

Analyzing I/O in Deep Learning Workloads

Hariharan Devarajan (devarajan1@llnl.gov)

Analyzing Parallel I/O
Birds of Feather

DfTracer

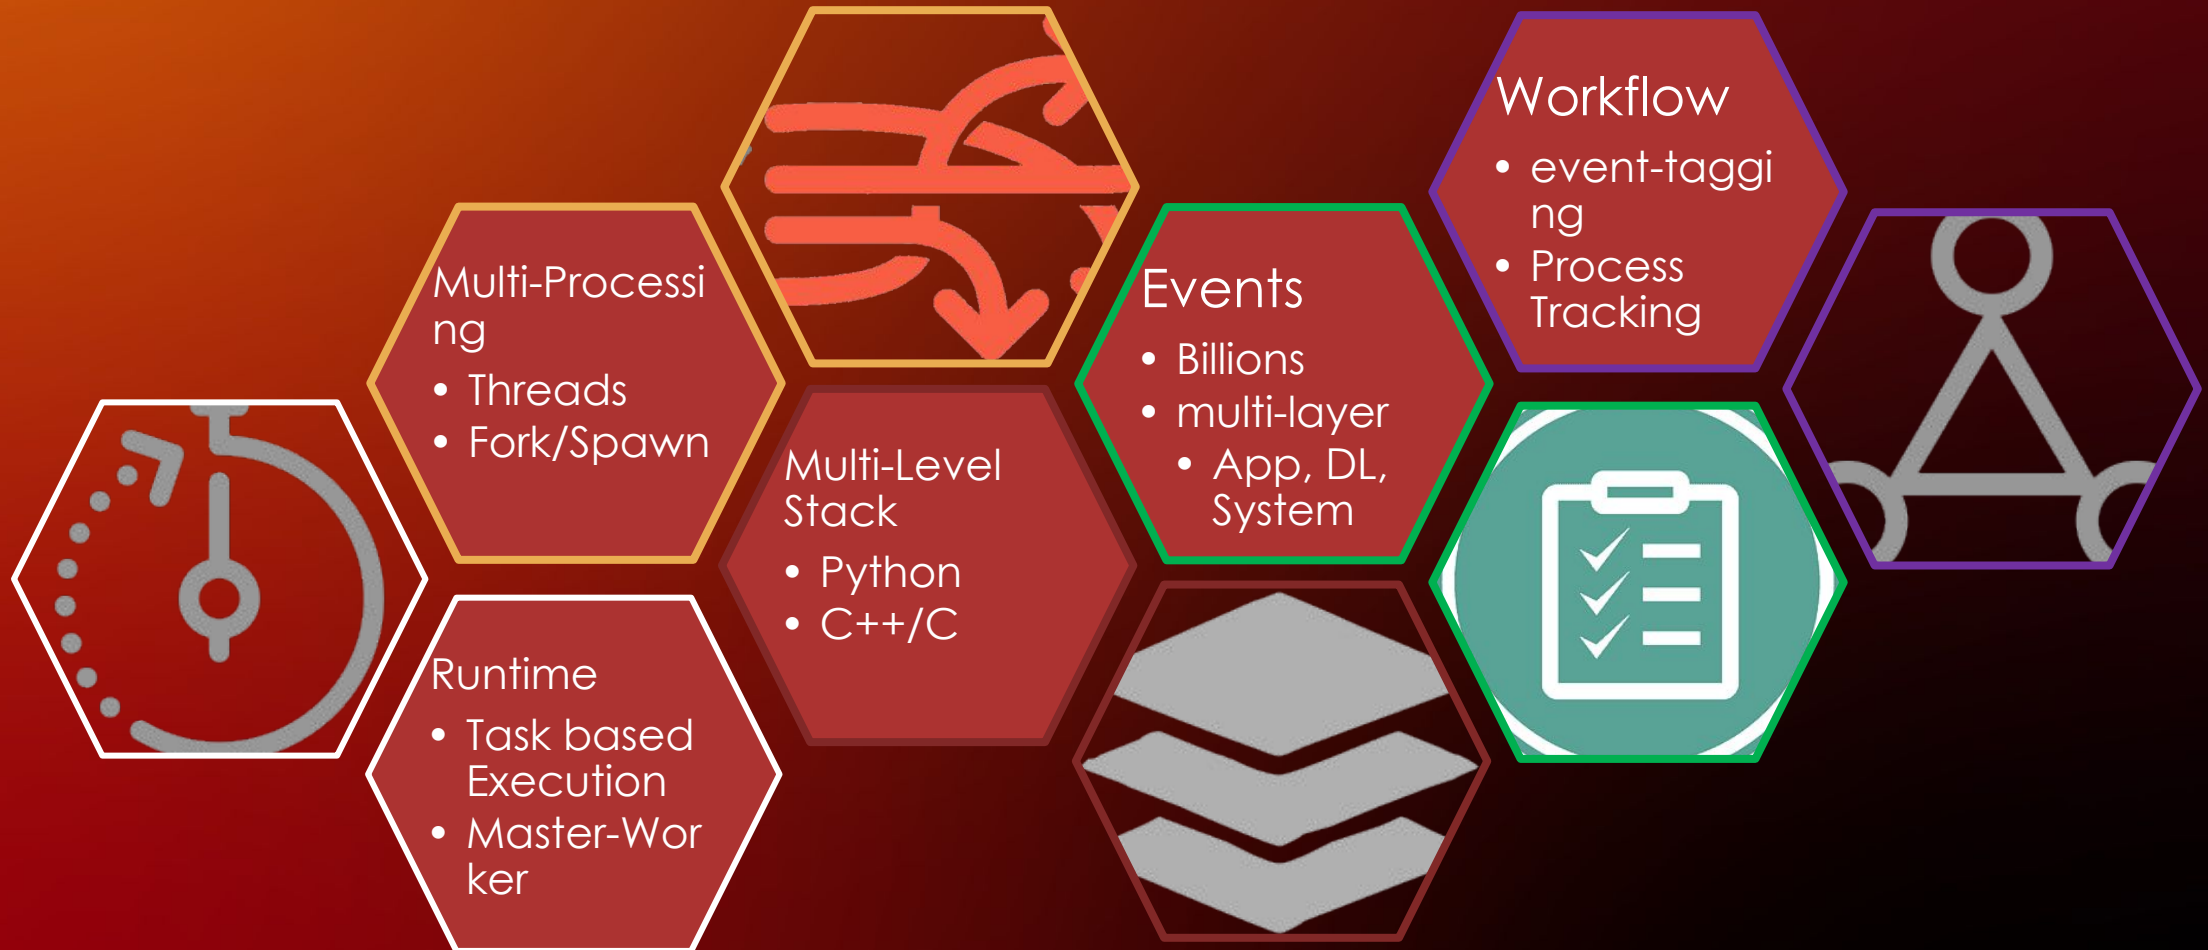


SCAN ME



SC24
Atlanta, GA | hpc creates.

Unique Features of DL Workloads



Recent Tools developed for I/O in AI



DFTracer: A multi-level dataflow tracer for capturing I/O calls from AI-driven workflows.




WisIO: An automatic bottleneck detector for large-scale workloads



DFAnalyzer: A multi-level analyzer for understanding I/O from multiple levels of the hierarchy



Visualization of Megatron Deepspeed with SCR (DFTracer + Perfetto)



Capture multiple data-centric layers of the workload

- Application layer (python code)
- System layer (POSIX and STDIO interface)



Optimizing Megatron Deepspeed with SCR (DFTracer + WisIO)

Difference between overall I/O and Unoverlapped or perceived I/O

```
Runtime: 9773.62 seconds (100%)
├── Application I/O Time: 4424.42 seconds (45.27%)
├── Checkpoint I/O Time: 4405.64 seconds (45.08%)
├── Compute Time: 7345.04 seconds (75.15%)
├── I/O Time: 2218.73 seconds (22.70%)
├── Read I/O Time: 29.65 seconds (0.30%)
├── Unoverlapped Application I/O Time: 3162.50 seconds (32.36%)
├── Unoverlapped Checkpoint I/O Time: 3162.50 seconds (32.36%)
├── Unoverlapped Compute Time: 5934.06 seconds (60.72%)
├── Unoverlapped I/O Time: 807.75 seconds (8.26%)
├── Unoverlapped Read I/O Time: 0.00 seconds (0.00%)
└── Unoverlapped Application Compute Time: 6083.12 seconds (62.24%)
```


```
Runtime: 8428.34 seconds (100%)
├── Application I/O Time: 1487.00 seconds (17.64%)
├── Checkpoint I/O Time: 1485.76 seconds (17.63%)
├── Compute Time: 7327.39 seconds (86.94%)
├── I/O Time: 938.97 seconds (11.14%)
├── Read I/O Time: 1.69 seconds (0.02%)
├── Unoverlapped Application I/O Time: 297.31 seconds (3.53%)
├── Unoverlapped Checkpoint I/O Time: 297.31 seconds (3.53%)
├── Unoverlapped Compute Time: 6469.93 seconds (76.76%)
├── Unoverlapped I/O Time: 81.50 seconds (0.97%)
├── Unoverlapped Read I/O Time: 0.00 seconds (0.00%)
└── Unoverlapped Application Compute Time: 6137.71 seconds (72.82%)
```

Reduced overall runtime is a result of increasing overlap between checkpoint I/O and Computation



DFAnalyzer: Multi-Level analysis of AI workloads

Stromer application from ANL



HDF5 Dataset with PyTorch Data Loader

We need to look at multiple layers and different metrics to understand performance.

Level	Item	Metric	Stromer
App	Compute	Num Instances Agg. Time	2920 15.04 min
	Data Loading	Num Instances Agg. Time	2920 4.32 min
Data Loader	Num Workers	Num Instances	32
	Preprocess	Num Instances Agg. Time	2920 2.66 min
	Get Item	Num Instances Agg. Time	2920 17.93 min
System Call	IO Size	Total	2.93 TB
	All POSIX	Num Instances Agg. Time Agg. Size	82756 16.19 min 2.93 TB
	Data > 4KB	Dur. wrt POSIX Num Instances Agg. Time Agg. Size Transfer Size Bandwidth	98.87 % 5904 16.05 min 2.93 TB 521.05 MB 0.88 GB/s
	Data <= 4KB	Dur. wrt POSIX Num Instances Agg. Time Agg. Size	0.60 % 29244 6.85 secs 6.03 MB
	Others	Dur. wrt POSIX Num Instances Agg. Time	0.53 % 47588 7.76 secs




Questions or Comments

DFTracer



Analyzer



Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC



Detailed Characterization and Analysis (DFTracer + WisIO)

```

POSIX I/O Characteristics
Runtime
I/O Time 9773.62 seconds
2218.73 seconds (100%)
├── Read: 612.49 seconds (27.61%)
├── Write: 1409.17 seconds (63.51%)
└── Metadata: 773.33 seconds (34.85%)
I/O Ops. 114,626,459 ops
├── Read - 24,729,700 ops (21.57%)
├── Write - 4,478,320 ops (3.91%)
└── Metadata - 65,177,974 ops (56.86%)
I/O Size 86.64 TiB
├── Read - 41.83 TiB (48.27%)
└── Write - 44.82 TiB (51.73%)
Read Requests 4 kiB→64 MiB - 24,729,700 ops
├── >64 MiB: 24,696,810 ops (99.87%)
└── [Others: 32,890 ops (0.13%)]
Write Requests 4 kiB→64 MiB - 4,478,320 ops
├── <4 kiB: 233,346 ops (5.21%)
├── >64 MiB: 4,227,808 ops (94.41%)
└── [Others: 17,166 ops (0.38%)]
Time Periods 151 time periods (Time Granularity: 65.00 seconds)
Access Pattern Sequential: 29,208,020 ops (100.00%) - Random: 0 ops (0.00%)
R: Read - W: Write - M: Metadata

```

```

POSIX I/O Characteristics
Runtime 9773.62 seconds
Unoverlapped I/O Time 807.75 seconds (100%)
├── Read: 148.93 seconds (18.44%)
├── Write: 229.07 seconds (28.36%)
└── Metadata: 294.69 seconds (36.48%)
Unoverlapped I/O Ops. 84,805,998 ops
├── Read - 21,718,805 ops (25.61%)
├── Write - 3,233,634 ops (3.81%)
└── Metadata - 57,117,804 ops (67.35%)
Unoverlapped I/O Size 47.29 TiB
├── Read - 26.41 TiB (55.83%)
└── Write - 20.89 TiB (44.17%)
Read Requests 4 kiB→64 MiB - 21,718,805 ops
├── >64 MiB: 21,717,137 ops (99.99%)
└── [Others: 1,668 ops (0.01%)]
Write Requests 4 kiB→64 MiB - 3,233,634 ops
├── >64 MiB: 3,218,930 ops (99.55%)
└── [Others: 14,704 ops (0.45%)]
Time Periods 36 time periods (Time Granularity: 65.00 seconds)
Access Pattern Sequential: 24,952,439 ops (100.00%) - Random: 0 ops (0.00%)
R: Read - W: Write - M: Metadata

```

Difference between overall I/O and Unoverlapped or perceived I/O



Megatron Deepspeed: Characterization and Analysis (DFTracer + WisIO)

WisIO

APP I/O Characteristics

```
App Time Runtime: 9773.62 seconds (100%)
├── Application I/O Time: 4424.42 seconds (45.27%)
├── Checkpoint I/O Time: 4405.64 seconds (45.08%)
├── Compute Time: 7345.04 seconds (75.15%)
├── I/O Time: 2218.73 seconds (22.70%)
├── Read I/O Time: 29.65 seconds (0.30%)
├── Time: 23268.77 seconds (238.08%)
├── Unoverlapped Application I/O Time: 3162.50 seconds (32.36%)
├── Unoverlapped Checkpoint I/O Time: 3162.50 seconds (32.36%)
├── Unoverlapped Compute Time: 5934.06 seconds (60.72%)
├── Unoverlapped I/O Time: 807.75 seconds (8.26%)
├── Unoverlapped Read I/O Time: 0.00 seconds (0.00%)
└── Unoverlapped Application Compute Time: 6083.12 seconds (62.24%)
```

Look at workloads with different lenses.



Details Characterization and Analysis with SCR (DFTracer + WisIO)

POSIX I/O Characteristics	
Runtime	8428.34 seconds
Unoverlapped I/O Time	81.50 seconds (100%) <ul style="list-style-type: none">Read: 3.27 seconds (4.02%)Write: 40.36 seconds (49.52%)Metadata: 22.56 seconds (27.68%)
Unoverlapped I/O Ops.	16,543,145 ops <ul style="list-style-type: none">Read - 1,224,090 ops (7.40%)Write - 73,755 ops (0.45%)Metadata - 13,926,874 ops (84.19%)
Unoverlapped I/O Size	4.20 TiB <ul style="list-style-type: none">Read - 6.47 GiB (0.15%)Write - 4.20 TiB (99.85%)
Read Requests	4 kiB→64 MiB - 1,224,090 ops <ul style="list-style-type: none">16-64 MiB: 177,796 ops (14.52%)>64 MiB: 1,037,568 ops (84.76%)[Others: 8,726 ops (0.71%)]
Write Requests	4 kiB→64 MiB - 73,755 ops <ul style="list-style-type: none">4-16 MiB: 13,701 ops (18.58%)>64 MiB: 59,919 ops (81.24%)[Others: 135 ops (0.18%)]
Time Periods	10 time periods (Time Granularity: 65.00 seconds)
Access Pattern	Sequential : 1,297,845 ops (100.00%) - Random : 0 ops (0.00%)

R: Read - **W**: Write - **M**: Metadata



Details Characterization and Analysis with SCR (DFTracer + WisIO)

```
WisIO
----- APP I/O Characteristics -----
App Time  Runtime: 8428.34 seconds (100%)
         Application I/O Time: 1487.00 seconds (17.64%)
         Checkpoint I/O Time: 1485.76 seconds (17.63%)
         Compute Time: 7327.39 seconds (86.94%)
         I/O Time: 938.97 seconds (11.14%)
         Read I/O Time: 1.69 seconds (0.02%)
         Time: 18436.94 seconds (218.75%)
         Unoverlapped Application I/O Time: 297.31 seconds (3.53%)
         Unoverlapped Checkpoint I/O Time: 297.31 seconds (3.53%)
         Unoverlapped Compute Time: 6469.93 seconds (76.76%)
         Unoverlapped I/O Time: 81.50 seconds (0.97%)
         Unoverlapped Read I/O Time: 0.00 seconds (0.00%)
         Unoverlapped Application Compute Time: 6137.71 seconds (72.82%)

----- POSIX I/O Characteristics -----
Runtime      8428.34 seconds
I/O Time    938.97 seconds (100%)
├── Read: 4.15 seconds (0.44%)
├── Write: 841.95 seconds (89.67%)
└── Metadata: 123.98 seconds (13.20%)

I/O Ops.    43,581,671 ops
├── Read - 1,983,942 ops (4.55%)
├── Write - 600,319 ops (1.38%)
└── Metadata - 20,758,464 ops (47.63%)

I/O Size    20.12 TiB
├── Read - 13.75 GiB (0.07%)
└── Write - 20.11 TiB (99.93%)

Read Requests 4 kiB→64 MiB - 1,983,942 ops
├── 16-64 MiB: 177,796 ops (8.96%)
├── >64 MiB: 1,791,744 ops (90.31%)
└── [Others: 14,402 ops (0.73%)]

Write Requests 4 kiB→64 MiB - 600,319 ops
├── <4 kiB: 282,240 ops (47.02%)
├── 4-16 MiB: 20,498 ops (3.41%)
├── >64 MiB: 297,446 ops (49.55%)
└── [Others: 135 ops (0.02%)]

Time Periods 130 time periods (Time Granularity: 65.00 seconds)
Access Pattern Sequential: 2,584,261 ops (100.00%) - Random: 0 ops (0.00%)

R: Read - W: Write - M: Metadata
```

```
----- POSIX I/O Characteristics -----
Runtime      8428.34 seconds
Unoverlapped I/O Time 81.50 seconds (100%)
├── Read: 3.27 seconds (4.02%)
├── Write: 40.36 seconds (49.52%)
└── Metadata: 22.56 seconds (27.68%)

Unoverlapped I/O Ops. 16,543,145 ops
├── Read - 1,224,090 ops (7.40%)
├── Write - 73,755 ops (0.45%)
└── Metadata - 13,926,874 ops (84.19%)

Unoverlapped I/O Size 4.20 TiB
├── Read - 6.47 GiB (0.15%)
└── Write - 4.20 TiB (99.85%)

Read Requests 4 kiB→64 MiB - 1,224,090 ops
├── 16-64 MiB: 177,796 ops (14.52%)
├── >64 MiB: 1,037,568 ops (84.76%)
└── [Others: 8,726 ops (0.71%)]

Write Requests 4 kiB→64 MiB - 73,755 ops
├── 4-16 MiB: 13,701 ops (18.58%)
├── >64 MiB: 59,919 ops (81.24%)
└── [Others: 135 ops (0.18%)]

Time Periods 10 time periods (Time Granularity: 65.00 seconds)
Access Pattern Sequential: 1,297,845 ops (100.00%) - Random: 0 ops (0.00%)

R: Read - W: Write - M: Metadata
```



Megatron Deepspeed (DFTracer + DFAnalyzer)

```
Allocation Scheduler Allocation Details
├── Nodes: 16
├── Processes: 256
├── Inthread allocations across nodes (includes d
│   ├── Compute: 128
│   └── I/O: 31067
├── Events Recorded: 135M
Dataset Description of Dataset Used
├── Files: 484081
I/O Behavior Behavior of Application
├── Split of Time in application
│   ├── Total Time: 9773.622 sec
│   ├── Overall App Level I/O: 4514.237 sec
│   ├── Unoverlapped App I/O: 3659.268 sec
│   ├── Unoverlapped App Compute: 6487.549 sec
│   ├── Compute: 7342.518 sec
│   ├── Overall I/O: 2338.010 sec
│   ├── Unoverlapped I/O: 1035.517 sec
│   └── Unoverlapped Compute: 6040.025 sec
```

- Checkpointing + Reading (App I/O)
- Overlapped I/O is Reading uses MMap
 - We cannot see kernel's page faults
- Checkpointing is Unoverlapped
 - It is synchronous wrt the computation.



Megatron Deepspeed w/wo SCR (WisIO)

Reduced overall runtime is a result of increasing overlap between checkpoint I/O and Computation

```
Runtime 9773.62 seconds
App Time App Time Breakdown
  | Compute Time: 7345.04 seconds (75.15%)
  | I/O Time: 2218.73 seconds (22.70%)
  | Application I/O Time: 4424.42 seconds (45.27%)
  | Read I/O Time: 29.65 seconds (0.30%)
  | Checkpoint I/O Time: 4405.64 seconds (45.08%)
  | Unoverlapped Application Compute Time: 6083.12 seconds (62.24%)
  | Unoverlapped Application I/O Time: 3162.50 seconds (32.36%)
  | Unoverlapped Compute Time: 5934.06 seconds (60.72%)
  | Unoverlapped I/O Time: 807.75 seconds (8.26%)
  | Unoverlapped Checkpoint I/O Time: 3162.50 seconds (32.36%)
  | Unoverlapped Read I/O Time: 0.00 seconds (0.00%)
```

```
Runtime 8428.34 seconds
App Time App Time Breakdown
  | Compute Time: 7327.39 seconds (86.94%)
  | I/O Time: 938.97 seconds (11.14%)
  | Application I/O Time: 1487.00 seconds (17.64%)
  | Read I/O Time: 1.69 seconds (0.02%)
  | Checkpoint I/O Time: 1485.76 seconds (17.63%)
  | Unoverlapped Application Compute Time: 6137.71 seconds (72.82%)
  | Unoverlapped Application I/O Time: 297.31 seconds (3.53%)
  | Unoverlapped Compute Time: 6469.93 seconds (76.76%)
  | Unoverlapped I/O Time: 81.50 seconds (0.97%)
  | Unoverlapped Checkpoint I/O Time: 297.31 seconds (3.53%)
  | Unoverlapped Read I/O Time: 0.00 seconds (0.00%)
```



Megatron Deepspeed with SCR

```
Allocation Scheduler Allocation Details
├── Nodes: 16
├── Processes: 256
├── Thread allocations across nodes (includes dy
│   ├── Compute: 128
│   └── I/O: 14784
└── Events Recorded: 64M

Dataset Description of Dataset Used
└── Files: 462527

I/O Behavior Behavior of Application
├── Split of Time in application
│   ├── Total Time: 8641.546 sec
│   ├── Overall App Level I/O: 1932.854 sec
│   ├── Unoverlapped App I/O: 993.126 sec
│   ├── Unoverlapped App Compute: 6392.047 sec
│   ├── Compute: 7331.776 sec
│   ├── Overall I/O: 1529.879 sec
│   ├── Unoverlapped I/O: 445.952 sec
└── Unoverlapped Compute: 6247.848 sec
```

- Performs more I/O first to SSD then async to PFS
- Improves overall time by 1132 secs.
- 2x speed up on Checkpointing
- Improved overlapping: from 20% to 50%.

