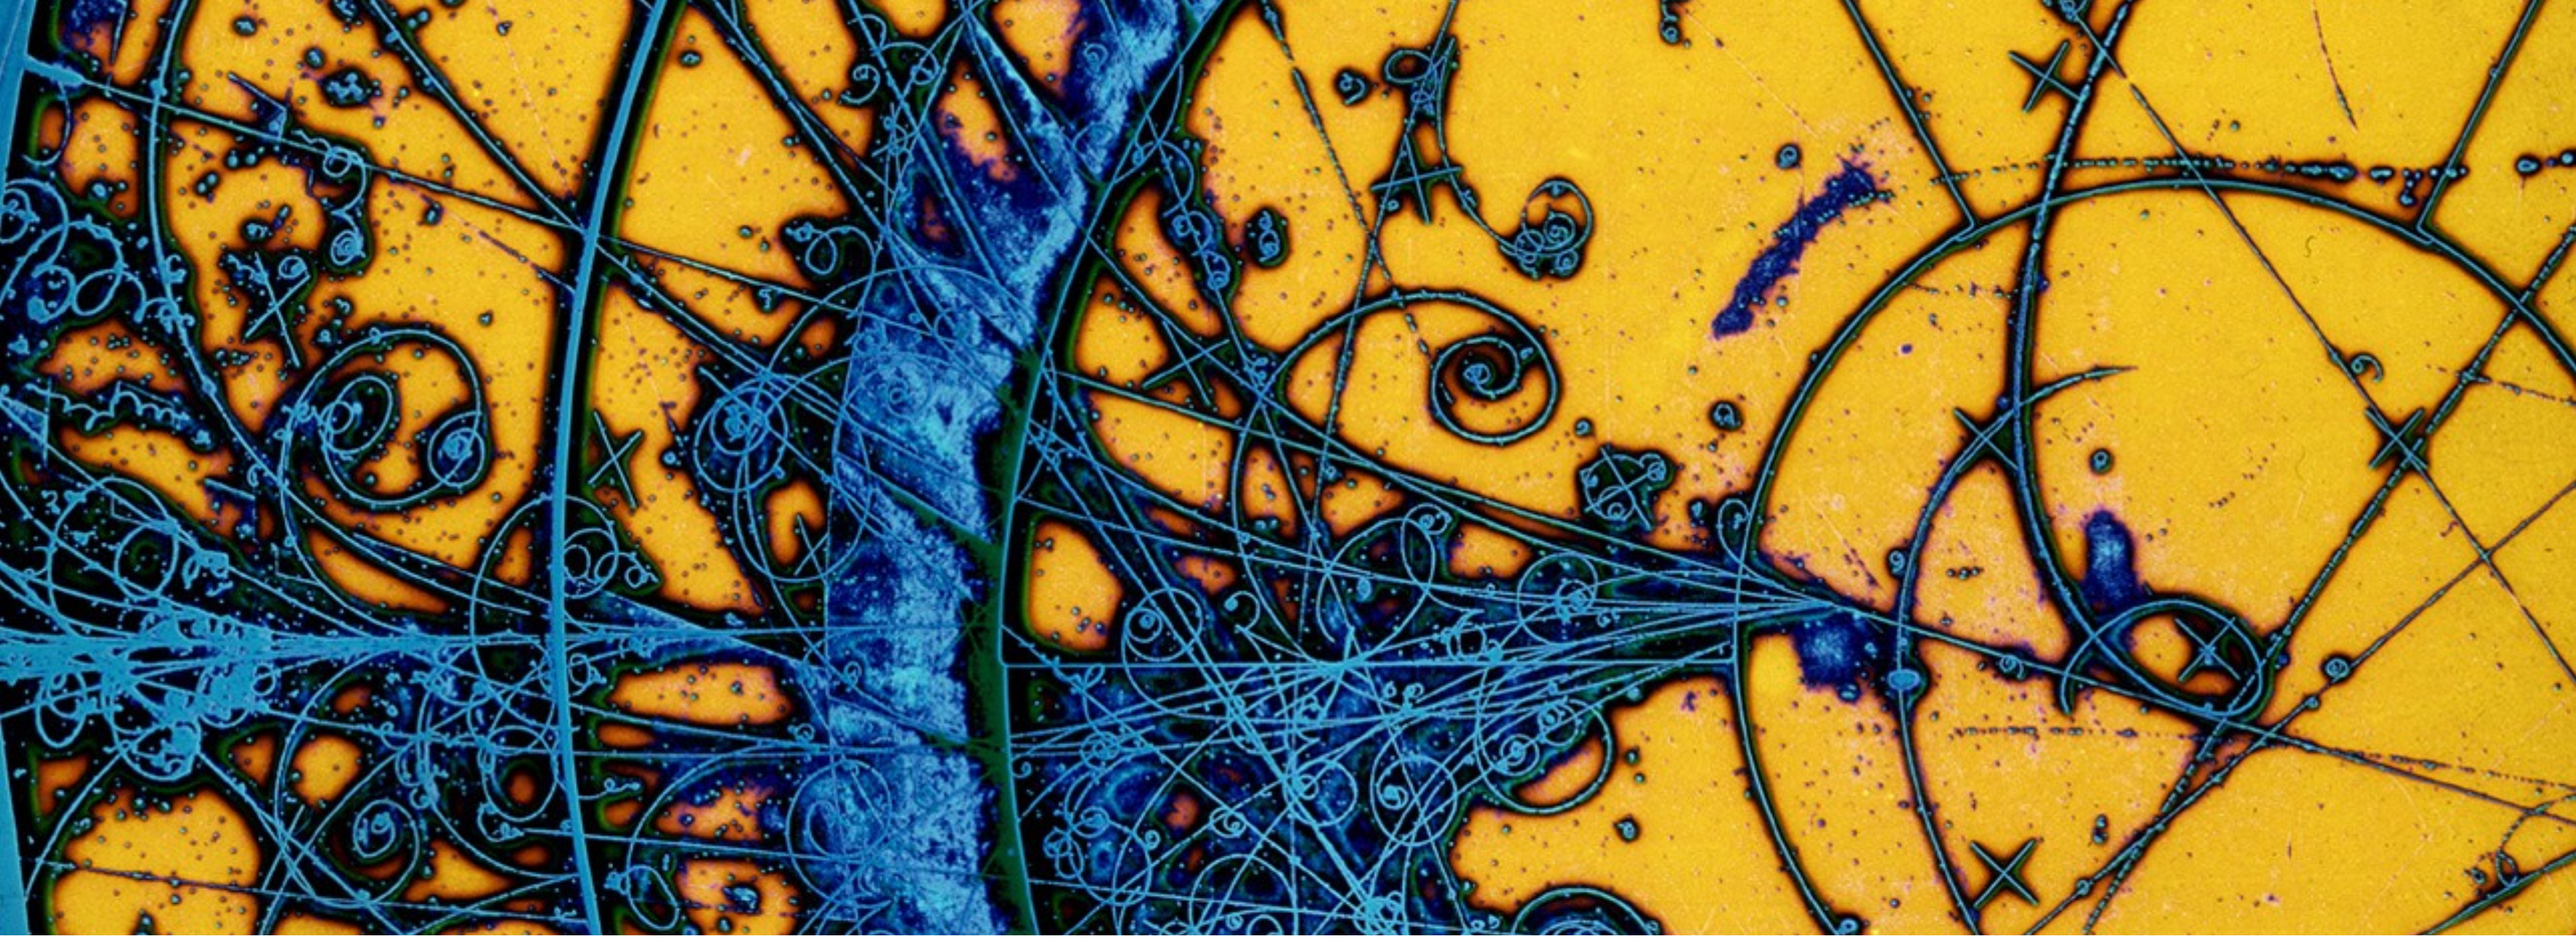

# LbMCSubmit: A new flexible and scalable request submission system for LHCb simulation

**Chris Burr** on behalf of the LHCb Computing and Simulation projects

9th May 2023

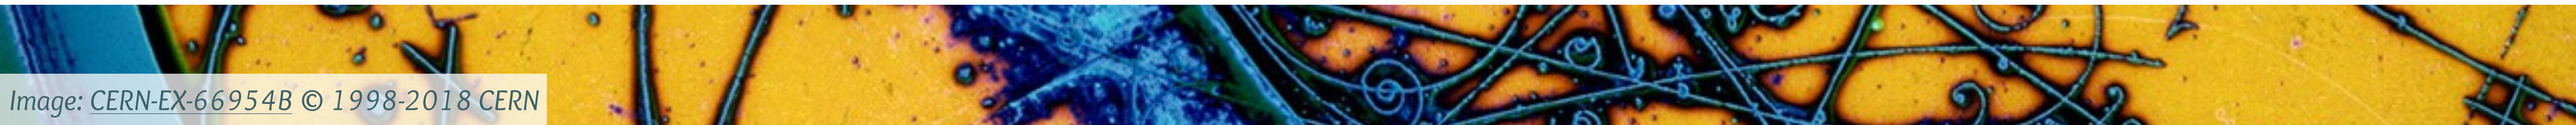

# Simulation in LHCb

- Current simulation requests have to span from 2011 to future upgrades
- Wide variety of needs - varied physics program and detector studies
  - Flavour physics analyses typically need many specific samples
  - Typically not much overlap between analyses
- Massive compute requirements
  - An increasing number of “fast” simulation options available
- Thousands of “basic” configurations (without considering decays)
  - Multiply by many thousands of “event types”
  - Effectively infinite possible customisations

**How can we handle this complexity without making things complicated?**

# Processing chain

- Simulated LHCb data is typically processed by the full processing chain
  - Typically ~8 separate applications to run
  - Lots of commonality but not trivial to isolate

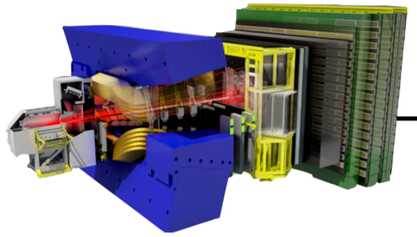

**Trigger**
Moore

**Reconstruction**
Brunel

**Stripping**
DaVinci

**Storage**
CHARM.MDST | DIMUON.DST | ...

**Digitisation**
Boole
MC

**Particle simulation**
Gauss
MC

**Generation**
Pythia/POWHEG/...

**Decay**
EvtGen

**Propagation**
Geant4

*Representative of Run 1+2 proton-proton running, exact dataflow varies by year and analyst needs*

# Past state

*See CHEP 2015: https://cds.cern.ch/record/2013233*

- Web application for submitting productions
  - Based around “models” which are prepared by experts
  - A model is made up of a series of “steps” (correspond to a executable + arguments)
  - Liaisons* clone the model and make small edits for specific productions

- Worked well for many years but shows problems
  - Large maintenance burden on already busy experts
  - Lack of frontend developer expertise for improvements

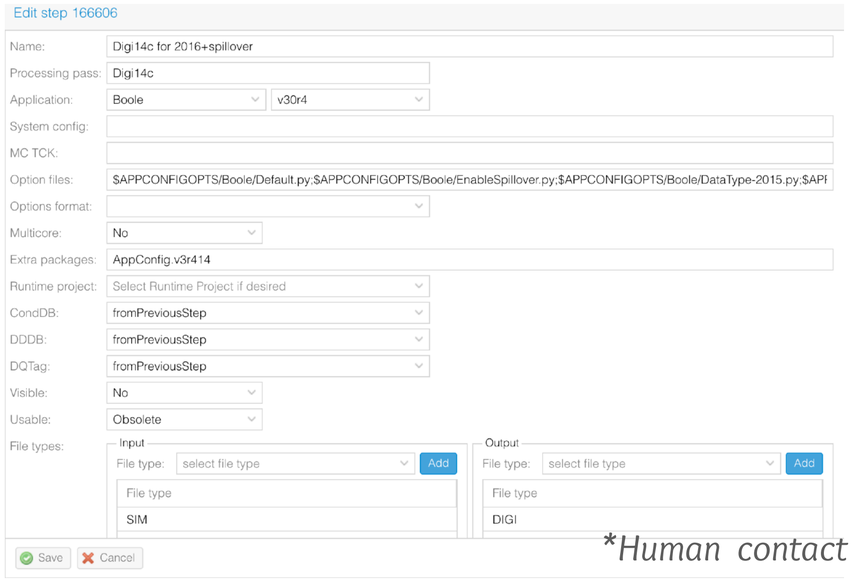


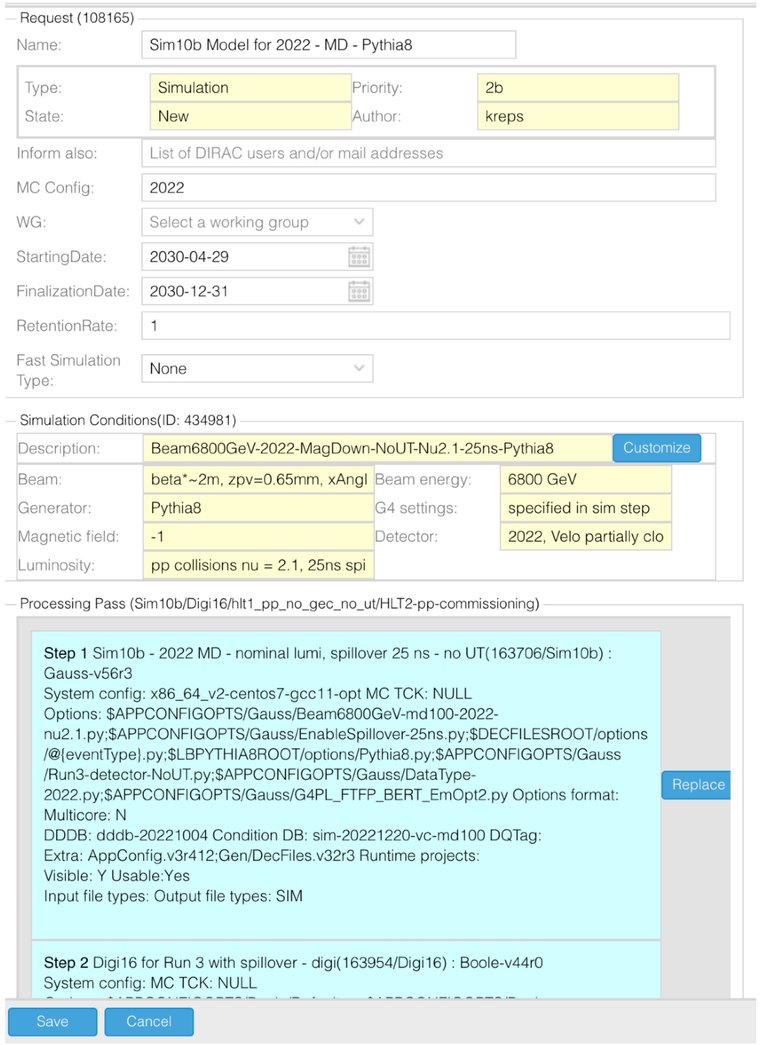


**Human contact between the simulation project and a specific area of the LHCb physics programme*

# What do we actually want?

- 95%+ of simulation requests can be summarised in a few words:

**I want 100,000 events per magnet polarity of**

**D+ → K- pi+ pi+**
**D*+ → D0 → (pi+ pi-) pi-**

**in 2012 and 2016 proton-proton conditions**

- Should result in 8 requests being created in LHCbDIRAC (the LHCb grid middleware)
  - 2 magnet polarities × 2 data taking years × 2 event types
- What does LHCbDIRAC actually need to know?

# What does LHCbDIRAC need to know?

```
- type: Simulation
  priority: 2a
  name: Ds2KKpi 2012 pp MagUp
  mc_config_version: '2012'
  sim_condition: Beam4000GeV-2012-MagUp-Nu2.5-Pythia8
  author: null
  inform:
  - auser
  - firstname.surname@cern.ch
  comment: ''
  wg: Charm
  retention_rate: 1.0
  event_types:
  - id: '21103030'
    num_events: 100000
    num_test_events: 10
  - id: '27163001'
    num_events: 100000
    num_test_events: 10
  steps:
  - name: Sim09l - 2012 - MagUp - Pythia8
    processing_pass: Sim09l
    visible: true
    application:
      name: Gauss
      version: v49r24
      binary_tag: x86_64-slc6-gcc48-opt
    options:
    - $APPCONFIGOPTS/Gauss/Sim08-Beam4000GeV-mu100-2012-nu2.5.py
    - $APPCONFIGOPTS/Gauss/DataType-2012.py
    - $APPCONFIGOPTS/Gauss/RICHRandomHits.py
    - $APPCONFIGOPTS/Gauss/NoPacking.py
    - $DECFILESROOT/options/@{eventType}.py
    - $LBPYTHIA8ROOT/options/Pythia8.py
    - $APPCONFIGOPTS/Gauss/G4PL_FTFP_BERT_EmNoCuts.py
    data_pkgs:
    - AppConfig.v3r412
    - Gen/DecFiles.v30r87
    input: []
    output:
    - type: SIM
      visible: false
    dbtags:
      DDDB: dddb-20170721-2
      CondDB: sim-20160321-2-vc-mu100
  - name: Digi14c for 2012
    processing_pass: Digi14c
    visible: false
    application: Boole/v30r4
    options:
    - $APPCONFIGOPTS/Boole/Default.py
    - $APPCONFIGOPTS/Boole/DataType-2012.py
    - $APPCONFIGOPTS/Boole/NoPacking.py
    - $APPCONFIGOPTS/Boole/Boole-SetOdinRndTrigger.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: SIM
      visible: false
    output:
    - type: DIGI
      visible: false
  - name: L0 emulation for 2012 - TCK 0x0045
    processing_pass: L0Trig0x0045
    visible: false
    application: Moore/v24r4
    options:
    - $APPCONFIGOPTS/L0App/L0AppSimProduction.py
    - $APPCONFIGOPTS/L0App/L0AppTCK-0x0045.py
    - $APPCONFIGOPTS/L0App/DataType-2012.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DIGI
      visible: false
    output:
    - type: DIGI
      visible: false
    options_format: l0app
  - name: TCK-0x409f0045 Flagged for 2012
    processing_pass: Trig0x409f0045
    visible: true
    application:
      name: Moore
      version: v14r8p1g1
      binary_tag: x86_64-slc5-gcc46-opt
    options:
    - $APPCONFIGOPTS/Moore/MooreSimProductionForSeparateL0AppStep.py
    - $APPCONFIGOPTS/Conditions/TCK-0x409f0045.py
    - $APPCONFIGOPTS/Moore/DataType-2012.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DIGI
      visible: false
    output:
    - type: DIGI
      visible: false
  - name: Reco14c for MC 2012
    processing_pass: Reco14c
    visible: true
    application:
      name: Brunel
      version: v43r2p13
      binary_tag: x86_64-slc5-gcc46-opt
    options:
    - $APPCONFIGOPTS/Brunel/DataType-2012.py
    - $APPCONFIGOPTS/Brunel/MC-WithTruth.py
    - $APPCONFIGOPTS/Brunel/Sim09-Run1.py
    - $APPCONFIGOPTS/Persistency/DST-multipleTCK-2012.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DIGI
      visible: false
    output:
    - type: DST
      visible: false
  - name: Stripping21NoPrescalingFlagged for 2012
    processing_pass: Stripping21NoPrescalingFlagged
    visible: true
    application: DaVinci/v36r1p5
    options:
    - $APPCONFIGOPTS/DaVinci/DV-Stripping21-Stripping-MC-NoPrescaling.py
    - $APPCONFIGOPTS/DaVinci/DV-RedoCaloPID-Stripping21.py
    - $APPCONFIGOPTS/DaVinci/DataType-2012.py
    - $APPCONFIGOPTS/DaVinci/InputType-DST.py
    - $APPCONFIGOPTS/DaVinci/DV-Stripping-MC-muDST.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DST
      visible: false
    output:
    - type: SIM
      visible: false
    dbtags:
      DDDB: dddb-20170721-2
      CondDB: sim-20170721-2-vc-md100
  - name: Digi14c for 2016+spillover
    processing_pass: Digi14c
    visible: false
    application: Boole/v30r4
    options:
    - $APPCONFIGOPTS/Boole/Default.py
    - $APPCONFIGOPTS/Boole/EnableSpillover.py
    - $APPCONFIGOPTS/Boole/DataType-2015.py
    - $APPCONFIGOPTS/Boole/Boole-SetOdinRndTrigger.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: SIM
      visible: false
    output:
    - type: DIGI
      visible: false
    options_format: l0app
  - name: TCK-0x5138160F (HLT1) Flagged for 2016
    processing_pass: Trig0x5138160F
    visible: false
    application: Moore/v25r4
    options:
    - $APPCONFIGOPTS/Moore/MooreSimProductionForSeparateL0AppStep2015.py
    - $APPCONFIGOPTS/Conditions/TCK-0x5138160F.py
    - $APPCONFIGOPTS/Moore/DataType-2016.py
    - $APPCONFIGOPTS/Moore/MooreSimProductionHlt1.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DIGI
      visible: false
    output:
    - type: DIGI
      visible: false
  - name: TCK-0x6139160F (HLT2) Flagged for 2016
    processing_pass: Trig0x6139160F
    visible: true
    application: Moore/v25r4
    options:
    - $APPCONFIGOPTS/Moore/MooreSimProductionForSeparateL0AppStep2015.py
    - $APPCONFIGOPTS/Conditions/TCK-0x6139160F.py
    - $APPCONFIGOPTS/Moore/DataType-2016.py
    - $APPCONFIGOPTS/Moore/MooreSimProductionHlt1.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DIGI
      visible: false
    output:
    - type: DST
      visible: false
  - name: Turbo lines (MC) for 2016 Turbo03a
    processing_pass: Turbo03a
    visible: true
    application: DaVinci/v41r5
    options:
    - $APPCONFIGOPTS/Turbo/Tesla_2016_LinesFromStreams_MC.py
    - $APPCONFIGOPTS/Turbo/Tesla_PR_Truth_2016.py
    - $APPCONFIGOPTS/Turbo/Tesla_Simulation_2016.py
    - $APPCONFIGOPTS/Turbo/Tesla_FilterMC.py
    data_pkgs:
    - AppConfig.v3r412
    - TurboStreamProd.v4r2p9
    input:
    - type: DST
      visible: false
    output:
    - type: DST
      visible: false
    options_format: Tesla
  - name: Stripping28r2NoPrescalingFlagged for 2016
    processing_pass: Stripping28r2NoPrescalingFlagged
    visible: true
    application: DaVinci/v44r10p5
    options:
    - $APPCONFIGOPTS/DaVinci/DV-Stripping28r2-Stripping-MC-NoPrescaling-DST.py
    - $APPCONFIGOPTS/DaVinci/DV-RedoCaloPID-Stripping_28_24.py
    - $APPCONFIGOPTS/DaVinci/DataType-2016.py
    - $APPCONFIGOPTS/DaVinci/InputType-DST.py
    - $APPCONFIGOPTS/DaVinci/DV-Stripping-MC-muDST.py
    - $APPCONFIGOPTS/DaVinci/DV-RawEventJuggler-4_3-to-4_3.py
    - $APPCONFIGOPTS/Conditions/Sim09-deuteron-nameFix.py
    data_pkgs:
    - AppConfig.v3r412
    - TMVAWeights.v1r16
    input:
    - type: DST
      visible: false
    output:
    - type: ALLSTREAMS.MDST
      visible: false
  - name: Merge for ALLSTREAMS.MDST
    processing_pass: merging
    visible: false
    application: DaVinci/v46r5
    options:
    - $APPCONFIGOPTS/Merging/DVMergeDST.py
    - $APPCONFIGOPTS/DaVinci/DataType-2016.py
    - $APPCONFIGOPTS/Merging/WriteFSR.py
    - $APPCONFIGOPTS/Merging/MergeFSR.py
    - $APPCONFIGOPTS/DaVinci/Simulation.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: ALLSTREAMS.MDST
      visible: true
    output:
    - type: ALLSTREAMS.MDST
      visible: true
    options_format: merge
- type: Simulation
  priority: 2a
  name: Ds2KKpi 2012 pp MagDown
  mc_config_version: '2012'
  sim_condition: Beam4000GeV-2012-MagDown-Nu2.5-Pythia8
  author: null
  inform:
  - auser
  - firstname.surname@cern.ch
  comment: ''
  wg: Charm
  retention_rate: 1.0
  event_types:
  - id: '21103030'
    num_events: 100000
    num_test_events: 10
  - id: '27163001'
    num_events: 100000
    num_test_events: 10
  steps:
  - name: Sim09l - 2012 - MagDown - Pythia8
    processing_pass: Sim09l
    visible: true
    application:
      name: Gauss
      version: v49r24
      binary_tag: x86_64-slc6-gcc48-opt
    options:
    - $APPCONFIGOPTS/Gauss/Sim08-Beam4000GeV-md100-2012-nu2.5.py
    - $APPCONFIGOPTS/Gauss/DataType-2012.py
    - $APPCONFIGOPTS/Gauss/RICHRandomHits.py
    - $APPCONFIGOPTS/Gauss/NoPacking.py
    - $DECFILESROOT/options/@{eventType}.py
    - $LBPYTHIA8ROOT/options/Pythia8.py
    - $APPCONFIGOPTS/Gauss/G4PL_FTFP_BERT_EmNoCuts.py
    data_pkgs:
    - AppConfig.v3r412
    - Gen/DecFiles.v30r87
    input: []
    output:
    - type: SIM
      visible: false
    dbtags:
      DDDB: dddb-20170721-2
      CondDB: sim-20160321-2-vc-md100
  - name: Digi14c for 2012
    processing_pass: Digi14c
    visible: false
    application: Boole/v30r4
    options:
    - $APPCONFIGOPTS/Boole/Default.py
    - $APPCONFIGOPTS/Boole/DataType-2012.py
    - $APPCONFIGOPTS/Boole/NoPacking.py
    - $APPCONFIGOPTS/Boole/Boole-SetOdinRndTrigger.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: SIM
      visible: false
    output:
    - type: DIGI
      visible: false
  - name: L0 emulation for 2012 - TCK 0x0045
    processing_pass: L0Trig0x0045
    visible: false
    application: Moore/v24r4
    options:
    - $APPCONFIGOPTS/L0App/L0AppSimProduction.py
    - $APPCONFIGOPTS/L0App/L0AppTCK-0x0045.py
    - $APPCONFIGOPTS/L0App/DataType-2012.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DIGI
      visible: false
    output:
    - type: DST
      visible: false
    options_format: l0app
  - name: TCK-0x409f0045 Flagged for 2012
    processing_pass: Trig0x409f0045
    visible: true
    application:
      name: Moore
      version: v14r8p1g1
      binary_tag: x86_64-slc5-gcc46-opt
    options:
    - $APPCONFIGOPTS/Moore/MooreSimProductionForSeparateL0AppStep.py
    - $APPCONFIGOPTS/Conditions/TCK-0x409f0045.py
    - $APPCONFIGOPTS/Moore/DataType-2012.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DIGI
      visible: false
    output:
    - type: DIGI
      visible: false
  - name: Reco14c for MC 2012
    processing_pass: Reco14c
    visible: true
    application:
      name: Brunel
      version: v43r2p13
      binary_tag: x86_64-slc5-gcc46-opt
    options:
    - $APPCONFIGOPTS/Brunel/DataType-2012.py
    - $APPCONFIGOPTS/Brunel/MC-WithTruth.py
    - $APPCONFIGOPTS/Brunel/Sim09-Run1.py
    - $APPCONFIGOPTS/Persistency/DST-multipleTCK-2012.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DIGI
      visible: false
    output:
    - type: DST
      visible: false
  - name: Stripping21NoPrescalingFlagged for 2012
    processing_pass: Stripping21NoPrescalingFlagged
    visible: true
    application: DaVinci/v36r1p5
    options:
    - $APPCONFIGOPTS/DaVinci/DV-Stripping21-Stripping-MC-NoPrescaling.py
    - $APPCONFIGOPTS/DaVinci/DV-RedoCaloPID-Stripping21.py
    - $APPCONFIGOPTS/DaVinci/DataType-2012.py
    - $APPCONFIGOPTS/DaVinci/InputType-DST.py
    - $APPCONFIGOPTS/DaVinci/DV-Stripping-MC-muDST.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: DST
      visible: false
    output:
    - type: ALLSTREAMS.MDST
      visible: false
  - name: Merge for ALLSTREAMS.MDST
    processing_pass: merging
    visible: false
    application: DaVinci/v46r5
    options:
    - $APPCONFIGOPTS/Merging/DVMergeDST.py
    - $APPCONFIGOPTS/DaVinci/DataType-2012.py
    - $APPCONFIGOPTS/Merging/WriteFSR.py
    - $APPCONFIGOPTS/Merging/MergeFSR.py
    - $APPCONFIGOPTS/DaVinci/Simulation.py
    data_pkgs:
    - AppConfig.v3r412
    input:
    - type: ALLSTREAMS.MDST
      visible: true
    output:
    - type: ALLSTREAMS.MDST
      visible: true
    options_format: merge
- type: Simulation
  priority: 2a
  name: Ds2KKpi 2016 pp MagDown
  mc_config_version: '2016'
  sim_condition: Beam6500GeV-2016-MagDown-Nu1.6-25ns-Pythia8
  author: null
  inform:
  - auser
  - firstname.surname@cern.ch
  comment: ''
  wg: Charm
  retention_rate: 1.0
```

# Managing the complexity

- That's a lot...
- Undesirable to move abstractions into LHCbDIRAC itself
  - Experts regularly need to manually provide that information for the O(%) special cases
  - LHCbDIRAC developers are disjoint from the Simulation experts
- LHCbDIRAC and Simulation are very small teams (<10 people), enables:
  - Simulation experts to focus on simulation
  - LHCbDIRAC experts to focus on grid/production activities
  - Analysts to focus on analysis
- Instead, start from a minimal YAML file which can be written by anyone

```yaml
sim-version: 09
name: Ds2KKpi
inform:
  - auser
  - firstname.surname@cern.ch
WG: Charm

samples:
  - event-types:
      - 23103005
      - 23103006
    data-types:
      - 2012
      - 2016
    num-events: 100_000
```

# How: Stages

- Previous YAML file is “stage 0”
- Apply a series of transformations to add more information at each stage
- Eventually get “stage 6” which can be directly executed/submitted
- Transformation between stages is performed using a Python package (`LbMCSubmit`)
- Benefits:
  - **Safe:** Non-expert users can only modify things which are deemed suitable
  - **Maintainable:** Python expertise common in LHCb
  - **Flexible:** Experts can **edit any stage** directly to handle special requests

# What are the stages?

- Very LHCb specific, targeted around how LHCb evolved in the last 15 years
- At a high level:
  1. Include defaults for optional parameters
  2. Convert to a list of requests
  3. Flatten each request into the final list of distinct requests
  4. Apply conditions specific information (beam info, geometry, trigger config)
  5. Apply special fixes and sanity checks
  6. Convert the high level configuration into a format LHCbDIRAC can execute

*Actual example: https://gitlab.cern.ch/lhcb-simulation/lbmcsubmit/-/blob/main/docs/transformations.md*

# Submission

- At this point the request is the same as the legacy system
  - Can we do better? - Take inspiration from LHCb's Analysis Productions*
- Use GitLab to submit the request to LHCbDIRAC
  - CI can run test jobs for each production
  - **Catch bugs before submission**
  - Opportunity for **review prior to submission**
  - Expose estimated resource usage to requesters
- Complements the existing testing procedure
  1. CI checks if the configuration works prior to submission
  2. LHCbDIRAC testing phase generates more events to catch rarer issues
  3. Full production launches

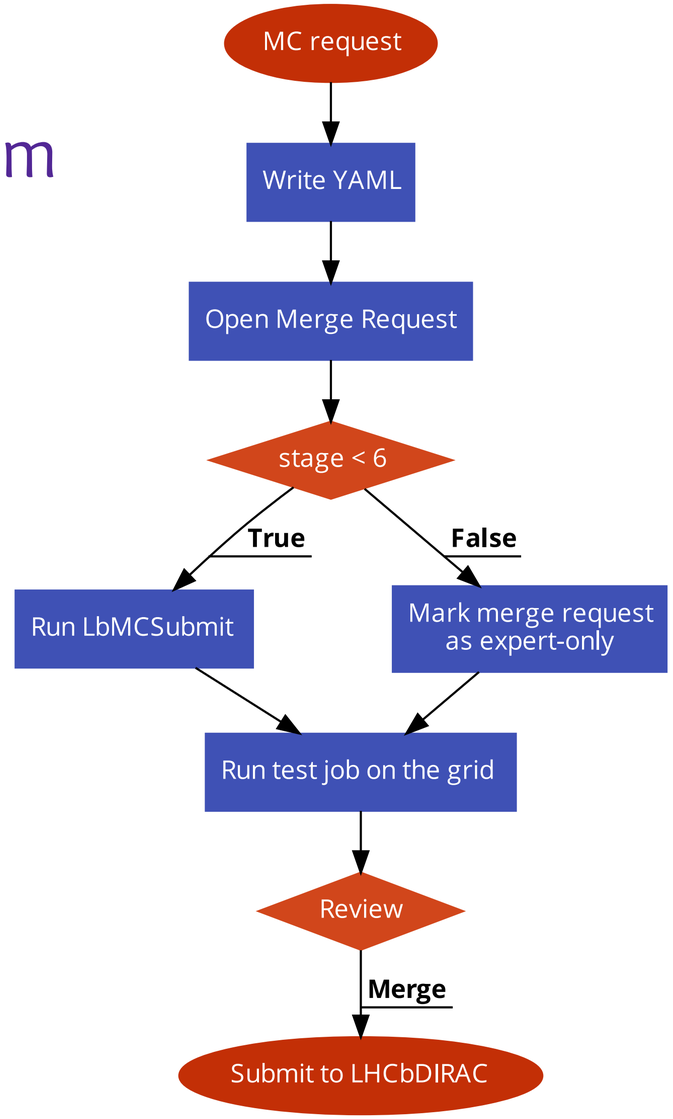


*Analysis Productions: A declarative approach to ntupling (9 May 2023, 17:45)

# CI and testing

*Overall summary*

LHCb MC Requests — Log out Chris Burr (christopher.burr@cern.ch)

Home
Pipelines
Documentation

MC Request Pipelines / 252 / Bc2pimumu

12 jobs triggered by rihender tested at commit bc77783d for lhcb-simulation/mc-requests!50

`Change event type, keeping same conditions as last Bc2pimumu request.`

**Looks good!**
12 jobs completed successfully.

## Jobs (12 total)

Filter jobs... | Retry Failed/Cancelled | Cancel Running

| | | Test job statistics | | | Production resource usage | | | |
|---|---|---|---|---|---|---|---|---|
| Status | Job Name | Generated | Stored | Size | Requested | Generate | Disk | CPU years (Gauss) |
| SUCCESS | request.yaml:Bc2pimumu 2018 pp MagUp:14113400 | 10 | 10 | 2.3 MB | 500,000 | 500,000 | 114.87 GB | 23 |
| SUCCESS | request.yaml:Bc2pimumu 2018 pp MagDown:14113400 | 10 | 10 | 2.66 MB | 500,000 | 500,000 | 133.01 GB | 22 |
| SUCCESS | request.yaml:Bc2pimumu 2017 pp MagUp:14113400 | 10 | 10 | 2.53 MB | 500,000 | 500,000 | 126.6 GB | 22 |
| SUCCESS | request.yaml:Bc2pimumu 2017 pp MagDown:14113400 | 10 | 10 | 2.59 MB | 500,000 | 500,000 | 129.5 GB | 22 |
| SUCCESS | request.yaml:Bc2pimumu 2016 pp MagUp:14113400 | 10 | 10 | 2.38 MB | 500,000 | 500,000 | 119.16 GB | 27 |
| SUCCESS | request.yaml:Bc2pimumu 2016 pp MagDown:14113400 | 10 | 10 | 2.84 MB | 500,000 | 500,000 | 141.94 GB | 20 |
| SUCCESS | request.yaml:Bc2pimumu 2015 pp MagUp:14113400 | 10 | 10 | 2.19 MB | 500,000 | 500,000 | 109.54 GB | 22 |
| SUCCESS | request.yaml:Bc2pimumu 2015 pp MagDown:14113400 | 10 | 10 | 2.45 MB | 500,000 | 500,000 | 122.27 GB | 22 |
| SUCCESS | request.yaml:Bc2pimumu 2012 pp MagUp:14113400 | 10 | 10 | 2.77 MB | 500,000 | 500,000 | 138.47 GB | 19 |
| SUCCESS | request.yaml:Bc2pimumu 2012 pp MagDown:14113400 | 10 | 10 | 3.01 MB | 500,000 | 500,000 | 150.51 GB | 10 |
| SUCCESS | request.yaml:Bc2pimumu 2011 pp MagUp:14113400 | 10 | 10 | 2.07 MB | 500,000 | 500,000 | 103.28 GB | 5 |
| SUCCESS | request.yaml:Bc2pimumu 2011 pp MagDown:14113400 | 10 | 10 | 1.73 MB | 500,000 | 500,000 | 86.52 GB | 17 |

## Configuration

Full Stage 6 | request.yaml

*Can view detailed information about each test*

LHCb MC Requests — Log out Chris Burr (christopher.burr@cern.ch)

Home
Pipelines
Documentation

MC Request Pipelines / 252 / Bc2pimumu / request.yaml:Bc2pimumu 2018 pp MagUp:14113400

500,000 events requested for RD
There have been 1 attempts to run this job.

**Looks good!**
Ran for 10 events in dirac:740737716
Production was submitted as 112557.

## Steps

| No. | Application | Data packages | DB tags | Options | ProcPass | Visible | Test info |
|---|---|---|---|---|---|---|---|
| 1 | Gauss/Gauss (x86_64-slc6-gcc48-opt) | AppConfig.v3r414 Gen/DecFiles.v30r90 | **DDDB:** dddb-20170721-3 **CondDB:** sim-20190430-vc-mu100 | | Sim09l | ✓ | |
| 2 | Boole/v30r4 | AppConfig.v3r414 | — | | Digi14c | ✗ | |
| 3 | Moore/v28r3p1 | AppConfig.v3r414 | — | | LOTrig0x18a4 | ✗ | |
| 4 | Moore/v28r3p1 | AppConfig.v3r414 | — | | Trig0x517a18a4 | ✗ | |
| 5 | Moore/v28r3p1 | AppConfig.v3r414 | — | | Trig0x617d18a4 | ✓ | |
| 6 | Brunel/v54r4 | AppConfig.v3r414 Det/SQLDDDB.v7r10 | — | | Reco18 | ✓ | |
| 7 | DaVinci/v44r7 | AppConfig.v3r414 TurboStreamProd.v4r2p10 | — | | Turbo05-WithTurcal | ✓ | |
| 8 | DaVinci/v44r11p1 | AppConfig.v3r414 TMVAWeights.v1r17 | — | | Stripping34r0p2NoPrescalingFlagged | ✓ | |

## Resource Consumption

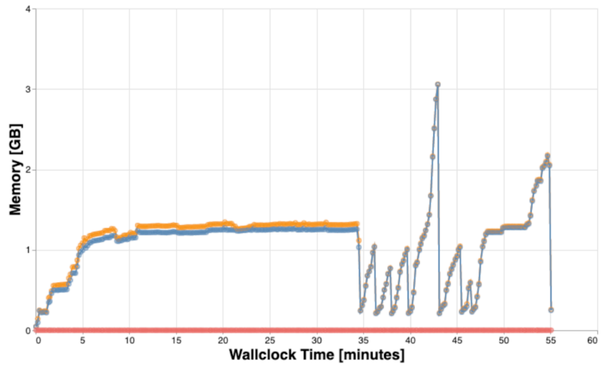


| | Maximum | Average | Largest Increase [per 30s] |
|---|---|---|---|
| Resident Set Size | 3.06 GB | 1.11 GB | 484.06 MB at 42 min |
| Proportional Set Size | 3.05 GB | 1.06 GB | 484.26 MB at 42 min |
| Swap | — | — | — at 0 min |

## Logs

Gauss_1.log | prodConf_Gauss_1.json | Boole_2.log | prodConf_Boole_2.json | Moore_3.log | prodConf_Moore_3.json | Moore_4.log | prodConf_Moore_4.json | Moore_5.log | prodConf_Moore_5.json | Brunel_6.log | prodConf_Brunel_6.json | DaVinci_7.log | prodConf_DaVinci_7.json | DaVinci_8.log | prodConf_DaVinci_8.json | DIRAC.log | prmon.txt

# Summary

- Recently replaced the LHCb Simulation requests system
- Essential to manage the complexity of supporting over a decade of data taking
  - Frees up expert time so they can focus on more impactful work
- Especially valuable non-mainstream areas of the physics programme
  - No need to prioritise which request “models” get updated frequency
- Simpler, safer, more flexible - pick three!

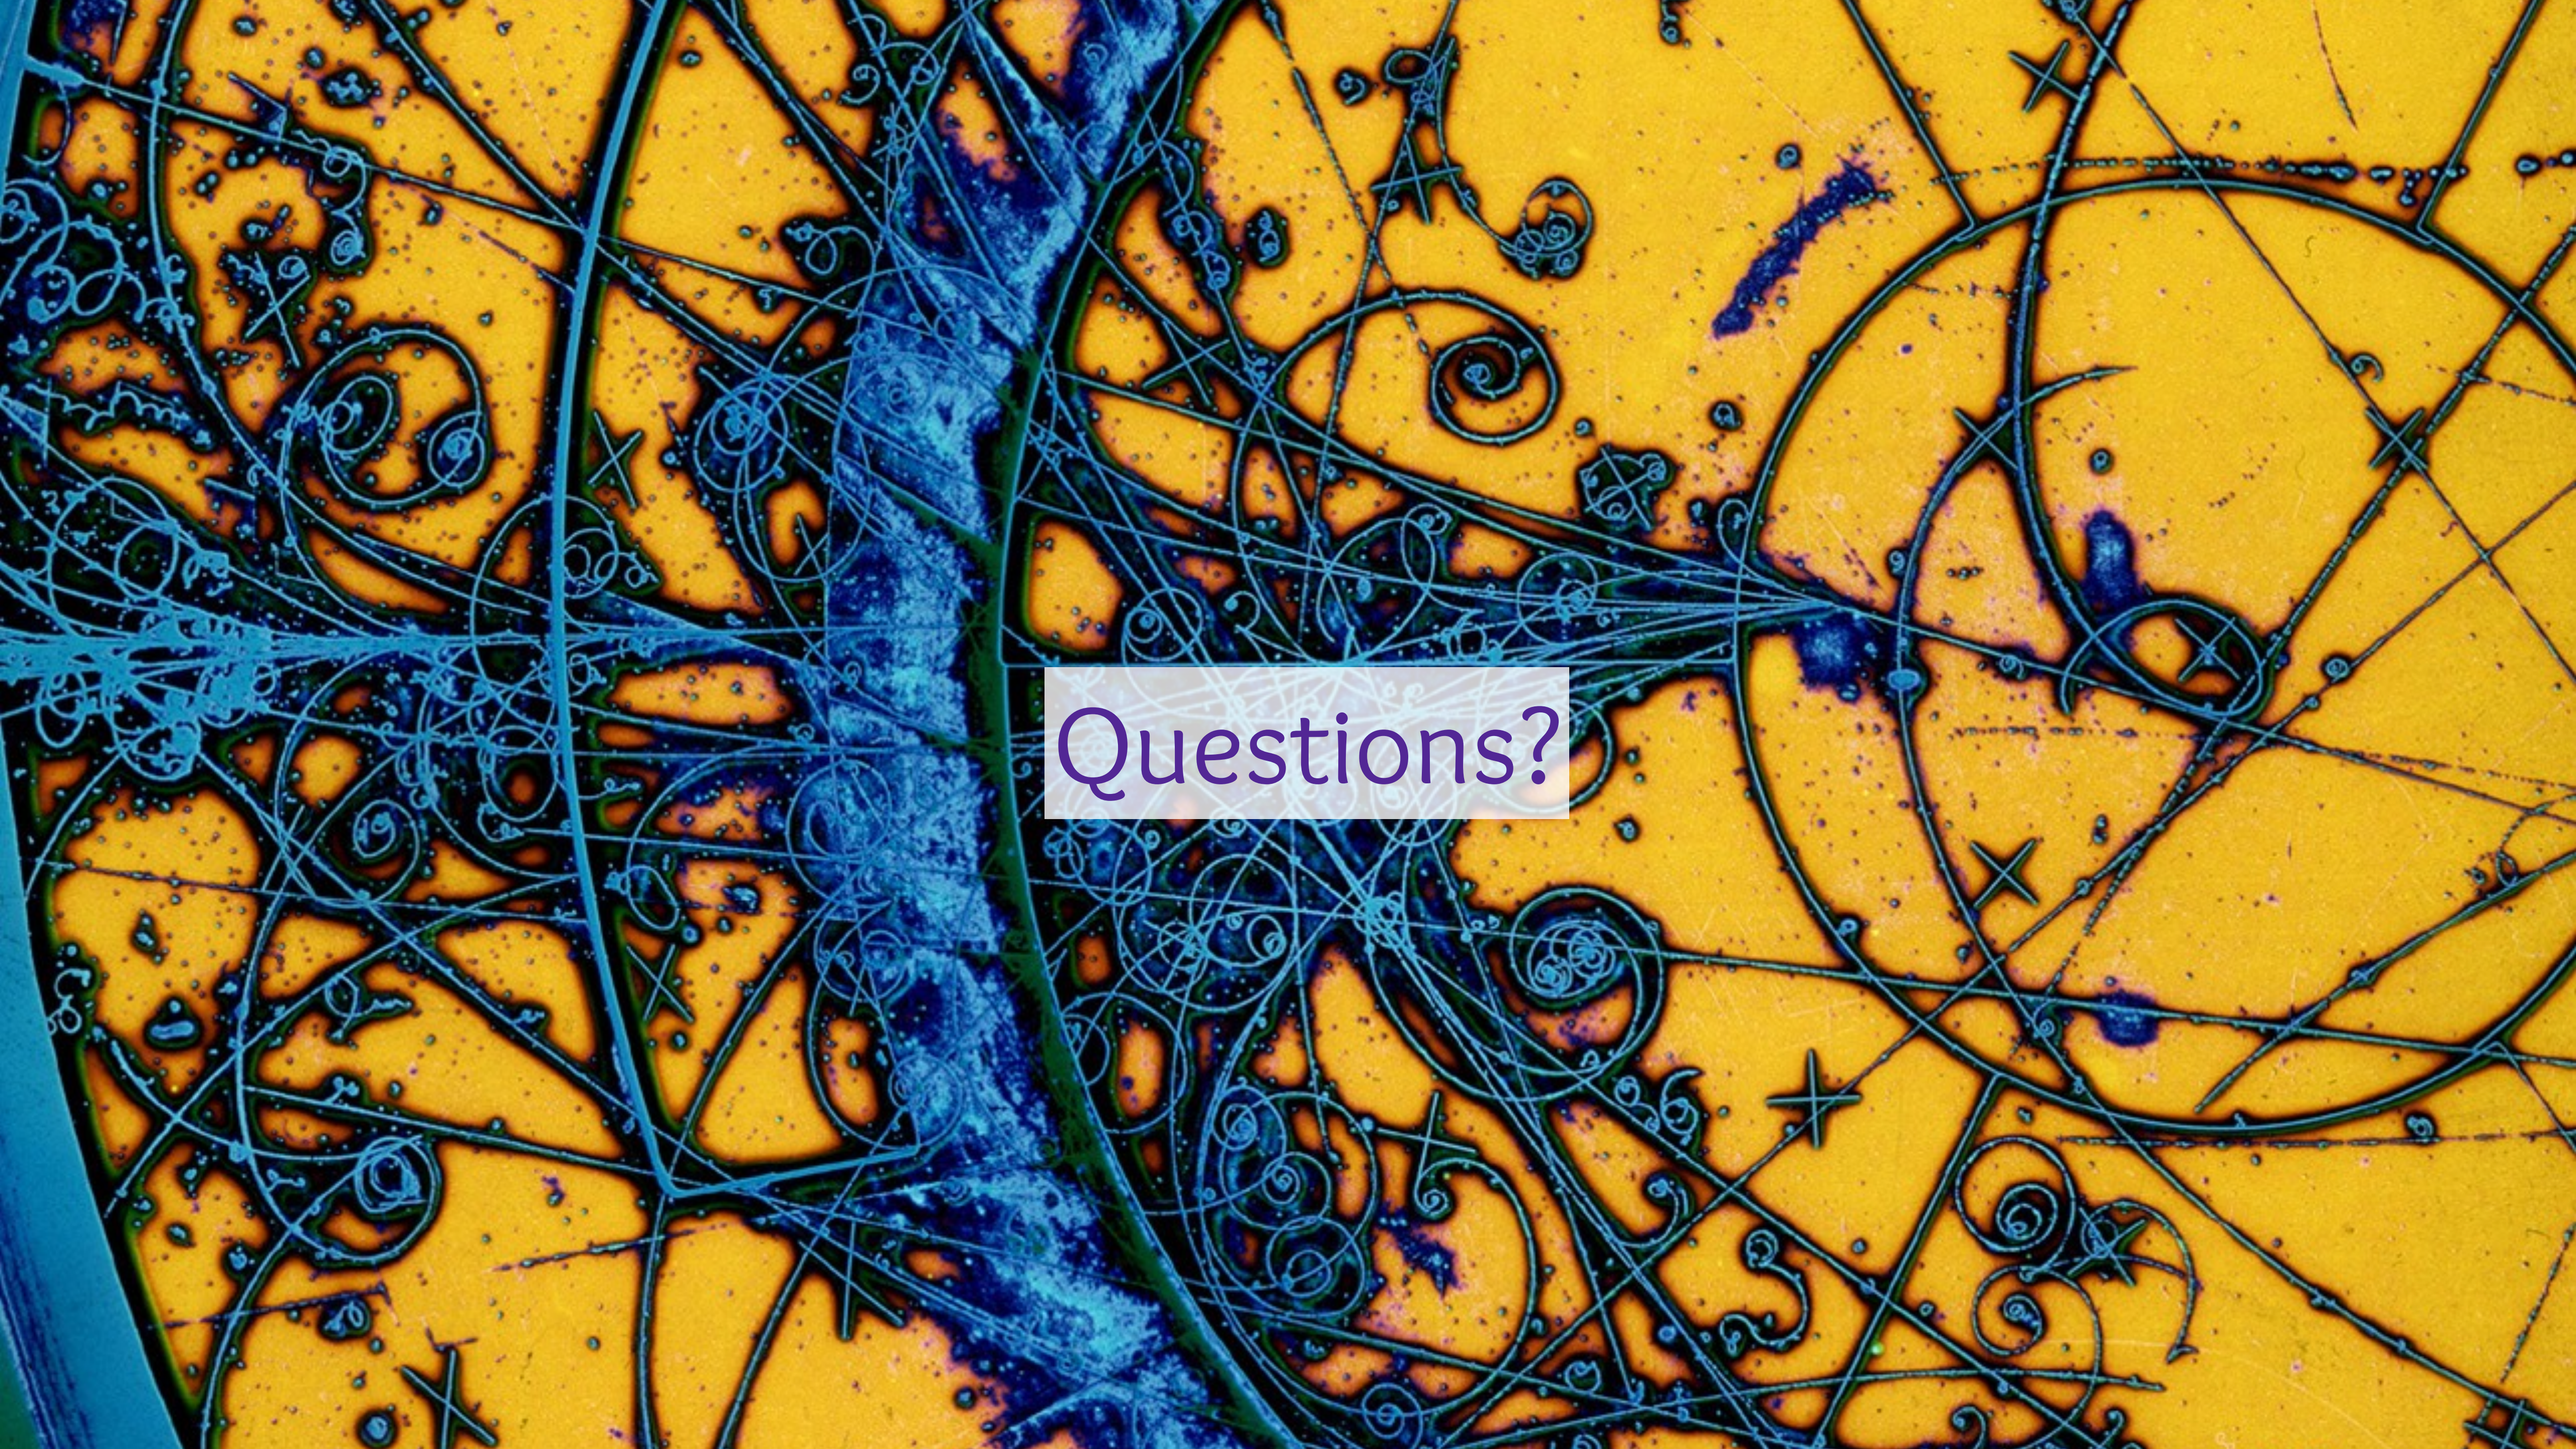
Questions?