

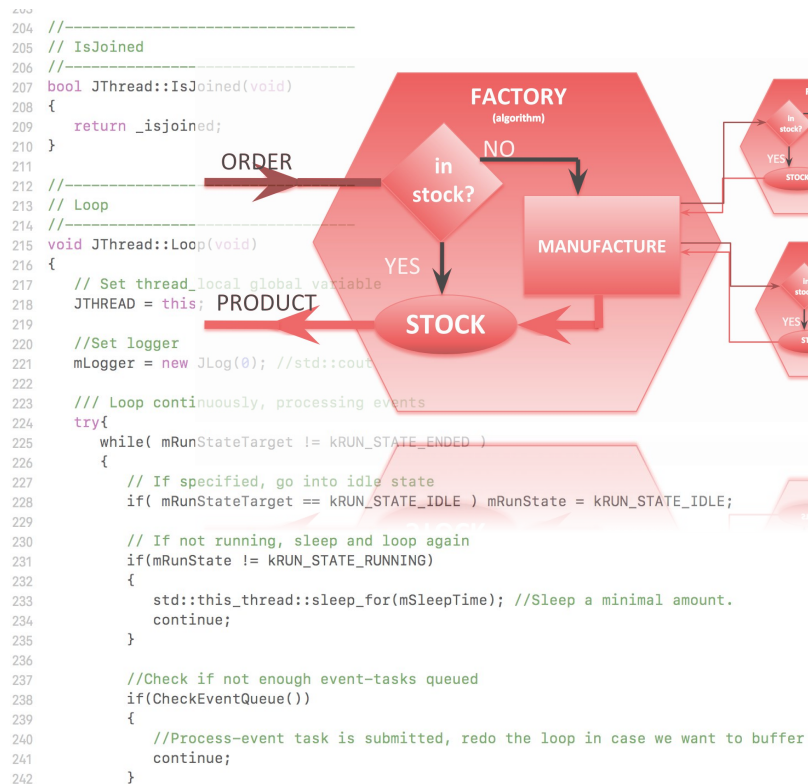
JANA2: Multi-threaded Event Reconstruction

David Lawrence
Jefferson Lab

September 21, 2022

HSF Framework Working Group

Jefferson Lab



The JANA Framework

- JANA is a multithreaded reconstruction framework project with nearly 2 decades of experience behind it
- JANA2 is a rewrite incorporating more modern coding and CS practices and improving on the original using lessons learned
 - Streaming DAQ and Heterogeneous hardware support strongly considered in redesign

Projects using JANA

- GlueX
- INDRA-ASTRA (*near-realtime calibrations using AI/ML*)
- BDX
- TriDAS (+ERSAP) + JANA2 Streaming DAQ

GlueX Computing Needs

	2017 (low intensity GlueX)	2018 (low intensity GlueX)	2019 (PrimEx)	2019 (high intensity GlueX)
Real Data	1.2PB	6.3PB	1.3PB	3.1PB
MC Data	0.1PB	0.38PB	0.16PB	0.3PB
Total Data	1.3PB	6.6PB	1.4PB	3.4PB
Real Data CPU	21.3Mhr	67.2Mhr	6.4Mhr	39.6Mhr
MC CPU	3.0Mhr	11.3Mhr	1.2Mhr	8.0Mhr
Total CPU	24.3PB	78.4Mhr	7.6Mhr	47.5Mhr

Anticipate 2018 data will be processed by end of summer 2019

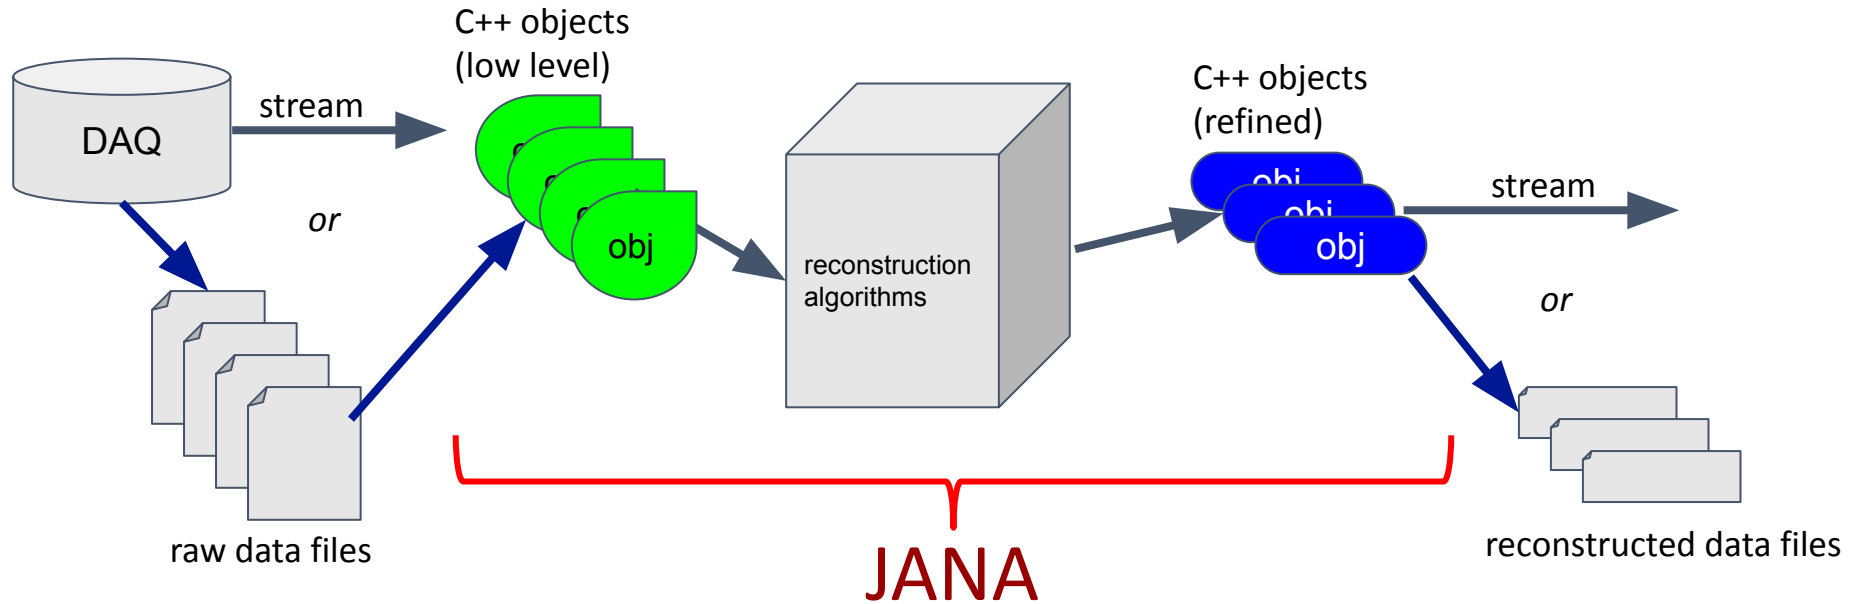
Projection for out-years of GlueX High Intensity running at 32 weeks/year

11/27/18

	Out - years (high intensity GlueX)
Real Data	16.2PB
MC Data	1.4PB
Total Data	17.6PB
Real Data CPU	125.6Mhr
MC CPU	36.5Mhr
Total CPU	162.1Mhr


Event size:
12-13kB

JANA's Role in Data Processing



JANA2 arrows separate sequential and parallel tasks

- CPU intensive event reconstruction will be done as a parallel arrow
- Other tasks (e.g. I/O) can be done as a sequential arrow
- Fewer locks in user code allows framework to better optimize workflow

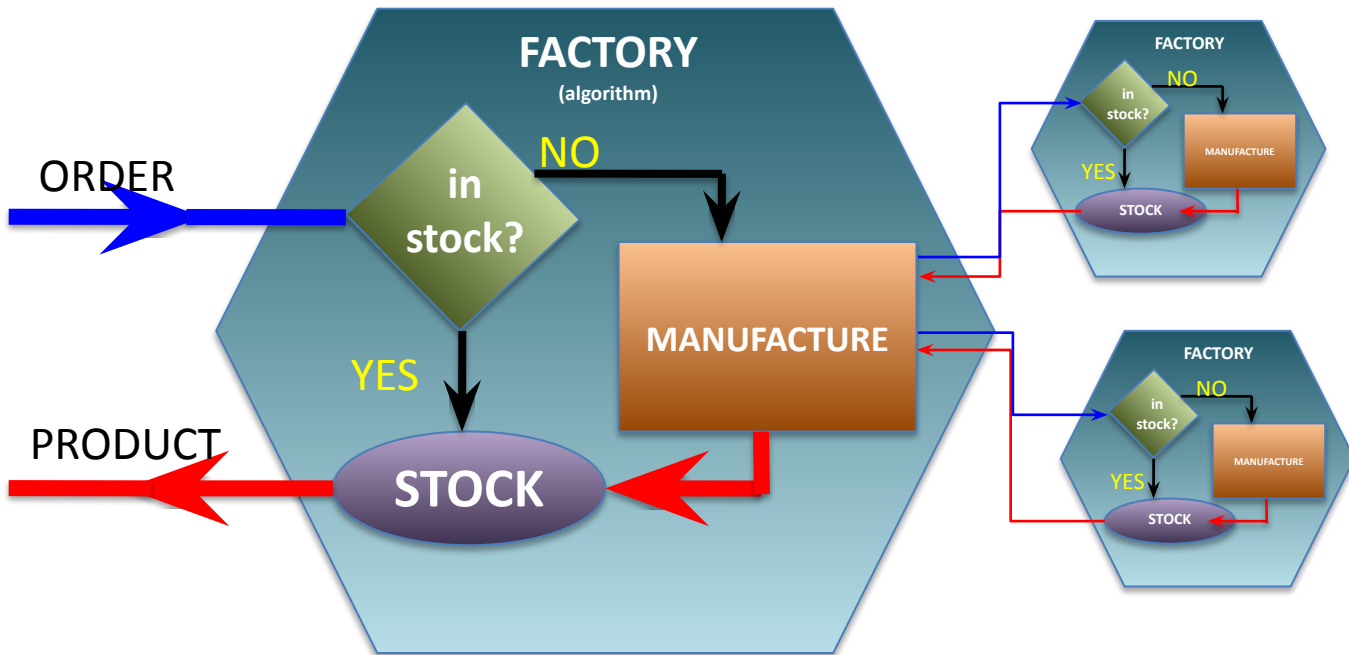


Reactive/Dataflow Programming

- Data is presented to arrow in the form of a queue
- Arrow transforms data and places it in downstream queue
- Minimal synchronization time spent in accessing queues
- Course tasks within arrow can eliminate most or all other synchronization points



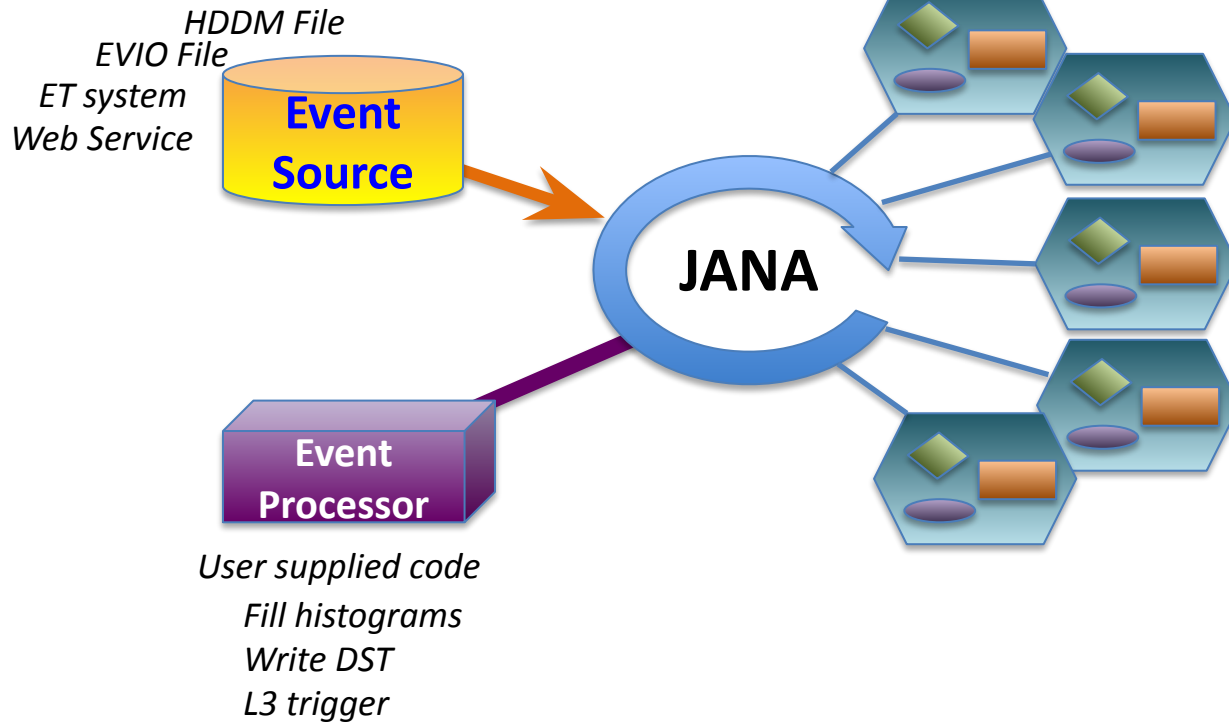
Factory Model



Data on demand = Don't do it unless you need it
Stock = Don't do it twice

**Conservation
of CPU cycles!**

Complete Event Reconstruction in JANA



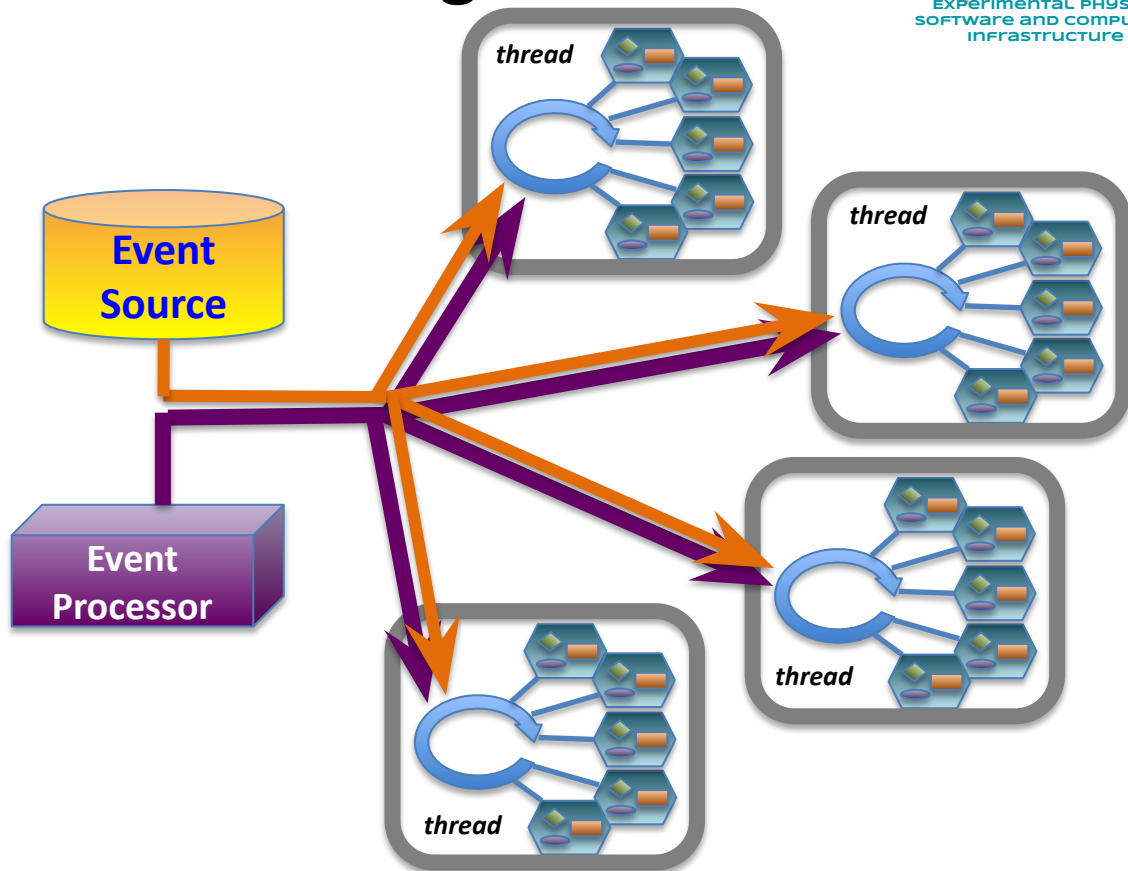
Framework has a layer that directs object requests to the factory that completes it

Multiple algorithms (factories) may exist in the same program that produce the same type of data objects

This allows the framework to easily redirect requests to alternate algorithms specified by the user at run time

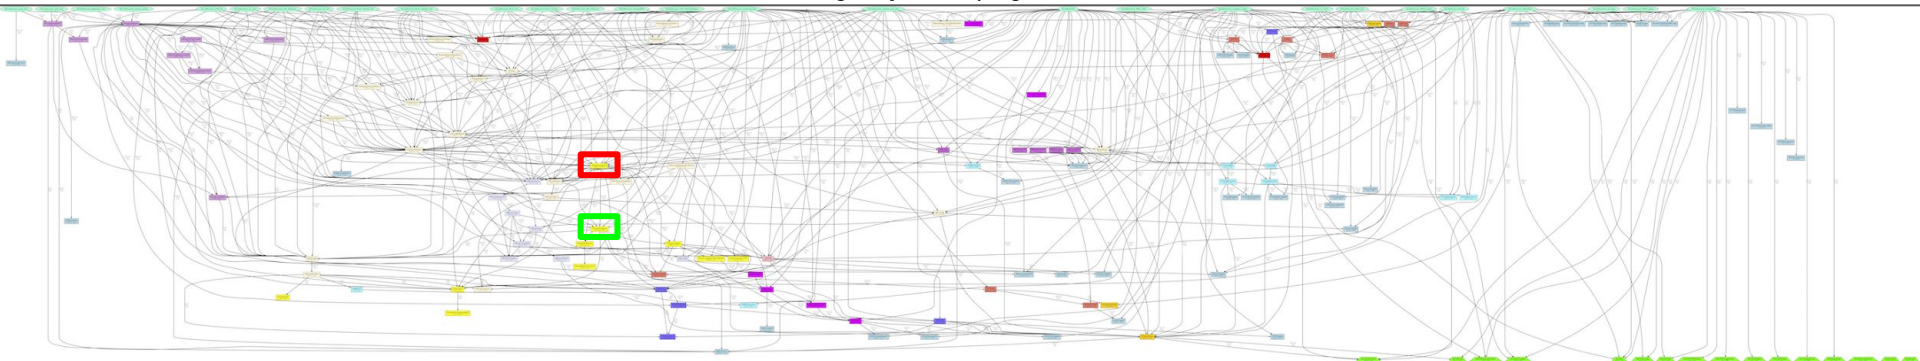
Multi-threading

- A complete set of factories is assigned to an event giving it exclusive use while that event is processed
- Factories only work with other factories in the same thread eliminating the need for expensive mutex locking within the factories
- All events are seen by all Event Processors (multiple processors can exist in a program)



Large experiments have complex call graphs

GlueX Reconstruction - automated rendering via janadot plugin



Run 42513:

Physics Production mode Trigger: FCAL_BCAL_PS_m9.conf

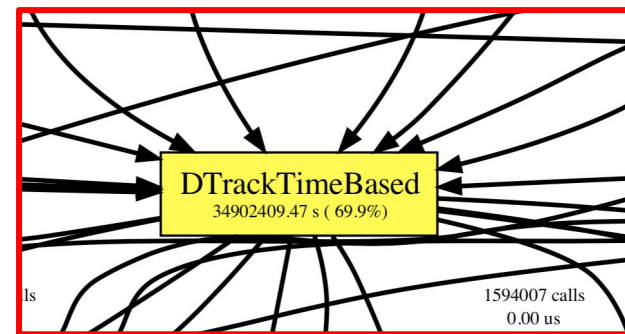

setup: hd_all.tsg

0/90 PERP 90

JD70-100 58um

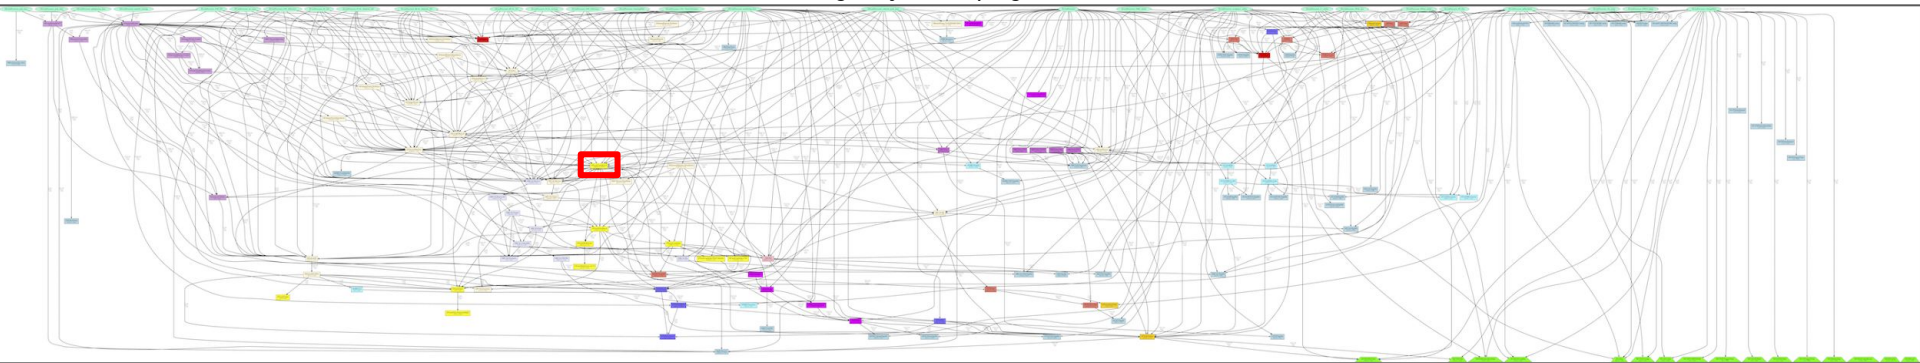
TPOL Be 75um

beam looks stable



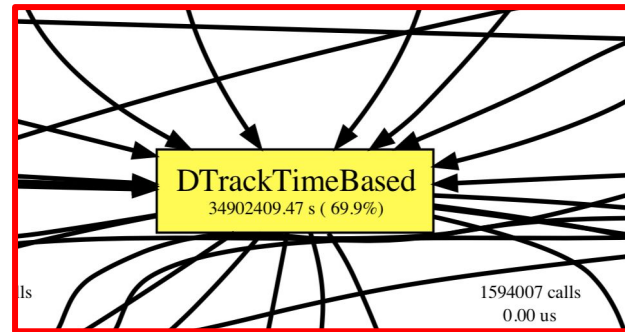
Large experiments have complex call graphs

GlueX Reconstruction - automated rendering via janadot plugin



Modular design:

- Factories (algorithms) need to know what they depend on
- Factories do **not** need to know what depends on them
- Dependencies do **not** need to be specified at higher level




Features Added in JANA2

- Better use of “modern” C++ features
 - thread model via C++ language (introduced in c++11)
 - lock guards
 - shared pointers
 - lambda functions
- Generalized use of threads (pool)
 - multiple queues
 - arrows (sequential or parallel)
- NUMA awareness
- Python API (both embedded and as an extension)

Multiple Affinity and Locality strategies

OS, chip type, memory architecture, and nature of job all can affect which model yields optimal performance



```
enum class  
AffinityStrategy {  
    None,  
    MemoryBound,  
    ComputeBound };
```

```
enum class  
LocalityStrategy {  
    Global,  
    SocketLocal,  
    NumaDomainLocal,  
    CoreLocal,  
    CpuLocal };
```

Configurable at run
time via Config.
Parameters

Gaudi Property = JANA Config. Parameter

Algorithms need parameters that can be configured at run time.

Gaudi stores configurable parameters in `Gaudi::Property` members

```
30 Gaudi::Property<double> m_timeResolution{this, "timeResolution", 10}; // todo : add units
31 Gaudi::Property<double> m_threshold{this, "threshold", 0. * Gaudi::Units::keV};

49 StatusCode sc = m_gaussDist.initialize(randSvc, Rndm::Gauss(0.0, m_timeResolution.value()));
```

JANA stores configurable parameters in central service and copies to local variables (e.g. members)

```
34 config.threshold = 0.0;
35 config.timeResolution = 8.0 * dd4hep::ns; // correct units?
36
37 app->SetDefaultParameter("RPOTS:ForwardRomanPotRawHits:threshold", config.threshold );
38 app->SetDefaultParameter("RPOTS:ForwardRomanPotRawHits:timeResolution", config.timeResolution );
39
```

JANA maintains list of all configuration parameters and their defaults.

- values can be dumped to config. file at end of job for use on subsequent jobs
- values that differ from defaults can be flagged
- values that have no implementation in code (typos) can be flagged

Connecting Algorithms

Gaudi: *reconstruction.py*
explicit piping

JANA: *factory C++*
implicit piping

```
289 ## Roman pots
290 ffi_romanpot_coll = SimTrackerHitsCollector(
291     "ffi_romanpot_coll",
292     inputSimTrackerHits=forward_romanpot_collections,
293     outputSimTrackerHits="ForwardRomanPotAllHits",
294 )
295 algorithms.append(ffi_romanpot_coll)
296 ffi_romanpot_digi = TrackerDigi(
297     "ffi_romanpot_digi",
298     inputHitCollection=ffi_romanpot_coll.outputSimTrackerHits,
299     outputHitCollection="ForwardRomanPotRawHits",
300     timeResolution=8,
301 )
302 algorithms.append(ffi_romanpot_digi)
303
304 ffi_romanpot_reco = TrackerHitReconstruction(
305     "ffi_romanpot_reco",
306     inputHitCollection=ffi_romanpot_digi.outputHitCollection,
307     outputHitCollection="ForwardRomanPotRecHits",
308 )
309 algorithms.append(ffi_romanpot_reco)
310
311 ffi_romanpot_parts = FarForwardParticles(
312     "ffi_romanpot_parts",
313     inputCollection=ffi_romanpot_reco.outputHitCollection,
314     outputCollection="ForwardRomanPotParticles",
315 )
316 algorithms.append(ffi_romanpot_parts)
```

```
46 //-----
47 // Process
48 void Process(const std::shared_ptr<const JEvent> &event) override {
49     // Get inputs
50     auto sim_hits1 = event->Get<edm4hep::SimTrackerHit>("ForwardRomanPotHits1");
51     auto sim_hits2 = event->Get<edm4hep::SimTrackerHit>("ForwardRomanPotHits2");
52     auto sim_hits_all = sim_hits1;
53     sim_hits_all.insert( sim_hits_all.end(), sim_hits2.begin(), sim_hits2.end() );
54
55     // Call Process for generic algorithm
56     auto rawhits = produce( sim_hits_all );
57
58     // Hand owner of algorithm objects over to JANA
59     Set(rawhits);
60 }
```

```
46 //-----
47 // Process
48 void Process(const std::shared_ptr<const JEvent> &event) override {
49     // Get inputs
50     auto rawhits = event->Get<edm4eic::RawTrackerHit>("ForwardRomanPotRawHits");
51
52     // Call Process for generic algorithm
53     std::vector<edm4eic::TrackerHit* > trackerhits;
54     for( auto rawhit : rawhits) trackerhits.push_back(produce(rawhit));
55
56     // Hand owner of algorithm objects over to JANA
57     Set(trackerhits);
58 }
```

Primary Key for Connecting Algorithms

Gaudi uses collection name (string) as the primary key

```

296 ffi_romanpot_digi = TrackerDigi(
297     "ffi_romanpot_digi",
298     inputHitCollection=ffi_romanpot_coll.outputSimTrackerHits,
299     outputHitCollection="ForwardRomanPotRawHits",
300     timeResolution=8,
301 )
    
```

python code

```

33 DataHandle<edm4hep::SimTrackerHitCollection> m_inputHitCollection{inputHitCollection, Gaudi::DataHandle::Reader,
34                                                                    this};
    
```

C++ code

JANA uses typeid as the primary key and string as secondary key ("tag")

```

50 auto sim_hits1 = event->Get<edm4hep::SimTrackerHit>("ForwardRomanPotHits1");
    
```

primary key

secondary key

*errors in primary key
are caught at compile
time.*

*errors in secondary key
are caught at run time.*

n.b. C++ linker does **not** link algorithms together. Run time list is searched using typeid

Summary

- JANA is a multithreaded framework project with nearly 2 decades of experience behind it
- JANA2 is a rewrite incorporating more modern coding and CS practices and improving on the original using lessons learned
 - Streaming DAQ and Heterogeneous hardware support strongly considered in redesign
- JANA2 has been selected for use with the EIC first detector (ePIC) and is currently being implemented there

Github: <https://github.com/JeffersonLab/JANA2>

Documentation: <https://jeffersonlab.github.io/JANA2/>

Example project: https://github.com/faustus123/EIC_JANA_Example

Publications:

<https://arxiv.org/abs/2202.03085> *Streaming readout for next generation electron scattering experiments*

<https://doi.org/10.1051/epiconf/202125104011> *Streaming Readout of the CLAS12 Forward Tagger Using TriDAS and JANA2*

<https://doi.org/10.1051/epiconf/202024501022> *JANA2 Framework for Event Based and Triggerless Data Processing*

<https://doi.org/10.1051/epiconf/202024507037> *Offsite Data Processing for the GlueX Experiment*

<https://iopscience.iop.org/article/10.1088/1742-6596/119/4/042018> *Multi-threaded event reconstruction with JANA*

<https://pos.sissa.it/070/062> *Multi-threaded event processing with JANA*

<https://iopscience.iop.org/article/10.1088/1742-6596/219/4/042011> *The JANA calibrations and conditions database API*

<https://iopscience.iop.org/article/10.1088/1742-6596/1525/1/012032> *JANA2: Multithreaded Event Reconstruction*

Backups

Python support in JANA2

As pure python script

python3 jana.py

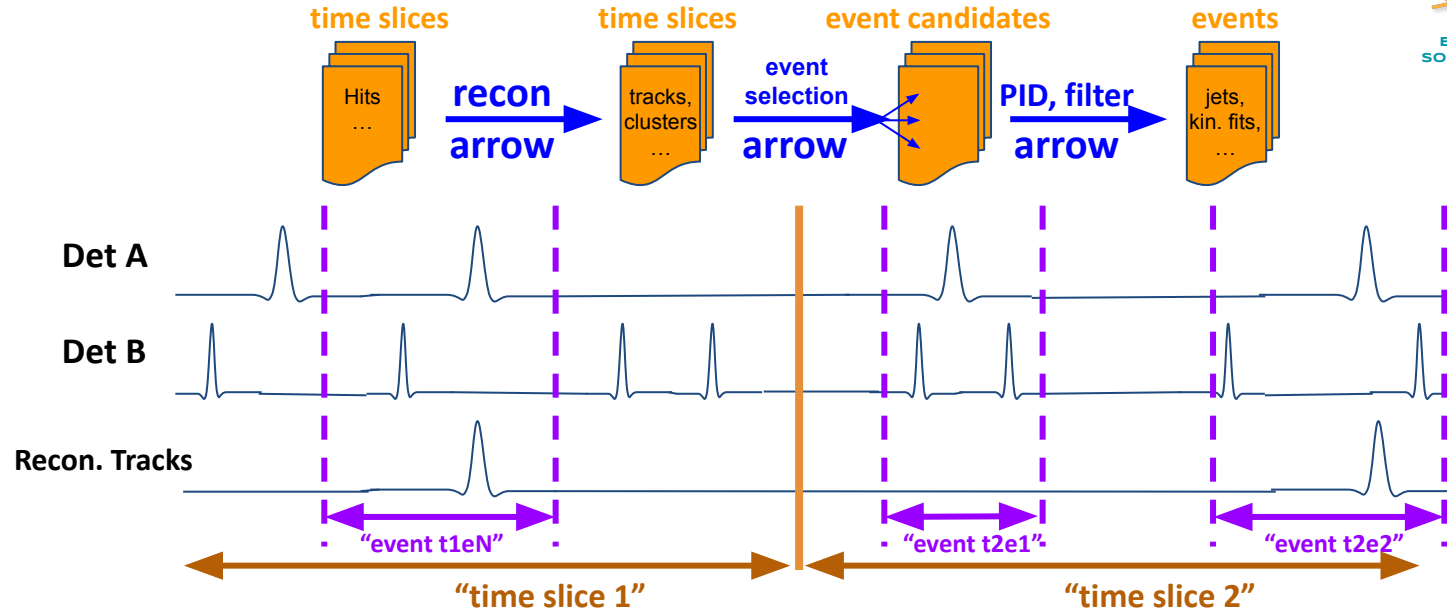
```
4 # This example JANA python script
5 import time
6 import jana
7
8 print('Hello from jana.py!!!')
9
10 # Turn off JANA's standard ticker so we can print our own updates
11 jana.SetTicker(False)
12
13 # Wait for 4 seconds before allowing processing to start
14 for i in range(1,5):
15     time.sleep(1)
16     print(" waiting ... %d" % (4-i))
17
18 # Start event processing
19 jana.Start()
20
21 # Wait for 5 seconds while processing events
22 for i in range(1,6):
23     time.sleep(1)
24     print(" running ... %d (Nevents: %d)" % (i, jana.GetNeventsProcessed()))
25
26 # Tell program to quit gracefully
27 jana.Quit()
```

As embedded interpreter

jana -PPLUGINS=janapy -PJANA_PYTHON_FILE=myfile.py

```
4 # This is a simple example JANA python script. It shows how to add plugins
5 # and set configuration parameters. Event processing will start once this
6 # script exits.
7 import jana
8
9 jana.AddPlugin('JTest')
10 jana.SetParameterValue('jana:nevents', 200)
11
12 jana.Run()
```

Streaming Data

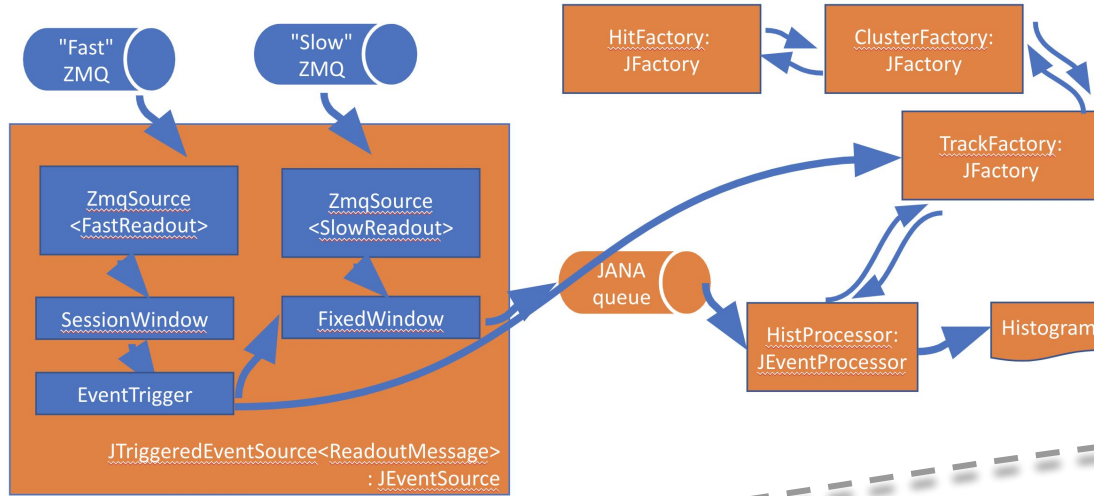


- **JANA2** has streaming readout features tested under multiple detector setups and in beam conditions*
- **EPSCI** has multiple experts working on streaming DAQ systems in same group as **JANA2** developers
- **EPSCI** works closely with the *JLab Fast Electronics Group* (Chris Cuevas) and partners routinely in performance testing in the DAQ Lab.

*Publications relevant to Streaming:

- <https://arxiv.org/abs/2202.03085> Streaming readout for next generation electron scattering experiments
- <https://doi.org/10.1051/epjconf/202125104011> Streaming Readout of the CLAS12 Forward Tagger Using TriDAS and JANA2
- <https://doi.org/10.1051/epjconf/202024501022> JANA2 Framework for Event Based and Triggerless Data Processing
- <https://doi.org/10.2172/1735849> Evaluation & Development of Algorithms & Techniques for Streaming Detector Readout

Streaming Readout



INDRA-ASTRA initiative:

- Software trigger
- Multi-flavored stream merging
- Event building

Support for Heterogeneous Hardware

- Sub-event level parallelism
 - Run ML on GPU or TPU



What the user needs to know:

```
auto tracks = jevent->Get<DTrack>();
```

```
for(auto t : tracks){
```

```
    // ... do something with const DTrack* t
```

```
}
```

```
vector<const *DTrack> tracks
```

Data on Demand => Software Trigger

**Event by event
decision on
whether to
activate a factory:**

Software triggers
may have multiple
“keep” or
“discard”
conditions that
may be probed in
order of CPU cost

```
// Getting hit objects is cheap so we check that first
auto NcaloHits = jevent->Get<CaloHit>().size();
if( NcaloHits > minCaloHits ){

    keep_event = true;

// Tracks factory only activated if not already keeping event
}else if( jevent->Get<Tracks>().size() > minTrackHits ) {

    keep_event = true;

}
```

If an alternate factory is desired:
(i.e. algorithm)

```
auto tracks = jevent->Get<DTrack>("MyTest");
```

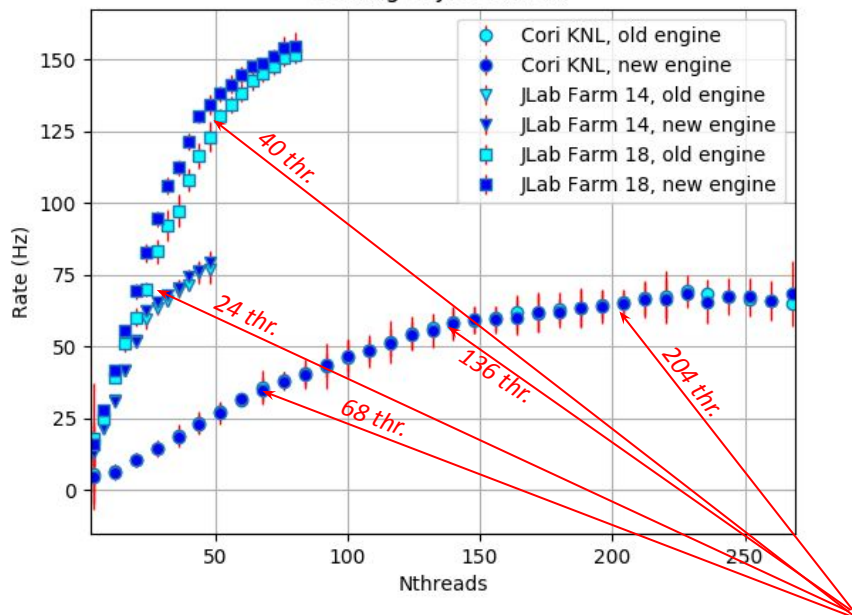
or, even better

set configuration parameter: **DTrack:DEFTAG=MyTest**


- Configuration parameters are set at run time
- NAME:DEFTAG is special and tells JANA to re-route ALL requests for objects of type NAME to the specified factory.

JANA2 Scaling Tests (JLab + NERSC)

Scaling of JTest (new)



JANA2 Block Size Scaling



kinks indicate hardware boundaries

TOPOLOGY STATUS

```

Thread team size [count]: 4
Total uptime [s]: 50.09
Uptime delta [s]: 0.5062
Completed events [count]: 587
Inst throughput [Hz]: 14
Avg throughput [Hz]: 11.7
Sequential bottleneck [Hz]: 335
Parallel bottleneck [Hz]: 11.9
Efficiency [0..1]: 0.986
  
```

Name	Status	Type	Par	Threads	Chunk	Thresh	Pending	Completed
dummy_evt_src	Running	Source	F	0	16	-	-	672
processors	Running	Sink	T	4	1	500	81	587

Name	Avg latency [ms/event]	Inst latency [ms/event]	Queue latency [ms/visit]	Queue visits [count]	Queue overhead [0..1]
dummy_evt_src	2.98	1.03	0.00415	42	8.71e-05
processors	337	321	0.00883	1450	6.48e-05

ID	Last arrow name	Useful time [ms]	Retry time [ms]	Idle time [ms]	Scheduler time [ms]	Scheduler visits [count]
0	processors	623	0	0	0.000576	76
1	processors	622	0	0	0.000624	138
2	processors	668	0	0	0.000553	131
3	processors	734	0	0	0.000606	125

JANA2 now has much better built-in diagnostics compared to the original JANA.

This helps pinpoint bottlenecks, especially in more complex systems

Boilerplate code generation

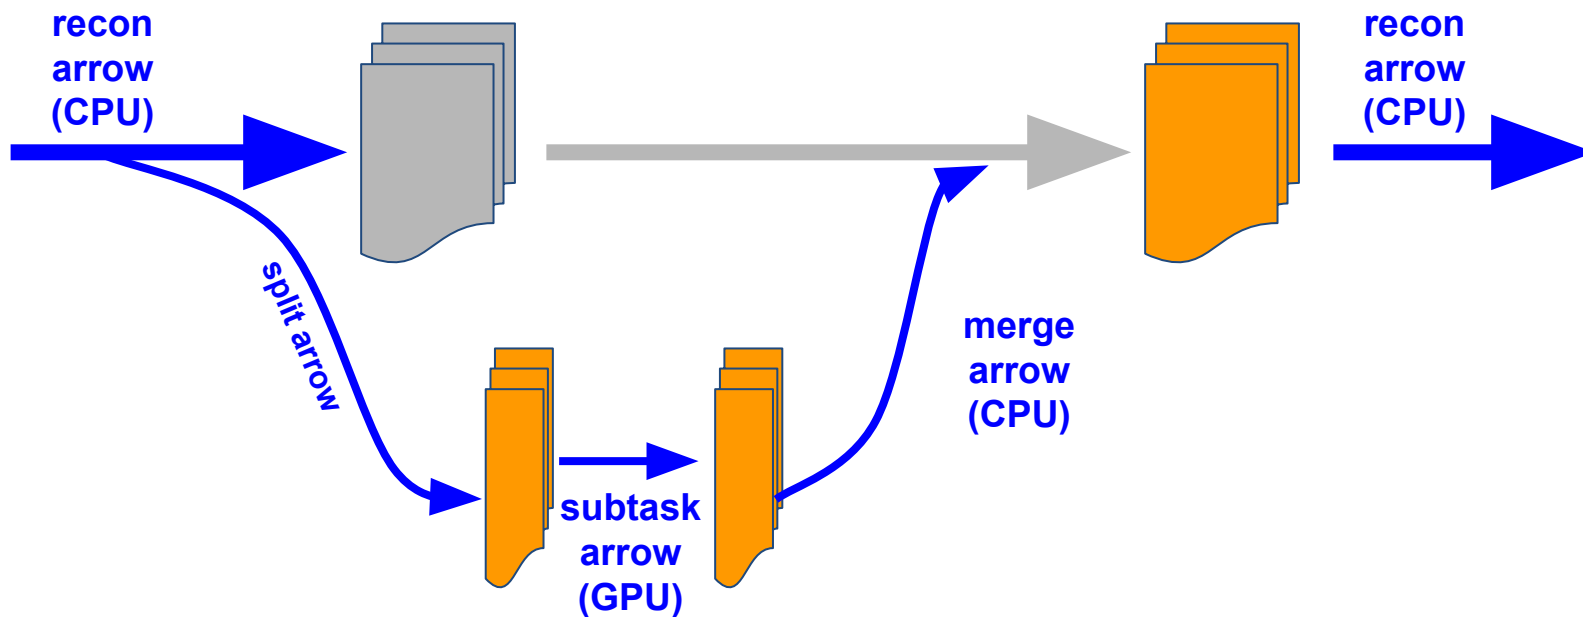
```
> jana-generate.py
Usage: jana-generate.py [-h|--help] [type] [args...]
      type: JObject JEventSource JEventProcessor RootEventProcessor JEventProcessorTest JFactory Plugin Project
```

```
> jana-generate.py --help
...
Plugin
Create a code skeleton for a plugin in its own directory. Takes the following positional arguments:
      name           The name of the plugin, e.g. "trk_eff" or "TrackingEfficiency"
[is_standalone] Is this a new project, or are we inside the source tree of an existing CMake project? (default=True)
[is_mini]       Reduce boilerplate and put everything in a single file? (default=True)
[include_root]  Include a ROOT dependency and stubs for filling a ROOT histogram? (default=True)

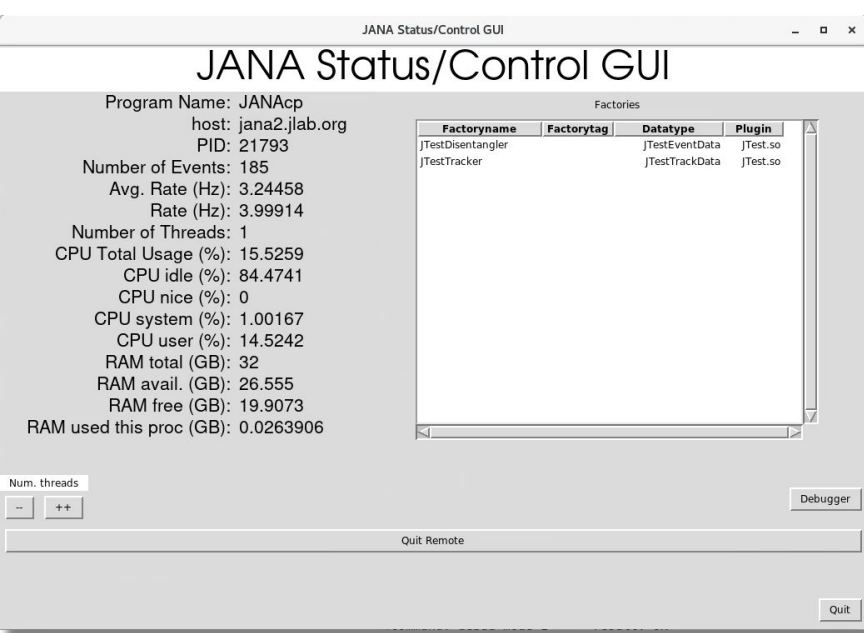
Example: `jana_generate.py Plugin TrackingEfficiency 1 0 0`
...
```

```
> jana-generate.py Plugin DaveTest
> ls DaveTest/
CMakeLists.txt  DaveTest.cc
> mkdir DaveTest/build
> cd DaveTest/build/
> cmake ..
...
> make install
[ 50%] Building CXX object CMakeFiles/DaveTest_plugin.dir/DaveTest.cc.o
[100%] Linking CXX shared library DaveTest.so
[100%] Built target DaveTest_plugin
Install the project...
-- Install configuration: ""
-- Installing: /Users/davidl/builds/JANA2/JANA2/plugins/DaveTest.so
```

Heterogeneous Hardware Support



Inspection Tools



JANA Status/Control GUI

Program Name: JANAcP
host: jana2.jlab.org
PID: 21793

Number of Events: 185
Avg. Rate (Hz): 3.24458
Rate (Hz): 3.99914
Number of Threads: 1
CPU Total Usage (%): 15.5259
CPU idle (%): 84.4741
CPU nice (%): 0
CPU system (%): 1.00167
CPU user (%): 14.5242
RAM total (GB): 32
RAM avail. (GB): 26.555
RAM free (GB): 19.9073
RAM used this proc (GB): 0.0263906

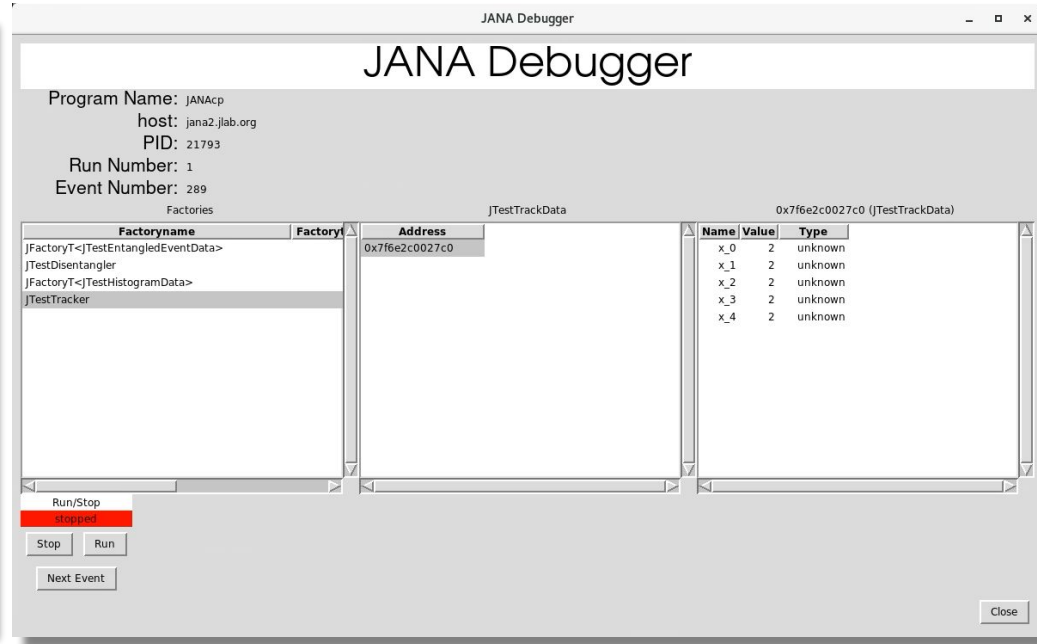
Factoryname	Factorytag	Datatype	Plugin
JTestDisentangler		JTestEventData	JTest.so
JTestTracker		JTestTrackData	JTest.so

Num. threads: -- ++

Debugger

Quit Remote

Quit



JANA Debugger

Program Name: JANAcP
host: jana2.jlab.org
PID: 21793
Run Number: 1
Event Number: 289

Factoryname	Factory	Address
JFactoryT-<JTestEntangledEventData>		0x7f6e2c0027c0
JTestDisentangler		
JFactoryT-<JTestHistogramData>		
JTestTracker		

Name	Value	Type
x_0	2	unknown
x_1	2	unknown
x_2	2	unknown
x_3	2	unknown
x_4	2	unknown

Run/Stop

Stop Run

Next Event

Close

```
> jana -Pplugins=JTest,janacontrol
```

← Add *janacontrol* plugin to any process

```
> jana-control.py [--host host] [--port port]
```

← Run GUI from remote (or same) node

JANA Command Line Debugging w/ gdb

```

davidl@jana2:JANA
File Edit View Search Terminal Help
Class name:  JTestParser
Sequential:  0

JANA: [INFO] Status: 0 events processed  0.0 Hz (0.0 Hz avg)

JANA: h
-----
Available commands
-----
pe  PrintEvent
pf  PrintFactories [filter_level <- {0,1,2,3}]
pfd PrintFactoryDetails fac_idx
po  PrintObjects fac_idx
po  PrintObject fac_idx obj_idx
pfp PrintFactoryParents fac_idx
pop PrintObjectParents fac_idx obj_idx
poa PrintObjectAncestors fac_idx obj_idx
vt  ViewAsTable
vj  ViewAsJson
x   Exit
h   Help
-----

JANA: p
```

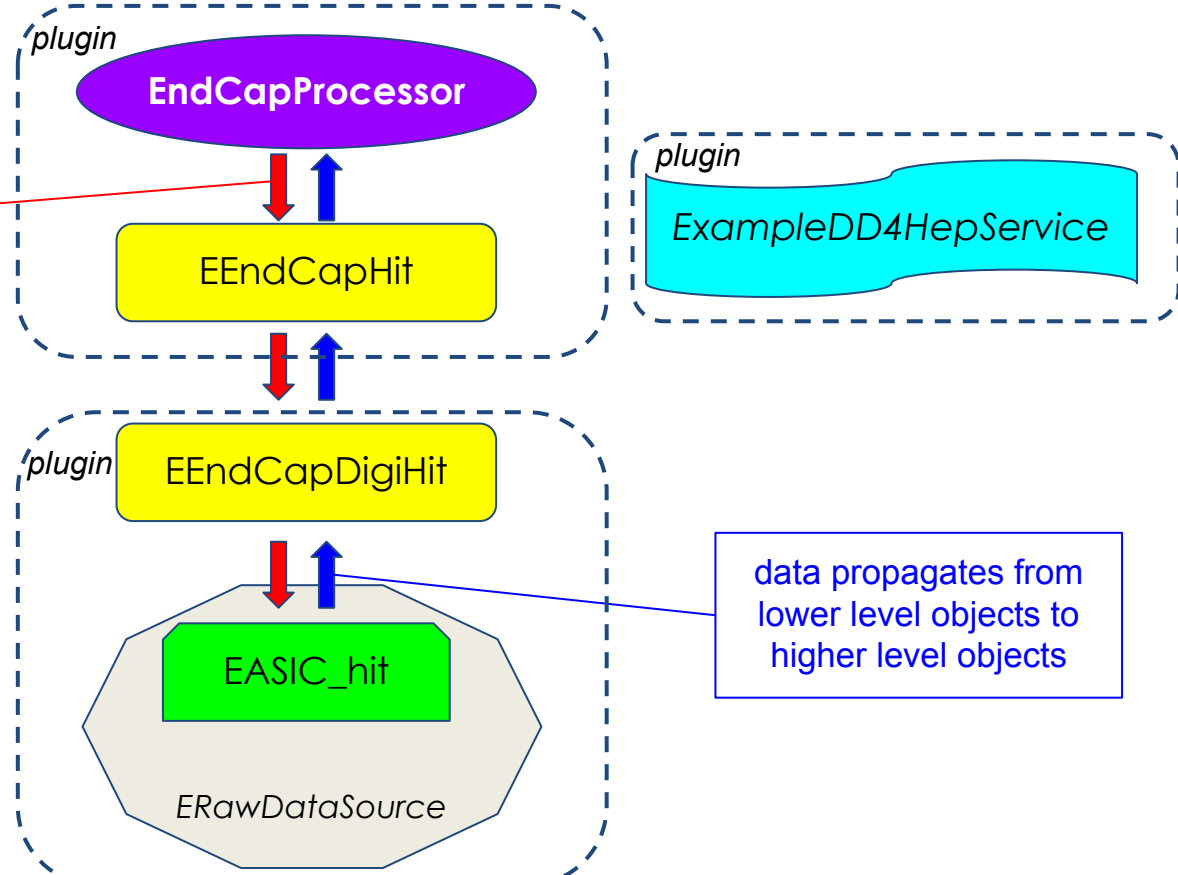
Certain JANA methods are written with the intention of being called from debugger.

This allows easier browsing from the framework point of view.

Example with Geometry Service

https://github.com/faustus123/EIC_JANA_Example

requests start with higher level objects and propagate to lower level objects

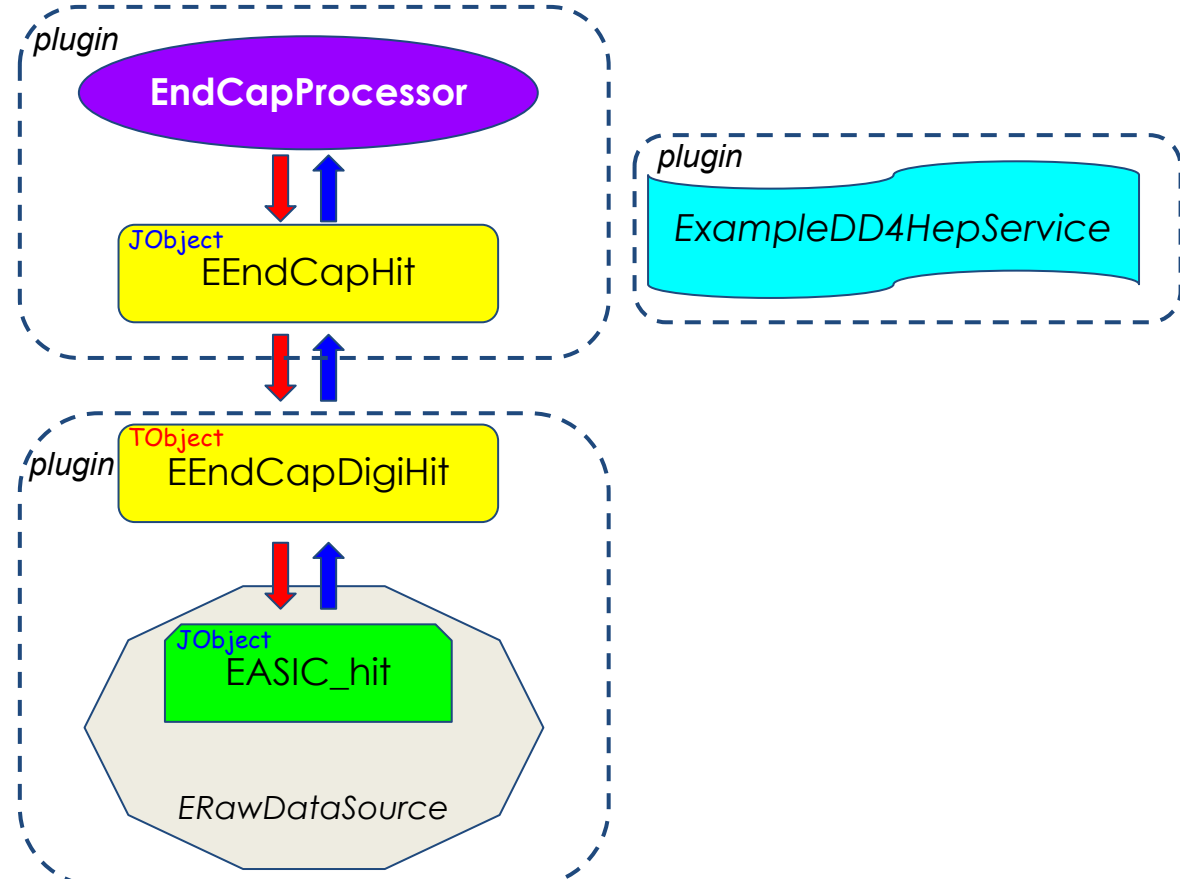


Wouter suggested example:
“...[take] a collection of hits and selecting those hits that are on a particular endcap tracking detector and have a position outside a minimum radial range.”

data propagates from lower level objects to higher level objects

Example with Mixed TObject and JObject

https://github.com/faustus123/EIC_JANA_Example/tree/TObject_example



JFactory_EEndCapHit

```
2 #ifndef _JFactory_EEndCapHit_h_
3 #define _JFactory_EEndCapHit_h_
4
5 #include <JANA/JFactoryT.h>
6 #include <ExampleDD4HepService/ExampleDD4HepService.h>
7 #include "EEndCapHit.h"
8
9 class JFactory_EEndCapHit : public JFactoryT<EEndCapHit> {
10
11     // Insert any member variables here
12
13 public:
14     JFactory_EEndCapHit();
15     void Init() override;
16     void ChangeRun(const std::shared_ptr<const JEvent> &event) override;
17     void Process(const std::shared_ptr<const JEvent> &event) override;
18
19 protected:
20     double min_radius;
21
22     const ExampleDD4HepService *geomservice=nullptr;
23
24 };
25
26 #endif // JFactory_EEndCapHit_h_
```

```
24 void JFactory_EEndCapHit::Init() {
25     auto app = GetApplication();
26
27     // Just for fun, create a configuration parameter named
28     // EndCap:min_radius so we can set the threshold at run time.
29     min_radius = 15.0;
30     app->SetDefaultParameter("EndCap:min_radius", min_radius, "The mini
31
32     /// Acquire geometry service pointer (see ExampleDD4HepService plugin)
33     geomservice = app->GetService<ExampleDD4HepService>().get();
34 }
```

boilerplate

added for this example

JFactory_EEndCapHit::Process

```
44 //-----
45 // Process
46 //-----
47 void JFactory_EEndCapHit::Process(const std::shared_ptr<const JEvent> &event) {
48
49     /// JFactories are local to a thread, so we are free to access and modify
50     /// member variables here. However, be aware that events are _scattered_to
51     /// different JFactory instances, not _broadcast_; this means that JFactory
52     /// instances only see _some_ of the events.
53
54     // The EEndCapDigiHit objects are made by a factory in the EICRawData plugin.
55     // That factory uses the low-level EASIC_hit objects coming from the event source
56     auto endcapdigiHits = event->Get<EEndCapDigiHit>();
57
58     // Loop over the EEndCapDigiHit objects and create calibrated hits
59     // objects with geometry info.
60     std::vector<EEndCapHit *> hits;
61     for( auto digihit : endcapdigiHits ){
62
63         auto pos = geomService->GetVTXPixelLocation( digihit->layer, digihit->chip, digihit->pixel );
64         auto r = pos.Perp();
65         if( r > min_radius ){
66
67             auto hit = new EEndCapHit();
68             hit->x = pos.X();
69             hit->y = pos.Y();
70             hit->z = pos.Z();
71             hit->t = ((double)digihit->t - 125.0)*2.50E-1; // Here we would apply calibrations read from DB
72             hits.push_back(hit);
73         }
74     }
75
76     /// Publish outputs
77     Set(hits);
78
79     // n.b. if we created additional types of objects we could also add them to the event using event->Insert() )
80 }
```

ExampleDD4HepService

```

26 class ExampleDD4HepService: public JService {
27
28
29
30
31
32 // The geometry service needs to be sensitive to the exact data being processed since subtle
33 // alignment changes or even significant changes to the detector could appear between one
34 // data set and the next. The most versatile system would allow data from multiple different
35 // geometry definitions to exist at the same time.
36 //
37 // For this to return the correct geometry, it needs information from the data stream itself
38 // on when it was acquired so it can access the correct DB. I do not try and add that
39 // complication here right now. I do demonstrate though that the JEvent reference would be
40 // passed in so that the needed info can be extracted. Note that this should not be called for
41 // every event, but rather from the ChangeRun method of a factory or processor indicating a
42 // new calibration region of the stream has been reached.
43 const dd4hep::Assembly* GetDD4hepAssembly(const std::shared_ptr<const JEvent> &event) const {
44
45     // Retrieve the correct Assembly based on when the given
46     // JEvent was acquired.
47
48     return _assembly;
49 }
50
51 // There is a lot of freedom in how this class could be organized. One is to simply provide a
52 // reference to the DD4hep Assembly object as above and let all of the algorithms speak "DD4hep".
53 // A more practical approach would be to augment that with some dedicated methods that answer
54 // common questions about the geometry for specific detectors. Here is an example of this:
55 TVector3 GetVTXPixelLocation( int layer, int chip, int pixel ) const {
56
57     // This is where the code to extract the location information given the layer,chip, and pixel
58     // values would reside. This could either be directly from the dd4hep reference or from some
59     // cached value.
60     assert( layer>=1 && layer<=9 );
61
62     double x = (double)chip*2.7; // Totally unrealistic. Just for demo
63     double y = (double)pixel*1.2; // Totally unrealistic. Just for demo
64     double z = z_layer[layer-1]; // Lookup table (this should actually be close to correct!)
65
66     return TVector3( x, y, z);
67 }
68
69 private:
70 dd4hep::Assembly *_assembly = nullptr;
71 double z_layer[9] = {-106.0, -79.0, -52.0, -25.0, 25.0, 49.0, 73.0, 106.0, 125.0};
72
73 };

```

Service is added to application with single line:

```

21 extern "C" {
22     void InitPlugin(JApplication *app) {
23         InitJANAPugin(app);
24         app->ProvideService( std::make_shared<ExampleDD4HepService>() );
25     }
26 }

```

The following are some notes I made a while back when trying to understand how JANA, Gaudi, and Fun4all approach the basic function of the framework. It is terribly incomplete, but may give some insight so I included it here in the backup slides.

Here I try and breakdown some example reconstruction code from ATHENA's juggler framework based on GAUDI. At the same time I try and compare this to what an equivalent JANA2 implementation would look like.

This is the first algorithm I looked at in the ATHENA repository and can be found here:

<https://eicweb.phy.anl.gov/EIC/juggler/-/blob/master/JugReco/src/components/SimpleClustering.cpp>

I looked at it first since the name "SimpleClustering" seemed like a good place to start.

SimpleClustering.cpp 6.21 KB

```
1  #include <algorithm>
2
3  #include "Gaudi/Property.h"
4  #include "GaudiAlg/GaudiAlgorithm.h"
5  #include "GaudiAlg/GaudiTool.h"
6  #include "GaudiAlg/Transformer.h"
7  #include "GaudiKernel/PhysicalConstants.h"
8  #include "GaudiKernel/RndmGenerators.h"
9  #include "GaudiKernel/ToolHandle.h"
10
11 #include "DDRec/CellIDPositionConverter.h"
12 #include "DDRec/Surface.h"
13 #include "DDRec/SurfaceManager.h"
14
15 #include "JugBase/DataHandle.h"
16 #include "JugBase/IGeoSvc.h"
17 #include "JugBase/UniqueID.h"
18
19 // Event Model related classes
20 #include "dd4pod/CalorimeterHitCollection.h"
21 #include "eicd/CalorimeterHitCollection.h"
22 #include "eicd/ClusterCollection.h"
23 #include "eicd/ProtoClusterCollection.h"
24 #include "eicd/RawCalorimeterHitCollection.h"
25
26 using namespace Gaudi::Units;
27
28 namespace Jug::Reco {
29
```

This is a preamble to the file. Nothing remarkable here.

```
30  /** Simple clustering algorithm.
31   *
32   * \ingroup reco
33   */
34  class SimpleClustering : public GaudiAlgorithm, AlgorithmIDMixin<> {
35  public:
```

Class is defined in implementation file in a Java-like way. This may be a stylistic choice, but definitely something allowed by GAUDI. Without a header file, the class cannot be directly used in code outside of this. Any use would have to come from properties of the class coming through one of its base classes.

```
177
178  DECLARE_COMPONENT(SimpleClustering)
179
180 } // namespace Jug::Reco
```

The class is declared to GAUDI by the DECLARE_COMPONENT call at the bottom of the file. This is defined through a few files but eventually gets to this file and the following line:

Gaudi/GaudiPluginService/include/Gaudi/PluginServiceV2.h

Registry::instance().add(id, { libraryName(), std::move(f), std::move(props) });

At this point I don't know if that is instantiating an object of this class or otherwise generating code that can be used to instantiate SimpleClustering objects later.

```
9 class SimpleClustering : public JFactoryT<Cluster> {
~
6 extern "C" {
7     void InitPlugin(JApplication *app) {
8         InitJANAPlugin(app);
9         app->Add(new JFactoryGeneratorT<SimpleClustering>());
10    }
11 }
```

The JANA equivalent here would be to create a class inheriting from JFactory and then report that to JANA by instantiating a JFactoryGenerator class via template.

JANA will use the JFactoryGenerator class to instantiate multiple SimpleClustering objects later.

```

30  /** Simple clustering algorithm.
31  *
32  * \ingroup reco
33  */
34  class SimpleClustering : public GaudiAlgorithm, AlgorithmIDMixin<> {
35  public:
36      using RecHits = eic::CalorimeterHitCollection;
37      using ProtoClusters = eic::ProtoClusterCollection;
38      using Clusters = eic::ClusterCollection;
39
40      DataHandle<RecHits>      m_inputHitCollection{"inputHitCollection", Gaudi::DataHandle::Reader, this};
41      DataHandle<ProtoClusters> m_outputProtoClusters{"outputProtoCluster", Gaudi::DataHandle::Writer, this};
42      DataHandle<Clusters>    m_outputClusters{"outputClusterCollection", Gaudi::DataHandle::Writer, this};
43
44      Gaudi::Property<std::string> m_mcHits{this, "mcHits", ""};
45
46      Gaudi::Property<double> m_minModuleEdep{this, "minModuleEdep", 5.0 * MeV};
47      Gaudi::Property<double> m_maxDistance{this, "maxDistance", 20.0 * cm};
48
49      /// Pointer to the geometry service
50      SmartIF<IGeoSvc> m_geoSvc;
51
52      // Monte Carlo particle source identifier
53      const int32_t m_kMonteCarloSource{uniqueID<int32_t>("mcparticles")};
54      // Optional handle to MC hits
55      std::unique_ptr<DataHandle<dd4pod::CalorimeterHitCollection>> m_inputMC;
56
57      SimpleClustering(const std::string& name, ISvcLocator* svcLoc)
58          : GaudiAlgorithm(name, svcLoc)
59            , AlgorithmIDMixin<>(name, info()) {
60          declareProperty("inputHitCollection", m_inputHitCollection, "");
61          declareProperty("outputProtoClusterCollection", m_outputClusters, "Output proto clusters");
62          declareProperty("outputClusterCollection", m_outputClusters, "Output clusters");
63      }

```

Convenience declarations

Data objects in Gaudi are contained in DataHandle templated classes. It looks like these wrappers are instantiated with a pointer to the algorithm object they belong to.

Gaudi Property objects look to similarly wrap variables in a class and register it with the Gaudi system. This will allow Gaudi to know and set these values externally.

The JANA equivalent to these properties are configuration parameters. It is not clear if Gaudi expects to change these after event processing has started, but in JANA they are not expected to change. A comparable JANA call would be:

```
double m_minModuleEdep = 5.0 * MeV;
app->SetDefaultParameter("minModuleEdep", m_minModuleEdep, "...");
```

typo?

Input and output objects are declared explicitly in the constructor. It is not clear why this is needed in addition to the DataHandle constructors above.

```

65 StatusCode initialize() override
66 {
67     if (GaudiAlgorithm::initialize().isFailure()) {
68         return StatusCode::FAILURE;
69     }
70     // Initialize the MC input hit collection if requested
71     if (m_mcHits != "") {
72         m_inputMC =
73         std::make_unique<DataHandle<dd4pod::CalorimeterHitCollection>>(m_mcHits, Gaudi::DataHandle::Reader, this);
74     }
75     m_geoSvc = service("GeoSvc");
76     if (!m_geoSvc) {
77         error() << "Unable to locate Geometry Service. "
78             << "Make sure you have GeoSvc and SimSvc in the right order in the configuration." << endmsg;
79         return StatusCode::FAILURE;
80     }
81     return StatusCode::SUCCESS;
82 }

```

Gaudi initialization method. This returns a value indicating if the initialization succeeds or fails.

Here, a string property of the class is used to determine if an input container should be made for MC hits.

```

14 void SimpleClustering::Init() {
15     auto app = GetApplication();
16     // Acquire any parameters
17     // app->GetParameter("parameter_name", m_destination);
18     // Acquire any services
19     // m_service = app->GetService<ServiceT>();
20     // Set any factory flags
21     // SetFactoryFlag(JFactory_Flags_t::NOT_OBJECT_OWNER);
22 }
23 }
24 }
25 }

```

JANA initialization method. Unlike Gaudi, JANA does not emit a return value. In JANA, Init() is only called at event processing time if/when an algorithm is first used and so it is assumed to be required. Fatal errors in the Init() method are expected to emit errors to the logging service and to tell the application to quit via a call to app->Quit(). One may also explicitly set an exit code with app->SetExitCode(val).


```

84  StatusCode execute() override
85  {
86      // input collections
87      const auto& hits = *m_inputHitCollection.get();
88      // Create output collections
89      auto& proto = *m_outputProtoClusters.createAndPut();
90      auto& clusters = *m_outputClusters.createAndPut();
91      // Optional MC data
92      const dd4pod::CalorimeterHitCollection* mcHits = nullptr;
93      if (m_inputMC) {
94          mcHits = m_inputMC->get();
95      }
96
97      std::vector<std::pair<uint32_t, eic::ConstCalorimeterHit>> the_hits;
98      std::vector<std::pair<uint32_t, eic::ConstCalorimeterHit>> remaining_hits;
99
100

```

This is the top of the execute() method which is called for every event for which the algorithm is active. The first lines are used to get the inputs for the algorithm and to create the output containers for the algorithm.

This mechanism uses the existence of a container that may or may not have been created in the init() method to determine whether to get the actual hits into the container.

JANA method that is called for every event.

```

3  void SimpleCluster_factory::Process(const std::shared_ptr<const JEvent> &jevent){
4
5      auto calohits = jevent->Get<DFCALHit>(); // Get input objects
6
7      // ... Create cluster objects ...
8      {
9          auto cluster = new DFCALCluster( a, b, c );
10         for( auto hit : myhits )cluster->AddAssociatedObject( hit );
11         Insert( cluster ); //pass ownership to framework
12     }
13 }

```

Input objects obtained as vector<const DFCALHit*> calohits

Algorithm creates cluster objects and “Inserts” them into the event using the Insert() method. One could also fill a local std::vector<> of pointers and publish those with the Set() method.

If the DFCALCluster class inherits from JObject, then the AssociatedObject mechanism can be used. This allows the framework to know about which hit objects were used to make the cluster.

Here is a comparison with Fun4All. This is taken from the following:

<https://github.com/ECCE-EIC/coresoftware/blob/master/offline/packages/CaloReco/RawClusterBuilderFwd.h>

I wanted to use another calorimeter clustering algorithm and this was the best I could locating with a quick search.

To start with, I should note that some of the code dealing with this is spread over a few classes:

RawClusterDefs	←	<i>Namespace. Defines RawClusterDefs::keytype</i>
RawCluster	←	<i>Inherits from PHObject</i>
RawClusterContainer	←	<i>Inherits from PHObject</i>
RawClusterBuilderFwd	←	<i>Inherits from SubsysReco</i>

```
1  #ifndef CALOBASE_RAWCLUSTERDEFS_H
2  #define CALOBASE_RAWCLUSTERDEFS_H
3
4  namespace RawClusterDefs
5  {
6      typedef unsigned int keytype;
7  }
8
9  #endif
```

This is just a namespace used to define the keytype used for the RawCluster objects. Presumably this is useful for object persistence since the unique id can be reproduced if the data were replayed.

JANA has removed support for object ids in JANA2. This is due to almost never being used in JANA1. This is likely due to the heavy use of pointers which also provide unique ids within the event, but don't require lookup tables to get at the object data.

```

14 class RawClusterContainer : public PHObject
15 {
16     public:
17         typedef std::map<RawClusterDefs::keytype, RawCluster *> Map;
18         typedef Map::iterator Iterator;
19         typedef Map::const_iterator ConstIterator;
20         typedef std::pair<Iterator, Iterator> Range;
21         typedef std::pair<ConstIterator, ConstIterator> ConstRange;
22
23         RawClusterContainer() {}
24         ~RawClusterContainer() override {}
25
26         void Reset() override;
27         int isValid() const override;
28         void identify(std::ostream &os = std::cout) const override;
29
30         ConstIterator AddCluster(RawCluster *clus);
31
32         RawCluster *getCluster(const RawClusterDefs::keytype id);
33         const RawCluster *getCluster(const RawClusterDefs::keytype id) const;
34
35         //! return all clusters
36         ConstRange getClusters(void) const;
37         Range getClusters(void);
38         const Map &getClustersMap() const { return _clusters; }
39         Map &getClustersMap() { return _clusters; }
40
41         unsigned int size() const { return _clusters.size(); }
42         double getTotalEdep() const;
43
44     protected:
45         Map _clusters;

```

The `RawClusterContainer` class is interesting because it really serves as a customized container class for `RawCluster` objects. It has several methods like `AddCluster`, `getCluster`, `getClusters`, ... that include the word “cluster” in their names. These do not seem to be doing anything special that any other container class would not already be doing. It is unclear why a more general (templated) container class is not used which could provide more uniformity in the code.

n.b. `getTotalEdep()` looks to be the only method that has functionality that would not be provided by a generic container class.

In JANA, the `JFactory` (i.e. algorithm) class that produces the data objects owns them and serves the combined purpose of the `RawClusterContainer` and `RawClusterBuilderFwd` classes. The `JFactory` class is actually a template itself where the template parameter is the specific type of primary data object the factory produces.

n.b. More than one object type can be produced by a `JFactory`. The supplementary types would use `Insert()` to add them to the event and would no longer be owned by the factory. This would make no difference to the end user. The emphasis on having a factory produce a single, primary object type is meant to encourage modularity in the overall design by having more, smaller algorithms.