Numerical Issues in Testing Linear Algebra Algorithms

Nick Higham
School of Mathematics
The University of Manchester

higham@ma.man.ac.uk
http://www.ma.man.ac.uk/~higham
@nhigham

SIAM Conference on
Computational Science and Engineering
Boston, February 2013

What Are We Testing?

- Algorithm (against other algorithms).
- Implementation.

- Accuracy or stability.
- Speed.

Normwise Errors for $X = f(A)$

Forward relative error: $$\frac{\|X - \widehat{X}\|}{\|X\|}.$$

Normwise backward error:

$$\min\left\{ \frac{\|\Delta A\|}{\|A\|} : \widehat{X} = f(A + \Delta A) \right\}.$$

Normally a scaled residual.

Residual: $$\frac{\|f^{-1}(\widehat{X}) - A\|}{\|A\|}.$$

May be several.

E.g., for A^{+}: $AXA - I$, $XAX - I$, $\ldots$.

Judging the Results

- **Normwise backward error**: compare with u.
- **Forward relative error:** compare with $\mathrm{cond}(f, A)u$.

For forward error, how do we compute “x”?

Randomization for Matrix Functions

$A \in \mathbb{C}^{n \times n}$ with $\|A\| = 1$.

Davies (2007): "approximate diagonalization":

```
tol = 1e-8;
[V,D] = eig(A + tol*randn(n));
F = V*diag(feval(fun,diag(D)))/V;
```

Davies' analysis suggests rel err $\approx 10^{-8}$.

Randomization for Matrix Functions

$A \in \mathbb{C}^{n \times n}$ with $\|A\| = 1$.

Davies (2007): "approximate diagonalization":

```
tol = 1e-8;
[V,D] = eig(A + tol*randn(n));
F = V*diag(feval(fun,diag(D)))/V;
```

Davies' analysis suggests rel err $\approx 10^{-8}$.

H & Relton (2013, in progress):

```
digits(d)  % E.g., d = 50.
tol = 10^(-d/2);
[V,D] = eig(vpa(A) + tol*vpa( rand(n) ));
F = double( V*diag(feval(fun,diag(D)))/V );
```

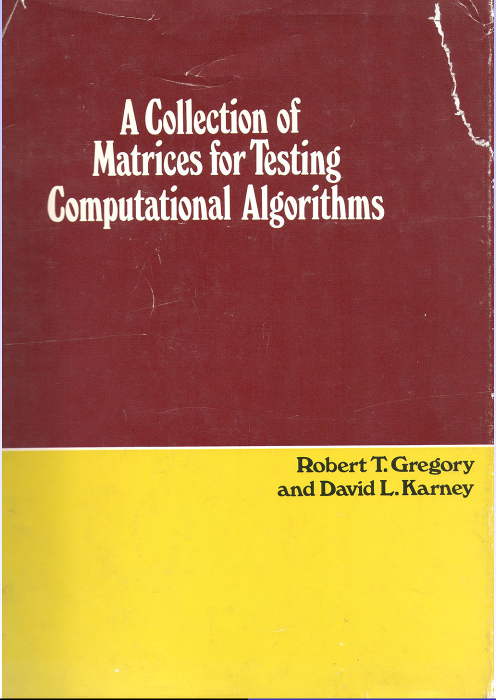
A Collection of
Matrices for Testing
Computational Algorithms
Robert T. Gregory
and David L. Karney

Test Problems

- Both artificial (difficult, wide-ranging) and “real-life”.
- Tool for **reproducible research**.
- **Linear algebra**
  - ACM Alg 694 (H, 1991), The Test Matrix Toolbox for MATLAB (H, 1993)
  - MATLAB `gallery` (63 matrices).
  - Matrix Market (Boisvert et al., 1997)
  - CONTEST (Taylor & D. J. Higham, 2009)
  - University of Florida Sparse Matrix Collection (Davis, 2011)
- **Optimization**: Cute (Bongartz et al., 1995), Cuter (Gould et al., 2003), . . .

Nonlinear Eigenvalue Problems (NEPs)

Let $F\colon \Omega \to \mathbb{C}^{m\times n}$ be analytic on open set $\Omega \subseteq \mathbb{C}$.

The **nonlinear eigenvalue problem**: Find scalars λ and nonzero $x, y \in \mathbb{C}^n$ satisfying $F(\lambda)x = 0$ and $y^*F(\lambda) = 0$.

In practice, elements of F most often **polynomial, rational or exponential functions of λ**.

Quadratic Eigenvalue Problem (QEP)

$$(\lambda^2 M + \lambda D + K)x = 0.$$

NLEVP Toolbox

NLEVP: A Collection of Nonlinear Eigenvalue Problems, T. Betcke, N. J. Higham, V. Mehrmann, C. Schröder, F. Tisseur, ACM TOMS, Feb 2013.

- v 1.0 (April 2008), v 2.0 (Nov 2010), v 3.0 (Dec 2011).
- 52 quadratic, polynomial, rational, nonlinear problems.
- Provided in the form of a MATLAB Toolbox.
- Problems from real-life applications + specifically constructed problems.

http://www.mims.manchester.ac.uk/research/numerical-analysis/nlevp.html

Sample of Quadratic Problems

$m \times n$ quadratic $Q(\lambda) = \lambda^2 M + \lambda D + K$.

Speaker box (`pep,qep,real,symmetric`).
$n = m = 107$. Finite element model of a speaker box.
$\|M\|_2 = 1$, $\|C\|_2 = 5.7 \times 10^{-2}$, $\|K\|_2 = 1.0 \times 10^7$.

Railtrack (`pep,qep,t-palindromic,sparse`).
$n = m = 1005$. Model of vibration of rail tracks under the excitation of high speed trains. $M = K^T$, $D = D^T$.

Surveillance (`pep,qep,real,nonsquare, nonregular`). $m = 21$, $n = 16$. From calibration of surveillance camera using human body as target (600-by-400 pixel camera).

Sample of Higher Degree Problems

Plasma drift (`pep`). Cubic polynomial ($n = 128, 512$) from modeling of drift instabilities in the plasma edge inside a Tokamak reactor.

Mirror (`pep, real, random`). Quartic, 9×9. homography-based method for calibrating a central cadioptric vision system

Butterfly (`pep, real, parameter-dependent, T-even, scalable, sparse`) quartic matrix polynomial with T-even structure.

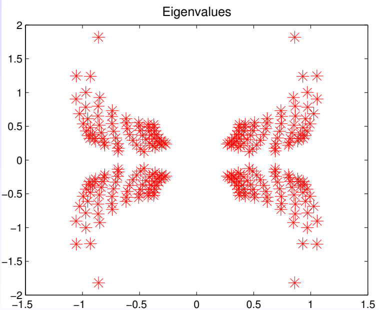

Usage (1)

```
>> nlevp('version')
This is the NLEVP collection of nonlinear
eigenvalue problems version 3.0. It was
released 22-Dec-2011 and contains 52 problems.

>> nlevp('query','problems')
ans =
    'acoustic_wave_1d'
     ...
>> coeffs = nlevp('omnicam2')
coeffs =
 [15x15 double] [15x15 double] [15x15 double]

>> [X,e] = polyeig(coeffs{:}); max(abs(e))
ans =
  3.6351e-001
```

Usage (2)

```
>> nlevp query sleeper
ans =
    'pep'
    'qep'
    'real'
    'scalable'
    'sparse'
    'symmetric'
    'proportionally-damped'
    'solution'
```

Usage (3)

```
>> pep = nlevp('query','pep');

>> qep = nlevp('query','qep');

>> pep_cubic_plus = setdiff(pep,qep)
pep_cubic_plus =
    'butterfly'
    'mirror'
    'orr_sommerfeld'
    'planar_waveguide'
    'plasma_drift'
    'relative_pose_5pt'
```

Residual

```
>> coeffs = nlevp('omnicam2');

>> [X,e] = polyeig(coeffs{:});

>> lam = e(end); x = X(:,end);
>> res = norm( nlevp('eval','omnicam2',lam)*x )
ans =
   1.3477e-30

>> berr = res / (norm(x)*( norm(coeffs{1}) ...
   + abs(lam)*norm(coeffs{2}) ...
   + abs(lam)^2*norm(coeffs{3})) )
berr =
   2.1897e-30
```

Tiny Relative Errors

Normwise relative errors (IEEE double, $u \approx 1.1 \times 10^{-16}$)

$$\frac{\|x - \widehat{x}\|_\infty}{\|x\|_\infty} = \frac{\max_i |x_i - \widehat{x}_i|}{\max_i |x_i|}$$

from a numerical experiment:

```
1.32e-22  3.39e-22  3.39e-21  8.67e-20
1.39e-18  4.36e-18  5.30e-18  5.83e-18
1.45e-17  3.76e-17  3.76e-17  4.27e-17
...
```

How can errors be $\ll u \approx 10^{-16}$?

Scalar Case

Base $\beta = 2$, $u = 2^{-t}$.

Theorem (Dingle & H, 2013)

If $x \neq 0$ and y are distinct normalized flpt numbers then $|x - y|/|x| \geq u$ and this lower bound is attainable.

Scalar Case

Base $\beta = 2$, $u = 2^{-t}$.

Theorem (Dingle & H, 2013)

If $x \neq 0$ and y are distinct normalized flpt numbers then $|x - y|/|x| \geq u$ and this lower bound is attainable.

Vector Case

But $\dfrac{\|x-y\|_\infty}{\|x\|_\infty} \ll u$ is possible.

$$x = \begin{bmatrix} 1 \\ 10^{-22} \end{bmatrix}, \quad y = \begin{bmatrix} 1 \\ 2 \times 10^{-22} \end{bmatrix}, \quad \frac{\|x-y\|_\infty}{\|x\|_\infty} = 10^{-22}.$$

Theorem (Dingle & H, 2013)

Let x, y with $x \neq 0$ be n-vectors of normalized floating point numbers. If $\|x-y\|_\infty/\|x\|_\infty < u$ then $x_k = y_k$ for all k such that $|x_k| = \|x\|_\infty$.

Relative Errors

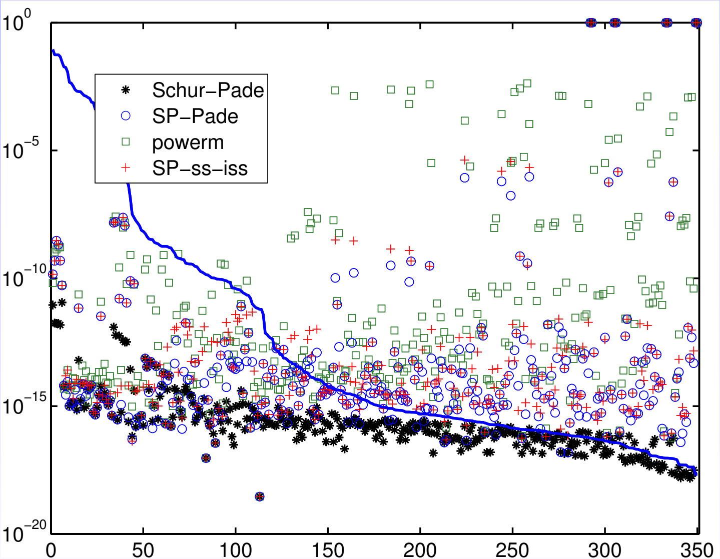

Performance Profiles

Dolan & Moré (2002).

For the given set of solvers and test problems, **plot**

x-axis: α

y-axis: probability that solver has error within factor α of smallest error over all solvers on the test set.

Performance Profile

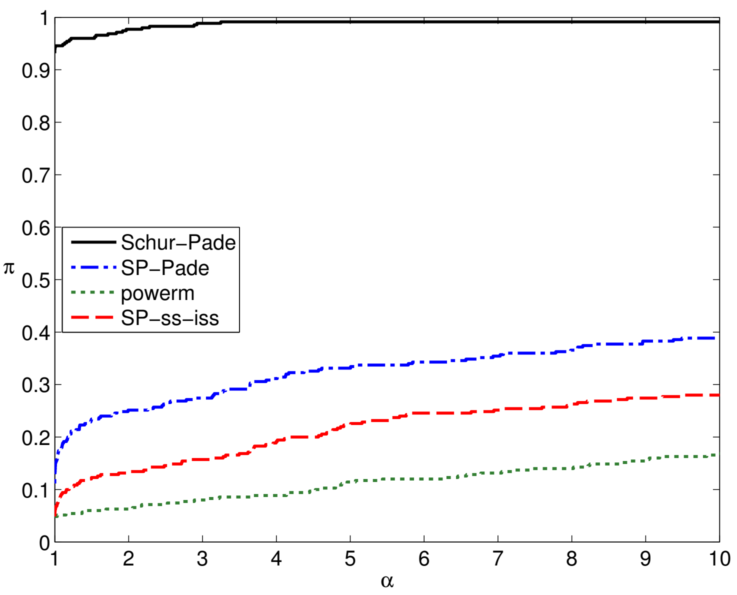

The Effect of Tiny Errors

Normwise relative errors (artificial):

Problem	Algorithm 1	Algorithm 2
1	4e-14	1e-16
2	6e-16	4e-16
3	1e-16	3e-16
4	9e-23	1e-17
5	6e-20	5e-17

Which algorithm is better?

Profile

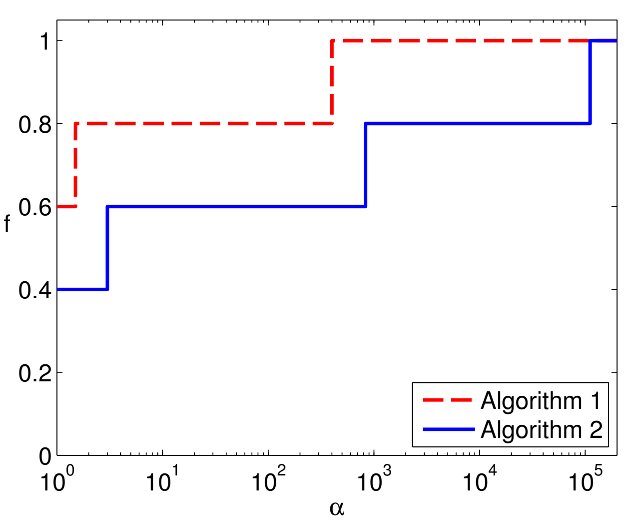

Alg 1	Alg 2
4e-14	1e-16
6e-16	4e-16
1e-16	3e-16
9e-23	1e-17
6e-20	5e-17

Round Small Errors up to u

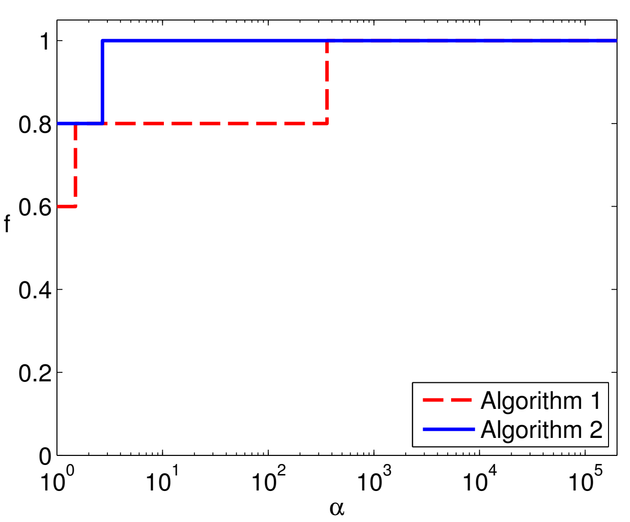

Alg 1	Alg 2
4e-14	**1e-16**
6e-16	4e-16
1e-16	3e-16
1e-16	**1e-16**
1e-16	**1e-16**

Add u to All Errors

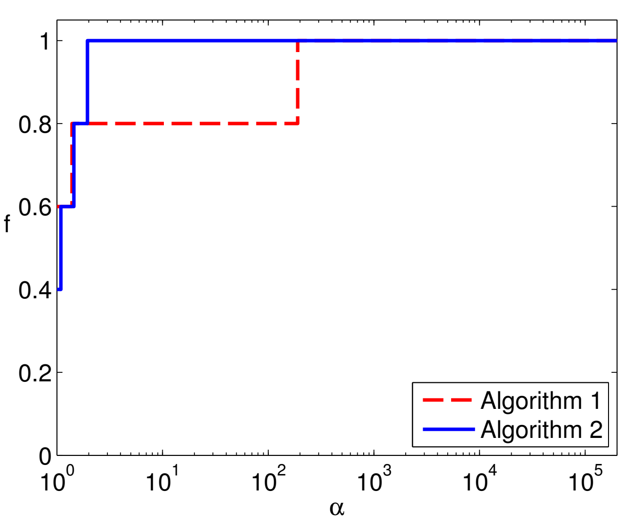

Alg 1	Alg 2
4e-14	2e-16
6e-16	4e-16
2e-16	3e-16
1.1e-16	1.2e-16
1.1e-16	1.6e-16

Transform the Data

Requirements

- Impose positive minimum.
- Do not decrease any error $< u$.
- Do not change any error $\geq u$.
- Preserve ordering of errors.

Transformation

- Map 0 to a (parameter).
  Typically, $a = u/20$.
- Map $[0, u]$ to $[a, u]$ linearly.
- Leave values $\geq u$ alone.

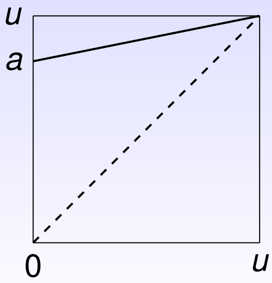

Before

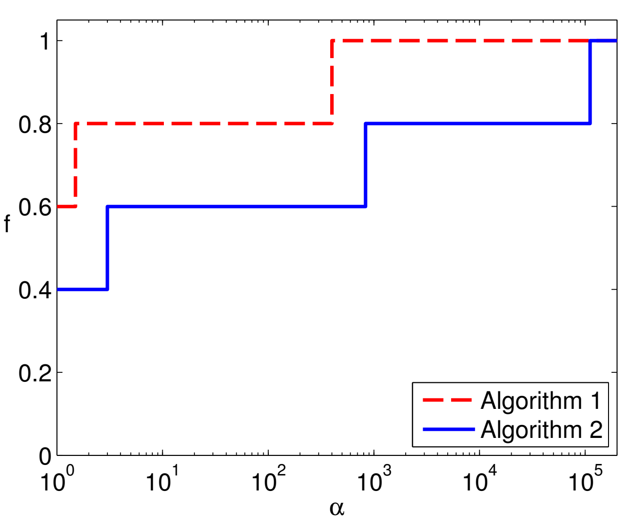

Alg 1	Alg 2
4e-14	1e-16
6e-16	4e-16
1e-16	3e-16
9e-23	1e-17
6e-20	5e-17

After

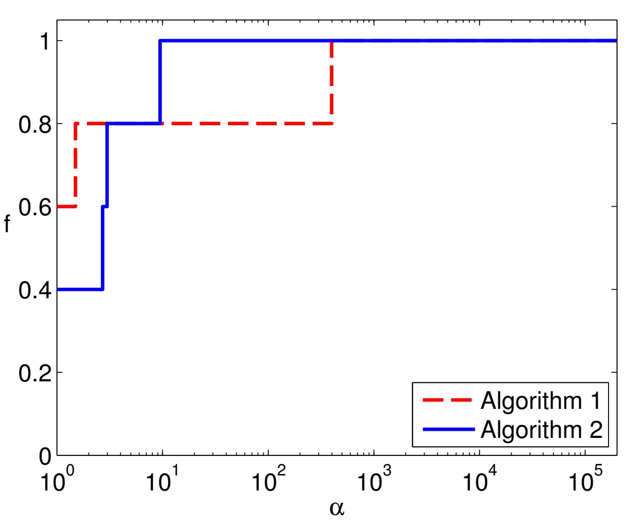

Alg 1	Alg 2
4.0e-14	**1.0e-16**
6.0e-16	4.0e-16
1.0e-16	3.0e-16
5.6e-18	**1.1e-17**
5.6e-18	**5.0e-17**

Matrix Exponential (Al-Mohy & H, 2011)

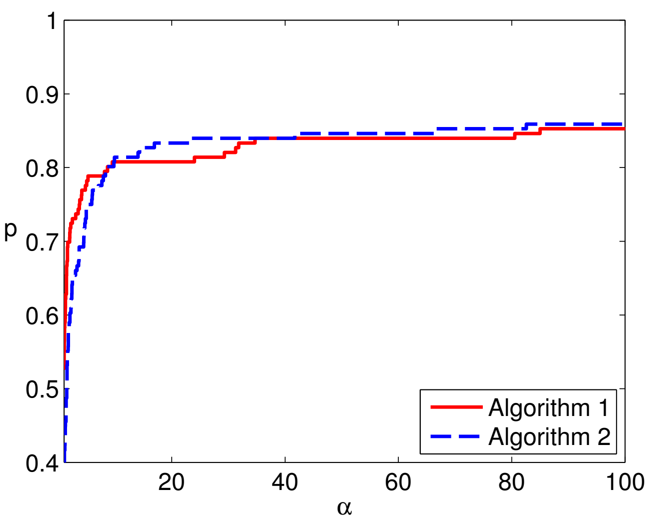

Matrix Exponential Transformed

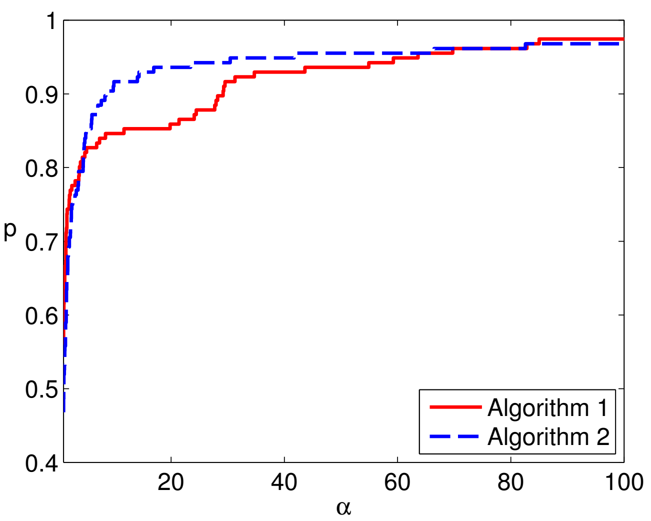

Finally

- Need for more test collections, with automated access by **problem property**.

- Tiny relative error vectors are possible and can skew performance profiles.

  Our transformation produces more **useful** profiles.

References I

A. H. Al-Mohy and N. J. Higham.
Computing the action of the matrix exponential, with an application to exponential integrators.
SIAM J. Sci. Comput., 33(2):488–511, 2011.

T. Betcke, N. J. Higham, V. Mehrmann, C. Schröder, and F. Tisseur.
NLEVP: A collection of nonlinear eigenvalue problems.
ACM Trans. Math. Software, 39(2):7:1–7:28, Feb. 2013.

References II

- R. F. Boisvert, R. Pozo, K. Remington, R. F. Barrett, and J. J. Dongarra.
  Matrix Market: A Web resource for test matrix collections.
  In R. F. Boisvert, editor, *Quality of Numerical Software: Assessment and Enhancement*, pages 125–136. Chapman and Hall, London, 1997.
- E. B. Davies.
  Approximate diagonalization.
  SIAM J. Matrix Anal. Appl., 29(4):1051–1064, 2007.
- T. A. Davis and Y. Hu.
  The University of Florida sparse matrix collection.
  ACM Trans. Math. Software, 38(1):1:1–1:25, 2011.

References III

N. J. Dingle and N. J. Higham.
Reducing the influence of tiny normwise relative errors on performance profiles.
ACM Trans. Math. Software, 39(4):24:1–24:11, 2013.

E. D. Dolan and J. J. Moré.
Benchmarking optimization software with performance profiles.
Math. Programming, 91:201–213, 2002.

References IV

N. J. Higham.
Accuracy and Stability of Numerical Algorithms.
Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edition, 2002.
ISBN 0-89871-521-0.
xxx+680 pp.

I. C. F. Ipsen.
Accurate eigenvalues for fast trains.
SIAM News, 37(9):1–2, Nov. 2004.

A. Taylor and D. J. Higham.
CONTEST: A controllable test matrix toolbox for MATLAB.
ACM Trans. Math. Software, 35(4):26:1–26:17, 2009.