Blue Waters Failure Analysis: Final Presentation

Team Gamma: Leon He and Elizabeth Dennis

Project Overview

Objective: To analyze 30,000 Blue Waters Data Points and observe:

1) Failure Type Probabilities
2) Conditional Probabilities, Failure Densities and Hazard Rates
3) Joint Distribution: Cores Used, Duration
4) Inter Arrival Distribution and CDF

Data Sample

	A	B	C	D	E	F	G	H	I	J	K
1	inter-arrival	duration	NodesNumbers	TasksPerNode	nodeType	coresNumber	exitStatus	failureType	JobScriptStart	coreHours	SuccessOrFailure
2		1602	25700	16	compute	411200	1	JOB_SCRIPT	SUCCESS	1829840	FAILURE
3	4393	2883	22629	32	compute	724128	0	JOB_EXEC_OK	SUCCESS	5799058.4	SUCCESS
4	522442	9024	4800	32	compute	153600	-10	EXCEED_CLOCK_	FAIL	3850240	FAILURE
5	72643	88	22629	32	compute	724128	1	JOB_SCRIPT	SUCCESS	177009.067	FAILURE
6	13663	2690	22629	32	compute	724128	0	JOB_EXEC_OK	SUCCESS	5410845.33	SUCCESS
7	171814	2	24576	16	compute	393216	1	JOB_SCRIPT	SUCCESS	2184.53333	FAILURE
8	93	3	12288	32	compute	393216	1	JOB_SCRIPT	SUCCESS	3276.8	FAILURE
9	15272	9821	22528	32	compute	720896	0	JOB_EXEC_OK	SUCCESS	19666443.4	SUCCESS
10	113	1019	8192	32	compute	262144	143	TERM_SIGNAL	SUCCESS	742013.156	FAILURE
11	69680	531	8256	32	compute	264192	0	JOB_EXEC_OK	SUCCESS	389683.2	SUCCESS
12	17378	3	4608	32	compute	147456	0	JOB_EXEC_OK	SUCCESS	1228.8	SUCCESS
13	47563	35	8192	32	compute	262144	137	FATAL_ERROR	SUCCESS	25486.2222	FAILURE
14	679	3620	4096	32	compute	131072	-10	EXCEED_CLOCK_	FAIL	1318001.78	FAILURE
15	13263	36	3000	32	compute	96000	1	JOB_SCRIPT	SUCCESS	9600	FAILURE
16	13227	39	5000	32	compute	160000	1	JOB_SCRIPT	SUCCESS	17333.3333	FAILURE
17	8449	8984	8192	32	compute	262144	0	JOB_EXEC_OK	SUCCESS	6541949.16	SUCCESS
18	4682	119	32	32	compute	1024	0	JOB_EXEC_OK	SUCCESS	338.488889	SUCCESS
19	621	225	864	2	compute	1728	0	JOB_EXEC_OK	SUCCESS	1080	SUCCESS
20	1315	6630	1350	16	compute	21600	-10	EXCEED_CLOCK_	FAIL	397800	FAILURE
21	9335	257	1375	4	compute	5500	0	JOB_EXEC_OK	SUCCESS	3926.38889	SUCCESS
22	782	827	22600	1	compute	22600	0	JOB_EXEC_OK	SUCCESS	51917.2222	SUCCESS
23	664	1258	22600	1	compute	22600	0	JOB_EXEC_OK	SUCCESS	78974.4444	SUCCESS

*Received from Catello Di Martino, Ph-D

Data Analysis Approach and Concepts used from Class

- Analyzed Data Using R because it is better for statistical analysis than Matlab
  - Correlation, Variance, Table functions
- Conditional Probability of Failure
  - Took the number of failures in an interval and divided by the total number of tasks within an interval
- Failure Density and Hazard Rate
  - Used methods that we learned in class to create tables like the ones in lecture

Data Analysis Approach and Concepts Used From the Class

- Failure Probabilities and Exit Code Polarity
  - Looked at the polarity of exit codes to determine Job Script Success and Failure
- Joint Distribution: Cores Used, Duration
  - Each variable was divided into intervals, a table was made with likelihoods and variables on respective x and y axes to determine the joint distribution.

Failure Probabilities

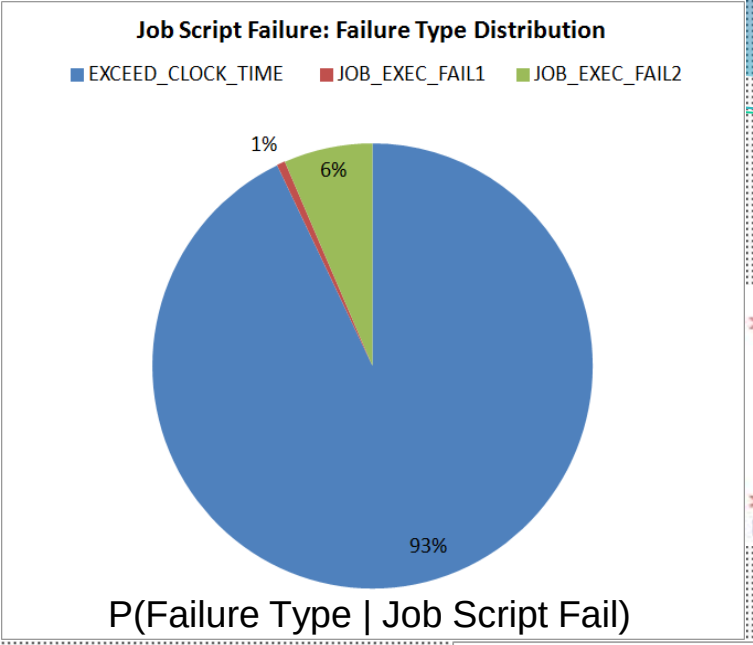

P(Failure Type | Job Script Fail)

```
> summary(Job_Script_Fail$failureType)
      CANNOT_EXEC  COMMAND_NOT_FOUND EXCEED_CLOCK_TIME        FATAL_ERROR  ILLEGAL_INSTRUCT
                0                 0              1364                  0                 0
    JOB_EXEC_FAIL1    JOB_EXEC_FAIL2       JOB_EXEC_OK        JOB_SCRIPT      LIMIT_EXCEED
                9                96                  0                  0                 0
       MISC_ERROR         SEG_FAULT       TERM_SIGNAL
                0                 0                  0
> nrow(Job_Script_Fail)
[1] 1469
```

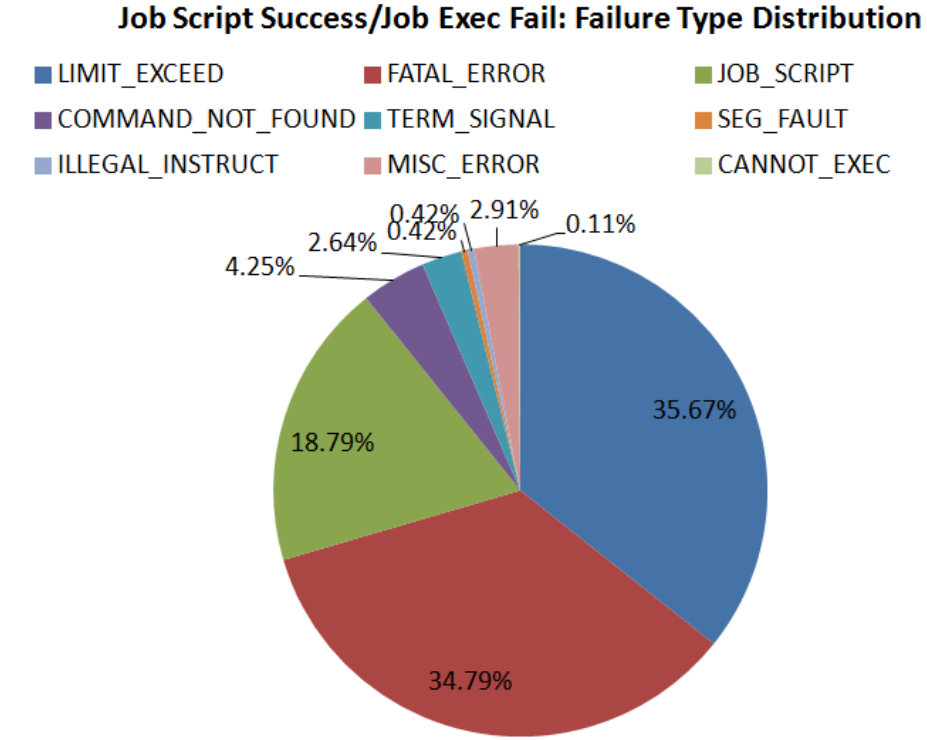

P(Failure Type | Job Script Success/Job Execute Fail)

Conditional Failure Probability, Failure Density, Hazard Rate

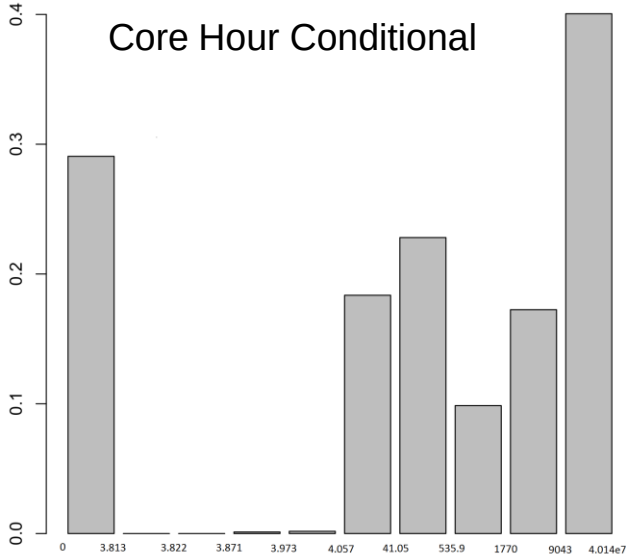

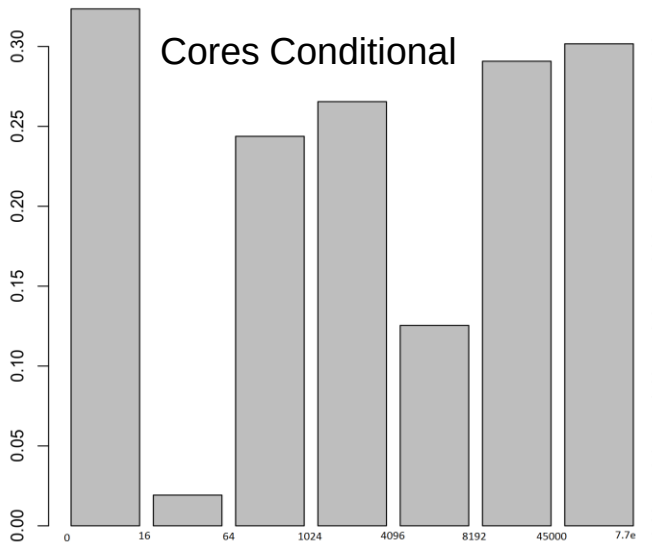

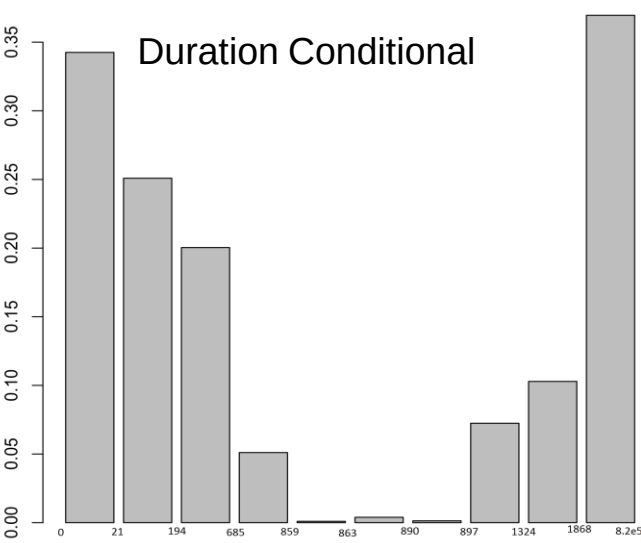

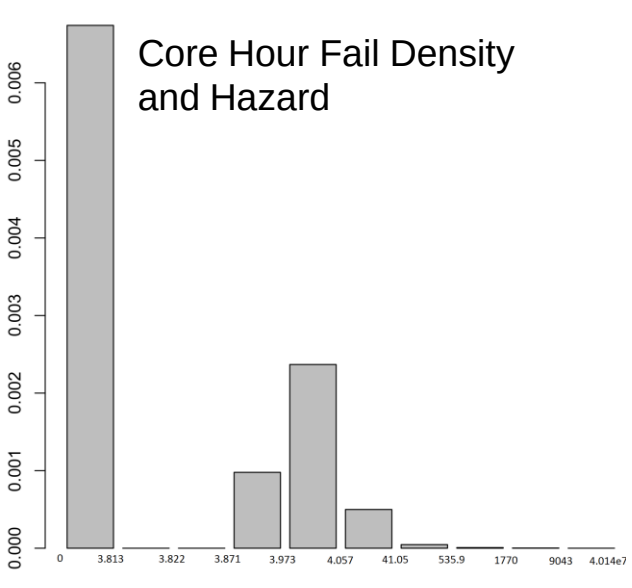

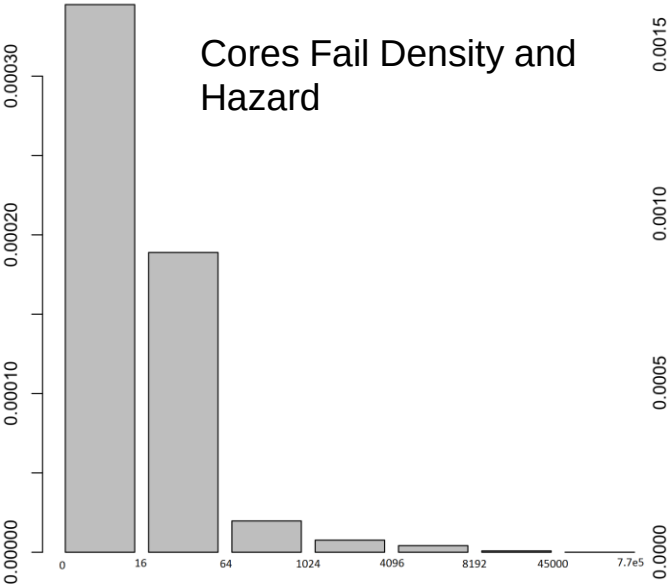

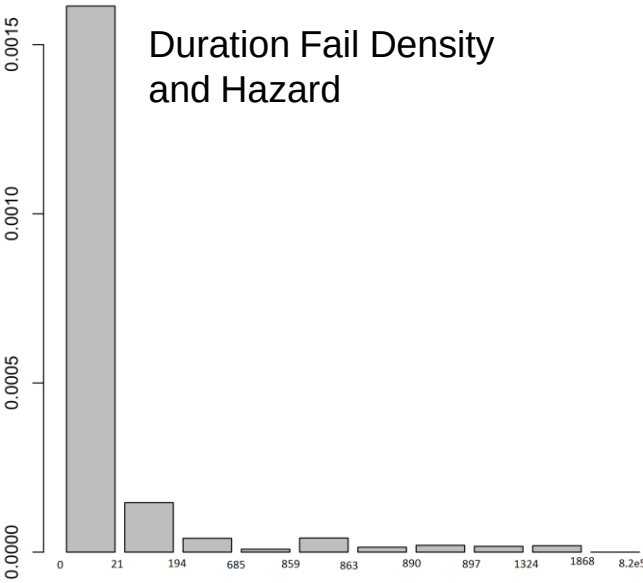

Duration, Cores and Core Hour Data

```
> cor(bluewaters$coresNumber, bluewaters$duration)
[1] 0.03065871
```

```
> givenDuration
       t1       t2 Failure Success Total P(F | t1 < t < t2)  failDensity       hazard
1     0.0     21.0    1017    1952  2969       0.3425395756 1.614286e-03 1.614286e-03
2    21.0    194.0     759    2266  3025       0.2509090909 1.462428e-04 1.513744e-04
3   194.0    685.0     602    2402  3004       0.2003994674 4.086897e-05 4.344066e-05
4   685.0    859.0      45     836   881       0.0510783201 8.620690e-06 9.362852e-06
5   859.0    863.0       5    5055  5060       0.0009881423 4.166667e-05 4.532763e-05
6   863.0    890.4      12    3049  3061       0.0039202875 1.459854e-05 1.588409e-05
7   890.4    897.0       4    2954  2958       0.0013522650 2.020202e-05 2.199059e-05
8   897.0   1324.0     220    2818  3038       0.0724160632 1.717408e-05 1.869729e-05
9  1324.0   1868.0     309    2694  3003       0.1028971029 1.893382e-05 2.077900e-05
10 1868.0 823535.0    1109    1892  3001       0.3695434855 4.498984e-08 4.993877e-08
> givenCores
      t1     t2 Failure Success Total P(F | t1 < t < t2)  failDensity       hazard
1     -1     16     176     368   544         0.32352941 3.450980e-04 3.450980e-04
5     16     64     272   13872 14144         0.01923077 1.888889e-04 1.900036e-04
6     64   1024     568    1762  2330         0.24377682 1.972222e-05 2.002121e-05
7   1024   4096     703    1945  2648         0.26548338 7.628038e-06 7.895430e-06
8   4096   8192     512    3570  4082         0.12542871 4.166667e-06 4.419929e-06
9   8192  45120     936    2283  3219         0.29077353 8.448873e-07 9.127668e-07
10 45120 773568     915    2118  3033         0.30168150 4.186984e-08 4.681158e-08
> givenCoreTime
            t1          t2 Failure Success Total P(F | t1 < t < t2)  failDensity       hazard
1   -0.2413889 3.813333e+00     820    2001  2821        0.290677065 6.741111e-03 6.741111e-03
2    3.8133333 3.822222e+00       0    4940  4940        0.000000000 0.000000e+00 0.000000e+00
3    3.8222222 3.871111e+00       0    1265  1265        0.000000000 0.000000e+00 0.000000e+00
4    3.8711111 3.973333e+00       3    2497  2500        0.001200000 9.782609e-04 1.005751e-03
5    3.9733333 4.057778e+00       6    3451  3457        0.001735609 2.368421e-03 2.435227e-03
6    4.0577778 4.104533e+01     554    2463  3017        0.183626119 4.992670e-04 5.134555e-04
7   41.0453333 5.358933e+02     684    2315  2999        0.228076025 4.607475e-05 4.830145e-05
8  535.8933333 1.770382e+03     296    2707  3003        0.098568099 7.992512e-06 8.583945e-06
9 1770.3822222 9.042800e+03     517    2480  2997        0.172505839 2.369684e-06 2.572295e-06
10 9042.8000000 4.013611e+07    1202    1799  3001        0.400533156 9.984947e-10 1.104530e-09
```

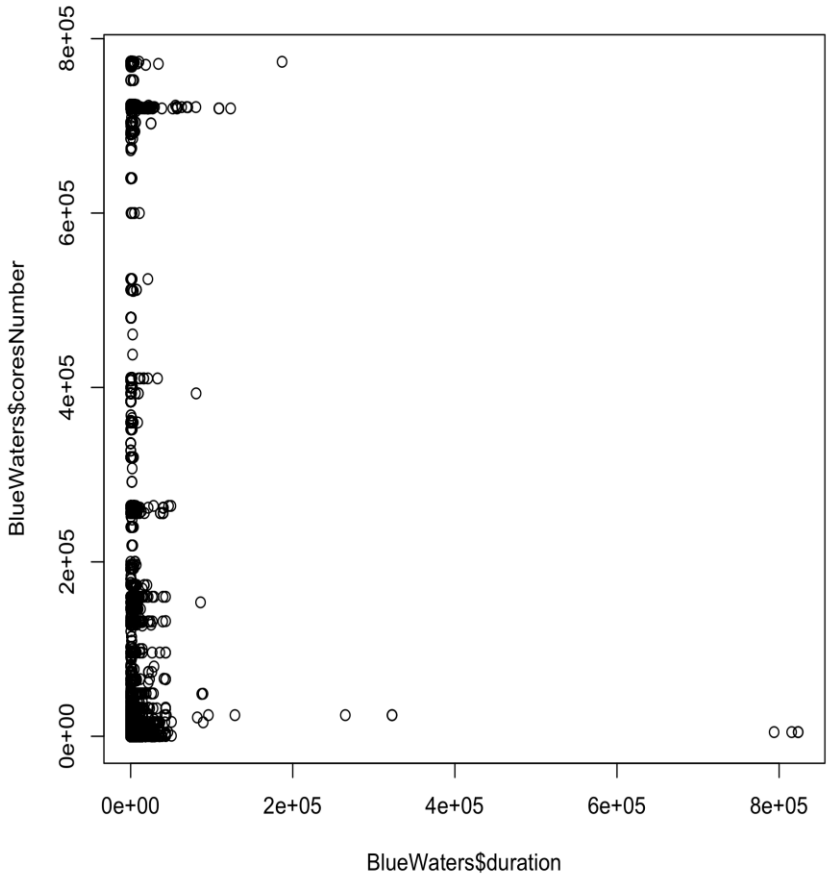

Joint Distribution: Number of Cores Used vs. Duration

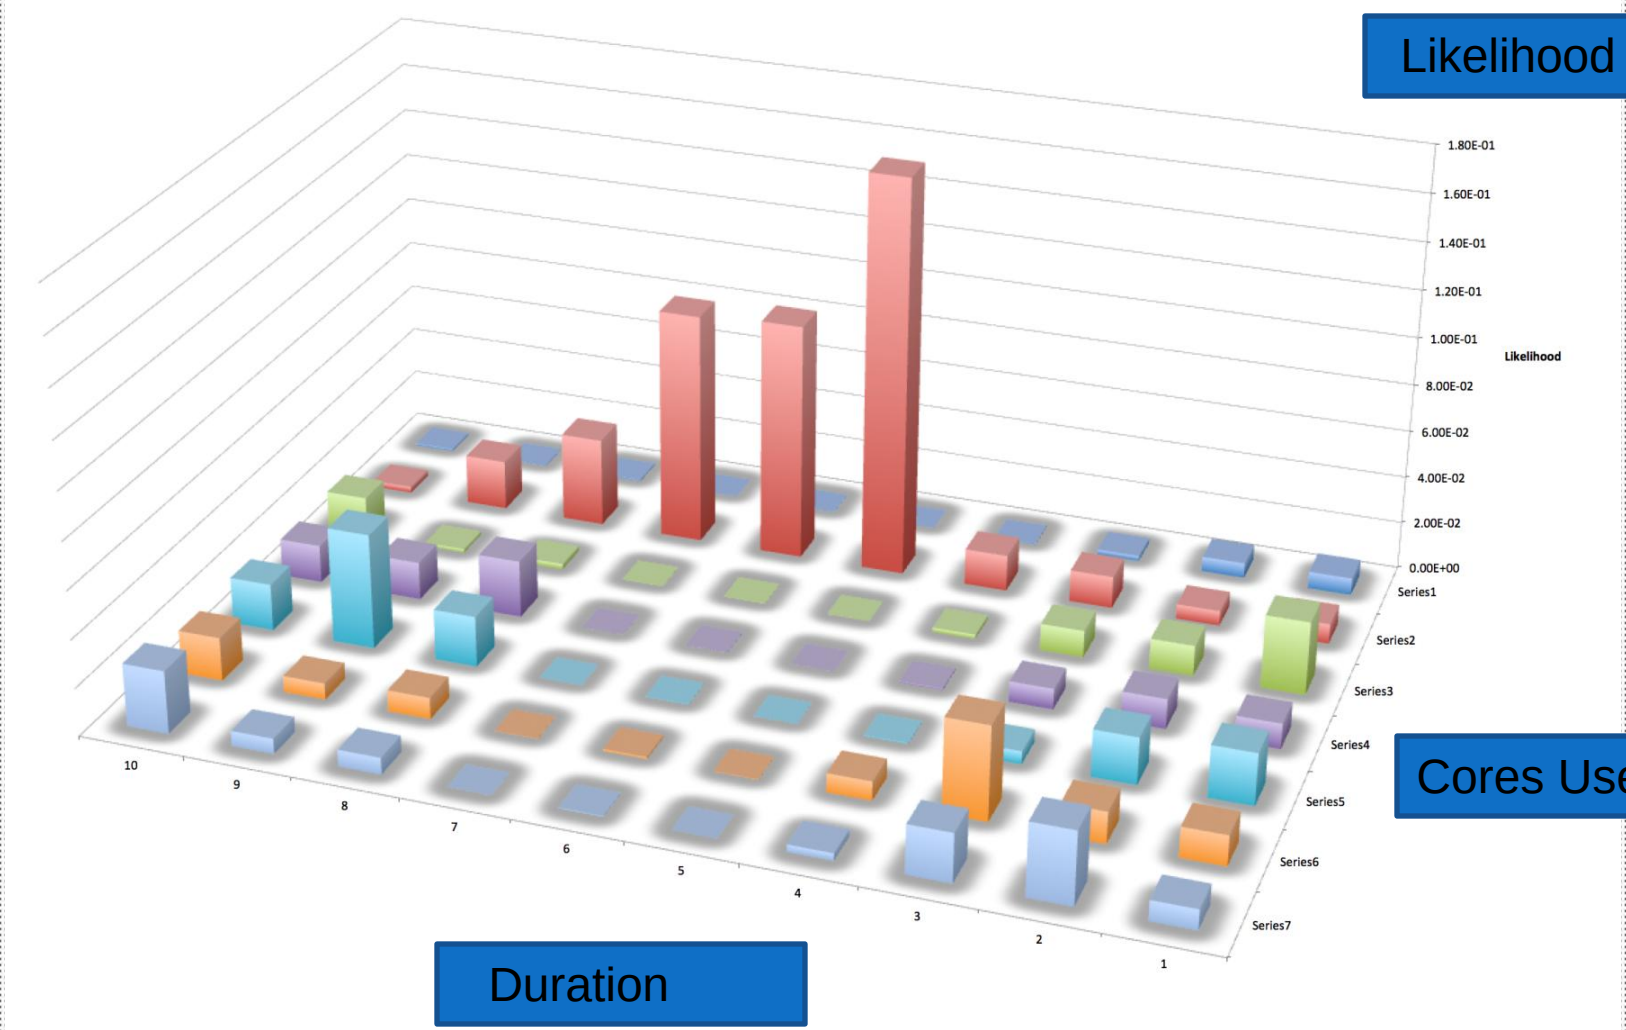

Interarrival Distribution and Problems with Data

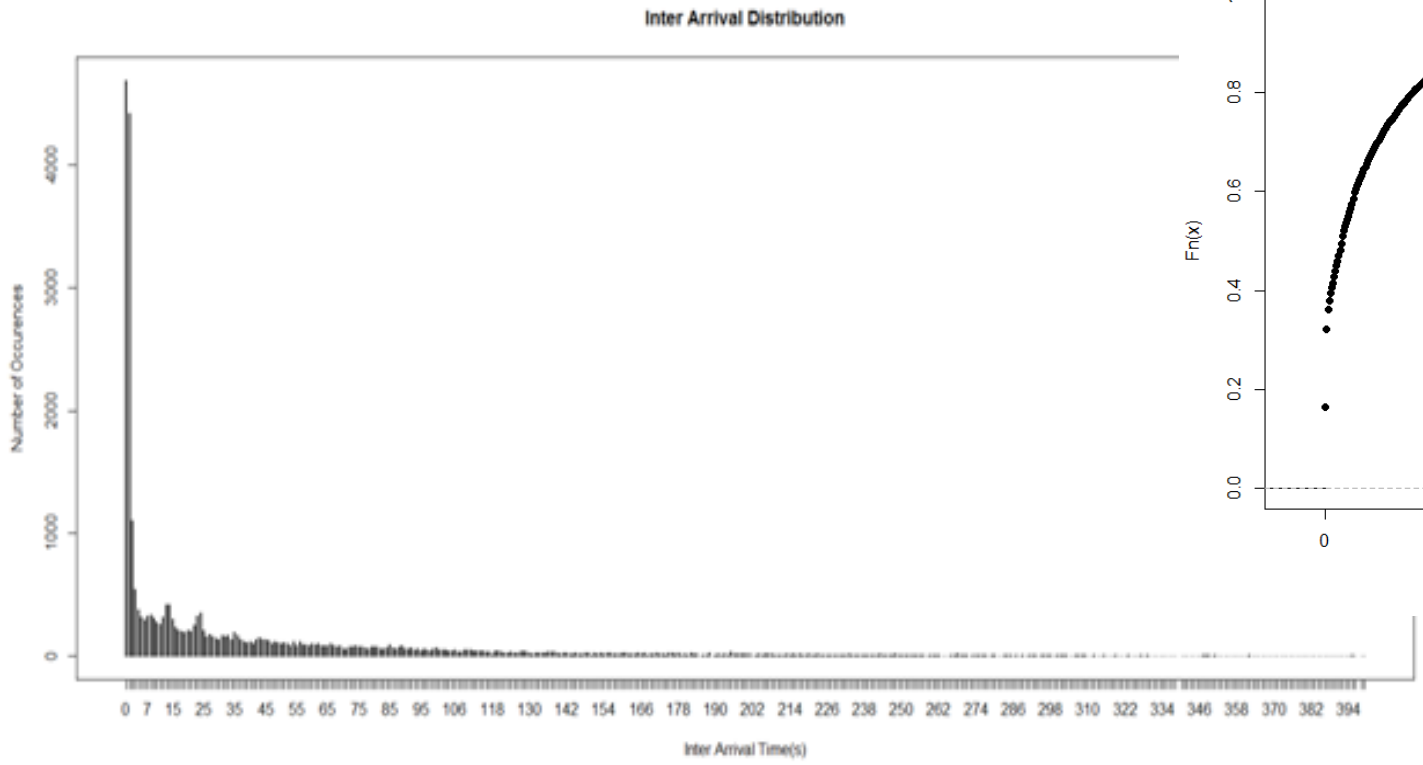

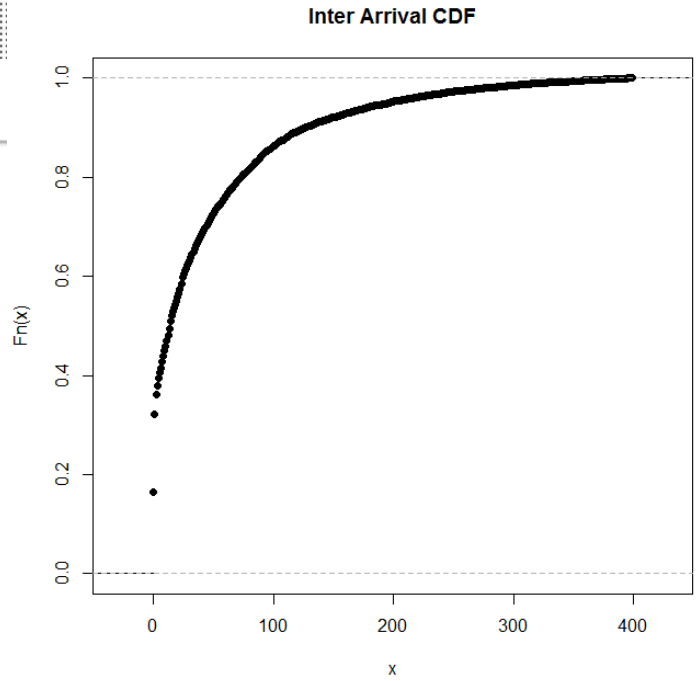


```
> cor(completebluewaters$usedVmem_kb, completebluewaters$usedMem_kb)
[1] 0.9841249
```

```
> cor(completebluewaters$coreHours, completebluewaters$usedVmem_kb)
[1] 0.004801629
> cor(completebluewaters$coreHours, completebluewaters$usedMem_kb)
[1] 0.00679183
```

Conclusions

- Two types of failures: Job Script Success Failures and Job Script Fail Failures
- The conditional probabilities show that as any of the variables increases, so does the probability of failure. The failure densities and hazard rates show that as the variables increase, the density and hazard decrease.
- Joint Distribution Graph shows the most likely combination of Cores (16-64 cores) and Duration (859-863 seconds)

Feelings about the Class

- More in class projects or discussion
- Class expected students to have previous knowledge of R language

Likes:

- Insight on Reliability Systems
- Hearing about the relation to what we did in class to real world examples

Questions?