Finite element methods in scientific computing

Wolfgang Bangerth, Colorado State University

Lecture 3.98:

The ideas behind the finite element method

Part 9: Sparsity as a key property

On linear systems resulting from the FEM

Recall:

- We find the coefficients U_j of the solution

$$u_h(x) = \sum_{j=1}^{N} U_j \varphi_j(x)$$

  by solving a linear system

$$AU = F$$

- The size of the linear system equals the number of coefficients U_j
- There may be many coefficients: 1000s to billions!

On linear systems resulting from the FEM

A few thousand unknowns:

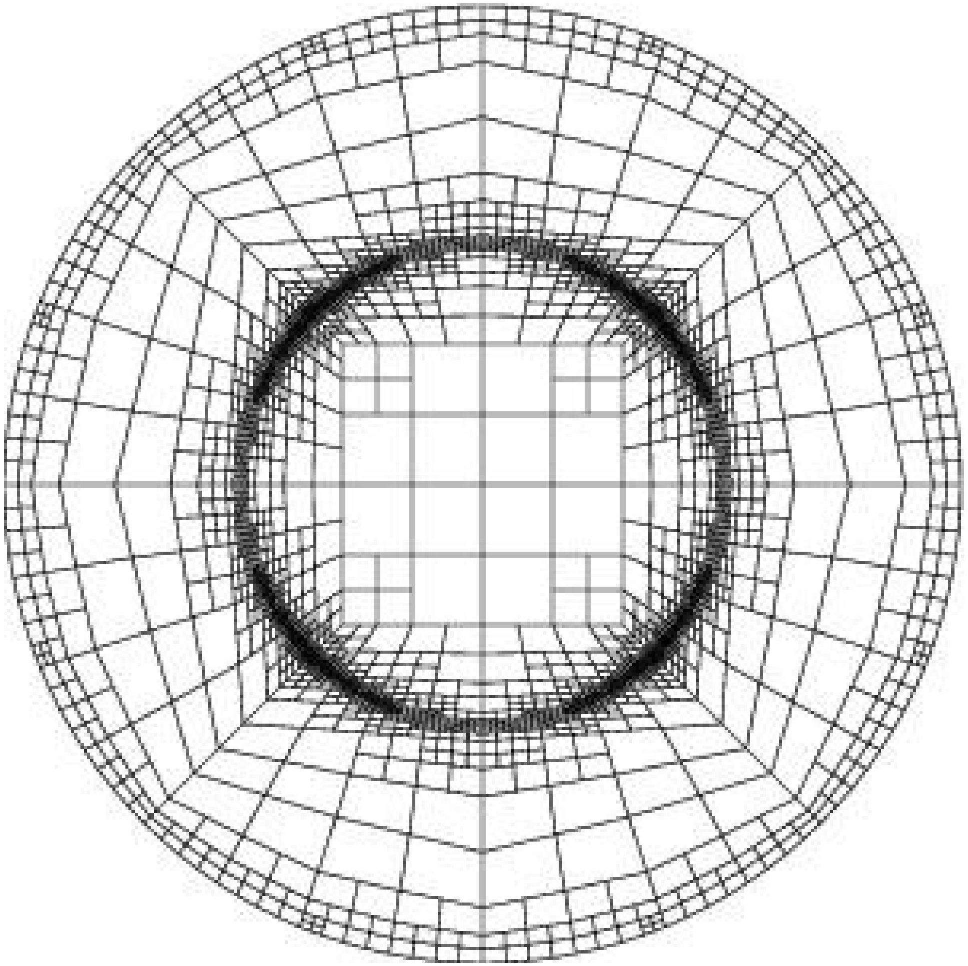

On linear systems resulting from the FEM

A few million unknowns:

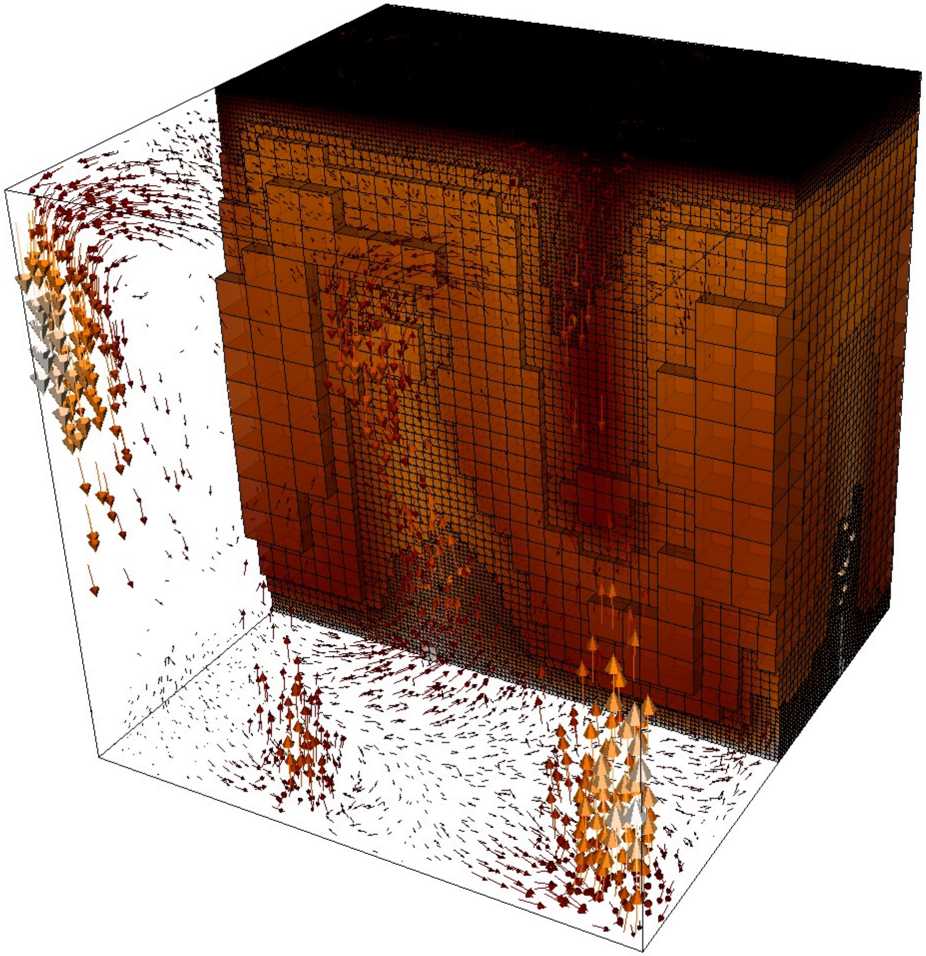

On linear systems resulting from the FEM

A few billion unknowns:

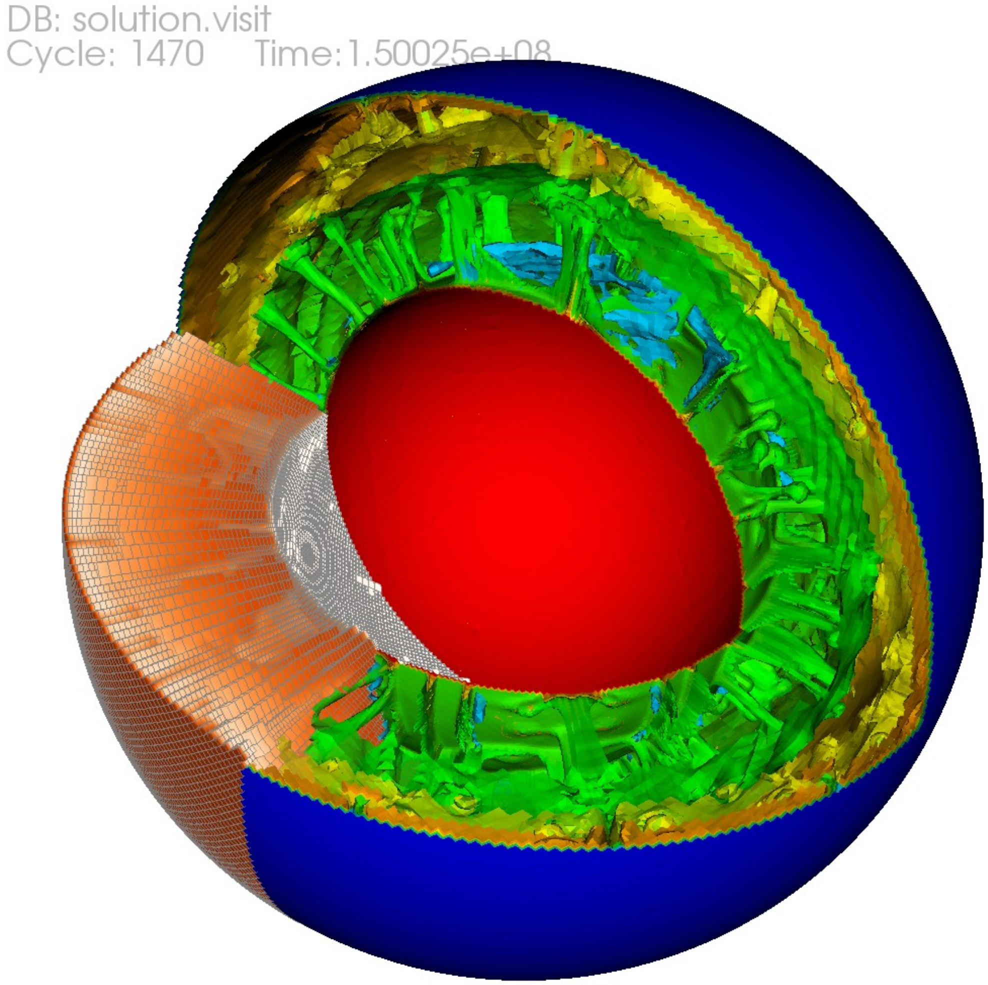

On linear systems resulting from the FEM

Question: How can we even imagine solving linear systems

$$AU = F$$

with millions or billions of unknowns?

Problem 1: With N unknowns, storing everything requires

$M = (N^2+N+N) * 8$ bytes of memory

Examples:

- $N=10^6 \rightarrow M=10{,}000$ GB → maybe possible
- $N=10^9 \rightarrow M=10^{10}$ GB → not possible

On linear systems resulting from the FEM

Question: How can we even imagine solving linear systems

$$AU = F$$

with millions or billions of unknowns?

Problem 2: With N unknowns, Gauss elimination takes

$C = 2/3\ N^3$ operations

Examples:

- $N=10^6 \rightarrow C=10^{18}$ operations = *30* years @ 10^9 ops/sec
- $N=10^9 \rightarrow C=10^{27}$ operations = irrelevant

On linear systems resulting from the FEM

Question: How can we even imagine solving linear systems

$$AU = F$$

with millions or billions of unknowns?

Answers:

- We *can't* solve general linear system of these sizes
- The FEM must be producing linear systems of a special kind that makes this feasible
- *How* we solve these linear systems: Lectures 34-38

On linear systems resulting from the FEM

Summary:

The FEM must be producing linear systems

$$AU = F$$

of a special kind that makes their solution feasible.

This property is *sparsity*:
Nearly all entries in the
finite element matrix *A* are zero!

This is not an accident: It is a design criterion of the FEM.
It is what makes the method successful!

The basis functions of the FEM

Recall: We chose the basis functions φ_j so that they are 1 at one of the nodes and 0 at all of the others.

Example for a 1d mesh:

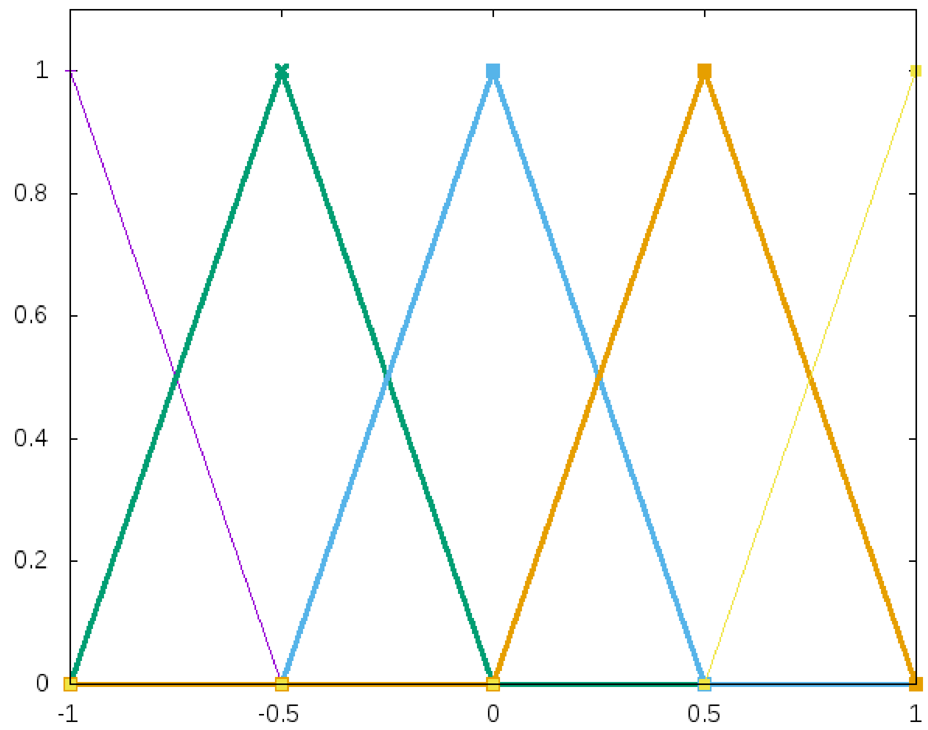

The basis functions of the FEM

Recall: We chose the basis functions φ_j so that they are 1 at one of the nodes and 0 at all of the others.

Example for a triangular 2d mesh:

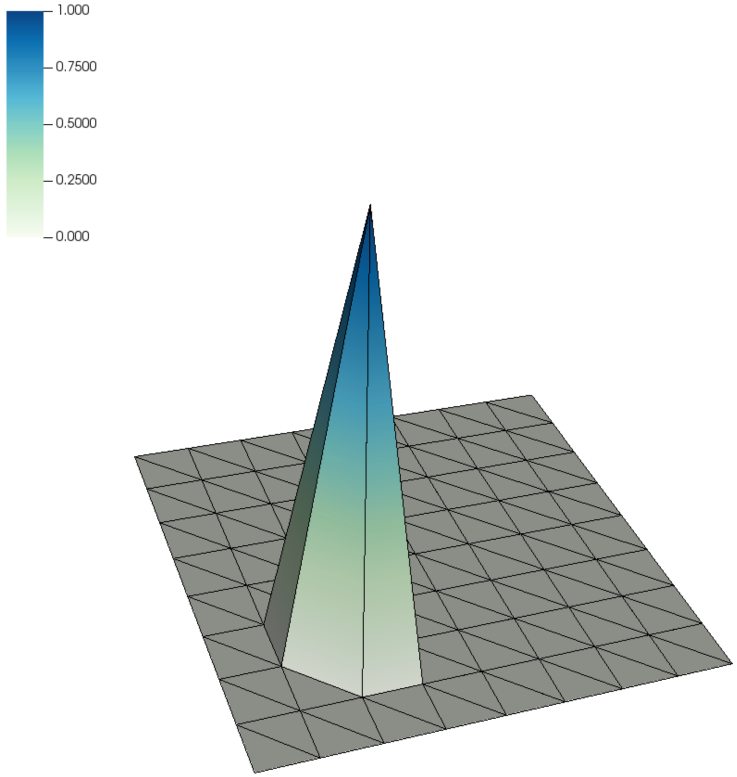

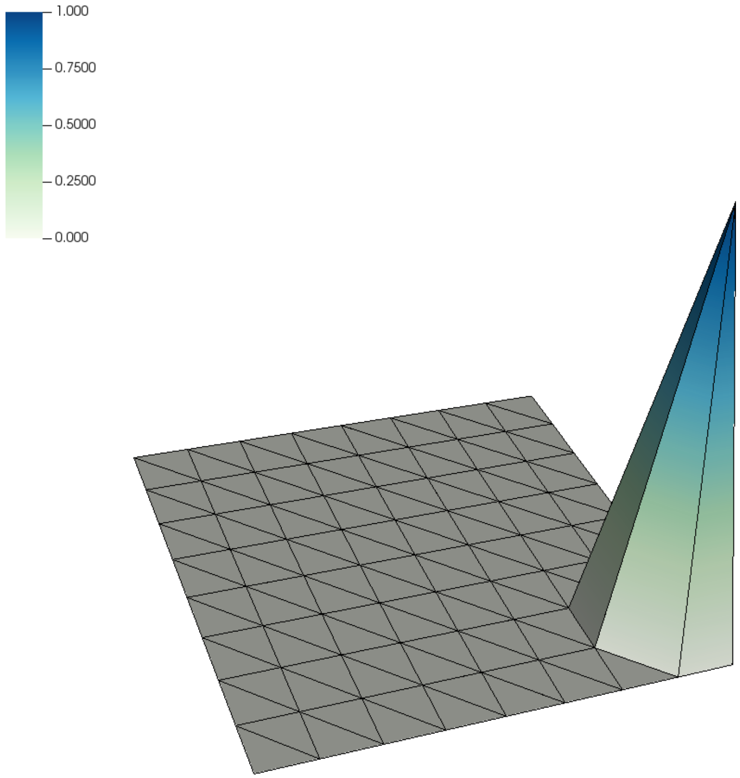

The basis functions of the FEM

Recall: We chose the basis functions φ_j so that they are 1 at one of the nodes and 0 at all of the others.

Example for a quadrilateral 2d mesh:

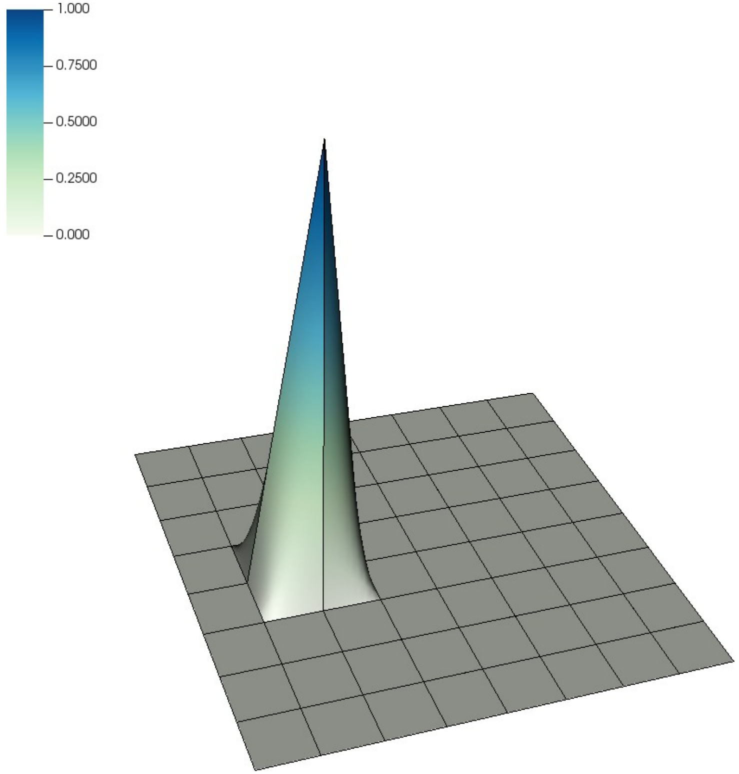

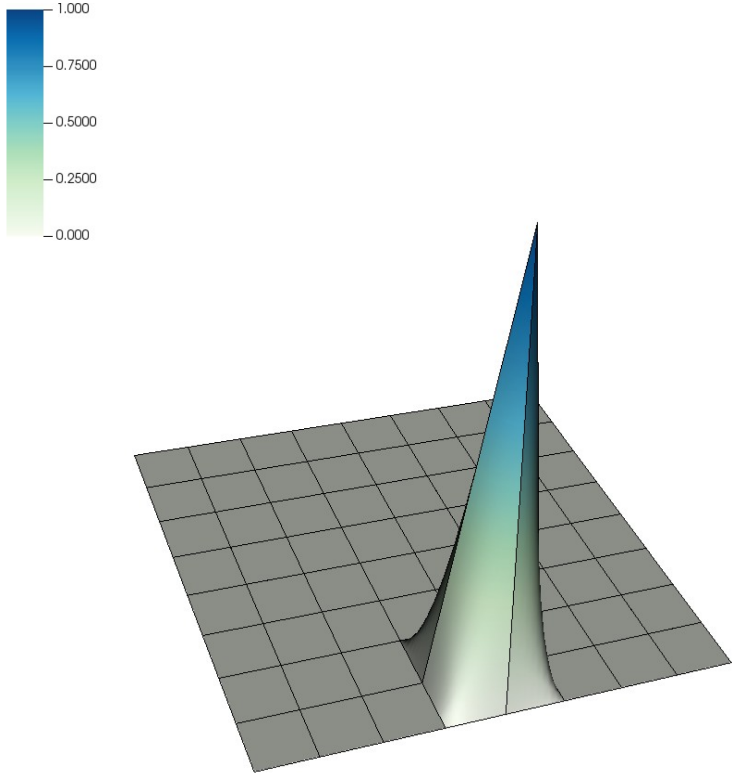

The entries of the matrix *A*

Also recall:

For the linear system corresponding to the Laplace equation,

$$AU = F$$

the matrix entries are defined by

$$A_{ij} = \int_\Omega \nabla \varphi_i(x) \cdot \nabla \varphi_j(x) \; dx$$

Important: A_{ij} is only nonzero if shape functions φ_i and φ_j are nonzero in regions that *overlap*!

This is only true if φ_i and φ_j are defined at vertices that are part of a common cell.

The entries of the matrix *A*

Example:

Assume that these are φ_{13} and φ_{42}:

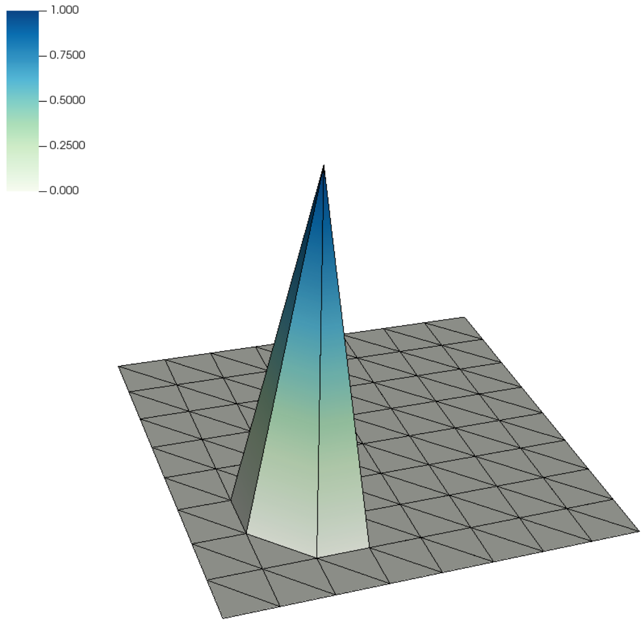

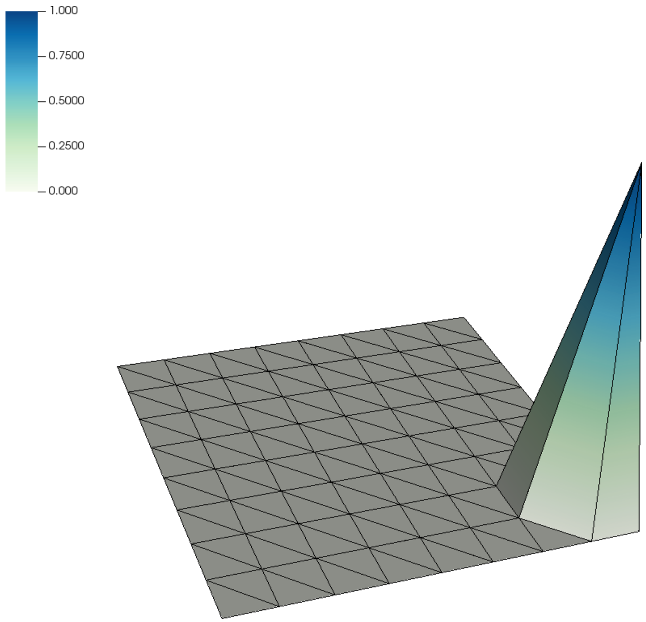

Then: $A_{13,42}=\int_\Omega \nabla\varphi_{13}(x)\cdot\nabla\varphi_{42}(x)\ dx \ = \ 0$

The entries of the matrix *A*

More specifically, for triangles:

Assume that this is φ_{13}:

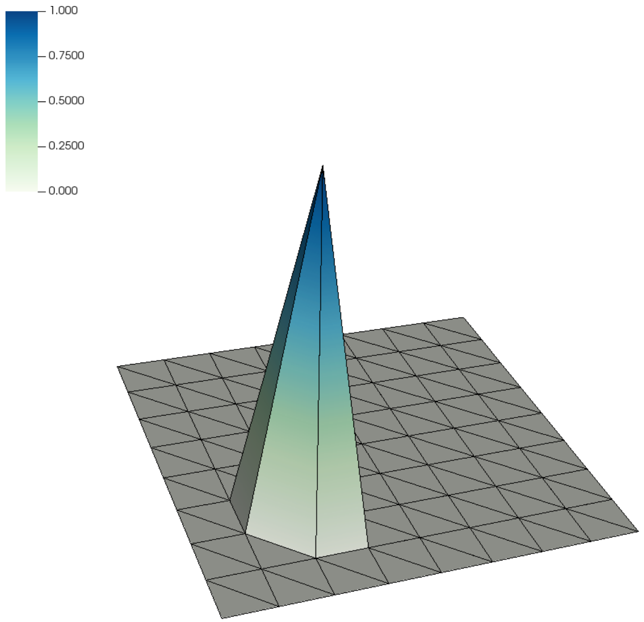

Then: $A_{13,j} \neq 0$ only if $j=13$ or if j is one of 6 adjacent vertices

→ At most 7 nonzero entries per row of A!

The entries of the matrix *A*

More specifically, for quadrilaterals:

Assume that this is φ_{13}:

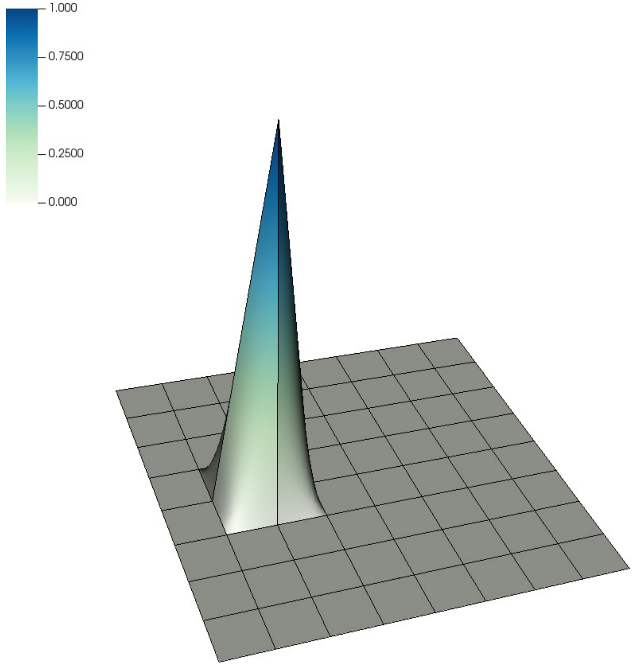

Then: $A_{13,j} \neq 0$ only if $j=13$ or if j is one of 8 adjacent vertices

→ At most 9 nonzero entries per row of *A*!

The entries of the matrix A

Finite element matrices are "sparse":

- The number of entries per row is always $\leq m$
- m depends on
  - the equation (i.e., weak form)
  - the polynomial degree of the shape functions
  - the dimension of the domain
- Typical values:
  - 2d Laplace, triangles, piecewise linears: $m=7$
  - 3d Stokes, hexahedra, Taylor-Hood elements: $m\approx 400$
- **But:** m does not depend on the number of unknowns N!

The entries of the matrix *A*

Finite element matrices are "sparse"! Examples of the "sparsity patterns" of matrices:

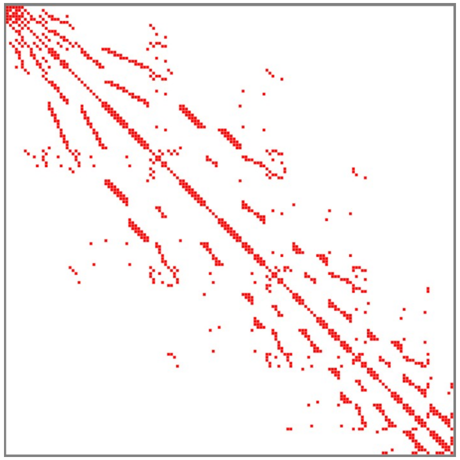

(from step-2)

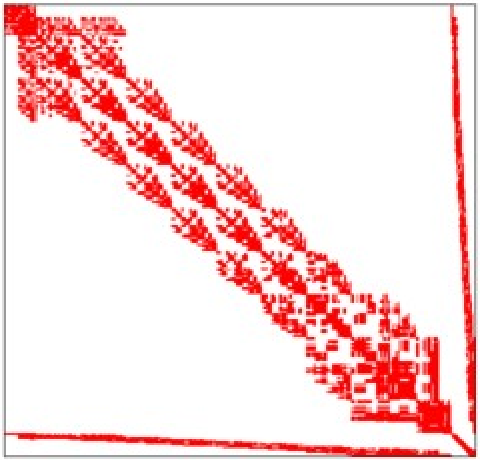

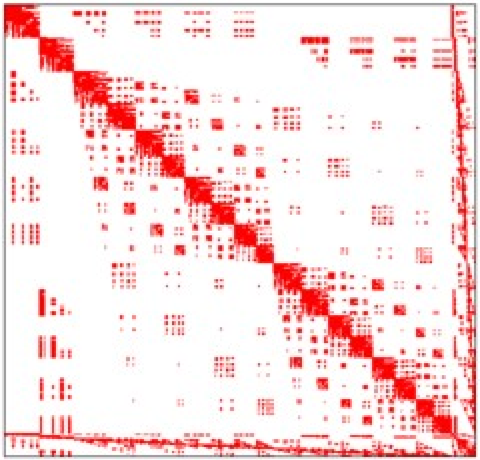

(from the DoFRenumbering namespace)

Note: The form of the sparsity pattern depends on how we enumerate our shape functions. But *m* does not!

The entries of the matrix *A*

Finite element matrices are "sparse"!

Consequence:

- Storing A requires at most mN memory locations, rather than N^2
- Matrix-vector product with A requires at most mN operations, rather than N^2

- There are algorithms that solve linear systems $AU=F$ using at most N matrix-vector products

There is hope for storing and solving even very large problems!

Finite element methods in scientific computing

Wolfgang Bangerth, Colorado State University