


Replacing the SSD with Better Code


Brian Bockelman

Written in 10 minutes, please excuse mistakes

Oh the woes of HEP!



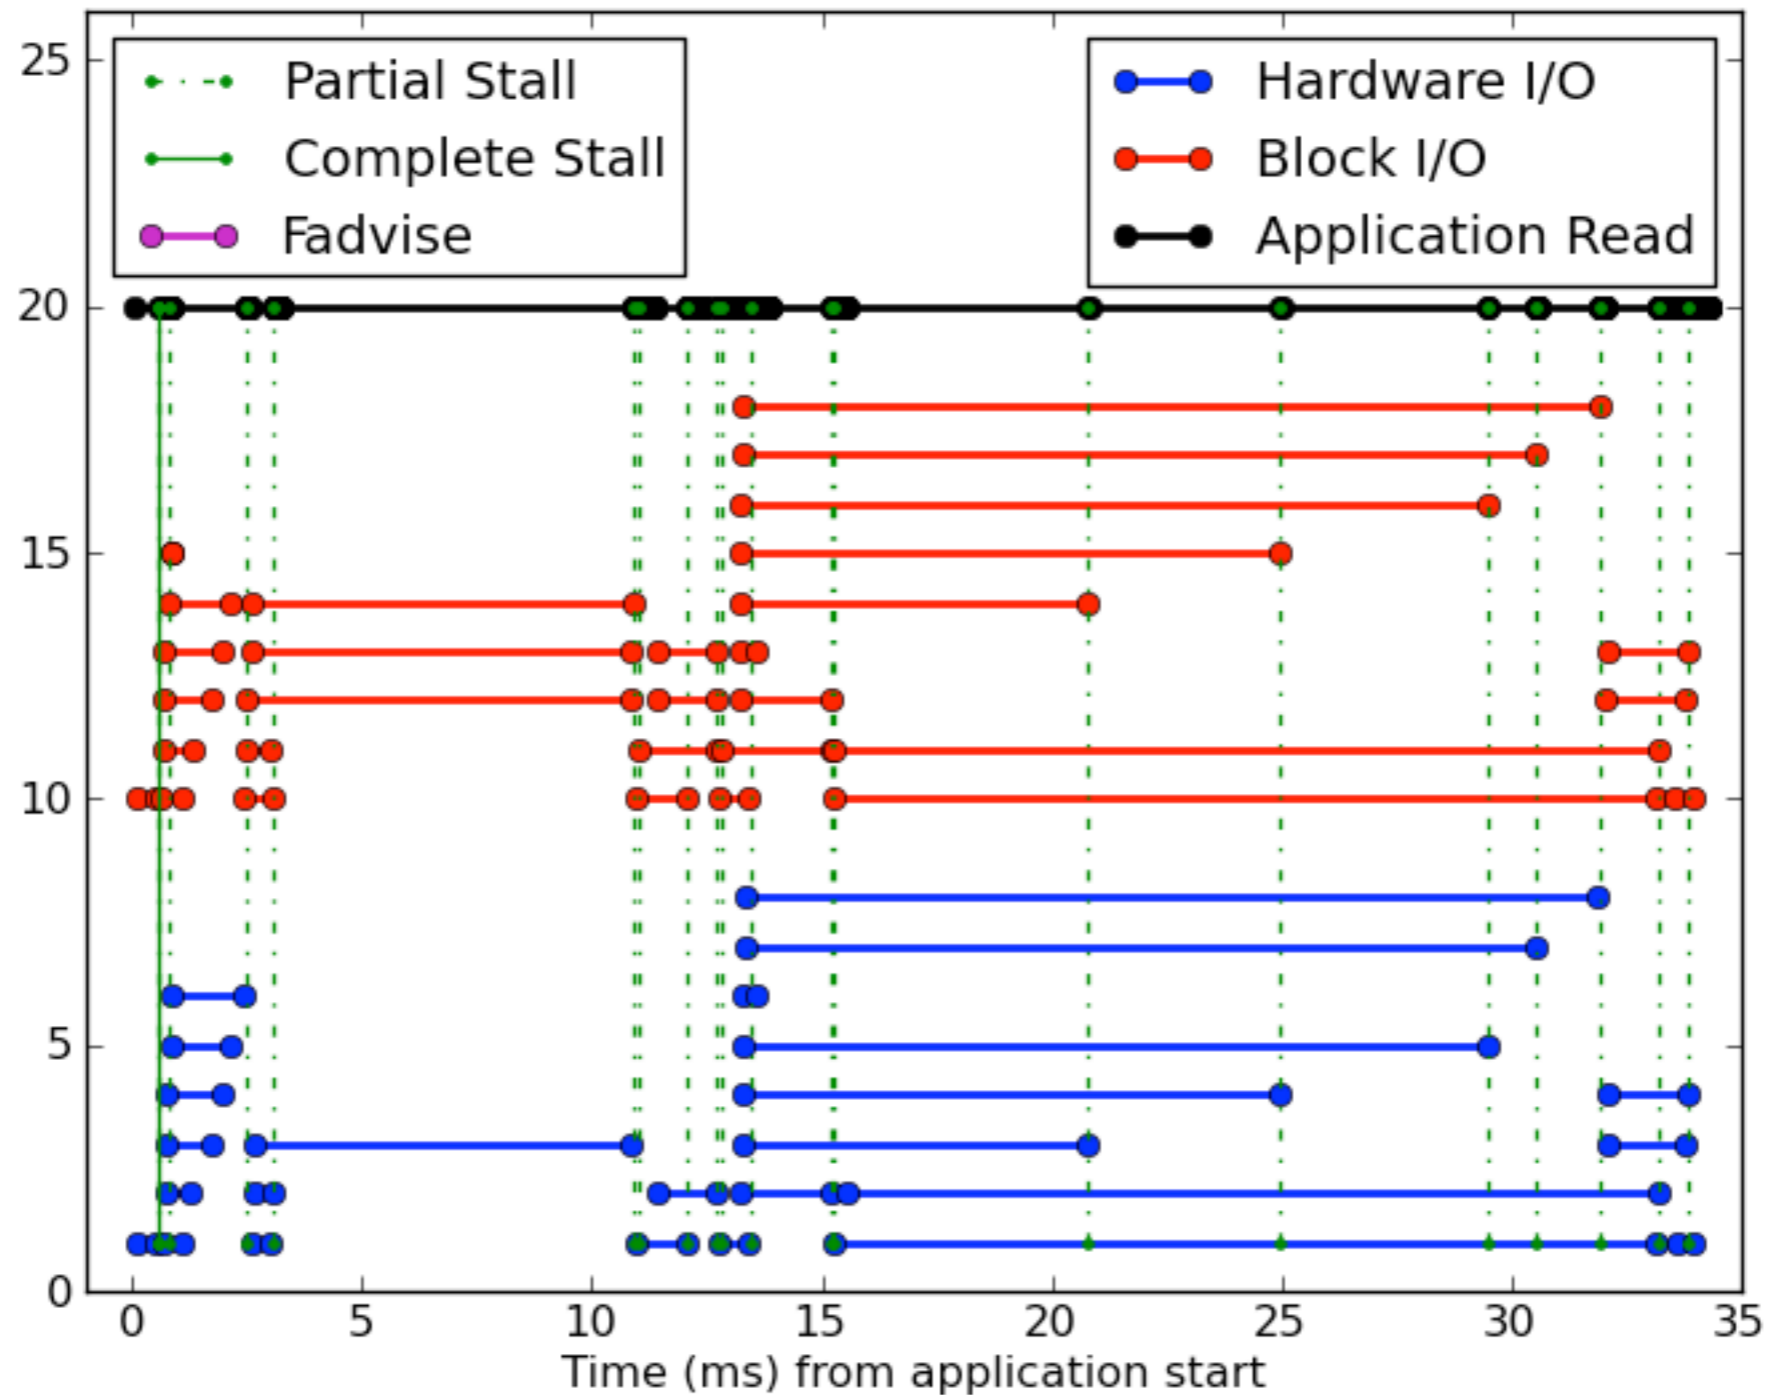
We need SSDs!




Measuring I/O

'cat' example


I/O Lifeline



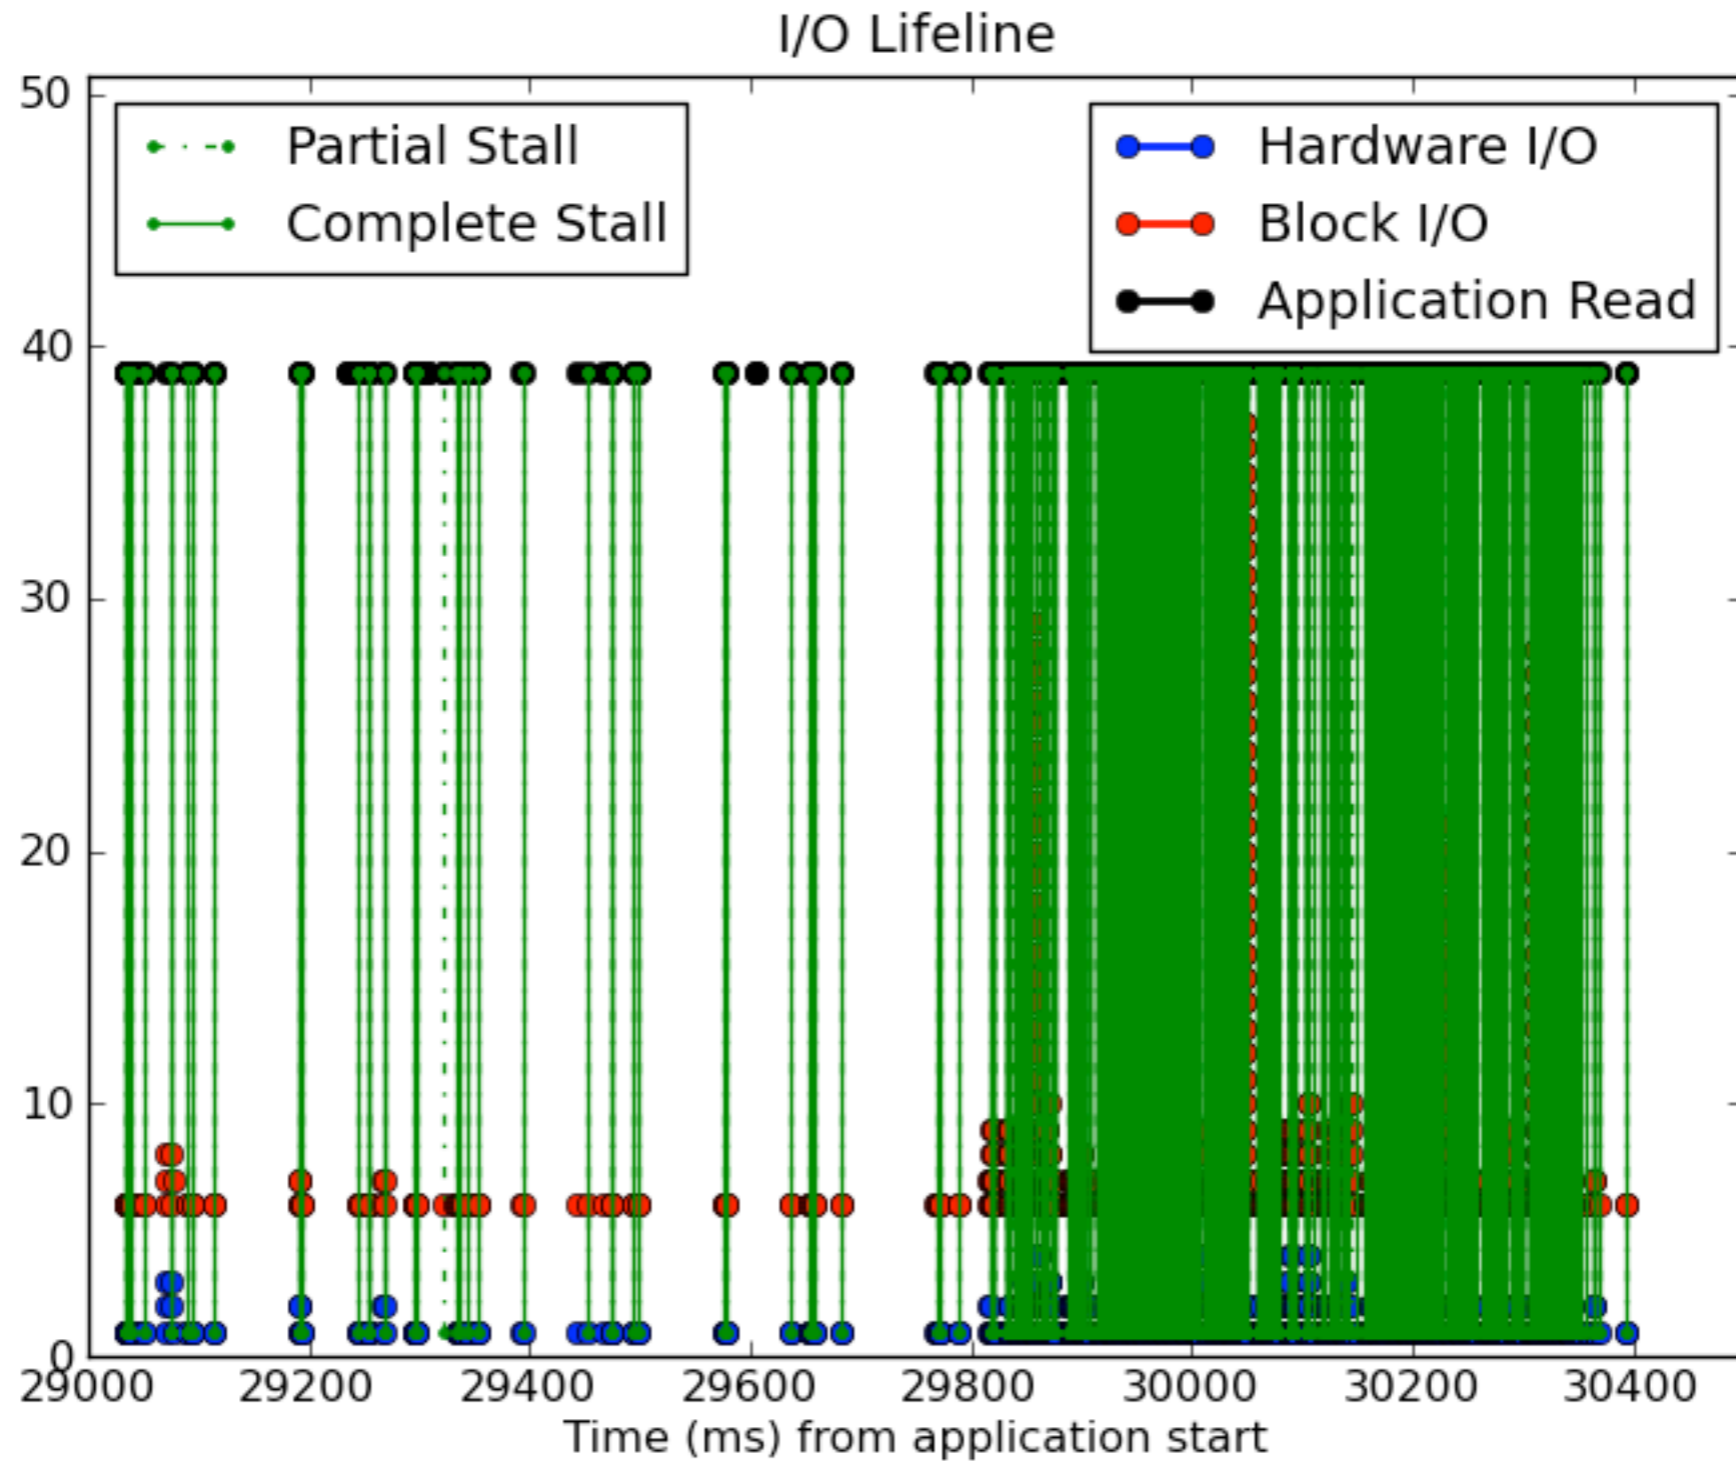
CMS Application




With App Prefetching



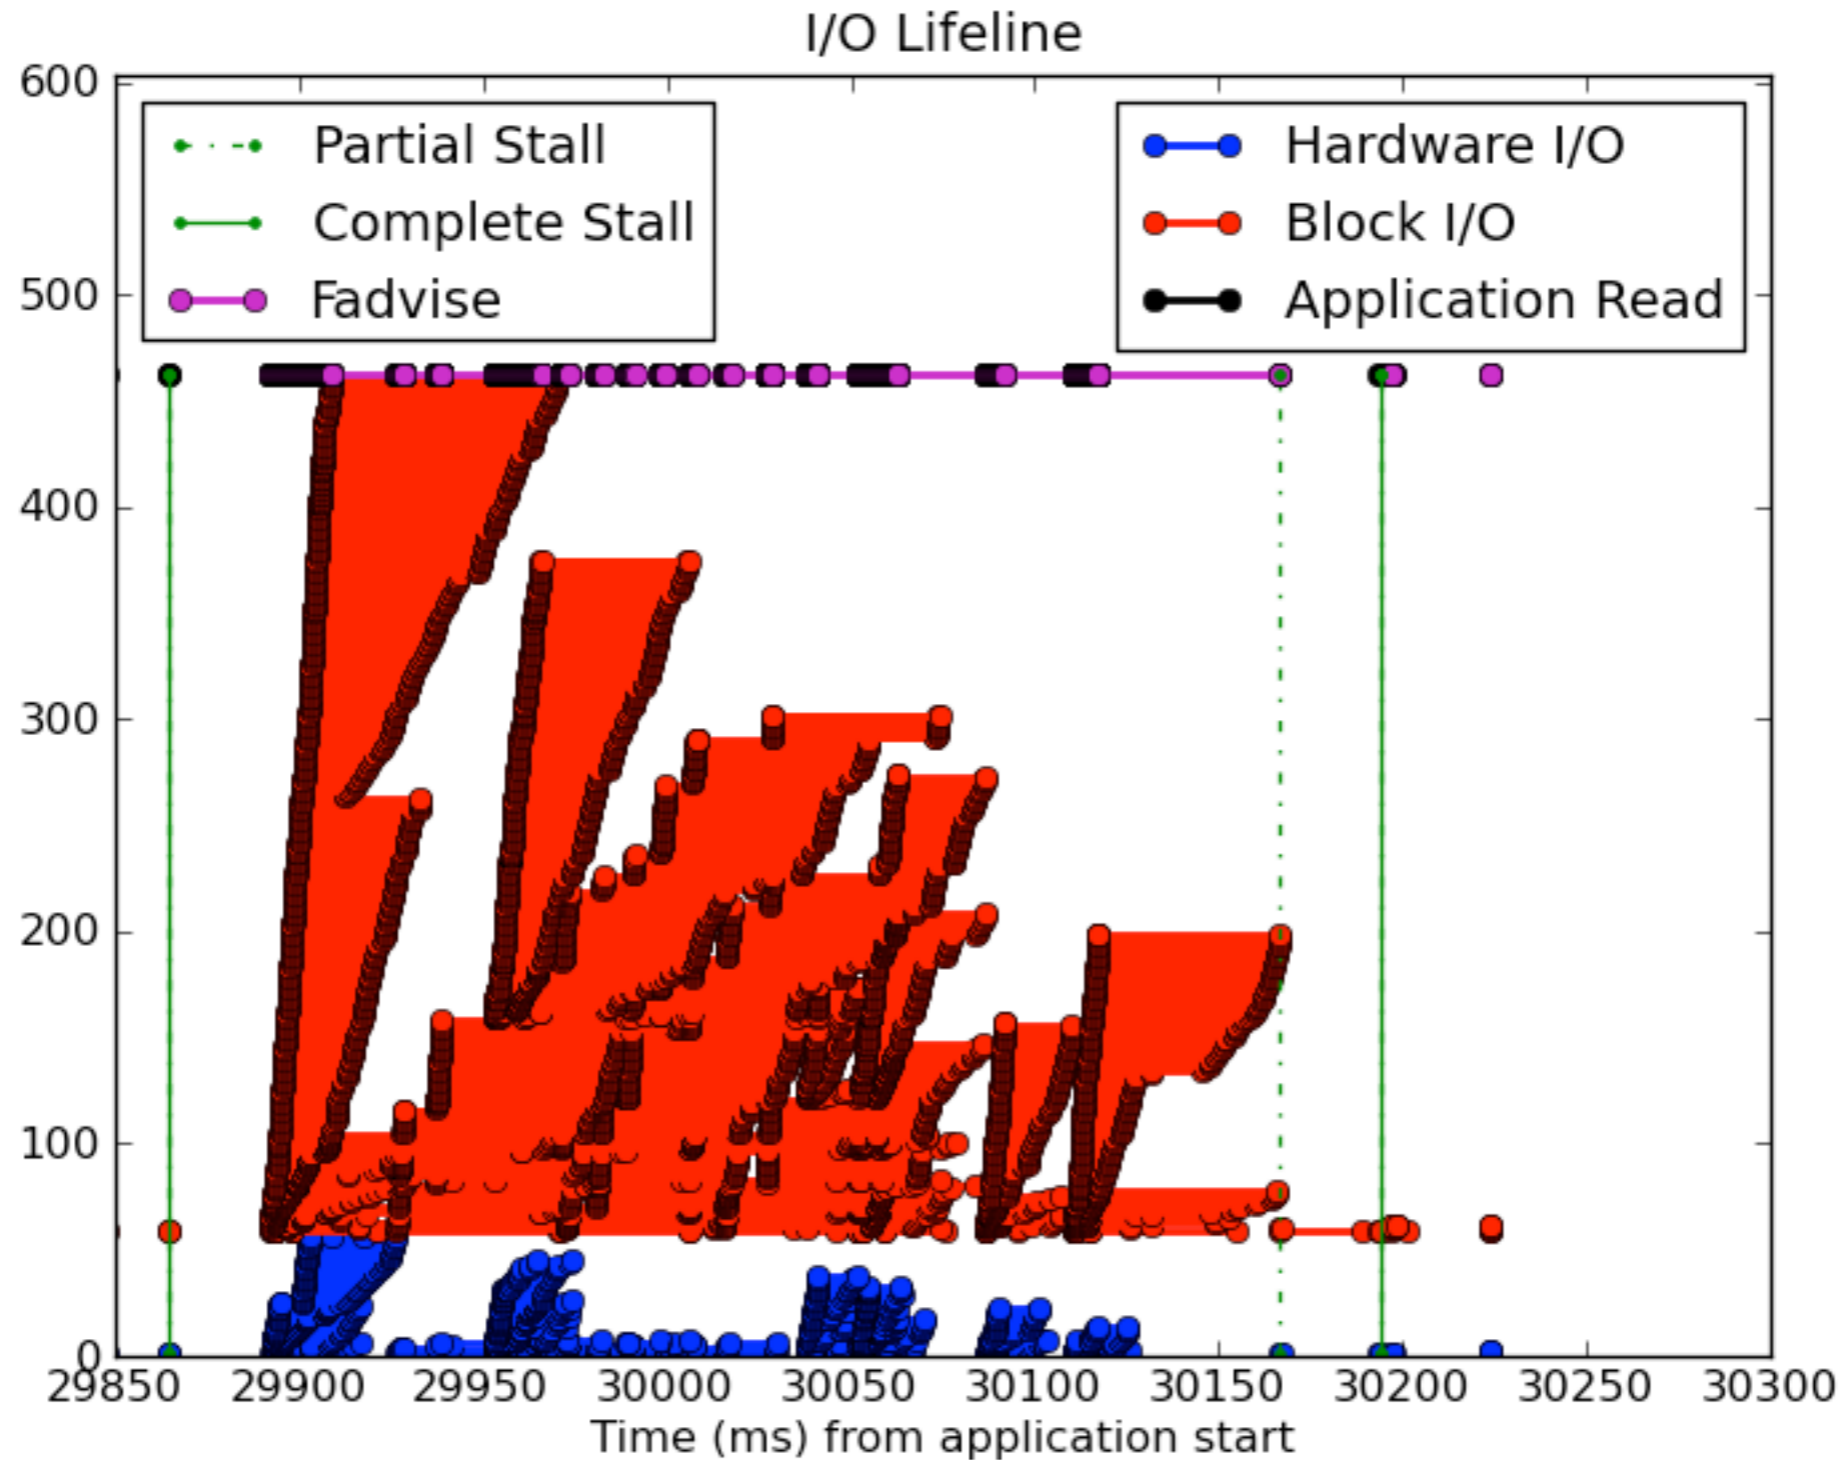
Zoom-in



With OS Prefetching

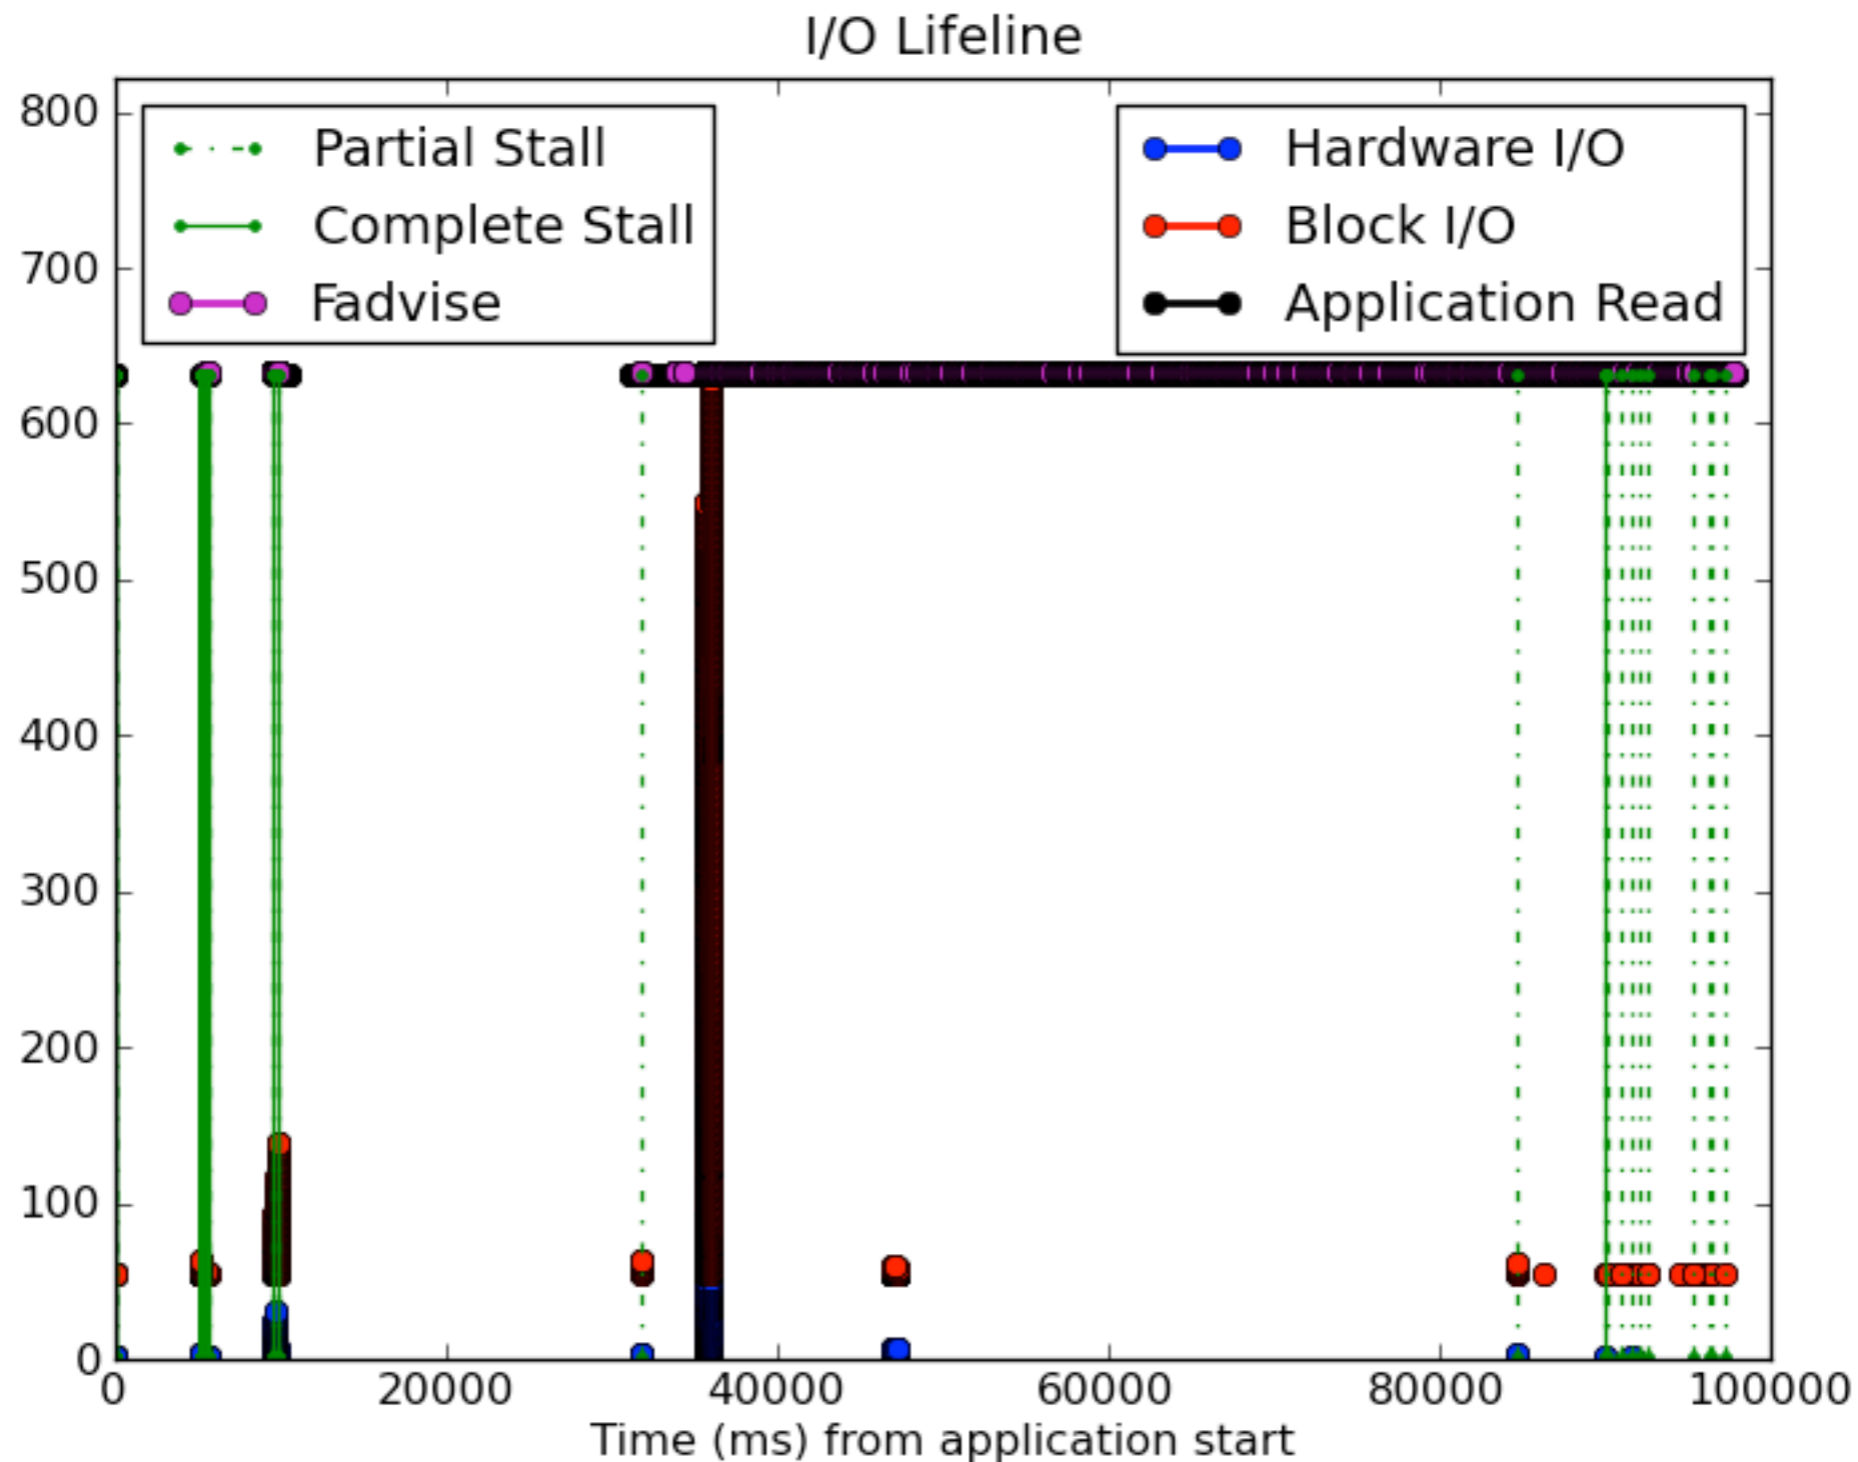


Zoom-In

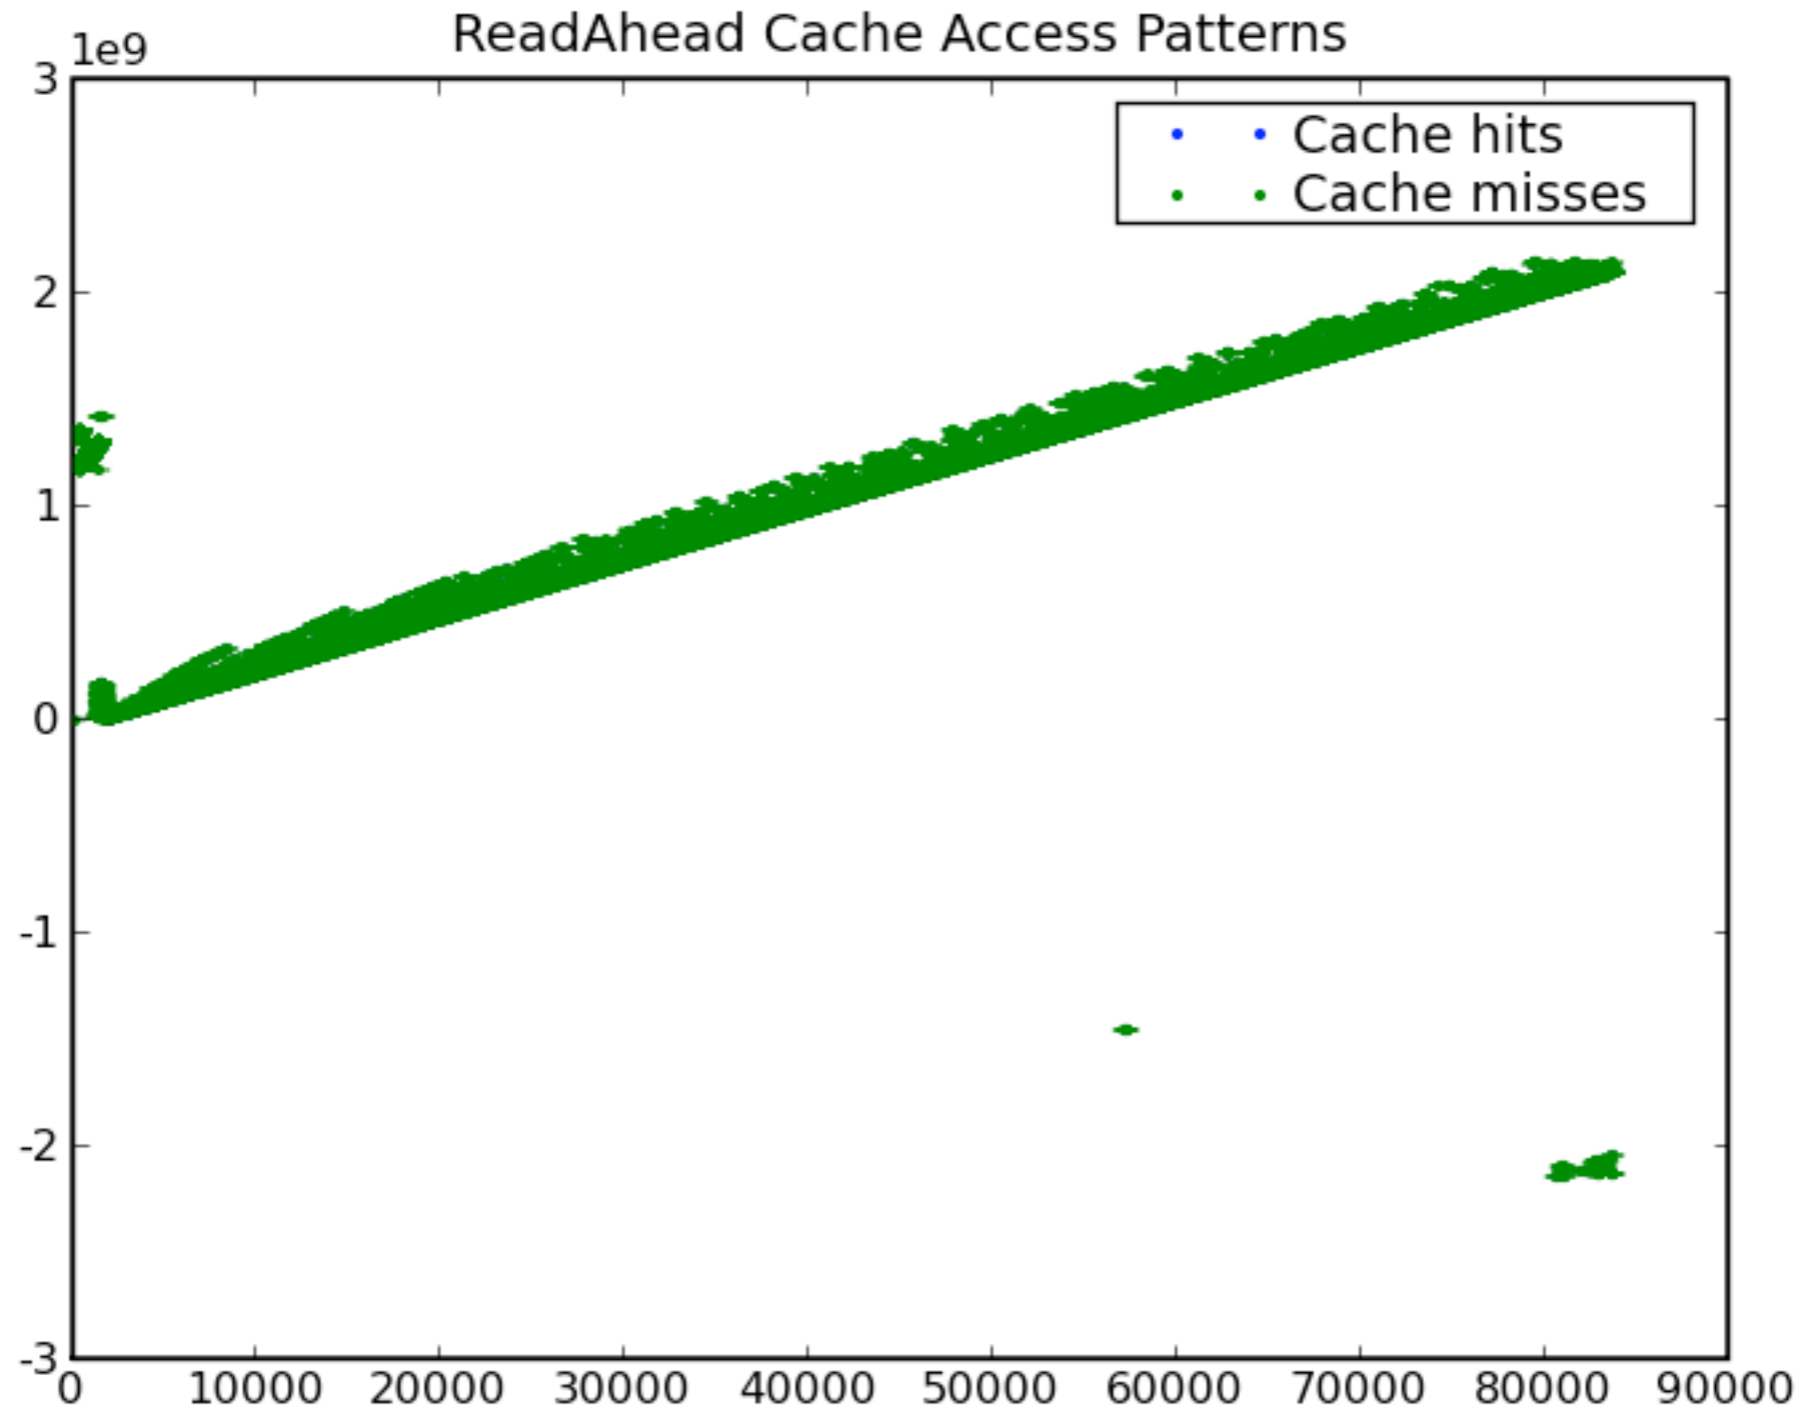


New Techniques

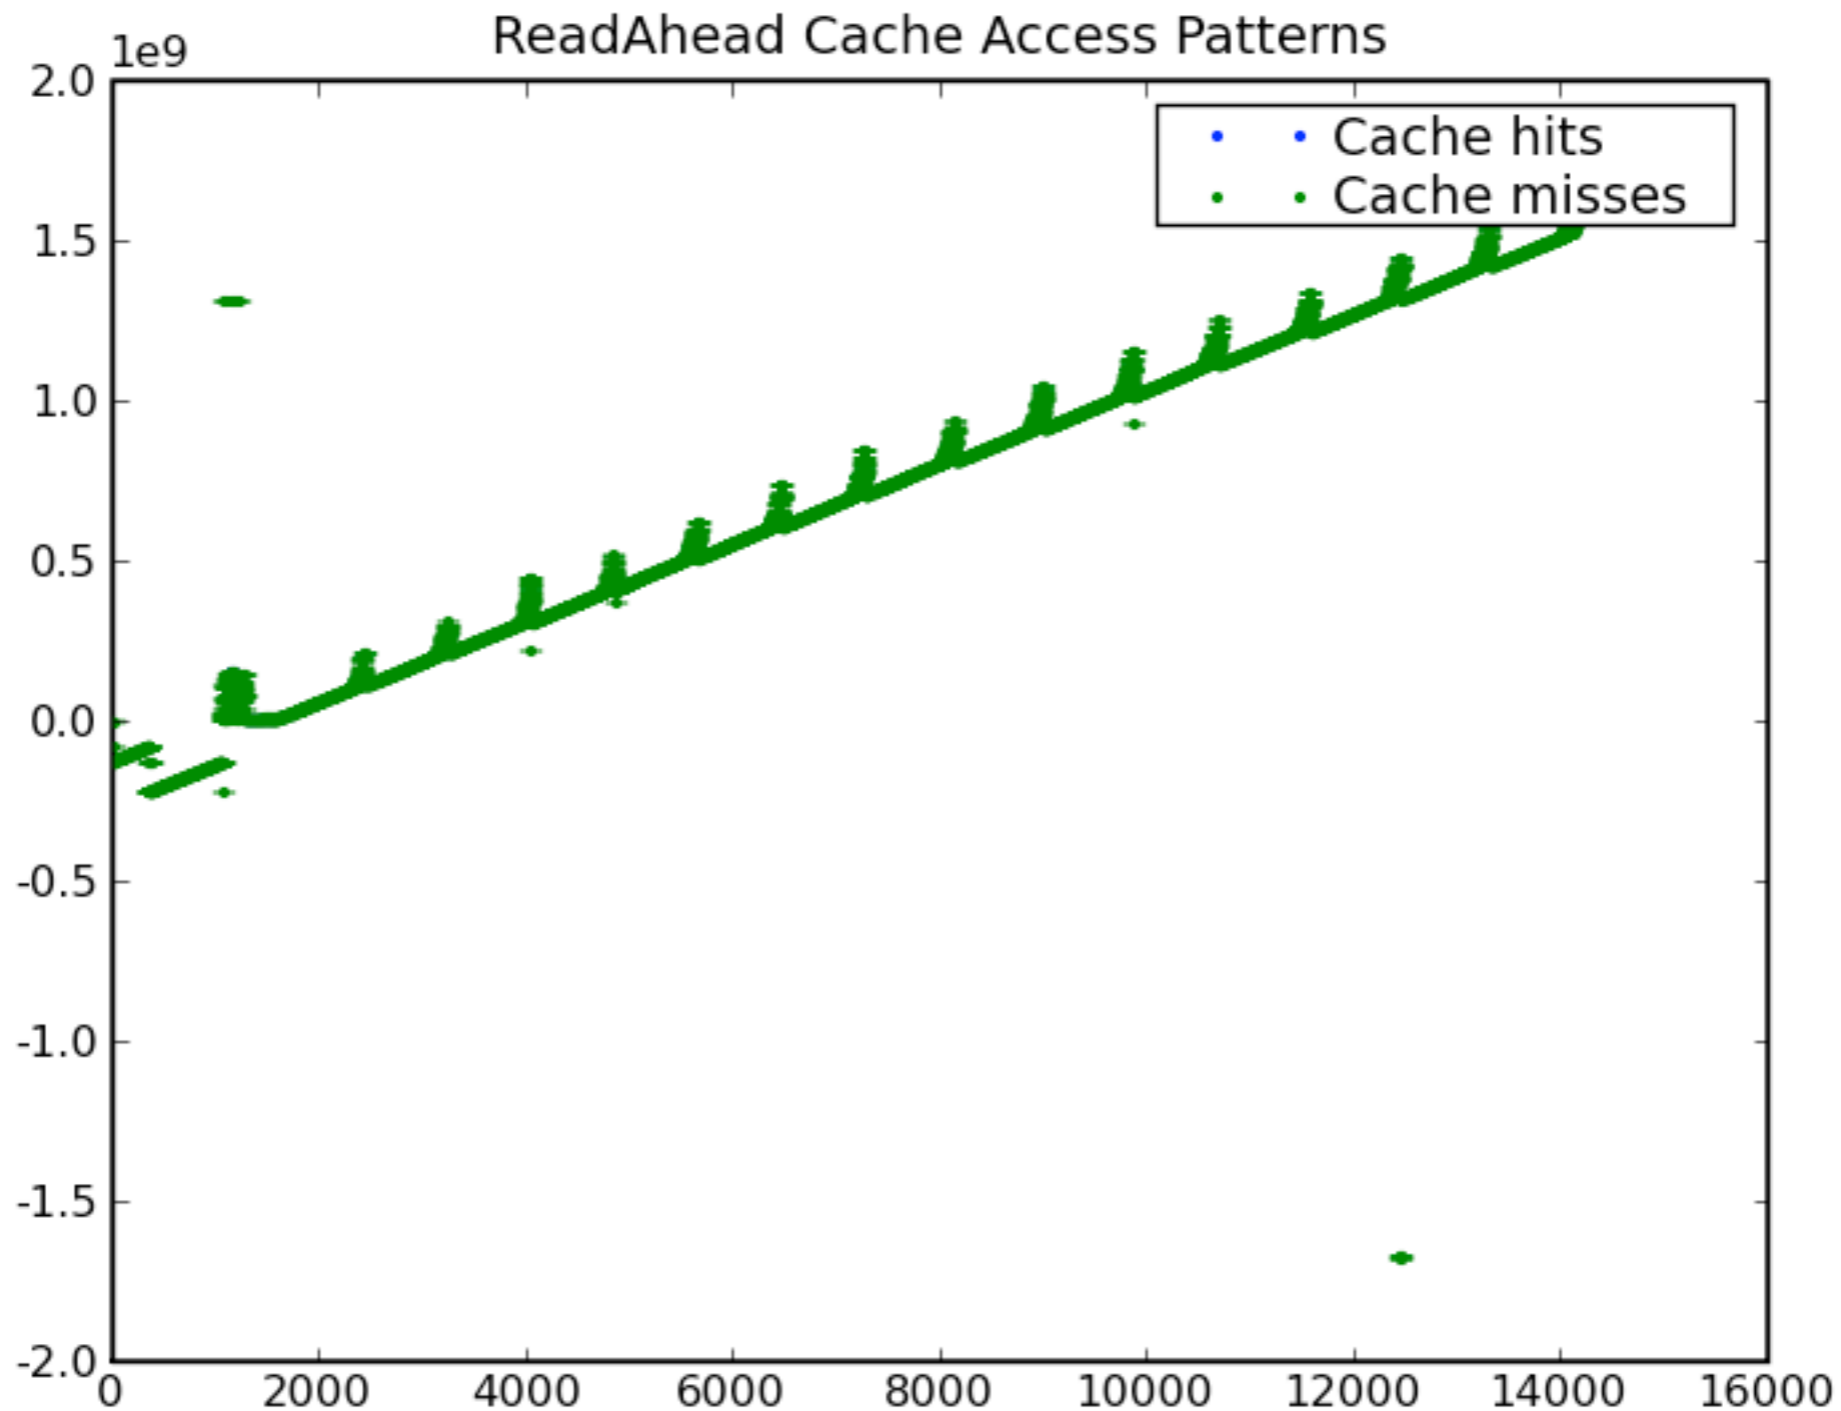
Based on the kernel traces, we were able to work further



Before (90,000 reads)



After (16,000 reads)



Future

We can now increase the latency - by a lot!

- Working with the ROOT team along these lines, we can get this workflow down to around ~5,000 reads.
- Because almost all HEP is based on the ROOT libraries, improvements they make causes huge downstream changes.
- Conclusion: It would have been fun to buy 1PB of SSDs. Fixing HEP was cheaper.