ETH zürich

Network Security Group

Model-Based Insights on the Performance, Fairness, and Stability of BBR

Simon Scherrer, Markus Legner, Adrian Perrig
ETH Zürich

Stefan Schmid
TU Berlin & Fraunhofer SIT

ANRP Award Talk, IETF 117, San Francisco

The journey of BBR development

2016:
First version presented by Google

The journey of BBR development

2016:
First version presented by Google

2017:
BBR enabled for YouTube

The journey of BBR development

2016:
First version presented by Google

2017:
BBR enabled for YouTube

2022:
40% of downstream Internet traffic uses BBR

The journey of BBR development

2016:
First version presented by Google

2017:
BBR enabled for YouTube

2019:
BBRv2 is presented at IETF

2022:
40% of downstream Internet traffic uses BBR

The journey of BBR development

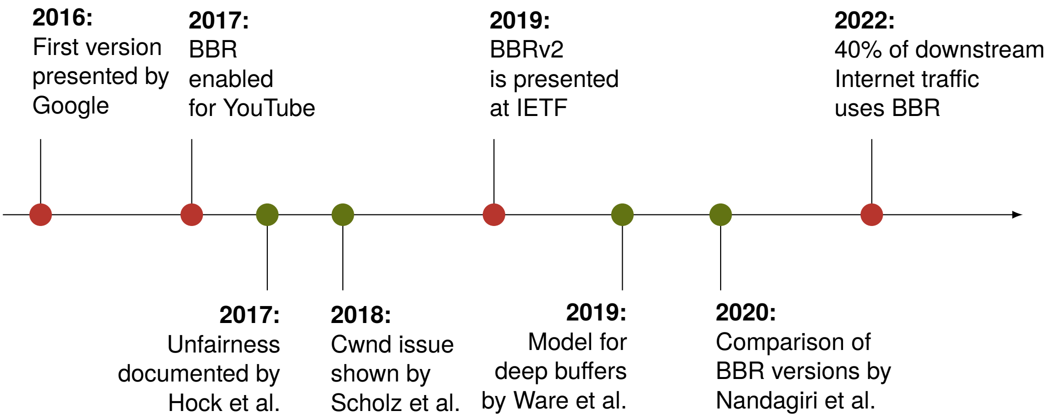

Previous research on BBR leaves open important questions

The approaches of prior research have limitations:

Previous research on BBR leaves open important questions

The approaches of prior research have limitations:

Experimental evaluations

Previous research on BBR leaves open important questions

The approaches of prior research have limitations:

Experimental evaluations

Scale-dependent cost
Experiment cost may be overwhelming for large-scale networks or high speeds

Previous research on BBR leaves open important questions

The approaches of prior research have limitations:

Experimental evaluations

Scale-dependent cost
Experiment cost may be overwhelming for large-scale networks or high speeds

Steady-state models

Previous research on BBR leaves open important questions

The approaches of prior research have limitations:

Experimental evaluations

Scale-dependent cost
Experiment cost may be overwhelming for large-scale networks or high speeds

Steady-state models

No expression of transient effects
Transient phenomena (e.g., convergence behavior) are ignored, although highly relevant

A fluid model can fill the gaps in BBR analysis

A fluid model can fill the gaps in BBR analysis

Fluid model

System of ordinary differential equations (ODEs) describing the joint dynamics of

A fluid model can fill the gaps in BBR analysis

Fluid model

System of ordinary differential equations (ODEs) describing the joint dynamics of

- the *congestion-control algorithms (CCAs)* running at the sending nodes (sending-rate evolution),

A fluid model can fill the gaps in BBR analysis

Fluid model

System of ordinary differential equations (ODEs) describing the joint dynamics of

- the *congestion-control algorithms (CCAs)* running at the sending nodes (sending-rate evolution),
- and the *network metrics* (evolution of queue length, loss rate, arrival rate, etc.).

A fluid model can fill the gaps in BBR analysis

Fluid model

System of ordinary differential equations (ODEs) describing the joint dynamics of

- the *congestion-control algorithms (CCAs)* running at the sending nodes (sending-rate evolution),
- and the *network metrics* (evolution of queue length, loss rate, arrival rate, etc.).

Fluid models enable:

A fluid model can fill the gaps in BBR analysis

Fluid model

System of ordinary differential equations (ODEs) describing the joint dynamics of

- the *congestion-control algorithms (CCAs)* running at the sending nodes (sending-rate evolution),
- and the *network metrics* (evolution of queue length, loss rate, arrival rate, etc.).

Fluid models enable:

Efficient evaluation (vs experiments)

A fluid model can fill the gaps in BBR analysis

Fluid model

System of ordinary differential equations (ODEs) describing the joint dynamics of

- the *congestion-control algorithms (CCAs)* running at the sending nodes (sending-rate evolution),
- and the *network metrics* (evolution of queue length, loss rate, arrival rate, etc.).

Fluid models enable:

Efficient evaluation (vs experiments)

Differential equations can be efficiently solved for a wide range of scenarios
(e.g., simulation cost is independent of flow rates)

A fluid model can fill the gaps in BBR analysis

Fluid model

System of ordinary differential equations (ODEs) describing the joint dynamics of

- the *congestion-control algorithms (CCAs)* running at the sending nodes (sending-rate evolution),
- and the *network metrics* (evolution of queue length, loss rate, arrival rate, etc.).

Fluid models enable:

Efficient evaluation (vs experiments)

Differential equations can be efficiently solved for a wide range of scenarios
(e.g., simulation cost is independent of flow rates)

Expression of transient effects (vs steady-state models)

A fluid model can fill the gaps in BBR analysis

Fluid model

System of ordinary differential equations (ODEs) describing the joint dynamics of

- the *congestion-control algorithms (CCAs)* running at the sending nodes (sending-rate evolution),
- and the *network metrics* (evolution of queue length, loss rate, arrival rate, etc.).

Fluid models enable:

Efficient evaluation (vs experiments)

Differential equations can be efficiently solved for a wide range of scenarios
(e.g., simulation cost is independent of flow rates)

Expression of transient effects (vs steady-state models)

Fluid models allow to investigate if/how the CCA converges to an equilibrium (stability analysis)

Our contribution: A BBR analysis based on a fluid model

Our contribution: A BBR analysis based on a fluid model

Fluid-model design

Formalization of BBR behavior

Design of new techniques

$$\dot{\tau}_i^{\min} = -\Gamma \quad \tau_i^{\min}(t) - \tau_i(t - d_i^{\mathrm{p}})$$

$$\dot{x}_i^{\mathrm{btl}} = \sigma\left(t_i^{\mathrm{pbw}} - T_i^{\mathrm{pbw}} + 0.01\right) \cdot \left(x_i^{\max} - x_i^{\mathrm{btl}}\right)$$

$$x_i^{\mathrm{dlv}} = \frac{x_i(t - d_i^{\mathrm{p}})}{y_\ell \quad t - d_{i,\ell}^{\mathrm{b}}} \cdot \begin{cases} C_\ell & \text{if } q_\ell(t - d_{i,\ell}^{\mathrm{b}}) > 0 \\ y_\ell\left(t - d_{i,\ell}^{\mathrm{b}}\right) & \text{otherwise} \end{cases}$$

Our contribution: A BBR analysis based on a fluid model

Fluid-model design

Formalization of BBR behavior

Design of new techniques

$$\dot{\tau}_i^{\text{min}} = -\Gamma \quad \tau_i^{\text{min}}(t) - \tau_i(t - d_i^{\text{p}})$$

$$\dot{x}_i^{\text{btl}} = \sigma\left(t_i^{\text{pbw}} - T_i^{\text{pbw}} + 0.01\right) \cdot \left(x_i^{\text{max}} - x_i^{\text{btl}}\right)$$

$$x_i^{\text{dlv}} = \frac{x_i(t - d_i^{\text{p}})}{y_\ell \quad t - d_{i,\ell}^{\text{b}}} \cdot \begin{cases} C_\ell & \text{if } q_\ell(t - d_{i,\ell}^{\text{b}}) > 0 \\ y_\ell\left(t - d_{i,\ell}^{\text{b}}\right) & \text{otherwise} \end{cases}$$

Experimental validation

Confirmation of prior insights

Generation of new insights

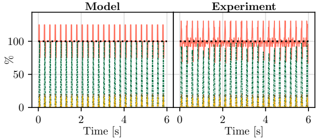

Our contribution: A BBR analysis based on a fluid model

Fluid-model design

Formalization of BBR behavior

Design of new techniques

$$\dot{\tau}_i^{\min} = -\Gamma \quad \tau_i^{\min}(t) - \tau_i(t - d_i^{\mathrm{p}})$$

$$\dot{x}_i^{\mathrm{btl}} = \sigma\left(t_i^{\mathrm{pbw}} - T_i^{\mathrm{pbw}} + 0.01\right) \cdot \left(x_i^{\max} - x_i^{\mathrm{btl}}\right)$$

$$x_i^{\mathrm{dlv}} = \frac{x_i(t - d_i^{\mathrm{p}})}{y_\ell \quad t - d_{i,\ell}^{\mathrm{b}}} \cdot \begin{cases} C_\ell & \text{if } q_\ell(t - d_{i,\ell}^{\mathrm{b}}) > 0 \\ y_\ell\left(t - d_{i,\ell}^{\mathrm{b}}\right) & \text{otherwise} \end{cases}$$

Experimental validation

Confirmation of prior insights

Generation of new insights

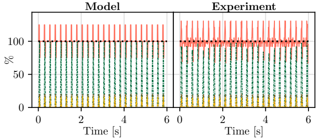

Theoretical stability analysis

Characterization of equilibria

Proof of asymptotic stability

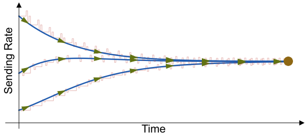

Our contribution: A BBR analysis based on a fluid model

Fluid-model design

Formalization of BBR behavior

Design of new techniques

$$\dot{\tau}_i^{\text{min}} = -\Gamma \quad \tau_i^{\text{min}}(t) - \tau_i(t - d_i^{\text{p}})$$

$$\dot{x}_i^{\text{btl}} = \sigma\left(t_i^{\text{pbw}} - T_i^{\text{pbw}} + 0.01\right) \cdot \left(x_i^{\text{max}} - x_i^{\text{btl}}\right)$$

$$x_i^{\text{dlv}} = \frac{x_i(t - d_i^{\text{p}})}{y_\ell \quad t - d_{i,\ell}^{\text{b}}} \cdot \begin{cases} C_\ell & \text{if } q_\ell(t - d_{i,\ell}^{\text{b}}) > 0 \\ y_\ell\left(t - d_{i,\ell}^{\text{b}}\right) & \text{otherwise} \end{cases}$$

Experimental validation

Confirmation of prior insights

Generation of new insights

Theoretical stability analysis

Characterization of equilibria

Proof of asymptotic stability

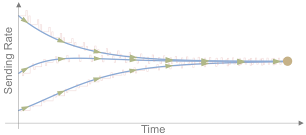

A primer on CCA fluid models: RENO [Low'02]

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

```
if ack_received then
  w ← w + 1/w
else // packet loss
  w ← w/2
```

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

if ack_received **then**

$w \leftarrow w + \frac{1}{w}$

else // packet loss

$w \leftarrow w/2$

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

```
if ack_received then
  w ← w + 1/w
else // packet loss
  w ← w/2
```

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

$$x(t) = \frac{w(t)}{\tau(t)}$$

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

if ack_received **then**

$w \leftarrow w + \frac{1}{w}$

else // packet loss

$w \leftarrow w/2$

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

$$x(t) = \frac{w(t)}{\tau(t)}$$

$$\dot{w}(t) = (1 - p\,(t-\tau)) \cdot x(t-\tau) \cdot \frac{1}{w(t)} \quad - \quad p\,(t-\tau) \cdot x(t-\tau) \cdot \frac{w(t)}{2}$$

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

```
if ack_received then
  w ← w + 1/w
else // packet loss
  w ← w/2
```

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

$$x(t) = \frac{w(t)}{\tau(t)}$$

$$\dot{w}(t) = \underbrace{(1 - p\,(t-\tau))}_{\text{Non-lost traffic (\%)}} \cdot x(t-\tau) \cdot \frac{1}{w(t)} \quad - \quad p\,(t-\tau) \cdot x(t-\tau) \cdot \frac{w(t)}{2}$$

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

```
if ack_received then
  w ← w + 1/w
else // packet loss
  w ← w/2
```

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

$$x(t) = \frac{w(t)}{\tau(t)}$$

$$\dot{w}(t) = \underbrace{\underbrace{(1 - p\,(t-\tau))}_{\text{Non-lost traffic (\%)}} \cdot x(t-\tau)}_{\text{Rate of incoming ACKs}} \cdot \frac{1}{w(t)} \quad - \quad p\,(t-\tau) \cdot x(t-\tau) \cdot \frac{w(t)}{2}$$

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

```
if ack_received then
  w ← w + 1/w
else // packet loss
  w ← w/2
```

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

$$x(t) = \frac{w(t)}{\tau(t)}$$

$$\dot{w}(t) = \underbrace{\underbrace{\underbrace{(1 - p\,(t-\tau))}_{\text{Non-lost traffic (\%)}} \cdot x(t-\tau)}_{\text{Rate of incoming ACKs}} \cdot \frac{1}{w(t)}}_{\text{Rate of cwnd increase}} \quad - \quad p\,(t-\tau) \cdot x(t-\tau) \cdot \frac{w(t)}{2}$$

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

```
if ack_received then
```
$w \leftarrow w + \frac{1}{w}$
```
else // packet loss
```
$w \leftarrow w/2$

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

$$x(t) = \frac{w(t)}{\tau(t)}$$

$$\dot{w}(t) = \underbrace{\underbrace{\underbrace{(1 - p\,(t-\tau))}_{\text{Non-lost traffic (\%)}} \cdot x(t-\tau)}_{\text{Rate of incoming ACKs}} \cdot \frac{1}{w(t)}}_{\text{Rate of cwnd increase}} \quad - \quad \underbrace{p\,(t-\tau)}_{\text{Loss rate (\%)}} \cdot x(t-\tau) \cdot \frac{w(t)}{2}$$

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

```
if ack_received then
  w ← w + 1/w
else // packet loss
  w ← w/2
```

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

$$x(t) = \frac{w(t)}{\tau(t)}$$

$$\dot{w}(t) = \underbrace{\underbrace{\underbrace{(1 - p\,(t-\tau))}_{\text{Non-lost traffic (\%)}} \cdot x(t-\tau)}_{\text{Rate of incoming ACKs}} \cdot \frac{1}{w(t)}}_{\text{Rate of cwnd increase}} \quad - \quad \underbrace{\underbrace{p\,(t-\tau)}_{\text{Loss rate (\%)}} \cdot x(t-\tau)}_{\text{Loss rate (segments/time)}} \cdot \frac{w(t)}{2}$$

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

if ack_received **then**
  $w \leftarrow w + \frac{1}{w}$
else // packet loss
  $w \leftarrow w/2$

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

$$x(t) = \frac{w(t)}{\tau(t)}$$

$$\dot{w}(t) = \underbrace{\underbrace{\underbrace{(1 - p(t-\tau))}_{\text{Non-lost traffic (\%)}} \cdot x(t-\tau)}_{\text{Rate of incoming ACKs}} \cdot \frac{1}{w(t)}}_{\text{Rate of cwnd increase}} - \underbrace{\underbrace{\underbrace{p(t-\tau)}_{\text{Loss rate (\%)}} \cdot x(t-\tau)}_{\text{Loss rate (segments/time)}} \cdot \frac{w(t)}{2}}_{\text{Rate of cwnd decrease}}$$

A primer on CCA fluid models: RENO [Low'02]

Reno control loop: Congestion-window size w

```
if ack_received then
  w ← w + 1/w
else // packet loss
  w ← w/2
```

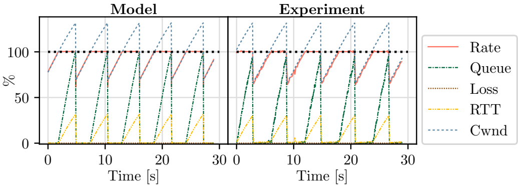

Fluid-model approximation: Congestion-window size w, sending rate x, RTT τ

$$x(t) = \frac{w(t)}{\tau(t)}$$

$$\dot{w}(t) = \underbrace{\underbrace{\underbrace{(1 - p(t-\tau))}_{\text{Non-lost traffic (\%)}} \cdot x(t-\tau)}_{\text{Rate of incoming ACKs}} \cdot \frac{1}{w(t)}}_{\text{Rate of cwnd increase}} - \underbrace{\underbrace{\underbrace{p(t-\tau)}_{\text{Loss rate (\%)}} \cdot x(t-\tau)}_{\text{Loss rate (segments/time)}} \cdot \frac{w(t)}{2}}_{\text{Rate of cwnd decrease}}$$

Representing BBR in a fluid model: Background

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

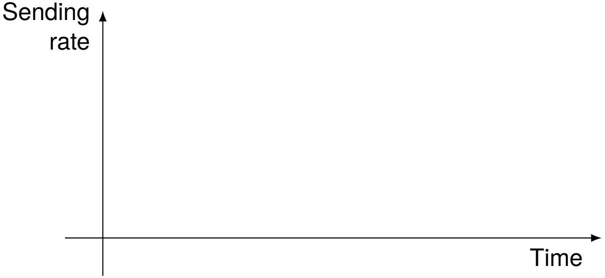

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

Probing periods of 8 MinRTT (phases)

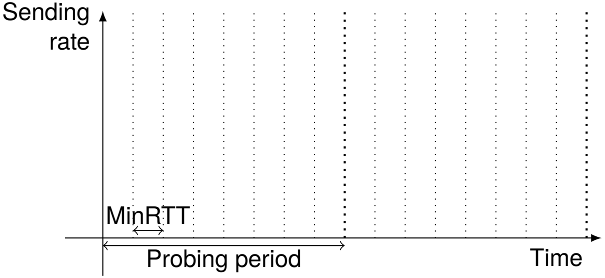

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

Probing periods of 8 MinRTT (phases)

Base rate during period is bottleneck-bandwidth estimate b

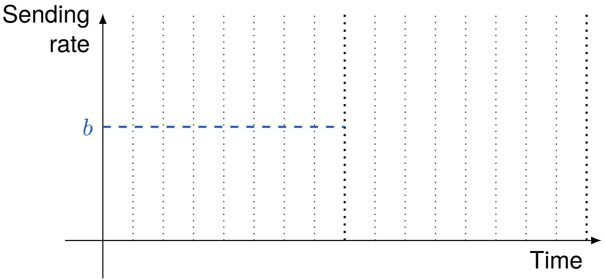

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

Probing periods of 8 MinRTT (phases)

Base rate during period is bottleneck-bandwidth estimate b

In a random phase, rate is raised to $5/4 \cdot b$

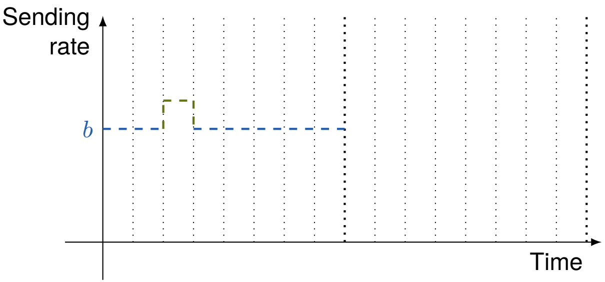

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

Probing periods of 8 MinRTT (phases)

Base rate during period is bottleneck-bandwidth estimate b

In a random phase, rate is raised to $5/4 \cdot b$

In next phase, rate is reduced to $3/4 \cdot b$

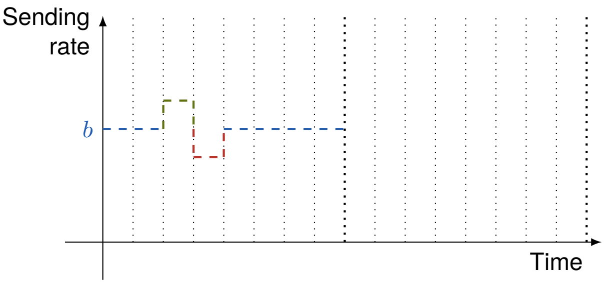

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

Probing periods of 8 MinRTT (phases)

Base rate during period is bottleneck-bandwidth estimate b

In a random phase, rate is raised to $^5/_4 \cdot b$

In next phase, rate is reduced to $^3/_4 \cdot b$

Delivery rate is observed during period

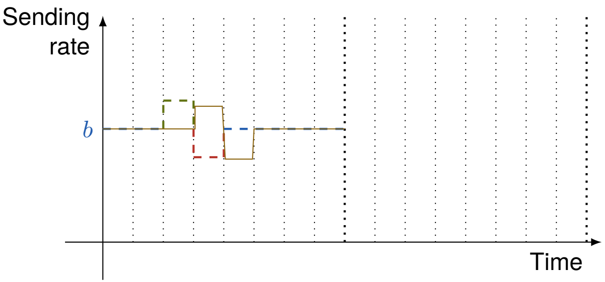

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

Probing periods of 8 MinRTT (phases)

Base rate during period is bottleneck-bandwidth estimate b

In a random phase, rate is raised to ${}^{5}/_{4} \cdot b$

In next phase, rate is reduced to ${}^{3}/_{4} \cdot b$

Delivery rate is observed during period

Maximum delivery rate is new bottleneck estimate b'

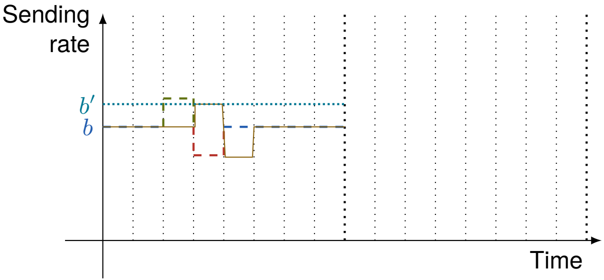

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

Probing periods of 8 MinRTT (phases)

Base rate during period is bottleneck-bandwidth estimate b

In a random phase, rate is raised to $5/4 \cdot b$

In next phase, rate is reduced to $3/4 \cdot b$

Delivery rate is observed during period

Maximum delivery rate is new bottleneck estimate b'

Process is repeated in next period

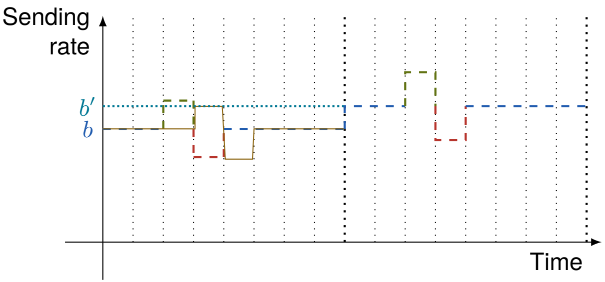

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

Probing periods of 8 MinRTT (phases)

Base rate during period is bottleneck-bandwidth estimate b

In a random phase, rate is raised to $5/4 \cdot b$

In next phase, rate is reduced to $3/4 \cdot b$

Delivery rate is observed during period

Maximum delivery rate is new bottleneck estimate b'

Process is repeated in next period

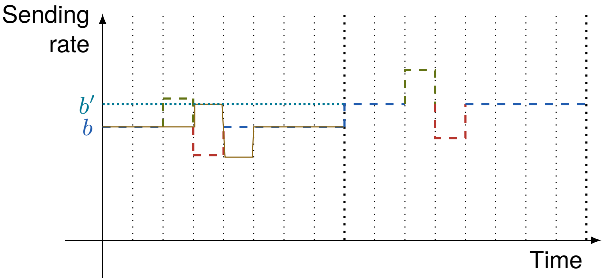

How to model this probing with (differential) equations?

Representing BBR in a fluid model: Background

BBRv1 bandwidth probing:

Probing periods of 8 MinRTT (phases)

Base rate during period is bottleneck-bandwidth estimate b

In a random phase, rate is raised to $\nicefrac{5}{4} \cdot b$

In next phase, rate is reduced to $\nicefrac{3}{4} \cdot b$

Delivery rate is observed during period

Maximum delivery rate is new bottleneck estimate b'

Process is repeated in next period

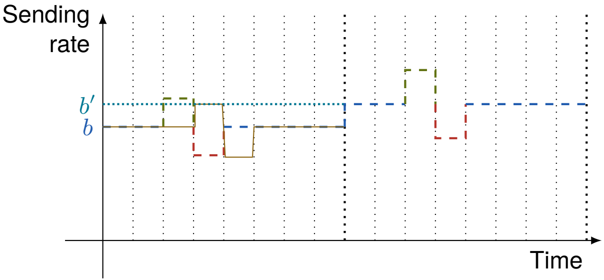

How to model this probing with (differential) equations?

Probing pulses? Random phases? Maximum tracking? Periodic adjustment?

Representing BBR in a fluid model: Probing pulses

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

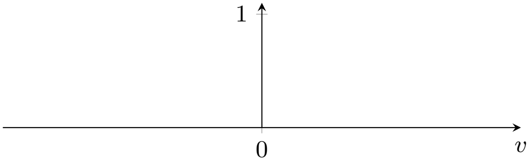

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

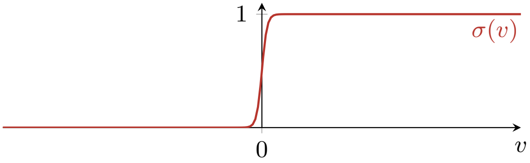

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

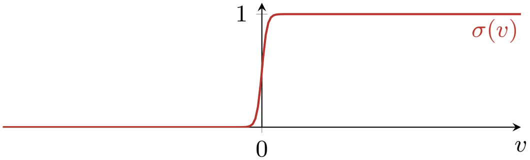

Pulse function

$$\Phi(t, \phi) = \sigma\left(t - \phi \cdot \tau_{\min}\right) \cdot \sigma\left((\phi + 1) \cdot \tau_{\min} - t\right)$$

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

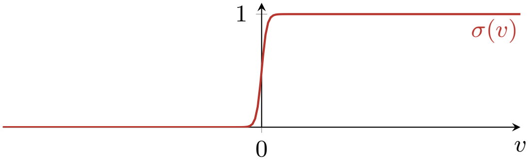

Pulse function

$$\Phi(t, \phi) = \sigma\left(t - \phi \cdot \tau_{\min}\right) \cdot \sigma\left((\phi + 1) \cdot \tau_{\min} - t\right)$$

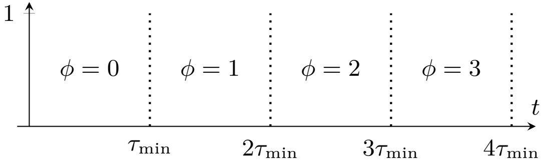

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

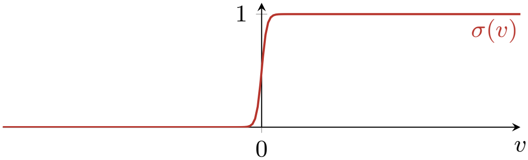

Pulse function

$$\Phi(t, \phi) = \sigma\left(t - \phi \cdot \tau_{\min}\right) \cdot \sigma\left((\phi + 1) \cdot \tau_{\min} - t\right)$$

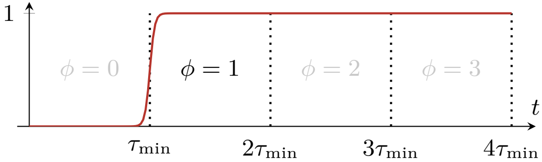

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

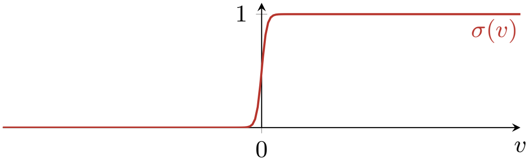

Pulse function

$$\Phi(t, \phi) = \sigma\left(t - \phi \cdot \tau_{\min}\right) \cdot \sigma\left((\phi + 1) \cdot \tau_{\min} - t\right)$$

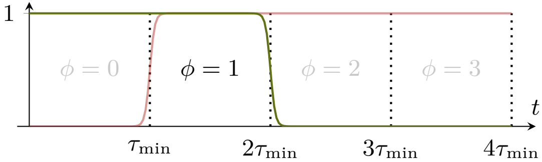

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

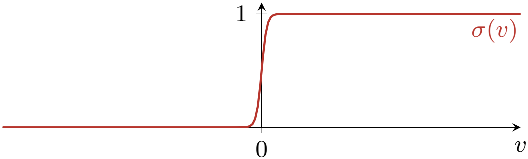

Pulse function

$$\Phi(t, \phi) = \sigma\left(t - \phi \cdot \tau_{\min}\right) \cdot \sigma\left((\phi + 1) \cdot \tau_{\min} - t\right)$$

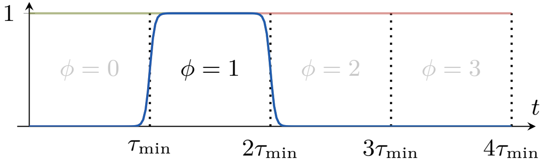

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

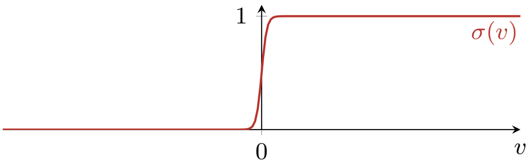

Pulse function

$$\Phi(t, \phi) = \sigma\left(t - \phi \cdot \tau_{\text{min}}\right) \cdot \sigma\left((\phi + 1) \cdot \tau_{\text{min}} - t\right)$$

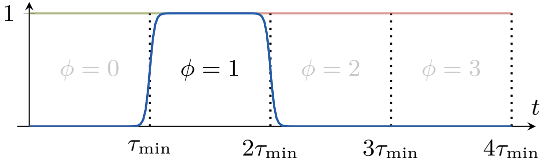

Pacing rate

$$x^{\text{pcg}}(t) = x^{\text{btl}}(t) \cdot \left(1 + {}^{1}/_{4}\Phi(t, \phi') - {}^{1}/_{4}\Phi(t, \phi' + 1)\right)$$

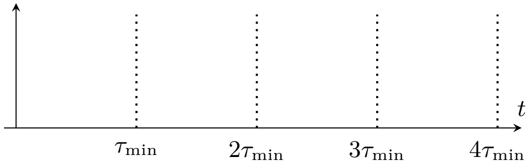

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

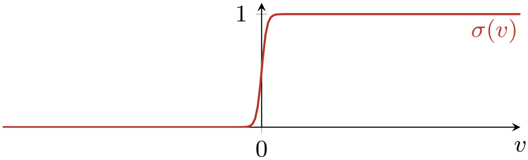

Pulse function

$$\Phi(t, \phi) = \sigma\left(t - \phi \cdot \tau_{\text{min}}\right) \cdot \sigma\left((\phi + 1) \cdot \tau_{\text{min}} - t\right)$$

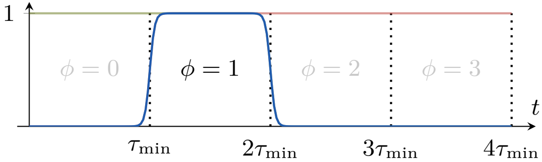

Pacing rate

$$x^{\text{pcg}}(t) = x^{\text{btl}}(t) \cdot \left(1 + \tfrac{1}{4}\Phi(t, \phi') - \tfrac{1}{4}\Phi(t, \phi' + 1)\right)$$

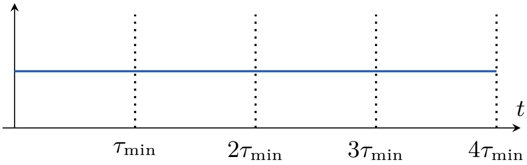

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

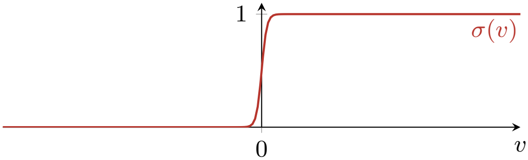

Pulse function

$$\Phi(t, \phi) = \sigma\left(t - \phi \cdot \tau_{\min}\right) \cdot \sigma\left((\phi + 1) \cdot \tau_{\min} - t\right)$$

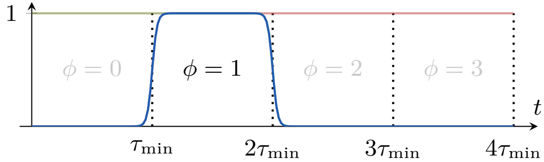

Pacing rate

$$x^{\mathrm{pcg}}(t) = x^{\mathrm{btl}}(t) \cdot \left(1 + {}^{1}\!/\!_{4}\Phi(t, \phi') - {}^{1}\!/\!_{4}\Phi(t, \phi' + 1)\right)$$

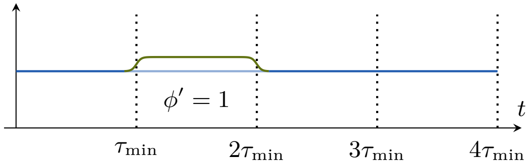

Representing BBR in a fluid model: Probing pulses

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

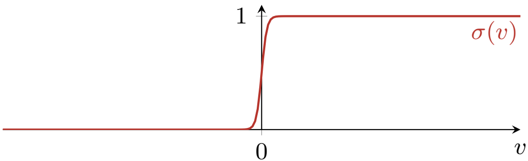

Pulse function

$$\Phi(t, \phi) = \sigma\left(t - \phi \cdot \tau_{\min}\right) \cdot \sigma\left((\phi + 1) \cdot \tau_{\min} - t\right)$$

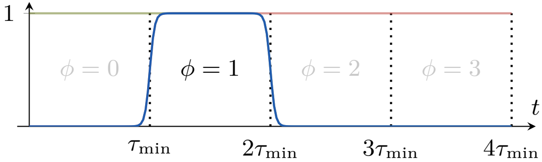

Pacing rate

$$x^{\text{pcg}}(t) = x^{\text{btl}}(t) \cdot \left(1 + {}^{1}/_{4}\Phi(t, \phi') - {}^{1}/_{4}\Phi(t, \phi' + 1)\right)$$

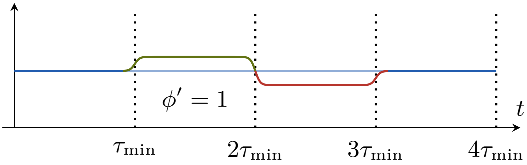

Representing BBR in a fluid model: *Randomized* probing pulses

Pacing rate of flow i

$$x_i^{\text{pcg}}(t) = x_i^{\text{btl}}(t) \cdot (1 + {}^{1}\!/\!{}_{4}\Phi(t, \phi_i) - {}^{1}\!/\!{}_{4}\Phi(t, \phi_i + 1))$$

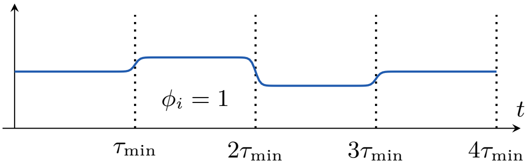

Representing BBR in a fluid model: *Randomized* probing pulses

Pacing rate of flow i

$$x_i^{\text{pcg}}(t) = x_i^{\text{btl}}(t) \cdot (1 + {}^1\!/_4 \Phi(t, \phi_i) - {}^1\!/_4 \Phi(t, \phi_i + 1))$$

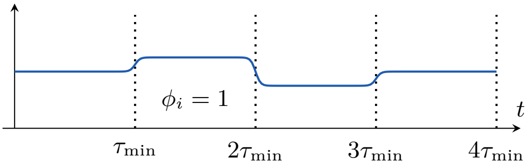

Achieve *intention* behind randomization
by deterministic means
$\implies$ Desynchronization

Representing BBR in a fluid model: *Randomized* probing pulses

Pacing rate of flow i

$$x_i^{\text{pcg}}(t) = x_i^{\text{btl}}(t) \cdot (1 + {}^1/_4 \Phi(t, \phi_i) - {}^1/_4 \Phi(t, \phi_i + 1))$$

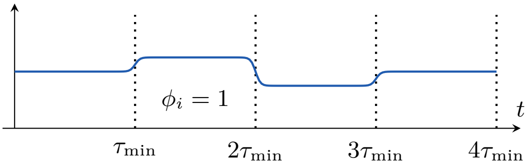

Achieve *intention* behind randomization
by deterministic means
$\implies$ Desynchronization

$$\forall i \in \mathbb{N}. \quad \phi_i = i \bmod 7$$

Representing BBR in a fluid model: *Randomized* probing pulses

Pacing rate of flow i

$$x_i^{\text{pcg}}(t) = x_i^{\text{btl}}(t) \cdot (1 + {}^1\!/\!{}_4\Phi(t, \phi_i) - {}^1\!/\!{}_4\Phi(t, \phi_i + 1))$$

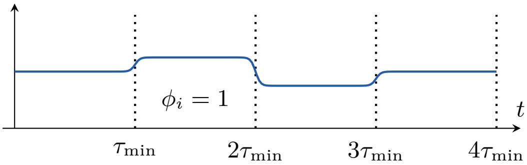

Achieve *intention* behind randomization
by deterministic means
$\implies$ Desynchronization

$$\forall i \in \mathbb{N}. \quad \phi_i = i \bmod 7$$

$$\phi_1 = 1 \qquad \phi_7 = 0 \qquad \phi_9 = 2$$

Representing BBR in a fluid model: *Randomized* probing pulses

Pacing rate of flow i

$$x_i^{\mathrm{pcg}}(t) = x_i^{\mathrm{btl}}(t) \cdot (1 + {}^1\!/\!{}_4\Phi(t, \phi_i) - {}^1\!/\!{}_4\Phi(t, \phi_i + 1))$$

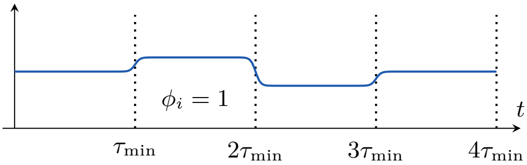

Achieve *intention* behind randomization
by deterministic means
$\implies$ Desynchronization

$$\forall i \in \mathbb{N}. \quad \phi_i = i \bmod 7$$

$$\phi_1 = 1 \qquad \phi_7 = 0 \qquad \phi_9 = 2$$

Desynchronized pacing rates
for flows 1, 7, and 10:

$$x_1^{\mathrm{pcg}}(t) \qquad x_7^{\mathrm{pcg}}(t) \qquad x_9^{\mathrm{pcg}}(t)$$

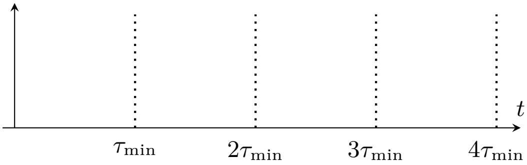

Representing BBR in a fluid model: *Randomized* probing pulses

Pacing rate of flow i

$$x_i^{\mathrm{pcg}}(t) = x_i^{\mathrm{btl}}(t)\cdot\left(1 + {}^1\!/_4\Phi(t,\phi_i) - {}^1\!/_4\Phi(t,\phi_i+1)\right)$$

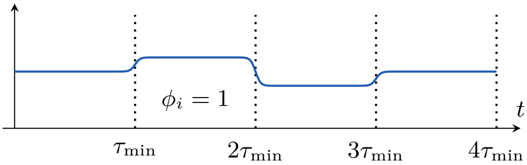

Achieve *intention* behind randomization
by deterministic means
$\implies$ Desynchronization

$$\forall i \in \mathbb{N}.\quad \phi_i = i \bmod 7$$

$$\phi_1 = 1 \qquad \phi_7 = 0 \qquad \phi_9 = 2$$

Desynchronized pacing rates
for flows 1, 7, and 10:

$$x_1^{\mathrm{pcg}}(t) \qquad x_7^{\mathrm{pcg}}(t) \qquad x_9^{\mathrm{pcg}}(t)$$

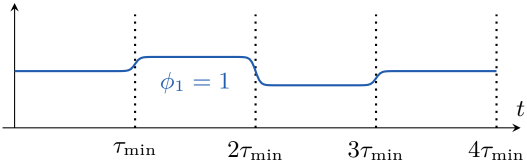

Representing BBR in a fluid model: *Randomized* probing pulses

Pacing rate of flow i

$$x_i^{\text{pcg}}(t) = x_i^{\text{btl}}(t) \cdot (1 + {}^1\!/\!{}_4\Phi(t, \phi_i) - {}^1\!/\!{}_4\Phi(t, \phi_i + 1))$$

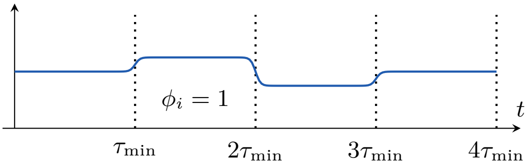

Achieve *intention* behind randomization
by deterministic means
$\implies$ Desynchronization

$$\forall i \in \mathbb{N}. \quad \phi_i = i \bmod 7$$

$$\phi_1 = 1 \qquad \phi_7 = 0 \qquad \phi_9 = 2$$

Desynchronized pacing rates
for flows 1, 7, and 10:

$$x_1^{\text{pcg}}(t) \qquad x_7^{\text{pcg}}(t) \qquad x_9^{\text{pcg}}(t)$$

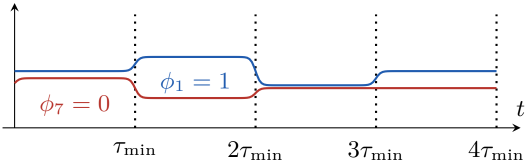

Representing BBR in a fluid model: *Randomized* probing pulses

Pacing rate of flow i

$$x_i^{\text{pcg}}(t) = x_i^{\text{btl}}(t) \cdot (1 + {}^1\!/\!_4\Phi(t,\phi_i) - {}^1\!/\!_4\Phi(t,\phi_i+1))$$

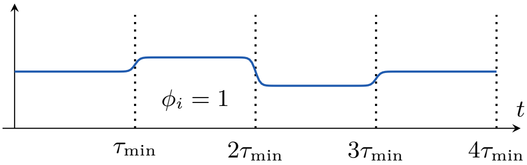

Achieve *intention* behind randomization
by deterministic means
$\implies$ Desynchronization

$$\forall i \in \mathbb{N}. \quad \phi_i = i \bmod 7$$

$\phi_1 = 1$ $\quad$ $\phi_7 = 0$ $\quad$ $\phi_9 = 2$

Desynchronized pacing rates
for flows 1, 7, and 10:

$x_1^{\text{pcg}}(t)$ $\quad$ $x_7^{\text{pcg}}(t)$ $\quad$ $x_9^{\text{pcg}}(t)$

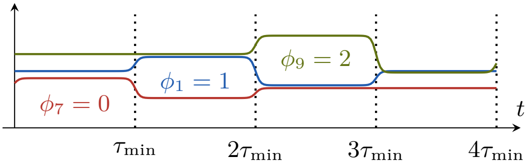

Representing BBR in a fluid model: Maximum tracking

Representing BBR in a fluid model: Maximum tracking

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

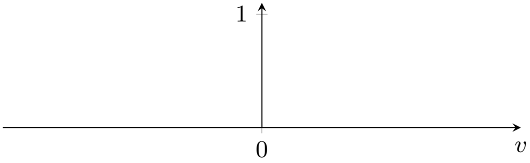

Representing BBR in a fluid model: Maximum tracking

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

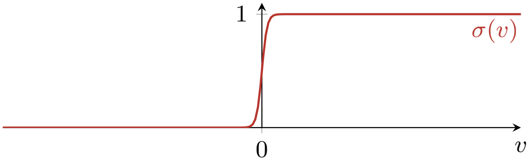

Representing BBR in a fluid model: Maximum tracking

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

Maximum function

$$\Gamma(v) = v \cdot \sigma(v)$$

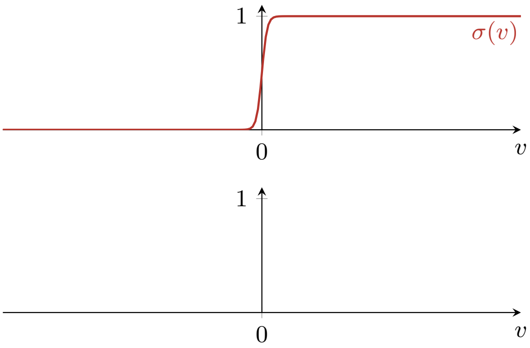

Representing BBR in a fluid model: Maximum tracking

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

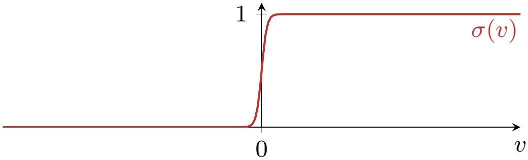

Maximum function

$$\Gamma(v) = v \cdot \sigma(v)$$

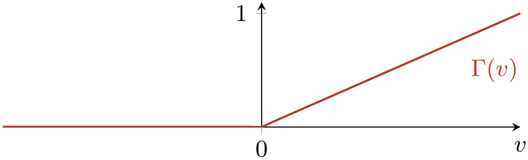

Representing BBR in a fluid model: Maximum tracking

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

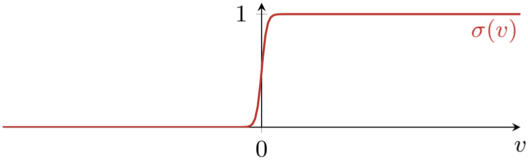

Maximum function

$$\Gamma(v) = v \cdot \sigma(v)$$

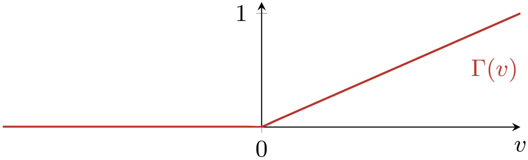

Tracking of maximum delivery rate

$$\dot{x}^{\max}(t) = \Gamma\left(x^{\mathrm{dlv}}(t) - x^{\max}(t)\right)$$

Representing BBR in a fluid model: Maximum tracking

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

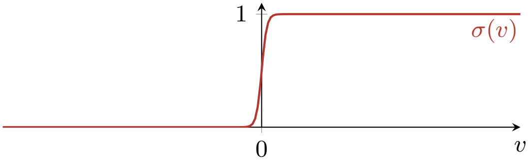

Maximum function

$$\Gamma(v) = v \cdot \sigma(v)$$

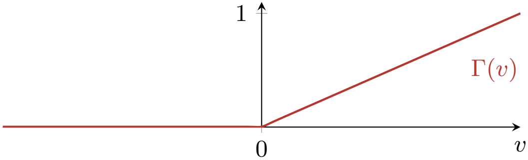

Tracking of maximum delivery rate

$$\dot{x}^{\max}(t) = \Gamma\left(x^{\text{dlv}}(t) - x^{\max}(t)\right)$$

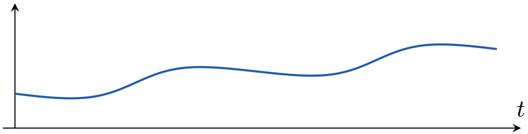

Representing BBR in a fluid model: Maximum tracking

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

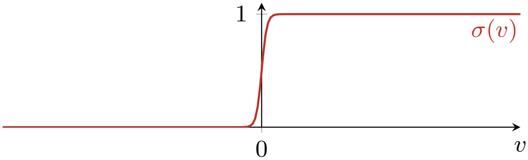

Maximum function

$$\Gamma(v) = v \cdot \sigma(v)$$

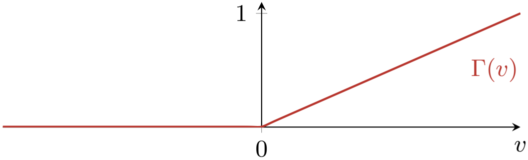

Tracking of maximum delivery rate

$$\dot{x}^{\max}(t) = \Gamma\left(x^{\text{dlv}}(t) - x^{\max}(t)\right)$$

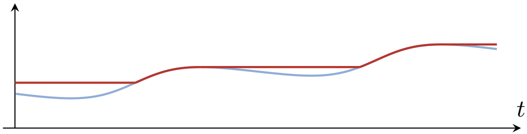

Representing BBR in a fluid model: Periodic adjustment

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

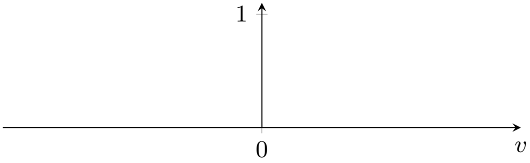

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

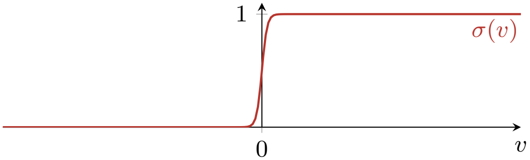

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

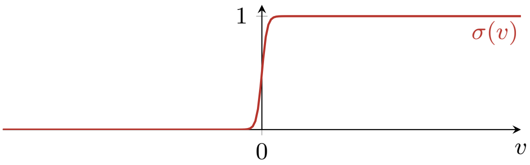

Pulse at period end

$$\Phi'(t) = \sigma\left(t - 7.9 \cdot \tau_{\min}\right) \;\cdot\; \sigma\left(8 \cdot \tau_{\min} - t\right)$$

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

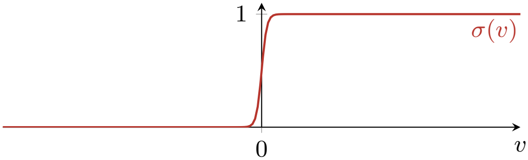

Pulse at period end

$$\Phi'(t) = \sigma\left(t - 7.9 \cdot \tau_{\min}\right) \cdot \sigma\left(8 \cdot \tau_{\min} - t\right)$$

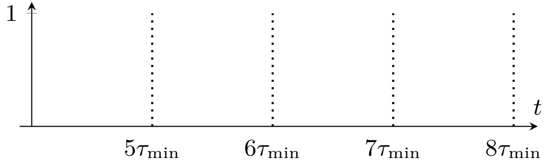

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

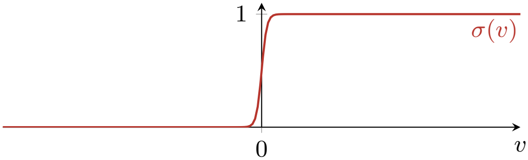

Pulse at period end

$$\Phi'(t) = \sigma\left(t - 7.9 \cdot \tau_{\min}\right) \cdot \sigma\left(8 \cdot \tau_{\min} - t\right)$$

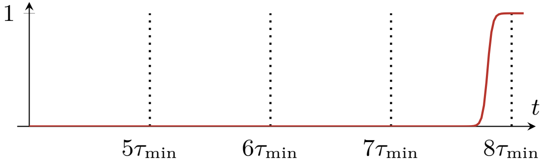

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

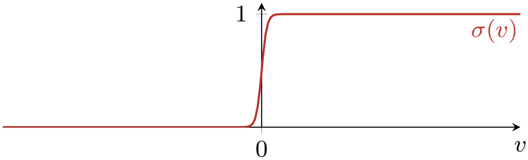

Pulse at period end

$$\Phi'(t) = \sigma\left(t - 7.9 \cdot \tau_{\min}\right) \cdot \sigma\left(8 \cdot \tau_{\min} - t\right)$$

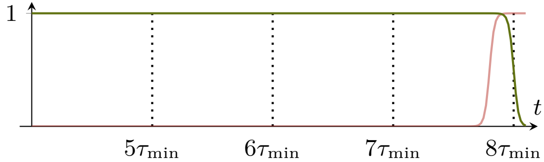

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

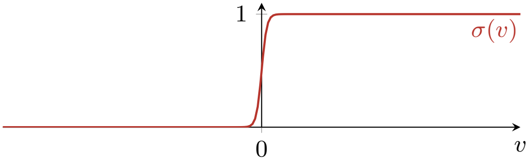

Pulse at period end

$$\Phi'(t) = \sigma\left(t - 7.9 \cdot \tau_{\min}\right) \cdot \sigma\left(8 \cdot \tau_{\min} - t\right)$$

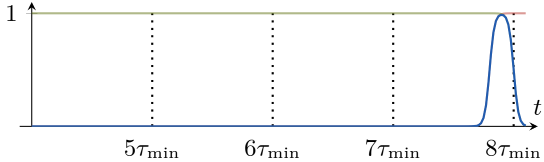

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

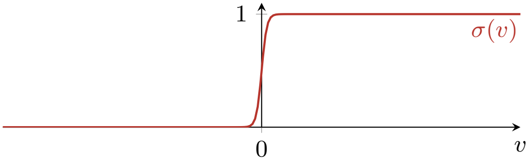

Pulse at period end

$$\Phi'(t) = \sigma\left(t - 7.9 \cdot \tau_{\min}\right) \cdot \sigma\left(8 \cdot \tau_{\min} - t\right)$$

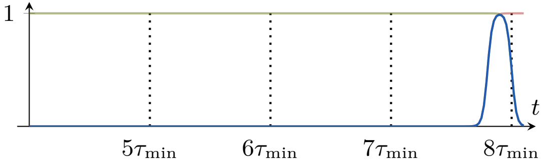

Update of bottleneck-bandwidth estimate

$$\dot{x}^{\text{btl}}(t) = \Phi'(t) \cdot \left(x^{\max}(t) - x^{\text{btl}}(t)\right)$$

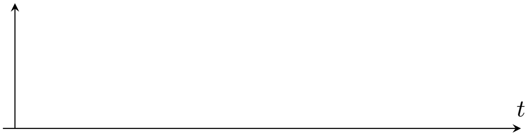

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

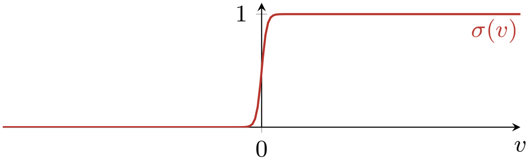

Pulse at period end

$$\Phi'(t) = \sigma\left(t - 7.9 \cdot \tau_{\text{min}}\right) \cdot \sigma\left(8 \cdot \tau_{\text{min}} - t\right)$$

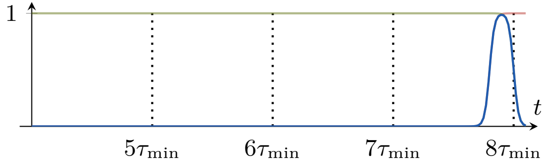

Update of bottleneck-bandwidth estimate

$$\dot{x}^{\text{btl}}(t) = \Phi'(t) \cdot \left(x^{\text{max}}(t) - x^{\text{btl}}(t)\right)$$

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

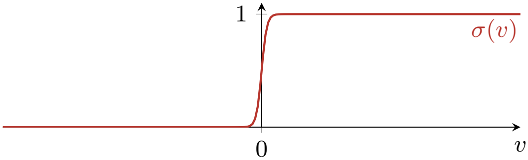

Pulse at period end

$$\Phi'(t) = \sigma\left(t - 7.9 \cdot \tau_{\min}\right) \cdot \sigma\left(8 \cdot \tau_{\min} - t\right)$$

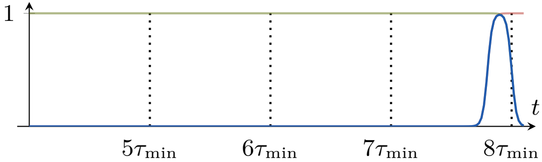

Update of bottleneck-bandwidth estimate

$$\dot{x}^{\text{btl}}(t) = \Phi'(t) \cdot \left(x^{\max}(t) - x^{\text{btl}}(t)\right)$$

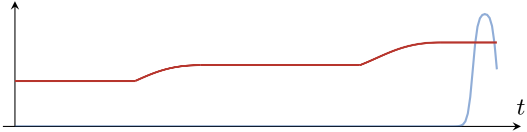

Representing BBR in a fluid model: Periodic adjustment

Sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-K \cdot v}}$$

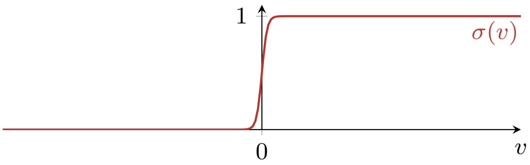

Pulse at period end

$$\Phi'(t) = \sigma\left(t - 7.9 \cdot \tau_{\min}\right) \cdot \sigma\left(8 \cdot \tau_{\min} - t\right)$$

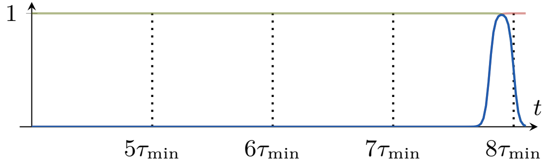

Update of bottleneck-bandwidth estimate

$$\dot{x}^{\mathrm{btl}}(t) = \Phi'(t) \cdot \left(x^{\max}(t) - x^{\mathrm{btl}}(t)\right)$$

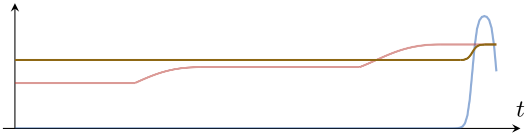

Representing BBR in a fluid model: End result

Network model

$$y_\ell = \sum_{i \in U_\ell} x_i\left(t - d^{\mathrm{f}}_{i,\ell}\right),$$

$$\dot{q}_\ell = (1 - p_\ell) \cdot y_\ell - C_\ell, \quad q_\ell(t) \in [0, B_\ell],$$

$$\tau_{\pi_i} = \sum_{\ell \in \pi_i} \tau_\ell = \sum_{\ell \in \pi_i} d_\ell + \frac{q_\ell}{C_\ell}.$$

$$p_\ell(t) = \sigma\left(y_\ell(t) - C_\ell\right) \cdot \left(1 - \frac{C_\ell}{y_\ell}\right) \cdot \left(\frac{q_\ell}{B_\ell}\right)^L$$

$$p_\ell = \frac{q_\ell}{B_\ell} \quad \in [0, 1].$$

$$p_{\pi_i}(t) = 1 - \prod_{\ell \in \pi} \left(1 - p_\ell\left(t + d^{\mathrm{f}}_{i,\ell}\right)\right) \approx \sum_{\ell \in \pi} p_\ell\left(t + d^{\mathrm{f}}_{i,\ell}\right).$$

$$x_i = \frac{w_i}{\tau_i}.$$

Basic BBR model

$$\dot{\tau}^{\mathrm{min}}_i = -\Gamma\left(\tau^{\mathrm{min}}_i(t) - \tau_i(t - d^{\mathrm{p}}_i)\right)$$

$$\Gamma(v) = v \cdot \sigma(v).$$

$$\Delta m^{\mathrm{prt}}_i = \sigma\left(t^{\mathrm{prt}}_i - T^{\mathrm{prt}}_i\right) \cdot \left((1 - m^{\mathrm{prt}}_i) - m^{\mathrm{prt}}_i\right)$$

$$T^{\mathrm{prt}}_i = m^{\mathrm{prt}}_i \cdot 0.2 + \left(1 - m^{\mathrm{prt}}_i\right) \cdot 10$$

$$\dot{t}^{\mathrm{prt}}_i = 1 - \sigma\left(t^{\mathrm{prt}}_i - T^{\mathrm{prt}}_i\right) \cdot t^{\mathrm{prt}}_i - \sigma\left(\tau^{\mathrm{min}}_i - \tau_i(t - d_\ell)\right) \cdot t^{\mathrm{prt}}_i$$

$$x_i = m^{\mathrm{prt}}_i \cdot \frac{w^{\mathrm{prt}}_i}{\tau_i} - \left(1 - m^{\mathrm{prt}}_i\right) \cdot x^{\mathrm{pbw}}_i$$

$$x^{\mathrm{pbw}}_i = \min\left(\frac{w^{\mathrm{pbw}}_i}{\tau_i}, x^{\mathrm{pcg}}_i\right)$$

$$\dot{t}^{\mathrm{pbw}}_i = 1 - \sigma\left(t^{\mathrm{pbw}}_i - T^{\mathrm{pbw}}_i\right) \cdot t^{\mathrm{pbw}}_i$$

$$x^{\mathrm{dlv}}_i = \frac{x_i(t - d^{\mathrm{p}}_i)}{y_\ell\left(t - d^{\mathrm{b}}_{i,\ell}\right)} \cdot \begin{cases} C_\ell & \text{if } q_\ell(t - d^{\mathrm{b}}_{i,\ell}) > 0 \\ y_\ell\left(t - d^{\mathrm{b}}_{i,\ell}\right) & \text{otherwise} \end{cases}$$

$$\dot{x}^{\mathrm{max}}_i = \Gamma(x^{\mathrm{dlv}}_i - x^{\mathrm{max}}_i) - \sigma(0.01 - t^{\mathrm{pbw}}_i) \cdot x^{\mathrm{max}}_i$$

$$\dot{v}_i = x_i - x^{\mathrm{dlv}}_i$$

BBRv1 model

$$\dot{x}^{\mathrm{btl}}_i = \sigma\left(t^{\mathrm{pbw}}_i - T^{\mathrm{pbw}}_i + 0.01\right) \cdot \left(x^{\mathrm{max}}_i - x^{\mathrm{btl}}_i\right)$$

$$\Phi_i(t, \phi) = \sigma\left(t^{\mathrm{pbw}}(t) - \phi \cdot \tau^{\mathrm{min}}_i\right) \cdot \sigma\left((\phi + 1) \cdot \tau^{\mathrm{min}}_i - t^{\mathrm{pbw}}\right)$$

$$x^{\mathrm{pcg}}_i = x^{\mathrm{btl}}_i \cdot \left(1 + \tfrac{1}{4} \cdot \Phi_i(t, \phi_i) - \tfrac{1}{4} \cdot \Phi_i(t, \phi_i + 1)\right)$$

$$w^{\mathrm{prt}}_i = 4$$

$$w^{\mathrm{pbw}}_i = 2 \cdot \overline{w}_i = 2 \cdot x^{\mathrm{btl}}_i \cdot \tau^{\mathrm{min}}_i$$

BBRv2 model

$$T^{\mathrm{pbw}}_i = \min\left(62 \cdot \tau^{\mathrm{min}}_i, 2 + \frac{i}{N}\right)$$

$$x^{\mathrm{pcg}}_i = x^{\mathrm{btl}}_i \cdot \left(1 + \tfrac{1}{4} \cdot \sigma\left(t^{\mathrm{pbw}}_i - \tau^{\mathrm{min}}_i\right) \cdot \left(1 - m^{\mathrm{dwn}}_i\right) - \tfrac{1}{4} \cdot m^{\mathrm{dwn}}_i\right)$$

$$\begin{aligned} \Delta m^{\mathrm{dwn}}_i = {} & (1 - m^{\mathrm{crs}}_i) \cdot (1 - m^{\mathrm{dwn}}_i) \cdot \sigma \;\; t^{\mathrm{pbw}}_i - \tau^{\mathrm{min}}_i \\ & \cdot \min\left(\sigma\left(v_i - 5/4 \cdot \overline{w}_i\right) + \sigma\left(p_{\pi} i - 0.02\right),\ 1\right) \\ & - m^{\mathrm{dwn}}_i \cdot \sigma\left(w^-_i - v_i\right) \end{aligned}$$

$$\Delta m^{\mathrm{crs}}_i = -\Delta m^{\mathrm{dwn}}_i - \sigma\left(t^{\mathrm{pbw}}_i - T^{\mathrm{pbw}}_i\right) \cdot m^{\mathrm{crs}}_i$$

$$\dot{x}^{\mathrm{btl}}_i = m^{\mathrm{dwn}}_i \cdot \left(\max\left(x^{\mathrm{max}}_i,\ x^{\mathrm{max}}_i(t - T^{\mathrm{pbw}})\right) - x^{\mathrm{btl}}_i\right)$$

$$\begin{aligned} \dot{w}^{\mathrm{hi}}_i = {} & (1 - m^{\mathrm{crs}}_i) \cdot \sigma\left(t^{\mathrm{pbw}}_i - \tau^{\mathrm{min}}_i\right) \cdot \sigma\left(v_i - w^{\mathrm{hi}}\right) \cdot 2^{t^{\mathrm{pbw}}_i / \tau^{\mathrm{min}}_i} \\ & - \sigma\left(p_{\pi_i} - 0.02\right) \cdot \frac{0.3}{\tau^{\mathrm{min}}_i} \cdot w^{\mathrm{hi}}_i \end{aligned}$$

$$\dot{w}^{\mathrm{lo}}_i = (1 - m^{\mathrm{crs}}_i) \cdot \left(w^-_i - w^{\mathrm{lo}}_i\right) - m^{\mathrm{crs}}_i \cdot \sigma(p_{\pi_i}) \cdot \frac{0.3 w^{\mathrm{lo}}_i}{\tau^{\mathrm{min}}_i}.$$

$$w^{\mathrm{pbw}}_i = \min\left(2 \cdot \overline{w}_i,\ (1 - m^{\mathrm{crs}}_i) \cdot w^{\mathrm{hi}}_i + m^{\mathrm{crs}}_i \cdot w^{\mathrm{lo}}_i\right)$$

$$w^{\mathrm{prt}}_i = \frac{\overline{w}_i}{2}$$

Representing BBR in a fluid model: End result

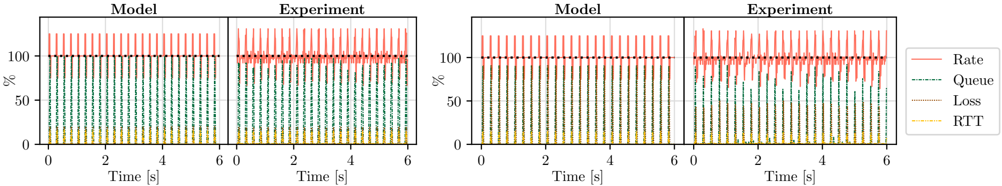

(a) Drop-tail (b) RED

BBRv1.

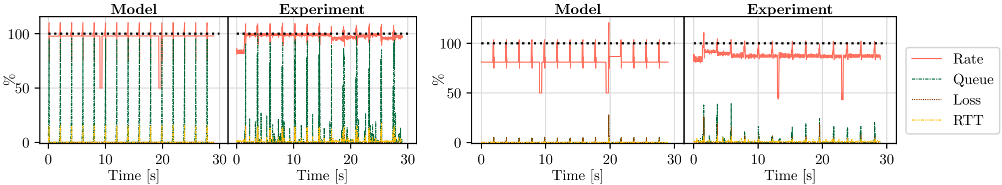

(a) Drop-tail (b) RED

BBRv2.

Our contribution: A BBR analysis based on a fluid model

Fluid-model design

Formalization of BBR behavior

Design of new techniques

$$\dot{\tau}_i^{\min} = -\Gamma\left(\tau_i^{\min}(t) - \tau_i(t-d_i^{\mathrm{p}})\right)$$

$$\dot{x}_i^{\mathrm{btl}} = \sigma\left(t_i^{\mathrm{pbw}} - T_i^{\mathrm{pbw}} + 0.01\right) \cdot \left(x_i^{\max} - x_i^{\mathrm{btl}}\right)$$

$$x_i^{\mathrm{dlv}} = \frac{x_i(t-d_i^{\mathrm{p}})}{y_\ell\left(t-d_{i,\ell}^{\mathrm{b}}\right)} \cdot \begin{cases} C_\ell & \text{if } q_\ell(t-d_{i,\ell}^{\mathrm{b}}) > 0 \\ y_\ell\left(t-d_{i,\ell}^{\mathrm{b}}\right) & \text{otherwise} \end{cases}$$

Experimental validation

Confirmation of prior insights

Generation of new insights

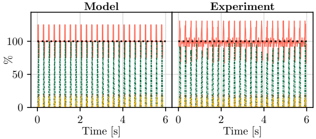

Theoretical stability analysis

Characterization of equilibria

Proof of asymptotic stability

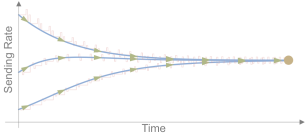

Experimental validation of BBR fluid model

Configuration

Topology
Dumbbell topology Single bottleneck

Congestion-control algorithms
Homogeneous or heterogeneous (balanced)

Experimental validation of BBR fluid model

Configuration

Topology

Dumbbell topology
Single bottleneck

Congestion-control algorithms

Homogeneous or heterogeneous (balanced)

Evaluation tools

Fluid-model simulator

Solution of differential equations (Method of steps)

Experiment environment

Emulation with Mininet
Load generation with iperf

Experimental validation of BBR fluid model

Configuration

Topology

Dumbbell topology
Single bottleneck

Congestion-control algorithms

Homogeneous or heterogeneous (balanced)

Evaluation tools

Fluid-model simulator

Solution of differential equations (Method of steps)

Experiment environment

Emulation with Mininet
Load generation with iperf

Result validation

Trace validation

Evolution of network metrics over time for single flow

Aggregate-result validation

Network metrics (aggregated over time) for multiple flows

Confirmation of prior insights: Unfairness of BBRv1

Previous insight: BBRv1 is unfair towards loss-sensitive CCAs in shallow buffers.

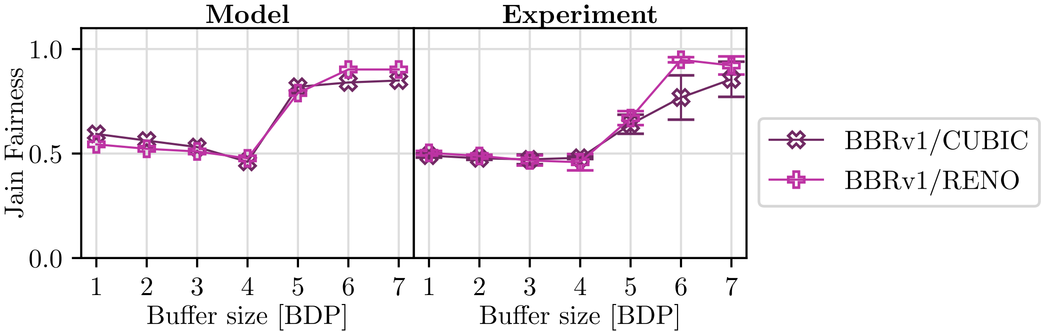

Confirmation of prior insights: Improved fairness in BBRv2

Previous insight: BBRv2 is quite fair to loss-based CCAs (under a drop-tail queuing discipline).

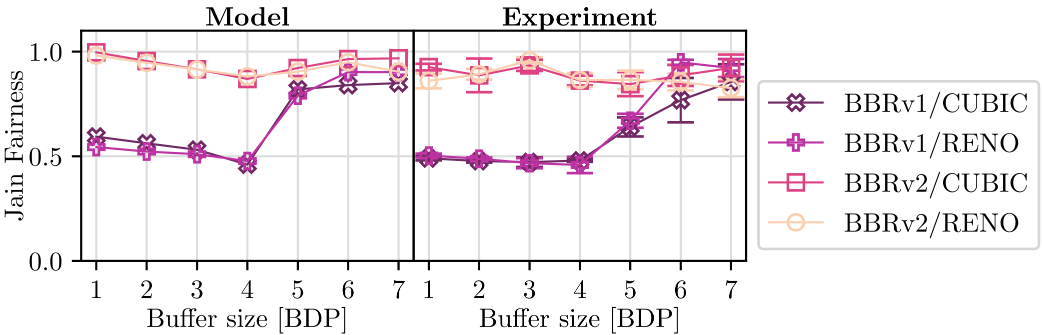

Generation of new insights: Limited fairness in BBRv2 under RED

New insight: BBRv2 is mildly unfair to loss-based CCAs under a RED queuing discipline.

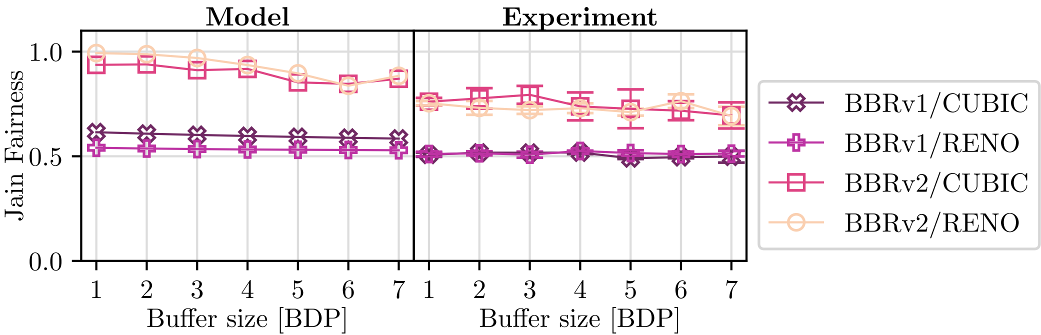

Confirmation of prior insights: High loss of BBRv1

Previous insight: BBRv1 leads to high loss in shallow buffers.

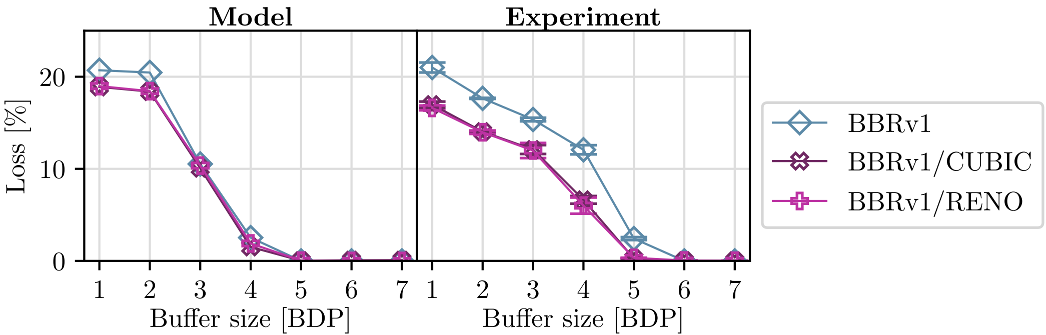

Confirmation of prior insights: Improved loss in BBRv2

Previous insight: BBRv2 leads to little loss (comparable to loss-based CCAs).

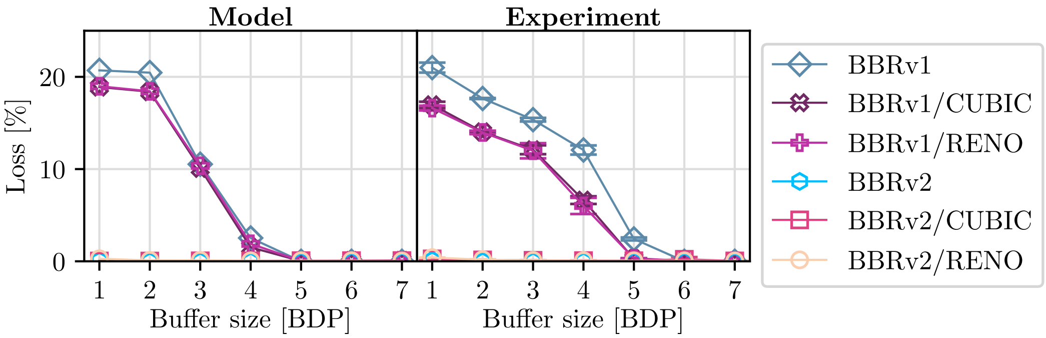

Generation of new insights: Bufferbloat in BBRv2

New insight: BBRv2 leads to intense queuing in large buffers.

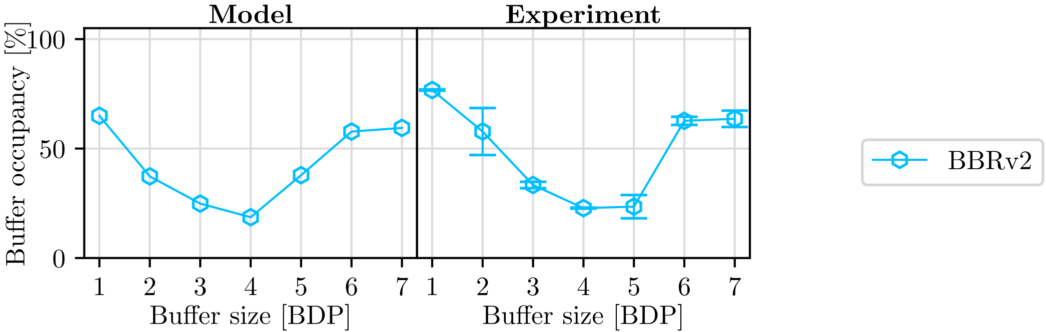

Generation of new insights: Bufferbloat in BBRv2

Generation of new insights: Bufferbloat in BBRv2

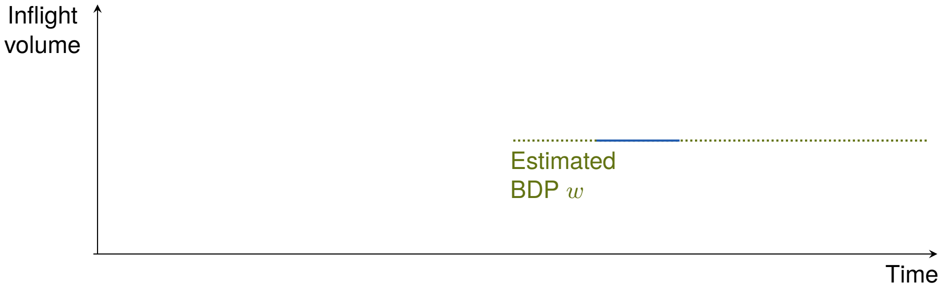

Generation of new insights: Bufferbloat in BBRv2

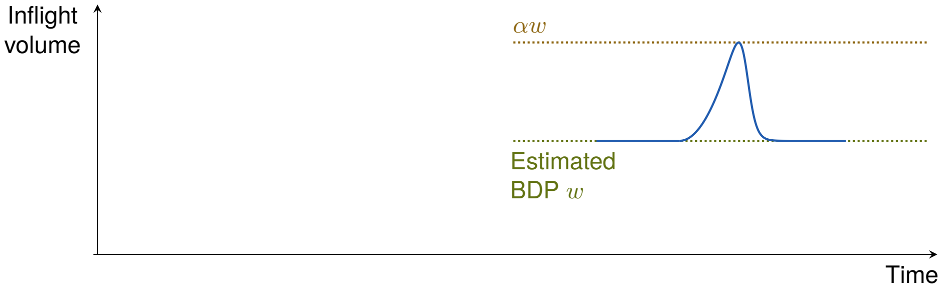

Generation of new insights: Bufferbloat in BBRv2

Generation of new insights: Bufferbloat in BBRv2

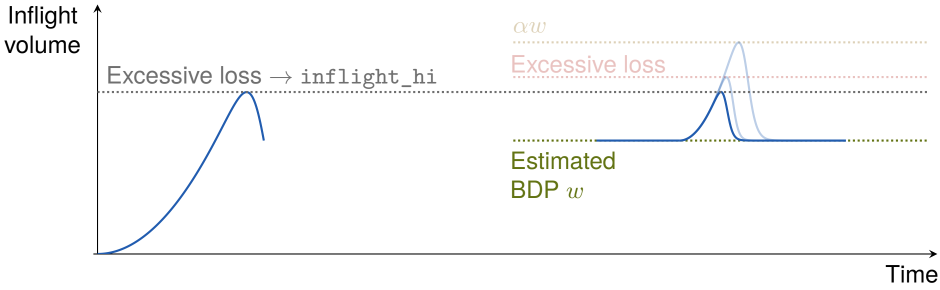

Generation of new insights: Bufferbloat in BBRv2

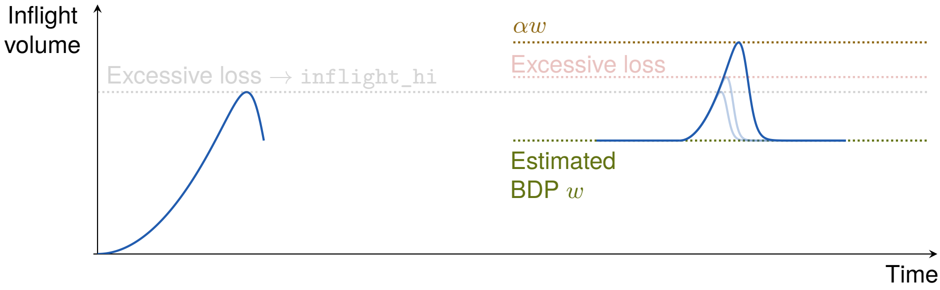

Large buffers disable loss-based safeguards

Generation of new insights: Bufferbloat in BBRv2

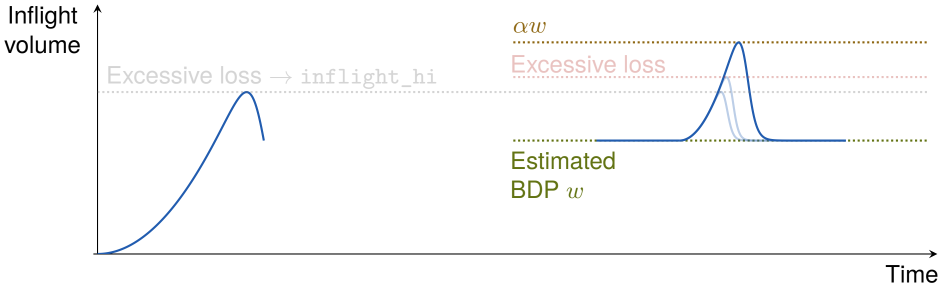

Large buffers disable loss-based safeguards

$\implies$ More aggressive probing $\implies$ Higher delivery rate
$\implies$ Higher estimated BDP $\implies$ Higher buffer utilization

Generation of new insights: Bufferbloat in BBRv2

Large buffers disable loss-based safeguards

$\implies$ More aggressive probing $\implies$ Higher delivery rate $\implies$ Higher estimated BDP $\implies$ Higher buffer utilization

Our fluid model reproduces this dynamic effect

Our contribution: A BBR analysis based on a fluid model

Fluid-model design

Formalization of BBR behavior

Design of new techniques

$$\dot{\tau}_i^{\min} = -\Gamma\left(\tau_i^{\min}(t) - \tau_i(t - d_i^{p})\right)$$

$$\dot{x}_i^{\text{btl}} = \sigma\left(t_i^{\text{pbw}} - T_i^{\text{pbw}} + 0.01\right) \cdot \left(x_i^{\max} - x_i^{\text{btl}}\right)$$

$$x_i^{\text{dlv}} = \frac{x_i(t - d_i^{p})}{y_\ell(t - d_{i,\ell}^{b})} \cdot \begin{cases} C_\ell & \text{if } q_\ell(t - d_{i,\ell}^{b}) > 0 \\ y_\ell\left(t - d_{i,\ell}^{b}\right) & \text{otherwise} \end{cases}$$

Experimental validation

Confirmation of prior insights

Generation of new insights

Theoretical stability analysis

Characterization of equilibria

Proof of asymptotic stability

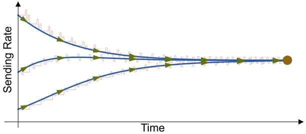

Theoretical stability analysis: Approach

Fluid model

Full fluid model
(used for simulation)

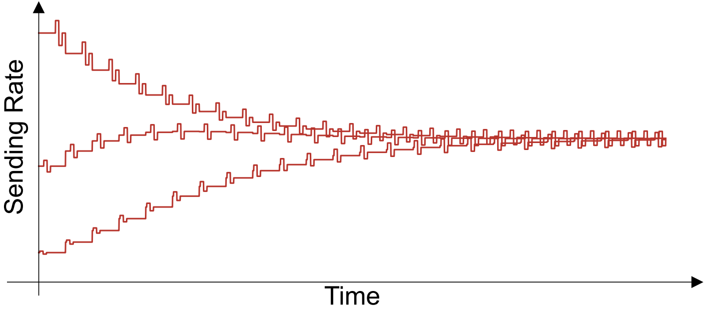

Theoretical stability analysis: Approach

Fluid model

Full fluid model
(used for simulation)

Reduced fluid model

High-level model
(macroscopic behavior)

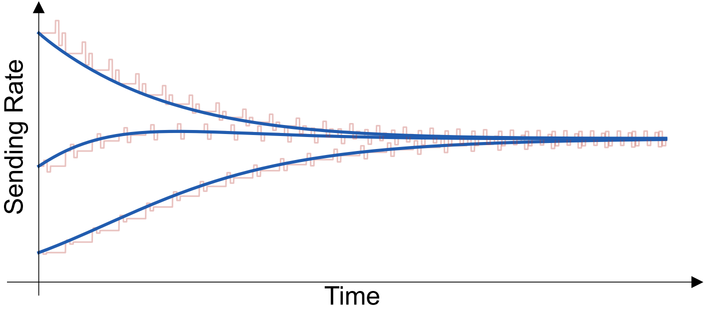

Theoretical stability analysis: Approach

Fluid model

Full fluid model
(used for simulation)

Reduced fluid model

High-level model
(macroscopic behavior)

Equilibria

Rate distribution &
queue length in steady state

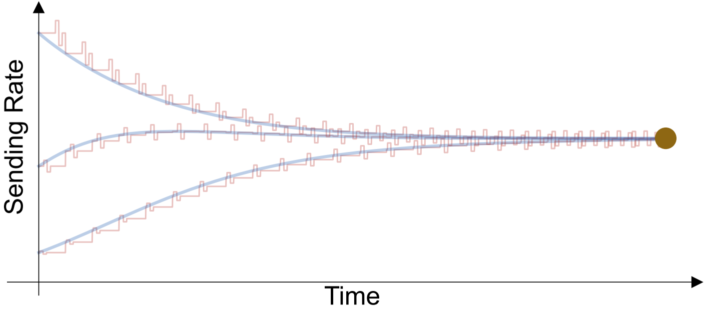

Theoretical stability analysis: Approach

Fluid model

Full fluid model
(used for simulation)

Reduced fluid model

High-level model
(macroscopic behavior)

Equilibria

Rate distribution &
queue length in steady state

Asymptotic stability

Proof of attractiveness
(Lyapunov method)

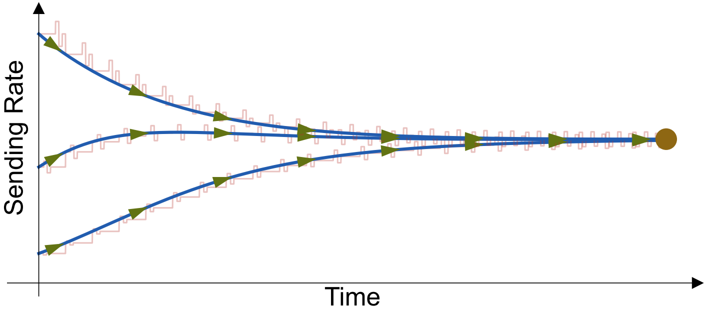

Theoretical stability analysis: Results

Theoretical stability analysis: Results

Equilibrium Type
BBRv1
BBRv2

Theoretical stability analysis: Results

Equilibrium Type
BBRv1
Deep buffers
Shallow buffers
BBRv2

Theoretical stability analysis: Results

Equilibrium Type	Uniqueness (Send Rates)
BBRv1	
Deep buffers	✗
Shallow buffers	
BBRv2	

Theoretical stability analysis: Results

Equilibrium Type	Uniqueness (Send Rates)	Guaranteed Fairness
BBRv1		
Deep buffers	✗	✗
Shallow buffers		
BBRv2		

Theoretical stability analysis: Results

Equilibrium Type	Uniqueness (Send Rates)	Guaranteed Fairness	Possible Fairness
BBRv1			
Deep buffers	✗	✗	✓
Shallow buffers			
BBRv2			

Theoretical stability analysis: Results

Equilibrium Type	Uniqueness (Send Rates)	Guaranteed Fairness	Possible Fairness	Loss Avoidance
BBRv1				
Deep buffers	✗	✗	✓	✓
Shallow buffers				
BBRv2				

Theoretical stability analysis: Results

Equilibrium Type	Uniqueness (Send Rates)	Guaranteed Fairness	Possible Fairness	Loss Avoidance	Asymptotic Stability
BBRv1					
Deep buffers	✗	✗	✓	✓	✓
Shallow buffers					
BBRv2					

Theoretical stability analysis: Results

Equilibrium Type	Uniqueness (Send Rates)	Guaranteed Fairness	Possible Fairness	Loss Avoidance	Asymptotic Stability
BBRv1					
Deep buffers	✗	✗	✓	✓	✓
Shallow buffers	✓	✓	✓	✗	✓
BBRv2					

Theoretical stability analysis: Results

Equilibrium Type	Uniqueness (Send Rates)	Guaranteed Fairness	Possible Fairness	Loss Avoidance	Asymptotic Stability
BBRv1					
Deep buffers	✗	✗	✓	✓	✓
Shallow buffers	✓	✓	✓	✗	✓
BBRv2					
Deep buffers	✗	✗	✓	✓	✓

Theoretical stability analysis: Results

Equilibrium Type	Uniqueness (Send Rates)	Guaranteed Fairness	Possible Fairness	Loss Avoidance	Asymptotic Stability
BBRv1					
Deep buffers	✗	✗	✓	✓	✓
Shallow buffers	✓	✓	✓	✗	✓
BBRv2					
Deep buffers	✗	✗	✓	✓	✓

Follow-up work: Stability does not hold if BBR competes with CUBIC!

Our contribution: A BBR analysis based on a fluid model

Fluid-model design

Formalization of BBR behavior

Design of new techniques

$$\dot{\tau}_i^{\min} = -\Gamma \quad \tau_i^{\min}(t) - \tau_i(t - d_i^{\mathrm{p}})$$

$$\dot{x}_i^{\mathrm{btl}} = \sigma\left(t_i^{\mathrm{pbw}} - T_i^{\mathrm{pbw}} + 0.01\right) \cdot \left(x_i^{\max} - x_i^{\mathrm{btl}}\right)$$

$$x_i^{\mathrm{dlv}} = \frac{x_i(t - d_i^{\mathrm{p}})}{y_\ell \quad t - d_{i,\ell}^{\mathrm{b}}} \cdot \begin{cases} C_\ell & \text{if } q_\ell(t - d_{i,\ell}^{\mathrm{b}}) > 0 \\ y_\ell\left(t - d_{i,\ell}^{\mathrm{b}}\right) & \text{otherwise} \end{cases}$$

Experimental validation

Confirmation of prior insights

Generation of new insights

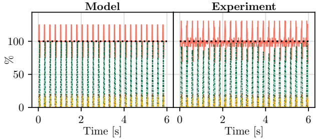

Theoretical stability analysis

Characterization of equilibria

Proof of asymptotic stability

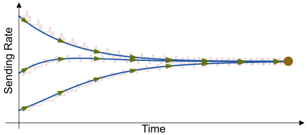

Conclusion

Conclusion

Fluid models

BBR & Congestion Control

Conclusion

Fluid models

Fluid models predict congestion-control behavior with **surprising accuracy** (qualitatively and quantitatively)

BBR & Congestion Control

Conclusion

Fluid models

Fluid models predict congestion-control behavior with **surprising accuracy** (qualitatively and quantitatively)

Fluid models are a **valuable complement** to experiments and steady-state models

BBR & Congestion Control

Conclusion

Fluid models

Fluid models predict congestion-control behavior with **surprising accuracy** (qualitatively and quantitatively)

Fluid models are a **valuable complement** to experiments and steady-state models

Fluid models might support **standardization efforts**, e.g., for parameter recommendation (see RFC 8312 for CUBIC)

BBR & Congestion Control

Conclusion

Fluid models

Fluid models predict congestion-control behavior with **surprising accuracy** (qualitatively and quantitatively)

Fluid models are a **valuable complement** to experiments and steady-state models

Fluid models might support **standardization efforts**, e.g., for parameter recommendation (see RFC 8312 for CUBIC)

BBR & Congestion Control

BBRv2 represents an **incomplete improvement** over BBRv1, e.g., regarding buffer usage

Conclusion

Fluid models

Fluid models predict congestion-control behavior with **surprising accuracy** (qualitatively and quantitatively)

Fluid models are a **valuable complement** to experiments and steady-state models

Fluid models might support **standardization efforts**, e.g., for parameter recommendation (see RFC 8312 for CUBIC)

BBR & Congestion Control

BBRv2 represents an **incomplete improvement** over BBRv1, e.g., regarding buffer usage

Difficulty of congestion control motivates proposals for **network-enforced resource allocation**, e.g., congestion shares [1], bandwidth reservation in SCION [2]

[1] Lloyd Brown, et al., On the Future of Congestion Control for the Public Internet, *HotNets 2020.*

[2] Giacomo Giuliari, et al., COLIBRI: A Cooperative Lightweight Inter-Domain Bandwidth-Reservation Infrastructure, *CoNEXT 2021.*

Conclusion

Fluid models

Fluid models predict congestion-control behavior with **surprising accuracy** (qualitatively and quantitatively)

Fluid models are a **valuable complement** to experiments and steady-state models

Fluid models might support **standardization efforts**, e.g., for parameter recommendation (see RFC 8312 for CUBIC)

BBR & Congestion Control

BBRv2 represents an **incomplete improvement** over BBRv1, e.g., regarding buffer usage

Difficulty of congestion control motivates proposals for **network-enforced resource allocation**, e.g., congestion shares [1], bandwidth reservation in SCION [2]

Efficient, fair, and stable Internet congestion control remains an **important research objective**

[1] Lloyd Brown, et al., On the Future of Congestion Control for the Public Internet, *HotNets 2020.*

[2] Giacomo Giuliari, et al., COLIBRI: A Cooperative Lightweight Inter-Domain Bandwidth-Reservation Infrastructure, *CoNEXT 2021.*

Our contribution: A BBR analysis based on a fluid model

Fluid-model design

Formalization of BBR behavior

Design of new techniques

$$\dot{\tau}_i^{\min} = -\Gamma \quad \tau_i^{\min}(t) - \tau_i(t - d_i^{\mathrm{p}})$$

$$\dot{x}_i^{\mathrm{btl}} = \sigma\left(t_i^{\mathrm{pbw}} - T_i^{\mathrm{pbw}} + 0.01\right) \cdot \left(x_i^{\max} - x_i^{\mathrm{btl}}\right)$$

$$x_i^{\mathrm{dlv}} = \frac{x_i(t - d_i^{\mathrm{p}})}{y_\ell \quad t - d_{i,\ell}^{\mathrm{b}}} \cdot \begin{cases} C_\ell & \text{if } q_\ell(t - d_{i,\ell}^{\mathrm{b}}) > 0 \\ y_\ell\left(t - d_{i,\ell}^{\mathrm{b}}\right) & \text{otherwise} \end{cases}$$

Experimental validation

Confirmation of prior insights

Generation of new insights

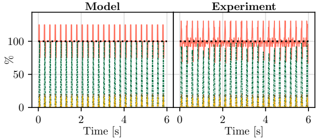

Theoretical stability analysis

Characterization of equilibria

Proof of asymptotic stability