_i A_i^T)^{-1}$$

and

$$P_{\mathcal{R}(A_i^T)} = A_i^+ A_i$$

Block Cimmino : The basic facts

- Block row projection method [Elfving (Numer. Math. 1980)]
- CG acceleration: iteration matrix is $\sum_{i=1}^p A_i^+ A_i = \sum_{i=1}^p P_{\mathcal{R}(A_i^T)}$

Acceleration

Apply CG to solve the SPD system

$$\sum_{i=1}^p A_i^+ A_i x = \sum_{i=1}^p A_i^+ b_i$$

Block Cimmino : The basic facts

- Block row projection method [Elfving (Numer. Math. 1980)]
- CG acceleration: iteration matrix is $\sum_{i=1}^p A_i^+ A_i = \sum_{i=1}^p P_{\mathcal{R}(A_i^T)}$
- Can also exploit 2nd and 3rd levels of parallelism (sparsity structure, BLAS3 Kernels)

$$\text{Projections: } \delta_i^{(k)} = A_i^+ b_i - P_{\mathcal{R}(A_i^T)} x^{(k)}$$

Block Cimmino : The basic facts

- Block row projection method [Elfving (Numer. Math. 1980)]
- CG acceleration: iteration matrix is $\sum_{i=1}^p A_i^+ A_i = \sum_{i=1}^p P_{\mathcal{R}(A_i^T)}$
- Can also exploit 2nd and 3rd levels of parallelism (sparsity structure, BLAS3 Kernels)

$$\text{Projections: } \delta_i^{(k)} = A_i^+ b_i - P_{\mathcal{R}(A_i^T)} x^{(k)}$$

Solve independently using MUMPS the systems for each partition

$$\begin{bmatrix} I & A_i^T \\ A_i & 0 \end{bmatrix} \begin{bmatrix} u_i \\ v_i \end{bmatrix} = \begin{bmatrix} 0 \\ b_i - A_i x \end{bmatrix}$$

$$\text{where : } u_i = A_i^+ (b_i - A_i x) = \delta_i$$

Issues with Block-Cimmino:

- Convergence is problem dependent
- Erratic convergence behaviour (plateaux based)
- Trail of small eigenvalues
- Multiple solves require a re-run of Block-CG (too expensive)

Issues with Block-Cimmino:

- Convergence is problem dependent
- Erratic convergence behaviour (plateaux based)
- Trail of small eigenvalues
- Multiple solves require a re-run of Block-CG (too expensive)

Proposed solution:

- Enforce numerical orthogonality between partitions by adding extra variables and constraints
- Extract a condensed smaller subsystem (similar to Schur complement techniques) that can be reused for efficient further solves

Issues with Block-Cimmino:

- Convergence is problem dependent
- Erratic convergence behaviour (plateaux based)
- Trail of small eigenvalues
- Multiple solves require a re-run of Block-CG (too expensive)


Proposed solution:

- Enforce numerical orthogonality between partitions by adding extra variables and constraints
- Extract a condensed smaller subsystem (similar to Schur complement techniques) that can be reused for efficient further solves

⇒ **Augmented Block Cimmino Distributed solver (ABCD solver)**

The augmentation process


- Partition the matrix with respect to its structure (not necessarily in block-tridiagonal form)



Illustrative example

The augmentation process


- Partition the matrix with respect to its structure (not necessarily in block-tridiagonal form)
- For each pair of partitions ($j > i$), expand with $C_{i,j} = A_i A_j^T$,



Illustrative example

The augmentation process


- Partition the matrix with respect to its structure (not necessarily in block-tridiagonal form)
- For each pair of partitions ($j > i$), expand with $C_{i,j} = A_i A_j^T$, and enforce numerical orthogonality



Illustrative example

The augmentation process

- Partition the matrix with respect to its structure (not necessarily in block-tridiagonal form)
- For each pair of partitions ($j > i$), expand with $C_{i,j} = A_i A_j^T$, and enforce numerical orthogonality to obtain $\bar{A} = [A \quad C]$



Illustrative example

The augmentation process

- Partition the matrix with respect to its structure (not necessarily in block-tridiagonal form)
- For each pair of partitions ($j > i$), expand with $C_{i,j} = A_i A_j^T$, and enforce numerical orthogonality to obtain $\bar{A} = [A \ C]$
- Add extra constraints to build an equivalent linear system :

$$\begin{bmatrix} A & C \\ 0 & I \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix},$$

where $y = 0$ ensures the same solution x .

The augmentation process

- Partition the matrix with respect to its structure (not necessarily in block-tridiagonal form)
- For each pair of partitions ($j > i$), expand with $C_{i,j} = A_i A_j^T$, and enforce numerical orthogonality to obtain $\bar{A} = [A \ C]$
- Add extra constraints to build an equivalent linear system :

$$\begin{bmatrix} A & C \\ 0 & I \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix},$$

where $y = 0$ ensures the same solution x .

Problem

the extra partition $Y = [0 \ I]$, linked to the constraints equations, is not orthogonal to the previous partitions in $\bar{A} = [A \ C]$.

The augmentation process

To enforce this orthogonality, we project the column vectors Y^T onto the null space of $\bar{A} = [A \ C]$ (orthogonal complement of $\mathcal{R}(\bar{A}^T)$) :


$$W^T = (I - P) Y^T,$$

where (as a result of the enforced orthogonality) :

$$P = P_{\mathcal{R}(\bar{A}^T)} = P_{\bigoplus_{i=1}^p \mathcal{R}(\bar{A}_i^T)} = \sum_{i=1}^p P_{\mathcal{R}(\bar{A}_i^T)}$$

We finally obtain $\begin{bmatrix} A & C \\ B & S \end{bmatrix}$, where $[B \ S] = W$, an augmented matrix with mutually numerically orthogonal partitions

The augmentation process



Illustrative example

The augmentation process

To keep the consistency within the solution of the new system :

$$\begin{bmatrix} A & C \\ B & S \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} b \\ f \end{bmatrix}$$

we compute the right hand side f as :

$$\begin{aligned} f &= [B \ S] \begin{bmatrix} x \\ 0 \end{bmatrix} = Y(I - P) \begin{bmatrix} x \\ 0 \end{bmatrix} \\ &= -YP \begin{bmatrix} x \\ 0 \end{bmatrix} \quad (\text{since } Y = [0 \ I]) \\ &= -Y\bar{A}^+ \bar{A} \begin{bmatrix} x \\ 0 \end{bmatrix} \\ f &= -Y\bar{A}^+ b \end{aligned}$$

Since all the partitions in the new equivalent linear system

$$\begin{bmatrix} A & C \\ B & S \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} b \\ f \end{bmatrix}$$

are mutually numerically orthogonal, the Cimmino iteration matrix becomes the Identity matrix, and the solution can be directly obtained as :

$$\begin{aligned} \begin{bmatrix} x \\ y \end{bmatrix} &= \bar{A}^+ b + W^+ f \\ &= \bar{A}^+ b - W^+ Y \bar{A}^+ b \\ &= \sum_{i=1}^p \bar{A}_i^+ b_i - W^+ Y \sum_{i=1}^p \bar{A}_i^+ b_i \end{aligned}$$

Computational Ingredients

Knowing that $W = [B \ S] = Y(I - P)$, with $Y = [0 \ I]$,
we have :

$$\begin{aligned} WW^T &= Y(I - P)(I - P)^T Y^T \\ &= Y(I - P)^2 Y^T \\ &= Y(I - P) Y^T \\ &= [B \ S] Y^T \\ &= S \end{aligned}$$

Computational Ingredients

Knowing that $W = [B \ S] = Y(I - P)$, with $Y = [0 \ I]$,
we have :

$$\begin{aligned} WW^T &= Y(I - P)(I - P)^T Y^T \\ &= Y(I - P)^2 Y^T \\ &= Y(I - P) Y^T \\ &= [B \ S] Y^T \\ &= S \end{aligned}$$

Therefore $S = Y(I - P) Y^T$ and is SPD.

Computational Ingredients

Knowing that $W = [B \ S] = Y(I - P)$, with $Y = [0 \ I]$, we have :

$$\begin{aligned} WW^T &= Y(I - P)(I - P)^T Y^T \\ &= Y(I - P)^2 Y^T \\ &= Y(I - P) Y^T \\ &= [B \ S] Y^T \\ &= S \end{aligned}$$

Therefore $S = Y(I - P) Y^T$ and is SPD.

And the pseudo inverse $W^+ = W^T(WW^T)^{-1}$ is given by

$$\begin{aligned} W^+ &= W^T S^{-1} \\ W^+ &= (I - P) Y^T S^{-1} \end{aligned}$$

The solution is thus given by :

$$\begin{aligned} \begin{bmatrix} x \\ y \end{bmatrix} &= \bar{A}^+ b + W^+ f \\ &= \bar{A}^+ b - (I - P) Y^T S^{-1} Y \bar{A}^+ b \end{aligned}$$

The solution is thus given by :

$$\begin{aligned} \begin{bmatrix} x \\ y \end{bmatrix} &= \bar{A}^+ b + W^+ f \\ &= \bar{A}^+ b - (I - P) Y^T S^{-1} Y \bar{A}^+ b \end{aligned}$$

which can be computed through the 4 following steps :

The solution is thus given by :

$$\begin{aligned} \begin{bmatrix} x \\ y \end{bmatrix} &= \bar{A}^+ b + W^+ f \\ &= \bar{A}^+ b - (I - P) Y^T S^{-1} Y \bar{A}^+ b \end{aligned}$$

which can be computed through the 4 following steps :

- Build $w = \bar{A}^+ b$ and then by simple restriction set $f = -Yw$

The solution is thus given by :

$$\begin{aligned} \begin{bmatrix} x \\ y \end{bmatrix} &= \bar{A}^+ b + W^+ f \\ &= \bar{A}^+ b - (I - P) Y^T S^{-1} Y \bar{A}^+ b \end{aligned}$$

which can be computed through the 4 following steps :

- Build $w = \bar{A}^+ b$ and then by simple restriction set $f = -Yw$
- Solve $Sz = f$ (S should be small enough)

The solution is thus given by :

$$\begin{aligned} \begin{bmatrix} x \\ y \end{bmatrix} &= \bar{A}^+ b + W^+ f \\ &= \bar{A}^+ b - (I - P) Y^T S^{-1} Y \bar{A}^+ b \end{aligned}$$

which can be computed through the 4 following steps :

- Build $w = \bar{A}^+ b$ and then by simple restriction set $f = -Yw$
- Solve $Sz = f$ (S should be small enough)
- Expand z and then project it onto the null space of \bar{A} viz.

$$u = (I - P) Y^T z$$

The solution is thus given by :

$$\begin{aligned} \begin{bmatrix} x \\ y \end{bmatrix} &= \bar{A}^+ b + W^+ f \\ &= \bar{A}^+ b - (I - P) Y^T S^{-1} Y \bar{A}^+ b \end{aligned}$$

which can be computed through the 4 following steps :

- Build $w = \bar{A}^+ b$ and then by simple restriction set $f = -Yw$
- Solve $Sz = f$ (S should be small enough)
- Expand z and then project it onto the null space of \bar{A} viz.

$$u = (I - P) Y^T z$$

- Then sum $w + u$ to obtain the solution $\begin{bmatrix} x \\ y \end{bmatrix}$ (where $y = 0$)

The solution is thus given by :

$$\begin{aligned} \begin{bmatrix} x \\ y \end{bmatrix} &= \bar{A}^+ b + W^+ f \\ &= \bar{A}^+ b - (I - P) Y^T S^{-1} Y \bar{A}^+ b \end{aligned}$$

which can be computed through the 4 following steps :

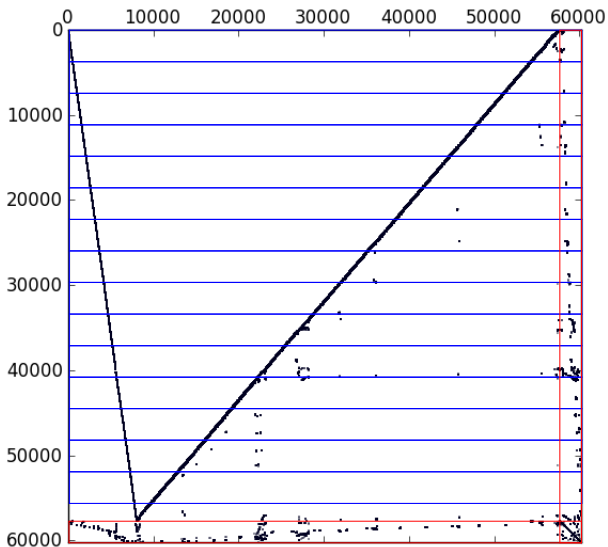
- Build $w = \bar{A}^+ b$ and then by simple restriction set $f = -Yw$
- Solve $Sz = f$ (S should be small enough)
- Expand z and then project it onto the null space of \bar{A} viz.

$$u = (I - P) Y^T z$$

- Then sum $w + u$ to obtain the solution $\begin{bmatrix} x \\ y \end{bmatrix}$ (where $y = 0$)

Note that we don't need to build B , only S is used

bayer01(57k) - 16 uniform parts - $N(S) = 2409$



R6 (132k), non-symmetric, 16 partitions on 32 cores

	BC (blk. size = 16)	ABCD (size $S = 6506$)
Fact.	0.2 s.	0.2s.
CG	(521itr) 107.6 s.	-
Augmentation	-	0.16s.
Build S	-	5.5s.
Fact S	-	1.2s.

bmw3_2 (227k) 16 partitions on 32 cores

	BC (blk. size = 1)	ABCD (size $S = 16695$)
Fact.	1.7 s.	1.97s.
CG	<i>(Failed)</i> 176.5 s.	-
Augmentation	-	0.6s.
Build S	-	40.0s.
Fact S	-	18.0s.

Hamrle3 (1.447M), non-symmetric, 64 partitions on 32 cores

	BC (blk. size = 4)	ABCD (size $S = 54608$)
Fact.	2.13 s.	3.4s.
CG	(615itr) 282.90 s.	-
Augmentation	-	4.6s.
Build S	-	145.4s.
Fact S	-	49.1s.

Hamrle3 (1.447M) 64 partitions on 32 cores - MUMPS_TRUNK

	BC (blk. size = 4)	ABCD (size $S = 54608$)
Fact.	2.13 s.	2.7s.
CG	(615itr) 282.90 s.	-
Augmentation	-	4.6s.
Build S	-	97.0s.
Fact S	-	47.1s.

Hamrle3 (1.447M) 64 partitions on 32 cores - MUMPS_TRUNK+ES

	BC (blk. size = 4)	ABCD (size $S = 54608$)
Fact.	2.13 s.	2.7s.
CG	(615itr) 282.90 s.	-
Augmentation	-	4.6s.
Build S	-	58.4s.
Fact S	-	43.1s.

Hamrle3 (1.447M) 64 partitions on 32 cores - MUMPS_TRUNK+ES

	BC (blk. size = 4)	ABCD (size $S = 54608$)
Fact.	2.13 s.	2.7s.
CG	(615itr) 282.90 s.	-
Augmentation	-	4.6s.
Build S	-	145s. → 58.4s.
Fact S	-	43.1s.

Thanks to **François-Henry ROUET** for the hints on **Exploit-Sparsity** feature (available in the future release of MUMPS)

Main issue: Size of S

- sparsity structure (preprocessing, permutations...)
- number and size of partitions
- interconnections between partitions

Matrix	N	Pts	Size of S	Ratio
Hamrle3	1,447,360	64	54,608	3.8%
R6	132,106	16	8,536	6.5%
ohne2	181,343	16	48,920	27%

Main issue: Size of S and its density

- sparsity structure (preprocessing, permutations...)
- number and size of partitions
- interconnections between partitions

Matrix	N	Pts	Size of S	Ratio	NZ(S)
Hamrle3	1,447,360	64	54,608	3.8%	90,105,183
R6	132,106	16	8,536	6.5%	9,836,139
ohne2	181,343	16	48,920	27%	221,956,502

Main issue: Size of S and its density

- sparsity structure (preprocessing, permutations...)
- number and size of partitions
- interconnections between partitions

Possible solutions

- Relax the augmentation process by reducing the number of columns in C and therefore reduce the size of S
- Avoid building S by using implicitly (MV products) in an iterative process (CG, S is SPD)

Target:

- A reduced size of S with respect to the size of A : better control of memory requirements

Issues:

- The augmented partitions \bar{A}_i lose "*partly*" their mutual numerical orthogonality
- $(I - P)$ is no longer explicitly available, and must be recovered via an iterative process

Relaxation of the augmentation process

Target:

- A reduced size of S with respect to the size of A : better control of memory requirements

Issues:

- The augmented partitions \bar{A}_i lose "partly" their mutual numerical orthogonality
- $(I - P)$ is no longer explicitly available, and must be recovered via an iterative process

Results on bayer01

Drop threshold	0	0.1	0.2	0.3	0.4	BC
Size of S	752	270	77	46	18	-
$w = \bar{A}^+ b$	1	103	282	466	1183	1700
AVG. iter per column	1	17	45	64	340	-
Total iterations to build S	752	4590	3465	2944	6130	-

Target:

- A reduced size of S with respect to the size of A : better control of memory requirements

Issues:

- The augmented partitions \bar{A}_i lose "*partly*" their mutual numerical orthogonality
- $(I - P)$ is no longer explicitly available, and must be recovered via an iterative process

In general :

- + Build a reduced size S
- + Outperforms regular Block-Cimmino after a few successive solves (in the bayer01's case after 15 solves with a drop of 0.3)
- Slower to build (less parallel advantages + iterations)

Iterative solution of $Sz = f$

Recall that $S = Y(I - P)Y^T$ is SPD, therefore the system $Sz = f$ can be solved using CG.

Iterative solution of $Sz = f$

Recall that $S = Y(I - P)Y^T$ is SPD, therefore the system $Sz = f$ can be solved using CG. In the CG iteration, S is used implicitly in the instruction:

$$\alpha_k = (r_k^T r_k) / (p_k^T S p_k)$$

Iterative solution of $Sz = f$

Recall that $S = Y(I - P)Y^T$ is SPD, therefore the system $Sz = f$ can be solved using CG. In the CG iteration, S is used implicitly in the instruction:

$$\alpha_k = \left(r_k^T r_k \right) / \left(p_k^T S p_k \right)$$

The matrix-vector product can be written as:

$$\begin{aligned} S p_k &= Y(I - P)Y^T p_k \\ &= p_k - Y P Y^T p_k \end{aligned}$$

Iterative solution of $Sz = f$

Recall that $S = Y(I - P)Y^T$ is SPD, therefore the system $Sz = f$ can be solved using CG. In the CG iteration, S is used implicitly in the instruction:

$$\alpha_k = (r_k^T r_k) / (p_k^T S p_k)$$

The matrix-vector product can be written as:

$$\begin{aligned} S p_k &= Y(I - P)Y^T p_k \\ &= p_k - Y P Y^T p_k \end{aligned}$$

Where $Y^T p_k = \begin{bmatrix} 0 \\ p_k \end{bmatrix}$ is to be projected by solving augmented systems using MUMPS.


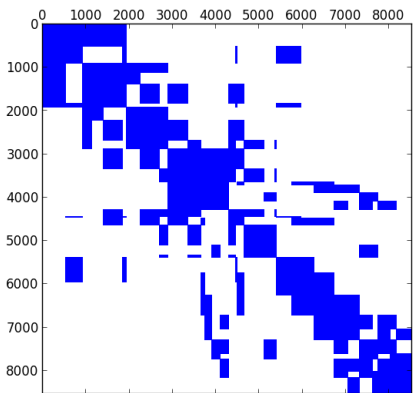
Iterative solution of $Sz = f$: Test

- Conjugate Gradient acceleration with stopping criteria 1×10^{-8} .
- A testing (rudimentary) preconditioner (partial build of S)

- Conjugate Gradient acceleration with stopping criteria 1×10^{-8} .
- A testing (rudimentary) preconditioner (partial build of S)

	size(S)	CG	PCG
bayer01	752	F	257
Hamrle3	54,608	1,911	1,238
R6	8,536	F	F
ohne2	48,920	32,301	8,066

The R6 case



The good side:

- Building S is fast (and could be faster), thanks MUMPS!
- ABCD solves block-Cimmino convergence issues

The issues:

- S can be really large and dense (problem dependent)
- Reducing the size of S can perform better than block-Cimmino in the long run.
- Iteratively solving $Sz = f$ is not there yet. Still looking for a preconditioner.