

# Swirles: Intel GPU Testbed at Cambridge

Dr Marek Szuba  
Faculty HPC Manager  
Faculty of Mathematics, University of Cambridge

Computing Insight UK 2024

# Overview

- 1 HPC at Cambridge Mathematics
- 2 The Swirles cluster
- 3 PVC Testbed
- 4 Summary

# Cambridge Faculty of Mathematics



- Departments:
  - Applied Mathematics and Theoretical Physics
    - home of Stephen Hawking Centre for Theoretical Cosmology
    - Intel oneAPI Centre of Excellence
  - Pure Mathematics and Mathematical Statistics
    - Statistical Laboratory
- Wide variety of research topics
  - computationally intensive include HEP, GR, fluid dynamics, ...

# HPC at Maths

- Goals:
  - *Complement* large systems such as central-IT clusters
  - Small (but powerful) development environment
  - No long-term storage
  - Environment familiar to Maths researchers
- Resources:
  - Slurm on Desktops
  - Co-located, managed compute systems
  - *Dedicated HPC clusters*

# HPC at Maths

- Goals:
  - *Complement* large systems such as central-IT clusters
  - Small (but powerful) development environment
  - No long-term storage
  - Environment familiar to Maths researchers
- Resources:
  - Slurm on Desktops
  - Co-located, managed compute systems
  - *Dedicated HPC clusters*


# The Swirles Cluster

- Planned 4Q 2022, delivered 2Q–3Q 2024
- Namesake: Bertha Swirles, Lady Jeffreys (1903–1999)
  - alumna and fellow of Girton College
  - accomplished quantum physicist
  - co-author of *Methods of Mathematical Physics* (“Jeffreys & Jeffreys”)
- High-density compute nodes with direct water cooling
- AMD Genoa and Intel Sapphire Rapids CPUs
- Nvidia Ampere, Nvidia Lovelace and Intel Ponte Vecchio GPUs
- 400 Gb/s interconnect (NDR InfiniBand)
- Three-node storage cluster with 360 TB of raw HDD storage
- OS: Ubuntu 22.04 LTS “Jammy Jellyfish”
- Funded in (large) part by ExCALIBUR

# Swirles Compute Nodes

- “Genoa” — 6 nodes, 192 Genoa cores / 1.5 TB DDR5 RAM per node
- “SPR” — 4 nodes, 112 SPR cores / 128 GB HBM + 1 TB DDR5 RAM per node
- “PVC” — 3 nodes, 112 SPR cores / 128 GB HBM + 1 TB DDR5 RAM / 4 fully interconnected Intel GPU Max 1550 (128 GB) GPUs per node
- “Lovelace” — 1 node, 48 Genoa cores / 1.5 TB DDR5 RAM / 4 Nvidia L40S GPUs
- “Ampere” — 3 nodes, 48 Ice Lake cores / 2 TB DDR4 RAM / 4 fully interconnected Nvidia A100 (40 GB SXM) GPUs

# Swirles Hardware







# Swirles Hardware



# Swirles Hardware



# Swirles PVC Testbed

- Hardware:
  - 3x 4x Intel GPU Max 1550
  - 2 tiles per GPU
  - 128 GB of HBM2e memory per GPU
  - SPR-HBM CPUs, DDR5 RAM, NDR IB, NVMe SSD scratch
- Software:
  - GNU/Linux
  - Intel runtime + *fork of* i915 driver
  - PSM3 OpenFabric provider
  - gcc, OpenBLAS, OpenMPI, Intel oneAPI Base+HPC Toolkits
- Access:
  - Faculty affiliates
  - Slurm with GPUs as GRES
  - Limited direct SSH access

# Ready, Set, Go!

## First impressions of operation

- Avoided the early-adopter penalty
- It Just Works! (TM)
  - Standard Linux DRI interfaces
  - Intel packages cover both OpenCL and Level Zero
  - Few non-FOSS components
  - Supported by recent versions of Slurm
  - SYCL work loads developed well in advance
- Successfully ran a training workshop in October
- No tile-level granularity in Slurm
  - ...ever?

# Ready, Set, Go!

## First impressions of operation

- Avoided the early-adopter penalty
- It Just Works! (TM)
  - Standard Linux DRI interfaces
  - Intel packages cover both OpenCL and Level Zero
  - Few non-FOSS components
  - Supported by recent versions of Slurm
  - SYCL work loads developed well in advance
- Successfully ran a training workshop in October
- No tile-level granularity in Slurm
  - ...ever?

# The Ugly Details

## Deployment

- *Discontinued by Intel*
- Thank Bob for Dawn!
- Highly limited documentation
  - driver + runtime: OK
  - fabric: fragmented + partly wrong
  - link calibration, memory issues *etc.* — all word of mouth
  - source-code dives not out of the question
- Scattered software:
  - driver + runtime — Intel Apt repo
  - PSM3 — manual download from Intel
  - xpumanager — manual download from Github
  - oneAPI tools — different Intel Apt repo *and* Conda repo

# The Ugly Details

## Deployment

- *Discontinued by Intel*
- Thank Bob for Dawn!
- Highly limited documentation
  - driver + runtime: OK
  - fabric: fragmented + partly wrong
  - link calibration, memory issues *etc.* — all word of mouth
  - source-code dives not out of the question
- Scattered software:
  - driver + runtime — Intel Apt repo
  - PSM3 — manual download from Intel
  - xpumanager — manual download from Github
  - oneAPI tools — different Intel Apt repo *and* Conda repo

# The Ugly Details

## Deployment

- *Discontinued by Intel*
- Thank Bob for Dawn!
- Highly limited documentation
  - driver + runtime: OK
  - fabric: fragmented + partly wrong
  - link calibration, memory issues *etc.* — all word of mouth
  - source-code dives not out of the question
- Scattered software:
  - driver + runtime — Intel Apt repo
  - PSM3 — manual download from Intel
  - xpumanager — manual download from Github
  - oneAPI tools — different Intel Apt repo *and* Conda repo



# Summary

- Working Intel PVC testbed at Cambridge Mathematics
- Deployed with no major issues
  - ...but only thanks to prior experience
- Already being put to good use
- No long-term future...
- Let's hope next-gen hardware retains current APIs!

**THANK YOU!**