



SOLID WORK-GROUP SYNCHRONIZATION ON CPUS

LLVM PERFORMANCE WORKSHOP @ CGO'23

**JOACHIM MEYER, SEBASTIAN HACK (SAARLAND UNIVERSITY) AND
AKSEL ALPAY, HOLGER FRÖNING, VINCENT HEUVELINE (HEIDELBERG UNIVERSITY)**



jmeyer@cs.uni-saarland.de



WHY SHOULD WE CARE?

OpenSYCL



ComputeCpp™





WORK-GROUP SYNCHRONIZATION IS RE-IMPLEMENTED TOO OFTEN

```
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  });
```




```
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

```
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

```
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



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!

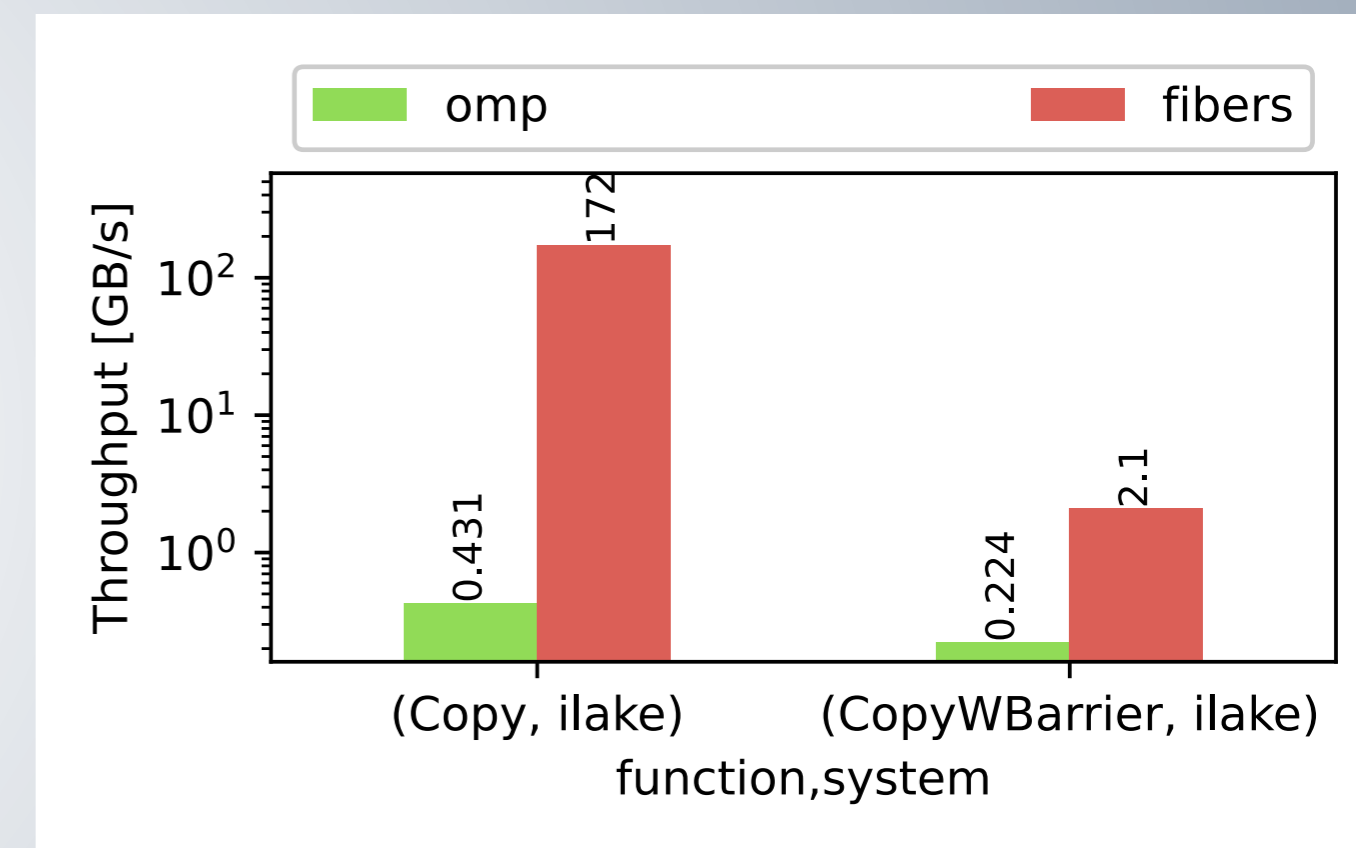
- 1 work-item : 1 thread
 - E.g. OpenMP parallel for + #pragma omp barrier
 - ⚡ Many threads → scheduling overhead
 - ⚡ No vectorization across work-items
- 1 work-item : 1 fiber
 - Lightweight threads + synchronization
 - Can optimize barrier-free kernels!
 - ⚡ Context-switch overhead
 - ⚡ Limited vectorization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

NO, 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})
```

*After all, this is a compiler workshop 🤖

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 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[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];
```

HOW TO AVOID STORING UNIFORM VALUES TO THOSE ARRAYS?

 Value shape analysis based on LLVM's SyncDependenceAnalysis

```
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     // ..
```


HOW TO PROPAGATE VALUE CONTIGUITY TO THE LLVM 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] = ..;
```

HOW TO IMPLEMENT/MAP SUB-GROUP ALGORITHMS?

- Outer-loop vectorization of wi-loops + intrinsics (RV)
- Hierarchically split again at sub-group "barriers"

Scalar kernel

```
1 // A
2 item.barrier();
3 // B
4 cond = group_all_of(sg, x < 32); // barrier
5 // C
```

Outer-loop vectorized (RV)

```
1 for(lid : items[0::4])
2   // A<4>
3   for(lid : items[0::4])
4     // B<4>
5     cond = rv_all(x < 32)
6     // C<4>
```

HOW TO IMPLEMENT/MAP SUB-GROUP ALGORITHMS?

- Outer-loop vectorization of wi-loops + intrinsics (RV)
- Hierarchically split again at sub-group "barriers"

Scalar kernel

```
1 // A
2 item.barrier();
3 // B
4 cond = group_all_of(sg, x < 32); // barrier
5 // C
```

Outer-loop vectorized (RV)

```
1 for(lid : items[0::4])
2   // A<4>
3   for(lid : items[0::4])
4     // B<4>
5     cond = rv_all(x < 32)
6     // C<4>
```

HOW TO IMPLEMENT/MAP SUB-GROUP ALGORITHMS?

- Outer-loop vectorization of wi-loops + intrinsics (RV)
- Hierarchically split again at sub-group "barriers"

Scalar kernel

```
1 // A
2 item.barrier();
3 // B
4 cond = group_all_of(sg, x < 32); // barrier
5 // C
```

Outer-loop vectorized (RV)

```
1 for(lid : items[0::4])
2   // A<4>
3   for(lid : items[0::4])
4     // B<4>
5     cond = rv_all(x < 32)
6     // C<4>
```


HOW TO IMPLEMENT/MAP SUB-GROUP ALGORITHMS?

■ Outer-loop vectorization of wi-loops + intrinsics (RV)

■ Hierarchically split again at sub-group "barriers"

Scalar kernel

```
1 // A
2 item.barrier();
3 // B
4 cond = group_all_of(sg, x < 32); // barrier
5 // C
```

Outer-loop vectorized (RV)

```
1 for(lid : items[0::4])
2   // A<4>
3 for(lid : items[0::4])
4   // B<4>
5   cond = rv_all(x < 32)
6   // C<4>
```

Hierarchical split

```
1 c[max_sg_size] = alloca ..;
2 for(sg : sub_groups)
3   for(sid : sg.items)
4     lid = sg.offset + sid;
5     // A
6 for(sg : sub_groups)
7   for(sid : sg.items)
8     lid = sg.offset + sid;
9     // B
10    c[sid] = x < 32
11    cond = group_all_of(sg, c)
12    for(sid : sg.items)
13      lid = sg.offset + sid;
14    // C
```

HOW TO IMPLEMENT/MAP SUB-GROUP ALGORITHMS?

■ Outer-loop vectorization of wi-loops + intrinsics (RV)

■ Hierarchically split again at sub-group "barriers"

Scalar kernel

```
1 // A
2 item.barrier();
3 // B
4 cond = group_all_of(sg, x < 32); // barrier
5 // C
```

Outer-loop vectorized (RV)

```
1 for(lid : items[0::4])
2   // A<4>
3 for(lid : items[0::4])
4   // B<4>
5   cond = rv_all(x < 32)
6   // C<4>
```

Hierarchical split

```
1 c[max_sg_size] = alloca ..;
2 for(sg : sub_groups)
3   for(sid : sg.items)
4     lid = sg.offset + sid;
5     // A
6 for(sg : sub_groups)
7   for(sid : sg.items)
8     lid = sg.offset + sid;
9     // B
10    c[sid] = x < 32
11    cond = group_all_of(sg, c)
12    for(sid : sg.items)
13      lid = sg.offset + sid;
14    // C
```

HOW TO IMPLEMENT/MAP SUB-GROUP ALGORITHMS?

■ Outer-loop vectorization of wi-loops + intrinsics (RV)

■ Hierarchically split again at sub-group "barriers"

Scalar kernel

```
1 // A
2 item.barrier();
3 // B
4 cond = group_all_of(sg, x < 32); // barrier
5 // C
```

Outer-loop vectorized (RV)

```
1 for(lid : items[0::4])
2   // A<4>
3 for(lid : items[0::4])
4   // B<4>
5   cond = rv_all(x < 32)
6   // C<4>
```







Hierarchical split

```
1 c[max_sg_size] = alloca ..;
2 for(sg : sub_groups)
3   for(sid : sg.items)
4     lid = sg.offset + sid;
5     // A
6 for(sg : sub_groups)
7   for(sid : sg.items)
8     lid = sg.offset + sid;
9     // B
10    c[sid] = x < 32
11    cond = group_all_of(sg, c)
12    for(sid : sg.items)
13      lid = sg.offset + sid;
14      // C
```


PHI NODES ARE HARD.

-  PoCL demotes PHI nodes to make transformations feasible
-  CBS has a single flow change → keep PHIs

SO, DEEP LOOP FISSION OR CBS?

| | DLF (PoCL) | CBS |
|----------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Correct barrier semantic |  |  |
| Expected maintenance cost |  |  |
| Linearized control-flow |  |  |
| Geomean speedup/fiber SYCL | 29 | 38 |
| Performance in OpenCL | on par | |



PROPOSAL

- Continuation-based synchronization as common infrastructure in LLVM upstream.
 - ⇒ No need to re-implement for new programming models / implementations
 - ⇒ New features benefit all implementations (sub-group support, uniformity analysis improvements, ...)



PROPOSAL

- Continuation-based synchronization as common infrastructure in LLVM upstream.
 - ⇒ No need to re-implement for new programming models / implementations
 - ⇒ New features benefit all implementations (sub-group support, uniformity analysis improvements, ...)
- Reuse our battle-tested open-source implementation from Open SYCL (formerly hipSYCL)
- Already used as fallback in PoCL upstream



DESIGN

- Add `llvm.spmd.barrier` intrinsic
- `cbs_kernel` annotation
- `llvm.spmd.group_size`, `llvm.spmd.local_id`
- CBS passes can be added to custom pipeline or extension points



THANK YOU FOR LISTENING!

LOOKING FORWARD TO QUESTIONS AND DISCUSSIONS

Contact: Joachim Meyer

jmeyer@cs.uni-saarland.de

github.com/OpenSYCL/OpenSYCL