Design Patterns and Best Practices

Daniel Cámpora | thematic CERN School of Computing

Table of Contents

Precision for Computing

When using decimal numbers, one typically uses floating point numbers.

The IEEE 754 standard defines single precision, double precision, and half precision:

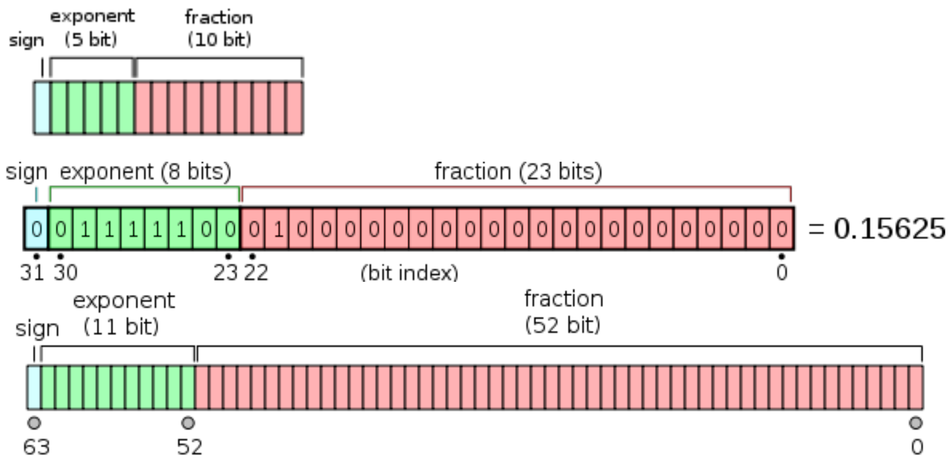

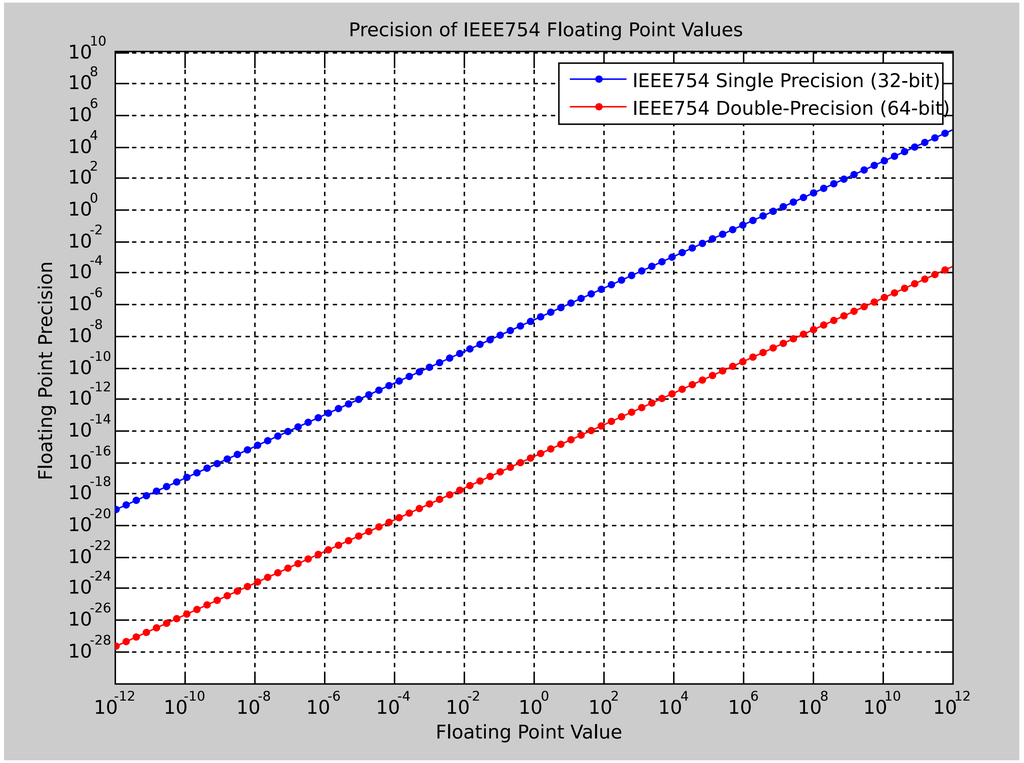

Still, how much precision is enough?

- Depends on your use case.
- Use the least you need.

What Precision Does Your Algorithm Need?

The answer may not be as simple as one-precision-fits-all. You should consider what precision you need for **storage** and **arithmetic**!

Speed

Arithmetic	**Storage**
Double precision	Double precision
Double precision	Single precision
Single precision	Single precision
Single precision	Half precision
Half precision	Half precision

Precision

Precision Matters (a Lot) in GPUs

This is especially true for GPUs, where you would have to go for scientific cards to be able to get good performance with double precision.

Theoretical performance	FP16 (TFLOPS)	FP32 (TFLOPS)	FP64 (TFLOPS)
RTX 4090 (consumer)	82.58	82.58	1.29
RTX 6000 ADA (professional)	91.06	91.06	1.42
H100 (scientific)	204.9	51.22	25.61

Floating Point Rounding

The standard describes four rounding modes:

- **round to nearest** (typically the default)
- round down
- round up
- round towards zero

In addition, Fused Multiply-Add (FMA) units add precision and performance when doing floating point operations… which changes slightly the result!

With no optimization flags, GPU compilers have FMAs turned on as opposed to CPU compilers. However, as a general rule *one should not expect FP bit-level precision across different architectures or compilers*, even if they run under the same standard!

Floating Point Rounding – An Example

Consider the dot product example:

$$\vec{a} = \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{bmatrix} \quad \vec{b} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \end{bmatrix} \quad \vec{a} \cdot \vec{b} = a_1 b_1 + a_2 b_2 + a_3 b_3 + a_4 b_4$$

Dot Product – Serial Approach

In the serial approach, every element is calculated sequentially.

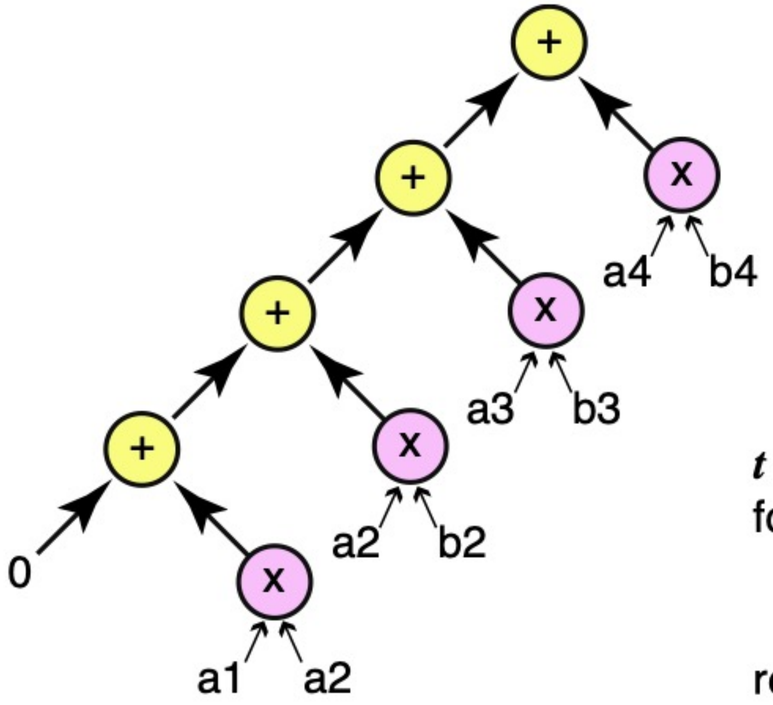

$t = 0$
for i from 1 to 4
$\quad p = \text{rn}\ (a_i \times b_i)$
$\quad t = \text{rn}\ (t + p)$
return t

Dot Product – FMA Method

Using FMAs, each subsequent multiplication and addition is done in one instruction.

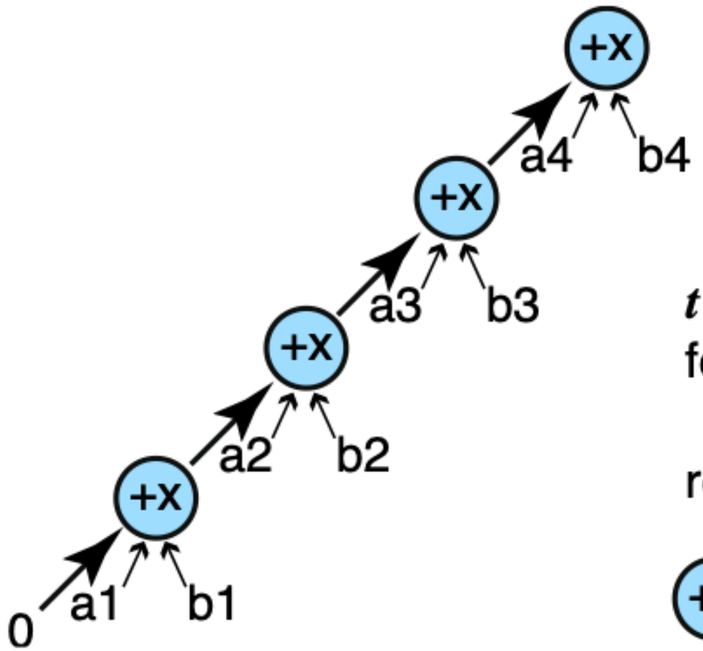

$t = 0$
for i from 1 to 4
$\quad t = \mathrm{rn}\,(a_i \times b_i + t)$
return t

+x Fused multiply-add

Dot Product – Parallel Method

The parallel method divides the problem such that each multiplication is done, followed by the additions, in a Divide and Conquer approach.

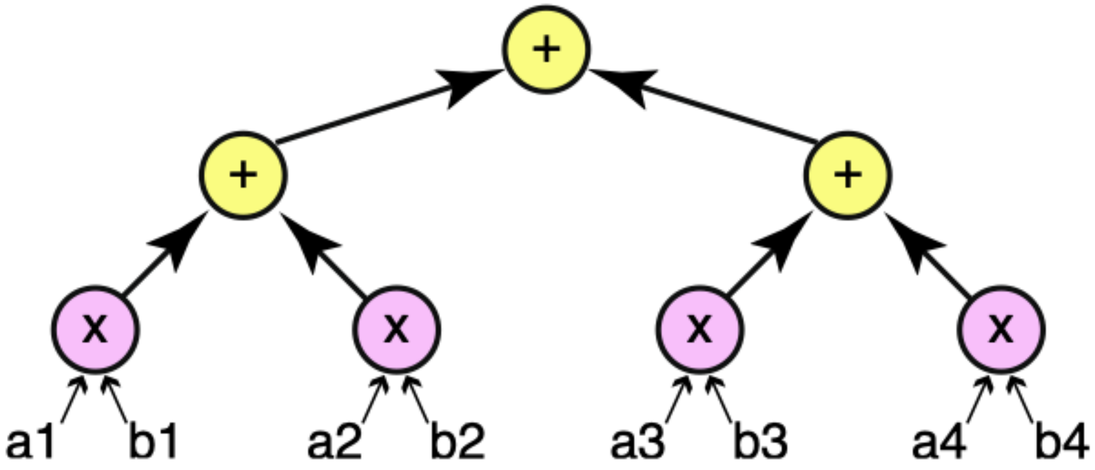

$p1 = \text{rn}\ (a_1 \times b_1)$
$p2 = \text{rn}\ (a_2 \times b_2)$
$p3 = \text{rn}\ (a_3 \times b_3)$
$p4 = \text{rn}\ (a_4 \times b_4)$
$s_{left} = \text{rn}\ (p_1 + p_2)$
$s_{right} = \text{rn}\ (p_3 + p_4)$
$t = \text{rn}\ (s_{left} + s_{right})$
return t

Results

The results for vectors:

```
a = [1.907607, -0.7862027, 1.147311, 0.9604002]
b = [-0.9355000, -0.6915108, 1.724470, -0.7097529]
```

is

method	result	float value
exact	.0559587528435...	0x3D65350158...
serial	.0559588074	0x3D653510
FMA	.0559587515	0x3D653501
parallel	.0559587478	0x3D653500

Compiler Flags

Last, you have some control over how your computation is done. You may want to consider *fast math*, which can impact performance and results quite substantially.

mode	**flags**

IEEE 754 mode (default)	`-ftz=false`
	`-prec-div=true`
	`-prec-sqrt=true`

fast mode	`-ftz=true`
	`-prec-div=false`
	`-prec-sqrt=false`

Bear in mind with fast mode:

- Denormals are flushed to zero.
- Division and square root are not computed to the nearest FP value.

A Practical Example

What is wrong with the following code?

```
__global__ void shared_memory_example(float* dev_array) {
  for (int i = threadIdx.x; i < 256; i += blockDim.x) {
    dev_array[i] = 1 / std::sqrt(2. + dev_array[i]);
  }
}
```

A Practical Example

What is wrong with the following code?

```
__global__ void shared_memory_example(float* dev_array) {
  for (int i = threadIdx.x; i < 256; i += blockDim.x) {
    dev_array[i] = 1 / std::sqrt(2. + dev_array[i]);
  }
}
```

Use compiler flag `-Wdouble-promotion` to avoid surprises!

Table of Contents

Register Spilling

Every thread has a maximum number of registers it can use:

- In GPUs, this limit is configurable (typically between 63 and 255).
- If this limit is surpassed, the kernel will use **local memory** as swap space.

It is “local” because each thread has its own private area. It is actually stored in **global memory** (yes, the slow one).

Register Spilling (2)

Developers have no control over the spilling process:

- Address of global memory where memory is swapped is resolved by compiler.
- Stores are cached in L1 memory.

Register Spilling (3)

Developers have no control over the spilling process:

- Address of global memory where memory is swapped is resolved by compiler.
- Stores are cached in L1 memory.

Spilling could hurt performance:

- Increases memory traffic.
- Increases instruction count.

Register Spilling (4)

Developers have no control over the spilling process:

- Address of global memory where memory is swapped is resolved by compiler.
- Stores are cached in L1 memory.

Spilling could hurt performance:

- Increases memory traffic.
- Increases instruction count.

But it is not always bad:

- If accesses are cached.
- If your code is not instruction-throughput limited.

How to Deal With Register Spilling

One can evaluate the impact of register spilling through profiling.

The developer has several tools to impact register spilling:

- Increase globally the amount of registers per kernel.
- Increase the amount of configurable L1 cache.
- Some compilers allow to specify *non-caching loads* for global memory.
- `__launch_bounds__` (HIP, CUDA) – Controls *maximum threads per block* and *minimum blocks per SM*. These two impact the number of registers in a kernel.

Write Single-Source Kernels

It is possible to organize the code with several header files, containing **definitions**, and source files, containing **implementations**.

However, doing so in GPU code heavily affects performance. The reason is that the compiler optimizes functions to use a number of registers, shared memory and threads, and it cannot perform that optimization if the compilation unit cannot see all code involved.

In other words, if your `__global__` function calls `__device__` functions either free standing or within structs, those should be defined in either:

- The same source file.
- A header file, either **templated** or **inlined**.

A Practical Use-case: The Velo Pixel Subdetector of LHCb

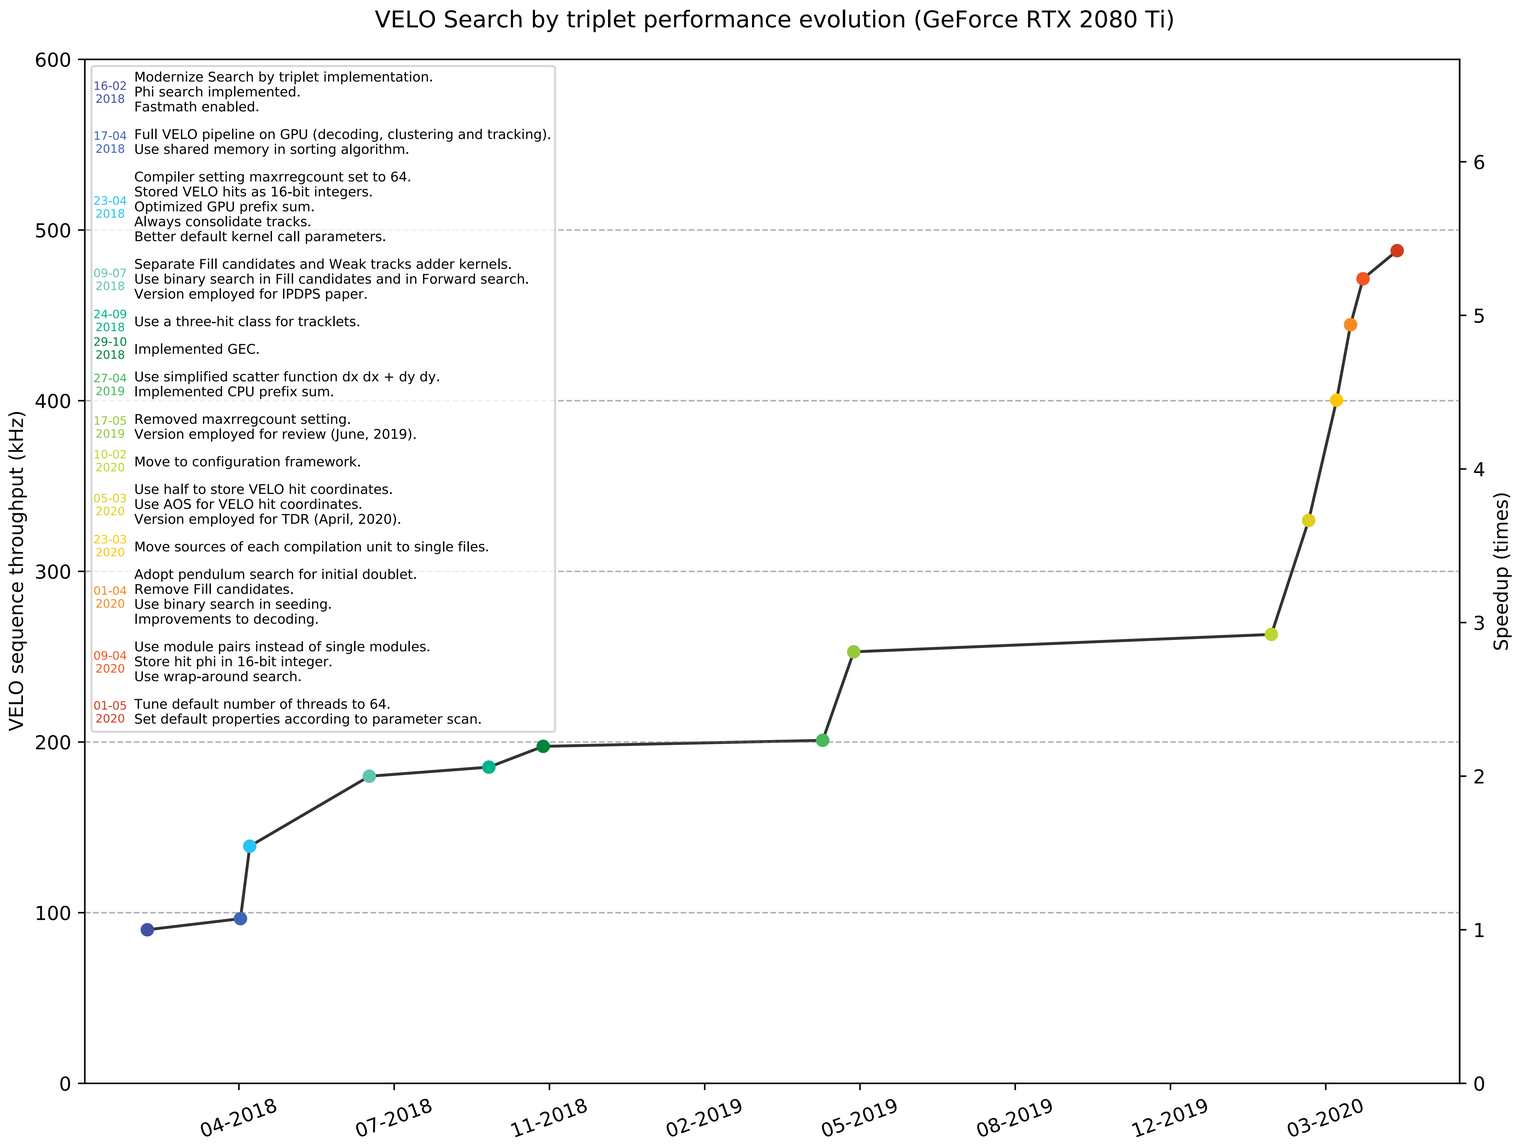

Table of Contents

OpenCL

OpenCL™ (Open Computing Language) is a multi-vendor open standard for general-purpose parallel programming of heterogeneous systems that include CPUs, GPUs, and other devices.

OpenCL 3.0 was released in September 2020:

- It has increased the flexibility by making every functionality from OpenCL 1.2 onwards **optional and queryable**.
- C++ for OpenCL adopts C++17.
- Unified specification.

OpenCL Compilation

OpenCL can either be compiled *offline* or *online*:

- **Offline compilation**: Kernel is pre-built with an OpenCL compiler. Pro: It runs with low invocation latency. Con: It is compiled for a specific architecture.
- **Online compilation**: Kernel source code is distributed instead. Pro: It can run on various architectures. Con: It needs to be JIT-compiled.

Offline compilation is supported through the clang compiler:

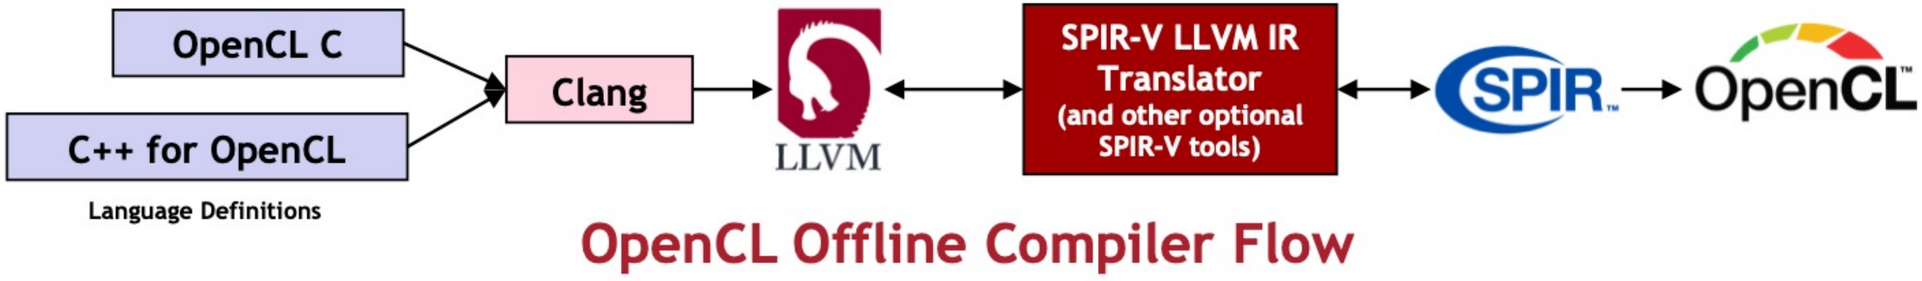

OpenCL Offline Compiler Flow

C++ for OpenCL

C++17 arrived with the release of *C++ for OpenCL*. The C++ support has some caveats:

2.1.1. Restrictions to C++ features

The following C++ language features are not supported:

- Virtual functions (C++17 `[class.virtual]`);
- References to functions including member functions (C++17 `[class.mfct]`);
- Pointers to class member functions (in addition to the regular non-member functions that are already restricted in OpenCL C);
- Exceptions (C++17 `[except]`);
- `dynamic_cast` operator (C++17 `[expr.dynamic.cast]`);
- Non-placement `new`/`delete` operators (C++17 `[expr.new]`/`[expr.delete]`);
- Standard C++ libraries (C++17 `[library]`).

Status of OpenCL as of 2024

- OpenCL hasn't changed since 2020.
- It gained back an implementation by NVIDIA.
- At the same time, Apple stopped supporting OpenCL (in favor of Metal), and so did AMD (in favor of ROCm).

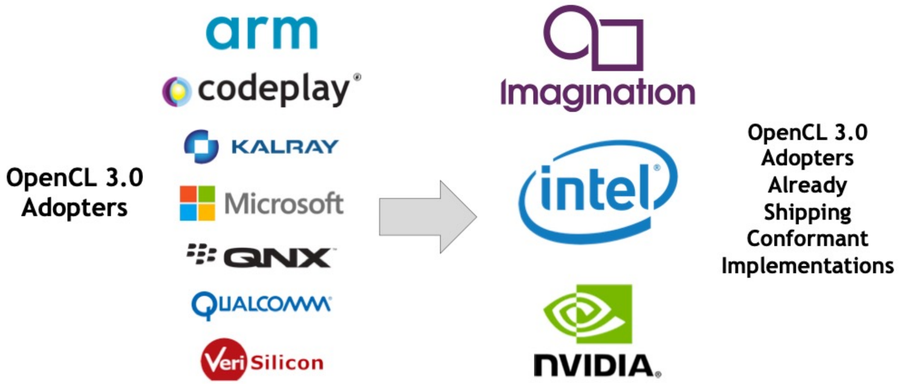

For more information on OpenCL check out:

- Khronos website on OpenCL
- IWOCL 2021 presentation

ROCm

ROCm is a platform that has appeared in recent years and is quickly evolving and adapting. It includes:

- Support of frameworks (Tensorflow / Pytorch)
- Libraries (MIOpen / Blas / RCCL)
- Programming model (HIP)
- Inter-connect (OCD)

HIP

HIP is a high performance, CUDA-like programming model that is built on an open and portable framework.

- It supports C++17.
- It is almost a 1:1 copy of CUDA – most of the time changes required are very minimal and non-intrusive.
- It supports AMD and NVIDIA targets.

Differences Between HIP and CUDA

- Library call prefix is hip instead of cuda.
- Warp size depends on GPU: 64 on AMD and 32 on NVIDIA.
- Profiling / debugging is not as advanced.
- Low-level calls are different, newer CUDA features are ahead of HIP support.
- Specialized hardware (tensor cores) is naturally not there.

For more information: ROCm docs.

SYCL

SYCL 2020's primary goal is to achieve closer convergence with ISO C++, furthering our work to bring parallel heterogeneous programming to modern C++ through open standards.

- It supports C++20, its intent is to become part of the standard.
- It attempts to support *everything* (CPUs, AMD GPUs, NVIDIA GPUs, Intel GPUs, Intel FPGAs).
- SYCL is built on top of OpenCL and SPIR– (the low-level representation shared by eg. Vulkan or OpenGL).

From the Creators of…

It is another standard developed by the Khronos group (same as OpenCL).

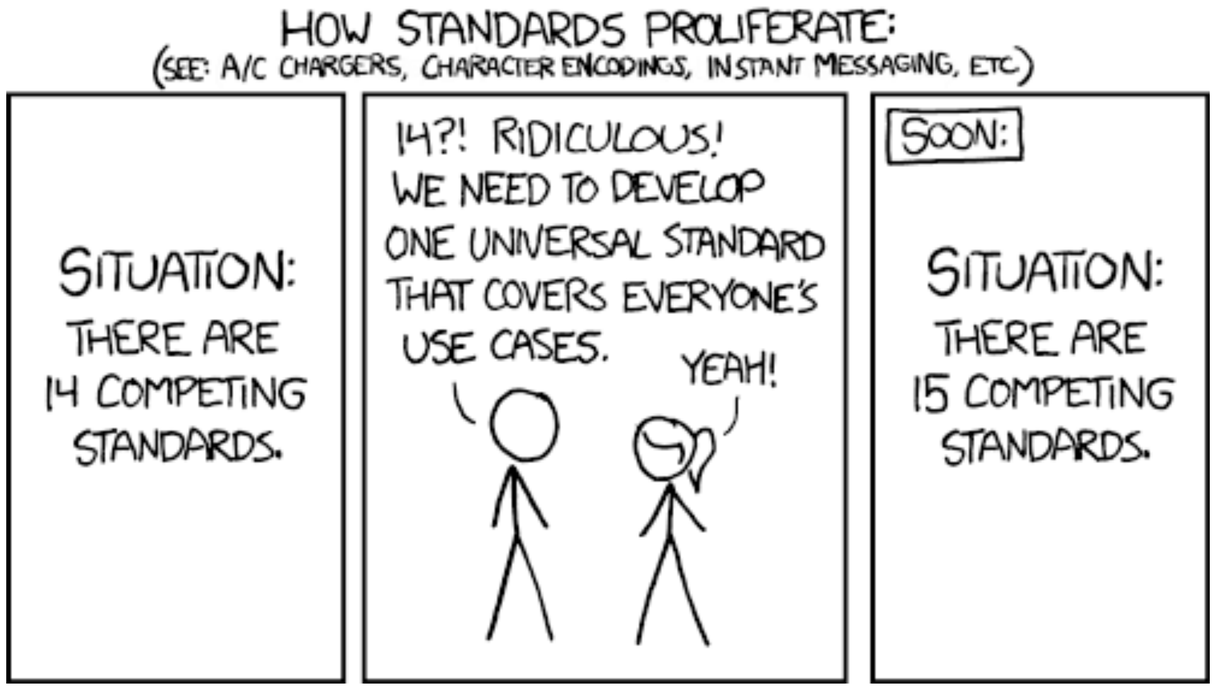

From XKCD comic Standards (https://xkcd.com/927/)

SYCL to Everything

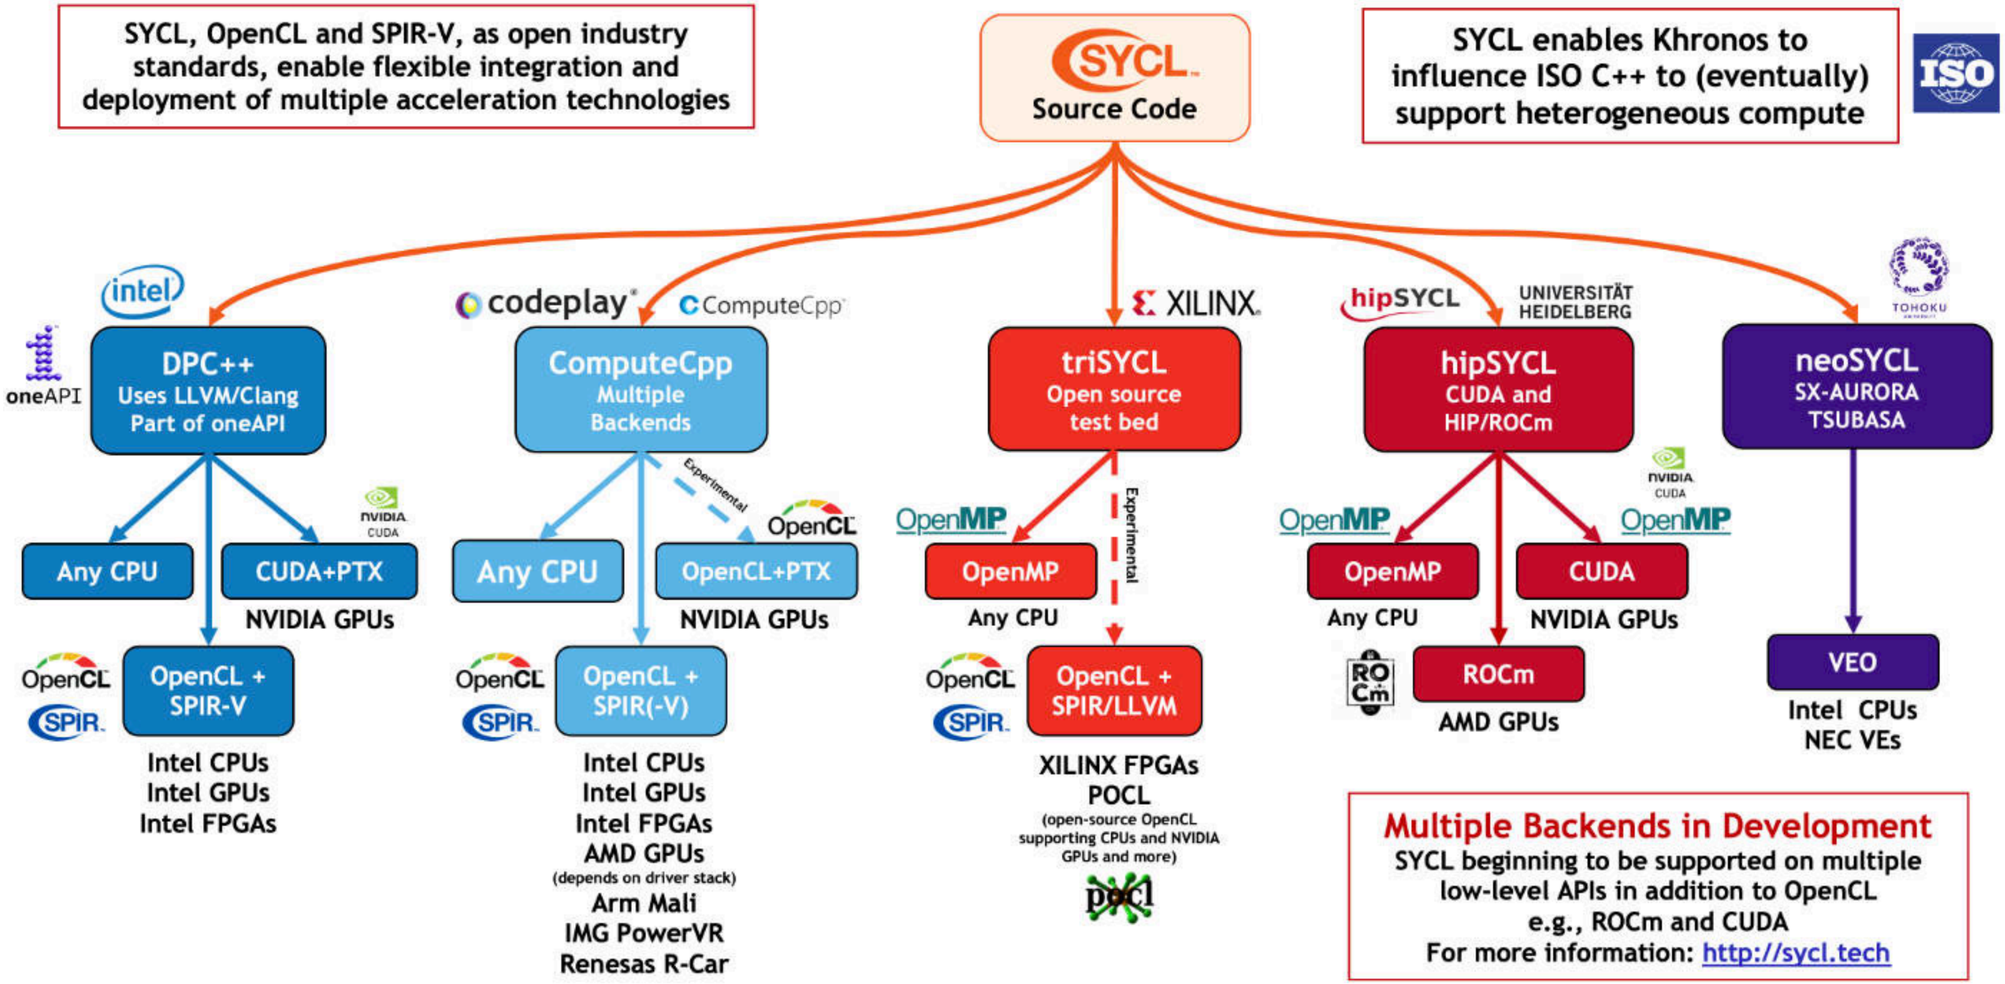

Status of SYCL

SYCL released and updated the SYCL 2020 specification.

- Intel supports SYCL as a first-class citizen through its release of OneAPI.
- Intel GPUs are supported.

- Syntax is not easily translatable from CUDA / HIP. Adapting requires work.
- There is no one-size-fits-all and there will never be.
- Given that CUDA is a low-level language, adapting to a higher level one may not be in your best interest if performance is your goal.

In spite of all that SYCL looks very interesting, the next few years will determine whether the language has wide adoption.

Table of Contents

Alpaka

The alpaka library is a header-only C++14 abstraction library for accelerator development.

- It acts as a middle layer that can target CPU, NVIDIA GPUs or AMD GPUs through a variety of backends.
- C++-style API which is optimized away by the compiler.

Alpaka Backends

Accelerator Back-end	Lib/API	Devices	Execution strategy grid-blocks	Execution strategy block-threads
Serial	n/a	Host CPU (single core)	sequential	sequential (only 1 thread per block)
OpenMP 2.0+ blocks	OpenMP 2.0+	Host CPU (multi core)	parallel (preemptive multitasking)	sequential (only 1 thread per block)
OpenMP 2.0+ threads	OpenMP 2.0+	Host CPU (multi core)	sequential	parallel (preemptive multitasking)
OpenMP 5.0+	OpenMP 5.0+	Host CPU (multi core)	parallel (undefined)	parallel (preemptive multitasking)
		GPU	parallel (undefined)	parallel (lock-step within warps)
OpenACC (experimental)	OpenACC 2.0+	Host CPU (multi core)	parallel (undefined)	parallel (preemptive multitasking)
		GPU	parallel (undefined)	parallel (lock-step within warps)
std::thread	std::thread	Host CPU (multi core)	sequential	parallel (preemptive multitasking)
Boost.Fiber	boost::fibers::fiber	Host CPU (single core)	sequential	parallel (cooperative multitasking)
TBB	TBB 2.2+	Host CPU (multi core)	parallel (preemptive multitasking)	sequential (only 1 thread per block)
CUDA	CUDA 9.0+	NVIDIA GPUs	parallel (undefined)	parallel (lock-step within warps)
HIP(clang)	HIP 4.0+	AMD GPUs	parallel (undefined)	parallel (lock-step within warps)

Alpaka Hello World Kernel

```
#include <alpaka/alpaka.hpp>
//! Prints "[x, y, z][gtid] Hello World" where tid is the global thread number.
struct HelloWorldKernel
{
  template<typename TAcc>
  ALPAKA_FN_ACC auto operator()(TAcc const& acc) const -> void
  {
    using Dim = alpaka::Dim<TAcc>;
    using Idx = alpaka::Idx<TAcc>;
    using Vec = alpaka::Vec<Dim, Idx>;
    using Vec1 = alpaka::Vec<alpaka::DimInt<1u>, Idx>;

    Vec const globalThreadIdx = alpaka::getIdx<alpaka::Grid, alpaka::Threads>(acc);
    Vec const globalThreadExtent = alpaka::getWorkDiv<alpaka::Grid, alpaka::Threads>(acc);
    Vec1 const linearizedGlobalThreadIdx = alpaka::mapIdx<1u>(globalThreadIdx, globalThreadExtent);
    printf(
      "[z:%u, y:%u, x:%u][linear:%u] Hello World\n",
    static_cast<unsigned>(globalThreadIdx[0u]),
    static_cast<unsigned>(globalThreadIdx[1u]),
    static_cast<unsigned>(globalThreadIdx[2u]),
    static_cast<unsigned>(linearizedGlobalThreadIdx[0u]));
  }
};
```

Should You Use Alpaka?

It depends on the priorities of your project:

- It provides portability across different platforms.
- Easier to maintain.
- There is a delay between appeareance of new features and Alpaka support.
- It remains a thin library, but doesn't give you the same control and flexibility as a low-level language (CUDA, HIP).
- Requires learning a different language extension which is not so widely adopted and departs from others (learning curve).

Materials on Alpaka:

- Repository, documentation, online tutorial.

Kokkos

Kokkos Core implements a programming model in C++ for writing performance portable applications targeting all major HPC platforms.

- Very similar to Alpaka conceptually, also a header-only library.
- C++ middle layer that targets CPU, NVIDIA, AMD platforms.

Kokkos Core Capabilities

Concept	Example
Parallel Loops	**parallel_for**(N, **KOKKOS_LAMBDA** (int i) { ...BODY… });
Parallel Reduction	**parallel_reduce**(**RangePolicy**<ExecSpace>(0,N), **KOKKOS_LAMBDA** (int i, double& upd) { …BODY... upd += ... }, Sum<>(result));
Tightly Nested Loops	**parallel_for**(**MDRangePolicy**<**Rank**<3> > ({0,0,0},{N1,N2,N3},{T1,T2,T3}, **KOKKOS_LAMBDA** (int i, int j, int k) {…BODY...});
Non-Tightly Nested Loops	**parallel_for**(**TeamPolicy**<**Schedule**<**Dynamic**>>(N, TS), **KOKKOS_LAMBDA** (Team team) { … COMMON CODE 1 ... **parallel_for**(**TeamThreadRange**(team, M(N)), [&] (int j) { ... INNER BODY... }); … COMMON CODE 2 ... });
Task Dag	**task_spawn**(**TaskTeam**(scheduler , priority), **KOKKOS_LAMBDA** (Team team) { … BODY });
Data Allocation	**View**<double**, Layout, MemSpace> a("A",N,M);
Data Transfer	**deep_copy**(a,b);
Atomics	atomic_add(&a[i],5.0); View<double*,MemoryTraits<AtomicAccess>> a(); a(i)+=5.0;
Exec Spaces	Serial, Threads, OpenMP, Cuda, HPX (experimental), ROCm (experimental)

Kokkos Reduction Kernel

```
struct squaresum {
  using value_type = int;
  KOKKOS_INLINE_FUNCTION
  void operator()(const int i, int& lsum) const {
    lsum += i * i; // compute the sum of squares
  }
};

int n = 10, sum = 0;
Kokkos::parallel_reduce(n, squaresum(), sum);
```

Should You Use Kokkos?

- It provides portability and maintainability.
- It provides some higher level functionality (eg. parallel reduction, execution policies).

- Documentation is scarce, learning curve.
- It provides a subset of low-level functionality of vendor-driven languages (CUDA, HIP, SYCL).

Materials on Kokkos:

- Repository.
- GTC presentation, video.

Bonus: Portability in CUDA / HIP?

- CUDA and HIP are very similar standards.
- Kernel language is practically the same for most use-cases.
- For hardware-specific optimizations (eg. tensor cores), you would have to implement a portable version by hand for portability.

- Utility functions (memcpy, memset, malloc, kernel invocation…) can be defined with macros or a hand-made middle language.
- In practice, making your own middleware just for utility functions is very little work.

- It is therefore possible to have one codabease with a low-maintainance self-developed wrapper just covering your utility / kernel needs.
- If targeting performance, it is always better to use the native solution.
- What about CPU execution?

Bonus: CUDA / HIP Running on CPU?

If the CUDA code satisfies that it produces the same result when invoked with a block dimension of `{1, 1, 1}` – or in other words:

- for-loops over threads are block-dimension strided.
- if-statements for a single thread refer to threads of index 0.

Then, with some macros and function definitions it is possible to compile the code for CPUs.

See the following presentation for details.

So, Standard or Middleware?

- Standards offer the best performance for their native platform.
- Middlewares offer portability.

- It is possible to obtain good performance on a middleware.
- Low-level functions not supported by the middleware will require your own implementation across vendors (high effort).

- It is possible to achieve portability between CUDA / HIP / CPU.
- You may want to focus on a single CPU backend if you do this (as opposed to the many variants offered by Alpaka for instance).
- You will maintain the portability layer.

What is your application's main target?

Table of Contents

Data Parallel vs Streaming Patterns

Data parallel patterns

- Map
- Farm
- Reduction
- Stencil

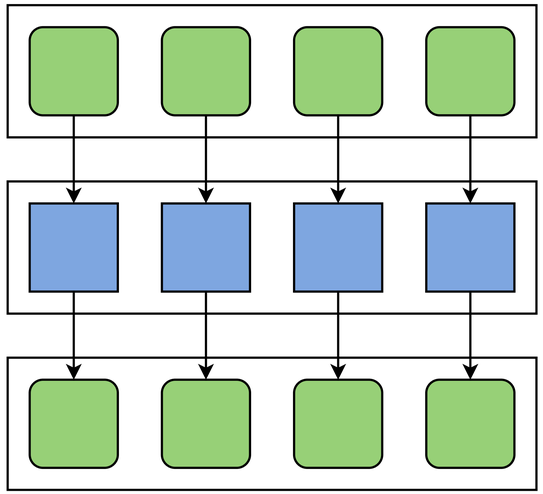

Streaming patterns

- Farm
- Pipeline

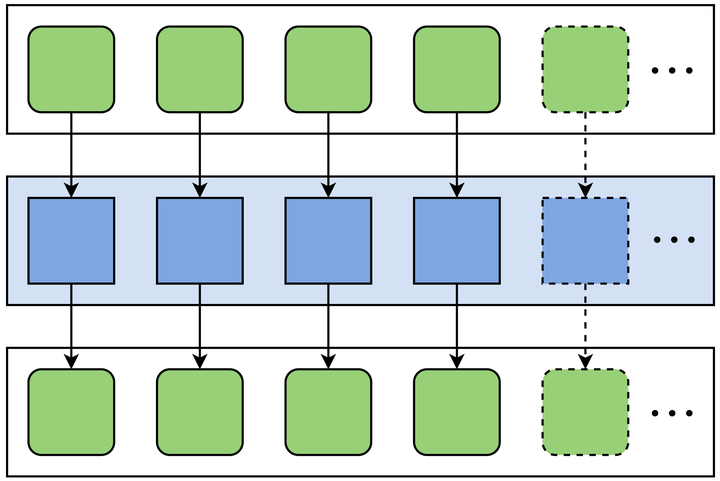

Data Parallel vs Streaming Parallel Patterns

- Size of the input + dependencies between items define which patterns to use.
- Data parallel patterns may not be efficient in streaming scenarios, and the other way around.
- For streaming patterns, there is usually one (or more) input items that distribute the input elements to working items as they come.

Map (parallel)

- Used on embarrassingly parallel collections of items.
- Same function applied to every item, all at the same “time”.
- Applicable if all items are independent.
- Usually good candidate for SIMD abstractions.

Used in
Ray tracing, Monte Carlo simulations

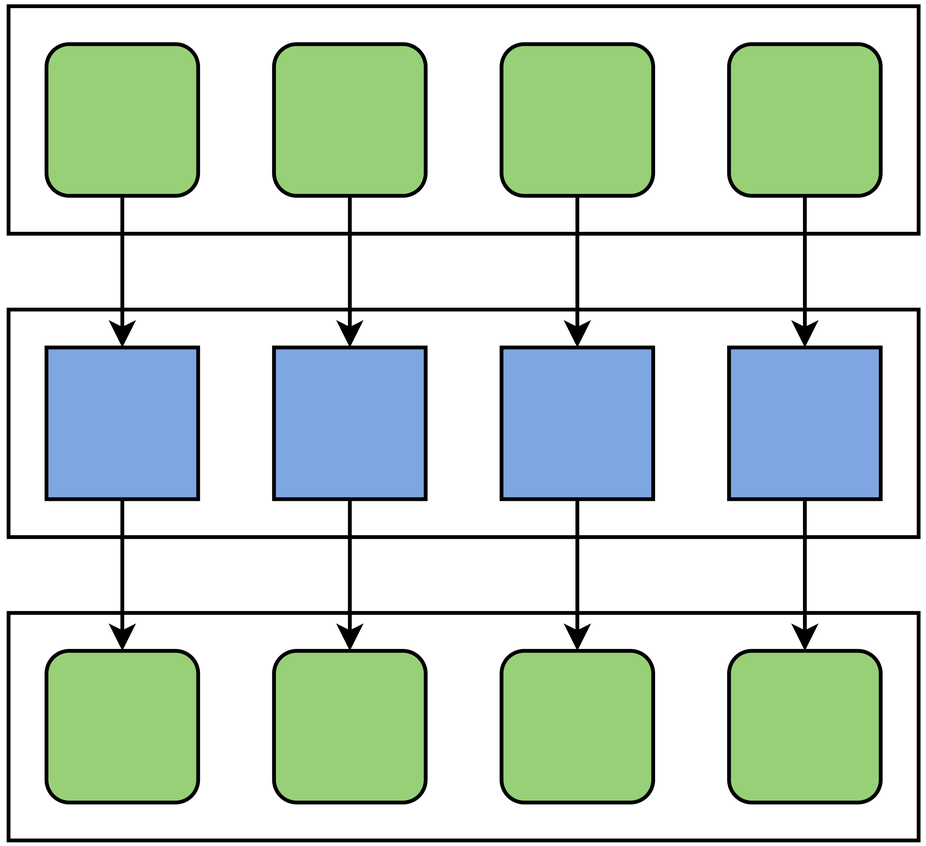

Stencil (parallel)

- When for every item of a collection, we need data from the neighbourhood items.
- Usually a fixed number of neighbourhood is accessed.
- Boundary conditions have to be taken into account.
- Data reuse in the implementation (cache).

Used in
Convolutional neural networks, signal filtering, image processing, grid methods.

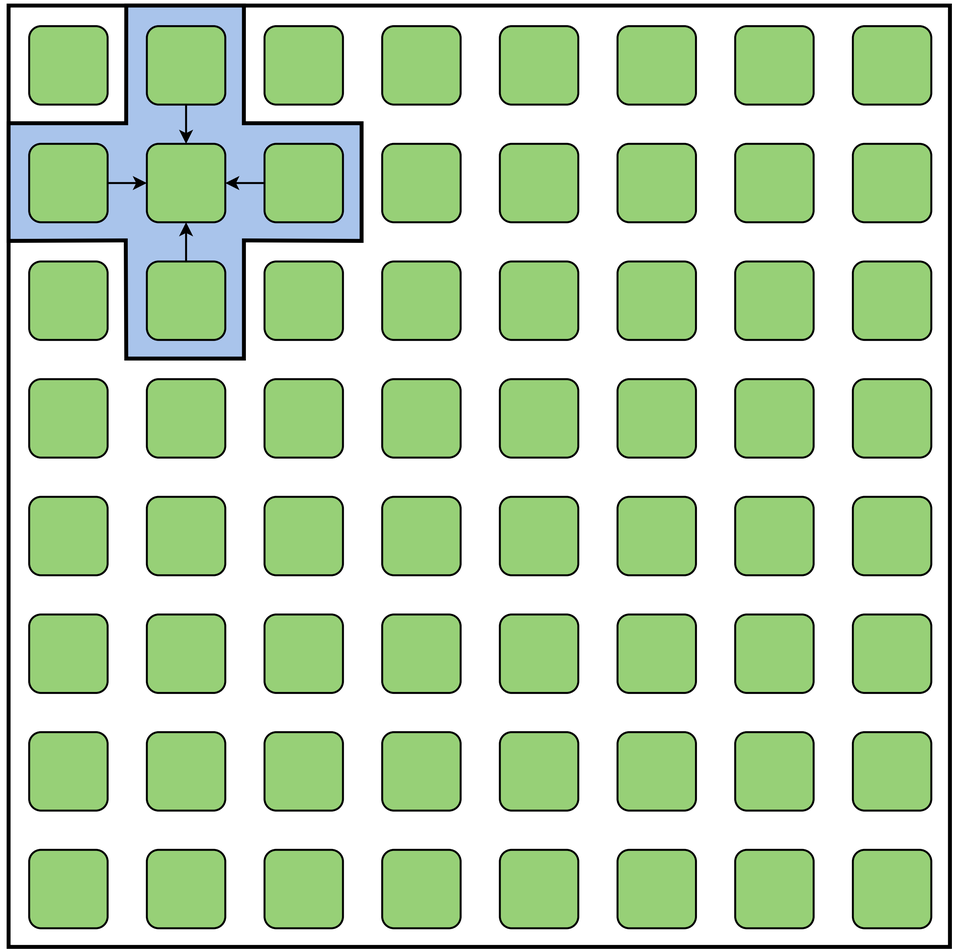

Stencil (parallel)

- When for every item of a collection, we need data from the neighbourhood items.
- Usually a fixed number of neighbourhood is accessed.
- Boundary conditions have to be taken into account.
- Data reuse in the implementation (cache).

Used in
Convolutional neural networks, signal filtering, image processing, grid methods.

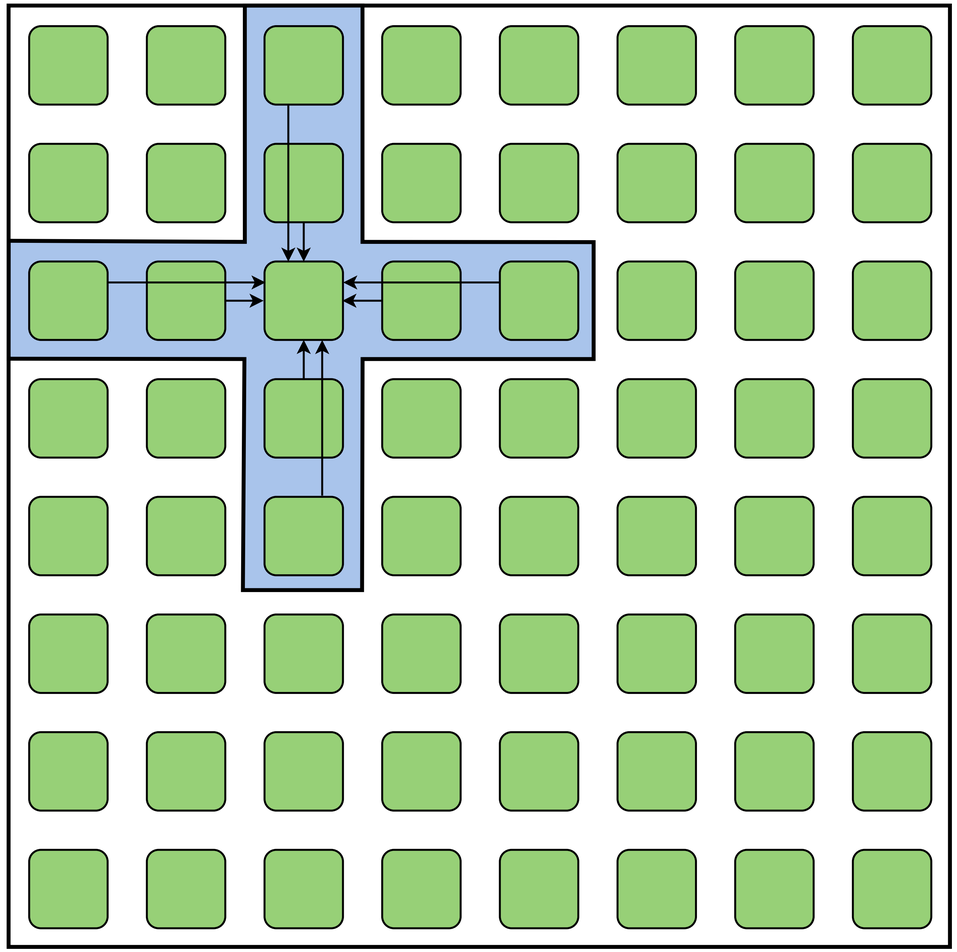

Farm (Parallel Streaming)

- Similar to map, but size of collection is not known in advance.
- Used for embarrassingly parallel computations in stream computations.
- There is at least one producer item.

Used in
HEP online trigger software.

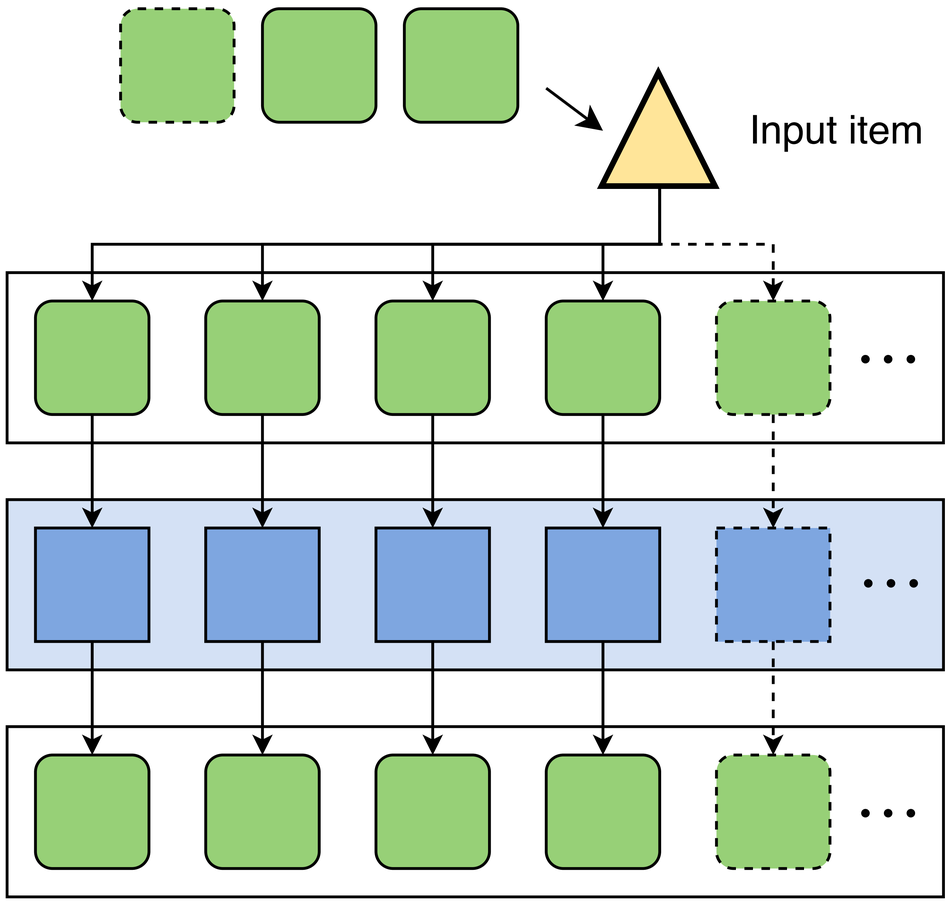

Pipeline (Streaming)

- Size of collection not needed in advance.
- Different steps run in parallel, but others may not be able to run in parallel.
- Different functions are applied in different steps, where the order is important.

Used in
Image filtering, signal processing, game engines.

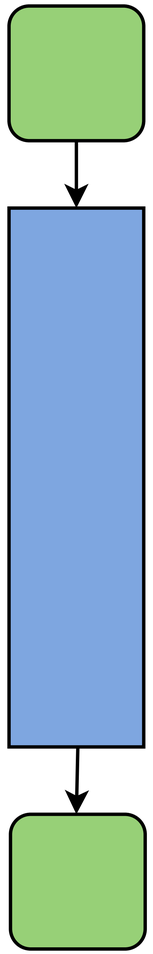

Pipeline (Streaming)

- Size of collection not needed in advance.
- Different steps run in parallel, but others may not be able to run in parallel.
- Different functions are applied in different steps, where the order is important.

Used in
Image filtering, signal processing, game engines.

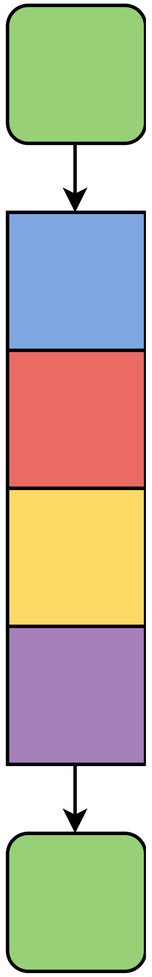

Pipeline (Streaming)

- Size of collection not needed in advance.
- Different steps run in parallel, but others may not be able to run in parallel.
- Different functions are applied in different steps, where the order is important.

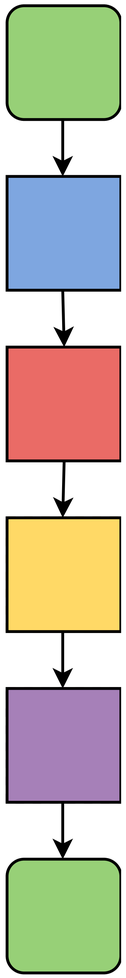

Used in

Image filtering, signal processing, game engines.

Pipeline (Streaming)

- Size of collection not needed in advance.
- Different steps run in parallel, but others may not be able to run in parallel.
- Different functions are applied in different steps, where the order is important.

Used in
Image filtering, signal processing, game engines.

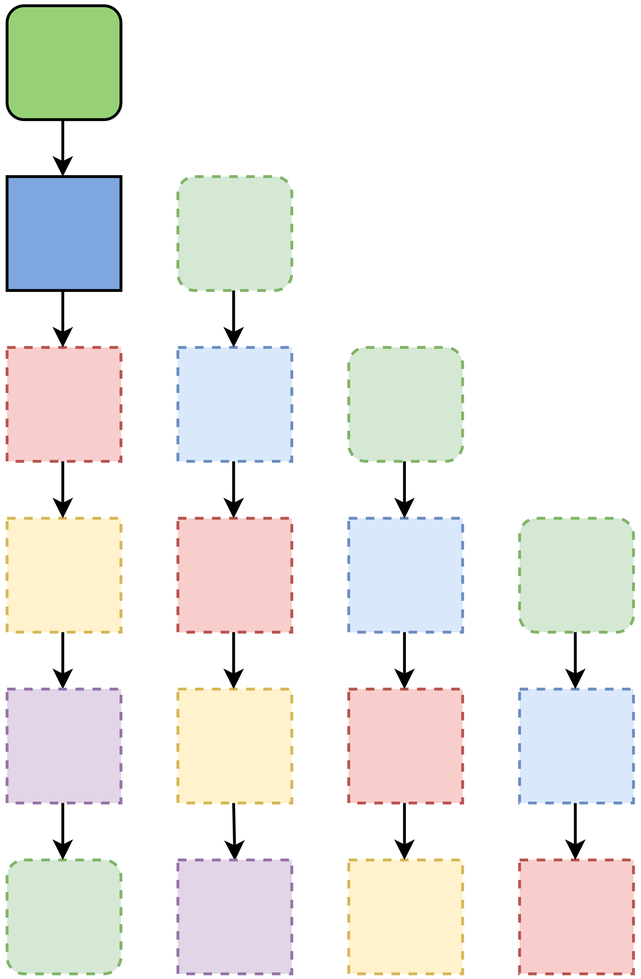

Pipeline (Streaming)

- Size of collection not needed in advance.
- Different steps run in parallel, but others may not be able to run in parallel.
- Different functions are applied in different steps, where the order is important.

Used in

Image filtering, signal processing, game engines.

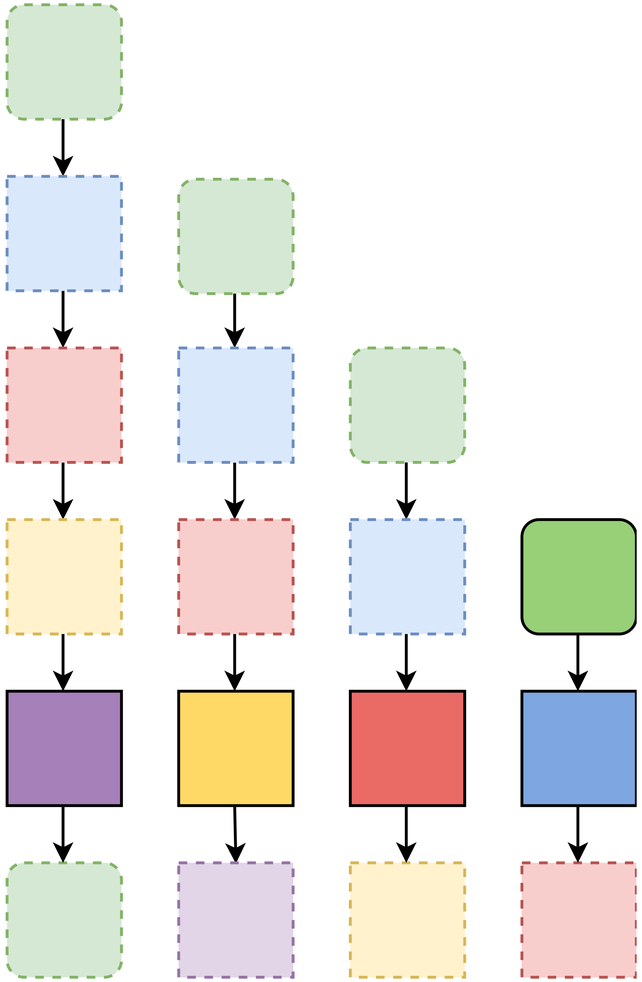

Reduction (Sequential)

- Combines a collection of items into one, with a defined operation.
- Many different partition options.
- Elements depend on each other, but are associative.

Used in

Matrix operations, computing of statistics on datasets.

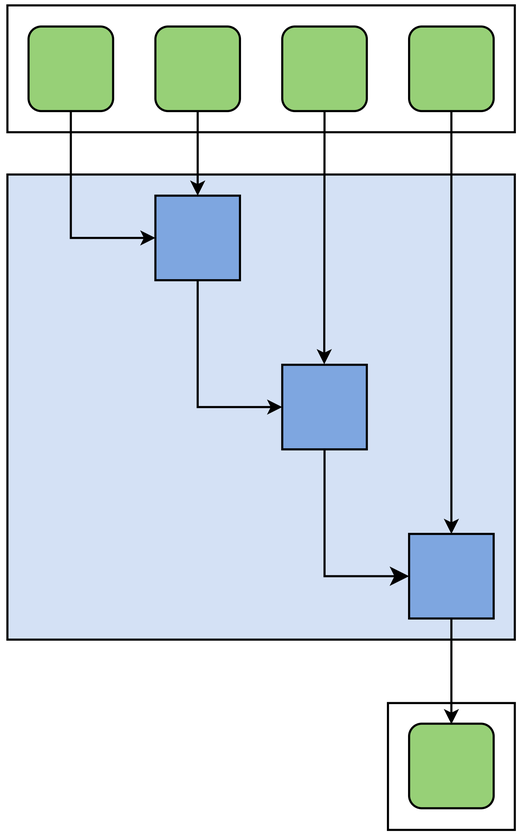

Reduction (Parallel)

- Combines a collection of items into one, with a defined operation.
- Many different partition options.
- Elements depend on each other, but are associative.

Used in

Matrix operations, computing of statistics on datasets.

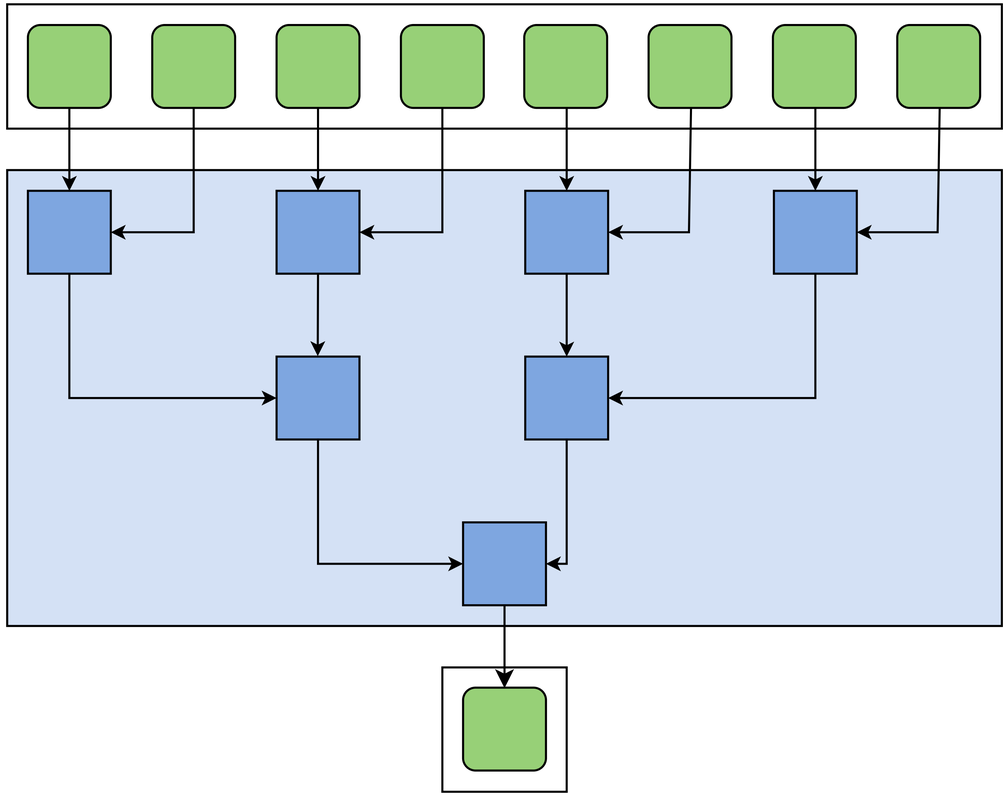

Prefix Sum Example

The prefix sum is a problem that consists in calculating the accumulated sum at every element of an array. For example:

- Number of tracks per event: `[10, 15, 32, 45, 24]`
- Prefix sum: `A = [0, 10, 25, 57, 102, 126]`

The prefix sum of an array of numbers is extremely useful. It provides:

- The accumulated sum of the entire array (last element).
- The offset of each element (on element `A[i]`).
- The size of each element (`A[i+1] - A[i]`).

Efficient GPU Prefix Sum: Blelloch Scan

The Blelloch scan consists in performing two sweeps of the data.

- *up-sweep* – the tree is traversed from leaves to root computing partial sums (**reduction**).

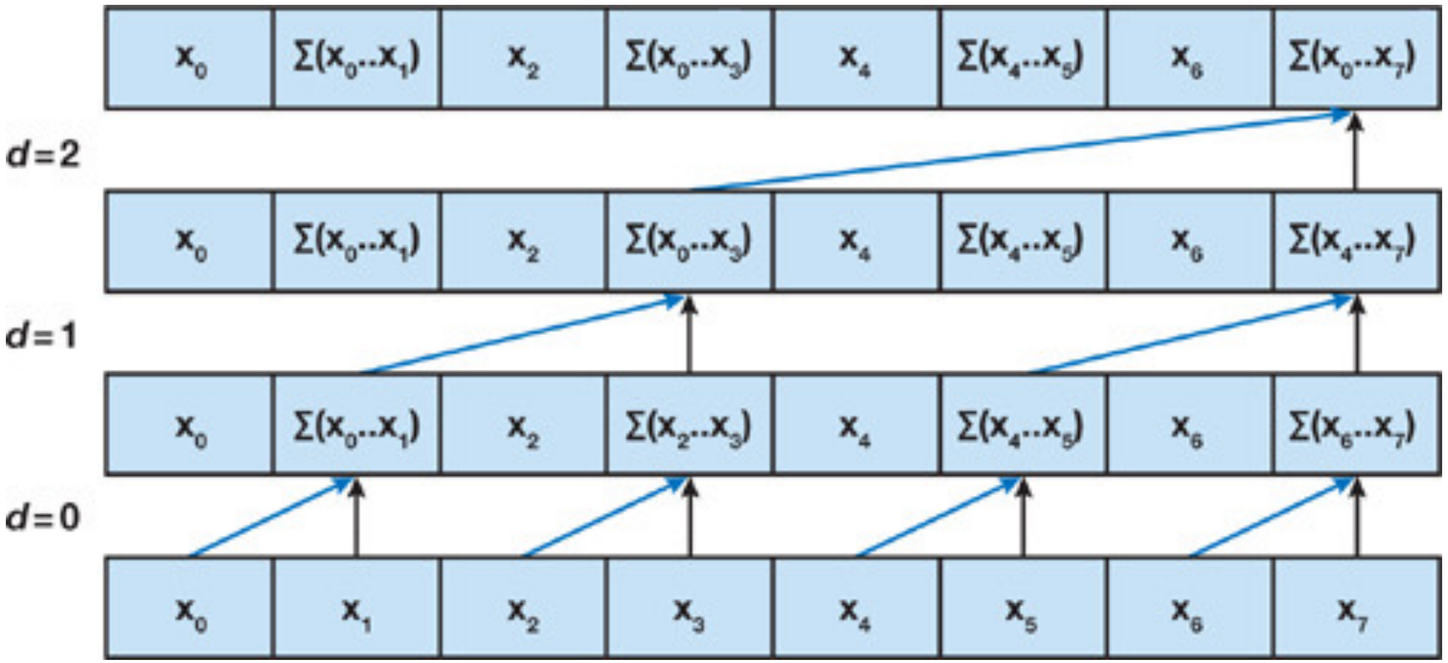

Efficient GPU Prefix Sum: Blelloch Scan (2)

The Blelloch scan consists in performing two sweeps of the data.

- *up-sweep* – the tree is traversed from leaves to root computing partial sums (**reduction**).
- *down-sweep* – the tree is traversed from root to leaves. Each node of the current level passes its values to the element on the left, and the sum of the former and current value on the right.

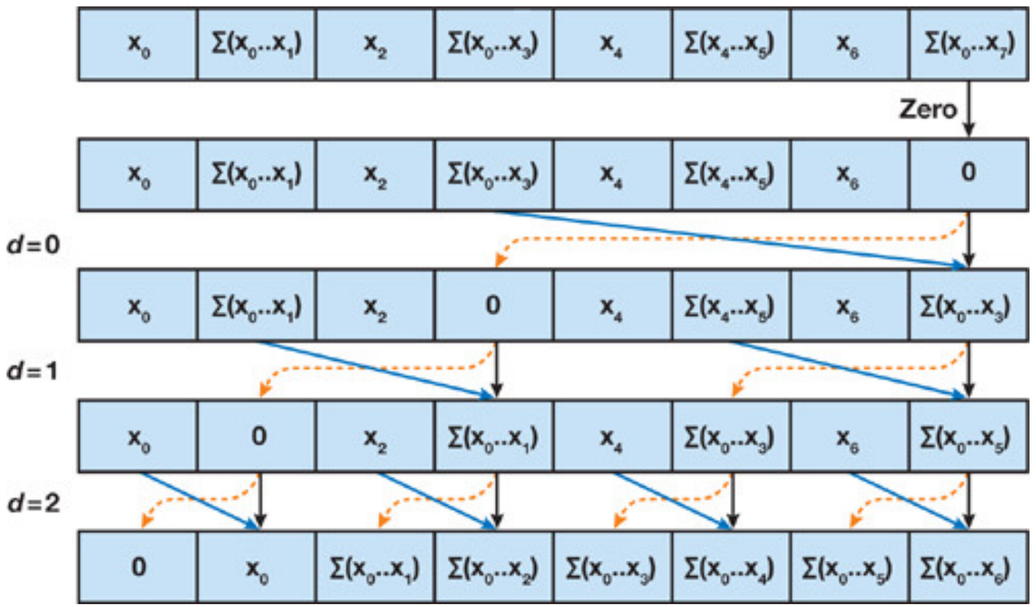

LHCb HLT1

The prefix sum is an essential tool of the LHCb HLT1 reconstruction.

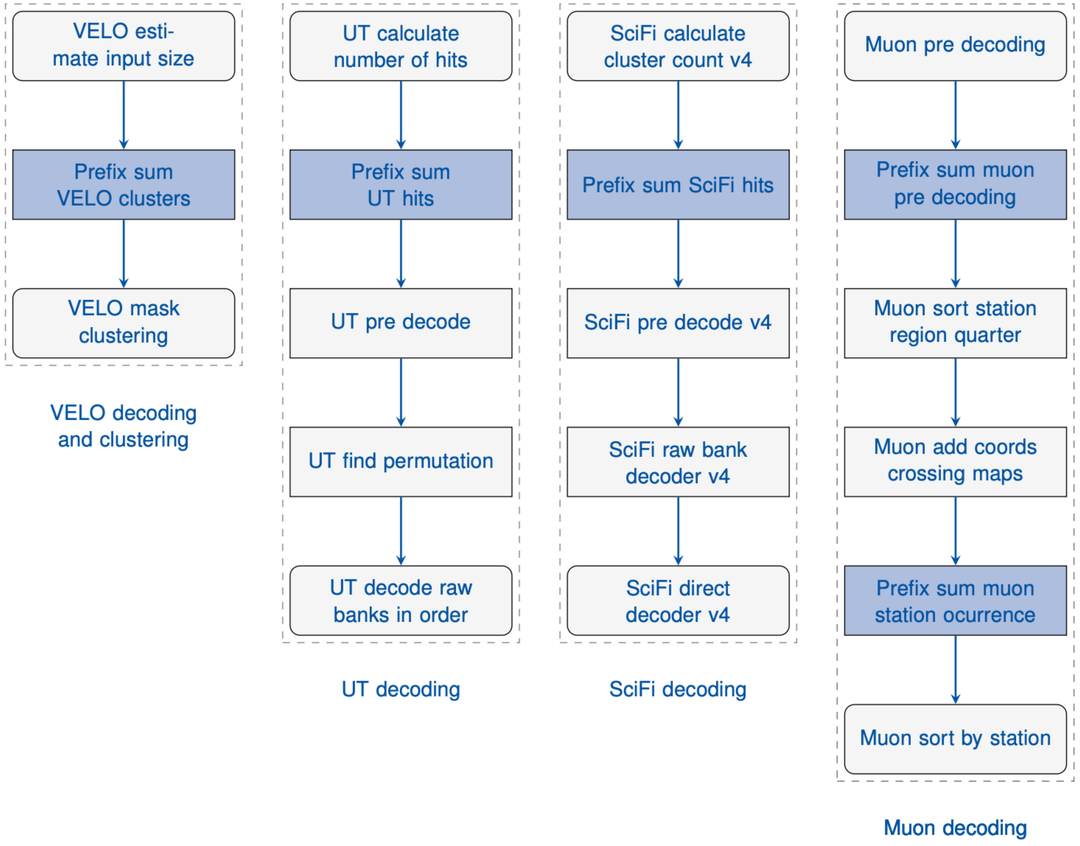

Table of Contents

Summary

- Precision affects performance, especially in GPUs.
- GPUs implement the IEEE754 standard, deviations are expected from compiler / architecture variability.
- Be mindful about register spilling.
- Prefer single-source kernels.

- Choose your standard wisely if targeting best performance.
- Consider middlewares if targeting best performance.
- It is possible to obtain portability with a standard, and to obtain performance with a middleware.

- Design patterns are a powerful high-level design tool.
- Know your patterns, design algorithms better.

Resources Used in the Talk

- GPU Teaching Kit on Accelerated Computing.
- Local Memory and Register Spilling by Paulius Micikevicius.
- Precision and Performance: Floating Point and IEEE 754 Compliance for NVIDIA GPUs, white paper.
- From sequential to parallel programming with patterns by Placido Fernandez .

Resources: Where to Go From Here

- If you're interested in AI, I suggest the course:
  - *TinyML and Efficient Deep Learning Computing* by Song Han

- To learn more about GPU Computing, join the NVIDIA dev program to get a free DLI course:
  - https://developer.nvidia.com/join-nvidia-developer-program

NVIDIA®