Careful convergence of congestion control from retained state with QUIC

draft-kuhn-tsvwg-careful-resume

Nicolas Kuhn (Thales Alenia Space)

Emile Stephan (Orange)

Gorry Fairhurst (University of Aberdeen)

Christian Huitema (Private Octopus Inc.)

Catching up with the activity

High BDP paths need a quick ramp-up to speed

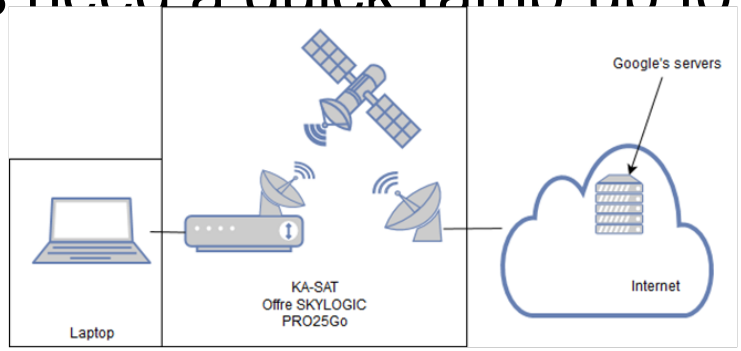

Target (1 object, 5.3MB)

Stores previous path characteristics (rate, rtt, etc).
Decides to use information to initialise a new connection
Careful to « undo » if any problem is found

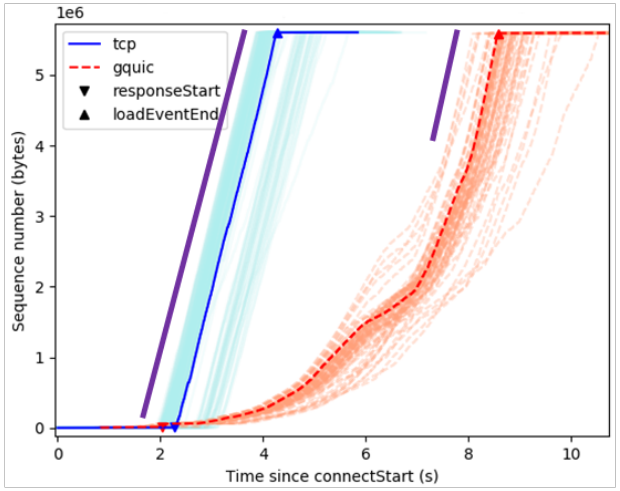

Google QUIC performance over a public SATCOM access
International Journal of Satellite Communications and Networking
THOMAS, L. ; DUBOIS, E. ; KUHN, N. ; LOCHIN, E. 2019

Changes in draft-kuhn-tsvwg-careful-resume-00

draft-kuhn-quic-0rtt-bdp

- concept developed

draft-kuhn-quic-careful-resume

- defined states

draft-kuhn-tsvwg-careful-resume-00

- major rewrite of the spec
- refactored content between careful-resume as a CC method & BDP Frame for receiver signalling

Thanks to comments/text from: John Border, Gabriel Montenegro, Patrick McManus, Ian Swett, Igor Lubashev, Robin Marx, Roland Bless,Franklin Simo, Tom Jones.

draft-kuhn-tsvwg-careful-resume Objectives

- Addresses careful convergence of Congestion Control (CC) in QUIC, providing a cautious method to enable fast startup for wide range of connections

- Reuses a set of computed CC parameters (bandwidth, capacity, RTT)
  - Based on previously observed path characteristics between the endpoints

- Defines requirements
  - How a sender utilizes parameters to provide opportunities for a new connection to more quickly get up to speed (i.e. utilize available capacity)
- Discusses how changes impact capacity at a shared network bottleneck and a safe response (e.g. in the case of inappropriate rate)

Careful Resume Phases

- **Observe Phase**
  - Information obtained from a previous session
- **Reconnaissance Phase**
  - Check consistency of previous session information
- **Unvalidated Phase**
  - Fast (but not instantaneous) convergence towards previous rate
- **Retreat Phase**
  - Stop using previously measured parameters
  - Re-initialised CC parameters to drain any queue built at the bottleneck during the Unvalidated Phase
- **Normal Phase**
  - Resume standard CC method

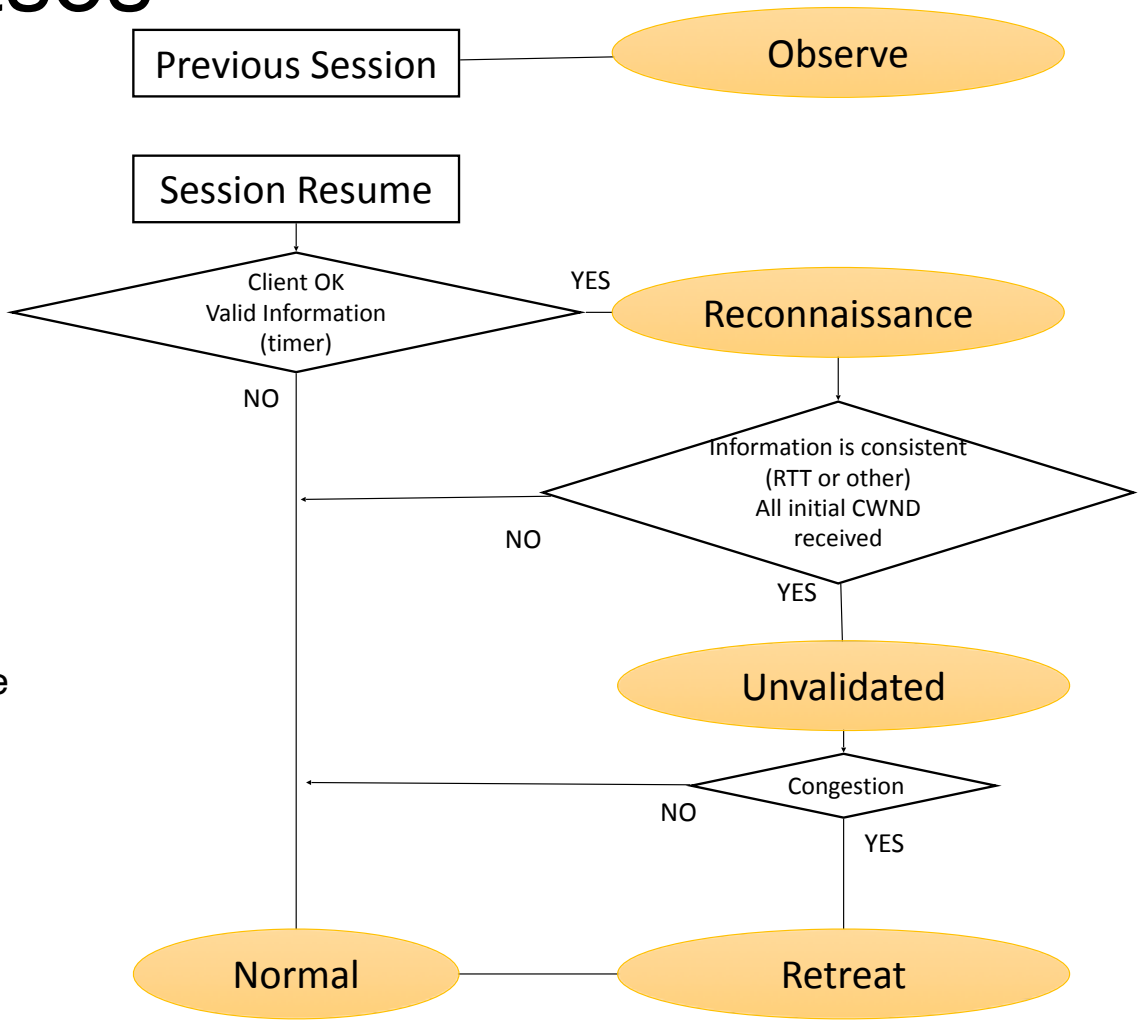

"Normal" Slow Start

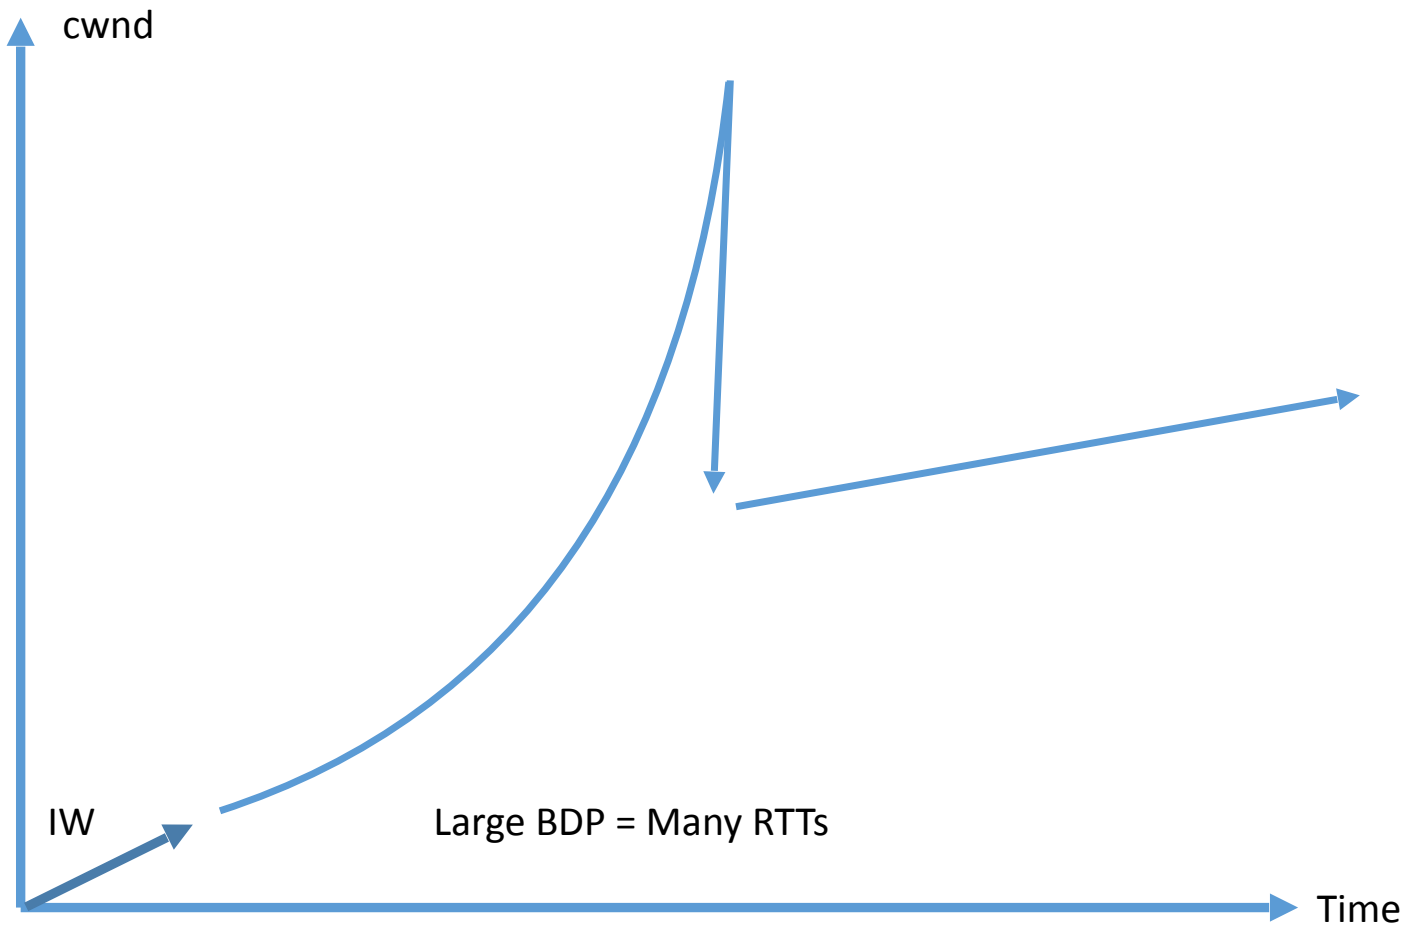

Why we need Discussion on Recommendations

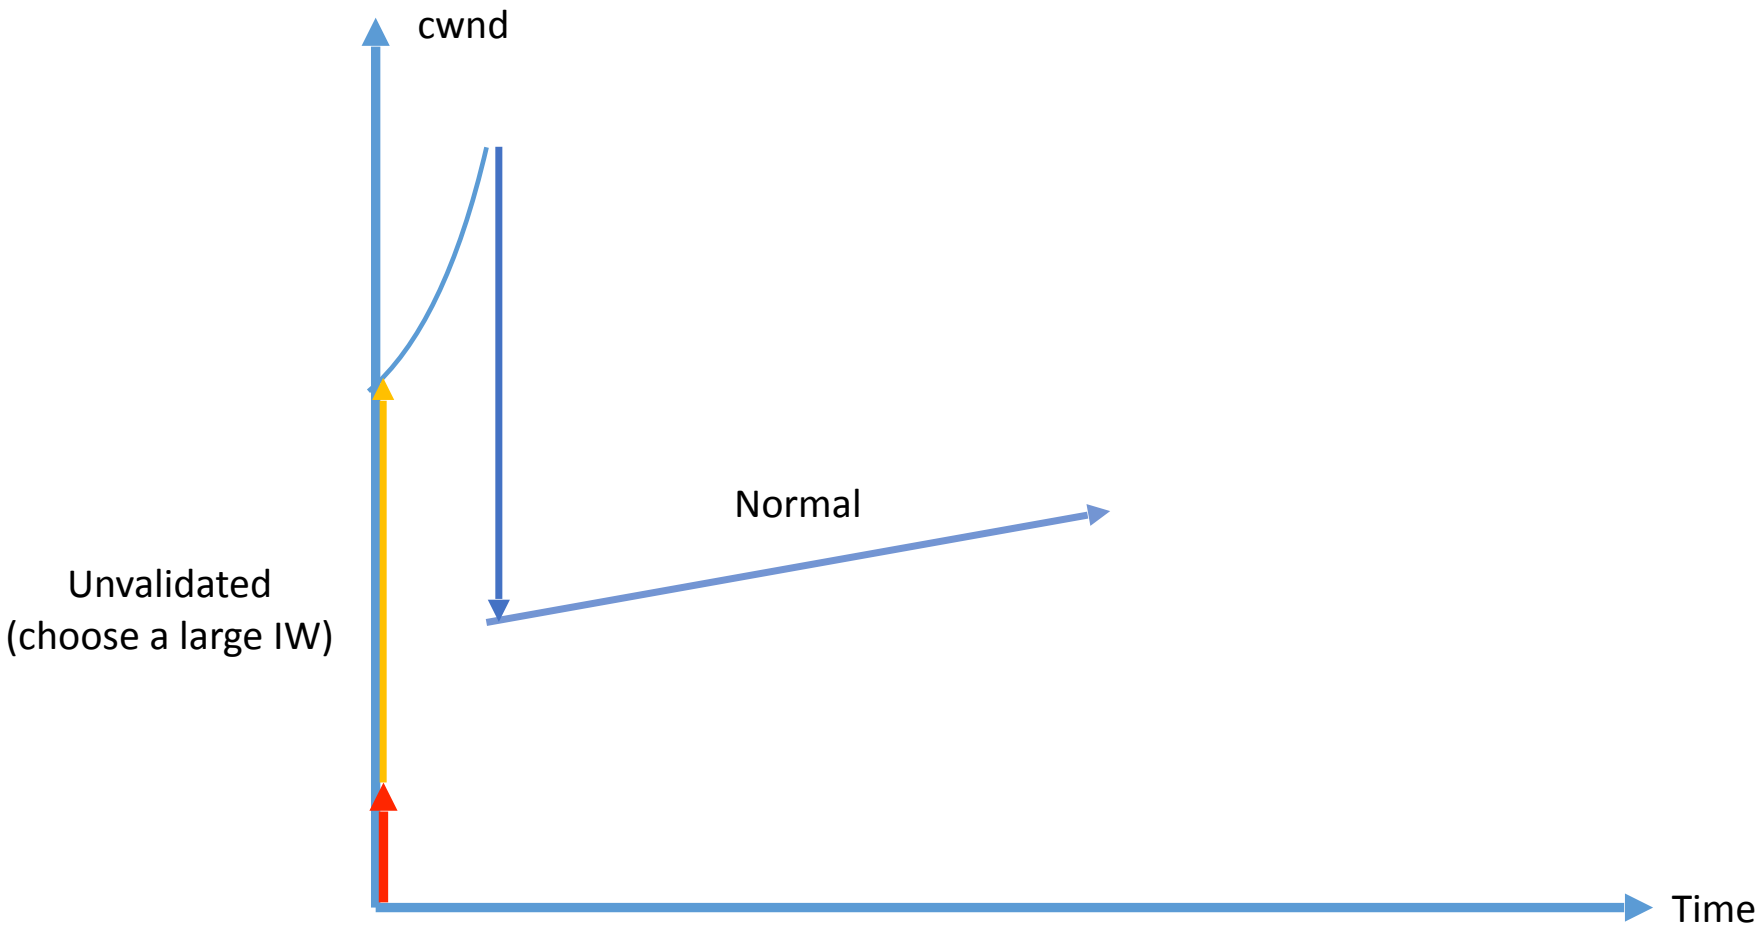

Careful Resume

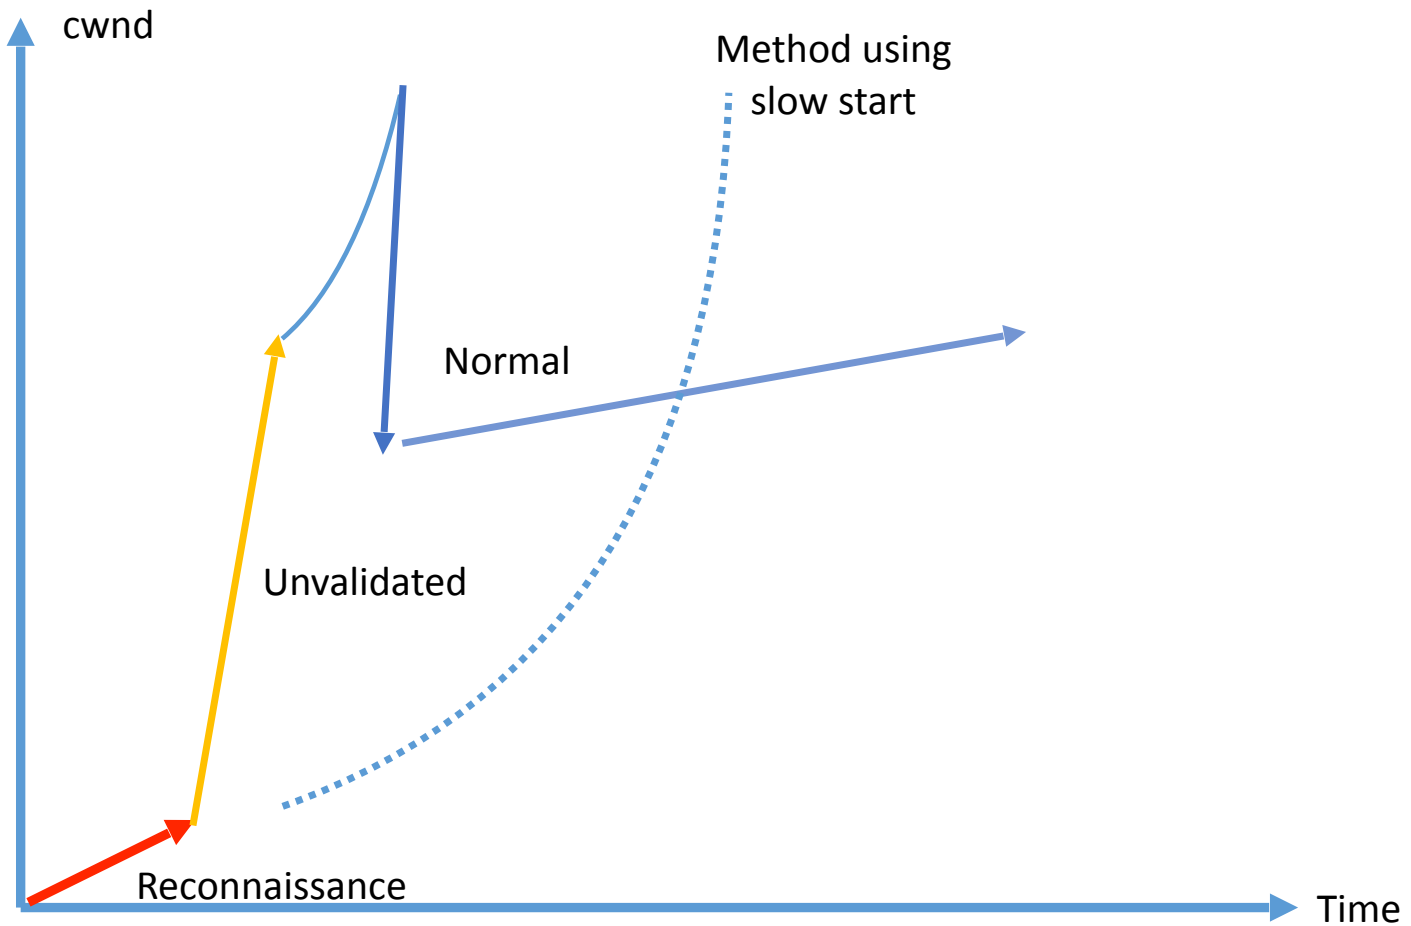

Careful Resume with Safe Retreat

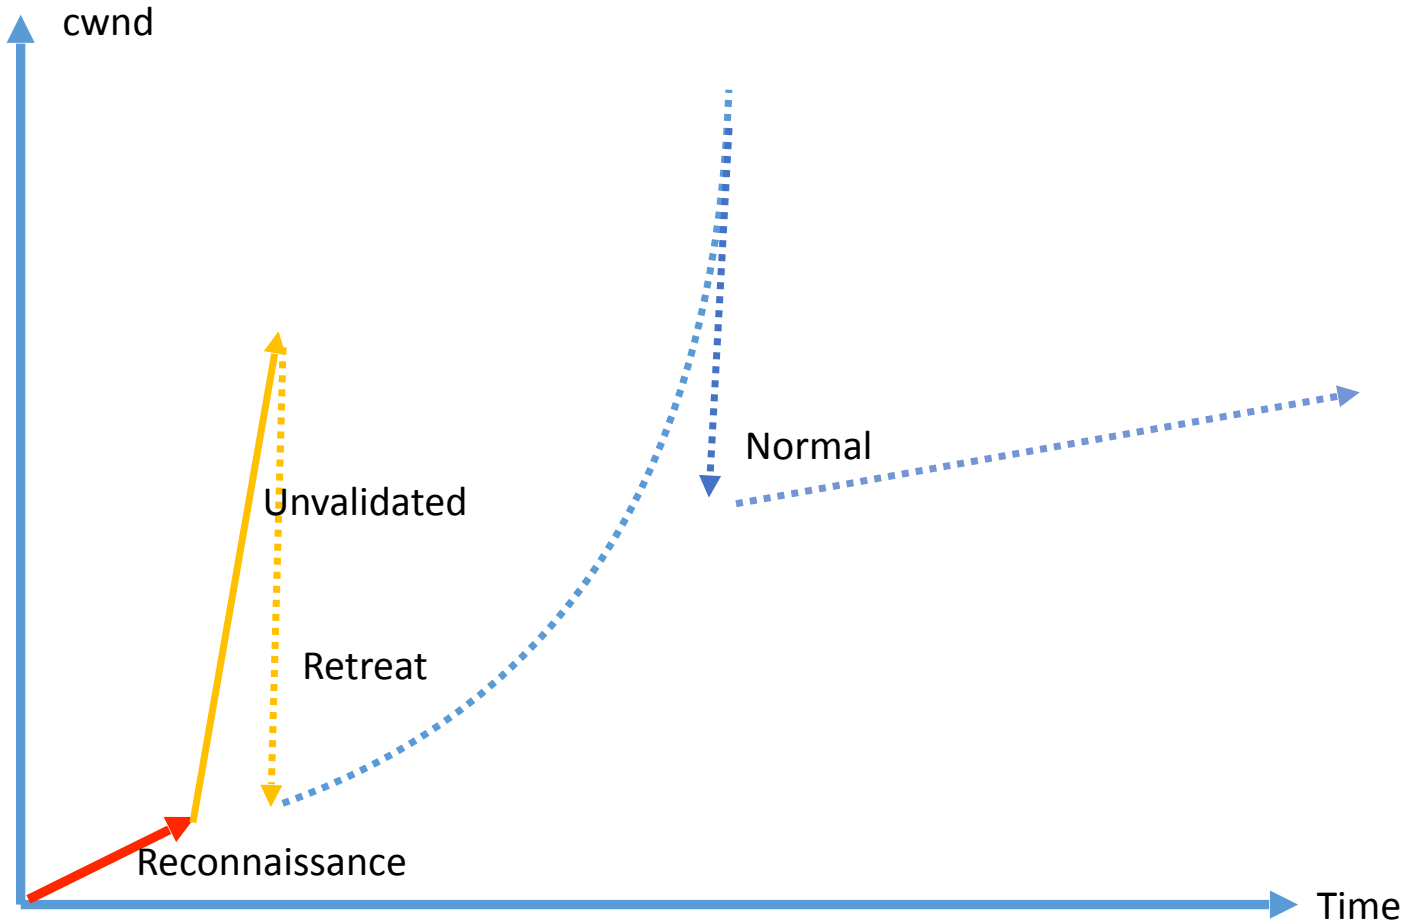

Comments on Validity of Saved Information

- *The sender MUST check the validity of any received saved_rtt and saved_bb parameters.*
  - *Might the path have changed?*
  - *Might the capacity have changed?*

QUIC BDP Frame Extension

draft-kuhn-quic-bdpframe-extension-01

Nicolas Kuhn (Thales Alenia Space)

Emile Stephan (Orange)

Gorry Fairhurst (University of Aberdeen)

Christian Huitema (Private Octopus Inc.)

draft-kuhn-quic-bdpframe-extension

- Builds on careful resume to allow:
  - The BDP information to be stored at the receiver
  - Releases the sender from needing local storage
  - Enables a receiver to advise when to use the method

- Reviews approaches (local storage, NEW_TOKEN, BDP_FRAME)
  - Focus on implementation details of BDP_FRAME in QUIC
  - Proposes a method for the server to protect the integrity of the BDP information returned by the receiver to the server

Discussion

- Is there interest in developing the careful resume method in TSVWG ?
- Will the WG consider adoption of draft-kuhn-tsvwg-careful-resume ?

Q&A

Additional material

draft-kuhn-tsvwg-careful-resume-00 Guidelines and requirements

- Observation Phase
  - The sender SHOULD NOT store and/or send CC parameter information related to an estimated bottleneck bandwidth, if the cwnd is not at least four times larger than the IW.
- Reconnaissance Phase
  - The sender MUST NOT send more than the IW in the first RTT of transmitted data [RFC9000].
  - The sender MUST compare the measured transport parameters (in particular current_rtt) of the new session with those of the previous session (in particular RTT).
  - The method MUST NOT be used when the path fails to be validated.
- Unvalidated Phase
  - A new connection MUST NOT directly use the previously measured saved_rtt and saved_bb to simply initialize a new flow to resume sending at the same rate.
  - Careful Resume MUST be robust to changes in network traffic, including the arrival of new traffic flows that compete for the bottleneck capacity.
  - The sender MUST check the validity of any received saved_rtt and saved_bb parameters, whether these are sent by a receiver or are stored at the sender.
  - The sender MUST NOT use the parameters unless the first IW packets when packets are detected as lost or acknowledgments indicate the packets were ECN CE-marked. These are indication of potential congestion and therefore the method MUST NOT be used.
  - The sender MUST implement the retreat method when packets are detected as lost or acknowledgments indicate the packets were ECN CE-marked. These are indication of potential congestion and therefore the method MUST NOT be used.
- And other requirements in the draft!