

Washington University in St. Louis
Washington University Open Scholarship

All Computer Science and Engineering Research

Computer Science and Engineering

Report Number: WUCSE-2016-003

2016-7

MERCATOR (Mapping EnumERATOR for CUDA) User's Manual

Authors: Stephen V. Cole and Jeremy Buhler

Welcome to the MERCATOR user's manual! MERCATOR is a CUDA/C++ system designed to assist you in writing efficient CUDA applications by automatically generating significant portions of the GPU-side application code. We hope you find it helpful; please feel free to contact the authors with any questions or feedback.

Follow this and additional works at: http://openscholarship.wustl.edu/cse_research



Part of the [Computer Engineering Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Cole, Stephen V. and Buhler, Jeremy, "MERCATOR (Mapping EnumERATOR for CUDA) User's Manual" Report Number: WUCSE-2016-003 (2016). *All Computer Science and Engineering Research*.
http://openscholarship.wustl.edu/cse_research/980

MERCATOR: Mapping EnumERATOR for CUDA ¹

User's Manual ²

Stephen V. Cole (svcole@wustl.edu)
Jeremy D. Buhler (jbuhler@wustl.edu)

Washington University in St. Louis

¹Funded by NSF CNS-1500173 and Exegy, Inc.

²The latest version of this manual is available at <http://sbs.wustl.edu/pubs/MercatorManual.pdf>

Contents

1	Introduction	1
1.1	Intended use case: modular irregular streaming applications	1
1.2	MERCATOR application terminology	2
1.3	System overview and basic workflow	3
2	MERCATOR system documentation	6
2.1	Initial input	6
2.1.1	Module type specification	6
2.1.2	Node specification	12
2.1.3	Edge specification	14
2.1.4	Data parameter type specification	17
2.2	System feedback	18
2.3	Full program	19
2.3.1	Module run function definitions	19
2.3.2	Application data I/O	22
2.3.3	MERCATOR application hooks	24
2.3.4	Restrictions	26
A	Appendix: example application code	26
A.1	Initial input specifications	26
A.2	Module type definitions	28
A.3	Driver code	29

Introduction

Welcome to the MERCATOR user's manual! MERCATOR is a CUDA/C++ system designed to assist you in writing efficient CUDA applications by automatically generating significant portions of the GPU-side application code. We hope you find it helpful; please feel free to contact the authors with any questions or feedback.

1.1 Intended use case: modular irregular streaming applications

MERCATOR facilitates the development of applications that are

- *streaming*: they take many small data items as inputs, process them once, and output or discard them;
- *modular*: they may be described by a (constrained) Data Flow Graph of processing steps (i.e., nodes) connected by directed data flow edges; and
- *irregular*: any node may filter an input item or replace it with new output item(s) to send downstream, and the data flow triggered by any given input is unknown at compile time.

MERCATOR's support for irregularity distinguishes it from systems that process Synchronous Data Flow (SDF) applications, in which the I/O rate(s) of each node are known *a priori*.




Figure 1.1: Example Data Flow Graph of an application with two filtering nodes that each accept integers and output integers. Dashed lines indicate that input items may be dropped by the filtering nodes, which makes the data flow of this application irregular. The `SOURCE` and `SINK` nodes shown represent input generation and output collection in the application.

Figure 1.1 shows the DFG of a simple filtering application.

1.2 MERCATOR application terminology

We use the following terminology to describe the topology of a MERCATOR application:

- A *module type* (or just *module*) is an entity whose type is defined by its function (code).
- A *module instance* is one occurrence of a module in an application's DFG (i.e. one *node*).
- An *edge* represents data flow between two module instances. Note that due to their support for irregular data flow, the exact data movement cardinalities of edges are not necessarily known at compile time.
- Data inputs flow into the application from a single `source node`, and outputs may be emitted to one or more `sink nodes`.

MERCATOR supports any modular streaming application whose topology obeys the following constraints:

1. Each instance may have multiple output edges but only a single input edge.

2. While input items may trigger the generation of multiple output items from a node, a maximum number of possible output items generated for each input item is known at compile time.

We use the following terminology to describe additional aspects of MERCATOR:


- Any code generated by MERCATOR is considered *system* code, while any code supplied by an application programmer (the *user* of the system) is considered *user* code.
- User-supplied *parameter data* may be associated with MERCATOR DFG objects at three granularities: a global set of data, module-type-specific data, and node-specific data. This allows nodes, module types, and the overall application to maintain state during execution.

When implemented in CUDA, the code for each module type is a separate function, with module instances of the same type sharing function code but operating on different input sets with potentially different parameter values. This sharing of code between instances is important in the CUDA context where execution happens in wide SIMD. In fact, one key contribution of MERCATOR’s execution model is to group work by module type rather than by node to take maximal advantage of CUDA’s wide-SIMD execution.


Example DFGs labeled with MERCATOR module types and instances are shown in Figure 1.2.

1.3 System overview and basic workflow

Given the terminology in the previous section, the application support provided by MERCATOR may be stated as follows: the MERCATOR *system* supplies CUDA code for a DFG’s *edges* and provides an interface between the *edges* and *nodes*, enabling a *user* to generate a fully-functional CUDA implementation of a DFG by supplying code for its *module types* only. We refer to this implementation as a ‘MERCATOR application’ or just ‘application,’ even though it is not a full runnable CUDA program without host-side driver code that launches



(a) DFG with two instances, each of a different module type indicated by its label and color.



(b) DFG with three instances of two module types. Subscripts in the type label distinguish instances of the same type.

Figure 1.2: Example DFGs. `SOURCE` and `SINK` nodes have been omitted for clarity. Note that when executing the graph in (b), `MERCATOR` may fire instances A_1 and A_2 simultaneously since they are of the same module type.

the device-side `MERCATOR` application. `MERCATOR` provides wrapper functions that perform these kernel calls and exposes the wrapper functions to the user via a CPU-side (host-side) API.

The user interacts with a `MERCATOR` application via this host-side API, which in addition to the function calls that initialize and run the `MERCATOR` application contains functions for providing inputs and gathering outputs from it. A `MERCATOR` application is therefore deployed in a fully user-defined context.

In order to generate edge code, `MERCATOR` must have information about an application's topology, which the user provides to the system in a well-defined specification format. In order to maintain a consistent interface between nodes and edges, `MERCATOR` provides the user with function headers for each module type in the application.

An overview of the `MERCATOR` workflow is shown in Figure 1.3. Each phase of the workflow is described in the system documentation that follows.

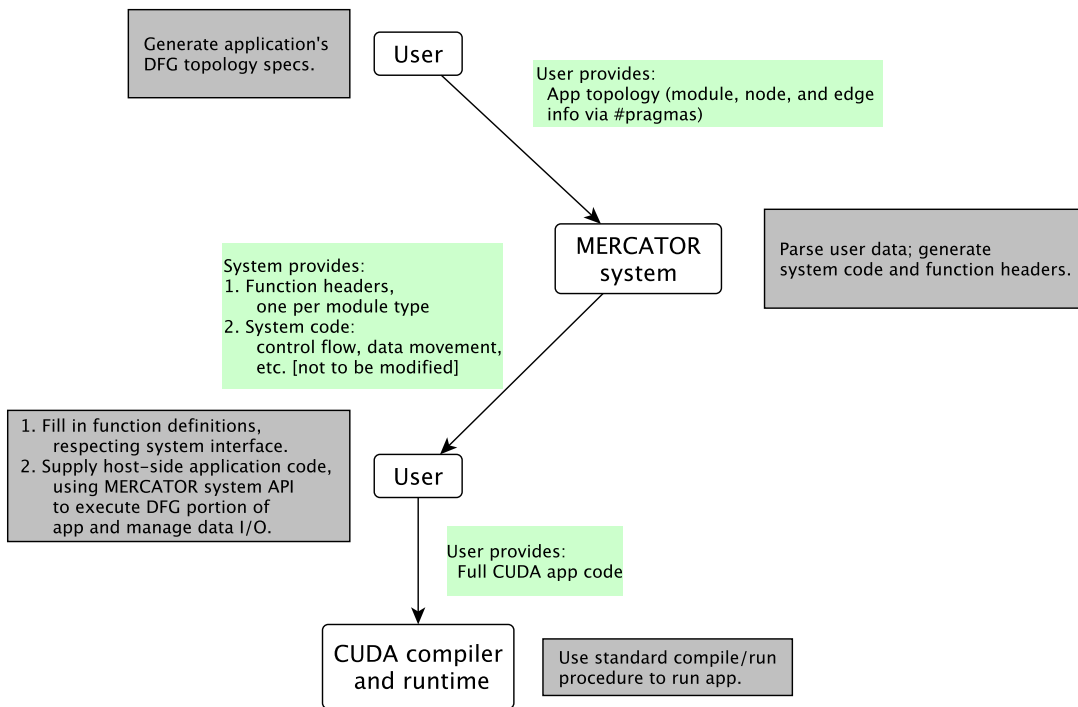


Figure 1.3: Workflow for writing a program incorporating a MERCATOR application.

MERCATOR system documentation

In the following instructions, we highlight the user’s responsibilities and the interface between the user and the MERCATOR system.

2.1 Initial input

We begin with the user’s initial input to the system, which consists of `#pragma` statements specifying the application’s topology, module type requirements, and user-defined parameter data types.

2.1.1 Module type specification

Information about each module type in the application– which may be multiply instantiated in its Data Flow Graph– is provided by the user as a *module specification*. Specifications for user-defined module types must conform to the following EBNF grammar. Note that here and in all following specifications, we use the shorthand `<typename>` to stand in for the set of the names of all CUDA-valid types, including user-defined types.

```

<module-spec> = "#pragma mtr module" <module-name>
                { <parameter-type> } "(" <input-stream> "->"
                <output-stream>{"," <output-stream> } "|"
                <elements> ":" <threads> ")"
<input-stream> = <io-type> <num-inputs>
<output-stream> = <stream-name> "<" <io-type> ">"
                [":" <outputs-per-input>]
<parameter-type> = <typename>
    <num-inputs> = "[" <int> "]"
<outputs-per-input> = <int> | ("?" [<int>])
<module-name> = <string>
<stream-name> = <string>
    <io-type> = <typename>
    <elements> = <int>
    <threads> = <int>

```

Explanation:

The specification consists of three main parts: a description of the data type and range of item cardinalities for each input and output stream, and the work assignment ratio. Notes on specific symbols follow.

<parameter-type> : type name of (optional) parameter data to be associated with this module type.

<num-inputs> : maximum number of input elements processed upon each module firing.

<outputs-per-input> : number of output elements produced per input element. A “?” indicates that the number of outputs per input is *a priori* unknown but bounded above by the integer specified in **<outputs-per-input>**. The absence of a “?” indicates a fixed known data rate. If no integer is provided, the number of outputs per input is assumed to be 1; for example, “?” is equivalent to “?1”.

<module-name> : chosen by user.

<stream-name> : chosen by user. Required to disambiguate output streams.

<elements> : <threads> : ratio of input elements to threads for processing.

<num-instances> : total number of instances of this module type in application's dataflow graph

When more output elements are produced than inputs consumed per firing, outputs are produced on a per-element basis; that is, each input element leads to the production of (up to) a fixed uniform number of output elements. This uniform bound is the integer specified in <outputs-per-input>.

The current version of MERCATOR only supports an <elements>:<threads> ratio with <elements> = 1. In other words, multiple threads may be assigned the same input item during execution of a module, but one thread may not be assigned multiple input items.

Examples:

1. Vector addition module:

```
struct VectorPair {
    Vector a, b;
};
#pragma mtr module addVec(VectorPair[128] ->
    outStream<Vector> | 1:1)
```

Interpretation: The module processes up to 128 input elements per firing, each input element being a struct containing two vectors and each output element being a single vector. The absence of a specified number of outputs per input indicates that each input produces a single output, no matter how many inputs are processed by a given firing. The module assigns one thread to each input element, which implies that the entire element-wise sum of the two vectors is computed by a single thread.

2. Vector addition module, different mapping ratio:

```
#pragma mtr module addVec(VectorPair[128] ->
    outputStream<Vector> | 1:4)
```

Interpretation: Here we assume the same definition of a `VectorPair` as above. This specification is identical to that of the previous example except that the module implementation assigns four threads to process each input vector pair. In this case, the user may arbitrarily divide the work of computing the vector sum among these four threads in the module function definition.

3. Filtering module:

```
#pragma mtr module evalThreshold<StateInfoStruct>(myItem_t[128] ->
    outputStream<myItem_t> : ? | 1:1)
```

Interpretation: The module processes up to 128 input elements per firing, each element being of a user-defined type, and produces either 0 or 1 output elements per input element. The module has associated with it a parameter data object of type `StateInfoStruct`.

4. Cloning module, single output stream:

```
#pragma mtr module cloneItem(myItem_t[128] ->
    outputStream<myItem_t> : 4 | 1:1)
```

Interpretation: The module processes up to 128 input elements per firing, each element being of a user-defined type, and produces 4 output elements per input element.

5. Cloning module, variable output production:

```
#pragma mtr module cloneItem(myItem_t[128] ->
    outputStream<myItem_t> : ?4 | 1:1)
```

Interpretation: Identical to previous example except that each input element produces between 0 and 4 output elements. In this case the MERCATOR infrastructure allocates space for the maximum number of output elements (i.e. 512) for each firing, and the user's module function definition is responsible for respecting this limit.

6. Module that combines cloning and filtering, multiple output streams:

```
#pragma mtr module cloneFilterItem(myItem_t[128] ->
outStream1<myItem_t>, outStream2<myItem_t>, outStream3<myItem_t> :?,
outStream4<myItem_t> :2 | 1:1)
```

Interpretation: The module's output behavior varies by stream: it sends a single element to `outStream1` and `outStream2`, sends 0 or 1 elements to `outStream3`, and sends 2 elements to `outStream4` per input element per firing.

Special types: SOURCE and SINK

In addition to the application module types designed by the user, MERCATOR provides two module types that must be included in an application: `SOURCE` and `SINK`. The single `SOURCE` module type serves as the data entry point into the application, taking input from a user-facing input buffer (to be discussed later) and feeding it to the conceptual root node of the application's DFG. The `SOURCE` module type is therefore always instantiated exactly once, and its input type and (single) output type always match the input type of the DFG's conceptual root node.

The `SINK` module type serves as a data exit point from an application, taking final output from user-defined nodes and passing it to user-facing output buffers (to be discussed later). Since the application's DFG may contain multiple nodes that produce final output, and each output stream may be of a different type, the `SINK` module type is parameterized by this output type and describes a category of types rather than a monolithic type. Each `SINK` type may be multiply instantiated, and its input and output type(s) will all match the final output type it is designed to pass through.

Although the implementation of `SOURCE` and `SINK` module types is part of MERCATOR's system code, the user must include `SOURCE` and `SINK` modules in the input specifications so that the system knows their input and output stream information and so that they may be appropriately included in the application topology. Since their functionality is limited,

however, the amount of information required from their specification is considerably less than from user-defined module types, and their specifications are therefore simpler.

SOURCE and SINK module types must conform to the following EBNF grammar:

```
<source-module-spec> = "#pragma mtr module " ( "SOURCE" | "SINK" )  
                      "<" <io-type> ">"  
                      <io-type> = <typename>
```

Here <io-type> is the input or output type processed by the module. As stated above, there can be only one SOURCE module specified for an application, but there may be multiple SINK modules specified if an application produces outputs of multiple types.

Examples:

1. SOURCE module:

```
#pragma mtr module SOURCE<int>
```

Interpretation: The SOURCE module, and hence the conceptual root of the application's DFG, accepts inputs of type int.

2. SINK modules:

```
#pragma mtr module SINK<int> (2.1)
```

```
#pragma mtr module SINK<MyItem> (2.2)
```

Interpretation: Two types of outputs may be produced by the application: integers and the user-defined type MyItem.

2.1.2 Node specification

Information about each module instance (i.e., node) in the application is provided by the user as a *node specification*. Node specifications must conform to the following EBNF grammar:

```
<node-spec> = "#pragma mtr node " <node-name>
              {<parameter-type>} ":" <module-name>
<node-name> = <string>
<parameter-type> = <typename>
<module-type-name> = <module-name>
```

Explanation:

The specification consists of two parts: a user-chosen name for the module instance, and the name of the module type of which it is an instance.

<node-name> : user-chosen name for this node that will be used to identify it throughout the application.

<parameter-type> : type name of (optional) parameter data to be associated with this node.

<module-type-name> : name of the module type of which this node is an instance; must match the **<module-name>** of some previously-declared module specification.

Example:

In the following example, the relevant module specifications are given, followed by the node specifications describing their instantiations.

```

filterStage1(myItem_t -> fs1outStream<myItem_t> : ? | 1:1)
filterStage2(myItem_t[128] -> fs2outStream<myItem_t> : ? | 1:1)
node1<ThresholdStruct> : filterStage1
node2<ThresholdStruct> : filterStage1
node3<ThresholdStruct> : filterStage2

```

Interpretation: `node1` and `node2` are instances of the module type `filterStage1`, and `node3` is an instance of the module type `filterStage2`. Each node has an associated parameter data object of type `ThresholdStruct`.

SOURCE and SINK nodes

The specification of `SOURCE` and `SINK` nodes must obey the following EBNF grammar:

```

<source-node-spec> = "#pragma mtr node " <node-name> ":SOURCE"
<sink-node-spec>   = "#pragma mtr node " <node-name> ":SINK<"<io-type>">"
<node-name>       = <string>
<io-type>         = <typename>

```

Explanation:

Module instances (nodes) of the `SOURCE` and `SINK` module types are specified in the same way as nodes of user-defined module types, with the keyword ‘`SOURCE`’ being used as the (unique) `SOURCE` module type, and the (possibly not unique) `SINK` module type being specified similarly but with its output type included.

Example:

In the following example, the relevant module specifications are given, followed by the node specifications describing their instantiations.

```
#pragma mtr module SOURCE
#pragma mtr module SINK<int>
#pragma mtr module SINK<MyItem>
...
#pragma mtr node sourceNode : SOURCE
#pragma mtr node resultNode1 : SINK<MyItem>
#pragma mtr node resultNode2 : SINK<MyItem>
#pragma mtr node resultNode3 : SINK<MyItem>
#pragma mtr node resultNode4 : SINK<int>
```

Interpretation: `sourceNode` is the single source node, the nodes `resultNode1-3` produce application results of type `MyItem`, and the node `resultNode4` produces application results of type `int`. The specifications of all user-defined module types and nodes has been omitted in this example.

2.1.3 Edge specification

An *edge specification* describes a dataflow relationship between two nodes; i.e., it represents an edge in the application's Data Flow Graph. Edge specifications must conform to the following EBNF grammar:

```

<edge-spec> = <us-node> ":@" <output-stream>
              "->" <ds-node>
<us-node>   = <node-name>
<ds-node>   = <node-name>
<output-stream> = <stream-name>

```

Explanation:

<us-node> : upstream node in connection; must match the **<node-name>** in a previously-declared node specification.

<ds-node> : downstream node in connection; must match the **<node-name>** in a previously-declared node specification.

<output-stream> : stream of **<us-node>** to be connected to input of **<ds-node>**; must match a **<stream-name>** in the specification of the module type of which **<us-node>** is an instance.

As a result of our single-input assumption for modules, each module may appear as the **<ds-module>** in only a single edge specification.

Note that since nodes of **SOURCE** and **SINK** module types are named in the same way as those of user-defined module types, edge specifications containing them need not have a distinct form. However, a user is expected to include the following **SOURCE** and **SINK** edge specifications in every application:

1. A single edge specification that explicitly connects the **SOURCE** node to the conceptual DFG root node.
2. An edge specification that explicitly connects each user-defined node that produces final application output to its associated **SINK** node.

Examples

In the following examples, the relevant module and node specifications are given, followed by the edge specification(s) describing the dataflow between the nodes.

1. Connection between two filtering modules:

```
filterStage1(myItem_t -> fs1outStream<myItem_t> : ? | 1:1)
filterStage2(myItem_t[128] -> fs2outStream<myItem_t> : ? | 1:1)
node1 : filterStage1
node2 : filterStage2
node1::fs1outStream -> node2
```

Interpretation: `node1` and `node2` are instances of module types `filterStage1` and `filterStage2` respectively. The elements output by `node1` on its `fs1outStream` stream are placed on the input worklist of `node2`. *Note:* It is assumed that `filterStage1` and `filterStage2` instantiate different filter code, rather than the same filter code parameterized differently (e.g., a simple threshold filter with different threshold values). An example of such a topology exists in the Viola-Jones filter cascade, in which different filters assay different features of a candidate face image.

2. Connection requiring named streams for disambiguation:

```
splitItems(myItem_t[128] -> siOutputStream1<int>:2,siOutputStream2<int>:3 | 1:1)
filterItems(int[128] -> fiOutputStream<int>:? | 1:1)
node1 : splitItems
node2 : filterItems
node1::siOutputStream1 -> node2
```

Interpretation: The elements in the `siOutputStream1` output stream of `node1` are placed on the input worklist of `node2`. *Note:* In this case naming module output streams enables disambiguation; both `siOutputStream1` and `siOutputStream2` have the same type, so inferring which stream to use to connect `node1` and `node2` would not be possible without a stream ID.

2.1.4 Data parameter type specification

As stated in the description of a MERCATOR application, user-defined parameter data may be associated with module types, nodes, and the application. The parameter data types of module types and nodes are included in their specification, but since there is no specification for the application, its parameter data type (if present) must be specified separately. This specification must conform to the following EBNF grammar:

```
<app-param-data-spec> = "#pragma mtr appdata " <parameter-name>  
    <parameter-name> = <typename>
```

Here `<parameter-name>` is the name of the application's parameter data type.

Example: LinearPipe application

The set of module type, node, and edge specifications for an example MERCATOR application are shown in the Appendix in Listing 1. This application, which we refer to as “LinearPipe,” implements the linear pipeline of filter nodes shown in the examples in Section 1 of this manual, with the additional property that all nodes in the pipeline may produce final output. This allows the application to expose elements that it filters out as well as those that pass through the filters at each stage. A diagram of the application's topology is shown in Figure 2.1. This application will serve as a running example for other pieces of user-generated code in the sections that follow.

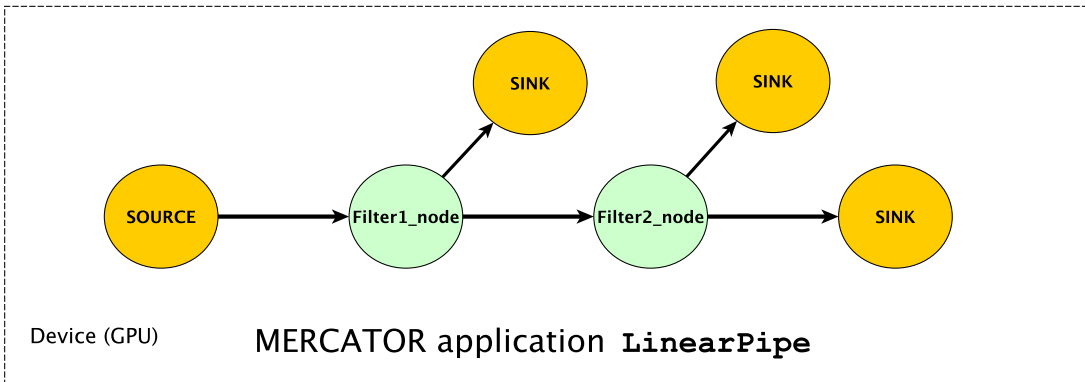


Figure 2.1: Diagram showing the DFG of the MERCATOR filtering application named `LinearPipe` described in the text. Note the existence of multiple `SINK` nodes.

2.2 System feedback

Based on the user’s initial input, the MERCATOR system will generate and return CUDA code of two types: system infrastructure code and module function headers.

The system infrastructure code is not intended to be modified by the user, but along with the user’s function definitions, it constitutes an essential part of the MERCATOR application to be integrated into the user’s final program code.

Module function headers

MERCATOR provides a set of module function headers written in CUDA that implicitly contain the user’s interface to the MERCATOR infrastructure code. One header is produced for each module type, and the function bodies are to be filled in by the user. The details of

writing function definitions are presented in the next section. Module function headers have the following format:

```
void run (  
    <input-type> inputItem,  
    <instance-type>* node)
```

Explanation:

`<input-type>` : type of elements in module’s input stream; determined by module specification previously given by the user.

`<instance-type>` : type of the current node being executed.

`node` : the current node being executed; used for API function calls in module definition.

Since MERCATOR’s scheduling strategy involves firing multiple instances of the same module simultaneously, the `<instance-type>` as well as the `inputItem` may vary across threads executing the same function call.

2.3 Full program

2.3.1 Module run function definitions

If there are more threads present on the GPU device than there are items to be operated on, MERCATOR ensures that the “extra” threads do not execute the `run()` function call. This ensures that each thread that executes the function will have a valid input item passed in through the `inputItem` parameter; however, it also implies that to avoid undefined behavior, module definitions must be written with no CUDA synchronization calls such as `__syncthreads()`.

The remaining paragraphs in this section describe MERCATOR API functions accessible to the user inside `run()` functions.

Accessing user-defined parameter data

Within the `run()` function, users may access parameter data using the following API function calls:

- Application level: `get_myAppData()`
- Module type level: `this->get_myModuleData()`
- Module type level: `this->get_myNodeData(node)`

Writing output

Output may be written from within the `run()` function using the following API call:

```
node->push(item, Out::streamName)
```

where `item` is the output item to be written, and `streamName` is the name given to the desired output stream by the user in the initial input specification. Here we reiterate that although many threads may be writing to the same output stream, no synchronization code should be included in the `run()` function; MERCATOR's system code ensures the integrity of the output streams.

Multiple threads per input element

When multiple threads are assigned to each input element (as dictated by the `<elements>:<threads>` ratio in the module type input specification), the following API call may be used to access a thread's offset relative to all threads assigned to the same input element:

```
myOffset(threadIdx.x)
```

This call will return an `int`, which can be used to process the appropriate partition of a shared input. For example, if many threads are assigned the same input string of which they are intended to process consecutive substrings, each thread may find its starting point relative to the input string with a call to `myOffset` .

Multiple output elements per input element

When multiple output elements may be written to an output stream for each input consumed, either by a single thread or in the case of multiple threads per input element, an output slot within the output stream must be specified for each `push` operation. This is done with an (optional) third parameter to the function, so that the function call will be of the following form:

```
node->push(item, Out::streamName, slotIdx)
```

where `slotIdx` is the desired slot within the output stream relative to all possible outputs generated by the current input.

For example, if each thread writes four outputs to the same stream during each `run()` function execution, those four writes may take the following form:

```
node->push(item, Out::streamName, 0) // the '0' parameter is optional
node->push(item, Out::streamName, 1)
node->push(item, Out::streamName, 2)
node->push(item, Out::streamName, 3)
```

In the case of multiple threads being assigned to the same input element, the `myOffset` function may be used to specify an output slot. For example, if multiple threads are assigned to the same input element and each thread writes a single output during each `run()` function execution, that write may take the following form:

```
node->push(item, Out::streamName, myOffset(threadIdx.x))
```


One possible set of function definitions for the headers produced for the `LinearPipe` application from the input specifications shown in Listing 1 is shown in Listing 2.

2.3.2 Application data I/O

The MERCATOR infrastructure code and the user’s module definitions may be incorporated into a host-side driver file via a single `#include` of the module header file. The resulting CUDA code may be submitted to any CUDA compiler (e.g., `nvcc`) for compilation and subsequently executed using any CUDA runtime.

We now turn to the portions of the MERCATOR user interface that reside on the CPU. We begin with the data interface between user code and the MERCATOR application. MERCATOR provides special input and output buffer types to manage host-device data transfers, and it is these buffers with which user code interacts to move data into and out of MERCATOR. These buffer types are parameterized by the type of data they hold.

Input buffers

Input buffers are created with a capacity and, optionally, with an initial data set; the API signature for the input buffer constructor is the following:

```
InputBuffer<io-type>(int capacity, io-type* inputElements)
```

where `inputElements` is an optional parameter.

This data set must have been allocated using CUDA’s “Managed Memory”— i.e., with a call to `cudaMallocManaged()` rather than the usual `cudaMalloc()`— so that the data will be accessible on both host and device. If the initial data set is not provided, elements must be added to the buffer one at a time using the API member call `add(nextElement)`.

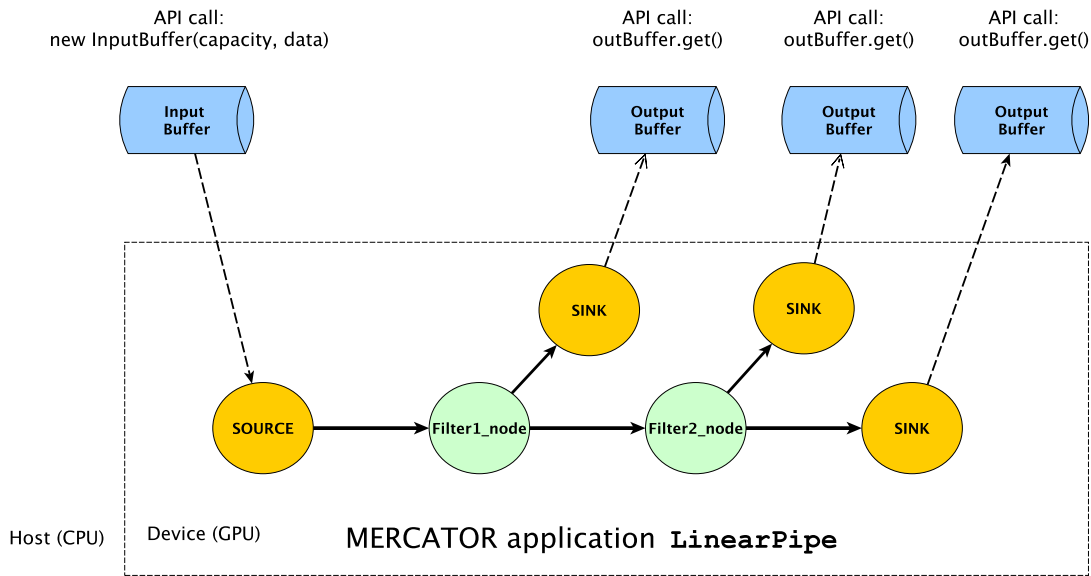


Figure 2.2: Diagram showing the LinearPipe DFG as well as its data interface with the host via input and output buffers.

Output buffers

Output buffers are created with a capacity; the API for the output buffer constructor is the following:

```
OutputBuffer<io-type>(int capacity)
```

Data elements may be retrieved from an output buffer one at a time with the API member call `get()`, which pops an item off the output buffer and returns it.

A diagram of the LinearPipe application along with its associated input and output buffers is shown in Figure 2.2.

2.3.3 MERCATOR application hooks

To set up and run the MERCATOR application from the user's host-side program, a host-side class called `App` in the `Mercator` namespace is provided, and API member calls on an object of this class perform the necessary operations.

Setup and buffer/node association

In order to connect the buffers to the appropriate data entry/exit points of the MERCATOR application— i.e., the appropriate `SOURCE` and `SINK` nodes— the user must specify which buffer gets associated with which node. This is done with an API call that takes the relevant node names and buffer objects as parameters. For the (single) input buffer, only a single name and a single buffer object pointer must be provided; in the case of the (possibly multiple) output buffers, a C++ STL vector of names and a vector of buffer object pointers is required. The format of the API member call to associate buffers with `SOURCE/SINK` nodes, then, is the following:

```
setup(inBuffer, sourceNodeName, outBufferVec, sinkNodeNameVec)
```

where `inBuffer` is a pointer to an `InputBuffer` object, `sourceNodeName` is the name (as a `string`) of the source node, and `outBufferVec` and `sinkNodeNameVec` are parallel vectors containing at each index an output buffer pointer and the name of a sink node that should be paired.

In addition to creating the buffer/node associations, this function initializes all necessary data structures on the GPU device. It is a blocking function call, so that when it returns, all GPU kernel calls have been completed and the MERCATOR application is primed for execution.

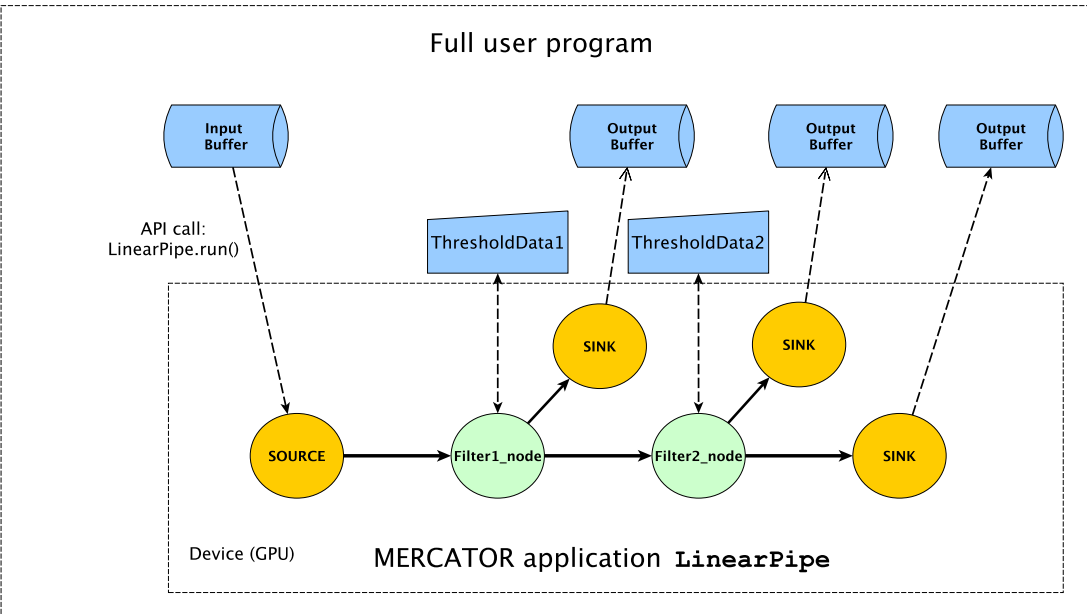


Figure 2.3: Diagram showing the `LinearPipe` DFG and its interface with host-side code and data, including I/O buffers and user-defined parameter data accessible from within the modules' `run()` functions. Two of the program's output buffers have been omitted for clarity.

Running the MERCATOR application

Once input and output buffers have been created, the input buffer has been initialized with data, and buffers have been associated with nodes using the `setup()` function call on the `Mercator::App` object, the MERCATOR application may be run from a user program with the API member call `run()`. This blocking function call encapsulates the entire execution of the MERCATOR application, so that when it returns, the input buffer will be empty and the output buffers will contain all program outputs assuming an error-free execution on the GPU device.

Once the application has been run, data may be gathered from the output buffers using the API calls described above.

The relevant parts of the user program for the `LinearPipe` application are shown in Listing 3, and a diagram showing the full program layout is shown in Figure 2.3.

2.3.4 Restrictions

The following restrictions on incorporating MERCATOR applications into user programs currently exist, though we plan to eliminate them in future versions of the system:

- Single MERCATOR app: Only a single MERCATOR application may be invoked from a single user program.
- Single-batch execution: All input and output buffers must fit into the GPU card's global memory simultaneously; multiple rounds of execution and I/O copying are not supported.

A Appendix: example application code

We show here all code required to create a simple filtering application in MERCATOR and access it from a user program.

A.1 Initial input specifications

Listing 1 shows the set of input specifications for the application.

Listing 1: Input specification for the "LinearPipe" MERCATOR application. The application contains 2 module types in addition to the SOURCE type and a single SINK type, and consists of a linear pipeline with shunts to SINK nodes off of each node in the pipeline.

```
/** App-level data type. */  
#pragma mtr appdata FilterAppData  
  
/** Module (i.e., module type) specs. */  
  
// Filter 1  
// Note the user-defined parameter data of type "Filter1State"  
#pragma mtr module Filter1<Filter1State> (
```

```

        int [10] -> accept<int>: ?4, reject<int>: ?4 | 1 : 1 )

// Filter 2
#pragma mtr module Filter2 (
    int [10] -> accept<int>: ?4, reject<int>: ?4 | 1 : 4 )

// SOURCE Module
#pragma mtr module SOURCE<int>

// SINK Module
#pragma mtr module SINK<int>

/**** Node (i.e., module instance) specs. ****/
#pragma mtr node sourceNode : SOURCE
#pragma mtr node Filter1_node : Filter1

// Note the user-defined parameter data of type "Filter2ThresholdStruct"
#pragma mtr node Filter2_node<Filter2ThresholdStruct> : Filter2
#pragma mtr node sinkNodeReject1 : SINK<int>
#pragma mtr node sinkNodeReject2 : SINK<int>
#pragma mtr node sinkNodeAccept : SINK<int>

/**** Edge specs. ****/

// SOURCE -> Filter1
#pragma mtr edge sourceNode::outStream -> Filter1_node

// Filter1 -> Filter2
#pragma mtr edge Filter1_node::accept -> Filter2_node

// Filter1 -> Rejects

```

```

#pragma mtr edge Filter1_node::reject -> sinkNodeReject1

// Filter2 -> SINK
#pragma mtr edge Filter2_node::accept -> sinkNodeAccept

// Filter2 -> Rejects
#pragma mtr edge Filter2_node::reject -> sinkNodeReject2

```

A.2 Module type definitions

Example module type definitions are shown in Listing 2.

Listing 2: Two sample module function bodies for the module types defined in Listing 1. `Filter1` conditionally writes four values either to the “accept” output stream or the “reject” output stream; note its use of `Filter1`’s parameter data, the global application parameter data, and the use of slot indices in the last parameter of the `push()` function. `Filter2` conditionally writes a single element to the `accept` or `reject` stream. Note its use of the `myOffset()` function to assign output to the correct slot, since more than one thread is given the same input.

```

///// module Filter1<UserDataExt>(int [10] -> accept<int> : ?4, reject<int>: ?
__device__
void TestApp::DerivedModule_Filter1::run(int inputItem,
    Mercator::ModuleInstance<int, DerivedModule_Filter1>* node)
{
    int stateNum = get_myModuleData()->stateNum;
    if(stateNum < 6) {
        int scaler = get_myAppData()->scale;
        int outputItem = inputItem * scaler;
        node->push(outputItem, Out::accept, 0); // 0-param is optional
        node->push(outputItem + 1, Out::accept, 1);
        node->push(outputItem + 2, Out::accept, 2);
        node->push(outputItem + 3, Out::accept, 3);
    }
}

```

```

    else {
        node->push(inputItem, Out::reject, 0);
        node->push(inputItem + 1, Out::reject, 1);
        node->push(inputItem + 2, Out::reject, 2);
        node->push(inputItem + 3, Out::reject, 3);
    }
}

///// module Filter2(int[10] -> accept<int>: ?4, reject<int>: ?4 | 1 : 4)
__device__
void TestApp::DerivedModule_Filter2::run(int inputItem,
Mercator::ModuleInstance<int, DerivedModule_Filter2>* node)
{
    if(inputItem < get_myNodeData(node)->thresh)
        node->push(inputItem, Out::accept, myOffset(threadIdx.x));
    else
        node->push(inputItem, Out::reject, myOffset(threadIdx.x));
}

```

A.3 Driver code

Driver code for a user program that invokes the `LinearPipe` application is shown in Listing 3.

Listing 3: User code that creates MERCATOR I/O buffers, fills the input buffers with consecutive integers, initializes and invokes the "LinearPipe" application, and reads the results from the output buffers.

```

/**
 * User program that acts as test harness
 * for MERCATOR filtering application "LinearPipe."
 *
 */

#include <iostream>

```



```

#include "userInput/TestApp.cuh" // headers for module run() fcns
#include "system_headers.cuh" // other MERCATOR system headers

int main()
{
    ////////// set up input buffer
    const int IN_BUFFER_CAPACITY = 128;

    int* inBufferData;
    cudaMallocManaged(&inBufferData,
        IN_BUFFER_CAPACITY * sizeof(int));
    Mercator::IO::InputBuffer<int>* inBuffer =
        new Mercator::IO::InputBuffer<int>(inBufferData,
            IN_BUFFER_CAPACITY);

    std::string sourceNodeName = "sourceNode";

    ////////// set up output buffers
    std::vector<Mercator::IO::BufferBase*> outBufferVec;
    std::vector<std::string> sinkNodeNameVec;

    // output buffers accommodates 4 outputs/input
    const int OUT_BUFFER_CAPACITY1 = 4 * IN_BUFFER_CAPACITY;
    const int OUT_BUFFER_CAPACITY2 = 4 * OUT_BUFFER_CAPACITY1;
    Mercator::IO::OutputBuffer<int>* outBuffer1 =
        new Mercator::IO::OutputBuffer<int>(OUT_BUFFER_CAPACITY1);
    outBufferVec.push_back(outBuffer1);
    std::string sinkNodeName1 = "sinkNodeReject1";
    sinkNodeNameVec.push_back(sinkNodeName1);

    Mercator::IO::OutputBuffer<int>* outBuffer2 =
        new Mercator::IO::OutputBuffer<int>(OUT_BUFFER_CAPACITY2);

```

```

outBufferVec.push_back(outBuffer2);
std::string sinkNodeName2 = "sinkNodeReject2";
sinkNodeNameVec.push_back(sinkNodeName2);

Mercator::IO::OutputBuffer<int>* outBuffer3 =
    new Mercator::IO::OutputBuffer<int>(OUT_BUFFER_CAPACITY2);
outBufferVec.push_back(outBuffer3);
std::string sinkNodeName3 = "sinkNodeAccept";
sinkNodeNameVec.push_back(sinkNodeName3);

//////// fill input buffer
for(int i=0; i < IN_BUFFER_CAPACITY; ++i)
    inBuffer->add(i);

//////// create app object
Mercator::App* LinearPipe = new Mercator::App();

//////// set node-, module-, app-level user data
FilterAppData* fullAppData = new FilterAppData();
LinearPipe->set_fullAppUserData(fullAppData);

Filter1State* filter1data = new Filter1State();
LinearPipe->set_moduleTypeUserData("Filter1", filter1data);

Filter2ThresholdStruct* filter2nodeData =
    new Filter2ThresholdStruct(4);
LinearPipe->set_moduleInstUserData("Filter2_node", filter2nodeData)

//////// initialize app object
LinearPipe->setup(inBuffer, sourceNodeName,
                outBufferVec, sinkNodeNameVec);

```

```
//////// run main function
LinearPipe->run();

//////// print contents of output buffers
outBuffer1->print();
outBuffer2->print();
outBuffer3->print();

cudaFree(inBufferData);
return 0;
}
```