I/O Workload Overview of the Applications on Intrepid Supercomputer

Pablo J. Pavan , Valéria S. Girelli,
Jean L. Bez, Francieli Z. Boito,
Philippe O. A. Navaux

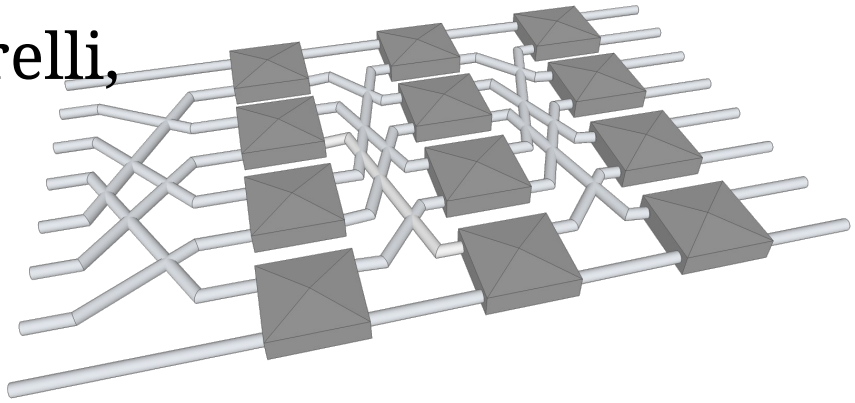

WSPPD 2018
September 05, 2018

Summary

- Motivation
- Goal
- Methodology
- Results
- Conclusion

Motivation

- HPC systems allow complex simulations
  - Medical
  - Oil and Gas Exploration
  - Weather forecasting
- I/O is an important part of HPC applications

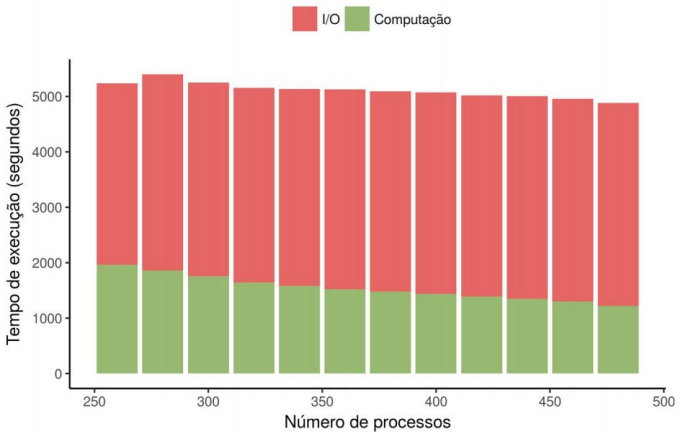

Data collected in the SantosDumond Supercomputer with the Olam application

Motivation

- Characterizing I/O behavior can contribute:
  - I/O scheduler reconfiguration
  - I/O Stack reconfiguration
  - Creating or improving burst-buffer
  - I/O Forwarding

GOAL

Investigate the I/O workload of the applications executed on the Intrepid Supercomputer

Methodology

- DARSHAN - HPC I/O Characterization Tool
  - Developed by Argonne L. C. F.
  - Profile I/O operations at the application level
  - Application information
    - Application identifier
    - Job identifier
    - User identifier
    - Number of MPI processes
    - Execution time

Methodology

- DARSHAN - HPC I/O Characterization Tool
  - Developed by Argonne L. C. F.
  - Profile I/O operations at the application level
  - Counters
    - Individual/Collective Access
    - Interface (POSIX and MPI-IO)
    - Access sizes
    - Data Transferred
    - Number of I/O operation
    - Time spent in I/O operations

Methodology

- Supercomputer Intrepid Blue Gene/P, Argonne USA
  - 91,994 captured jobs during 2012
  - Anonymized information

Methodology

- Our process to analyze the DARSHAN data

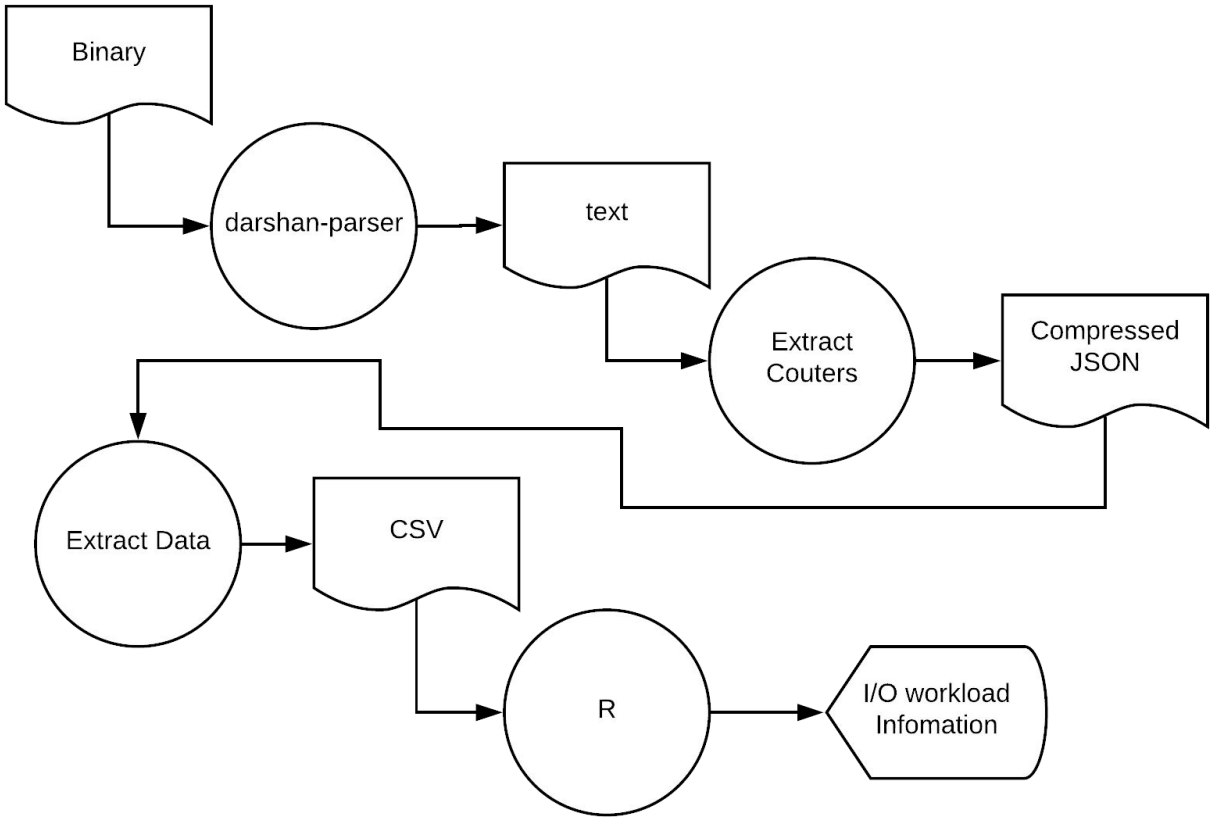

Results

- Overview of the 2012 dataset

Results

- Overview of the 2012 dataset
- The analysis of the applications whose I/O percentage was above 50%

Results

- Overview of the 2012 dataset
- The analysis of the applications whose I/O percentage was above 50%
- Top ten most executed applications during 2012

Results

- Overview of the 2012 dataset
  - 91,994 jobs resulted in 26,034 different applications
  - Executions with a high I/O percentage do not run for long periods of time
    - 45,000 seconds executing
    - 2,500 seconds performing I/O

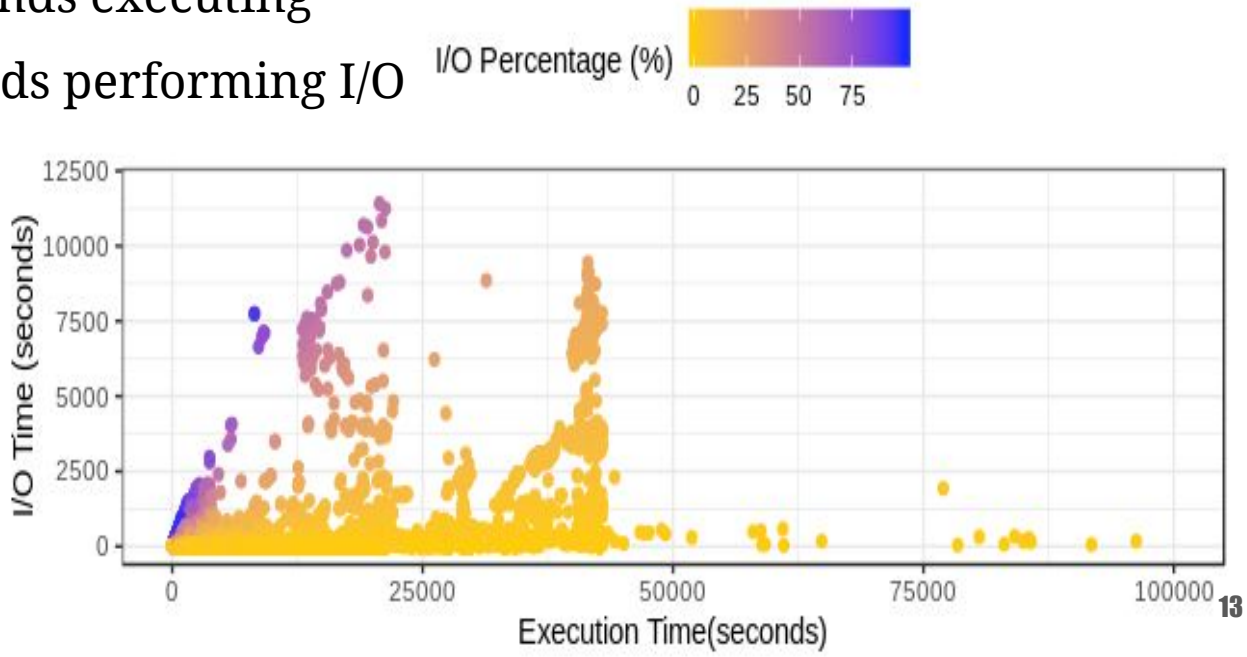

Results

- Overview of the 2012 dataset
  - 91,994 jobs resulted in 26,034 different applications
  - Executions with a high I/O percentage do not run for long periods of time
    - 45,000 seconds executing
    - 2,500 seconds performing I/O

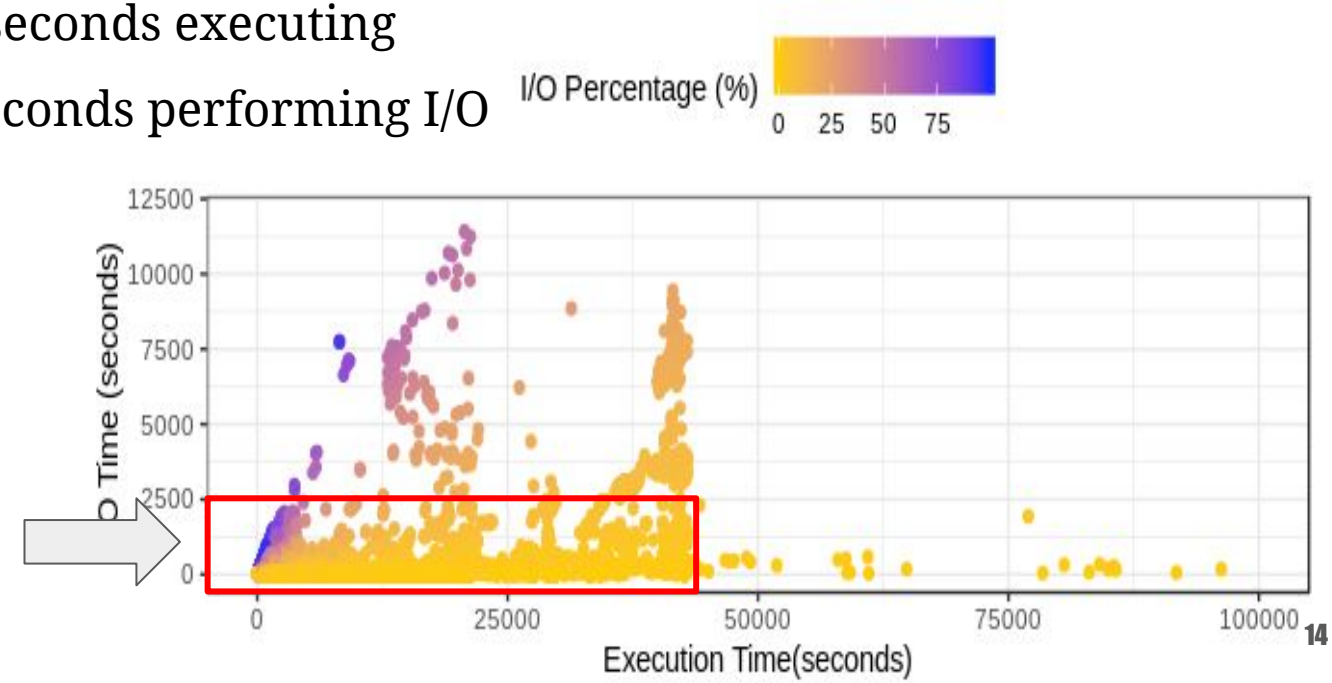

- 85,710 jobs
- 10,050 jobs have an I/O percentage above 50%

Results

- The analysis of the applications whose I/O percentage was above 50%
  - We selected three applications

Exec (Anonymised)	Jobs with > 50% of the time in I/O operations
1176110786	**980**
1338247359	**898**
902685977	**761**
Total	2.639

Results

- The analysis of the applications whose I/O percentage was above 50%
- Exec 1176110786
  - 1,024 or 2048 MPI processes
  - Executed 980 times
  - From June 06, 2012 to December 18, 2012

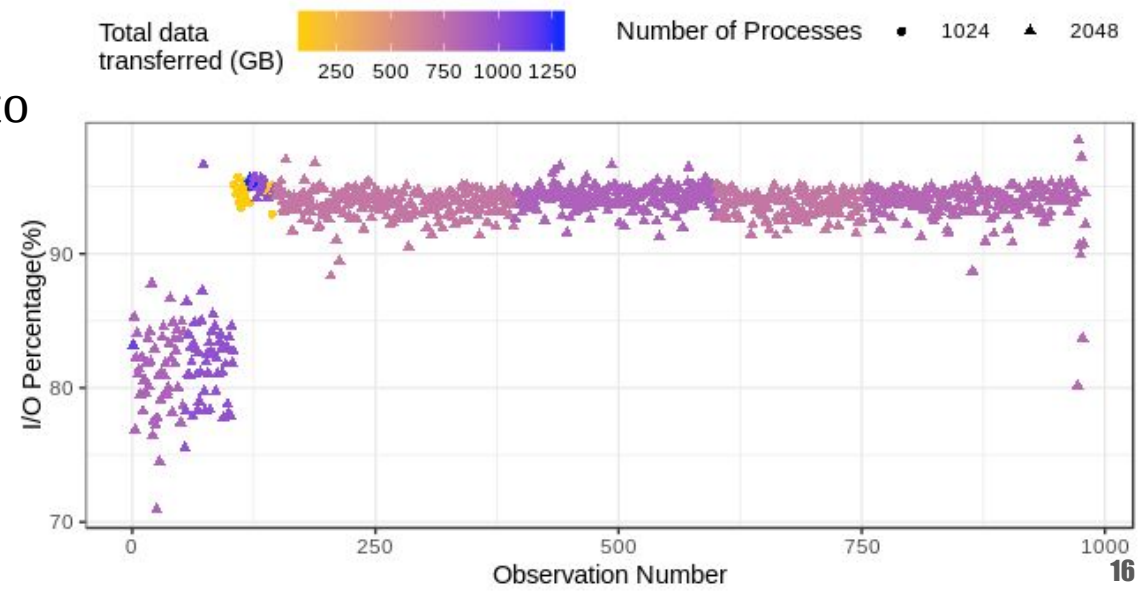

Results

- The analysis of the applications whose I/O percentage was above 50%
- Exec 1176110786
  - 1,024 or 2048 MPI processes
  - Executed 980 times
  - From June 06, 2012 to December 18, 2012

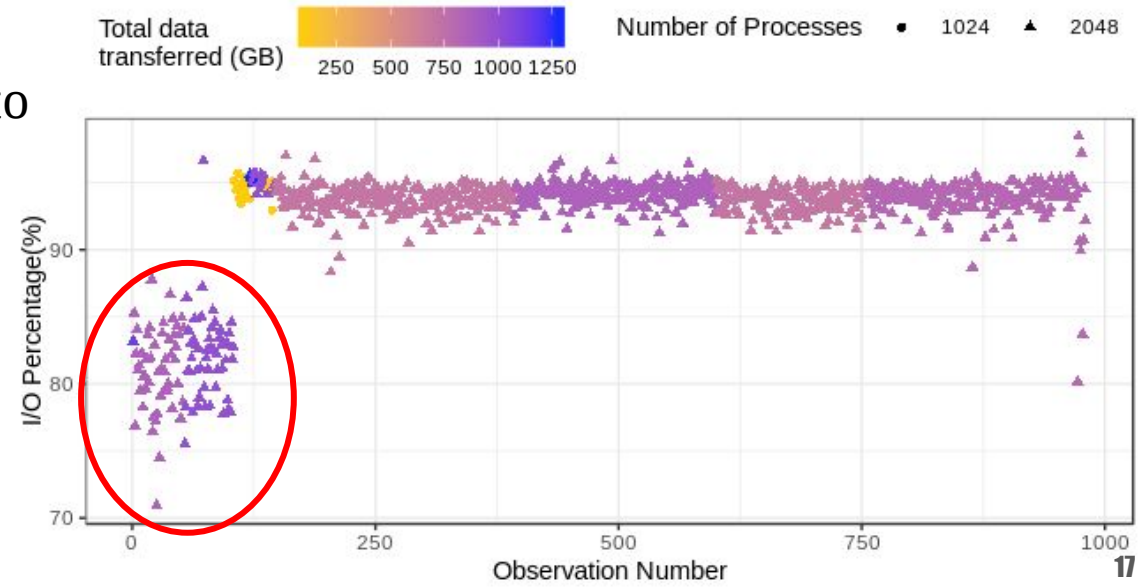

Results

- The analysis of the applications whose I/O percentage was above 50%
- Exec 1176110786
  - 1,024 or 2048 MPI processes
  - Executed 980 times
  - From June 06, 2012 to December 18, 2012

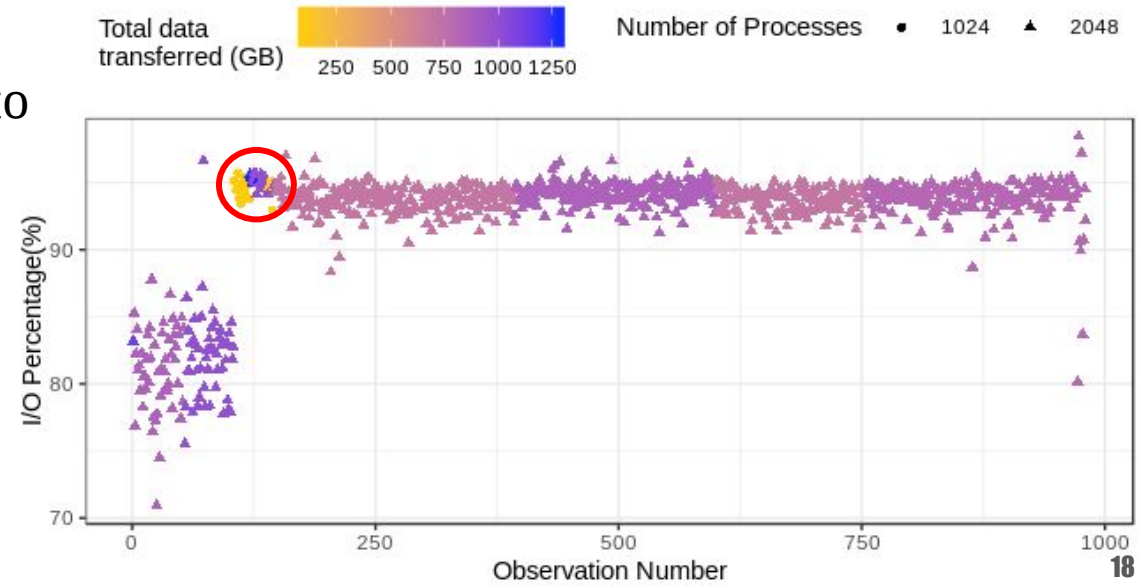

Results

- The analysis of the applications whose I/O percentage was above 50%
- Exec 902685977
  - 32, 64 and 128 MPI processes
  - Executed 763 times
  - From October 23, 2012 to October 30, 2012

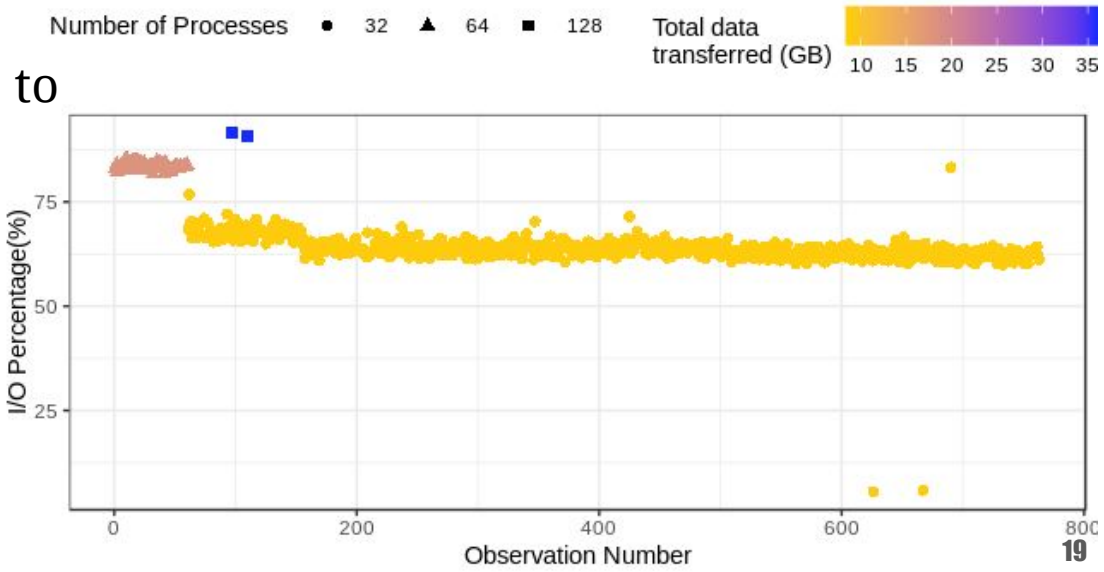

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Exec (Anonymised)	**Number of Observations**	Execution time (seconds)	I/O Time (seconds)	I/O %	Data transferred (GB)
685531913	**8.016**	1.370	3,12	0,22	0,003
1330277471	**5.176**	1.314	536,63	40,84	5,69
931947437	**4.922**	2.906	985,65	33,91	0,013
3069475893	**3.191**	2.808	3,92	0,13	0,012
1074553177	**2.690**	1.489	533,55	35,83	1,12
1648769576	**2.588**	1.993	191,82	9,62	7,59
1633035531	**2.259**	42.003	8.176,28	19,46	6223,78
2425255765	**1.339**	3.720	2.967,61	79,77	2,85
3475271559	**1.303**	1.512	310,14	20,51	2,25
1338247359	**1.051**	2.415	1.924,92	79,70	17,80

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Exec (Anonymised)	Number of Observations	Execution time (seconds)	**I/O Time (seconds)**	I/O %	Data transferred (GB)
685531913	8.016	1.370	**3,12**	0,22	0,003
1330277471	5.176	1.314	**536,63**	40,84	5,69
931947437	4.922	2.906	**985,65**	33,91	0,013
3069475893	3.191	2.808	**3,92**	0,13	0,012
1074553177	2.690	1.489	**533,55**	35,83	1,12
1648769576	2.588	1.993	**191,82**	9,62	7,59
1633035531	2.259	42.003	**8.176,28**	19,46	6223,78
2425255765	1.339	3.720	**2.967,61**	79,77	2,85
3475271559	1.303	1.512	**310,14**	20,51	2,25
1338247359	1.051	2.415	**1.924,92**	79,70	17,80

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Exec (Anonymised)	Number of Observations	Execution time (seconds)	I/O Time (seconds)	**I/O %**	Data transferred (GB)
685531913	8.016	1.370	3,12	**0,22**	0,003
1330277471	5.176	1.314	536,63	**40,84**	5,69
931947437	4.922	2.906	985,65	**33,91**	0,013
3069475893	3.191	2.808	3,92	**0,13**	0,012
1074553177	2.690	1.489	533,55	**35,83**	1,12
1648769576	2.588	1.993	191,82	**9,62**	7,59
1633035531	2.259	42.003	8.176,28	**19,46**	6223,78
2425255765	1.339	3.720	2.967,61	**79,77**	2,85
3475271559	1.303	1.512	310,14	**20,51**	2,25
1338247359	1.051	2.415	1.924,92	**79,70**	17,80

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Exec (Anonymised)	Number of Observations	Execution time (seconds)	I/O Time (seconds)	**I/O %**	Data transferred (GB)
685531913	8.016	1.370	3,12	**0,22**	0,003
1330277471	5.176	1.314	536,63	**40,84**	5,69
931947437	4.922	2.906	985,65	**33,91**	0,013
3069475893	3.191	2.808	3,92	**0,13**	0,012
1074553177	2.690	1.489	533,55	**35,83**	1,12
1648769576	2.588	1.993	191,82	**9,62**	7,59
1633035531	2.259	42.003	8.176,28	**19,46**	6223,78
2425255765	1.339	3.720	2.967,61	**79,77**	2,85
3475271559	1.303	1.512	310,14	**20,51**	2,25
1338247359	1.051	2.415	1.924,92	**79,70**	17,80

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Exec (Anonymised)	Number of Observations	Execution time (seconds)	I/O Time (seconds)	**I/O %**	Data transferred (GB)
685531913	8.016	1.370	3,12	**0,22**	0,003
1330277471	5.176	1.314	536,63	**40,84**	5,69
931947437	4.922	2.906	985,65	**33,91**	0,013
3069475893	3.191	2.808	3,92	**0,13**	0,012
1074553177	2.690	1.489	533,55	**35,83**	1,12
1648769576	2.588	1.993	191,82	**9,62**	7,59
1633035531	2.259	42.003	8.176,28	**19,46**	6223,78
2425255765	1.339	3.720	2.967,61	**79,77**	2,85
3475271559	1.303	1.512	310,14	**20,51**	2,25
1338247359	1.051	2.415	1.924,92	**79,70**	17,80

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Exec (Anonymised)	Number of Observations	Execution time (seconds)	I/O Time (seconds)	I/O %	Data transferred (GB)
685531913	8.016	1.370	3,12	**0,22**	0,003
1330277471	5.176	1.314	536,63	⇨ **40,84**	5,69
931947437	4.922	2.906	985,65	⇨ **33,91**	0,013
3069475893	3.191	2.808	3,92	**0,13**	0,012
1074553177	2.690	1.489	533,55	⇨ **35,83**	1,12
1648769576	2.588	1.993	191,82	**9,62**	7,59
1633035531	2.259	42.003	8.176,28	⇨ **19,46**	6223,78
2425255765	1.339	3.720	2.967,61	⇨ **79,77**	2,85
3475271559	1.303	1.512	310,14	⇨ **20,51**	2,25
1338247359	1.051	2.415	1.924,92	⇨ **79,70**	17,80

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Exec (Anonymised)	Number of Observations	Execution time (seconds)	I/O Time (seconds)	I/O %	**Data transferred (GB)**
685531913	8.016	1.370	3,12	0,22	**0,003**
1330277471	5.176	1.314	536,63	40,84	**5,69**
931947437	4.922	2.906	985,65	33,91	**0,013**
3069475893	3.191	2.808	3,92	0,13	**0,012**
1074553177	2.690	1.489	533,55	35,83	**1,12**
1648769576	2.588	1.993	191,82	9,62	**7,59**
1633035531	2.259	42.003	8.176,28	19,46	**6223,78**
2425255765	1.339	3.720	2.967,61	79,77	**2,85**
3475271559	1.303	1.512	310,14	20,51	**2,25**
1338247359	1.051	2.415	1.924,92	79,70	**17,80**

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Exec (Anonymised)	Number of Observations	Execution time (seconds)	I/O Time (seconds)	I/O %	**Data transferred (GB)**
685531913	8.016	1.370	3,12	0,22	**0,003**
1330277471	5.176	1.314	536,63	40,84	**5,69**
931947437	4.922	2.906	985,65	33,91	**0,013**
3069475893	3.191	2.808	3,92	0,13	**0,012**
1074553177	2.690	1.489	533,55	35,83	**1,12**
1648769576	2.588	1.993	191,82	9,62	**7,59**
1633035531	2.259	42.003	8.176,28	19,46	**6223,78**
2425255765	1.339	3.720	2.967,61	79,77	**2,85**
3475271559	1.303	1.512	310,14	20,51	**2,25**
1338247359	1.051	2.415	1.924,92	79,70	**17,80**

Results

- Top ten most executed applications during the year
  - The total executions count was 32,535
  - Which represents 35.36%

Exec (Anonymised)	Number of Observations	Execution time (seconds)	I/O Time (seconds)	I/O %	**Data transferred (GB)**
685531913	8.016	1.370	3,12	0,22	**0,003**
1330277471	5.176	1.314	536,63	40,84	**5,69**
931947437	4.922	2.906	985,65	33,91	**0,013**
3069475893	3.191	2.808	3,92	0,13	**0,012**
1074553177	2.690	1.489			
1648769576	2.588	1.993			
1633035531	2.259	42.003			
2425255765	1.339	3.720			
3475271559	1.303	1.512	310,14	20,51	**2,25**
1338247359	1.051	2.415	1.924,92	79,70	**17,80**

This case can indicate:
- Problem with the program level
- I/O contention

Conclusion

- We analyzed the results of I/O workload found during the year 2012 on the Intrepid supercomputer
- We obtained **91,973** submitted jobs that represent **26,034** different applications

Conclusion

- Applications that had spend **more than 50%** of their **time in I/O**
- We could notice **different behaviors**:
  - MPI process numbers
  - Transferred data
  - Execution period during the year

Conclusion

- Applications that had **more than 50%** of spend **time I/O**
- We could notice **different behaviors**:
  - MPI process numbers
  - Transferred data
  - Execution period during the year
- **Top ten** most executed applications
  - Representing 35.36% of the total number of executions observed
  - In seven applications:
    - More than 19% of their execution time spent in I/O time

Conclusion

- As future work
  - We intend to **extend this analysis** of the **characterization of the intervals** for **all the jobs of Intrepid** seeking to **group applications** that have a **similar I/O behavior**

Conclusion

- As future work
  - We intend to **extend this analysis** of the **characterization of the intervals** for **all the jobs of Intrepid** seeking to **group applications** that have a **similar I/O behavior**
- We also intend to **aggregate data** from **multiple applications** to have a **global view** of the use of **I/O in a supercomputer**

Acknowledgment

- This research received funding from the **Petrobras project**, grant n. 2016/00133-9. It was also supported by **PIBIC CNPq-UFRGS** and **PROBIC FAPERGS-UFRGS**
- This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357

I/O Workload Overview of the Applications on Intrepid Supercomputer

Pablo J. Pavan , Valéria S. Girelli,
Jean L. Bez, Francieli Z. Boito,
Philippe O. A. Navaux

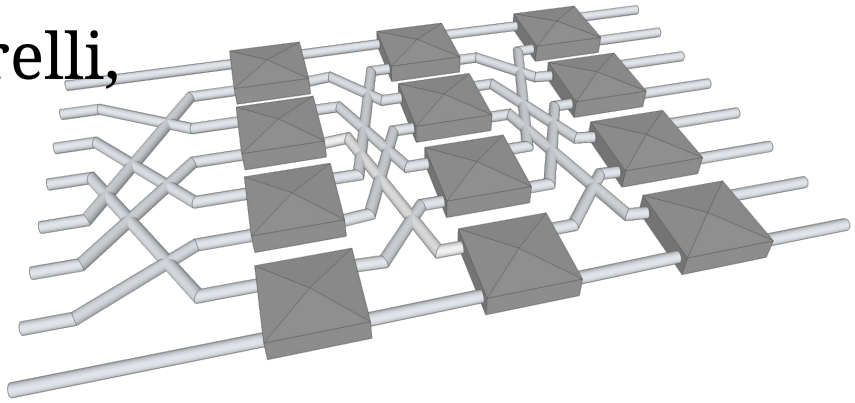

WSPPD 2018
September 05, 2018