ATLAS Simulation Profiling

Zach Marshall & John Apostolakis
for ATLAS Simulation team

23 September 2013

Overview and Reminders

- Last detailed benchmarking of the simulation was ~Dec 2009
  - Report is at https://cds.cern.ch/record/1247601
- Details of the 'standard candle' – the type of simulation run
  - All on 50 ttbar events with FTFP_BERT (what we will be running)
  - Full Geant4 simulation – not Integrated Simulation Framework (ISF)
  - Geant4 9.6 patch02
- Types of Simulation benchmarking
- Results
- Next steps

Types of benchmarking

- Simulation benchmarking comes in at least four flavors:
  - CPU profiling per function – first today
  - PHYSICS profiling – second today
  - Memory allocation per function – soon
  - Memory alloc/free per function – soon
- That last one is very important and often forgotten
  - The biggest performance gains and losses we got in ATLAS came from changes to the physics of the simulation!
- Platform: gcc 4.7, 64-bit, SLC6
  - (should be the "fastest" production setup we have)
- We have a new fancy simulation now, so let's use it!

What a profile looks like

The Tower of .. cachegrind

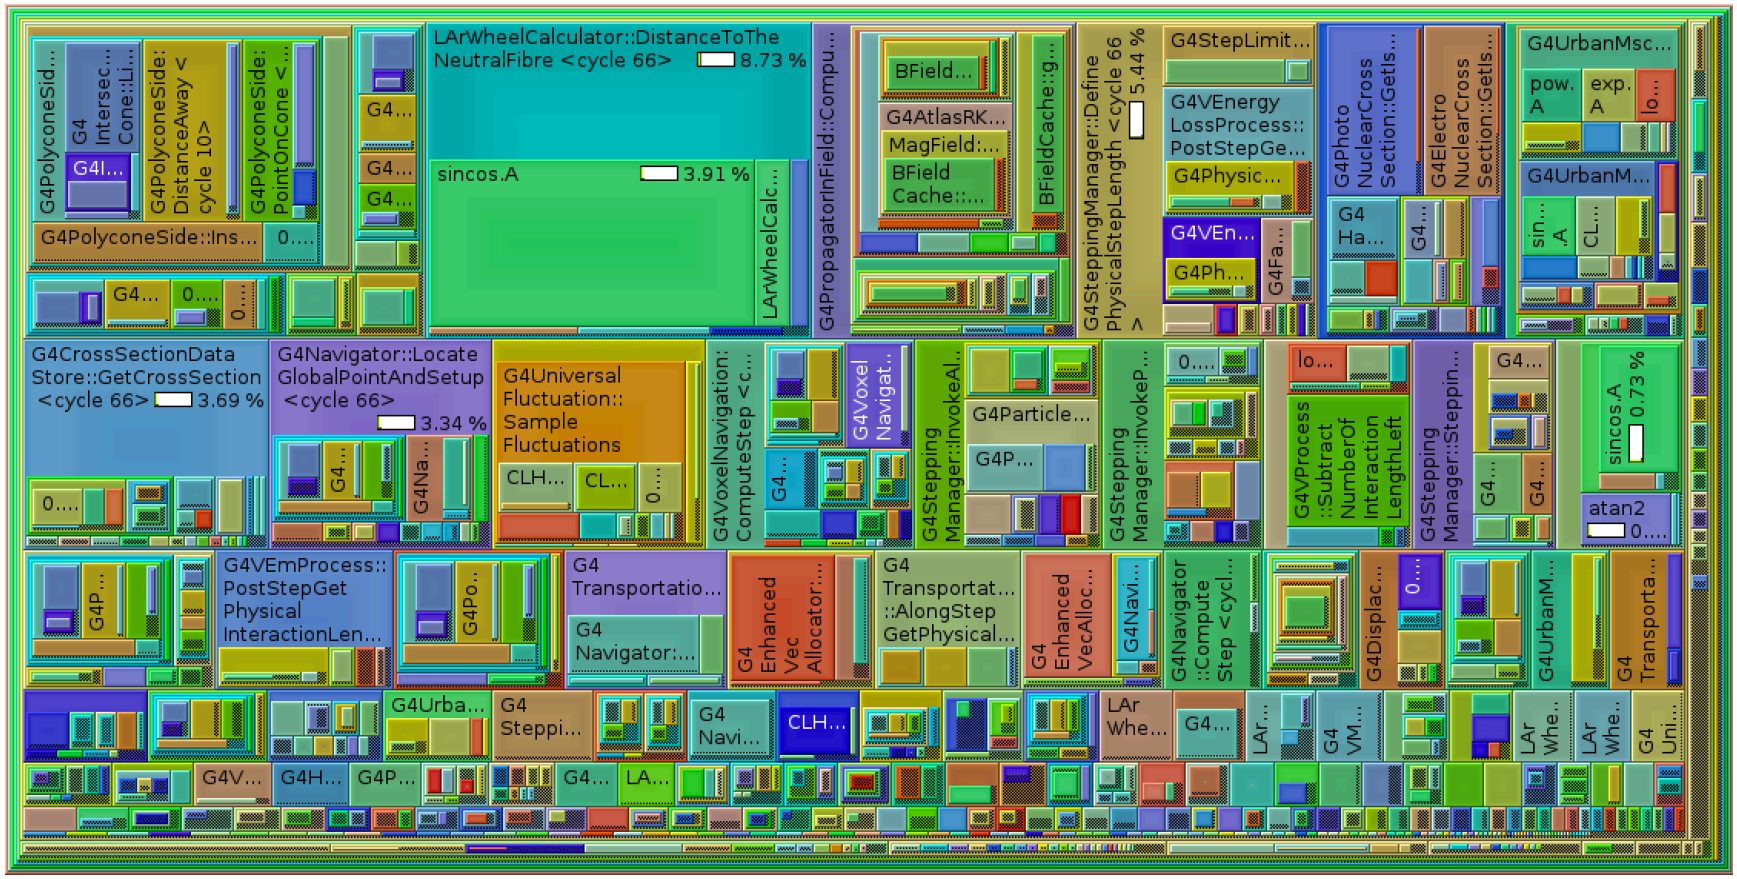

Second, a new Profile

B-field is under 2% of the total sim time – was ~15% !

Extensive effort – and great result.

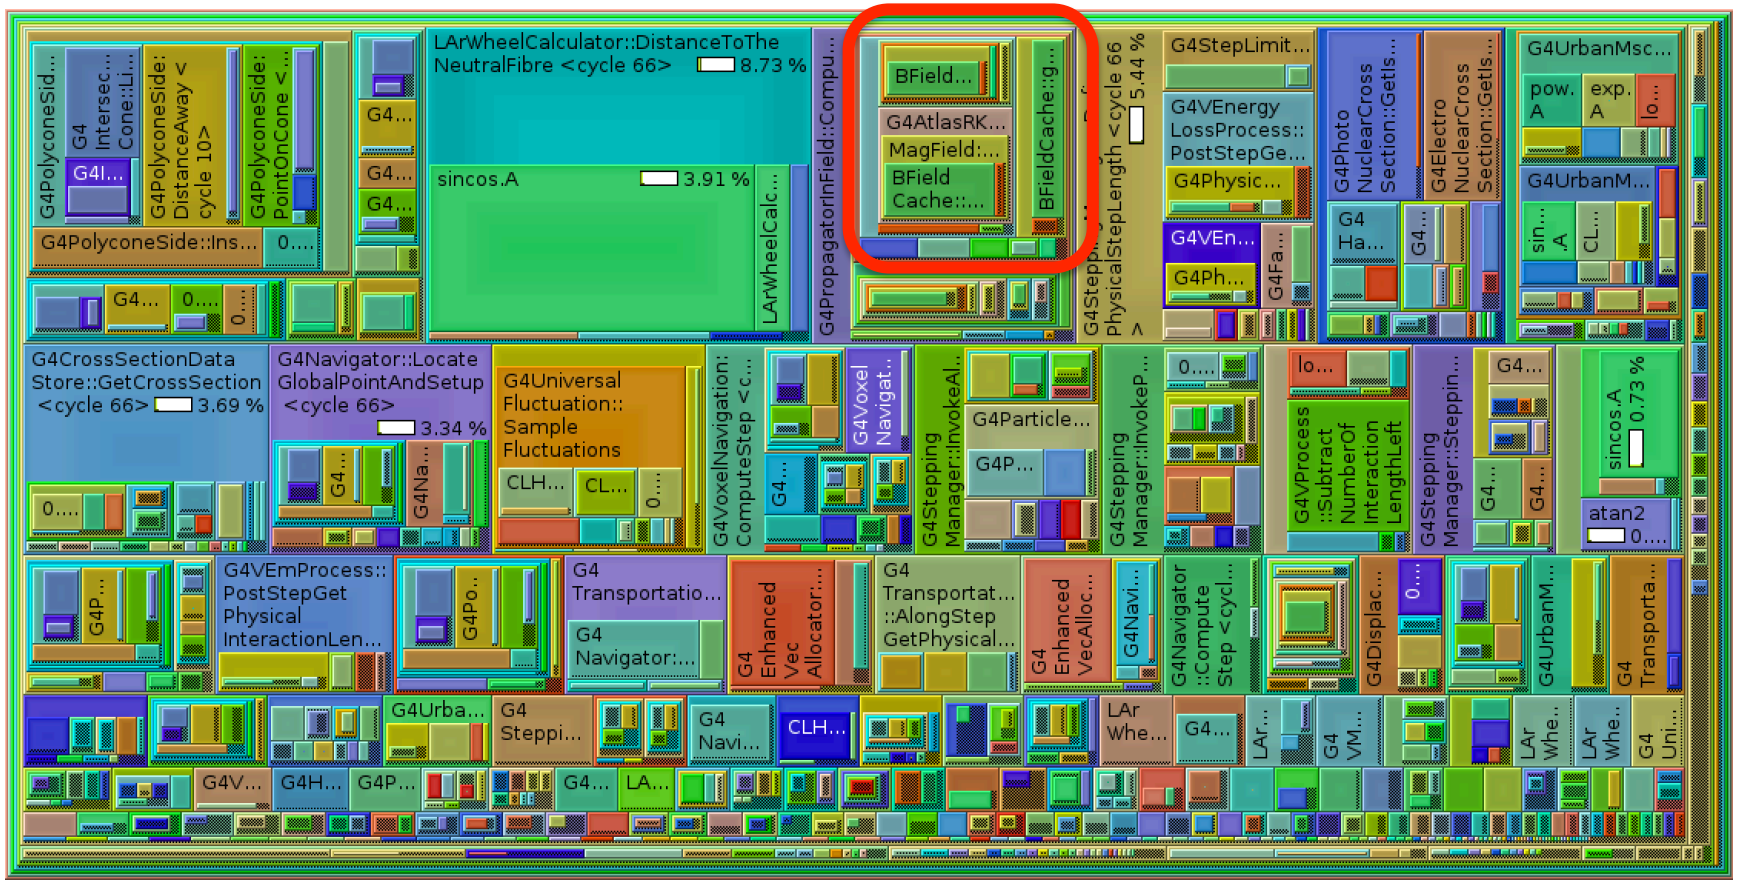

Second, a new Profile

Both the best Stepper and a 2nd best are both used .. ?

That's an (user) error.

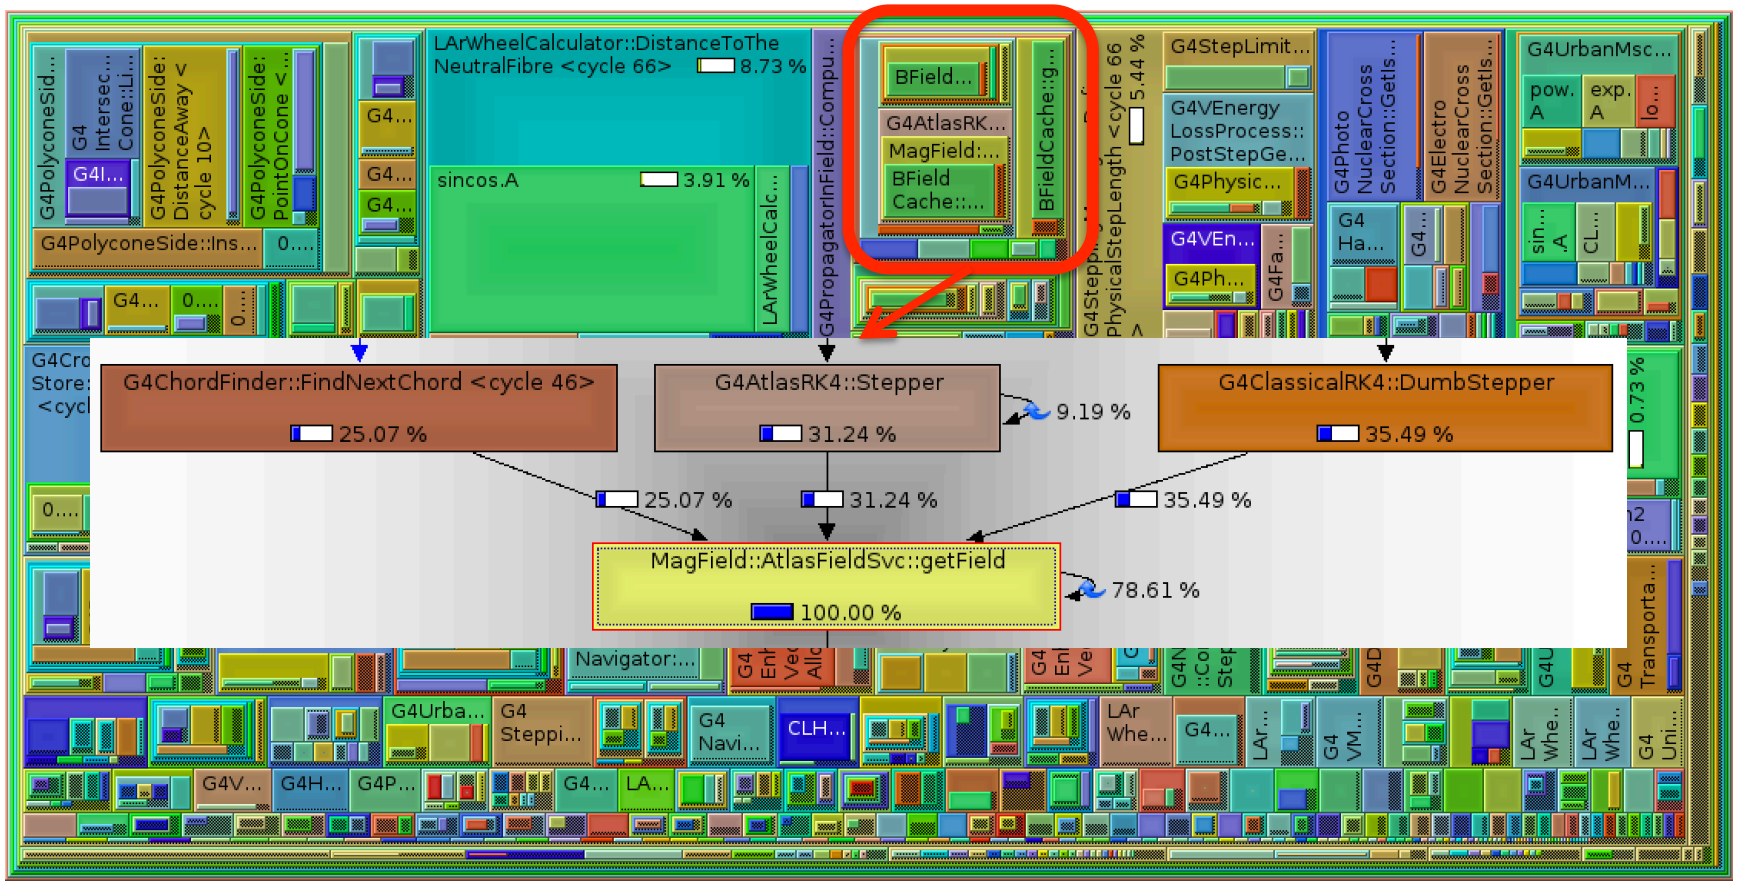

Second, a new Profile

A costly special ATLAS solid

This is a all the EMEC (>20% in a few big methods)

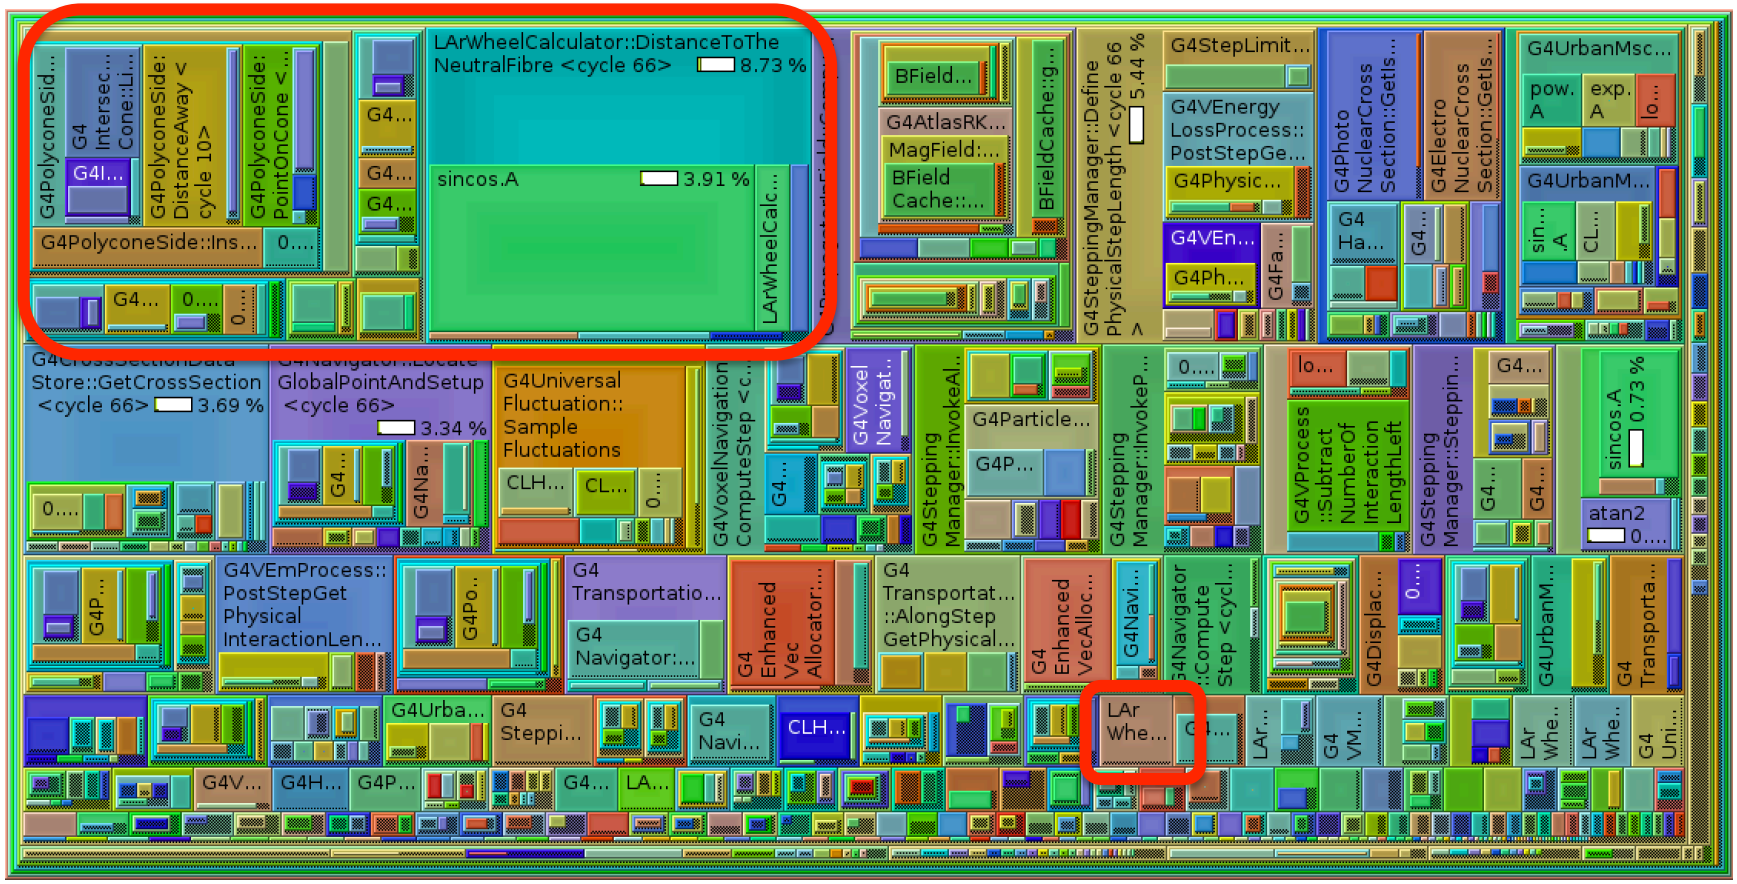

Second, a new Profile

The rest is almost all G4! UserActions are negligible

Even if G4 were perfect, ATLAS is setting the parameters (cuts,..)

Is it possible to find the 'minimum' time for Geant4 to be 'this accurate' ?

- What couldbe done faster or skipped (in G4) ?
- What could not requested by ATLAS ?

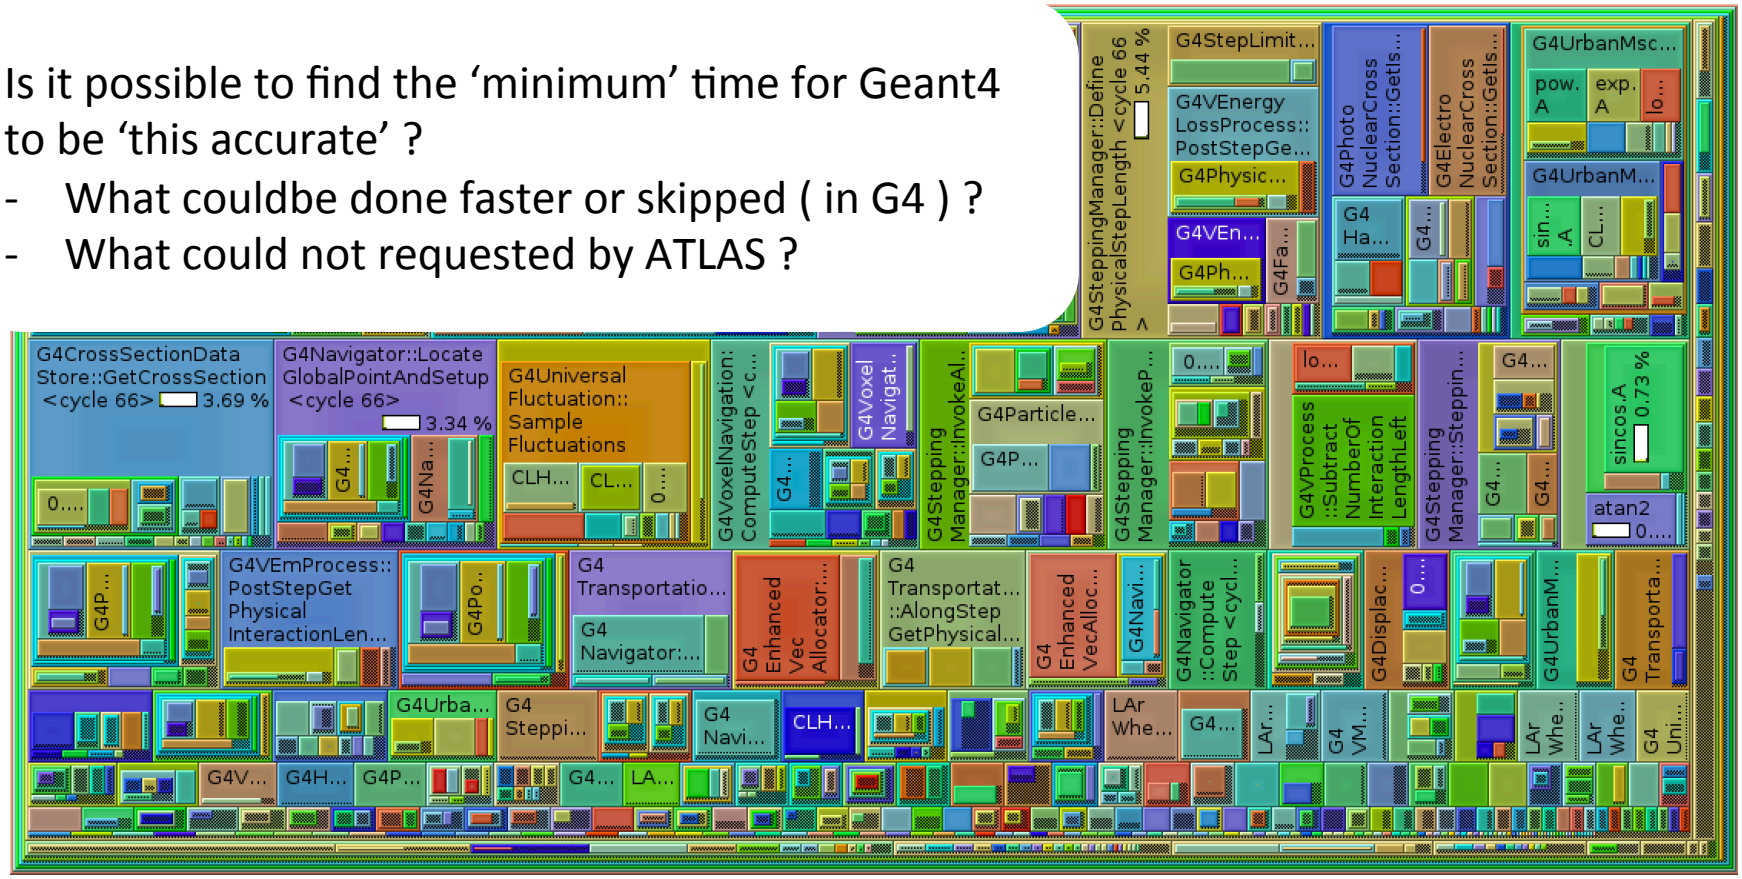

Second, a new Profile

On the other hand: **This is almost all G4!!**

Still, there are some **oddities** in here… Will be looking into them for a while.

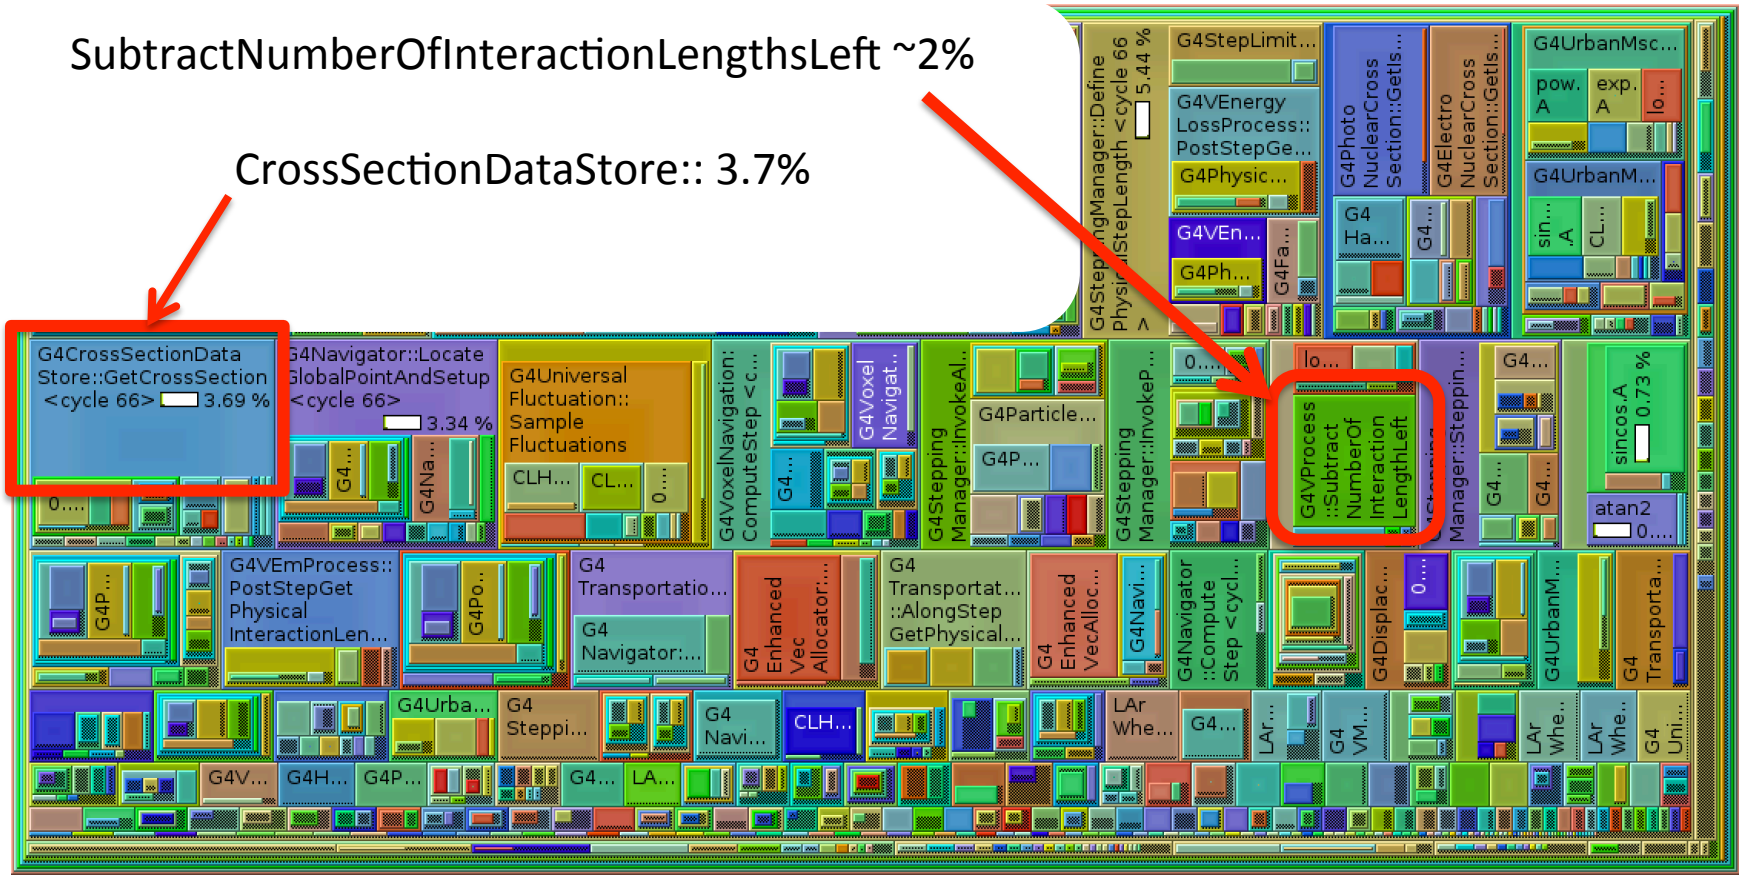

Physics Profiling: Really Fun Plots

- This kind of profiling consists of taking all the steps in the simulation and classifying them:
  - By location, by energy, by particle type, by process…
- And then looking at them and trying to find CRAZY STUFF
  - It's really a lot of fun – like a Rorschach test for physicists…
- Note: there are a lot of steps
  - That means I directly make histograms – so correlations are hard until I know exactly what I'm looking for

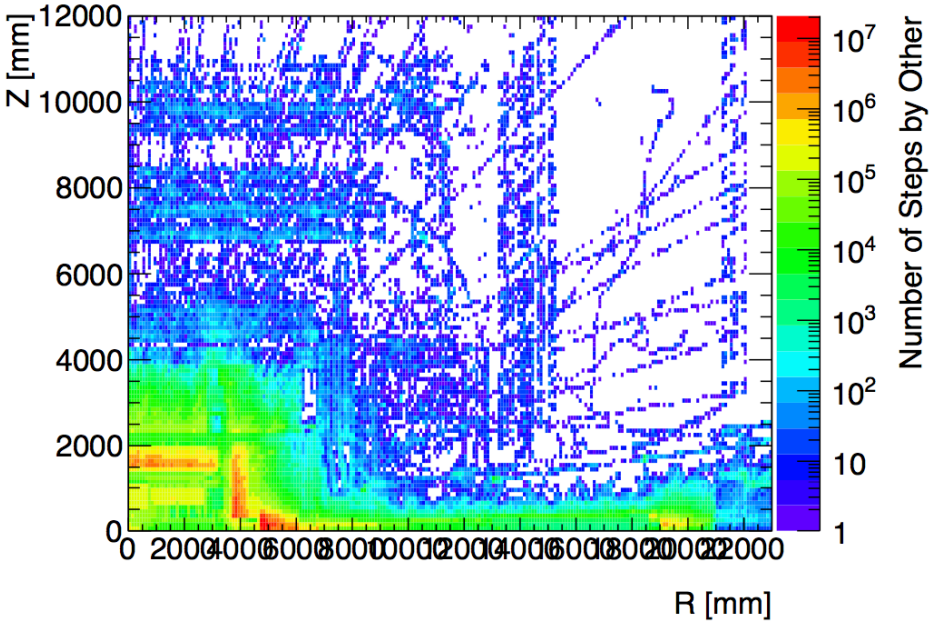

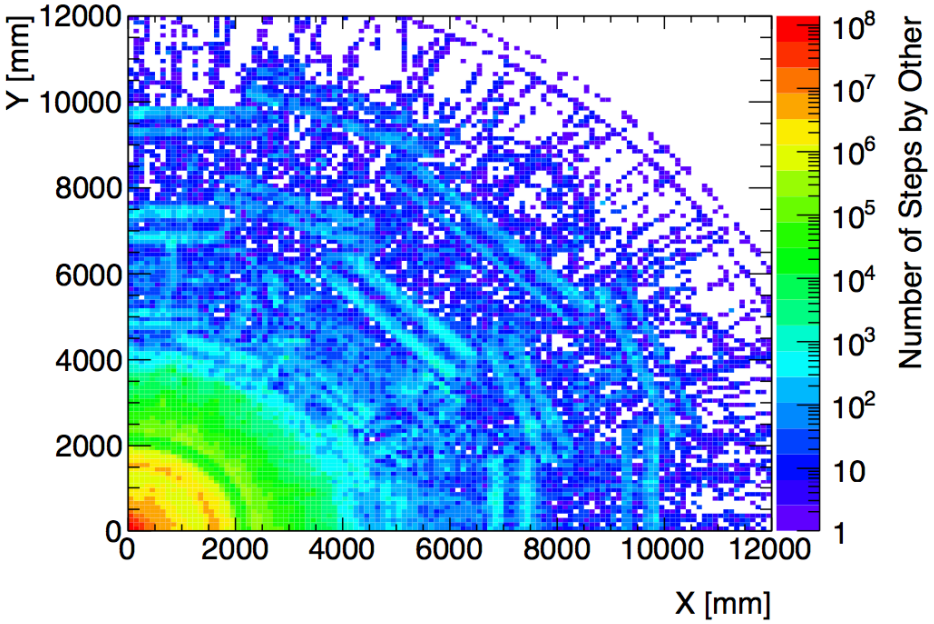

Bug: Low Energy Photons

- In hadronic (non-decay) interactions, a number of quite low energy (~eV) photons are produced
- No physics interactions are enabled for these photons, so they fly out of the detector
- Flying out of the detector takes them quite some time, so we waste *hundreds of thousands of steps* on them

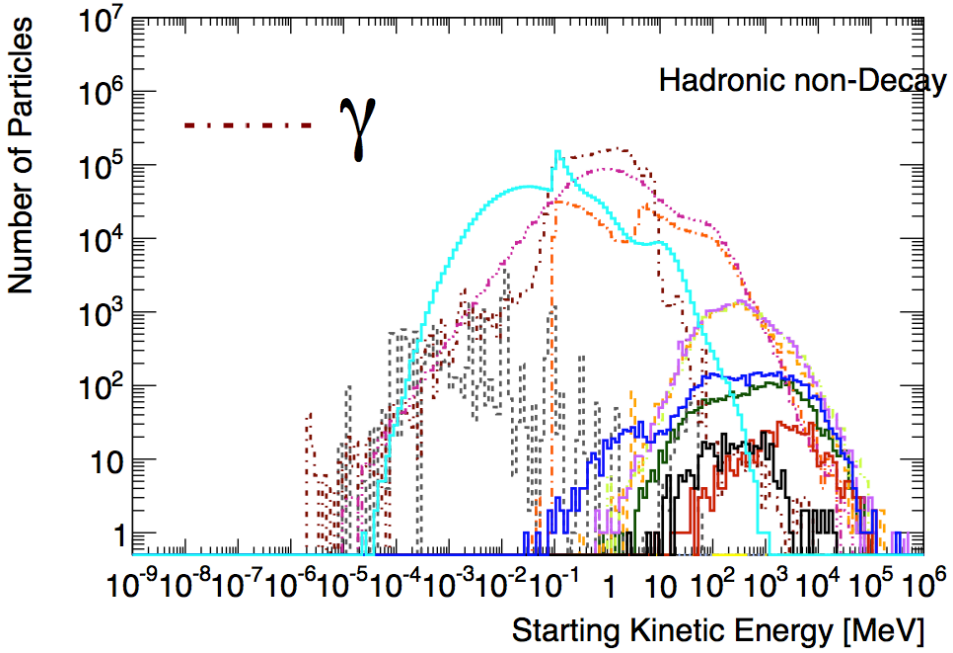

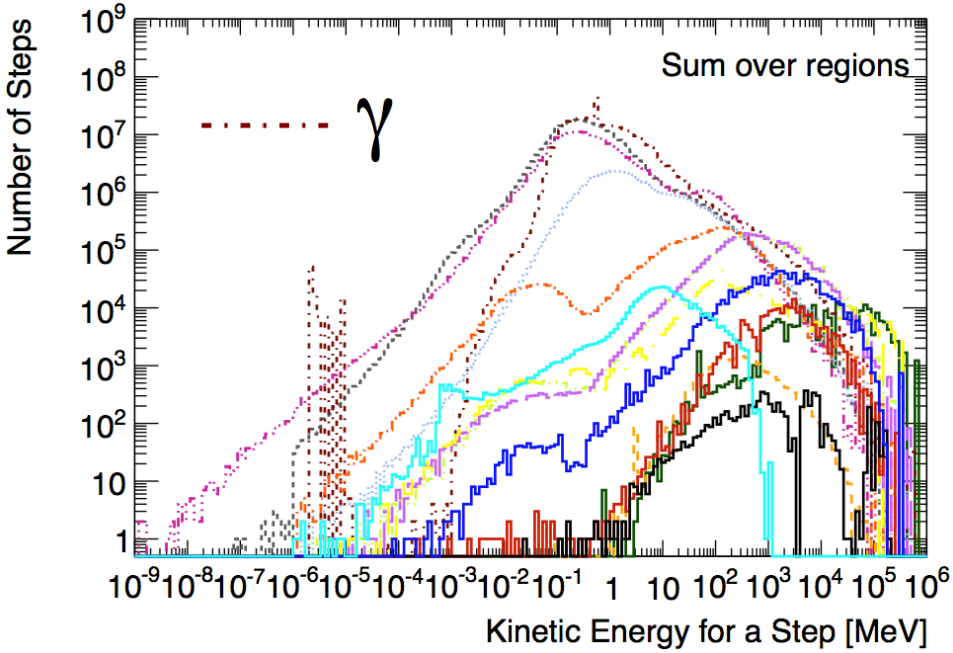

Bug: Electron Propagation

- In the EM and hadronic barrel calorimeters, there are a significant number of electrons propagating <100 fm in a step
  - Re-running now to try to drop the x-range of the histogram (batch is slow)
- There is *not* a correspondingly large number of electrons with a *total* track length below 100 pm
  - These are steps in a track, not single steps before the electron dies
- I'll bet good money this will turn out to be a bug, though I don't know where yet

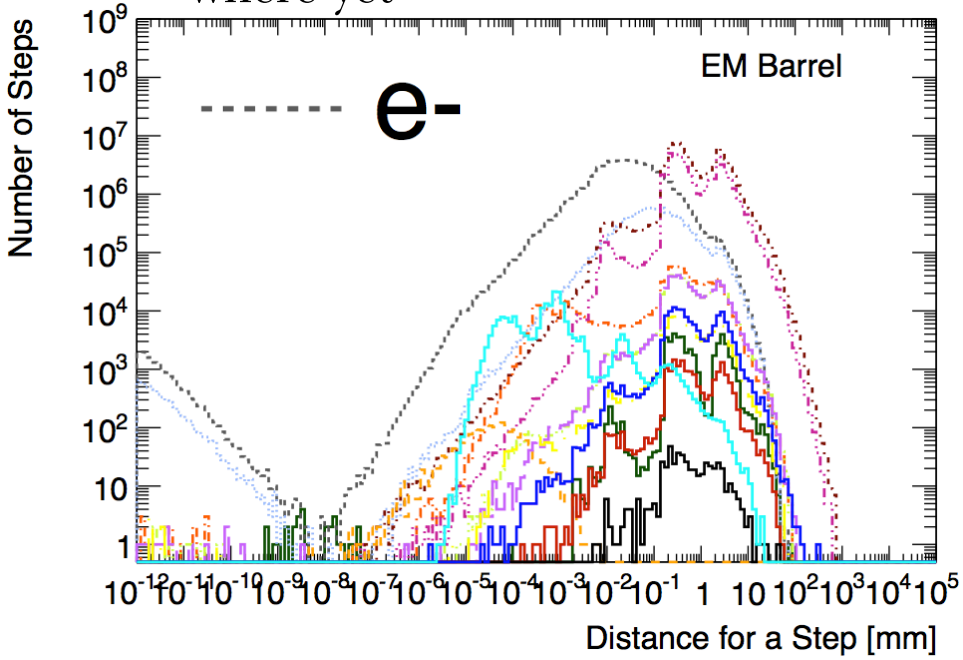

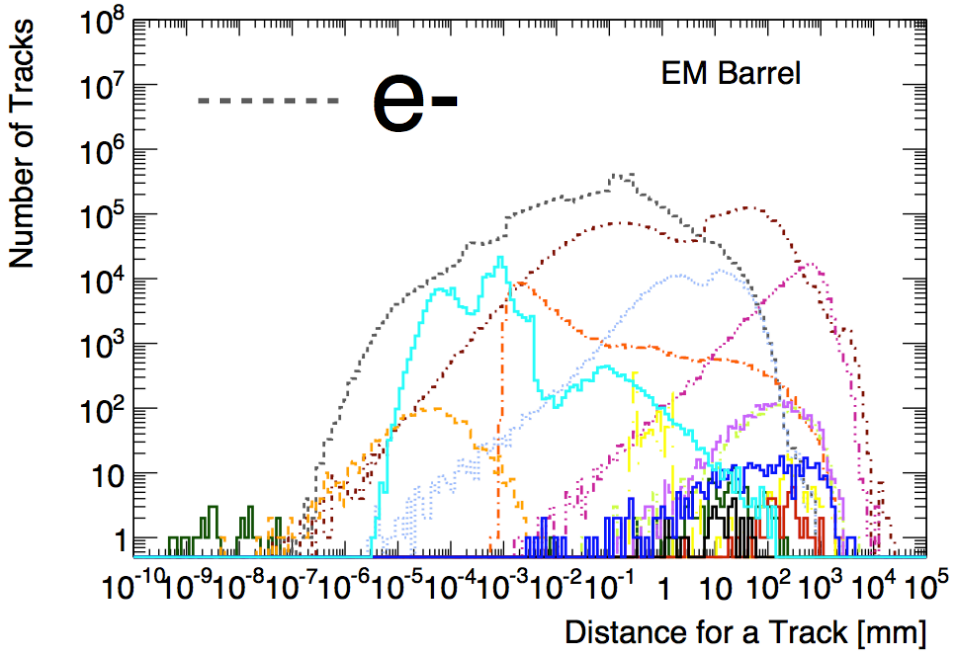

Feature: Neutrino Propagation

- Spot the production thresholds
- Notice a lot of high energy "others" – those are largely neutrinos that we are still propagating
  - Thought to be "conservative" back in the day. We'll start killing those now and see what we get back.

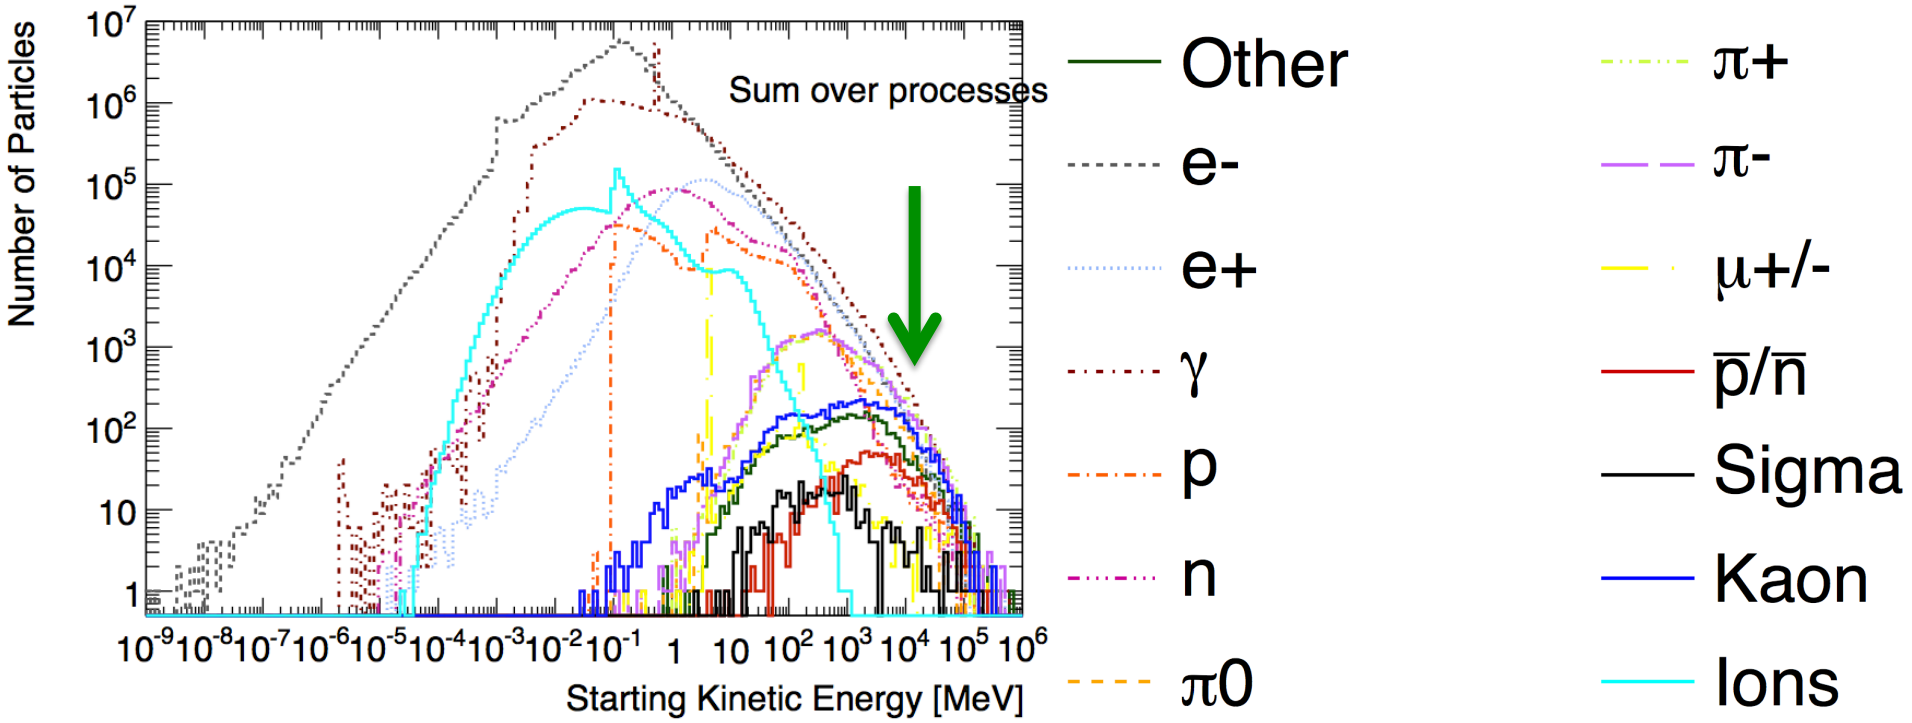

Feature: Low-energy Electrons

- Despite the production threshold in some processes, we have a surprisingly large number of *super* low energy electrons
- **Who wants to guess the process??**

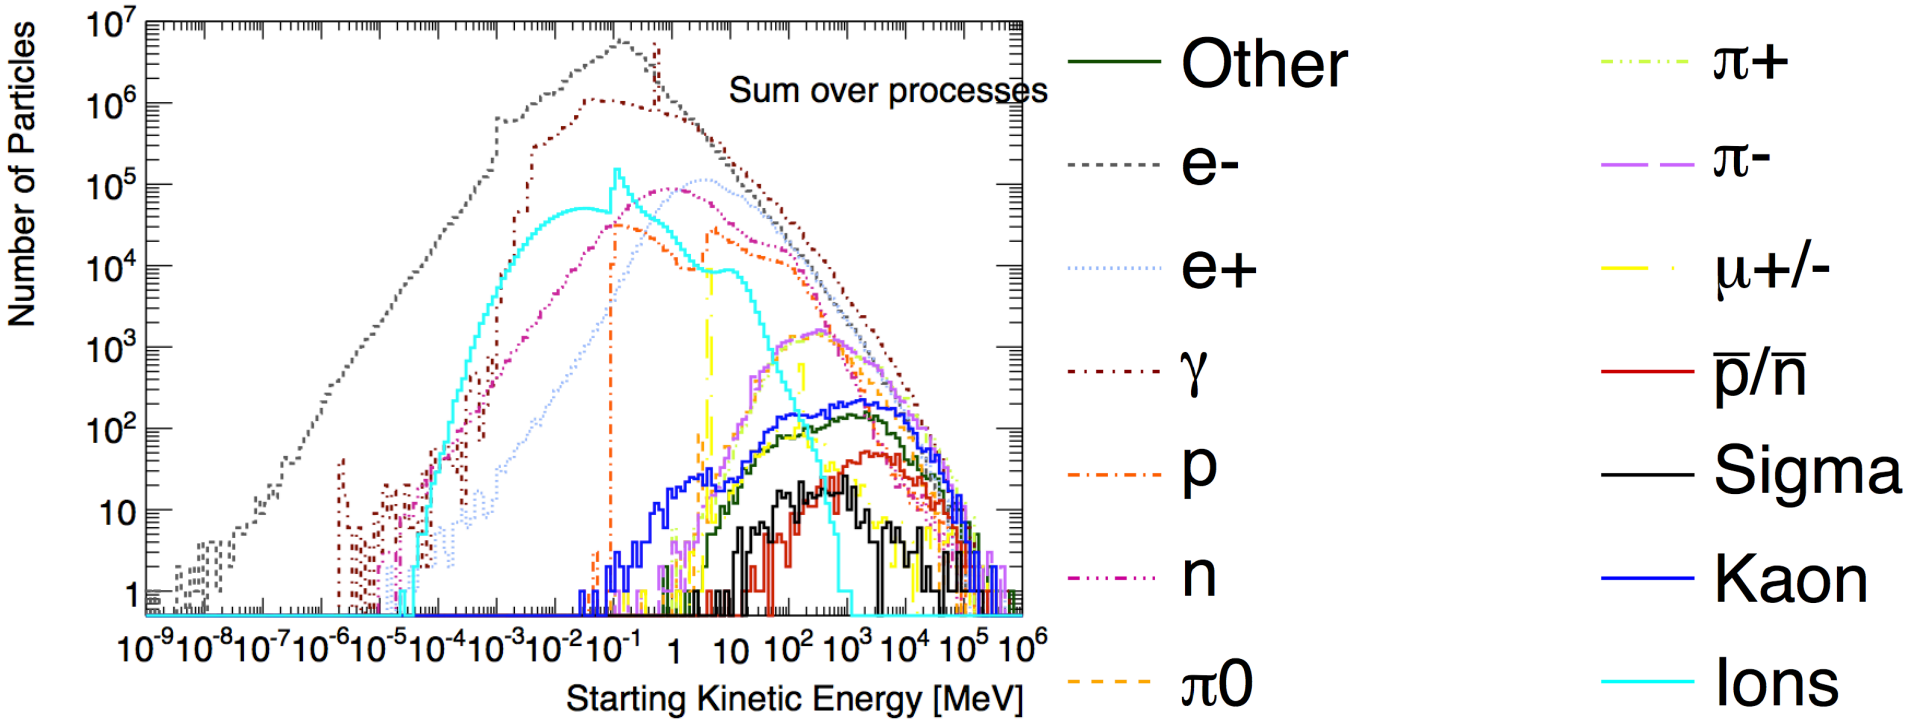

Feature: Low-energy Electrons

- Despite the production threshold in some processes, we have a surprisingly large number of *super* low energy electrons
- **Photoelectric effect and compton scattering!**
- Can we just ask these processes to respect range cuts?

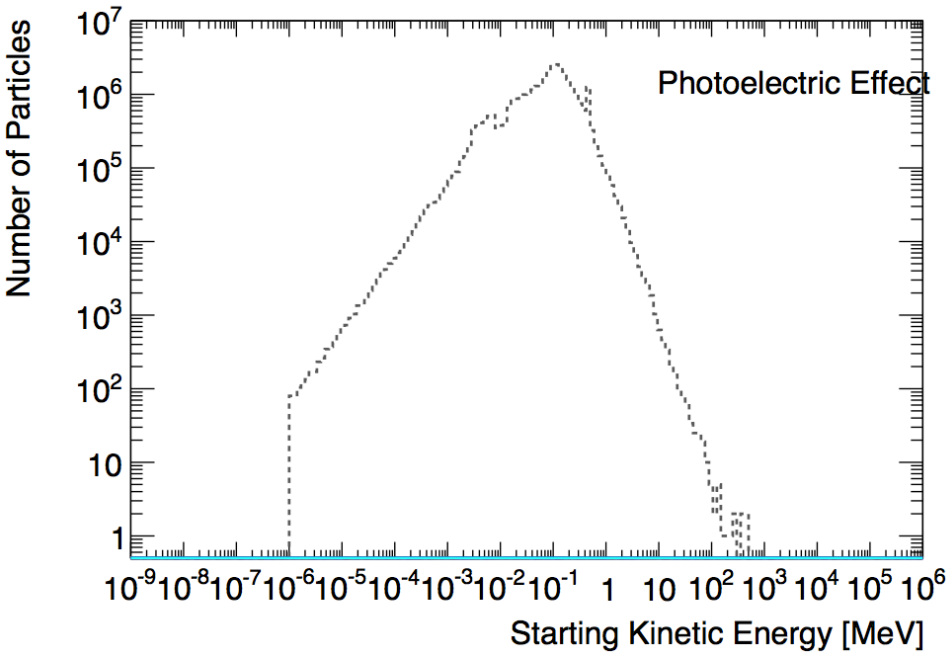

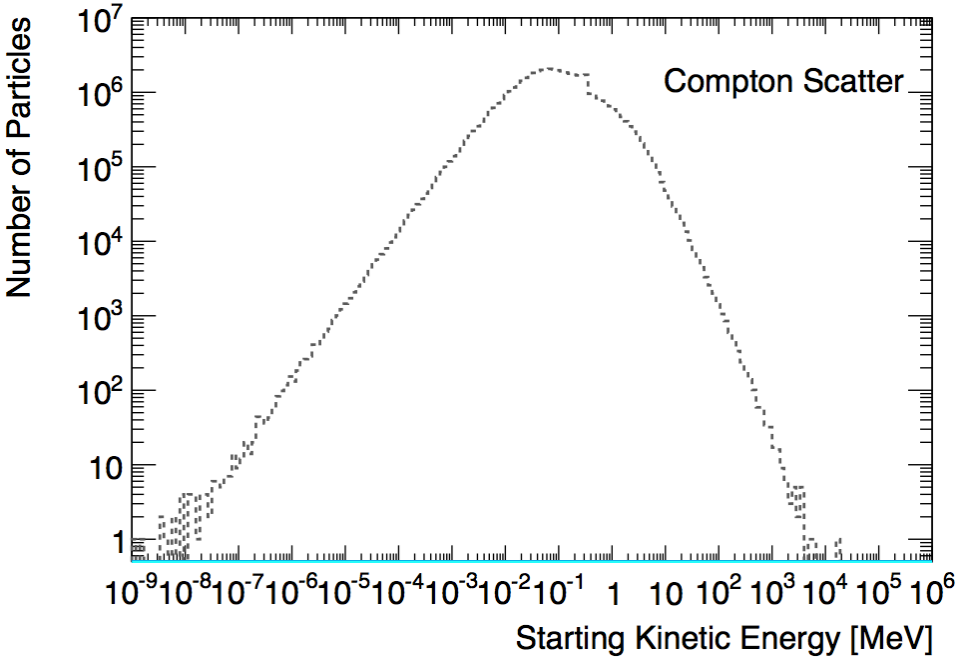

Feature: Longest tracks

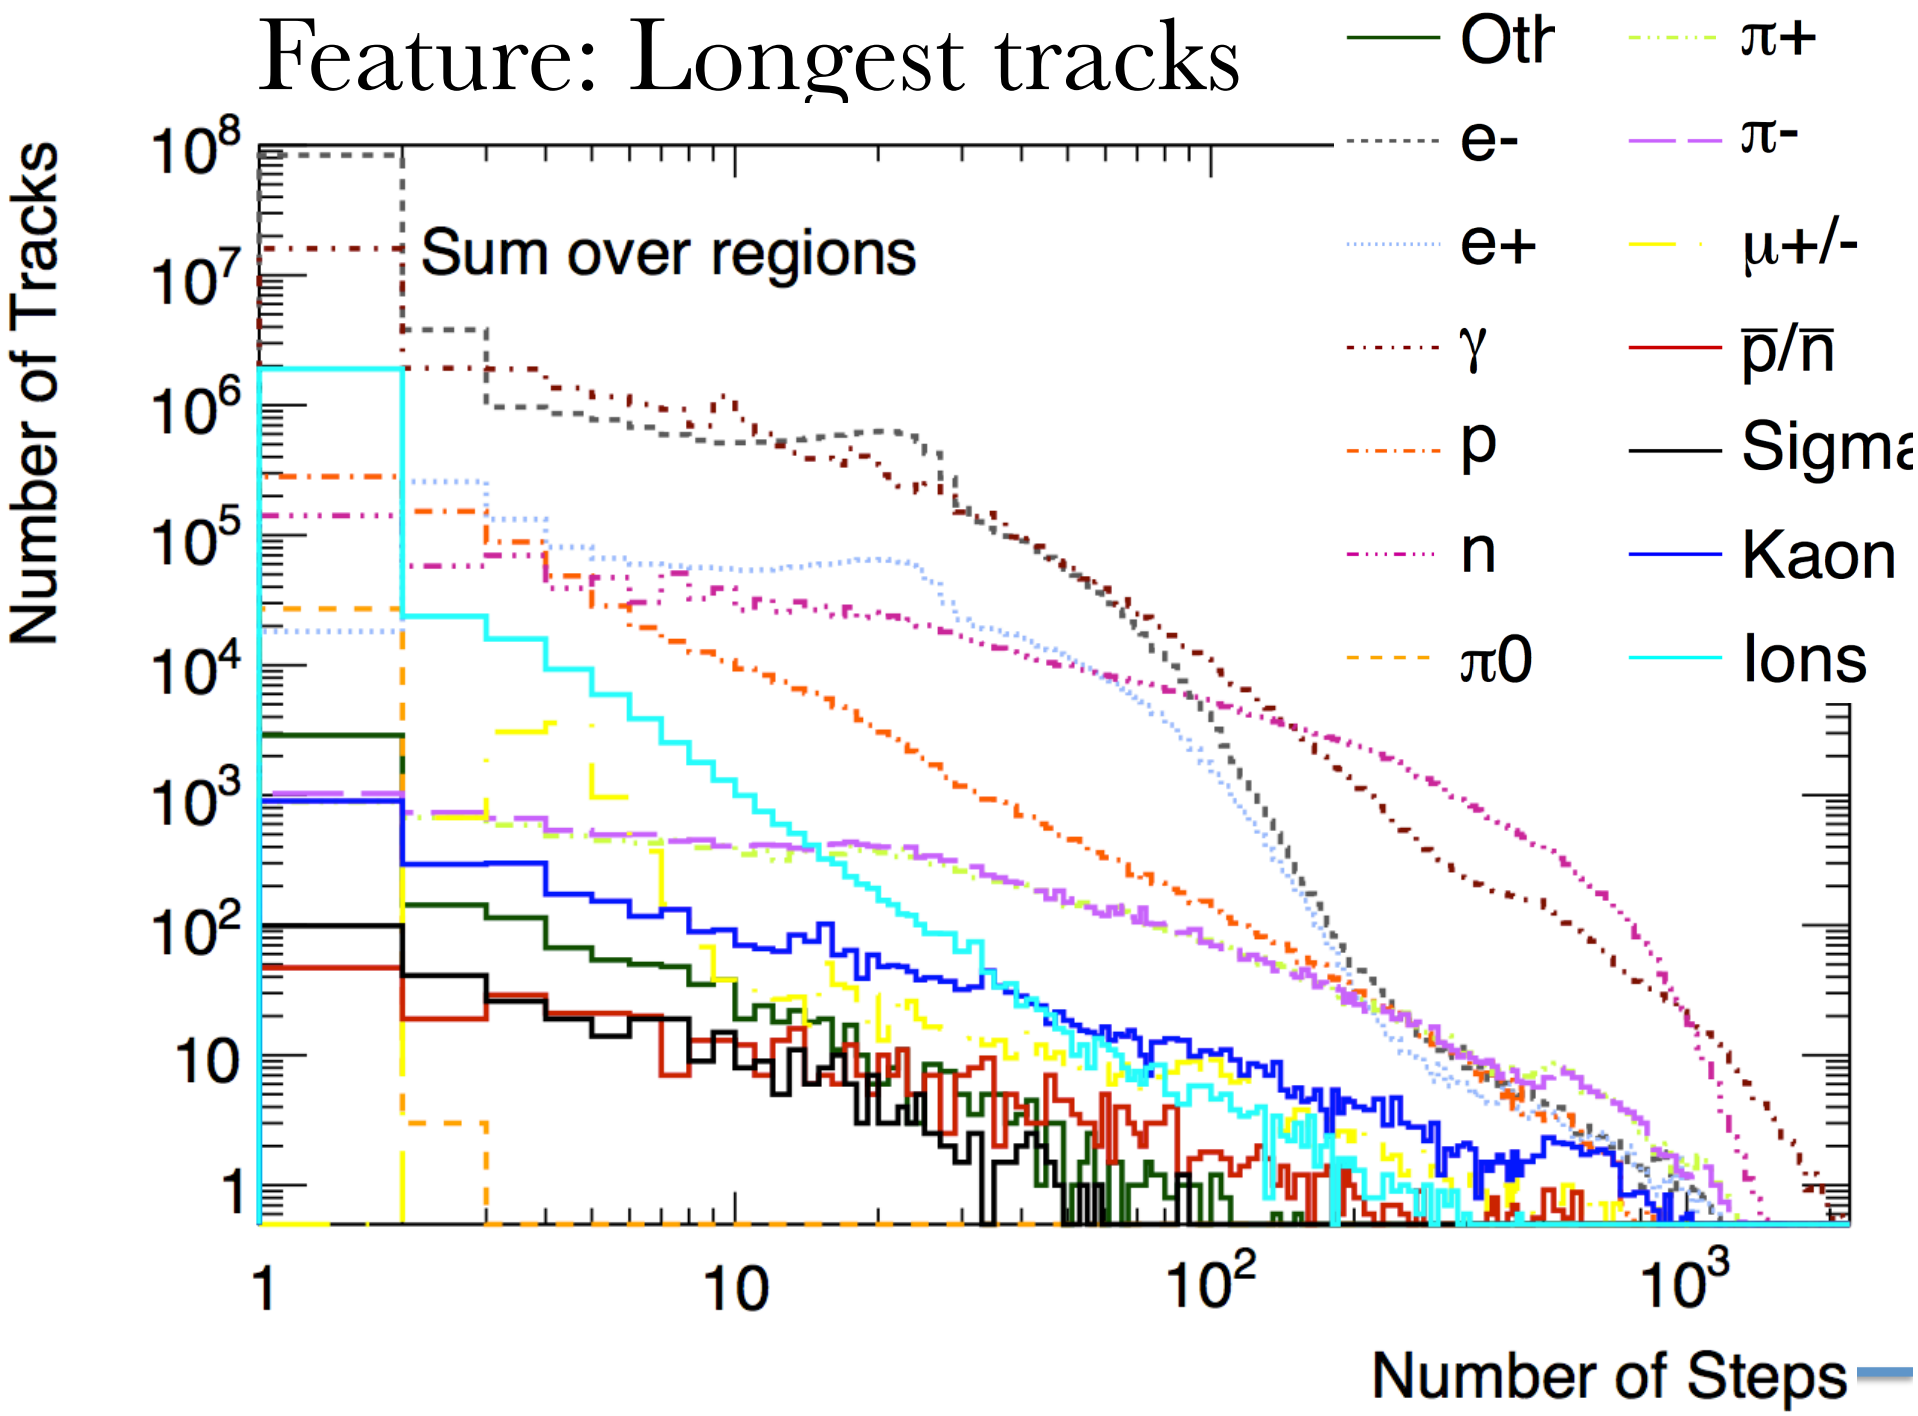

Feature: Long tracks – first reactions

- Probably the long photon tracks are the very low energy ones
- Still leaves a lot of other tracks taking a large number of steps
- Some of these are probably getting stuck – looper killer is a good thing – but we should see about hunting them down again, depending on how serious an issue it turns out to be
- Also, if the neutrons are low energy, might just kick them out…

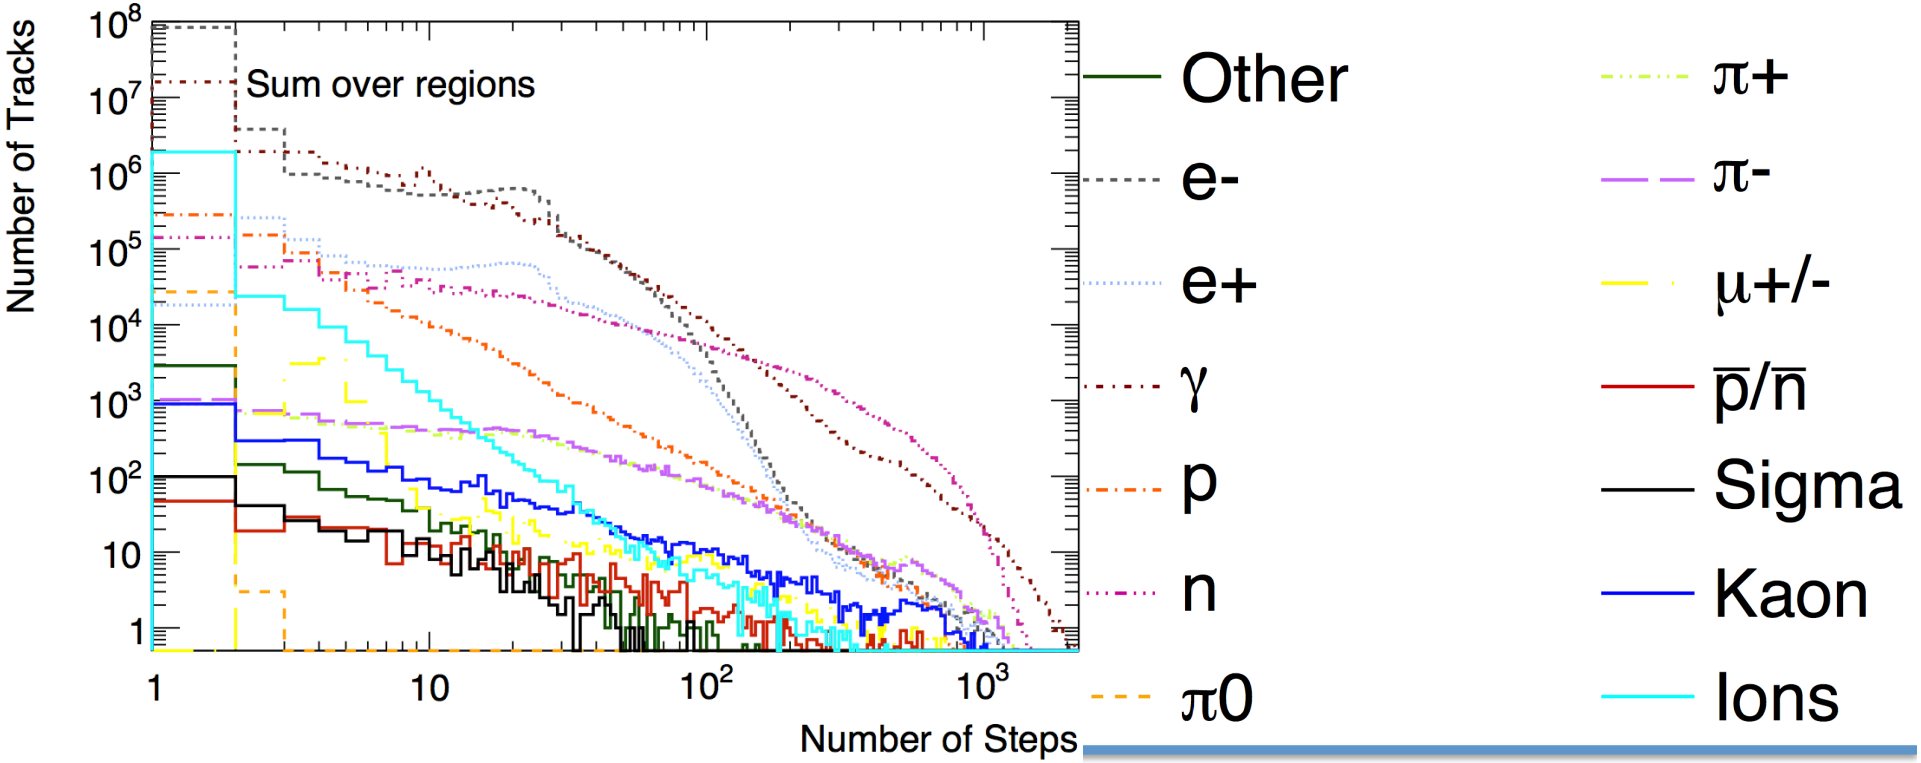

Newest data: steps of lowest E tracks

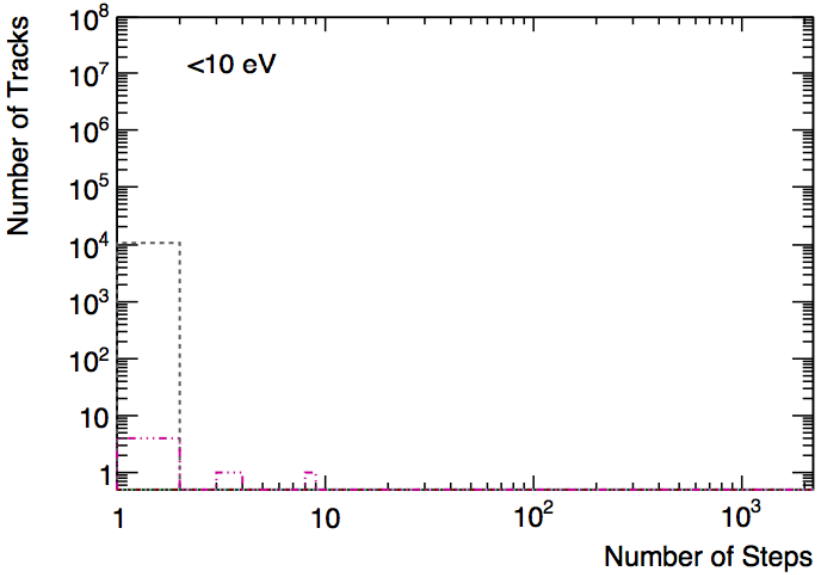

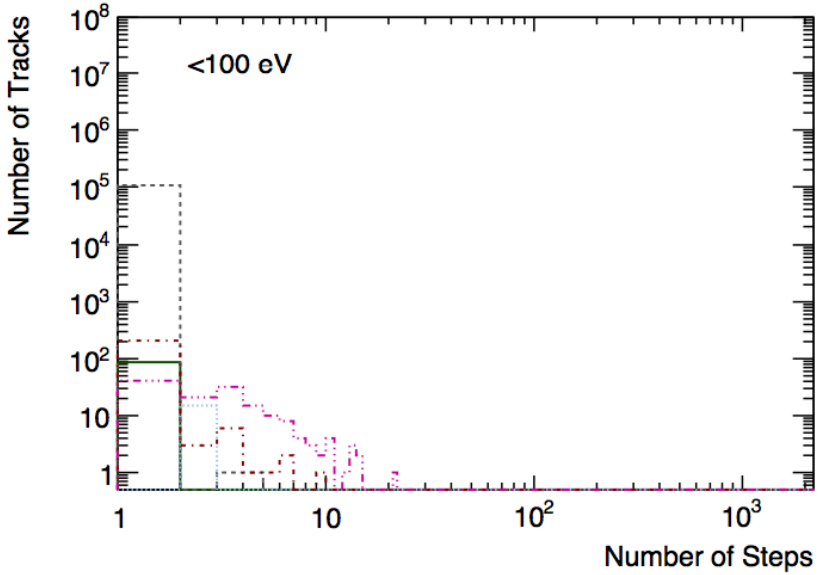

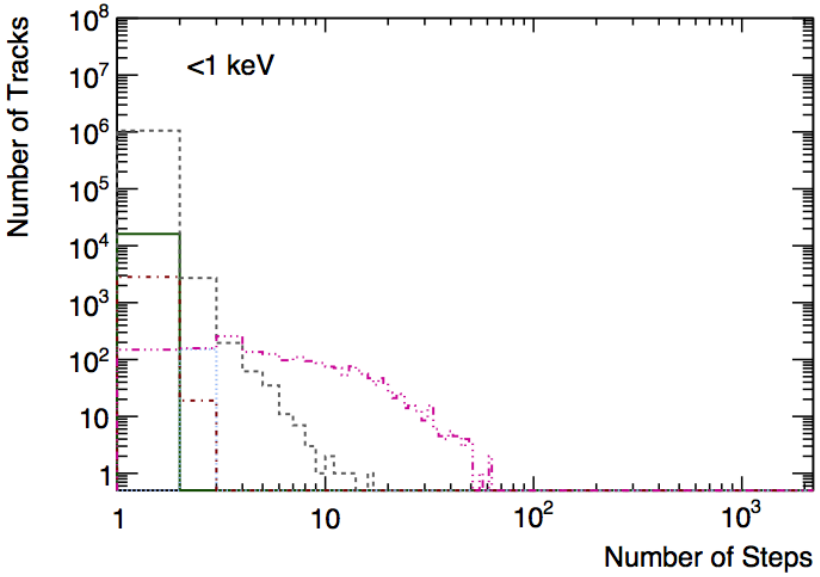

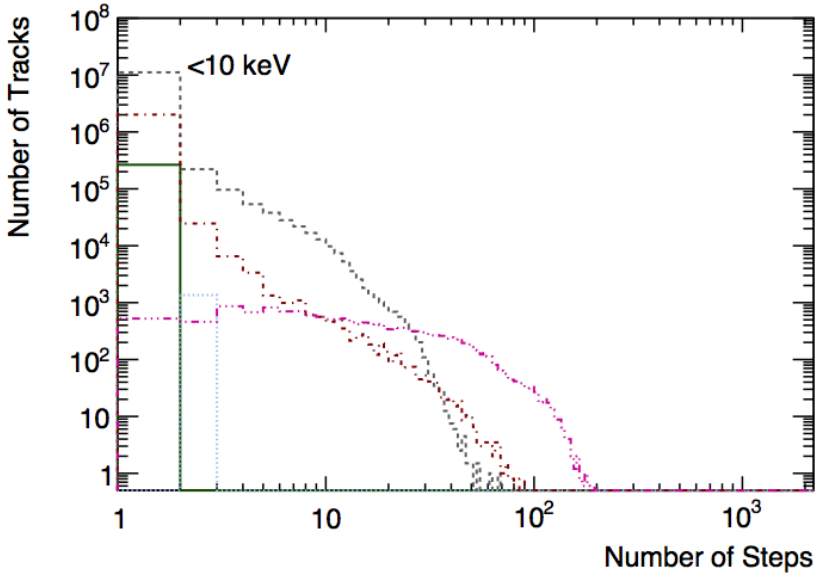

Steps – the 'medium energy' tracks

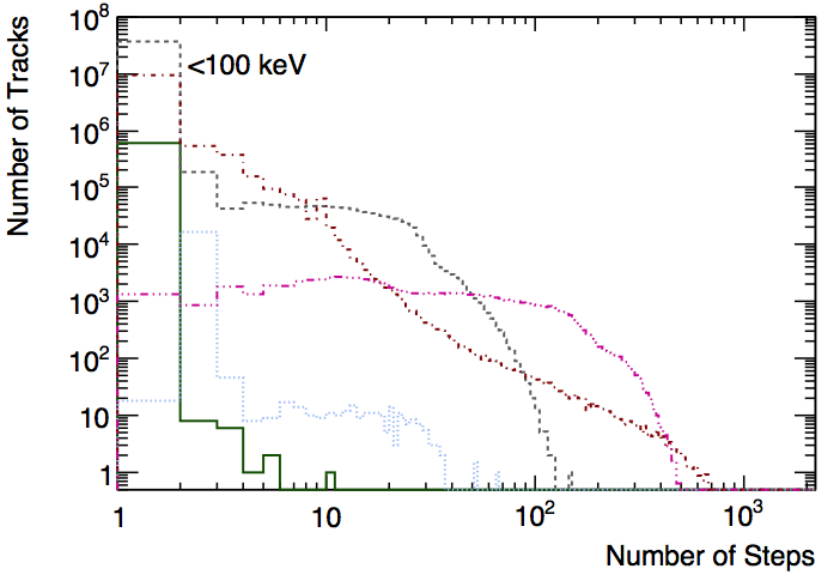

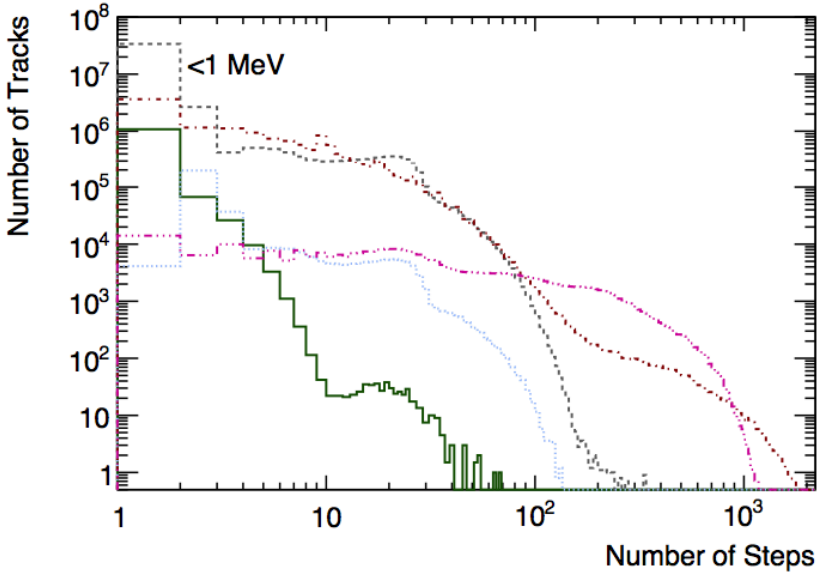

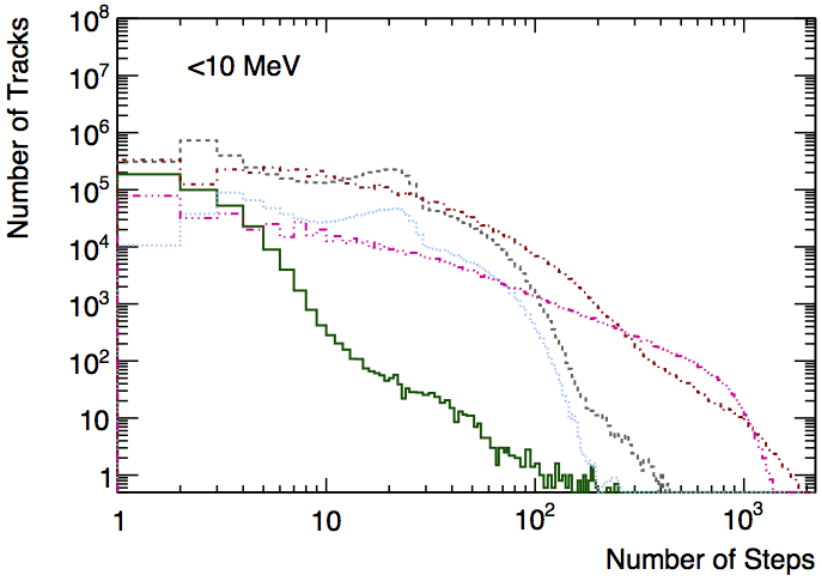

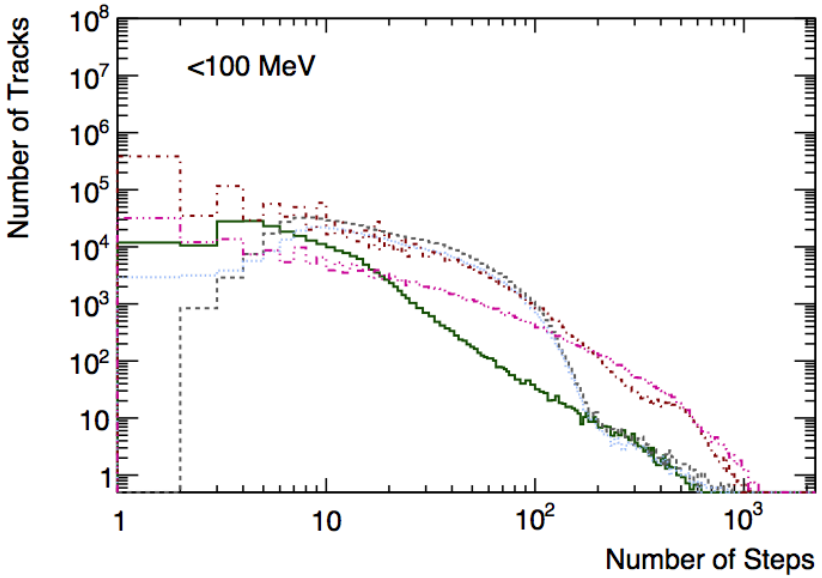

Steps – the highest energy tracks

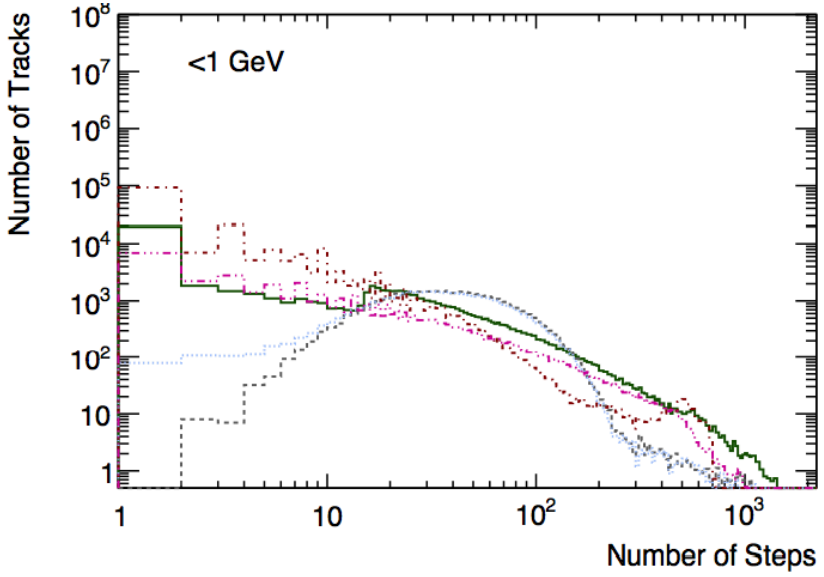

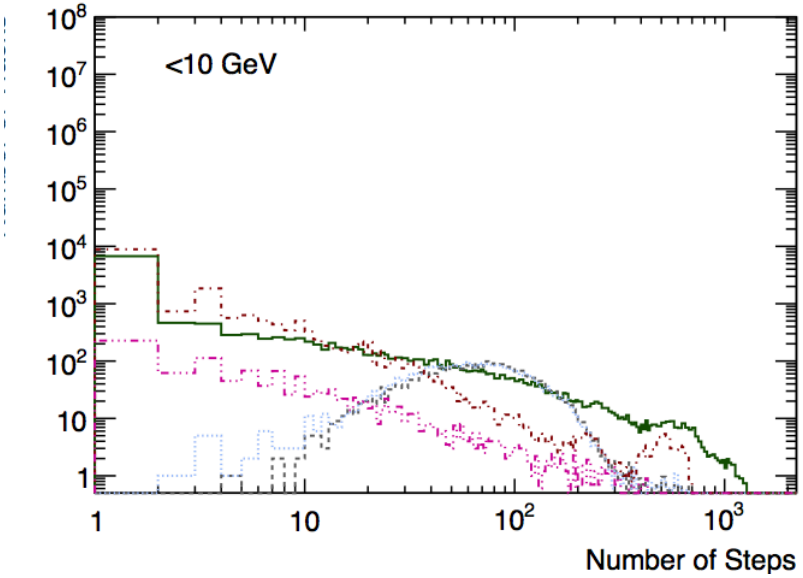

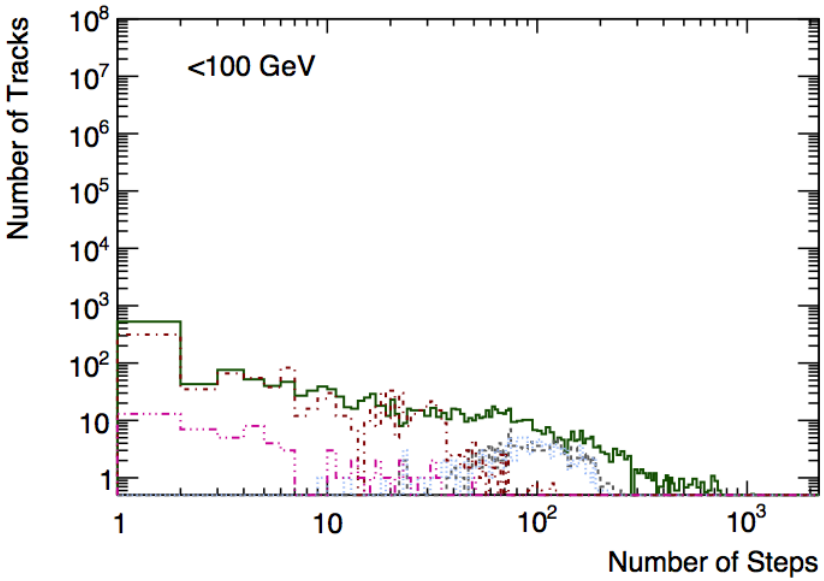

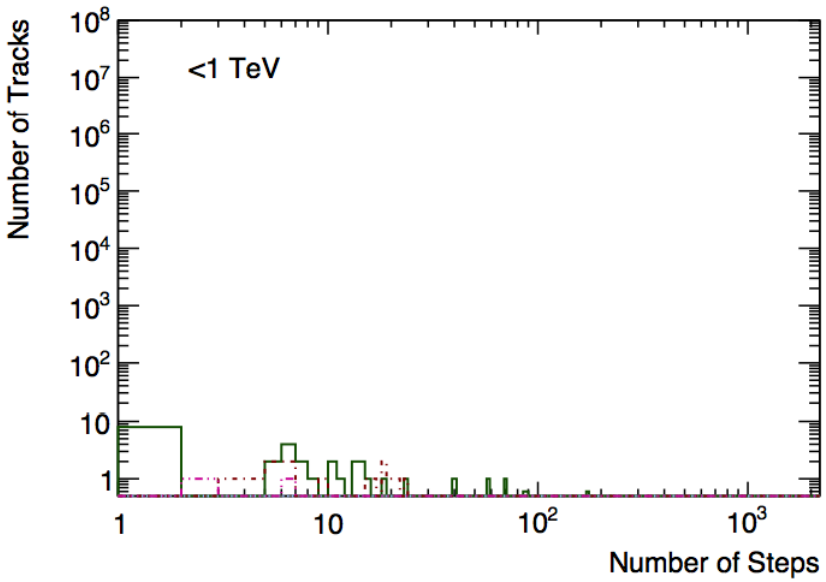

Summary and Future Work

- CPU and physics benchmarking done so far
- Lots and lots of interesting plots to look at
- Several problems (and solutions) already identified
  - Almost all of these are directly portable to ISF – and affect that simulation as well – so we gain for the present and the future
- Will begin iterating with G4 developers ~now to discuss whether these should be solved centrally or within ATLAS
  - Expect more feedback from ATLAS to G4 …
- Will keep mining the profiling for more useful tricks
- Will get memory profiling up and running in order to see where improvements can be made there
  - We expect to significantly reduce the footprint of the simulation

BACKUP SLIDES

First, a Bug!

- Initial profile with gathena showed a big do_system hotspot
- Makes little to no sense
- Traced to gathena not liking a system call for memory checking
- Removed the Memory snooper, no more hotspot! We should do this by default in the future – this is now functionality maintained by core sftw

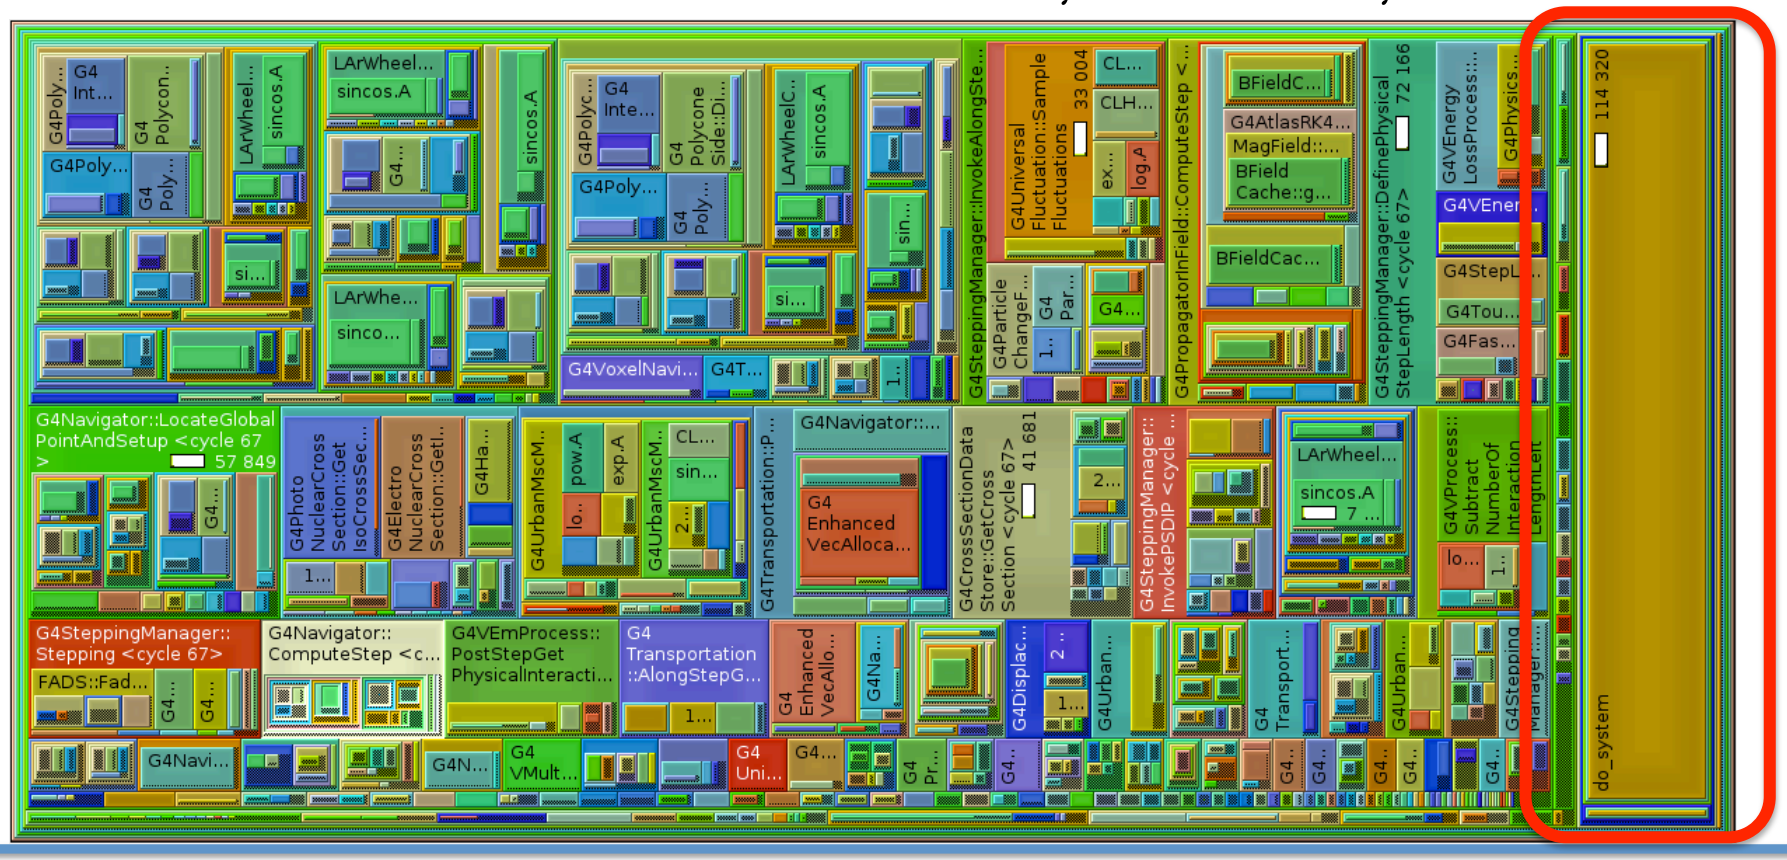

Key to next slides – colors for particles