Network dimensioning for elastic traffic based on flow-level QoS

Pasi Lassila and Jorma Virtamo

Networking Laboratory
Helsinki University of Technology

Introduction (1)

- Dimensioning problem
  - Given the network (nodes, links, routing) and the traffic demands
  - Determine link capacities such that a given criterion is fulfilled
  - Connection with QoS: the criterion should be a meaningful performance metric for the network and the applications
- "Good old" circuit switched networks
  - QoS determined by blocking probability
  - Dimensioning based on classical Erlang formula
  - Erlang formula is insensitive to call holding time distribution $\Rightarrow$ robustness
- The current best effort Internet
  - No established QoS metric has been defined
  - Majority of Internet traffic is elastic (controlled by TCP)
  - For elastic traffic, a natural QoS metric is the flow throughput and dimensioning could be based on that

Introduction (2)

- The "multi-service" Internet
  - Existing QoS architectures (IntServ/DiffServ) have proven problematic
    - QoS is control realized "at the packet level" assuming substantial sized buffers but the performance of the mechanisms always heavily depends on the packet level statistics
    - Making mechanisms measurement based and adaptive helps but also adds complexity
    - QoS parameters: packet delay, packet loss
  - New QoS architecture based on flow-aware approach [J. Roberts et al.]
    - Traffic consists essentially of realtime traffic and elastic traffic
    - Routers are bufferless $\Rightarrow$ eliminates complex self-similarity effects at the packet level
    - QoS control realized through flow-level control of arriving flows (access control)
    - QoS parameters: flow throughput and service accessibility
  - Dimensioning the multi-service Internet
    - Independent of the chosen QoS architecture, traffic can be largely categorized in two types: streaming and elastic
    - Dimensioning should take into account both types of traffic
    - For the flow aware approach the dimensioning problem and the QoS control are directly related, but for the IntServ/DiffServ approach the connection is obviously indirect

Introduction (3)

- Scope of this study:
  - We only consider elastic traffic for which dimensioning is naturally based on flow throughput

- Theoretical framework
  - We apply models based on the notion of balanced fairness
    - Idealized bandwidth sharing scheme that allows explicit evaluation of throughput
    - Minimal assumptions on traffic: session arrivals are Poisson
    - The models are insensitive, i.e., performance only depends on the load
    - Balanced fairness approximates max-min fairness and proportional fairness (which are approximated by TCP), for which performance can not be easily evaluated
  - Thus, balanced fairness is used as a computationally efficient tool that allows robust dimensioning based on flow-level throughput

- Practical use of the results:
  - Solutions can be used as a sort of “educated guess” for how much bandwidth is needed to have a given throughput performance

Related work

- General remarks
  - Given performance requirement can be satisfied by a continuum of possible solutions for the link capacities
  - Thus, an optimization formulation is required to choose the “best” solution
  - Main differences in the system model: static/stochastic
- “Classic” Bertsekas & Gallager square-root method
  - System model: open network of M/M/1 queues
  - Task: minimize link costs s.t. total mean packet delay equals some target delay
  - Problem: as such the network model is a packet-level model
- Multi-commodity flow optimization models [M. Pioro et al.]
  - Task: maximize flow allocation according to some notion of fairness when the allowed overall network cost is given (budget constraint)
  - Formulations for max-min fairness, proportional fairness, “robust” networks
  - Problem: network model treats traffic as fluid and ignores the dynamic/stochastic nature of traffic, no notion of offered traffic
- (Link dimensioning: robust link dimensioning using PS models)

Dimensioning problem decomposition

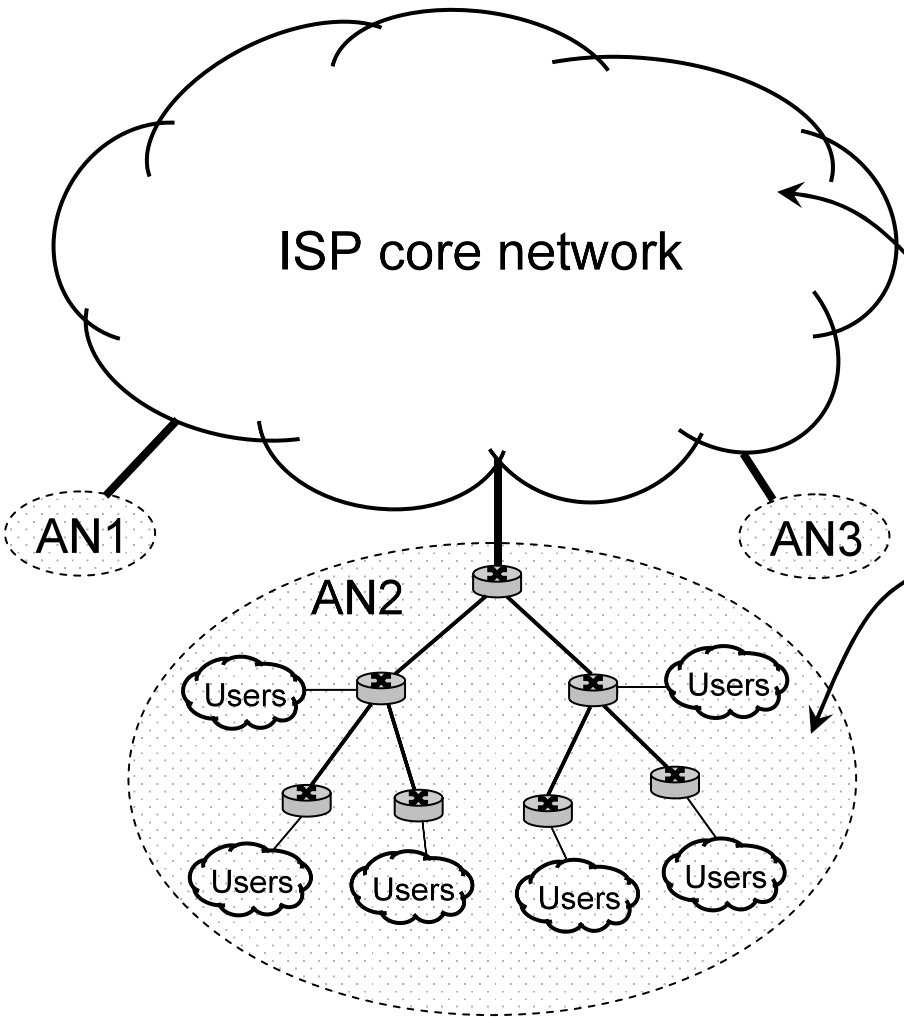

- Detailed dimensioning applied only where bandwidth is scarce,
  - i.e., in the access networks (ANs)
- ISP core network dimensioned such that it is “transparent” to the users
  - Link loads < 0.5 (rough guess)
- ANs dimensioned based on per flow throughput
  - Utilizing new results on balanced fairness

Optimization problem for access networks (1)

- ANs have multi-level tree topology
  - *N* classes (routes), *J* links
- Optimization formulation

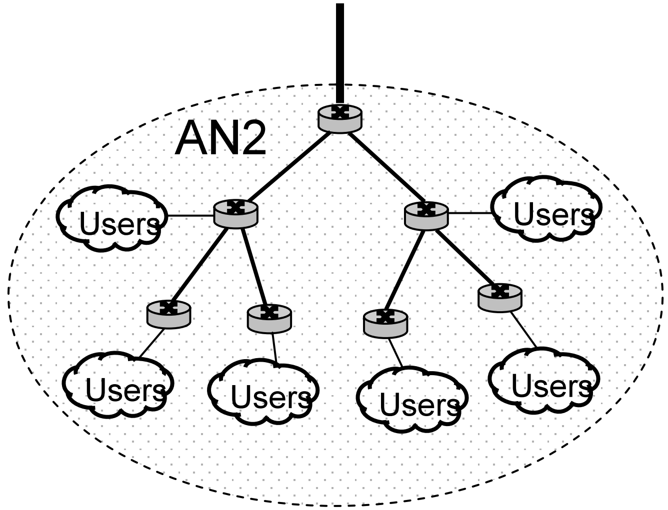

C_j = capacity of link *j*

ρ_i = load of class *i* (bit/s)

$\gamma_i(\mathbf{C})$= througput of class *i*

$T_i(\mathbf{C})$= througput of class *i*

F_j = classes that use link *j*

<u>Version 1</u>

$$\min \quad \sum_{j=1}^{J} C_j$$

$$\text{s.t.} \quad \sum_{i \in F_j} \rho_i < C_j \ , \ j = 1, \ldots, J$$

$$\gamma_i(\mathbf{C}) \geq \gamma_i^{\min}, \ i = 1, \ldots, N$$

<u>Version 2</u>

$$\min \quad \sum_{j=1}^{J} C_j$$

$$\text{s.t.} \quad \sum_{i \in F_j} \rho_i < C_j \ , \ j = 1, \ldots, J$$

$$T_i(\mathbf{C}) \leq T_i^{\max}, \ i = 1, \ldots, N$$

Optimization problem for access networks (2)

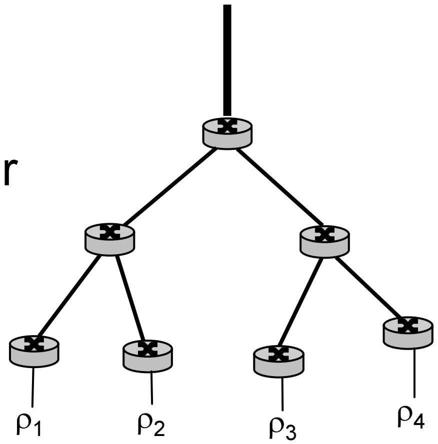

- The throughput $\gamma_i(\mathbf{C})$ can be computed
  - exactly using an efficient recursive algorithm, or
  - approximately with store-and-forward lower bound or
  - the more accurate parking-lot approximation
- Note on using the store-and-forward bound
  - Conservative bound (in some sense "safe")
  - Dimensioning problem becomes equivalent to the Bertsekas-Gallager square-root method
    - In the square-root method, the model gives an exact dimensioning rule for the packet delay a network of M/M/1 queues
    - In our case, it gives an approximate dimensioning rule for the flow-level performance
- Other "spices"
  - Possible to have per-flow rate limitations (SF-bound or exact solution)
  - Cost function can be arbitrary (also nonlinear)

Example: homogeneous tree with *n* branches

- Optimal solution using the SF-bound:

$$C_0 = n\rho + \gamma^{\min}\left(1+\sqrt{n}\right) \quad , \quad C_1 = \rho + \gamma^{\min}\frac{1+\sqrt{n}}{\sqrt{n}}$$

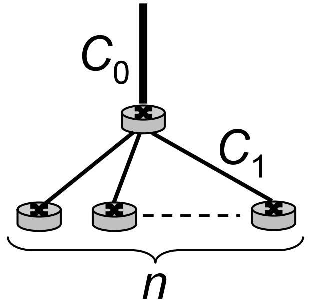

- Optimal costs for various $\gamma^{\min}$ as a function of *n*

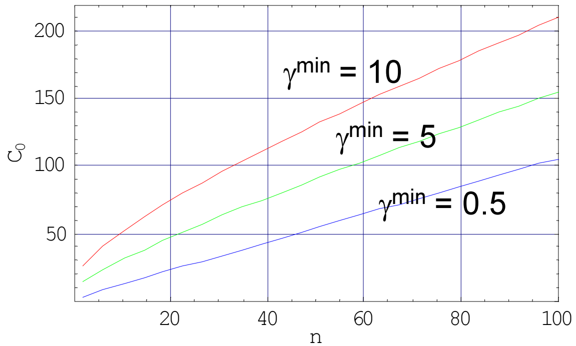

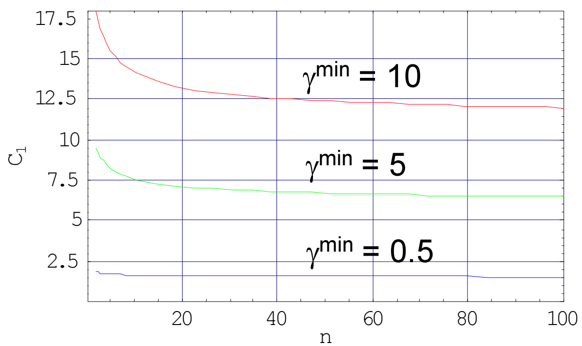

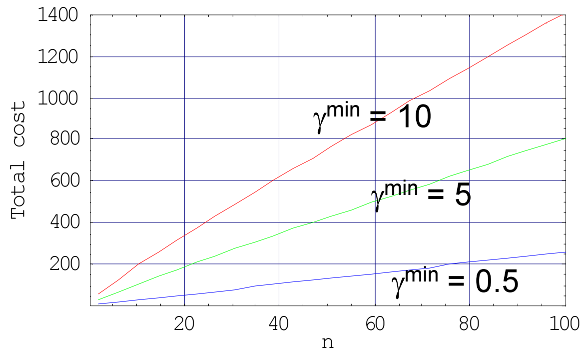

- Relative error between SF dimensioning and exact solution

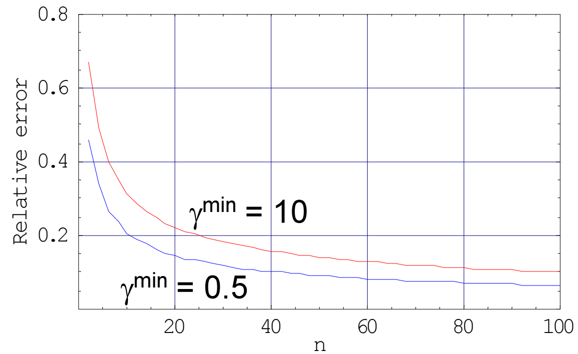

Conclusions

- Models based on balanced fairness allow robust dimensioning for elastic traffic
- The dimensioning can be based on a natural flow-level QoS requirement, i.e., the per flow throughput
- Using the SF-bound for approximating the performance, the optimization problem is also numerically simple
  - Efficient exact solution algorithms can be used to verify the “actual” performance
- Future research
  - Dimensioning of networks with both elastic and streaming traffic (performance bounds for such networks are available)
  - Dimensioning of interference limited wireless networks with multi-hop radio links (RANs)