Aalto University
School of Electrical
Engineering

How Impatience Affects the Performance and Scalability of P2P Video-on-Demand Systems

Samuli Aalto, Pasi Lassila
Aalto University, Finland
Petri Savolainen, Sasu Tarkoma
Helsinki Institute for Information Technology, Finland

MAMA'11
8 June 2011
San Jose, CA

Outline

- Peer-to-peer systems
- Models for P2P file sharing
- Fluid model for P2P VoD without impatience
- Models for P2P VoD with impatience

Fundamental principle

- Client/Server (CS) paradigm
  - Clients download content from servers
    - Clear distinction between the two roles
  - Service capacity remains the same, while load increases
  - Offered load bounded by the stability limit (for sure!)
- Peer-to-peer (P2P) systems
  - Peers download pieces of content from other peers/seeds and simultaneously upload downloaded pieces to other peers
    - Blurring of roles
  - Service capacity scales with the offered load
  - No stability limit (for sure?)

Applications

- P2P file sharing
  - Retrieve the whole file as soon as possible
  - Retrieve pieces in any order
- P2P streaming
  - Retrieve pieces at least at playback rate
  - Retrieve pieces in almost sequential order
- P2P video-on-demand (VoD)
  - Retrieve the whole file
  - Retrieve pieces at least at playback rate
  - Retrieve pieces in almost sequential order

Performance issues

- Scalability
  – Is the steady-state number of peers finite for any load?
- Stability
  – If not: Where is the stability limit for the load?
- Performance
  – If stable: Is the performance sufficient?
- Performance scalability
  – Is the performance sufficient for any load?

Outline

- Peer-to-peer systems
- Models for P2P file sharing
- Fluid model for P2P VoD without impatience
- Models for P2P VoD with impatience

Models for P2P file sharing

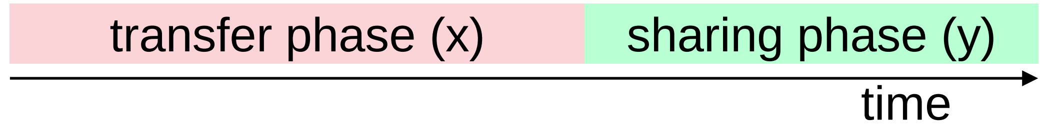

- Life span of a peer consists of two sequential phases:
  - file transfer phase, during which the peers are called leechers
  - sharing phase, during which the peers are called seeds
- Model by Qiu and Srikant (2004):
  - deterministic fluid model
    - nonlinear system of differential equations
  - describing system dynamics when sharing a single file
- Model by Menasche et al. (2009):
  - stochastic queueing model
    - utilizing $M/G/\infty$ queues (self-scaling property!)

Fluid model

- Switched nonlinear system:

$$x'(t) = \lambda - \theta \cdot x(t) - \phi(t)$$
$$y'(t) = \phi(t) - \gamma \cdot y(t)$$
$$\phi(t) = \min\{cx(t), \mu(\eta x(t) + y(t) + k)\}$$

NOTE!

For file sharing application:
$\eta \approx 1$

  – Unique steady-state solution either download constrained or upload constrained (depending on parameters)

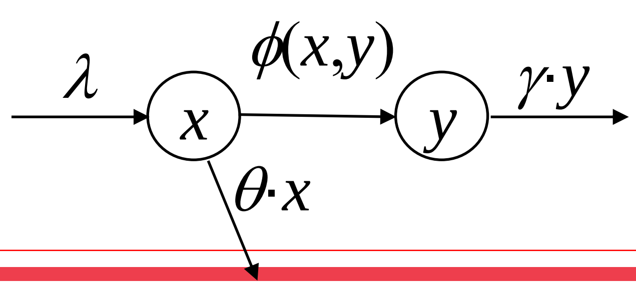

Deterministic model vs. Stochastic simulations

$y(t)$

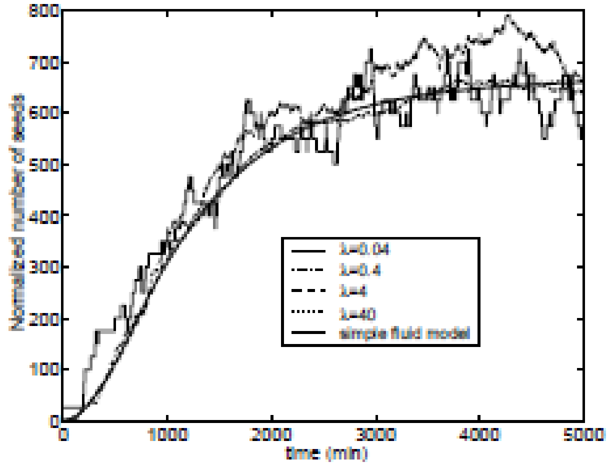

Figure 1: Experiment 1 : The evolution of the number of seeds as a function of time

$x(t)$

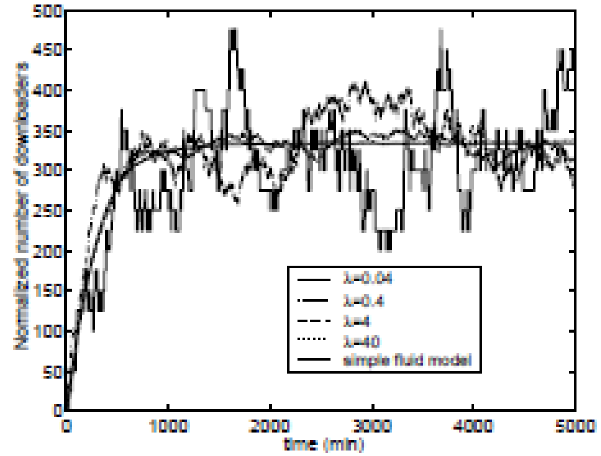

Figure 2: Experiment 1 : The evolution of the number of downloaders as a function of time

Source: Qiu and Srikant (2004)

Conclusions

- Scalability
  - System scalable for any $\eta > 0$
- Stability
  - Consequently, system stable for any $\lambda > 0$
- Performance
  - The mean file transfer time is

$$T = \frac{x}{\lambda} \leq \max\{\frac{1}{c}, \frac{1}{\eta}(\frac{1}{\mu} - \frac{1}{\gamma})\} \leq \max\{\frac{1}{c}, \frac{1}{\eta\mu}\} \approx \frac{1}{\mu}$$

  - Thus, no real problems in performance if reasonable upload rate with respect to the file size

NOTE!

For file sharing application:

$\eta \approx 1$

Outline

- Peer-to-peer systems
- Models for P2P file sharing
- Fluid model for P2P VoD without impatience
- Models for P2P VoD with impatience

Model for P2P VoD

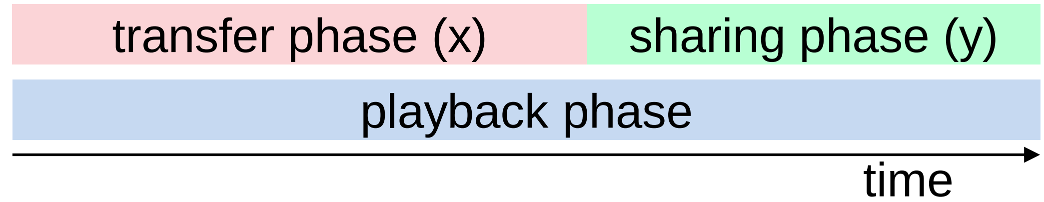

- Life span of a peer consists of two overlapping phases:
  - file transfer phase
  - playback phase
- Model by Aalto et al. (2010):
  - deterministic fluid model
    - system of differential equations
  - describing system dynamics when sharing a single video file
  - model takes explicitly into account the playback phase
  - worst case scenario:
    altruistic peers leave as soon as the playback phase is over;
    selfish peers leave already after the transfer phase

Fluid model (without impatience)

- Switched nonlinear system:

$$x'(t) = \lambda - \phi(t)$$

$$y'(t) = \varsigma \cdot \phi(t) - \frac{y(t)}{z - x(t)/\lambda}$$

$$\phi(t) = \min\{cx(t), \mu(\eta x(t) + y(t) + k)\}$$

NOTE!

For VoD application:

$\eta < 1$

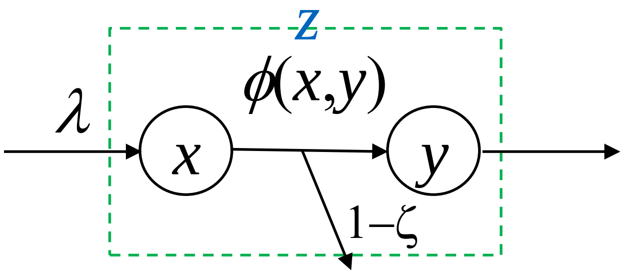

Steady-state synthesis (based on fluid model and stochastic simulations)

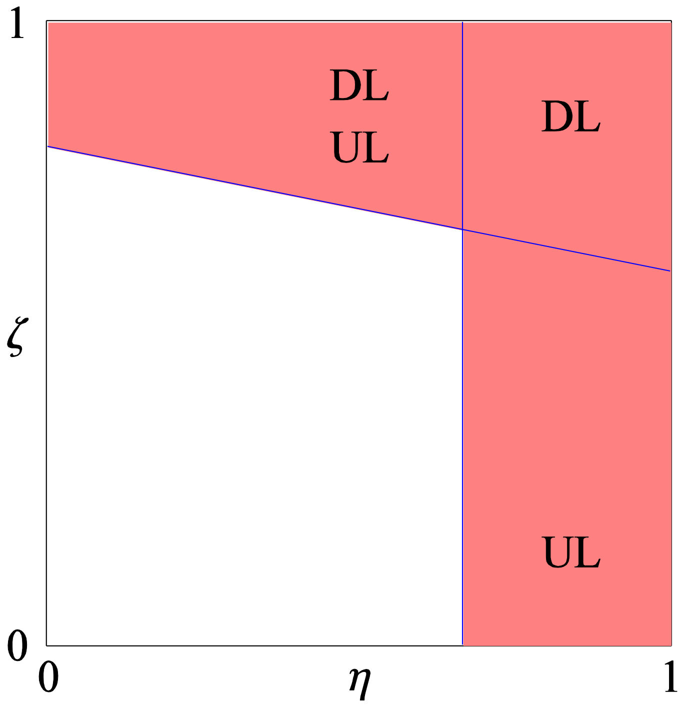

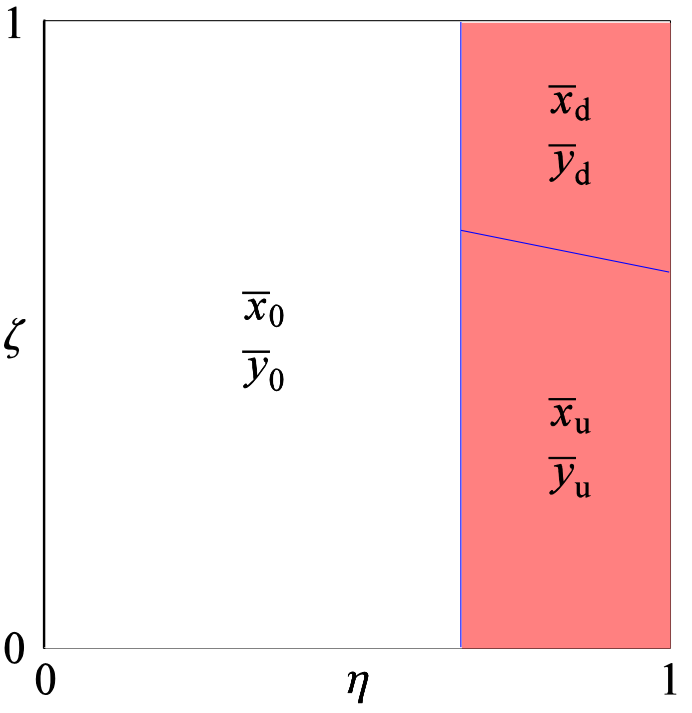

Fluid model vs. Stochastic and BitTorrent simulations

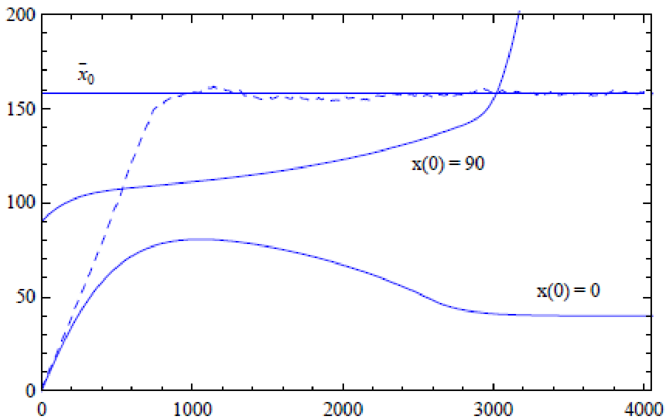

Fig. 5. Comparison of the fluid model (solid smooth lines) against the stochastic model (dashed line) with $\eta = 0.5$ and $\zeta = 0.8$.

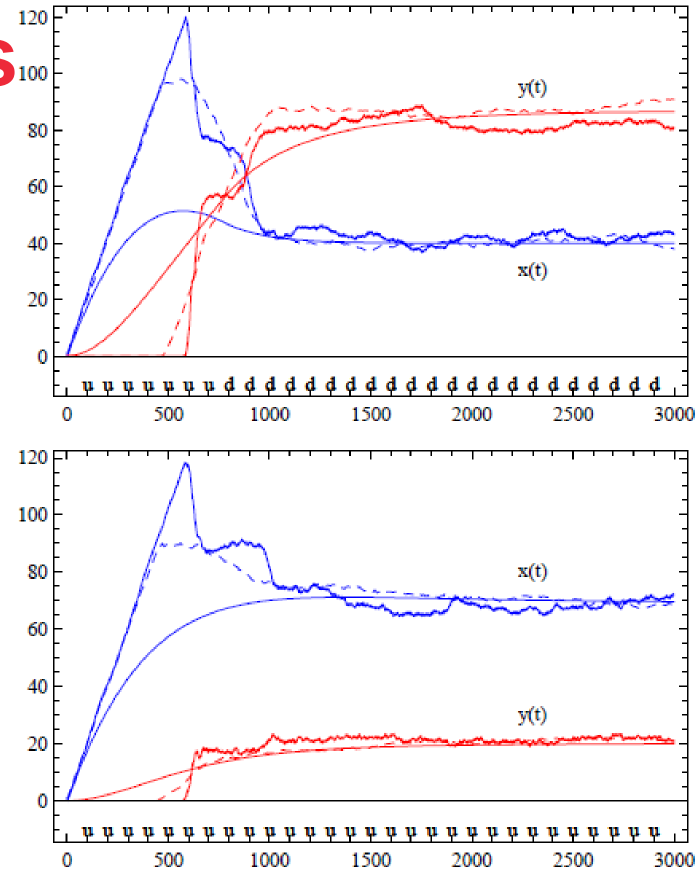

Fig. 4. Comparison of the fluid model (solid smooth line) against the stochastic model (dashed line) and the BitTorrent simulation (solid jagged line) with $\zeta = 0.9$ (upper panel) and $\zeta = 0.3$ (lower panel).

Performance threshold

- If

$$\eta > \frac{1}{z}\left(\frac{1}{\mu} - \frac{k}{\lambda}\right)$$

then transfer rate > playback rate,
i.e. sufficient playback quality

- But if

$$\eta < \frac{1}{z}\left(\frac{1}{\mu} - \frac{k}{\lambda}\right)$$

then transfer rate < playback rate,
i.e. playback quality problems

Conclusions

- Scalability
  - System scalable for any $\eta > 0$: $x \leq \lambda/(\eta\mu)$ for small η
- Stability
  - Consequently, system stable for any $\lambda > 0$
- Performance
  - Playback quality problems if η is too small: $\eta < 1/(z\mu) - k/(z\lambda)$
- Performance scalability
  - Performance scales if η is sufficiently large: $\eta > 1/(z\mu)$
  - Necessary condition for that: $\mu > 1/z$

Outline

- Peer-to-peer systems
- Models for P2P file sharing
- Fluid model for P2P VoD without impatience
- Models for P2P VoD with impatience

Fluid model (with impatience)

- Switched nonlinear system:

$$x'(t) = \lambda - \theta \cdot x(t) - \phi(t)$$

$$y'(t) = \varsigma \cdot \phi(t) - \frac{y(t)}{z - x(t)/\phi(t)}$$

$$\phi(t) = \min\{cx(t), \mu(\eta x(t) + y(t) + k)\}$$

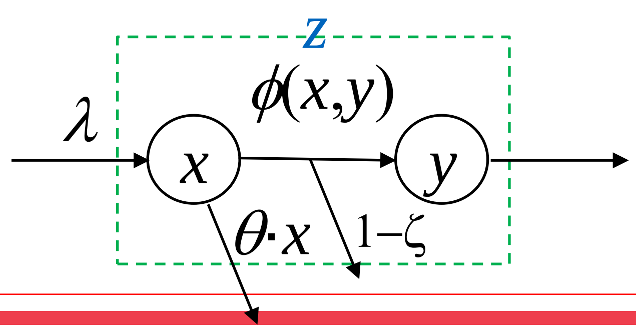

Fluid model vs. Stochastic simulations

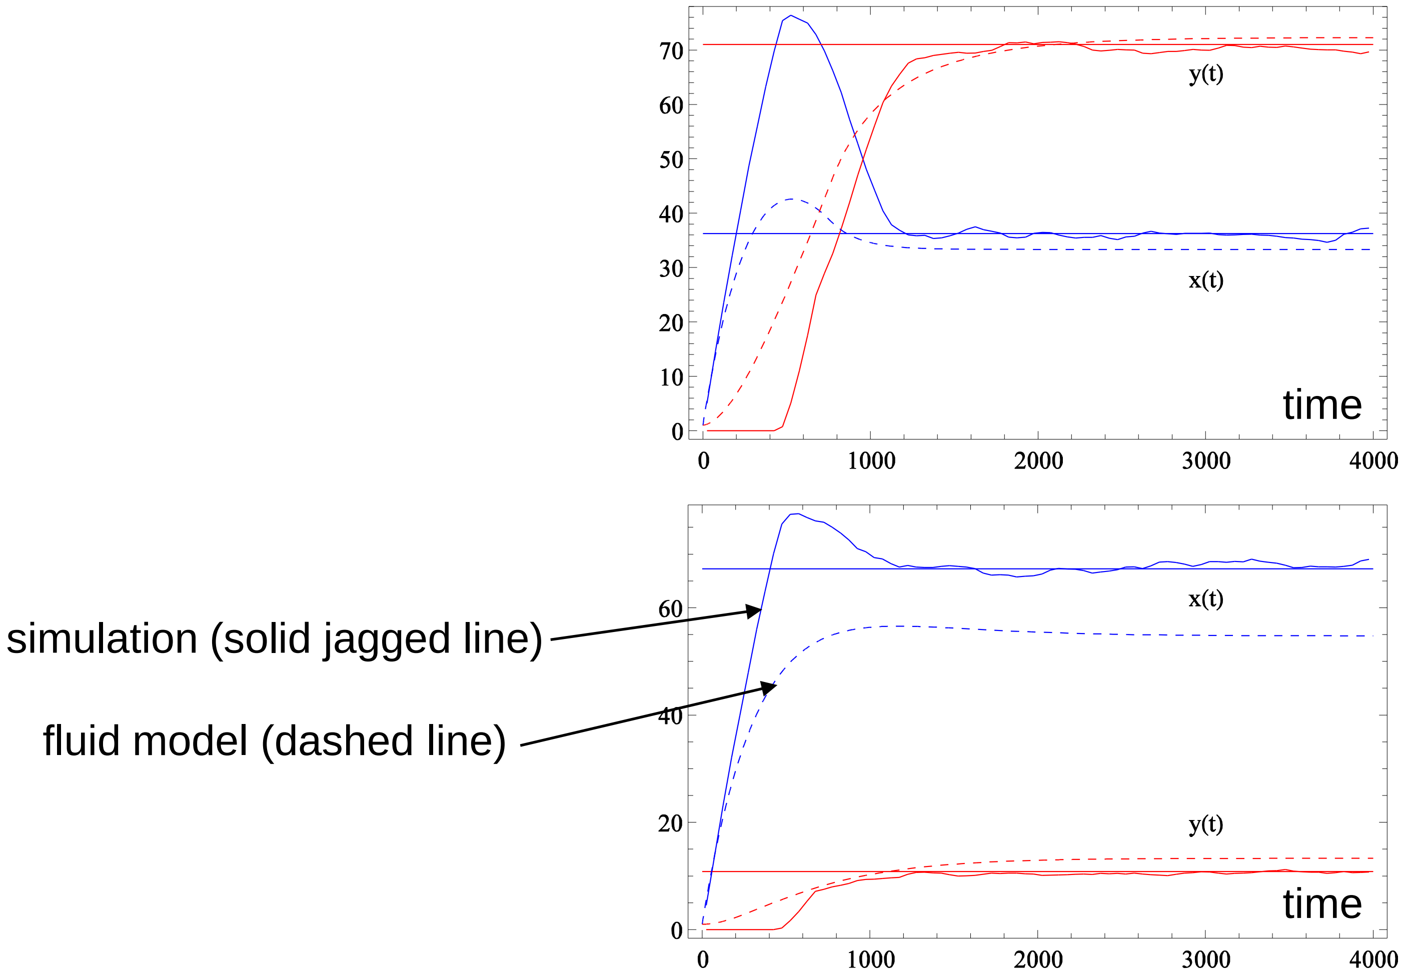

Approximative queueing model

- Pure download constrained case ($\mu \to \infty$):
  - utilizing M/G/∞ queues (self-scaling property!)
  - cf. Menasche et al. (2009)

$$\phi(t) = \min\{cx(t), \mu(\eta x(t) + y(t) + k)\} = cx(t)$$

$$x_{\mathrm{d}} = E[X] = \lambda \cdot E[\min\{A, \tfrac{1}{c}\}] = \frac{\lambda}{\theta}(1 - e^{-\theta / c})$$

$$y_{\mathrm{d}} = E[Y] = \lambda \cdot P\{A > \tfrac{1}{c}\} \cdot \varsigma \cdot (z - \tfrac{1}{c}) = \lambda e^{-\theta / c} \varsigma (z - \tfrac{1}{c})$$

Approximative queueing model (cont.)

- Pure upload constrained case ($c \to \infty$):
  - utilizing M/G/∞ queues (self-scaling property!)
  - cf. Menasche et al. (2009)

$$\phi(t) = \min\{cx(t), \mu(\eta x(t) + y(t) + k)\} \approx \mu(\eta x_{\mathrm{u}} + y_{\mathrm{u}} + k)$$

$$\tilde{\mu} = \mu(\eta + \frac{y_{\mathrm{u}} + k}{x_{\mathrm{u}}})$$

$$x_{\mathrm{u}} = E[X] = \lambda \cdot E[\min\{A, \tfrac{1}{\tilde{\mu}}\}] = \frac{\lambda}{\theta}(1 - e^{-\theta/\tilde{\mu}})$$

$$y_{\mathrm{u}} = E[Y] = \lambda \cdot P\{A > \tfrac{1}{\tilde{\mu}}\} \cdot \varsigma \cdot (z - \tfrac{1}{\tilde{\mu}}) = \lambda e^{-\theta/\tilde{\mu}} \varsigma (z - \tfrac{1}{\tilde{\mu}})$$

Approximative queueing model vs. Stochastic simulations

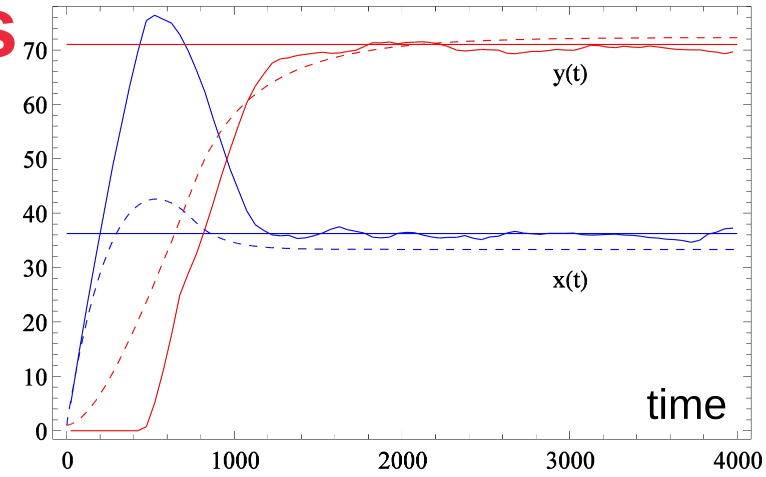

queueing model (horizontal line)

simulation (solid jagged line)

fluid model (dashed line)

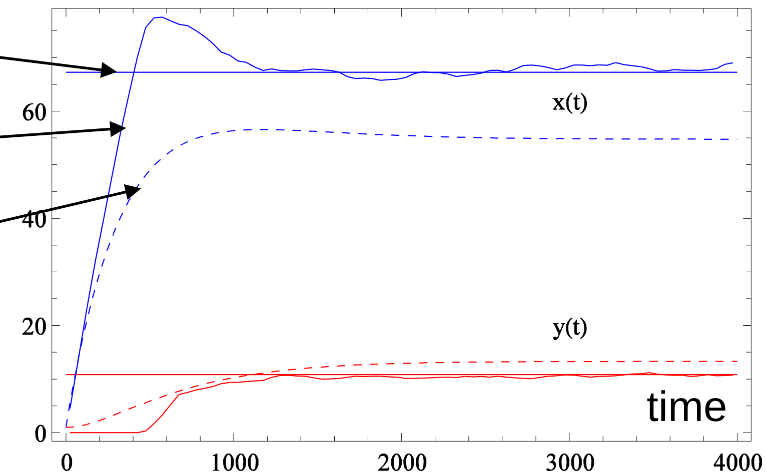

Impact of the impatience parameter

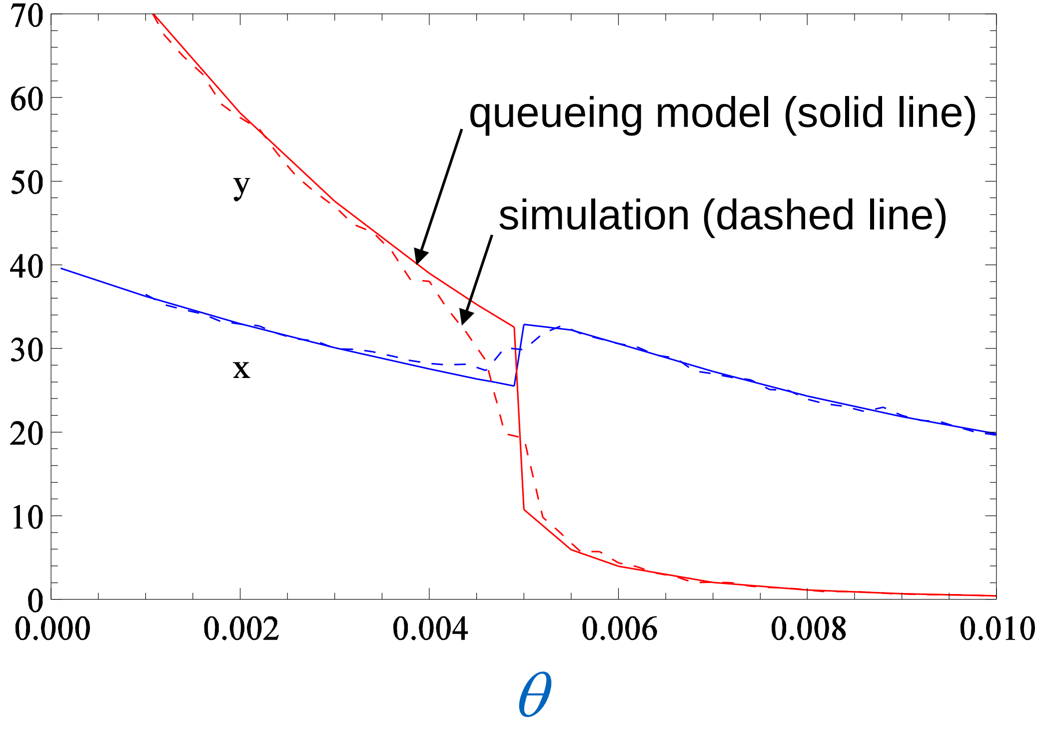

Performance threshold

- Approximative queueing model shows a qualitatively different behavior when η is below a certain threshold
- The critical value η_0 is determined by requiring that the (approximate) transfer rate in the upload constrained case equals the playback rate

$$\tilde{\mu} = \frac{1}{z} \quad \Rightarrow \quad \eta_0 = \frac{1}{z}\left(\frac{1}{\mu} - \frac{k\theta z}{\lambda(1-e^{-\theta z})}\right)$$

$$x_0 = x_u \,|_{\varsigma=0}$$

$$y_0 = y_u \,|_{\varsigma=0} = 0$$

Approximative queueing model vs. Stochastic simulations

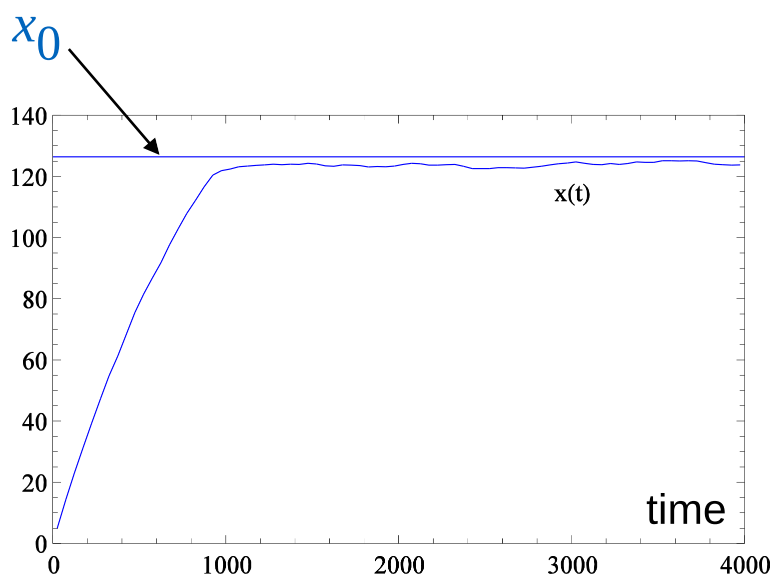

$\eta < \eta_0$

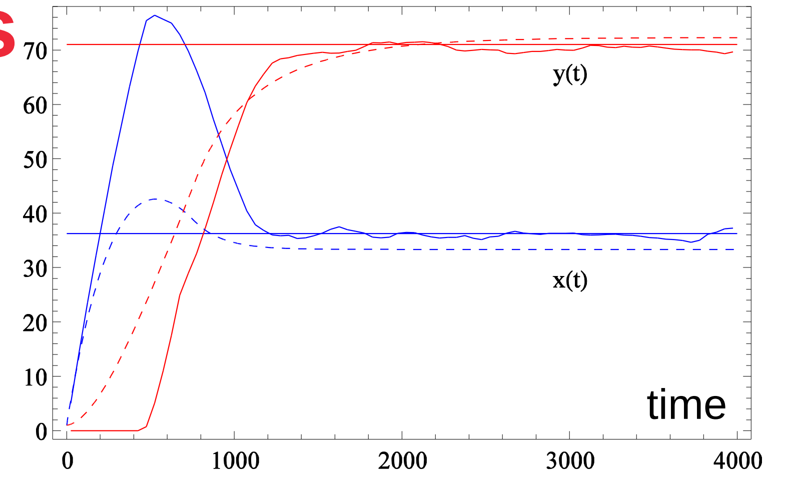

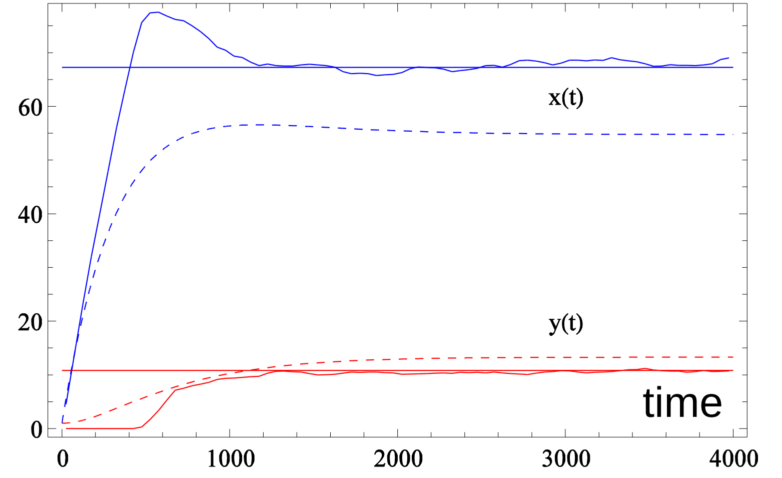

Aalto University
School of Electrical
Engineering

Performance threshold (cont.)

- If

$$\eta > \eta_0 = \frac{1}{z}\left(\frac{1}{\mu} - \frac{k\theta z}{\lambda(1-e^{-\theta z})}\right)$$

then transfer rate > playback rate,
i.e. sufficient playback quality

- But if

$$\eta < \eta_0 = \frac{1}{z}\left(\frac{1}{\mu} - \frac{k\theta z}{\lambda(1-e^{-\theta z})}\right)$$

then transfer rate < playback rate,
i.e. playback quality problems

Conclusions

$$\eta_0 = \frac{1}{z}\left(\frac{1}{\mu} - \frac{k\theta z}{\lambda(1-e^{-\theta z})}\right) \le \frac{1}{z}\left(\frac{1}{\mu} - \frac{k}{\lambda}\right) = \eta_0 \Big|_{\theta=0}$$

- Thus, the most stringent conditions concerning the playback quality are related to the case with the least amount of impatience: $\theta = 0$

Conclusions (cont.)

- Scalability
  - System scalable for any $\eta > 0$
- Stability
  - Consequently, system stable for any $\lambda > 0$
- Performance
  - Playback quality problems if η is too small: $\eta < \eta_0$
- Performance scalability
  - Performance scales if η is sufficiently large: $\eta > 1/(z\mu)$
  - Necessary condition for that: $\mu > 1/z$

The End