

A decorative header at the top of the slide features four overlapping spheres. From left to right, they are light green, light blue, light red, and light yellow. The spheres are partially cut off by the top edge of the slide.

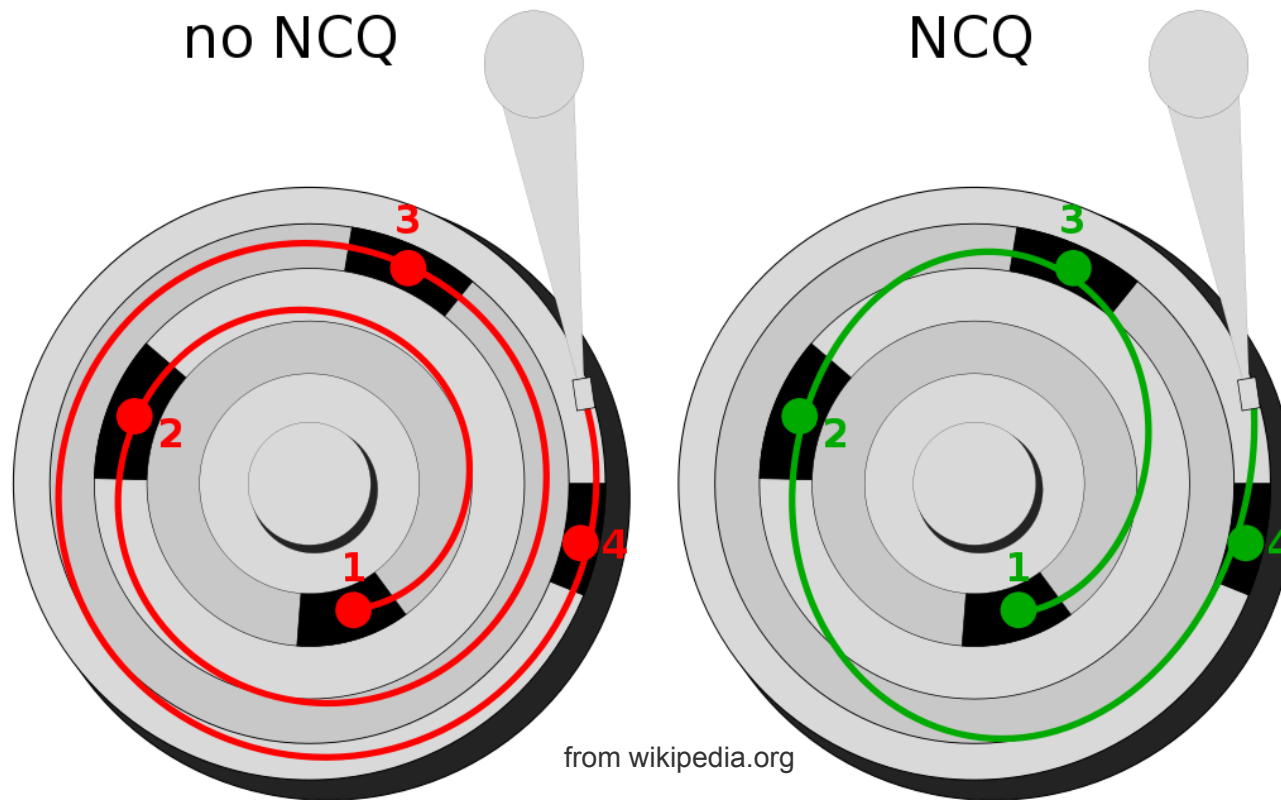
Native Command Queuing

Gwendal Grignou



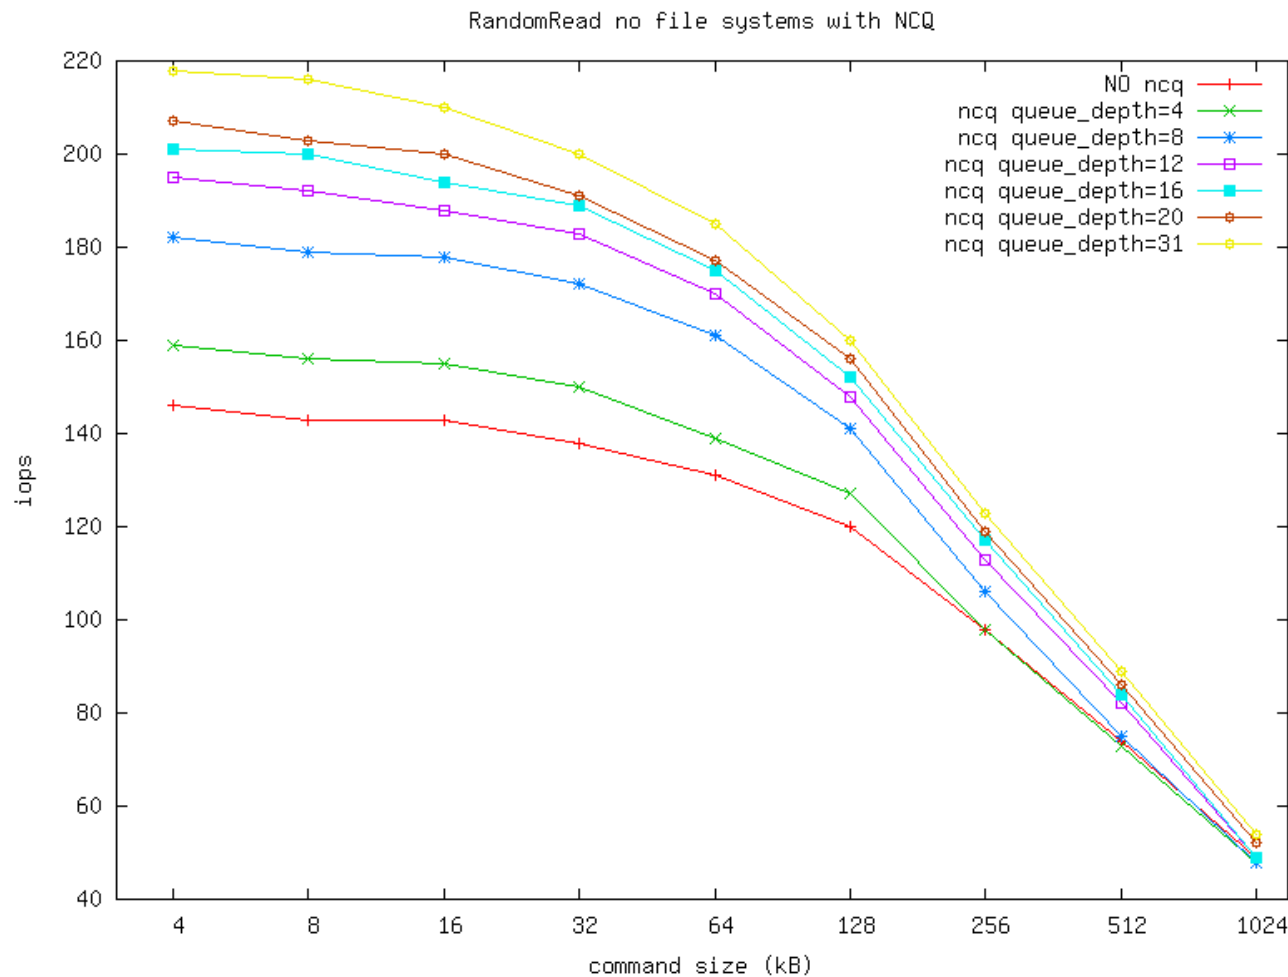
Introduction

- Improving latency by completing more IO per platter rotation
 - Reordering IO allowed
 - wider actuator arm movements



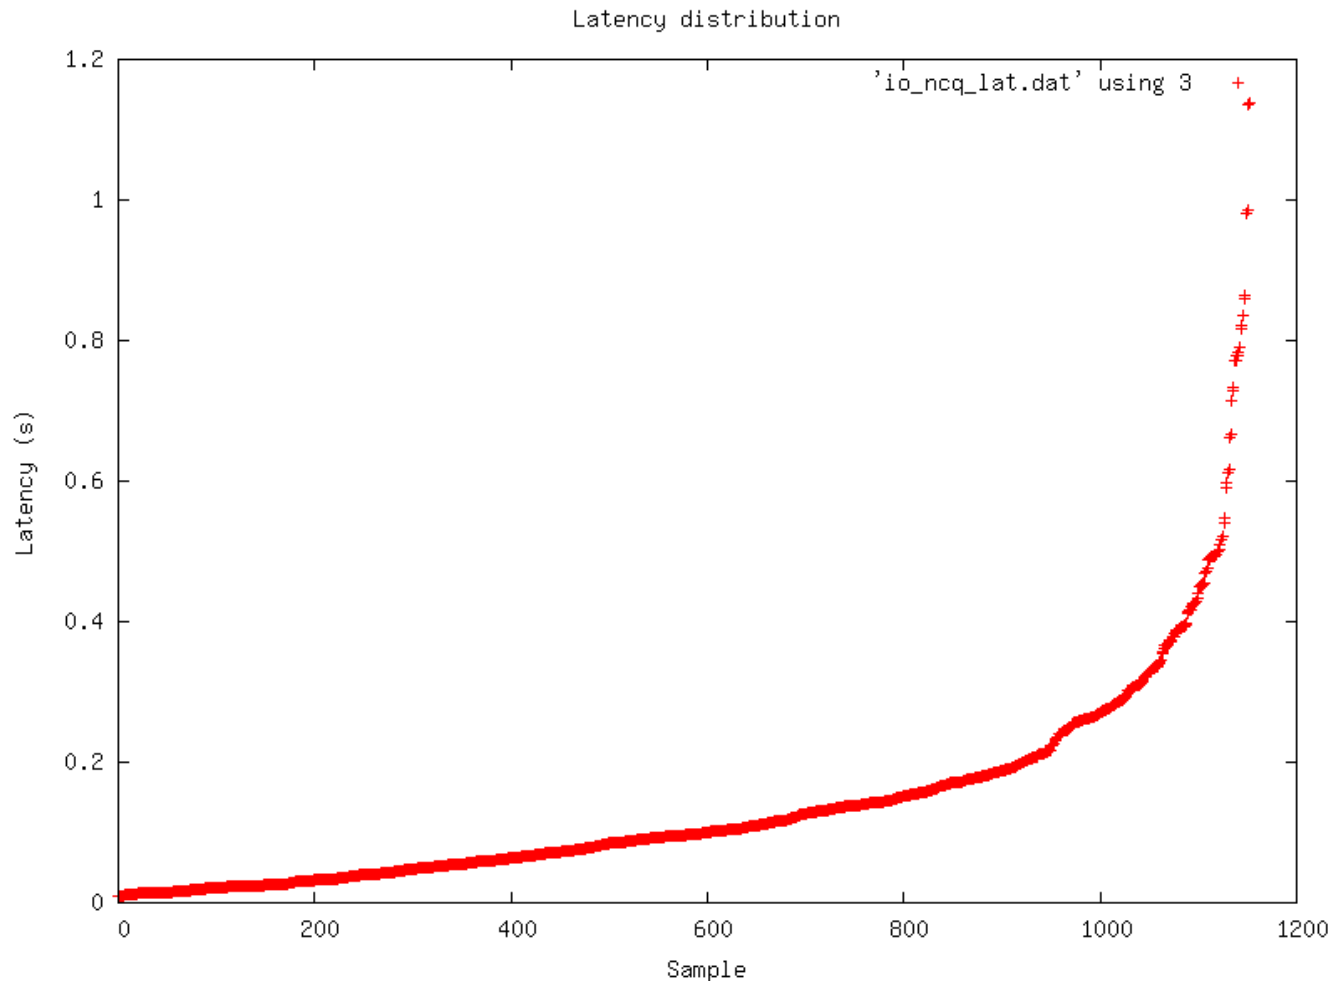
Current Status

- It works!
 - +50% IOps in some case



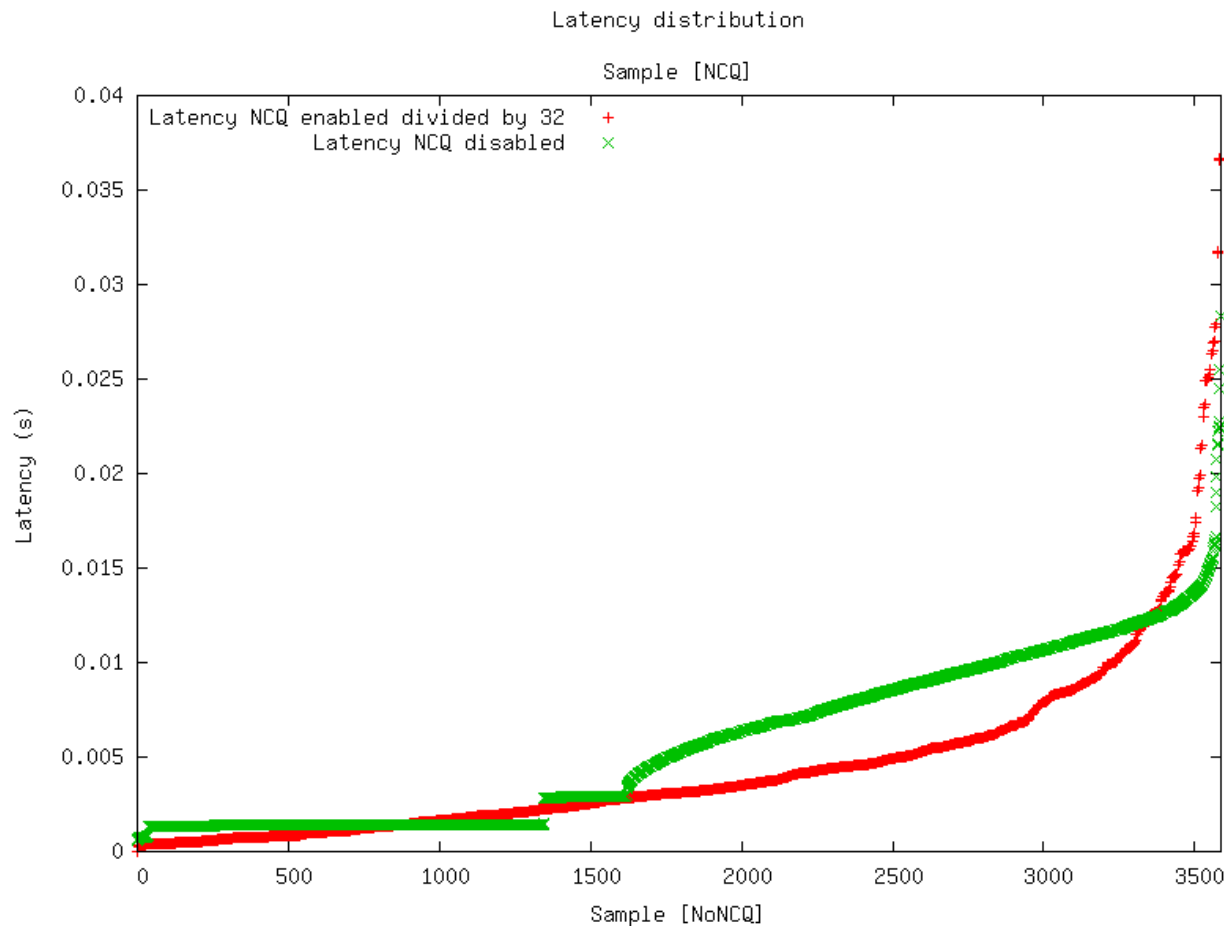
Current Status

- Degraded performance on real life application
 - Some commands take 1-2 seconds to be completed.



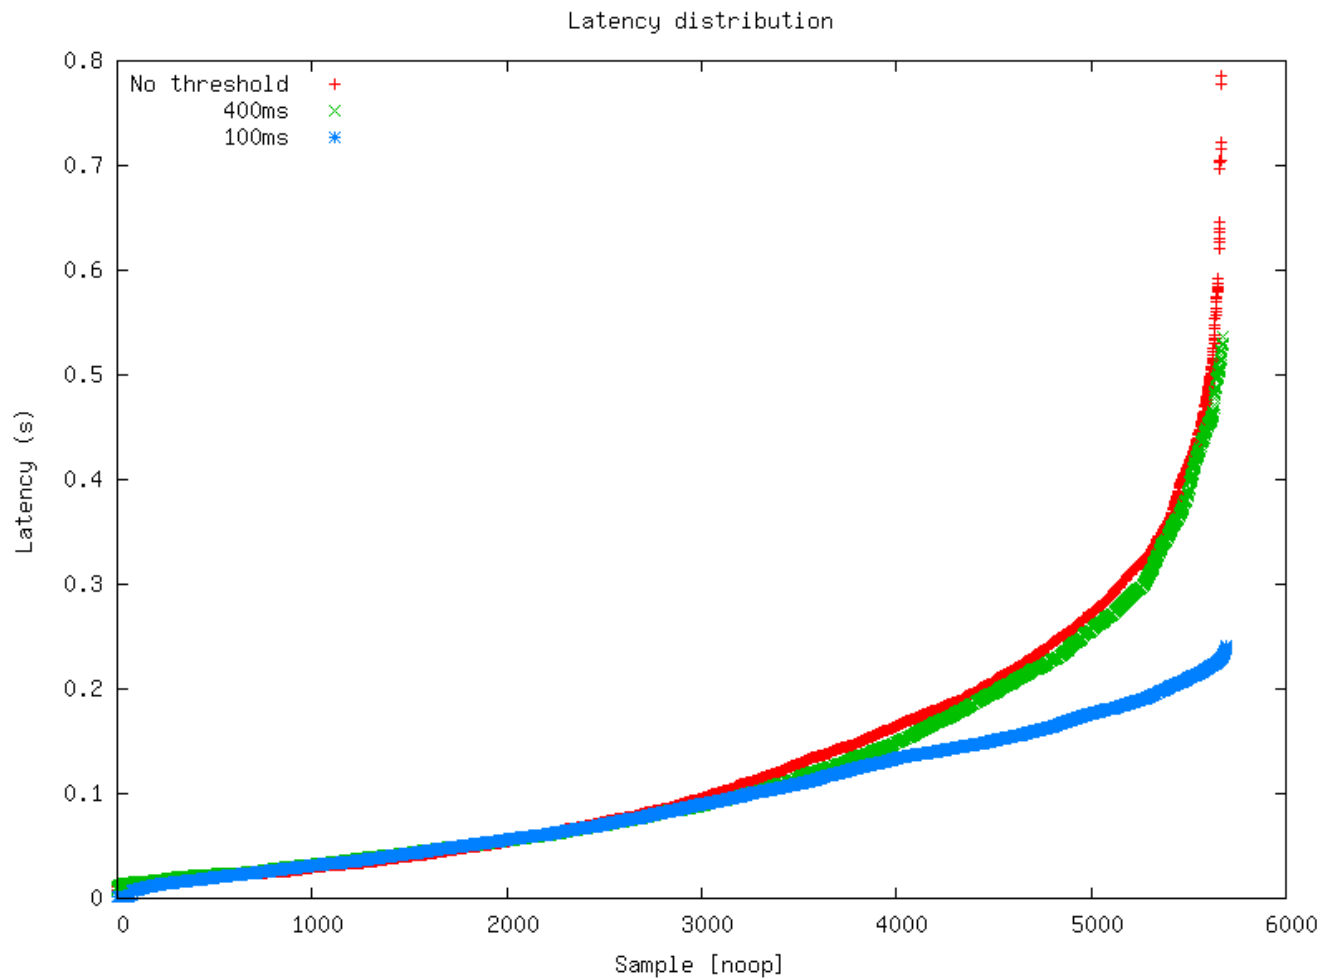
Current Status

- Degraded performance on real life application
 - earlier tests report better Median, worst Mean



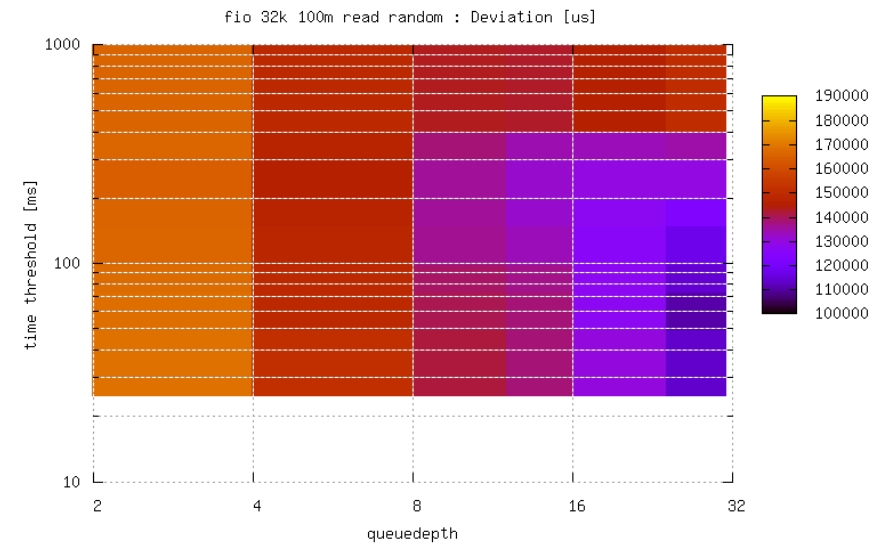
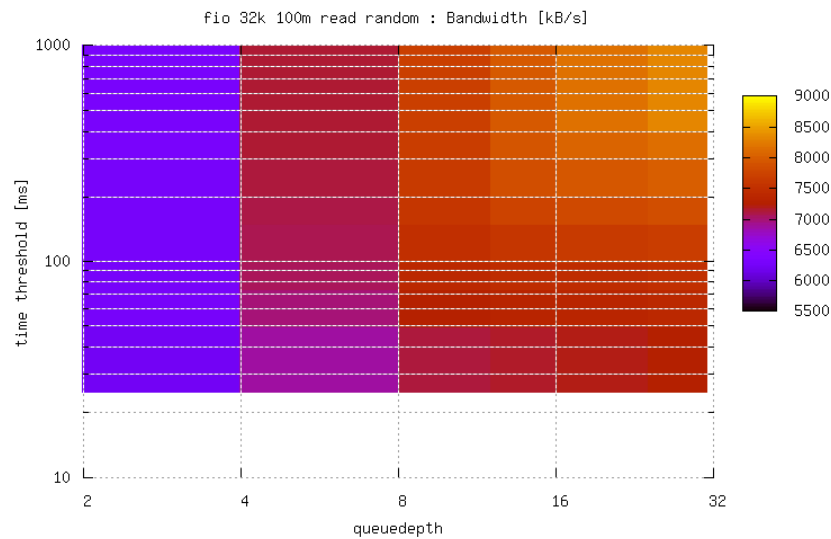
Current Status

- Code implemented to limit queue size and put a limit on IO completion time.




Current Status

- Impact on application
 - Raw performance slightly degraded
 - Better deviation



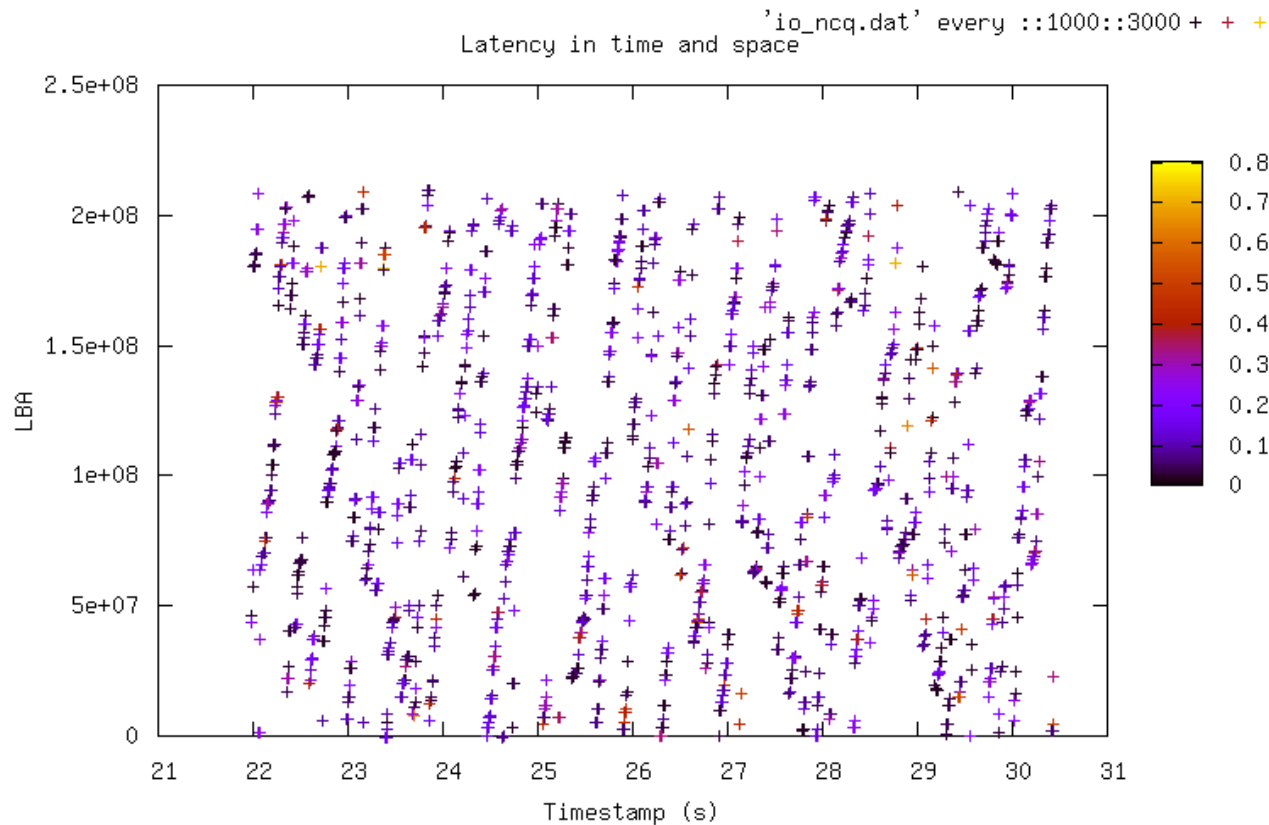
Risks

- NCQ increases our dependancy on vendors
 - Disks, PM, HBA
 - Behavior differences depending on models/fwrev



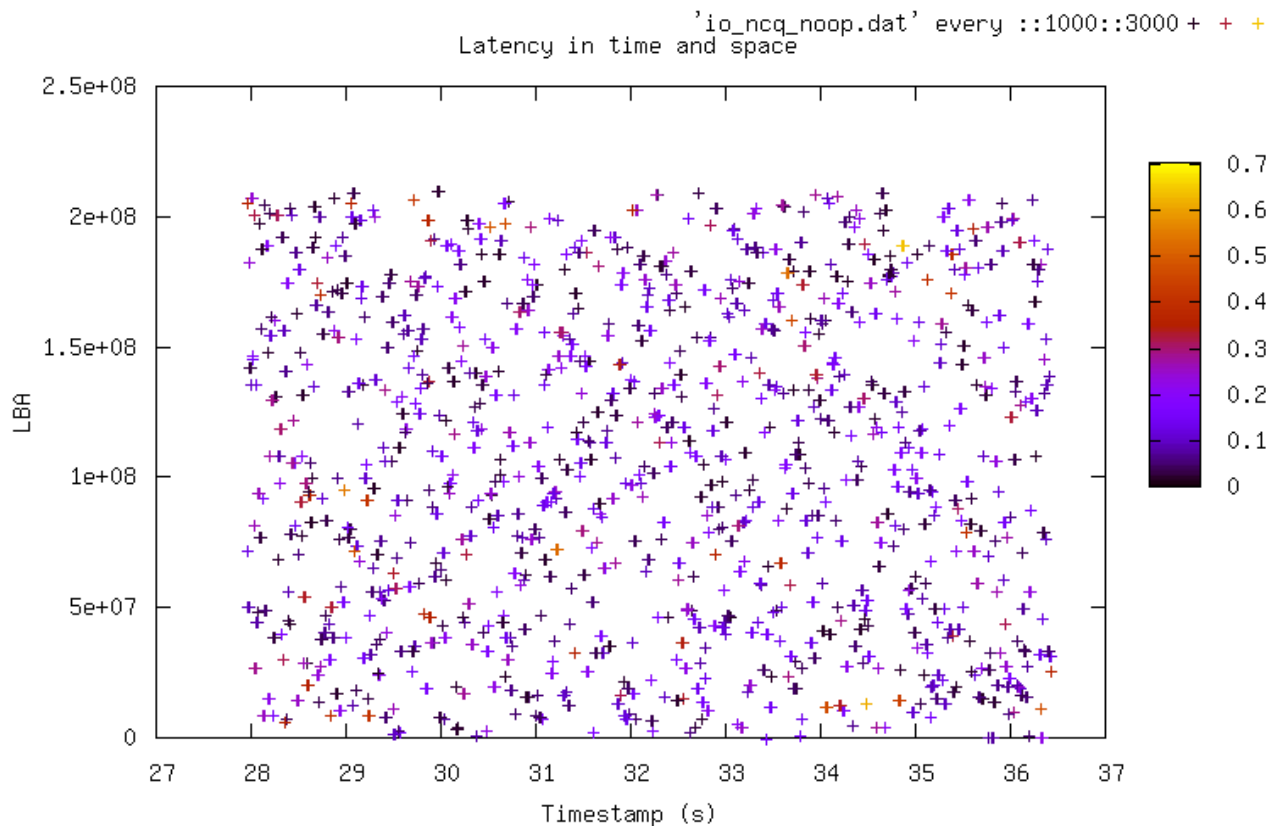
Scheduler Impact

- Random Test:
 - Using AS



Scheduler Impact

- Random Test:
 - Using NOOP



Longer Term

- NCQ in the host?
- Use Drive Topology
 - Not always provided, not always accurate
- Test the drive to figure out topology
- Modify Scheduler to use the topology?
 - too complex for marginal gain.

4.4.2 Cylinder allocation

Table 4: Cylinder allocation table

Zone	Logical cyl. (OD)	Logical cyl. (ID)	1000GB model #sec
0	0	8319	1680
1	8320	16639	1632
2	16640	24959	1620
3	24960	32639	1600
4	32640	39679	1560
5	39680	47359	1520
6	47360	51199	1520
7	51200	56319	1488
8	56320	61439	1440
9	61440	66559	1440
10	66560	71679	1392
11	71680	78719	1360
12	78720	81919	1360
13	81920	87039	1320
14	87040	92031	1280
15	92032	97023	1248
16	97024	101887	1200
17	101888	106367	1200
18	106368	110719	1152
19	110720	114175	1120
20	114176	118015	1080
21	118016	121727	1080
22	121728	125823	1040
23	125824	129919	1008
24	129920	134015	960
25	134016	137855	912
26	137856	140671	912
27	140672	143231	880
28	143232	146303	880
29	146304	147583	840

Data cylinder

This cylinder contains the user data which can be sent and retrieved via read/write commands and a spare area for reassigned data.

Spare cylinder

The spare cylinder is used by Hitachi Global Storage Technologies manufacturing and includes data sent from a defect location.