

LA-UR-23-28509

Approved for public release; distribution is unlimited.

Title: SYSTEM AND JOB SCHEDULING SIMULATION FOR ENHANCING PRODUCTION HPC

Author(s): Hafener, Vivian Erica
Senator, Steven Terry
Jones, William M.
Walker, Craig S.
Debardeleben, Nathan A.

Intended for: Slurm User Group, 2023-09-11/2023-09-13 (Provo, Utah, United States)

Issued: 2023-07-26



SYSTEM AND JOB SCHEDULING SIMULATION FOR ENHANCING PRODUCTION HPC

Principal Developers: **Craig Walker** (CCU staff), **Vivian Hafener** (LANL staff, RIT student, presenter)

*Former Contributors: Nicklaus Przybylski (LANL post-bacc / NGP), Braeden Slade (LANL post-bacc / storage)
Gavin Bailey (Boeing / PhD student)*

LANL POCs: **Steve Senator** (HPC-ENV) and **Nathan DeBardleben** (HPC-DES)

CCU Faculty Mentor: **William M. Jones**

Work funded under LANL contract #C3132/SR489790212

Executed: 11-JAN-2023

September 12-13th, 2023

Slurm User Group

Brigham Young University



Bottomline Upfront – Takeaways

Bottomline Upfront – Takeaways

- Specific experiments with specific LANL workloads
 - Resilience and reliability
 - Node sharing / packing
 - DST scheduling



Bottomline Upfront – Takeaways

- Specific experiments with specific LANL workloads
 - Resilience and reliability
 - Node sharing / packing
 - DST scheduling
- Tooling / software
 - Quantify user- and system-centric metrics
 - Applied to systems of arbitrary size
 - Workloads of interest behind the fence



```
25 def check_db():
26     if not os.path.isfile(FILE_NAME):
27         db.create_all()
28
29
30 @app.route("/")
31 def home():
32     check_db()
33     all_books = db.session.query(Book).all()
34     return render_template("index.html", books=all_books)
35
36 @app.route("/edit", methods=['GET', 'POST'])
37 def edit():
38
39     if request.method == 'POST':
40         book_id = request.form['id']
41         book_to_update = Book.query.get(book_id)
42         book_to_update.rating = request.form['rating']
43         db.session.commit()
44         return redirect(url_for('home'))
```



Bottomline Upfront – Takeaways

- Specific experiments with specific LANL workloads
 - Resilience and reliability
 - Node sharing / packing
 - DST scheduling
- Tooling / software
 - Quantify user- and system-centric metrics
 - Applied to systems of arbitrary size
 - Workloads of interest behind the fence
- Interested in feedback on ways to improve / expand our work



```
25 def check_db():
26     if not os.path.isfile(FILE_PATH):
27         db.create_all()
28
29 @app.route("/")
30 def home():
31     check_db()
32     all_books = db.session.query(Book).all()
33     return render_template("index.html", books=all_books)
34
35 @app.route("/edit", methods=['GET', 'POST'])
36 def edit():
37     if request.method == 'POST':
38         book_id = request.form['id']
39         book_to_update = Book.query.get(book_id)
40         book_to_update.rating = request.form['rating']
41         db.session.commit()
42         return redirect(url_for('home'))
```



How did this all start?

What is BatSim?


How did we modify BatSim?

What are some questions that these tools can help answer?

Simulation-based framework to explore impact to performance of large-scale systems due to degraded reliability in DRAM systems



Simulation-based framework to explore impact to performance of large-scale systems due to degraded reliability in DRAM systems

(primarily looking at soft errors)



500 0 500 1500
(kilometers)
NAD27/*DNAG

Cosmic-ray Exposure (nGy/hr)



The Universe is Hostile to Computers by Veritasium

What is Batsim ?





realistic, Based on SimGrid, multi-processed, on-going active development


Pierre-François Dutot, Michael Mercier, Millian Poquet, Olivier Richard. *"Batsim: a Realistic Language-Independent Resources and Jobs Management Systems Simulator,"* **20th Workshop on Job Scheduling Strategies for Parallel Processing**, May 2016.

Real

RJMS





Batsim Architecture



real platform


Batsim simulation

decision maker (RJMS + adaptor)




batsim protocol

batsim


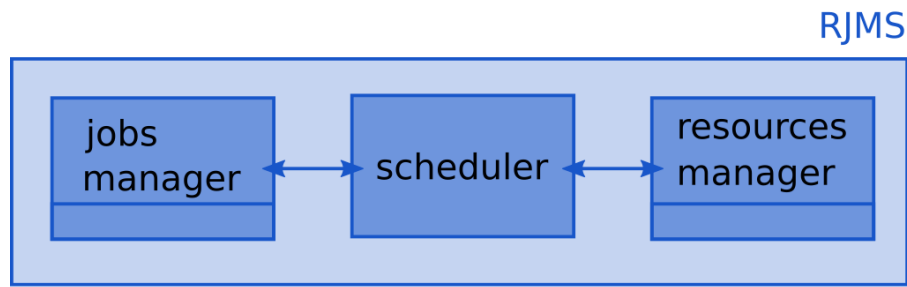


results



simulated platform


Real




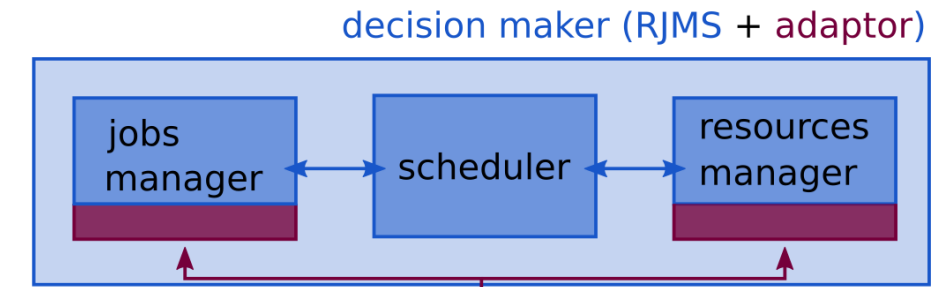
Batsim Architecture

Needed to add:

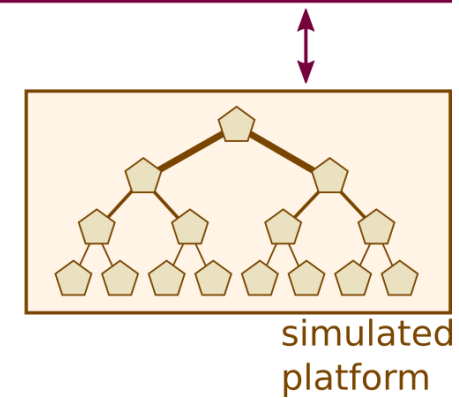


fault model
checkpoint-restart



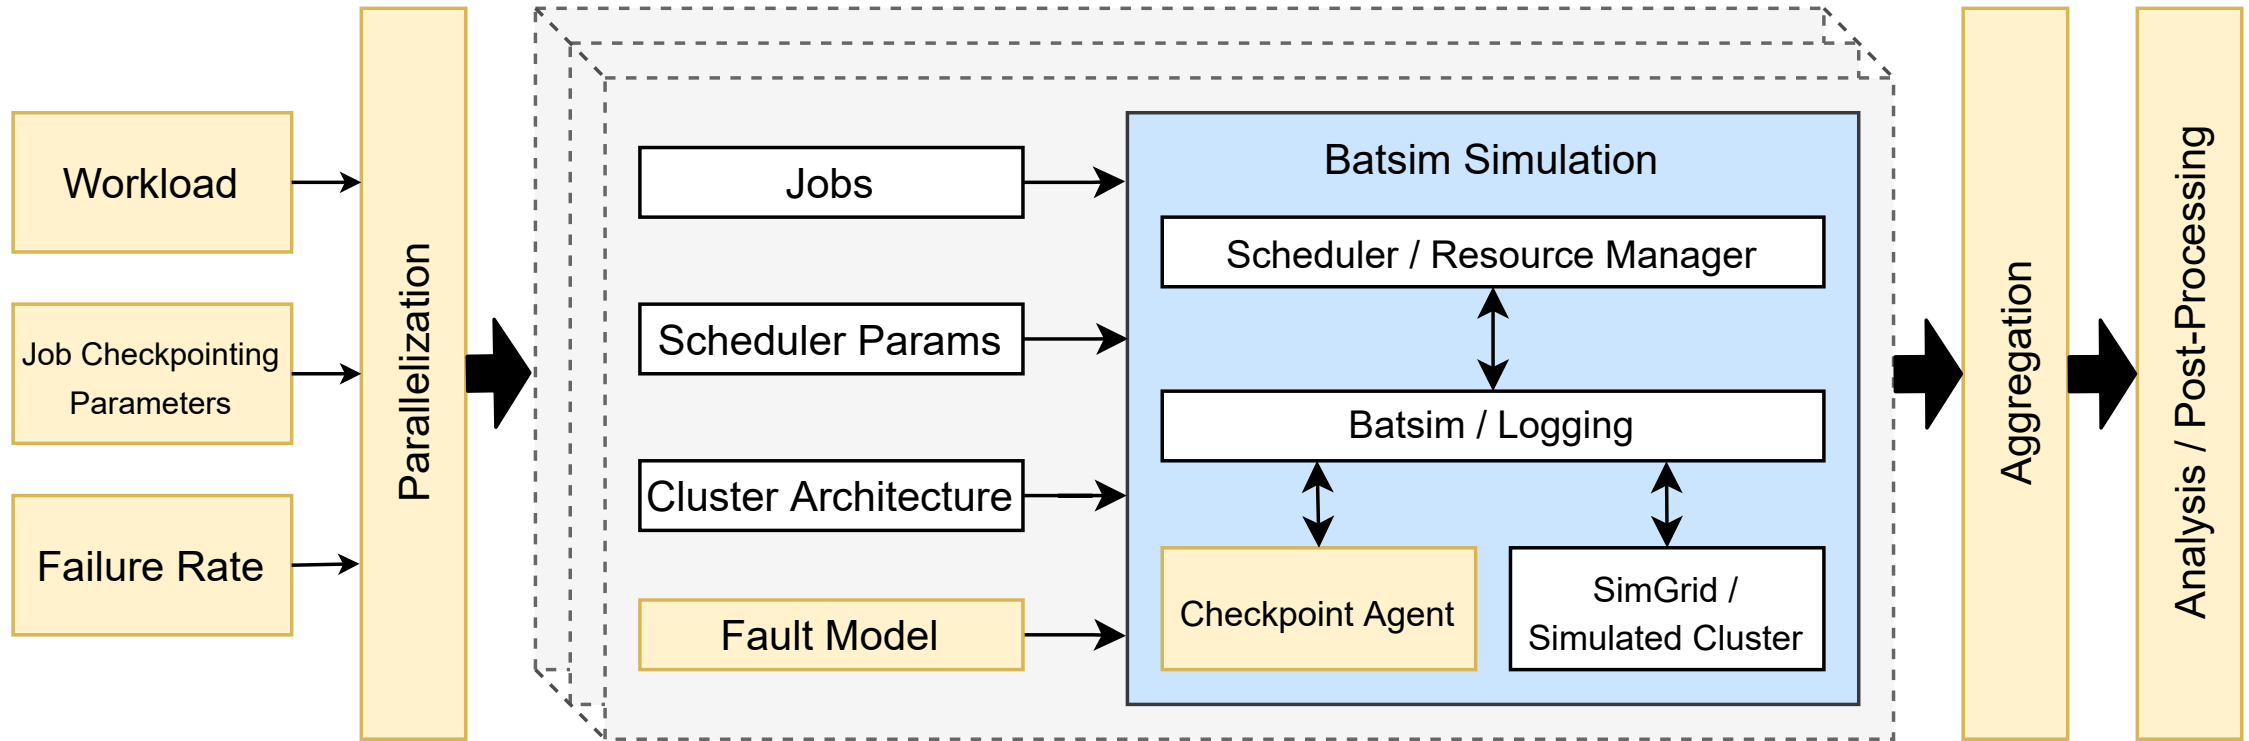
Batsim simulation



batsim protocol

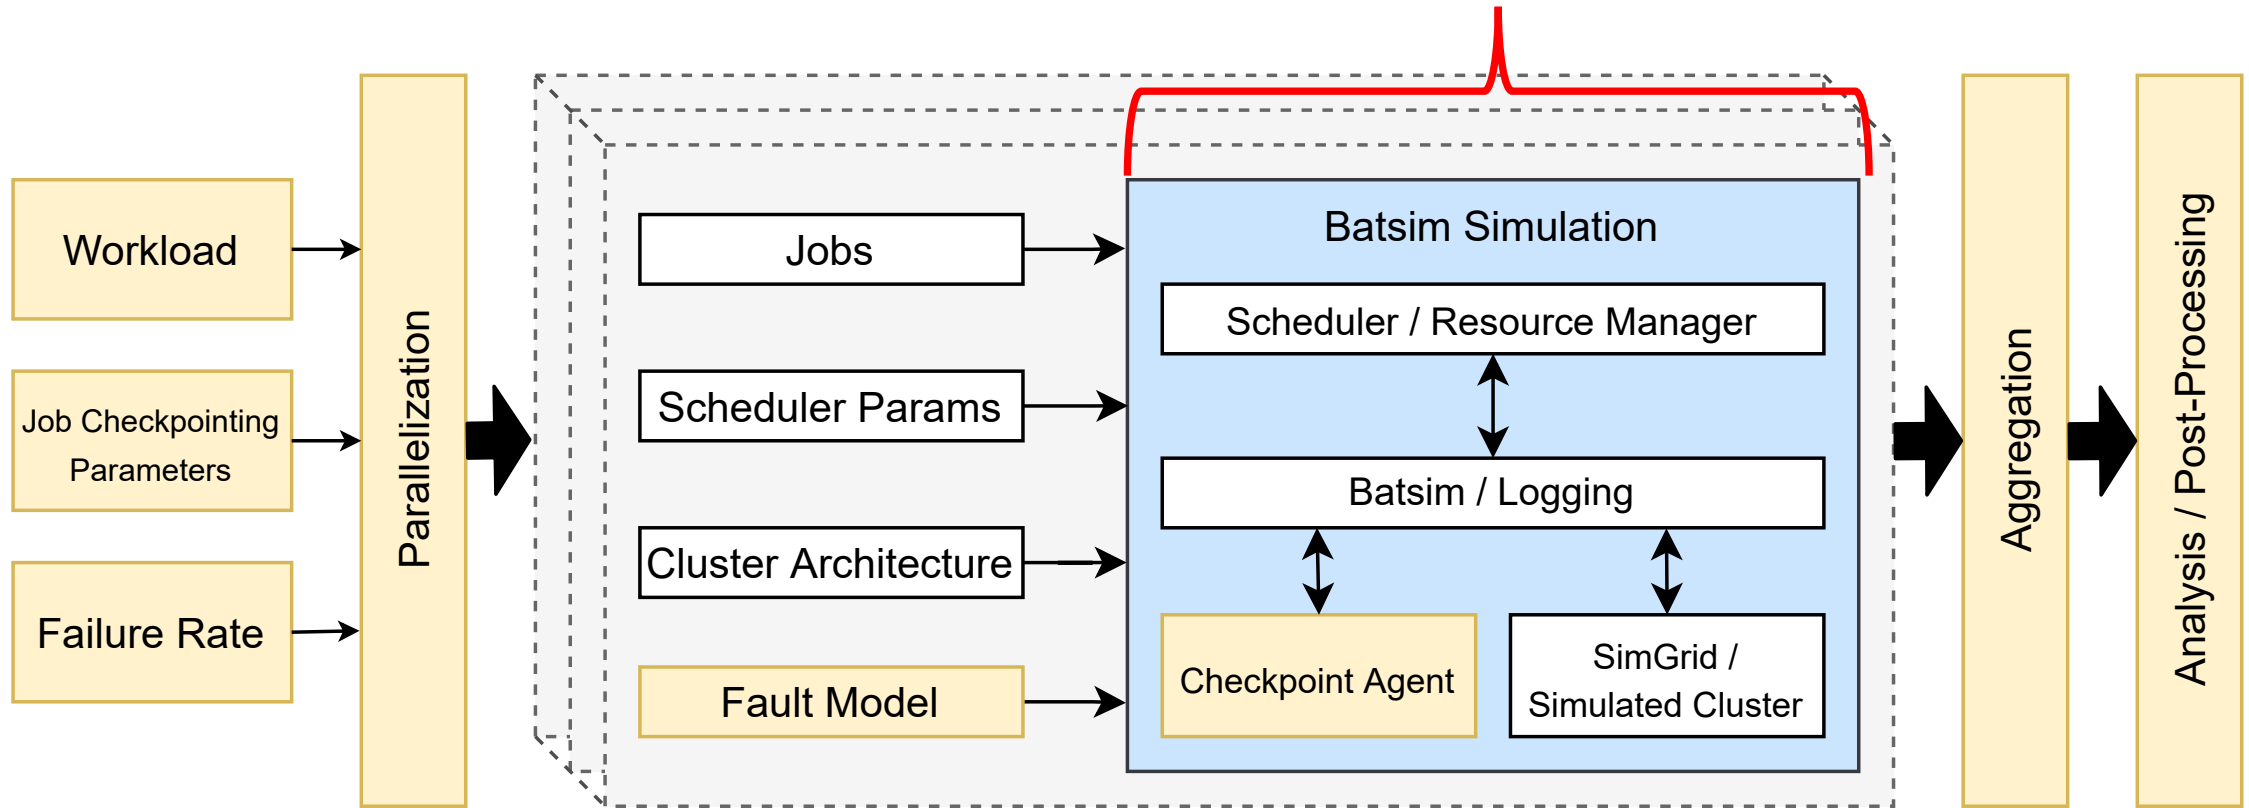


Modifications to Batsim



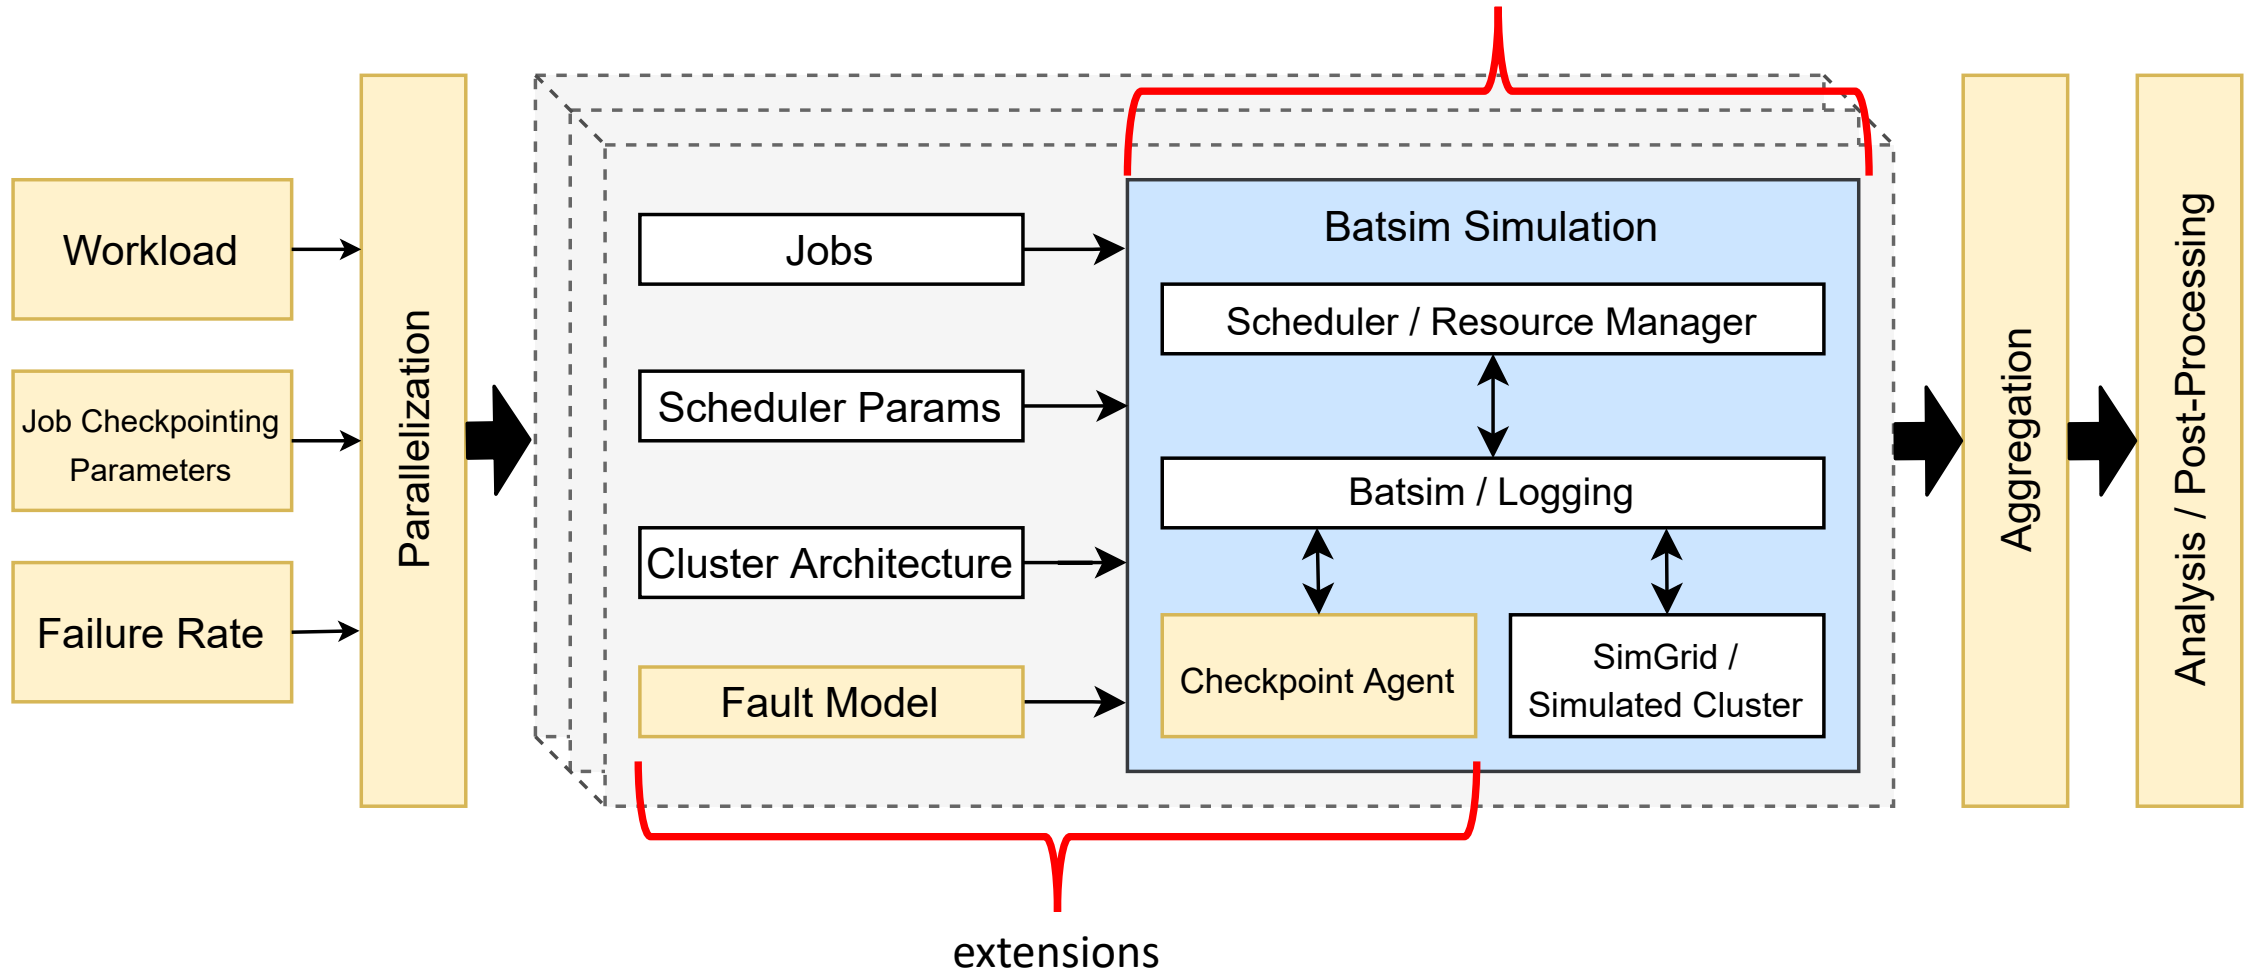
Modifications to Batsim

existing Batsim application



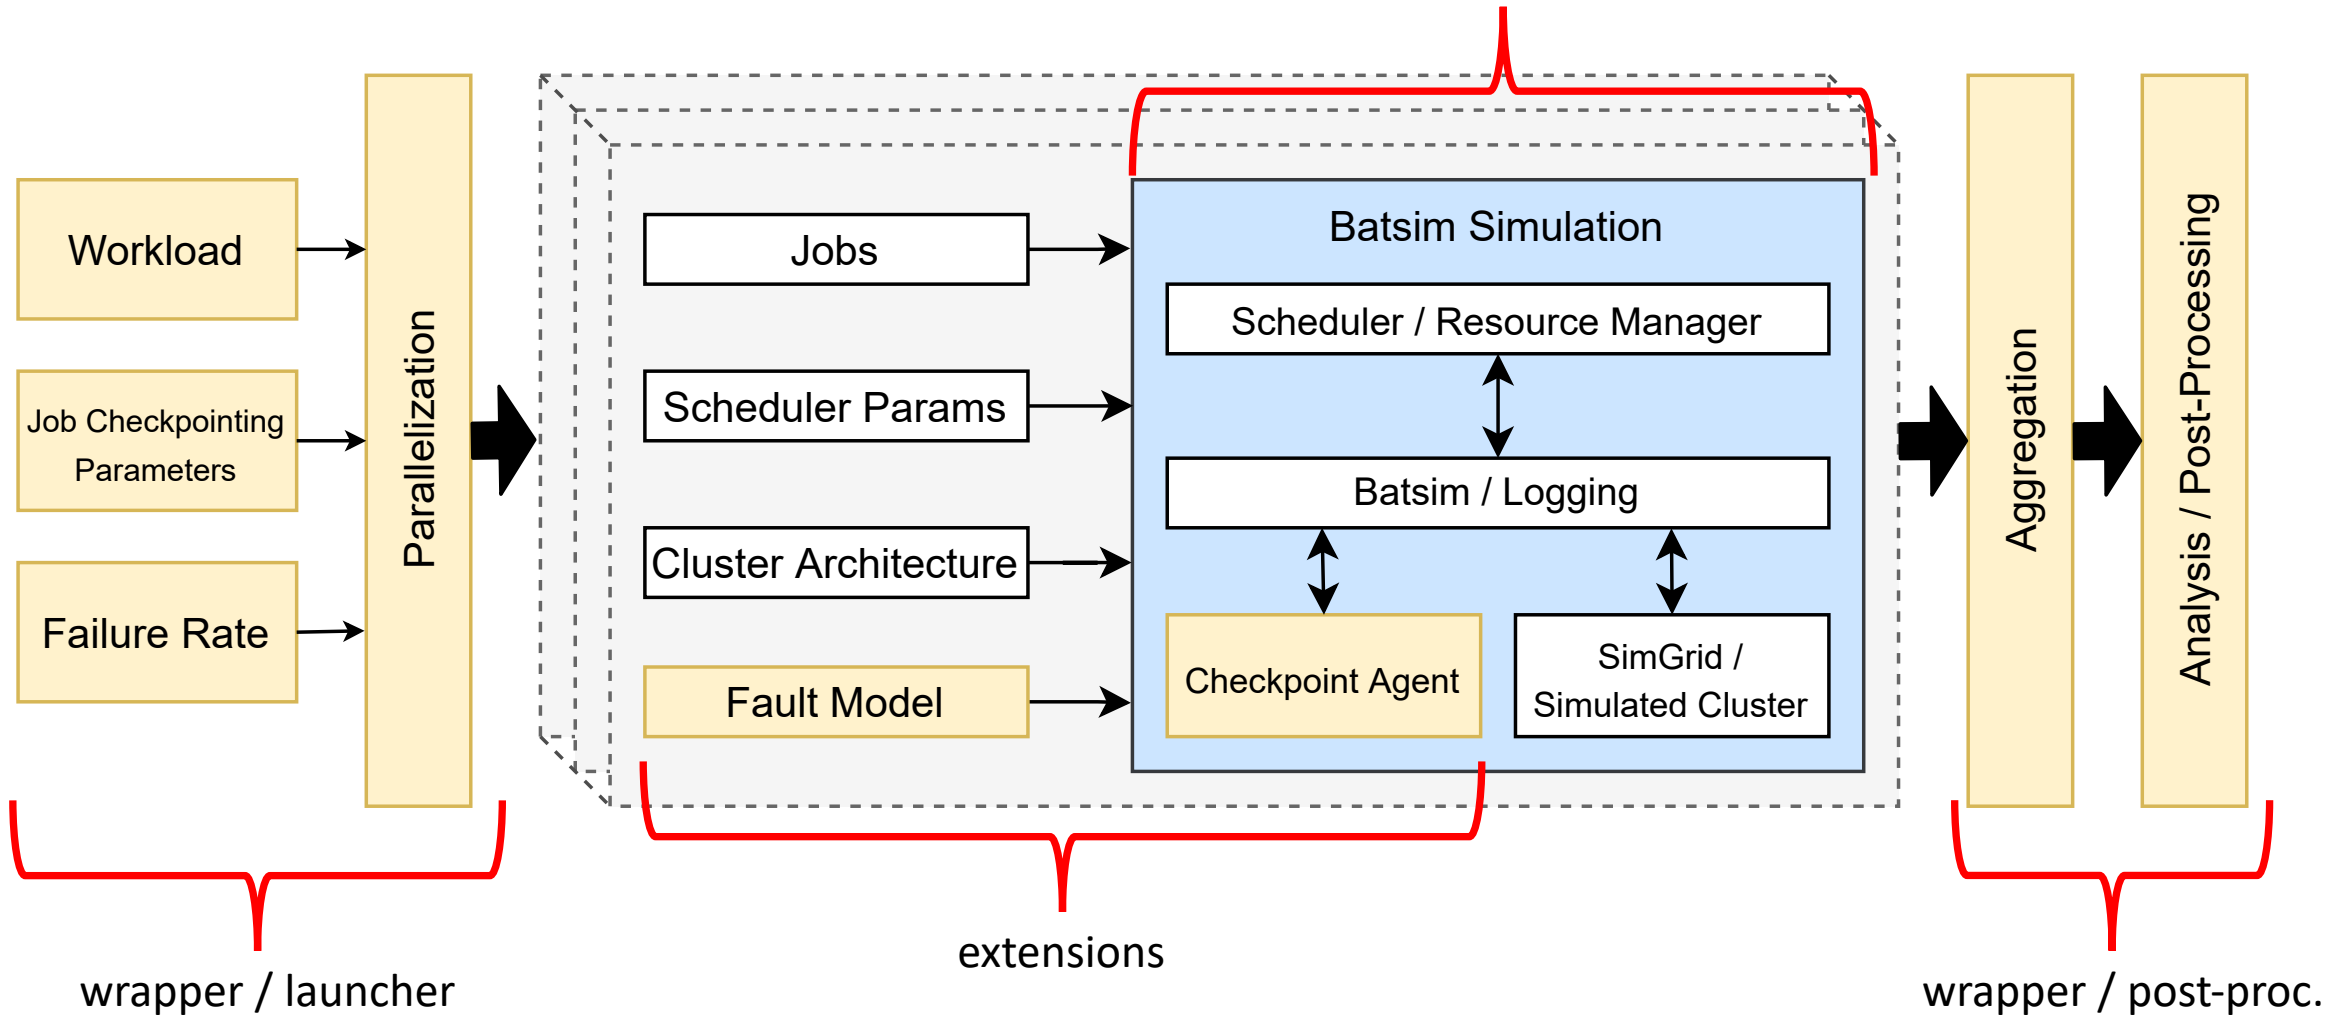
Modifications to Batsim

existing Batsim application



Modifications to Batsim

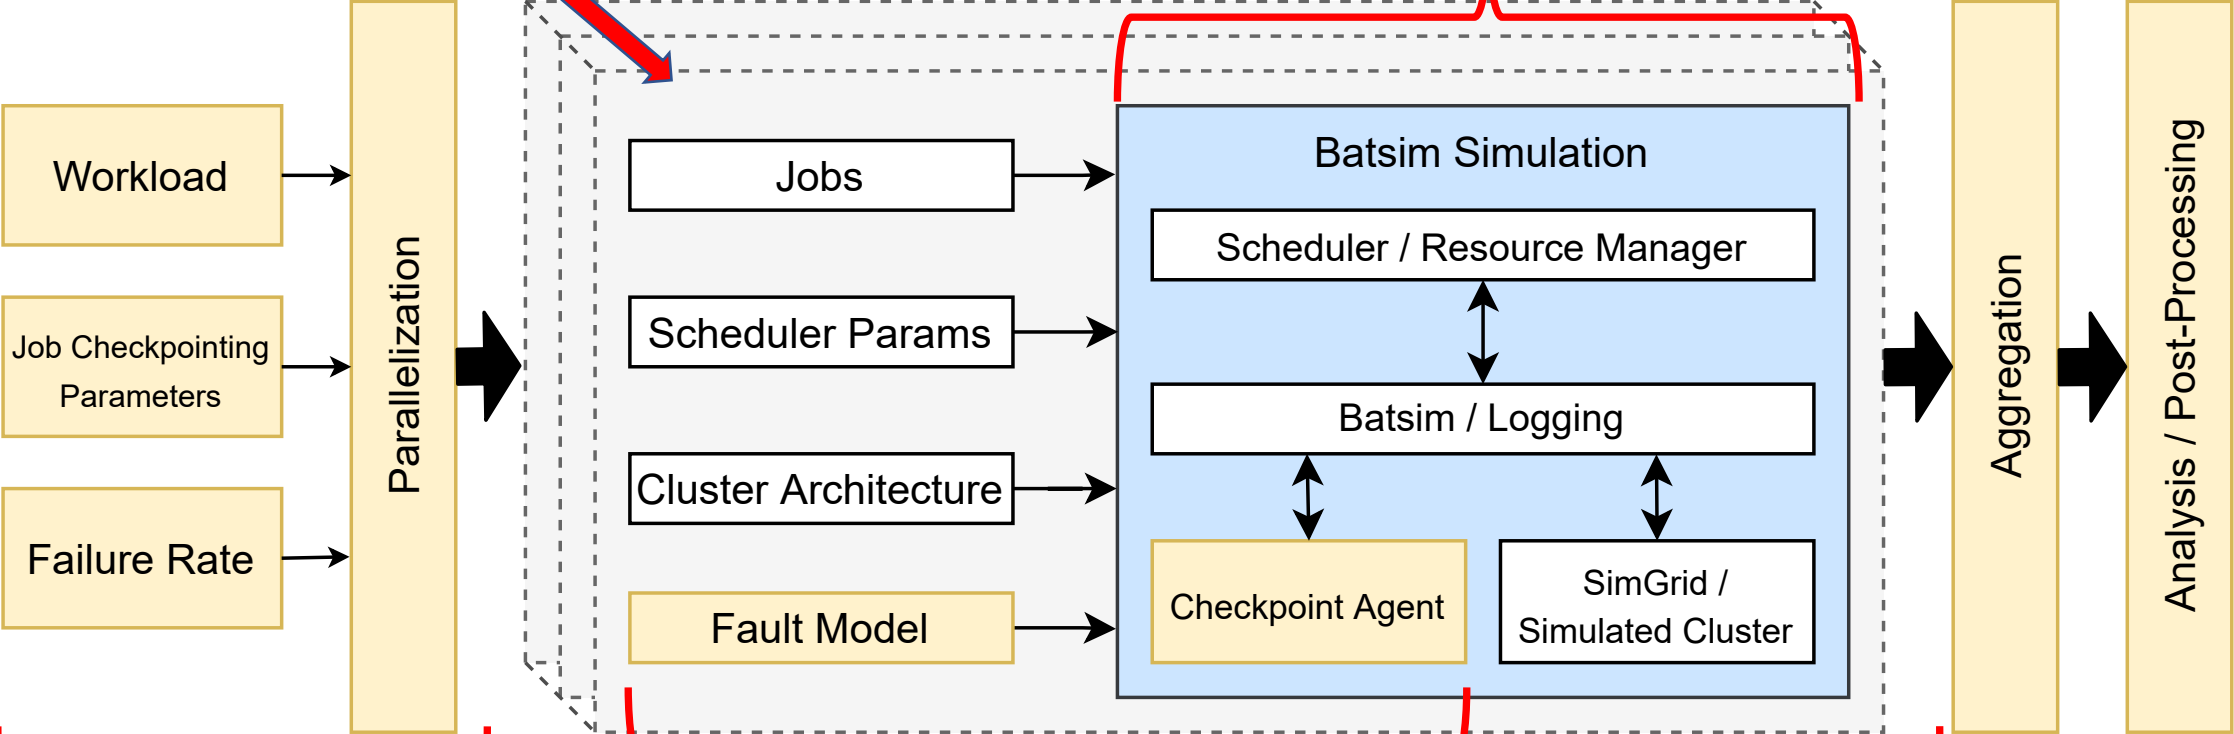
existing Batsim application



Modifications to Batsim

parallelize for large-scale
Monte Carlo
parameter sweeps

existing Batsim application



wrapper / launcher

extensions

wrapper / post-proc.

Experimentation

Experimentation

LANL Cluster Grizzly

Experimentation

LANL Cluster Grizzly

11 months of real job log data from 2018
180K jobs

Experimentation

LANL Cluster Grizzly

11 months of real job log data from 2018
180K jobs

DESCRIPTION OF INPUT WORKLOADS USED IN PERFORMANCE ANALYSIS

workload	description
WL1	All jobs are one node wide and 24 hours long. This represents the analog of WL6, where all jobs span the entire cluster.
WL2	Jobs are based on 11 months of data from a capacity cluster at with 1490 nodes. This workload contained very few large jobs, with roughly 48% of the jobs requiring one node, and less than one hour of wall time, with the remaining jobs distributed in duration and width.
WL3	Job widths are uniformly distributed from one to size of cluster, with durations also uniformly distributed from one to 24 hours.
WL4	Job widths are uniformly distributed from 32 to size of cluster, with durations from 6 to 24 hours.
WL5	Job widths are divided into two bins: 512 or the entire cluster, 1490, with durations ranging from one to 24 hours, according to the distributions of durations present in the original WL2 workload.
WL6	Jobs are all 1490 nodes wide with durations of 24 hours. This represents a “worst-case” scenario from a reliability point of view, as any node failure will result in a job failure that spans the entire cluster.

Experimentation

LANL Cluster Grizzly

11 months of real job log data from 2018
180K jobs
very few “large jobs”

explore a range of scenarios
capability to capacity

DESCRIPTION OF INPUT WORKLOADS USED IN PERFORMANCE ANALYSIS

workload	description
WL1	All jobs are one node wide and 24 hours long. This represents the analog of WL6, where all jobs span the entire cluster.
WL2	Jobs are based on 11 months of data from a capacity cluster at with 1490 nodes. This workload contained very few large jobs, with roughly 48% of the jobs requiring one node, and less than one hour of wall time, with the remaining jobs distributed in duration and width.
WL3	Job widths are uniformly distributed from one to size of cluster, with durations also uniformly distributed from one to 24 hours.
WL4	Job widths are uniformly distributed from 32 to size of cluster, with durations from 6 to 24 hours.
WL5	Job widths are divided into two bins: 512 or the entire cluster, 1490, with durations ranging from one to 24 hours, according to the distributions of durations present in the original WL2 workload.
WL6	Jobs are all 1490 nodes wide with durations of 24 hours. This represents a “worst-case” scenario from a reliability point of view, as any node failure will result in a job failure that spans the entire cluster.

Experimentation

LANL Cluster Grizzly

11 months of real job log data from 2018
180K jobs
very few “large jobs”

explore a range of scenarios
capability to capacity

WL1 – purely capacity

DESCRIPTION OF INPUT WORKLOADS USED IN PERFORMANCE ANALYSIS

workload	description
WL1	All jobs are one node wide and 24 hours long. This represents the analog of WL6, where all jobs span the entire cluster.
WL2	Jobs are based on 11 months of data from a capacity cluster at with 1490 nodes. This workload contained very few large jobs, with roughly 48% of the jobs requiring one node, and less than one hour of wall time, with the remaining jobs distributed in duration and width.
WL3	Job widths are uniformly distributed from one to size of cluster, with durations also uniformly distributed from one to 24 hours.
WL4	Job widths are uniformly distributed from 32 to size of cluster, with durations from 6 to 24 hours.
WL5	Job widths are divided into two bins: 512 or the entire cluster, 1490, with durations ranging from one to 24 hours, according to the distributions of durations present in the original WL2 workload.
WL6	Jobs are all 1490 nodes wide with durations of 24 hours. This represents a “worst-case” scenario from a reliability point of view, as any node failure will result in a job failure that spans the entire cluster.



Experimentation

LANL Cluster Grizzly

11 months of real job log data from 2018
180K jobs
very few “large jobs”

explore a range of scenarios
capability to capacity

WL1 – purely capacity
WL6 – purely capability

DESCRIPTION OF INPUT WORKLOADS USED IN PERFORMANCE ANALYSIS

workload	description
WL1	All jobs are one node wide and 24 hours long. This represents the analog of WL6, where all jobs span the entire cluster.
WL2	Jobs are based on 11 months of data from a capacity cluster at with 1490 nodes. This workload contained very few large jobs, with roughly 48% of the jobs requiring one node, and less than one hour of wall time, with the remaining jobs distributed in duration and width.
WL3	Job widths are uniformly distributed from one to size of cluster, with durations also uniformly distributed from one to 24 hours.
WL4	Job widths are uniformly distributed from 32 to size of cluster, with durations from 6 to 24 hours.
WL5	Job widths are divided into two bins: 512 or the entire cluster, 1490, with durations ranging from one to 24 hours, according to the distributions of durations present in the original WL2 workload.
WL6	Jobs are all 1490 nodes wide with durations of 24 hours. This represents a “worst-case” scenario from a reliability point of view, as any node failure will result in a job failure that spans the entire cluster.



Experimentation

LANL Cluster Grizzly

11 months of real job log data from 2018
180K jobs
very few “large jobs”

explore a range of scenarios
capability to capacity

WL1 – purely capacity
WL6 – purely capability
WL{3-5} -- mixture

DESCRIPTION OF INPUT WORKLOADS USED IN PERFORMANCE ANALYSIS

workload	description
WL1	All jobs are one node wide and 24 hours long. This represents the analog of WL6, where all jobs span the entire cluster.
WL2	Jobs are based on 11 months of data from a capacity cluster at with 1490 nodes. This workload contained very few large jobs, with roughly 48% of the jobs requiring one node, and less than one hour of wall time, with the remaining jobs distributed in duration and width.
WL3	Job widths are uniformly distributed from one to size of cluster, with durations also uniformly distributed from one to 24 hours.
WL4	Job widths are uniformly distributed from 32 to size of cluster, with durations from 6 to 24 hours.
WL5	Job widths are divided into two bins: 512 or the entire cluster, 1490, with durations ranging from one to 24 hours, according to the distributions of durations present in the original WL2 workload.
WL6	Jobs are all 1490 nodes wide with durations of 24 hours. This represents a “worst-case” scenario from a reliability point of view, as any node failure will result in a job failure that spans the entire cluster.



Tradeoff between reliability and performance


Tradeoff between reliability and performance

1

Baseline system
reliable DRAM

Tradeoff between reliability and performance


1




N nodes

Tradeoff between reliability and performance

1




N nodes




N nodes

Tradeoff between reliability and performance

1



N nodes



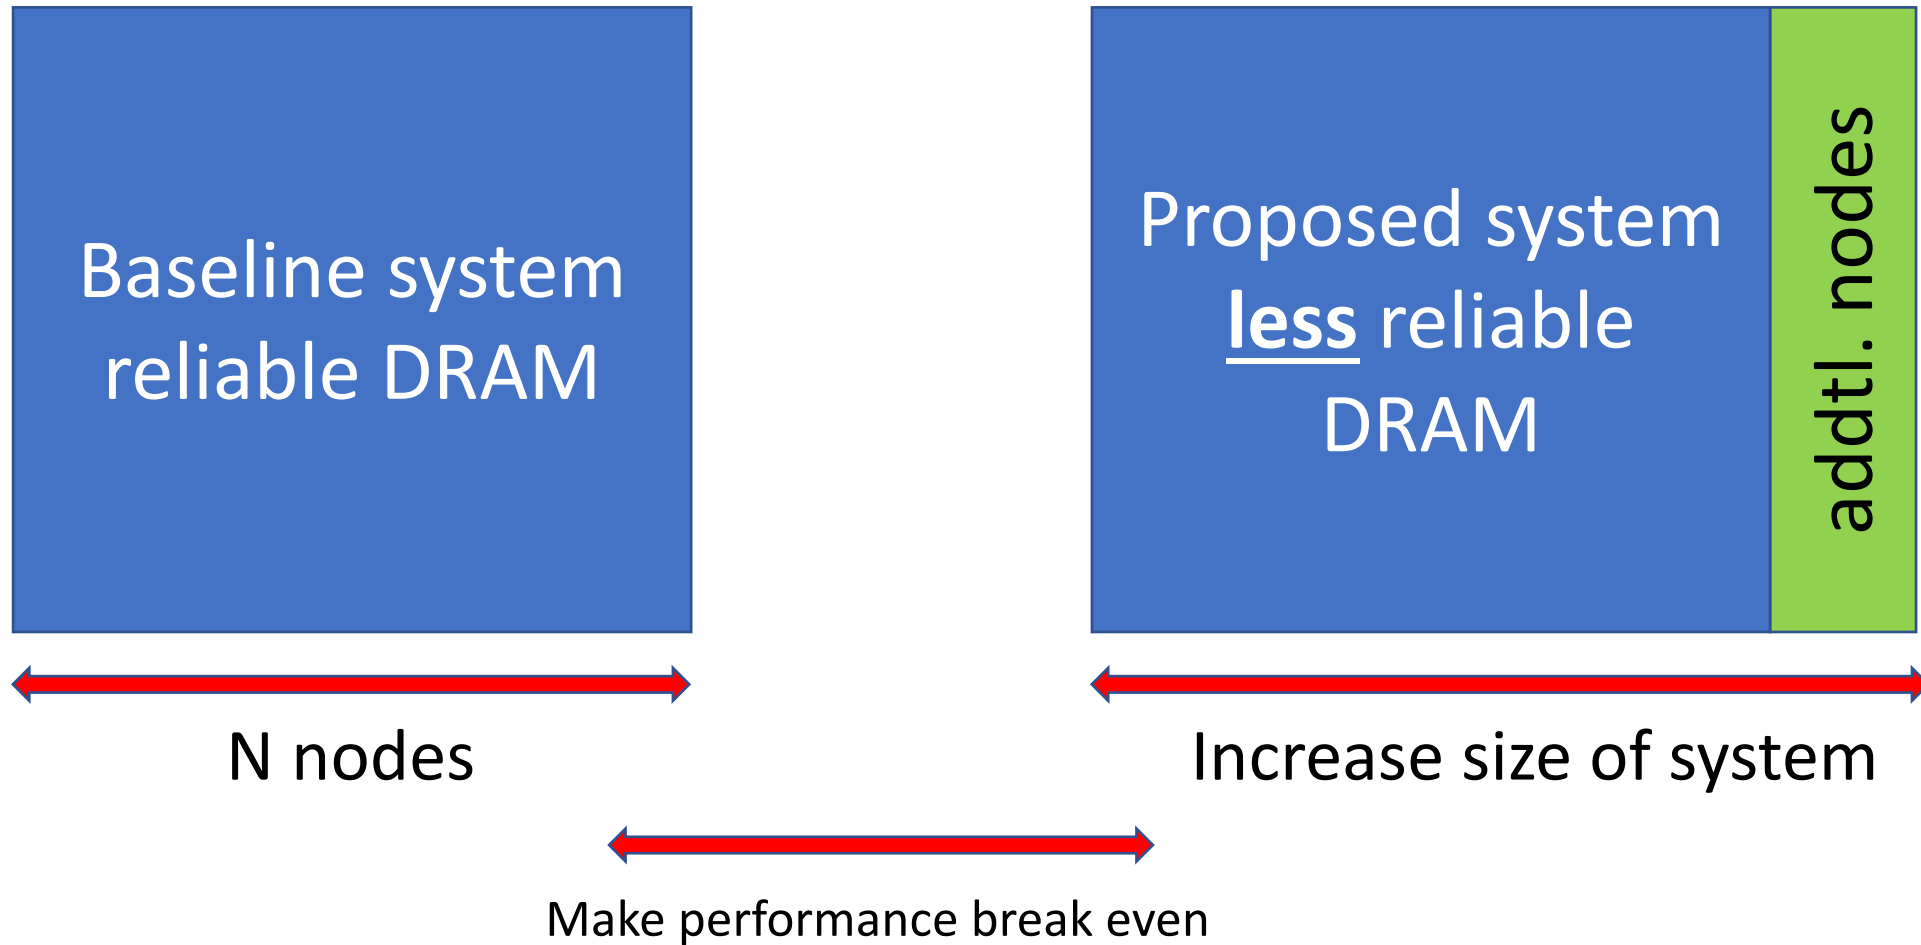
N nodes



Compare performance

Tradeoff between reliability and performance

2




First experimental campaign

- Grizzly WL2 (original) – 1490 nodes
- Assuming 13 day overall MTBF
 - **comparable to Trinity (20K nodes) 1 day MTBF**
- Assume proposed system w/ **32x less reliable DRAM**
- Assume entire workload arrives at $t = 0$
 - evaluate makespan
- Random, treat as Monte Carlo
 - → extremely large number of overall simulations conducted


Results

Change in Makespan as a Function of Cluster Size Relative to Baseline System (32x Less Reliable)



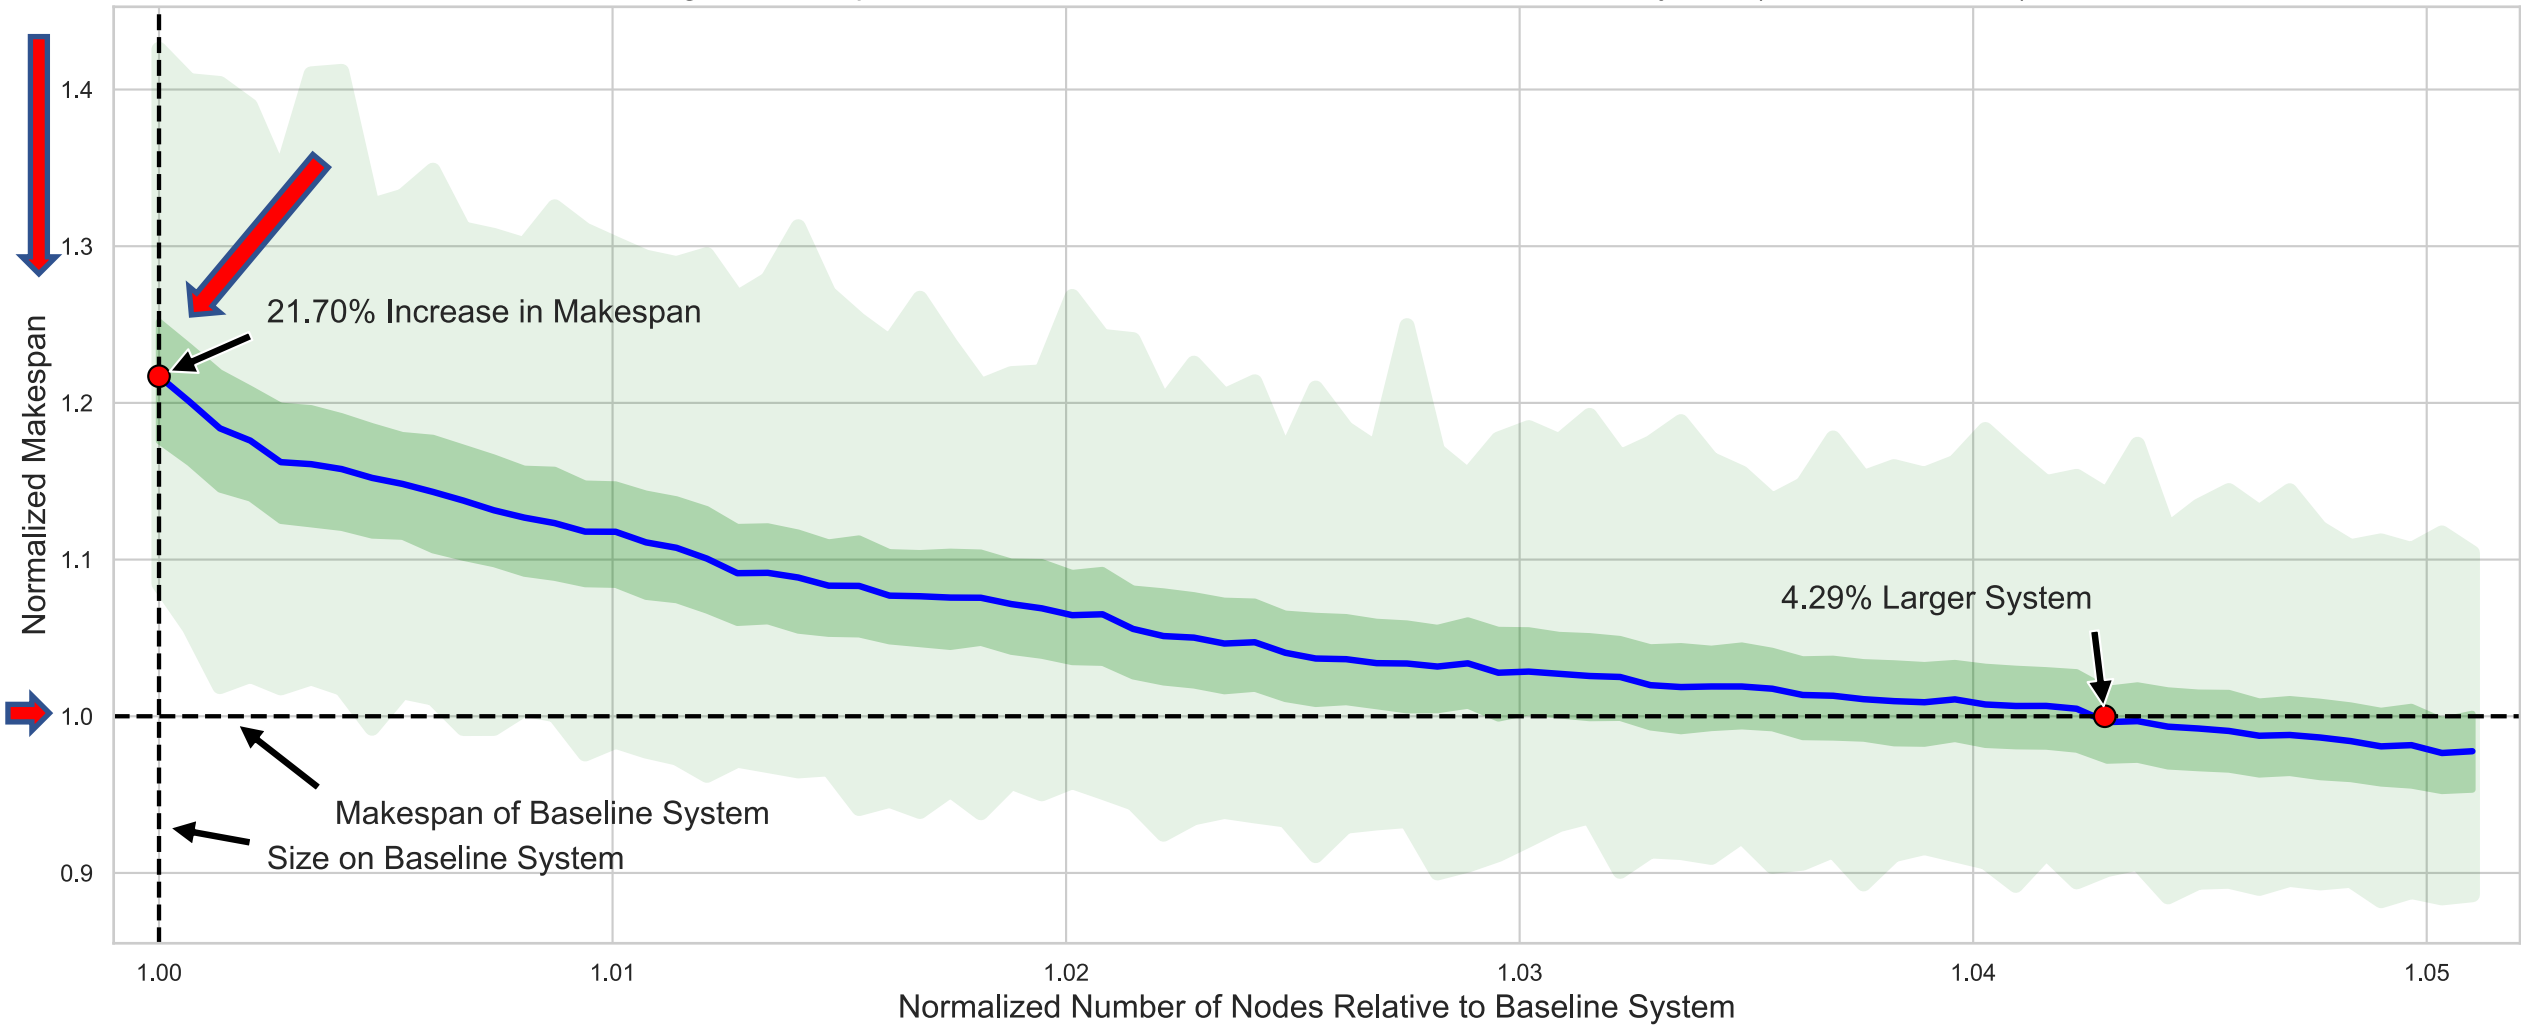
Results

Change in Makespan as a Function of Cluster Size Relative to Baseline System (32x Less Reliable)




Results

Change in Makespan as a Function of Cluster Size Relative to Baseline System (32x Less Reliable)




Results

Change in Makespan as a Function of Cluster Size Relative to Baseline System (32x Less Reliable)




Impact to Makespan for different workloads and reliability factors

Degraded System Compared to Baseline System - Change in Makespan



Impact to size (break-even point) for different workloads and reliability factors

Change in Degraded Cluster Size to Break Even With Original Baseline Performance




Experimentation


#nodes	1490
SMTTF	~1.6 days
#WLs	6
#jobs / WL	30K (<i>varied</i>)
#trials	400

Results (WL2 – Grizzly)


Jobs are based on 11 months of data from a capacity cluster at with 1490 nodes. This workload contained very few large jobs, with roughly 48% of the jobs requiring one node, and less than one hour of wall time, with the remaining jobs distributed in duration and width.



O(minutes)






O(hours)



O(days)




Results (WL2 – Grizzly)

Jobs are based on 11 months of data from a capacity cluster at with 1490 nodes. This workload contained very few large jobs, with roughly 48% of the jobs requiring one node, and less than one hour of wall time, with the remaining jobs distributed in duration and width.






Results (WL2 – Grizzly)

Jobs are based on 11 months of data from a capacity cluster at with 1490 nodes. This workload contained very few large jobs, with roughly 48% of the jobs requiring one node, and less than one hour of wall time, with the remaining jobs distributed in duration and width.




Results (WL2 – Grizzly)

Jobs are based on 11 months of data from a capacity cluster at with 1490 nodes. This workload contained very few large jobs, with roughly 48% of the jobs requiring one node, and less than one hour of wall time, with the remaining jobs distributed in duration and width.




Results (WL6)


Jobs are all 1490 nodes wide with durations of 24 hours. This represents a “worst-case” scenario from a reliability point of view, as any node failure will result in a job failure that spans the entire cluster.



O(minutes)




O(hours)




O(days)

Results (WL1)


All jobs are one node wide and 24 hours long. This represents the analog of WL6, where all jobs span the entire cluster.



O(minutes)




O(hours)



O(days)


Results (WL1)

All jobs are one node wide and 24 hours long. This represents the analog of WL6, where all jobs span the entire cluster.




O(minutes)

Negligible impact



O(hours)

Negligible impact



O(days)

Negligible impact

Thrust Areas

Thrust Areas

Analyze workloads to (automatically / real-time?) classify them as “production” or “development”

What are the salient metrics?

Thrust Areas

Analyze workloads to (automatically / real-time?) classify them as “production” or “development”

What are the salient metrics?

Incorporate reservations into BatSim.

Backfilling ahead of reservations, what are impacts to performance?

Thrust Areas

Analyze workloads to (automatically / real-time?) classify them as “production” or “development”

What are the salient metrics?

Incorporate reservations into BatSim.

Backfilling ahead of reservations, what are impacts to performance?

Incorporate “partitions” into BatSim

What are the optimal relative sizes of the ‘standard’ and ‘debug’ partitions?

Thrust Areas

Analyze workloads to (automatically / real-time?) classify them as “production” or “development”

What are the salient metrics?

Incorporate reservations into BatSim.

Backfilling ahead of reservations, what are impacts to performance?

Incorporate “partitions” into BatSim

What are the optimal relative sizes of the ‘standard’ and ‘debug’ partitions?

Incorporate ‘node sharing’ into BatSim

How much share-packing is required to accommodate fewer nodes of greater resources?

Thrust Areas

Analyze workloads to (automatically / real-time?) classify them as “production” or “development”

What are the salient metrics?

Incorporate reservations into BatSim.

Backfilling ahead of reservations, what are impacts to performance?

Incorporate “partitions” into BatSim

What are the optimal relative sizes of the ‘standard’ and ‘debug’ partitions?

Incorporate ‘node sharing’ into BatSim

How much share-packing is required to accommodate fewer nodes of greater resources?

Production-driven goals

SLUG / BYU September 12-13, 2023

Thrust Areas

Analyze workloads to (automatically / real-time?) classify them as “production” or “development”

What are the salient metrics?

Incorporate reservations into BatSim.

Backfilling ahead of reservations, what are impacts to performance?

Incorporate “partitions” into BatSim

What are the optimal relative sizes of the ‘standard’ and ‘debug’ partitions?

Incorporate ‘node sharing’ into BatSim

How much share-packing is required to accommodate fewer nodes of greater resources?

stats / basic ML

Thrust Areas

Analyze workloads to (automatically / real-time?) classify them as “production” or “development”

What are the salient metrics?

Incorporate reservations into BatSim.

Backfilling ahead of reservations, what are impacts to performance?

Incorporate “partitions” into BatSim

What are the optimal relative sizes of the ‘standard’ and ‘debug’ partitions?

Incorporate ‘node sharing’ into BatSim

How much share-packing is required to accommodate fewer nodes of greater resources?

stats / basic ML

modifications to BatSim

Next Step

**Incorporate
'node sharing'
into BatSim**

How much share-
packing is
required to
accommodate
fewer nodes of
greater
resources?

Next Step

View node as collection of cores / PEs

**Incorporate
'node sharing'
into BatSim**

How much share-
packing is
required to
accommodate
fewer nodes of
greater
resources?

Next Step

View node as collection of cores / PEs

Prior work: job were just
“rectangles”, i.e. didn't use CPU
model

**Incorporate
'node sharing'
into BatSim**

How much share-
packing is
required to
accommodate
fewer nodes of
greater
resources?

Next Step

View node as collection of cores / PEs

Prior work: job were just “rectangles”, i.e. didn’t use CPU model

Model to address impact to baseline runtimes with concurrent jobs on same node

Incorporate ‘node sharing’ into BatSim

How much share-packing is required to accommodate fewer nodes of greater resources?

(simplifying) Assumptions

Assume all 1-node jobs in the Grizzly workload represent ‘serial’ jobs that these jobs only really need 1 “core”

(simplifying) Assumptions

Assume all 1-node jobs in the Grizzly workload represent ‘serial’ jobs that these jobs only really need 1 “core”

All others are true, parallel (multi-node) jobs fully use a node (or nodes)


(simplifying) Assumptions

Assume all 1-node jobs in the Grizzly workload represent ‘serial’ jobs that these jobs only really need 1 “core”

All others are true, parallel (multi-node) jobs fully use a node (or nodes)


Assume that 1-core jobs can ‘fill’ up to X% of cores (configurable) of a node, without “interaction” – not addressing “interaction” at all at this point

WORKLOAD 2 4-core 30,000 jobs



EASY BF (later conservative)



WI 4 core fcfs



FCFS

WORKLOAD 2 4-core 30,000 jobs

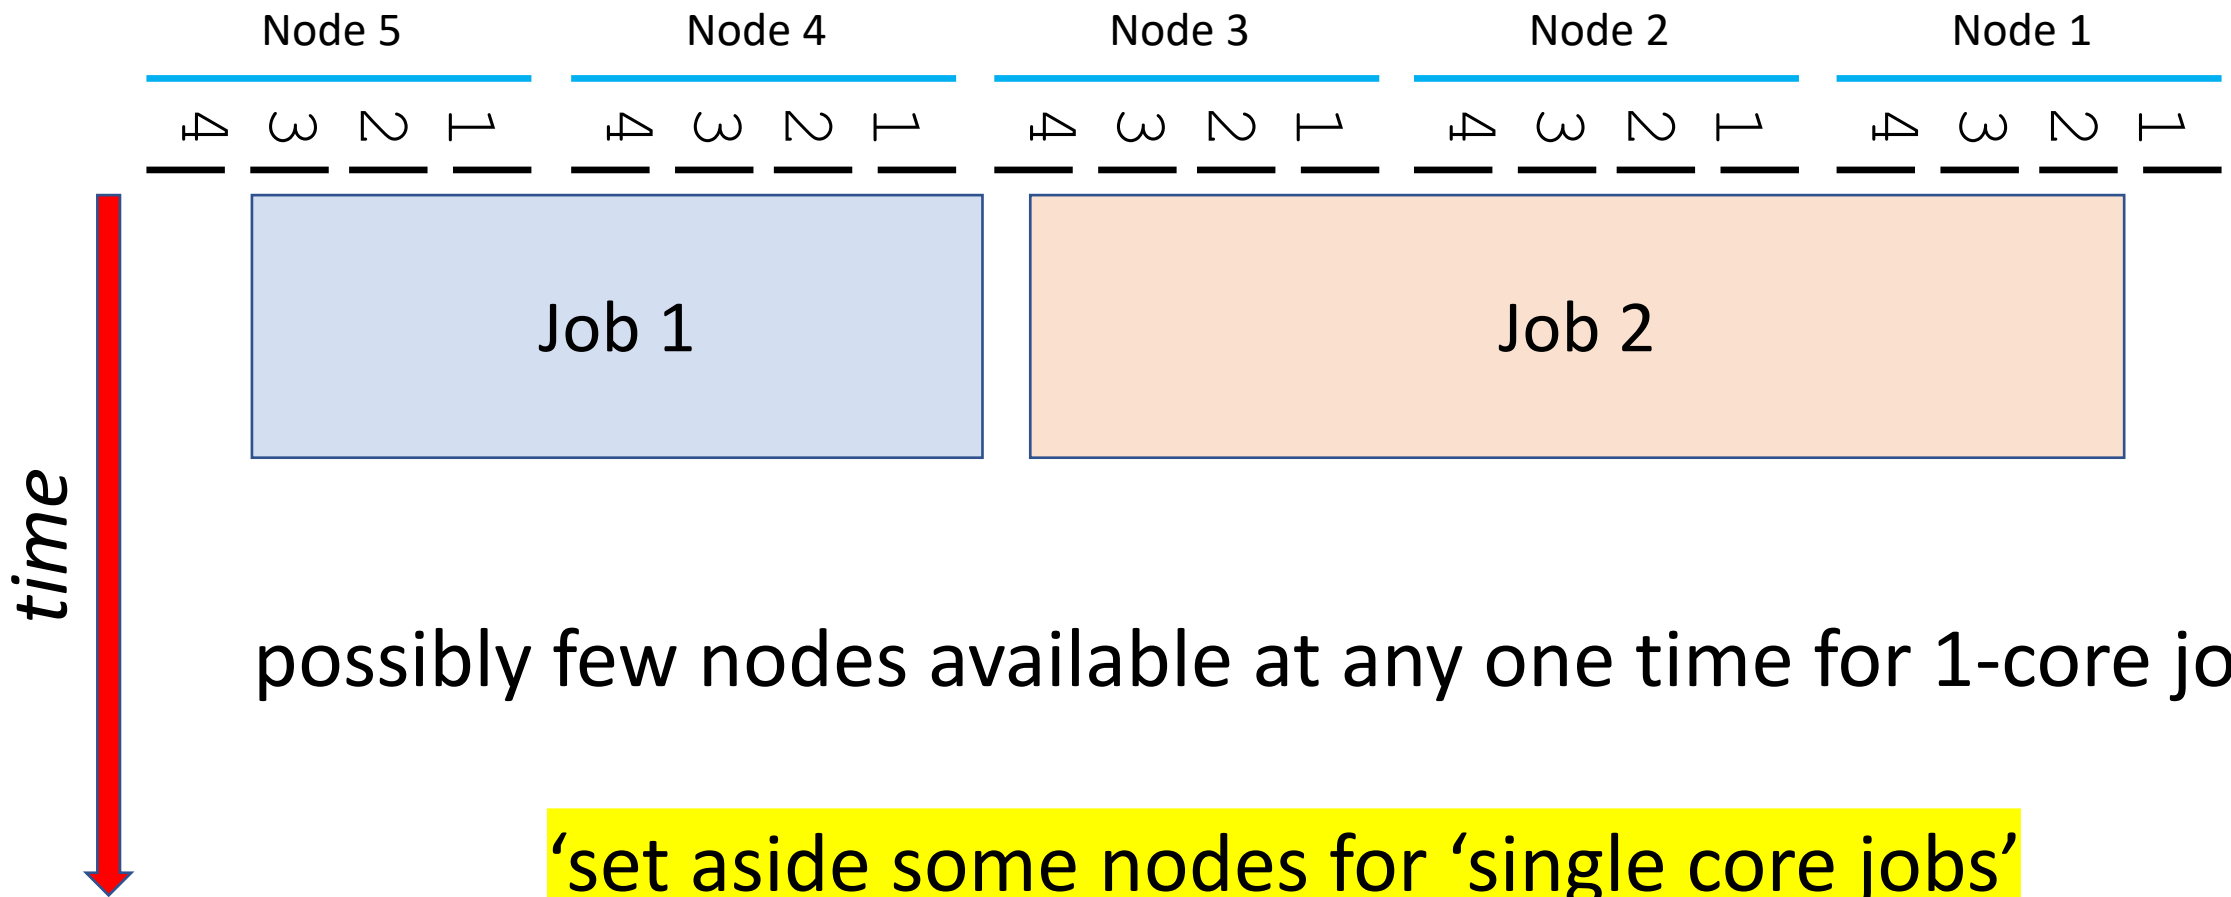
WI2 4 core fcfs



EASY BF (later conservative)



FCFS



possibly few nodes available at any one time for 1-core jobs

**'set aside some nodes for 'single core jobs'
partitions**

congruent with debug scenarios





The next slides:


The next slides have 4 graphs differing in the amount of cores.

They are all for 300:exp interarrival time


They are all for the bin $(0,1]$ ie 1 resource jobs

Assumed that Grizzly was 1500 nodes, with up to 10 “held back” exclusively for single core (packable) jobs.





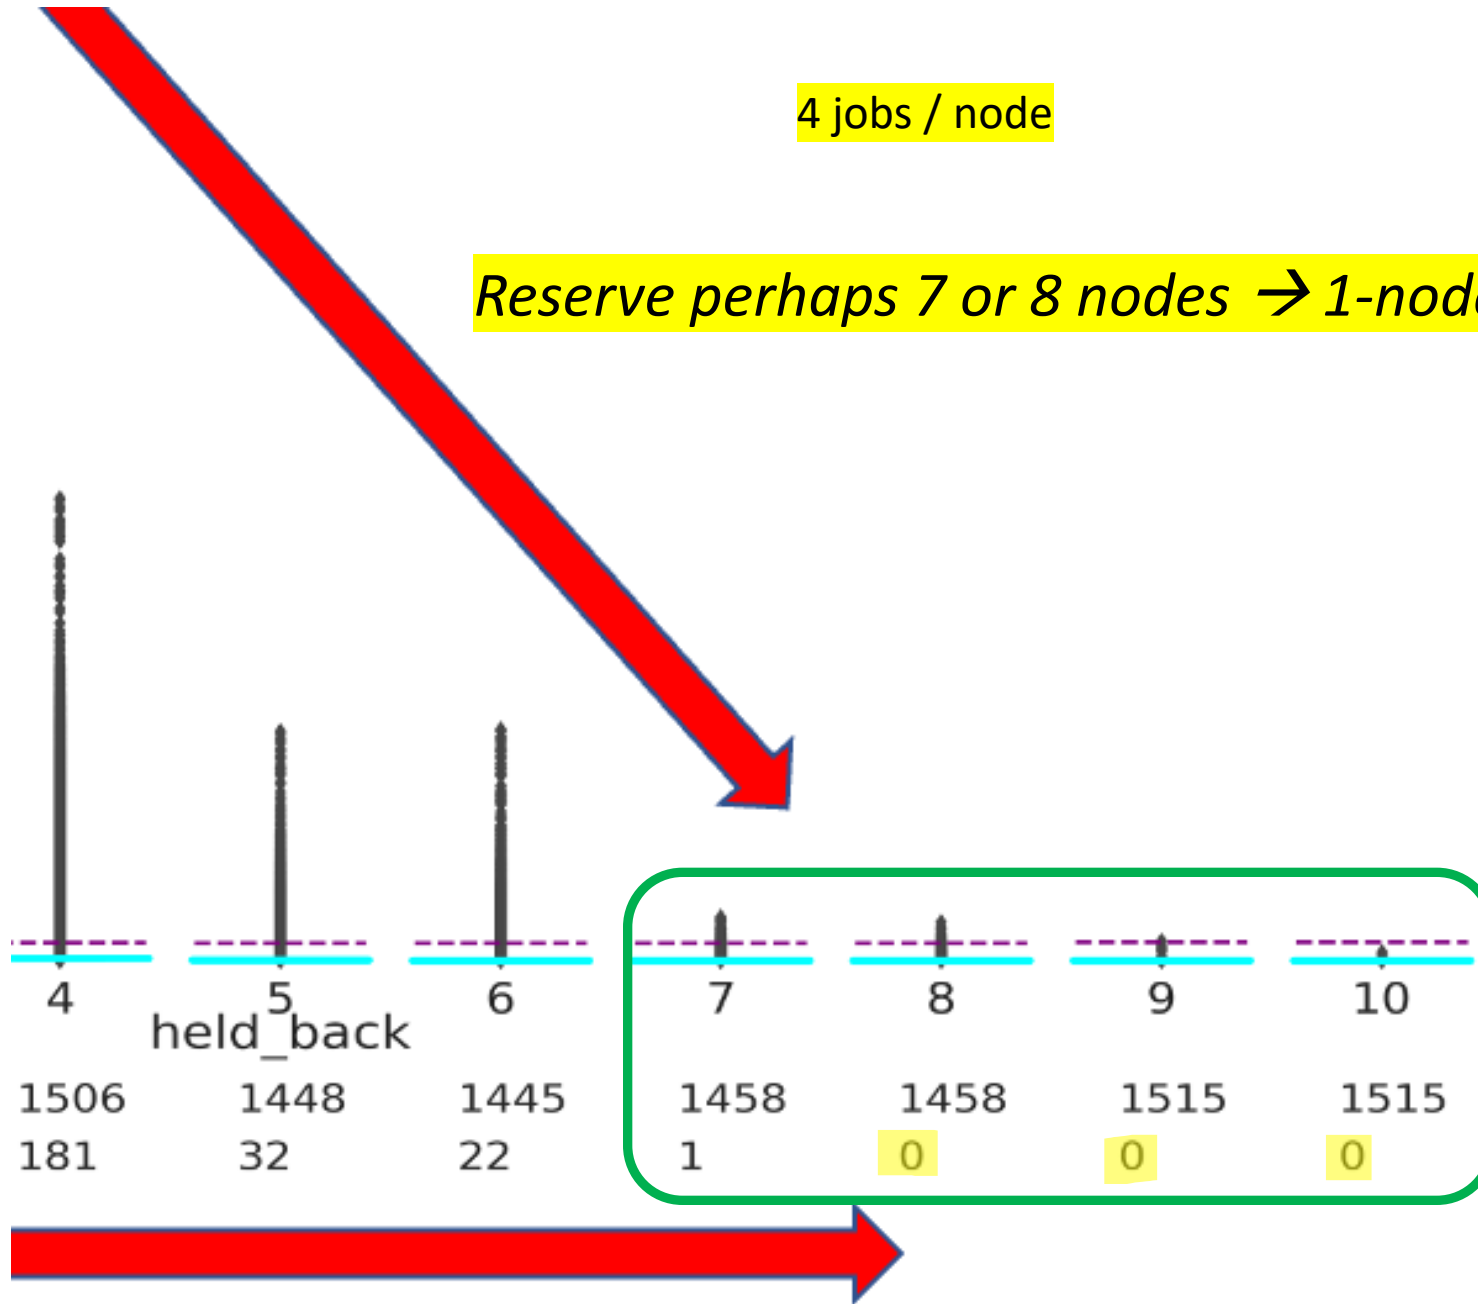
4 jobs / node



SLUG / BYU September 12-13, 2023

4 jobs / node

Reserve perhaps 7 or 8 nodes → 1-node jobs never wait



Thrust Areas

Analyze workloads to (automatically / real-time?) classify them as “production” or “development”

What are the salient metrics?

Incorporate reservations into BatSim.

Backfilling ahead of reservations, what are impacts to performance?

Incorporate “partitions” into BatSim

What are the optimal relative sizes of the ‘standard’ and ‘debug’ partitions?

Incorporate ‘node sharing’ into BatSim

How much share-packing is required to accommodate fewer nodes of greater resources?

stats / basic ML

modifications to BatSim

Incorporate Reservations into BatSim

- Different scenarios
 - schedule reservations “in the future”
 - before jobs have been scheduled – jobs ‘flow’ around these reservations
 - conflict between scheduled jobs and future reservation
 - conflict between running jobs and future (current) reservation
 - schedule reservations
- Policies
 - when a conflict happens:
 - only reschedule impacted jobs
 - reschedule all jobs so that the implicit priorities are respected
 - impacted jobs – progress they make deducted from requested wall times?
- Maintain compatibility with existing improvements
 - checkpointing / faults / jobs restarting




Some initial investigation

- Compare and contrast periodic reservations
 - Span the full cluster
 - Staggered reservations


“Full”

Nodes



“Staggered”

Nodes




Time

SLUG / BYU September 12-13, 2023


“Full”

Nodes



“Staggered”

Nodes




Time

SLUG / BYU September 12-13, 2023

“Full”



Nodes




“1 subdivision”

“Staggered”



Nodes




“2 subdivisions”

Time

SLUG / BYU September 12-13, 2023

“Staggered”


Nodes



“2 subdivisions”

“Staggered”

Nodes




“4 subdivisions”

Time

SLUG / BYU September 12-13, 2023

“Staggered”


Nodes



“2 subdivisions”

“Staggered”

Nodes




“4 subdivisions”

Time

SLUG / BYU September 12-13, 2023

“Staggered”


Nodes



“2 subdivisions”

“Staggered”

Nodes




“4 subdivisions”

Time

SLUG / BYU September 12-13, 2023

“Staggered”


Nodes



“2 subdivisions”

“Staggered”

Nodes




“4 subdivisions”

Time

SLUG / BYU September 12-13, 2023

“Staggered”

Nodes




“4 subdivisions”

“Staggered”



Nodes



“4 subdivisions”

“Staggered”



Nodes



“4 subdivisions”

2 days

Parameter Sweeps

Parameter Sweeps

Reservation durations: 1, 15, 60, 240, and 480 minutes (5 cases)

Parameter Sweeps

Reservation durations: 1, 15, 60, 240, and 480 minutes (5 cases)
Number of subdivisions: 1, 2, 4, 8 (4 cases)

Parameter Sweeps

Reservation durations:	1, 15, 60, 240, and 480 minutes (5 cases)
Number of subdivisions:	1, 2, 4, 8 (4 cases)
Interval:	1 month (across entire res. cycle), {2, 4, 8} days (4 cases)

Parameter Sweeps

Reservation durations: 1, 15, 60, 240, and 480 minutes (5 cases)
Number of subdivisions: 1, 2, 4, 8 (4 cases)
Interval: 1 month (across entire res. cycle), {2, 4, 8} days (4 cases)

Parameter Sweeps

Reservation durations: 1, 15, 60, 240, and 480 minutes (5 cases)
Number of subdivisions: 1, 2, 4, 8 (4 cases)
Interval: 1 month (across entire res. cycle), {2, 4, 8} days (4 cases)

80 total combinations

Parameter Sweeps

Reservation durations: 1, 15, 60, 240, and 480 minutes (5 cases)
Number of subdivisions: 1, 2, 4, 8 (4 cases)
Interval: 1 month (across entire res. cycle), {2, 4, 8} days (4 cases)

80 total combinations

Grizzly 2018 logs

Grizzly 2022 logs

Parameter Sweeps

Reservation durations: 1, 15, 60, 240, and 480 minutes (5 cases)
Number of subdivisions: 1, 2, 4, 8 (4 cases)
Interval: 1 month (across entire res. cycle), {2, 4, 8} days (4 cases)

80 total combinations

Grizzly 2018 logs → 25 Monte Carlo runs per experiment
Grizzly 2022 logs

Parameter Sweeps

Reservation durations: 1, 15, 60, 240, and 480 minutes (5 cases)
Number of subdivisions: 1, 2, 4, 8 (4 cases)
Interval: 1 month (across entire res. cycle), {2, 4, 8} days (4 cases)

80 total combinations


Grizzly 2018 logs → 25 Monte Carlo runs per experiment

Grizzly 2022 logs


Very resource intensive – 50-100GB per experiment,
3-8 days to complete.

As compared to <20GB and 5 minutes for earlier experiments.


Time: 00:480:00, 2018 full year




Time: 00:480:00, 2018 full year




Time: 00:480:00, 2018 full year




Time: 00:480:00, 2018 full year




Time: 00:480:00, 2018 full year




Time: 00:480:00, 2018 full year




Time: 00:480:00, 2018 full year




Time: 00:480:00, 2018 full year




Time: 00:480:00, 2018 full year



Time: 00:480:00, 2018 full year




Drill down on impact to jobs in the 4-day interval simulations with 8 subdivisions with reservations of length 480 minutes



	n=1,712,900	n=143,450	n=696,550	n=241,225	n=227,600	n=125,975	n=252,700	n= 32,225	n= 20,675	n= 3,825	n= 4,425	Overall Mean	Overall Max	Overall Min	Overall Avg Max	Overall Avg Min	Overall Std
Bin Mean:	$\mu = 1,945.59$	$\mu = 7,133.92$	$\mu = 18,458.80$	$\mu = 14,855.02$	$\mu = 16,075.87$	$\mu = 20,140.29$	$\mu = 25,629.46$	$\mu = 25,641.38$	$\mu = 33,384.33$	$\mu = 33,255.89$	$\mu = 36,156.35$	$\mu = 10,189.99$	$\mu = 287,624.00$	$\mu = 0.00$	$\mu = 287,624.00$	$\mu = 0.00$	$\mu = 24,202.54$
Bin Max:	$\tau = 173,477.00$	$\tau = 174,750.00$	$\tau = 223,274.00$	$\tau = 236,643.00$	$\tau = 239,189.00$	$\tau = 234,199.00$	$\tau = 258,079.00$	$\tau = 273,620.00$	$\tau = 287,624.00$	$\tau = 163,756.00$	$\tau = 241,535.00$	$\tau = 287,624.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 287,624.00$	$\tau = 0.00$	$\tau = 24,202.54$
Bin Min:	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$	$\iota = 0.00$
Bin Avg Max:	$\tau = 173,477.00$	$\tau = 174,750.00$	$\tau = 223,274.00$	$\tau = 236,643.00$	$\tau = 239,189.00$	$\tau = 234,199.00$	$\tau = 258,079.00$	$\tau = 273,620.00$	$\tau = 287,624.00$	$\tau = 163,756.00$	$\tau = 241,535.00$	$\tau = 287,624.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 287,624.00$	$\tau = 0.00$	$\tau = 24,202.54$
Bin Avg Min:	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$	$\tau = 0.00$
Bin Std:	$\sigma = 8,711.15$	$\sigma = 16,602.94$	$\sigma = 30,464.97$	$\sigma = 26,342.08$	$\sigma = 28,949.69$	$\sigma = 32,947.96$	$\sigma = 37,166.00$	$\sigma = 36,149.40$	$\sigma = 41,883.84$	$\sigma = 35,339.87$	$\sigma = 46,287.12$	$\sigma = 10,189.99$	$\sigma = 287,624.00$	$\sigma = 0.00$	$\sigma = 287,624.00$	$\sigma = 0.00$	$\sigma = 24,202.54$
Bin Mean:	$\mu = 32.43$ min	$\mu = 1.98$ hrs	$\mu = 5.13$ hrs	$\mu = 4.13$ hrs	$\mu = 4.47$ hrs	$\mu = 5.59$ hrs	$\mu = 7.12$ hrs	$\mu = 7.12$ hrs	$\mu = 9.27$ hrs	$\mu = 9.24$ hrs	$\mu = 10.04$ hrs	$\mu = 2.83$ hrs	$\mu = 3.33$ days	$\mu = 0.00$ sec	$\mu = 3.33$ days	$\mu = 0.00$ sec	$\mu = 6.72$ hrs
Bin Max:	$\tau = 2.01$ days	$\tau = 2.02$ days	$\tau = 2.58$ days	$\tau = 2.74$ days	$\tau = 2.77$ days	$\tau = 2.71$ days	$\tau = 2.99$ days	$\tau = 3.17$ days	$\tau = 3.33$ days	$\tau = 1.90$ days	$\tau = 2.80$ days	$\tau = 3.33$ days	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 3.33$ days	$\tau = 0.00$ sec	$\tau = 943$
Bin Min:	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec	$\iota = 0.00$ sec
Bin Avg Max:	$\tau = 2.01$ days	$\tau = 2.02$ days	$\tau = 2.58$ days	$\tau = 2.74$ days	$\tau = 2.77$ days	$\tau = 2.71$ days	$\tau = 2.99$ days	$\tau = 3.17$ days	$\tau = 3.33$ days	$\tau = 1.90$ days	$\tau = 2.80$ days	$\tau = 3.33$ days	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 3.33$ days	$\tau = 0.00$ sec	$\tau = 943$
Bin Avg Min:	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec	$\tau = 0.00$ sec
Bin Std:	$\sigma = 2.42$ hrs	$\sigma = 4.61$ hrs	$\sigma = 8.46$ hrs	$\sigma = 7.32$ hrs	$\sigma = 8.04$ hrs	$\sigma = 9.15$ hrs	$\sigma = 10.32$ hrs	$\sigma = 10.04$ hrs	$\sigma = 11.63$ hrs	$\sigma = 9.82$ hrs	$\sigma = 12.86$ hrs	$\sigma = 10,189.99$	$\sigma = 287,624.00$	$\sigma = 0.00$	$\sigma = 287,624.00$	$\sigma = 0.00$	$\sigma = 24,202.54$





n=1,712,900 n=143,450 n=696,550 n=241,225 n=227,600 n=125,975 n=252,700 n= 32,225 n= 20,675 n= 3,825 n= 4,425

Bin Mean: $\mu =$
 Bin Max: $t =$
 Bin Min: $i =$
 Bin Avg Max: $\bar{x} =$
 Bin Avg Min: $\bar{y} =$
 Bin Std: $\sigma =$

$\mu =$ 1,945.59	$\mu =$ 7,133.92	$\mu =$ 18,458.80	$\mu =$ 14,855.02	$\mu =$ 16,075.87	$\mu =$ 20,140.29	$\mu =$ 25,629.46	$\mu =$ 25,641.38	$\mu =$ 33,384.33	$\mu =$ 33,255.89	$\mu =$ 36,156.35
$t =$ 173,477.00	$t =$ 174,750.00	$t =$ 223,274.00	$t =$ 236,643.00	$t =$ 239,189.00	$t =$ 234,199.00	$t =$ 258,079.00	$t =$ 273,620.00	$t =$ 287,624.00	$t =$ 163,756.00	$t =$ 241,535.00
$i =$ 0.00	$i =$ 0.00	$i =$ 0.00	$i =$ 0.00	$i =$ 0.00	$i =$ 0.00	$i =$ 0.00	$i =$ 0.00	$i =$ 0.00	$i =$ 0.00	$i =$ 0.00
$\bar{x} =$ 173,477.00	$\bar{x} =$ 174,750.00	$\bar{x} =$ 223,274.00	$\bar{x} =$ 236,643.00	$\bar{x} =$ 239,189.00	$\bar{x} =$ 234,199.00	$\bar{x} =$ 258,079.00	$\bar{x} =$ 273,620.00	$\bar{x} =$ 287,624.00	$\bar{x} =$ 163,756.00	$\bar{x} =$ 241,535.00
$\bar{y} =$ 0.00	$\bar{y} =$ 0.00	$\bar{y} =$ 0.00	$\bar{y} =$ 0.00	$\bar{y} =$ 0.00	$\bar{y} =$ 0.00	$\bar{y} =$ 0.00	$\bar{y} =$ 0.00	$\bar{y} =$ 0.00	$\bar{y} =$ 0.00	$\bar{y} =$ 0.00
$\sigma =$ 8,711.15	$\sigma =$ 16,602.94	$\sigma =$ 30,464.97	$\sigma =$ 26,342.08	$\sigma =$ 28,949.69	$\sigma =$ 32,947.96	$\sigma =$ 37,166.00	$\sigma =$ 36,149.40	$\sigma =$ 41,883.84	$\sigma =$ 35,339.87	$\sigma =$ 46,287.12
$\mu =$ 32.43 min	$\mu =$ 1.98 hrs	$\mu =$ 5.13 hrs	$\mu =$ 4.13 hrs	$\mu =$ 4.47 hrs	$\mu =$ 5.59 hrs	$\mu =$ 7.12 hrs	$\mu =$ 7.12 hrs	$\mu =$ 9.27 hrs	$\mu =$ 9.24 hrs	$\mu =$ 10.04 hrs
$t =$ 2.01 days	$t =$ 2.02 days	$t =$ 2.58 days	$t =$ 2.74 days	$t =$ 2.77 days	$t =$ 2.71 days	$t =$ 2.99 days	$t =$ 3.17 days	$t =$ 3.33 days	$t =$ 1.90 days	$t =$ 2.80 days
$i =$ 0.00 sec	$i =$ 0.00 sec	$i =$ 0.00 sec	$i =$ 0.00 sec	$i =$ 0.00 sec	$i =$ 0.00 sec	$i =$ 0.00 sec	$i =$ 0.00 sec	$i =$ 0.00 sec	$i =$ 0.00 sec	$i =$ 0.00 sec
$\bar{x} =$ 2.01 days	$\bar{x} =$ 2.02 days	$\bar{x} =$ 2.58 days	$\bar{x} =$ 2.74 days	$\bar{x} =$ 2.77 days	$\bar{x} =$ 2.71 days	$\bar{x} =$ 2.99 days	$\bar{x} =$ 3.17 days	$\bar{x} =$ 3.33 days	$\bar{x} =$ 1.90 days	$\bar{x} =$ 2.80 days
$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec	$\bar{y} =$ 0.00 sec
$\sigma =$ 2.42 hrs	$\sigma =$ 4.61 hrs	$\sigma =$ 8.46 hrs	$\sigma =$ 7.32 hrs	$\sigma =$ 8.04 hrs	$\sigma =$ 9.15 hrs	$\sigma =$ 10.32 hrs	$\sigma =$ 10.04 hrs	$\sigma =$ 11.63 hrs	$\sigma =$ 9.82 hrs	$\sigma =$ 12.86 hrs

Overall Mean: 10,189.986382
 Binned Mean

Overall Mean $\mu =$ 10,189.99
 Overall Max $t =$ 287,624.00
 Overall Min $i =$ 0.00
 Overall Avg Max $\bar{x} =$ 287,624.00
 Overall Avg Min $\bar{y} =$ 0.00
 Overall Std $\sigma =$ 24,202.54


Overall Mean $\mu =$ 2.83 hrs
 Overall Max $t =$ 3.33 days
 Overall Min $i =$ 0.00 sec
 Overall Avg Max $\bar{x} =$ 3.33 days
 Overall Avg Min $\bar{y} =$ 0.00 sec
 Overall Std $\sigma =$ 6.72 hrs




What's happening 'around' these reservations?

What's happening 'around' these reservations?

Determine impact to different job classes





SLUG / BYU September 12-13, 2023

Evalys

Evalys - Overview

pypi v4.0.6

Offline or online cluster monitoring



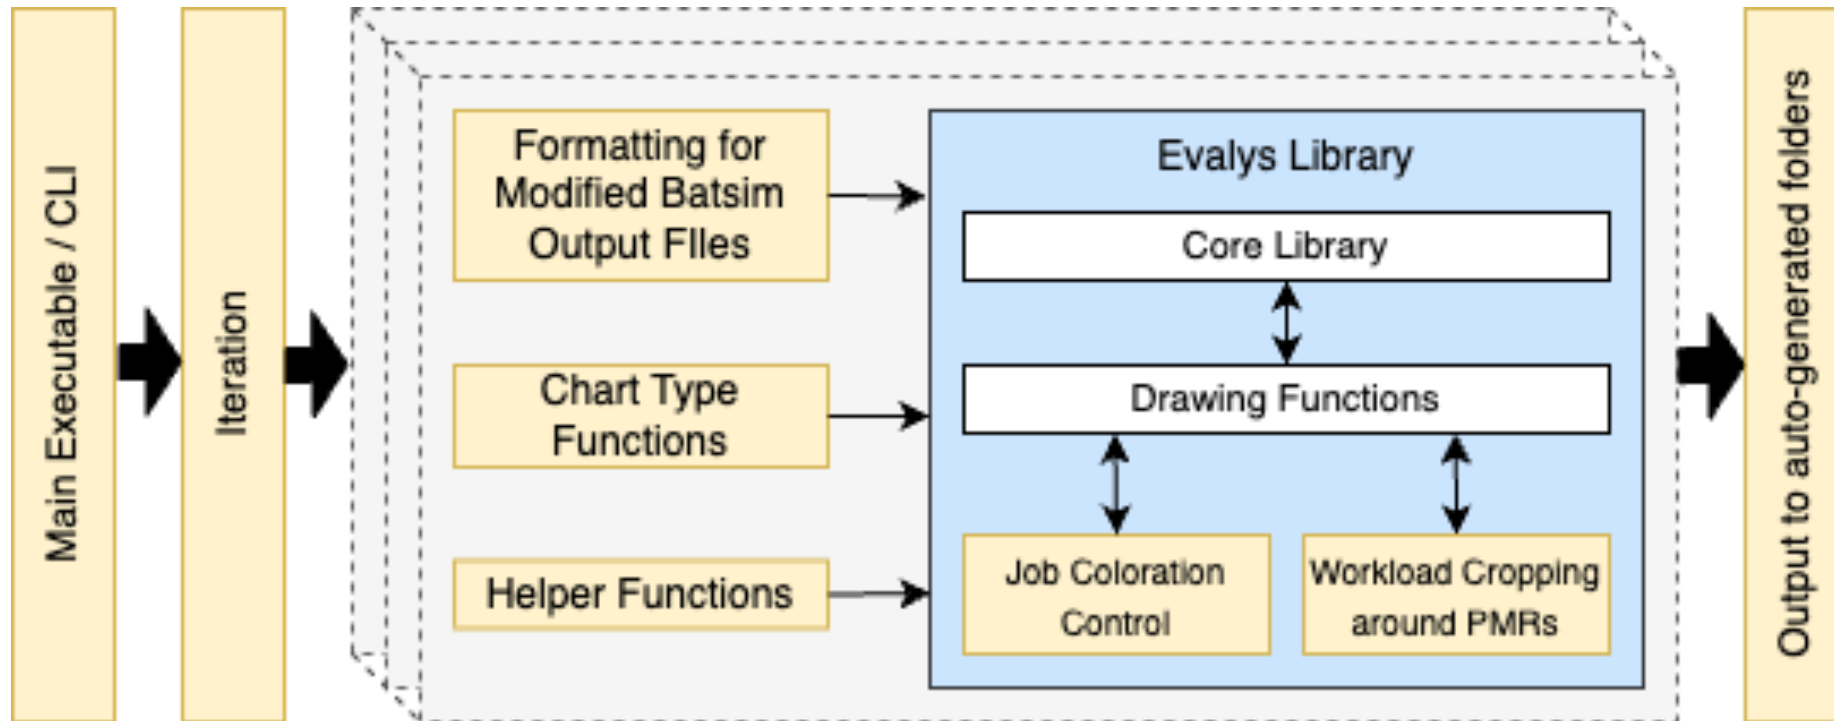
Features

- Load and all [Batsim](#) outputs files
 - Compute and plot free slots
 - Simple Gantt visualisation
 - Compute utilisation / queue
 - Compute fragmentation
 - Plot energy and machine state
- Load SWF workload files from [Parallel Workloads Archive](#)
 - Compute standard scheduling metrics
 - Show job details
 - Extract periods with a given mean utilisation



from gallery


Modifications to Evalys



12-hour

Window from 5486400.0-7344000.0+-169200S


Full




12-hour

Window from 5486400.0-7344000.0+-169200S

Full




12-hour
Full




12-hour
Full

Reservation from 7300800-7344000+-169200S




12-hour
Full

Reservation from 7300800-7344000+-169200S






12-hour
Full


Reservation from 7100800-7344000+-169200S




12-hour Full



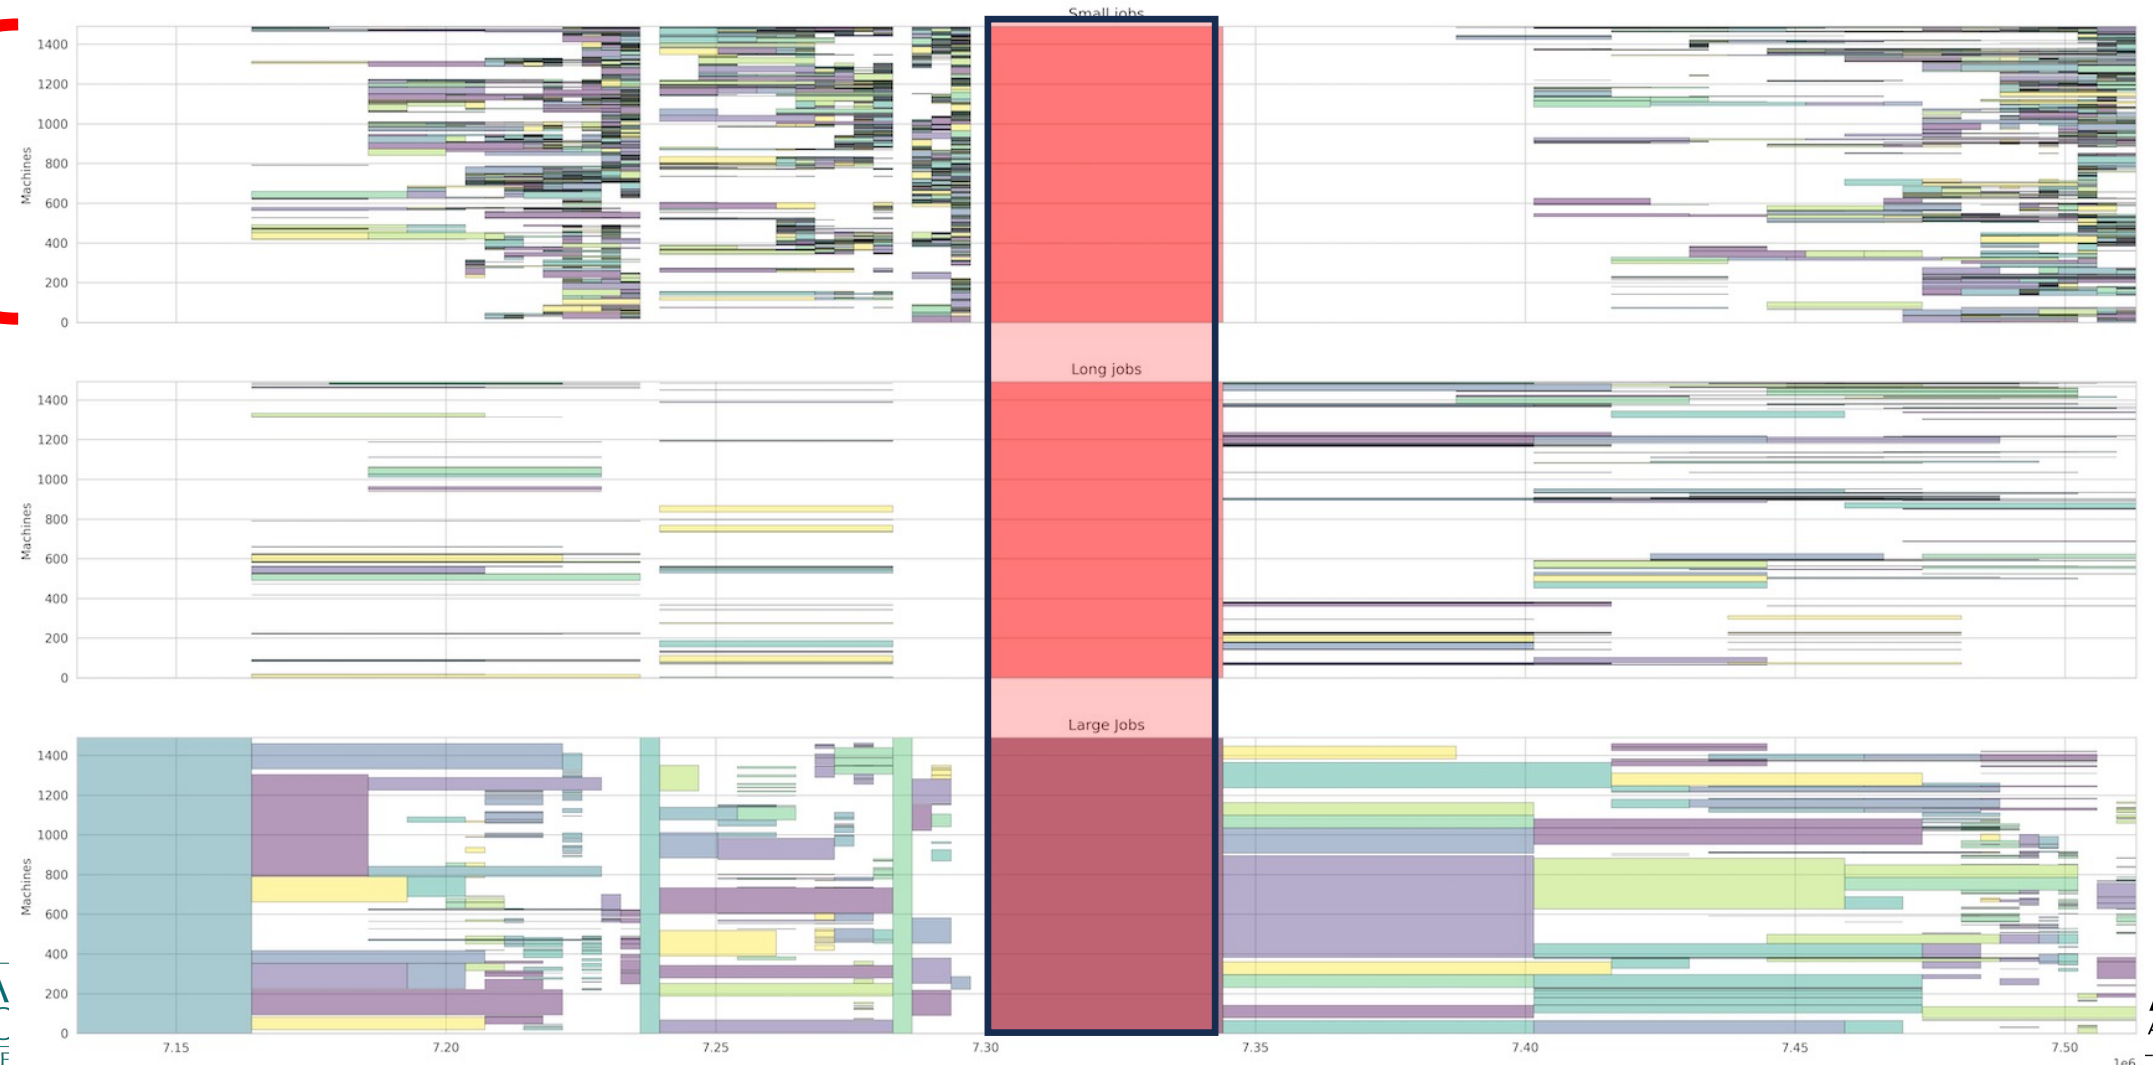
12-hour Full




12-hour Full



small




12-hour Full




small



long



12-hour Full



small




long



large



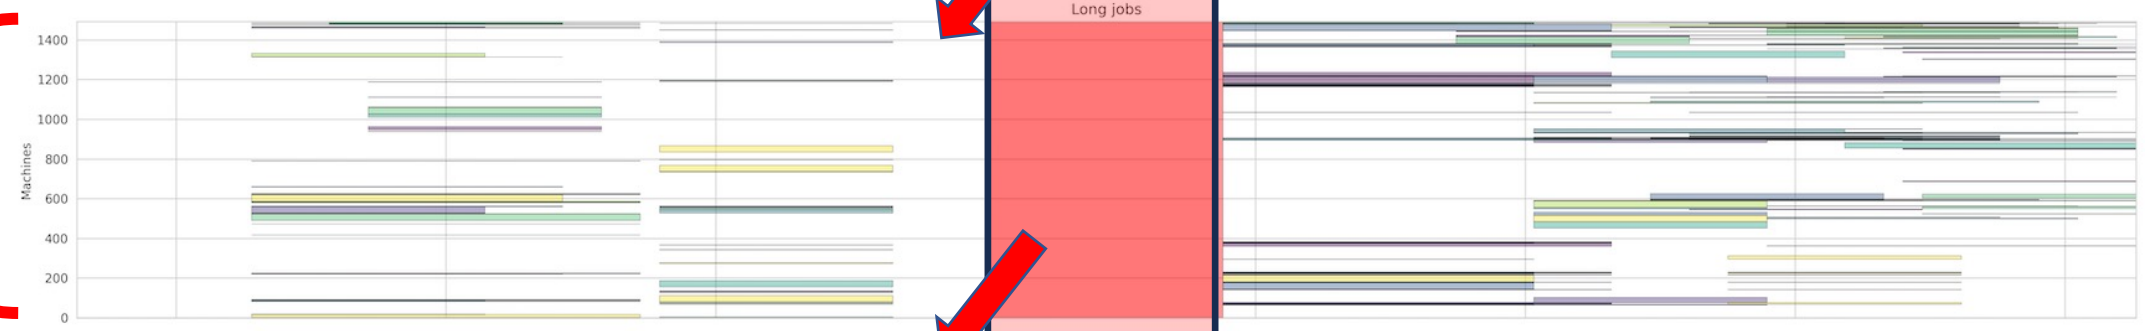
12-hour Full



small




long



large




12-hour Full




small



long



large



12-Full




12-Half

12-hour

Half


Window from 5486400.0-7344000.0+-169200S






12-hour











Half

Reservation from 6091200-6134400+-169200S




12-hour Half







Reservation from 2635200-2664000+-1692005




Reservation from 5227200-5256000+-1692005




Reservation from 7819200-7848000+-1692005




Reservation from 10411200-10440000+-1692005




Reservation from 13003200-13032000+-1692005




Reservation from 15595200-15624000+-1692005




Reservation from 18187200-18216000+-1692005



Reservation from 20779200-20808000+-1692005



Reservation from 23371200-23400000+-1692005



Reservation from 25963200-25992000+-1692005

