Data Oriented Design

Generic Programming

Johannes de Fine Licht

University of Copenhagen (DK)

Thematic CERN School of Computing 2014

The problem at hand

Navigation in detector geometry

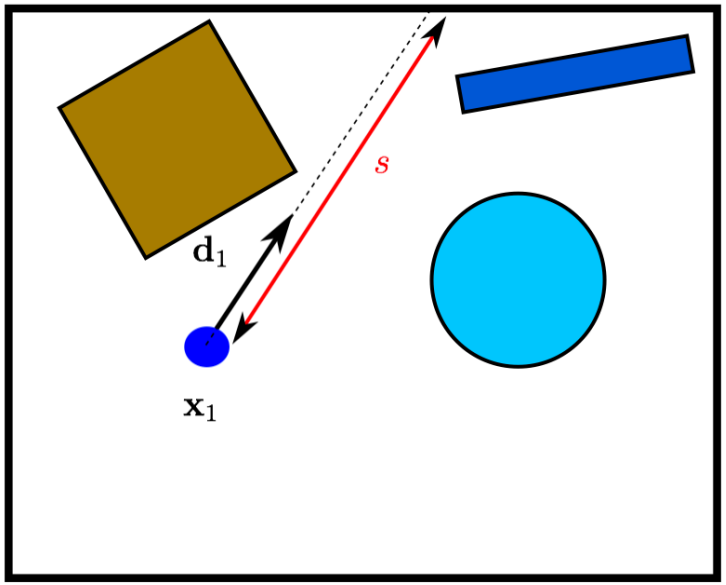

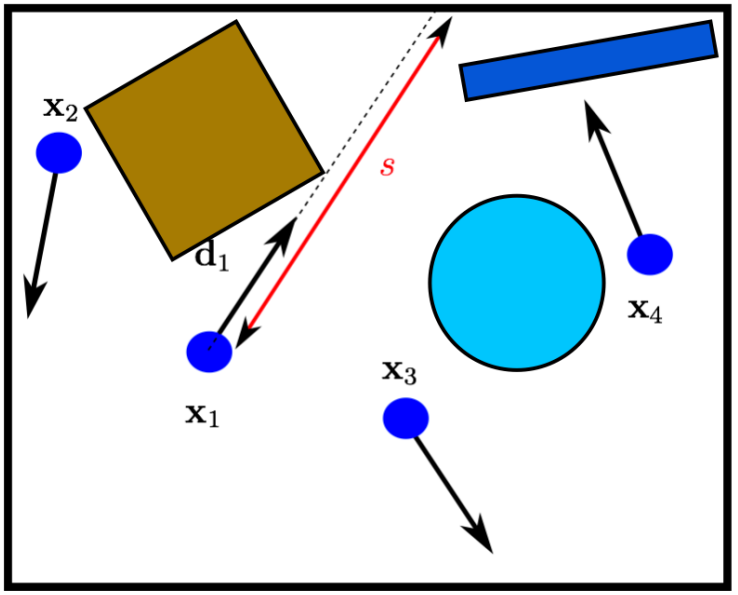

Geant4: single particles navigated in sequence.

GeantV: vectors of particles exploiting SIMD operations.

Borrowed from F. Carminati, "*Towards a high performance detector geometry library on CPU and GPU*", Annual Concurrency Forum 2014

Generic programming

- We want to exploit CPU **vector instructions** (SSE, AVX) and accelerators such as **GPUs** and the **Xeon Phi**
- Writing implementations for each backend would be an enormous effort, and is hopeless in terms of **maintainability**
- Our goal was to write **generic kernels**, that could exploit various **backends** (platforms, libraries, hardware)

Vc for wrapping intrinsics

http://code.compeng.uni-frankfurt.de/projects/vc/

- Vc is a project developed in Germany by Matthias Kretz
- The library wraps vector instruction intrinsics in easy-to-use classes, resulting in **portable** *explicitly vectorized* code

```
#include <Vc/Vc>
using namespace Vc;

double_v add(double_v const &a, double_v const &b) {
  // Performs 2 (SSE), 4 (AVX) or 8 (MIC) additions
  return a + b;
}
```

Abstracted kernels

- Write algorithms in a way where we can plug in the backend
- Compiles to CPU vector instruction intrinsics through Vc
- Compiles to CUDA kernel code executable on GPU

```
template <class Double_t, class Bool_t>
Bool_t BoxContains(const double dimensions[3],
                   const Double_t point[3]) {
  Bool_t contains[3];
  for (int i = 0; i < 3; ++i) {
    contains[i] = abs(point[i]) < dimensions[i];
  }
  return contains[0] && contains[1] && contains[2];
}
```

Maximizing hardware occupancy

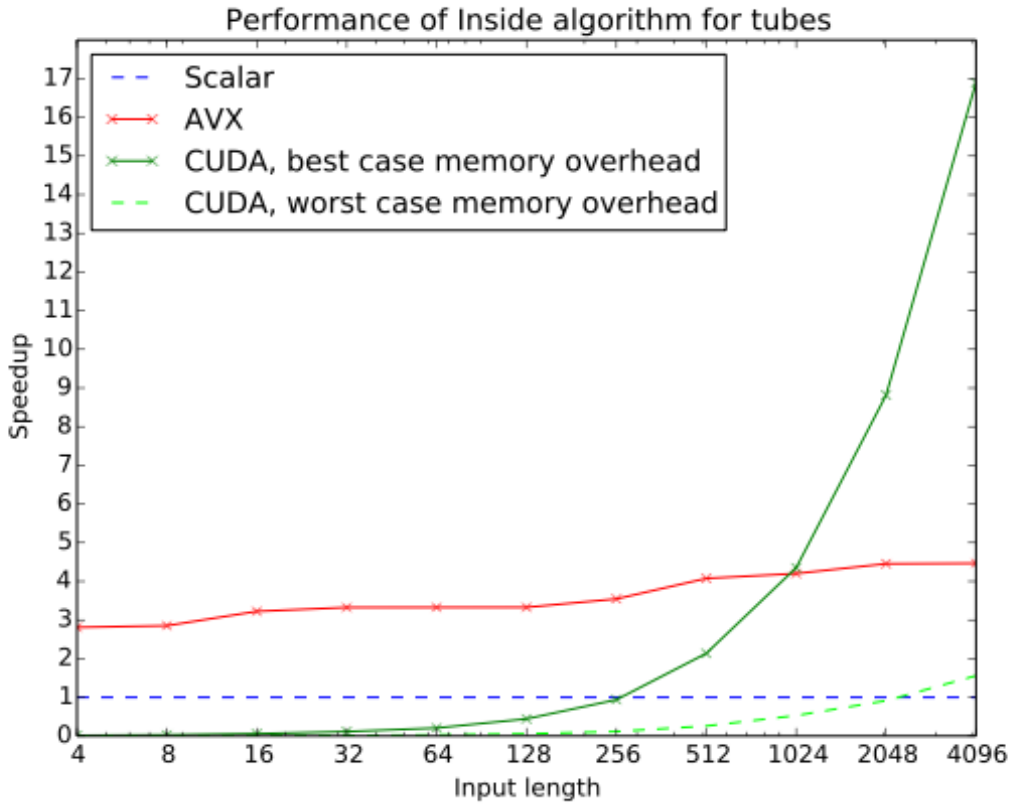

(Generic algorithm implemented by Georgios Bitzes.)

- **Single code base means we can develop algorithms and intrinsics separately**
- **With a good scheduler, we can dispatch work to the optimal processor at runtime**
- **We even gain from other effects; with AVX we exceed 4x speedup at high input size…**

Acknowledgements

- **Geometry team**
  - John Apostolakis (CERN/SFT)
  - Marilena Bandiermonte (CERN/SFT)
  - Georgios Bitzes (CERN/openlab)
  - Gabriele Cosmo (CERN/SFT)
  - Laurent Duhem (Intel)
  - Andrei Gheata (CERN/SFT)
  - Guilherme Lima (Fermilab)
  - Tatiana Nikitina (CERN/SFT)
  - Sandro Wenzel (CERN/SFT)
- **GeantV team at CERN (Federico Carminati) and Fermilab (Philippe Canal)**
- **tCSC 2014, Andrzej Nowak**