A Study of Multiple Platforms' I/O Workload

HUONG LUU, WILLIAM GROPP, MARIANNE WINSLETT

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

We use lightweight I/O profiler to collect applications' I/O characteristics.

Darshan is enabled by default on Mira, Edison and Blue Waters.

	Mira	Edison	Blue Waters
Architecture	BG/Q	Cray XC30	Cray XE6/XK7
Peak Flops	10 PF	2.57 PF	13.34 PF
Cores	768K	130K	792K+59K smx
Total Storage	24 PB	7.56 PB	26.4 PB
Peak I/O Throughput	240 GB/s	168 GB/s	963 GB/s
File System	GPFS	Lustre	Lustre
Time period	Jan – Sep'15	Jan – Sep'**14**	Jan – Sep'15
Number of jobs	***152K***	***703K***	***570K***
Number of apps	***455***	***1183***	***1657***

Table 1: Target platforms' characteristics

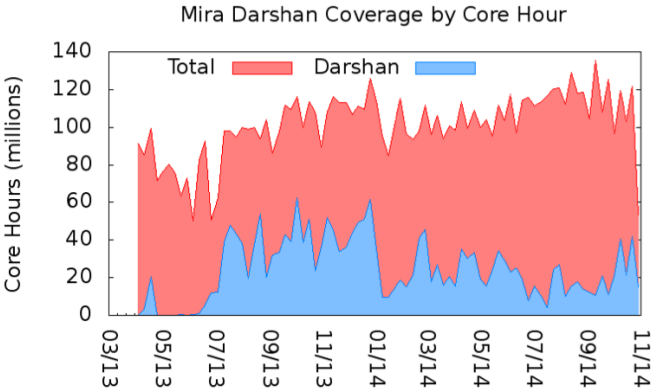

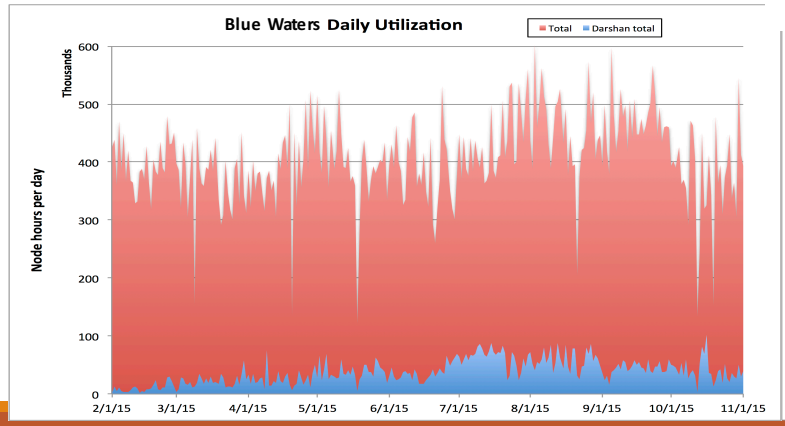

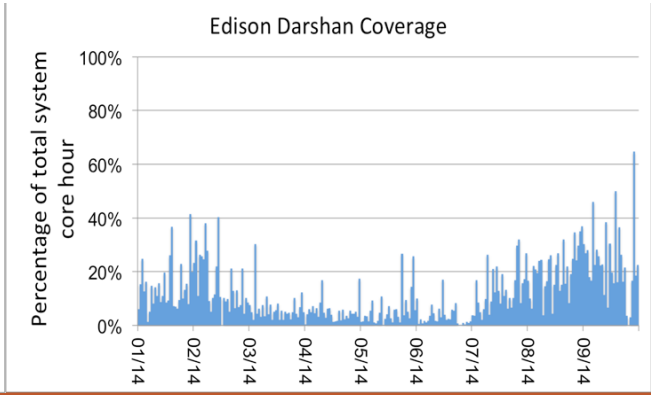

These platforms have similarly distributed I/O workload.

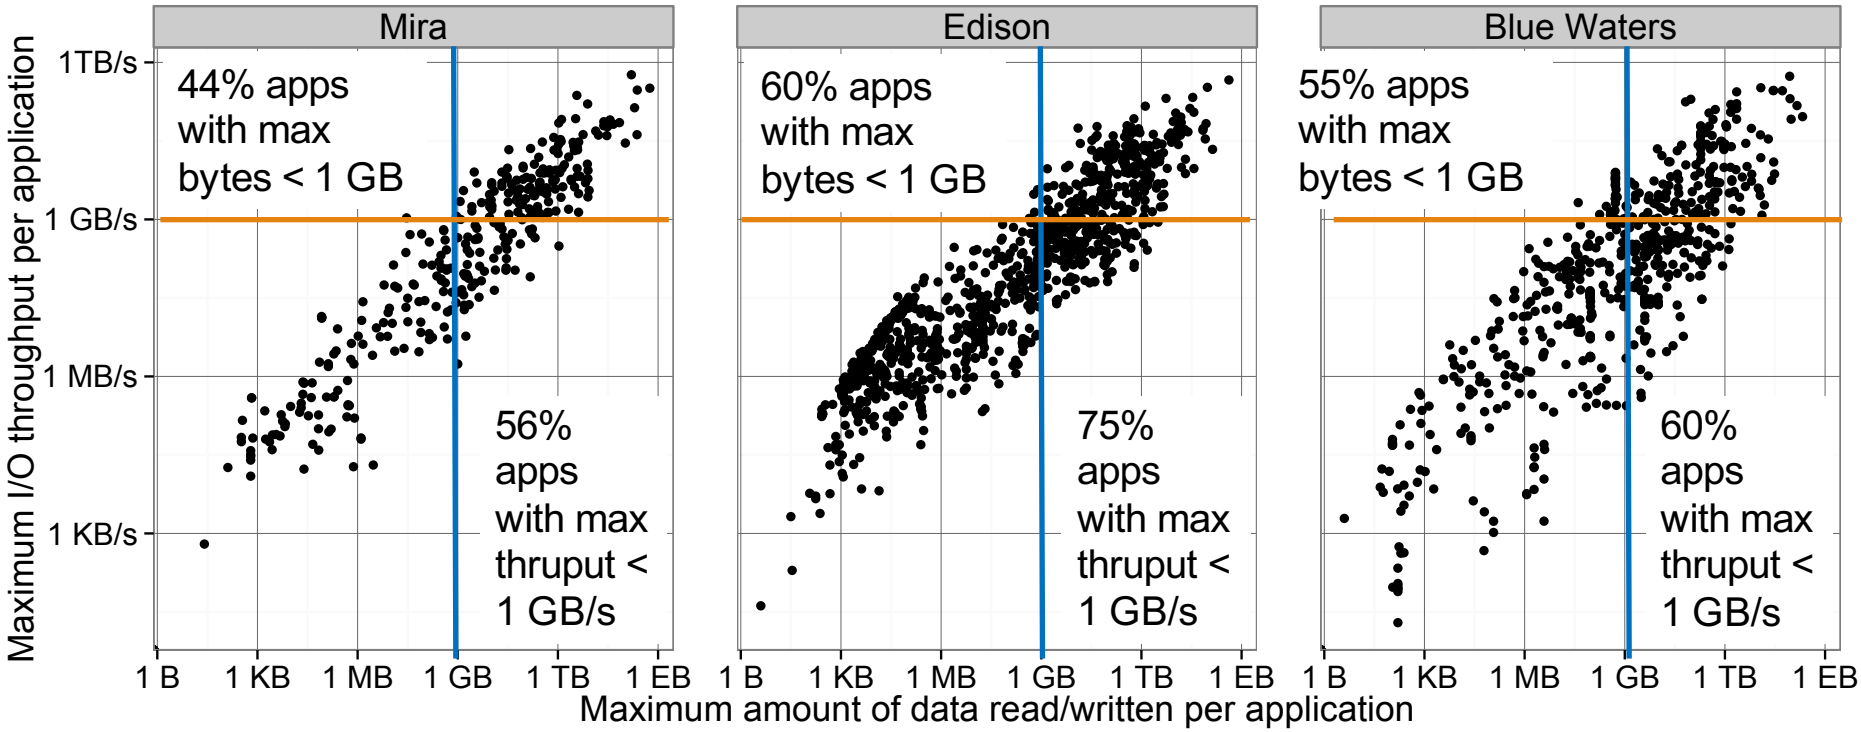

Job Size Distribution per I/O Throughput

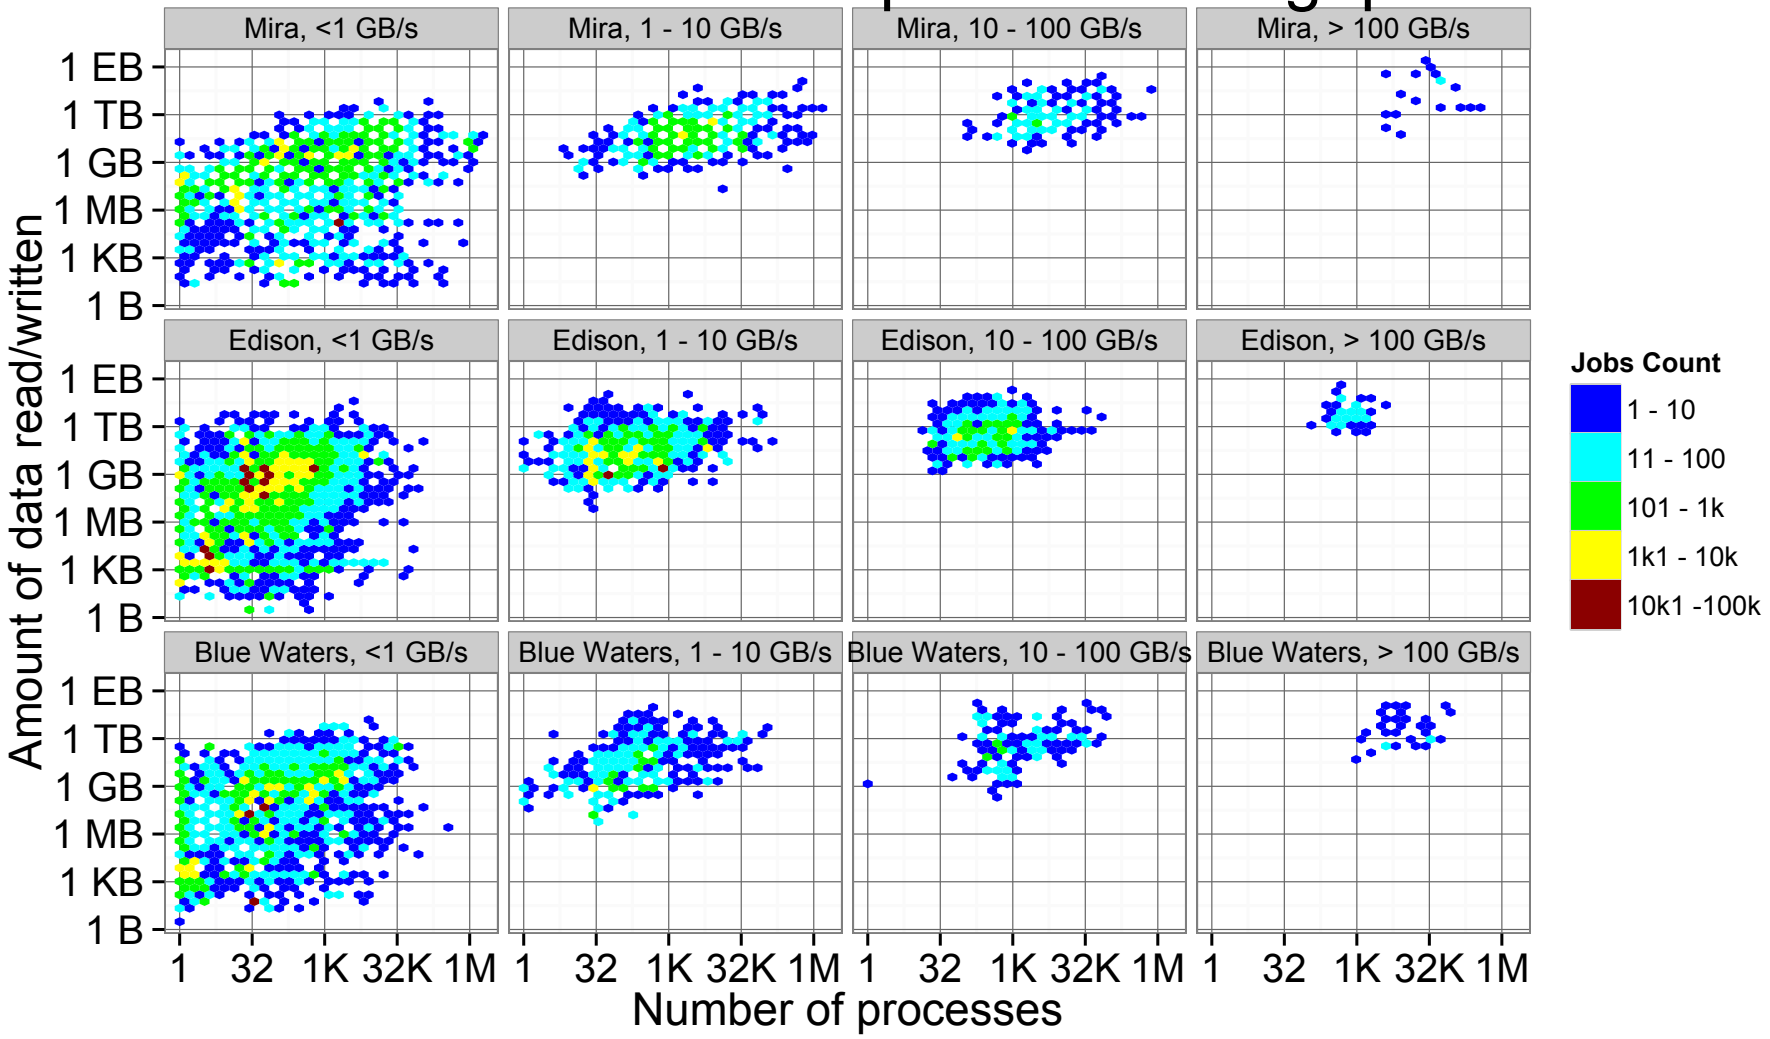

Choices of I/O interface

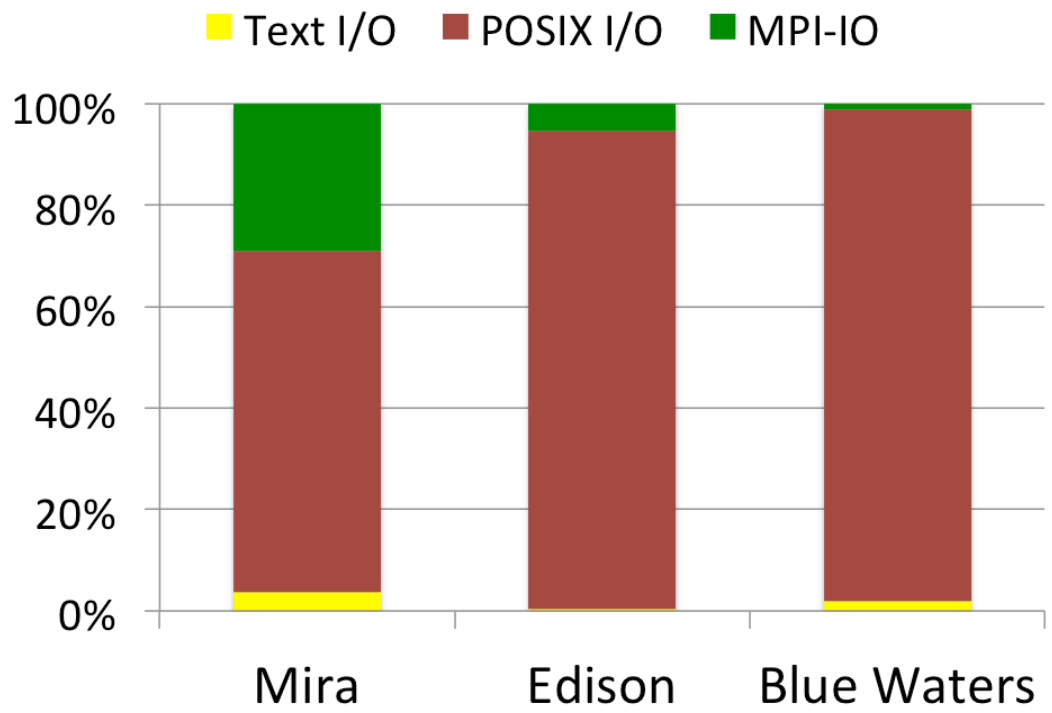

Choices of I/O interface

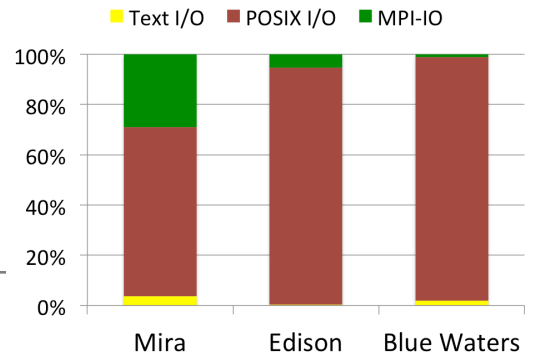

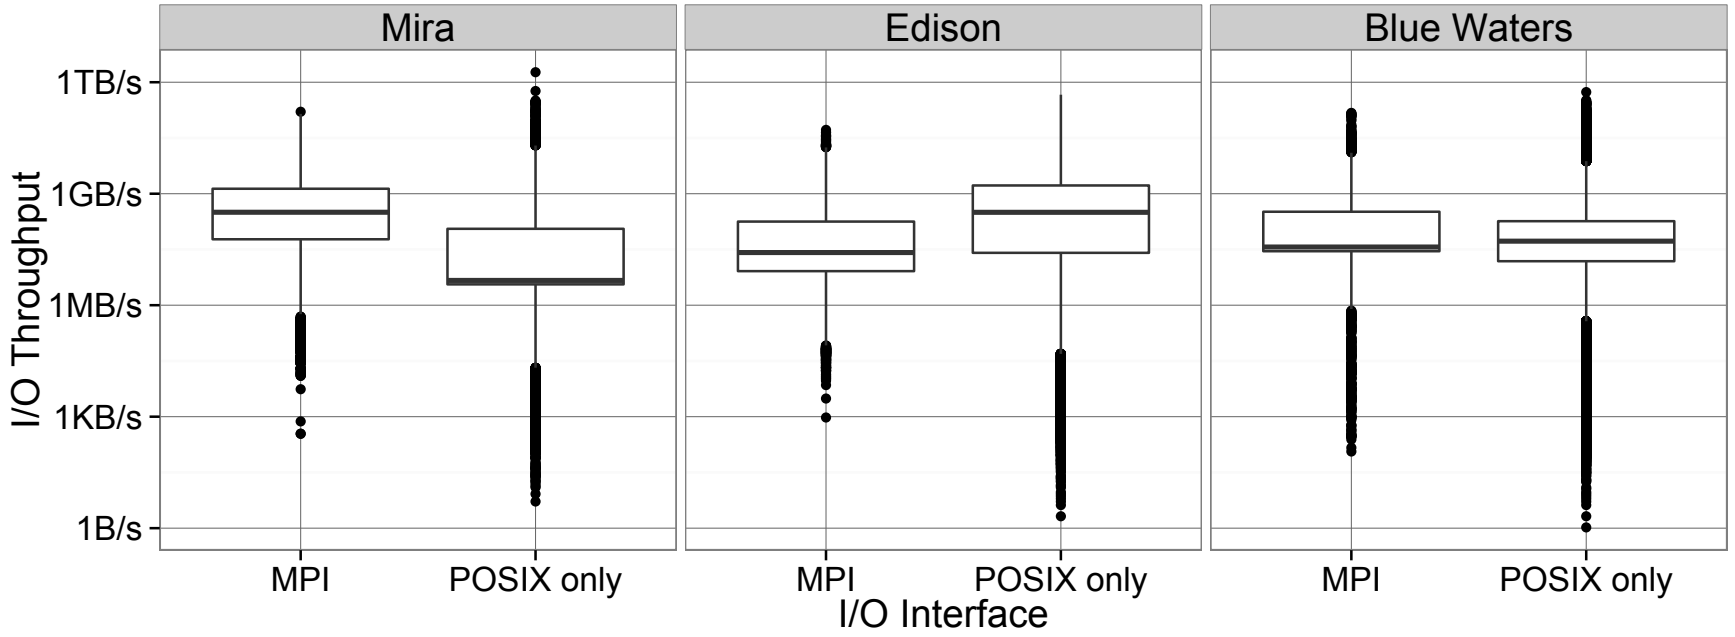

Open Questions

1. What would you most like to see in these graphs that we don't currently show?
   - Additional analysis
   - Collecting additional data to improve the interpretation of the data
2. Is there a difference in I/O performance between Lustre and GPFS based systems?
   - Does this reflect the performance capabilities and differences of Lustre and GPFS, the I/O approach of the applications workload, or is just a co-incidence?

Thank you very much for your attention!