

HyperFlow: An Efficient Dataflow Architecture for Multi CPU-GPU Systems



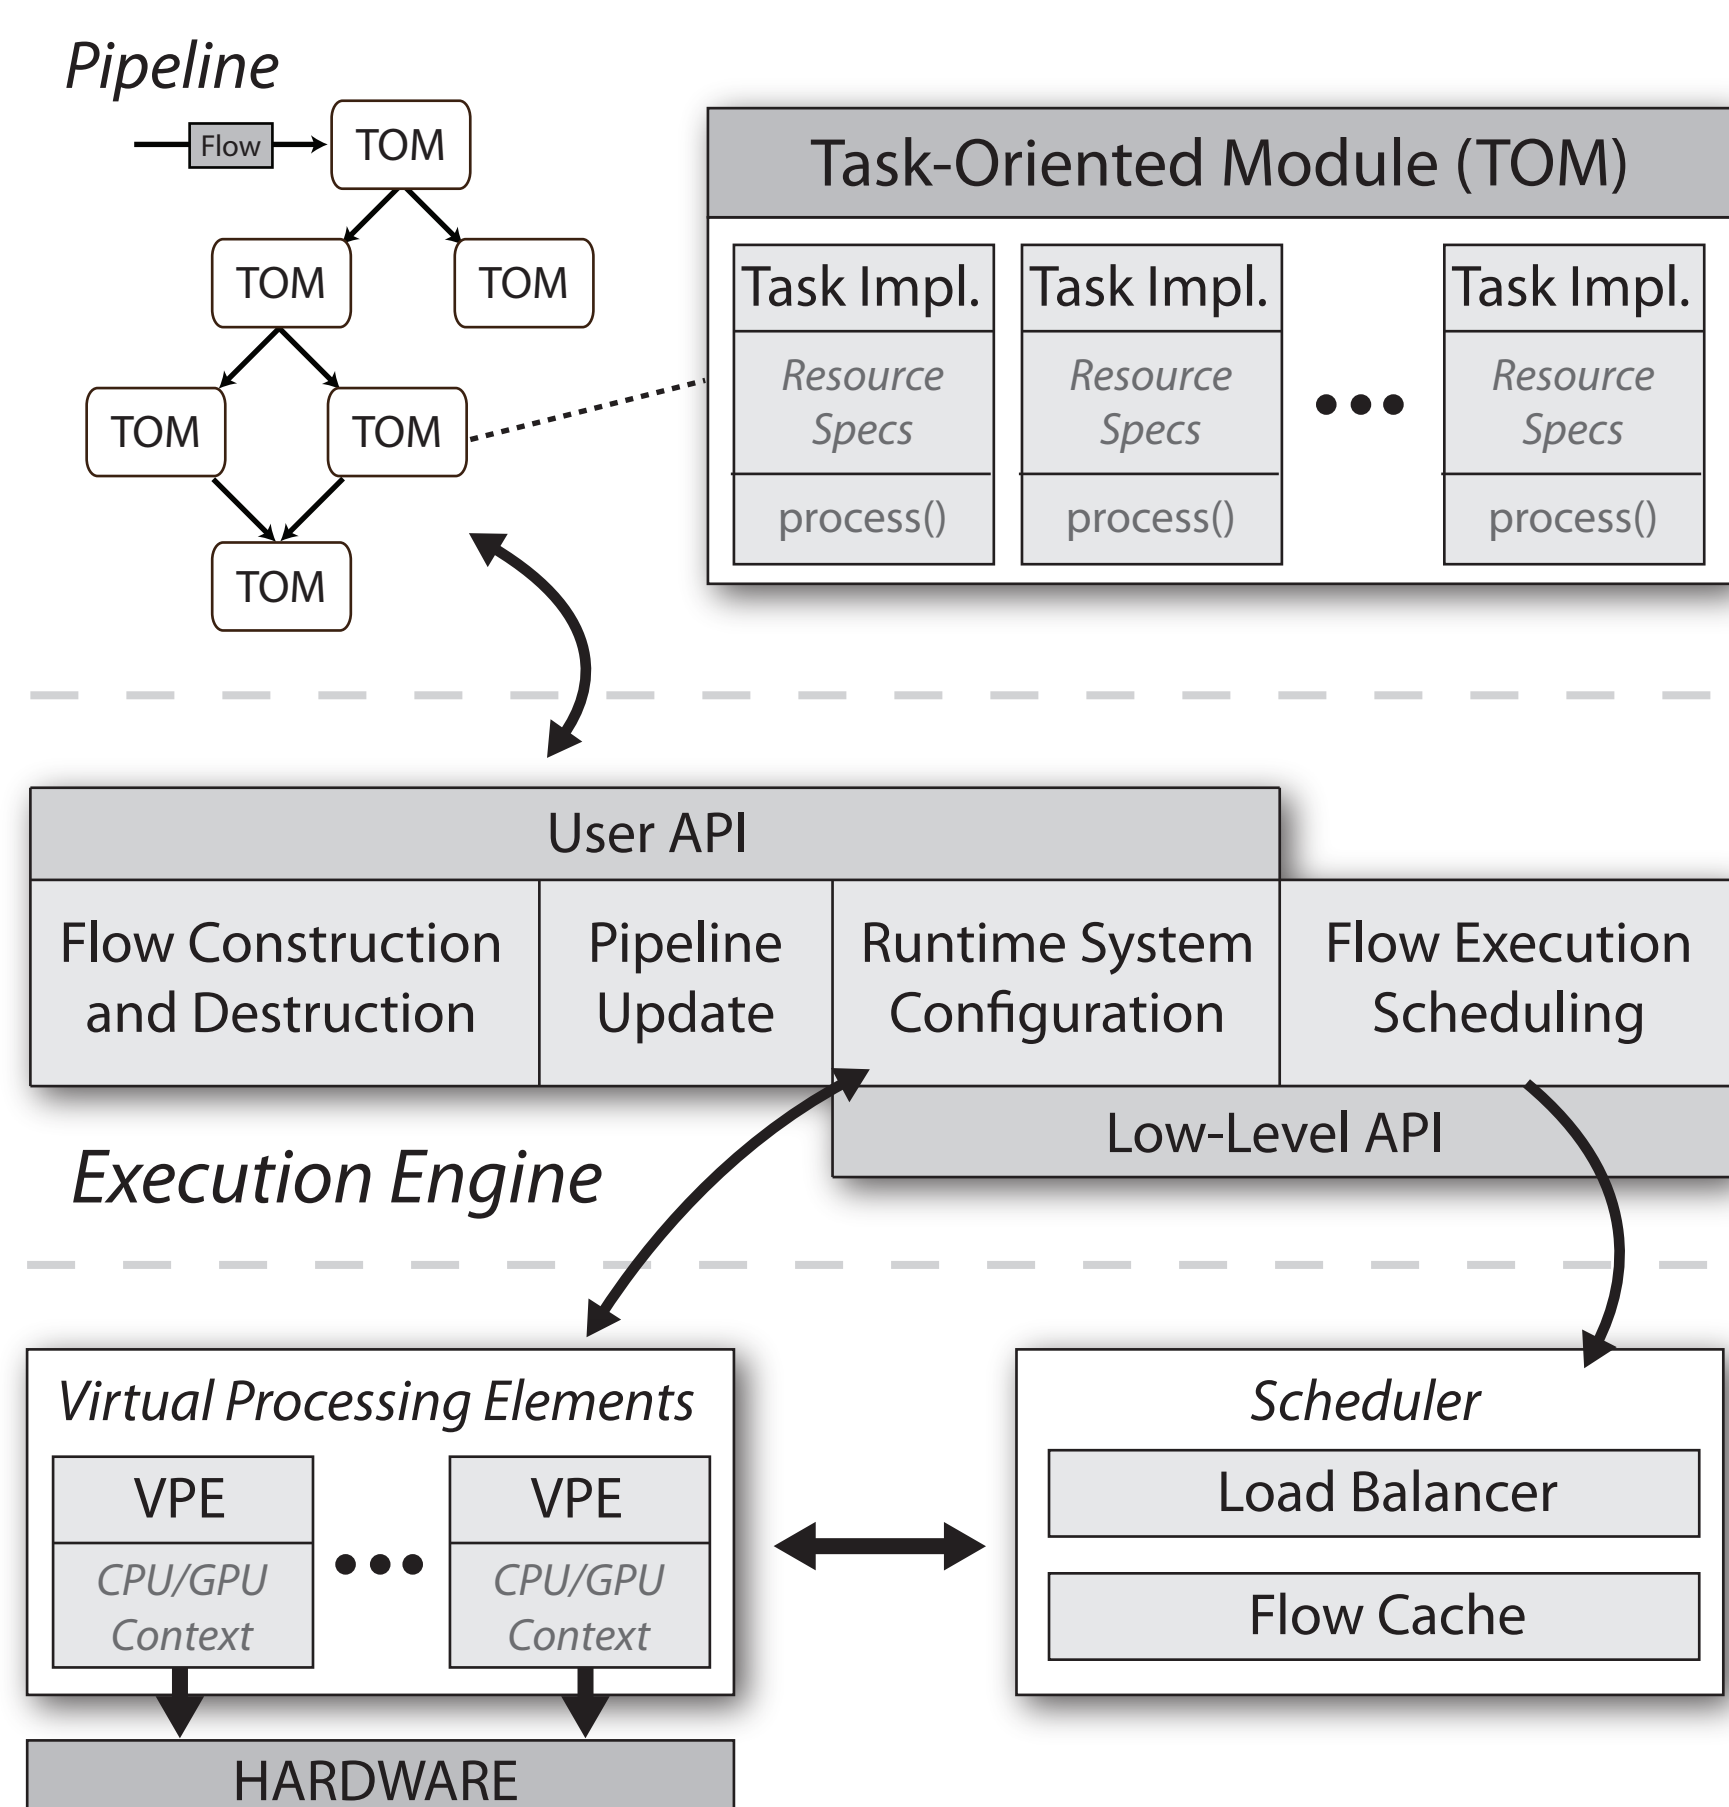
Huy Vo¹, Daniel Osmari², Luiz Scheidegger², João Comba², Jason Shepherd³ and Cláudio Silva¹

¹ University of Utah, USA ² UFRGS, Brazil ³ Sandia National Laboratories, USA

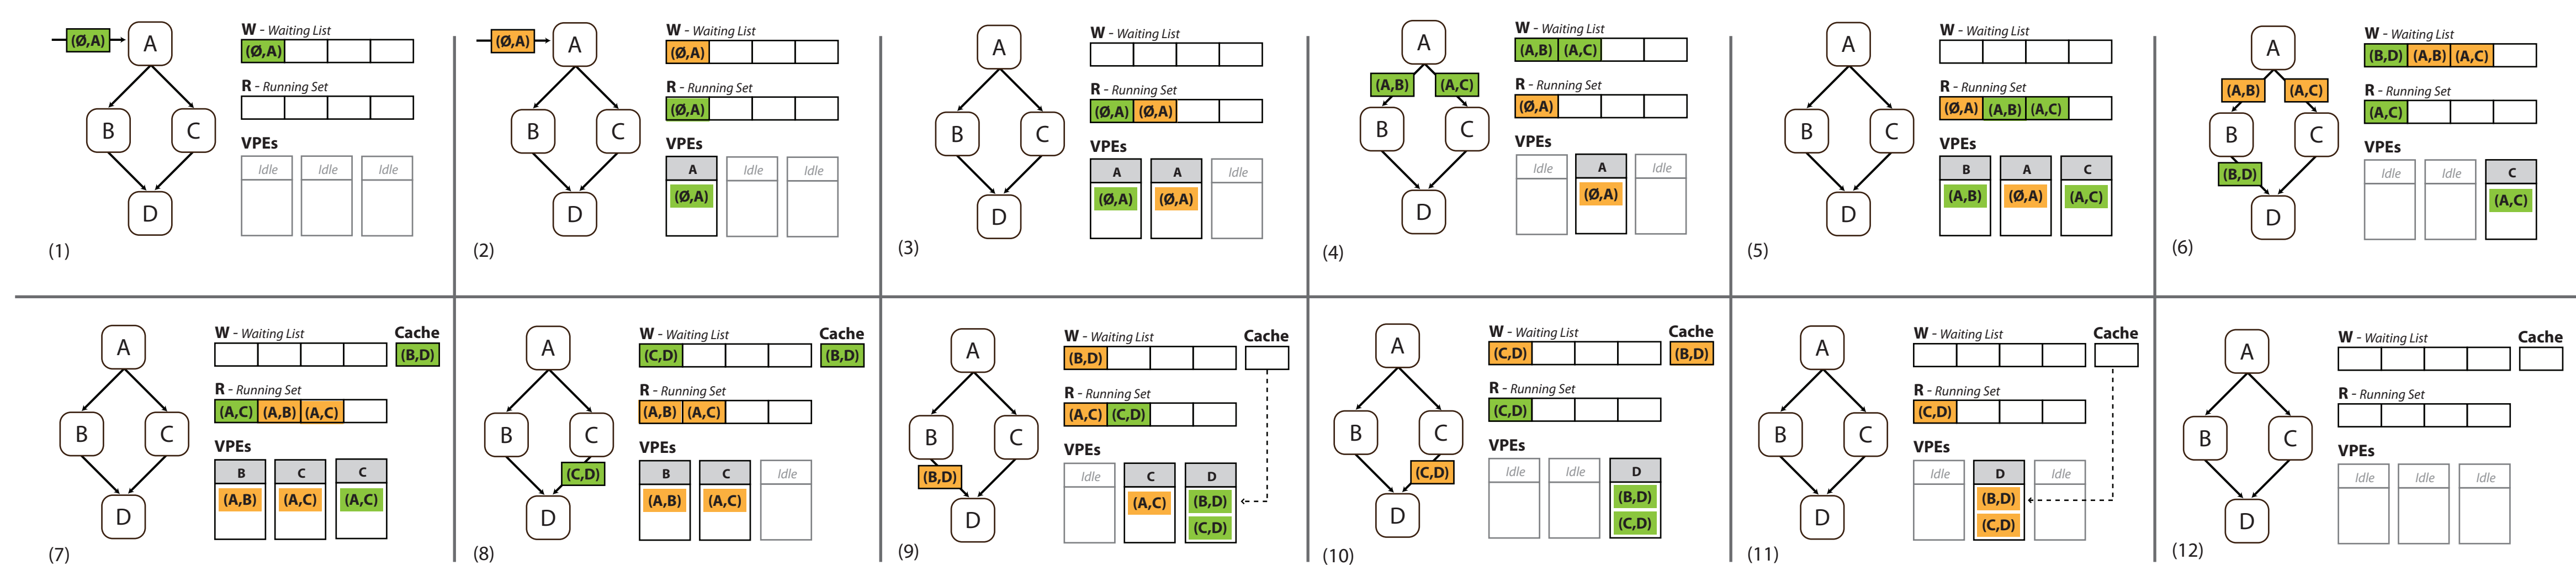
WHY HYPERFLOW?

- Dataflow is widely used in scientific computing
- Current dataflow systems lack parallelism support:
 - Mostly designed for single-core machines
 - Some only support data-parallelism with MPI
 - Ad-hoc for multi-core and (multi-)GPU platforms
- There exists parallel frameworks such as TBB, StarPU Stream SDK, etc. but difficult to apply to dataflow due to being tightly coupled with task-graph execution.

SYSTEM ARCHITECTURE





PIPELINE EXECUTION



SHOW CASE 1 - HIGH-THROUGHPUT IMAGE PROCESSING

Edge Detection Pipeline



Machine Specs:

- 2 x Xeon w5580, 8-core
24GB RAM
Quadro 5800
- Intel 970i, 4-core
6GB RAM
GTX 295, Tesla C1060

Timing Results

Input: 312 images ~ 6GB

#CPU Cores	#GPU Cores	Time (s)
1	0	3261
2	0	1747
4	0	1049
8	0	668
8	1	440
8	2	272
4	3	151

TASK-ORIENTED MODULE (TOM)

- Hold meta-data for its task and parameters
- Manage multiple implementations for a specific task
- Each implementation specifies its own resources

```

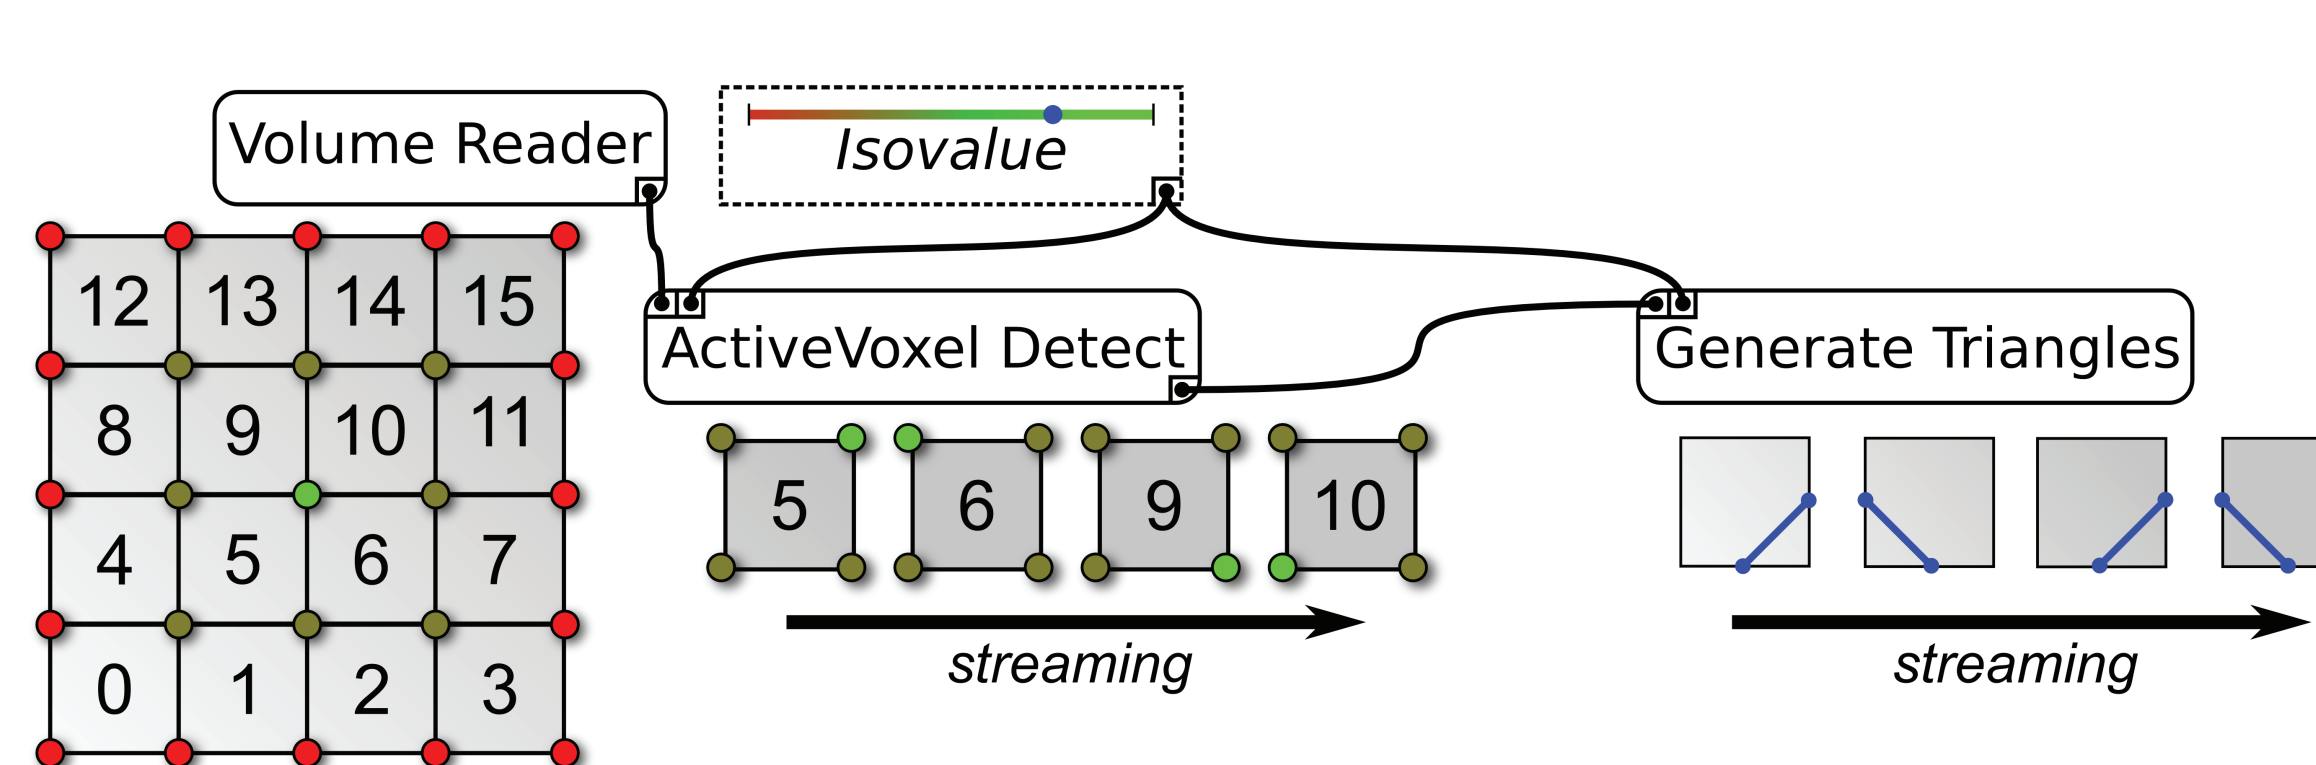
1class TaskOrientedModule {
2public:
3    TaskOrientedModule(int numberOfInputPorts,
4                       int numberOfOutputPorts);
5    void addImplementation(TaskImplementation* impl);
6    static void connect(
7        TaskOrientedModule* srcModule, int srcPort,
8        TaskOrientedModule* dstModule, int dstPort
9    );
10   // Additional parameters if subclassed
11   ...
12};
13
14// Construction of the pipeline
15TaskOrientedModule Source(0, 1, "Image Reader")
16Source.addImplementation(new SourceImpl());
17
18TaskOrientedModule Filter(1, 1, "Gaussian Blur")
19Filter.addImplementation(new FilterImpl());
20
21TaskOrientedModule Sink(1, 0, "Viewer")
22Sink.addImplementation(new SinkImpl());
23
24TaskOrientedModule::connect(&Source, 0, &Filter, 0);
25TaskOrientedModule::connect(&Filter, 0, &Sink, 0);
    
```

SHOW CASE 2 - HYBRID ISOSURFACE EXTRACTION

The pipeline starts with a dataset reader that loads blocks of the volume and stream them to a TOM classifier. The classifier module traverses each block to generate a list of active voxels, i.e., only voxels that may contain the desired isosurface (voxels that do not contain v can be safely discarded). The list of active voxels is progressively streamed to a triangle generation TOM, which is designed to inspect each voxel and, based on a standard Marching Cubes case table, construct a set of triangles that approximates the desired isosurface inside that voxel. This set of triangles can also be streamed to other modules for further post-processing.

Using HyperFlow, it becomes very easy to optimize the usage of CPU and GPU computational resources. We implemented the dataset reader to divide the input volume into blocks. The execution scheduler automatically instantiates separate, concurrent TOMs to perform voxel classification on each block. Since there is no data dependency between different instances, this parallelism greatly improves algorithm performance. The TOMs designed to construct the actual isosurface can also be trivially parallelized by HyperFlow.

Isosurface Extraction Pipeline



Timing Results

Input: 1 billion voxels

#CPU Cores	#GPU Cores	Time (s)
1	1	132
1	2	101
2	1	125
2	2	72
4	1	127
4	2	71
4	3	49
8	1	113
8	2	69

MULTI-MEDIUM DATA INTERFACE

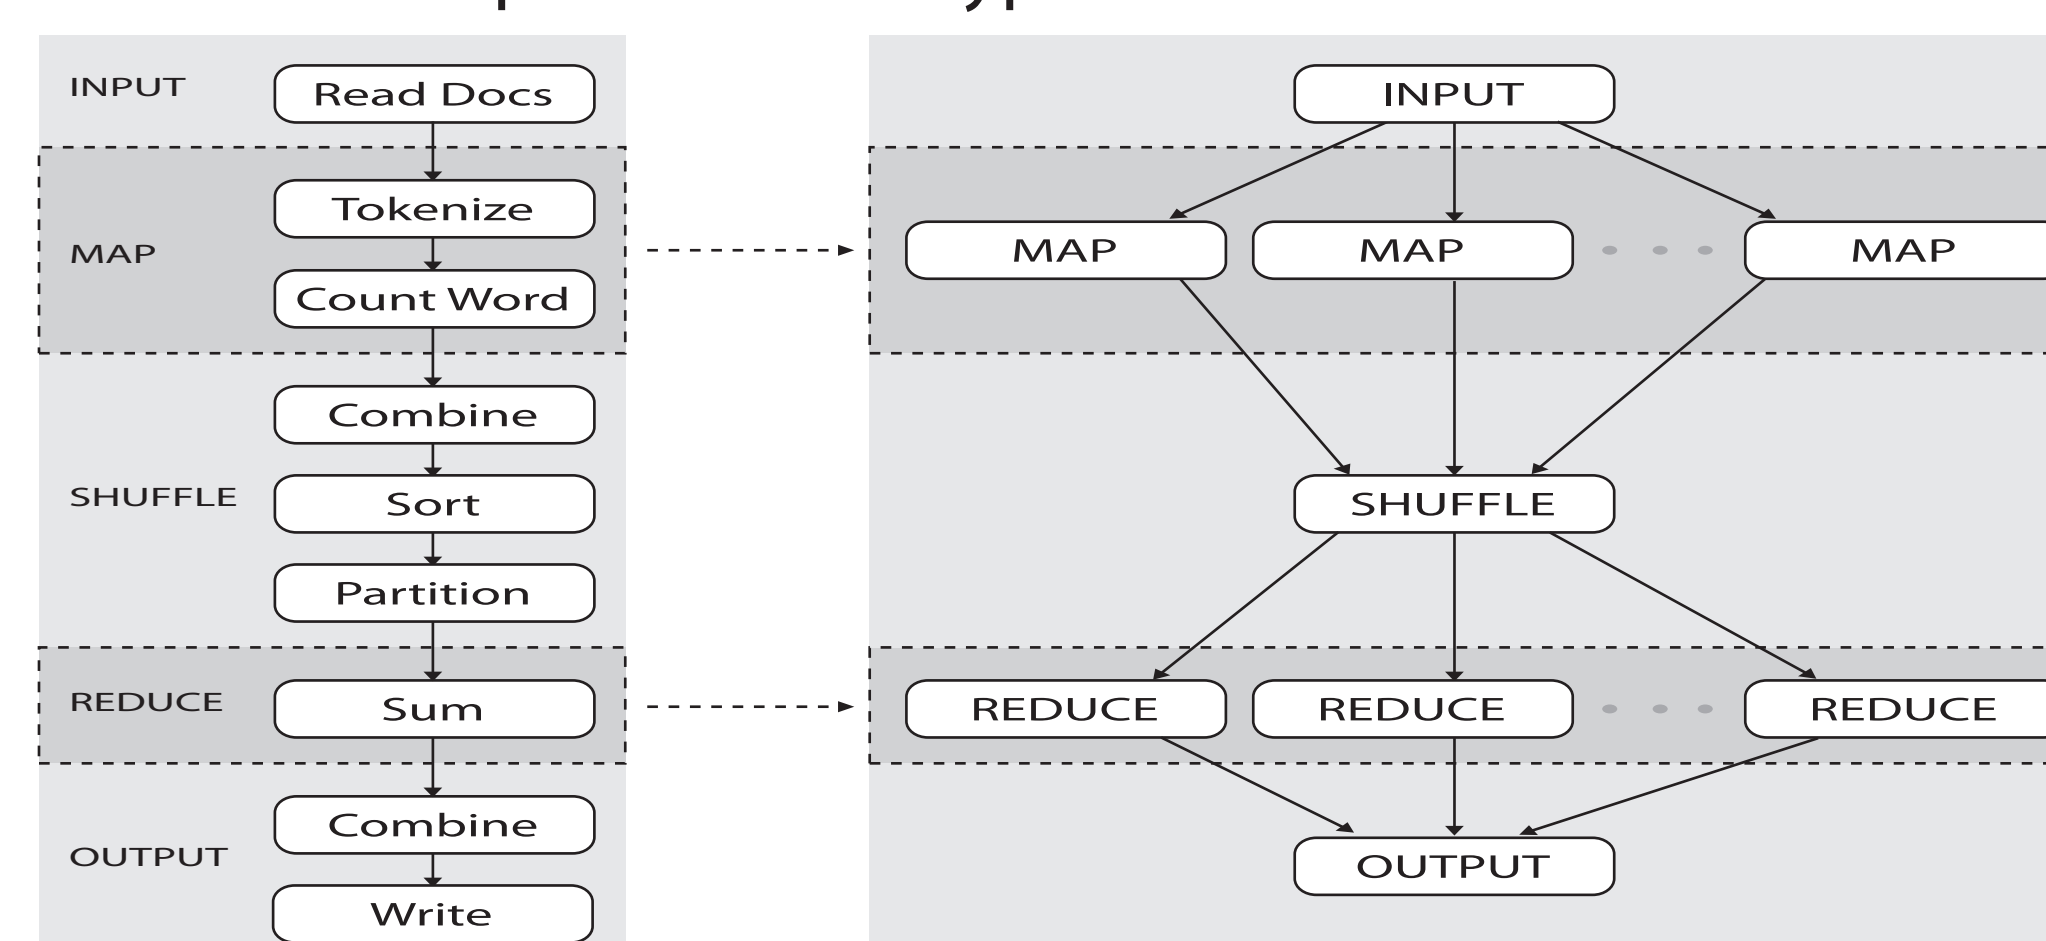
- Data explicitly denotes which platform to reside in
- Data transferring is triggered automatically but conversions are done by user-define functions

```

1class DImageData : public Data {
2public:
3    DImageData(BaseImage *img, DMTYPE medium) {
4        // Store image with the medium type
5        // ...
6    }
7    virtual Data* createInstance(DMTYPE targetMedium) {
8        // Convert from CPU Image to CUDA Image
9        if (this->medium==DM_CPU_MEMORY &&
10            targetMedium==DM_GPU_MEMORY) {
11            CUDAImage *newImage = new CUDAImage();
12            cudaMemcpy(newImage, this->image,
13                      cudaMemcpyHostToDevice);
14            return new DImageData(newImage, DM_GPU_MEMORY);
15        }
16        // Convert from CUDA Image to CPU Image
17        else if (this->medium==DM_GPU_MEMORY &&
18                targetMedium==DM_CPU_MEMORY) {
19            CPUImage *newImage = new CPUImage();
20            cudaMemcpy(newImage, this->image,
21                      cudaMemcpyDeviceToHost);
22            return new DImageData(newImage, DM_CPU_MEMORY);
23        }
24        // Use default transfer methods for CPU-CPU, etc.
25        else
26            return Data::createInstance(medium);
27    }
28};
    
```

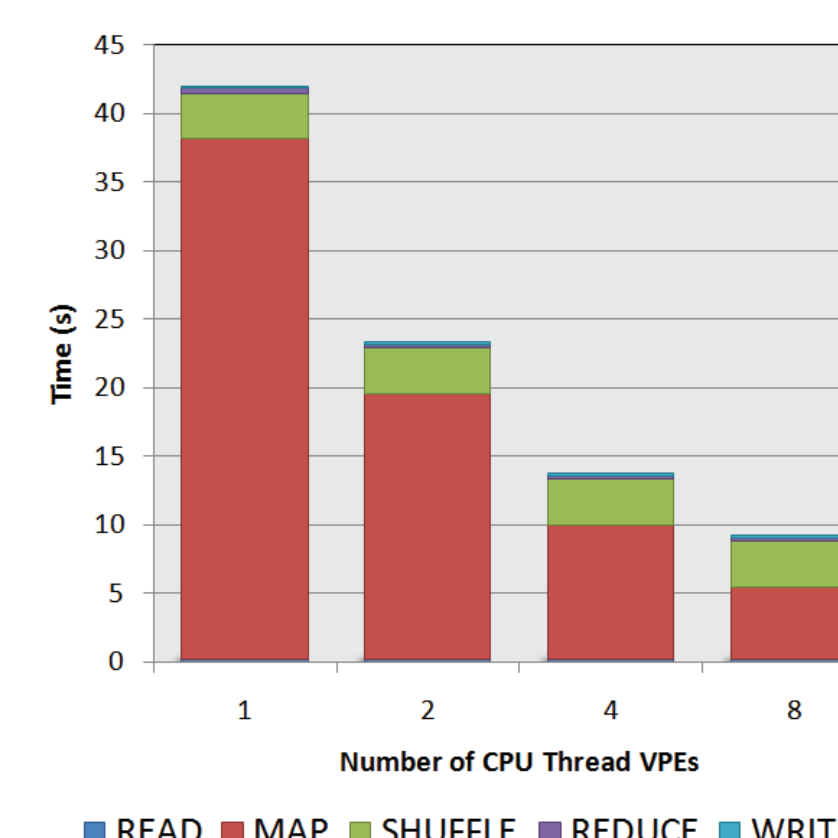
SHOW CASE 3 - MAP/REDUCE WORD COUNT

Map/Reduce in HyperFlow



Timing Results

Input: 83 million words



```

1class IReader: public TaskImplementation {
2    virtual int process(ExecutionContext *ctx,
3                      TaskData *input, TaskData *output) {
4        Data *data[N];
5        // Read input file and place them into data[]
6        ...
7        // Then construct flows downstream
8        for (int i=0; i<N; i++) {
9            // We are not passing any reference flow so that
10           // newly created flows will have unique id
11           Flow *flow = ctx->createFlowForOutputPort(0);
12           flow->data = data[i];
13           output->addFlow(flow);
14       }
15       return 0;
16   };
17};
18
19class MCombine: public TaskOrientedModule {
20public:
21    int numberOfMappers;
22    SafeCounter doneCount;
23    Mutex mutex;
24    ListOfData data;
25};
26
27class ICombine: public TaskImplementation {
28    virtual int process(ExecutionContext *ctx,
29                      TaskData *input, TaskData *output) {
30        MCombiner *M = (MCombiner*)ctx->module;
31        // Locking resource to accumulate
32        M->mutex.lock();
33        M->data.push_back(input->flows[0]->data);
34        M->mutex.unlock();
35    }
36    if (M->doneCounter->increase() < M->numberOfMappers)
37        return -1;
38    // Process data in module->data
39    ...
40    return 0;
41};
42
43};
    
```

Two execution scenarios need to be addressed: (1) an upstream module, which executes sequentially, but generates output for concurrent execution downstream and (2) a downstream module which can only be executed sequentially and has to collect flows from upstream modules executed in parallel. Scenario (1) occurs when data is moved from the reader to mappers, as well as from the shuffling phase to reducers. This is necessary to enable data parallelism in the Map and Reduce phases. Scenario (2) happens when the execution comes back from the Map and Reduce to the Shuffling and Writer. HyperFlow supports both scenarios: (1) an upstream can generate multiple flows with unique identifications to trigger data-parallelism in the pipeline; (2) TOM can limit task implementations to run sequentially by using a global lock. These are illustrated in the source code on the right.