New Trends in Polynomial Homotopy Continuation
a priori step size control

Jan Verschelde[†]

University of Illinois at Chicago
Department of Mathematics, Statistics, and Computer Science
http://www.math.uic.edu/~jan
janv@uic.edu

2021 SIAM Conference on Applied Algebraic Geometry
Minisymposium on New Trends in Polynomial System Solving
16-20 August 2021, online

[†]Supported by the National Science Foundation, grant DMS 1854513.

Outline

1. Problem Statement
   - step size control
   - numerical continuation
2. A Priori Step Size Control
   - schematic overview
   - the Fabry-Hesse-Newton-Padé predictor
   - cost analysis and computational results
3. Trends in Polynomial Homotopy Continuation
   - parallel algorithms
   - solving power series systems

problem statement

To solve a polynomial system, we apply homotopy continuation, in two stages:

1. A homotopy method constructs a family of polynomial systems, connecting the input system to a system with known solutions.
2. A continuation method tracks the solution paths originating at the known solutions leading to the solutions of the input system.

Consider homotopies in one single parameter and assume

1. no singularity on each path, and
2. no diverging paths.

Problem: determine the step size of the path tracker.

- Too small: inefficient.
- Too large: jump off the path, possibly onto another path.

tracking a curve

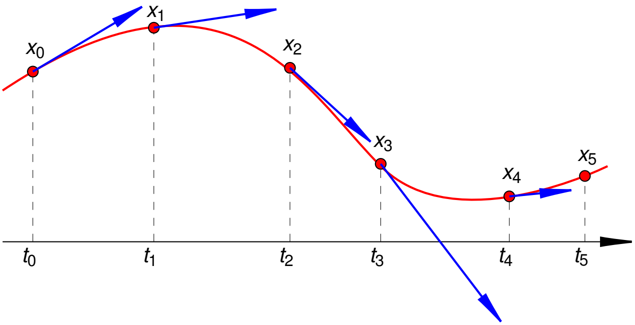

the main references for this talk

- N. Bliss and J. Verschelde. **The method of Gauss–Newton to compute power series solutions of polynomial homotopies.** *Linear Algebra and its Applications*, 542:569–588, 2018.
- S. Telen, M. Van Barel, and J. Verschelde. **A Robust Numerical Path Tracking Algorithm for Polynomial Homotopy Continuation.** *SIAM Journal on Scientific Computing* 42(6):A3610–A3637, 2020.
- S. Telen, M. Van Barel, and J. Verschelde. **Robust numerical tracking of one path of a polynomial homotopy on parallel shared memory computers.** In the *Proceedings of the 22nd International Workshop on Computer Algebra in Scientific Computing (CASC 2020)*, pages 563–582. Springer-Verlag, 2020.

numerical continuation

Numerical continuation applies an adaptive step size control in a predictor-corrector method with double precision arithmetic.

Alternatives to numerical continuation:

- interval or ball arithmetic
  [Kearfott and Xing, 1994], [Lecerf and van der Hoeven, 2016];
- symbolic deformation methods
  [Jeronimo, Matera, Solernó, and Waissbein, 2009],
  [Hauenstein, Safey El Din, Schost, Vu, 2021];
- certified homotopy tracking
  [Beltrán and Leykin, 2013], [Xu, Burr, and Yap, 2018].

Padé approximants as predictors

1. H. Schwetlick and J. Cleve. Higher order predictors and adaptive steplength control in path following algorithms. *SIAM Journal on Numerical Analysis*, 24(6):1382–1393, 1987.
2. A. Trias. The holomorphic embedding load flow method. In *2012 IEEE Power and Energy Society General Meeting*, pages 1–8. IEEE, 2012.
3. A. Trias and J. L. Martin. The holomorphic embedding loadflow method for DC power systems and nonlinear DC circuits. *IEEE Transactions on Circuits and Systems*, 63(2):322–333, 2016.

The holomorphic embedding load flow method takes the poles of the Padé approximants into account in its step size control.

a posteriori and a priori step size control

To solve a polynomial system $f(x) = 0$, a typical homotopy is

$$h(x,t) = \gamma(1-t)g(x) + t\,f(x) = 0, \quad t \in [0,1], \ \text{random } \gamma \in \mathbb{C}.$$

An *a posteriori* step size control uses feedback loops.

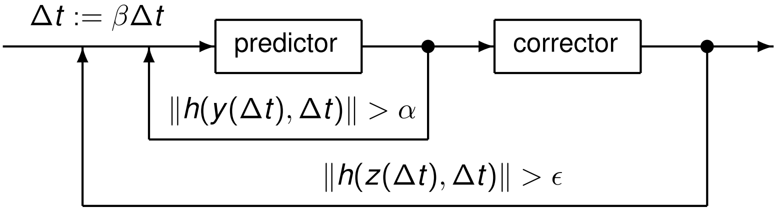

Extreme choices for α and ϵ (not recommended):

- If $\alpha \leq \epsilon$, then the corrector is not needed.
- If $\alpha = \infty$, then the first feedback loop does never happen.

Setting 0.5 for β cuts the step size Δ in half.

An *a priori* step size control does not need feedback loops.

two causes of path jumping

Curves are far apart, with high curvature:

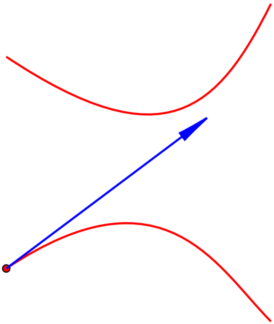

Curves are close to each other, with low curvature:

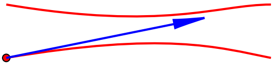

detecting nearby singularities

Applying the ratio theorem of Fabry, we can detect singular points based on the coefficients of the Taylor series.

Theorem (the ratio theorem, Fabry 1896)

If for the series $x(t) = c_0 + c_1 t + c_2 t^2 + \cdots + c_n t^n + c_{n+1} t^{n+1} + \cdots$, *we have* $\lim_{n \rightarrow \infty} c_n / c_{n+1} = z$, *then*

- *z is a singular point of the series, and*
- *it lies on the boundary of the circle of convergence of the series.*

Then the radius of this circle is less than $|z|$.

The ratio c_n / c_{n+1} is the pole of Padé approximants of degrees $[n/1]$ (n is the degree of the numerator, with linear denominator).

error analysis of a lower triangular block Toeplitz solver

Solving $(A_0 + A_1 t + A_2 t^2 + \cdots + A_d t^d)(x_0 + x_1 t + x_2 t^2 + \cdots + x_d t^d)$
$= (b_0 + b_1 t + b_2 t^2 + \cdots + b_d t^d)$

leads to a lower triangular block system:

$$\begin{bmatrix} A_0 & & & & \\ A_1 & A_0 & & & \\ A_2 & A_1 & A_0 & & \\ \vdots & \vdots & \vdots & \ddots & \\ A_d & A_{d-1} & A_{d-2} & \cdots & A_0 \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ \vdots \\ x_d \end{bmatrix} = \begin{bmatrix} b_0 \\ b_1 \\ b_2 \\ \vdots \\ b_d \end{bmatrix}.$$

Cost to solve: $O(n^3) + O(dn^2)$.

Let κ be the condition number of A_0. Let $\|A_0\| = \|x_0\| = 1$, $\|x_d\| \approx \rho^d$.

In our context, $\rho \approx 1/R$, where R is the convergence radius.

If $\|A_d\| \approx \rho^d$, then $\dfrac{\|\Delta x_d\|}{\|x_d\|} \approx \kappa^{d+1}\epsilon_{\rm mach}$, and accuracy is lost.

estimating the distance to the nearest path

Consider a Taylor series expansion of the homotopy at one path, truncated after degree 2, to estimate the distance to the nearest path.

The distance $\|\Delta \mathbf{z}\|$ to the nearest path is estimated by

$$\eta = \frac{2\sigma_n(J_h)}{\sqrt{\sigma_{1,1}^2 + \sigma_{2,1}^2 + \cdots + \sigma_{n,1}^2}} \lesssim \|\Delta \mathbf{z}\|,$$

where

- $\sigma_n(J_h)$ is the smallest singular value of the Jacobian matrix,
- $\sigma_{i,1}$ is the largest singular value of the Hessian matrix at the i-th polynomial in the homotopy h.

With Padé approximants p_i/q_i we compute an estimate for the error e_0:

$$\left\| x(\Delta t) - \left(\frac{p_1(\Delta t)}{q_1(\Delta t)}, \ldots, \frac{p_n(\Delta t)}{q_n(\Delta t)} \right) \right\| \approx \|e_0\| \, |\Delta t|^k,$$

where k is determined by the degrees of the Padé approximants.

schematic summary of a priori step size control

$$(h, J_h, z_{t^*}^{(i)}, t^*)$$

Newton	Differentiate
$\mathbf{x}(t)$ power series	$\mathcal{H}_i$ Hessians
Padé	SVD
$\frac{p_j(t)}{q_j(t)}$ $\|e_0\|$	$\sigma_n(J_h)$ $\sigma_1(\mathcal{H}_k)$
$\Delta t_2 = \beta_2 D$ ‘pole distance’	$\Delta t_1 = \sqrt[k]{\frac{\beta_1 \eta}{\|e_0\|}}$ ‘curvature bound’

$$\Delta t = \min(\Delta t_1, \Delta t_2, t_{\text{end game}} - t^*)$$

The values β_1 and β_2 are experimentally defined tolerances.

cost analysis

For n variables,

- the cost of the linear algebra is $O(n^4)$,
- the cost to differentiate and evaluate n Hessians is 2n times the cost of computing the Jacobian,
- for power series truncated at degree n, the cost overhead factor of Newton's method is $O(n \log(n))$.

Relative to a posteriori step size control,
the cost overhead of a priori step size control is $O(n \log(n))$.

Use parallel computers to offset the cost overhead.

computational results

Available in PHCpack since v2.4.72, released 1 September 2019. To track a large number of paths, a static workload distribution message passing implementation ran on a 44-core workstation.

Two benchmarks:

1. 1,048,576 paths defined by 20 quadrics, one linear equation, the `katsura-20` benchmark from computational physics. About 66 solutions have a large condition number of about 10^7.

   HOM4PS-2.0para [Li, Tsai, Parallel Computing 2009] reported 4 path jumpings in their runs on `katsura-20`.

2. 1,594,297 paths defined by 13 cubic equations, in `noon-13`, arising in a model of a neural network.

All runs were done in double precision, no path jumpings occurred.

Homogeneous coordinate formulations are important.

a pipelined algorithm to solve matrix series

We solve $\mathbf{A}(t)\mathbf{x}(t) = \mathbf{b}(t)$ for series $\mathbf{x}(t)$,
given $\mathbf{A}(t) = A_0 + A_1 t + A_2 t^2 + \cdots$ and $\mathbf{b}(t) = b_0 + b_1 t + b_2 t^2 + \cdots$.
For example, for series truncated at degree 2:

$$\begin{array}{rcl} A_0 x_0 & = & b_0 \\ A_0 x_1 & = & b_1 - A_1 x_0 \\ A_0 x_2 & = & b_2 - A_2 x_0 - A_1 x_1 \end{array}$$

1. $F = \text{Factor}(A_0)$; $x_0 = \text{Solve}(F, b_0)$
2. for k from 1 to d do
    1. update b_ℓ with $b_\ell - A_\ell \mathbf{x}_k$ simultaneously, for ℓ from k to d
    2. $x_k = \text{Solve}(F, b_k)$

With d threads, the speedup is then $1 + \dfrac{d(d-1)}{4(d+1)}$.
As $d \rightarrow \infty$, this ratio equals $1 + d/4$.

accelerated polynomial evaluation and differentiation

Evaluating polynomials at power series of degree $d = 152$ in deca double precision (CAMPARY software), on five different GPUs. The last line is the wall clock time for all convolution and addition kernels. All units are milliseconds.

	C2050	K20C	P100	V100	RTX 2080
convolution	12947.26	11290.22	1060.03	634.29	10002.32
addition	10.72	11.13	1.37	0.77	5.01
sum	12957.98	11301.35	1061.40	635.05	10007.34
wall clock	12964.00	11309.00	1066.00	640.00	10024.00

- The $12964/640 \approx 20.26$ is for the V100 over the oldest C2050.
- Compare the ratio of the wall clock times for P100 over V100 $1066/640 \approx 1.67$ with the ratios of theoretical double peak performance of the V100 of the P100: $7.9/4.7 \approx 1.68$.

IPDPSW 2021

solving power series systems

Solving with power series $\Rightarrow$ solving systems of power series.

Consider input coefficients as power series given up to some degree:

$$\left\{ \begin{array}{rcl} \sin(t) - y & = & 0 \\ x^2 + y^2 & = & 1. \end{array} \right.$$

The solution $x = \cos(t)$ is obtained as a power series,
running Newton's method on

$$\left\{ \begin{array}{rcl} t - 1/6t^3 + 1/120t^5 - 1/5040t^7 - y & = & 0 \\ x^2 + y^2 & = & 1 \end{array} \right.$$

yields $x = 1 - 1/2t^2 + 1/24t^4 - 1/720t^6 + 1/40320t^8$.

Laurent series are needed if $\sin(t)$ and $\cos(t)$ are flipped in the input.

conclusions

Trends in Polynomial Homotopy Continuation:

- a priori step size control,
- parallel algorithms,
- multiple double arithmetic,
- systems of power series.

GNU GPL software at
`https://github.com/janverschelde/PHCpack`.

Link to a prerecording of this talk:
`https://youtu.be/AUFpEkZYtLI`.