Using condition numbers to assess numerical quality in HPC applications

Marc Baboulin

Inria Saclay / Université Paris-Sud, France

INRIA - Illinois Petascale Computing Joint Laboratory
9th workshop, June 12-14 2013

Motivation

Questions for HPC applications:

- How to measure the difficulty of solving the problem accurately ?
- Impact of errors in numerical algorithms ?
- Indicator for numerical “quality” ?
- Implementation in HPC libraries ?

Outline

1. Condition numbers

Outline

Outline

Outline

Outline

1. Condition numbers
2. Least squares conditioning
3. Numerical experiments
4. Conclusion

Perturbation analysis scheme

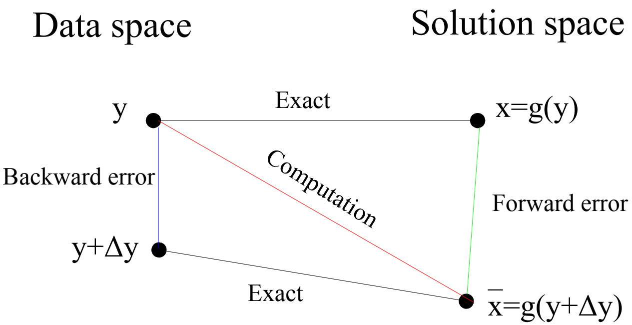

- Approach based on backward error analysis (Wilkinson)
- Notion of sensitivity of a solution to change in data (Turing, Rice)

Tools for error analysis (Wilkinson)

- Forward error: $\|x - \overline{x}\|$ (absolute)
  relative$\rightarrow$ independent to scaling
- Backward error (of a solution): distance between exact and perturbed problem $\rightarrow$ measure perturbation on data
  $\eta(\overline{x}) = \inf\{\|\Delta y\| : \overline{x} = g(y + \Delta y)\}$
- Condition number (of a problem): effect on the solution of small change in data $\rightarrow$ measure error amplification
  $K(y) = \lim_{\delta \rightarrow 0} \sup_{0 < \|\Delta y\| \leqslant \delta} \frac{\|g(y+\Delta y) - g(y)\|}{\|\Delta y\|}$

Up to first order: $\eta(\overline{x}) \times K(y) \approx \|x - \overline{x}\|$

Condition number

Assume g is differentiable,
From Taylor's theorem: $x - \overline{x} = g'(y).\Delta y + \mathcal{O}(\|\Delta y\|^2)$,
Condition number of g at y is :

$$K(y) = |||g'(y)||| = \max_{z \neq 0} \frac{\|g'(y).z\|}{\|z\|}.$$

- First order approximation
- Can be normalized with $K(y)\|y\|/\|g(y)\|$
- CN depends on metrics chosen to measure errors

Measuring errors on data

Normwise CN: use classical norms (e.g $\|\cdot\|_p,\ p = 1, 2, \infty$ or $\|\cdot\|_F$)

Example: $x = g(A, b),\ $ (e.g., $x = A^{-1}b$ or $x = A^{\dagger}b$)

product norm $\|(\Delta A, \Delta b)\|_E = \sqrt{\alpha^2 \|\Delta A\|^2_{\text{F or 2}} + \beta^2 \|\Delta b\|_2^2}$

- $\alpha = \beta = 1$ (choice for this talk)
- $\alpha = \frac{1}{\|A\|_{\text{F or 2}}}$ and $\beta = \frac{1}{\|b\|_2}$

Measuring errors on data

Componentwise CN: use metrics that take into account matrix structure like sparsity or scaling

- minimize amplification of errors resulting in minimal condition number
- **Example:** $x = g(A, b)$, (e.g., $x = A^{-1}b$ or $x = A^{\dagger}b$)
  product norm $\|(\Delta A, \Delta b)\|_E = \min\{\omega, |\Delta A| \leqslant \omega|A|, |\Delta b| \leqslant \omega|b|\}$
- max of relative perturbation for each data component

Deriving condition numbers using dual techniques

Property: If $J : E \rightarrow G$ linear, $J^* : G \rightarrow E$ then $|||J||| = |||J^*|||$.

1. choose norms $\|\cdot\|_E$ and $\|\cdot\|_G$ and determine dual norms,
2. determine the derivative $g'(y)$,
3. determine the adjoint operator $g'(y)^*$,
4. compute $K(y) = \max_{\|x\|_{G*}=1} \|g'(y)^*.x\|_{E*}$.

Working on dual space $\rightarrow$ **maximization over a space of smaller dimension**

Details in [MB, Gratton, BIT 2009]

Outline

Application to linear least squares

Linear least squares (LLS), full rank:

$$\mathbf{Ax} \simeq \mathbf{b} \quad (A \in \mathbb{R}^{m \times n}, m > n)$$

Solution is: $x = A^{\dagger} b = (A^T A)^{-1} A^T b$
We study the sensitivity of

$$\mathbf{g}(\mathbf{A}, \mathbf{b}) = \mathbf{x} \quad \rightarrow \quad \kappa_{LS}$$

or

$$\mathbf{g}(\mathbf{A}, \mathbf{b}) = \mathbf{e_i^T x} \quad \rightarrow \quad \kappa_i$$

- Choice of norm: $\|(\Delta A, \Delta b)\| = \sqrt{\|\Delta A\|_{\mathrm{F}}^2 + \|\Delta b\|_2^2}$
- Square linear system = special case of LLS
- Generalized to $\mathbf{g}(\mathbf{A}, \mathbf{b}) = \mathbf{L^T x}$
  [MB, Gratton, SIMAX 2007 and SIMAX 2011]

LLS condition numbers

Assume *R* factor available (from QR decomposition):

Conditioning of the solution x:

$$\kappa_{LS} = \|R^{-1}\|_2 \left(\|R^{-1}\|_2^2 \|r\|_2^2 + \|x\|_2^2 + 1\right)^{\frac{1}{2}}.$$

Conditioning of a solution component x_i:

$$\kappa_i = \left(\|R^{-1}R^{-T}e_i\|_2^2 \|r\|_2^2 + \|R^{-T}e_i\|_2^2 (\|x\|_2^2 + 1)\right)^{\frac{1}{2}}.$$

If errors only on b:

$$\kappa_{LS} = \left\|R^{-1}\right\|_2, \qquad \kappa_i = \left\|R^{-T}e_i\right\|_2$$

Statistical interpretation

Statistical model:

- $Ax = b + \epsilon$ with $E(\epsilon) = 0$ and $var(\epsilon) = \sigma^2 I$
- Variance-covariance: $C = \sigma^2 (A^T A)^{-1}$
- c_{ii}: variance of each x_i
  $c_{ii} = \sigma^2 \|e_i^T A^\dagger\|_2^2 = \sigma^2 \left\|R^{-T} e_i\right\|_2^2$
- $c_{ij},\ i \neq j$: covariance between x_i and x_j
  $C_i = \sigma^2 (A^T A)^{-1} e_i = \sigma^2 R^{-1} (R^{-T} e_i)$
- σ^2 is estimated by $\frac{1}{m-n} \|r\|_2^2$

Condition numbers:

- $\kappa_{LS} = \frac{\|C\|_2^{1/2}}{\sigma_b} ((m-n)\|C\|_2 + \|x\|_2^2 + 1)^{1/2}$
- $\kappa_i = \frac{1}{\sigma_b} ((m-n)\|C_i\|_2^2 + c_{ii}(\|x\|_2^2 + 1))^{1/2}$

Algorithms in [MB, Dongarra, Gratton, Langou, NLAA 2009] and [MB, Dongarra, Lacroix, AMMCS 2013]

Componentwise condition numbers

- **Computable formula**:
  $\kappa_{LS} = \| \sum_{j=1}^{n} |(A^T A)^{-1} (e_j r^T - x_j A^T)||A(:,j)| + |A^\dagger||b| \|_\infty$
  [MB, Gratton, BIT 2009]
- **If R is available**:
  $(A^T A)^{-1} (e_j r^T - x_j A^T) = R^{-1} R^{-T} (e_j r^T - x_j A^T)$
- **With (Sca)LAPACK**: 2 triangular solves with multiple RHS.
- When $m = n$ (**linear system**):
  $\kappa = \| |A^{-1}|(|A||x| + |b|) \|_\infty$ [Higham, 2002]

Statistical condition estimation

Objective: less flops in computing condition numbers

Algorithm:

- Generate q random orthogonal vectors
- For $j = 1$ to q

  Compute $\kappa_j = \left(\|R^{-1}R^{-T}z_j\|_2^2\|r\|_2^2 + \|R^{-T}z_i\|_2^2(\|x\|_2^2 + 1)\right)^{\frac{1}{2}}$
- Compute $\bar{\kappa}_{LS} = \frac{\omega_q}{\omega_n}\sqrt{\sum_{j=1}^{q}\kappa_j^2}$ with $\omega_q = \sqrt{\frac{2}{\pi(q-\frac{1}{2})}}$

Cost : $2qn^2$ flops (if $n \gg q$).

See [MB, Gratton, Lacroix, Laub, 2012]

Accuracy: For $q = 3$, $Pr\left(\kappa_{LS}/10 \leqslant \bar{\kappa}_{LS} \leqslant 10\kappa_{LS}\right) \approx 99.9\%$.

Experiments:

$cond_2(A)$	50	10^3	10^5	10^7	10^8	10^{10}
$\|r\|_2 = 10^{-10}$	3.32	1.46	1.19	1.10	1.03	1.07
$\|r\|_2 = 10^{-5}$	3.33	1.45	1.18	1.07	1.09	1.05
$\|r\|_2 = 1$	3.36	1.45	1.19	1.19	1.05	1.15
$\|r\|_2 = 10^5$	3.33	1.24	1.04	1.05	1.05	1.02
$\|r\|_2 = 10^{10}$	1.44	1.07	1.09	1.00	1.01	1.07

Ratio $\bar{\kappa}_{LS}/\kappa_{LS}$ for $q = 2$

100 random problems of size: $10^4 \times 2.5 \cdot 10^3$

$(cond_2(A) = \|A\|_2 \left\|A^\dagger\right\|_2, \quad r = b - Ax)$

Componentwise statistical condition estimation

Algorithm:

- For $j = 1$ to q
  Generate $S_j \in \mathbb{R}^{n \times n}$, $g_j \in \mathbb{R}^n$, $h_j \in \mathbb{R}^n$ with entries in $\mathcal{N}(0, 1)$
  Compute $u_j = R^{-1}(g_j - S_j x + \|Ax - b\|_2 R^{-T} h_j)$
- Compute vector $\bar{\kappa}_{CW} = \frac{\sum_{i=1}^{q} |u_j|}{q\omega_p\sqrt{p}}$ with $\omega_q = \sqrt{\frac{2}{\pi(q-\frac{1}{2})}}$ and $p = m(n+1)$

Cost (flops) $\simeq 2qn^2$ (2 $n \times n$ triangular solves with q RHS).

Componentwise statistical condition estimation

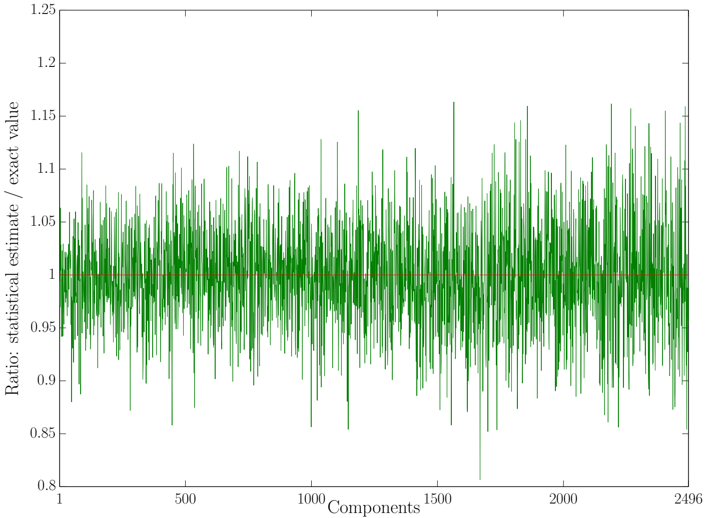

Ratio $\bar{\kappa}_i/\kappa_i$.

Outline

Computing condition numbers with some HPC libraries

condition number	linear algebra operation	Routine	flops count
κ_{LS}	singular values of R	DSYEVD	$\mathcal{O}(n^3)$
$\bar{\kappa}_{LS}$	generate random orthogonal vectors 2 triangular solves	DTRSV	$\mathcal{O}(n^2)$
κ_i	$R^T y = e_i$ and $Rz = y$	DTRSV	$2n^2$
all $\kappa_i, \; i = 1, n$	$RY = I$ and compute YY^T	DPOTRI	$2n^3/3$
all $\bar{\kappa}_i$	generate random vectors 2 triangular solves	DTRSV	$\mathcal{O}(n^2)$

Computation of least squares conditioning with (Sca)LAPACK and MAGMA
(cost for solution = $2mn^2$)

informatics mathematics Inria

Performance results

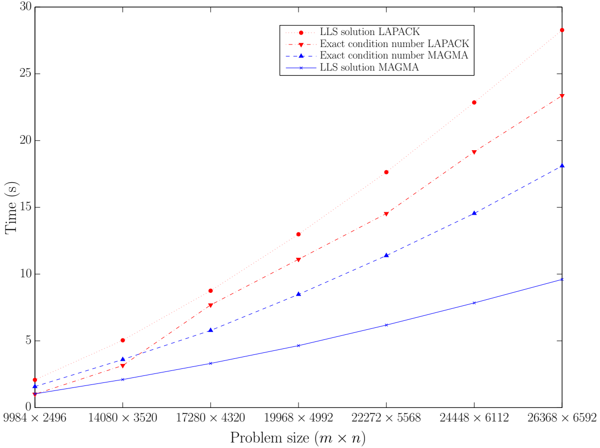

Time for LLS solution and condition number using LAPACK (MKL) and MAGMA

Intel Xeon E5645 2 $\times$ 6 cores @ 2.4 GHz - GPU C2075 @ 1.15 GHz

MAGMA: 3 times faster for the solution and 1.3 times for the conditioning

but, contrary to LAPACK, CN is twice more expensive than solution

Performance results (MAGMA)

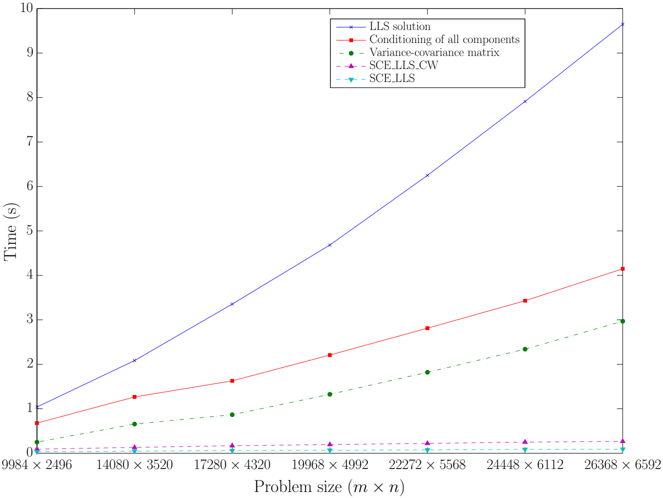

Computation of LLS condition numbers with MAGMA.

Intel Xeon E5645 2 × 6 cores @ 2.4 GHz - GPU C2075 @ 1.15 GHz

Physical application

- Earth's gravity field coefficients

$$V(r,\theta,\lambda) = \frac{GM}{R}\sum_{l=0}^{l_{max}}\left(\frac{R}{r}\right)^{l+1}\sum_{m=0}^{l}\overline{P}_{lm}(\cos\theta)\left[\overline{C}_{lm}\cos m\lambda + \overline{S}_{lm}\sin m\lambda\right]$$

$\overline{C}_{lm}, \overline{S}_{lm}$?

- Solution computed using incremental least squares solver (QR) based on ScaLAPACK(90, 000 unknowns, 2.6 millions obs.)
- Condition number of each x_i:

$$\kappa_i^{(rel)} = \left\|e_i^T R^{-1}\right\|_2 \|b\|_2 / |x_i|,$$

- Possibility of regularization (special case of Tikhonov) by performing the QR factorization of $\begin{pmatrix} R \\ D \end{pmatrix}$, with $D = diag(0,\cdots,0,\alpha,\cdots,\alpha)$, $\alpha \propto 10^{-5}/l_{max}^2$

Numerical results

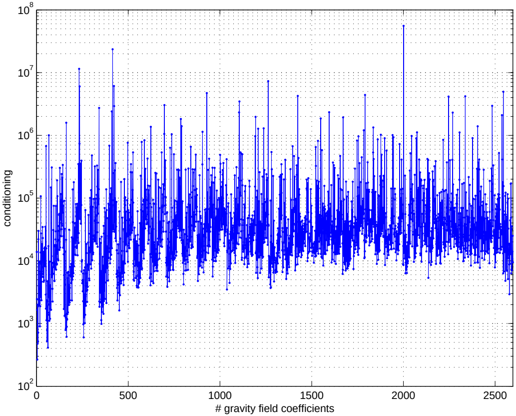

Amplitude of relative condition numbers for gravity field coefficients.

Numerical results

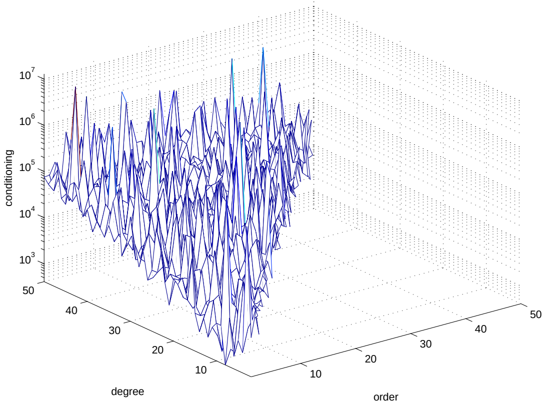

Conditioning of spherical harmonic coefficients $\overline{C}_{\ell m}$ $(2 \leqslant \ell \leqslant 50\ ,\ 1 \leqslant m \leqslant 50)$.

Numerical results

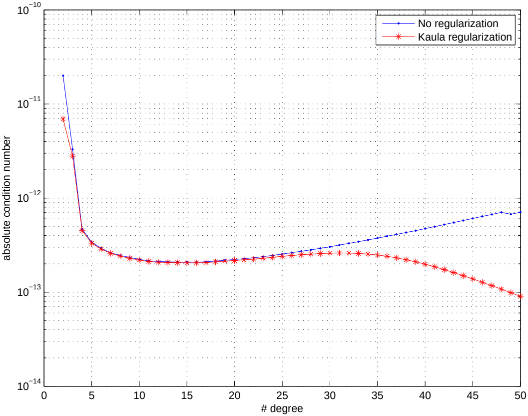

Effect of regularization on zonal coefficients $\overline{C}_{\ell 0}$ $(2 \leqslant \ell \leqslant 50)$.

Conclusion

- Exact expressions, statistical estimates and algorithms for computing condition numbers of least squares and linear systems
- With **statistical estimates**, the computational cost is $\mathcal{O}(n^2)$ (to be compared with $\mathcal{O}(mn^2)$ for the solution process and $\mathcal{O}(n^3)$ for the "exact" conditioning)
- Can be also applied to **linear systems**
- Implementation for HPC public domain libraries: (Sca)LAPACK, MAGMA
- For the GPU version, starting collaboration with Karl Rupp (Argonne National Laboratory)
- Visit of PhD student Yushan Wang to Argonne in August 2013

References for this talk

[1] M. Baboulin, S. Gratton, R. Lacroix, A. J. Laub
Efficient computation of condition estimates for linear least squares problems.
LAPACK Working Note 273 (2012).

[2] M. Baboulin, S. Gratton,
A contribution to the conditioning of the total least squares problem.
SIAM J. Matrix Analysis and Appl., Vol. 32, No 3, pp. 685-699 (2011).

[3] M. Baboulin, S. Gratton,
Using dual techniques to derive componentwise and mixed condition numbers for a linear function of a linear least squares solution.
BIT Numerical Mathematics, Vol. 49, No1, pp. 3-19 (2009).

[4] M. Baboulin, J. Dongarra, S. Gratton, J. Langou,
Computing the conditioning of the components of a linear least squares solution.
Numerical Linear Algebra with Applications, Vol. 16, No7, pp. 517-533 (2009).

[5] M. Arioli, M. Baboulin, S. Gratton,
A partial condition number for linear least-squares problems.
SIAM J. Matrix Analysis and Appl., Vol. 29, No 2, pp. 413-433 (2007).

[6] M. Baboulin, G. Balmino, S. Bruinsma, S. Gratton, J-C Marty,
Gravity field parameter estimation using QR factorization.

In proceedings of European Geosciences Union General Assembly (EGU 2008).

Collaborators for this talk:

- Mario Arioli (Rutherford Appleton Laboratory, UK)
- Jack Dongarra (U. Tennessee, USA)
- Serge Gratton (CERFACS, France)
- Rémi Lacroix (Inria, France)
- Julien Langou (U.C. Denver, USA)
- Alan Laub (UCLA, USA)