



### NOvA

NOvA is a second generation long-baseline neutrino oscillation experiment situated in the Fermilab NuMI beam line. NOvA consists of a 300-ton near detector and a 14-kiloton far detector, separated by 810 km, which measure the  $\nu_e$  and  $\nu_\mu$  content of the NuMI beam before and after oscillation



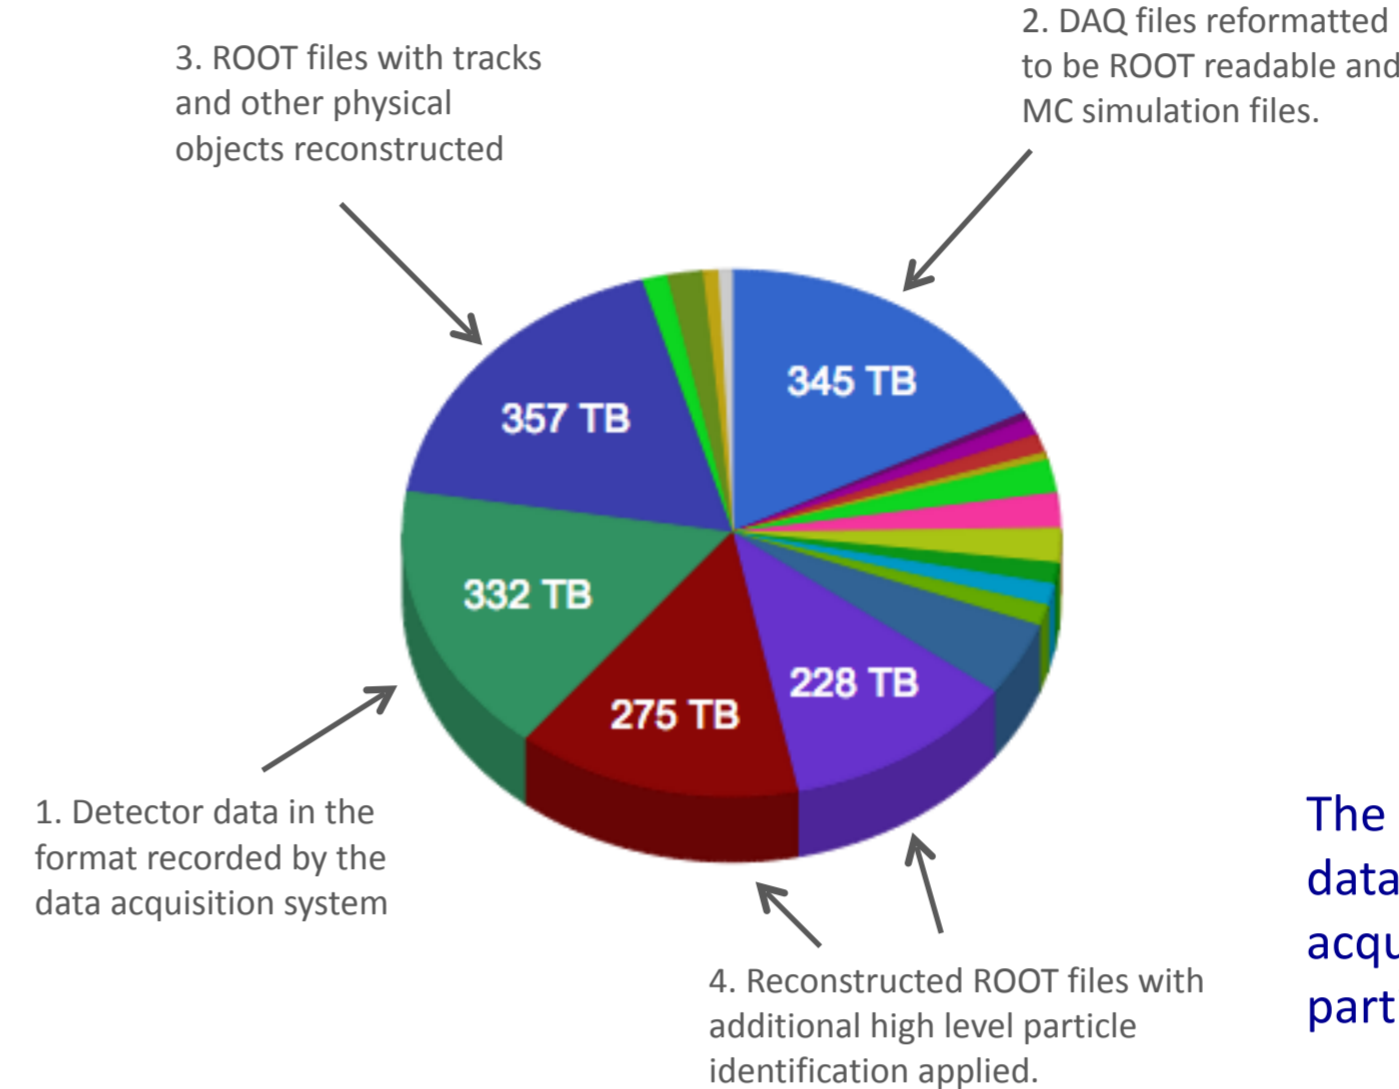
The NOvA Far Detector under construction. The completed detector measures 15m high x 15m wide x 60m long.



NOvA will make high precision measurements of the neutrino mixing parameters and aim to determine the neutrino mass hierarchy.

### Data & software

NOvA has already produced more than 1 million Monte Carlo and detector generated files amounting to more than 1 PB in size. This is divided between a number of parallel data streams such as far and near detector beam spills, cosmic ray backgrounds, several data-driven triggers and over 20 different Monte Carlo configurations.



Each of these data streams must be processed through the appropriate steps of the rapidly evolving, multi-tiered, interdependent NOvA software framework. In total there are greater than 12 individual software tiers, each of which performs a different function and can be configured differently depending on the input stream.

The total size of NOvA recorded data subdivided according to the different data tiers. These tiers are separated into logical steps such as: data-acquisition, offline file formatting, reconstruction of physics objects and particle ID. Each tier progressively adds higher level information into files.

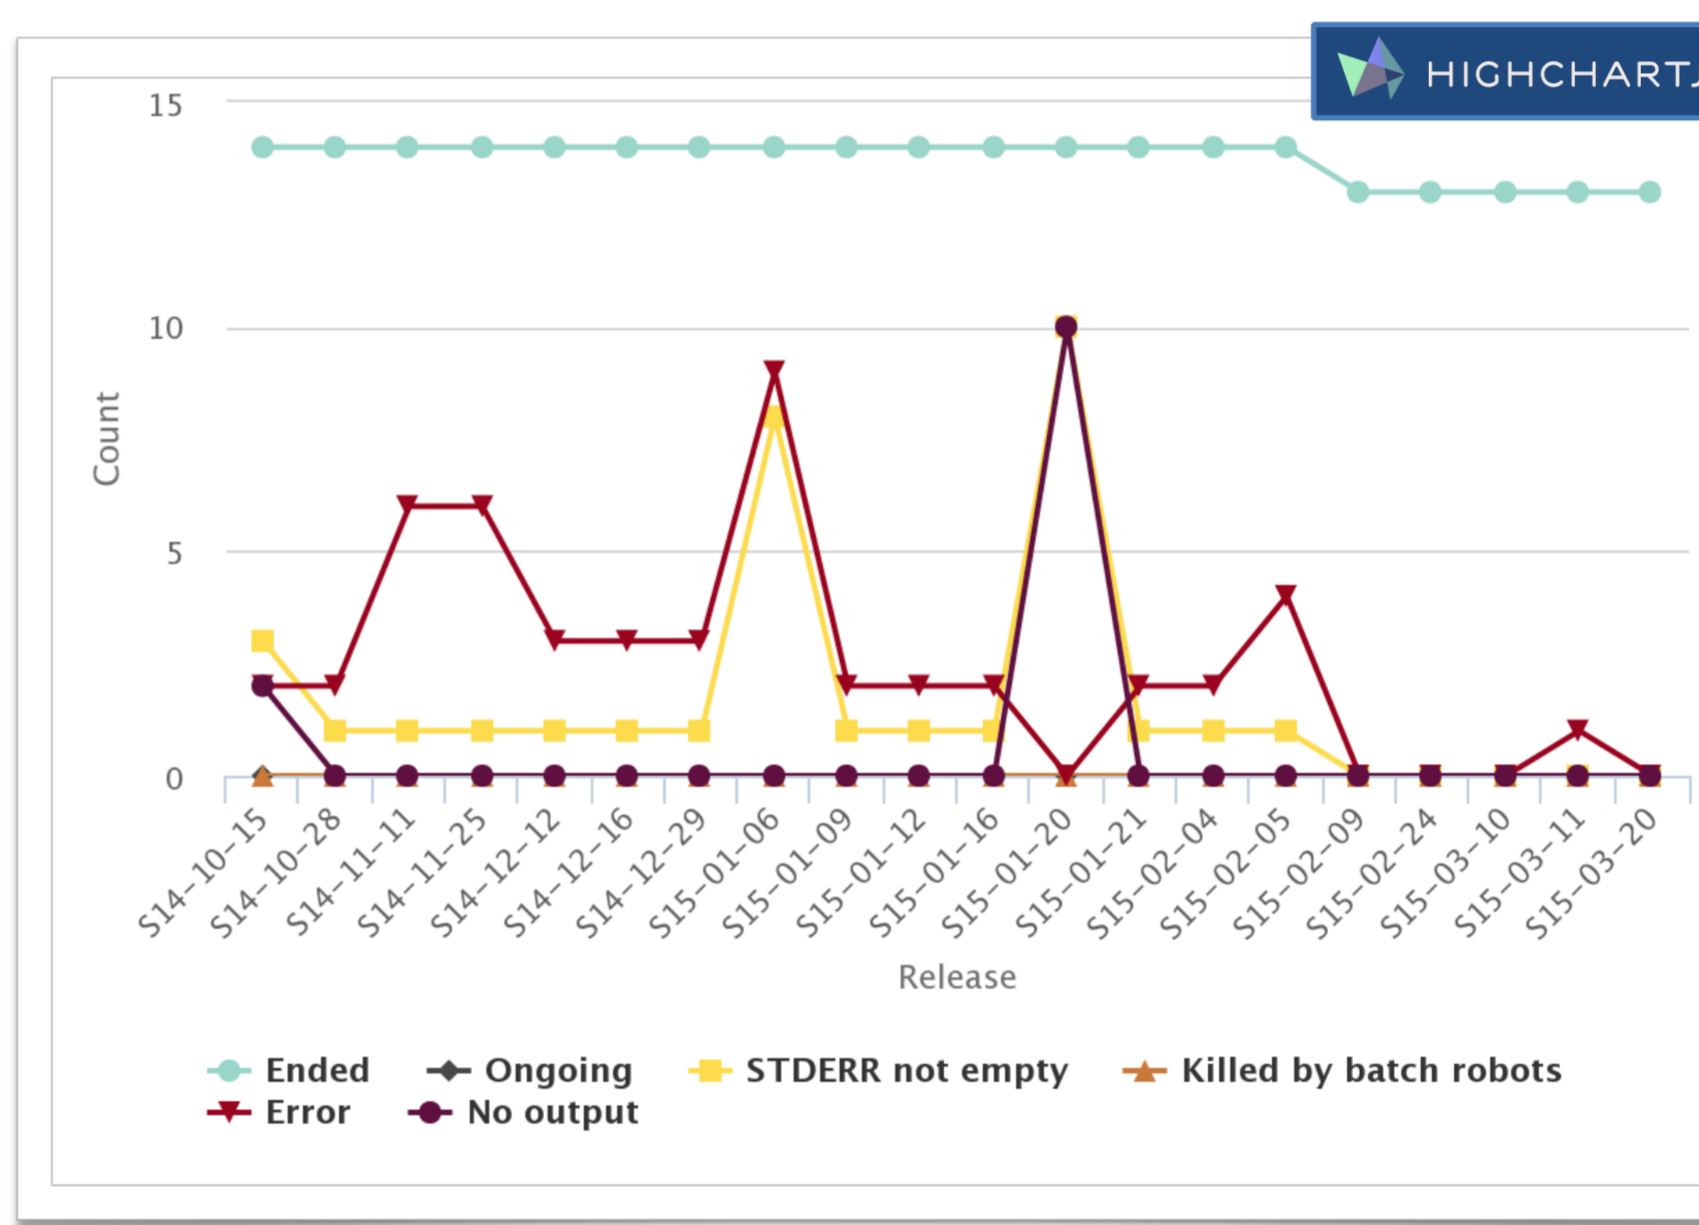
### The testing framework

In order to regularly test and validate that all of these software stages are working correctly NOvA has designed a powerful, modular testing framework that enables detailed validation and benchmarking to be performed in a fast, efficient and accessible way with minimal expert knowledge.

The core of this system is a novel series of python modules which wrap, monitor and handle the underlying C++ software framework and then report the results to a slick front-end web-based interface. This interface utilises modern, cross-platform, visualisation libraries to render the test results in a meaningful way. They are fast and flexible, allowing us to cater, easily, for new tests and datasets.

Tier	In evt	User CPU (/in evt) [s]	Memory	DB queries	Query time [s]	Child	Events (efficiency)	Size (/out evt)
cry log, fcl, metrics	20	576.60 (28.83)	651.51 MB	2	0.01	osmics_gen.root	20 (100 [%])	53.91 MB (2.7 MB)
pchits log, fcl, metrics	20	21.06 (1.05)	372.58 MB	7	0.21	clist_reco.root	20 (100 [%])	2.66 MB (136.35 KB)
reco log, fcl, metrics	-	-	-	-	-	tstop_reco.root	19 (95 [%])	194.75 KB (10.25 KB)
- failed, got return code: 65	-	-	-	-	-	timingcal.root	1 (5 [%])	84.48 KB (84.48 KB)

An example of the web output displaying the results of a number of tests. In this case a problem has been identified with the reconstruction.



An example of a test summary graphic showing the collected results of many tests on static snapshots of the software framework over five months.

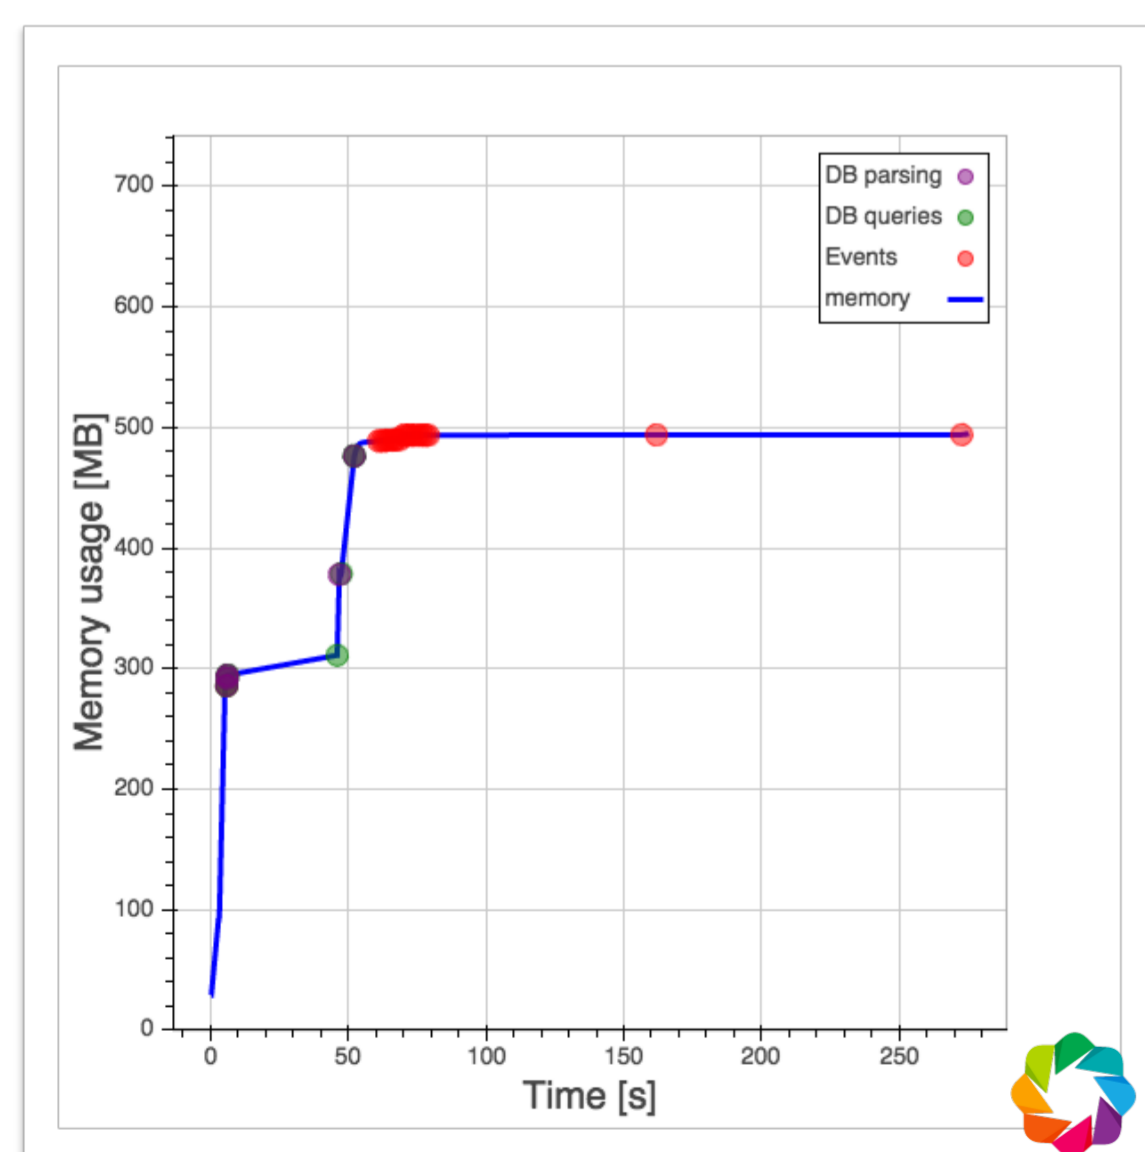
In total upwards of 14 individual streams are regularly tested amounting to over 70 individual software processes, producing over 25 GB of output files. These include nightly tests of the current state of the art software setup as well as automated tests of any static releases created.

The rigour enforced through this flexible testing framework enables NOvA to rapidly verify configurations, results and software and thus ensure that data is available for physics analysis in a timely and robust manner.

FD_data_cosmics:raw2root:	cmd, log, out, err	✓
FD_data_NuMI:raw2root:	cmd, log, out, err	✓
FD_cosmics	cmd, log, out, err	✗
FD_genie_FHC_nonswap	cmd, log, out, err	✓
FD_genie_FHC_swap	cmd, log, out, err	✓
FD_genie_FHC_tau	cmd, log, out, err	✓
FD_genie_FHC_nonswap	cmd, log, out, err	✓
FD_genie_RHC_swap	cmd, log, out, err	✓
FD_genie_RHC_tau	cmd, log, out, err	✓
ND_data_cosmics:raw2root:	cmd, log, out, err	✓
ND_data_NuMI:raw2root:	cmd, log, out, err	✓
ND_cosmics	cmd, log, out, err	✗
ND_genie_FHC_nonswap	cmd, log, out, err	✗
ND_genie_FHC_swap	cmd, log, out, err	✗
ND_genie_RHC_nonswap	cmd, log, out, err	✗

An example of overview table presented at the top of each test summary. This provides a visual overview as well as all of the batch process logs.

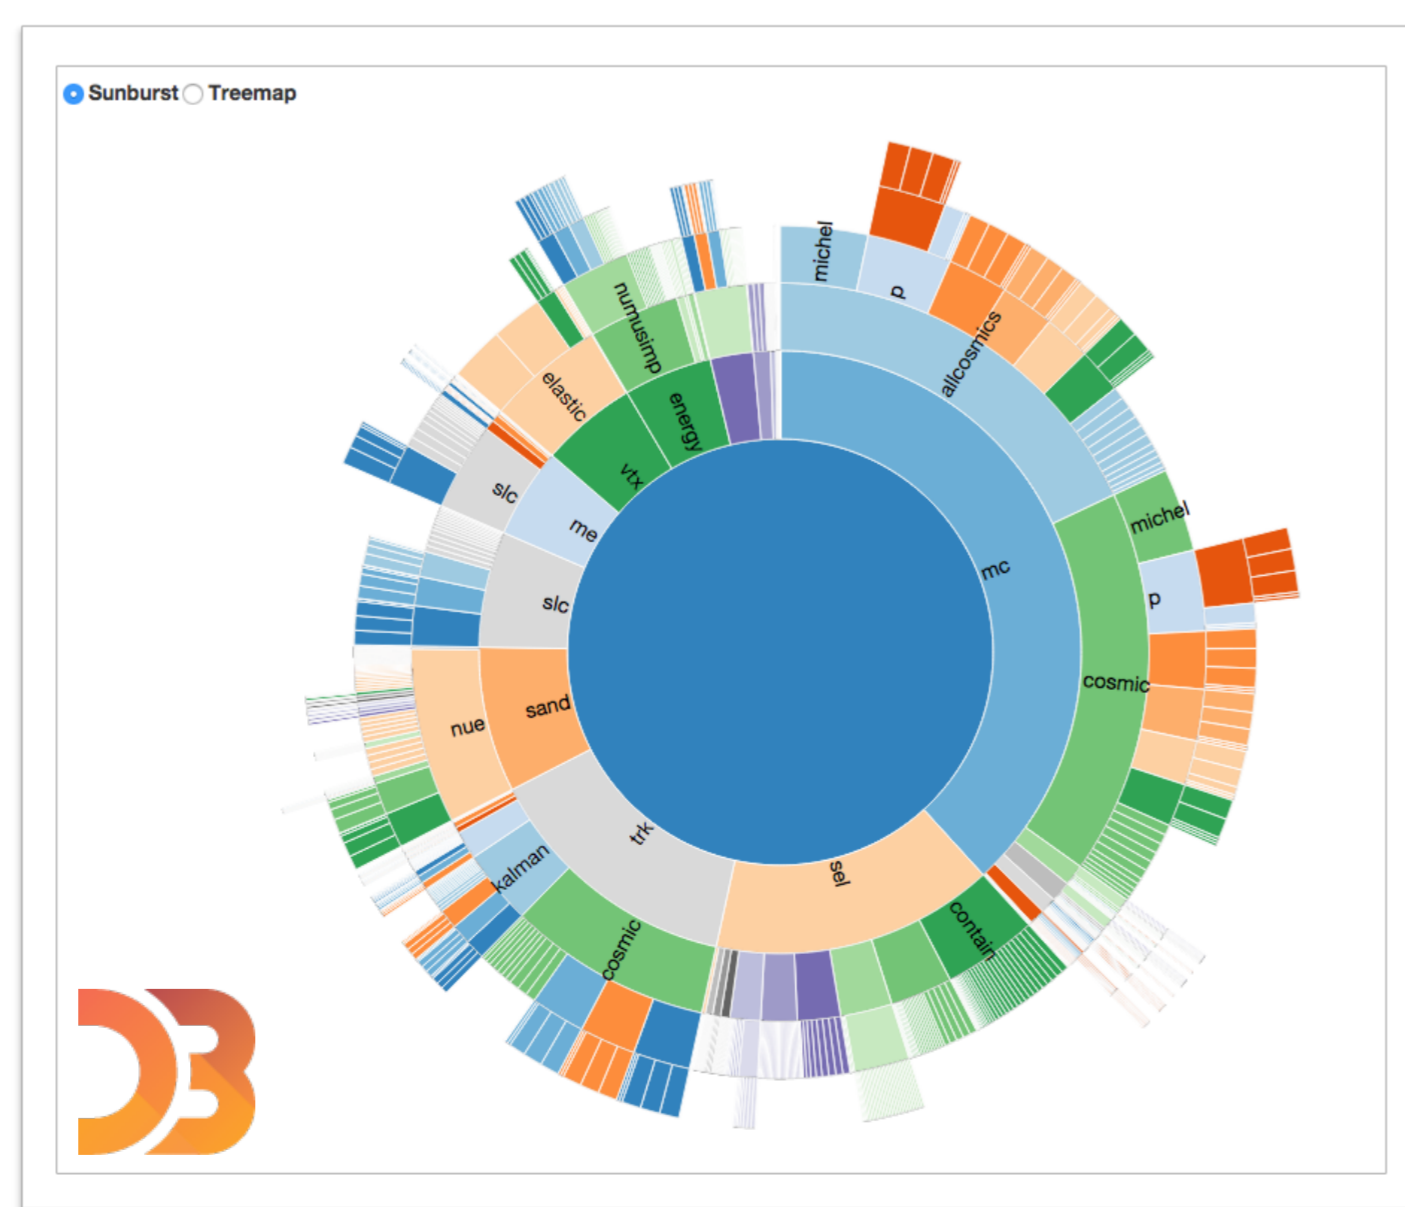
### Visualising the results



The memory usage of one test visualised using the Bokeh python library. The coloured circles are displayed at times when important processing steps occur. These provide the user with contextual information on mouse hover.

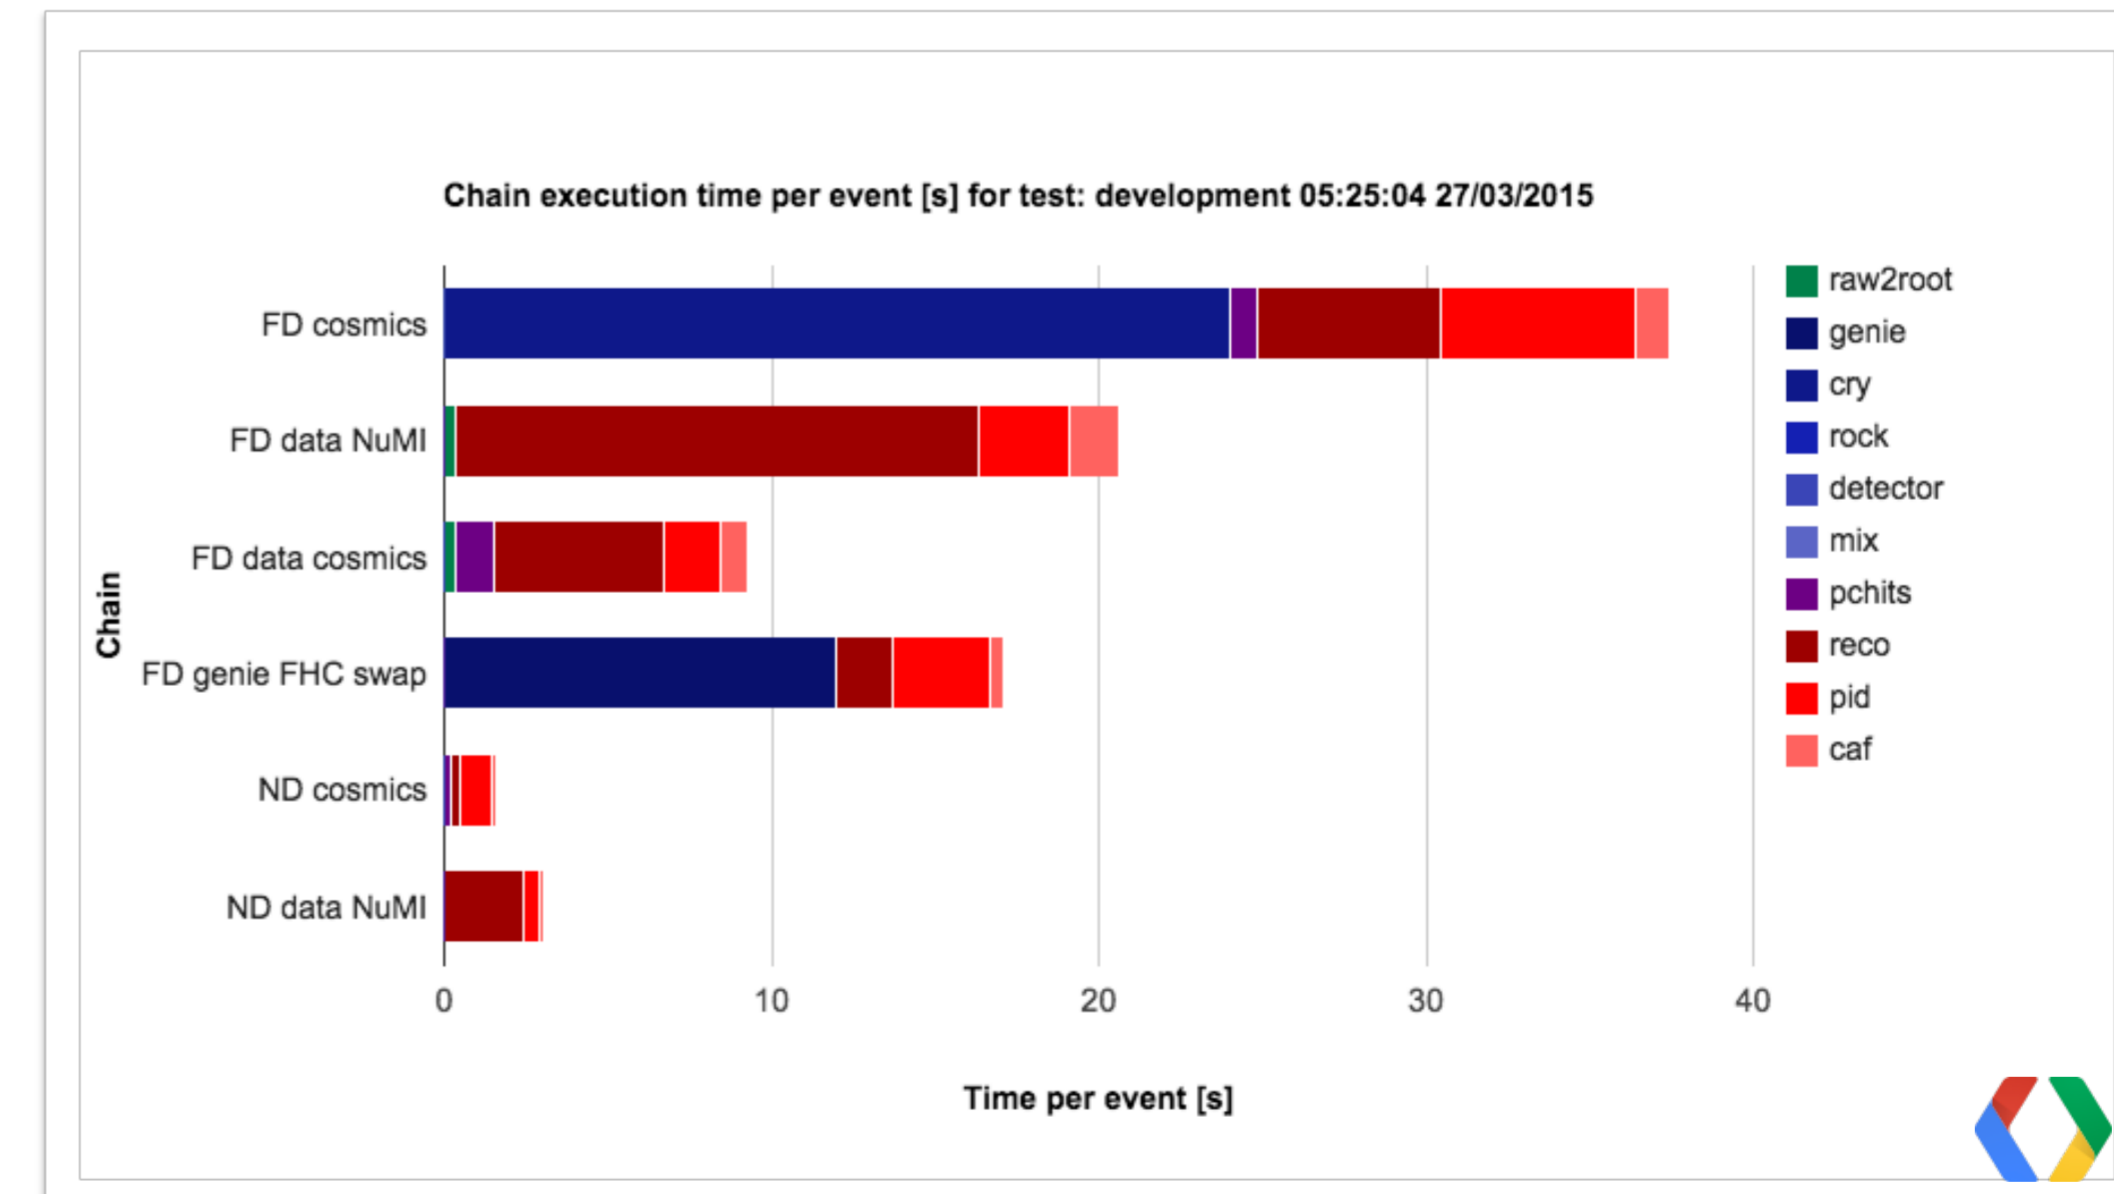
The results are interactive graphs and charts that allow the experimenter to contextualise important features, such as spikes in CPU or memory usage associated with events occurring within the processing or the point of failure of a given process.

The outputs of the tests are analysed using a variety of web-based visualisation tools in order to convey meaningful information to the experimenter. Examples of these include CPU and memory usage, and a detailed breakdown of the size and content of created files.



A D3 based graphic displaying the size of one data format produced during tests. This graphic allows users to drill down to study the particular portion of the data format they are concerned with.

### Predicting future needs



A bar chart, produced using google charts, showing the recorded time per event for a number of different tiers run on a number of different data streams. As with the other graphics this is interactive enabling the user to select processing tiers and data streams at will.

The benchmarking information reported by the tests is utilised to forecast the future computational needs of the experiment in terms of integrated CPU usage, memory footprint and storage volume needed.

This information allows NOvA to rapidly and accurately predict their future needs and to plan production campaigns accordingly.