A Review of Scaling Behaviors in Internet Traffic

Steve UHLIG

Computer Science and Engineering Dept.
Université Catholique de Louvain, Belgium

Agenda

- Internet traffic : a layered perspective
- What we know about scaling in Internet traffic
- Scaling viewed by a networker
- Conclusions

Internet traffic : a layered perspective

The Internet stack

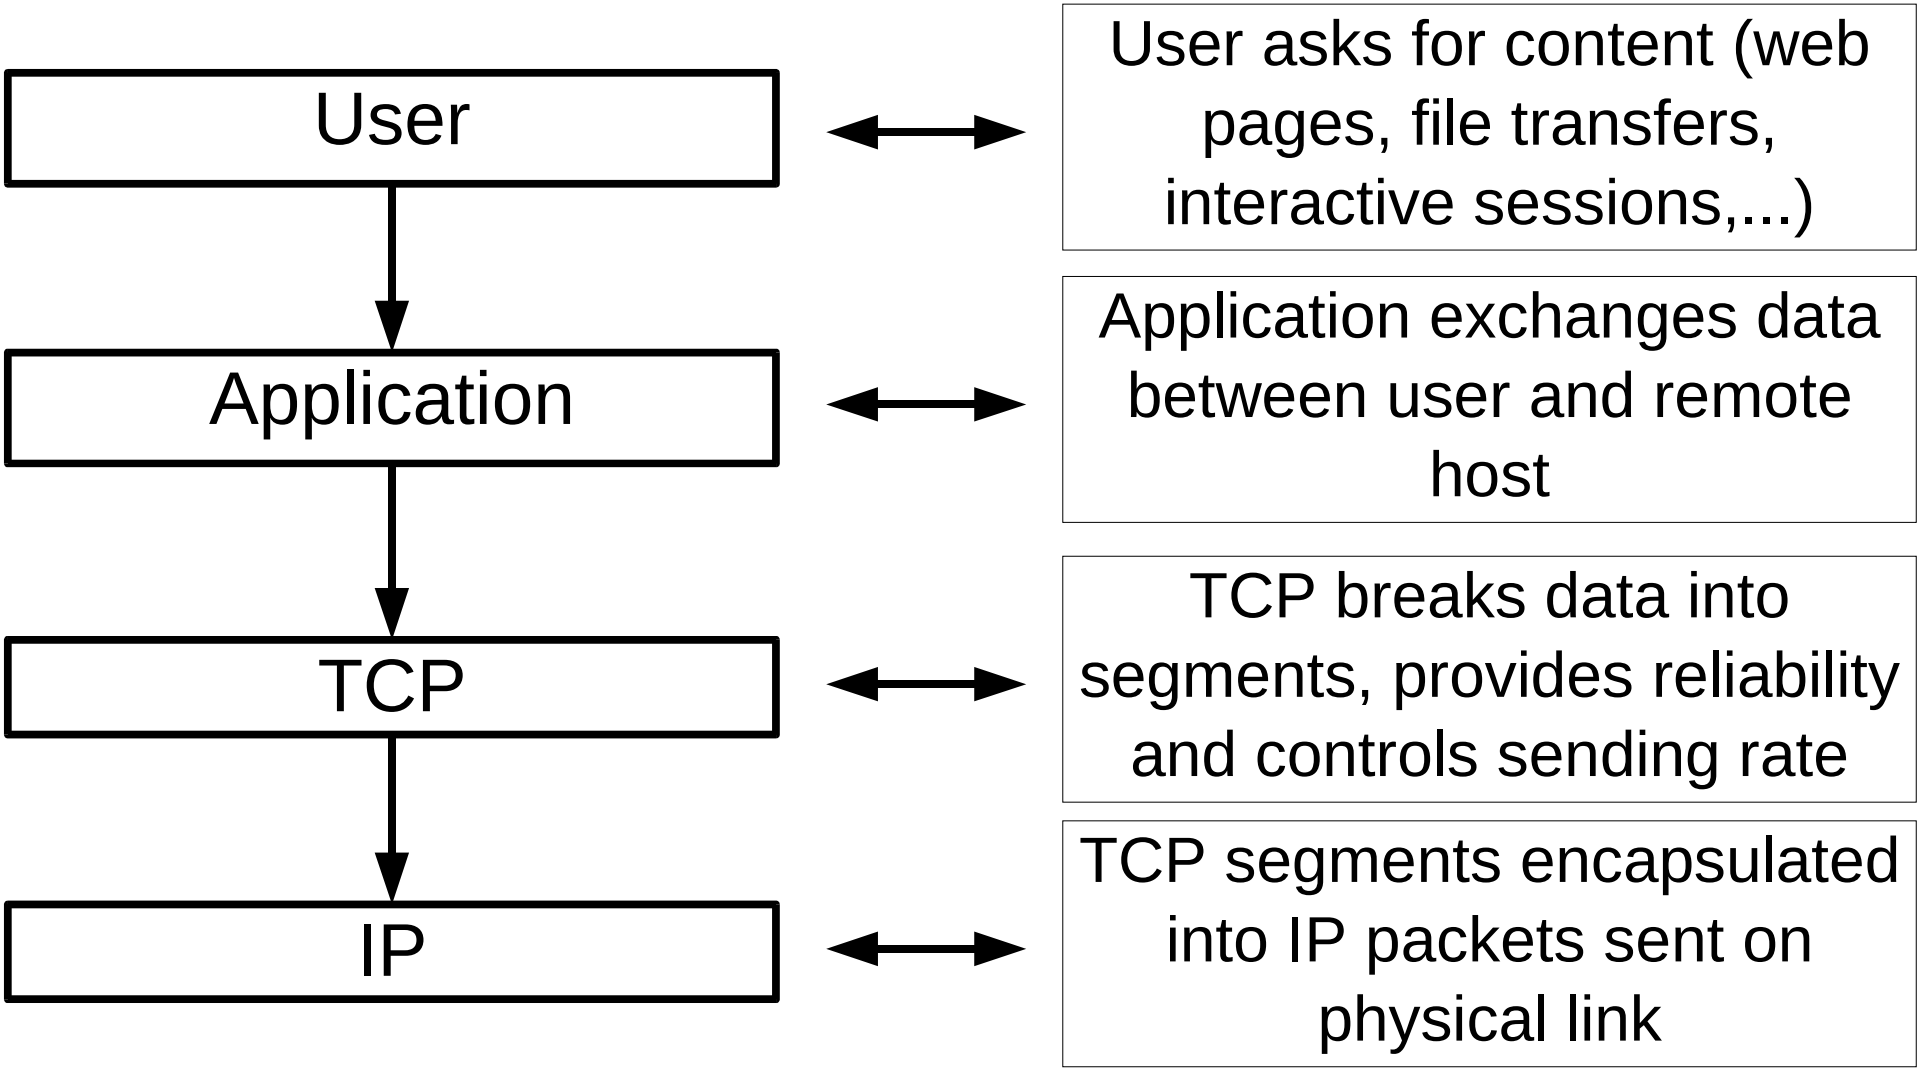

The Internet stack illustrated

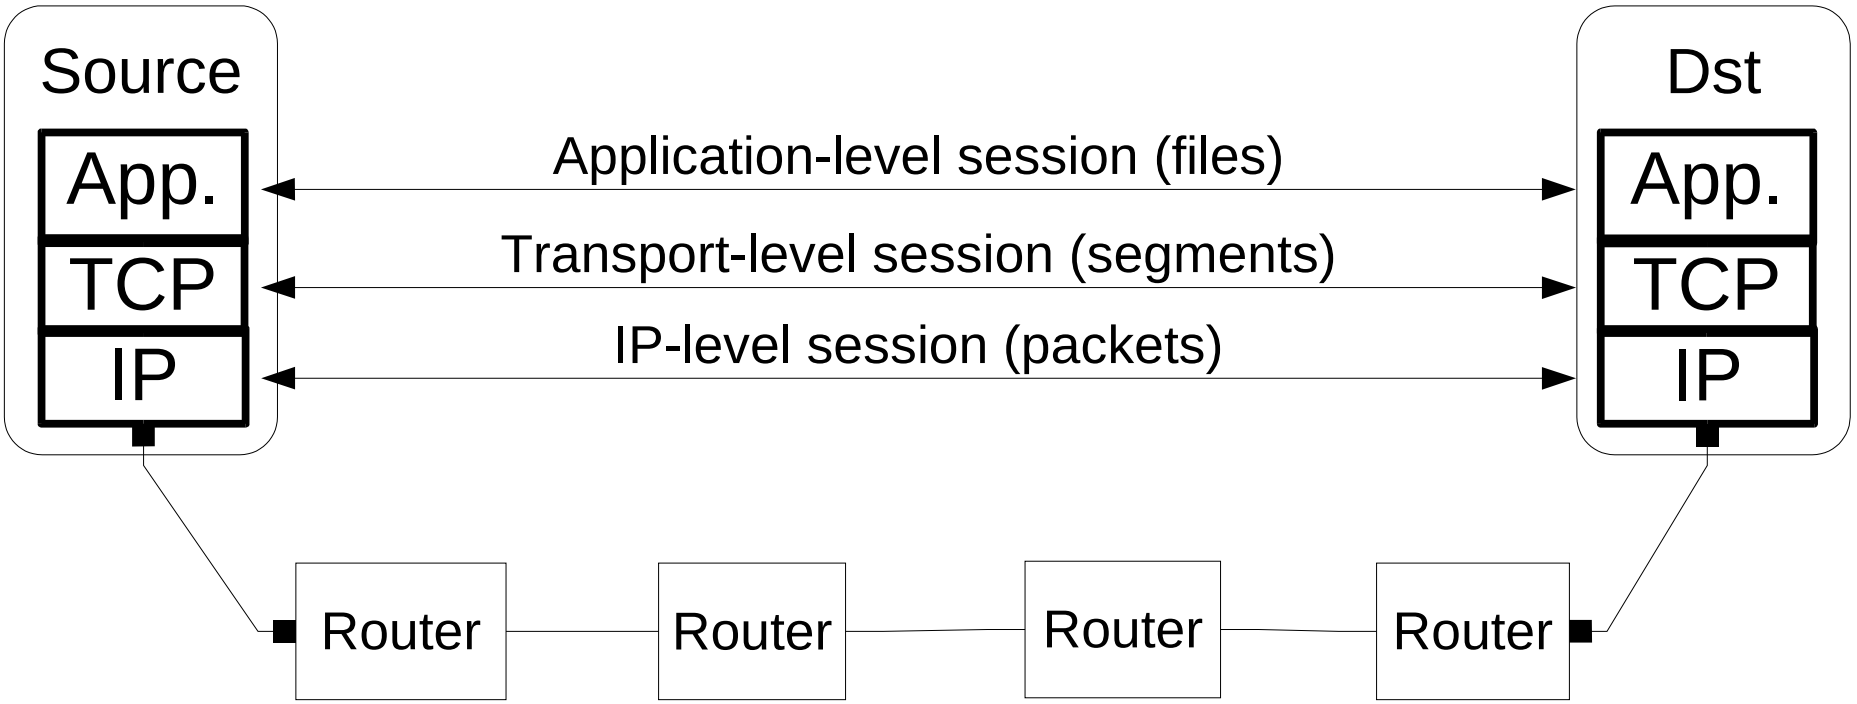

Application-transport interface

1. Application decides to exchange data between source and destination hosts

2. Several TCP connections established to exchange each "object"

3. Single IP session as seen from the outside (with ON/OFF activity periods)

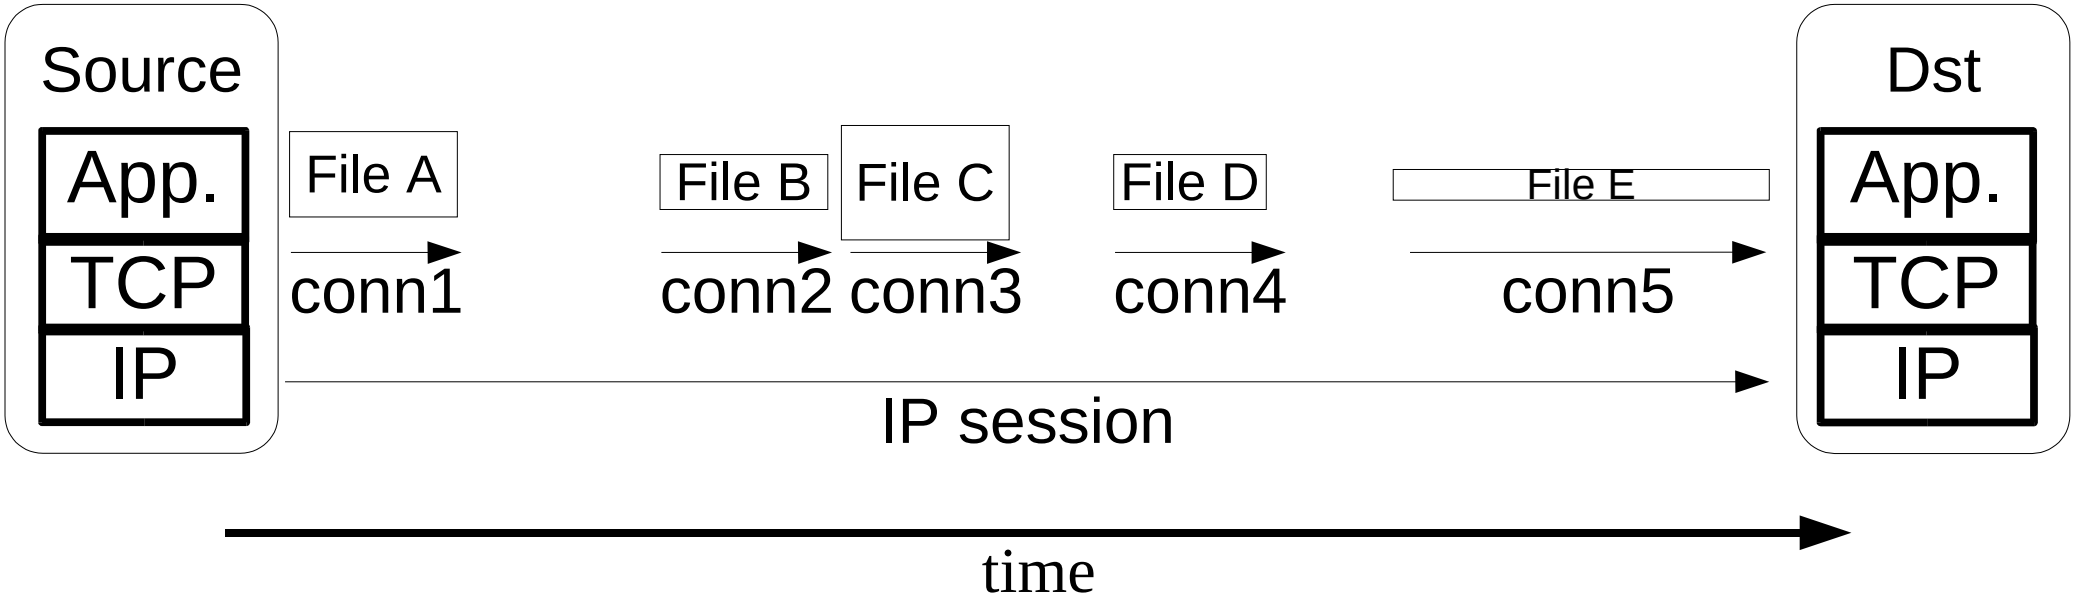

Transport-IP interface

1. Application sends data to transport layer
2. TCP breaks data into segments
3. TCP state machine decides how and when segments are to be sent to IP layer
4. IP layer sends packets over network

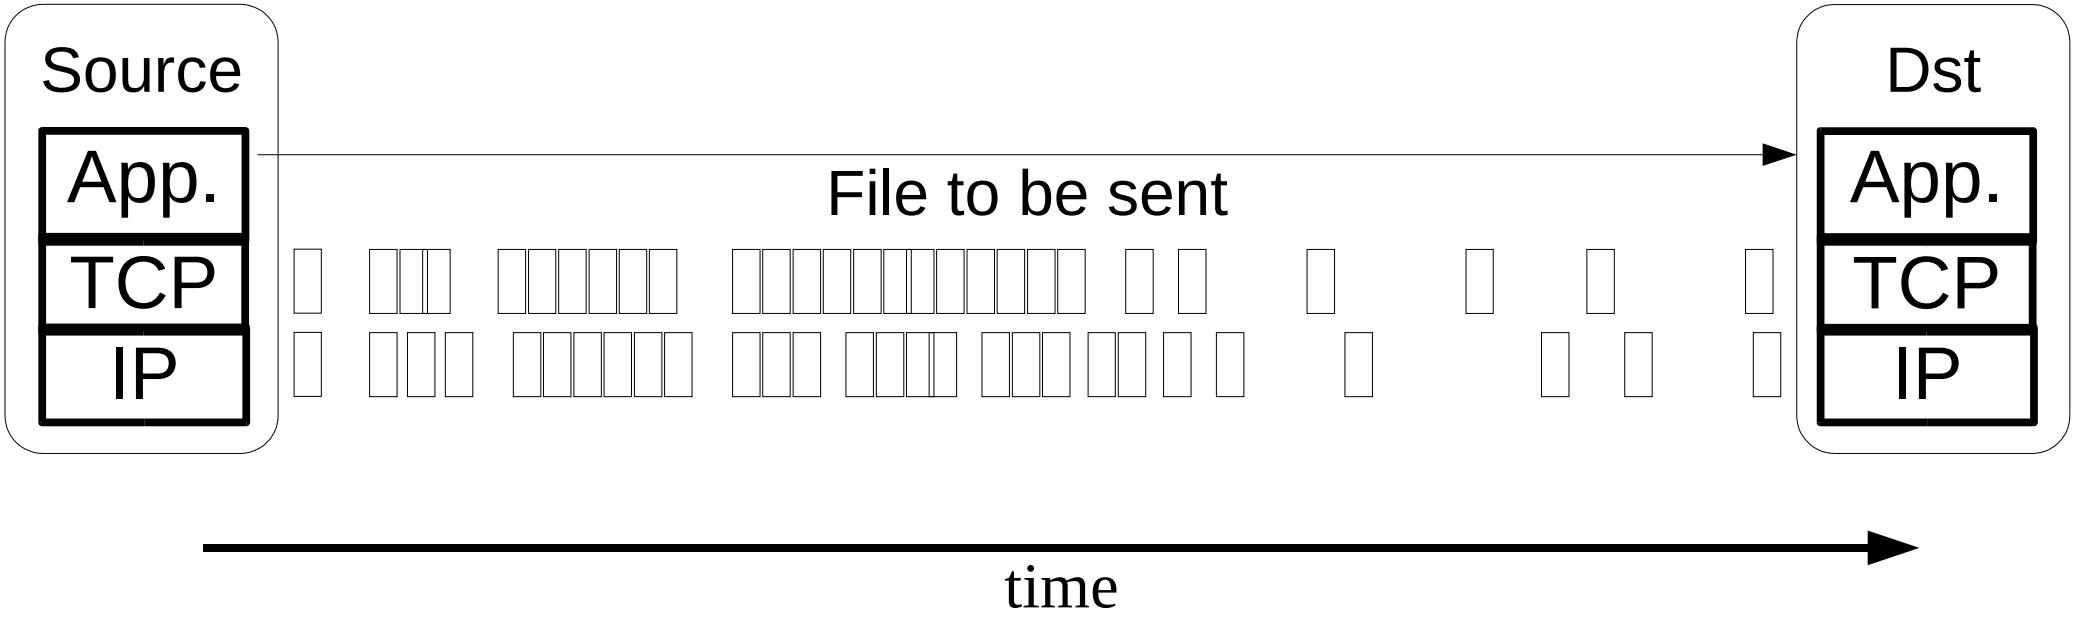

What we know about scaling in Internet traffic

Application-layer scaling

- **Heavy-tails in data flows exchanged between applications :**
  - **Power-laws in the web objects sizes [Crovella et al., 1996]**
  - **Generate second-order self-similarity (FBM or Lévy motion) : independent sources with heavy-tailed ON/OFF activity periods [Taqqu et al., 1997] [Willinger et al., 2001]**
  - **see [Downey, 2001] for a recent review of long-tailed distributions in applications**

TCP-level scaling

- **Breaking of application-level flows into segments by TCP :**
  - **conservative cascades [Feldmann et al., 1998] [Gilbert et al., 1999] [Veitch et al., 2000] (see Pierre Chainais and Patrice Abry's presentations)**
- **Scaling in arrivals of TCP flows :**
  - **second-order scaling [Feldmann, 2000]**
  - **high-order scaling and non-stationarity [Uhlig, 2004]**

IP-level scaling

- **Failure of Poisson assumption in IP traffic [Paxson et al., 1995]**
- **Self-similarity at the source activity level (hosts and networks) [Uhlig et al., 2001]**
- **Scaling at IP level does not depend on TCP flow arrivals [Hohn et al., 2002]**
- **Non-stationary Poisson IP traffic below seconds and LRD beyond [Karagianis et al., 2004]**

Scaling viewed by a networker

Application-level scaling

- **Distributional properties are an invariant of the Internet.**
- **Nothing to be done since nobody has control over the distribution of the objects' size to be transferred across the network.**
- **Self-similarity due to applications has to be lived with.**

TCP-level scaling

Cascade behavior of TCP traffic :

- **Valid at timescales up to a few RTT's [Uhlig, 2003]**
- **Valid for aggregation of TCP flows (not individual)**
- **Of what practical use is the cascade model ? What can be empirically detected about TCP flows ? (interesting open question)**

TCP-level scaling

Arrivals of TCP flows :

- **no scaling (old) or multifractal scaling (recent) below seconds**
- **LRD between seconds and minutes (due to TCP flows within user sessions ?)**
- **self-similarity (not LRD) at timescales larger than minutes (user sessions ?)**

- **Aggregate users behavior is difficult to control and should preferably not be controlled at all.**

IP-level scaling

Arrivals of IP flows :

- **Origin of scaling in IP flows arrivals ?**
  - **aggregate user behavior**
  - **is heavy-tailed ON/OFF activity the explanation ? possible.**
  - **still we aren't sure about why IP flows contain scaling at timescales beyond seconds**

Macroscopic aspects

- **Is network topology involved in large timescales scaling of the traffic ?**
- **Current traffic models care only about what is seen over one physical link.**
- **What about interactions between hosts, networks, ...?**
- **Is scaling at large timescales an emergent feature of the traffic ?**

The future

- **What are the causes of scaling ? Current state-of-the-art provides some answers.**
- **What can be done about scaling from a networking viewpoint ? (for now, nothing)**
- **Descriptive models are not enough for networking guys.**
- **They want predictive and physical models of network traffic.**

Conclusions

- **Presence of scaling behaviors in network traffic is not questionable.**
- **Several “causes” have already been identified.**
- **Identifying scaling is useless by itself (except for writing papers).**
- **Now the focus should be on the physics of scaling and “what can be done about it ?”**