PBoT: An Analytical Model of P2P Traffic Based on Transactions

TELEFÓNICA I+D

Date: 2nd October 2008

Francisco Javier Ramón Salguero (fjrs@tid.es)
Gerardo García de Blas (ggdb@tid.es)
Adrián Maeso Martín-Carnerero (ammc@tid.es)
María L. García Osma
José Enríquez Gabeiras

Index

Index

01 Motivation

- What are we looking for:
  - How performance (download time) and load are affected by:
    - Changes in available BW → Access links
    - Changes in traffic demand:
      - Pattern of requests
      - Content size
      - Number of parallel downloads
- Why do we need a simplified model?
  - Per application models are not adequate for network planning
    - Too many parameters → hard to estimate and compute
    - Each application requires a whole different model!

Index

02 General analytical model for P2P traffic

Model for content requests

- P2P content requests are not a completely exogenous parameter
  - **The faster the contents can be downloaded, the greater motivation for the users to make new requests.**
    - An indeterminate amount of potential P2P traffic demand may be dormant
- **Model suppositions** for the "average user":
  1. **Each user needs a certain amount of time to consume a downloaded content.**
     - "Consumption" does not always imply an actual use of the content, but any actions that lead users to stop sharing a given file (e.g. copies to other supports, deletions from the hard disk, etc.)
  2. **In the long term, a user will not request a new content until, at least, a previously downloaded content has been "consumed" (stability).**
     - Otherwise, contents would be queued ad infinitum and shared (stored) forever
  3. **Each user will not exceed a maximum number of simultaneous pending downloads.**
     - In the long term, users would be actually replacing the "active" requests by new ones after finishing the download and "consumption"

02 General analytical model for P2P traffic

Model for content requests

- E.g. Model for 2 simultaneous requests per user:

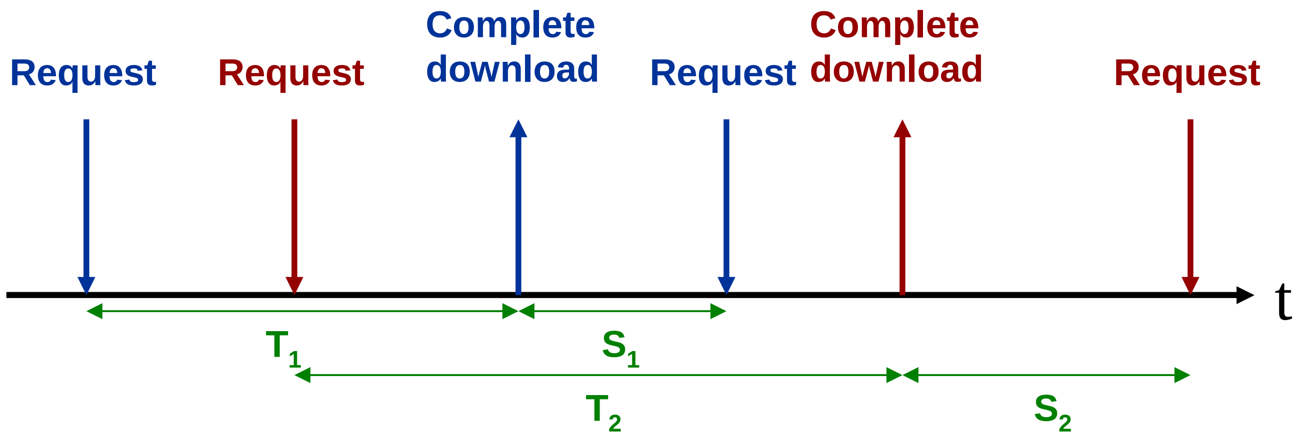

Model parameters:

- *T*: Average total download time of a complete file
- *S*: Average consumption time
  - Waiting time (since a download is complete until a new request is made)
- N_{pet}: average maximum number of active requests per user

02 General analytical model for P2P traffic

General simplified queuing model for P2P systems

- A P2P network can be seen as a complex system which
  - receives requests from users at a given global rate: λ
  - serves them at a global rate, using users' uplinks: **μ**

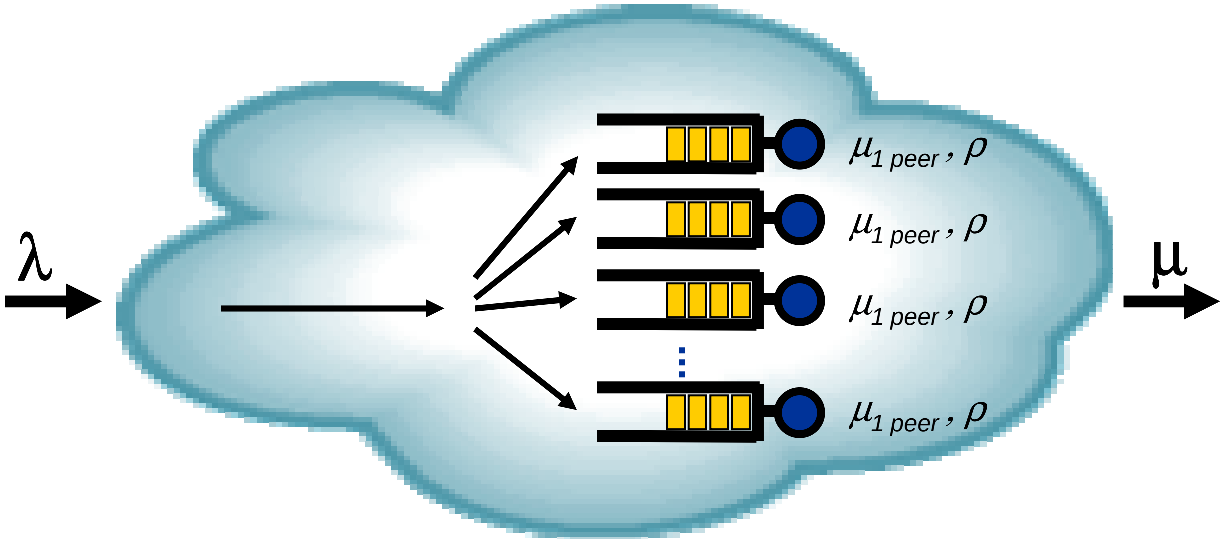

- **Each peer can be seen both as contributor to the global request rate and as server** with a bitrate equal to its uplink capacity.

02 General analytical model for P2P traffic

General simplified queuing model for P2P systems

- The P2P application determines:
  - The use of these “servers”
  - The queuing disciplines behind them, and
  - The way the jobs are distributed among servers

- What can be said about these queuing systems (in general)?
  - What is the average number of occupied servers?
  - What is the relation between λ and μ?
  - What happens in the most popular P2P applications?

02 General analytical model for P2P traffic

General simplified queuing model for P2P systems

- According to Little's Law, the average number of occupied servers (peers' uplinks) is

$$\bar{n} = \frac{\lambda}{\mu_{peer}} = \frac{\lambda}{\mu} \cdot N_{users}$$

  - *ñ*: average number of occupied "servers", i.e. peers' uplinks
  - N_{users}: number of users in the P2P system
  - μ_{peer}: serving rate per "server" (peer). In practice, it is a representation of uplink bandwidth.

- Besides, the occupation probability of a single server (ρ) can be calculated as

$$\rho = \frac{\bar{n}}{N_{users}}$$

- Hence

$$\lambda = \rho \cdot \mu \qquad \textbf{[1]}$$

02 General analytical model for P2P traffic

General simplified queuing model for P2P systems

- On the other hand, we know that the occupation probability of a peer's uplink can be approximated by the following expression

$$\rho = 1 - \left[\frac{\ln(m)}{m} \right]^{N_{users}}$$

m: Number of chunks per file

D. Qiu and R. Srikant, "Modeling and Performance Analysis of BitTorrent-Like Peer-Peer Networks", in Proc. of SIGCOMM 2004, September 2004.

- However, even with few users and chunks, ρ is very close to 1
  - E.g.
    - 5 users
Content size: 5 MBytes ➔ ρ = **0.99992**
Chunk size: 256 KBytes
    - 5 users
Content size: 10 MBytes ➔ ρ = **0.999993**
Chunk size: 256 KBytes
- Hence, from [1] we can safely assume that $\lambda \approx \mu$

02 General analytical model for P2P traffic

PBoT Model

- Taking into account the request model…

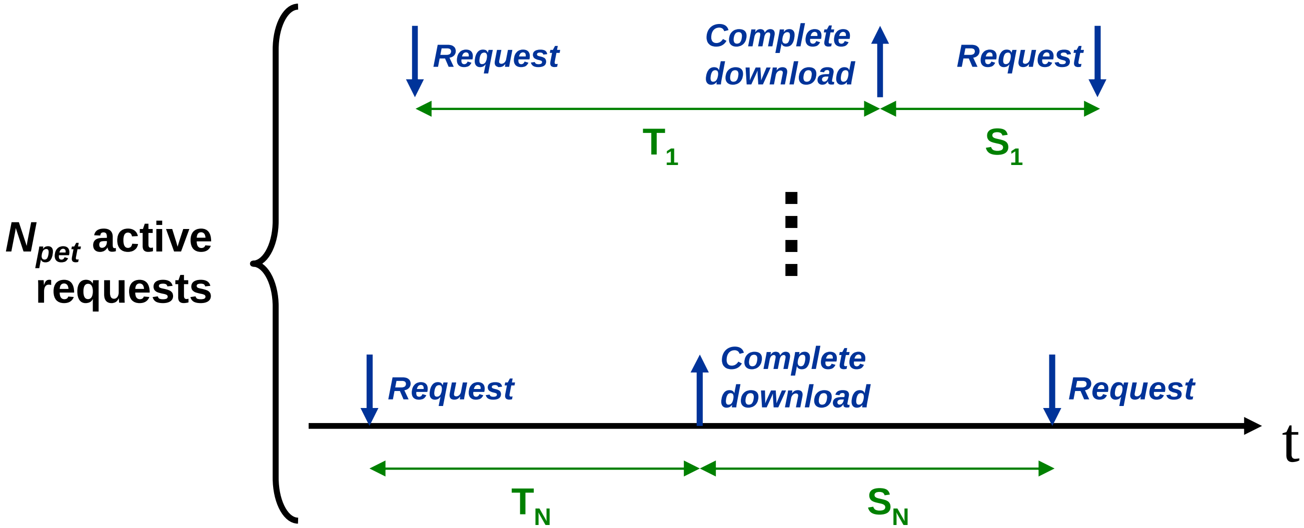

- … we can trivially calculate λ:

$$\lambda = \frac{N_{pet}}{T+S} \cdot N_{users}$$

02 General analytical model for P2P traffic

PBoT Model

- On the other hand, assuming that most users in the P2P system have their uplink capacities saturated, we have that:

$$\mu = \frac{\mathrm{BW}_{\mathrm{up,\,user}}}{F} \cdot \mathrm{N}_{\mathrm{users}}$$

F: File size

- Finally, taking into account that, in equilibrium, λ ≈ μ, we get the final expression of the PBoT model:

$$T + S = \frac{\mathrm{N}_{\mathrm{pet}} \cdot F}{BW_{up,user}}$$

02 General analytical model for P2P traffic

Assuming that the PBoT model is right…

$$T + S = \frac{N_{pet} \cdot F}{BW_{up,user}}$$

1. Equilibrium does not depend on the number of users in the system
2. Users have influence on the point of equilibrium via:
   - Average "consumption time" (*S*)
   - Average number of active contents (N_{pet})
   - Average content size (*F*)
3. Network operators have influence on the equilibrium via uplink BW
   - Side effect on offered load in bps $\lambda = \frac{BW_{up,user}}{F} \cdot N_{users}$
4. In normal working conditions and stationary state, $T + S = \text{constant}$
   - **Direct linear tradeoff between average download time and average silence time**

Index

03 Model validation

Introduction

- Are we missing something relevant in our model?
  - Let’s validate it by simulating an actual P2P system
- ***Burrito* simulator**:

  M.L. García et al., “Burrito: A simulation tool for P2P networks”, in Proceedings 3rd International Workshop on Internet Performance, Simulation, Monitoring and Measurement IPS-MoMe 2005

  - Freely available
  - It emulates a large **eMule** network, including all the queuing and rating systems
  - **Any other simulators or P2P applications would have done!**
- Parameters:
  - S: avg.=1 day, exponential
  - N_{pet}: avg.=7, binomial
  - F = 700 MBytes
  - $BW_{up,user}$ = 128 kbps
  - $BW_{down,user}$ = 512 kbps
  - N_{user} = 1000
- **Confidence interval: 95 %**

03 Model validation

Typical evolution of download times

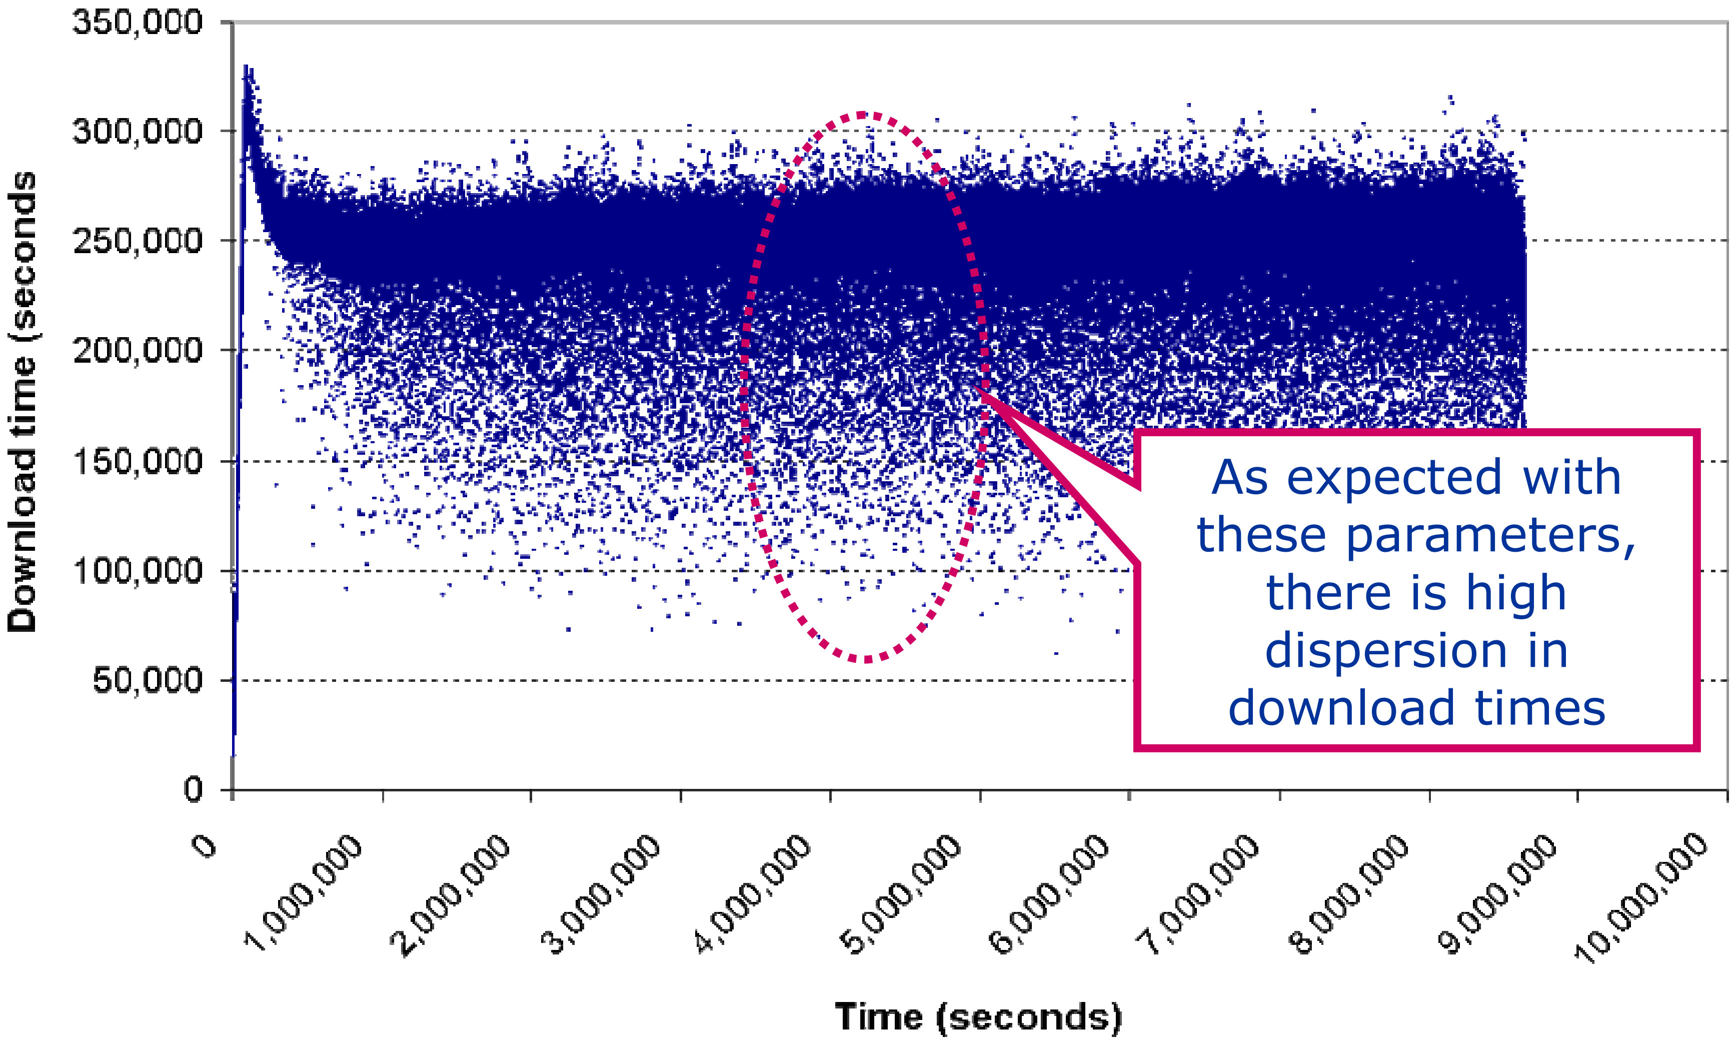

03 Model validation

Download time vs. number of users

$$T = \frac{N_{pet} \cdot F}{BW_{up,user}} - S$$

- Is the equilibrium independent on the number of users?

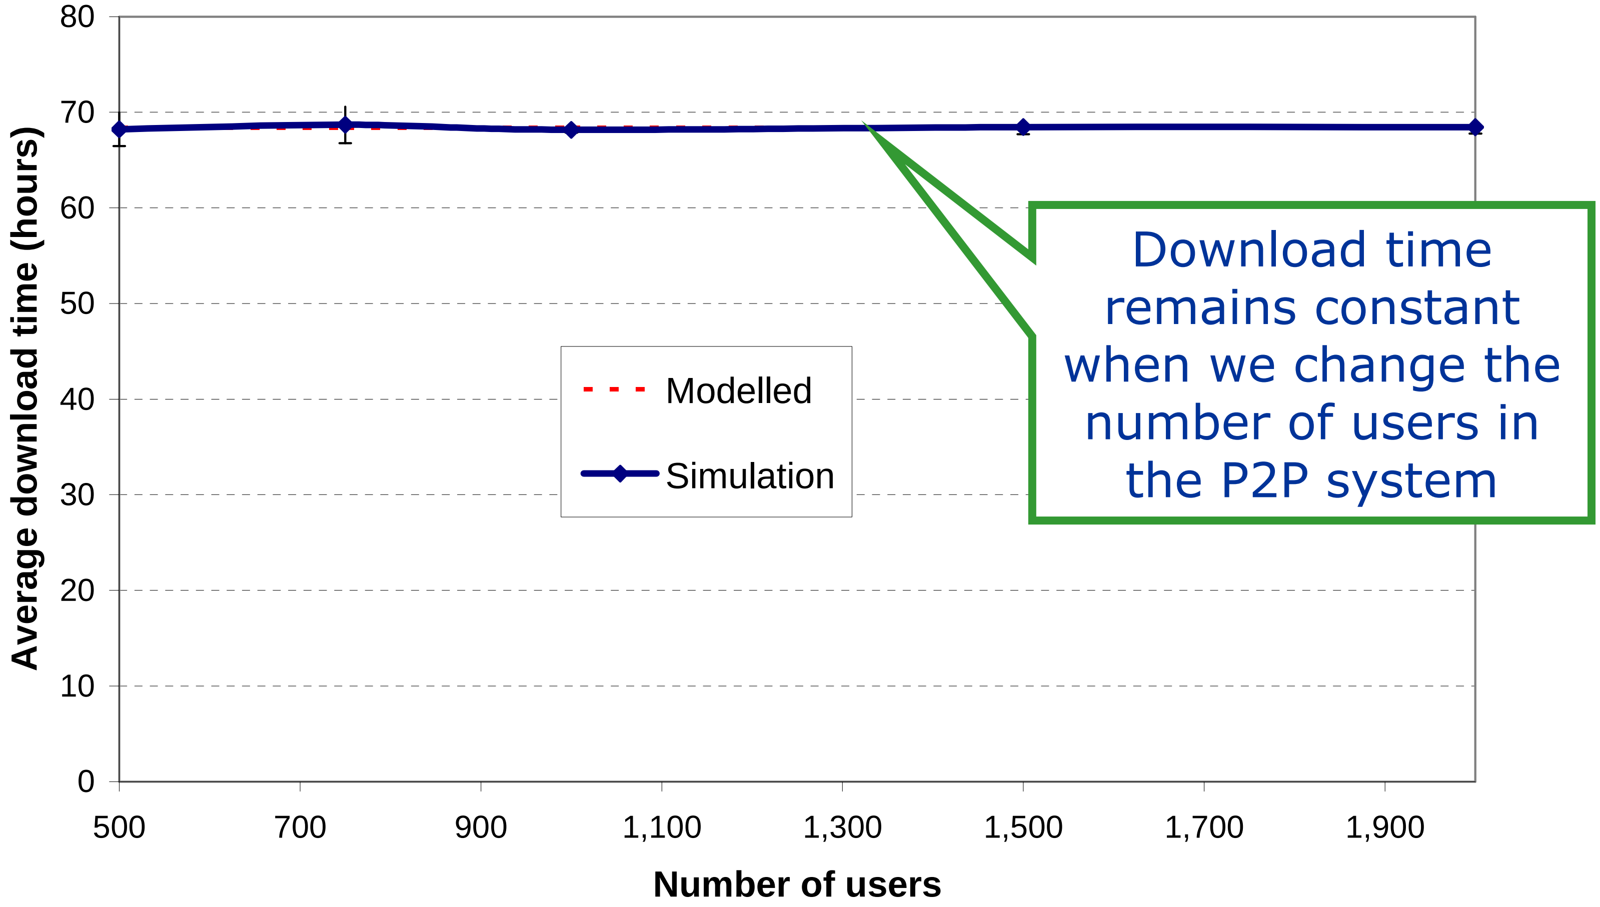

03 **Model validation**

Download time vs. downstream BW

$$T = \frac{N_{pet} \cdot F}{BW_{up,user}} - S$$

- Is the equilibrium related to downlink bandwidth in any unexpected way?

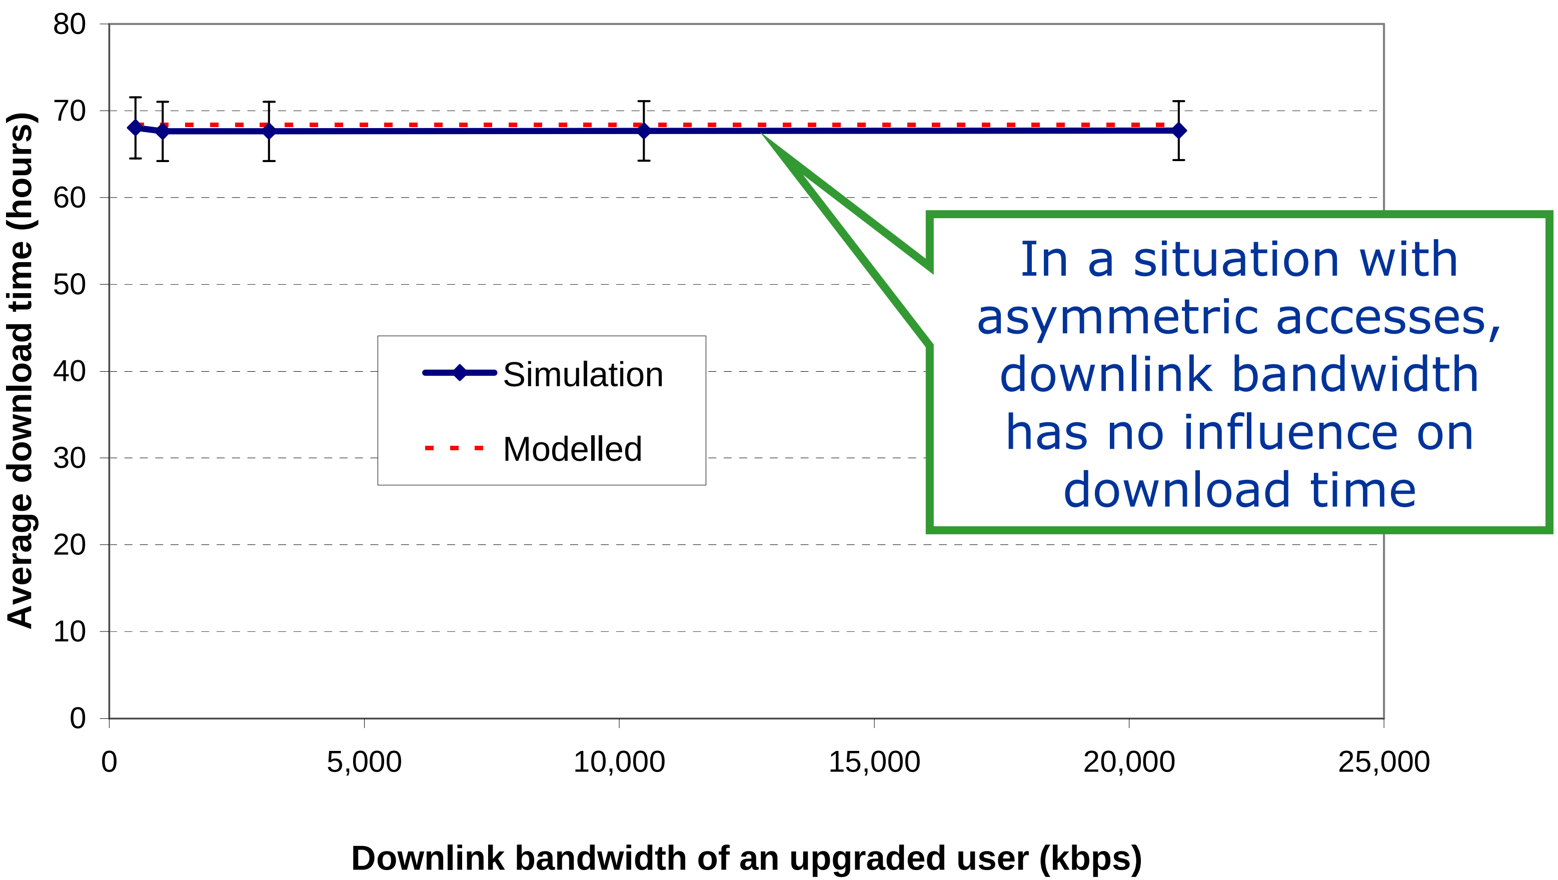

Model validation

Download time vs. silence time

$$T + S = \frac{N_{pet} \cdot F}{BW_{up,user}}$$

- Is there a linear tradeoff between T and S, as predicted by the model?

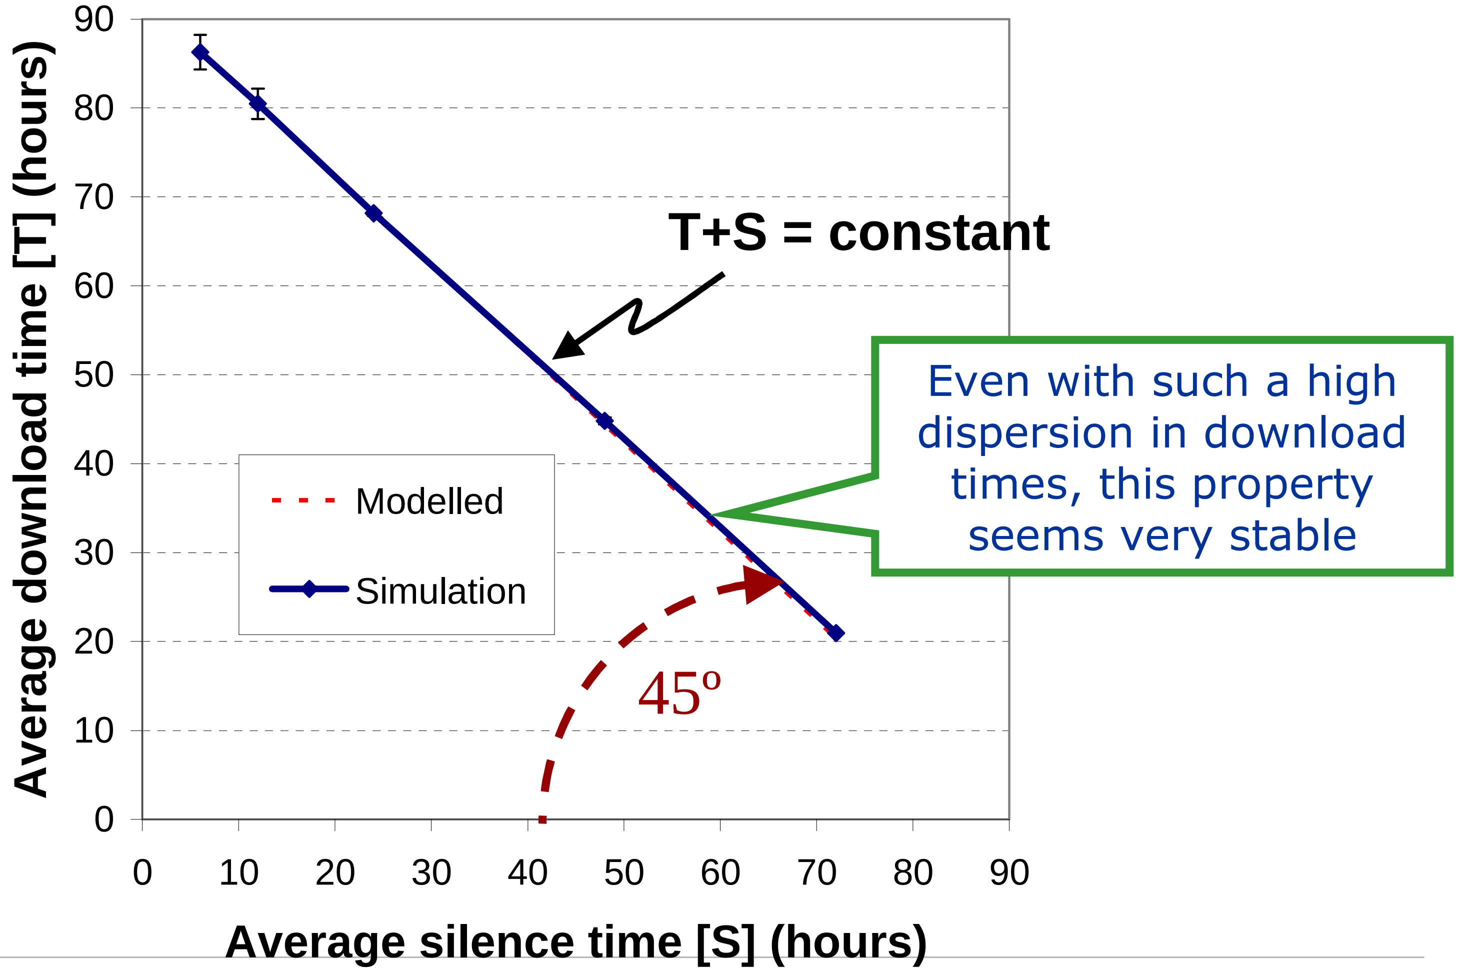

03 Model validation

Download time vs. simultaneous requests

$$T = \left(\frac{F}{BW_{up,user}} \right) \cdot N_{pet} - S$$

- Is there a simple linear relationship between download time and average number of simultaneous requests per user?

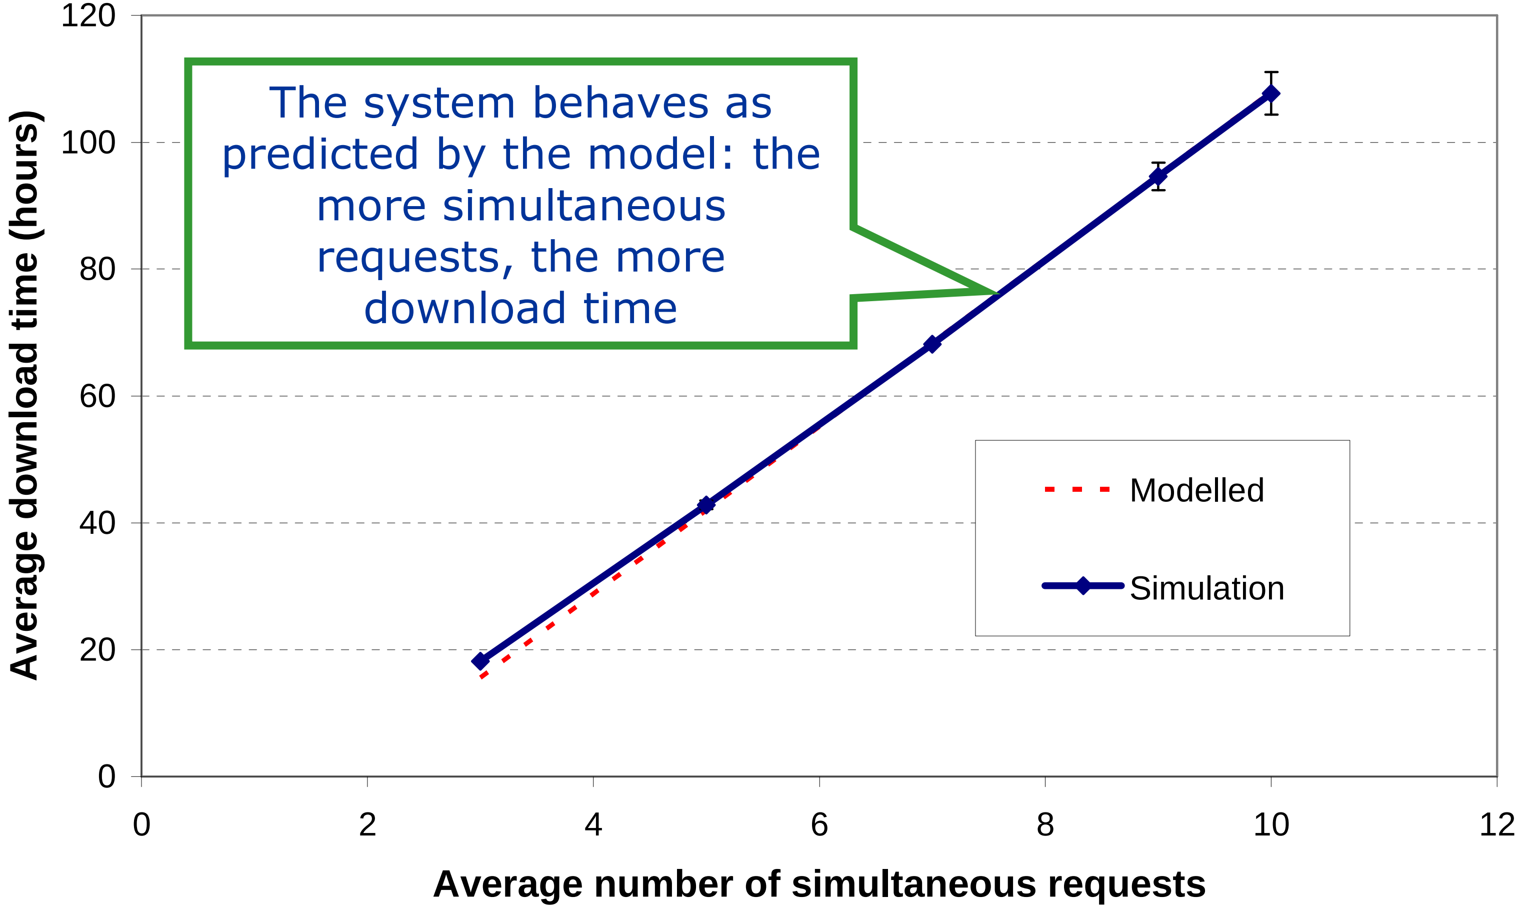

03 Model validation

Download time vs. content size

$$T = \left(\frac{N_{pet}}{BW_{up,user}} \right) \cdot F - S$$

- Is there a simple linear relationship between download time and average content size?

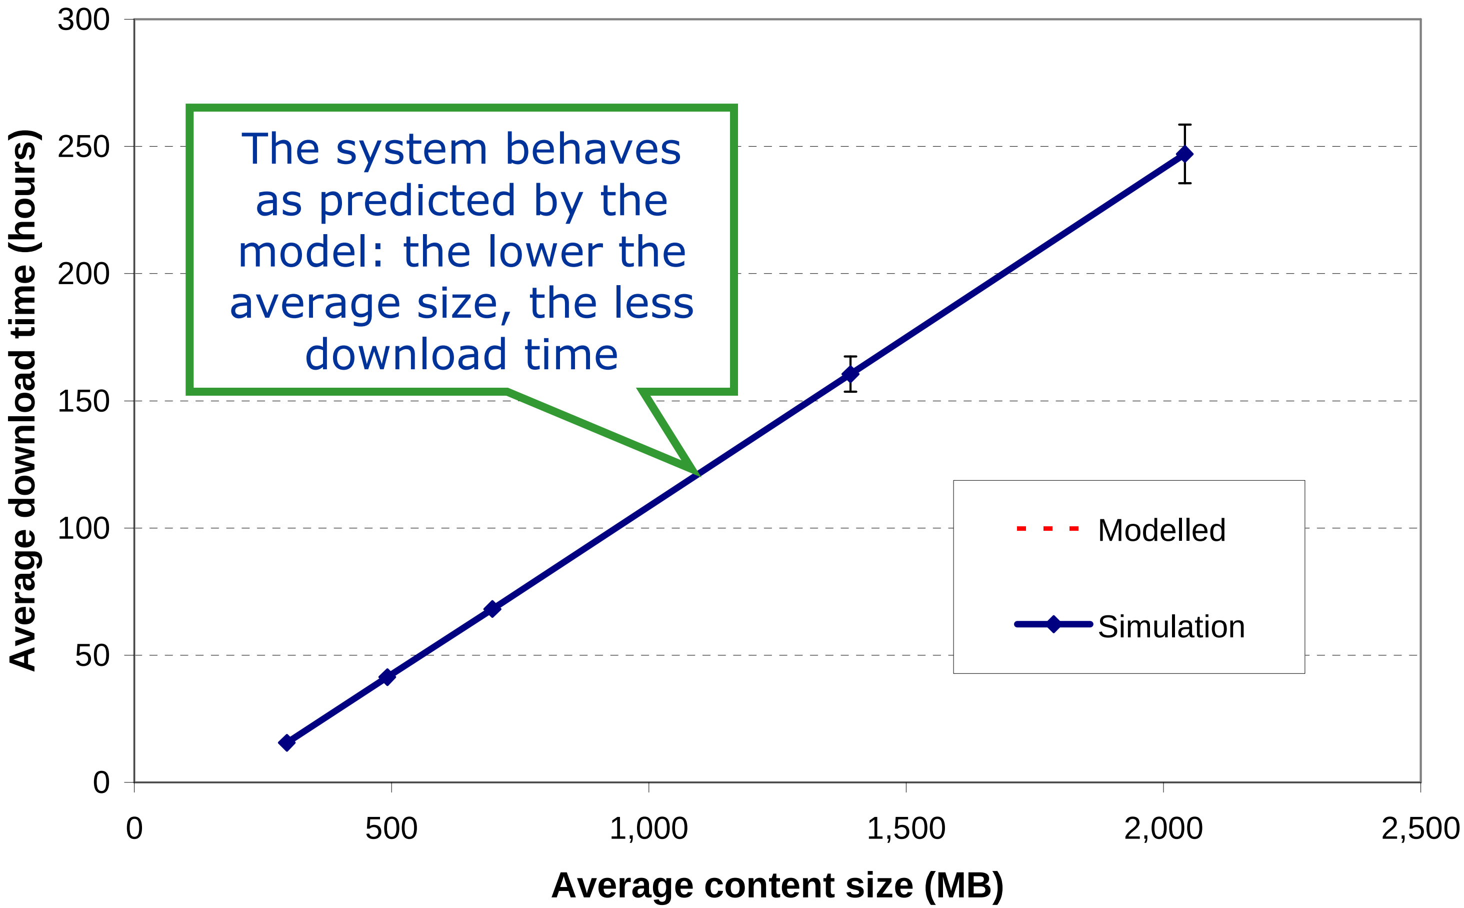

03 Model validation

Download time vs. uplink BW

$$T = \left(N_{pet} \cdot F\right) \cdot \frac{1}{BW_{up,user}} - S$$

$$\left(while\ \lambda \approx \mu\right)$$

- Is there an hyperbolic relationship between download time and uplink bandwidth while the latter is the bottleneck?

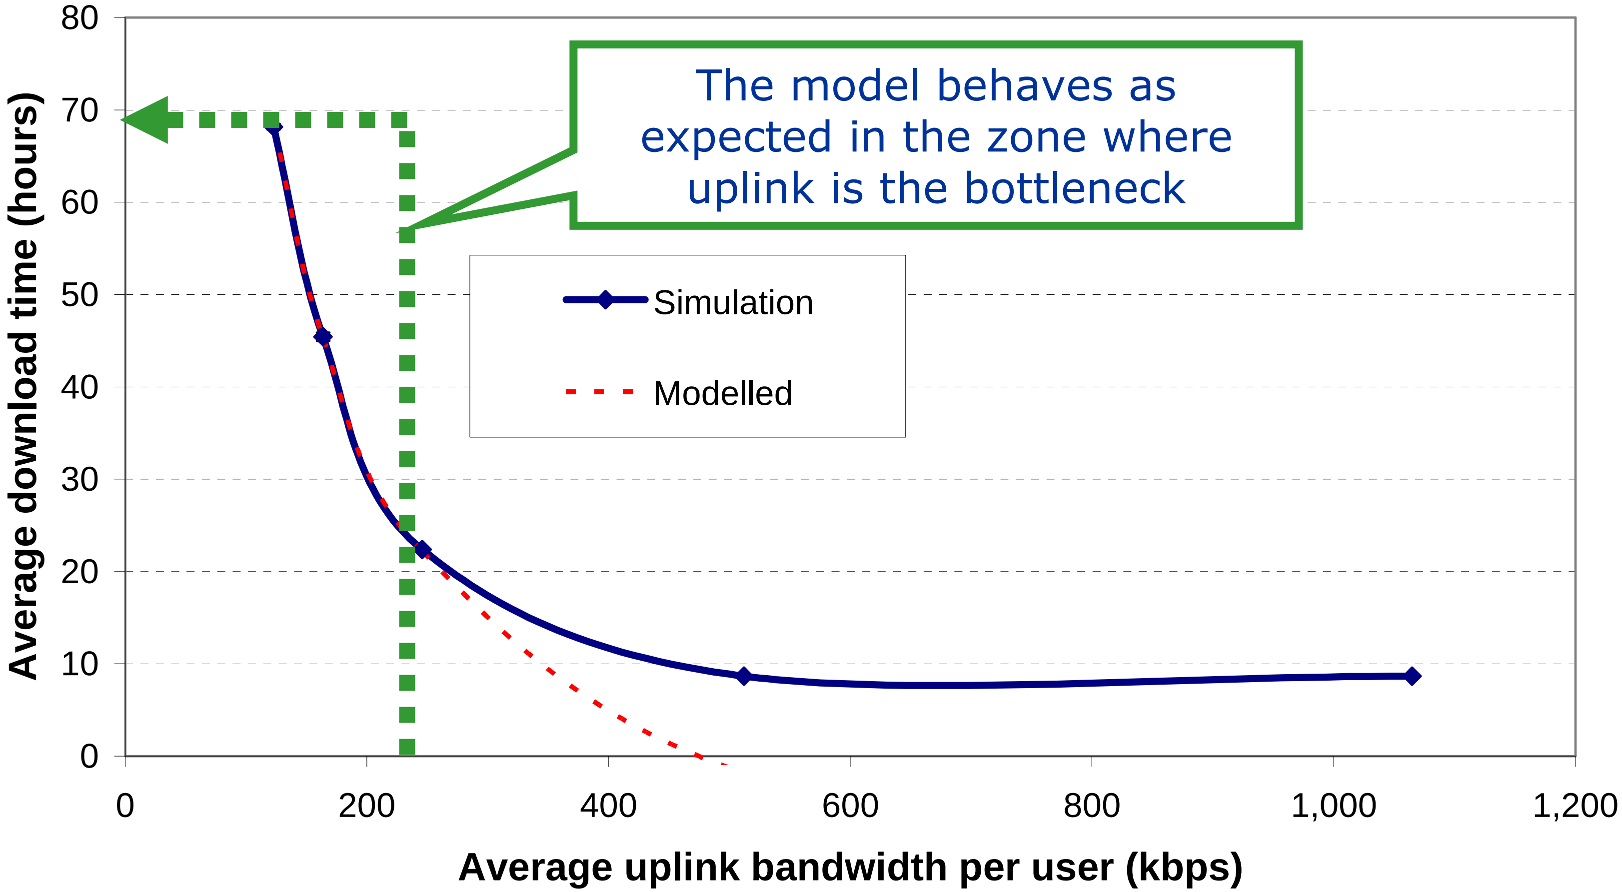

03 Model validation

Download time vs. uplink BW

$$T = \left(N_{pet} \cdot F\right) \cdot {}^{1}\!/_{BW_{up,user}} - S$$

$$\left(while\ \lambda \approx \mu\right)$$

- Is there an hyperbolic relationship between download time and uplink bandwidth while the latter is the bottleneck?

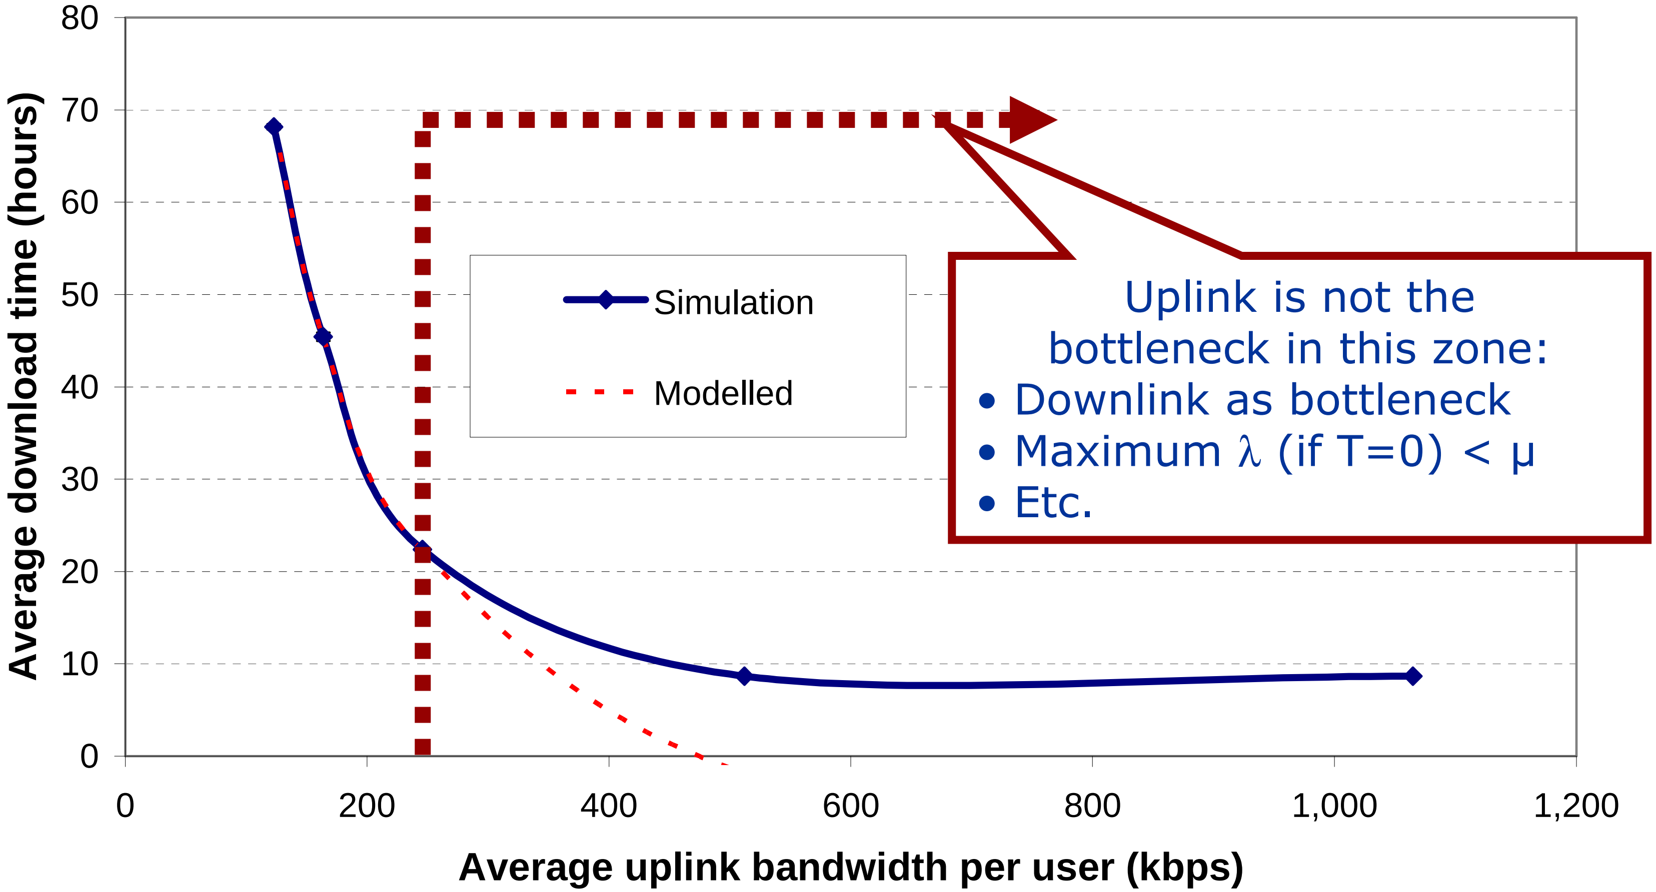

Index

04 Conclusions

- The PBoT model presents many **practical advantages**:
  - **Accurate results** for the most relevant parameters (download time, offered load, etc.)
    - Excellent predictions in eMule simulations
  - **Compact and simple expression**
    - Very few parameters, feasible to estimate
    - Just basic operations: +, -, ×, ÷
  - **Only rough approximations of averages** (silence time, content size, etc.) **are required**
    - No assumptions about statistical distributions at all!
  - **Valid to any efficient P2P system** (with very light requirements)

04 Conclusions

- Consequences and practical recommendations:
  - **The point of equilibrium does not depend on the number of users in the system**
    - Just a minimum to guarantee ρ≈1
  - Parameters in the equation can be grouped trivially, so that **the influence of the agents can be easily identified**

$$T = \frac{N_{pet} \cdot F}{BW_{up,user}} - S$$

➔ Users, P2P service provider

➔ Network operator

04 Conclusions

- Consequences and practical recommendations:
  - Users (and sometimes P2P service providers indirectly) are able to modify the point of equilibrium via *S*, N_{pet}, or *F*
    - However, **these strategies must be coordinated** or will become ineffective due to the emergence of a typical **Prisoner's Dilemma** situation
      - **Each user has incentives to increase its N_{pet}… but a general increase of N_{pet} is bad for all**
    - **The preferred enforcement point is the very P2P application** (coordinated effect)
      - Limitations for parallel downloads ➔ Indirect control of *Npet*
      - Specialization per contents (music, video, etc.) ➔ Indirect control of content size (*F*)

04 Conclusions

- Consequences and practical recommendations:
  - Network operators can change the equilibrium by changing average uplink BW…
    - … with an automatic increase in P2P traffic!

$$P2P\ traffic = \lambda \cdot F \approx \mathrm{N_{users}} \cdot \mathrm{BW_{up,user}}$$

- In the short-middle term (when *F*, *Npet*, and uplink BW are stable), it can be assumed that

$$\boxed{T + S = \mathrm{constant}}$$

  - There is a direct linear trade-off between average download time and average silence time.
  - **With the same average load (eq. same T+S), QoE can be very different!**

Index

05 Next steps

- New experiments with other P2P systems ➔ BitTorrent
  - Simulator in development
- Characterize real P2P traffic using the main parameters of the PBoT model
- Model generalization
  - Scenarios where users have symmetrical or heterogeneous access connections
  - Including effects of special peers (e.g. seeders)
    - There is already an analytical expression
    - To be validated by simulation