Computing and Software for Big Science paper

Sean Wilkinson
University of Texas at Arlington
24 April 2019

Status

Note: see https://indico.cern.ch/event/812706/ links.

- We are close! The end is near!
- It looks like a paper, but it does not read like a paper yet.
- Section 4 (effect on Titan) content is finished.
  - “Inconclusive” results require careful handling.
  - As usual, this is most of what I will focus on.

Optimism

- It already looks like a paper.
- When content is approved, I can make this read like a paper in short order, I promise.
- Most of the content has been approved.
  - Remember the “X to write, Y to check” stuff?
- ⇒ We are nearly done! The end is near!

Section 4

- There have been very substantial changes to Section 4 since the last TIM.
- Spoiler: still haven’t really found any effects.
- I need everyone’s brilliant minds to check this.
- I apologize in advance to those who have had to sit through this already!

Short version

- I have only ever found evidence that is suggestive of certain interpretations.
- Everything in this slide show has already been committed into the draft repository.
- If approved by others, I am ready to close this case.

Introduction

- Basic history about project
- Specifics on Titan which may belong in Section 3
- “The goal of CSC108 has been to consume idle resources on Titan which would otherwise have gone to waste, while making a good-faith effort not to disturb the rest of Titan’s ecosystem.”

Subsection: “Compression study”

- Needs a more sophisticated name
- Study was rescheduling (without reordering) 3 years of log traces with and without CSC108, to test “displacement” due to CSC108.
- Algorithm is shown in paper but omitted here because the text was really small.

Plot to show successful consumption of idle resources

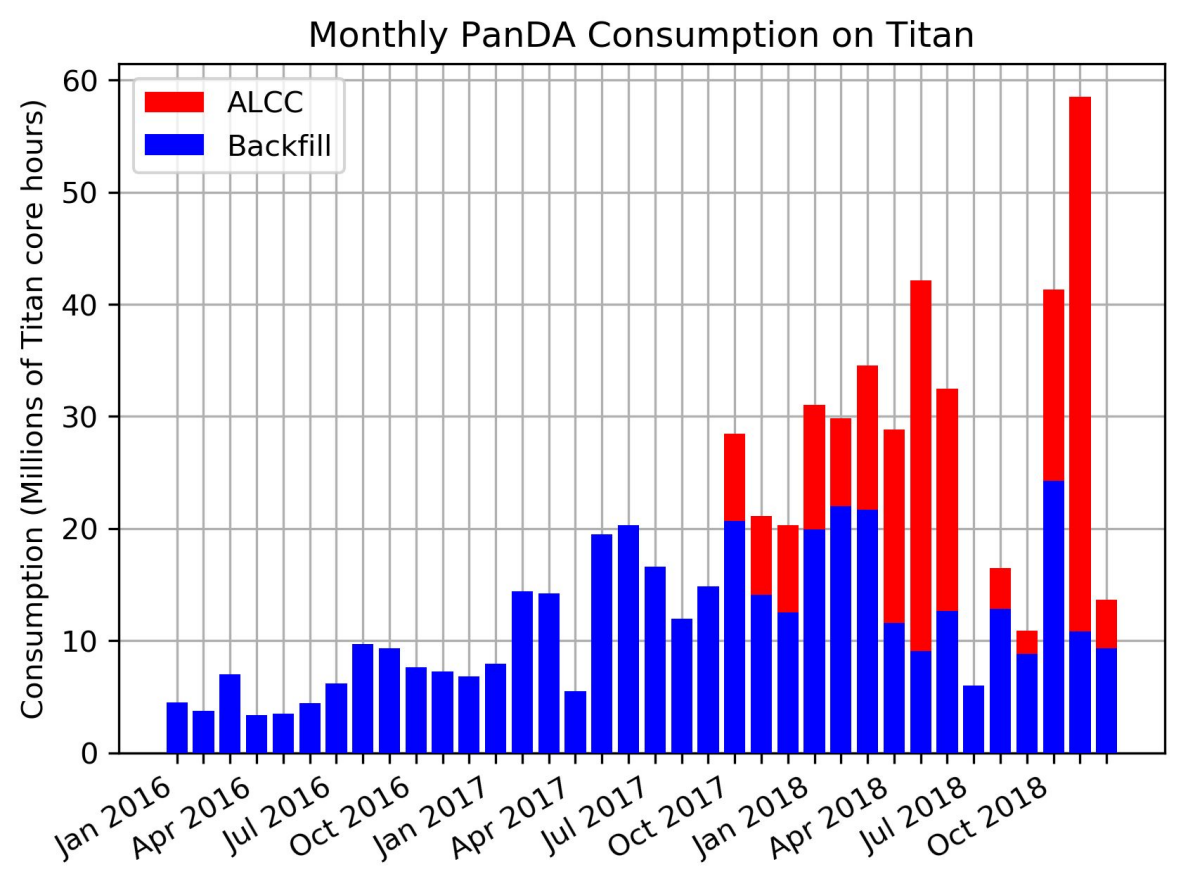

Plot to suggest that there is competition for resources

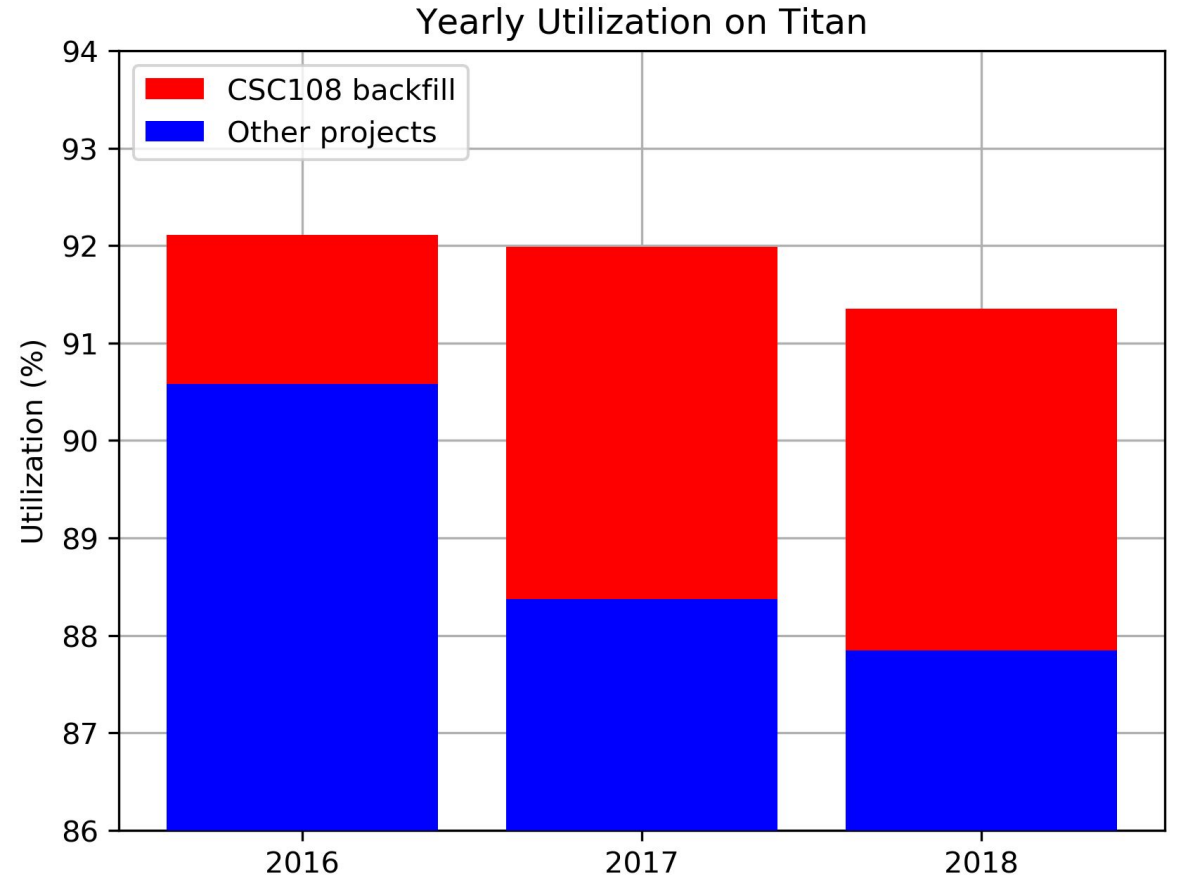

Table of results from the compression study

	Without CSC108	With CSC108	Percent change
Time to completion (days)	1021.2	1034.5	1.30
Throughput (jobs completed per day)	1324.93	1515.19	14.36
Utilization (percent)	92.36	94.15	1.94

Results of “compression study”

- “The results, which are shown in Table 2, suggest that the hypothesis that CSC108 has no effect on Titan should be rejected.”
- “More importantly, however, these results suggest that CSC108 has successfully consumed idle resources which would otherwise have gone to waste.”

Subsection: Simple linear relationships

- Data now use the three years of traces along with daily availability data for Titan provided by OLCF.
- Methods are Ordinary Least Squares (OLS) linear regression, focusing on throughput and utilization, while separating CSC108 jobs by bin and checking goodness of fit with R^2.

Figure 7a (shown here alone for clarity); R^2 goodness of fit: 0.0040

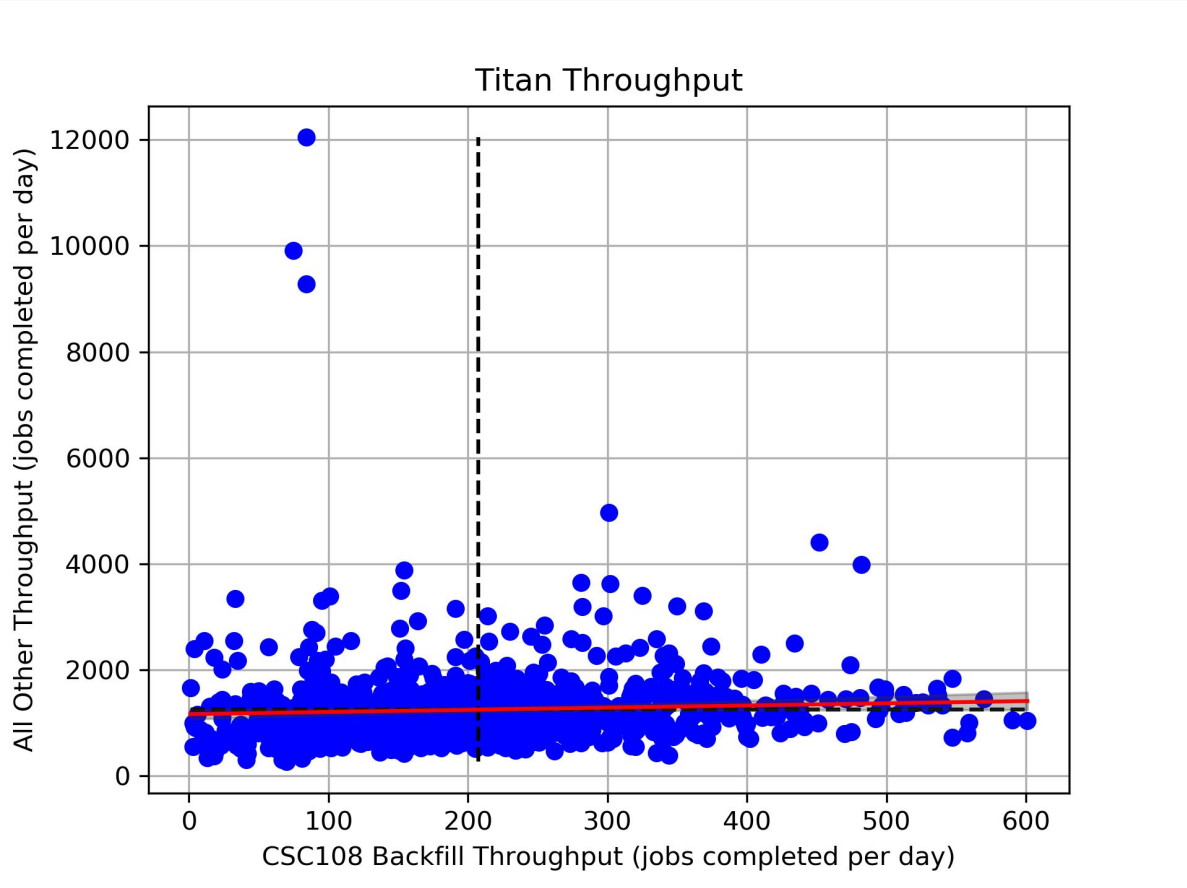

Figure 7b (shown here alone for clarity); R^2 goodness of fit: 0.0005

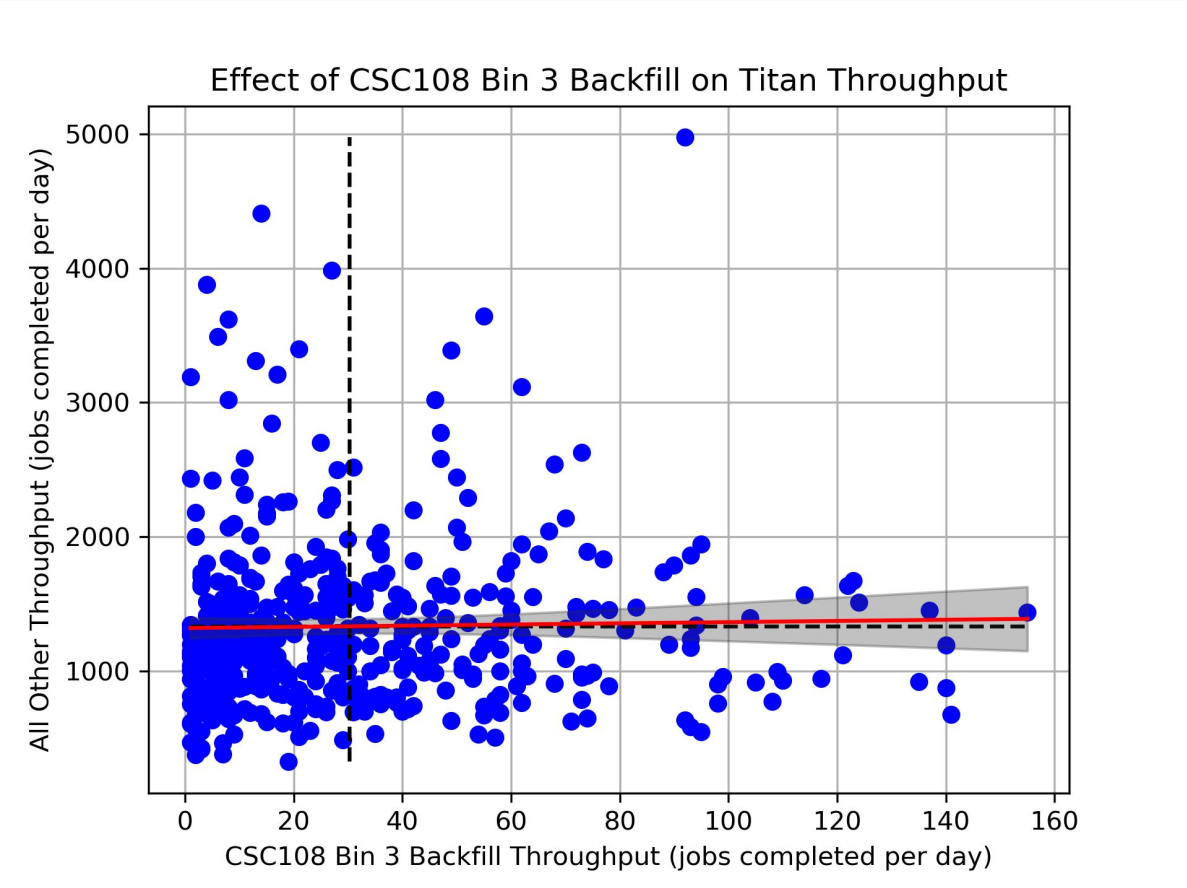

Figure 7c (shown here alone for clarity); R^2 goodness of fit: 0.0027

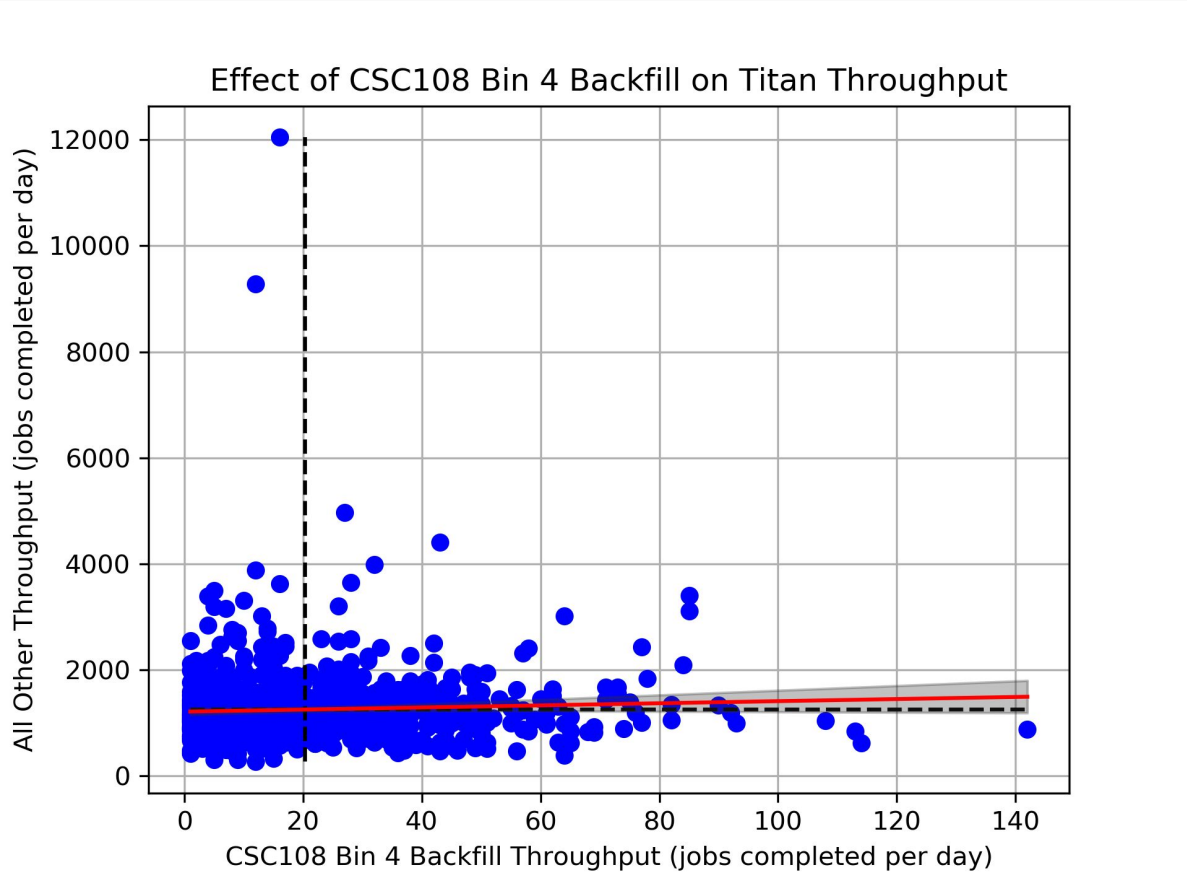

Figure 7d (shown here alone for clarity); R^2 goodness of fit: 0.0018

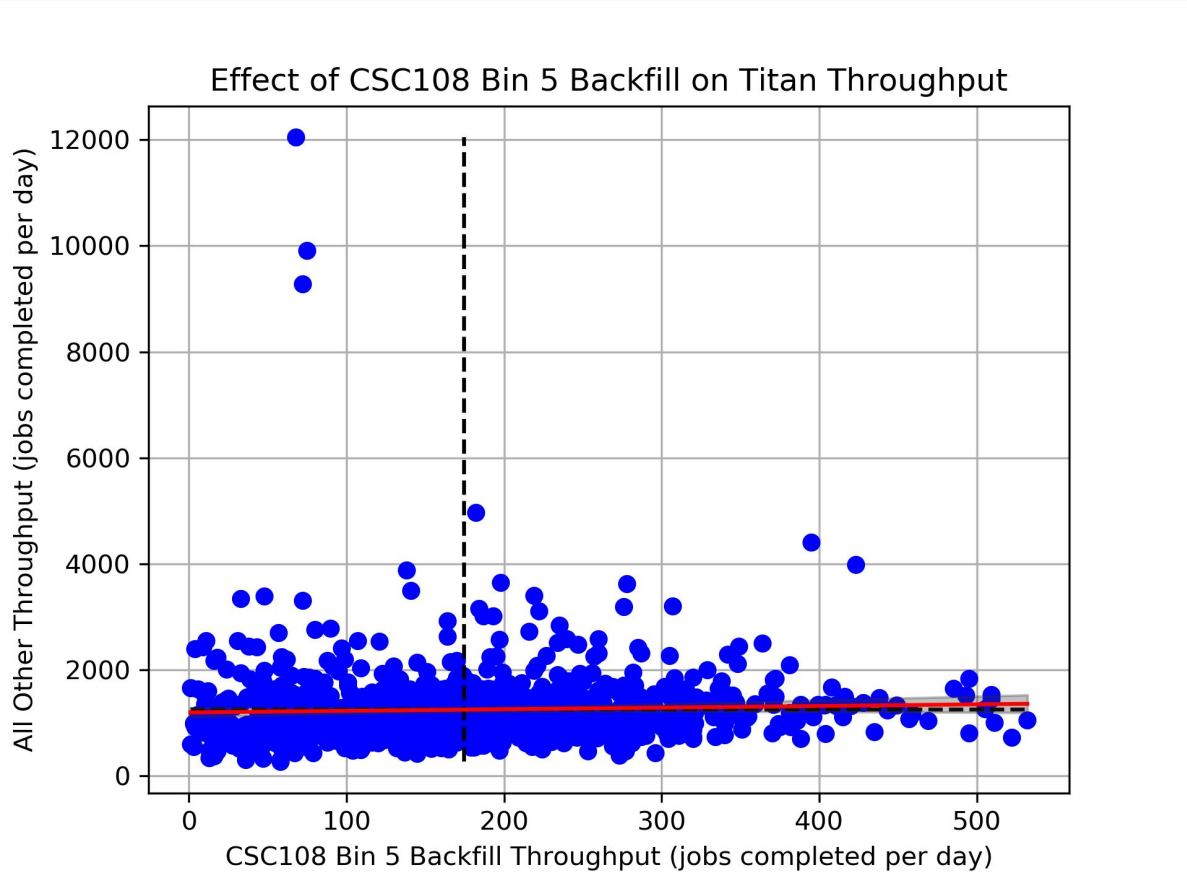

Table of model parameters and goodness of fit for throughput relationships

Figure	OLCF Bin	Slope	Y intercept	R^2
7a	All	0.4106	1164.2561	0.0040
7b	3	0.4419	1322.0784	0.0005
7c	4	1.9819	1211.3384	0.0027
7d	5	0.3072	1195.6684	0.0018

Figure 8a (shown here alone for clarity); R^2 goodness of fit: 0.0330

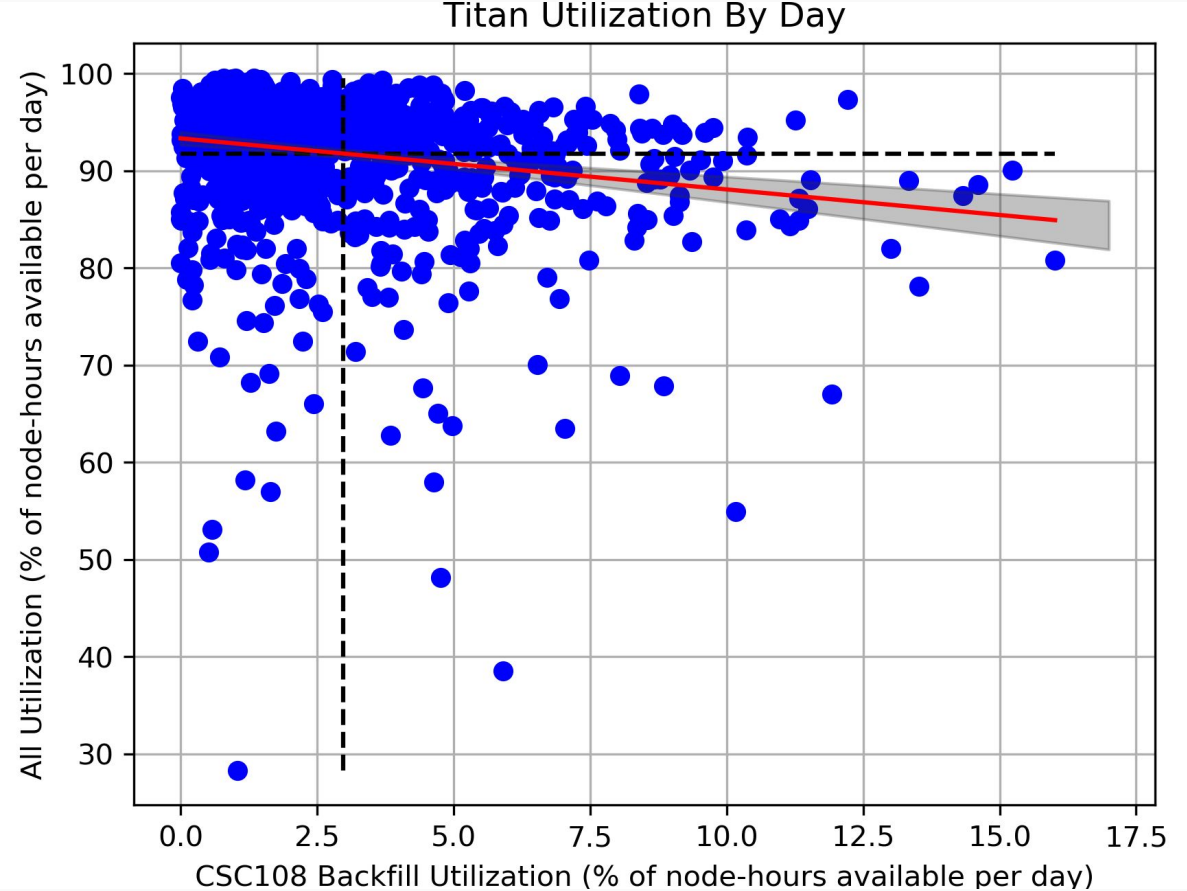

Figure 8b (shown here alone for clarity); R^2 goodness of fit: 0.1359

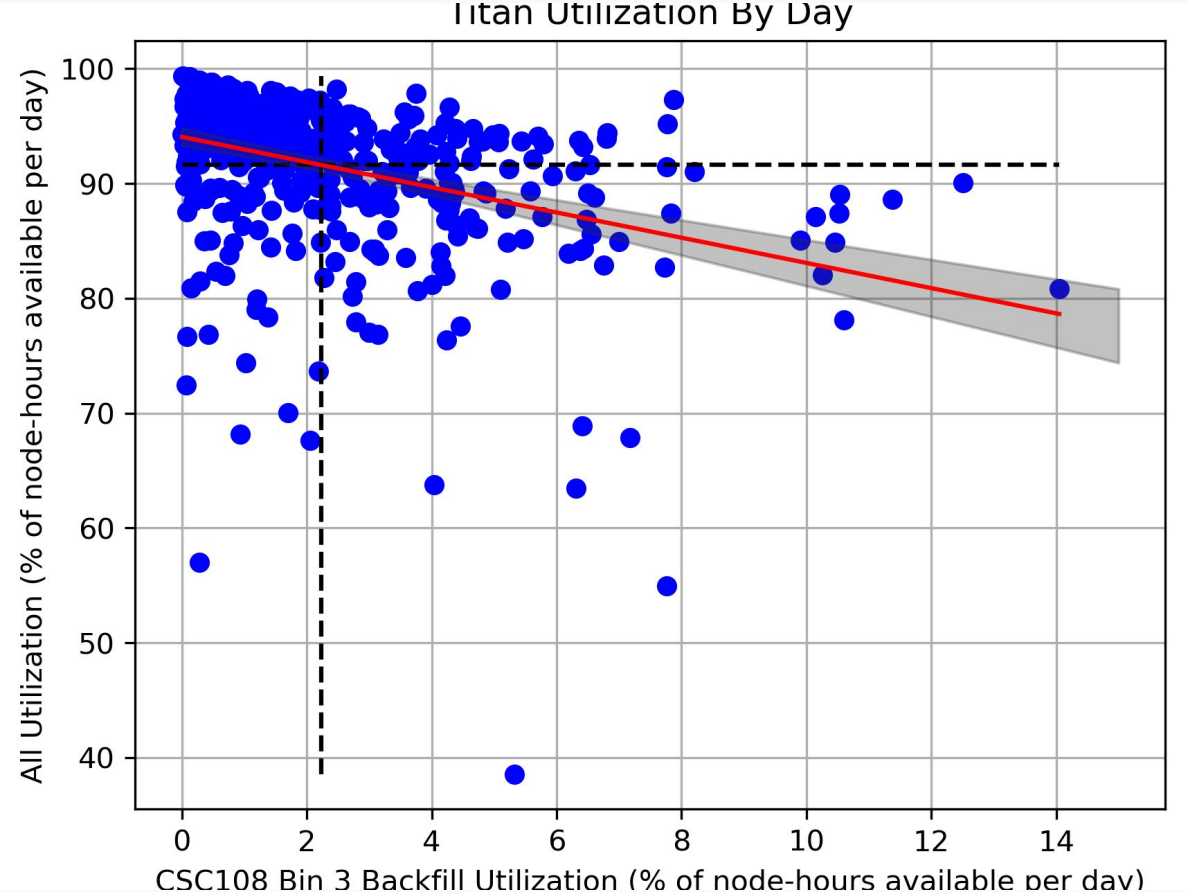

Figure 8c (shown here alone for clarity); R^2 goodness of fit: 0.0378

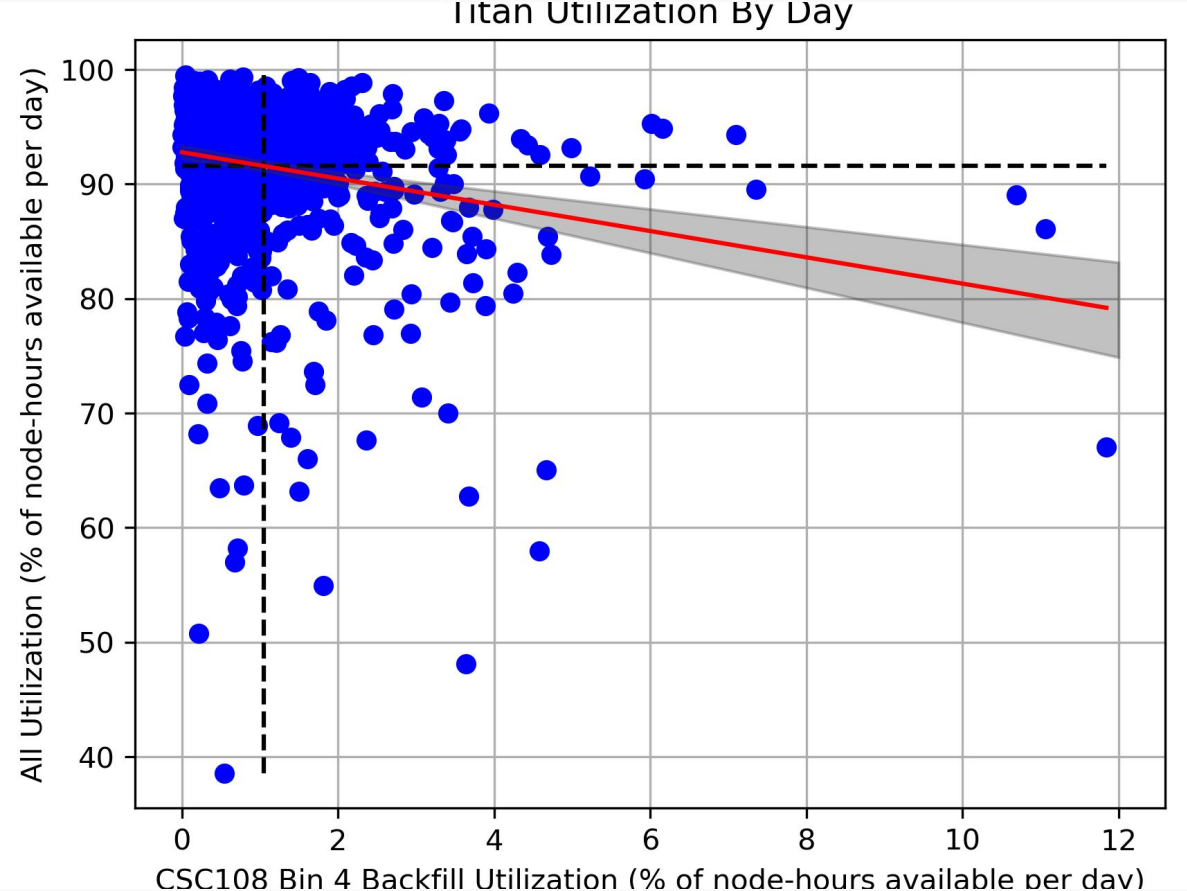

Figure 8d (shown here alone for clarity); R^2 goodness of fit: 0.1046

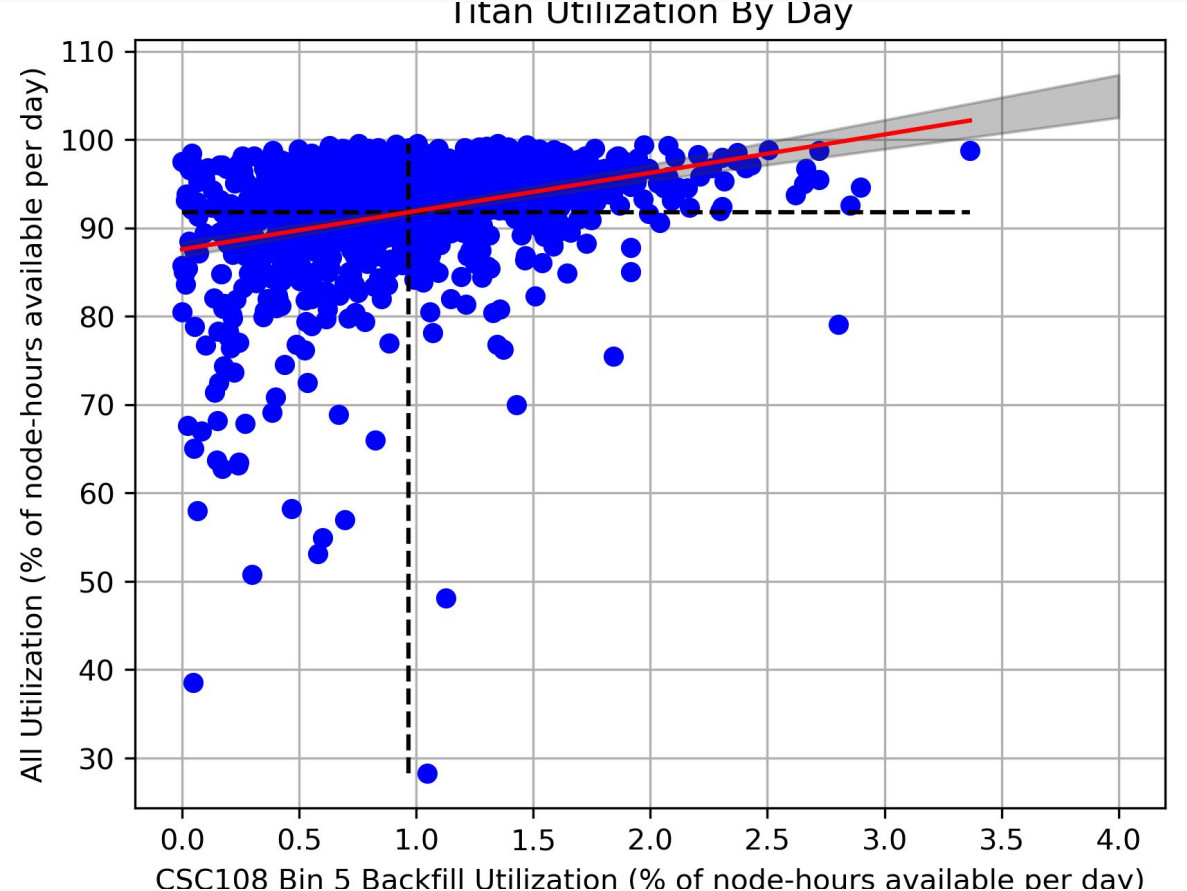

Table of model parameters and goodness of fit for utilization relationships

Figure	OLCF Bin	Slope	Y intercept	R^2
8a	All	-0.5258	93.3404	0.0330
8b	3	-1.0977	94.0609	0.1359
8c	4	-1.1472	92.7870	0.0378
8d	5	4.3328	87.5839	0.1046

Results for simple linear relationships

- Throughput increases across all bins, but fits are poor.
- Utilization decreases except for bin 5, but all fits are poor.
- It’s not easy to write about inconclusive results. I did what I thought was best, but I seriously appreciate input on how it can be improved or even rewritten in the draft.

Subsection: Blocking probability

- Data now also includes polling data from Moab.
- Formal definitions are improved but do not use equations.
- We now consider wait times as a third indicator.
- I argue that blocking probability can be used as an indicator for times of competition for resources.

Aside about naming

For the purposes of our discussion today, I have not changed the name of the concept we have been calling “blocking probability”. This is because we need to focus on logic right now. But in the paper, we probably need to change the name, because blocking probability is a technical term in telecommunication stuff.

Formal definition of blocking probability

Let C_i be the abstract resources in use by CSC108 at the i^{th} sample point in time, and let U_i be the unused (idle) resources remaining on Titan. We then define a boolean B_i representing a “block” to be 1 if there exists at least one job at the i^{th} sample point which requests $(C_i + U_i)$ resources or less when C_i is non-zero; we define B_i to be zero otherwise. Summing B_i over all i gives a count of sample points at which a block occurred, and dividing that count by the number of total sample points yields a quantity we call a “blocking probability”. The blocking probability is a rational number between 0 and 1.

Intuition behind blocking probability

It represents the proportion of samples in which a block occurred. The idea here is that when blocking probability increases, the system is experiencing greater competition for its resources. Blocking probability does not predict the probability that a particular job will be blocked, but rather the probability that a given sample will contain a block.

One-dimensional blocking

- Spatial blocking indicates insufficient total nodes.
- Temporal blocking indicates insufficient total wall time.
- “Due to CSC108” means at least one blocked job would be unblocked if CSC108’s resources were available:
  - “Spatial due to CSC108” refers to CSC108’s nodes.
  - “Temporal due to CSC108” is the same for wall time.

Figure 9a (shown here alone for clarity)

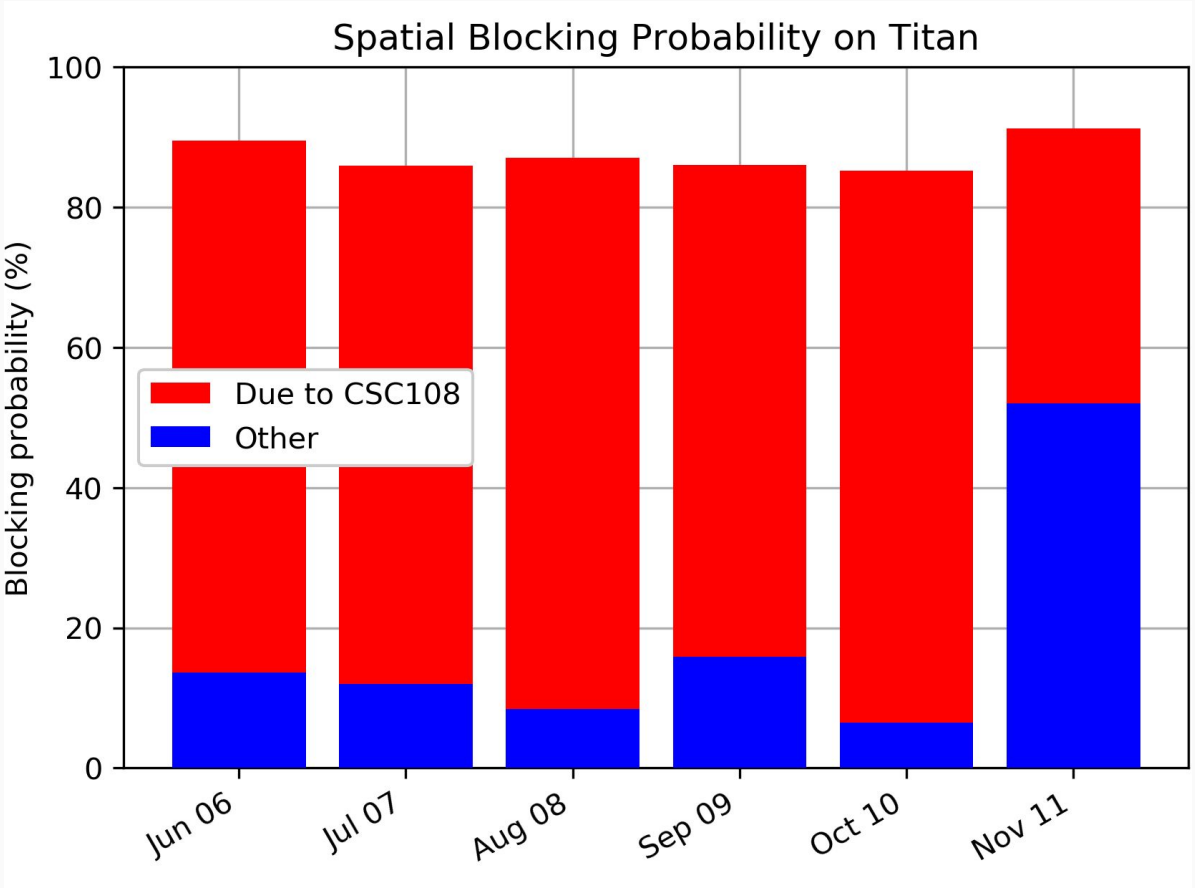

Figure 9b (shown here alone for clarity)

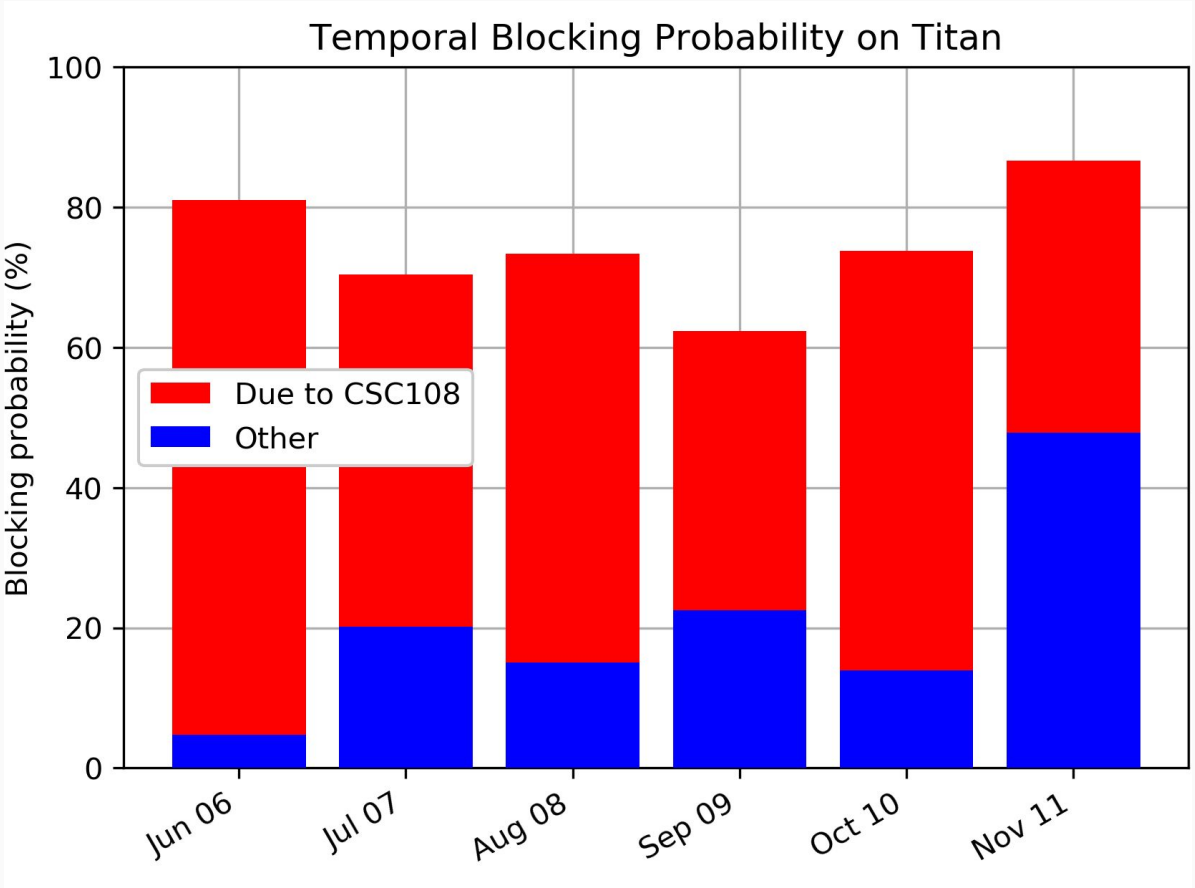

Aside on previous two graphs

- I presented this material to a fresh audience at Oak Ridge National Lab recently, and they found the stacked bars misleading.
- I agree with them.
- I forgot to remake the plots before writing these slides.

Spatial vs Temporal Blocking on Titan; R^2 goodness of fit: 0.4410

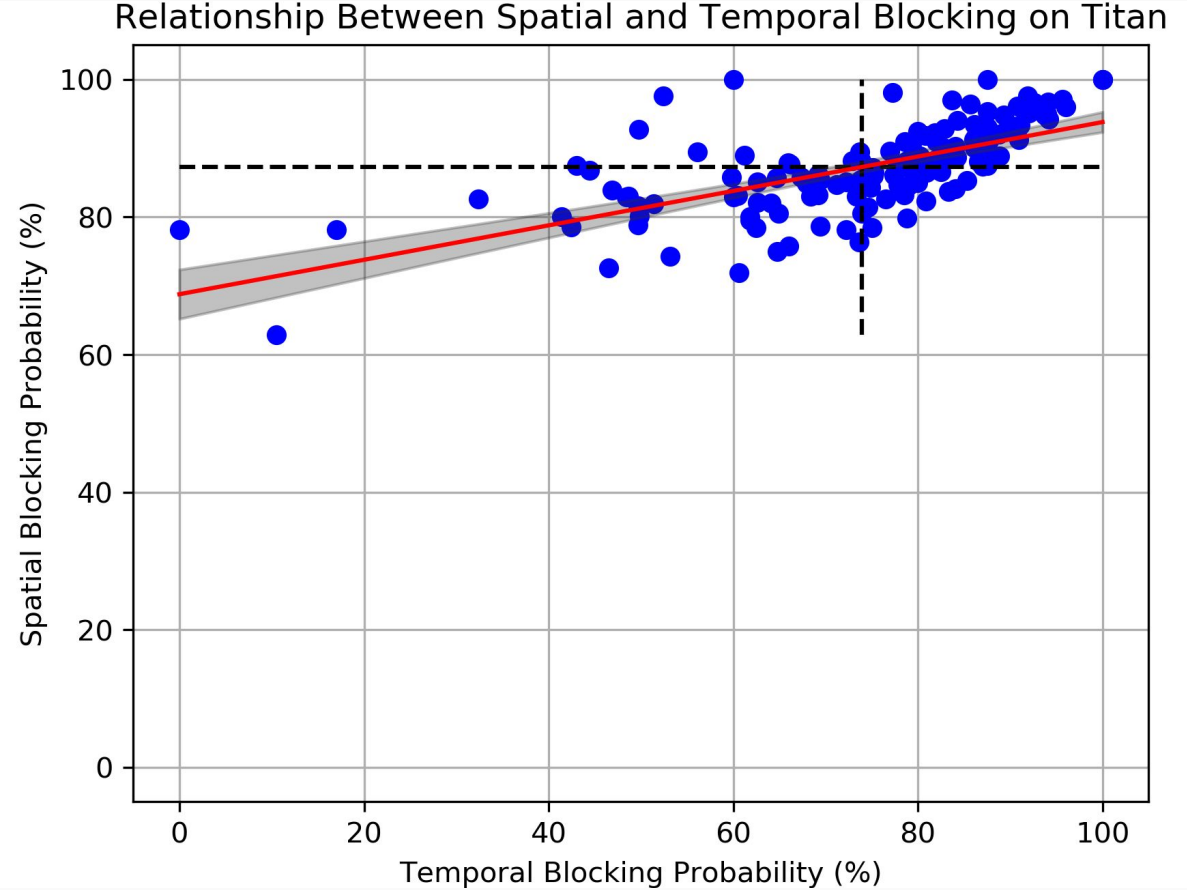

Figure 11a (shown here alone for clarity); R^2 goodness of fit: 0.0737

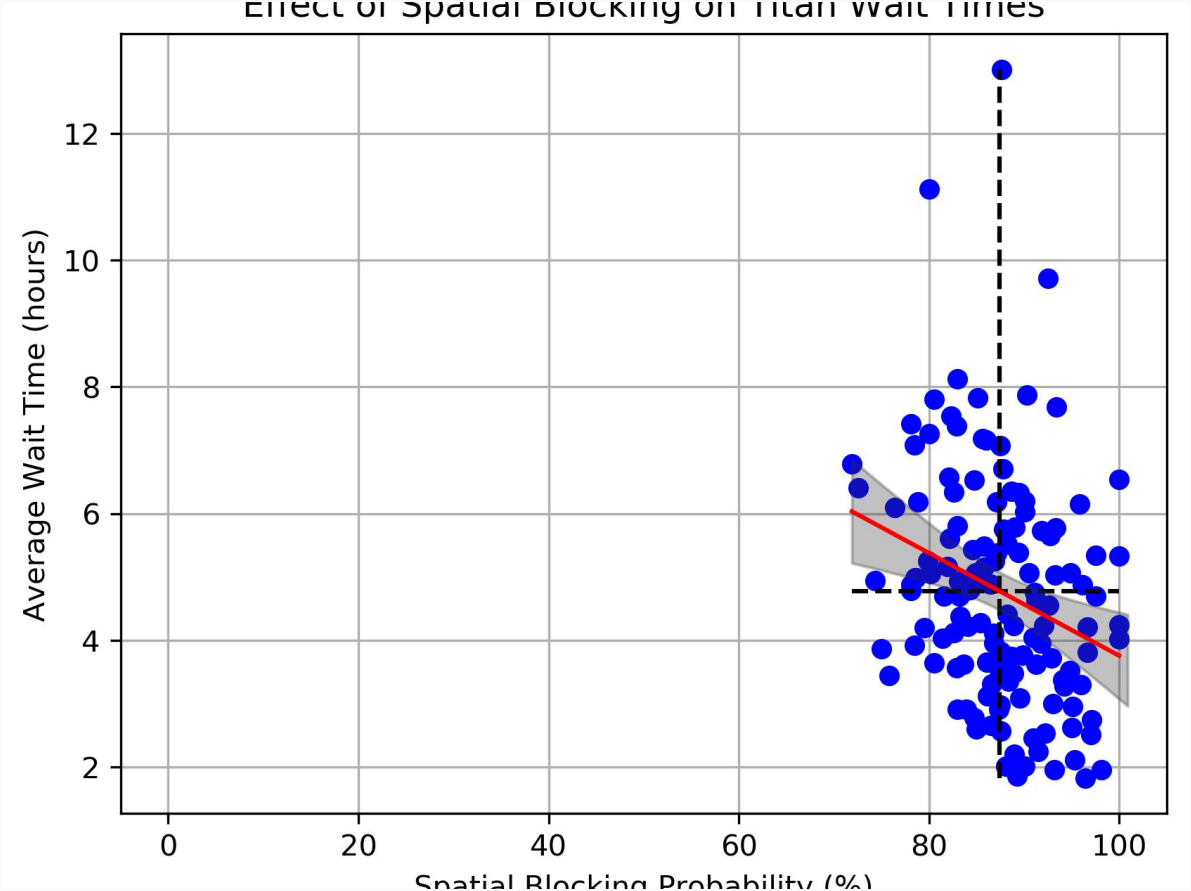

Figure 11b (shown here alone for clarity); R^2 goodness of fit: 0.1265

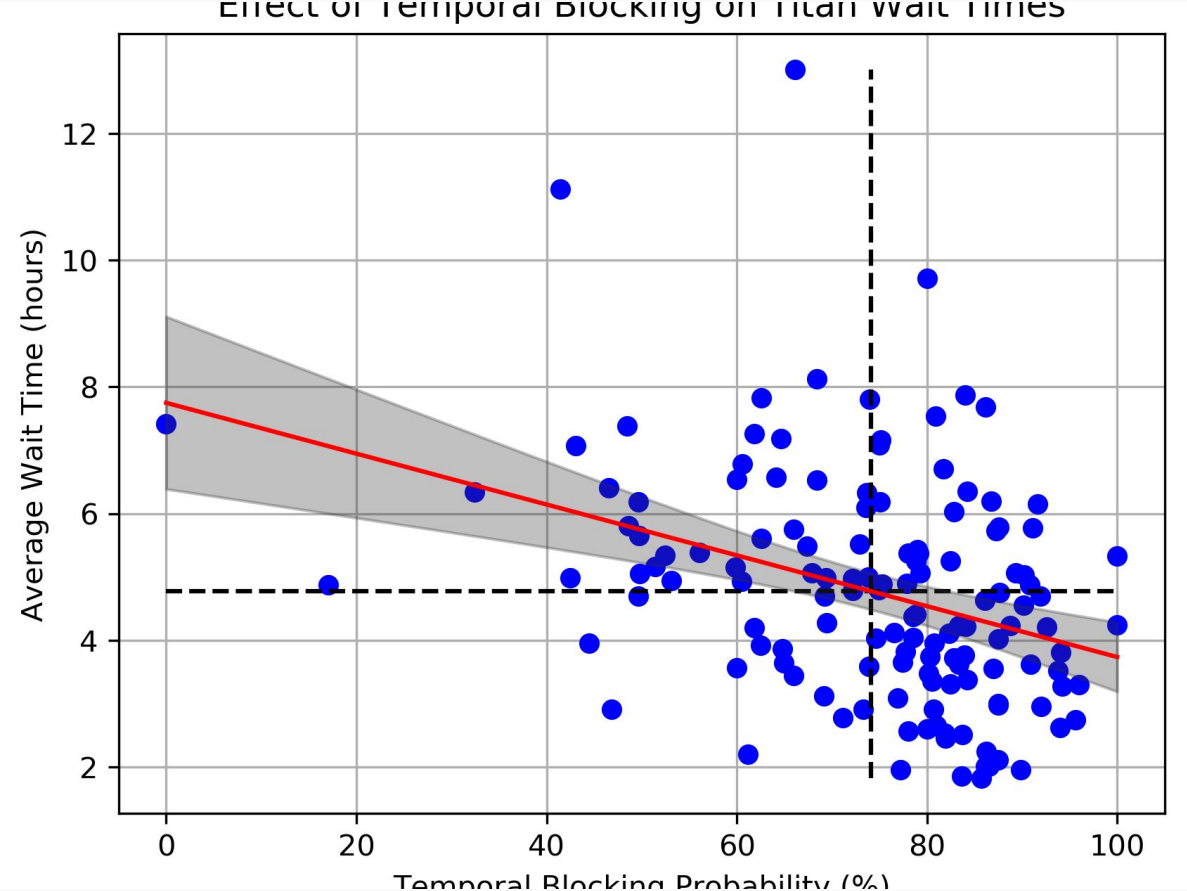

Figure 11c (shown here alone for clarity); R^2 goodness of fit: 0.0509

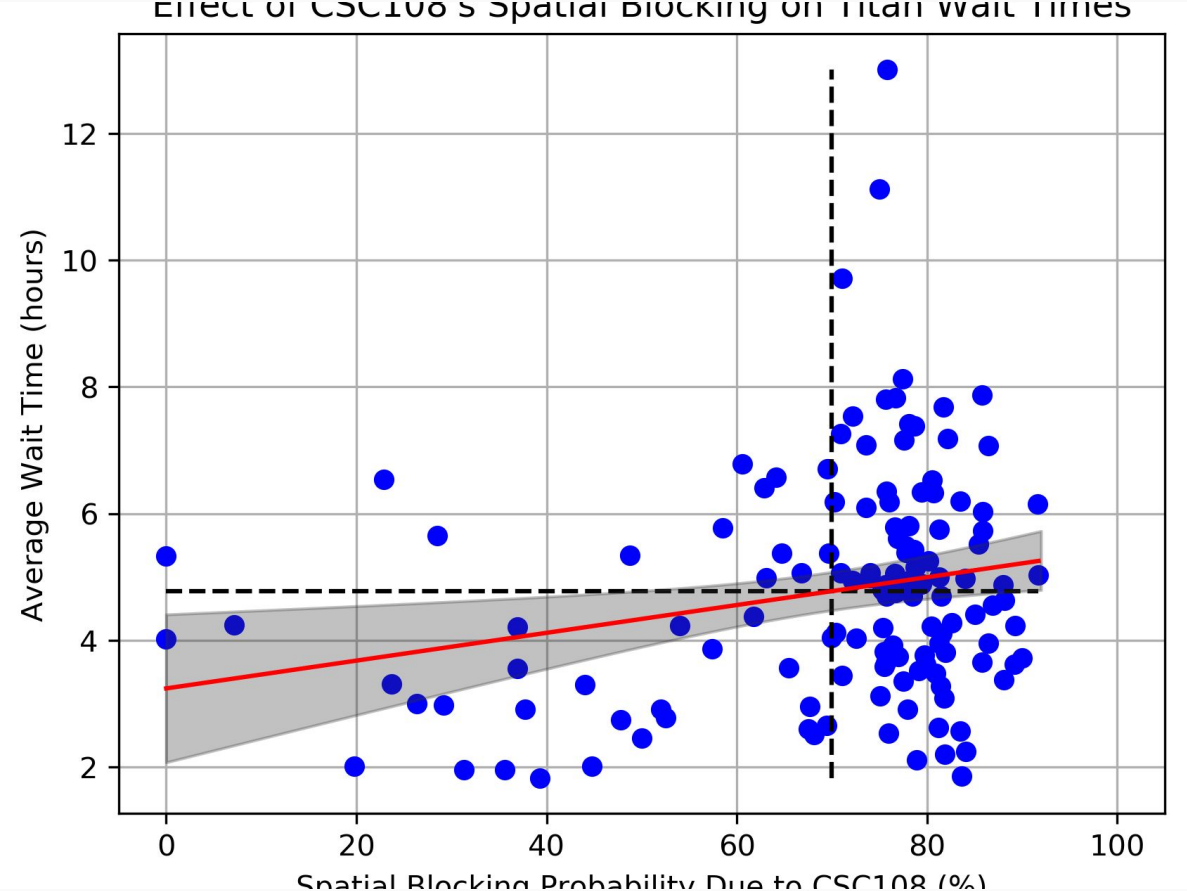

Figure 11d (shown here alone for clarity); R^2 goodness of fit: 0.0147

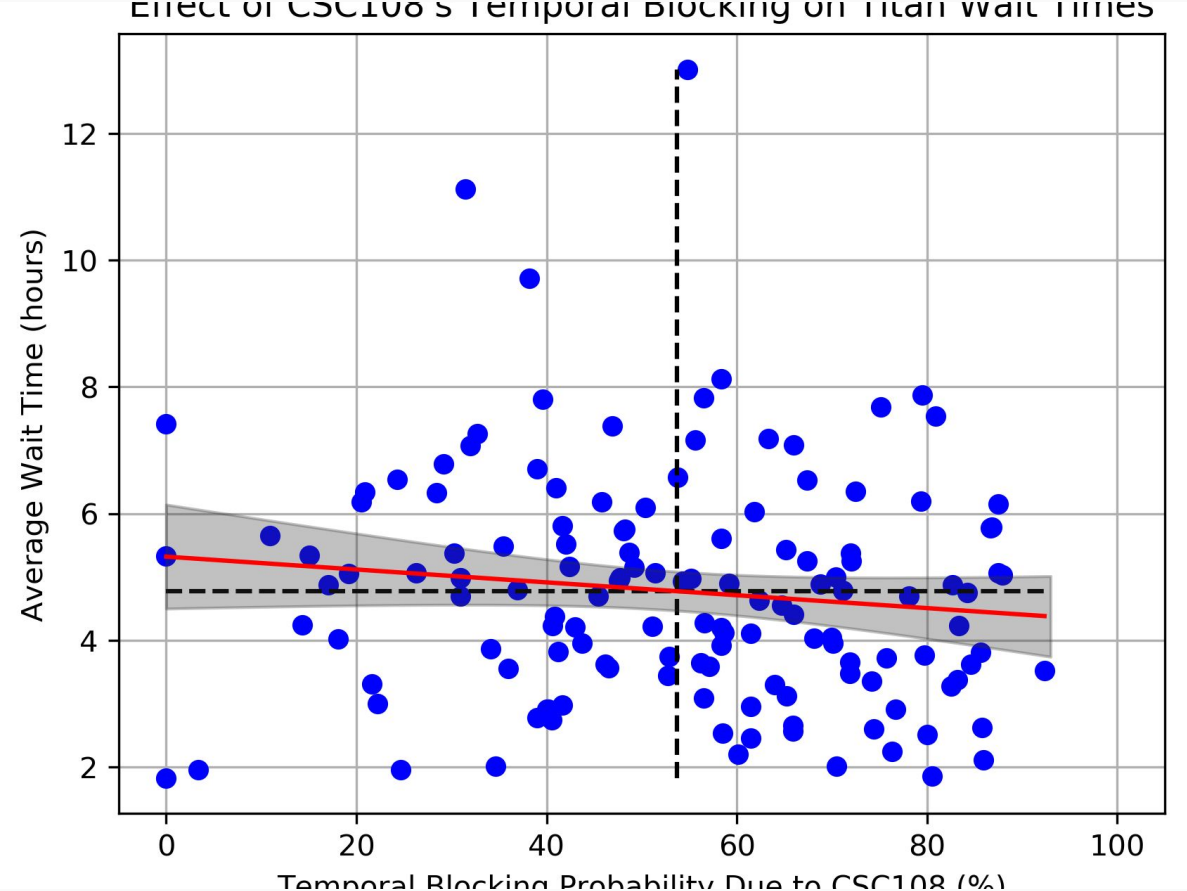

Table of model parameters et al. for average wait time vs blocking relationships

Figure	Slope	Y intercept	R^2
11a	-0.0810	11.8610	0.0737
11b	-0.0401	7.7491	0.1265
11c	0.0219	3.2420	0.0509
11d	-0.0102	5.3217	0.0147

Figure 12a (shown here alone for clarity); R^2 goodness of fit: 0.0122

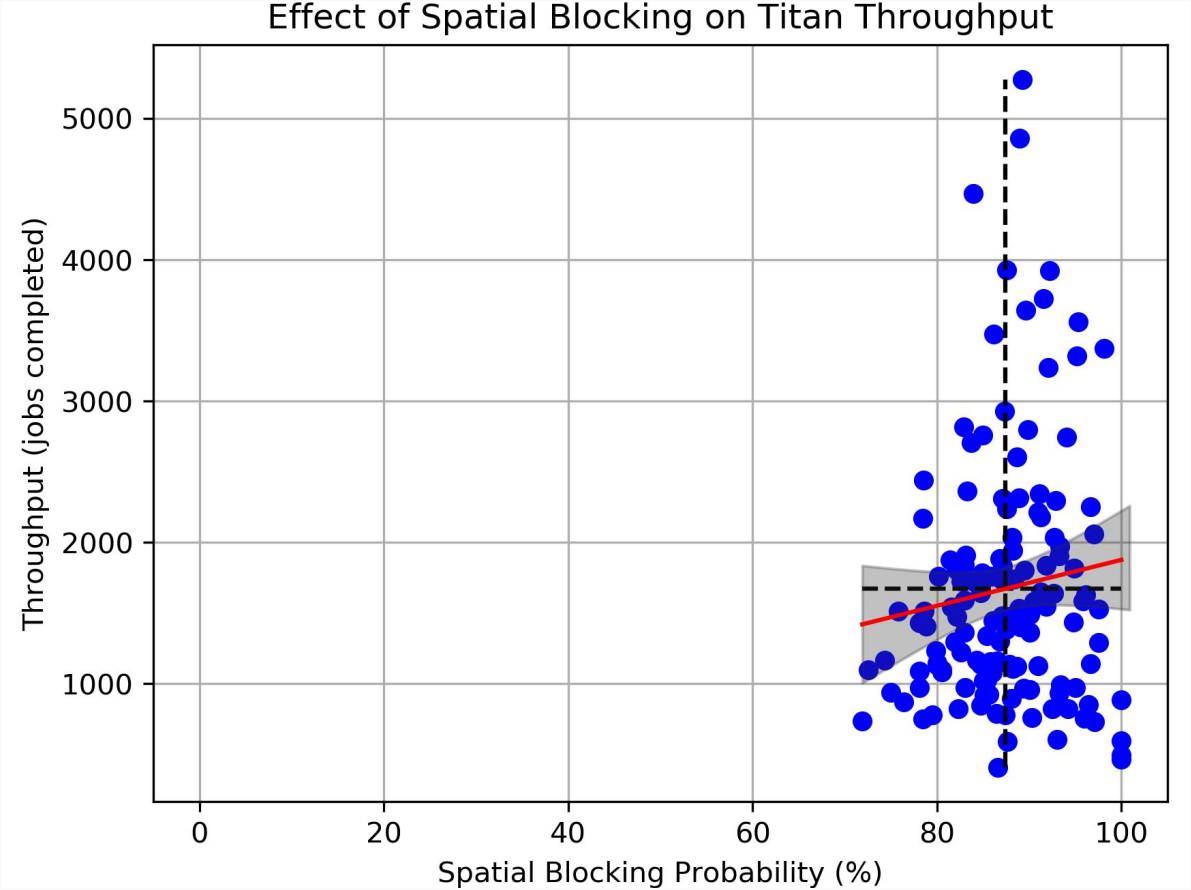

Figure 12b (shown here alone for clarity); R^2 goodness of fit: 0.0010

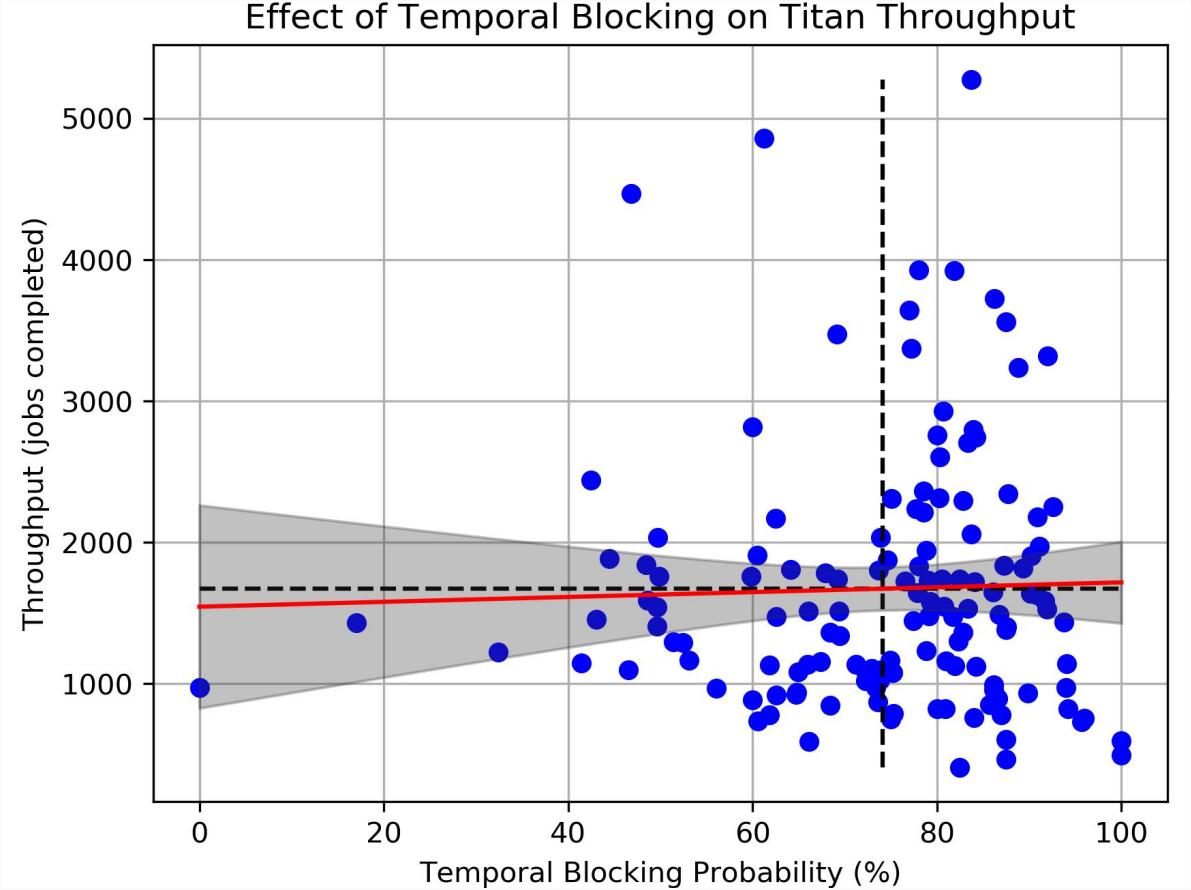

Figure 12c (shown here alone for clarity); R^2 goodness of fit: 0.0790

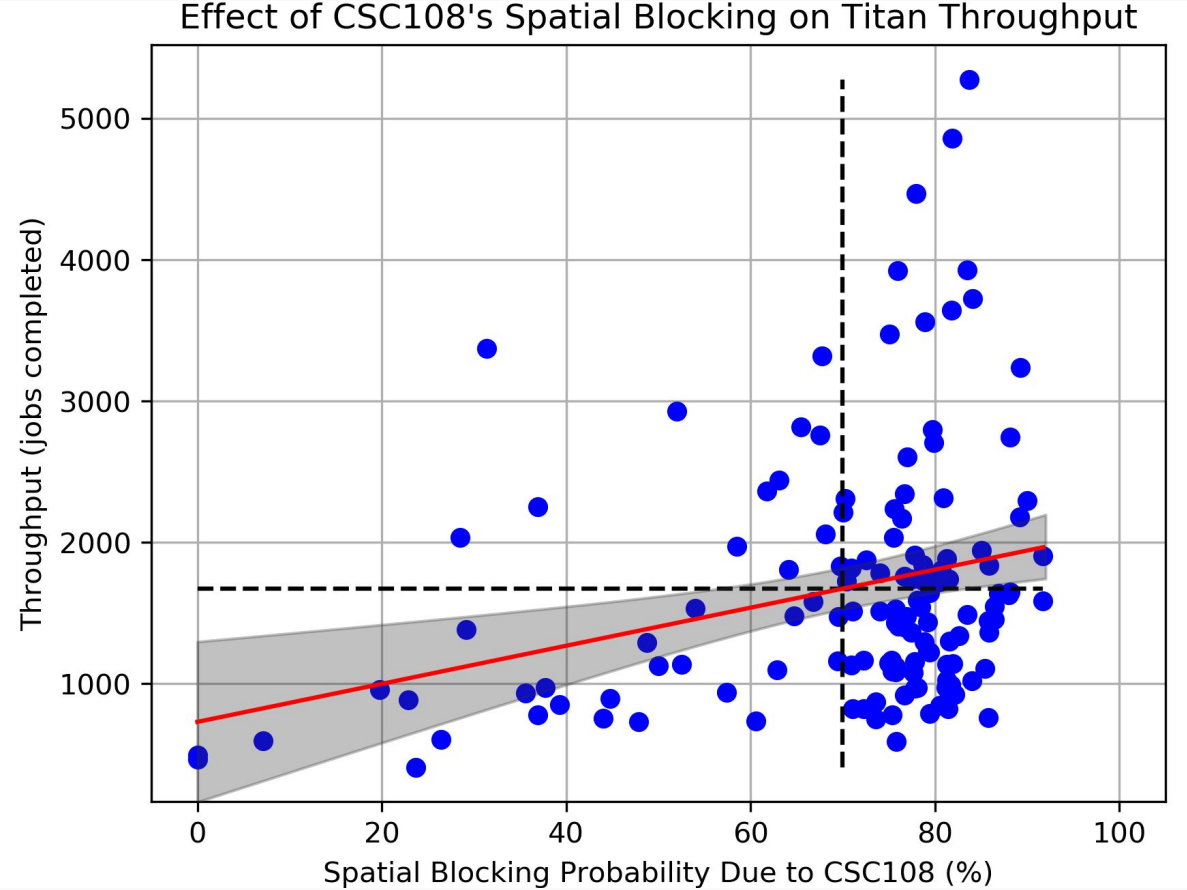

Figure 12d (shown here alone for clarity); R^2 goodness of fit: 0.0587

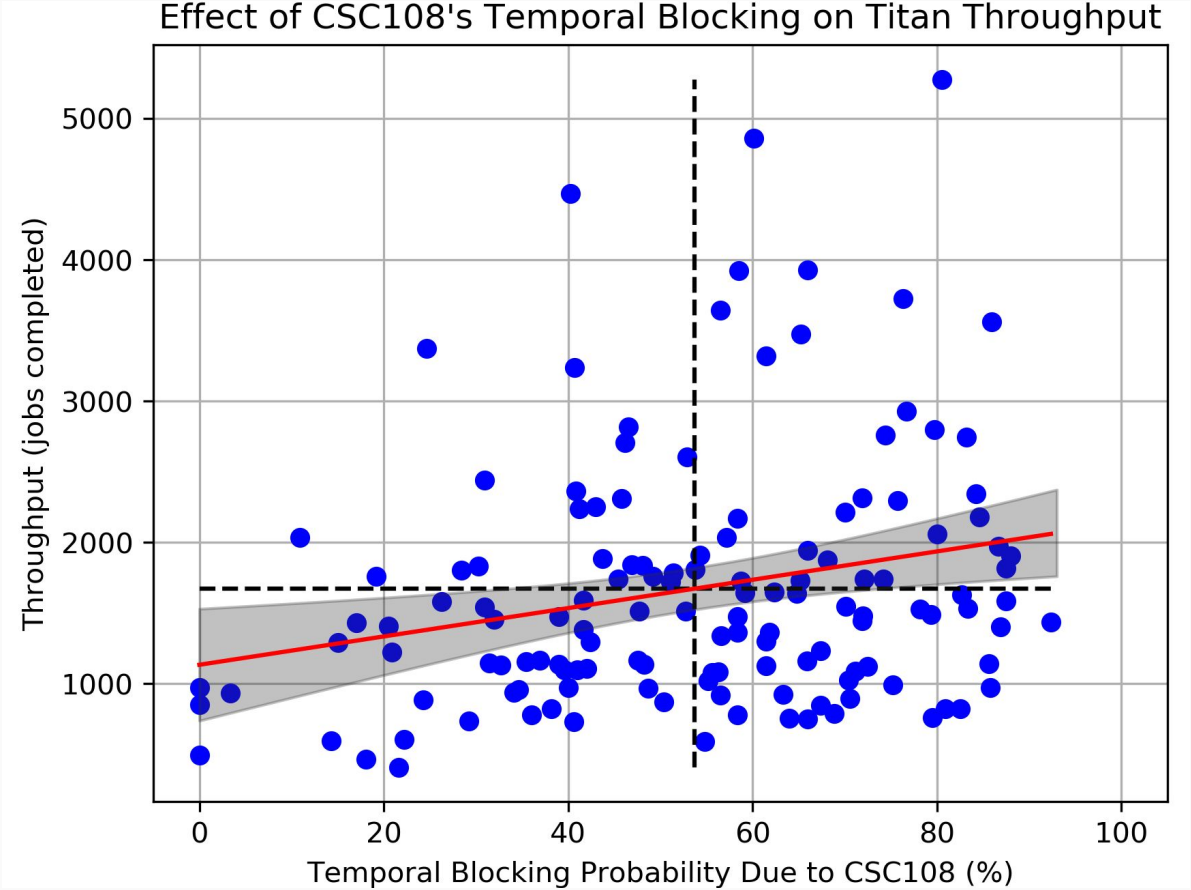

Table of model parameters et al. for throughput vs blocking relationships

Figure	Slope	Y intercept	R^2
12a	16.2402	252.3652	0.0122
12b	1.7196	1544.9669	0.0010
12c	13.4683	730.0687	0.0790
12d	10.0245	1134.0212	0.0587

Figure 13a (shown here alone for clarity); R^2 goodness of fit: 0.1543

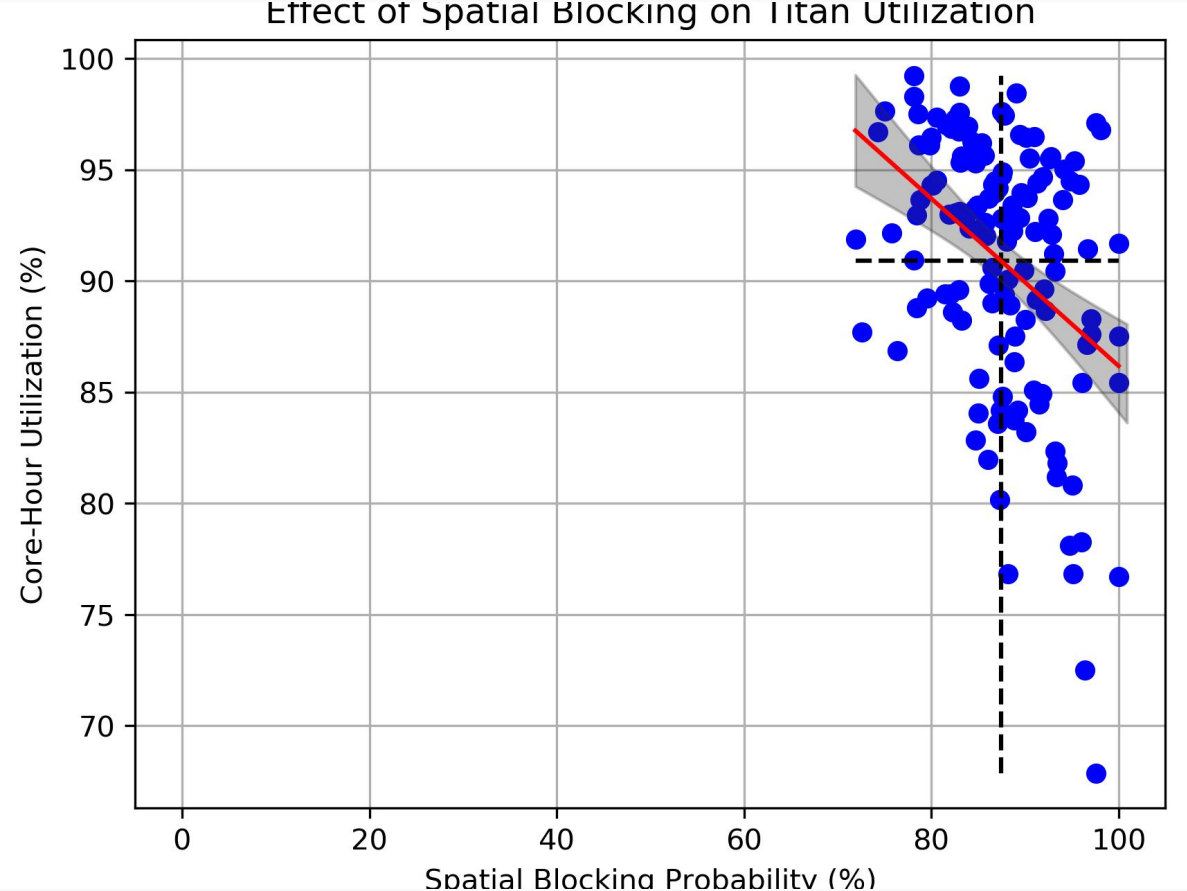

Figure 13b (shown here alone for clarity); R^2 goodness of fit: 0.2084

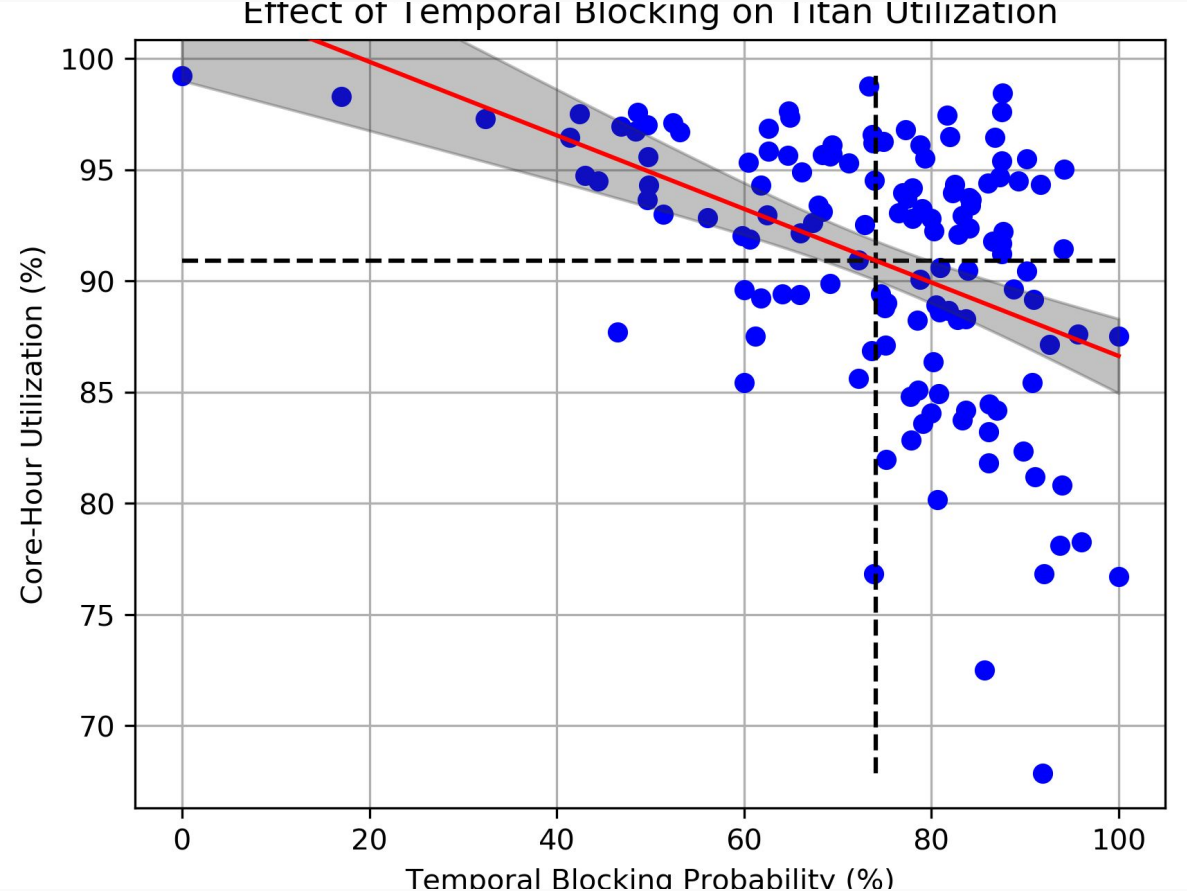

Figure 13c (shown here alone for clarity); R^2 goodness of fit: 0.0391

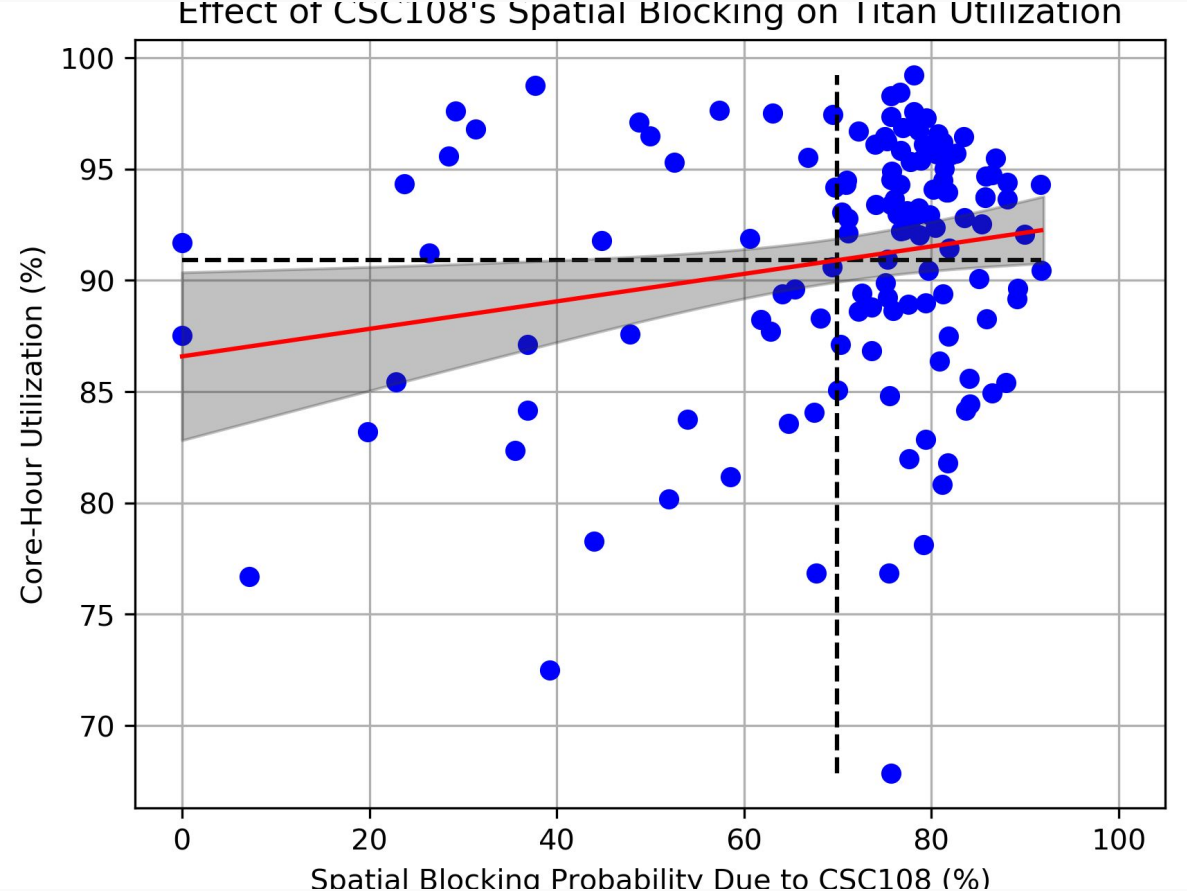

Figure 13d (shown here alone for clarity); R^2 goodness of fit: 0.0370

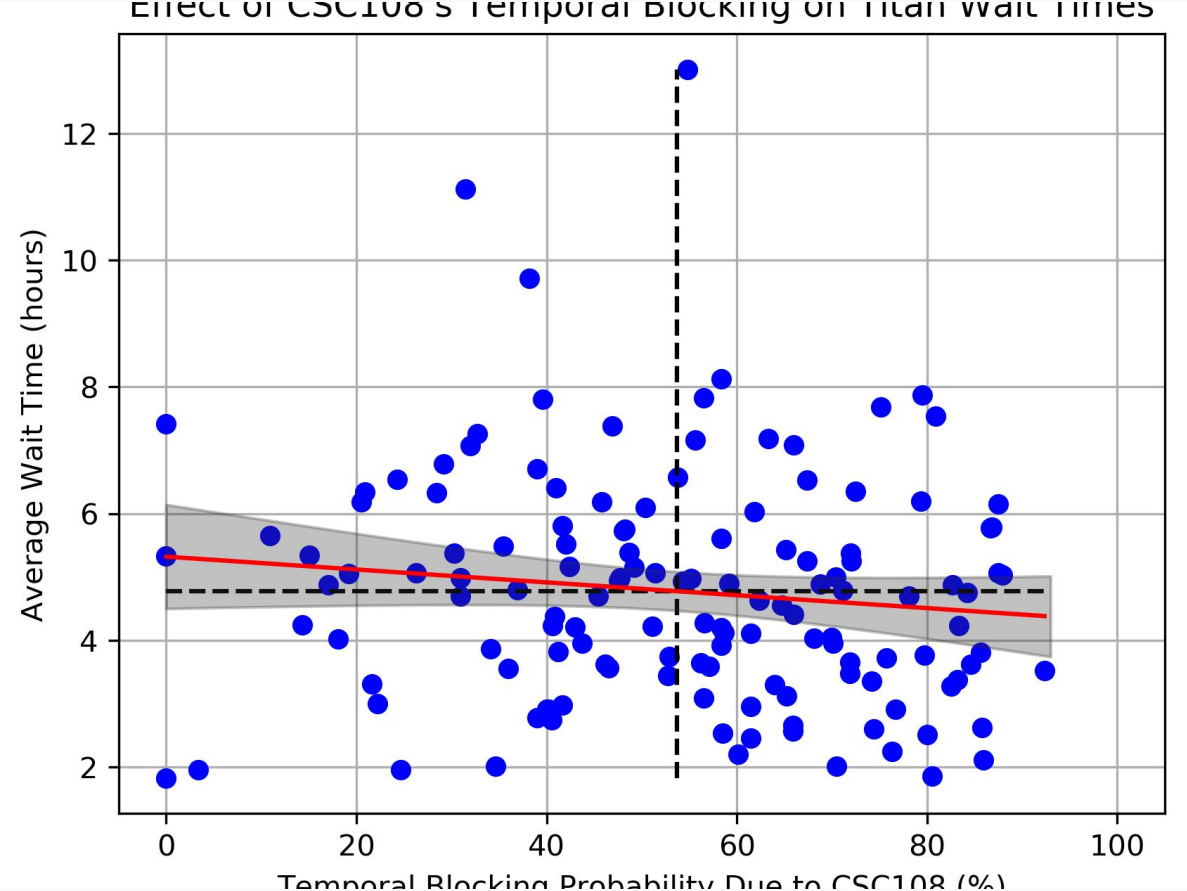

Table of model parameters et al. for utilization vs blocking relationships

Figure	Slope	Y intercept	R^2
13a	-0.3766	123.8332	0.1543
13b	-0.1654	103.1603	0.2084
13c	0.0617	86.5830	0.0391
13d	-0.0518	93.6845	0.0370

Results for blocking probability

- Wait times: only “spatial due to CSC108” increases.
- Throughput: all increase.
- Utilization: only “spatial due to CSC108” increases.
- Goodness of fit are all extremely poor, which really weakens what I am able to say regarding the results anyway.

Overall results suggest that...

- CSC108 has successfully accomplished the goal of consuming idle resources which would otherwise have gone to waste.
- CSC108 increases wait times (negative impact) but increases throughput (positive) and utilization (positive), too.

Results suggest that… (continued)

- Goodness of fit were uniformly poor; there was no relationship found anywhere where R^2 was “good”.
- “Interestingly, the inability to find simple relationships by using blocking probability suggests that users’ judging system performance by monitoring the batch queue is similarly incapable.”

Bottom line

- “In any case, the difficulty in confirming any impact may simply provide evidence that the CSC108 project has impacted Titan minimally, at least with respect to the indicators used.”
- I haven’t found anything really satisfying, one way or the other, and I’m ready to wrap this up.

Draining

- I introduced the concept of draining, and then I basically blamed it for complicating things and suggested that we study this further by finding some kind of signature to indicate draining mode vs non-draining mode.
- This might be a terrible thing to have done, which is why I’m telling you I did it. Co-authors == co-conspirators.

Questions?