How long does it take to compute the eigenvalues of a random symmetric matrix?

Govind Menon (Brown University)

Joint work with:

Christian Pfrang (Ph.D, Brown 2011)

Percy Deift, Tom Trogdon (Courant Institute)

Enrique Pujals (IMPA).

The guiding question

How do eigenvalue algorithms perform on large, random matrices?*

** this is* ***theoretical*** *numerical analysis (at least, so far), I'm not a professional numerical linear algebraist....*

Three stories

(1) Eigenvalue algorithms as dynamical systems (flashback to the 1980s).

(2) Random matrices and empirical universality (with Pfrang, Deift, Trogdon).

(3) A few explanations (with Pujals).

Part 1. Eigenvalue algorithms as dynamical systems

The QR factorization

Assume L is a real, symmetrix matrix. The Gram-Schmidt procedure may be viewed as the matrix factorization

$$L = QR$$

Q is an orthogonal matrix

R is an upper triangular matrix

The naive QR algorithm

The QR algorithm is an iterative scheme to compute the eigenvalues of L.
There are two steps:

Step 1: Compute QR factors at iterate k.

$$L_k = Q_k R_k$$

Step 2: Intertwine these factors to obtain the next iterate

$$L_{k+1} = R_k Q_k.$$

Observe that $$L_{k+1} = Q_k^T L_k Q_k.$$

The unshifted QR algorithm generates a sequence of <u>isospectral matrices</u> that typically converges to the desired diagonal matrix of eigenvalues.

The practical QR algorithm

(1) Reduce the initial matrix to tridiagonal form.
(2) Use a shift while iterating as follows.

At the k-th step compute the QR factorization: $L_k - \mu_k I = Q_k R_k.$

Then set

$$L_{k+1} = R_k Q_k + \mu_k I.$$

Again, we find

$$L_{k+1} = Q_k^T L_k Q_k.$$

The shift parameter is often chosen to be the eigenvalue of the lower right 2x2 block that is closer to the right-hand corner entry of the matrix (Wilkinson).

Jacobi matrices

Assume L is tridiagonal

$$L = \begin{pmatrix} a_1 & b_1 & 0 & \cdots & \cdots \\ b_1 & a_2 & b_2 & & \cdots \\ 0 & b_2 & a_3 & \ddots & 0 \\ \vdots & \vdots & \ddots & \ddots & b_{n-1} \\ 0 & 0 & \cdots & b_{n-1} & a_n \end{pmatrix}$$

and that the off-diagonal entries are positive.

This is a practical assumption: we may reduce a full matrix to a Jacobi matrix with the same spectrum by Householder reflections.

Spectral and inverse spectral theory (Stieltjes).

The spectral data consist of the eigenvalues of L and the top row of U:

$$L = U\Lambda U^T, \quad u = (1, 0, \ldots, 0)\, U.$$

Given L, clearly we may find Λ and u.

Stieltjes established the converse: given Λ and u we may reconstruct L.

$$\sum_{j=1}^{n} \frac{u_j^2}{z - \lambda_j} = \cfrac{1}{z - a_n + \cfrac{b_{n-1}^2}{z - a_{n-1} + \cfrac{b_{n-2}^2}{z - a_{n-2} + \cdots}}}.$$

QR as a dynamical system: the phase space

Assume L is a Jacobi matrix with distinct eigenvalues, say

$$L = U\Lambda U^T.$$

The QR iterates are <u>isospectral</u> Jacobi matrices.

$$L_k = U_k \Lambda U_k^T, \quad L_0 = L.$$

A more sophisticated form of the inverse spectral theorem reveals that the isospectral manifold is a convex polytope --*the permutahedron*.

The iterates of the QR algorithm live in the interior of this polytope.

Phase space and QR iterates for 3x3 matrices

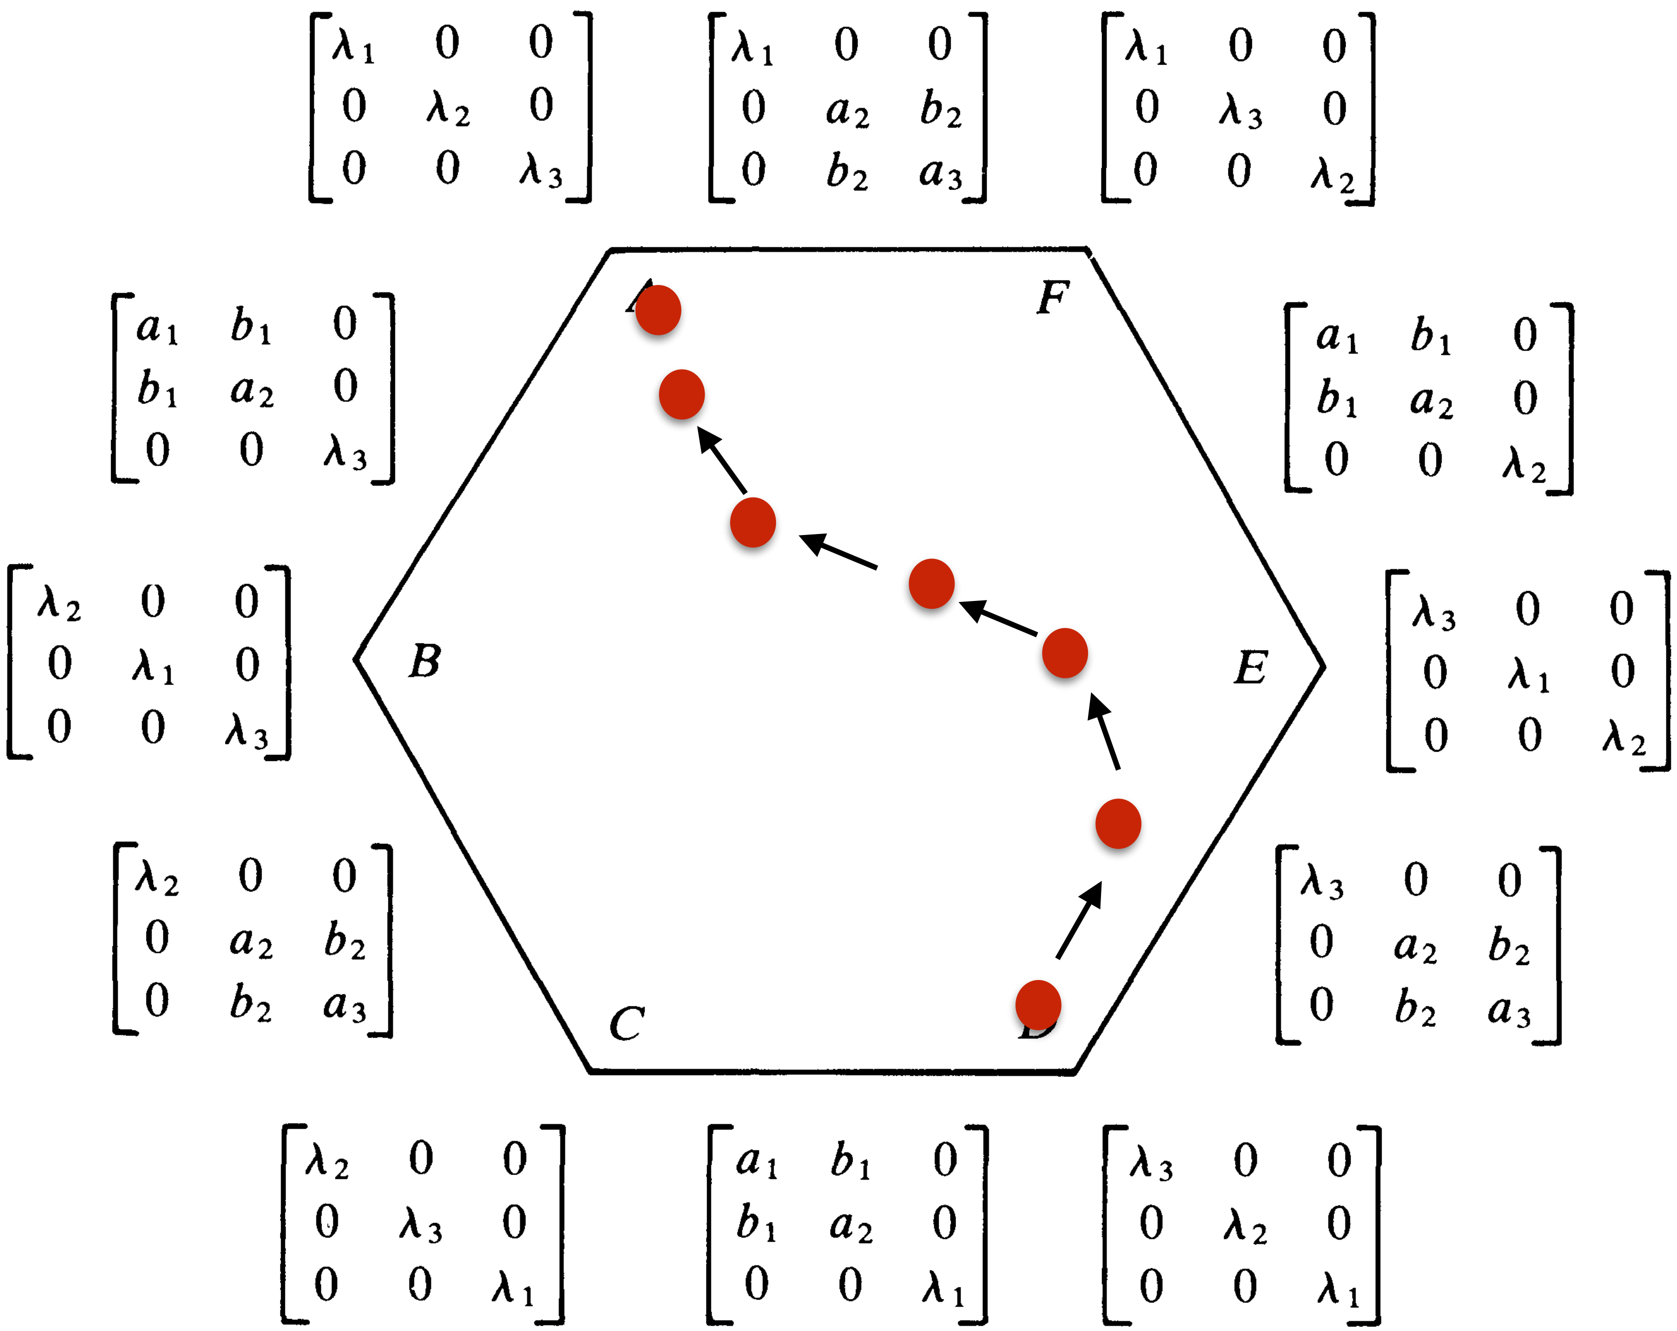

Convention: The eigenvalues are labelled in decreasing order: $\lambda_1 > \lambda_2 > \ldots > \lambda_n$

Deift, Nanda, Tomei (SIAM J. Numer. Anal., 1983).

The Toda lattice

Hamiltonian system of n particles with unit mass on the line. Each interacts with its neighbors through an exponential potential.

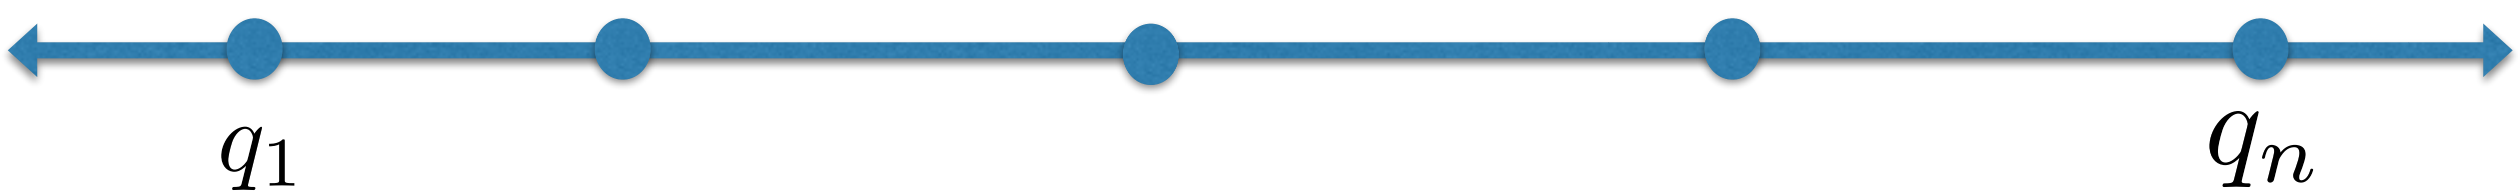

$$H(p,q) = \frac{1}{2}\sum_{j=1}^{n} p_j^2 + \sum_{j=1}^{n-1} e^{q_j - q_{j+1}}$$

$$\ddot{q}_j = e^{q_{j-1} - q_j} - e^{q_j - q_{j+1}}$$

Particle system to tridiagonal matrices (Flaschka, Manakov)

Set $$a_j = -\frac{1}{2}p_j, \quad b_j = \frac{1}{2}e^{(q_j - q_{j+1})/2}.$$

$$L = \begin{pmatrix} a_1 & b_1 & 0 & \cdots & \cdots \\ b_1 & a_2 & b_2 & & \cdots \\ 0 & b_2 & a_3 & \ddots & 0 \\ \vdots & \vdots & \ddots & \ddots & b_{n-1} \\ 0 & 0 & \cdots & b_{n-1} & a_n \end{pmatrix}$$

Jacobi

$$B = \begin{pmatrix} 0 & b_1 & 0 & \cdots & \cdots \\ -b_1 & 0 & b_2 & & \cdots \\ 0 & -b_2 & 0 & \ddots & 0 \\ \vdots & \vdots & \ddots & \ddots & b_{n-1} \\ 0 & 0 & \cdots & -b_{n-1} & 0 \end{pmatrix}$$

Tridiagonal, skew-symmetric

$$H(p,q) = \frac{1}{2}\sum_{j=1}^{n} p_j^2 + \sum_{j=1}^{n-1} e^{q_j - q_{j+1}}$$

$$\dot{q} = \frac{\partial H}{\partial p}, \quad \dot{p} = -\frac{\partial H}{\partial q}$$

$$H(L) = \frac{1}{2}\mathrm{Trace}(L^2)$$

$$\dot{L} = BL - LB := [B, L]$$

Moser's solution formula (1975)

The Toda lattice may be solved explicitly by inverse spectral theory. Suppose we know the spectral data for the initial matrix. If we denote

$$L(t) = U(t)\Lambda U(t)^T, \quad u(t) = (1, 0, \ldots, 0)\, U(t)$$

Then

$$v(t) = u(0)e^{\Lambda t}, \quad u(t) = \frac{v(t)}{\|v(t)\|}.$$

Thus, L(t) may be recovered by Stieltjes inverse spectral mapping.

Completely integrable Hamiltonian flows on Jacobi matrices

Consider a scalar function G with derivative g and extend it to a function on matrices.

$$L = U\Lambda U^T, \quad G(L) = UG(\Lambda)U^T.$$

Fundamental fact: the Lax equation

$$\dot{L} = [P_s(g(L)), L],$$

defines a completely integrable Hamiltonian flow, with Hamiltonian

$$H(L) = \mathrm{Tr}\ G(L).$$

P_s is the projection onto skew-symmetric matrices.

Complete integrability: all these flows commute, and all of them may be solved by Moser's recipe.

Symes theorem (1980)

The iterates of the QR algorithm are exactly the same as the solutions to the QR flow evaluated at integer times!

The Hamiltonian for the QR flow is the <u>spectral entropy</u>.

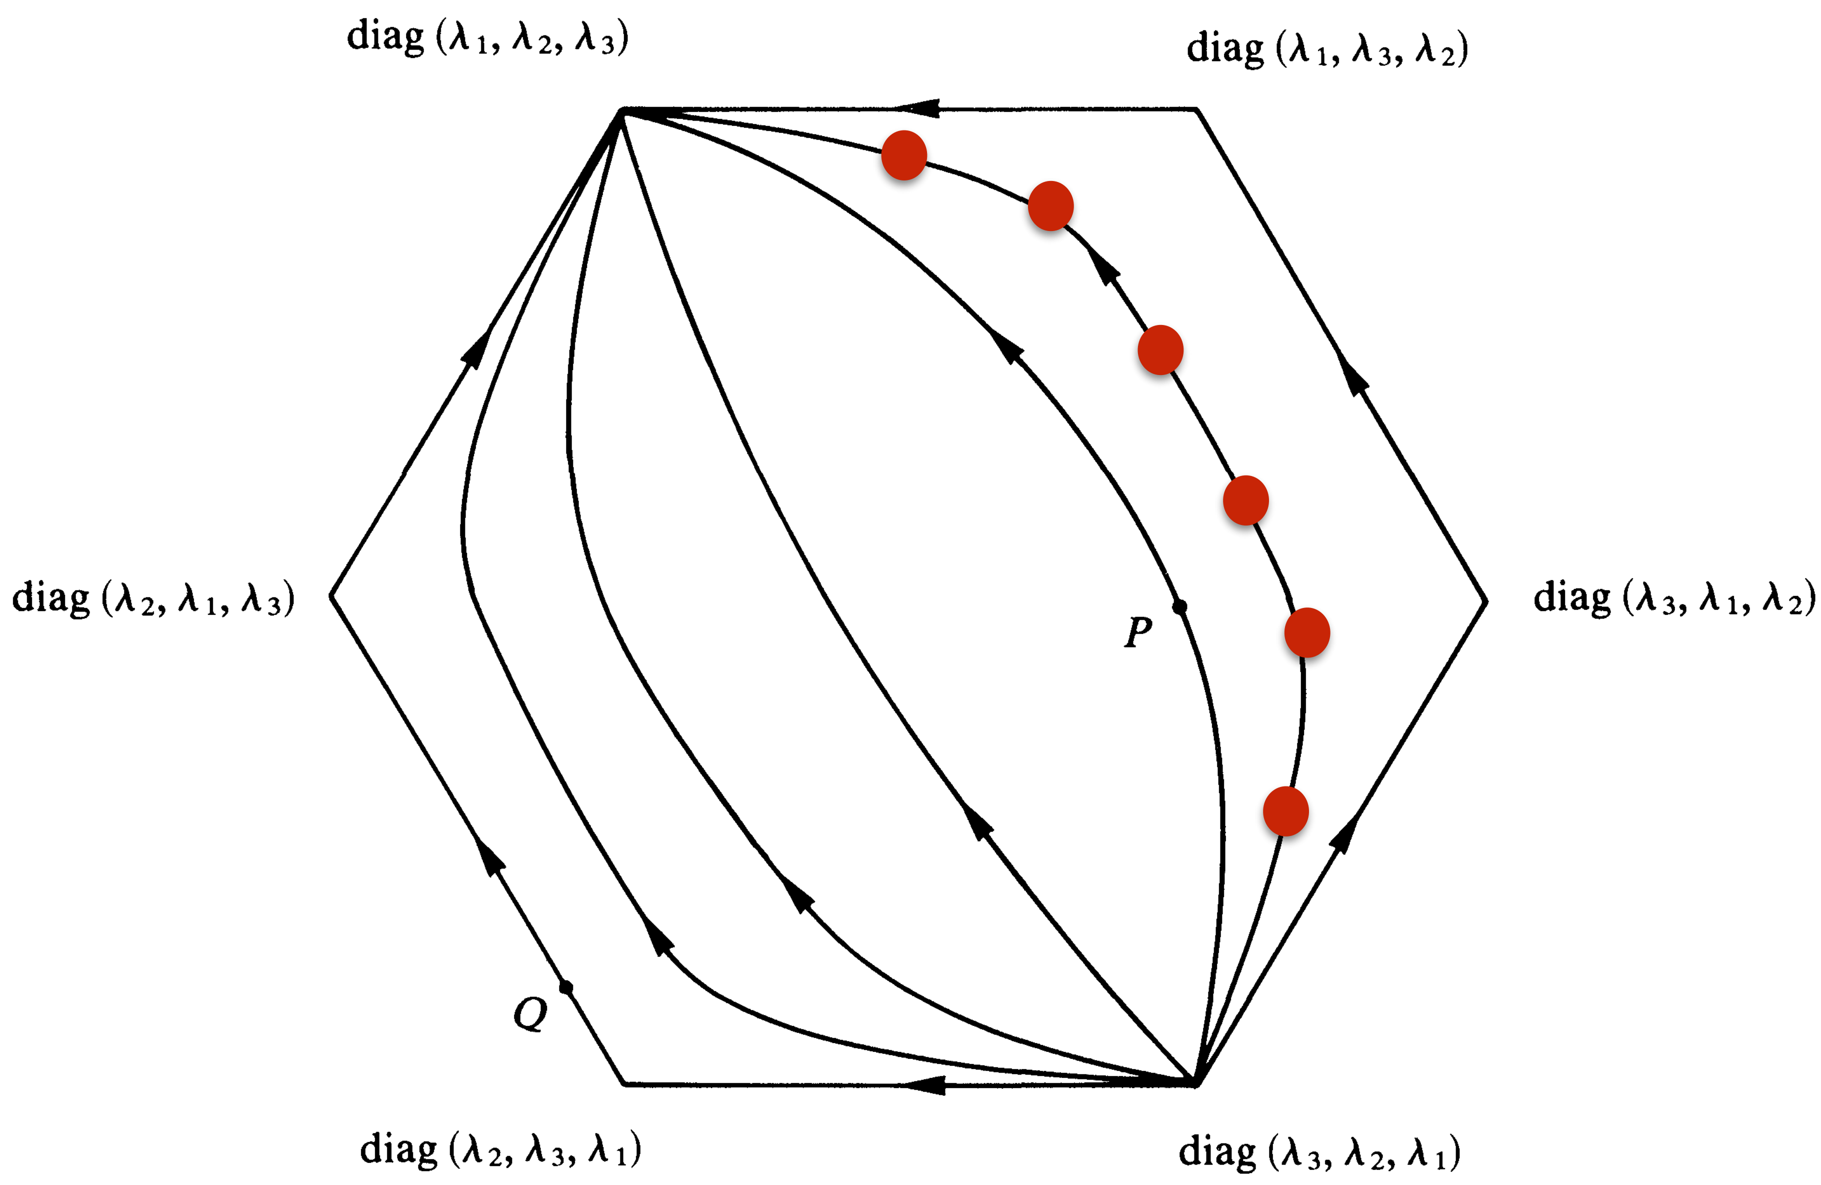

Hamiltonian = algorithm

The Toda algorithm:

$$G(x) = \frac{x^2}{2}, \quad H(L) = \frac{1}{2}\mathrm{Tr}\, L^2.$$

The QR algorithm:

$$G(x) = x\left(\log x - 1\right), \quad H(L) = \mathrm{Tr}\left(L \log L - L\right).$$

The signum algorithm (Pfrang, Deift, M.):

$$G(x) = |x|, \quad H(L) = \mathrm{Tr}\,|L|.$$

Summary (part 1)

Unexpected connections between dynamical systems and fundamental iterative algorithms were discovered in the 1980s.

Many numerical linear algebra algorithms were shown to be time-1 maps of completely integrable Hamiltonian flows by Deift, Demmel, Li, Tomei and Nanda. Moody Chu and Watkins investigated these ideas, from a numerical linear algebra viewpoint.

Certain sorting and dynamic programming algorithms, can also be recast as gradient flows (Brockett, Bloch, Fabyusovich). This includes, for example, a flow that solves the assignment problem -- so "purely discrete" problems can be solved by "continuous" methods.

Bayer and Lagarias found gradient flows that underlie Karmarkar's algorithm and other interior point methods.

Part 2. Random matrices and empirical universality

Motivation

An important feature of many numerical algorithms is that ``typical behaviour" is often much better than ``worst case" behaviour. It is interesting to try to quantify this. Some examples:

1) Testing Gaussian elimination on random matrices: (Goldstine, von Neumann 1947, Demmel 1988, Edelman, 1988).

2) Average runtime for the simplex method (Smale, 1983).

3) Smoothed analysis (Spielman, Sankar, Teng, 2004).

Deflation as a stopping criterion

When computing eigenvalues we must choose a stopping criterion to decide on convergence. Typically, we try to split the matrix L into blocks as follows:

$$L = \begin{pmatrix} L_{11} & L_{12} \\ L_{21} & L_{22} \end{pmatrix}, \quad \tilde{L} = \begin{pmatrix} L_{11} & 0 \\ 0 & L_{22} \end{pmatrix}.$$

We call this <u>deflation</u>. For a given tolerance, we define the <u>deflation time</u>

$$T_\varepsilon = \min\{k \geq 0 | \max_j |\lambda_j(L_k) - \lambda_j(\tilde{L}_k)| < \varepsilon\}$$

Christian Pfrang's thesis (2011)

Empirically investigates the performance of three eigenvalue algorithms on random matrices from different ensembles.

The algorithms:

(1) QR -- with and without shifts.
(2) The Toda algorithm.
(3) The signum algorithm.

The random matrix ensembles:

(1) Gaussian Orthogonal Ensemble (GOE).
(2) Gaussian Wigner ensemble.
(3) Bernoulli ensemble.
(4) Hermite-1 ensemble.
(5) Jacobi uniform ensemble (JUE).
(6) Uniform doubly stochastic Jacobi ensemble (UDSJ)

The numerical experiments

Choose:

(1) an algorithm (QR, Toda, signum)

(2) a deflation tolerance in the range: 1e-2 -- 1e-12

(3) a matrix size between 10--200

(4) an ensemble (1--6)

Compute the deflation time and deflation index for a large number of random samples (typically 5000--10000). In this manner, he generates empirical distributions of the deflation time and deflation index.

(Various computational tricks are used to streamline this process).

Examples of empirical distributions

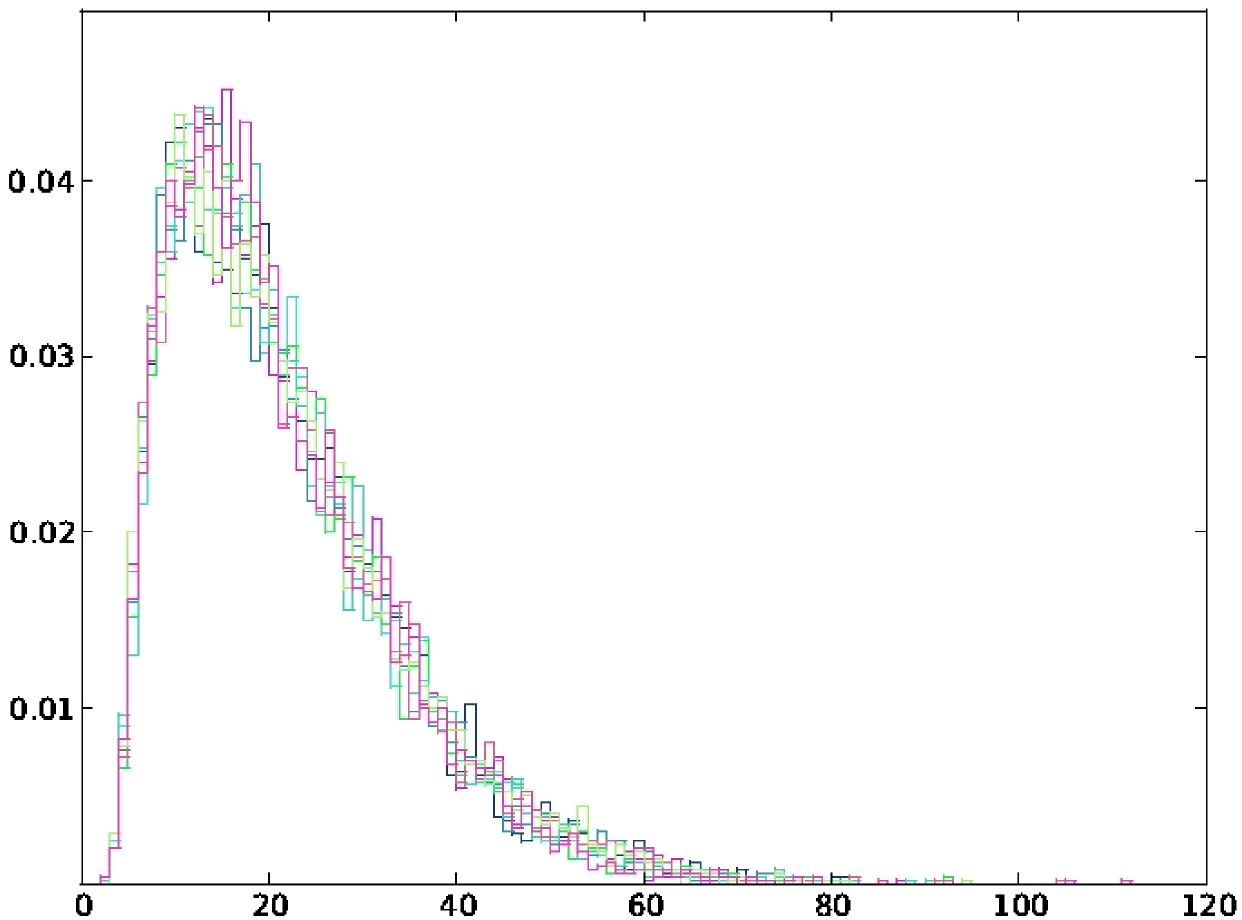

unshifted QR, GOE data

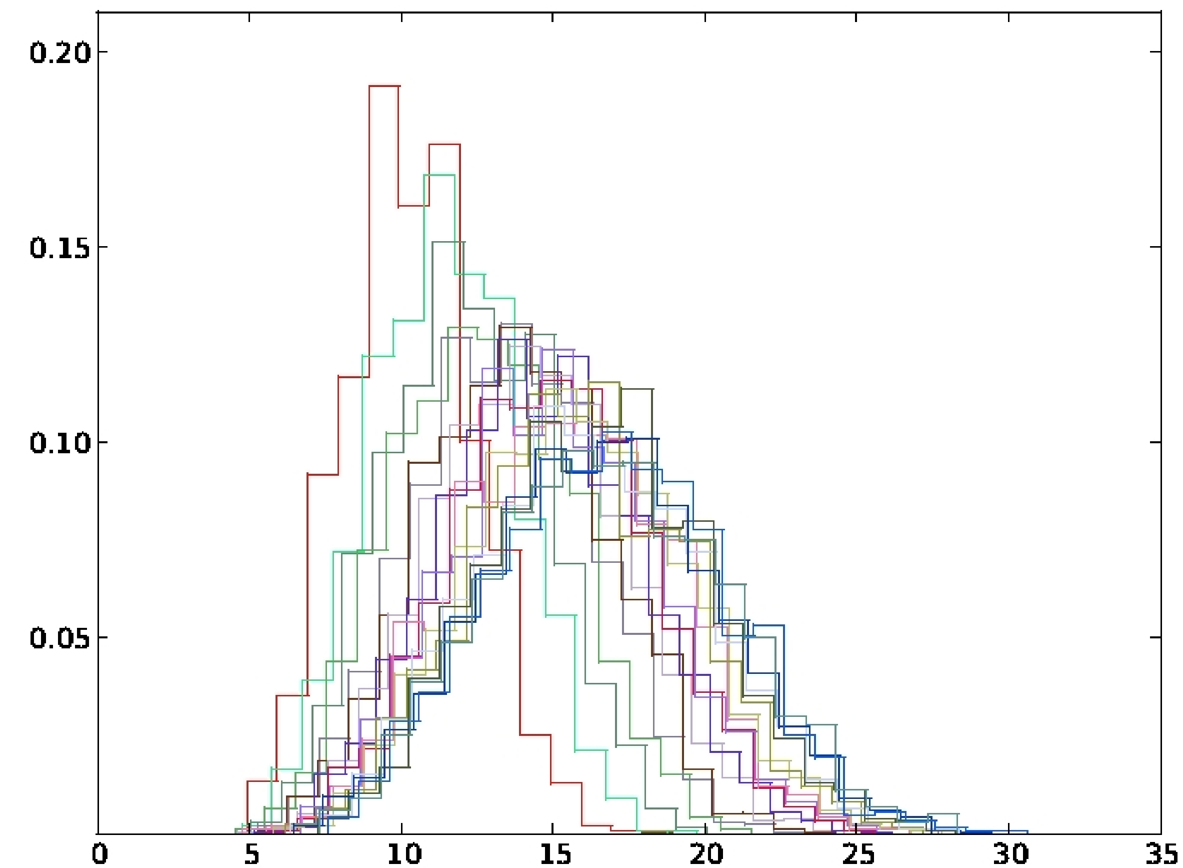

Toda, GOE data

These are histograms of the deflation time for a fixed tolerance (1 e-8) and matrix sizes that range from 10, 30, 50, ..., 190.

Empirical universality

For each algorithm, the rescaled (zero mean, unit variance) empirical distributions collapse onto a single curve that depends only on the algorithm (i.e. not on the ensemble, matrix size or deflation tolerance).

Ensembles (1)-(4) have the Wigner law as limiting spectral distribution. Ensembles (5)-(6) do not. However, we find that the QR algorithm with Wilkinson shift yields the same distribution in fluctuations for these ensembles too.

Empirical deflation time statistics: unshifted QR on GOE

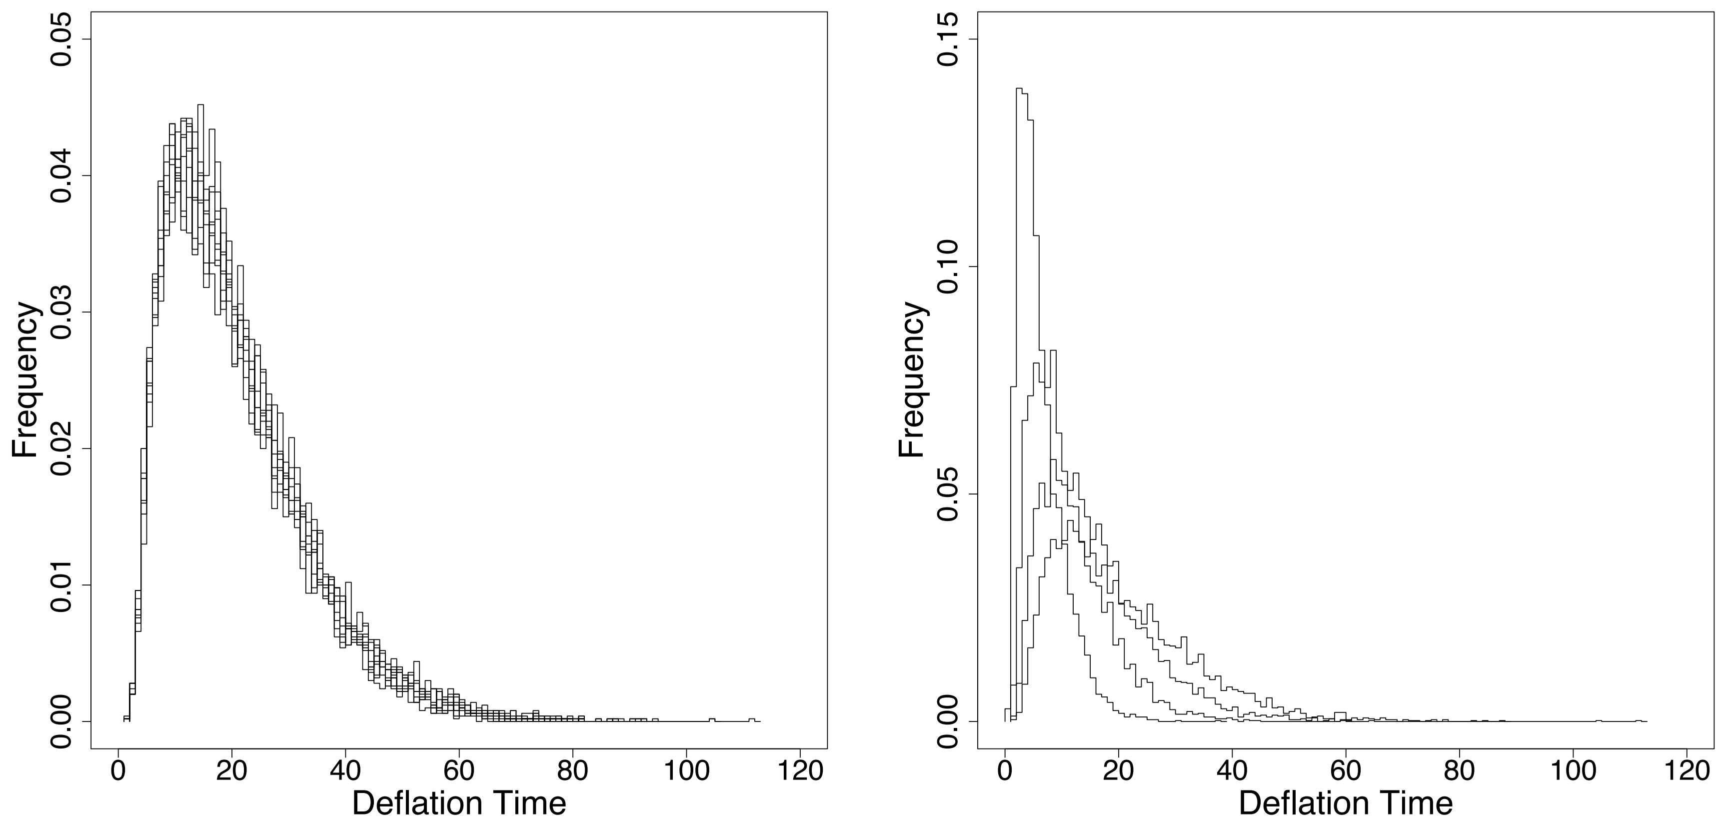

(a) (b)

Figure 1: **The QR algorithm applied to GOE.** (a) Histogram for the empirical frequency of $\tau_{n,\epsilon}$ as n ranges from $10, 30, \ldots, 190$ for a fixed deflation tolerance $\epsilon = 10^{-8}$. The curves (there are 10 of them plotted one on top of another) do not depend significantly on n. (b) Histogram for empirical frequency of $\tau_{n,\epsilon}$ when $\epsilon = 10^{-k}$, $k = 2, 4, 6, 8$ for fixed matrix size $n = 190$. The curves move to the right as ϵ decreases.

Empirical universality for QR with shifts

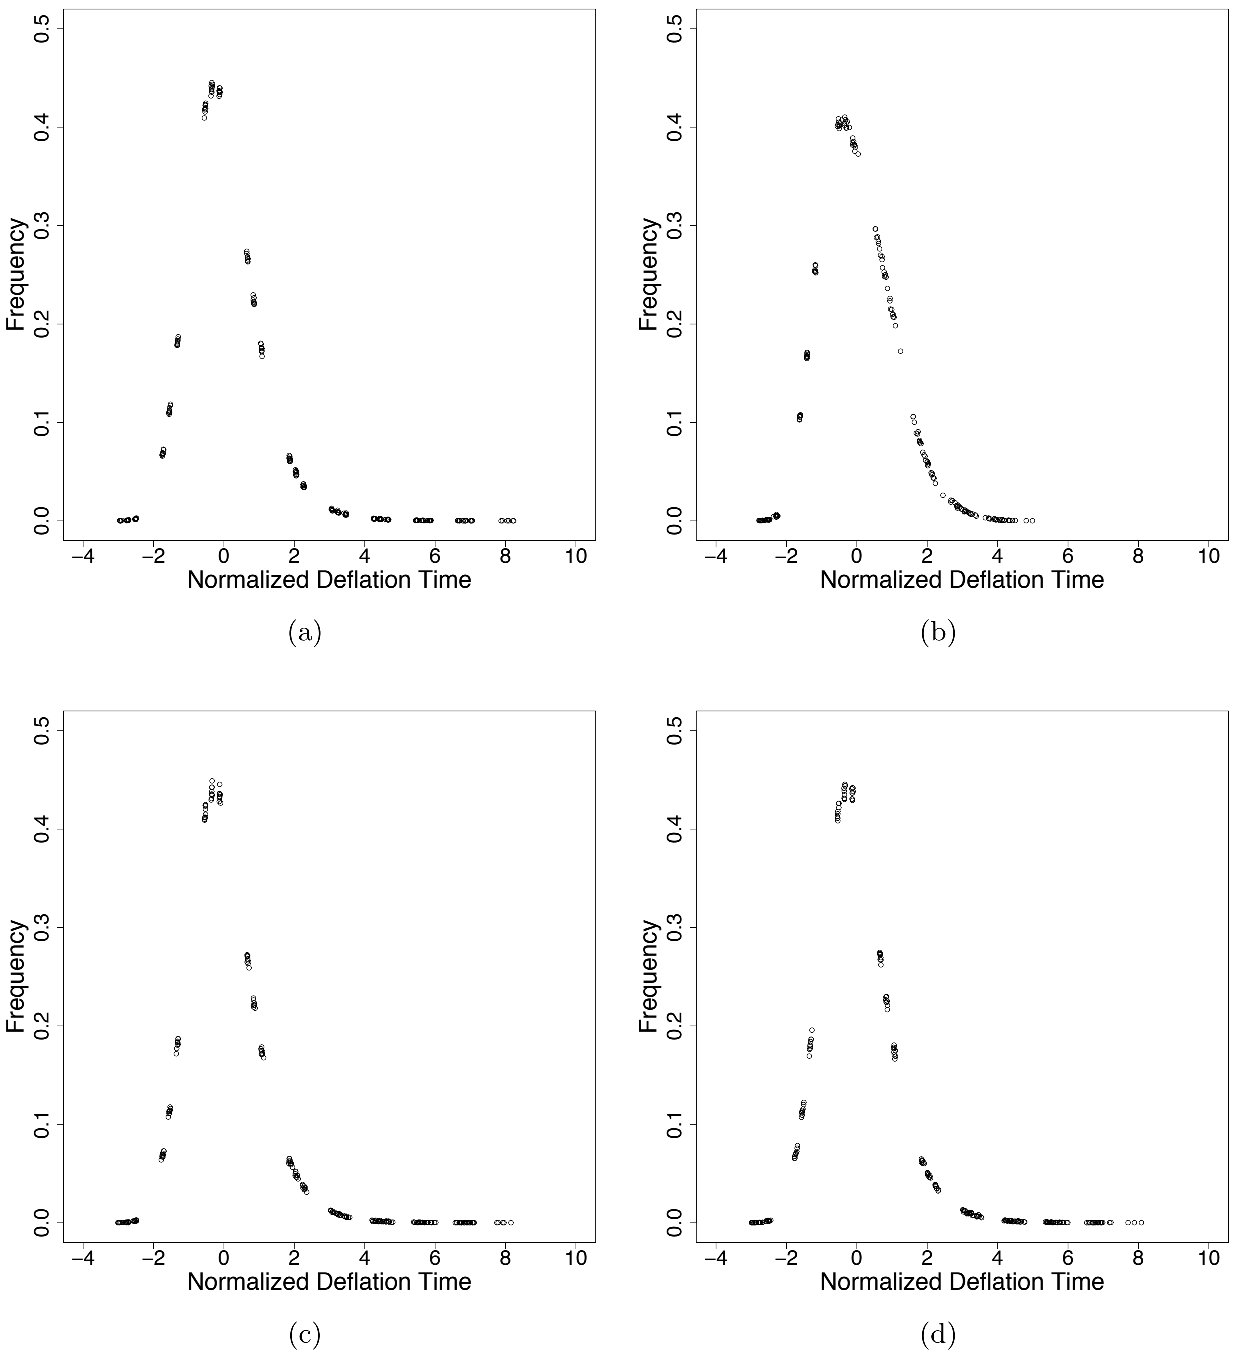

The QR algorithm with Wilkinson shift takes very few iterations to deflate. This behavior is insensitive to matrix size.

Nevertheless, by choosing very small deflation thresholds, we can begin to see a universal empirical distribution for deflation times.

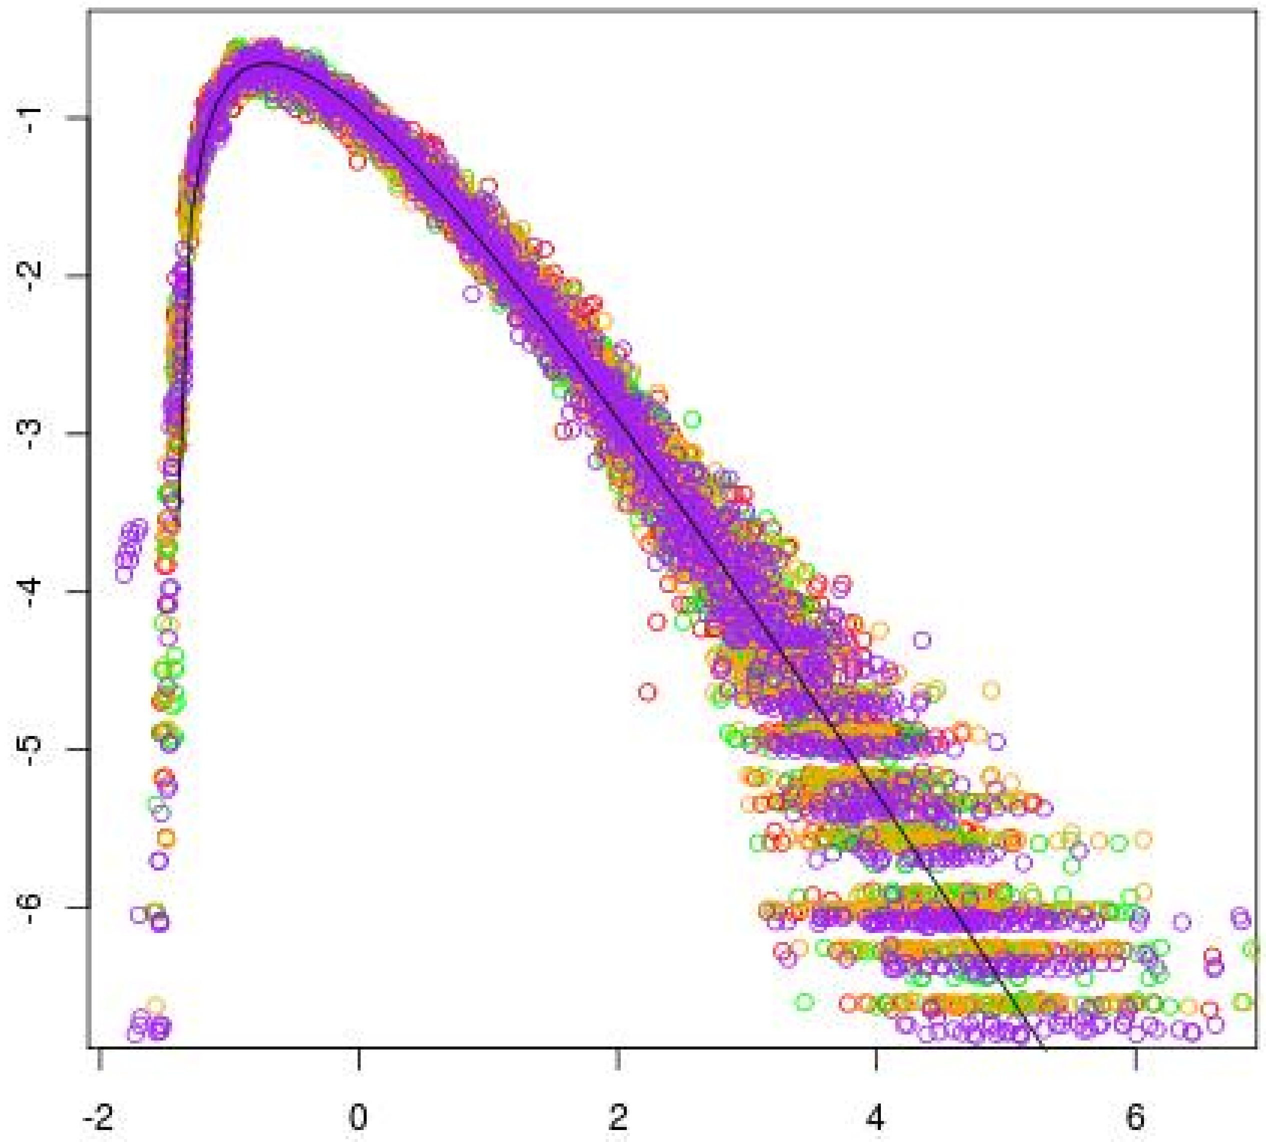

unshifted QR.
Fit is Gamma(2,1) with
exponential tail.

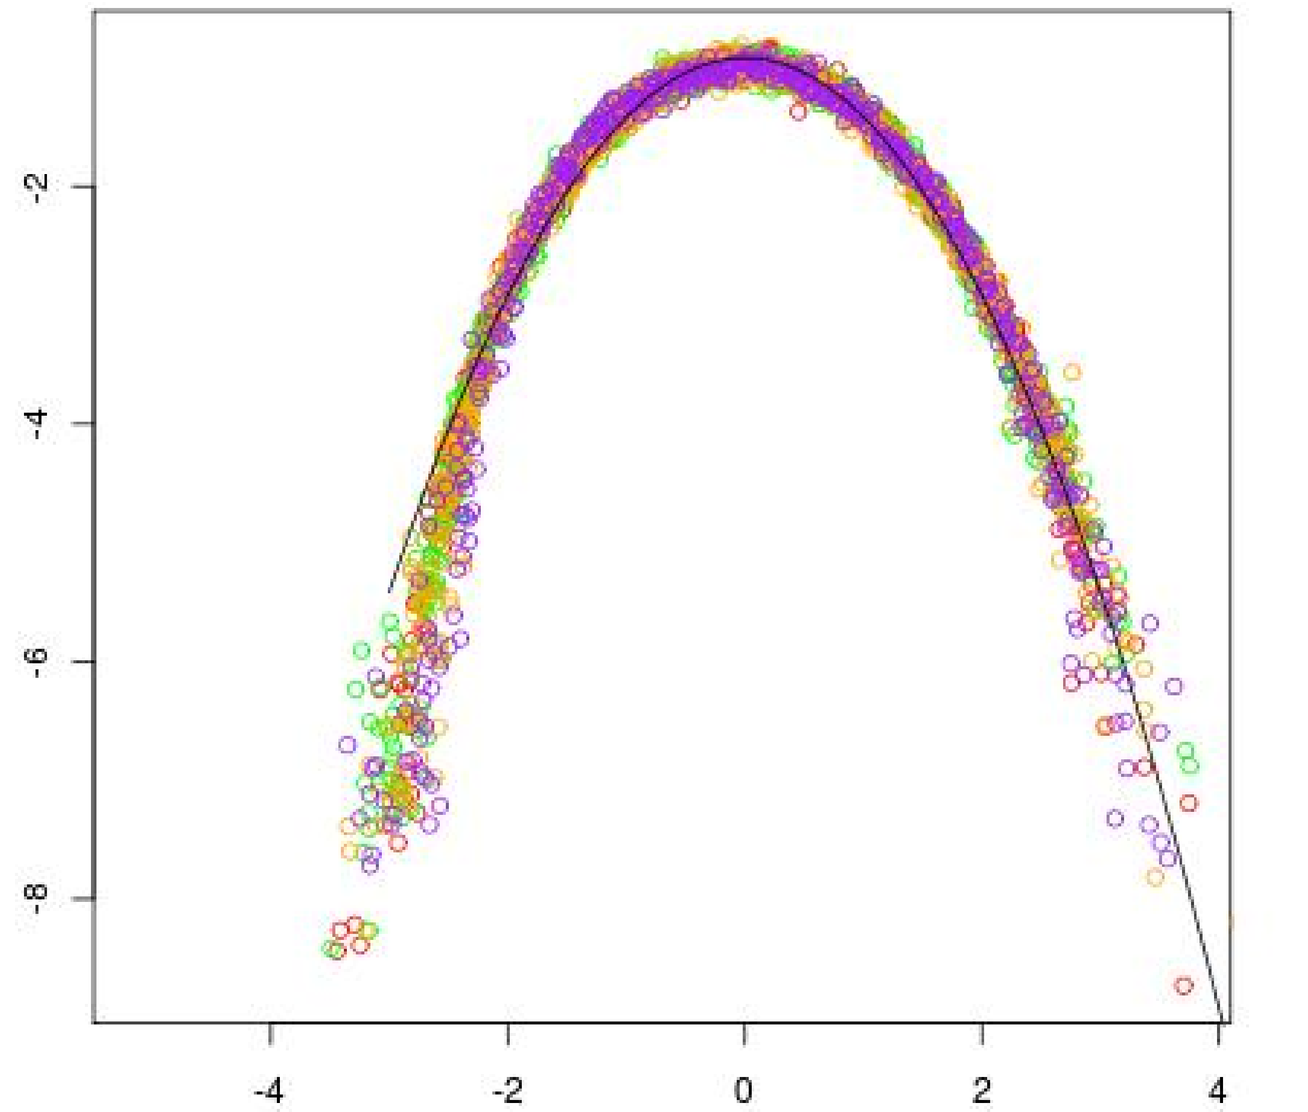

Toda.
Fit is standard normal.

Rescaled histograms for all matrix sizes, deflation tolerance, and
ensembles (1)--(4) (GOE, Bernoulli, Gaussian Wigner, Hermite-1).

Summary (part 2)

Numerical experiments reveal universal fluctuations in the distribution of the deflation time for three Hamiltonian eigenvalue algorithms.

Such "universality in halting times" is not restricted to these algorithms. Tom Trogdon performed extensive numerical computations on several algorithms, and found similar universality. These include a critical scaling regime for the conjugate gradient method, and an interesting decision time algorithm.

These results are reported in Deift, M. , Olver, Trogdon (PNAS, 2014).

Part 3. The largest gap conjecture, gradient flows on the permutahedron and sorting networks

The largest gap conjecture

Denote the largest gap in the spectrum by

$$g_n = \max_{1 \leq j \leq n-1} (\lambda_j - \lambda_{j+1})$$

Our guess: The observed "universal distribution" for the Toda flow is the edge-scaling limit of $1/g_n$. QR and signum are more subtle.

The largest gap conjecture: numerical evidence

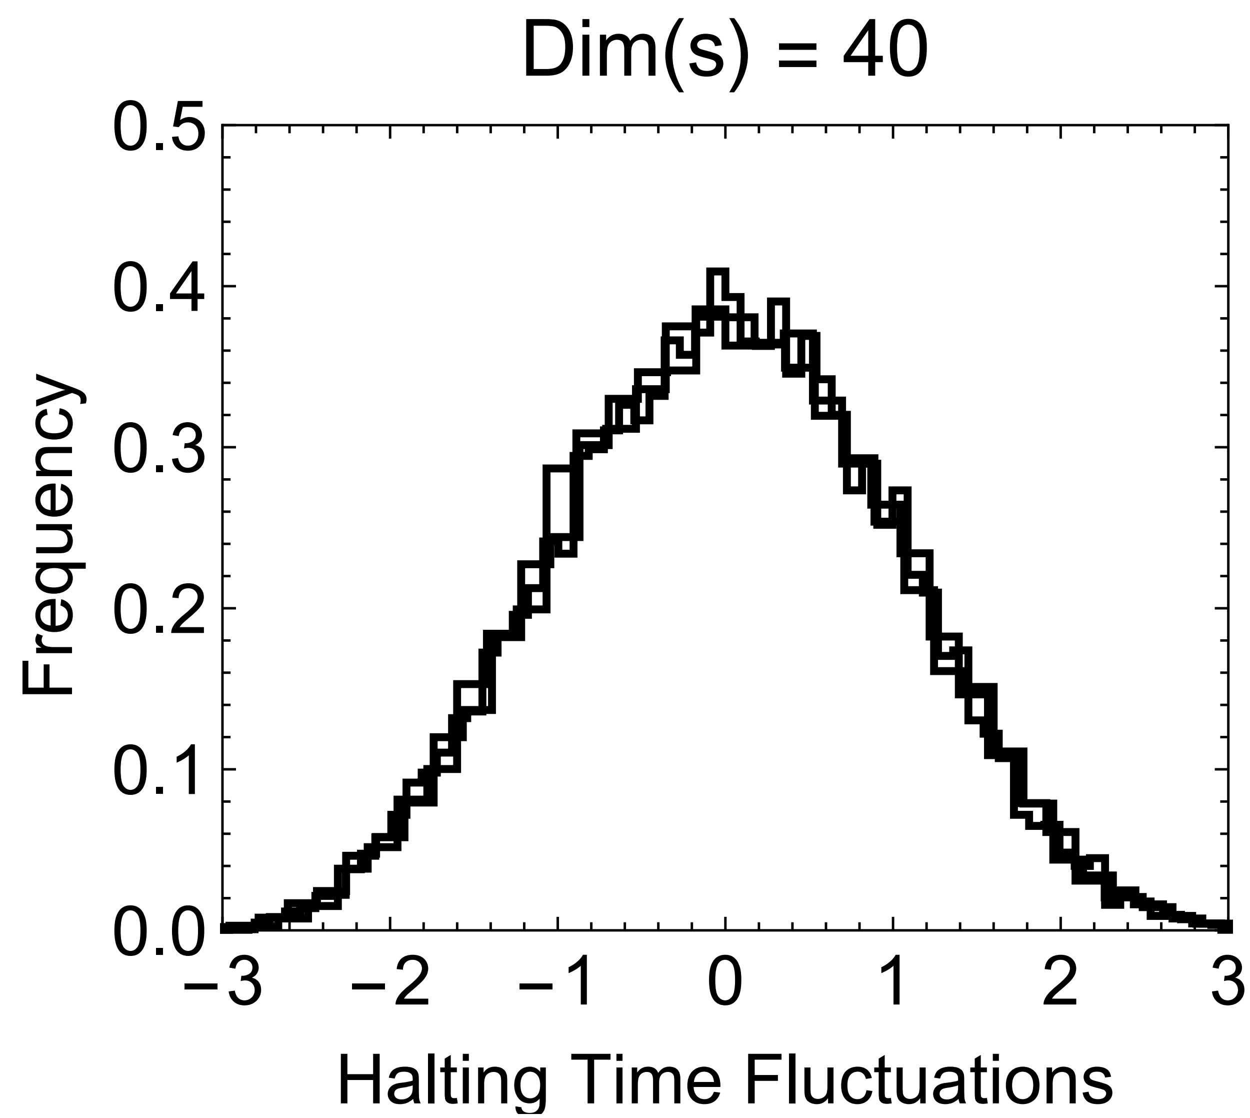

Overlay of statistics of the largest gap and observed deflation time.

The largest gap via stochastic Airy

The edge-scaling limit of the spectrum is given by the Stochastic Airy operator.

$$-\frac{d^2\varphi}{dx^2} + \left(x + 2\dot{b}\right)\varphi = \lambda\varphi, \quad 0 < x < \infty, \quad \varphi(0) = \lim_{x\to\infty}\varphi(x) = 0.$$

Heuristically, the spectrum of stochastic Airy is a random perturbation of the spectrum of the Airy operator, i.e. minus the zeros of the Airy function

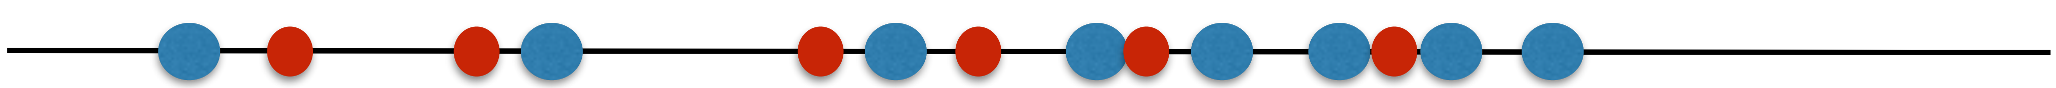

Eigenvalues of Airy. Eigenvalues of stochastic Airy.

An explicit description of the law of the largest gap is not known.
Even the fact that it is well-defined requires some care...

Why the largest gap?

Conventional wisdom: Toda rates are determined by the smallest gap, since the smallest gap determines the rate of convergence to equilibrium (Moser).

But convergence to equilibrium is not the same as deflation.

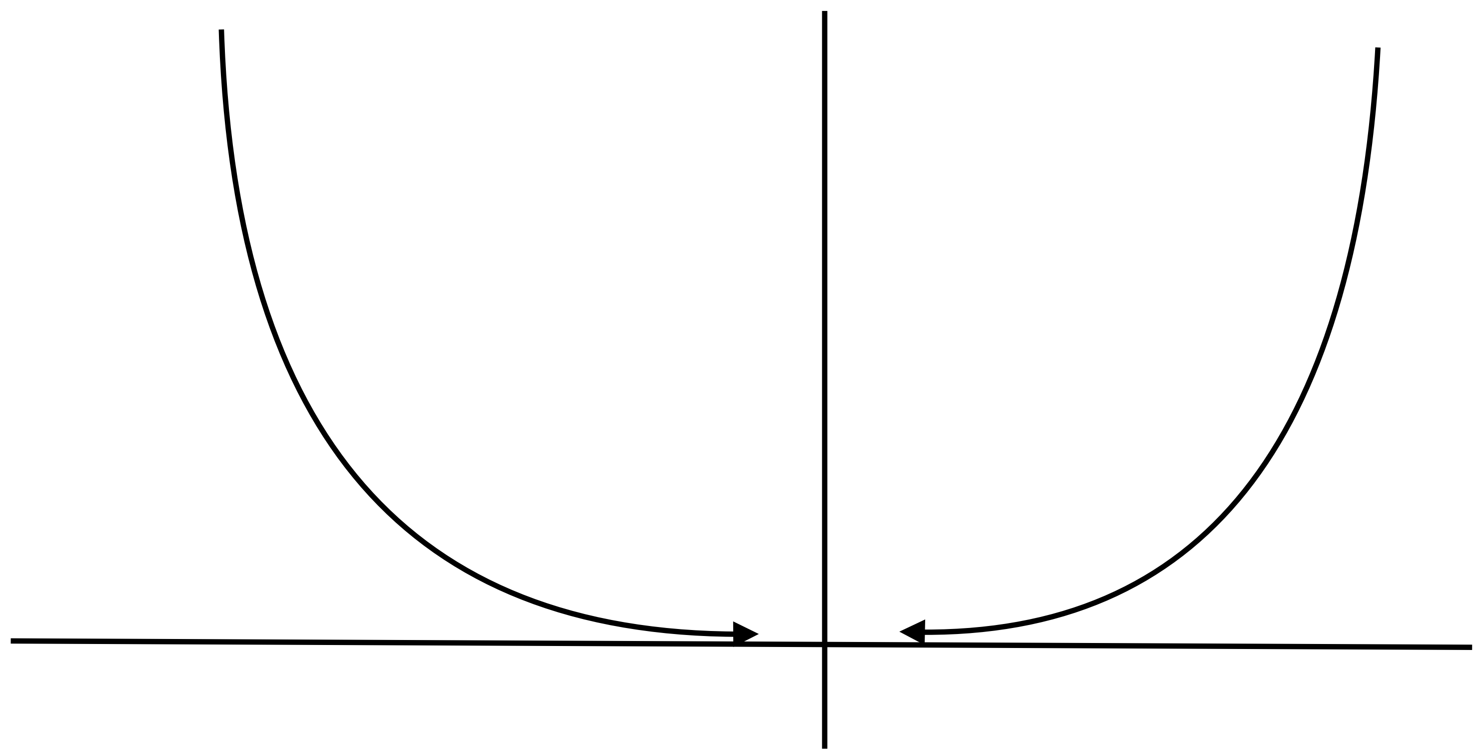

$$\dot{x} = -\mu x, \quad \dot{y} = -\lambda y, \quad 0 < \mu \ll \lambda.$$

Convergence to equilibrium occurs at slow rate μ

Approach to x-axis occurs at fast rate λ

Understanding deflation (M, Pujals)

Main new idea: Combine analysis of Toda with combinatorics of permutations.

Invariant manifolds:

Each (n-k)-face on the polytope is invariant.

Each such face corresponds to block diagonal matrices with k blocks.

$=3$

$$L = \begin{pmatrix} L_1 & 0 & 0 & \cdots & 0 \\ 0 & L_2 & 0 & & \cdots \\ 0 & 0 & L_3 & 0 & \cdots \\ & \ddots & \ddots & \vdots & 0 \\ 0 & & \cdots & 0 & L_k \end{pmatrix}$$

Toda phase space for 4x4 matrices

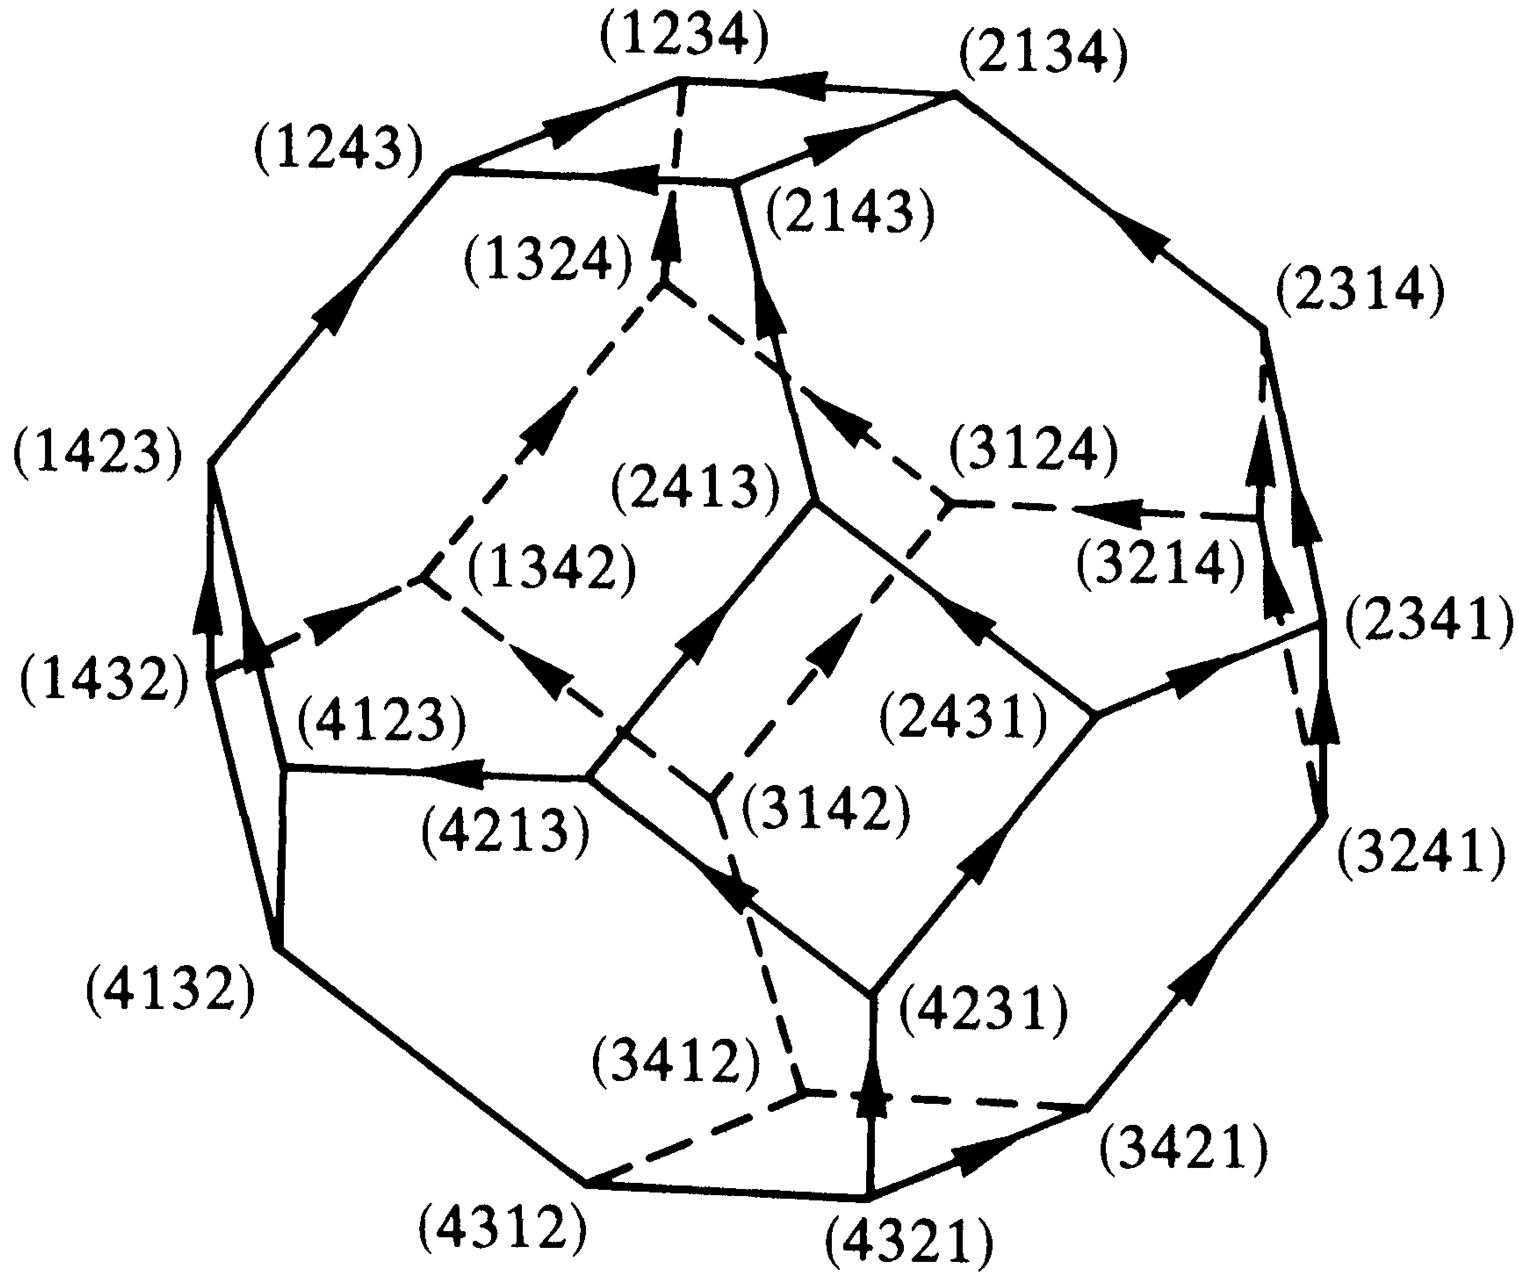

Deift, Nanda, Tomei (SIAM J. Numer. Anal., 1983).

Normal hyperbolicity (M, Pujals)

Main observation : order relations between eigenvalues of blocks imply normal hyperbolicity of invariant manifolds.

Simplest example (enough to understand largest gap):

The stable equilibrium lies at the intersection of (n-1) invariant manifolds.

Each such manifold is of the form $b_k = 0, \quad 1 \leq k \leq n-1.$

The normal rate of attraction for any such manifold is at least $\lambda_k - \lambda_{k+1}.$

Easy consequence of Schur-Horn theorem for Toda.

Analogous statement requires more care for other flows.

From deflation to global dynamics...

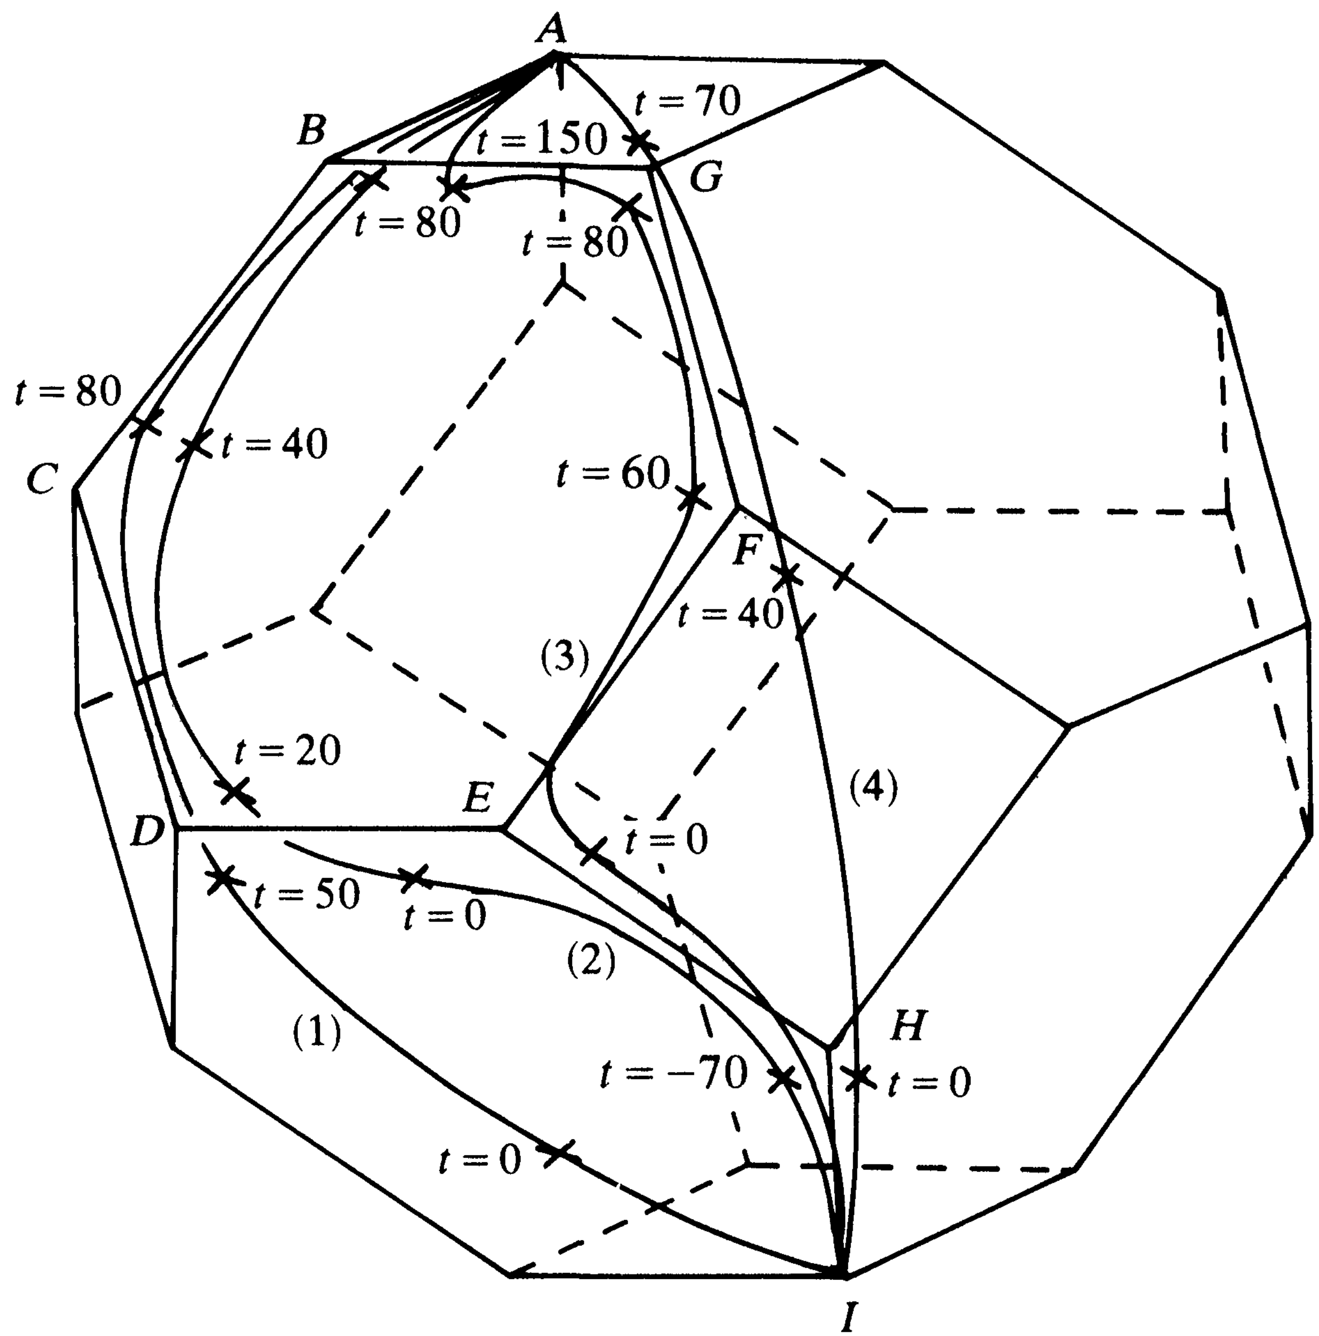

Deift, Nanda, Tomei (SIAM J. Numer. Anal., 1983).

Sorting networks

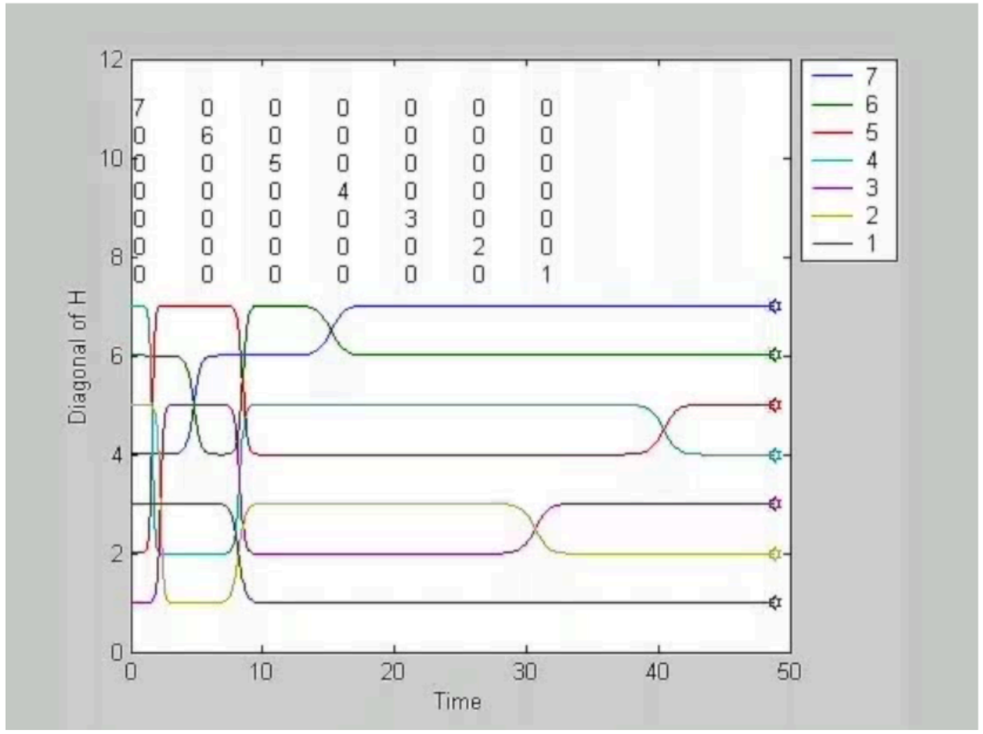

Projection onto diagonal of a Jacobi matrix evolving by Toda (Roger Brockett).

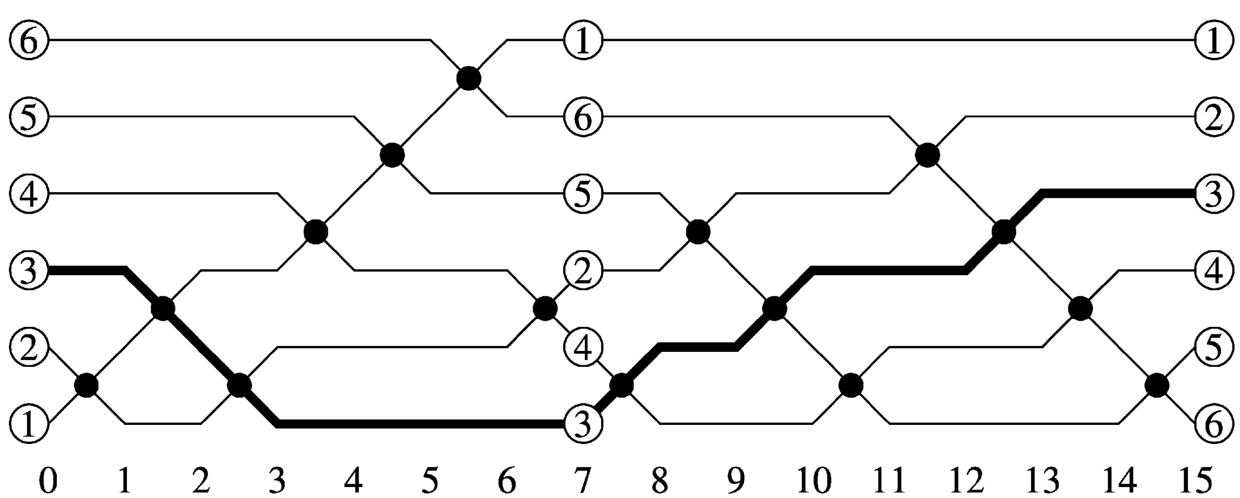

Angel, Holroyd, Romik, Virag-- Uniform sorting networks (2007)

From random matrix theory to a sorting network

With probability 1, each choice of initial condition determines a sorting network.

Law of the sorting network is determined by the scattering theorem for Toda (and related flows).

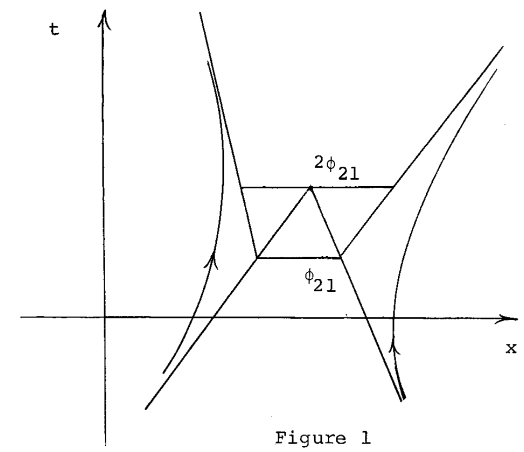

Figure 1

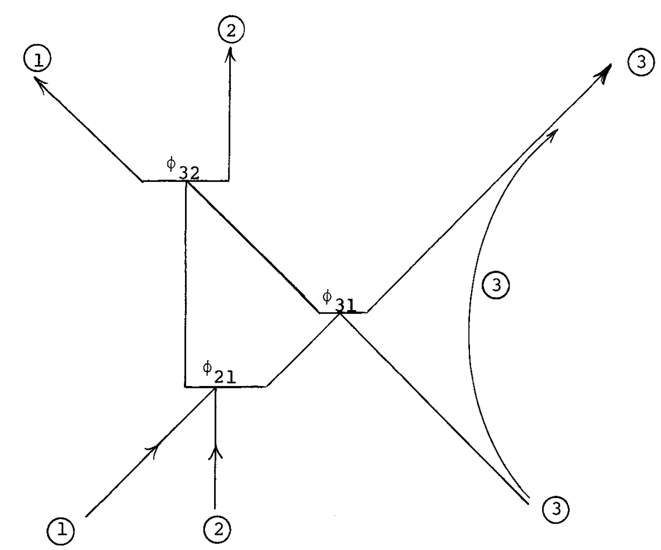

Figure 2

J. Moser -- Finitely many mass-points on the line (1976)

Summary

(1) Several iterative eigenvalue algorithms are tied to integrable systems.

(2) Empirical universality for fluctuations of deflation times emerges from universality of gap statistics for Tracy-Widom point process.

(3) Each algorithm and matrix ensemble determines a class of random sorting networks.

It seems reasonable to hope that universality could hold for the sorting networks (this has not been tested).