

Extending the LLVM/Clang Framework for OpenMP Metadirective Support

Seminar Supercomputer: Forschung und Innovation

Marcel Robohm

Arbeitsbereich Wissenschaftliches Rechnen
Fachbereich Informatik
Fakultät für Mathematik, Informatik und Naturwissenschaften
Universität Hamburg

2021-12-14

Gliederung

- 1 Motivation
- 2 Grundlagen
- 3 Metadirektiven
- 4 Dynamische Direktiven
- 5 Fazit
- 6 Literatur

Motivation [MMC20]

- Code besser portierbar
 - Dynamisch Konfiguration nach Hardware bestimmen
- Erfüllung des OpenMP-5.0-Standards [ea18]

Was ist OpenMP?

- Parallelisierung mit geteiltem Speicher
- Direkt in C(++)-Code mittels `#pragma omp`
- Zusätzlich Prozeduren in `<omp.h>`

Beispiel

```
1  size_t length = 1000;
2  int* output = new int[length];
3
4
5  for (int i = 0; i < length; i++) {
6      output[i] = heavyCalculation(i);
7  }
```

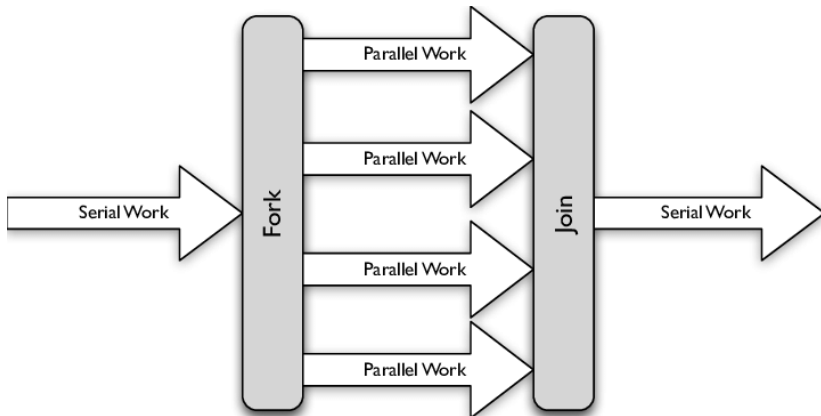
Einfache for-Schleife

Beispiel

```
1  size_t length = 1000;
2  int* output = new int[length];
3
4  #pragma omp parallel for
5  for (int i = 0; i < length; i++) {
6      output[i] = heavyCalculation(i);
7  }
```

Einfache for-Schleife mit Annotation

Beispiel - Fork-Join [Bec15]

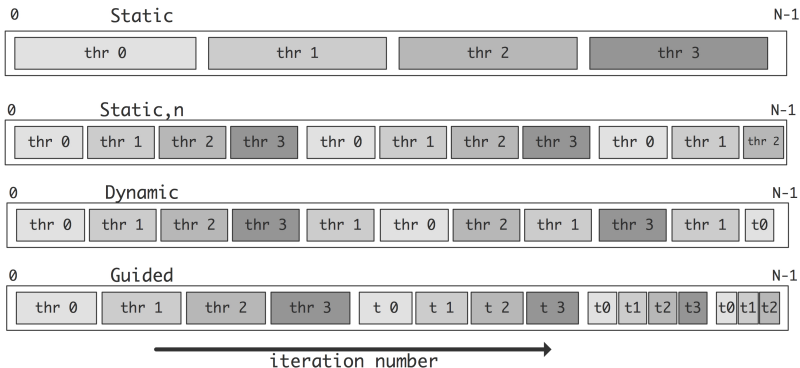


Beispiel


```
1  size_t length = 1000;
2  int* output = new int[length];
3
4  #pragma omp parallel for schedule(dynamic, 1)
5  for (int i = 0; i < length; i++) {
6      output[i] = heavyCalculation(i);
7  }
```

Einfache for-Schleife mit Annotation und explizitem Scheduling

Scheduling - Schema [EvdGC16]




Scheduling - Beispiel [ZuAD]



Auslagerung

- Auslagerung (Offloading) zur Steigerung der Leistung durch Spezialhardware
- Zuweisung von Teams

OpenMP - Teams [Lar18]



Auslagerung - Beispiel


```
1  size_t length = 1000;
2  int* output = new int[length];
3  int* inputA = new int[length];
4  int* inputB = new int[length];
5
6  #pragma omp target teams distribute parallel for
   ↪ map(to:inputA,inputB), map(from:output)
7  for (int i = 0; i < length; i++) {
8      output[i] = inputA[i] + inputB[i];
9  }
```

Auslagerung der Rechenlast mittels target

Herausforderung

- Heterogene Hardware
 - Speicher vs. Prozessorleistung
- Spezialhardware
 - GPU
 - FPGA
- Heterogene Einschränkungen (Stromverbrauch)

Adaption notwendig! [BCM⁺16]



Beispiel

```
1  int v1[N], v2[N], v3[N];
2  #if defined(nvptx)
3      #pragma omp target teams distribute \
4          parallel for map(to:v1,v2) map(from:v3)
5  for (int i = 0; i < N; i++)
6      v3[i] = v1[i] * v2[i];
7  #else
8      #pragma omp target \
9          parallel for map(to:v1,v2) map(from:v3)
10 for (int i = 0; i < N; i++)
11     v3[i] = v1[i] * v2[i];
12 #endif
```

Adaption ohne Metadirektiven

Beispiel

```
1 int v1[N], v2[N], v3[N];
2 #pragma omp target data map(to:v1,v2) map(from:v3)
3 #pragma omp metadirective \
4     when(device=arch("nvptx"): target teams
5         ↪ distribute parallel for) \
6     default(target parallel for)
7 for(int i = 0; i < N; i++)
   v3[i] = v1[i] * v2[i];
```

Adaption mit Metadirektiven

Syntax

```
#pragma omp metadirective [metclause[[,] metclause]
    ↪ ... ] new-line \
```

Metadirektive

```
when(context-selector-specification:
    ↪ [directive-variant]) \
default (directive-variant) \
```

metclause

```
directive-name [clause[ [,] clause] ... ]
```

directive-variant

context-selector-specification

- Gerät
 - Art (host, nohost, cpu, gpu, fpga)
 - Befehlssatzarchitektur (x64, x86)
 - Architektur (nvptx, nvptx64, gcn)
- Implementation
 - Anbieter (nvidia, amd)
 - Erweiterungen

Syntax

```
1  for (idev=0; idev < omp_get_num_devices(); idev++)
2      #pragma omp target device(idev)
3      #pragma omp metadirective \
4          when(implementation={vendor(nvidia)}),
5              ↪ device={arch("kepler")}: \
6                  teams num_teams(512) thread_limit(32) ) \
7          when(implementation={vendor(amd)}),
8              ↪ device={arch("fiji")}: \
9                  teams num_teams(512) thread_limit(64) ) \
10         default(teams)
11     #pragma omp distribute parallel for
12     for(i = 0; i < N; i++)
13         work_on_chunk(idev, i);
```

Umfangreiches Beispiel

Dynamische Adaption (bisher)

```
1  if(N > 1000) {
2      #pragma omp target teams distribute parallel for
3      for(int i = 0; i < N; i++)
4          heavyCalculation(i);
5  } else if(N > 100 && N <= 1000) {
6      #pragma omp parallel for
7      for(int i = 0; i < N; i++)
8          heavyCalculation(i);
9  } else {
10     for(int i = 0; i < N; i++)
11         heavyCalculation(i);
12 }
```



Dynamische Adaption ohne Metadirektiven

Dynamische Adaption (neu)


```
1 #pragma omp metadirective \  
2   when(user={condition(N > 1000)}: target teams  
   ↪ distribute parallel for) \  
3   when(user={condition(N > 100 && N <= 1000)}:  
   ↪ parallel for) \  
4   default()  
5 for(int i = 0; i < N; i++)  
6   heavyCalculation(i);
```

Dynamische Adaption mit Metadirektiven


Speicherbedarf



Performanz



a: A Summit node with 2 IBM Power9 CPUs (128 threads) and NVIDIA Volta V100 GPU



b: A Seawulf node with 2 Intel Xeon E5-2683v3 CPUs (28 threads) and NVIDIA Tesla K80 GPU

Fazit

- Erfüllung des OpenMP 5.0 Standards
 - Statische Metadirektiven
- Einführung dynamischer Metadirektiven

Literatur

- [BCM⁺16] Md Abdullah Shahneous Bari, Nicholas Chaimov, Abid M Malik, Kevin A Huck, Barbara Chapman, Allen D Malony, and Osman Sarood. Arcs: Adaptive runtime configuration selection for power-constrained openmp applications. In *2016 IEEE international conference on cluster computing (CLUSTER)*. IEEE, 2016.
- [Bec15] David Beckingsale. *Towards Scalable Adaptive Mesh Refinement on Future Parallel Architectures*. PhD thesis, 01 2015.
- [ea18] O. Consortium et al. Openmp specification version 5.0. 2018.
- [EvdGC16] Victor Eijkhout, Robert van de Geijn, and Edmond Chow. *Introduction to High Performance Scientific Computing*. Zenodo, Apr 2016.
- [Lar18] Jeff Larkin. Openmp on gpus, first experiences and best practices. *NVIDIA GTC*, Mar 2018.
- [MMC20] Alok Mishra, Abid M Malik, and Barbara Chapman. Extending the llvm/clang framework for openmp metadirective support. In *2020 IEEE/ACM 6th Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC) and Workshop on Hierarchical Parallelism for Exascale Computing (HiPar)*. IEEE, 2020.
- [ZulAD] Zulan. Openmp dynamic vs guided scheduling. *Stackoverflow*, Mar 27AD.