

# Accelerator Service

Matti, Salvatore

# Previous problem(s)

- Continued the Service prototyping
  - Scheduling part almost ready
    - With a mock simulating the GPU
    - Everything ran on CPU, “GPU” work launched with `std::async`
- Finished testing unified memory with Nvidia profiler
  - Example works, can observe page transfers with the profiler
- Started looking into CUDA stream callbacks

# Next problem(s)

- Continue prototyping of Service
  - Should be finished on the afternoon
  - Next step, either
    - Use real GPU workload
    - Test to not using a Service for scheduling
      - Should be doable and avoids any global state
- Continue testing CUDA stream callbacks