

SIZING BURST-BUFFERS EFFICIENTLY

Guillaume Aupy, Olivier Beaumont, Lionel Eyraud-Dubois

Inria & University of Bordeaux, France



Scheduling Workshop, Berkeley, June 2018

IO congestion in HPC systems:

- ▶ HPC applications are generating lots of data for PFS.
- ▶ Idea is to use a buffer when the I/O bandwidth is fully occupied
- ▶ The buffer can be emptied at a later time.

STORAGE BANDWIDTH UTILIZATION OF A MAJOR HPC PRODUCTION STORAGE SYSTEM

- 99% of the time < 33% of max
- 70% of the time < 5% of max

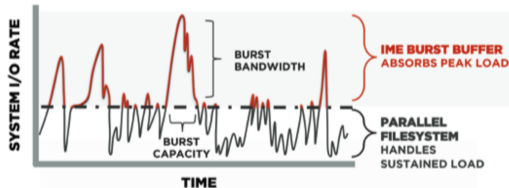


Figure: Burst-buffers to absorb IO peaks

Source: DDN ad material.

SOME USE OF BURST-BUFFERS

Historically, Burst-Buffers were attached to IONodes (ION), used as buffers when the I/O Bandwidth was not enough (Gordon@SDSC).

But many other possible uses:

- ▶ For temporary data that may eventually not be needed (e.g. fault-tolerance)
- ▶ For intermediate data (e.g. BigData on HPC machine, In-situ/In-transit)
- ▶ For other uses?

SOME USE OF BURST-BUFFERS

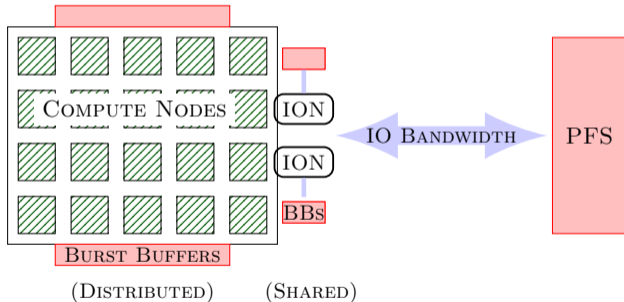
Historically, Burst-Buffers were attached to IONodes (ION), used as buffers when the I/O Bandwidth was not enough (Gordon@SDSC).

But many other possible uses:

- ▶ For temporary data that may eventually not be needed (e.g. fault-tolerance)
- ▶ For intermediate data (e.g. BigData on HPC machine, In-situ/In-transit)
- ▶ For other uses?

How do we choose the *right* amount of Burst-Buffers for each use?

A FEW QUESTIONS TO ANSWER (I)



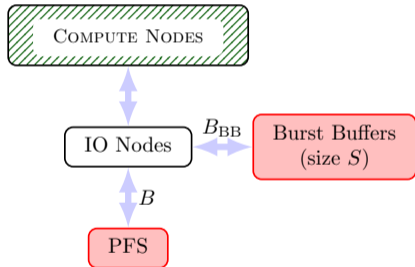
Application Models:

- ▶ Compute and I/O behaviors, buffer needs?
- ▶ Performance model of application?

Distributed Buffers:

- ▶ How to partition the buffers amongst applications?
- ▶ Location of data?

A FEW QUESTIONS TO ANSWER (II)



Centralized buffers for I/O management:

- ▶ What bandwidth B_{BB} ?
- ▶ What size S ?
- ▶ What filling/emptying policy?

Filling policies

- ▶ Use as a **Cache** to PFS.
 - ▶ Pro: More efficient, lower latency
 - ▶ Cons: If SSD-based, limited write-life
- ▶ Use as a **Buffer** when too many simultaneous I/O calls

FILLING/EMPTYING POLICY

Filling policies

- ▶ Use as a **Cache** to PFS.
 - ▶ Pro: More efficient, lower latency
 - ▶ Cons: If SSD-based, limited write-life
- ▶ Use as a **Buffer** when too many simultaneous I/O calls

Emptying policies

- ▶ When some I/O bandwidth is available, empty as much as possible.
- ▶ When some I/O bandwidth is available, AND Burst-Buffers are at least $T\%$ full, empty as much as possible.

APPLICATION MODEL

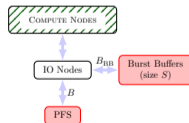
We consider a unit time characteristic of the system.

Applications: At any time unit, application \mathcal{A}_i sends data:

- ▶ with probability p_i
- ▶ at bandwidth b_i .

Machine is characterized by:

- ▶ The Burst Buffer size S
- ▶ Its expected IO load: $\text{EXPECTEDLOAD} = \sum_i p_i b_i$;
- ▶ Its bandwidth to PFS: B



WHAT BANDWIDTH B_{BB} ?

X_i : random variable indicating whether \mathcal{A}_i is sending I/Os.

→ $X_i = 1$ with proba p_i and 0 with $1 - p_i$.

Instant bandwidth $X = \sum_i b_i X_i$

WHAT BANDWIDTH B_{BB} ?

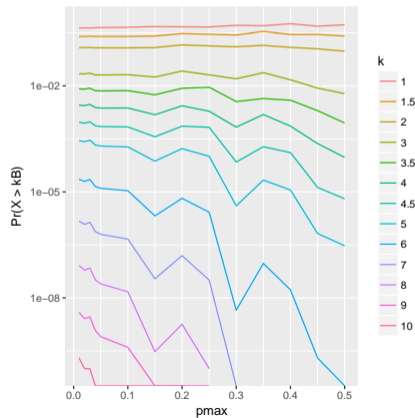
X_i : random variable indicating whether \mathcal{A}_i is sending I/Os.

→ $X_i = 1$ with proba p_i and 0 with $1 - p_i$.

$$\text{Instant bandwidth } X = \sum_i b_i X_i$$

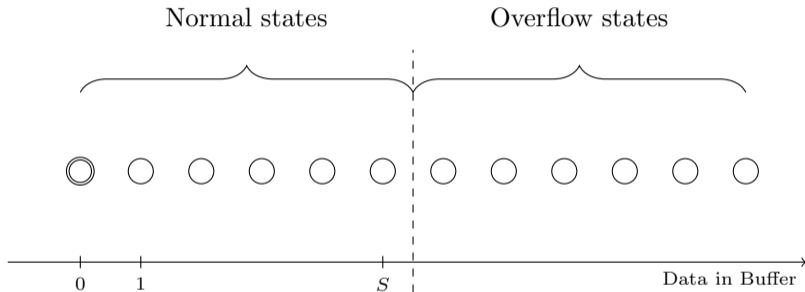
Simulation setup: we fix p_{max} .

- ▶ While $\sum_i p_i b_i < B$:
- ▶ Create a new application with p_i (resp. b_i) chosen uniformly at random in $[0, p_{\text{max}}]$ (resp. $[0, B]$).



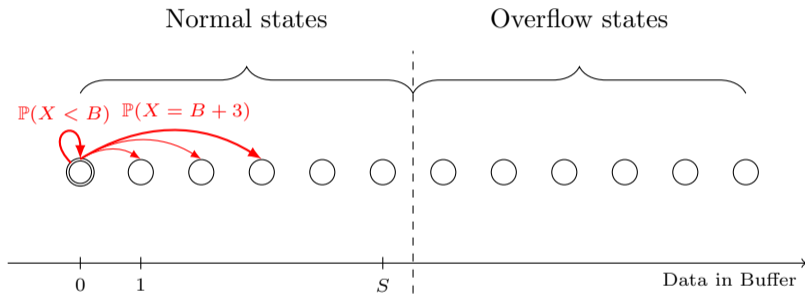
MODELING WITH MARKOV CHAINS

Platform model: when buffer full, stall all applications for one time unit



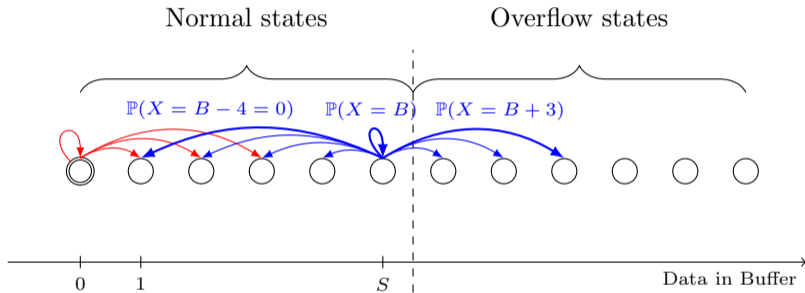
MODELING WITH MARKOV CHAINS

Platform model: when buffer full, stall all applications for one time unit



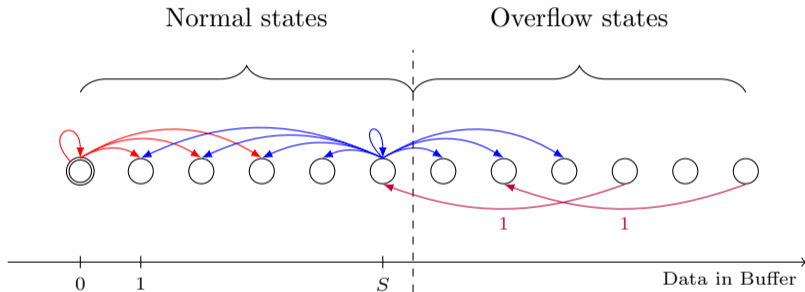
MODELING WITH MARKOV CHAINS

Platform model: when buffer full, stall all applications for one time unit



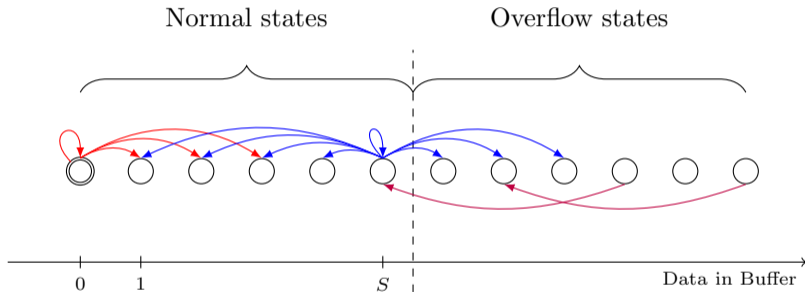
MODELING WITH MARKOV CHAINS

Platform model: when buffer full, stall all applications for one time unit



MODELING WITH MARKOV CHAINS

Platform model: when buffer full, stall all applications for one time unit

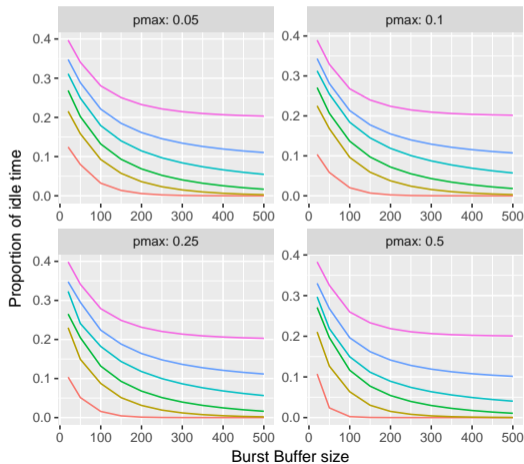


Results

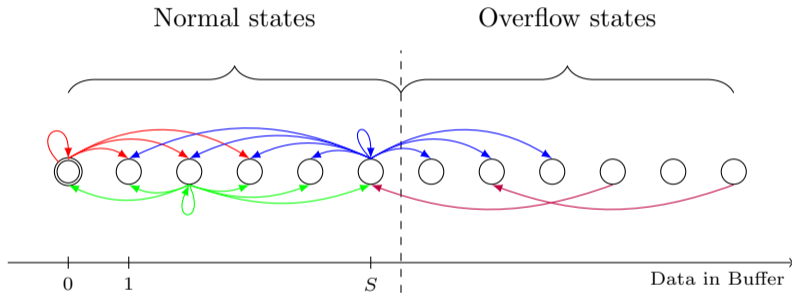
- ▶ Aperiodic and irreducible MC: unique stationary probability π
- ▶ Fraction of time spent idle is $\sum_{s \in \text{overflow states}} \pi_s$

WHAT SIZE S ?

Proportion of idle time as a function of buffer size, for different values of p_{\max} and stress $\alpha = \frac{\text{EXPECTEDLOAD}}{B}$ ($B = 100$)

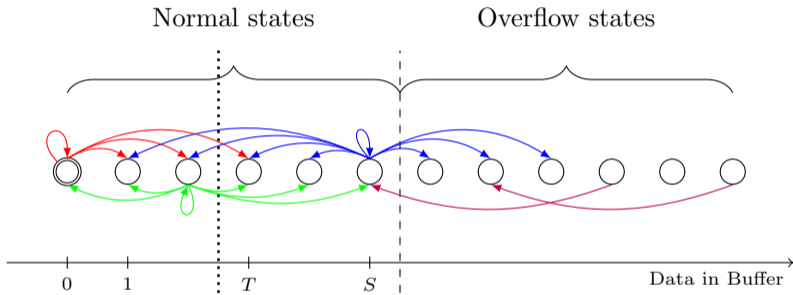


Lazy Emptying [Cluster 2017]: Only empty the burst buffer when its load reaches a threshold T .



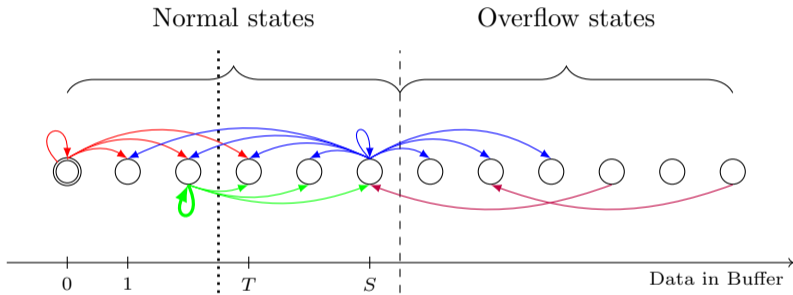
WHAT POLICY?

Lazy Emptying [Cluster 2017]: Only empty the burst buffer when its load reaches a threshold T .

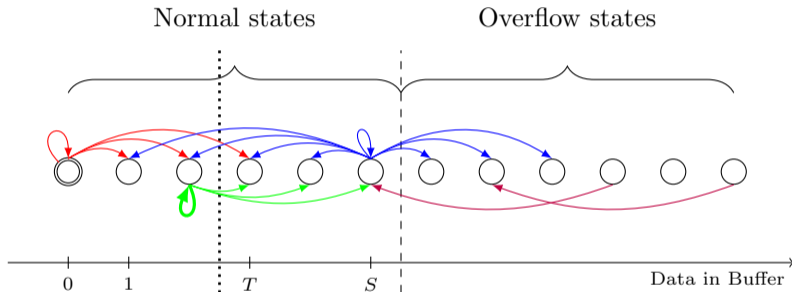


WHAT POLICY?

Lazy Emptying [Cluster 2017]: Only empty the burst buffer when its load reaches a threshold T .



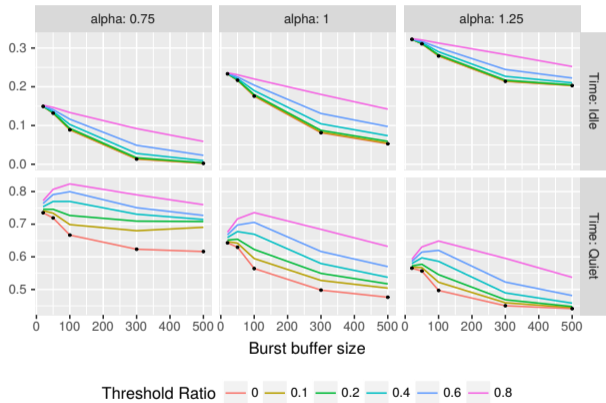
Lazy Emptying [Cluster 2017]: Only empty the burst buffer when its load reaches a threshold T .



- ▶ Still unique stationary distribution π^{LAZY}
- ▶ QUIET time (no data sent from buffer):

$$\sum_{s \in \text{normal state}} \pi_s^{\text{LAZY}} \cdot \sum_{t \geq s} P^{\text{LAZY}}(s, t)$$

WHAT POLICY?



Lazy Emptying: Only empty the burst buffer when its load reaches a threshold T (black dots are for $T = 0$, $p_{\max} = 0.1$).

COMPARISON WITH A DIFFERENT MODEL

Pseudo-periodic model: application \mathcal{A}_i

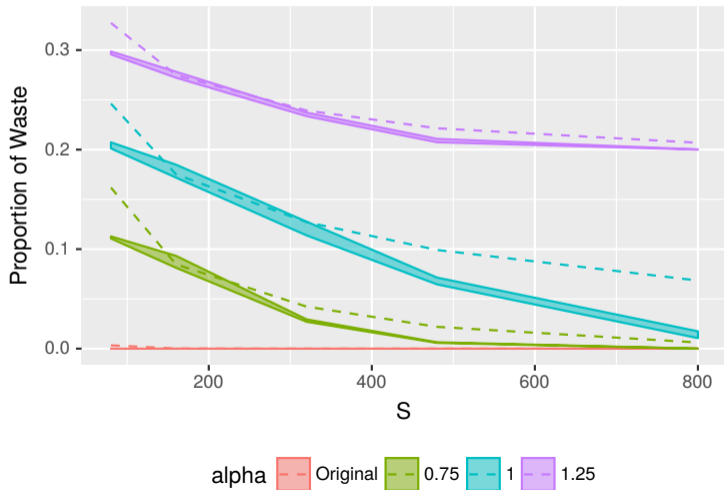
- ▶ has a **pseudo-period** d_i
- ▶ **sends data** at bandwidth b_i during $p_i d_i \pm 25\%$
- ▶ **computes** during $(1 - p_i) d_i \pm 25\%$

Evaluated on applications from APEX data set

[LANL Tech. Report, 2016]

Workflow	EAP	LAP	Silverton	VPIC
Number of Instances	13	4	2	1
b_i (GB/s)	160	80	160	160
d_i Period (s)	5671	12682	15005	4483
Checkpoint time (s)	20	25	280	23,4
$p_i (\times 10^{-3})$	3.51	1.97	18.7	5.11

COMPARISON WITH A DIFFERENT MODEL



- ▶ Increase stress α by scaling up p_i values
- ▶ Close behavior despite very different model

- ▶ Tractable model for dimensioning Burst Buffers
 - ▶ Which size for a given stress?
 - ▶ Emptying threshold: 20-40% is a reasonable choice
 - ▶ Validated against a different model
- ▶ Further questions
 - ▶ Other application models (maybe not Markovian)
 - ▶ Characterization of I/O patterns
 - ▶ Improve platform model (congestion, distributed BB)
- ▶ Open for criticism, remarks, suggestions, and collaborations!