

Research Methods in computer science

Fall 2013

Lecture 12

Omprakash Gnawali

October 8, 2013

Agenda

Research Conference Updates

Data Aggregation

HW6 feedback

HW8

Sparklines

“small, high resolution graphics
embedded in a context of words,
numbers, images”

"data-intense, design-simple,
word-sized graphics”

-- Edward Tufte

Example sparklines in small multiple

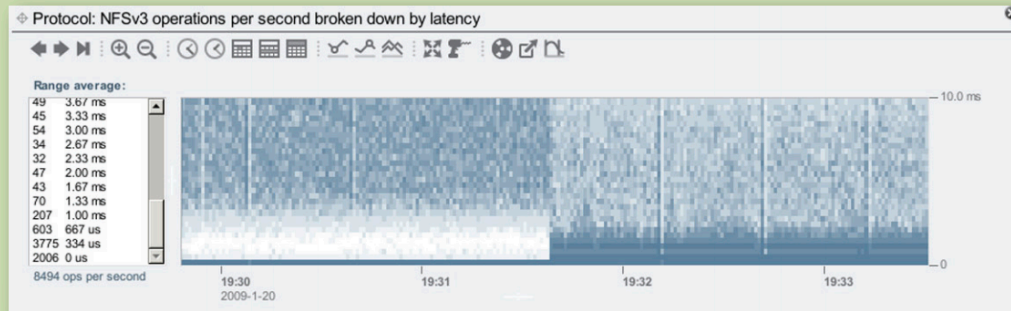
Index Day		Value	Change
Dow Jones		10765.45	-32.82
S&P 500		1256.92	-8.10

Example sparklines of the [Dow Jones Industrial Average](#) and [S&P 500](#) on February 7, 2006

wikipedia

FIGURE 1

NFS Latency When Enabling SSD-based Cache Devices



Latency distribution of an NFS read workload and the effect on NFS latency when using an additional layer of flash-memory-based cache

This and next four heatmaps by Brendan Gregg, 2010

FIGURE 2

Synchronous Writes to a Striped Pool of Disks

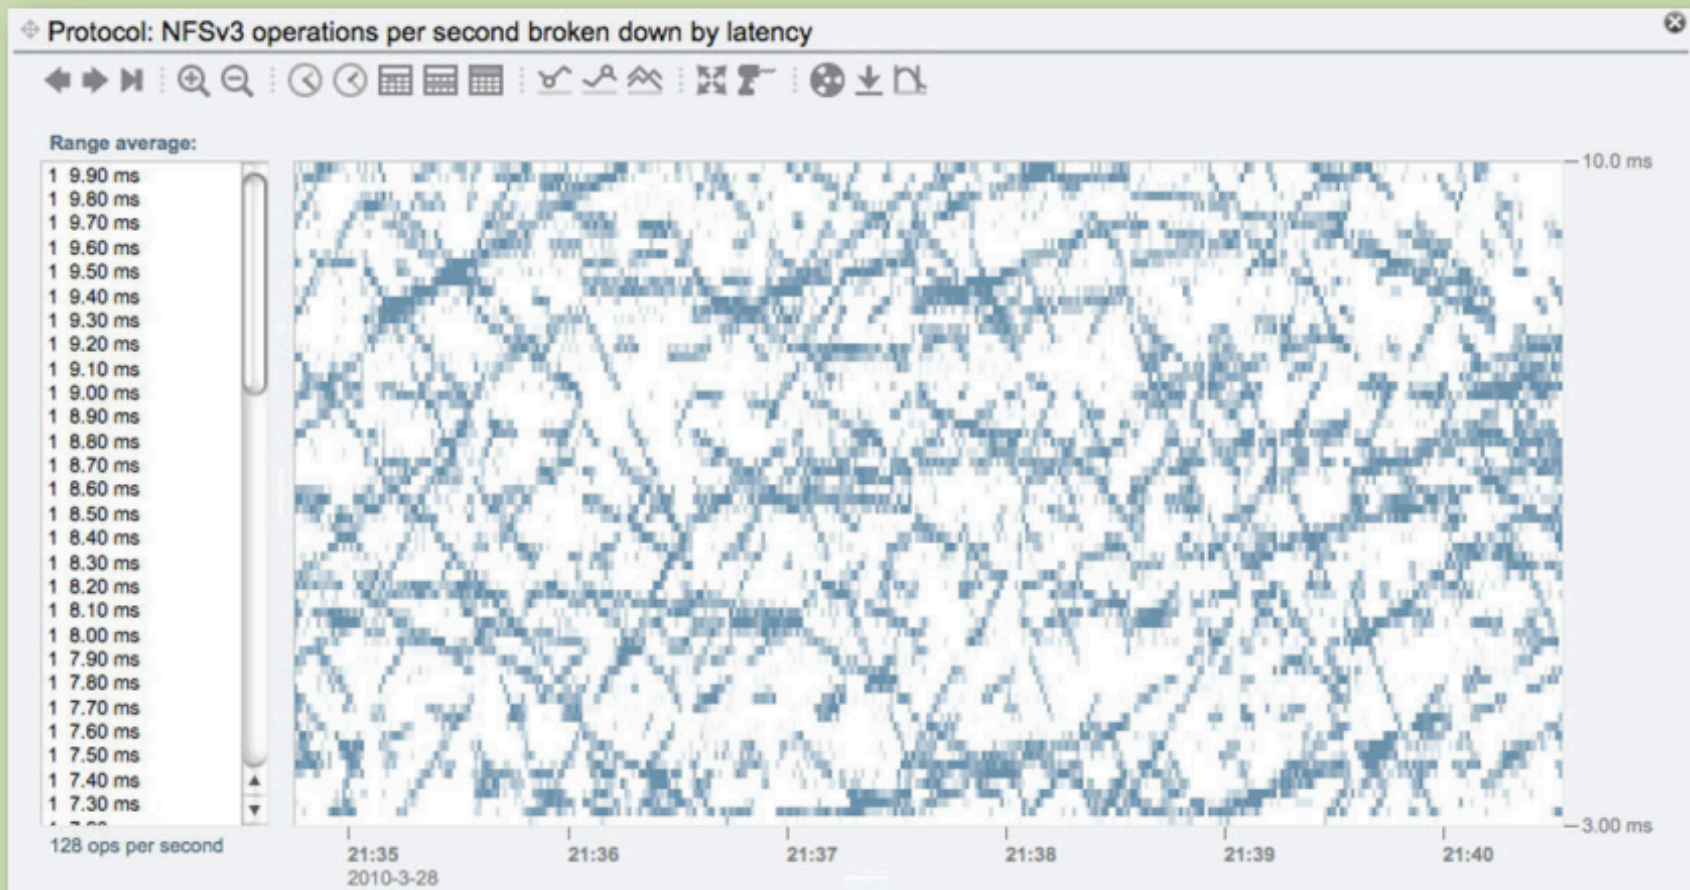


FIGURE 3

Single Disk Latency from a Striped Pool of Disks

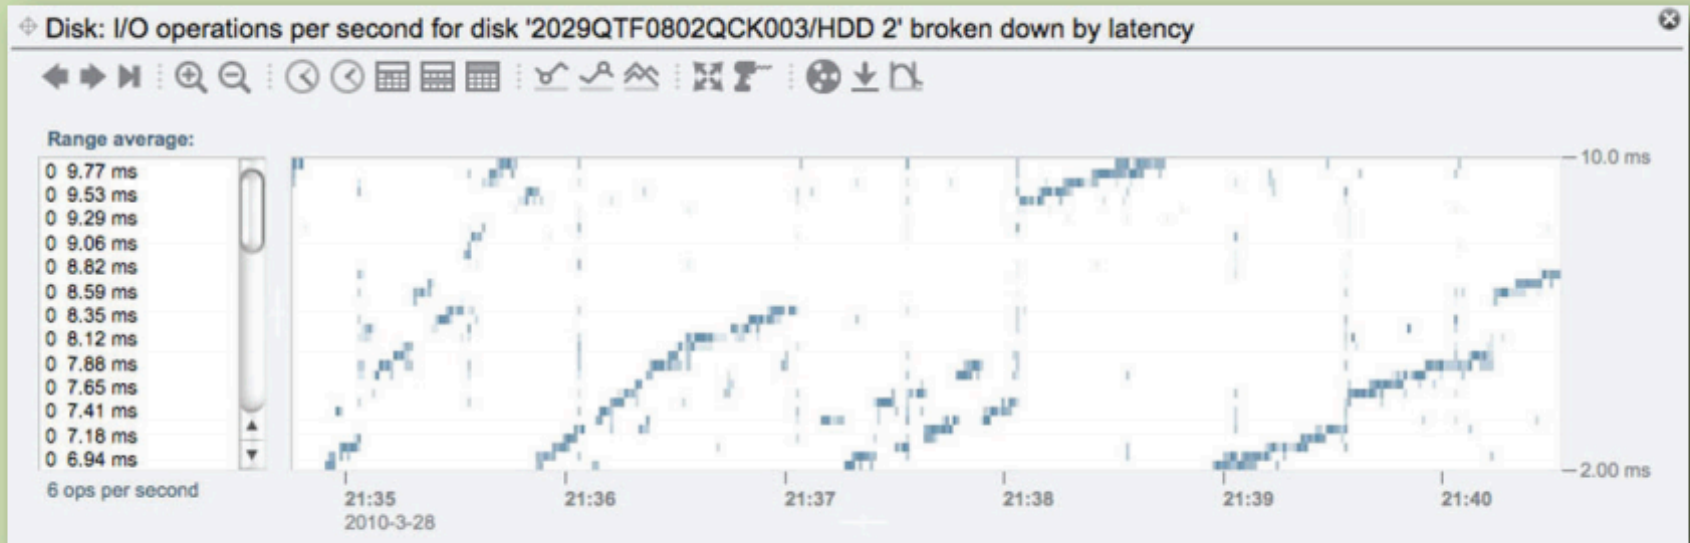
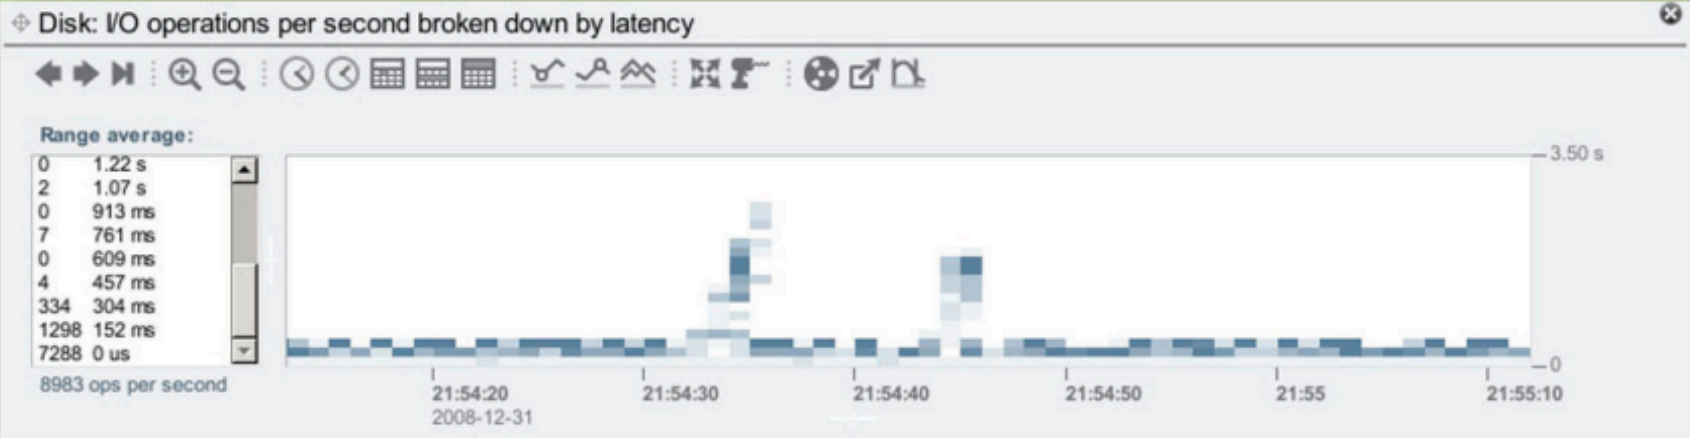


FIGURE 9

High-latency I/O



Outliers?

Data Aggregation


overall statistics

#node total_sent uniq_rcv success_rate total_rcv repeated_rcv (frac) repeated_rcvthl (frac) minseq maxseq


3 814 814 1.000 815 1 0.001, 0.001, 3, 816
4 826 826 1.000 831 5 0.006, 0.006, 3, 828
5 826 826 1.000 837 11 0.013, 0.012, 3, 828
6 819 819 1.000 821 2 0.002, 0.002, 2, 820
9 821 821 1.000 829 8 0.010, 0.010, 2, 822
10 816 816 1.000 819 3 0.004, 0.004, 2, 817
13 827 827 1.000 830 3 0.004, 0.004, 3, 829
16 811 804 0.991 814 10 0.012, 0.012, 4, 814
17 813 813 1.000 823 10 0.012, 0.012, 3, 815
18 836 836 1.000 840 4 0.005, 0.005, 2, 837
19 816 816 1.000 817 1 0.001, 0.001, 3, 818
20 829 829 1.000 834 5 0.006, 0.006, 3, 831
21 825 825 1.000 828 3 0.004, 0.004, 3, 827
22 793 596 0.752 601 5 0.008, 0.007, 17, 809
23 816 816 1.000 821 5 0.006, 0.006, 3, 818
24 838 838 1.000 841 3 0.004, 0.004, 4, 841
25 813 813 1.000 817 4 0.005, 0.005, 3, 815
26 810 810 1.000 812 2 0.002, 0.002, 2, 811
27 818 818 1.000 823 5 0.006, 0.006, 4, 821
28 834 834 1.000 836 2 0.002, 0.002, 3, 836
31 815 815 1.000 817 2 0.002, 0.002, 3, 817
32 797 44 0.055 44 0 0.000, 0.000, 14, 810
39 816 816 1.000 821 5 0.006, 0.006, 4, 819
40 822 804 0.978 886 82 0.093, 0.093, 4, 825
41 829 829 1.000 833 4 0.005, 0.005, 3, 831
42 825 825 1.000 828 3 0.004, 0.004, 4, 828
45 832 832 1.000 835 3 0.004, 0.004, 4, 835
46 816 815 0.999 822 7 0.009, 0.009, 3, 818
47 812 812 1.000 813 1 0.001, 0.001, 3, 814
48 818 818 1.000 820 2 0.002, 0.002, 2, 819
49 815 815 1.000 818 3 0.004, 0.004, 4, 818
50 831 831 1.000 836 5 0.006, 0.006, 4, 834
53 818 817 0.999 822 5 0.006, 0.006, 3, 820
55 833 833 1.000 837 4 0.005, 0.005, 3, 835
56 822 822 1.000 823 1 0.001, 0.001, 3, 824
57 826 826 1.000 829 3 0.004, 0.004, 3, 828
58 818 818 1.000 825 7 0.008, 0.008, 4, 821


...

How to analyze/visualize the reliability (4th column)





What is wrong with this plot?

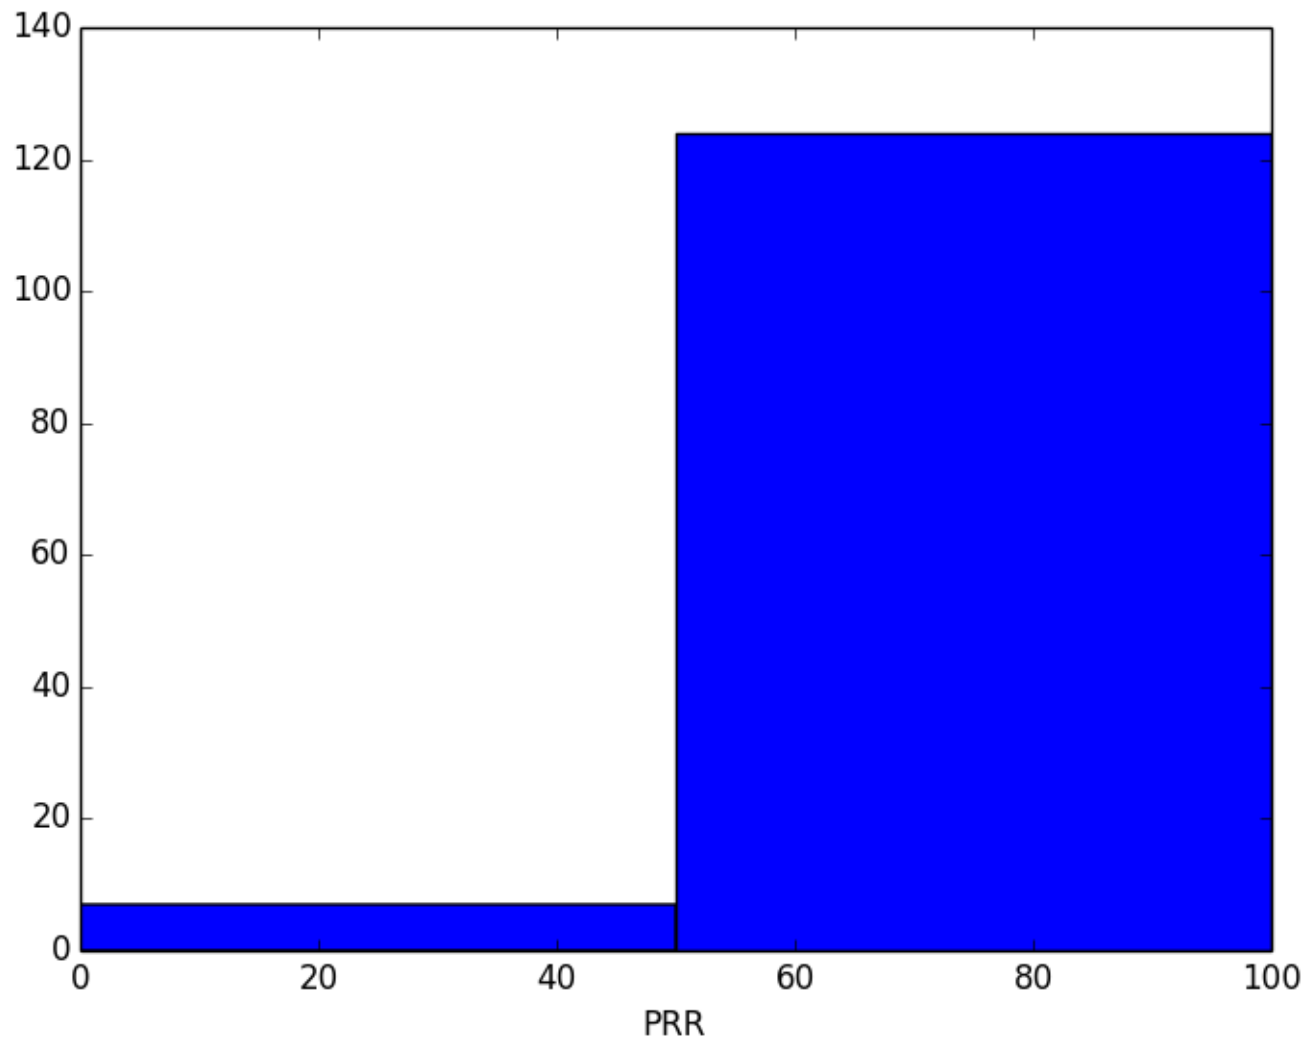






What is the right bucket size?

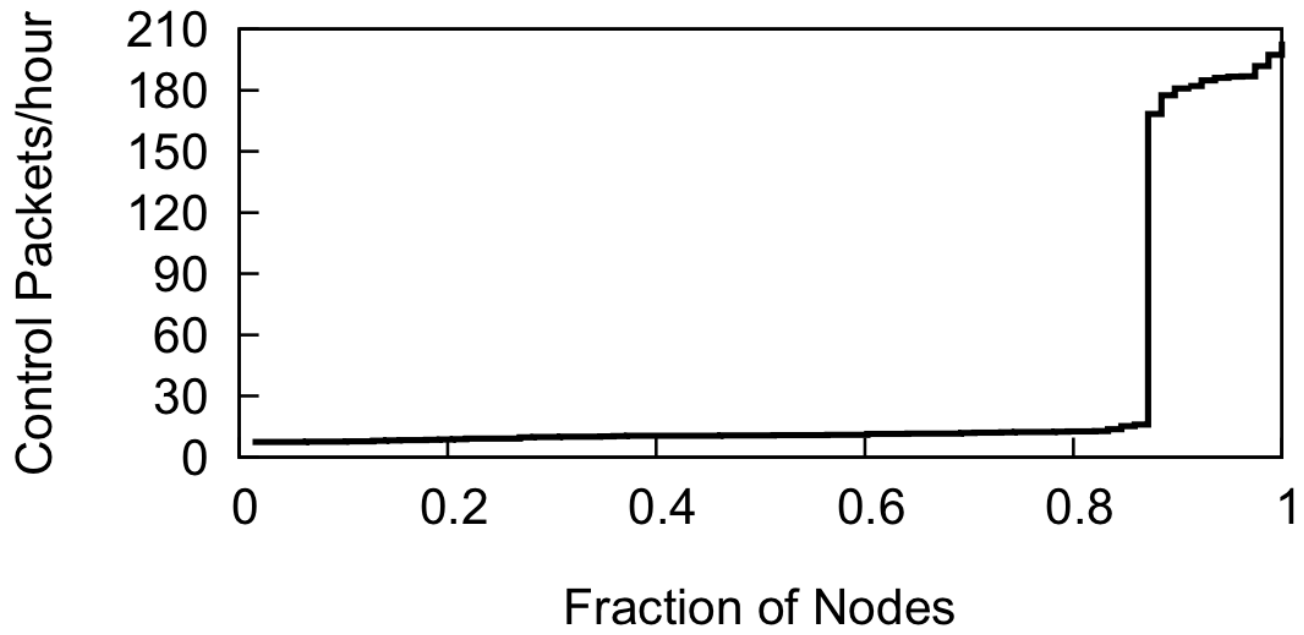


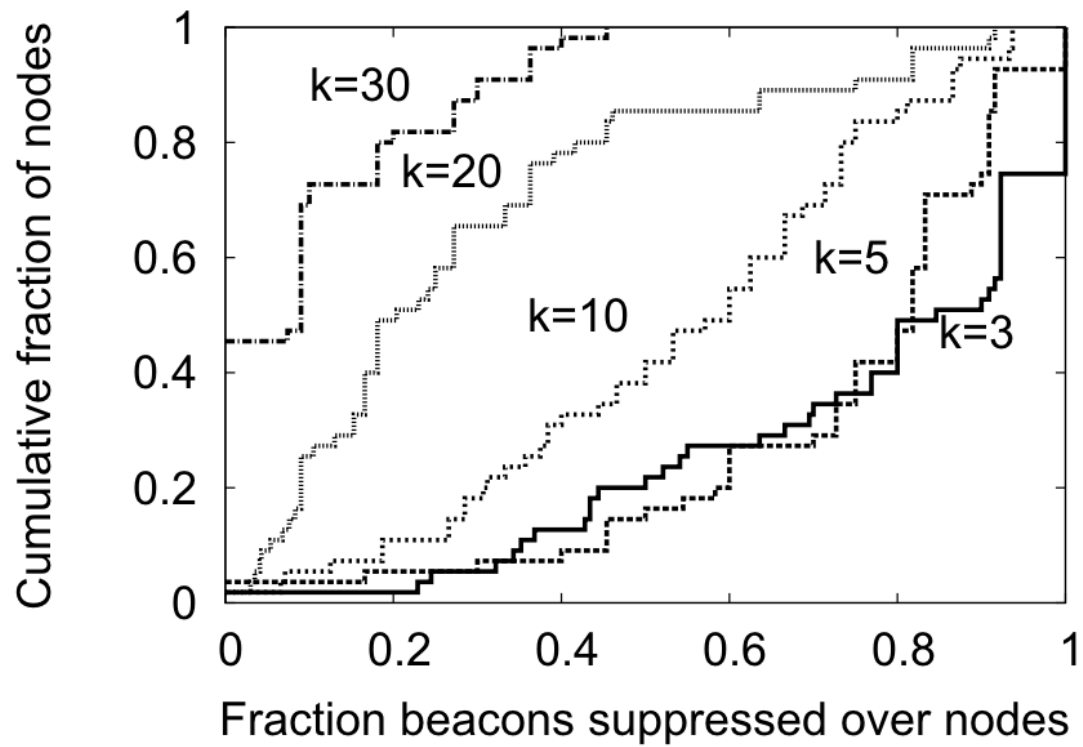












HW8

- Snapshot of your 6-page research paper. The paper must be entirely written by you.