## Automatic distribution of OpenMP kernels using CATO

# Jannek Squar Tim Jammer Michael Blesel Michael Kuhn Thomas Ludwig 2021-03-25

Universität Hamburg squar@informatik.uni-hamburg.de



#### Motivation

Design

Implementation

Outlook

Jannek Squar

## 1. Increasing core count of manycore architecture

- $\rightarrow$  **2** peak performance **2**
- ightarrow Parallelisation techniques on shared memory needed
  - First choice: OpenMP
- 2. Increasing gap between memory and CPU performance in manycore architecture
  - $\rightarrow$   $\bigcirc$  memory <u>CPU core</u>  $\bigcirc$  favours CPU-bound applications
  - $\rightarrow~$  Limited maximal problem size per node
  - $\rightarrow~\mbox{Parallelisation}$  techniques on distributed memory needed
    - First choice: MPI

- 1. Focus on shared memory parallelisation techniques (e.g. OpenMP)
- 2. Limited problem size to fit single node memory
  - $\rightarrow \,$  Limited horizontal scalability
- 3. Trivial distribution approach: Fragment problem space
  - Execute new process for each sub-problem
    - $\rightarrow \;$  Additional overhead
    - $\rightarrow$  Sub-problems must be independent

## **Motivation**

- Better explanatory power through model run on extended problem size
  - 🖸 Problem area 🖸
  - 🛛 Resolution 🖸
  - 🖸 Grid dimension 🖸
- OpenMP:
  - Easy to use ...
  - ... but limited to shared memory
- MPI:
  - More difficult to use ...
  - ... but allows distributed computing

#### Goals:

- Automatic distribution of input problem
- Minimal data redundancy

Jannek Squar

- Distributed Shared Memory
  - PGAS [Pad11] (e.g. UPC, Coarray Fortran, TreadMarks [ACD+96])
  - XcalableMP [Xca], OpenMPD [LSB07]
  - SSI [BCJ01], Cluster OpenMP [Hoe06]
- Existing MPI approaches [BE05, SCC15, MMPS08, HFE11, AT06]

Not suitable to achieve our goal, because:

- Require changes to code or OS/FS
- Focus on optimisation of performance rather than of problem size
- Some not actively maintained anymore

Motivation

## Design

Implementation

Outlook

Jannek Squar

**Design | Frameworks** 

MPI

- Message Passing Interface
- Communication between processes
- User must pay attention to:
  - memory consistency
  - synchronisation
- Established in HPC



Figure 1: Distributed computing [Ble21]



- Modular open source compiler infrastructure
- Abstraction layer: Intermediate Representation (IR)



Figure 2: Modular compiler infrastructure (based on [Lat])

LLVM (II)





Figure 3: LLVM workflow: Frontend  $\rightarrow$  Optimizer  $\rightarrow$  backend (based on [Sam])

Jannek Squar

```
int *counter = (int*)malloc(1*sizeof(int));
*counter = 0;
#pragma omp parallel
{
    printf("Hello from thread %d\n", omp_get_thread_num());
    #pragma omp critical
    (*counter)**;
}
```

Listing 1: Example C application (extract)

Jannek Squar

2 3

4

5

6 7

8

## Intermediate Representation Language (II)

```
%1 = alloca i32*, align 8
 1
        %3 = call noalias i8* @malloc(i64 4) #3
 2
        %4 = hitcast i8* %3 to i32*
 3
 4
         store i32* %4, i32** %1, align 8
         [...]
 5
         call void ( [...] ) @__kmpc_fork_call( [...] @.omp_outlined. to void ( [...] )), i32** %1)
 6
 7
         [...]
        ret i32 0
8
 9
10
11
    define internal void @.omp outlined.([...]) #0 {
12
         [...]
13
        %8 = call i32 @omp get thread num()
        %9 = call i32 (i8*, ...) @printf([...] i32 %8)
14
         [...]
15
         call void @__kmpc_critical(%ident_t* @0, i32 %11, [8 x i32]* @.gomp_critical_user_.var)
16
17
        %12 = load i32*. i32** %7. align 8
18
        %13 = load i32. i32* %12. align 4
        %14 = add nsw i32 %13. 1
19
        store i32 %14. i32* %12. align 4
20
         call void @ kmpc end critical(%ident t* @0. i32 %11. [8 x i32]* @.gomp critical user .var)
21
         ret void
22
23
```

#### CATO (compiler assisted source transformation of OpenMP kernels)

- Main component: LLVM transform pass
  - Local installation + Robustness
- Derive information about suitable data distribution from OpenMP kernels
  - $\rightarrow~$  Shared memory might be a candidate for distribution
- Automated code replacement
  - Probably less scalable...
  - ... but focus on eased optimisation of memory usage and horizontal scaling
  - Eased usage of advanced features of MPI-3
  - ightarrow Can be used by untrained personnel

Motivation

Design

Implementation

Outlook

Jannek Squar

- 1. Replace compiler call with CATO wrapper script
- 2. (optional) Provide expert knowledge by adding CATO pragmas
- 3. LLVM frontend translates original code into IR
- 4. CATO analyses and transforms IR
- 5. LLVM backend translates IR into machine code
- 6. Execute binary via mpiexec or srun

#### **Benchmark Setup**

Partdiff



Figure 4: Example output of partdiff

- Partial differential equation solver
  - Jacobi method
  - (Gauß-Seidel method)
- Testmachine (2 nodes):
  - + 2  $\times$  Intel Xeon X5650 @ 2.67GHz
  - 6 cores per CPU with hyperthreading
  - 12GB RAM
  - 1Gb Ethernet
  - Ubuntu@18.04, Linux@4.15, MPICH@3.3.1

Jannek Squar

## **Scaling Behaviour**



Figure 5: Strong scaling

Figure 6: Weak scaling

Jannek Squar

Motivation

Design

Implementation

Outlook

Jannek Squar

Previous and current related topics (BSc/MSc):

- 2017: Initial tool implementation [Ble17]
- 2018: Improve implementation and consideration of patterns [Jam18]
- 2019: Static code analysis for MPI verification [HB19]
- TBD: Performance comparison of two-sided and one-sided MPI communication
- TBD: Static and dynamic code analysis to derive communication pattern

Memory Consumption and I/O backends

- · Avoid sequential data initialisation by master process
  - Avoid peak memory consumption before distribution
  - (Potential speedup by parallelised I/O phase)
- Limitation: Application must allocate memory itself
  - I/O library handles memory: CATO not feasible!
  - I/O receives pointer by application: CATO replaces I/O calls
- $\Rightarrow$  Implementation of NetCDF4 backend is ongoing

- Optimise memory chasing via user tree
- Reinsert OpenMP threads
- Automatise communication pattern recognition
  - Heuristics
  - Reuse well-established MPI versions of communication pattern
  - Make use of more advanced MPI constructs
- Move into public repository https://github.com/jsquar/cato

- Functional prototype of CATO
- Automatic replacement of OpenMP kernels through equivalent MPI operations
- Focus on increasing input size through distribution
  - Chasing and handling shared memory
  - Improve horizontal scaling

 $\Rightarrow$  Reach problem size beyond the possibility of OpenMP

#### **Table of Contents**

#### References

Backup Slides

#### References

- [ABC<sup>+</sup>06]
   Krste Asanovic, Ras Bodik, Bryan Christopher Catanzaro, Joseph James Gebis, Parry Husbands, Kurt Keutzer, David A Patterson,

   William Lester Plishker, John Shalf, Samuel Webb Williams, et al., The landscape of parallel computing research: A view from berkeley,

   Tech. report, Technical Report UCB/EECS-2006-183, EECS Department, University of California, Berkeley, 2006.
- [ACD<sup>+</sup>96] Cristiana Amza, Alan L Cox, Sandhya Dwarkadas, Pete Keleher, Honghui Lu, Ramakrishnan Rajamony, Weimin Yu, and Willy Zwaenepoel, Treadmarks: Shared memory computing on networks of workstations, Computer **29** (1996), no. 2, 18–28.
- [AT06] Dieter An Mey and Irene Tedjo, Adaptive Integration Form OpenMP to MPI.
- [BCJ01] Rajkumar Buyya, Toni Cortes, and Hai Jin, Single system image, Int. J. High Perform. Comput. Appl. 15 (2001), no. 2, 124–135.
- [BE05] Ayon Basumallik and Rudolf Eigenmann, Towards automatic translation of openmp to MPI, Proceedings of the 19th Annual International Conference on Supercomputing, ICS 2005, Cambridge, Massachusetts, USA, June 20-22, 2005 (Arvind and Larry Rudolph, eds.), ACM, 2005, pp. 189–198.
- [Ble17] Michael Blesel, Compiler assisted translation of OpenMP to MPI using LLVM, Online https://wr.informatik.uni-hamburg.de/\_media/research:theses: michael\_blesel\_compiler\_assisted\_translation\_of\_openmp\_to\_mpi\_using\_llvm.pdf, 10 2017.
- [Ble21] Michael Blesel, Message passing safety and correctness checks at compile time using Rust, Master's thesis, Universität Hamburg, 02 2021.
- [Col04] Phillip Colella, Defining software requirements for scientific computing.
- [HB19] Marcel Heing-Becker, Verification of one-sided MPI communication code using static analysis in LLVM, Master's thesis, Universität Hamburg, 02 2019.
- [HDT<sup>+</sup>15] Torsten Hoefler, James Dinan, Rajeev Thakur, Brian Barrett, Pavan Balaji, William Gropp, and Keith Underwood, *Remote Memory Access Programming in MPI-3*, ACM Transactions on Parallel Computing **2** (2015), no. 2, 1–26.

#### References

- [HFE11] Khaled Hamidouche, Joel Falcou, and Daniel Etiemble, A framework for an automatic hybrid mpi+openmp code generation, 2011 Spring Simulation Multi-conference, SpringSim '11, Boston, MA, USA, April 03-07, 2011. Volume 6: Proceedings of the 19th High Performance Computing Symposia (HPC) (Layne T. Watson, Gary W. Howell, William I. Thacker, and Steven Seidel, eds.), SCS/ACM, 2011, pp. 48–55.
- [Hoe06] Jay P. Hoeflinger, Extending openmp to clusters, White Paper, Intel Corporation (2006).
- [Jam18] Tim Jammer, Characterization and translation of OpenMP use cases to MPI using LLVM, Master's thesis, Universität Hamburg, 12 2018.
- [Lat] Chris Lattner, The architecture of open source applications, http://www.aosabook.org/en/llvm.html.
- [LLV18] LLVM Project, Writing an LLVM Pass, 2018.
- [LSB07] Jinpil Lee, Mitsuhisa Sato, and Taisuke Boku, Design and implementation of openmpd: An openmp-like programming language for distributed memory systems, A Practical Programming Model for the Multi-Core Era, 3rd International Workshop on OpenMP, IWOMP 2007, Beijing, China, June 3-7, 2007, Proceedings (Barbara M. Chapman, Weimin Zheng, Guang R. Gao, Mitsuhisa Sato, Eduard Ayguadé, and Dongsheng Wang, eds.), Lecture Notes in Computer Science, vol. 4935, Springer, 2007, pp. 143–147.
- [MMPS08] Daniel Millot, Alain Muller, Christian Parrot, and Frédérique Silber-Chaussumier, STEP: A distributed openmp for coarse-grain parallelism tool, OpenMP in a New Era of Parallelism, 4th International Workshop, IWOMP 2008, West Lafayette, IN, USA, May 12-14, 2008, Proceedings (Rudolf Eigenmann and Bronis R. de Supinski, eds.), Lecture Notes in Computer Science, vol. 5004, Springer, 2008, pp. 83–99.
- [Pad11] David A. Padua (ed.), *Encyclopedia of parallel computing*, Springer, 2011.
- [Sam] Adrian Sampson, Llvm for grad students, http://www.cs.cornell.edu/~asampson/blog/llvm.html.
- [SCC15] Albert Saà-Garriga, David Castells-Rufas, and Jordi Carrabina, OMP2MPI: automatic MPI code generation from openmp programs, CoRR abs/1502.02921 (2015).
- [Xca] XcalableMP Specification Working Group, Xcalablemp language specification.

#### **Table of Contents**

References

Backup Slides

## **Considered Communication Patterns**

- Classification[Jam18] of OpenMP kernels based on 7+6 dwarves
  - · Important patterns for science and engineering
  - Similar communication and data movement patterns
  - Similar computation patterns
- ightarrow Provide Equivalence Class templates
  - Preserve semantics of OpenMP kernel, equivalent behaviour
  - One-sided MPI-3 communication operations
  - Written in C++
  - Shared object to load during tool-runtime

#### 7+6 dwarves

## Original dwarves: [Col04]

- Dense Linear Algebra
- Sparse Linear Algebra
- Spectral Methods
- N-Body Methods
- Structured Grids
- Unstructured Grids
- Monte Carlo

## Additional dwarves: [ABC<sup>+</sup>06]

- Combinational Logic
- Graph Traversal
- Dynamic Programming
- Backtrack & Branch and Bound
- Graphical Models
- Finite State Machine

## Workflow - CATO (I)

- 1. MPI initialisation
- 2. Memory allocation
- 3. Memory initialisation
- 4. Microtask execution
- 5. Memory interaction
- 6. MPI finalisation

```
1 int main()
2 {
       int* a = (int*)malloc(sizeof(int)*4);
 3
       int* b = (int*)malloc(sizeof(int)*4):
       a[0] = 0;
       a[1] = 1:
       a[2] = 2:
 8
 9
       a[3] = 3:
10
       #pragma omp parallel for
11
12
13
            for(int i = 0; i < 4; i + +)
14
15
                b[i] = a[i]:
16
17
18
        printf("[%d,%d,%d]\n", b[0], b[1], b[2], b[3]);
19
20
       free(a):
21
       free(b):
22
23 }
```

Listing 2: Original code

### **Step 1: MPI Initialisation**

```
14 define dso_local i32 @main() #0 {
15 entry:
16 %a = alloca i32*, align 8
17 %b = alloca i32*, align 8
18 %call = call noalias i8* @malloc(i64 16) #4
19 %0 = bitcast i8* %call to i32*
20 store i32* %0, i32** %a, align 8
```

Listing 3: Original IR

```
14 define dso local i32 @main() #0 {
15 entry:
    \%0 = alloca i32
16
    \%1 = alloca i 32
17
    \% = alloca i32
18
    \%3 = alloca i32
10
    \%4 = alloca i32
20
    \%5 = alloca i32
21
    \%6 = alloca i32
22
    \%7 = alloca i 32
23
24
    \%8 = alloca i32
    call void @ Z15cato initializeb(i1 ↔
25
         \hookrightarrow false)
    br label %entry.split
26
27
28 entry.split:
    %a = alloca i32*, align 8
29
    %b = alloca i32*, align 8
30
```

Listing 4: Transformed IR

#### **Step 2: Memory Allocation**

14 define dso\_local i32 @main() #0 {

- 15 entry:
- 16 %a = alloca i32\*, align 8
- 17 %b = alloca i32\*, align 8
- 18 %call = call noalias i8\* @malloc(i64 16) #4
- 19 %0 = bitcast i8\* %call to i32\*
- 20 store i32\* %0, i32\*\* %a, align 8
- 21 %call1 = call noalias i8\* @malloc(i64 16) #4
- 22 %1 = bitcast i8\* %call1 to i32\*
- 23 store i32\* %1, i32\*\* %b, align 8
- 24 %2 = load i32\*, i32\*\* %a, align 8

Listing 5: Original IR

```
28 entry.split:
29
    %a = alloca i32*, align 8
    %b = alloca i32*, align 8
30
    %call = call i8* ↔
31
        \hookrightarrow 16. i32 1275069445, i32 1)
    \%9 =  bitcast i8* %call to i32*
32
    store i32* %9, i32** %a, align 8
33
  %call1 = call i8* ↔
34
        \leftrightarrow a Z22allocate shared memorylii(i64 \leftrightarrow
        \leftrightarrow 16, i32 1275069445, i32 1)
    %10 = hitcast i8* %call1 to i32*
35
    store i32* %10. i32** %b. align 8
36
    %11 = load i32*. i32** %a. align 8
37
```

Listing 6: Transformed IR

#### **Step 3: Memory Initialisation**

#### Backup Slides | Example

%arravidx = getelementptr inbounds i32, i32\* %2, i64 0 25 store i32 0, i32\* %arravidx, align 4 26 %3 = load i32\*, i32\*\* %a, align 8 27 %arravidx2 = getelementptr inbounds i32, i32\* %3, i64 1 28 store i32 1. i32\* %arravidx2. align 4 29 %4 = load i32\*, i32\*\* %a, align 8 30 %arrayidx3 = getelementptr inbounds i32, i32\* %4, i64 2 31 store i32 2. i32\* %arravidx3. align 4 32 %5 = load i32\*. i32\*\* %a. align 8 33 %arravidx4 = getelementptr inbounds i32. i32\* %5. i64 3 34 store i32 3. i32\* %arravidx4. align 4 35

Listing 7: Original IR

%arravidx = getelementptr inbounds i32, i32\* %11, i64 0 38 %12 = bitcast i32\* %11 to i8\* 20 store i32 0, i32\* %0 40 %13 = bitcast i32\* %0 to i8\* 61 call void (i8\*, i8\*, i32, ...) ↔ ↔ @ Z30shared memory sequential storePvS iz(i8\* %12, i8\* %13, ↔  $\rightarrow$  i32 1, i64 0) %14 = load i32\*, i32\*\* %a, align 8 44 %arrayidx2 = getelementptr inbounds i32, i32\* %14, i64 1 %15 = bitcast i32\* %14 to i8\* 45 store i32 1, i32\* %1 46 %16 = bitcast i32\* %1 to i8\* 47 call void (i8\*, i8\*, i32, ...) ↔ 48  $\leftrightarrow$  a Z30shared memory sequential storePvS iz(i8\* %15, i8\* %16,  $\leftrightarrow$  $\leftrightarrow$  i32 1, i64 1) %17 = load i32\*, i32\*\* %a, align 8 %arravidx3 = getelementptr inbounds i32, i32\* %17, i64 2 50 %18 = bitcast i32\* %17 to i8\* 51 store i32 2. i32\* %2 52 %19 = bitcast i32\* %2 to i8\* 52 call void (i8\*, i8\*, i32, ...) ↔ 54 ↔ @ Z30shared memory sequential storePyS iz(i8\* %18, i8\* %19, ↔  $\leftrightarrow$  i32 1, i64 2) %20 = load i32\*. i32\*\* %a. align 8 55 %arravidx4 = getelementptr inbounds i32, i32\* %20, i64 3 56 %21 = bitcast i32\* %20 to i8\* 57 store i32 3. i32\* %3 58 %22 = hitcast i32\* %3 to i8\* 50 call void (i8\*, i8\*, i32, ...) ↔ 60 ↔ @\_Z30shared\_memory\_sequential\_storePvS\_iz(i8\* %21, i8\* %22, ↔)  $\leftrightarrow$  i32 1, i64 3)

Listing 8: Transformed IR

Listing 9: Original IR

call void @.omp\_outlined.(i32\* null, i32\* null, i32\*\* %b, i32\*\* %a)
call void @\_Z11mpi\_barrierv()

Listing 10: Transformed IR

#### Step 4: Microtask Execution (II)

```
store i32 0, i32* %.omp.lb, align 4
82
    store i32 3. i32* %.omp.ub. align 4
83
    store i32 1. i32* %.omp.stride. align 4
84
    store i32 0. i32* %.omp.is last. align 4
85
    %2 = load i32*, i32** %.global tid.,addr, align 8
86
    %3 = load i32, i32* %2, align 4
87
    call void @ kmpc for static init 4(%struct.ident t* @0, i32 %3, i32 34, i32* ↔
88
         \hookrightarrow %.omp.is last, i32* %.omp.lb, i32* %.omp.ub, i32* %.omp.stride, i32 1, \leftrightarrow
         \rightarrow i32 1)
```

Listing 11: Original IR

```
130 store i32 0, i32* %.omp.lb, align 4
131 store i32 3, i32* %.omp.ub, align 4
132 store i32 1, i32* %.omp.stride, align 4
133 store i32 0, i32* %.omp.is_last, align 4
134 %2 = load i32*, i32** %.global_tid..addr, align 8
135 %3 = load i32, i32* %2, align 4
136 call void ∂_Z26modify_parallel_for_boundsPiS_i(i32* %.omp.lb, i32* %.omp.ub, ↔
132 1)
```

Listing 12: Transformed IR

#### **Step 5: Memory Interaction**

#### **Backup Slides | Example**

- 37
   %6 = load i32\*, i32\*\* %b, align 8

   38
   %arrayidx5 = getelementptr inbounds i32, i32\* %6, i64 0

   39
   %7 = load i32, i32\* %arrayidx5, align 4

   40
   %8 = load i32\*, i32\*\* %b, align 8

   41
   %arrayidx6 = getelementptr inbounds i32, i32\* %8, i64 1
- 42 %9 = load i32, i32\* %arrayidx6, align 4
- 43 %10 = load i32\*, i32\*\* %b, align 8
- 44 %arrayidx7 = getelementptr inbounds i32, i32\* %10, i64 2
- 45 %11 = load i32, i32\* %arrayidx7, align 4
- 46 %12 = load i32\*, i32\*\* %b, align 8
- 47 %arrayidx8 = getelementptr inbounds i32, i32\* %12, i64 3
- 48 %13 = load i32, i32\* %arrayidx8, align 4
- 49 %call9 = call i32 (i8\*, ...) @printf(i8\* getelementptr ↔
  - $\hookrightarrow$  inbounds ([15 x i8], [15 x i8]\* @.str.1, i64 0, i64  $\leftrightarrow$ 
    - $\hookrightarrow$  0), i32 %7, i32 %9, i32 %11, i32 %13)

Listing 13: Original IR

call void (i8\*, i8\*, i32, ...) ↔ 81 ↔ @ Z29shared memory sequential loadPvS iz(i8\* %34. i8\* %35. i32 ↔  $\leftrightarrow$  1, i64 2) %36 = bitcast i8\* %35 to i32\* %37 = load i32, i32\* %36 %38 = load i32\*, i32\*\* %b, align 8 84 %arravidx8 = getelementptr inbounds i32, i32\* %38, i64 3 20 %39 = bitcast i32\* %38 to i8\* %40 = bitcast i32\* %7 to i8\* call void (i8\*, i8\*, i32, ...) ↔ → @ Z29shared memory sequential loadPvS iz(i8\* %39, i8\* %40, i32 ↔  $\leftrightarrow$  1, i64 3) %41 = bitcast i8\* %40 to i32\* %42 = load i32, i32\* %41 90 %call9 = call i32 (i8\*. ...) @printf(i8\* getelementptr inbounds ([15 ↔ 91 → x i8], [15 x i8]\* @.str.1, i64 0, i64 0), i32 %27, i32 %32, ↔  $\rightarrow$  i32 %37, i32 %42)

Listing 14: Transformed IR

#### **Step 6: MPI Finalisation**

```
50 %14 = load i32*, i32** %a, align 8
51 %15 = bitcast i32* %14 to i8*
52 call void @free(i8* %15) #4
53 %16 = load i32*, i32** %b, align 8
54 %17 = bitcast i32* %16 to i8*
55 call void @free(i8* %17) #4
56 ret i32 0
```

Listing 15: Original IR

```
92 %43 = load i32*, i32** %a, align 8
```

```
93 %44 = bitcast i32* %43 to i8*
```

```
94 call void @_Z18shared_memory_freePv(i8* %44)
```

```
95 %45 = load i32*, i32** %b, align 8
```

```
96 %46 = bitcast i32* %45 to i8*
```

```
97 call void @_Z18shared_memory_freePv(i8* %46)
```

```
98 store i32 0, i32* %8
```

```
99 br label %cato_finalize
```

```
100
```

```
101 cato_finalize:
```

```
102 call void @_Z13cato_finalizev()
```

```
103 %47 = load i32, i32* %8
```

```
104 ret i32 %47
```

Listing 16: Transformed IR

## LLVM Function Pass[LLV18]

```
#include "llvm/Pass.h"
 1
    #include "llvm/IR/Function.h"
 2
    #include "llvm/Support/raw ostream.h"
 3
 4
 5
     using namespace llvm;
6
    namespace
 7
         struct printFunction : public FunctionPass
 8
 9
             static char ID:
10
11
             printFunction() : FunctionPass(ID) {}
12
             bool runOnFunction(Function &F) override
13
14
                 errs() << "Hello: ":
15
                 errs().write_escaped(F.getName()) << '\n';</pre>
16
                 return false;
17
18
19
         }:
20
21
22
    char printFunction:: ID = 0:
     static RegisterPass<printFunction> X("hello", "Hello World Pass", false, false);
23
```

Scopeparallel, single, master, task, sections, forClausesnum\_threads, private, firstprivate, threadprivate, shared,<br/>reductionSynchronisationbarrier, critical, atomicSchedulingstatic, dynamic, guided, auto, runtimeFunctionsget\_thread\_num, get\_num\_threads<br/>MiscMisctarget, simd

<sup>&</sup>lt;sup>1</sup>colour key: not planned, planned, work in progress

#### **Extended Weak Scaling**

#### Backup Slides | LLVM



Figure 7: Using 10 nodes for weak scaling [Jam18]

### Active Target Synchronisation (I)

#### Backup Slides | LLVM



Figure 8: Collective Fence Synchronisation [HDT+15]

### Active Target Synchronisation (II)

#### Backup Slides | LLVM



Figure 9: GATS Synchronisation [HDT+15]

### Project-Example - SAGA-GIS (I)

#### **Backup Slides | LLVM**



## Project-Example - SAGA-GIS (II)

#### Backup Slides | LLVM



Downscaling

27kmightarrow1km



### Project-Example - SAGA-GIS (III)

```
1 for(int y=0; y<Get_NY() && Set_Progress(y); y++)
2 {
3 #pragma omp parallel for
4 for(int x=0; x<Get_NX(); x++)
5 {
6 int Belt = 0; // no data
7 if( !pL->is_NoData(x, y) )
8 {
9 [...]
10 }
11 pBelt->Set_Value(x, y, Belt);
12 }
```

Listing 17: "Thermic Belt Classification"

Total occurrences in all SAGA-GIS tools:

- 322× **#pragma** omp parallel **for** 
  - $\hookrightarrow$  [private(...)] [reduction(...)]
  - 1× **#pragma** omp critical