A Modeling Framework to Understand the Tussle between ISPs and Peer-to-Peer File Sharing Users

Michele Garetto – Univ. of Torino
Daniel Figueiredo – Fed. Univ. of Rio de Janeiro
Rossano Gaeta – Univ. of Torino
Matteo Sereno – Univ. of Torino

IFIP Performance 2007

UNIVERSITÀ DEGLI STUDI DI TORINO
ALMA UNIVERSITAS TAURINENSIS

Laboratory for modeling, analisys and development of networks and computing systems

An Internet Tale

- Once upon a time...

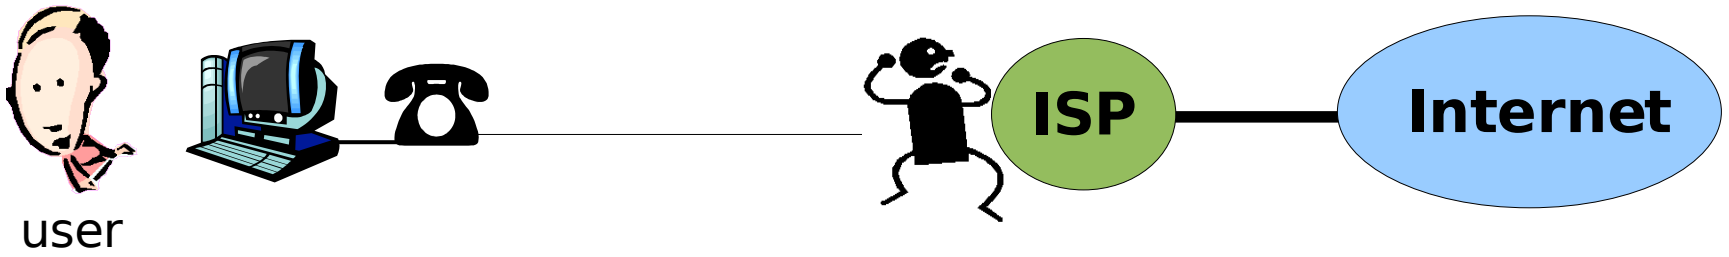

  - user unhappy (“world wide wait”)
  - ISP unhappy (little revenue)

- Then came broadband access...

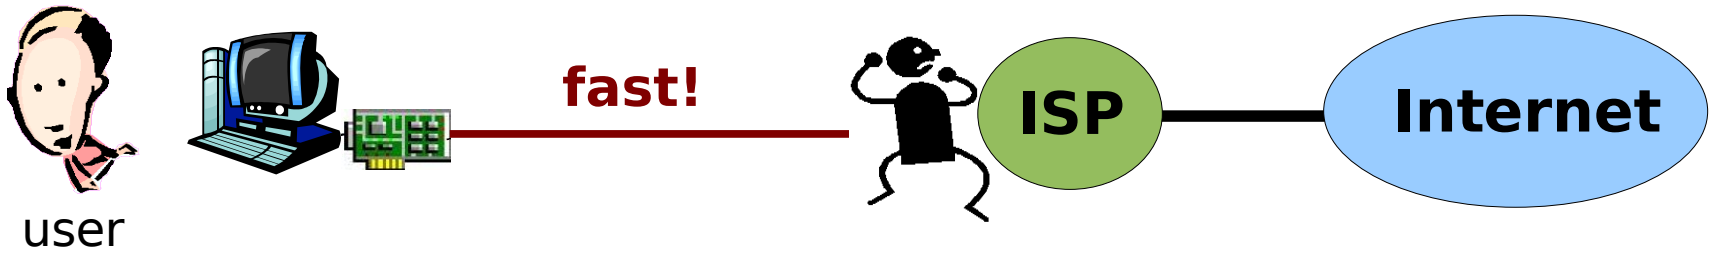

- And they lived happily ever after...

The Villain Arrives

- P2P file sharing application (Kazza, Bittorrent, Emule, etc)

- **users love it!**
- good and free content, overnight downloads

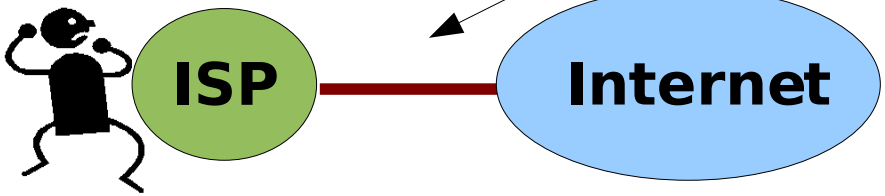

- **ISPs hate it!**
- users using their link
- Internet link utilization gone wild
- degrades all subscribers
- more bandwidth costs money!

Taking Care of The Villain

- **Mafia style!**

- seriously threaten application developers!
- doesn't seem to work (Napster story)

Is it Really a Villain?

user

- Users love it!
- Driving force for broadband adoption
- Increased revenue for ISPs

Some Other Options

- User **unfriendly** ideas
  - increase subscription cost
  - volume based pricing
  - block / shape P2P traffic
- User **friendly** ideas
  - acquire more bandwidth
  - network caching
  - application-layer redirection

What should the ISP do?

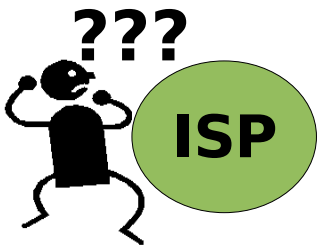

The Real Thing (Data)

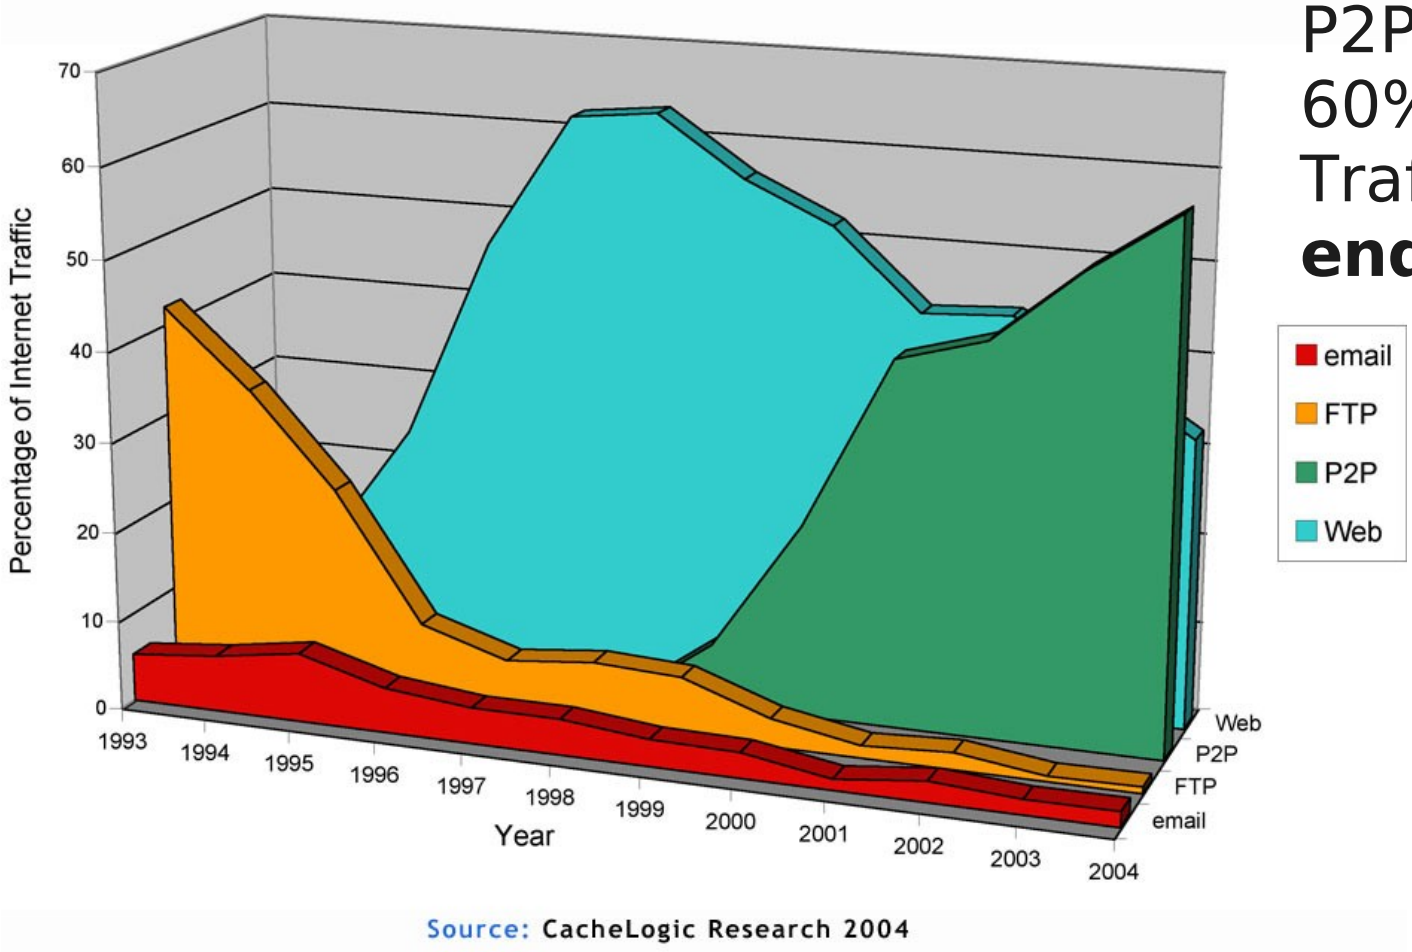

P2P represented 60% of Internet Traffic at the **end of 2004!**

Our contribution

- Modeling framework to analyze interactions between P2P file sharing users (their traffic) and ISP
  - economic + performance models
- Basic insights about system dynamics
- Used to evaluate different strategies to manage P2P traffic

Meet the Players

user

- generates queries
- quality of service expectations
- what's hot, what's not

- **goal**: to make money!
- sets subscription price
- controls bandwidth
- influences P2P app. behavior

- P2P application
- locates object

Network

- network architecture
- protocols

System Setting

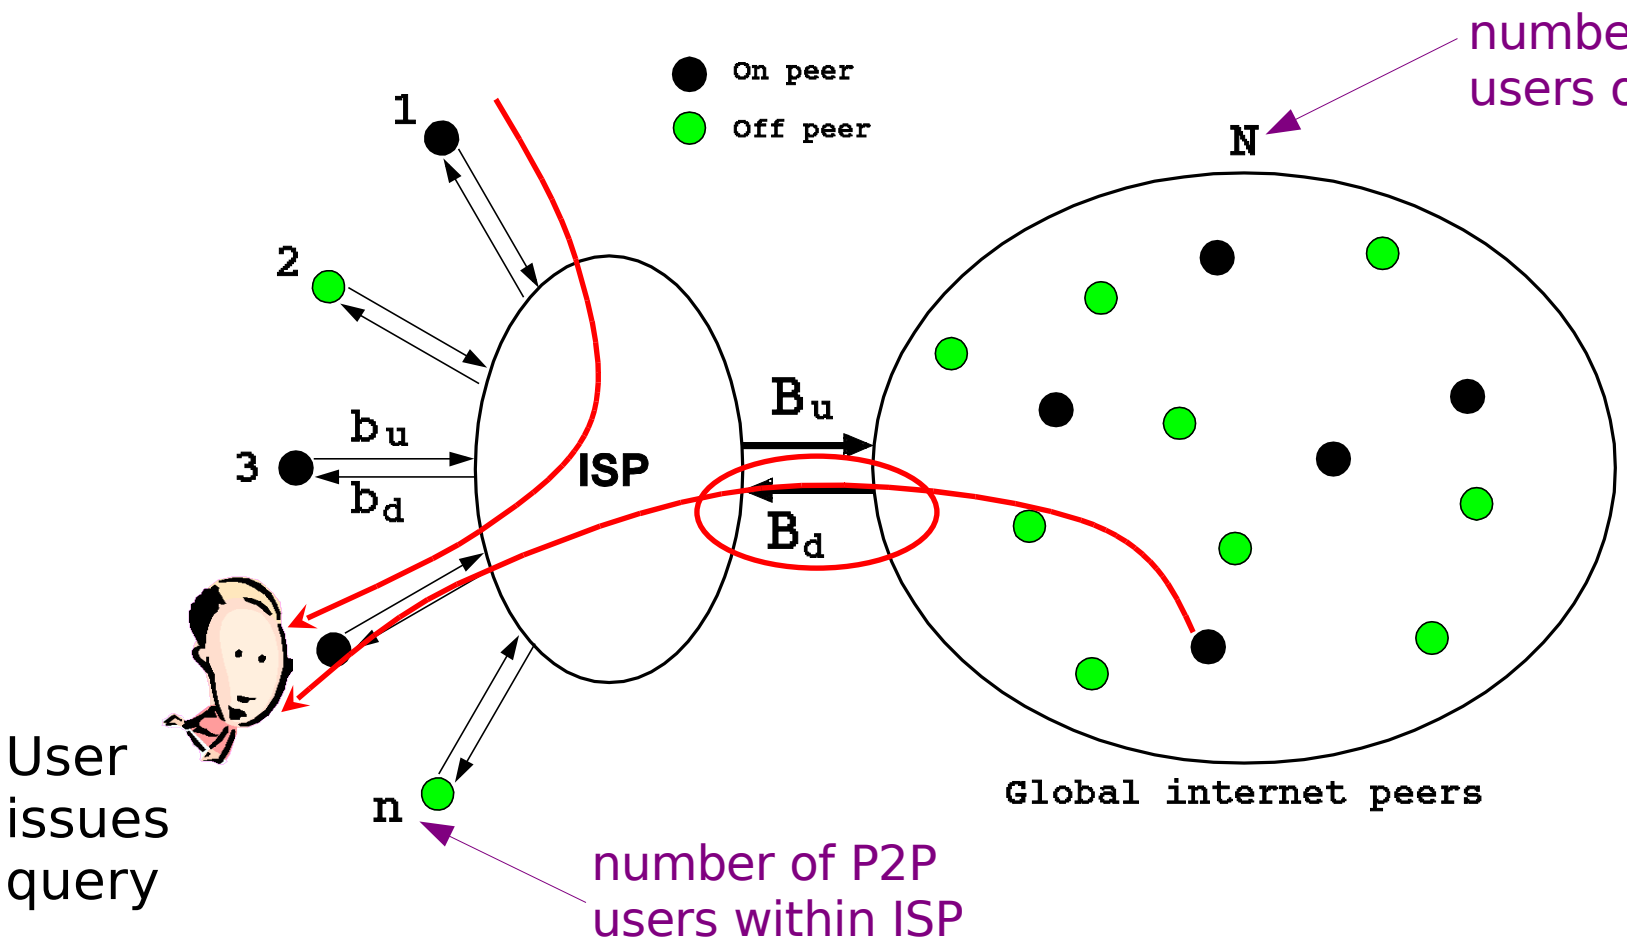

- B_d → constrained resource for ISP
- Outside download consumes B_d

Simple System Model

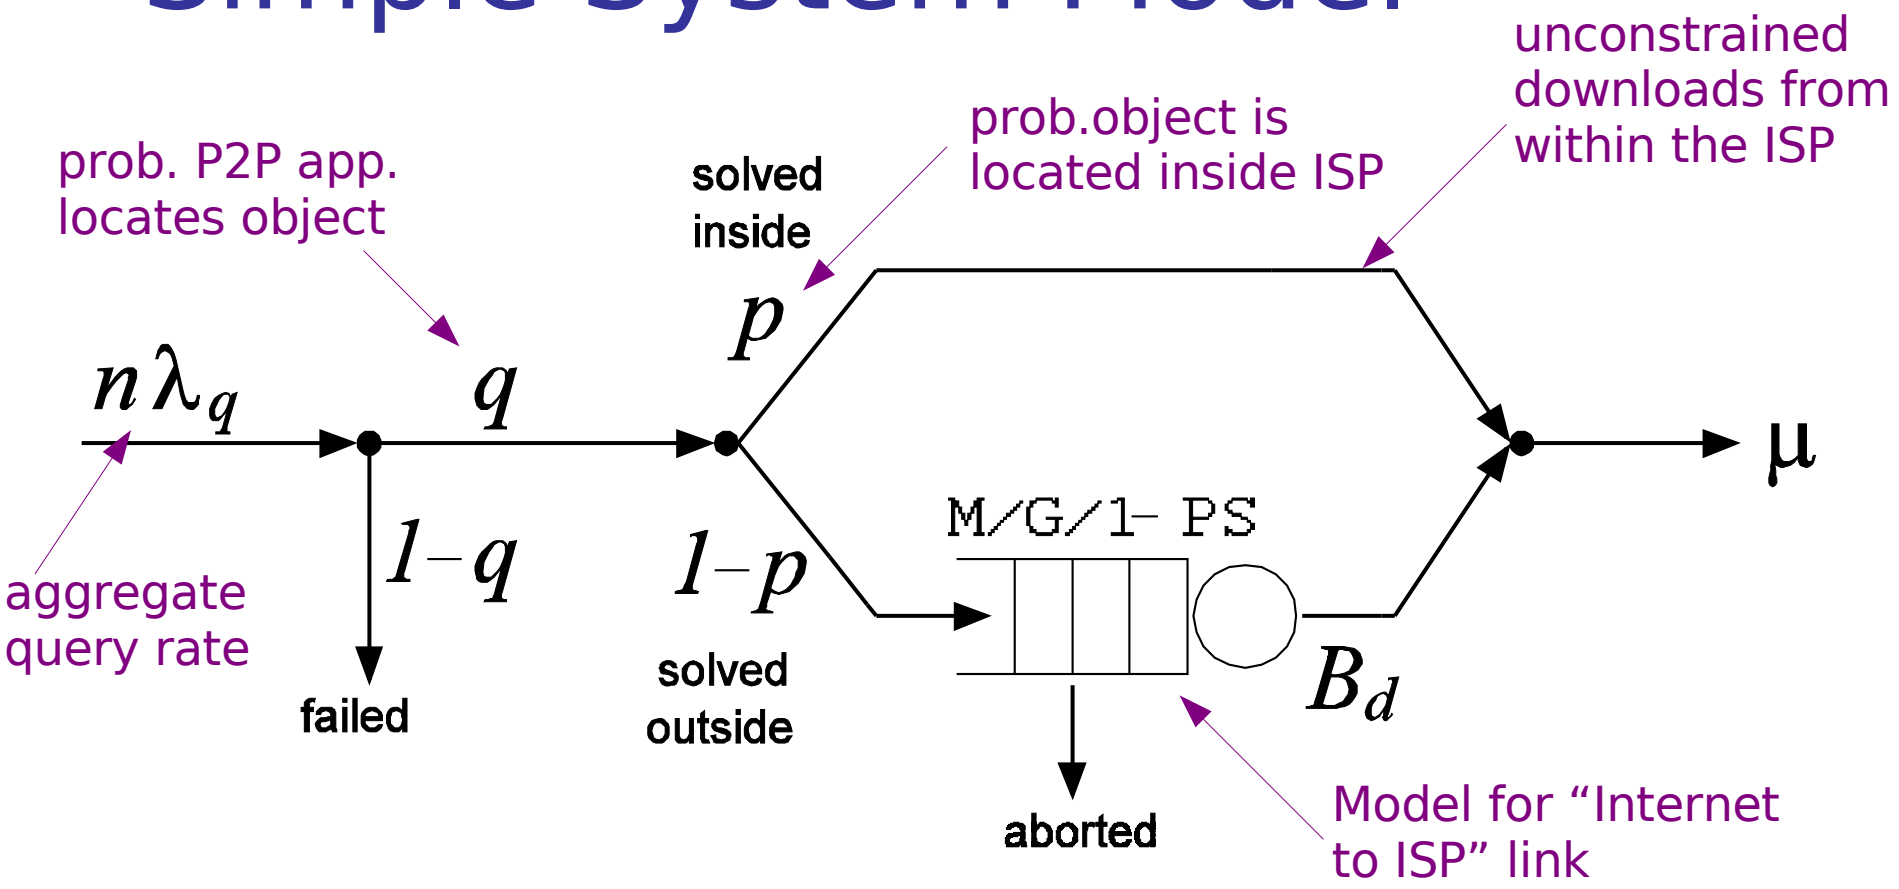

object retrieval probability: $\sigma = \dfrac{\mu}{n\lambda_q}$

User Utility Function

- Satisfaction model for user i

user

$$U_i(\sigma, c) = \log(\alpha_i \sigma + 1) - c$$

shape parameter

probability of successful object retrieval

subscription cost

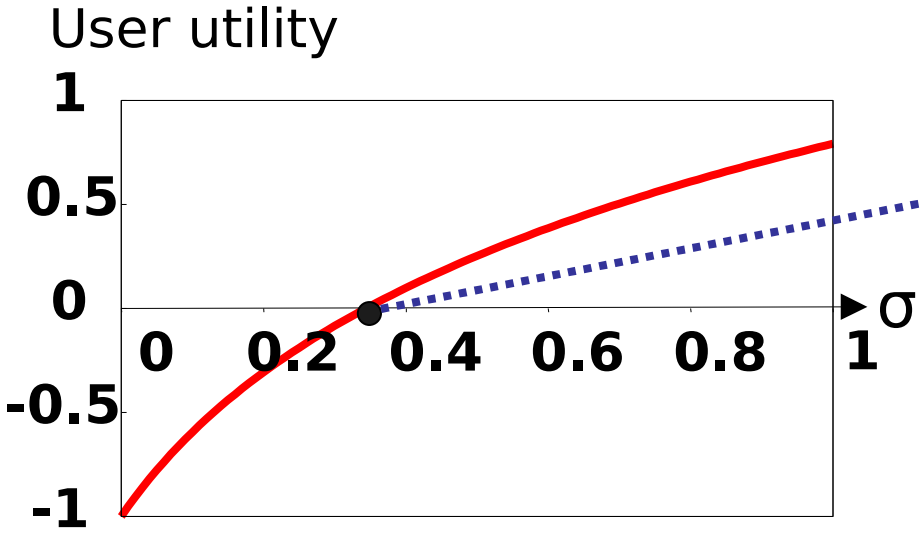

$\sigma \min_i$

Minimum service level for user i

ISP Utility Function

- Profit for ISP (revenue - costs)

$$U_{\text{ISP}}(B_d, c) = \sum_{i=1}^{n} c\mathbb{1}_{(U_i(\sigma,c)\geq 0)} - (\beta_2 B_d + \beta_1)$$

revenue from subscribers' fee

cost per unit of external bandwidth

fixed charge

- The ISP starts service only if

$$U_{\text{ISP}}(B_d, c) > 0$$

Modeling Traffic Locality

- Probability there exist at least one internal replica of object replicated r times in the system

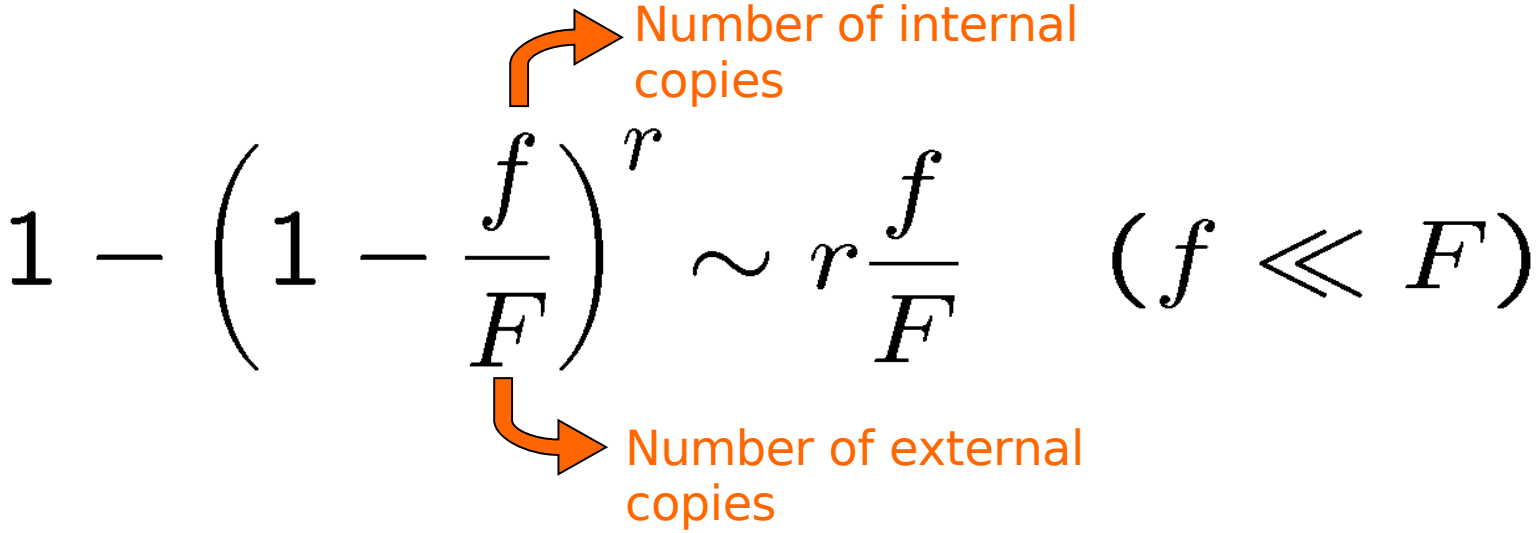

- Probability to download from internal replica

$$p \sim \gamma r f / F$$

locality parameter

Analytical Results

- How much bandwidth should the ISP buy to minimally satisfy the users?

$$B_{min} = max[0, n\lambda_q(\sigma_{min} - q\gamma r n/N)]$$

  - identical users and n >> N

- Non-linear behavior (on *n*)
  - more users, more locality, less BW needed
  - can be zero if *n* large enough
- May not yield profit
  - too few users, too costly to satisfy them
- Dependent on multiple parameters

Impact of Object Replication (r)

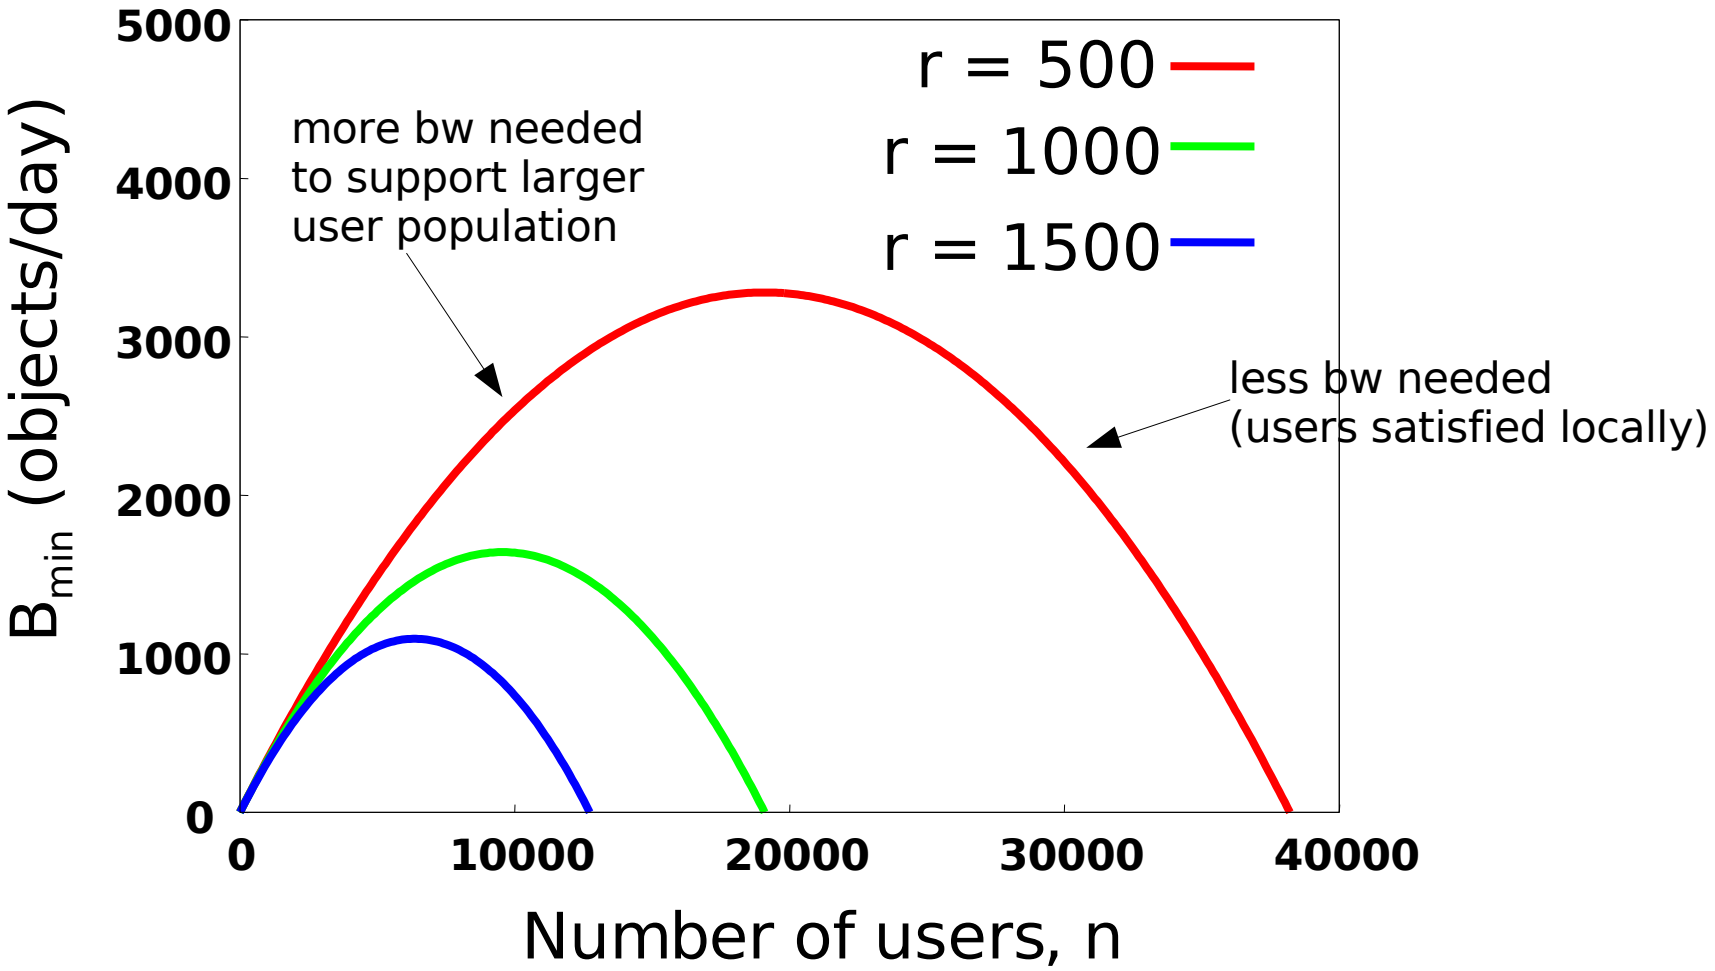

- more replicas, better locality, lower B_{min}

Impact of Subscription Cost (c)

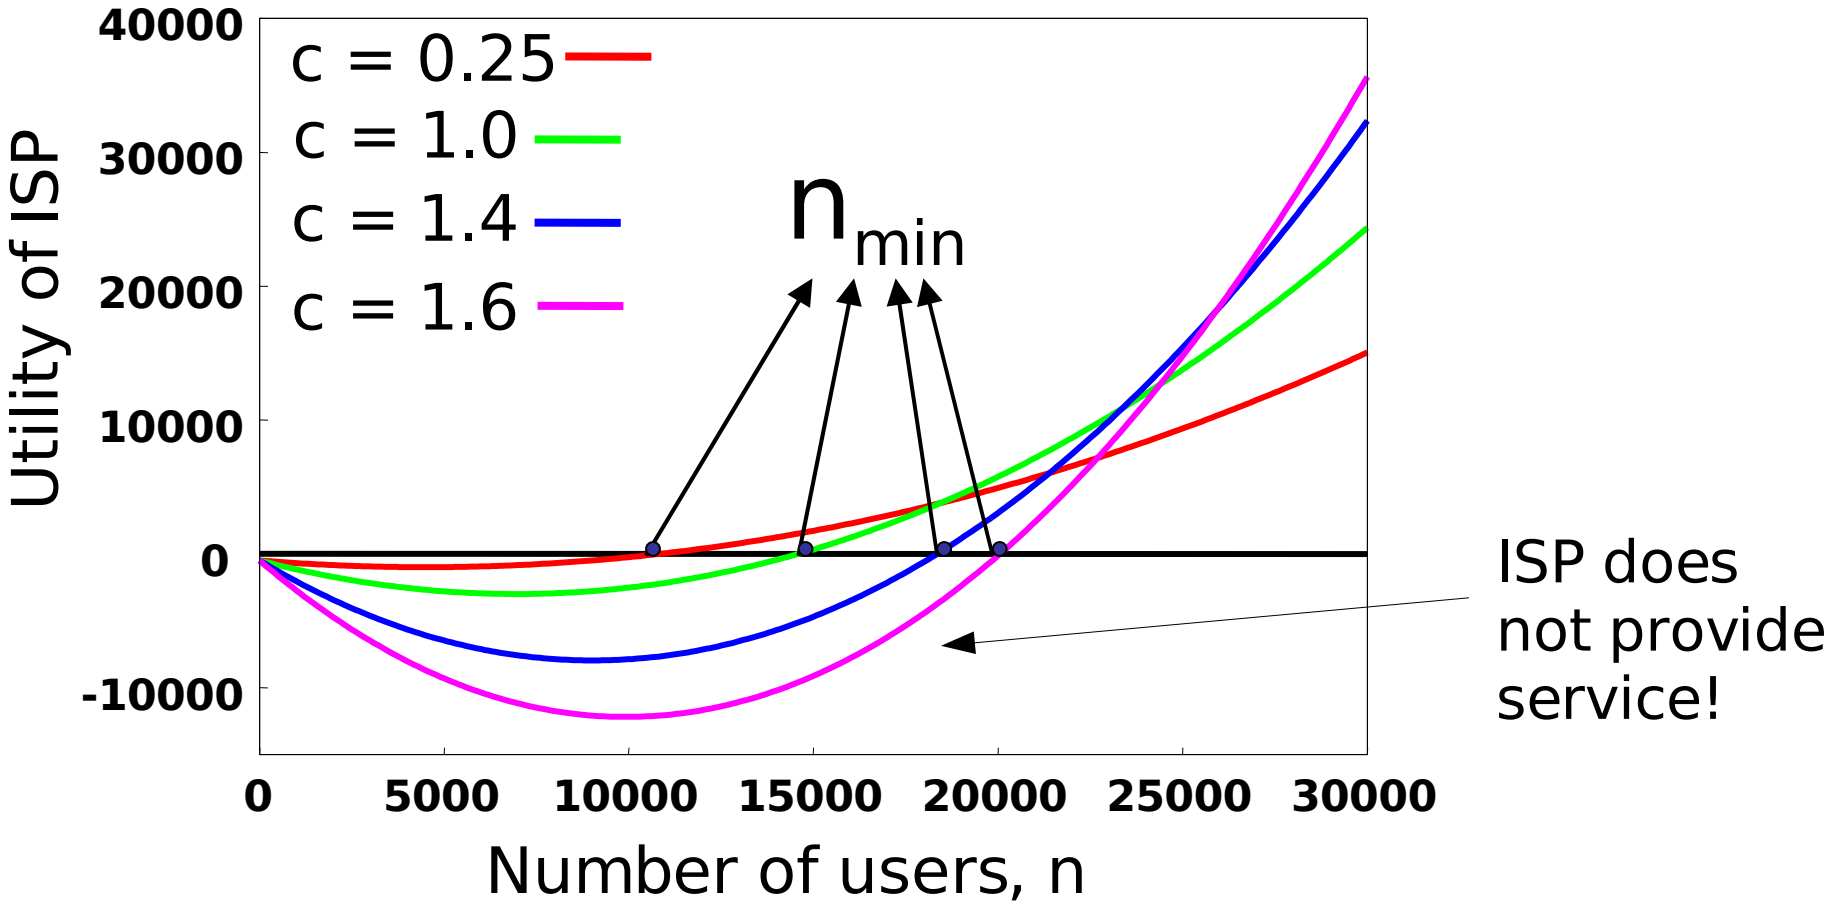

- critical mass of users, n_{min}
- lower cost, more profit earlier, less profit later

Critical Mass of Users, n_{min}

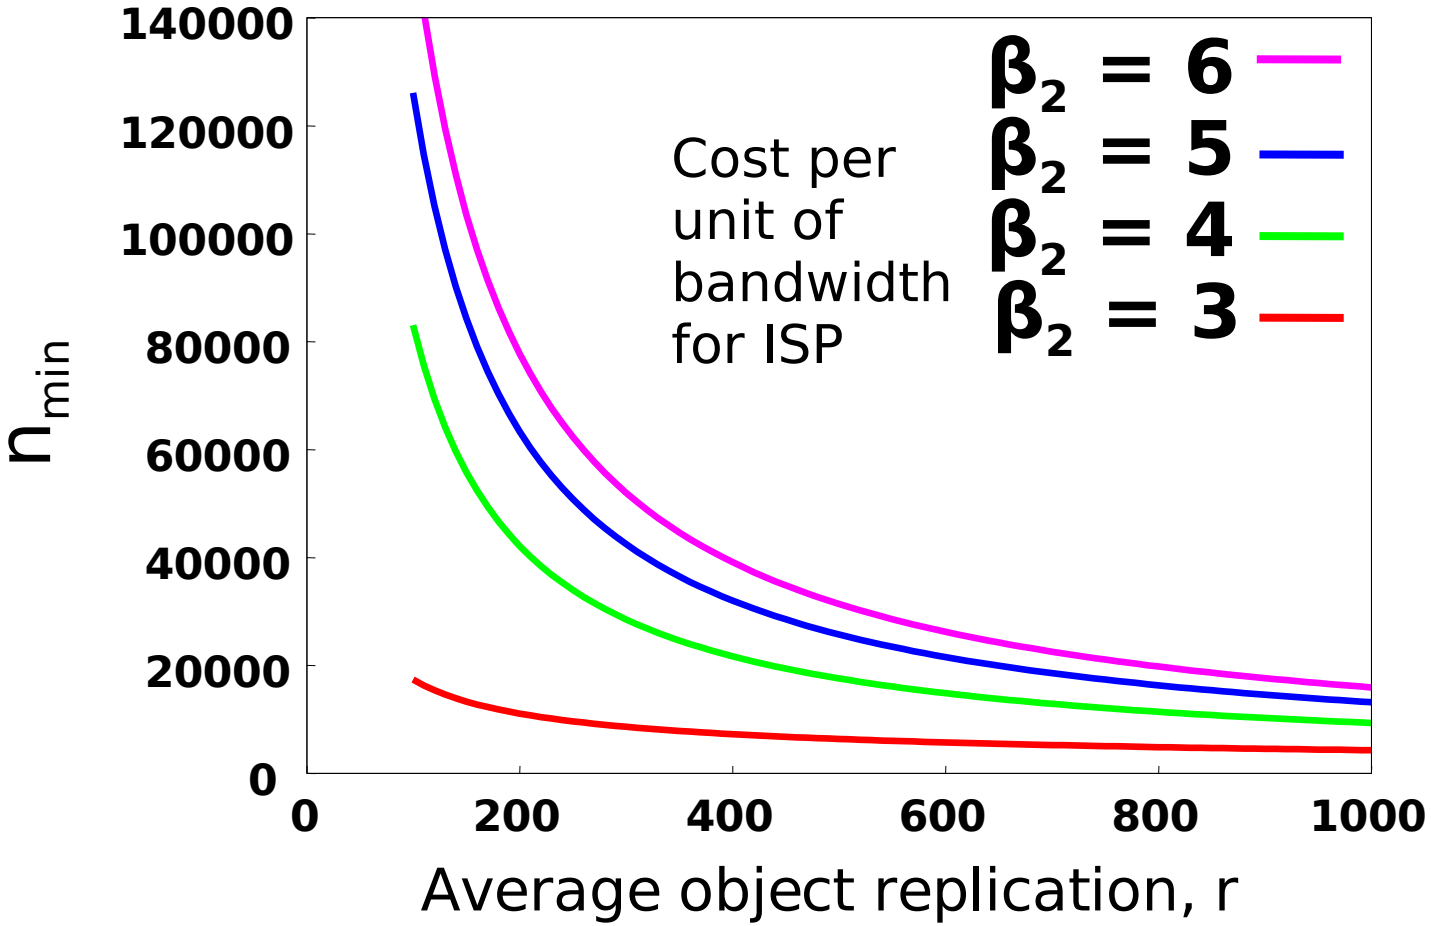

- higher bw cost for ISP, higher critical mass
- large influence of number of replicas

Model Refinements

- Simple model
  - users' access bandwidth are unconstrained
  - object replication is a parameter
  - all objects are identical (no popularity)
  - users availability identical
- Refined model relax these assumptions
  - propose object popularity and replication model

Object Popularity and Replication Model

- Temporal evolution of object popularity
- Objects' popularities evolve differently
- Objects continuously introduced and removed by users

Number of replicas of an object at time t?

- Analytical technique based on Poisson shot noise process

Example

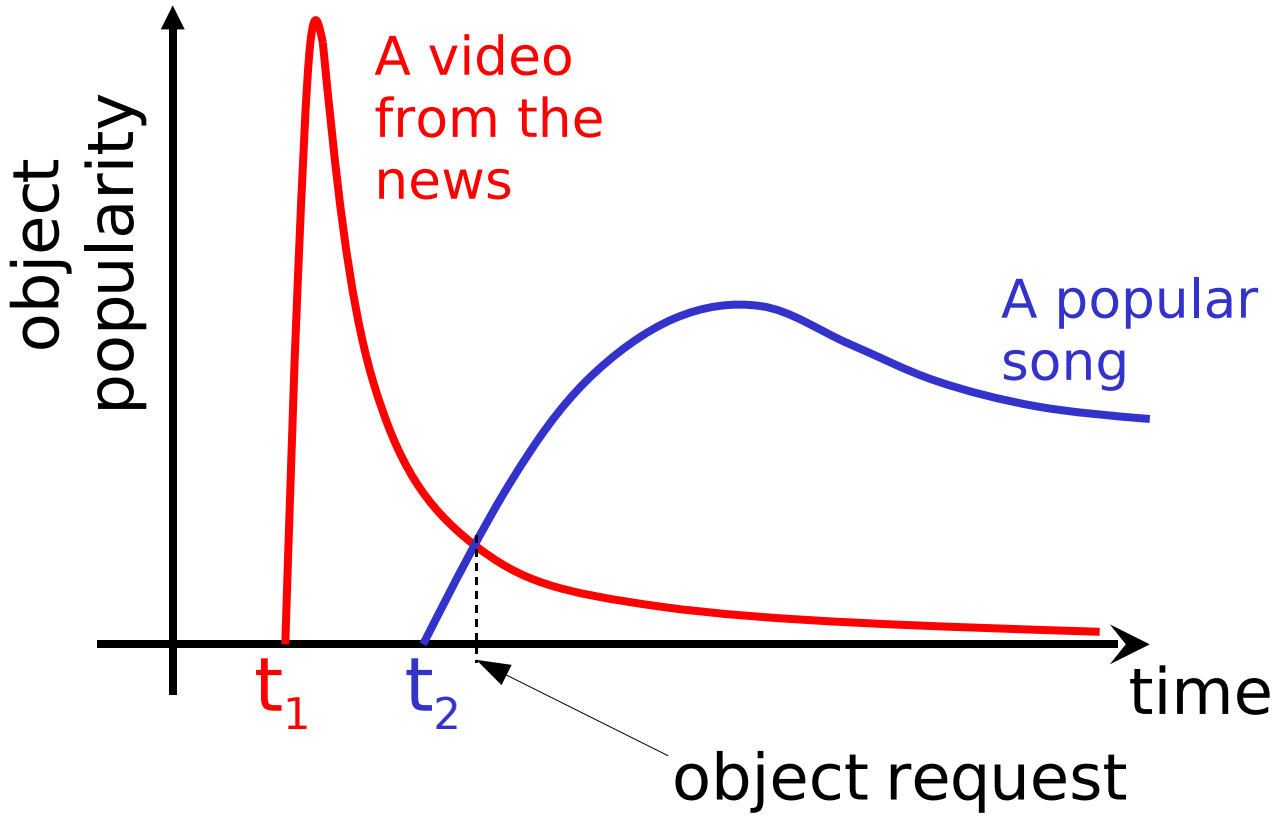

- at request time, both have same popularity, but news has more replicas

Limited Bandwidth Refinements

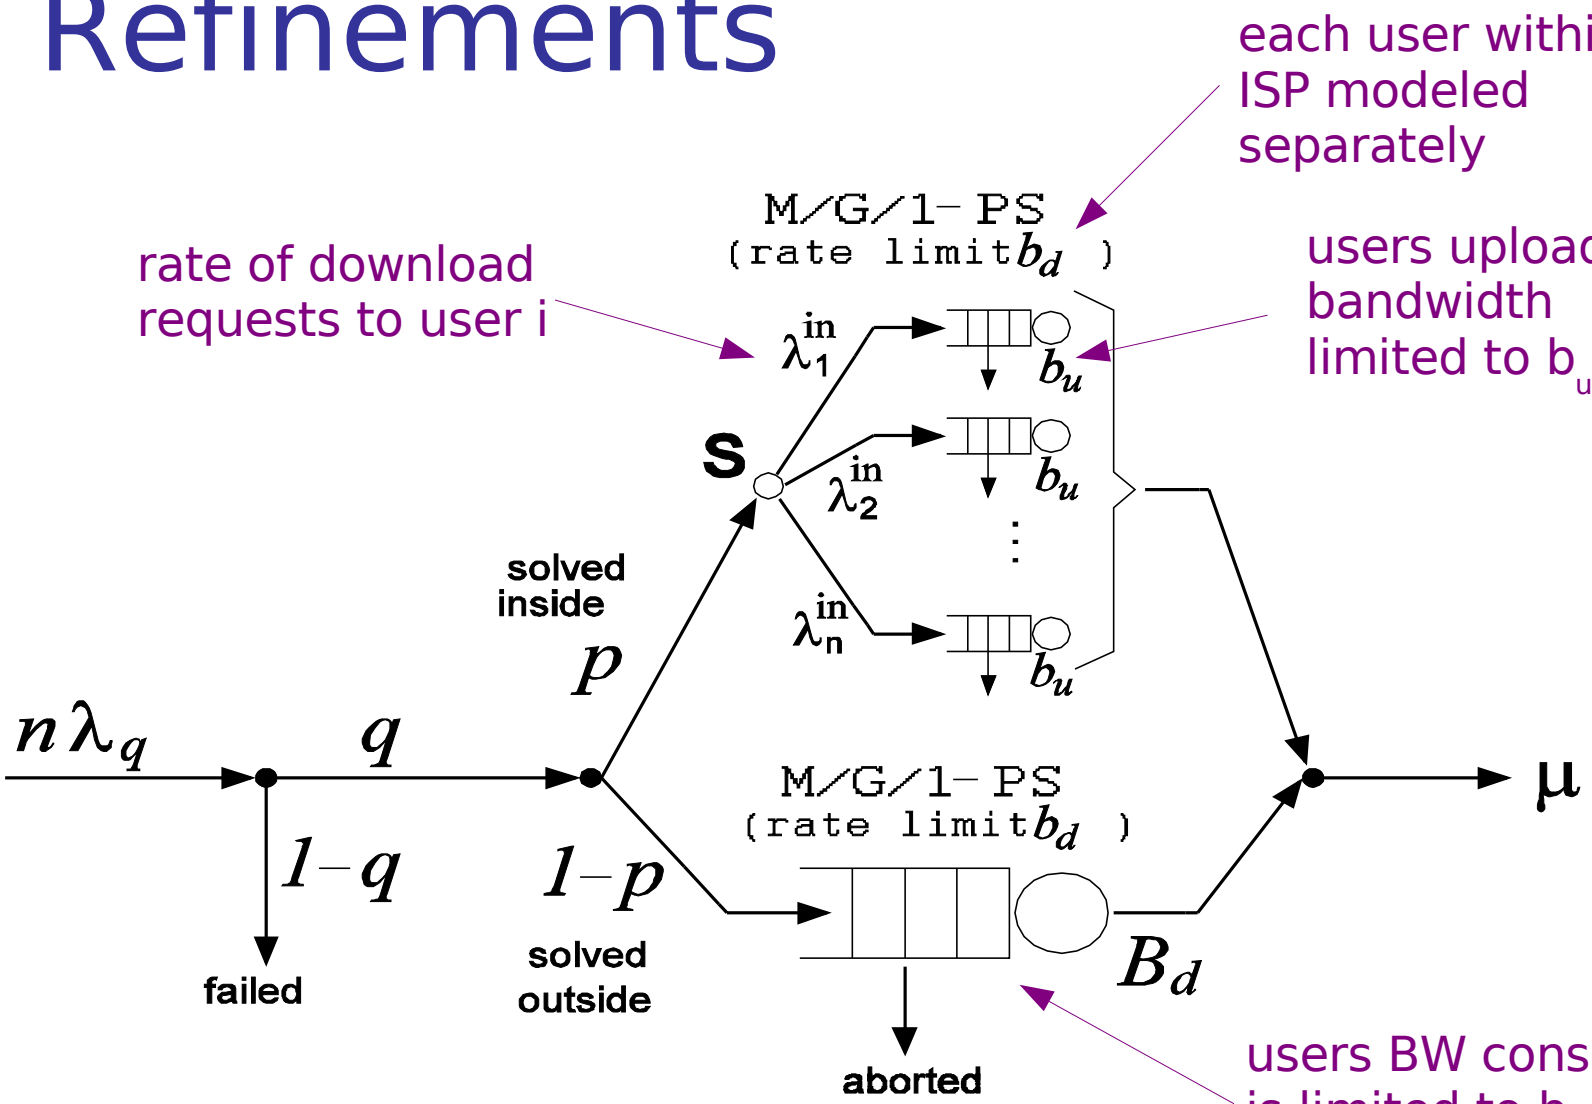

Results from Refined Model

- Degenerate to simple model
  - when parameters set appropriately

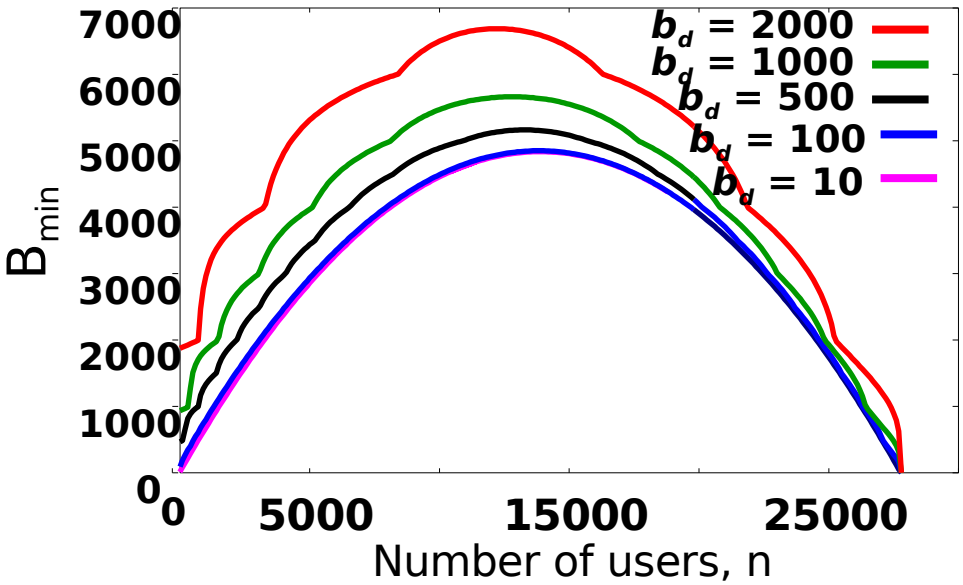

- Other interesting insights
  - influence of limited upload bandwidth
  - upload/download bandwidth asymmetry
  - object popularity and replication
  - influence of user impatience

Impact of asymmetric access bandwidths

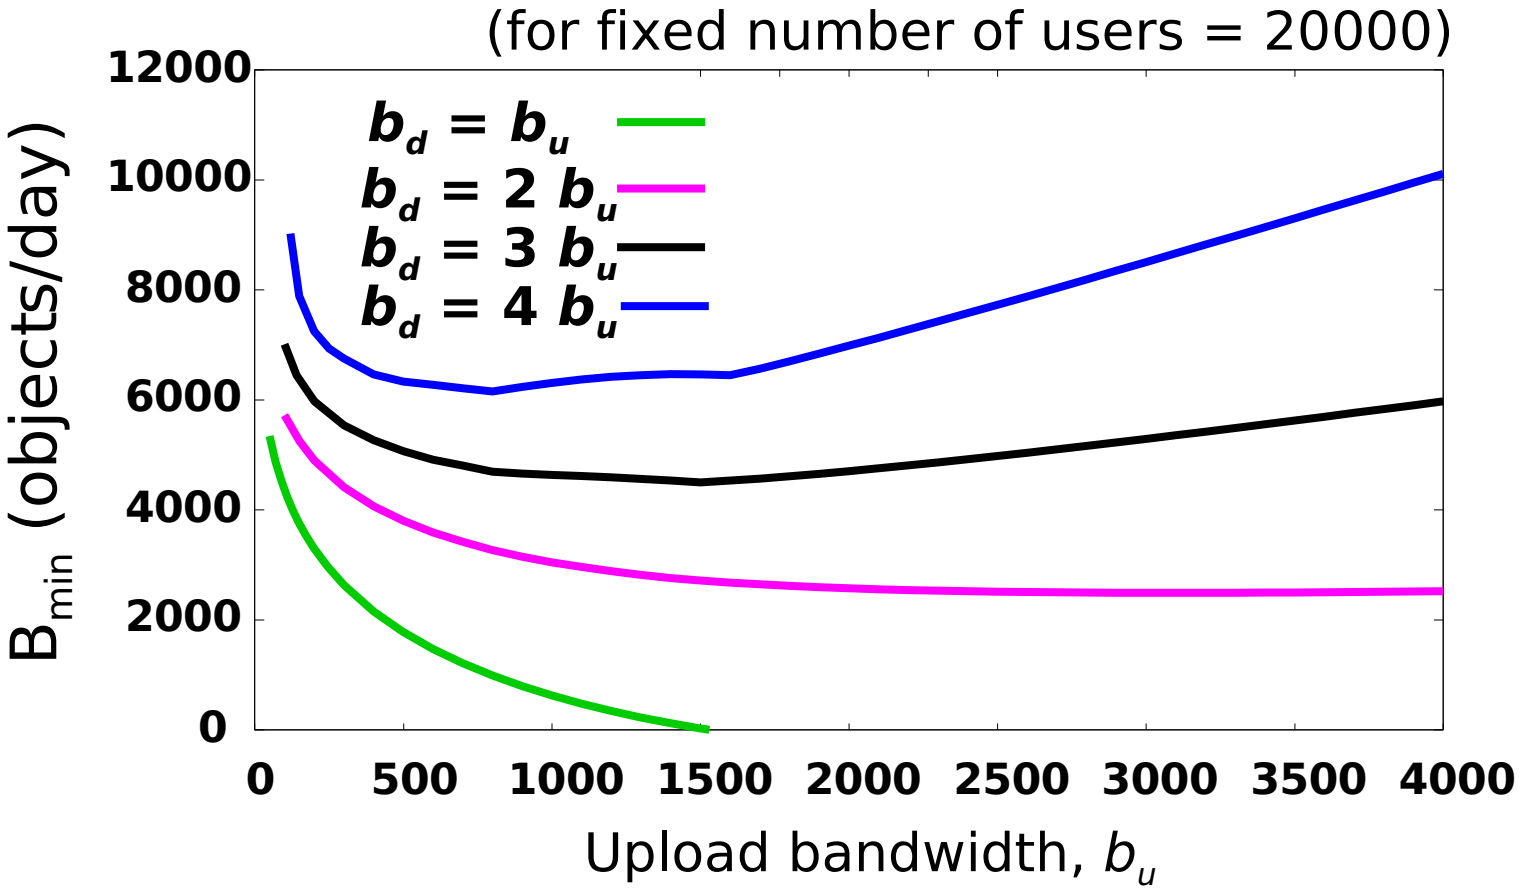

- cost for ISP *increases* as ratio increases
- better if upload BW is *greater* than download

Conclusions

- Development of simple analytical model
  - economics + performance
  - interaction between P2P users (their traffic) and ISP
  - insights into strategy for ISP to manage its traffic
- Model for object popularity and replication
  - of independent interest
- Future work
  - Multiple ISPs competing with each other

THE END

- Thank you!
- Questions? Comments?

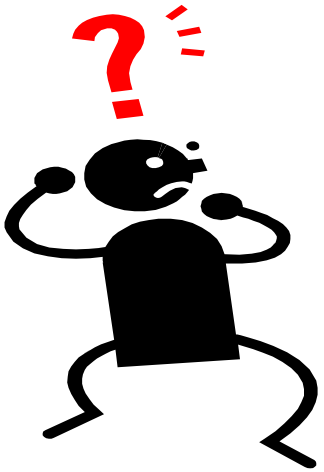

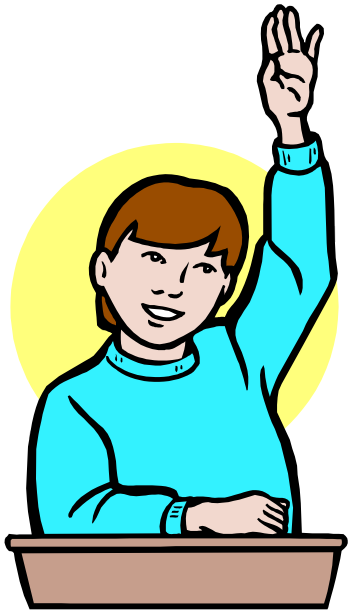

Impact on Service Providers

P2P is driving consumer broadband uptake
…and broadband is driving P2P uptake

- P2P is the dominant protocol
  - In excess of 92% of P2P traffic crosses t
  - P2P protocols will aggressively consum capacity
  - Due to P2P's symmetrical nature on av capacity is consumed by P2P
- P2P affects QoS levels for ALL subscr
- Service Providers can not afford to bl
- ISPs must intelligently manage P2P - doesn't work

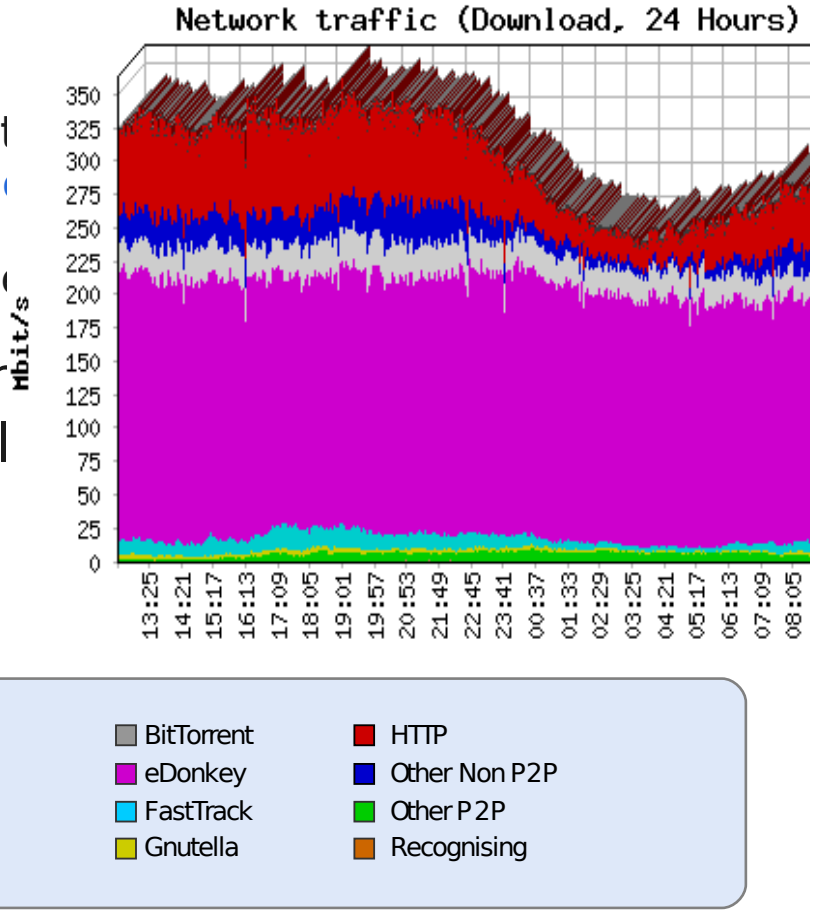

The ISP perspective vs P2P:
threat or opportunity ?

- P2P traffic: friend or foe ?
  - friend: driving force for adoption of broadband access by the users
  - foe: overwhelming amount of traffic
- What is the best strategy to manage P2P traffic in my network ?
  - Try to kill it ?
  - Do nothing ?
  - Educate it ? How ?

Strategies to manage P2P traffic

- Acquire more bandwidth
- Block P2P traffic
- Traffic shaping (e.g., priority to non-P2P)
- Pricing schemes based on user traffic volumes / bandwidth caps
- Network caching / customized P2P application within ISP
- Application-layer redirection of P2P traffic

Results

- Assumption: n identical user($\sigma_{\min} = 1/3$)
- N = 50 millions
- $\lambda_q = 1$ request rate by user (object/day)
- $\lambda_o = 10^{-5}$ introduction of new contents by user (object/day)

Minimum required external bandwidth:

$$B_{\min} = \max\left[0, n\lambda_q\left(\sigma_{\min} - \sigma_1\gamma r\frac{n}{N}\frac{\lambda_o + \lambda_q\sigma_{\min}}{\lambda_o + \lambda_q\hat{\sigma}}\right)\right]$$

The impact of efficacy in exploiting traffic locality (γ)

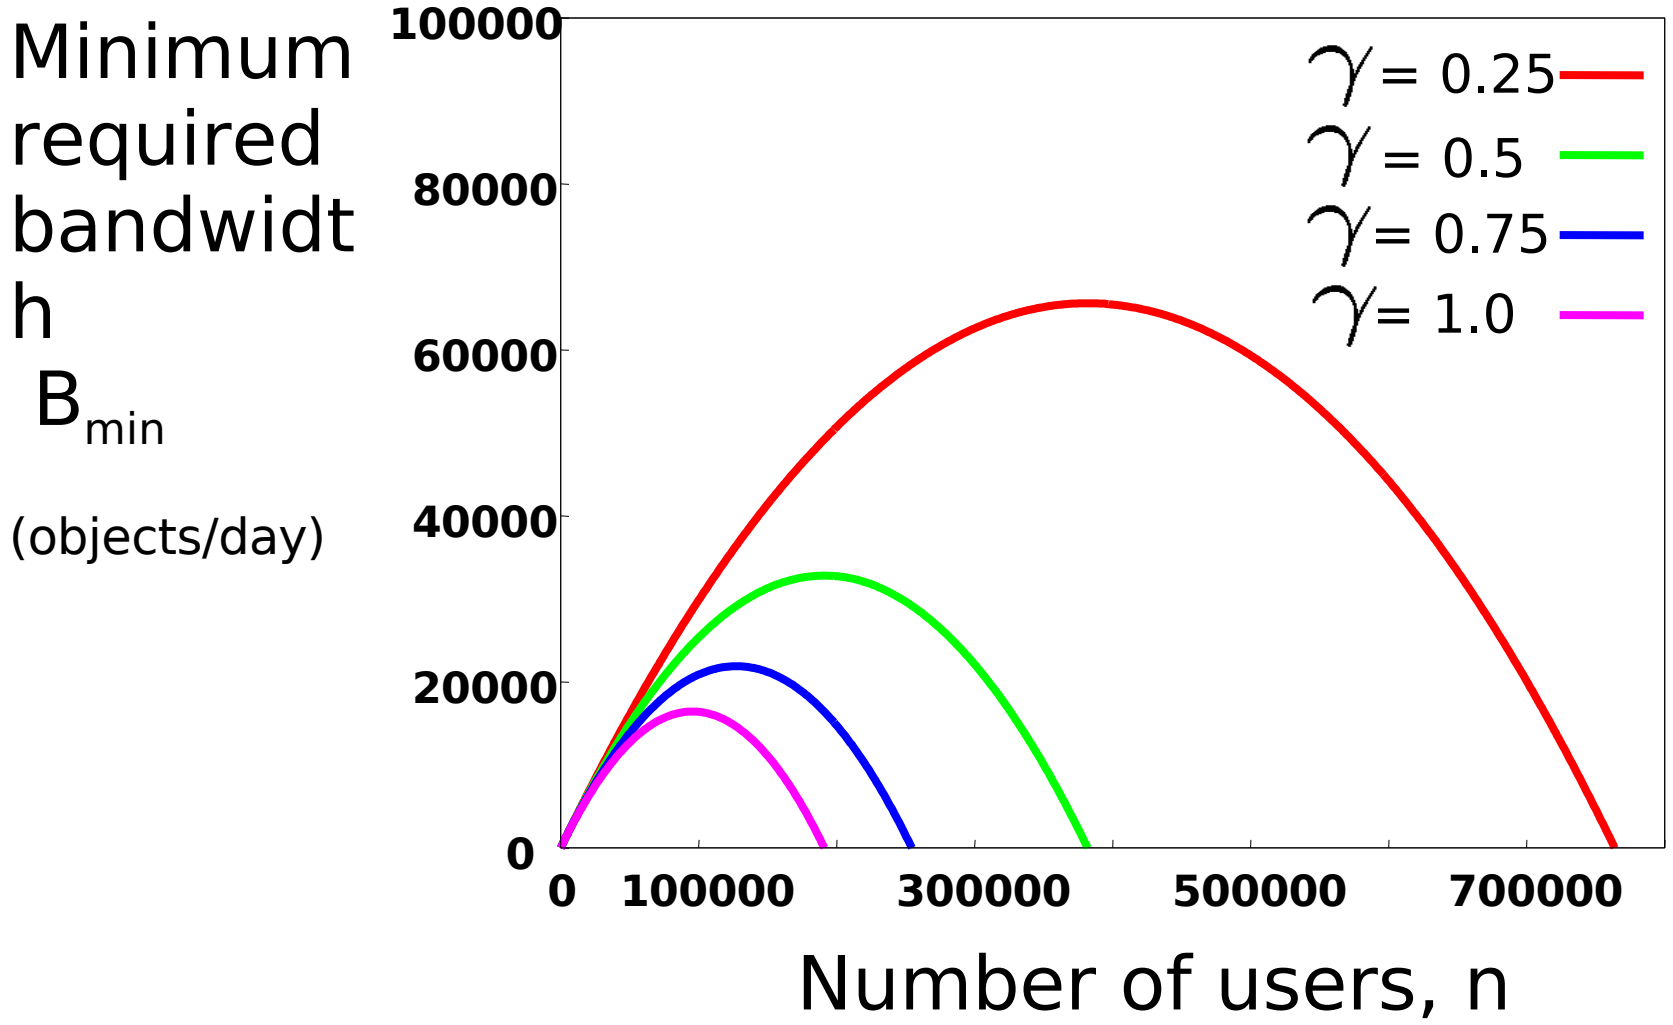

Model refinements

- Impact of finite bandwidth of the users
  - In case of constant traffic load, the cost for the ISP increases if it provides more download bandwidth to the users (!)
  - The system performance is strongly affected by the upload bandwidth of the users, that should be larger than or equal to the download bandwidth (contrary to common practices, e.g., ADSL lines !)

The impact of b_d