A Contract and Balancing Mechanism for Sharing Capacity in a Communication Network

Richard Steinberg
University of Cambridge

Joint with

Eddie Anderson, AGSM
Frank Kelly, University of Cambridge

July 2005

Problem

Problem

The owner of a communication network wishes to sell bandwidth to users

Problem

The owner of a communication network wishes to sell bandwidth to users

- Bandwidth: Amount of data that can be transmitted per unit time via fibre optic cables (specified in *bits per second*, bps)

Problem

The owner of a communication network wishes to sell bandwidth to users

- Bandwidth: Amount of data that can be transmitted per unit time via fibre optic cables (specified in *bits per second*, bps)

- Users are those who require a dedicated network

Problem

The owner of a communication network wishes to sell bandwidth to users

- Bandwidth: Amount of data that can be transmitted per unit time via fibre optic cables (specified in *bits per second*, bps)

- Users are those who require a dedicated network
  - small telecoms companies

Problem

The owner of a communication network wishes to sell bandwidth to users

- Bandwidth: Amount of data that can be transmitted per unit time via fibre optic cables (specified in *bits per second*, bps)

- Users are those who require a dedicated network
  - small telecoms companies
  - business that host network applications (e.g., Blockbuster Entertainment on demand)

Problem

The owner of a communication network wishes to sell bandwidth to users

- Bandwidth: Amount of data that can be transmitted per unit time via fibre optic cables (specified in *bits per second*, bps)

- Users are those who require a dedicated network
  - small telecoms companies
  - business that host network applications (e.g., Blockbuster Entertainment on demand)
  - other (e.g., Merrill Lynch, for one-off video broadcast for its analysts)

The Customers' Dilemma

The Customers' Dilemma

Customers can either:

The Customers' Dilemma

Customers can either:

(1) Buy capacity for only their regular needs

The Customers’ Dilemma

Customers can either:

(1) Buy capacity for only their regular needs

- But then they’ll have problems when usage spikes

The Customers’ Dilemma

Customers can either:

(1) Buy capacity for only their regular needs

- But then they’ll have problems when usage spikes

(2) Buy sufficient capacity to serve maximum demand

The Customers' Dilemma

Customers can either:

(1) Buy capacity for only their regular needs

- But then they'll have problems when usage spikes

(2) Buy sufficient capacity to serve maximum demand

- But then they waste money on access they're not using

Two types of communication networks

Two types of communication networks

- Circuit-switched networks

Two types of communication networks

- Circuit-switched networks
  - A specific path to destination is obtained for the message

Two types of communication networks

- Circuit-switched networks
  - A specific path to destination is obtained for the message
  - No other information travels along the lines involved

Two types of communication networks

- Circuit-switched networks
  - A specific path to destination is obtained for the message
  - No other information travels along the lines involved
  - Example: the public telephone network

Two types of communication networks

- Circuit-switched networks
  - A specific path to destination is obtained for the message
  - No other information travels along the lines involved
  - Example: the public telephone network

- Packet-switched (IP) networks

Two types of communication networks

- Circuit-switched networks
  - A specific path to destination is obtained for the message
  - No other information travels along the lines involved
  - Example: the public telephone network

- Packet-switched (IP) networks
  - Data is broken down into small units, called packets

Two types of communication networks

- Circuit-switched networks
  - A specific path to destination is obtained for the message
  - No other information travels along the lines involved
  - Example: the public telephone network

- Packet-switched (IP) networks
  - Data is broken down into small units, called packets
  - Packets individually routed & re-assembled at destination

Two types of communication networks

- Circuit-switched networks
  - A specific path to destination is obtained for the message
  - No other information travels along the lines involved
  - Example: the public telephone network

- Packet-switched (IP) networks
  - Data is broken down into small units, called packets
  - Packets individually routed & re-assembled at destination

Our focus will be on packet-switched networks.

Sharing capacity

Sharing capacity

- IP networks have congestion control mechanisms to allow users to share capacity…

Sharing capacity

- IP networks have congestion control mechanisms to allow users to share capacity…

  … and thus absorb random fluctuations in users' demands

Sharing capacity

- IP networks have congestion control mechanisms to allow users to share capacity…

  … and thus absorb random fluctuations in users' demands

- The rate at which a source sends packets to the destination is controlled by TCP

Sharing capacity

- IP networks have congestion control mechanisms to allow users to share capacity…

  … and thus absorb random fluctuations in users' demands

- The rate at which a source sends packets to the destination is controlled by TCP

  … implemented as software on computers on IP networks

Sharing capacity

- IP networks have congestion control mechanisms to allow users to share capacity…

  … and thus absorb random fluctuations in users' demands

- The rate at which a source sends packets to the destination is controlled by TCP

  … implemented as software on computers on IP networks

- TCP shares bandwidth amongst flows

How TCP operates

How TCP operates

- Each time a packet arrives at the destination, an acknowledgement packet is sent back to the source

How TCP operates

- Each time a packet arrives at the destination, an acknowledgement packet is sent back to the source
- When the destination becomes overloaded, one or more packets are lost

How TCP operates

- Each time a packet arrives at the destination, an acknowledgement packet is sent back to the source
- When the destination becomes overloaded, one or more packets are lost
  - the source knows that the packet was lost because it failed to receive an acknowledgement packet

How TCP operates

- Each time a packet arrives at the destination, an acknowledgement packet is sent back to the source
- When the destination becomes overloaded, one or more packets are lost
  - the source knows that the packet was lost because it failed to receive an acknowledgement packet
  - the source slows down

How TCP operates

- Each time a packet arrives at the destination, an acknowledgement packet is sent back to the source

- When the destination becomes overloaded, one or more packets are lost
  - the source knows that the packet was lost because it failed to receive an acknowledgement packet
  - the source slows down

- TCP gradually increases the sending rate until it again receives an indication of congestion

How TCP operates

- Each time a packet arrives at the destination, an acknowledgement packet is sent back to the source
- When the destination becomes overloaded, one or more packets are lost
  - the source knows that the packet was lost because it failed to receive an acknowledgement packet
  - the source slows down
- TCP gradually increases the sending rate until it again receives an indication of congestion
- In this way, TCP serves to discover and utilise whatever bandwidth is available and to share it amongst flows

Structure of a Packet

Structure of a Packet

- All packets have three parts

Structure of a Packet

- All packets have three parts
  - IP header

Structure of a Packet

- All packets have three parts
  - IP header
  - TCP header

Structure of a Packet

- All packets have three parts
  - IP header
  - TCP header
  - data from “application layer”

Structure of a Packet

- All packets have three parts
  - IP header
  - TCP header
  - data from “application layer”

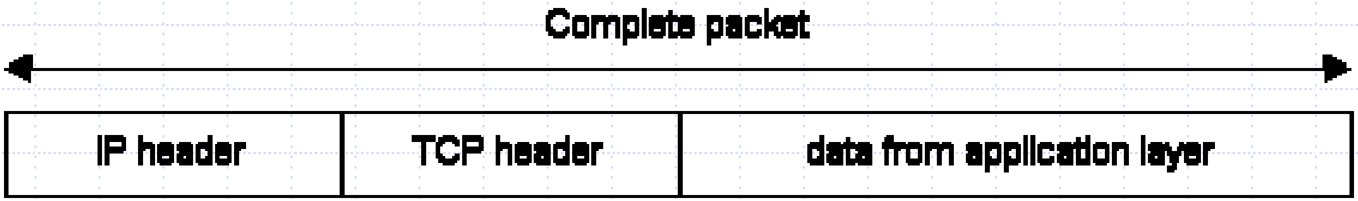

Structure of a Packet

- All packets have three parts
  - IP header
  - TCP header
  - data from "application layer"

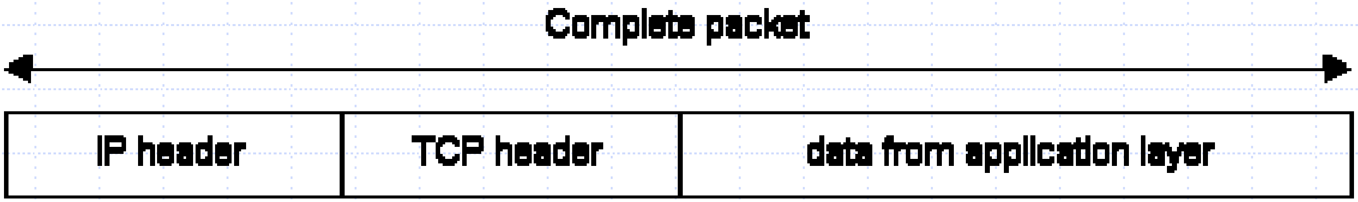

- There is currently one bit in IP header reserved for some "future use"

Structure of a Packet

- All packets have three parts
  - IP header
  - TCP header
  - data from “application layer”

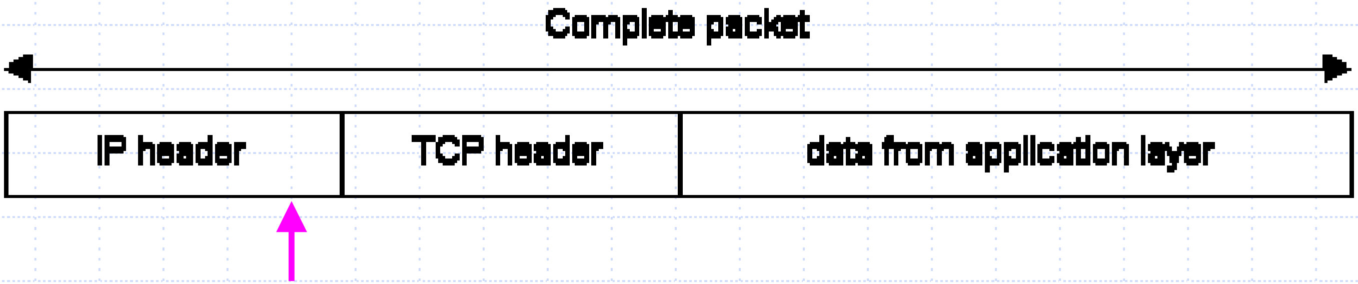

- There is currently one bit in IP header reserved for some “future use” (can be set to 0 or 1)

Is there a better way to signal congestion?

Is there a better way to signal congestion?

Two problems with dropped packets

Is there a better way to signal congestion?

Two problems with dropped packets

(1) Wasteful of system resources

Is there a better way to signal congestion?

Two problems with dropped packets

(1) Wasteful of system resources

- Dropped packets need to be re-sent

Is there a better way to signal congestion?

Two problems with dropped packets

(1) Wasteful of system resources

- Dropped packets need to be re-sent

(2) The congestion signal is delayed

Is there a better way to signal congestion?

Two problems with dropped packets

(1) Wasteful of system resources

- Dropped packets need to be re-sent

(2) The congestion signal is delayed

- Until the point when packets begin to be dropped, users are unaware of imminent congestion problems

Is there a better way to signal congestion?

Two problems with dropped packets

(1) Wasteful of system resources

- Dropped packets need to be re-sent

(2) The congestion signal is delayed

- Until the point when packets begin to be dropped, users are unaware of imminent congestion problems

A better way

Is there a better way to signal congestion?

Two problems with dropped packets

(1) Wasteful of system resources

- Dropped packets need to be re-sent

(2) The congestion signal is delayed

- Until the point when packets begin to be dropped, users are unaware of imminent congestion problems

A better way

There is a 'Proposed Standard' for congestion control...

ECN: Explicit Congestion Notification

ECN: Explicit Congestion Notification

- Packets encountering long queues can have that extra bit set to indicate congestion

ECN: Explicit Congestion Notification

- Packets encountering long queues can have that extra bit set to indicate congestion
- Users detecting ECN marks should respond by reducing transmission rates

ECN: Explicit Congestion Notification

- Packets encountering long queues can have that extra bit set to indicate congestion
- Users detecting ECN marks should respond by reducing transmission rates
- Thus, ECN marks are designed to be used to control congestion

ECN: Explicit Congestion Notification

- Packets encountering long queues can have that extra bit set to indicate congestion
- Users detecting ECN marks should respond by reducing transmission rates
- Thus, ECN marks are designed to be used to control congestion
- Our idea: ECN marks can be also used to control the supply of capacity

“Contract and Balancing Mechanism” (CBM)

(Anderson, Kelly and Steinberg 2005)

"Contract and Balancing Mechanism" (CBM)

(Anderson, Kelly and Steinberg 2005)

- Owner of link sells contracts for usage over a period

"Contract and Balancing Mechanism" (CBM)

(Anderson, Kelly and Steinberg 2005)

- Owner of link sells contracts for usage over a period
- At end of period, users participate in a balancing process

"Contract and Balancing Mechanism" (CBM)

(Anderson, Kelly and Steinberg 2005)

- Owner of link sells contracts for usage over a period
- At end of period, users participate in a balancing process
- Users make or receive payments based on:

"Contract and Balancing Mechanism" (CBM)

(Anderson, Kelly and Steinberg 2005)

- Owner of link sells contracts for usage over a period
- At end of period, users participate in a balancing process
- Users make or receive payments based on:
  - proportion of marks they received, compared with

"Contract and Balancing Mechanism" (CBM)

(Anderson, Kelly and Steinberg 2005)

- Owner of link sells contracts for usage over a period
- At end of period, users participate in a balancing process
- Users make or receive payments based on:
  - proportion of marks they received, compared with
  - capacity they had contracted for.

“Contract and Balancing Mechanism” (CBM)

Example

“Contract and Balancing Mechanism” (CBM)

Example

- Two users contract for part of 300 Mbps line for month

"Contract and Balancing Mechanism" (CBM)

Example

- Two users contract for part of 300 Mbps line for month
  - User A contracts for 100 Mbps

“Contract and Balancing Mechanism” (CBM)

Example

- Two users contract for part of 300 Mbps line for month
  - User A contracts for 100 Mbps
  - User B contracts for 200 Mbps

“Contract and Balancing Mechanism” (CBM)

Example

- Two users contract for part of 300 Mbps line for month
  - User A contracts for 100 Mbps
  - User B contracts for 200 Mbps

- Rates charged by the owner

“Contract and Balancing Mechanism” (CBM)

Example

- Two users contract for part of 300 Mbps line for month
  - User A contracts for 100 Mbps
  - User B contracts for 200 Mbps

- Rates charged by the owner
  - contractual rate: c dollars per megabit

“Contract and Balancing Mechanism” (CBM)

Example

- Two users contract for part of 300 Mbps line for month
  - User A contracts for 100 Mbps
  - User B contracts for 200 Mbps

- Rates charged by the owner
  - contractual rate: *c* dollars per megabit
  - balancing charge: b dollars per mark

“Contract and Balancing Mechanism” (CBM)

<u>Example</u>

- Two users contract for part of 300 Mbps line for month
  - User A contracts for 100 Mbps
  - User B contracts for 200 Mbps

- Rates charged by the owner
  - contractual rate: c dollars per megabit
  - balancing charge: b dollars per mark

- At beginning of the month, owner receives

"Contract and Balancing Mechanism" (CBM)

Example

- Two users contract for part of 300 Mbps line for month
  - User A contracts for 100 Mbps
  - User B contracts for 200 Mbps

- Rates charged by the owner
  - contractual rate: c dollars per megabit
  - balancing charge: b dollars per mark

- At beginning of the month, owner receives
  - $100c$ dollars from User A

“Contract and Balancing Mechanism” (CBM)

<u>Example</u>

- Two users contract for part of 300 Mbps line for month
  - User A contracts for 100 Mbps
  - User B contracts for 200 Mbps

- Rates charged by the owner
  - contractual rate: c dollars per megabit
  - balancing charge: b dollars per mark

- At beginning of the month, owner receives
  - $100c$ dollars from User A
  - $200c$ dollars from User B

During the month...

During the month...

- Number of marks received by each user is recorded:

During the month...

- Number of marks received by each user is recorded:

  z_A = number of marks received by User A

During the month...

- Number of marks received by each user is recorded:

  z_A = number of marks received by User A

  z_B = number of marks received by User B

At end of month...

At end of month...

- If User A received exactly 1/3 the total marks, then no further payments are made

At end of month...

- If User A received exactly 1/3 the total marks, then no further payments are made

- If User A received more than 1/3 the total marks, then A will pay B the balancing charge times the excess:

At end of month...

- If User A received exactly 1/3 the total marks, then no further payments are made

- If User A received more than 1/3 the total marks, then A will pay B the balancing charge times the excess:

$$\beta\,[z_A - \alpha\,(z_A + z_B)]$$

At end of month...

- If User A received exactly 1/3 the total marks, then no further payments are made

- If User A received more than 1/3 the total marks, then A will pay B the balancing charge times the excess:

$$\beta\,[z_A - \alpha\,(z_A + z_B)]$$

- If User A received less than 1/3 the total marks, then A will <u>receive</u> from B the above amount.

A word about ECN marking...

A word about ECN marking...

- Recall the basic principle of ECN marking:

A word about ECN marking...

- Recall the basic principle of ECN marking:

  Packets encountering long queues have a bit set to indicate congestion

A word about ECN marking...

- Recall the basic principle of ECN marking:

  Packets encountering long queues have a bit set to indicate congestion

- We haven’t discussed the point at which a queue is considered to be `long’

A word about ECN marking...

- Recall the basic principle of ECN marking:

  Packets encountering long queues have a bit set to indicate congestion

- We haven't discussed the point at which a queue is considered to be `long'

- However, this is not very significant, as we only need to know the proportion of ECN marks received

Player’s optimal contract amount

Player's optimal contract amount

Theorem: Suppose the users can each predict their own usage over the period.

Player’s optimal contract amount

Theorem: Suppose the users can each predict their own usage over the period.

Then a user’s optimal contract quantity is a weighted average of his anticipated usage over the period.

Player's optimal contract amount

Theorem: Suppose the users can each predict their own usage over the period.

Then a user's optimal contract quantity is a weighted average of his anticipated usage over the period.

Corollary: Users who are:

Player's optimal contract amount

Theorem: Suppose the users can each predict their own usage over the period.

Then a user's optimal contract quantity is a weighted average of his anticipated usage over the period.

Corollary: Users who are:

- price insensitive, or

Player's optimal contract amount

Theorem: Suppose the users can each predict their own usage over the period.

Then a user's optimal contract quantity is a weighted average of his anticipated usage over the period.

Corollary: Users who are:

- price insensitive, or
- have a constant usage over time

Player’s optimal contract amount

Theorem: Suppose the users can each predict their own usage over the period.

Then a user’s optimal contract quantity is a weighted average of his anticipated usage over the period.

Corollary: Users who are:

- price insensitive, or
- have a constant usage over time

should contract for precisely their anticipated usage.

Conclusions

Conclusions

Two benefits of the scheme

Conclusions

Two benefits of the scheme

- If a user has constant bandwidth requirements, i.e., never delivers more than his contracted capacity…

Conclusions

Two benefits of the scheme

- If a user has constant bandwidth requirements, i.e., never delivers more than his contracted capacity…
  - Then he will only receive--never be required to make--payments in the balancing process.

Conclusions

Two benefits of the scheme

- If a user has constant bandwidth requirements, i.e., never delivers more than his contracted capacity...
  - Then he will only receive--never be required to make--payments in the balancing process.

- If capacity of a link is fully contracted, the network owner receives no additional payments from process.

Conclusions

Two benefits of the scheme

- If a user has constant bandwidth requirements, i.e., never delivers more than his contracted capacity...

  - Then he will only receive--never be required to make--payments in the balancing process.

- If capacity of a link is fully contracted, the network owner receives no additional payments from process.

  - Thus, he is indifferent to short term behaviour of network.

Conclusions

Two benefits of the scheme

- If a user has constant bandwidth requirements, i.e., never delivers more than his contracted capacity…
  - Then he will only receive--never be required to make--payments in the balancing process.

- If capacity of a link is fully contracted, the network owner receives no additional payments from process.
  - Thus, he is indifferent to short term behaviour of network.

Conclusions

Two benefits of the scheme

- If a user has constant bandwidth requirements, i.e., never delivers more than his contracted capacity…

  - Then he will only receive--never be required to make--payments in the balancing process.

- If capacity of a link is fully contracted, the network owner receives no additional payments from process.

  - Thus, he is indifferent to short term behaviour of network.