Using TCP/IP Traffic shaping to achieve iSCSI service predictability

Paper presentation

Jarle Bjørgeengen

University of Oslo / USIT

November 11, 2010

Outline

- About resource sharing in storage devices
- Lab setup / job setup
- Experiment illustrating the problem
- One half of the solution: the throttle
- Live demo
  - The throttle
  - Part two of the solution: the controller
- How the controller works
- Conclusion and future work

General problem of sharing resources

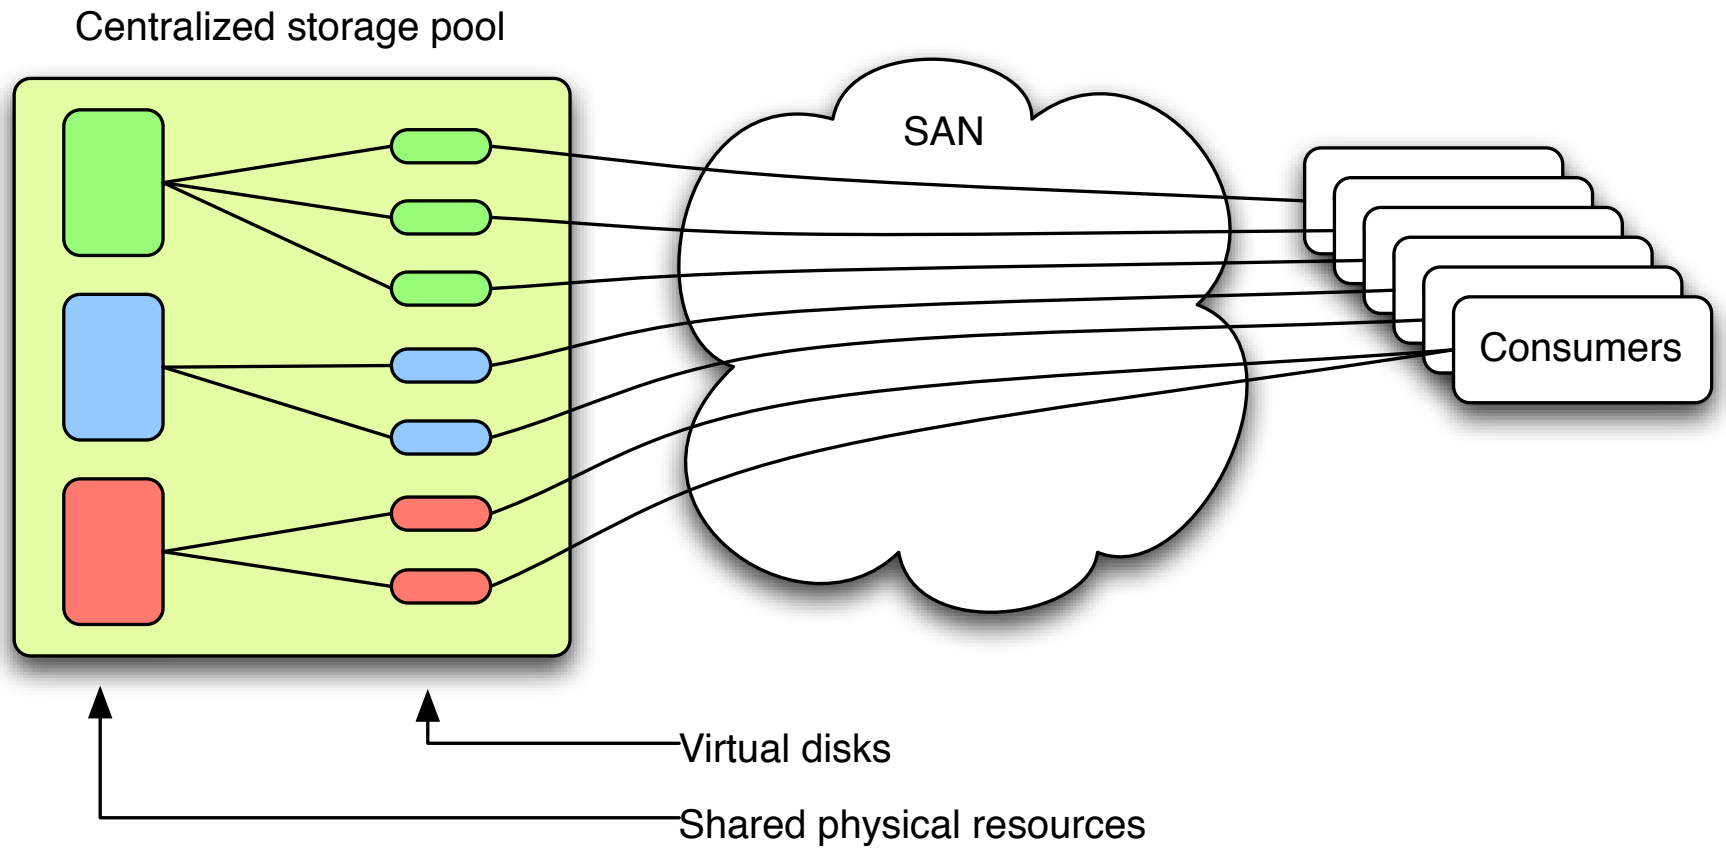

Free competition causes unpredictable I/O performance for any given consumer.

Lab setup

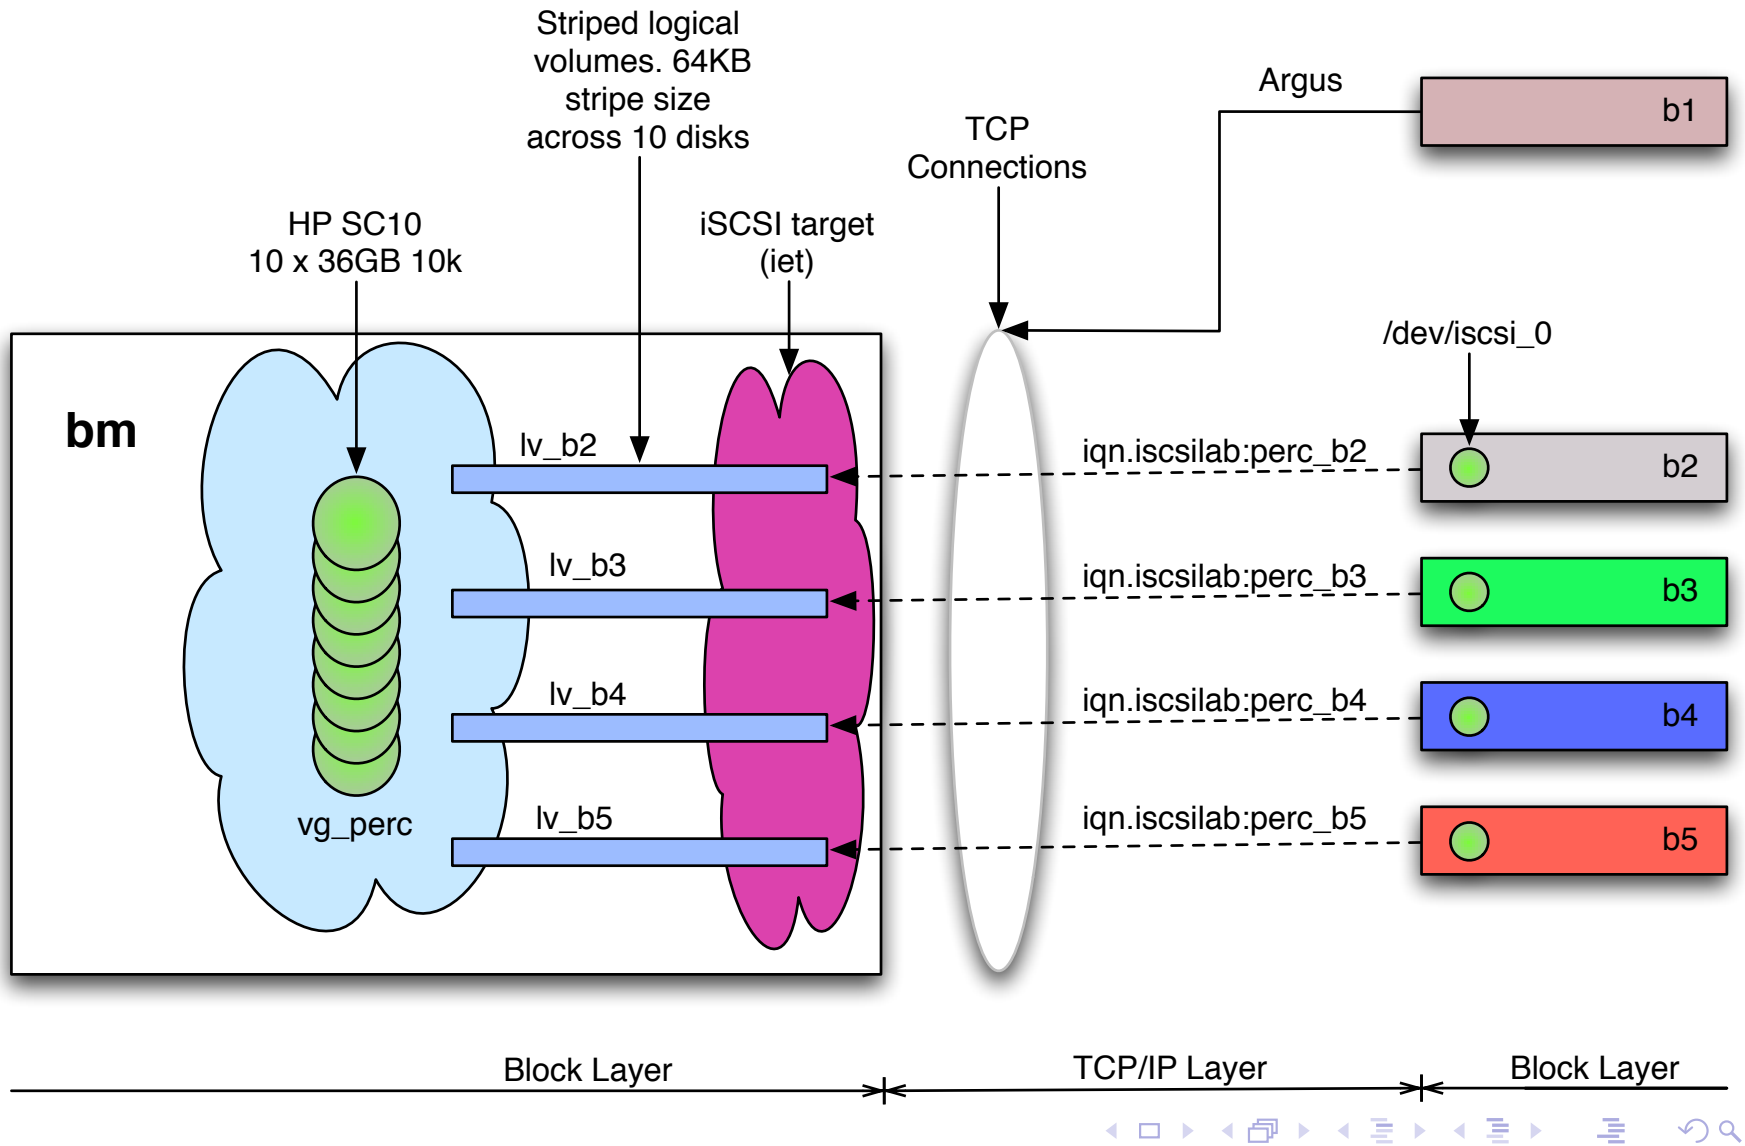
Striped logical volumes. 64KB stripe size across 10 disks
Argus
TCP Connections
b1
HP SC10 10 x 36GB 10k
iSCSI target (iet)
/dev/iscsi_0
bm
lv_b2
iqn.iscsilab:perc_b2
b2
lv_b3
iqn.iscsilab:perc_b3
b3
lv_b4
iqn.iscsilab:perc_b4
b4
vg_perc
lv_b5
iqn.iscsilab:perc_b5
b5
Block Layer
TCP/IP Layer
Block Layer

Is read response time affected by write activity ?

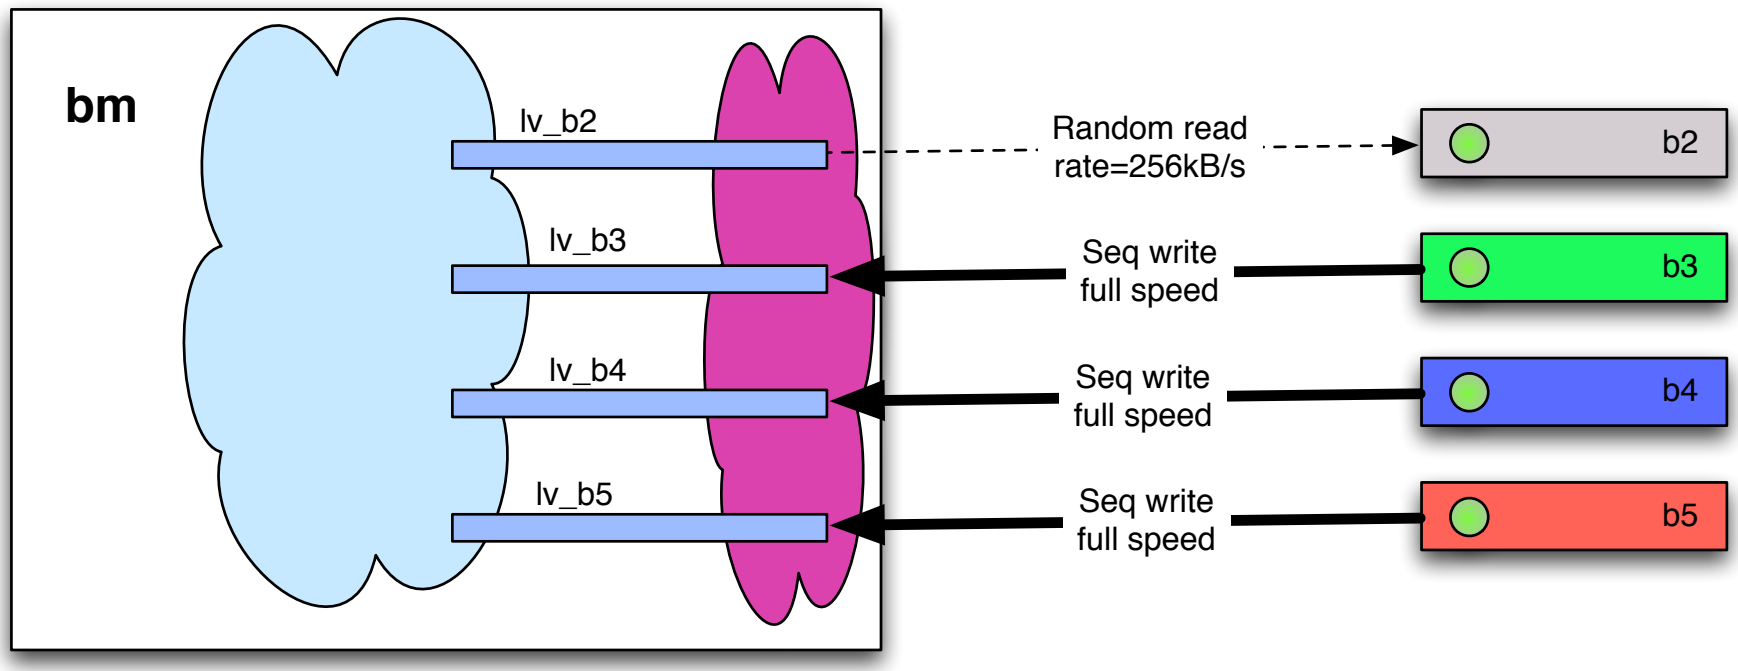

The Answer is yes

- Long response times adversely affect application service availability.

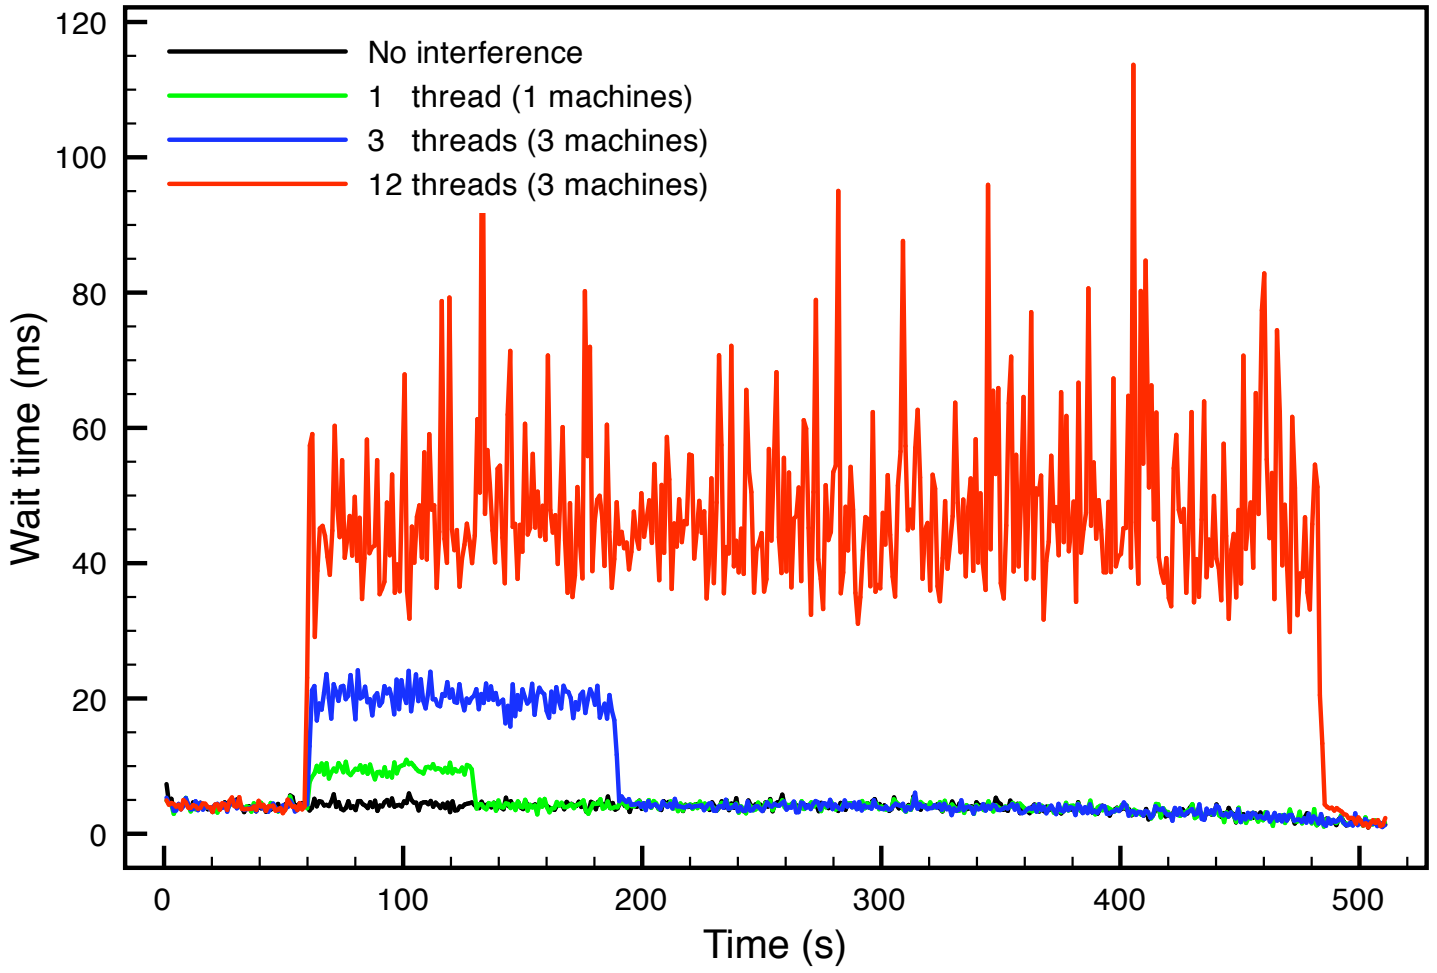

Throttling method

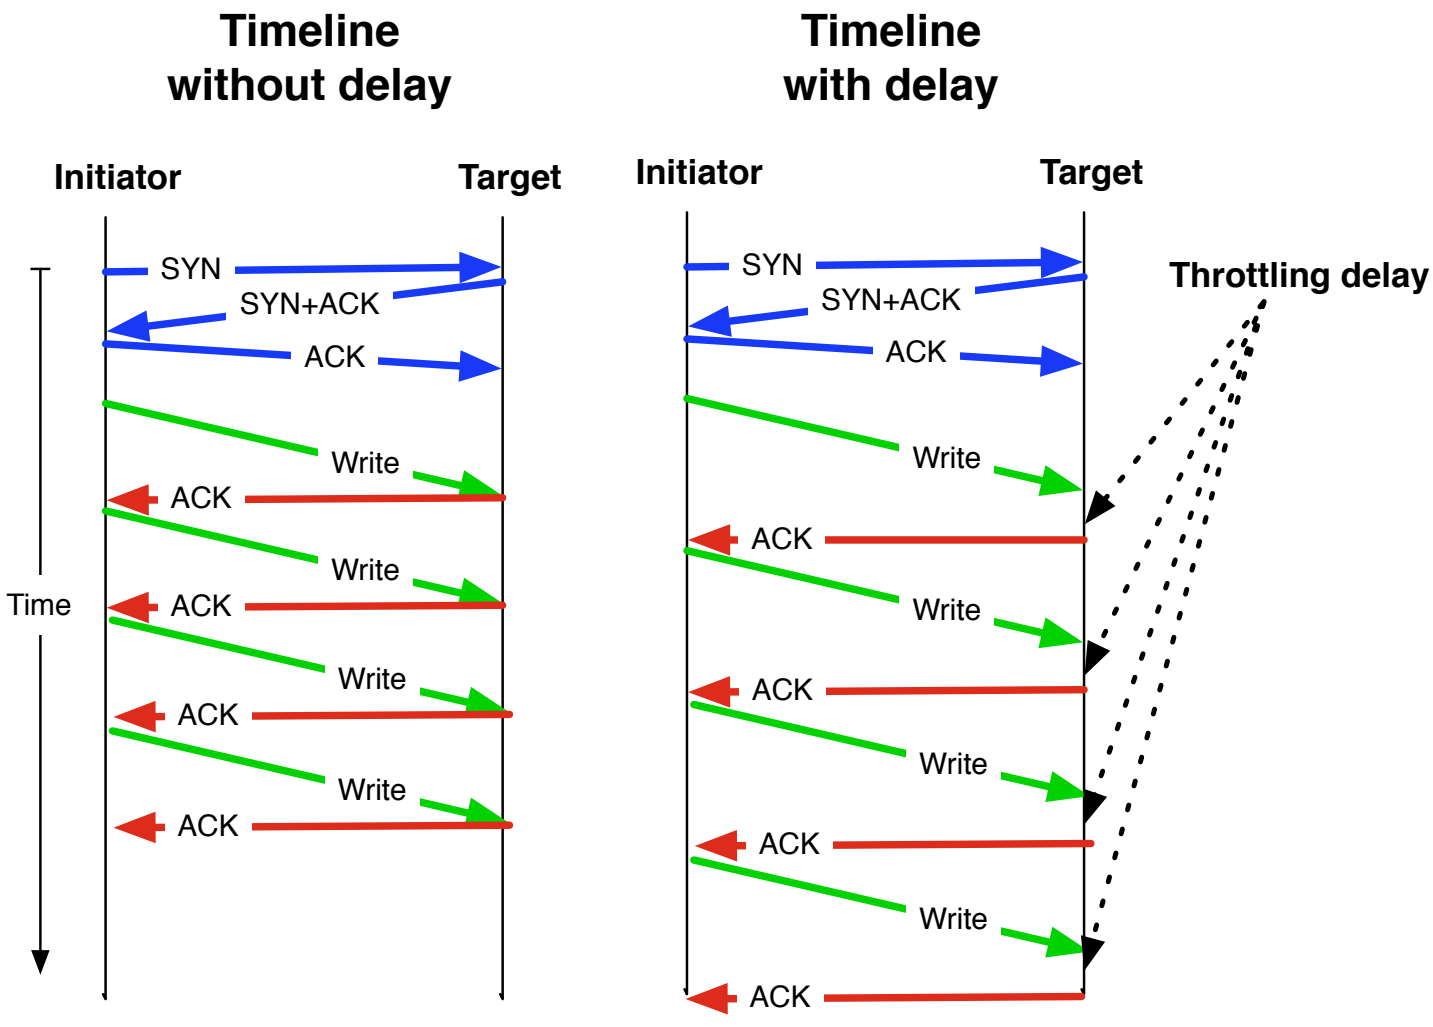

Relation between packet delay and average rate

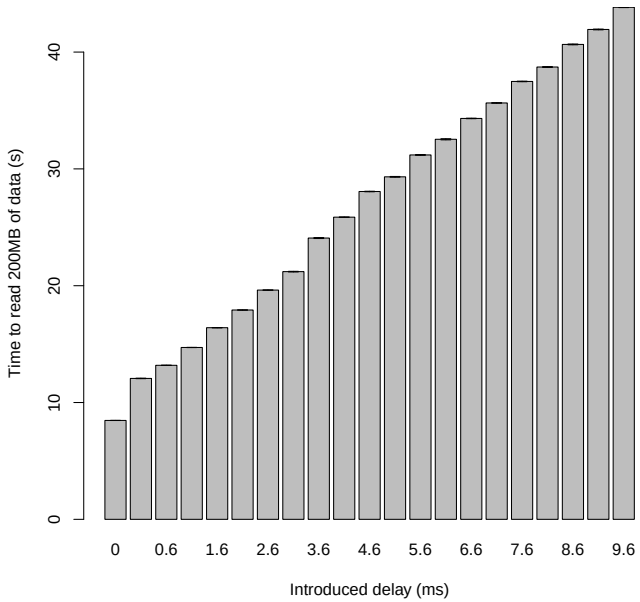

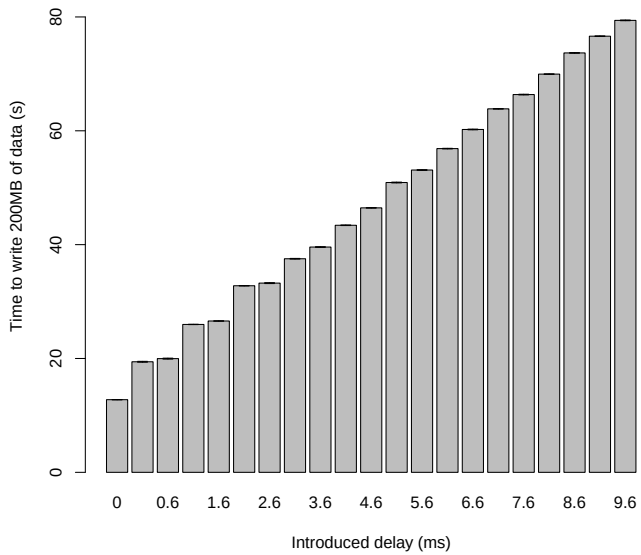

- Write rate 15 MB/s - 2.5 MB/s
- Read rate 22 MB/s - 5 MB/s

Managing consumers

- Need to operate on sets of consumers (throttlable={10.0.0.243,10.0.0.244})
- Ipset: One rule to match them all

```
ipset -N $throttlable  ipmap --network 10.0.0.0/24
ipset -A $throttlable 10.0.0.243
ipset -A $throttlable 10.0.0.244
iptables --match-set $throttlable dst -j MARK --set-mark $mark
```

- The mark is a step in the range of available packet delays

Live demonstration

- Manual throttling and QoS specification
- An automatic QoS policy and automated throttling

Dynamic throttling decision

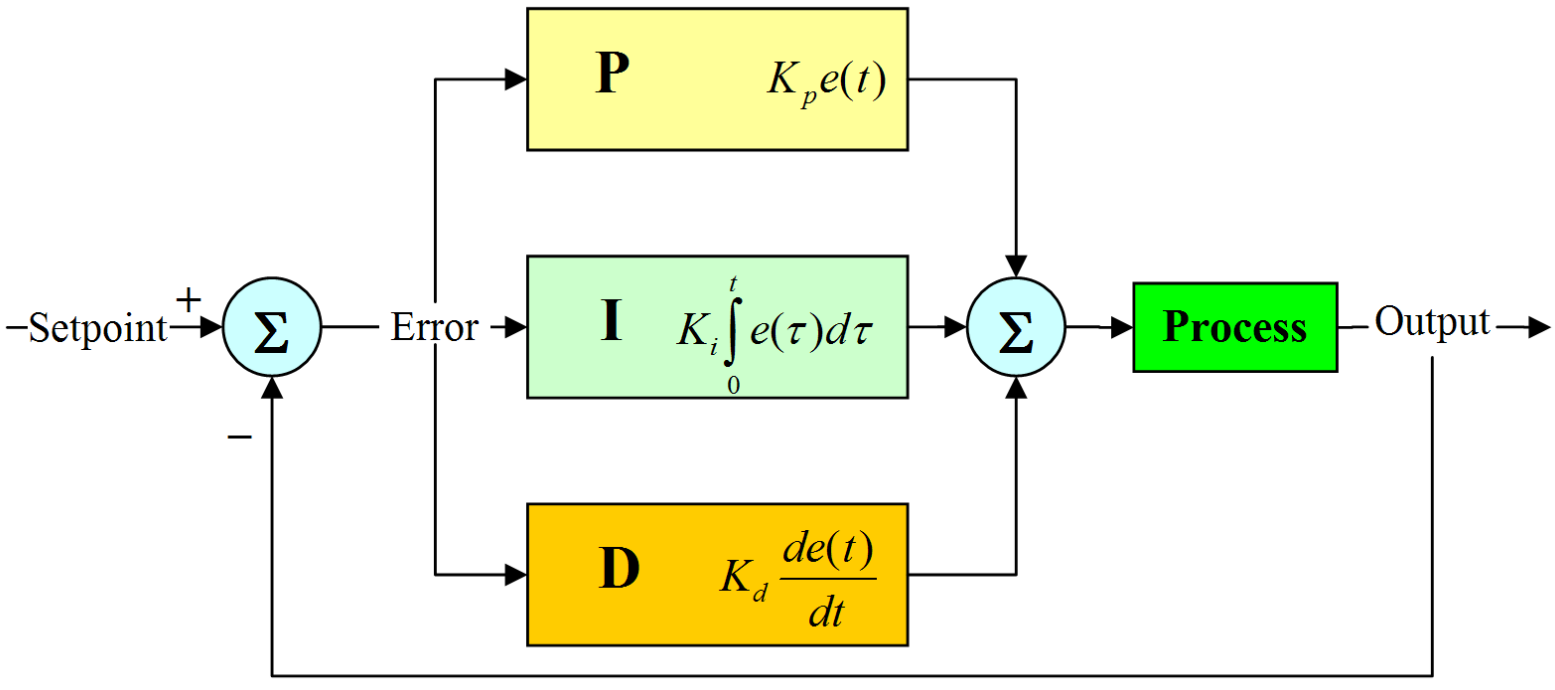

Figure: Block diagram of a PID controller. Created by SilverStar(at)en.wikipedia. Licensed under the terms of Creative Commons Attribution 2.5 Generic.

Modified PID function

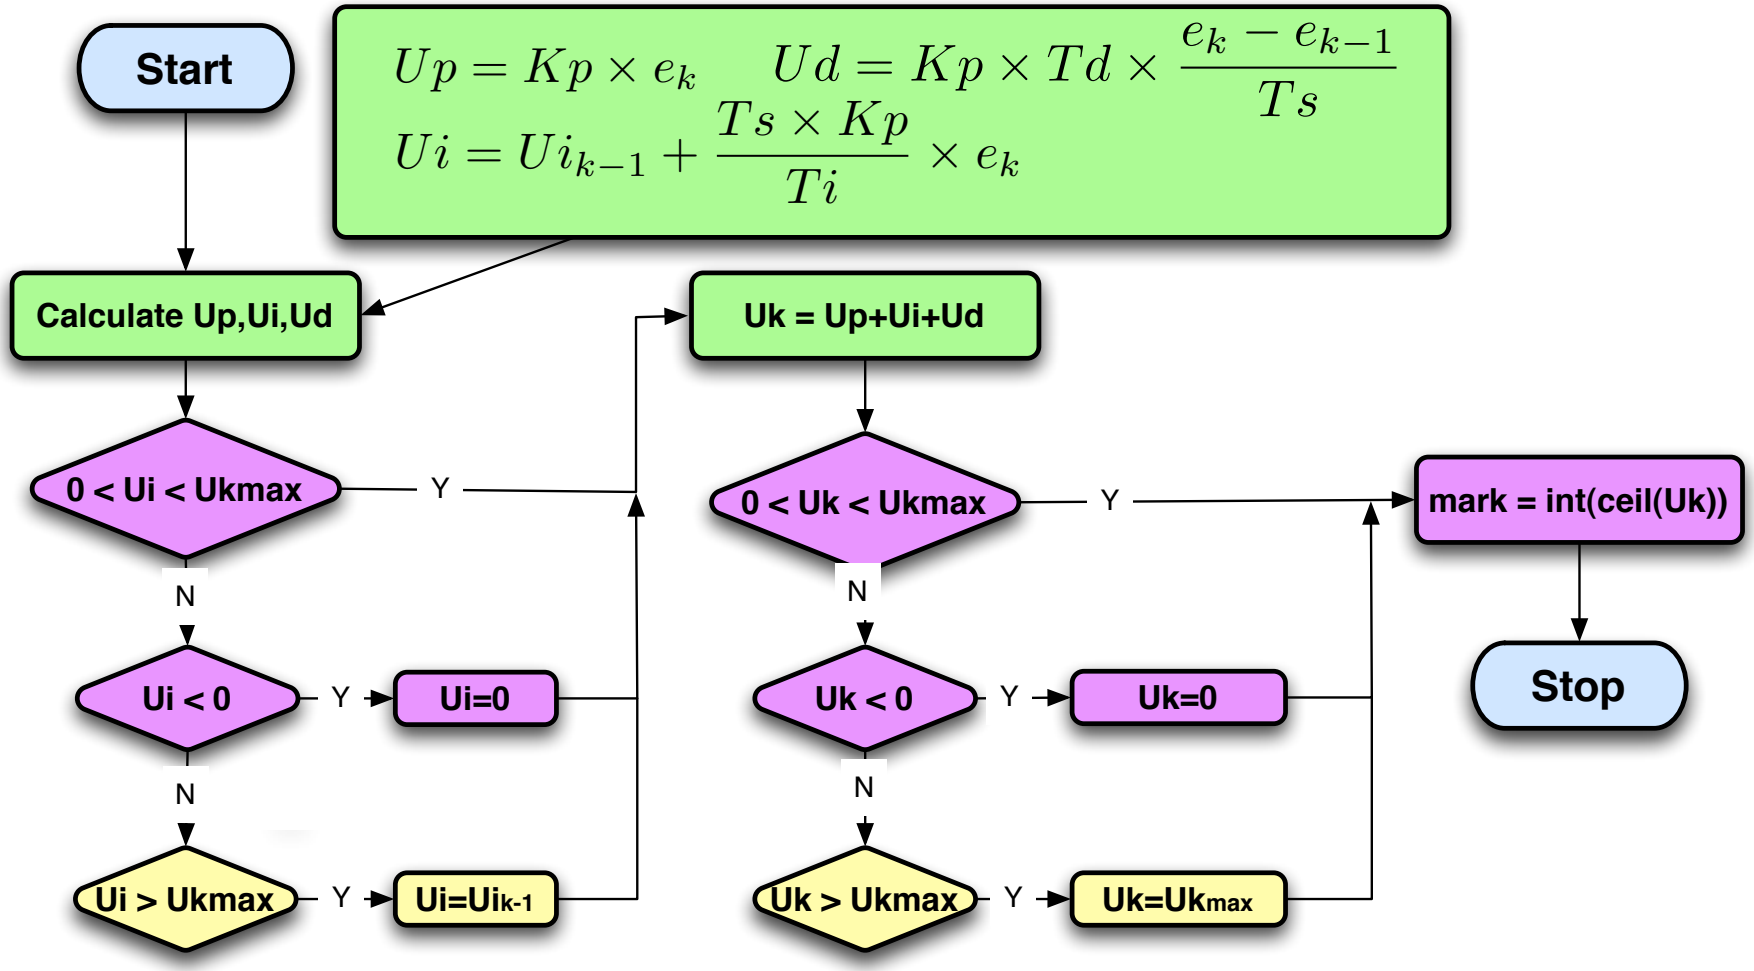

The completely automated approach

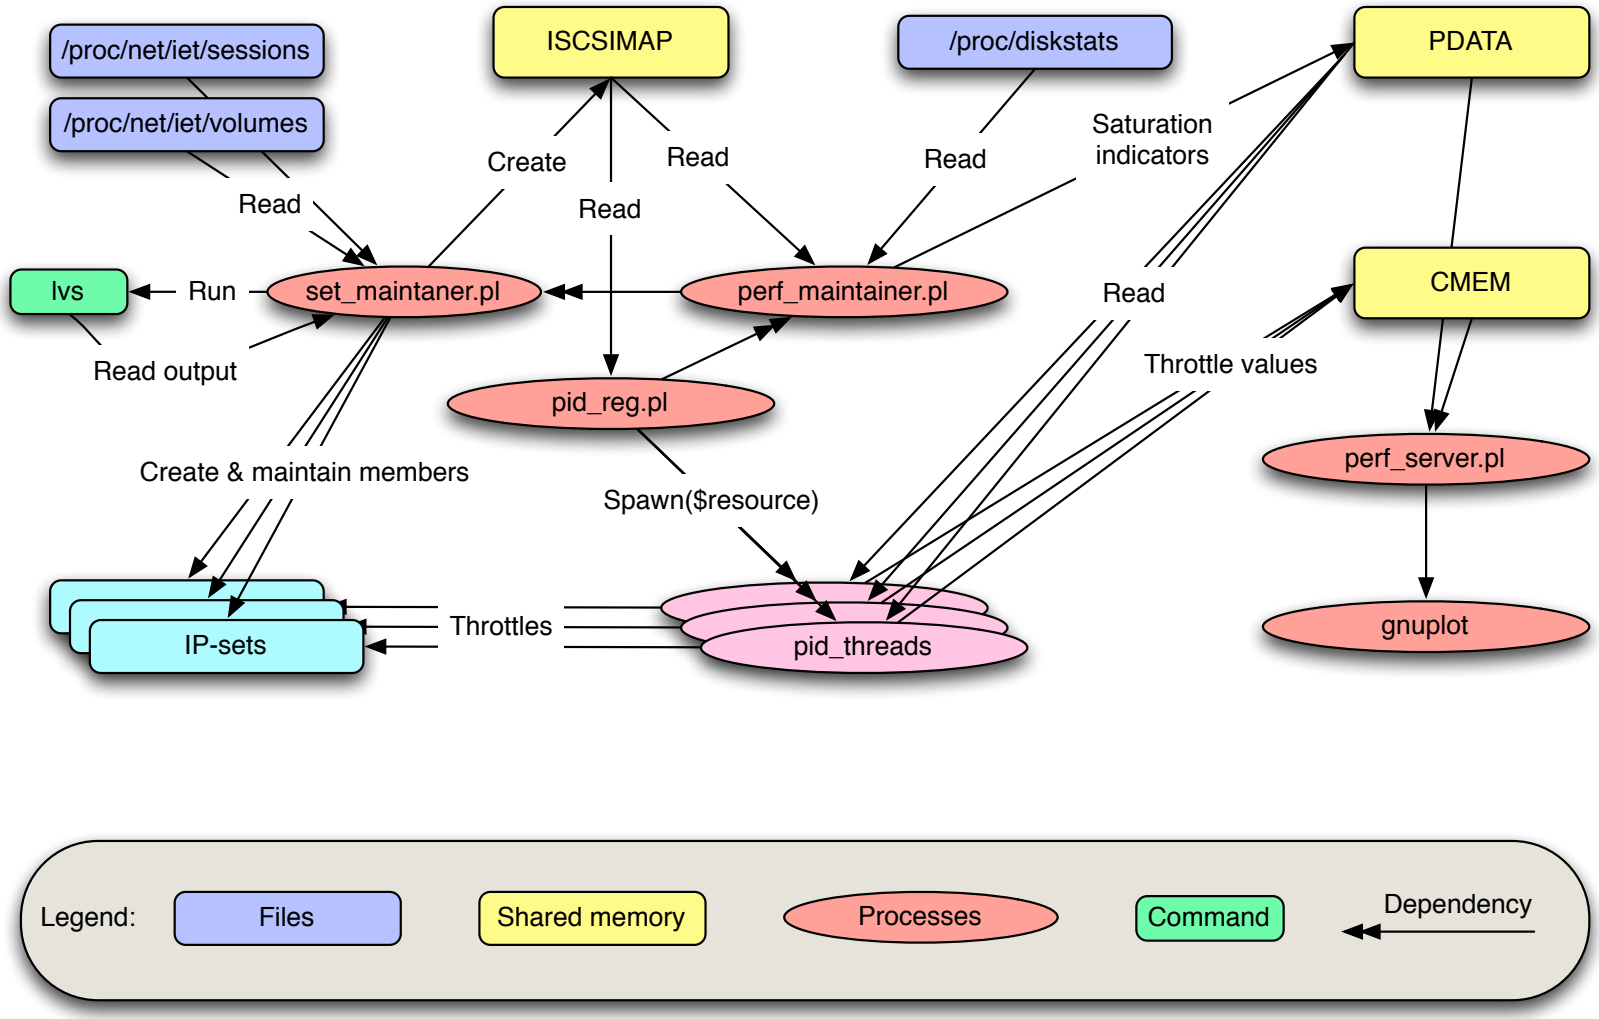

Impact

- The packet delay throttle is very efficient
  - Solves the throttling need completely for iSCSI (likely other TCP based storage networks too)
- The modified PID controller is consistently keeping response time low in spite of rapidly changing load interference.
- The concept is widely applicable

Impact

- The packet delay throttle is very efficient
  - Solves the throttling need completely for iSCSI (likely other TCP based storage networks too)
- The modified PID controller is consistently keeping response time low in spite of rapidly changing load interference.
- The concept is widely applicable

Impact

- The packet delay throttle is very efficient
  - Solves the throttling need completely for iSCSI (likely other TCP based storage networks too)
- The modified PID controller is consistently keeping response time low in spite of rapidly changing load interference.
- The concept is widely applicable

Future work

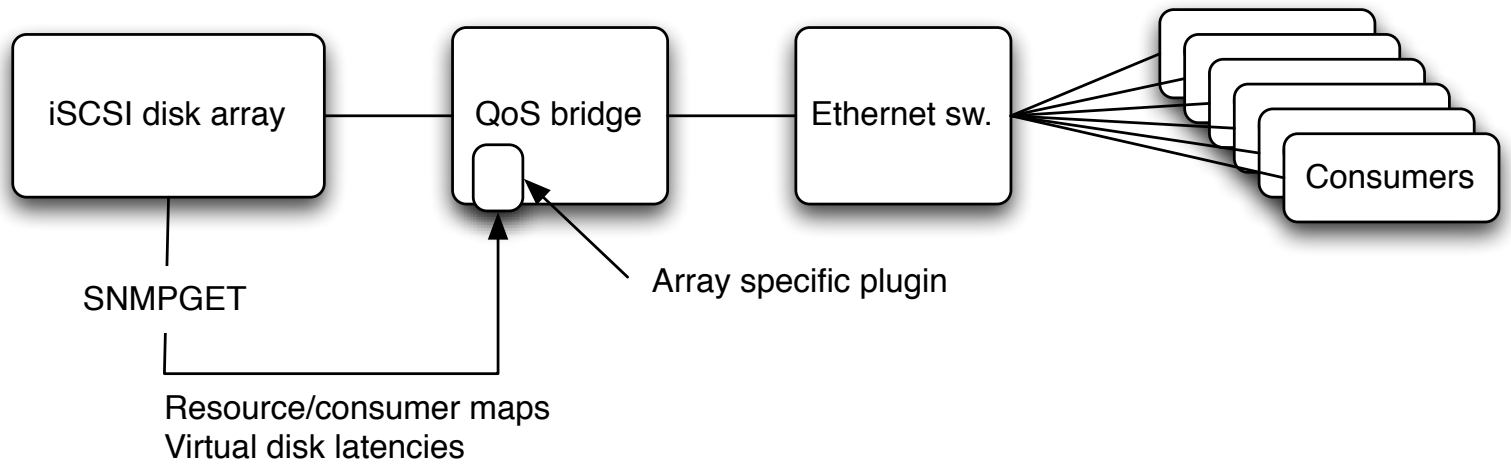

- Packet delay throttle with other algorithms
- PID controller with other throttles

Thanks for the attention !

Overhead

- Negligeble overhead introduced by TC filters
- Differences measured 20 times
- t-test 99% confidence shows 0.4% / 1.7 %• overhead for read/write (worst case)

Is response time improved by throttling ?

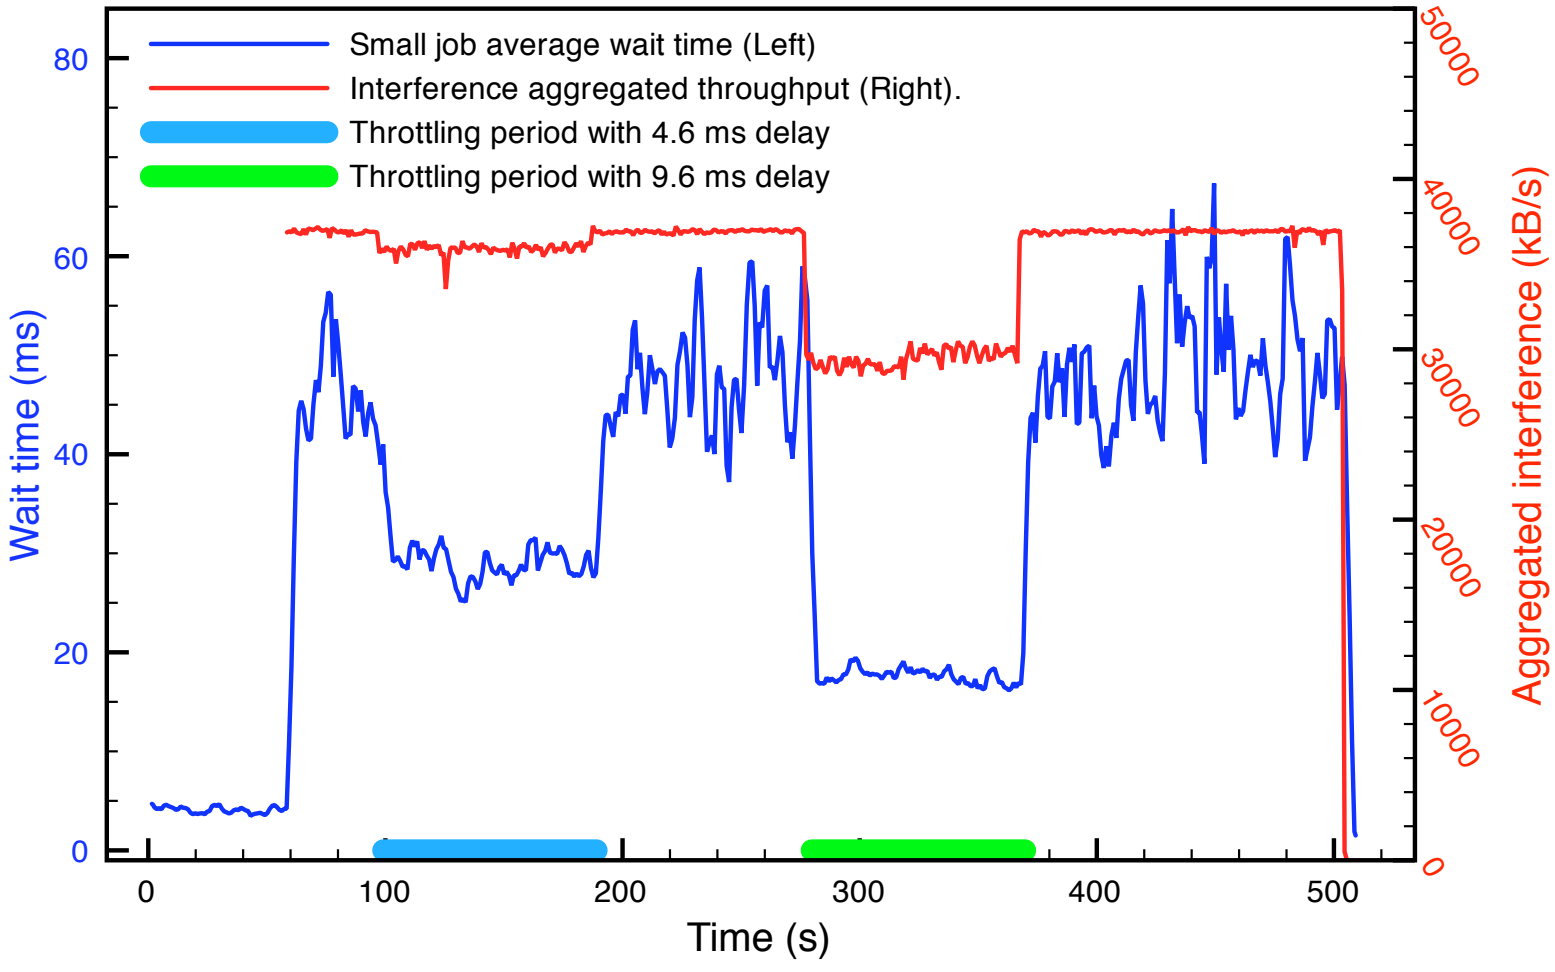

Automatically controlled wait time

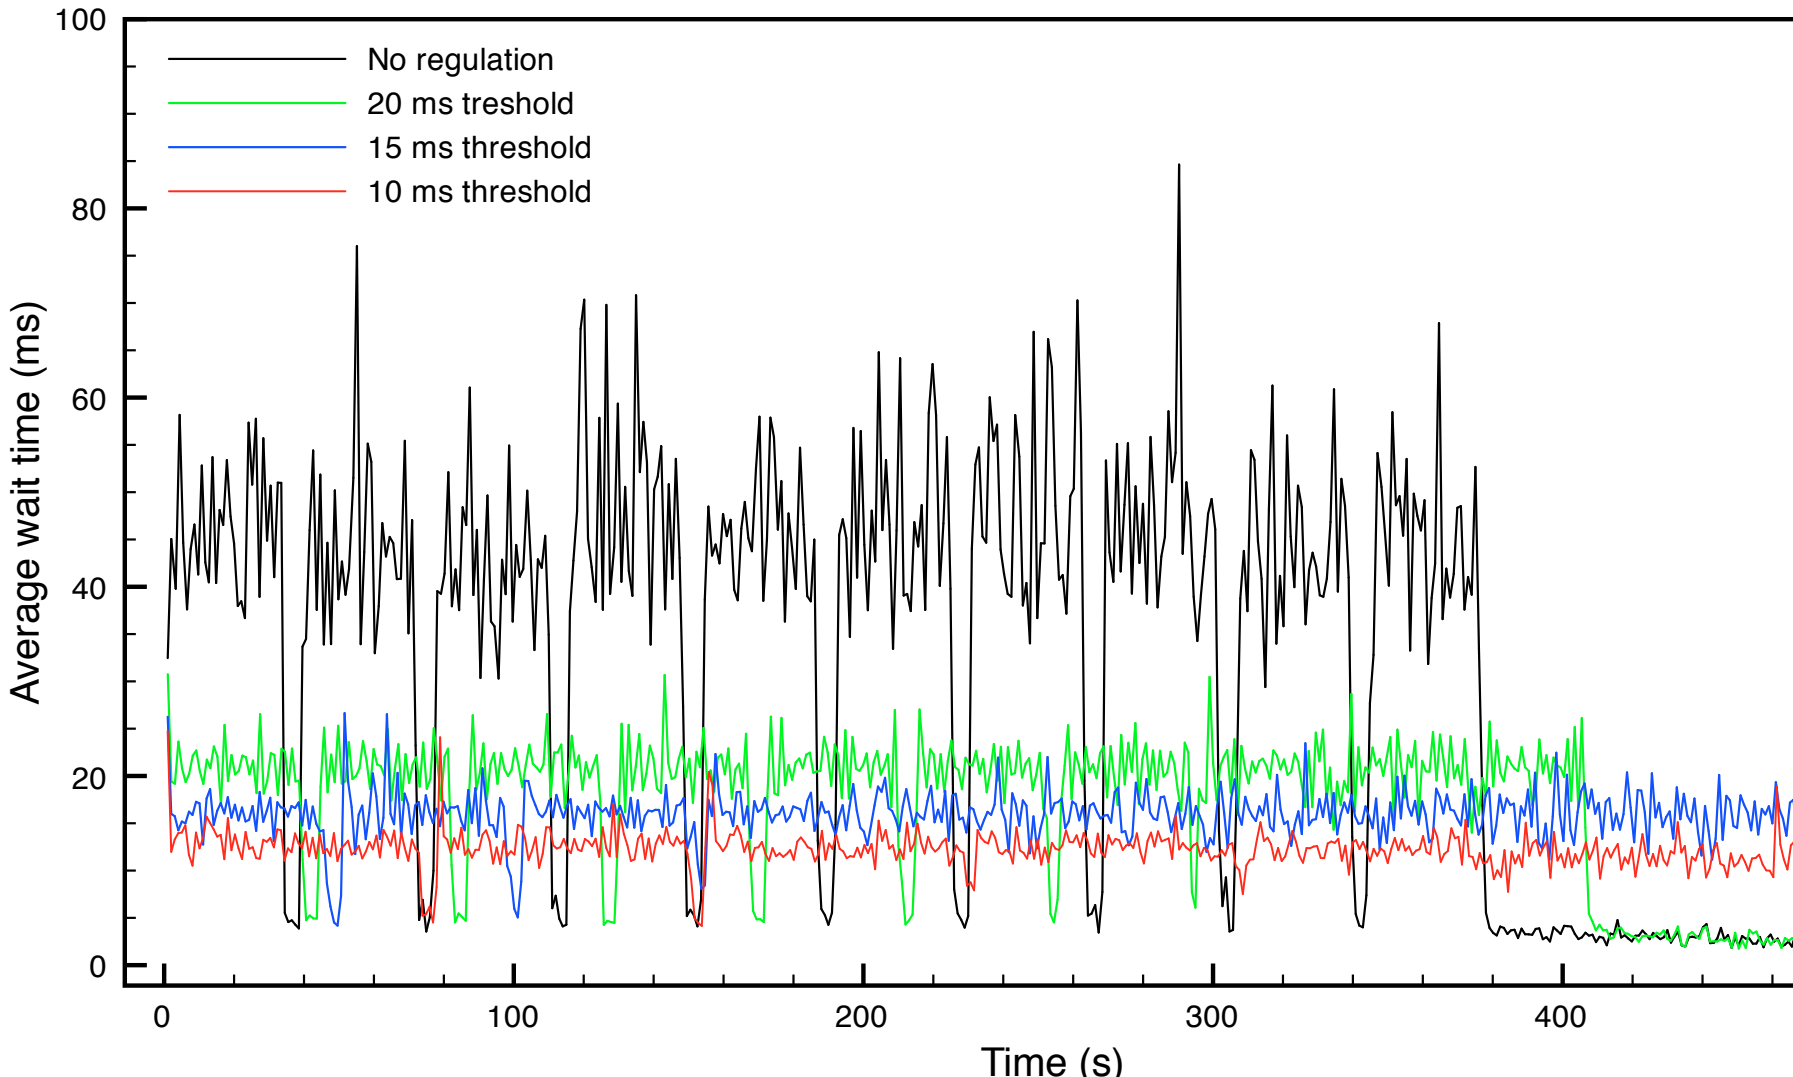

The throttled rates

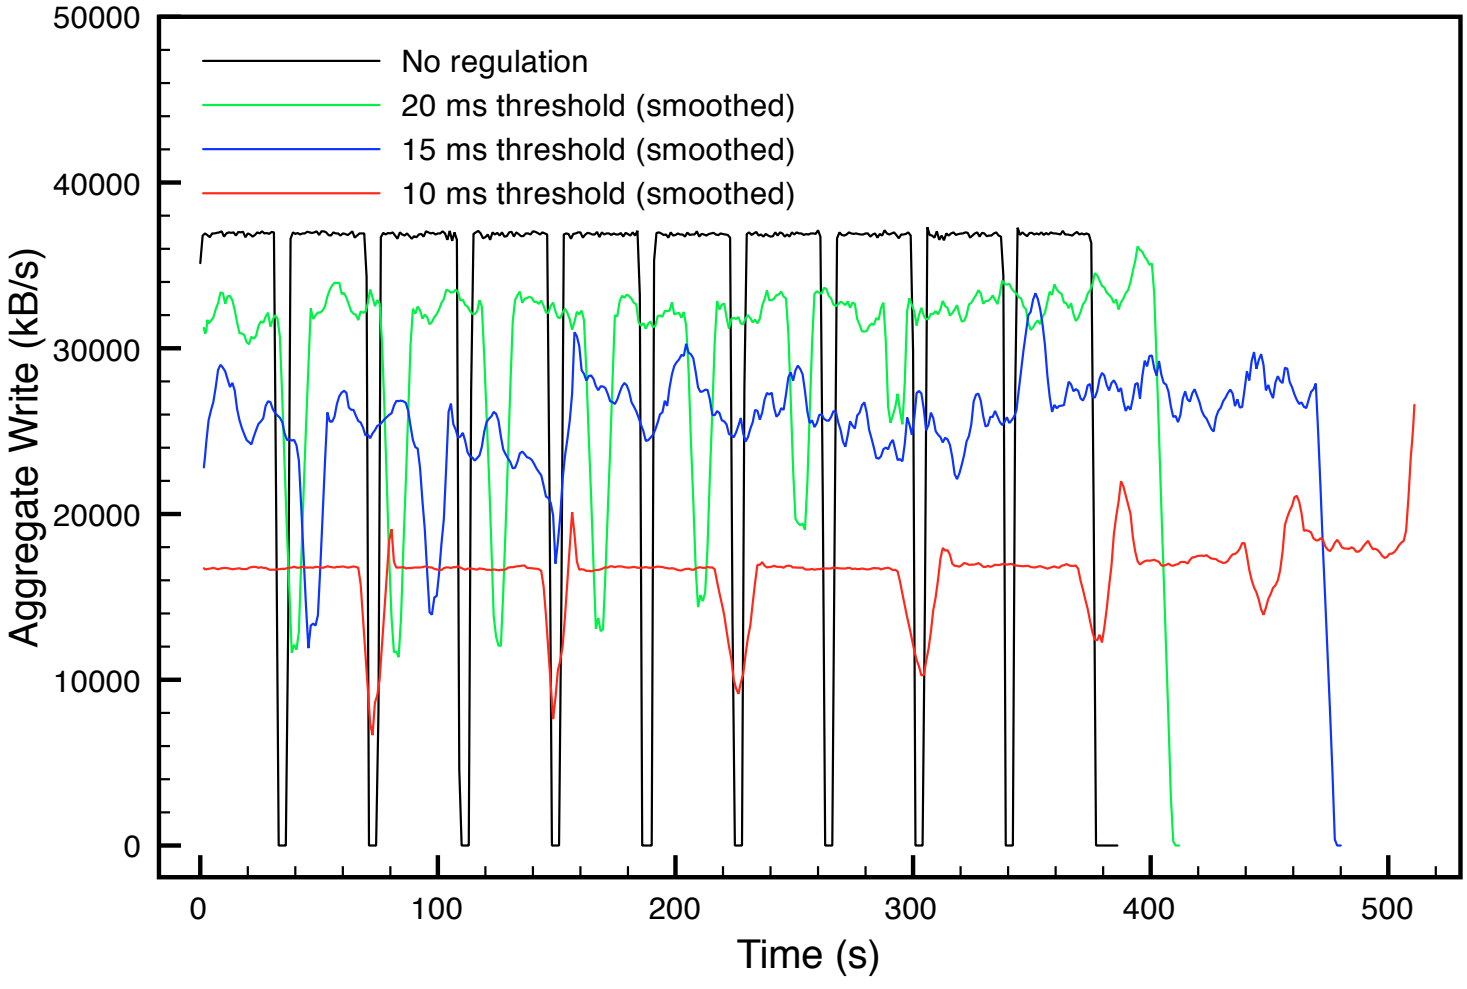

Exposing the throttling value

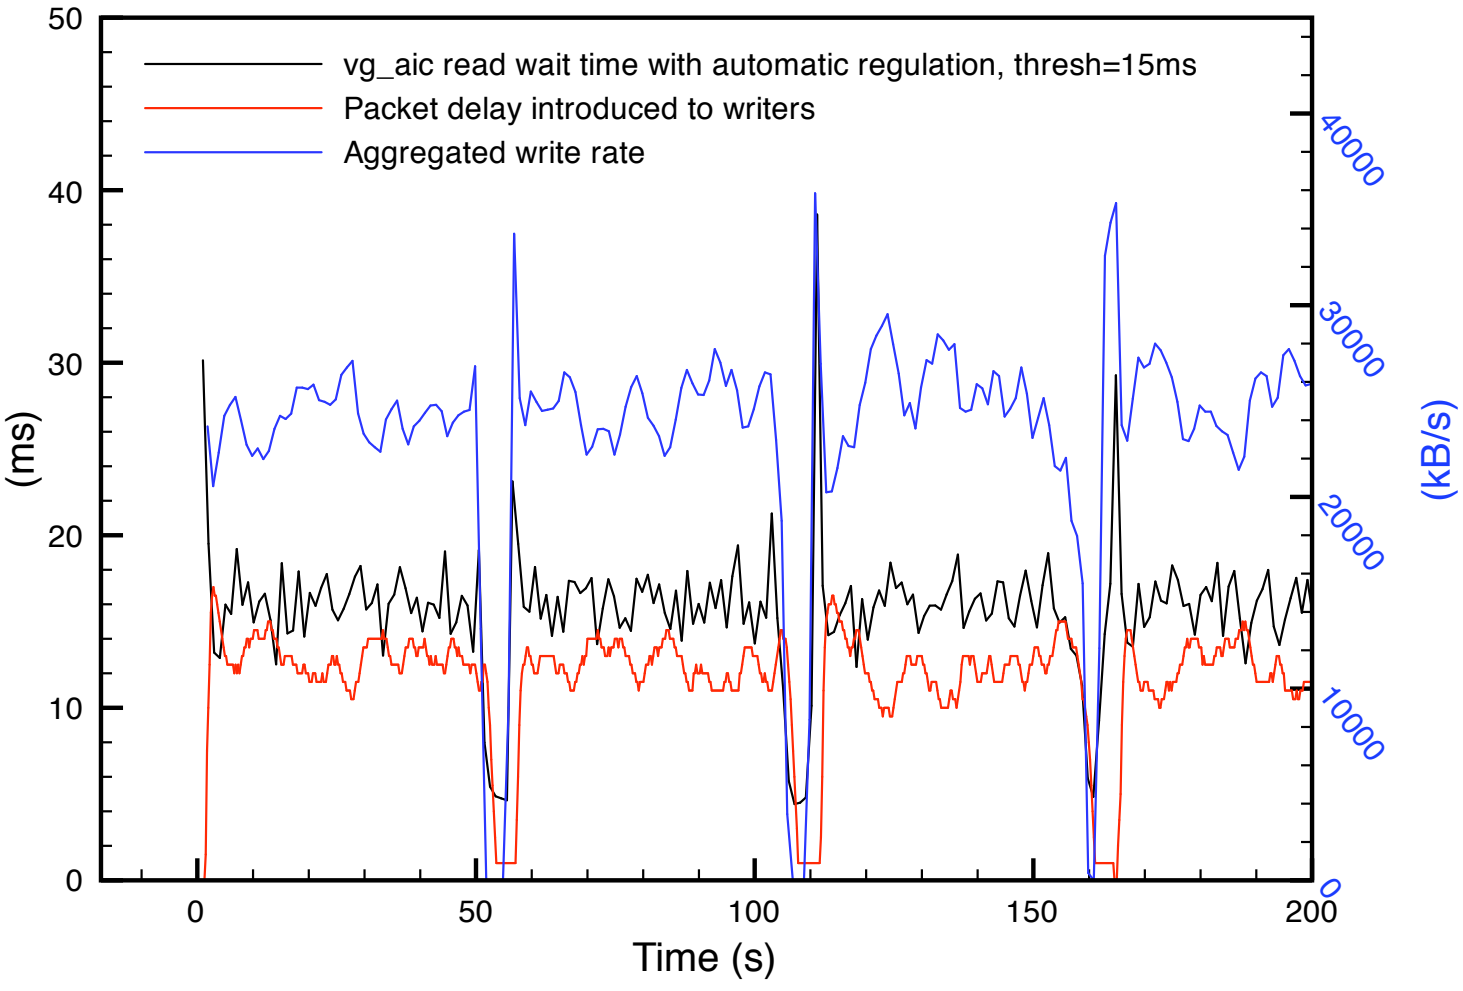

Effect of the packet delay throttle: Reads

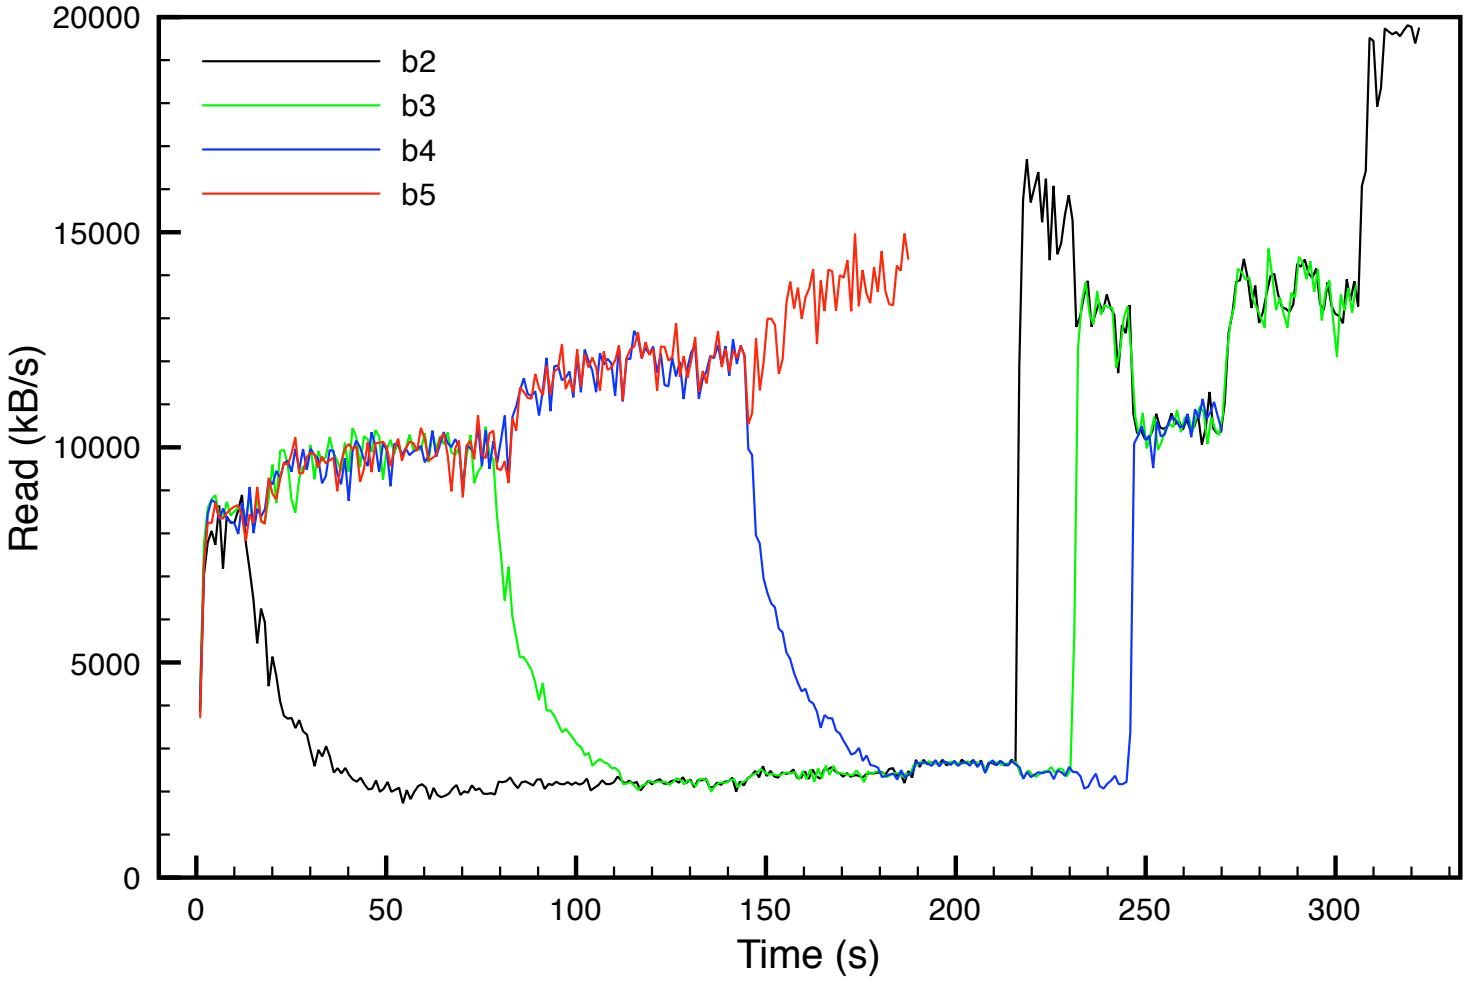

Effect of the packet delay throttle: Writes

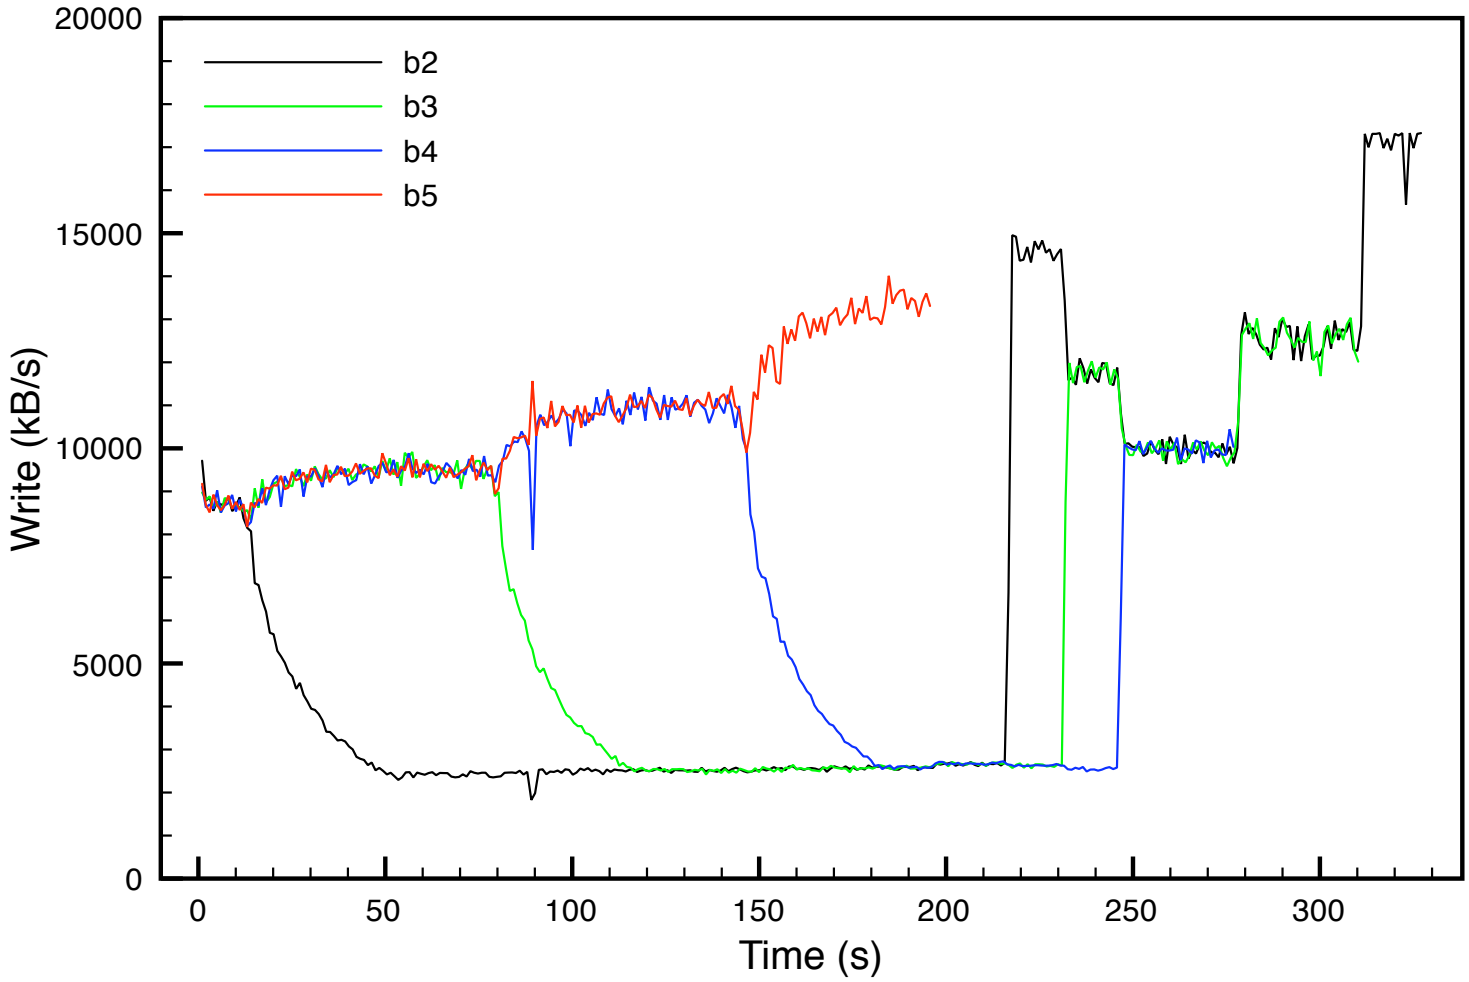

The tc delay queues

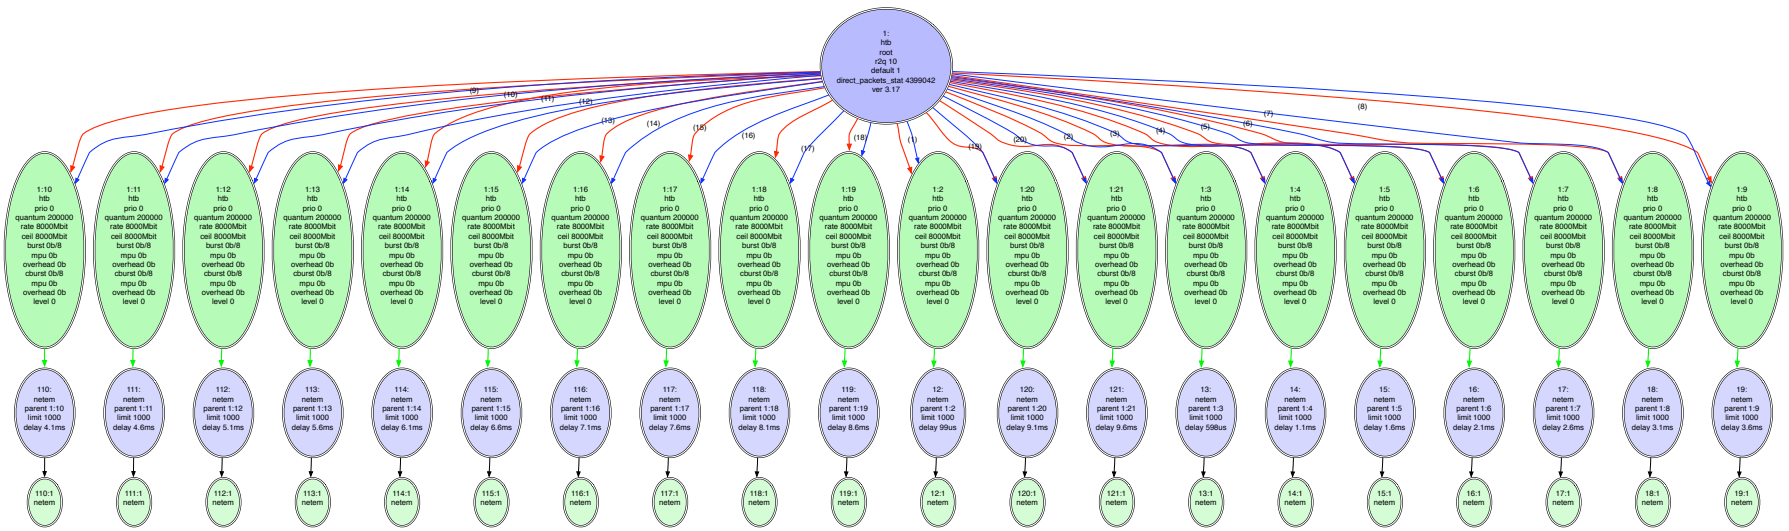

The tc bandwidth queues

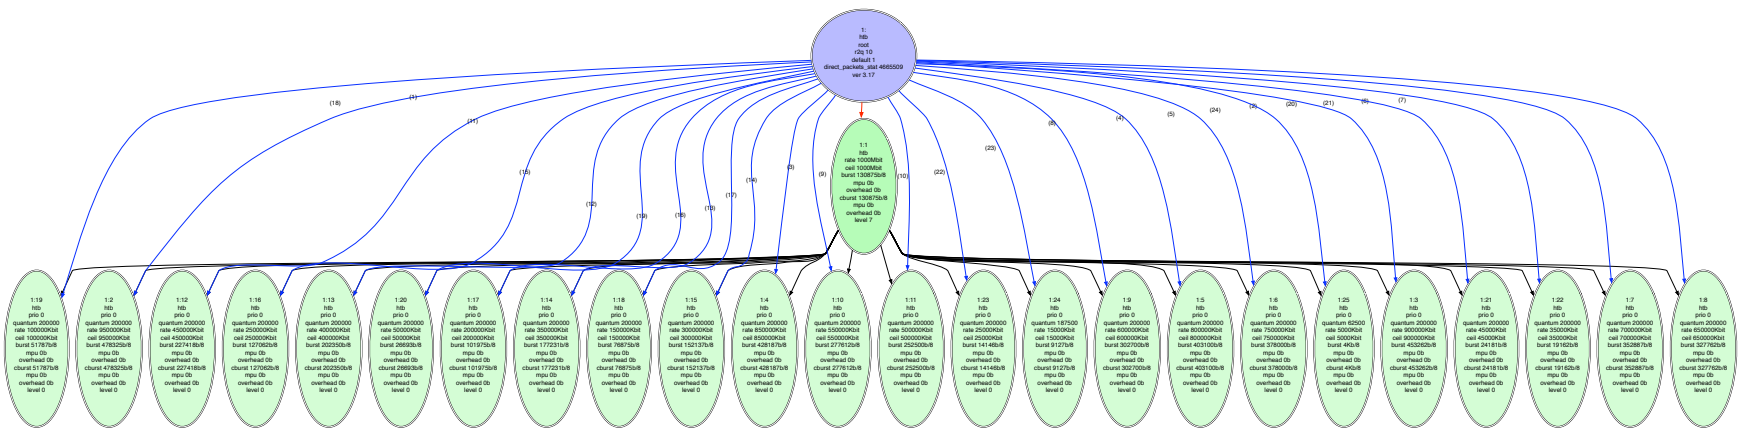

Input signal

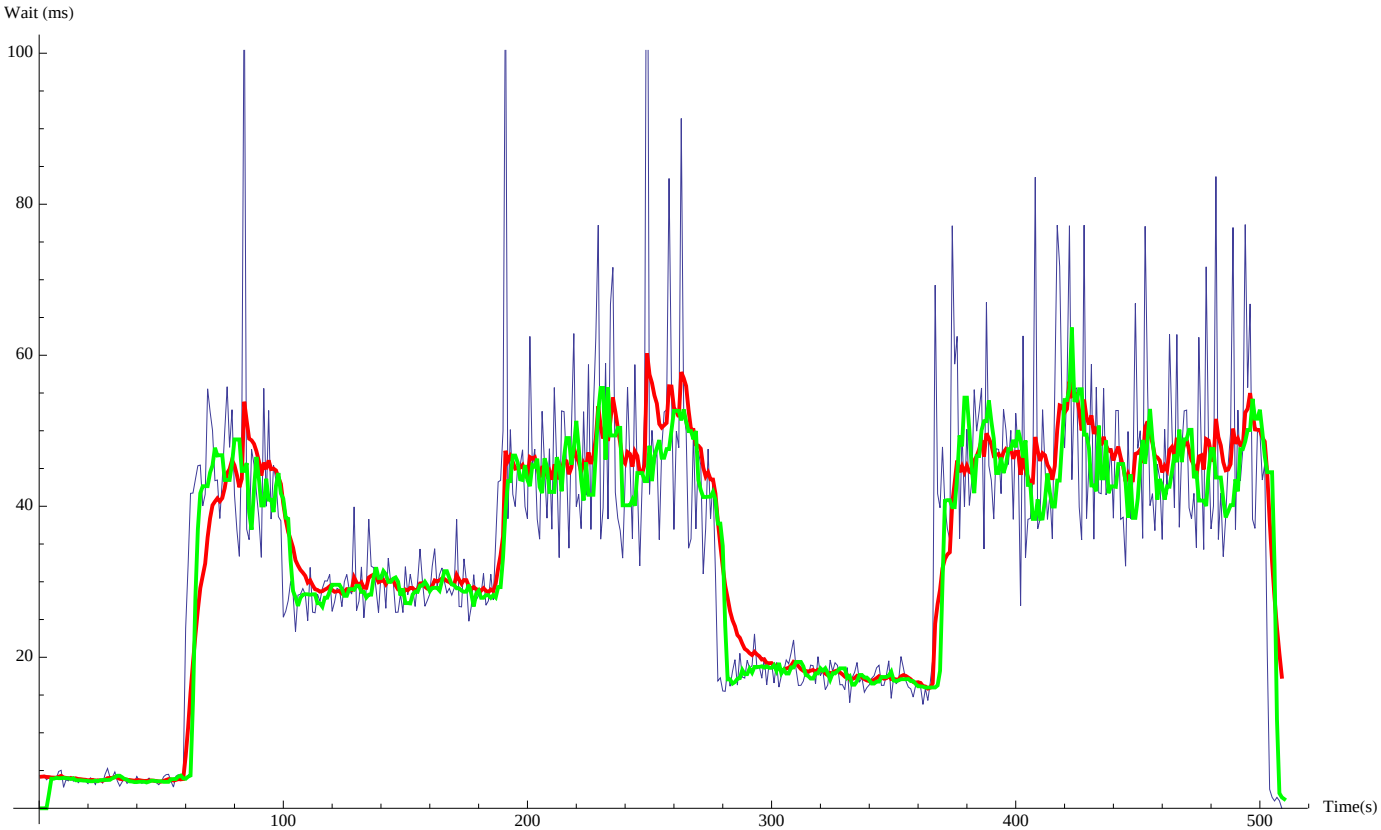

- Red: Exponential Weighted Moving Average (EWMA)
- Green: Moving median
- $L_{(t)} = I_{(t)}\alpha + L_{(t-1)}(1-\alpha)$
- EWMA, also called low pass filter

$$u(t) = \overbrace{\underbrace{K_p e(t)}_{\text{Proportional}} + \underbrace{\frac{K_p}{T_i}\int_0^t e(\tau)d\tau}_{\text{Integral}} + \underbrace{K_p T_d e'(t)}_{\text{Derivative}}}^{\textit{Continous}}$$

$$u_k = \underbrace{\underbrace{u_{k-1}}_{\textit{Previous}} + \underbrace{K_p(1+\frac{T}{T_i})e_k - K_p e_{k-1} + \frac{K_p T_d}{T}(e_k - 2e_{k-1} + e_{k-2})}_{\textit{Delta}}}_{\textit{Incremental form}}$$

$$u_k = \underbrace{\underbrace{K_p e_k}_{\textit{Proportional}} + \underbrace{u_{i(k-1)} + \frac{K_p T}{Ti} e_k}_{\textit{Integral}} + \underbrace{\frac{K_p T_d}{T}(e_k - e_{k-1})}_{\textit{Derivative}}}_{\textit{Absolute form}}$$