Playing with the Bandwidth Conservation Law

Orange Labs

Farid Benbadis, Nidhi Hegde, Fabien Mathieu, Diego Perino, Research & Development
P2P 2008, Aachen

Bandwidth conservation law in a P2P system?

- Common assumption: download speed = average upload
- And in real life?
  - Joost offers 700kbps download for 120kbps upload (6/1)
  - BitTorrent: 300/1 for some personal experiments…

Sources	Clients	Reçoit à	Envoie à	Restant
106 (712)	8 (171)	131,5 ko/s	789 o/s	1h 24m
375 (4007)	124 (2384)	922,5 ko/s	2,9 ko/s	5m 12s

- But the bandwidth conservation is a ``physical'' law

) maybe applying it correctly can give some insight

Roadmap

- Model
  - Bandwidth, just bandwidth, only bandwidth
  - Leechers, seeders, servers
- The closed model
  - Some nice formulas
  - Application: is Joost scalable?
- The open model
  - Other formulas (nice too!)
  - Application: maximum download speed in BitTorrent

What matters in a P2P system?

- Access … width
- ~~Core~~
- ~~Later~~
- ~~Jitte~~
- ~~Availability~~
- ~~Allocation schemes~~
- ~~Routing~~
- …

BANDWIDTH

The Leechers, the Seeders and the Servers

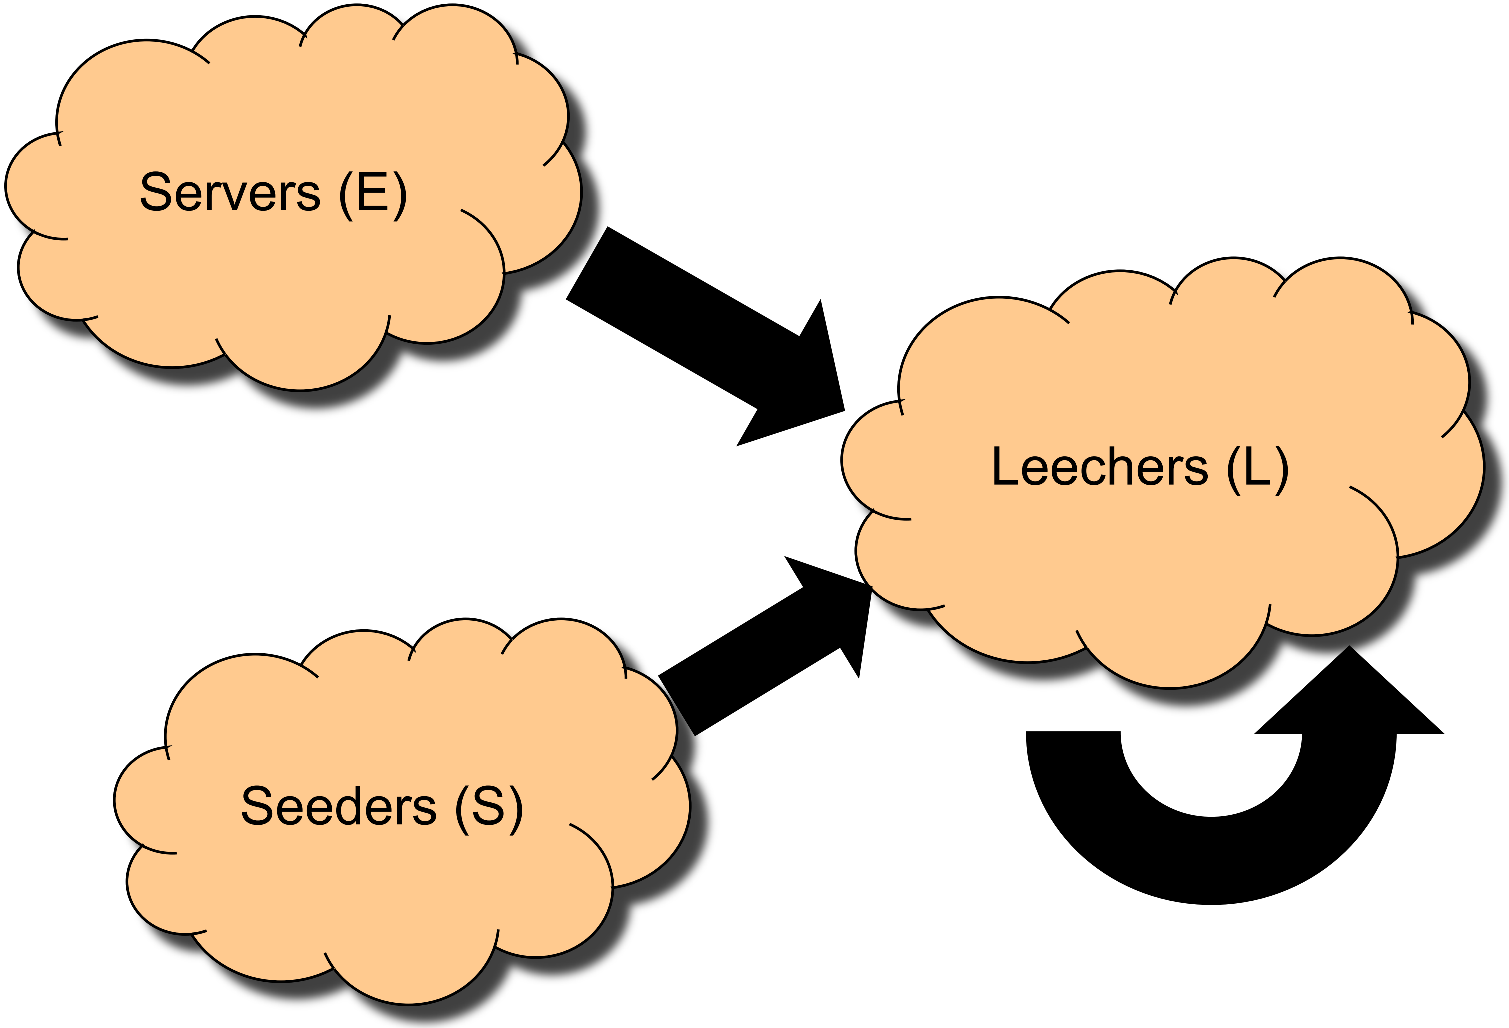

Closed Systems,
looking for scalability

A first bandwidth conservation law

- We assume populations and upload capacity are fixed

$$N_L, \bar{u_L}, N_S, \bar{u_S}, U_E$$

- If bandwidth is fairly distributed (same d for all), we have

$$d = \min\left(d_{\max}, \bar{u_L} + \beta \bar{u_S} + \frac{U_E}{N_L}\right),$$

with $\beta = \frac{N_S}{N_L}$

Streaming point of view

- For a given r let $\alpha_X = \frac{\bar{u_X}}{r}$ and $N_E = \frac{U_E}{r}$
- If $d \geq r$, the streamrate can be achieved, otherwise…

- Case 1: $\alpha_L + \beta\alpha_S \geq 1$
  - System is scalable (it works independently o N_L),
  - Servers are unnecessary (except for bootstrap/security)

- Cas 2 : $\alpha_L + \beta\alpha_S < 1$
  - N_L is bounded by $\frac{N_E}{1-(\alpha_L+\beta\alpha_S)}$

  ! the server capacity is leveraged by $\frac{1}{1-(\alpha_L+\beta\alpha_S)}$

Back to Joost

- The upload/download ratio gives $\alpha = \frac{1}{6}$
- If $\beta = 0$ (no seeders), P2P increases the server capacity by 20%
- If $\beta = 1$, this becomes 50% (most probable behavior)
- If $\beta = 2$, 100%
- …
- Scalability is achieved only if $\beta \geq 5$ all the time!
- Open question: is Joost more CDN or P2P?

Open Systems
Stationary process

Leechers/Seeders Life Cycle

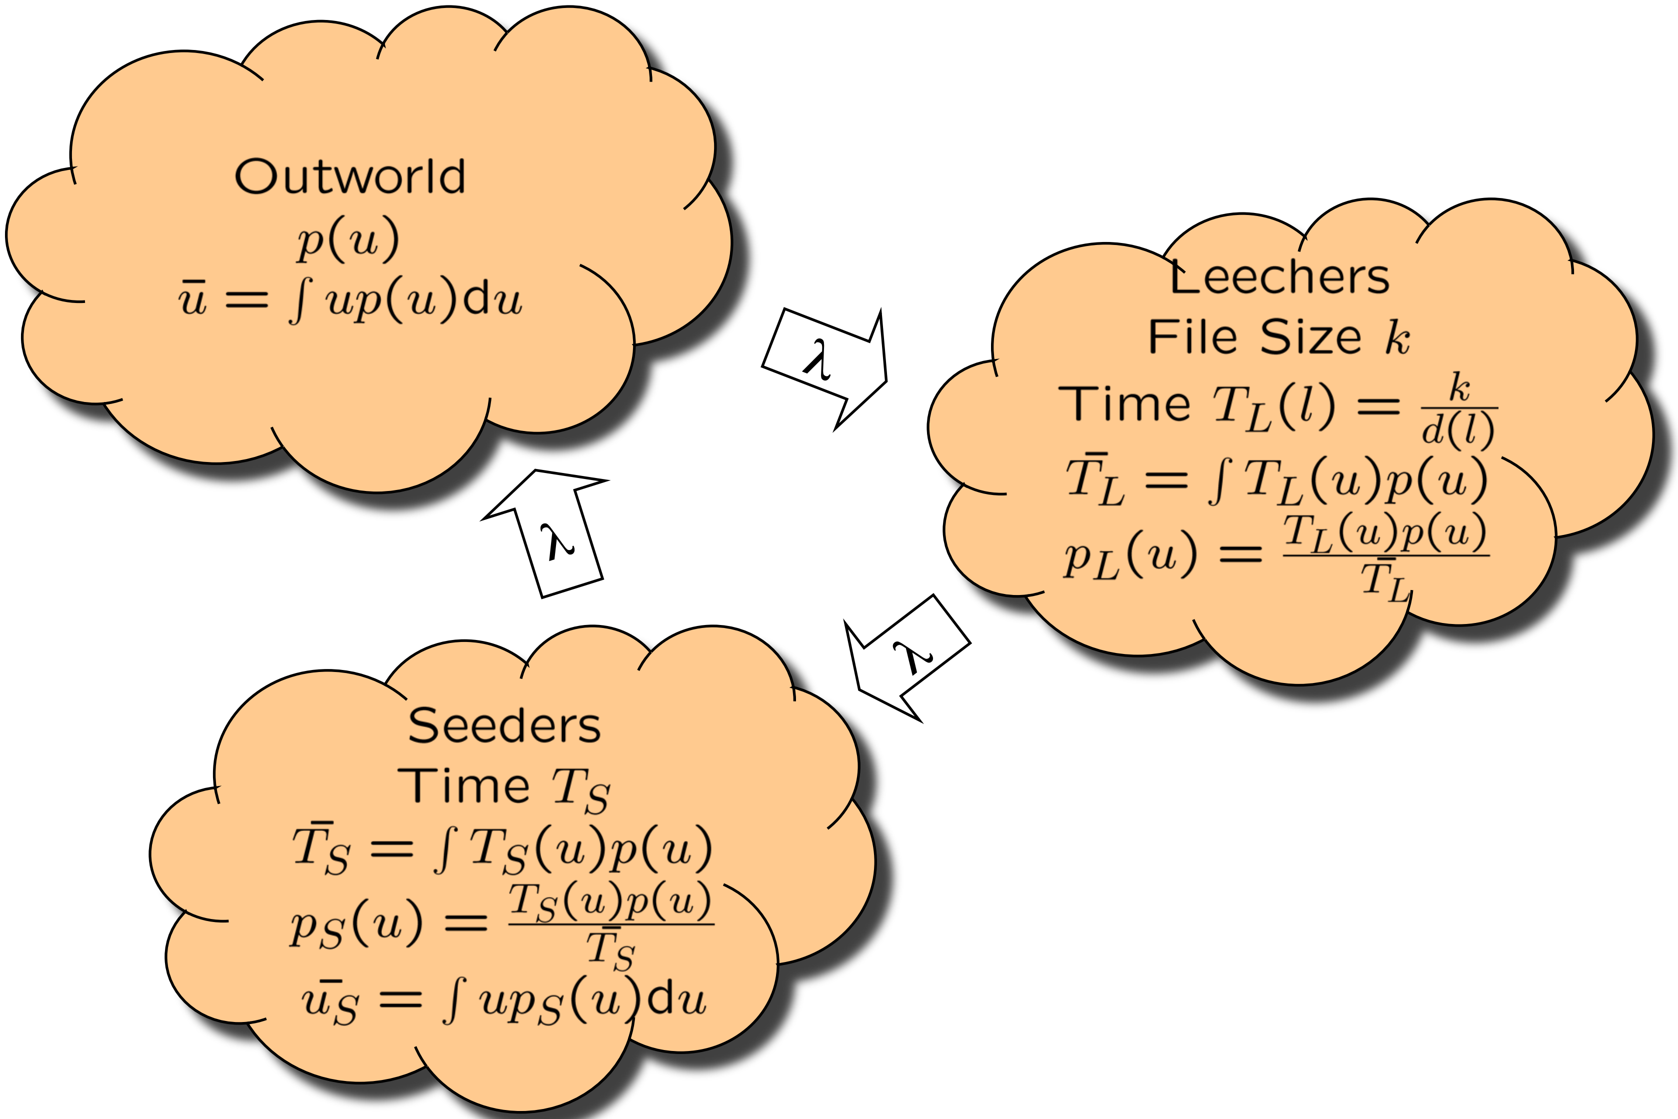

Overprovisioning condition

- We use $d = \min\left(d_{\text{max}}, \bar{u_L} + \beta\bar{u_S} + \frac{U_E}{N_L}\right)$, and introduce the intensity.

- This gives a sufficient condition for achieving d_{max} in any optimal allocation scheme:

$$\bar{T_S}\bar{u_S} + \frac{U_E}{\lambda} > k(1 - \frac{\bar{u}}{d_{\text{max}}}).$$

- In particular:
  - if $\bar{T_S}\bar{u_S} + \frac{U_E}{\lambda} \geq k$, any target rate can be achieved,
  - if $\bar{T_S}\bar{u_S} \geq k$, any target rate can be achieved, for any λ, U_E.

Back to BitTorrent

- No server and scalability mean that we want the $\bar{T_S}\bar{u_S} \geq k$ condition;
- $\longrightarrow$ seeding time must be greater than $\frac{k}{u_S}$;
- In other word, if you seed the time needed to get the file when there is no seeder...
- ...then you download in no time!

Back to BitTorrent

- Interpretation: a system where you pay your bandwidth debt after your download can work.
- Question: how can you force the users to seed?
- First answer: ask politely (sometimes it works)
- Second answer: private tracker and share-ratio policy.

Share-ratio enforcement

- Some systems force a minimal download/upload share ratio $\tau<1$ from their participants
- Users that do not meet the share-ratio have ``trouble''
- This gives an equation for T_S :

$$T_S(u) = \max(0, \frac{\tau k}{u} - T_L(u))$$

- No easy way to solve this in the general case, but one can
  - Consider specific cases,
  - Find singularities,
  - Approximate the solution,
  - Use a numeric solver

) performance studies can be made

Example with heterogeneous bandwidth distribution and Tit-for-Tat

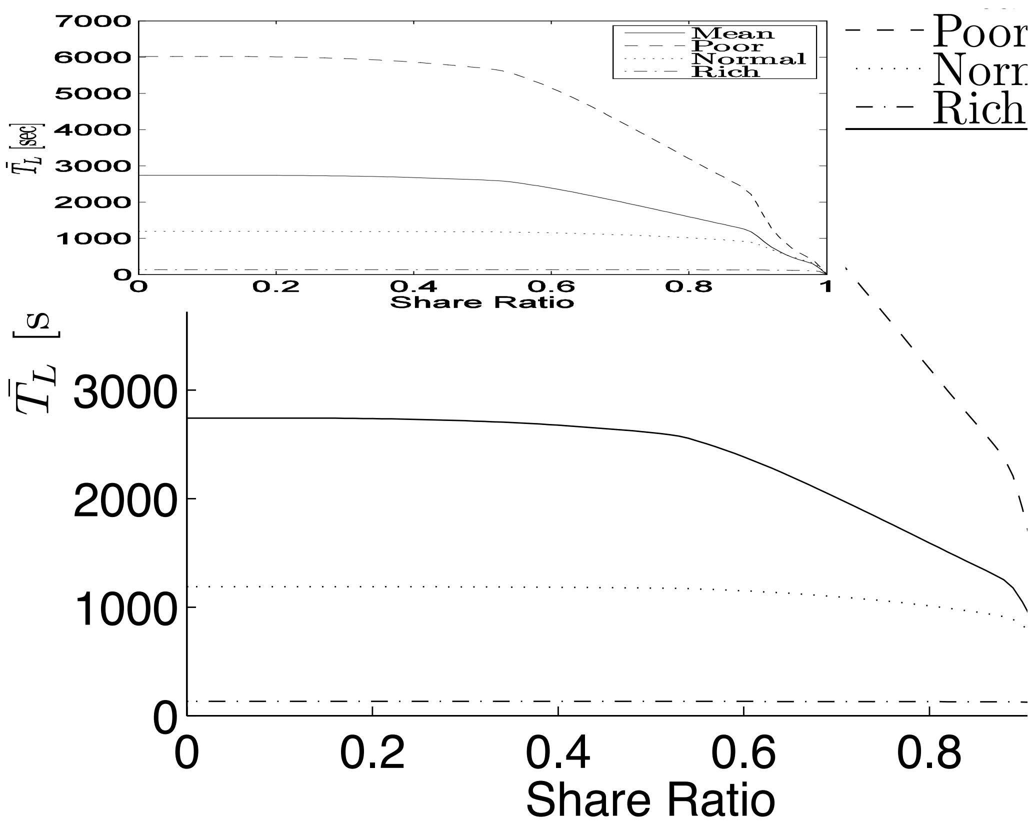

Back to the ``personal experiment''

Sources	Clients	Reçoit à	Envoie à	Restant
106 (712)	8 (171)	131,5 ko/s	789 o/s	1h 24m
375 (4007)	124 (2384)	922,5 ko/s	2,9 ko/s	5m 12s

- It was a private tracker!

- It also was very popular content…

…but that's another story!

Conclusion

- Bandwidth conservation is simple,
- It applies to all bandwidth-consuming application,
- It gives answers on what we can expect from a P2P (or hybrid) system

- For more applications, play with the article
  - Jetlaging with the bandwidth conservation law,
  - Free-riding with the bandwidth conservation law
  - …

Thank you very many!

Bonus 1:
worldwide convolution

Rush hour

- In practice, leeching (and therefore β) is a function of time
- The dimensioning must be performed in the rush hour

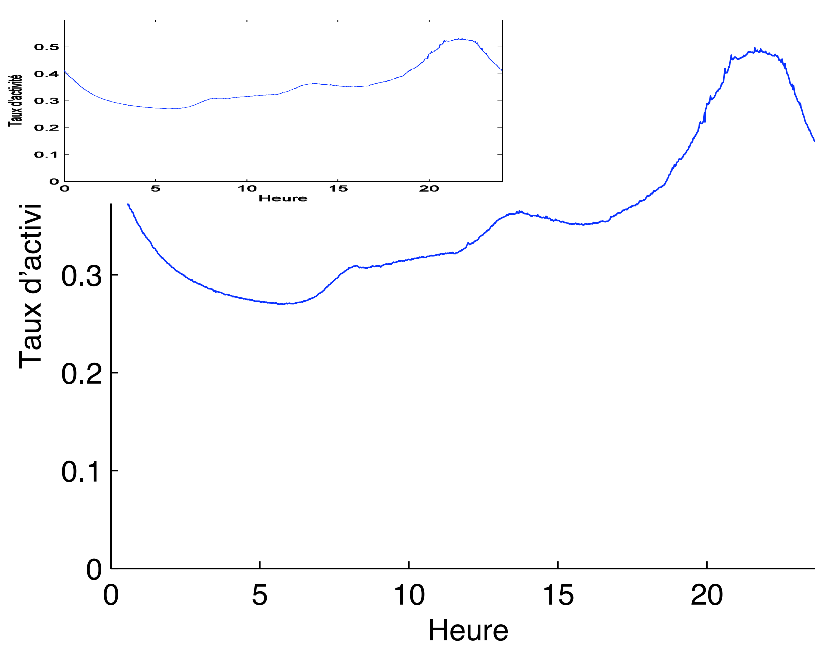

Rush hour – geographical smoothing

- In a worldwide context, the rush hour is timeshifted
- ! Worldwide peaks are smoothed by convolution

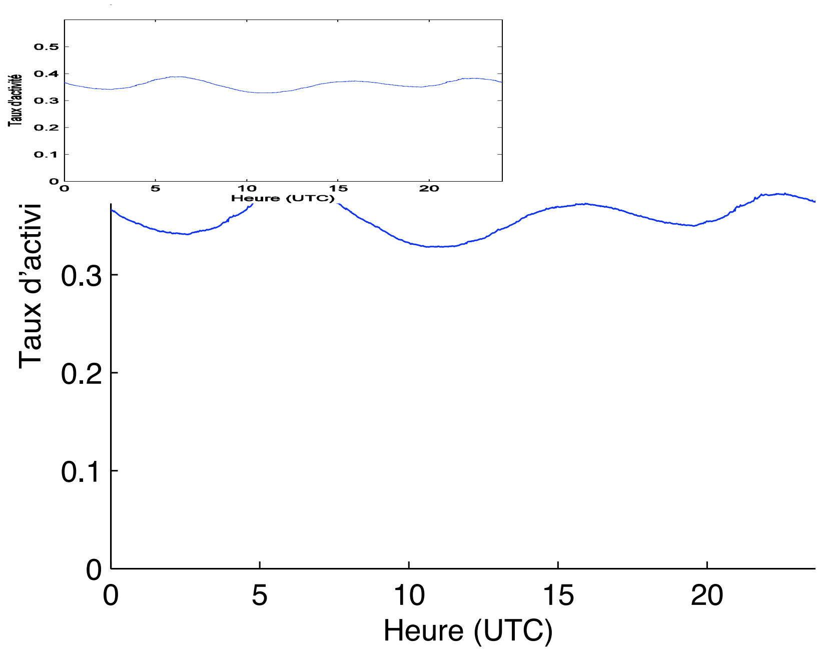

Pros and cons of worldwide convolution

- Use the smoothed peak in place of the real one
  - More clients,
  - More streamrate,
  - Less servers
- Each time the local dimensioning rules are broken, extra-bandwidth MUST come from transatlantic/transcontinental links
  - Latency issues,
  - Core networks gets loaded,
  - Isn't managing local connections the current network trend?

Bonus 2:
streaming in the open model

Target streamrate

- We want to know if a given rate r can be achieved...

- If $\bar{T_S}\bar{u_S} + \frac{U_E}{\lambda} \geq k$, overprovisioning is achieved.

- Otherwise, the achievability condition is

$$\lambda k \leq \frac{U_E}{1 - \alpha_L - \beta\alpha_S}$$

- The intensity $\frac{U_E}{k}$ that can be supported by servers alone is leveraged by $\frac{1}{1-\alpha_L-\beta\alpha_S}$

Introducing Tit-for-Tat

- For a proper understanding of share-ratio in BitTorrent, we need a Tit-for-Tat model first
- Assume some exchanges are reciprocity-driven
- *d* is not constant anymore, T_L neither, and $p_L \neq p$
- Several models try to describe Tit-for-Tat. The simplest gives

$$d(u) = \gamma u + (1 - \gamma)\bar{u_L} + \frac{\bar{T_S}}{\bar{T_L}}\bar{u_S} + \frac{U_E}{\lambda \bar{T_L}}$$

- γ is the Tit-for-Tat coefficient
- No easy way to solve this in the general case, but one can
  - Consider specific cases,
  - Find singularities,
  - Approximate the solution,
  - Use a numeric solver

Bonus 3:
Freeriders in BitTorrent

Theoretical application: Freeriders in the Sky

- Assume a proportion p_f of newcomers are freeriders (no upload at all).
- Question: how can a stationary regime exist in presence of freeriders?
- Answer:

$$p_f \leq (1-\gamma) + \gamma \left(\frac{\bar{T_S} \bar{u_S}}{k} + \frac{U_E}{\lambda k} \right)$$

- Beyond the critical value, the system cannot treat the freeriders, which accumulate