Motivating Participation in Internet Routing Overlays

Dave Levin Randolph Baden Cristian Lumezanu
Neil Spring Bobby Bhattacharjee

University of Maryland

Internet Routing

BGP routing

End-users

Internet Routing

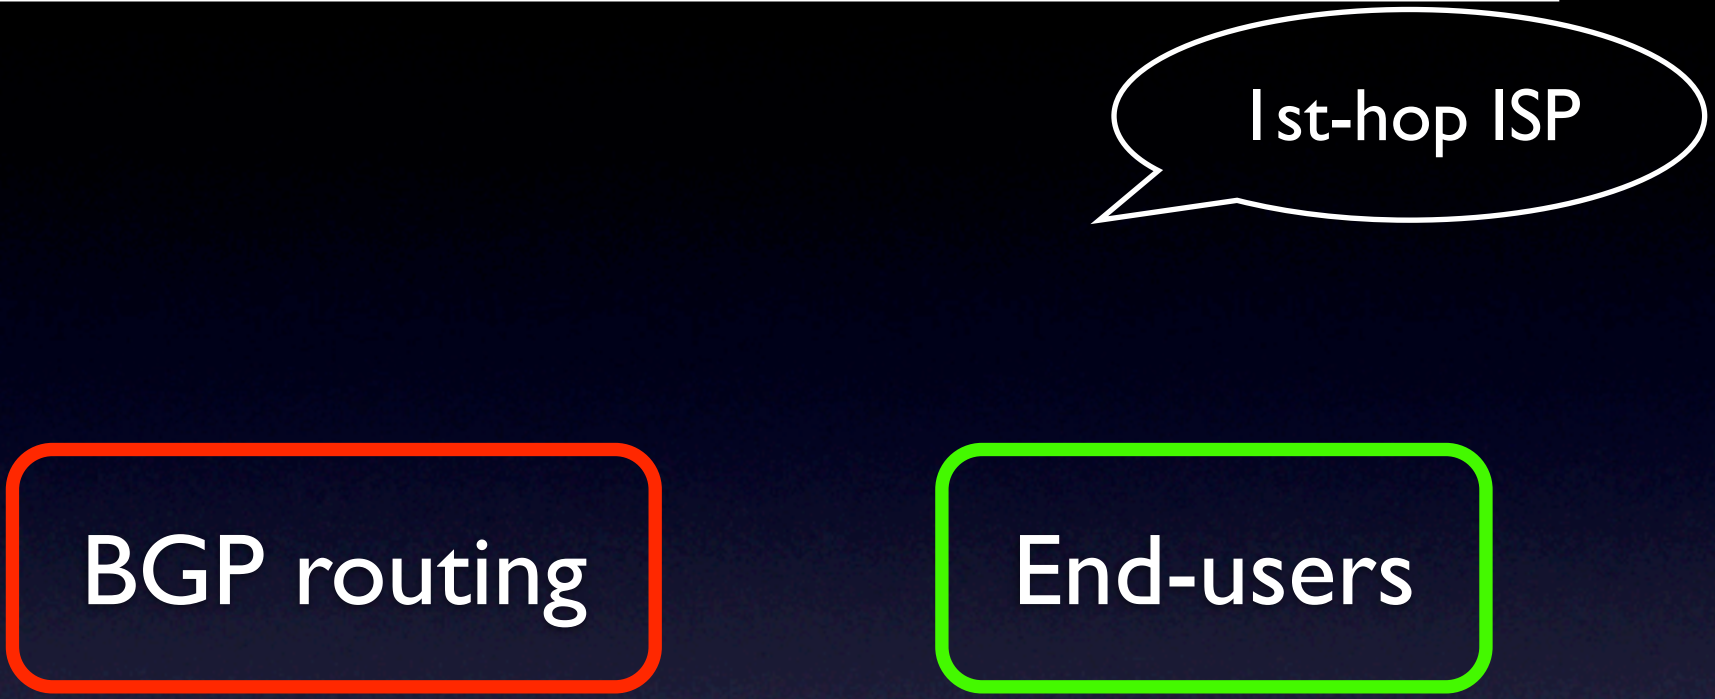

Internet Routing

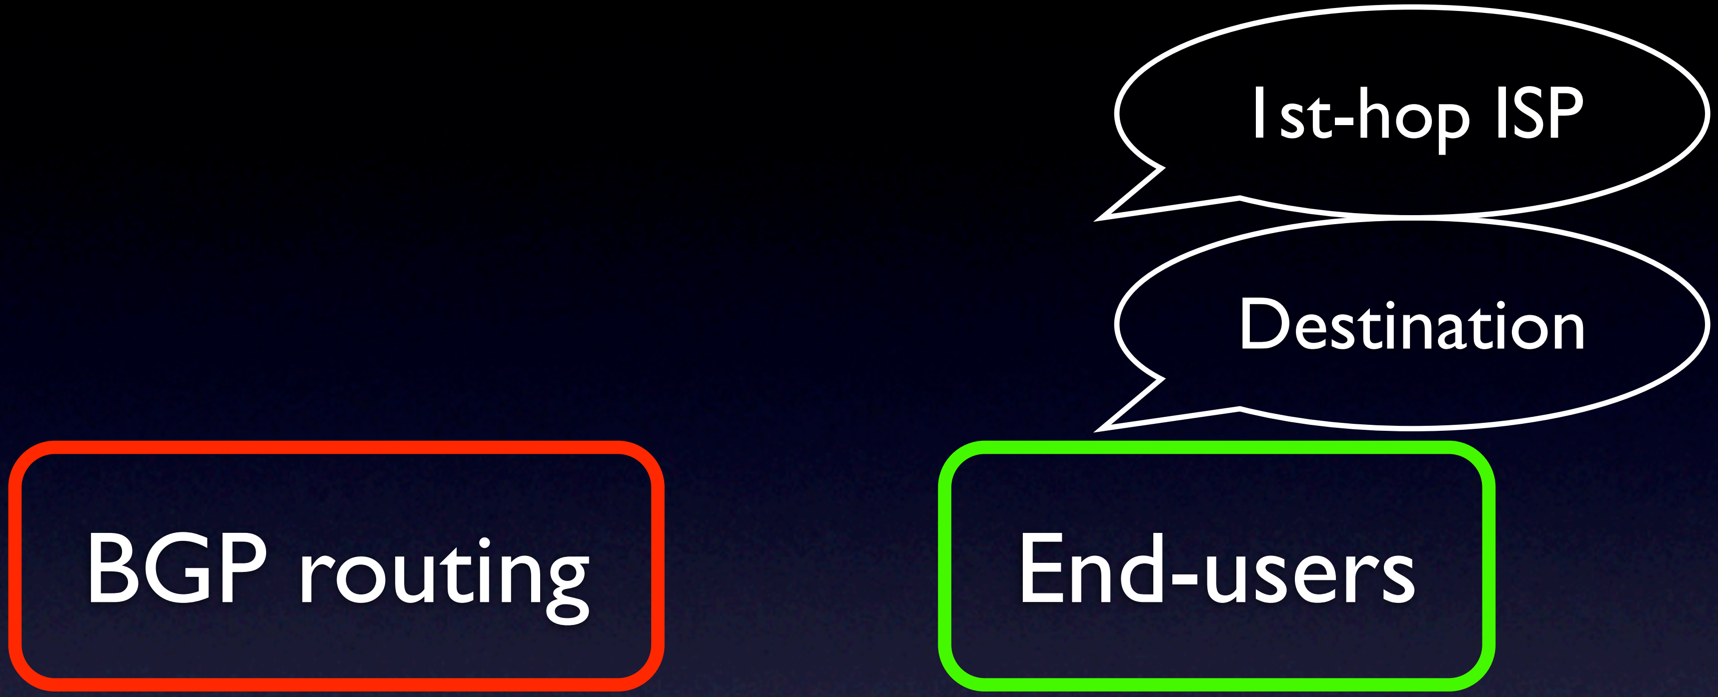

1st-hop ISP

Destination

BGP routing

End-users

Internet Routing

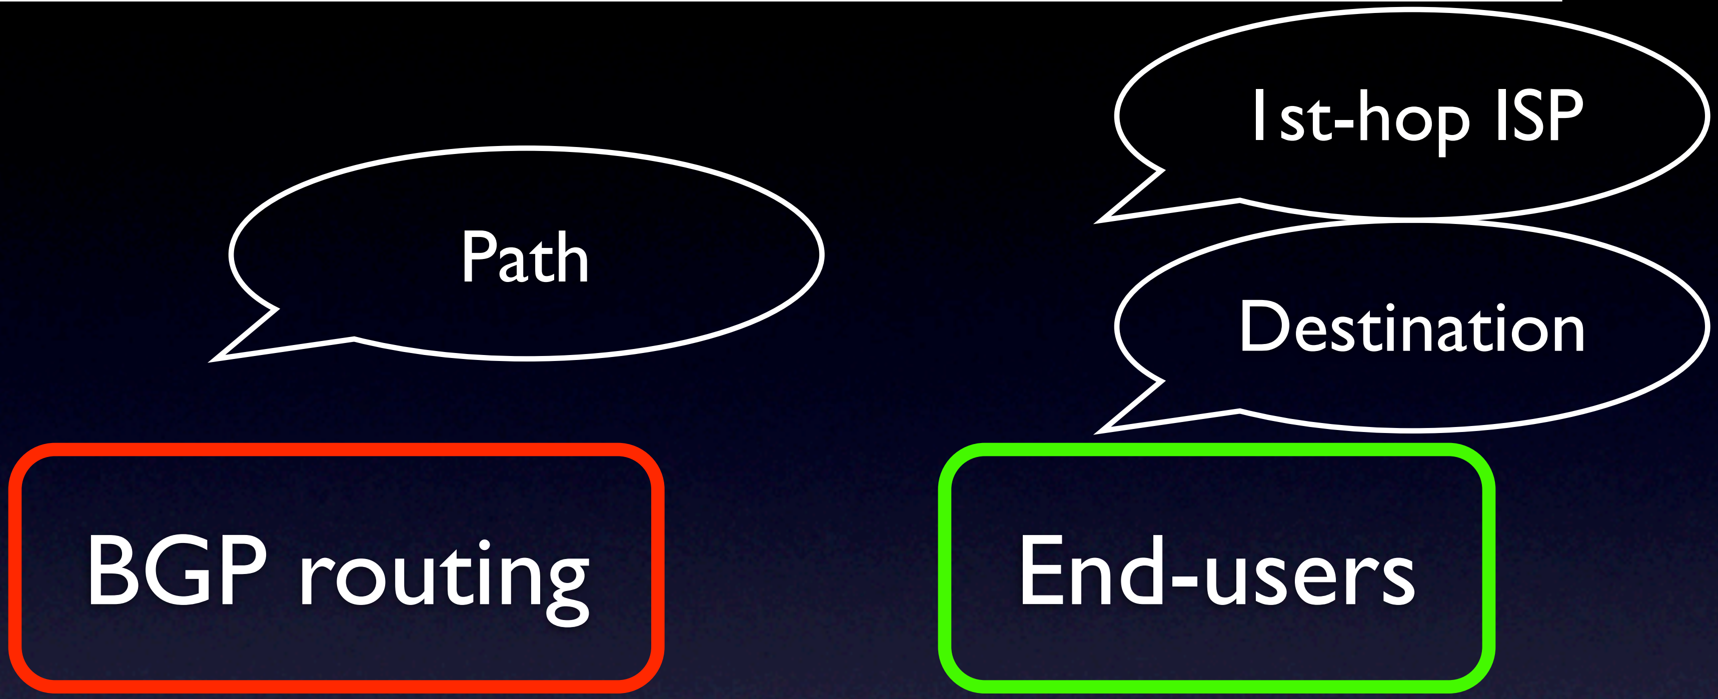

Path

1st-hop ISP

Destination

BGP routing

End-users

Internet Routing

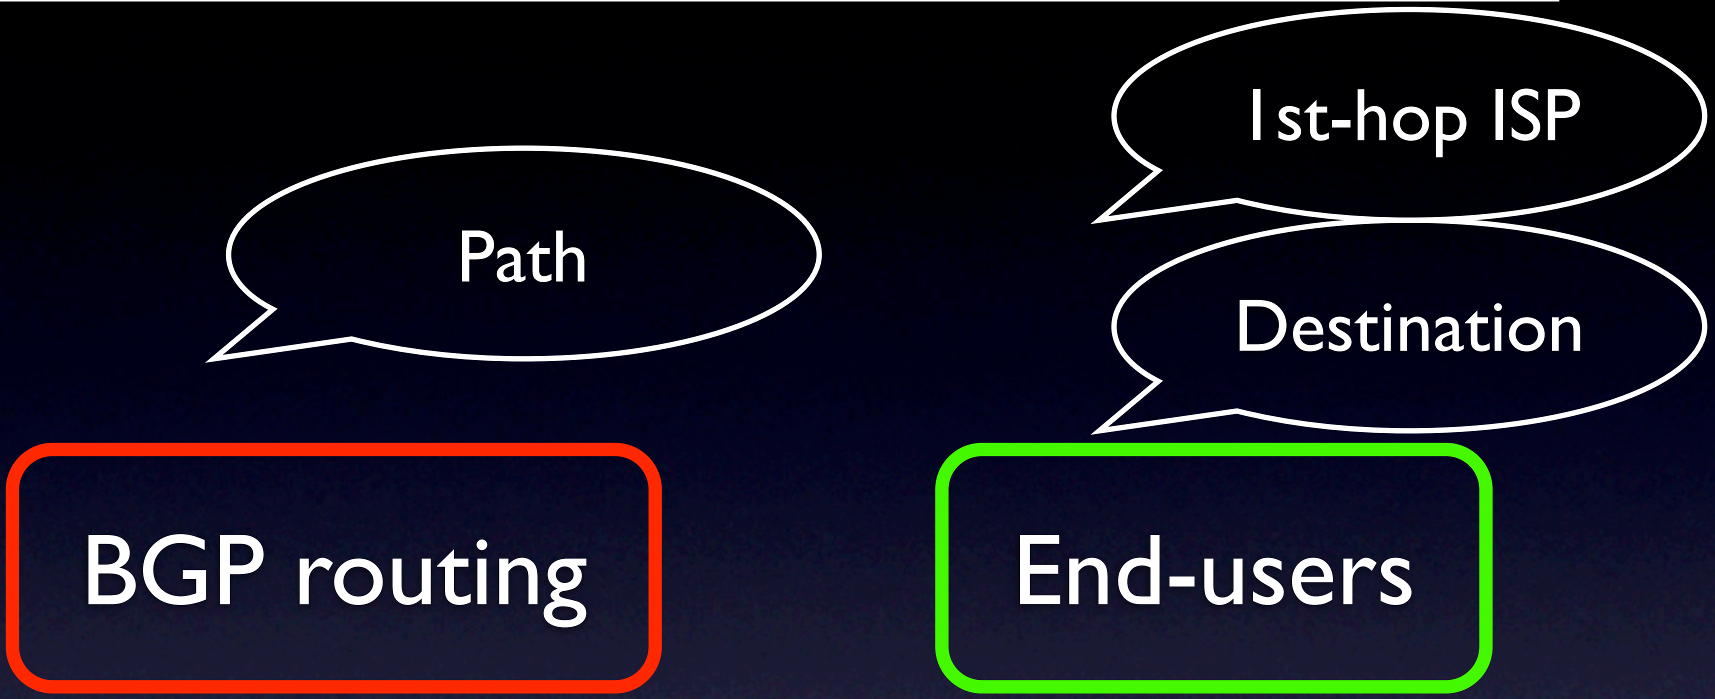

Path

1st-hop ISP

Destination

BGP routing

End-users

Correctness Connectivity Connectivity

Internet Routing

1st-hop ISP

Path

Destination

	BGP routing	End-users
Correctness	Connectivity	Connectivity
Optimization	-Least work -Business practices	-Lowest latency -Boycott ISPs

Internet Routing

Path

1st-hop ISP

Destination

	BGP routing	End-users
Correctness	Connectivity	Connectivity
Optimization	-Least work -Business practices	-Lowest latency -Boycott ISPs

Goal: Give choice to the users

Internet Routing Overlays

Source

S

Ds

Destination

Internet Routing Overlays

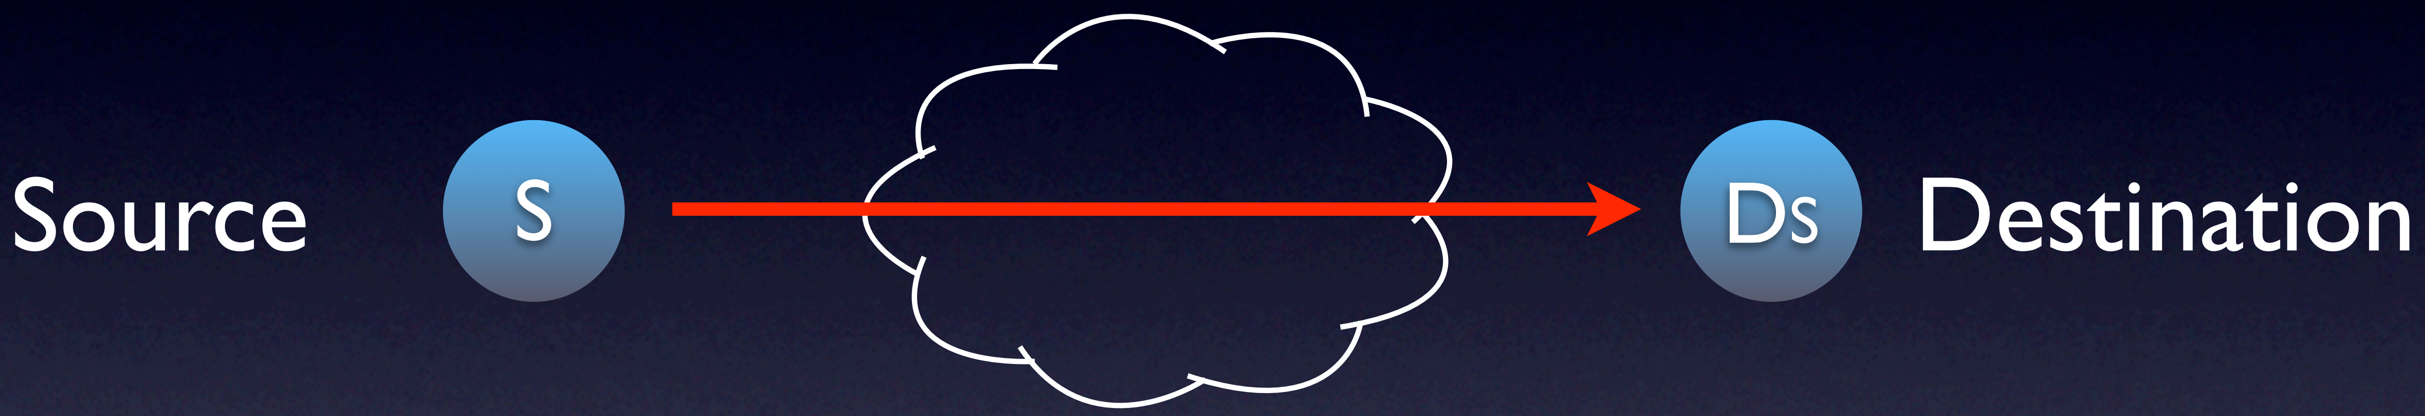

Internet Routing Overlays

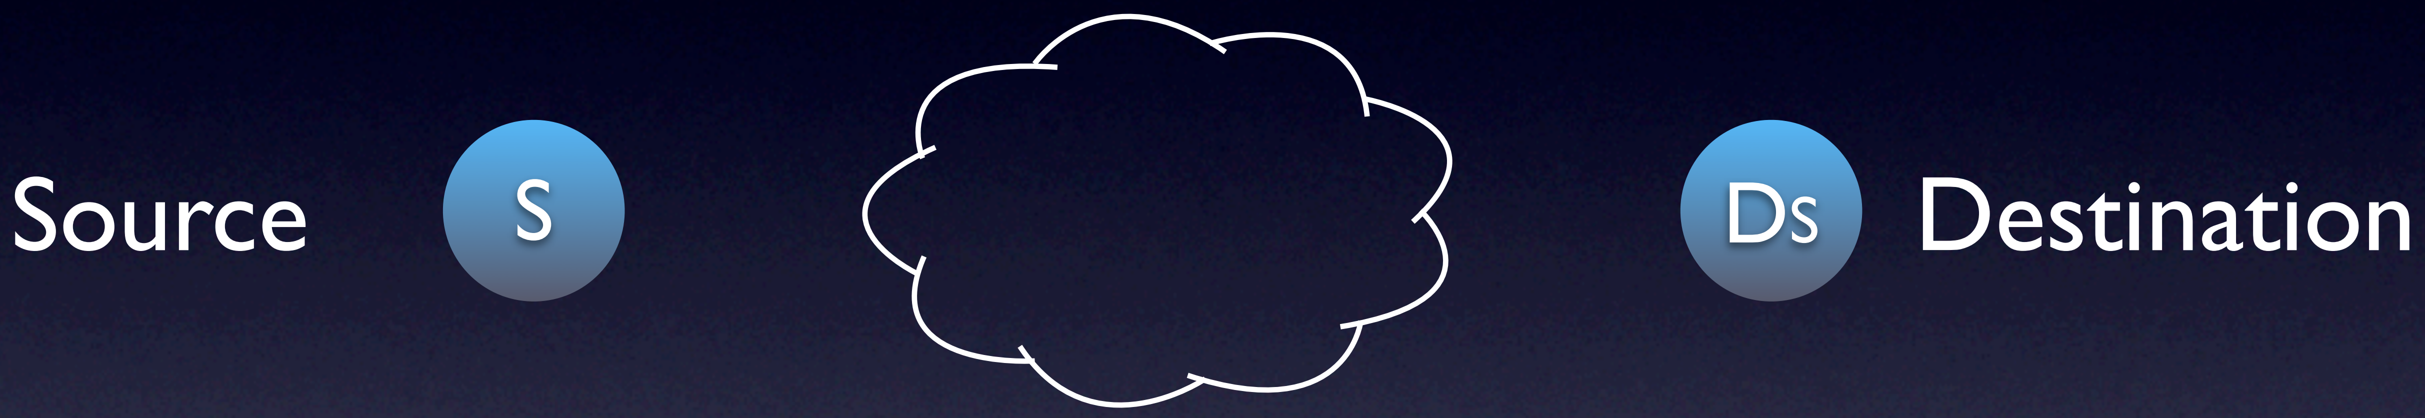

Source

S

Ds

Destination

Internet Routing Overlays

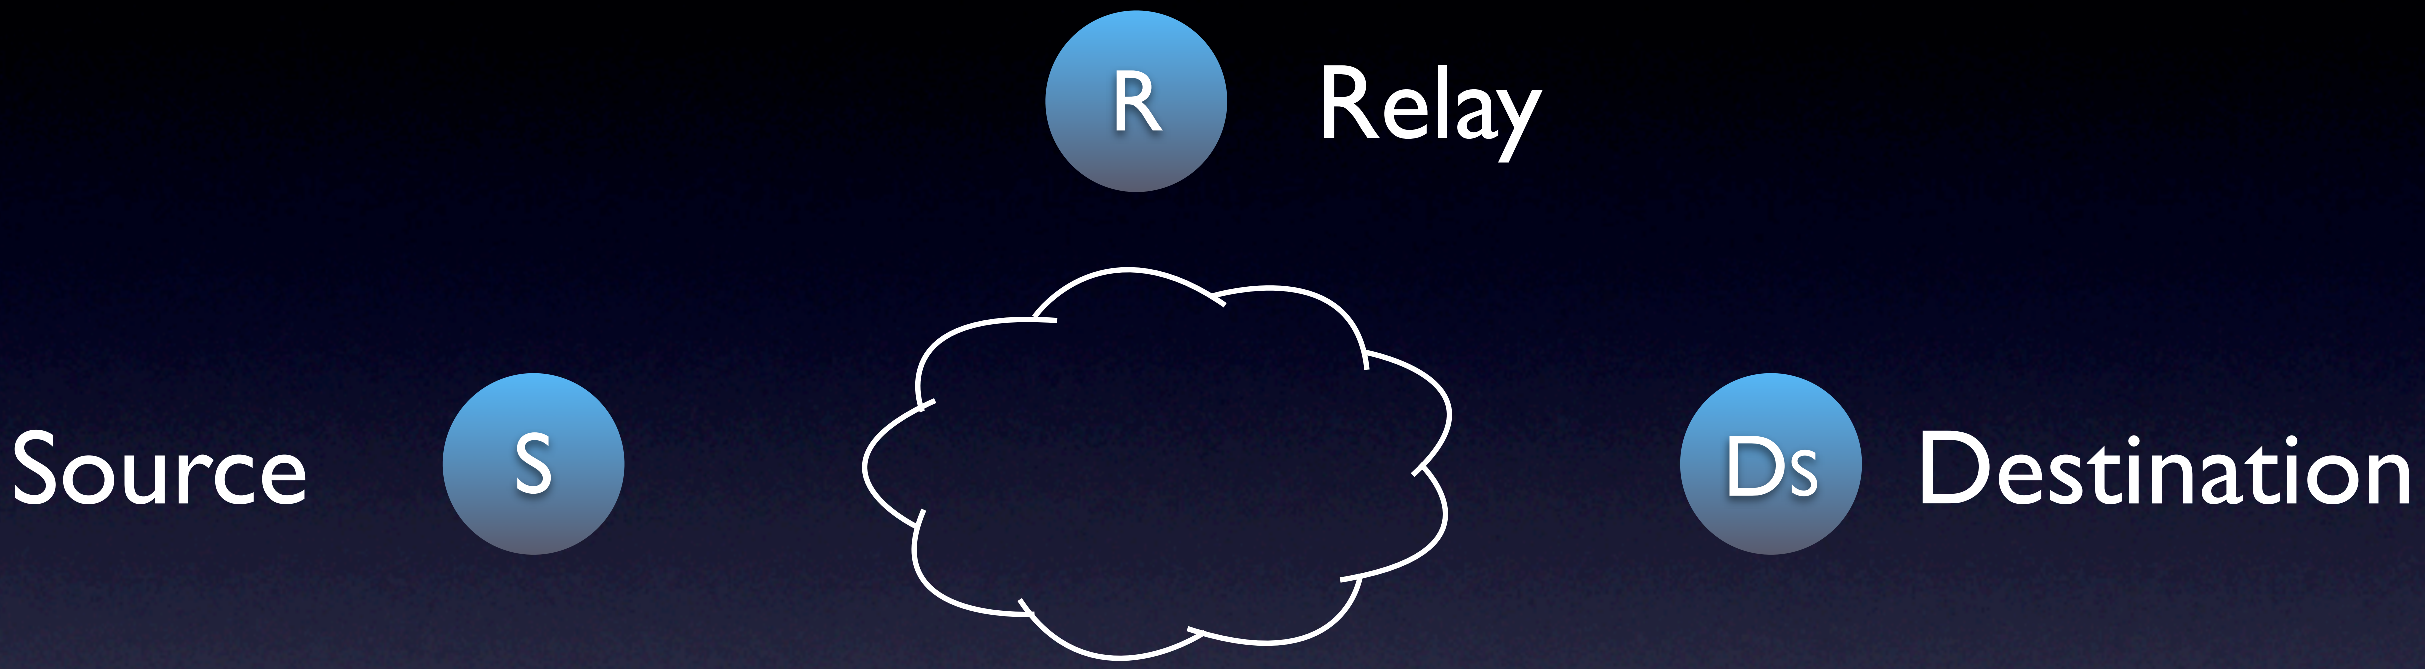

Internet Routing Overlays

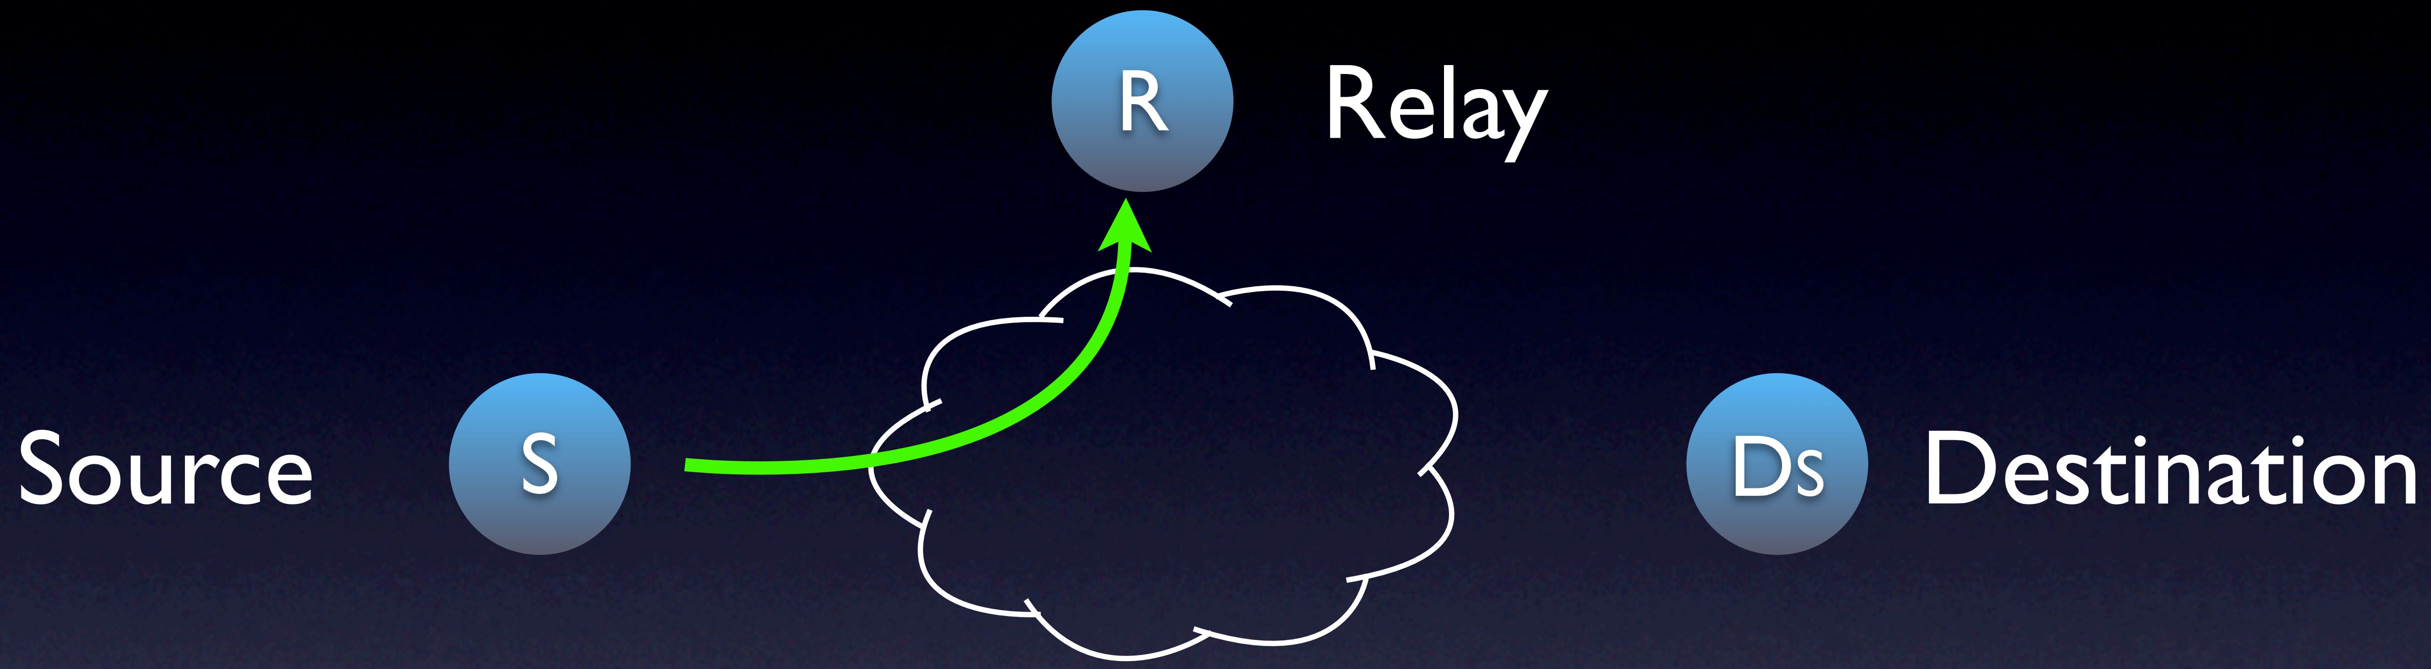

Internet Routing Overlays

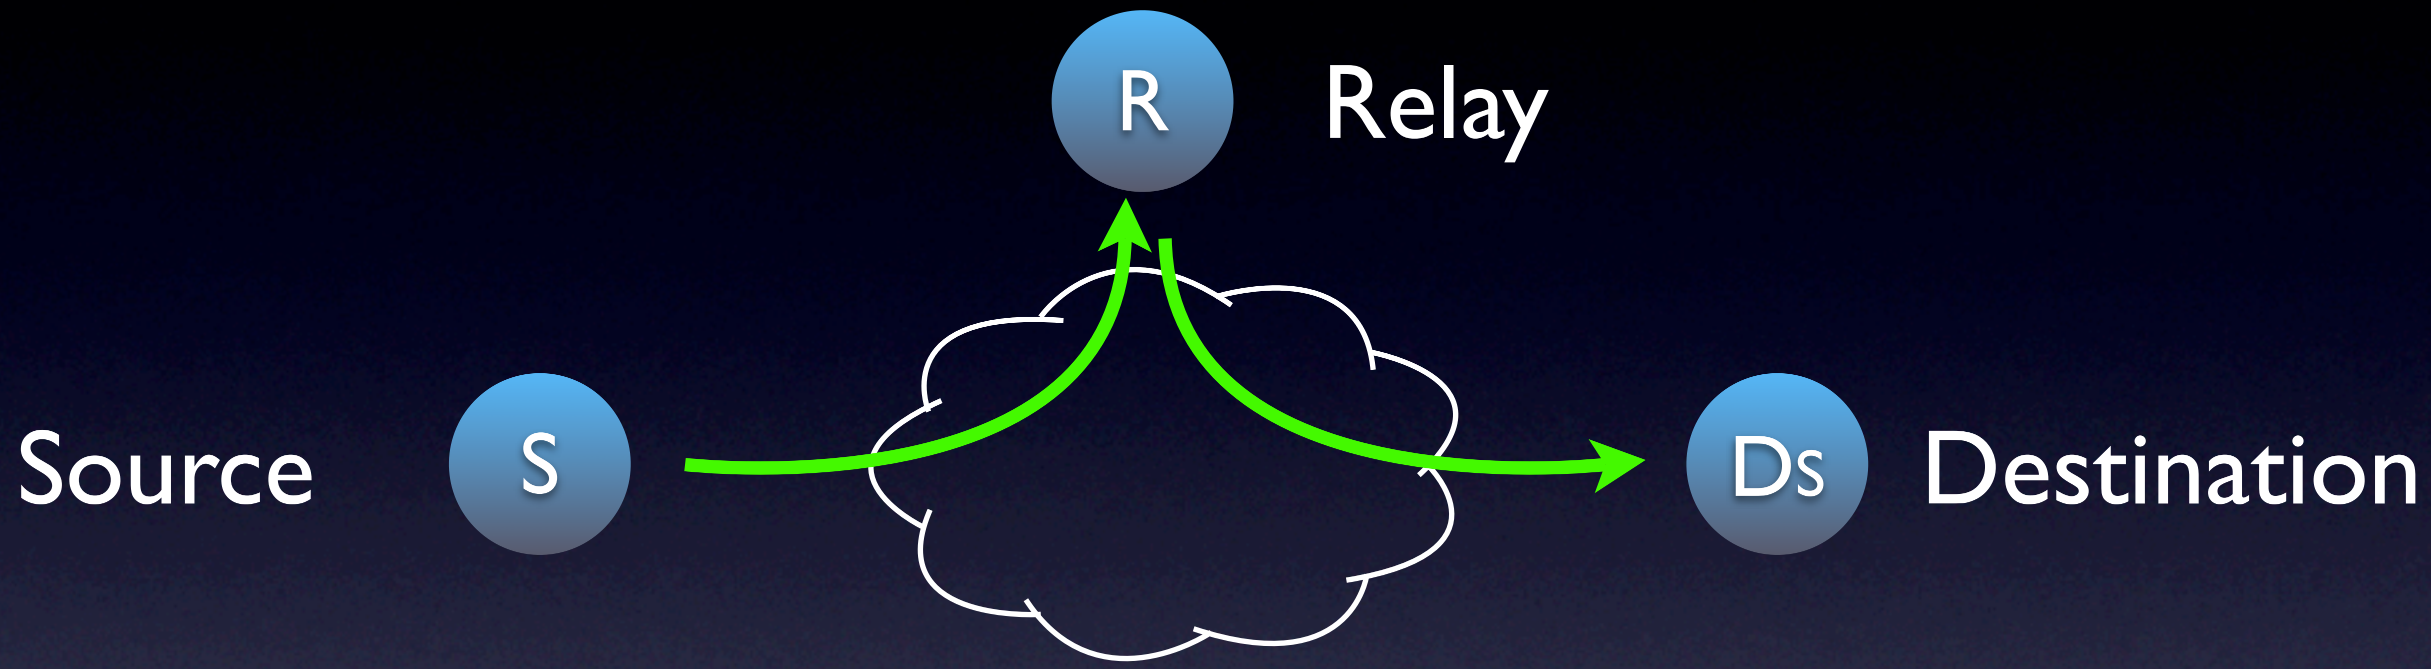

Relay

R

Source

S

Ds

Destination

Internet Routing Overlays

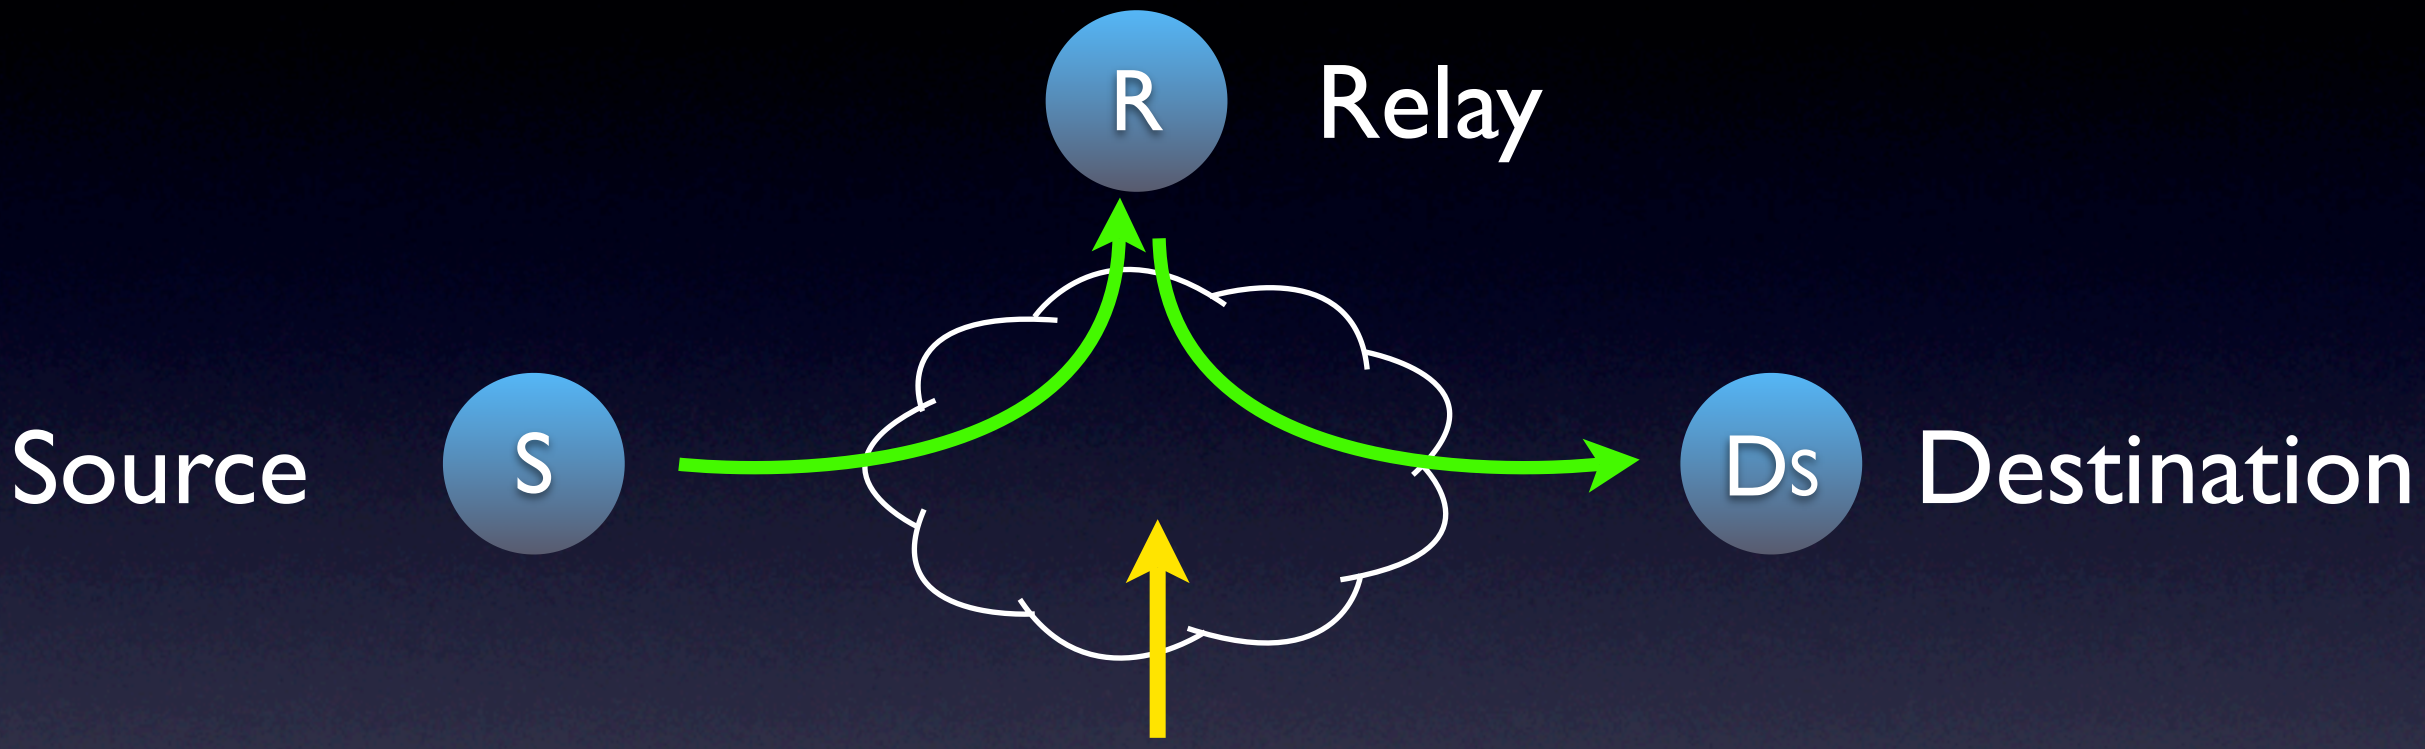

Provide end-user some choice over path

Internet Routing Overlays

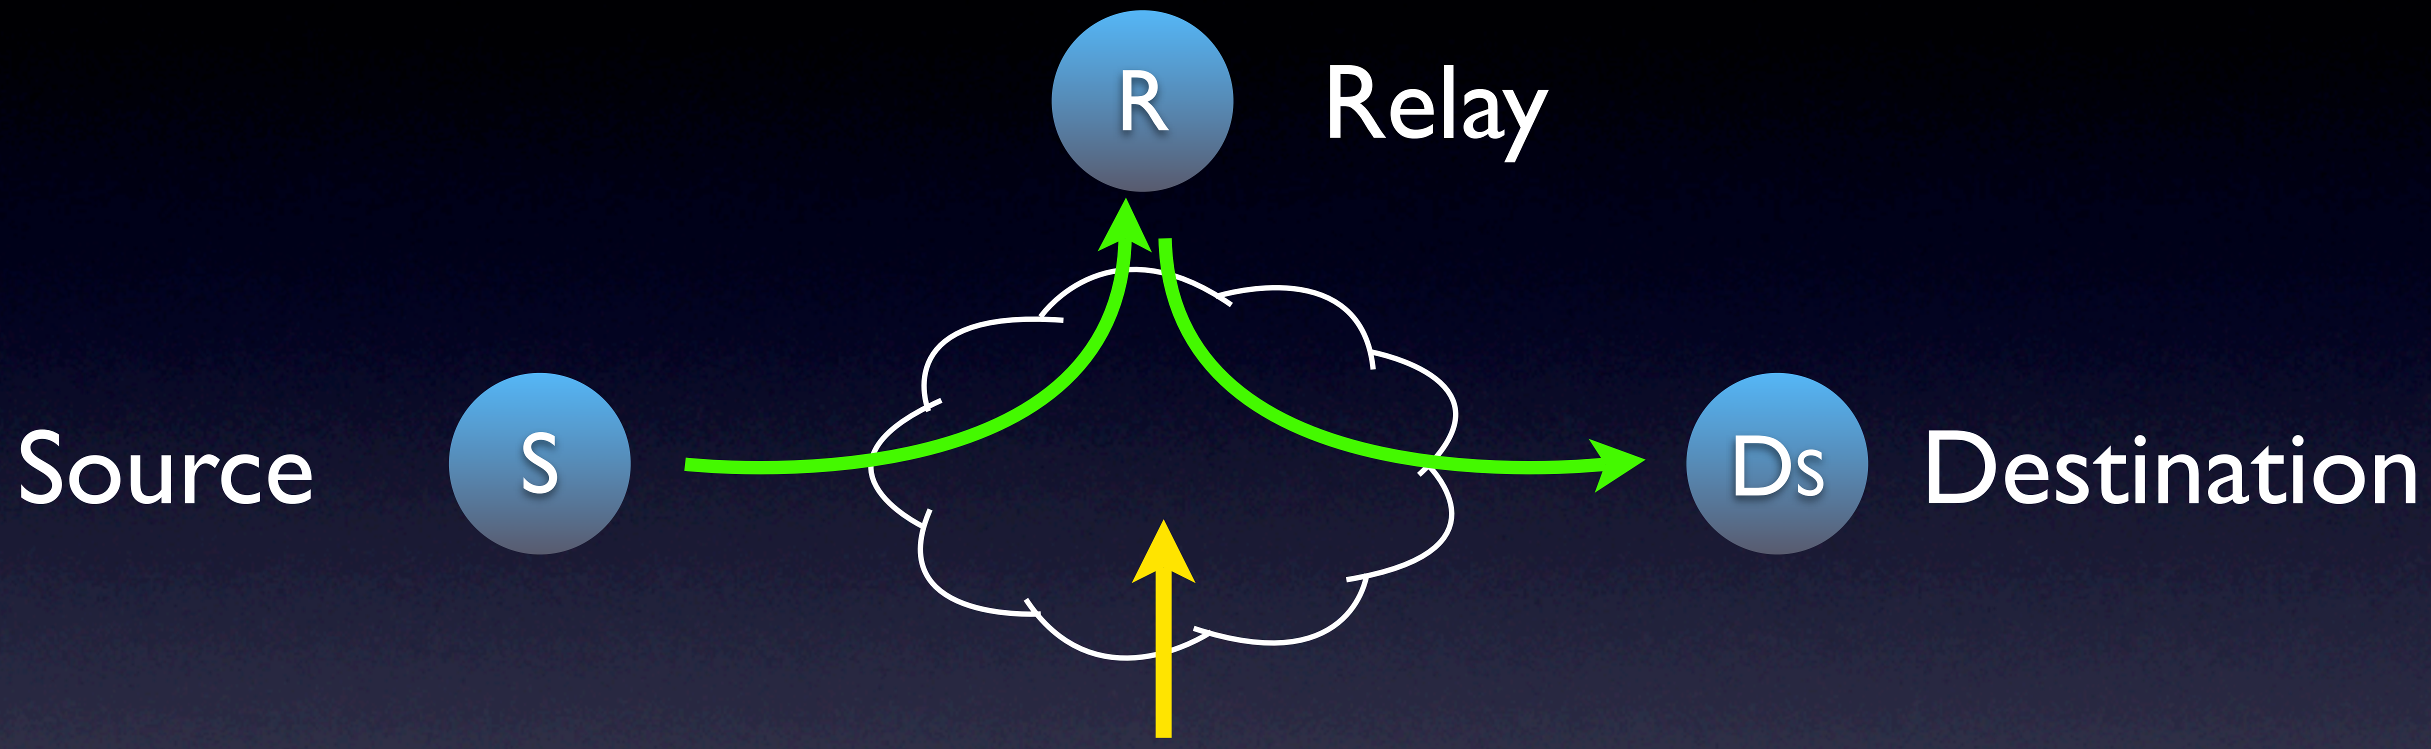

Provide end-user some choice over path

Route around failures quickly

SOSR, RON

Internet Routing Overlays

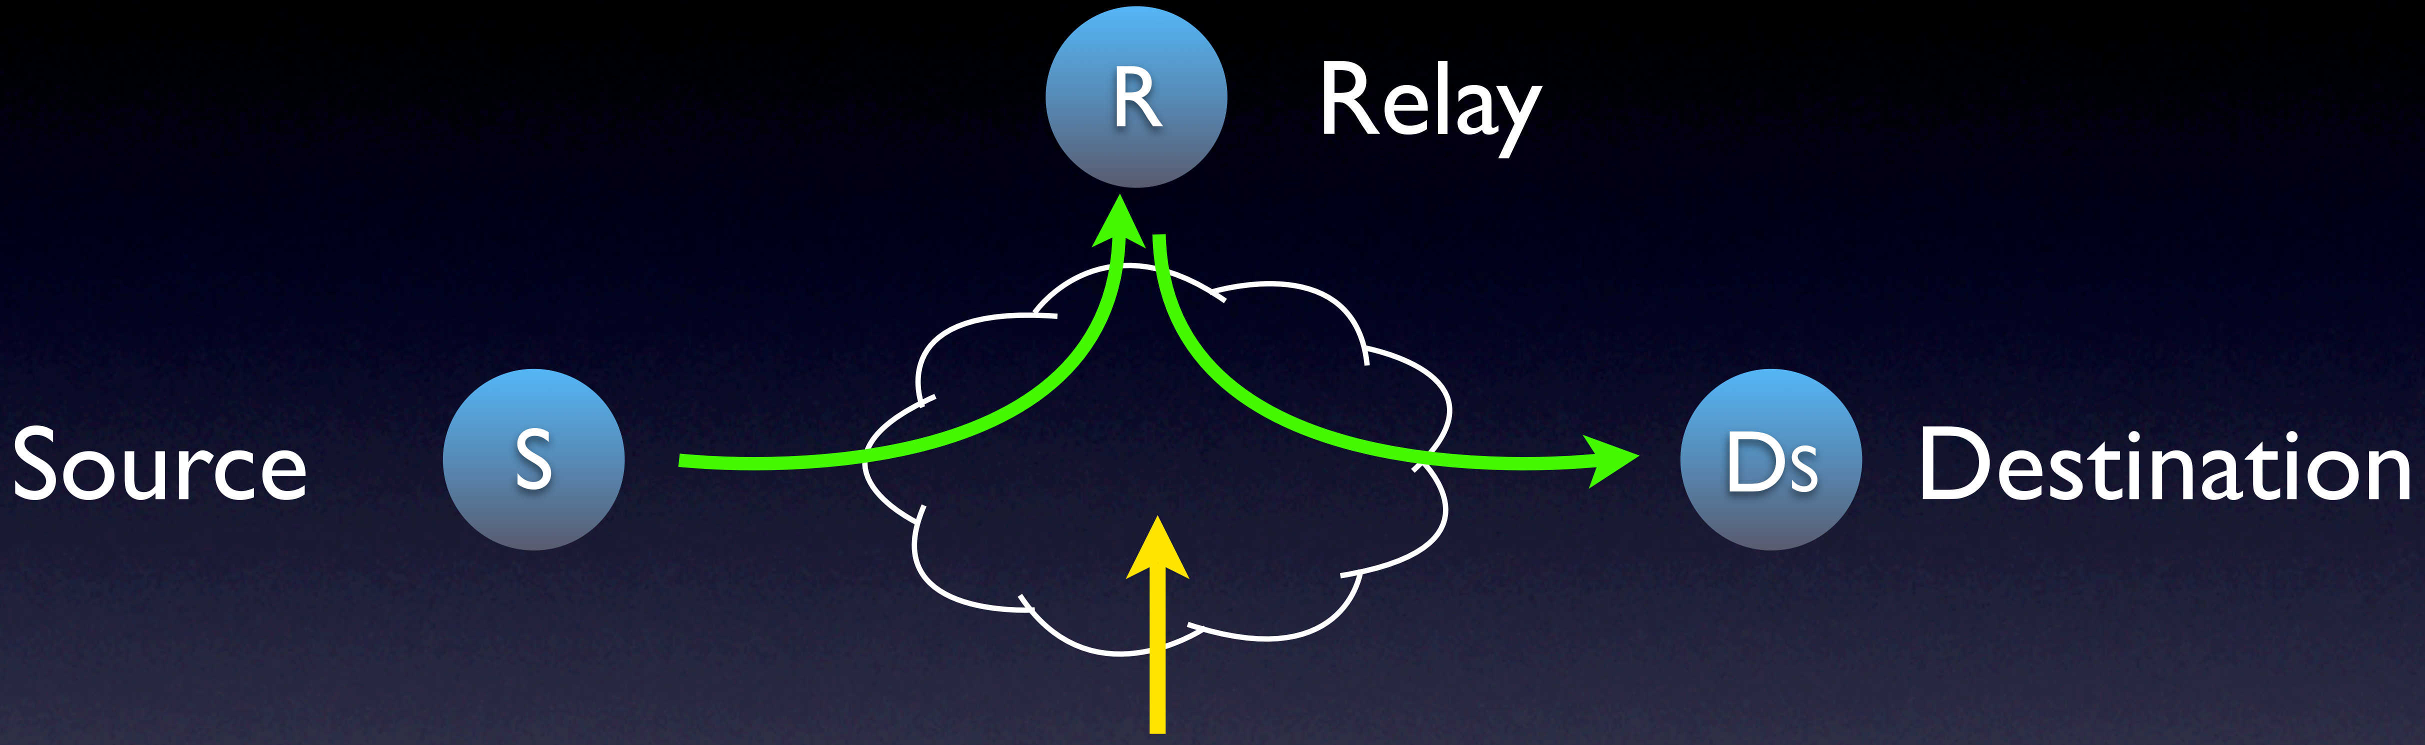

Provide end-user some choice over path

Route around failures quickly

SOSR, RON

Latency reduction

PeerWise

PeerWise

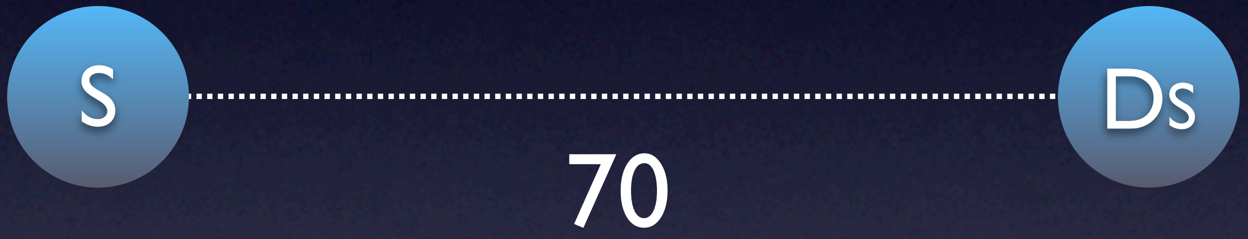

PeerWise

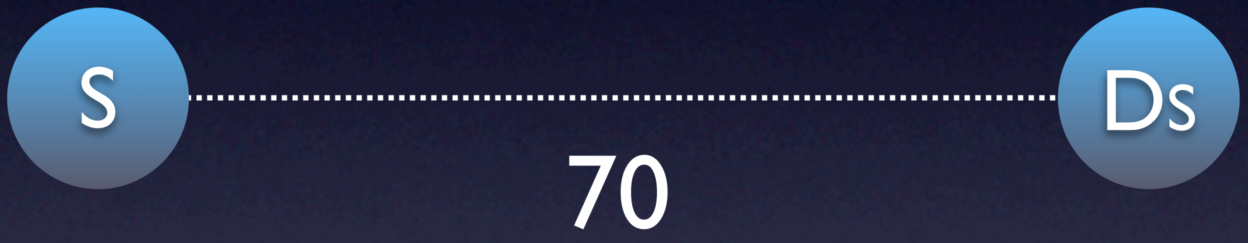

Triangle inequality violations

PeerWise

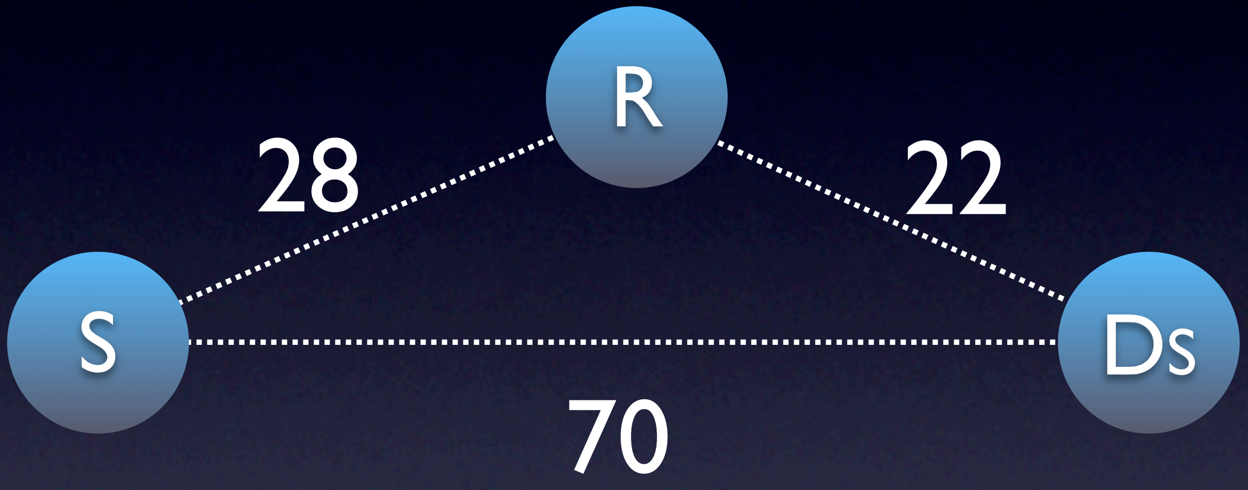

Triangle inequality violations

PeerWise

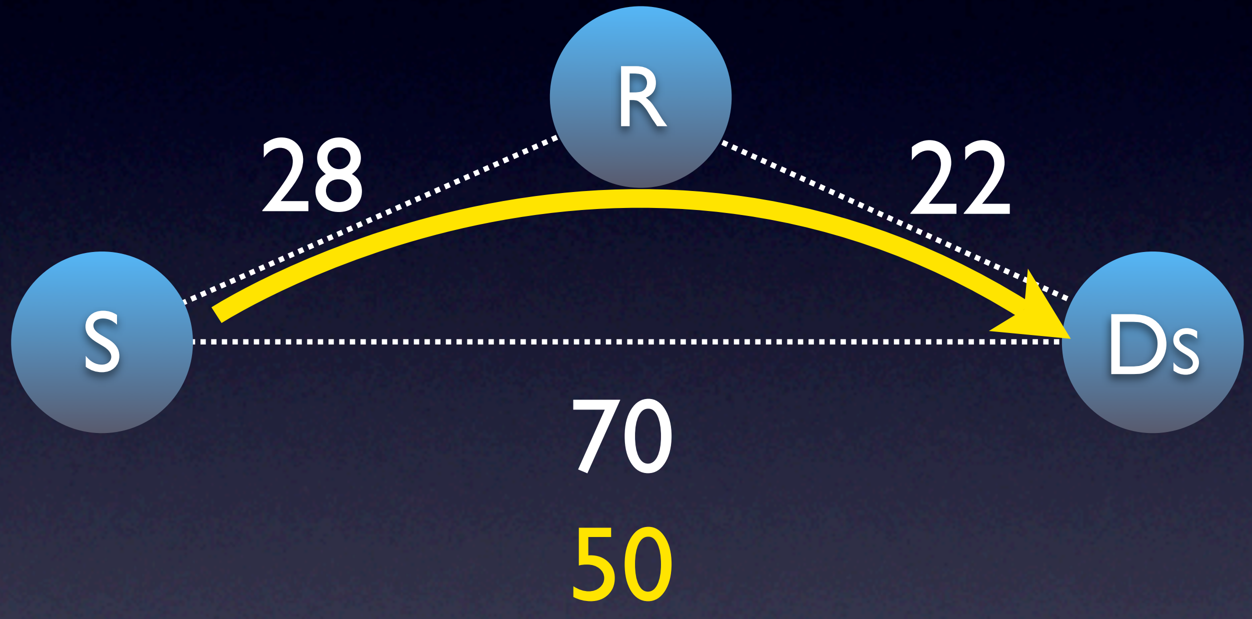

Triangle inequality violations

Peerings

PeerWise

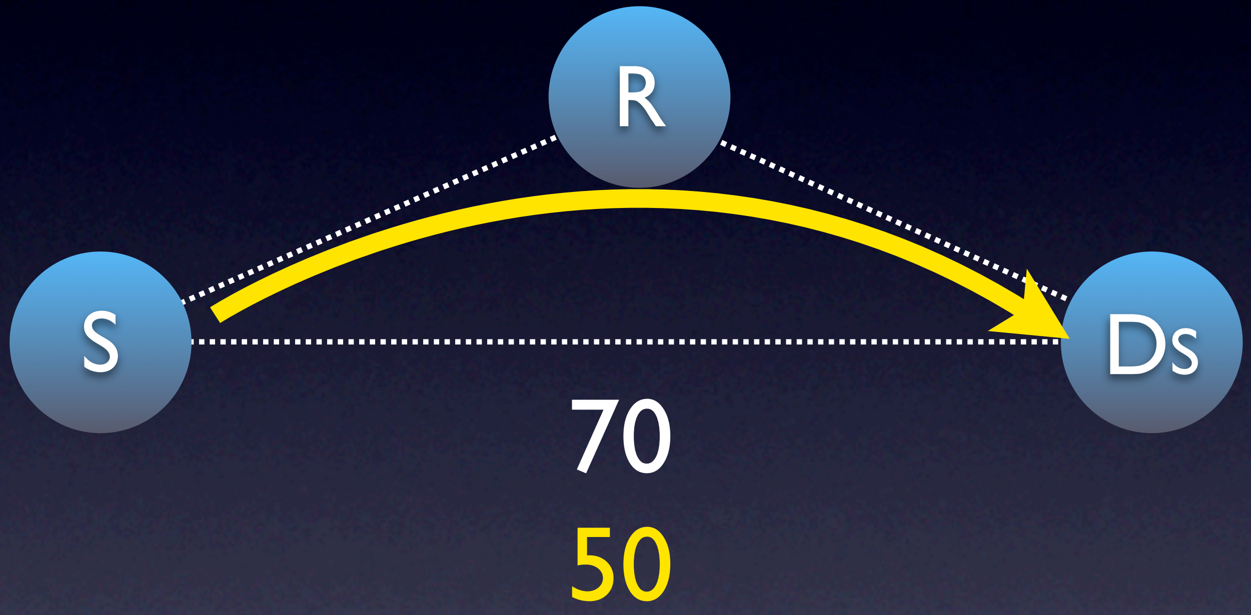

Triangle inequality violations

Peerings

PeerWise

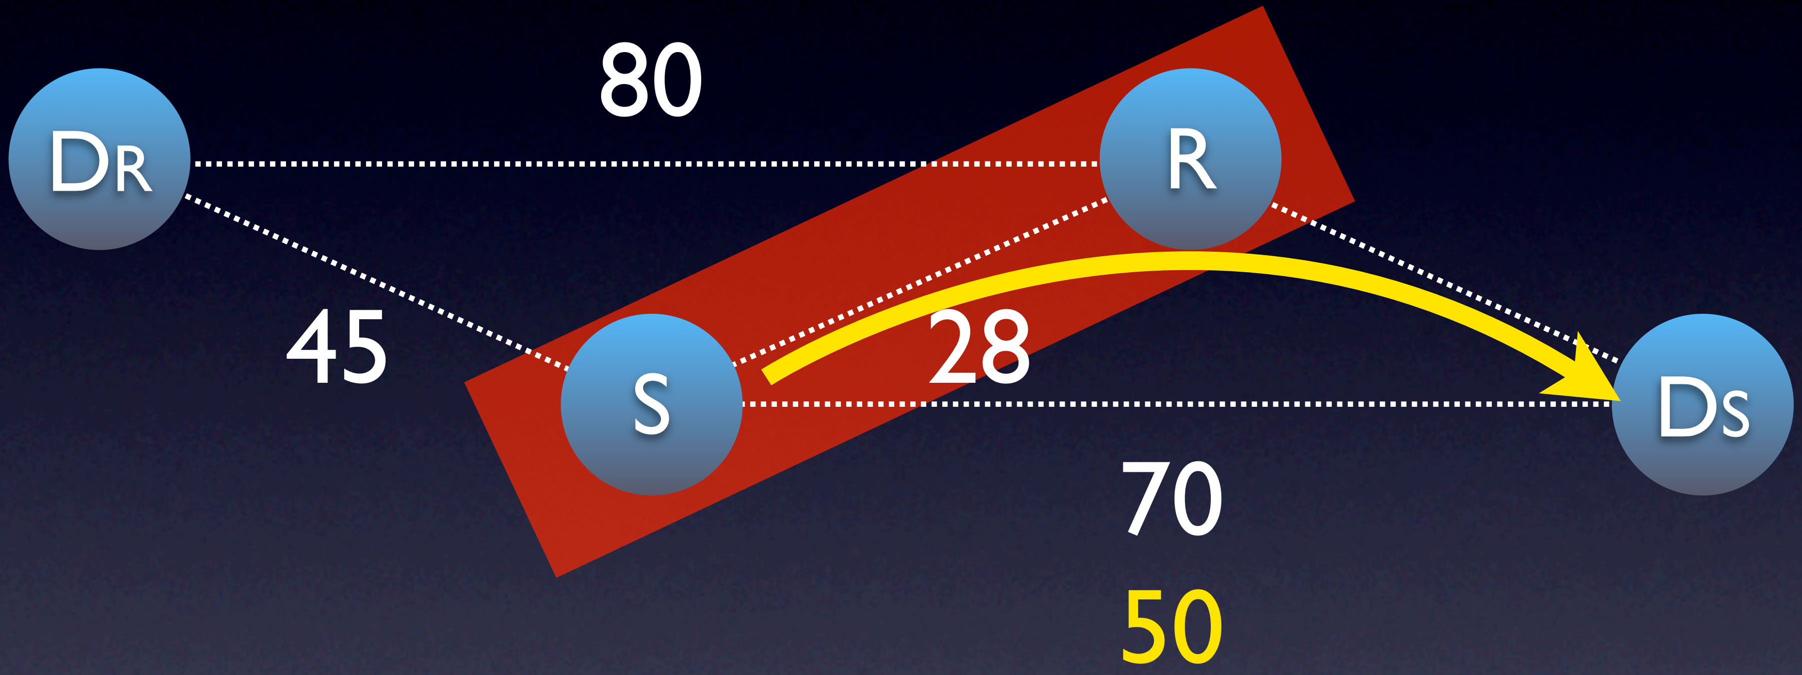

Triangle inequality violations

Peerings

PeerWise

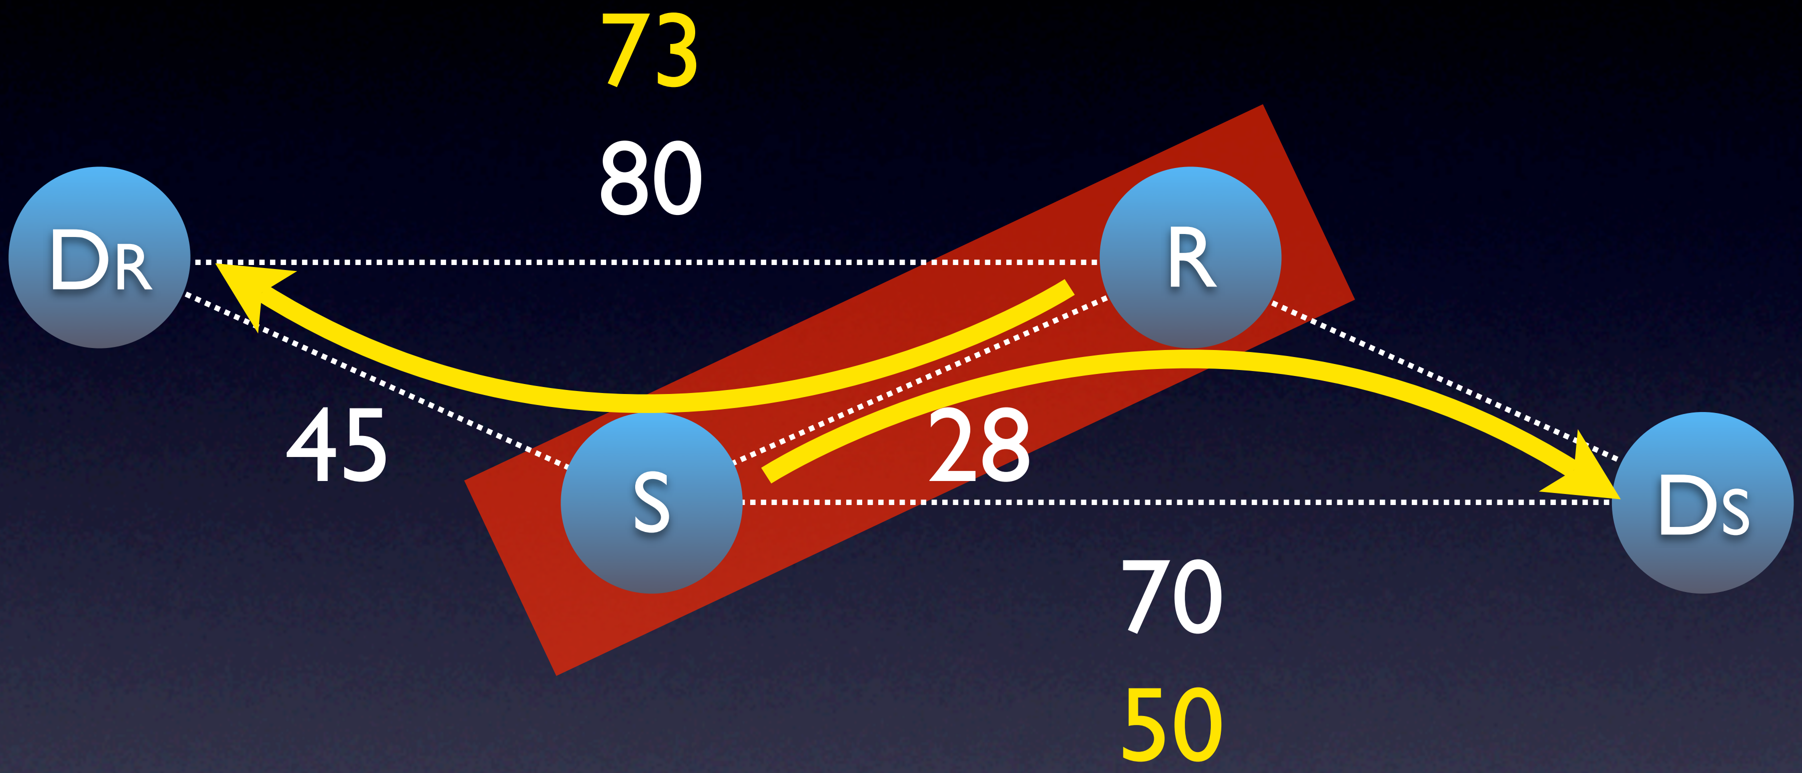

Triangle inequality violations

Peerings

PeerWise

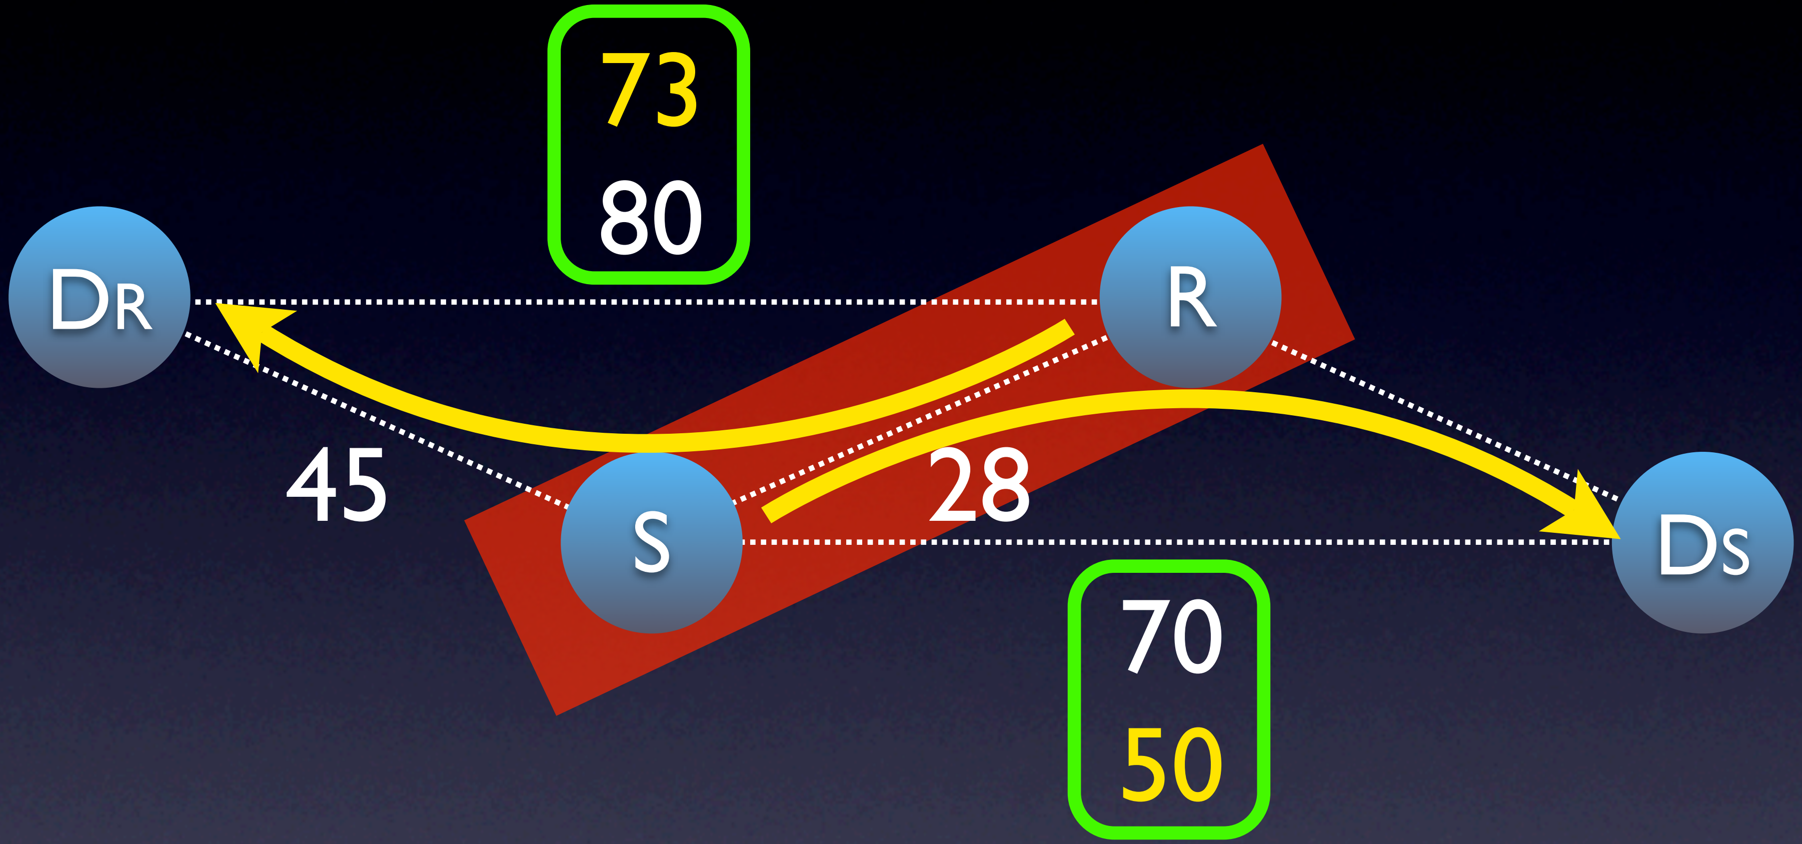

Triangle inequality violations

Peerings

Mutual benefit

PeerWise

Triangle inequality violations

Peerings

Mutual benefit

PeerWise

Triangle inequality violations

Peerings

Mutual benefit

PeerWise

Triangle inequality violations

Peerings

Mutual benefit

PeerWise

Triangle inequality violations

Found via errors in network embedding systems (Vivaldi)

Peerings

PeerWise reduces latency

Mutual benefit

Mutual interest is common

PeerWise

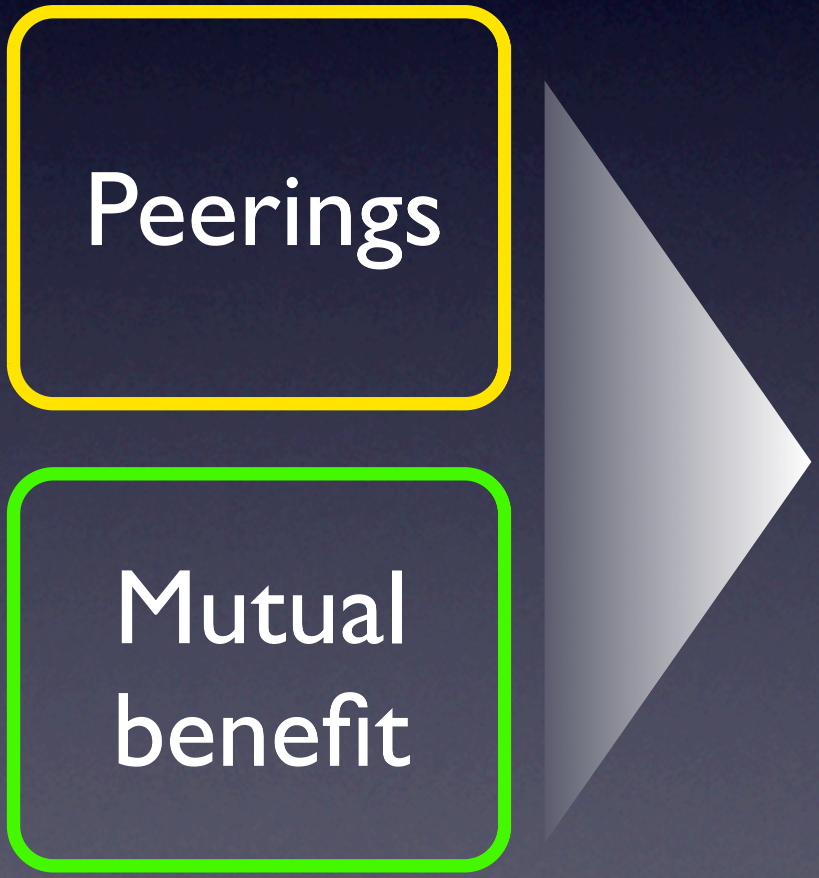

Peerings

Mutual benefit

How do we prolong interaction?

Promoting cooperation

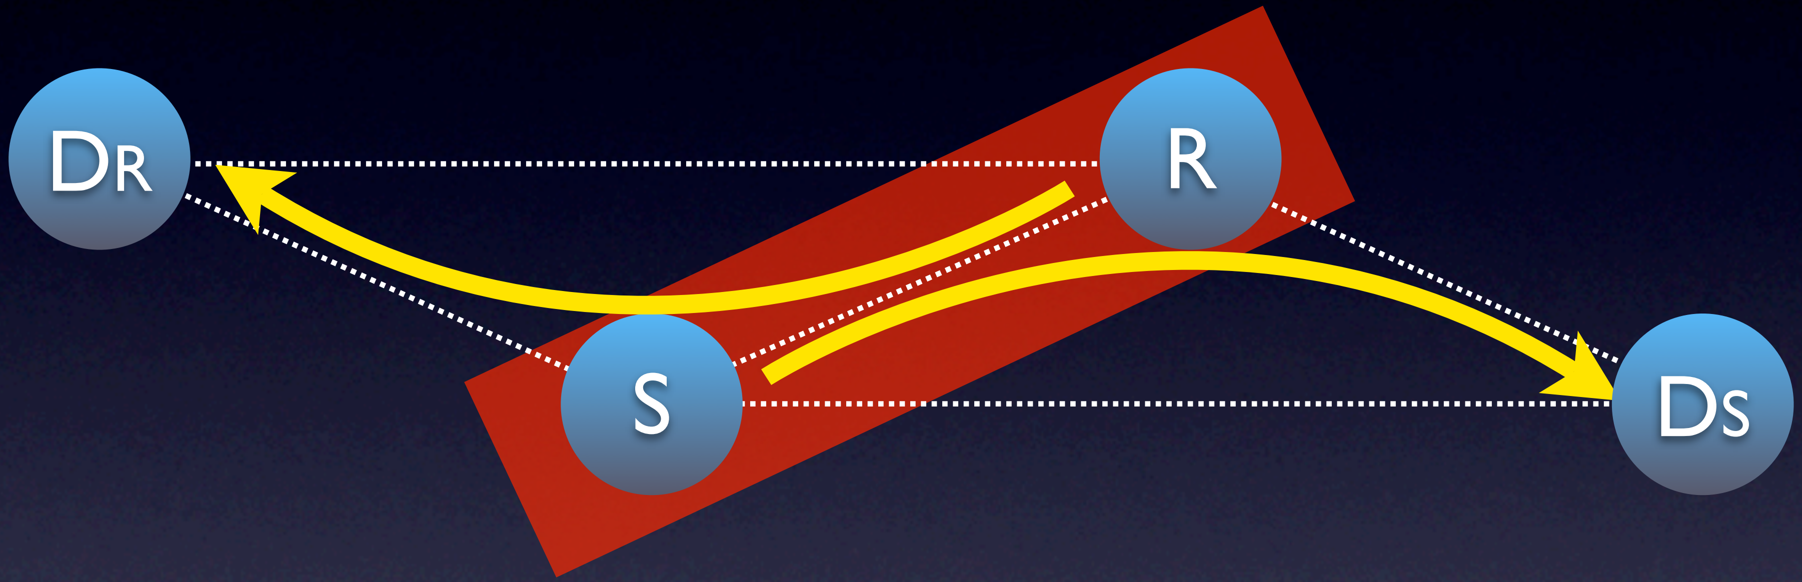

Promoting cooperation

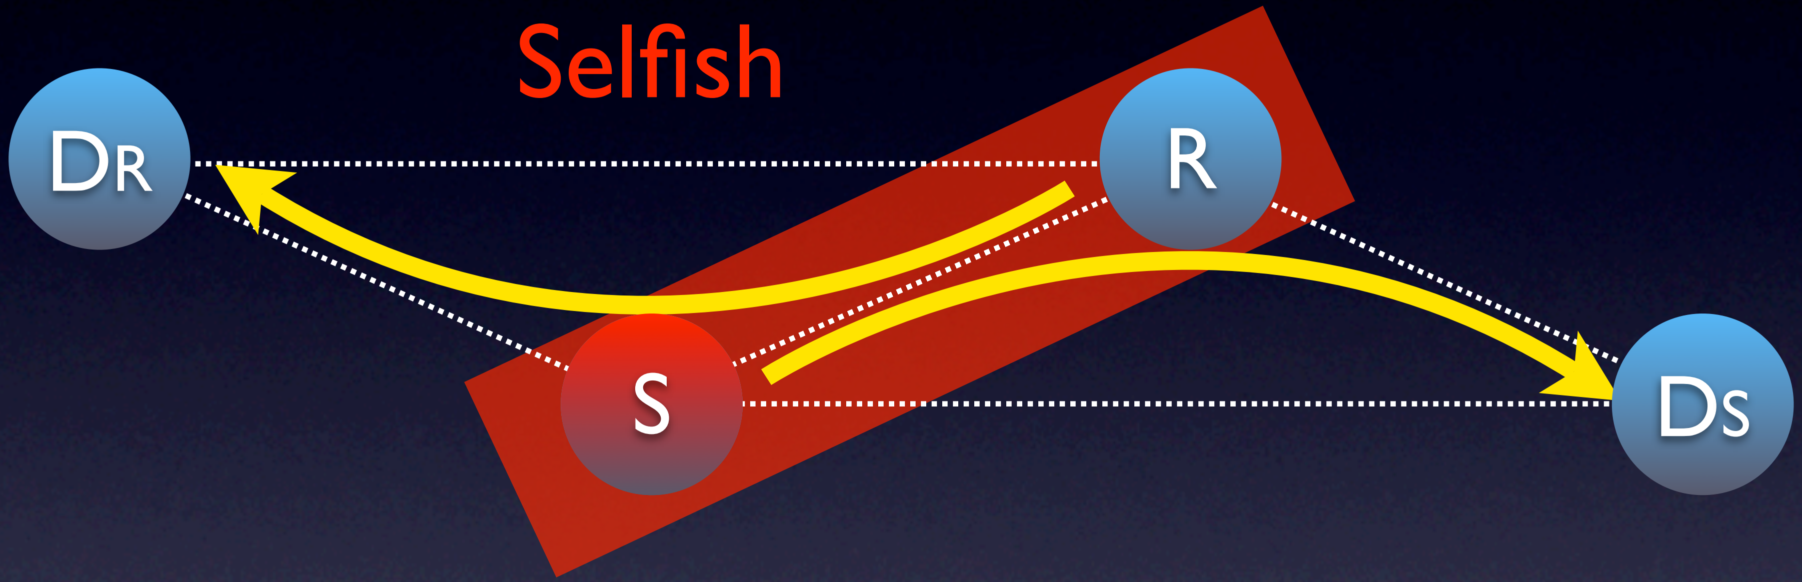

Promoting cooperation

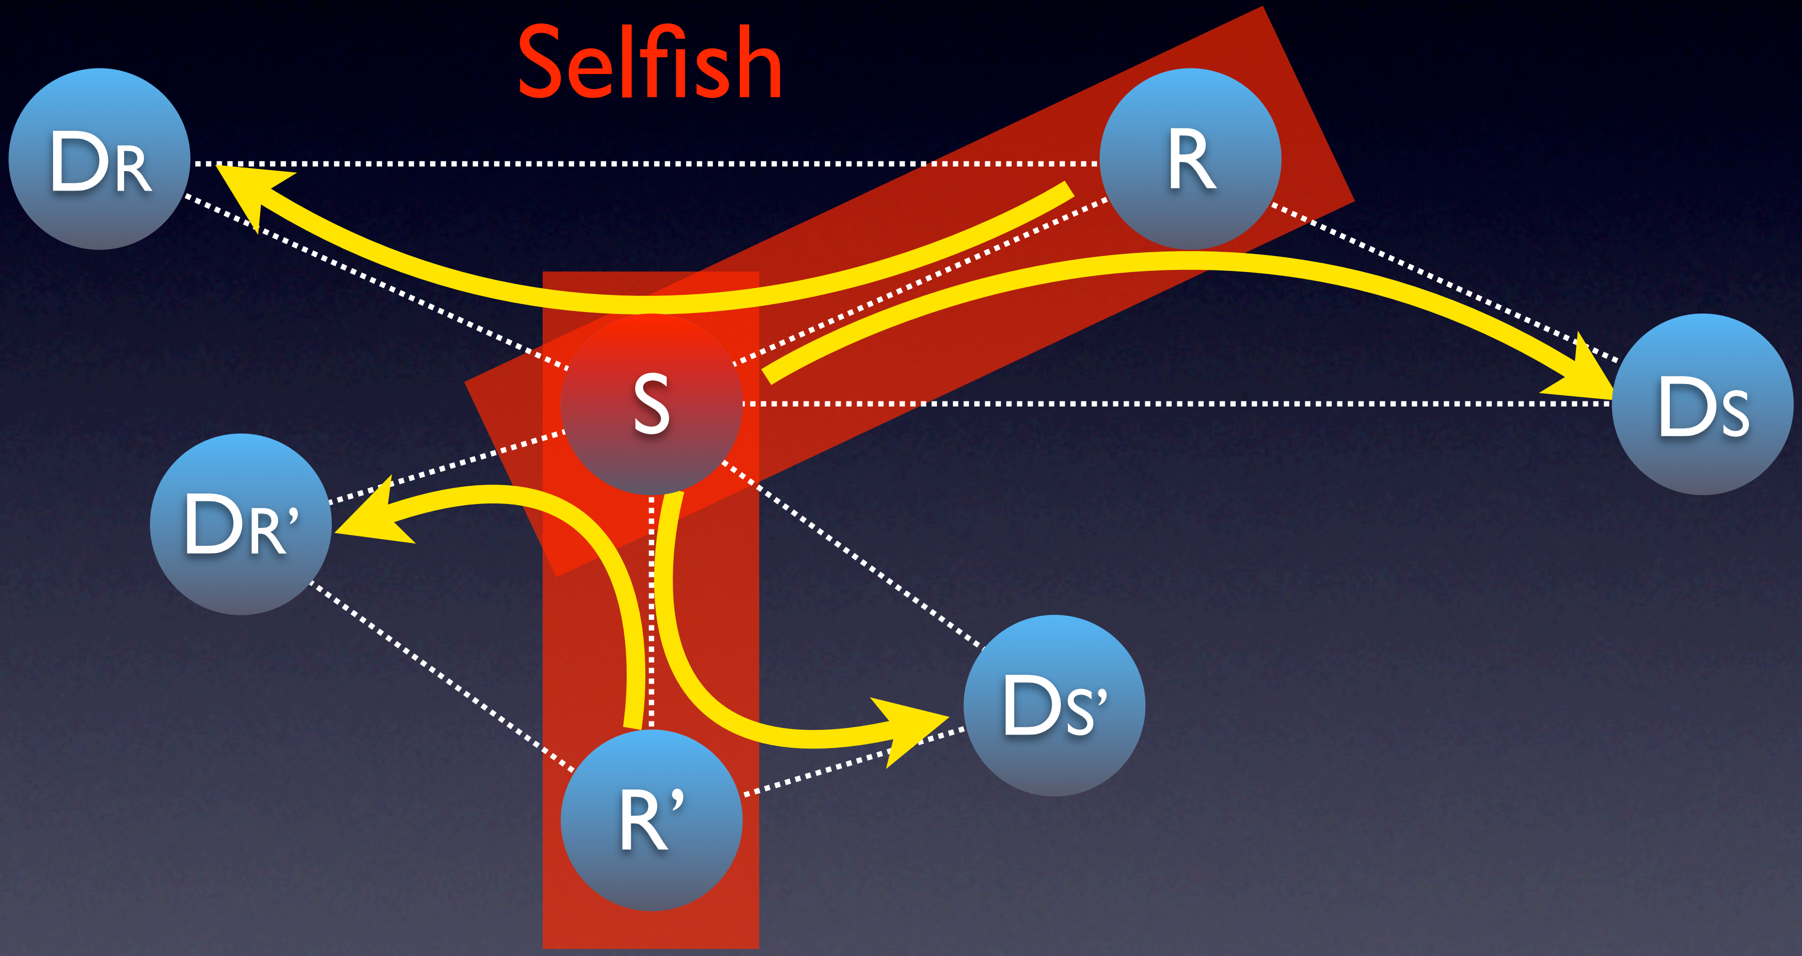

Promoting cooperation

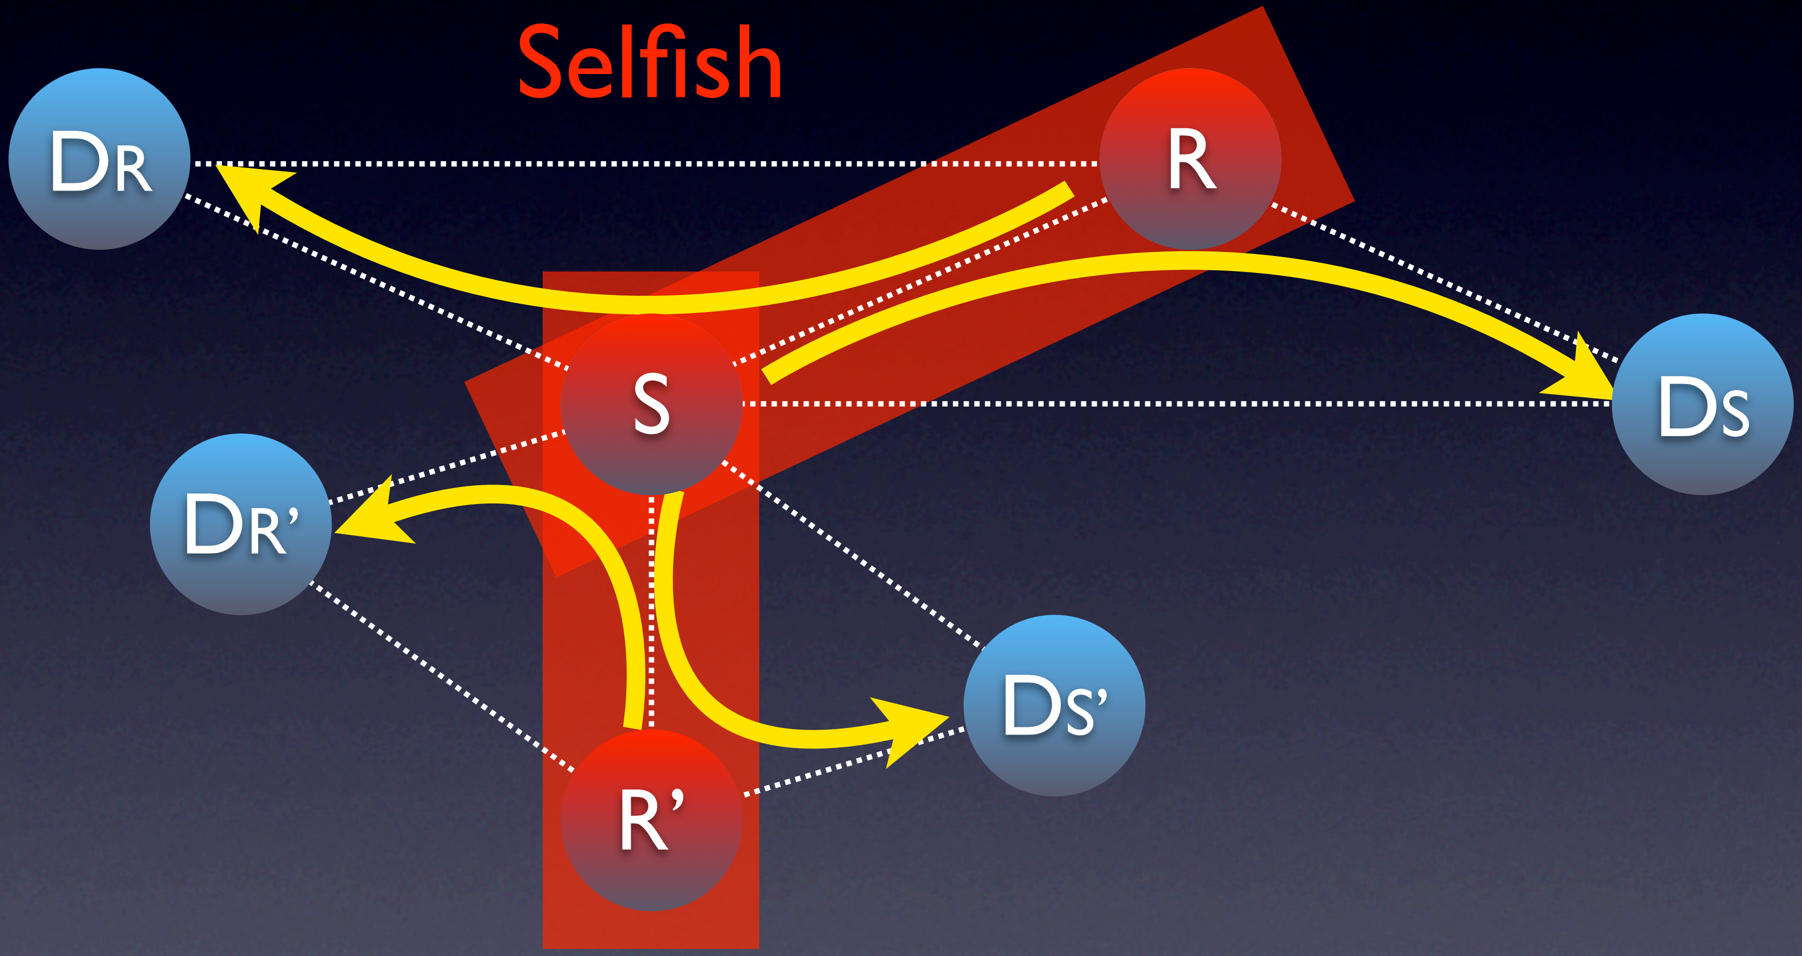

Goals of a selfish PeerWise peer

- Maximize benefit
  - Reduce latency as much as possible
- Minimize cost
  - Conserve bandwidth for your own flows
- Selfishly select neighbors
  - Cannot relay for everyone

Goals of a selfish PeerWise peer

- Maximize benefit
  - Reduce latency as much as possible

- Minimize cost
  - Conserve bandwidth for your own flows

- Selfishly select neighbors
  - Cannot relay for everyone

Goals of a selfish PeerWise peer

- Maximize benefit
  - Reduce latency as much as possible
- Minimize cost
  - Conserve bandwidth for your own flows
- Selfishly select neighbors
  - Cannot relay for everyone

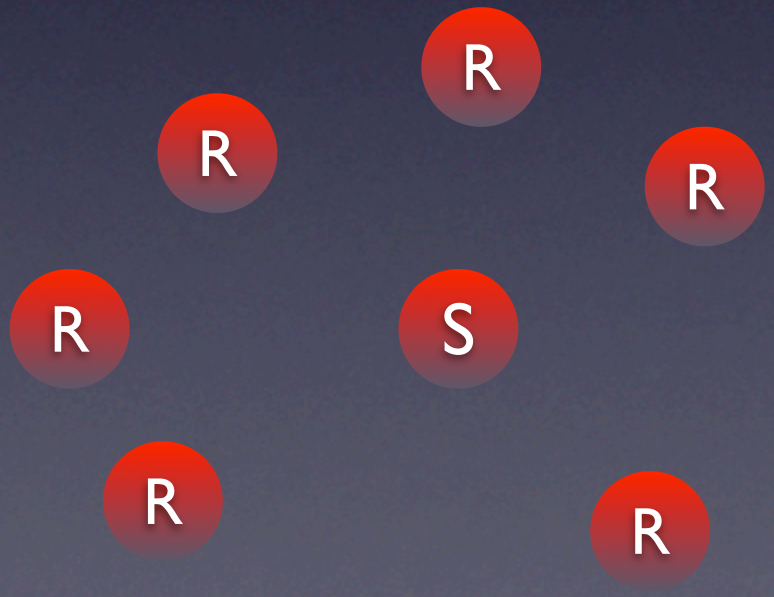

Goals of a selfish PeerWise peer

- Maximize benefit
  - Reduce latency as much as possible

- Minimize cost
  - Conserve bandwidth for your own flows

- Selfishly select neighbors
  - Cannot relay for everyone

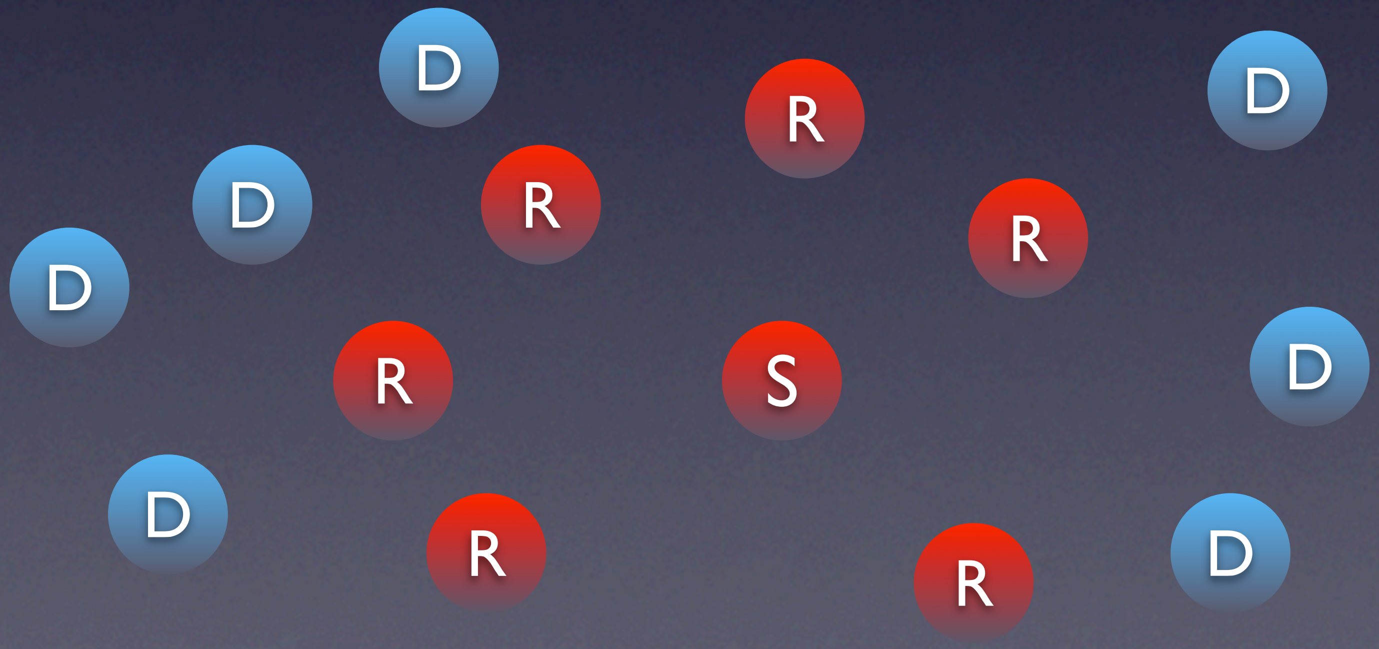

Goals of a selfish PeerWise peer

- Maximize benefit
  - Reduce latency as much as possible
- Minimize cost
  - Conserve bandwidth for your own flows
- Selfishly select neighbors
  - Cannot relay for everyone

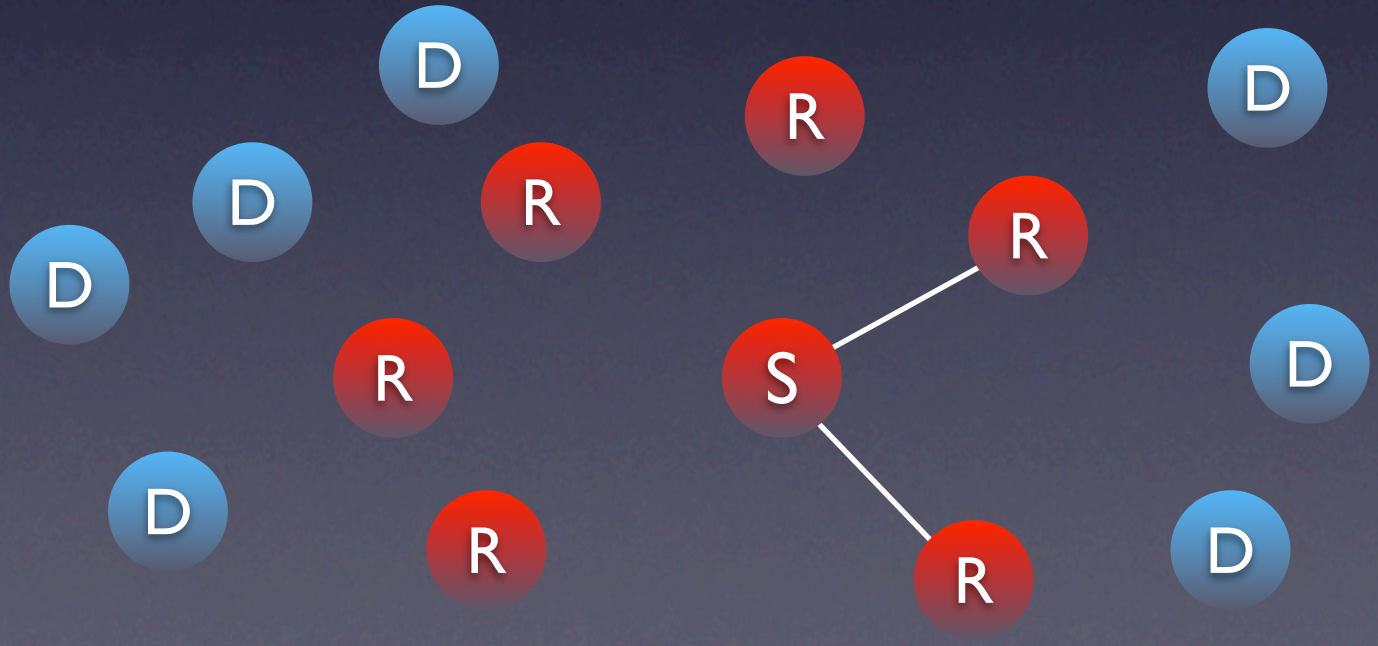

Goals of a selfish PeerWise peer

- Maximize benefit
  - Reduce latency as much as possible
- Minimize cost
  - Conserve bandwidth for your own flows
- Selfishly select neighbors
  - Cannot relay for everyone

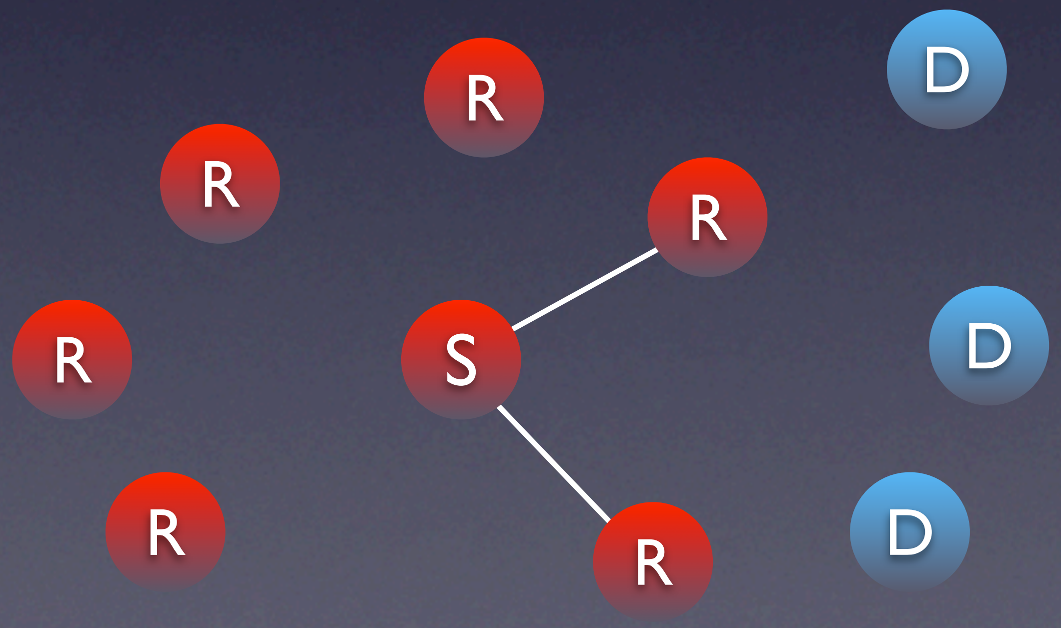

Haven’t I heard this before?

- Peers selfishly want to get without giving
- System built on symmetric interest
- Neighbor selection

Haven’t I heard this before?

- Peers selfishly want to get without giving
- System built on symmetric interest
- Neighbor selection

Not quite

Contributions

- The selfish routing overlay problem
- Mechanisms to achieve long-lived peerings
- Simulation study

Contributions

- The selfish routing overlay problem
- Mechanisms to achieve long-lived peerings
- Simulation study

Peers can opt out

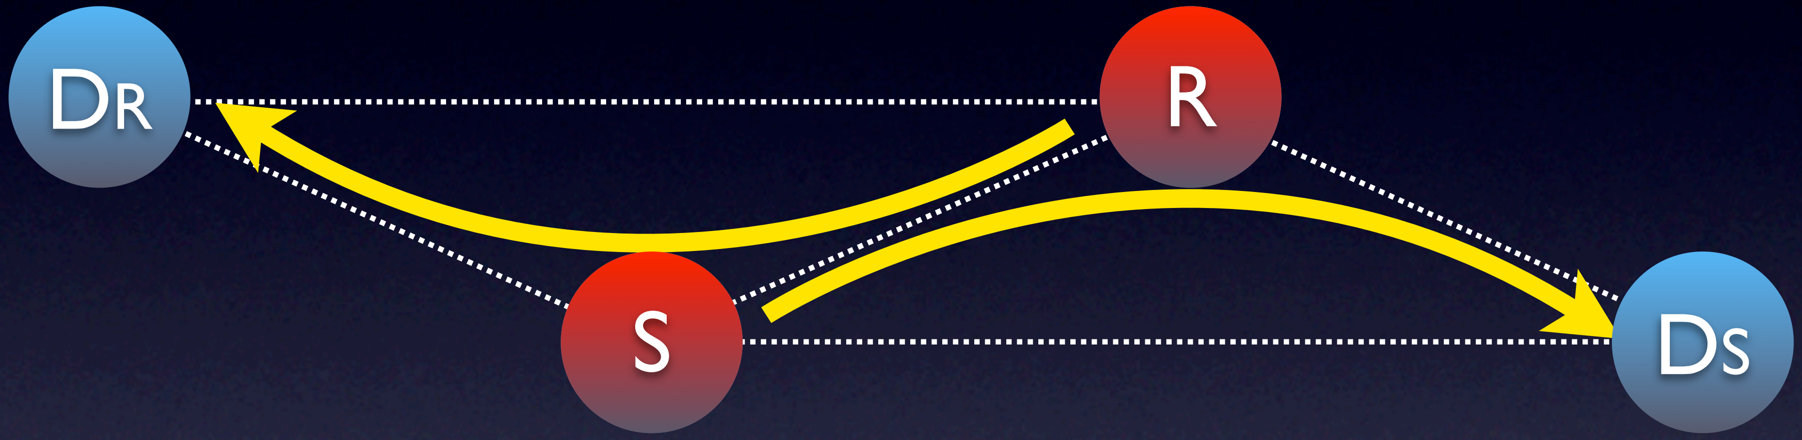

Peers can opt out

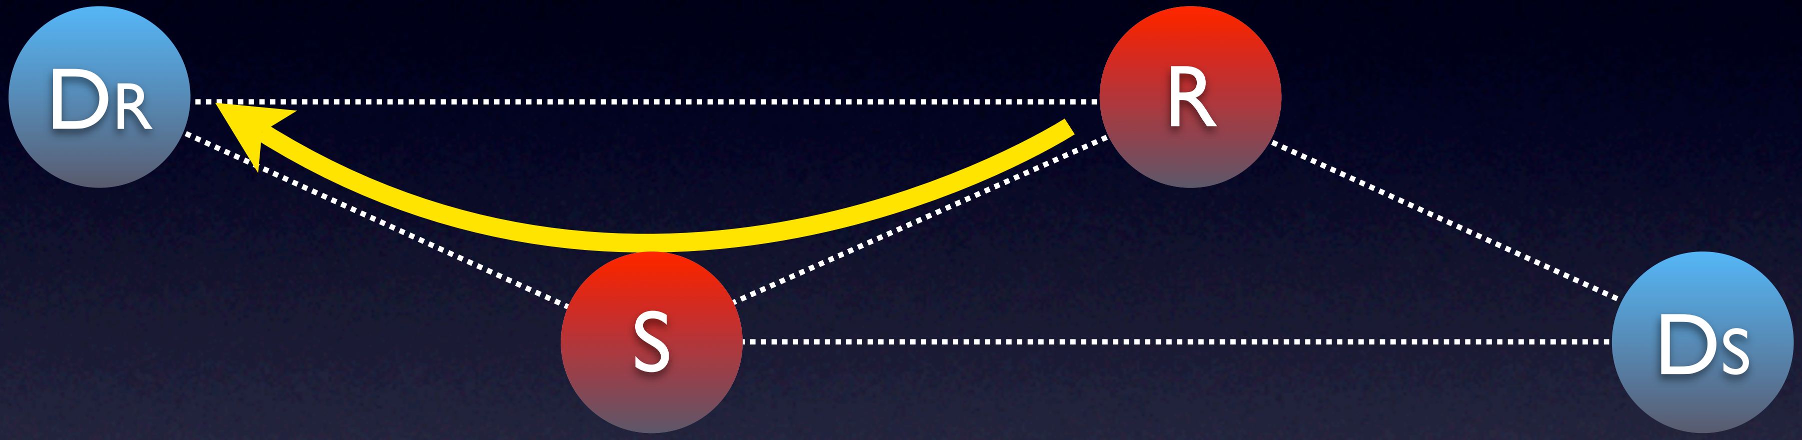

Peers can opt out

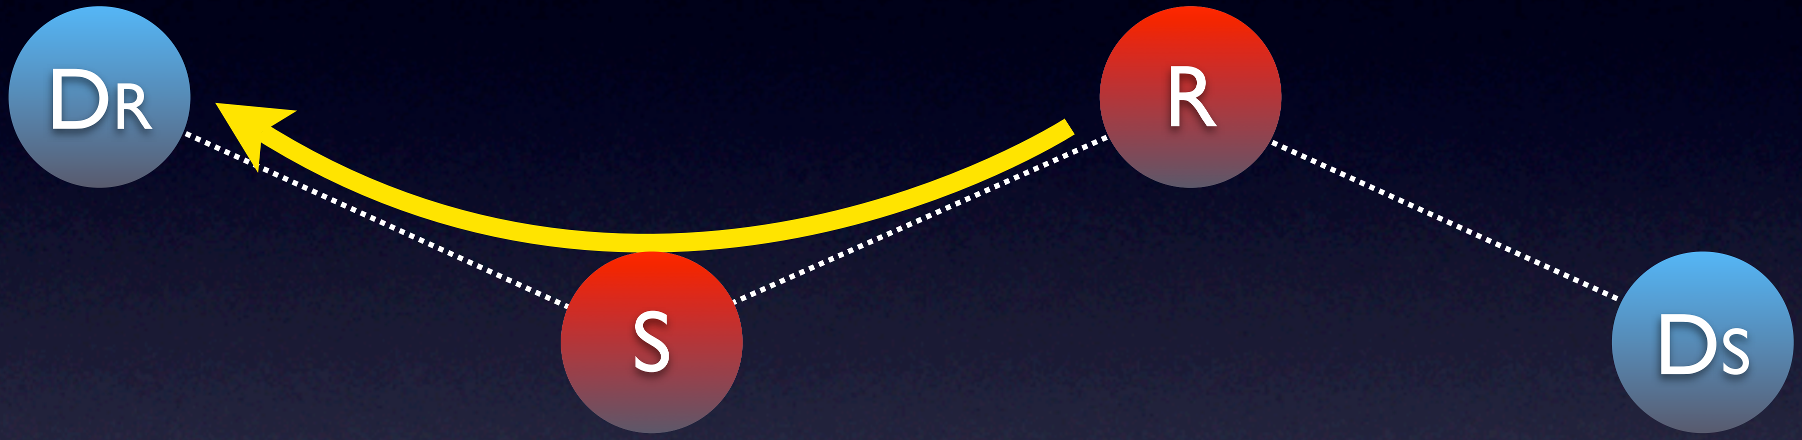

Peers can opt out

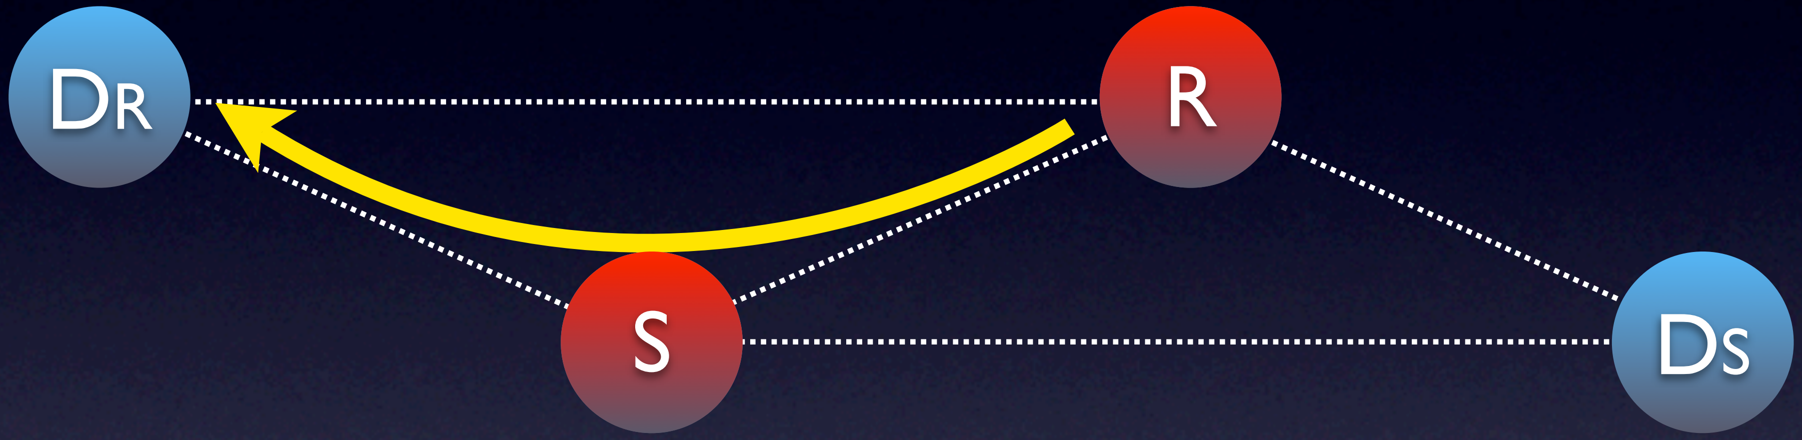

Peers can opt out

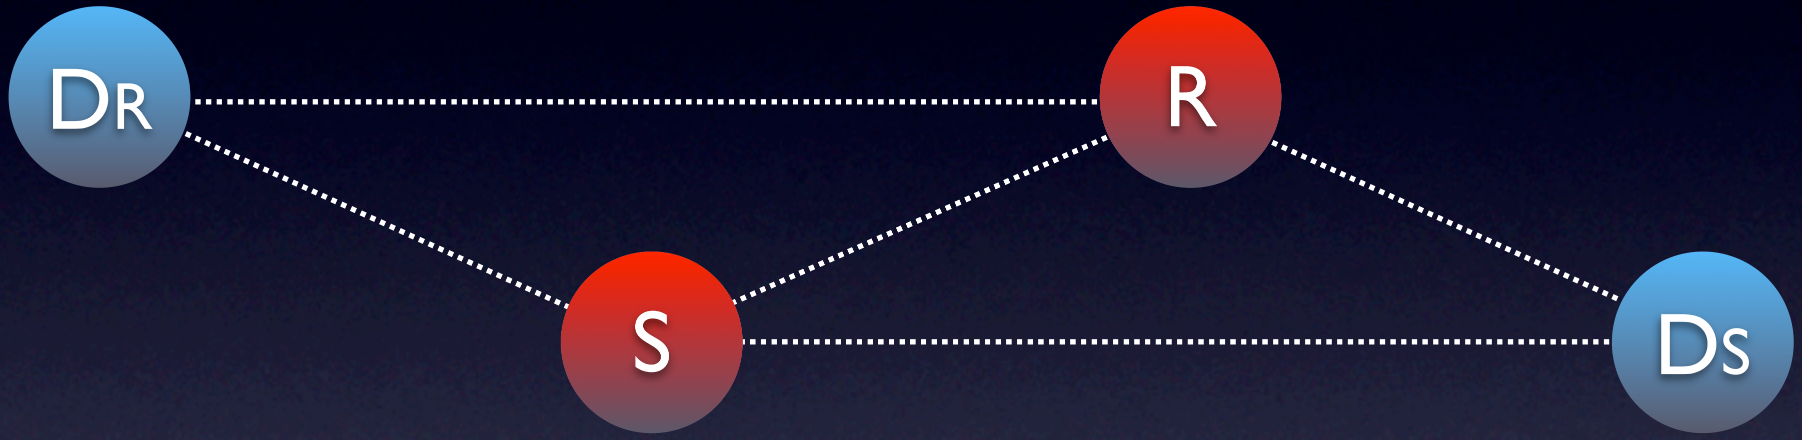

Peers can opt out

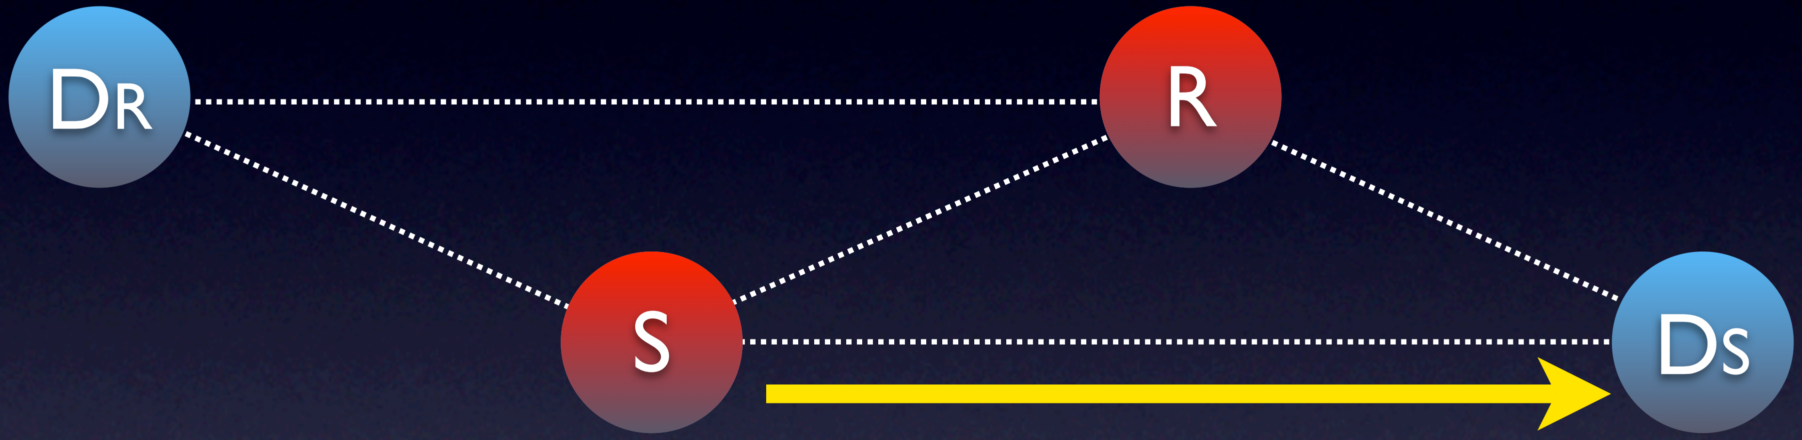

Peers can opt out

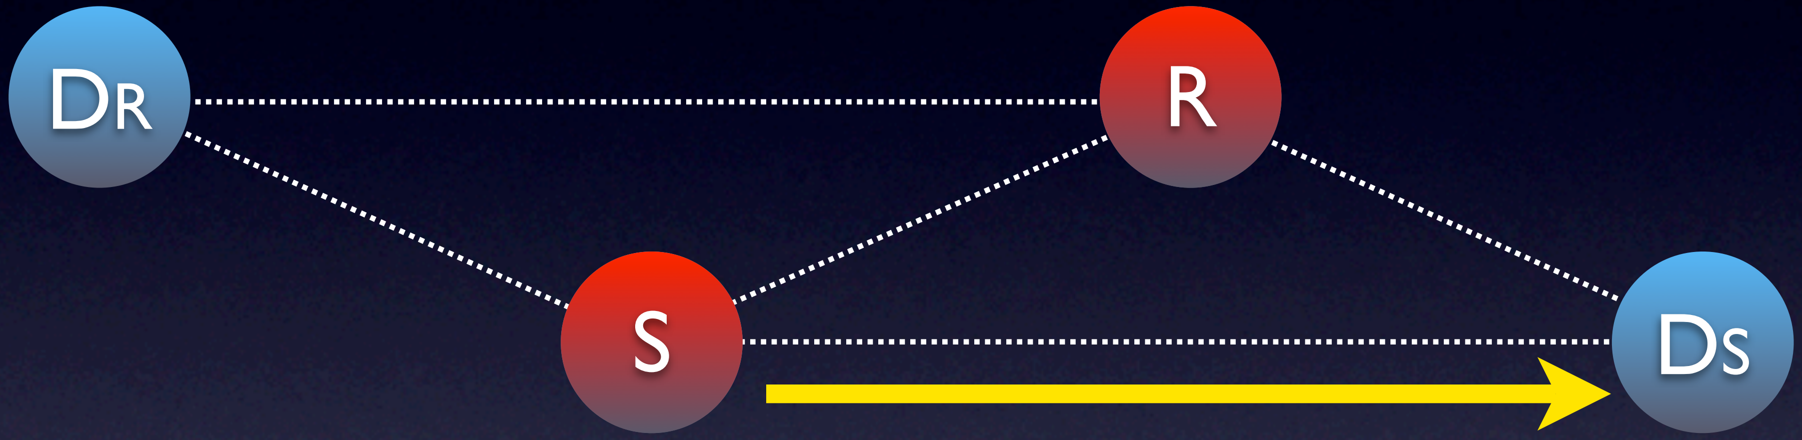

Keep users in the system

Not all peers benefit from one another

Not all peers benefit from one another

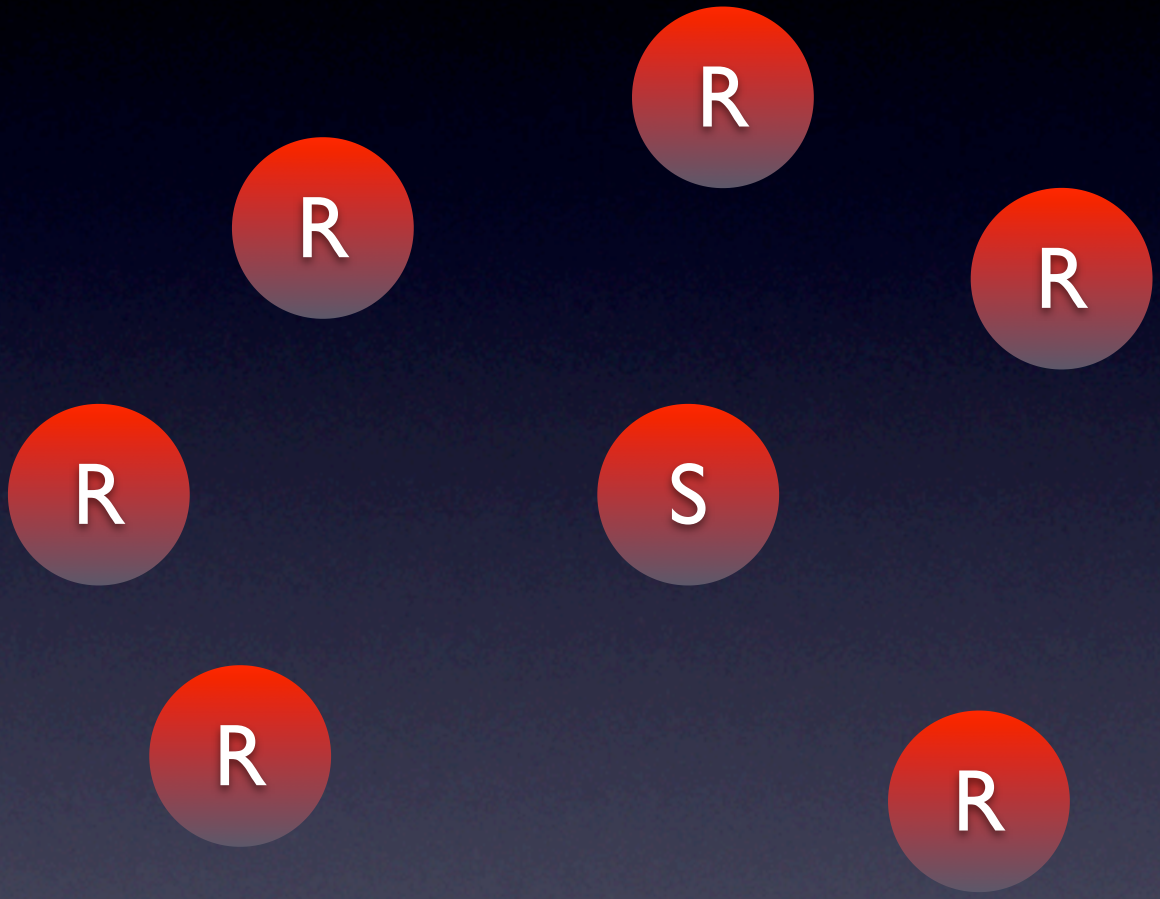

Not all peers benefit from one another

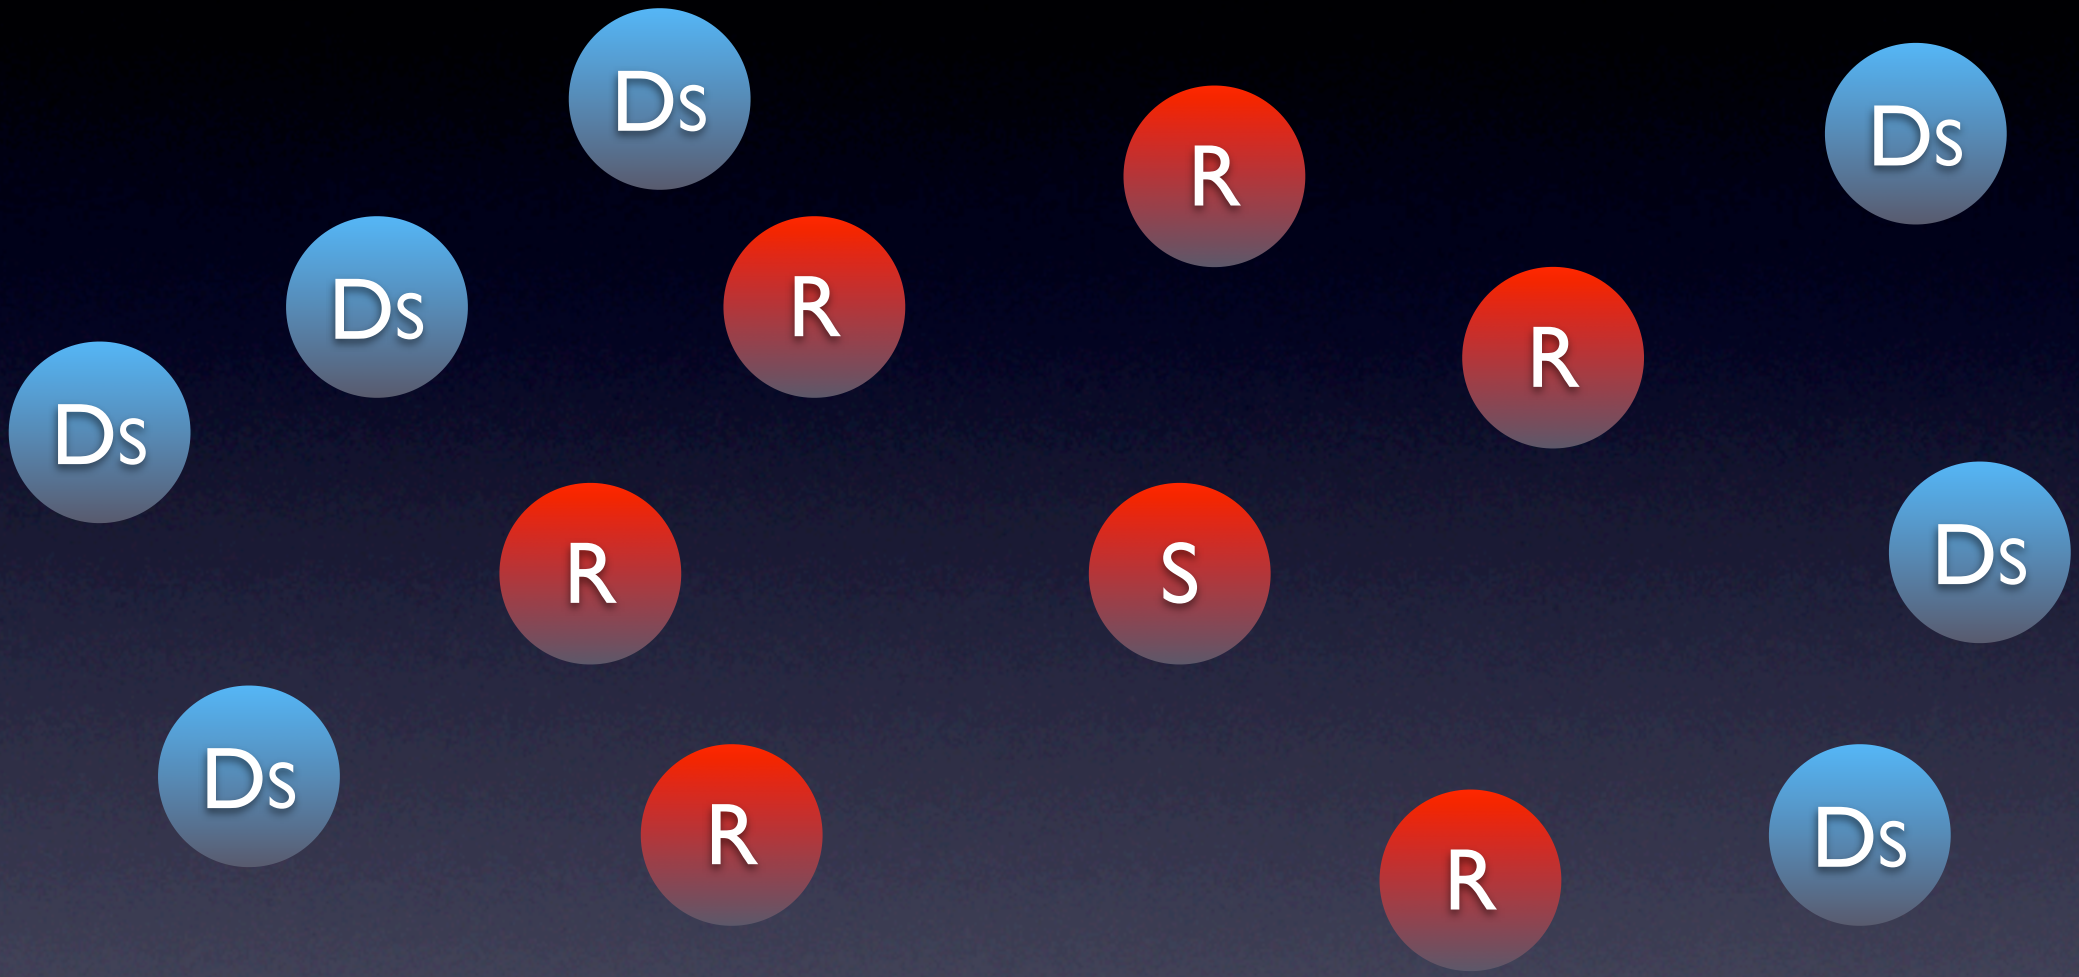

Not all peers benefit from one another

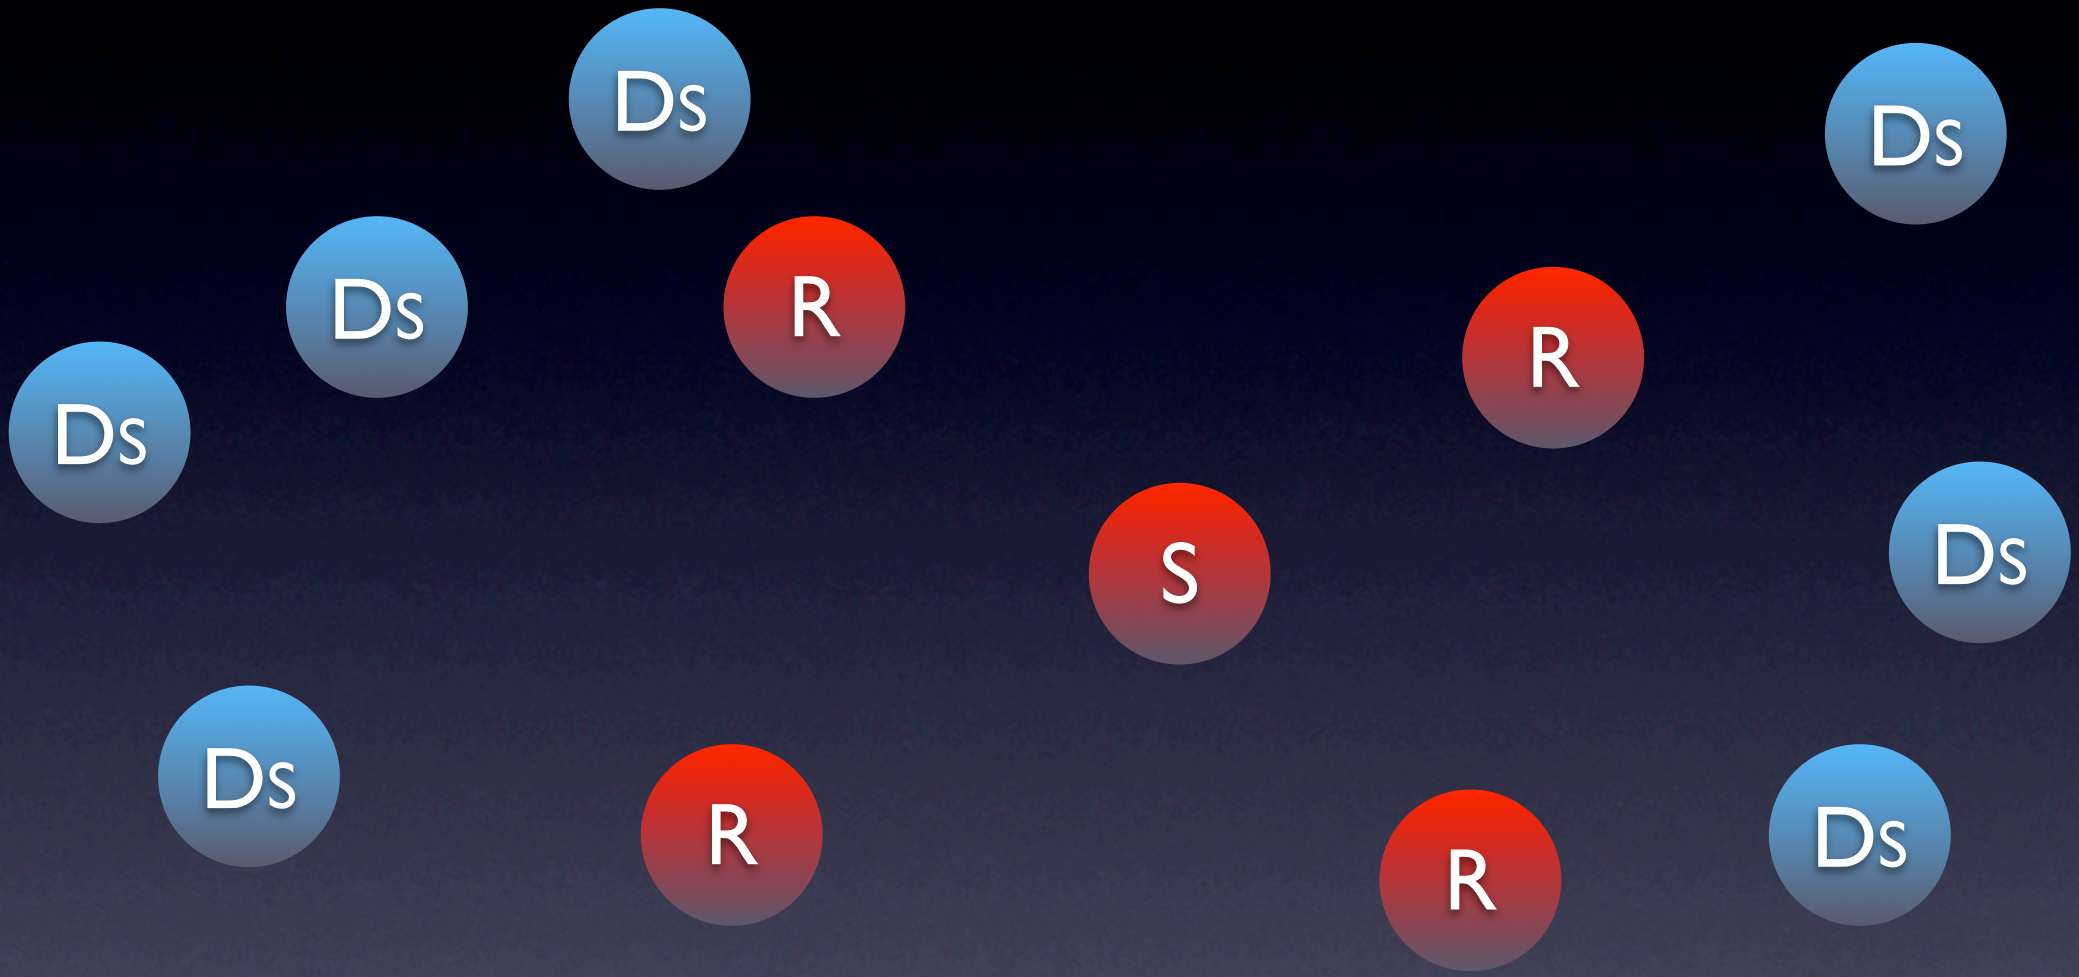

Not all peers benefit from one another

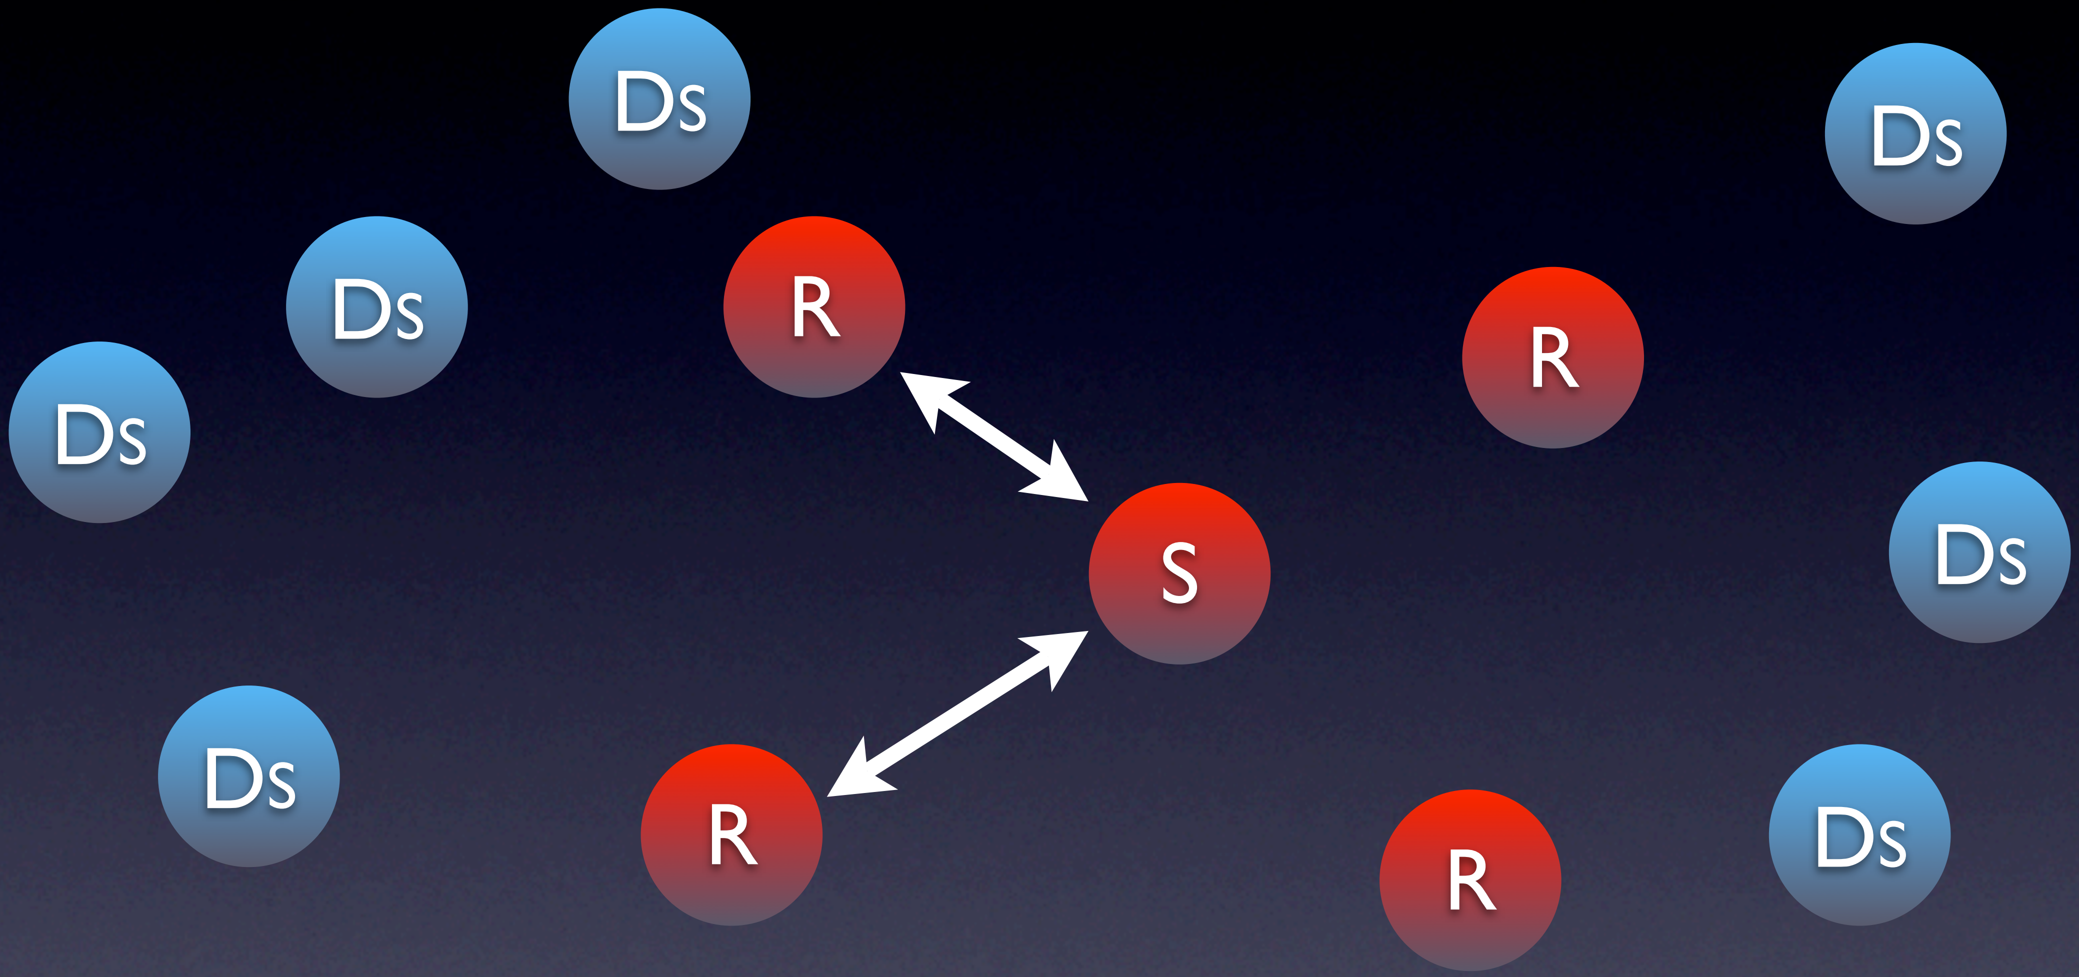

Not all peers benefit from one another

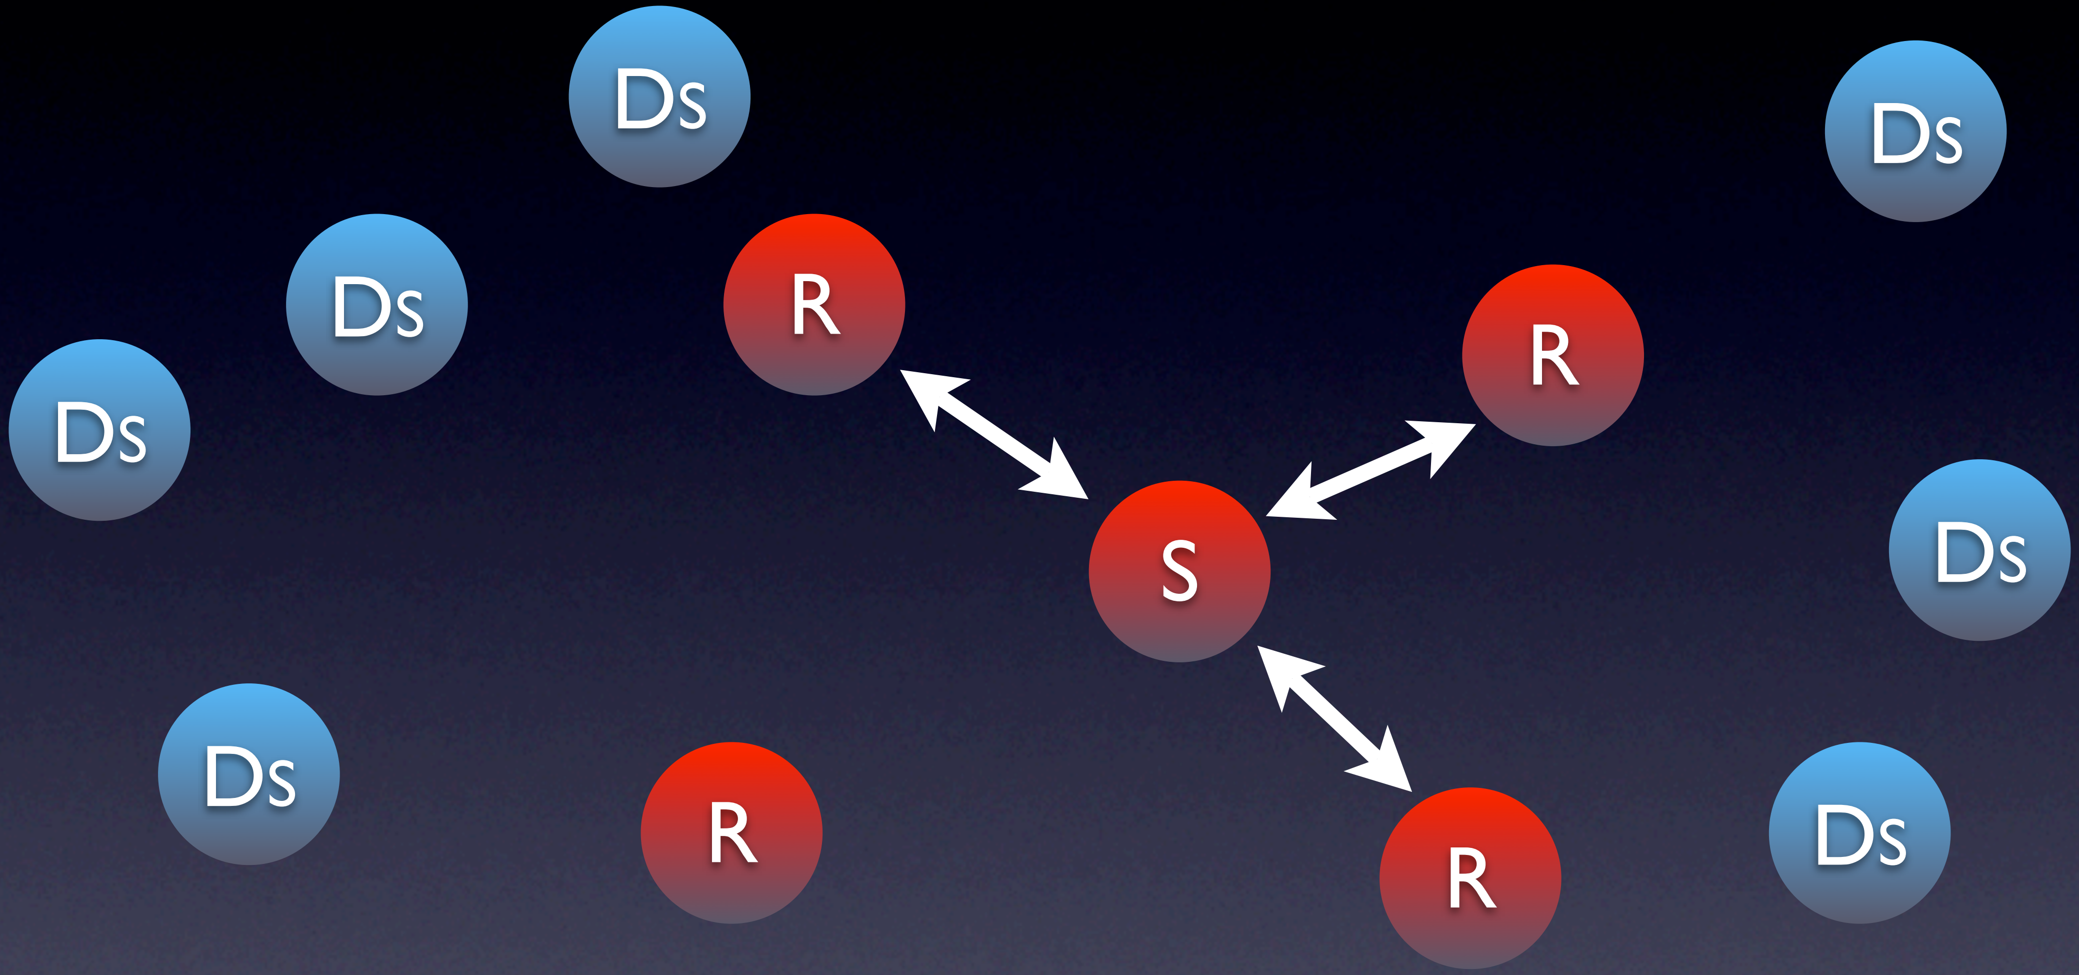

Not all peers benefit from one another

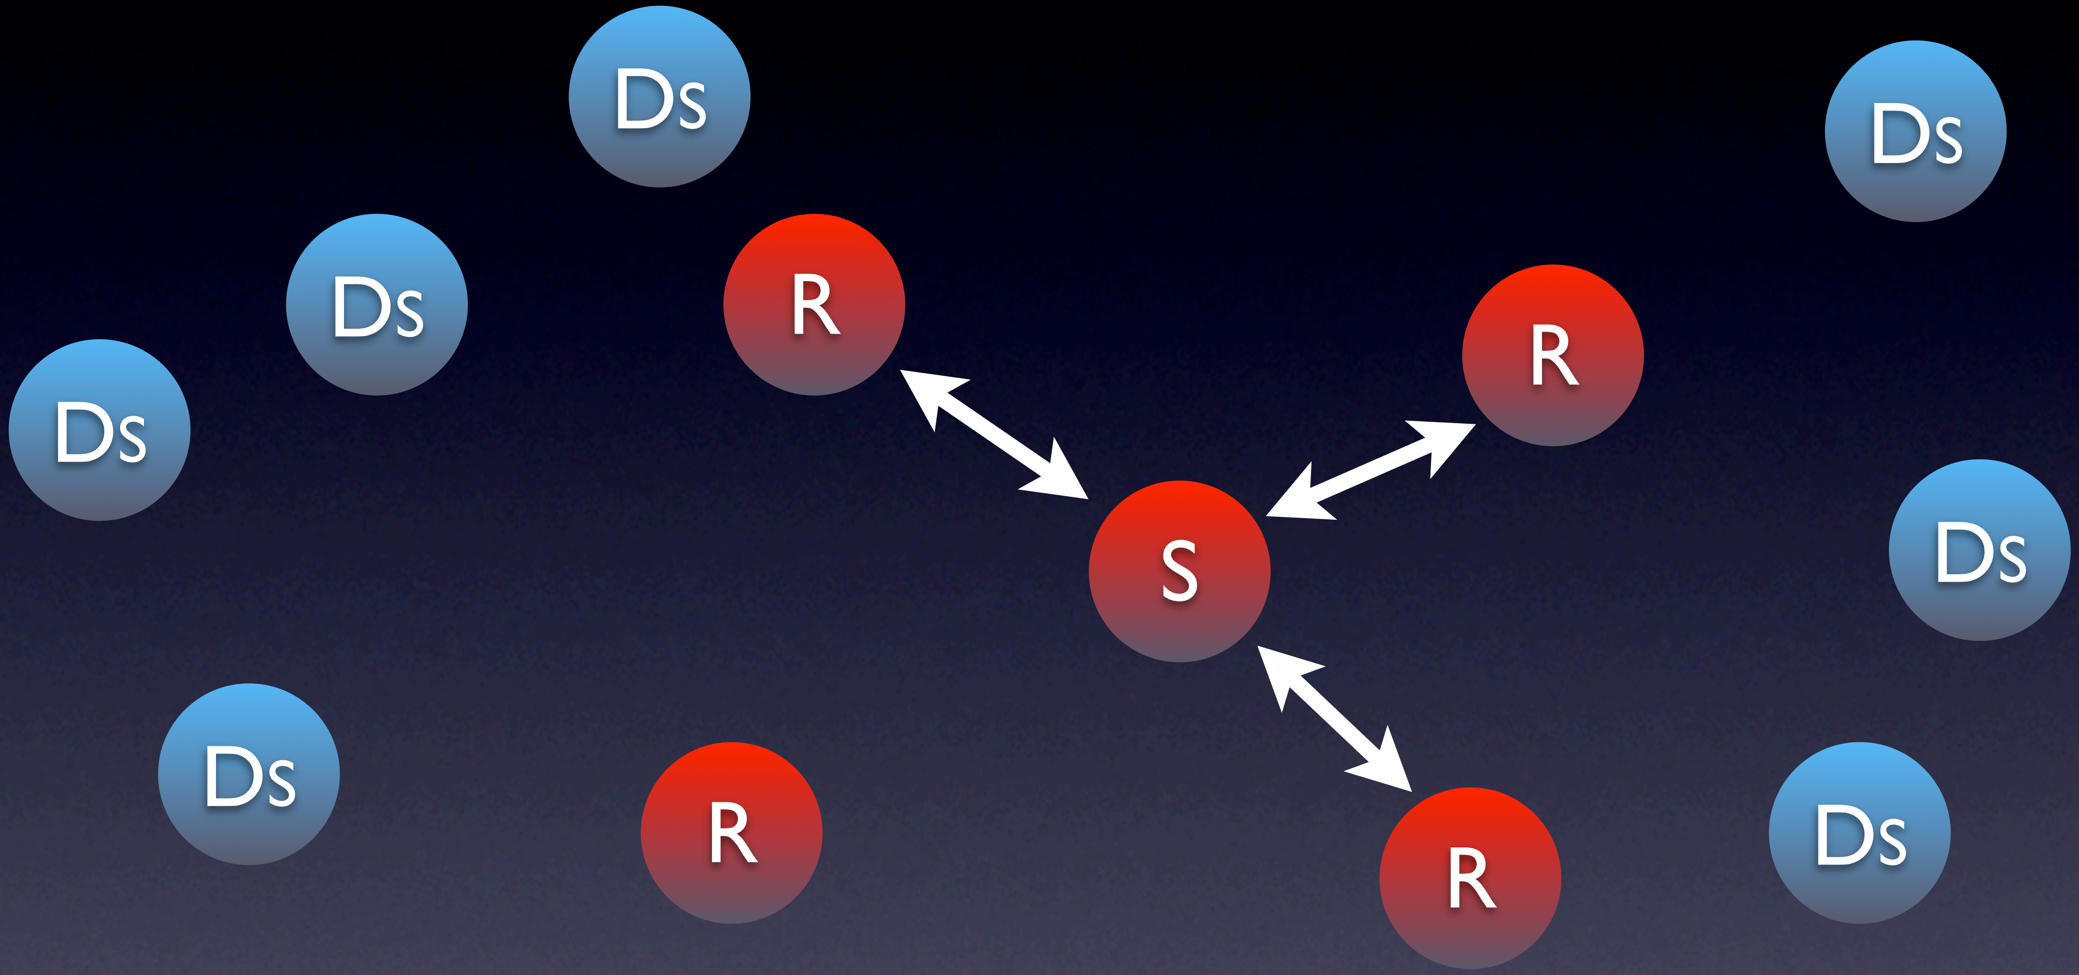

Benefit as many peers as possible

Users have varying demand

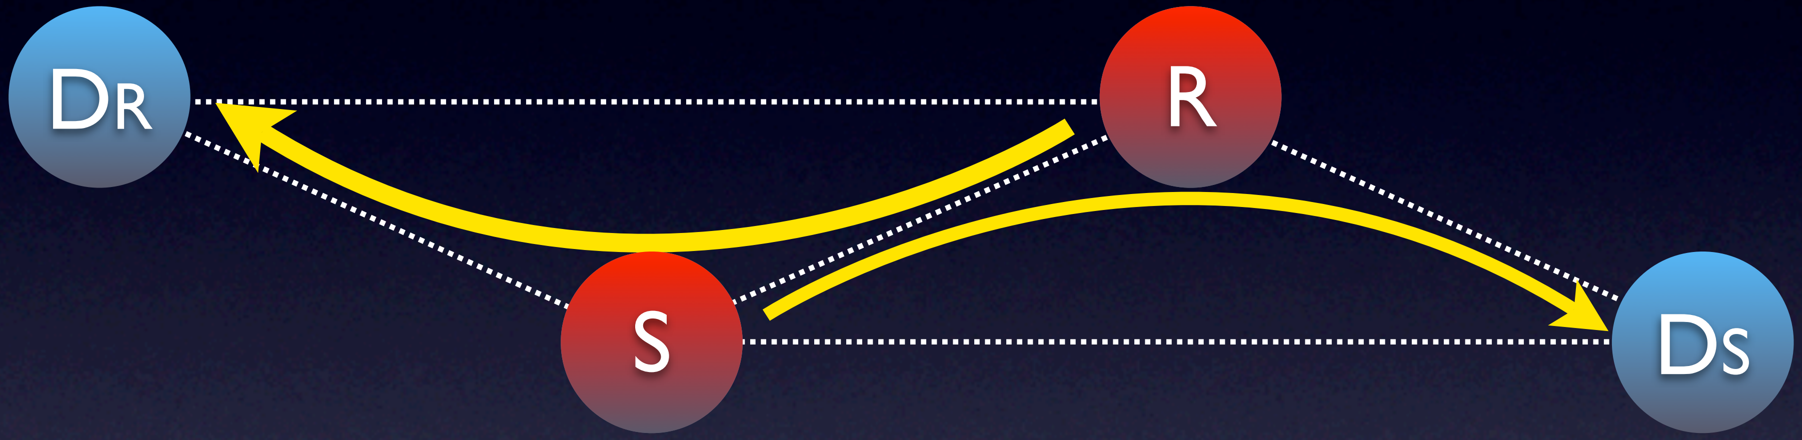

Users have varying demand

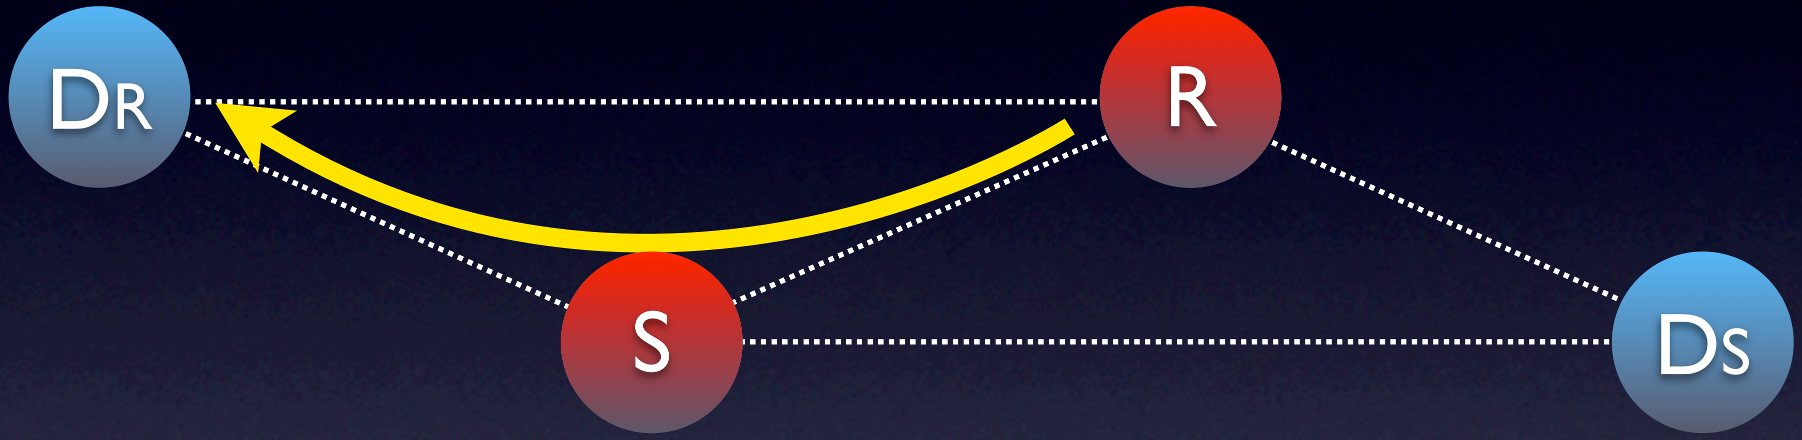

Users have varying demand

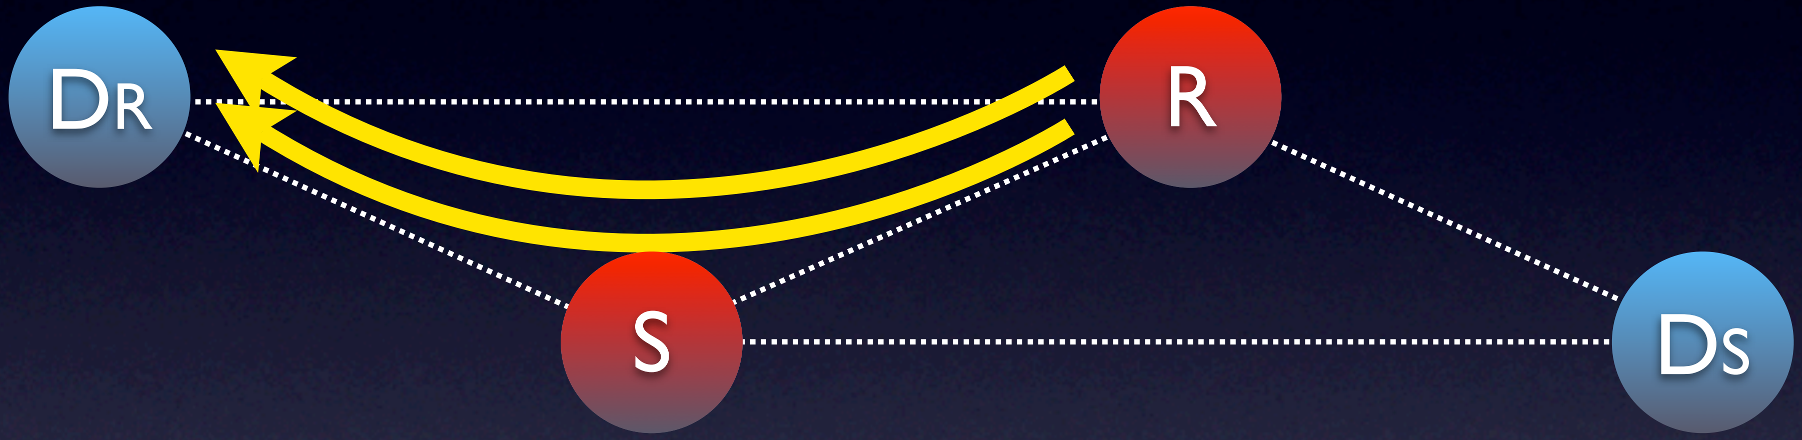

Users have varying demand

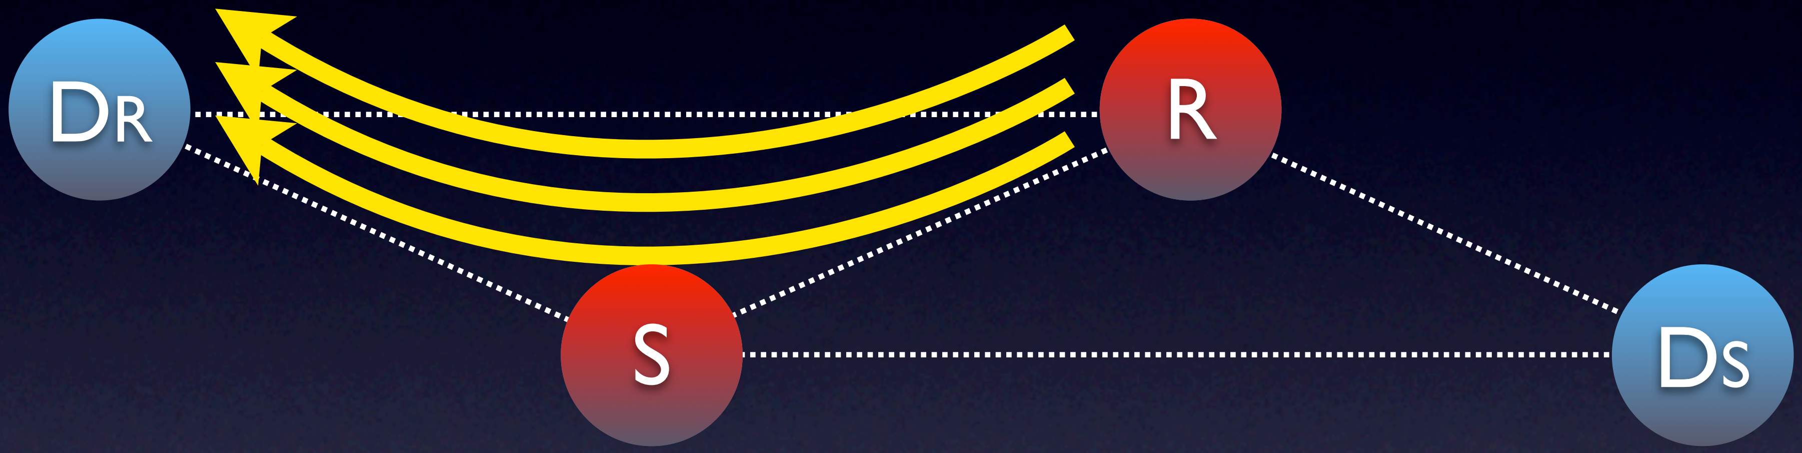

Users have varying demand

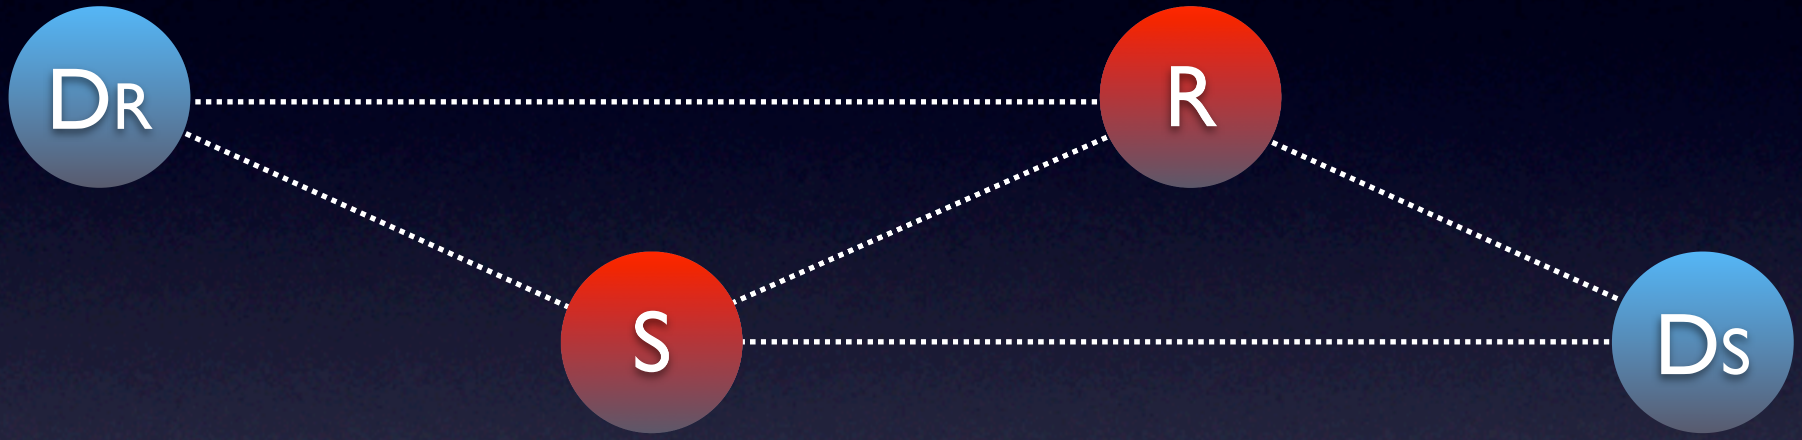

Users have varying demand

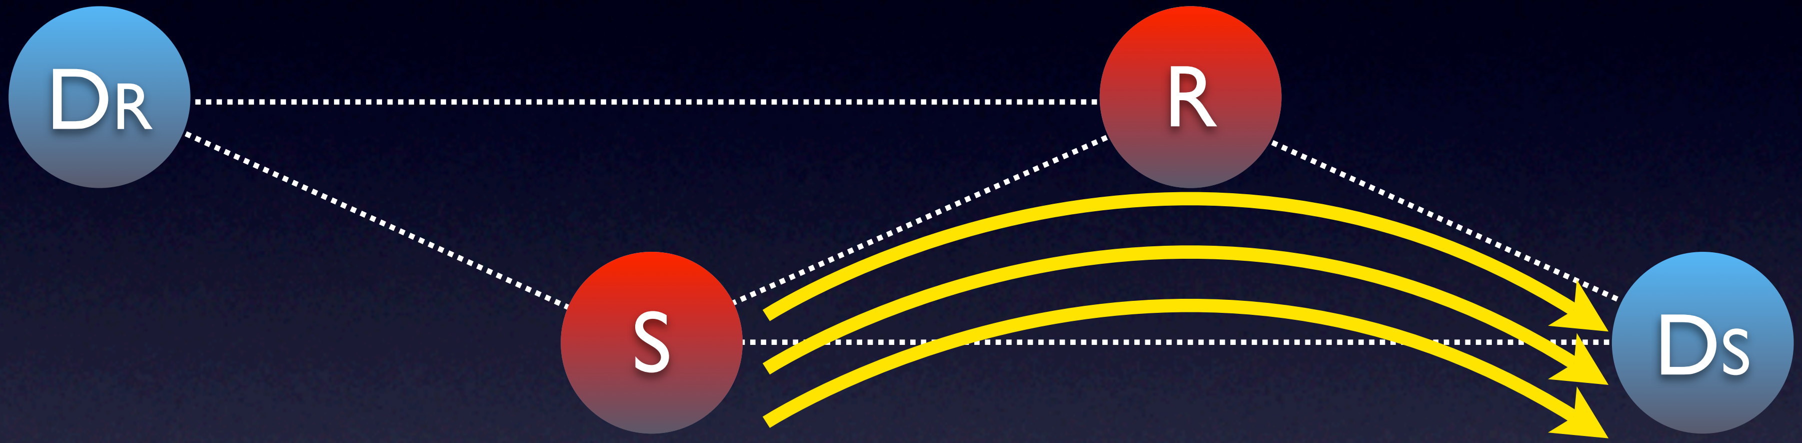

Users have varying demand

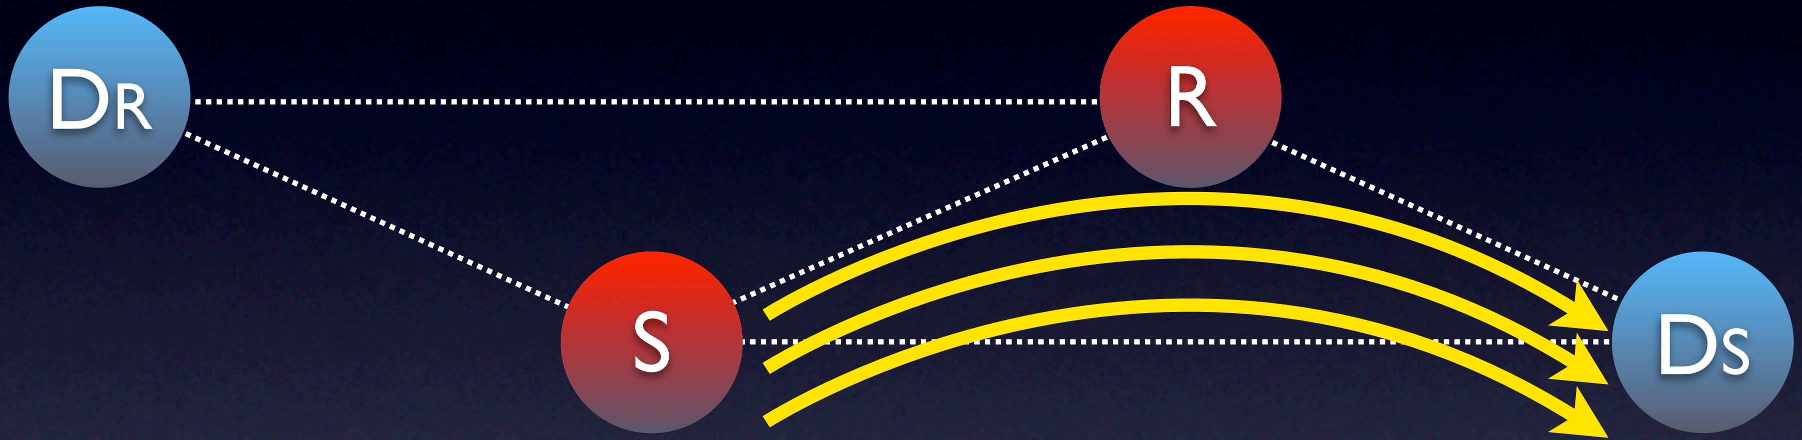

Long-lived peerings preferred

The Selfish Routing Overlay Problem

Users can opt out

Keep users in the system

Not all peers benefit from one another

Benefit as many peers as possible

Users have varying demand

Long-lived peerings preferred

Contributions

- The selfish routing overlay problem
- Mechanisms to achieve long-lived peerings
- Simulation study

Contributions

- The selfish routing overlay problem
- Mechanisms to achieve long-lived peerings
- Simulation study

Our approach

Our approach

Let peers communicate their demands

Our approach

Maximize benefit

Latency reduction

Minimize cost

Long-term bandwidth consumption

Short-term burstiness

Let peers communicate their demands

Our approach

Latency reduction

Long-term bandwidth consumption

Short-term burstiness

Let peers communicate their demands

Our approach

Latency reduction

Long-term bandwidth consumption

Short-term burstiness

Service-Level Agreements

Let peers communicate their demands

Our approach

Latency reduction

Long-term bandwidth consumption

Short-term burstiness

Service-Level Agreements

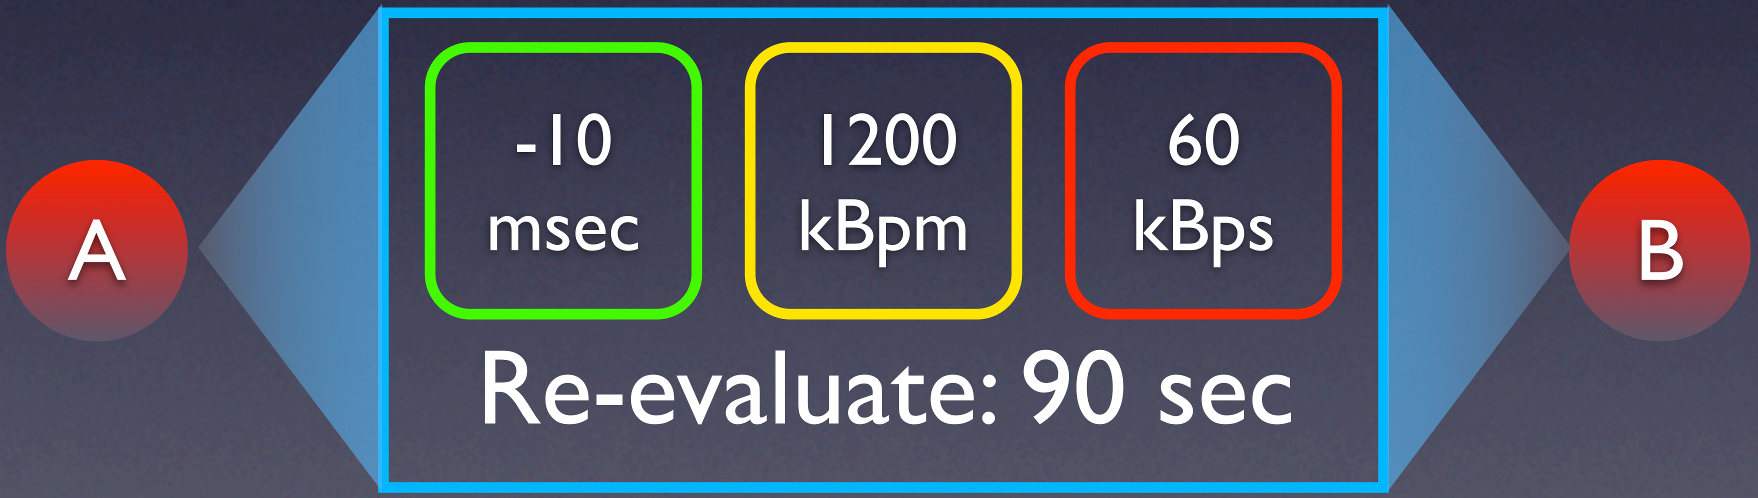

Our approach

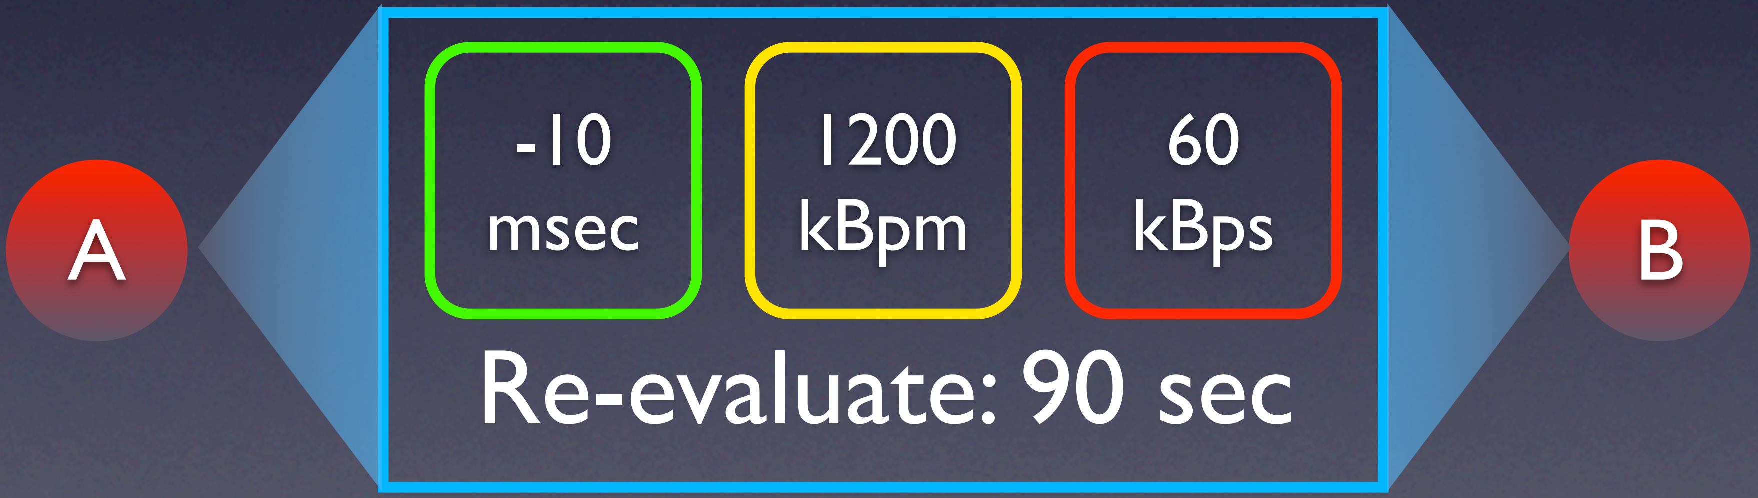

Service-Level Agreements

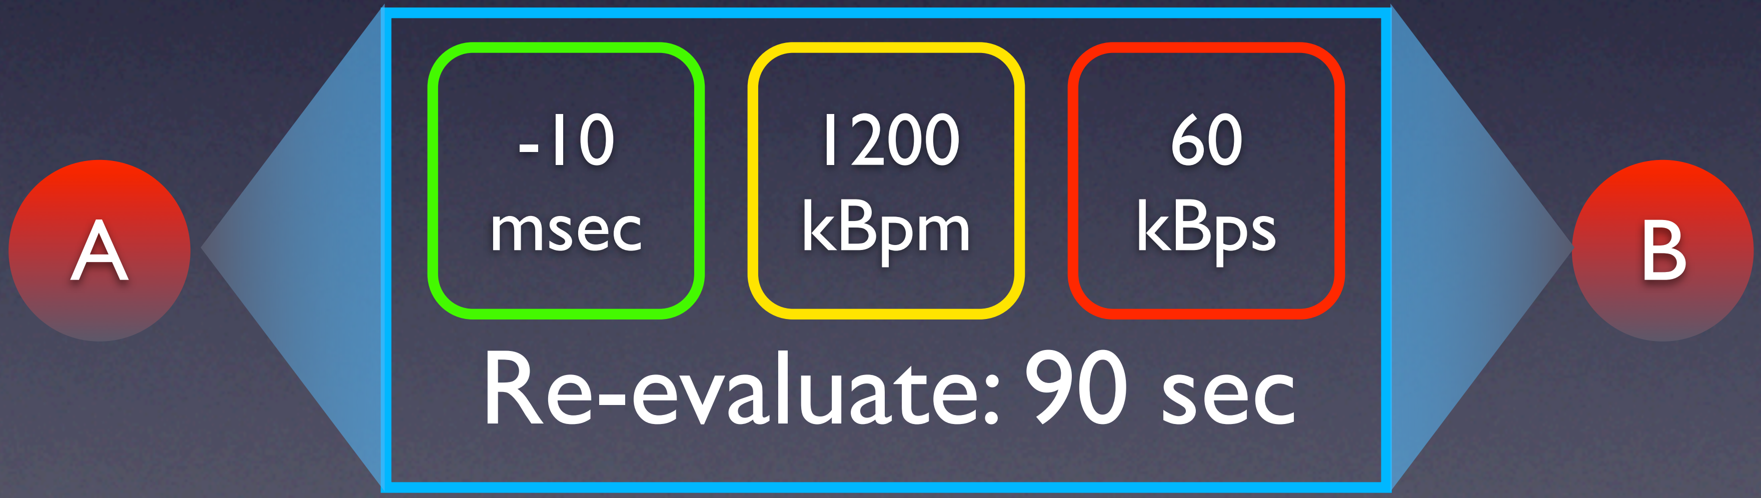

Service-Level Agreements

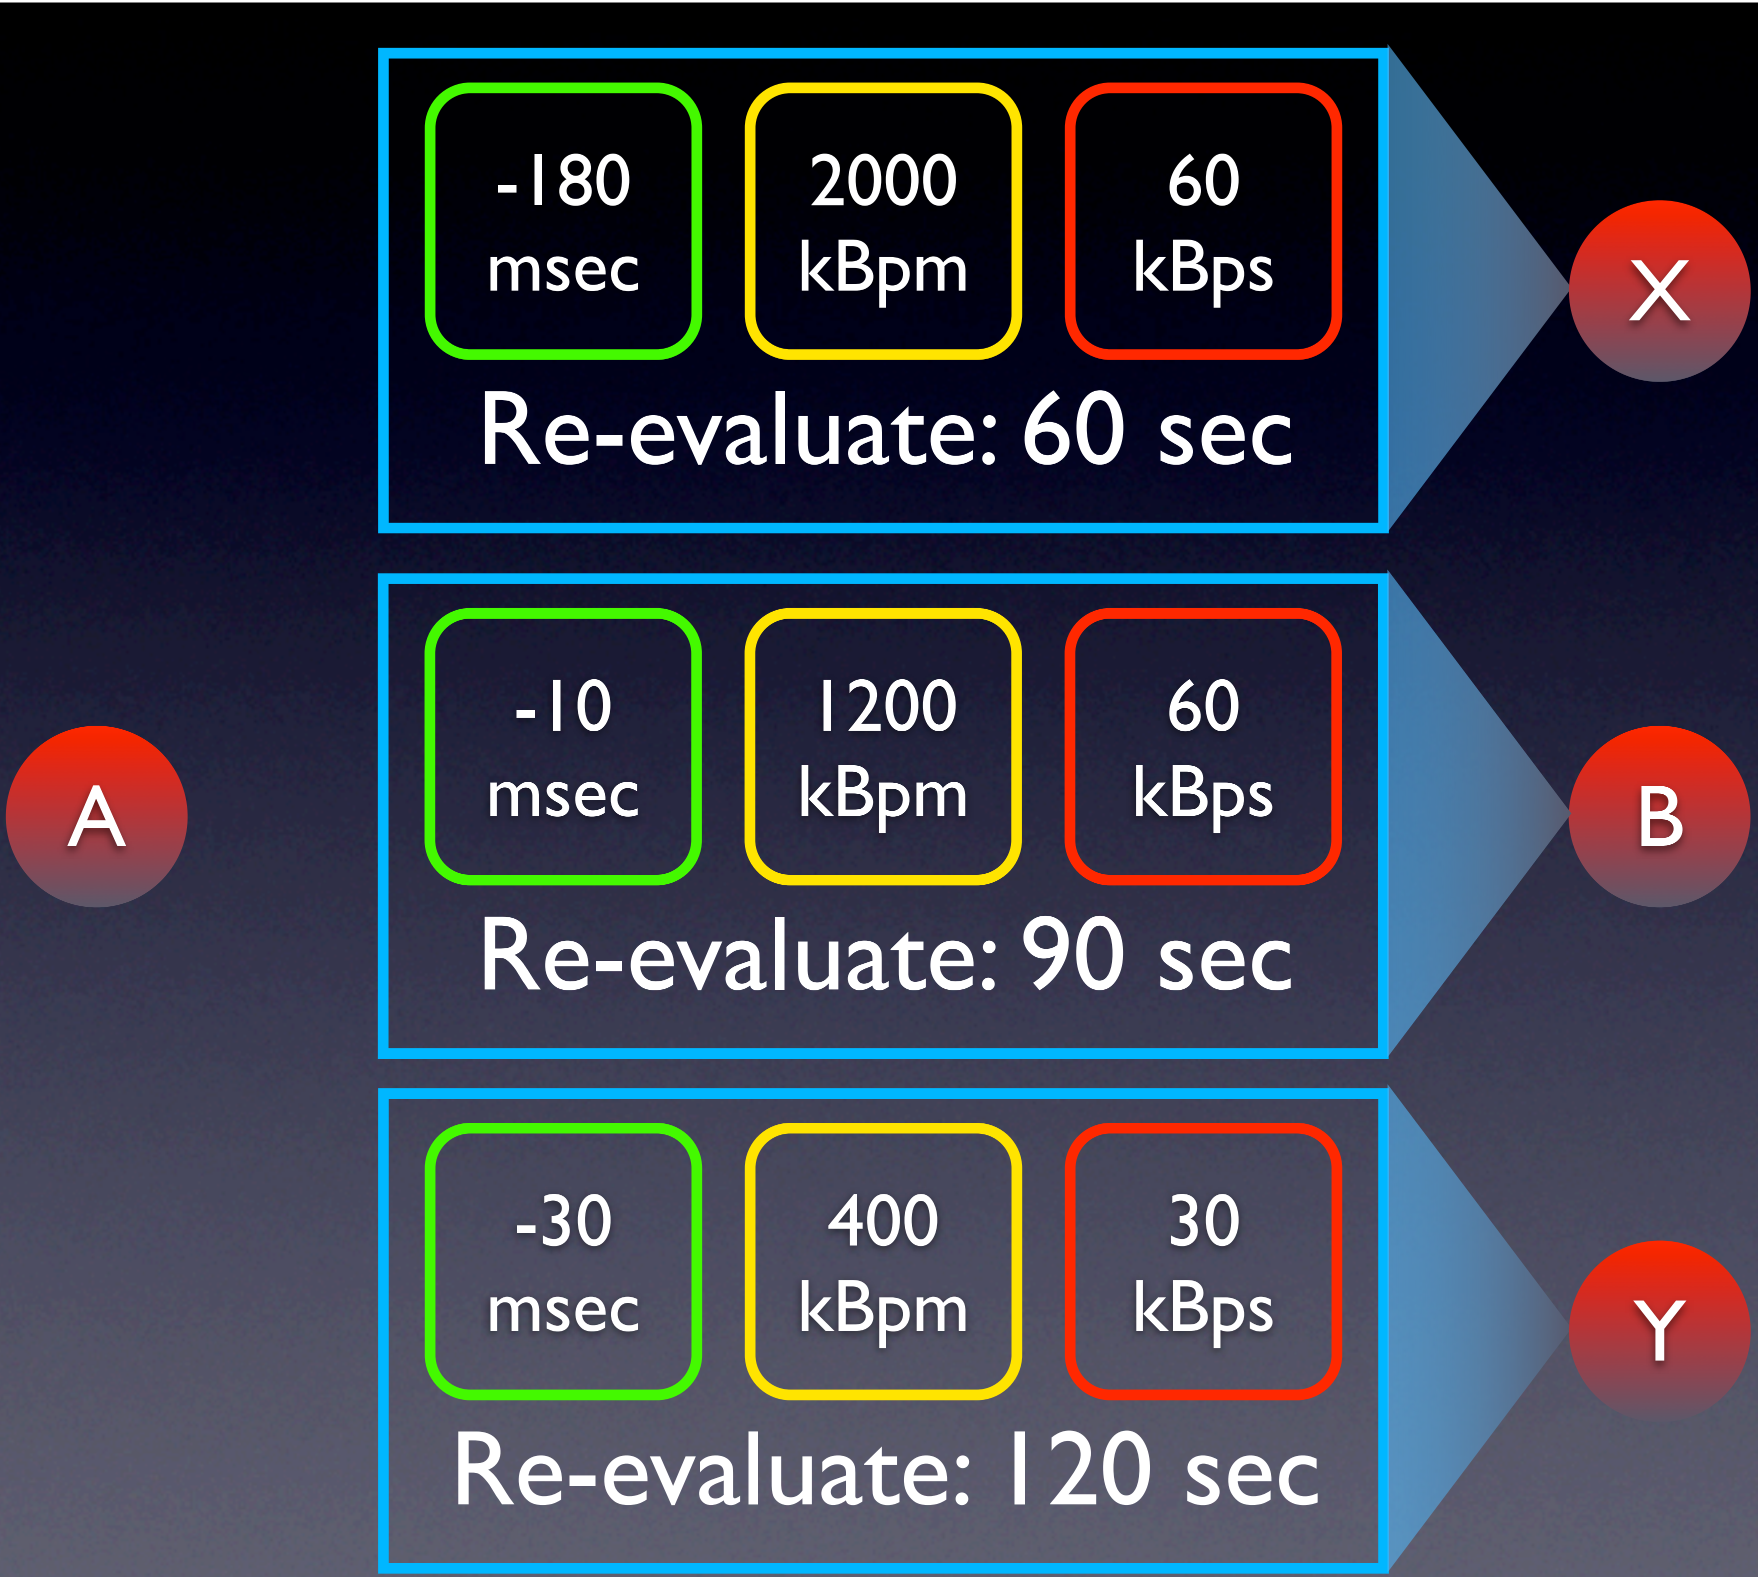

-180 msec | 2000 kBpm | 60 kBps

Re-evaluate: 60 sec

-10 msec | 1200 kBpm | 60 kBps

Re-evaluate: 90 sec

-30 msec | 400 kBpm | 30 kBps

Re-evaluate: 120 sec

Service-Level Agreements

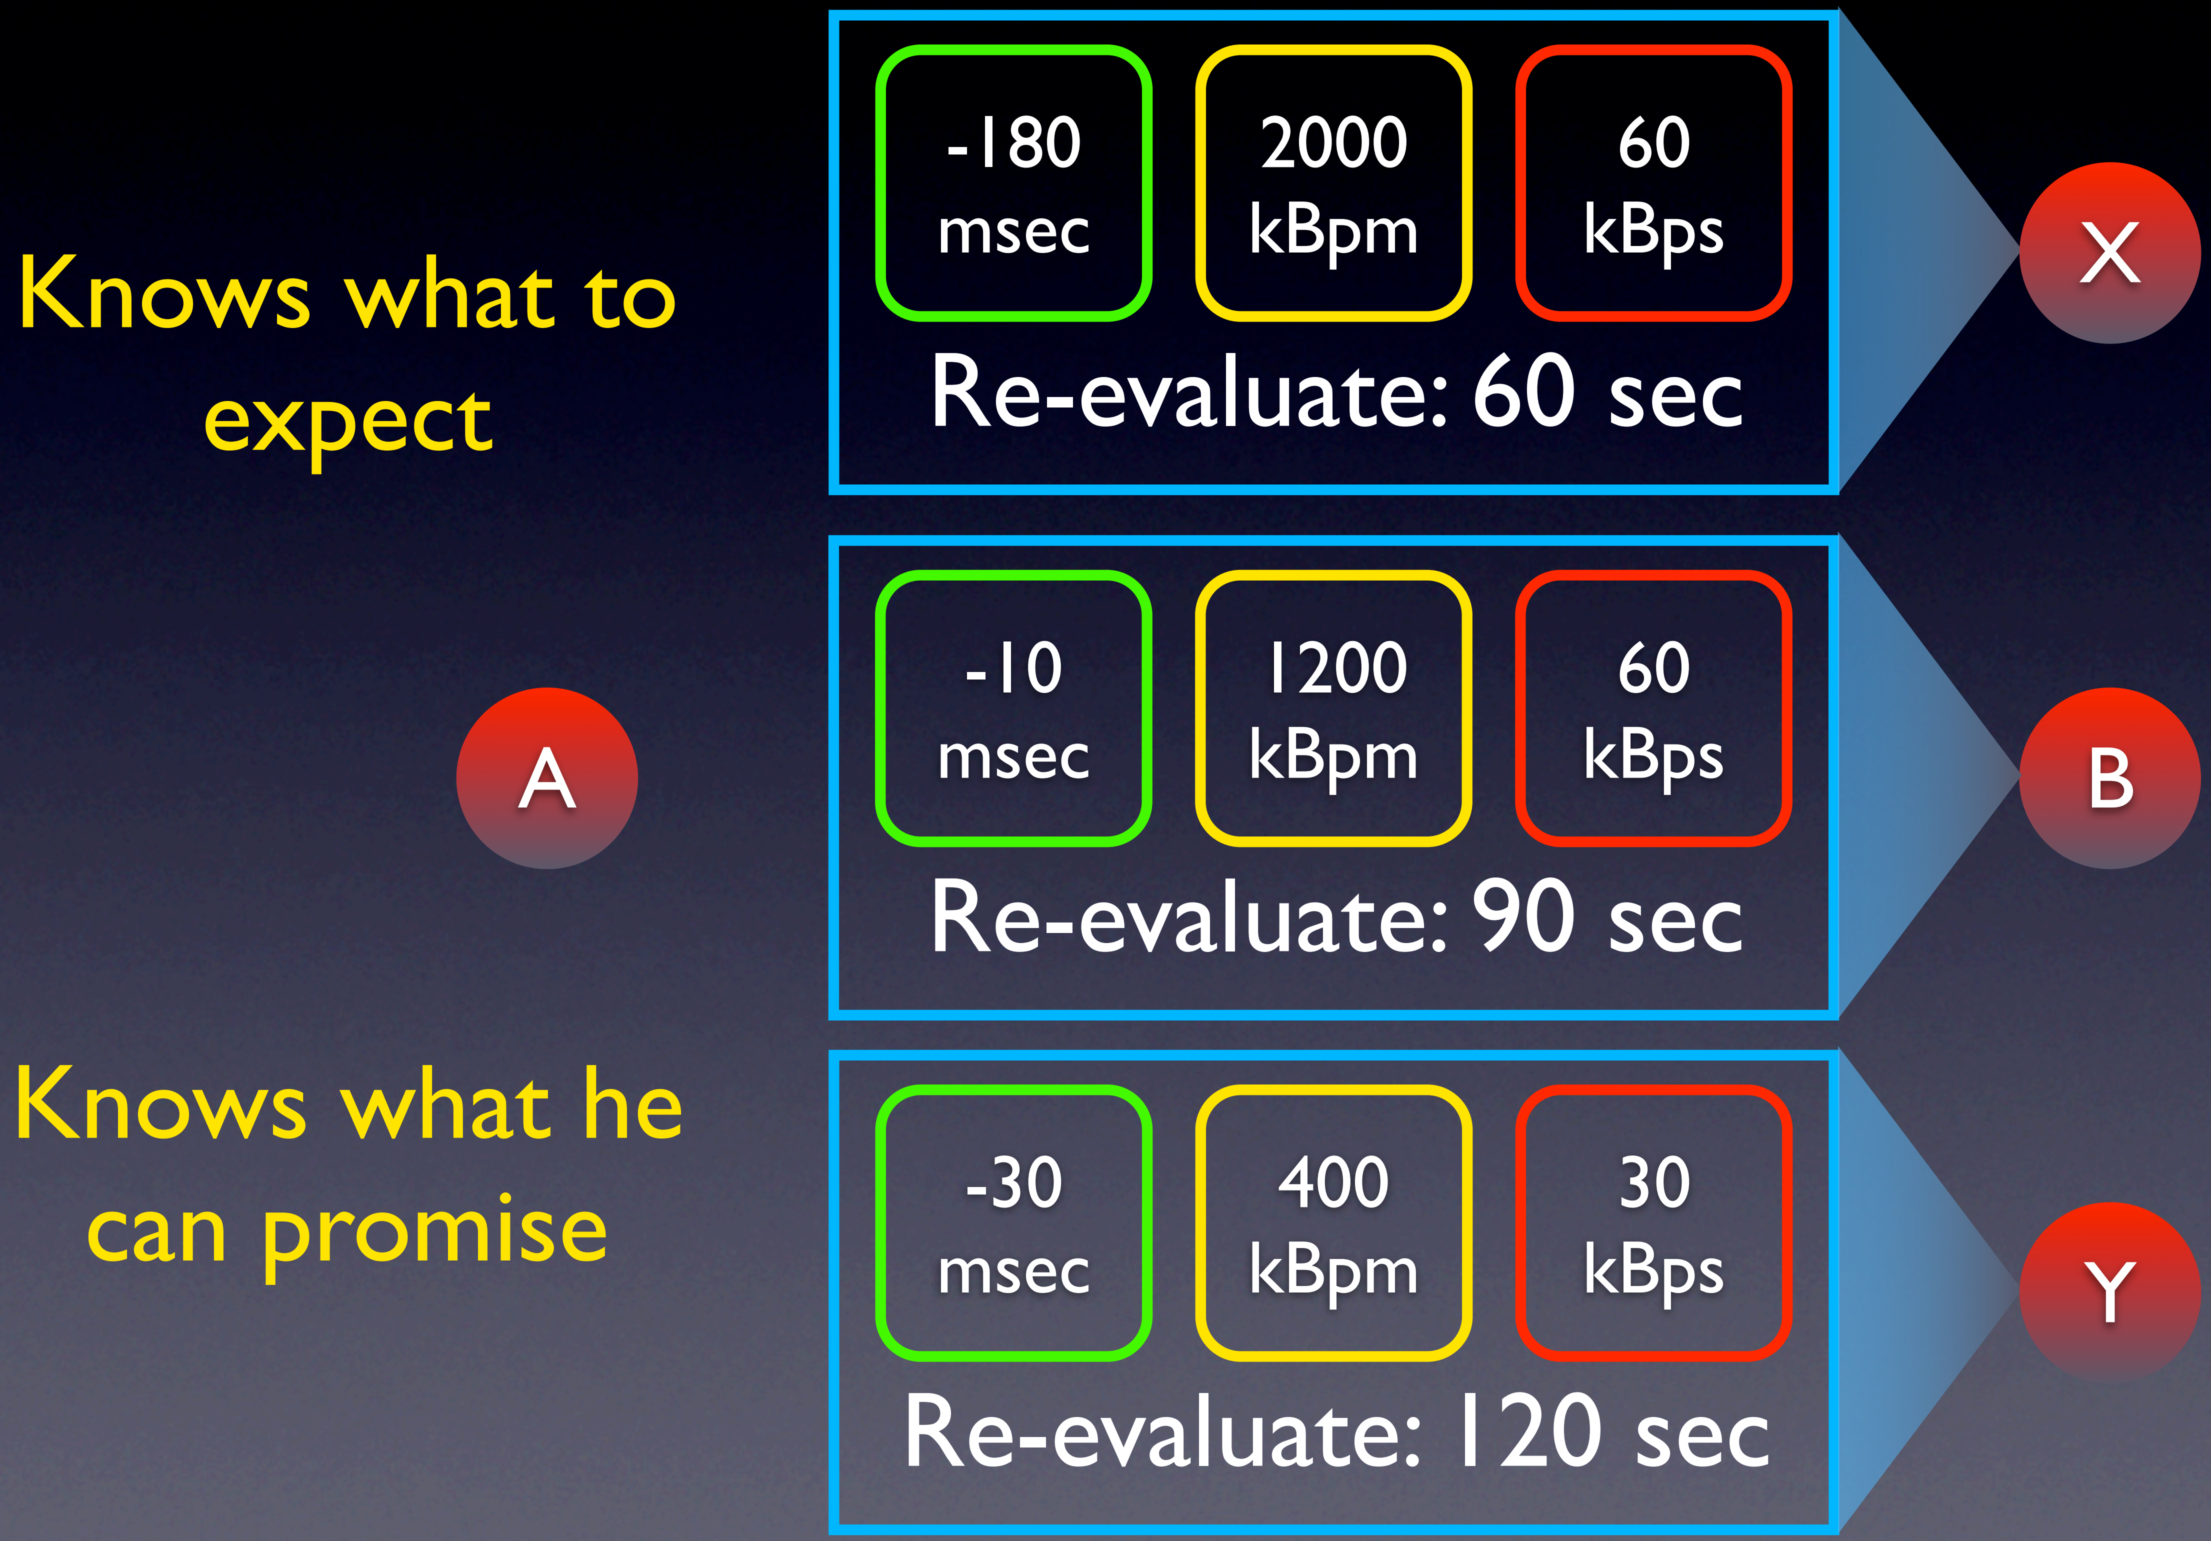

Knows what to expect

-180 msec | 2000 kBpm | 60 kBps

Re-evaluate: 60 sec

X

A

-10 msec | 1200 kBpm | 60 kBps

Re-evaluate: 90 sec

B

Knows what he can promise

-30 msec | 400 kBpm | 30 kBps

Re-evaluate: 120 sec

Y

Why violate an SLA?

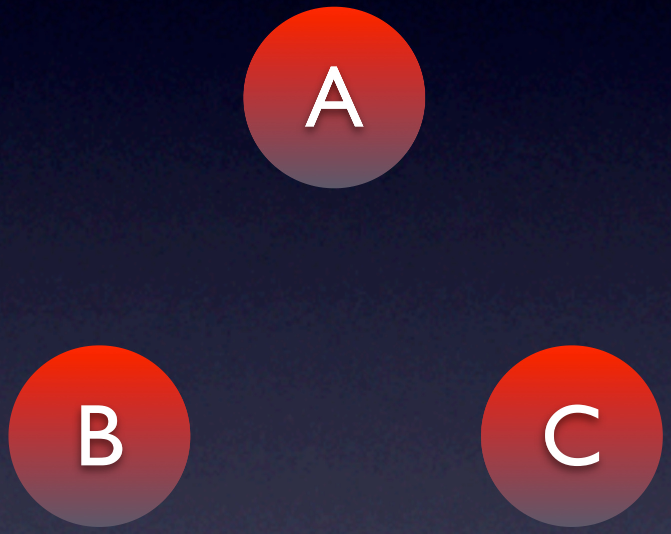

Why violate an SLA?

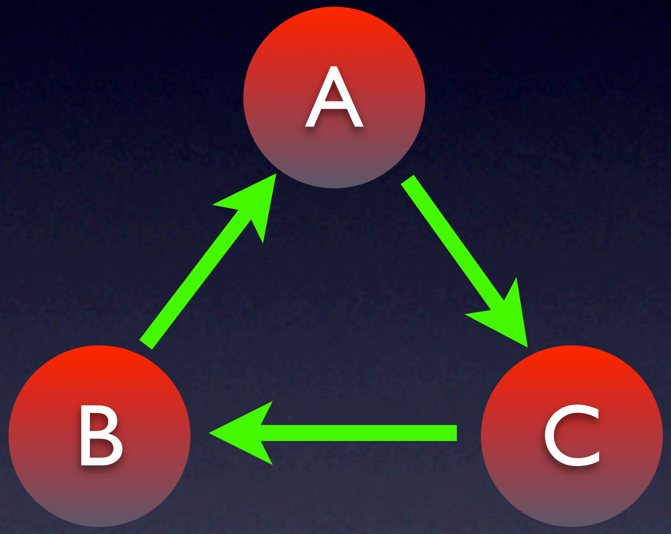

Preferred

Why violate an SLA?

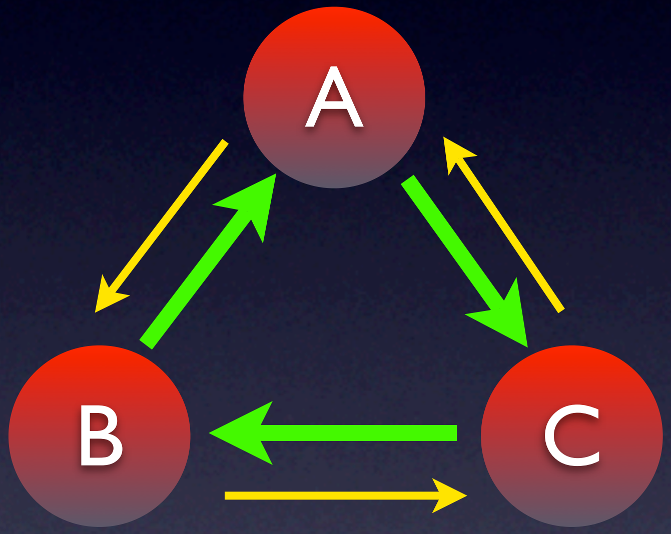

Preferred Acceptable

Why violate an SLA?

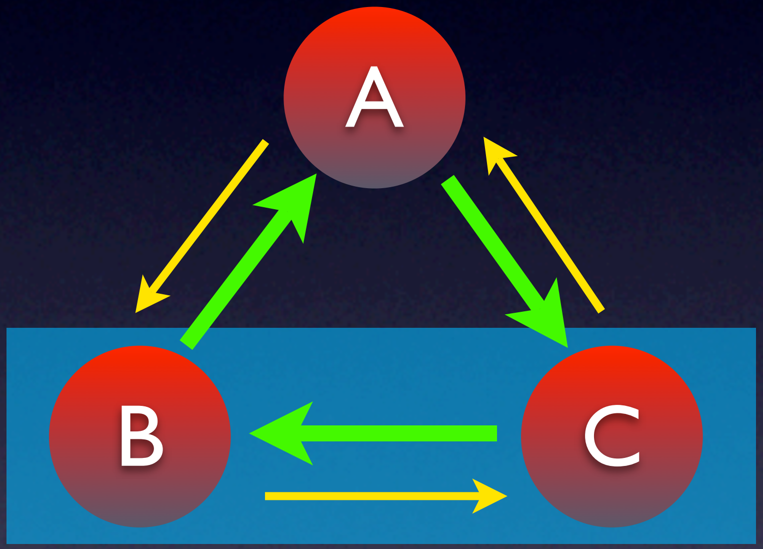

Preferred Acceptable SLA

Why violate an SLA?

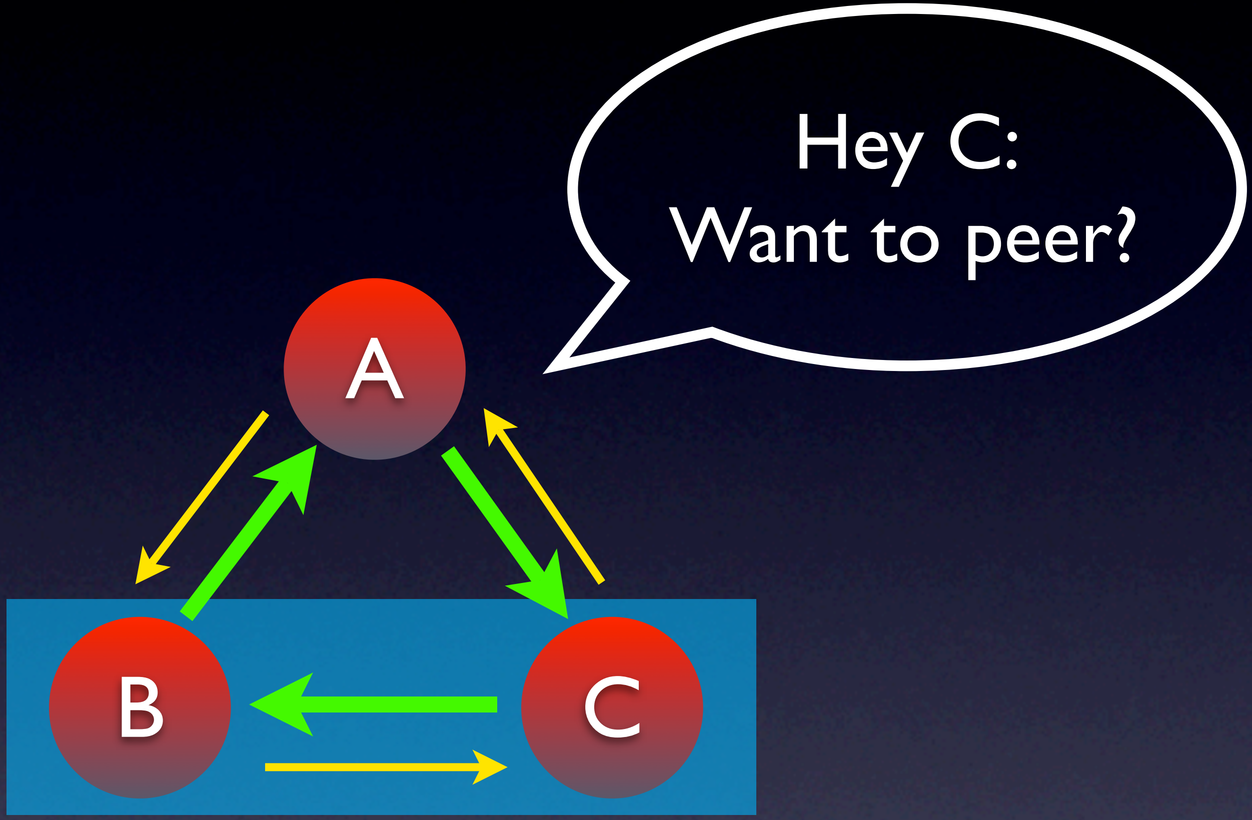

Preferred Acceptable SLA

Why violate an SLA?

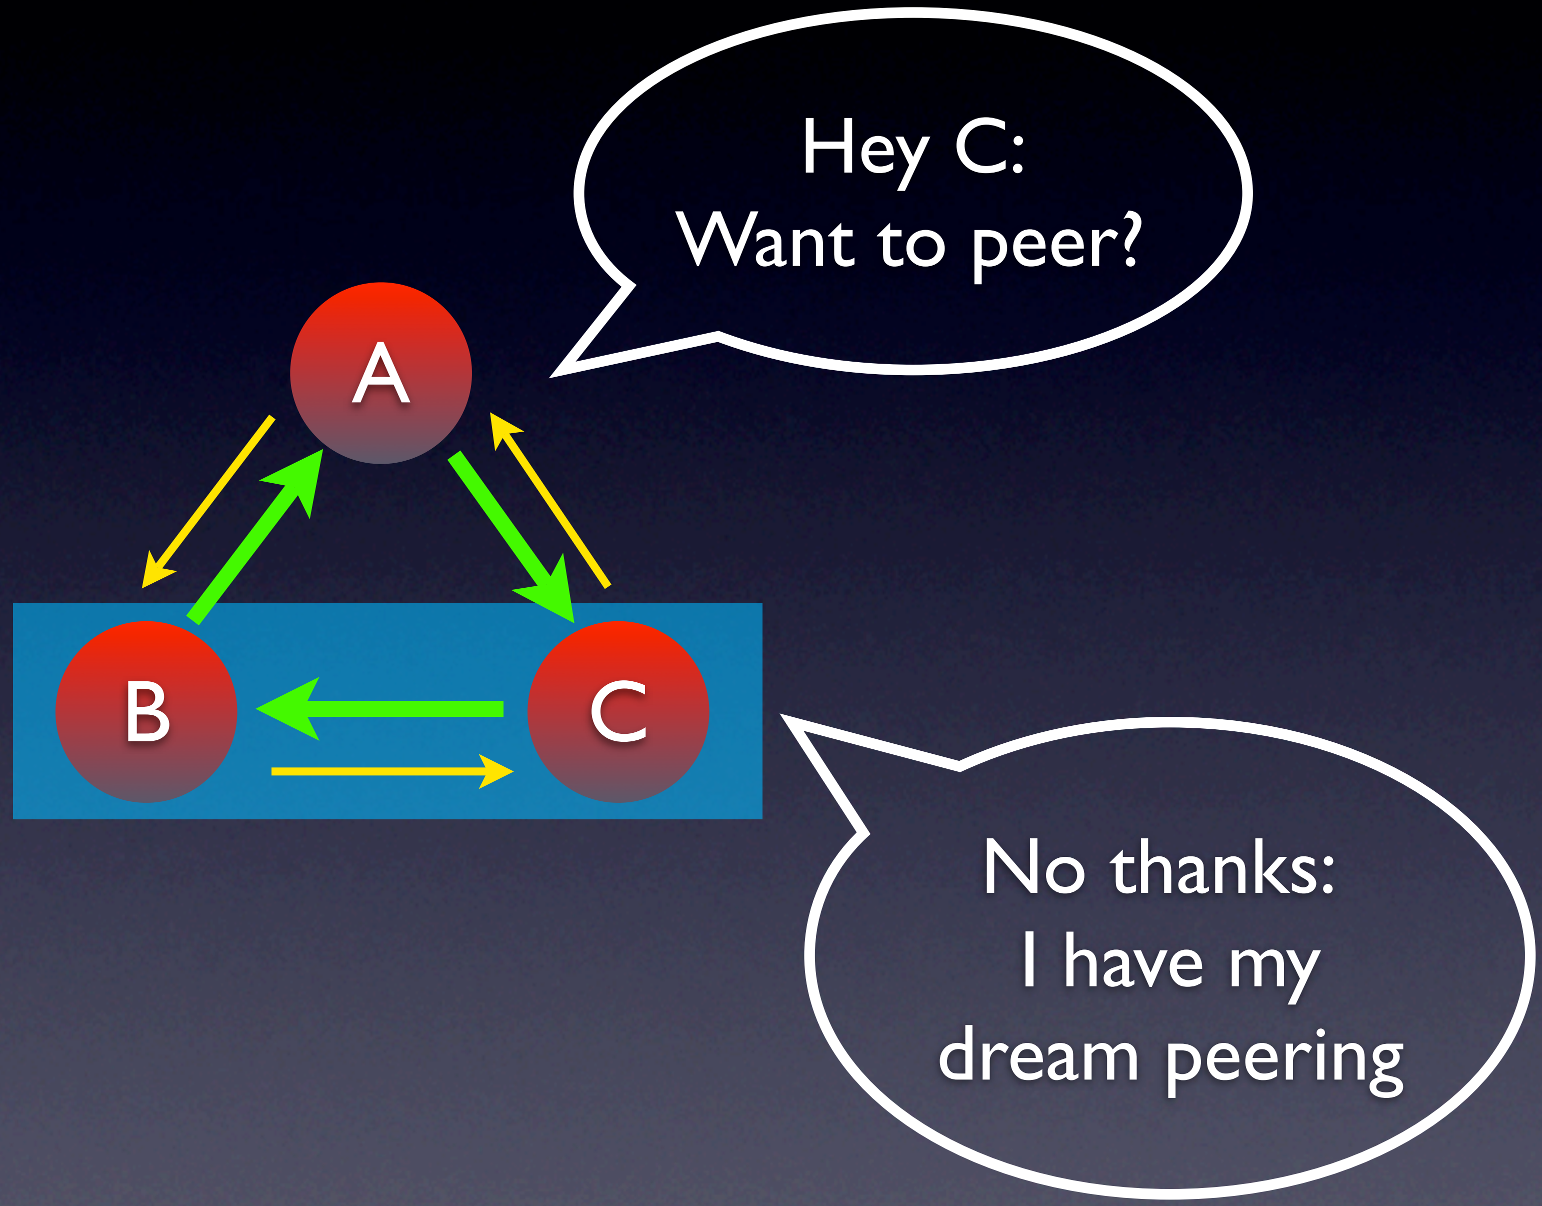

Preferred Acceptable SLA

Why violate an SLA?

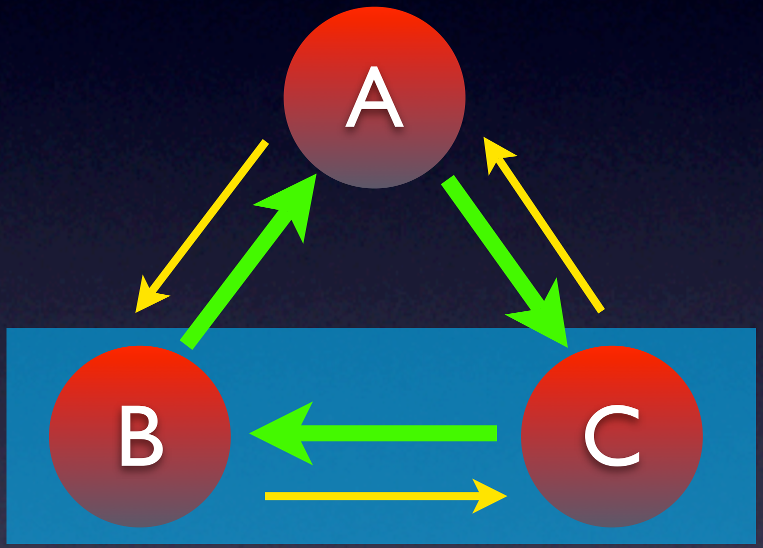

Preferred Acceptable SLA

Why violate an SLA?

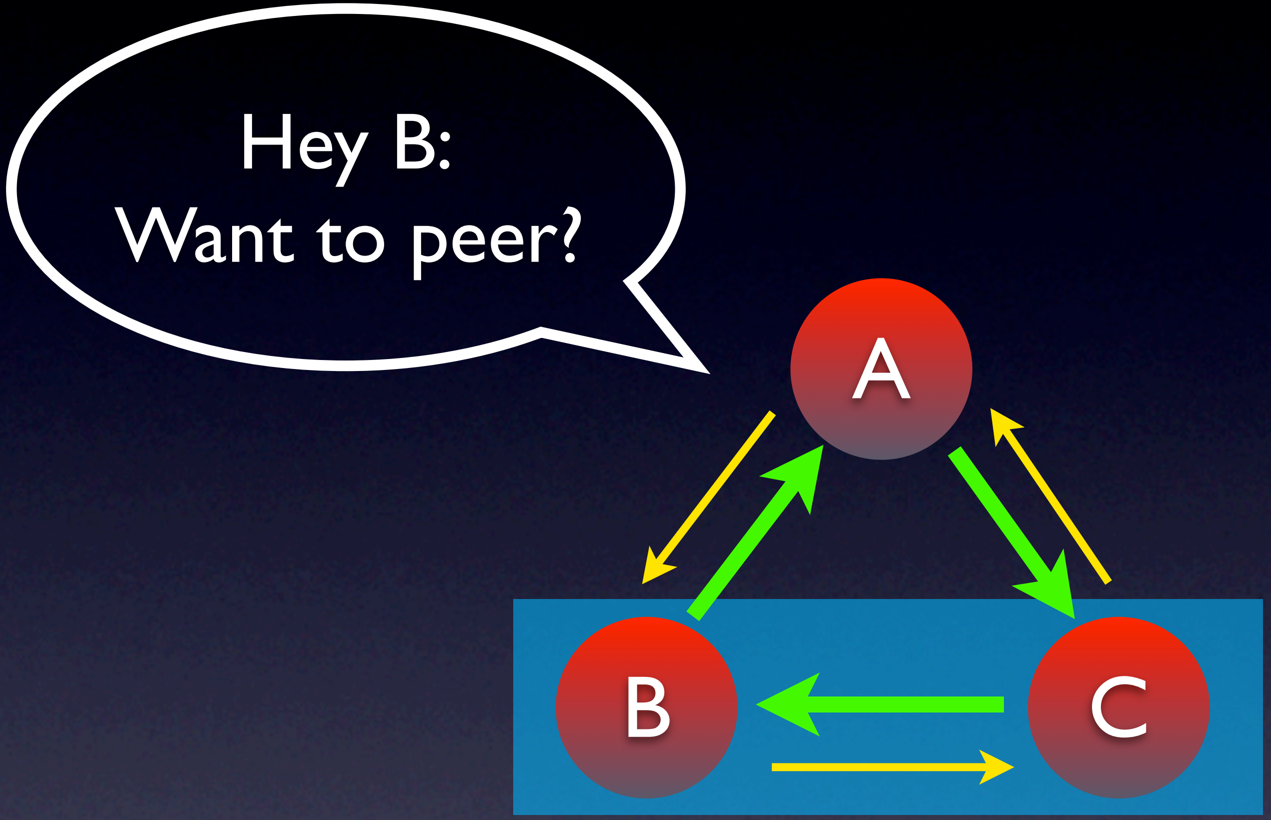

Preferred Acceptable SLA

Why violate an SLA?

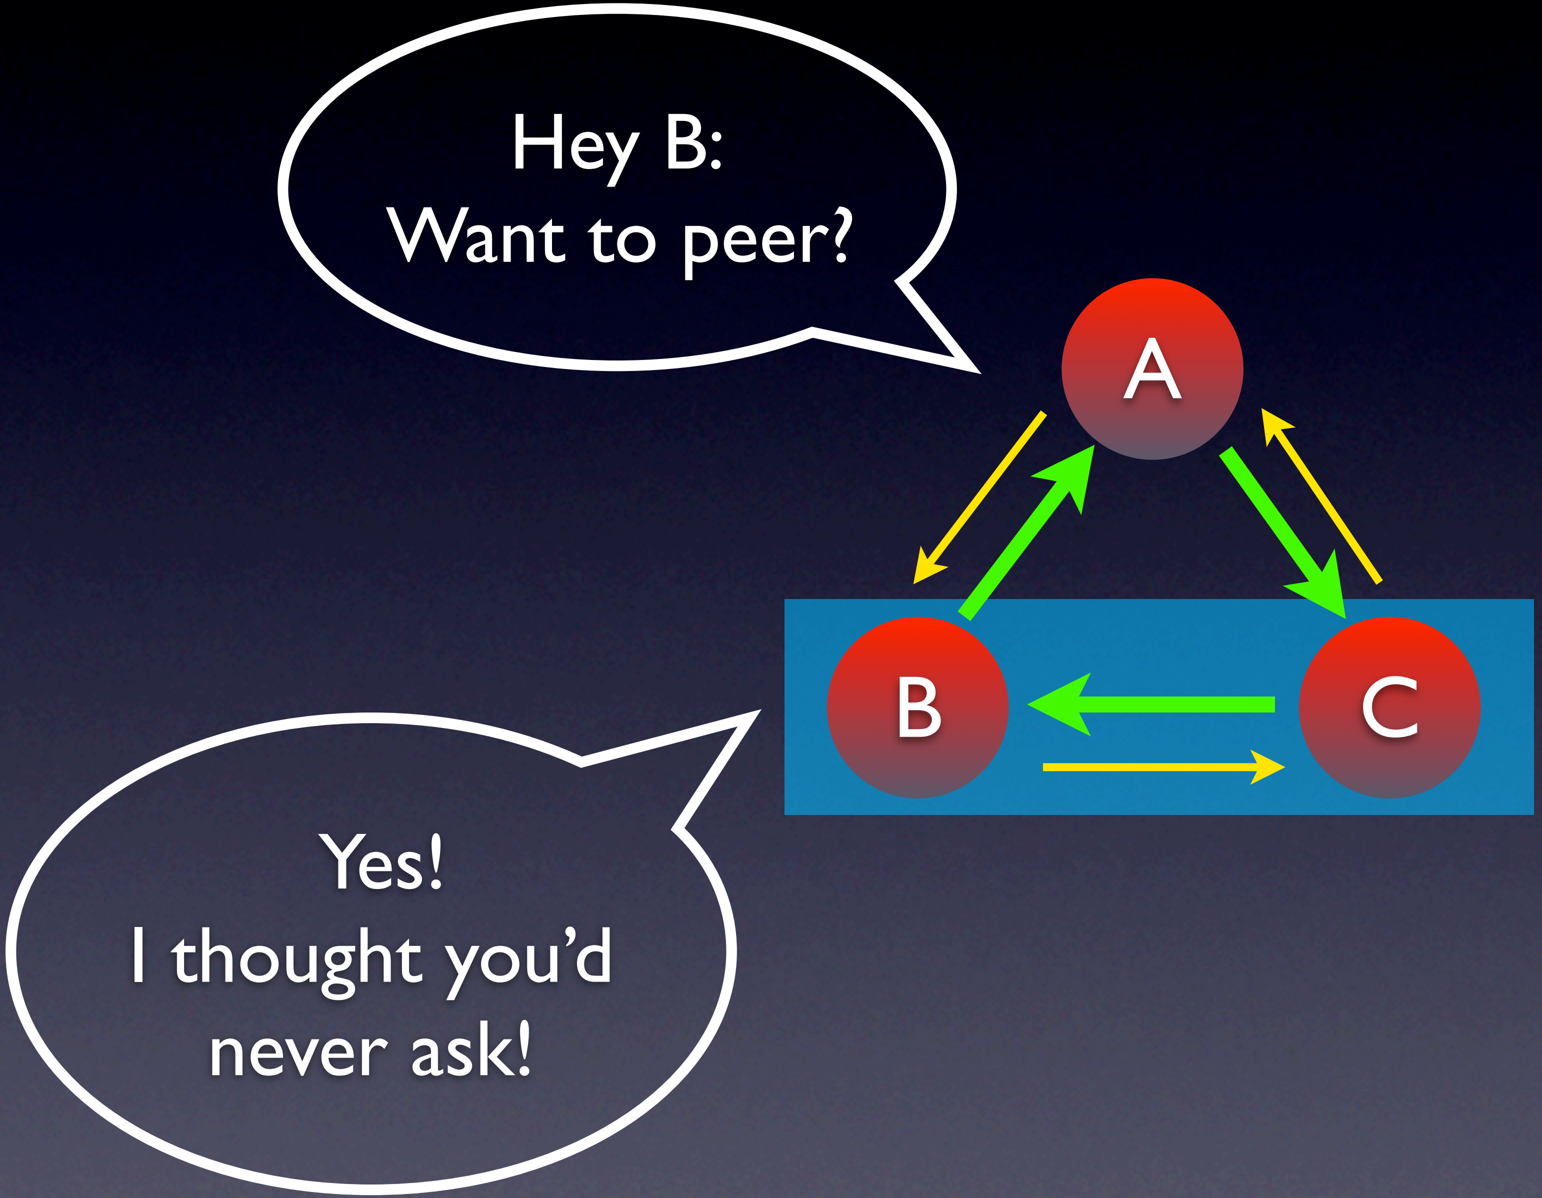

Preferred Acceptable SLA

Why violate an SLA?

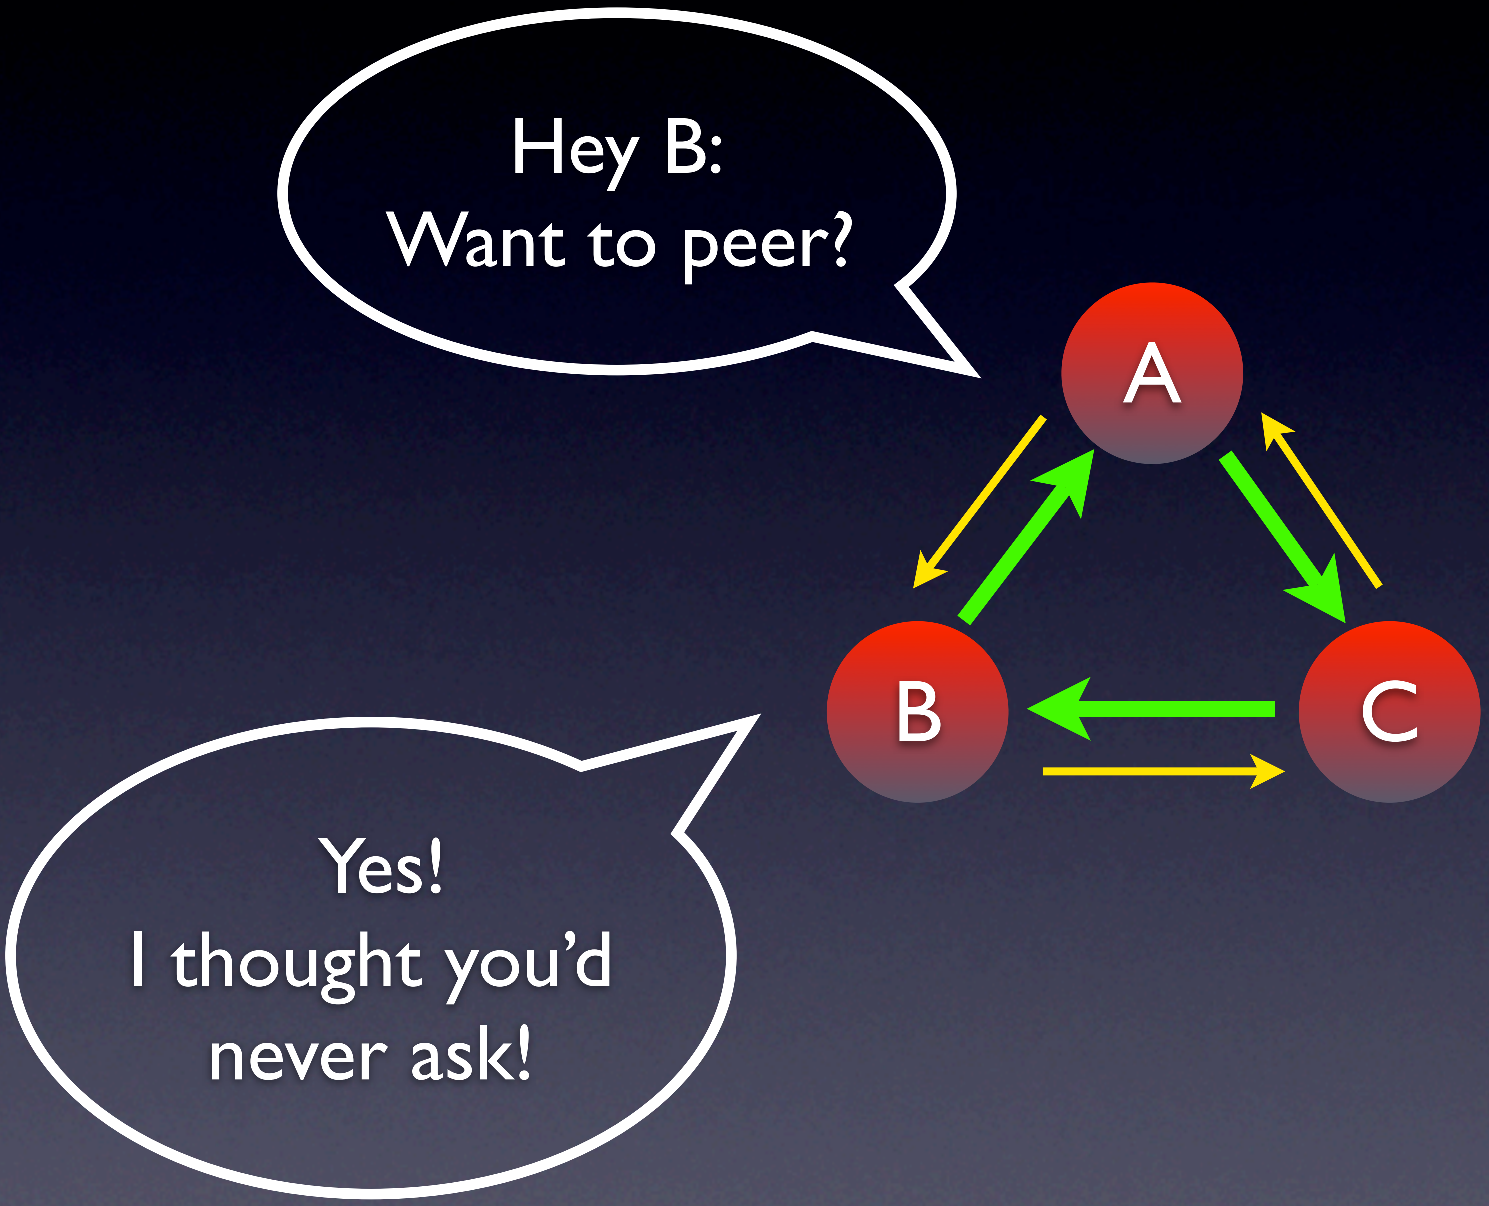

Preferred Acceptable SLA

Why violate an SLA?

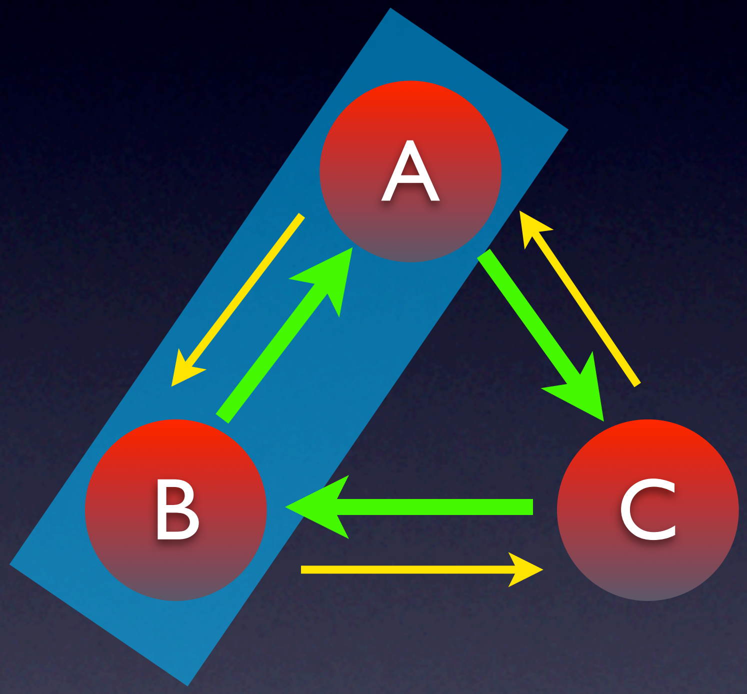

Preferred Acceptable SLA

Why violate an SLA?

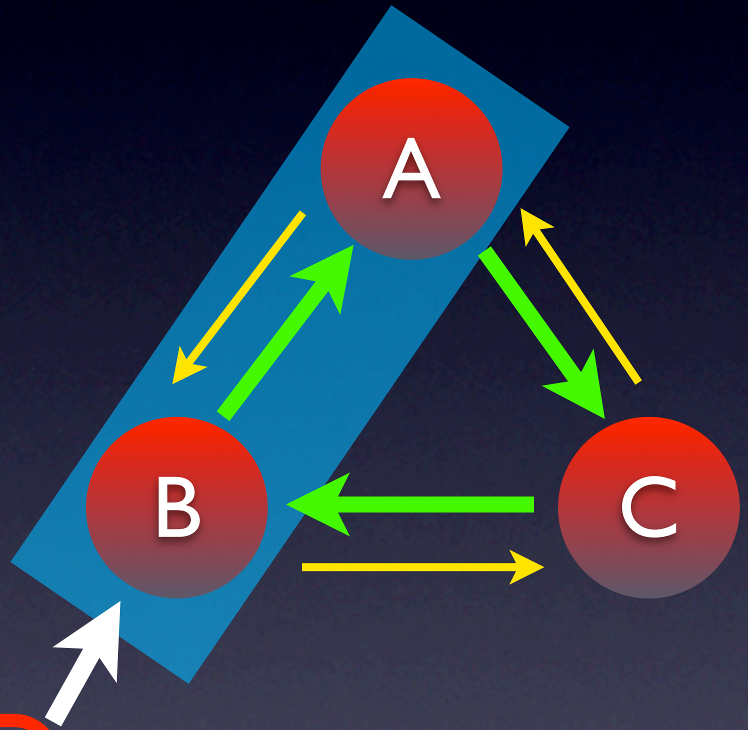

Violates SLA with C for an SLA with A

Preferred Acceptable SLA

Why violate an SLA?

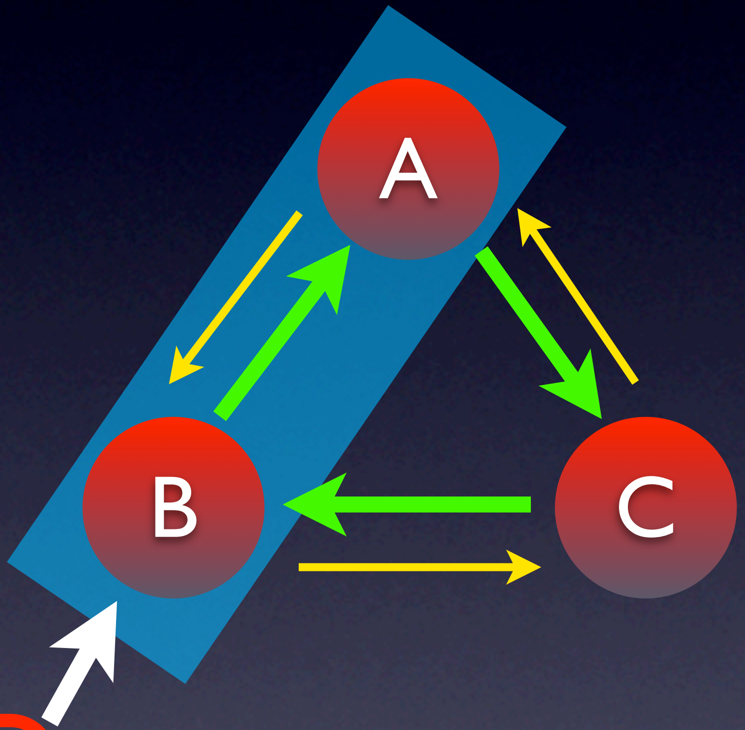

Violates SLA with C for an SLA with A

Hey A:
Want to peer?

Preferred Acceptable SLA

Why violate an SLA?

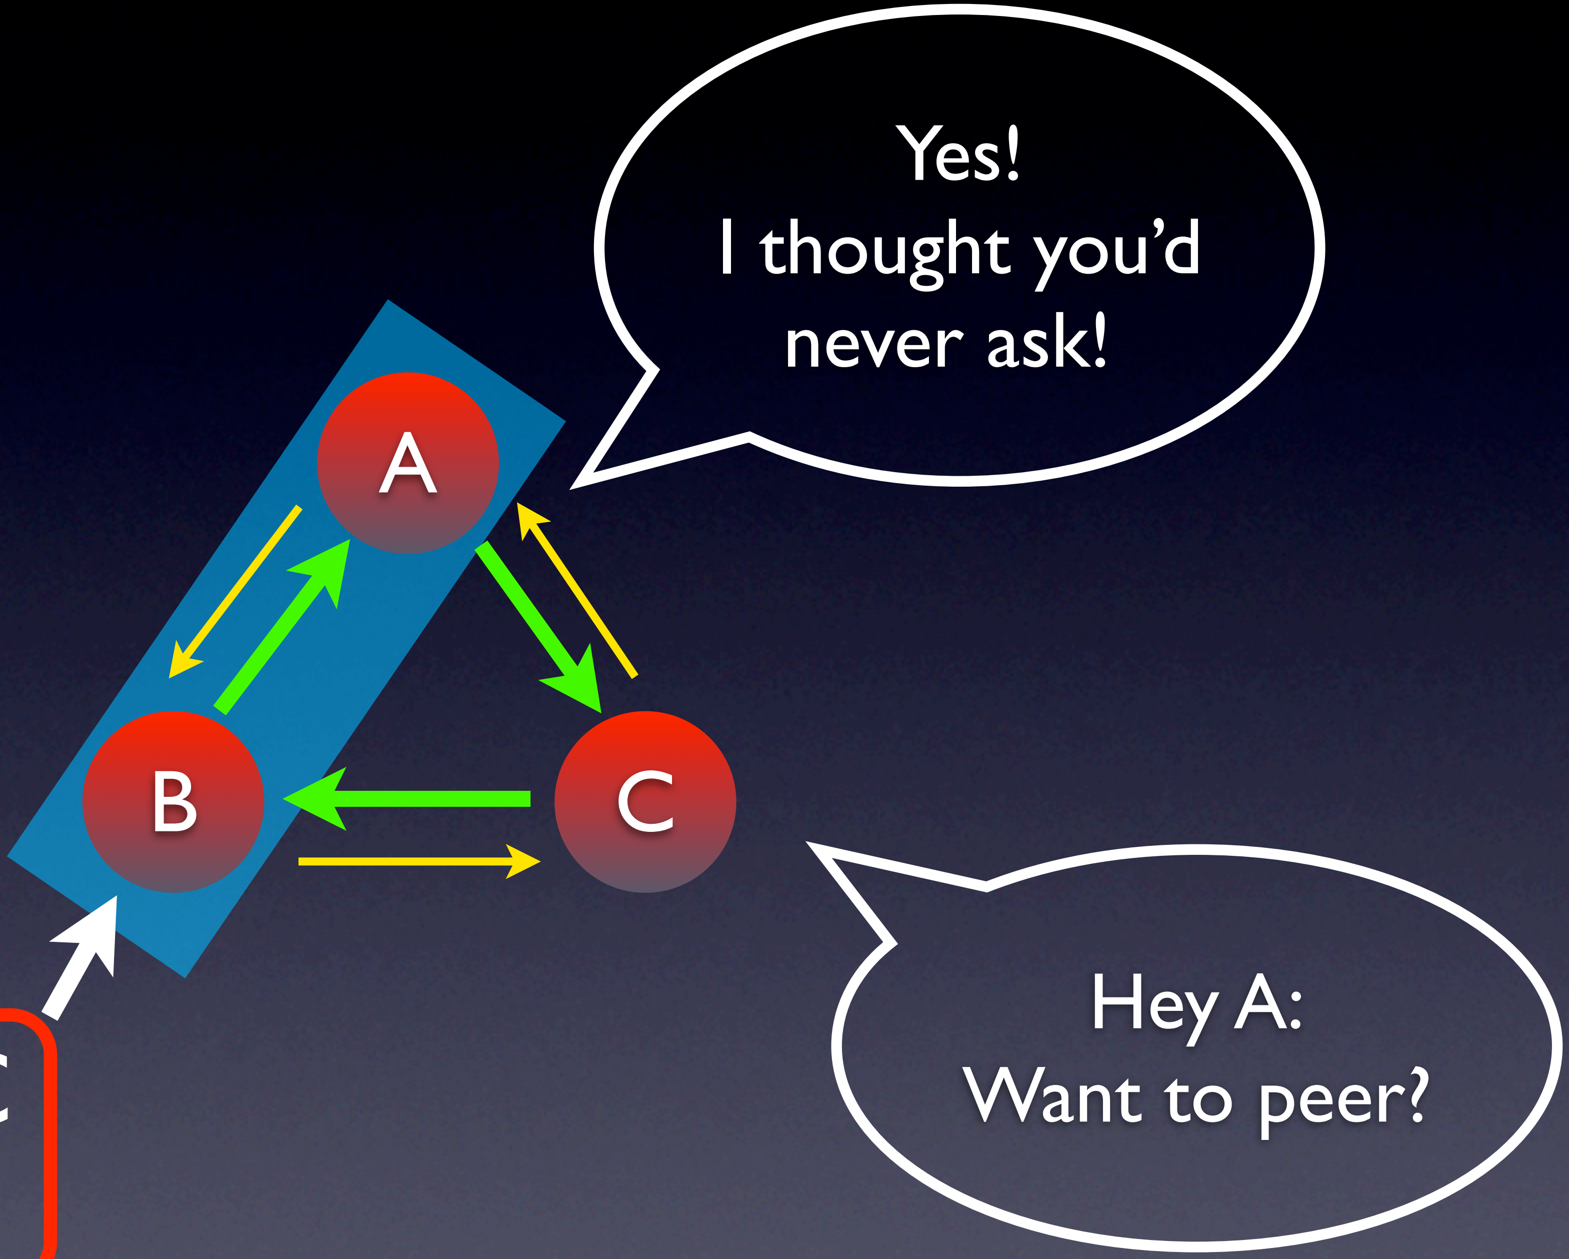

Violates SLA with C for an SLA with A

Preferred Acceptable SLA

Why violate an SLA?

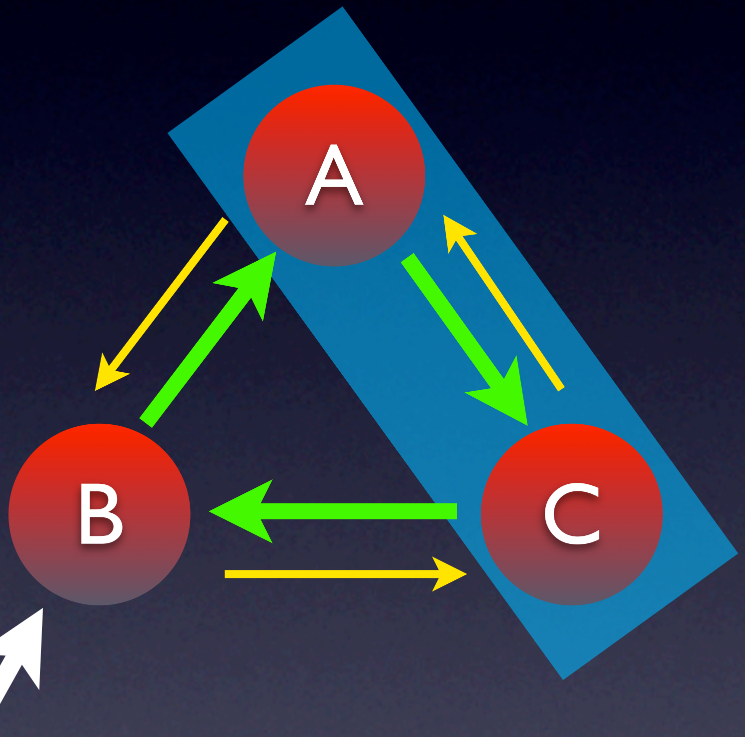

Violates SLA with C for an SLA with A

Preferred Acceptable SLA

Why violate an SLA?

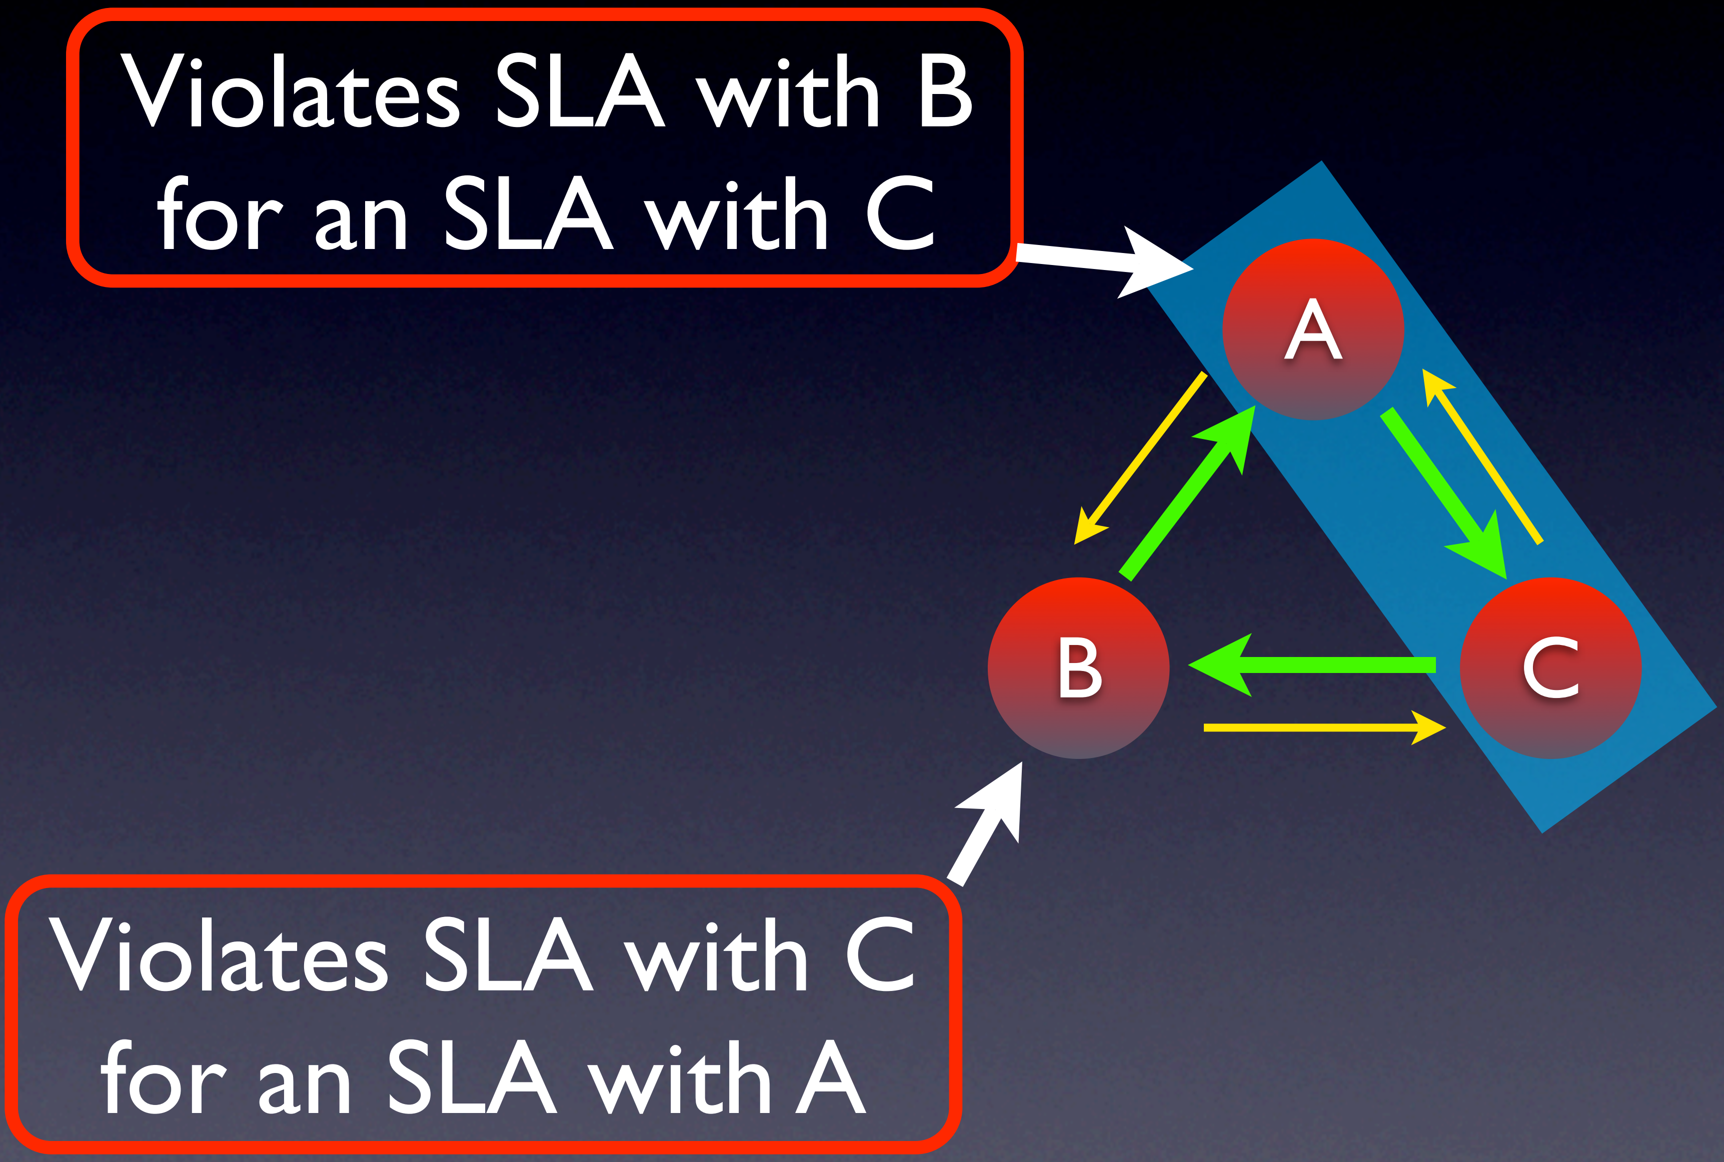

Violates SLA with B for an SLA with C

Violates SLA with C for an SLA with A

Preferred Acceptable SLA

Why violate an SLA?

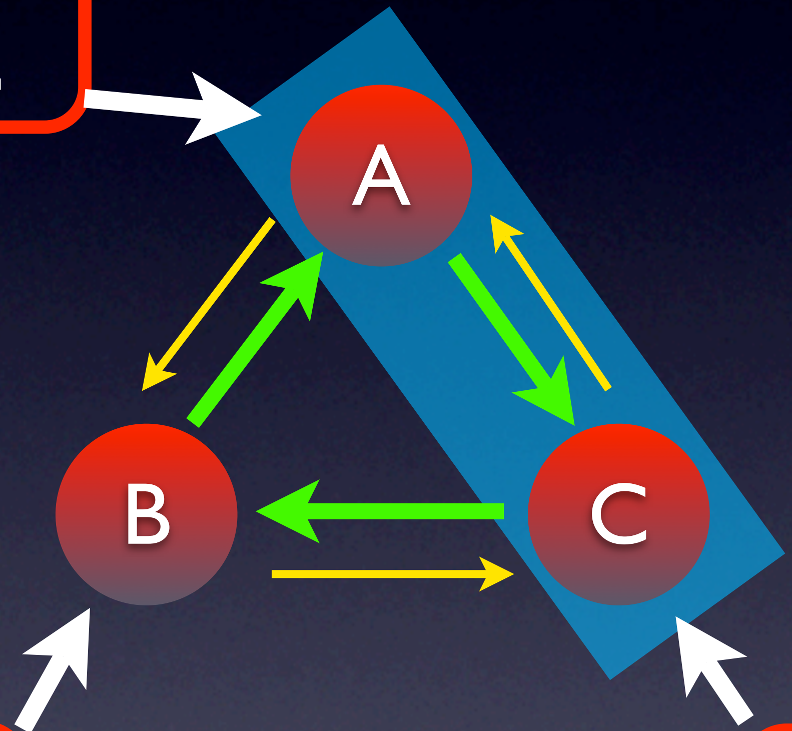

Violates SLA with B for an SLA with C

Violates SLA with C for an SLA with A

Violates SLA with B for an SLA with C

Preferred   Acceptable   SLA

Reacting to an SLA Violation

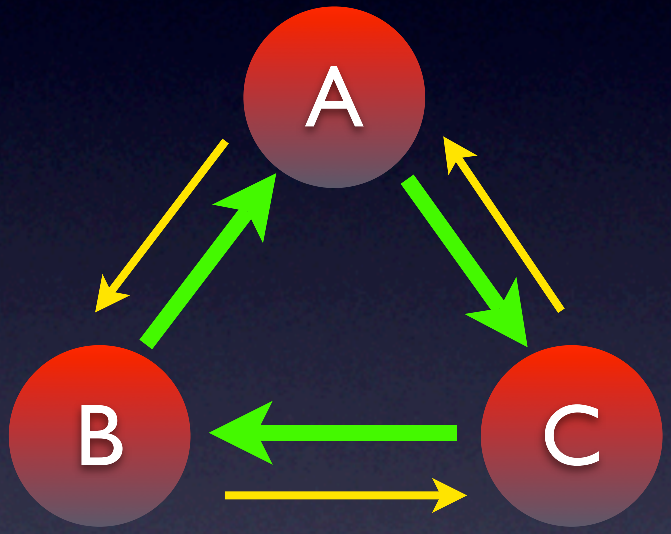

Reacting to an SLA Violation

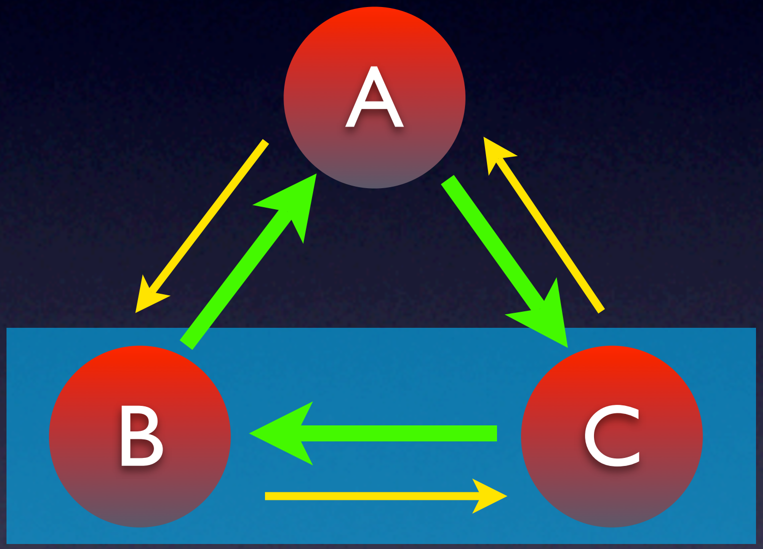

Reacting to an SLA Violation

Violates SLA with C for an SLA with A

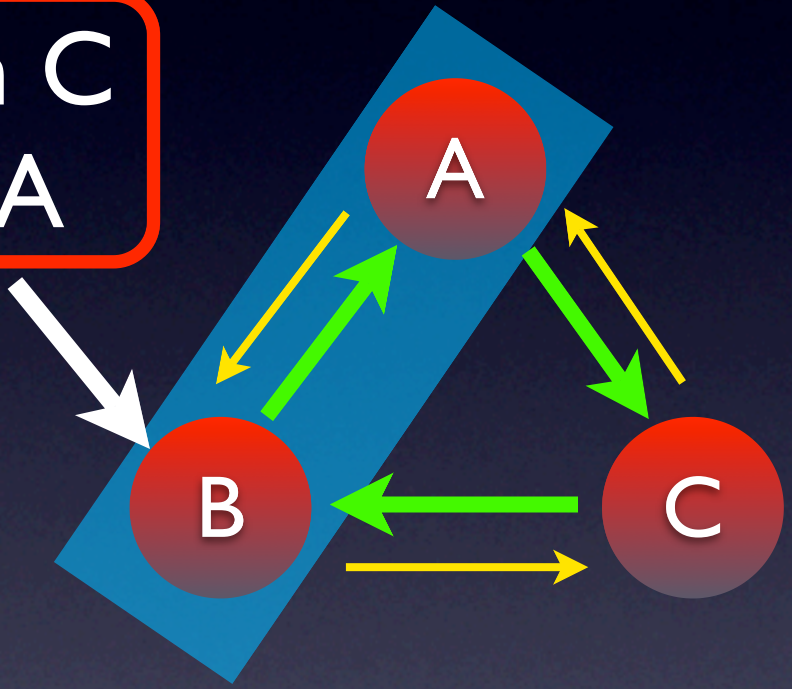

Reacting to an SLA Violation

Violates SLA with C for an SLA with A

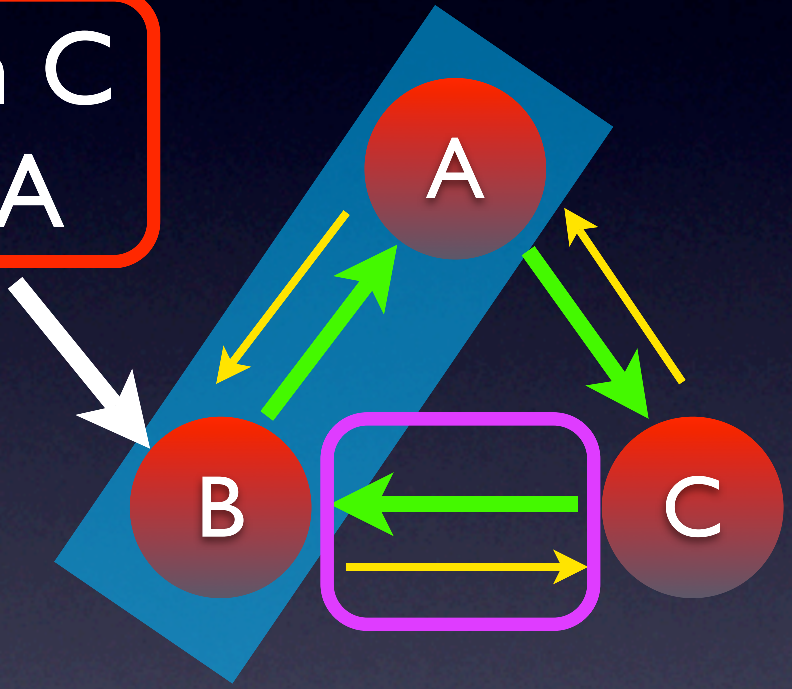

Why prefer someone who is willing to violate agreements with you?

Reacting to an SLA Violation

Violates SLA with C for an SLA with A

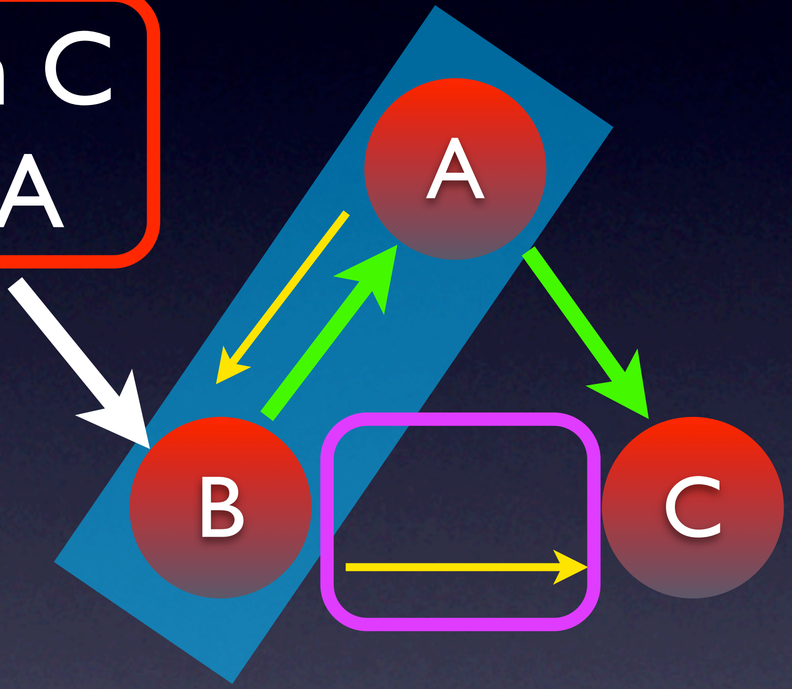

Why prefer someone who is willing to violate agreements with you?

Reacting to an SLA Violation

Violates SLA with C for an SLA with A

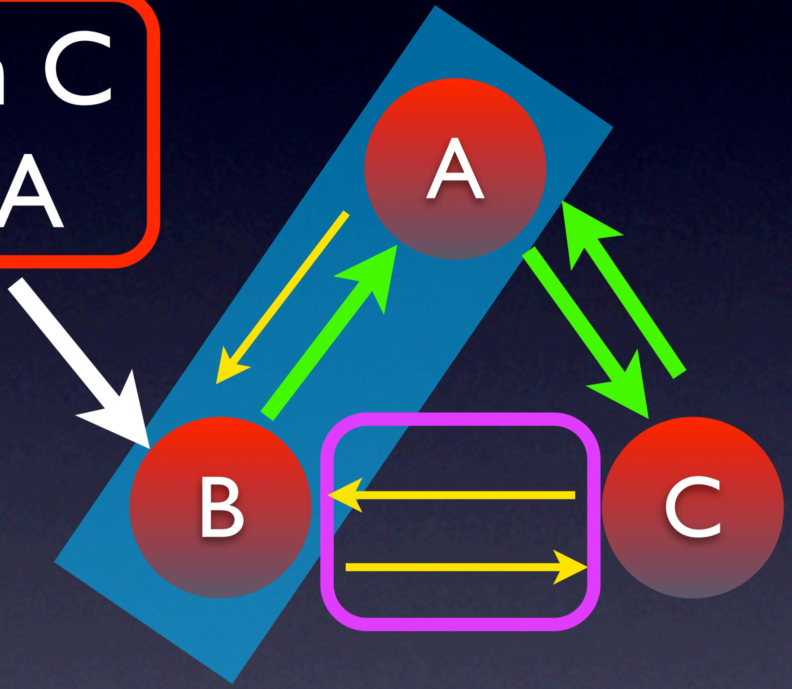

Why prefer someone who is willing to violate agreements with you?

Reacting to an SLA Violation

Violates SLA with C for an SLA with A

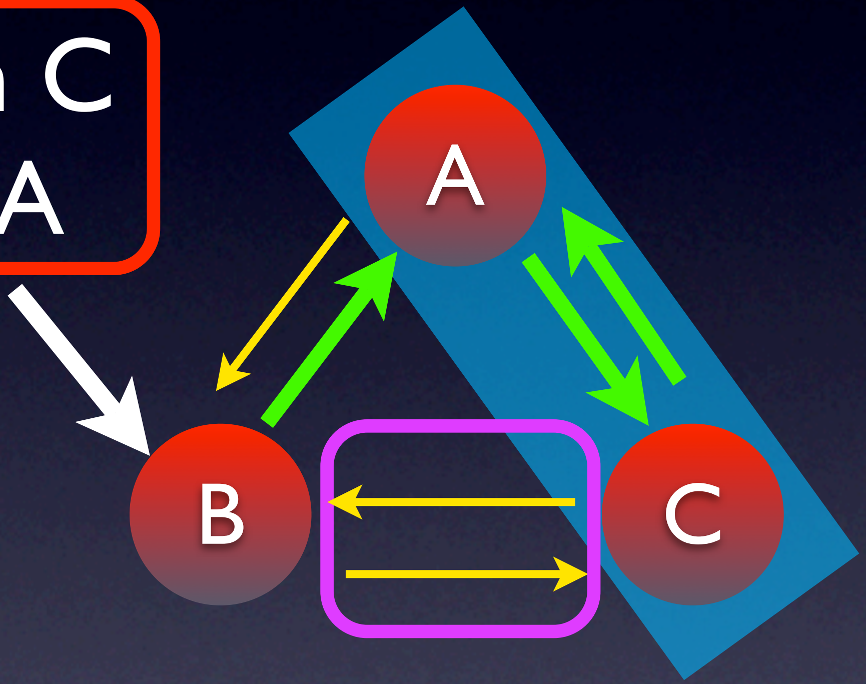

Why prefer someone who is willing to violate agreements with you?

Reacting to an SLA Violation

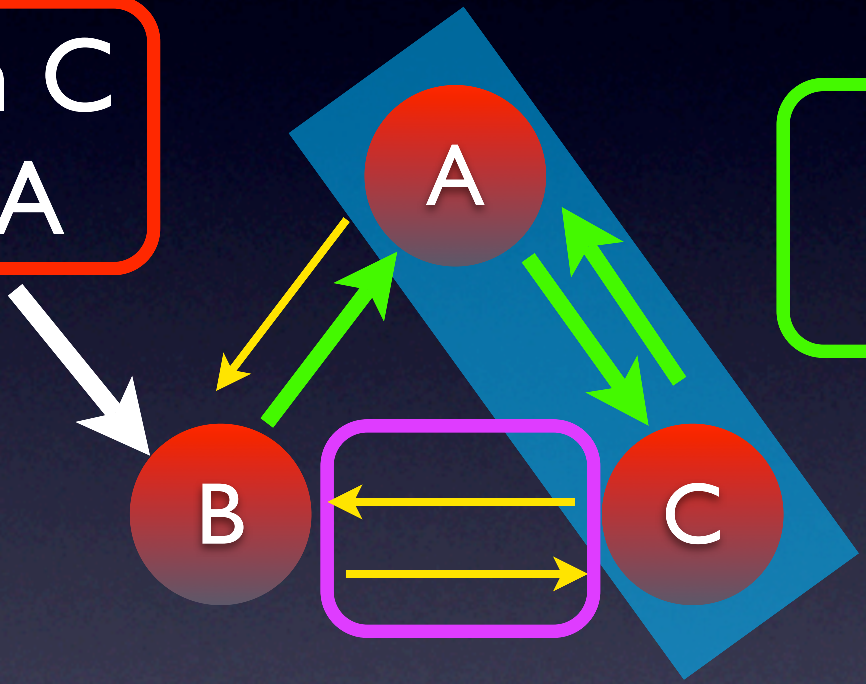

Violates SLA with C for an SLA with A

Long-lived peering

Why prefer someone who is willing to violate agreements with you?

A Reaction to SLA Violation

- Maintain per-neighbor confidence
  - What's the likelihood my neighbor will forward my packets?

- Let the confidence influence the selection of peerings

A Reaction to SLA Violation

- Maintain per-neighbor confidence
  - What’s the likelihood my neighbor will forward my packets?

- Let the confidence influence the selection of peerings

j cooperated

$c_{i,j} = c_{i,j} + c_{step}$

A Reaction to SLA Violation

- Maintain per-neighbor confidence
  - What’s the likelihood my neighbor will forward my packets?

- Let the confidence influence the selection of peerings

j cooperated

$$c_{i,j} = c_{i,j} + c_{step}$$

Local policy

A Reaction to SLA Violation

- Maintain per-neighbor confidence
  - What’s the likelihood my neighbor will forward my packets?

- Let the confidence influence the selection of peerings

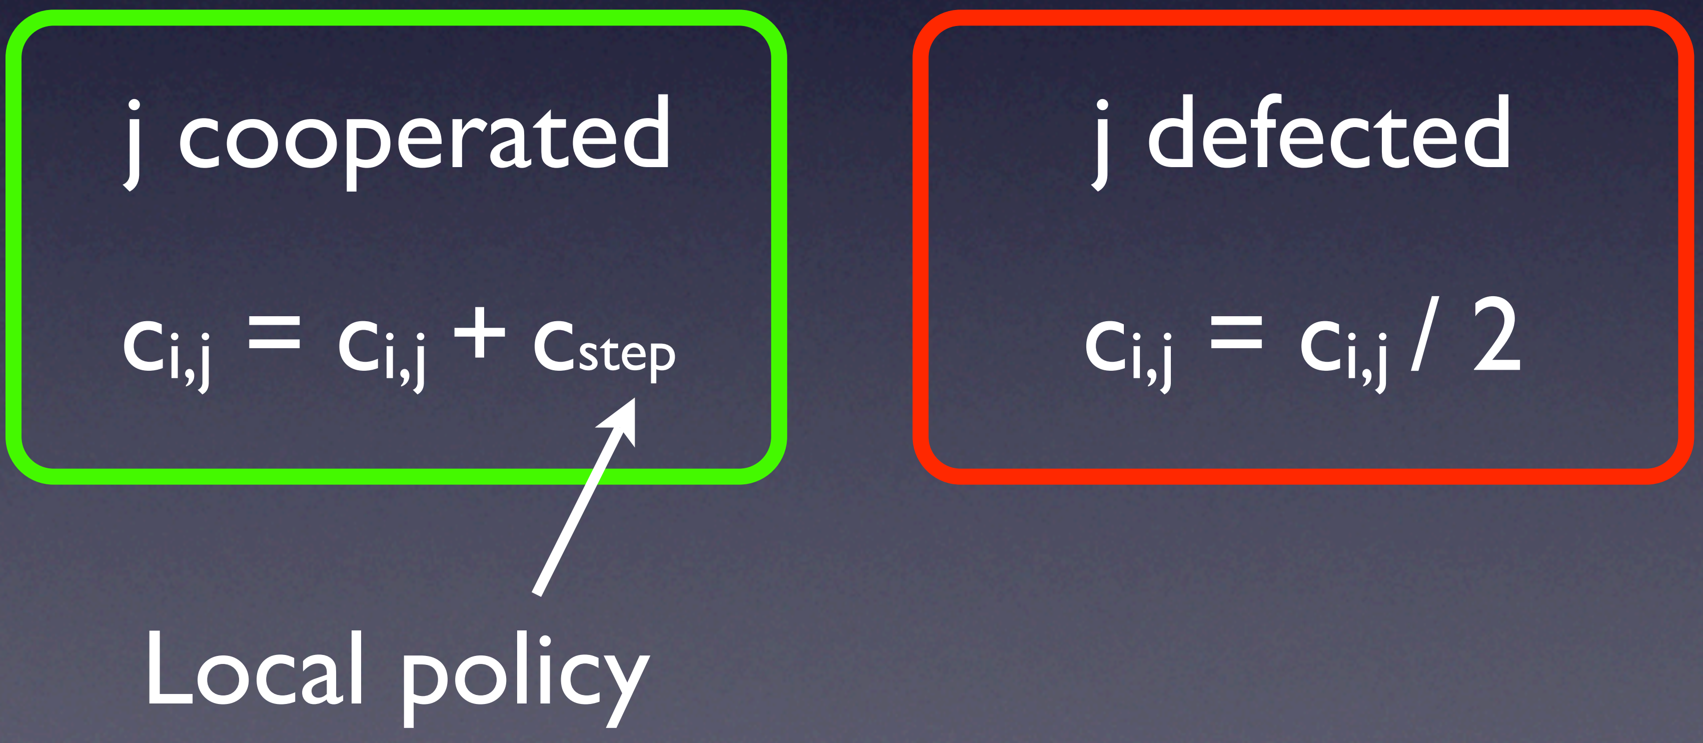

j cooperated

$c_{i,j} = c_{i,j} + c_{step}$

Local policy

j defected

$c_{i,j} = c_{i,j} / 2$

A Reaction to SLA Violation

- Maintain per-neighbor confidence
  - What’s the likelihood my neighbor will forward my packets?

- Let the confidence influence the selection of peerings

j cooperated

$c_{i,j} = c_{i,j} + c_{step}$

j defected

$c_{i,j} = c_{i,j} / 2$

No interaction

$c_{i,j} = c_{i,j} + \frac{c_{step}}{s_{i,j} + 2}$

Local policy

A Reaction to SLA Violation

- Maintain per-neighbor confidence
  - What’s the likelihood my neighbor will forward my packets?

- Let the confidence influence the selection of peerings

j cooperated

$$c_{i,j} = c_{i,j} + c_{step}$$

j defected

$$c_{i,j} = c_{i,j} / 2$$

No interaction

$$c_{i,j} = c_{i,j} + \frac{c_{step}}{s_{i,j} + 2}$$

Local policy

Number of times *j* violated an SLA with *i*

A Reaction to SLA Violation

j cooperated

$$c_{i,j} = c_{i,j} + c_{step}$$

j defected

$$c_{i,j} = c_{i,j} / 2$$

No interaction

$$c_{i,j} = c_{i,j} + \frac{c_{step}}{s_{i,j} + 2}$$

A Reaction to SLA Violation

j cooperated

$c_{i,j} = c_{i,j} + c_{step}$

Confidence builds slowly

j defected

$c_{i,j} = c_{i,j} / 2$

but can diminish quickly.

No interaction

$$c_{i,j} = c_{i,j} + \frac{c_{step}}{s_{i,j} + 2}$$

Time heals all wounds.

Contributions

- The selfish routing overlay problem
- Mechanisms to achieve long-lived peerings
- Simulation study

Contributions

- The selfish routing overlay problem
- Mechanisms to achieve long-lived peerings
- Simulation study

Evaluation

- Does this mechanism promote cooperation?
- How sensitive are peerings to confidence?
- Is there a correct amount of confidence to have?

Peering lifetimes

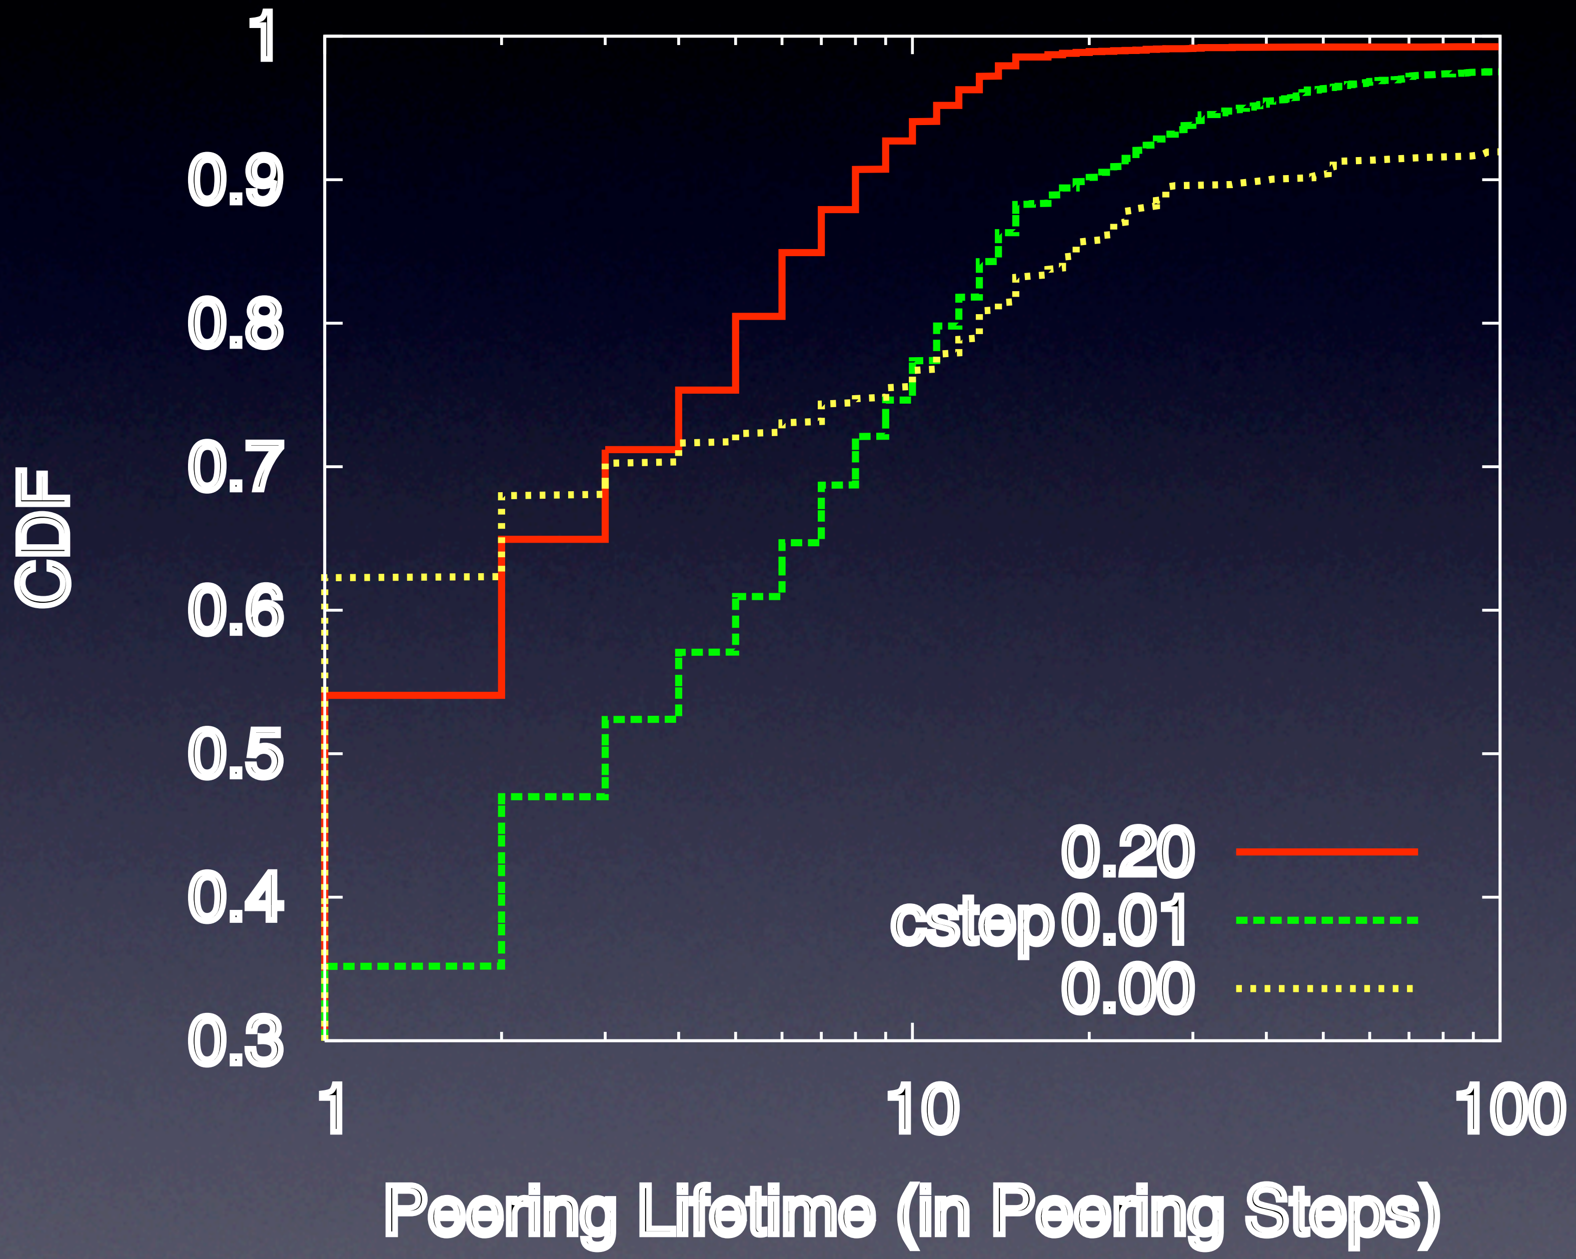

Peerings last longer if you never allow others to regain your confidence

Number of peerings

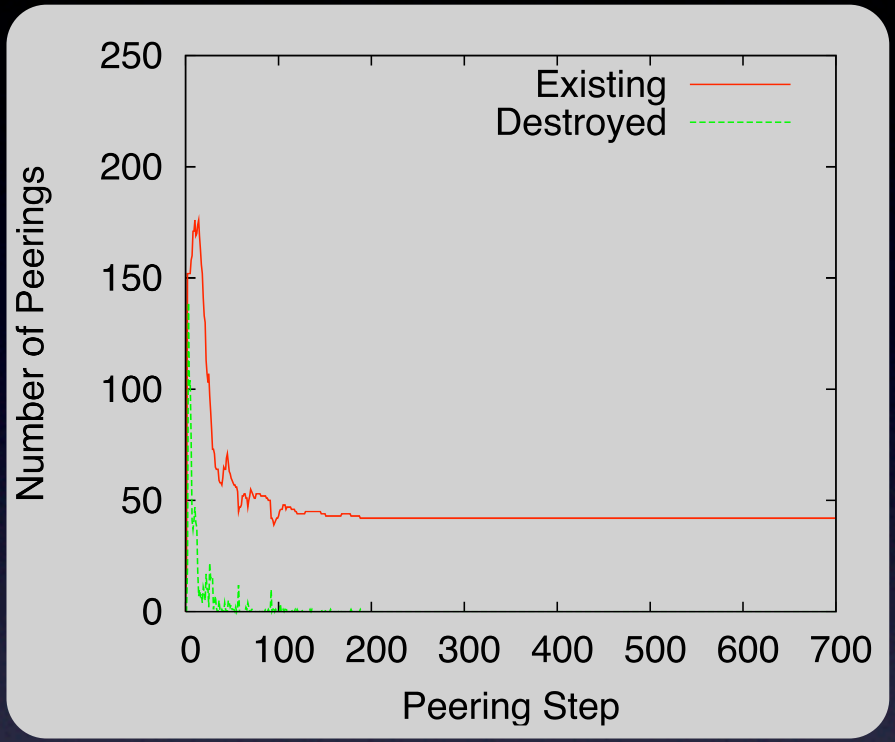

$C_{step} = 0$

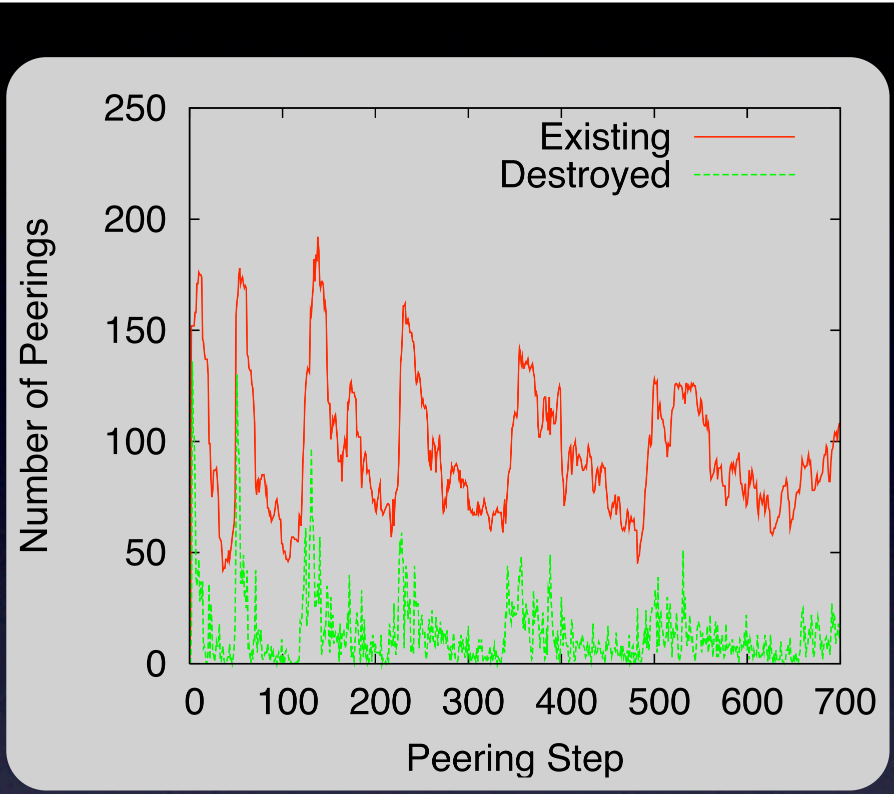

$C_{step} = 0.01$

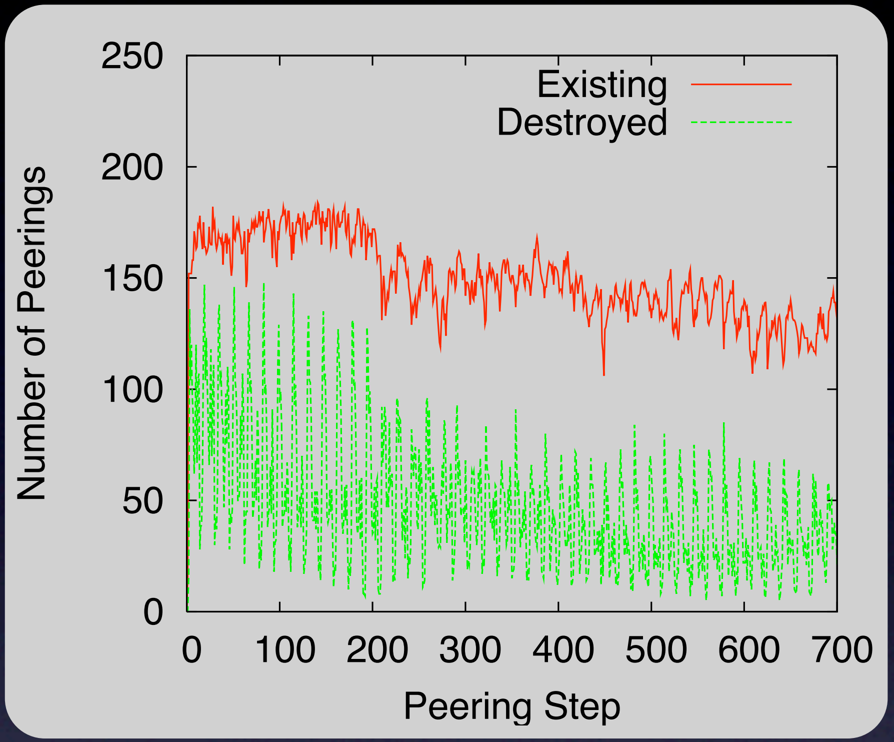

$C_{step} = 0.2$

Number of peerings

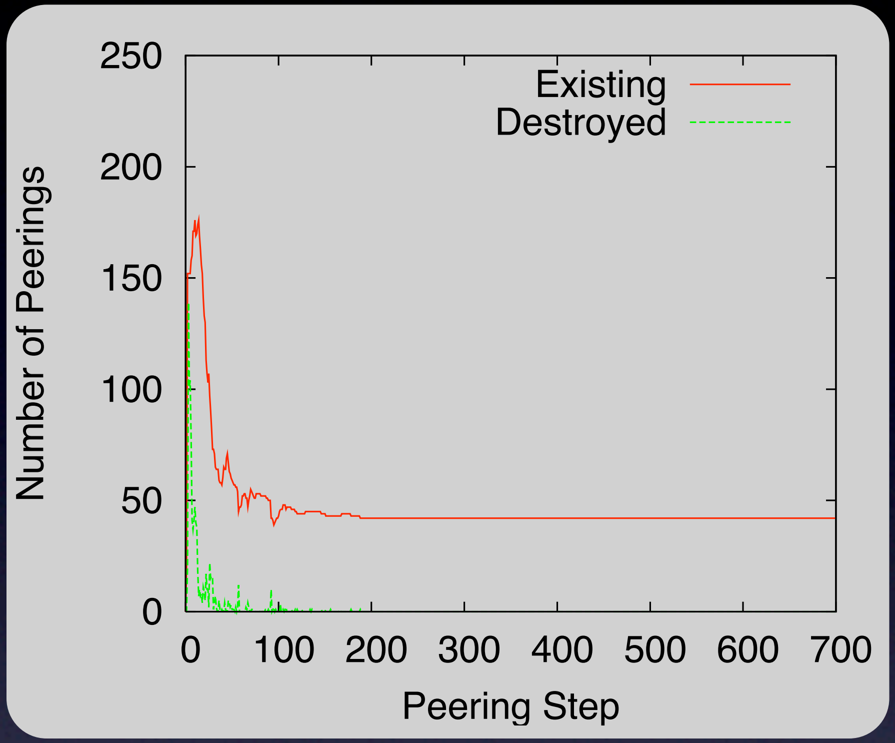

$C_{step} = 0$

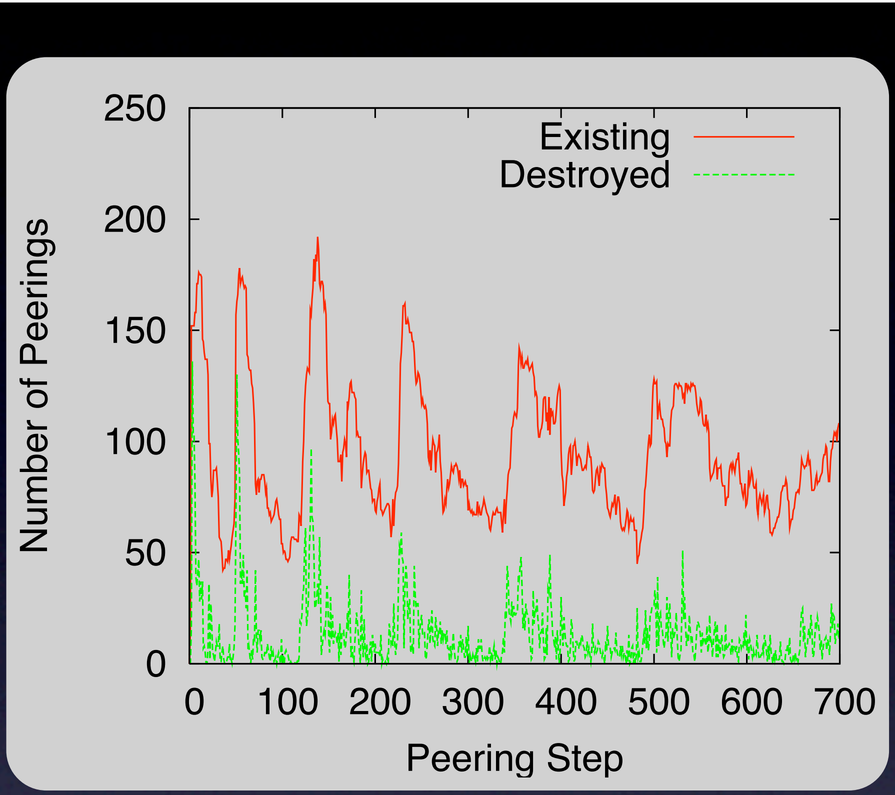

$C_{step} = 0.01$

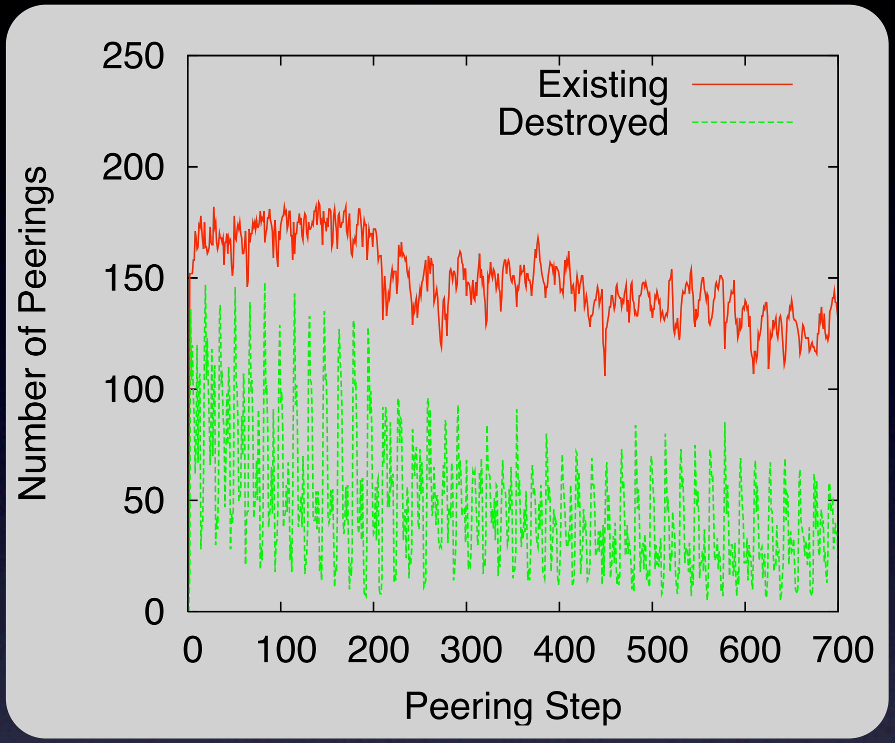

$C_{step} = 0.2$

Number of peerings

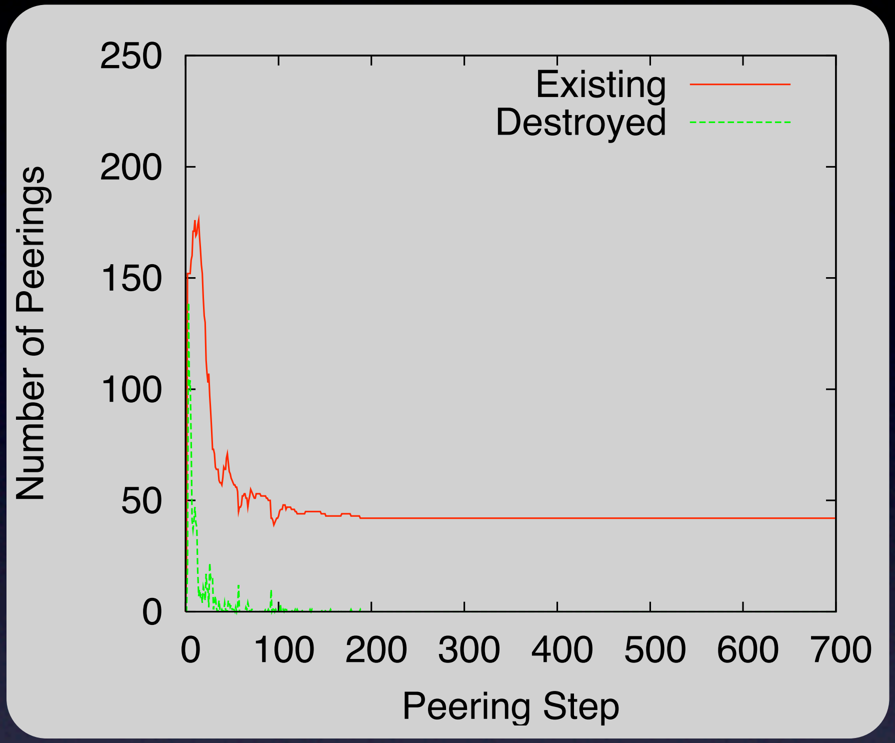

$C_{step} = 0$

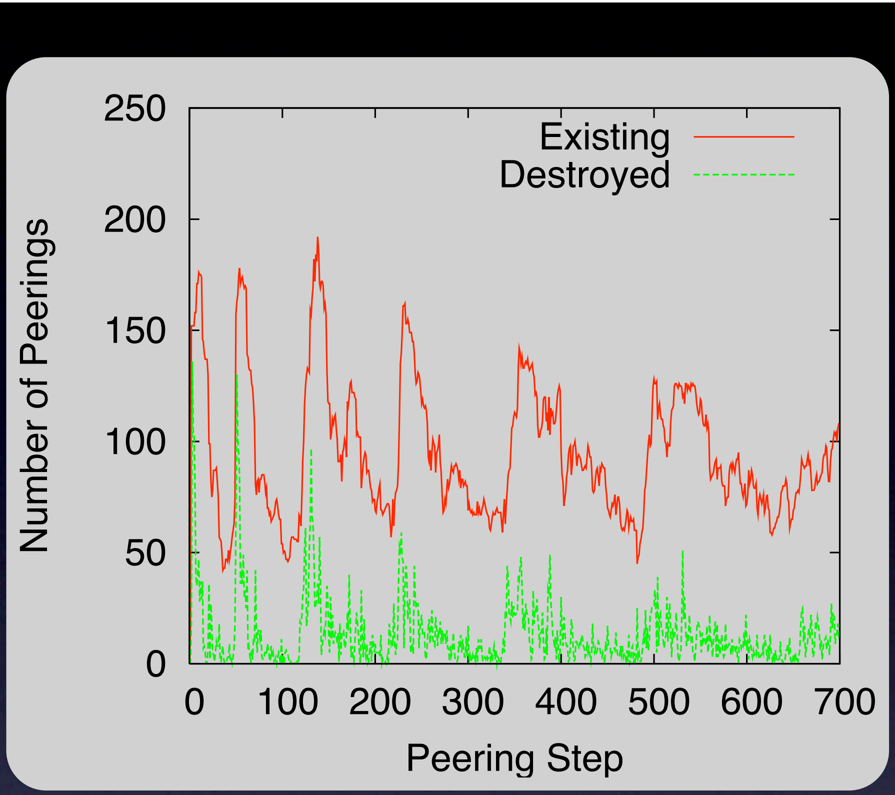

$C_{step} = 0.01$

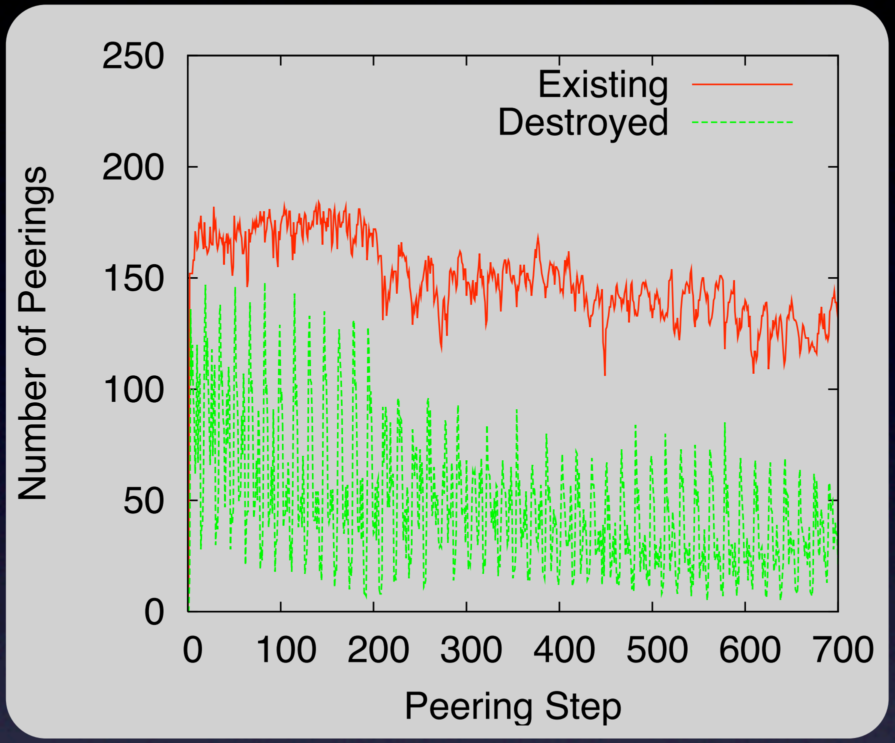

$C_{step} = 0.2$

More peerings, shorter life

Motivating Participation

- Per-peer, dynamic SLAs
  - Express demands
  - Express expectations

- Avoid, don't punish
  - Find someone else or fall back on direct path

- Explicitly incorporate confidence in others
  - Grows slowly
  - Goes away quickly
  - Time heals all wounds

Overselling Resources

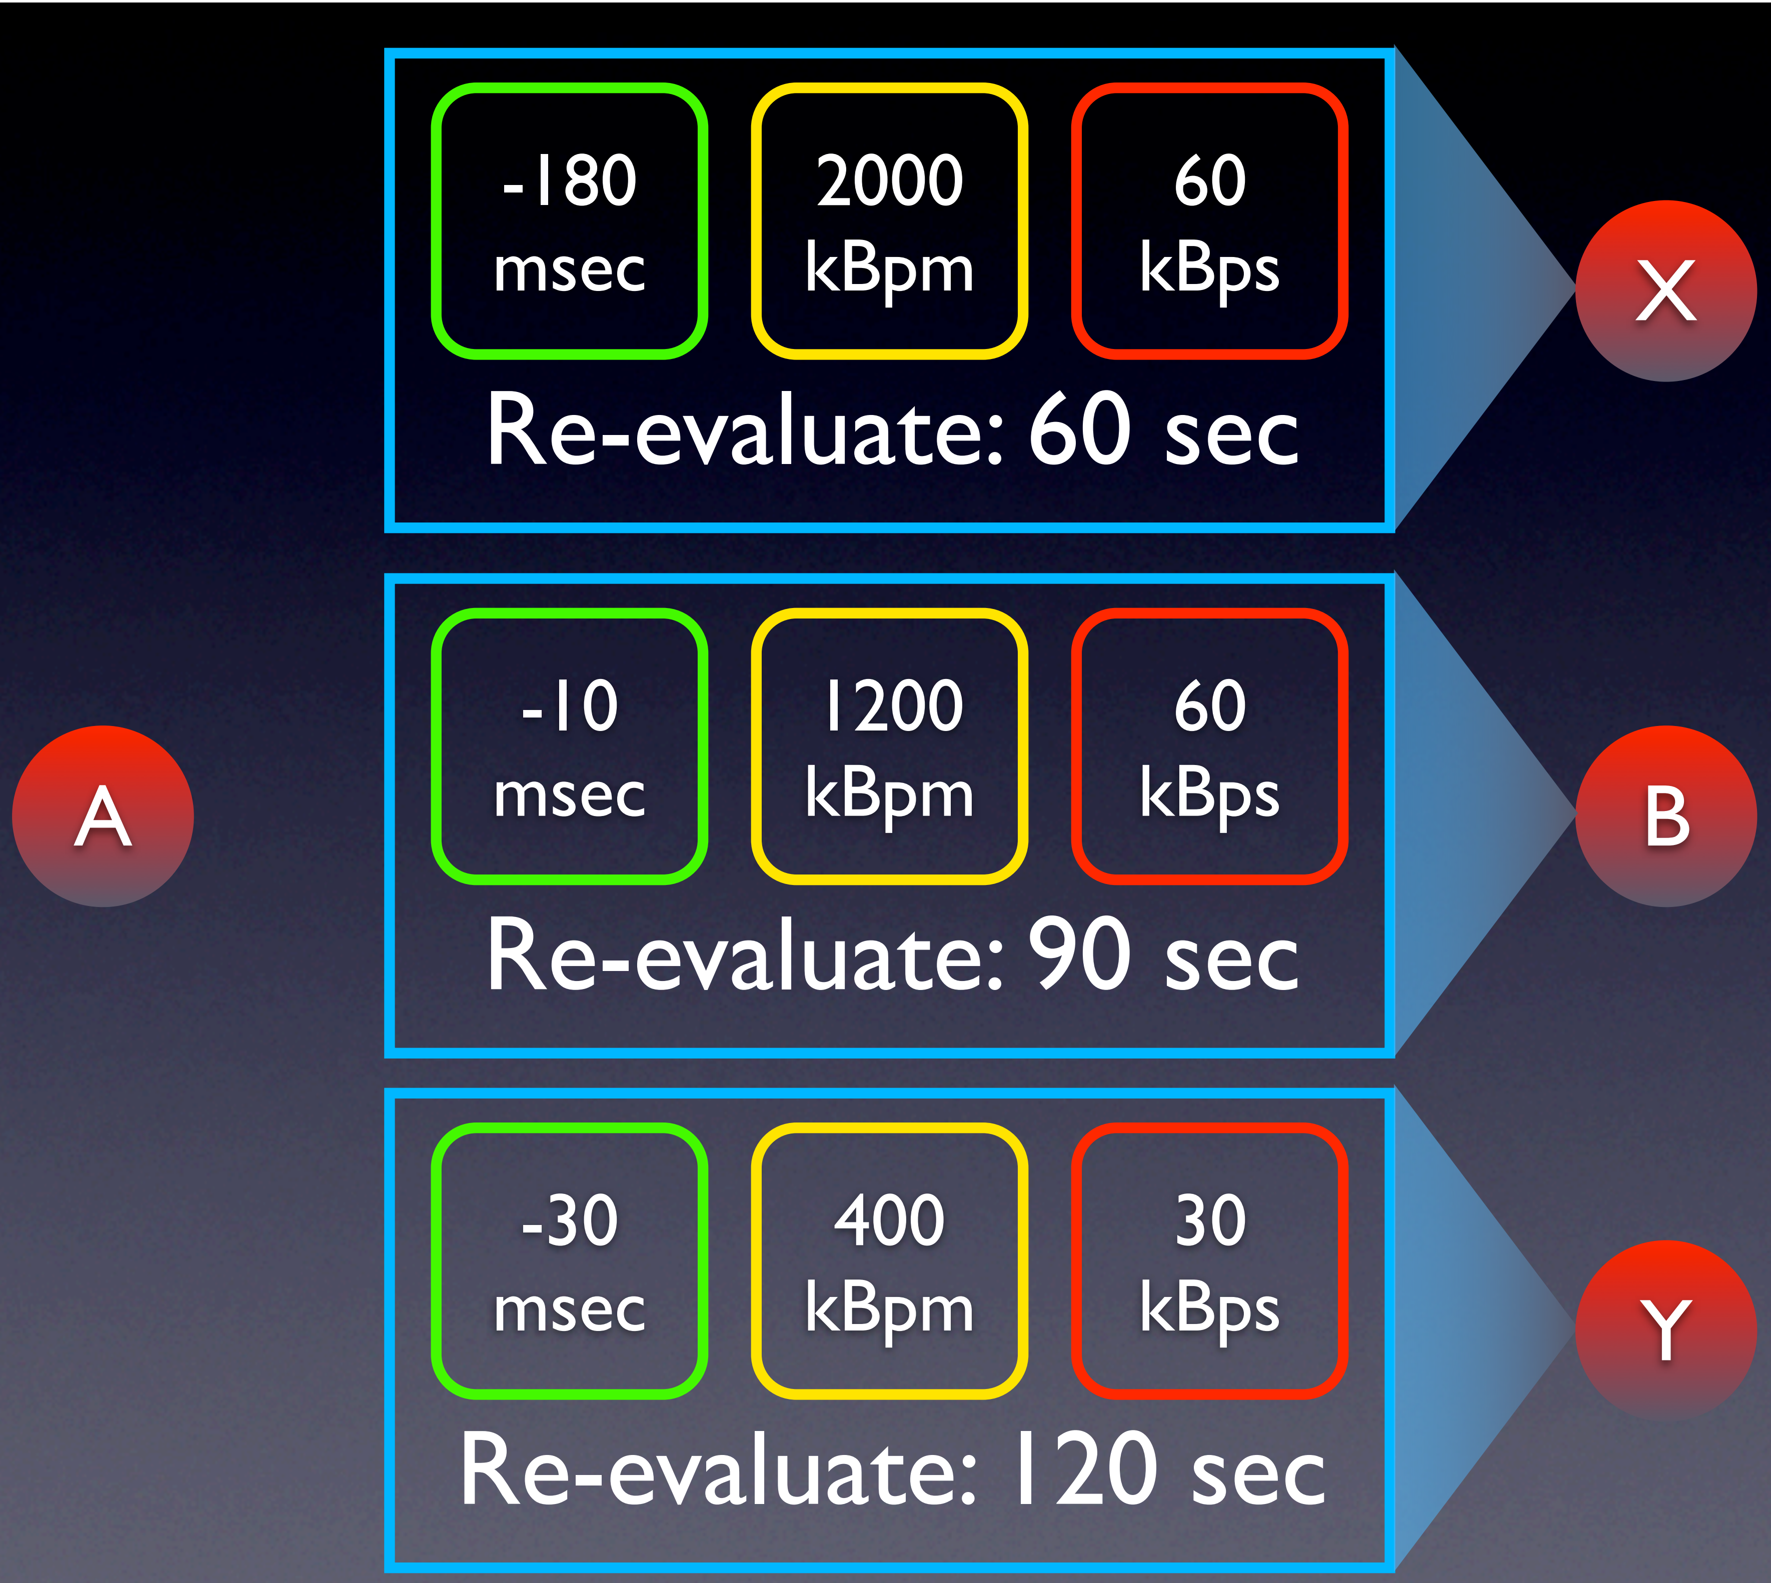

A

-180 msec | 2000 kBpm | 60 kBps

Re-evaluate: 60 sec

X

-10 msec | 1200 kBpm | 60 kBps

Re-evaluate: 90 sec

B

-30 msec | 400 kBpm | 30 kBps

Re-evaluate: 120 sec

Y

Overselling Resources

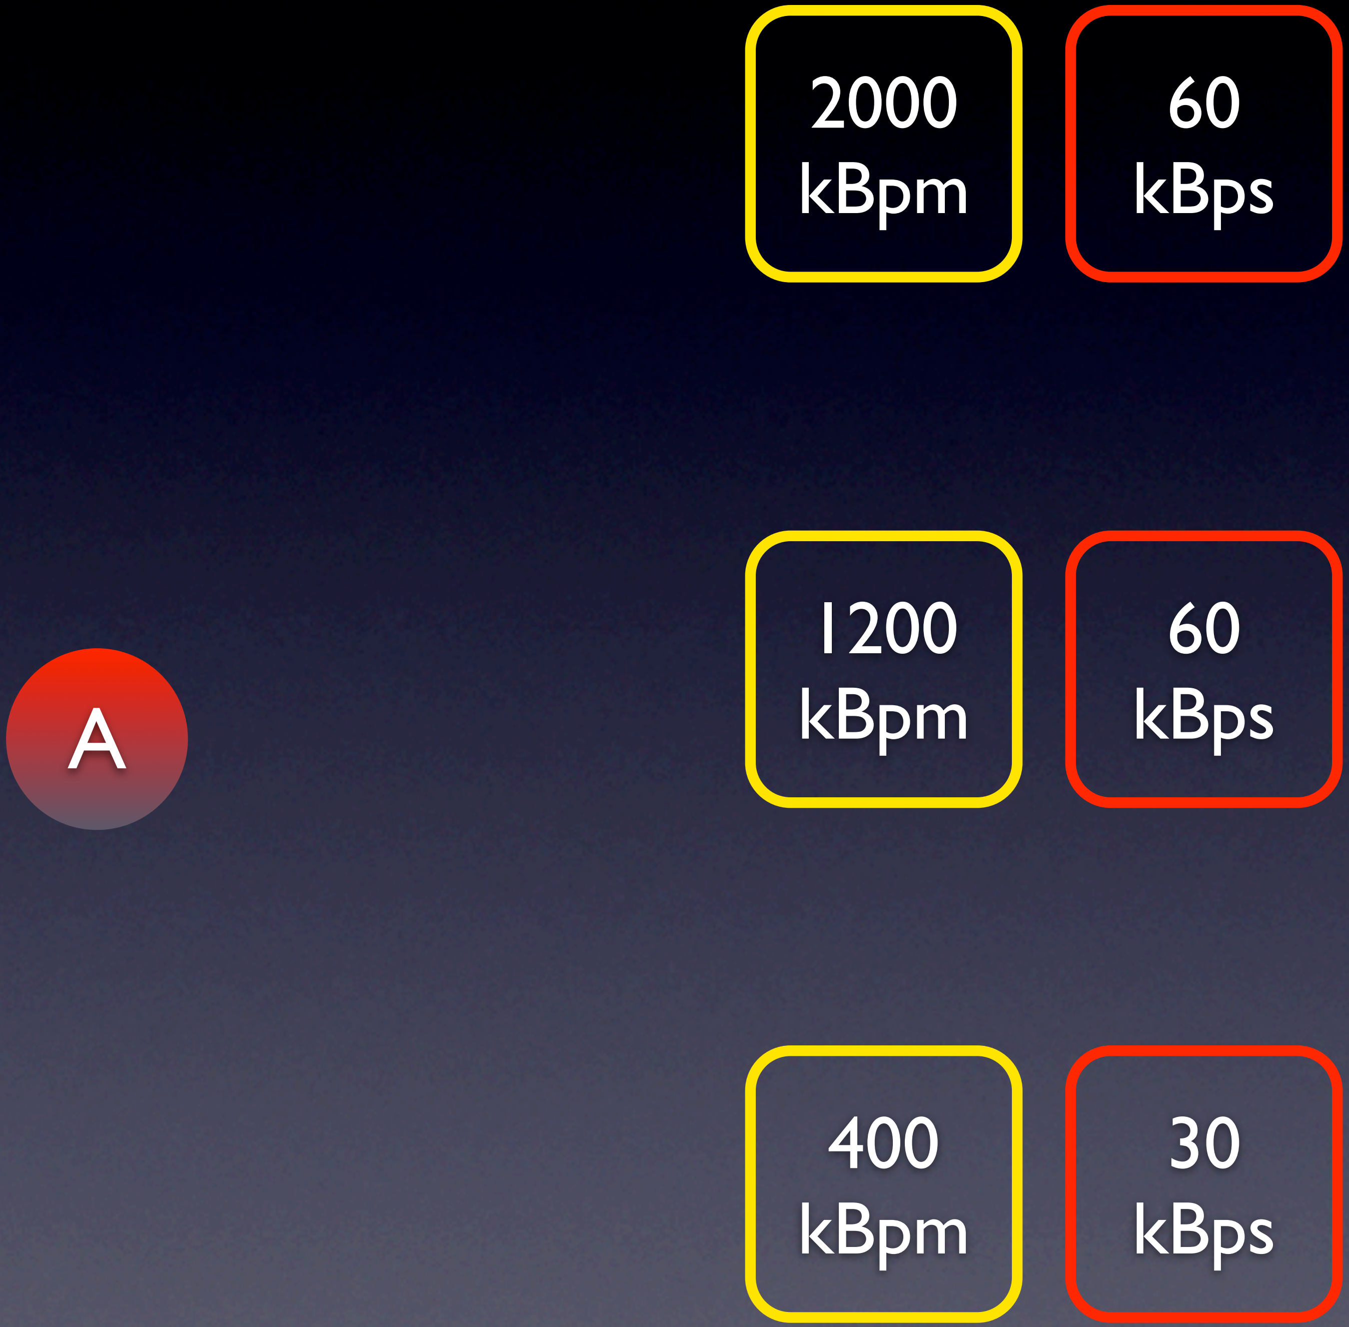

2000 kBpm

60 kBps

1200 kBpm

60 kBps

400 kBpm

30 kBps

Overselling Resources

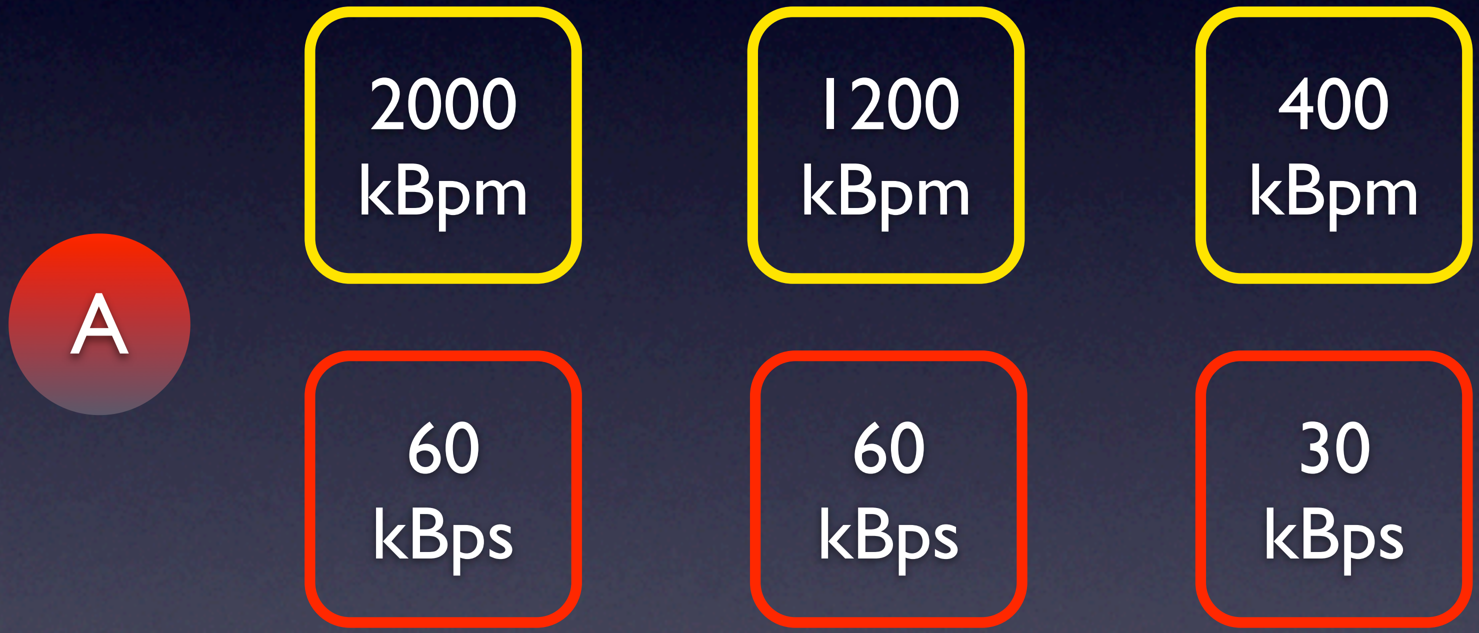

A

2000 kBpm

1200 kBpm

400 kBpm

60 kBps

60 kBps

30 kBps

Overselling Resources

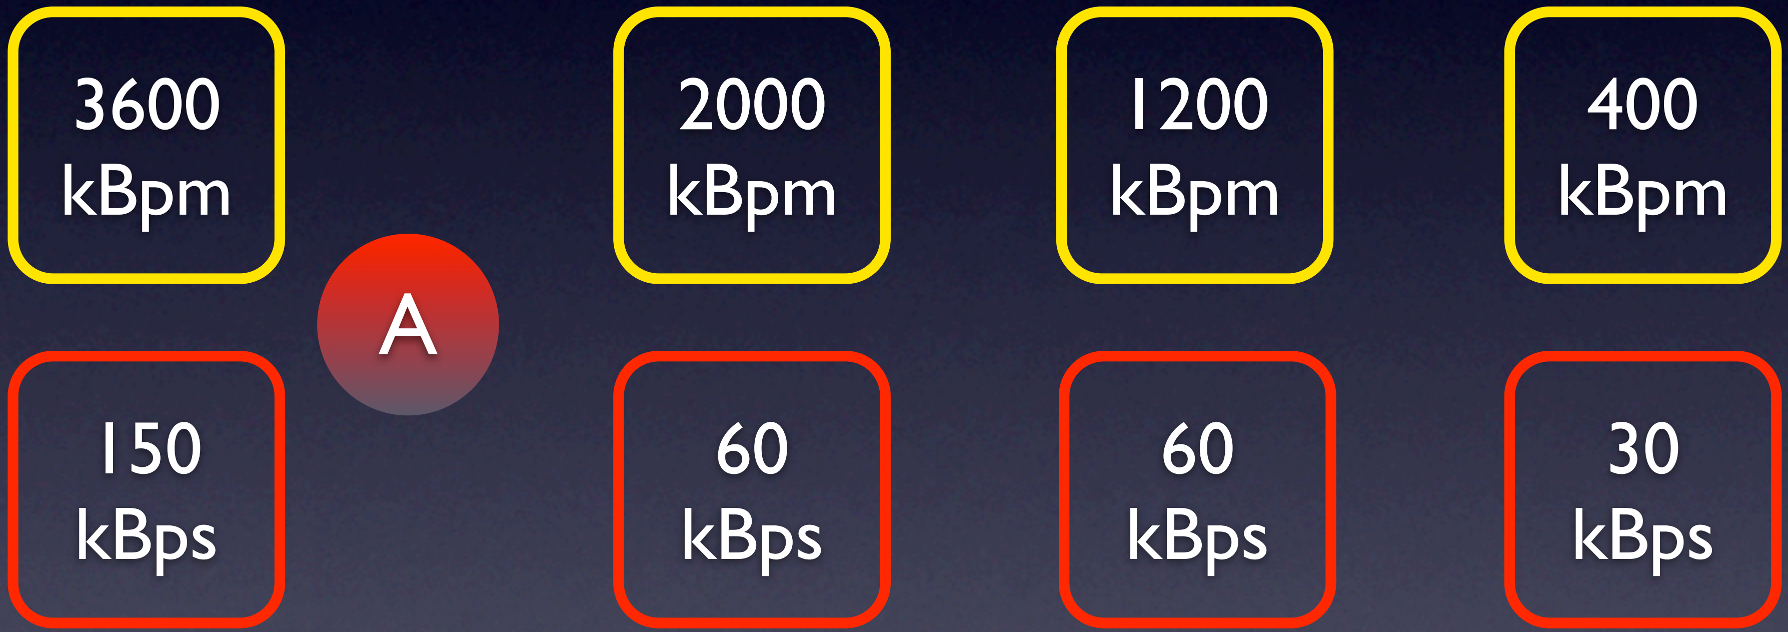

3600 kBpm

2000 kBpm

1200 kBpm

400 kBpm

A

150 kBps

60 kBps

60 kBps

30 kBps

Overselling Resources

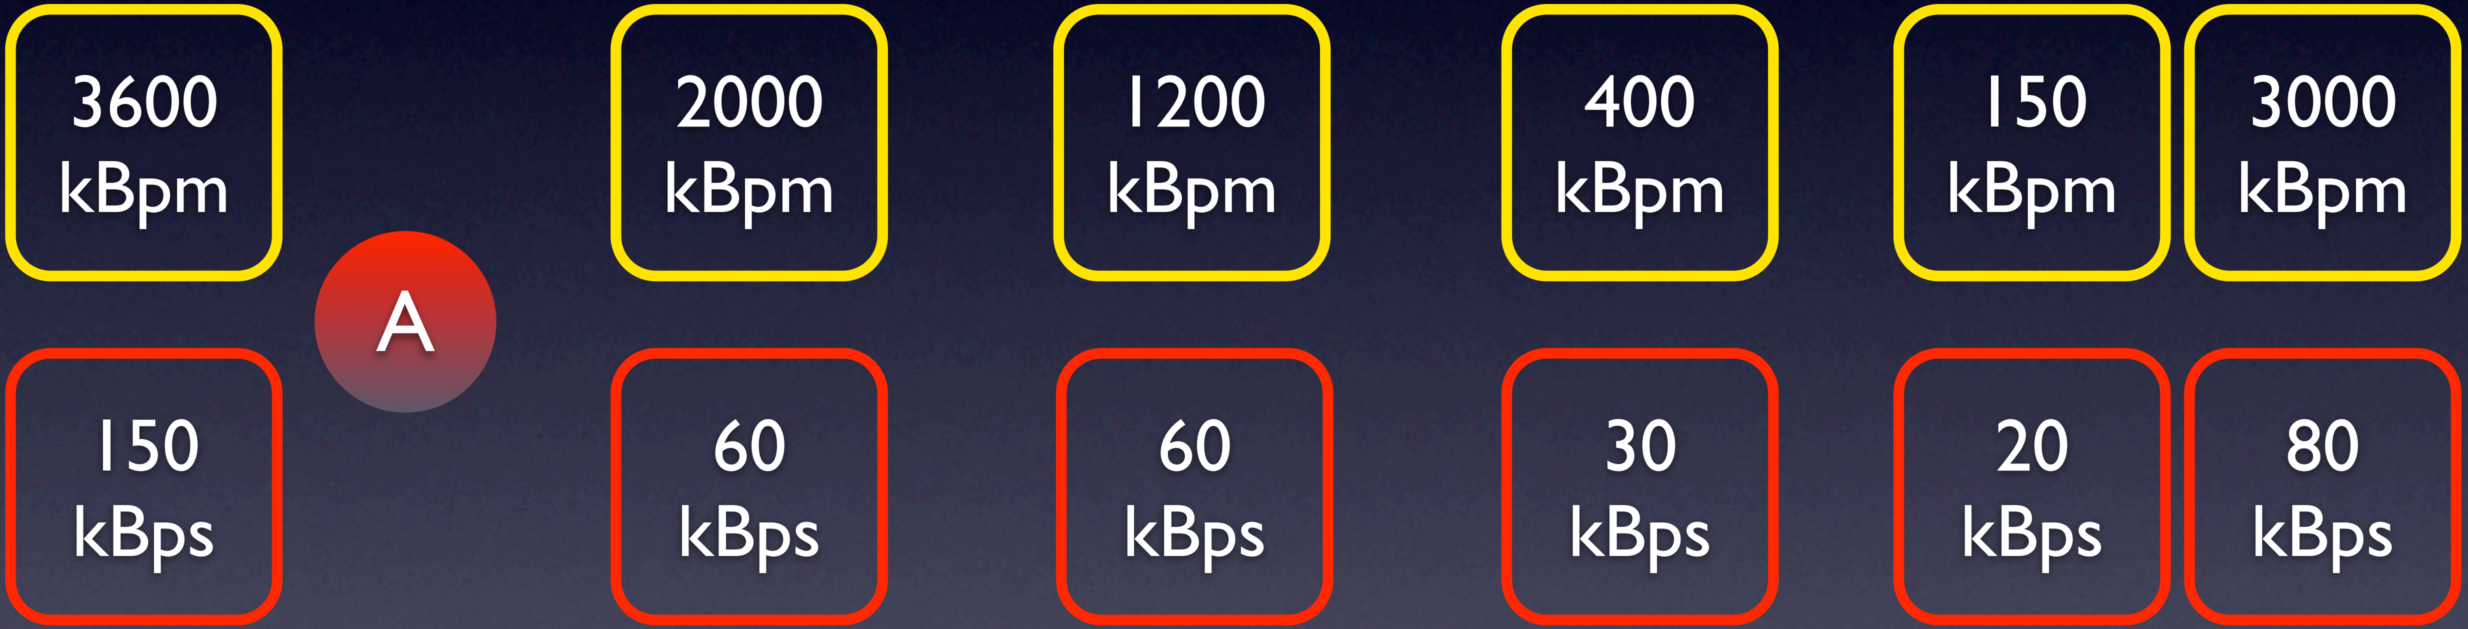

3600 kBpm

2000 kBpm

1200 kBpm

400 kBpm

150 kBpm

3000 kBpm

A

150 kBps

60 kBps

60 kBps

30 kBps

20 kBps

80 kBps

Re-selling others' resources

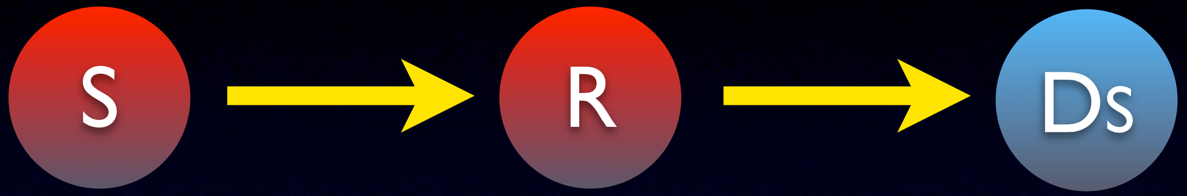

-30 msec

Re-selling others' resources

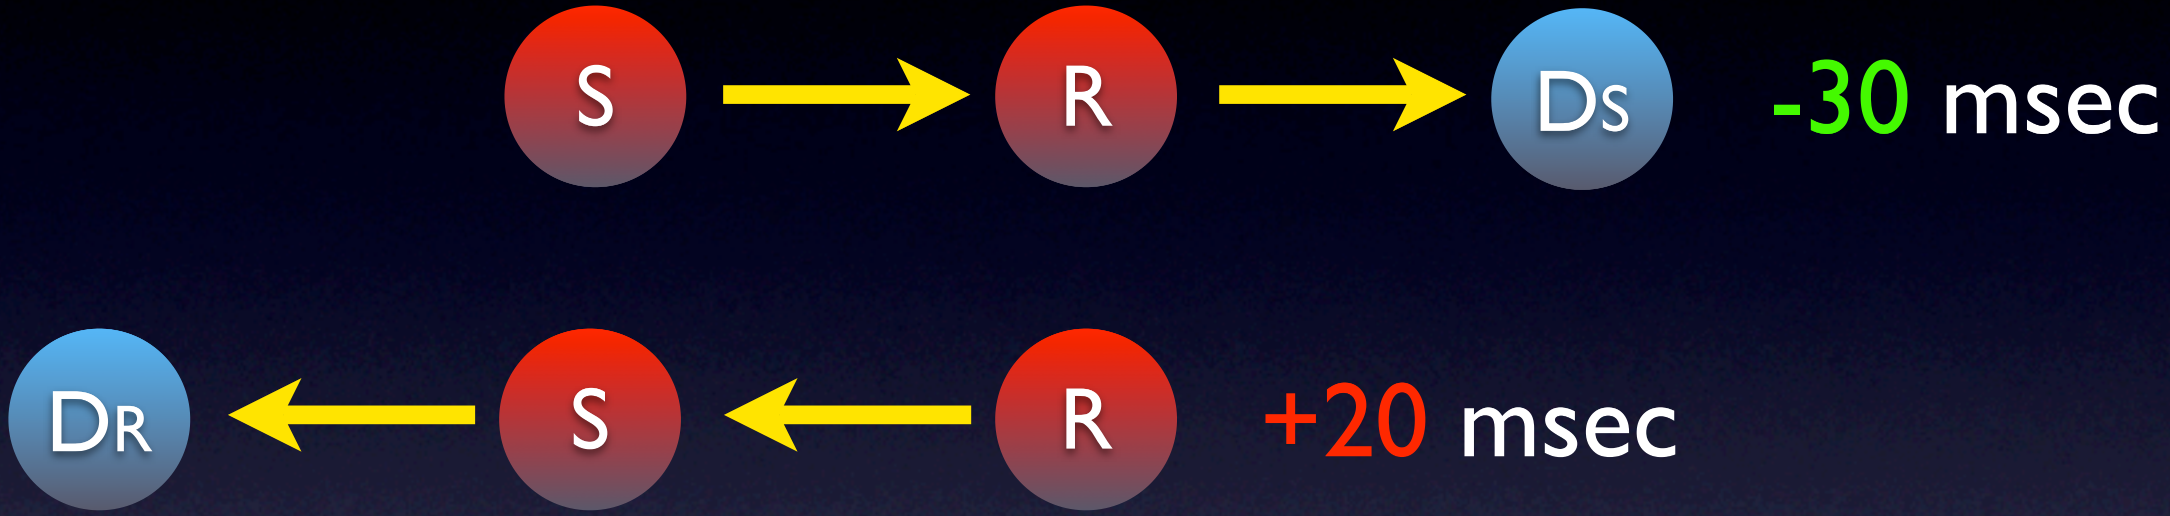

Re-selling others' resources

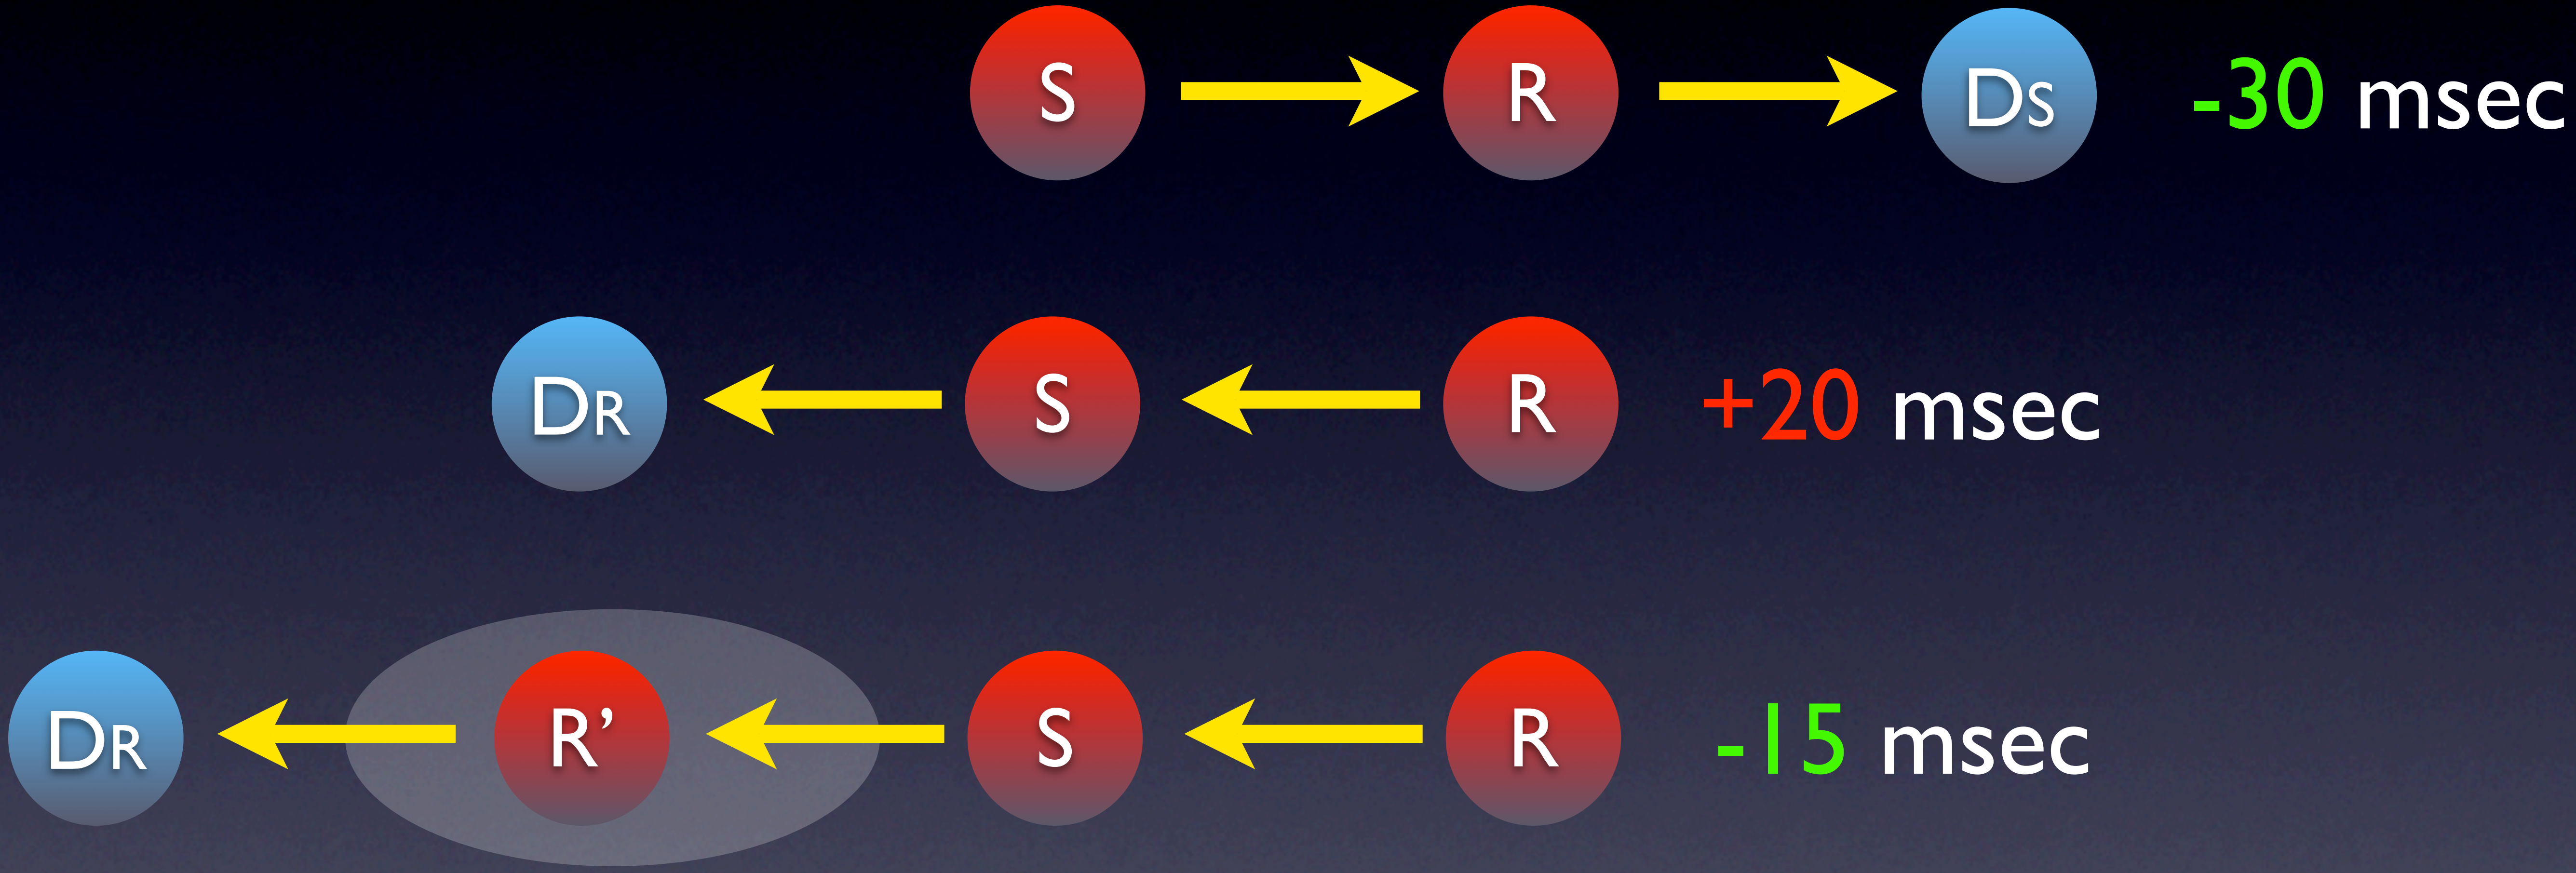

Re-selling others' resources

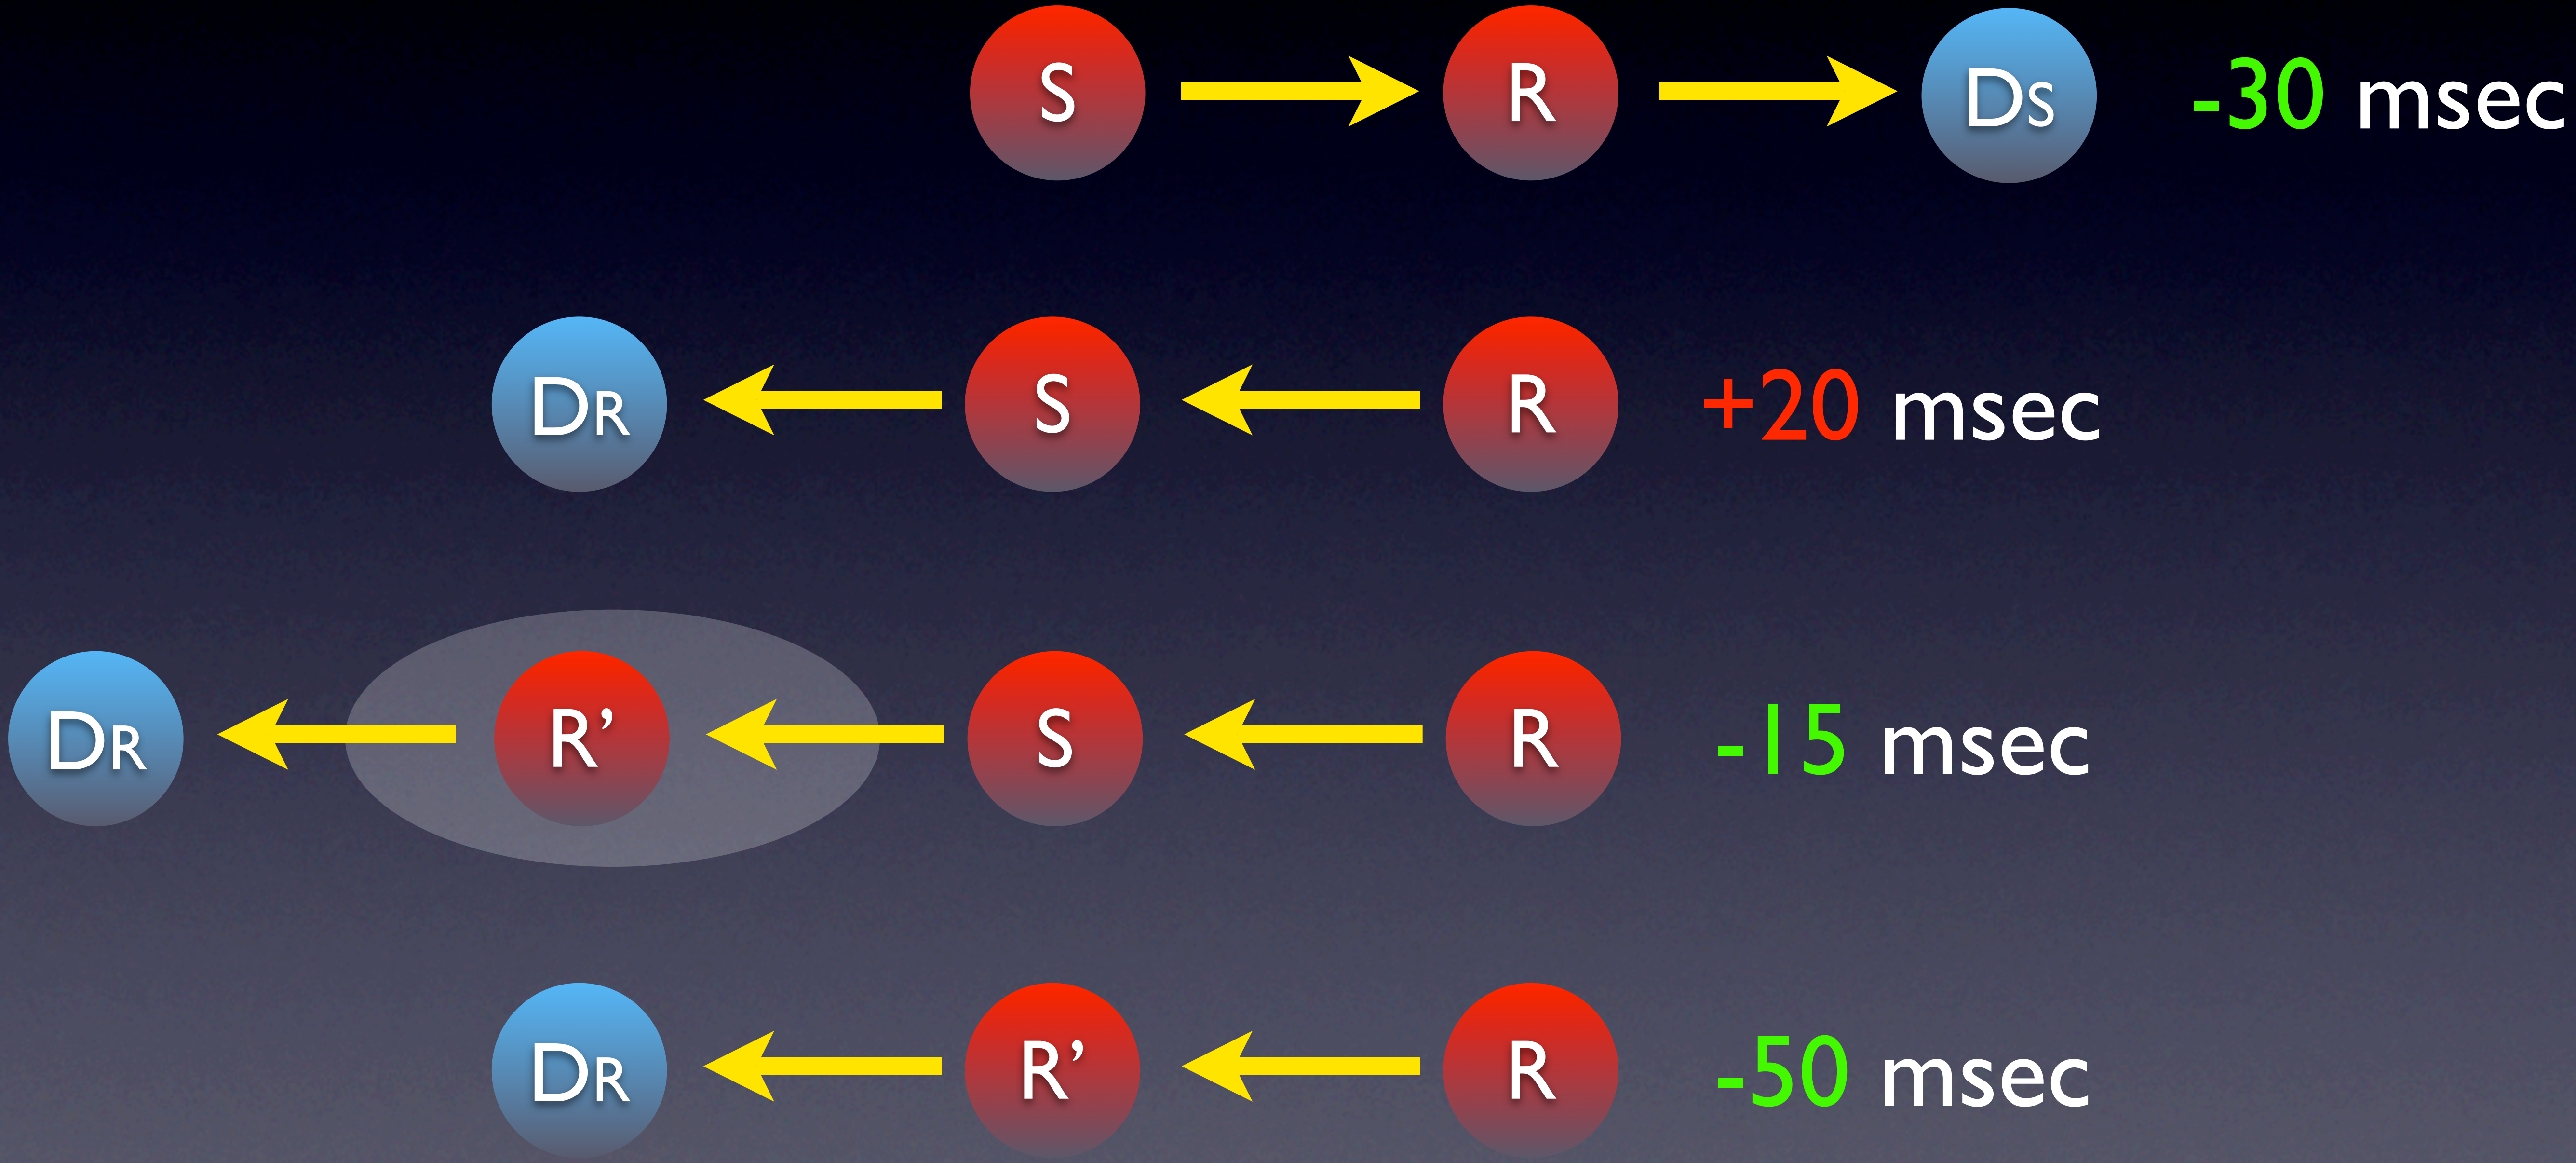

Summary

- The selfish routing overlay problem
  - Distinct from other systems' incentives
  - Goal: Long-lived peerings
- SLAs for long-lived peerings
  - Avoid, don't punish
  - Confidence to react to violations
- Emergent behaviors

http://www.cs.umd.edu/~nspring/peerwise.html