**Advanced Topics from Scientific Computing**
**TU Berlin Winter 2022/23**
**Notebook 07**

**Jürgen Fuhrmann**

```
begin
    using LinearAlgebra
    using BenchmarkTools
    using Pkg
end
```

Direct solution of linear systems of equations
LU Factorization for dense matrices
LU Factorization for sparse matrices
1D PDE problems
2D PDE problems
Complexity estimate for sparse direct solvers
One more thing: how to create a sparse matrix
Final remarks

# Direct solution of linear systems of equations

## LU Factorization for dense matrices

Let aus create a matrix and solve the corresponding linear system:

```
n1 = 1000
```
```
n1=1000
```

```
A1 = 1000×1000 Matrix{Float64}:
 0.693248    0.232691  0.748356   0.283001   …  0.667931    0.204756   0.306197
 0.967957    0.82514   0.8328     0.237919      0.97181     0.201167   0.395003
 0.29409     0.510149  0.990261   0.518299      0.138954    0.262015   0.641518
 0.783424    0.955413  0.30494    0.881054      0.995383    0.451897   0.1188
 0.522539    0.995946  0.627975   0.638433      0.00522948  0.956318   0.743116
 0.785752    0.873038  0.501976   0.883699   …  0.700803    0.756529   0.865197
 0.329995    0.230966  0.526016   0.105139      0.303009    0.0174611  0.336172
 ⋮                                           ⋱
 0.00357113  0.364658  0.233611   0.141952      0.307982    0.495757   0.564955
 0.348737    0.373673  0.506648   0.438111   …  0.287645    0.637436   0.357411
 0.124254    0.723455  0.973879   0.482691      0.623194    0.133019   0.487457
 0.213532    0.357382  0.470023   0.0166586     0.536065    0.299495   0.666153
 0.888758    0.528283  0.995849   0.564579      0.605873    0.599814   0.603267
 0.683135    0.615527  0.0306897  0.397419      0.0191046   0.321058   0.605768
```
```
A1=rand(n1,n1)
```

```
x1 =
 [0.0835718, 0.570492, 0.457436, 0.800853, 0.197794, 0.697688, 0.132689, 0.0889008, 0.06345
```
```
x1=rand(n1)
```

```
b1 =
 [244.736, 246.721, 260.593, 251.688, 255.152, 264.101, 257.239, 254.639, 248.813, 249.344,
```
```
b1=A1*x1
```

```
2.18614015778939\57e-12
```
```
norm(A1\b1-x1, Inf)
```

The "\" operator provides a default solver for linear systems of equations based on the **LU factorization** of the matrix into an upper and a lower triangular matrix, and the subsequent solution of the triangular systems:

$$A = LU$$
$$Ly = b$$
$$Ux = y$$

This approach is equivalent to the Gaussian elimination process. The algorithm is improved by stability enhancing **pivoting** - reordering of the system of equations such that divisions by small main diagonal elements is avoided.

We can demonstrate this process:

```
lu1 = LU{Float64, Matrix{Float64}, Vector{Int64}}
      L factor:
      1000×1000 Matrix{Float64}:
       1.0         0.0         0.0         …  0.0        0.0        0.0       0.0
       0.0518669   1.0         0.0            0.0        0.0        0.0       0.0
       0.969827   -0.940419    1.0            0.0        0.0        0.0       0.0
       0.248192   -0.0403483   0.00202699     0.0        0.0        0.0       0.0
       0.468948    0.562996   -0.488007       0.0        0.0        0.0       0.0
       0.779174   -0.640945    0.939828    …  0.0        0.0        0.0       0.0
       0.969857   -0.0762975   0.374435       0.0        0.0        0.0       0.0
       ⋮                                   ⋱
       0.427535   -0.138868    0.215564       0.0        0.0        0.0       0.0
       0.956484   -0.710936    0.364323    …  0.0        0.0        0.0       0.0
       0.354957    0.173867   -0.16407        1.0        0.0        0.0       0.0
       0.462897    0.450184    0.052842      -0.925787   1.0        0.0       0.0
       0.149157    0.459136   -0.304943      -0.320925  -0.629319   1.0       0.0
       0.462981   -0.125944    0.557686       0.578528   0.144964  -0.256556  1.0
      U factor:
      1000×1000 Matrix{Float64}:
       0.99804  0.924528  0.424142  0.311441  …   0.0333858   0.136      0.0555658
       0.0      0.937384  0.916286  0.406637      0.642873    0.76253    0.633635
       0.0      0.0       1.31225   0.326986      0.984465    0.760908   0.946926
       0.0      0.0       0.0       0.898911      1.00976     0.26399    0.568143
       0.0      0.0       0.0       0.0           0.366248    0.489084   0.627371
       0.0      0.0       0.0       0.0       …  -0.723896    0.55752   -0.0855989
       0.0      0.0       0.0       0.0           1.24888    -0.57103    0.104092
       ⋮                                      ⋱
       0.0      0.0       0.0       0.0           0.0820403  -2.4489    -2.5543
       0.0      0.0       0.0       0.0       …   4.22801    -2.91279   -5.37484
       0.0      0.0       0.0       0.0           4.08454     5.25878   -1.80898
       0.0      0.0       0.0       0.0           6.40455     3.42185    3.11135
       0.0      0.0       0.0       0.0           0.0         8.04945    2.83873
       0.0      0.0       0.0       0.0           0.0         0.0       -2.35751
```

```julia
lu1=lu(A1)
```

Extract the pivoting permutation:

```
p =
 [994, 773, 325, 570, 948, 736, 2, 429, 777, 508, 203, 420, 757, 423, 78, 421, 45, 353, 829,
```

```julia
p=lu1.p
```

Permute the right hand side vector:

```
b1_permuted =
 [263.805, 255.389, 254.666, 254.681, 245.603, 259.258, 246.721, 247.357, 260.931, 251.208,
```

```julia
b1_permuted=b1[p]
```

Solve the triangular systems with L and U

```
y =
 [263.805, 241.706, 226.125, 198.501, 103.531, -80.4637, 0.575317, 184.254, -135.836, 23.31
```

```julia
y=lu1.L\b1_permuted
```

```
x1_lu =
 [0.0835718, 0.570492, 0.457436, 0.800853, 0.197794, 0.697688, 0.132689, 0.0889008, 0.06345
```

```julia
x1_lu=lu1.U\y
```

```
2.1861401577893957e-12
```

```julia
norm(x1-x1_lu, Inf)
```

These steps are combined in the "\" operator for LU factorizations

```
 [0.0835718, 0.570492, 0.457436, 0.800853, 0.197794, 0.697688, 0.132689, 0.0889008, 0.06345
```

```julia
A1\b1
```

```
 [0.0835718, 0.570492, 0.457436, 0.800853, 0.197794, 0.697688, 0.132689, 0.0889008, 0.06345
```

```julia
inv(A1)*b1
```

- LU factorization takes the most time in this approach, triangular solves are fast
- LU factorization is significantly faster than matrix inversion
- Symmetry of the matrix can be utilized to speed up the resulting method is called Cholesky factorization
- For standard floating point types, Julia uses highly optimized versions of LAPACK and BLAS
  - Same for python/numpy and many other coding environments

# LU Factorization for sparse matrices

As we focus in this course on partial differential equations, we need discuss matrices which evolve from the discretization of PDEs.

- Are there any structural or numerical patterns in these matrices we can take advantage of with regard to memory and time complexity when solving linear systems ?

In this lecture we introduce a relatively simple "drosophila" problem which we will use do discuss these issues.

For the start we use simple structured disceretization grids and a finite difference approach to the discretization. Later, this will be generalized to more general grids and to finite element and finite volume discretization methods.

## 1D PDE problems

Assume a one-dimensional rod

- Heat source $f(x)$
- $v_L, v_R$: ambient temperatures
- $\alpha$: boundary heat transfer coefficient
- Second order boundary value problem in $\Omega = [0,1]$:

$$\begin{cases} -u''(x) & = f(x) \text{ in } \Omega \\ -u'(0) + \alpha(u(0) - v_L) & = 0 \\ u'(1) + \alpha(u(1) - v_R) & = 0 \end{cases}$$

- The solution $u$ describes the equilibrium temperature distribution. Behind the second derivative is Fouriers law and the continuity equation
- In math, the boundary conditions are called "Robin" or "third kind". They describe a heat in/outflux proportional to the difference between rod end temperature and ambient temperature
- Fix a number of discretization points N
- Let $h = \frac{1}{N-1}$
- Let $x_i = (i-1)h \;\; i = 1 \dots N$ be discretization points

```
N0=10;
```

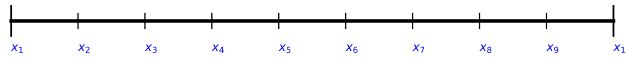


We can approximate continuous functions $f$ by piecewise linear functions defined by the values $f_i = f(x_i)$. Using more points yields a better approximation:

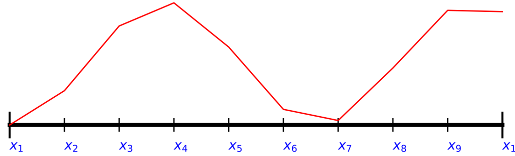


- Let $u_i$ approximations for $u(x_i)$ and $f_i = f(x_i)$
- We can use a finite difference approximation to approximate $u'(x_{i+\frac{1}{2}}) \approx \frac{u_{i+1} - u_i}{h}$
- Same approach for second derivative: $u''(x_i) = \frac{u'(x_{i+\frac{1}{2}}) - u'(x_{i-\frac{1}{2}})}{h}$
- Finite difference approximation of the PDE:

$$\begin{aligned} -u'(0) + \alpha(u(0) - v_L) &\approx \frac{1}{2h}(u_0 - u_2) + \alpha(u_1 - v_L) = 0 \\ -u''(x_i) - f(x_i) &\approx \frac{-u_{i+1} + 2u_i - u_{i-1}}{h^2} - f_i = 0 && (i = 1 \dots N) \; (*) \\ u'(1) + \alpha(u(1) - v_R) &\approx \frac{1}{2h}(u_{N+1} - u_{N-1}) + \alpha(u_N - v_R) = 0 \end{aligned}$$

- Here, we introduced "mirror values" $u_0$ and $u_{N+1}$ in order to approximate the boundary conditions accurately, such that the finite difference formulas used to approximate $u'(0)$ or $u'(1)$ are centered around these values.
- After rearranging, these values can be expressed via the boundary conditions:

$$\begin{aligned} u_0 &= u_2 + 2h\alpha(u_1 - v_L) \\ u_{N+1} &= u_{N-1} + 2h\alpha(u_N - v_L) \end{aligned}$$

- Finally, they can be replaced in $(*)$

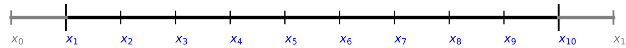


The resulting discretization matrix is

$$A = \begin{pmatrix} \alpha + \frac{1}{h} & -\frac{1}{h} & & & & & \\ -\frac{1}{h} & \frac{2}{h} & -\frac{1}{h} & & & & \\ & -\frac{1}{h} & \frac{2}{h} & -\frac{1}{h} & & & \\ & \ddots & \ddots & \ddots & \ddots & & \\ & & -\frac{1}{h} & \frac{2}{h} & -\frac{1}{h} & & \\ & & & -\frac{1}{h} & \frac{2}{h} & -\frac{1}{h} & \\ & & & & -\frac{1}{h} & \frac{1}{h} + \alpha \end{pmatrix}$$

Outside of the three diagonals, the entries are zero.

The right hand side is:

$$f = \begin{pmatrix} \frac{h}{2} f_1 + \alpha v_L \\ h f_2 \\ h f_3 \\ \vdots \\ h f_{N-2} \\ h f_{N-1} \\ \frac{h}{2} f_N + \alpha v_R \end{pmatrix}$$

We want to solve $Au = f$

Let us define functions assembling these:

heatmatrix1d! (generic function with 1 method)

```
function heatmatrix1d!(A,N;α=1)
    h=1/(N-1)
    A[1,1]=1/h+α
    for i=2:N-1
        A[i,i]=2/h
    end
    for i=1:N-1
        A[i,i+1]=-1/h
    end
    for i=2:N
        A[i,i-1]=-1/h
    end
    A[N,N]=1/h+α
    A
end
```

heatrhs1d (generic function with 1 method)

```
function heatrhs1d(N;vl=0,vr=0,func=x->0,α=1)
    h=1/(N-1)
    F=zeros(N)
    F[1]=h/2*func(0)+α*vl
    for i=2:N-1
        F[i]=h*func((i-1)*h)
    end
    F[N]=h/2*func(1)+α*vr
    F
end
```

α = 100

```
α=100
```

N1 = 100

```
N1=100
```

A1_heat =

```
100×100 Matrix{Float64}:
 199.0  -99.0    0.0    0.0    0.0    0.0  …    0.0    0.0    0.0    0.0    0.0
 -99.0  198.0  -99.0    0.0    0.0    0.0       0.0    0.0    0.0    0.0    0.0
   0.0  -99.0  198.0  -99.0    0.0    0.0       0.0    0.0    0.0    0.0    0.0
   0.0    0.0  -99.0  198.0  -99.0    0.0       0.0    0.0    0.0    0.0    0.0
   0.0    0.0    0.0  -99.0  198.0  -99.0       0.0    0.0    0.0    0.0    0.0
   0.0    0.0    0.0    0.0  -99.0  198.0  …    0.0    0.0    0.0    0.0    0.0
   0.0    0.0    0.0    0.0    0.0  -99.0       0.0    0.0    0.0    0.0    0.0
   ⋮                                  ⋮    ⋱    ⋮
   0.0    0.0    0.0    0.0    0.0    0.0     -99.0    0.0    0.0    0.0    0.0
   0.0    0.0    0.0    0.0    0.0    0.0  …  198.0  -99.0    0.0    0.0    0.0
   0.0    0.0    0.0    0.0    0.0    0.0     -99.0  198.0  -99.0    0.0    0.0
   0.0    0.0    0.0    0.0    0.0    0.0       0.0  -99.0  198.0  -99.0    0.0
   0.0    0.0    0.0    0.0    0.0    0.0       0.0    0.0  -99.0  198.0  -99.0
   0.0    0.0    0.0    0.0    0.0    0.0       0.0    0.0    0.0  -99.0  199.0
```

```
A1_heat=heatmatrix1d!(zeros(N1,N1),N1,α=α)
```

b1_heat =

```
[0.00505051, 0.010101, 0.010101, 0.010101, 0.010101, 0.010101, 0.010101, 0.010101, 0.01010
```

```
b1_heat=heatrhs1d(N1,func=x->1,α=α)
```

u1_heat =

```
[0.005, 0.00999949, 0.0148969, 0.0196924, 0.0243858, 0.0289771, 0.0334665, 0.0378538, 0.04
```

```
u1_heat=A1_heat\b1_heat
```



allN = [20, 40, 80, 160, 320, 640, 1280, 2560, 5120, 10240]

times_full =

```
[1.2533e-5, 1.5776e-5, 4.1468e-5, 0.000155477, 0.000830601, 0.00497244, 0.0334687, 0.22564
```

```
times_full= let
    times=[]
  for N∈allN
        A=heatmatrix1d!(zeros(N,N),N,α=α)
        b=heatrhs1d(N,α=α)
        t=@elapsed u=A\b
        push!(times,t)
  end
    times
end
```

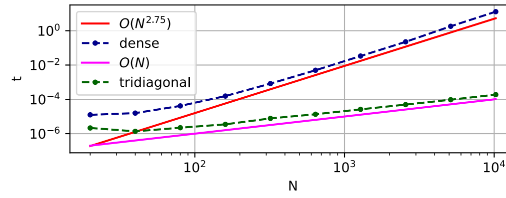


The complexity of the LU factorization of a dense matrix is estimated with $O(N^{2.75})$

However, we can store only three diagonals of the matrix and apply methods tailored to this case:

Try tridiagonal ☑

```
[2.147e-6, 1.366e-6, 2.222e-6, 3.502e-6, 7.768e-6, 1.353e-5, 2.6204e-5, 4.9427e-5, 9.513e-
```

```
if try_tridiagonal
   times_tridiagonal= let
    times=[]
  for N∈allN
       A=Tridiagonal(zeros(N-1),zeros(N),zeros(N-1))
       A=heatmatrix1d!(A,N,α=α)
       b=heatrhs1d(N,α=α)
       t=@elapsed u=A\b
       push!(times,t)
  end
    times
  end
end
```

We learn, that in this case, solution time is O(N), much better.

One caveat: Never calculate the inverse of a matrix from a PDE, as it neither tridiagonal nor sparse anymore!

```
10×10 Matrix{Float64}:
 0.00990196  0.00881264  0.00772331  …  0.00227669  0.00118736  9.80392e-5
 0.00881264  0.106731    0.0935379      0.0275732   0.0143803   0.00118736
 0.00772331  0.0935379   0.179352       0.0528698   0.0275732   0.00227669
 0.00663399  0.080345    0.154056       0.0781663   0.0407662   0.00336601
 0.00554466  0.067152    0.128759       0.103463    0.0539591   0.00445534
 0.00445534  0.0539591   0.103463    …  0.128759    0.067152    0.00554466
 0.00336601  0.0407662   0.0781663      0.154056    0.080345    0.00663399
 0.00227669  0.0275732   0.0528698      0.179352    0.0935379   0.00772331
 0.00118736  0.0143803   0.0275732      0.0935379   0.106731    0.00881264
 9.80392e-5  0.00118736  0.00227669     0.00772331  0.00881264  0.00990196
```

```
let
    A=Tridiagonal(zeros(N0-1),zeros(N0),zeros(N0-1))
    A=heatmatrix1d!(A,N0,α=α)
    inv(A)
end
```

## 2D PDE problems

Just pose the heat problem in a 2D domain $\Omega = (0,1) \times (0,1)$:

$$\begin{cases} -\frac{\partial^2 u}{\partial x^2} - \frac{\partial^2 u}{\partial y^2} & = f(x,y) \text{ in } \Omega \\ \frac{\partial u}{\partial n} + \alpha(u - v) & = 0 \text{ on } \partial\Omega \end{cases}$$

We use 2D regular discretization $n \times n$ grid with grid points $x_{ij} = ((i-1)h, (j-1)h)$. The finite difference approximation yields:

$$\frac{-u_{i-1,j} - u_{i,j-1} + 4u_{ij} - u_{i+1,j} - u_{i,j+1}}{h^2} = f_{ij}$$

This just comes from summing up the 1D finite difference formula for the $x$ and $y$ directions.

We do not discuss the boundary conditions here.

The $n \times n$ grid leads to an $n^2 \times n^2$ matrix!

```
plotgrid2d (generic function with 1 method)
```

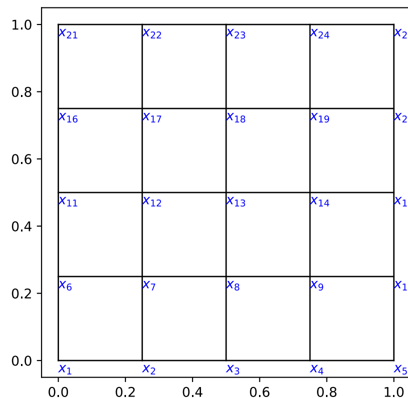


```
plotgrid2d(5)
```

Matrix and right hand side assembly inspired by the finite volume method which will be covered later in the course. The result is the same as for the finite difference method with the mirror trick for the boundary condition.

```
heatmatrix2d! (generic function with 1 method)
```

```julia
function heatmatrix2d!(A,n;α=1)
    function update_pair(A,v,i,j)
        A[i,j]+=-v
        A[j,i]+=-v
        A[i,i]+=v
        A[j,j]+=v
    end
    N=n^2
    h=1.0/(n-1)
    l=1
        for j=1:n
            for i=1:n
                if i<n
                    update_pair(A,1.0,l,l+1)
                end
                if i==1|| i==n
                    A[l,l]+=α
                end
                if j<n
                    update_pair(A,1,l,l+n)
                end
                if j==1|| j==n
                    A[l,l]+=α
                end
                l=l+1
            end
        end
    A
end
```

```
heatrhs2d (generic function with 1 method)
```

```julia
function heatrhs2d(n; rhs=(x,y)->0,bc=(x,y)->0,α=1.0)
    h=1.0/(n-1)
    x=collect(0:h:1)
    y=collect(0:h:1)
    N=n^2
    f=zeros(N)
    for i=1:n-1
        for j=1:n-1
            ij=(j-1)*n+i
            f[ij]+=h^2/4*rhs(x[i],y[j])
            f[ij+1]+=h^2/4*rhs(x[i+1],y[j])
            f[ij+n]+=h^2/4*rhs(x[i],y[j+1])
            f[ij+n+1]+=h^2/4*rhs(x[i+1],y[j+1])
        end
    end

    for i=1:n
        ij=i
        fac=h
        if i==1 || i==n
            fac=h/2
        end
        f[ij]+=fac*α*bc(x[i],0)
        ij=i+(n-1)*n
        f[ij]+=fac*α*bc(x[i],1)
    end
    for j=1:n
        fac=h
        if j==1 || j==n
            fac=h/2
        end
        ij=1+(j-1)*n
        f[ij]+=fac*α*bc(0,y[j])
        ij=n+(j-1)*n
        f[ij]+=fac*α*bc(1,y[j])
    end
    f
end
```

```
n = 20
```

```julia
n=20
```

```
b2 =
[0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0.0, 0
```

```julia
b2=heatrhs2d(n,rhs=(x,y)->sin(3*π*x)*sin(3*π*y),α=α)
```

```
A2 = 400×400 Matrix{Float64}:
 202.0   -1.0    0.0    0.0    0.0    0.0  …    0.0    0.0    0.0    0.0    0.0
  -1.0  103.0   -1.0    0.0    0.0    0.0       0.0    0.0    0.0    0.0    0.0
   0.0   -1.0  103.0   -1.0    0.0    0.0       0.0    0.0    0.0    0.0    0.0
   0.0    0.0   -1.0  103.0   -1.0    0.0       0.0    0.0    0.0    0.0    0.0
   0.0    0.0    0.0   -1.0  103.0   -1.0       0.0    0.0    0.0    0.0    0.0
   0.0    0.0    0.0    0.0   -1.0  103.0  …    0.0    0.0    0.0    0.0    0.0
   0.0    0.0    0.0    0.0    0.0   -1.0       0.0    0.0    0.0    0.0    0.0
   ⋮                                  ⋮     ⋱    ⋮
   0.0    0.0    0.0    0.0    0.0    0.0      -1.0    0.0    0.0    0.0    0.0
   0.0    0.0    0.0    0.0    0.0    0.0  …  103.0   -1.0    0.0    0.0    0.0
   0.0    0.0    0.0    0.0    0.0    0.0      -1.0  103.0   -1.0    0.0    0.0
   0.0    0.0    0.0    0.0    0.0    0.0       0.0   -1.0  103.0   -1.0    0.0
   0.0    0.0    0.0    0.0    0.0    0.0       0.0    0.0   -1.0  103.0   -1.0
   0.0    0.0    0.0    0.0    0.0    0.0       0.0    0.0    0.0   -1.0  202.0
```

```julia
A2=heatmatrix2d!(zeros(n^2,n^2),n,α=α)
```

```julia
using Tables
```

| | Column1 | Column2 | Column3 | Column4 | Column5 | Column6 | Column7 | Column8 |
|---|---|---|---|---|---|---|---|---|
| 1 | 202.0 | -1.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| 2 | -1.0 | 103.0 | -1.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| 3 | 0.0 | -1.0 | 103.0 | -1.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| 4 | 0.0 | 0.0 | -1.0 | 103.0 | -1.0 | 0.0 | 0.0 | 0.0 |
| 5 | 0.0 | 0.0 | 0.0 | -1.0 | 103.0 | -1.0 | 0.0 | 0.0 |
| 6 | 0.0 | 0.0 | 0.0 | 0.0 | -1.0 | 103.0 | -1.0 | 0.0 |
| 7 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | -1.0 | 103.0 | -1.0 |
| 8 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | -1.0 | 103.0 |
| 9 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | -1.0 |
| 10 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |

more

```
Tables.table(A2)
```

u2 =
[1.28855e-7, 1.30143e-5, 2.28006e-5, 2.71143e-5, 2.49065e-5, 1.6704e-5, 4.48045e-6, -8.819

```
u2=A2\b2
```

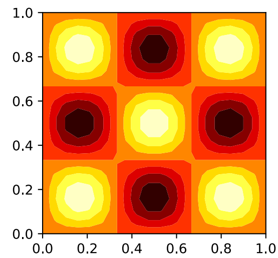


n2d = [8, 11, 15, 21, 29, 41, 57, 81]

```
n2d=Int[ceil(5*sqrt(2)^i) for i=1:8]
```

times_full2d =
[4.1288e-5, 0.00012452, 0.000438527, 0.00219039, 0.0108819, 0.0671273, 0.475402, 3.31411]

```
times_full2d= let
    times=[]
  for n∈n2d
       A=heatmatrix2d!(zeros(n^2,n^2),n,α=α)
       b=heatrhs2d(n,α=α)
       t=@elapsed u=A\b
       push!(times,t)
  end
    times
end
```

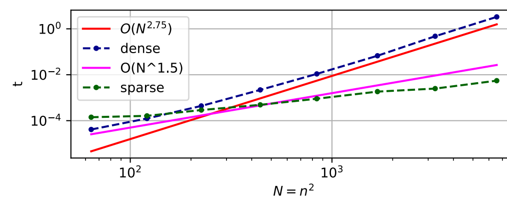


The matrix of this system has 5 nonzero diagonals. While it is possible to store just these five diagonals, there is not much software around which handles this case. But we can use the possibility to store it as a **sparse** matrix.

Try sparse ☑

```
using SparseArrays
```

[0.000141471, 0.00016211, 0.000290197, 0.000489928, 0.000903914, 0.00184265, 0.00249679, 0

```
if try_sparse
times_sparse2d= let
    times=[]
  for n∈n2d
       A0=spzeros(n^2,n^2)
       A=heatmatrix2d!(A0,n,α=α)
       b=heatrhs2d(n,α=α)
       t=@elapsed u=A\b
       push!(times,t)
  end
    times
end
end
```

In Julia, sparse matrices are stored in Compressed Column Storage (CSC) format.

```
A = 5×5 Matrix{Float64}:
 1.0  0.0  5.0   8.0   0.0
 2.0  4.0  0.0   9.0   0.0
 3.0  0.0  6.0   0.0  11.0
 0.0  0.0  7.0  10.0   0.0
 0.0  0.0  0.0   0.0  12.0
```

```
A=Float64[1 0  5  8 0;
          2 4  0  9  0;
          3 0  6  0  11;
          0 0  7 10 0;
          0 0  0  0  12]
```

```
As = 5×5 SparseMatrixCSC{Float64, Int64} with 12 stored entries:
 1.0   ⋅   5.0   8.0    ⋅
 2.0  4.0   ⋅    9.0    ⋅
 3.0   ⋅   6.0    ⋅   11.0
  ⋅    ⋅   7.0  10.0    ⋅
  ⋅    ⋅    ⋅     ⋅   12.0
```

```
As=sparse(A)
```

The matrix is stored using three vectors:

`nzval` contains the values of the nonzero entries, stored contiguously column by column:

```
[1.0, 2.0, 3.0, 4.0, 5.0, 6.0, 7.0, 8.0, 9.0, 10.0, 11.0, 12.0]
```

```
As.nzval
```

`rowval` contains the row indices of the nonzero entries

```
[1, 2, 3, 2, 1, 3, 4, 1, 2, 4, 3, 5]
```

```
As.rowval
```

`colptr` points to the start of a column in `rowval` and `nzval`

```
[1, 4, 5, 8, 11, 13]
```

```
As.colptr
```

So we can access the nonzeros and the rowvals column by column:

```
[[1.0, 2.0, 3.0], [4.0], [5.0, 6.0, 7.0], [8.0, 9.0, 10.0], [11.0, 12.0]]
```

```
[As.nzval[As.colptr[i]:As.colptr[i+1]-1] for i=1:size(A,1)]
```

```
[[1, 2, 3], [2], [1, 3, 4], [1, 2, 4], [3, 5]]
```

```
[As.rowval[As.colptr[i]:As.colptr[i+1]-1] for i=1:size(A,1)]
```

Sparse direct solvers allow to handle LU factorizations for sparse matrices. For 2D problems, the complexity of the LU factorization is $O(N^{1.5})$

```
Np = 100000
```

```
Np=100_000
```

```
dim = 2
```

```
dim=2
```

```
np = 317
```

```
np=Int(ceil(Np^(1/dim)))
```

```
100489×100489 SparseMatrixCSC{Float64, Int64} with 501177 stored entries:
```

```
if dim==1
    Ap=heatmatrix1d!(spzeros(Np,Np),np)
elseif dim==2
    Ap=heatmatrix2d!(spzeros(np^2,np^2),np)
end
```

```
LUp = SuiteSparse.UMFPACK.UmfpackLU{Float64, Int64}
      L factor:
      100489×100489 SparseMatrixCSC{Float64, Int64} with 3089774 stored entries:

      U factor:
      100489×100489 SparseMatrixCSC{Float64, Int64} with 3089774 stored entries:
```

```
LUp=lu(Ap)
```

```
12.330071012835784
```

```
nnz(LUp)/nnz(Ap)
```

We observe that for space dimension>1 the number of nonzero entries of a sparse LU factorization is significantly larger than the number of nonzeros of the original matrix and depends on the ordering of the unknowns. This phenomenon is called **fill-in**.

Solution steps with sparse direct solvers

1. Pre-ordering
   - Decrease amount of non-zero elements generated by fill-in by re-ordering of the matrix
   - Several, graph theory based heuristic algorithms exist
   - Julia uses resonable defaults with UMFPACK
2. Symbolic factorization
   - If pivoting is ignored, the indices of the non-zero elements are calculated and stored
   - Most expensive step wrt. computation time
3. Numerical factorization
   - Calculation of the numerical values of the nonzero entries
   - Moderately expensive, once the symbolic factors are available
4. Upper/lower triangular system solution
   - Fairly quick in comparison to the other steps

- Separation of steps 2 and 3 allows to save computational costs for problems where the sparsity structure remains unchanged, e.g. time dependent problems on fixed computational grids
- With pivoting, steps 2 and 3 have to be performed together, and pivoting can increase fill-in
- Instead of pivoting, *iterative refinement* may be used in order to maintain accuracy of the solution

## Complexity estimate for sparse direct solvers

- Complexity estimates depend on storage scheme, reordering etc.
- Sparse matrix - vector multiplication has complexity $O(N)$
- Some estimates can be given from graph theory for discretizations of heat equation with $N = n^d$ unknowns on close to cubic grids in space dimension $d$
- sparse LU factorization:

| $d$ | *work* | *storage* |
|---|---|---|
| 1 | $O(N) \mid O(n)$ | $O(N) \mid O(n)$ |
| 2 | $O(N^{\frac{3}{2}}) \mid O(n^3)$ | $O(N \log N) \mid O(n^2 \log n)$ |
| 3 | $O(N^2) \mid O(n^6)$ | $O(N^{\frac{4}{3}}) \mid O(n^4)$ |

- triangular solve: work dominated by storage complexity

| $d$ | *work* |
|---|---|
| 1 | $O(N) \mid O(n)$ |
| 2 | $O(N \log N) \mid O(n^2 \log n)$ |
| 3 | $O(N^{\frac{4}{3}}) \mid O(n^4)$ |

Source: J. Poulson, Fast parallel solution of heterogeneous 3D time-harmonic wave equations (PhD thesis, UT Austin, 2012)

## One more thing: how to create a sparse matrix

```
N3 = 200
N3=200
```

A sparse matrix `A` in Julia can be updated just by writing into `A[i,j]`, updating the nonzero entries is done automatically. So start with a sparse matrix with no nonzero entries and just write into it...

```
BenchmarkTools.Trial: 9 samples with 1 evaluation.
 Range (min … max):  546.810 ms … 619.819 ms  ┊ GC (min … max): 0.00% … 0.00%
 Time  (median):     598.617 ms               ┊ GC (median):    0.00%
 Time  (mean ± σ):   595.068 ms ±  20.878 ms  ┊ GC (mean ± σ):  0.00% ± 0.00%
```

```
 547 ms            Histogram: frequency by time            620 ms <

 Memory estimate: 6.22 MiB, allocs estimate: 30.
```

```
@benchmark heatmatrix2d!(spzeros(N3^2,N3^2),N3)
```

```
BenchmarkTools.Trial: 124 samples with 1 evaluation.
 Range (min … max):  38.695 ms … 46.067 ms  ┊ GC (min … max): 0.00% … 0.00%
 Time  (median):     40.159 ms              ┊ GC (median):    0.00%
 Time  (mean ± σ):   40.314 ms ±  1.139 ms  ┊ GC (mean ± σ):  0.81% ± 1.35%
```

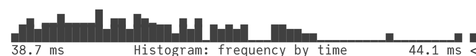

```
 38.7 ms          Histogram: frequency by time          44.1 ms <

 Memory estimate: 34.30 MiB, allocs estimate: 66.
```

```
let
    A=heatmatrix2d!(spzeros(N3^2,N3^2),N3)
    b=rand(N3^2)
    @benchmark $A\$b
end
```

Matrix build-up is much more expensive than solution of the linear system.

... Re-arranging the internal structure is connected to shifting and re-allocating th the arrays many times.

A frequent recommendation ist to use the "coordinate" or "triplet" format as an intermediate: Just collect in three arrays I, J, A all updates of matrix entries and pass them to `sparse`:

heatmatrix2d_coo (generic function with 1 method)

```
function heatmatrix2d_coo(n;α=1)
    I=Int[]
    J=Int[]
    A=Float64[]
    function addentry(i,j,v)
        push!(I,i)
        push!(J,j)
        push!(A,v)
    end
    function update_pair(i,j,v)
        addentry(i,j,-v)
        addentry(j,i,-v)
        addentry(i,i,v)
        addentry(j,j,v)
    end
    N=n^2
    h=1.0/(n-1)
    l=1
        for j=1:n
            for i=1:n
                if i<n
                    update_pair(l,l+1,1)
                end
                if i==1|| i==n
                    addentry(l,l,α)
                end
                if j<n
                    update_pair(l,l+n,1)
                end
                if j==1|| j==n
                    addentry(l,l,α)
                end
                l=l+1
            end
        end
    sparse(I,J,A)
end
```

true

```
heatmatrix2d!(spzeros(N3^2,N3^2),N3)==heatmatrix2d_coo(N3)
```

```
BenchmarkTools.Trial: 511 samples with 1 evaluation.
 Range (min … max):  8.256 ms … 13.592 ms  ┊ GC (min … max): 0.00% … 9.22%
 Time  (median):     9.340 ms              ┊ GC (median):    0.00%
 Time  (mean ± σ):   9.781 ms ±  1.369 ms  ┊ GC (mean ± σ):  4.33% ± 6.76%
```

```
 8.26 ms        Histogram: frequency by time        13.2 ms <

 Memory estimate: 24.51 MiB, allocs estimate: 57.
```

```
@benchmark heatmatrix2d_coo(N3)
```

This approach requires to modify the structure of the assembly loop. If we run through this loop once this is ok. If one wants to update the nonzero entries, one needs to implement this loop twice. Moreover, one loses the intuitive way of writing into a matrix.

The ExtendableSparse.jl packages provides a remedy. It uses a linked list internal representation to build up the matrix and hides it behind the intuitive way of writing into a matrix. A `flush!` method sets up a SparseMatrixCSC structure after assembly.

```
using ExtendableSparse
```

```
BenchmarkTools.Trial: 482 samples with 1 evaluation.
 Range (min … max):   8.766 ms …  13.052 ms  ┊ GC (min … max): 0.00% … 17.28%
 Time  (median):     10.402 ms               ┊ GC (median):    0.00%
 Time  (mean ± σ):   10.367 ms ± 942.968 μs  ┊ GC (mean ± σ):  1.81% ±  4.31%
```

```
 8.77 ms         Histogram: frequency by time         12.6 ms <

 Memory estimate: 11.54 MiB, allocs estimate: 30.
```

```
@benchmark begin
    Aext=heatmatrix2d!(ExtendableSparseMatrix(N3^2,N3^2),N3)
    flush!(Aext)
end
```

This approach can be made faster by using `updateindex!` instead of `+=`. In Julia issue 15630 it is proposed that the compiler should detect this situation.

heatmatrix2d_update! (generic function with 1 method)

```
function heatmatrix2d_update!(A,n;α=1)
    function update_pair(A,v,i,j)
        updateindex!(A,+,-v,i,j)
        updateindex!(A,+,-v,j,i)
        updateindex!(A,+,v,i,i)
        updateindex!(A,+,v,j,j)
    end
    N=n^2
    h=1.0/(n-1)
    l=1
        for j=1:n
            for i=1:n
                if i<n
                    update_pair(A,1.0,l,l+1)
                end
                if i==1|| i==n
                    updateindex!(A,+,α,l,l)
                end
                if j<n
                    update_pair(A,1.0,l,l+n)
                end
                if j==1|| j==n
                    updateindex!(A,+,α,l,l)
                end
                l=l+1
            end
        end
    A
end
```

```
BenchmarkTools.Trial: 707 samples with 1 evaluation.
 Range (min … max):  6.433 ms …   9.019 ms  ┊ GC (min … max): 0.00% … 24.87%
 Time  (median):     6.799 ms               ┊ GC (median):    0.00%
 Time  (mean ± σ):   7.066 ms ± 584.933 μs  ┊ GC (mean ± σ):  2.49% ±  5.63%
```

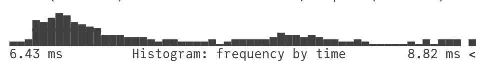

```
  6.43 ms         Histogram: frequency by time        8.82 ms <

 Memory estimate: 11.54 MiB, allocs estimate: 30.
```

```
@benchmark begin
    Aext=heatmatrix2d_update!(ExtendableSparseMatrix(N3^2,N3^2),N3)
    flush!(Aext)
end
```

# Final remarks

- As a rule, direct solution of linear systems of equations is implemented via LU factorization
- Matrices from finite difference methods for PDEs are sparse. True also for finite elements and finite volume methods.
- LU factorizations from sparse matrices suffer from fill-in: LU factors tend to have more nonzeor entries than the original matrices. In 3D significantly so.
- Inverses of matrices from PDEs tend to be full matrices
- The Julia `\` operator by default maps to the UMFPACK sparse direct solver from the Suitesparse collection by T. Davis.
- Other sparse direct solvers (e.g. the thread-parallel Pardiso) are available as Julia packages.