Enhancing TCP to support Rate-Limited Traffic

ICCRG/TCPM

draft-fairhurst-tcpm-newcwv-05.txt

G. Fairhurst, A. Sathiaseelan

University of Aberdeen

Scotland, UK

http://trac.tools.ietf.org/group/irtf/trac/wiki/ICCRG_newcwv

Congestion Window Validation

RFC2861 had a good motivation (protect the network)

- However, too conservative for apps to benefit.
- Not widely implemented or used.

Propose to obsolete RFC 2861, and define something else.

IETF diff between -03 and -04

ICCRG feedback

Used term rate-limited in all places.

Added justification and minor changes suggested on the list.

Added text to tie-in with more accurate ECN marking.

Added ref to Hug01 (but did not specify pacing)

IETF diff between -04 and -05

Fixed issue for infrequent large bursts:

- Non-Validated Period (NVP)
- Introduced pipeACK, to replace FlightSize
- This reflects actual acknowledged usage

Changed NVP entry to pipeACK < ½*cwnd

Changed NVP exit conditions:

- pipeACK leaves NVP after pipe was acknowledged.
- Removed need for hysteresis.

Key Features

Differentiate between Validated & Non-Validated Phases

Validated: Standard behaviour

Non-Validated: Updated behaviour

- ssthresh adjusted
- Different rate reduction for loss (D-R)/2
- cwnd does not increase
- cwnd decreases after NVP

cwnd tracks pipeACK (I)

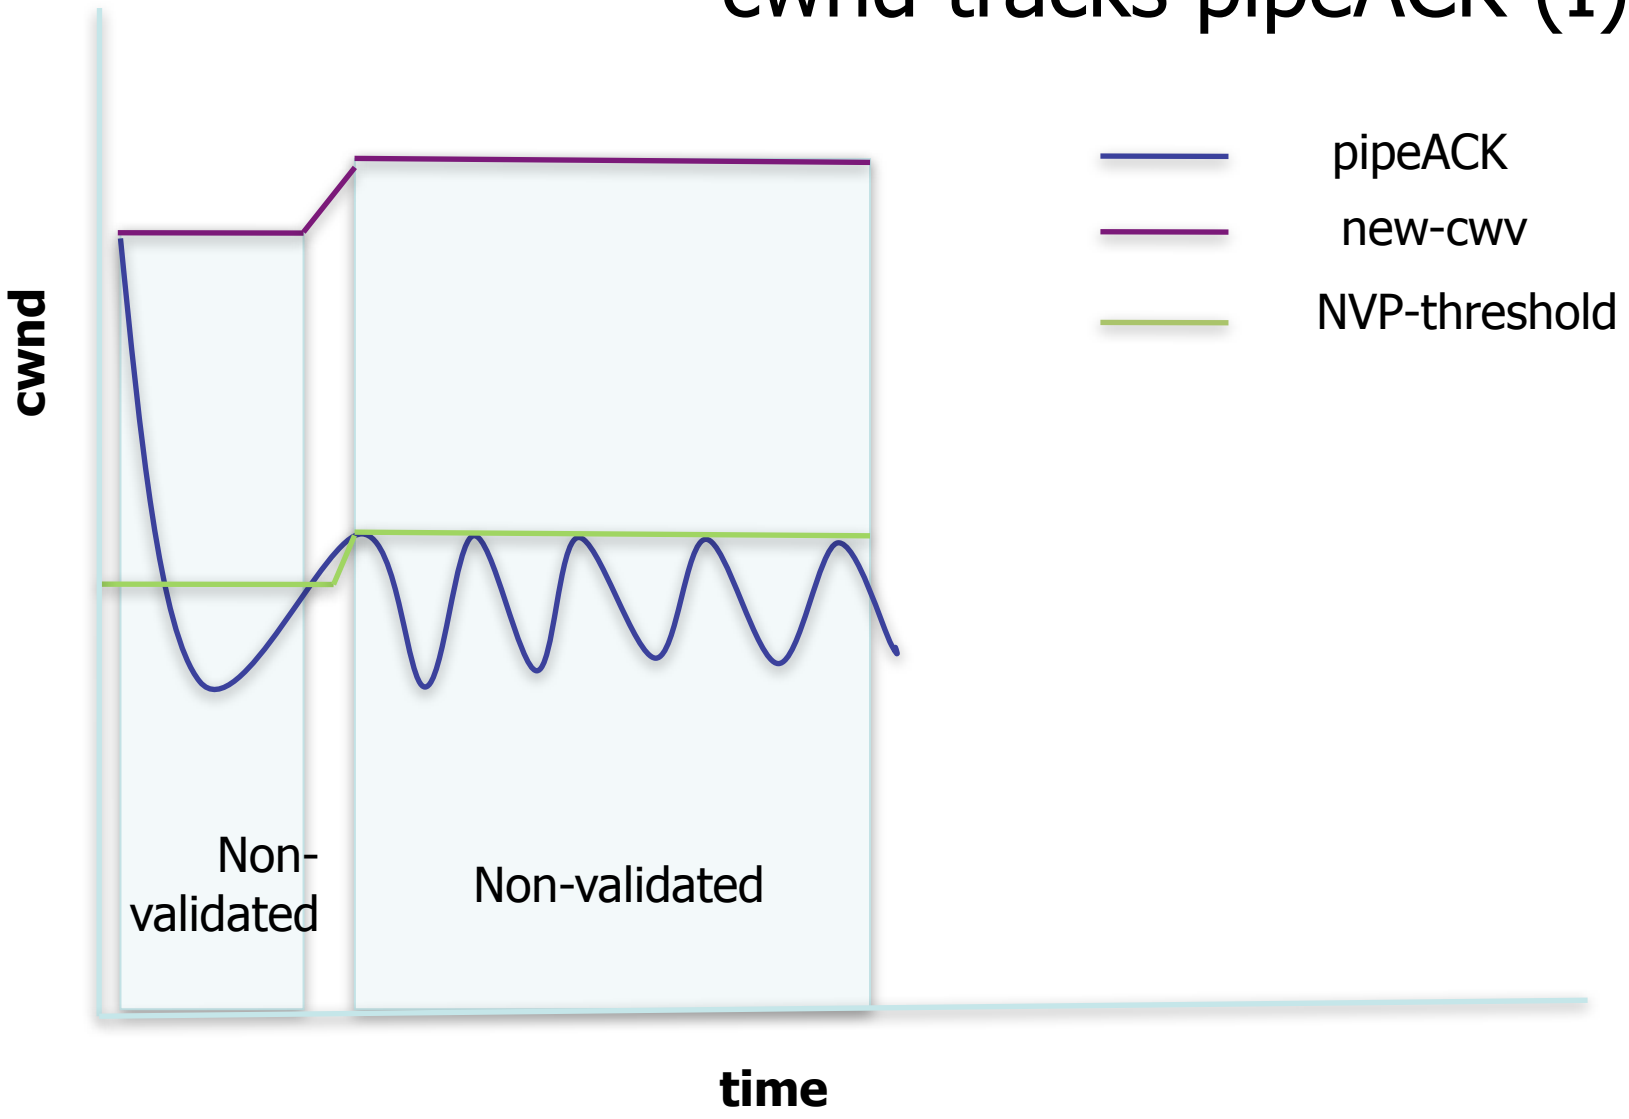

Varying pipeACK
With new-cwv, the cwnd does not grow beyond 2*pipeACK

cwnd tracks pipeACK (II)

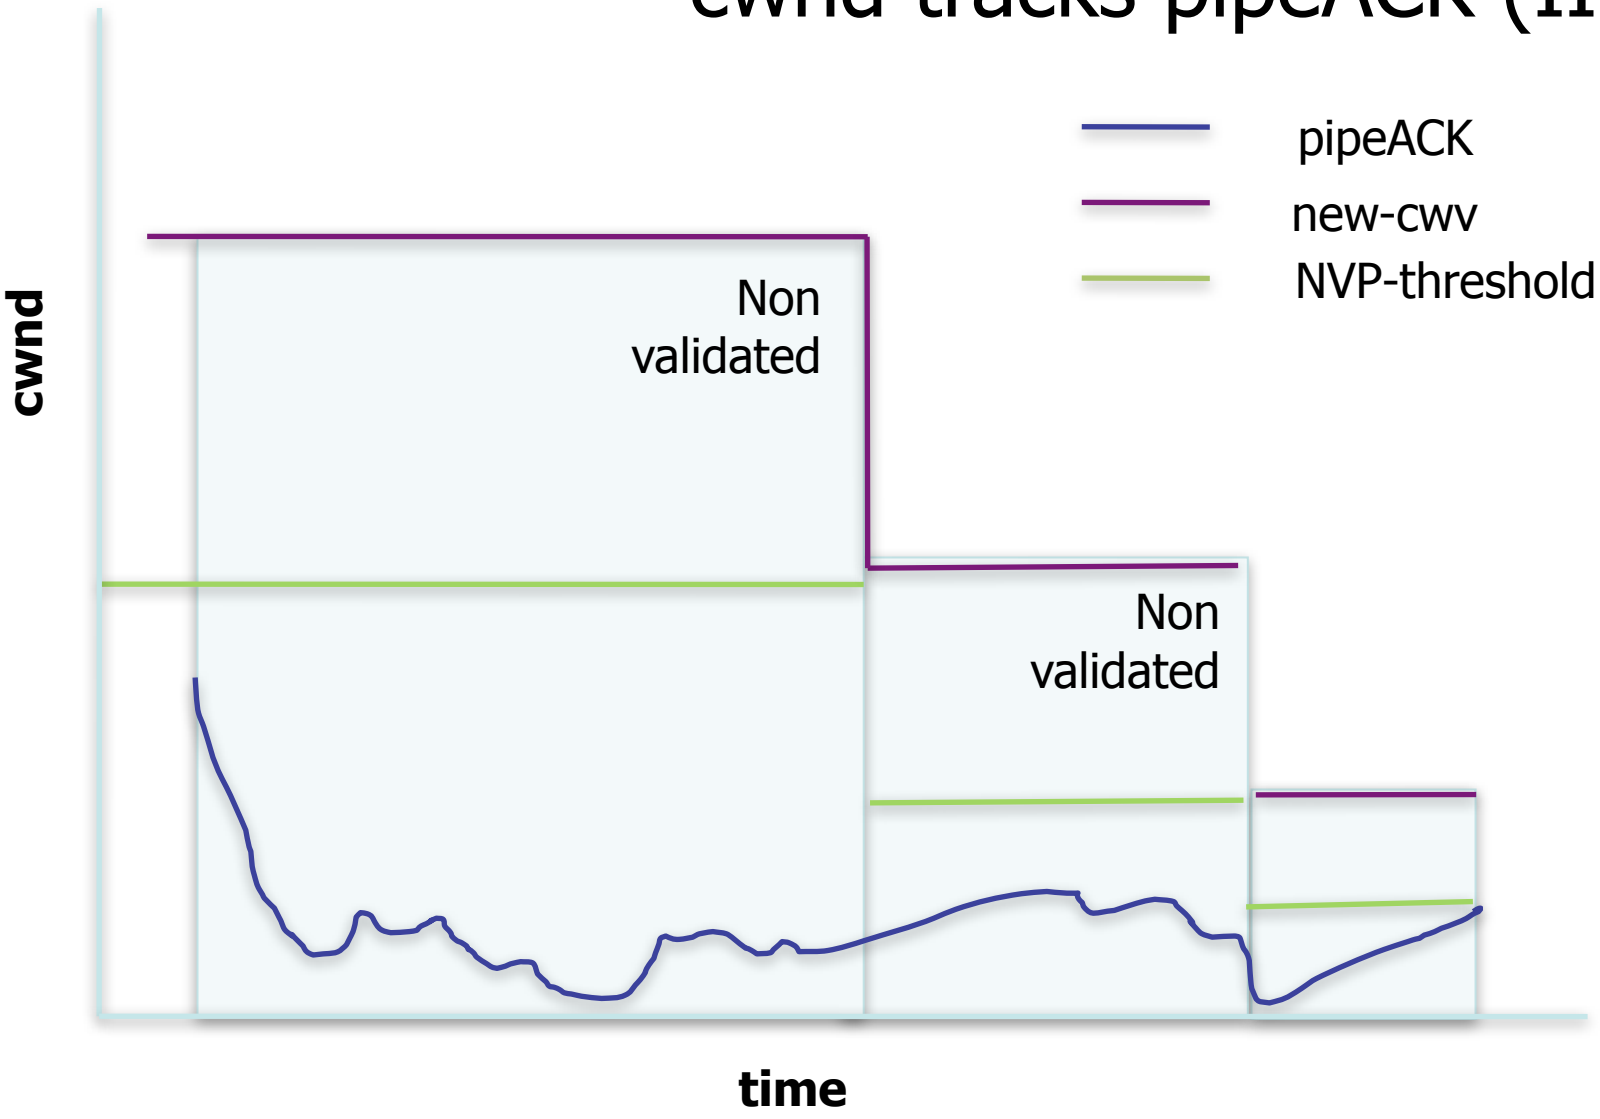

Varying pipeACK (around ½ cwnd)
new-cwv behaviour reduces cwnd after 5 minutes by ½

cwnd tracks pipeACK (III)

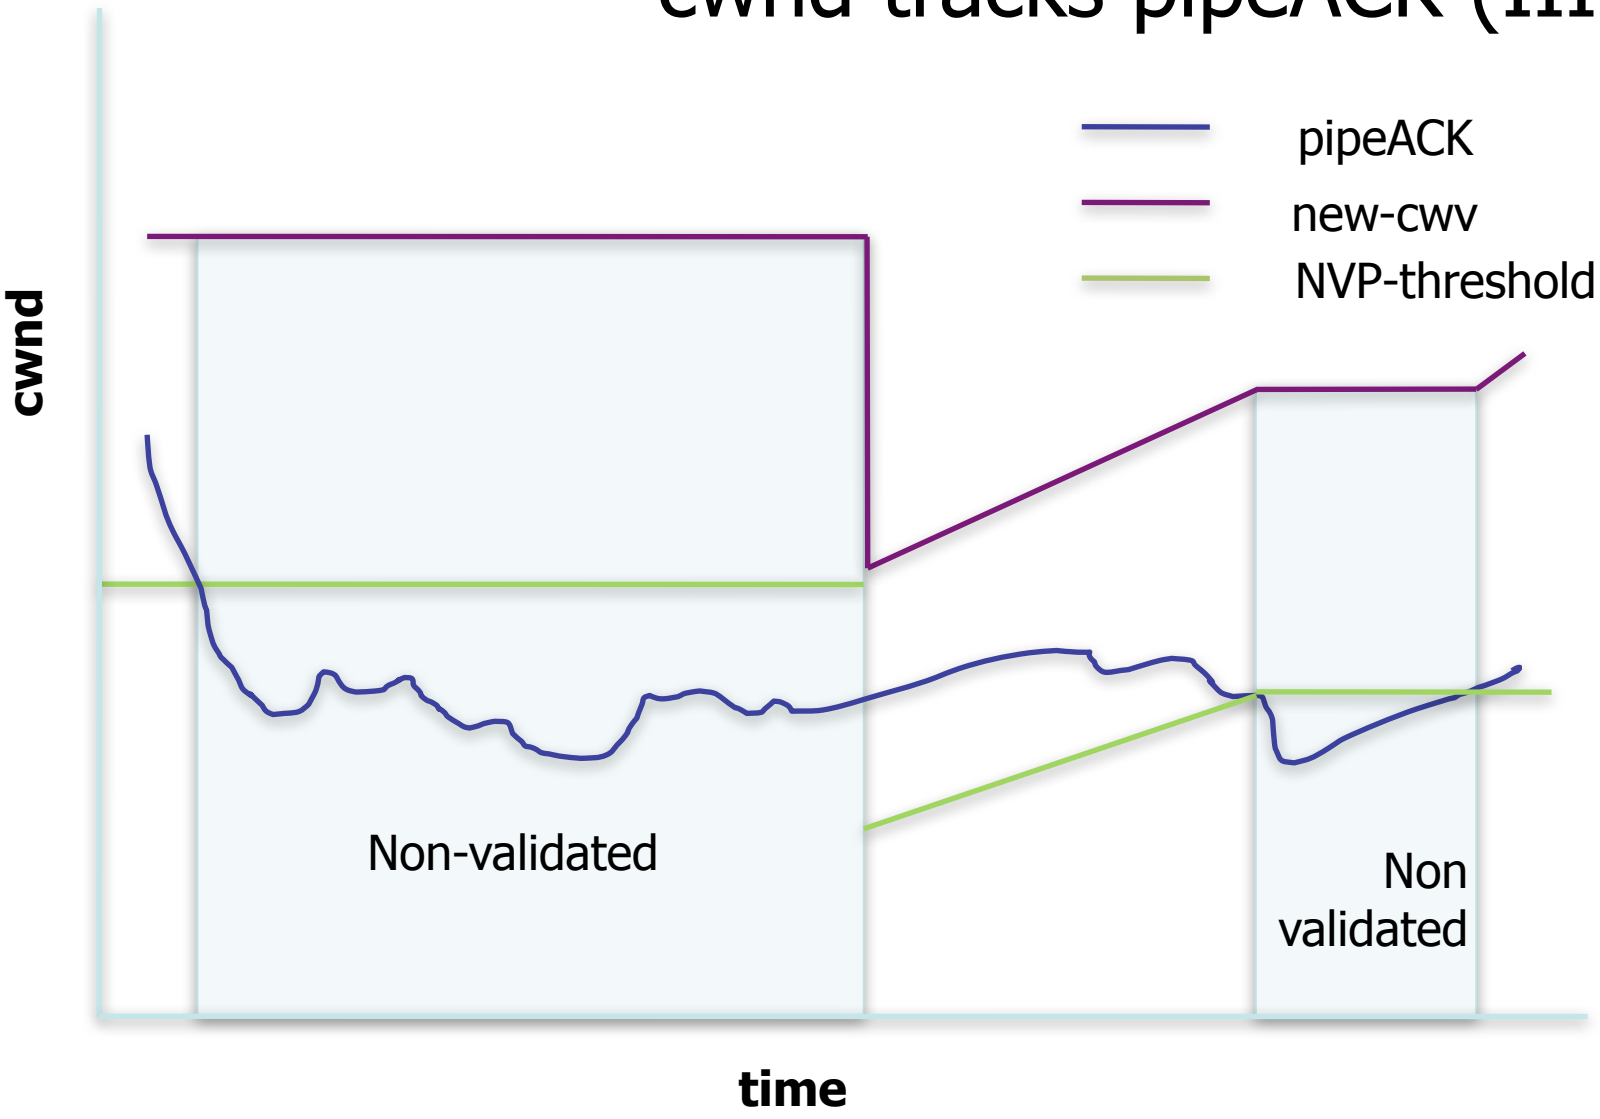

Varying pipeACK (around ½ cwnd)
new-cwv behaviour tracks pipe

Why is NVP 5 mins?

There is no “magic number”

Characteristic idle periods ~ few secs to few minutes

Network paths are relatively stable for several minutes

TCP default user timeout of 5 minutes - how long transmitted data may be unacknowledged before closed.

Expected to be sufficient for common apps

Updates planned for -06

Require reset of pipeACK after congestion

Added comment on effect of congestion after a short burst (M. Allman)

Correction of minor typos to improve consistency

Next Steps

We think this is a problem we should address

Outstanding issues:

- IW has similarities, but is different (see draft)
- Laminar is different (but this proposal for STD TCP)
- Tail loss can also be an issue for bursty apps

We think this is a useful starting point for this work

We would like to see this adopted as a TCPM work item!

http://trac.tools.ietf.org/group/irtf/trac/wiki/ICCRG_newcwv