





# **D2.2.A: Requirements Analysis**

Grant Agreement:	287767
Project Acronym:	CARP
Project Name:	Correct and Efficient Accelerator Programming
Instrument:	Small or medium scale focused research project (STREP)
Thematic Priority:	Alternative Paths to Components and Systems
Start Date:	1 December 2011
Duration:	36 months
Document Type <sup>1</sup> :	D (Deliverable)
Document Distribution <sup>2</sup> :	CO (Confidential)
Document Code <sup>3</sup> :	CARP-ARM-RP-001
Version:	v1.0
Editor (Partner):	A. Lokhmotov (ARM)
Contributors:	ARM, REAL, RIGHT
Workpackage(s):	WP2
Reviewer(s):	RWTHA
Due Date:	30 September 2012
Submission Date:	10 November 2012
Number of Pages:	258

Copyright © 2014 by the CARP Consortium.

 $<sup>^{1}</sup>$ MD = management document; TR = technical report; D = deliverable; P = published paper; CD = communication/dissemination.  $^{2}$ PU = Public; PP = Restricted to other programme participants (including the Commission Services); RE = Restricted to a group specified by the consortium (including the Commission Services); CO = Confidential, only for members of the consortium (including the Commission Services).

<sup>&</sup>lt;sup>3</sup>This code is constructed as described in the Project Handbook.



# **D2.2.A: Requirements Analysis**

A. Kravets<sup>1</sup>, G. Kouveli<sup>1</sup>, A. Lokhmotov<sup>1</sup>, E. Hajiyev<sup>2</sup>, L. Marák<sup>2</sup>, T. Virolainen <sup>3</sup>

<sup>1</sup>ARM, <sup>2</sup>REAL, <sup>3</sup>RIGHT

Date	Version	Author	Modification	
2012-09-06	0.1	A. Kravets (ARM),	First version	
		G. Kouveli (ARM),		
		A. Lokhmotov (ARM)		
2012-09-21	0.2	E. Hajiyev (REAL)	Changes from REAL	
2012-10-03	0.3	T. Virolainen (RIGHT)	Changes from RIGHT	
2012-10-08	0.4	E. Hajiyev (REAL)	Updates following comments	
			from ARM	
2012-10-31	0.5	A. Kravets (ARM),	Updates following comments	
		G. Kouveli (ARM)	from ICL	
2012-11-08	0.6	T. Virolainen (RIGHT)	Updates following comments	
			from ARM	
2012-11-09	1.0	A. Kravets (ARM),	First version sent to the European	
		G. Kouveli (ARM),	Commission	
		A. Lokhmotov (ARM)		

## **REVISION HISTORY**

# APPROVALS

Role	Name	Partner	Date
Workpackage Leader	A. Lokhmotov	ARM	
Project Manager	A.F. Donaldson	ICL	



# Contents

1	Exec	cutive Su	ımmary	5
2	Bend	enchmark Analysis		
	2.1	Definiti	ions	6
		2.1.1	Statements	6
		2.1.2	Work-item identifiers	6
		2.1.3	Data structures	7
		2.1.4	Iteration mapping	8
		2.1.5	Memory access mapping	8
		2.1.6	Dependences	9
		2.1.7	Code abstractions	10
	2.2	SHOC	Level 1	12
		2.2.1	Triad	12
		2.2.2	Stencil2D	17
		2.2.3	Reduction	27
		2.2.4	Scan	32
		2.2.5	Sort	43
		2.2.6	MD	55
		2.2.7	SGEMM	60
		2.2.8	Sparse matrix-vector multiplication	70
	2.3	Rodinia	1	81
		2.3.1	Back Propagation	81
		2.3.2	Breadth-First Search	93
		2.3.3	CFD Solver	101
		2.3.4	Gaussian Elimination	112
		2.3.5	Heart Wall	120
		2.3.6	Hot Spot	123
		2.3.7	K-means	130
		2.3.8	LavaMD	140
		2.3.9	Leukocyte Tracking	148
		2.3.10	LU Decomposition	158
		2.3.11	k-Nearest Neighbors	170
		2.3.12	Needleman-Wunsch	176
		2.3.13	Particle Filter	188
		2.3.14	PathFinder	203
		2.3.15	Speckle Reducing Anisotropic Diffusion	209
		2.3.16	Stream Cluster	225
	2.4	Parboil	· · · · · · · · · · · · · · · · · · ·	230
		2.4.1	Breadth-First Search	230
		2.4.2	SGEMM	230
		2.4.3	Sparse Matrix-Vector Multiplication	230
		2.4.4	Stencil	230
		2.4.5	Other benchmarks	231



2.5	RealEyes		
	2.5.1	Arithmetic Operations	
	2.5.2	GEMM – General Matrix Multiply	
	2.5.3	Image Statistics	
	2.5.4	Image Normalization	
	2.5.5	Image (Data) Repacking	
2.6	Basem	narkCL	
	2.6.1	Histogram equalization	
	2.6.2	SPH fluid	



# **1** Executive Summary

Software for accelerated systems is predominantly written using low-level APIs, such as OpenCL and CUDA. Often developers compromise on performance portability (by tuning only to a specific architecture *e.g.* NVIDIA's Fermi) and correctness portability (by relying on certain architectural features *e.g.* warp-based execution). We aim to address the performance portability problem by designing PENCIL, a platform-neutral compute intermediate language. PENCIL will be used as both a target language for compilers from domain-specific languages (DSLs) and an efficiency-layer language for expert programmers, in particular, library implementers. Advanced code generation techniques based on the polyhedral model will translate PENCIL code into efficient platform-specific code.

To gather design requirements for PENCIL, we have analyzed two prominent GPGPU benchmark suites, SHOC [8]( $\S2.2$ ) and Rodinia [5, 6]( $\S2.3$ ), as well as validation cases from Realeyes ( $\S2.5$ ). We focused our analysis on two levels: a high-level description of a benchmark (a mathematical formulation of the problem, abstract data structures, iteration domains, dependences) and low-level implementation details (concrete data structures, partitioned iteration domains, memory access patterns, target-specific optimisations).

The high-level description may only specify the desired result, not the algorithm, so in general is unsuitable for representing accelerator workloads. On the other hand, the low-level description may be too implementation-specific. Therefore, we aim to design PENCIL to be somewhere in between these two levels.

Our analysis is the largest effort to date evaluating access/execute specifications [12]. The execute specifications describe the iteration spaces, dependences and partitioning, while the access specifications describe the memory locations accessed on each iteration. Given the emphasis on iteration space partitioning and data movement in accelerator programming, we believe incorporating such specifications is crucial to success of an accelerator programming model.

The detailed analysis also led to discovering several bugs in both Rodinia and SHOC.





# 2 Benchmark Analysis

We analyze several benchmark suites.

For each benchmark we describe the algorithm and data structures at both high level (mathematical formulation) and low level (OpenCL implementation). See  $\S2.2.1$  for an introductory example.

# 2.1 Definitions

### 2.1.1 Statements

**Definition 1** A statement is a sequence of operations under linear control flow that can be defined in terms of its side-effects.

According to this definition, a maximal linear sequence of operations (*basic block*) can be considered a single statement and associated with a single iteration of an iteration space.

## 2.1.2 Work-item identifiers

For each dimension i = 0, 1, 2:

•  $\Lambda_i$  denotes the local work size, *i.e.* the number of unique local-ids in dimension *i*:

$$\Lambda_i = \text{get\_local\_size}(i) \tag{2.1}$$

•  $\lambda_i$  denotes the local-id in dimension  $i, 0 \leq \lambda_i < \Lambda_i$ :

$$\lambda_i = \text{get_local_id}(i)$$
 (2.2)

•  $N_i$  denotes the number of work-groups in dimension *i*:

$$N_i = \text{get_num_groups}(i)$$
 (2.3)

•  $v_i$  denotes the work-group-id in dimension  $i, 0 \le v_i < N_i$ :

$$v_i = \text{get}_{group}(i)$$
 (2.4)

•  $\Gamma_i$  denotes the global work size, *i.e.* the number of unique global-ids in dimension *i*:

$$\Gamma_i = \text{get_global_size}(i)$$
 (2.5)

$$\Gamma_i = N_i \Lambda_i \tag{2.6}$$

•  $\Delta_i$  denotes the global-id offset in dimension *i*:

$$\Delta_i = \texttt{get\_global\_offset}(i) \tag{2.7}$$

•  $\gamma_i$  denotes the global-id in dimension i,  $\Delta_i \leq \gamma_i < \Gamma_i + \Delta_i$ :

$$\gamma_i = \text{get_global_id}(i)$$
 (2.8)

 $\gamma_i = \Lambda_i \nu_i + \lambda_i + \Delta_i \tag{2.9}$ 





# **2.1.3 Data structures**

#### Abstract data structures

**Range** A monotonic sequence  $\{R_i\}$  of integer numbers denoted as start : end : stride, where

 $R_i = \text{start} + i \times \text{stride}, 0 \le i < |R| = 1 + (\text{end} - \text{start})/\text{stride}$ 

If start  $\leq$  end, start : end is a shorthand for start : end : 1; otherwise, start : end is a shorthand for start : end : -1.

**Vector** An ordered collection of numbers. The *i*<sup>th</sup> element of vector  $\vec{A}$  is denoted as  $A_i$ ,  $0 \le i < |\vec{A}|$ , where  $|\vec{A}|$  is the number of elements in vector  $\vec{A}$  (the *size* of vector  $\vec{A}$ ).

Slice For range R = start: end: stride and vector  $\vec{A}$ , slice  $\vec{A}_R$  denotes vector  $\vec{A'}$  of |R| elements:

$$A_i' = A_{R_i}, 0 \le i < |\mathbf{R}|$$

A[] is a shorthand for A[0: N-1], where N is the size of  $\vec{A}$ .

Matrix A two-dimensional table of numbers:

$$M:\begin{pmatrix} M_{0,0} & M_{0,1} & \cdots & M_{0,n-1} \\ M_{1,0} & M_{1,1} & \cdots & M_{1,n-1} \\ \vdots & \vdots & \ddots & \vdots \\ M_{m-1,0} & M_{m-1,1} & \cdots & M_{m-1,n-1} \end{pmatrix}$$

where *m* is the number of rows and *n* is the number of columns in the table. The element in row *i* and column *j* is denoted as  $M_{i,j}$ .

**Graph** For a graph G = (V, E) the following notation is used:

- *V* denotes the set of vertices.
- *E* denotes the set of edges.
- $E^{\nu}$  denotes the set of edges connected to vertex  $\nu \in V$ .
- $e_h$  denotes the source vertex (*head*) of edge  $e \in E$ .
- $e_t$  denotes the sink vertex (*tail*) of edge  $e \in E$ .
- $|v| = |E^v|$  denotes the degree of vertex  $v \in V$ .

#### **Concrete data structures**

*N*-dimensional array A systematic arrangement of data objects, each identified by *N*-element tuple of integer numbers (index). A[i0][i1]..[iN] denotes the element of array A with index (i0, i1, ..., iN).

Buffer A one-dimensional array. The only device data structure apart from Image.

```
CARP-ARM-RP-001-v1.0
```





### **Two-dimensional image (texture)**

# 2.1.4 Iteration mapping

- *I* denotes the iteration space in high-level descriptions.
- *P* denotes the iteration space in low-level descriptions.
- $f: I \rightarrow P$  denotes the mapping from *I* to *P*.
- $f^{-1}: P \to I$  denotes the inverse mapping from P to I:  $f^{-1} \circ f = id_I$ .

# 2.1.5 Memory access mapping

In high-level descriptions:

- $M_r(i)$  denotes the set of uniform memory locations read on iteration  $i \in I$ .
- $M_w(i)$  denotes the set of uniform memory locations written on iteration  $i \in I$ .
- $M_{mr}(i)$  denotes the set of uniform memory locations possibly read on iteration  $i \in I$ .
- $M_{mw}(i)$  denotes the set of uniform memory locations possibly written on iteration  $i \in I$ .
- $F_r(i) = M_r(i) \cup M_{mr}(i)$
- $F_w(i) = M_w(i) \cup M_{mw}(i)$
- $F_{rw}(i) = F_r(i) \cup F_w(i)$

In low-level descriptions:

- $H_r(p)$  denotes the set of host memory locations read on iteration  $p \in P$ .
- $H_w(p)$  denotes the set of host memory locations written on iteration  $p \in P$ .
- $H_{mr}(p)$  denotes the set of host memory locations possibly read on iteration  $p \in P$ .
- $H_{mw}(p)$  denotes the set of host memory locations possibly written on iteration  $p \in P$ .
- $G_r(p)$  denotes the set of global memory locations read on iteration  $p \in P$ .
- $G_w(p)$  denotes the set of global memory locations written on iteration  $p \in P$ .
- $G_{mr}(p)$  denotes the set of global memory locations possibly read on iteration  $p \in P$ .
- $G_{mw}(p)$  denotes the set of global memory locations possibly written on iteration  $p \in P$ .
- $L_r(p)$  denotes the set of local memory locations read on iteration  $p \in P$ .
- $L_w(p)$  denotes the set of local memory locations written on iteration  $p \in P$ .
- $L_{mr}(p)$  denotes the set of local memory locations possibly read on iteration  $p \in P$ .
- $L_{mw}(p)$  denotes the set of local memory locations possibly written on iteration  $p \in P$ .



- $C_r(p)$  denotes the set of constant memory locations read on iteration  $p \in P$ .
- $C_{mr}(p)$  denotes the set of constant memory locations possibly read on iteration  $p \in P$ .
- $P_r(p)$  denotes the set of private memory locations read on iteration  $p \in P$ .
- $P_w(p)$  denotes the set of private memory locations written on iteration  $p \in P$ .
- $P_{mr}(p)$  denotes the set of private memory locations possibly read on iteration  $p \in P$ .
- $P_{mw}(p)$  denotes the set of private memory locations possibly written on iteration  $p \in P$ .

# 2.1.6 Dependences

The notation above can be used for a single statement S. For example,  $M_r(S)$  denotes the set of uniform memory locations read by statement S.

# **Definition 2 (Memory-based data dependence)**

We say that there is a memory-based data dependence from statement  $S_1$  to statement  $S_2$  (equivalently,  $S_2$  depends on  $S_1$ ) and write  $S_1 \rightarrow S_2$ , if and only if:

- 1. statement  $S_1$  is executed before statement  $S_2$ ,
- 2.  $S_1$  and  $S_2$  access the same memory location and at least one of them writes into it:  $((F_r(S_1) \cap F_w(S_2)) \cup (F_w(S_1) \cap F_r(S_2)) \cup (F_w(S_1) \cap F_w(S_2))) \neq \emptyset.$

Dependences are traditionally classified according to the relative order of memory accesses.

# Definition 3 (Classes of memory-based dependences)

**True (or flow, or read-after-write) dependence:** *Statement*  $S_1$  *writes into a memory location from which statement*  $S_2$  *later reads (denoted as*  $S_1 \xrightarrow{t} S_2$ )*:* 

S1: A = ... S2: ... = A

**Anti (or write-after-read) dependence:** Statement  $S_1$  reads from a memory location into which statement  $S_2$  later writes (denoted as  $S_1 \xrightarrow{a} S_2$ ):

S1:  $\ldots$  = A S2: A =  $\ldots$ 

**Output (or write-after-write) dependence:** Both statements  $S_1$  and  $S_2$  write into the same memory location (denoted as  $S_1 \xrightarrow{o} S_2$ ):

S1: A = ... S2: A = ...

If there exist a dependence between  $S_1$  on iteration  $(i_0^1, \ldots, i_{k_1}^1) \in I_1$  and  $S_2$  on iteration  $(i_0^2, \ldots, i_{k_2}^2) \in I_2$ , we write:

$$I_1: (i_0^1, \ldots, i_{k_1}^1) \to I_2: (i_0^2, \ldots, i_{k_2}^2)$$

The set of all the dependences to statements in iteration space I is denoted as  $\delta_I$ .





# 2.1.7 Code abstractions

### Host code abstractions

- ALLOCATEHOSTMEMORY (size): Allocate host array of size elements.
- ALLOCATEBUFFER (size, prop[, hptr]): Allocate device buffer of size elements with access property prop, which can be one of the following:
  - READ\_ONLY: the buffer can only be read on device.
  - WRITE\_ONLY: the buffer can only be written on device.
  - READ\_WRITE: the buffer can be read and written on device.

Additional properties can also be specified, if hptr is given:

- USE\_HOST\_PTR: The slice hptr[0:(size 1)] should be used as buffer storage (default).
- COPY\_HOST\_PTR: The slice hptr[0:(size 1)] should be copied to created device buffer.

This function corresponds to the clCreateBuffer OpenCL API function.

- ALLOCATEIMAGE2D (size, prop[, hptr]): Allocate two dimensional device image of size = (size1, size2) elements with access property prop. The set of properties is the same as for the ALLOCATEBUFFER, but READ\_WRITE is not allowed. This function corresponds to the clCreateImage2D OpenCL API function.
- ALLOCATEIMAGE3D (size,prop[,hptr]): Allocate three dimensional device image of size = (size1,size2,size3) elements with access property prop. The set of properties is the same as for the ALLOCATEBUFFER, but READ\_WRITE is not allowed. This function corresponds to the clCreateImage3D OpenCL API function.
- BUILDPROGRAM (source): Build program from source. This function corresponds to a call to the clCreateProgramWithSource followed by call to the clBuildProgram OpenCL API functions.
- BUILDKERNEL (program, name): Build kernel called name from program obtained from a BUILDPROGRAM call. This function corresponds to the clCreateKernel OpenCL API function.
- GETKERNELINFO (kernel, prop): For kernel, query the value of property prop, which can be:
  - MAX\_WORK\_SIZE: The maximum work-group size the given kernel can be launched with (*cf.* CL\_DEVICE\_MAX\_WORK\_GROUP\_SIZE).
- COPYTODEVICE (dst, src, blocking, {events}): Copy host buffer src to device buffer dst. If Boolean blocking is true, the operation will be synchronous, *i.e.* the call will return after the copy is complete; otherwise, the call will return immediately without



waiting for the operation to complete. Wait for all events in the {events} set to complete before starting the copy. This function corresponds to the <code>clEnqueueWriteBuffer</code> OpenCL API function.<sup>1</sup>

- COPYTOHOST (src,dst,blocking, {events}): Copy device buffer src to host buffer dst. If Boolean blocking is true, the operation will be synchronous, *i.e.* the call will return after the copy is complete; otherwise, the call will return immediately without waiting for the operation to complete. Wait for all events in the {events} set to complete before starting the copy. This function corresponds to the clEnqueueReadBuffer OpenCL API function.
- COPYBUFFER (src,dst, {events}): Copy device buffer src to device buffer dst. Wait for all events in the {events} set to complete before starting the copy. This function corresponds to the clenqueueCopyBuffer OpenCL API function.
- SETKERNELARGUMENTS (kernel, arg0, arg1, ..., argN): Use arg0, arg1, ...argN as arguments for kernel. This function corresponds to a sequence of clSetKernelArgument OpenCL API calls (one for each argument).
- ENQUEUEKERNEL (kernel,gws,lws,{events}): Execute kernel on device, using gws as the global work size and lws as the local work size. Wait for all events in the {events} set to complete before launching kernel. If lws is **none**, the local work size is determined by the runtime. This function corresponds to the clenqueueNDRangeKernel OpenCL API function.

### **Device code abstractions**

- barrier({flags}): All work-items in a work-group must execute this function before any are allowed to continue execution. The {flags} argument can be a logical or of the two values:
  - LOCAL\_MEM: Ensure correct ordering of operations to local memory.
  - GLOBAL\_MEM: Ensure correct ordering of operations to global memory.

This function corresponds to the OpenCL built-in synchronization function barrier.

<sup>&</sup>lt;sup>1</sup>There is a difference in the semantics of a blocking call to this function and a blocking all to clEnqueueWriteBuffer. However, we have not encountered a case where the difference is significant for the host code abstractions provided here.





# 2.2 SHOC Level 1

For this document, we only analyze the Level 1 benchmarks. The Level 0 benchmarks are not representative of real workloads (benchmarking the OpenCL driver and platform). Due to time constraints, analysis of the Level 2 benchmark (S3D) and of the BFS and FFT Level 1 benchmarks is omitted. S3D is very complex (consisting of 27 kernels), but has regular control and memory accesses. The BFS implementation is based on a parallelization technique that is not valid for OpenCL, requiring synchronization between work-groups.<sup>2</sup> Another implementation of the BFS benchmark is described as part of the Rodinia suite. We believe an FFT implementation would be hard to derive automatically from a high-level description.

**Caveat** The SHOC level 1 benchmarks use in-order command queues, which ensures commands are completed in the order they are submitted. Therefore, the event-based synchronization in host code is redundant.

# 2.2.1 Triad

The Triad benchmark streams data from host main memory to device global memory and back: by design, there is no temporal data reuse.

# **High-level description**

Abstract data structures  $\vec{A}, \vec{B}$  and  $\vec{C}$  are *N*-element vectors of floating point numbers.

**Computation**  $\vec{C} = \vec{A} + s\vec{B}$ 

**Iteration space** The iteration space is one dimensional:

$$I = \{i : 0 \le i < N\}$$

Dependences No dependences exist between iterations.

**Memory access mapping** On each iteration  $i \in I$ , one element from each of vectors  $\vec{A}$  and  $\vec{B}$  is read, and one element of vector  $\vec{C}$  is written:

$$M_r(i) = \{A_i, B_i\}$$
  
 $M_w(i) = \{C_i\}$ 

### Low-level implementation details

**Device code** The kernel computes one output element per work-item. See Algorithm 1.

<sup>&</sup>lt;sup>2</sup>The SHOC implementers are working on a replacement version.



Algorithm 1 The Triad device code abstraction. Input: global float A[] Input: global float B[] Input: float s Output: global float C[] 1: kernel TRIAD (A,B,C,s) 2:  $C[\gamma] \leftarrow A[\gamma] + s \times B[\gamma]$ 3: end kernel

**Host code** The one-dimensional iteration space *I* is partitioned into *blocks* of the same size T. The implementation is *double-buffered*: for each array two buffers are used. The device computation for one block is overlapped<sup>3</sup> with transferring the inputs for the next block from host and the outputs of the previous block to host.

The operation is repeated with the following block sizes:  $blockSizes = \{2^{14}, 2^{15}, \dots, 2^{22}\}$  (resulting in 64 KiB, 128 KiB, ..., 16 MiB for single-precision floating-point data).

#### Host data structures

- float  $h[N]: A_x \to h[x], B_x \to h[x], C_x \to h[x]$
- float dA[2][N]:  $A_x \rightarrow dA[(x/T)\&2][x\&T]$
- float dB[2][N]:  $B_x \rightarrow dB[(x/T)$  %2][x%T]
- float dC[2][N]:  $C_x \rightarrow dC[(x/T) \& 2][x\& T]$

### **Device data structures**

- global float  $A[N]: A_x \to A[x \text{ST}]$
- global float  $B[N]: B_x \to B[x \text{ st}]$
- global float  $C[N]: C_x \to C[x & T]$

**Input datasets** The input data is random floating point numbers in the range [0.0, 10.0).

**Partitioned iteration space** A given block size *T* results in the following iteration space:

$$P = \{(b, \gamma) : 0 \le b < N/T, 0 \le \gamma < \Gamma = T\}$$

The following function maps each iteration  $i \in I$  to iteration  $(b, \gamma) \in P$ :

$$f: i \to (i/T, i\%T)$$

The following function gives the inverse mapping:

$$f^{-1}: (b, \gamma) \to T \times b + \gamma$$

CARP-ARM-RP-001-v1.0

<sup>&</sup>lt;sup>3</sup>At least conceptually: see caveat in the beginning of  $\S 2.2$ .



Algorithm 2 The Triad host code abstraction. 1: procedure ALLOCATEMEMORY() ▷ Allocate host memory and device buffers.  $h \leftarrow ALLOCATEHOSTMEMORY(N)$ 2: 3:  $dA[0] \leftarrow ALLOCATEBUFFER(N, READ_ONLY)$  $dA[1] \leftarrow ALLOCATEBUFFER(N, READ_ONLY)$ 4: 5:  $dB[0] \leftarrow ALLOCATEBUFFER(N, READ_ONLY)$  $dB[1] \leftarrow ALLOCATEBUFFER(N, READ_ONLY)$ 6:  $dC[0] \leftarrow ALLOCATEBUFFER(N, WRITE_ONLY)$ 7:  $dC[1] \leftarrow ALLOCATEBUFFER(N, WRITE_ONLY)$ 8: 9: end procedure 10: **procedure** INIT(h[0: N-1])  $\triangleright$  Initialize host memory. for *i* in 0: (N/2 - 1) do 11:  $h[i] \leftarrow RAND(0.0, 10.0)$ 12: 13:  $h[i+N/2] \leftarrow h[i]$ end for 14: 15: end procedure 16: **procedure** VERIFY(h[0: N-1]) ▷ Verify results. for *i* in 0 : (N/2 - 1) do 17: 18: if  $h[i] \neq h[i+N/2]$  then ERROR() 19: 20: end if end for 21: 22: end procedure 23: procedure MAIN ALLOCATEMEMORY() 24: ⊳ See Algorithm 1. 25: ▷ Build two copies of 26:  $kernel[0] \leftarrow BUILDKERNEL(program, "triad")$  $kernel[1] \leftarrow BUILDKERNEL(program, "triad")$  $\triangleright$  the triad kernel. 27:  $maxLocalWorkSize \leftarrow GetKernelInfo(MAX_WORK_SIZE)$ 28:  $\Lambda \leftarrow MAX(128, maxLocalWorkSize)$  $\triangleright$  Local work size  $\leq 128$ . 29: for T in blockSizes do 30: INIT(h[]) 31:  $\Gamma \gets \mathtt{T}$ ▷ Global work size. 32:  $\texttt{toDeviceA}[0] \leftarrow \text{COPYTODEVICE}(\texttt{dA}[0][0:\texttt{T}-1],\texttt{h}[0:\texttt{T}-1],\texttt{false},\emptyset)$ 33:  $toDeviceB[0] \leftarrow COPYTODEVICE(dB[0][0:T-1],h[0:T-1],false,\emptyset)$ 34: ▷ Start first block transfer to device. SETKERNELARGUMENTS(kernel[0], dA[0], dB[0], dC[0], 1.75f) 35:  $\texttt{nqKernel}[0] \leftarrow \texttt{ENQUEUEKERNEL}(\texttt{kernel}[0], \Gamma, \Lambda, \{\texttt{toDeviceA}[0], \texttt{toDeviceB}[0]\})$ 36:  $\triangleright$  Start first kernel execution, wait for writes to dA[0][] and dB[0][]. 37: if T < N then  $toDeviceA[1] \leftarrow COPYTODEVICE(dA[1][0:T-1],h[T:2 \times T-1],false,\emptyset)$ 38:  $\texttt{toDeviceB}[1] \leftarrow \text{COPYTODEVICE}(\texttt{dB}[1][0:\texttt{T}-1],\texttt{h}[\texttt{T}:2\times\texttt{T}-1],\texttt{false},\emptyset)$ 39: > Start second block transfer to device. end if 40:



```
Algorithm 2 The Triad host code abstraction (continued).
41:
             b \leftarrow 1
                                                                                          ▷ Block index.
             i \leftarrow 0
                                                                                     ⊳ Host array index.
42:
             while i < N do
43:
                 \texttt{this} \leftarrow b\%2
                                                                           \triangleright 0 if b is even, 1 otherwise.
44:
                 \texttt{that} \leftarrow 1 - \texttt{this}
                                                                           \triangleright 1 if b is even, 0 otherwise.
45:
                 toHostC[that]
                                              COPYTOHOST(dC[that]]0 : T-1], h[i : i +
                                      \leftarrow
46:
    T-1], false, {nqKernel[that]})
                                                       ▷ Start transfer of previous block from device.
                 i \leftarrow i + T
                                                                                        ▷ Increase index.
47:
                 if i < N then
48:
                      SETKERNELARGUMENTS(kernel[this], dA[this], dB[this], dC[this],
49:
     1.75f
                      eventsToWaitOn \leftarrow \{toDeviceA[this], toDeviceB[this]\}
50:
                      if b > 1 then
51:
                          eventsToWaitOn \leftarrow eventsToWaitOn \cup \{toHostC[this]\}
52:
53:
                      end if
                      nqKernel[this] \leftarrow ENQUEUEKERNEL(kernel[this],
                                                                                                Γ,
                                                                                                        Λ,
54:
     eventsToWaitOn)
                 end if
55:
                 if i + T < N then
56:
                      \texttt{toDeviceA[that]} \leftarrow \texttt{COPYTODEVICE}(\texttt{dA[that]}[0:\texttt{T}-1],\texttt{h}[i+\texttt{T}:i+2\times
57:
    T-1, false, {nqKernel[that]})
                      \texttt{toDeviceB[that]} \leftarrow \texttt{COPYTODEVICE}(\texttt{dB[that]}[0:\texttt{T}-1],\texttt{h}[i+\texttt{T}:i+2\times
58:
    T-1, false, {nqKernel[that]})
                                                               ▷ Start transfer of next block to device.
                 end if
59:
                 b \leftarrow b + 1
                                                                                ▷ Increase block index.
60:
61:
             end while
             WAITFOREVENTS({toHostC[that]})
                                                                  ▷ Wait for last transfer from device.
62:
63:
             VERIFY(h[])
                                                                                         ▷ Verify results.
         end for
64:
65: end procedure
```





Global memory access mapping Each work-item performs two reads and one write:

$$G_r(\gamma) = \{ \mathbb{A}[\gamma], \mathbb{B}[\gamma] \}$$
  
 $G_w(\gamma) = \{ \mathbb{C}[\gamma] \}$ 

### **Performance metrics**

- 1. Speed (*GFLOP/s*): the total number of floating point operations (2*N*) divided by the total time for all data transfers plus calculations.
- 2. Bandwidth (*GB/s*): the total number of memory operations (2N reads and N writes) divided by the total time for all data transfers plus calculations.

Data transfers between host and device are expected to dominate the execution time, and better performance is expected for larger block sizes.

**Verification mechanism** The two halves of the input buffers are initialized with the same data, so the two halves of the output buffer are expected to contain the same results. This is checked after the operation.

**Target-specific optimisations applied** Allocating memory for the host buffer can be done either using the C++ new operator (or equivalent methods *e.g.* malloc), or using clCreateBuffer with the flag ALLOC\_HOST\_PTR. This benchmark uses the latter method, which is good practice for NVIDIA cards because the allocated pages get pinned to main memory, enabling copies to device global memory to use DMAs instead of keeping the host busy.

**Target-specific optimization opportunities** Kernel code can be easily vectorized, which would improve the performance of this benchmark on vector-based architectures.

On unified memory systems, the operations for reading and writing the buffers would result in unnecessary memory copies, which would give misleading memory bandwidth figures. On such systems, it would be beneficial to replace calls to clEnqueueReadBuffer and clEnqueueWriteBuffer with calls to clEnqueueMapBuffer.





# 2.2.2 Stencil2D

### **High-level description**

The benchmark performs a 9-point two-dimensional (2D) stencil on floating point data.

Typically, the computation is repeated until convergence is achieved, which implies a variable number of iterations depending on the initial conditions. However, in this benchmark, the computation is repeated only for a fixed number of time steps T.

**Abstract data structures** The benchmark operates on a three-dimensional array *A* of floating point values, where the first dimension *t* corresponds to the time step and the other two dimensions (i, j) correspond to the position in a two dimensional grid of size  $(R+2) \times (C+2)$ .

**Computation** The following computation takes place at every time step t for each inner element of array A:

$$A_{t,i,j} = w_0 \times A_{t-1,i,j} + w_c \times (A_{t-1,i-1,j} + A_{t-1,i,j-1} + A_{t-1,i+1,j} + A_{t-1,i,j+1}) + w_d \times (A_{t-1,i-1,j-1} + A_{t-1,i-1,j+1} + A_{t-1,i+1,j-1} + A_{t-1,i+1,j+1})$$

**Iteration space** 

$$I = \{(t, i, j) : 0 \le t < T, 1 \le i \le R + 1, 1 \le j \le C + 1\}$$

Dependences

$$\begin{split} \delta_{I} &= I: (t-1,i-1,j-1) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i-1,j) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i-1,j+1) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i,j-1) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i,j) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i,j+1) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i+1,j-1) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i+1,j-1) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i+1,j) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i+1,j+1) \xrightarrow{t} I: (t,i,j) \land \\ I: (t-1,i+1,j+1) \xrightarrow{t} I: (t,i,j) \end{split}$$

**Memory access mapping** On each iteration, there are nine reads and one write:

$$M_{r}(t,i,j) = \{A_{t-1,i,j}, A_{t-1,i-1,j}, A_{t-1,i,j-1}, A_{t-1,i+1,j}, A_{t-1,i,j+1}, A_{t-1,i-1,j-1}, A_{t-1,i-1,j+1}, A_{t-1,i+1,j-1}, A_{t-1,i+1,j+1}\}$$
  

$$M_{w}(t,i,j) = \{A_{t,i,j}\}$$

#### Low-level implementation details

The benchmark operates on floating point data of type **TYPE**, where **TYPE** is **float** for single precision and **double** for double precision.





**Device code** There are two kernels in this benchmark:

- CopyRect () This kernel is similar in functionality to the clEnqueueCopyBufferRect API call of the OpenCL 1.1 specification, implementations of which were not widely available when the benchmark was written. It copies the left and right halo columns from one device buffer to another. See Algorithm 3.
- Stencil() Each kernel launch performs one time step. See Algorithm 4.

```
Algorithm 3 The Stencil2D device code abstraction (CopyRect).
Output: global TYPE dst[]
                                                                              ▷ Offset for dst.
Input: int doffset
Input: int dpitch
                                                                               ▷ Pitch for dst.
Input: global TYPE src[]
                                                                              ▷ Offset for src.
Input: int soffset
                                                                               ▷ Pitch for src.
Input: int spitch
                                                                 ▷ Number of columns to copy.
Input: int width
Input: int height
                                                                     ▷ Number of rows to copy.
 1: kernel COPYRECT (dst,doffset,dpitch,src,soffset,spitch,width,height)
 2:
        r \leftarrow v_0 \times \Lambda_0 + \lambda_0
                                                                                       \triangleright \mathbf{r} = \gamma_0
 3:
        if r < height then
 4:
            for c in 0: (width -1) do
               dst[r \times dpitch + doffset + c] \leftarrow src[r \times spitch + soffset + c]
 5:
            end for
 6:
        end if
 7:
 8: end kernel
```

Host code See Algorithm 5.

**Host data structures** Conceptually, the benchmark operates on a three-dimensional array, the first dimension being time. However, keeping all the intermediate results is unnecessary. Instead, the computation can be done by alternating between two arrays, one for the current time step and one for the previous time step.

- **TYPE** data[Rp \* Cp]:  $A_{0,x,y} \rightarrow \text{data}[x * Cp + y]$ ,  $A_{T-1,x,y} \rightarrow \text{data}[x * Cp + y]$
- **TYPE** buffer0[Rp \* Cp]:  $A_{t,x,y} \rightarrow \text{buffer0}[x * Cp + y]$ , if t is even
- **TYPE** buffer1[Rp \* Cp]:  $A_{t,x,y} \rightarrow \text{buffer1}[x * Cp + y]$ , if t is odd

Device data structures For the Stencil() kernel:

- global TYPE din[Rp \* Cp]:  $A_{t,x,y} \rightarrow \text{din}[x * Cp + y]$
- global TYPE dout[Rp \* Cp]:  $A_{t,x,y} \rightarrow \text{dout}[x * Cp + y]$





```
Algorithm 4 The Stencil2D device code abstraction (Stencil).
Input: int row
                                                                             ▷ Row index, including halo.
Input: int col
                                                                        ▷ Column index, including halo.
                                                                             ▷ Row pitch, including halo.
Input: int rowPitch
  1: function GETOFFSET(row, col, rowPitch)
 2:
         return row × rowPitch + col
 3: end function
Input: global TYPE din[]
                                                ▷ Alignment required for rows (determines padding).
Input: int alignment
                                                                         ▷ Weights for stencil operation.
Input: TYPE w0, wc, wd
Input: int LR
                                                                       ▷ Number of rows per work-item.
Output: global TYPE dout[]
Local: local TYPE ld[LRp×LCp]
                                                               \triangleright LRp = LR + 2, LCp = LC + 2 = \Lambda_1 + 2
 4: kernel STENCIL (din, dout, alignment, w0, wc, wd)
 5:
         \mathtt{r} \leftarrow v_0 	imes \mathtt{LR}
                                                                      \triangleright Row in the original grid, \Lambda_0 = 1.
         c \leftarrow v_1 \times \Lambda_1 + \lambda_1
                                                                           ▷ Column in the original grid.
 6:
 7:
         C \leftarrow N_1 \times \Lambda_1 + 2
                                                                         ▷ Total columns including halo.
         C_p \leftarrow C
 8:
         if \neg C_p : alignment then
 9:
                                                                                       ▷ Misaligned pitch.
10:
             C_p \leftarrow \lceil C_p / \texttt{alignment} \rceil \times \texttt{alignment}
                                                                                      \triangleright Next aligned pitch.
         end if
11:
         for lr in 0: (LR + 1) do \triangleright Copy data in global memory to local memory (excluding left
12:
     and right halo).
             lidx \leftarrow GETOFFSET(\lambda_0 + lr, \lambda_1 + 1, \Lambda_1 + 2)
13:
14:
             gidx \leftarrow GETOFFSET(r+lr, c+1, C_p)
15:
             ld[lidx] \leftarrow din[gidx]
         end for
16:
         if \lambda_1 = 0 then
                                                                       \triangleright Copy left halo to local memory.
17:
             for lr in 0 : (LR + 1) do
18:
19:
                  lidx \leftarrow GETOFFSET(\lambda_0 + lr, \lambda_1, \Lambda_1 + 2)
                  gidx \leftarrow GETOFFSET(r+lr,c,C_p)
20:
                  ld[lidx] \leftarrow din[gidx]
21:
22:
             end for
         end if
23:
         if \lambda_1 = \Lambda_1 - 1 then
                                                                     ▷ Copy right halo to local memory.
24:
25:
             for lr in 0 : (LR + 1) do
                  lidx \leftarrow GETOFFSET(\lambda_0 + lr, \lambda_1 + 2, \Lambda_1 + 2)
26:
                  gidx \leftarrow GETOFFSET(r+lr, c+2, C_p)
27:
                  ld[lidx] \leftarrow din[gidx]
28:
             end for
29:
         end if
30:
                                                                       ▷ Wait until transfers are finished.
         BARRIER(CLK_LOCAL_MEM_FENCE)
31:
```



Alg	orithm 4 The Stencil2D device code abstraction (Stencil) (	continued).
32:	for lr in $0$ : (LR $- 1$ ) do	⊳ Stencil operation.
33:	$\texttt{cidx} \gets \texttt{GETOFFSET}(\lambda_0 + 1 + \texttt{lr}, \lambda_1 + 1, \Lambda_1 + 2)$	
34:	$\texttt{nidx} \gets \texttt{GETOFFSET}(\lambda_0 + \texttt{lr}, \lambda_1 + 1, \Lambda_1 + 2)$	
35:	$\texttt{sidx} \gets \texttt{GETOFFSET}(\lambda_0 + 2 + \texttt{lr}, \lambda_1 + 1, \Lambda_1 + 2)$	
36:	$\texttt{eidx} \gets \texttt{GETOFFSET}(\lambda_0 + 1 + \texttt{lr}, \lambda_1 + 2, \Lambda_1 + 2)$	
37:	$\texttt{widx} \gets \texttt{GETOFFSET}(\lambda_0 + +1\texttt{lr}, \lambda_1, \Lambda_1 + 2)$	
38:	$\texttt{neidx} \gets \texttt{GETOFFSET}(\lambda_0 + \texttt{lr}, \lambda_1 + 2, \Lambda_1 + 2)$	
39:	$\texttt{seidx} \gets \texttt{GETOFFSET}(\lambda_0 + 2 + \texttt{lr}, \lambda_1 + 2, \Lambda_1 + 2)$	
40:	$\texttt{nwidx} \gets \texttt{GETOFFSET}(\lambda_0 + \texttt{lr}, \lambda_1, \Lambda_1 + 2)$	
41:	$\texttt{swidx} \gets \texttt{GETOFFSET}(\lambda_0 + 2 + \texttt{lr}, \lambda_1, \Lambda_1 + 2)$	
42:	$\texttt{centerValue} \gets \texttt{ld}[\texttt{cidx}]$	
43:	$\texttt{cardinalValues} \gets \texttt{ld[nidx]} + \texttt{ld[sidx]} + \texttt{ld[eidx]}$	] + ld[widx]
44:	$\texttt{diagonalValues} \gets \texttt{ld}[\texttt{neidx}] + \texttt{ld}[\texttt{seidx}] + \texttt{ld}[\texttt{nw}]$	idx] + ld[swidx]
45:	$dout[GETOFFSET(r + 1 + lr, c + 1, C_p)] \leftarrow w0$	$\times$ centerValue + wc $\times$
	cardinalValues + wd  imes diagonalValues	
46:	end for	
47:	end kernel	

Alg	gorithm 5 The Stencil2D host code abs	traction.
1:	<b>procedure</b> INITIALIZEGRID( $M[R+2]$ )	$[C+2]$ , haloVal) $\triangleright$ Initialize stencil grid. <sup><i>a</i></sup>
2:	<b>for i in</b> 1 : <b>R do</b>	▷ Initialize grid, excluding halo.
3:	<b>for</b> j <b>in</b> 1 : C <b>do</b>	
4:	$M[i][j] \leftarrow RAND(0.0,1.0)$	
5:	end for	
6:	end for	
7:	<b>for</b> $j$ <b>in</b> $0 : (C+1)$ <b>do</b>	▷ Initialize top and bottom halo rows.
8:	$\texttt{M}[0][\texttt{j}] \gets \texttt{haloVal}$	
9:	$\mathtt{M}[\mathtt{R}+1][\mathtt{j}] \leftarrow \mathtt{haloVal}$	
10:	end for	
11:	<b>for i in</b> 1 : <b>R do</b>	▷ Initialize left and right halo columns.
12:	$\texttt{M[i]}[0] \gets \texttt{haloVal}$	
13:	$\mathtt{M}[\mathtt{i}][\mathtt{C}+1] \leftarrow \mathtt{haloVal}$	
14:	end for	
15:	end procedure	

<sup>a</sup>When M is passed to this function, it is actually a flattened, one-dimensional array. Two-dimensional indexing is used for convenience.



```
Algorithm 5 The Stencil2D host code abstraction (continued).
 16: procedure PERFORMSTENCIL(input[R+2][C+2], T, w0, wc, wd)
                                                                                                                                                                                                         ▷ Perform stencil computation on host.<sup>a</sup>
                             tmp \leftarrow ALLOCATEMEMORY((R+2) \times (C+2))
  17:
                             for t in 0: (T-1) do
  18:
 19:
                                           \texttt{tmp}[] \leftarrow \texttt{input}[]
                                                                                                                                                                                                                                                                                                               ⊳ Copy grid.
                                           for i in 1 : R do
 20:
                                                         for j in 1 : C do
 21:
 22:
                                                                        oldCenterValue \leftarrow tmp[i][j]
                                                                        \texttt{oldNSEWValues} \leftarrow \texttt{tmp}[\texttt{i} - 1][\texttt{j}] + \texttt{tmp}[\texttt{i} + 1][\texttt{j}] + \texttt{tmp}[\texttt{i}][\texttt{j} - 1] + \texttt{tmp}[\texttt{i} - 1][\texttt{j}] + \texttt{tmp}[\texttt{i} - 1][\texttt{j} - 1] + \texttt{tmp}[\texttt{j} -
 23:
                tmp[i][j+1]
                                                                       oldDiagonalValues \leftarrow tmp[i-1][j-1] + tmp[i+1][j-1] + tmp[i-1]
 24:
                1][j+1] + tmp[i+1][j+1]
                                                                       input[i][j] \leftarrow w0 \times oldCenterValue + wc \times oldNSEWValues + wd \times
 25:
                oldDiagonalValues
                                                         end for
 26:
 27:
                                           end for
                             end for
 28:
 29: end procedure
 30: procedure VALIDATERESULT(expected[0 : N \times M - 1], actual[0 : N \times M - 1]
                1], relErrorThreshold)
                                                                                                                                                                                                                                                                                            ▷ Validate results.
                             for i in 0 : (N - 1) do
 31:
                                           for j in 0 : (M - 1) do
 32:
 33:
                                                          expVal \leftarrow expected[i \times M + j]
                                                          actualVal \leftarrow actual[i \times M + j]
 34:
                                                         delta \leftarrow |actualVal - expVal|
 35:
                                                         if expVal \neq 0 then
 36:
                                                                       relError \leftarrow delta/expVal
 37:
                                                         else
 38:
                                                                       \texttt{relError} \gets 0
                                                                                                                                                                                                                                                                                                    ▷ Potential bug.
 39:
                                                         end if
 40:
 41:
                                                         if relError > relErrorThreshold then ERROR()
                                                         end if
 42:
 43:
                                           end for
 44:
                             end for
 45: end procedure
```

 $^{a}$ When *input* is passed to this function, it is actually a flattened, one-dimensional array. Two-dimensional indexing is used for convenience.



Algorithm 5 Stencil2D : Host code abstrac	tion (continued)
Input: int R	▷ Number of rows in the stencil grid.
Input: int C	▷ Number of columns in the stencil grid.
Input: int T	⊳ Time steps.
Input: float haloVal	▷ Value for halo points of the grid.
Input: TYPE w0	▷ Weights for the stencil operation.
Input: TYPE WC	
Input: TYPE wd	
Input: int LR	▷ Number of rows per work-item.
Input: int LC	$\triangleright \Lambda_1$
46: $\texttt{alignment} \leftarrow 16$	▷ Architecture-specific?
47: $R_p \leftarrow R+2$	
48: $C_p \leftarrow C + 2$	
49: <b>if</b> $\neg C_p$ : alignment <b>then</b>	▷ Misaligned pitch.
50: $C_p \leftarrow \lceil C_p / \texttt{alignment} \rceil \times \texttt{alignment}$	ent ▷ Next aligned pitch.
51: <b>end if</b>	
52: expected $\leftarrow$ AllocateMemory(R <sub>p</sub>	$(\mathbf{x} \times \mathbf{C}_{\mathbf{p}})$
53: INITIALIZEGRID(expected, haloVal	)
54: PERFORMSTENCIL(expected,T)	
55: actual $\leftarrow$ ALLOCATEMEMORY(R <sub>p</sub> $\times$	C <sub>p</sub> )
56: INITIALIZEGRID(actual, haloVal)	
57: stencilProgram $\leftarrow$ BUILDPROGRAM	M("stencil2d.cl")
58: $copyRectKernel \leftarrow BUILDKERNEL($	<pre>stencilProgram,"CopyRect")</pre>
59: stencilKernel $\leftarrow$ BUILDKERNEL(s	tencilProgram,"Stencil")
$\texttt{60: buffer0} \gets \texttt{ALLOCATEBUFFER}(\texttt{R}_p \times$	$C_p$ , READ_WRITE) $\triangleright$ Allocate device buffers.
$\texttt{61: buffer1} \gets \texttt{ALLOCATEBUFFER}(\texttt{R}_p \times$	C <sub>p</sub> ,READ_WRITE)
62: $in \leftarrow buffer0$	$\triangleright$ Pointers to the buffers.
63: out $\leftarrow$ buffer1	



Algorithm 5 Stencil2D : Host code abstraction (continued)

64: writeEvent  $\leftarrow$  COPYTODEVICE(in, actual, False,  $\emptyset$ ) 65:  $COPYBUFFER(in[0:C_p-1], out[0:C_p-1], {writeEvent}) \triangleright Copy top and bottom halo Copy top and Copy t$ rows.  $\texttt{66: COPYBUFFER}(\texttt{in}[(\texttt{R}_p-1)\times\texttt{C}_p:\texttt{R}_p\times\texttt{C}_p-1],\texttt{out}[(\texttt{R}+1)\times\texttt{C}_p:\texttt{R}_p\times\texttt{C}_p-1],\emptyset)$  $\triangleright$ Synchronization here is problematic. 67: SETKERNELARGUMENTS(copyRectKernel, out,  $0, C_p, in, 0, C_p, 1, R_p) \triangleright$  Copy left and right halo columns. 68:  $cwEvent \leftarrow ENQUEUEKERNEL(copyRectKernel, R_p, none, \emptyset)$ 69: SETKERNELARGUMENTS(copyRectKernel,out,  $C + 1, C_p, in, C + 1, C_p, 1, R$ ) 70: ceEvent  $\leftarrow$  ENQUEUEKERNEL(copyRectKernel, R<sub>p</sub>, **none**,  $\emptyset$ ) 71: WAITFOREVENTS({cwEvent, ceEvent}) 72: prevKernelEvent  $\leftarrow \emptyset$ 73: **for** t **in** 1 : (T) **do** SETKERNELARGUMENTS(stencilKernel, in, out, alignment, w0, wc, wd) 74: 75:  $kernelEvent \leftarrow ENQUEUEKERNEL(stencilKernel, (R/LR, C), (1, LC), prevKernelEvent)$  $prevKernelEvent \leftarrow \{kernelEvent\}$ 76: 77: if in = buffer0 then ⊳ Swap buffers.  $\texttt{in} \leftarrow \texttt{buffer1}$ 78: 79:  $\texttt{out} \leftarrow \texttt{buffer0}$ 80: else 81:  $\texttt{in} \gets \texttt{buffer0}$  $\texttt{out} \leftarrow \texttt{buffer1}$ 82: end if 83: 84: end for 85: COPYTOHOST(in, actual, True, prevKernelEvent) ▷ Buffers have been swapped, so in is the buffer we should copy to host. 86: VALIDATERESULT(expected, actual)



• local TYPE ld[LRp \* LCp]:

 $A_{t,x,y} \rightarrow \operatorname{ld}[(x - 1) \ \text{\ } \operatorname{LR} \ + \ 1][(y - 1) \ \text{\ } \operatorname{LC} \ + \ 1], \text{ if } 1 \leq x \leq R \wedge 1 \leq y \leq C$ 

The mapping of elements of A[] to the halo elements of ld[] is omitted, as it is fairly complicated.<sup>4</sup>

**Input datasets** There are four classes for the sizes of the datasets on which the benchmark operates: classes 1, 2, 3 and 4 work on arrays of dimensions  $512 \times 512$ ,  $1024 \times 1024$ ,  $2048 \times 2048$  and  $4096 \times 4096$  respectively (the dimensions do not include the halo, for which additional memory is allocated). The computation is repeated for random single and double precision floating point inputs in the range [0.0, 1.0).

**Command line parameters** The following command line parameters are available:

- seed: The seed for the random number generator.
- haloVal: The value used for the halo elements; the default is 0.
- weight-center, weight-cardinal, weight-diagonal: The weights used for stencil computation; the defaults are  $(w_0, w_c, w_d) = (0.25, 0.15, 0.05)$ .
- lsize: The parameters for partitioning the iteration space; the defaults are (LR, LC) = (8, 256).
- val-threshold: The relative error threshold used for validation; the default is 0.01.

**Partitioned iteration space** The following partitioned iteration space is used:

$$P = \{(t, v_0, v_1, \lambda_0, \lambda_1, l): 0 \le t < T, \\ 0 \le v_0 < N_0 = R/LR, \\ 0 \le v_1 < N_1 = C/LC, \\ \lambda_0 = 0, 0 \le \lambda_1 < \Lambda_1 = LC, \\ 0 < l < LR\}$$

which means the following hold:

$$\begin{array}{lll} (\Gamma_0, \Gamma_1) & = & (R/LR, C) \\ (\gamma_0, \gamma_1) & = & (\nu_0, \nu_1 \times LC + \lambda_1) \end{array}$$

Every iteration (t, i, j) of the original iteration space maps to this new iteration according to the following function:

$$f: (t,i,j) \rightarrow (t,i/LR,j/LC,0,j\% LC,i\% LR)$$

The inverse mapping is given by the function:

$$f^{-1}: (t, \mathbf{v}_0, \mathbf{v}_1, \lambda_0, \lambda_1, l) \to (t, \mathbf{v}_0 \times LR + l, \mathbf{v}_1 \times LC + \lambda_1)$$

CARP-ARM-RP-001-v1.0

<sup>&</sup>lt;sup>4</sup>There are eight different cases: four for the corner halo elements, one for each of the top and bottom halo rows and one for each of the left and right halo columns.





No special care (e.g. row padding, special handling of remaining elements) is taken for the case when C is not divided exactly by LC. Therefore we assume in this subsection that C is a multiple of LC.

For copying data between global and local memory, a slightly different iteration space is used:

$$P = \{(t, v_0, v_1, \lambda_0, \lambda_1, l): 0 \le t < T, \\ 0 \le v_0 < R/LR, \\ 0 \le v_1 < C/LC, \\ 0 \le \lambda_0 < 1, 0 \le \lambda_1 < LC, \\ 0 \le l < LR + 2\}$$

Note that the only difference is two additional iterations for the innermost index.

**Device memory access mapping** Each work-item executing the Stencil() kernel performs the following reads and writes:

• For copying data between global and local memory:

• For the actual stencil computation:

$$\begin{split} L_r(t, \mathbf{v}_0, \mathbf{v}_1, \lambda_0, \lambda_1, l) &= & \{ \texttt{ld}[(\lambda_0 + l + 1) \times (\Lambda_1 + 2) + \lambda_1 + 1], \\ & \texttt{ld}[(\lambda_0 + l) \times (\Lambda_1 + 2) + \lambda_1 + 1], \\ & \texttt{ld}[(\lambda_0 + l + 2) \times (\Lambda_1 + 2) + \lambda_1 + 1], \\ & \texttt{ld}[(\lambda_0 + l + 1) \times (\Lambda_1 + 2) + \lambda_1 + 2], \\ & \texttt{ld}[(\lambda_0 + l) \times (\Lambda_1 + 2) + \lambda_1 + 2], \\ & \texttt{ld}[(\lambda_0 + l) \times (\Lambda_1 + 2) + \lambda_1 + 2], \\ & \texttt{ld}[(\lambda_0 + l) \times (\Lambda_1 + 2) + \lambda_1 + 2], \\ & \texttt{ld}[(\lambda_0 + l) \times (\Lambda_1 + 2) + \lambda_1 + 2], \\ & \texttt{ld}[(\lambda_0 + l + 2) \times (\Lambda_1 + 2) + \lambda_1 + 2], \\ & \texttt{ld}[(\lambda_0 + l + 2) \times (\Lambda_1 + 2) + \lambda_1], \\ & \texttt{ld}[(\lambda_0 + l + 2) \times (\Lambda_1 + 2) + \lambda_1] \} \\ G_w(t, \mathbf{v}_0, \mathbf{v}_1, \lambda_0, \lambda_1, l) &= & \{\texttt{dout}[(N_0 \times LR + \lambda_0 + l + 1) \times \mathbb{C}_p + \gamma_1 + 1] \} \end{split}$$





**Performance metrics** Performance is evaluated based on a time measurement that includes the transfers of data between host and device, one execution of CopyRect and *T* executions of Stencil.

**Validation mechanism** The results from a run on device are compared again a reference run on host. Since the comparison is between floating point numbers, the relative error is compared against a threshold (val-threshold).

The validation mechanism contains a potential bug: when the expected value is zero, the relative error is also set to zero regardless of the actual value, missing potential errors.

**Target-specific optimisations** The kernel uses local memory. This increases performance when local memory is implemented on-chip (as on desktop GPUs). This, however, might decrease performance when local memory is in the same physical space as global memory, since it implies redundant memory copies and barriers.

**Target-specific optimization opportunities** Vectorization can potentially improve performance if the penalty for unaligned memory access is tolerable.





# 2.2.3 Reduction

This benchmark performs a reduction operation on a vector of floating point numbers.

### **High-level description**

The reduction operation is defined as follows:

$$O = \sum_{i} I_i, 0 \le i < N$$

### Abstract data structures

- $\vec{I}$  is an input *N*-element vector.
- $\vec{R}$  is an intermediate *B*-element vector.
- *O* is an output value.

**Computation** Range [0: (N-1)] is partitioned into *B* disjoint ranges by a monotonically increasing sequence  $\{i_k\}: 0 = i_0 < i_1 < ... < i_B = N$ :

$$[0:(N-1)] = \bigcup_{b=0}^{B-1} [i_b:(i_{b+1}-1)]$$

Reduction (2.10) is performed in two steps:

• Reduction on blocks:

$$R_b = \sum_j I_j, \ 0 \le b < B \land i_b \le j < i_{b+1}$$

• Reduction of intermediate results:

$$O = \sum_{b} R_b, 0 \le b < B$$

Reduction is a well-known parallel primitive, which is assumed to be associative and commutative (*i.e.* additions can be performed in any order). Since the iteration space, dependences and memory accesses depend on how this primitive is actually implemented, we only analyze the low-level implementation.

### Low-level implementation details

The benchmark operates on floating point data of type **TYPE**, where **TYPE** is **float** for single precision and **double** for double precision.

The first step of the algorithm is performed on the device, whereas the second step is performed on the host.

**Device code** See Algorithm 6.

CARP-ARM-RP-001-v1.0



```
Algorithm 6 The Reduction device code abstraction.
Input: global const TYPE I[N]
Output: global TYPE R[B]
                                                                                                ▷ Partial reduction results.
Local: local TYPE lr[\Lambda_0]
  1: kernel REDUCE (I,R)
 2:
           lr[\lambda_0] \leftarrow 0
           for i in [(2 	imes \Lambda_0 	imes v_0 + \lambda_0) : (N-1) : (2 	imes \Lambda_0 	imes N_0)] do
 3:
                                                                       ▷ Reduce multiple elements per work-item.
                \texttt{lR}[\lambda_0] \leftarrow \texttt{lR}[\lambda_0] + \texttt{I}[\texttt{i}] + \texttt{I}[\texttt{i} + \Lambda_0]
 4:
           end for
 5:
           BARRIER(CLK_LOCAL_MEM_FENCE)
 6:
 7:
           for s \leftarrow \Lambda_0/2 until s \le 0 step s \leftarrow \lfloor s/2 \rfloor do \triangleright Tree-like reduction of partial results.
                if \lambda_0 < s then
 8:
                     \texttt{lR}[\lambda_0] \leftarrow \texttt{lR}[\lambda_0] + \texttt{lR}[\lambda_0 + \texttt{s}]
 9:
10:
                end if
                BARRIER(CLK_LOCAL_MEM_FENCE)
11:
           end for
12:
           if \lambda_0 = 0 then
                                                                                       ▷ Write result to global memory.
13:
                \mathtt{R}[\textit{v}_0] \gets \mathtt{l}\mathtt{R}[0]
14:
           end if
15:
16: end kernel
```



Algorithm 7 The Reduction host code abstraction.







### Host code See Algorithm 7.

#### Host data structures

- TYPE hI[N]:  $I_x \to$  hI[x]
- TYPE hR[B]:  $R_x \rightarrow hR[x]$

### **Device data structures**

- global const TYPE  $I[N]: I_x \to I[X]$
- global TYPE  $R[B]: R_x \to R[x]$
- local TYPE  $lr[\Lambda_0]$ : intermediate results

**Input datasets** The input sizes considered are  $2^{20}, 2^{23}, 2^{25}$  and  $2^{26}$  bytes of data, which result in  $N \in \{2^{18}, 2^{21}, 2^{23}, 2^{24}\}$  for single-precision and  $N \in \{2^{17}, 2^{20}, 2^{22}, 2^{23}\}$  for double-precision.

The *i*<sup>th</sup> element of the input array is initialized with the value i%3, resulting in a pattern of  $\{0, 1, 2, 0, 1, 2, ...\}$ .

Partitioned iteration space The algorithm (reduction) uses the following iteration spaces:

• Algorithm 6, line 2:

$$P_{R_1} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256\}$$

• Algorithm 6, lines 3–5:

$$P_{R_2} = \{ (v_0, \lambda_0, i) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256, \\ i = 2 \times \Lambda_0 \times v_0 + k \times 2 \times \Lambda_0 \times N_0, 0 \le k < (N - 2 \times \Lambda_0 \times v_0 - \lambda_0) / (2 \times \Lambda_0 \times N_0) \}$$

• Algorithm 6, lines 7–12:

$$P_{R_3} = \{ (v_0, \lambda_0, s) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256, s = \lfloor \Lambda_0/2^k \rfloor, k \in \{1, 2, \dots, \log_2 \Lambda_0\} \}$$

• Algorithm 6, lines 13–15:

$$P_{R_4} = \{ (v_0, \lambda_0) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256 \}$$

#### **Device memory access mapping**

• For  $P_{R_1}$ :





• For  $P_{R_2}$ :

• For  $P_{R_3}$ :

$$G_r(\mathbf{v}_0, \lambda_0, s) = \mathbf{0}$$

$$G_w(\mathbf{v}_0, \lambda_0, s) = \mathbf{0}$$

$$L_r(\mathbf{v}_0, \lambda_0, s) = \{ \mathbf{lR}[\lambda_0], \mathbf{lR}[\lambda_0 + s] \}, \lambda_0 < s$$

$$L_w(\mathbf{v}_0, \lambda_0, s) = \{ \mathbf{lR}[\lambda_0] \}, \lambda_0 < s$$

$$L_r(\mathbf{v}_0, \lambda_0, s) = \mathbf{0}, \lambda_0 \ge s$$

$$L_w(\mathbf{v}_0, \lambda_0, s) = \mathbf{0}, \lambda_0 \ge s$$

• For  $P_{R_4}$ :

$$\begin{array}{rcl} G_r(v_0,0) &= & \emptyset \\ G_w(v_0,0) &= & \{ \mathbb{R}[v_0] \} \\ L_r(v_0,0) &= & \{ \mathbb{1R}[0] \} \\ L_w(v_0,0) &= & \emptyset \\ G_r(v_0,\lambda_0) &= & \emptyset, \lambda_0 \neq 0 \\ G_w(v_0,\lambda_0) &= & \emptyset, \lambda_0 \neq 0 \\ L_r(v_0,\lambda_0) &= & \emptyset, \lambda_0 \neq 0 \\ L_w(v_0,\lambda_0) &= & \emptyset, \lambda_0 \neq 0 \end{array}$$

### **Performance metrics**

- Bandwidth (*GB/s*) based on the execution time: the total size of all the memory operations divided by the total kernel execution time.
- Bandwidth (*GB/s*) based on total time: the total size of all the memory operations divided by the total time for all data transfers plus calculations.
- Parity: the transfer (upload and download) time divided by the kernel execution time.

**Validation mechanism** The results of the computation are verified against the results of the same operation performed on the host.

**Target-specific optimisations** The global and local work sizes are independent of the problem size and hardcoded in the benchmark code, which suggests their choice is target-specific.

Work items within a work-group with consecutive local ids access consecutive memory locations, which improves the memory system performance (caching, coalescing).

**Target-specific optimization opportunities** For a vector target architecture, vectorization of the additions in the device code would significantly improve performance.





# 2.2.4 Scan

This benchmark performs a scan (prefix sum) operation on a vector of floating point numbers.

### **High-level description**

The scan operation is defined as follows:

$$O_i = \sum_{j=0}^{i} I_j, 0 \le i < N$$
(2.10)

## Abstract data structures

- $\vec{I}$  is an input *N*-element vector.
- $\vec{O}$  is an output *N*-element vector.
- $\vec{R}$  is an intermediate *B*-element vector.
- $\vec{T}$  is an intermediate *B*-element vector.

**Computation** Range [0: (N-1)] is partitioned into *B* disjoint ranges by a monotonically increasing sequence  $\{i_k\}: 0 = i_0 < i_1 < ... < i_B = N$ :

$$[0:(N-1)] = \bigcup_{b=0}^{B-1} [i_b:(i_{b+1}-1)]$$

Scan (2.10) is performed in three steps:

• Segmented Reduction:

$$R_b = \sum_j I_j, \ 0 \le b < B \land i_b \le j < i_{b+1}$$

• Top Scan:

$$T_0 = 0, T_b = \sum_k R_k, \ 0 \le k < b < B$$

• Bottom Scan:

$$O_i = T_b + \sum_j I_j, \ 0 \le b < B \land i_b \le j \le i < i_{b+1}$$

Reductions and scans are well-known parallel primitives, which are assumed to be associative and commutative (*i.e.* additions can be performed in any order). Since the iteration space, dependences and memory accesses depend on how these primitives are actually implemented, we only analyze the low-level implementation.





### Low-level implementation details

The benchmark operates on floating point data of type **TYPE** and **VECTYPE**, where **TYPE** is **float** and **VECTYPE** is **float4** for single precision, **TYPE** is **double** and **VECTYPE** is **double4** for double precision.

**Device code** Each step of the high-level algorithm maps to one device kernel:

- Segmented Reduction: Algorithm 8.
- Top Scan: Algorithm 10 (calls Algorithm 9).
- Bottom Scan: Algorithm 11 (calls Algorithm 9).

### Host code See Algorithm 12.

### Host data structures

- TYPE hI[N]:  $I_x \to$  hI[x]
- TYPE hO[N]:  $O_x \to hO[x]$
- TYPE dI[N]:  $I_x \to dI[x]$
- TYPE  $dO[N]: O_x \to dO[x]$
- TYPE dr[b]:  $R_x \to dr[x], T_x \to dr[x]$

### **Device data structures**

- Segmented Reduction:
  - global const TYPE  $I[N]: I_X \to I[X]$
  - global TYPE R[B]:  $R_{\chi} \rightarrow$  R[x]
  - local TYPE  $lR[\Lambda_0]$ : intermediate results
- Top Scan:
  - global TYPE  $R[B]: R_x \to R[x], T_x \to R[x]$
  - local TYPE  $1S[2 \times \Lambda_0]$ : intermediate results.
- Bottom Scan:
  - global const TYPE  $I[N]: I_X \to I[X]$
  - global const TYPE  $T[B]: T_x \to T[x]$
  - global TYPE  $O[N]: O_X \to O[X]$
  - local TYPE  $ls[2 \times \Lambda_0]$ : intermediate results
  - local TYPE localSeed: intermediate result



```
Algorithm 8 The Scan (Segmented Reduction) device code abstraction.
Input: global const TYPE I[N]
Output: global TYPE R[B]
Local: local TYPE lR[\Lambda_0]
                                                                                       ▷ Partial reduction results.
  1: kernel REDUCE (I,R)
                                                                                                  \triangleright N_0 = B, v_0 = b
 2:
          blockSize \leftarrow ||N/4|/N_0| \times 4
                                                                                     ▷ Alignment of 4 elements.
                                          ▷ All work-groups but the last work on blockSize elements.
                                                                                              \triangleright blockStart = i_{V_0}
 3:
          \texttt{blockStart} \leftarrow v_0 \times \texttt{blockSize}
          if v_0 = N_0 - 1 then
                                                                ▷ Last work-group gets remaining elements.
 4:
               \texttt{blockEnd} \gets N
 5:
                                                                                                 \triangleright blockEnd = i_N
          else
 6:
 7:
               blockEnd \leftarrow blockStart + blockSize
                                                                                              \triangleright blockEnd = i_{v_0+1}
          end if
 8:
          \texttt{sum} \gets 0
 9:
          for i in [(blockStart + \lambda_0) : (blockEnd - 1) : \Lambda_0] do \triangleright Reduce multiple elements
10:
     per work-item.
               \texttt{sum} \leftarrow \texttt{sum} + \texttt{I}[\texttt{i}]
11:
          end for
12:
          lR[\lambda_0] \leftarrow sum
                                                                      ▷ Store partial result into local memory.
13:
          BARRIER(CLK_LOCAL_MEM_FENCE)
14:
15:
          for \mathbf{s} \leftarrow \Lambda_0/2 ensuremath to \mathbf{s} > 0 step \mathbf{s} \leftarrow |\mathbf{s}/2| do \triangleright Tree-like reduction of partial
     results.
               if \lambda_0 < s then
16:
                   lR[\lambda_0] \leftarrow lR[\lambda_0] + lR[\lambda_0 + s]
17:
18:
               end if
               BARRIER(CLK_LOCAL_MEM_FENCE)
19:
          end for
20:
          if \lambda_0 = 0 then
                                                                               ▷ Write result to global memory.
21:
               \mathbb{R}[v_0] \leftarrow \mathbb{lR}[0]
22:
          end if
23:
24: end kernel
```



function performs a parallel Kogge-Stone style scan<sup>a</sup> on these values. The *i*<sup>th</sup> element of the exclusive scan is returned to work-item with  $\lambda_0 = i$ . Input: TYPE v Local: local TYPE  $1S[2 \times \Lambda_0]$ 1: **function** SCANLOCALMEM(v)  $ls[\lambda_0] \leftarrow 0$ ▷ Initialize first half of the array with zeros. 2: 3:  $ls[\Lambda_0 + \lambda_0] \leftarrow v$ ▷ Initialize second half of the array with input values. BARRIER(CLK\_LOCAL\_MEM\_FENCE) 4:  $\triangleright$  Use 2 ×  $\Lambda_0$  local memory elements to eliminate control statements. for  $i \leftarrow 1$  ensuremath to  $i < \Lambda_0$  step  $i \leftarrow i \times 2$  do 5:  $\mathtt{t} \leftarrow \mathtt{lS}[\Lambda_0 + \lambda_0 - \mathtt{i}]$ 6: BARRIER(CLK\_LOCAL\_MEM\_FENCE) 7:  $\mathtt{lS}[\Lambda_0 + \lambda_0] \leftarrow \mathtt{lS}[\Lambda_0 + \lambda_0] + \mathtt{t}$ 8: BARRIER(CLK\_LOCAL\_MEM\_FENCE) 9: end for 10: return  $ls[\Lambda_0 + \lambda_0 - 1]$ ▷ Return exclusive scan elements. 11: 12: end function "http://en.wikipedia.org/wiki/Kogge-Stone\_adder""

 $\triangleright$  Each work-item  $\lambda_0$  in a work-group calls this function with its value  $\forall$ . The

Algorithm 9 The Scan (scanLocalMem) device code abstraction.

Algorithm 10 The Scan (top scan) device code abstraction.

```
Input: global TYPE R[B]
  1: kernel TOPSCAN (R)
            \mathtt{v} \leftarrow 0
 2:
            if \lambda_0 < B then
 3:
 4:
                   \mathbf{v} \leftarrow \mathtt{R}[\lambda_0]
  5:
            end if
            \texttt{r} \gets \texttt{SCANLOCALMEM}(\texttt{v})
 6:
            if \lambda_0 < B then
 7:
                   \mathtt{R}[\lambda_0] \leftarrow \mathtt{r}
 8:
            end if
 9:
10: end kernel
```



```
Algorithm 11 The Scan (bottom scan) device code abstraction.
Input: global const TYPE I[N]
Input: global const TYPE T[B]
Output: global TYPE O[N]
Local: local TYPE localSeed
  1: kernel BOTTOMSCAN (I,T,0)
            VECTYPE * in4 \leftarrow (VECTYPE*)I
  2:
            \texttt{VECTYPE} * \texttt{out4} \leftarrow (\texttt{VECTYPE} *) \texttt{0}
  3:
            blockSize \leftarrow ||N/4|/N_0|
                                                                                                    ▷ Alignment of 4 elements.
  4:
                                                 ▷ All work-groups but the last work on blockSize elements.
            \texttt{blockStart} \leftarrow v_0 	imes \texttt{blockSize}
  5:
                                                                                                              \triangleright blockStart = i_{v_0}
            if v_0 = N_0 - 1 then
                                                                           ▷ Last work-group gets remaining elements.
  6:
  7:
                 blockEnd \leftarrow N
                                                                                                                  \triangleright blockEnd = i_N
  8:
            else
                 \texttt{blockEnd} \gets \texttt{blockStart} + \texttt{blockSize}
                                                                                                              \triangleright blockEnd = i_{v_0+1}
  9:
            end if
 10:
            \texttt{seed} \leftarrow \texttt{T}[v_0]
11:
            for w in [\texttt{blockStart}:(\texttt{blockEnd}-1):\Lambda_0] do
12:
                 if \mathtt{w}+\lambda_0 < \mathtt{blockEnd} then
 13:
                       \mathtt{v} \gets \mathtt{in4}[\mathtt{w} + \lambda_0]
 14:
 15:
                 else
                       v \leftarrow \vec{0}
16:
                 end if
17:
                 \texttt{v.y} \leftarrow \texttt{v.y} + \texttt{v.x}
 18:
19:
                 v.z \leftarrow v.z + v.y
                 \texttt{v.w} \leftarrow \texttt{v.w} + \texttt{v.z}
20:
                 \texttt{r} \leftarrow \texttt{SCANLOCALMEM}(\texttt{v.w})
21:
22:
                 \texttt{v.x} \gets \texttt{v.x} + \texttt{r} + \texttt{seed}
23:
                 v.y \leftarrow v.y + r + \texttt{seed}
                 \texttt{v.z} \gets \texttt{v.z} + \texttt{r} + \texttt{seed}
24:
                 \texttt{v.w} \leftarrow \texttt{v.w} + \texttt{r} + \texttt{seed}
25:
                 if \mathtt{w} + \lambda_0 < \mathtt{blockEnd} then
26:
                       \mathtt{out4}[\mathtt{w} + \lambda_0] \leftarrow \mathtt{v}
27:
                 end if
28:
                 if \lambda_0 = \Lambda_0 - 1 then
29:
30:
                       \texttt{localSeed} \leftarrow \texttt{v.w}
31:
                 end if
                  BARRIER(CLK_LOCAL_MEM_FENCE)
32:
                  \texttt{seed} \leftarrow \texttt{localSeed}
33:
34:
            end for
35: end kernel
```


Algorithm 12 The Scan host code abstraction.			
Inp	put: const int N $\triangleright$ N	Number of input and output elements.	
1:	: $hI \leftarrow AllocateMemory(N)$		
2:	2: $hO \leftarrow ALLOCATEMEMORY(N)$		
3:	3: <b>for</b> $i$ <b>in</b> $0: (N-1)$ <b>do</b>		
4:	$h: hI[i] \leftarrow i\%3$	▷ To aid validation.	
5:	5: $hO[i] \leftarrow -1$		
6:	5: end for		
7:	7: globalWorkSize $\leftarrow 16384$		
8:	3: localWorkSize $\leftarrow 256$		
9:	): numWorkGroups $\leftarrow$ globalWorkSize/localWorkS	ize $\triangleright$ numWorkGroups = $B = 64$	
10:	): $dI \leftarrow AllocateBuffer(N, READ_WRITE)$		
11:	: $dO \leftarrow AllocateBuffer(N, READ_WRITE)$		
12:	$2: dR \leftarrow AllocateBuffer(numWorkGroups, READ_WF)$	RITE)	
13:	B: COPYTODEVICE(dI,hI, <b>True</b> ,∅)		
14:	$\texttt{: scanProgram} \gets \texttt{BUILDPROGRAM}(\texttt{"scan.cl"})$		
15:	5: reduceKernel $\leftarrow$ BUILDKERNEL(scanProgram,")	reduce")	
16:	5: topScanKernel $\leftarrow$ BUILDKERNEL(scanProgram,	'topScan'')	
17:	$: \texttt{bottomScanKernel} \gets \texttt{BUILDKERNEL}(\texttt{scanProgr})$	am,"bottomScan")	
18:	B: SETKERNELARGUMENTS(reduceKernel,dI,dR)		
19:	e: SETKERNELARGUMENTS(topScanKernel,dR)		
20:	: SETKERNELARGUMENTS(bottomScanKernel,dI,	dR, d0)	
21:	${\tt ENQUEUEKERNEL}({\tt reduceKernel,globalWorkSize,localWorkSize}, \emptyset)$		
22:	$E_{NQUEUEKERNEL}(\texttt{topScanKernel}, \texttt{localWorkSize}, \texttt{localWorkSize}, \emptyset)$		
23:	B: ENQUEUEKERNEL(bottomScanKernel,globalWork	$rkSize,localWorkSize,\emptyset)$	
24:	k: COPYTOHOST(d0,h0, <b>True</b> ,∅)		





**Input datasets** The input data is of sizes  $2^{20}, 2^{23}, 2^{25}$  and  $2^{26}$  bytes, which results in  $N \in \{2^{18}, 2^{21}, 2^{23}, 2^{24}\}$  for single-precision and  $N \in \{2^{17}, 2^{20}, 2^{22}, 2^{23}\}$  for double-precision. The *i*<sup>th</sup> element of I[] is initialized with the value *i*%3, resulting in a pattern of 0, 1, 2, 0, 1, 2, ...

#### **Partitioned iteration spaces**

Segmented Reduction, Algorithm 8:

• Lines 2–9:

$$P_{R_1} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256\}$$

• Lines 10–12:

$$P_{R_2} = \{ (v_0, \lambda_0, i) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256, \\ i = i_{v_0} + \lambda_0 + k \times \Lambda_0, 0 \le k < (i_{v_0+1} - i_{v_0} - \lambda_0) / \Lambda_0 \}$$

• Lines 13–14:

$$P_{R_3} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256\}$$

• Lines 15–20:

$$P_{R_4} = \{ (v_0, \lambda_0, s) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256, \\ s = \lfloor \Lambda_0 / 2^k \rfloor, k \in \{1, 2, \dots, \log_2 \Lambda_0\} \}$$

• Lines 21–23:

$$P_{R_5} = \{ (v_0, \lambda_0) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256 \}$$

scanLocalMem, Algorithm 9:5

• Lines 2–4:

$$P_{SL_1} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0, 0 \le \lambda_0 < \Lambda_0\}$$

• Lines 5–10:

$$\begin{array}{ll} P_{SL_2} &=& \{(v_0,\lambda_0,i): 0 \leq v_0 < N_0, 0 \leq \lambda_0 < \Lambda_0, \\ & i = 2^k, k \in \{0,1,2,\ldots, \log_2 \Lambda_0 - 1\}\} \end{array}$$

• Line 11:

$$P_{SL_3} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0, 0 \le \lambda_0 < \Lambda_0\}$$

Top Scan, Algorithm 10:

<sup>&</sup>lt;sup>5</sup>scanLocalMem can be viewed as a blackbox.





• Lines 2–9:

$$P_{TS_1} = \{(v_0, \lambda_0) : v_0 = 0, N_0 = 1, 0 \le \lambda_0 < \Lambda_0 = 256\}$$

Bottom Scan, Algorithm 11:

• Lines 2–12:

$$P_{BS_1} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256\}$$

• Lines 13–36:

$$P_{BS_2} = \{ (v_0, \lambda_0, w) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256, \\ w = i_{v_0} + k \times \Lambda_0, 0 \le k < (i_{v_0+1} - i_{v_0})/\Lambda_0 \}$$

# Device memory access mapping

Segmented Reduction:

• For  $P_{R_1}$ :

$$G_r(v_0, \lambda_0) = \emptyset$$
  

$$G_w(v_0, \lambda_0) = \emptyset$$
  

$$L_r(v_0, \lambda_0) = \emptyset$$
  

$$L_w(v_0, \lambda_0) = \emptyset$$

• For  $P_{R_2}$ :

$$G_r(v_0, \lambda_0, i) = \{ I[i] \}$$
  

$$G_w(v_0, \lambda_0, i) = \emptyset$$
  

$$L_r(v_0, \lambda_0, i) = \emptyset$$
  

$$L_w(v_0, \lambda_0, i) = \emptyset$$

• For  $P_{R_3}$ :

$$G_r(v_0, \lambda_0) = \emptyset$$
  

$$G_w(v_0, \lambda_0) = \emptyset$$
  

$$L_r(v_0, \lambda_0) = \emptyset$$
  

$$L_w(v_0, \lambda_0) = \{ lR[\lambda_0] \}$$

• For  $P_{R_4}$ :





• For  $P_{R_5}$ :

$$G_{r}(v_{0},0) = \emptyset$$

$$G_{w}(v_{0},0) = \{R[v_{0}]\}$$

$$L_{r}(v_{0},0) = \{IR[0]\}$$

$$L_{w}(v_{0},0) = \emptyset$$

$$G_{r}(v_{0},\lambda_{0}) = \emptyset, \lambda_{0} \neq 0$$

$$G_{w}(v_{0},\lambda_{0}) = \emptyset, \lambda_{0} \neq 0$$

$$L_{r}(v_{0},\lambda_{0}) = \emptyset, \lambda_{0} \neq 0$$

$$L_{w}(v_{0},\lambda_{0}) = \emptyset, \lambda_{0} \neq 0$$

scanLocalMem:

• For  $P_{SL_1}$ :

$$\begin{array}{lll} L_r(\mathbf{v}_0, \boldsymbol{\lambda}_0) &= & \boldsymbol{\emptyset} \\ L_w(\mathbf{v}_0, \boldsymbol{\lambda}_0) &= & \{ \texttt{lS}[\boldsymbol{\lambda}_0], \texttt{lS}[\boldsymbol{\lambda}_0 + \boldsymbol{\Lambda}_0] \} \end{array}$$

• For  $P_{SL_2}$ :

$$\begin{split} L_r(\mathbf{v}_0, \lambda_0, i) &= \{ \texttt{lS}[\lambda_0 + \Lambda_0] \}, \texttt{lS}[\lambda_0 + \Lambda_0 - i] \\ L_w(\mathbf{v}_0, \lambda_0, i) &= \{ \texttt{lS}[\lambda_0 + \Lambda_0] \} \end{split}$$

• For  $P_{SL_3}$ :

$$\begin{array}{lcl} L_r(\nu_0,\lambda_0) &=& \{ \mathtt{lS}[\lambda_0+\Lambda_0-1] \} \\ L_w(\nu_0,\lambda_0) &=& \emptyset \end{array}$$

Top Scan:

• For  $P_{TS_1}$ :

Bottom Scan:

• For  $P_{BS_1}$ :

```
G_r(\mathbf{v}_0, \lambda_0) = \{\mathsf{T}[\mathbf{v}_0]\}
G_w(\mathbf{v}_0, \lambda_0) = \emptyset
L_r(\mathbf{v}_0, \lambda_0) = \emptyset
L_w(\mathbf{v}_0, \lambda_0) = \emptyset
```



• For  $P_{BS_2}$ :

# **Performance metrics**

- Bandwidth (*GB/s*) based on execution time: the size of the input array divided by the total kernel execution time.
- Bandwidth (*GB/s*) based on total time: the size of the input array divided by the total time for data transfers plus kernel execution.
- Parity: the data transfer time divided by the kernel execution time.

**Validation mechanism** The results of the computation are verified against the results of the same operation performed on the host. What is interesting, however, is that the validation is done through an immediate comparison and not taking into account the relative error. This can result into false positives when detecting errors, as floating point addition is not associative and therefore changing the order of the operations can make the results vary between the host and device calculations.

# **Target-specific optimisations**

• An implicit assumption becomes obvious from the above description: the local work size used for all three kernels has to be greater than or equal to the number of work-groups used for the first and third kernel:

$$\Lambda_0 \ge N_0 = \Gamma_0 / \Lambda_0$$

The local work size also needs to be a power of two, in order for the tree-like reductions to work properly. Other than that, the kernels do not seem to rely on particular local work sizes, even though it is mentioned in comments in the code that these kernels would only work with the local work size of 256. In practice, with the global work size of 16384 that is hardcoded in the benchmark, we can reduce the local work size to 128 and still get correct results.





- Another point is that when the input array is partitioned between work-groups, care is taken that the blocks have a size that is a multiple of 4. For an input array with N elements, and a number of work-groups given by  $N_0 = \Gamma_0/\Lambda_0$ , each block will have a size of  $\lfloor \lfloor N/4 \rfloor / N_0 \rfloor \times 4$ , with any extra elements assigned to the last work-group. This could lead to significantly imbalanced loads if the input array is small and the number of work-groups is large. It also gives some insight to the selection of  $\Gamma_0 = 16384$  for the benchmark, with respect to the input array sizes used. To keep the number of elements per work-group and per-thread meaningful, the value of  $\Gamma_0$  will need to scale down when the problem size is reduced.
- The choice of the number four as an alignment factor is not clearly documented, and is probably another architecture specific choice.

### **Target-specific optimization opportunities**

- Vectorization could improve the performance for vector architectures.
- The choice of different work sizes can also be a significant architecture-dependent choice especially since the sizes are hardcoded in the benchmark code and do not depend on the input size.
- More importantly, launching exactly as many work-items as necessary (*B*) for the top scan, rather than re-using the same local work-size as for the other two kernels, the complicated assumptions regarding the global and local work sizes would be eliminated, and the control statements in topScan could be eliminated.





### 2.2.5 Sort

This benchmark performs a radix sort on unsigned integers.

#### **High-level description**

Given an input *N*-element integer vector  $\vec{I}$ , the result of sorting  $\vec{I}$  gives an *N*-element vector  $\vec{O}$  such that  $\vec{O}$  is a permutation of  $\vec{I}$  and the following holds:

$$O_i \le O_{i+1}, \ 0 \le i < N-1 \tag{2.11}$$

If the radix is r and the integer bitwidth is w, each element can be represented using  $N_d = \lfloor w/r \rfloor$  digits.

Least significant digit (LSD) radix sort performs  $N_d$  steps: at each step d, the elements are sorted according to the value of their  $d^{\text{th}}$  least significant digit.

#### Abstract data structures

- *I* is an  $N_d \times N$ -element matrix, where  $I_0$  is the original input vector.
- *O* is an  $N_d \times N$ -element matrix, where  $O_{N_d-1}$  is the final output vector.
- *C* is an intermediate  $B \times 2^r$ -element matrix.
- *S* is an intermediate  $B \times 2^r$ -element matrix.
- *X* is an intermediate *N*-element vector.

**Computation** Range [0: (N-1)] is partitioned into *B* disjoint ranges by a monotonically increasing sequence  $\{i_k\}: 0 = i_0 < i_1 < ... < i_B = N$ :

$$[0:(N-1)] = \bigcup_{b=0}^{B-1} [i_b:(i_{b+1}-1)]$$

We define for convenience:  $digit(a,d,r) \stackrel{\text{def}}{=} \lfloor a/2^{d \times r} \rfloor \% 2^r$ , where *a* is an integer element, *d* is a digit position and *r* is the radix.

For every digit  $d \in \{0, 1, ..., \lfloor w/r \rfloor\}$  the following three steps are performed:

• Segmented Reduction:

$$C_{b,v} = |\{I_{d,j}: \texttt{digit}(I_{d,j}, d, r) = v, i_b \le j < i_{b+1}\}|, \ 0 \le b < B \land 0 \le v < 2^r$$

• Top Scan:

$$S_{b,0} = \sum_{k} C_{k,0}, \ 0 \le k < b < B$$
  
$$S_{b,v} = \sum_{k} C_{k,v} + S_{B-1,v-1}, \ 0 \le k < b < B \land 0 \le v < 2^{r}$$





• Bottom Scan:

Finally, the buffers are swapped:

$$I_{d+1} = O_d, \ 0 \le d < N_d - 1$$

Reductions and scans are well-known parallel primitives, which are assumed to be associative and commutative (*i.e.* additions can be performed in any order). Since the iteration space, dependences and memory accesses depend on how these primitives are actually implemented, we only analyze the low-level implementation.

#### Low-level implementation details

**Device code** Each step of the high-level algorithm maps to one device kernel:

- Segmented Reduction: Algorithm 13.
- Top Scan: Algorithm 14 (calls Algorithm 9, which is reused from the Scan benchmark).
- Bottom Scan: Algorithm 15 (calls Algorithm 9).

#### Host code See Algorithm 16.

#### Host data structures

- uint hI[N]:  $I_{0,x} \rightarrow$  hI[x]
- uint hO[N]:  $O_{N_d-1,x} \to hO[x]$
- uint dI[N]:  $I_{2k,x} \rightarrow dI[x], O_{2k+1,x} \rightarrow dI[x]$
- uint dO[N]:  $O_{2k,x} \rightarrow dO[x], I_{2k+1,x} \rightarrow dO[x],$
- uint dC[B]:  $C_{b,v} \rightarrow dC[v * N_0 + b], S_{b,v} \rightarrow dC[v * N_0 + b]$

#### **Device data structures**

- Segmented Reduction:
  - global const uint  $I[N]: I_{d,x} \to I[x]$
  - global uint C[B \* 16]:  $C_{b,v} \rightarrow C[v * N_0 + b]$
  - local uint  $lC[\Lambda_0]$ : intermediate results
  - private int pC[16]: intermediate results
- Top Scan:

CARP-ARM-RP-001-v1.0



Algorithm 13 The Sort (Segmented Reduction) device code abstraction.

Input: global const uint I[N]  $\triangleright$  Perform this step for the  $d^{\text{th}}$  least significant digit. Input: int d **Output:** global uint C[B \* 16]  $\triangleright$  r = 4, so each digit can have  $2^4 = 16$  different values. ▷ Partial reduction results. **Local:** local uint  $lC[\Lambda_0]$ 1: kernel REDUCE (I,d,C)  $blockSize \leftarrow \lfloor \lfloor N/4 \rfloor / N_0 \rfloor \times 4$  $\triangleright N_0 = B, v_0 = b$  $\triangleright$  Alignment of 4 elements. 2: ▷ All work-groups but the last work on blockSize elements. 3:  $\texttt{blockStart} \gets \textit{v}_0 \times \texttt{blockSize}$  $\triangleright$  blockStart =  $i_{v_0}$ ▷ Last work-group gets remaining elements. 4: if  $v_0 = N_0 - 1$  then  $\texttt{blockEnd} \leftarrow N$  $\triangleright$  blockEnd =  $i_N$ 5: 6: else  $blockEnd \leftarrow blockStart + blockSize$ 7:  $\triangleright$  blockEnd =  $i_{v_0+1}$ 8: end if  $pC[0:15] \leftarrow \vec{0}$ ▷ Private histogram vector of size 16. 9: for i in  $[(blockStart + \lambda_0) : (blockEnd - 1) : \Lambda_0]$  do  $\triangleright$  Reduce multiple elements 10: per work-item.  $\texttt{digit} \leftarrow |\texttt{I}[\texttt{i}]/2^{d \times 4} |\%16$  $\triangleright r = 4.$ 11:  $pC[digit] \leftarrow pC[digit] + 1$ 12: end for 13: **for** v **in** [0 : 15] **do** ▷ Iterate over all possible digit values. 14:  $lC[\lambda_0] \leftarrow pC[v]$ ▷ Store partial result into local memory. 15: 16: BARRIER(CLK\_LOCAL\_MEM\_FENCE) 17: for  $s \leftarrow \Lambda_0/2$  ensuremath to  $s \le 0$  step  $s \leftarrow |s/2|$  do ▷ Tree-like reduction of partial results. if  $\lambda_0 < s$  then 18:  $lC[\lambda_0] \leftarrow lC[\lambda_0] + lC[\lambda_0 + s]$ 19: 20: end if BARRIER(CLK\_LOCAL\_MEM\_FENCE) 21: 22: end for if  $\lambda_0 = 0$  then 23:  $\triangleright$  Write result to global memory.  $C[v \times N_0 + v_0] \leftarrow 1C[0]$ 24: 25: end if end for 26: 27: end kernel



Algorithm 14 The Sort (Top Scan) device code abstraction.		
Inpu	<b>f: global</b> uint C[B * 16]	
Loca	l: int localSeed	
1: <b>k</b>	kernel TOPSCAN (C)	
2:	$\texttt{localSeed} \leftarrow 0$	
3:	BARRIER(CLK_LOCAL_MEM_FENCE)	
4:	<b>for</b> v <b>in</b> [0 : 15] <b>do</b>	▷ Iterate over all possible digit values.
5:	$c \gets 0$	
6:	if $\lambda_0 < B$ then	
7:	$\mathtt{c} \gets \mathtt{C}[B \times \mathtt{v} + \lambda_0]$	
8:	end if	
9:	$\texttt{r} \leftarrow \texttt{SCANLOCALMEM}(\texttt{c})$	
10:	if $\lambda_0 < B$ then	
11:	$\mathtt{C}[B imes \mathtt{v}+\lambda_0] \leftarrow \mathtt{r}+\mathtt{localSeed}$	
12:	end if	
13:	if $\lambda_0 = B - 1$ then	▷ Last thread to do useful work.
14:	$\texttt{localSeed} \gets \texttt{localSeed} + \texttt{r} + \texttt{c}$	
15:	end if	
16:	BARRIER(CLK_LOCAL_MEM_FENCE)	
17:	end for	
18: <b>e</b>	end kernel	

- global uint C[B \* 16]:  $C_{b,v} \to$  C[v \*  $N_0$ + b],  $S_{b,v} \to$  C[v \*  $N_0$ + b]

- local uint  $ls[2 \times \Lambda_0]$ : intermediate results (scanLocalMem)
- local int localSeed: intermediate result
- Bottom Scan:
  - global const uint  $I[N]: I_x \to I[x]$
  - global const uint S[B  $\star$  16]:  $S_{b,v} \to$  S[v  $\star$   $N_0 +$  b]
  - global uint  $O[N]: O_{d,x} \to O[x]$
  - local uint  $ls[2 \times \Lambda_0]$ : intermediate results (scanLocalMem)
  - local uint lSeed[16]: intermediate results
  - local uint lC[16]: intermediate results
  - private int pC[16]: intermediate results

**Input datasets** The input data is of sizes  $2^{20}$ ,  $2^{23}$ ,  $2^{25}$  and  $2^{26}$  bytes, which results in  $N \in \{2^{18}, 2^{21}, 2^{23}, 2^{24}\}$  for single-precision and  $N \in \{2^{17}, 2^{20}, 2^{22}, 2^{23}\}$  for double-precision.

The  $i^{\text{th}}$  element of I[] is initialized with the value i%16, so all input values are in the range [0, 16). This can affect the validation of the results of this benchmark (see below).

#### **Partitioned iteration spaces**

Segmented Reduction, Algorithm 13:



Algorithm 15 The Sort (Bottom Scan) device code abstraction.

```
Input: global const uint I[N]
                                                                                                     \triangleright The input array.
Input: global const uint S[B * 16]
                                                                                           \triangleright The results of top scan.
                                                        \triangleright Perform this step for the d^{\text{th}} least significant digit.
Input: int d
                                                                                                   \triangleright The output array.
Output: global uint O[N]
Local: local uint lSeed[16]
                                                                 ▷ Cached elements of S used by work-group.
                                           ▷ Histogram of digit values seen so far, local for work-group.
Local: local uint lC[16]
  1: kernel BOTTOMSCAN (I,S,d,O)
 2:
          uint4*in4 \leftarrow (uint4*)I
          blockSize \leftarrow ||N/4|/N_0|
 3:
                                                                                        \triangleright Alignment of 4 elements.
                                           ▷ All work-groups but the last work on blockSize elements.
          \texttt{blockStart} \leftarrow \textit{v}_0 	imes \texttt{blockSize}
 4:
                                                                                                 \triangleright blockStart = i_{v_0}
 5:
          if v_0 = N_0 - 1 then
                                                                  ▷ Last work-group gets remaining elements.
               \texttt{blockEnd} \leftarrow N
                                                                                                     \triangleright blockEnd = i_N
  6:
 7:
          else
               \texttt{blockEnd} \gets \texttt{blockStart} + \texttt{blockSize}
 8:
                                                                                                 \triangleright blockEnd = i_{v_0+1}
 9:
          end if
          if \lambda_0 < 16 then
10:
               \texttt{lSeed}[\lambda_0] \leftarrow \texttt{S}[\lambda_0 \times N_0 + v_0]
11:
                                                                                                                  \triangleright S_{v_0,\lambda_0}
               lC[\lambda_0] \leftarrow 0
12:
                                                             ▷ Initialize work-group local histogram to zeros.
          end if
13:
          BARRIER(CLK_LOCAL_MEM_FENCE)
14:
          pC[0:15] \leftarrow 0
15:
                                                                           ▷ Private histogram vector of size 16.
16:
          for w in [blockStart: (blockEnd -1): \Lambda_0] do
               if w + \lambda_0 < \texttt{blockEnd} then
17:
                    \texttt{val} \leftarrow \texttt{in4}[\texttt{w} + \lambda_0]
18:
                    \texttt{digit.x} \gets |\texttt{val.x}/2^{d \times 4}|\%16
                                                                                                 \triangleright digit(val.x, d, 4).
19:
                    digit.y \leftarrow |val.y/2^{d \times 4}|\%16
                                                                                                 \triangleright digit(val.y, d, 4).
20:
                    \texttt{digit.z} \leftarrow |\texttt{val.z}/2^{d \times 4}|\%16
                                                                                                 \triangleright digit(val.z, d, 4).
21:
                    digit.w \leftarrow |val.w/2^{d \times 4}|\%16
                                                                                                  \triangleright digit(val.w,d,4).
22:
                    pC[digit.x] \leftarrow pC[digit.x] + 1
23:
                    pC[digit.y] \leftarrow pC[digit.y] + 1
24:
                    pC[digit.z] \leftarrow pC[digit.z] + 1
25:
                    \texttt{pC}[\texttt{digit.w}] \gets \texttt{pC}[\texttt{digit.w}] + 1
26:
               end if
27:
               for v in [0 : 15] do
                                                                          ▷ Iterate over all possible digit values.
28:
                    pC[v] \leftarrow SCANLOCALMEM(pC[v])
29:
                    BARRIER(CLK_LOCAL_MEM_FENCE)
30:
31:
               end for
```



Algorithm 15 The Sort (Bottom Scan) device code abstraction (continued).		
32:	${f if}$ w $+ \lambda_0 < { t blockEnd}$ then	
33:	$\texttt{idx} \gets \texttt{pC}[\texttt{digit.x}] + \texttt{lSeed}[\texttt{digit.x}] + \texttt{lC}[\texttt{digit.x}]$	
34:	$\texttt{out}[\texttt{idx}] \leftarrow \texttt{v.x}$	
35:	$\texttt{pC}[\texttt{digit.x}] \gets \texttt{pC}[\texttt{digit.x}] + 1$	
36:	$\texttt{idx} \gets \texttt{pC}[\texttt{digit.y}] + \texttt{lSeed}[\texttt{digit.y}] + \texttt{lC}[\texttt{digit.y}]$	
37:	$\texttt{out}[\texttt{idx}] \gets \texttt{v.y}$	
38:	$\texttt{pC}[\texttt{digit.y}] \gets \texttt{pC}[\texttt{digit.y}] + 1$	
39:	$\texttt{idx} \gets \texttt{pC}[\texttt{digit.z}] + \texttt{lSeed}[\texttt{digit.z}] + \texttt{lC}[\texttt{digit.z}]$	
40:	$\texttt{out}[\texttt{idx}] \leftarrow \texttt{v.z}$	
41:	$\texttt{pC}[\texttt{digit.z}] \gets \texttt{pC}[\texttt{digit.z}] + 1$	
42:	$\texttt{idx} \gets \texttt{pC}[\texttt{digit.w}] + \texttt{lSeed}[\texttt{digit.w}] + \texttt{lC}[\texttt{digit.w}]$	
43:	$\texttt{out}[\texttt{idx}] \gets \texttt{v.w}$	
44:	$\texttt{pC}[\texttt{digit.w}] \gets \texttt{pC}[\texttt{digit.w}] + 1$	
45:	end if	
46:	$BARRIER(CLK\_LOCAL\_MEM\_FENCE) \qquad \qquad \triangleright Make sure lC is no longer needed.$	
47:	if $\lambda_0 = \Lambda_0 - 1$ then	
48:	<b>for</b> v <b>in</b> $[0:15]$ <b>do</b> $\triangleright$ Update work-group local histogram.	
49:	$\texttt{lC}[\texttt{v}] \gets \texttt{lC}[\texttt{v}] + \texttt{pC}[\texttt{v}]$	
50:	end for	
51:	end if	
52:	BARRIER(CLK_LOCAL_MEM_FENCE)	
53:	end for	
54:	end kernel	



Algorithm 16 The Sort host code abstraction. 1: procedure VALIDATERESULT(A[N]) for i in 0: (N-2) do 2: if A[i] > A[i+1] then 3: 4: ERROR() end if 5: end for 6: 7: end procedure Input: N ▷ Number of input and output elements. 8:  $hI \leftarrow ALLOCATEMEMORY(N)$ 9:  $h0 \leftarrow ALLOCATEMEMORY(N)$ 10: for i in 0: (N-1) do  $hI[i] \leftarrow i\%16$ ▷ Inadequate input range may leave errors undetected. 11:  $h0[i] \leftarrow -1$ 12: 13: end for 14: globalWorkSize  $\leftarrow 16384$ 15: localWorkSize  $\leftarrow 256$ 16: numWorkGroups ← globalWorkSize/localWorkSize  $\triangleright$  numWorkGroups = B 17:  $dI \leftarrow ALLOCATEBUFFER(N, READ_WRITE)$ 18:  $d0 \leftarrow ALLOCATEBUFFER(N, READ_WRITE)$ 19:  $dC \leftarrow AllocateBuffer(numWorkGroups \times 16, READ_WRITE)$ 20: COPYTODEVICE(dI,hI,**True**,∅) 21: sortProgram ← BUILDPROGRAM("sort.cl") 22: reduceKernel  $\leftarrow$  BUILDKERNEL(sortProgram, "reduce") 23: topScanKernel  $\leftarrow$  BUILDKERNEL(sortProgram, "topScan") 24: bottomScanKernel  $\leftarrow$  BUILDKERNEL(sortProgram, "bottomScan") 25: SETKERNELARGUMENTS(topScanKernel,dC) 26: **for** d **in** [0:7] do  $\triangleright$  Sorting 32-bit integers with r = 4 in 8 steps. if d%2 = 0 then 27:  $\triangleright$  d is even. 28: SETKERNELARGUMENTS(reduceKernel, dI, d, dC) SETKERNELARGUMENTS(bottomScanKernel, dI, dC, d, d0) 29: else  $\triangleright$  d is odd. 30: SETKERNELARGUMENTS(reduceKernel, d0, d, dC) 31: SETKERNELARGUMENTS(bottomScanKernel, d0, dC, d, dI) 32: end if 33: ENQUEUEKERNEL(reduceKernel,globalWorkSize,localWorkSize,Ø) 34: ENQUEUEKERNEL(topScanKernel,localWorkSize,localWorkSize,Ø) 35: ENQUEUEKERNEL(bottomScanKernel,globalWorkSize,localWorkSize, $\emptyset$ ) 36: 37: end for 38: COPYTOHOST(dI,h0,**True**,∅)



• Lines 2–9:

$$P_{R_1} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256\}$$

• Lines 10–13:

$$P_{R_2} = \{ (v_0, \lambda_0, i) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256, \\ i = i_{v_0} + \lambda_0 + k \times \Lambda_0, 0 \le k < (i_{v_0+1} - i_{v_0} - \lambda_0) / \Lambda_0 \}$$

• Lines 14–26:<sup>6</sup>

$$P_{R_3} = \{ (v_0, \lambda_0, v, s) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256, \\ 0 \le v < 16, s = \lfloor \Lambda_0 / 2^k \rfloor, k \in \{0, 1, \dots, \log_2 \Lambda_0 + 1\} \}$$

scanLocalMem, Algorithm 9:

• Lines 2–4:

$$P_{SL_1} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0, 0 \le \lambda_0 < \Lambda_0\}$$

• Lines 5–10:

$$egin{array}{rcl} P_{SL_2} &=& \{(m{v}_0, \lambda_0, i): 0 \leq m{v}_0 < N_0, 0 \leq \lambda_0 < \Lambda_0, \ &i=2^k, k \in \{0,1,\ldots, \log_2\Lambda_0-1\}\} \end{array}$$

• Line 11:

$$P_{SL_3} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0, 0 \le \lambda_0 < \Lambda_0\}$$

Top Scan, Algorithm 14:

• Lines 2–3:

$$P_{TS_1} = \{(v_0, \lambda_0) : v_0 = 0, N_0 = 1, 0 \le \lambda_0 < \Lambda_0 = 256\}$$

• Lines 4–17:

$$P_{TS_2} = \{(v_0, \lambda_0, v) : v_0 = 0, N_0 = 1, 0 \le \lambda_0 < \Lambda_0 = 256, 0 \le v < 16\}$$

Bottom Scan, Algorithm 15:

• Lines 2–15:

$$P_{BS_1} = \{(v_0, \lambda_0) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256\}$$

• Lines 16–53:<sup>7</sup>

$$P_{BS_2} = \{ (v_0, \lambda_0, w) : 0 \le v_0 < N_0 = 64, 0 \le \lambda_0 < \Lambda_0 = 256, \\ w = i_{v_0} + k \times \Lambda_0, 0 \le k < (i_{v_0+1} - i_{v_0}) / \Lambda_0 \}$$

<sup>6</sup>These lines contain an imperfect loop nest. For convenience, we analyze the inner loop by adding two extra iterations, one before and one after the actual iterations. The non-perfectly nested statements in lines 15–16 then belong to the first innermost loop iteration and the statements in lines 23–25 belong to the last inner loop iteration.

<sup>&</sup>lt;sup>7</sup>These lines contain an imperfect loop nest. For simplicity, we summarize the memory accesses for the inner loops of lines 28–31 and 48–50 for every iteration of the outer loop.





# Device memory access mapping

Segmented Reduction:

• For  $P_{R_1}$ :

$$G_r(v_0, \lambda_0) = \emptyset$$

$$G_w(v_0, \lambda_0) = \emptyset$$

$$L_r(v_0, \lambda_0) = \emptyset$$

$$L_w(v_0, \lambda_0) = \emptyset$$

$$P_r(v_0, \lambda_0) = \emptyset$$

$$P_w(v_0, \lambda_0) = \{pC[0:15]\}$$

• For  $P_{R_2}$ :

$$G_{r}(v_{0}, \lambda_{0}, i) = \{I[i]\}$$

$$G_{w}(v_{0}, \lambda_{0}, i) = \emptyset$$

$$L_{r}(v_{0}, \lambda_{0}, i) = \emptyset$$

$$L_{w}(v_{0}, \lambda_{0}, i) = \emptyset$$

$$P_{m}r(v_{0}, \lambda_{0}, i) = \{pC[0:15]\}$$

$$P_{m}w(v_{0}, \lambda_{0}, i) = \{pC[0:15]\}$$

• For  $P_{R_3}$ :

$$G_r(v_0, \lambda_0, v, \Lambda_0) = \emptyset$$

$$G_w(v_0, \lambda_0, v, \Lambda_0) = \emptyset$$

$$L_r(v_0, \lambda_0, v, \Lambda_0) = \emptyset$$

$$L_w(v_0, \lambda_0, v, \Lambda_0) = \{ lC[\lambda_0] \}$$

$$P_r(v_0, \lambda_0, v, \Lambda_0) = \{ pC[v] \}$$

$$P_w(v_0, \lambda_0, v, \Lambda_0) = \emptyset$$

CARP-ARM-RP-001-v1.0



$$\begin{array}{rcl} G_r(v_0, \lambda_0, v, 0) &=& \emptyset, \lambda_0 \neq 0 \\ G_w(v_0, 0, v, 0) &=& \{ \mathsf{C}[v \times N_0 + v_0] \} \\ G_w(v_0, \lambda_0, v, 0) &=& \emptyset, \lambda_0 \neq 0 \\ L_r(v_0, 0, v, 0) &=& \{ \mathsf{lC}[0] \} \\ L_r(v_0, \lambda_0, v, 0) &=& \emptyset \\ L_w(v_0, \lambda_0, v, 0) &=& \emptyset \\ P_r(v_0, \lambda_0, v, 0) &=& \emptyset \\ P_w(v_0, \lambda_0, v, 0) &=& \emptyset \end{array}$$

scanLocalMem:<sup>8</sup>

• For  $P_{SL_1}$ :

$$\begin{array}{lll} L_r(\mathbf{v}_0,\boldsymbol{\lambda}_0) &= & \boldsymbol{\emptyset} \\ L_w(\mathbf{v}_0,\boldsymbol{\lambda}_0) &= & \{ \texttt{lS}[\boldsymbol{\lambda}_0],\texttt{lS}[\boldsymbol{\lambda}_0+\boldsymbol{\Lambda}_0] \} \end{array}$$

• For  $P_{SL_2}$ :

 $\begin{array}{lll} L_r(\nu_0,\lambda_0,i) &=& \{ \texttt{lS}[\lambda_0+\Lambda_0] \},\texttt{lS}[\lambda_0+\Lambda_0-i] \\ L_w(\nu_0,\lambda_0,i) &=& \{ \texttt{lS}[\lambda_0+\Lambda_0] \} \end{array}$ 

• For  $P_{SL_3}$ :

$$\begin{array}{lll} L_r(\mathbf{v}_0, \boldsymbol{\lambda}_0) &=& \{ \mathtt{lS}[\boldsymbol{\lambda}_0 + \boldsymbol{\Lambda}_0 - 1] \} \\ L_w(\mathbf{v}_0, \boldsymbol{\lambda}_0) &=& \boldsymbol{\emptyset} \end{array}$$

Top Scan:

• For  $P_{TS_1}$ :

$$egin{array}{rll} G_r(m{v}_0,m{\lambda}_0)&=&m{0}\ G_w(m{v}_0,m{\lambda}_0)&=&m{0}\ L_r(m{v}_0,m{\lambda}_0)&=&m{0}\ L_w(m{v}_0,m{\lambda}_0)&=&\{ ext{localSeed}\} \end{array}$$

• For  $P_{TS_2}$ :

$$egin{aligned} G_r(m{v}_0,m{\lambda}_0,m{v}) &= egin{cases} \{ \mathsf{C}[B imesm{v}+m{\lambda}_0] \}, & m{\lambda}_0 < B \ m{ extsft{\emptyset}}, & m{\lambda}_0 \geq B \ m{ extsft{\emptyset}}, & m{ extsft{\lambda}}_0 \geq B \ m{ extsft{g}}, & m{ extsft{\lambda}}_0 < B \ m{ extsft{\emptyset}}, & m{ extsft{\lambda}}_0 < B \ m{ extsft{\emptyset}}, & m{ extsft{\lambda}}_0 \geq B \ m{ extsft{\xi}}, & m{ extsft{\xi}}_0 = m{ extsft{\xi}}, & m{ extsft{\xi}}_0 = m{ extsft{\xi}}, & m{ extsft{\xi}}_0 = m{ extsft{\xi}}_0, & m{ extsft{\xi}}$$

 $<sup>^8 \</sup>verb+scanLocalMem+$  can be viewed as a blackbox function.



$$\begin{split} L_r(v_0,\lambda_0,v) &= \begin{cases} \{ \texttt{lS}[\lambda_0+\Lambda_0], & \texttt{lS}[\lambda_0+\Lambda_0-1], & \texttt{ls}[\lambda_0+\Lambda_0-i], & \texttt{localSeed} \}, \\ & i = 2^k, k \in \{0,1,\dots, \log_2\Lambda_0-1\}, \lambda_0 < B \\ \{ \texttt{lS}[\lambda_0+\Lambda_0], & \texttt{lS}[\lambda_0+\Lambda_0-1], & \texttt{ls}[\lambda_0+\Lambda_0-i] \}, \\ & i = 2^k, k \in \{0,1,\dots, \log_2\Lambda_0-1\}, \lambda_0 \geq B \end{cases} \\ L_w(v_0,\lambda_0,v) &= \begin{cases} \{\texttt{localSeed},\texttt{lS}[\lambda_0],\texttt{lS}[\lambda_0+\Lambda_0]\}, & \lambda_0 = B-1 \\ \{\texttt{lS}[\lambda_0],\texttt{lS}[\lambda_0+\Lambda_0]\}, & \lambda_0 \neq B-1 \end{cases} \end{split}$$

Bottom Scan:

• For  $P_{BS_1}$ :

$$egin{aligned} G_r(v_0,\lambda_0) &= egin{cases} \{ \mathbf{S}[\lambda_0 imes N_0+v_0] \}, &\lambda_0 < 16\ \emptyset, &\lambda_0 \ge 16 \ \end{bmatrix} \ G_w(v_0,\lambda_0) &= \ \emptyset \ L_r(v_0,\lambda_0) &= \ \emptyset \ \end{bmatrix} \ L_w(v_0,\lambda_0) &= \ \{ \mathbf{ISeed}[\lambda_0],\mathbf{IC}[\lambda_0] \}, &\lambda_0 < 16\ \emptyset, &\lambda_0 \ge 16 \ \end{bmatrix} \ P_r(v_0,\lambda_0) &= \ \emptyset \ P_w(v_0,\lambda_0) &= \ \{ \mathbf{pC}[0:15] \} \end{aligned}$$

• For  $P_{BS_2}$ :

$$\begin{split} G_r(v_0,\lambda_0,w) &= \begin{cases} \{ \mathrm{in}[w+\lambda_0+i] \}, & i \in \{0,1,2,3\}, {}^9 w+\lambda_0 < i_{v_0+1} \\ \emptyset, & w+\lambda_0 \geq i_{v_0+1} \end{cases} \\ G_{mw}(v_0,\lambda_0,w) &= \begin{cases} \{ \mathrm{D}[0:(N-1)] \}, & w+\lambda_0 < i_{v_0+1} \\ \emptyset, & w+\lambda_0 \geq i_{v_0+1} \end{cases} \\ G_w(v_0,\lambda_0,w) &= & \emptyset, w+\lambda_0 \geq i_{v_0+1} \end{cases} \\ G_w(v_0,\lambda_0,w) &= & \begin{cases} \{ \mathrm{IS}[\lambda_0+\Lambda_0] \}, & \mathrm{IS}[\lambda_0+\Lambda_0-1], & \mathrm{IS}[\lambda_0+\Lambda_0-i] \}, \\ & i=2^k, k \in \{0,1,\ldots,\log_2\Lambda_0-1\}, & \lambda_0 \neq \Lambda_0-1 \\ \{ \mathrm{IS}[\lambda_0+\Lambda_0] \}, & \mathrm{IS}[\lambda_0+\Lambda_0-1], & \mathrm{IS}[\lambda_0+\Lambda_0-i], & \mathrm{IC}[0:15] \}, \\ & i=2^k, k \in \{0,1,\ldots,\log_2\Lambda_0-1\}, & \mathrm{IC}[0:15] \}, \end{cases} \\ L_w(v_0,\lambda_0,w) &= \begin{cases} \{ \mathrm{IS}[\mathrm{ed}[0:15] \}, & w+\lambda_0 < i_{v_0+1} \\ \emptyset, & w+\lambda_0 \geq i_{v_0+1} \\ \emptyset, & w+\lambda_0 \geq i_{v_0+1} \end{cases} \\ L_w(v_0,\lambda_0,w) &= \begin{cases} \{ \mathrm{IS}[\lambda_0], \mathrm{IS}[\lambda_0+\Lambda_0] \}, & \lambda_0 \neq \Lambda_0-1 \\ \{ \mathrm{IC}[0:15], \mathrm{IS}[\lambda_0], \mathrm{IS}[\lambda_0+\Lambda_0] \}, & \lambda_0 = \Lambda_0-1 \\ \end{bmatrix} \\ P_r(v_0,\lambda_0,w) &= \begin{cases} \mathrm{PC}[0:15] \} \\ P_w(v_0,\lambda_0,w) &= \\ \{ \mathrm{PC}[0:15] \} \end{cases} \end{split}$$





### **Performance metrics**

- Bandwidth (*GB/s*) based on execution time: the size of the input array divided by the total kernel execution time.
- Bandwidth (*GB/s*) based on total time: the size of the input array divided by the total time for data transfers plus kernel execution.
- Parity: the data transfer time divided by the kernel execution time.

**Validation mechanism** The results of the computation are checked to ensure that every element of the output array is less than or equal to the next element. The validation is not adequate, however, as the input array contains only values in the range [0, 16), which can be represented using a single 4-bit digit, which leaves errors such as the one shown in Algorithm 16, line 38, undetected.

**Target-specific optimisations and optimization opportunities** For target-specific optimisations and optimization opportunities refer to  $\S2.2.4$ , as this benchmark is closely related to Scan.





# 2.2.6 MD

### **High-level description**

This benchmark calculates the Lennard-Jones potential for particles in a cube, using a neighborlist algorithm.

## Abstract data structures

- *p* is an input *P* element array of coordinate vectors  $\vec{p} = (p_x, p_y, p_z)$ .
- *f* is an output *P* element array of forces  $\vec{f} = (f_x, f_y, f_z)$ .
- *n* is an input  $P \times N$  array of integer indices:  $n_{i,j}$  is the *j*<sup>th</sup> neighbor particle of the *i*<sup>th</sup> particle.

**Computation** For every particle, the force is calculated based on its coordinates and the coordinates of its N neighbor particles.

for i in 0 : (P-1) do  $\vec{f}_i \leftarrow 0$ for j in 0 : (N-1) do  $\vec{d} \leftarrow \vec{p}_i - \vec{p}_{n_{i,j}}$   $r \leftarrow ||\vec{d}||$ if r < R then  $c \leftarrow lj_1/r^{14} - lj_2/r^8$ end if end for end for end for

**Iteration space** The iteration space is two-dimensional:

$$I = \{(i, j) : 0 \le i < P, \ 0 \le j < N\}$$

**Dependences** The outer loop carries no dependences. The dependence carried by the inner loop is reduction.

### Memory access mapping

$$M_r(i,j) = \{ \vec{p}_i, n_{i,j}, p_{n_{i,j}} \}$$
  
 $M_w(i,j) = \{ \vec{f}_i \}$ 

#### Low-level implementation details

The benchmark operates on floating point data of type **TYPE** and **VECTYPE**, where **TYPE** is **float** and **VECTYPE** is **float4** for single precision, and **TYPE** is **double** and **VECTYPE** is **double4** for double precision.

CARP-ARM-RP-001-v1.0





**Device code** See Algorithm 17.

```
Algorithm 17 The MD device code abstraction.
Input: global VECTYPE positions[P]
Input: global int neighbors[P * N]
Output: global VECTYPE forces[P]
  1: kernel COMPUTE_LJ_FORCE (positions, neighbors, forces)
           ipos \leftarrow positions[\gamma]
 2:
           f \leftarrow \vec{0}
 3:
           for j in 0: (N-1) do
 4:
                                                                                                              ▷ Coalesced read.
                jidx \leftarrow neighbors[j \times P + \gamma]
 5:
                 jpos \leftarrow positions[jidx]
                                                                                                          \triangleright Uncoalesced read.
 6:
                dx \leftarrow ipos.x - jpos.x
 7:
                \mathtt{dy} \gets \mathtt{ipos.y} - \mathtt{jpos.y}
 8:
                \texttt{dz} \gets \texttt{ipos.z} - \texttt{jpos.z}
 9:
                r2 \leftarrow dx^2 + dy^2 + dz^2
10:
                R \leftarrow 4
                                                                                                   \triangleright R is the cutoff distance.
11:
                if r_2 < R^2 then
12:
                      \texttt{r2inv} \leftarrow 1.0/\texttt{r2}
13:
                      \texttt{r6inv} \gets \texttt{r2inv} \times \texttt{r2inv} \times \texttt{r2inv}
14:
                                                                                                 > lj1 and lj2 are constants.
                      c \leftarrow r2inv \times r6inv \times (lj1 \times r6inv - lj2)
15:
16:
                      \texttt{f.x} \gets \texttt{f.x} + \texttt{c} \times \texttt{dx}
                      \texttt{f.y} \gets \texttt{f.y} + \texttt{c} \times \texttt{dy}
17:
                      \texttt{f.z} \gets \texttt{f.z} + \texttt{c} \times \texttt{dz}
18:
19:
                end if
           end for
20:
21:
           forces[\gamma] \leftarrow f
22: end kernel
```

Host code See Algorithm 18.

#### Host data structures

- **VECTYPE** h\_positions[P]:  $\vec{p_o} \rightarrow h_{positions[o]}$
- **VECTYPE** h\_forces[P]:  $\vec{f}_o \rightarrow$  h\_forces[0]
- int h\_neighbors[P \* N]:  $n_{p,o} \rightarrow$  h\_neighbors[P \* o + p]

### **Device data structures**

- global VECTYPE positions[P]:  $\vec{p_o} \rightarrow \text{positions[o]}$
- global VECTYPE forces[P]:  $\vec{f}_o \rightarrow$  forces[0]
- global int neighbors [P \* N]:  $n_{p,o} \rightarrow \text{neighbors}$  [P \* o + p]



```
Algorithm 18 The MD host code abstraction.
Input: const int P
                                                                           \triangleright Number of particles.
                                                       ▷ Maximum number of nearest neighbors.
Input: const int N
                                                     ▷ Cutoff distance (for neighbor calculation).
Input: const int cutoff
                                                                 ▷ Constants for LJ computation.
Input: const int lj1, lj2
Input: const int epsilon
                                                                  ▷ Relative error for validation.
 1: function DIST2(pos1, pos2)
        return (pos1.x - pos2.x)^2 + (pos1.y - pos2.y)^2 + (pos1.z - pos2.z)^2
 2:
 3: end function
 4: procedure INSERTINORDER(currDist, currList, j, distIJ, N) \triangleright Insert a particle in
    the neighbor list if the particle is close enough.
        if distIJ < MAX(currDist) then > Compare with the furthest neighbor already in
 5:
    the list.
            for i in 0: (N-1) do
 6:
                if distIJ < currDist[i] then
 7:
 8:
                    INSERT(currDist,i,distIJ)
                                                                         \triangleright Insert in neighbor lists.
 9:
                    INSERT(currList,i,j)
                    RESIZE(currDist, N)
                                                                                       ⊳ Trim lists.
10:
                    RESIZE(currList, N)
11:
                end if
12:
            end for
13:
        end if
14:
15: end procedure
16: procedure BUILDNEIGHBORLIST(positions[P], neighbors[N][P])
                                                                                       ▷ Initialize
    neighbor list.<sup>a</sup>
        for i in 0: P - 1 do
17:
            currList \leftarrow \{-1, \ldots, -1\}
                                                                                    \triangleright N elements.
18:
            currDist \leftarrow \{\infty, \dots, \infty\}
19:
                                                                                    \triangleright N elements.
            for j in 0 : P − 1 do
20:
                if i \neq j then
21:
                    distIJ \leftarrow DIST2(positions[i], positions[j])
22:
                    INSERTINORDER(currDist, currList, j, distIJ, N)
23:
24:
                end if
                idx \leftarrow 0
25:
                for n in currList do
26:
                    neighbors[idx][i] \leftarrow n
27:
                    \mathtt{idx} \gets \mathtt{idx} + 1
28:
29:
                end for
            end for
30:
        end for
31:
32: end procedure
```

<sup>&</sup>lt;sup>*a*</sup>When passed to this function, neighbors is actually a flattened, one-dimensional array. Two-dimensional indexing is used for convenience.



Algorithm 18 The MD host code abstraction (continued).

```
33: procedure VALIDATEMD(positions[P], neighbors[P \times N], forces[P])
       expected_forces[] \leftarrow HOSTMD(positions[], neighbors[])
34:
                            ▷ Host code is nearly identical to device code, and hence omitted.
       for i in 0: P - 1 do
35:
           diff.x \leftarrow (forces[i].x - expected_forces[i].x)/forces[i].x
36:
           diff.y \leftarrow (forces[i].y - expected_forces[i].y)/forces[i].y
37:
           diff.z \leftarrow (forces[i].z - expected_forces[i].z)/forces[i].z
38:
           err \leftarrow \sqrt{diff.x^2} + \sqrt{diff.y^2} + \sqrt{diff.z^2}
39:
           if err > 3 \times epsilon then ERROR()
40 \cdot
41:
           end if
42.
       end for
43: end procedure
44: h_positions \leftarrow ALLOCATEMEMORY(P)
                                                                     ▷ Allocate host memory.
45: h_forces \leftarrow ALLOCATEMEMORY(P)
46: h_neighbors \leftarrow ALLOCATEMEMORY(P \times N)
47: d_positions \leftarrow ALLOCATEBUFFER(P, READ_WRITE)
                                                                    ▷ Should be READ_ONLY.
48: d_neighbors \leftarrow ALLOCATEBUFFER(P \times N, READ_WRITE)
                                                                    ▷ Should be READ_ONLY.
49: d_forces \leftarrow ALLOCATEBUFFER(P, READ_WRITE)
                                                                  ▷ Should be WRITE_ONLY.
                                ▷ Initialize particle positions with random floating-point data.
50: for i in 0 : P − 1 do
       h_{positions[i]} x \leftarrow RAND(0,20)
51:
       \texttt{h\_positions[i].y} \gets \texttt{RAND}(0,\!20)
52:
53:
       h_{positions[i].z \leftarrow RAND(0,20)
54: end for
55: BUILDNEIGHBORLIST(h_positions,h_neighbors)
                                                                     \triangleright Initialize neighbor list.
56: COPYTODEVICE(d_positions, h_positions, True, Ø)
57: COPYTODEVICE(d_neighbors,h_neighbors,True,∅)
58: MDProgram ← BUILDPROGRAM("MD.cl")
                                                                             ▷ Algorithm 17.
59: MDKernel ← BUILDKERNEL(MDProgram, "compute_lj_forces")
60: SETKERNELARGUMENTS(MDKernel, d_positions, d_neighbors, d_forces)
61: \Gamma \leftarrow P
62: \Lambda \leftarrow 128
                                                                 ▷ Hardcoded local work size.
63: kernelEvent \leftarrow ENQUEUEKERNEL(MDKernel,\Gamma, \Lambda, \emptyset)
64: WAITFOREVENTS({kernelEvent})
65: COPYTOHOST(d_forces, h_forces, True, \emptyset)
66: VALIDATEMD(h_positions,h_neighbors,h_forces)
```





**Input datasets** The benchmark operates on randomly generated data,<sup>10</sup> with N = 128 and  $P = \{12288, 24576, 36864, 73728\}$ . The positions are uniformly distributed in a  $20.0 \times 20.0 \times 20.0$  cube. The neighbor list is calculated based on the positions and the cutoff distance.

#### **Partitioned iteration space**

$$P = \{(\gamma, j) : 0 \le \gamma < P, 0 \le j < N\}$$
$$f : (i, j) \rightarrow (i, j)$$
$$f^{-1} : (\gamma, j) \rightarrow (\gamma, j)$$

**Device memory access mapping** 

$$\begin{array}{lll} G_r(\gamma) &=& \{\texttt{positions}[\gamma]\} \\ G_r(\gamma,j) &=& \{\texttt{neighbors}[j \times P + \gamma],\texttt{positions}[\texttt{neighbors}[j \times P + \gamma]]\} \\ G_w(\gamma) &=& \{\texttt{forces}[\gamma]\} \end{array}$$

#### **Performance metrics**

- Speed (*GFLOP/s*): the total number of floating point operations  $(8 \times P \times N + 13 \times N_{pairs})$ , where  $N_{pairs}$  is the number of pairs of particles within the cutoff distance) divided by the total time for all data transfers plus calculations.
- Speed (*GFLOP/s*) based on the kernel execution time only: the total number of floating point operations divided by the kernel execution time.
- Bandwidth (*GB/s*) based on the transfer time: the total size of all the memory operations divided by the total time for all data transfers.
- Bandwidth (*GB/s*) based on total time: the total size of all the memory operations divided by the total time for all data transfers plus calculations.
- Parity: the transfer (upload and download) time divided by the kernel execution time.

**Validation mechanism** The results of running the computation on the device are compared against running the same computation on the host (see Algorithm 18).

**Target-specific optimisations** The elements of neighbors [] are stored in column-major order to enable coalesced memory reads.

Target-specific optimization opportunities The device code is explicitly scalarized.

<sup>&</sup>lt;sup>10</sup>Using random data could potentially lead to non-reproducibility of results. In this case, however, the pseudorandom number generator is seeded with a fixed value.





#### 2.2.7 SGEMM

#### **High-level description**

The benchmark performs a matrix-matrix multiplication on floating-point data.

Abstract data structures A, B and C are matrices of size  $N \times N$ .

Computation

$$C_{i,j} = \alpha \times \sum_{k=0}^{N-1} (A_{i,k} \times B_{k,j}) + \beta \times C_{i,j}, 0 \le i < N, 0 \le j < N$$

**Iteration space and dependences** The iteration space is three-dimensional:  $I = \{(i, j, k) : 0 \le i < N, 0 \le j < N, 0 \le k < N\}.$ 

One of the many ways in which the above operation can be performed is the following, which we will assume for dependence analysis: Considering, for simplicity, the statement for

```
for i in 0: N-1 do
for j in 0: N-1 do
C_{i,j} \leftarrow \beta \times C_{i,j}
for k in 0: N-1 do
C_{i,j} \leftarrow C_{i,j} + \alpha \times A_{i,k} \times B_{k,j}
end for
end for
end for
```

initialization of  $C_{i,j}$  to happen at (i, j, -1), the dependences that arise are:

$$\delta_{I} = I: (i, j, k-1) \xrightarrow{i} I: (i, j, k),$$
  

$$I: (i, j, k-1) \xrightarrow{a} I: (i, j, k),$$
  

$$I: (i, j, k-1) \xrightarrow{o} I: (i, j, k)$$

These dependences are related to the order in which the additions are performed. The calculation of the matrix-matrix multiply requires essentially one dot product calculation per element of the output matrix, which is a reduction and can be implemented in many ways for parallelism. Obviously, some of the dependences can be removed by introducing new intermediate variables.

When describing the mapping of the above code to the device, we will be concerned with the multiplication operations. The mapping between the additions is not always meaningful, as there can be additions between different intermediate results.

#### Memory access mapping

$$M_{r}(i, j, -1) = \{C_{i,j}\}$$

$$M_{w}(i, j, -1) = \{C_{i,j}\}$$

$$M_{r}(i, j, k) = \{C_{i,j}, A_{i,k}, B_{k,j}\}$$

$$M_{w}(i, j, k) = \{C_{i,j}\}$$



Once more, the above mapping depends on the implementation of the matrix-matrix multiply, which will be significantly different for parallel implementations.

#### Low-level implementation details

The benchmark operates on floating point data of type **TYPE**, where **TYPE** is **float** for single precision and **double** for double precision.

**Device code** The matrix-matrix multiplication is performed using two different, independent kernels:

- sgemmNT(): N and T mean that the matrix A is non-transposed and the matrix B is transposed. See Algorithm 20.
- sgemmNN(): N and T mean that the matrices A and B are both non-transposed. See Algorithm 21.

The meaning of transposed and non-transposed is more complicated by the fact that the matrices are assumed to be stored in column-major order. So, non-transposed means a matrix is stored in column-major order and transposed means it is stored in row-major order.

Alg	porithm 19 SGEMM benchmark : Device code (SAXPY operation)
1:	<b>procedure</b> SAXPY(A,B[0:15],C[0:15])
2:	$\mathtt{C}[0] \gets \mathtt{C}[\mathtt{0}] + \mathtt{A} \times \mathtt{B}[0]$
3:	$\texttt{C}[1] \gets \texttt{C}[\texttt{0}] + \texttt{A} \times \texttt{B}[1]$
4:	$\mathtt{C}[2] \gets \mathtt{C}[\mathtt{0}] + \mathtt{A} \times \mathtt{B}[2]$
5:	$\mathtt{C[3]} \leftarrow \mathtt{C[0]} + \mathtt{A} \times \mathtt{B[3]}$
6:	$C[4] \leftarrow C[0] + A  imes B[4]$
7:	$C[5] \leftarrow C[0] + A  imes B[5]$
8:	$\texttt{C[6]} \gets \texttt{C[0]} + \texttt{A} \times \texttt{B[6]}$
9:	$C[7] \leftarrow C[0] + A  imes B[7]$
10:	$\texttt{C[8]} \leftarrow \texttt{C[0]} + \texttt{A} \times \texttt{B[8]}$
11:	$\texttt{C}[9] \gets \texttt{C}[\texttt{0}] + \texttt{A} \times \texttt{B}[9]$
12:	$\texttt{C}[10] \gets \texttt{C}[\texttt{0}] + \texttt{A} \times \texttt{B}[10]$
13:	$\texttt{C}[11] \leftarrow \texttt{C}[\texttt{0}] + \texttt{A} \times \texttt{B}[11]$
14:	$\texttt{C}[12] \leftarrow \texttt{C}[\texttt{0}] + \texttt{A} \times \texttt{B}[12]$
15:	$\texttt{C}[13] \leftarrow \texttt{C}[\texttt{0}] + \texttt{A} \times \texttt{B}[13]$
16:	$C[14] \leftarrow C[0] + A  imes B[14]$
17:	$\texttt{C}[15] \gets \texttt{C}[\texttt{0}] + \texttt{A} \times \texttt{B}[15]$
18:	end procedure

**Host code** Since the invocation of both kernels is straightforward and exactly the same, only the invocation of sgemmNT is shown in Algorithm 22.

**Host data structures** All matrices are assumed to be stored in column-major order. When invoking the sgemmNT () kernel:



Algorit	hm 20 SGEMM benchmark : Device code (sgemmNT ker	nel)
Input:	global const TYPE A[]	
Input:	int lda	
Input:	global const TYPE B[]	
Input:	int ldb	
Output	: global TYPE C[]	
Input:	int ldc	
Input:	int k	
Input:	TYPE alpha	
Input:	TYPE beta	
Local:	<pre>local TYPE bs[4][16]</pre>	
1: <b>ke</b> i	<b>mel</b> SGEMMNT (A,lda,B,ldb,C.ldc,k,alpha,beta)	
2:	$\mathtt{id} \leftarrow \lambda_0 + \lambda_1 \times 16$	
3:	$\texttt{Ap} \leftarrow \texttt{A} + v_0 \times \texttt{64} + \texttt{id}$	▷ Pointer arithmetic.
4:	$ ext{Bp} \leftarrow v_1  imes 16 + \lambda_0 + \lambda_1  imes  ext{ldb}$	
5:	$\texttt{Cp} \leftarrow v_0  imes 64 + \texttt{id} + v_1  imes 16  imes \texttt{ldc}$	
6:	$\texttt{counter} \leftarrow 0$	
7:	for $i$ in $[0:3]$ do	
8:	$\texttt{a}[\texttt{i}] \gets \texttt{Ap}[\texttt{i} \times \texttt{lda}]$	▷ private TYPE a[4]
9:	end for	
10:	$b \leftarrow Bp[0]$	▷ <b>private TYPE</b> b
11:	$\texttt{Ap} \leftarrow \texttt{Ap} + 4 \times \texttt{lda}$	
12:	$\texttt{Bp} \leftarrow \texttt{Bp} + 4  imes \texttt{ldb}$	
13:	$\texttt{counter} \leftarrow \texttt{counter} + 4 \times \texttt{ldb}$	
14:	for i in [0:15] do	
15:	$c[i] \leftarrow 0.0$	▷ private TYPE c[16]
16:	end for	



Algorithm 20 SGEMM benchmark : Device code (sgemmNT kernel) (continued)		
17:	repeat	
18:	<b>for i in</b> [0 : 3] <b>do</b>	
19:	$\texttt{as[i]} \gets \texttt{a[i]]}$	▷ private TYPE as[4]
20:	end for	
21:	$\mathtt{bs}[\lambda_0][\lambda_1] \gets \mathtt{b}$	
22:	BARRIER(CLK_LOCAL_MEM_FENCE)	
23:	for i in [0:3] do	
24:	$\texttt{a[i]} \gets \texttt{Ap[i \times \texttt{lda}]}$	
25:	end for	
26:	$\mathtt{b} \gets \mathtt{Bp}[0]$	
27:	SAXPY(as[0], bs[0], c)	
28:	SAXPY(as[1], bs[1], c)	
29:	SAXPY(as[2],bs[2],c)	
30:	SAXPY(as[3], bs[3], c)	
31:	$\texttt{Ap} \gets \texttt{Ap} + 4 \times \texttt{lda}$	
32:	$\texttt{Bp} \gets \texttt{Bp} + 4 \times \texttt{ldb}$	
33:	$\texttt{counter} \gets \texttt{counter} + 4 \times \texttt{ldb}$	
34:	BARRIER(CLK_LOCAL_MEM_FENCE)	▷ Synchronize so that bs[] won't get
	overwritten.	
35:	${f until} \ {f counter} \ge {f k}  imes {f ldb}$	
36:	$\mathtt{bs}[\lambda_0][\lambda_1] \gets \mathtt{b}$	
37:	BARRIER(CLK_LOCAL_MEM_FENCE)	
38:	SAXPY(as[0],bs[0],c)	
39:	SAXPY(as[1],bs[1],c)	
40:	SAXPY(as[2],bs[2],c)	
41:	SAXPY(as[3],bs[3],c)	
42:	for i in $[0:15]$ do	
43:	$\texttt{Cp}[0] \gets \texttt{alpha} \times \texttt{c}[\texttt{i}] + \texttt{beta} \times \texttt{Cp}[0]$	
44:	$\texttt{Cp} \gets \texttt{Cp} + \texttt{ldc}$	
45:	end for	
46:	end kernel	



```
Algorithm 21 SGEMM benchmark : Device code (sgemmNN kernel)
Input: global const TYPE A[]
Input: int lda
Input: global const TYPE B[]
Input: int ldb
Output: global TYPE C[]
Input: int ldc
Input: int k
Input: TYPE alpha
Input: TYPE beta
Local: local TYPE bs[16][17]
  1: kernel SGEMMNN (A,lda,B,ldb,C.ldc,k,alpha,beta)
 2:
           \texttt{id} \leftarrow \lambda_0 + \lambda_1 \times 16
           \texttt{Ap} \gets \texttt{A} + \textit{v}_0 \times 64 + \texttt{id}
                                                                                                      ▷ Pointer arithmetic.
 3:
           \texttt{Bp} \leftarrow \lambda_0 + (\lambda_1 + \nu_1 \times 16) \times \texttt{ldb}
 4:
           \texttt{Cp} \leftarrow v_0 \times 64 + \texttt{id} + v_1 \times 16 \times \texttt{ldc}
 5:
           for i in [0:15] do
 6:
                c[i] \leftarrow 0.0
 7:
                                                                                                ▷ private TYPE c[16]
           end for
 8:
 9:
           \texttt{counter} \gets 0
           repeat
10:
                for i in [0:3] do
11:
12:
                     a[i] \leftarrow Ap[i \times lda]
                                                                                                  ▷ private TYPE a[4]
                end for
13:
14:
                bs[\lambda_0][\lambda_1] \leftarrow Bp[0 \times ldb]
                \mathtt{bs}[\lambda_0][\lambda_1+4] \leftarrow \mathtt{Bp}[4 \times \mathtt{ldb}]
15:
                \mathtt{bs}[\lambda_0][\lambda_1+8] \leftarrow \mathtt{Bp}[8 \times \mathtt{ldb}]
16:
                bs[\lambda_0][\lambda_1 + 12] \leftarrow Bp[12 \times ldb]
17:
                BARRIER(CLK_LOCAL_MEM_FENCE)
18:
                \texttt{Ap} \gets \texttt{Ap} + 4 \times \texttt{lda}
19:
                SAXPY(as[0], bs[0], c)
20:
                a[0] \leftarrow Ap[0 \times lda]
21:
                SAXPY(as[1], bs[1], c)
22:
23:
                a[1] \leftarrow Ap[1 \times lda]
                SAXPY(as[2], bs[2], c)
24:
25:
                a[2] \leftarrow Ap[2 \times 1da]
                SAXPY(as[3], bs[3], c)
26:
                a[3] \leftarrow Ap[3 \times 1da]
27:
                \texttt{Ap} \gets \texttt{Ap} + 4 \times \texttt{lda}
28:
29:
                SAXPY(as[0], bs[4], c)
                \texttt{a}[0] \gets \texttt{Ap}[0 \times \texttt{lda}]
30:
                SAXPY(as[1], bs[5], c)
31:
                a[1] \leftarrow Ap[1 \times lda]
32:
                SAXPY(as[2], bs[6], c)
33:
                a[2] \leftarrow Ap[2 \times 1da]
34:
35:
                SAXPY(as[3], bs[7], c)
                a[3] \leftarrow Ap[3 \times 1da]
36:
```

CARP-ARM-RP-001-v1.0



Alg	orithm 21 SGEMM benchmark : Device code (sgemmNN kernel) (continued)	
37:	$\texttt{Ap} \leftarrow \texttt{Ap} + 4 \times \texttt{lda}$	
38:	SAXPY(as[0], bs[8], c)	
39:	$\texttt{a}[0] \gets \texttt{Ap}[0 \times \texttt{lda}]$	
40:	SAXPY(as[1], bs[9], c)	
41:	$\texttt{a}[1] \gets \texttt{Ap}[1 \times \texttt{lda}]$	
42:	SAXPY(as[2], bs[10], c)	
43:	$\texttt{a}[2] \gets \texttt{Ap}[2 \times \texttt{lda}]$	
44:	SAXPY(as[3], bs[11], c)	
45:	$\texttt{a[3]} \leftarrow \texttt{Ap[3 \times \texttt{lda}]}$	
46:	$\texttt{Ap} \gets \texttt{Ap} + 4 \times \texttt{lda}$	
47:	SAXPY(as[0], bs[12], c)	
48:	SAXPY(as[1], bs[13], c)	
49:	SAXPY(as[2],bs[14],c)	
50:	SAXPY(as[3], bs[15], c)	
51:	$\mathtt{Bp} \leftarrow \mathtt{Bp} + 16$	
52:	$\texttt{counter} \gets \texttt{counter} + 16$	
53:	BARRIER(CLK_LOCAL_MEM_FENCE) > Synchronize so that bs[] won't	get
	overwritten.	
54:	<b>until</b> counter $\geq$ k	
55:	for i in $[0:15]$ do	
56:	$\texttt{Cp}[0] \gets \texttt{alpha} \times \texttt{c[i]} + \texttt{beta} \times \texttt{Cp}[0]$	
57:	$\texttt{Cp} \leftarrow \texttt{Cp} + \texttt{ldc}$	
58:	end for	
59:	end kernel	



Algorithm 22 SGEMM : Host code abstraction			
Inp	put: N	▷ Size of matrices	
1:	: $A \leftarrow AllocateMemory(N \times N)$	▷ Allocate host memory.	
2:	: $B \leftarrow AllocateMemory(N \times N)$		
3:	: $C \leftarrow AllocateMemory(N \times N)$		
4:	: for i in $[0: N-1]$ do		
5:	: <b>for</b> j <b>in</b> $[0: N-1]$ <b>do</b>		
6:	$: \qquad \mathbf{A}[\mathbf{i} \times N + \mathbf{j}] \leftarrow 0.5 + 1.5 \times \text{Rand}(0.0, 1.0)$		
7:	$: \qquad B[i \times N + j] \leftarrow 0.5 + 1.5 \times RAND(0.0, 1.0)$		
8:	: $C[i \times N + j] \leftarrow 0.0$		
9:	end for		
10:	end for		
11:	: Agpu $\leftarrow$ AllocateBuffer( $N \times N$ , READ_WRITE)	▷ Allocate device buffers.	
12:	: Bgpu $\leftarrow$ AllocateBuffer( $N \times N$ , READ_WRITE)		
13:	: Cgpu $\leftarrow$ AllocateBuffer( $N \times N$ , READ_WRITE)		
14:	$: \texttt{writeEvent} \leftarrow \texttt{COPYTODEVICE}(\texttt{Agpu}[],\texttt{A}[],\texttt{True},\emptyset)$		
15:	$: \texttt{writeEvent} \leftarrow \texttt{COPYTODEVICE}(\texttt{Bgpu}[],\texttt{B}[],\texttt{True}, \emptyset)$		
16:	$: \texttt{writeEvent} \leftarrow \texttt{COPYTODEVICE}(\texttt{Cgpu}[],\texttt{C}[],\texttt{True}, \emptyset)$		
17:	$: \texttt{sgemmProgram} \gets \texttt{BUILDPROGRAM}(\texttt{"sgemm.cl"})$		
18:	$: \texttt{sgemmKernel} \gets \texttt{BUILDKERNEL}(\texttt{sgemmProgram}, \texttt{`sgemmNT''})$	)	
19:	: SETKERNELARGUMENTS(sgemmKernel,Agpu,N,Bgpu,N,Cgpu,N,N, $lpha,eta$ )		
20:	$\texttt{kernelEvent} \gets \texttt{ENQUEUEKERNEL}(\texttt{sgemmKernel}, (N/4, N/4), (16, 4), \emptyset)$		
21:	: WAITFOREVENTS({kernelEvent})		
22:	: COPYTOHOST(Cgpu[],C[], <b>True</b> ,Ø)		



- TYPE  $A[N \star N]: A_{x,y} \to A[y \star N + x],$
- TYPE B[N  $\star$  N]:  $B_{x,y} \to$  B[x  $\star$  N + y],
- TYPE  $C[N * N]: C_{x,y} \rightarrow C[y * N + x],$
- **TYPE** Agpu[N \* N]:  $A_{x,y} \rightarrow \text{Agpu[y * N + x]}$ ,
- **TYPE** Bgpu[N \* N]:  $B_{x,y} \rightarrow$  Bgpu[x \* N + y],
- **TYPE** Cgpu[N \* N]:  $C_{x,y} \rightarrow$  Cgpu[y \* N + x]

When invoking the  ${\tt sgemmNN}$  () kernel:

- TYPE A[N \* N]:  $A_{x,y} \rightarrow A[y * N + x]$ ,
- TYPE B[N \* N]:  $B_{x,y} \rightarrow B[y * N + x]$ ,
- TYPE  $C[N * N]: C_{x,y} \rightarrow C[y * N + x],$
- **TYPE** Agpu[N \* N]:  $A_{x,y} \rightarrow \text{Agpu[y * N + x]}$ ,
- **TYPE** Bgpu[N \* N]:  $B_{x,y} \rightarrow$  Bgpu[y \* N + x],
- **TYPE** Cgpu[N \* N]:  $C_{x,y} \rightarrow$  Cgpu[y \* N + x]

**Device data structures** For the sgemmNT () kernel:

- global const TYPE  $A[N * N]: A_{x,y} \rightarrow A[y * N + x],$
- global const TYPE  $B[N * N]: B_{x,y} \rightarrow B[x * N + y],$
- global TYPE  $C[N * N]: C_{x,y} \rightarrow C[y * N + x],$
- local TYPE bs[4][16]:  $B_{x,y} \rightarrow$  bs[x % 4][y % 16]

For the sgemmNN () kernel:

- global const TYPE A[N \* N]:  $A_{x,y} \rightarrow$  A[y \* N + x],
- global const TYPE  $B[N * N]: B_{x,y} \rightarrow B[y * N + x],$
- global TYPE  $C[N * N]: C_{x,y} \rightarrow C[y * N + x],$
- local TYPE bs[16][17]:  $B_{x,y} \rightarrow$  bs[x % 16][y % 16]

**Input datasets** There are four classes for the sizes of the datasets on which the benchmark operates: classes 1, 2, 3 and 4 work on arrays of dimensions  $256 \times 256$ ,  $1024 \times 1024$ ,  $2048 \times 2048$  and  $4096 \times 4096$  respectively for single-precision data, and on arrays of dimensions  $128 \times 128$ ,  $512 \times 512$ ,  $1024 \times 1024$  and  $2048 \times 2048$  respectively for double-precision data.

The computation is repeated for random single and double precision floating point inputs, initialized as shown in Algorithm 22.





**Partitioned iteration space** For the mapping between iteration spaces, we consider the mapping of the multiplication operations, as the additions are a reduction and can be performed in different ways.

$$P = \{ (v_0, v_1, \lambda_0, \lambda_1) : 0 \le v_0 < N_0 = N/64, 0 \le v_1 < N_1 = N/16, 0 \le \lambda_0 < \Lambda_0 = 16, 0 \le \lambda_1 < \Lambda_1 = 4 \}$$

The following function maps every iteration  $(i, j, k) \in I$  to iteration  $(v_0, v_1, \lambda_0, \lambda_1) \in P$ :

$$f: (i, j, k) \to (i/64, j/16, (i\%64)\%16, (i\%64)/16)$$

The following (one-to-many) relation gives the inverse mapping:

$$f^{-1}: (v_0, v_1, \lambda_0, \lambda_1) \to (64 \times v_0 + \lambda_0 + 16 \times \lambda_1, 16 \times v_1, p), p \in \{0, 1, 2, \dots, N-1\}$$

#### **Device memory access mapping**

• For the sgemmNT () kernel:

• For the sgemmNN () kernel:

**Performance metrics** The performance of each algorithm is evaluated using the following metrics:

- Speed (*GFLOP/s*) based on total execution time: the total number of floating point operations  $(2 \times N^3)$  divided by the total time for all data transfers plus calculations.
- Speed (*GFLOP/s*) based on kernel execution time only: the total number of floating point operations  $(2 \times N^3)$  divided by the kernel execution time.
- Parity: the transfer (upload and download) time divided by the kernel execution time.





Validation mechanism No validation mechanism is used for this benchmark.

### **Target-specific optimisations**

- The kernels use local memory, which will increase performance for systems with nonunified memory architectures.
- The local array in the sgemmNN() kernel is padded by one column, which is a common technique for NVIDIA cards to reduce bank conflicts in the shared (local) memory.

**Target-specific optimization opportunities** There are obvious vectorization opportunities missed (e.g. SAXPY operation).





### 2.2.8 Sparse matrix-vector multiplication

This benchmark multiplies a sparse matrix by a dense vector.

#### **High-level description**

$$\vec{y} = A \times \vec{x}$$

#### Abstract data structures

- A is an input  $N \times N$  sparse matrix, with Nz non-zero elements.
- $\vec{x}$  is an input *N*-element vector.
- $\vec{y}$  is an output *N*-element vector.

A sparse matrix is a matrix populated primarily with zeros. Since storing a sparse matrix as a two-dimensional array is not space-efficient, many specialized representations have been developed. This benchmark uses two formats, Compressed Sparse Row (CSR) and Ellpack-R, which will be described in separate subsections.

#### **Compressed Sparse Row (CSR)**

#### **Data structures**

- vals is an *Nz*-element array that stores the non-zero elements of *A*.
- cols an *Nz*-element array that stores the column indices of the non-zero elements.
- rows is an (N + 1)-element array: rows<sub>i</sub> is the index in vals where the first element of the  $i^{\text{th}}$  row of A is stored, and rows<sub>N</sub> = Nz. If row i only contains zeros, rows<sub>i</sub> = rows<sub>i+1</sub>.

#### Computation

$$y_i = \sum_{j=\mathrm{rows}_i}^{\mathrm{rows}_{i+1}-1} \mathrm{vals}_j \times x_{\mathrm{cols}_j}$$

#### **Iteration space**

$$I = \{(i, j) : 0 \le i < N, rows_i \le j < rows_{i+1}\}$$

**Dependences** The outer loop carries no dependences. The dependence carried by the inner loop is reduction.

#### Memory access mapping

$$M_r(i,j) = \{ \text{rows}_i, \text{rows}_{i+1}, \text{vals}_j, \text{cols}_j, x_{\text{cols}_j}, y_i \}$$
$$M_w(i,j) = \{ y_i \}$$

#### Ellpack-R

CARP-ARM-RP-001-v1.0





**Data structures** Let  $Nz_R$  be the maximum number of non-zero elements per row, calculated over all the rows of A.

- rl in an *N*-element array: rl<sub>i</sub> is the number of non-zero elements in the  $i^{th}$  row.
- vals is an  $N \times Nz_R$ -element array that stores the values of non-zero elements of A. The array is stored in column-major order: vals $[i + j \times N]$  is the  $j^{\text{th}}$  non-zero element of the  $i^{\text{th}}$  row, if  $j < rl_i$ , or zero if  $j \ge rl_i$ .
- cols is an *N* × *Nz*<sub>*R*</sub>-element array that stores the column indices corresponding to the non-zero elements stored in vals.

Computation

$$y_i = \sum_{j=0}^{\mathrm{rl}_i - 1} \mathrm{vals}_{i+j \times N} \times x_{\mathrm{cols}_{i+j \times N}}$$

**Iteration space** 

$$I = \{(i, j) : 0 \le i < N, 0 \le j < rl_i\}$$

**Dependences** The outer loop carries no dependences. The dependence carried by the inner loop is reduction.

#### Memory access mapping

$$M_r(i, j) = \{ rl_i, vals_{i+j \times N}, cols_{i+j \times N}, x_{cols_{i+j \times N}}, y_i \}$$
  
$$M_w(i, j) = \{ y_i \}$$

#### Low-level implementation details

The benchmark operates on floating point data of type **TYPE**, where **TYPE** is **float** for single precision and **double** for double precision.

#### **Compressed Sparse Row (CSR)**

**Device code** See Algorithm 23 and Algorithm 24 (based on [2, 16]).

Host code See Algorithm 25, Algorithm 26, Algorithm 27.

#### Host data structures

- **TYPE** hrows[N+1]: rows<sub>i</sub> → hrows[i]
- **TYPE** hvals[Nz]: vals<sub>i</sub> → hvals[i]
- **TYPE** hcols[Nz]:  $cols_i \rightarrow hcols[i]$
- **TYPE** hin[N]:  $x_i \rightarrow$  hin[i]
- **TYPE** hout [N]:  $y_i \rightarrow \text{hout}[i]$



```
Algorithm 23 The SpMV : spmvCsrScalar device code abstraction.
Input: global const TYPE vals[Nz]
Input: global const TYPE cols[Nz]
Input: global const TYPE rows[N+1]
Input: global const TYPE in [N]
Output: global TYPE out [N]
 1: kernel SPMVCSRSCALAR (vals, cols, rows, in, out)
 2:
        if \gamma_0 < N then
            \mathtt{t} \gets 0
 3:
            for j in rows[\gamma_0] : (rows[\gamma_0 + 1] - 1) do
 4:
 5:
                t \leftarrow t + vals[j] \times in[cols[j]]
            end for
 6:
            \operatorname{out}[\gamma_0] \leftarrow t
 7.
        end if
 8:
 9: end kernel
```

Device data structures For spmvCsrScalar:

- global const TYPE vals[Nz]: vals $i \rightarrow vals[i]$
- global const TYPE  $cols[Nz]: cols_i \rightarrow cols[i]$
- global const TYPE rows[N+1]: rowsi→ rows[i]
- global const TYPE  $in[N]: x_i \rightarrow in[i]$
- global TYPE out[N]: y<sub>i</sub> → out[i]

For spmvCsrVector:

- global const TYPE vals[Nz]: vals[i]
- global const TYPE  $cols[Nz]: cols_i \rightarrow cols[i]$
- global const TYPE rows[N+1]:  $rows_i \rightarrow rows[i]$
- global const TYPE  $in[N]: x_i \rightarrow in[i]$
- local TYPE ls[128]: intermediate results
- global TYPE out[N]:  $y_i \rightarrow \text{out[i]}$

**Input datasets** A file in the Matrix Market format<sup>11</sup> can be provided for initializing the sparse matrix *A*. The details of this initialization are omitted in the host code abstraction.

If no such file is provided, the matrix is initialized randomly as follows. The matrix sizes considered are  $N \times N$ , where  $N \in \{1024, 8192, 12288, 16384\}$ . The number of non-zero elements is  $Nz = N^2/100$ . vals and  $\vec{x}$  contain pseudo-random values between zero and a user-defined maximum value. rows and cols are also randomly initialized,<sup>12</sup> as shown in Algorithm 26.

<sup>&</sup>lt;sup>11</sup>http://math.nist.gov/MatrixMarket/

<sup>&</sup>lt;sup>12</sup>Using random data could potentially lead to non-reproducibility of results. While the pseudo-random number generator is seeded with the same fixed value on all platforms, the results may still vary depending on the platform-specific generator implementation.


```
Algorithm 24 The SpMV : spmvCsrVector device code abstraction.
Input: global const TYPE vals[Nz]
Input: global const TYPE cols [Nz]
Input: global const TYPE rows[N+1]
Input: global const TYPE in [N]
Output: global TYPE out [N]
Local: local TYPE ls[128]
                                                                                            \triangleright \Lambda_0 \leq 128 holds.
  1: kernel SPMVCSRVECTOR (vals, cols, rows, in, out)
     ▷ Despite the name of the kernel, the kernel does not contain vector operations. Each row is
     split in multiple "vectors", that are multiplied in parallel by V threads.
 2:
         id \leftarrow \lambda_0 \% V \triangleright V is the vector size. id is the position of the element that the work-item
     is calculating for each vector of the row.
 3:
         i \leftarrow v_0 \times |\Lambda_0/V| + \lambda_0/V
                                                                                         \triangleright Row to operate on.
         ls[\lambda_0] \leftarrow 0
 4:
         if i < N then
 5:
              \mathtt{t} \gets 0
 6:
              for k in 0: (\lceil (rows[i+1] - rows[i] - id)/V \rceil - 1) do \triangleright V work-items calculate
 7:
     the result for row i in parallel.
                   j \leftarrow k \times V + id
 8:
                  t \leftarrow t + vals[j] \times in[cols[j]]
 9:
10:
              end for
                                                          \triangleright Reduction of partial sums – assumes V = 32.
              ls[\lambda_0] \leftarrow t
11:
              BARRIER(CLK_LOCAL_MEM_FENCE)
12:
              if id < 16 then
13:
                   ls[\lambda_0] \leftarrow ls[\lambda_0] + ls[\lambda_0 + 16]
14:
              end if
15:
              BARRIER(CLK_LOCAL_MEM_FENCE)
16:
              if id < 8 then
17:
                   \texttt{ls}[\lambda_0] \leftarrow \texttt{ls}[\lambda_0] + \texttt{ls}[\lambda_0 + 8]
18:
              end if
19:
20:
              BARRIER(CLK_LOCAL_MEM_FENCE)
              if id < 4 then
21:
                   ls[\lambda_0] \leftarrow ls[\lambda_0] + ls[\lambda_0 + 4]
22:
              end if
23:
              BARRIER(CLK_LOCAL_MEM_FENCE)
24:
25:
              if id < 2 then
                   ls[\lambda_0] \leftarrow ls[\lambda_0] + ls[\lambda_0 + 2]
26:
              end if
27:
              BARRIER(CLK_LOCAL_MEM_FENCE)
28:
              if id < 1 then
29:
                  ls[\lambda_0] \leftarrow ls[\lambda_0] + ls[\lambda_0 + 1]
30:
              end if
31:
              BARRIER(CLK_LOCAL_MEM_FENCE)
32:
              if id = 0 then
33:
                   \operatorname{out}[i] \leftarrow \operatorname{ls}[\lambda_0]
34:
35:
              end if
         end if
36:
37: end kernel
                                                                                           10 November 2012
CARP-ARM-RP-001-v1.0
                                                       73
```



Algorithm 25 The SpMV : CSR host code abstraction.1: procedure VALIDATERESULT(cpuC[N], gpuC[N])2: for i in 0 : (N-1) do3: error  $\leftarrow |cpuC[i] - gpuC[i]|/cpuC[i]$ 4: if error  $\geq$  threshold then5: ERROR()6: end if

- 7: **end for**
- 8: end procedure

17: end procedure

```
9: procedure SPMVCPU(vals[Nz], cols[Nz], rows[N+1], in[N], out[N])

10: for i in 0: (N-1) do

11: t \leftarrow 0

12: for j in rows[i]: (rows[i+1]-1) do

13: t \leftarrow t + vals[j] \times in[cols[j]]

14: end for

15: out[i] \leftarrow t

16: end for
```

Algorithm 26 The SpMV : CSR host code abstraction.

```
1: procedure INITRANDOMMATRIX(cols[Nz], rows[N+1])
         \texttt{count} \gets 0
 2:
         P \leftarrow Nz/N^2
 3:
                             ▷ Probability that the current element gets assigned a non-zero value.
         \texttt{fillRemaining} \gets False
 4:
         for i in 0: (N-1) do
 5:
             rows[i] \leftarrow count
 6:
             for j in 0: (N-1) do
 7:
                 \texttt{Nrem} \gets N^2 - (\texttt{i} \times N + j)
                                                                                    ▷ Remaining elements.
 8:
                 if Nrem \leq Nz - \text{count} then
 9:
10:
                      fillRemaining \leftarrow True \triangleright All remaining elements need to be non-zero.
                 end if
11:
                 if (\text{count} < Nz \land \text{RAND}(0.0, 1.0) \le P) \lor \text{fillRemaining then}
12:
13:
                      cols[count] \leftarrow j
                      \texttt{count} \leftarrow \texttt{count} + 1
14:
                  end if
15:
16:
             end for
17:
         end for
         rows[N] \leftarrow Nz
18:
19: end procedure
20: procedure FILLRANDOM(m[Nz])
         for i in 0 : (Nz - 1) do
21:
             m[i] \leftarrow MAX_VAL \times RAND(0.0, 1.0)
22:
         end for
23:
24: end procedure
```

```
CARP-ARM-RP-001-v1.0
```



Algorithm 27 The SpMV : CSR host code abstraction.         Input: int N         > Problem size.		
2: $hin \leftarrow AllocateMemory(N)$		
3: hout $\leftarrow$ ALLOCATEMEMORY(N)		
4: refout $\leftarrow$ ALLOCATEMEMORY(N)		
5: hrows $\leftarrow$ AllocateMemory $(N+1)$		
6: $hcols \leftarrow AllocateMemory(Nz)$		
7: hvals $\leftarrow$ ALLOCATEMEMORY( $Nz$ )		
8: FILLRANDOM(hin)		
9: INITRANDOMMATRIX(hcols,hrows)		
10: FILLRANDOM(hvals)		
11: SPMVCPU(hvals,hcols,hrows,hin,refout)	▷ Obtaining reference output.	
12: din $\leftarrow$ AllocateBuffer( $N$ ,READ_WRITE)	▷ Should be READ_ONLY.	
13: dout $\leftarrow$ ALLOCATEBUFFER( $N$ , READ_WRITE)		
14: drows $\leftarrow$ AllocateBuffer( $N + 1$ , READ_WRITE)	▷ Should be READ_ONLY.	
15: $dvals \leftarrow AllocateBuffer(Nz, READ_WRITE)$	▷ Should be READ_ONLY.	
16: $dcols \leftarrow AllocateBuffer(Nz, READ_WRITE)$	▷ Should be READ_ONLY.	
17: COPYTODEVICE(din,hin, <b>True</b> ,∅)		
18: COPYTODEVICE(drows, hrows, True, Ø)		
<pre>19: COPYTODEVICE(dvals,hvals,True,∅)</pre>		
20: COPYTODEVICE(dcols,hcols, <b>True</b> ,∅)		
21: $spmvProgram \leftarrow BUILDPROGRAM("spmv.cl")$		
22: $csrScalarKernel \leftarrow BUILDKERNEL(spmvProgram,"$	${\tt spmvCsrScalar"})$	
23: SETKERNELARGUMENTS(csrScalarKernel,vals,co	ols,rows,in,out)	
24: csrVectorKernel $\leftarrow$ BUILDKERNEL(spmvProgram,"	spmvCsrVector")	
25: SETKERNELARGUMENTS(csrVectorKernel,vals,co	ols,rows,in,out)	
26: ENQUEUEKERNEL(csrScalarKernel, $(\Gamma_0 = N), (\Lambda_0 = N)$	= 128),Ø)	
27: COPYTOHOST(dout, hout, <b>True</b> , Ø)		
28: VALIDATERESULT(refout, hout)		
29: COPYTODEVICE(dout, hin, <b>True</b> , Ø)	▷ Clobber previous results.	
30: $\Gamma_0 \leftarrow N \times V$	$\triangleright \lor$ is the vector size ( $V = 32$ ).	
31: $\Lambda_{max} \leftarrow GetKernelInfo(csrVectorKernel,MAX_WOF)$	RK_SIZE)	
32: $\Lambda_0 \leftarrow \max(\mathbb{V}, \mathbb{V} \times \lfloor \min(128, \Lambda_{\max})/\mathbb{V} \rfloor)  \triangleright \text{ If } \Lambda_{\max} < \mathbb{V} =$	= 32, enqueueKernel will fail.	
33: ENQUEUEKERNEL(csrVectorKernel, $\Gamma_0, \Lambda_0, \emptyset$ )		
34: COPYTOHOST(dout, hout, <b>True</b> , ∅)		

35: VALIDATERESULT(refout, hout)





**Partitioned iteration space** For spmvCsrScalar:

$$P = \{(\gamma_0, j) : 0 \le \gamma_0 < N, \text{rows}[\gamma_0] \le j < \text{rows}[\gamma_0 + 1]\}$$

$$f: (i, j) \rightarrow (i, j)$$
  
 $f^{-1}: (\gamma_0, j) \rightarrow (\gamma_0, j)$ 

For spmvCsrVector:

$$\begin{split} P &= \{ (\mathbf{v}_0, \lambda_0, k) : 0 \leq \mathbf{v}_0 < N_0 = N/w, 0 \leq k < \Lambda_0 = w \times V, w \in \mathbb{N}, \\ &0 \leq k < \lceil (\operatorname{rows}[r+1] - \operatorname{rows}[r] - \lambda_0 \% \mathbb{V}) / \mathbb{V} \rceil, \\ &r = \mathbf{v}_0 \times (\Lambda_0 / \mathbb{V}) + \lfloor \lambda_0 / \mathbb{V} \rfloor \} \end{split}$$

$$\begin{aligned} f:(i,j) &\to \quad \lfloor (i \times \mathbb{V})/\Lambda_0 \rfloor, \mathbb{V} \times (i \mathscr{B}(\Lambda_0/\mathbb{V})) + j \mathscr{B} \mathbb{V}, \lfloor j/\mathbb{V} \rfloor ) \\ f^{-1}:(\nu_0,\lambda_0,k) &\to \quad (\nu_0 \times (\Lambda_0/\mathbb{V}) + \lfloor \lambda_0/\mathbb{V} \rfloor, k \times \mathbb{V} + \lambda_0 \mathscr{B} \mathbb{V}) \end{aligned}$$

We do not describe the imperfectly nested statements as separate iteration spaces.

Device memory access mapping For spmvCsrScalar:

• Before the main loop body:

$$G_r(\gamma_0) = \{ \mathtt{rows}[\gamma_0], \mathtt{rows}[\gamma_0+1] \}$$

• Main loop body:

$$G_r(\gamma_0, j) = \{ vals[j], cols[j], in[cols[j]] \}$$

• After the main loop body:

$$G_w(\gamma_0) = \{ \mathtt{out}[\gamma_0] \}$$

For spmvCsrVector:

• Before the main loop body:

$$\begin{array}{lll} G_r(\mathbf{v}_0, \mathbf{\lambda}_0) &=& \{\texttt{rows}[r], \texttt{rows}[r+1]\}, \\ && r = \mathbf{v}_0 \times (\mathbf{\Lambda}_0/\mathbf{V}) + \lfloor \mathbf{\lambda}_0/\mathbf{V} \rfloor, r < N \\ L_w(\mathbf{v}_0, \mathbf{\lambda}_0) &=& \{\texttt{ls}[\mathbf{\lambda}_0]\}, \end{array}$$

• Main loop body:

$$\begin{array}{ll} G_r(\mathbf{v}_0, \lambda_0, k) &=& \{\texttt{vals}[j], \texttt{cols}[j], \texttt{in}[\texttt{cols}[j]]\}, j = k \times \mathtt{V} + \lambda_0 \% \mathtt{V} \\ & r = \mathbf{v}_0 \times (\Lambda_0/\mathtt{V}) + \lfloor \lambda_0/\mathtt{V} \rfloor, r < N \end{array}$$



• After the main loop body:

$$L_w(\mathbf{v}_0, \mathbf{\lambda}_0) \ = \ \{\texttt{ls}[\mathbf{\lambda}_0]\},$$

$$L_r(\nu_0,\lambda_0) = \begin{cases} \{ ls[\lambda_0], ls[\lambda_0+16], ls[\lambda_0+8], ls[\lambda_0+4], ls[\lambda_0+2], ls[\lambda_0+1] \}, & \lambda_0 = 0\\ \{ ls[\lambda_0], ls[\lambda_0+16], ls[\lambda_0+8], ls[\lambda_0+4], ls[\lambda_0+2] \}, & \lambda_0 = 1\\ \{ ls[\lambda_0], ls[\lambda_0+16], ls[\lambda_0+8], ls[\lambda_0+4] \}, & 2 \le \lambda_0 < 4\\ \{ ls[\lambda_0], ls[\lambda_0+16], ls[\lambda_0+8] \}, & 4 \le \lambda_0 < 8\\ \{ ls[\lambda_0], ls[\lambda_0+16] \}, & 8 \le \lambda_0 < 16 \end{cases}$$

$$G_w(v_0, \lambda_0) = \{ \mathtt{out}[r] \}, r = v_0 imes (\Lambda_0 / \mathtt{V}) + \lfloor \lambda_0 / \mathtt{V} 
floor, r < N$$

## **Performance metrics**

- Speed (*GFLOP/s*) based on calculations: the total number of floating point operations  $(2 \times Nz)$  divided by the time for calculations.
- Speed (*GFLOP/s*) based on total time: the total number of floating point operations  $(2 \times Nz)$  divided by the total time for all data transfers plus calculations.

**Validation mechanism** The results of the computation are verified against the results of the same sparse-matrix vector multiplication performed on the host, using the CSR format.

## Target-specific optimisations

- When the device supports images,  $\vec{x}$  is stored using an image type. This enables the use of texture memory on NVIDIA graphics cards, which will have different effects on performance for different architectures, even for different generations of NVIDIA graphics cards.
- The two kernels are run two times: for the first run, the arrays rows, vals and cols are initialized as shown in the host code abstraction; for the second run, vals and cols are padded so that each row of values is a multiple of a padding factor. rows is adjusted accordingly. (Details of this padding are omitted in the host code representation).
- In spmvCsrVector, V work-items within a work-group with consecutive local ids access consecutive elements of the arrays vals and cols, which improves the memory system performance (caching, coalescing).

**Target-specific optimization opportunities** Vectorization could be applied to the device code, however the indirect accesses to  $\vec{x}$  cannot be vectorized.

Sparse-matrix vector multiplication is a well-studied problem with various different formats and corresponding algorithms. Choosing the right data representation to implement is a target-specific choice.





**Ellpack-R** 

**Device code** See Algorithm 28 (based on [9]).

```
Algorithm 28 The SpMV : spmvEllpackR device code abstraction.
Input: global const TYPE vals [N \times NzR]
Input: global const TYPE cols[N \times NzR]
Input: global const TYPE rl[N]
Input: global const TYPE in[N]
Output: global TYPE out [N]
 1: kernel SPMVELLPACKR (vals, cols, rl, in, out)
 2:
         if \gamma_0 < N then
             \texttt{t} \gets 0
 3:
             for j in 0 : (rl[\gamma_0] - 1) do
 4:
                 \texttt{idx} \leftarrow \gamma_0 + \texttt{j} \times N
                                                                              ▷ Column-major storage.
 5:
                 t \leftarrow t + vals[idx] \times in[cols[idx]]
 6:
             end for
 7:
             \operatorname{out}[\gamma_0] \leftarrow t
 8:
 9:
         end if
10: end kernel
```

Host code See Algorithm 25, Algorithm 26, Algorithm 29.

### Host data structures

- TYPE hrl[N]:  $rl_i \rightarrow hrl[i]$
- **TYPE** hvals\_cm[N \* NzR]: vals\_ $i \rightarrow hvals_cm[i]$
- **TYPE** hcols\_cm[N ★ NzR]: cols<sub>i</sub> → hcols\_cm[i]
- **TYPE** hin[N]:  $x_i \rightarrow$  hin[i]
- **TYPE** hout[N]:  $y_i \rightarrow$  hout[i]

#### **Device data structures**

- global const TYPE vals[N \* NzR]: vals $i \rightarrow vals[i]$
- global const TYPE cols[N \* NzR]:  $cols_i \rightarrow cols[i]$
- global const TYPE  $rl[N]: rl_i \rightarrow rl[i]$
- global const TYPE  $in[N]: x_i \rightarrow in[i]$
- global TYPE out[N]: y<sub>i</sub> → out[i]

**Input datasets** The input datasets used are the same as for the CSR implementation.



Alg	orithm 29 The SpMV : Ellpack-R host code abstraction.	
1:	$Nz \leftarrow N^2/100$	▷ 1% of elements is non-zero.
2:	$hin \leftarrow ALLOCATEMEMORY(N)$	
3:	$\texttt{hout} \leftarrow \texttt{ALLOCATEMEMORY}(N)$	
4:	$\texttt{refout} \leftarrow \texttt{ALLOCATEMEMORY}(N)$	
5:	$hrows \leftarrow ALLOCATEMEMORY(N+1)$	
6:	$\texttt{hcols} \leftarrow \texttt{ALLOCATEMEMORY}(Nz)$	
7:	$hvals \leftarrow ALLOCATEMEMORY(Nz)$	
8:	FILLRANDOM(hin)	
9:	INITRANDOMMATRIX(hcols,hrows)	
10:	FILLRANDOM(hvals)	
11:	${\tt SPMvCPU}({\tt hvals}, {\tt hcols}, {\tt hrows}, {\tt hin}, {\tt refout})$	▷ Obtaining reference output.
12:	$\texttt{hrl} \leftarrow \texttt{ALLOCATEMEMORY}(N)$	
13:	$NzR \leftarrow 0$	
14:	for i in $0: (N-1)$ do	
15:	$\texttt{hrl}[\texttt{i}] \gets \texttt{hrows}[\texttt{i}+\texttt{1}] - \texttt{hrows}[\texttt{i}]$	
16:	if $hrl[i] > NzR$ then	
17:	$NzR \leftarrow \texttt{hrl}[\texttt{i}]$	
18:	end if	
19:	end for	
20:	$\texttt{hvals\_cm} \leftarrow \texttt{ALLOCATEMEMORY}(N \times NzR)$	
21:	$\texttt{hcols\_cm} \leftarrow \texttt{ALLOCATEMEMORY}(N \times NzR)$	
22:	$CONVERTTOCOLUMN MAJOR (\texttt{hvals}, \texttt{hcols}, \texttt{hrows}, \texttt{hvals}, \texttt{hcols}, \texttt{hcols}, \texttt{hrows}, \texttt{hvals}, \texttt{hcols}, \texttt{hrows}, \texttt{hvals}, \texttt{hcols}, \texttt{hrows}, \texttt{hvals}, \texttt{hcols}, \texttt{hcols}, \texttt{hrows}, \texttt{hvals}, \texttt{hcols}, \texttt{hrows}, \texttt{hcols}, \texttt{hrows}, \texttt{hvals}, \texttt{hcols}, \texttt{hrows}, \texttt{hvals}, \texttt{hcols}, \texttt{hcols}, \texttt{hrows}, \texttt{hvals}, \texttt{hcols}, \texttt$	Ls_cm,hcols_cm,hrl) ▷
	Details omitted.	
23:	$\texttt{din} \leftarrow \texttt{ALLOCATEBUFFER}(N, \texttt{READ\_WRITE})$	▷ Should be READ_ONLY.
24:	$\texttt{dout} \leftarrow \texttt{ALLOCATEBUFFER}(N, \texttt{READ\_WRITE})$	
25:	$drl \leftarrow AllocateBuffer(N, READ_WRITE)$	▷ Should be READ_ONLY.
26:	$dvals\_cm \leftarrow AllocateBuffer(N \times NzR, READ\_WRITE)$	▷ Should be READ_ONLY.
27:	$dcols\_cm \leftarrow AllocateBuffer(N \times NzR, READ\_WRITE)$	▷ Should be READ_ONLY.
28:	COPYTODEVICE(din,hin, <b>True</b> ,∅)	
29:	COPYTODEVICE(drl,hrl,True,∅)	
30:	COPYTODEVICE(dvals_cm, hvals_cm, <b>True</b> , Ø)	
31:	COPYTODEVICE(dcols_cm, hcols_cm, <b>True</b> , ∅)	
32:	<pre>spmvProgram</pre>	
33:	ellpackrKernel ← BUILDKERNEL(spmvProgram,"spm	vEllpackR")
34:	SETKERNELARGUMENTS(ellpackrKernel,vals,cols,	rows,in,out)
35:	ENQUEUEKERNEL(ellpackrKernel, $(\Gamma_0 = N), (\Lambda_0 = 123)$	8),Ø)
36:	COPYTOHOST(dout, hout, <b>True</b> , Ø)	
37:	VALIDATERESULT(refout, hout)	





## Partitioned iteration space

$$P = \{(\gamma_0, j) : 0 \le \gamma_0 < N, 0 \le j < rl[\gamma_0]\}$$
$$f: (i, j) \rightarrow (i, j)$$
$$f^{-1}: (\gamma_0, j) \rightarrow (\gamma_0, j)$$

We do not describe the imperfectly nested statements as separate iteration spaces.

## **Device memory access mapping**

• Before the main loop body:

$$G_r(\gamma_0) = \{ \mathtt{rl}[\gamma_0] \}$$

• Main loop body:

$$G_r(\gamma_0, j) = \{ \texttt{vals}[\gamma_0 + j imes N], \texttt{cols}[\gamma_0 + j imes N], \texttt{in}[\texttt{cols}[\gamma_0 + j imes N]] \}$$

• After the main loop body:

$$G_w(\gamma_0) = \{ \mathtt{out}[\gamma_0] \}$$

**Performance metrics** The performance metrics used are the same as for the CSR implementation.

**Validation mechanism** The validation mechanism used is the same as for the CSR implementation.

**Target-specific optimisations** In case the device supports images,  $\vec{x}$  is stored using an image type. This enables the use of texture memory on NVIDIA graphics cards, which will have different effects on performance for different architectures, even for different generations of NVIDIA graphics cards.

Storing the vals and cols arrays in column-major format causes work items with consecutive global ids to access consecutive memory locations, which improves the memory system performance (caching, coalescing).

**Target-specific optimization opportunities** Refer to the target-specific optimization opportunities section for the CSR implementation.



## 2.3 Rodinia

The Rodinia 2.1 benchmark suite[5, 6] from Virginia university.

## 2.3.1 Back Propagation

The benchmark uses the Back Propagation machine-learning algorithm to train the weights of edges in an artificial layered neural network [15].

## **High-level description**

Abstract data structures The benchmark operates over the artificial layered neural network  $P = (I, H, O, W^I, W^H)$ :

- *I* is an  $(N_I + 1)$  element layer of the input nodes.
- *H* is an  $(N_H + 1)$  element layer of the hidden nodes.
- *O* is an  $(N_O + 1)$  element layer of the output nodes.
- $W^I$  is an  $(N_I + 1) \times (N_H + 1)$  input matrix of weights from the input layer to the hidden layer.
- $W^H$  is an  $(N_H + 1) \times (N_O + 1)$  input matrix of weights from the hidden layer to the output layer.
- $W^{I'}$  is an  $(N_I + 1) \times (N_H + 1)$  intermediate matrix of weights from the input layer to the hidden layer.
- $W^{H'}$  is an  $(N_H + 1) \times (N_O + 1)$  intermediate matrix of weights from the hidden layer to the output layer.
- $\vec{S}$  is an  $(N_I + 1)$  element input vector of values for the input layer.
- $\vec{T}$  is an  $(N_O + 1)$  element input vector of target values for the output layer.
- $\vec{D}^O$  is an  $(N_O + 1)$  element intermediate vector of delta values for the output layer.
- $\vec{D}^H$  is an  $(N_H + 1)$  element intermediate vector of delta values for the hidden layer.

**Computation** The following *squash* function is used for Forward Propagation:

$$F(x) = (1 + e^{-x})^{-1}$$

Initially input nodes layer *I* contains the values from input vector  $\vec{S}$ 

• Initialize:

I = S  $I_0 = 1$   $W^{H'} = ||0||$  $W^{I'} = ||0||$ 



• Compute the hidden nodes from the input nodes (Forward Propagation):

$$H_0 = 1$$
$$H_i = F(\sum_{0 \le j \le N_I} I_j \times W_{j,i}^I), 0 < i \le N_H$$

• Compute the output nodes from the hidden nodes (Forward Propagation):

$$O_i = F(\sum_{0 \le j \le N_H} H_j \times W_{j,i}^H), 0 < i \le N_O$$

• Compute the deltas between the target and actual output vectors (Backward Propagation):

$$D_{i}^{O} = O_{i} \times (1 - O_{i}) \times (O_{i} - T_{i}), 0 < i \le N_{O}$$

• Compute the "deltas" for the hidden nodes (Backward Propagation):

$$D_i^H = H_i \times (1 - H_i) \times \left(\sum_{0 < j \le N_O} D_j^O \times W_{i,j}^H\right), 0 < i \le N_H$$

• Update the hidden-to-output weights (Weights Adjustment):

$$W_{i,j}^{H'} = 0.3 \times D_j^O \times H_i + 0.3 \times W_{i,j}^{H'}, 0 \le i \le N_H, 0 < j \le N_O$$

$$W_{i,j}^{H} = W_{i,j}^{H} + W_{i,j}^{H'}, 0 \le i \le N_{H}, 0 < j \le N_{O}$$

• Update the input-to-hidden weights (Weights Adjustment):

$$W_{i,j}^{I'} = 0.3 \times D_j^H \times I_i + 0.3 \times W_{i,j}^{I'}, 0 \le i \le N_I, 0 < j \le N_H$$

$$W_{i,j}^{I} = W_{i,j}^{I} + W_{i,j}^{I'}, 0 \le i \le N_{I}, 0 < j \le N_{H}$$

#### **Iteration spaces**

• Compute the hidden nodes from the input nodes:

$$I_H = \{(i, j) : 0 < i \le N_H, 0 \le j \le N_I\}$$

$$I_H: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 0 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j \end{pmatrix} \le \begin{pmatrix} N_H\\ N_I\\ -1\\ 0 \end{pmatrix}$$





• Compute the output nodes from the hidden nodes:

$$I_O = \{(i, j) : 0 < i \le N_O, 0 \le j \le N_H\}$$

$$I_{O}: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 0 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j \end{pmatrix} \leq \begin{pmatrix} N_{O}\\ N_{H}\\ -1\\ 0 \end{pmatrix}$$

• Compute the delta between the target and actual output:

$$I_{D_O} = \{(i) : 0 < i \le N_O\}$$

$$I_{D_O}: \begin{pmatrix} 1\\ -1 \end{pmatrix} \times (i) \leq \begin{pmatrix} N_O\\ -1 \end{pmatrix}$$

• Compute the "delta" for the hidden nodes:

$$I_{D_H} = \{(i, j) : 0 < i \le N_H, 0 < j \le N_O\}$$

$$I_{D_{H}}: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 0 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j \end{pmatrix} \leq \begin{pmatrix} N_{H}\\ N_{O}\\ -1\\ -1 \end{pmatrix}$$

• Update the hidden-to-output weights:

$$I_{W_H} = \{(i, j) : 0 \le i \le N_H, 0 < j \le N_O\}$$

$$I_{W_{H}}: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 0 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j \end{pmatrix} \leq \begin{pmatrix} N_{H}\\ N_{O}\\ 0\\ -1 \end{pmatrix}$$

• Update the input-to-hidden weights:

$$I_{W_I} = \{(i, j) : 0 \le i \le N_I, 0 < j \le N_H\}$$

$$I_{W_{I}}: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 0 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j \end{pmatrix} \leq \begin{pmatrix} N_{I}\\ N_{H}\\ 0\\ -1 \end{pmatrix}$$

$$I = I_H \cup I_O \cup I_{D_O} \cup I_{D_H} \cup I_{W_H} \cup I_{W_I}$$





## Dependences

$$\begin{split} \delta_{I_O} &= I_H : (j, (1:N_I)) \xrightarrow{t} I_O : (i,j) \\ \delta_{I_{D_O}} &= I_O : (i, (1:N_H)) \xrightarrow{t} I_{D_O} : (i) \\ \delta_{I_{D_H}} &= I_H : (i, (1:N_I)) \xrightarrow{t} I_{D_H} : (i,j) \land \\ &I_{D_O} : (j) \xrightarrow{t} I_{D_H} : (i,j) \\ \delta_{I_{W_H}} &= I_H : (i, (1:N_I)) \xrightarrow{t} I_{W_H} : (i,j) \land \\ &I_{D_O} : (j) \xrightarrow{t} I_{W_H} : (i,j) \\ \delta_{I_{W_I}} &= I_{D_H} : (j, (1:N_O)) \xrightarrow{t} I_{W_I} : (i,j) \end{split}$$

## Memory access mapping

• Compute the hidden nodes from the input nodes: for  $(i, j) \in I_H$ :

$$M_r(i,j) = \{I_j, W_{j,i}^I\}$$
  
 $M_w(i,j) = \{H_i\}$ 

• Compute the output nodes from the hidden nodes: for  $(i, j) \in I_O$ :

$$egin{array}{rll} M_r(i,j) &=& \{H_j,W_{j,i}^H\}\ M_w(i,j) &=& \{O_i\} \end{array}$$

• Compute the delta between the target and actual output: for  $(i) \in I_{D_0}$ :

$$M_r(i) = \{O_i, T_i\}$$
  
 $M_w(i) = \{D_i^O\}$ 

• Compute the delta for the hidden nodes: for  $(i, j) \in I_{D_H}$ :

$$M_r(i,j) = \{H_i, W^H_{i,j}, D^O_j\}$$
  
 $M_w(i,j) = \{D^H_i\}$ 

• Update hidden to output weights: for  $(i, j) \in I_{W_H}$ :

$$M_{r}(i,j) = \{D_{j}^{O}, W_{i,j}^{H'}, H_{i}, W_{i,j}^{H}\}$$
$$M_{w}(i,j) = \{W_{i,j}^{H'}, W_{i,j}^{H}\}$$

• Update input to hidden weights: for  $(i, j) \in I_{W_i}$ :

$$\begin{aligned} M_r(i,j) &= \{D_j^H, W_{i,j}^{I'}, I_i, W_{i,j}^I\} \\ M_w(i,j) &= \{W_{i,j}^{I'}, W_{i,j}^I\} \end{aligned}$$



```
Algorithm 30 The Forward Propagation device code abstraction.
Input: global float dI[]
Input: global float dWI[]
Input: int Nh
Output: global float dH[]
Local: local float dlI[]
Local: local float dlW[]
  1: kernel BPKERNEL1 (dI,dWI,Nh,dH)
          \texttt{xidx} \leftarrow \lambda_0 + 1
 2:
          \texttt{yidx} \gets \gamma_1 + 1
 3:
 4:
          idx \leftarrow (Nh+1) \times yidx + xidx
          if \lambda_0 = 0 then
  5:
               dll[\lambda_1] \leftarrow dl[yidx]
 6:
 7:
          end if
  8:
          BARRIER(CLK_LOCAL_MEM_FENCE)
          \texttt{dlW}[\lambda_1 \times \Lambda_0 + \lambda_0] \leftarrow \texttt{dWI}[\texttt{idx}]
 9:
                                                                                     ▷ Unnecessary memory access.
          BARRIER(CLK_LOCAL_MEM_FENCE)
                                                                                     ▷ Unnecessary memory barrier.
10:
          \mathtt{dlW}[\lambda_1 \times \Lambda_0 + \lambda_0] \leftarrow \mathtt{dlW}[\lambda_1 \times \Lambda_0 + \lambda_0] \times \mathtt{dlI}[\lambda_1]
11:
          BARRIER(CLK_LOCAL_MEM_FENCE)
12:
          for i in 1 : \log_2(\Lambda_1) do
13:
               if \lambda_1 \mod (2^i) = 0 then
14:
                     \texttt{dlW}[\lambda_1 \times \Lambda_0 + \lambda_0] \leftarrow \texttt{dlW}[\lambda_1 \times \Lambda_0 + \lambda_0] + \texttt{dlW}[(\lambda_1 + 2^{i-1}) \times \Lambda_0 + \lambda_0]
15:
               end if
16:
               BARRIER(CLK_LOCAL_MEM_FENCE)
17:
          end for
18:
          dWI[idx] \leftarrow dlW[\lambda_1 \times \Lambda_0 + \lambda_0]
                                                          Unnecessary memory access (see device memory
19:
      access mapping for more details).
          BARRIER(CLK_LOCAL_MEM_FENCE)
                                                                                     ▷ Unnecessary memory barrier.
20:
          if \lambda_0 = 0 then
21:
               \mathtt{dH}[\gamma_1] \gets \mathtt{dlW}[\lambda_1 \times \Lambda_0 + \lambda_0]
22:
          end if
23:
24: end kernel
```



```
Algorithm 31 The Weights Adjustment device code abstraction.
Input: global float dDH[]
Input: global float dI[]
Input/Output: global float dWI[]
Input/Output: global float dWI0[]
Input: int Nh
Output: global int dCout[]
 1: kernel BPKERNEL2 (dDH, dI, dWI, dWIo, Nh)
         \texttt{xidx} \gets \lambda_0 + 1
 2:
 3:
         yidx \leftarrow \gamma_1 + 1
 4:
         idx \leftarrow (Nh+1) \times yidx + xidx
         \texttt{dWIo[idx]} \leftarrow 0.3 \times \texttt{dDH[xidx]} \times \texttt{dI[yidx]} + 0.3 \times \texttt{dWIo[idx]}
 5:
         dWI[idx] \leftarrow dWI[idx] + dWIo[idx]
 6:
         BARRIER(CLK_LOCAL_MEM_FENCE)
 7:
 8:
         if \lambda_1 = 0 \land v_1 = 0 then
             dWIo[xidx] \leftarrow 0.3 \times dDH[xidx] + 0.3 \times dWIo[xidx]
 9:
10:
             dWI[xidx] \leftarrow dWI[xidx] + dWIo[xidx]
         end if
11:
12: end kernel
```

## Low-level implementation details

Only Forward Propagation and Weights Adjustment are implemented as OpenCL kernels (both only for the input-to-hidden computation).

**Device code** See Algorithms 30 and 31.

Host code See Algorithm 32.

## Host data structures

- float hI[Ni+1]:  $I_i \rightarrow$  hI[i]
- float hH[Nh+1]:  $H_i \rightarrow$  hH[i]
- float hO[No+1]:  $O_i \rightarrow$  hO[i]
- float hT[No+1]:  $T_i \rightarrow hT[i]$
- float hDH[Nh+1]:  $D_i^H \rightarrow \text{hDH[i]}$
- float hDO[No+1]:  $D_i^O \rightarrow hDO[i]$
- float hWI[Ni+1][Nh+1]:  $W_{i,j}^I \rightarrow hWI[i][j]$
- float hWH[Nh+1][No+1]:  $W_{i,i}^H \rightarrow hWH[i][j]$
- float hWIo[Ni+1][Nh+1]:  $W_{i,i}^{I'} \rightarrow hWIo[i][j]$
- float hWHo[Nh+1][No+1]:  $W_{i,i}^{H'} \rightarrow$  hWHo[i][j]

```
CARP
```



Algorithm 32 The Back Propagation host code abstraction. 1: **procedure** ALLOCATEMEMORY(Ni,Nh,No) ▷ Allocate host memory and device buffers.  $hI \leftarrow AllocateHostMemory(Ni+1)$ 2:  $hH \leftarrow AllocateHostMemory(Nh+1)$ 3:  $hO \leftarrow AllocateHostMemory(No+1)$ 4:  $hT \leftarrow AllocateHostMemory(No+1)$ 5:  $hDH \leftarrow ALLOCATEHOSTMEMORY(Nh+1)$ 6:  $hDO \leftarrow ALLOCATEHOSTMEMORY(No+1)$ 7:  $hWI \leftarrow AllocateHostMemory((Ni+1) \times (Nh+1))$ 8:  $hWH \leftarrow AllocateHostMemory((Nh+1) \times (No+1))$ 9: 10:  $hWIo \leftarrow AllocateHostMemory((Ni+1) \times (Nh+1))$ 11:  $hWHo \leftarrow AllocateHostMemory((Nh+1) \times (No+1))$  $hPS \leftarrow ALLOCATEHOSTMEMORY(Nh \times N_1)$ 12:  $dI \leftarrow ALLOCATEBUFFER(Ni + 1, READ_WRITE)$ 13: 14:  $dH \leftarrow ALLOCATEBUFFER(Nh \times N_1, READ_WRITE)$  $dDH \leftarrow ALLOCATEBUFFER(Nh + 1, READ_WRITE)$ 15:  $dWI \leftarrow AllocateBuffer((Ni + 1) \times (Hh + 1), READ_WRITE)$ 16:  $dWIo \leftarrow AllocateBuffer((Ni+1) \times (Hh+1), READ_WRITE)$ 17: 18: end procedure 19: procedure INITIALIZEMEMORY(Ni, Nh, No) ▷ Initialize host memory.  $hI \leftarrow RANDOMELEMENTS(Ni+1)$ 20: 21:  $hT \leftarrow RANDOMELEMENTS(No + 1)$  $hWI \leftarrow RANDOMELEMENTS((Ni+1) \times (Nh+1))$ 22: 23:  $hWH \leftarrow RANDOMELEMENTS((Nh+1) \times (No+1))$  $hWIo \leftarrow ZEROELEMENTS((Ni + 1) \times (Nh + 1))$ 24:  $hWHo \leftarrow ZEROELEMENTS((Nh+1) \times (No+1))$ 25: 26: end procedure 27: procedure FORWARDPROPOGATION(hH,hO,hWH,Nh,No) 28:  $h0[0] \leftarrow 1$ for j in 1 : No do 29: 30:  $\texttt{sum} \gets 0$ for k in 0: Nh do 31: 32:  $\texttt{sum} \leftarrow \texttt{sum} + \texttt{hWH}[\texttt{k}][\texttt{j}] \times \texttt{hH}[\texttt{k}]$ 33: end for  $Ho[i] \leftarrow (1 + e^{sum})^{-1}$ 34: end for 35: 36: end procedure 37: procedure OUTPUTERROR(hDO, hT, hO, No) for j in 1 : No do 38: 39: gen  $\leftarrow$  h0[j] 40:  $trg \leftarrow hT[j]$  $hDO[j] \leftarrow sum + gen \times (1 - gen) \times (trg - gen)$ 41: 42: end for 43: end procedure



```
Algorithm 32 The Back Propagation host code abstraction (continues).
44: procedure HIDDENERROR(hDH, hDO, hWH, hH, Nh, No)
         for j in 1 : Nh do
45:
             h \leftarrow hH[j]
46:
             \texttt{sum} \gets 0
47:
             for j in 1 : No do
48:
49:
                 \texttt{sum} \leftarrow \texttt{sum} + \texttt{hWH}[\texttt{j}][\texttt{k}] \times \texttt{hDO}[\texttt{k}]
             end for
50:
51:
             hDH[j] \leftarrow h \times (1-h) \times sum
52:
         end for
53: end procedure
54: procedure ADJUSTWEIGHTS(hD0, hH, hWH, hWHo, Nh, No)
55:
         hH[0] \leftarrow 1
         for j in 1 : No do
56:
57:
             for k in 0: Nh do
                 hWHo[k][j] \leftarrow 0.3 \times hDO[j] \times hH[k] + 0.3 \times hWHo[k][j]
58:
                 hWH[k][j] \leftarrow hWH[k][j + hWHo[k][j]
59:
             end for
60:
         end for
61:
62: end procedure
Input: int Ni
                                                                              ▷ Number of input nodes.
Input: int Nh
                                                                            ▷ Number of hidden nodes.
Input: int No
                                                                            ▷ Number of output nodes.
63: ALLOCATEMEMORY(Ni, Nh, No)
64: INITIALIZEMEMORY(N,M,file)
65: bpProgram \leftarrow BUILDPROGRAM("bpProgram.cl")
66: bpKernel1 ← BUILDKERNEL(bpProgram, "bpKernel1")
67: bpKernel2 ← BUILDKERNEL(bpProgram, "bpKernel2")
68: COPYTODEVICE(dI,hI,True,∅)
69: COPYTODEVICE(dWI, hWI, True, ∅)
70: SETKERNELARGUMENTS(bpKernel1,dI,dWI,Nh,dH)
71: ENQUEUEKERNEL(bpKernel1, (16, Ni), (16, 16), Ø)
72: COPYTOHOST(dH, hPS, True, ∅)
73: for j in 1 : Nh do
         \texttt{sum} \gets 0
74:
         for k in 0: (N_1 - 1) do
75:
             \texttt{sum} \leftarrow \texttt{sum} + \texttt{hPS}[\texttt{k} \times \texttt{Nh} + \texttt{j} - 1]
76:
         end for
77:
         \texttt{sum} \leftarrow \texttt{sum} + \texttt{hWI}[0][j]
78:
         hH[j] \leftarrow (1 + e^{sum})^{-1}
79:
80: end for
```



Algorithm 32 The Back Propagation host code abstraction (continues).

- 81: FORWARDPROPOGATION(hH, h0, hWH, Nh, No)
- 82: OUTPUTERROR(hDO, hT, hO, No)
- 83: HIDDENERROR(hDH, hDO, hWH, hH, Nh, No)
- 84: ADJUSTWEIGHTS(hDO, hH, hWH, hWHo, Nh, No)
- 85: COPYTODEVICE(dDH, hDH, **True**, ∅)
- 86: COPYTODEVICE(dWI, hWI, **True**, **Ø**)
- 87: COPYTODEVICE(dWIo,hWIo, **True**, **Ø**)
- 88: SETKERNELARGUMENTS(bpKernel2,dDH,dI,dWI,dWIo,Nh)
- 89: ENQUEUEKERNEL(bpKernel1,(16,Ni),(16,16),∅)
- 90: COPYTOHOST(dI,hI,**True**,∅)

▷ Unnecessary memory coping.

91: CopyToHost(dWI, hWI, True,  $\emptyset$ )

## Device data structures

- global float dI[Ni+1]:  $I_i \rightarrow$  dI[i]
- global float dH[Nh \* ( $N_1$ )]:  $H_i \rightarrow dH[i + Nh * g], 0 \le g < N_1$
- global float dDH[Nh+1]:  $D_i^H \rightarrow dDH[i]$
- global float dWI[(Ni+1) \* (Nh+1)]:  $W_{i,j}^I \rightarrow dWI[i * (Nh+1) + j]$
- global float dWIo[(Ni+1) \* (Nh+1)]:  $W_{i,j}^{I'} \to$  dWI[i \* (Nh+1) + j]
- local float dll[ $\Lambda_1$ ]:  $I_i \rightarrow dll[(i 1) % (\Lambda_1)]$
- local float dlw[( $\Lambda_1$ ) \* ( $\Lambda_0$ )]:  $W_{i,j}^I \rightarrow \text{dlw}[((i 1) \% (\Lambda_0)) * (\Lambda_1) + (j 1)]$

**Input datasets** The benchmark uses a randomly generated network with the following parameters:

- The number of input nodes: 65536 (can be changed via command line).
- The number of hidden nodes: 16 (Cannot be changed).
- The number of output nodes: 1 (Cannot be changed).

### **Partitioned iteration space**

• Compute the hidden nodes from the input nodes (device):

$$I_{H} = \{(\gamma_{1}, \lambda_{0}, \lambda_{1}, i) : 0 < \gamma_{1} \leq (\Gamma_{1}), 0 \leq \lambda_{0} < (\Lambda_{0}), 0 \leq \lambda_{1} < (\Lambda_{1}), 1 \leq i \leq \log_{2}(\Lambda_{1})\}$$

$$\Lambda_0 = N_H = 16$$
  

$$\Lambda_1 = 16$$
  

$$\Gamma_1 = N_I / \Lambda_1$$



$$f(i,j) \rightarrow ((i-1)/\Lambda_1, j-1, (i-1)\%\Lambda_1, 1: \log_2(\Lambda_1))$$

$$f^{-1}(\gamma_1, \lambda_0, \lambda_1, i) \to (\gamma_1 \times \Lambda_1 + \lambda_1 + 1, \lambda_0 + 1)$$

• Compute the output nodes from the hidden nodes (host):

$$P_O = \{(i, j) : 0 < i \le N_O, 0 \le j \le N_H\}$$

$$f(i,j) \to (i,j)$$
 
$$f^{-1}(i,j) \to (i,j)$$

• Compute the delta between the target and actual output (host):

$$P_{D_O} = \{(i) : 0 < i \le N_O\}$$

$$f(i) \rightarrow (i)$$
  
 $f^{-1}(i) \rightarrow (i)$ 

• Compute the delta for the hidden nodes (host):

$$P_{D_H} = \{(i, j) : 0 < i \le N_H, 0 < j \le N_O\}$$

$$f(i,j) \to (i,j)$$
$$f^{-1}(i,j) \to (i,j)$$

• Update hidden to output weights (host):

$$P_{W_H} = \{(i, j) : 0 \le i \le N_H, 0 < j \le N_O\}$$

 $f(i,j) \to (i,j)$ 

$$f^{-1}(i,j) \to (i,j)$$

• Update input to hidden weights (device):

$$P_{W_I} = \{(\gamma_1, \lambda_0) : 0 \le \gamma_1 < N_I, 0 \le \lambda_0 < N_H\}$$
$$f(i, j) \to (i - 1, j - 1)$$

$$f^{-1}(\gamma_1,\lambda_0) \to (\gamma_1+1,\lambda_0+1)$$



**Device memory access mapping** Each work-item executing bpKernel1 performs the follow-ing reads and writes:

$$\begin{array}{lll} G_r(\gamma_1,\lambda_0,\lambda_1,\mathtt{i}) &=& \{\mathtt{dI}[\gamma_1+1],\mathtt{dWI}[(\mathtt{Nh}+1)\times(\gamma_1+1)+\lambda_0+1],\\ && \mathtt{dII}[\lambda_1],\mathtt{dIW}[\lambda_1\times\Lambda_0+\lambda_0],\mathtt{dIW}[(\lambda_1+2^{i-1})\times\Lambda_0+\lambda_0]\}\\ G_w(\gamma_1,\lambda_0,\lambda_1,\mathtt{i}) &=& \{\mathtt{dII}[\lambda_1],\mathtt{dIW}[\lambda_1\times\Lambda_0+\lambda_0],\mathtt{dH}[\gamma_1],\mathtt{dWI}[(\mathtt{Nh}+1)\times(\gamma_1+1)+\lambda_0+1]\}\end{array}$$

 $dWI[(Nh+1) \times (\lambda_1 + 1) + \lambda_0 + 1]$  memory access is unnecessary since the dWI buffer would be overwritten by host after the kernel execution is finished.

Each work-item executing the bpKernel2 performs the following reads and writes:

$$\begin{array}{lll} G_r(\gamma_1,\lambda_0,\lambda_1,\mathtt{i}) &=& \{\mathtt{dWIo}[(\mathtt{Nh}+1)\times(\gamma_1+1)+\lambda_0+1], \mathtt{dWI}[(\mathtt{Nh}+1)\times(\gamma_1+1)+\lambda_0+1], \\ && \mathtt{dWIo}[\lambda_0+1], \mathtt{dWI}[\lambda_0+1], \mathtt{dDH}[\lambda_0+1], \mathtt{dI}[\gamma_1+1]\} \\ G_w(\gamma_1,\lambda_0,\lambda_1,\mathtt{i}) &=& \{\mathtt{dWIo}[(\mathtt{Nh}+1)\times(\gamma_1+1)+\lambda_0+1], \mathtt{dWI}[(\mathtt{Nh}+1)\times(\gamma_1+1)+\lambda_0+1], \\ && \mathtt{dWIo}[\lambda_0+1], \mathtt{dWI}[\lambda_0+1]\} \end{array}$$

#### Host memory access mapping

• Compute the output nodes from the hidden nodes: for  $(i, j) \in I_O$ :

$$\begin{array}{lll} H_r(i,j) &=& \{\mathtt{h}\mathtt{H}[\mathtt{j}],\mathtt{h}\mathtt{W}\mathtt{H}[\mathtt{j}][\mathtt{i}]\} \\ H_w(i,j) &=& \{\mathtt{h}\mathtt{O}[\mathtt{i}]\} \end{array}$$

• Compute the delta between the target and actual output: for  $(i) \in I_{D_0}$ :

$$\begin{array}{lll} H_r(i) &=& \{\texttt{hO}[\texttt{i}],\texttt{hT}[\texttt{i}]\} \\ H_w(i) &=& \{\texttt{hDO}[\texttt{i}]\} \end{array}$$

• Compute the delta for the hidden nodes: for  $(i, j) \in I_{D_H}$ :

• Update hidden to output weights: for  $(i, j) \in I_{W_H}$ :

$$\begin{split} H_r(i,j) &= \{\texttt{hDO[j]},\texttt{hH[i]},\texttt{hWHo[i][j]},\texttt{hWH[i][j]}\} \\ H_w(i,j) &= \{\texttt{hWHo[i][j]},\texttt{hWH[i][j]}\} \end{split}$$

**Performance metrics** The performance of the algorithm implementation is evaluated using the following metrics:

- The enqueue time for the bpKernel1 launch.
- The total time for each kernel launch (bpKernel1 and bpKernel2).
- The total time for all kernel launches (the whole computation time).



**Validation mechanism** The benchmark supports two implementations of the algorithm – parallel (OpenCL based) and sequential. To validate the OpenCL version the outputs of the both implementations can be compared.

## **Target-specific optimizations**

- **Blocking.** The input data is processed in blocks having the number of columns equal to the number of work items per work group. Consecutive work items of the same group access consecutive memory addresses which improves the memory subsystem performance (caches, coalescing).
- **Taking overhead into account.** Although the hidden to output forward propagation and delta calculation can be implemented as OpenCL kernels, the input data size is not sufficient (the overhead for the kernels launches and memory transfer is greater then the benefits from the parallel computations).
- **Tree reduction.** The OpenCL code implements parallel reduction (sum) using a treebased approach, which requires additional memory barriers.





## 2.3.2 Breadth-First Search

The benchmark performs breadth-first traversal on an ordered unweighted graph computing the shortest path to each node of the graph from the initial node.

#### **High-level description**

Abstract data structures The benchmark operates on the following data objects:

- G = (V, E) is an input graph with vectors of vertices V and edges E.
- *C* is an output vector of costs (|V| = |C|).
- $v_0$  is the initial vertex.

**Computation** The algorithm computes the cost of the shortest path for each node according to the following rules:

$$S_0 = \{v_0\}$$
$$S_{\tau+1} = \{v : v \notin S_j, 0 \le j \le \tau \land \exists u \in S_\tau : (u, v) \in E\}$$

$$C_j = \tau, \forall v_j \in S_{\tau}$$

The algorithm terminates for  $\tau$  such that  $S_{\tau+1} = \emptyset$ .

Iteration spaces and dependences This algorithm requires two iteration spaces:

• Selecting next front  $(S_{\tau+1})$ :

$$I_{S} = \{ (\tau, v, e) : \tau \ge 0, v \in S_{\tau}, e \in E^{v} \}$$

• Marking the selected front as visited:

$$I_M = \{(\tau, v) : \tau \ge 0, v \in S_{\tau+1}\}$$

The computation generates the following dependences:

$$\begin{split} \delta_{I_{S}} &= I_{S} : (\{\tau' \in 0 : (\tau - 1)\}, \{v' \in S_{\tau'}\}, \{e' \in E^{v'}\}) \xrightarrow{t} I_{S} : (\tau, v, e) \land \\ I_{S} : (\{\tau' \in 0 : (\tau - 1)\}, \{v' \in S_{\tau'}\}, \{e' \in E^{v'}\}) \xrightarrow{o} I_{S} : (\tau, v, e) \land \\ I_{M} : (\tau - 1, \{v' \in S_{\tau - 1}\}) \xrightarrow{t} I_{S} : (\tau, v, e) \\ \delta_{I_{M}} &= I_{S} : (\{\tau' \in 0 : \tau\}, \{v' \in S_{\tau'}\}, \{e' \in E^{v'}\}) \xrightarrow{t} I_{M} : (\tau, v) \end{split}$$





## Memory access mapping

• Selecting next front: for  $(\tau, v, e) \in I_S$ :

$$egin{array}{rcl} M_r( au, v, e) &=& \{e_h, v = e_t\} \ M_w( au, v, e) &=& \{v\} \ M_{mw}( au, v, e) &=& \{e_t\} \end{array}$$

• Marking the selected front as visited: for  $(\tau, v) \in I_M$ :

$$egin{array}{rcl} M_r( au,v,e) &=& \{v\} \ M_w( au,v,e) &=& \{v\} \end{array}$$

## Low-level implementation details

Each kernel launch computes next front  $S_{\tau+1}$  [11].

**Device code** See Algorithms 33 and 34. In Algorithm 33 different work items can write to the same memory location. But since all vertices in the same front have the same cost these memory writes do not cause data races.

```
Algorithm 33 The Breadth-First Search (New Front Selection) device code abstraction
Input: global Vertex dV[]
                                                                                 ▷ Array of all nodes
Input: global int dE[]
                                                                                 ▷ Array of all edges
                                                                    ▷ True for current front vertices.
Input/Output: global bool dM[]
Input/Output: global bool dU[]
                                                                       \triangleright True for next front vertices.
                                                                          ▷ True for visited vertices.
Input/Output: global bool dS[]
Input/Output: global int dC[]
Input: int N
 1: kernel BFSKERNEL1 (dV, dE, dM, dU, dS, dC, N)
        if (\gamma < \mathbb{N} \land d\mathbb{M}[\gamma]) then
 2:
            dM[\gamma] \leftarrow false
                                             ▷ Remove vertex being processed from current front.
 3:
            for e in dV[\gamma].offset : (dV[\gamma].offset + dV[\gamma].length - 1) do
 4:
 5:
                 v \leftarrow dE[e]
                if dS[v] = false then
 6:
                     dC[v] \leftarrow dC[\gamma] + 1
 7:
                     dU[v] \leftarrow true
 8:
                 end if
 9:
10:
            end for
        end if
11:
12: end kernel
```

**Host code** The host code executes bfsKernel1 and bfsKernel2 until no new nodes are marked as visited. See Algorithm 35.



```
Algorithm 34 The Breadth-First Search (Front Status Update) device code abstraction
Input/Output: global bool dM[]
Input/Output: global bool dU[]
Input/Output: global bool dS[]
Input/Output: global int dF[]
Input: int N
  1: kernel BFSKERNEL2 (dM, dU, dS, dF, N)
         if (\gamma < \mathbb{N} \land d\mathbb{U}[\gamma]) then
 2:
             dM[\gamma] \leftarrow true
 3:
             dS[\gamma] \leftarrow true
 4:
 5:
             dU[\gamma] \leftarrow false
             \mathtt{dF}[0] \gets \textbf{true}
 6:
         end if
 7:
 8: end kernel
```

### Host data structures

```
• struct Vertex:
```

```
struct Vertex
{
    int offset;
    int length;
};
```

- struct Vertex  $hV[Nv]: V_i \rightarrow hV[i]$
- int hE[Ne]:  $E^{V_i} \rightarrow hE[hV[i].offset: (hV[i].offset + hV[i].length 1)]$
- int hC[Nv]:  $C_i \rightarrow$  hC[i]
- bool hM[Nv]:  $V_i \rightarrow$  hM[i]
- **bool** hU[Nv]:  $V_i \rightarrow$  hU[i]
- **bool** hS[Nv]:  $V_i \rightarrow$  hS[i]

The graph is represented by array of vertices hV[] and array of edges hE[]. Each edge is represented by an integer, which is the index of its tail vertex in array hV[].

Each vertex is represented by a structure, which describes the slice in array hE[], which corresponds to edges for which this vertex is the head.

## **Device data structures**

```
• struct Vertex:
```

```
struct Vertex
{
    int offset;
    int length;
};
```

• struct Vertex  $dV[Nv]: V_i \rightarrow dV[i]$ 



Algorithm 35 The Breadth-First Search host code abstraction. 1: **procedure** ALLOCATEMEMORY(Nv, Ne) ▷ Allocate host memory and device buffers.  $hV \leftarrow AllocateHostMemory(Nv)$ 2:  $hE \leftarrow ALLOCATEHOSTMEMORY(Ne)$ 3:  $hM \leftarrow ALLOCATEHOSTMEMORY(Nv)$ 4:  $hU \leftarrow ALLOCATEHOSTMEMORY(Nv)$ 5:  $hS \leftarrow AllocateHostMemory(Nv)$ 6: 7:  $hC \leftarrow ALLOCATEHOSTMEMORY(Nv)$  $hCref \leftarrow ALLOCATEMEMORY(Nv)$ 8:  $dV \leftarrow ALLOCATEBUFFER(Nv, READ_ONLY, hV)$ 9: 10:  $dE \leftarrow AllocateBuffer(Ne, READ_ONLY, hE)$ 11:  $dM \leftarrow AllocateBuffer(Nv, READ_WRITE, hM)$  $dU \leftarrow AllocateBuffer(Nv, READ_WRITE, hU)$ 12:  $dS \leftarrow ALLOCATEBUFFER(Nv, READ_WRITE, hS)$ 13: 14:  $dC \leftarrow ALLOCATEBUFFER(Nv, READ_WRITE, hC)$ 15:  $dF \leftarrow AllocateBuffer(1, READ_WRITE, hC)$ 16: end procedure 17: **procedure** RESETMEMORY(Nv, v0) ▷ Reset temporary memory for reuse. 18: for i in 0 : (Nv - 1) do  $hM[i] \leftarrow false$ 19:  $hU[i] \leftarrow false$ 20:  $hS[i] \leftarrow false$ 21: end for 22: 23:  $hM[v0] \leftarrow true$  $hS[v0] \leftarrow true$ 24: 25: end procedure 26: **procedure** INITIALIZEMEMORY(file, Nv, Ne, v0) ▷ Initialize host memory. 27: RESETMEMORY(Nv, v0) for i in 0 : (Nv - 1) do 28:  $hC[i] \leftarrow -1$ 29:  $hCref[i] \leftarrow -1$ 30: 31: end for 32:  $hV \leftarrow READFROMFILE(file, Nv)$  $\texttt{hE} \leftarrow \texttt{READFROMFILE}(\texttt{file}, \texttt{Ne})$ 33: 34:  $hC[v0] \leftarrow 0$  $hCref[v0] \leftarrow 0$ 35: 36: end procedure 37: procedure COMPARERESULTS(hC,hCref,Nv)  $\texttt{error} \gets \textbf{false}$ 38: for i in 0: (NV - 1) do 39: if  $hC[i] \neq hCref[i]$  then 40:  $error \leftarrow true$ 41: end if 42: end for 43: 44: end procedure



```
Algorithm 35 The Breadth-First Search host code abstraction (continues).
 1: procedure BFsHost(hV, hE, hM, hU, hS, hC, Nv)
 2:
        repeat
 3:
            hF \leftarrow false
            for k in 0 : (Nv - 1) do
 4:
                if hM[k]) then
 5:
                    hM[k] \leftarrow false
 6:
                    for i in hV[k].offset : (hV[k].offset + hV[k].length - 1) do
 7:
                        id \leftarrow hE[i]
 8.
                        if hS[id] = false then
 9:
10:
                            hC[id] \leftarrow hC[k] + 1
                            hU[id] \leftarrow true
11:
                        end if
12:
                    end for
13:
                end if
14:
            end for
15:
            for v in 0 : (Nv - 1) do
16:
                if hM[v]) then
17:
                    \texttt{hM}[\texttt{v}] \leftarrow \textbf{true}, \texttt{hS}[\texttt{v}] \leftarrow \textbf{true}, \texttt{hF} \leftarrow \textbf{true}
18:
19:
                    hU[v] \leftarrow false
                end if
20:
21:
            end for
        until hF = false
22:
23: end procedure
Input: int Nv
                                                          ▷ Number of vertices in the input graph.
Input: int Ne
                                                            ▷ Number of edges in the input graph.
Input: int v0 = 0
                                                                                    ▷ Initial vertex.
Input: const char * file
                                                                  ▷ File with pre-generated graph.
24: ALLOCATEMEMORY(Nv, Ne)
25: INITIALIZEMEMORY(file, Nv, Ne, v0)
26: bfsProgram ← BUILDPROGRAM("bfsProgram.cl")
27: bfsKernel1 \leftarrow BUILDKERNEL(bfsProgram, "bfsKernel1")
28: bfsKernel2 ~ BUILDKERNEL(bfsProgram, "bfsKernel2")
29: COPYTODEVICE(dV, hV, True, ∅)
30: COPYTODEVICE(dE, hE, True, \emptyset)
31: COPYTODEVICE(dM, hM, True, ∅)
32: COPYTODEVICE(dU, hU, True, ∅)
33: COPYTODEVICE(dS,hS, True, ∅)
34: COPYTODEVICE(dC, hC, True, \emptyset) > Probably a bug, since CL_MEM_USE_HOST_PTR is used
    when creating buffers.
```





Algorithm 55 The Dieduti-First Search nost code abstraction (continues)
---

- 35: repeat
- 36:  $hF \leftarrow false$
- 37: COPYTODEVICE( $dF, hF, True, \emptyset$ )
- 38: SETKERNELARGUMENTS(bfsKernel1,dV,dE,dM,dU,dS,dC,Nv)
- 39: ENQUEUEKERNEL(bfsKernel1,(Nv,Nv), none,∅)
- 40: SETKERNELARGUMENTS(bfsKernel2, dM, dU, dS, dF, Nv)
- 41: ENQUEUEKERNEL(bfsKernel2, (Nv, Nv), none,  $\emptyset$ )
- 42: COPYTOHOST( $dF, hF, True, \emptyset$ )
- 43: **until** hF = false
- 44: COPYTOHOST(dC,hC, **True**, ∅)
- 45: ResetMemory(Nv, v0)
- 46: BFSHOST(hV,hE,hM,hU,hS,hCref,Nv)
- 47: COMPARERESULTS(hC,hCref,Nv)
  - int dE[Ne]:  $E^{V_i} \rightarrow dE[dV[i].offset:(dV[i].offset + dV[i].length 1)]$
  - int dC[Nv]:  $C_i \rightarrow dC[i]$
  - bool  $dM[Nv]: V_i \to dM[i]$
  - bool dU[Nv]:  $V_i \rightarrow dU[i]$
  - **bool** dS[Nv]:  $V_i \rightarrow$  dS[i]

Input datasets The benchmark uses a set of pre-generated graphs with the following sizes:

- 4096 vertices; 24576 edges
- 65536 vertices; 393216 edges
- 1000000 vertices; 5999970 edges

#### **Partitioned iteration space**

• Selecting next front (device):

$$P_{S} = \{(\tau, \gamma, i): 0 \le \tau, 0 \le \gamma < \Gamma = \texttt{Nv}, \texttt{dV}[\gamma].\texttt{offset} \le i < \texttt{dV}[\gamma].\texttt{offset} + \texttt{dV}[\gamma].\texttt{length}\}$$

$$f(\tau, \nu = V_i, e = E_i^{\nu} = E_k) \to (\tau, i, k)$$

$$f^{-1}(\tau,\gamma,i) \to (\tau,V_{\gamma},E_i)$$

• Marking the selected front as visited (device):

$$P_M = \{(\tau, \gamma) : 0 \le \tau, 0 \le \gamma < \Gamma = \mathtt{Nv}\}$$



$$f(\tau, \nu = V_i) \rightarrow (\tau, i)$$
  
 $f^{-1}(\tau, \gamma) \rightarrow (\tau, V_{\gamma})$ 

• Selecting next front (host):

$$P_{S} = \{(\tau, l, i) : 0 \leq \tau, 0 \leq l < \texttt{Nv}, \texttt{hV}[l].\texttt{offset} \leq i < \texttt{hV}[l].\texttt{offset} + \texttt{hV}[l].\texttt{length}\}$$

$$f(\tau, \nu = V_i, e = E_j^{\nu} = E_k) \to (\tau, i, k)$$

$$f^{-1}(\tau,l,i) \to (\tau,V_l,E_i)$$

• Marking the selected front as visited (host):

$$P_M = \{(\tau, l) : 0 \le \tau, 0 \le l < \mathbb{N} \mathsf{v}\}$$

 $f(\tau, v = V_i) \to (\tau, i)$ 

$$f^{-1}(\tau,\gamma) \to (\tau,V_{\gamma})$$

**Device memory access mapping** Each work-item executing the bfsKernel1 performs the following reads and writes:

Each work-item executing the bfsKernel2 performs the following reads and writes:

$$\begin{array}{lll} G_r(\tau,\gamma) &=& \{\mathrm{dM}[\gamma]\} \\ G_w(\tau,\gamma) &=& \{\mathrm{dM}[\gamma],\mathrm{dS}[\gamma],\mathrm{dU}[\gamma]\} \end{array} \end{array}$$





**Host memory access mapping** On each iteration on the host corresponding to the next front selection the following memory accesses are performed:

On each iteration on the host corresponding to the marking the selected front as visited the following memory accesses are performed:

 $\begin{array}{lll} H_r(\tau, {\bf k}) &=& \{ {\rm d} {\bf M}[{\bf k}] \} \\ H_w(\tau, {\bf k}) &=& \{ {\rm d} {\bf M}[{\bf k}], {\rm d} {\bf S}[{\bf k}], {\rm d} {\bf U}[{\bf k}] \} \end{array}$ 

**Performance metrics** The performance of the algorithm implementation is evaluated using the following metrics:

- The total time for each kernel launch (computation time for each front).
- The total time for all kernel launches (the whole computation time).

**Validation mechanism** The benchmark runs two implementations of the algorithm – parallel (using OpenCL) and sequential. In order to validate the OpenCL version the outputs of the both implementations are compared.

### **Target-specific optimizations**

• Synchronization-free kernels. Although there are output dependences between iterations performed by different working items in parallel, no synchronization primitives are used. In general such implementation causes data races and invalid results. In given implementation different work items can write to the same memory location, as long as they write the same values, which ensures the correctness of the implementation.

### **Target-specific optimizations opportunities**

• Reduce number of global memory accesses. In the current algorithm implementation the same vertex can be marked as processed multiple times (by multiple working items). Although this doesn't cause data races (see explanation above), additional memory accesses can affect the performance of the code. They can be avoided by caching the results in work group local memory and performing the global memory update by a single work item.





## 2.3.3 CFD Solver

### **High-level description**

The benchmark solves three-dimensional Euler equations for compressible flow [7]:

$$\frac{d}{dt} \int_{\Omega} u d\Omega + \int_{\Gamma} F \cdot n d\Gamma = 0, \qquad (2.12)$$

where

$$u = \begin{cases} \rho \\ \rho v_{x} \\ \rho v_{y} \\ \rho v_{z} \\ \rho e \end{cases}, F = \begin{cases} \rho v_{x} & \rho v_{y} & \rho v_{z} \\ \rho v_{x}^{2} + p & \rho v_{x} v_{y} & \rho v_{x} v_{z} \\ \rho v_{y} v_{x} + p & \rho v_{y}^{2} + p & \rho v_{y} v_{z} \\ \rho v_{z} v_{x} + p & \rho v_{z} v_{y} & \rho v_{z}^{2} + p \\ v_{x} (\rho e + p) & v_{y} (\rho e + p) & v_{z} (\rho e + p) \end{cases},$$
(2.13)

and

$$p = (\gamma - 1)\rho \left[ e - \frac{1}{2} ||v||^2 \right]$$
 (2.14)

- $\rho$  denotes the density.
- $v_x, v_y, v_x$  denote the x, y, z velocities.
- *e* denotes the total energy.
- *p* denotes the pressure.
- $\gamma$  denotes the ratio of the heats.

The continuous equations (2.12), (2.13) and (2.14) are discretized using a cell-centered, finite-volume scheme of the form:

$$\operatorname{vol}_{i} \frac{du_{i}}{dt} = R_{i} = -\sum_{\text{faces}} \|s\| \left[ \frac{1}{2} (f_{i} + f_{j}) - \beta \cdot \lambda_{\max} \cdot (u_{i} - u_{j}) \right]$$
(2.15)

where

$$f_i = \frac{s}{\|s\|} \cdot F_i, \quad \lambda_{\max} = \|v\| + c \tag{2.16}$$

- $vol_i$  denotes the volume of the  $i^{th}$  element.
- *s* denotes the face normal.
- *j* is the index of the neighboring element.
- $\beta$  is the parameter, controlling the amount of artificial viscosity.
- *c* is the speed of sound.





## Abstract data structures

- A is an input N element areas array.
- *C* is an input  $N \times N_{nb}$  elements neighbor array.
- $C_{i,j}$  the  $j^{th}$  neighbor of the  $i^{th}$  cell.
- *U* is an intermediate *N* elements cell array.
- $U^f$  is an input initial value of the U array elements.
- *N* is an intermediate  $N \times N_{nb}$  elements normal array.
- $N_{i,j}$  the  $j^{th}$  normal of the  $i^{th}$  cell.
- *F* is an intermediate *N* element fluxes array.
- *S* is an intermediate *N* element step factors array.

**Computation** The following function are used for computation:

- ComputeStepFactor(a, u). Pure function, which computes the step factor.
- ComputeFlux(c, n, u, U). Function, which computes the flux. It might read any element of the U array.
- ComputeTimeStep(u, s, f, j). Pure function, which computes new value of the  $U_i$  variable.

On every iteration k in 0 :  $(K_{max} - 1)$  the following computations are performed:

• Compute step factor:

$$S_i = \texttt{ComputeStepFactor}(A_i, U_i), 0 \le i < N$$

- For each dimension j in 0 : (RK 1):
  - Compute flux:

$$F_i = \texttt{ComputeFlux}(C_i, N_i, U_i, U), 0 \le i < N$$

- Compute time step:

$$U_i = \texttt{ComputeTimeStep}(U_i, S_i, F_i, j), 0 \le i < N$$

### **Iteration spaces**

• Compute step factor:

$$I_{S} = \{ (k,i) : 0 \le k < k_{\max}, 0 \le i < N \}$$

• Compute flux:

$$I_F = \{(k, j, i, n) : 0 \le k < k_{\max}, 0 \le j < \texttt{RK}, 0 \le i < N, 0 \le n < N_{nb}\}$$

• Compute time step:

$$I_T = \{(k, j, i) : 0 \le k < k_{\max}, 0 \le j < \mathtt{RK}, 0 \le i < N\}$$





## Dependences

## Memory access mapping

• Compute step factor:

$$M_r(k,i) = \{A_i, U_i\}$$
$$M_w(k,i) = \{S_i\}$$

• Compute flux:

$$egin{array}{rll} M_r(k,j,i,n) &= \{C_{i,n},N_{i,n},U_i\}\ M_{mr}(k,j,i,n) &= \{U\}\ M_w(k,j,i,n) &= \{F_i\} \end{array}$$

• Compute time step:

$$M_r(k, j, i) = \{U_i, S_i, F_i\}$$
  
 $M_w(k, j, i) = \{U_i\}$ 

### Low-level implementation details

In low level implementation each cell is represented as five element slice of the [hd]U array. See Host/Device data structures for more details.

**Device code** See Algorithm 36, Algorithm 37, Algorithm 38, Algorithm 39, Algorithm 40.

 Algorithm 36 The CFD device code abstraction (Memset).

 Input: char Val

 Output: global char dM[]

 1: kernel CFDKERNEL1 (Val,dM)

 2:  $dM[\gamma] \leftarrow Val$  

 3: end kernel

Host code See Algorithm 41.



```
Algorithm 37 The CFD device code abstraction (Initialize).

Input: global float dUf[]

Input: global int N

Output: global float dU[]

1: kernel CFDKERNEL2 (dUf, N, dUf)

2: for j in 0: (Nvar - 1) do

3: dU[\gamma + j \times N] \leftarrow dUf[j]

4: end for

5: end kernel
```

```
Algorithm 38 The CFD device code abstraction (Compute step factor).
Input: global float dU[]
Input: global float dA[]
Input: global int N
Output: global float dS[]
  1: kernel CFDKERNEL3 (dU, dA, N, dS)
 2:
          \rho \leftarrow dU[\gamma + N \times 0]
          \mathtt{Ux} \leftarrow \mathtt{dU}[\gamma + \mathtt{N} \times 1]
 3:
          Uy \leftarrow dU[\gamma + N \times 2]
 4:
          Uz \leftarrow dU[\gamma + N \times 3]
 5:
          \texttt{Ue} \leftarrow \texttt{dU}[\gamma + \texttt{N} \times 4]
 6:
          dS[\gamma] \leftarrow COMPUTESTEPFACTOR(dA[\gamma], \rho, Ux, Uy, Uz, Ue)
 7:
 8: end kernel
```





Algorithm 39 The CFD device code abstraction (Compute flux). Input: global float dC[] Input: global float dN[] Input: global float dU[] Input: global int N Output: global float dF[] 1: kernel CFDKERNEL4 (dC, dN, dU, N, dF)  $U_{\rho} \leftarrow dU[\gamma + N \times 0]$ 2:  $U_x \leftarrow dU[\gamma + N \times 1]$ 3:  $U_v \leftarrow dU[\gamma + N \times 2]$ 4:  $U_z \leftarrow dU[\gamma + N \times 3]$ 5: 6:  $U_e \leftarrow dU[\gamma + N \times 4]$ ▷ Extract cell data.  $\mathtt{f}_{\rho} \leftarrow \text{GetInitialFluxDensity}()$ 7:  $f_x \leftarrow GetInitialFluxMomentumX()$ 8:  $f_v \leftarrow GetInitialFluxMomentumY()$ 9:  $f_z \leftarrow GetInitialFluxMomentumZ()$ 10:  $f_e \leftarrow GetInitialFluxEnergy()$ ▷ Get initial flux values. 11: for j in 0 : (Nnb - 1) do 12: 13:  $\operatorname{idx} \leftarrow \operatorname{dC}[\gamma + j \times \mathbb{N}]$ ▷ Get neighbor index.  $N_x \leftarrow dN[\gamma + N \times (j + 0 \times Nnb)]$ 14:  $N_v \leftarrow dN[\gamma + N \times (j + 1 \times Nnb)]$ 15:  $N_z \leftarrow dN[\gamma + N \times (j + 2 \times Nnb)]$  $\triangleright$  Extract normal data. 16:  $C_{\rho} \leftarrow dU[idx + N \times 0]$ 17:  $C_x \leftarrow dU[idx + N \times 1]$ 18: 19:  $C_v \leftarrow dU[idx + N \times 2]$  $C_z \leftarrow dU[idx + N \times 3]$ 20:  $C_e \leftarrow dU[idx + N \times 4]$ ▷ Extract neighboring cell data. 21: 22:  $f_{\rho} \leftarrow UPDATEFLUXDENSITY(U_{\rho}, U_x, U_y, U_z, U_e, C_{\rho}, C_x, C_y, C_z, C_e, N_x, N_y, N_z)$  $\mathbf{f}_{x} \leftarrow UPDATEFLUXMOMENTUMX(\mathbf{U}_{\rho}, \mathbf{U}_{x}, \mathbf{U}_{y}, \mathbf{U}_{z}, \mathbf{U}_{e}, \mathbf{C}_{\rho}, \mathbf{C}_{x}, \mathbf{C}_{y}, \mathbf{C}_{z}, \mathbf{C}_{e}, \mathbf{N}_{x}, \mathbf{N}_{y}, \mathbf{N}_{z})$ 23:  $\mathbf{f}_{y} \leftarrow \text{UPDATEFLUXMOMENTUMY}(\mathbf{U}_{\rho}, \mathbf{U}_{x}, \mathbf{U}_{y}, \mathbf{U}_{z}, \mathbf{U}_{e}, \mathbf{C}_{\rho}, \mathbf{C}_{x}, \mathbf{C}_{y}, \mathbf{C}_{z}, \mathbf{C}_{e}, \mathbf{N}_{x}, \mathbf{N}_{y}, \mathbf{N}_{z})$ 24:  $\mathbf{f}_z \leftarrow \text{UPDATEFLUXMOMENTUMZ}(\mathbf{U}_{\rho}, \mathbf{U}_x, \mathbf{U}_y, \mathbf{U}_z, \mathbf{U}_e, \mathbf{C}_{\rho}, \mathbf{C}_x, \mathbf{C}_y, \mathbf{C}_z, \mathbf{C}_e, \mathbf{N}_x, \mathbf{N}_y, \mathbf{N}_z)$ 25:  $f_e \leftarrow UPDATEFLUXENERGY(U_{\rho}, U_x, U_y, U_z, U_e, C_{\rho}, C_x, C_y, C_z, C_e, N_x, N_y, N_z) \triangleright Update$ 26: flux values. end for 27:  $d\mathbf{F}[\gamma + 0 \times N] \leftarrow \mathbf{f}_{\rho}$ 28:  $d\mathbf{F}[\gamma + 1 \times N] \leftarrow \mathbf{f}_x$ 29:  $dF[\gamma + 2 \times N] \leftarrow f_{\nu}$ 30:  $dF[\gamma + 3 \times N] \leftarrow f_z$ 31:  $dF[\gamma + 4 \times N] \leftarrow f_e$ 32: 33: end kernel



Algorithm 40 The CFD device code abstraction (Compute time step). Input: global float dUold[] Input: global float dS[] Input: global float dF[] Input: global int N Input: global int j Output: global float dU[] 1: **kernel** CFDKERNEL5 (dUold, dS, dF, Nj, dU) factor  $\leftarrow dF[\gamma]/(RK + 1 - j)$ 2:  $\texttt{dU}[\gamma + \texttt{N} \times 0] \leftarrow \texttt{dUold}[\gamma + \texttt{N} \times 0] + \texttt{factor} \times \texttt{dF}[\gamma + \texttt{N} \times 0]$ 3: 4:  $\texttt{dU}[\gamma + \texttt{N} \times 1] \leftarrow \texttt{dUold}[\gamma + \texttt{N} \times 1] + \texttt{factor} \times \texttt{dF}[\gamma + \texttt{N} \times 1]$ 5:  $d\mathtt{U}[\gamma+\mathtt{N}\times 2] \leftarrow d\mathtt{Vold}[\gamma+\mathtt{N}\times 2] + \mathtt{factor}\times d\mathtt{F}[\gamma+\mathtt{N}\times 2]$  $d\mathtt{U}[\gamma+\mathtt{N}\times3] \leftarrow d\mathtt{Vold}[\gamma+\mathtt{N}\times3] + \mathtt{factor}\times d\mathtt{F}[\gamma+\mathtt{N}\times3]$ 6:  $dU[\gamma + N \times 4] \leftarrow dUold[\gamma + N \times 4] + factor \times dF[\gamma + N \times 4]$ 7: 8: end kernel

### Host data structures

```
• float hA[N]: A_i \rightarrow dA[i]
    • float hN[N * Nnb * 3]: N_{i,j} \rightarrow
       hN[(i + (j + 0 * Nnb) * N):(i + (j + 2 * Nnb) * N):(N * Nnb)]
          -N_{i,j}.x \rightarrow hN[i + (j + 0 * Nnb) * N]
          - N_{i,i}, y \rightarrow hN[i + (j + 1 * Nnb) * N]
          - N_{i,j} \cdot z \rightarrow hN[i + (j + 2 * Nnb) * N]
    • int hC[N * Nnb]: C_{i,j} \rightarrow hC[i + j ( N]
    • float hUf[5]: U^f \rightarrow hUf
          - U^f.\rho \rightarrow huf[0]
          - U^f.x \rightarrow hUf[1]
          - U^f.y \rightarrow hUf[2]
          - U^f.z \rightarrow hUf[3]
          - U^f.e \rightarrow huf[4]
Device data structures
    • global float dA[N]: A_i \rightarrow dA[i]
    • global float dN[N * Nnb * 3]: N_{i,j} \rightarrow
       dN[(i + (j + 0 * Nnb) * N):(i + (j + 2 * Nnb) * N):(N * Nnb)]
          - N_{i,j}.x \rightarrow dN[i + (j + 0 * Nnb) * N]
          - N_{i,j}. y \rightarrow dN[i + (j + 1 * Nnb) * N]
          - N_{i,j} \cdot z \rightarrow dN[i + (j + 2 * Nnb) * N]
```





Algorithm 41 The CFD host code abstraction. 1: procedure ALLOCATEMEMORY(N, Nnb) ▷ Allocate host memory and device buffers.  $hA \leftarrow AllocateHostMemory(N)$ 2: 3:  $hN \leftarrow ALLOCATEHOSTMEMORY(N \times Nnb \times 3)$  $hC \leftarrow AllocateHostMemory(N \times Nnb)$ 4: 5:  $hUf \leftarrow ALLOCATEHOSTMEMORY(5)$  $dA \leftarrow AllocateBuffer(N, Nnb, READ_WRITE)$ 6:  $dN \leftarrow AllocateBuffer(N \times Nnb \times 3, READ_WRITE)$ 7:  $dF \leftarrow AllocateBuffer(N \times 5, READ_WRITE)$ 8:  $dC \leftarrow ALLOCATEBUFFER(N \times Nnb, READ_WRITE)$ 9:  $dS \leftarrow AllocateBuffer(N, READ_WRITE)$ 10:  $dU \leftarrow AllocateBuffer(N \times 5, READ_WRITE)$ 11:  $dUold \leftarrow AllocateBuffer(N \times 5, READ_WRITE)$ 12: 13.  $dUf \leftarrow ALLOCATEBUFFER(5, READ_WRITE)$ 14: end procedure 15: **procedure** INITIALIZEMEMORY(file) ▷ Initialize host memory.  $hA \leftarrow READFROMFILE(file)$ 16:  $hN \leftarrow READFROMFILE(file)$ 17: 18:  $hC \leftarrow READFROMFILE(file)$ 19:  $hUf \leftarrow GENERATEINITIAL()$ 20: end procedure  $\triangleright$  Number of cells. Input: int N Input: int Ni ▷ Number of iterations. Input: const char \* file  $\triangleright$  File with pre generated data. 21: Nnb  $\leftarrow 4$ 22: ALLOCATEMEMORY(N, Nnb) 23: INITIALIZEMEMORY(file) 24: COPYTODEVICE( $dA, hA, True, \emptyset$ ) 25: COPYTODEVICE(dC, hC, **True**, ∅) 26: COPYTODEVICE(dN, hN, **True**, ∅) 27: COPYTODEVICE(dUf, hUf, True,  $\emptyset$ ) 28: cfdProgram ← BUILDPROGRAM("cfdProgram.cl") 29: cfdKernel1  $\leftarrow$  BUILDKERNEL(cfdProgram,"cfdKernel1") 30: cfdKernel2 ← BUILDKERNEL(cfdProgram, "cfdKernel2") 31: cfdKernel3 ← BUILDKERNEL(cfdProgram, "cfdKernel3") 32: cfdKernel4 ~ BUILDKERNEL(cfdProgram, "cfdKernel4") 33: cfdKernel5 ← BUILDKERNEL(cfdProgram, "cfdKernel5") 34: SETKERNELARGUMENTS(cfdKernel2, N, dU, dUf) 35: ENQUEUEKERNEL(sradKernel2, (N), (192), ∅) 36: SETKERNELARGUMENTS(cfdKernel2, N, dUold, dUf) 37: ENQUEUEKERNEL(sradKernel2,(ℕ),(192),∅) 38: SETKERNELARGUMENTS(cfdKernel2, N, dF, dUf) 39: ENQUEUEKERNEL(sradKernel2,(N),(192),∅) 40: SETKERNELARGUMENTS(cfdKernel1,0,dS) 41: ENQUEUEKERNEL(sradKernel2, (N×sizeof(float)), (192), ∅)





Algorithm 41 The CFD host code abstraction (continues).		
42:	for $i in 0 : (Ni - 1) do$	
43:	$\texttt{RK} \leftarrow 3$	
44:	$COPYBUFFER(dU, dUold, \emptyset)$	
45:	SETKERNELARGUMENTS(cfdKernel3,dU,dA,N,dS)	
46:	$EnQUEUEKERNEL(\mathtt{sradKernel3},(\mathtt{N}),(192), \emptyset)$	
47:	for $j in 0 : (RK - 1) do$	
48:	SETKERNELARGUMENTS(cfdKernel4,dC,dN,dU,N,dF)	
49:	$ENQUEUEKERNEL(\mathtt{sradKernel4},(\mathtt{N}),(192), \emptyset)$	
50:	SETKERNELARGUMENTS(cfdKernel5,dUold,dS,dF,Nj,dU)	
51:	$EnQUEUEKerNEL(\mathtt{sradKernel5},(\mathtt{N}),(192), \emptyset)$	
52:	end for	

53: **end for** 

```
• global float dF[N \star 5]: F_i \rightarrow
   dF[(i + 0 * N):(i + 4 * N):N]
      - F_i.\rho \rightarrow dF[i + 0 * N]
      - F_i.x \rightarrow dF[i + 1 * N]
      - F_i.y \rightarrow dF[i + 2 * N]
      - F_{i.z} \rightarrow dF[i + 3 * N]
      - F_i.e \rightarrow dF[i + 4 * N]
• global int dC[N * Nnb]: C_{i,j} \rightarrow dC[i + j ( N]
• global float dS[N]: S_i \rightarrow dS[i]
• global float dUf[5]: U^f \to dUf
      - U^f.
ho \to duf[0]
      - U^f.x \rightarrow \operatorname{dUf}[1]
      - U^f.y \rightarrow duf[2]
      - U^f.z \rightarrow duf[3]
      - U^f.e \rightarrow dUf[4]
• global float dUold[N * 5]: U_i \rightarrow dUold[(i + 0 * N):(i + 4 * N):N]
      - U_i.
ho \rightarrow dUold[i + 0 * N]
      - U_i . x \rightarrow dUold[i + 1 * N]
      - U_i.y \rightarrow dUold[i + 2 * N]
      - U_i.z \rightarrow dUold[i + 3 * N]
      - U_i.e \rightarrow \text{dUold[i + 4 * N]}
• global float dU[N * 5]: U_i \rightarrow
   dU[(i + 0 * N):(i + 4 * N):N]
```




-  $U_i \cdot \rho \rightarrow du[i + 0 * N]$ -  $U_i \cdot x \rightarrow du[i + 1 * N]$ -  $U_i \cdot y \rightarrow du[i + 2 * N]$ -  $U_i \cdot z \rightarrow du[i + 3 * N]$ -  $U_i \cdot e \rightarrow du[i + 4 * N]$ 

**Input datasets** The benchmark uses pre-generated input datasets with following number of elements (*N*): 97046, 193474, 232536.

For all input data sets the number of neighboring elements  $(N_{nb})$  is 4.

## **Partitioned iteration space**

• Compute step factor (cfdKernel3):

$$P_{S} = \{ (k, \gamma) : 0 \le k < N_{i}, 0 \le \gamma < \Gamma = N \}$$

$$f(k,i) \to (k,i)$$

$$f^{-1}(k,\gamma) \to (k,\gamma)$$

• Compute flux (cfdKernel4):

$$P_F = \{(k, j, \gamma, n) : 0 \le k < N_i, 0 \le j < \texttt{RK}, 0 \le \gamma < \Gamma = N, 0 \le n < N_{nb}\}$$

$$f(k, j, i, n) \rightarrow (k, j, i, n)$$

$$f^{-1}(k, j, \gamma, n) \to (k, j, \gamma, n)$$

• Compute time step (cfdKernel5):

$$P_T = \{(k, j, \gamma) : 0 \le k < N_i, 0 \le j < \mathsf{RK}, 0 \le \gamma < \Gamma = N\}$$

$$f(k, j, i) \rightarrow (k, j, i)$$

$$f^{-1}(k, j, \gamma) \to (k, j, \gamma)$$

CARP-ARM-RP-001-v1.0





## **Device memory access mapping**

• Compute step factor (cfdKernel3):

 $\begin{array}{lll} G_r(k,\gamma) &=& \{ \mathtt{dA}[\gamma], \mathtt{dU}[\gamma+\mathtt{N}\times 0], \mathtt{dU}[\gamma+\mathtt{N}\times 1], \mathtt{dU}[\gamma+\mathtt{N}\times 2], \mathtt{dU}[\gamma+\mathtt{N}\times 3], \mathtt{dU}[\gamma+\mathtt{N}\times 4] \} \\ G_w(k,\gamma) &=& \{ \mathtt{dS}[\gamma] \} \end{array}$ 

• Compute flux (cfdKernel4):

• Compute time step (cfdKernel5):

$$\begin{array}{lll} G_r(k,j,\gamma) &=& \{\mathrm{dF}[\gamma+0\times N], \mathrm{dF}[\gamma+1\times N], \mathrm{dF}[\gamma+2\times N], \mathrm{dF}[\gamma+3\times N], \mathrm{dF}[\gamma+4\times N], \\ && \mathrm{dUold}[\gamma+\mathrm{N}\times 0], \mathrm{dUold}[\gamma+\mathrm{N}\times 1], \mathrm{dUold}[\gamma+\mathrm{N}\times 2], \\ && \mathrm{dUold}[\gamma+\mathrm{N}\times 3], \mathrm{dUold}[\gamma+\mathrm{N}\times 4], \mathrm{dS}[\gamma] \} \\ G_w(k,j,\gamma) &=& \{\mathrm{dU}[\gamma+\mathrm{N}\times 0], \mathrm{dU}[\gamma+\mathrm{N}\times 1], \mathrm{dU}[\gamma+\mathrm{N}\times 2], \mathrm{dU}[\gamma+\mathrm{N}\times 3], \mathrm{dU}[\gamma+\mathrm{N}\times 4] \} \end{array}$$

**Performance metrics** The performance of the algorithm implementation is evaluated using the following metrics:

- The total execution time.
- The OpenCL context creation time.
- The OpenCL context release time.
- Device memory allocation time.
- Device memory free time.
- The time to transfer data from host to device.
- The time to transfer data from device to host.
- The time to transfer data from device to device.
- The OpenCL kernels total execution time.
- The OpenCL kernels total compilation time.

**Validation mechanism** The benchmark has no built-in validation mechanisms. The benchmark can write the solution to the file, which can be validated by external tools (not supplied).





**Target-specific optimizations** The following target-specific optimization are used in the algorithm implementation:

• Coalescing-friendly memory access (probably CUDA only). In order to achieve coalesced memory access the input data buffer is logically transposed. For CUDA devices this transformation makes the data access in kernel coalesced. In other hand this greatly reduces the data locality. Although for CUDA devices the advantages of the coalesced memory access overrides the disadvantages of the non-localized memory access, this might not be the same for OpenCL. Since OpenCL standard doesn't require the consecutive memory accesses within threads of the same working group to be coalesced, this transformation might cause a performance degradation.

#### **Target-specific optimizations opportunities**

• **Buffer swapping elimination.** Device to device memory coping (*dU* to *dUold*) can be eliminated by either using only one buffer (as described in height level description) or using the logical I/O buffers alternation (see PathFinder target specific optimizations section for more details).



# 2.3.4 Gaussian Elimination

The benchmark uses Gaussian Elimination to solve a linear system of equations  $A\vec{x} = \vec{b}$ :

$$\begin{pmatrix} A_{0,0} & A_{0,1} & \cdots & A_{0,n-1} \\ A_{1,0} & A_{1,1} & \cdots & A_{1,n-1} \\ \vdots & \vdots & \ddots & \vdots \\ A_{n-1,0} & A_{n-1,1} & \cdots & A_{n-1,n-1} \end{pmatrix} \times \begin{pmatrix} x_0 \\ x_1 \\ \vdots \\ x_{n-1} \end{pmatrix} = \begin{pmatrix} b_0 \\ b_1 \\ \vdots \\ b_{n-1} \end{pmatrix}$$
(2.17)

#### **High-level description**

**Abstract data structures** The benchmark operates on the following floating point data objects:

- A is an input  $n \times n$  matrix of coefficients.
- *M* is an intermediate  $n \times n$  matrix of multipliers.
- $\vec{b}$  is an input *n*-element vector of constant terms.
- $\vec{x}$  is an output *n*-element vector of solutions.

**Computation** The algorithm consists of two steps:

1. Forward Elimination: Matrix A is reduced to triangular form by elementary rows transformations using intermediate matrix M; the same transformations are applied to vector  $\vec{b}$  to preserve the original solution:  $A\vec{x} = \vec{b} \Rightarrow A'\vec{x} = \vec{b}'$ .

$$\begin{pmatrix} A'_{0,0} & A'_{0,1} & \cdots & A'_{0,n-1} \\ 0 & A'_{1,1} & \cdots & A'_{1,n-1} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & A'_{n-1,n-1} \end{pmatrix} \times \begin{pmatrix} x_0 \\ x_1 \\ \vdots \\ x_{n-1} \end{pmatrix} = \begin{pmatrix} b'_0 \\ b'_1 \\ \vdots \\ b'_{n-1} \end{pmatrix}$$
(2.18)

for k in 0: (n-2) do  $M_{i,k} = A_{i,k}/A_{k,k}, \quad k < i < n$   $A_{i,j} = A_{i,j} - M_{i,k} \times A_{k,j}, \quad k < i < n, k \le j < n$   $b_i = b_i - M_{i,k} \times b_k, \quad k < i < n$ end for

M: Compute multiplier matrix
 A: Modify coefficient matrix
 b: Modify constant term vector

This step requires  $O(n^3)$  operations.

2. Back Substitution: Triangular system (2.18) is solved from  $x_{n-1}$  to  $x_0$ :

This step requires  $O(n^2)$  operations.



for i in (n-1):0 do  $x_i = \left(b'_i - \sum_{i < j < n} A'_{i,j} x_j\right) / A'_{i,i}$  end for

Iteration spaces and dependences This algorithm requires several iteration spaces:

• Compute multiplier matrix:

$$I_M = \{(k,i) : 0 \le k < n-1, k < i < n\}$$

$$I_M: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 1 & -1 \end{pmatrix} \times \binom{k}{i} \le \binom{n-2}{n-1} \\ 0\\ -1 \end{pmatrix}$$

• Modify coefficient matrix:

$$I_A = \{(k, i, j) : 0 \le k < n - 1, k < i < n, k \le j < n\}$$

$$I_A: \begin{pmatrix} 1 & 0 & 0\\ 0 & 1 & 0\\ 0 & 0 & 1\\ -1 & 0 & 0\\ 1 & -1 & 0\\ 1 & 0 & -1 \end{pmatrix} \times \begin{pmatrix} k\\ i\\ j \end{pmatrix} \le \begin{pmatrix} n-2\\ n-1\\ n-1\\ -1\\ -1\\ 0 \end{pmatrix}$$

• Modify constant term vector:

$$I_b = \{(k,i) : 0 \le k < n-1, k < i < n\}$$

$$I_b: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 1 & -1 \end{pmatrix} \times \begin{pmatrix} k\\ i \end{pmatrix} \le \begin{pmatrix} n-2\\ n-1\\ 0\\ -1 \end{pmatrix}$$

• Compute solution (Back Substitution):

$$I_x = \{(i, j) : 0 \le i < n, i < j < n\}$$

$$I_{x}: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 1 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j \end{pmatrix} \leq \begin{pmatrix} n-1\\ n-1\\ 0\\ -1 \end{pmatrix}$$

CARP-ARM-RP-001-v1.0



$$I = I_A \cup I_M \cup I_b \cup I_x$$

The computation generates the following dependences:

$$\begin{split} \delta_{I_A} &= I_M : (j,i) \xrightarrow{a} I_A : (k,i,j) \land \\ &I_M : (i,(i+1:n)) \xrightarrow{a} I_A : (k,i,i) \land \\ &I_M : (k,i) \xrightarrow{t} I_A : (k,i,j) \land \\ &I_A : (k-1,i,j) \xrightarrow{t} I_A : (k,i,j) \land \\ &I_A : (k-1,k,j) \xrightarrow{t} I_A : (k,i,j) \land \\ &I_A : (k-1,i,j) \xrightarrow{a} I_A : (k,i,j) \land \\ &I_A : (k-1,i,j) \xrightarrow{t} I_b : (k,i) \land \\ &I_b : (k-1,i) \xrightarrow{t} I_b : (k,i) \land \\ &I_b : (k-1,i) \xrightarrow{a} I_b : (k,i) \land \\ &I_b : (k-1,i) \xrightarrow{t} I_b : (k,i) \land \\ &I_b : (k-1,i) \xrightarrow{t} I_b : (k,i) \land \\ &I_b : (k-1,i) \xrightarrow{t} I_b : (k,i) \land \\ &I_b : (k-1,i) \xrightarrow{t} I_b : (k,i) \land \\ &I_b : (k-1,i) \xrightarrow{t} I_a : (i,j) \land \\ &I_A : (i,i,j) \xrightarrow{t} I_x : (i,j) \land \\ &I_b : (i,i+1:n) \xrightarrow{t} I_x : (i,j) \land \end{split}$$

#### Memory access mapping

• Compute multiplier matrix: for  $(k,i) \in I_M$ :

$$egin{array}{rll} M_r(k,i) &=& \{A_{i,k},A_{k,k}\} \ M_w(k,i) &=& \{M_{i,k}\} \end{array}$$

• Modify coefficient matrix: for  $(k, i, j) \in I_A$ :

$$egin{array}{rll} M_r(k,i,j) &=& \{A_{k,j},A_{i,j},M_{i,k}\} \ M_w(k,i,j) &=& \{A_{i,j}\} \end{array}$$

• Modify constant vector: for  $(k, i) \in I_b$ :

$$M_r(k,i) = \{b_k, b_i, M_{i,k}\}$$
  
$$M_w(k,i) = \{b_i\}$$

• Compute solution (Back Substitution): for  $(i, j) \in I_x$ :

$$M_r(i,j) = \{x_j, A_{i,j}\}$$
  
 $M_r(i) = \{b_i, A_{i,i}\}$   
 $M_w(i,j) = \{x_i\}$ 





#### Low-level implementation details

Only Forward Elimination (step 1) is implemented using OpenCL. Back Substitution (step 2) is implemented using a sequential host loop.

**Device code** Each kernel launch performs one iteration of forward elimination. See Algorithms 42 and 43.

Algorithm 42 The Gaussian Elimination (Coefficient computations) device code abstraction.

```
Input: global float dA[]

Input: int n

Input: int k

Output: global float dM[]

1: kernel GEKERNEL1 (dM, dA, n, k)

2: i \leftarrow \gamma + k + 1

3: if i < n then

4: dM[n \times i + k] \leftarrow dA[n \times i + k]/dA[n \times k + k]

5: end if

6: end kernel
```

```
Algorithm 43 The Gaussian Elimination (Forward Elimination) device code abstraction.
Input/Output: global float dA[]
Input/Output: global float db[]
Input: global float dM[]
Input: int n
Input: int k
  1: kernel GEKERNEL2 (dM, dA, db, n, k)
 2:
           i \leftarrow \gamma_0 + k + 1
 3:
           j \leftarrow \gamma_1 + k
           if (i < n) and (j < n) then
 4:
                \mathtt{dA}[\mathtt{n}\times\mathtt{i}+\mathtt{j}] \gets \mathtt{dA}[\mathtt{n}\times\mathtt{i}+\mathtt{j}] - \mathtt{dA}[\mathtt{n}\times\mathtt{k}+\mathtt{j}] \times \mathtt{dM}[\mathtt{n}\times\mathtt{i}+\mathtt{k}]
 5:
                if \gamma_1 = 0 then
 6:
                     db[i] \leftarrow db[i] - dM[n \times i + k] \times db[k]
  7:
                end if
 8.
           end if
 9:
10: end kernel
```

**Host code** The outer loop is executed on host, while the inner loops are executed as kernels on device. See Algorithm 44.

## Host data structures

- float  $hA[n * n]: A_{y,x} \rightarrow hA[y * n + x]$
- float  $hM[n * n]: M_{y,x} \rightarrow hM[y * n + x]$



Algorithm 44 The Gaussian Elimination (Forward Elimination) host code abstraction. ▷ Allocate host memory and device buffers. 1: **procedure** ALLOCATEMEMORY(n)  $hA \leftarrow ALLOCATEHOSTMEMORY(n^2)$ 2:  $hM \leftarrow ALLOCATEHOSTMEMORY(n^2)$ 3:  $hb \leftarrow ALLOCATEHOSTMEMORY(n)$ 4:  $hx \leftarrow ALLOCATEHOSTMEMORY(n)$ 5:  $dA \leftarrow AllocateBuffer(n^2, READ_WRITE)$ 6:  $dM \leftarrow AllocateBuffer(n^2, READ_WRITE)$ 7:  $db \leftarrow ALLOCATEBUFFER(n, READ_WRITE)$ 8: 9: end procedure 10: procedure INITIALIZEMEMORY(file) ▷ Initialize host memory.  $hA \leftarrow READFROMFILE(file, n^2)$ 11:  $\texttt{hb} \gets \texttt{READFROMFILE}(\texttt{file},\texttt{n})$ 12: 13: end procedure 14: **procedure** BACKSUBSTITUTION(hA, hb, hx, n) for i in (n-1): 0 do 15:  $hx[i] \leftarrow hb[i]$ 16: for j in (n-1): (i+1) do 17:  $hx[i] \leftarrow hx[i] - hA[i \times n + j] * hx[j]$ 18: end for 19:  $hx[i] \leftarrow hx[i]/hA[i \times n+i]$ 20: end for 21: 22: end procedure Input: int n ▷ Coefficient matrix number of rows (columns). Input: const char \* file ▷ File with pre-generated coefficient matrix and constant term. 23: ALLOCATEMEMORY(n) 24: INITIALIZEMEMORY(file) 25: geProgram  $\leftarrow$  BUILDPROGRAM("geProgram.cl") 26: geKernel1  $\leftarrow$  BUILDKERNEL(geProgram, "geKernel1") 27: geKernel2 ← BUILDKERNEL(geProgram, "geKernel2") 28: COPYTODEVICE(dA, hA, **True**, ∅) 29: COPYTODEVICE(db, hb, **True**, ∅) 30: COPYTODEVICE(dM, hM, **True**, ∅) ▷ In-order execution model is used, so no additional synchronization required. 31: for k in 0 : (n-2) do SetKernelArguments(geKernel1,dM,dA,n,k)32: 33: ENQUEUEKERNEL(geKernel1, (n), none,  $\emptyset$ ) 34: SETKERNELARGUMENTS(geKernel2, dM, dA, db, n, k) ENQUEUEKERNEL(geKernel2, (n, n), none, ∅) 35: 36: end for 37: COPYTOHOST( $dA, hA, True, \emptyset$ ) 38: COPYTOHOST(db, hb, **True**,  $\emptyset$ ) 39: COPYTOHOST(dM, hM, **True**, ∅) 40: BACKSUBSTITUTION(hA, hb, hx, n)





- float  $hb[n]: b_y \to hb[y]$
- float  $hx[n]: x_y \to hx[y]$

#### **Device data structures**

- global float  $dA[n * n]: A_{y,x} \rightarrow dA[y * n + x]$
- global float  $dM[n * n]: M_{y,x} \rightarrow dM[y * n + x]$
- global float  $db[n]: b_y \rightarrow db[y]$
- global float  $dx[n]: x_y \rightarrow dx[y]$

**Input datasets** The benchmark uses pre-generated matrices of single precision floating point numbers in range [-0.9, 0.9] with the following values of *n*:  $\{3, 4, 16, 208, 1024\}$ . It also contains a script for generating matrices and vectors for any *n*.

### **Partitioned iteration space**

• Intermediate multiplier matrix calculation (geKernel1):

$$P_M = \{(k, \gamma) : 0 \le k < n, 0 \le \gamma < (\Gamma = n) - k - 1\}$$

$$f(k,i) \to (k,i-k-1)$$

$$f^{-1}(k,\gamma) \to (k,k+\gamma+1)$$

• Coefficient matrix and constant vector modification (geKernel2):

$$P_{A} = \{(k, \gamma_{0}, \gamma_{1}) : 0 \le k < n, 0 \le \gamma_{1} < (\Gamma_{0} = n) - k - 1, 0 \le \gamma_{1} < (\Gamma_{1} = n) - k\}$$

$$f(k,i,j) \to (k,i-k-1,j-k)$$

$$f^{-1}(k, \gamma_0, \gamma_1) \rightarrow (k, k + \gamma_0 + 1, k + \gamma_1)$$

• Back Substitution (host code):

$$P_x = \{(i, j) : 0 \le i < n, i < j < n\}$$
$$f(i, j) \to (i, j)$$
$$f^{-1}(i, j) \to (i, j)$$

CARP-ARM-RP-001-v1.0





**Device memory access mapping** Each work-item executing the geKernel1 performs the following reads and writes:

$$\begin{array}{lll} G_r(\mathbf{k}, \gamma) &=& \{ \mathtt{dA}[\mathbf{n} \times (\gamma + \mathbf{k} + 1) + \mathbf{k}], \mathtt{dA}[\mathbf{n} \times \mathbf{k} + \mathbf{k}] \} \\ G_w(\mathbf{k}, \gamma) &=& \{ \mathtt{dM}[\mathbf{n} \times (\gamma + \mathbf{k} + k) + \mathbf{k}] \} \end{array}$$

Each work-item executing the geKernel2 performs the following reads and writes:

**Host memory access mapping** On each iteration on the host the following memory accesses are performed:

**Performance metrics** The performance is evaluated using the following metrics:

- The total time for all kernel launches (computation time).
- The total memory transfer time from host to device.
- The total memory transfer time from device to host.

The performance is evaluated only for the part of the algorithm implemented in OpenCL (Forward Elimination).

**Validation mechanism** The benchmark has no built-in validation mechanisms. The benchmark can print the transformed matrices and the solution for the original system, which can be validated by external tools (not supplied).

**Target-specific optimizations** The following target-specific optimization are used in the algorithm implementation:

• **Consecutive memory access.** Work items with consecutive global ids in dimension 1 access consecutive memory locations which improves the memory system performance (caching, coalescing).

CARP-ARM-RP-001-v1.0





## Target-specific optimizations opportunities

• The intermediate multiplier array *M*[][] can be reduced to one dimensional array *m*[], so that this algorithm:

for k in 0: (n-2) do  $M_{i,k} = A_{i,k}/A_{k,k}$  k < i < n  $A_{i,j} = A_{i,j} - A_{k,j} \times M_{i,k}$   $k < i < n, k \le j < n$   $b_i = b_i - M_{i,k} \times b_k$  k < i < nend for

becomes:

for k in 0: (n-2) do  $m_i = A_{i,k}/A_{k,k}$  k < i < n  $A_{i,j} = A_{i,j} - m_i \times A_{k,j}$   $k < i < n, k \le j < n$   $b_i = b_i - m_i \times b_k$  k < i < nend for

This transformation decreases total memory consumption by almost 50% and increases locality.



# 2.3.5 Heart Wall

# **High-level description**

The benchmark tracks the movements of the mouse heart walls on the series of ultrasound images. See [21] for more details.

Abstract data structures The benchmark operates on the following data objects:

- *S* is an input  $N_f$  video sequence (array of frames).
- $S_i$  is an input  $N_c \times N_r$  frame.
- D is an input  $N_{endo}$  element array of endo-cardial surfaces coordinates.
  - $D_i^r$  denotes row coordinate of the  $i^{th}$  surface.
  - $D_i^c$  denotes column coordinate of the  $i^{th}$  surface.
- P is an input  $N_{epi}$  element array of epi-cardial surfaces coordinates (same as D).
- D' is an intermediate  $N_{endo} \times N_f$  array of endo-cardial surfaces coordinates.
- P' is an intermediate  $N_{epi} \times N_f$  array of endo-cardial surfaces coordinates.
- DT is intermediate  $(N_{\text{DT}}) \times N_{\text{endo}}$  buffer of floating point values.
- PT is intermediate  $(N_{\text{PT}}) \times N_{\text{epi}}$  buffer of floating point values.
- IN2 is intermediate  $(N_{IN2}) \times (N_{epi} + N_{endo})$  buffer of floating point values.
- C is intermediate  $(N_{\rm C}) \times (N_{\rm epi} + N_{\rm endo})$  buffer of floating point values.
- IN2P is intermediate  $(N_{\text{IN2P}}) \times (N_{\text{epi}} + N_{\text{endo}})$  buffer of floating point values.
- IN2S is intermediate  $(N_{\text{IN2S}}) \times (N_{\text{epi}} + N_{\text{endo}})$  buffer of floating point values.
- IN2SC is intermediate  $(N_{\text{IN2SC}}) \times (N_{\text{epi}} + N_{\text{endo}})$  buffer of floating point values.
- IN2SCS is intermediate  $(N_{\text{IN2SCS}}) \times (N_{\text{epi}} + N_{\text{endo}})$  buffer of floating point values.
- IN2SQR is intermediate  $(N_{\text{IN2SQR}}) \times (N_{\text{epi}} + N_{\text{endo}})$  buffer of floating point values.
- IN2SQR2 is intermediate  $(N_{\text{IN2SQR2}}) \times (N_{\text{epi}} + N_{\text{endo}})$  buffer of floating point values.
- INSQR is intermediate  $(N_{\text{INSQR}}) \times (N_{\text{epi}} + N_{\text{endo}})$  buffer of floating point values.
- M is intermediate  $(N_{\rm M}) \times (N_{\rm epi} + N_{\rm endo})$  buffer of floating point values.
- + MC is intermediate  $(N_{\rm MC}) \times (N_{\rm epi} + N_{\rm endo})$  buffer of floating point values.





Computation The algorithm tracks the movements of the surfaces: GENERATETEMPLATE( $S_0$ ) for F in S do  $\triangleright$  For each frame. for d in D do  $\triangleright$  For each inner point (Endo-cardial surface). TRACKINNERPOINTMOVEMENT(F, d) end for for p in P do  $\triangleright$  For each outer point (Epi-cardial surface). TRACKOUTERPOINTMOVEMENT(F, p) end for end for

All intermediate buffers are used within initialization (GenerateTemplate) and tracking (TrackInnerPointMovement, TrackOuterPointMovement) procedures.

#### Low-level implementation details

The benchmark device code consists of a massive kernel (about 1.2k loc). The iteration space of the kernel is quite complex (20 loop nests with depth from 1 to 4). The further analysis briefly covers the host/device data structures.

#### Host data structures

- int hS[Nx \* Ny]:  $S_{k,i,j} \rightarrow$  hS[i \* Nx + j]
- int hDC[Nendo]:  $D_i^c \rightarrow$  hDC[i]
- int hDR[Nendo]:  $D_i^r \rightarrow hDR[i]$
- int hPC[Nepi]:  $D_i^c \rightarrow$  hPC[i]
- int hPR[Nepi]:  $D_i^r \rightarrow hPR[i]$
- int hDCT[Nendo \* Nf]:  $D'_{k,i}^c \rightarrow$  hDCT[i \* Nf + k]
- int hDRT[Nendo \* Nf]:  $D'_{k,i}^r \rightarrow \text{hDRT[i * Nf + k]}$
- int hPCT[Nepi \* Nf]:  $P'^c_{k,i} \rightarrow$  hPCT[i \* Nf + k]
- int hPRT[Nepi \* Nf]:  $P'^r_{k,i} \rightarrow$  hPRT[i \* Nf + k]

#### **Device data structures**

- int dS[Nx \* Ny]:  $S_{k,i,j} \rightarrow dS[i * Nx + j]$
- int dDC[Nendo]:  $D_i^c \rightarrow dDC[i]$
- int dDR[Nendo]:  $D_i^r \rightarrow dDR[i]$
- int dPC[Nepi]:  $D_i^c \rightarrow dPC[i]$
- int dPR[Nepi]:  $D_i^r \rightarrow dPR[i]$
- int dDCT[Nendo \* Nf]:  $D'_{k,i}^c \rightarrow dDCT[i * Nf + k]$



- int dDRT[Nendo \* Nf]:  $D'_{k,i}^r \rightarrow dDRT[i * Nf + k]$
- int dPCT[Nendo \* Nf]:  $P'_{k,i}^c \rightarrow dPCT[i * Nf + k]$
- int dPRT[Nendo \* Nf]:  $P'_{k,i}^r \rightarrow dPRT[i * Nf + k]$
- int dC[Nc  $\star$  (Nendo + Nepi)]:  $C_i \rightarrow dC[i]$
- int dIN2P[Nin2p \* (Nendo + Nepi)]:  $IN2P_i \rightarrow dIN2P[i]$
- int dIN2S[Nin2s \* (Nendo + Nepi)]:  $IN2S_i \rightarrow dIN2S[i]$
- int dIN2SC[Nin2sc  $\star$  (Nendo + Nepi)]: IN2SC<sub>i</sub>  $\rightarrow$  dIN2SC[i]
- int dIN2SCS[Nin2scs \* (Nendo + Nepi)]:  $IN2SCS_i \rightarrow dIN2SCS[i]$
- int dIN2SQR[Nin2sqr \* (Nendo + Nepi)]:  $IN2SQR_i \rightarrow dIN2SQR[i]$
- int dIN2SQR2[Nin2sqr2 \* (Nendo + Nepi)]:  $IN2SQR2_i \rightarrow dIN2SQR2[i]$
- int dINSQR[Ninsqr \* (Nendo + Nepi)]:  $INSQR_i \rightarrow dINSQR[i]$
- int dM[Nm \* (Nendo + Nepi)]:  $M_i \rightarrow dM[i]$
- int dMC[Nmc \* (Nendo + Nepi)]:  $MC_i \rightarrow dMC[i]$

**Input datasets** The benchmark uses pre generated input data set with following parameters:

- Number of frames to process  $(N_f)$ . Can be adjusted via command line. The default value is 5.
- Size of the frame  $(N_c \times N_r) = (656 \times 744)$ .
- Number of endo-cardial surfaces  $(N_{endo}) = 20$ .
- Number of epi-cardial surfaces  $(N_{epi}) = 31$ .

**Performance metrics** The performance of the algorithm implementation is evaluated using the following metrics:

- Total time to read input data from file.
- Total time for OpenCL initialization.
- Total time for Device memory allocation and Host-Device memory transfer.
- Total time for all Heart Wall kernel launches.
- Total time for Device memory deallocation.

**Validation mechanism** The benchmark has no validation mechanisms.





# 2.3.6 Hot Spot

The benchmark estimates the temperature of a processor based on its two-dimensional floorplan and simulated power measurements using a system of differential equations [13].

#### **High-level description**

Abstract data structures The benchmark operates on the following data objects:

- *P* is an input  $N \times M$  matrix of dissipated power values.
- $T^0$  is an input  $N \times M$  matrix of initial temperature values.
- $T^k$  is an intermediate  $N \times M$  matrix of average temperature values on  $k^{\text{th}}$  time step  $(0 \le k < k_{\text{max}})$ .
- $T^{k_{\text{max}}}$  is an output  $N \times M$  matrix of average temperature values.

**Computation** The algorithm iterates over the temperature grid:

$$\begin{array}{ll} T_{i,j}^{k+1} & = & T_{i,j}^k + S_{\texttt{div}} \times (P_{i,j} + (T_{i+1,j}^k + T_{i-1,j}^k - 2 \times T_{i,j}^k) \times R_y^{-1} \\ & & + (T_{i,j+1}^k + T_{i,j-1}^k - 2 \times T_{i,j}^k) \times R_x^{-1} + (T_{\texttt{base}} - T_{i,j}^k) \times R_z^{-1} ) \end{array}$$

The following constants are used:

 $T_{\text{base}} = 80.0$  Ambient temperature, assuming no package at all  $C_t = 0.0005$  Chip parameter  $C_f = 0.5$  Capacitance fitting factor  $C_h = 0.016$  Chip height  $C_w = 0.016$  Chip width  $G_h = C_h/N$  Grid step height  $G_w = C_w/M$  Grid step width  $K_{si} = 100$  $S_h = 1.75 \times 10^6$  $PD_{max} = 3 \times 10^6$  Maximum power density  $S_{\max} = PD_{\max}/(C_f \times C_t \times S_h)$  $Cap = C_f \times S_h \times C_t \times G_w \times G_h$  $S_{\text{div}} = 10^{-3}/(S_{\text{max}} \times \text{Cap})$  $R_x = G_w/(2 \times K_{si} \times C_t \times G_h)$  $R_v = G_h / (2 \times K_{si} \times C_t \times G_w)$  $R_{z} = C_t / (K_{si} \times G_h \times G_w)$ 





**Iteration spaces and dependences** On each iteration the algorithm updates the temperature values for each cell in the grid:

$$I = \{(k, i, j) : 0 \le k < k_{\max}, 0 \le i < N, 0 \le j < M\}$$

$$I: \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix} \times \begin{pmatrix} k \\ i \\ j \end{pmatrix} \le \begin{pmatrix} K_{\max} - 1 \\ N - 1 \\ M - 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

The computation generates the following dependences:

$$\begin{split} \delta_I &= I: (k-1,i,j) \xrightarrow{t} I: (k,i,j) \land \\ I: (k-1,i-1,j) \xrightarrow{t} I: (k,i,j) \land \\ I: (k-1,i+1,j) \xrightarrow{t} I: (k,i,j) \land \\ I: (k-1,i,j-1) \xrightarrow{t} I: (k,i,j) \land \\ I: (k-1,i,j+1) \xrightarrow{t} I: (k,i,j) \land \end{split}$$

**Memory access mapping** For  $(k, i, j) \in I$ :

$$M_{r}(i,j,k) = \{P_{i,j}, T_{i,j}^{k}, T_{i-1,j}^{k}, T_{i+1,j}^{k}, T_{i,j-1}^{k}, T_{i,j+1}^{k}\}$$
$$M_{w}(i,j,k) = \{T_{i,i}^{k+1}\}$$

#### Low-level implementation details

**Device code** Each kernel launch performs S time steps on the entire grid. See Algorithm 45.

Host code See Algorithm 46.

Host data structures

- float hP[N \* M]:  $P_{i,j} \rightarrow$  hP[i \* M + j]
- float hT[N \* M]:  $T^k_{i,j} \rightarrow$  hT[i \* M + j]

#### **Device data structures**

- global float dP[N \* M]:  $P_{i,j} \rightarrow dP[i * M + j]$
- global float dTin[N \* M]:  $T^{k+1}_{i,j} \rightarrow \mathrm{dTin}\,[\mathrm{i}~*~\mathrm{M}~+~\mathrm{j}]$
- global float dTout[N \* M]:  $T^k_{i,j} \rightarrow \text{dTout[i * M + j]}$



```
Algorithm 45 The Hot Spot device code abstraction.
Input: int Iter
                                                                                             \triangleright Number of iterations.
Input: global float dP[]
                                                                            ▷ Power value for each cell in a grid.
Input: global float dTin[]
                                                                                        ▷ Initial temperature values.
Output: global float dTout[]
                                                                                   ▷ Resulting temperature values.
Input: int N
                                                                                     \triangleright Number of rows in the grid.
Input: int M
                                                                                ▷ Number of columns in the grid.
Input: int Br
                                                                                          \triangleright Number of border rows.
                                                                                     \triangleright Number of border columns.
Input: int Bc
Input: float Cap
                                                                                      ▷ Device parameter constant.
Input: float Rx
                                                                                      ▷ Device parameter constant.
Input: float Ry
                                                                                      ▷ Device parameter constant.
Input: float Rz
                                                                                      ▷ Device parameter constant.
                                                                                      ▷ Device parameter constant.
Input: float Step
Local: local float dlP [\Lambda_1] [\Lambda_0]
                                                                                ▷ Temporary power values buffer.
Local: local float dlTin[\Lambda_1][\Lambda_0]
                                                                         ▷ Temporary input temperature values.
Local: local float dlTout [\Lambda_1] [\Lambda_0]
                                                                        ▷ Temporary output temperature values.
  1: kernel HSKERNEL (Iter, dP, dTin, dTout, N, M, Br, Bc, Cap, Rx, Ry, Rz, Step)
 2:
          \texttt{Sdiv} \leftarrow \texttt{Step}/\texttt{Cap}
          \texttt{Tbase} \leftarrow 80
 3:
 4:
          \texttt{smallBlockCols} \leftarrow \Lambda_0 - 2 \times \texttt{Iter}
          \texttt{smallBlockRows} \leftarrow \Lambda_1 - 2 \times \texttt{Iter}
 5:
          \texttt{blkXmin} \leftarrow v_0 \times \texttt{smallBlockCols} - \texttt{Bc}
 6:
 7:
          \texttt{blkYmin} \leftarrow v_1 \times \texttt{smallBlockRows} - \texttt{Br}
 8:
          \texttt{blkXmax} \leftarrow \texttt{blkX} + \Lambda_0 - 1
 9:
          \texttt{blkYmax} \gets \texttt{blkY} + \Lambda_1 - 1
10:
          xidx \leftarrow blkXmin + \lambda_0
          yidx \leftarrow blkYmin + \lambda_1
11:
          index \leftarrow \mathbb{N} \times \text{yidx} + \text{xidx} > \text{Probably a bug, the global row number yidx should be}
12:
     multiplied by the number of columns (M) rather then number of rows (N).
13:
          if (yidx in (0: (N-1))) \land (xidx in (0: (M-1))) then
               dlP[\lambda_1][\lambda_0] \leftarrow dP[index]
14:
               dlTin[\lambda_1][\lambda_0] \leftarrow dTin[index]
15:
          end if
16:
          BARRIER(CLK_LOCAL_MEM_FENCE)
17:
18:
          X\min \leftarrow \max(0, -\texttt{blk}X\min)
          Ymin \leftarrow max(0, -blkYmin)
19:
          \texttt{Xmax} \gets \min(\Lambda_0 - 1, \Lambda_0 - 1 - (\texttt{blkXmax} - \texttt{N} + \texttt{1}))
20:
          \texttt{Ymax} \gets \min(\Lambda_1 - 1, \Lambda_1 - 1 - (\texttt{blkYmax} - \texttt{N} + \texttt{1}))
21:
          \mathbb{N} \leftarrow \max(\mathtt{Ymin}, \lambda_1 - 1)
22:
          \mathtt{S} \gets \min(\Lambda_1 + 1, \mathtt{Ymax})
23:
          W \leftarrow \max(\texttt{Xmin}, \lambda_0 - 1)
24:
          \texttt{E} \leftarrow \min(\Lambda_0 + 1, \texttt{Xmax})
25:
```



Algorithm 45 The Hot Spot device code abstraction (continues).	
26: <b>for</b> i in 0: $(Iter - 1)$ do	
27: $isComputed \leftarrow false$	
28: <b>if</b> $(\lambda_0 \text{ in } (i+1) : \Lambda_0 - i - 2) \land (\lambda_0 \text{ in } \text{Xmin} : \text{Xmax}) \land (\lambda_1 \text{ in } (i+1) : \Lambda_1 - i - 2)$	$\wedge$
$(\lambda_1 \text{ in } \mathtt{Ymin}: \mathtt{Ymax})$ then	
29: isComputed $\leftarrow$ true	
30: $dlTout[\lambda_1][\lambda_0] \leftarrow dlTin[\lambda_1][\lambda_0] + Sdiv \times (dlP[\lambda_1][\lambda_0] + (dlTin[S][\lambda_0])$	+
$\texttt{dlTin}[\texttt{N}][\lambda_0] - 2 \times \texttt{dlTin}[\lambda_1][\lambda_0]) \times \texttt{Ry}^{-1} + (\texttt{dlTin}[\lambda_1][\texttt{E}] + \texttt{dlTin}[\lambda_1][\texttt{W}] - 2$	$\times$
$\mathtt{dlTin}[\lambda_1][\lambda_0]) \times \mathtt{Rx}^{-1} + (\mathtt{Tbase} - \mathtt{dlTin}[\lambda_1][\lambda_0]) \times \mathtt{Rz}^{-1}) \qquad \qquad \triangleright  \mathtt{See}  (2.19)$	9).
31: end if	
32: BARRIER(CLK_LOCAL_MEM_FENCE)	
33: <b>if</b> $i = Iter - 1$ <b>then</b>	
34: break	
35: <b>end if</b>	
36: if isComputed then	
37: $dlTin[\lambda_1][\lambda_0] \leftarrow dlTout[\lambda_1][\lambda_0]$	
38: end if	
39: BARRIER(CLK_LOCAL_MEM_FENCE)	
40: <b>end for</b>	
41: end kernel	

where s is the number of time steps performed by a single kernel launch.

**Input datasets** The benchmark uses a pre-generated  $512 \times 512$  initial temperature and power values grids. The total number of iterations  $K_{\text{max}}$  and the number of iterations processed by a single kernel launch *S* can be specified via the command line. By default  $K_{\text{max}} = 60$  and S = 1.

# **Partitioned iteration space**

$$P = \{(r, v_1, v_0, \lambda_1, \lambda_0, i) : 0 \le r < K_{\max}, 0 \le v_1 < N_1, 0 \le v_0 < N_0, 0 \le \lambda_1 < \Lambda_1, 0 \le \lambda_0 < \Lambda_0\}$$

$$N_1 = N$$
$$N_0 = M$$



```
Algorithm 46 The Hot Spot host code abstraction
 1: procedure AllocateMemory(N,M)
                                                        ▷ Allocate host memory and device buffers.
 2:
         hP \leftarrow AllocateHostMemory(N \times M)
 3:
         hT \leftarrow AllocateHostMemory(N \times M)
         dP \leftarrow ALLOCATEBUFFER(N \times M, READ_ONLY, hP)
 4:
         dTin \leftarrow ALLOCATEBUFFER(N \times M, READ_ONLY, hT)
 5:
         dTout \leftarrow ALLOCATEBUFFER(N \times M, READ_ONLY)
 6:
         hTt[0] \leftarrow dTin
 7:
         \mathtt{hTt}[1] \gets \mathtt{dTout}
 8:
 9: end procedure
10: procedure INITIALIZEMEMORY(N,M,file)
                                                                             ▷ Initialize host memory.
         hP \leftarrow READFROMFILE(file, N \times M)
11:
         hT \leftarrow READFROMFILE(file, N \times M)
12:
13: end procedure
Input: int N
                                                                     ▷ Number of rows in input array.
Input: int M
                                                                 ▷ Number of columns in input array.
Input: int Imax
                                                                               ▷ Number of iterations.
Input: int Iter > Number of iterations for each kernel launch (see Algorithm 45 for more
    details).
Input: const char * file
                                                                       ▷ File with pre-generated data.
14: size \leftarrow n \times m
15: AllocateMemory(N,M)
16: INITIALIZEMEMORY(N,M,file)
17: hsProgram ← BUILDPROGRAM("hsProgram.cl")
18: hsKernel ← BUILDKERNEL(hsProgram, "hsKernel")
19: src, ret \leftarrow 0, 1
20: for r in 0: (Imax - 1): Iter do
         it \leftarrow min(Iter, Imax - r - 1)
21:
         \texttt{Cap} \leftarrow \texttt{Cap}
                                                                      ▷ See the Computation section.
22:
         \mathtt{Rx} \leftarrow \mathtt{R}_{\mathtt{x}}
                                                                      ▷ See the Computation section.
23:
24:
         Ry \leftarrow R_y
                                                                      ▷ See the Computation section.
         \mathtt{Rz} \leftarrow \mathtt{R_z}
                                                                      ▷ See the Computation section.
25:
         \texttt{Step} \leftarrow 10^{-3}/\texttt{S}_{max}
                                                                      ▷ See the Computation section.
26:
         SETKERNELARGUMENTS(pfKernel, it, dP, hTt[src], hTt[ret], N, M, Iter, Iter,
27:
    Cap, Rx, Ry, Rz, Step)
         ENQUEUEKERNEL(hsKernel, (M/(16 - 2 \times \text{Iter}) \times 16, N/(16 - 2 \times \text{Iter}) \times
28:
    16), (16, 16), \emptyset)
         \texttt{src}, \texttt{ret} \leftarrow \texttt{ret}, \texttt{src}
29:
30: end for
31: COPYTOHOST(hTt[ret],hT, True, ∅)
```



$$\begin{array}{rcl} f(k,i,j) & \to & ((i/S) \times S, i/(\Lambda_1 - S \times 2), j/(\Lambda_0 - S \times 2), \\ & & i \bmod (\Lambda_1 - S \times 2) + S, j \bmod (\Lambda_0 - S \times 2) + S, k \bmod S \end{array}$$

$$f^{-1}(r, v_1, v_0, \lambda_1, \lambda_0, i) = (r + i, v_1 * (\Lambda_1 - S \times 2) - S + \lambda_1, v_0 * (\Lambda_0 - S \times 2) - S + \lambda_0)$$

**Device memory access mapping** Each work-item performs the following reads and writes:

The global memory access offset is probably wrong for non-square matrices (M should be used instead of N as a multiplier).

**Performance metrics** The performance is evaluated using the following metric:

• The total time for all kernel launches (computation time).

Validation mechanism The benchmark has no validation mechanisms.

## **Target-specific optimizations**

- **Blocking.** The input data is processed in blocks having the number of columns equal to the number of work items per work group. Consecutive work items of the same group access consecutive memory addresses which improves the memory subsystem performance (caches, coalescing).
- Ghost zone[14]. The computation of the (i, j)<sup>th</sup> element on the k<sup>th</sup> iteration requires the (i, j − 1)<sup>th</sup>, (i, j + 1)<sup>th</sup>, (i − 1, j)<sup>th</sup>, (i + 1, j)<sup>th</sup> elements from the previous iteration. Since these data can be produced by another work group, the ghost zone technique is used to avoid data races: Each kernel launch performs an S iterations over input data. Given a block with Λ rows and Λ columns the (Λ − S × 2)<sup>2</sup> elements of the block can be computed on the last iteration by a single work group without communicating with other groups. Therefor, the whole grid is divided into non-overlapping blocks with (Λ − 2 × S) rows and (Λ − 2 × S) columns. Each group allocates a local buffer for Λ<sup>2</sup> elements, which is used to compute the resulting (Λ − 2 × S)<sup>2</sup> elements without affecting the other groups' data. This technique allows each work group to compute the result independently, although the elements located in the ghost areas are computed twice.





• **I/O buffers alternation.** The output of the kernel execution on the current iteration is used as kernel input for the next one. To avoid additional memory transfers, the input and output buffers are logically swapped between the kernel launches:

```
...
kernel(buffptr[0], buffptr[1])
kernel(buffptr[1], buffptr[0])
kernel(buffptr[0], buffptr[1])
...
```

The 2-element array buffptr contains the pointers to the device memory buffers.





# 2.3.7 K-means

#### **High-level description**

The benchmark uses k-means clustering to partition a cluster of objects into several sub-clusters based on objects' features.

### Abstract data structures

- *P* is an input  $N_e$  element vector of objects. Each object is an  $N_f$  element vector of features (floating point values).
- *S* is an output set of sub-clusters. Each object from *P* belongs to a single sub-cluster from *S*.
- *C* is an intermediate vector of the centroids.
- Q is an intermediate vector of clusterization costs.

**Computation** The algorithm partitions a cluster of objects into several sub-clusters  $(S = \{S_i\} : K_{\min} \le |S| \le K_{\max})$  based on objects' features. Each object is associated with a sub-cluster, having the nearest centroid, according to the Euclidean distance.

• *C*<sub>A</sub> denotes the centroid of the sub-cluster A:

$$C_A = \frac{\sum_{p \in A} p}{|A|}$$

•  $Q_S$  denotes the quality (cost) of the partitioning:

$$Q_S = \sum_{S_i \in S} \sum_{p \in S_i} \|p - C_{S_i}\|$$

The algorithm tries to determine the minimal-cost partitioning:

$$S = \min_{Q_S}(S^k : \mathbf{K}_{\min} \le k = |S^k| \le \mathbf{K}_{\max})$$

For each partitioning size (number of clusters) k in  $K_{min}$ :  $K_{max}$ , the algorithm starts with a random partitioning of this size and iteratively recalculates partitioning until a fixed point is reached.

- $S_{i,\tau}^k$  denotes the *i*<sup>th</sup> sub-cluster on the  $\tau^{\text{th}}$  iteration.
- $S_i^k$  denotes the *i*<sup>th</sup> sub-cluster when the fixed point is reached. set.
- $S^k = \{S_i^k\}$
- 1. Initialize centroids:

$$C_{S^k_{i,0}} = \texttt{RandomCentroid}(), 0 \leq i < k$$





2. Recalculate clusters:

$$S_{i,\tau}^{k} = \{ p : \| p - C_{S_{i,\tau}^{k}} \| < \| p - C_{S_{j,\tau-1}^{k}} \|, \forall j \neq i \}, 0 \le i < k, 1 \le \tau$$

3. Recalculate centroids:

$$C_{S_{i,\tau}^k} = \frac{\sum_{p \in S_{i,\tau}^k} p}{|S_{i,\tau}^k|}, 0 \le i < k, 1 \le \tau$$

4. Repeat steps 2 and 3 (increasing  $\tau$ ) until a fixed point is reached:

$$S^k = \{S^k_{i, au} : 0 \le i < k, S^k_{i, au} = S^k_{i, au+1}\}$$

5. Calculate the cost:

$$Q_{S_k} = \sum_{S_i^k \in S^k} \sum_{p \in S_i^k} \|p - C_{S_i^k}\|$$

The final partitioning is the one having the minimal cost:

$$S = \min_{Q_{S_k}} S^k$$

### **Iteration spaces**

• Initialize centroids:

$$I_{I} = \{(k, c, f) : K_{\min} \le k \le K_{\max}, c \in C_{0}^{k}, f \in c\}$$

• Recalculate clusters:

$$I_{S} = \{(k, \tau, p, s, f) : K_{\min} \le k \le K_{\max}, 1 \le \tau, p \in F, s \in \{S_{i,\tau}^{k} : 0 \le i < k\}, f \in p\}$$

• Recalculate centroids:

$$I_P = \{(k, \tau, s, p, f) : K_{\min} \le k \le K_{\max}, 1 \le \tau, s \in \{S_{i,\tau}^k : 0 \le i < k\}, p \in s, f \in p\}$$

• Calculate cost:

$$I_C = \{(k, s, p, f) : K_{\min} \le k \le K_{\max}, s \in S^k, p \in s, f \in p\}$$

The  $\tau$  iteration variable represents the repeat-until loop, which terminates when a fixed point is reached.

## Dependences

$$\begin{array}{lll} \delta_{I_{I}} &=& \emptyset \\ \delta_{I_{S}} &=& I_{I}(k,C_{s},f) \xrightarrow{t} I_{S}(k,1,p,s,f) \wedge \\ && I_{P}(k,\tau-1,s,s,f) \xrightarrow{t} I_{S}(k,\tau,p,s,f) \\ \delta_{I_{P}} &=& I_{S}(k,\tau,p,\{S_{i,\tau}^{k}:0\leq i< k\},f) \xrightarrow{t} I_{P}(k,\tau,s,p,f) \\ \delta_{I_{C}} &=& I_{S}(k,\{\tau:\tau>0\},s,p,f) \xrightarrow{t} I_{C}(k,s,p,f) \end{array}$$





## Memory access mapping

• Initialize centroids:

$$M_w(k,c,f) = \{c\}$$

• Recalculate clusters:

$$M_r(k, \tau, p, s, f) = \{p, C_s\}$$
  
 $M_w(k, \tau, p, s, f) = \{s\}$ 

• Recalculate centroids:

$$\begin{array}{lll} M_r(k,\tau,s,p,f) &=& \{p\} \\ M_w(k,\tau,s,p,f) &=& \{C_s\} \end{array}$$

• Calculate cost:

$$M_r(k,s,p,f) = \{p,C_s\}$$
  
$$M_w(k,s,p,f) = \{Q_{S_k}\}$$

#### Low-level implementation details

**Device code** See Algorithms 47 and 48.

**Host code** See Algorithm 49. The algorithm is repeated Niter times for each k.<sup>13</sup>

## Host data structures

- global float hP[Ne][Nf]:  $P_{p,f} \rightarrow$  hP[p][f]
- global int hM[Ne]: For every element p hM[p] denotes the cluster the element belongs to before kernel execution.
- global int hM2[Ne]: For every element p hM[p] denotes the cluster the element belongs to after kernel execution.
- global float  $hL[k]: ||S_{i,\tau}^k|| \to hL[i]$
- global float  $hC[K][Nf]: C_{c,f} \rightarrow hC[c][f]$  before kernel execution.
- global float hC2[K][Nf]:  $C_{c,f} \rightarrow$  hC2[c][f] after kernel execution.

#### **Device data structures**

- global float dP[Ne \* Nf]:  $P_{p,f} \rightarrow dP[p * Nf + f]$
- global float dP2[Ne \* Nf]:  $P_{p,f} \rightarrow dP[f * Np + p]$
- global float dC[K \* Nf]:  $C_{c,f} \rightarrow dC[c * Nf + f]$
- global int dM[Ne]: For every element p dM[p] denotes the cluster the element belongs to.

<sup>&</sup>lt;sup>13</sup>Probably for benchmarking purposes since it's not required by the algorithm.



```
Algorithm 47 The K-Means device code abstraction (Clusters Recalculation).
Input: int Ne
Input: int Nf
Input: int K
                                                                                          \triangleright Number of clusters.
Input: global float dP[]
Input: global float dC[]
Output: int dM[]
 1: kernel KMKERNEL1 (dP, dC, dM, Ne, Nf, K)
          if \gamma < \text{Ne} then
 2:
 3:
              \mathtt{idx} \leftarrow 0
              \texttt{min} \leftarrow \texttt{MAX\_FLOAT}
 4:
 5:
              for i in 0 : (K−1) do
 6:
                   \texttt{ans} \gets 0
                   for f in 0: (Nf - 1) do
 7:
                        ans \leftarrow ans + (dP[f \times Ne + \gamma] - dC[i \times Nf + f])^2
 8:
 9:
                   end for
10:
                   if ans < min then
11:
                       \texttt{min} \leftarrow \texttt{ans}
                        \mathtt{idx} \gets \mathtt{i}
12:
                   end if
13:
14:
              end for
              dM[\gamma] \leftarrow idx
15:
          end if
16:
17: end kernel
```

Algorithm 48 The K-Means device code abstraction (Matrix Transposition).

```
Input: int Ne

Input: int Nf

Input: global float dP[]

Output: global float dP2[]

1: kernel KMKERNEL2 (dP, dP2, Ne, Nf)

2: for f in 0: (Nf - 1) do

3: dP2[f × Ne + \gamma] \leftarrow dP[\gamma × Nf + f]

4: end for

5: end kernel
```



```
Algorithm 49 The K-Means host code abstraction.
 1: procedure ALLOCATEMEMORYINIT(Ne, Nf)
        hP \leftarrow AllocateHostMemory(Ne \times Nf)
 2:
 3:
        hM \leftarrow AllocateHostMemory(Ne)
 4: end procedure
 5: procedure ALLOCATEMEMORYITER(Ne, Nf, K)
        dP \leftarrow AllocateBuffer(Ne \times Nf, READ_WRITE)
 6:
 7:
        dP2 \leftarrow AllocateBuffer(Ne \times Nf, READ_WRITE)
 8:
        dC \leftarrow ALLOCATEBUFFER(K \times Nf, READ_WRITE)
 9:
        dM \leftarrow ALLOCATEBUFFER(Ne, READ_WRITE)
        kmProgram \leftarrow BUILDPROGRAM("kmProgram.cl")
10:
        kmKernel1 \leftarrow BUILDKERNEL(kmProgram, "kmKernel1")
11:
        kmKernel2 \leftarrow BUILDKERNEL(kmProgram, "kmKernel2")
12:
13:
        COPYTODEVICE(dP, hP, True, ∅)
        SETKERNELARGUMENTS(kmKernel2, dP, dP2, Nf, K)
14:
        ENQUEUEKERNEL(kmKernel1, (Ne), none, ∅)
15:
16: end procedure
17: procedure INITIALIZEMEMORY(Ne,Nf,file)
        hP \leftarrow READFROMFILE(file, Ne \times Nf)
18:
19: end procedure
20: function KMEANS(Nf,Ne,K)
21:
        COPYTODEVICE(dC, hC, True, \emptyset)
        SETKERNELARGUMENTS(kmKernel1, dP, dC, dM, Ne, Nf, K)
22:
        ENQUEUEKERNEL(kmKernel1, (Ne), none, ∅)
23:
24:
        COPYTOHOST(dM, hM2, True, ∅)
        \texttt{delta} \gets 0
25:
26:
        for i in 0 : (Ne - 1) do
            \texttt{Cid} \leftarrow \texttt{hM2[i]}
27:
            hL[Cid] \leftarrow hL[Cid] + 1
28:
            if hM2[i] \neq hM[i] then
29:
               \texttt{delta} \gets \texttt{delta} + 1
30:
               hM[i] \leftarrow hM2[i]
31:
32:
            end if
            for j in 0 : (Nf - 1) do
33:
               hC2[Cid][j] \leftarrow hC2[Cid][j] + hP[i][j]
34:
            end for
35:
        end for
36:
        return delta
37:
38: end function
```



```
Algorithm 49 The K-Means host code abstraction (continues).
39: function DISTANCE(P1, P2, S)
40:
          \texttt{sum} \gets 0
          for i in 0 : (S - 1) do
41:
               \texttt{sum} \leftarrow \texttt{sum} + (\texttt{P2[i]} - \texttt{P1[i]})^2
42:
          end for
43:
          return sum
44:
45: end function
46: function FINDNEARESTPOINT(P,Nf,Objs,S)
47:
          \texttt{min} \gets \texttt{FLT}\_\texttt{MAX}
          for i in 0 : (S - 1) do
48:
               dist \leftarrow DISTANCE(P, Objs[i], Nf)
49:
               \min \leftarrow \min(\min, \texttt{dist})
50:
51:
               \texttt{idx} \gets \texttt{argmin}(\texttt{min},\texttt{dist})
          end for
52:
          return idx
53:
54: end function
55: function RMSERROR(Nf, Ne, K)
          \texttt{sum} \gets 0
56:
          for i in 0 : (Ne -1) do
57:
               idx \leftarrow FindNearestPoint(hP[i], Nf, hC, K)
58:
59:
               \texttt{sum} \gets \texttt{sum} + \texttt{DISTANCE}(\texttt{hP}[\texttt{i}],\texttt{hC}[\texttt{idx}],\texttt{Nf})
          end for.State return \sqrt{\frac{\text{sum}}{\text{Ne}}}
60:
61: end function
```



Alge	orithm 49 The K-Means host code abstrac	tion (continues).
Inp	ut: int Niter	▷ Number of iterations (for benchmarking).
Inp	ut: int Ne	▷ Number of elements in the input set.
Inp	ut: int Nf	▷ Number of features in each element.
Inp	ut: int Kmax	▷ Maximal number of clusters.
Inp	ut: int Kmin	▷ Minimal number of clusters.
Inp	ut: const char * file	▷ File with pre-generated elements.
62:	AllocateMemoryInit(Ne,Nf)	
63:	INITIALIZEMEMORY(Ne, Nf, file)	
64:	$\texttt{RmseMin} \gets \texttt{FLT}\_\texttt{MAX}$	
65:	for K in Kmin : Kmax do	
66:	AllocateMemoryIter(Ne,Nf,K)	
67:	for iter in $0$ : (Niter $-1$ ) do	
68:	$hC \leftarrow AllocateHostMemory$	$(\texttt{Nf} \times \texttt{K})$
69:	$\texttt{hC2} \leftarrow \texttt{AllocateHostMemory}$	Y(Nf  imes K)
70:	$\texttt{hL} \leftarrow \texttt{AllocateHostMemory}$	(K)
71:	$\texttt{hM2} \leftarrow \texttt{AllocateHostMemory}$	Y(Ne)
72:	for c in $0$ : $(K-1)$ do	
73:	for f in $0$ : $(Nf - 1)$ do	
74:	$\texttt{hC}[c][f] \gets \texttt{RANDOMFEAT}$	URE()
75:	end for	
76:	end for	
77:	$\texttt{delta} \gets 0$	
78:	$\texttt{threshold} \gets 0.001 \triangleright \texttt{Probably a}$	bug, since delta is less than threshold only when
	delta is 0.	
79:	repeat	
80:	$\texttt{delta} \leftarrow \texttt{KMEANS}(\texttt{Nf}, \texttt{Ne}, \texttt{K})$	
81:	for i in $0$ : $(K-1)$ do	
82:	for j in $0: (Nf-1)$ do	
83:	if $hL[i] > 0$ then	
84:	$\texttt{hC[i][j]} \leftarrow \texttt{hC2[i][j]}$	j]/hL[i]
85:	end if	
86:	$\texttt{hC2[i][j]} \gets 0$	
87:	end for	
88:	$\texttt{hL}[\texttt{i}] \gets 0$	
89:	end for	
90:	until delta < threshold	
91:	$rmse \leftarrow RMSERROR(Nf, Ne, K)$	
92:	if rmse < RmseMin then	
93:	$\texttt{RmseMin} \leftarrow \texttt{rmse}$	
94:	BestK — K	
95:	$\texttt{BestC} \leftarrow \texttt{hC}$	
96:	end if	
97:	end for	
98:	end for	





**Input datasets** The benchmark uses the following pre generated input data sets:

- 100 objects with 34 features.
- 204800 objects with 34 features.
- 494020 objects with 34 features.
- 819200 objects with 34 features.

Each feature is a floating point number.

### **Partitioned iteration space**

• Initialize centroids (Host):

$$P_I = \{(k, i, c, f) : K_{\min} \le k \le K_{\max}, 0 \le i < \texttt{Niter}, 0 \le c < K, 0 \le f < \texttt{Nf}\}$$

$$f(k, c = C_i, f = c_j) \rightarrow (k, 0: (\texttt{Niter} - 1), i, j)$$

$$f^{-1}(k,i,c,f) \rightarrow (k,C_c,c_f)$$

• Recalculate clusters (kmKernel1):

 $P_{Sd} = \{(k, i, \tau, \gamma, c, f) : K_{\min} \leq k \leq K_{\max}, 0 \leq i < \texttt{Niter}, 0 \leq \tau, 0 \leq \gamma < \Gamma = N_e, 0 \leq c < K, 0 \leq f < \texttt{Nf} \}$ 

$$f(k,\tau,p=P_i,s=C_j,f=p_l) \rightarrow (k,0:(\texttt{Niter}-1),\tau,i,j,l)$$

$$f^{-1}(k,i,\tau,\gamma,c,f) \rightarrow (k,\tau,p=P_{\gamma},C_c,p_f)$$

• Recalculate clusters (Host):

$$P_{Sh} = \{(k, i, \tau, p, f) : K_{\min} \le k \le K_{\max}, 0 \le i < \texttt{Niter}, 0 \le \tau, 0 \le p < N_e, 0 \le f < \texttt{Nf}\}$$

$$f(k,\tau,p=P_i,s,f=p_l) \to (k,0:(\texttt{Niter}-1),\tau,i,l)$$

$$f^{-1}(k,i,\tau,e,f) \rightarrow (k,\tau,p=P_e,C_{hM[e]},p_f)$$

• Recalculate centroids (Host):

$$P_P = \{ (k, i, \tau, c, f) : K_{\min} \le k \le K_{\max}, 0 \le i < \texttt{Niter}, 0 \le \tau, 0 \le p < K, 0 \le f < \texttt{Nf} \}$$

$$f(k,\tau,s=C_i,p,f=p_l) \to (k,0:(\texttt{Niter}-1),\tau,i,l)$$

$$f^{-1}(k,i,\tau,c,f) \rightarrow (k,\tau,c,P,c_f)$$

CARP-ARM-RP-001-v1.0





• Calculate cost (Host):

$$P_C = \{(k,i,p,s,f): K_{\min} \leq k \leq K_{\max}, 0 \leq i < \texttt{Niter}, 0 \leq p < N_e, 0 \leq s < K, 0 \leq f < \texttt{Nf}\}$$

$$f(k,\tau,s=C_i,p=P_j,f=p_l) \rightarrow (k,0:(\texttt{Niter}-1),\tau,j,i,l)$$

$$f^{-1}(k,i,\tau,e,c,f) \to (k,\tau,C_{hM[e]},C_c,p=P_e,p_f)$$

#### **Device memory access mapping**

• kmKernel1:

$$\begin{array}{lll} G_r(\texttt{k},\gamma,\texttt{i},\texttt{f}) &=& \{\texttt{dP2}[\texttt{f}\times\texttt{Ne}+\gamma],\texttt{dC}[\texttt{i}\times\texttt{Nf}+\texttt{f}]\}\\ G_w(\texttt{k},\gamma) &=& \{\texttt{dM}[\gamma]\} \end{array}$$

• kmKernel2:

$$\begin{array}{lll} G_r({\tt k},\gamma,{\tt f}) &=& \{{\tt dP}[\gamma\times{\tt Nf}+{\tt f}]\}\\ G_w({\tt k},\gamma,{\tt f}) &=& \{{\tt dP2}[{\tt f}\times{\tt Ne}+\gamma]\} \end{array}$$

### Host memory access mapping

• Initialize centroids:

$$G_w(k,i,c,f) = \{hC[c][f]\}$$

• Recalculate clusters:

• Recalculate centroids:

$$\begin{array}{lll} G_r(k,i,\tau,c,f) &=& \{\mathtt{hL}[\mathtt{c}],\mathtt{hC2}[\mathtt{c}][\mathtt{f}]\} \\ G_w(k,i,\tau,p,f) &=& \{\mathtt{hL}[\mathtt{c}],\mathtt{hC2}[\mathtt{c}][\mathtt{f}],\mathtt{hC2}[\mathtt{c}][\mathtt{f}]\} \end{array}$$

• Calculate cost:

$$G_r(k,i,p,c,f) = \{hP[p][f],hC[c][f]\}$$

**Performance metrics** The performance of the algorithm implementation is evaluated by measuring the total time required to perform clusterization for the given set of objects (both host and device parts).

**Validation mechanism** The benchmark has no built-in validation mechanisms. The benchmark can print the traceback values which can be validated by external tools (not supplied).





## **Target-specific optimizations**

• **Coalescing-friendly memory access (probably CUDA only).** In order to achieve coalesced memory access the input data buffer is logically transposed. For CUDA devices this transformation makes the data access in kernel coalesced. In other hand this greatly reduces the data locality. Although for CUDA devices the advantages of the coalesced memory access overrides the disadvantages of the non-localized memory access, this might not be the same for OpenCL. Since OpenCL standard does not require the consecutive memory accesses within threads of the same working group to be coalesced, this transformation might cause a performance degradation.

### **Target-specific optimizations opportunities**

• Vector operations. The device code can benefit from using vector operations.





# 2.3.8 LavaMD

## **High-level description**

The benchmark calculates particle potential and relocation within a large 3D space based on the mutual forces between particles.

**Abstract data structures** The benchmark operates on particles in a 3D space. The space is divided into boxes. Each box contains  $N_{pb}$  particles. Each non-boundary box is surrounded by 26 neighbor boxes. The number of neighbors for the boundary boxes is smaller (17 for sides and 11 for edges and 7 for corners). Each particle is defined as a quintuple (x, y, z, v, q):

- *x* denotes the *x* coordinate of the particle.
- *y* denotes the *y* coordinate of the particle.
- *z* denotes the *z* coordinate of the particle.
- *v* denotes the velocity of the particle.
- q denotes the charge of the particle.

The benchmark operates the on the following data structures:

- *B* is an  $N_{bx} \times N_{by} \times N_{bz}$  input three dimensional array of boxes. Each box contains  $N_{pb}$  particles.
- *B*<sub>*i*,*j*,*k*</sub> is the *N*<sub>*pb*</sub> element input array of initial configurations of the particles, located in the box with indices (*i*, *j*,*k*) in the array *B*.
- P = (x, y, z, v, q) is the particle configuration.
- $N_b$  is a set of neighbor boxes for the box b including the b.
- $F_{i,j,k}$  is the  $N_{pb}$  element output array of updated configurations of the particles, located in the box with indices (i, j, k) in the array *B*.





**Computation** The new configurations  $F_{i,j,k}$  for the particles in the box  $B_{i,j,k}$  are calculated as follows for each particle  $P_s = (x, y, z, v, q)$  in  $B_{i,j,k}$ 

$$\begin{split} F_{i,j,k} &= \{P'_{s}, 0 \leq s < N_{pb}\} \\ P'_{s} &= (x',y',z',v',q') \\ x' &= \sum_{B' \in N_{B_{i,j,k}}} \sum_{C \in B'} 2 \times e^{-\frac{v + C_{v} - (x,y,z) \cdot (C_{x},C_{y},C_{z})}{2}} \times (x - C_{x}) \times C_{q} \\ y' &= \sum_{B' \in N_{B_{i,j,k}}} \sum_{C \in B'} 2 \times e^{-\frac{v + C_{v} - (x,y,z) \cdot (C_{x},C_{y},C_{z})}{2}} \times (y - C_{y}) \times C_{q} \\ z' &= \sum_{B' \in N_{B_{i,j,k}}} \sum_{C \in B'} 2 \times e^{-\frac{v + C_{v} - (x,y,z) \cdot (C_{x},C_{y},C_{z})}{2}} \times (z - C_{z}) \times C_{q} \\ v' &= \sum_{B' \in N_{B_{i,j,k}}} \sum_{C \in B'} e^{-\frac{v + C_{v} - (x,y,z) \cdot (C_{x},C_{y},C_{z})}{2}} \times C_{q} \\ q' &= q \end{split}$$

**Iteration space and dependences** For each particle in the each box the algorithm iterates over all particles in the neighbor boxes:

$$I = \{(b, p, n, c) : b \in B, p \in b, n \in N_b, c \in n\}$$

The only dependences present in this computation are the dependences introduced by the "reduce" operation (for summation).

#### Memory access mapping

$$M_r(b, p, n, c) = \{b, p, n, c\}$$
$$M_w(b = B_{i,j,k}, p = P_s, n, c) = \{p' = P'_s \in F_{i,j,k}\}$$

#### Low-level implementation details

The benchmark launches  $N_{bx} \times N_{by} \times N_{bz}$  work groups *i.e.* one group for each box [20, 19].

**Device code** See Algorithm 50.

Host code See Algorithm 51.

#### Host data structures

```
• struct Neighbor
    struct Neighbor
    {
        int x, y, z; //Coordinates.
        int number; //Global neighbor box id.
        long offset; //Box offset in Particles array.
    }
```

CARP-ARM-RP-001-v1.0



```
Algorithm 50 LavaMD device code abstraction.
Input: global struct Box dB[]
Input: global struct Particle dP[]
Input: global float dQ[]
Input: int S
Input: int Npb
                                                                                            \triangleright S = N_{bx} \times N_{by} \times N_{bz}
Output: global struct Particle dF[]
Local: local Particle dlA[\Lambda]
Local: local Particle dlB[A]
Local: local Float dlo[A]
  1: kernel LMDKERNEL (dB, dP, dQ, S, dF, Npb)
 2:
          if \gamma < S then
               \texttt{alpha} \leftarrow 0.5
 3:
 4:
               first_i \leftarrow dB[\gamma].offset
               for wtx in \lambda : (Npb - 1) : \Lambda do
 5:
                    dlA[wtx] \leftarrow dP[first_i + wtx]
 6:
 7:
               end for
               BARRIER(CLK_LOCAL_MEM_FENCE)
 8:
               for k in 0 : (dB[\gamma].nn) do
 9:
                   if k = 0 then
10:
                        \texttt{ptr} \gets \gamma
11:
12:
                   else
                        ptr \leftarrow dB[\gamma].nei[k-1].number
13:
                   end if
14:
                   \texttt{first}_i \gets \texttt{dB}[\texttt{ptr}].\texttt{offset}
15:
                   for wtx in \lambda : (Npb -1) : \Lambda do
16:
                        dlB[wtx] \leftarrow dP[first_1 + wtx]
17:
18:
                        dlQ[wtx] \leftarrow dQ[first_i + wtx]
                   end for
19:
20:
                   BARRIER(CLK_LOCAL_MEM_FENCE)
                   for wtx in \lambda : (Npb - 1) : \Lambda do
21:
                        for j in 0 : (Npb-1) do
22:
                             r2 \leftarrow dlA[wtx].v + dlB[wtx].v + dlA[wtx].(x,y,z) \cdot dlB[wtx].(x,y,z)
23:
                             u2 \leftarrow 2 \times alpha^2 \times r22
24:
                             v_{i,j} \leftarrow e^{-u^2}
25:
                             \texttt{fs} \gets 2 \times \texttt{v}_{\texttt{i},\texttt{j}}
26:
                             dx \leftarrow dlA[wtx].x - dlB[j].x
27:
                             \texttt{fx}_{\texttt{i},\texttt{i}} \gets \texttt{fs} \times \texttt{dx}
28:
                             dy \leftarrow dlA[wtx].y - dlB[j].y
29:
30:
                             fy_{i,j}fs \times dy
                             dz \leftarrow dlA[wtx].z - dlB[j].z
31:
                             fz_{i,i}fs \times dz
32:
```



Algorithm 50 LavaMD device code abstraction (continues).				
33:	$\texttt{dF}[\texttt{first}_{\texttt{i}} + \texttt{wtx}].\textit{v} \leftarrow \texttt{dlQ}[\texttt{j}] \times \texttt{v}_{\texttt{i},\texttt{j}}$			
34:	$dF[\texttt{first}_i + \texttt{wtx}].x \leftarrow dlQ[j]  imes \texttt{fx}_{i,j}$			
35:	$\mathtt{dF}[\mathtt{first}_{\mathtt{i}} + \mathtt{wtx}].y \leftarrow \mathtt{dlQ}[\mathtt{j}] \times \mathtt{fy}_{\mathtt{i},\mathtt{j}}$			
36:	$dF[first_i + wtx].z \leftarrow dlQ[j] \times fz_{i,j}$			
37:	end for			
38:	BARRIER(CLK_LOCAL_MEM_FENCE)			
39:	end for			
40:	BARRIER(CLK_LOCAL_MEM_FENCE)			
41:	end for			
42:	end if			
43:	end kernel			

```
• struct Box
            typedef struct Box
              int x, y, z; //Coordinates.
              int number; //Global box id.
              long offset; //Box offset.
              int nn; //Number of neighbors.
              struct Neighbor nei[26]; //Neighbors.
            };
   • struct Particle
            struct Particle
              int x;
              int y;
              int z;
              int v;
            }
   • struct Box hB[Nbx * Nby * Nbz]: B_{i,j,k} \to hB[i * Nby * Nbx + j * Nbx + k]
   • struct Particle hP[Nbx * Nby * Nbz * Npb]:
     B_{i,j,k}[s] \rightarrow hP[(i * Nby * Nbx + j * Nbx + k) * Npb + s]
   • struct Particle hF[Nbx * Nby * Nbz * Npb]:
     F_{i,j,k}[s] \rightarrow hP[(i * Nby * Nbx + j * Nbx + k) * Npb + s]
   • float hQ[Nbx * Nby * Nbz * Npb]:
     B_{i,j,k}[s].q \rightarrow hq[(i * Nby * Nbx + j * Nbx + k) * Npb + s]
Device data structures
```

```
• struct Neighbor
    struct Neighbor
    {
        int x, y, z; //Coordinates.
```



Algorithm 51 LavaMD host code abstraction

```
1: procedure ALLOCATEMEMORY(Nbx, Nby, Nbz, Npb) > Allocate host memory and device
     buffers.
         hB \leftarrow AllocateHostMemory(Nbx \times Nby \times Nbz)
 2:
 3:
         hP \leftarrow ALLOCATEHOSTMEMORY(Nbx \times Nby \times Nbz \times Npb)
         hF \leftarrow AllocateHostMemory(Nbx \times Nby \times Nbz \times Npb)
 4:
         hQ \leftarrow ALLOCATEHOSTMEMORY(Nbx \times Nby \times Nbz \times Npb)
 5:
         \texttt{dB} \leftarrow \texttt{AllocateBuffer}(\texttt{Nbx} \times \texttt{Nby} \times \texttt{Nbz}, \texttt{READ\_WRITE})
 6:
         dP \leftarrow ALLOCATEBUFFER(Nbx \times Nby \times Nbz \times Npb, READ_WRITE)
 7:
         dF \leftarrow AllocateBuffer(Nbx \times Nby \times Nbz \times Npb, READ_WRITE)
 8:
         dQ \leftarrow AllocateBuffer(Nbx \times Nby \times Nbz \times Npb, READ_WRITE)
 9:
10: end procedure
11: procedure INITIALIZEMEMORY(Nbx, Nby, Nbz, Npb)
                                                                              \triangleright Initialize host memory.
12:
         hP \leftarrow RANDOMELEMENTS(Nbx \times Nby \times Nbz \times Npb)
13:
         \texttt{hQ} \leftarrow \texttt{RANDOMELEMENTS}(\texttt{Nbx} \times \texttt{Nby} \times \texttt{Nbz} \times \texttt{Npb})
         hF \leftarrow ZEROELEMENTS(Nbx \times Nby \times Nbz \times Npb)
14:
         hB \leftarrow INITELEMENTS(Nbx \times Nby \times Nbz)
15:
16: end procedure
Input: int Nbx
Input: int Nby
                                                                                            ⊳ Cube size.
Input: int Nbz
Input: int Npb
                                                                   \triangleright Number of particles in each box.
17: ALLOCATEMEMORY(Nbx, Nby, Nbz, Npb)
18: INITIALIZEMEMORY(Nbx, Nby, Nbz, Npb)
19: lmdProgram ← BUILDPROGRAM("lmdProgram.cl")
20: lmdKernel ← BUILDKERNEL(lmdProgram, "lmdKernel")
21: COPYTODEVICE(dB, hB, True, ∅)
22: COPYTODEVICE(dP, hP, True, \emptyset)
23: COPYTODEVICE(dF, hF, True, ∅)
24: COPYTODEVICE(dQ, hQ, True, \emptyset)
25: S \leftarrow Nbx \times Nby \times Nbz
26: Lsize \leftarrow 128
27: SETKERNELARGUMENTS(lmdKernel, dB, dP, dQ, S, dF, Npb)
28: ENQUEUEKERNEL(lmdKernel, (Lsize × S), (Lsize), ∅)
29: COPYTOHOST(dF,hF,True,∅)
```




```
int number; //Global neighbor box id.
               long offset; //Box offset in Particles array.
             }
   • struct Box
             typedef struct Box
               int x, y, z; //Coordinates.
               int number; //Global box id.
               long offset; //Box offset.
               int nn; //Number of neighbors.
               struct Neighbor nei[26]; //Neighbors.
             };
   • struct Particle
             struct Particle
               int x;
               int y;
               int z;
               int v;
             }
   • global struct Box dB[Nbx * Nby * Nbz]: B_{i,j,k} \rightarrow dB[i * Nby * Nbx + j * Nbx + k]
   • global struct Particle dP[Nbx * Nby * Nbz * Npb]:
      B_{i,j,k}[s] \rightarrow dP[(i * Nby * Nbx + j * Nbx + k) * Npb + s]
   • global struct Particle dF[Nbx * Nby * Nbz * Npb]:
      F_{i,j,k}[s] \rightarrow dP[(i * Nby * Nbx + j * Nbx + k) * Npb + s]
   • global float dQ[Nbx * Nby * Nbz * Npb]:
      B_{i,j,k}[s].q 
ightarrow 	ext{dQ[(i * Nby * Nbx + j * Nbx + k) * Npb + s]}
   • local float dlA[Npb]: B_{i,j,k}[s] \rightarrow dlQ[s]
   • local float dlB[Npb]: B_{i,i,k}[s] \rightarrow dlQ[s]
   • local float dlQ[Npb]: B_{i,j,k}[s].q \rightarrow dlQ[s]
Input datasets The benchmark uses a randomly generated 10 \times 10 \times 10 cube of boxes. Each
```

box contains 100 particles with randomly generated attributes.

## Partitioned iteration space

$$S = N_{bx} \times N_{by} \times N_{bz}$$
  

$$i = \gamma/(Nbx \times Nby)$$
  

$$j = (\gamma/Nbx) \mod Nby$$
  

$$k = \gamma \mod Nbx$$





$$P = \{(\mathbf{v}, \lambda, \mathbf{k}, \mathtt{wtx}, \mathtt{j}) : 0 \leq \gamma < \Gamma = S, 0 \leq \lambda < \Lambda = 128, 0 \leq k < |N_{B_{i,j,k}}|, \mathtt{wtx} \text{ in } \lambda : (N_{pb} - 1) : \Lambda, 0 \leq j < N_{pb}\}$$

$$f(B_{i,j,k}, B_{i,j,k}[s], (B_{i',j',k'} = N_{B_{i,j,k}}[l]), B_{i',j',k'}[s']) \rightarrow (i \times N_{bx} \times N_{by} + j \times N_{bx} + k, s \mod (\Lambda), l, (s/\Lambda) \times \Lambda, s')$$

$$f^{-1}(\mathbf{v}, \mathbf{\lambda}, \mathbf{k}, \mathtt{wtx}, \mathtt{j}) \to (B_{i,j,k}, B_{i,j,k}[\mathtt{wtx} + \mathbf{\lambda}], N_{B_{i,j,k}}[\mathtt{k}], N_{B_{i,j,k}}[\mathtt{k}][\mathtt{j}])$$

**Device memory access mapping** The function, to obtain the neighbor index in the global array based on the its index in local array.

ptr(k, v) = dB[v].nei[k-1].number

Each work-item performs the following reads and writes:

**Performance metrics** The performance is evaluated using the following metric:

- Total OpenCL initialization time:
  - OpenCL context creation.
  - OpenCL command queue creation.
  - OpenCL program compilation.
  - OpenCL lmdKernel compilation.
- Device memory allocation time.
- Host to device memory transfer time.
- Kernel (ImdKernel) execution time.
- Device to host memory transfer time.
- Device memory freeing time.

**Validation mechanism** The benchmark has no validation mechanisms.

## **Target-specific optimisations**

- **Consecutive memory access.** Work items with consecutive global ids in dimension 0 access consecutive memory locations which improves the memory system performance (caching, coalescing).
- **Branching elimination.** In order to avoid heavily branching kernel code, the neighbors information is computed on the host side as a benchmark input (code, which computes this information uses a lot of conditional jumps, which degrade the performance on the device).





## Target-specific optimization opportunities

• Dynamic local work size. One iteration variable wtx and one loop can be eliminated from the kernel code if the local work size is equal to the number of particles per box. In current implementation  $\Lambda = 128$  and  $N_{pb} = 100$  which results in additional boundary checking (currently performed via while loop).





# 2.3.9 Leukocyte Tracking

## **High-level description**

The benchmark detects and tracks rolling white blood cells (leukocytes) in the input video stream. Cells are detected in the first video frame and tracked in following frames. See [4] for more details.

The description below is applied to the following parts of the algorithm:

- Cells detection:
  - Image dilation calculation.
  - Gradient Inverse Coefficient of Variation (GICOV) score computation.
- Cells tracking:
  - Motion Gradient Vector Flow (MGVF) computation see [17] for more details.

Only these parts of the algorithm are implemented as OpenCL kernels *i.e.* are the most interesting from GPGPU benchmark analysis point of view.

Abstract data structures The benchmark operates on the following data objects:

- Image dilation:
  - *I* is an input  $I_n \times I_m$  frame (image).
  - *M* is an input  $M_n \times M_m$  dilate mask.
  - *D* is an intermediate  $I_n \times I_m$  dilated image.
- GICOV score computation:
  - $N_c$  denotes the number of ellipses to try.
  - $N_p$  denotes the number of sample points.
  - X is an input  $I_n \times I_m$  matrix of image X-gradients.
  - *Y* is an input  $I_n \times I_m$  matrix of image Y-gradients.
  - *G* is an intermediate  $I_n \times I_m$  image GICOV score array.
  - T is an intermediate  $N_p$  angles array.
  - C is an intermediate  $N_p \cos array$ .
  - S is an intermediate  $N_p$  sin array.
  - Tx is an intermediate  $N_c \times N_p$  offset array.
  - Ty is an intermediate  $N_c \times N_p$  offset array.
- MGVF computation:
  - *I* is an input  $I_n \times I_m$  frame (image).
  - *W* is an intermediate  $I_n \times I_m$  MGVF array.





## Computation

• Image dilation calculation:

$$D_{i,j} = \max_{0 \le x < M_n, 0 \le y < M_m, M_{x,y} \ne 0} D_{i+x-M_n/2, j+y-M_m/2}, 0 \le i < I_n, 0 \le j < I_m$$

• GICOV score computation (with radius *R*):

$$T_i = 2 \times \pi \times i/N_p, 0 \le i < N_p$$
  

$$C_i = \cos(T_i), 0 \le i < N_p$$
  

$$S_i = \sin(T_i), 0 \le i < N_p$$

For every point (i, j):  $R + 2 \le i < I_n - R - 2, R + 2 \le j < I_m - R - 2$ :

$$\begin{aligned} \mathsf{Gt}_{k,n} &= X_{j+\mathsf{Ty}_{k,n},i+Tx_{k,n}} \times C_n + Y_{j+\mathsf{Ty}_{k,n},i+Tx_{k,n}} \times S_n, 0 \leq k < N_c, 0 \leq n < N_p \\ M_k &= \left(\sum_n \mathsf{Gt}_{k,n}\right) / N_p \\ V_k &= \left(\sum_n (\mathsf{Gt}_{k,n} - \mathsf{M}_k)^2\right) / (N_p - 1) \\ G_{j,i} &= \max_k M_k^2 / V_k \end{aligned}$$

- MGVF computation (for every cell c in  $0: (N_c 1)$ ):
  - Initialize MGVF array:

$$W^0 = I$$

- Iteratively compute the MGVF array (*W*):

$$\begin{split} W_{i,j}^{\tau+1} &= W_{i,j}^{\tau} + \frac{\mu}{\lambda} \sum_{l=-1}^{1} \sum_{m=-1}^{1} H_e \left( (lv^y + mv^x) \times (W_{i+l,j+m}^{\tau} - W_{i,j}^{\tau}) \right) \times (W_{i+l,j+m}^{\tau} - W_{i,j}^{\tau}) \\ &- \frac{1}{\lambda} I_{i,j} (W_{i,j}^{\tau} - I_{i,j}) \end{split}$$

– Compute new  $W^{\tau}$  till fixed point is reached:

$$W = W^{\tau}, \|W^{\tau} - W^{\tau-1}\| < \varepsilon$$

– The  $H_e()$  function:

$$H_e(x) \to (\arctan(x) \times \frac{1}{\pi}) + \frac{1}{2}$$

## **Iteration spaces**

• Image dilation calculation:

$$I_D = \{(i, j, x, y) : 0 \le i < I_n, 0 \le j < I_m, 0 \le x < M_n, 0 \le y < M_m\}$$

• GICOV score computation (with radius *R*):

$$I_G = \{(i, j, k, n) : R + 1 \le i < I_n - R - 2, R + 1 \le j < I_m - R - 2, 0 \le k < N_c, 0 \le n < N_p\}$$

• MGVF computation:

$$I_M = \{ (c, \tau, i, j) : 0 \le c < N_c, 0 \le \tau, 0 \le i < I_n, 0 \le j < I_m \}$$

CARP-ARM-RP-001-v1.0





**Dependences** The dependencies are introduced by the other parts of the algorithm (not included in this description).

#### Memory access mapping

• Image dilation calculation:

$$M_r(i, j, x, y) = \{D_{i+x-S_n/2, j+y-S_m/2}, M_{x,y}\}$$
  
 $M_w(i, j, x, y) = \{D_{i,j}\}$ 

• GICOV score computation (with radius *R*):

$$egin{array}{rll} M_r(i,j,k,n) &= \{ { t Tx}_{k,n}, { t Ty}_{k,n}, C_n, S_n, { t Gt}_{k,n} \} \ M_{mr}(i,j,k,n) &= \{ X,Y \} \ M_w(i,j,k,n) &= \{ { t Gt}_{k,n}, G_{j,i} \} \end{array}$$

• MGVF computation:

$$\begin{aligned} M_r(c,\tau,i,j) &= \{ I_{i,j}, W^{\tau}_{(i-1):(i+1),(j-1):(j+1)} \} \\ M_w(c,\tau,i,j) &= \{ W^{\tau}_{i,j} \} \end{aligned}$$

#### Low-level implementation details

**Device code** See Algorithm 52, Algorithm 53, Algorithm 54.

#### **Device data structures**

- int dI[In \* Im]:  $I_{i,j} \rightarrow d$ I[i \* Im + j]
- int dW[In \* Im]:  $W_{i,j} \rightarrow dW[i * Im + j]$
- int dD[In \* Im]:  $D_{i,j} \rightarrow$  dD[i \* Im + j]
- int dM[Mn \* Mm]:  $M_{i,j} \rightarrow$  dM[i \* Mm + j]
- int dX[In \* Im]:  $X_{i,j} \rightarrow dX[i * Im + j]$
- int dY[In \* Im]:  $Y_{i,j} \rightarrow dY[i * Im + j]$
- int dG[In \* Im]:  $G_{i,j} \rightarrow dG[i * Im + j]$
- int dT[Np]:  $T_i \rightarrow dT[i]$
- int dC[Np]:  $C_i \rightarrow dC[i]$
- int dS[Np]:  $S_i \rightarrow$  dS[i]



```
Algorithm 52 The Leukocyte device code abstraction (GICOV calculation).
Input: int Im
Input: global float dX[]
Input: global float dY[]
Input: constant float dS[]
Input: constant float dC[]
Input: constant int dTx[]
Input: constant int dTy[]
Output: global float dG[]
  1: kernel LKKERNEL1 (Im, dX, dX, dS, dC, dTx, dTy, dG)
  2:
            i \leftarrow v + R + 2
            j \leftarrow \lambda + R + 2
  3:
            \mathtt{G}_{max} \gets 0
  4:
            for k in 0 : (N_c - 1) do
  5:
                  \texttt{sum} \gets 0
  6:
  7:
                  \texttt{M2} \leftarrow 0
  8:
                  \texttt{mean} \gets 0
                  for n in 0: (N_p - 1) do
  9:
                       \mathbf{y} \leftarrow \mathbf{j} + \mathbf{dTy}[\mathbf{k} \times N_p + \mathbf{n}]
 10:
                       \mathbf{x} \leftarrow \mathbf{i} + \mathbf{dTx}[\mathbf{k} \times N_p + \mathbf{n}]
11:
                       \texttt{addr} \gets \texttt{x} \times \texttt{Im} + \texttt{y}
12:
13:
                       p \leftarrow dX[addr] \times dC[n] + dY[addr] \times dS[n]
14:
                       \texttt{sum} \gets \texttt{sum} + \texttt{p}
 15:
                       \texttt{delta} \gets \texttt{p}-\texttt{mean}
                       mean \leftarrow mean + delta/(n+1)
 16:
                       M2 \leftarrow M2 + \texttt{delta}/(n+1)
17:
                  end for
18:
                  \texttt{mean} \leftarrow \texttt{sum}/\texttt{N}_p
19:
20:
                  \texttt{var} \gets \texttt{M2}/(\texttt{N}_\texttt{p}-1)
                  \texttt{G}_{max} \gets max(\texttt{mean}^2/\texttt{var},\texttt{G}_{max})
21:
            end for
22:
            \mathtt{dG}[\mathtt{i} \times \mathtt{Im} + \mathtt{j}] \gets \mathtt{G}_{max}
23:
24: end kernel
```



Algorithm 53 The Leukocyte device code abstraction (Dilation calculation). Input: int Im Input: int In Input: int Mm Input: int Mn Input: global constant float dM[] Input: global float dI[] Output: global float dD[] 1: kernel LKKERNEL2 (Im, In, Mm, Mn, dM, dI, dD)  $\mathtt{C_i} \gets \mathtt{Mm}/2$ 2:  $C_j \leftarrow Mn/2$ 3:  $\mathtt{i} \gets \gamma \bmod \mathtt{Im}$ 4:  $j \leftarrow \gamma / \texttt{Im}$ 5: 6:  $\texttt{max} \gets 0$ 7: for  $el_i$  in 0: (Mm - 1) do  $\texttt{y} \gets \texttt{i} - \texttt{C}_\texttt{i} + \texttt{el}_\texttt{i}$ 8: if  $y \geq 0 \wedge y < \texttt{Im}$  then 9: for  $el_1$  in 0: (Mn-1) do 10:  $x \gets j - \texttt{C}_j + \texttt{el}_j$ 11: 12: if  $x \ge 0 \land x < \texttt{In} \land \texttt{dM}[\texttt{el}_i \times \texttt{Mn} + \texttt{el}_j] \ne 0$  then 13:  $\texttt{addr} \gets \texttt{x} \times \texttt{Im} + \texttt{y}$ 14:  $\texttt{tmp} \leftarrow \texttt{dI}[\texttt{addr}]$  $\texttt{max} \gets \max(\texttt{max},\texttt{tmp})$ 15: 16: end if end for 17: end if 18: 19: end for  $\texttt{dD}[\texttt{i} \times \texttt{In} + \texttt{j}] \gets \texttt{max}$ 20: 21: end kernel





Algorithm 54 The Leukocyte device code abstraction (MGVF calculation).

```
Input: global int dI[]
Input: constant int dIO[]
Input: constant int dMO[]
Input: constant int dNO[]
Input: float vx
Input: float vy
Input: float e
Input: float cutoff
Input: int Imax
Input/Output: global float dW[]
Local: float dlw[]
Local: float dlB
Local: int Conv
  1: kernel LKKERNEL3 (dI,dIO,dMO,dNO,vx,vy,e,cutoff,Imax,dW)
 2:
          I \leftarrow dIO[v]
 3:
          dpW \leftarrow \&(dIW[I])
 4:
          dpI \leftarrow \&(dI[I])
          \mathtt{m} \leftarrow \&(\mathtt{dMO}[v])
 5:
          n \leftarrow \&(dNO[v])
 6:
          \texttt{max} \gets (\texttt{m}+\texttt{n}+\Lambda-1)/\Lambda
 7:
          for tb in 0: (\max - 1) do
 8:
 9:
               \texttt{offset} \gets \texttt{tb} \times \Lambda
               \mathtt{i} \gets (\lambda + \texttt{offset}) / \mathtt{n}
10:
               j \leftarrow (\lambda + \texttt{offset}) \mod n
11:
               if i < m then
12:
                    dlW[i \times n + j] \leftarrow dpW[i \times n + j]
13:
               end if
14:
          end for
15:
16:
          BARRIER(CLK_LOCAL_MEM_FENCE)
          if \lambda = 0 then
17:
18:
               \texttt{flag} \leftarrow true
          end if
19:
          BARRIER(CLK_LOCAL_MEM_FENCE)
20:
          \texttt{tmod} \gets \lambda \text{ mod } \texttt{n}
21:
          \texttt{tbmod} \leftarrow \Lambda \mod \texttt{n}
22:
23:
          \texttt{iter} \gets 0
24:
          while \texttt{flag} \land \texttt{iter} < \texttt{Imax} \ \textbf{do}
               \texttt{diff} \gets 0
25:
26:
               j \leftarrow \text{tmod} - \text{tbmod} \triangleright Possible bug here: i is not initialized (value from the previous
     loop is used).
```



Algorithm	<b>n 54</b> The Leukocyte device code abstraction (IMGVF calculation continues).
27:	for tb in $0: (\max - 1)$ do
28:	$oldi \leftarrow i$
29:	$\texttt{oldj} \gets \texttt{j}$
30:	$\texttt{offset} \gets \texttt{tb} \times \Lambda$
31:	$\texttt{i} \leftarrow (\lambda + \texttt{offset}) / \texttt{n}$
32:	$\texttt{j} \leftarrow \texttt{j} + \texttt{tbmod}$
33:	if $j \ge n$ then
34:	$j \leftarrow j-n$
35:	end if
36:	$\mathtt{nval} \gets 0$
37:	$\texttt{oval} \leftarrow 0$
38:	if $i < m$ then
39:	$\texttt{oval} \leftarrow \texttt{FETCH2D}(\texttt{dlW},\texttt{i},\texttt{j})$
40:	$\texttt{iU} \gets \max(0, i-1)$
41:	$\texttt{iD} \gets \min(\texttt{m}-1,i+1)$
42:	$\texttt{jL} \gets \max(0, j-1)$
43:	$\mathtt{jR} \gets \min(\mathtt{n}-1,j+1)$
44:	$\texttt{U} \leftarrow \texttt{Fetch2d}(\texttt{dlW},\texttt{iU},\texttt{j})$
45:	$D \leftarrow Fetch2d(dlW, iD, j)$
46:	$L \leftarrow Fetch2d(dlW, i, jL)$
47:	$R \leftarrow Fetch2d(dlW, i, jR)$
48:	$\texttt{UR} \leftarrow \texttt{FETCH2D}(\texttt{dlW},\texttt{iU},\texttt{jR})$
49:	$DR \leftarrow Fetch2d(dlW, iD, jR)$
50:	$\texttt{UL} \leftarrow \texttt{FETCH2D}(\texttt{dlW},\texttt{iU},\texttt{jL})$
51:	$\texttt{DL} \leftarrow \texttt{FETCH2D}(\texttt{dlW},\texttt{iD},\texttt{jL})$
52:	$\texttt{nval} \leftarrow \texttt{COMPUTENEWVALUE}(\texttt{U},\texttt{D},\texttt{L},\texttt{R},\texttt{UR},\texttt{DR},\texttt{UL},\texttt{DL},\texttt{oval})$
53:	$\texttt{vI} \gets \texttt{dpI}[\texttt{i} \times \texttt{n} + \texttt{j}]$
54:	$\texttt{nval} \gets \texttt{nval} - 1/4 \times \texttt{vI} \times (\texttt{nval} - \texttt{vI})$
55:	end if
56:	if $tb > 0$ then
57:	$\texttt{offset} \gets (\texttt{tb}-1) \times \Lambda$
58:	if oldi < m then
59:	$\texttt{dlW}[\texttt{oldi} \times n + \texttt{oldj}] \gets \texttt{dlB}[\lambda]$
60:	end if
61:	end if
62:	if $tb < max - 1$ then
63:	$\texttt{dlB}[\boldsymbol{\lambda}] \gets \texttt{nval}$
64:	else
65:	if i < m then
66:	$dlW[i  imes n + j] \leftarrow nval$
67:	end if
68:	end if
69:	$\texttt{diff} \leftarrow \texttt{diff} +  \texttt{nval} - \texttt{oval} $
70:	BARRIER(CLK_LOCAL_MEM_FENCE)
71:	end for

CARP-ARM-RP-001-v1.0



Aigu	<b>Gruin 54</b> The Leukocyte device code abstraction (hvio vir calculation continues).
72:	$\texttt{dlB}[\lambda] \gets \texttt{diff}$
73:	if $\lambda \geq 2^8$ then
74:	$\texttt{dlB}[\lambda-2^8] \leftarrow \texttt{dlB}[\lambda-2^8] + \texttt{dlB}[\lambda]$
75:	end if
76:	$\texttt{th} \gets 2^7$
77:	while $ au > 0$ do
78:	if $\lambda <  ext{th}$ then
79:	$\mathtt{dlB}[\lambda] \gets \mathtt{dlB}[\lambda] + \mathtt{dlB}[\lambda + \mathtt{th}]$
80:	end if
81:	BARRIER(CLK_LOCAL_MEM_FENCE)
82:	$\texttt{th} \gets \texttt{th}/2$
83:	end while
84:	if $\lambda = 0$ then
85:	$\texttt{mean} \gets \texttt{dlB}[0] / (\texttt{m} \times \texttt{n})$
86:	if mean < cutoff then
87:	$\texttt{flag} \gets \textbf{false}$
88:	end if
89:	end if
90:	BARRIER(CLK_LOCAL_MEM_FENCE)
91:	$\texttt{iter} \gets \texttt{iter} + 1$
92:	end while
93:	for tb in $0:(\max - 1)$ do
94:	$\texttt{offset} \gets \texttt{tb} \times \Lambda$
95:	$\texttt{i} \leftarrow (\boldsymbol{\lambda} + \texttt{offset}) / \texttt{n}$
96:	$\mathtt{j} \leftarrow (\lambda + \mathtt{offset}) mod \mathtt{n}$
97:	if i < m then
98:	$dpW[i  imes n + j] \leftarrow dlW[i  imes n + j]$
99:	end if
100:	end for
101:	end kernel

155

# Algorithm 54 The Leukocyte device code abstraction (IMGVF calculation continues).





- int dTx[Nc \* Np]:  $Tx_{i,j} \rightarrow dTx[i * Np + j]$
- int dTy[Nc \* Np]:  $Ty_{i,j} \rightarrow dTy[i * Np + j]$
- int dIO[Nc]: Intermediate offsets buffer for image buffer dI
- int dNO[Nc]: Intermediate boundaries buffer for image buffer dI
- int dMO[Nc]: Intermediate boundaries buffer for image buffer dI
- int dlw[41 \* 81]: Local buffer for dw.
- int dlB[A]: Local buffer to compute reduction.

**Input datasets** The benchmark uses pre generated input video file with following parameters:

- Number of frames to process. Can be adjusted via command line. The default value is 10.
- Size of the video frame *i.e.* single image  $(I_n \times I_m) = (640 \times 480)$ .

## **Partitioned iteration space**

• Image dilation calculation (lkKernel2):

$$P_D = \{(\gamma, i, j) : 0 \le \gamma < \Gamma = I_m \times I_n, \}$$

$$f(i, j, x, y) \rightarrow (i \times I_m + j, x, y)$$

$$f^{-1}(\gamma, i, j) \to (\gamma/I_m, \gamma \mod I_m, i, j)$$

• GICOV score computation (lkKernel1):

$$P_G = \{(v, \lambda, k, n) : 0 \le v < N = (I_n - 2 \times R), 0 \le \lambda < \Lambda = (I_n - 2 \times R), 0 \le k < N_c, 0 \le n < N_p\}$$

$$f(i,j,k,n) \rightarrow (i-R-2,j-R-2,k,n)$$

$$f^{-1}(\mathbf{v}, \lambda, k, n) \rightarrow (\mathbf{v} + \mathbf{R} + 2, \lambda + \mathbf{R} + 2, k, n)$$

• MGVF computation (lkKernel3):

$$P_G = \{ (v, \lambda, \tau, t) : 0 \le v < N = Nc, 0 \le \lambda < \Lambda = 256, 0 \le \tau, 0 \le t < I_m \times I_n / \Lambda \}$$

$$f(c, \tau, i, j) \to (c, (i \times n + j) \mod \Lambda, \tau, (i \times n + j)/\Lambda)$$

$$f^{-1}(\mathbf{v}, \lambda, \tau, t) \to (\mathbf{v}, \tau, (\lambda + t \times \Lambda) / I_n, (\lambda + t \times \Lambda) \mod I_n)$$





## **Device memory access mapping**

• Image dilation calculation (lkKernel2):

$$G_r(\gamma, i, j) = \{ dM[i \times M_n + j], dI[(\gamma/I_m - M_n/ + j) \times I_m + \gamma \mod I_m - M_m/2 + i], \}$$
  
$$G_w(\gamma i, j) = \{ dI[(\gamma \mod I_m) \times I_n + (\gamma/I_m)] \}$$

• GICOV score computation (lkKernel1):

 $\begin{aligned} G_r(\mathbf{v}, \lambda, k, n) &= \{ \mathrm{dTy}[k \times N_p + n], \mathrm{dTx}[k \times N_p + n], \mathrm{dS}[n], \mathrm{dC}[n], \} \\ G_{mr}(\mathbf{v}, \lambda, k, n) &= \{ \mathrm{dX}[], \mathrm{dY}[] \} \\ G_w(\mathbf{v}, \lambda, k, n) &= \{ \mathrm{dG}[(\mathbf{v} + R + 2) \times I_m + (\lambda + R + 2)] \} \end{aligned}$ 

• MGVF computation (lkKernel3):

**Performance metrics** The performance of the algorithm implementation is evaluated using the following metrics:

- GICOV score computation time.
- Dilation computation time.
- MGVF computation time.
- Total tracking time.
- Total application execution time.

**Validation mechanism** The benchmark has no built-in validation mechanisms. The benchmark can print number of detected cells, which can be validated by external tools (not supplied).

## **Target-specific optimizations**

- **Consecutive memory access.** Work items with consecutive global ids access consecutive memory locations which improves the memory system performance (caching, coalescing).
- Local memory usage. In order to decrease the number of global memory accesses each work group allocates local buffers, which are used to perform all computations. This technique also increases data locality since the data layout can be changed while copying from global to local memory and vice versa.





# 2.3.10 LU Decomposition

#### **High-level description**

The benchmark computes the *LU* Decomposition of a square matrix *M*:

 $M = L \times U$ 

where L is a lower triangular matrix and U is an upper triangular matrix.

**Abstract data structures** The benchmark operates on the following floating point data objects:

- *M* is an input  $n \times n$  square matrix.
- *L* is an output  $n \times n$  square lower triangular matrix.
- U is an output  $n \times n$  square upper triangular matrix.

**Computation** The matrices are computed as follows:

$$U_{i,j} = M_{i,j} - \sum_{k=0}^{i-1} L_{i,k} \times U_{k,j} \quad 0 \le i < n, i \le j < n$$
$$L_{j,i} = \frac{1}{U_{i,i}} \times (M_{j,i} - \sum_{k=0}^{i-1} L_{j,k} \times U_{k,i}) \quad 0 \le i < n, i < j < n$$

**Iteration spaces** 

• Compute *U*:

$$I_U = \{(i, j, k) : 0 \le i < n, i \le j < n, 0 \le k < i\}$$

$$I_U: \begin{pmatrix} 1 & 0 & 0\\ 0 & 1 & 0\\ -1 & 0 & 1\\ -1 & 0 & 0\\ 1 & -1 & 0\\ 0 & 0 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j\\ k \end{pmatrix} \leq \begin{pmatrix} n-1\\ n-1\\ -1\\ 0\\ 0\\ 0 \end{pmatrix}$$

• Compute *L*:

$$I_L = \{(i, j, k) : 0 \le i < n, i < j < n, 0 \le k < i\}$$

$$I_{L}: \begin{pmatrix} 1 & 0 & 0\\ 0 & 1 & 0\\ -1 & 0 & 1\\ -1 & 0 & 0\\ 1 & -1 & 0\\ 0 & 0 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j\\ k \end{pmatrix} \leq \begin{pmatrix} n-1\\ n-1\\ -1\\ 0\\ -1\\ 0 \end{pmatrix}$$

CARP-ARM-RP-001-v1.0





## Dependences

$$\begin{split} \delta_{I_U} &= I_U : (k, j, 0 : (k-1)) \xrightarrow{t} I_U : (i, j, k) \land \\ &I_L : (k, i, 0 : (k-1)) \xrightarrow{t} I_U : (i, j, k) \\ \delta_{I_L} &= I_U : (i, i, 0 : (i-1)) \xrightarrow{t} I_L : (i, j, k) \land \\ &I_U : (k, i, 0 : (k-1)) \xrightarrow{t} I_L : (i, j, k) \land \\ &I_L : (k, j, 0 : (k-1)) \xrightarrow{t} I_L : (i, j, k) \end{split}$$

## Memory access mapping

• Compute *U*:

$$M_r(i, j, k) = \{M_{i,j}, L_{i,k}, U_{k,j}\}$$
  
 $M_w(i, j, k) = \{U_{i,j}\}$ 

• Compute *L*:

$$M_r(i, j, k) = \{U_{i,i}, U_{k,i}, M_{j,i}, L_{j,k}\}$$
  
 $M_w(i, j, k) = \{L_{j,i}\}$ 

## Low-level implementation details

The computation is performed by three OpenCL kernels, which update the input matrix in-place. Figure 2.1 shows the elements updated by each kernel (offset is incremented by block size B = 16 on each iteration).

**Device code** See Algorithm 55, Algorithm 56 and Algorithm 57.

Host code See Algorithm 58.

## Host data structures

• global float hM[n \* n]:  $M_{i,j}, L_{i,j}, U_{i,j} \rightarrow$  hM[i \* n + j]

**Device data structures** See Figure 2.2.

- global float dM[n \* n]:  $M_{i,j}, L_{i,j}, U_{i,j} \rightarrow dM[i * n + j]$
- local float dlD[B \* B]:  $M_{i,j}, L_{i,j}, U_{i,j} \rightarrow dlD[(i \ B) \ * B + j \ B], 0 \le i < n, ((i/B) \times B) \le j < ((i/B) + 1) \times B)$  (only blocks on the main diagonal).
- local float dlr[B \* B]:  $M_{i,j}, L_{i,j}, U_{i,j} \rightarrow dlr[(i \ \& B) \ * B + j \ \& B]$
- local float dlC[B \* B]:  $M_{i,j}, L_{i,j}, U_{i,j} \rightarrow dlC[(i \& B) * B + j \& B]$



```
Algorithm 55 The LU Decomposition (Diagonal elements) device code abstraction.
Input/Output: global float dM[]
Local: local float dlD[]
Input: int n
Input: int offset
 1: kernel LUDKERNEL1 (dM, n, offset)
          \texttt{dMoffset} \gets \texttt{offset} \times \texttt{n} + \texttt{offset}
 2:
          for i in 0 : (B - 1) do
 3:
                                                                ▷ Prefetch global data to local buffers (dlD).
               dlD[i \times B + \lambda] \leftarrow dM[dMoffset + i \times n + \lambda]
 4:
          end for
 5:
          BARRIER(CLK_LOCAL_MEM_FENCE)
 6:
          for i in 0 : (B - 1) do
                                                                              ▷ Compute diagonal block (dlD).
 7:
 8:
               if \lambda > i then
                    for j in 0 : (i - 1) do
 9:
                        dlD[\lambda \times B + i] \leftarrow dlD[\lambda \times B + i] - dlD[\lambda \times B + j] \times dlD[j \times B + i]
10:
                        dlD[\lambda \times B + i] \leftarrow dlD[\lambda \times B + i]/dlD[i \times B + i]
11:
                   end for
12:
               end if
13:
               BARRIER(CLK_LOCAL_MEM_FENCE)
14:
               if \lambda > i then
15:
16:
                   for j in 0 : i do
                        \mathtt{dlD}[(\mathtt{i}+1)\times\mathtt{B}+\lambda]\leftarrow\mathtt{dlD}[(\mathtt{i}+1)\times\mathtt{B}+\lambda]-\mathtt{dlD}[(\mathtt{i}+1)\times\mathtt{B}+\mathtt{j}]\times
17:
     dlD[j \times B + \lambda]
                   end for
18:
               end if
19:
               BARRIER(CLK_LOCAL_MEM_FENCE)
20:
21:
          end for
          \texttt{dMoffset} \gets (\texttt{offset} + 1) \times \texttt{n} + \texttt{offset}
22:
          for i in 0 : (B - 1) do
                                                              ▷ Write local buffers to global memory (dlD).
23:
               dM[dMoffset + \lambda] \leftarrow dlD[i \times B + i \times n + \lambda]
24:
          end for
25:
26: end kernel
```



```
Algorithm 56 The LU Decomposition (Perimeter Elements) device code abstraction.
Input/Output: global float dM[]
Local: local float dlR[]
Local: local float dlC[]
Local: local float dlD[]
Input: int n
Input: int offset
  1: kernel LUDKERNEL2 (dM, n, offset)
          if \lambda < B then
 2:
                                                                          ▷ Prefetch global data to local buffers.
               idx \leftarrow \lambda
 3:
 4:
               \texttt{dMoffset} \gets \texttt{offset} \times \texttt{n} + \texttt{offset}
               for i in 0 : (B/2 - 1) do
                                                                                  \triangleright Prefetch dlD (rows 0:(B/2 - 1)).
 5:
                    dlD[i \times B + idx] \leftarrow dM[dMoffset + i \times n + idx]
 6:
               end for
 7:
               \texttt{dMoffset} \gets \texttt{offset} \times \texttt{n} + \texttt{offset}
 8:
               for i in 0 : (B - 1) do
                                                                                                         ▷ Prefetch dlR.
 9:
                    dlR[i \times B + idx] \leftarrow dM[dMoffset + (v+1) \times B + i \times n + idx]
10:
               end for
11:
          else
12:
               \mathtt{idx} \gets \lambda - \mathtt{B}
13:
               \texttt{dMoffset} \gets (\texttt{offset}) \times \texttt{n} + \texttt{offset}
14:
15:
               for i in (B/2) : (B-1) do
                                                                                \triangleright Prefetch dlD (rows (B/2):(B-1).
                    \texttt{dlD}[\texttt{i} \times \texttt{B} + \texttt{idx}] \leftarrow \texttt{dM}[\texttt{dMoffset} + \texttt{i} \times \texttt{n} + \texttt{idx}]
16:
               end for
17:
               \texttt{dMoffset} \gets (\texttt{offset} + (\nu + 1) \times \texttt{B}) \times \texttt{n} + \texttt{offset}
18:
               for i in 0 : (B - 1) do
                                                                                                         ▷ Prefetch dlC.
19:
20:
                    dlC[i \times B + idx] \leftarrow dM[dMoffset + i \times n + idx]
               end for
21:
          end if
22:
23:
          BARRIER(CLK_LOCAL_MEM_FENCE)
```



Algorithm 56 The LU Decomposition (Perimeter Elements) device code abstraction (continues). if  $\lambda < B$  then ▷ Compute data in local memory. 24:  $\mathtt{idx} \gets \lambda$ 25: for i in 1 : (B - 1) do 26: for j in 0 : (i - 1) do $\triangleright$  Compute dlR. 27:  $dlR[i \times B + idx] = dlR[i \times B + idx] - dlD[i \times B + j] \times dlR[j \times B + idx]$ 28: end for 29: 30: end for else 31:  $\mathtt{idx} \gets \lambda - \mathtt{B}$ 32: for i in 0 : (B - 1) do 33: for j in 0 : (i - 1) do⊳ Compute dlC. 34:  $dlC[idx \times B + i] = dlC[idx \times B + i] - dlC[idx \times B + j] \times dlD[j \times B + i]$ 35: 36: end for  $dlC[idx \times B + i] = dlC[idx \times B + i]/dlD[i \times B + i]$ 37: 38: end for end if 39: BARRIER(CLK\_LOCAL\_MEM\_FENCE) 40: 41: if  $\lambda < B$  then ▷ Write local buffers to global memory.  $\mathtt{idx} \leftarrow \lambda$ 42: 43:  $dMoffset \leftarrow (offset) \times n + offset$ for i in 1 : (B - 1) do ⊳ Write dlR. 44:  $dM[dMoffset + (v+1) \times B + i \times n + idx] \leftarrow dlR[i \times B + idx]$ 45: end for 46: else 47:  $\mathtt{idx} \leftarrow \lambda - \mathtt{B}$ 48:  $\texttt{dMoffset} \leftarrow (\texttt{offset} + (\nu + 1) \times \texttt{B}) \times \texttt{n} + \texttt{offset}$ 49: for i in 0 : (B - 1) do ⊳ Write dlC. 50:  $dM[dMoffset + i \times n + idx] \leftarrow dlC[i \times B + idx]$ 51: end for 52: end if 53: 54: end kernel





Figure 2.1: The LU Decomposition partitioned iteration space.

**Input datasets** The benchmark uses pre-generated matrices of single-precision floating point values  $\in (0, 1)$  of the following sizes:  $64 \times 64$ ,  $256 \times 256$ ,  $512 \times 512$ ,  $2048 \times 2048$ .

# Partitioned iteration space

• Compute diagonal elements (ludKernel1):

$$egin{aligned} P_1 &= \{(\texttt{offset}, \lambda, \texttt{i}, \texttt{j}): & \texttt{offset} \ \mathbf{in} \ 0: (\texttt{n}-\texttt{B}-1): \texttt{B}, \ & 0 \leq \texttt{i} < \texttt{B}, \ & \texttt{i} < \lambda < \Lambda = \texttt{B}, \ & 0 \leq \texttt{j} \leq \texttt{i} \} \end{aligned}$$



```
Algorithm 57 The LU Decomposition (Internal Elements) device code abstraction.
```

Input/Output: global float dM[] Local: local float dlR[] Local: local float dlC[] Input: int n Input: int offset 1: **kernel** LUDKERNEL3 (dM, n, offset)  $\texttt{globalRow} \leftarrow \texttt{offset} + (v_1 + 1) \times \texttt{B}$ 2:  $globalCol \leftarrow offset + (v_0 + 1) \times B$ 3:  $\mathtt{dlR}[\lambda_1 \times \mathtt{B} + \lambda_0] \gets \mathtt{dM}[(\mathtt{offset} + \lambda_1) \times \mathtt{n} + \mathtt{globalCol} + \lambda_0]$ 4:  $\texttt{dlC}[\lambda_1 \times \texttt{B} + \lambda_0] \leftarrow \texttt{dM}[(\texttt{globalRow} + \lambda_1) \times \texttt{n} + \texttt{offset} + \lambda_0]$ 5: BARRIER(CLK\_LOCAL\_MEM\_FENCE) 6:  $\texttt{sum} \gets 0$ 7: for i in 0 : (B - 1) do 8: 9:  $\texttt{sum} \leftarrow \texttt{sum} + \texttt{dlC}[\lambda_1 \times \texttt{B} + \texttt{i}] \times \texttt{dlR}[\texttt{i} \times \texttt{B} + \lambda_0]$ end for 10:  $\texttt{dM}[(\texttt{globalRow} + \lambda_1) \times \texttt{n} + \texttt{globalCol} + \lambda_0] \leftarrow \texttt{dM}[(\texttt{globalRow} + \lambda_1) \times \texttt{n} +$ 11:  $globalCol + \lambda_0] - sum$ 12: end kernel

$$\begin{split} I_U: f(i, j, k) &\to ((\texttt{offset} = ((i \text{ mod } B) \times B)), i - \texttt{offset}, j - \texttt{offset}, k - \texttt{offset}) \\ I_L: f(i, j, k) &\to ((\texttt{offset} = ((j \text{ mod } B) \times B)), j - \texttt{offset}, i - \texttt{offset}, k - \texttt{offset}) \end{split}$$

$$\begin{array}{rcl} f^{-1}(\texttt{offset},\lambda,\texttt{i},\texttt{j}) & \to & P_U:(\texttt{offset}+\lambda,\texttt{offset}+\texttt{i},\texttt{offset}+\texttt{j}) \land \\ & & P_L:(\texttt{offset}+\texttt{i},\texttt{offset}+\lambda,\texttt{offset}+\texttt{j}) \end{array}$$

• Compute perimeter elements (ludKernel2):

$$\begin{array}{ll} P_2 &= \{(\texttt{offset}, v, \lambda, \texttt{i}, \texttt{j}): & \texttt{offset} \ \textbf{in} \ \texttt{0} : (\texttt{n}-\texttt{B}-\texttt{1}): \texttt{B}, \\ & 0 \leq v < N = (\texttt{n}-\texttt{offset})/\texttt{B}-\texttt{1}, \\ & 0 \leq \lambda < \Lambda = 2 \times \texttt{B}, \\ & 0 \leq \texttt{i} < \texttt{B}, \\ & 0 \leq \texttt{j} < \texttt{i} \} \end{array}$$

$$\begin{split} I_U: f(i, j, k) &\to ((\texttt{offset} = ((k \text{ mod } B) \times B)), (j - \texttt{offset})/B - 1, i - \texttt{offset} - 1, k - \texttt{offset}) \\ I_L: f(i, j, k) &\to ((\texttt{offset} = ((k \text{ mod } B) \times B)), (i - \texttt{offset})/B - 1, j - \texttt{offset} - 1, k - \texttt{offset}) \end{split}$$



Alg	orithm 58 The LU Decomposition host code abstraction.
1:	<b>procedure</b> ALLOCATEMEMORY(n) > Allocate host memory and device buffers.
2:	$\texttt{hM} \leftarrow \texttt{AllocateHostMemory}(\texttt{n} \times \texttt{n})$
3:	$\texttt{dM} \leftarrow \texttt{ALLOCATEBUFFER}(\texttt{n} \times \texttt{n}, \texttt{READ\_WRITE})$
4:	end procedure
5:	<b>procedure</b> INITIALIZEMEMORY(n) > Initialise host memory.
6:	$\texttt{hM} \gets \texttt{READFROMFILE}(\texttt{n} \times \texttt{n})$
7:	end procedure
Inp	<b>ut:</b> int n ▷ Input matrix size.
8:	ALLOCATEMEMORY(n)
9:	INITIALIZEMEMORY(n)
10:	$\texttt{ludProgram} \gets \texttt{BUILDPROGRAM}(\texttt{`'ludProgram.cl''})$
11:	$\texttt{ludKernel1} \leftarrow \texttt{BUILDKERNEL}(\texttt{ludProgram}, \texttt{``ludKernel1''})$
12:	$\texttt{ludKernel2} \gets \texttt{BUILDKERNEL}(\texttt{ludProgram},\texttt{`'ludKernel2''})$
13:	$\texttt{ludKernel3} \leftarrow \texttt{BUILDKERNEL}(\texttt{ludProgram},\texttt{`'ludKernel3''})$
14:	$B \leftarrow 16$ $\triangleright$ Block size.
15:	COPYTODEVICE(dM, hM, <b>True</b> , ∅)
16:	for offset in $0:(n-1-B):B$ do
17:	SETKERNELARGUMENTS(ludKernel1,dM,n,offset)
18:	$ENQUEUEKERNEL(ludKernel1, (B), (B), \emptyset)$
19:	SETKERNELARGUMENTS(ludKernel2,dM,n,offset)
20:	$ENQUEUEKERNEL(\mathtt{ludKernel2}, (\mathtt{B}  imes \mathtt{2}  imes ((\mathtt{n-offset})/\mathtt{B} - \mathtt{1})), (\mathtt{B}  imes \mathtt{2}), \emptyset)$
21:	SETKERNELARGUMENTS(ludKernel3,dM,n,offset)
22:	$\texttt{ENQUEUEKERNEL(ludKernel3, (B \times ((n - \texttt{offset})/B - 1), B \times ((n - \texttt{offset})/B - 1)), (B, B), \emptyset) = \texttt{ENQUEUEKERNEL(ludKernel3, (B \times ((n - \texttt{offset})/B - 1), B \times ((n - \texttt{offset})/B - 1))))}$
23:	end for
24:	COPYTOHOST(dM, hM, <b>True</b> , ∅)





Figure 2.2: The LU Decomposition partitioned iteration space.

 $\begin{array}{ll} f^{-1}(\texttt{offset}, \mathbf{v}, \lambda, \texttt{i}, \texttt{j}) & \to & P_U : (\texttt{offset} + 1 + i, \texttt{offset} + (\mathbf{v} + 1) \times B + \lambda, \texttt{offset} + j) \land \\ & P_L : (\texttt{offset} + (\mathbf{v} + 1) \times B + \lambda, \texttt{offset} + 1 + i, \texttt{offset} + j) \end{array}$ 





• Compute internal elements (ludKernel3):

$$\begin{array}{ll} P_3 &= \{(\texttt{offset}, v_0, v_1, \lambda_0, \lambda_1, \texttt{i}): & \texttt{offset in } 0: (\texttt{n}-\texttt{B}-1):\texttt{B}, \\ & 0 \leq v_0 < N_0 = (\texttt{n}-\texttt{offset})/\texttt{B}-1, \\ & 0 \leq v_1 < N_1 = (\texttt{n}-\texttt{offset})/\texttt{B}-1, \\ & 0 \leq \lambda_0 < \Lambda_0 = \texttt{B}, \\ & 0 \leq \lambda_1 < \Lambda_1 = \texttt{B}, \\ & 0 \leq \texttt{i} < \texttt{B} \} \end{array}$$

$$\begin{split} I_U: f(i,j,k) &\to \quad ((\texttt{offset} = ((k \bmod B) \times B)), (j - \texttt{offset})/B - 1, (i - \texttt{offset})/B - 1, \\ & (j - \texttt{offset}) \bmod B, (i - \texttt{offset}) \bmod B, k - \texttt{offset}) \\ I_L: f(i_L, j_L, k_L) &\to \quad ((\texttt{offset} = ((k \bmod B) \times B)), (i - \texttt{offset})/B - 1, (j - \texttt{offset})/B - 1, \\ & (i - \texttt{offset}) \bmod B, (j - \texttt{offset}) \bmod B, k - \texttt{offset}) \end{split}$$

$$\begin{array}{ll} f^{-1}(\texttt{offset}, \textit{v}_0, \textit{v}_1, \lambda_0, \lambda_1, \texttt{i}) & \rightarrow & P_U: (\texttt{offset} + (\textit{v}_1 + 1) \times \textit{B} + \lambda_1, \\ & \texttt{offset} + (\textit{v}_0 + 1) \times \textit{B} + \lambda_0, \texttt{offset} + \texttt{i}) \land \\ & P_L: (\texttt{offset} + (\textit{\gamma}_0 + 1) \times \textit{B} + \lambda_0, \\ & \texttt{offset} + (\textit{v}_1 + 1) \times \textit{B} + \lambda_1, \texttt{offset} + \texttt{i}) \end{array}$$

## **Device memory access mapping**

• Compute diagonal elements (ludKernel1):

CARP-ARM-RP-001-v1.0





• Compute perimeter elements (ludKernel2):

$$\begin{array}{lll} G_r(\texttt{offset}, v, \lambda, \texttt{i}, \texttt{j}) &= & \{\texttt{dM}[\texttt{offset} \times (\texttt{n}+1) + \texttt{n} \times \texttt{i} + \lambda]_{\lambda < \texttt{B}}, \\ & \texttt{dM}[\texttt{offset} \times (\texttt{n}+1) + \texttt{n} \times \texttt{i} + \lambda + (v+1) \times \texttt{B}]_{\lambda < \texttt{B}}, \\ & \texttt{dM}[(\texttt{offset} + \texttt{B}/2) \times \texttt{n} + \texttt{offset} + \texttt{n} \times \texttt{i} + \lambda - \texttt{B}]_{\lambda \geq \texttt{B}}, \\ & \texttt{dM}[(\texttt{offset} + (v+1) \times \texttt{B}) \times \texttt{n} + \texttt{offset} + \texttt{n} \times \texttt{i} + \lambda - \texttt{B}]_{\lambda \geq \texttt{B}}, \\ & \texttt{dM}[(\texttt{offset} \times (\texttt{n}+1) + \texttt{n} \times \texttt{i} + \lambda + (v+1) \times \texttt{B}]_{\lambda < \texttt{B}}, \\ & \texttt{dM}[(\texttt{offset} \times (\texttt{n}+1) + \texttt{n} \times \texttt{i} + \lambda + (v+1) \times \texttt{B}]_{\lambda < \texttt{B}}, \\ & \texttt{dM}[(\texttt{offset} + (v+1) \times \texttt{B}) \times \texttt{n} + \texttt{offset} + \texttt{n} \times \texttt{i} + \lambda - \texttt{B}]_{\lambda \geq \texttt{B}}, \\ & \texttt{dM}[(\texttt{offset} + (v+1) \times \texttt{B}) \times \texttt{n} + \texttt{offset} + \texttt{n} \times \texttt{i} + \lambda - \texttt{B}]_{\lambda \geq \texttt{B}}, \\ & \texttt{dM}[(\texttt{offset} + (v+1) \times \texttt{B}) \times \texttt{n} + \texttt{offset} + \texttt{n} \times \texttt{i} + \lambda - \texttt{B}]_{\lambda \geq \texttt{B}}, \\ & \texttt{dIR}[\texttt{i} \times \texttt{B} + \lambda]_{\lambda < \texttt{B}}, \texttt{dIR}[\texttt{i} \times \texttt{B} + \texttt{j}]_{\lambda < \texttt{B}}, \\ & \texttt{dIR}[\texttt{j} \times \texttt{B} + \lambda]_{\lambda < \texttt{B}}, \texttt{dIC}[(\lambda - \texttt{B}) \times \texttt{B} + \texttt{j}]_{\lambda \geq \texttt{B}}, \\ & \texttt{dIC}[(\lambda - \texttt{B}) \times \texttt{B} + \texttt{i}]_{\lambda \geq \texttt{B}}, \texttt{dIC}[(\lambda - \texttt{B}) \times \texttt{B} + \texttt{j}]_{\lambda < \texttt{B}}, \\ & \texttt{dIC}[\texttt{i} \times \texttt{B} + \lambda]_{\lambda < \texttt{B}}, \texttt{dID}[\texttt{i} \times \texttt{B} + \lambda]_{\lambda < \texttt{B}}, \\ & \texttt{dID}[\texttt{i} \times \texttt{B} + \lambda]_{\lambda < \texttt{B}}, \texttt{dIC}[\texttt{i} \times \texttt{B} + \lambda]_{\lambda < \texttt{B}}, \\ & \texttt{dID}[\texttt{i} \times \texttt{B} + \lambda]_{\lambda < \texttt{B}}, \texttt{dIC}[\texttt{i} \times \texttt{B} + \lambda]_{\lambda < \texttt{B}}, \\ & \texttt{dID}[\texttt{i} \times \texttt{B} + \lambda - \texttt{B}]_{\lambda \geq \texttt{B}}, \\ & \texttt{dIC}[(\lambda - \texttt{B}) \times \texttt{B} + \texttt{i}]_{\lambda \geq \texttt{B}} \} \end{aligned}$$

• Compute internal elements (ludKernel3):

**Performance metrics** The performance of the algorithm implementation is evaluated as the total time required by the OpenCL implementation to perform:

- 1. Write data to device.
- 2. Perform computations on device.
- 3. Read data from device.

**Validation mechanism** The benchmark multiplies matrix *L* by matrix *U* and compares the result to the original matrix with error tolerance  $10^{-4}$ .

**Target-specific optimizations** The following target-specific optimization are used in the algorithm implementation:

• **Consecutive memory access.** Work items with consecutive global ids access consecutive memory locations which improves the memory system performance (caching, coalescing).





• Local memory usage. In order to decrease the number of global memory accesses each work group allocates local buffers, which are used to perform all computations. This technique also increases data locality since the data layout can be changed while copying from global to local memory and vice versa.





# 2.3.11 k-Nearest Neighbors

Give a point *B* and a set of points  $\{P_i\}$  in a two-dimensional Euclidean space, find *k* nearest points from  $\{P_i\}$  to *B*.

## **High-level description**

Abstract data structures The benchmark operates on the following objects:

- Each point is a pair (x, y), where x is the latitude and y is the longitude of the point.
- *P* is an input array of *n* points.
- *D* is an intermediate array of distances.
- *B* is an input point.
- *R* is an output array of *k* points nearest to *B*.

**Computation** The algorithm consists of two steps:

1. Distance Calculation: For each point  $P_i$ ,  $0 \le i < n$ , the distance from  $P_i$  to B is calculated:

$$D_i = \sqrt{(P_i^x - B^x)^2 + (P_i^y - B^y)^2}$$

2. Neighbors Selection: k points with the minimal distance from B are selected:

$$R = \{p : |R| = k \land \forall p \in R, p' \notin R \to D_p \le D_{p'}\}$$

**Iteration domain and dependences** This algorithm requires several iteration spaces:

• Distance Calculation:

$$I_D = \{(i) : 0 \le i < n\}$$
$$I_D : \begin{pmatrix} 1 \\ -1 \end{pmatrix} \times (i) \le \begin{pmatrix} n-1 \\ 0 \end{pmatrix}$$

• Neighbors Selection:

$$I_R = \{(i, j) : 0 \le i < k, i \le j < n\}$$

$$I_R: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 1 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j \end{pmatrix} \le \begin{pmatrix} k-1\\ n-1\\ 0\\ -1 \end{pmatrix}$$

$$I = I_D \cup I_R$$

The computation generates the following dependences:



$$\begin{split} \delta_{I_R} &= I_R : (i,i:(j-1)) \xrightarrow{t} I_R : (i,j) \land \\ &I_R : (0:(i-1),n-1) \xrightarrow{t} I_R : (i,j) \land \\ &I_R : (0:(i-1),0:(n-1)) \xrightarrow{t} I_R : (i,j) \land \\ &I_D : (i:j) \xrightarrow{t} I_R : (i,j) \land \\ &I_D : (0:(n-1)) \xrightarrow{t} I_R : (i,n-1) \land \\ &I_D : (0:(n-1)) \xrightarrow{a} I_R : (i,n-1) \end{split}$$

**Memory access mapping** The following data accesses are performed in each of the iteration spaces:

• Distance calculation: for  $(i) \in I_D$ :

$$M_r(i) = \{P_i\}$$
  
 $M_w(i) = \{D_i\}$ 

• Neighbors Selection: for  $(i, j) \in I_R$ :

$$\begin{array}{rcl} M_r(i,j) &=& \{D_j\} \\ M_{mr}(i,j) &=& \{D[i:j]\} \\ M_w(i,n-1) &=& \{D_i,P_i\} \\ M_{mw}(i,n-1) &=& \{P[],D[]\} \end{array}$$

## Low-level implementation details

Only Distance Calculation (step 1) is implemented using OpenCL. Neighbors Selection (step 2) is implemented using a sequential host loop.

**Device code** Each work item calculates the distance from one point to the base point. See Algorithm 59.

Host code See Algorithm 60.

## Host data structures

```
• struct point:
    struct point
    {
        long y;
        long x;
    };
```

• struct point hP[n]:  $P_u \to$  hP[u]



Algorithm 59 The Nearest Neighbors (distance calculation) device code abstraction

Input: global struct point dP[] Output: global int dD[] Input: int n Input: int x Input: int y 1: kernel NNKERNEL (dP, dD, n, x, y) 2: if  $\gamma < n$  then 3: dD[ $\gamma$ ]  $\leftarrow \sqrt{(x - dP[\gamma].x)^2 + (y - dP[\gamma].y)^2}$ 4: end if 5: end kernel

- struct point  $hR[n]: R_u \to hR[u]$
- long hD[n]:  $D_u \to$  hD[u]
- struct point hB:  $D \to$  hB

## **Device data structures**

```
• struct point:
```

```
struct point
{
    long y;
    long x;
};
```

- struct point dP[n]:  $P_u \rightarrow dP[u]$
- struct point  $dR[n]: R_u \to dR[u]$
- long dD[n]:  $D_u \to dD[u]$

**Input datasets** The benchmark uses a set of pre-generated arrays of points. Each set contains 10691 points.

## **Partitioned iteration space**

• Distance Calculation (nnKernel):

$$P_D = \{(\gamma) : 0 \le \gamma < \Gamma = n\}$$
  
 $f(i) \to (i)$   
 $f^{-1}(\gamma) \to (\gamma)$ 



Algorithm 60 The Nearest Neighbors (distance calculation) host code abstraction ▷ Allocate host memory and device buffers. 1: **procedure** ALLOCATEMEMORY(n,k) 2:  $hP \leftarrow AllocateHostMemory(n)$  $hD \leftarrow ALLOCATEHOSTMEMORY(n)$ 3:  $hR \leftarrow AllocateHostMemory(k)$ 4:  $dP \leftarrow AllocateBuffer(n, READ_OLNY)$ 5:  $dD \leftarrow ALLOCATEBUFFER(n, WRITE_ONLY)$ 6: 7: end procedure 8: procedure INITIALIZEMEMORY(file)  $\triangleright$  Initialize host memory. 9:  $hP \leftarrow READFROMFILE(file, n)$ 10: end procedure 11: procedure SORTNEIGHBORS(hP, hD, k, n) for i in 0: (k-1) do 12: 13:  $\texttt{min} \gets \texttt{i}$ for j in i : (n-1) do 14: if hD[j] < hD[min] then 15:  $\min \leftarrow j$ 16: end if 17: end for 18: SWAP(hP[min], hP[i]) 19: SWAP(hD[min], hD[i]) 20: end for 21: 22: end procedure Input: int n ▷ Number of points in the input set. Input: int k ▷ Number of points in the output set. Input: struct point hB ⊳ Base point. ▷ File with pre-generated locations data. Input: const char \* file 23: ALLOCATEMEMORY(n,k) 24: INITIALIZEMEMORY(file) 25: nnProgram ← BUILDPROGRAM("nnProgram.cl") 26:  $nnKernel \leftarrow BUILDKERNEL(nnProgram, "nnKernel")$ 27: COPYTODEVICE(dP, hP, **True**, ∅) 28: SETKERNELARGUMENTS(nnKernel, dP, dD, n, B.x, B.y) 29: ENQUEUEKERNEL(nnKernel, (n), none, ∅) 30: COPYTOHOST(dD, hD, **True**,  $\emptyset$ ) 31: SORTNEIGHBORS(hP,hD,k,n) 32:  $hR \leftarrow hP[0:k-1]$ 





• Neighbors Section (Host code):

$$P_x = \{(i, j) : 0 \le i < k, i < j < n\}$$
$$f(i, j) \to (i, j)$$
$$f^{-1}(i, j) \to (i, j)$$

**Device memory access mapping** Each work-item performs the following reads and writes:

$$G_r(\gamma) = \{dP[\gamma]\}$$
  
 $G_w(\gamma) = \{dD[\gamma]\}$ 

Host memory access mapping

**Performance metrics** The performance of the algorithm implementation is evaluated using the following metrics:

- The total time for the kernel execution (computation time).
- The total memory transfer time from host to device.
- The total memory transfer time from device to host.

The performance is evaluated only for the part of the algorithm implemented in OpenCL (distance calculation).

**Validation mechanism** The benchmark has no built-in validation mechanisms. The benchmark can print the transformed matrices and the solution for the original system, which can be validated by external tools (not supplied).

**Target-specific optimizations** The following target-specific optimization are used in the algorithm implementation:

• **Consecutive memory access.** Work items with consecutive global ids in dimension 0 access consecutive memory locations which improves the memory system performance (caching, coalescing).

CARP-ARM-RP-001-v1.0





## **Target-specific optimizations opportunities**

• **Parallelizing Neighbors Selection.** Neighbors Selection (step 2) can be implemented using OpenCL for selecting a minimal element of *D*. The algorithm can be modified to avoid swap operations:

1: **procedure** SORTNEIGHBORS(hP,hD,hR,k,n)

for i in 0: (k-1) do 2:  $\texttt{min} \gets \texttt{0}$ 3: for j in i : (n-1) do 4: if hD[j] < hD[min] then 5:  $\texttt{min} \gets \texttt{j}$ 6: 7: end if end for 8:  $hD[min] \leftarrow MAX_DIST$ 9:  $hR[i] \leftarrow hP[min]$ 10: 11: end for 12: end procedure

The MAX\_DIST value can be defined as  $\sqrt{(x_{max} - x_{min})^2 + (y_{max} - y_{min})^2} + 1$ , where

- $x_{max} = 180$  (maximal longitude)
- $x_{min} = -180$  (minimal longitude)
- $y_{max} = 90$  (maximal latitude)
- $y_{min} = -90$  (minimal latitude)

*i.e.* the value that is guaranteed to be greater then any value, produced by distance calculation.





# 2.3.12 Needleman-Wunsch

## **High-level description**

The benchmark uses Needleman-Wunsch algorithm to align two DNA sequences. Having two DNA sequences A = aabc and B = abbc the alignment can be represented as follows:

The score of the alignment is determined using the similarity matrix *S*. For the alignment below the score would be:

$$S(a,*) + S(a,a) + S(b,b) + S(*,b) + S(c,c)$$

Abstract data structures The benchmark operates on the following data objects:

- *A* is *N* elements input DNA sequences (vector).
- *B* is *N* elements input DNA sequences (vector).
- *S* is an input similarity matrix.
- *C* is an  $(N+1) \times (N+1)$  intermediate score matrix.
- *T* is an  $(N+1) \times (N+1)$  intermediate traceback matrix:

$$T_{i,j} = \{\mathbf{up}, \mathbf{left}, \mathbf{diag}, \mathbf{done}\}$$

- *P* is an input penalty value.
- *R* is an output alignment (\* denotes gap in the sequence):

$$R_i = (a,b) : a \in A \cup \{*\}, b \in B \cup \{*\}$$

**Computation** The Needleman-Wunsch algorithm uses the dynamic programming technique to determine the alignment with maximal score:

$$\begin{array}{rcl} T_{0,0} &=& {\bf done} \\ T_{0,j} &=& {\bf left}, 0 \leq j \leq N \\ T_{i,0} &=& {\bf up}, 0 \leq i \leq N \\ C_{0,j} &=& -j \times P, 1 \leq j \leq N \\ C_{i,0} &=& -i \times P, 1 \leq i \leq N \end{array}$$

For every  $(i, j) : 1 \le i \le N, 1 \le j \le N$ :



$$T_{i,j} = \begin{cases} \mathbf{diag} & q_{i,j}^{d} = \max(q_{i,j}^{d}, q_{i,j}^{u}, q_{i,j}^{l}) \\ \mathbf{up} & q_{i,j}^{u} = \max(q_{i,j}^{d}, q_{i,j}^{u}, q_{i,j}^{l}) \\ \mathbf{left} & q_{i,j}^{l} = \max(q_{i,j}^{d}, q_{i,j}^{u}, q_{i,j}^{l}) \end{cases}$$

The alignment is determined using the traceback matrix:

$$i \leftarrow n$$
  

$$j \leftarrow n$$
  
while  $i \neq 0 \lor j \neq 0$  do  
if  $T_{i,j} =$  diag then  
 $i \leftarrow i-1$   
 $j \leftarrow j-1$   
 $R \leftarrow R + (A_i, B_j)$   
else if  $T_{i,j} =$  left then  
 $j \leftarrow j-1$   
 $R \leftarrow R + (*, B_j)$   
else if  $T_{i,j} =$  up then  
 $i \leftarrow i-1$   
 $R \leftarrow R + (A_i, *)$   
end if  
end while

# Iteration spaces and dependences

• Forward computation:

$$I_F = \{(i,j) : 1 \le i \le n, 1 \le j \le n\}$$

$$I_F: \begin{pmatrix} 1 & 0\\ 0 & 1\\ -1 & 0\\ 0 & -1 \end{pmatrix} \times \begin{pmatrix} i\\ j \end{pmatrix} \le \begin{pmatrix} n\\ n\\ -1\\ -1 \end{pmatrix}$$

• Traceback:

$$I_T = \{(\tau) : 0 \le \tau \le 2 \times n\}$$

The computation generates the following dependences:

$$\begin{split} \delta_{I_F} &= I_F : (i-1,j-1) \xrightarrow{t} I_F : (i,j) \land \\ &I_F : (i-1,j) \xrightarrow{t} I_F : (i,j) \land \\ &I_F : (i,j-1) \xrightarrow{t} I_F : (i,j) \land \\ \delta_{I_T} &= I_F : (1:n,1:n) \xrightarrow{t} I_T : (\tau) \land \\ &I_T : (\tau-1) \xrightarrow{t} I_T : (\tau) \end{split}$$





## **Memory access mapping** For $(i, j) \in I_F$ :

$$M_r(i,j) = \{C_{i-1,j-1}, C_{i-1,j}, C_{i,j-1}, A_{i-1}, B_{i-1}, S\}$$
  
$$M_w(i,j) = \{C_{i,j}, T_{i,j}\}$$

For  $(\tau) \in I_T$ :

$$egin{array}{rcl} M_r( au) &=& \{T_{i( au),j( au)}\}\ M_w( au) &=& \{R\} \end{array}$$

#### Low-level implementation details

Only the first part of the algorithm (intermediate score matrix computation) is implemented using OpenCL. The abstract iteration spaces is divided into two concrete iteration spaces, which process the upper-left and the lower-right parts of the score matrix.

**Device code** See Algorithms 61 and 62.

Host code See Algorithm 63.

#### Host data structures

- int hC[(n + 1) \* (n+1)]:  $C_{i,j} \rightarrow$  hC[i \* (n+1) + j]
- int hT[(n + 1) \* (n+1)]:  $T_{i,j} \rightarrow$  hT[i \* (n+1) + j]
- int hS[(n + 1) \* (n+1)]:  $S_{A_i,B_j} \rightarrow$  hS[i \* (n+1) + j]

**Device data structures** B is the block size (B = 16).

- global int dC[(n + 1) \* (n+1)]:  $C_{i,j} \rightarrow dC[i * (n+1) + j]$
- global int dS[(n + 1) \* (n+1)]:  $S_{A_i,B_i} \rightarrow dS[i * (n+1) + j]$
- local int dlC[(B + 1) \* (B+1)]:  $C_{i,j} \rightarrow dlC[(i \& B) \& B + j \& B]$  $C_{i,j}|_{i \mod B=0, j \mod B=0} \rightarrow dlC[B \& B + B]$  (Boundary).
- local int dls[B \* B]:  $S_{A_i,B_i} \rightarrow dls[(i \ \ B) \ \ * \ B + j \ \ B]$

**Input datasets** The benchmark uses randomly generated pair of 2048 elements DNA sequence. The size of the sequences can be changed via command line.



```
Algorithm 61 The Needleman Wunsch (Upper-left) device code abstraction.
Input: global float dS[]
Input/Output: global float dC[]
Input: local float dlC[]
Input: local float dls[]
Input: int P
                                                                                                      \triangleright Penalty.
Input: int Cols
                                                                                                         \triangleright n+1.
  1: kernel NWKERNEL1 (dS,dC,P,Cols)
         B \leftarrow 16
 2:
 3:
         idx_x \leftarrow v
 4:
         idx_v \leftarrow N - v - 1
         idx \leftarrow Cols \times B \times idx_v + B \times idx_x + \lambda + Cols + 1
 5:
         idx_n \leftarrow Cols \times B \times idx_v + B \times idx_x + \lambda + 1
 6:
 7:
         idx_w \leftarrow Cols \times B \times idx_v + B \times idx_x + Cols
         idx_{nw} \leftarrow Cols \times B \times idx_y + B \times idx_x
 8:
         if \lambda = 0 then
 9:
              dlC[0] \leftarrow dC[idx_{nw}]
10:
         end if
11:
         BARRIER(CLK_LOCAL_MEM_FENCE)
12:
13:
         for i in 0 : (B - 1) do
              dlS[i \times B + \lambda] \leftarrow dS[idx + Cols \times i]
14:
15:
         end for
16:
         BARRIER(CLK_LOCAL_MEM_FENCE)
                                                                            ▷ Unnecessary memory barrier.
         dlC[(\lambda + 1) \times (B + 1)] \leftarrow dC[idx_w + Cols \times \lambda]
17:
         BARRIER(CLK_LOCAL_MEM_FENCE)
                                                                            ▷ Unnecessary memory barrier.
18:
19:
         dlC[\lambda + 1] \leftarrow dC[idx_n]
20:
         BARRIER(CLK_LOCAL_MEM_FENCE)
21:
         for m in 0 : (B - 1) do
              if \lambda < m then
22:
                  tidx_x \leftarrow \lambda + 1
23:
                  \mathtt{tidx}_v \leftarrow \mathtt{m} - \lambda + 1
24:
25:
                  q_{nw} \leftarrow dlC[(tidx_v - 1) \times (B+1) + tidx_x - 1] + dlS[(tidx_v - 1) \times B + tidx_x - 1]
     1]
                  q_w \leftarrow dlC[(tidx_v) \times (B+1) + tidx_x - 1] - P
26:
                   q_l \leftarrow dlC[(tidx_v - 1) \times (B + 1) + tidx_x] - P
27:
28:
                  dlC[tidx_y \times (B+1) + tidx_x] \leftarrow max(q_{nw}, q_w, q_l)
29:
              end if
30:
              BARRIER(CLK_LOCAL_MEM_FENCE)
         end for
31:
32:
         BARRIER(CLK_LOCAL_MEM_FENCE)
                                                                            ▷ Unnecessary memory barrier.
```



```
Algorithm 61 The Needleman Wunsch (Upper-left) device code abstraction (continues).
```

for m in (B-2): 0: (-1) do 33: if  $\lambda \leq m$  then 34:  $\mathtt{tidx}_x \leftarrow \lambda + \mathtt{B} - \mathtt{m}$ 35:  $tidx_v \leftarrow B - \lambda$ 36:  $\mathbf{q}_{nw} \leftarrow \mathtt{dlC}[(\mathtt{tidx}_y - 1) \times (\mathtt{B} + 1) + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1) \times \mathtt{B} + \mathtt{tidx}_x - 1] + \mathtt{dlS}[(\mathtt{tidx}_y - 1] + \mathtt{dlS} + \mathtt{dlS}$ 37: 1]  $q_w \leftarrow dlC[(tidx_v) \times (B+1) + tidx_x - 1] - P$ 38:  $q_l \leftarrow dlC[(tidx_y - 1) \times (B + 1) + tidx_x] - P$ 39:  $dlC[tidx_y \times (B+1) + tidx_x] \leftarrow max(q_{nw}, q_w, q_l)$ 40: end if 41: BARRIER(CLK\_LOCAL\_MEM\_FENCE) 42: 43: end for **for** i **in** 0 : (B - 1) **do** 44: 45:  $dC[idx + Cols \times i] \leftarrow dlC[(i+1) \times (B+1) + \lambda + 1]$ end for 46: 47: end kernel

## Partitioned iteration space

• Upper-left part (nwKernel1), upper-left part of the block:

$$P_U = \{(b, v, \lambda, m) : 0 \le b \le n/B, 0 \le v < \Gamma = b, 0 \le \lambda < \Lambda = B, \lambda \le m < B\}$$

$$f(i,j) \to ((i-1)/B + (j-1)/B + 1, (j-1)/B, (j-1) \bmod (B), (i+j-2) \bmod (B))$$

$$f^{-1}(b, \mathbf{v}, \lambda, m) \rightarrow ((b - \mathbf{v} - 1) \times B + m - \lambda + 1, B \times \mathbf{v} + \lambda + 1)$$

• Upper-left part (nwKernel1), lower-right part of the block:

$$P_U = \{(b, v, \lambda, m) 0 \le b \le n/B, 0 \le v < \Gamma = b, 0 \le \lambda < \Lambda = B, \lambda \le m < B - 1\}$$

$$f(i,j) \rightarrow ((i-1)/B + (j-1)/B + 1, (j-1)/B, i \bmod (B) \times -1, (i+j) \bmod (B) \times -1)$$

$$f^{-1}(b, \mathbf{v}, \lambda, m) \to ((b - \mathbf{v} - 1) \times B + B - \lambda, B \times \mathbf{v} + \lambda + B - m)$$

• Lower-right part (nwKernel2), upper-left part of the block:

$$P_U = \{(b, v, \lambda, m) | \le b < n/B, 0 \le v < \Gamma = b, 0 \le \lambda < \Lambda = B, \lambda \le m < B\}$$

$$\begin{array}{rcl} f(i,j) & \to & (2 \times n/B - ((i-1)/B + (j-1)/B), n/B - 1 - (i-1)/B, \\ & (j-1) \ \mathrm{mod} \ (B), (i+j-2) \ \mathrm{mod} \ (B)) \end{array}$$

CARP-ARM-RP-001-v1.0


Algorithm 62 The Needleman Wunsch (Lower-right) device code abstraction.

```
Input: global float dS[]
Input/Output: global float dC[]
Input: local float dlC[]
Input: local float dls[]
Input: int P
                                                                                                                  ⊳ Penalty.
Input: int Cols
                                                                                                                     \triangleright n+1.
Input: int Bw
                                                                                                                      \triangleright n/B
  1: kernel NWKERNEL2 (dS, dC, P, Cols, Bw)
          B \leftarrow 16
 2:
 3:
          \operatorname{idx}_x \leftarrow \nu + \operatorname{Bw} - N
 4:
          idx_v \leftarrow Bw - v - 1
          \mathtt{idx} \leftarrow \mathtt{Cols} \times \mathtt{B} \times \mathtt{idx}_v + \mathtt{B} \times \mathtt{idx}_x + \lambda + \mathtt{Cols} + 1
 5:
          \operatorname{idx}_n \leftarrow \operatorname{Cols} \times \operatorname{B} \times \operatorname{idx}_v + \operatorname{B} \times \operatorname{idx}_x + \lambda + 1
 6:
          idx_w \leftarrow Cols \times B \times idx_v + B \times idx_x + Cols
 7:
          idx_{nw} \leftarrow Cols \times B \times idx_v + B \times idx_x
 8:
 9:
          if \lambda = 0 then
               dlC[0] \leftarrow dC[idx_{nw}]
10:
          end if
11:
          BARRIER(CLK_LOCAL_MEM_FENCE)
12:
          for i in 0 : (B - 1) do
13:
14:
               dlS[i \times B + \lambda] \leftarrow dS[idx + Cols \times i]
          end for
15:
          BARRIER(CLK_LOCAL_MEM_FENCE)
                                                                                     ▷ Unnecessary memory barrier.
16:
17:
          dlC[(\lambda + 1) \times (B + 1)] \leftarrow dC[idx_w + Cols \times \lambda]
          BARRIER(CLK_LOCAL_MEM_FENCE)
18:
                                                                                     ▷ Unnecessary memory barrier.
19:
          dlC[\lambda + 1] \leftarrow dC[idx_n]
          BARRIER(CLK_LOCAL_MEM_FENCE)
20:
          for m in 0 : (B - 1) do
21:
22:
               if \lambda \leq m then
                    tidx_x \leftarrow \lambda + 1
23:
                    tidx_v \leftarrow m - \lambda + 1
24:
                     q_{nw} \leftarrow dlC[(tidx_v - 1) \times (B+1) + tidx_x - 1] + dlS[(tidx_v - 1) \times B + tidx_x - 1]
25:
      1]
                    q_w \leftarrow dlC[(tidx_v) \times (B+1) + tidx_x - 1] - P
26:
                     q_l \leftarrow dlC[(tidx_v - 1) \times (B + 1) + tidx_x] - P
27:
                     dlC[tidx_v \times (B+1) + tidx_x] \leftarrow max(q_{nw}, q_w, q_l)
28:
29:
               end if
               BARRIER(CLK_LOCAL_MEM_FENCE)
30:
          end for
31:
32:
          BARRIER(CLK_LOCAL_MEM_FENCE)
                                                                                     ▷ Unnecessary memory barrier.
```



```
Algorithm 62 The Needleman Wunsch (Lower-right) device code abstraction (continues).
```

33: for m in (B-2): 0: (-1) do if  $\lambda \leq m$  then 34:  $\mathtt{tidx}_{\mathtt{x}} \gets \lambda + \mathtt{B} - \mathtt{m}$ 35:  $\texttt{tidx}_v \gets \texttt{B} - \lambda$ 36:  $\mathtt{q}_{\mathit{nw}} \gets \mathtt{dlC}[(\mathtt{tidx}_y-1)\times(\mathtt{B}+1)+\mathtt{tidx}_x-1]+\mathtt{dlS}[(\mathtt{tidx}_y-1)\times\mathtt{B}+\mathtt{tidx}_x-1]$ 37: 1]  $q_w \leftarrow dlC[(tidx_v) \times (B+1) + tidx_x - 1] - P$ 38:  $q_l \leftarrow dlC[(tidx_y - 1) \times (B + 1) + tidx_x] - P$ 39:  $dlC[tidx_v \times (B+1) + tidx_x] \leftarrow max(q_{nw}, q_w, q_l)$ 40: 41: end if 42: BARRIER(CLK\_LOCAL\_MEM\_FENCE) end for 43: 44: for i in 0 : (B - 1) do 45:  $\texttt{dC}[\texttt{idx} + \texttt{Cols} \times \texttt{i}] \leftarrow \texttt{dlC}[(\texttt{i} + 1) \times (\texttt{B} + 1) + \lambda + 1]$ end for 46: 47: end kernel

$$f^{-1}(b, \mathbf{v}, \lambda, m) \rightarrow ((n/B - \mathbf{v} - 1) \times B + m - \lambda + 1, B \times (\mathbf{v} + n/B - b) + \lambda + 1)$$

• Lower-right part (nwKernel2), lower-right part of the block:

$$P_U = \{(b, v, \lambda, m) | \le b < n/B, 0 \le v < \Gamma = b, 0 \le \lambda < \Lambda = B, \lambda \le m < B - 1\}$$

$$f(i,j) \to (2 \times n/B - ((i-1)/B + (j-1)/B), n/B - 1 - (i-1)/B, , i \bmod (B) \times -1, (i+j) \bmod (B) \times -1)$$

$$f^{-1}(b, \mathbf{v}, \lambda, m) \to ((n/B - \mathbf{v} - 1) \times B + B - \lambda, B \times (\mathbf{v} + n/B - b) + \lambda + B - m)$$

### **Device memory access mapping**

- Upper-left part (nwKernel1):
  - Global to local memory transfer:



Algorithm 63 The Needleman Wunsch host code abstraction.

```
1: procedure TRACEBACK(n)
             \mathtt{i} \gets \mathtt{n}
 2:
 3:
             \mathtt{j} \gets \mathtt{n}
             while \mathtt{i} \neq 0, \mathtt{j} \neq 0 do
                                                                                     ▷ Bug. Logical 'OR' should be used instead.
 4:
                   if i = 0 then
 5:
 6:
                          nw = limit
                          n = limit
 7:
                          w = hC[j-1]
 8:
                   else if j = 0 then
 9:
                          nw = limit
10:
11:
                          n = hC[(i-1) \times (n+1)]
                          w = limit
12:
                   else
13:
                          \texttt{nw} = \texttt{hC}[(\texttt{i}-1) \times (\texttt{n}+1) + \texttt{j}-1]
14:
                          \mathtt{n} = \mathtt{h}\mathtt{C}[(\mathtt{i} - 1) \times (\mathtt{n} + 1) + \mathtt{j}]
15:
                          w = hC[i \times (n+1) + j - 1]
16:
17:
                   end if
                   \texttt{nw}_{\texttt{new}} \leftarrow \texttt{nw} + \texttt{hS}[\texttt{i} \times (\texttt{n}+1) + \texttt{j}]
18:
                   \mathtt{w}_{\mathtt{new}} \gets \mathtt{w} - \mathtt{P}
19:
                   \mathtt{n_{new}} \gets \mathtt{n} - \mathtt{P}
20:
                   \texttt{trace} = \max(\texttt{nw}_{\texttt{new}}, \texttt{n}_{\texttt{new}}, \texttt{w}_{\texttt{new}})
21:
22:
                   if trace = nw_{new} then
23:
                          \mathtt{i} \gets \mathtt{i} - 1
                          j \leftarrow j - 1
24:
                   else if trace = w_{new} then
25:
                          \texttt{j} \leftarrow \texttt{j}-\texttt{l}
26:
27:
                   else
28:
                          i \leftarrow i - 1
                   end if
29:
             end while
30:
31: end procedure
```



```
Algorithm 63 The Needleman Wunsch host code abstraction (continues).
32: procedure ALLOCATEMEMORY(n)
                                                   ▷ Allocate host memory and device buffers.
        hC \leftarrow AllocateHostMemory((n+1) \times (n+1))
33:
        hS \leftarrow ALLOCATEHOSTMEMORY(n \times n)
34:
        dC \leftarrow ALLOCATEBUFFER((n+1) \times (n+1), READ_WRITE)
35:
        dS \leftarrow AllocateBuffer(n \times n, READ_WRITE)
36:
37: end procedure
38: procedure INITIALIZEMEMORY(n, P)
                                                                      ▷ Initialize host memory.
        for i in 1:n do
39:
            \texttt{hC[i]} \gets \texttt{i} \times \texttt{P}
40:
            \texttt{hC}[\texttt{i} \times (\texttt{n}+1)] \leftarrow \texttt{i} \times \texttt{P}
41:
42:
        end for
        A \leftarrow RANDOMELEMENTS(n)
                                                                           ▷ Temporary vector.
43:
        B \leftarrow RANDOMELEMENTS(n)
                                                                           ▷ Temporary vector.
44:
        for i in 0: (n-1) do
45:
            for j in 0 : (n-1) do
46:
47:
               hS[i \times n + j] \leftarrow S[A[i]][B[j]]
            end for
48:
        end for
49:
50: end procedure
Input: int n
                                                                         ▷ Input sequence size.
Input: int P
                                                                                      ⊳ Penalty.
Input: int S[24][24]
                                                                            ▷ Similarity matrix.
51: ALLOCATEMEMORY(n)
52: INITIALIZEMEMORY(n,P)
53: nwProgram ← BUILDPROGRAM("ludProgram.cl")
54: nwKernel1 ~ BUILDKERNEL(nwProgram, "nwKernel1")
55: nwKernel2 ← BUILDKERNEL(nwProgram, "nwKernel2")
56: B \leftarrow 16
                                                                                   ▷ Block size.
57: COPYTODEVICE(dS,hS,True,∅)
58: COPYTODEVICE(dC, hC, True, ∅)
59: for blk in 1 : (n/B) do
        SETKERNELARGUMENTS(nwKernel1, dS, dC, P, n + 1)
60:
61:
        ENQUEUEKERNEL(nwKernel1, (B \times blk), (B), \emptyset)
62: end for
63: for blk in (n/B-1): 1: (-1) do
        SetKernelArguments(nwKernel2, dS, dC, P, n + 1, B/B)
64:
        ENQUEUEKERNEL(nwKernel2, (B \times blk), (B), \emptyset)
65:
66: end for
67: COPYTOHOST(dC,hC, True, ∅)
68: TRACEBACK(n)
```





– Upper-left part of the block:

$$\begin{array}{lll} L_r(b,\mathbf{v},\lambda,m) &=& \{ \mathtt{dlC}[(m-\lambda)\times(B+1)+\lambda], \mathtt{dlC}[(m-\lambda+1)\times(B+1)+\lambda], \\ && \mathtt{dlC}[(m-\lambda)\times(B+1)+\lambda+1], \mathtt{dlS}[(m-\lambda)\times B+\lambda] \} \\ L_w(b,\mathbf{v},\lambda,m) &=& \{ \mathtt{dlC}[(m-\lambda+1)\times(B+1)+\lambda+1], \} \end{array}$$

- Lower-right part of the block:

- Local to global memory transfer:

$$\begin{array}{lll} G_w(b,\mathbf{v},\lambda) &=& \{ \mathrm{dC}[((b-\mathbf{v}-1)\times B\times (n+1)+\mathbf{v}\times B+\lambda+n+2+i\times (n+1)] \} \\ L_r(b,\mathbf{v},\lambda) &=& \{ \mathrm{dlC}[(\mathrm{i}+1)\times (B+1)+\lambda+1] \} \end{array}$$

- Lower-right part (nwKernel2):
  - Global to local memory transfer:

- Upper-left part of the block:

$$\begin{array}{lll} L_r(b,\mathbf{v},\lambda,m) &=& \{ \mathtt{dlC}[(m-\lambda)\times(B+1)+\lambda], \mathtt{dlC}[(m-\lambda+1)\times(B+1)+\lambda],\\ && \mathtt{dlC}[(m-\lambda)\times(B+1)+\lambda+1], \mathtt{dlS}[(m-\lambda)\times B+\lambda] \} \\ L_w(b,\mathbf{v},\lambda,m) &=& \{ \mathtt{dlC}[(m-\lambda+1)\times(B+1)+\lambda+1], \} \end{array}$$





- Lower-right part of the block:

$$\begin{array}{lll} L_r(b,\mathbf{v},\lambda,m) &=& \{ \mathtt{dlC}[(B-\lambda-1)\times(B+1)+\lambda+B-m-1], \\ && \mathtt{dlC}[(B-\lambda)\times(B+1)+\lambda+B-m-1], \\ && \mathtt{dlC}[(B-\lambda-1)\times(B+1)+\lambda+B-m], \\ && \mathtt{dlS}[(B-\lambda-1)\times B+\lambda+B-m-1], \} \\ L_w(b,\mathbf{v},\lambda,m) &=& \{ \mathtt{dlC}[(B-\lambda)\times(B+1)+\lambda+B-m] \} \end{array}$$

- Local to global memory transfer:

$$G_w(b, \mathbf{v}, \lambda) = \{ dC[((n/B - \mathbf{v} - 1) \times B \times (n+1) + (\mathbf{v} + n/B - b) \times B + \lambda + n + 2 + i \times (n+1)] \}$$
  
$$L_r(b, \mathbf{v}, \lambda) = \{ dlC[(i+1) \times (B+1) + \lambda + 1] \}$$

**Host memory access mapping** On each iteration on the host the following memory accesses are performed:

$$H_r(\tau) = H_r(\mathbf{i}(\tau), \mathbf{j}(\tau)) = \{ hC[(\mathbf{i}-1) \times (\mathbf{n}+1) + \mathbf{j}-1], hC[(\mathbf{i}-1) \times (\mathbf{n}+1) + \mathbf{j}], \\ hC[\mathbf{i} \times (\mathbf{n}+1) + \mathbf{j}-1], hS[\mathbf{i} \times (\mathbf{n}+1) + \mathbf{j}] \}$$

**Performance metrics** The performance is evaluated using the following metrics:

- The total time for all nwKernel1 launches.
- The total time for all nwKernel2 launches.
- The total time for all kernel launches (computation time).
- The total queuing time.

**Validation mechanism** The benchmark has no built-in validation mechanisms. The benchmark can print the traceback values which can be validated by external tools (not supplied).

**Target-specific optimizations** The following target-specific optimization are used in the algorithm implementation:

- **Consecutive memory access.** Work items with consecutive global ids in dimension 1 access consecutive memory locations which improves the memory system performance (caching, coalescing).
- Local memory usage. In order to decrease the number of global memory accesses each work group allocates a local buffer, which is used to perform all computations. This technique also increases data locality since the data layout can be changed while copying from global to local memory and vise versa.





# **Target-specific optimizations opportunities**

- **Global memory usage.** In some OpenCL implementations local and global memories resident in the same memory spaces, so copying data from global to local memory just introduces additional overhead.
- Avoiding unnecessary synchronizations. Not all memory barriers used in the algorithm implementation are necessary (see device code abstraction). Removing such barriers can decrease kernel execution time.





# 2.3.13 Particle Filter

## **High-level description**

The benchmark uses particle filter algorithm to track object movement in the video sequence between consecutive frames.

Abstract data structures The benchmark operates on the following data objects:

- V is an  $N_f$ -frame input video sequence.
- $F = V_i$  is a  $N_x \times N_y$  video frame.
- $N_p$  is an input number of particles to be used for approximation.
- *R* is an input radius of the particle.
- $P[N_p]$  is an intermediate array of particles coordinates within the video frame.
- $P_i^x$  is the *x* coordinate of the *i*<sup>th</sup> particle.
- $L[N_p]$  is an intermediate array of particles likelihoods.
- $S[N_p]$  is an intermediate array of point sets.
- $W[N_p]$  is an intermediate array of particles weights.
- $C[N_p], U[N_p]$  are intermediate arrays.
- E = 100 is an empty point value.
- O = 228 is an object point value.

**Computation** The initial data is initialized as follows:

$$P_i = (N_x/2, (N_y/2)), 0 \le i < N_p$$

For each frame  $F = V_k$ ,  $0 \le k < N_f$  the following computations are performed:

1. Particles locations recalculation:

$$P_i = (P_i^x + 1 + 5 \times \operatorname{rand}(), P_i^y - 2 + 2 \times \operatorname{rand}()), 0 \le i < N_p$$

2. Points sets generation:

$$S_i = \{(x, y) : (x - P_i^x)^2 + (y - P_i^y)^2 \le R^2\}$$

3. Likelihood calculation:

$$L_{i} = \left(\sum_{(x,y):\in S_{i}} ((I_{x,y} - E)^{2} - (I_{x,y} - O)^{2})/50\right) / |S_{i}|$$



4. Weights calculation:

$$W_i = \exp(L_i) / N_p, 0 \le i < N_p$$

5. Weights normalization:

$$egin{array}{rcl} S_w &=& \displaystyle{\sum_{0 \leq i < N_p} W_i} \ W_i &=& \displaystyle{W_i / S_w, 0 \leq i < N_p} \end{array}$$

6. Intermediate arrays (C, U) calculation:

$$\begin{array}{rcl} U_{\texttt{tmp}} &=& \textbf{rand}() \\ U_i &=& i/N_p + U_{\texttt{tmp}}, 0 \leq i < N_p \\ C_i &=& \displaystyle\sum_{0 \leq j \leq i} W_j, 0 \leq i < N_p \end{array}$$

7. Particles locations update:

$$\begin{array}{lll} P' &=& P \\ P_i^x &=& {P'}_j^x : C_j > U_i \wedge C_{j-1} \le U_i, 0 \le i < N_p \\ P_i^y &=& {P'}_j^y : C_j > U_i \wedge C_{j-1} \le U_i, 0 \le i < N_p \end{array}$$

The **rand**() function call generates the normally distributed random numbers.

Iteration spaces and dependences This algorithm requires several iteration spaces:

1. Particles locations recalculation:

$$I_P = \{(k,i) : 0 \le k < N_f, 0 \le i < N_p\}$$

2. Points sets generation:

$$I_{S} = \{(k,i) : 0 \le k < N_{f}, 0 \le i < N_{p}\}$$

3. Likelihood calculation:

$$I_L = \{(k,i) : 0 \le k < N_f, 0 \le i < N_p\}$$

4. Weights calculation:

$$I_W = \{(k,i) : 0 \le k < N_f, 0 \le i < N_p\}$$

5. Weights sum calculation:

$$I_{S_w} = \{(k,i) : 0 \le k < N_f, 0 \le i < N_p\}$$





6. Weights normalization:

$$I_N = \{(k,i) : 0 \le k < N_f, 0 \le i < N_p\}$$

7. Intermediate array U calculation:

$$I_U = \{(k,i) : 0 \le k < N_f, 0 \le i < N_p\}$$

8. Intermediate array *C* calculation:

$$I_C = \{(k, i, j) : 0 \le k < N_f, 0 \le i < N_p, 0 \le j \le i\}$$

9. Particles locations update:

$$I_F = \{(k, i, j) : 0 \le k < N_f, 0 \le i < N_p, 0 \le j < N_p\}$$

The computation generates the following dependences:

$$\begin{split} \delta_{I_{P}} &= I_{F} : (k-1,i,0:(N_{p}-1)) \xrightarrow{t} I_{P} : (k,i) \wedge \\ &I_{S}(k-1,i) \xrightarrow{a} I_{P} : (k,i) \\ \delta_{I_{S}} &= I_{P} : (k,i) \xrightarrow{t} I_{S} : (k,i) \wedge \\ &I_{L}(k-1,i) \xrightarrow{a} I_{S} : (k,i) \\ \delta_{I_{L}} &= I_{S} : (k,i) \xrightarrow{t} I_{L} : (k,i) \wedge \\ &I_{W}(k-1,i) \xrightarrow{a} I_{L} : (k,i) \\ \delta_{I_{W}} &= I_{L} : (k,i) \xrightarrow{t} I_{W} : (k,i) \wedge \\ &I_{N}(k-1,i) \xrightarrow{a} I_{W} : (k,i) \\ \delta_{I_{S_{w}}} &= I_{W} : (k,i) \xrightarrow{t} I_{S_{w}} : (k,i) \\ \delta_{I_{S_{w}}} &= I_{W} : (k,i) \xrightarrow{t} I_{S_{w}} : (k,i) \\ \delta_{I_{N}} &= I_{W} : (k,i) \xrightarrow{t} I_{N} : (k,i) \wedge \\ &I_{N}(k-1,0:(N_{p}-1)) \xrightarrow{a} I_{S_{w}} : (k,i) \\ \delta_{I_{N}} &= I_{F} : (k-1,i,0:(N_{p}-1)) \xrightarrow{a} I_{N} : (k,i) \\ \delta_{I_{U}} &= I_{F} : (k-1,i,0:(N_{p}-1)) \xrightarrow{a} I_{U} : (k,i) \wedge \\ &I_{F}(k-1,0:(N_{p}-1),i) \xrightarrow{a} I_{C} : (k,i,j) \wedge \\ &I_{F}(k-1,0:(N_{p}-1),i) \xrightarrow{a} I_{C} : (k,i,j) \wedge \\ &I_{C} : (k,j,0:j) \xrightarrow{t} I_{F}(k,i,j) \wedge \\ &I_{C} : (k,j,0:j) \xrightarrow{t} I_{F}(k,i,j) \wedge \\ &I_{P} : (k,i) \xrightarrow{a} I_{F}(k,i,0:(N_{p}-1)) \end{split}$$

### Memory access mapping

• Particles locations recalculation:

$$M_r(k,i) = \{P_i\}$$
  
 $M_w(k,i) = \{P_i\}$ 





• Points sets generation:

$$M_r(k,i) = \{P_i\}$$
  
 $M_w(k,i) = \{S_i\}$ 

• Likelihood calculation:

$$M_r(k,i) = \{S_i\}$$
  
 $M_w(k,i) = \{L_i\}$ 

• Weights calculation:

$$M_r(k,i) = \{L_i\}$$
  
 $M_w(k,i) = \{W_i\}$ 

• Weights sum calculation:

$$M_r(k,i) = \{W_i\}$$

• Weights normalization:

$$M_r(k,i) = \{W_{0:(0-N_p-1)}\}$$
  
 $M_w(k,i) = \{W_i\}$ 

• Intermediate array U calculation:

$$M_w(k,i) = \{U_i\}$$

• Intermediate array *C* calculation:

$$M_r(k,i,j) = \{W_j\}$$
  
 $M_w(k,i,j) = \{C_i\}$ 

• Particles locations update:

$$M_r(k, i, j) = \{C_j, U_i\}$$
  
 $M_w(k, i, j) = \{P_i\}$ 

### Low-level implementation details

The detailed description of the algorithm implementation can be found in [10]. The optimized implementation is described above, *i.e.* all computations are performed on device.

**Device code** See Algorithm 64, Algorithm 65, Algorithm 66 and Algorithm 67.

Host code See Algorithm 68.



```
Algorithm 64 The Particle Filter (Likelihood recalculation) device code abstraction.
  1: function CALCLIKELIHOODSUM(dV,dS,SSize,γ)
         \texttt{sum} \gets 0
 2:
 3:
         for x in 0: (SSize -1) do
               \texttt{sum} \leftarrow \texttt{sum} + ((\texttt{dV}[\texttt{dS}[\texttt{SSize} \times \gamma + \texttt{x}]] - 100)^2 - (\texttt{dV}[\texttt{dS}[\texttt{SSize} \times \gamma + \texttt{x}]] - 100)^2
 4:
     (228)^2)/50
 5:
         end for
         return sum
 6:
 7: end function
Input: int Np
                                                                                         ▷ Number of particles.
Input: int Ny
Input: int Nf
                                                                                                 \triangleright Size of the S_i.
Input: int SSize
Input: int VSize
                                                                                                 \triangleright N_f \times N_x \times N_y
Input: int k
Input: global int dS[]
Input: global int dV[]
Input: global int dSbase[]
Output: global float dX[]
Output: global float dY[]
Input: global float dX2[]
Input: global float dY2[]
Input/Output: global float dW[]
Output: global int dL[]
Output: global int dSums[]
Local: local int dlB[]
 8: kernel PFKERNEL1 (Np, Ny, Nf, SSize, VSize, k, dS, dV, dSbase, dX, dY, dW, dL, dSums, dX2, dY2)
 9:
         if \gamma < Np then
              dX[\gamma] \leftarrow dX2[\gamma]
10:
              dY[\gamma] \leftarrow dY2[\gamma]
11:
              dX[\gamma] \leftarrow dX[\gamma] + 1 + 5 \times \text{RAND}()
12:
              dY[\gamma] \leftarrow dY[\gamma] - 2 + 2 \times RAND()
13:
         end if
14:
         BARRIER(CLK_GLOBAL_MEM_FENCE)
15:
16:
         if \gamma < Np then
              for y in 0: (SSize -1) do
17:
                   idxX \leftarrow dX[\gamma] + dSbase[2 \times y + 1]
18:
                   idxY \leftarrow dY[\gamma] + dSbase[2 \times y]
19:
                   dS[\gamma \times SSize + y] \leftarrow idxX \times Ny \times Nf + IdxY \times Nf + k
20:
                   if dS[\gamma \times SSize + y] \ge VSize then
21:
                       dS[\gamma \times SSize + y] \leftarrow 0
22:
                   end if
23:
              end for
24:
25:
              dL[\gamma] \leftarrow CALCLIKELIHOODSUM(dV, dS, SSize, \gamma)
26:
              dL[\gamma] \leftarrow dL[\gamma]/SSize
              dW[\gamma] \leftarrow dW[\gamma] \times \exp(dL[\gamma])
27:
         end if
28:
```



Algorithm 64 The Particle Filter (Likelihood recalculation) device code abstraction (continues).

29:	BARRIER(CLK_GLOBAL_MEM_FENCE CLK_LOCAL_MEM_FENCE)
30:	if $\gamma < Np$ then
31:	$\texttt{dlB}[\lambda] \gets \texttt{dW}[\gamma]$
32:	end if
33:	BARRIER(CLK_LOCAL_MEM_FENCE)
34:	$\mathbf{s} \leftarrow \Lambda/2$
35:	while $s > 0$ do
36:	if $\lambda < s$ then
37:	$\texttt{dlB}[\lambda] \gets \texttt{dlB}[\lambda] + \texttt{dlB}[\lambda + \texttt{s}]$
38:	end if
39:	BARRIER(CLK_GLOBAL_MEM_FENCE)
40:	if $\lambda = 0$ then
41:	$\texttt{dSums}[\boldsymbol{\nu}] \gets \texttt{dlB}[0]$
42:	end if
43:	$\mathbf{s} \leftarrow \mathbf{s}/2$
44:	end while
45:	end kernel

Algorithm 65 The Particle Filter (Sum) device code abstr	raction.
Input: int Np	▷ Number of particles.
Input/Output: global float dSums[]	
1: <b>kernel</b> PFKERNEL2 (Np,dSums)	
2: <b>if</b> $\gamma = 0$ <b>then</b>	
3: $sum \leftarrow 0$	
4: <b>for</b> x <b>in</b> $0: (Np/512 - 1)$ <b>do</b>	
5: $sum \leftarrow sum + dSums[x]$	
6: end for	
7: $dSums[0] \leftarrow sum$	
8: end if	
9: end kernel	



Algorithm 66 The Particle Filter (Normalize Weights) device code abstraction.

```
1: procedure CDFCALC(dC, dW, Np)
        \texttt{dC}[0] \gets \texttt{dW}[0]
 2:
        for x in 0 : (Np - 1) do
 3:
            dC[x] \leftarrow dW[x] + dC[x-1]
 4:
 5:
        end for
 6: end procedure
Input: int Np
                                                                              ▷ Number of particles.
Input: global float dSums[]
Output: global float dC[]
Output: global float dU[]
Input/Output: global float dW[]
 7: kernel PFKERNEL3 (Np,dW,dSums,dC,dU)
 8:
        if \lambda = 0 then
 9:
            \texttt{sum} \leftarrow \texttt{dSums}[0]
10:
        end if
        BARRIER(CLK_LOCAL_MEM_FENCE)
11:
        if \gamma < Np then
12:
13:
            dW[i] \leftarrow dW[i]/sum
        end if
14:
        BARRIER(CLK_GLOBAL_MEM_FENCE)
15:
16:
        if \gamma = 0 then
            CDFCALC(dC,dW,Np)
17:
            dU[0] \leftarrow RAND() / Np
18:
19:
        end if
        BARRIER(CLK_GLOBAL_MEM_FENCE)
20:
21:
        if \lambda = 0 then
            \mathtt{u1} \leftarrow \mathtt{dU}[0]
22:
        end if
23:
        BARRIER(CLK_LOCAL_MEM_FENCE)
24:
        if \gamma < Np then
25:
            dU[i] \leftarrow u1 + i/Np
26:
        end if
27:
28: end kernel
```



```
Algorithm 67 The Particle Filter (Find Index) device code abstraction.
Input: global float dU[]
Input: global float dX[]
Input: global float dY[]
Input: global float dC[]
Output: global float dX2[]
Output: global float dY2[]
Output: global float dW[]
Input: int Np
                                                                                  ▷ Number of particles.
 1: kernel PFKERNEL4 (Np, dU, dC, dX, dY, dX2, dY2, dW)
         if \gamma < Np then
 2:
             \mathtt{idx} \gets -1
 3:
 4:
             for x in 0 : (Np - 1) do
                 if dC[x] \le u[\gamma] then
 5:
                      \mathtt{idx} \gets \mathtt{x}
 6:
                      break
 7:
                 end if
 8:
 9:
             end for
10:
             if idx = -1 then
                  \mathtt{idx} \gets \mathtt{Np} - 1
11:
12:
             end if
             dX2[\gamma] \leftarrow dX[idx]
13:
             \texttt{dY2}[\gamma] \gets \texttt{dY}[\texttt{idx}]
14:
             \texttt{dW}[\gamma] \leftarrow 1/\texttt{Np}
15:
         end if
16:
         BARRIER(CLK_GLOBAL_MEM_FENCE)
17:
18: end kernel
```



Algorithm 68 The Particle Filter host code abstraction.				
1:	<pre>procedure AllocateMemory(Nx,Ny,Nf,Np,SSize)</pre>	▷ Allocate host memory and		
	device buffers.			
2:	$\texttt{hV} \leftarrow \texttt{AllocateHostMemory}(\texttt{Nx} \times \texttt{Ny} \times \texttt{Nf})$			
3:	$\texttt{hX} \leftarrow \texttt{AllocateHostMemory(Np)}$			
4:	$\texttt{hY} \leftarrow \texttt{AllocateHostMemory(Np)}$			
5:	$\texttt{hW} \leftarrow \texttt{AllocateHostMemory(Np)}$			
6:	$\texttt{hSbase} \leftarrow \texttt{ALLOCATEHOSTMEMORY}(\texttt{SSize} \times 2)$			
7:	$\mathtt{dX} \leftarrow \mathtt{AllocateBuffer(Np, \mathtt{READ\_WRITE})}$			
8:	$\mathtt{dY} \leftarrow \mathtt{AllocateBuffer(Np, \mathtt{READ\_WRITE})}$			
9:	$\mathtt{dX2} \leftarrow \mathtt{AllocateBuffer(Np, \mathtt{READ\_WRITE})}$			
10:	$\texttt{dY2} \leftarrow \texttt{ALLOCATEBUFFER(Np, \texttt{READ}\_\texttt{WRITE})}$			
11:	$\texttt{dC} \leftarrow \texttt{ALLOCATEBUFFER}(\texttt{Np},\texttt{READ\_WRITE})$			
12:	$\texttt{dU} \leftarrow \texttt{ALLOCATEBUFFER}(\texttt{Np}, \texttt{READ\_WRITE})$			
13:	$\texttt{dL} \leftarrow \texttt{ALLOCATEBUFFER}(\texttt{Np},\texttt{READ\_WRITE})$			
14:	$\texttt{dW} \leftarrow \texttt{AllocateBuffer(Np, \texttt{READ}\_WRITE)}$			
15:	$\texttt{dV} \leftarrow \texttt{AllocateBuffer}(\texttt{Nx} \times \texttt{Ny} \times \texttt{Nf}, \texttt{READ\_WRITE})$			
16:	$\texttt{dSbase} \leftarrow \texttt{ALLOCATEBUFFER}(\texttt{SSize} \times 2, \texttt{READ\_WRITE})$	)		
17:	$\texttt{dS} \leftarrow \texttt{ALLOCATEBUFFER}(\texttt{SSize} \times \texttt{Np}, \texttt{READ\_WRITE})$			
18:	$\texttt{dSums} \leftarrow \texttt{ALLOCATEBUFFER(Np, \texttt{READ}\_\texttt{WRITE})}$			
19:	end procedure			
20.	procedure INITIALIZEMEMORY(Ny Ny Nf No SSize R)	⊳ Initialize host memory		
20. 21·	for i in $0 \cdot (Nn - 1)$ do	Finitualize nost memory.		
21. 22·	$hX[i] \leftarrow Nx/2$			
22. 23.	$hY[i] \leftarrow Ny/2$			
23. 24·	$hW[i] \leftarrow 1/Np$			
25:	end for			
26:	$hV \leftarrow RANDOMELEMENTS(Nx \times Nv \times Nf)$			
27:	$idx \leftarrow 0$			
28:	for $\mathbf{x}$ in $0:(2 \times R - 2)$ do			
29:	for y in $0: (2 \times R - 2)$ do			
30:	if $(x - R + 1)^2 + (y - (R + 1)^2 < R^2$ then			
31:	$hSbase[idx] \leftarrow y$			
32:	$\texttt{hSbase}[\texttt{idx}+1] \leftarrow \texttt{x}$			
33:	$\mathtt{idx} \leftarrow \mathtt{idx} + 2$			
34:	end if			
35:	end for			
36:	end for			
37:	end procedure			



Algorithm 68 The Particle Filter host code abstraction (continues). Input: int Nx ▷ Input video sequence width. Input: int Ny ▷ Input video sequence height. Input: int Nf ▷ Input video sequence number of frames. Input: int Np ▷ Number of particles to use for approximation. Input: int R ▷ Target object radius. 38: SSize  $\leftarrow$  GETCIRCLEAREA(R) 39:  $VSize \leftarrow Nx \times Ny \times Nf$ 40: ALLOCATEMEMORY(Nx, Ny, Nf, Np, SSize) 41: INITIALIZEMEMORY(Nx, Ny, Nf, Np, SSize, R) 42: pfProgram ← BUILDPROGRAM("pfProgram.cl") 43: pfKernel1  $\leftarrow$  BUILDKERNEL(pfProgram, "pfKernel1") 44: pfKernel2 ← BUILDKERNEL(pfProgram, "pfKernel2") 45:  $pfKernel3 \leftarrow BUILDKERNEL(pfProgram, "pfKernel3")$ 46: pfKernel4 ← BUILDKERNEL(pfProgram, "pfKernel4") 47: WGsize  $\leftarrow 512$  $\triangleright$  Work group size. 48: NWG  $\leftarrow$  (Np+WGsize -1)/WGsize ▷ Work groups number. 49: COPYTODEVICE( $dV, hV, True, \emptyset$ ) 50: COPYTODEVICE(dW, hW, **True**, ∅) 51: COPYTODEVICE(dSbase, hSbase, **True**, ∅) 52: COPYTODEVICE(dX2, hX, **True**,  $\emptyset$ ) 53: COPYTODEVICE( $dY2, hY, True, \emptyset$ ) 54: width  $\leftarrow$  IMG\_MAX\_WIDTH 55: height  $\leftarrow Np/(4 \times IMG_MAX_WIDTH)$ 56: for k in 1 : (Nf - 1) do 57: SETKERNELARGUMENTS(pfKernel1,Np,Ny,Nf,SSize,VSize,k,dS,dV, dSbase, dX, dY, dX2, dY2, dW, dL, dSums) ENQUEUEKERNEL(pfKernel1, (WGsize  $\times$  NWG), (WGsize),  $\emptyset$ ) 58: SETKERNELARGUMENTS(pfKernel2, Np, dSums) 59: ENQUEUEKERNEL(pfKernel2,(WGsize × NWG),(WGsize),Ø) 60: SETKERNELARGUMENTS(pfKernel3, Np, dW, dSums, dC, dU)61: ENQUEUEKERNEL(pfKernel3,(WGsize × NWG),(WGsize),Ø) 62: SETKERNELARGUMENTS(pfKernel4, Np, dU, dC, dX, dY, dX2, dY2, dW) 63. 64: ENQUEUEKERNEL(pfKernel4, (WGsize  $\times$  NWG), (WGsize),  $\emptyset$ ) 65: end for 66: COPYTOHOST(dW, hW, **True**, ∅) 67: COPYTOHOST(dX,hX, **True**, ∅)

68: CopyToHost( $dY,hY,True,\emptyset$ )





# Host data structures

- int hV[Nx \* Ny \* Nf]:  $V_{k,i,j} \rightarrow$  hV[i \* Ny \* Nf + j \* Nf + k]
- float hW[Np]:  $W_i \rightarrow hW[i]$
- float  $hX[Np]: P_i^x \to hX[i]$
- float hY[Np]:  $P_i^y \to$  hY[i]
- int hSbase[|S| \* 2]: Base values for S<sub>i</sub> generation.

## **Device data structures**

- global int dV[Nx \* Ny \* Nf]:  $V_{k,i,j} \rightarrow dV[i * Ny * Nf + j * Nf + k]$
- global float  $dW[Np]: W_i \rightarrow dW[i]$
- global float  $dX[Np]: P_i^x \to dX[i]$
- global float dY[Np]:  $P_i^y \rightarrow dY[i]$
- global float dX2[Np]:  $P_i^{\prime X} \rightarrow dX[i]$
- global float dY2[Np]:  $P'^y_i \rightarrow$  dY[i]
- global float  $dL[Np]: L_i \rightarrow dL[i]$
- global float  $dC[Np]: C_i \rightarrow dC[i]$
- global float  $dU[Np]: U_i \rightarrow dU[i]$
- global int dSbase[|S| \* 2]: Base values for  $S_i$  generation.
- global int dS[|S| \* Np]:  $S_i \rightarrow dI[|S| * i: (|S| * (i + 1) 1)]$
- global float dSums[Np]: Intermediate array used for parallel sum calculation.
- local float dlB[ $\Lambda$ ]:  $W_i \rightarrow$  dlB[i % ( $\Lambda$ )]

IMG\_MAX\_WIDTH denotes OpenCL image maximal width.

**Input datasets** The benchmark uses randomly generated video sequence, which consists of 10 frames with size  $128 \times 128$ . It uses 10000 particles to approximate the target object. The size of the video sequences and the number of particles used for approximation can be changed via command line.





## **Partitioned iteration space**

• Particles locations recalculation (pfKernel1):

$$P_P = \{(k, \gamma) : 0 \le k < N_f, 0 \le \gamma < \Gamma = N_p\}$$

$$f(k,i) \rightarrow (k,i)$$

$$f^{-1}(k,\gamma) \to (k,\gamma)$$

• Points sets generation (pfKernel1):

$$P_{S} = \{(k, \gamma, y) : 0 \le k < N_{f}, 0 \le \gamma < \Gamma = N_{p}, 0 \le y < \|S_{i}\|\}$$

$$f(k,i) \to (k,i,0:(||S_i||-1))$$

 $f^{-1}(k,\gamma,y) \to (k,\gamma)$ 

• Likelihood calculation (pfKernel1):

$$P_L = \{(k, \gamma, y) : 0 \le k < N_f, 0 \le \gamma < \Gamma = N_p, 0 \le y < \|S_i\|\}$$

$$f(k,i) \to (k,i,0:(||S_i||-1))$$
$$f^{-1}(k,\gamma,y) \to (k,\gamma)$$

• Weights calculation (pfKernel1):

$$P_W = \{ (k, \gamma) : 0 \le k < N_f, 0 \le \gamma < \Gamma = N_p \}$$

$$f^{-1}(k,i) \to (k,i)$$

$$f^{-1}(k, \gamma) \to (k, \gamma)$$

• Weights partial sum calculation (pfKernel1):

$$P_{S_w} = \{(k, v, \lambda, s) : 0 \le k < N_f, 0 \le v < N = N_p / \Lambda, 0 \le \lambda < \Lambda, s \in \{s_i = 2^i, 0 \le i < \log_2(\Lambda)\}\}$$

There is no mapping between abstract and concrete iteration spaces in that case, since the reduction operation (summation) is parallelized assuming that the sum operation is commutative and distributed *i.e.* the order of additions is completely different.





• Weights sum calculation (pfKernel2):

$$P_W = \{(k, x) : 0 \le k < N_f, 0 \le x < N_p / \Lambda\}$$
$$f^{-1}(k, i) \to (k, i / \Lambda)$$

$$f^{-1}(k,x) \to (k, (x \times \Lambda) : (x \times \Lambda + \Lambda - 1))$$

• Weights normalization (pfKernel3):

$$P_N = \{(k, \gamma) : 0 \le k < N_f, 0 \le \gamma < \Gamma = N_p\}$$
  
 $f(k, i) \to (k, i)$ 

$$f^{-1}(k,\gamma) \to (k,\gamma)$$

• Intermediate array *U* calculation (pfKernel3):

$$P_U = \{(k, \gamma) : 0 \le k < N_f, 0 \le \gamma < \Gamma = N_p\}$$

 $f(k,i) \to (k,i)$ 

$$f^{-1}(k, \gamma) \to (k, \gamma)$$

• Intermediate array *C* calculation (pfKernel3):

$$P_C = \{(k, j) : 0 \le k < N_f, 0 \le j < N_p\}$$
$$f(k, i, j) \to (k, j)$$

$$f^{-1}(k,j) \to (k,j:(N_p-1),j)$$

• Particles locations update (pfKernel4):

$$P_F = \{(k, \gamma, x) : 0 \le k < N_f, 0 \le \gamma < \Gamma = N_p, 0 \le x < N_p\}$$
  
 $f(k, i, j) \rightarrow (k, i, j)$   
 $f^{-1}(k, \gamma, j) \rightarrow (k, \gamma, j)$ 





## **Device memory access mapping**

• Particles locations recalculation (pfKernel1):

$$G_r(k, \gamma) = \{ dX[\gamma], dY[\gamma], \}$$
  
$$G_w(k, \gamma) = \{ dX[\gamma], dY[\gamma], \}$$

• Points sets generation (pfKernel1):

 $\begin{array}{lll} G_r(k,\gamma,y) &=& \{\texttt{dX}[\gamma],\texttt{dY}[\gamma],\texttt{dSbase}[2\times\texttt{y}],\texttt{dSbase}[2\times\texttt{y}+1]\}\\ G_w(k,\gamma,y) &=& \{\texttt{dS}[\gamma\times\texttt{SSize}+\texttt{y}]\} \end{array}$ 

• Likelihood calculation (pfKernel1):

$$\begin{array}{lll} G_r(k,\gamma,y) &=& \{ \texttt{dV}[\texttt{dS}[\gamma \times \texttt{SSize} + \texttt{y}]] \} \\ G_w(k,\gamma,y) &=& \{\texttt{dL}[\gamma] \} \end{array}$$

• Weights calculation (pfKernel1):

$$\begin{array}{lll} G_r(k,\gamma) &=& \left\{ \mathrm{d}\mathbb{W}[\gamma], \mathrm{d}\mathbb{L}[\gamma] \right\} \\ G_w(k,\gamma) &=& \left\{ \mathrm{d}\mathbb{W}[\gamma], \mathrm{d}\mathbb{Y}[\gamma] \right\} \end{array}$$

• Weights partial sum calculation (pfKernel1):

• Weights sum calculation (pfKernel2):

$$\begin{array}{lll} G_r(k,x) &=& \{\texttt{dSums}[\mathtt{x}]\}\\ G_w(k,x) &=& \{\texttt{dSums}[0]\} \end{array}$$

• Weights normalization (pfKernel3):

$$egin{array}{rll} G_r(k,\gamma) &=& \{ \mathtt{dW}[\gamma] \} \ G_w(k,\gamma) &=& \{ \mathtt{dW}[\gamma] \} \end{array}$$





• Intermediate array *U* calculation (pfKernel3):

$$G_w(k,\gamma) = \{ \mathtt{dU}[\gamma] \}$$

• Intermediate array *C* calculation (pfKernel3):

$$\begin{array}{lll} G_r(k,j) &=& \{ \mathtt{dW}[\mathtt{j}], \mathtt{dC}[\mathtt{j}-1] \} \\ G_w(k,j) &=& \{ \mathtt{dC}[\mathtt{j}] \} \end{array}$$

• Particles locations update (pfKernel4):

$$G_r(k, \gamma, x) = \{d[x], dU[\gamma]\}$$
  

$$G_{mr}(k, \gamma, x) = \{dX[], dY[]\}$$
  

$$G_w(k, \gamma, x) = \{dX2[\gamma], dY2[\gamma], dW[\gamma]\}$$

**Performance metrics** The performance is evaluated using the following metrics:

- The total time for initial host to device memory transfer.
- The total time for all kernel launches (computation time).
- The time for dW device to host transfer.
- The time for dX device to host transfer.
- The time for dY device to host transfer.
- The total time, required by the OpenCL implementation to perform.

Validation mechanism The benchmark has no validation mechanisms.

# Target-specific optimizations

- Avoiding memory transfers between host and device. In order to minimize number of memory transfers between host and device the parts of algorithm, which must be executed sequentially have been implemented as OpenCL kernels. Although such implementation is slower, that the host-implementation it doesn't allows keeping all data in device memory only. See [10] for more details.
- **Image memory usage.** Image memory usage for some data objects (*C* and partial sums) could increase the performance of the device code.

# Target-specific optimizations opportunities

• Avoiding unnecessary synchronizations. Not all memory barriers used in the algorithm implementation are necessary (see device code abstraction). Removing such barriers can decrease kernel execution time. Also in some cases global memory fences can be replaced by local ones.





# 2.3.14 PathFinder

The benchmark uses dynamic programming to find a path in the 2D grid from the bottom row to the top row with the minimum accumulated weight. On each step the path moves straight ahead or diagonally ahead.

## **High-level description**

**Abstract data structures** The benchmark operates on the following floating point data objects:

- W is an input  $n \times m$  grid of weights where n is the number of rows and m is the number of columns.
- *C* is an intermediate  $n \times m$  matrix of accumulated weights (costs).

**Computation** The algorithm iterates over rows, picking for each node a neighboring node in the previous row with the smallest accumulated weight.

$$C_{0,j} = W_{0,j} \quad 0 \le j \le m - 1$$

 $C_{i,j} = \min(C_{i-1,j-1}, C_{i-1,j}, C_{i-1,j+1}) + W_{i,j}, 0 < i < n, 0 < j \le m - 2$ 

$$C_{i,0} = \min(C_{i-1,0}, C_{i-1,1}) + W_{i,0}, 0 < i < n$$

$$C_{i,m-1} = \min(C_{i-1,m-2}, C_{i-1,m-1}) + W_{i,m-1}, 0 < i < n$$

$$R = \min_{0 \le j < m} (C_{n-1,j})$$

**Iteration spaces and dependences** The algorithm iterates over the 2D grid starting from the second row:

$$I = \{(i, j) : 1 \le i < n, 0 \le j < m\}$$
$$I : \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ -1 & 0 \\ 0 & -1 \end{pmatrix} \times \begin{pmatrix} i \\ j \end{pmatrix} \le \begin{pmatrix} n-1 \\ m-1 \\ -1 \\ 0 \end{pmatrix}$$

The computation generates the following dependences:

$$\delta_{I} = I: (i-1, j-1) \xrightarrow{t} I: (i, j) \land$$
$$I: (i-1, j) \xrightarrow{t} I: (i, j) \land$$
$$I: (i-1, j+1) \xrightarrow{t} I: (i, j)$$





**Memory access mapping** For  $(i, j) \in I$ :

$$\begin{split} M_r(i,j)|_{j=0} &= \{W_{i,j}, C_{i-1,j}, C_{i-1,j+1}\}\\ M_r(i,j)|_{0 < j < m-1} &= \{W_{i,j}, C_{i-1,j}, C_{i-1,j-1}, C_{i-1,j+1}\}\\ M_r(i,j)|_{j=m-1} &= \{W_{i,j}, C_{i-1,j}, C_{i-1,j-1}\}\\ M_w(i,j) &= \{C_{i,j}\} \end{split}$$

#### Low-level implementation details

**Device code** See Algorithm 69.

Host code See Algorithm 70.

#### Host data structures

- int hW[n \* m]:  $W_{i,j} \rightarrow hW[i * m + j]$
- int hR[m]:  $C_{n-1,j} \rightarrow hR[j]$

#### **Device data structures**

- global int dW[n \* m]:  $W_{i,j} \rightarrow dW[i * m + j]$
- global int dCin[m]:  $C_{i,j} \rightarrow dCin[j]$
- global int dCout[m]:  $C_{i,j} \rightarrow dCout[j]$
- local int dlP[ $\Lambda$ ]:  $C_{i,j} \rightarrow dlP[j \& (\Lambda S * 2) + S]$
- local int  $dlr[\Lambda]: C_{i,j} \rightarrow dlr[j % (\Lambda S * 2) + S]$

where s is the number of rows in a slice processed by a single kernel launch.

**Input datasets** The benchmark uses a randomly generated  $100 \times 100000$  table of integer weights.

#### **Partitioned iteration space**

$$P = \{ (r, v, \lambda, i) : 0 \le r < n, 0 \le v < N = (n \times m) / \Lambda, 0 \le \lambda < \Lambda, 0 < i < S \}$$

$$f(i,j) = ((i/S) \times S, j/(\Lambda - S \times 2), j \mod (\Lambda - S \times 2) + S, i \mod S)$$

$$f^{-1}(r, \mathbf{v}, \lambda, i) = (r + i, \mathbf{v} \times (\Lambda - S \times 2) - S + \lambda)$$



```
Algorithm 69 The Path Finder device code abstraction
Require: S \times 2 < \Lambda
Input: global int dW[]
Input: global int dCin[]
Output: global int dCout[]
Local: local int dlP[\Lambda]
Local: local int dlr[\Lambda]
                                                                            ▷ Number of rows in a slice.
Input: int S
                                                                                  ▷ Number of columns.
Input: int m
Input: int r
                                                                                               \triangleright Base row.
Input: int b
                                                                          ▷ Number of border columns.
 1: kernel PFKERNEL (S, dW, dCin, dCout, m, r, b)
         smallBlockCols \leftarrow \Lambda - (S \times 2)
                                                       ▷ Calculate the actual size of the output block.
 2:
         blkXmin \leftarrow (smallBlockCols \times \gamma) - b \triangleright Starting position of the block (including the
 3:
    ghost zone).
         blkXmax \leftarrow blkXmin + \Lambda - 1 \triangleright End position of the block (including the ghost zone).
 4:
         xidx \leftarrow blkXmin + \lambda
                                                                                   ⊳ Column coordinate.
 5:
         validXmin \leftarrow max(-blkXmin,0) \triangleright Index of the first valid element within the block
 6:
         validXmax \leftarrow min(\Lambda - 1 - (blkXmax - m + 1), \Lambda - 1)
                                                                                \triangleright Index of the last valid
 7:
     element within the block
         W \leftarrow \max(validXmin, \lambda - 1)
 8:
         E \leftarrow \min(\texttt{validXmax}, \lambda + 1)
                                                      > Diagonal elements adjusted according to the
 9:
     boundaries.
         isValid \leftarrow \lambda \in (validXmin: validXmax - 1)  > Check whether work item should
10:
     perform any work (only data within the valid range should be processed).
         if xidx \in 0: m-1 then
11:
             dlP[\lambda] \leftarrow dCin[xidx]
12:
         end if
13:
         BARRIER(CLK_LOCAL_MEM_FENCE)
14:
         for i in 0 : (S - 1) do
15:
16:
             \texttt{isComputed} \leftarrow \textbf{false}
             if \lambda \in (i+1: \Lambda - i - 2) \land is Valid then
17:
                  \texttt{isComputed} \leftarrow true
18:
                  left, up, right \leftarrow dlP[W], dlP[\lambda], dlP[E]
19:
                  shortest \leftarrow min(left, up, right)
20:
21:
                  index \leftarrow m \times (r+i) + xidx
22:
                  dlR[\lambda] \leftarrow shortest + dW[index]
             end if
23:
24:
             BARRIER(CLK_LOCAL_MEM_FENCE)
25:
             if isComputed then
26:
                  dlP[\lambda] \leftarrow dlR[\lambda]
             end if ▷ Probably a bug. Local memory barrier required at this point to ensure that
27:
     dlP has been fully updated before next iteration.
         end for
28:
         if isComputed then
29:
30:
             dCout[xidx] \leftarrow dlR[\lambda]
         end if
31:
32: end kernel
CARP-ARM-RP-001-v1.0
```

```
205
```



Algorithm 70 The Path Finder host code abstraction

```
1: procedure ALLOCATEMEMORY(n,m)
                                                          ▷ Allocate host memory and device buffers.
         hW \leftarrow ALLOCATEHOSTMEMORY(n \times m)
 2:
         hR \leftarrow ALLOCATEHOSTMEMORY(m)
 3:
         dW \leftarrow ALLOCATEBUFFER((n-1) \times m, READ_OLNY, hW + m)
 4:
         \texttt{dCin} \leftarrow \texttt{AllocateBuffer(m, \texttt{READ}\_\texttt{WRITE}, \texttt{hW})}
 5:
         dCout \leftarrow ALLOCATEBUFFER(m, READ_WRITE)
 6:
         \texttt{hC}[0] \gets \texttt{dCin}
 7:
         \mathtt{hC}[1] \gets \mathtt{dCout}
 8:
 9: end procedure
10: procedure INITIALIZEMEMORY(n,m)
                                                                               ▷ Initialize host memory.
11:
         \texttt{hW} \leftarrow \texttt{RANDOMELEMENTS}(\texttt{n} \times \texttt{m})
12: end procedure
Input: int n
                                                                      ▷ Number of rows in input array.
                                                                  ▷ Number of columns in input array.
Input: int m
Input: int b
                       ▷ Height of the computational block (see Algorithm 69 for more details).
13: size \leftarrow n \times m
14: b \leftarrow S
15: ALLOCATEMEMORY(n,m)
16: INITIALIZEMEMORY(n,m)
17: pfProgram ← BUILDPROGRAM("pfProgram.cl")
18: pfKernel ← BUILDKERNEL(pfProgram, "pfKernel")
19: \texttt{src}, \texttt{ret} \leftarrow 1, 0
20: for r in 0 : (n-1) : b do
         \texttt{src}, \texttt{ret} \leftarrow \texttt{ret}, \texttt{src}
21:
         S \leftarrow \min(b, n - r - 1)
22:
         SETKERNELARGUMENTS(pfKernel,S,dW,hC[src],hC[ret],m,r,b)
23:
         ENQUEUEKERNEL(pfKernel,(size), none, ∅)
                                                                          \triangleright Probably a bug, since only
24:
     (m/(\Lambda - 2 \times S) \times \Lambda) threads perform actual computations.
25: end for
26: COPYTOHOST(hC[ret], hR, True, ∅)
```





**Device memory access mapping** Each work-item performs the following reads and writes:

**Performance metrics** The performance is evaluated using the following metric:

• The total time for all kernel launches (computation time).

**Validation mechanism** The benchmark has no validation mechanisms.

## Target-specific optimizations

- **Blocking.** The input data is processed in blocks having the number of columns equal to the number of work items per work group. Consecutive work items of the same group access consecutive memory addresses which improves the memory subsystem performance (caches, coalescing).
- **Ghost zone [14].** The computation of the  $j^{th}$  element on the  $i^{th}$  iteration requires the  $(j-1)^{th}$  and  $(j+1)^{th}$  elements from the previous iteration. Since these data can be produced by another work group, the ghost zone technique is used to avoid data races: Each kernel launch processes an *N* rows slice of input data. Given a block with *N* rows and *M* columns the  $M N \times 2$  elements of the last row of the block can be computed by a single work group without communicating with other groups. Therefor, the whole slice is divided into non-overlapping blocks with *N* rows and  $(\Lambda 2 \times N)$  columns. Each group allocates a local buffer for  $\Lambda$  elements, which is used to compute the resulting  $(\Lambda 2 \times N)$  elements without affecting the other groups' data. This technique allows each work group to compute the result independently, although the elements located in the ghost areas are computed twice.
- **I/O buffers alternation.** The output of the kernel execution on the current iteration is used as kernel input for the next one. In order to avoid additional memory transfers, the input and output buffers are logically swapped between the kernel launches:

...
kernel(buffptr[0], buffptr[1])
kernel(buffptr[1], buffptr[0])
kernel(buffptr[0], buffptr[1])
...

The 2-element array buffptr contains the pointers to the device memory buffers.

### Target-specific optimizations opportunities





• Work items number adjusting. In the current implementation each kernel launch spawns  $n \times m$  work items, while only  $((m \times \Lambda)/(\Lambda - 2 \times S))$  perform actual computation. Using the precisely required number of work items for each kernel launch should decrease the overhead for kernel launch.





# 2.3.15 Speckle Reducing Anisotropic Diffusion

## **High-level description**

The benchmark uses Speckle Reducing Anisotropic Diffusion (SRAD) algorithm to remove from an image locally correlated noise, known as speckles, while preserving important image features. See [21] and [22] for more details.

Abstract data structures The benchmark operates on the following data objects:

- *I* is a compressed  $N_r \times N_c$  image.
- *C* is an intermediate  $N_r \times N_c$  array of floats.
- *N* is an intermediate  $N_r \times N_c$  array of floats.
- *S* is an intermediate  $N_r \times N_c$  array of floats.
- *W* is an intermediate  $N_r \times N_c$  array of floats.
- *E* is an intermediate  $N_r \times N_c$  array of floats.

# Computation

• Decompression:

$$I_{i,j} = \exp(I_{i,j}/255)$$

- Perform continuous image processing (*N<sub>i</sub>* iterations):
  - Calculate Statistics (Mean and Standard Deviation):

$$M = \left(\sum_{0 \le i < N_r, 0 \le j < N_c} I_{i,j}\right) / (N_r \times N_c)$$
$$Q = \left(\left(\sum_{0 \le i < N_r, 0 \le j < N_c} I_{i,j}^2\right) / (N_r \times N_c) - M^2\right) / M^2$$





- SRAD1. For each  $(i, j) : 0 \le i < N_r, 0 \le j < N_r$ :

- SRAD2. For each 
$$(i, j) : 0 \le i < N_r, 0 \le j < N_r$$
:

$$C_N = C_{i,j}$$

$$C_S = C_{\min(i+1,N_r-1),j}$$

$$C_W = C_{i,j}$$

$$C_E = C_{i,\min(j+1,N_r-1)}$$

$$D = (C_N \times N_{i,j} + C_S \times S_{i,j} + C_W \times W_{i,j} + C_E \times E_{i,j})/4$$

$$I_{i,j} = I_{i,j} + D \times \Lambda$$

• Compress image:

$$I_{i,j} = \log(I_{i,j}) * 255$$

 $\Lambda$  is an input parameter equal to 0.5.

# **Iteration spaces**

• Decompression:

$$I_E = \{(i, j) : 0 \le i < N_r, 0 \le j < N_c\}$$

• Statistics:

$$I_{S} = \{(k, i, j) : 0 \le k < N_{i}, 0 \le i < N_{r}, 0 \le j < N_{c}\}$$

• SRAD1:

$$I_1 = \{(k, i, j) : 0 \le k < N_i, 0 \le i < N_r, 0 \le j < N_c\}$$

• SRAD2:

$$I_2 = \{(k, i, j) : 0 \le k < N_i, 0 \le i < N_r, 0 \le j < N_c\}$$

• Compression:

$$I_C = \{(i, j) : 0 \le i < N_r, 0 \le j < N_c\}$$





# Dependences

$$\begin{split} \delta_{I_{S}} &= I_{E}: (i,j) \stackrel{t}{\to} I_{S}: (0,i,j) \land \\ I_{2}: (k-1,i,j) \stackrel{t}{\to} I_{S}: (k,i,j) \\ \delta_{I_{1}} &= I_{E}: (i,j) \stackrel{t}{\to} I_{1}: (0,i,j) \land \\ I_{E}: (\max(i-1,0),j) \stackrel{t}{\to} I_{1}: (0,i,j) \land \\ I_{E}: (\min(i+1,N_{r}-1),j) \stackrel{t}{\to} I_{1}: (0,i,j) \land \\ I_{E}: (i,\max(j-1,0)) \stackrel{t}{\to} I_{1}: (0,i,j) \land \\ I_{E}: (i,\min(N_{c}-1,j+1)) \stackrel{t}{\to} I_{1}: (0,i,j) \land \\ I_{2}: (k-1,i,j) \stackrel{t}{\to} I_{1}: (k,i,j) \land \\ I_{2}: (k-1,\min(i+1,N_{r}),j) \stackrel{t}{\to} I_{1}: (0,i,j) \land \\ I_{2}: (k-1,i,\min(i+1,N_{r}),j) \stackrel{t}{\to} I_{1}: (0,i,j) \land \\ I_{2}: (k-1,i,\min(N_{c},j+1)) \stackrel{t}{\to} I_{1}: (k,i,j) \land \\ I_{2}: (k-1,0: (N_{r}-1),0: (N_{c}-1)) \stackrel{t}{\to} I_{1}: (k,i,j) \land \\ I_{5}: (k-1,0: (N_{r}-1),0: (N_{c}-1)) \stackrel{t}{\to} I_{1}: (k,i,j) \land \\ I_{1}: (k,i,j) \stackrel{t}{\to} I_{2}: (k,i,j) \land \\ I_{1}: (k,i,j) \stackrel{t}{\to} I_{2}: (k,i,j) \land \\ I_{1}: (k,\min(i+1,N_{r}-1),j) \stackrel{t}{\to} I_{2}: (k,i,j) \land \\ I_{1}: (k,i,\min(j+1,N_{c}-1)) \stackrel{t}{\to} I_{2}: (k,i,j) \land \\ I_{1}: (k,i,\min(j+1,N_{c}-1)) \stackrel{t}{\to} I_{2}: (k,i,j) \land \\ I_{1}: (k,i,\min(j+1,N_{c}-1)) \stackrel{t}{\to} I_{2}: (k,i,j) \land \end{split}$$

# Memory access mapping

• Image Extraction:

$$M_r(i,j) = \{I_{i,j}\}$$
  
 $M_w(i,j) = \{I_{i,j}\}$ 

• Statistics Calculation:

$$M_r(k,i,j) = \{I_{i,j}\}$$

• SRAD1:

$$M_r(k,i,j) = \{I_{i,j}, I_{\max(i-1,0),j}, I_{i,\max(j-1,0)}, I_{\min(i+1,N_r-1),j}, I_{i,\min(j+1,N_r-1)}\}$$
  
$$M_w(k,i,j) = \{C_{i,j}, N_{i,j}, S_{i,j}, E_{i,j}, W_{i,j}\}$$





• SRAD2:

• Image Compression:

$$M_r(i,j) = \{I_{i,j}\}$$
  
 $M_w(i,j) = \{I_{i,j}\}$ 

### Low-level implementation details

**Device code** 

- Decompression see Algorithm 71.
- Statistics (calculated using reduction):
  - Prepare data for reduction see Algorithm 72.
  - Statistics see Algorithm 73.
- SRAD1 see Algorithm 74.
- SRAD2 see Algorithm 75.
- Compression see Algorithm 76.

Algorithm 71 The SRAD device code abstraction (Decompression).

```
Input: int Ne

Input/Output: global float dI[]

1: kernel SRADKERNEL1 (Ne,dI)

2: if \gamma < Ne then

3: dI[\gamma] \leftarrow exp(dI[\gamma]/255)

4: end if

5: end kernel
```

Host code See Algorithm 77.

### Host data structures

- float hI[Nr \* Nc]:  $I_{i,j} \rightarrow$  hI[j \* Nr + i]
- global int hIN[Nr]:hIN[i] = max(i-1,0)
- global int hIS[Nr]:hIS[i] =  $min(i+1, N_r-1)$
- global int hJW[Nc]:hJW[j] = max(j-1,0)
- global int hJE[Nc]:hJE[j] =  $\min(j+1, N_c-1)$

 $\triangleright N_e = N_r \times N_c.$ 



Algorithm 72 The SRAD device code abstraction (Prepare).  $\triangleright N_e = N_r \times N_c.$ Input: int Ne Input: global float dI[] Output: global float dSums[] Output: global float dSums2[] 1: kernel SRADKERNEL2 (Ne,dI,dSums,dSums2) 2: if  $\gamma < Ne$  then  $dSums[\gamma] \leftarrow dI[\gamma]$ 3:  $dSums2[\gamma] \leftarrow dI[\gamma]^2$ 4: 5: end if 6: end kernel

```
Algorithm 73 The SRAD device code abstraction (Statistics).
                                                                                                  \triangleright N_e = N_r \times N_c.
Input: int Ne
Input: int No
                                                                               ▷ Number of elements in array.
Input: int Step
Input: int Dim
Input/Output: global float dSums[]
Input/Output: global float dSums2[]
Local: local float dlSums[]
Local: local float dlSums2[]
  1: kernel SRADKERNEL3 (Ne, No, Step, Dim, dSums, dSums2)
 2:
          \texttt{Nf} \leftarrow \Lambda - (\texttt{Dim} \times \Lambda - \texttt{No})
                                                                     ▷ Number of elements in the last block.
          if \gamma < No then
 3:
              dlSums[\lambda] \leftarrow dSums[\gamma \times Step]
 4:
              dlSums2[\lambda] \leftarrow dSums2[\gamma \times Step]
 5:
          end if
 6:
          if Nf = \Lambda then
 7:
              i \leftarrow 2
 8:
              while i \leq \Lambda do
 9:
10:
                   if i \mid (\lambda + 1) then
                        dlSums[\lambda] \leftarrow dlSums[\lambda] + dlSums[\lambda - i/2]
11:
                        dlSums2[\lambda] \leftarrow dlSums2[\lambda] + dlSums2[\lambda - i/2]
12:
13:
                   end if
14:
                   BARRIER(CLK_LOCAL_MEM_FENCE)
15:
                   i \leftarrow 2 \times i
              end while
16:
              if \lambda = \Lambda - 1 then
17:
                   dSums[v \times Step \times \Lambda] \leftarrow dlSums[\lambda]
18:
                   dSums2[v \times Step \times \Lambda] \leftarrow dlSums2[\lambda]
19:
20:
              end if
```



```
Algorithm 73 The SRAD device code abstraction (Reduction continues).
21:
            else
22:
                 if v \neq (Dim - 1) then
                       i \leftarrow 2
23:
                       while i \leq \Lambda do
24:
                             if i \mid (\lambda + 1) then
25:
                                  dlSums[\lambda] \leftarrow dlSums[\lambda] + dlSums[\lambda - i/2]
26:
                                  dlSums2[\lambda] \leftarrow dlSums2[\lambda] + dlSums2[\lambda - i/2]
27:
                             end if
28:
29:
                             BARRIER(CLK_LOCAL_MEM_FENCE)
30:
                             i \leftarrow 2 \times i
                       end while
31:
                       if \lambda = \Lambda - 1 then
32:
                             dSums[v \times Step \times \Lambda] \leftarrow dlSums[\lambda]
33:
                             dSums2[v \times Step \times \Lambda] \leftarrow dlSums2[\lambda]
34:
35:
                       end if
                 else
36:
                       \mathtt{i} \gets 2
37:
                       while i \leq \Lambda do
38:
                            if \mathtt{Nf} \geq \mathtt{i} then
39:
40:
                                  \texttt{Df} \gets \texttt{i}
                             end if
                                                                                      \triangleright Df is the nearest power of 2 for Nf.
41:
                             \mathtt{i} \gets 2 \times \mathtt{i}
42:
43:
                       end while
                       \mathtt{i} \leftarrow 2
44:
                       while i \leq Df do
45:
                             if i \mid (\lambda + 1) \land \lambda < Df then
46:
                                  dlSums[\lambda] \leftarrow dlSums[\lambda] + dlSums[\lambda - i/2]
47:
                                  \mathtt{dlSums2}[\lambda] \gets \mathtt{dlSums2}[\lambda] + \mathtt{dlSums2}[\lambda - \mathtt{i}/2]
48:
                             end if
49:
                             \mathtt{i} \gets 2 \times \mathtt{i}
50:
                       end while
51:
                       if \lambda = \Lambda - 1 then
52:
                             for i in (\mathbf{v} \times \mathbf{A} + \mathtt{Df}) : (\mathbf{v} \times \mathbf{A} + \mathtt{Nf} - 1) do
53:
                                  dlSums[\lambda] \leftarrow dlSums[\lambda] + dSums[i]
54:
                                  dlSums2[\lambda] \leftarrow dlSums2[\lambda] + dSums2[i]
55:
                             end for
56:
                             dSums[v \times Step \times \Lambda] \leftarrow dlSums[\lambda]
57:
                             dSums2[v \times Step \times \Lambda] \leftarrow dlSums2[\lambda]
58:
59:
                       end if
                 end if
60:
            end if
61:
62: end kernel
```



```
Algorithm 74 The SRAD device code abstraction (SRAD1).
                                                                                                       \triangleright N_e = N_r \times N_c.
Input: int Ne
Input: float L
Input: int Nr
Input: float Q
Input: global int dIN[]
Input: global int dIS[]
Input: global int dJW[]
Input: global int dJE[]
Input: global float dI[]
Output: global float dN[]
Output: global float dS[]
Output: global float dE[]
Output: global float dW[]
Output: global float dC[]
  1: kernel SRADKERNEL4 (Ne,L,Nr,Q,dIN,dIS,dJW,dJE,dI,dN,dS,dW,dE,dC)
 2:
          \texttt{row} \gets \gamma \bmod \texttt{Nr}
 3:
          \texttt{col} \leftarrow \gamma/\texttt{Nr}
 4:
          if \gamma < Ne then
               dJc \leftarrow dI[\gamma]
  5:
               dlN \leftarrow dI[dIN[row] + Nr \times col] - dJc
 6:
               dlS \leftarrow dI[dIS[row] + Nr \times col] - dJc
 7:
               dlE \leftarrow dI[row + Nr \times dIW[col]] - dJc
 8:
               \texttt{dlW} \leftarrow \texttt{dI}[\texttt{row} + \texttt{Nr} \times \texttt{dIE}[\texttt{col}]] - \texttt{dJc}
 9:
               dG2 \leftarrow (dlN^2 + dlS^2 + dlE^2 + dlW^2)/dJc^2
10:
               dL \leftarrow (dlN + dlS + dlE + dlW)/dJc
11:
               dNum \leftarrow 0.5 \times dG2 - dL^2/16
12:
               \texttt{dDen} \leftarrow 1 + \texttt{dL}/4
13:
               \texttt{dQ} \gets \texttt{dNum}/\texttt{dDen}^2
14:
               \texttt{dDen} \gets (\texttt{dQ}-\texttt{Q})/(\texttt{Q}\times(\texttt{Q}+1))
15:
               dlC \leftarrow 1/(1 + dDen)
16:
               \mathtt{dlC} \leftarrow \max(\mathtt{dlC}, 0)
17:
               dlC \leftarrow min(dlC, 1)
18:
               dN[\gamma] \leftarrow dlN
19:
               dS[\gamma] \leftarrow dlS
20:
               dW[\gamma] \leftarrow dlW
21:
               dE[\gamma] \leftarrow dlE
22:
               dC[\gamma] \leftarrow dlC
23:
24:
          end if
25: end kernel
```



```
Algorithm 75 The SRAD device code abstraction (SRAD2).
Input: int Ne
                                                                                                                \triangleright N_e = N_r \times N_c.
Input: float L
Input: int Nr
Input: float Q
Input: global int dIS[]
Input: global int dJE[]
Input: global float dN[]
Input: global float dS[]
Input: global float dE[]
Input: global float dW[]
Input: global float dC[]
Input/Output: global float dI[]
  1: kernel SRADKERNEL5 (Ne,L,Nr,Q,dIS,dJE,dI,dN,dS,dW,dE,dC)
           \texttt{row} \leftarrow \gamma \mod \texttt{Nr}
 2:
           \texttt{col} \leftarrow \gamma/\texttt{Nr}
 3:
           if \gamma < \text{Ne} then
 4:
 5:
                dCN \leftarrow dC[\gamma]
                dCS \leftarrow dC[dIS[row] + Nr \times col]
 6:
 7:
                dCW \leftarrow dC[\gamma]
                \texttt{dCE} \gets \texttt{dC}[\texttt{row} + \texttt{Nr} \times \texttt{dIE}[\texttt{col}]]
 8:
                \texttt{dD} \leftarrow \texttt{dCN} \times \texttt{dN}[\gamma] + \texttt{dCS} \times \texttt{dS}[\gamma] + \texttt{dCW} \times \texttt{dW}[\gamma] + \texttt{dCE} \times \texttt{dE}[\gamma]
 9:
10:
                \mathtt{dI}[\gamma] \leftarrow \mathtt{dI}[\gamma] + \mathtt{L} \times \mathtt{dD}/4
11:
           end if
12: end kernel
```

 Algorithm 76 The SRAD device code abstraction (Compress).

 Input: int Ne
  $\triangleright N_e = N_r \times N_c$ .

 Input/Output: global float dI[]
 1: kernel SRADKERNEL6 (Ne,dI)

 2:
 if  $\gamma < Ne$  then

 3:
 dI[ $\gamma$ ]  $\leftarrow \log(dI[\gamma]) \times 255$  

 4:
 end if

 5:
 end kernel


```
Algorithm 77 The SRAD host code abstraction.
 1: procedure ALLOCATEMEMORY(Nr,Nc)
                                                          ▷ Allocate host memory and device buffers.
 2:
         \texttt{hI} \leftarrow \texttt{AllocateHostMemory}(\texttt{Nr} \times \texttt{Nc})
 3:
         hIN \leftarrow ALLOCATEHOSTMEMORY(Nr)
 4:
         hIS \leftarrow AllocateHostMemory(Nr)
 5:
         hJW \leftarrow ALLOCATEHOSTMEMORY(Nc)
         \texttt{hJE} \leftarrow \texttt{AllocateHostMemory(Nc)}
 6:
 7:
         dI \leftarrow ALLOCATEBUFFER(Nr \times Nc, READ_WRITE)
 8:
         dN \leftarrow AllocateBuffer(Nr \times Nc, READ_WRITE)
         dS \leftarrow AllocateBuffer(Nr \times Nc, READ_WRITE)
 9:
         dW \leftarrow AllocateBuffer(Nr \times Nc, READ_WRITE)
10:
         dE \leftarrow ALLOCATEBUFFER(Nr \times Nc, READ_WRITE)
11:
         dC \leftarrow AllocateBuffer(Nr \times Nc, READ_WRITE)
12:
         dSums \leftarrow AllocateBuffer(Nr \times Nc, READ_WRITE)
13:
14:
         dSums2 \leftarrow AllocateBuffer(Nr \times Nc, READ_WRITE)
         dIN \leftarrow ALLOCATEBUFFER(Nr, READ_WRITE)
15:
16:
         dIS \leftarrow ALLOCATEBUFFER(Nr, READ_WRITE)
         dJW \leftarrow ALLOCATEBUFFER(Nc, READ_WRITE)
17:
         dJE \leftarrow AllocateBuffer(Nc, READ_WRITE)
18:
19: end procedure
20: procedure INITIALIZEMEMORY(Ni,Nh)
                                                                                ▷ Initialize host memory.
21:
         hI \leftarrow READIMAGE(Ni, Nh)
         for i in 0 : (Nr - 1) do
22:
             hIN[i] \leftarrow i - 1
23:
             \mathtt{hIS[i]} \gets \mathtt{i+l}
24:
         end for
25:
         \texttt{hIN}[0] \leftarrow 0
26:
         hIS[Nr-1] \leftarrow Nr-1
27:
         for i in 0: (Nc - 1) do
28:
             \texttt{hIW}[\texttt{i}] \gets \texttt{i} - 1
29:
             hIE[i] \leftarrow i+1
30:
31:
         end for
         \texttt{hIW}[0] \leftarrow 0
32:
         \texttt{hIE}[\texttt{Nr}-1] \leftarrow \texttt{Nc}-1
33:
34: end procedure
```

# CARP



```
Algorithm 77 The SRAD host code abstraction (continues).
Input: int Nr
                                                                ▷ Number of rows to process.
Input: int Nc
                                                                 \triangleright Number of cols to process.
Input: int Ni
                                                                      ▷ Number of iterations.
Input: float L
                                                                         ▷ Lambda multiplier.
35: ALLOCATEMEMORY(Nr,Nc)
36: INITIALIZEMEMORY(Nr,Nc)
37: COPYTODEVICE(dI,hI,True,∅)
38: COPYTODEVICE(dIN, hIN, True, ∅)
39: COPYTODEVICE(dIS,hIS,True, ∅)
40: COPYTODEVICE(dJW,hJW, True, ∅)
41: COPYTODEVICE(dJE, hJE, True, ∅)
42: \texttt{Ne} \leftarrow \texttt{Nr} \times \texttt{Nc}
43: LWS \leftarrow 512
44: NWG \leftarrow Ne/LWS
45: if NWG \times LWS < Ne then
        NWG \leftarrow NWG + 1
46.
47: end if
48: GWS \leftarrow LWS \times NWG
49: sradProgram ← BUILDPROGRAM("sradProgram.cl")
50: sradKernel1 \leftarrow BUILDKERNEL(sradProgram, "sradKernel1")
51: sradKernel2 ~ BUILDKERNEL(sradProgram, "sradKernel2")
52: sradKernel3 \leftarrow BUILDKERNEL(sradProgram, "sradKernel3")
53: sradKernel4 \leftarrow BUILDKERNEL(sradProgram, "sradKernel4")
54: sradKernel5 \leftarrow BUILDKERNEL(sradProgram, "sradKernel5")
55: sradKernel6 \leftarrow BUILDKERNEL(sradProgram, "sradKernel6")
56: SETKERNELARGUMENTS(sradKernel1, Ne, dI)
57: ENQUEUEKERNEL(sradKernel1,(GWS),(LWS),∅)
58: for k in 0 : (Ni - 1) do
        SETKERNELARGUMENTS(sradKernel2, Ne, dI, dSums, dSums2)
59:
60:
        ENQUEUEKERNEL(sradKernel2, (GWS), (LWS), Ø)
        \texttt{No} \gets \texttt{Ne}
61:
62:
        \texttt{Step} \leftarrow 1
        \texttt{Dim} \gets \texttt{NWG}
63:
```

 $\textbf{64:} \qquad \texttt{GWS2} \leftarrow \texttt{GWS}$ 

# CARP



Alge	orithm 77 The SRAD host code abstraction (continues).
65:	while $\texttt{Dim} \neq 0$ do
66:	SETKERNELARGUMENTS(sradKernel3,No,Step,Dim,dSums,dSums2)
67:	$EnQUEUEKernel(\mathtt{sradKernel3},(\mathtt{GWS2}),(\mathtt{LWS}), \emptyset)$
68:	$\texttt{No} \leftarrow \texttt{Dim}$
69:	if $Dim = 1$ then
70:	$\texttt{Dim} \gets 0$
71:	else
72:	$\texttt{Step} \gets \texttt{Step} \times \texttt{LWS}$
73:	$\texttt{NWG} \gets \texttt{Dim}/\texttt{LWS}$
74:	if $NWG \times LWS < Dim$ then
75:	$\texttt{NWG} \gets \texttt{NWG} + 1$
76:	end if
77:	$\texttt{Dim} \gets \texttt{NWG}$
78:	$\texttt{GWS2} \gets \texttt{NWG} \times \texttt{LWS}$
79:	end if
80:	end while
81:	$COPYTOHOST(dSums, T, True, \emptyset)$
82:	$COPYTOHOST(dSums2, T2, True, \emptyset)$
83:	$M \leftarrow T/Ne$
84:	$V \leftarrow T2/Ne - M^2$
85:	$\mathtt{Q} \leftarrow \mathtt{V}/\mathtt{M}^2$
86:	SETKERNELARGUMENTS(sradKernel4, Ne, L, Nr, Q, dIN, dIS, dJW, dJE, dI, dN, dS, dW, dE, dC)
87:	$EnQUEUEKernel(\mathtt{sradKernel4},(\mathtt{GWS}),(\mathtt{LWS}), \emptyset)$
88:	SetKerneLArgumentS(sradKernel5, Ne, L, Nr, Q, dIS, dJE, dI, dN, dS, dW, dE, dC)
89:	$EnQUEUEKernel(\mathtt{sradKernel5},(\mathtt{GWS}),(\mathtt{LWS}), \emptyset)$
90:	end for
91:	SETKERNELARGUMENTS(sradKernel6, Ne, dI)
92:	$ENQUEUEKERNEL(sradKernel6, (GWS), (LWS), \emptyset)$
93:	COPYTOHOST(dI,hI, <b>True</b> ,∅)





### **Device data structures**

- global float dI[Nr \* Nc]:  $I_{i,j} \rightarrow$  dI[j \* Nr + i]
- global float  $dN[Nr * Nc]: N_{i,i} \rightarrow dN[j * Nr + i]$
- global float dS[Nr \* Nc]:  $S_{i,i} \rightarrow dS[j * Nr + i]$
- global float dW[Nr \* Nc]:  $W_{i,j} \rightarrow dW[j * Nr + i]$
- global float dE[Nr \* Nc]:  $E_{i,j} \rightarrow dE[j * Nr + i]$
- global float dC[Nr \* Nc]:  $C_{i,j} \rightarrow dC[j * Nr + i]$
- global int dIN[Nr]:dIN[i] = max(i-1,0)
- global int dIS[Nr]:dIS[i] =  $min(i+1, N_r-1)$
- global int dJW[Nc]:dJW[j] = max(j-1,0)
- global int dJE[Nc]:dJE[j] =  $\min(j+1, N_c-1)$
- global float dSums[Nr \* Nc]: Intermediate buffer, used for reduction operation.
- global float dSums2[Nr \* Nc]: Intermediate buffer, used for reduction operation.
- local float dlSums[A]: Intermediate buffer, used for reduction operation.
- local float dlSums2[A]: Intermediate buffer, used for reduction operation.

**Input datasets** The benchmark uses pre-generated  $502 \times 498$  input image. Command line parameters:

- The number of iterations  $(N_i)$ . The default value is 100.
- Lambda multiplier ( $\Lambda$ ). The default value is 0.5.
- The Number of rows  $(N_r)$ . The default value is 502.
- The Number of columns ( $N_c$ ). The default value is 498. The benchmark processes an  $N_r \times N_c$  area of the input image staring from the top left corner.

### **Partitioned iteration space**

• Image Extraction (sradKernel1):

$$P_E = \{(\gamma) : 0 \le \gamma < \Gamma = N_e\}$$

$$f(i,j) \to (j \times N_r + i)$$

$$f^{-1}(\gamma) \to (\gamma \mod (N_r), \gamma/N_r)$$





• Statistics Calculation (sradKernel2):

$$P_S = \{(k, \gamma) : 0 \le k < N_i, 0 \le \gamma < \Gamma = N_e\}$$

$$f(k,i,j) \rightarrow (k,j \times N_r + i)$$

$$f^{-1}(k,\gamma) \to (k,\gamma \mod (N_r),\gamma/N_r)$$

• Statistics Calculation (sradKernel3):

$$P_{R} = \{(k, s, \gamma, i) : 0 \le k < N_{i}, s \in \{s_{l} = \Lambda^{l}, 0 \le l < \log_{\Lambda} N_{e}, 0 \le \gamma < \Gamma = N_{e}/s, i = \{i_{l} = 2^{l}, 1 \le l \le \log_{2} \Lambda\}$$

There is no mapping between abstract and concrete iteration spaces in this case, since the reduction operation (summation) is parallelized assuming that the sum operation is commutative and distributive *i.e.* the order of additions is changed.

• SRAD1 (sradKernel4):

$$\begin{split} P_1 &= \{(k,\gamma): 0 \leq k < N_i, 0 \leq \gamma < \Gamma = N_e\} \\ &\qquad f(k,i,j) \to (k,j \times N_r + i) \\ &\qquad f^{-1}(k,\gamma) \to (k,\gamma \bmod (N_r),\gamma/N_r) \end{split}$$
• SRAD2 (sradKernel5):

$$P_2 = \{ (k, \gamma) : 0 \le k < N_i, 0 \le \gamma < \Gamma = N_e \}$$

 $f(k, i, j) \rightarrow (k, j \times N_r + i)$ 

$$f^{-1}(k,\gamma) \to (k,\gamma \mod (N_r),\gamma/N_r)$$

• Image Compression (sradKernel6):

$$P_C = \{(\gamma) : 0 \le \gamma < \Gamma = N_e\}$$

$$f(i,j) \to (j \times N_r + i)$$

$$f^{-1}(\gamma) \to (\gamma \mod (N_r), \gamma/N_r)$$





## **Device memory access mapping**

• Image Extraction (sradKernel1):

$$G_r(\gamma) = \{ dI[\gamma] \}$$
  
 $G_w(\gamma) = \{ dI[\gamma] \}$ 

• Statistics Calculation (sradKernel2):

 $egin{array}{rll} G_r(\gamma) &=& \{ \mathtt{dI}[\gamma] \} \ G_w(\gamma) &=& \{ \mathtt{dSums}[\gamma] \mathtt{dSums2}[\gamma] \} \end{array}$ 

• Statistics Calculation (sradKernel3):

 $\begin{array}{lll} G_r(k,s,\gamma=\nu\times\Lambda+\lambda,i) &=& \{\mathrm{dSums}[\gamma\times\mathrm{s}],\mathrm{dSums2}[\gamma\times\mathrm{s}]\}\\ G_w(k,s,\gamma=\nu\times\Lambda+\lambda,i) &=& \{\mathrm{dSums}[\Lambda\times\mathrm{s}\times\nu],\mathrm{dSums2}[\Lambda\times\mathrm{s}\times\nu]\}\\ L_r(k,s,\gamma=\nu\times\Lambda+\lambda,i) &=& \{\mathrm{dISums}[\lambda],\mathrm{dISums2}[\lambda],\mathrm{dISums}[\lambda+\mathrm{i}/2],\mathrm{dISums2}[\lambda+\mathrm{i}/2]\}\\ L_r(k,s,\gamma=\nu\times\Lambda+\lambda,i) &=& \{\mathrm{dISums}[\lambda],\mathrm{dISums2}[\lambda]\} \end{array}$ 

• SRAD1 (sradKernel4):

• SRAD2 (sradKernel5):

$$\begin{array}{lll} G_r(k,\gamma) &=& \{ \mathtt{dIS}[\gamma \bmod \mathtt{Nr}], \mathtt{dJE}[\gamma/\mathtt{Nr}], \mathtt{dI}[\gamma], \mathtt{dC}[\gamma], \mathtt{dN}[\gamma], \mathtt{dS}[\gamma], \mathtt{dW}[\gamma], \mathtt{dE}[\gamma], \\ && \mathtt{dC}[\mathtt{Nr} \times (\gamma/\mathtt{Nr}) + \min(\gamma \bmod \mathtt{Nr} + 1, \mathtt{Nr} - 1)], \\ && \mathtt{dC}[\mathtt{Nr} \times (\min(\gamma/\mathtt{Nr}) - 1, \mathtt{Nc} - 1) + \gamma \bmod \mathtt{Nr}], \} \\ G_w(k,\gamma) &=& \{ \mathtt{dI}[\gamma] \} \end{array}$$

• Image Compression (sradKernel6):

$$egin{array}{rll} G_r(\gamma) &=& \{ \mathtt{dI}[\gamma] \} \ G_w(\gamma) &=& \{ \mathtt{dI}[\gamma] \} \end{array}$$





**Performance metrics** The performance of the algorithm implementation is evaluated using the following metrics:

- The time to read an image.
- Host arrays (N, S, W, E) initialization time.
- Total time required by OpenCL implementation to perform compilation:
  - Context and Command Queue creation.
  - Device memory allocation.
  - Host to device memory transfer.
  - Kernel launches.
  - Device to host memory transfer.

**Validation mechanism** The benchmark has no built-in validation mechanisms. The benchmark can write output image to the file, which can be validated by external tools (not supplied).

## **Target-specific optimizations**

• **Tree reduction.** The OpenCL code implements parallel reduction (sum) using a treebased approach, which requires additional memory barriers.

## **Target-specific optimizations opportunities**

• Halo usage. Current algorithm implementation uses 4 additional arrays (N, S, W, E) to avoid having boundary checks. While this approach eliminates such checks, it also increases the memory consumption and makes the global memory access non regular (from the OpenCL compiler point of view). Another approach to avoiding boundary checks is to use one element width halo for each buffer. See Figure 2.3 for more details. OpenCL image data type can also be used.

# CARP





Figure 2.3: Halo for SRAD.



# 2.3.16 Stream Cluster

# **High-level description**

The benchmark finds a predetermined number of medians in the input stream of points, so that each point is assigned to its nearest median. The quality of the clustering is measured as a sum of square distances (SSQ metric).

Abstract data structures The benchmark operates on the following objects:

- *S* is an input  $N_e$  elements stream (vector) of Points divided into  $N_c$  elements chunks. Number of chunks is  $(N_s = N_e/N_c)$ .
- *P* is a point, characterized by the following features:
  - $P^{w}$  denotes the weight of the point *P*.
  - $P^c$  denotes the  $N_f$  vector of the coordinates of the point P.
- $C^i$  is an intermediate set of clusters centers  $(K_{\min} \le ||C^i|| \le K_{\max})$  for the  $i^{th}$  chunk.
- *C* is an output set of clusters centers  $(K_{\min} \le ||C|| \le K_{\max})$  for the whole stream *P*.

**Computation** The algorithm reads point by chunks from the input stream and updates the set of clusters to minimize the SSQ: On the first step the intermediate  $C_i$  is calculated for the each chunk  $C_i$ . On the second step C is generated as a set of clusters centers for the  $\{C_i\}$  set.

- Generate  $C_i$  for each input chunk  $S_i$ :
  - Randomly select initial centers.
  - Try to improve selected centers by opening new centers (opening a new center can cause some opened centers to be closed). See below for more details.
- Generate C from  $C_i$ : The same algorithm is applied to the set of centers obtained on the step one.

In order to determine which which centers should be opened the cost of opening a new center should be calculated. Only this part of the algorithm is implemented as OpenCL kernel (the rest of the document provides the description of this part only):

$$Q_i = \left(\sum_{0 \le j < N_c} \|P_i - P_j\|\right) \times P_i^w, 0 \le i < N_c$$

If the new cost  $Q_i$  is less than original cost of the point  $P_i$ , than it should be marked as a new center.

## **Iteration space**

$$I = \{(i, j) : 0 \le i < N_c, 0 \le j < N_c\}$$

**Dependences** The dependencies are introduced by the other parts of the algorithm (not included in this description).





## Memory access mapping

$$M_r(i,j) = \{P_i,P_j\}$$
$$M_w(i,j) = \{Q_i\}$$

## Low-level implementation details

The implementation is based on the PARSEC Benchmark Suite implementation [3].

**Device code** See Algorithm 78.

## **Device data structures**

- global struct Point dP[Nc]:  $S_i \rightarrow dP[i \ \% \ Nc]$
- global float dPC[Nc \* Nf]:  $S_i^c[j] \rightarrow dPC[j * Nc + i % Nc]$
- global float dW[Kmax \* Nc]: Intermediate buffer using for reduction.
- global int dC[Nc]:  $C_i \rightarrow dC[i]$
- global bool dM[Nc]:  $S_i \rightarrow dC[i \ \& Nc]$  If dM[i] is true then  $S_i$  is marked as center.
- local float dlPC[Nf]: $S_i^c[j] \rightarrow dPC[j]$

**Input datasets** The benchmark can either use randomly generated stream of points, or read the stream from file. The following parameters can be specified via command line:

- Minimal number of centers allowed  $(K_{\min})$ . The default value is 10.
- Maximal number of centers allowed  $(K_{\text{max}})$ . The default value is 20.
- Dimension of each data point  $(N_f)$ . The default value is 256.
- Number of data points in the input stream  $(N_e)$ . The default value is 65536.
- Number of data points in each chunk  $(N_c)$ . The default value is 65536.
- Maximal number of intermediate centers. The default value is 1000.

Be default the randomly generated stream of points is used.



```
Algorithm 78 Stream Cluster device code abstraction.
Input: global struct Point dP[]
Input: global float dPC[]
Input/Output: global float dW[]
Input: global int dC[]
Output: global bool dM[]
Local: local float dlPC[]
Input: int Nc
Input: int Nf
Input: int X
                                                                                                    ⊳ Current center.
Input: int K
                                                                                              ▷ Number of centers.
  1: kernel SCKERNEL (dP, dPC, dW, dC, dM, Nc, Nf, X, K)
 2:
          if \gamma < Nc then
               if \lambda = 0 then
 3:
                   for i in 0: (Nf - 1) do
 4:
 5:
                        dlPC[i] \leftarrow dPC[i \times Nc + X]
                   end for
 6:
 7:
               end if
               BARRIER(CLK_LOCAL_MEM_FENCE)
 8:
 9:
               \texttt{cost} \gets 0
               for i in 0 : (Nf - 1) do
10:
                    \texttt{cost} \leftarrow \texttt{cost} + (\texttt{dPC}[\texttt{i} \times \texttt{Nc} + \gamma] - \texttt{dlPC}[\texttt{i}])^2
11:
12:
               end for
               \texttt{cost} \leftarrow \texttt{cost} \times \texttt{dP}[\gamma].\texttt{weight}
13:
               \texttt{cost0} \gets \texttt{dP}[\gamma].\texttt{cost}
14:
               \texttt{base} \leftarrow \gamma \times (\texttt{K} + 1)
15:
               if cost < cost0 then
16:
                    dM[\gamma] \leftarrow true
17:
                    \texttt{dW}[\texttt{base} + \texttt{K}] \gets \texttt{cost} - \texttt{cost0}
18:
19:
               else
                    addr \leftarrow base + dC[dP[\gamma].assign]
20:
                    dW[addr] \leftarrow dW[addr] + cost0 - cost
21:
               end if
22:
          end if
23:
24: end kernel
```





# **Partitioned iteration space**

$$P = \{(x, \gamma) : 0 \le x < N_c, 0 \le \gamma < N_c\}$$
$$f(i, j) \to (i, j)$$
$$f^{-1}(x, \gamma) \to (x, \gamma)$$

# Device memory access mapping

**Performance metrics** The performance of the algorithm implementation is evaluated using the following metrics:

- Total time for OpenCL initialization.
- Total time for Host to Device memory transfers.
- Total time for Device to Host memory transfers.
- Total time for Device memory allocation.
- Total time for Device memory deallocation.
- Total time for all scKernel launches.

**Validation mechanism** The benchmark has no built-in validation mechanisms. The benchmark can print the computed centers, which can be validated by external tools (not supplied).

**Target-specific optimizations** The following target-specific optimization are used in the algorithm implementation:

• Coalescing-friendly memory access (probably CUDA only). In order to achieve coalesced memory access the input data buffer is logically transposed. For CUDA devices this transformation makes the data access in kernel coalesced. In other hand this greatly reduces the data locality. Although for CUDA devices the advantages of the coalesced memory access overrides the disadvantages of the non-localized memory access, this might not be the same for OpenCL. Since OpenCL standard doesn't require the consecutive memory accesses within threads of the same working group to be coalesced, this transformation might cause a performance degradation.





## **Target-specific optimizations opportunities**

• Vector operations usage. Distance calculation in scKernel can be vectorized, which would improve the performance on vector-based architectures.





# 2.4 Parboil

The Parboil benchmark suite,<sup>14</sup> developed at the University of Illinois at Urbana-Champaign, was the first CUDA benchmarking suite. It was updated [18] to include OpenCL versions of the benchmarks. Unlike Rodinia and SHOC, Parboil includes both base (naïve) implementations and implementations optimized for NVIDIA hardware.

# 2.4.1 Breadth-First Search

The BFS implementation in Parboil is similar to the one in Rodinia (see §2.3.2). The difference is that the Parboil implementation assigns costs to each edge, which requires locks to ensure the new frontier is correctly computed. The locks are implemented using the following OpenCL built-in functions: atom\_min, atom\_xchg and atom\_add.

# 2.4.2 SGEMM

The base implementation of SGEMM in Parboil uses Algorithm 79. This unoptimized kernel is equivalent to the optimized sgemmNT kernel in SHOC (§2.2.7).

Algorithm 79 Parboil SGEMM device code abstraction.

```
Input: global const float dA[]
Input: global const float dB[]
Input: int Na
Input: int Nb
Input: int Nc
Input: int k
Input: float alpha
Input: float beta
Input/Output: global float dC[]
  1: kernel SGEMMNT (dA, dB, Na, Nb, Nc, k, alpha, beta, dC)
           \texttt{sum} \leftarrow 0
  2:
           for i in 0:(\mathtt{k}-1) do
  3:
                 \texttt{sum} \leftarrow \texttt{sum} + \texttt{dA}[\gamma_0 + \texttt{i} \times \texttt{Na}] \times \texttt{dB}[\gamma_1 + \texttt{i} \times \texttt{Nb}]
  4:
  5:
           end for
           \texttt{dC}[\gamma_0 + \gamma_1 \times \texttt{Nc}] \leftarrow \texttt{dC}[\gamma_0 + \gamma_1 \times \texttt{Nc}] \times \texttt{beta} + \texttt{alpha} \times \texttt{sum}
  6:
  7: end kernel
```

# 2.4.3 Sparse Matrix-Vector Multiplication

The implementation of SpMV in Parboil uses the Ellpack format, while the implementations of SpMV in SHOC use the Compressed Sparse Row (CSR) and Ellpack-R formats (§2.2.8).

# 2.4.4 Stencil

The base implementation of a 7-point 3D stencil operation in Parboil uses Algorithm 80. The Stencil implementation in SHOC is that of a 9-point 2D stencil operation (§2.2.2).

<sup>14</sup>http://impact.crhc.illinois.edu/parboil.aspx



```
Algorithm 80 The Parboil base Stencil device code abstraction.
 1: procedure INDEX3D(Nx,Ny,i,j,k)
         return i + Nx \times (j + Ny \times k)
 2:
 3: end procedure
Input: global const dAin[]
Input: int Nx, Ny, Nz
Input: float c0, c1
Output: global float dAout[]
 4: kernel STENCILKERNEL1 (dAin, Nx, Ny, Nz, c0, c1, dAout)
 5:
         i \leftarrow \gamma_0 + 1
 6:
         j \leftarrow \gamma_1 + 1
 7:
        \mathbf{k} \leftarrow \gamma_2 + 1
         if i < Nx - 1 then
 8:
             dAout[INDEX3D(Nx, Ny, i, j, k)] \leftarrow
 9:
10:
                c1 \times \sum_{i'=i\pm 1} \sum_{j'=j\pm 1} \sum_{k'=k\pm 1} dA[INDEX3D(Nx, Ny, i', j', k')] +
                c0 \times dA[INDEX3D(Nx,Ny,i,j,k)]
11:
12:
         end if
13: end kernel
```

# 2.4.5 Other benchmarks

The other Parboil benchmarks [18] are:

- Cutoff-limited Coulombic Potential
- MRI Cartesian Gridding
- Histogram
- Lattice-Boltzman Method simulation
- MRI Non-Cartesian Q Matrix Calculation
- Sum of Absolute Differences
- Two-Point Angular Correlation Function





Figure 2.4: The conceptual schema of the image representation of the image in the memory. Image *S* is a sub-image of image *I*.

# 2.5 RealEyes

### 2.5.1 Arithmetic Operations

In image processing, per pixel arithmetic operations are an important common use case. The operation might even seem simplistic, but it is important from the compiler point-of-view for two reasons. First, arithmetic operations take lots of computing time, so it is desirable to complete them in parallel. Second, because they can be accelerated considerably using local memory. In this section we introduce some standard concepts of image storage as well as the importance of the local accelerator memory even when performing seemingly trivial tasks.

Initially, when an image I is allocated, it is stored in memory as a continuous array. In many cases, however, we want to work only with sub-pictures, which are usually rectangular or cubical subsets of the original image. In this case, to avoid the need for reallocating and copying the image into a new array, we create the sub-picture by defining a set of properties. On Fig. 2.4 we can see an image representation in the memory. The properties describing the image S are: the number of rows (rows(S)), the number of columns (cols(S)), the distance between two consecutive lines (step(S)) and the offset (offset(S)). These values are derived from the original image and the desired properties of the sub-image.

Let's take an example. Consider an image *I* of size  $\{rows(I) \times cols(I)\}$ . If we want to find a pixel (x, y) in *I* we can do so as:

$$I_{x,y} = I[\operatorname{rows}(I) \times x + y]$$
(2.19)

Now we want to consider an image  $S \subset I$  of size  $\{rows(S) \times cols(S)\}$ , where  $S_{0,0} = I_{x_o,y_o}$ . If





we want to find pixel (q, w) in *S*, we can do that as:

$$S_{q,w} = I[\texttt{offset}(S) + \texttt{step}(S) \times q + w]$$
(2.20)

where  $offset(S) = rows(I) \times x_0 + y_0$ . To simplify the formalism, we consider  $S_{0,0} = I[offset(S)]$  and

$$S_{q,w} = S[\operatorname{step}(S) \times q + w] \tag{2.21}$$

We apply the image processing operators on images like *S*. In this section we consider arithmetic and trigonometric operators. These are used to transform the images into different fields. Simple operators are, for example, sin, cos, exp, or any closed combination of thereof. For the sake of simplicity, we only consider the exp operator; the other combinations follow similarly.

## Notation

Through this section we are using the following notations:

- $\lambda_1$  the current row in the 2D iteration as returned by the OpenCL framework
- $\lambda_0$  the current column in the 2D iteration as returned by the OpenCL framework
- $\tau_{\lambda_1}$  the current row of the thread within the group of threads
- $\tau_{\lambda_0}$  the current column of the thread within the group of the threads
- $v_1$  the current row of the group within the groups of the kernel
- $v_0$  the current column of the group within the groups of the kernel
- rows(.) the number of rows of the image
- cols(.) the number of columns of the image
- $N_1$  the number group rows (the number of groups is  $N_1 \times N_0$ )
- $N_0$  the number of group columns

### Benchmark

The general equation can be stated as:

$$\underline{\mathbf{R}} = \exp \underline{\mathbf{A}} \tag{2.22}$$

### **High-level Description**

The algorithm, operates on the following data:

- $\underline{\mathbf{A}}$  is the input matrix of size  $m \times n$ .
- **<u>R</u>** is the result matrix of size  $m \times n$ .

**Computation** We implement two versions of the algorithm. In the naïve implementation we implement the arithmetic operation directly, whereas in the optimized calculation we use the local memory.



#### Algorithm 81 Arithmetic Operator

1: kernel EXP (image, result)

**Input:**  $\lambda_1, \lambda_0, \tau_{\lambda_1}, \tau_{\lambda_0}, v_1, v_0, N_1, N_0$ 

- 2:  $q_{\text{image}} \leftarrow \text{step}(\text{image}) \cdot \lambda_1 + \lambda_0 \triangleright$  We compute the image specific index for the current thread.
- 3:  $q_{\texttt{result}} \leftarrow \texttt{step}(\texttt{result}) \cdot \lambda_1 + \lambda_0 \qquad \triangleright$  We compute the result specific index for the current thread.
- 4:  $\operatorname{result}[q_{\operatorname{result}}] \leftarrow \exp(\operatorname{image}[q_{\operatorname{image}}])$
- 5: end kernel

### Algorithm 82 Optimized Arithmetic Operator

1: **kernel** EXP (image, result)

**Input:**  $\lambda_1, \lambda_0, \tau_{\lambda_1}, \tau_{\lambda_0}, v_1, v_0, \text{rows}, \text{cols}, N_1, N_0$ **Local:**  $\text{buff}[N_0 \times N_1]$ 

- 2:  $q_{\text{image}} \leftarrow \text{step}(\text{image}) \cdot \lambda_1 + \lambda_0 \triangleright$  We compute the image specific index for the current thread.
- 3:  $q_{\text{buff}} \leftarrow N_0 \cdot \tau_{\lambda_1} + \tau_{\lambda_0}$   $\triangleright$  We compute the buffer specific index for the current thread.
- 4:  $\operatorname{buff}[q_{\operatorname{buff}}] \leftarrow \operatorname{image}[q_{\operatorname{image}}]$
- 5: BARRIER(CLK\_LOCAL\_MEM\_FENCE) ▷ At this point we have filled up the local memory with the input image data.
- 6:  $q_{\text{result}} \leftarrow \text{step}(\text{result}) \cdot \lambda_1 + \lambda_0 \qquad \triangleright$  We calculate the result specific index for the current thread.
- 7:  $\operatorname{result}[q_{\operatorname{result}}] \leftarrow \exp(\operatorname{buff}[q_{\operatorname{buff}}])$
- 8: end kernel



**Naïve Implementation** The naive implementation can be observed in algorithm 81. In the optimized implementation we make sure, that all the reads and writes are coalesced. The optimized implementation can be observed in algorithm 82.

**Dependences** The computation is embarrassingly parallel: there are no dependences beyond the availability of the input.

**Complexity** The arithmetic operations have  $O(rows(I) \times cols(I))$  complexity.

**Memory Access Mapping and Complexity** This algorithm is a 2D iteration. The iteration variables are  $\lambda_1$  and  $\lambda_0$ . In the light of the above, the memory access mapping can be written as:

$$G_r(\lambda_1, \lambda_0) = \operatorname{image}(\lambda_1, \lambda_0)$$
 (2.23)

$$G_w(\lambda_1, \lambda_0) = \texttt{result}(\lambda_1, \lambda_0) \tag{2.24}$$

**Computational Complexity** The computational complexity is linear with the number of pixels  $O(\Lambda_1 \times \Lambda_0)$ 

**Optimized Implementation** In the optimized implementation we are using the local cache to read and write the memory in a coalesced way. This approach also works, and is beneficial also if the size of the types differ, that is to say if we do type-casting during the operation.

# 2.5.2 GEMM – General Matrix Multiply

The benchmark uses a simple general matrix multiplication algorithm to test the effects of the local memory on the computation performance. The general equation can be stated as:

$$\underline{\underline{\mathbf{R}}} = \alpha \underline{\underline{\mathbf{A}}} \cdot \underline{\underline{\mathbf{B}}} + \beta \underline{\underline{\mathbf{C}}}$$
(2.25)

or

$$\underline{\underline{\mathbf{C}}} = \alpha \underline{\underline{\mathbf{A}}} \cdot \underline{\underline{\mathbf{B}}} + \beta \underline{\underline{\mathbf{C}}}$$
(2.26)

### **High-level Description**

The benchmark operates on he following floating point data objects:

- <u>A</u> is the first multiplier matrix of size  $m \times n$
- **<u>B</u>** is the second multiplier matrix of size  $n \times o$
- $\underline{\mathbf{C}}$  is the addend matrix of size  $m \times o$
- **<u>R</u>** is the result matrix of size  $m \times o$

Depending on the application the multiplication can either be implemented in-place of the  $\underline{\underline{C}}$  matrix or using the functional programming paradigm creating a result  $\underline{\underline{R}}$  matrix.



**Computation** The algorithm uses the definition of the matrix multiplication<sup>15</sup>:

$$\underline{\mathbf{R}}_{i,j} = \beta \underline{\mathbf{C}}_{i,j} + \sum_{q=0}^{n-1} \alpha \underline{\mathbf{A}}_{i,q} \underline{\mathbf{B}}_{q,o}$$
(2.27)

We compare two implementations. The first is the naïve implementation, whereas the second is optimized for the local memory usage and coalesced reading.

## **Naïve Implementation**

The naïve implementation can be found in algorithm 83.

Algo	orithm 83 GEMM Naïve Implementation		
Inpu	lt: global float A[]		
Inpu	l <b>t: global float</b> B[]		
Inpu	l <b>f: global float</b> C[]		
Out	<pre>put: global float R[]</pre>		
CLV	$^{\prime}$ ar: $\lambda_{1}$ , $\lambda_{0}$	⊳ current rov	w: $\lambda_1$ , current column: $\lambda_0$
1: <b>I</b>	kernel GEMM_NAIVE (A, B, C, R)		
2:	$\texttt{sub} \gets 0$		
3:	for $q$ in $0\ldots \mathtt{cols}(\mathtt{A})-1$ do		
4:	$\mathtt{sub} \gets \mathtt{sub} + \mathtt{A}[\lambda_1  imes \mathtt{cols}(\mathtt{A}) + q] \cdot \mathtt{B}[q  imes \mathtt{co}]$	$ls(B) + \lambda_0]$	$\triangleright$ we calculate the
1	multiplicand		
5:	end for		
6:	$R[\lambda_1  imes \mathtt{cols}(R) + \lambda_0] \leftarrow \mathtt{sub} + \mathtt{C}[\lambda_1  imes \mathtt{cols}(\mathtt{C})]$	$+\lambda_0]$	▷ we calculate the linear
(	combination		
7:	end kernel		

**Computational Dependencies** The computation does not generate any dependencies beyond the dependencies on the availability of the input.

**Memory Access Mapping** It is important to note that most of the memory accesses of the naïve kernel are non-coalesced, which means that the computation time cannot be estimated with the global memory bandwidth. (There is only one iteration.)

$$G_r(\lambda_1,\lambda_0) = A(\lambda_1,\cdot) \cup B(\cdot,\lambda_0) \cup C(\lambda_1,\lambda_0) G_w(\lambda_1,\lambda_0) = R(\lambda_1,\lambda_0)$$
(2.28)

**Complexity** The row of matrix *A* and the column of matrix *B* is accessed in the calculation of each final point. This results in a memory access complexity of:

uncoalesced reads : $m \times n \times o$	(2.29)
---	--------

- coalesced reads :  $m \times n \times o$  (2.30)
- coalesced writes :  $m \times o$  (2.31)
  - local writes : 1 (2.32)
  - local reads : 1 (2.33)

 $<sup>^{15}</sup>$ we use the informatics notation in the algorithms, where the indices start from 0





 $O(m \times n \times o)$ . This complexity is to be understood for uncoalesced global memory access count.

**Computation Complexity** The basic algorithm has an  $O(m \times n \times o)$  complexity.

## **Optimized Implementation**

The optimized implementation is using a local "cache"-like buffer which is much faster than the global memory. On top of this the global memory accesses are solely coalesced. The memory is accessed linearly from the input images.

The algorithm creates a window of the size of the work group of threads. The windows are read into the local memory where the window related multiplications are performed. After the work on the window is finished, the window is moved along the multiplication axis.

The optimized implementation can be observed in algorithm 84.

**Computational Dependencies** In the optimized computation we have introduced two new local memories: buffA and buffB. The buffers hold the patch of the input images on which the current calculation is performed.

$$\delta_{I_R} = R \xrightarrow{t} (\operatorname{sub} \cup C) \bigwedge$$
(2.34)

$$\operatorname{sub} \xrightarrow{t} \operatorname{buffA} \cup \operatorname{buffB}$$
 (2.35)

**Memory Access Mapping and Complexity** The complete data is accessed during the iteration.

$$G_r = \{A \cup B \cup C\} \tag{2.36}$$

$$G_w = \{R\}\tag{2.37}$$

Matrices *A* and *B* are accessed strictly once in coalesced order. The local matrices are accessed by each result pixel.

coalesced reads : 
$$m \times n + n \times o$$
 (2.38)

coalesced writes : 
$$m \times o$$
 (2.39)

local writes :  $m \times n + n \times o$  (2.40)

local reads : 
$$m \times n \times o$$
 (2.41)

**Computational Complexity** The basic algorithm has an  $O(m \times n \times o)$  complexity.

### **Optimization Opportunities**

**Target-specific Optimizations** The size of the work-group can be tuned for the target device.

**Optimal Alignment** If the data is aligned to a special value (row and cols are divisible by 4 or 64), we can set up kernel without boundary checking. This can accelerate the kernels. Also if the image cannot be embedded into an aligned memory space we can still cut up the image into 4 sub-images. One image with sizes divisible by 64 and three non-symmetric matrices. In this case two differently optimized kernels can be called on the sub images accelerating the computation time.



```
Algorithm 84 GEMM Optimized Implementation
Input: global float A[]
Input: global float B[]
Input: global float C[]
Output: global float R[]
                                                    \triangleright current row: \lambda_1, current column: \lambda_0, thread row: \tau_{\lambda_1}, thread
CLVar: \lambda_1, \lambda_0, \tau_{\lambda_1}, \tau_{\lambda_0}, v_1, v_0
       column: \tau_{\lambda_0}, the row of the group: v_1, the column of the group: v_0
Local: local buffA[16(16+1)]
                                                                         ▷ We allocate +1 to avoid memory banks conflicts
Local: local buffB[16(16+1)]
                                                                         \triangleright We allocate +1 to avoid memory banks conflicts
  1: kernel GEMM_OPTIMIZED (A, B, C, R)
             d \leftarrow 16
  2:
                                                       \triangleright block size; We are working with 16 \times 16 threads per group.
                                                                                  \triangleright the first element of the A matrix to process
  3:
             A_{\texttt{begin}} \leftarrow d \cdot v_1 \cdot \texttt{cols}(A)
  4:
             B_{\texttt{begin}} \leftarrow d \cdot v_0
                                                                                  ▷ the first element of the B matrix to process
  5:
             \texttt{step}(\texttt{A}) \leftarrow d
             step(B) \leftarrow d \cdot cols(B)
  6:
             N_{\texttt{blocks}} \leftarrow (\texttt{cols}(\texttt{A}) + d - 1) \mod d
                                                                                           ▷ the number of total blocks to process
  7:
             sub \leftarrow 0
                                                                                                                                          \triangleright \texttt{sub} \leftarrow 0
  8:
  9:
             \mathtt{A}_{\texttt{offset}} \leftarrow \tau_{\lambda_1} + \tau_{\lambda_0} \cdot \texttt{cols}(\mathtt{A})
 10:
             B_{\texttt{offset}} \leftarrow \tau_{\lambda_1} + \tau_{\lambda_0} \cdot \texttt{cols}(B)
             for block in 0 \dots N_{\text{blocks}} - 1 do
 11:
                   if (block \cdot d + \tau_{\lambda_1} < cols(A) \land v_1 \cdot d + \tau_{\lambda_0} < rows(A)) then
 12:
                          \texttt{buffA}[\tau_{\lambda_0} \cdot (d+1) + \tau_{\lambda_1}] \leftarrow \texttt{A}[\texttt{A}_{\texttt{begin}} + \texttt{A}_{\texttt{offset}}]
 13:
                   else
 14:
                         \texttt{buffA}[\tau_{\lambda_0} \cdot (d+1) + \tau_{\lambda_1}] \leftarrow 0
 15:
                   end if
 16:
                   if (block \cdot d + \tau_{\lambda_0} < rows(B) \land v_0 \cdot d + \tau_{\lambda_1} < cols(B)) then
 17:
                          \texttt{buffB}[\tau_{\lambda_1} \cdot (d+1) + \tau_{\lambda_0}] \leftarrow \texttt{B}[\texttt{B}_{\texttt{begin}} + \texttt{B}_{\texttt{offset}}]
 18:
                   else
 19:
                         \texttt{buffB}[\tau_{\lambda_1} \cdot (d+1) + \tau_{\lambda_0}] \leftarrow 0
20:
                   end if
21 \cdot
                   BARRIER(CLK_LOCAL_MEM_FENCE)
22:
23:
                   for i in 0...20 - 1 do
                          \mathtt{sub} \leftarrow \mathtt{sub} + \mathtt{buffA}[\tau_{\lambda_1} \cdot d + i] \cdot \mathtt{buffB}[\tau_{\lambda_0} \cdot d]
24:
25:
                   end for
                   BARRIER(CLK_LOCAL_MEM_FENCE)
26:
27:
                   A_{\texttt{begin}} \leftarrow A_{\texttt{begin}} + \texttt{step}(A)
                   B_{\text{begin}} \leftarrow B_{\text{begin}} + \text{step}(B)
28:
29:
             end for
             if \lambda_1 < \texttt{rows}(A) \bigwedge \lambda_0 < \texttt{cols}(B) then
30:
                   R[\lambda_1 \times \texttt{cols}(B) + \lambda_0] \leftarrow \texttt{sub} + C[\lambda_1 \times \texttt{cols}(C) + \lambda_0]
31:
             end if
32:
33: end kernel
```





Figure 2.5: The conceptual schema of the image repacking. Images  $A_1 \dots A_4$  are repacked as columns of image *I*.

**SIMD Enforcement** On architectures which do support SIMD instructions, the data can be collected into a local aligned transposed buffer. The multiplication can then be performed using the SIMD instructions. Some architectures might support SIMD operations.

# 2.5.3 Image Statistics

This benchmark extracts statistical properties of an image: maximum value, minimum value and the sum of the image's pixels. To simplify calculation, we work on several images at once and the images are re-packed into a big image as the big image's columns. The concept can be seen on fig. 2.5.

# **Input Data Structures**

- $I_{0...n-1}$  is the set of input images (all images have the same size)
- sample is an image containing the repacked *I* images
- M is an intermediate array containing intermediate results

We repack the images  $I_q$  into a single image sample the following way:

$$\mathtt{sample}(q,w) = I_w(q \div \mathtt{cols}(I_w), q \mod \mathtt{cols}(I_w)) \tag{2.42}$$

The M structure is an array of records, each containing a maximum value, a minimum value and a sum value. These are noted by  $\max(\mathtt{sample}[i])$ ,  $\min(\mathtt{sample}[i])$  and  $\sup(\mathtt{sample}[i])$  respectively.

In the end of the calculation M will contain the following values:

$$\forall i \in \{0...\,\texttt{cols}(\texttt{sample}) - 1\} \tag{2.43}$$

$$\max(\mathbb{M}[\mathtt{i}]) = \max\{\mathtt{sample}(*,i)\} = \max\{I_i\}$$
(2.44)

$$\min(\mathbb{M}[\mathbf{i}]) = \min\{\mathtt{sample}(*, i)\} = \min\{I_i\}$$
(2.45)

$$\operatorname{sum}(\mathtt{M}[\mathtt{i}]) = \sum_{p \in \{\mathtt{sample}(*,i)\}} p = \sum_{p \in I_i} p$$
(2.46)





In the remainder of this discussion, for brevity, we consider the maximum value only, the rest of the values can be treated similar way.

It is difficult to extract the maximum on an parallel accelerator efficiently. In the naïve algorithm we consider each pixel p and compare it with a candidate value c. If p > c then we replace the candidate value with p. In a parallel accelerator all the comparisons would be carried out once, so in the end we would choose an undefined value. On the other hand, if we would make the comparison atomically, then we would end up with a sequential algorithm. The solution is to work by tiles, and collect the statistical properties of the patches in a temporary array of records.

**Computation** The computation is done in two phases. In the first phase we divide the image into tiles and we extract the maximum value of each group into a contracted array (M). In the second phase we iterate on the array M, contracting it further in each iteration until in the end we are left with a single line containing the maximum value of each image. For the simplicity of the implementation we choose the patches of size  $(k \times cols(sample))$ . This way we only contract the rows and not the columns. The algorithm can be observed in algorithm 85.

Algorithm 85 Image Statistics Conceptual Schema

Input: sample **Input:**  $\Lambda_1$ ▷ block size; number of rows in the group Input:  $\Lambda_0$ ▷ block size; number of columns in the group 1: CREATEARRAY(M, rows(sample)  $\div \Lambda_1$ , cols(sample)) 2: for q in 0... rows $(sample) \div \Lambda_1 - 1$  do  $\triangleright$  We initialize the M array. for w in  $0 \dots \Lambda_0 - 1$  do 3: 4:  $\max(\mathbb{M}(q,w)) \leftarrow \max\{ \mathtt{sample}(q\Lambda_1 \dots q(\Lambda_1+1), w) \} \triangleright \mathsf{We} \text{ contract the image by}$ extracting the maximum value from each tile end for 5: 6: **end for** 7:  $r \leftarrow rows(sample) \div \Lambda_1$ 8:  $s \leftarrow 1$ 9: while r > 1 do  $\triangleright$  We keep exponentially contracting the image until we get a single line containing the maxes for all the images in the sample 10: for q in  $0 \dots \operatorname{rows}(M) \div (\Lambda_1 s) - 1$  do for w in  $0 \dots \Lambda_0 - 1$  do 11:  $\max(\mathsf{M}(q \ast s, w)) \leftarrow \max\left\{\bigcup_{i \in 0...\Lambda_{1}-1} \mathsf{M}((q+i)s, w)\right\}$ 12: end for 13: end for 14: 15:  $r \leftarrow r \div \Lambda_1$ 16:  $s \leftarrow s\Lambda_1$ 17: end while

Now we are going to discuss the dependencies in the memory accesses. In the above schema the maximum values are collected in the first line of the patch. This way the patches can be further contracted independently. In each iteration *i* we access the contracted image or array of



CARP

records, so the memory access evolves as follows:

$$M_r(0) = \texttt{sample} \tag{2.47}$$

$$M_{r}(i > 0) = \bigcup_{q \in \frac{\operatorname{rows}(M)}{\Lambda_{1}^{i}}} \operatorname{M}(q\Lambda_{1}^{i}, *)$$
(2.48)

$$M_w(i) = \bigcup_{q \in \frac{\operatorname{rows}(M)}{\Lambda_1^{i+1}}}^{1} \mathbb{M}(q\Lambda_1^{i+1}, *)$$
(2.49)

If we consider the memory access complexity, we have to note that all the global reads and writes are coalesced. The memory complexity is as follows:

coalesced reads:  $\frac{\text{cols(sample)} \cdot \text{rows(sample)}}{\Lambda_1^i}$  (2.50)

coalesced writes : 
$$\frac{\text{cols(sample)} \cdot \text{rows(sample)}}{\Lambda_1^{i+1}}$$
(2.51)

local reads: 
$$\frac{\text{cols}(\text{sample}) \cdot \text{rows}(\text{sample})}{\Lambda_1^i}$$
(2.52)

local writes : 
$$\frac{\text{cols(sample)} \cdot \text{rows(sample)}}{\Lambda_1^{i+1}}$$
(2.53)

#### **Low-level Implementation Details**

The first and the second phase of the algorithm are implemented as two separate kernels. The kernels are then called in a wrapper function. The implementation can be seen in algorithm 86.

As presented above, after making the copy to the local memory of the group of threads, only one thread per column continues the computation. As we work simultaneously *n* images, this results working with one thread per  $(d_{\tau}\Lambda_0) \div n$ . The optimal number of images<sup>16</sup> is  $(d_{\tau}\Lambda_0) \div 2$ .

**Computational Dependencies** In the kernel every information us cached locally before it is processed. This way global memory access is always coalesced.

$$\delta_{I_R} = M \xrightarrow{t} \text{buff} \bigwedge$$
(2.54)

$$buff \xrightarrow{l} sample$$
 (2.55)

**Low-level Memory Mapping** From the above implementation, we can see that the calculation does not depend on the host's memory once the image is on the device. The device memory access evolves as follows:

$$G_r(0) = \texttt{sample} \tag{2.56}$$

$$G_r(i > 0) = \bigcup_{q \in \frac{\operatorname{rows}(M)}{\Lambda_1^i}} M(q\Lambda_1^i, *)$$
(2.57)

$$G_w(i) = \bigcup_{q \in \frac{\operatorname{rows}(M)}{\Lambda_1^{i+1}}}^{1} \mathbb{M}(q\Lambda_1^{i+1}, *)$$
(2.58)

<sup>&</sup>lt;sup>16</sup>from the point of view of thread occupancy



```
Algorithm 86 Image Statistics Implementation Details
  1: procedure STATISTICS(sample, \Lambda_1, \Lambda_0, M)
             MAXESPATCH( sample, M, rows(sample) \div \Lambda_1, cols(sample) )
  2:
  3:
             r \leftarrow \texttt{rows}(\texttt{sample}) \div \Lambda_1
             s \leftarrow 1
  4:
             while r > 1 do
  5:
  6:
                   MAXESITERATEPATCH(M, r, cols(sample), s)
  7:
                   r \leftarrow r \div \Lambda_1
  8:
                   s \leftarrow s \Lambda_1
             end while
  9:
 10: end procedure
 11: kernel MAXESPATCH (sample, M, M<sub>row</sub>, M<sub>col</sub>)
Input: \lambda_1, \lambda_0, \tau_{\lambda_1}, \tau_{\lambda_0}, v_1, v_0, v_1, \Lambda_0
 12:
             buff[1024]
             T \leftarrow \tau_{\lambda_1} \times \Lambda_0 + \tau_{\lambda_0}
 13:
             \texttt{buff}[T] \leftarrow \texttt{sample}[\lambda_1 \times \texttt{cols}(\texttt{sample}) + \lambda_0]
14:
 15:
             BARRIER(CLK_LOCAL_MEM_FENCE)
             if \tau_{\lambda_1} == 0 then
 16:
                   v \leftarrow \texttt{buff}[\tau_{\lambda_0}]
 17:
                   for \mu_{row} in 1 \dots \Lambda_1 do
 18:
                         I \leftarrow \mu_{\texttt{row}} \times \Lambda_0 + \tau_{\lambda_0}
 19:
20:
                         if v > buff[I] then
                                v \leftarrow \texttt{buff}[I]
21:
22:
                          end if
23:
                   end for
                   \mathtt{M}[v_1 \times \mathtt{cols}(\mathtt{sample}) + \lambda_0] \gets v
24:
             end if
25:
26: end kernel
27: kernel MAXESITERATEPATCH (M, M_{row}, M_{col}, s)
Input: \lambda_1, \lambda_0, \tau_{\lambda_1}, \tau_{\lambda_0}, v_1, v_0, v_1, \Lambda_0
28:
             buff[1024]
             T \leftarrow \tau_{\lambda_1} \times \Lambda_0 + \tau_{\lambda_0}
29:
             \texttt{buff}[T] \leftarrow \texttt{M}[s \cdot \lambda_1 \times \texttt{cols}(\texttt{sample}) + \lambda_0]
30:
             BARRIER(CLK_LOCAL_MEM_FENCE)
31:
             if \tau_{\lambda_1} == 0 then
32:
33:
                   v \leftarrow \texttt{buff}[\tau_{\lambda_0}]
                   for \mu_{row} in 1 \dots \Lambda_1 do
34:
35:
                         I \leftarrow \mu_{\texttt{row}} \times \Lambda_0 + \tau_{\lambda_0}
                         if v > buff[I] then
36:
37:
                               v \leftarrow \texttt{buff}[I]
38:
                         end if
                   end for
39:
40:
                   \mathbb{M}[sv_1 \times \texttt{cols}(\texttt{sample}) + \lambda_0] \leftarrow v
             end if
41:
42: end kernel
```





## **Target-specific Optimization**

The statistical extraction relies on the exact size and grouping if the threads during the computation. If it's possible during the repacking, the sample size can be adjusted to be optimal compared to the specific hardware.

# 2.5.4 Image Normalization

This benchmark performs image normalization by scaling and shifting. To optimize the GPU occupancy, the normalization is performed parallel on several images of the same size.

## **Input Data Structures**

- $I_{0...n-1}$  is the set of input images (all have the same size)
- sample is an image containing the repacked *I* images
- *M* is an array of records containing the scaling and the shifting constant

We repack the images  $I_q$  into a single image sample the following way:

$$sample(q, w) = I_w(q \div cols(I_w), q \mod cols(I_w))$$
(2.59)

High-level Description The normalization can be described with the following formula:

$$\mathtt{sample}(q, w) = \mathtt{scale}(M[w]) \cdot \mathtt{sample}(q, w) + \mathtt{shift}(M[w]) \tag{2.60}$$

**High-level Memory Access Pattern** The memory access can be expressed the following way:

$$M_r = \text{sample}$$
 (2.61)

$$M_w = \text{sample}$$
 (2.62)

Low-level Implementation Details The algorithm is implemented in the following kernel 87.

Algorithm 87 Image Normalization Implementation Details1: kernel MAXESPATCH (sample, M)Input:  $\rho, \kappa, \tau_{\rho}, \tau_{\kappa}, g_{\rho}, g_{\kappa}, d_{\rho}, d_{\kappa}$ 2: shift[1024]3: scale[1024]4: if  $\tau_{\rho} == 0$  then5: shift[ $\tau_{\kappa}$ ]  $\leftarrow$  shift(M[ $\kappa$ ])6: scale[ $\tau_{\kappa}$ ]  $\leftarrow$  scale(M[ $\kappa$ ])

9: 
$$I \leftarrow \rho \times \operatorname{cols}(M) + \kappa$$

10:  $\operatorname{sample}[I] \leftarrow \operatorname{scale}[\tau_{\kappa}] \cdot \operatorname{sample}[I] + \operatorname{shift}[\tau_{\kappa}]$ 





**Computational Dependencies** The input data is read into a local buffer. As the buffers are reused this optimizes the global memory reads. The full dependency looks as follows:

$$\delta_{I_R} = \operatorname{sample}(\cdot, \lambda_0) \xrightarrow{t} \operatorname{shift}[\lambda_0] \bigwedge$$
(2.63)

$$\operatorname{sample}(\cdot,\lambda_0) \xrightarrow{t} \operatorname{scale}[\lambda_0] \bigwedge$$
(2.64)

$$\operatorname{scale}[\lambda_0] \xrightarrow{t} M[\lambda_0] \bigwedge$$
(2.65)

$$\operatorname{shift}[\lambda_0] \xrightarrow{t} M[\lambda_0]$$
 (2.66)

**Computational Complexity** The basic algorithm has linear  $O(rows(sample) \cdot cols(sample))$  complexity.

## 2.5.5 Image (Data) Repacking

In machine learning algorithms we often need to extract higher level information from our data. For the algorithms to work on multiple types of data, the input data has to be standardized. The *de facto* standard is to organize the input data into matrices. In the special case of face tracking, one sample consists of the collection of patches around a candidate point. We, therefore, need to repack the patches from around a candidate point into a single matrix line-by-line.

This benchmark operator performs a matrix reorganization based on predefined permutation rules.

#### **High-level Description**

The benchmark operates on an input image I. On this image we define patches P as

$$P_{x,y,s} = I\left(\left[x - \frac{s}{2}, x + \frac{s}{2}\right], \left[y - \frac{s}{2}, y + \frac{s}{2}\right]\right)$$
(2.67)

where x, y is the center of the patch and s is the size of the patch. We create a collection of tiles for every given patch, that fits into the image. These patches are then packed line-by-line into the result matrix. That is, for every patch  $P_i$ , we have

$$R_{i,q} = P_i[q] = (P_i)_{(q \div s, q \mod s)}$$
(2.68)

where R is the repacked image.

**Computation** The algorithm we implement in this benchmark uses the reduction for the repacking of the data. In the first part of the algorithm, we read the content of the patch into a local buffer and in the second part we feed the content of the buffer into the result matrix. As the patches and the matrix lines are non-continuous, we over-fill the buffer so we could fetch all the pixel which are necessary to finish the line of the matrix. The schema of the computation can be observed in fig. 2.5.5.

Implementation Pseudo-code is shown in algorithm 88.

**Computational Dependences** The computation does not generate any dependencies beyond the availability of the input.



```
Algorithm 88 Filling the Sample Matrix with Patches
  1: kernel FILLPATCH (image, patchSize, result)
Input: \lambda_1, \lambda_0, \tau_{\lambda_1}, \tau_{\lambda_0}, v_1, v_0, \Lambda_1, \Lambda_0
Local: buff[1024]
  2:
            row_{start}(result) \leftarrow v_1 \cdot \Lambda_1
            col_{start}(result) \leftarrow v_0 \cdot \Lambda_0
  3:
  4:
            row_{start}(image) \leftarrow row_{start}(result) mod patchSize
  5:
            col_{start}(image) \leftarrow col_{start}(result) + row_{start}(result) \div patchSize
            row_{end}(image) \leftarrow (row_{start}(result) + \Lambda_1 - 1) \div patchSize
  6:
             col_{end}(image) \leftarrow (col_{start}(result) + \Lambda_0 - 1) + (row_{start}(result) + \Lambda_1 - 1)
  7:
       mod patchSize
  8:
             \tau \leftarrow \tau_{\lambda_1} \cdot \Lambda_0 + \tau_{\lambda_0}
  9:
            row_{current}(image) \leftarrow row_{start}(image) + (col_{start}(image) + \tau) \div cols(image)
            col_{current}(image) \leftarrow (col_{start}(image) + \tau) \div cols(image)
 10:
 11:
             \tau_{\texttt{current}}(\texttt{image}) \leftarrow \texttt{row}_{\texttt{current}}(\texttt{image}) \cdot \texttt{cols}(\texttt{image}) + \texttt{col}_{\texttt{current}}(\texttt{image})
             \texttt{step} \gets \textbf{false}
 12:
            if row_{current}(image) = row_{end}(image) \wedge col_{current}(image) > col_{end}(image) then
 13:
                  \texttt{step} \leftarrow \textbf{true}
 14 \cdot
            end if
 15:
            if row_{current}(image) > row_{end}(image) then
 16:
 17:
                  \texttt{step} \leftarrow \textbf{true}
            end if
 18:
 19:
            if !step then
20:
                  \texttt{buff}[\tau] \leftarrow \texttt{image}[\tau_{\texttt{current}}(\texttt{image})]
            end if
21:
22:
             BARRIER(CLK_LOCAL_MEM_FENCE)
23:
            row_{current}(result) \leftarrow row_{start}(result) + \tau_{\lambda_1}
            col_{current}(result) \leftarrow col_{start}(result) + \tau_{\lambda_0}
24:
25:
             \tau_{\texttt{current}}(\texttt{result}) \leftarrow \texttt{row}_{\texttt{current}}(\texttt{result}) \cdot \texttt{cols}(\texttt{result}) + \texttt{col}_{\texttt{current}}(\texttt{result})
            \lambda_1(\text{image}) \leftarrow \text{row}_{\text{current}}(\text{result}) \div \text{patchSize}
26:
             \lambda_0(\texttt{image}) \gets \texttt{col}_{\texttt{current}}(\texttt{result}) + \texttt{row}_{\texttt{current}}(\texttt{result}) \hspace{0.2cm} \texttt{mod} \hspace{0.2cm} \texttt{patchSize}
27:
28:
             \tau_{\texttt{buff}} = (\texttt{row}_{\texttt{image}}) - \texttt{row}_{\texttt{start}}(\texttt{image}) \cdot \texttt{cols} + (\lambda_0(\texttt{image}) - \texttt{col}_{\texttt{start}}(\texttt{image}))
            if row<sub>current</sub>(image) > rows(result) then
29:
                  return
30:
            end if
31:
32:
             result[\tau_{current}(result)] \leftarrow buff[\tau_{buff}]
33: end kernel
```





Figure 2.6: The schema of the reorganization of the patch image for the MLP training. All the input patches will be repacked into a single R matrix.

**Memory Access Mapping and Complexity** All of the global memory accesses are coalesced and we assure reading the input and the output only once. The access map is as follows:

$$G_r = \{I\}\tag{2.69}$$

$$G_w = \{R\} \tag{2.70}$$

**Computational Complexity** The basic algorithm has  $O(rows(I) \times cols(I) \times patchSize)$  complexity.

## **Optimization Opportunities**

The key of the optimization is the use of local memory for non-coalesced reads and writes. Therefore it is crucial to exploit this local buffering programming paradigm in the implementation.

# 2.6 BasemarkCL

## 2.6.1 Histogram equalization

### Description

This benchmark performs a histogram equalization to an image. All color channels are equalized independently.

Threads in a local work-group all read a single color value of the input image and increment the corresponding value in the local histogram array using atomic operations. After this the local histogram array is summed into the global histogram array by again avoiding memory race condition by using atomic operations. Cumulative sum is applied to the global histogram. On the final kernel threads read a single pixel value from input image, use the cumulative histogram value on that pixel value to calculate an output color and writes it into output image.





## **Data structures**

- *W* is the width of the image.
- *H* is the height of the image.
- *HR* histogram array for red channel
- *HG* histogram array for green channel
- *HB* histogram array for blue channel
- HA histogram array for alpha channel
- img is the 2 dimensional image.

## Iteration space and dependencies

$$I = \{(i, j) : 0 \le i < W, 0 \le j < H\}$$

## Memory access mapping

A pixel is read from the image and the corresponding indices in the histogram arrays are incremented. Histogram calculation: for  $(i, j) \in I_D$ :

### **Device specific data structures**

- histogram is an array of 1024 unsigned integers
- local\_histogram is an array of 1024 unsigned integers in local memory
- img is the 2d RGBA image with 8bits per color channel

## **Device code abstraction**

```
CARP-ARM-RP-001-v1.0
```

# CARP



```
Input/Output: global uint histogram[]
```

```
Local: uint csum
```

```
kernel GENERATECUMULATIVEDISTRIBUTION (histogram)

csum \leftarrow 0

for x in 0: 255 do

csum \leftarrow csum + histogram[x + \gamma_0 \times 255]

histogram[x + \gamma_0 \times 256] \leftarrow csum

end for
```

```
end kernel
```

```
Input/Output: global uint histogram[]
kernel ZERODISTRIBUTION (histogram)
for x in 0 : M do
histogram[x + \gamma_0 \times 256] \leftarrow 0
end for
end kernel
```

# Host code abstraction

```
procedure HISTOGRAM(inputImage,outputImage,histogram)
    SETKERNELARGUMENTS(zeroDistribution,histogram)
    ENQUEUEKERNEL(zeroDistribution,4)
```





```
SETKERNELARGUMENTS(generateHistogram, inputImage, histogram)
ENQUEUEKERNEL(generateHistogram, W, H)
SETKERNELARGUMENTS(generateCumulativeDistribution, histogram)
ENQUEUEKERNEL(generateCumulativeDistribution, 4)
SETKERNELARGUMENTS(applyHistogram, inputImage, outputImage, histogram)
ENQUEUEKERNEL(applyHistogram, W, H)
end procedure
```

## **Input datasets**

This benchmark operates on a set of 5 2048x2048 input images with 4 8-bit color channels (RGBA).

## **Partitioned iteration space**

 $P = \{(\gamma_0, \gamma_1, \lambda_0, \lambda_1, i) : 0 \le \gamma_0 < \Gamma_0 = 2048, 0 \le \gamma_1 < \Gamma_1 = 2048, 0 \le i < B = 256\}$ 

The device is allowed to choose it's local workgroup size therefore the local id:s can run between any limit that divides the global size exactly.

## Device memory access mapping

• Histogram calculation

$$egin{array}{rll} G_r(\gamma_0,\gamma_1)&=&\{ extsf{img}_{\gamma_0,\gamma_1}\}\ L_{rw}(\gamma_0,\gamma_1)&=&\{ extsf{limg}_{\gamma_0,\gamma_1}]\}\ G_{rw}(i)&=&=\{ extsf{histogram}[i]\} \end{array}$$

• Apply histogram:

• Zero distribution:  $0 \le \gamma 0 < 4$ 

$$G_w(\gamma_0, i) = \{\texttt{histogram}[i + \gamma_0 imes B]\}$$

• Generate cumulative histogram:  $0 \le \gamma 0 < 4$ 

$$G_{rw}(\gamma_0, i) = \{\texttt{histogram}[i + \gamma_0 imes B]\}$$

## **Performance metrics**

The performance is evaluated by taking the geometric mean of inverse frame times thus giving the geometric average of frames per second.





# Validation mechanism

The benchmark has no built-in validation mechanism. The validation is visual only.

# 2.6.2 SPH fluid

## **High-level description**

The benchmark performs a smoothed particle hydrodynamics fluid simulation [1].

# Data structures

The benchmark operates on particles in 3D space. Density of the fluid on a location of a particle depends on the amount particles within a distance h from the particle under consideration. Viscous and pressure dependent forces acting upon a particle are calculated by examining the neighboring particles within distance h. The space is divided into a regular grid of voxels where particles reside in order to speed up the neighboring particle search.

The benchmark operates on following data structures and has following constants:

- M is the amount of particles.
- V is the amount of voxels in the three dimensional space.
- keys is an array containing M tuples of values containing the index of a voxel and the index of a particle residing in said voxel.
- voxelParticles is an array of V scalar values containing the smallest particle index of all particles in the voxel.
- sortedVelocity is an array of M vectors containing the velocity of a particle in sorted order.
- sortedDensity is an array of M scalar values containing the density at the particles position in sorted order.
- position is an array of M vectors containing the position of each particle.
- velocity is an array of M vectors containing the velocity of each particle.
- acceleration is an array of M vectors containing the acceleration applied for each particle.

# Computation

Voxel index for each particle is calculated for each particle into the array keys using position such that  $keys[\gamma_0] = (VoxelIndexAt(position[\gamma_0]), \gamma_0)$ . The keys array is then sorted so that the smallest voxel index resides at the start of the array. After sorting the arrays sortedPosition and sortedVelocity are updated so that sortedPosition[\gamma\_0] = position[keys[\gamma\_0].y]. Using the sorted keys array the voxelParticles is updated so that voxelParticles\_i contains the index of the first particle residing in said voxel so that the sortedVelocity and sortedPosition can be indexed using that information.





Density and acceleration of each particle is then calculated by spawning a thread for each particle, looking up the voxel index of said particle using sortedPosition array and then looping trough neighboring voxels and the voxel the particle currently resides in. Neighboring particles in each voxel are found by looking the index of the first particle residing in said voxel by using the voxelParticles array and then increasing the index until keys<sub> $\gamma_0$ </sub> is different from the previous iteration, after which a second voxel is examined. Arrays sortedPosition, sortedDensity and sortedVelocity are accessed using the index. At the end acceleration for a particle is written into acceleration array by looking up the required index from the keys array.

As a last step a simple Eulerian integration is used to update the data in position and velocity arrays by using the information in acceleration array. If particle crosses the boundary of the world it bounces inelastically and the position gets clamped to the edge.

### **Iteration space**

#### Memory access mapping

• updateKeys: for  $(i) \in I_D$ :

$$egin{array}{rll} M_r(i) &=& \{ extsf{position}_i\}\ M_w(i) &=& \{ extsf{keys}_i\} \end{array}$$

• sortPostProcess: for  $(i) \in I_D$ :

 $egin{array}{rll} M_r(i) &=& \{ {\tt position}_i \} \ M_w(i) &=& \{ {\tt sortedPosition}_i, {\tt sortedVelocity}_i \} \end{array}$ 

### **Device code abstraction**

```
Input: int4 position
```

```
Input: uint4 voxelCount
```

- 1: **function** INDEXFROMPOSITION(position, voxelCount) ▷ Returns the voxel on which a particle resides based on the current physical location
- 2: return voxelCount.x × voxelCount.y × position.z × voxelCount.x × position.y × position.x
- 3: end function

```
Input: float4 position
Input: uint4 voxelCount
Input: float h
 4: function GETPOSCOORD (position, voxelCount)
           \texttt{minx} \leftarrow \texttt{FLOOR}(\frac{\texttt{position.x}}{\texttt{L}})
 5:
          \texttt{miny} \leftarrow \texttt{FLOOR}(\underbrace{\overset{n}{\overset{position.y}{\flat}}}_{\texttt{b}})
 6:
          \min z \leftarrow FLOOR(\frac{\text{position.}z}{r})
  7:
 8:
           posCoord.x \leftarrow CLAMP(minx, 0, voxelCount.x - 1)
           posCoord.y \leftarrow CLAMP(miny, 0, voxelCount.y - 1)
 9:
10:
           posCoord.z \leftarrow CLAMP(minz, 0, voxelCount.z - 1)
11:
           \texttt{posCoord.w} \leftarrow 0
```





```
return posCoord
12:
13: end function
Input: global uint2 keys[]
Input: global float4 sortedPosition[]
Input: global int voxelParticles[]
Input: uint4 voxelCount
Input: input float h
Input: float m
Output: float sortedDensity[]
Local: totalDensity
Local: ownVoxel
Local: ownPosition
14: kernel CALCULATEDENSITY (keys, sortedPosition, voxelParticles, voxelCount, h, m)
15:
        ownVoxel \leftarrow GETPOSCOORD(sortedPosition[\gamma_0], voxelCount, h)
        ownVoxel \leftarrow CLAMP(ownVoxel, (1, 1, 1, 1), (voxelCount.x - 2, voxelCount.y - 2, voxelCount.z - 2)
16:
    (2,0))
        ownPosition \leftarrow sortedPosition[\gamma_0]
17:
        totalDensity \leftarrow 0
18:
        for z in -1:1 do
19:
           for y in -1:1 do
20:
               for x in -1:1 do
21:
                   pid \leftarrow voxelParticles[INDEXFROMPOSITION(ownVoxel + (x, y, z, 0), voxelCount)]
22:
                   while pid < \Gamma_0 \land \text{keys}[\text{pid}].x = \text{index } \mathbf{do}
23:
                       \texttt{distance} \gets \texttt{ownPosition} - \texttt{sortedPosition}[\texttt{pid}]
24:
                       if distance < h then
25:
                          totalDensity 
    totalDensity + W(distance, h)
26:
27:
                      end if
28:
                      pid \leftarrow pid + 1
                   end while
29:
               end for
30.
           end for
31:
32:
        end for
        \texttt{sortedDensity}[\gamma_0] \leftarrow \texttt{totalDensity} \times \texttt{m}
33:
34: end kernel
Input: global uint2 keys[]
Input: global float4 sortedPosition[]
Input: global int voxelParticles[]
Input: global float sortedDensity[]
Input: global float4 sortedVelocity[]
Input: uint4 voxelCount
Input: float h
Input: float m
Input: float u
Output: global float acceleration[]
Local: float4 pressureAcceleration
```
## CARP



```
Local: float4 viscousAcceleration
Local: int4 ownVoxel
Local: float4 ownPosition
 35: kernel CALCULATEACCELERATION (keys, sortedPosition, voxelParticles, sortedDenisty, sortedVeloc
                       ownVoxel \leftarrow GETPOSCOORD(sortedPosition[\gamma_0], voxelCount, h)
 36:
                       ownVoxel \leftarrow CLAMP(ownVoxel, (1, 1, 1, 1), (voxelCount.x - 2, voxelCount.y - 2, voxelCount.z - 2)
 37:
             (2,0))
                       ownPosition \leftarrow sortedPosition[\gamma_0]
 38:
 39:
                       pressureAcceleration \leftarrow 0
 40:
                       viscousAcceleration \leftarrow 0
                       for z in -1:1 do
 41:
 42:
                                  for y in -1:1 do
                                             for x in -1:1 do
 43:
                                                         pid \leftarrow voxelParticles[INDEXFROMPOSITION(ownVoxel + (x, y, z, 0), voxelCount)]
 44:
                                                         while pid < \Gamma_0 \land \text{keys}[\text{pid}].x = \text{index } \mathbf{do}
 45:
                                                                    distance \leftarrow ownPosition - sortedPosition[pid]
 46:
                                                                    if distance \leq h then
                                                                                                                                                                          ▷ W2 and W3 refer simply to purely
 47:
             mathematical manipulations on the values supplied there.
                                                                              pressureAcceleration \leftarrow pressureAcceleration + W2(sortedDensity[pid], distance dist
 48:
                                                                              viscousAcceleration \leftarrow viscousAcceleration + W3(sortedVelocity[pid], distance viscousAcceleration + V3(sortedVelocity[pid], distance viscousAcceleration + V3(sortedVelocity[pi
 49:
                                                                    end if
 50:
 51:
                                                                   \texttt{pid} \gets \texttt{pid} + 1
 52:
                                                         end while
                                             end for
 53:
                                  end for
 54:
                       end for
 55:
                       acceleration[keys[\gamma_0].y] \leftarrow pressureAcceleration+viscousAcceleration+
 56:
             gravity
 57: end kernel
Input/Output: global float4 position[]
Input/Output: global float4 velocity[]
Input: global float4 acceleration[]
Input: uint4 voxelCount
Input: float h
Input: float dt
 58: kernel INTEGRATE (position, velocity, acceleration, voxelCount, h, dt)
 59:
                       \texttt{newVelocity} \leftarrow \texttt{velocity}[\gamma_0] + \texttt{acceleration}[\gamma_0] \times \texttt{dt}
                       position[\gamma_0] \leftarrow position[\gamma_0] + newVelocity \times dt
 60:
                       if position[\gamma_0].x < 0 \lor position[\gamma_0].x > (voxelCount.x - 1) \times h then
 61:
                                  position [\gamma_0].x \leftarrow CLAMP(position [\gamma_0].x, 0, (voxelCount.x - 1) \times h)
 62:
                                  \texttt{newVelocity.x} \leftarrow -\texttt{newVelocity.x}
 63:
                       end if
 64:
                       if position[\gamma_0].y < 0 \lor position[\gamma_0].y > (voxelCount.y - 1) \times h then
 65:
                                  position[\gamma_0].y \leftarrow CLAMP(position[\gamma_0].y, 0, (voxelCount.y - 1) \times h)
 66:
 67:
                                  newVelocity.y \leftarrow -newVelocity.y
 68:
                       end if
```

CARP-ARM-RP-001-v1.0

## CARP



```
if position[\gamma_0] . z < 0 \lor position[\gamma_0] . z > (voxelCount. z - 1) \times h then
69:
70:
            position[\gamma_0].z \leftarrow CLAMP(position[\gamma_0].z, 0, (voxelCount.z - 1) \times h)
            newVelocity.z \leftarrow -newVelocity.z
71:
72:
        end if
        velocity[\gamma_0] \leftarrow newVelocity
73:
74: end kernel
Input: global float4 position[]
Input: uint4 voxelCount
Input: float h
Output: global uint2 keys[]
75: kernel UPDATEKEYS (position, voxelCount, h, keys)
        posCoord \leftarrow GETPOSCOORD(position[\gamma_0], voxelCount, h)
76:
77:
        \text{keys}[\gamma_0] \leftarrow (\text{voxelCount.x} \times \text{voxelCount.y} \times \text{posCoord.z} + \text{voxelCount.x} \times \text{posCoord.y} +
    posCoord.x, \gamma_0
78: end kernel
Input/Output: global uint2 keys[]
Input: uint stage
Input: uint pass
79: kernel BITONICSORT (keys, stage, pass)
        \texttt{pairDistance} \gets 2^{(\texttt{stage-pass})}
80:
        leftID \leftarrow (\gamma_0 \& (pairDistance - 1)) \mid ((\gamma_0 \& \sim (pairDistance - 1)) \times 2)
81:
        direction \leftarrow ((\gamma_0 \gg \text{stage}) \& 1) == 1?0:1
82:
        rightID \leftarrow leftID + pairDistance
83:
        left \leftarrow keys[leftID]
84:
        right \leftarrow keys[rightID]
85:
        larger \leftarrow left.x > right.x?left:right
86:
87:
        smaller \leftarrow left.x > right.x?right: left
        keys[leftID] ← direction?smaller: larger
88:
        keys[rightID] \leftarrow direction?larger: smaller
89:
90: end kernel
Input: global uint2 keys[]
Input: global float4 position[]
Input: global float4 velocity[]
Output: global float4 sortedPosition[]
Output: global float4 sortedVelocity[]
91: kernel SORTPOSTPROCESS (keys, position, velocity, sortedPosition, sortedVelocity)
        sortedPosition[\gamma_0] \leftarrow position[keys[\gamma_0].y]
92:
93:
        sortedVelocity[\gamma_0] \leftarrow velocity[keys[\gamma_0].y]
94: end kernel
Output: global int voxelParticles[]
95: kernel CLEARVOXELS (voxelParticles)
96:
        voxelParticles[\gamma_0] \leftarrow -1
97: end kernel
Input: global uint2 keys[]
```

CARP-ARM-RP-001-v1.0

### CARP



```
Output: global int voxelParticles[]
98: kernel SETVOXELS (keys, voxelParticles)
        if \gamma_0 = 0 then
99.
             voxelParticles[keys[0].x] \leftarrow 0
100:
         return
         end if
101:
         if keys[\gamma_0 - 1].x \neq keys[\gamma_0].x then
102:
103:
             voxelParticles[keys[\gamma_0].x] \leftarrow \gamma_0
104:
         end if
105: end kernel
```

### Host code abstraction

```
procedure INITIALIZE(keys, voxelParticles, sortedVelocity, sortedPosition,
position, velocity, acceleration )
                                      \triangleright Run only once at start of the program
  SETKERNELARGUMENTS(updateKeys, keys, position, voxelCount, h)
  SETKERNELARGUMENTS(sortPostProcess, keys, position, velocity, sortedPosition,
sortedVelocity)
  SETKERNELARGUMENTS(clearVoxels, voxelParticles)
  SETKERNELARGUMENTS(setVoxels, voxelParticles, keys)
  SETKERNELARGUMENTS(calculateDensity, keys, sortedPosition, voxelParticles,
sortedDensity, voxelCount, h, m)
  SETKERNELARGUMENTS(calculateAcceleration, keys, sortedPosiition,
voxelParticles, sortedDensity, sortedVelocity, acceleration, voxelCount,
h, m, u)
  SETKERNELARGUMENTS(integrate, position, velocity, acceleration,
voxelCount, h, dt)
  SETKERNELARGUMENTS(clearVelocityAcceleration, velocity, acceleration)
  ENQUEUEKERNEL(clearVelocityAcceleration, M)
  ENQUEUEKERNEL(updateKeys, M)
  SORT(keys)
  ENQUEUEKERNEL(sortPostProcess, M)
  ENQUEUEKERNEL(clearVoxels, V)
  ENQUEUEKERNEL(setVoxels, M)
end procedure
procedure SORT(keys)
  for stage in 0:log2(M)-1 do
     for pass in 0:stage do
        SETKERNELARGUMENTS(bitonicsort, keys, stage, pass)
        ENQUEUEKERNEL(bitonicSort, M/2)
     end for
  end for
end procedure
procedure UPDATESCENE
                                                   ▷ Called once per frame
  ENQUEUEKERNEL(updateKeys, M)
  SORT(keys)
```

CARP-ARM-RP-001-v1.0

10 November 2012





ENQUEUEKERNEL(sortPostProcess, M) ENQUEUEKERNEL(clearVoxels, V) ENQUEUEKERNEL(setVoxels, M) ENQUEUEKERNEL(calculateDensity, M) ENQUEUEKERNEL(calculateAcceleration, M) ENQUEUEKERNEL(integrate, M) end procedure

#### **Partitioned iteration space**

#### **Device memory access mapping**

#### **Performance metrics**

The performance is evaluated by taking the geometric mean of inverse frame times thus giving the geometric average of frames per second.

#### Validation mechanism

The benchmark does not contain any internal validation mechanisms. The validation is visual only.



# **Bibliography**

- [1] Alan Heirich. Smoothed particle hydrodynamics, 2010.
- [2] N. Bell and M. Garland. Implementing sparse matrix-vector multiplication on throughputoriented processors. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, SC '09, pages 18:1–18:11, New York, NY, USA, 2009. ACM.
- [3] C. Bienia, S. Kumar, J. P. Singh, and K. Li. The parsec benchmark suite: characterization and architectural implications. In *Proceedings of the 17th international conference on Parallel architectures and compilation techniques*, PACT '08, pages 72–81, New York, NY, USA, 2008. ACM.
- [4] M. Boyer, D. Tarjan, S. T. Acton, and K. Skadron. Accelerating leukocyte tracking using cuda: A case study in leveraging manycore coprocessors. In *Proceedings of the 2009 IEEE International Symposium on Parallel&Distributed Processing*, IPDPS '09, pages 1–12, Washington, DC, USA, 2009. IEEE Computer Society.
- [5] S. Che, M. Boyer, J. Meng, D. Tarjan, J. W. Sheaffer, S.-H. Lee, and K. Skadron. Rodinia: A benchmark suite for heterogeneous computing. In *Proceedings of the 2009 IEEE International Symposium on Workload Characterization (IISWC)*, IISWC '09, pages 44–54, Washington, DC, USA, 2009. IEEE Computer Society.
- [6] S. Che, J. W. Sheaffer, M. Boyer, L. G. Szafaryn, L. Wang, and K. Skadron. A characterization of the rodinia benchmark suite with comparison to contemporary cmp workloads. In *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC'10)*, IISWC '10, pages 1–11, Washington, DC, USA, 2010. IEEE Computer Society.
- [7] A. Corrigan, F. Camelli, R. Löhner, and J. Wallin. Running unstructured grid cfd solvers on modern graphics hardware. In *19th AIAA Computational Fluid Dynamics Conference*, number AIAA 2009-4001 in AIAA Conference Proceedings, June 2009.
- [8] A. Danalis, G. Marin, C. McCurdy, J. S. Meredith, P. C. Roth, K. Spafford, V. Tipparaju, and J. S. Vetter". The Scalable Heterogeneous Computing (SHOC) benchmark suite. In D. R. Kaeli and M. Leeser, editors, *Proceedings of 3rd Workshop on General Purpose Processing on Graphics Processing Units, GPGPU 2010, Pittsburgh, Pennsylvania, USA, March 14, 2010*, volume 425 of ACM International Conference Proceeding Series, pages 63–74. ACM, 2010.
- [9] F. Vázquez and E. M. Garzón and J. A. Martínez and J. J. Fernández. The sparse matrix vector product on GPUs. Technical report, University of Almeria, 2009.
- [10] M. A. Goodrum, M. J. Trotter, A. Aksel, S. T. Acton, and K. Skadron. Parallelization of particle filter algorithms. In *Proceedings of the 2010 international conference on Computer Architecture*, ISCA'10, pages 139–149, Berlin, Heidelberg, 2012. Springer-Verlag.
- [11] P. Harish and P. Narayanan. Designing Perspectively-Correct Multiplanar Displays. *IEEE Transactions on Visualization and Computer Graphics*, 99:1–1, 2012.

CARP-ARM-RP-001-v1.0



- [12] L. W. Howes, A. Lokhmotov, A. F. Donaldson, and P. H. Kelly. Deriving efficient data movement from decoupled access/execute specifications. In *Proceedings of the* 4th International Conference on High Performance and Embedded Architectures and Compilers (HiPEAC), volume 5409 of Lecture Notes in Computer Science, pages 168–182. Springer, 2009.
- [13] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. R. Stan. Hotspot: acompact thermal modeling methodology for early-stage VLSI design. *IEEE Trans. Very Large Scale Integr. Syst.*, 14(5):501–513, May 2006.
- [14] J. Meng and K. Skadron. Performance modeling and automatic ghost zone optimization for iterative stencil loops on gpus. In *Proceedings of the 23rd international conference on Supercomputing*, ICS '09, pages 256–265, New York, NY, USA, 2009. ACM.
- [15] T. M. Mitchell. Machine Learning. McGraw-Hill, New York, 1997.
- [16] Muthu Manikandan Baskaran and Rajesh Bordawekar. Optimizing Sparse Matrix-Vector Multiplication on GPUs. Technical report, IBM, 2008.
- [17] N. Ray and S. T. Acton. Motion gradient vector flow: an external force for tracking rolling leukocytes with shape and size constrained active contours. *IEEE Trans. Med. Imaging*, 23(12):1466–1478, 2004.
- [18] J. Stratton, C. Rodrigues, I. Sung, N. Obeid, L. Chang, N. Anssari, G. Liu, and W. Hwu. Parboil: A revised benchmark suite for scientific and commercial throughput computing. *Center for Reliable and High-Performance Computing*, 2012.
- [19] F. H. Streitz, J. N. Glosli, M. V. Patel, B. Chan, R. K. Yates, B. R. de Supinski, J. Sexton, and A. Gunnels. 100+ TFlop Solidification Simulations on BlueGene/L. In *Proceedings* of the 2005 Supercomputing Conference (SC 05), Seattle, WA, 2005.
- [20] L. G. Szafaryn, T. Gamblin, B. deSupinski, and K. Skadron. Experiences with Achieving Portability across Heterogeneous Architectures. In WOLFHPC workshop at 25th International Conference on Supercomputing, Tucson, AZ, 2010.
- [21] L. G. Szafaryn, K. Skadron, and J. J. Saucerman. Experiences accelerating matlab systems biology applications. In *Proceedings of the Workshop on Biomedicine in Computing: Systems, Architectures, and Circuits (BiC)*, June 2009.
- [22] Y. Yu and S. Acton. Speckle reducing anisotropic diffusion. In *IEEE Transactions on Image Processing*, pages 1260–1270, 2002.