



πλατφόρμα μαθηματικών μεθόδων και λογισμικού
για την επίλυση προβλημάτων **πολλαπλών πεδίων**
σε σύγχρονες υπολογιστικές αρχιτεκτονικές

Αριθμητικές Μέθοδοι Collocation

Απεικόνιση σε Σύγχρονες Υπολογιστικές Αρχιτεκτονικές



European Union
European Social Fund




MINISTRY OF EDUCATION & RELIGIOUS AFFAIRS
MANAGING AUTHORITY

Co-financed by Greece and the European Union



Hermite Collocation Method

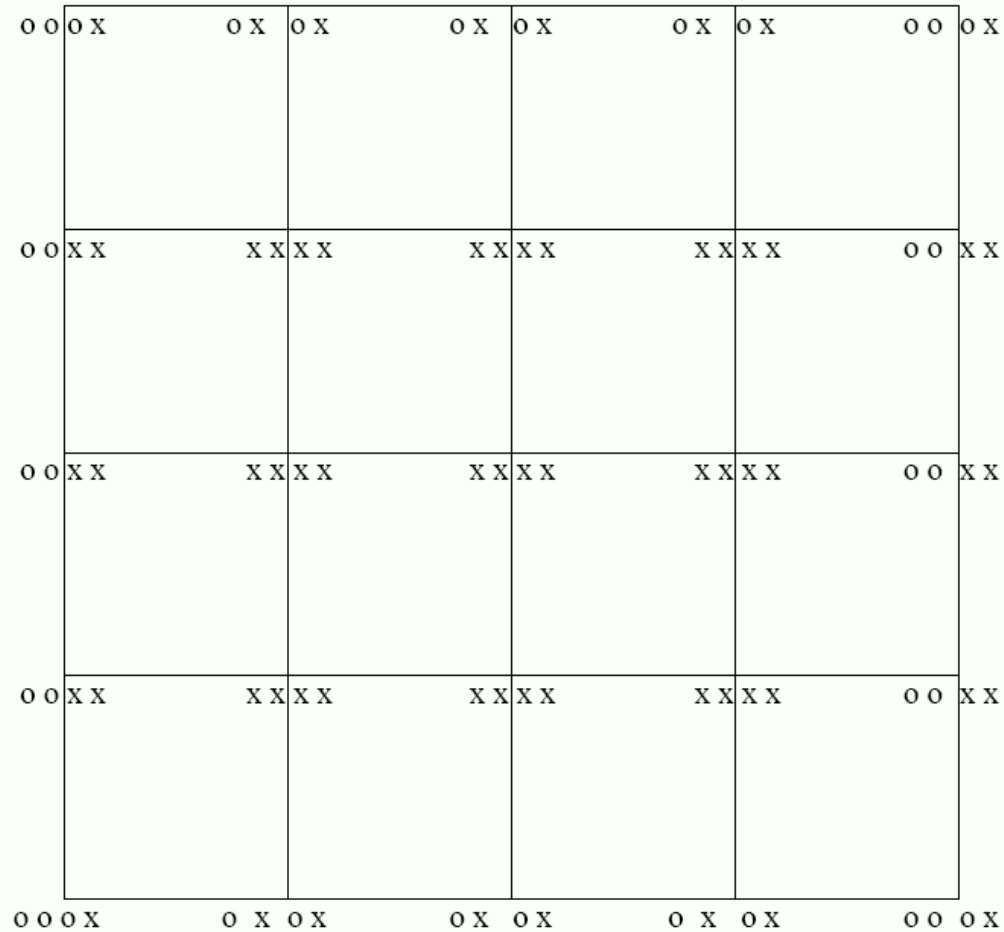
$$\begin{aligned} \text{BVP} \quad \mathcal{L} u(x, y) &= f(x, y) \quad , \quad (x, y) \in \Omega \\ \mathcal{B} u(x, y) &= g(x, y) \quad , \quad (x, y) \in \partial\Omega \end{aligned}$$




$$\begin{cases} u_n(x_i, y_j) = a_{2i-1, 2j-1} \quad , \quad h \frac{\partial}{\partial y} u_n(x_i, y_j) = a_{2i-1, 2j} \\ h \frac{\partial}{\partial x} u_n(x_i, y_j) = a_{2i, 2j-1} \quad , \quad h^2 \frac{\partial^2}{\partial x \partial y} u_n(x_i, y_j) = a_{2i, 2j} \end{cases}$$




Red – Black Collocation Linear system



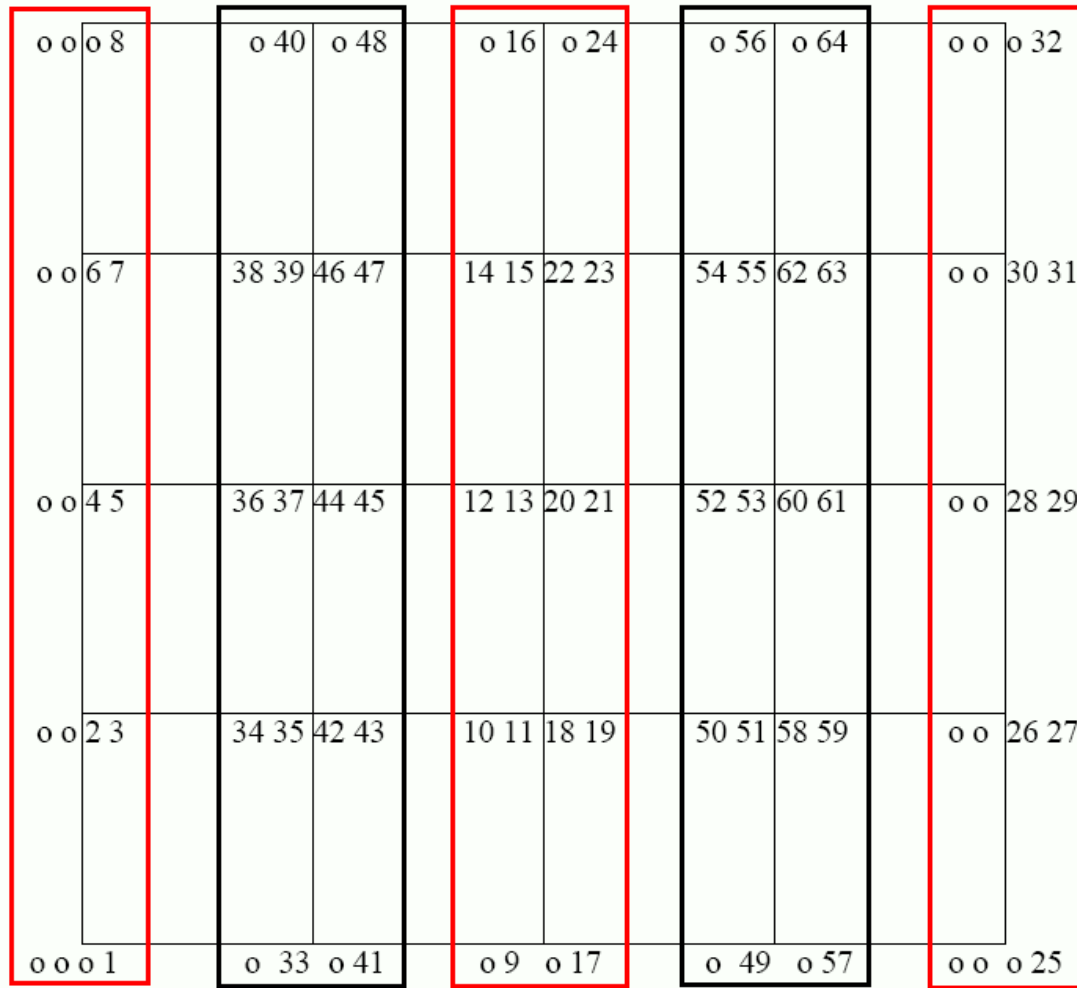
Red – Black Collocation Linear system



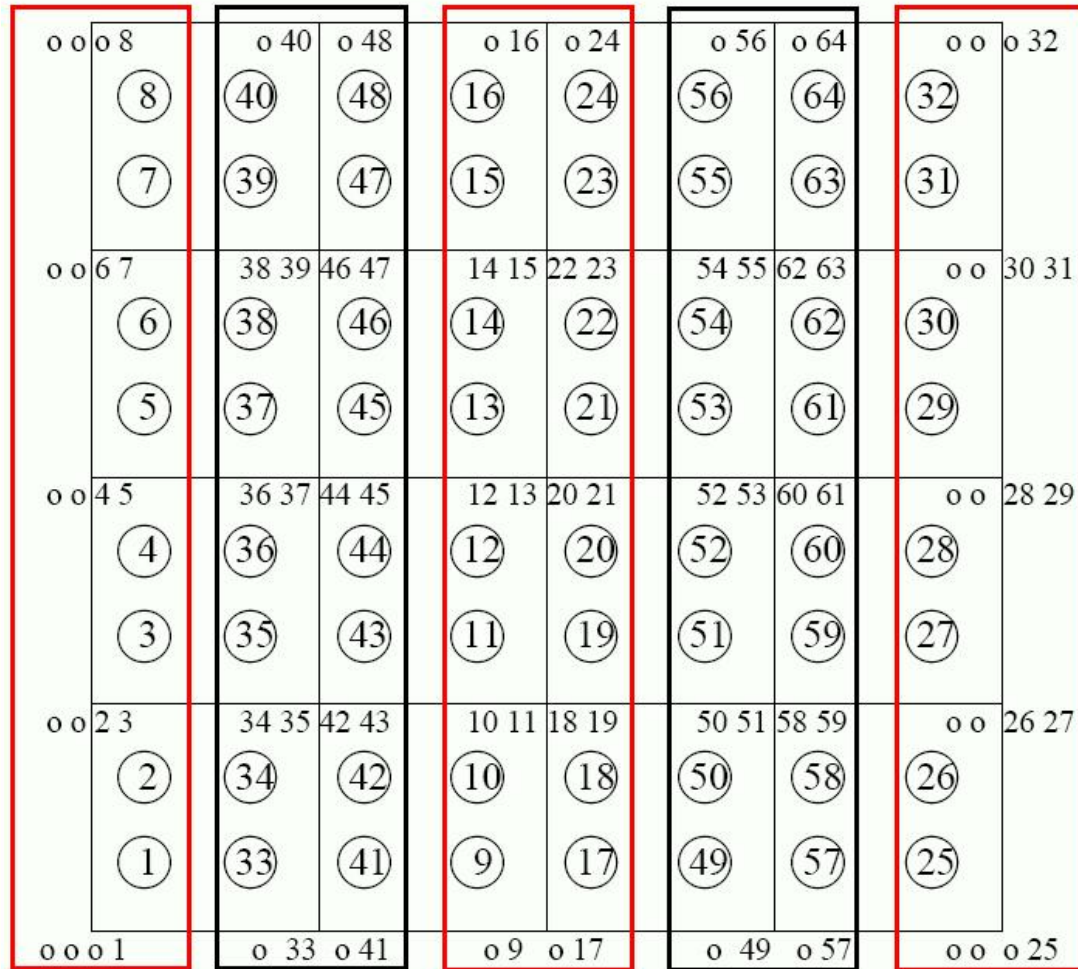
Red – Black Collocation Linear system



Red – Black Collocation Linear system




Red – Black Collocation Linear system



Model Problem

$$\begin{aligned} \nabla^2 u(x, y) - \lambda u(x, y) &= f(x, y) \quad , \quad (x, y) \in \Omega \\ u(x, y) &= g(x, y) \quad , \quad (x, y) \in \partial\Omega \quad \text{with} \quad \lambda \geq 0 \end{aligned}$$

$$A = \begin{bmatrix} D_R & H_B \\ H_R & D_B \end{bmatrix}$$



Helmholtz Collocation Matrix

$$A = \begin{bmatrix} D_R & H_B \\ H_R & D_B \end{bmatrix}$$

$$D_R = 2 \operatorname{diag} \left[\frac{1}{2} A_2 \quad A_1 \quad A_2 \quad \cdots \quad A_1 \quad A_2 \quad -\frac{1}{2} A_2 \right]$$

$$D_B = 2 \operatorname{diag} [A_1 \quad A_2 \quad \cdots \quad A_1 \quad A_2]$$

$$H_B = - \begin{bmatrix} A_3 & -A_4 & O & O & \cdots & O & O & O & O \\ A_3 & A_4 & A_3 & -A_4 & \cdots & O & O & O & O \\ -A_3 & -A_4 & A_3 & -A_4 & \cdots & O & O & O & O \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ O & O & O & O & \cdots & A_3 & A_4 & A_3 & -A_4 \\ O & O & O & O & \cdots & -A_3 & -A_4 & A_3 & -A_4 \\ O & O & O & O & \cdots & O & O & A_3 & A_4 \end{bmatrix}$$

$$H_R = - \begin{bmatrix} A_4 & A_3 & -A_4 & O & O & \cdots & O & O & O & O & O \\ -A_4 & A_3 & -A_4 & O & O & \cdots & O & O & O & O & O \\ O & A_3 & A_4 & A_3 & -A_4 & \cdots & O & O & O & O & O \\ O & -A_3 & -A_4 & A_3 & -A_4 & \cdots & O & O & O & O & O \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots \\ O & O & O & O & O & \cdots & A_3 & A_4 & A_3 & -A_4 & O \\ O & O & O & O & O & \cdots & -A_3 & -A_4 & A_3 & -A_4 & O \\ O & O & O & O & O & \cdots & O & O & A_3 & A_4 & -A_4 \\ O & O & O & O & O & \cdots & O & O & -A_3 & -A_4 & -A_4 \end{bmatrix}$$



Red – Black Collocation Linear system

$$A_i = \begin{bmatrix} a_2 & a_3 & -a_4 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 & 0 \\ a_4 & a_1 & -a_2 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 & 0 \\ 0 & a_1 & a_2 & a_3 & -a_4 & \cdots & 0 & 0 & 0 & 0 & 0 \\ 0 & a_3 & a_4 & a_1 & -a_2 & \cdots & 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & a_1 & a_2 & a_3 & -a_4 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & a_3 & a_4 & a_1 & -a_2 & 0 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & a_1 & a_2 & -a_4 \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 0 & a_3 & a_4 & -a_2 \end{bmatrix}$$

| | a_1 | a_2 | a_3 | a_4 |
|-------|-------|------------|-------|------------|
| A_1 | r^+ | s^+ | q | t^+ |
| A_2 | s^+ | u^+ | t^- | ϵ |
| A_3 | q | t^- | r^- | s^- |
| A_4 | t^+ | ϵ | s^- | u^- |

with

$$\begin{aligned} \epsilon &= -\frac{\lambda}{24n_s^2}, \quad q = 24 + 22\epsilon, \\ r^\pm &= 86\epsilon - 24 \pm (48\epsilon - 18)\sqrt{3}, \\ s^\pm &= 13\epsilon - 12 \pm (7\epsilon - 8)\sqrt{3}, \\ t^\pm &= 5\epsilon + 3 \pm (\epsilon + 1)\sqrt{3}, \\ u^\pm &= 2\epsilon - 3 \pm (\epsilon - 2)\sqrt{3}. \end{aligned}$$



Iterative
+
Parallel

∞ Ο collocation πίνακας είναι μεγάλης διάστασης, αραιός και δεν έχει κάποια ιδιαίτερα χρήσιμη ιδιότητα (π.χ. συμμετρικός, θετικά διευθετιμένος)



Iterative Solution

$$A = \begin{bmatrix} D_R & H_B \\ H_R & D_B \end{bmatrix}$$

$$A = D_A - L_A - U_A$$

Όπου

$$D_A = \begin{bmatrix} D_R & O \\ O & D_B \end{bmatrix}, \quad L_A = \begin{bmatrix} O & O \\ -H_R & O \end{bmatrix}, \quad U_A = \begin{bmatrix} O & -H_B \\ O & O \end{bmatrix}$$

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_R \\ \mathbf{x}_B \end{bmatrix} \quad \text{and} \quad \mathbf{b} = \begin{bmatrix} \mathbf{b}_R \\ \mathbf{b}_B \end{bmatrix} .$$



Iterative Solution

$$M_1^{-1} A M_2^{-1} M_2 \mathbf{x} = M_1^{-1} \mathbf{b}$$

where $M_1 = D_A - L_A = D_A(I - D_A^{-1}L_A)$

and $M_2 = I - D_A^{-1}U_A$



$$\begin{bmatrix} I & O \\ O & S \end{bmatrix} \begin{bmatrix} \mathbf{x}_R + D_R^{-1}H_R\mathbf{x}_B \\ \mathbf{x}_B \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{b}}_R \\ \hat{\mathbf{b}}_B \end{bmatrix}$$

where $S = D_B - H_R D_R^{-1} H_B$

$$\hat{\mathbf{b}}_R = D_R^{-1} \mathbf{b}_R \quad \text{and} \quad \hat{\mathbf{b}}_B = \mathbf{b}_B - H_R \hat{\mathbf{b}}_R$$



Eigenvalues of Collocation matrix



Schur Complement Iterative Solution

S1: Solve $D_R \hat{\mathbf{b}}_R = \mathbf{b}_R$

S2: Evaluate $\hat{\mathbf{b}}_B = \mathbf{b}_B - H_R \hat{\mathbf{b}}_R$

S3: Solve with BiCGSTAB $S \mathbf{x}_B = \hat{\mathbf{b}}_B$

S4: Evaluate $\hat{\mathbf{x}}_B = H_B \mathbf{x}_B$

S5: Solve $D_R \hat{\mathbf{x}}_R = \hat{\mathbf{x}}_B$

S6: Evaluate $\mathbf{x}_R = \hat{\mathbf{b}}_R - \hat{\mathbf{x}}_R$



Parallel Iterative Solution of Collocation Linear system on Shared Memory Architectures

- Ομοιόμορφη κατανομή φόρτου εργασίας μεταξύ core threads
- Ελαχιστοποίηση κύκλων αδρανών core threads
- Ελαχιστοποίηση κόστους επικοινωνίας




case of $n_s = 2p$


| | | | | |
|---------|-------------|-------------|-------------|-------------|
| o o 8 | o 40 o 48 | o 16 o 24 | o 56 o 64 | o o o 32 |
| (8) | (40) (48) | (16) (24) | (56) (64) | (32) |
| (7) | (39) (47) | (15) (23) | (55) (63) | (31) |
| o o 6 7 | 38 39 46 47 | 14 15 22 23 | 54 55 62 63 | o o o 30 31 |
| (6) | (38) (46) | (14) (22) | (54) (62) | (30) |
| (5) | (37) (45) | (13) (21) | (53) (61) | (29) |
| o o 4 5 | 36 37 44 45 | 12 13 20 21 | 52 53 60 61 | o o o 28 29 |
| (4) | (36) (44) | (12) (20) | (52) (60) | (28) |
| (3) | (35) (43) | (11) (19) | (51) (59) | (27) |
| o o 2 3 | 34 35 42 43 | 10 11 18 19 | 50 51 58 59 | o o o 26 27 |
| (2) | (34) (42) | (10) (18) | (50) (58) | (26) |
| (1) | (33) (41) | (9) (17) | (49) (57) | (25) |
| o o o 1 | o 33 o 41 | o 9 o 17 | o 49 o 57 | o o o 25 |



$x_1^{(R)}$



$x_2^{(R)}$



$x_3^{(R)}$




$x_4^{(R)}$



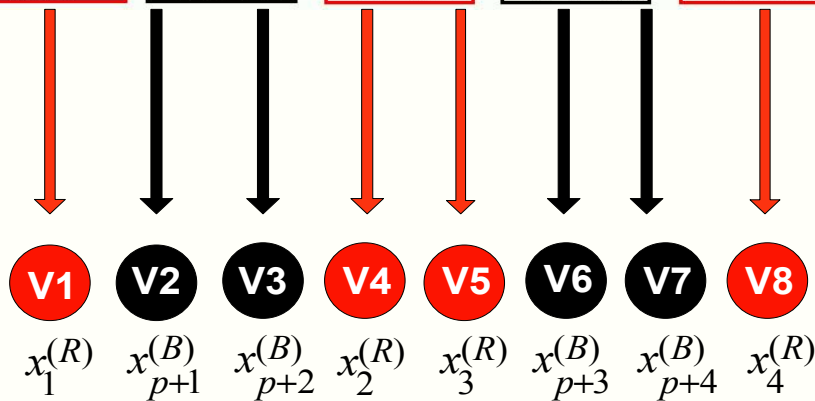
case of $n_s = 2p$

| | | | | |
|---------|-------------|-------------|-------------|-------------|
| o o 8 | o 40 o 48 | o 16 o 24 | o 56 o 64 | o o o 32 |
| (8) | (40) (48) | (16) (24) | (56) (64) | (32) |
| (7) | (39) (47) | (15) (23) | (55) (63) | (31) |
| o o 6 7 | 38 39 46 47 | 14 15 22 23 | 54 55 62 63 | o o o 30 31 |
| (6) | (38) (46) | (14) (22) | (54) (62) | (30) |
| (5) | (37) (45) | (13) (21) | (53) (61) | (29) |
| o o 4 5 | 36 37 44 45 | 12 13 20 21 | 52 53 60 61 | o o o 28 29 |
| (4) | (36) (44) | (12) (20) | (52) (60) | (28) |
| (3) | (35) (43) | (11) (19) | (51) (59) | (27) |
| o o 2 3 | 34 35 42 43 | 10 11 18 19 | 50 51 58 59 | o o o 26 27 |
| (2) | (34) (42) | (10) (18) | (50) (58) | (26) |
| (1) | (33) (41) | (9) (17) | (49) (57) | (25) |
| o o o 1 | o 33 o 41 | o 9 o 17 | o 49 o 57 | o o o 25 |




case of $n_s = 2p$


| | | | | |
|---------|-------------|-------------|-------------|-----------|
| o o 8 | o 40 o 48 | o 16 o 24 | o 56 o 64 | o o o 32 |
| (8) | (40) (48) | (16) (24) | (56) (64) | (32) |
| (7) | (39) (47) | (15) (23) | (55) (63) | (31) |
| o o 6 7 | 38 39 46 47 | 14 15 22 23 | 54 55 62 63 | o o 30 31 |
| (6) | (38) (46) | (14) (22) | (54) (62) | (30) |
| (5) | (37) (45) | (13) (21) | (53) (61) | (29) |
| o o 4 5 | 36 37 44 45 | 12 13 20 21 | 52 53 60 61 | o o 28 29 |
| (4) | (36) (44) | (12) (20) | (52) (60) | (28) |
| (3) | (35) (43) | (11) (19) | (51) (59) | (27) |
| o o 2 3 | 34 35 42 43 | 10 11 18 19 | 50 51 58 59 | o o 26 27 |
| (2) | (34) (42) | (10) (18) | (50) (58) | (26) |
| (1) | (33) (41) | (9) (17) | (49) (57) | (25) |
| o o o 1 | o 33 o 41 | o 9 o 17 | o 49 o 57 | o o o 25 |



| | | | | | | | |
|---|----|----|----|----|----|----|----|
| 8 | 40 | 48 | 16 | 24 | 56 | 64 | 32 |
| 7 | 39 | 47 | 15 | 23 | 55 | 63 | 31 |
| 6 | 38 | 46 | 14 | 22 | 54 | 62 | 30 |
| 5 | 37 | 45 | 13 | 21 | 53 | 61 | 29 |
| 4 | 36 | 44 | 12 | 20 | 52 | 60 | 28 |
| 3 | 35 | 43 | 11 | 19 | 51 | 59 | 27 |
| 2 | 34 | 42 | 10 | 18 | 50 | 58 | 26 |
| 1 | 33 | 41 | 9 | 17 | 49 | 57 | 25 |



Mapping into a fixed size Architecture of N Cores



case of $k = 2p/N$ even


$2k$ virtual threads

$$x_l^{(R)}$$

$$l = (j-1)k + 1, \dots, jk$$

$$x_l^{(B)}$$

$$l = 2p + (j-1)k + 1, \dots, 2p + jk$$



$$j = 1, \dots, N$$



Parallel Schur Complement Iterative Solution

- S1: Solve in parallel on host $D_R \hat{\mathbf{b}}_R = \mathbf{b}_R$
- S2: Send matrices A_3 and A_4 to GPU
- S3: Evaluate in parallel on GPU $\hat{\mathbf{b}}_B = \mathbf{b}_B - H_R \hat{\mathbf{b}}_R$
- S4: Solve in parallel with BiCGSTAB $S \mathbf{x}_B = \hat{\mathbf{b}}_B$
- S5: Evaluate in parallel on GPU $\hat{\mathbf{x}}_B = H_B \mathbf{x}_B$
- S6: Solve in parallel on host $D_R \hat{\mathbf{x}}_R = \hat{\mathbf{x}}_B$
- S7: Evaluate in parallel on host $\mathbf{x}_R = \hat{\mathbf{b}}_R - \hat{\mathbf{x}}_R$



Parallel BICGSTAB

Choose initial approximation $x^{(0)}$ of the solution x_B

$$r^{(0)} = S_B - Sx^{(0)}$$

Choose \hat{r} (usually $\hat{r} = r^{(0)}$)

for $i = 1, 2, \dots$

$$\rho_{i-1} = \hat{r}^T r^{(i-1)}$$

if $\rho_{i-1} = 0$ method fails

if $i = 1$

$$p^{(1)} = r^{(0)}$$

else

$$\beta_{i-1} = \frac{\rho_{i-1}}{\rho_{i-2}} \frac{\alpha_{i-1}}{\omega_{i-1}}$$

$$p^{(i)} = r^{(i-1)} + \beta_{i-1} (p^{(i-1)} - \omega_{i-1} v^{(i-1)})$$

endif

$$v^{(i)} = S p^{(i)}$$

$$\alpha_i = \frac{\rho_{i-1}}{\hat{r}^T v^{(i)}}$$

$$s = r^{(i-1)} - \alpha_i v^{(i)}$$

if $\|s\|$ is small enough then

$$x_B^{(i)} = x_B^{(i-1)} + \alpha_i p^{(i)} \text{ stop}$$

$$t = S s$$

$$\omega_i = \frac{s^T t}{t^T t}$$

$$x_B^{(i)} = x_B^{(i-1)} + \alpha_i p^{(i)} + \omega_i s$$

Check for Convergence

if $\omega_i = 0$ stop

$$r^{(i)} = s - \omega_i t$$

end



Parallel BICGSTAB

Evaluation of $\mathbf{t} = S\mathbf{p}$

S1: Send \mathbf{p} from host to GPU

S2: Evaluate in parallel on GPU $\mathbf{t} = H_B \mathbf{p}$

S3: Send \mathbf{t} from GPU to host

S4: Solve in parallel on host $D_R \mathbf{s} = \mathbf{t}$

S5: Send \mathbf{s} from host to GPU

S6: Evaluate in parallel on GPU $\mathbf{q} = H_R \mathbf{s}$

S7: Send \mathbf{q} from GPU to host


S8: Evaluate in parallel on host $\mathbf{t} = D_B \mathbf{p} - \mathbf{q}$



The Dirichlet Helmholtz Problem

$$u(x, y) = 10\varphi(x)\varphi(y) \quad , \quad (x, y) \in [0, 1] \times [0, 1]$$

$$\text{with } \varphi(x) = (x^2 - x)e^{-100(x-0.1)^2}$$



HP SL390s - Tesla M2070 GPUs

HP SL390s



+ 2 x



6 core Xeon@2.8GHz
24GB memory
Oracle Linux 6.3 x64
PGI 13.5 Fortran
PCI-e gen2 x16

TECHNICAL SPECIFICATIONS

| | Tesla M2070 / M2075 |
|--|---------------------|
| Peak double precision floating point performance | 515 Gigaflops |
| Peak single precision floating point performance | 1030 Gigaflops |
| CUDA cores | 448 |
| Memory size (GDDR5) | 6 GigaBytes |
| Memory bandwidth (ECC off) | 150 GBytes/sec |



The Portland Group



Realization on HP SL390s Tesla GPU machine

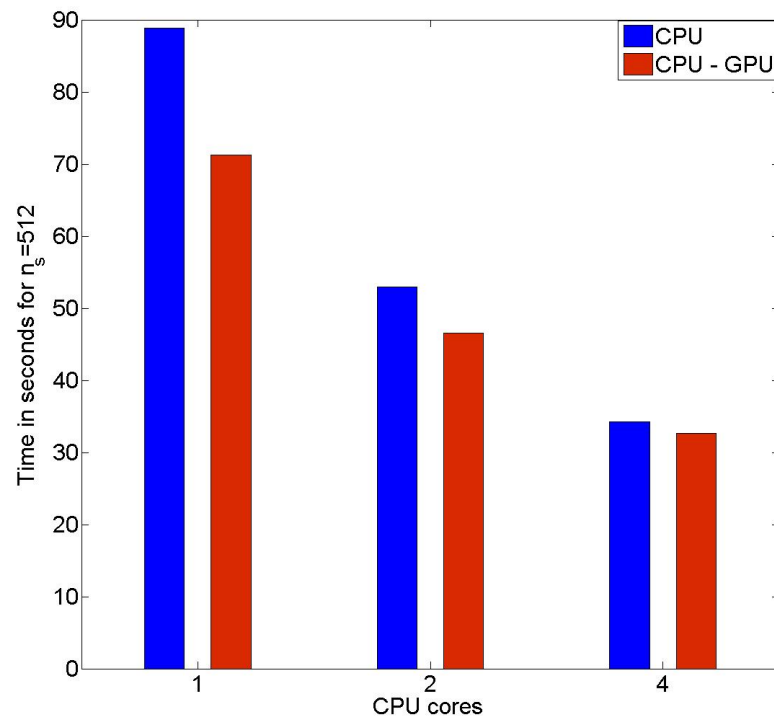

Iterations / Error measurements

| n_s | BiCGSTAB Iterations | $\ \mathbf{b} - \mathbf{A}\mathbf{x}_n \ _2$ |
|-------|------------------------|---|
| 256 | 294 | 6.06e-11 |
| 512 | 589 | 2.85e-11 |
| 1024 | 1161 | 1.39e-11 |
| 2048 | 3726 | 9.59e-12 |





Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης





Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης



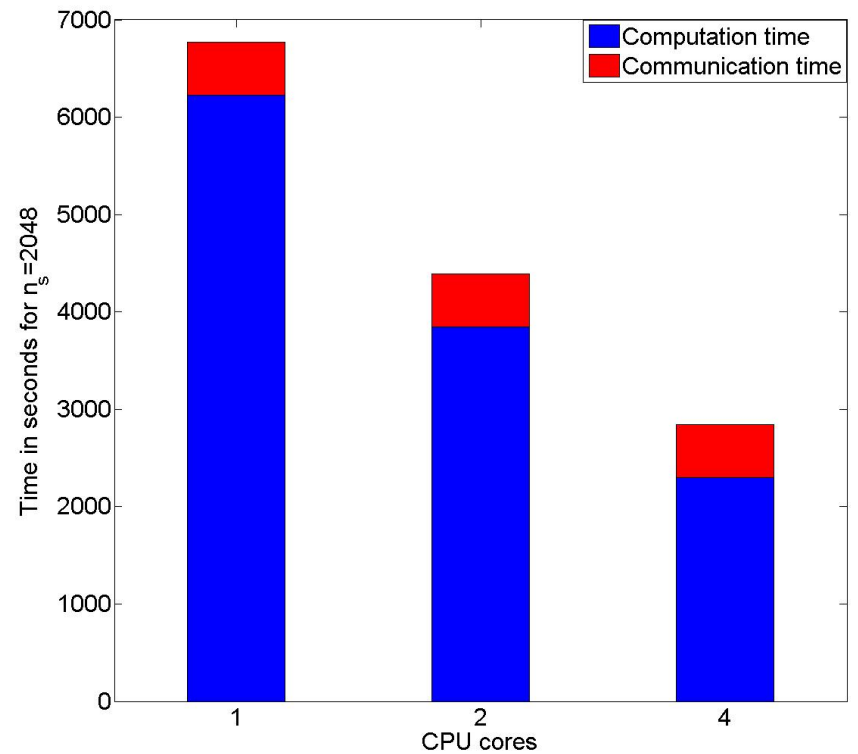

Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης



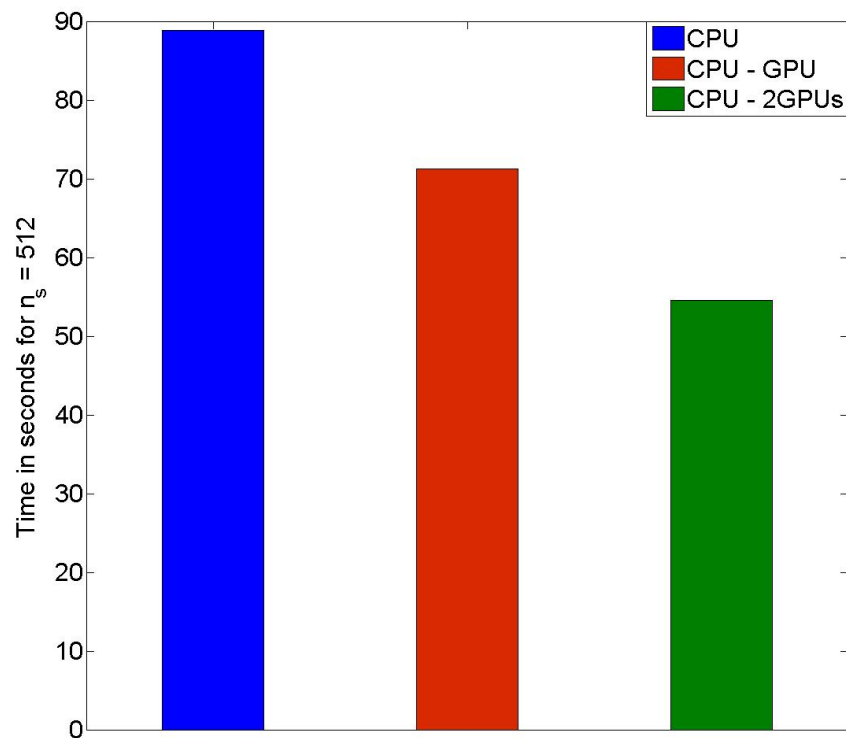
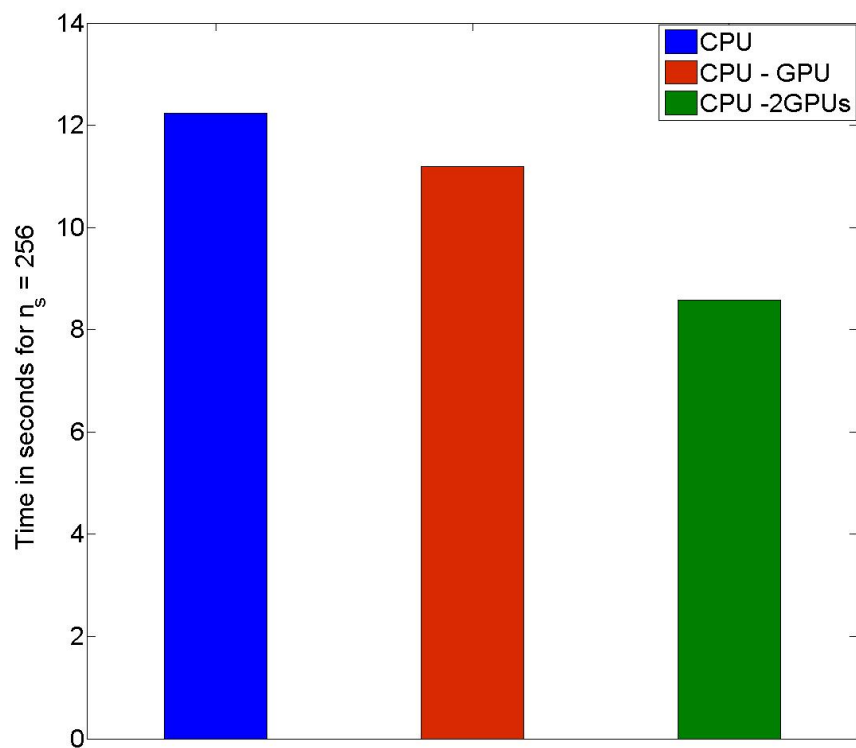
Realization on HP SL390s Tesla GPU machine

Time measurements



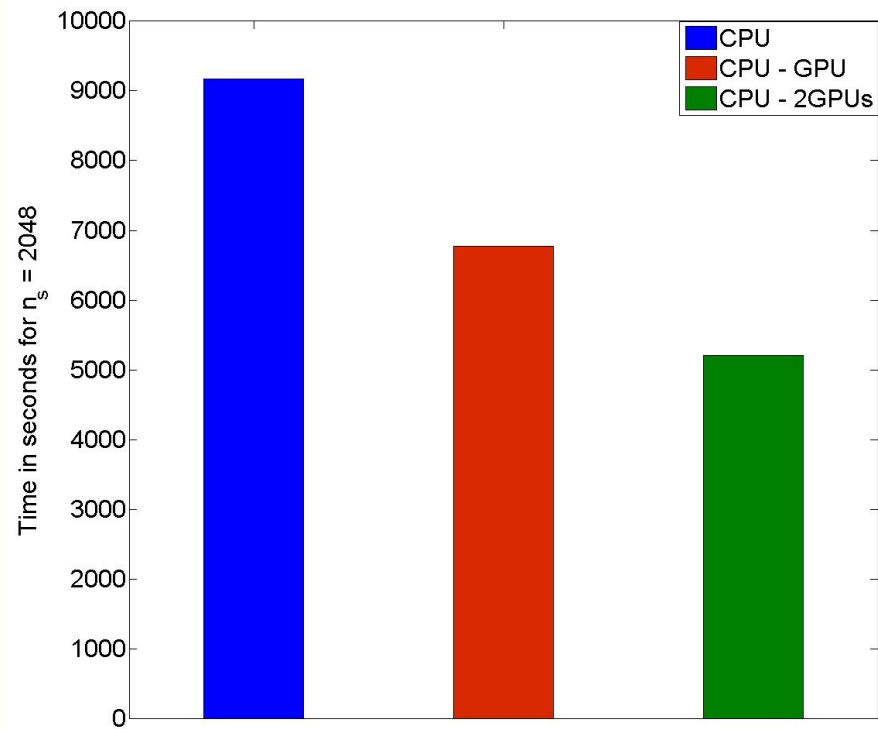
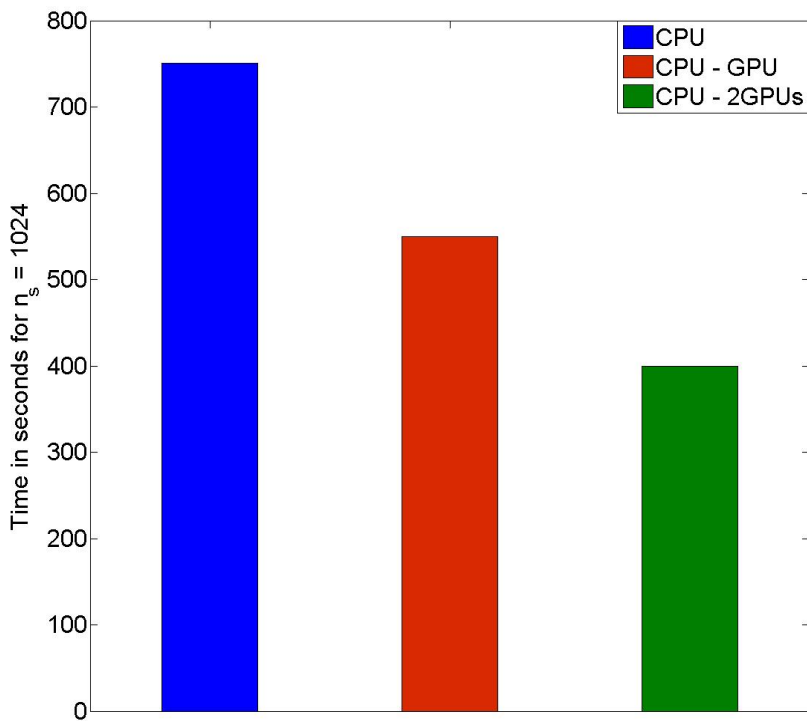
Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης



Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης



Συμπεράσματα

- Σχεδίαση νέου παράλληλου αλγορίθμου της Schur complement μεθόδου με την BiCGSTAB για την Hermite Collocation αριθμητική μέθοδο.
- Υλοποίηση σε αρχιτεκτονικές υπολογισμών Κοινής Μνήμης για multi-core μηχανήματα με επιταχυντές (GPUs).
- Επιτάχυνση μέχρι 30%.



Επόμενα βήματα

- Κατασκευή αποδοτικού αλγορίθμου της μεθόδου Discontinuous Hermite Collocation για αρχιτεκτονικές με Multiprocessor /Grid μηχανήματα με επιταχυντές υπολογισμών.



Semi-Implicit Discontinuous Hermite Collocation Method

DHC system of ODEs


$$(C_x^0 \otimes C_y^0) \dot{\mathbf{a}} = (D_x C_x^2 \otimes C_y^0) \mathbf{a} + (D_x C_x^0 \otimes C_y^2) \mathbf{a}$$

where C_x^* and C_y^* denotes the 1d Discontinuous and Continuous Collocation matrices

$$D_x = \text{diag}(\gamma, \dots, \gamma, 1, \dots, 1, \gamma, \dots, \gamma)$$

$$A_0 := C_x^0 \otimes C_y^0$$

$$B := D_x C_x^2 \otimes C_y^0 + D_x C_x^0 \otimes C_y^2$$




$$A_0 \dot{\mathbf{a}} = B \mathbf{a}$$



Semi-Implicit Discontinuous Hermite Collocation Method

$$A_0 \dot{\mathbf{a}} = B \mathbf{a}$$



DHC method coupled
with an optimal 2-step
and 3rd order DIRK

$$A \mathbf{a}^{(n,1)} = A_0 \mathbf{a}^{(n)}$$

$$A \mathbf{a}^{(n,2)} = A_0 \mathbf{a}^{(n)} + \tau(1 - 2\lambda)B \mathbf{a}^{(n,1)} .$$

$$A_0 \mathbf{a}^{(n+1)} = A_0 \mathbf{a}^{(n)} + \frac{\tau}{2} [B \mathbf{a}^{(n,1)} + B \mathbf{a}^{(n,2)}]$$

where $A := A_0 - \tau\lambda B$, τ is the time spacing and $\lambda = \frac{1}{2} + \frac{\sqrt{3}}{2}$



Semi-Implicit Discontinuous Hermite Collocation Method

```
INPUT  $\mathbf{a}_{old}$  ,  $B$  ,  $A_0$  ,  $A$  ,  $A_b$   
for  $t = \tau$  to  $t_{max}$  with step  $\tau$  do  
  compute  $\mathbf{a}_0 = A_0 \mathbf{a}_{old}$   
  if  $t \leq 2\tau$  then  
    solve  $A_b \mathbf{a}_{new} = \mathbf{a}_0$  with BiCGSTAB  
  else  
    solve  $A \mathbf{a}_1 = \mathbf{a}_0$  with BiCGSTAB  
    compute  $\mathbf{a}_3 = \mathbf{a}_0 - \tau \frac{\sqrt{3}}{3} B \mathbf{a}_1$   
    solve  $A \mathbf{a}_2 = \mathbf{a}_3$  with BiCGSTAB  
    compute  $\mathbf{a}_2 = \mathbf{a}_0 + \frac{\tau}{2} B (\mathbf{a}_1 + \mathbf{a}_2)$   
    solve  $A_0 \mathbf{a}_{new} = \mathbf{a}_2$  with BiCGSTAB  
  endif  
  compute  $\mathbf{a}_{old} = \mathbf{a}_{new}$   
enddo
```



Preconditioned BiCGSTAB Method...



```
Compute  $r^{(0)} = b - Ax^{(0)}$  for some initial guess  $x^{(0)}$ 
Choose  $\tilde{r}$  (for example,  $\tilde{r} = r^{(0)}$ )
for  $i = 1, 2, \dots$ 
   $\rho_{i-1} = \tilde{r}^T r^{(i-1)}$ 
  if  $\rho_{i-1} = 0$  method fails
  if  $i = 1$ 
     $p^{(i)} = r^{(i-1)}$ 
  else
     $\beta_{i-1} = (\rho_{i-1}/\rho_{i-2})(\alpha_{i-1}/\omega_{i-1})$ 
     $p^{(i)} = r^{(i-1)} + \beta_{i-1}(p^{(i-1)} - \omega_{i-1}v^{(i-1)})$ 
  endif
  solve  $M\hat{p} = p^{(i)}$ 
   $v^{(i)} = A\hat{p}$ 
   $\alpha_i = \rho_{i-1}/\tilde{r}^T v^{(i)}$ 
   $s = r^{(i-1)} - \alpha_i v^{(i)}$ 
  check norm of  $s$ ; if small enough: set  $x^{(i)} = x^{(i-1)} + \alpha_i \hat{p}$  and stop
  solve  $M\hat{s} = s$ 
   $t = A\hat{s}$ 
   $\omega_i = \tilde{r}^T s / \tilde{r}^T t$ 
   $x^{(i)} = x^{(i-1)} + \alpha_i \hat{p} + \omega_i \hat{s}$ 
   $r^{(i)} = s - \omega_i t$ 
  check convergence; continue if necessary
  for continuation it is necessary that  $\omega_i \neq 0$ 
end
```

$$A \sim M = \text{ilu}(A)$$





Eigenvalue Distribution ...

Matrix **A**



Eigenvalue Distribution ...

Matrix A_0



Parallel BICGSTAB

- ✓ All basic linear algebra operations are performed on the GPU
- ✓ The preconditioning procedure $Mz = t$ with $M=LU$ is performed on the CPU

Step 1 : GPU sends to CPU vector t



Step 2 : CPU solution of $Ly = t$

Step 3 : CPU solution of $Uz = y$

Step 4 : CPU sends to GPU vector z



The test problem ...



HP SL390s - Tesla M2070 GPUs

HP SL390s



+



6 core Xeon@2.8GHz
24GB memory
Oracle Linux 6.3 x64
PGI 14.5 Cuda Fortran
Cuda toolkit 6.0
PCI-e gen2 x16

TECHNICAL SPECIFICATIONS

| | Tesla M2070 / M2075 |
|--|---------------------|
| Peak double precision floating point performance | 515 Gigaflops |
| Peak single precision floating point performance | 1030 Gigaflops |
| CUDA cores | 448 |
| Memory size (GDDR5) | 6 GigaBytes |
| Memory bandwidth (ECC off) | 150 GBytes/sec |

The Portland Group



NVIDIA® CUDA™
Parallel Programming and Computing Platform



Development tools ...

- MatLab R2012b
- PGI 14.5 Cuda Fortran
Cuda toolkit 6.0 - cuBLAS - cuSRARSE
for GPU operations
SparseKit for CPU operations



Development tools ...

```
module bicgstabGPU
contains
  subroutine bicgstabGPU(n,x,b,A,L,U,iA,iL,jA,jL,jU,
+      nzA,descrA,handle,error,r,rh,pi,ph,
+      t,s,sh,ui,istep,temp,ttt)
  implicit real*8 (a-h,o-z)
  real*8 ttt(n),L(*),U(*),temp(n)
  integer*8 handle,descrA,iA,jA,A,
+      x,b,r,rh,pi,ph,s,sh,ui,t
  integer cusparse_dcsrnmv,iL(n+1),jL(*),jU(*),
+  cuda_memcpy_c2fort_real,cuda_memcpy_fort2c_real

  imaxstep=istep
  tol=error
  istep=0
  dnrmb=cublas_dnrmb2(n,b,1)
  call cublas_dcopy(n,b,1,x,1)
  istatus=cusparse_dcsrnmv(handle,0,n,n,nzA,1.0d0,descrA,A,iA,jA,
+      x,0.0d0,r)
```



Applied Mathematics
and Computers Laboratory



Development tools ...

999

```
call cublas_dscal(n,-1.0d0,r,1)
call cublas_daxpy(n,1.0d0,b,1,r,1)
call cublas_dcopy(n,r,1,rh,1)
continue
istep=istep+1
if (istep.gt.1) roip2=roip1
  roip1=cublas_ddot(n,rh,1,r,1)
if (istep.eq.1) then
  call cublas_dcopy(n,r,1,pi,1)
else
  bi=(roip1/roip2)*(ai/wi)
  call cublas_daxpy(n,-wi,ui,1,pi,1)
  call cublas_dcopy(n,r,1,t,1)
  call cublas_daxpy(n,bi,pi,1,t,1)
  call cublas_dcopy(n,t,1,pi,1)
endif
endif
```

C-----



Development tools ...

```
icudaStat = cuda_memcpy_c2fort_real(ttt,pi,n*8,2)
call lsol(n,temp,ttt,L,jL,iL)
call udsol(n,ttt,temp,U,jU)
icudaStat3 = cuda_memcpy_fort2c_real(ph,ttt,n*8,1)
```

C

```
istatus=cusparsolve(handle,0,n,n,nzA,1.0d0,descrA,A,iA,jA,
+      ph,0.0d0,ui)
ai=roip1/cublas_ddot(n,rh,1,ui,1)
call cublas_dcopy(n,r,1,s,1)
call cublas_daxpy(n,-ai,ui,1,s,1)
if (cublas_dnrm2(n,s,1).lt.1.0d-25) then
call cublas_daxpy(n,ai,ph,1,x,1)
  print*,'BiCGSTAB incomplete'
else
```

C

```
icudaStat = cuda_memcpy_c2fort_real(ttt,s,n*8,2)
call lsol(n,temp,ttt,L,jL,iL)
call udsol(n,ttt,temp,U,jU)
icudaStat3 = cuda_memcpy_fort2c_real(sh,ttt,n*8,1)
```



Development tools ...

```
istatus=cusparsolve(handle,0,n,n,nzA,1.0d0,descrA,A,iA,jA,  
+ sh,0.0d0,t)  
wi=cublas_ddot(n,t,1,s,1)/cublas_ddot(n,t,1,t,1)  
call cublas_daxpy(n,ai,ph,1,x,1)  
call cublas_daxpy(n,wi,sh,1,x,1)  
call cublas_daxpy(n,-wi,t,1,s,1)  
call cublas_dcopy(n,s,1,r,1)  
dnrmr=cublas_dnrm2(n,s,1)  
error=dnrmr/dnrmb  
c print*,error,istep  
if (wi.ne.0.0d0.and.error.gt.tol.and.istep.lt.imaxstep)  
+ goto 999  
endif  
return  
end  
end module
```




Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης

Finite Elements : 400

Unknowns : 1.600

Dofs : 6.400




Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης

Finite Elements : 1.600

Unknowns : 6.400

Dofs : 25.600




Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης

Finite Elements : 6.400

Unknowns : 25.600

Dofs : 102.400




Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης

Finite Elements : 25.600

Unknowns : 102.400

Dofs : 409.600



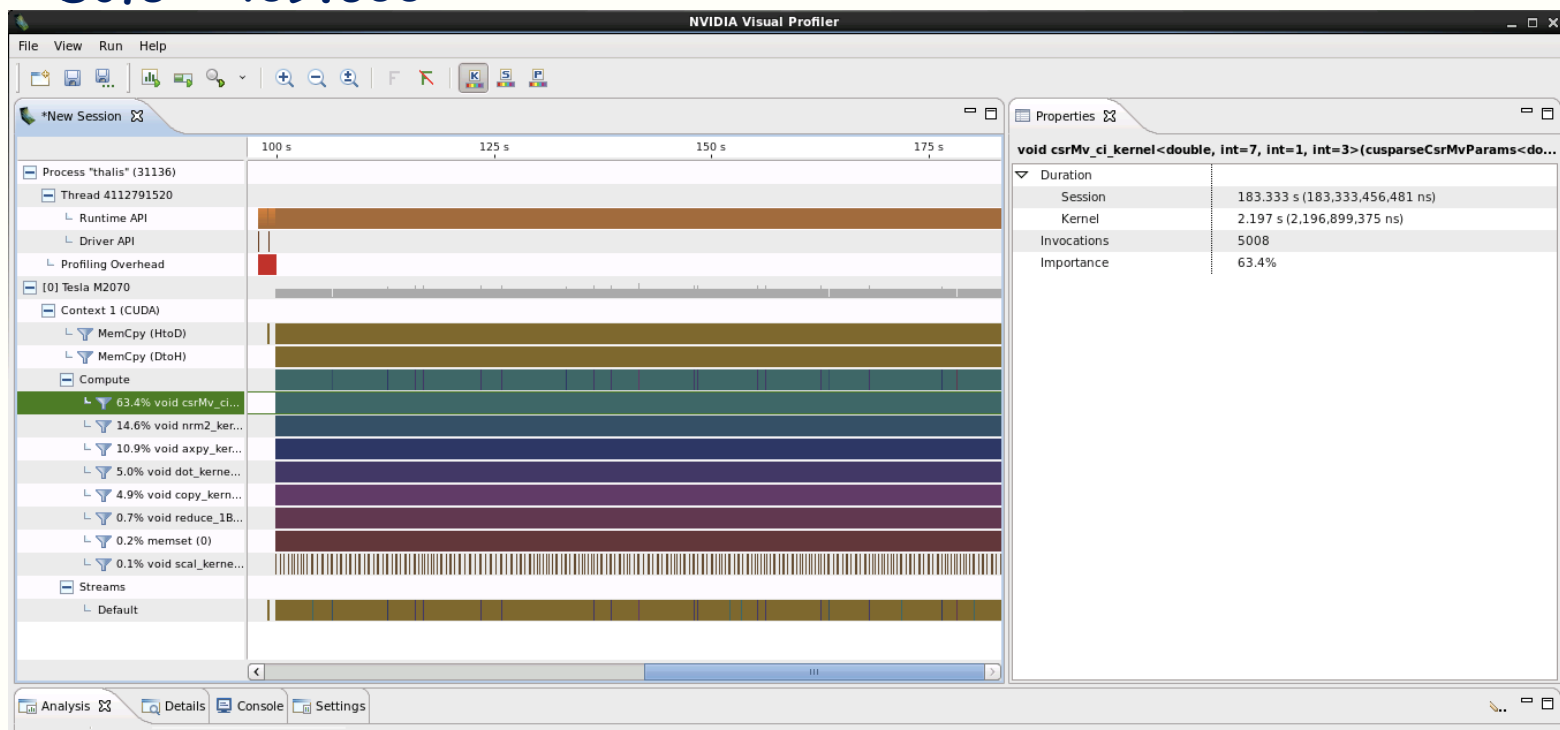
Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης

Finite Elements : 25.600

Unknowns : 102.400

Dofs : 409.600




Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης

Finite Elements : 25.600

Unknowns : 102.400

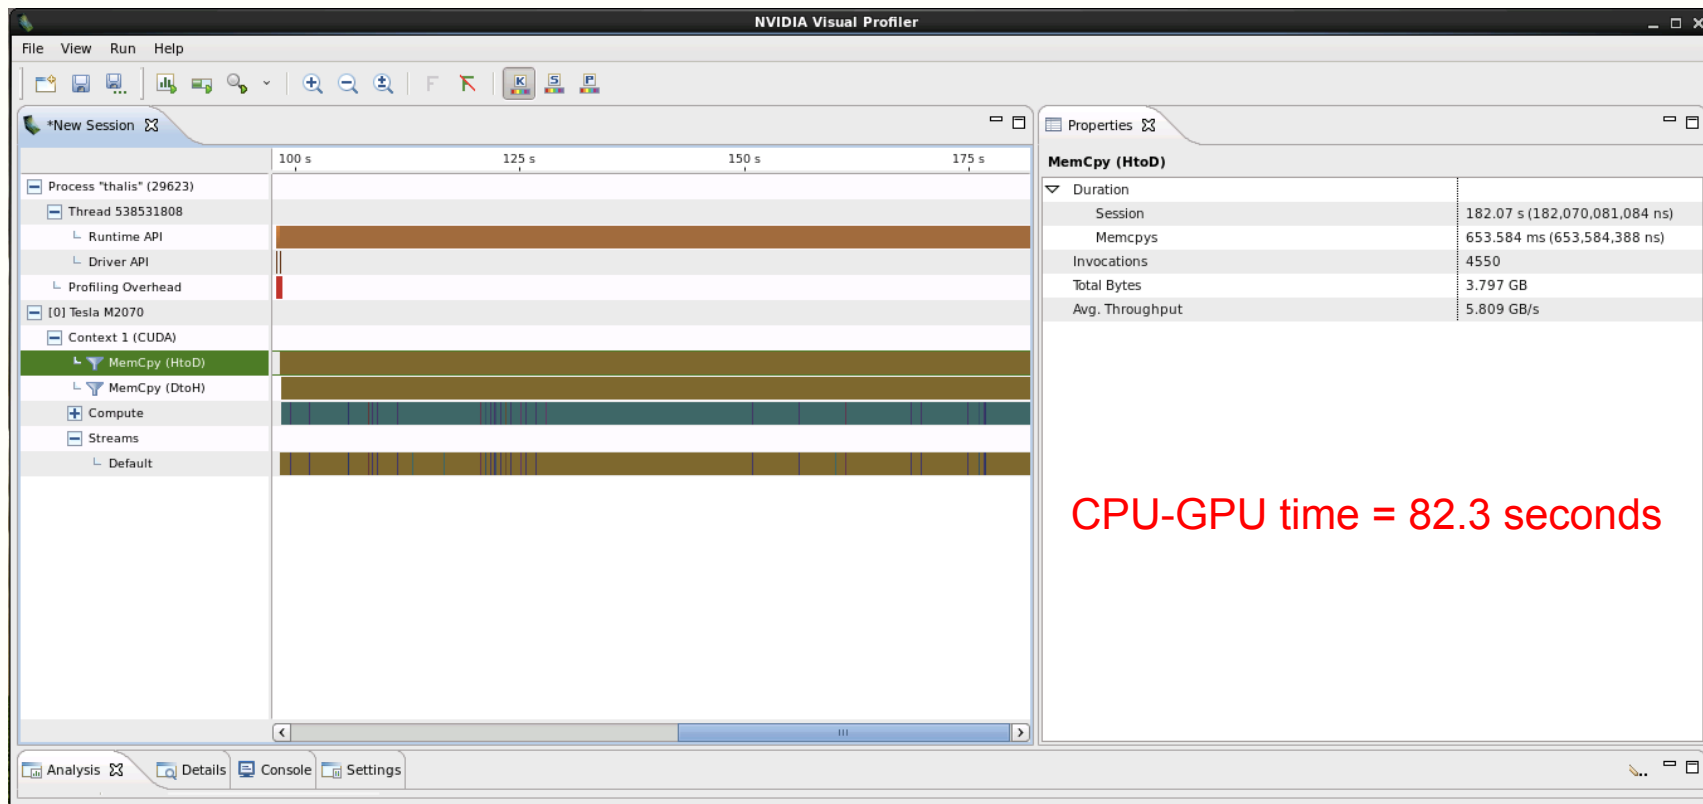
Dofs : 409.600



Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης

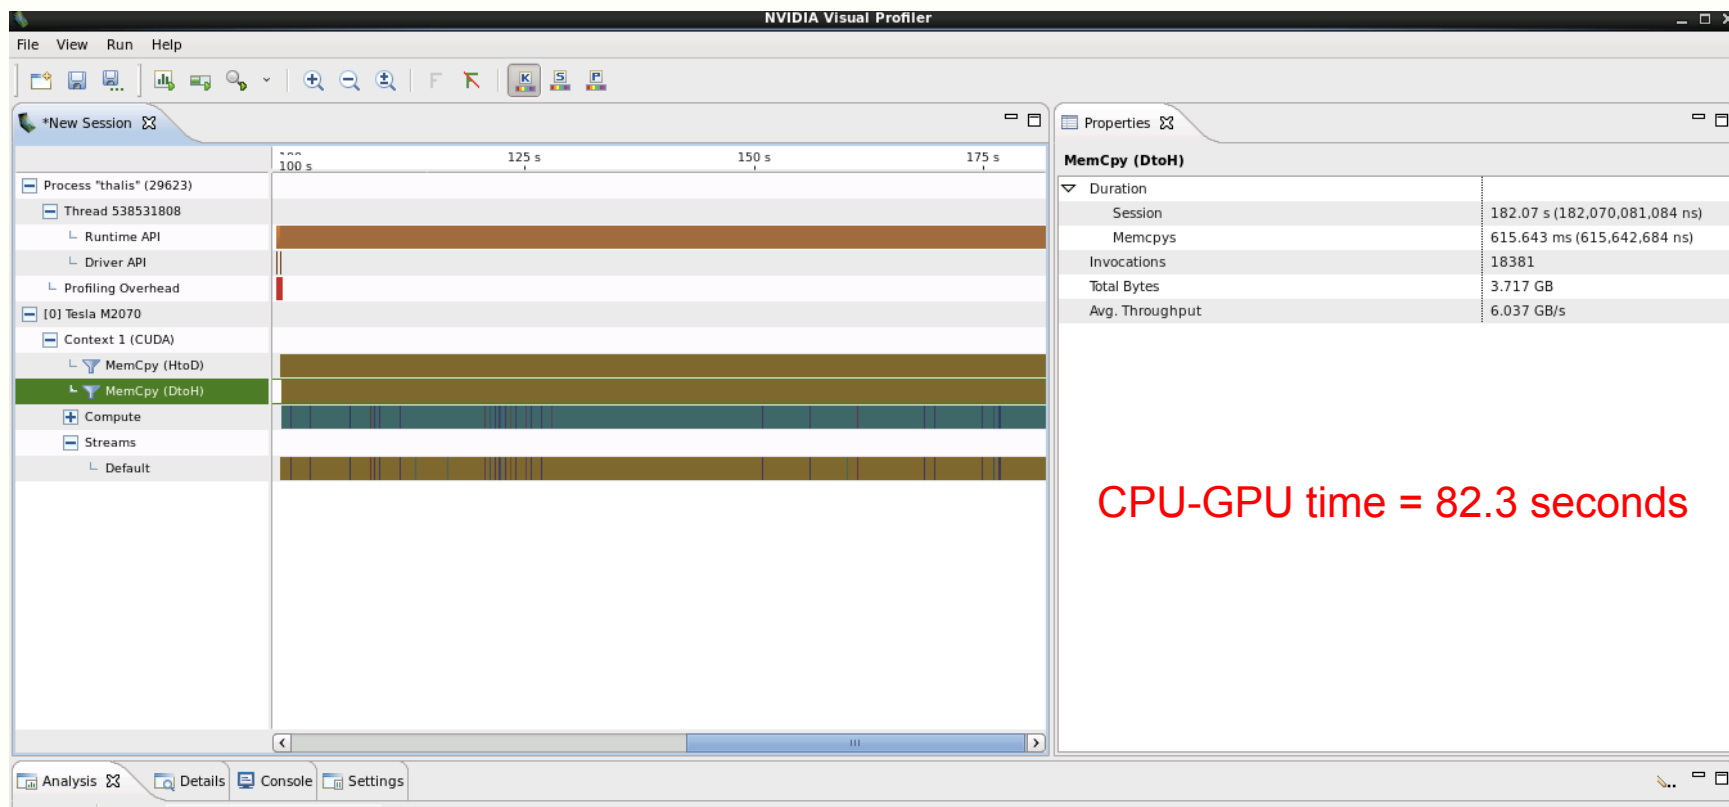
Finite Elements : 25.600 Unknowns : 102.400 Dofs : 409.600



Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις χρόνου εκτέλεσης

Finite Elements : 25.600 Unknowns : 102.400 Dofs : 409.600



Συμπεράσματα

- Σχεδίαση νέου παράλληλου αλγορίθμου της αριθμητικής μεθόδου Discontinuous Hermite Collocation.
- Υλοποίηση σε αρχιτεκτονικές υπολογισμών Κοινής Μνήμης για μηχανήματα με επιταχυντές (GPUs) .
- Παρατηρήσαμε επιτάχυνση $2x$ για πυκνές διακριτοποιήσεις σε σχέση με την MatLab multithread υλοποίηση.



Semi-Implicit Discontinuous Hermite Collocation Method στην περίπτωση θεραπείας

Create matrices $A_0, A_u, A_c, A_r, A_{rc}, B_u, B_c, B_r,$
 B_{rc} and initial vector \mathbf{a}_{old}

for $t = dt$ **to** t_{max} **with time step** dt

Compute $\mathbf{a}_0 = A_0 \mathbf{a}_{old}$

if *radiotherapy* **then**

Solve $A_r \mathbf{a}_1 = \mathbf{a}_0$

Compute $\mathbf{a}_0 = \mathbf{a}_0 - dt \frac{\sqrt{3}}{3} B_r \mathbf{a}_1$

Solve $A_r \mathbf{a}_2 = \mathbf{a}_0$

Compute $\mathbf{a}_2 = \mathbf{a}_0 + \frac{dt}{2} B_r (\mathbf{a}_1 + \mathbf{a}_2)$

else if *chemotherapy* **then**

Solve $A_c \mathbf{a}_1 = \mathbf{a}_0$

Compute $\mathbf{a}_0 = \mathbf{a}_0 - dt \frac{\sqrt{3}}{3} B_c \mathbf{a}_1$

Solve $A_c \mathbf{a}_2 = \mathbf{a}_0$

Compute $\mathbf{a}_2 = \mathbf{a}_0 + \frac{dt}{2} B_c (\mathbf{a}_1 + \mathbf{a}_2)$

else if *radiotherapy and chemotherapy* **then**

Solve $A_{rc} \mathbf{a}_1 = \mathbf{a}_0$

Compute $\mathbf{a}_0 = \mathbf{a}_0 - dt \frac{\sqrt{3}}{3} B_{rc} \mathbf{a}_1$

Solve $A_{rc} \mathbf{a}_2 = \mathbf{a}_0$

Compute $\mathbf{a}_2 = \mathbf{a}_0 + \frac{dt}{2} B_{rc} (\mathbf{a}_1 + \mathbf{a}_2)$

else

Solve $A_u \mathbf{a}_1 = \mathbf{a}_0$

Compute $\mathbf{a}_0 = \mathbf{a}_0 - dt \frac{\sqrt{3}}{3} B_u \mathbf{a}_1$

Solve $A_u \mathbf{a}_2 = \mathbf{a}_0$

Compute $\mathbf{a}_2 = \mathbf{a}_0 + \frac{dt}{2} B_u (\mathbf{a}_1 + \mathbf{a}_2)$

endif

Solve $A_0 \mathbf{a}_{new} = \mathbf{a}_2$



Compute $\mathbf{a}_{old} = \mathbf{a}_{new}$

endfor





Eigenvalue Distribution ...

Matrix A_0



Eigenvalue Distribution ...

Matrix A_{rc}



Υλοποίηση σε HP SL390s Tesla GPU μηχανήμα

Μετρήσεις Speedup

| n_s | MatLab - CPU | MatLab - GPU | CPU - GPU |
|-----------|--------------|--------------|-----------|
| 40 x 40 | 5.2 | 5.6 | 1.1 |
| 80 x 80 | 5.9 | 7.9 | 1.4 |
| 160 x 160 | 5.3 | 6.6 | 1.3 |



Δημοσιεύσεις

- Athanasakis I., Vilanakis N. and Mathioudakis E., "Solving Discontinuous Collocation Equations for a Class of Brain Tumor Models on GPUs", Lecture Notes in Engineering and Computer Science, (2217), pp. 529-534, IAENG, 2015.
- Vilanakis N, Mathioudakis E., "Parallel iterative solution of the Hermite Collocation equations on GPUs II", Journal of Physics: Conference Series, vol. 490, 012097, 2014.
- Mathioudakis E., Vilanakis N., Papadopoulou E., Saridakis Y., "Parallel iterative solution of the hermite collocation equations on GPUs", Lecture Notes in Engineering and Computer Science, vol. 2205, IAENG, pp. 1281-1286, 2013.

