

# LHCb Recent Efforts on Compute and Software Optimization

---

Stefan Roiser

WLCG Workshop, Manchester

20 June '17




# Content

- “Fast Stop”
- Code Vectorization
- How to involve the collaboration for the upgrade

# Fast Stop – Motivation

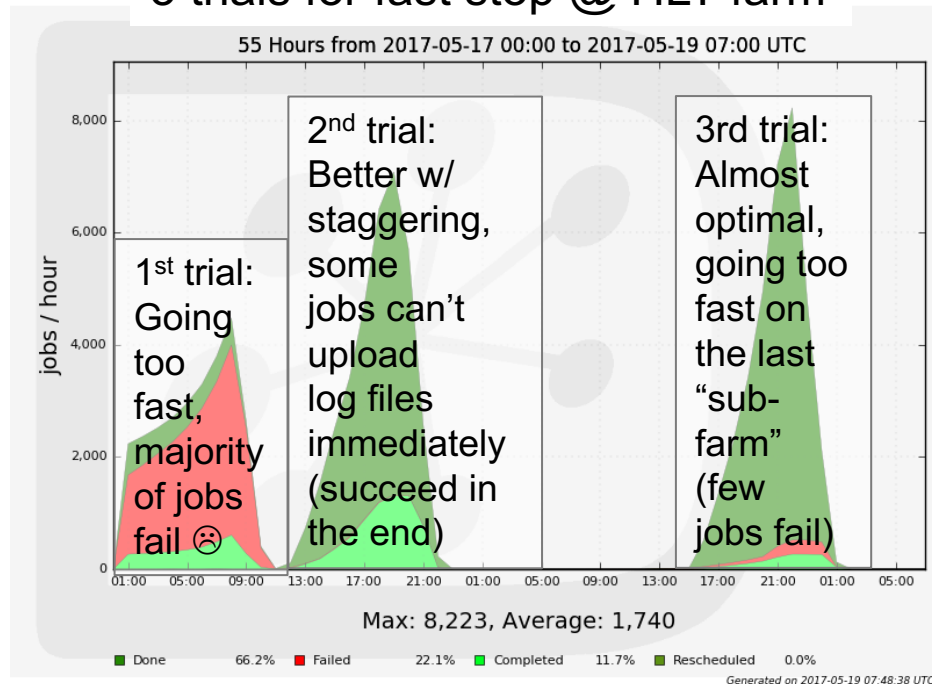
- E.g. draining of the HLT farm 5 – 6 February
  - Minimum 24 hours draining of the resource
  - Lot of more work possible but not available for payload execution



# Fast Stop – Implementation

- Send user defined signal to (Simulation) job which will tear down gracefully after the processing of the current event. Takes 1 – 5 minutes.
- NB: Limited by local network and storage bandwidth. In case of HLT limited, therefore needed to stagger the stoppage
  - Pull out from 40k running jobs within 2 hours

## 3 trials for fast stop @ HLT farm

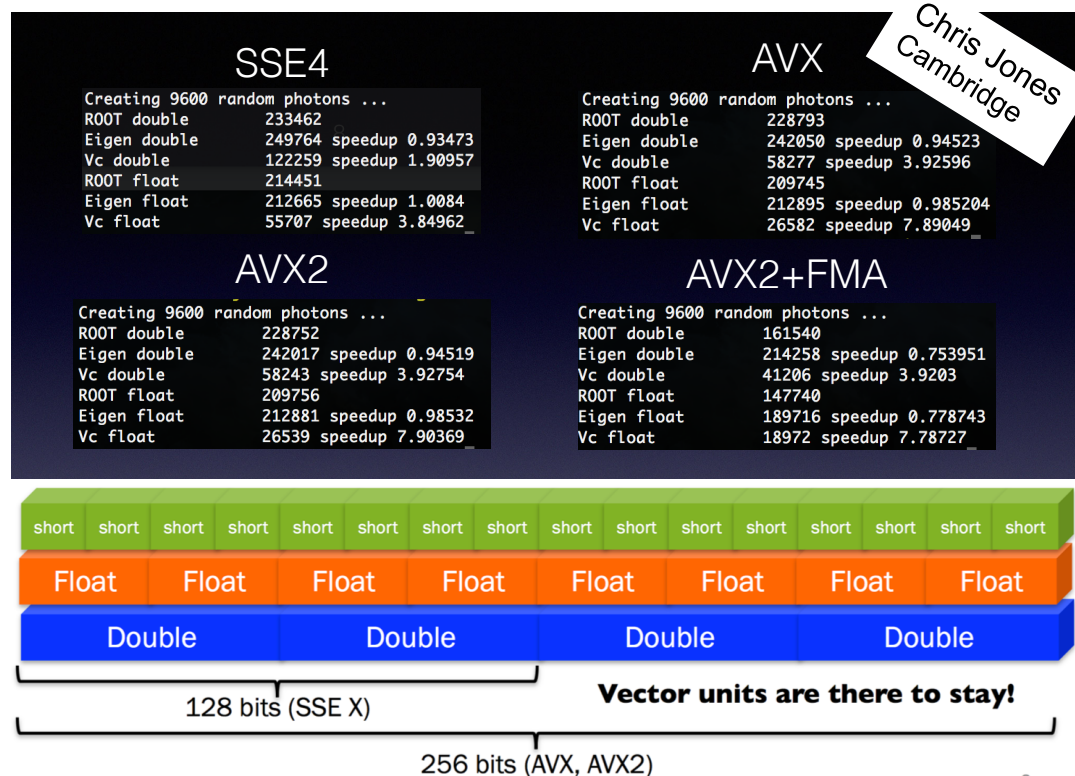


# Fast Stop – Use Cases

- Fast draining of Sites
  - stop workflows on WN, rack, row, whole site. Use the site to produce data until the "last minute" before a downtime.
- Sharing resources e.g. HLT farm, CERN disk servers, etc
  - allow concurrent usage of offline (simulation) and other (e.g. trigger) workflows. In case more resources are needed → stop offline simulation at short notice
- Next step:
  - Include into machine job features mechanism. Set machine/job feature "shutdown\_time" → jobs will tear down gracefully specified time

# Code Vectorization – Motivation

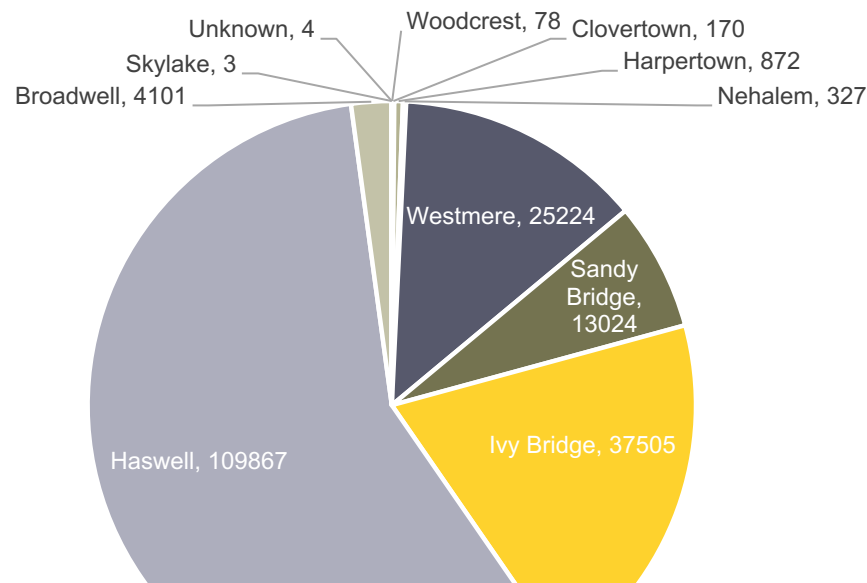
- LHCb needs to increase the software performance by several factors for Run3 (2021)
- One optimization dimension is to exploit more SIMD instructions
- E.g. RICH photon reconstruction shows perfect speedup, e.g. factor 8 on AVX2/float



# How old is the grid?

- CPU architecture of grid worker nodes LHCb jobs landed on in May
- 2017 LHCb software stack built for SSE3 and SSE4.2
  - SSE4.2 build cannot run on 1120 “older than Nehalem” worker nodes
  - Solution: Dirac executes “correct” build depending on WN capability
- Plan to have minimum requirement on SSE4.2 for builds in 2018
  - Need “late matching” of payloads

Architectures of Individual Worker Nodes, LHCb Jobs, May '17



<i>Tic/Toc</i>	<i>Vectorization</i>	<i>Code Name</i>	<i>Release Date</i>
Toc	SSE3	Woodcrest, Clovertown	Jun.06
Tic		Harpertown	Nov.07
Toc	SSE4.2	Nehalem	Nov.08
Tic		Westmere	Jan.10
Toc	AVX	Sandy Bridge	Jan.11
Tic		Ivy Bridge	Apr.12
Toc	AVX2	Haswell	Jun.13
Tic		Broadwell	Sep.14
Toc	AVX512	Skylake	Aug.15

# How old is the grid?

- CPU architecture of grid worker nodes LHCb jobs landed on in May
- 2017 LHCb software stack built for SSE3 and SSE4.2


- SSE4.2 build cannot run on 1120 “older than Nehalem” work

- Solution: Dirac execut build depending

- Plan to have requirement on SSE4.2 in 2018

- Need “late matching” of payloads

Architectures of Individual Worker Nodes, LHCb Jobs, May '17



**IMPORTANT: Sites providing WNs through virtualization need to export the CPU capabilities properly onto the VM**

Tic/Toc	Vectorization	Code Name	Release Date
Toc	SSE3	Woodcrest, Clovertown	Jun.06
Tic		Harpertown	Nov.07
Toc	SSE4.2	Nehalem	Nov.08
Tic		Westmere	Jan.10
Toc	AVX	Sandy Bridge	Jan.11
Tic		Ivy Bridge	Apr.12
Toc	AVX2	Haswell	Jun.13
Tic		Broadwell	Sep.14
Toc	AVX512	Skylake	Aug.15



# Getting the Whole Experiment on Board

- The LHCb Computing Project will provide the foundation for the upgrade
  - Task-parallel, multi-threaded Gaudi framework
  - New implementations for conditions data, detector description, event model
- ... but the majority of code is being converted by physicists working in the sub-systems

Training of the collaboration is essential for the success of the Run 3 upgrade program !!!!!!!!!!!!!!!



Bi-monthly training and hackathon weeks organized by the Computing Project



# E.g. 7<sup>th</sup> Software Upgrade Hackathon, 19 – 23 June



## 7th hackathon of software for the upgrade

19 Jun 2017, 14:00 → 23 Jun 2017, 13:15 Europe/Zurich

3862-1-001 - Till Moritz Karbach Room (CERN)

Concezio Bozzi (CERN and INFN Ferrara), Sebastian Ponce (CERN), Stefan Roiser (CERN)

**Description** The 7th hackathon of software for the LHCb upgrade will be held during the week of June 19th, at point 8, in the Moritz Karbach conference room just above the LHCb control room.

For logistic purposes, participants are required to register to this event.

The first session on monday afternoon is devoted to talks and tutorials, tentatively on git and performance measurements. We can also re-play presentations and tutorials on other topics, depending on specific requests. Participants are encouraged to make their choices when registering. A video connection will be available for this session.

The hackathon starts on tuesday and continues on wednesday and thursday, with summary of the work done at the end of each day. A wrap-up session is scheduled on friday morning.

The topics to be hacked will depend on the choices made by the participants upon registration.

Possible examples in the domain of vectorization and its impact on code speedup:

- measure degree of vectorization/autovectorization of functional algorithms, with modified callgrind and/or intel tools
- try to vectorize the main ones and measure improvements
- work on using autovectorization in the common geometry code
- convert to functional and improve more algorithms

Other items:

- compare logs in Gaudi Hive
- Parallel I/O with root
- automation of performance measurements
- conditions
- new Gaudi DataHandles
- multi-events algorithms

Videoconference Rooms

Upgrade\_hackathon Join Till Moritz Karbach Room

Registration

Participants Register

MONDAY, 19 JUNE

14:00 → 14:15 Introduction / setup 15m

14:15 → 19:00 Hackathon: Tutorials

14:15 **Tutorial -- the GIT tool and best practices.** 30m

Speaker: Sebastian Ponce (CERN)

gitintro.pdf

14:45 **Tutorial -- Tools for measuring code performance** 15m

Speaker: Monir Hadji (Universite Claude Bernard-Lyon I (FR))

measure-perf-hacka...

16:00 **Coffee / tea** 30m

16:30 **emacs and vim extensions for the new framework** 15m

Speaker: Adam Davis (Tsinghua University (CN))

emacs-skeleton.pdf

16:45 **Tutorial -- the new Gaudi functional framework** 15m

Speaker: Sebastian Ponce (CERN)

LHCbFutureHandsO...

18:00 **Hackathon setup** 15m

HackathonIntro.pdf

C++ course @  
<http://cern.ch/sponce/C++Course>