

Initial Experiences in Porting a GPU Graph Analysis Workload to Chapel

*PAUL SATHRE, ATHARVA GONDHALEKAR,
& WU-CHUN FENG*
JUNE 2, 2023



VIRGINIA TECH™

Why GPU via Chapel?

Productivity for non-GPU-experts


How would you rather add vectors on a GPU?

How would you rather add vectors on a GPU?

How I learned and spent 10+ years

OpenCL (via MetaCL)

```
1. __kernel void vecAddKernel(__global float *A, __global float *B, __global float *C, int
nelem) {
2.   size_t tid = get_global_id(0);
3.   if (tid < nelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
8.   meta_set_acc(-1, metaModePreferOpenCL);
9.   cl_device_id dev;
10.  cl_platform_id plat;
11.  cl_context ctx;
12.  cl_command_queue q;
13.  meta_get_state_OpenCL(&plat, &dev, &ctx, &q);
14.  cl_mem A, B, C;
15.  A = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
16.  B = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
17.  C = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
18.  clEnqueueWriteBuffer(q, A, CL_FALSE, 0, sizeof(float) * nelem, A_h, 0, NULL, NULL);
19.  clEnqueueWriteBuffer(q, B, CL_TRUE, 0, sizeof(float) * work, B_h, 0, NULL, NULL);
20.  size_t local[3] = {256, 1, 1};
21.  size_t global[3] = {(nelem / local[0]) + (nelem % local[0] ? 1 : 0) * local[0], 1, 1};
22.  metacl_vecAdd_vecAddKernel(q, &global, &local, NULL, false, NULL, &A,
&B, &C, nelem);
23.  //Copy buffers
24.  clEnqueueReadBuffer(q, C, CL_TRUE, 0, sizeof(float) * work, C_h, 0, NULL, NULL);
25.  clFinish(q);
26.  //Release buffers
27.  clReleaseMemObject(A);
28.  clReleaseMemObject(B);
29.  clReleaseMemObject(C);
30. }
```



Portable

Least Programmable

Most Programmable⁴

How would you rather add vectors on a GPU?

How I learned and spent 10+ years

How I've been doing it recently (on FPGA)

OpenCL (via MetaCL)



SYCL

```

1. __kernel void vecAddKernel(__global float *A, __global float *B, __global float *C, int
   nelelem) {
2.     size_t tid = get_global_id(0);
3.     if (tid < nelelem) {
4.         C[tid] = A[tid] + B[tid];
5.     }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelelem) {
8.     meta_set_acc(-1, metaModePreferOpenCL);
9.     cl_device_id dev;
10.    cl_platform_id plat;
11.    cl_context ctx;
12.    cl_command_queue q;
13.    meta_get_state_OpenCL(&plat, &dev, &ctx, &q);
14.    cl_mem A, B, C;
15.    A = clCreateBuffer(ctx, NULL, sizeof(float) * nelelem, NULL, NULL);
16.    B = clCreateBuffer(ctx, NULL, sizeof(float) * nelelem, NULL, NULL);
17.    C = clCreateBuffer(ctx, NULL, sizeof(float) * nelelem, NULL, NULL);
18.    clEnqueueWriteBuffer(q, A, CL_FALSE, 0, sizeof(float) * nelelem, A_h, 0, NULL, NULL);
19.    clEnqueueWriteBuffer(q, B, CL_TRUE, 0, sizeof(float) * work, B_h, 0, NULL, NULL);
20.    size_t local[3] = {256, 1, 1};
21.    size_t global[3] = ((nelem / local[0]) + (nelem % local[0] ? 1 : 0)) * local[0], 1, 1;
22.    metacl_vecAdd_vecAddKernel(q, &global, &local, NULL, false, NULL, &A,
   &B, &C, nelelem);
23.    //Copy buffers
24.    clEnqueueReadBuffer(q, C, CL_TRUE, 0, sizeof(float) * work, C_h, 0, NULL, NULL);
25.    clFinish(q);
26.    //Release buffers
27.    clReleaseMemObject(A);
28.    clReleaseMemObject(B);
29.    clReleaseMemObject(C);
30. }
    
```

```

1. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelelem) {
2.     sycl::queue myQueue;
3.     sycl::buffer<float> A(A_h, nelelem, sycl::property::buffer::use_host_ptr());
4.     sycl::buffer<float> B(B_h, nelelem, sycl::property::buffer::use_host_ptr());
5.     sycl::buffer<float> C(C_h, nelelem, sycl::property::buffer::use_host_ptr());
6.     C.set_write_back(true);
7.     sycl::range<1> local{256};
8.     sycl::range<1> global{((nelem / local.get(0)) + (nelem % local.get(0) ? 1 :
   0)) * local.get(0)};
9.     myQueue.submit([&](sycl::handler &cgh) { //GPU submit
10.         auto A_acc = A.get_access<sycl::access::mode::read>(cgh,
   sycl::range<1>{size_t}nelem);
11.         auto B_acc = B.get_access<sycl::access::mode::read>(cgh,
   sycl::range<1>{size_t}nelem);
12.         auto C_acc = C.get_access<sycl::access::mode::discard_write>(cgh,
   sycl::range<1>{size_t}nelem);
13.         cgh.parallel_for(sycl::nd_range<1>{global, local}, [=](sycl::nd_item<1>
   tid_info) {
14.             size_t tid = tid_info.get_global_linear_id();
15.             if (tid < nelelem) {
16.                 C_acc[tid] = A_acc[tid] + B_acc[tid];
17.             }
18.         });
19.     });
20.     myQueue.wait();
21. }
    
```



Portable

Portable

Least Programmable

Most Programmable⁵

How would you rather add vectors on a GPU?


How I learned and spent 10+ years

OpenCL (via MetaCL)

```

1. __kernel void vecAddKernel(__global float *A, __global float *B, __global float *C, int
   nelem) {
2.   size_t tid = get_global_id(0);
3.   if (tid < nelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
8.   meta_set_acc<I, metaModePreferOpenCL>;
9.   cl_device_id dev;
10.  cl_platform_id plat;
11.  cl_context ctx;
12.  cl_command_queue q;
13.  meta_get_state_OpenCL(&plat, &dev, &ctx, &q);
14.  cl_mem A, B, C;
15.  A = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
16.  B = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
17.  C = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
18.  clEnqueueWriteBuffer(q, A, CL_FALSE, 0, sizeof(float) * nelem, A_h, 0, NULL, NULL);
19.  clEnqueueWriteBuffer(q, B, CL_TRUE, 0, sizeof(float) * work, B_h, 0, NULL, NULL);
20.  size_t local[3] = {256, 1, 1};
21.  size_t global[3] = ((nelem / local[0]) + (nelem % local[0] ? 1 : 0)) * local[0], 1, 1;
22.  metacl_vecAdd_vecAddKernel(q, &global, &local, NULL, false, NULL, &A,
   &B, &C, nelem);
23.  //Copy buffers
24.  clEnqueueReadBuffer(q, C, CL_TRUE, 0, sizeof(float) * work, C_h, 0, NULL, NULL);
25.  clFinish(q);
26.  //Release buffers
27.  clReleaseMemObject(A);
28.  clReleaseMemObject(B);
29.  clReleaseMemObject(C);
30. }

```




How I've been doing it recently (on FPGA)

SYCL

```

1. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
2.   sycl::queue myQueue;
3.   sycl::buffer<float> A(A_h, nelem, sycl::property::buffer::use_host_ptr());
4.   sycl::buffer<float> B(B_h, nelem, sycl::property::buffer::use_host_ptr());
5.   sycl::buffer<float> C(C_h, nelem, sycl::property::buffer::use_host_ptr());
6.   C.set_write_back(true);
7.   sycl::range<1> local{256};
8.   sycl::range<1> global{((nelem / local.get(0)) + (nelem % local.get(0) ? 1 :
   0)) * local.get(0)};
9.   myQueue.submit([&](sycl::handler &cgh) { //GPU submit
10.    auto A_acc = A.get_access<sycl::access::mode::read>(cgh,
   sycl::range<1>{{(size_t)nelem}});
11.    auto B_acc = B.get_access<sycl::access::mode::read>(cgh,
   sycl::range<1>{{(size_t)nelem}});
12.    auto C_acc = C.get_access<sycl::access::mode::discard_write>(cgh,
   sycl::range<1>{{(size_t)nelem}});
13.    cgh.parallel_for(sycl::nd_range<1>{global, local}, [=](sycl::nd_item<1>
   tid_info) {
14.      size_t tid = tid_info.get_global_linear_id();
15.      if (tid < nelem) {
16.        C_acc[tid] = A_acc[tid] + B_acc[tid];
17.      }
18.    });
19.  });
20.  myQueue.wait();
21. }

```




How most GPU kernels are written

CUDA

```

1. __global__ void vecAddKernel(float *A, float *B, float *C, int32_t nelem) {
2.   size_t tid = blockDim.x * blockIdx.x + threadIdx.x;
3.   if (tid < nelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
8.   float *A, *B, *C;
9.   int32_t work = hi-lo+1;
10.  cudaMalloc(&A, sizeof(float) * nelem);
11.  cudaMalloc(&B, sizeof(float) * nelem);
12.  cudaMalloc(&C, sizeof(float) * nelem);
13.  cudaMemcpy(A, A_h, sizeof(float) * nelem, cudaMemcpyHostToDevice);
14.  cudaMemcpy(B, B_h, sizeof(float) * nelem, cudaMemcpyHostToDevice);
15.  dim3 block = {256, 1, 1};
16.  dim3 grid = ((nelem / block.x) + (nelem % block.x ? 1 : 0), 1, 1);
17.  vecAddKernel<<<grid, block>>>(A, B, C, nelem);
18.  cudaMemcpy(C_h, C, sizeof(float) * nelem, cudaMemcpyDeviceToHost);
19.  cudaFree(dA);
20.  cudaFree(dB);
21.  cudaFree(dC);
22. }

```



Not Portable

Least Programmable

Most Programmable⁶

How would you rather add vectors on a GPU?


How I learned and spent 10+ years

OpenCL (via MetaCL)

```

1. __kernel void vecAddKernel(__global float *A, __global float *B, __global float *C, int
nelem) {
2.   size_t tid = get_global_id(0);
3.   if (tid < nelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
8.   meta_set_acc<I, metaModePreferOpenCL>;
9.   cl_device_id dev;
10.  cl_platform_id plat;
11.  cl_context ctx;
12.  cl_command_queue q;
13.  meta_get_state_OpenCL(&plat, &dev, &ctx, &q);
14.  cl_mem A, B, C;
15.  A = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
16.  B = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
17.  C = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
18.  clEnqueueWriteBuffer(q, A, CL_FALSE, 0, sizeof(float) * nelem, A_h, 0, NULL, NULL);
19.  clEnqueueWriteBuffer(q, B, CL_TRUE, 0, sizeof(float) * work, B_h, 0, NULL, NULL);
20.  size_t local[3] = {256, 1, 1};
21.  size_t global[3] = ((nelem / local[0]) + (nelem % local[0] ? 1 : 0)) * local[0], 1, 1;
22.  metacl_vecAdd_vecAddKernel(q, &global, &local, NULL, false, NULL, &A,
&B, &C, nelem);
23.  //Copy buffers
24.  clEnqueueReadBuffer(q, C, CL_TRUE, 0, sizeof(float) * work, C_h, 0, NULL, NULL);
25.  clFinish(q);
26.  //Release buffers
27.  clReleaseMemObject(A);
28.  clReleaseMemObject(B);
29.  clReleaseMemObject(C);
30. }

```



Portable


How I've been doing it recently (on FPGA)

SYCL

```

1. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
2.   sycl::queue myQueue;
3.   sycl::buffer<float> A(A_h, nelem, sycl::property::buffer::use_host_ptr());
4.   sycl::buffer<float> B(B_h, nelem, sycl::property::buffer::use_host_ptr());
5.   sycl::buffer<float> C(C_h, nelem, sycl::property::buffer::use_host_ptr());
6.   C.set_write_back(true);
7.   sycl::range<1> local{256};
8.   sycl::range<1> global{((nelem / local.get(0)) + (nelem % local.get(0) ? 1 :
0)) * local.get(0)};
9.   myQueue.submit([&](sycl::handler &cgh) { //GPU submit
10.    auto A_acc = A.get_access<sycl::access::mode::read>(cgh,
sycl::range<1>{(size_t)nelem});
11.    auto B_acc = B.get_access<sycl::access::mode::read>(cgh,
sycl::range<1>{(size_t)nelem});
12.    auto C_acc = C.get_access<sycl::access::mode::discard_write>(cgh,
sycl::range<1>{(size_t)nelem});
13.    cgh.parallel_for(sycl::nd_range<1>{global, local}, [=](sycl::nd_item<1>
tid_info) {
14.      size_t tid = tid_info.get_global_linear_id();
15.      if (tid < nelem) {
16.        C_acc[tid] = A_acc[tid] + B_acc[tid];
17.      }
18.    });
19.  });
20.  myQueue.wait();
21. }

```



Portable


How most GPU kernels are written

CUDA

```

1. __global__ void vecAddKernel(float *A, float *B, float *C, int32_t nelem) {
2.   size_t tid = blockDim.x * blockIdx.x + threadIdx.x;
3.   if (tid < nelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
8.   float *A, *B, *C;
9.   int32_t work = hi-lo+1;
10.  cudaMalloc(&A, sizeof(float) * nelem);
11.  cudaMalloc(&B, sizeof(float) * nelem);
12.  cudaMalloc(&C, sizeof(float) * nelem);
13.  cudaMemcpy(A, A_h, sizeof(float) * nelem, cudaMemcpyHostToDevice);
14.  cudaMemcpy(B, B_h, sizeof(float) * nelem, cudaMemcpyHostToDevice);
15.  dim3 block = {256, 1, 1};
16.  dim3 grid = ((nelem / block.x) + (nelem % block.x ? 1 : 0), 1, 1);
17.  vecAddKernel<<<grid, block>>>(A, B, C, nelem);
18.  cudaMemcpy(C_h, C, sizeof(float) * nelem, cudaMemcpyDeviceToHost);
19.  cudaFree(dA);
20.  cudaFree(dB);
21.  cudaFree(dC);
22. }

```



Not Portable


How we get to talk about today =)

Chapel

```

1. use GPU;
2. proc vecAdd(A_h: [] real(32),
B_h: [] real(32), C_h: [] real(32)) {
3.   on here.gpus[0] { //GPU locale
4.     var A: [A_h.domain] real(32);
5.     var B: [B_h.domain] real(32);
6.     var C: [C_h.domain] real(32);
7.     A = A_h; //copy-to
8.     B = B_h; //copy-to
9.     C = A + B; //GPU add
10.    C_h = C; //copy-from
11.  }
12. }

```



Portable

OR

Least Programmable

Most Programmable


How would you rather add vectors on a GPU?

How I learned and spent 10+ years

OpenCL (via MetaCL)

```

1. __kernel void vecAddKernel(__global float *A, __global float *B, __global float *C, int
   nelem) {
2.   size_t tid = get_global_id(0);
3.   if (tid < nelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
8.   meta_set_acc<-I, metaModePreferOpenCL>;
9.   cl_device_id dev;
10.  cl_platform_id plat;
11.  cl_context ctx;
12.  cl_command_queue q;
13.  meta_get_state_OpenCL(&plat, &dev, &ctx, &q);
14.  cl_mem A, B, C;
15.  A = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
16.  B = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
17.  C = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
18.  clEnqueueWriteBuffer(q, A, CL_FALSE, 0, sizeof(float) * nelem, A_h, 0, NULL, NULL);
19.  clEnqueueWriteBuffer(q, B, CL_TRUE, 0, sizeof(float) * work, B_h, 0, NULL, NULL);
20.  size_t local[3] = {256, 1, 1};
21.  size_t global[3] = {(nelem / local[0]) + (nelem % local[0] ? 1 : 0), local[1], 1};
22.  meta_cl_vecAdd_vecAddKernel(q, &global, &local, NULL, false, NULL, &A,
   &B, &C, nelem);
23.  //Copy buffers
24.  clEnqueueReadBuffer(q, C, CL_TRUE, 0, sizeof(float) * work, C_h, 0, NULL, NULL);
25.  clFinish(q);
26.  //Release buffers
27.  clReleaseMemObject(A);
28.  clReleaseMemObject(B);
29.  clReleaseMemObject(C);
30. }
    
```




Portable

How I've been doing it recently (on FPGA)

SYCL

```

1. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
2.   sycl::queue myQueue;
3.   sycl::buffer<float> A(A_h, nelem, sycl::property::buffer::use_host_ptr());
4.   sycl::buffer<float> B(B_h, nelem, sycl::property::buffer::use_host_ptr());
5.   sycl::buffer<float> C(C_h, nelem, sycl::property::buffer::use_host_ptr());
6.   C.set_write_back(true);
7.   sycl::range<1> local{256};
8.   sycl::range<1> global{((nelem / local.get(0)) + (nelem % local.get(0) ? 1 :
   0)) * local.get(0)};
9.   myQueue.submit([&](sycl::handler &cgh) { //GPU submit
10.    auto A_acc = A.get_access<sycl::access::mode::read>(cgh,
   sycl::range<1>{(size_t)nelem});
11.    auto B_acc = B.get_access<sycl::access::mode::read>(cgh,
   sycl::range<1>{(size_t)nelem});
12.    auto C_acc = C.get_access<sycl::access::mode::discard_write>(cgh,
   sycl::range<1>{(size_t)nelem});
13.    cgh.parallel_for(sycl::nd_range<1>{global, local}, [=](sycl::nd_item<1>
   tid_info) {
14.      size_t tid = tid_info.get_global_linear_id();
15.      if (tid < nelem) {
16.        C_acc[tid] = A_acc[tid] + B_acc[tid];
17.      }
18.    });
19.  });
20.  myQueue.wait();
21. }
    
```




Portable

How most GPU kernels are written

CUDA

```

1. __global__ void vecAddKernel(float *A, float *B, float *C, int32_t nelem) {
2.   size_t tid = blockDim.x * blockIdx.x + threadIdx.x;
3.   if (tid < nelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
8.   float *A, *B, *C;
9.   int32_t work = hi-lo+1;
10.  cudaMalloc(&A, sizeof(float) * nelem);
11.  cudaMalloc(&B, sizeof(float) * nelem);
12.  cudaMalloc(&C, sizeof(float) * nelem);
13.  cudaMemcpy(A, A_h, sizeof(float) * nelem, cudaMemcpyHostToDevice);
14.  cudaMemcpy(B, B_h, sizeof(float) * nelem, cudaMemcpyHostToDevice);
15.  dim3 block = {256, 1, 1};
16.  dim3 grid = {(nelem / block.x) + (nelem % block.x ? 1 : 0), 1, 1};
17.  vecAddKernel<<<grid, block>>>(A, B, C, nelem);
18.  cudaMemcpy(C_h, C, sizeof(float) * nelem, cudaMemcpyDeviceToHost);
19.  cudaFree(dA);
20.  cudaFree(dB);
21.  cudaFree(dC);
22. }
    
```




Not Portable

How we get to talk about today =)

Chapel

```

1. use GPU;
2. proc vecAdd(A_h: [] real(32),
   B_h: [] real(32), C_h: [] real(32)) {
3.   on here.gpus[0] { //GPU locale
4.     var A: [A_h.domain] real(32);
5.     var B: [B_h.domain] real(32);
6.     var C: [C_h.domain] real(32);
7.     A = A_h; //copy-to
8.     B = B_h; //copy-to
9.     C = A + B; //GPU add (promoted)
10.    C_h = C; //copy-from
11.  }
12. }
    
```



Portable

OR

Kernel & Launch

Least Programmable

Most Programmable

How would you rather add vectors on a GPU?


How I learned and spent 10+ years

OpenCL (via MetaCL)

```

1. __kernel void vecAddKernel(__global float *A, __global float *B, __global float *C, int
   nelelem) {
2.   size_t tid = get_global_id(0);
3.   if (tid < nelelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelelem) {
8.   meta_set_acc<-I, metaModePreferOpenCL>;
9.   cl_device_id dev;
10.  cl_platform_id plat;
11.  cl_context ctx;
12.  cl_command_queue q;
13.  meta_get_state_OpenCL(&plat, &dev, &ctx, &q);
14.  cl_mem A, B, C;
15.  A = clCreateBuffer(ctx, NULL, sizeof(float) * nelelem, NULL, NULL);
16.  B = clCreateBuffer(ctx, NULL, sizeof(float) * nelelem, NULL, NULL);
17.  C = clCreateBuffer(ctx, NULL, sizeof(float) * nelelem, NULL, NULL);
18.  clEnqueueWriteBuffer(q, A, CL_FALSE, 0, sizeof(float) * nelelem, A_h, 0, NULL, NULL);
19.  clEnqueueWriteBuffer(q, B, CL_TRUE, 0, sizeof(float) * work, B_h, 0, NULL, NULL);
20.  size_t local[3] = {256, 1, 1};
21.  size_t global[3] = {(nelem / local[0]) + (nelem % local[0] ? 1 : 0) * local[0], 1, 1};
22.  metacl_vecAdd_vecAddKernel(q, &global, &local, NULL, false, NULL, &A,
   &B, &C, nelelem);
23.  //Copy buffers
24.  clEnqueueReadBuffer(q, C, CL_TRUE, 0, sizeof(float) * work, C_h, 0, NULL, NULL);
25.  clFinish(q);
26.  //Release buffers
27.  clReleaseMemObject(A);
28.  clReleaseMemObject(B);
29.  clReleaseMemObject(C);
30. }

```



Portable


How I've been doing it recently (on FPGA)

SYCL

```

1. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelelem) {
2.   sycl::queue myQueue;
3.   sycl::buffer<float> A(A_h, nelelem, sycl::property::buffer::use_host_ptr());
4.   sycl::buffer<float> B(B_h, nelelem, sycl::property::buffer::use_host_ptr());
5.   sycl::buffer<float> C(C_h, nelelem, sycl::property::buffer::use_host_ptr());
6.   C.set_write_back(true);
7.   sycl::range<1> local(256);
8.   sycl::range<1> global{((nelem / local.get(0)) + (nelem % local.get(0) ? 1 :
   0)) * local.get(0)};
9.   myQueue.submit([&](sycl::handler &cgh) { //GPU submit
10.    auto A_acc = A.get_access<sycl::access::mode::read>(cgh,
   sycl::range<1>{global});
11.    auto B_acc = B.get_access<sycl::access::mode::read>(cgh,
   sycl::range<1>{global});
12.    auto C_acc = C.get_access<sycl::access::mode::discard_write>(cgh,
   sycl::range<1>{global});
13.    cgh.parallel_for(sycl::nd_range<1>{global, local}, [=](sycl::nd_item<1>
   tid_info) {
14.      size_t tid = tid_info.get_global_linear_id();
15.      if (tid < nelelem) {
16.        C_acc[tid] = A_acc[tid] + B_acc[tid];
17.      }
18.    });
19.  });
20.  myQueue.wait();
21. }

```



Portable


How most GPU kernels are written

CUDA

```

1. __global__ void vecAddKernel(float *A, float *B, float *C, int32_t nelelem) {
2.   size_t tid = blockDim.x * blockIdx.x + threadIdx.x;
3.   if (tid < nelelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelelem) {
8.   float *A, *B, *C;
9.   int32_t work = hi-lo+1;
10.  cudaMalloc(&A, sizeof(float) * nelelem);
11.  cudaMalloc(&B, sizeof(float) * nelelem);
12.  cudaMalloc(&C, sizeof(float) * nelelem);
13.  cudaMemcpy(A, A_h, sizeof(float) * nelelem, cudaMemcpyHostToDevice);
14.  cudaMemcpy(B, B_h, sizeof(float) * nelelem, cudaMemcpyHostToDevice);
15.  dim3 block = {256, 1, 1};
16.  dim3 grid = {(nelem / block.x) + (nelem % block.x ? 1 : 0), 1, 1};
17.  vecAddKernel<<<grid, block>>>(A, B, C, nelelem);
18.  cudaMemcpy(C_h, C, sizeof(float) * nelelem, cudaMemcpyDeviceToHost);
19.  cudaFree(dA);
20.  cudaFree(dB);
21.  cudaFree(dC);
22. }

```



Not Portable


How we get to talk about today =)

Chapel

```

1. use GPU;
2. proc vecAdd(A_h: [] real(32), B_h: [] real(32), C_h: [] real(32)) {
3.   on here.gpus[0] { //GPU locale
4.     var A: [A_h.domain] real(32);
5.     var B: [B_h.domain] real(32);
6.     var C: [C_h.domain] real(32);
7.     A = A_h; //copy-to
8.     B = B_h; //copy-to
9.     C = A + B; //GPU add
10.    C_h = C; //copy-from
11.  }
12. }

```



Portable

OR

Device Allocate / Free

Least Programmable

Most Programmable


How would you rather add vectors on a GPU?

How I learned and spent 10+ years

OpenCL (via MetaCL)

```

1. __kernel void vecAddKernel(__global float *A, __global float *B, __global float *C, int
nelem) {
2.   size_t tid = get_global_id(0);
3.   if (tid < nelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
8.   meta_set_accr(-1, metaModePreferOpenCL);
9.   cl_device_id dev;
10.  cl_platform_id plat;
11.  cl_context ctx;
12.  cl_command_queue q;
13.  meta_get_state_OpenCL(&plat, &dev, &ctx, &q);
14.  cl_mem A, B, C;
15.  A = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
16.  B = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
17.  C = clCreateBuffer(ctx, NULL, sizeof(float) * nelem, NULL, NULL);
18.  clEnqueueWriteBuffer(q, A, CL_FALSE, 0, sizeof(float) * nelem, A_h, 0, NULL, NULL);
19.  clEnqueueWriteBuffer(q, B, CL_TRUE, 0, sizeof(float) * work, B_h, 0, NULL, NULL);
20.  size_t local[3] = {256, 1, 1};
21.  size_t global[3] = ((nelem / local[0]) + (nelem % local[0] ? 1 : 0)) * local[0], 1, 1);
22.  metacl_vecAdd_vecAddKernel(q, &global, &local, NULL, false, NULL, &A,
&B, &C, nelem);
23.  //Copy buffers
24.  clEnqueueReadBuffer(q, C, CL_TRUE, 0, sizeof(float) * work, C_h, 0, NULL, NULL);
25.  clFinish(q);
26.  //Release buffers
27.  clReleaseMemObject(A);
28.  clReleaseMemObject(B);
29.  clReleaseMemObject(C);
30. }
    
```




Portable

How I've been doing it recently (on FPGA)

SYCL

```

1. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
2.   sycl::queue myQueue;
3.   sycl::buffer<float> A(A_h, nelem, sycl::property::buffer::use_host_ptr());
4.   sycl::buffer<float> B(B_h, nelem, sycl::property::buffer::use_host_ptr());
5.   sycl::buffer<float> C(C_h, nelem, sycl::property::buffer::use_host_ptr());
6.   C.set_write_back(true);
7.   sycl::range<1> local(256);
8.   sycl::range<1> global{((nelem / local.get(0)) + (nelem % local.get(0) ? 1 :
0)) * local.get(0)};
9.   myQueue.submit([&](sycl::handler &cgh) { //GPU submit
10.    auto A_acc = A.get_access<sycl::access::mode::read>(cgh,
sycl::range<1>{local.get(0)});
11.    auto B_acc = B.get_access<sycl::access::mode::read>(cgh,
sycl::range<1>{local.get(0)});
12.    auto C_acc = C.get_access<sycl::access::mode::discard_write>(cgh,
sycl::range<1>{local.get(0)});
13.    cgh.parallel_for(sycl::nd_range<1>{global, local}, [=](sycl::nd_item<1>
tid_info) {
14.      size_t tid = tid_info.get_global_linear_id();
15.      if (tid < nelem) {
16.        C_acc[tid] = A_acc[tid] + B_acc[tid];
17.      }
18.    });
19.  });
20.  myQueue.wait();
21. }
    
```




Portable

How most GPU kernels are written

CUDA

```

1. __global__ void vecAddKernel(float *A, float *B, float *C, int32_t nelem) {
2.   size_t tid = blockDim.x * blockIdx.x + threadIdx.x;
3.   if (tid < nelem) {
4.     C[tid] = A[tid] + B[tid];
5.   }
6. }
7. void vecAdd(float *A_h, float *B_h, float *C_h, int32_t nelem) {
8.   float *A, *B, *C;
9.   int32_t work = hi-lo+1;
10.  cudaMalloc(&A, sizeof(float) * nelem);
11.  cudaMalloc(&B, sizeof(float) * nelem);
12.  cudaMalloc(&C, sizeof(float) * nelem);
13.  cudaMemcpy(A, A_h, sizeof(float) * nelem, cudaMemcpyHostToDevice);
14.  cudaMemcpy(B, B_h, sizeof(float) * nelem, cudaMemcpyHostToDevice);
15.  dim3 block = {256, 1, 1};
16.  dim3 grid = ((nelem / block.x) + (nelem % block.x ? 1 : 0), 1, 1);
17.  vecAddKernel<<<grid, block>>>(A, B, C, nelem);
18.  cudaMemcpy(C_h, C, sizeof(float) * nelem, cudaMemcpyDeviceToHost);
19.  cudaFree(dA);
20.  cudaFree(dB);
21.  cudaFree(dC);
22. }
    
```




Not Portable

How we get to talk about today =)

Chapel

```

1. use GPU;
2. proc vecAdd(A_h: [] real(32),
B_h: [] real(32), C_h: [] real(32)) {
3.   on here.gpus[0] { //GPU locale
4.     var A: [A_h.domain] real(32);
5.     var B: [B_h.domain] real(32);
6.     var C: [C_h.domain] real(32);
7.     A = A_h; //copy-to
8.     B = B_h; //copy-to
9.     C = A + B; //GPU add
10.    C_h = C; //copy-from
11.  }
12. }
    
```



Portable

OR

X-fer
CPU → GPU
GPU → CPU

Most Programmable¹⁰


Least Programmable

Why *Graph Analysis on GPU* via Chapel?


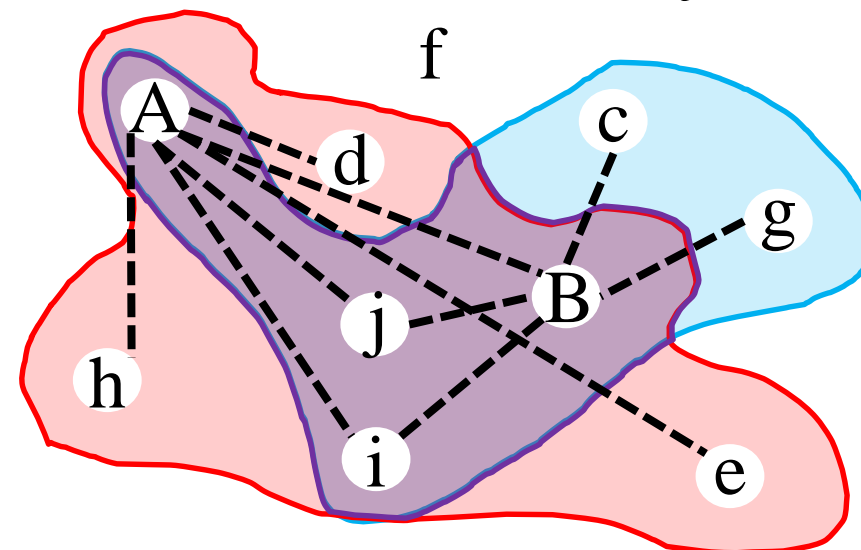
Understand Chapel's programmability and GPU performance on irregular applications relative to proven approaches

Graph Workload: Edge-connected Jaccard similarity

- Intersection over union for all edge-connected pairs' 1-hop neighborhoods
 - Essentially a batch of $|E|$ set-intersections
 - Lots of indirect access
- Why do we care?
 - *Strength* of similarity between known pairs
 - *Proxy* for more complex intersection algo's
 - Recommendations, Community detection, ...
 - But mostly: *How amenable is GPU Chapel to irregular algorithms?*
- What's our approach? Port two existing (CUDA/SYCL) kernel pipelines
 - Edge-centric: homegrown, 1 pair per thread \rightarrow good at near-hypersparse graphs
 - Vertex-centric: from *legacy* CuGraph, 3D, $n=8$ threads per pair \rightarrow better on denser graphs



A. Fender, et al, 2017. "Parallel Jaccard and Related Graph Clustering Techniques." (*Scala '17*).



$$\text{Jaccard Similarity}(A,B) = \frac{\text{Intersection}(A,B)}{\text{Union}(A,B)}$$

```
Edges = {(a,b), (a,d), (a,e), (a,h), (a,i),
         (a,j), (b,c), (b,g), (b,i), (b,j)}
forall e in Edges {
    JS(e) = intersect(e)/union(e)
}
```



How is the *programmability*?

Write how you know *first*, then try to make idiomatic

- Still “Chapel how a C programmer would write it”
- Currently only used promoted-array kernels once (vertex-centric *Fill*), but more opportunities exist
 - Edge-centric reverse-edge preprocessing (*Scan*)
 - intersect-over-union array division (*Weights*)
- C-like generic compressed sparse row (CSR) structure could be revisited
- But still modest source-lines improvement
- C-style Chapel: 756 lines
 - Edge-centric: 150
 - Vertex-Centric: 230
 - Everything else: 376
- CUDA: 1212
 - All Kernels: 395
 - Everything else: 817




```
template <typename vertex_t, typename
__global__ void jaccard_ec (vertex_t
*weight_j) {
    ... //thread-private vars
    tid = blockIdx.x * blockDim.x + thr
    if (tid < e) {
        ... // choose smaller vert as 're
        for (i = csrPtr[ref]; i < csrPtr[
            ref_col = csrInd[i];
            ... // bin-search for ref_col b
            ... // accumulate matches in we
        }
        ...
    }
}
```



```
forall i in indices.domain {
    assertOnGpu(); //Fail if this can
    ... // thread-private vars
    ... // choose smaller vert as 're
    for j in offsets[refs]..<(offsets
        ref_col = indices[j];
        ... // bin-search for ref_col b
        ... // accumulate matches in we
    }
}
```

Vertex-Centric: More complex

- Vertex-centric intersect didn't initially GPU-ize




```

var intersect : [0..<numEdges] atomic real(32);
forall Z in srcIters by gridDim.z*blockDim.z {
  forall Y in destIters by gridDim.y*blockDim.y {
    forall X in isectIters by gridDim.x*blockDim.x {
      ... // bin-search
      intersect[writeAddr].add(1.0);
    }
  }
}

```

Vertex-Centric: More complex



- Vertex-centric intersect didn't initially GPU-ize
 - 3D grid/block → Chapel pre-1.31 currently only supports 1D forall
 - Had to linearize to 1D with a *very large* index space

VC-CPU-Naïve

```
var intersect : [0..
```

VC-CPU-Linearized

```
var intersect : [0..
```

(↑ pseudo-code, see Github for real)


```

var intersect : [0..

```


```

var intersect : [0..

```

Vertex-Centric: More complex

- Vertex-centric intersect didn't initially GPU-ize
 - 3D grid/block → Chapel pre-1.31 currently only supports 1D forall
 - Had to linearize to 1D with a *very large* index space
 - for-by loops → Chapel by clause doesn't GPU-ize yet
 - GPU codes often increment by the thread count to keep co-executing threads aligned to memory
 - Had to replace with while-count



(↑ pseudo-code, see Github for real)

```

var intersect : [0..<numEdges] atomic real(32);
forall Z in srcIters by gridDim.z*blockDim.z {
  forall Y in destIters by gridDim.y*blockDim.y {
    forall X in isectIters by gridDim.x*blockDim.x {
      ... // bin-search
      intersect[writeAddr].add(1.0);
    }
  }
}

```

VC-CPU-Linearized

```

var intersect : [0..<numEdges] atomic real(32);
forall id in srcIters*destIters*isectIters {
  var nd_id : 3*int = get_ND_ID(id);
  var zCount = nd_id(2);
  while (zCount < zMax) {
    var yCount = nd_id(1);
    while (yCount < yMax) {
      var xCount = nd_id(0);
      while (xCount < xMax) {
        ... // bin-search
        intersect[writeAddr].add(1.0);
        xCount += gridDim.x*blockDim.x; }
      yCount += gridDim.y*blockDim.y; }
    zCount += gridDim.z*blockDim.z; }
}

```

VC-GPU


```

//replace atomic type .add() w/ extern call
ex_atomicAdd(c_ptrTo(intersect), writeAddr, 1.0)

```

(↑ pseudo-code, see Github for real)

Vertex-Centric: More complex



- Vertex-centric intersect didn't initially GPU-ize
 - 3D grid/block → Chapel pre-1.31 currently only supports 1D forall
 - Had to linearize to 1D with a *very large* index space
 - for-by loops → Chapel by clause doesn't GPU-ize yet
 - GPU codes often increment by the thread count to keep co-executing threads aligned to memory
 - Had to replace with while-count
 - Accumulate via atomicAdd → Chapel atomics don't GPU-ize yet
 - Had to call CUDA's via extern C

```

var intersect : [0..

```

VC-CPU-Linearized

```

var intersect : [0..

```

VC-GPU


```

//replace atomic type .add() w/ extern call
extern atomicAdd(c_ptrTo(intersect), writeAddr, 1.0)

```

(↑ pseudo-code, see Github for real)

Vertex-Centric: More complex



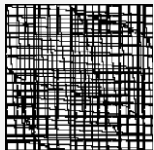


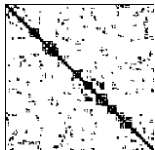
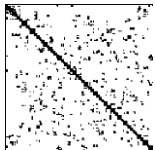

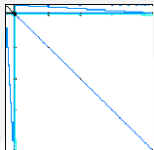
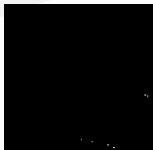
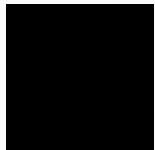
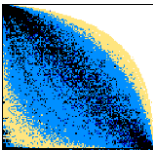
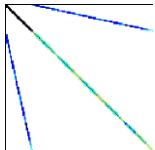
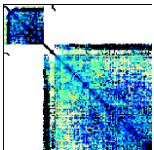
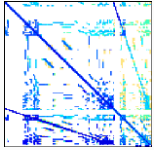
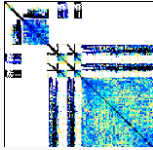


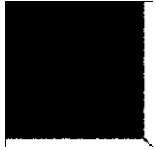
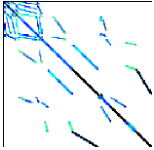
- Vertex-centric intersect didn't initially GPU-ize
 - 3D grid/block → Chapel pre-1.31 currently only supports 1D forall
 - Had to linearize to 1D with a *very large* index space
 - for-by loops → Chapel by clause doesn't GPU-ize yet
 - GPU codes often increment by the thread count to keep co-executing threads aligned to memory
 - Had to replace with while-count
 - Accumulate via atomicAdd → Chapel atomics don't GPU-ize yet
 - Had to call CUDA's via extern C
- But we *were* able to manage it
 - And Chapel mapped intermediate kernels to CPU → supporting *incremental validation*



How well did it perform?

Test Data

Sparsest to densest

<p>Kmer_A2a Protein k-mers V: 171M E: 361M Avg: 2.11 Range: 39 Std. Dev.: 0.56 Gini Index: 0.055</p> 	<p>Europe_osm European roads V: 50.9M E: 108M Avg: 2.12 Range: 12 Std. Dev.: 0.48 Gini Index: 0.085</p> 	<p>Road_usa US roads V: 23.9M E: 57.7M Avg: 2.41 Range: 8 Std. Dev.: 0.93 Gini Index: 0.211</p> 	<p>Road-roadNet-CA California roads V: 1.96M E: 5.52M Avg: 2.82 Range: 11 Std. Dev.: 0.99 Gini Index: 0.185</p> 	<p>Road-roadNet-PA Pennsylvania roads V: 1.09M E: 3.08M Avg: 2.83 Range: 8 Std. Dev.: 1.02 Gini Index: 0.188</p> 	<p>Delaunay_n24 Random Triangulations V: 16.8M E: 101M Avg: 6.00 Range: 23 Std. Dev.: 1.34 Gini Index: 0.122</p> 
<p>circuit5M Large circuit V: 5.56M E: 54.0M Avg: 9.71 Range: 1.29M Std. Dev.: 1357 Gini Index: 0.577</p> 	<p>Soc-LiveJournal1 Social network V: 4.85M E: 85.7M Avg: 17.7 Range: 20.3K Std. Dev.: 52.0 Gini Index: 0.711</p> 	<p>Wikipedia-20070206 Web page links V: 3.57M E: 84.8M Avg: 23.8 Range: 188K Std. Dev.: 255 Gini Index: 0.759</p> 	<p>GL7d19 Voronoi differentials V: 1.96M E: 74.6M Avg: 38.2 Range: 134 Std. Dev.: 6.73 Gini Index: 0.088</p> 	<p>dielFilterV2real Dielectric resonator V: 1.16M E: 47.4M Avg: 40.9 Range: 104 Std. Dev.: 16.1 Gini Index: 0.201</p> 	<p>Sc-Idoor Large door V: 952K E: 41.5M Avg: 43.6 Range: 76 Std. Dev.: 14.8 Gini Index: 0.183</p> 
<p>Stokes VLSI process sim. V: 11.4M E: 516M Avg: 45.1 Range: 1728 Std. Dev.: 61.8 Gini Index: 0.392</p> 	<p>Sc-msdoor Medium Door V: 416K E: 18.8M Avg: 45.1 Range: 76 Std. Dev.: 13.7 Gini Index: 0.166</p> 	<p>Ca-coauthors-dblp Coauthorship V: 540K E: 30.5M Avg: 56.4 Range: 3298 Std. Dev.: 66.2 Gini Index: 0.544</p> 	<p>Soc-orkut Social network V: 3.00M E: 213M Avg: 71.0 Range: 27.5K Std. Dev.: 140 Gini Index: 0.558</p> 	<p>Hollywood-2009 Costarring Actors V: 1.14M E: 113M Avg: 98.9 Range: 11.5K Std. Dev.: 272 Gini Index: 0.750</p> 	<p>HV15R CFD of engine fan V: 2.02M E: 325M Avg: 161 Range: 491 Std. Dev.: 47.8 Gini Index: 0.155</p> 

Graph data and images CC-BY-4.0 from the SparseSuite Matrix Collection (<https://sparse.tamu.edu/>).

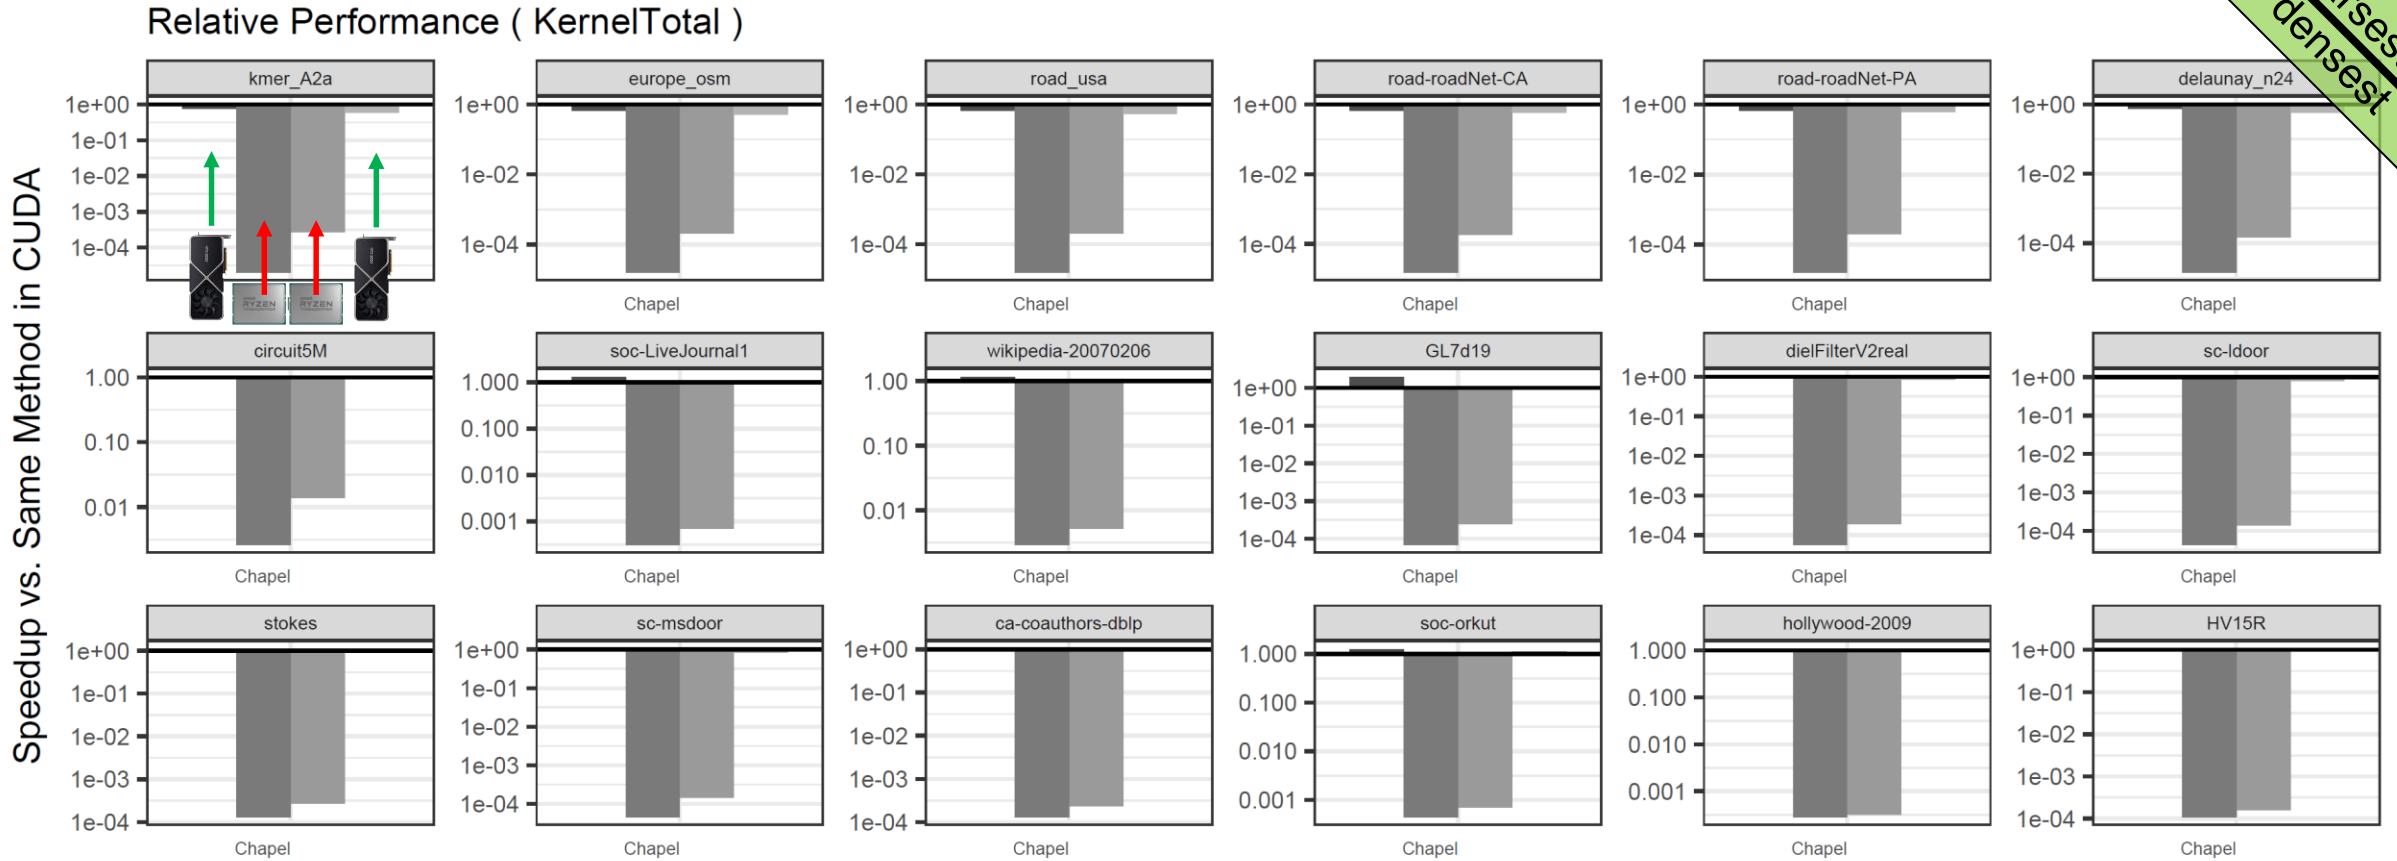
Preprocessed CSR binary files: <https://chrec.cs.vt.edu/SYCL-Jaccard/HPEC22-Data/index.html>

Fallback of running GPU forall on CPU is *functionally-portable*

- Great feature for validation, but don't expect GPU-tuned impl's to be *performance-portable* to CPU

Sparsest to densest

Higher is better!



CPU: AMD Threadripper 3960X
GPU: Nvidia RTX 3090
CUDA: 11.6 / driver 510.108.03
Chapel: pre-1.31 (d7664c9d81)


Chapel Version EC-GPU VC-CPU-Naive VC-CPU-Linearized VC-GPU

Fallback of running GPU forall on CPU is *functionally-portable*

- Great feature for validation, but don't expect GPU-tuned impl's to be *performance-portable* to CPU
- But if we remove the fallback CPU columns...

Sparsest to densest

Higher is better!

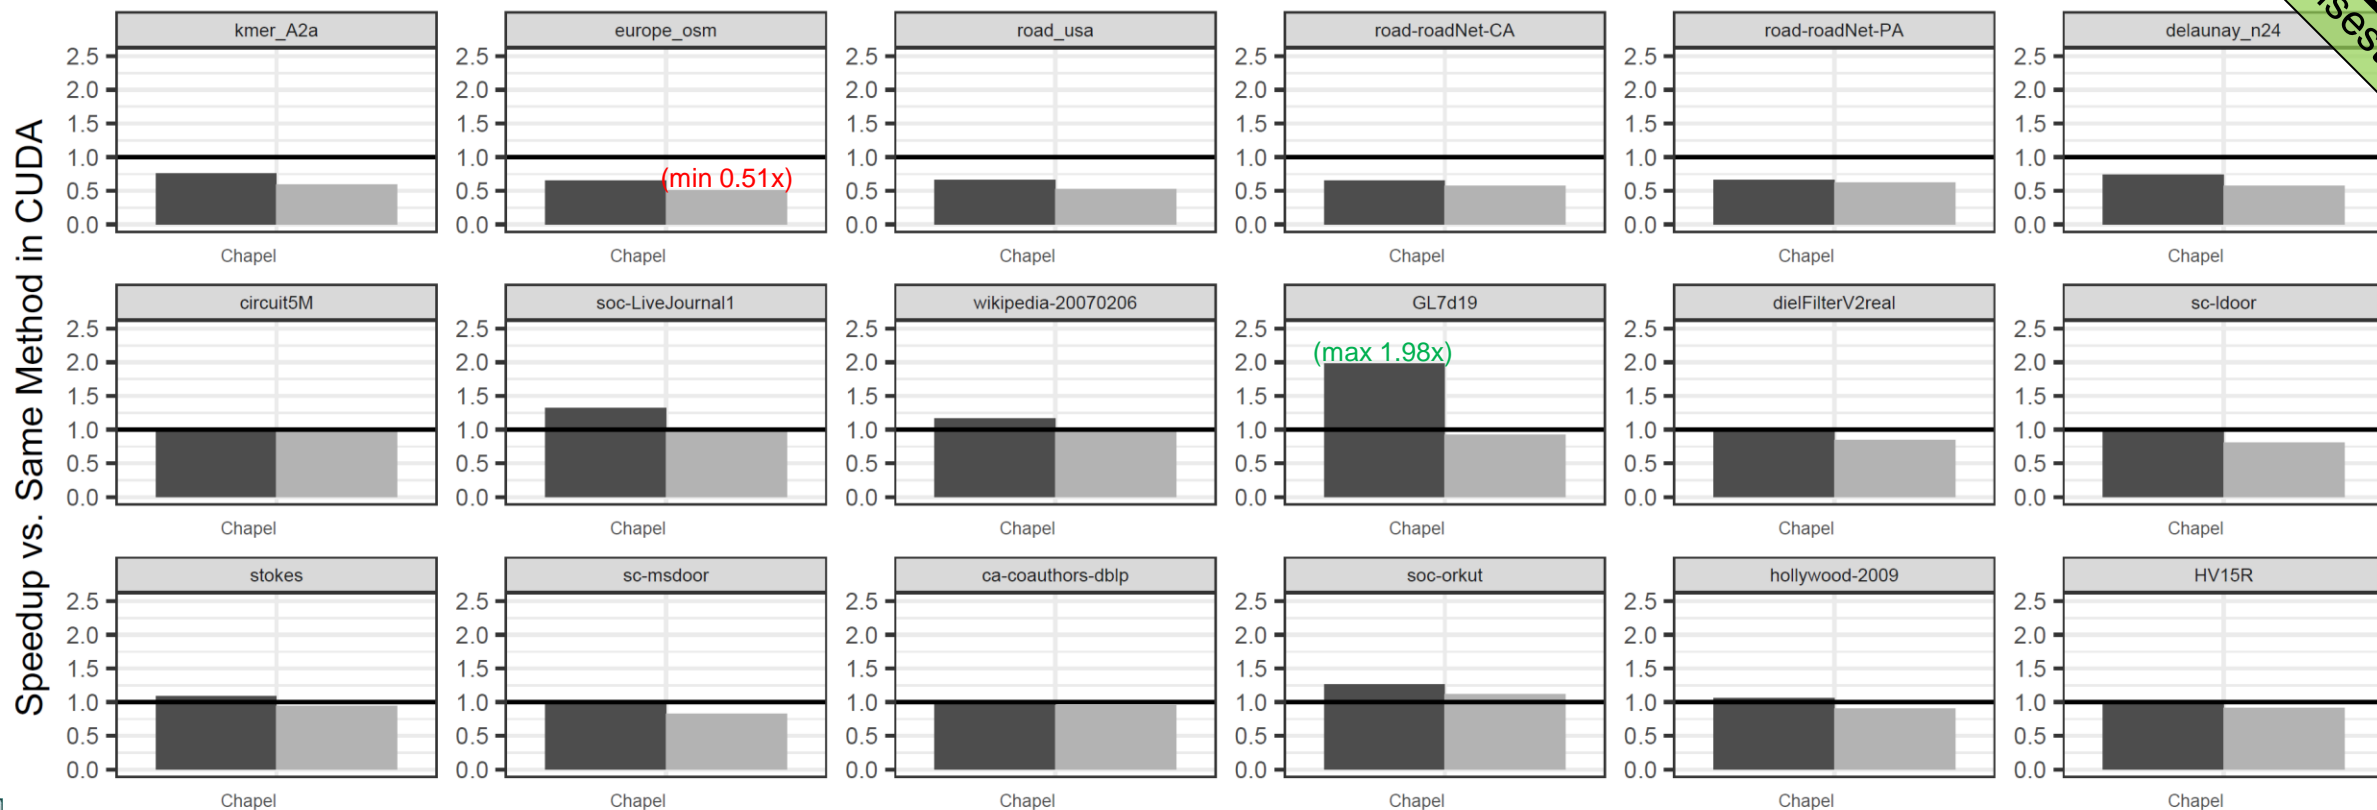


CPU: AMD Threadripper 3960X
GPU: Nvidia RTX 3090
CUDA: 11.6 / driver 510.108.03
Chapel: pre-1.31 (d7664c9d81)

Chapel Version EC-GPU VC-CPU-Naive VC-CPU-Linearized VC-GPU

GPU Performance is pretty good right out of the box

Relative Performance (KernelTotal)



Higher is better!

Sparsest to densest

CPU: AMD Threadripper 3960X CUDA: 11.6 / driver 510.108.03
 GPU: Nvidia RTX 3090 Chapel: pre-1.31 (d7664c9d81)

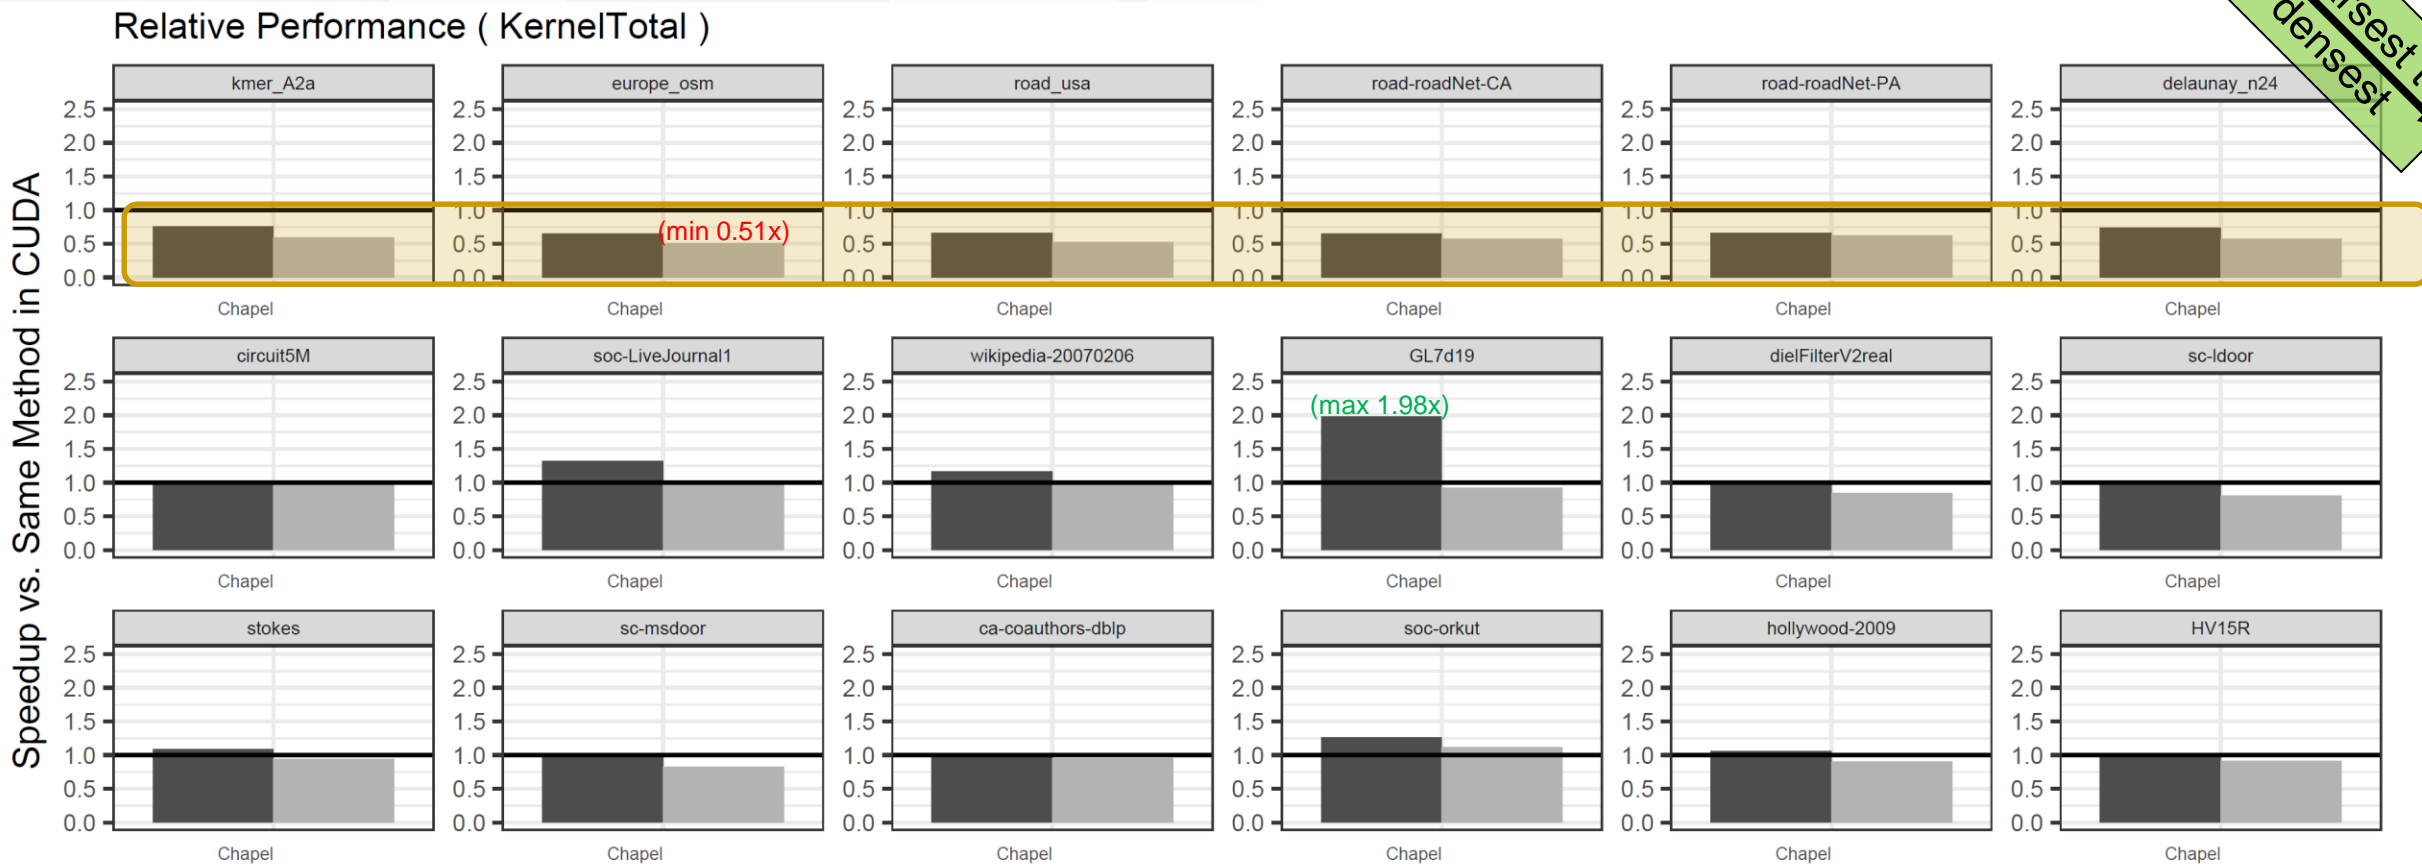
Chapel Version EC-GPU VC-GPU

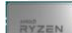

GPU Performance is pretty good right out of the box

- Room for tuning on the sparse end

Higher is better!

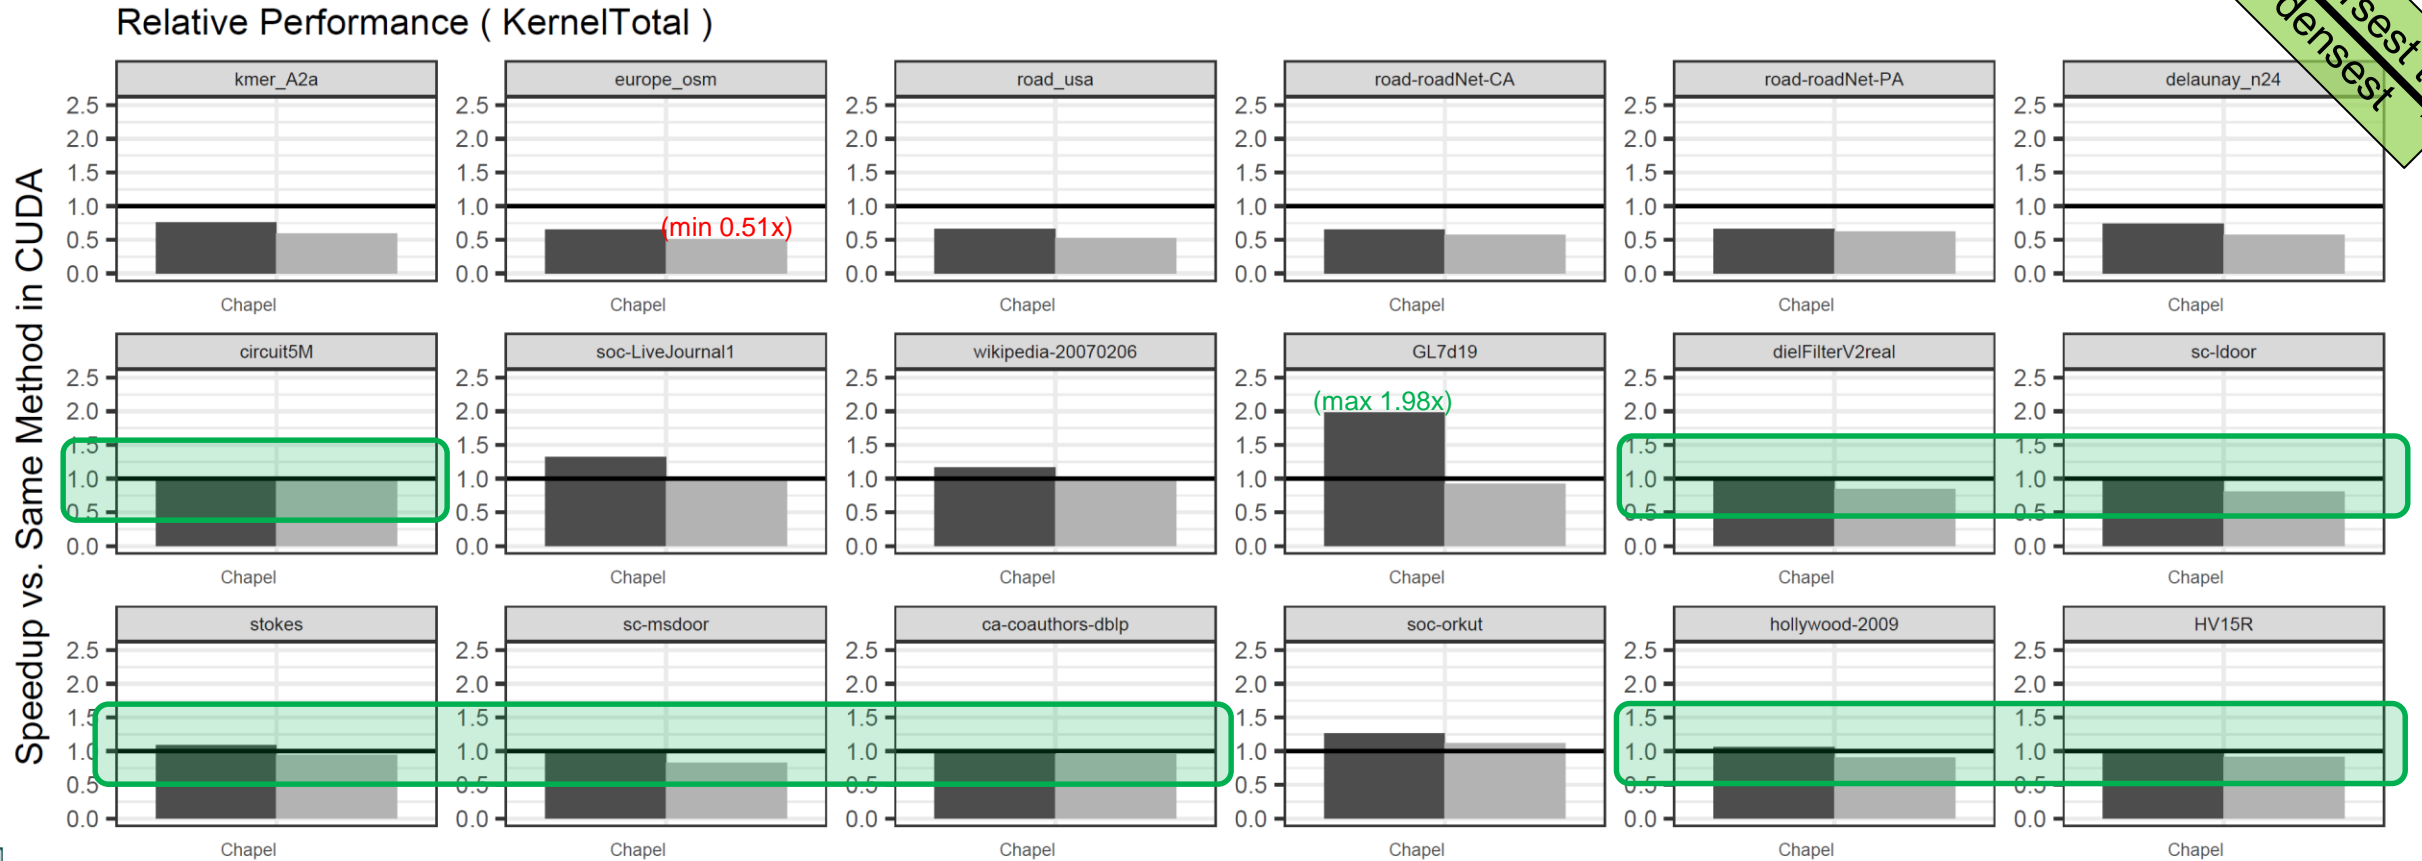
Sparsest to densest



 CPU: AMD Threadripper 3960X CUDA: 11.6 / driver 510.108.03
 GPU: Nvidia RTX 3090 Chapel: pre-1.31 (d7664c9d81)

GPU Performance is pretty good right out of the box

- Room for tuning on the sparse end, but approaching parity on the denser end



CPU: AMD Threadripper 3960X CUDA: 11.6 / driver 510.108.03
 GPU: Nvidia RTX 3090 Chapel: pre-1.31 (d7664c9d81)


Chapel Version EC-GPU VC-GPU

GPU Performance is pretty good right out of the box

- Room for tuning on the sparse end, but approaching parity on the denser end
- Chapel *outperforms* CUDA on a few inputs → Future work: *Why?* → tune CUDA

Sparsest to densest

Relative Performance (KernelTotal)



Higher is better!

CPU: AMD Threadripper 3960X
GPU: Nvidia RTX 3090
CUDA: 11.6 / driver 510.108.03
Chapel: pre-1.31 (d7664c9d81)

Chapel Version EC-GPU VC-GPU

What's next?

Future Work

- Understand performance gap and gains vs. CUDA
- Make more Chapel-idiomatic
 - Kernels: use array promotion
 - Generic CSR: abstract base class
- Implement other JS algorithms / optimizations
 - Consider distribution

Wishlist

- Chapel version of *clang-format*
- A pragma or attribute to name my kernels (for profiling & other tools)
 - Instead of picking out the right `chpl_gpu_kernel##` from assembly
- Native 2D/3D loops
- Warp/wavefront & sync primitives

In Conclusion / Q&A

- We need more *programmable* GPU languages → Chapel looks good to this GPU dev!
- Chapel was reasonable to port to with a “C on GPUs” background
 - Kernels: Embarrassingly-parallel is easy, thread-collaboration less-so but achievable *today*
- GPU Chapel’s programmability is less verbose than CUDA
 - “Chapel like a C programmer”: kernels take 96% as many lines, and the whole program only 62%
 - Should improve as we familiarize → revisit kernels for promotion and generic CSR representation
- GPU Chapel’s performance is slightly slower than CUDA, but within a shout
 - Between **0.51x** and **1.98x** performance, geo-mean **0.87x**
- I am excited about Chapel and would recommend as an easier route to custom GPU kernels



GitHub

Code: <https://github.com/vtsynergy/Chapel-Examples>
Input data: <https://chrec.cs.vt.edu/SYCL-Jaccard/HPEC22-Data/index.html>




The work detailed herein has been supported in part by NSF I/UCRC CNS-1822080 via the NSF Center for Space, High-performance, and Resilient Computing (SHREC).

30

Backup Slides

Rudimentary *Explicit* RTTI: How we did it (C-style)



```
record CSR_header {
  // Disk binary format of bitfield
  // runtime type descriptor flags
}
```

```
record CSR_descriptor {
  // Chapel bool `var` members
  // to describe runtime type
}
```

```
proc foo(...) {
  var myDesc = readHeader(file) :
  CSR_descriptor;
  var myCSR = makeCSR(myDesc);
  readArrays(file, myCSR);
  ...
}
```

Kernels and I/O func's "ladder-resolve" like this

```
//Resolve one runtime `var` to `param` per overload
proc fooOnCSR(in h : CSR_handle) {
  //resolve 1st param
  if (h.var1) then fooOnCSR(h, true) else fooOnCSR(h, false);
}
private proc fooOnCSR(in h : CSR_handle, param parm1 : bool) {
  ... // resolve 2nd param }
  ... // resolve 3rd..Nth params
private proc fooOnCSR(in h : CSR_handle, param parm1 : bool, ...
param parmN : bool) {
  doFooOnCSR(CSR(parm1, ..., parmN), h);
}


//Generic worker function
private proc doFooOnCSR(type csr_type : unmanaged CSR(?), in h :
CSR_handle) {
  assert(csr_type == h.desc, ...);
  // Do Foo
}
```

```
proc makeCSR(in desc :
CSR_descriptor) : CSR_handle
{
  ... // resolve and allocate
  concrete instantiation of
  CSR generic class
}
```

```
record CSR_handle {
  var desc : CSR_descriptor;
  var data : c_void_ptr;
}
```

```
class CSR {
  // Chapel bool `param` members
  // to instantiate concrete type
  // Chapel parameterized arrays to
  // contain vertex, edge, weight data
}
```

Rudimentary RTTI: How we'd redo it (Chapel/OO-style)



```
record CSR_header {
  // Disk binary format of bitfield
  // runtime type descriptor flags
}
```

```
class CSR_base {
  // Chapel bool `var` members
  // to instantiate concrete type
}
```

```
proc foo(...) {
  var myDesc = readHeader(file) :
  CSR_base;
  var myCSR = makeCSR(myDesc);
  readArrays(file, myCSR);
  ...
}
```

Kernels and I/O func's "ladder-resolve" like this

```
//Resolve one runtime `var` to `param` per overload
proc fooOnCSR(in h : CSR_base) {
  //resolve 1st param
  if (h.var1) then fooOnCSR(h, true) else fooOnCSR(h, false);
}
private proc fooOnCSR(in h : CSR_base, param parm1 : bool) {
  ... // resolve 2nd param }
  ... // resolve 3rd..Nth params
private proc fooOnCSR(in h : CSR_base, param parm1 : bool, ...
param parmN : bool) {
  doFooOnCSR(CSR_arrays(parm1, ..., parmN), h);
}

//Generic worker function
private proc doFooOnCSR(type csr_type : unmanaged CSR_arrays(?),
in h : CSR_base) {
  assert(csr_type.var1 == h.var1, ...);
  // Do Foo
}
```

```
proc makeCSR(in desc :
CSR_base) : unmanaged
CSR_base {
  ... // resolve and allocate
  concrete instantiation of
  CSR_arrays generic class
}
```

```
class CSR_arrays : CSR_base {
  // Chapel parameterized arrays to
  // contain vertex, edge, weight data
}
```