

Solving Stiff Equations

Chris Rackauckas

November 1, 2019

This tutorial is for getting into the extra features for solving stiff ordinary differential equations in an efficient manner. Solving stiff ordinary differential equations requires specializing the linear solver on properties of the Jacobian in order to cut down on the $O(n^3)$ linear solve and the $O(n^2)$ back-solves. Note that these same functions and controls also extend to stiff SDEs, DDEs, DAEs, etc.

0.1 Code Optimization for Differential Equations

0.1.1 Writing Efficient Code

For a detailed tutorial on how to optimize one's `DifferentialEquations.jl` code, please see the [Optimizing DiffEq Code tutorial](#).

0.1.2 Choosing a Good Solver

Choosing a good solver is required for getting top notch speed. General recommendations can be found on the solver page (for example, the [ODE Solver Recommendations](#)). The current recommendations can be simplified to a Rosenbrock method (`Rosenbrock23` or `Rodas5`) for smaller (<50 ODEs) problems, ESDIRK methods for slightly larger (`TRBDF2` or `KenCarp4` for <2000 ODEs), and Sundials `CVODE_BDF` for even larger problems. `lsoda` from [LSODA.jl](#) is generally worth a try.

More details on the solver to choose can be found by benchmarking. See the [DiffEqBenchmarks](#) to compare many solvers on many problems.

0.1.3 Check Out the Speed FAQ

See [this FAQ](#) for information on common pitfalls and how to improve performance.

0.1.4 Setting Up Your Julia Installation for Speed

Julia uses an underlying BLAS implementation for its matrix multiplications and factorizations. This library is automatically multithreaded and accelerates the internal linear algebra of `DifferentialEquations.jl`. However, for optimality, you should make sure that the number of BLAS threads that you are using matches the number of physical cores and not the number of logical cores. See [this issue for more details](#).

To check the number of BLAS threads, use:

```
ccall( (:openblas_get_num_threads64_, Base.libblas_name), Cint, ())
```

4

If I want to set this directly to 4 threads, I would use:

```
using LinearAlgebra
LinearAlgebra.BLAS.set_num_threads(4)
```

Additionally, in some cases Intel's MKL might be a faster BLAS than the standard BLAS that ships with Julia (OpenBLAS). To switch your BLAS implementation, you can use [MKL.jl](#) which will accelerate the linear algebra routines. Please see the package for the limitations.

0.1.5 Use Accelerator Hardware

When possible, use GPUs. If your ODE system is small and you need to solve it with very many different parameters, see the [ensembles interface](#) and [DiffEqGPU.jl](#). If your problem is large, consider using a [CuArray](#) for the state to allow for GPU-parallelism of the internal linear algebra.

0.2 Speeding Up Jacobian Calculations

When one is using an implicit or semi-implicit differential equation solver, the Jacobian must be built at many iterations and this can be one of the most expensive steps. There are two pieces that must be optimized in order to reach maximal efficiency when solving stiff equations: the sparsity pattern and the construction of the Jacobian. The construction is filling the matrix J with values, while the sparsity pattern is what J to use.

The sparsity pattern is given by a prototype matrix, the `jac_prototype`, which will be copied to be used as J . The default is for J to be a `Matrix`, i.e. a dense matrix. However, if you know the sparsity of your problem, then you can pass a different matrix type. For example, a `SparseMatrixCSC` will give a sparse matrix. Additionally, structured matrix types like `Tridiagonal`, `BandedMatrix` (from [BandedMatrices.jl](#)), `BlockBandedMatrix` (from [BlockBandedMatrices.jl](#)), and more can be given. `DifferentialEquations.jl` will internally use this matrix type, making the factorizations faster by utilizing the specialized forms.

For the construction, there are 3 ways to fill J :

- The default, which uses normal finite/automatic differentiation
- A function `jac(J,u,p,t)` which directly computes the values of J
- A `colorvec` which defines a sparse differentiation scheme.

We will now showcase how to make use of this functionality with growing complexity.

0.2.1 Declaring Jacobian Functions

Let's solve the Rosenbrock equations:

$$dy_1 = -0.04y_1 + 10^4 y_2 y_3 \quad (1)$$

$$dy_2 = 0.04y_1 - 10^4 y_2 y_3 - 3 * 10^7 y_2^2 \quad (2)$$

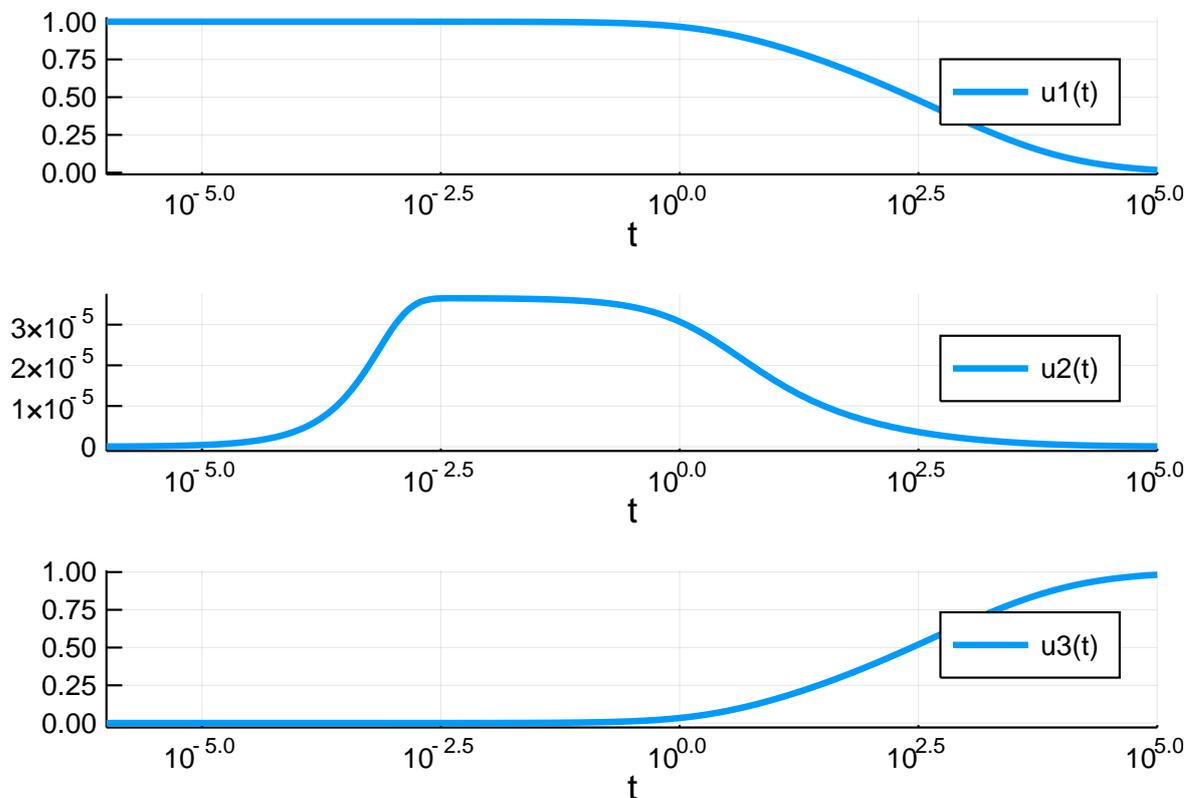
$$dy_3 = 3 * 10^7 y_3^2 \quad (3)$$

$$(4)$$

In order to reduce the Jacobian construction cost, one can describe a Jacobian function by using the `jac` argument for the `ODEFunction`. First, let's do a standard `ODEProblem`:

```
using DifferentialEquations
function rober(du,u,p,t)
    y_1,y_2,y_3 = u
    k_1,k_2,k_3 = p
    du[1] = -k_1*y_1+k_3*y_2*y_3
    du[2] = k_1*y_1-k_2*y_2^2-k_3*y_2*y_3
    du[3] = k_2*y_2^2
    nothing
end
prob = ODEProblem(rober, [1.0,0.0,0.0], (0.0,1e5), (0.04,3e7,1e4))
sol = solve(prob,Rosenbrock23())
```

```
using Plots
plot(sol, xscale=:log10, tspan=(1e-6, 1e5), layout=(3,1))
```



```
using BenchmarkTools
@btime solve(prob)
```

```

316.900  $\mu$ s (3062 allocations: 161.81 KiB)
retcode: Success
Interpolation: Automatic order switching interpolation
t: 115-element Array{Float64,1}:
  0.0
  0.0014148468219250373
  0.0020449182545311173
  0.0031082402716566307
  0.004077787050059496
  0.005515332443361059
  0.007190040962774541
  0.009125372578778032
  0.011053912492732977
  0.012779077276958607
  ⋮
47335.56357690261
52732.01292853374
58693.72991412389
65278.000210850696
72548.20206513454
80574.5643369749
89435.05301092885
99216.41264599326
100000.0
u: 115-element Array{Array{Float64,1},1}:
 [1.0, 0.0, 0.0]
 [0.9999434113193613, 3.283958829839966e-5, 2.3749092340286502e-5]
 [0.9999182177783585, 3.55426801363446e-5, 4.6239541505020656e-5]
 [0.999875715036629, 3.6302469334849744e-5, 8.798249403609506e-5]
 [0.9998369766077329, 3.646280308115459e-5, 0.00012656058918590176]
 [0.9997795672444667, 3.646643085642237e-5, 0.0001839663246768369]
 [0.9997127287139348, 3.6447279992896e-5, 0.00025082400607228316]
 [0.9996355450022019, 3.6366816179962866e-5, 0.00032808818161818775]
 [0.9995586925734838, 3.6018927453312764e-5, 0.00040528849906290045]
 [0.9994899965196854, 3.468694637786026e-5, 0.000475316533936808]
  ⋮
 [0.03394368168613229, 1.404798439362035e-7, 0.9660561778340258]
 [0.031028975539652698, 1.280360743781007e-7, 0.9689708964242754]
 [0.02835436357223889, 1.1668209524677941e-7, 0.9716455197456683]
 [0.025901326001934923, 1.0632276689411095e-7, 0.9740985676753005]
 [0.023652545345805354, 9.687112514942483e-8, 0.9763473577830714]
 [0.021591862129552664, 8.824767963573306e-8, 0.9784080496227692]
 [0.019704225538717677, 8.037977048382674e-8, 0.9802956940815135]
 [0.017975641463053707, 7.320098240041474e-8, 0.9820242853359655]
 [0.017850566233695766, 7.268384360678819e-8, 0.9821493610824623]

```

Now we want to add the Jacobian. First we have to derive the Jacobian $\frac{df_i}{du_j}$ which is $J[i, j]$. From this we get:

```

function rober_jac(J,u,p,t)
  y_1,y_2,y_3 = u
  k_1,k_2,k_3 = p
  J[1,1] = k_1 * -1
  J[2,1] = k_1
  J[3,1] = 0
  J[1,2] = y_3 * k_3
  J[2,2] = y_2 * k_2 * -2 + y_3 * k_3 * -1
  J[3,2] = y_2 * 2 * k_2

```

```

J[1,3] = k_3 * y_2
J[2,3] = k_3 * y_2 * -1
J[3,3] = 0
nothing
end
f = ODEFunction(rober, jac=rober_jac)
prob_jac = ODEProblem(f, [1.0, 0.0, 0.0], (0.0, 1e5), (0.04, 3e7, 1e4))

```

```
@btime solve(prob_jac)
```

239.200 μ s (2599 allocations: 153.11 KiB)

retcode: Success

Interpolation: Automatic order switching interpolation

t: 115-element Array{Float64,1}:

```

0.0
0.00141448468219250373
0.0020449182545311173
0.0031082402716566307
0.004077787050059496
0.005515332443361059
0.007190040962774541
0.009125372578778032
0.011053912492732977
0.012779077276958607

```

⋮

```

45964.060340548356
51219.40381376205
57025.01899700374
63436.021374561584
70513.1073617524
78323.14229130604
86939.82338876331
96444.41085674686
100000.0

```

100000.0

u: 115-element Array{Array{Float64,1},1}:

```

[1.0, 0.0, 0.0]
[0.9999434113193613, 3.283958829839966e-5, 2.3749092340286502e-5]
[0.9999182177783585, 3.55426801363446e-5, 4.6239541505020656e-5]
[0.999875715036629, 3.6302469334849744e-5, 8.798249403609506e-5]
[0.9998369766077329, 3.646280308115459e-5, 0.00012656058918590176]
[0.9997795672444667, 3.646643085642237e-5, 0.0001839663246768369]
[0.9997127287139348, 3.6447279992896e-5, 0.00025082400607228316]
[0.9996355450022019, 3.6366816179962866e-5, 0.00032808818161818775]
[0.9995586925734838, 3.6018927453312764e-5, 0.00040528849906290045]
[0.9994899965196854, 3.468694637786026e-5, 0.000475316533936808]

```

⋮

```

[0.03478048133177493, 1.4406682005231008e-7, 0.9652193746014031]
[0.03179591062189176, 1.313038656880417e-7, 0.9682039580742408]
[0.029057356622057315, 1.1966100432939363e-7, 0.9709425237169371]
[0.02654597011713668, 1.0904070990251299e-7, 0.9734539208421517]
[0.024244118287194777, 9.935385522693504e-8, 0.9757557823589477]
[0.022135344621501105, 9.05190025093182e-8, 0.9778645648594945]
[0.02020432071854, 8.246174295748071e-8, 0.9797955968197154]
[0.018436796681356796, 7.511410189106845e-8, 0.9815631282045397]
[0.01785426048218692, 7.269900678199638e-8, 0.9821456668188047]

```

0.2.2 Automatic Derivation of Jacobian Functions

But that was hard! If you want to take the symbolic Jacobian of numerical code, we can make use of `ModelingToolkit.jl` to symbolicify the numerical code and do the symbolic calculation and return the Julia code for this.

```
using ModelingToolkit
de = modelingtoolkitize(prob)
ModelingToolkit.generate_jacobian(de...) [2] # Second is in-place

:((##MTIIPVar#918, u, p, t)->begin
    #= C:\Users\accou\.julia\packages\ModelingToolkit\czHtj\src\utils
.jl:65 =#
    #= C:\Users\accou\.julia\packages\ModelingToolkit\czHtj\src\utils
.jl:66 =#
    let (x_1, x_2, x_3, α_1, α_2, α_3) = (u[1], u[2], u[3], p[1], p[2], p[3]
])
        ##MTIIPVar#918[1] = α_1 * -1
        ##MTIIPVar#918[2] = α_1
        ##MTIIPVar#918[3] = 0
        ##MTIIPVar#918[4] = x_3 * α_3
        ##MTIIPVar#918[5] = x_2 * α_2 * -2 + x_3 * α_3 * -1
        ##MTIIPVar#918[6] = x_2 * 2 * α_2
        ##MTIIPVar#918[7] = α_3 * x_2
        ##MTIIPVar#918[8] = α_3 * x_2 * -1
        ##MTIIPVar#918[9] = 0
    end
    #= C:\Users\accou\.julia\packages\ModelingToolkit\czHtj\src\utils
.jl:67 =#
    nothing
end)
```

which outputs:

```
:((##MTIIPVar#376, u, p, t)->begin
    #= C:\Users\accou\.julia\packages\ModelingToolkit\czHtj\src\utils.jl:65 =#
    #= C:\Users\accou\.julia\packages\ModelingToolkit\czHtj\src\utils.jl:66 =#
    let (x_1, x_2, x_3, α_1, α_2, α_3) = (u[1], u[2], u[3], p[1], p[2], p[3])
        ##MTIIPVar#376[1] = α_1 * -1
        ##MTIIPVar#376[2] = α_1
        ##MTIIPVar#376[3] = 0
        ##MTIIPVar#376[4] = x_3 * α_3
        ##MTIIPVar#376[5] = x_2 * α_2 * -2 + x_3 * α_3 * -1
        ##MTIIPVar#376[6] = x_2 * 2 * α_2
        ##MTIIPVar#376[7] = α_3 * x_2
        ##MTIIPVar#376[8] = α_3 * x_2 * -1
        ##MTIIPVar#376[9] = 0
    end
    #= C:\Users\accou\.julia\packages\ModelingToolkit\czHtj\src\utils.jl:67 =#
    nothing
end)
```

Now let's use that to give the analytical solution Jacobian:

```
jac = eval(ModelingToolkit.generate_jacobian(de...) [2])
f = ODEFunction(rober, jac=jac)
prob_jac = ODEProblem(f, [1.0, 0.0, 0.0], (0.0, 1e5), (0.04, 3e7, 1e4))

ODEProblem with uType Array{Float64,1} and tType Float64. In-place: true
timespan: (0.0, 100000.0)
u0: [1.0, 0.0, 0.0]
```

0.2.3 Declaring a Sparse Jacobian

Jacobian sparsity is declared by the `jac_prototype` argument in the `ODEFunction`. Note that you should only do this if the sparsity is high, for example, 0.1% of the matrix is non-zeros, otherwise the overhead of sparse matrices can be higher than the gains from sparse differentiation!

But as a demonstration, let's build a sparse matrix for the Rober problem. We can do this by gathering the I and J pairs for the non-zero components, like:

```
I = [1,2,1,2,3,1,2]
J = [1,1,2,2,2,3,3]
using SparseArrays
jac_prototype = sparse(I,J,1.0)
```

3×3 SparseArrays.SparseMatrixCSC{Float64,Int64} with 7 stored entries:

```
[1, 1] = 1.0
[2, 1] = 1.0
[1, 2] = 1.0
[2, 2] = 1.0
[3, 2] = 1.0
[1, 3] = 1.0
[2, 3] = 1.0
```

Now this is the sparse matrix prototype that we want to use in our solver, which we then pass like:

```
f = ODEFunction(rober, jac=jac, jac_prototype=jac_prototype)
prob_jac = ODEProblem(f, [1.0,0.0,0.0], (0.0,1e5), (0.04,3e7,1e4))
```

```
ODEProblem with uType Array{Float64,1} and tType Float64. In-place: true
timespan: (0.0, 100000.0)
u0: [1.0, 0.0, 0.0]
```

0.2.4 Automatic Sparsity Detection

One of the useful companion tools for `DifferentialEquations.jl` is `SparsityDetection.jl`. This allows for automatic declaration of Jacobian sparsity types. To see this in action, let's look at the 2-dimensional Brusselator equation:

```
const N = 32
const xyd_brusselator = range(0,stop=1,length=N)
brusselator_f(x, y, t) = (((x-0.3)^2 + (y-0.6)^2) <= 0.1^2) * (t >= 1.1) * 5.
limit(a, N) = a == N+1 ? 1 : a == 0 ? N : a
function brusselator_2d_loop(du, u, p, t)
    A, B, alpha, dx = p
    alpha = alpha/dx^2
    @inbounds for I in CartesianIndices((N, N))
        i, j = Tuple(I)
        x, y = xyd_brusselator[I[1]], xyd_brusselator[I[2]]
        ip1, im1, jp1, jm1 = limit(i+1, N), limit(i-1, N), limit(j+1, N), limit(j-1, N)
        du[i,j,1] = alpha*(u[im1,j,1] + u[ip1,j,1] + u[i,jp1,1] + u[i,jm1,1] - 4u[i,j,1]) +
                    B + u[i,j,1]^2*u[i,j,2] - (A + 1)*u[i,j,1] + brusselator_f(x, y, t)
        du[i,j,2] = alpha*(u[im1,j,2] + u[ip1,j,2] + u[i,jp1,2] + u[i,jm1,2] - 4u[i,j,2]) +
                    A*u[i,j,1] - u[i,j,1]^2*u[i,j,2]
    end
end
p = (3.4, 1., 10., step(xyd_brusselator))
```

```
(3.4, 1.0, 10.0, 0.03225806451612903)
```

Given this setup, we can give an example input and output and call `sparsity!` on our function with the example arguments and it will kick out a sparse matrix with our pattern, that we can turn into our `jac_prototype`.

```
using SparsityDetection, SparseArrays
input = rand(32,32,2)
output = similar(input)
sparsity_pattern = sparsity!(brusselator_2d_loop,output,input,p,0.0)
```

```
Explored path: SparsityDetection.Path{Bool[], 1}
```

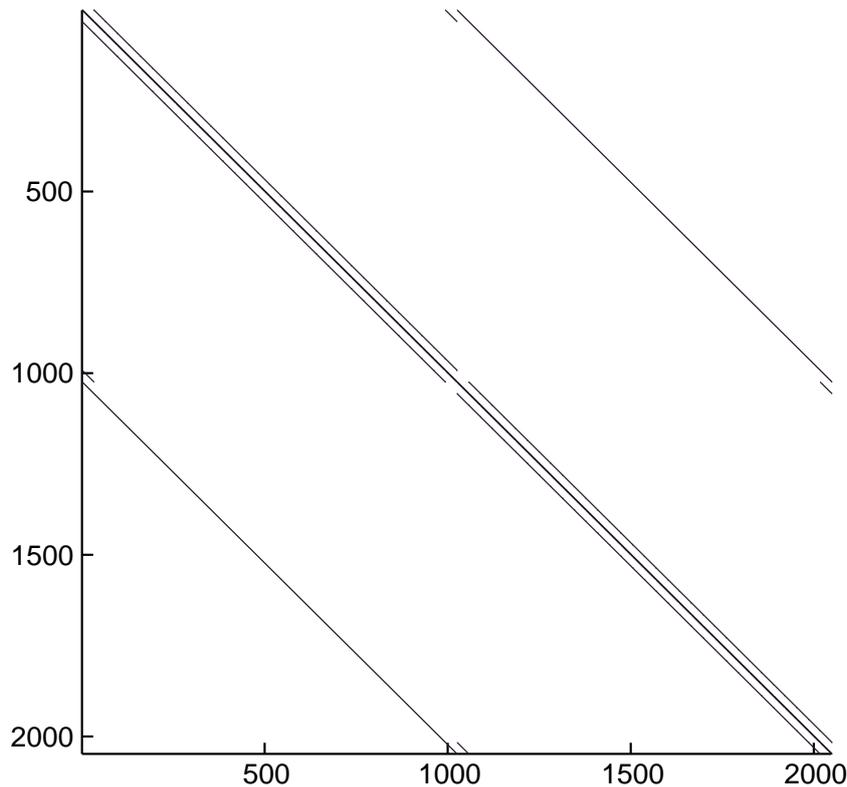
```
jac_sparsity = Float64.(sparse(sparsity_pattern))
```

```
2048×2048 SparseArrays.SparseMatrixCSC{Float64,Int64} with 12288 stored entries:
```

```
[1 , 1] = 1.0
[2 , 1] = 1.0
[32 , 1] = 1.0
[33 , 1] = 1.0
[993 , 1] = 1.0
[1025, 1] = 1.0
[1 , 2] = 1.0
[2 , 2] = 1.0
[3 , 2] = 1.0
⋮
[2015, 2047] = 1.0
[2046, 2047] = 1.0
[2047, 2047] = 1.0
[2048, 2047] = 1.0
[1024, 2048] = 1.0
[1056, 2048] = 1.0
[2016, 2048] = 1.0
[2017, 2048] = 1.0
[2047, 2048] = 1.0
[2048, 2048] = 1.0
```

Let's double check what our sparsity pattern looks like:

```
using Plots
spy(jac_sparsity,markersize=1,colorbar=false,color=:deep)
```



That's neat, and would be tedious to build by hand! Now we just pass it to the `ODEFunction` like as before:

```
f = ODEFunction(brusselator_2d_loop; jac_prototype=jac_sparsity)
```

```
(::DiffEqBase.ODEFunction{true,typeof(Main.WeaveSandBox1.brusselator_2d_loop),LinearAlgebra.UniformScaling{Bool},Nothing,Nothing,Nothing,SparseArrays.SparseMatrixCSC{Float64,Int64},Nothing,Nothing,Nothing,Nothing,Nothing}) (generic function with 7 methods)
```

Build the `ODEProblem`:

```
function init_brusselator_2d(xyd)
    N = length(xyd)
    u = zeros(N, N, 2)
    for I in CartesianIndices((N, N))
        x = xyd[I[1]]
        y = xyd[I[2]]
        u[I,1] = 22*(y*(1-y))^(3/2)
        u[I,2] = 27*(x*(1-x))^(3/2)
    end
    u
end
u0 = init_brusselator_2d(xyd_brusselator)
prob_ode_brusselator_2d = ODEProblem(brusselator_2d_loop,
                                     u0, (0., 11.5), p)

prob_ode_brusselator_2d_sparse = ODEProblem(f,
                                             u0, (0., 11.5), p)
```

```
ODEProblem with uType Array{Float64,3} and tType Float64. In-place: true
timespan: (0.0, 11.5)
```

```
u0: [0.0 0.12134432813715876 ... 0.1213443281371586 0.0; 0.0 0.12134432813715
876 ... 0.1213443281371586 0.0; ... ; 0.0 0.12134432813715876 ... 0.1213443281371
```

```
586 0.0; 0.0 0.12134432813715876 ... 0.1213443281371586 0.0]
```

```
[0.0 0.0 ... 0.0 0.0; 0.14892258453196755 0.14892258453196755 ... 0.14892258453196755 0.14892258453196755; ... ; 0.14892258453196738 0.14892258453196738 ... 0.14892258453196738 0.14892258453196738; 0.0 0.0 ... 0.0 0.0]
```

Now let's see how the version with sparsity compares to the version without:

```
@btime solve(prob_ode_brusselator_2d,save_everystep=false)
```

```
26.606 s (7315 allocations: 70.12 MiB)
```

```
@btime solve(prob_ode_brusselator_2d_sparse,save_everystep=false)
```

```
14.089 s (367199 allocations: 896.99 MiB)
```

```
retcode: Success
```

```
Interpolation: 1st order linear
```

```
t: 2-element Array{Float64,1}:
```

```
0.0
```

```
11.5
```

```
u: 2-element Array{Array{Float64,3},1}:
```

```
[0.0 0.12134432813715876 ... 0.1213443281371586 0.0; 0.0 0.12134432813715876
```

```
... 0.1213443281371586 0.0; ... ; 0.0 0.12134432813715876 ... 0.1213443281371586
```

```
0.0; 0.0 0.12134432813715876 ... 0.1213443281371586 0.0]
```

```
[0.0 0.0 ... 0.0 0.0; 0.14892258453196755 0.14892258453196755 ... 0.14892258453196755 0.14892258453196755; ... ; 0.14892258453196738 0.14892258453196738 ... 0.14892258453196738 0.14892258453196738; 0.0 0.0 ... 0.0 0.0]
```

```
[3.2183315970074036 3.2183043434767553 ... 3.2184226343677738 3.2183712473417185; 3.2183804713733872 3.2183499447177057 ... 3.2184831183646856 3.218425047282479; ... ; 3.218246108233481 3.2182241729222354 ... 3.2183185170391946 3.2182778079052787; 3.2182863194790094 3.218261945024488 ... 3.218367227674788 3.218321653767132]
```

```
[2.364108254063361 2.364109732940303 ... 2.364103502720394 2.3641061660225517; 2.364105345047017 2.3641069231419443 ... 2.3641002347797833 2.3641031002634882; ... ; 2.364113451334332 2.3641147252834216 ... 2.364109297958111 2.3641116159339757; 2.3641109923384915 2.364112358364487 ... 2.3641065653101885 2.3641090439583214]
```

0.2.5 Declaring Color Vectors for Fast Construction

If you cannot directly define a Jacobian function, you can use the `colorvec` to speed up the Jacobian construction. What the `colorvec` does is allows for calculating multiple columns of a Jacobian simultaneously by using the sparsity pattern. An explanation of matrix coloring can be found in the [MIT 18.337 Lecture Notes](#).

To perform general matrix coloring, we can use [SparseDiffTools.jl](#). For example, for the Brusselator equation:

```
using SparseDiffTools
colorvec = matrix_colors(jac_sparsity)
@show maximum(colorvec)
```

```
maximum(colorvec) = 12
```

```
12
```

This means that we can now calculate the Jacobian in 12 function calls. This is a nice reduction from 2048 using only automated tooling! To now make use of this inside of the ODE solver, you simply need to declare the colorvec:

```
f = ODEFunction(brusselator_2d_loop;jac_prototype=jac_sparsity,
               colorvec=colorvec)
prob_ode_brusselator_2d_sparse = ODEProblem(f,
      init_brusselator_2d(xyd_brusselator),
      (0.,11.5),p)
@btime solve(prob_ode_brusselator_2d_sparse,save_everystep=false)

2.391 s (19037 allocations: 881.07 MiB)
retcode: Success
Interpolation: 1st order linear
t: 2-element Array{Float64,1}:
 0.0
11.5
u: 2-element Array{Array{Float64,3},1}:
 [0.0 0.12134432813715876 ... 0.1213443281371586 0.0; 0.0 0.12134432813715876
 ... 0.1213443281371586 0.0; ... ; 0.0 0.12134432813715876 ... 0.1213443281371586
 0.0; 0.0 0.12134432813715876 ... 0.1213443281371586 0.0]

 [0.0 0.0 ... 0.0 0.0; 0.14892258453196755 0.14892258453196755 ... 0.14892258453
 196755 0.14892258453196755; ... ; 0.14892258453196738 0.14892258453196738 ... 0
 .14892258453196738 0.14892258453196738; 0.0 0.0 ... 0.0 0.0]

 [3.2183373918177796 3.2183101409241526 ... 3.2184284167956267 3.218377034566
 1604; 3.2183862623036537 3.218355740108313 ... 3.218488896905827 3.2184308308
 09056; ... ; 3.218251904608678 3.2182299624517134 ... 3.2183243097118095 3.2182
 835995190024; 3.2182921103674285 3.218267738694001 ... 3.2183730157748163 3.2
 183274412034346]

 [2.3641011711912463 2.364102627665652 ... 2.364096424152248 2.364099082779794
 5; 2.3640982676790627 2.3640998304296703 ... 2.3640931617281944 2.36409602465
 74303; ... ; 2.364106344376436 2.3641076180295504 ... 2.364102206048339 2.36410
 45205022344; 2.364103899515714 2.3641052552245445 ... 2.3640994754056486 2.36
 41019485955153]
```

Notice the massive speed enhancement!

0.3 Defining Linear Solver Routines and Jacobian-Free Newton-Krylov

A completely different way to optimize the linear solvers for large sparse matrices is to use a Krylov subspace method. This requires choosing a linear solver for changing to a Krylov method. Optionally, one can use a Jacobian-free operator to reduce the memory requirements.

0.3.1 Declaring a Jacobian-Free Newton-Krylov Implementation

To swap the linear solver out, we use the `linsolve` command and choose the GMRES linear solver.

```

@btime
solve(prob_ode_brusselator_2d,TRBDF2(linsolve=LinSolveGMRES()),save_everystep=false)

242.700 s (1266048 allocations: 120.80 MiB)

@btime
solve(prob_ode_brusselator_2d_sparse,TRBDF2(linsolve=LinSolveGMRES()),save_everystep=false)

4.756 s (1327264 allocations: 59.92 MiB)
retcode: Success
Interpolation: 1st order linear
t: 2-element Array{Float64,1}:
 0.0
11.5
u: 2-element Array{Array{Float64,3},1}:
 [0.0 0.12134432813715876 ... 0.1213443281371586 0.0; 0.0 0.12134432813715876
 ... 0.1213443281371586 0.0; ... ; 0.0 0.12134432813715876 ... 0.1213443281371586
 0.0; 0.0 0.12134432813715876 ... 0.1213443281371586 0.0]

 [0.0 0.0 ... 0.0 0.0; 0.14892258453196755 0.14892258453196755 ... 0.14892258453
 196755 0.14892258453196755; ... ; 0.14892258453196738 0.14892258453196738 ... 0
 .14892258453196738 0.14892258453196738; 0.0 0.0 ... 0.0 0.0]

 [2.8494040430340677 2.849376568123844 ... 2.849495874352271 2.84944397101885
 77; 2.8494535304517883 2.8494226751421077 ... 2.849557062218097 2.84949828863
 504; ... ; 2.8493164846505232 2.849294110741412 ... 2.8493903195873105 2.849349
 0728548774; 2.849357928360968 2.84933329062441 ... 2.8494396335090886 2.84939
 36648688254]

 [2.8157264541468283 2.8157283534566693 ... 2.8157208829524296 2.8157236606184
 397; 2.8157225956336194 2.815724834275517 ... 2.815716958084277 2.81571990149
 71726; ... ; 2.815734632998308 2.8157368388547357 ... 2.8157282527277308 2.8157
 31663143054; 2.815730494353417 2.815732379564653 ... 2.8157247313047327 2.815
 7277764523414]

```

For more information on linear solver choices, see the [linear solver documentation](#).

On this problem, handling the sparsity correctly seemed to give much more of a speedup than going to a Krylov approach, but that can be dependent on the problem (and whether a good preconditioner is found).

We can also enhance this by using a Jacobian-Free implementation of $f'(x)*v$. To define the Jacobian-Free operator, we can use [DiffEqOperators.jl](#) to generate an operator `JacVecOperator` such that `Jv*v` performs $f'(x)*v$ without building the Jacobian matrix.

```

using DiffEqOperators
Jv = JacVecOperator(brusselator_2d_loop,u0,p,0.0)

DiffEqOperators.JacVecOperator{Float64,typeof(Main.WeaveSandBox1.brusselator_2d_loop),Array{ForwardDiff.Dual{DiffEqOperators.JacVecTag,Float64,1},3},Array{ForwardDiff.Dual{DiffEqOperators.JacVecTag,Float64,1},3},Array{Float64,3},NTuple{4,Float64},Float64,Bool}(Main.WeaveSandBox1.brusselator_2d_loop,ForwardDiff.Dual{DiffEqOperators.JacVecTag,Float64,1}[Dual{DiffEqOperators.JacVecTag}(0.0,0.0) Dual{DiffEqOperators.JacVecTag}(0.12134432813715876,0.12134432813715876) ... Dual{DiffEqOperators.JacVecTag}(0.1213443281371586,0.1213443281371586) Dual{DiffEqOperators.JacVecTag}(0.0,0.0); Dual{DiffEqOperators.JacVecTag}(0.0,0.0) Dual{DiffEqOperators.JacVecTag}(0.12134432813715876,0.12134432813715876) ... Dual{DiffEqOperators.JacVecTag}(0.1213443281371586

```



```
f = ODEFunction(brusselator_2d_loop;jac_prototype=Jv)
prob_ode_brusselator_2d_jacfree = ODEProblem(f,u0,(0.,11.5),p)
@btime
solve(prob_ode_brusselator_2d_jacfree,TRBDF2(linsolve=LinSolveGMRES()),save_everystep=false)
```

```
3.427 s (1875298 allocations: 78.86 MiB)
retcode: Success
Interpolation: 1st order linear
t: 2-element Array{Float64,1}:
 0.0
 11.5
u: 2-element Array{Array{Float64,3},1}:
 [0.0 0.12134432813715876 ... 0.1213443281371586 0.0; 0.0 0.12134432813715876
 ... 0.1213443281371586 0.0; ... ; 0.0 0.12134432813715876 ... 0.1213443281371586
 0.0; 0.0 0.12134432813715876 ... 0.1213443281371586 0.0]

 [0.0 0.0 ... 0.0 0.0; 0.14892258453196755 0.14892258453196755 ... 0.14892258453
 196755 0.14892258453196755; ... ; 0.14892258453196738 0.14892258453196738 ... 0
 .14892258453196738 0.14892258453196738; 0.0 0.0 ... 0.0 0.0]
```

```
[2.7872216645408567 2.787194432792592 ... 2.78731308303355 2.787261467045361
5; 2.787271339364228 2.787240720815802 ... 2.787374377179153 2.78731605405600
74; ... ; 2.787134161549321 2.7871118187949984 ... 2.7872072238860723 2.7871659
77632712; 2.7871755020101205 2.7871508886342986 ... 2.7872566948955084 2.7872
10735234632]
```

```
[2.8988126677437585 2.8988142936416157 ... 2.8988075464551772 2.8988105556623
86; 2.898808902249186 2.8988104514436563 ... 2.898803969323616 2.898806883740
06; ... ; 2.898820028584711 2.898821666296394 ... 2.898814592161897 2.898817604
8750383; 2.8988163685403467 2.8988181996160387 ... 2.8988111330962316 2.89881
40808038274]
```

0.3.2 Adding a Preconditioner

The [linear solver documentation](#) shows how you can add a preconditioner to the GMRES. For example, you can use packages like [AlgebraicMultigrid.jl](#) to add an algebraic multigrid (AMG) or [IncompleteLU.jl](#) for an incomplete LU-factorization (iLU).

```
using AlgebraicMultigrid
pc = aspreconditioner(ruge_stuben(jac_sparsity))
@btime
solve(prob_ode_brusselator_2d_jacfree,TRBDF2(linsolve=LinSolveGMRES(P1=pc)),save_everystep=false)
```

```
2.597 s (233048 allocations: 139.27 MiB)
retcode: Success
Interpolation: 1st order linear
t: 2-element Array{Float64,1}:
 0.0
 11.5
u: 2-element Array{Array{Float64,3},1}:
 [0.0 0.12134432813715876 ... 0.1213443281371586 0.0; 0.0 0.12134432813715876
 ... 0.1213443281371586 0.0; ... ; 0.0 0.12134432813715876 ... 0.1213443281371586
 0.0; 0.0 0.12134432813715876 ... 0.1213443281371586 0.0]

 [0.0 0.0 ... 0.0 0.0; 0.14892258453196755 0.14892258453196755 ... 0.14892258453
 196755 0.14892258453196755; ... ; 0.14892258453196738 0.14892258453196738 ... 0
```

```
.14892258453196738 0.14892258453196738; 0.0 0.0 ... 0.0 0.0]
```

```
[3.5273952159283844e10 -1.4265682748702106e10 ... 9234.374594756042 13421.86  
8437681665; -7.091075675799031e9 6.51451873695435e9 ... 9234.400545337947 134  
21.868410996974; ... ; 13421.868025883945 9234.400562276434 ... 9234.4001922958  
72 13421.86842409367; 13421.868438369747 9234.37496117112 ... 9234.3749749346  
17 13421.868424050659]
```

```
[66730.63093229767 -115820.52698935539 ... 16462.92400611659 16458.1794290617  
3; 8.043448946694581e6 1.307043107719831e7 ... 11331.237739674985 11326.51840  
7046895; ... ; 11326.51842066477 11331.237738901911 ... 11331.237752373656 1132  
6.518406581376; 16458.179429033426 16462.923993307235 ... 16462.923992815315  
16458.179429539887]
```

0.4 Using Structured Matrix Types

If your sparsity pattern follows a specific structure, for example a banded matrix, then you can declare `jac_prototype` to be of that structure and then additional optimizations will come for free. Note that in this case, it is not necessary to provide a `colorvec` since the color vector will be analytically derived from the structure of the matrix.

The matrices which are allowed are those which satisfy the [ArrayInterface.jl](#) interface for automatically-colorable matrices. These include:

- Bidiagonal
- Tridiagonal
- SymTridiagonal
- BandedMatrix ([BandedMatrices.jl](#))
- BlockBandedMatrix ([BlockBandedMatrices.jl](#))

Matrices which do not satisfy this interface can still be used, but the matrix coloring will not be automatic, and an appropriate linear solver may need to be given (otherwise it will default to attempting an LU-decomposition).

0.5 Sundials-Specific Handling

While much of the setup makes the transition to using Sundials automatic, there are some differences between the pure Julia implementations and the Sundials implementations which must be taken note of. These are all detailed in the [Sundials solver documentation](#), but here we will highlight the main details which one should make note of.

Defining a sparse matrix and a Jacobian for Sundials works just like any other package. The core difference is in the choice of the linear solver. With Sundials, the linear solver choice is done with a Symbol in the `linear_solver` from a preset list. Particular choices of note are `:Band` for a banded matrix and `:GMRES` for using GMRES. If you are using Sundials, `:GMRES` will not require defining the `JacVecOperator`, and instead will always make use of

a Jacobian-Free Newton Krylov (with numerical differentiation). Thus on this problem we could do:

```
using Sundials
# Sparse Version
@btime solve(prob_ode_brusselator_2d_sparse,CVODE_BDF(),save_everystep=false)

29.142 s (51388 allocations: 3.20 MiB)

# GMRES Version: Doesn't require any extra stuff!
@btime
solve(prob_ode_brusselator_2d,CVODE_BDF(linear_solver=:GMRES),save_everystep=false)

350.144 ms (61058 allocations: 3.63 MiB)
retcode: Success
Interpolation: 1st order linear
t: 2-element Array{Float64,1}:
 0.0
 11.5
u: 2-element Array{Array{Float64,3},1}:
 [0.0 0.12134432813715876 ... 0.1213443281371586 0.0; 0.0 0.12134432813715876
 ... 0.1213443281371586 0.0; ... ; 0.0 0.12134432813715876 ... 0.1213443281371586
 0.0; 0.0 0.12134432813715876 ... 0.1213443281371586 0.0]

 [0.0 0.0 ... 0.0 0.0; 0.14892258453196755 0.14892258453196755 ... 0.14892258453
 196755 0.14892258453196755; ... ; 0.14892258453196738 0.14892258453196738 ... 0
 .14892258453196738 0.14892258453196738; 0.0 0.0 ... 0.0 0.0]

 [0.45369441125092624 0.45367162922766396 ... 0.45377307354145824 0.453728249
 24331306; 0.45372813444006976 0.45370139820263283 ... 0.45382031508907966 0.4
 537681622154197; ... ; 0.4536347409999057 0.4536184243336325 ... 0.453690734603
 503 0.4536589378647838; 0.4536631791063342 0.4536436405637919 ... 0.453729310
 5001047 0.45369169445940305]

 [5.023428953606044 5.023425514309876 ... 5.02343972583798 5.0234337753788845;
 5.023442660236476 5.023439873077652 ... 5.02345101637559 5.023446317614284;
 ... ; 5.023404093671991 5.023399216246354 ... 5.023419229667771 5.0234107290209
 42; 5.023415926060523 5.023411776722086 ... 5.02342895844194 5.02342180621704
 3]
```

Details for setting up a preconditioner with Sundials can be found at the [Sundials solver page](#).

0.6 Handling Mass Matrices

Instead of just defining an ODE as $u' = f(u, p, t)$, it can be common to express the differential equation in the form with a mass matrix:

$$Mu' = f(u, p, t)$$

where M is known as the mass matrix. Let's solve the Robertson equation. At the top we wrote this equation as:

$$dy_1 = -0.04y_1 + 10^4 y_2 y_3 \quad (5)$$

$$dy_2 = 0.04y_1 - 10^4 y_2 y_3 - 3 * 10^7 y_2^2 \quad (6)$$

$$dy_3 = 3 * 10^7 y_3^2 \quad (7)$$

$$(8)$$

But we can instead write this with a conservation relation:

$$dy_1 = -0.04y_1 + 10^4 y_2 y_3 \quad (9)$$

$$dy_2 = 0.04y_1 - 10^4 y_2 y_3 - 3 * 10^7 y_2^2 \quad (10)$$

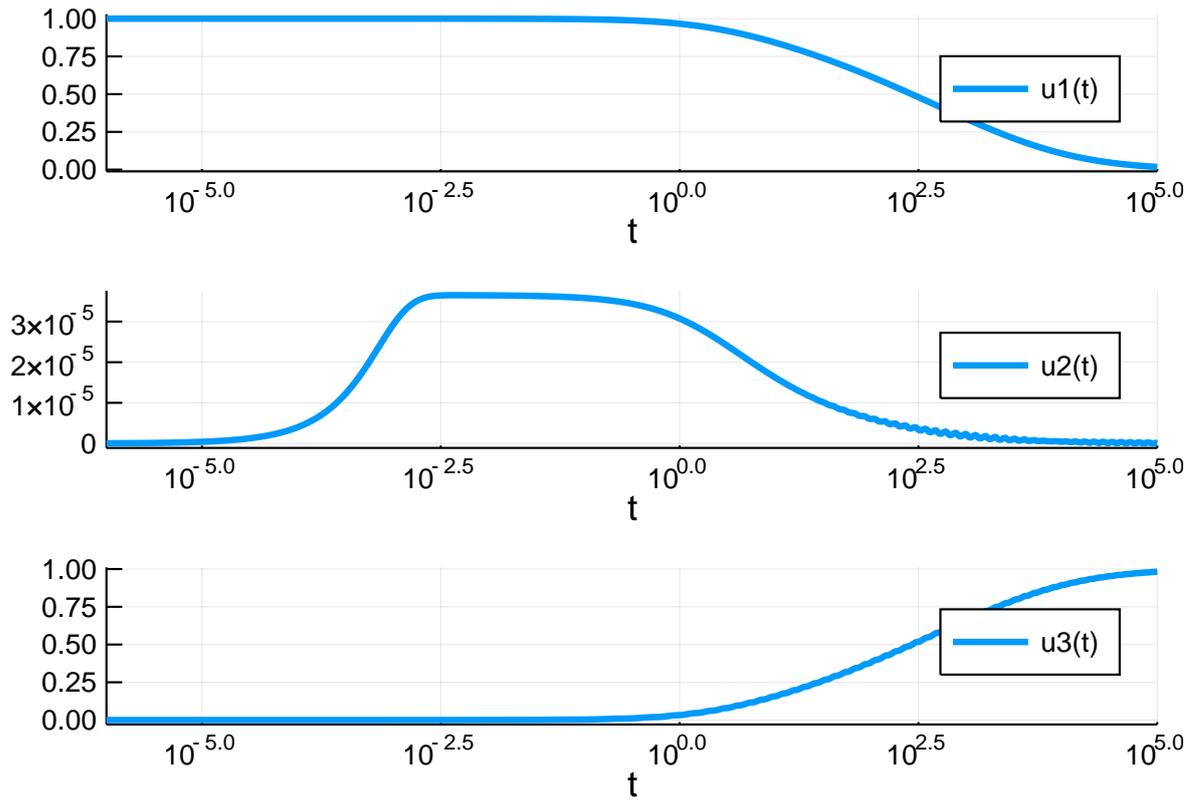
$$1 = y_1 + y_2 + y_3 \quad (11)$$

$$(12)$$

In this form, we can write this as a mass matrix ODE where M is singular (this is another form of a differential-algebraic equation (DAE)). Here, the last row of M is just zero. We can implement this form as:

```
using DifferentialEquations
function rober(du,u,p,t)
    y_1,y_2,y_3 = u
    k_1,k_2,k_3 = p
    du[1] = -k_1*y_1+k_3*y_2*y_3
    du[2] = k_1*y_1-k_2*y_2^2-k_3*y_2*y_3
    du[3] = y_1 + y_2 + y_3 - 1
    nothing
end
M = [1. 0 0
      0 1. 0
      0 0 0]
f = ODEFunction(rober,mass_matrix=M)
prob_mm = ODEProblem(f, [1.0,0.0,0.0], (0.0,1e5), (0.04,3e7,1e4))
sol = solve(prob_mm,Rodas5())

plot(sol, xscale=:log10, tspan=(1e-6, 1e5), layout=(3,1))
```



Note that if your mass matrix is singular, i.e. your system is a DAE, then you need to make sure you choose [a solver that is compatible with DAEs](#)