



GEANT4
A SIMULATION TOOLKIT



More on Multithreading

I. Semeniouk
LLR, CNRS – Ecole Polytechnique

Credits:

I. Hrivnacova(IJCLab), A. Dotti, M. Asai (SLAC)

Geant4 Tutorial,
23 - 27 November 2020, Orsay

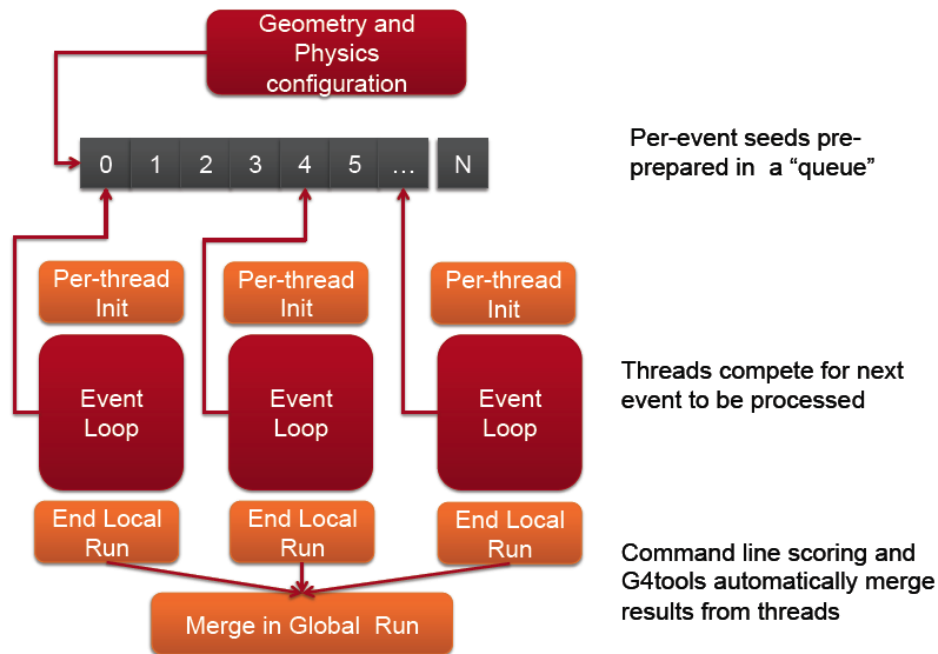
Outline

- Event data reduction
- Migration of a sequential Geant4 application to MT
- Results with Geant4 10.2 MT
 - Geant4 on MIC architecture
 - Scalability, memory reduction , reproducibility
- GPU and external frameworks

Event Data Reduction

Multi-threading in Geant4

- General design choice: event level parallelism via multi-threading (POSIX based, in 10.5 migration from POSIX threading to C++11 threading, since 10.7 Geant4 use PTL tasking library)
 - Each worker thread proceeds independently
 - Initializes its state from a master thread
 - Identifies its part of the work (events)
 - Generates hits in its own hits-collection
 - Geant4 automatically performs reductions (accumulation) when using scorers, G4Run derived classes or g4tools



Run Action

- In some users applications, [UserRunAction](#) is used to accumulate data from events and to calculate the result values for the whole run
 - E.g. in basic example B1, an energy deposited in a selected volume is accumulated event by event and a total dose is computed in the `EndOfRunAction()` method
- In multi-threading mode, the events are accumulated in [UserRunAction](#) objects instantiated **on workers** and the quantities accumulated on workers need to be merged in the [UserRunAction](#) **on master**
- This merging of the data accumulated on workers should be performed with use of [G4Run](#) or [G4Accummulable<T>](#) class

Run Action - Sequential

- An example of a run action used to accumulate data from events:
[MyRunAction](#) class
 - The run action class is the only action which is instantiated besides workers also on master

```
class RunAction : public G4UserRunAction
{
public:
    RunAction();
    virtual ~RunAction();

    virtual void BeginOfRunAction(const G4Run*);
    virtual void EndOfRunAction(const G4Run*);

    void AddEdep (G4double e)
    { fEdep += e; fEdep2 += e*e;};

private:
    G4double fEdep;
    G4double fEdep2;
}
```

sequential

Data accumulated
during event processing

Run Action + G4Run

- Separate data representing accounted data (if present) from your `RunAction` class in a new `Run` class (derived from `G4Run`)

```
class RunAction : public G4UserRunAction
{
public:
    RunAction();
    virtual ~RunAction();

    virtual G4Run* GenerateRun();
    virtual void BeginOfRunAction(const G4Event*);
    virtual void EndOfRunAction(const G4Event*);

    void AddEdep (G4double e);

private:
    Run* fRun;
}
```

MT (10.0)

```
class Run : public G4Run
{
public:
    Run();
    virtual ~Run();

    void AddEdep (G4double e)
    { fEdep += e; fEdep2 += e*e;};

    virtual void Merge(const G4Run*);

private:
    G4double fEdep;
    G4double fEdep2;
}
```

Run Action + G4Run (2)

- Implementation of new or changed functions:

```
G4Run* RunAction::GenerateRun()  
{  
    fRun = new Run();  
    return fRun;  
}  
  
void RunAction::AddEdep (G4double edep)  
{  
    fRun->AddEdep(edep);  
}
```

MT (10.0)

```
void Run::Merge(const G4Run* localRun)  
{  
    fEdep += localRun->fEdep;  
    fEdep2 += localRun->fEdep2;  
}
```

This function is called by the master run instance for each worker localRun instance

Data in master Run object

Data in worker Run object

See basic/B3b, B4b examples

Accumulables

- Classes for users “accumulables” management were added in 10.2 release
 - Accumulables are named variables registered to the accumulable manager, which provides the access to them by name and performs their merging in multi-threading mode
 - To better reflect the meaning of these objects, the classes base name "Parameter" used in 10.2 was changed in "Accumulable" in 10.3
- `G4Accumulabe<T>` - ready for use, for simple numeric types (double, int)
- Users can also define their own accumulables derived from the `G4VAccumulable` base class
 - Tested with `std::map<G4String, G4int>` used for processes counting in TestEm* examples

Accumulables (2)

- The accumulables are registered to [G4AccumulableManager](#)
 - Performs their merging in multi-threading mode according to their [MergeMode](#)
 - Provides the access to them by name
- Demonstrated in the basic examples B1 and B3a

Run Action + G4Accumulable

```
class RunAction : public G4UserRunAction
{
public:
    ...

    void AddEdep (G4double e)
    { fEdep += e; fEdep2 += e*e;};
```

sequential

MT (10.3)

```
private:
    G4double fEdep;
    G4double fEdep2;
}
```

Data accumulated
during event processing

```
#include "G4Accumulable.hh"
...
class RunAction : public G4UserRunAction
{
public:
    ...
    void AddEdep (G4double edep)
    { fEdep += e; fEdep2 += e*e;};
    // ...
private:
    G4Accumulable<G4double> fEdep;
    G4Accumulable<G4double> fEdep2;
};
```

Run Action + G4Accumulable (2)

```
#include "G4AccumulableManager.hh"
```

```
...
```

```
RunAction::RunAction()
```

```
: G4UserRunAction(),
```

```
  fEdep(0.),
```

```
  fEdep2(0.)
```

```
{
```

```
  //Register parameter to the parameter manager
```

```
  G4AccumulableManager* accManager = G4AccumulableManager::Instance();
```

```
  accManager->RegisterAccumulable(fEdep);
```

```
  accManager->RegisterAccumulable(fEdep2);
```

```
}
```

The accumulable are initialized with a name (optional) and a value

The accumulables not created via the manager have to be registered to it

```
void RunAction::EndOfRunAction(const G4Run* run) {
```

```
  ...
```

```
  // Merge parameters
```

```
  G4AccumulableManager* accManager = G4AccumulableManager::Instance();
```

```
  accManager->Merge();
```

```
  ...
```


```
}
```

The call to Merge() may be not necessary in future

Migrating Sequential Geant4 Application to MT

Migration to MT

Migration of a sequential application to MT is a 5-steps process:

1. Move user actions instantiation to new [G4UserActionInitialization](#) class
2. Use [G4RunManagerFactory](#) in your `main()` function (10.7.1)
~~Use [G4MTRunManager](#) in your `main()` function (old)~~
3. Split [DetectorConstruction::Construct\(\)](#) in two: SD and Field go in new method [ConstructSDandField\(\)](#)
-  4. Use [G4Run](#) to accumulate run data, implement [G4RunAction::Merge\(\)](#) method, or alternatively [G4Accumulables](#)
5. If you use anywhere [G4Allocator](#) (typically for hits), transform them to be [G4ThreadLocal](#)

More details can be found Geant4 documentation and a short “howto” in the TWiki migration page:

<https://twiki.cern.ch/twiki/bin/view/Geant4/QuickMigrationGuideForGeant4V10>

Moving from threads to tasks

- Introduced Geant4 10.7
 - Adds a new ‘task-oriented’ capability to adapt to frameworks of LHC experiments which are task oriented
 - Includes a native C++ implementation of the ‘task model’
 - Includes an (Intel) Thread Building Block ‘TBB task mode’ – the Geant4 installation must find & use TBB, (by configuring cmake with `-DGEANT4_USE_TBB=ON`)
- There is now a variety of RunManagers
 - Sequential (`G4RunManager`)
 - ‘Old-style’ Multi-threading (`G4MTRunManager`)
 - `G4TaskRunManager` in ‘native’ mode
 - `G4TaskRunManager` in TBB mode
- new class `G4RunManagerFactory` can be used to create any of these.
- Final version of tasking will be provided in 10.7 release, the default RunManager since 11.0

Moving from threads to tasks

G4RunManagerFactory

- The class G4RunManagerFactory can be used to create RunManager of any type
- The default RunManager
 - G4MTRunManager (before 10.0)
 - G4TaskManager (since 11.0)

- The ways to select RunManager Type

```
// [Option #1] enum class G4RunManagerType: // Default, Serial, MT, Tasking, TBB
```

```
auto* runMgr =
```

```
G4RunManagerFactory::CreateRunManager( G4RunManagerType::Default, 4);
```

```
// [Option #2] string: "default", "serial", "mt", "task", "tbb"
```

```
auto* runMgr = G4RunManagerFactory::CreateRunManager( "default", 4);
```

```
// [Option #3] Environment VARIABLES
```

```
export G4FORCE_RUN_MANAGER_TYPE=Serial | MT | Task | TBB
```

```
export G4FORCENUMBEROFTHREADS=4
```


Results With Geant4 10.x MT

Reproducibility

- Geant4 Version ≥ 10.0 guarantees strong reproducibility
- Given a setup and the random number engine status it is possible to reproduce any given event independently of the number of threads or the order in which events are processed
- Note: (optional) radioactive decay module breaks this in MT, Geant4 MT experts are currently working on a fix
- This does not mean the results are wrong!
- Simulation results are equivalent between Sequential and MT

MIC Architecture

- Geant4 has been ported to compile and run on Intel Xeon Phi (aka MIC)
- It requires Intel compiler (not free) and RTE
- **61 cores (x4 ways hyper-threading), w/ max 16GB of RAM**

The screenshot displays the Intel Xeon Phi Coprocessor Platform Status Panel and a Terminal window. The status panel shows system utilization, core temperature, memory usage, and power consumption. The terminal window shows the execution of Geant4 commands and the output of the simulation.

Intel Xeon Phi Coprocessor Platform Status Panel

Cards: 4 Memory: 64 GB Cores: 244

Utilization View (All Devices)

System [%] 46
User [%]

Average Core Utilization: 00:52:43 - 00:53:03 - 00:53:23

Average Core Temperature °C: 46

Total Memory Usage: 1.123 GB

Total Power Usage: 433 Watts

mic2: Core Histogram View

Individual Core Usage

System [%] 61
User [%]

Terminal


```
File Edit View Terminal Tags Help
FinalRange(m)= 0.01, dRoverRange= 0.1, Integral: 1, fluct: 1, linkLossLimit= 0.02
Stopping Power data for 17 ion/material pairs
===== EM models for the G4Region DefaultRegionForTheWorld =====
BraggIon : Emin= 0 eV Emax= 2 MeV
BetheBloch : Emin= 2 MeV Emax= 10 TeV


mic: for alpha SubType= 10
RangeFactor= 0.2, stepLimitType: 0, latDisplacement: 0
===== EM models for the G4Region DefaultRegionForTheWorld =====
UrbanMsc : Emin= 0 eV Emax= 10 TeV Table with 10 bins
ionIoni: for alpha SubType= 2
dE/dx and range tables from 100 eV to 10 TeV in 77 bins
Lambda tables from threshold to 10 TeV in 77 bins, spline: 1
finalRange(m)= 0.01, dRoverRange= 0.1, Integral: 1, fluct: 1, linkLossLimit= 0.02
===== EM models for the G4Region DefaultRegionForTheWorld =====
BraggIon : Emin= 0 eV Emax= 7.9452 MeV
BetheBloch : Emin= 7.9452 MeV Emax= 10 TeV

mic: for anti_proton SubType= 10
RangeFactor= 0.2, step limit type: 0, lateralDisplacement: 1, polarAngleLimit(deg)= 100
===== EM models for the G4Region DefaultRegionForTheWorld =====
WentzelVIUnit : Emin= 0 eV Emax= 10 TeV Table with 77 bins Emin= 100 eV Emax= 10 TeV
```

Memory consumption on Intel Xeon Phi

CMS geometry (GDML), π^- 50 GeV (FTFP_BERT), B field (4T) - Xeon Phi 3120A







Intel Xeon Phi™ 3120A

Throughput in sequential mode

CPU performances (50 GeV Pb/LAr)



Throughput in Sequential Mode



GPU and External Frameworks

Heterogeneous Parallelism

- MPI

- MPI works together with MT
- The examples of MPI parallelism with Geant4 MT are provided in Geant4 [examples/extended/parallel/MPI](#)
- New features in this category expected in the future: Geant4 MT experts are currently evaluating extensions !

- TBB

- Intel Thread Building Block (TBB): task based parallelism framework
 - <https://www.intel.com/content/www/us/en/developer/tools/oneapi/onetbb.html>
- Freely available for Linux/Mac/WIN
- Expression of interest by some LHC experiment
- One example is provided in Geant4 [examples/extended/parallel/TBB](#)
- Since 10.7 available as special run manager via PTL

G4Opticks

G4Opticks (part of Opticks): interfaces Geant4 user code with Opticks.

It defines a hybrid workflow where generation and tracing of optical photons is offloaded to Opticks (GPU/device) at stepping level when a certain amount of photons is reached. Geant4 (CPU/host) handles all other particles.

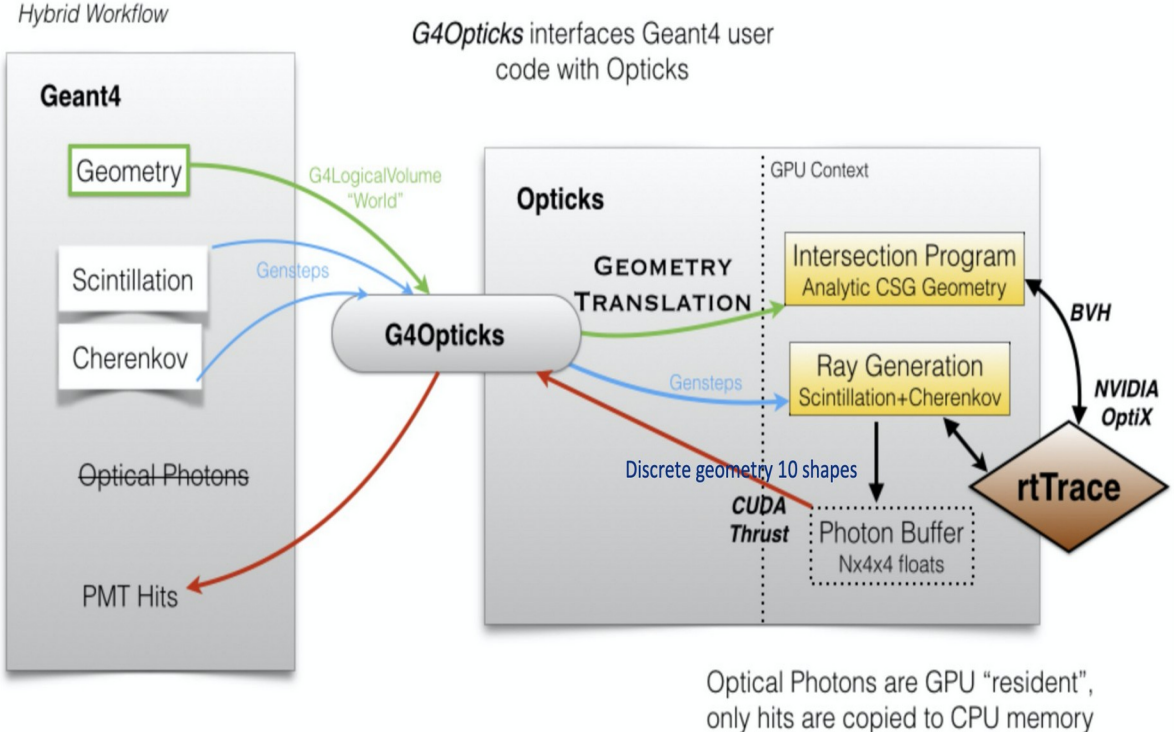
The Geant4 Cerenkov and Scintillation (C/S) processes are only used to calculate the number of optical photons to be generated at a given step and to provide all necessary quantities to generate the photons on the GPU.

The information collected is the so called GenStep which is different for Cerenkov and Scintillation (C/S).

Photon Hits are collected at the end of the G4Opticks call and added to the event hits collection.

Use NVIDIA® hardware (some with RTX: raytracing hardware acceleration) and software (CUDA, OptiX).

Figure from Simon's presentation



An advanced Geant4 example: CaTS Calorimeter and Tracker Simulation






Performance:

Hardware:

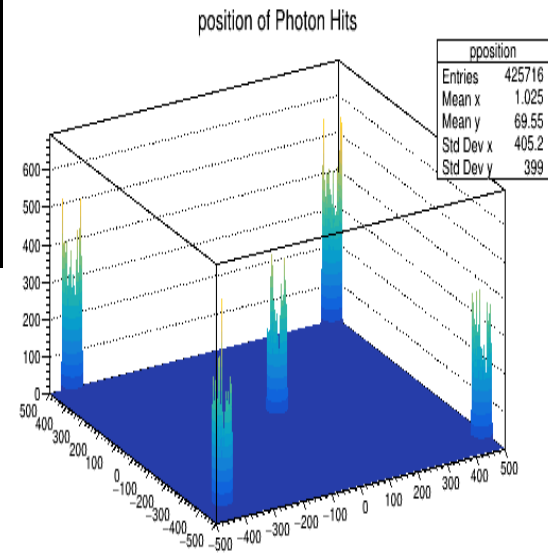
CP U	Intel(R) Core i7-9700K 3.6GHz 32 GB memory.
GP U	GeForce RTX 2070 CUDA Driver Version /11.3 CUDA Capability: 7.5 VRAM: 7981 Mbytes Cores: 2304



Integration of Opticks and Geant4 (CaTS)

From Hans Wenzel presentation
26TH Geant4 Collaboration Meeting

September 16th 2021



Timing results (Geant4 10.7.p01):

Geant4 optical physics	2438 sec/event
G4Opticks, RNGmax ¹ 10	6.45 sec/event
G4Opticks RTX enabled, RNGmax ¹ 10	2.72 sec/event
G4Opticks, RNGmax ¹ 100	6.86 sec/event
G4Opticks RTX enabled, RNGmax ¹ 100	2.87 sec/event

1) Memory pre allocated for pre-initialized (at installation) curandState files to load.


**Geant4/(Geant4 + Opticks) comparison:
2438/6.45 = 378 (x 2.4 ~ 900 with RTX) x speed up**

RTX Ray tracing hardware acceleration is usually not available on HPC platforms

AdePT

Accelerated demonstrator of electromagnetic Particle Transport

- Open Source <https://github.com/apt-sim/AdePT>
- External physics: G4HepEm and geometry: VecGeom
- Previously validated simulation results on GPU against Geant4



Andri Gheata