



UNIVERSITÀ DEGLI STUDI DI PISA

Dipartimento di Matematica
Corso di Laurea Triennale in Matematica

Laboratorio Computazionale

Parallel Homotopy Continuation in Julia

Studente: Francesco Minnocci
Matricola: 600455

ANNO ACCADEMICO 2022 - 2023

Contents

1	Introduction	2
2	Homotopy Continuation	2
2.1	Choosing the homotopy	2
2.1.1	Gamma trick	3
2.2	Tracking down the roots	3
2.2.1	Predictor: Euler's method	4
2.2.2	Corrector: Newton's method	4
2.2.3	Adaptive step size	4
3	Testing the method	5
4	Possible Improvements	5
4.1	Homogenized Coordinates	5
4.2	Predictor-Corrector	5
5	Appendix A: Results	5
5.1	Single- and Multi-threaded	5
5.2	Parallelization	7
6	Appendix B: Implementation	7
6.1	Julia code	7
6.2	Hardware	11

1 Introduction

Homotopy Continuation is a numerical method for solving systems of polynomial equations. It is based on the idea of "deforming" a given system of equations into a simpler one whose solutions are known, and then tracking the solutions of the original system as the deformation is undone.

In this project, the method will be implemented in the Julia programming language, which is particularly suited for scientific computing. The primary source for this report is [1], where the method is explained in much more detail.

2 Homotopy Continuation

We will only consider *square* systems of polynomial equations, i.e. systems of n polynomial equations in n variables, although over- or under-determined systems can often be solved by reducing them to square systems, by respectively choosing a suitable square subsystem or squaring it by adding equations. Moreover, we will restrict ourselves to systems which have isolated solutions, i.e. zero-dimensional varieties.

There are many ways to choose the "simpler" system, from now on called a *start system*, but in general we can observe that, by Bezout's theorem, a system $F = (f_1, \dots, f_n)$ has at most $D := d_1 \dots d_n$ solutions, where d_i is the degree of $f_i(x_1, \dots, x_n)$.

Therefore, we can build a start system of the same size and whose polynomials have the same degrees, but whose solutions are easy to find, and thus can be used as starting points for the method.

For instance, the system $G = (g_1, \dots, g_n)$, where

$$g_i(x_1, \dots, x_n) = x_i^{d_i} - 1,$$

is such a system, since its zero locus is obtained by combining the d_i -th roots of unity in each variable, which are exactly D points:

$$\left\{ \left(e^{\frac{k_1}{d_1} 2\pi i}, \dots, e^{\frac{k_n}{d_n} 2\pi i} \right), \text{ for } 0 \leq k_i \leq d_i - 1 \text{ and } i = 1, \dots, n \right\}.$$

2.1 Choosing the homotopy

The deformation between the original system and the start system is a *homotopy*, for instance the convex combination of F and G

$$H(x, t) = (1 - t)F(x) + tG(x), \tag{1}$$

where $x := (x_1, \dots, x_n)$ and $t \in [0, 1]$. This is such that the roots of $H(x, 0) = G(x)$ are known, and the roots of $H(x, 1) = F(x)$ are the solutions of the original system (the reason why we place the start system at $t = 0$ and the original system at $t = 1$ is that we need higher numerical precision for the solutions of the original system, and there are more floating point numbers near to $t = 0$; see [1], p. 33). Therefore, we can implicitly define a curve $z(t)$ in \mathbb{C}^n by the equation

$$H(z(t), t) = 0, \tag{2}$$

so that in order to approximate the roots of F it is enough to numerically track $z(t)$.

To do so, we derive the expression (2) with respect to t , and get the *Dauidenko Differential Equation*

$$\frac{\partial H}{\partial z} \frac{dz}{dt} + \frac{\partial H}{\partial t} = 0,$$

where $\frac{\partial H}{\partial z}$ is the Jacobian matrix of H with respect to z :

$$\frac{\partial H}{\partial z} = \begin{pmatrix} \frac{\partial H_1}{\partial z_1} & \cdots & \frac{\partial H_1}{\partial z_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial H_n}{\partial z_1} & \cdots & \frac{\partial H_n}{\partial z_n} \end{pmatrix}.$$

This can be rewritten as

$$\dot{z} = -\frac{\partial H^{-1}}{\partial z} \frac{\partial H}{\partial t}. \quad (3)$$

This is a system of n first-order differential equations, which can be solved numerically for $z(t)$ as an initial value problem, which is called *path tracking*.

2.1.1 Gamma trick

While (1) is a fine choice of a homotopy, it's not what it's called a *good homotopy*: in order to ensure that the solution paths $z(t)$ for different roots

- have no singularities, i.e. never cross each other for $t > 0$ (at $t = 0$ F could have singular solutions), and
- don't go to infinity for $t \rightarrow 0$ (as F could have a solution at infinity),

we can employ the *Gamma trick*: this consists in modifying the linear homotopy (1) by substituting the parameter $t \in [0, 1]$ with a complex curve $q(t)$ connecting 0 and 1, such as

$$q(t) = \frac{\gamma t}{\gamma t + (1 - t)},$$

where $\gamma \in (0, 1)$ is a random complex parameter.

This is a "probability one" procedure, i.e. for any particular system we can choose γ outside of a finite amount of rays through the origin to ensure that we get a good homotopy, basically because of the finiteness of the branch locus of the homotopy. After substituting, we have

$$H(x, t) = \frac{(1 - t)}{\gamma t + (1 - t)} F(x) + \frac{\gamma t}{\gamma t + (1 - t)} G(x),$$

and by clearing denominators, we get our final choice of homotopy:

$$H(x, t) = (1 - t)F(x) + \gamma t G(x). \quad (4)$$

2.2 Tracking down the roots

We then need to track down individual roots, following the solution paths from a root z_0 of the start system by solving the initial value problem associated to the Dauidenko differential equation (3) with starting value z_0 and t ranging from 1 to 0.

This will be done numerically, by using a first-order predictor-corrector tracking method, whose typical iteration goes like this:

- **Predictor:** we first apply Euler's method to get an approximation \tilde{z}_i of the next value of the solution path;
- **Corrector:** we then use Newton's method to correct \tilde{z}_i using equation (2), so that it becomes a good approximation z_i of the next value of the solution path.

In the following sections, we go into more detail on each of these steps.

2.2.1 Predictor: Euler's method

Recall that Euler's method consists in approximating the solution of the initial value problem associated to a system of first-order ordinary differential equations

$$\begin{cases} \dot{z} = f(z, t) \\ z(t_0) = z_0 \end{cases}$$

by the sequence of points $(z_i)_{i \in \mathbb{N}}$ defined by the recurrence relation

$$z_{i+1} = z_i + h \cdot f(z_i, t_i),$$

where h is the step size. In the case of the Davidenko equation (3), we have

$$f(z, t) = - \left(\frac{\partial H}{\partial z}(z, t) \right)^{-1} \frac{\partial H}{\partial t}(z, t)$$

and $t_0 = 1$, since we are tracking from 1 to 0. For the same reason, we set

$$t_{i+1} = t_i - h.$$

2.2.2 Corrector: Newton's method

Since we want to solve

$$H(z, t) = 0,$$

we can use Newton's method to improve the approximation \tilde{z}_i obtained by Euler's method. This is done by moving towards the root of the tangent line of H at the current approximation, or in other words through the iteration

$$z_{i+1} = z_i - \left(\frac{\partial H}{\partial z}(z_i, t_{i+1}) \right)^{-1} H(z_i, t_{i+1}),$$

where this time $z_0 = \tilde{z}_i$, with \tilde{z}_i and t_{i+1} obtained from the i -th Euler step.

Usually, only a few steps of Newton's method are needed; we chose a fixed number of 5 iterations. At which point, we use the final value of the Newton iteration as the starting value for the next Euler step.

2.2.3 Adaptive step size

In order to improve the efficiency of the method, we will use an adaptive step size, which is based on the norm of the residual of Newton's iteration. If the desired accuracy is not reached (say, when the norm of $H(z_i, t_i)$ is bigger than 10^{-8}), then we halve the step size; if instead we have 5 "successful" iterations in a row, we double the step size.

3 Testing the method

To test the method’s scalability, we first launched it on a single-threaded machine, then one a multi-threaded one, and finally parallelized it on a Cluster.

The latter was done by using the Julia package *Distributed.jl* to parallelize the tracking of the roots on separate nodes, and the `SlurmClusterManager` package, which allows to run Julia code using the Slurm workload manager.

In order to scale the method to larger systems, we also implemented a random polynomial generator which can be found in [random-poly.jl](#); this was used to evaluate the performance of the parallel implementation, by generating square systems of polynomials with normally distributed coefficients, each polynomial having total degree less or equal to a fixed maximum degree.

The single-threaded machine and multi-threaded tests (which used the `@threads` macro from the *Threads.jl* package on the root tracking for loop in the file [solve.jl](#)) were run in order to visualize the real solutions of small (2x2) systems: here, multi-threaded runs didn’t improve the performance on these smaller systems, as the overhead of multi-threading was too big compared to the actual computation time.

However, when testing a parallel implementation on larger randomly generated systems we observed an improvement in execution times on larger systems compared to the single-node runs, as we show in the [Results](#) section.

The Julia implementation for the tests described above can be found in [Appendix B](#), while the hardware specifications are listed in [Appendix A](#).

4 Possible Improvements

4.1 Homogenized Coordinates

Since our start systems have the maximum number of solutions for its degree, some of them might converge to a point at infinity of our original system. In our current implementation, we waste time by tracking them until reaching the maximum number of iterations.

To better treat such cases, we could view the system inside an affine patch of the projective plane, and using homogenized coordinates detect when a solution is going to infinity. This would involve homogenizing both systems and modifying the path-tracking algorithm for the detection of a point going to infinity.

4.2 Predictor-Corrector


Our (un)specific choice of predictor could be unsuitable for badly-conditioned systems; other software implementations of the homotopy continuation method use more accurate and numerically stable predictors, such as Runge-Kutta methods [\[3\]](#).

5 Appendix A: Results

5.1 Single- and Multi-threaded



Below are the plots of four different 2x2 systems for the single- (laptop) and multi- (desktop) threaded runs, with the real solutions being shown in **red**:

Single-threaded

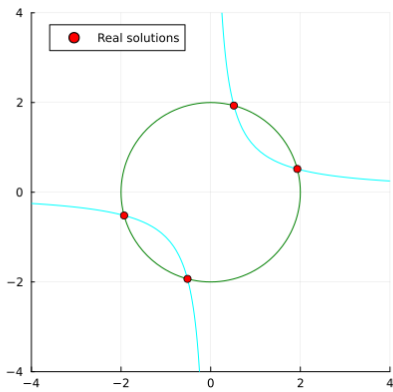



$$\begin{cases} x^3 + 5x^2 - y - 1 \\ 2x^2 - y - 1 \end{cases}$$

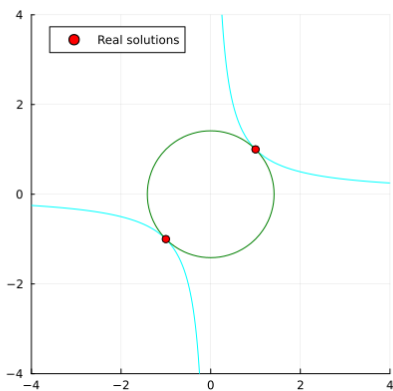
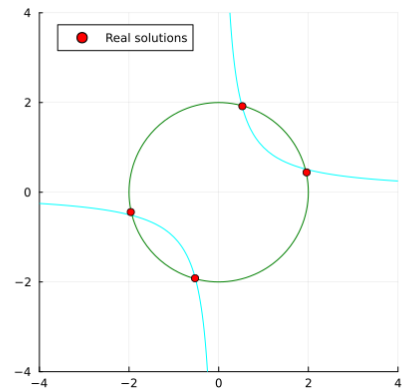
Multithreaded



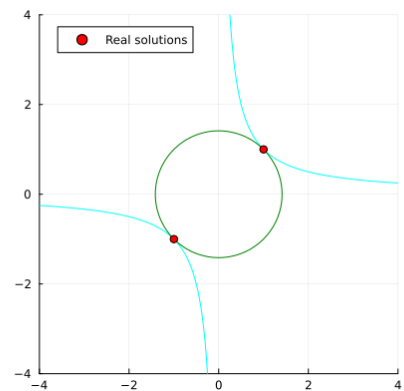
$$\begin{cases} x^2 + 2y \\ y - 3x^3 \end{cases}$$



$$\begin{cases} x^2 + y^2 - 4 \\ xy - 1 \end{cases}$$



$$\begin{cases} x^2 + y^2 - 2 \\ xy - 1 \end{cases}$$



5.2 Parallelization

The following figure compares the execution times of the `solve` function in `solve.jl` on the cluster, on a single node and on 20 nodes (using 1 or 2 threads per node).

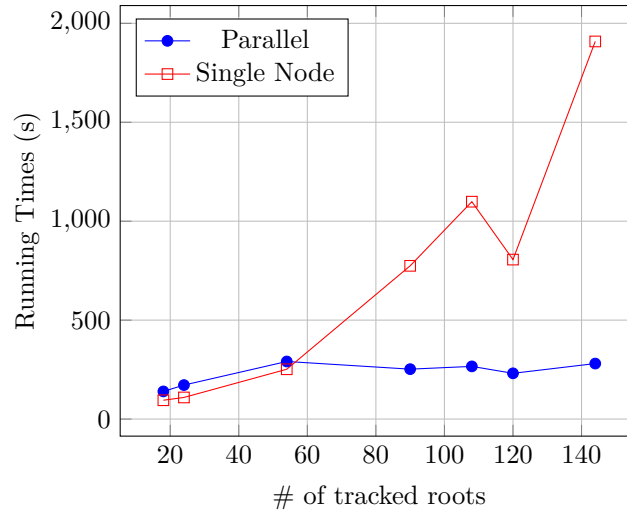


Figure 1: Performance comparison of parallel path tracking on a cluster.

As we can see from the plot, the parallel implementation appears to scale well with the number of tracked roots, and is faster than the single-node implementation for larger systems.

6 Appendix B: Implementation

6.1 Julia code

Listing 1: `solve.jl`

```
1 # External deps
2 using LinearAlgebra
3 using TypedPolynomials
4 using Distributed
5 # using SlurmClusterManager
6 # slurm_manager = SlurmManager()
7 # addprocs(slurm_manager)
8
9 # Local deps
10 include("random-poly.jl")
11 include("plot.jl")
12 using .RandomPoly
13 using .Plot
14 @everywhere begin
15     include("start-system.jl")
16     include("homotopy.jl")
17     include("euler-newton.jl")
18     include("adapt-step.jl")
19 end
20 # Macros defined in an @everywhere block aren't available inside it
21 @everywhere begin
22     using .StartSystem
23     using .Homotopy
24     using .EulerNewton
```



```

25 using .AdaptStep
26 end
27
28 # ciao sto facendo delle modifiche al codice
29
30 @everywhere function compute_root(H, r, maxsteps=200)
31     t = 1.0
32     step_size = 0.001
33     x0 = r
34     m = 0
35     steps = 0
36
37     while t > 0 && steps < maxsteps
38         x0 = en_step(H, x0, t, step_size)
39         (m, step_size) = adapt_step(H, x0, t, step_size, m)
40         t -= step_size
41         steps += 1
42     end
43     return (x0, steps)
44 end
45
46 # Main homotopy continuation loop
47 function solve(F, G, roots)
48     H = homotopy(F, G)
49
50     result = Array{Future}(undef, length(roots))
51     for i in eachindex(roots)
52         result[i] = @spawnat :any compute_root(H, roots[i])
53     end
54
55     sols = Array{ComplexF64,2}(undef, length(roots), length(F))
56     steps = Array{Int64}(undef, length(roots))
57     for i in eachindex(roots)
58         (solution, step_array) = fetch(result[i])
59         sols[i, :] = solution
60         steps[i] = step_array
61     end
62
63     return (sols, steps)
64 end
65
66 # @polyvar x y
67 # C = [x^3 - y + 5x^2 - 10, 2x^2 - y - 10]
68 # Q = [x^2 + 2y, y - 3x^3]
69 # F = [x*y - 1, x^2 + y^2 - 4]
70 # T = [x*y - 1, x^2 + y^2 - 2]
71
72 R = random_system(6, 6)
73 println("System: ", R)
74 (G, roots)=start_system(R)
75 println("Number of roots: ", length(roots))
76
77 # Parallel execution
78 println("PARALLEL")
79 @time begin
80     (sol, steps) = solve(R, G, roots)
81 end
82 println("Number of steps: ", steps)
83 # converting sR to array of arrays instead of a matrix
84 sol = [sol[i, :] for i in 1:length(sol[:, 1])]
85 sol = filter(u -> imag(u[1]) < 0.1 && imag(u[2]) < 0.1, sol)
86 sol = map(u -> real.(u), sol)
87 vars = variables(R)
88 println("Solutions: ", sol)
89 println("Norms (lower = better): ", [norm([f(vars => s) for f in R]) for s in
    sol])
90
91 # # Single execution
92 # println("SINGLE")
93 # wait(rmprocs(workers()))
94 # @time begin
95 #     (sol, steps) = solve(R, G, roots)
96 # end

```

```

97 # println("Number of steps: ", steps)
98 # # converting sR to array of arrays instead of a matrix
99 # sol = [sol[i, :] for i in 1:length(sol[:, 1])]
100 # sol = filter(u -> imag(u[1]) < 0.1 && imag(u[2]) < 0.1, sol)
101 # sol = map(u -> real(u), sol)
102 # vars = variables(R)
103 # println("Solutions: ", sol)
104 # println("Norms (lower = better): ", [norm([f(vars => s) for f in R]) for s
    in sol])
105
106 # # See https://github.com/kleinhenz/SlurmClusterManager.jl/issues/11
107 # finalize(slurm_manager)
108
109 # Plotting the system and the real solutions
110 # ENV["GKSwstype"] = "nul"
111 # plot_real(sC, C, 6, 12, "1")
112 # plot_real(sQ, Q, 2, 2, "2")
113 # plot_real(sF, F, 4, 4, "3")
114 # plot_real(sT, T, 4, 4, "4")
115 # plot_real(sol, R, 5, 5, "random")

```

Listing 2: start-system.jl

```

1 module StartSystem
2   using TypedPolynomials
3
4   export start_system
5
6   # Define start system based on total degree
7   function start_system(F)
8     degrees = [maxdegree(p) for p in F]
9     G = [x_i^d - 1 for (d, x_i) in zip(degrees, variables(F))]
10    r = [[exp(2im*pi/d)^k for k=0:d-1] for d in degrees]
11    roots = vec([collect(root) for root in collect(Iterators.product(r...))])
12    return (G, roots)
13  end
14 end

```

Listing 3: homotopy.jl

```

1 module Homotopy
2   export homotopy
3
4   # Define a straight-line homotopy between the two systems
5   function homotopy(F, G)
6     γ = cis(2π * rand())
7     function H(t)
8       return [(1 - t) * f + γ * t * g for (f, g) in zip(F, G)]
9     end
10    return H
11  end
12 end

```

Listing 4: homogenize.jl

```

1 module Homogenize
2   using TypedPolynomials
3
4   export homogenize, homogenized_start_system
5
6   function homogenize(F)
7     @polyvar h

```

```

8     return [sum([h^(maxdegree(p)-maxdegree(t))*t for t in p.terms]) for p in F
9 ]
10 end
11 function homogenized_start_system(F)
12     degrees = [maxdegree(p) for p in F]
13     @polyvar h
14     G = [x_i^d - h^d for (d, x_i) in zip(degrees, variables(F))]
15     r = [[exp(2im*pi/d)^k for k=0:d-1] for d in degrees]
16     roots = vec([vcat(collect(root), 1) for root in collect(Iterators.product(r
17     ...))])
18     return (G, roots)
19 end

```

Listing 5: euler-newton.jl

```

1 module EulerNewton
2 using LinearAlgebra
3 using TypedPolynomials
4
5 export en_step
6
7 # Euler-Newton predictor-corrector
8 function en_step(H, x, t, step_size)
9
10     # Predictor step
11     vars = variables(H(1))
12     # Jacobian of H evaluated at (x,t)
13     JH = [jh(vars=>x) for jh in differentiate(H(t), vars)]
14     #  $\partial H/\partial t = \nabla G - F = H(1) - H(0)$  for our homotopy; it doesn't depend on t
15      $\delta H_{\delta t} = [dh(vars=>x) for dh in H(1) - H(0)]$ 
16      $\Delta x = JH \setminus -\delta H_{\delta t}$ 
17     xh = x +  $\Delta x$  * step_size
18
19     # Corrector step
20     JHh = differentiate(H(t-step_size), vars)
21     for _ in 1:5
22         JH = [jh(vars=>xh) for jh in JHh]
23          $\Delta x = JH \setminus -[h(vars=>xh) for h in H(t-step_size)]$ 
24         xh = xh +  $\Delta x$ 
25     end
26
27     return xh
28 end
29 end

```

Listing 6: adapt-step.jl

```

1 module AdaptStep
2 using LinearAlgebra
3 using TypedPolynomials
4
5 export adapt_step
6
7 # Adaptive step size
8 function adapt_step(H, x, t, step, m)
9      $\Delta = \text{norm}([h(\text{variables}(H(t))=>x) for h in H(t-step)])$ 
10     if  $\Delta > 1e-10$ 
11         step = 0.5 * step
12         m = 0
13     else
14         m+=1
15         if (m == 4)
16             step = 2 * step
17             m = 0

```

```

18     end
19     end
20
21     return (m, step)
22 end
23 end

```

Listing 7: random-poly.jl

```

1 module RandomPoly
2     export random_system
3
4     using TypedPolynomials
5     using Random
6     using Distributions
7
8     # Random polynomial of degree n in m variables
9     function random_poly(n, m)
10        x = [TypedPolynomials.Variable{Symbol("x[$i]")}() for i in 1:m]
11
12        monomial_powers=collect(Iterators.product([0:n for _ in 1:m]...))
13        monomials = [prod(x.^i) for i in monomial_powers if sum(i) <= n && sum(i) != 0]
14
15        return sum(map(m -> rand(Normal()) * m, monomials))
16    end
17
18    # Generate a system of m random polynomials in m variables
19    # of degree d_i randomly chosen between 1 and max_degree
20    function random_system(m, max_degree)
21        d = rand(1:max_degree, m)
22        random_polys = [random_poly(d[i], m) for i in 1:m]
23
24        return random_polys
25    end
26 end

```

Listing 8: plot.jl

```

1 module Plot
2     using Plots, TypedPolynomials
3
4     export plot_real
5
6     function plot_real(solutions, F, h, v, name)
7         plot(xlim = (-h, h), ylim = (-v, v), aspect_ratio = :equal)
8         contour!(-h:0.1:h, -v:0.1:v, (x,y)->F[1](variables(F)=>[x,y]), levels=[0],
9             cbar=false, color=:cyan)
10        contour!(-h:0.1:h, -v:0.1:v, (x,y)->F[2](variables(F)=>[x,y]), levels=[0],
11            cbar=false, color=:green)
12        scatter!([real(sol[1]) for sol in solutions], [real(sol[2]) for sol in
13            solutions], color = "red", label = "Real solutions")
14
15        png(joinpath("./plots", "solutions" * name))
16    end
17 end

```

6.2 Hardware

For the single-threaded runs, the code was executed on a laptop with an Intel Core i7-3520M CPU @ 3.60GHz and 6 GB of RAM.

The multithreaded runs were tested on a desktop with an AMD FX-8350 CPU @ 4.00GHz with 4 cores and 8 threads, and 12 GB of RAM.

Finally, the parallel computations were run on a cluster with 20 nodes, each having a CPU @ 1.008GHz with 4 Performance cores, 2 efficiency cores and 4 GB of RAM.

References

- [1] Bates, Daniel J. *Numerically solving polynomial systems with Bertini*. SIAM, Society for Industrial Applied Mathematics, 2013.
- [2] <https://docs.julialang.org/en/v1/stdlib/Distributed>
- [3] Breiding, P., Timme, S. (2018). HomotopyContinuation.jl: A Package for Homotopy Continuation in Julia. In: Davenport, J., Kauers, M., Labahn, G., Urban, J. (eds) *Mathematical Software – ICMS 2018*. Lecture Notes in Computer Science, vol 10931. Springer, Cham.