Row scaling as a preconditioner for certain nonsymmetric linear systems with discontinuous coefficients

Dan Gordon

Computer Science
University of Haifa

Rachel Gordon

Aerospace Engineering
The Technion

Discontinuous coefficients

- Discontinuous coefficients arise in several types of linear systems
- Example: PDEs model certain physical phenomena in heterogeneous media
- Common approach: Domain Decomposition (DD)
- Problem: May be difficult when:
  - Unstructured grid
  - Complicated boundaries between subdomains

Row and/or column scaling is well-known

- van der Sluis 1969: effect of scaling on condition number
- Widlund 1971: "well-scaled ADI methods give good rates of convergence when the coefficients of elliptic problems vary very much in magnitude"
- Duff & van der Vorst 1998: "on vector machines, diagonal scaling is often competitive with other approaches"
- Graham & Hagger 1999: "diagonal scaling has been observed in practical computations to be very effective as a preconditioner for problems with discontinuous coefficients"
- Gambolati et al. 2003: use the least square logarithm scaling on rows and columns for geomechanics problems with discontinuous coefficients

This work

- Type of problems considered:
  - nonsymmetric
  - discontinuous coefficients
  - small to moderate convection terms
- We consider only a particular type of row scaling
- **Geometric row scaling** – GRS(p): divide each equation by the L_p-norm of its vector of coefficients
- We actually use only GRS(2), but GRS(1) gives similar results
- GRS will be used to denote GRS(2)
- Note: on symmetric problems, diagonal scaling is usually 2-sided to preserve symmetry

GRS(2) is inherent in some algorithms:

- Kaczmarz (1937) – inherently sequential
- Cimmino (1938) – inherently parallel
- CGMN: CG acceleration of Kaczmarz (Björck & Elfving, 1979)
- CG acceleration of Cimmino $\equiv$ CGNR + GRS
- CARP: a block-parallel version of Kaczmarz (G&G 2005)
- CARP-CG: CG acceleration of CARP (G&G)

Problem 0

- Suggested by Nachtigal, Reddy & Trefethen, 1992
- System matrix: $A = \mathrm{diag}(1, 4, 9, \ldots, n^2)$
- Algorithms CGNR, CGS and GMRES make little progress until the nth iteration.
- But...
- Apply GRS – and you get I.

Problem 1

- Problem F2DB, from Saad's book, 2003.
- 2-dimensional PDE on the unit square:

$$-\frac{\partial}{\partial x}(au_x) - \frac{\partial}{\partial y}(bu_y) + \frac{\partial(du)}{\partial x} + \frac{\partial(eu)}{\partial y} = h,$$

$$a(x,y) = b(x,y) = \begin{cases} 10^3 & \text{if } \frac{1}{4} < x,\ y < \frac{3}{4}, \\ 1 & \text{otherwise} \end{cases}$$

- $d(x,y) = 10(x+y)$, $e(x,y) = 10(x-y)$.
- Dirichlet boundary condition: $u = 0$; discretization: 128×128
- RHS h is immaterial: $b = Ae$, where A is the system matrix and $e = (1, \ldots, 1)^T$.

Problem 1: basic eigenvalue information

- Discretization: 40×40

Matrix	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\max}/\lambda_{\min}$	No. of e-values in first of 100 intervals
Original	1.87E-2	7.96E+3	4.25E+5	1126
With GRS	7.07E-6	1.78E+0	2.52E+5	4
Cont. coef.	1.17E+1	8.00E+3	6.81E+2	8

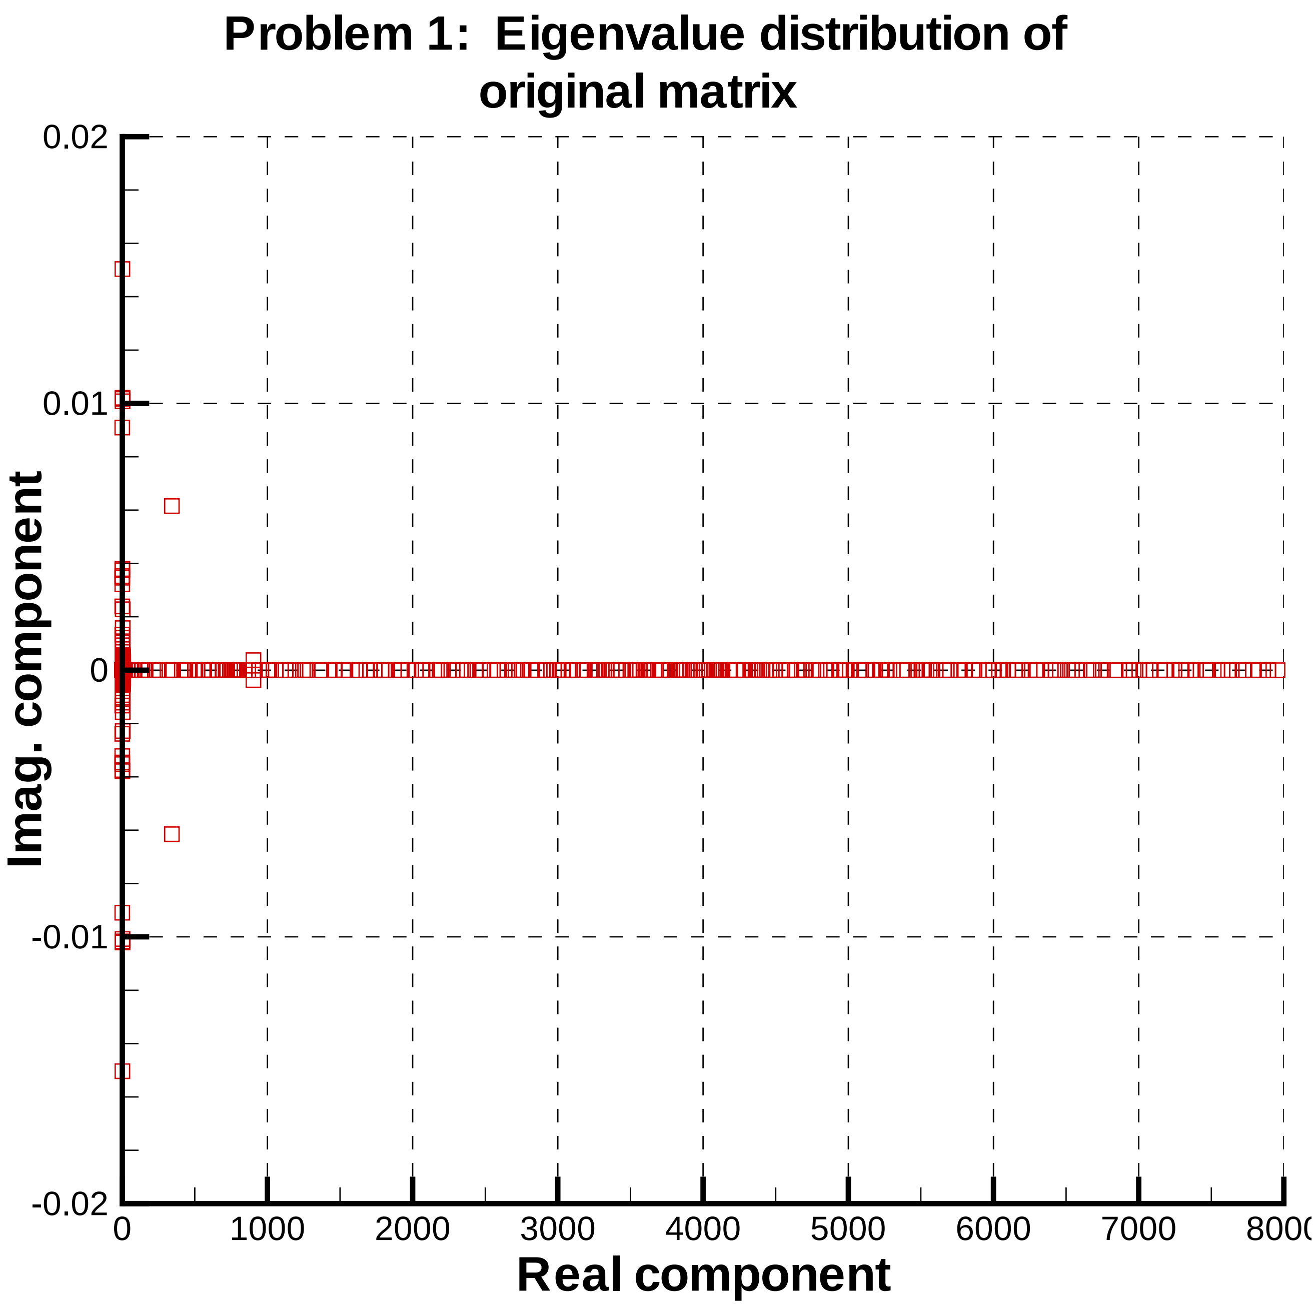
Problem 1: Eigenvalue distribution of original matrix
0.02
0.01
0
-0.01
-0.02
Imag. component
0
1000
2000
3000
4000
5000
6000
7000
800(
Real component

Eigenvalue distribution

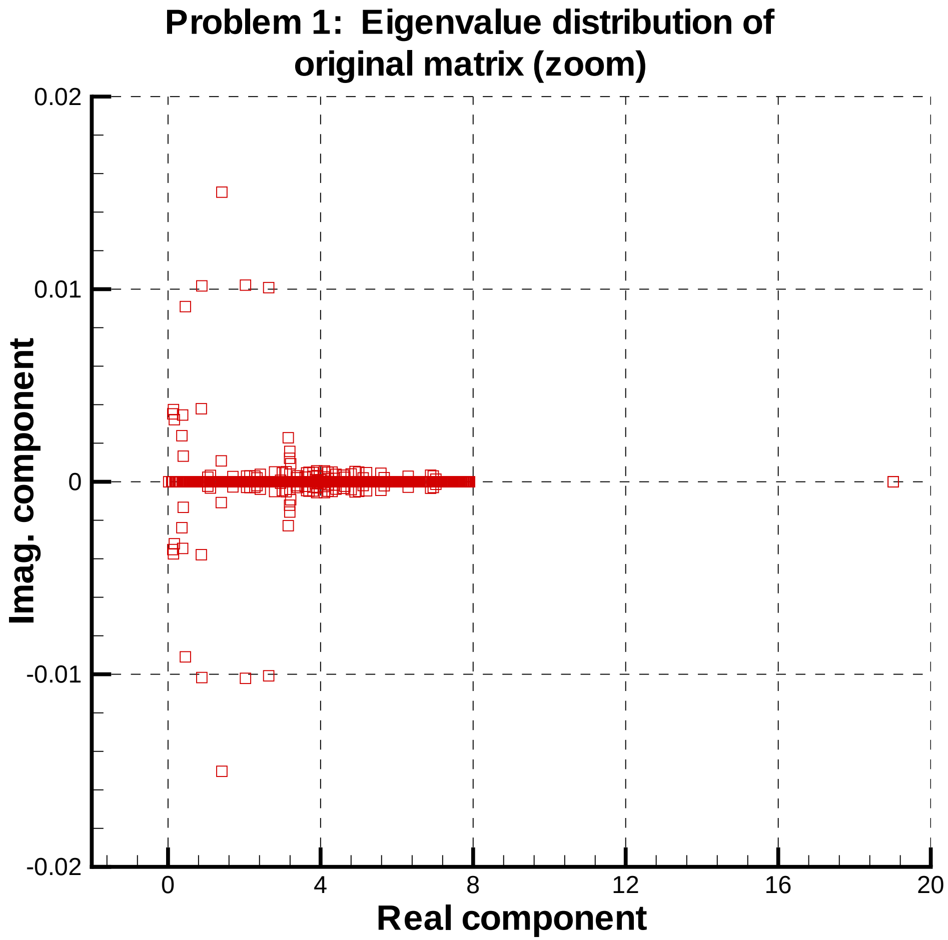

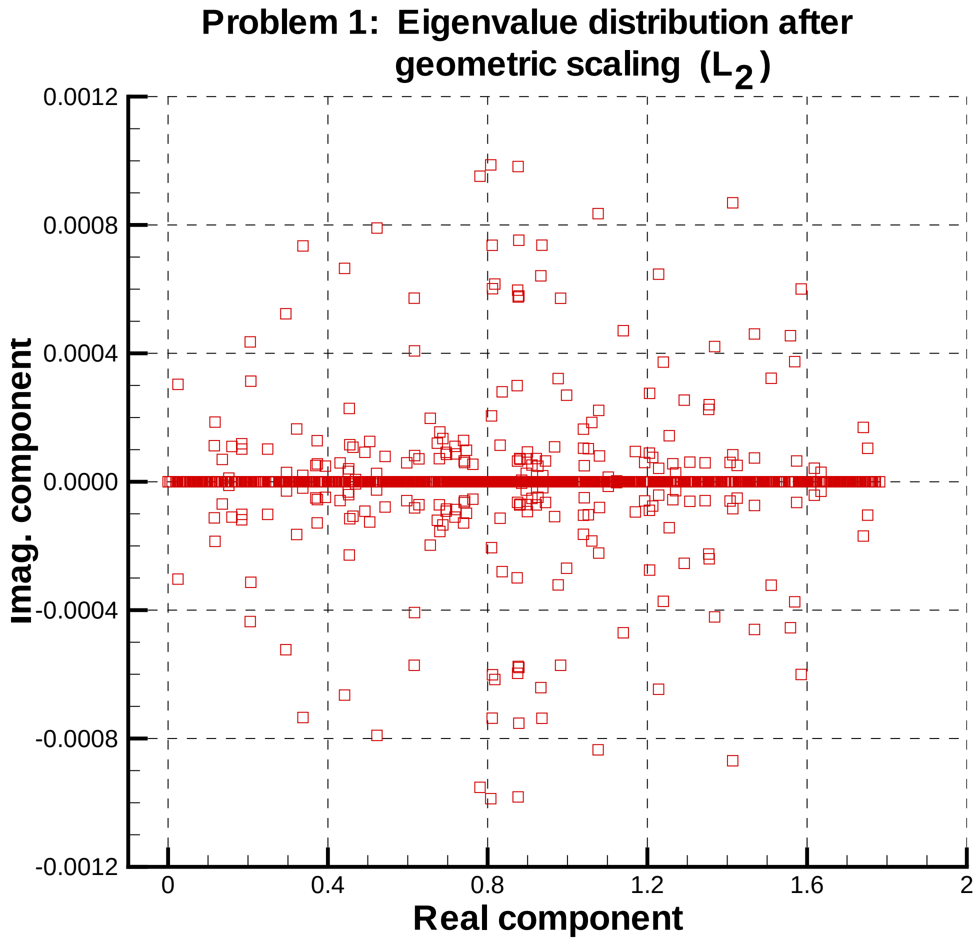

Problem 1: no. of iterations & runtimes (sec.)

Method	rel-res $=10^{-4}$	rel-res $=10^{-7}$	rel-res $=10^{-10}$
Bi-CGSTAB with GRS	no conv. 91 (0.30)	no conv. 299 (0.99)	no conv. 361 (1.19)
Bi-CGSTAB+ILU(0) with GRS	31 (0.23) 30 (0.23)	107 (0.67) 90 (0.59)	142 (0.88) 130 (0.81)
GMRES	converged to 3.8×10^{-2}		
with GRS	265 (0.85)	converged to 1.1×10^{-5}	
GMRES+ILU(0)	converged to 3.9×10^{-3}		
with GRS	39 (0.23)	converged to 1.1×10^{-5}	

Problem 2

- Based on Example 2 from van der Vorst 1992 (Bi-CGSTAB paper)
- with additional convection terms, making it nonsymmetric
- PDE on the unit square: $-\dfrac{\partial(D(x,y)u_x)}{\partial x} - \dfrac{\partial(D(x,y)u_y)}{\partial y} + au_x + bu_y = 1$
- Convection terms: $a = b = 200$; discretization: 150×150

Additional details:

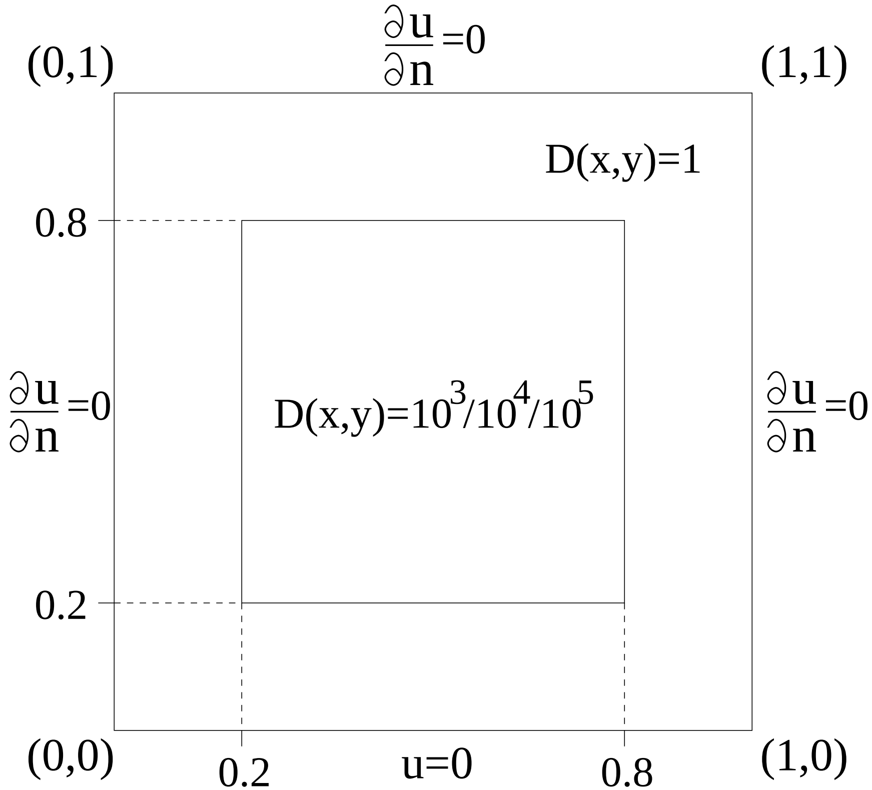

Problem 2: basic eigenvalue data

- Discretization: 40×40

Matrix	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\max}/\lambda_{\min}$	e-values around $x=0$
Original $(D=10^3)$	3.33E-1	1.34E+7	4.02E+7	1058
With GRS	9.15E-5	1.78E+0	1.95E+4	8
Continuous coef.	1.48E-2	1.34E+7	9.09E+8	130

Problem 2: eigenvalue distribution

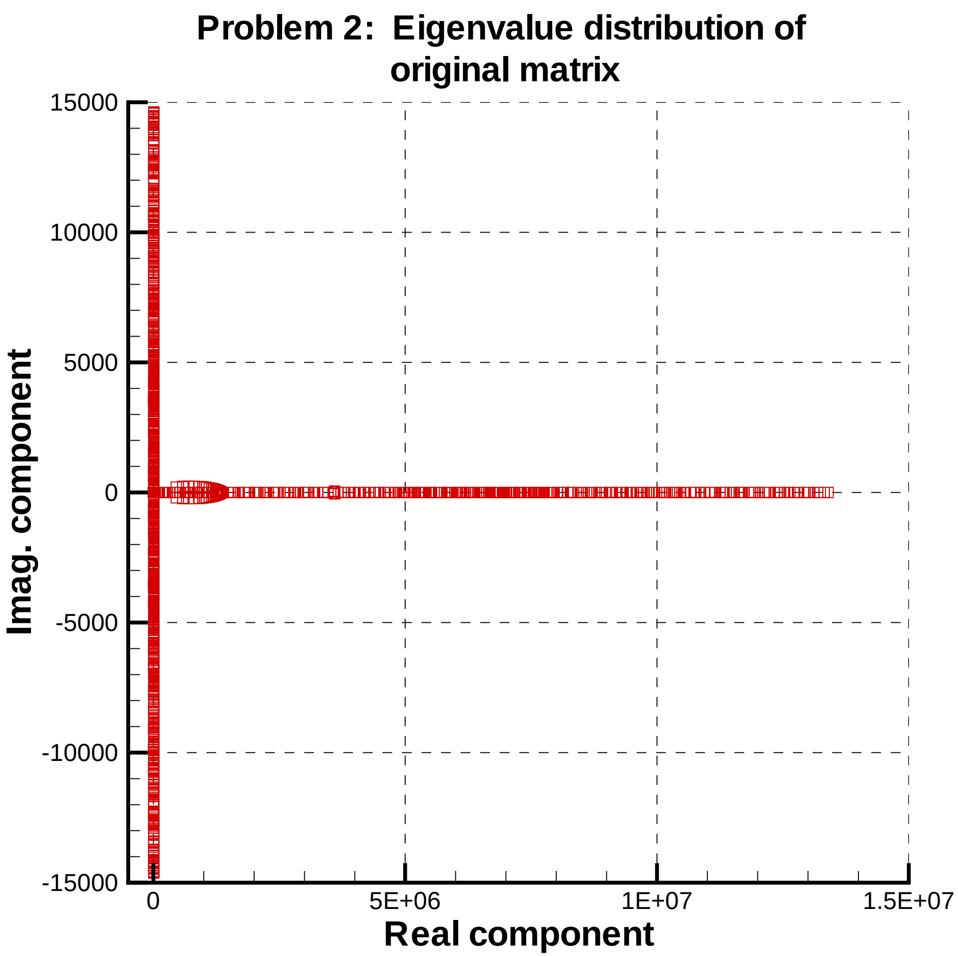

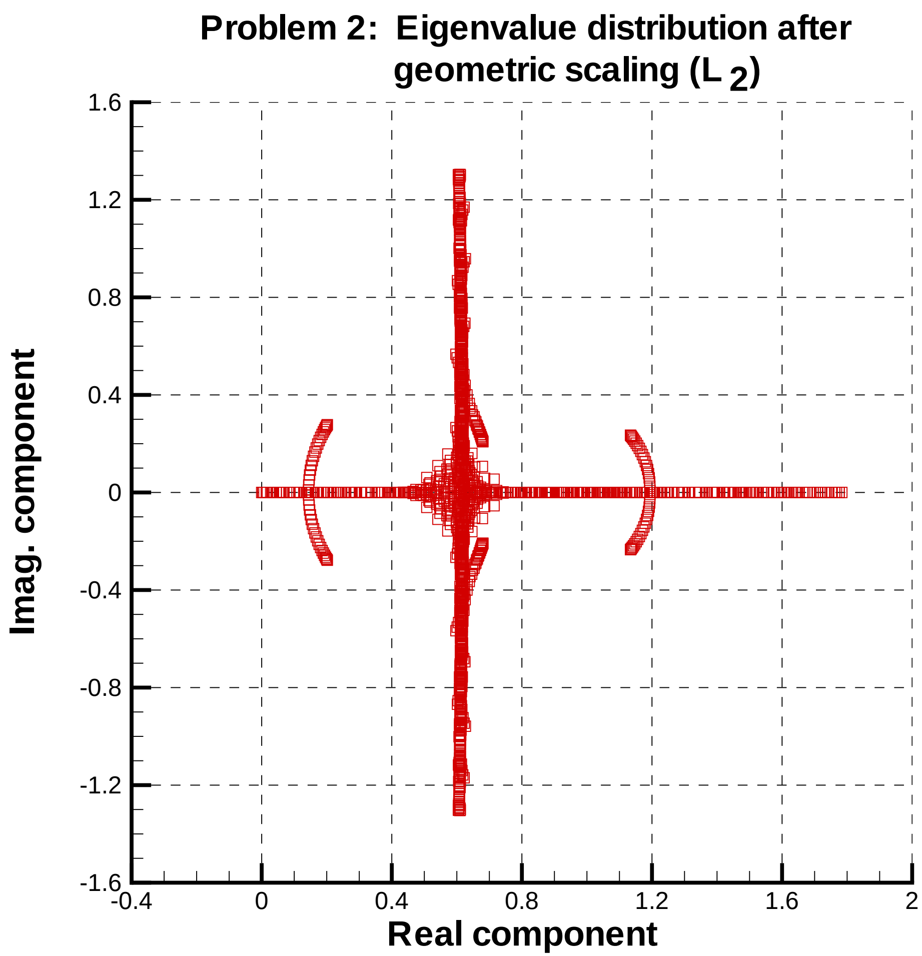

Problem 2: iterations & runtimes (sec.) with $D = 10^5$

Method	rel-res $=10^{-4}$	rel-res $=10^{-7}$	rel-res $=10^{-10}$
Bi-CGSTAB with GRS	no conv. 224 (0.82)	no conv. 730 (2.65)	no conv. conv. to 1.69×10^{-9}
Bi-CGSTAB+ILU(0) with GRS	96 (1.15) 21 (0.23)	192 (1.56) 125 (0.98)	conv. to 5.16×10^{-10} 217 (1.68)
GMRES	converged to 1.37		
with GRS	357 (1.38)	converged to 2.19×10^{-5}	
GMRES+ILU(0)	converged to 0.90		
with GRS	40 (0.32)	converged to 1.89×10^{-5}	

Problem 3 (well-known difficult problem)

- Example 4 from van der Vorst 1992 (Bi-CGSTAB paper)
- PDE on the unit square: $-\dfrac{\partial(A(x,y)u_x)}{\partial x} - \dfrac{\partial(A(x,y)u_y)}{\partial y} + B(x,y)u_x = F$
- Convection: $B(x,y) = 2\exp(2(x^2+y^2))$; discretization: 128×128

Additional details:

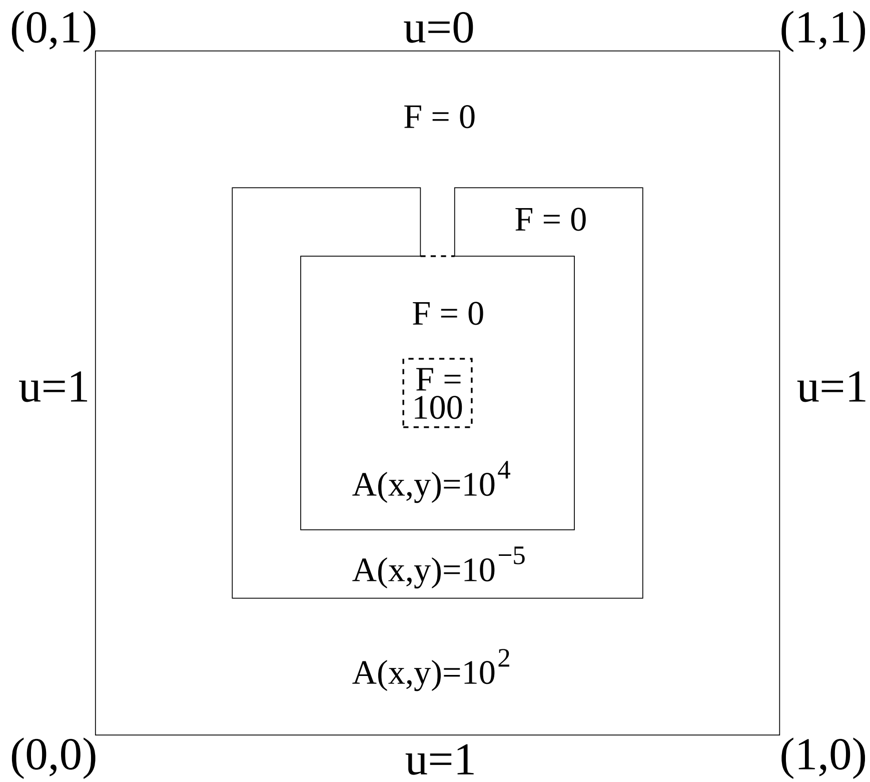

Problem 3: basic eigenvalue data

- Based on 40×40 discretization

Matrix	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\max}/\lambda_{\min}$	No. of e-values in first interval
Original	6.97E-5	7.93E+4	1.14E+9	1284
With GRS	1.21E-4	1.78E+0	1.47E+4	167

Problem 3: eigenvalue distribution

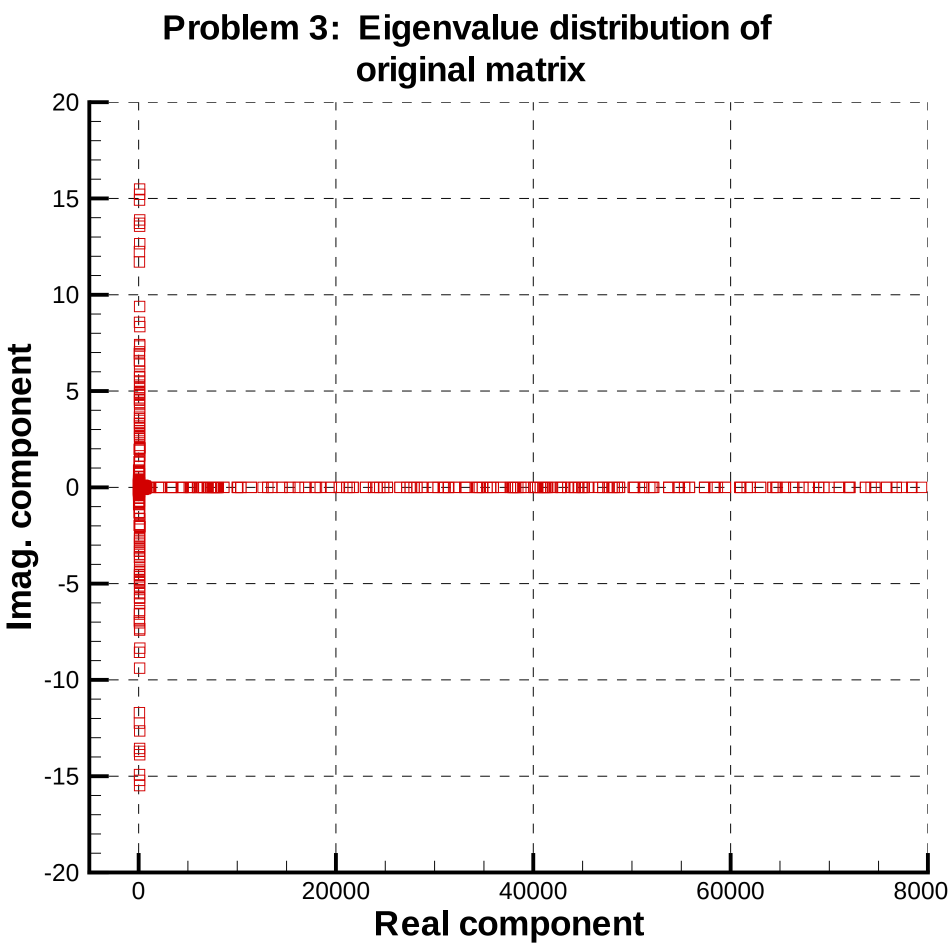

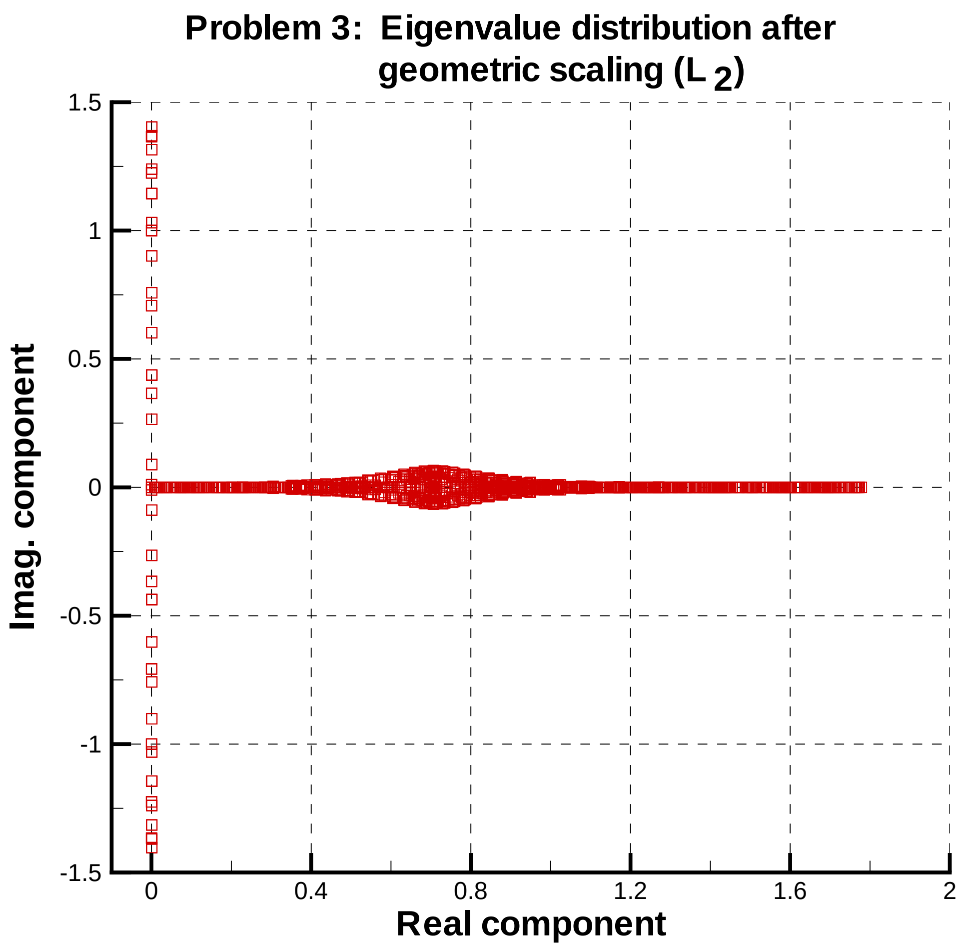

Problem 3: no. iterations & runtimes (sec.)

- Note: nothing converged without ILU(0)

Method	rel-res $=10^{-4}$	rel-res $=10^{-7}$	rel-res $=10^{-10}$
Bi-CGSTAB+ILU(0) with GRS	93 (0.60) 67 (0.44)	124 (0.79) 90 (0.59)	152 (0.95) 112 (0.72)
GMRES+ILU(0) with GRS	no conv. 338 (1.58)	no conv. 1008 (4.61)	no conv. 1683 (7.71)

Problem 4

- Based on a 3D (symmetric) problem of Graham & Hagger 1999
- With added convection terms
- PDE: $-\frac{\partial}{\partial x}(au_x) - \frac{\partial}{\partial y}(au_y) - \frac{\partial}{\partial z}(au_z) + du_x + eu_y + fu_z \; = \; 0$

$$\text{where } \; a(x,y,z) \; = \; \begin{cases} D & \text{if } \frac{1}{3} < x,\ y,\ z < \frac{2}{3}, \\ 1 & \text{otherwise.} \end{cases}$$

- $D = 10^4$ (original value) and $D = 10^6$ were tested
- Convection terms: $d = e = f = 100$
- Dirichlet boundary conditions ($u = 1$ on $z = 0$, $u = 0$ elsewhere)
- Discretization: $40\times40\times40$ and $80\times80\times80$

Problem 4: basic eigenvalue data, $D = 10^4$

- Discretization: $12 \times 12 \times 12$

Matrix	$\lambda_{\min}$	$\lambda_{\max}$	$\lambda_{\max}/\lambda_{\min}$	No. of e-values around $x=0$
Original	3.62E+0	1.09E+5	3.02E+4	1131
With GRS	8.83E-5	1.92E+0	2.17E+4	1
Cont. coef. ($D = 10^4$)	1.74E+3	1.18E+5	6.78E+1	1

Problem 4: eigenvalue distribution

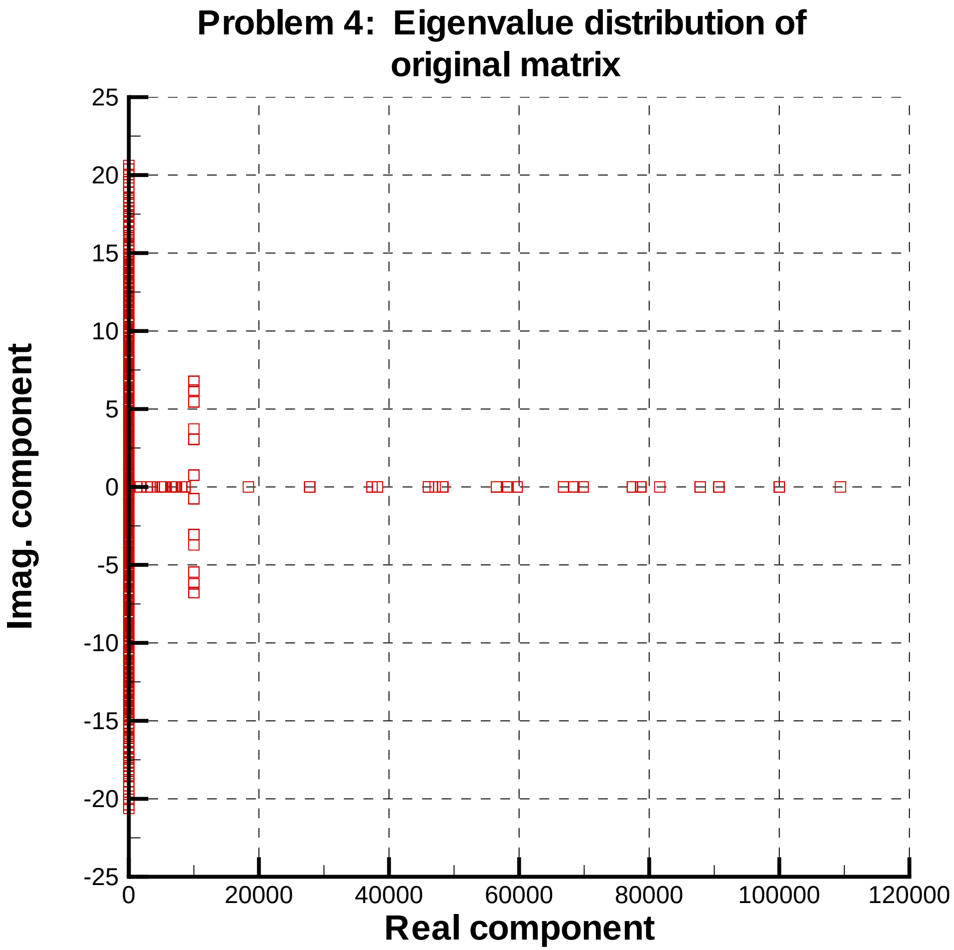

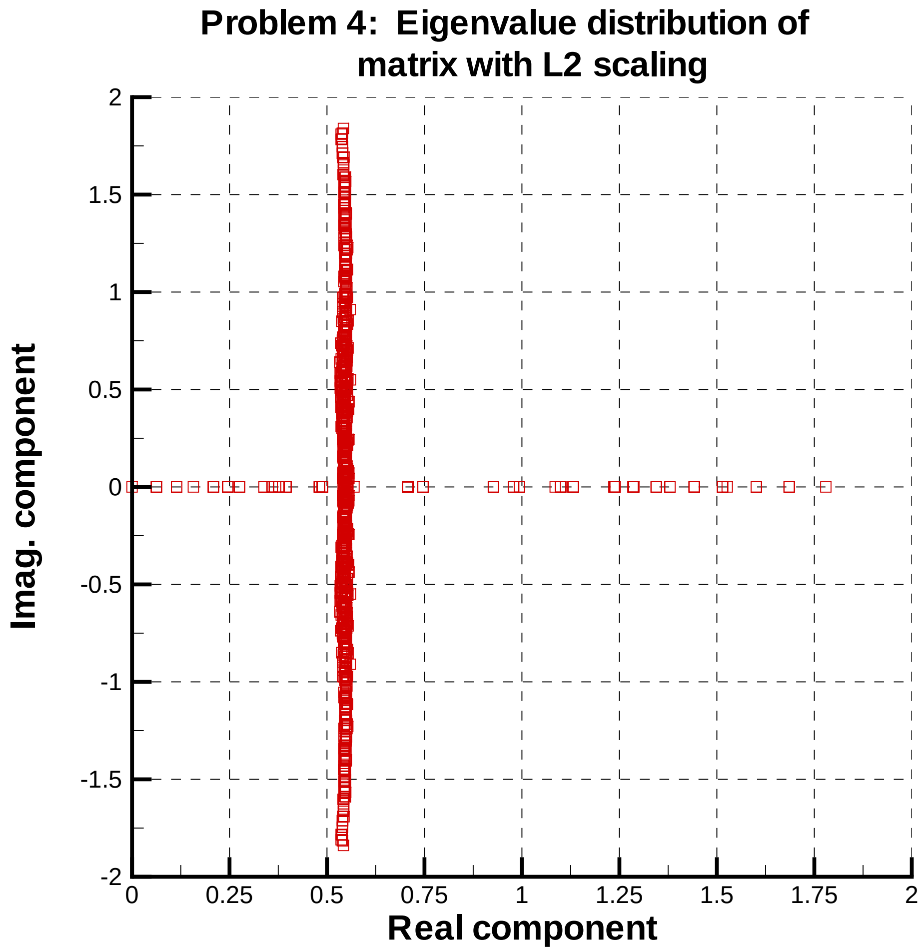

Problem 4: runtimes (sec.)

- $D = 10^6$, grid: $80 \times 80 \times 80$

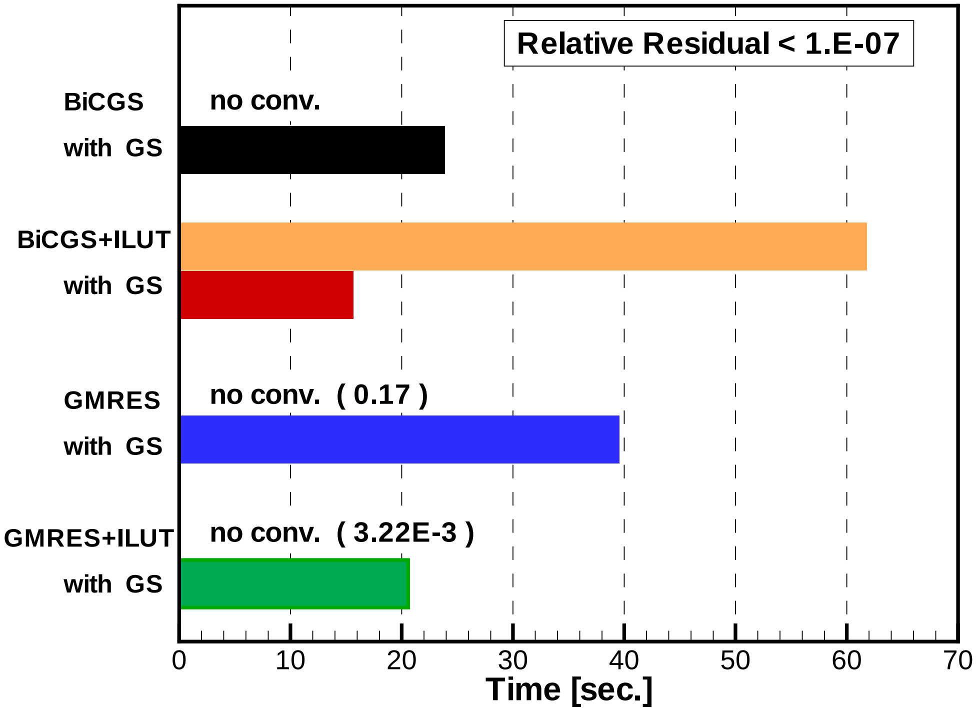

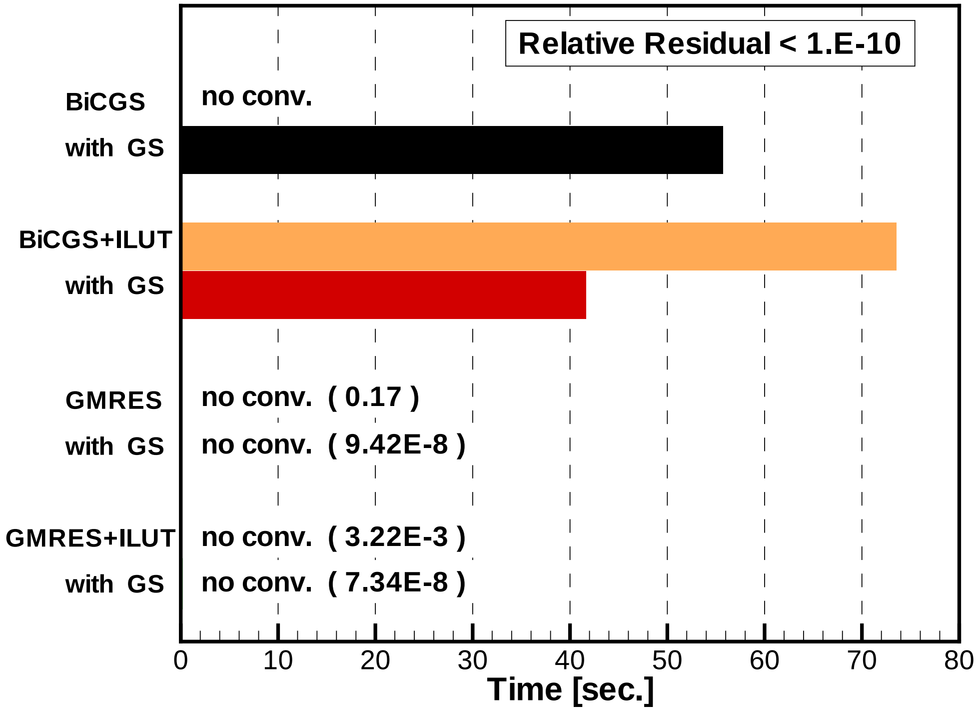

Summary of convergence data

Method	Problem 1	Problem 2	Problem 3	Problem 4
Bi-CGSTAB with GRS	$-\ -\ -$ $+\ +\ +$	$-\ -\ -$ 7.2×10^{-10}	$-\ -\ -$ $-\ -\ -$	$-\ -\ -$ $+\ +\ +$
Bi-CGSTAB+ILU(0) with GRS	$+\ +\ +$ $+\ *\ *$	$+\ +\ -$ $*\ *\ +$	$+\ +\ +$ $*\ *\ *$	$+\ +\ +$ $*\ *\ *$
GMRES with GRS	3.8×10^{-2} 1.1×10^{-5}	1.37 2.2×10^{-5}	$-\ -\ -$ $-\ -\ -$	0.27 1.8×10^{-5}
GMRES+ILU(0) with GRS	3.9×10^{-3} 1.1×10^{-5}	0.90 1.9×10^{-5}	$-\ -\ -$ $+\ +\ +$	0.29 1.25×10^{-5}

Note: '$-$' means no convergence, '$+$' means convergence, '$*$' means better convergence. The numbers indicate the best relative error obtained.

Problem 4: degradation with increased convection

Method / Convection:	100	200	500	1000
Bi-CGSTAB with GRS	— 10^{-10}	— 10^{-10}	— 10^{-10}	— —
Bi-CGSTAB+ILU(0) with GRS	10^{-10} 10^{-10}	10^{-4} 10^{-4}	— —	— —
GMRES with GRS	— 10^{-4}	— 10^{-4}	— 10^{-4}	— —
GMRES+ILU(0) with GRS	— 10^{-4}	— —	— —	— —
Note: '—' means no convergence. The numbers indicate which relative error goal was obtained.				

Problem 4: solutions for large convection

Convection:	**100**			**200**			**500**			**1000**		
Convergence goal:	10^{-4}	10^{-7}	10^{-10}	10^{-4}	10^{-7}	10^{-10}	10^{-4}	10^{-7}	10^{-10}	10^{-4}	10^{-7}	10^{-10}
Bi-CGSTAB+GRS	1.0	2.8	3.6	2.8	7.2	9.9	9.7	21.0	38.2	—	—	—
Bi-CGSTAB+ILU(0)	1.4	1.6	1.7	2.3	—	—	—	—	—	—	—	—
with GRS	**0.8**	**1.4**	**1.7**	**1.9**	—	—	—	—	—	—	—	—
GMRES+GRS	2.1	—	—	2.2	—	—	2.8	—	—	—	—	—
GMRES+ILU(0)+GRS	**0.8**	—	—	—	—	—	—	—	—	—	—	—
CGNR+GRS	5.1	10.4	11.3	4.8	9.5	10.2	6.0	12.4	13.4	7.3	15.0	16.2
CGNR+GCS	9.4	10.4	11.3	8.7	9.4	10.2	11.3	12.4	13.4	13.7	14.9	16.1
CGMN	2.0	4.3	4.6	**1.9**	**4.6**	**5.0**	**2.0**	**5.8**	**6.3**	**2.5**	**6.9**	**7.6**

- Times in seconds; minimal times in boldface
- Note: CGMN $\equiv$ CARP-CG on one processor

Thank you