IDAP:

IRIS-HEP Data Analysis Pipeline

Alexander Held (University of Wisconsin–Madison)
Oksana Shadura (University Nebraska–Lincoln)

March 15, 2024 https://indico.cern.ch/event/1388620/

This work was supported by the U.S. National Science Foundation (NSF) cooperative agreements OAC-1836650 and PHY-2323298 (IRIS-HEP).

“200 Gbps analysis”

What this means in practice

- DOMA G3.6 : “*Demonstrate analyses running at **200 Gbps** as part of the Analysis Grand Challenge*”

- “HL-LHC scale” (*our definition*): process **10% of 200 TB dataset in 20 min**
  - This is **133 Gbps**
  - Assume 2 kB event size -> **83 M events / second**
  - Assume 25 kHz / core -> **need 3340 cores** (5 MB/s per core)

Where are we at currently?

Plots are from May 2023

See https://indico.cern.ch/e/agc-demonstration

> 200 kHz event rates for simple I/O dominated tasks (simple plots, e.g. for analysis optimization) — XCache crucial!

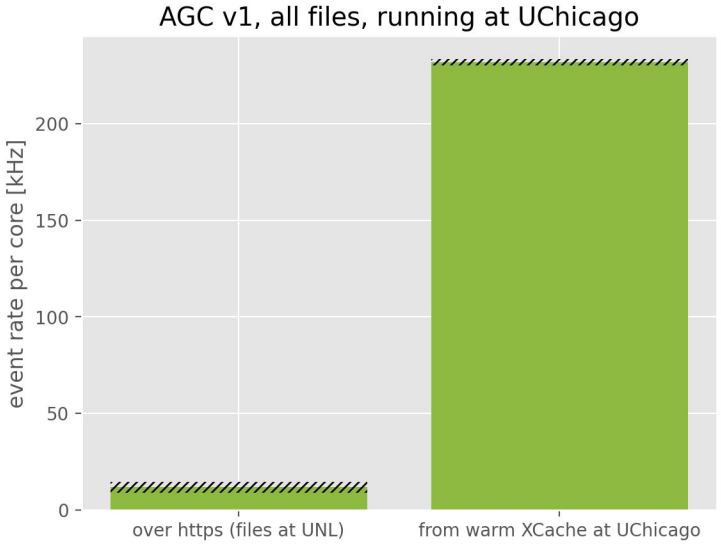

~ 20 kHz event rate for more computation-heavy task (more realistic for full analysis) -> expect slightly less at full scale

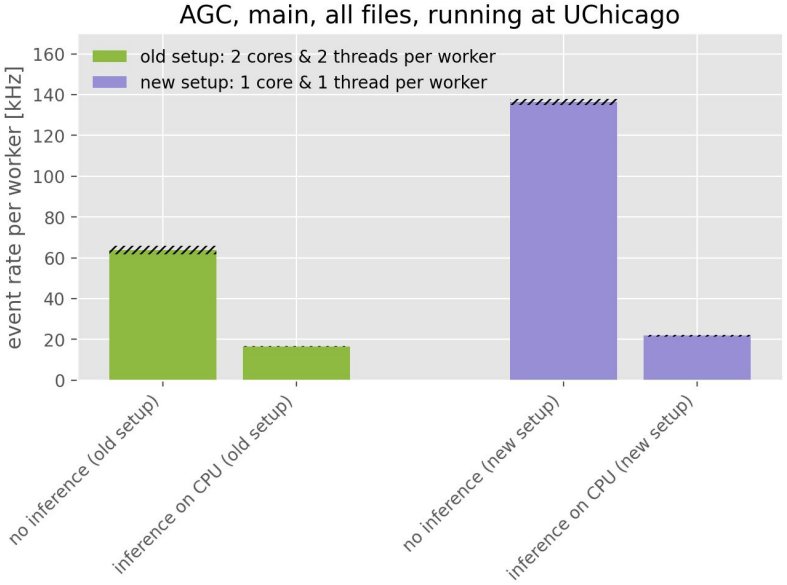

Scaling

Plots are from May 2023

- AGC v1, UNL
  - Identify bottlenecks

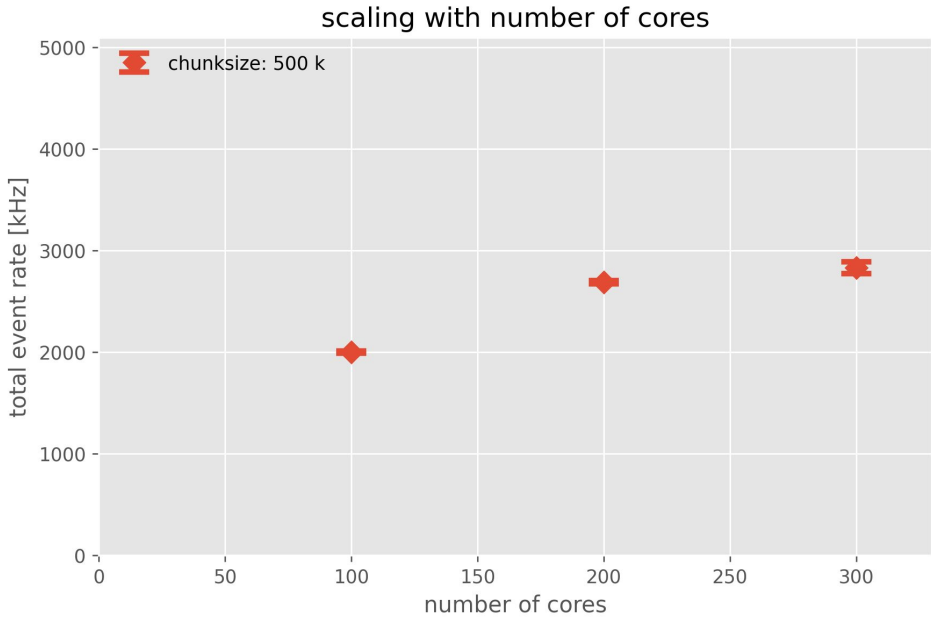

How to achieve our goals?

- **Improve analysis tools ecosystem**
  - **Scalability:** can run tools in distributed way (e.g. tuning Dask at scale)
  - **I/O efficiency:**
    - no I/O bottlenecks: uproot5/issues/1157 (e.g. Nick Smith)
    - Decompression bottleneck (how many cores for 200 Gbps?)
- **Tuning DOMA**
  - See Brian's talk
- **Efficient infrastructure setup**
  - Hardware provisioning:
    - network / computing resources
  - Facility "layer" setup and investigating scaling
    - E.g. tails with individual jobs taking too long / scheduling efficiency

Next steps towards the goal

- Brainstorm list of to-do items for different areas: AS, DOMA, SSL
  - Feel free to keep track in **this google doc**

- See what could be achieved through various fellow projects this year

- Target to show first results at CHEP (?)

IDAP meetings moving forward

- Bi-weekly meetings: **next IDAP on April 5** (https://indico.cern.ch/event/1394124/)
  - **Save the dates! -** April 19, May 3, May 17 (likely to cancel), May 31, June 14, June 28
- Envision **round-table updates** from focus areas / projects on IDAP-related progress
- **Format of a topical meeting** on subject spanning multiple focus areas
  - **Next topic:** G6.1 and G6.4 “ServiceX in production” / new frontend
  - Other ideas about future topics:
    - Workflow management (idea for April 19)
    - User environment handling
    - Coffea 2024 + Dask
    - …

We have scheduled next **IRIS-HEP Demo day** on 24 March 2024:
https://indico.cern.ch/event/1394151/

Backup

For IRIS-HEP v2: split AGC into 2 components

AGC (as it was in IRIS-HEP v1)

IRIS-HEP Data Analysis Pipeline (IDAP) **project**

with focus on delivering
integration test for IRIS-HEP
(SSL, DOMA, AS Areas)

Analysis Grand Challenge (AGC)
project

which <u>will continue</u> building
community

Yearly benchmarking exercises

- Year 1 goal: **stable analysis pipeline** at scale with 30 simultaneous users
- Subsequent years: iterative scaling to HL-LHC needs

getting ready for HL-LHC

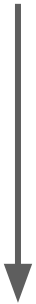

Timeline	Fraction of HL-LHC dataset processed in 1h
Year 2	20% (40 TB)
Year 3	50% (100 TB)
Year 4	75 % (150 TB)
Year 5	100% (200 TB)

IDAP/AGC Milestones and Deliverables - Year 1

ID	Description	Date	WBS
G8.1	Execute technical (internal) demo event	Y6Q2	WBS 8.1
G8.2	Release AGC v2 (more CPU-intense ML & systematics)	Y6Q2	WBS 8.1
G8.3	Demonstrate stable, multi-user, distributed AGC execution (B1)	Y6Q3	WBS 8.1
G8.4	Blueprint defining future “reference analyses”	Y6Q4	WBS 8.2
G8.5	Yearly performance investigation of all AGC tools / services (resulting in report)	Y7Q4	WBS 8.1
G8.6	Yearly AGC (IDAP) training event at community event (joint with SSC) on coffea-casa	Y6Q4	WBS 8.2

Demo day on March 1

AGC (not IDAP)

Should start already this year

Yearly benchmarking exercise

Timeline	Benchmark
Year 1	IRIS-HEP Data Analysis pipeline at scale (30 simultaneous user running full AGC example with at least 30 workers +100% service reliability at SSL facilities)
Year 2	IRIS-HEP Data Analysis pipeline with 20% of HL-LHC dataset (40 TB), completed within 1 h
Year 3	IRIS-HEP Data Analysis pipeline with 50% of HL-LHC dataset (100 TB), completed within 1 h
Year 4	IRIS-HEP Data Analysis pipeline with 75 % of HL-LHC dataset (150 TB), completed within 1 h
Year 5	IRIS-HEP Data Analysis pipeline with 100% of HL-LHC dataset (200 TB), completed within 1 h