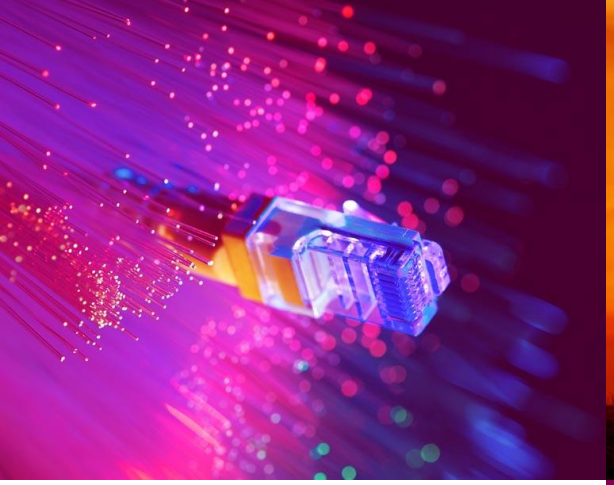

LUMI-Hackathon - Elmer

2024-09-26

CSC – Suomalainen tutkimuksen, koulutuksen, kulttuurin ja julkishallinnon ICT-osaamiskeskus

Hackathon

- How many people involved from ELMER team?

- 4 developers (J. Kataja, J. Ruokalainen, P. Råback, T. Zwinger) – underlined names present in Brussels.
- Worked with Emanuele Vitali (LUST) and Thierry Braconnier, George Markomanolis, Samuel Antaro

Getting Elmer ready for GPUs

- **POP3 audit** with MPI + OMP SIMD CPU version (baseline for next audit using GPU offload)
- Working through **code**-base to reach **compatibility** with several involved compiler suites (gcc, clang, CrayCE) – LUMI needs Cray-Fortran to **enable OMP target** offloading!
- Working on different **interfaces for offloading linear system** solution step to GPUs (M250 and A100)

Elmer

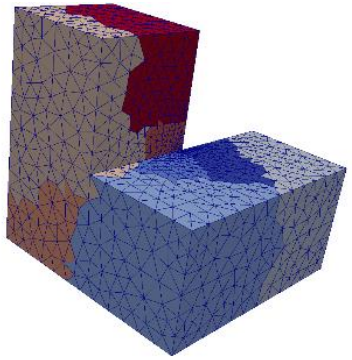

MPI – domain decomposition

OpenMP

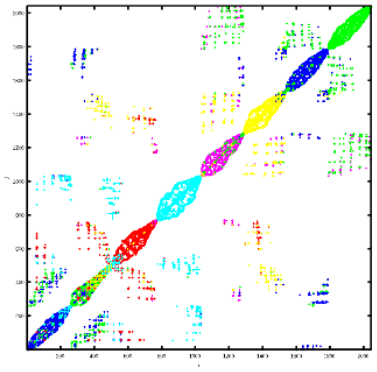

Target offload for **matrix assembly**

rocALUTION

AMD ROCm

Sparse matrix solution library for particular GPU

Goals for Hackathon

- Profile and measure the current code and interpret the profiling results
  - Optimize accordingly
- Build a system so that user may provide material (c) and force (k) parameters
  - Precompute to nodes on CPU w/ ElmerLibrary
- Keep CSR matrix on GPU and call rocALUTION solvers
- Explore compilers (CCE 16/17/18 and AMDclang/flang)

Goals for Hackathon

- Profile and measure the current code and interpret the profiling results
  - Optimize accordingly
- Build a system so that user may provide material (c) and force (k) parameters
  - Precompute to nodes on CPU w/ ElmerLibrary
- Keep CSR matrix on GPU and call rocALUTION solvers
- Explore compilers (CCE 16/17/18 and AMDclang/flang)

Compiler testing: Compile – (run) - crash

- Tested **AMDclang/flang**:
- Error on LLVM level?
- Difficult to get reproducer

```
[ 64%] Building Fortran object fem/src/CMakeFiles/
elmersolver.dir/GeneralUtils.F90.o
cd /scratch/project_462000007/tzwinger/cpe/
build_amd_Elmer_amd_lumi-hackathon-hacks_7345c18e7/fem/src && /
pfs/lustrep2/scratch/project_465001361/elmer/rocm-
afar-5891-0.5/bin/amdflang -DCONTIG="" -DELMER_HAVE_MPIF_HEADER
-DHAVE_EXECUTECOMMANDLINE -DUSE_ARPACK -DUSE_ISO_C_BINDINGS -
Delmersolver_EXPORTS -I/scratch/project_462000007/tzwinger/cpe/
build_amd_Elmer_amd_lumi-hackathon-hacks_7345c18e7 -I/scratch/
project_462000007/tzwinger/cpe/elmerfem/contrib/lua-5.1.5/src -
I/scratch/project_462000007/tzwinger/cpe/
build_amd_Elmer_amd_lumi-hackathon-hacks_7345c18e7/fem/src -I/
scratch/project_462000007/tzwinger/cpe/elmerfem/fhutiter/src -
I/scratch/project_462000007/tzwinger/cpe/
build_amd_Elmer_amd_lumi-hackathon-hacks_7345c18e7/fem/src/
binio -I/scratch/project_462000007/tzwinger/cpe/
build_amd_Elmer_amd_lumi-hackathon-hacks_7345c18e7/fhutiter/src
-I/scratch/project_462000007/tzwinger/cpe/elmerfem/umfpack/src/
umfpack/include -fopenmp -O2 -g -DNDEBUG -J../../fmodules -fPIC
-c /scratch/project_462000007/tzwinger/cpe/elmerfem/fem/src/
GeneralUtils.F90 -o CMakeFiles/elmersolver.dir/
GeneralUtils.F90.o
/tmp/GeneralUtils-12a2dd.ll:10066:24: error: use of undefined
value '%.U2939'
 10066 |         %161 = load i64, i64* %.U2939, align 8, !tbaa
!27623, !dbg !27588
       |                                ^
1 error generated.
make[2]: *** [fem/src/CMakeFiles/elmersolver.dir/
build.make:257: fem/src/CMakeFiles/elmersolver.dir/
GeneralUtils.F90.o] Error 1
```

Compiler testing: Compile – (run) - crash

- Cray Compilers:
- Only working combination @LUMI:
  - We then use CCE17 runtime

```
COMPILER="cce/16.0.1"
MPI="cray-mpich/8.1.27"
echo "loaded modules:"
module load $COMPILER $MPI
module load LUMI/23.09 partition/G
rocm/5.4.6
```

Conclusions: of hackathon

- Forget current stack CCE17 for compilation
- Workarounds for CCE18 – investment into next system update (Emanuele, Thierry, Thomas)
  - Main issue when passing pointer and receiving as object in subroutine (allowed by standard – object should behave as associated to the object behind the pointer)
  - In CCE17 – internal compiler error + segfault at runtime
  - In CCE18 – segfault; -O2 doesn't work – segfault in runtime
  - Workaround: callee function dummy argument changed to pointer instead of object

OpenMP kernels

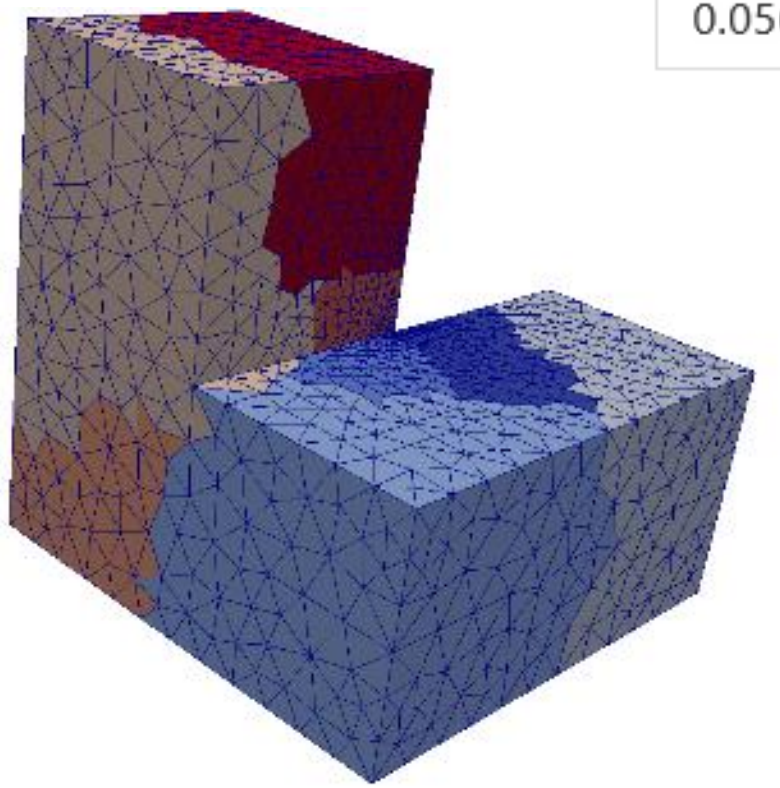

μs per element:

cpu_fast	cpu_slow	gpu_threadlimit	gpu_nothreadlimit
0.0565	0.567	0.0727	0.165

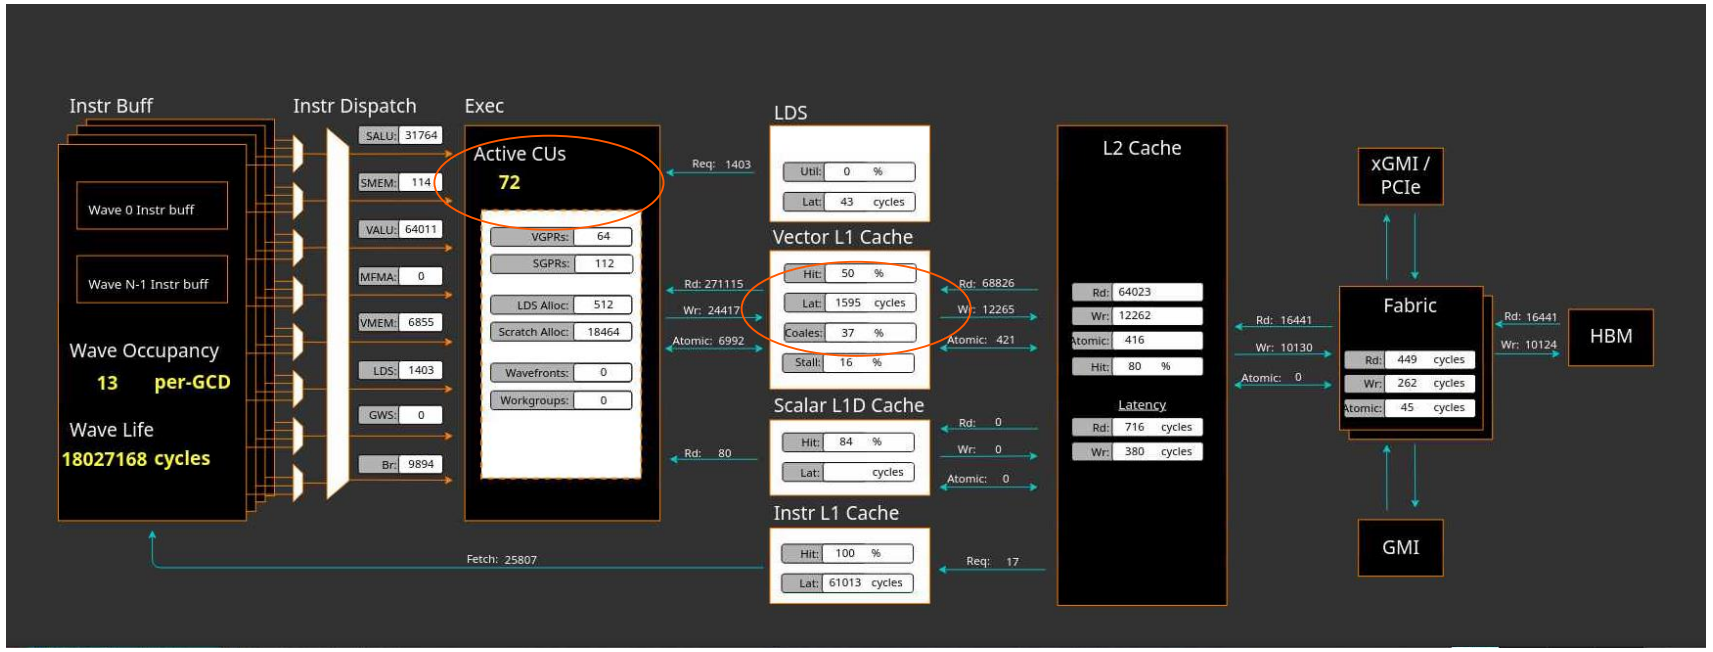

gpu_threadlimit: normal gpu offloading constructs + `num_teams(220) thread_limit(64)`

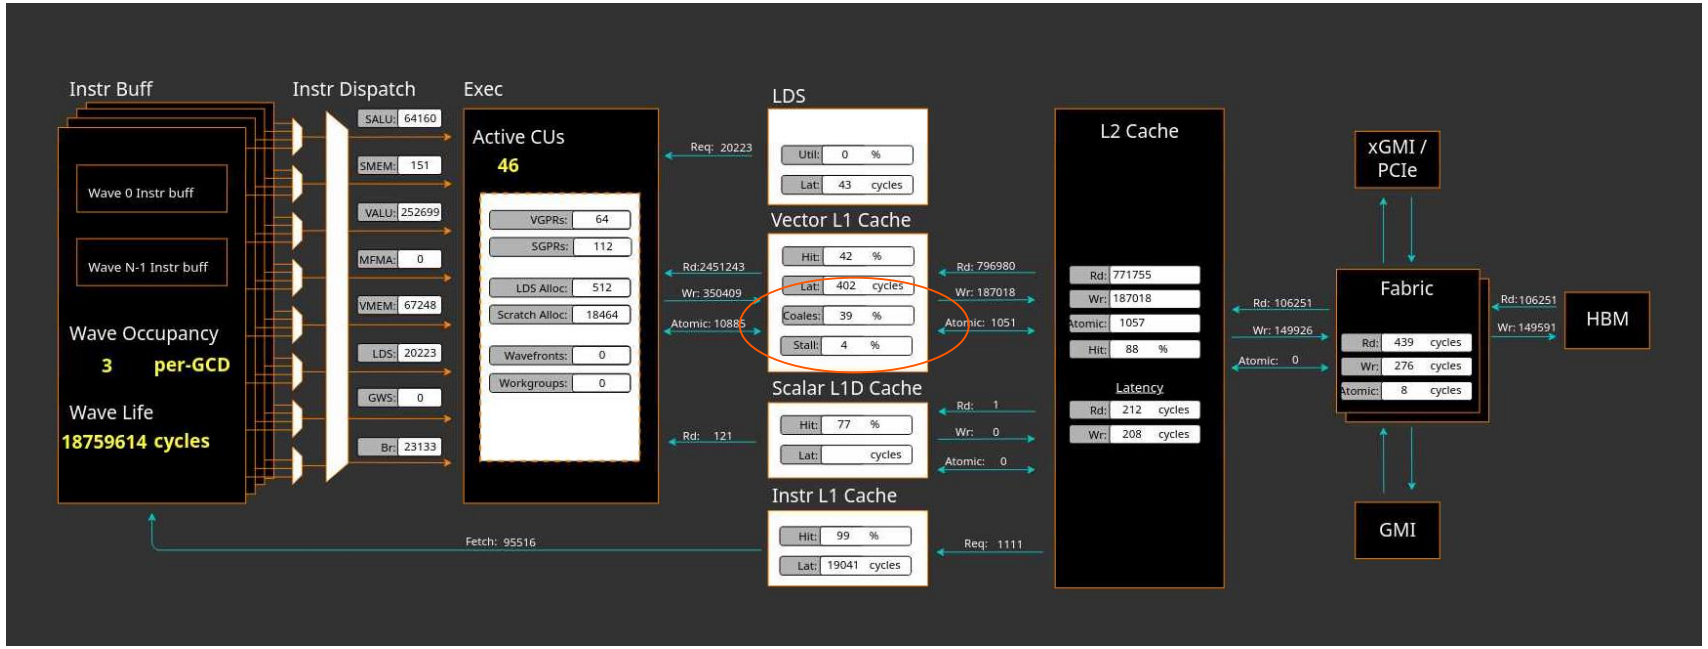

Coalesced is low (access to memory) – room for improvement on data-layout

10.1 Overall Instruction Mix

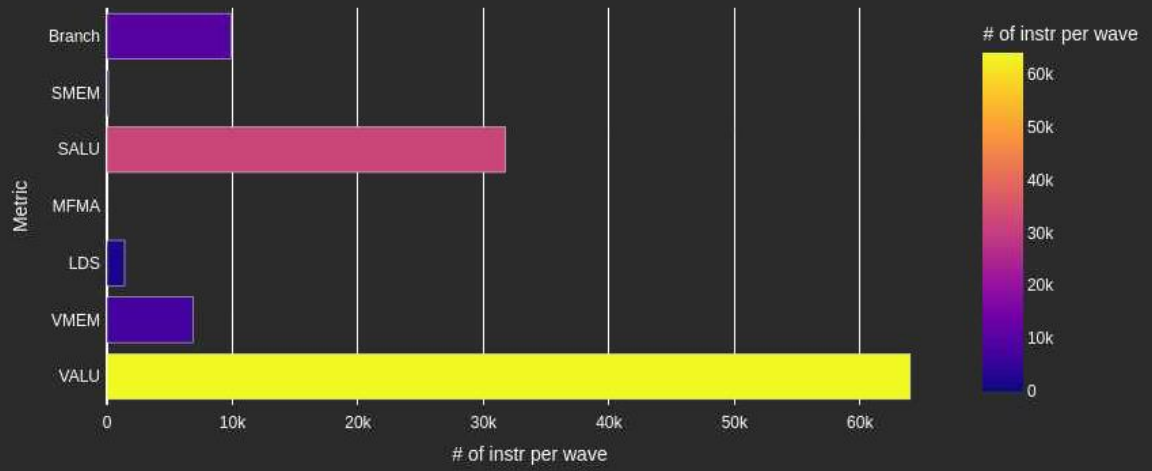

10.2 VALU Arithmetic Instr Mix

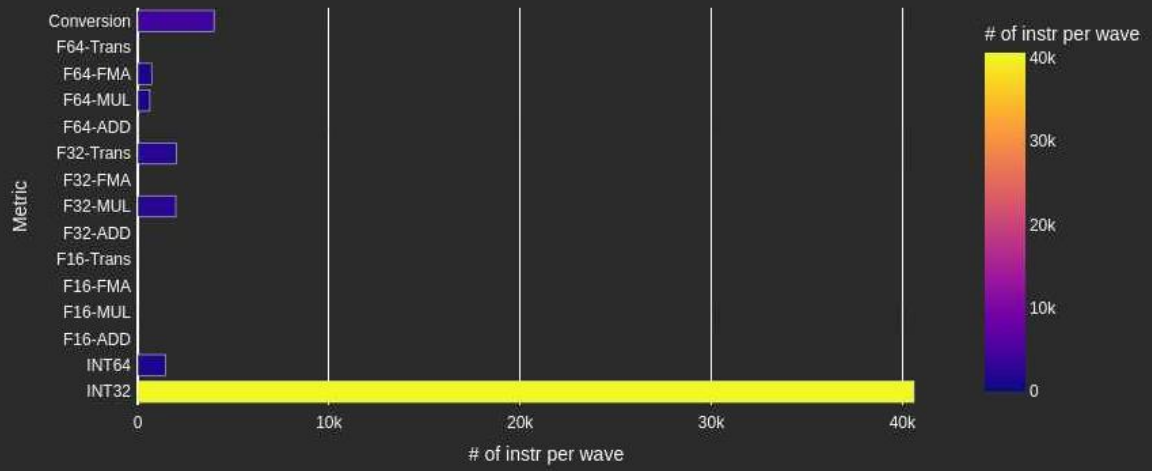

10.1 Overall Instruction Mix

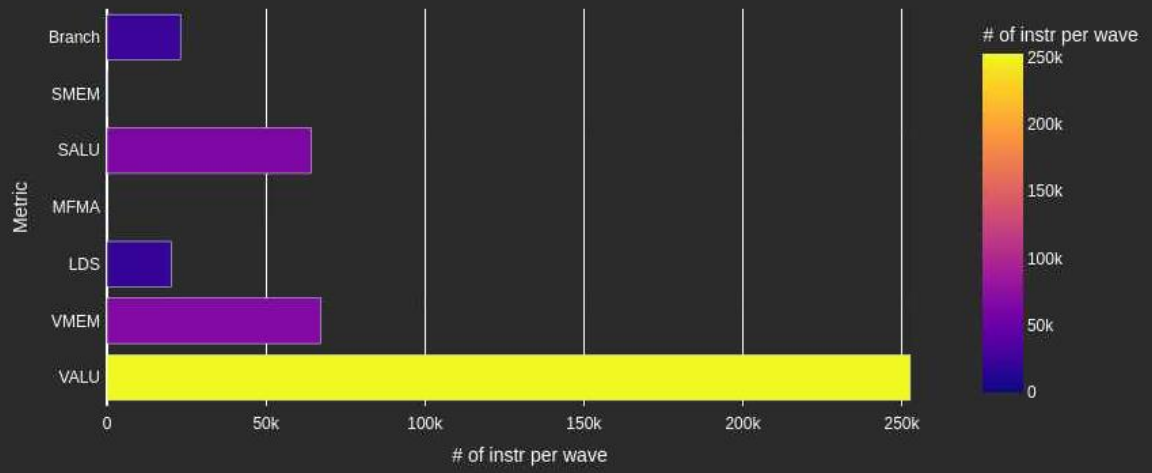

10.2 VALU Arithmetic Instr Mix

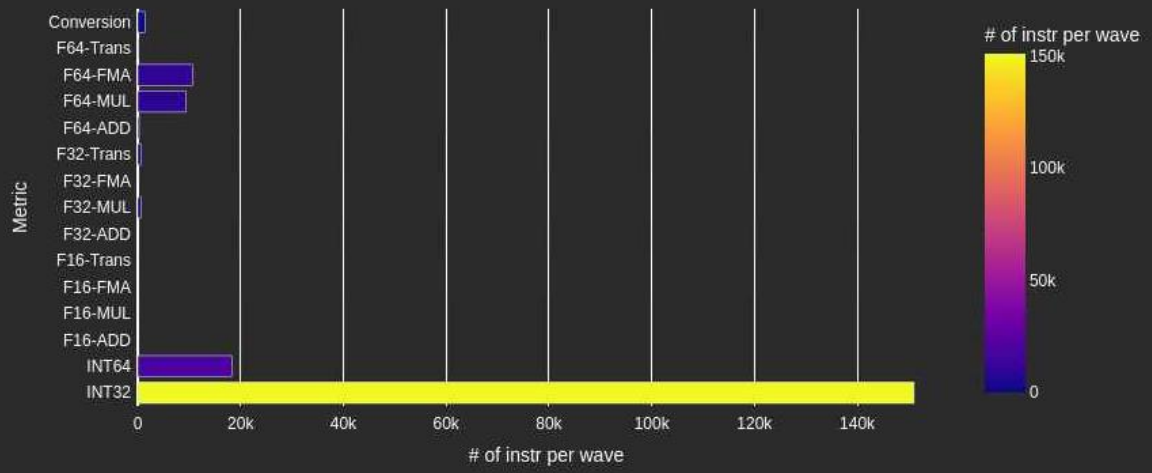

A lot of integer operations

```
183     #if 1
184         do i = 1,nd
185           do j = 1,nd
186             colind = 0
187             do k = rows(l2g(elem,i)), (rows(l2g(elem,i)+1)-1)
188               colind = colind + merge(k, 0, cols(k) == l2g(elem,j))
189             end do
190
191     #ifdef DEBUGPRINT
192             if (round < 3) then
193               print *, colind, stiff(i,j), l2g(elem,i), l2g(elem,j)
194             end if
195     #endif
196             values(colind) = values(colind) + stiff(i,j)
197             ! val_inds(elem,(i-1)*nd+j) = colind
198           end do
199           rhs(l2g(elem,i)) = rhs(l2g(elem,i)) + force(i)
200         end do
201     #endif
```

rocALUTION

- Run without MPI works (single task, single node multiple GPUs)
- Run with MPI failed
- Identified issue: missing device to host copy of results vector
- We think we fixed it – too little time to show results