MAELSTROM: FIRST BENCHMARK RESULTS

ISC 2022

01.06.2022 | Stepan Nassyr

- 6 Weather&Climate Applications using ML
- Dedicated versions of MAELSTROM applications benchmarked on
  - Jülich system: JUWELS (mostly Booster, but also Cluster)
  - E4 system: Intel+NVIDIA system
- Benchmarks run by application owners, guided by WP3 systems staff
- Metrics selected in cooperation
- Objective: Assess status, identify points of improvement, study hardware
- Final goal: Provide bespoke W&C ML system design; fitting W&C ML applications

Report in deliverable 3.4 (Released 12.04.2022)

METRICS

- Time-related
  - Total runtime
  - Total training time
  - Loading data time
  - Training time per epoch (avg, min, max)
  - Training time per iteration (avg, min, max)
  - Training time of first epoch
  - Model saving time
- Learning-related: Final loss (training, validation)
- Energy-related:
  - GPU power draw (max)
  - Energy consumption (GPU, node)
- (Additional metrics provided by tools influence analysis)

APP 1: BLEND CITIZEN OBSERVATIONS AND NUMERICAL WEATHER FORECASTS

- Developer: Norwegian Meteorological Institute (MetNor)
- Libraries used: TensorFlow 2.x
- Performance Insights:
  - Loading data dominates runtime (~⅔ of runtime)
  - JWC to JWB 2x performance uplift, E4 slower due to I/O performace
  - Mostly stable results over various experiments; first epoch always slower (JUWLS: 1.6 ×; E4: 2 ×)
- To be investigated:
  - Energy measurements unreliable (need to scale benchmark)
  - E4 training 10% faster 3 out of 5 times (Power/thermal limits?)

JUWELS Booster: Total time spent

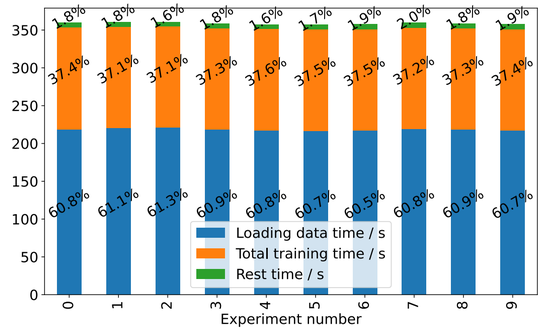

E4: Epoch comparison

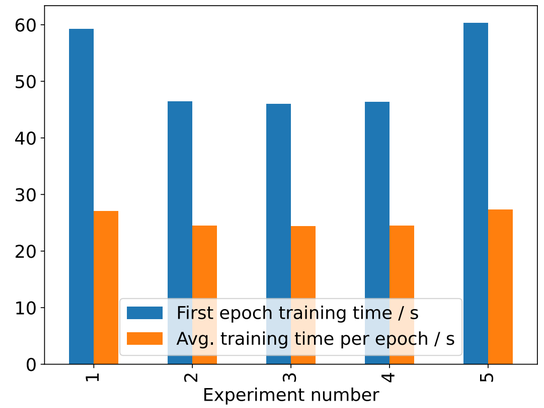

APP 3: NEURAL NETWORK EMULATORS FOR FASTER WEATHER FORECAST MODELS & DATA ASSIMILATION

- Developer: ECMWF
- Libraries used: TensorFlow 2.x
- Multiple configurations explored
- Performance Insights:
  - Training dominates runtime (> 97% of runtime)
  - Training coupled to I/O (Streaming Data)
  - JWC to JWB 2x performance uplift, 4x less energy
  - First multi-GPU (2) experiments. Impr. only when also increasing batch size from 512 to 1024
- To be investigated:
  - nocache/reboot evidence for I/O dependence
  - Scaling to multi-GPU requires tuning

JUWELS Booster: Total time spent

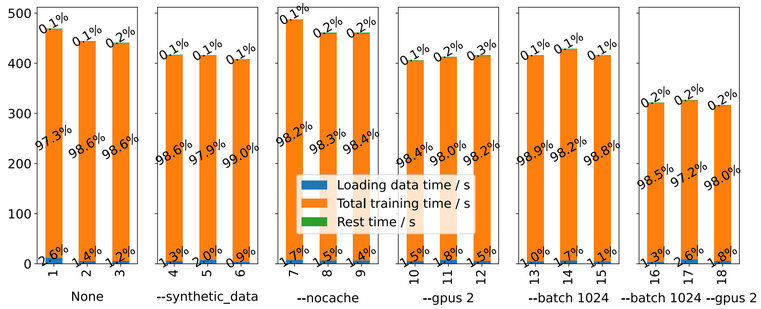

E4: Epoch comparison

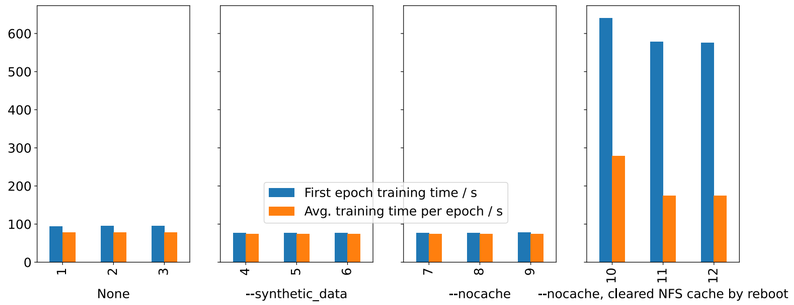

APP 4: IMPROVED ENSEMBLE PREDICTIONS IN FORECAST POST-PROCESSING

- Developer: ETH Zürich
- Libraries used: TF 1.14 transitioning to PyTorch 1.8
- Performance Insights:
  - Large runtime
  - Training dominates runtime (> 87% of runtime)
  - Loading data time and unaccounted time small but significant
  - Good statistics, good GPU usage (single)
- To be investigated:
  - 1 outlier had significantly higher runtime (all 3: I/O, training and unaccounted)

JUWELS Booster: Max GPU power draw

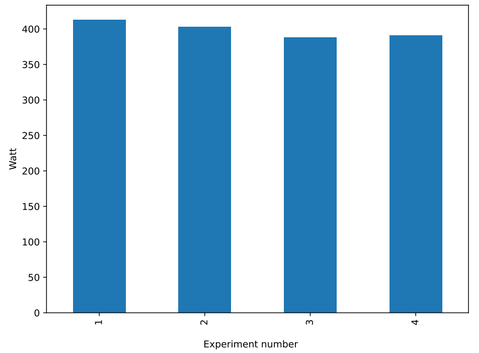

E4: Total time spent

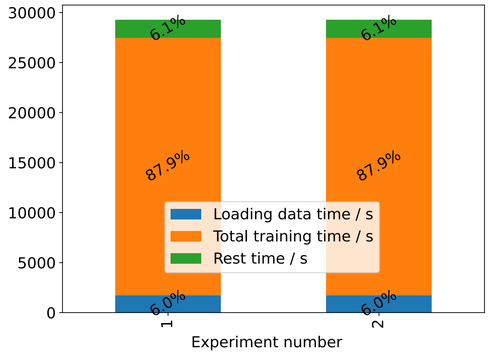

APP 5: IMPROVED LOCAL WEATHER PREDICTIONS IN FORECAST POST-PROCESSING

- Developer: Forschungszentrum Jülich
- Libraries used: TensorFlow 2.3.1/2.5 + Keras
  - Benchmark dataset scaled up from 75s to 1500s runtime, improving benchmark quality
    - Small benchmark shows curious behaviours
    - Evident i.e. in the Epoch ratio
  - JWC to JWB again 2x performance, 4x less power
  - Training dominates runtime (> 92%/98% of runtime for small/large DS)
  - Additional inference benchmarks performed

JUWELS Booster: Epoch ratio (Small)

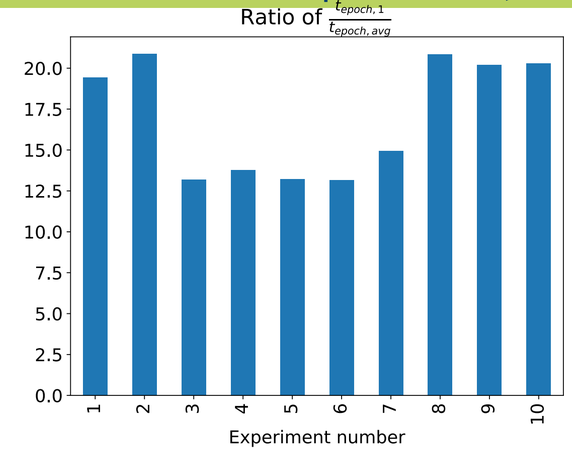

JUWELS Booster: Epoch ratio (Large)

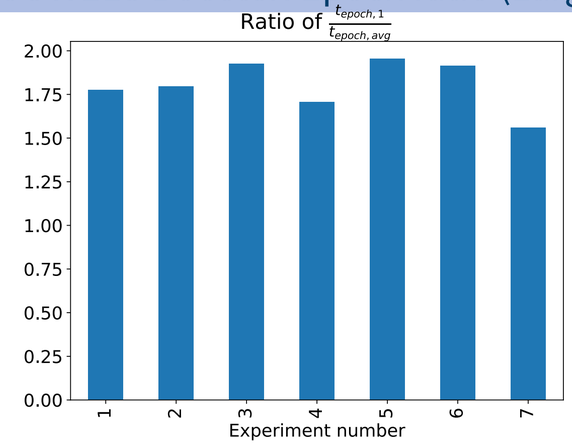

MAELSTROM

JÜLICH Forschungszentrum

APP 2: INCORPORATE SOCIAL MEDIA DATA INTO PREDICTION FRAMEWORK

- Developer: 4cast

Work on AP2 was faced with challenges regarding data acquisition and a staff bottleneck. The dedicated personnel only started working on it in April 2022. Therefore, benchmarking data for AP2 could not be provided until the deadline of the deliverable.

APP 6: BESPOKE WEATHER FORECASTS TO SUPPORT ENERGY PRODUCTION IN EUROPE

- Developer: 4cast
- No NN/GPU usage initially, but TF2 or PyTorch based approach in development
- Benchmarks only for CPU performance

JUWELS Booster: Total time spent

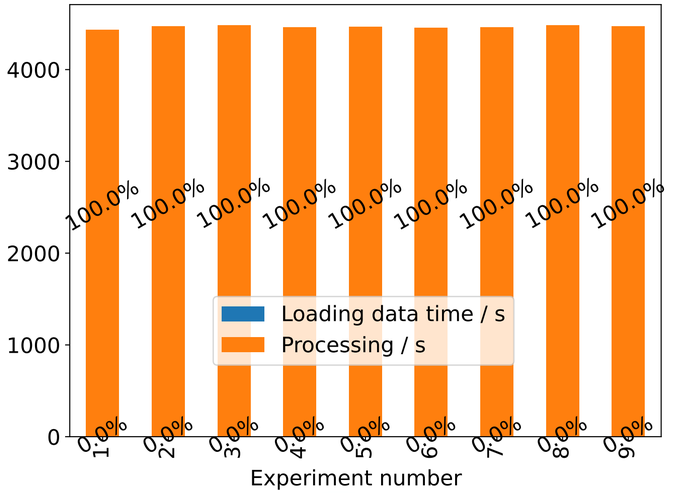

MAELSTROM

JÜLICH Forschungszentrum

CONCLUSION

- Examples shown of selected MAELSTROM application benchmarks
- Applications × Configurations × Hardware = Many data points
- Investigation ongoing, already many specific (and interesting!) features identified
- Also spotted curiosities for further investigation
- Much more data and results then presented here!
- See Deliverable D3.4 on maelstrom-eurohpc.eu

NEXT STEPS

- Feedback sessions with application developers
- Profiling
  - Tool choice
  - Detailed runtime breakdown
  - Communication
  - Memory access
- Developing profile-guided improvements

THANK YOU FOR YOUR ATTENTION!

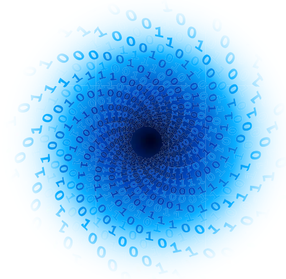

MAELSTROM

E4 COMPUTER ENGINEERING

ECMWF

4cast

ETH zürich

Meteorologisk institutt

JÜLICH Forschungszentrum

EuroHPC Joint Undertaking

Federal Ministry of Education and Research

"The MAELSTROM project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955513. The JU receives support from the European Union's Horizon 2020 research and innovation programme and United Kingdom, Germany, Italy, Luxembourg, Switzerland, Norway".

MAELSTROM

JÜLICH Forschungszentrum