# ElmerSolver Input File (SIF) Explained 

Elmer Team

CSC - IT Center for Science Ltd.

## Contents

- Elmer Modules
- Syntax of SIF
- Parameters, etc.
- Sections of SIF:
- Header
- Constants
- Simulation
- Solver
- Body
- Equation
- Body Force - Material
- Initial Condition
- Boundary Condition
- Tables and Arrays
- MATC
- User Defined Functions


## Elmer - Modules



## Elmer - Modules



## Sections of SIF

- The SIF is structured into sections
oHeader
oConstants
oSimulation
oSolver
oBody
oEquation
oBody Force
-Material
oInitial Condition
oBoundary Condition

The contents of each section is between the keyword above and an End-statement

## Sections of SIF: Header

- Declares search paths for main directories

```
Header
    Mesh DB "." "dirname"
    Include Path "includename"
    Results Directory "resultdir"
End
```

- preceding path + directory name of mesh database

Mesh directory under dirname

- Shared objects, etc. under includename
- different output directory resultdir By default under mesh-directory


## Sections of SIF: Constants

- Declares simulation-wide constants

```
Constants
    Gas Constant = Real 8.314E00
    Gravity(4) = 0 -1 0 9.81
```

- a casted scalar constant
- Gravity vector, an array with a registered name (special setup for certain solvers)


## Sections of SIF: Simulation

- Declares details of the simulation:

```
Simulation
Coordinate System = "Cartesian 2D"
Coordinate Mapping(3) = Integer 1 2 3
Coordinate Scaling(3) = Real 1.0 1.0 0.001
Simulation Type = "Transient"
Output Intervals(2) = 10 1
```

- choices: Cartesian\{1D,2D,3D\}, Polar\{2D,3D\}, Cylindric, Cylindric Symmetric, Axi Symmetric
- Permute, if you want to interchange directions in mesh
- That would scale the $3^{\text {rd }}$ direction by 1/1000
- Steady State, Transient or Scanning
- Interval of results being written to disk


## Sections of SIF: Simulation

- Declares details of the simulation:

```
Steady State Max Iterations = 10
Steady State Min Iterations = 2
Timestepping Method = "BDF"
BDF Order = 1
Timestep Intervals(2) = 10 100
Timestep Sizes(2) = 0.1 1.0
Output File = "name.result"
Post File = "name.vtu"
```

- How many min/max rounds on one timelevel/in a steady state simulation (see later)
- Choices: BDF, Newmark or Crank-Nicholson
- This would be implicit Euler
- Has to match array dimension of Timestep Sizes
- The length of one time step
- Contains data for restarting
- Contains output data for ParaView (vtu)
- alternatively, suffix.ep would produce ElmerPost legacy output


## Sections of SIF: Simulation

- Declares details of the simulation:

```
Restart File = "previous.result"
Restart Position = 10
Restart Time = 100
Initialize Dirichlet Condition = False
Restart Before Initial Conditions = True
```

Max Output Level $=5$
End

- Restart from this file at file-entry (not necessarilly timestep!) no. 10 and set time to 100 time-units
- Default is True. If false, Dirichlet conditions are called at Solver execution and not at beginning
- Default is False. IfTrue, then Initial Condition can overwrite previous results
- Level of verbosity:

```
1 = errors
    3 = warnings
    4 = default
10 = all (sometimes too much) information
```


## Sections of SIF: Solver

- Declares a physical model to be solved
Solver 3
Equation = "Navier-Stokes"
Exec Solver = "Always"
Linear System Solver = "Iterative"
Linear System Iterative Method = "BiCGStab"
Linear System Convergence Tolerance = 1.0e-6
Linear System Abort Not Converged = True
Linear System Preconditioning = "ILU2"
- Numbering from 1 (priority)
- The name of the equation
- Always (default), Before/After Simulation/Timestep/Saving
- Choices: Iterative, Direct, MultiGrid
- Lots of choices here, if
- Convergence criterion
- If not True (default) continues simulation in any case
- Pre-conditioning method


## Sections of SIF: Solver

## - Declares a physical model to be solved

| Nonlinear System Convergence Tolerance $=1.0 \mathrm{e}-5$ | - Convergence criterion for non-linear problem |
| :---: | :---: |
| Nonlinear System Max Iterations $=20$ | - The maximum rounds |
| Nonlinear System Min Iterations $=1$ | - The minimum rounds |
| Nonlinear System Newton After Iterations=10 | - Switch from fixed-point to Newton scheme after 10 iterations ... |
| Nonlinear System Newton AfterTolerance=1.0e-3 | - ... or after this criterion (NV.: has to be smaller than convergence criterion ot hit) |
| Steady State Convergence Tolerance $=1.0 \mathrm{e}-3$ | - The convergence on the time-level |
| Stabilization Method = Stabilized | - advection needs stabilization. Alternatives: Bubbles, VMS, P2/P1 |

## Sections of SIF: Solver



1. Timestep Intervals
2. Steady State Max Iterations
3. Nonlinear Max Iterations
4. Linear System Max Iterations
5. Linear System Convergence Tolerance $\epsilon_{l}$
6. Nonlinear System Convergence Tolerance
$\epsilon_{n}$
7. Steady State Convergence Tolerance $\quad \epsilon_{S}$
8. 

## Sections of SIF: Solver

## Before Simulation



## Sections of SIF: Body

- Declares a physical model to be solved Body 2

Name = "pipe"
Equation $=2$
Material $=2$
Body Force $=1$
Initial Condition $=2$

- Numbering from 1 to number of bodies
- Identifier of the body
- The assigned set of equations
- The assigned material section
- The assigned body force
- The assigned initial condition


## End

## Sections of SIF: Body

- Each Body has to have an Equation and Material assigned
- Body Force, Initial Condition optional
- Two bodies can have the same

Material/Equation/
Body Force/Initial Condition section assigned


## Sections of SIF: Equation

- Declares set of solvers for a body


## Equation 2

Active Solvers(2) = 13

Convection $=$ Computed

- Numbering from 1 to number of equation sets
- Declares the solvers (according to their numbers) to be solved within this set
- Important switch to account for convection term. Alternatives: None and Constant (needs
Convection Velocity to be declared in the Material section)


## Sections of SIF: Body Force

- Declares body forces and bulk and execution conditions for a body

Body Force 3
Flow Body Force $1=0.0$
Flow Body Force $2=-9.81$

MyVariable $=$ Real 0.0

Heat Source $=1.0$
End

- Numbering from 1 to number of body forces
- Gravity pointing in negative $x$-direction applied to NavierStokes solver
- A Dirichlet condition for a variable set within the body
- Heat source for the heat equation


## Sections of SIF: Material

- Declares set of material parameters for body

```
Material 1
    Density = 1000.0
    Heat Conductivity (3,3) = 1 0 0\
                                    O 1 0\
                                    0 0 2
    Viscosity = Variable Temperature
        Real MATC "viscosity(tx)"
    Heat Capacity = Variable Temperature
    Procedure "filename" "functionname"
```

    MyMaterialParameter \(=\) Real 0.0
    End
Elmer course CSC, May 2018

End

- Numbering from 1 to number of material
- Always declare a density
- Parameters can be arrays
- Or MATC functions of other variables
- Or Fortran functions with/without dependency on input variables
- Non-keyword DB parameters have to be casted


## Sections of SIF: Initial Condition

- Declares initial conditions for a body (by default restart values are used)

```
Initial Condition 2
    Velocity 1 = Variable Coordinate 2
        Real MATC "42.0*(1.0 - tx/100.0)"
    Velocity 2 = 0.0
    Velocity 3 = Variable Coordinate 3
        Procedure "filename" "functionname"
    MyVariable = Real 20.0
End
```

- Numbering from 1 to number of IC's
- Initial condition as a MATC function of a variable ..
- ... and as a constant value
- ... and as a user function
- Non-keyword DB parameters have to be casted


## Sections of SIF: Boundary Condition

## - Declares conditions at certain boundaries

```
Boundary Condition 3
    Target Boundaries(2) = 14
Velocity 1 = Variable Coordinate 2
        Real MATC "42.0*(1.0 - tx/100.0)"
Velocity 2 = 0.0
Velocity 3 = Variable Coordinate 3
        Procedure "filename" "functionname"
Normal-Tangential Velocity = Logical True
End
```

- Numbering from 1 to number of BC's
- The assigned mesh boundaries
- Variable as a MATC function and ...
... as a constant
... as a user function
- Set velocities in normal-tangential system

```
End
```


## Tables and Arrays

- Tables (piecewise linear or cubic):
- Arrays:

Expresions:

```
Density = Variable Temperature
Real cubic
        0 900
    273 1000
    300 1020
    4 0 0 1 0 0 0
End
```

Target Boundaries (3) = 5710
MyParamterArray (3,2) = Real 1 2\}
3 4
56
OneThird = Real \$1.0/3.0

## MATC

- Syntax close to C
- Even if-conditions and loops
- Can be use for on-the-fly functions inside the SIF
- Documentation on web-pages
- Do not use with simple numeric expressions:

> OneThird = Real \$1.0/3.0
is much faster than
OneThird = Real MATC "1.0/3.0"

## MATC

- Use directly in section:

```
Heat Capacity = Variable Temperature
    Real MATC "2.1275E3 + 7.253EO*(tx - 273.16)"
```

- Even with more than one dependency:

Temp $=$ Variable Latitude, Coordinate 3
Real MATC "49.13 + 273.16-0.7576*tx(0)-7.992E-03*tx(1)"

- Or declare functions (somewhere in SIF, outside a section)

```
$ function stemp(X) {\
    _stemp = 49.13 + 273.16-0.7576*X(0) - 7.992E-03*X(1)\
```

\}
being called by:

# Temp = Variable Latitude, Coordinate 3 <br> Real MATC "stemp(tx)" 

## User Defined Functions (UDF)

- Written in Fortran 90
- Dynamically linked to Elmer
- Faster, if more complicated computations involved
- Compilation command elmerf90

```
$ elmerf90 myUDF.f90 -o myUDF.so
```

- Call from within section:

MyVariable = Variable Temperature
Real Procedure "myUDF.so" "myRoutine"

## User Defined Functions (UDF)

- Example: $\quad \rho(T[K])=1000.0 \cdot\left[1-1 \times 10^{-4} \cdot(T-273.15)\right]$

```
FUNCTION getdensity( Model, N, T ) RESULT(dens)
    USE DefUtils !important definitions
    IMPLICIT None
    TYPE (Model_t) : : Model
    INTEGER :: N
    REAL (KIND=dp) :: T, dens
    dens = 1000.0_dp*(1.0_dp - 1.0d-04*(T - 273.0_dp))
END FUNCTION getdensity
```

- Definitions loaded from DefUtils
- Header: Model access-point to all ElmerSolver inside data; Node number $\mathbf{N}$; input value $\mathbf{T}$



## Elmer

# Software Development Practices APIs for Solver and UDF 

ElmerTeam<br>CSC - IT Center for Science, Finland CSC, 2018

## Elmer programming languages

- Fortrango (and newer)
- ElmerSolver (~"300,000 lines of which ~50\% in DLLs)
- C++
- ElmerGUI (~18,000 lines)
- ElmerSolver (~15,000 lines)
- C
- ElmerGrid (~30,000 lines)
- MATC (~11,000 lines)
- ElmerPost (~45,000 lines)


## Tools for Elmer development

- Programming languages
- Fortrango (and newer), C, C++
- Compilation
- Compiler (e.g. gnu), configure, automake, make, (cmake)
- Editing
o emacs, vi, notepad++,...
- Code hosting (git)
- https://github.com/ElmerCSC
- Consistency tests
- Currently around 450
- Code documentation
- Doxygen


## Elmer libraries

- ElmerSolver
- Required: Matc, Hutlter, Lapack, Blas, Umfpack (GPL)
-Optional: Arpack, Mumps, Hypre, Pardiso, Trilinos, SuperLU, Cholmod, NetCDF, HDF5, ...
- ElmerGUI
o Required: Ot, ElmerGrid, Netgen
-Optional:Tetgen, OpenCASCADE, VTK, OVT


## Elmer licenses

- ElmerSolver library is published under LGPL
o Enables linking with all license types
olt is possible to make a new solver even under proprierity license
- Note: some optional libraries may constrain this freedom due to use of GPL licences
- Most other parts of Elmer published under GPL
- Derived work must also be under same license ("copyleft")
- Proprierity modules linked with ElmerSolver may be freely licensed if they are not derived work
- Note that you must not violete licences of other libraries


## Elmer version control at GitHub

- In 2015 the official version control of Elmer was
transferred from svn at sf.net to git hosted at GitHub
- Git offers more flexibility over svn
- Distributed version control system
- Easier to maintain several development branches
- More options and hence also steeper learning curve
o Developed by Linus Torvalds to host Linux kernel development
- GitHub is a portal providing Git and some additional servives
- Management of user rights
- Controlling pull requests


## Directory listinq of elmerfem/trunk with TortoiseGIT:

| Name | Date modified | Type | Size |  |
| :---: | :---: | :---: | :---: | :---: |
| (2. buildtools | 3.11.2016 11:56 | File folder |  |  |
| d. cmake | 3.11.2016 11:56 | File folder |  |  |
| d. cpack | 3.11.2016 11:56 | File folder |  |  |
| ( eio | 3.11.2016 11:56 | File folder |  |  |
| (8. elmergrid | 3.11.2016 11:56 | File folder |  |  |
| (8) ElmerGUI | 3.11.2016 11:56 | File folder |  | ElmerGrid mesh manipulation <br> ElmerGUI graphical user interface |
| ( ElmerGUIlogger | 3.11.2016 11:56 | File folder |  |  |
| d. ElmerGUItester | 3.11.2016 11:56 | File folder |  |  |
| d. elmerice | 3.11.2016 11:56 | File folder |  |  |
| d elmerparam | 3.11.2016 11:56 | File folder |  | Elmer/ICE community developments <br> ElmerParam optimization module |
| ( fem | 3.11.2016 11:57 | File folder |  | ElmerSolver library and modules |
| d. fhutiter | 3.11.2016 11:57 | File folder |  |  |
| (6) front | 3.11.2016 11:57 | File folder |  | HUTiter Krylov methods library |
| (d) hutiter | 3.11.2016 11:57 | File folder |  | cerFront: Initial user interface (obsolite) |
| d. license_texts | 3.11.2016 11:57 | File folder |  |  |
| (d) matc | 3.11.2016 11:57 | File folder |  |  |
| d mathlibs | 3.11.2016 11:57 | File folder |  | MATC library |
| (2. meshgen2d | 3.11.2016 11:57 | File folder |  | Mesh2D (Delaunay triangularization,obsolite) |
| d. misc | 3.11.2016 11:57 | File folder |  |  |
| d. post | 3.11.2016 11:57 | File folder |  |  |
| d. umfpack | 3.11.2016 11:57 | File folder |  | ElmerPost: Initial postprocessor (obsolite) |
| d. utils | 3.11.2016 11:57 | File folder |  |  |
| 0 | 3.11.2016 11:56 | Text Document | 1 KB |  |
| Q CMakeLists | 3.11.2016 11:56 | Text Document | 13 KB |  |
| - README | 3.11.2016 11:56 | Text Document | 2 KB |  |

## Cmake build system

- During 2014-2015 Elmer was migrated from gnu autotools into cmake
- Cmake offers several advantages
- Enables cross compilation for diffirent platforms
(e.g. Intel MICs)
- More standardizes installation scripts
- Straight-forward package creation for many systems
(using cpack)
- Great testing utility with ctest - now also in parallel
- Transition to cmake required significant code changes
- ISO C-bindings \& many changes in APIs
- Backward compatibility in compilation lost


## Compiling fresh Elmer source from GitHub

\# clone the git repository.
\$ git clone https://www.github.com/ElmerCSC/elmerfem
\# Switch to devel branch (currently the default branch)
\$ cd elmerfem
\$ git checkout devel
\$cd..
\# create build directory
\$ mkdir build
\$cd build

```
$ cmake -DWITH_ELMERGUI:BOOL=FALSE -
DWITH_MPI:BOOL=FALSE-
DCMAKE_INSTALL_PREFIX=../install ../elmerfem
```

\$ cmake <flags>
\# You can tune the compilation parameters graphically with \$ ccmake or \$cmake-gui .
\$ make install
\# or alternatively compile in parallel (4 procs) \$ make -j4 install

## Consistency tests

- Utilize ctest system to run a set of Elmer cases
- Upon success each case writes 1 to file TEST.PASSED, and on failure o, respectively
- There are more than 58 o consistency tests (May 2018)
- Located under fem/tests
- Each time a significant commit is made the tests are run with the fresh version
- Aim: even devel version is a stable
o New tests for each major new feature
- The consistency tests provide a good starting point for taking some Solver into use
o cut-paste from sif file


## Executing the consistency tests of Elmer

```
>ctest -j4 -LE elmerice
    Start 143: mgdyn_torus_harmonic
        Start 304: ThermalActuator
        Start 344: RotatingBCMagnetoDynamicsGeneric
    1/310 Test #344: RotatingBCMagnetoDynamicsGeneric ... Passed 43.18 sec
        Start 293: mgdyn lamstack lowfreq harmonic
    2/310 Test #304: ThermalActuator ................... Passed 59.78 sec
        Start 222: mgdyn_transient_loss
    3/310 Test #293: mgdyn_lamstack_lowfreq_harmonic .... Passed 21.80 sec
        Start 322: mgdyn_bh
308/310 Test \#46: CoupledPoisson7
```

Passed
Passed

Passed
0.38 sec
0.38 sec
6.34 sec

```
100% tests passed, 0 tests failed out of 310
Total Test time (real) = 365.62 sec
```


## Doxygen - WWW documentation



## Doxygen - Example in code

## - Special comment indicators: !> and <!

```
!------------------------------------------------------------------------
!> For example, one may compute the the heat flux as the negative gradient of temperature
!> field multiplied by the heat conductivity.
!> \ingroup Solvers
!----------------------------------------------------------------------------------------------------------
SUBROUTINE FluxSolver( Model,Solver,dt,Transient )
USE CoordinateSystems
USE DefUtils
IMPLICIT NONE
```



```
REAL(KIND=dp) :: dt !< Timestep size for time dependent simulations
LOGICAL :: Transient !< Steady state or transient simulation
```



```
! Local variables
TYPE(ValueList t),POINTER :: SolverParams
```


## Doxygen - Example in WWW

subroutine FluxSolver ( TYPE(Model_t) Model,

$$
\begin{array}{ll}
\text { TYPE(Solver_t) } & \text { Solver, } \\
\text { REAL(KIND=dp) } & \mathrm{dt}, \\
\text { LOGICAL } & \text { Transient }
\end{array}
$$

)

Subroutine for computing fluxes and gradients of scalar fields. For example, one may compute the the heat flux as the negative gradient of temperature field multiplied by the heat conductivity.

## Parameters:

Solver Linear \& nonlinear equation solver options
Model All model information (mesh, materials, BCs, etc...)
dt Timestep size for time dependent simulations
Transient Steady state or transient simulation
References BulkAssembly().
Here is the call graph for this function:


## Installers

- Fresh Windows installers
o Currently only 64 bit version
- Also a parallel version with msmpi
- http://www.nic.funet.fi/pub/sci/physics/elmer/bin/ windows/
- Some version available also at sf.net
- Elmer for Debian \& Ubuntu etc. at launchpad
- Nightly builds from Git repository
- To install
\$ sudo apt-add-repository ppa:elmer-csc-ubuntu/elmer-csc-ppa

\$ sudo apt-get update
\$ sudo apt-get install elmerfem-csc


## Compilation of a DLL module

- Applies both to Solvers and User Defined Functions (UDF)
- Assumes that there is a working compile environment that provides "elmerf90" script
- Comes with the Windows installer, and Linux packages
- Generated automatically when ElmerSolver is compiled
elmerf90 MySolver. F90 -o MySolver.so


## User defined function API

```
!-----------------------------------------------------------
!> Standard API for UDF
!-------------------------------------------------------------
FUNCTION MyProperty( Model, n, t ) RESULT(f)
!------------------------------------------------------------
    USE DefUtils
    IMPLICIT NONE
```

```
!-----------------------------------------------------------
```

!-----------------------------------------------------------
TYPE(Model_t) :: Model !< Handle to all data
TYPE(Model_t) :: Model !< Handle to all data
INTEGER :: n !< Current node
INTEGER :: n !< Current node
REAL(KIND=dp) :: t !< Parameter(s)
REAL(KIND=dp) :: t !< Parameter(s)
REAL(KIND=dp) :: f !< Parameter value at node
REAL(KIND=dp) :: f !< Parameter value at node
!-----------------------------------------------------------
!-----------------------------------------------------------
Actual code

```
    Actual code
```


## Function API

MyProperty $=$ Variable time
"MyModule" "MyProperty"

- User defined function (UDF) typically returns a real valued property at a given point
- It can be located in any section that is used to fetch these values from a list
- Boundary Condition, Initial Condition, Material,...


## Solver API

```
!--------------------------------------------------------------
!> Standard API for Solver
!------------------------------------------------------------
SUBROUTINE MySolver( Model,Solver,dt,Transient )
!-----------------------------------------------------------
    USE DefUtils
    IMPLICIT NONE
!-----------------------------------------------------------
    TYPE(Solver_t) :: Solver !< Current solver
    TYPE(Model_t) :: Model !< Handle to all data
    REAL(KIND=dp) :: dt !< Timestep size
    LOGICAL :: Transient !< Time-dependent or not
!------------------------------------------------------------
    Actual code
```


## Solver API

Solver 1
Equation = "MySolver"
Procedure = "MyModule" "MySolver"
End

- Solver is typically a FEM implementation of a physical equation
- But it could also be an auxiliary solver that does something completely different
- Solver is usually called once for each coupled system iteration


## Elmer - High level abstractions

- The quite good success of Elmer as a multiphysics code may be addressed to certain design choices
- Solver is an asbtract dynamically loaded object
- Parameter value is an abstract property fecthed from a list
- The abstractions mean that new solvers may be implemented without much need to touch the main library
-Minimizes need of central planning
oSeveral applications fields may live their life quite independently (electromagnetics vs. glaceology)
- MATC - a poor man's Matlab adds to flexibility as algebraic expressions may be evalueted on-the-fly


## Solver as an abstract object

- Solver is an dynamically loaded object (.dll or .so)
- May be developed and compiled seperately
- Solver utilizes heavily common library utilities
- Most common ones have interfaces in DefUtils
- Any solver has a handle to all of the data
- Typically a solver solves a weak form of a differential equation
- Currently ~6o different Solvers, roughly half presenting physical phenomena
- No upper limit to the number of Solvers
-Often cases include ~10 solvers
- Solvers may be active in different domains, and even meshes
- The menu structure of each solver in ElmerGUI may be defined by an .xml file


## Property as an abstract object



- Properties are saved in a list structure by their name
- Namespace of properties is not fixed, they may be introduced in the command file
o E.g. "MyProperty = Real 1.23"adds a property "MyProperty" to a list structure related to the solver block
- In code parameters are fetched from the list
o E.g. "val = GetReal( Material,'MyProperty', Found)" retrieves the above value 1.23 from the list
- A "Real" property may be any of the following
- Constant value
- Linear or cubic dependence via table of values
- Expression given by MATC (MatLab-type command language)
- User defined functions with arbitrary dependencies
- Real vector or tensor
- As a result solvers may be weakly coupled without any a priori defined manner
- There is a price to pay for the generic approach but usually it is less than $10 \%$
- SOLVER. KEYWORDS file may be used to give the types for the keywords in the command file


## Code structure

- Elmer code structure has evolved over the years
- There has been no major restructuring operations
- Ufortunately there is no optimal hierarchy and the number of subroutines is rather large
- ElmerSolver library consists of more than $\sim 40$ modules
- There are all-in-all around 1050 SUBROUTINES and 650 FUNCTIONS (both internal and external)
- To ease the learning curve the most important routines for basic use have been collected into module DefUtils.Fgo


## DefUtils

- DefUtils module includes wrappers to the basic tasks common to standard solvers
oE.g. "DefaultDirichlet()" sets Dirichlet boundary conditions to the given variable of the Solver
oE.g. "DefaultSolve ()" solves linear systems with all available direct, iterative and multilevel solvers, both in serial and parallel
- Programming new Solvers and UDFs may usually be done without knowledge of other modules


## DefUtils - some functions

## Public Member Functions

| TYPE(Solver_t) function, pointer | GetSolver () |
| :---: | :---: |
| TYPE(Matrix_t) function, pointer | GetMatrix (USolver) |
| TYPE(Mesh_t) function, pointer | GetMesh (USolver) |
| TYPE(Element_t) function, pointer | GetCurrentElement (Element) |
| INTEGER function | GetElementIndex (Element) |
| INTEGER function | GetNOFActive (USolver) |
| REAL (KIND = dp) function | GetTime () |
| INTEGER function | GetTimeStep () |
| INTEGER function | GetTimeStepInterval () |
| REAL (KIND = dp) function | GetTimestepSize () |
| REAL (KIND = dp) function | GetAngularFrequency (ValueList, Found) |
| INTEGER function | GetCoupledIter () |
| INTEGER function | GetNonlinIter () |
| INTEGER function | GetNOFBoundaryElements (UMesh) |
| subroutine | GetScalarLocalSolution ( x , name, UElement, USolver, tStep) |
| subroutine | GetVectorLocalSolution ( x , name, UElement, USolver, tStep) |
| INTEGER function | GetNofEigenModes (name, USolver) |
| subroutine | GetScalarLocalEigenmode ( x , name, UElement, USolver, NoEigen, ComplexPart) |
| subroutine | GetVectorLocalEigenmode ( $x$, name, UElement, USolver, NoEigen, ComplexPart) |
| CHARACTER(LEN=MAX_NAME_LEN) function | GetString (List, Name, Found) |
| INTEGER function | GetInteger (List, Name, Found) |
| LOGICAL function | GetLogical (List, Name, Found) |
| recursive REAL (KIND $=$ dp) function | GetConstReal (List, Name, Found, x, y, z) |
| recursive REAL (KIND=dp) function | GetCReal (List, Name, Found) |
| recursive $\operatorname{REAL}($ KIND $=\mathrm{dp}$ ) <br> function, dimension(:), pointer | GetReal (List, Name, Found, UElement) |

## Modules related to linear algebra

BandMatrix.F90<br>BandwidthOptimize.F90<br>BlockSolve.F9o<br>cholmod.c<br>CircuitUtils.Fgo<br>ClusteringMethods.Fgo<br>CRSMatrix.F90<br>DirectSolve.F90<br>EigenSolve.F90<br>IterativeMethods.F90<br>IterSolve.F9o<br>LinearAlgebra.F9o<br>LUDecomposition.F90<br>MGPrec.F90<br>Multigrid.F90<br>Smoothers.F90<br>SolveBand.F90<br>SolveHypre.c<br>SolverUtils.F90<br>SolveSBand.F9o<br>SolveSuperLU.c<br>SolveTrilinos.cxx

## Modules related to space and time discretization

ElementDescription.F90
ElementUtils.F9o
H1ElementBasisFunctions.F90
PElementBase.F90
PElementMaps.F90
TimeIntegrate.F90

## Historical modules including physics

Differentials.F90
DiffuseConvectiveAnisotropic.F9o
DiffuseConvectiveGeneralAnisotropic.F90
ExchangeCorrelations.F9o
MaxwellAxiS.Fgo
Maxwell.F9o
MaxwellGeneral.Fgo
NavierStokesCylindrical.F9o
NavierStokes.F9o
NavierStokesGeneral.F90
Stress.F9o
StressGeneral.F90
VelocityUpdate.F90
Walls.F9o

## Example: Poisson equation

$$
-\nabla^{2} \phi=\rho
$$

- Implemented as an dynamically linked solver
- Available under tests/1dtests
- Compilation by:

Elmerf90 Poisson.F90 -o Poisson.so

- Execution by:

ElmerSolver case.sif

- The example is ready to go massively parallel and with all a plethora of elementtypes in 1D, 2D and 3D


## Poisson equation: code Poisson.Fgo

!> Solve the Poisson equation -|nabla|cdot|nabla $\backslash$ phi $=\mid$ rho
SUBROUTINE PoissonSolver( Model,Solver,dt,TransientSimulation )

## USE DefUtils <br> IMPLICIT NONE

!Initialize the system and do the assembly:
CALL DefaultInitialize()
active $=$ GetNOFActive()
DO $\mathrm{t}=1$, active
Element => GetActiveElement( t )
$\mathrm{n}=$ GetElementNOFNodes()

## LOAD $=0.0$ do

BodyForce => GetBodyForce()
IF (ASSOCIATED(BodyForce) )
Load(1:n) = GetReal( BodyForce, 'Source', Found )
Get element local matrix and rhs vector:

CALL LocalMatrix ( STIFF, FORCE, LOAD, Element, n )
Update global matrix and rhs vector from local contribs
CALL DefaultupdateEquations( STIFF, FORCE )
END DO

CALL DefaultFinishAssembly()
CALL DefaultDirichletBCs()
Norm = DefaultSolve()

CONTAINS

SUBROUTINE LocalMatrix (STIFF, FORCE, LOAD, Element, n )

```
CALL GetElementNodes(Nodes )
    STIFF = 0.odo
    FORCE = 0.0do
    ! Numerical integration:
P = GaussPoints( Element)
DO t=1,IP % n
    ! Basis function values & derivatives at the integration point:
    stat = ElementInfo( Element, Nodes, IP % U(t), IP % V(t), & 
        IP% W(t), detJ, Basis, dBasisdx )
    !The source term at the integration point:
    LoadAtIP = SUM(Basis(1:n) * LOAD(1:n) )
    ! Finally, the elemental matrix & vector:
    STIFF(1:n,1:n) = STIFF(1:n,1:n) +IP % s(t) * DetJ * &
        MATMUL(dBasisdx, TRANSPOSE( dBasisdx ))
    FORCE(1:n) = FORCE(1:n) + IP % s(t) * DetJ * LoadAtIP * Basis(1:n)
END DO
END SUBROUTINE LocalMatrix
END SUBROUTINE PoissonSolver
```


## Poisson equation: command file case.sif

Check Keywords "Warn"

Header
Mesh DB "." "mesh"
End

Simulation
Coordinate System = "Cartesian"
Simulation Type = Steady State
Steady State Max Iterations = 50 End

Body 1
Equation $=1$
Body Force =1
End

Equation 1
Active Solvers(1) = 1
End

## Solver 1

Equation = "Poisson"
Variable = "Potential"
Variable DOFs = 1
Procedure = "Poisson" "PoissonSolver"
Linear System Solver = "Direct"
Linear System Direct Method = umfpack
Steady State Convergence Tolerance $=1 \mathrm{e}-09$

## Poisson equation: source term, examples

Constant source:

```
Source = 1.0
```

Source dependeing piecewise linear on x :

```
Source = Variable Coordinate 1
```

    Real
        0.00 .0
        1.03 .0
        2.04 .0
    End
    Source depending on $x$ and $y$ :

```
Source = Variable Coordinate
    Real MATC "sin(2*pi*tx(0))*cos(2*pi(tx(1))"
```

Source depending on anything

Source = Variable Coordinate 1
Procedure "Source" "MySource"

## Poisson equation: ElmerGUI menus

```
<?xml version='1.0' encoding='UTF-8'?>
<!DOCTYPE edf>
<edf version="1.0" >
<PDE Name="Poisson" >
<Name>Poisson</Name>
```


## <BodyForce>

```
<Parameter Widget="Label" > <Name> Properties </Name> </Parameter>
<Parameter Widget="Edit" >
<Name> Source </Name>
<Type> String </Type>
<Whatis> Give the source term. </Whatis>
</Parameter>
</BodyForce>
<Solver>
<Parameter Widget="Edit" >
<Name> Procedure </Name>
<DefaultValue> "Poisosn" "PoissonSolver" </DefaultValue>
</Parameter>
<Parameter Widget="Edit">
<Name> Variable </Name>
<DefaultValue> Potential</DefaultValue>
</Parameter>
</Solver>
<BoundaryCondition>
<Parameter Widget="Label" > <Name> Dirichlet conditions </Name> </Parameter>
<Parameter Widget="Edit">
<Name> Potential </Name>
<Whatis> Give potential value for this boundary. </Whatis>
</Parameter>
</BoundaryCondition>

\section*{Development tools for ElmerSolver}
- Basic use
-Editor (emacs, vi, notepad++, jEdit, ...)
o elmerfgo script
- Advanced
o Editor
osvn client
- Compiler suite (gfortran, ifort, pathfgo, pgfgo,...)
- Documentation tools (Doxygen, LaTeX)
- Debugger (gdb)
- Profiling tools
-...

\section*{Elmer - some best practices}
- Use version control when possible
olf the code is left to your own local disk, you might as well not write it at all
- Do not fork! (userbase of 1000's)
- Always make a consistency test for a new feature
- Always be backward compatible
- If not, implement a warning to the code
- Maximize the level of abstraction
- Essential for multiphysics software
-E.g. any number of physical equations, any number of computational meshes, any number of physical or numerical parameters - without the need for recompilation


\title{
Mesh related features in Elmer
}

\author{
ElmerTeam \\ CSC - IT Center for Science, Finland
}

CSC, 2018

\section*{Outline}
- Supported element types
- Shapes
- Basic functions
- Mesh generation within ElmerSolver
- Mesh multiplication
- Mesh extrusion
- Adaptivity - very limited
- Mesh deformation \& movement
- Mesh projectors
- Mapping between meshes
- Mortar finite elements

\section*{ElmerSolver - Finite element shapes}
- All standard shaper of Finite Elements are supported
ooD: point
o1D: segment
o2D: triangles, quadrilaterals
o3D: tetraherdons, wedges, pyramids, hexahedrons
- Meshes may have mixed element types
- There may be also several meshes in same simulation

\section*{ElmerSolver - basis functions}

\section*{- Element families}
oNodal (up to 2-4th degree)
op-elements (up to 10th degree)
oEdge \& face-elements
\[
\begin{aligned}
& \text { oH(div) - often associated with"face" elements) } \\
& \text { oH(curl) - often associated with "edge" elements) }
\end{aligned}
\]
- Formulations
-Galerkin, Discontinuous Galerkin
oStabilization
oResidual free bubbles


\section*{ElmerSolver-internal mesh generation}
- Internal mesh division

○ \(2^{\wedge} D^{\prime} M^{\wedge} n\)-fold problem-size
- Known as "Mesh Multiplication"
- Simple inheritance of mesh grading
- Internal mesh extrusion

- Extruded given number of layers
- Idea is to remove bottle-necks from mesh generation
- These can also be performed on a parallel level

- Limited by generality since the internal meshing
features cannot increase the geometry description

Mesh multiplication example
\begin{tabular}{|l|l|}
\hline \begin{tabular}{l} 
Mesh \\
Levels
\end{tabular} & \begin{tabular}{l} 
Number of \\
Elements
\end{tabular} \\
\hline 1 & 7920 \\
\hline 2 & 63360 \\
\hline 3 & 506880 \\
\hline 4 & 4055040 \\
\hline
\end{tabular}


\section*{Limitations of mesh multiplication}
- Standard mesh multiplication does not increase geometric accuracy
- Polygons retain their shape
- Mesh multiplication could be made to honor boundary shapes but this is not currently done
- Optimal mesh grading difficult to achieve
- The coarsest mesh level does not usually have sufficient information to implement fine level grading


\section*{ElmerSolver - Internal mesh extrusion}
- Start from an initial 2D (1D) mesh and then extrude into 3D (2D)
- Mesh density may be given by arbitrary function
- Implemented also for partitioned meshes
- Extruded lines belong to the same partition by construction!
- There are many problems of practical problems where the mesh extrusion of a initial 2D mesh provides a good solution
-One such field is glasiology where glaciers are thin, yet the 2D approach is not always sufficient in accurary


Extruded Mesh Levels \(=21\) Extruded Mesh Density = Variable Coordinate 1 Real MATC "1+10*tx"

\section*{ElmerSolver - Internal extrusion example}


Design Alvar
Aalto, 1936


2D mesh by Gmsh


3D internally extruded mesh

\section*{Summary: Alternatives for increasing mesh resolution}
- Use of higher order nodal elements
- Elmer supports 2nd to 4th order nodal elements
oUnfortunately not all preprocessing steps are equally well supported for higher order elements
- E.g. Netgen output supported only for linear elements
- Use of hierarhical p-element basis functions
- Support up to 10th degree polynomials
o In practice Element = p:2, or p:3
o Not supported in all Solvers
- Mesh multiplication
- Subdivision of elements by splitting

\section*{ElmerSolver - Mesh deformation}
- Meshes may be internally deformed
- MeshUpdate solver uses linear elasticity equation to deform the mesh
- RigidMeshMapper uses rigid deformations and their smooth transitions to deform the mesh
- Deforming meshes have number of uses
- Deforming structures in multiphysics simultion
- E.g. fluid-structure interaction, ALE
- Rotating \& sliding structures
- Geometry optimization
- Mesh topology remains unchanged

\section*{Mapping \& Projectors}
- Ensuring continuity between conforming and nonconforming meshes
o For boundary and bulk meshes
- On-the-fly interpolation (no matrix created) - Mapping of finite element data
- from mesh to mesh
- From boundary to boundary
- Creation of interpolation and projection matrices
- Strong continuity, interpolation: \(x_{l}=P x_{r}\)
- Weak continuity, Mortar projector: \(Q x_{l}-P x_{r}=0\)


Tie contact in linear elasticity using mortar finite elements

\section*{Example: Mesh utilities applied to rotational problems}
- Rigid body movement may be used to implement rotation
- One of several contact pairs are used to define mortar projectors that ensure continuity of soluton
- Most important application area has been the simulation of electrical machines


\section*{Concluding remarks on internal meshing features}
- Internal meshing features can be used to resolve number of challenges related to meshes
- Accuracy
- I/O bottle-necks
- Continuity requirements
- Multiphysics coupling
- Deforming or moving computational domains

\title{
Post-processing utilities within ElmerSolver
}

\author{
ElmerTeam \\ CSC - IT Center for Science, Finland \\ CSC, 2018
}

\section*{Postprocessing utilities in ElmerSolver}
- Saving data
- FEM data
- Line data
- Scalars data
- Grid data
- Computing data
- Derived fields (gradient, curl, divecgence,...)
- Data reduction \& filtering
- Creating fields of material properties
- The functionality is usually achieved by use of atomic auxialiry solvers

\section*{Computing derived fields}
- Many solvers have internal options or dedicated post-processing solvers for computing derived fields
- E.g. stress fields by the elasticity solvers
- E.g. MagnetoDynamicsCalcFields
- Elmer offers several auxiliary solvers that may be used in a more generic way
- SaveMaterials: makes a material parameter into field variable
o StreamlineSolver: computes the streamlines of 2D flow
\(\bigcirc\) FluxSolver: given potential, computes the flux \(q=-c \nabla \phi\)
- VorticitySolver: computes the vorticity of flow, \(w=\nabla \times \phi\)
- PotentialSolver: given flux, compute the potential \(-c \nabla \phi=q\)
- FilterTimeSeries: compute filtered data from time series (mean, fourier coefficients, ...)

○...

\section*{Derived nodal data}
- By default Elmer operates on distributed fields but sometimes nodal values are of interest
- Multiphysics coupling may also be performed alternatively using nodal values for computing and setting loads
- Elmer computes the nodal loads from \(A x-b\) where \(A\), and \(b\) are saved before boundary conditions are applied
oCalculate Loads = True
- This is the most consistant way of obtaining boundary loads
- Note: the nodal data is really pointwise
o expressed in units N, C, W etc.
(rather than \(\mathrm{N} / \mathrm{m}^{\wedge} 2, \mathrm{C} / \mathrm{m}^{\wedge} 2, W / \mathrm{m}^{\wedge} 2\) etc.)
- For comparison with distributed data divided by the \(\sim\) size of the surface elements

\section*{Derived lower dimensional data}
- Derived boundary data
- SaveLine: Computes fluxes on-the-fly
- Derived lumped (or oD) data
- SaveScalars: Computes a large number of different quantities on-the-fly
- FluidicForce: compute the fluidic force acting on a surface
- ElectricForce: compute the electrostatic froce using the Maxwell stress tensor
- Many solvers compute lumped quantities internally for later use (Capacitance, Lumped spring, ...)

\section*{Exporting FEM data: ResultOutputSolve}
- Currently recommened format is VTU
oXML based unstructured VTK
- Has the most complete set of features
- Old ElmerPost format (with suffix .ep) is becoming obsolite
- Simple way to saveVTU files: Post File = file.vtu
- ResultOutputSolve offers additionally several formats
o vtk: Visualization tookit legacy format
o vtu: Visualization tookit XML format
- Gid: GiD software from CIMNE: http://gid.cimne.upc.es
- Gmsh: Gmsh software: http://www.geuz.org/gmsh
-Dx: OpenDx software

\section*{Exporting 2D/3D data: ResultOutputSolve}

An example shows how to save data in unstructured XML VTK (.vtu) files to directory "results" in single precision binary format.
```

Solver n
Exec Solver = after timestep
Equation = "result output"
Procedure = "ResultOutputSolve""ResultOutputSolver"
Output File Name = "case"
Output Format = String "vtu"
Binary Output = True
Single Precision = True
End

```

\section*{Saving 1D data: SaveLine}
- Lines of interest may be defined on-the-fly
- Data can either be saved in uniform 1D grid, or where element faces and lines intersect
- Flux computation using integration points on the boundary not the most accurate
- By default saves all existing field variables

\section*{Saving 1D data: SaveLine...}
```

Solver n
Equation = "SaveLine"
Procedure = File "SaveData" "SaveLine"
Filename = "g.dat"
File Append = Logical True
Polyline Coordinates(2,2) = Real 0.0 1.0 0.0 2.0
End
Boundary Condition m
Save Line = Logical True
End

```

\section*{Computing and saving oD data: SaveScalars}

Operators on bodies
- Statistical operators
- Min, max, min abs, max abs, mean, variance, deviation, rms
- Integral operators (quadratures on bodies)
- volume, int mean, int variance, int rms
- Diffusive energy, convective energy, potential energy

Operators on boundaries
- Statistical operators
- Boundary min, boundary max, boundary min abs, max abs, mean, boundary variance, boundary deviation, boundary sum, boundary rms
- Min, max, minabs, maxabs, mean
- Integral operators (quadratures on boundary)
- area
- Diffusive flux, convective flux

Other operators
- nonlinear change, steady state change, time, timestep size,...

\section*{Saving oD data: SaveScalars...}
```

Solver n
Exec Solver = after timestep
Equation = String SaveScalars
Procedure = File "SaveData" "SaveScalars"
Filename = File "f.dat"
Variable 1 = String Temperature
Operator 1 = String max
Variable 2 = String Temperature
Operator 2 = String min
Variable 3 = String Temperature
Operator 3 = String mean
End
Boundary Condition m
Save Scalars = Logical True
End

```

\section*{Slots for executing postprocessing solvers}
- Often the postprocessing solver need to computed only at desired slots, not at every time-step or coupled system iteration
- The execution is controlled by the "Exec Solver" keyword
- Exec Solver = before simulation
- Exec Solver = after simulation
- Exec Solver = before timesteo
- Exec Solver = after timestep
- Exec Solver = before saving
- Exec Solver = after saving
- The before/after saving slot is controlled by the output intervals
- Derived solvers often use the "before saving" slot
- Data is often saved with the "after saving" slot

\section*{Case: TwelveSolvers}

Natural convection with ten auxialiary solvers

\section*{Case: Motivation}
- The purpose of the example is to show the flexibility of the modular structure
- The users should not be afraid to add new atomistic solvers to perform specific tasks
- A case of 12 solvers is rather rare, yet not totally unrealitistic

\section*{Case: preliminaries}
- Square with hot wall on right and cold wall on left
- Filled with viscous fluid
- Bouyancy modeled with

Boussinesq approximation
- Temperature difference initiates a convection roll

Cold wall

\section*{Case: 12 solvers}
1. HeatSolver
2. FlowSolver
3. FluxSolver: solve the heat flux
4. StreamSolver: solve the stream function
5. VorticitySolver: solve the vorticity field (curl of vector field)
6. DivergenceSolver: solve the divergence
7. ShearrateSolver: calculate the shearrate
8. IsosurfaceSolver: generate an isosurface at given value
9. ResultOutputSolver: write data
10. SaveGridData: save data on uniform grid
11. SaveLine: save data on given lines
12. SaveScalars: save various reductions

Primary fields for natural convection


Pressure


Temperature

\section*{Derived fields for Navier-Stokes solution}


Shearrate field


Stream function



\section*{Derived fields for heat equation}

- Nodal loads only occur at boundaries (nonzero heat source)
- Nodal loads are associated to continuous heat flux by element size factor

Heat flux
Nodal heat loads

Visualization in differen postprocessors


GiD


Paraview

\section*{Example: total flux}

- Saved by SaveScalars
- Two ways of computing the total flux give different approximations
- When convergence is reached the agreement is good


\section*{Example: boundary flux}

- Saved by SaveLine
- Three ways of computing the boundary flux give different approximations
- At the corner the nodal flux should be normalized using only \(h / 2\)


\section*{Example, saving boundaries in .sif file}

Solver 2
Exec Solver = Always
Equation = "result output"
Procedure = "ResultOutputSolve"
"ResultOutputSolver"
Output File Name = case
Vtu Format = Logical True
Save Boundaries Only = Logical True End


\section*{Example, File size in Swiss Cheese}
- Memory consumption of vtu-files (for Paraview) was studied in the "swiss cheese" case
- The ResultOutputSolver with different flags was used to write output in parallel
- Saving just boundaries in single precision binary format may save over \(90 \%\) in files size compared to full data in ascii
- With larger problem sizes the benefits are amplified
\begin{tabular}{|c|c|c|c|}
\hline Binary output & Single Prec. & Only bound. & Bytes/node \\
\hline- & X & - & 376.0 \\
\hline X & - & - & 236.5 \\
\hline X & X & - & 184.5 \\
\hline X & - & X & 67.2 \\
\hline X & X & X & 38.5 \\
\hline
\end{tabular}

\section*{Manually editing the command files}
- Only the most important solvers and features are supported by the GUI
- Minor modifications are most easily done by manual manipulation of the files
- The tutorials, test cases and documentation all include usable sif file pieces
- Use your favorite text editor (emacs, notepad++,...) and copy-paste new definitions to your .sif file
- If your additiones were sensible you can rerun your case
- Note: you cannot read in the changes made in the .sif file

\section*{Exercise}
- Study the command file with 12 solvers
- Copy-paste an appropriate solver from there to some existing case of your own
- ResultOutputSolver forVTU output
- StreamSolver, VorticitySolver, FluxSolver,...
- Note: Make sure that the numbering of Solvers is consistant
o Solvers that involve finite element solution you need to activate by Active Solvers
- Run the modified case
- Visualize results in Paraview in different ways

\section*{Using tests as a starting point}

- There are over 500 consistancy tests that come with the Elmer distribution
- The hope is to minimize the propability of new bugs
- The tests are small for speedy computation
- Step-by-step instructions
1. Goto tests at \$ELMER_HOME/tests
2. Choose a test case relevant to you (by name, or by grep)
- Look in Models manual for good search strings
3. Copy the tests to your working directory
4. Edit the sif file
- Activate the output writing: Post File
- Make the solver more verbose: Max Output Level
5. Run the case (see runtest.cmake for the meshing procedure)
- Often just: ElmerSolver
6. Open the result file to see what you got
7. Modify the case and rerun etc.

\section*{Conclusions}
- It is good to think in advance what kind of data you need
-3D volume and 2D surface data
- Derived fields
o1D line data
○oD lumped data
- Internal strategies may allow better accuracy than doing the analysis with external postprocessing software
- Consistent use of basis functions to evaluate the data
- Often the same reduction operations may be done also at later stages but with significantly greater effort


\title{
Parallel computing with Elmer
}

\author{
ElmerTeam \\ CSC - IT Center for Science, Finland
}

CSC, 2018

\section*{Algorithm scalability}

- Before going into parallel computation let's study where the bottle-necks will appear in the serial

\section*{system}
- Each algorithm/procedure has a characteristic scaling law that sets the lower limit to how the solution time \(t\) increases with problem size \(n\)
oThe parallel implementation cannot hope to beat this

\(n\)
- Targeting very large problems the starting point should be nearly optimal (=linear) algorithm!

\section*{Poisson equation at "Winkel"}

Mesh generation
alpha betcita Gmsh 21.4
- Mesh generation is cheapest
- Success of various iterative methods determined mainly by preconditioning strategy
- Best preconditioner is clustering multigrid method (CMG)
- For simple Poisson almost all preconditioners work reasonable well
- Direct solvers differe significantly in scaling
Linear solver
\begin{tabular}{ll} 
alpha & beta \\
178.30 & 1.09 \\
180.22 & 1.10 \\
175.20 & 1.10
\end{tabular}
BiCgStab + ILUO 192.501.13
\begin{tabular}{lll}
\(\cdots\) \\
\(C G\) & 282.07 & 1.16
\end{tabular}
Idrs (4) + vanka 295.181 .16
\begin{tabular}{lll}
\(C G+\operatorname{diag}\) & 257.98 & 1.17 \\
BiCgStab \((4)+\) diag & 290.11 & 1.19
\end{tabular}
\begin{tabular}{lrr} 
MUMPS (PosDef) & 4753.99 & 1.77 \\
MUMPS & 12088.74 & 1.93 \\
umfpack & 74098.48 & 2.29
\end{tabular}

\section*{Motivation for using optimal linear solvers}
- Comparison of algorithm scaling in linear elasticity between different preconditioners
- ILUı vs. block preconditioning (Gauss-Seidel) with agglomeration multigrid for each component
- At smallest system performance about the same
- Increasing size with \(8^{\wedge} 3=512\) gives the block solver scalability of O(~1.03) while ILU1 fails to converge
\begin{tabular}{|l|l|l|l|l|}
\hline & \multicolumn{2}{|l|}{ BiCGstab(4)+ILU1 } & \multicolumn{2}{|l|}{ GCR+BP(AMG) } \\
\hline \#dofs & T(s) & \#iters & T(s) & \#iters \\
\hline 7,662 & 1.12 & 36 & 1.19 & 34 \\
\hline 40,890 & 11.77 & 76 & 6.90 & 45 \\
\hline 300,129 & 168.72 & 215 & 70.68 & 82 \\
\hline \(2,303,472\) & \(>21,244^{*}\) & \(>5000^{*}\) & 756.45 & 116 \\
\hline
\end{tabular}


Simulation Peter Råback, CSC.
* No convergence was obtained

\section*{Parallel computing concepts}

\section*{Computer architectures}
- Shared memory
- All cores can access the whole memory
- Distributed memory
- All cores have their own memory
- Communication between cores is needed in order to access the memory of other cores
- Current supercomputers combine the distributed and shared memory (within nodes) approaches


\section*{Programming models}
- Threads (pthreads, OpenMP)
- Can be used only in shared memory computer
- Limited parallel scalability
- Simpler or less explicit programming
- Message passing (MPI)
- Can be used both in distributed and shared memory computers
- Programming model allows good parallel scalability
- Programming is quite explicit
- Massively parallel FEM codes use typically MPI as the main parallelization strategy
- As does Elmer!

\section*{Weak vs. strong parallel scaling}

\section*{Strong scaling}
- How the solution time \(T\) varies with the number of processors \(P\) for a fixed total problem size.
- Optimal case: \(P \times T=\) const.
- A bad algorithm may have excellent strong scaling
- Typically \(10^{4-105}\) dofs needed in FEM for good strong scaling

\section*{Weak scaling}
- How the solution time \(T\) varies with the number of processors \(P\) for a fixed problem size per processor.
- Optimal case: \(T=\) const.
- Weak scaling is limited by algorithmic scaling


\section*{Serial workflow of Elmer}
- All steps in the workflow are serial
- Typically solution of the linear system is the main bottle-neck
- For larger problems bottle-necks starts to appear in all phases of the serial workflow


\section*{Basic Parallel workflow of Elmer}
- Addiational partition step using ElmerGrid
- Both assembly and solution is done in parallel using MPI
- Assembly is trivially parallel


\section*{ElmerGrid partitioning commands}

Basic volume mesh partitioning options
(geometric partitioning and Metis graph partitiong)
```

-partition int[3] : the mesh will be partitioned in cartesian main directions
-partorder real[3] : in the 'partition' method set the direction of the ordering
-partcell int[3] : the mesh will be partitioned in cells of fixed sizes
-partcyl int[3] : the mesh will be partitioned in cylindrical main directions
-metis int : mesh will be partitioned with Metis using mesh routines
-metiskway int : mesh will be partitioned with Metis using Kway routine
-metisrec int : mesh will be partitioned with Metis using Recursive routine
-metiscontig : enforce that the metis partitions are contiguous
-partdual : use the dual graph in partition method (when available)

```

There are additional flags to control the partitioning of contact boundaries and halo elements.

\section*{ElmerGrid partitioning examples}
- ElmerGrid 22 mesh -partcell \(n_{x} n_{y} n_{z}\)
- Partition elements in a uniform grid based on the bounding box
- Number of partitions may be lower than the product if there are empty cells
- Does not quarantee that partitions are of same size
- ElmerGrid 22 mesh -partition \(n_{x} n_{y} n_{z}\)
- Partition elements recursively in the main coordinate directions
- Partitions are of same size
- Goodness depends heavily on the geometry
- ElmerGrid 22 mesh -metisrec \(n\)
- Partition elements using a recursive routine of Metis
- Cannot beat the geometric strategy for some ideal shapes
- Robust in that partitioning is always reasonable

Mesh partitioning with ElmerGrid - structured mesh


Mesh partitioning with ElmerGrid - unstructured mesh


\section*{Mesh structure of Elmer}


\section*{Serial}
meshdir/
- mesh.header
size info of the mesh
- mesh. nodes
node coordinates
- mesh.elements
bulk element defs
- mesh.boundary
boundary element defs with reference to parents

\section*{Parallel}
meshdir/partitioning.N/
- mesh.n.header
- mesh.n.nodes
- mesh.n.elements
- mesh.n.boundary
- mesh.n.shared
information on shared nodes
for each in \([0, N-1]\)

\section*{Parallel linear solvers in Elmer}


\section*{Iterative}
- HUTITER
- Krylov methods initially coded at HUT
- Hypre
- Krylov solvers
- Algebraic multigrid: BoomerAMG
- Truly parallel ILU and Parasails preconditioning
- Trilinos
- Krylov solvers
- Algebraic multigrid: ML

○ ..
- ESPRESO
- FETI library of IT4 http://espreso.it4i.cz/

\section*{Direct}
- MUMPS
- Direct solver that may work when averything else fails
- MKL Pardiso
- Comes with the Intel MKL library
- Multihreaded

\section*{MUMPS}

\section*{Partitioning and matrix structure}



Contiguous parallel nümbering used
- Shared nodes result to need for communication.
- Each dof has just one owner partiotion and we know the neighbours for
- Owner partition usually handles the full row
- Results to point-to-point communication in MPI
- Matrix structure sets challenges to efficient preconditioners in parallel
- It is more difficult to implement algorithms that are sequential in nature, e.g. ILU
- Krylov methods require just matrix vector product, easy!
- Communication cannot be eliminated. It reflects the local interactions of the underlying PDE

\section*{Partitioning and matrix structure - unstructured mesh}


22


Metis partitioning into 8
- Partitioning should try to minimize communication
- Relative fraction of shared nodes goes as \(\mathrm{N}^{\wedge}(-1 / \mathrm{DIM})\)
- For vector valued and high order problems more communication with same dof count

\section*{Differences in serial and parallel algorithms}
- Some algorithms are slightly different in parallel
- ILU in ElmerSolver library is performed only blockwise which may result to inferior convergence
- Diagonal and vanka preconditions are exactly the same in parallel


\section*{Parallel computation in ElmerGUI}
- If you have parallel environment it can also be used interactively via ElmerGUI
- Calls ElmerGrid automatically for partiotioning (and fusing)


\section*{Parallel postprocessing using Paraview}
- Use ResultOutputSolver to save data to .vtu files
- The operation is almost the same for parallel data as for serial data
- There is a extra file . pvtu that holds is a wrapper for the parallel .vtu data of each partition

\section*{Summary: Files in serial vs. parallel solution}

\section*{Serial}
- Serial mesh files
- Command file (.sif) may be given as an inline parameter
- Execution with

ElmerSolver [case.sif]
- Writes results to one file

\section*{Parallel}
- Partitioned mesh files
- ELMERSOLVER_STARTINFO is always needed to define the command file (.sif)
- Execution with
mpirun -np N ElmerSolver_mpi
- Calling convention is platform dependent
- Writes results to \(N\) files + 1 wrapper file

\section*{Example: Weak scaling of Elmer (FETI)}
\begin{tabular}{|l|l|l|l|}
\hline \#Procs & Dofs & Time \((\mathbf{s})\) & Efficiency \\
\hline 8 & 0.8 & 47.80 & - \\
\hline 64 & 6.3 M & 51.53 & 0.93 \\
\hline 125 & 12.2 M & 51.98 & 0.92 \\
\hline 343 & 33.7 M & 53.84 & 0.89 \\
\hline 512 & 50.3 M & 53.90 & 0.89 \\
\hline 1000 & 98.3 M & 54.54 & 0.88 \\
\hline 1331 & 131 M & 55.32 & 0.87 \\
\hline 1728 & 170 M & 55.87 & 0.86 \\
\hline 2197 & 216 M & 56.43 & 0.85 \\
\hline 2744 & 270 M & 56.38 & 0.85 \\
\hline 3375 & 332 M & 57.24 & 0.84 \\
\hline
\end{tabular}

Solution of Poisson equation with FETI method where local problem (of size 32^3=32,768 nodes) and coarse problem (distributed to 10 partitions) is solved with MUMPS. Simulation with Cray XC (Sisu) by Juha Ruokolainen, CSC, 2013.

\section*{Block preconditioner: Weak scaling of 3D driven-cavity}
\begin{tabular}{|l|l|l|l|}
\hline Elems & Dofs & \#procs & Time (s) \\
\hline \(34^{\wedge} 3\) & 171,500 & 16 & 44.2 \\
\hline \(43^{\wedge} 3\) & 340,736 & 32 & 60.3 \\
\hline \(54^{\wedge} 3\) & 665,500 & 64 & 66.7 \\
\hline \(68^{\wedge} 3\) & \(1,314,036\) & 128 & 73.6 \\
\hline \(86^{\wedge} 3\) & \(2,634,012\) & 256 & 83.5 \\
\hline \(108^{\wedge} 3\) & \(5,180,116\) & 512 & 102.0 \\
\hline \(13^{\wedge} 3\) & \(9,410,548\) & 1024 & 106.8 \\
\hline
\end{tabular}


Velocity solves with Hypre: CG + BoomerAMG preconditioner for the 3D driven-cavity case (Re=100) on Cray XC (Sisu).
Simulation Mika Malinen, CSC, 2013.

Scalability of edge element AV solver for end-windings


\begin{tabular}{|l|l|l|}
\hline \#Procs & Time(s) & \(\mathrm{T}_{2 \mathrm{P}} / \mathrm{T}_{\mathrm{P}}\) \\
\hline 4 & 1366 & - \\
\hline 8 & 906 & 1.5 \\
\hline 16 & 260 & 3.5 \\
\hline 32 & 122 & 2.1 \\
\hline 64 & 58.1 & 2.1 \\
\hline 128 & 38.2 & 1.8 \\
\hline 256 & 18.1 & 2.1 \\
\hline
\end{tabular}

Magnetic field strength (left) and electric potential (right) of an electrical engine end-windings. Meshing M. Lyly, ABB. Simulation (Cray XC, Sisu) J. Ruokolainen, CSC.

\section*{Coupled model for electrical machines}
- Monolithic parallel linear system including
- Electric scalar potential using nodal elements
- Magnetic vector potential using edge elements (in 3D)
- Mortar projector for the nodal dofs \(P_{v}\) (for conductors)
- Mortar projector for the edge dofs \(P_{a}\) (in 3D)
- Current conditions for case driven by external circuit (few rather dense rows)
\(\left(\begin{array}{ccccc}V_{v} & V_{a} & P_{v}^{T} & 0 & 0 \\ A_{v} & A_{a} & 0 & P_{a}^{T} & N \\ P_{v} & 0 & 0 & 0 & 0 \\ 0 & P_{a} & 0 & 0 & 0 \\ 0 & S & 0 & 0 & R\end{array}\right)\left(\begin{array}{c}v \\ a \\ \lambda_{v} \\ \lambda_{a} \\ i\end{array}\right)=\left(\begin{array}{c}f_{a} \\ f_{v} \\ 0 \\ 0 \\ V_{\text {ext }}\end{array}\right)\)
- Solved with Krylov method, e.g. GCR or BiCGStab(I)
- Hybrid preconditioning strategy
- Vector potential with diagonal
- Scalar potential \& mortar projectors with ILU
- Electrical circuits either with ILU or MUMPS
- Still some challenges on robustness!

\section*{Hybrid partitioning scheme}
- The linear system arising from the electromagentic problem must be solved together with the continuity constraints
- To minimize communication (and coding) effort we partition the mesh cleverly
- Electrical machines have always rotating interface:
 Partition the interface elements so that opposing element layers on the cylinder are always within the same partition
- Unstructured surface meshes are treated similarly except halo elements are also saved on the boundary
- Other elements are partitioned with Metis

- Local mortar conditions much easier to deal with!

\section*{Parallel workflow for meshing bottle-necks}
- Large meshes may be finilized at the parallel level


\section*{Mesh Multiplication}
- Split elements edges after partitioning at parallel level o effectively eliminating memory and I/O bottle-necks - Each multiplication creates \(2^{\wedge}\) DIM-fold number of elements - Does not increase accuracy of geometry presentation - May inherit mesh grading
- CPU time used in neglible
\begin{tabular}{|l|l|l|l|l|l|}
\hline Mesh & \#splits & \#elems & \#procs & \begin{tabular}{l} 
T_center \\
\((\mathbf{s})\)
\end{tabular} & \begin{tabular}{l} 
T_graded \\
\((\mathbf{s})\)
\end{tabular} \\
\hline A & 2 & 4 M & 12 & 0.469 & 0.769 \\
\hline & 2 & 4 M & 128 & 0.039 & 0.069 \\
\hline & 3 & 32 M & 128 & 0.310 & 0.549 \\
\hline B & 2 & 4.20 M & 12 & 0.369 & \\
\hline & 2 & 4.20 M & 128 & 0.019 & \\
\hline & 3 & 33.63 M & 128 & 0.201 & \\
\hline
\end{tabular}

Mesh A: structured, 62500 hexahedrons
Mesh B: unstructured, 65689 tetrahedrons

\section*{Overcoming bottle-necks in postprocessing}
- Visualization
- Paraview and Visit excellent tools for parallel visualization
- Access to all data is often an overkill
- Reducing data
- Saving only boundaries
- Uniform point clouds
- A priori defined isosurfaces
- Using coarser meshes for output when hierarchy of meshes exist
- Extracting data
- Dimensional reduction (3D -> 2D)
- Averaging over time
- Integrals over BCs \& bodies
- More robust I/O
- Not all cores should write to disk in massively parallel simulations
\begin{tabular}{|c|c|c|c|}
\hline \begin{tabular}{c} 
Binary \\
output
\end{tabular} & \begin{tabular}{c} 
Single \\
Prec.
\end{tabular} & \begin{tabular}{c} 
Only \\
bound.
\end{tabular} & \begin{tabular}{c} 
Bytes/ \\
node
\end{tabular} \\
\hline- & X & - & 376.0 \\
\hline X & - & - & 236.5 \\
\hline X & X & - & 184.5 \\
\hline X & - & X & 67.2 \\
\hline X & X & X & 38.5 \\
\hline
\end{tabular}
- HDF5+XDML output available for Elmer, mixed experiences

\section*{Hybridization of the Finite Element code}
- The number of cores in CPUs keep increasing but the clock speed has stagnated
- Significant effort has been invested for the hybrization of Elmer
- Assembly process has been multithreaded and vectorized
○"Coloring" of element to avoid race conditions
- Speed-up of assembly for typical elements varies between 2 to 8.
- As an accompanion the multitreaded assembly requires multithreaded linear \({ }_{63}\) solvers.

Multicore speedup, \(\mathrm{P}=\mathbf{2}\)
128 threads on KNL, 24 threads on HSW
\begin{tabular}{lcccc}
\hline \begin{tabular}{l} 
Element (\#ndofs, \\
\#quadrature \\
points)
\end{tabular} & Speedup & \multicolumn{3}{c}{\begin{tabular}{c} 
Optimized local \\
matrix
\end{tabular}} \\
& KNL & HSW & KNL & HSW \\
Line (3, 4) & 0.7 & 2.0 & 4.2 M & 14.5 M \\
\hline Triangle (6, 16) & 2.5 & 3.9 & 2.6 M & 6.5 M \\
\hline Quadrilateral (8, 16) & 2.8 & 4.0 & 2.6 M & 6.6 M \\
\hline Tetrahedron (10, 64) & 7.9 & 6.3 & 1.0 M & 1.5 M \\
\hline Prism (15, 64) & 8.3 & 5.8 & 0.8 M & 0.9 M \\
\hline Hexahedron (20, 64) & 7.2 & 5.8 & 0.6 M & 0.9 M \\
\hline
\end{tabular}

Speed-up assembly process for poisson equation using 2nd order p-elements. Juhani Kataja, CSC, IXPUG Annual Spring Conference 2017.

\section*{Recipes for resolving scalability bottle-necks}
- Finalize mesh on a parallel level (no I/O)
- Mesh multiplication or parallel mesh generation
- Use algorithms that scale well
- E.g. Multigrid methods
- If the initial problem is difficult to solve effectively divide it into simpler sub-problems
- One component at a time -> block preconditioners
- GCR + Block Gauss-Seidel + AMG + SGS
- One domain at a time -> FETI
- Splitting schemes (e.g. Pressure correction in CFD)
- Analyze results on-the-fly and reduce the amount of data for visualization

\section*{Future outlook}
- Deeper integration of the workflow
o Heavy pre- and postprocessing internally or via API
- Cheaper flops from new multicore environments
o Interesting now also for the finite element solvers
- Usable via reasonable programming effort; attention to algorithms and implementation
- Complex physics introduces always new bottle-necks
- Rotating boundary conditions in parallel...```

