

The 11th International workshop on OpenCL and SYCL

IWOCL & SYCLcon 2023



Implementation Techniques for SPMD Kernels on CPUs

Joachim Meyer (Saarland University)

Aksel Alpay, Holger Fröning, Vincent Heuveline (Heidelberg University)
Sebastian Hack (Saarland University)



CONTENT

- Recap: what are SPMD kernels?



CONTENT

- Recap: what are SPMD kernels?
- 4 approaches to implementing SPMD kernels on CPUs:



CONTENT

- Recap: what are SPMD kernels?
- 4 approaches to implementing SPMD kernels on CPUs:
 - library-only:
 - 1 thread : 1 work-item
 - 1 fiber : 1 work-item

CONTENT

- Recap: what are SPMD kernels?
- 4 approaches to implementing SPMD kernels on CPUs:
 - library-only:
 - 1 thread : 1 work-item
 - 1 fiber : 1 work-item
 - compiler supported:
 - deep loop fission (DLF)
 - continuation-based synchronization (CBS)

CONTENT

- Recap: what are SPMD kernels?
- 4 approaches to implementing SPMD kernels on CPUs:
 - library-only:
 - 1 thread : 1 work-item
 - 1 fiber : 1 work-item
 - compiler supported:
 - deep loop fission (DLF)
 - continuation-based synchronization (CBS)
- Semantic flaw in DLF

CONTENT

- Recap: what are SPMD kernels?
- 4 approaches to implementing SPMD kernels on CPUs:
 - library-only:
 - 1 thread : 1 work-item
 - 1 fiber : 1 work-item
 - compiler supported:
 - deep loop fission (DLF)
 - continuation-based synchronization (CBS)
- Semantic flaw in DLF
- An important optimization: propagating contiguity

CONTENT

- Recap: what are SPMD kernels?
- 4 approaches to implementing SPMD kernels on CPUs:
 - library-only:
 - 1 thread : 1 work-item
 - 1 fiber : 1 work-item
 - compiler supported:
 - deep loop fission (DLF)
 - continuation-based synchronization (CBS)
- Semantic flaw in DLF
- An important optimization: propagating contiguity
- Performance results: library-only < DLF \approx CBS

SINGLE-PROGRAM MULTIPLE-DATA KERNEL



```
1 cgh.parallel_for(sycl::nd_range<1>{global_size, group_size},
2   [=](sycl::nd_item<1> item) noexcept {
3     const auto lid = item.get_local_id(0);
4     scratch[lid] = acc[item.get_global_id()];
5     for(size_t i = group_size / 2; i > 0; i /= 2) {
6       item.barrier();
7       if(lid < i) scratch[lid] += scratch[lid + i];
8     }
9
10    if(lid == 0) acc[item.get_global_id()] = scratch[lid];
11  });
```

SINGLE-PROGRAM MULTIPLE-DATA KERNEL



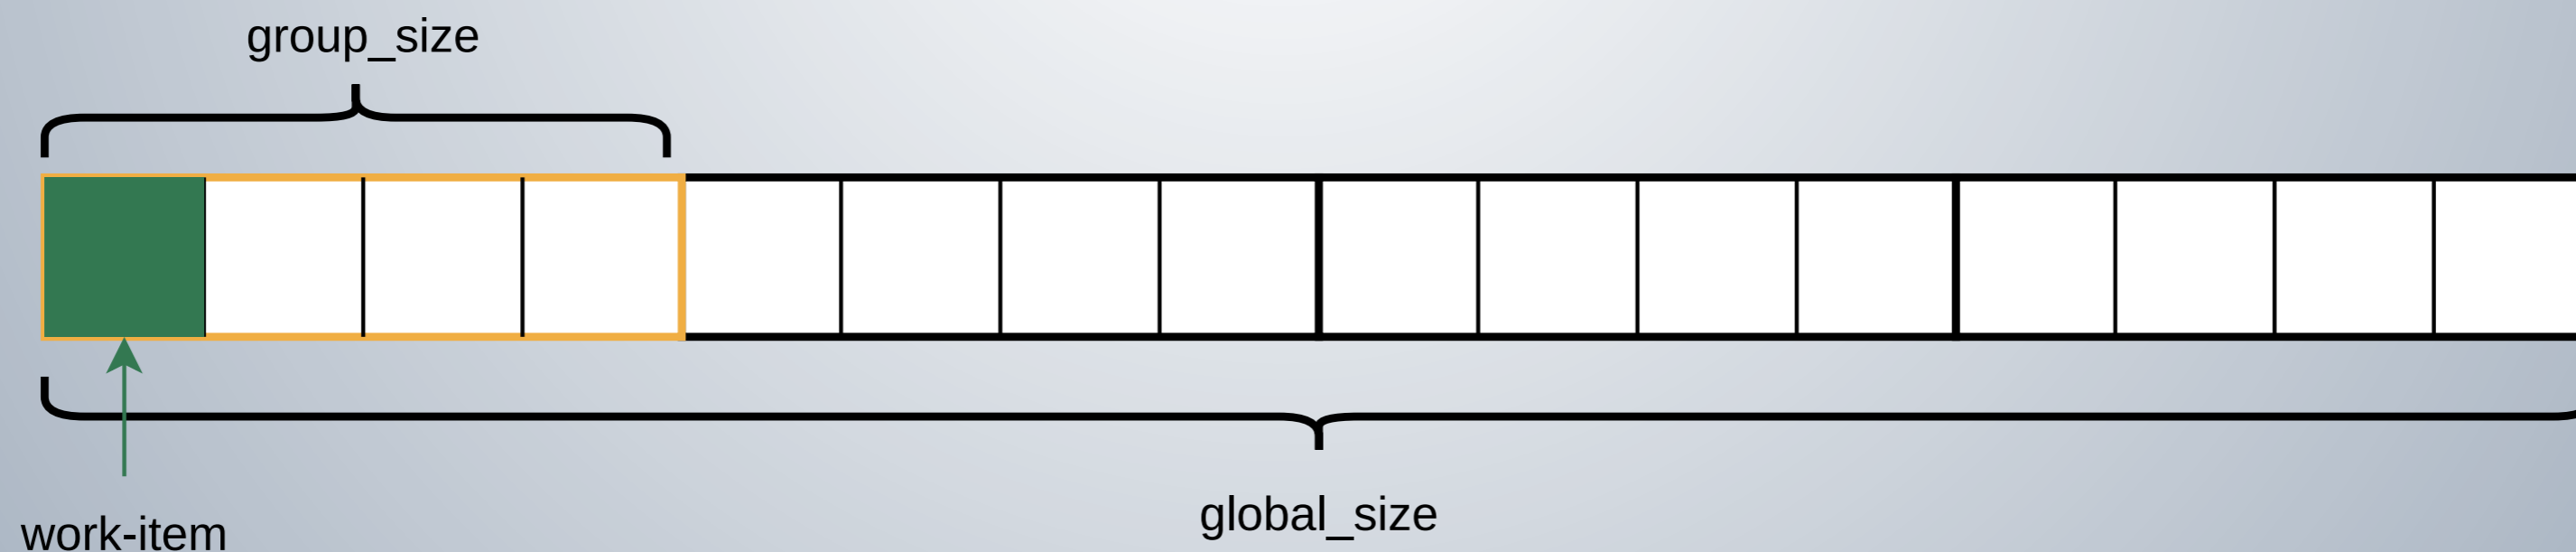
```
1 cgh.parallel_for(sycl::nd_range<1>{global_size, group_size},
2 [=](sycl::nd_item<1> item) noexcept {
3     const auto lid = item.get_local_id(0);
4     scratch[lid] = acc[item.get_global_id()];
5     for(size_t i = group_size / 2; i > 0; i /= 2) {
6         item.barrier();
7         if(lid < i) scratch[lid] += scratch[lid + i];
8     }
9
10    if(lid == 0) acc[item.get_global_id()] = scratch[lid];
11 });
```



work-item

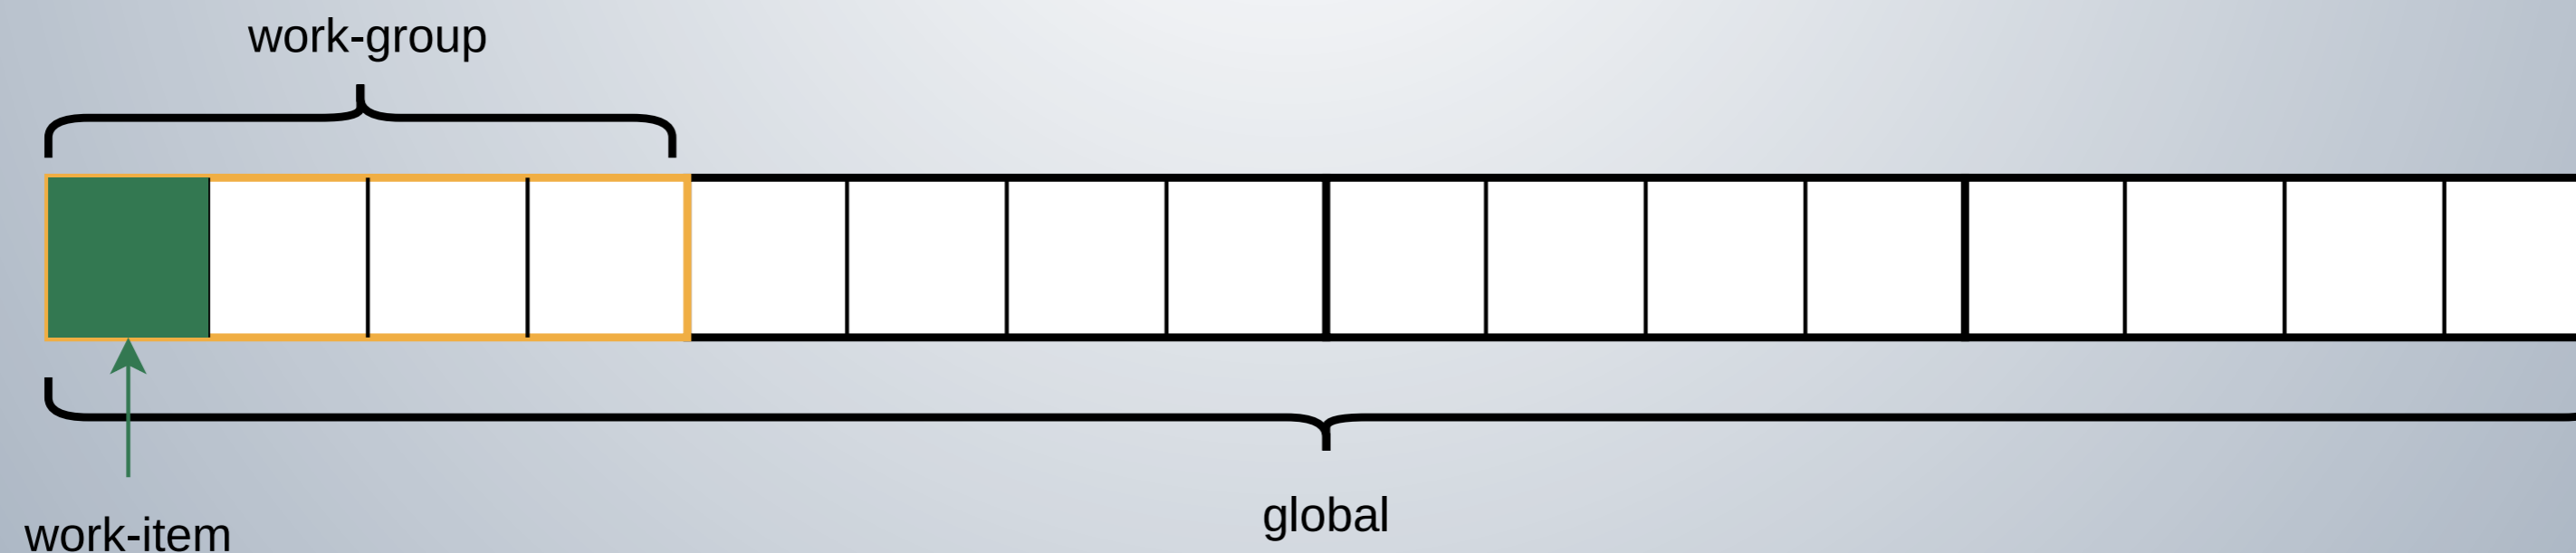
SINGLE-PROGRAM **MULTIPLE-DATA** KERNEL

```
1 cgh.parallel_for(sycl::nd_range<1>{global_size, group_size},
2 [=](sycl::nd_item<1> item) noexcept {
3     const auto lid = item.get_local_id(0);
4     scratch[lid] = acc[item.get_global_id()];
5     for(size_t i = group_size / 2; i > 0; i /= 2) {
6         item.barrier();
7         if(lid < i) scratch[lid] += scratch[lid + i];
8     }
9
10    if(lid == 0) acc[item.get_global_id()] = scratch[lid];
11 });
```



THE BIG BLOCKER ON CPUS

```
1 cgh.parallel_for(sycl::nd_range<1>{global_size, group_size},
2   [=](sycl::nd_item<1> item) noexcept {
3     const auto lid = item.get_local_id(0);
4     scratch[lid] = acc[item.get_global_id()];
5     for(size_t i = group_size / 2; i > 0; i /= 2) {
6       item.barrier();
7       if(lid < i) scratch[lid] += scratch[lid + i];
8     }
9
10    if(lid == 0) acc[item.get_global_id()] = scratch[lid];
11  });
```





THIS IS "SIMPLE" ON GPUS

- Execution of many (mostly) independent threads
→ Forward-Progress guarantees
- Hardware support for work-group barriers



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

hipSYCL



UNIVERSITÄT
DES
SAARLANDES

SIC Saarland Informatics
Campus

HOW TO MAP THIS TO CPUS?



CONCURRENCY!

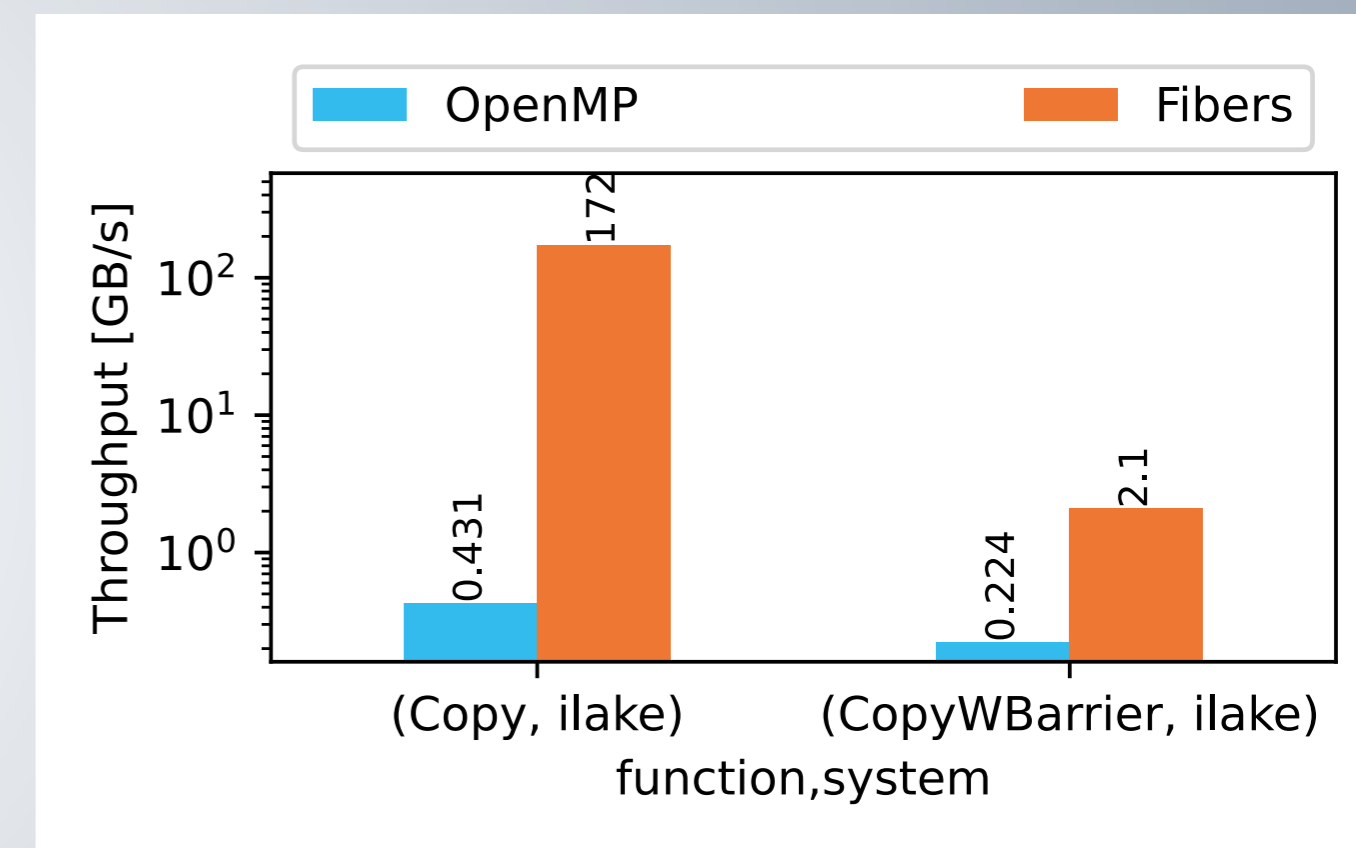
- 1 work-item : 1 thread
 - E.g. OpenMP parallel for + #pragma omp barrier
 - ⚡ Many threads → scheduling overhead
 - ⚡ No vectorization across work-items

CONCURRENCY!

- 1 work-item : 1 thread
 - E.g. OpenMP parallel for + #pragma omp barrier
 - ⚡ Many threads → scheduling overhead
 - ⚡ No vectorization across work-items
- 1 work-item : 1 fiber
 - Lightweight threads + synchronization
 - Can optimize barrier-free kernels!
 - ⚡ Context-switch overhead
 - ⚡ Limited vectorization across work-items

CONCURRENCY!

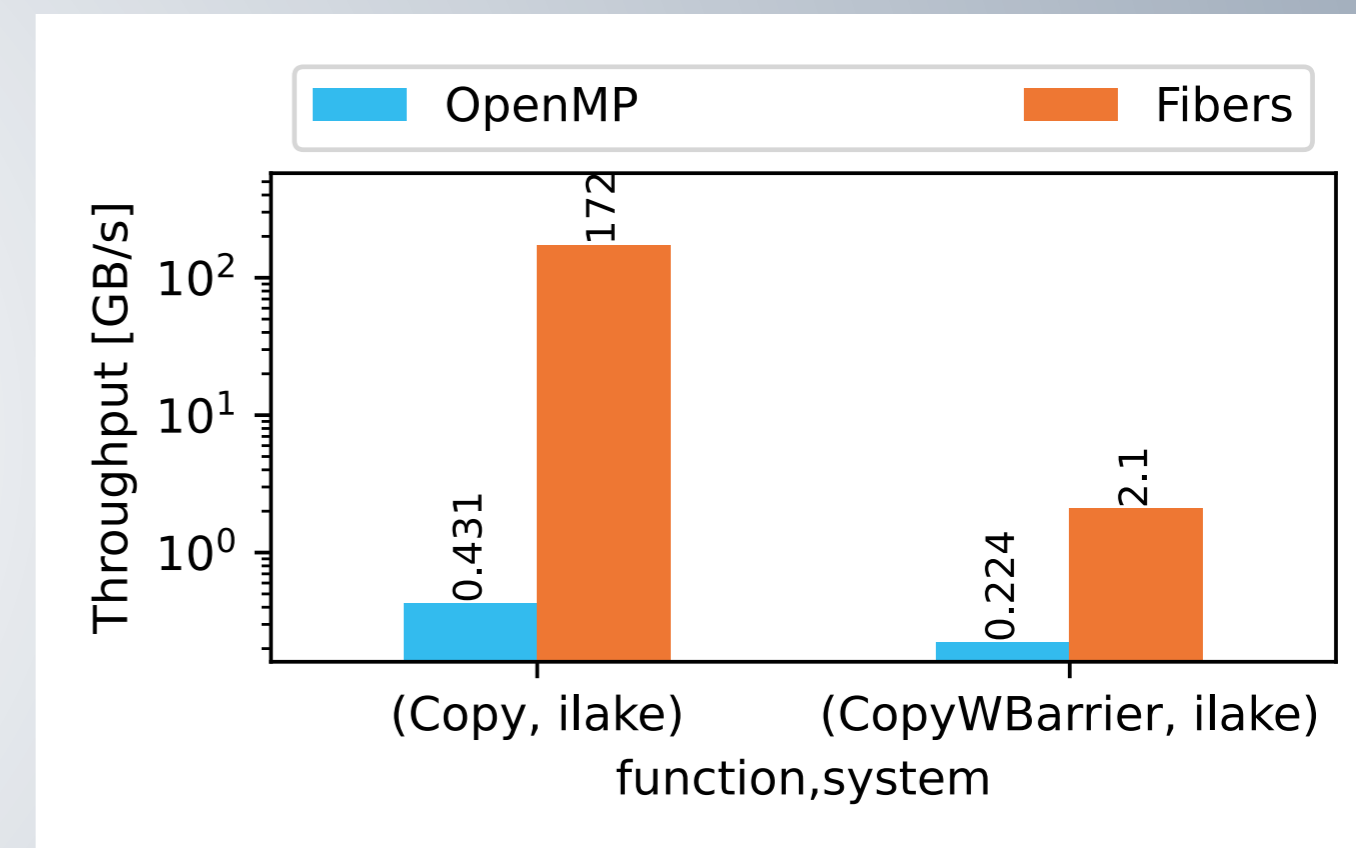
- 1 work-item : 1 thread
 - E.g. OpenMP parallel for + #pragma omp barrier
 - ⚡ Many threads → scheduling overhead
 - ⚡ No vectorization across work-items
- 1 work-item : 1 fiber
 - Lightweight threads + synchronization
 - Can optimize barrier-free kernels!
 - ⚡ Context-switch overhead
 - ⚡ Limited vectorization across work-items



github.com/UoB-HPC/BabelStream

CONCURRENCY!

- 1 work-item : 1 thread
 - E.g. OpenMP parallel for + #pragma omp barrier
 - ⚡ Many threads → scheduling overhead
 - ⚡ No vectorization across work-items
- 1 work-item : 1 fiber
 - Lightweight threads + synchronization
 - Can optimize barrier-free kernels!
 - ⚡ Context-switch overhead
 - ⚡ Limited vectorization across work-items



github.com/UoB-HPC/BabelStream

✓ Can be implemented without dedicated compiler support!

WANT PERFORMANCE? USE THE COMPILER!

- Threading on the **work-group level**
 - Loop over work-items in a single thread
 - Compiler extension to **split kernel** at barriers
- ✓ Vectorization across work-items possible
 - ✓ Improves performance over library-only by up to several orders of magnitude

```
#pragma omp parallel for
for(group : groups)
  #pragma omp simd
  for(item : itemsInGroup)
    kernel_before_barrier(nd_item{group, item})
  // implicit synchronization
  #pragma omp simd
  for(item : itemsInGroup)
    kernel_after_barrier(nd_item{group, item})
```

DEEP LOOP FISSION (POCL)

```
1 [=](sycl::nd_item<1> item) {
2     const auto lid = item.get_local_id(0);
3     scratch[lid] = acc[item.get_global_id()]; // A
4     item.barrier();
5     for(size_t i = group_size / 2; i > 0; i /= 2) {
6         if(lid < i) scratch[lid] += scratch[lid + i]; // B
7         item.barrier();
8     }
9
10    if(lid == 0) acc[item.get_global_id()] = scratch[lid]; // C
11 }
```

```
1 for(lid : items[0:])
2     // A
3     // barrier
4     for(i = group_size / 2; i > 0; i /= 2)
5         // B (lid = 0)
6         for(lid : items[1:])
7             // B
8             // barrier
9         for(lid : items[0:])
10            // C
```

DEEP LOOP FISSION (POCL)

```
1 [=](sycl::nd_item<1> item) {
2     const auto lid = item.get_local_id(0);
3     scratch[lid] = acc[item.get_global_id()]; // A
4     item.barrier();
5     for(size_t i = group_size / 2; i > 0; i /= 2) {
6         if(lid < i) scratch[lid] += scratch[lid + i]; // B
7         item.barrier();
8     }
9
10    if(lid == 0) acc[item.get_global_id()] = scratch[lid]; // C
11 }
```

```
1 for(lid : items[0:])
2     // A
3     // barrier
4     for(i = group_size / 2; i > 0; i /= 2)
5         // B (lid = 0)
6         for(lid : items[1:])
7             // B
8             // barrier
9         for(lid : items[0:])
10            // C
```

DEEP LOOP FISSION (POCL)

```
1 [=](sycl::nd_item<1> item) {
2     const auto lid = item.get_local_id(0);
3     scratch[lid] = acc[item.get_global_id()]; // A
4     item.barrier();
5     for(size_t i = group_size / 2; i > 0; i /= 2) {
6         if(lid < i) scratch[lid] += scratch[lid + i]; // B
7         item.barrier();
8     }
9
10    if(lid == 0) acc[item.get_global_id()] = scratch[lid]; // C
11 }
```

```
1 for(lid : items[0:])
2     // A
3     // barrier
4     for(i = group_size / 2; i > 0; i /= 2)
5         // B (lid = 0)
6         for(lid : items[1:])
7             // B
8             // barrier
9         for(lid : items[0:])
10            // C
```

DEEP LOOP FISSION (POCL)

```
1 [=](sycl::nd_item<1> item) {
2   const auto lid = item.get_local_id(0);
3   scratch[lid] = acc[item.get_global_id()]; // A
4   item.barrier();
5   for(size_t i = group_size / 2; i > 0; i /= 2) {
6     if(lid < i) scratch[lid] += scratch[lid + i]; // B
7     item.barrier();
8   }
9
10  if(lid == 0) acc[item.get_global_id()] = scratch[lid]; // C
11 }
```

```
1 for(lid : items[0:])
2   // A
3   // barrier
4   for(i = group_size / 2; i > 0; i /= 2)
5     // B (lid = 0)
6     for(lid : items[1:])
7       // B
8       // barrier
9   for(lid : items[0:])
10    // C
```

DEEP LOOP FISSION (POCL)

```
1 [=](sycl::nd_item<1> item) {
2     const auto lid = item.get_local_id(0);
3     scratch[lid] = acc[item.get_global_id()]; // A
4     item.barrier();
5     for(size_t i = group_size / 2; i > 0; i /= 2) {
6         if(lid < i) scratch[lid] += scratch[lid + i]; // B
7         item.barrier();
8     }
9
10    if(lid == 0) acc[item.get_global_id()] = scratch[lid]; // C
11 }
```

```
1 for(lid : items[0:])
2     // A
3     // barrier
4 for(i = group_size / 2; i > 0; i /= 2)
5     // B (lid = 0)
6     for(lid : items[1:])
7         // B
8         // barrier
9 for(lid : items[0:])
10    // C
```


DEEP LOOP FISSION – SEMANTIC PROBLEM

```
1 [=](sycl::nd_item<1> item) noexcept {
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier();
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier();
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 }
```

```
1 for(lid : items[0:])
2   // A
3   // barrier
4   for(i : [0,1])
5     // B (lid = 0)
6     for(lid : items[1:])
7       // B
8       if(i < 2)
9         // barrier
10    for(lid : items[0:])
11      // C
```

DEEP LOOP FISSION – SEMANTIC PROBLEM

```
1 [=](sycl::nd_item<1> item) noexcept {
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier();
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier();
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 }
```

```
1 for(lid : items[0:])
2   // A
3   // barrier
4   for(i : [0,1])
5     // B (lid = 0)
6     for(lid : items[1:])
7       // B
8       if(i < 2)
9         // barrier
10  for(lid : items[0:])
11    // C
```

CONTINUATION-BASED SYNCHRONIZATION TO THE RESCUE



```
1 [=](sycl::nd_item<1> item) noexcept { // 0
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier(); // 1
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier(); // 2
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 } // -1
```

Continuation-based Synchronization

```
1
2
3
4
5   case 0:
6
7       // A
8       // barrier
9
10  case 1:
11
12      i = 0
13      while(i < 2 + lid)
14          // B
15          if(i < 2) // barrier
16              i++;
17          // C
18
19
20  case 2:
21
22      i++;
23      while(i < 2 + lid)
24          // B
25          if(i < 2) // barrier
26              i++;
27          // C
28
29
30
31
```

CONTINUATION-BASED SYNCHRONIZATION TO THE RESCUE



```
1 [=](sycl::nd_item<1> item) noexcept { // 0
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier(); // 1
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier(); // 2
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 } // -1
```

Continuation-based Synchronization

```
1
2
3
4
5   case 0:
6
7     // A
8     // barrier
9
10    case 1:
11
12      i = 0
13      while(i < 2 + lid)
14        // B
15        if(i < 2) // barrier
16          i++;
17        // C
18
19
20    case 2:
21
22      i++;
23      while(i < 2 + lid)
24        // B
25        if(i < 2) // barrier
26          i++;
27        // C
28
29
30
31
```

CONTINUATION-BASED SYNCHRONIZATION TO THE RESCUE



```
1 [=](sycl::nd_item<1> item) noexcept { // 0
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier(); // 1
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier(); // 2
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 } // -1
```

Continuation-based Synchronization

```
1
2
3
4
5   case 0:
6
7       // A
8       // barrier
9
10  case 1:
11
12      i = 0
13      while(i < 2 + lid)
14          // B
15          if(i < 2) // barrier
16              i++;
17          // C
18
19
20  case 2:
21
22      i++;
23      while(i < 2 + lid)
24          // B
25          if(i < 2) // barrier
26              i++;
27          // C
28
29
30
31
```

CONTINUATION-BASED SYNCHRONIZATION TO THE RESCUE



```
1 [=](sycl::nd_item<1> item) noexcept { // 0
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier(); // 1
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier(); // 2
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 } // -1
```

Continuation-based Synchronization

```
1
2
3
4
5   case 0:
6
7       // A
8       // barrier
9
10  case 1:
11
12      i = 0
13      while(i < 2 + lid)
14          // B
15          if(i < 2) // barrier
16              i++;
17          // C
18
19
20  case 2:
21
22      i++;
23      while(i < 2 + lid)
24          // B
25          if(i < 2) // barrier
26              i++;
27          // C
28
29
30
31
```

CONTINUATION-BASED SYNCHRONIZATION TO THE RESCUE



```
1 [=](sycl::nd_item<1> item) noexcept { // 0
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier(); // 1
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier(); // 2
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 } // -1
```

Continuation-based Synchronization

```
1 i[items] = alloca ..;
2
3
4
5   case 0:
6     for(lid : items[0:])
7       // A
8       // barrier
9
10    case 1:
11      for(lid : items[0:])
12        i[lid] = 0
13        while(i[lid] < 2 + lid)
14          // B
15          if(i[lid] < 2) // barrier
16            i[lid]++;
17          // C
18
19
20    case 2:
21      for(lid : items[0:])
22        i[lid]++;
23        while(i[lid] < 2 + lid)
24          // B
25          if(i[lid] < 2) // barrier
26            i[lid]++;
27          // C
28
29
30
31
```

CONTINUATION-BASED SYNCHRONIZATION TO THE RESCUE



```
1 [=](sycl::nd_item<1> item) noexcept { // 0
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier(); // 1
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier(); // 2
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 } // -1
```

Continuation-based Synchronization

```
1 i[items] = alloca ..;
2 next = 0;
3
4
5   case 0:
6     for(lid : items[0:])
7       // A
8       next = 1;
9
10  case 1:
11    cont1: for(lid : items[0:])
12      i[lid] = 0
13      while(i[lid] < 2 + lid)
14        // B
15        if(i[lid] < 2) next = 2; goto cont1;
16        i[lid]++;
17        // C
18        next = -1;
19
20  case 2:
21    cont2: for(lid : items[0:])
22      i[lid]++;
23      while(i[lid] < 2 + lid)
24        // B
25        if(i[lid] < 2) next = 2; goto cont2;
26        i[lid]++;
27        // C
28        next = -1;
29
30
31
```


CONTINUATION-BASED SYNCHRONIZATION TO THE RESCUE



```
1 [=](sycl::nd_item<1> item) noexcept { // 0
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier(); // 1
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier(); // 2
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 } // -1
```

Continuation-based Synchronization

```
1 i[items] = alloca ..;
2 next = 0;
3 while(next != -1) {
4   switch(next) {
5     case 0:
6       for(lid : items[0:])
7         // A
8         next = 1;
9       break;
10    case 1:
11      cont1: for(lid : items[0:])
12        i[lid] = 0
13        while(i[lid] < 2 + lid)
14          // B
15          if(i[lid] < 2) next = 2; goto cont1;
16          i[lid]++;
17          // C
18          next = -1;
19      break;
20    case 2:
21      cont2: for(lid : items[0:])
22        i[lid]++;
23        while(i[lid] < 2 + lid)
24          // B
25          if(i[lid] < 2) next = 2; goto cont2;
26          i[lid]++;
27          // C
28          next = -1;
29      break;
30  }
31 }
```

CONTINUATION-BASED SYNCHRONIZATION TO THE RESCUE



```
1 [=](sycl::nd_item<1> item) noexcept { // 0
2   const auto lid = item.get_local_id(0);
3
4   scratch[lid] = acc[item.get_global_id()]; // A
5   item.barrier(); // 1
6
7   for(size_t i = 0; i < 2 + lid; i++) {
8     scratch[lid] += i; // B
9     // only call the barrier if all work-items still run the loop.
10    if(i < 2) item.barrier(); // 2
11  }
12  acc[item.get_global_id()] = scratch[lid]; // C
13 } // -1
```

Deep Loop Fission

```
1 for(lid : items[0:])
2   // A
3   // barrier
4   for(i : [0,1])
5     // B (lid = 0)
6     for(lid : items[1:])
7       // B
8       if(i < 2)
9         // barrier
10    for(lid : items[0:])
11      // C
```

Continuation-based Synchronization

```
1 i[items] = alloca ..;
2 next = 0;
3 while(next != -1) {
4   switch(next) {
5     case 0:
6       for(lid : items[0:])
7         // A
8         next = 1;
9       break;
10    case 1:
11      cont1: for(lid : items[0:])
12        i[lid] = 0
13        while(i[lid] < 2 + lid)
14          // B
15          if(i[lid] < 2) next = 2; goto cont1;
16          i[lid]++;
17          // C
18          next = -1;
19      break;
20    case 2:
21      cont2: for(lid : items[0:])
22        i[lid]++;
23        while(i[lid] < 2 + lid)
24          // B
25          if(i[lid] < 2) next = 2; goto cont2;
26          i[lid]++;
27          // C
28          next = -1;
29      break;
30  }
31 }
```

HOW ARE WORK-ITEM PRIVATE VALUES STORED?

 **Dynamically-sized stack arrays with large alignment (64)**

```
1 value = global[offset + lid];  
2 item.barrier();  
3 use(value);
```

⇒

```
1 value[items] = alloca ..;  
2 case 1:  
3   for(lid : items[0:])  
4     value1 = global[offset + lid];  
5     value[lid] = value1;  
6 case 2:  
7   for(lid : items[0:])  
8     value2 = value[lid];  
9     use(value2);
```

AVOID STORING UNIFORM VALUES TO THOSE ARRAYS

Value shape analysis

```
1 offset[items] = alloca ..;  
2 case 1:  
3   for(lid : items[0:])  
4     offset1 = 0; // uniform  
5     offset[lid] = offset1;  
6 case 2:  
7  
8   for(lid : items[0:])  
9     offset2 = offset[lid];
```

⇒

```
1 offset = alloca ..;  
2 case 1:  
3   for(lid : items[0:])  
4     offset1 = 0; // uniform  
5     offset = offset1;  
6 case 2:  
7   offset2 = offset;  
8   for(lid : items[0:])  
9     // ..
```

PROPAGATE VALUE CONTIGUITY TO THE OPTIMIZER

-  Value shape analysis + trace cont values to uniform values & wi-index
+ replicate trace after barrier

```
1 idx[items] = alloca ..;
2 case 1:
3   for(lid : items[0:])
4     idx1 = offset1 + lid; // contiguous
5     idx[lid] = idx1
6 case 2:
7
8   for(lid : items[0:])
9     idx2 = idx[lid];
10    // is this a contiguous access?
11    ptr[idx2] = ..;
```

⇒

```
1 offset = alloca ..; // uniform
2 case 1:
3   for(lid : items[0:])
4     idx1 = offset1 + lid; // contiguous
5     offset = offset1
6 case 2:
7   offset2 = offset
8   for(lid : items[0:])
9     idx2 = offset2 + lid;
10    // this is a contiguous access!
11    ptr[idx2] = ..;
```

PROPAGATE VALUE CONTIGUITY TO THE OPTIMIZER

- ✓ Value shape analysis + trace cont values to uniform values & wi-index
+ replicate trace after barrier

```
1 idx[items] = alloca ..;  
2 case 1:  
3   for(lid : items[0:])  
4     idx1 = offset1 + lid; // contiguous  
5     idx[lid] = idx1  
6 case 2:  
7  
8   for(lid : items[0:])  
9     idx2 = idx[lid];  
10    // is this a contiguous access?  
11    ptr[idx2] = ..;
```

GEOMEAN SPEEDUP

⇒

POCL: 7%,

HIPSYCL: 17%

```
1 offset = alloca ..; // uniform  
2 case 1:  
3   for(lid : items[0:])  
4     idx1 = offset1 + lid; // contiguous  
5     offset = offset1  
6 case 2:  
7   offset2 = offset  
8   for(lid : items[0:])  
9     idx2 = offset2 + lid;  
10    // this is a contiguous access!  
11    ptr[idx2] = ..;
```

HOW ABOUT PERFORMANCE? EVERYWHERE?

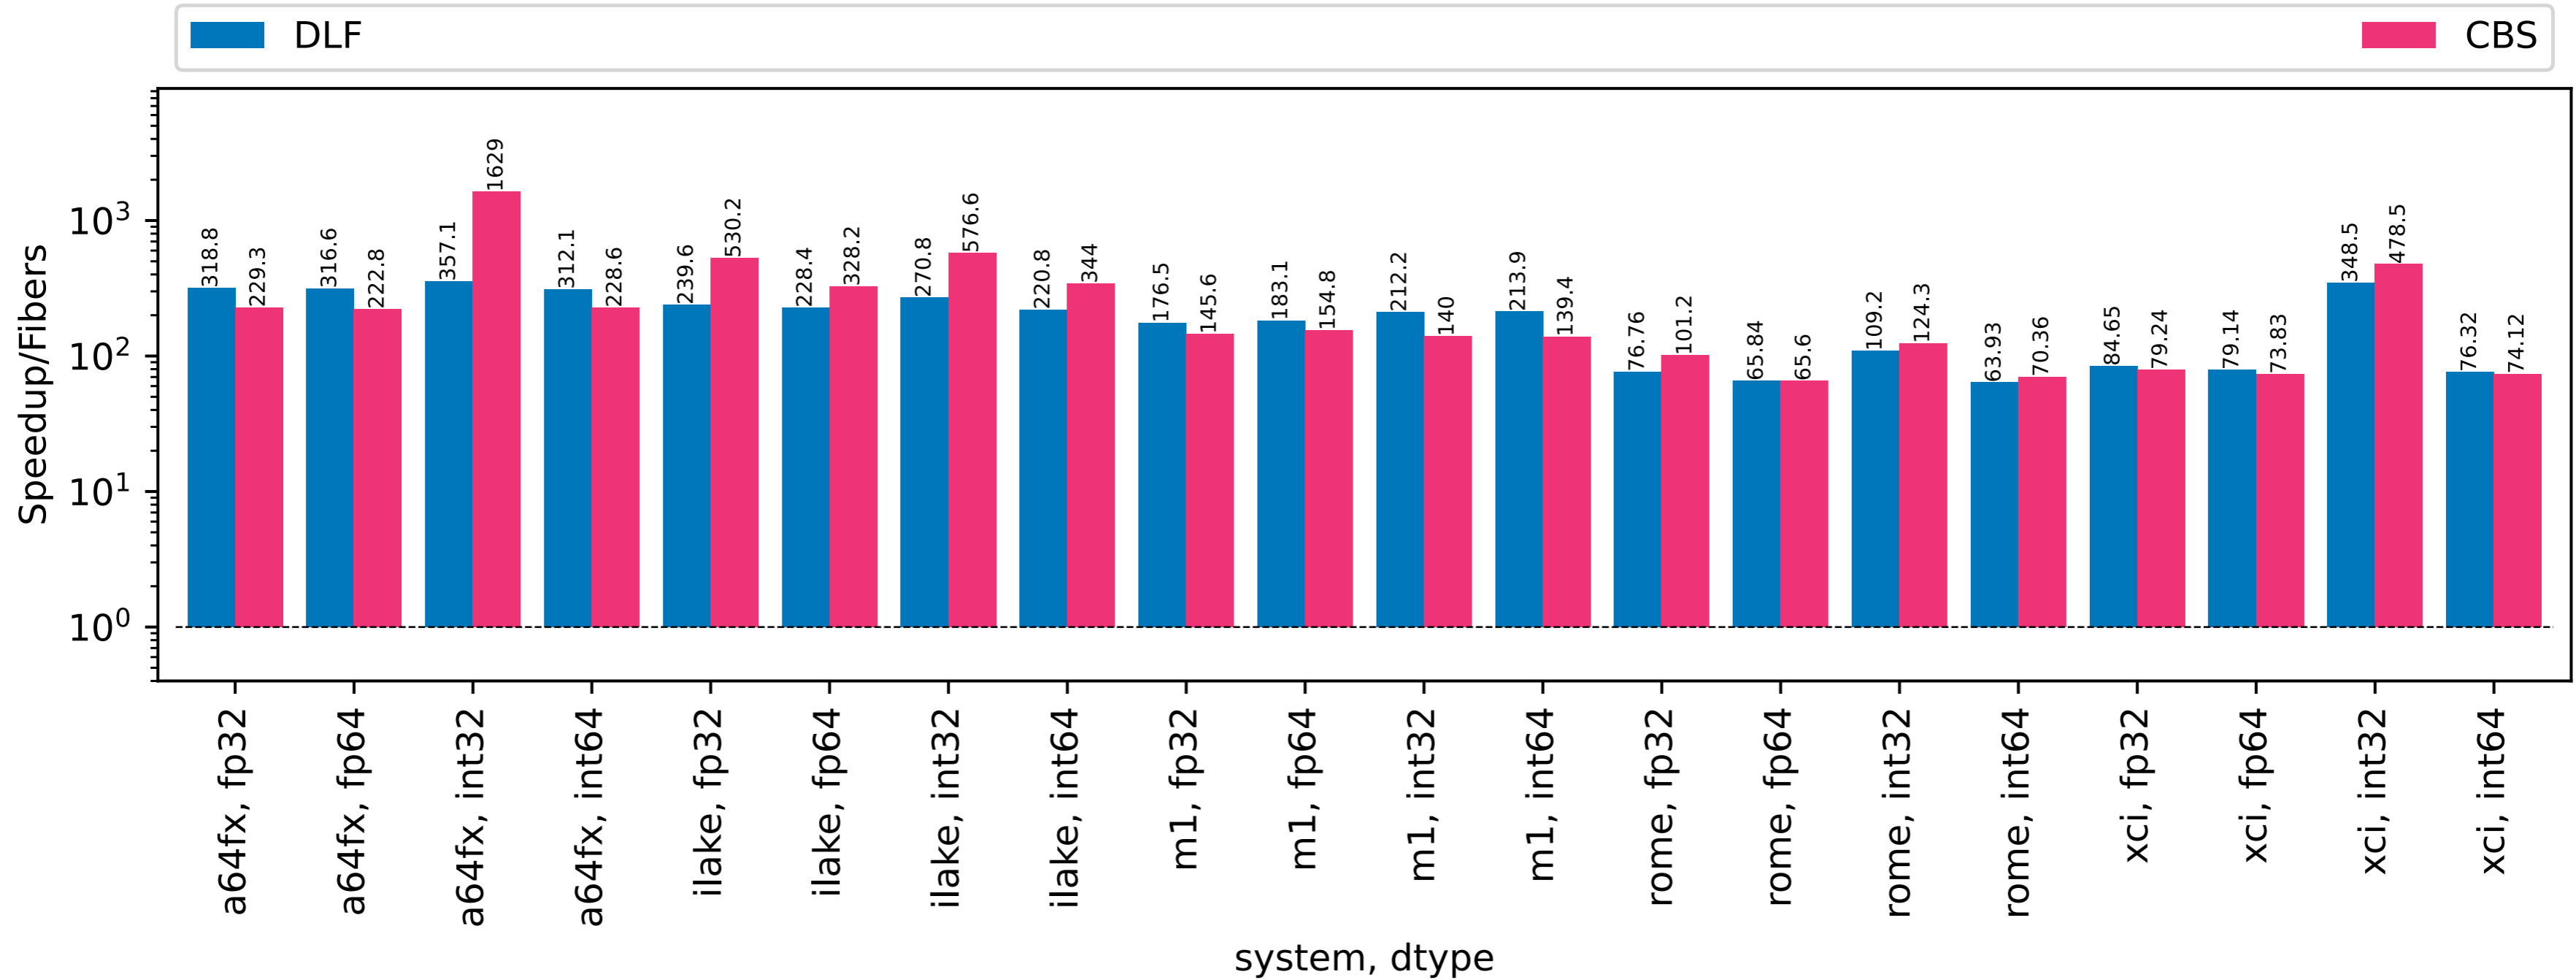
BENCHMARK – SYSTEMS

- Fujitsu A64FX "a64fx"
 - 1x 1.8GHz 48-core CPU
 - 512bit SVE
- Marvell ThunderX2 "xci"
 - 2x 2.1GHz 32-core, 128-threads CPU
 - NEON
- Mac Mini M1 "m1"
 - 1x 4e+4p-core CPU
 - NEON AdvSIMD
- Intel Xeon Gold 6338 "ilake"
 - 2x 2.00GHz 32-core, 64-threads Icelake CPU
 - AVX512
- AMD Epyc 7442 "rome"
 - 2x 2.25GHz 64-core, 128-threads Rome CPU
 - AVX2

MASSIVE SPEEDUPS OF COMPILER APPROACHES



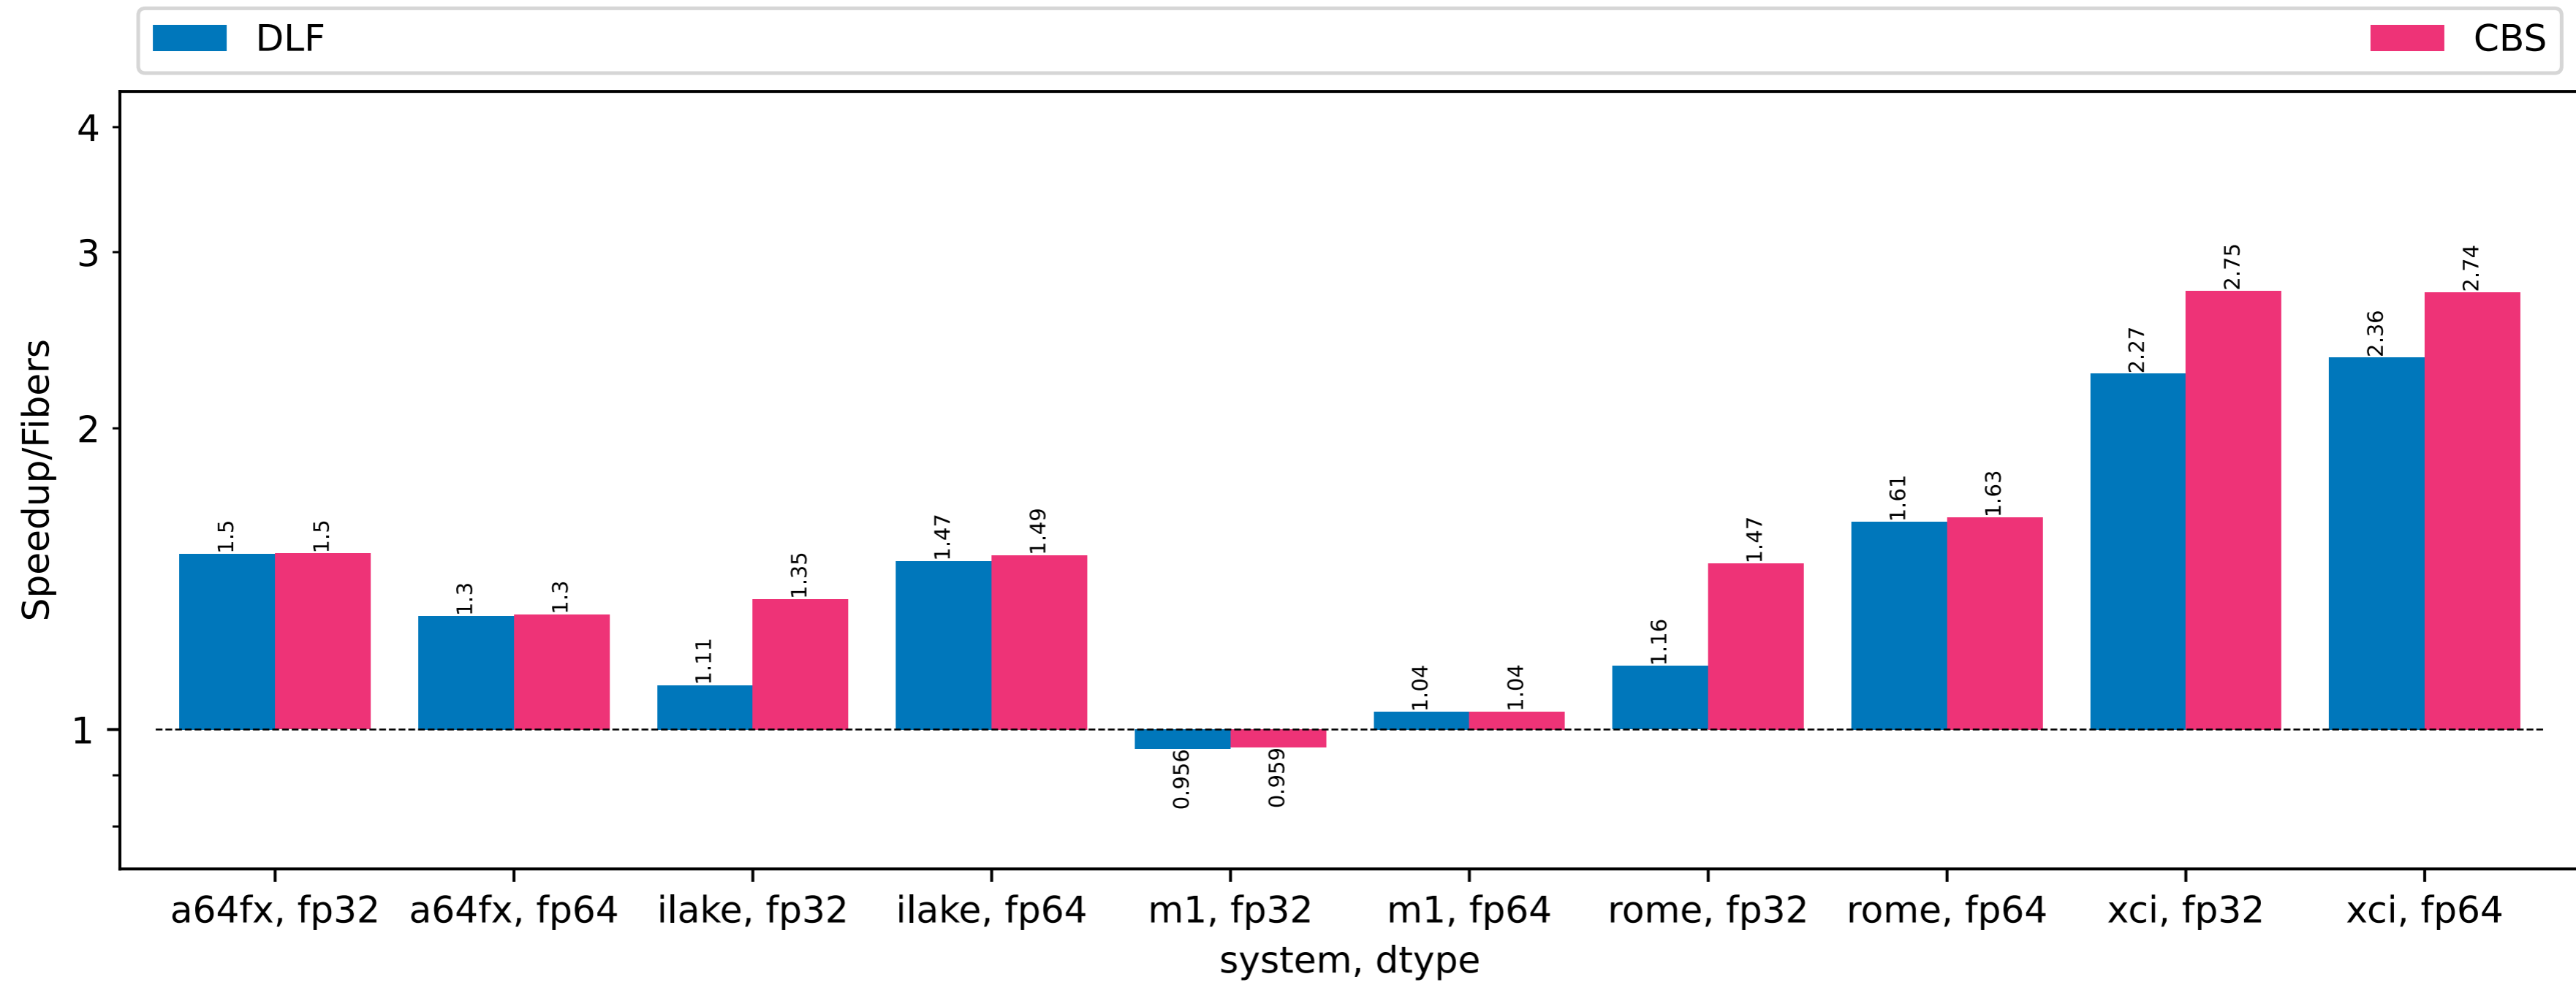
SYCL-Bench Reduction in hipSYCL



github.com/bcosenza/sycl-bench

FIBERS OK FOR HIGH COMPUTE/BARRIER RATIO

SYCL-Bench N-body in hipSYCL



github.com/bcosenza/sycl-bench



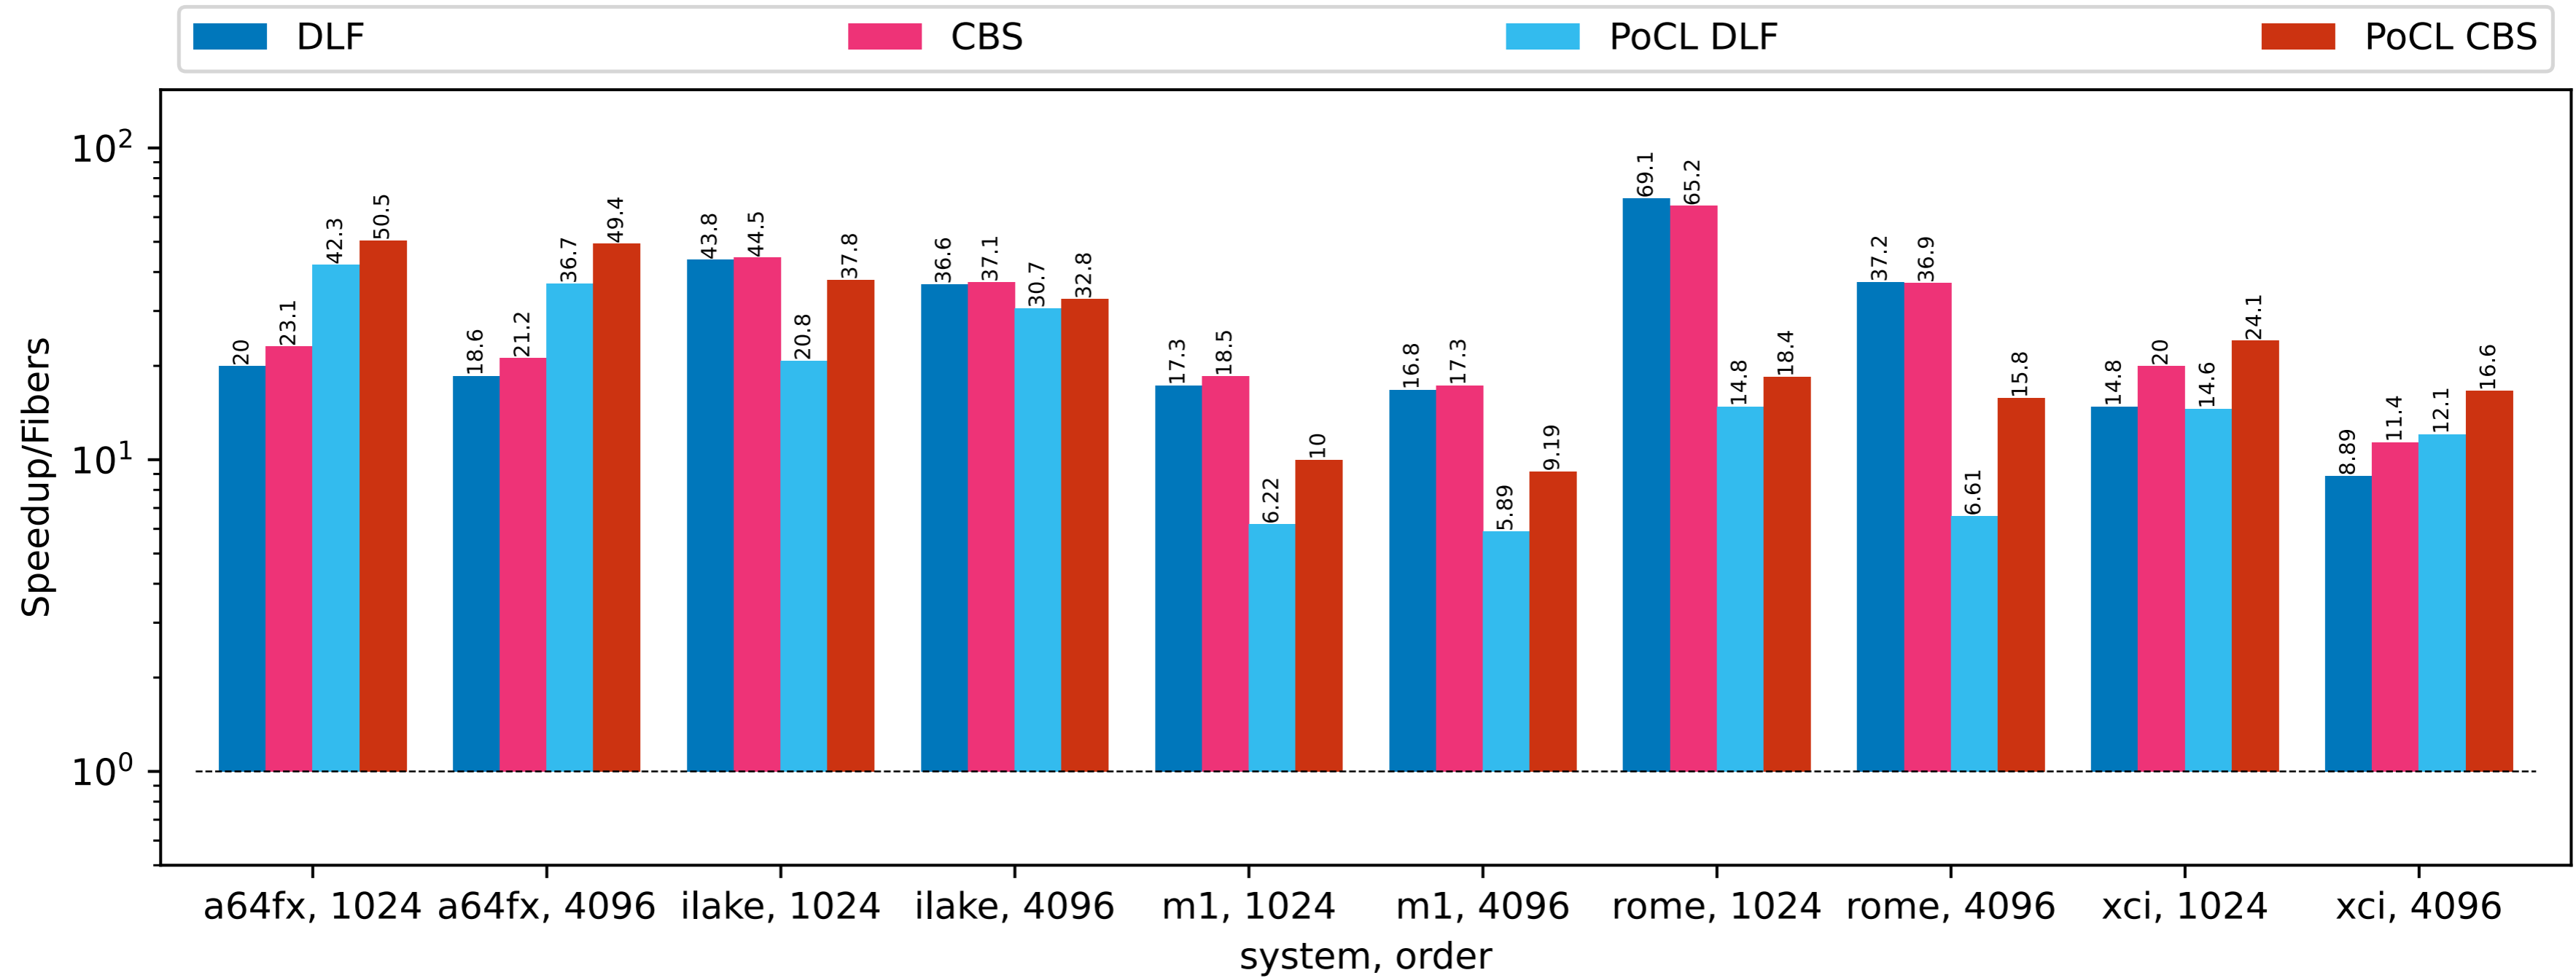
CBS IS COMPETITIVE

hipSYCL summary

	DLF	CBS
Geomean speedup/fibers	29.2	37.7
Number of Best	13	39

PROGRAMMING MODEL INDEPENDENT

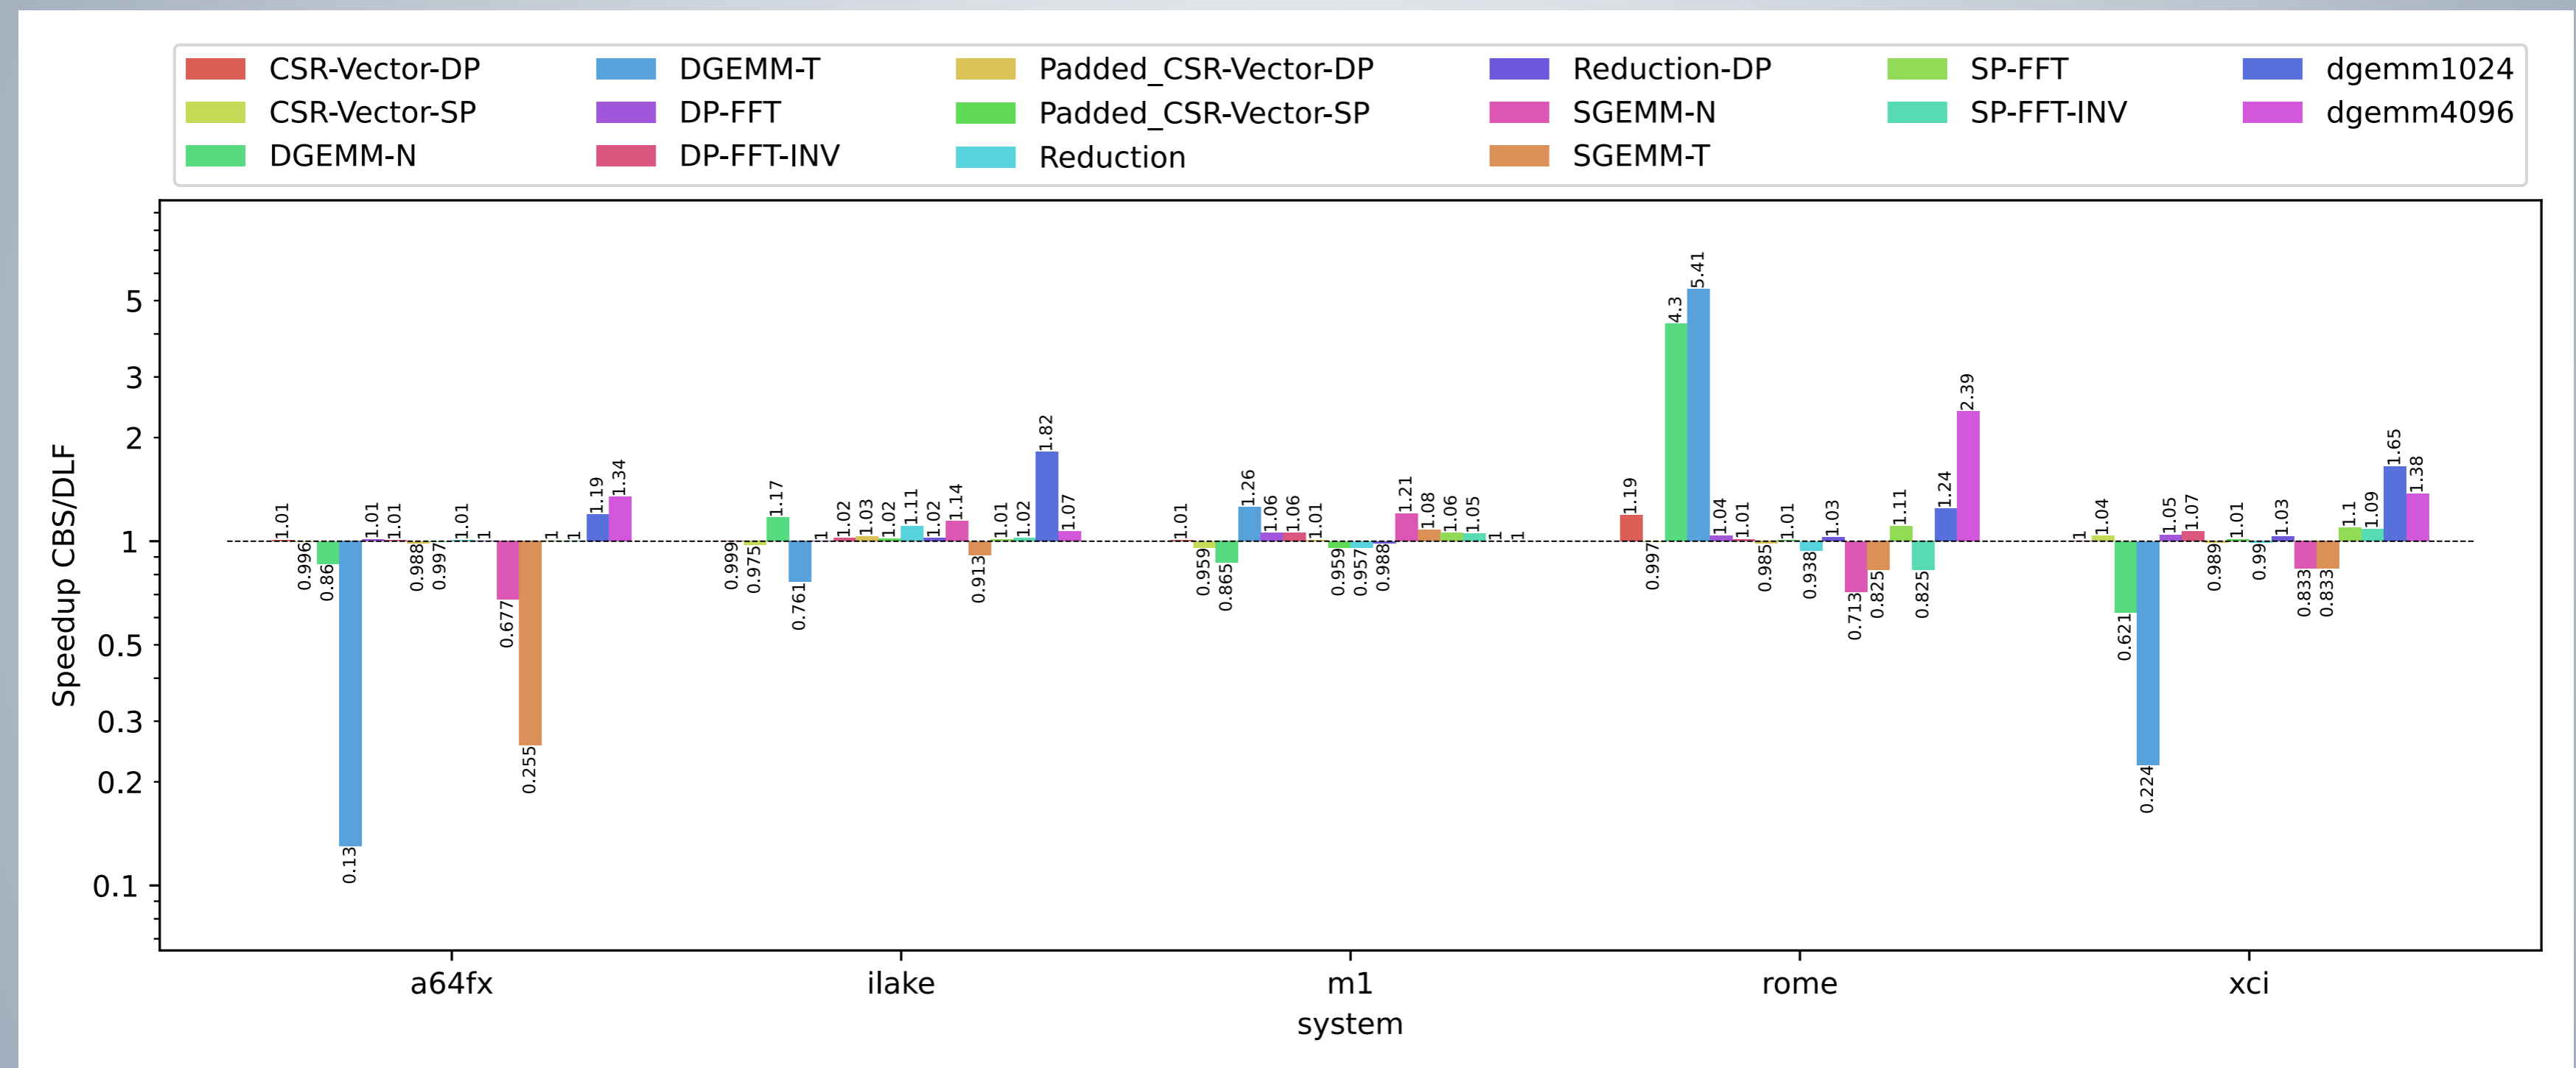
SYCL DGEMM in hipSYCL & PoCL



github.com/UoB-HPC/sycl_dgemm

ARM HPC NOT AS HAPPY WITH CBS

Performance in PoCL



github.com/vetter/shoc
github.com/UoB-HPC/sycl_dgemm

DLF AND CBS COMPARABLE IN OPENCIL

PoCL summary



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386



UNIVERSITÄT
DES
SAARLANDES

SIC Saarland Informatics
Campus

	a64fx	ilake	m1	rome	xc1	overall
Geomean CBS/DLF	0.81	1.05	1.03	1.27	0.93	1.00
CBS #of Best in 16	9	12	9	10	10	41/64

CONCLUSION

	OpenMP	Fibers	DLF	CBS
Library-only	✓	✓	✗	✗
Performance barrier-free	✗	✓	✓	✓
Performance with barrier	✗	✗	✓	✓
Performance on HPC ARM	✗	✗	✓	■
Covering full barrier semantic	✓	✓	✗	✓



THANK YOU FOR LISTENING!

LOOKING FORWARD TO QUESTIONS AND DISCUSSIONS

Contact: Joachim Meyer

jmeyer@cs.uni-saarland.de

github.com/OpenSYCL/OpenSYCL

github.com/pocl/pocl



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

hipSYCL



UNIVERSITÄT
DES
SAARLANDES

SIC Saarland Informatics
Campus

BACKUP

COMPILER APPROACHES SIMILAR TO HIERARCHICAL `PARALLEL_FOR`

Hierarchical `parallel_for`

```
1 h.parallel_for_work_group(  
2   range<1>(groups), [=](group<1> grp) {  
3     grp.parallel_for_work_item( [=](h_item<1> item) {  
4       before_barrier(..);  
5     });  
6     // implicit work-group barrier  
7     grp.parallel_for_work_item( [=](h_item<1> item) {  
8       after_barrier(..);  
9     });  
10  });
```

`nd_range parallel_for` after kernel splitting

```
1 #pragma omp parallel for  
2 for(group : groups)  
3   #pragma omp simd  
4   for(item : itemsInGroup)  
5     kernel_before_barrier(nd_item{group, item})  
6   // implicit synchronization  
7   #pragma omp simd  
8   for(item : itemsInGroup)  
9     kernel_after_barrier(nd_item{group, item})
```

COMPILER APPROACHES SIMILAR TO HIERARCHICAL `PARALLEL_FOR`

Hierarchical `parallel_for`

```
1 h.parallel_for_work_group(  
2   range<1>(groups), [=](group<1> grp) {  
3     grp.parallel_for_work_item( [=](h_item<1> item) {  
4       before_barrier(..);  
5     });  
6     // implicit work-group barrier  
7     grp.parallel_for_work_item( [=](h_item<1> item) {  
8       after_barrier(..);  
9     });  
10  });
```

`nd_range parallel_for` after kernel splitting

```
1 #pragma omp parallel for  
2 for(group : groups)  
3   #pragma omp simd  
4   for(item : itemsInGroup)  
5     kernel_before_barrier(nd_item{group, item})  
6   // implicit synchronization  
7   #pragma omp simd  
8   for(item : itemsInGroup)  
9     kernel_after_barrier(nd_item{group, item})
```

COMPILER APPROACHES SIMILAR TO HIERARCHICAL `PARALLEL_FOR`

Hierarchical `parallel_for`

```
1 h.parallel_for_work_group(  
2   range<1>(groups), [=](group<1> grp) {  
3     grp.parallel_for_work_item( [=](h_item<1> item) {  
4       before_barrier(..);  
5     });  
6     // implicit work-group barrier  
7     grp.parallel_for_work_item( [=](h_item<1> item) {  
8       after_barrier(..);  
9     });  
10  });
```

`nd_range parallel_for` after kernel splitting

```
1 #pragma omp parallel for  
2 for(group : groups)  
3   #pragma omp simd  
4   for(item : itemsInGroup)  
5     kernel_before_barrier(nd_item{group, item})  
6   // implicit synchronization  
7   #pragma omp simd  
8   for(item : itemsInGroup)  
9     kernel_after_barrier(nd_item{group, item})
```

COMPILER APPROACHES SIMILAR TO HIERARCHICAL `PARALLEL_FOR`

Hierarchical `parallel_for`

```
1 h.parallel_for_work_group(  
2   range<1>(groups), [=](group<1> grp) {  
3     grp.parallel_for_work_item( [=](h_item<1> item) {  
4       before_barrier(..);  
5     });  
6     // implicit work-group barrier  
7     grp.parallel_for_work_item( [=](h_item<1> item) {  
8       after_barrier(..);  
9     });  
10  });
```

`nd_range parallel_for` after kernel splitting

```
1 #pragma omp parallel for  
2 for(group : groups)  
3   #pragma omp simd  
4   for(item : itemsInGroup)  
5     kernel_before_barrier(nd_item{group, item})  
6 // implicit synchronization  
7 #pragma omp simd  
8 for(item : itemsInGroup)  
9   kernel_after_barrier(nd_item{group, item})
```

