Stanford Information Networks Group

The κ Factor: Inferring Protocol Performance Using Inter-link Reception Correlation

Kannan Srinivasan, Mayank Jain, Jung Il Choi, Tahir Azim, Edward S Kim, Philip Levis and Bhaskar Krishnamachari

MobiCom 2010

23rd Sep 2010

Spatial Independence Assumption

Losses on different links are independent

- after a link failure, routing protocols choose the next shortest path forwarder
- simulators explicitly generate channel states independently

Spatial Independence Assumption

Losses on different links are independent

- after a link failure, routing protocols choose the next shortest path forwarder
- simulators explicitly generate channel states independently

When is this assumption safe?

Why does it matter?

Inter-link (Spatial) Correlation

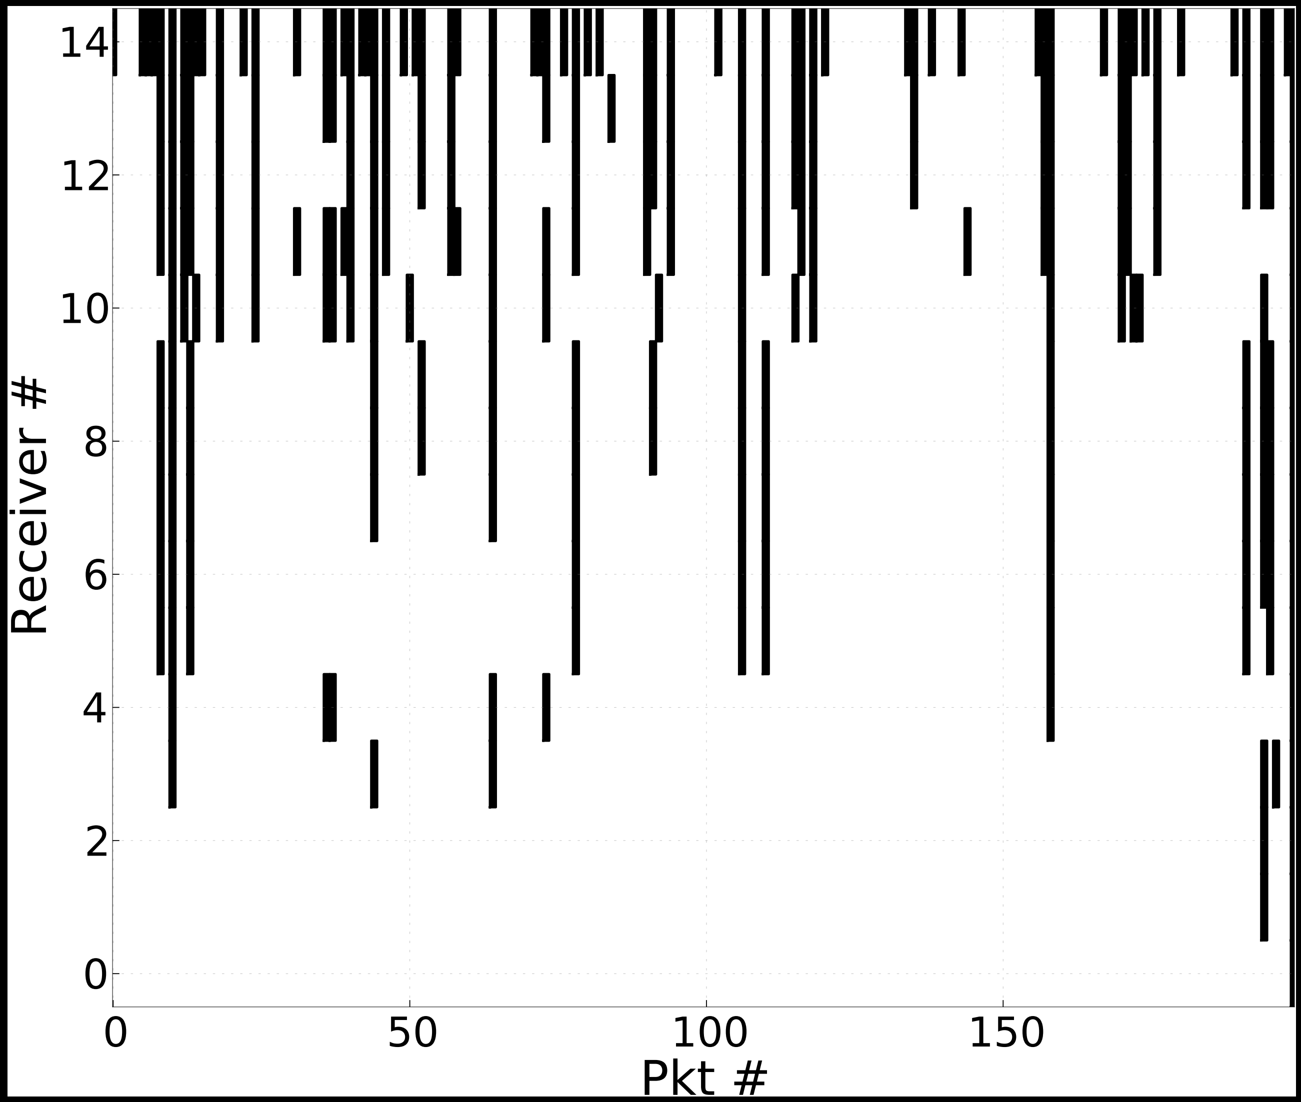

(a) Real Trace, Mirage

Inter-link (Spatial) Correlation

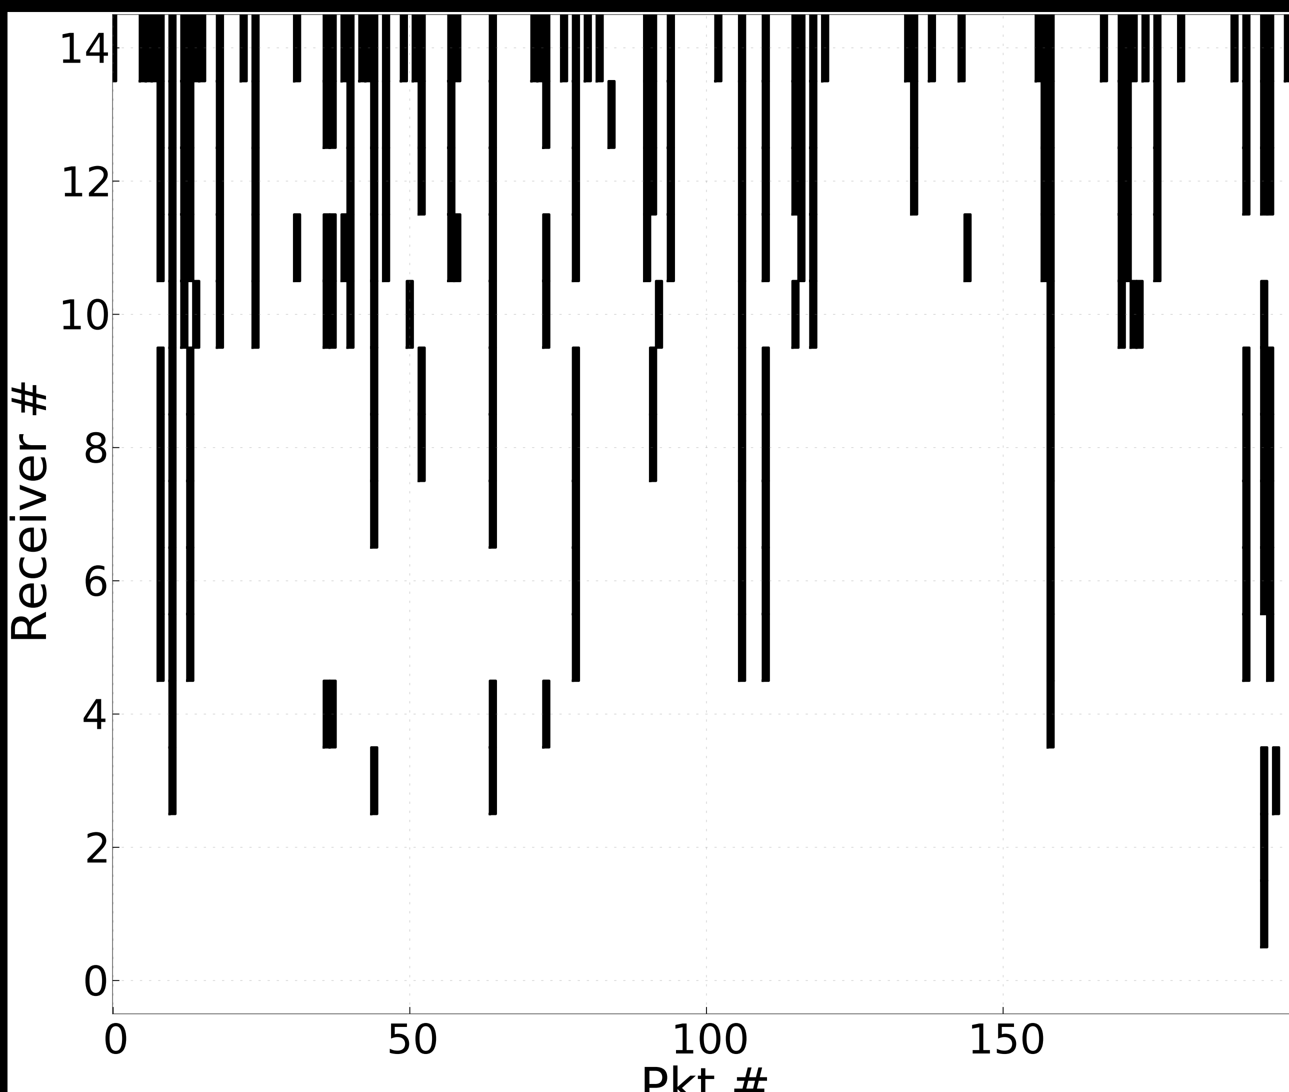

(a) Real Trace, Mirage

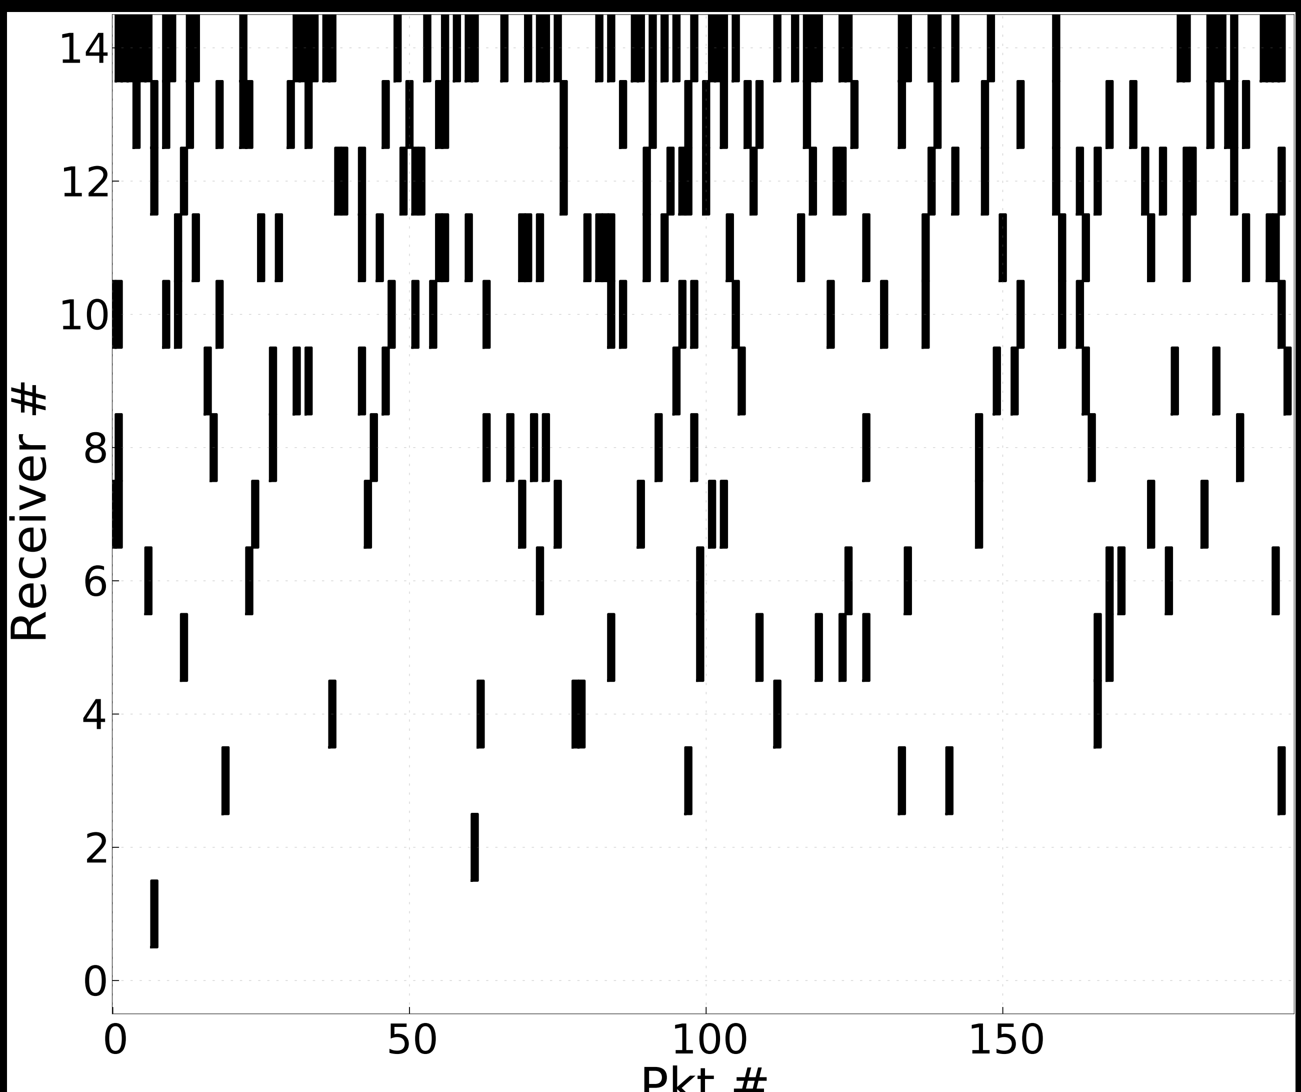

(b) Synthetic (Independent) Trace
(Every link has the same packet reception ratio (PRR) as in real trace)

Inter-link (Spatial) Correlation

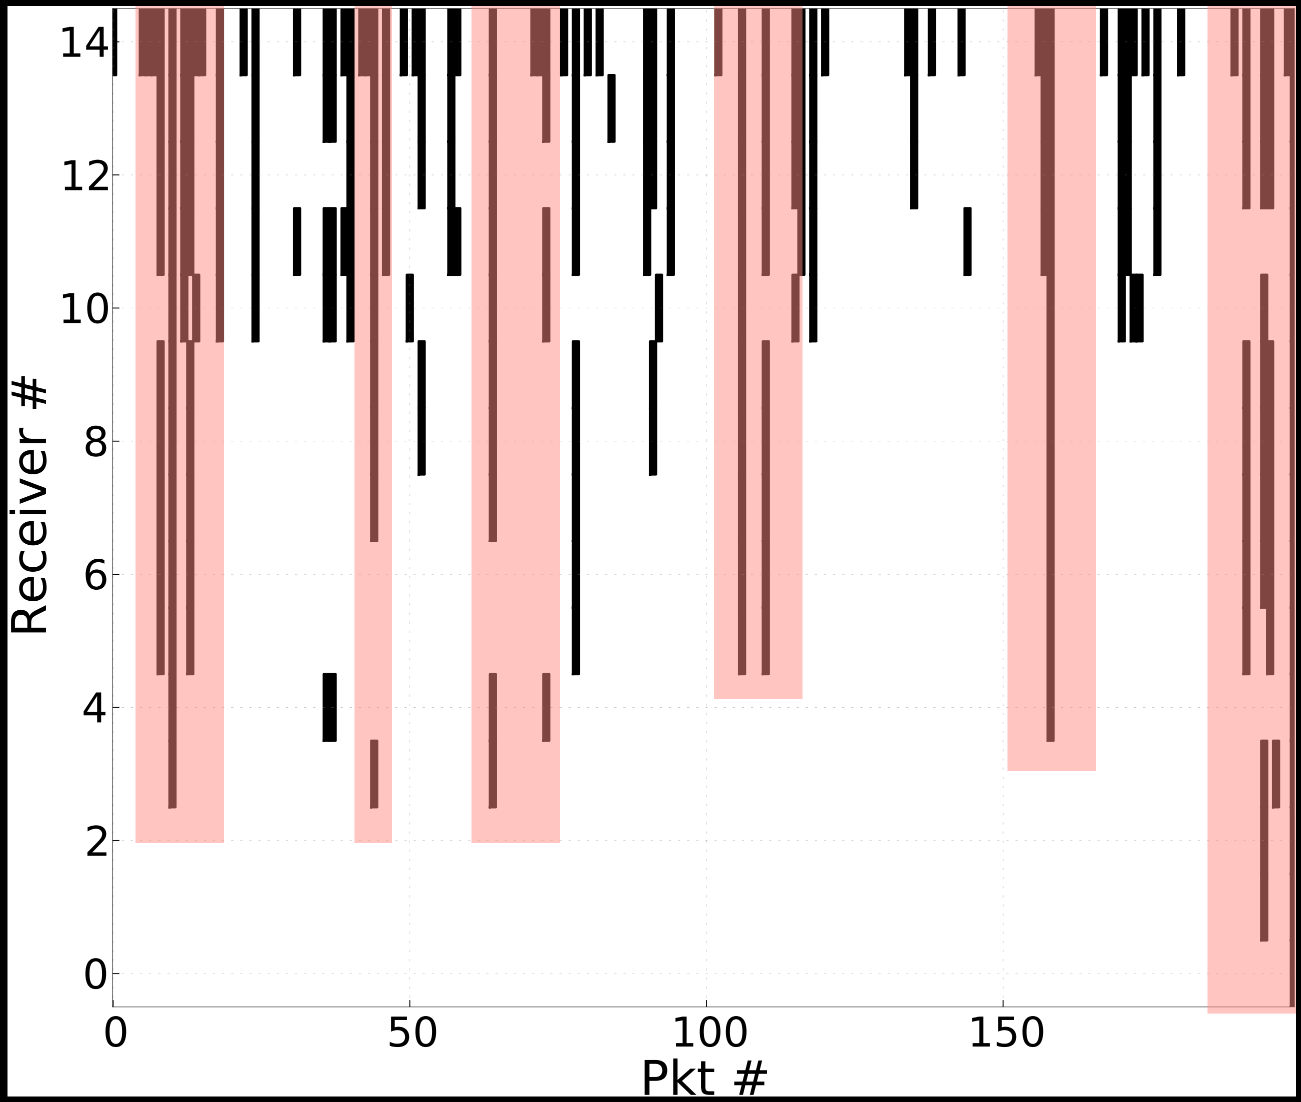

(a) Real Trace, Mirage

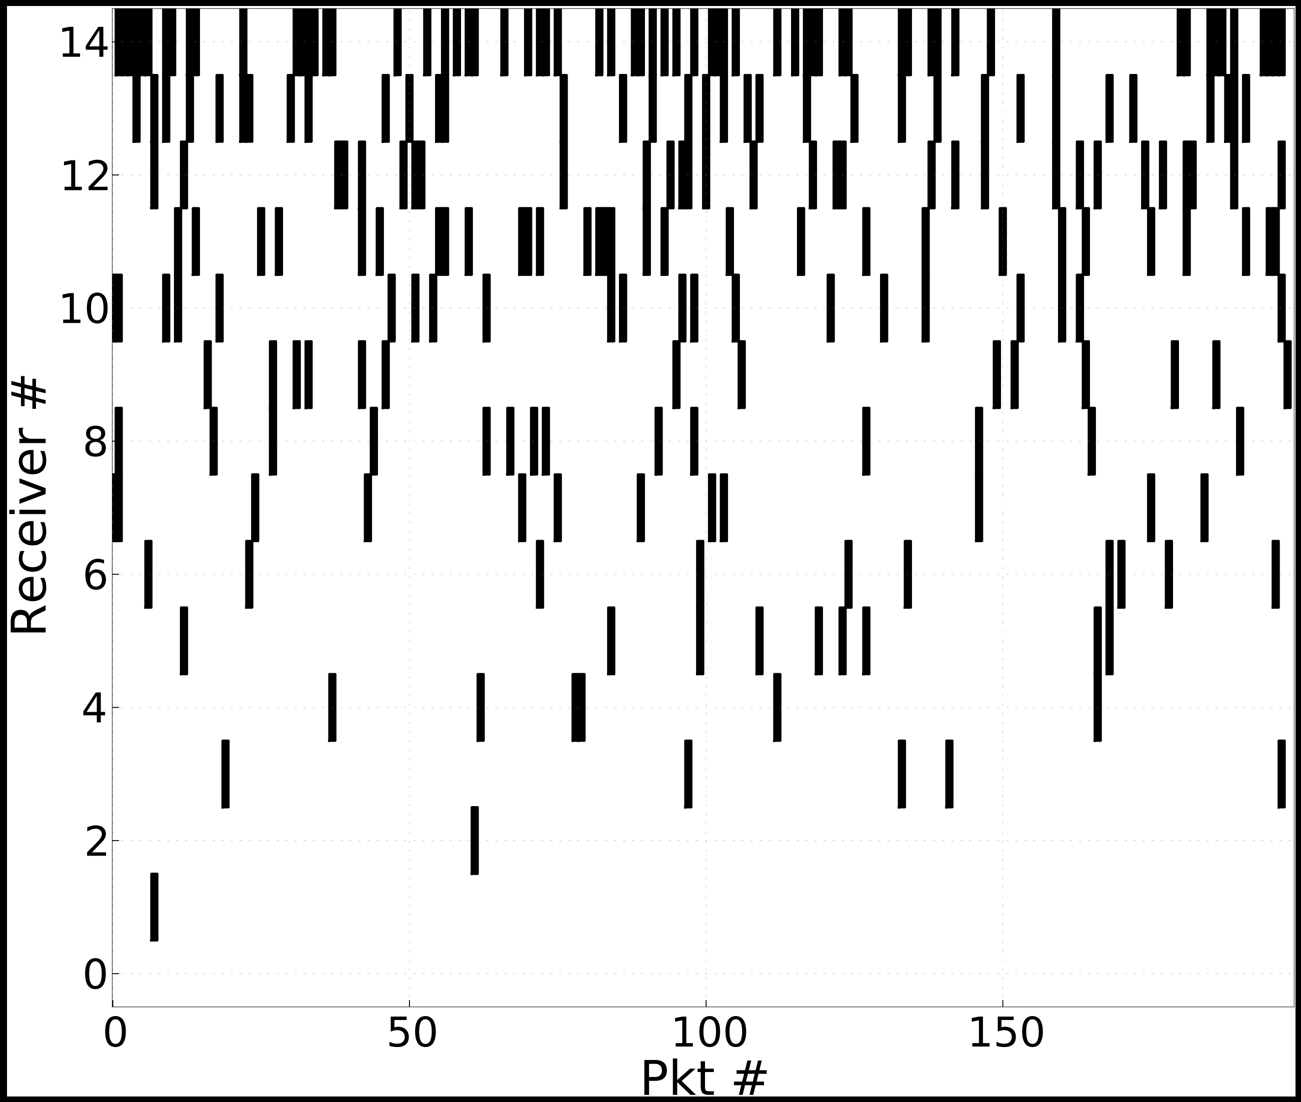

(b) Synthetic (Independent) Trace

Losses are well aligned (correlated)

So what?

Traditional Routing

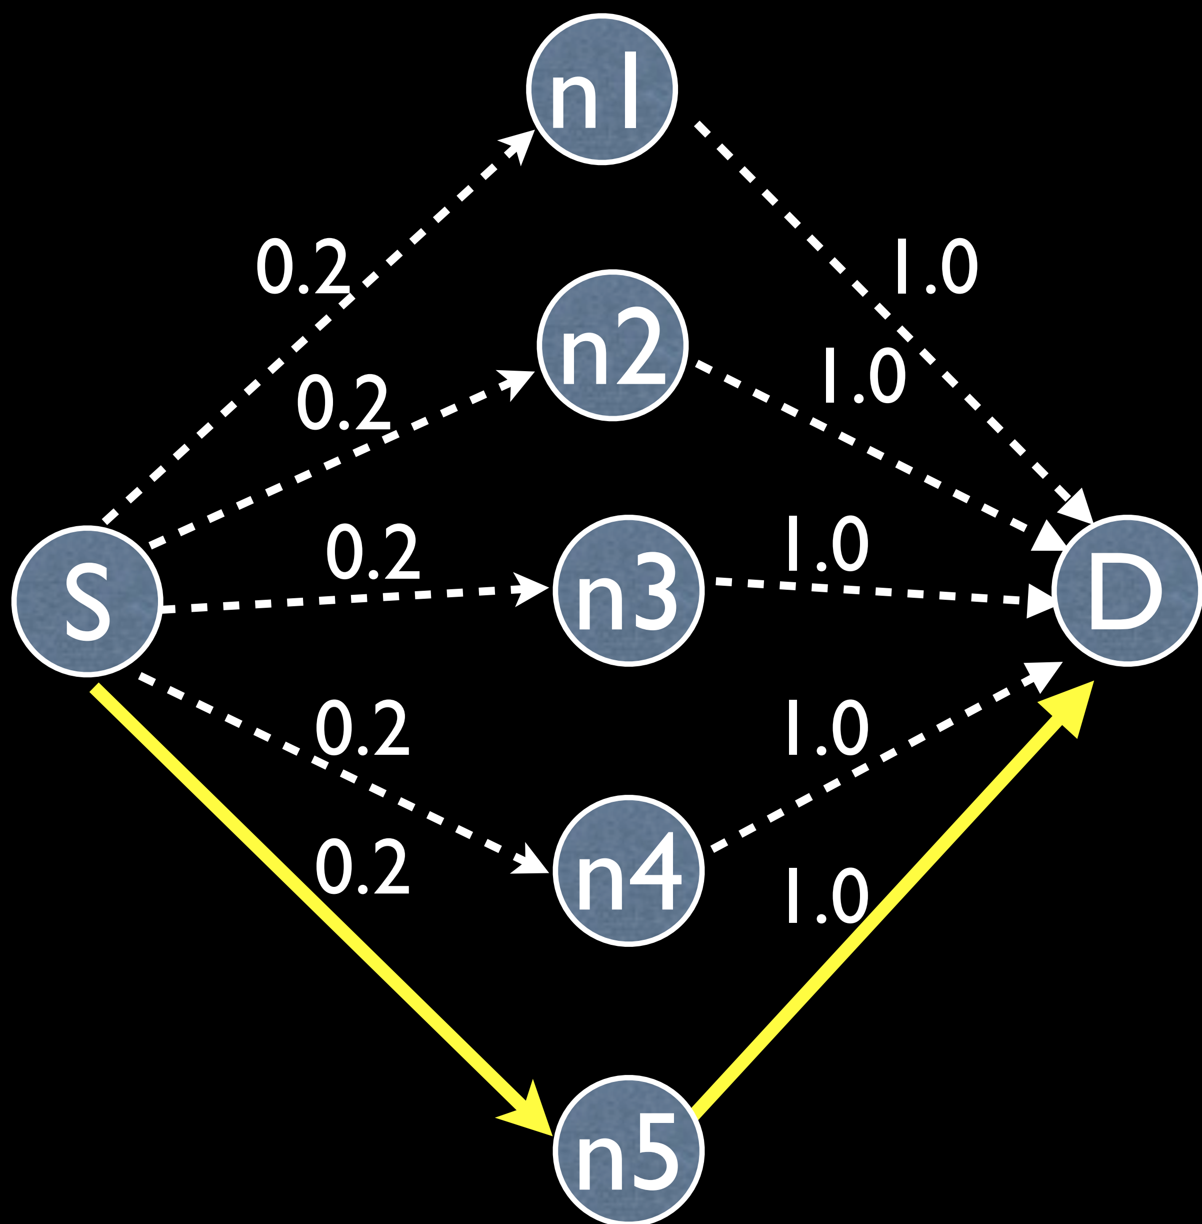

- Edge weights are PRRs
- S selects n5 as next-hop
- ETX: expected number of transmissions
  - $\text{ETX} = \frac{1}{0.2} + \frac{1}{1.0} = 6.0$

So what?

Opportunistic Routing

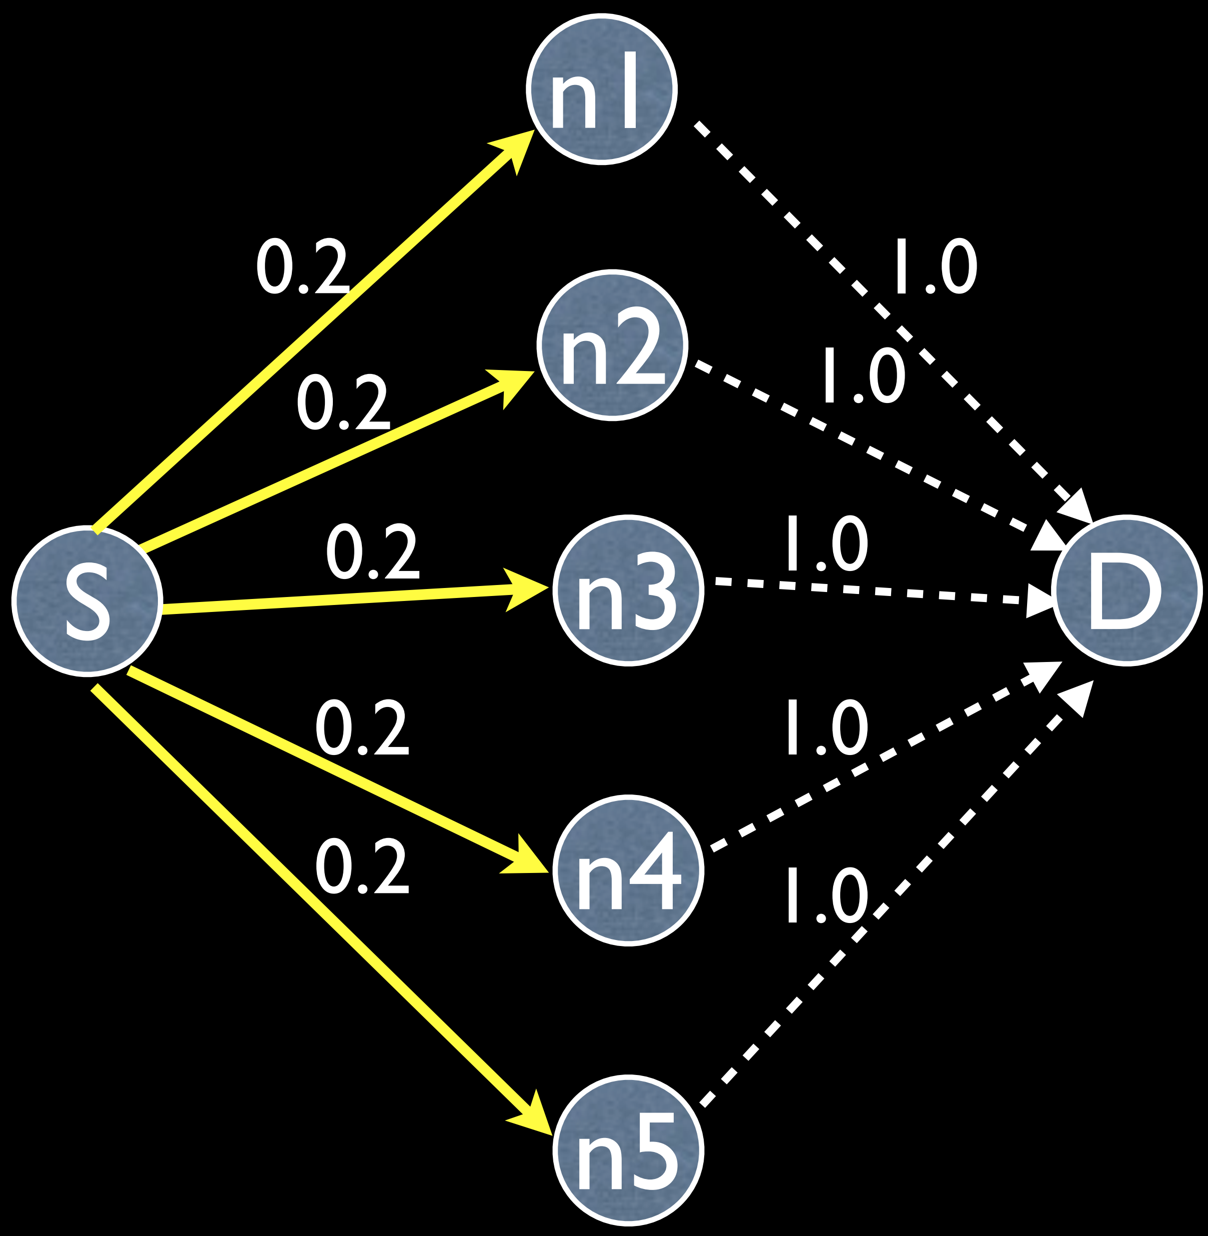

- S lists n1-n5 as next-hops
  - S stops as soon as at least one of the next hop nodes receives

So what?

Opportunistic Routing

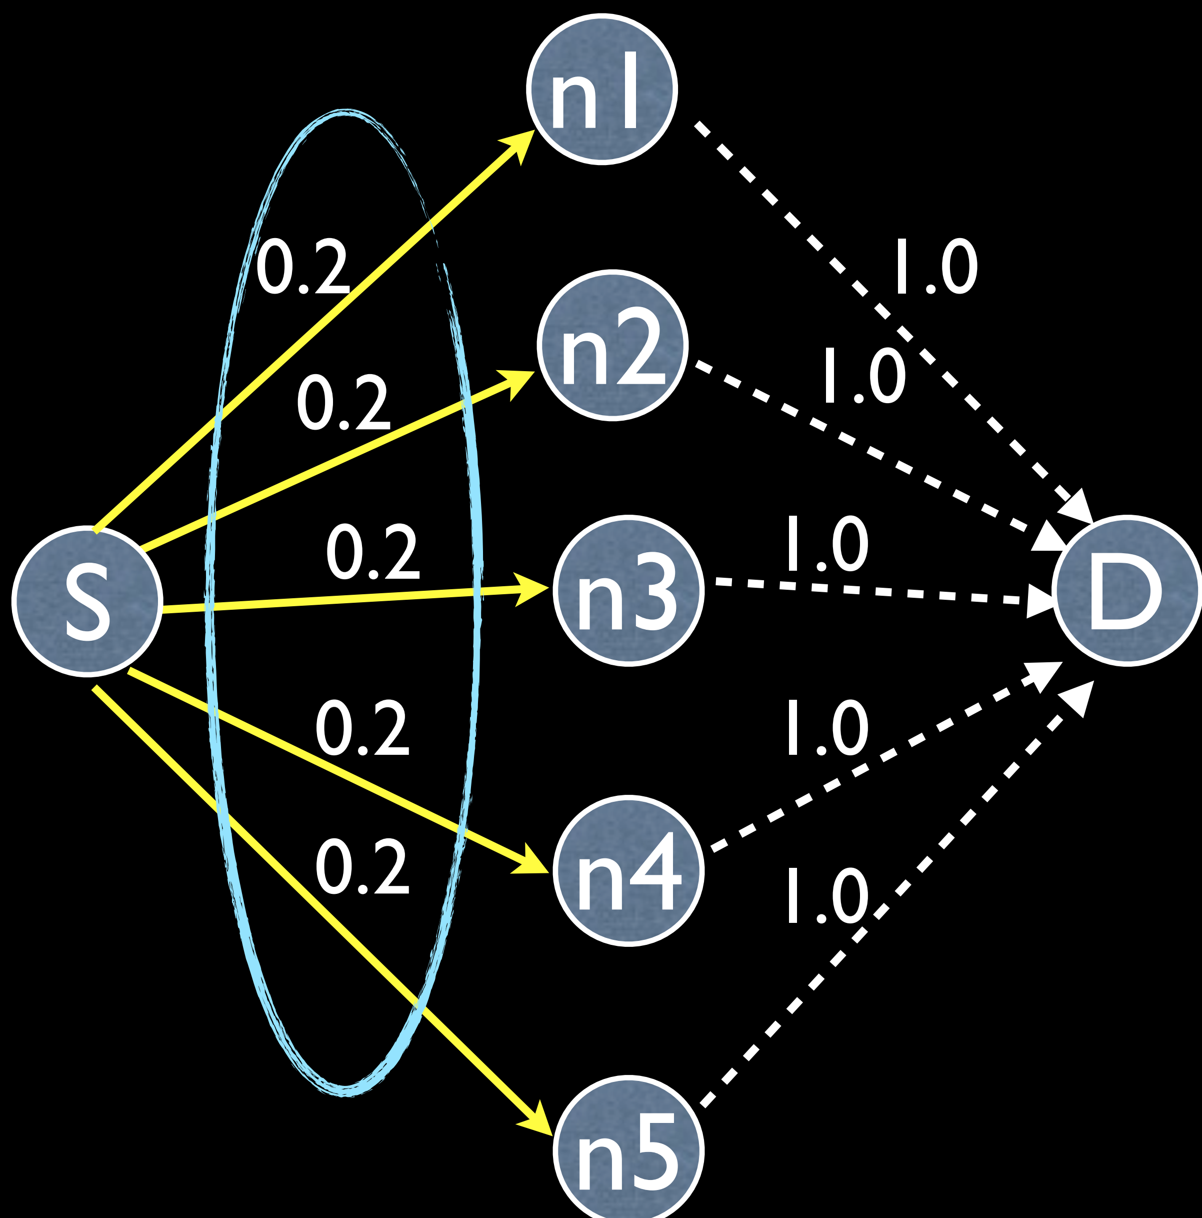

Independent:

$$\mathbf{ETX} = \frac{1}{1-(1-0.2)^5} + 1 = 2.49$$

So what?

Opportunistic Routing

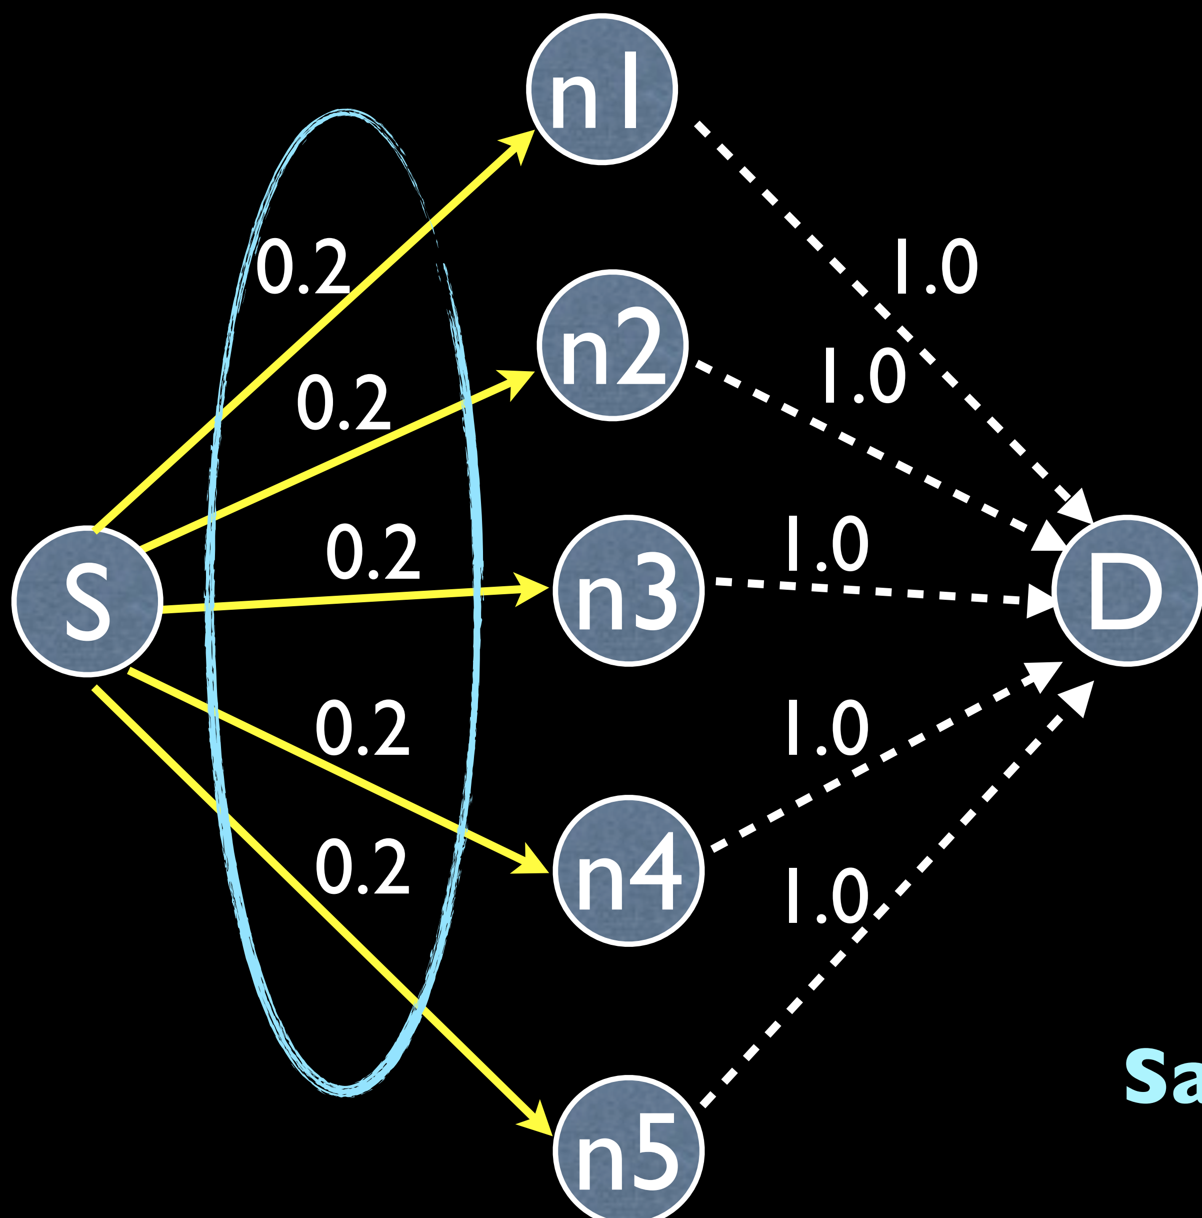

Independent:

$$\mathbf{ETX} = \frac{1}{1-(1-0.2)^5} + 1 = 2.49$$

Perfectly Correlated:

$$\mathbf{ETX} = \frac{1}{1-(1-0.2)} + 1 = 6.0$$

Same cost as traditional routing (without coordination cost)!

Correlation has implications to protocol performance

So far

- Spatial correlation assumption does not always hold true
  - a measured network: 70% of link pairs are highly correlated
- The degree of correlation has implications to protocol performance

Problem Statement

Need a good way to measure spatial correlation to understand its implications to protocol performance

- existing metrics conflate correlation with link pair PRRs

Research Contributions

- Present a new metric: κ
- Show how well network coding protocols perform, based on κ
- Show κ 's ability to predict opportunistic routing protocol performance (in paper)
  - perfect prediction when a node has 2 potential forwarders
  - more than 2 forwarders: perfect prediction for most of the nodes

Outline

- Desired Metric Properties
- The κ Metric
- κ 's Usefulness
- Open Questions

Outline

- Desired Metric Properties
- The κ Metric
- κ 's Usefulness
- Open Questions

Desired Metric Properties

a) A scalar with a finite range: [-1,1]

- >0: positive correlation
- <0: negative correlation

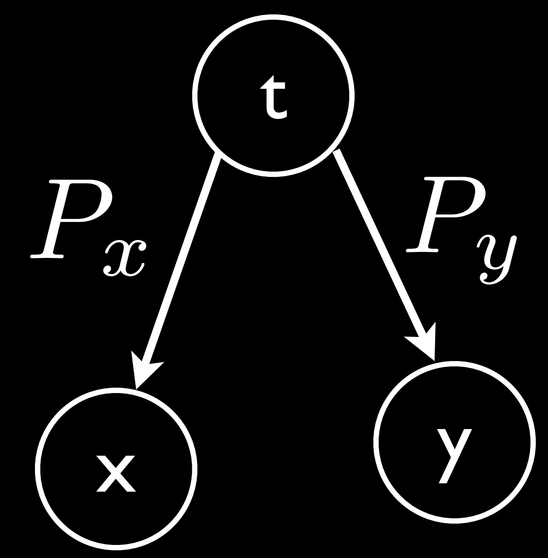

Desired Metric Properties

a) A scalar with a finite range: [-1,1]

b) Symmetric

- Metric(x,y) = Metric(y,x)

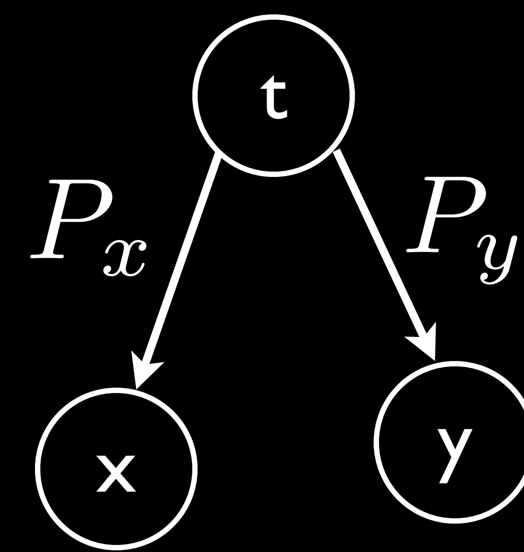

Desired Metric Properties

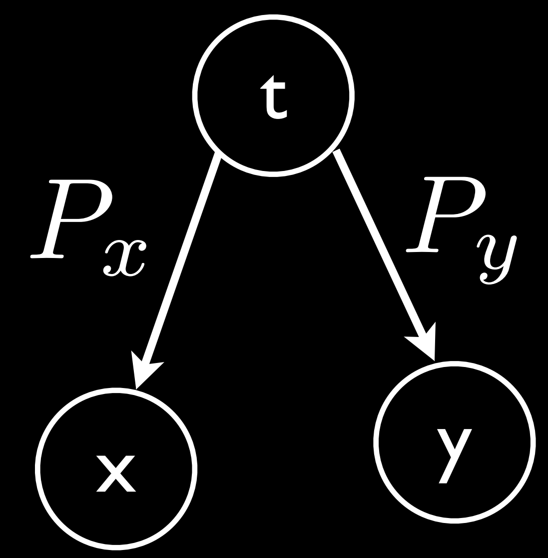

a) A scalar with a finite range: [-1,1]

b) Symmetric

c) Irrespective of PRRs:
   1 for perfectly positively correlated link pair
   -1 for perfectly negatively correlated link pair

Desired Metric Properties

a) A scalar with a finite range: [-1,1]

b) Symmetric

c) Irrespective of PRRs:

1 for perfectly positively correlated link pair

-1 for perfectly negatively correlated link pair

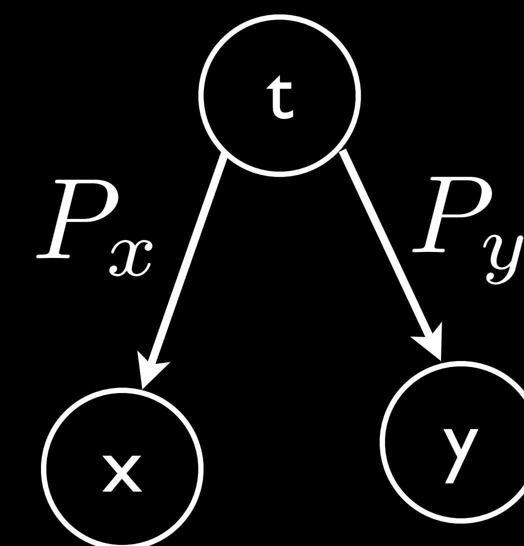

Not a made-up property!

Perfect Positive Correlation (Metric = 1)

Same PRR Link Pair

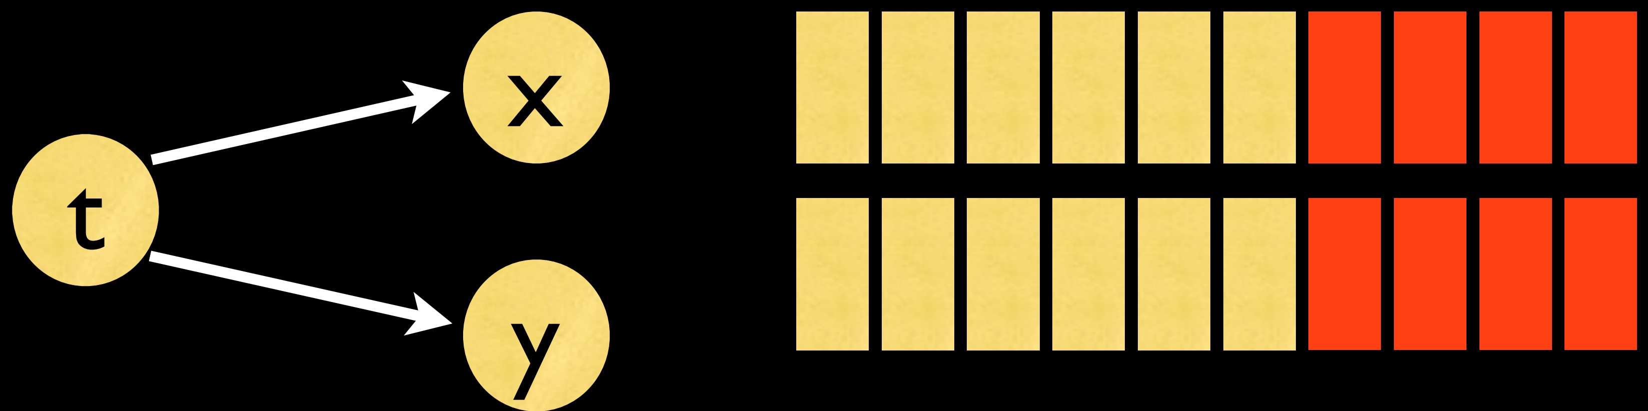

Opportunistic routing should choose x or y, but not both

Perfect Positive Correlation (Metric = 1)

Same PRR Link Pair

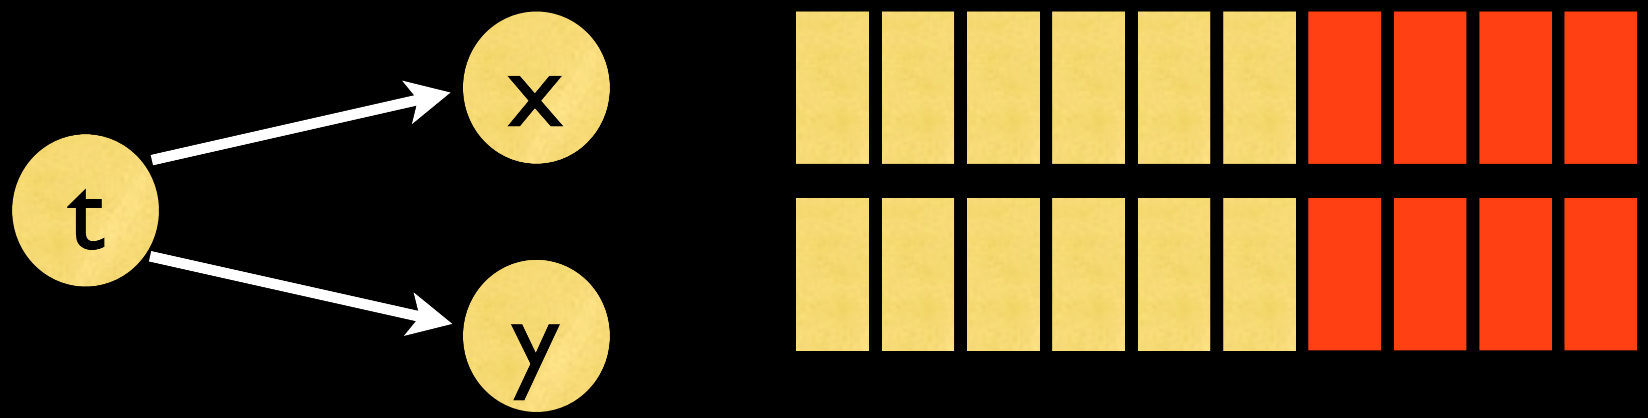

Different PRR Link Pair

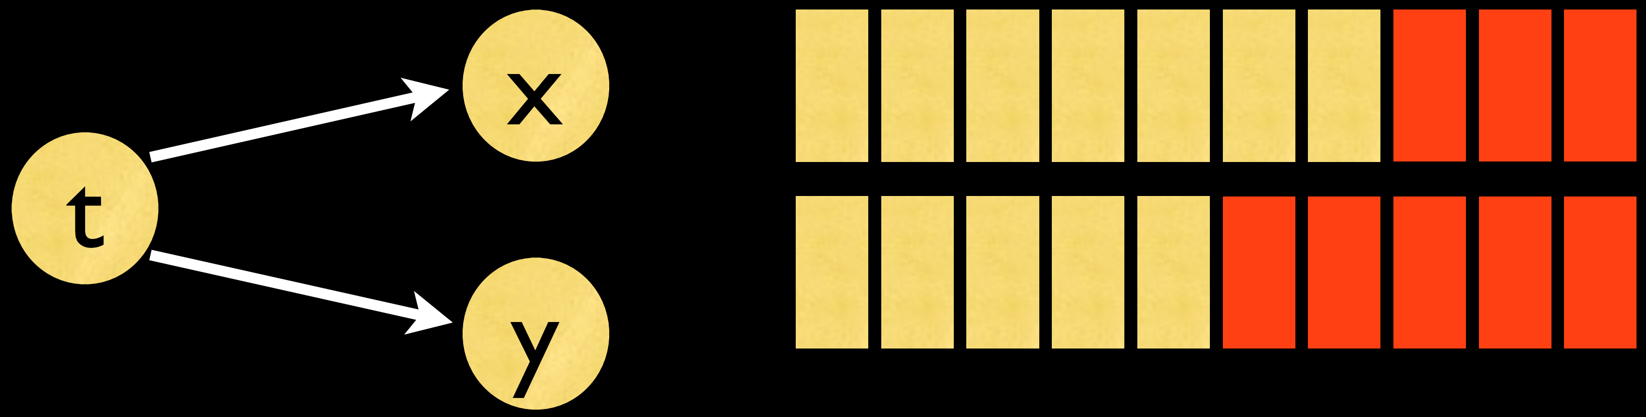

Opportunistic routing should only choose x

Perfect Negative Correlation (Metric = -1)

Sum of Link Pair PRRs = 1

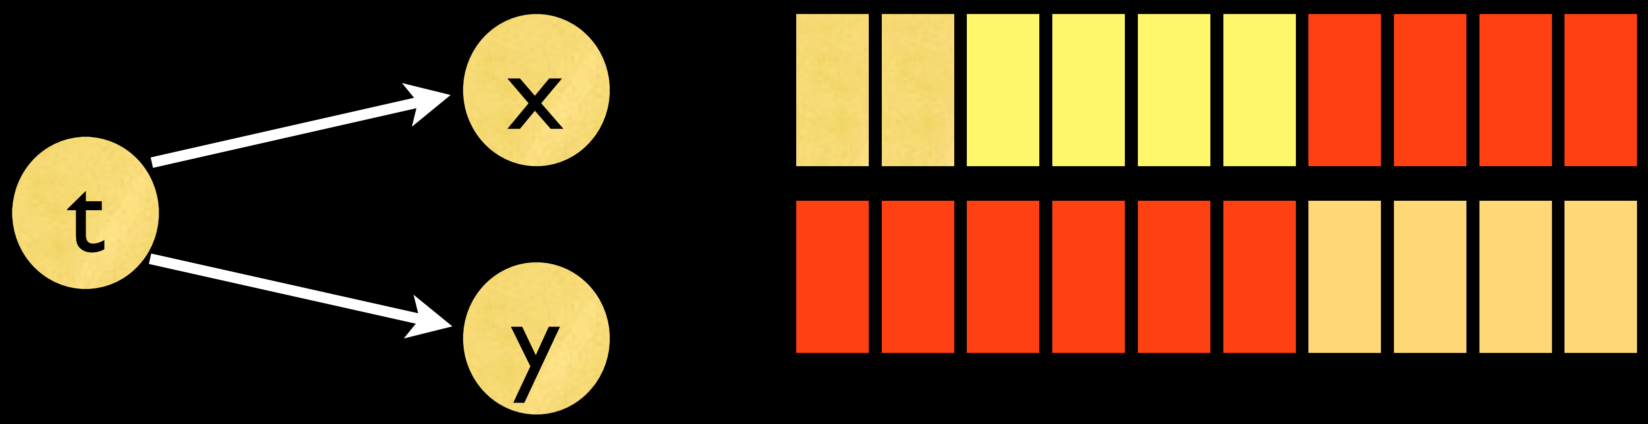

Every packet succeeds on only one link

Perfect Negative Correlation (Metric = -1)

Sum of Link Pair PRRs = 1

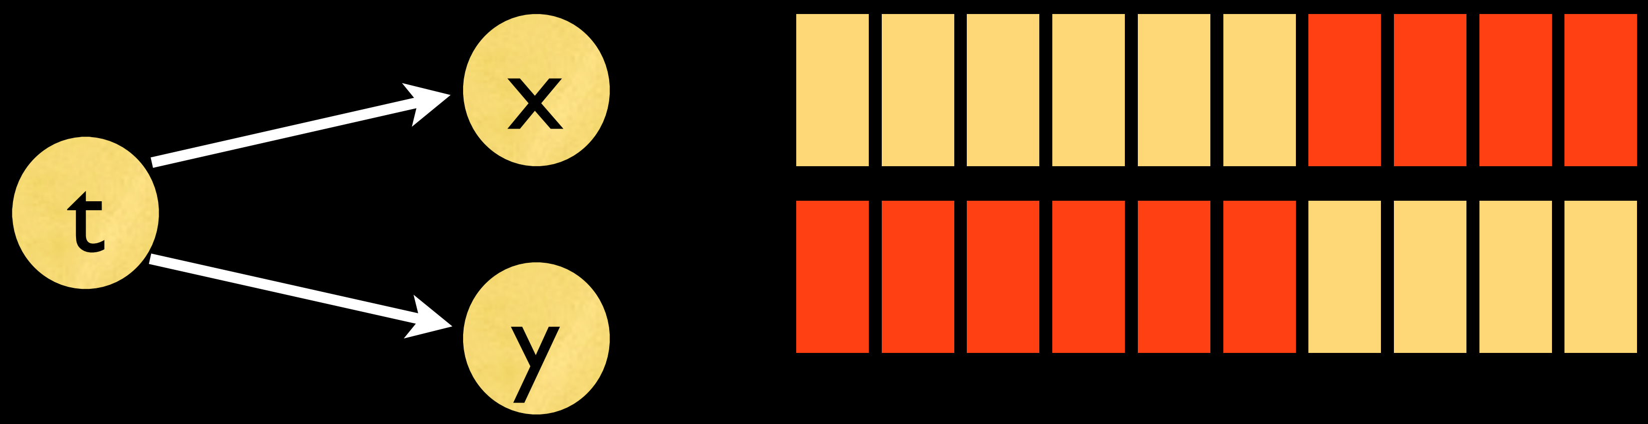

Sum of Link Pair PRRs > 1

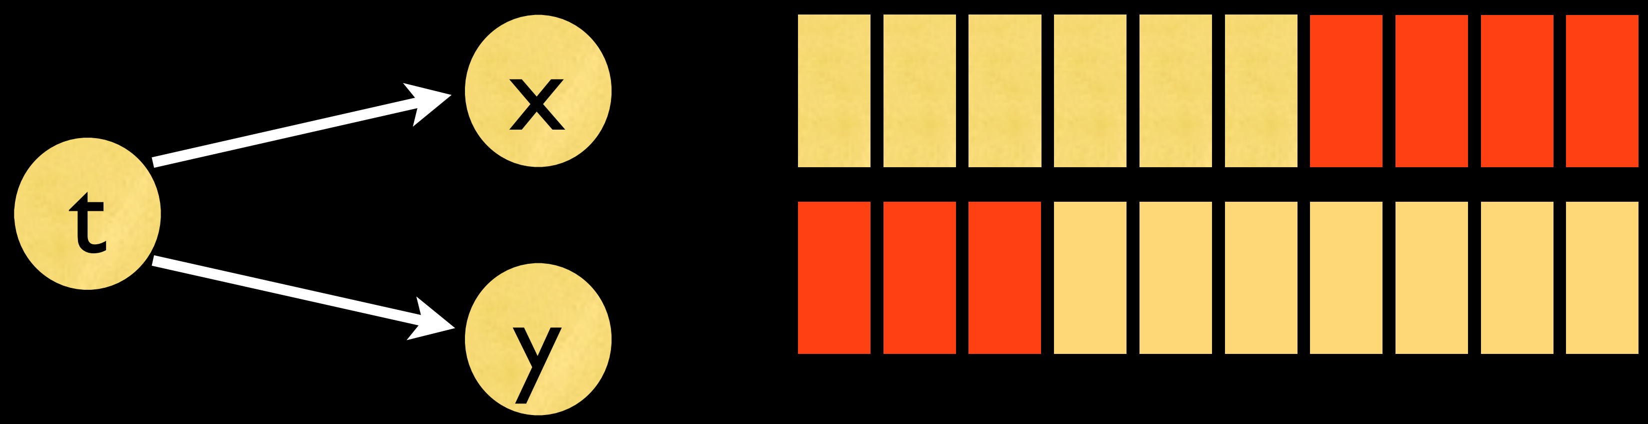

Every packet succeeds at one or both the links

Perfect Negative Correlation (Metric = -1)

Sum of Link Pair PRRs < 1

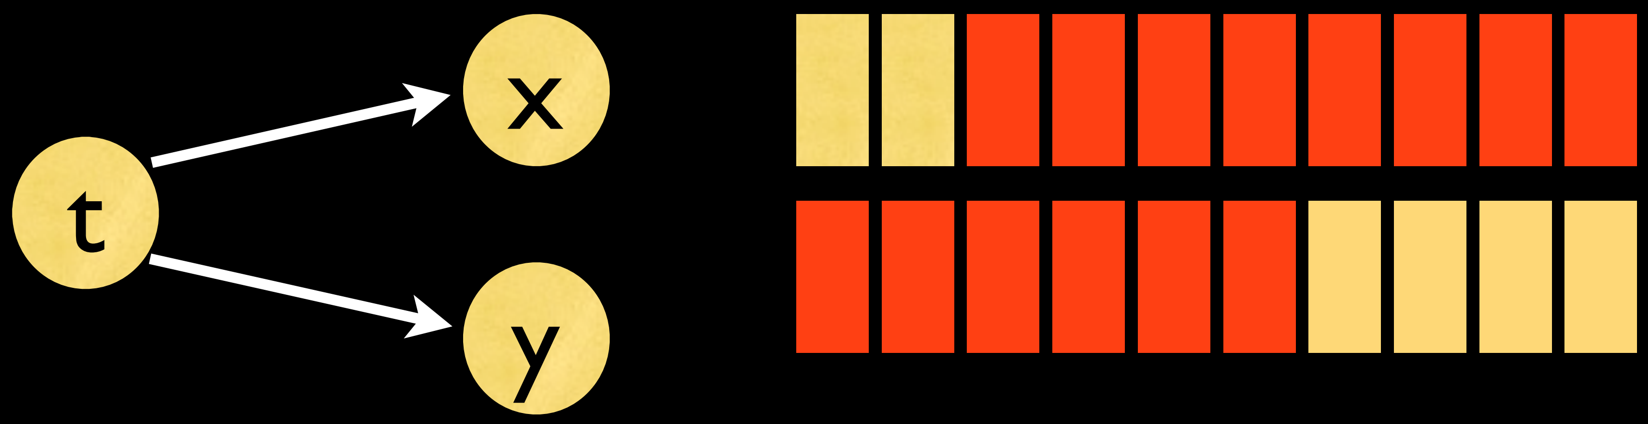

Opportunistic routing benefits most, given the two PRRs

Desired Metric Properties

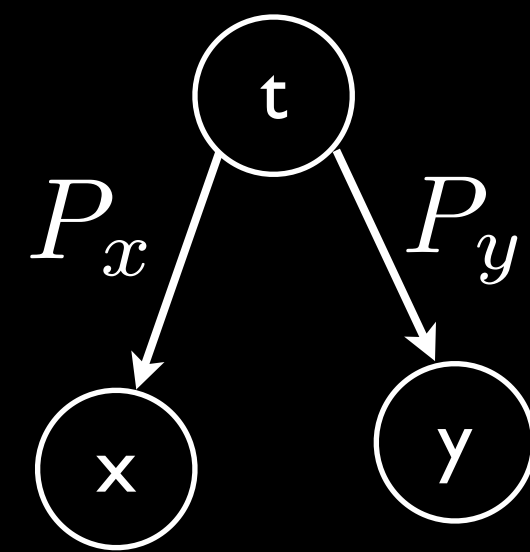

a) A scalar with a finite range: [-1,1]

b) Symmetric

c) Irrespective of PRRs:
1 for perfectly positively correlated link pair
-1 for perfectly negatively correlated link pair

An existing metric: χ

A recent inter-link correlation metric [1,2,3]:

$$\chi = P(x=0|y=0) - P(x=0)$$

$\chi = 0 \Rightarrow$ losses are independent

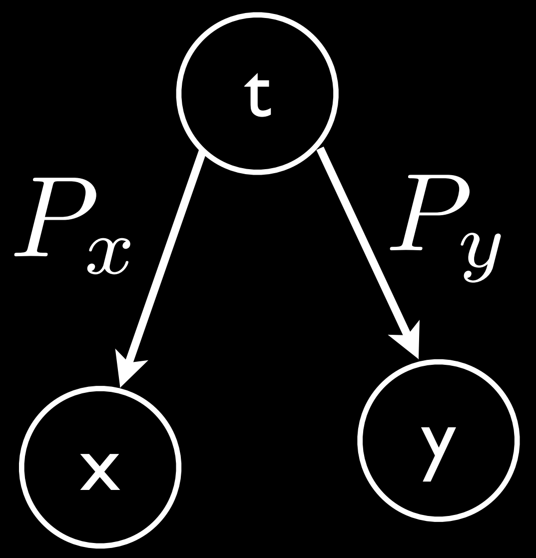

[1] A. Miu, G. Tan, H. Balakrishnan and J. Apostolopoulos, "Divert: fine-grained path selection for wireless LANs," MobiSys 2004.

[2] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall and J. Zahorjan, "Measurement-based models for delivery and interference in static wireless networks," SIGCOMM CCR 2006.

[3] R. Laufer, H. D.-Ferriere and L. Kleinrock, "Multirate Anypath Routing in Wireless Mesh Networks," INFOCOM 2009.

χ is not the desired metric

$$\chi = P(x=0|y=0) - P(x=0)$$

✓ a) A scalar with a finite range of [-1,1]

✗ b) Symmetric

✗ c) Irrespective of PRRs:

1 for perfectly positively correlated link pair

-1 for perfectly negatively correlated link pair

χ is not symmetric

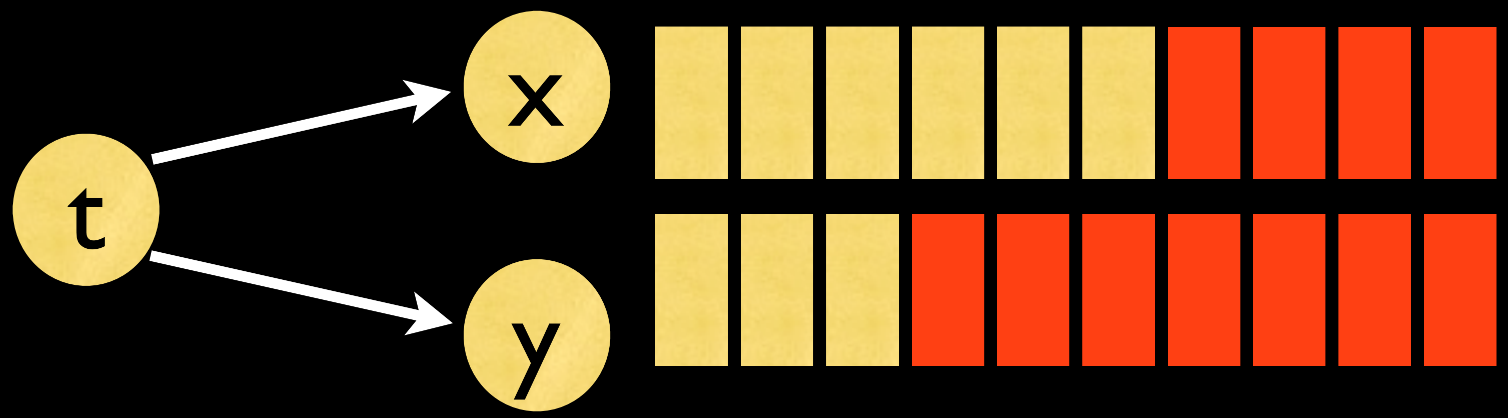

$$\chi_{x,y} = P(x=0|y=0) - P(x=0)$$

$$= 4/7 - 4/10 = 0.17$$

χ is not symmetric

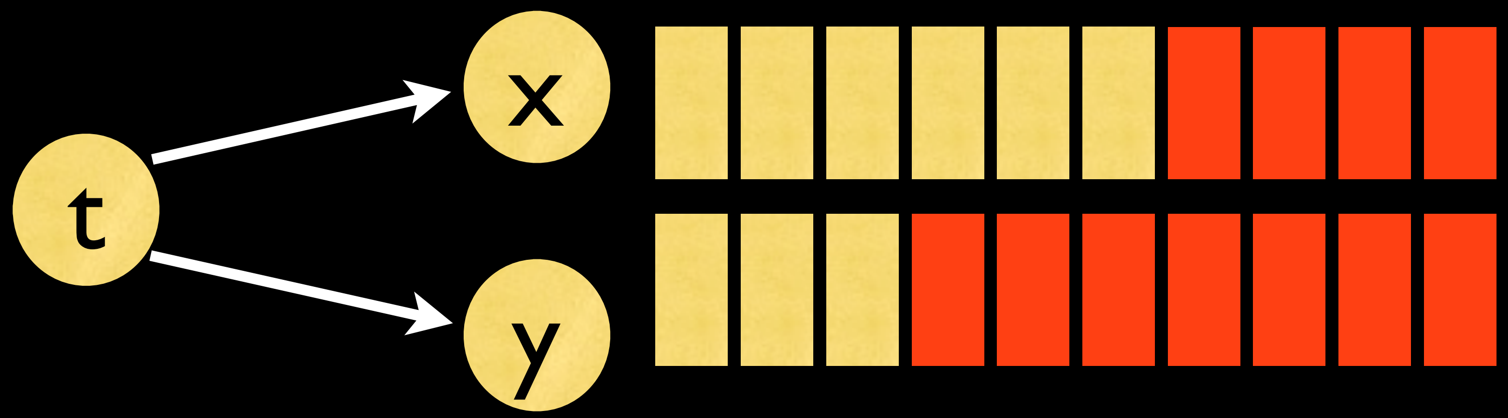

$$\chi_{x,y} = P(x=0|y=0) - P(x=0)$$

$$= 4/7 - 4/10 = 0.17$$

$$\chi_{y,x} = P(y=0|x=0) - P(y=0)$$

$$= 1 - 7/10 = 0.3$$

$$\chi_{x,y} \neq \chi_{y,x}$$

χ does not satisfy property (c)

c) 1 for perfectly positively correlated link pair

For the same PRR case (P(x=1) = P(y=1)):

$$\chi = P(x=0|y=0) - P(x=0)$$

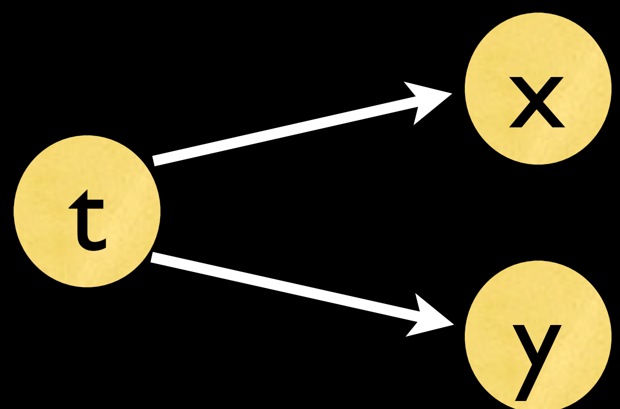

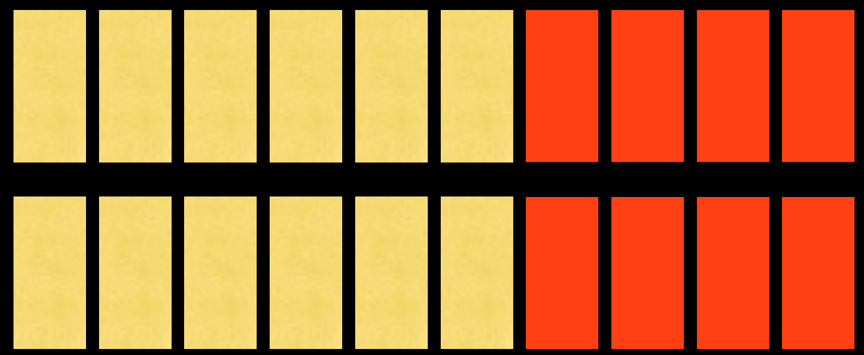

P(x=0|y=0) = 1

$$\chi = P(x=1) \neq 1$$

χ looks independent for low PRR link pairs

χ does not satisfy property (c)

c) 1 for perfectly positively correlated link pair

For the same PRR case (P(x=1) = P(y=1)):

$$\chi = P(x=0|y=0) - P(x=0)$$

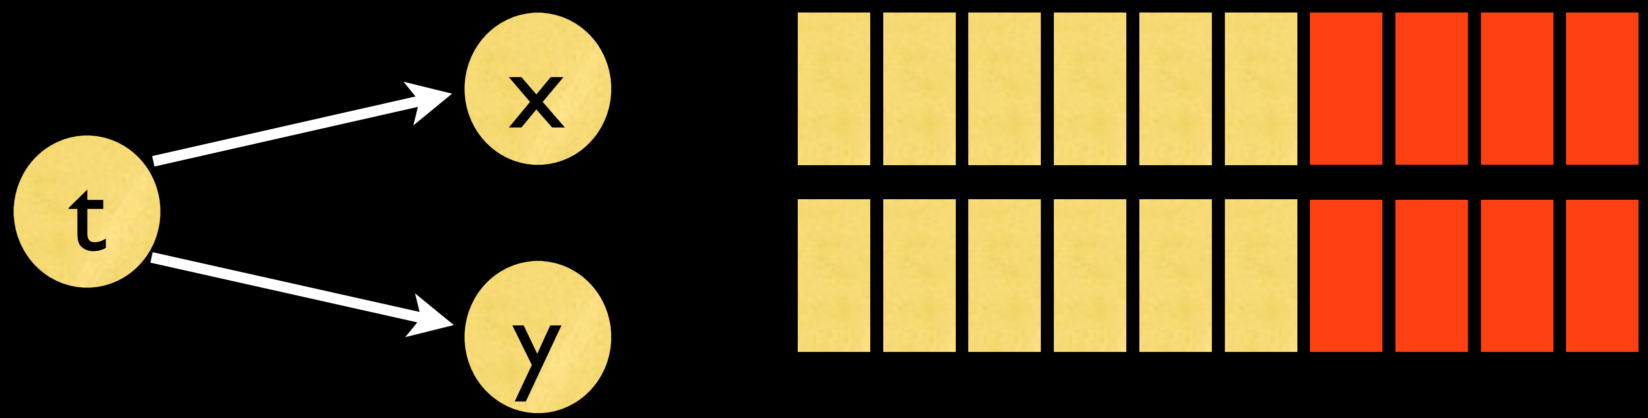

$P(x=0|y=0) = 1$

$$\chi = P(x=1) \neq 1$$

χ is not the desired metric

Cross-correlation Index: ρ

$$\rho = \begin{cases} \dfrac{P(x=1,y=1) - P(x=1).P(y=1)}{\sqrt{P(x=1)P(x=0)P(y=1)P(y=0)}}, & \prod_{a \in \{0,1\}} P(x=a)P(y=a) \neq 0 \\ 0, \text{ otherwise} \end{cases}$$

Cross-correlation Index: ρ

$$\rho = \begin{cases} \dfrac{P(x=1,y=1) - P(x=1).P(y=1)}{\sqrt{P(x=1)P(x=0)P(y=1)P(y=0)}}, & \prod_{a\in\{0,1\}} P(x=a)P(y=a) \neq 0 \\ 0, \text{otherwise} \end{cases}$$

✓ a) A scalar with a finite range of [-1,1]

✓ b) Symmetric

✗ c) Irrespective of PRRs:

1 for perfectly positively correlated link pair

-1 for perfectly negatively correlated link pair

ρ does not satisfy property (c)

c) 1 for perfectly positively correlated link pair

For the same PRR case ($P_x(1) = P_y(1)$):

$$\rho = \frac{P_{x,y}(1,1) - P_x^2(1)}{P_x(1).P_x(0)}$$

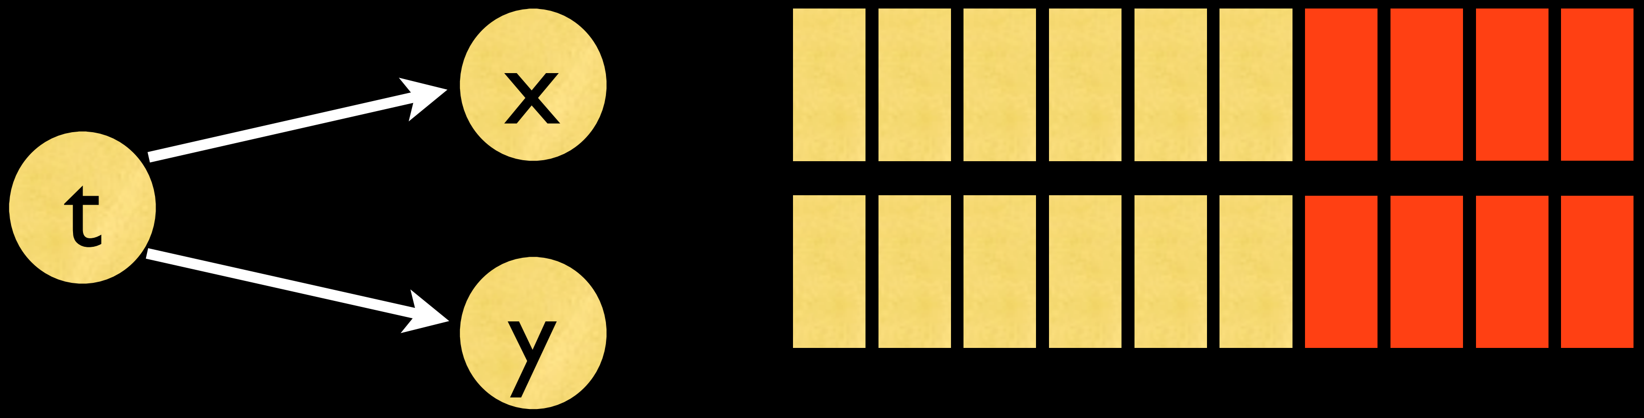

$P_{x,y}(1,1) = P_x(1)$

$\rho = 1$

Works when PRRs are same

ρ does not satisfy property (c)

c) 1 for perfectly positively correlated link pair

For the different PRR case ($P_x(1) \neq P_y(1)$):

$$P_{x,y}(1,1) = \min(P_x(1), P_y(1))$$

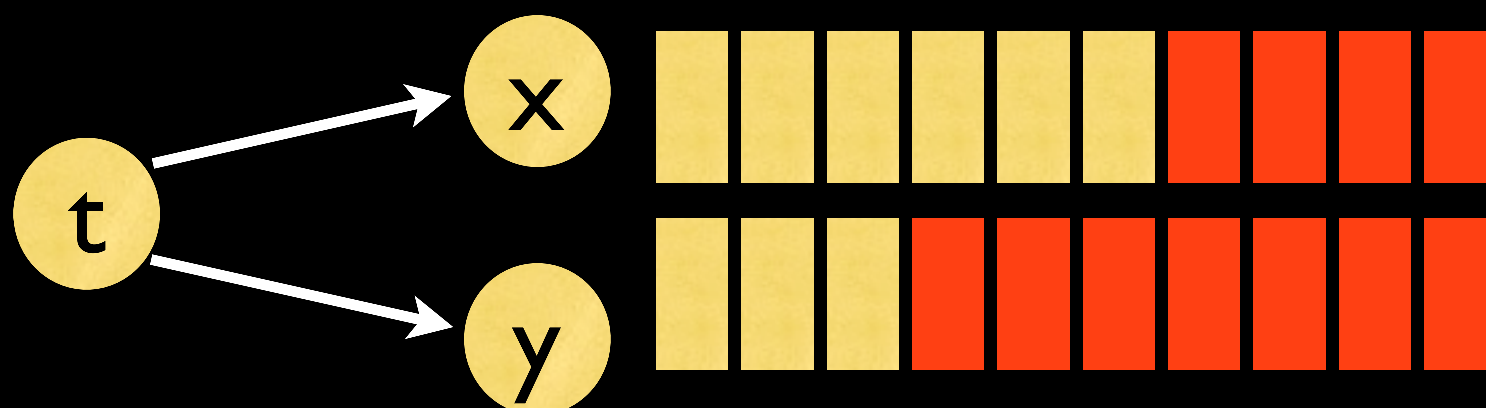

$$\rho = \sqrt{\frac{P_x(0).P_y(1)}{P_x(1).P_y(0)}} \neq 1$$

Does NOT work when PRRs are different

ρ does not satisfy property (c)

- Similarly, for perfectly negatively correlated link pairs:
  - is -1: only when PRRs sum to 1
  - does not work for other cases

ρ is not the desired metric

Outline

- Desired Metric Properties
- The κ Metric
- κ 's Usefulness
- Open Questions

New Metric: κ

- ρ almost satisfied all the desired properties
- Normalizing ρ satisfies all the properties:

$$\kappa = \begin{cases} \dfrac{\rho}{\rho_{max}} & \text{, if } \rho>0 \\ \dfrac{-\rho}{\rho_{min}} & \text{, if } \rho<0 \\ 0 & \text{, otherwise} \end{cases}$$

New Metric: κ

$$\kappa = \begin{cases} \dfrac{\rho}{\rho_{max}} & , \text{ if } \rho>0 \\ \dfrac{-\rho}{\rho_{min}} & , \text{ if } \rho<0 \\ 0 & , \text{ otherwise} \end{cases}$$

$\kappa \in [-1.0, 1.0], \forall\ Px, Py \in (0, 1)$

$\kappa = 0$: independent pairs

$\kappa > 0$: positively correlated pairs

- 1: perfectly positively correlated pairs

$\kappa < 0$: negatively correlated pairs

- -1: perfectly negatively correlated pairs

κ on Mirage

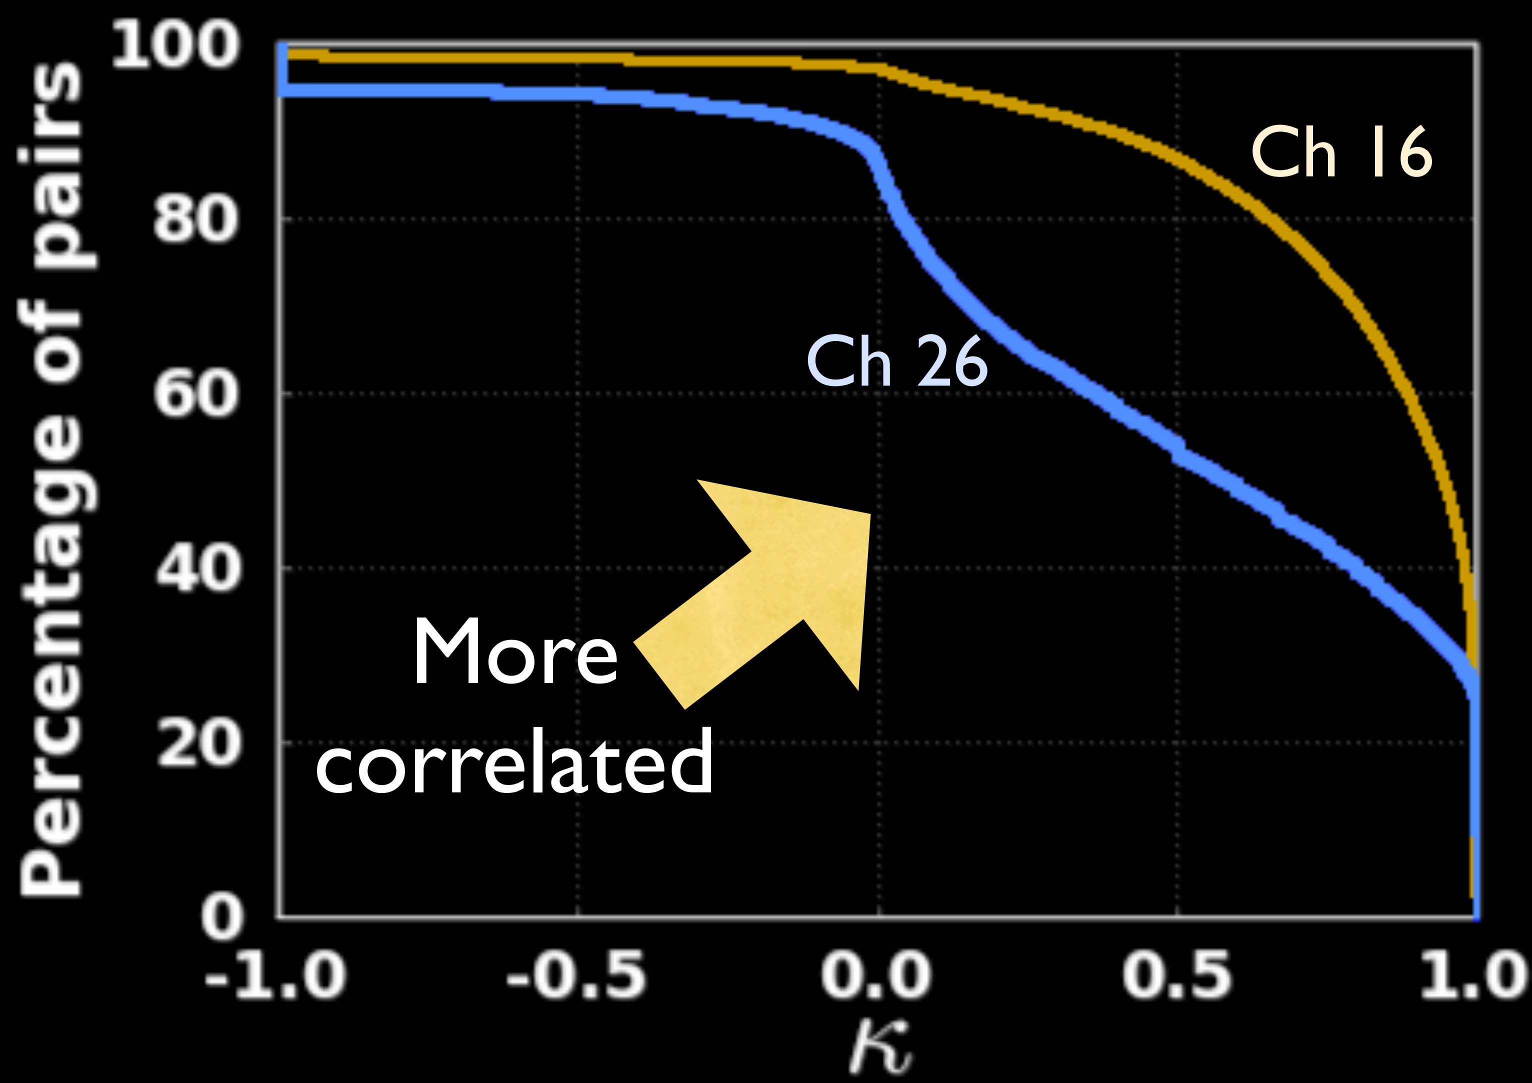

802.15.4 (Mirage)

WiFi (802.11) and 802.15.4 Spectrum

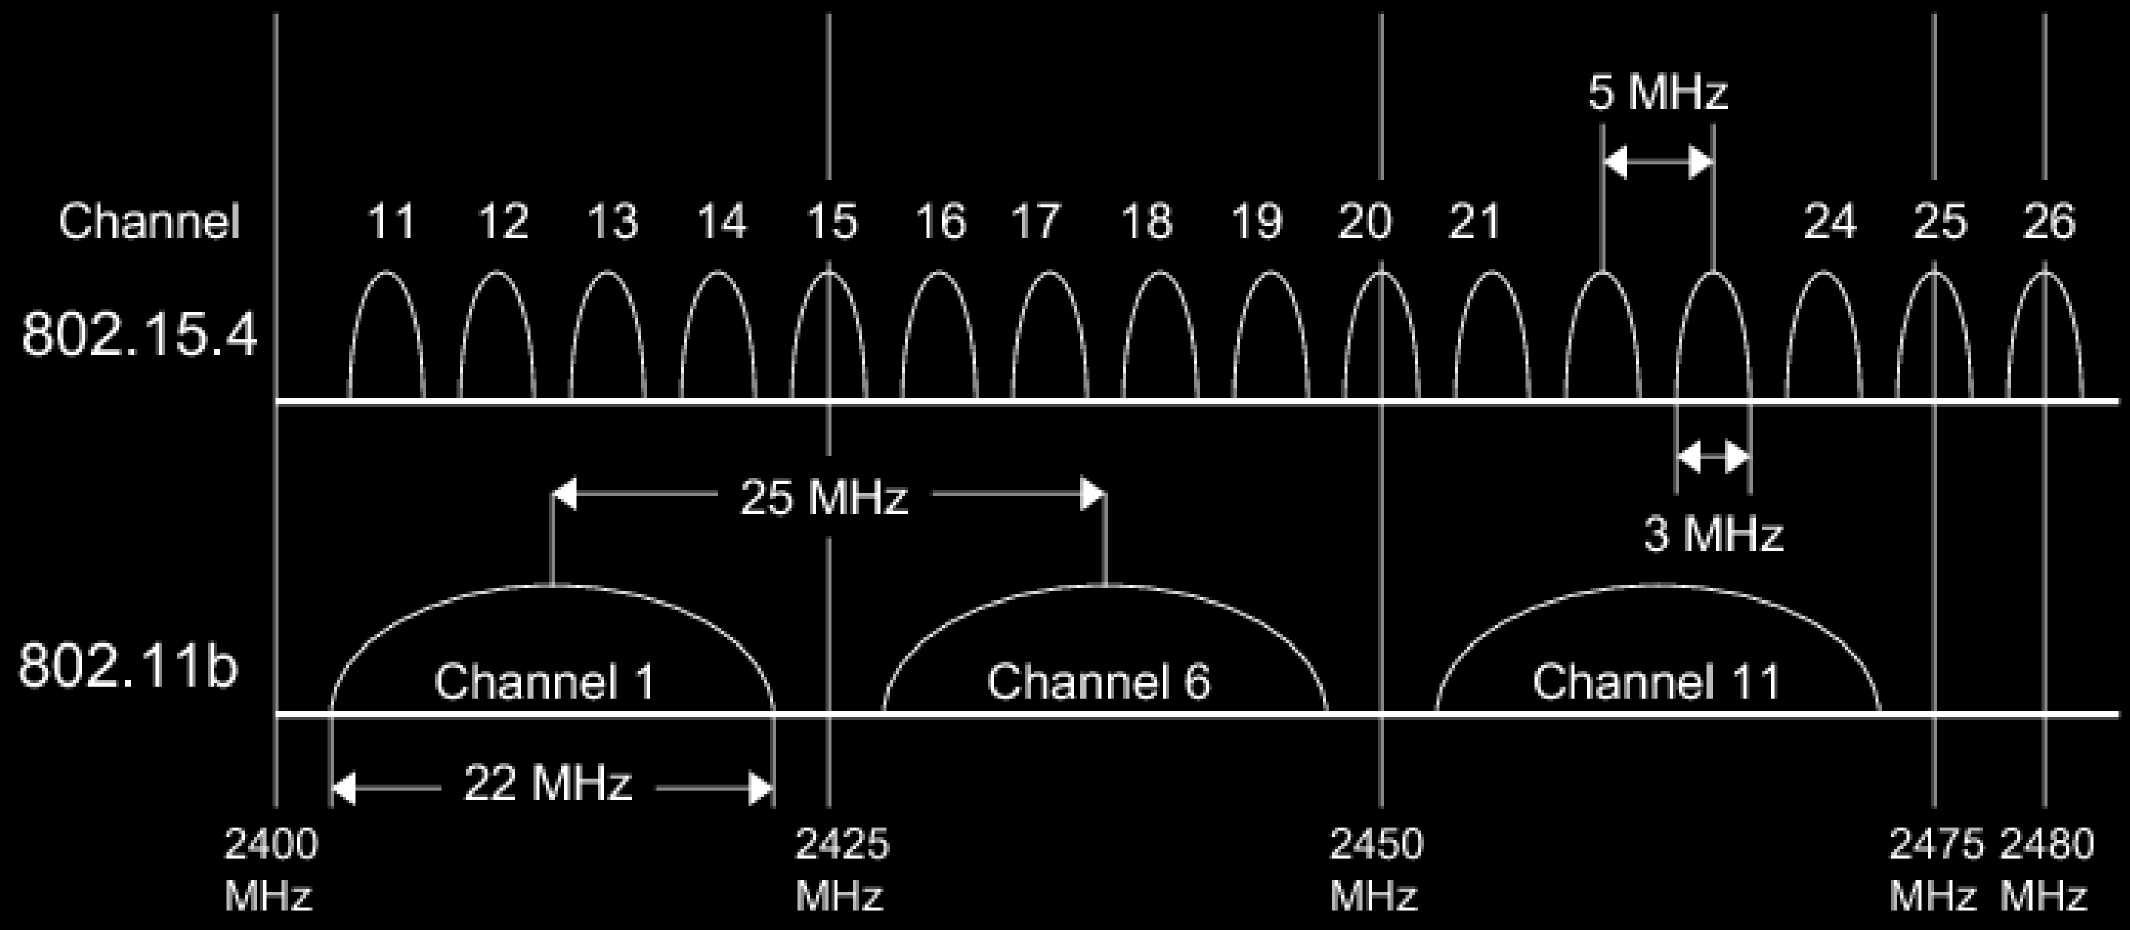

WiFi (802.11) and 802.15.4 Spectrum

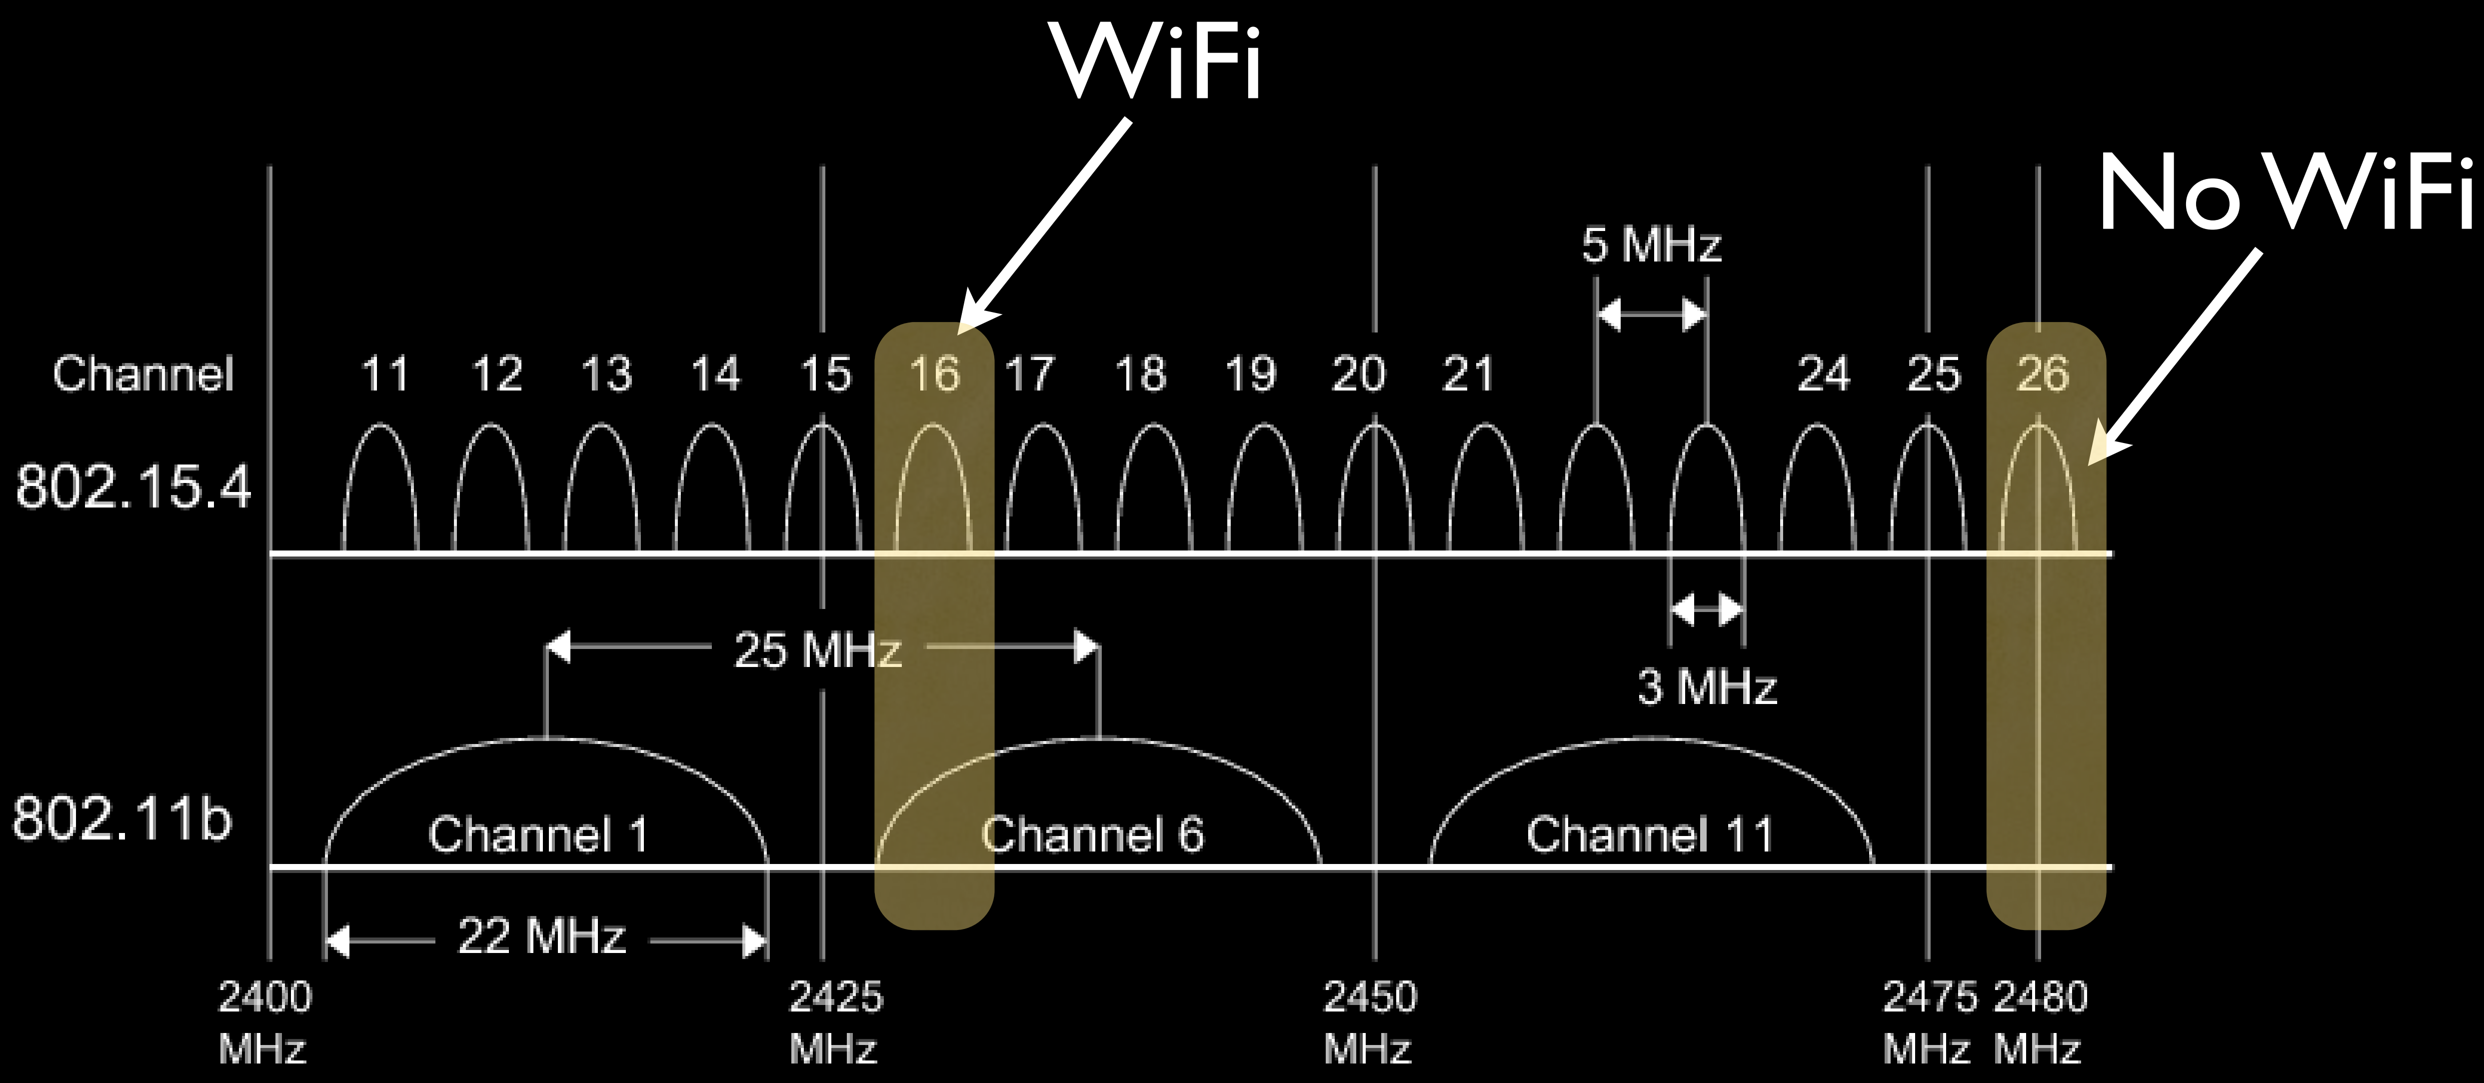

Outline

- Desired Metric Properties
- The κ Metric
- κ 's Usefulness
- Open Questions

How useful is κ?

Dissemination:

- Deliver large data to all nodes
  Eg. SPIN, RBP
- Deluge (standard protocol)
- Rateless Deluge (network coding)
- Compare the total time for dissemination

Dissemination

Deluge:

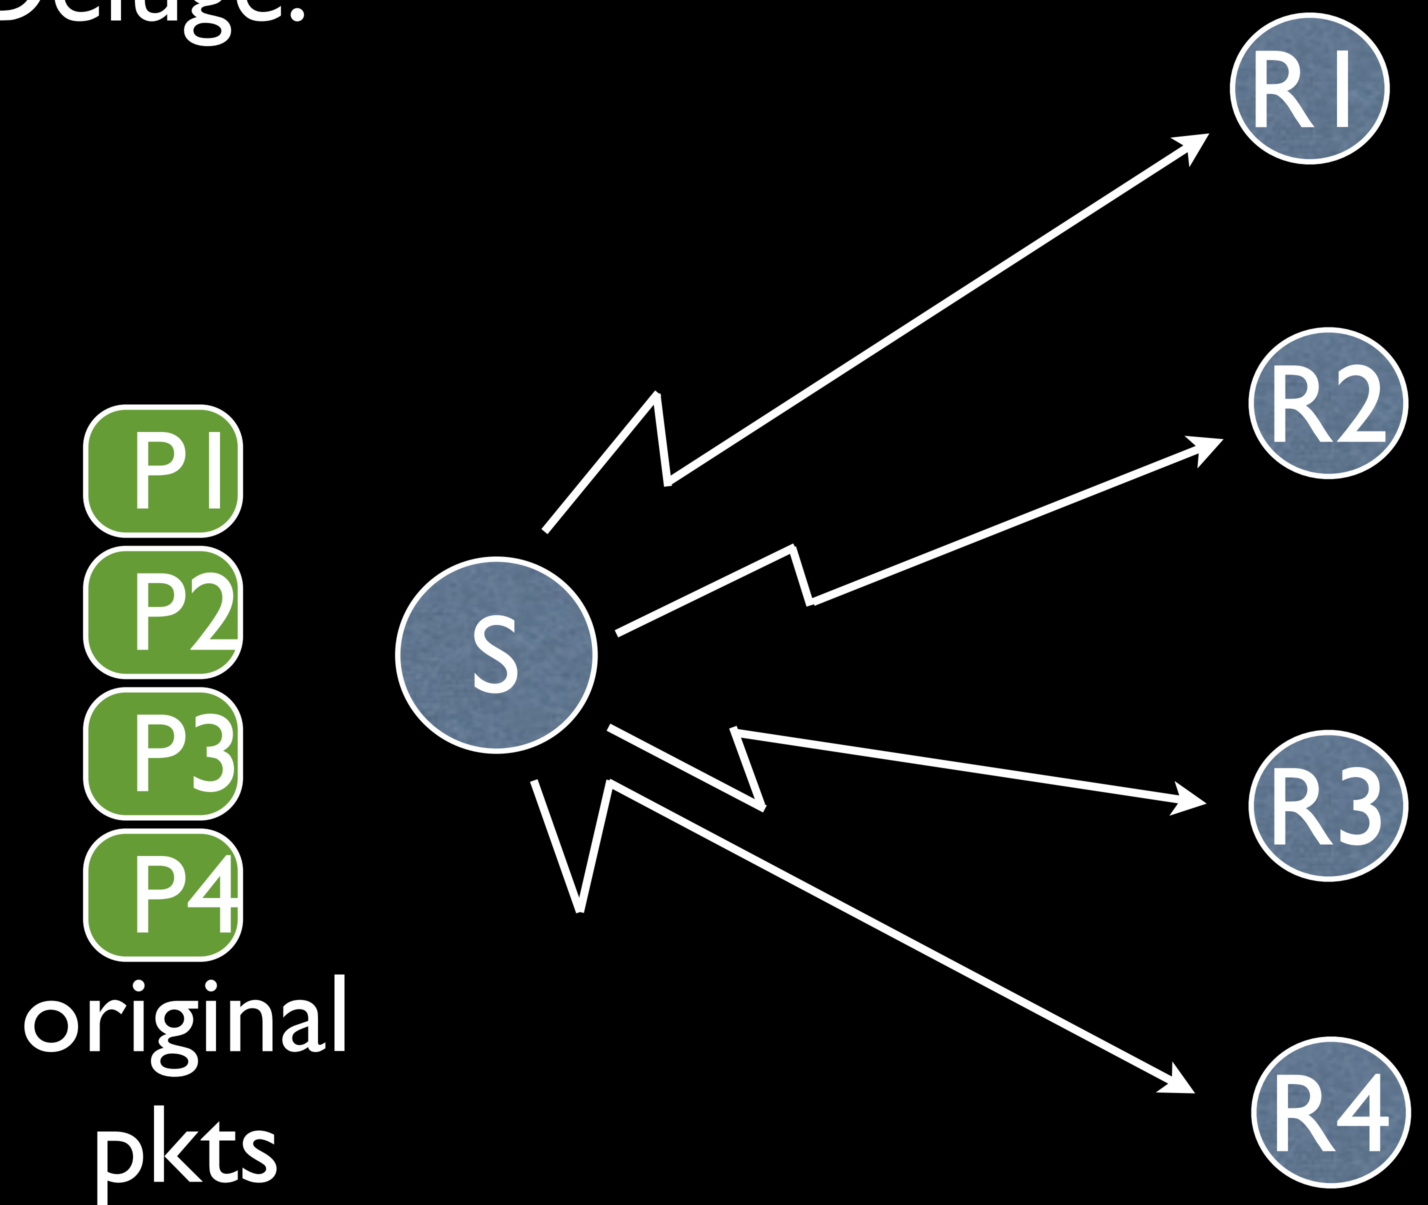

Number of transmissions from S: 4

Dissemination

Deluge:

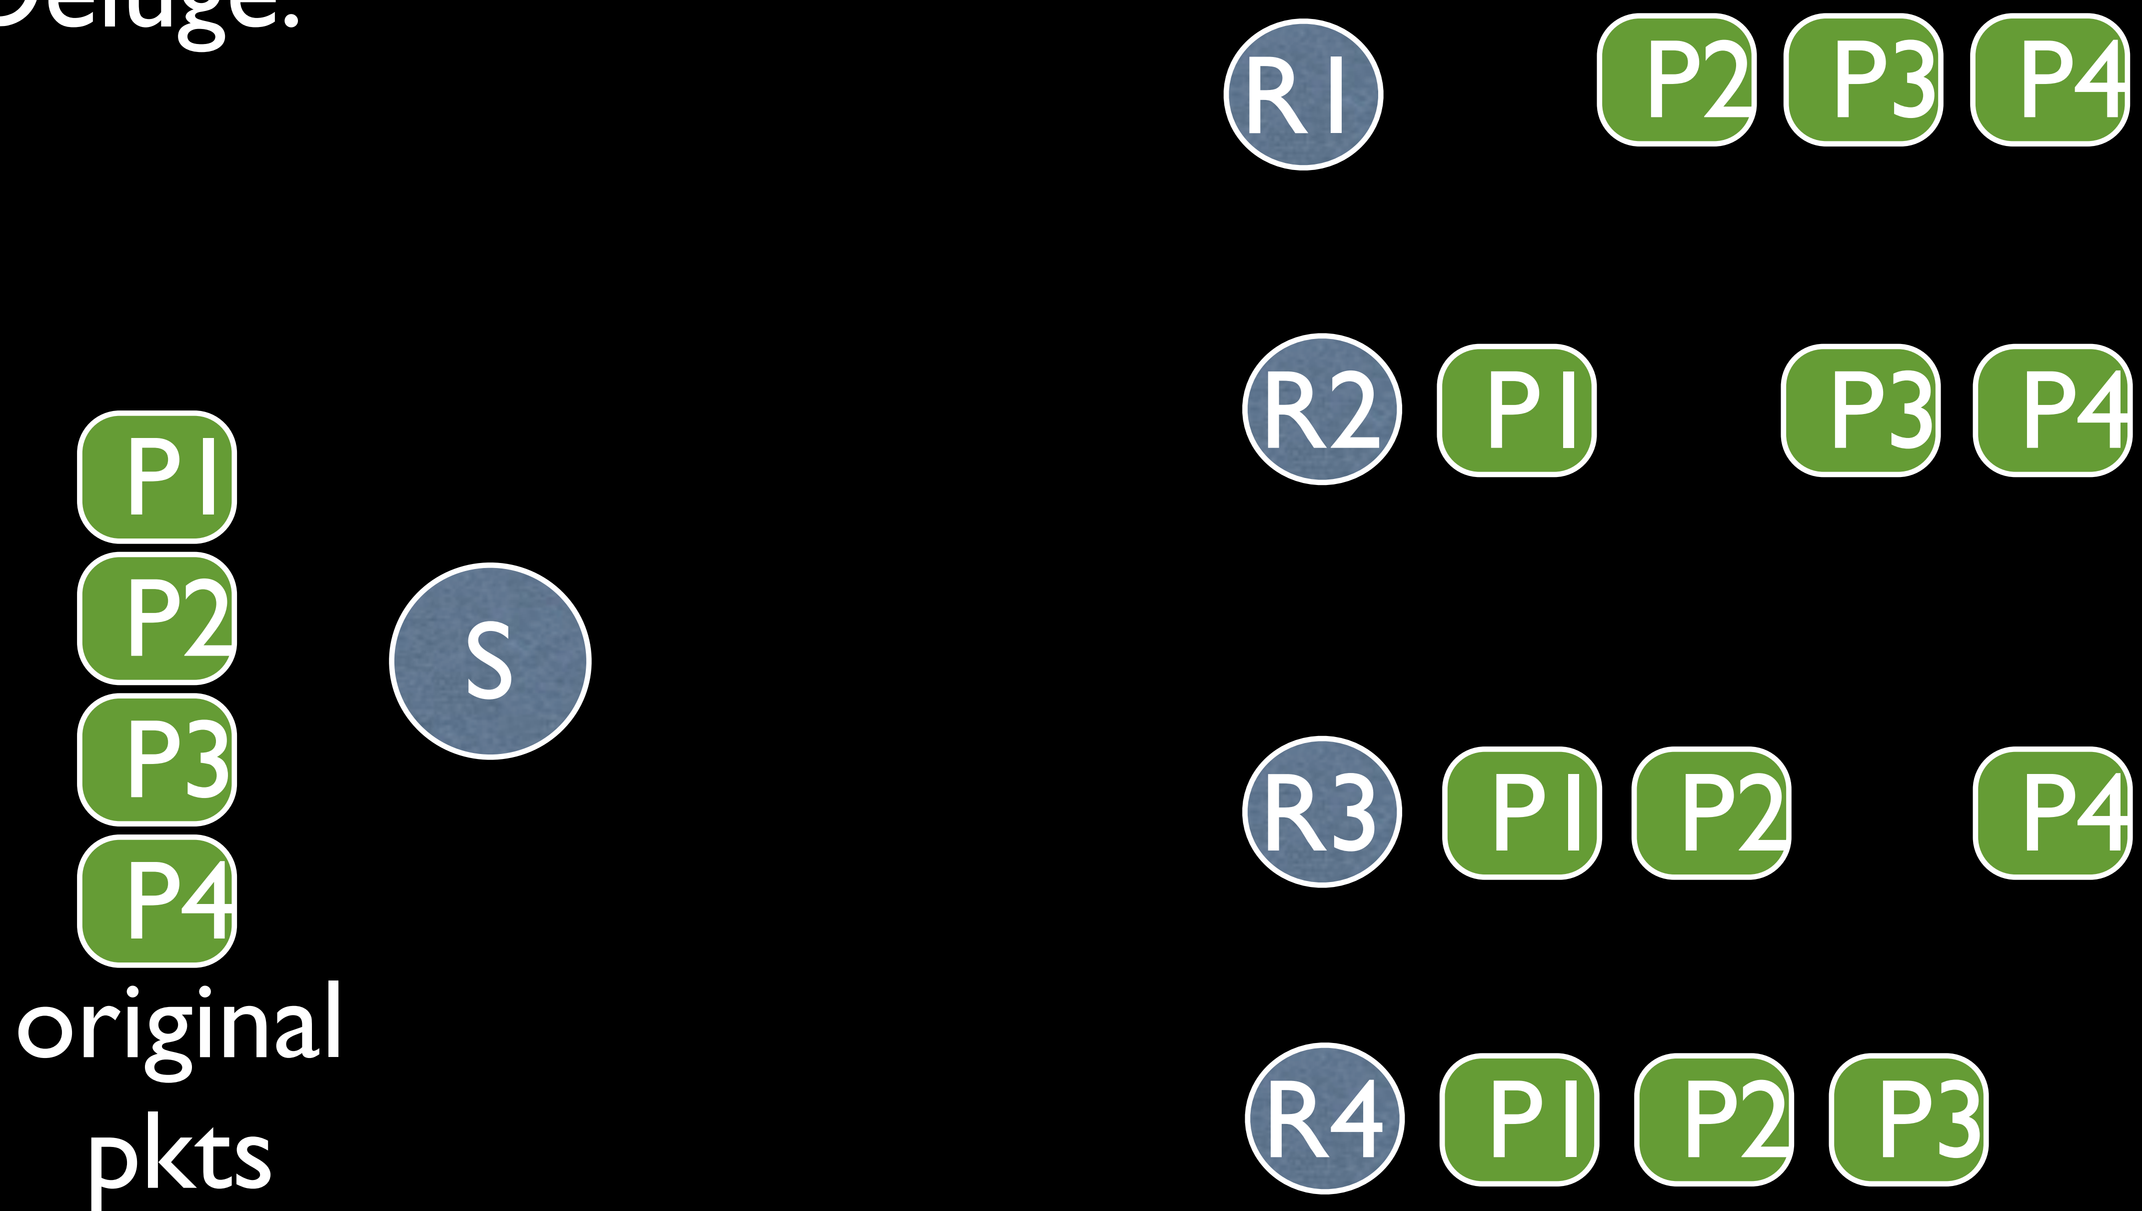

Number of transmissions from S: 4

Dissemination

Deluge:

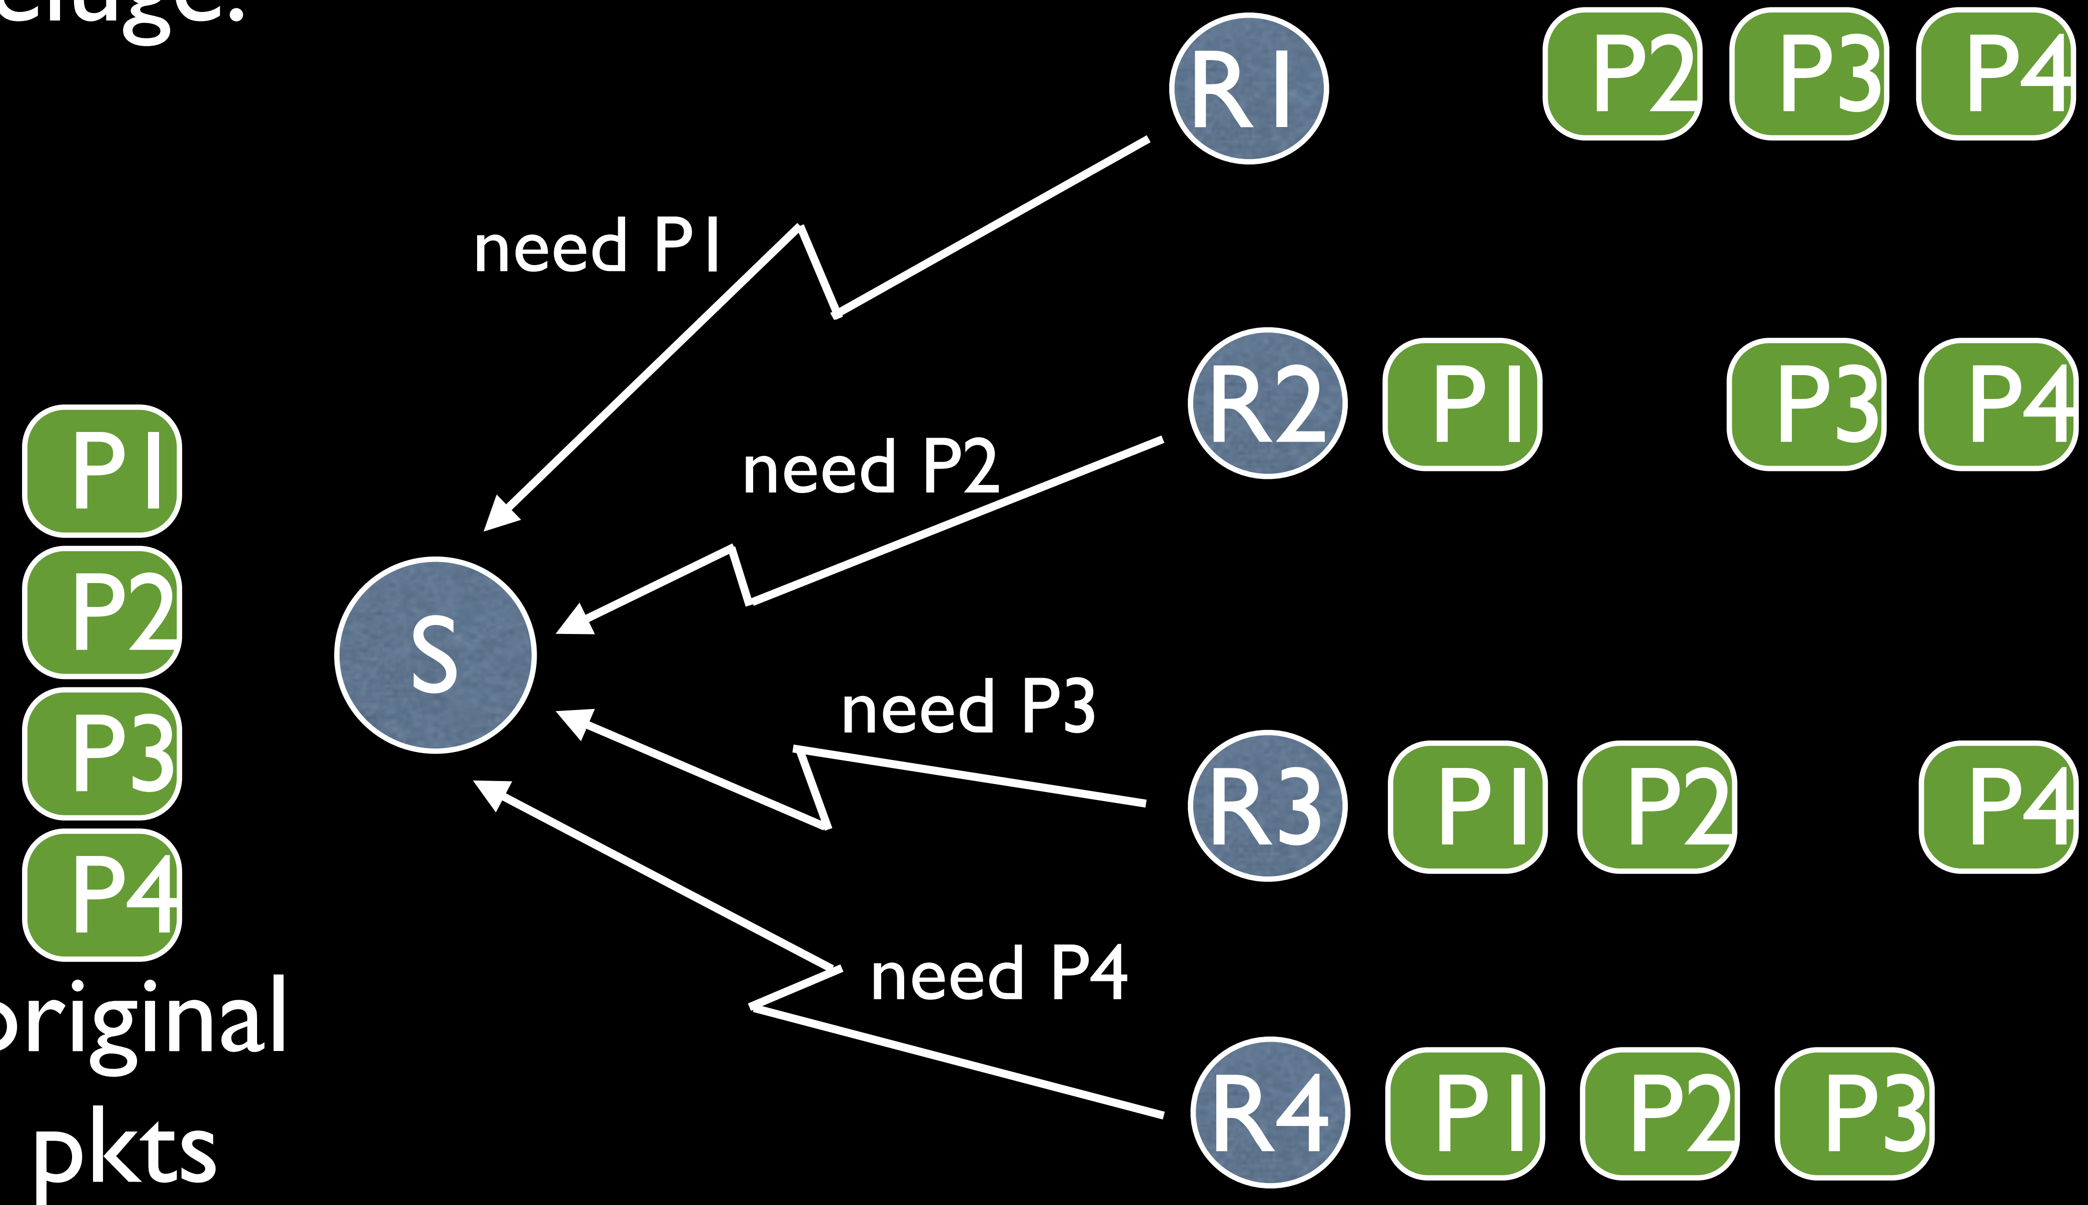

Number of transmissions from S: 4

Dissemination

Deluge:

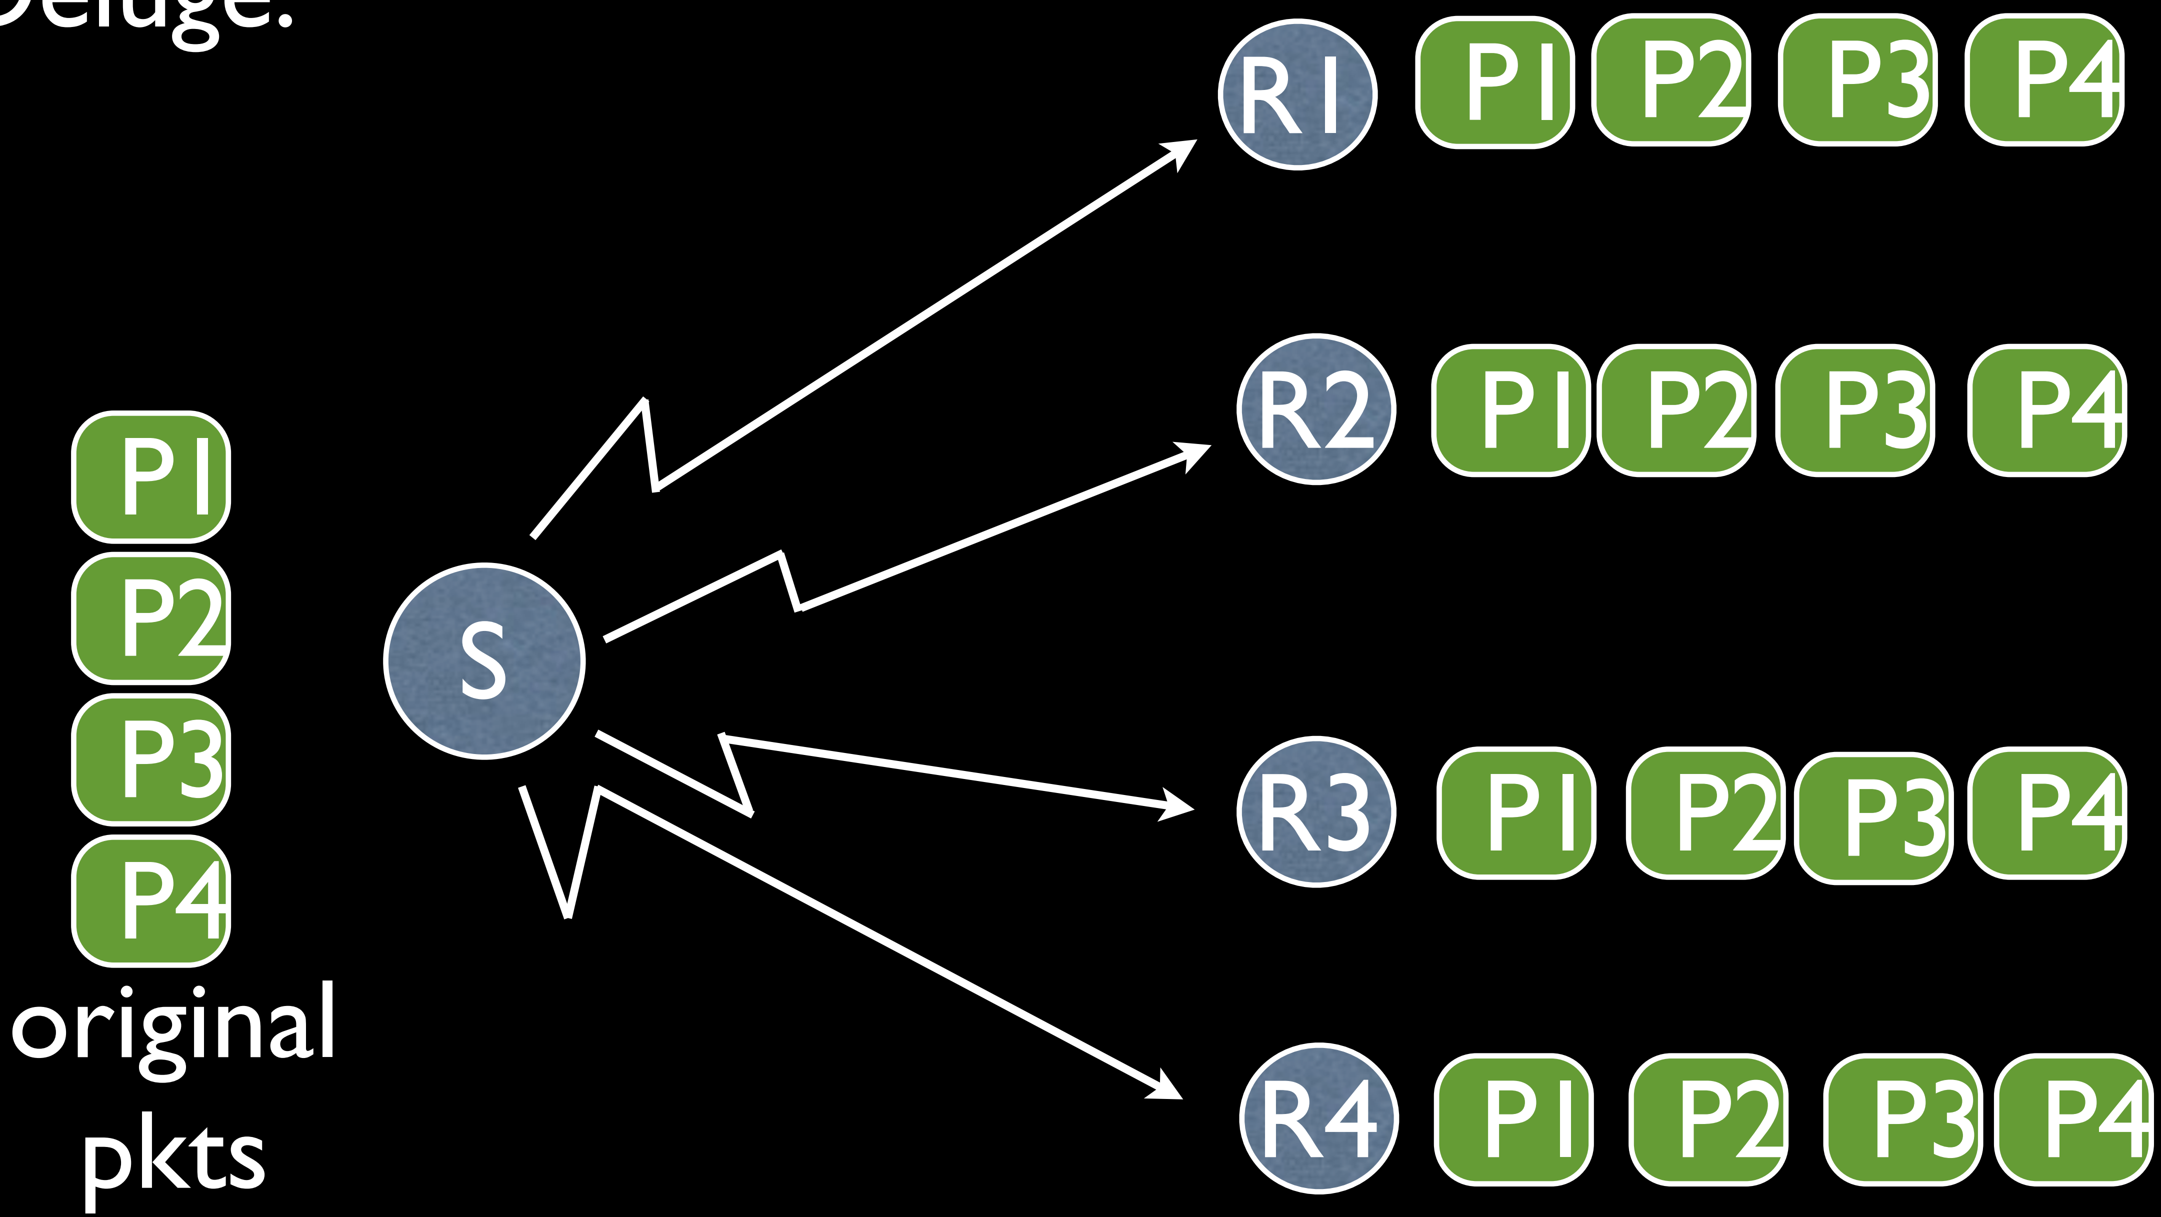

Number of transmissions from S: 4+4 = 8

Dissemination with Network Coding

Rateless Deluge:

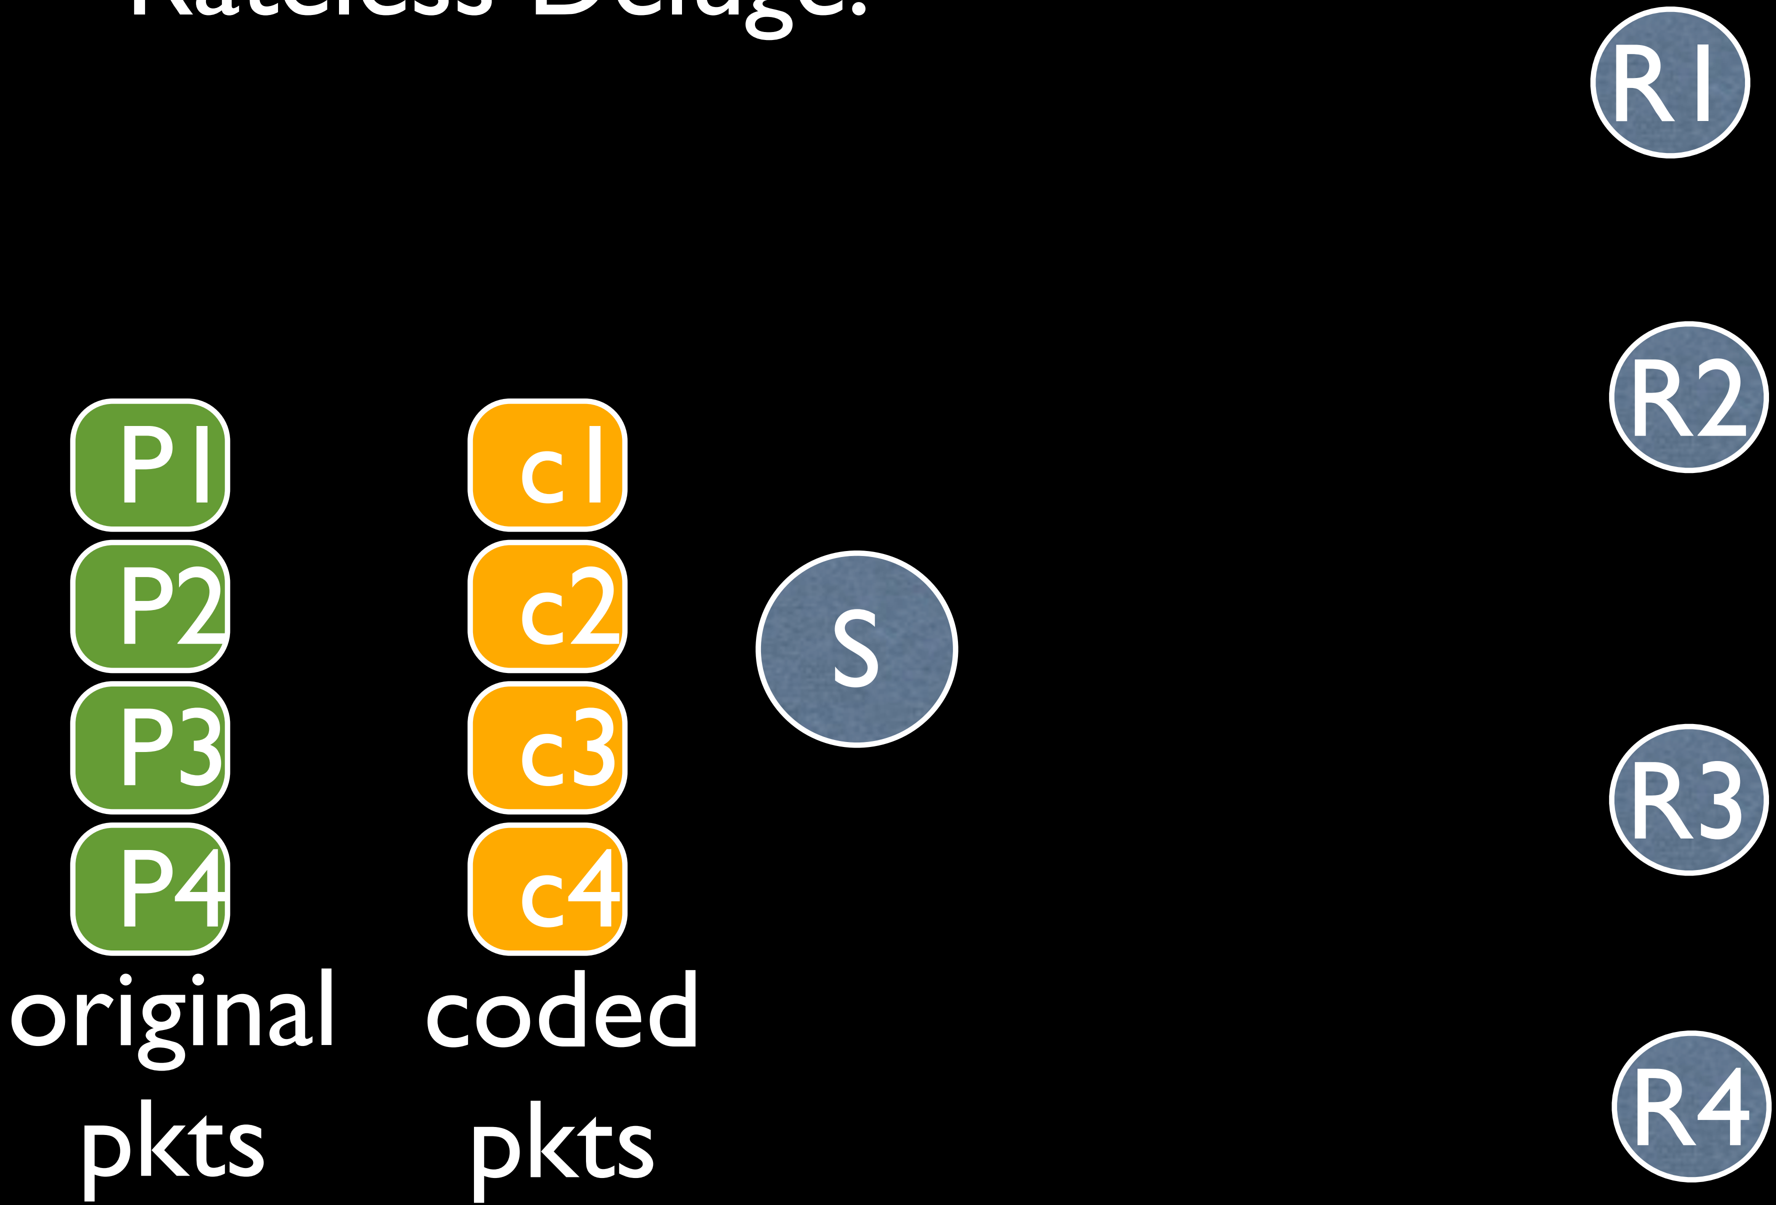

Dissemination with Network Coding

Rateless Deluge:

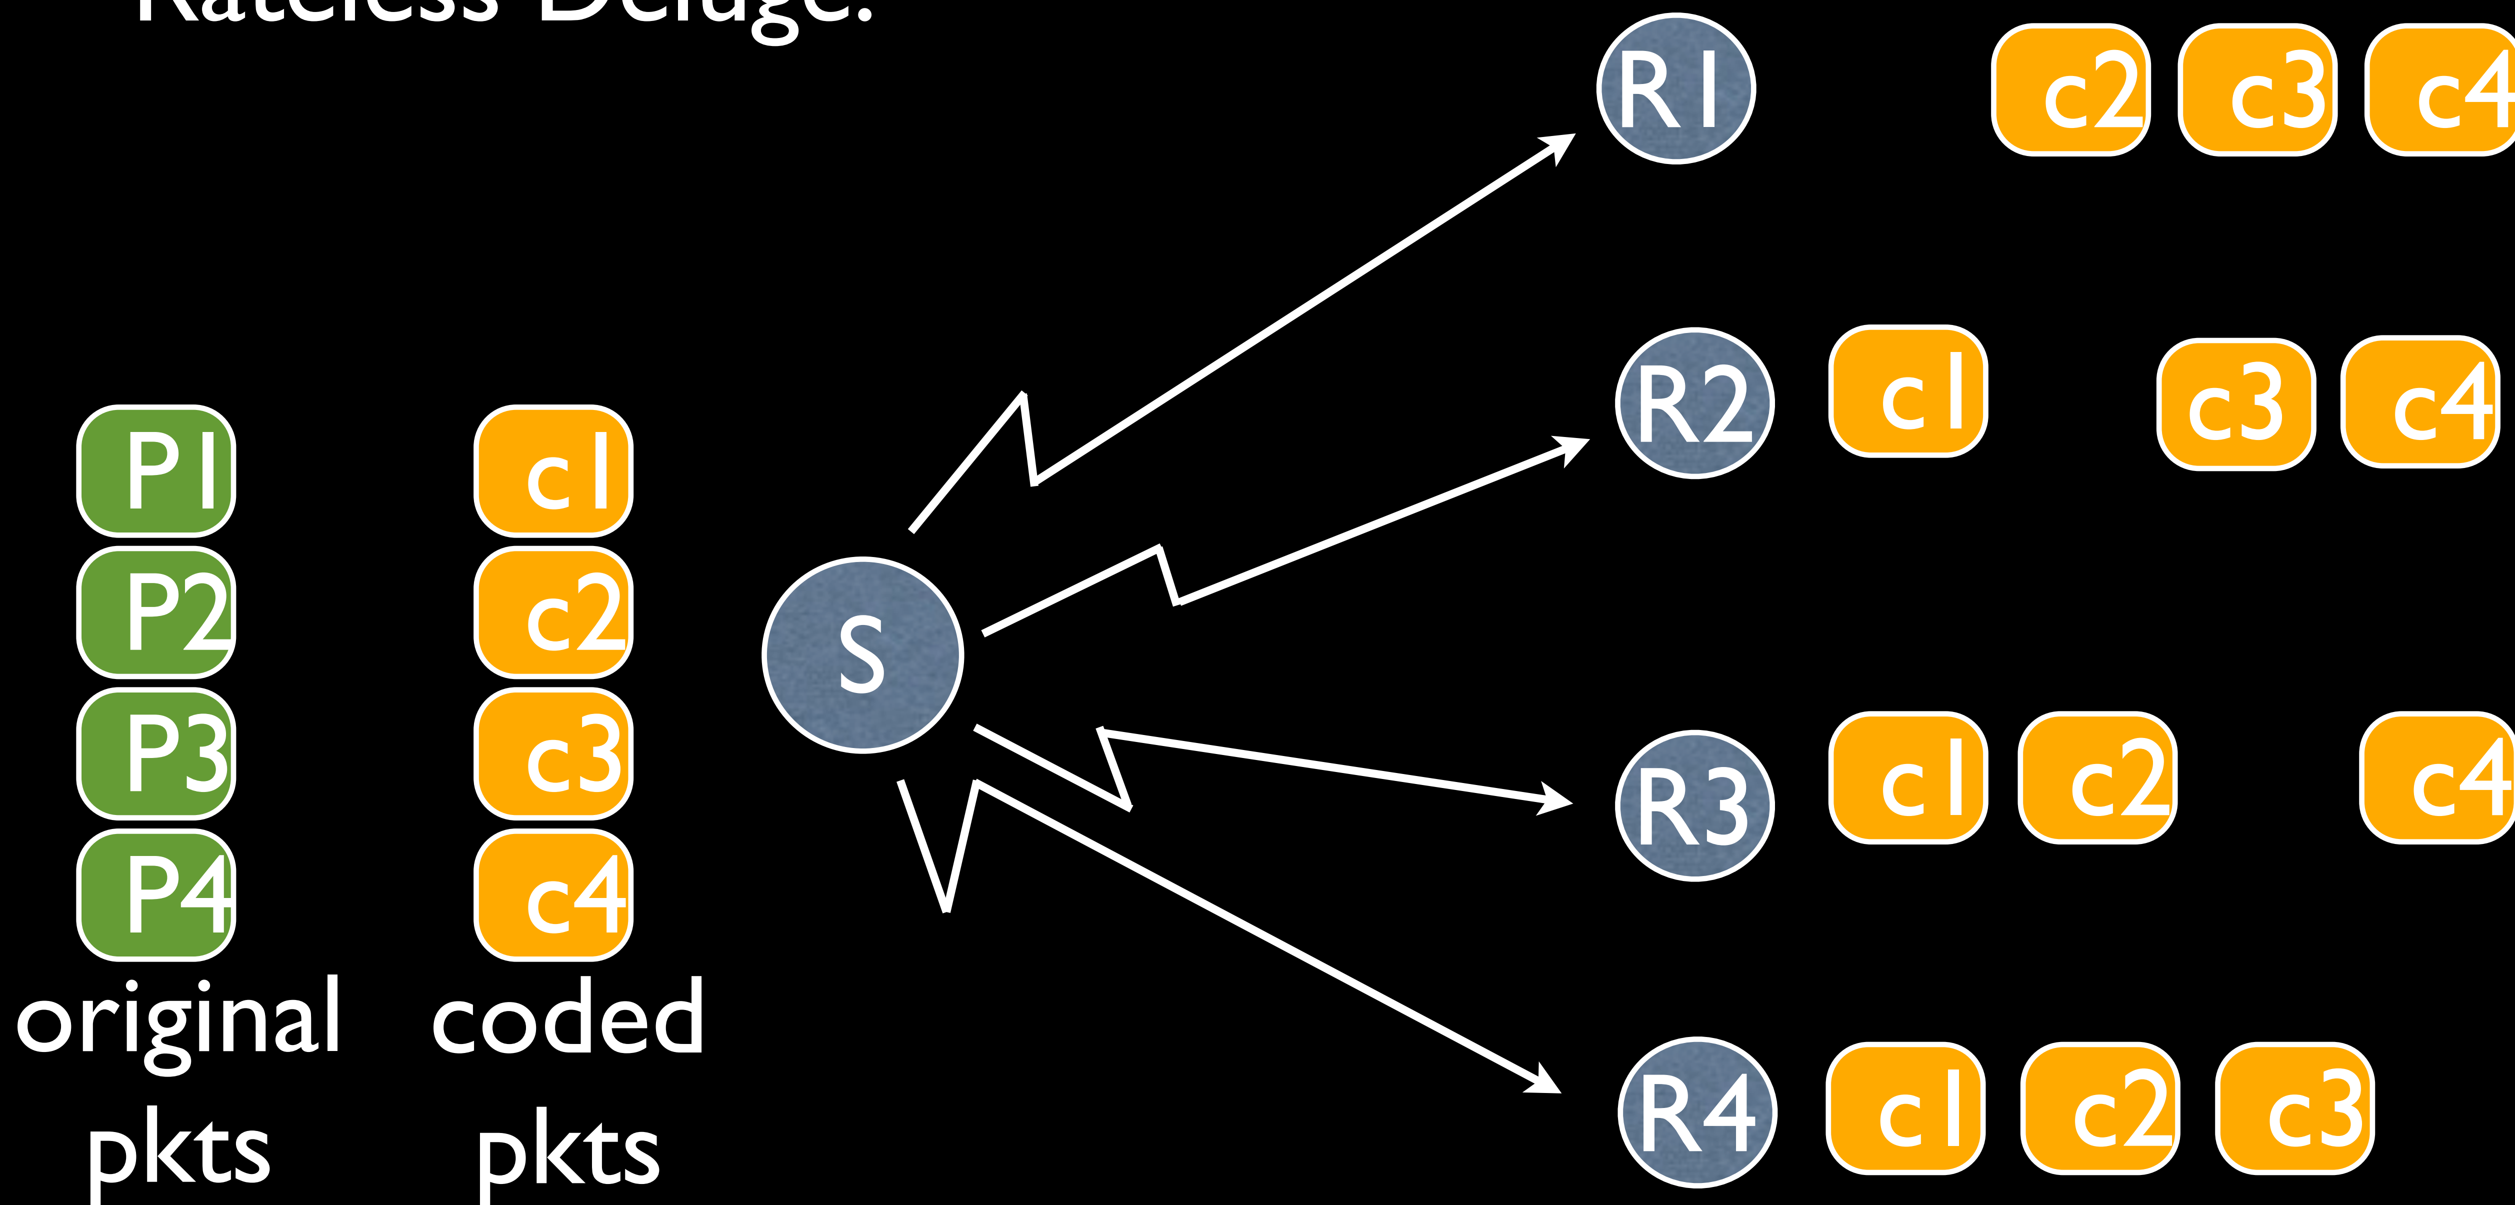

Number of transmissions from S: 4

Dissemination with Network Coding

Rateless Deluge:

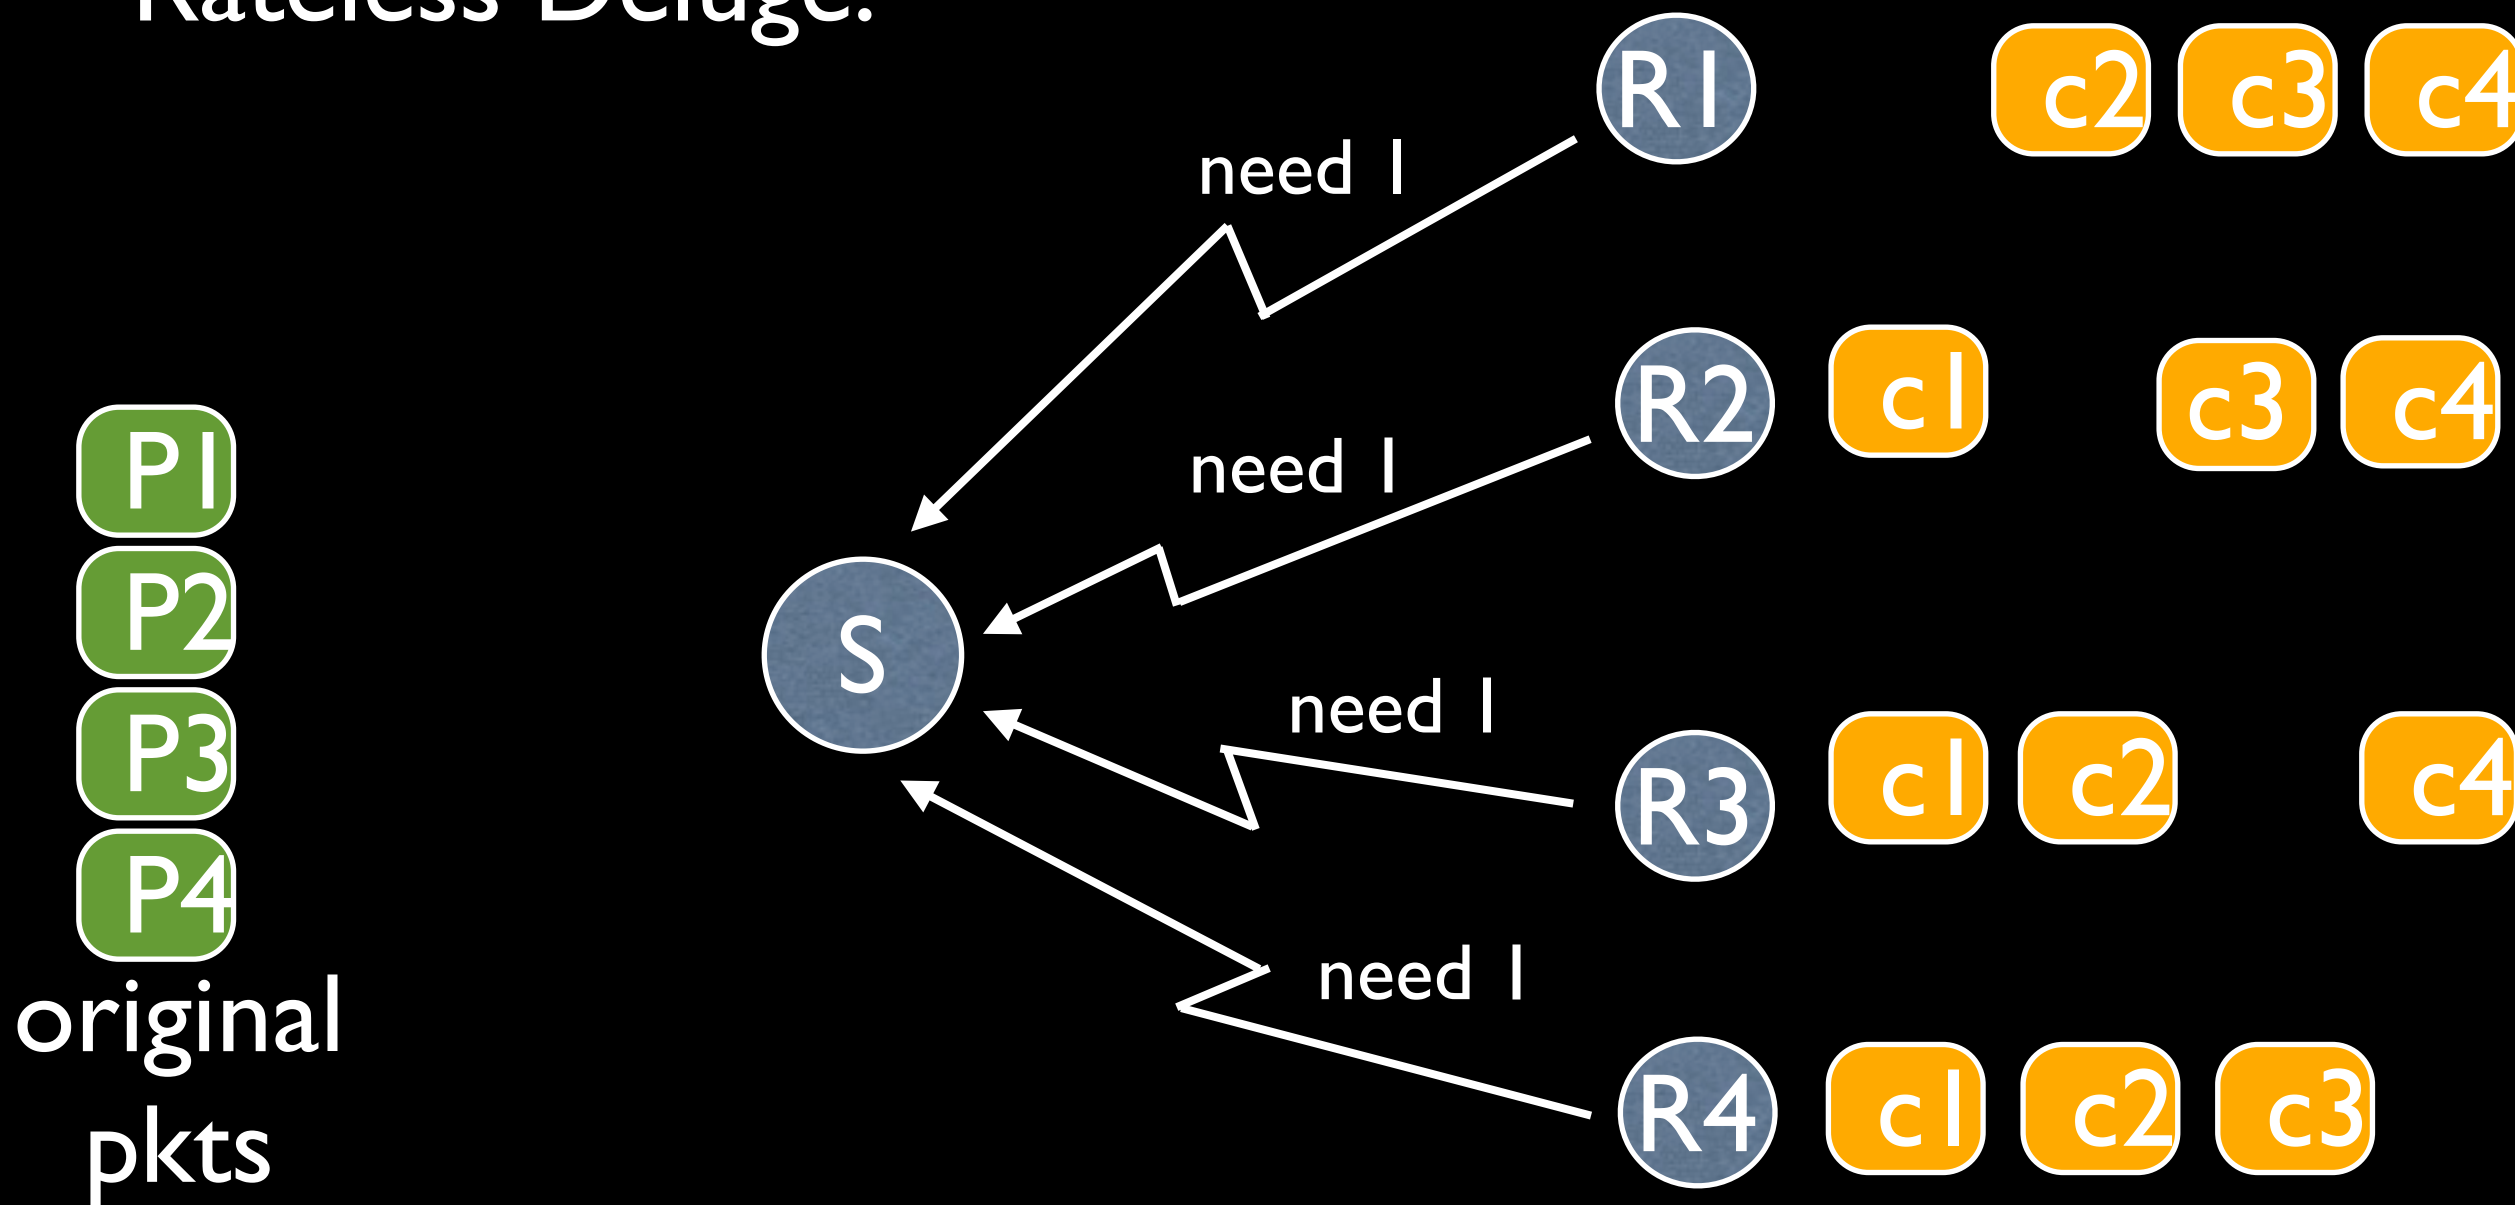

Number of transmissions from S: 4

Dissemination with Network Coding

Rateless Deluge:

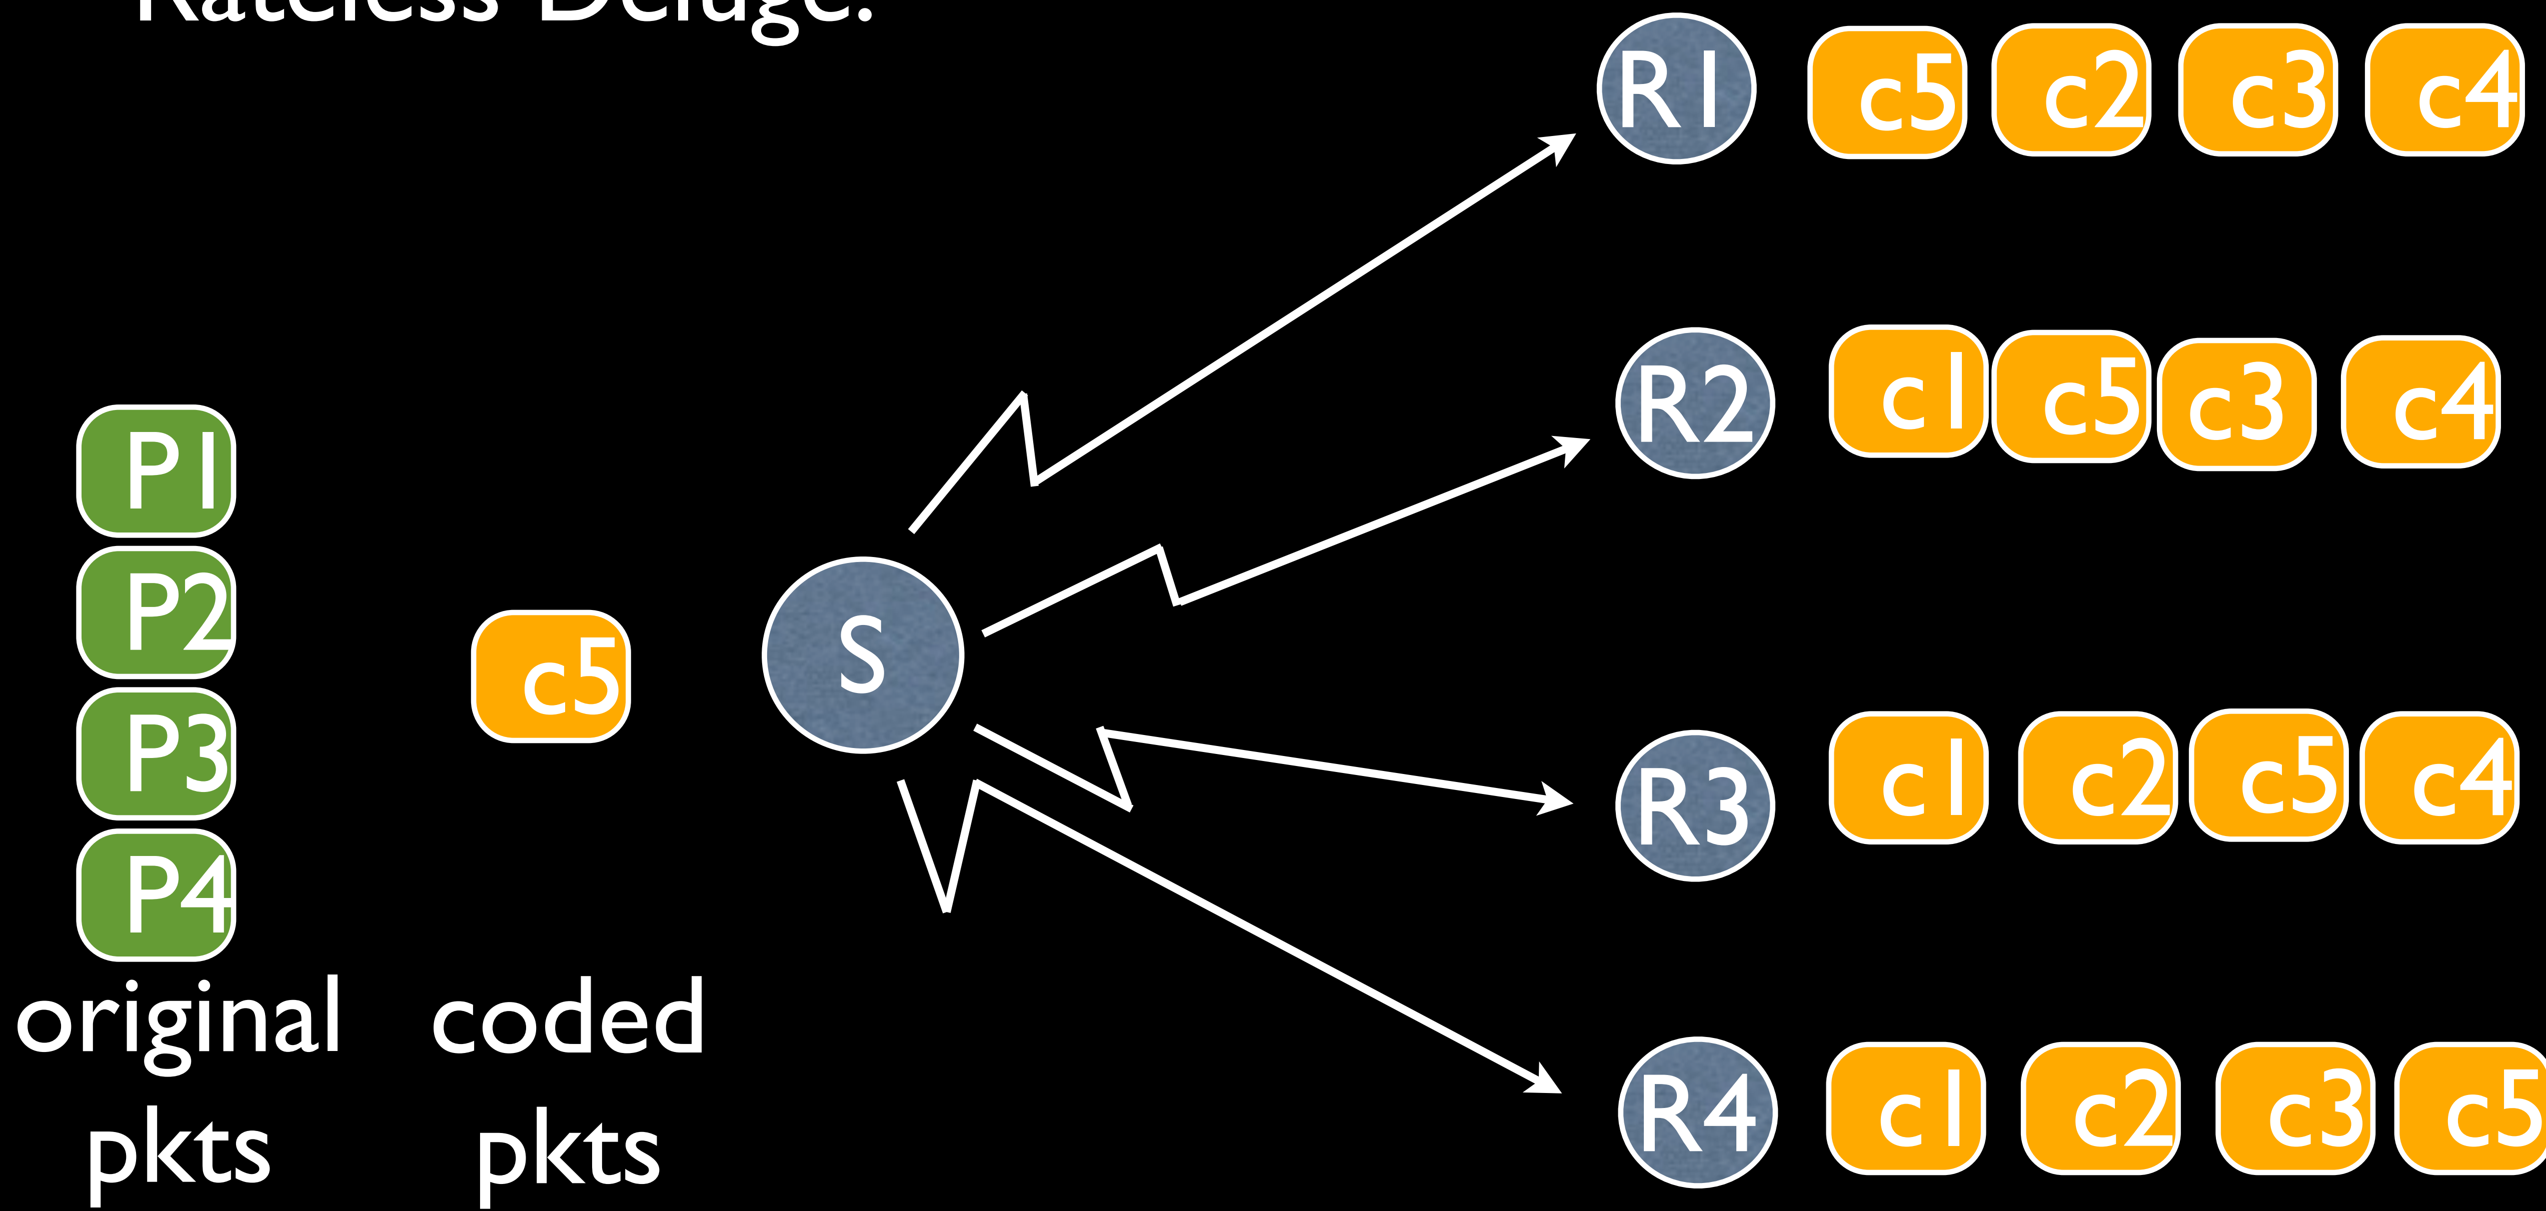

Number of transmissions from S:4+1=5

Deluge vs Rateless Deluge

Correlation Type	Deluge (# of pkts)	Rateless Deluge (# of pkts)
Perfect Negative	8	5

Rateless Deluge is great!

Deluge vs Rateless Deluge

Correlation Type	Deluge (# of pkts)	Rateless Deluge (# of pkts)
Perfect Negative	8	5
Perfect Positive	5	5

Total Dissemination Time

Correlation Type	Deluge (# of pkts)	Rateless Deluge (# of pkts)
Perfect Negative	8	5
Perfect Positive	5	5

- Total dissemination time:

  Deluge = 5p

  Rateless = 5p + 5c

p = time to send packets, c = time to code

In this case, Deluge is better!

A Controlled Study

- 1 transmitter at highest tx power
- 7 single-hop receivers with perfect reception
- Independent Losses:
  Receivers randomly (with prob. P_r) drop packets
- Correlated Losses:
  Transmitter randomly (with prob. P_t) drops packets from tx queue
- Vary P_t and P_r to vary spatial correlation and PRR of links

A Controlled Study

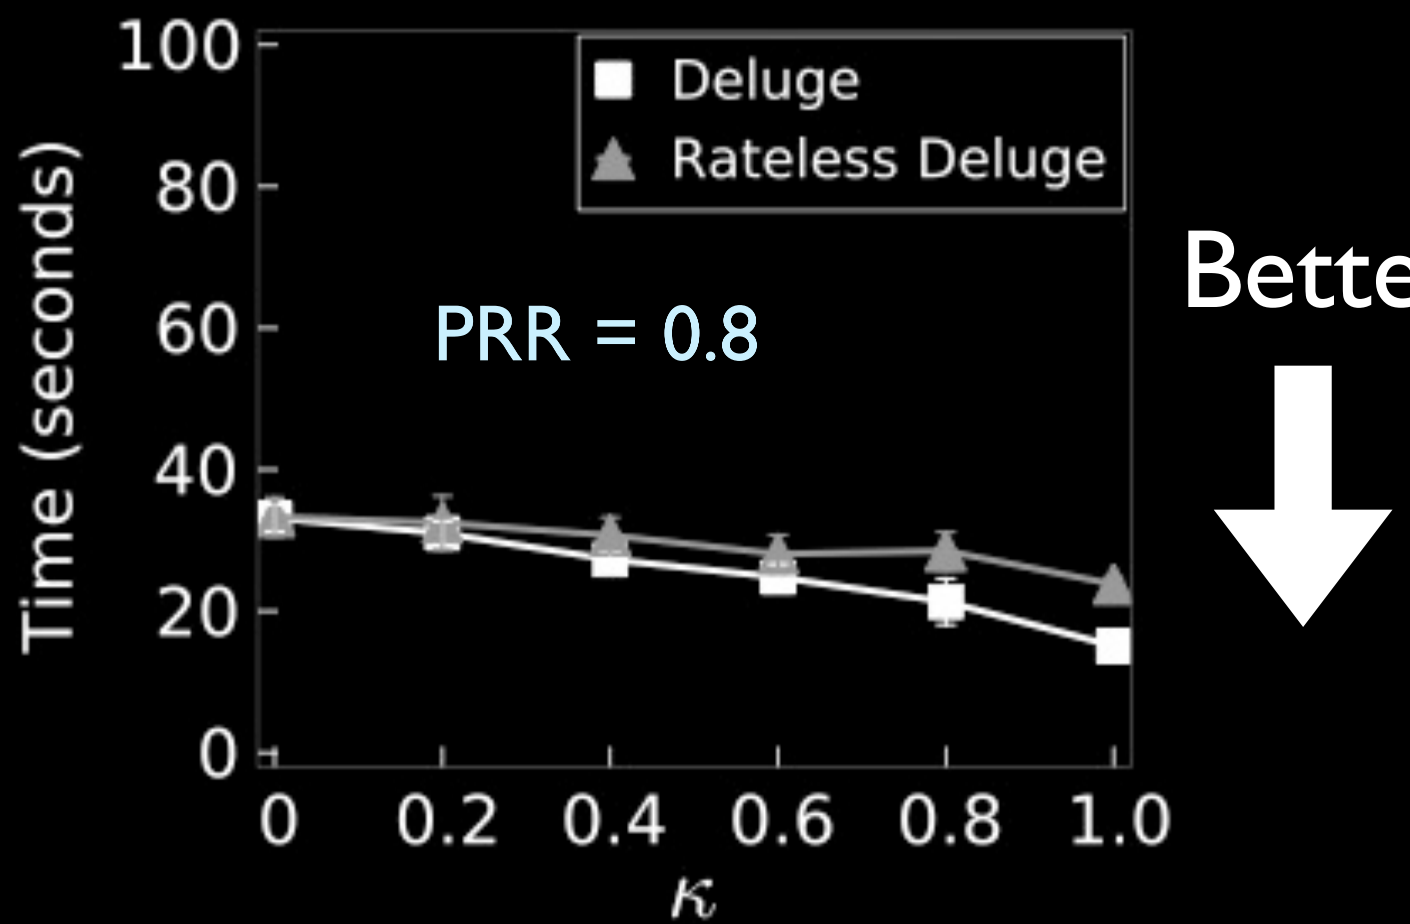

When PRR is high, Deluge is better!

A Controlled Study

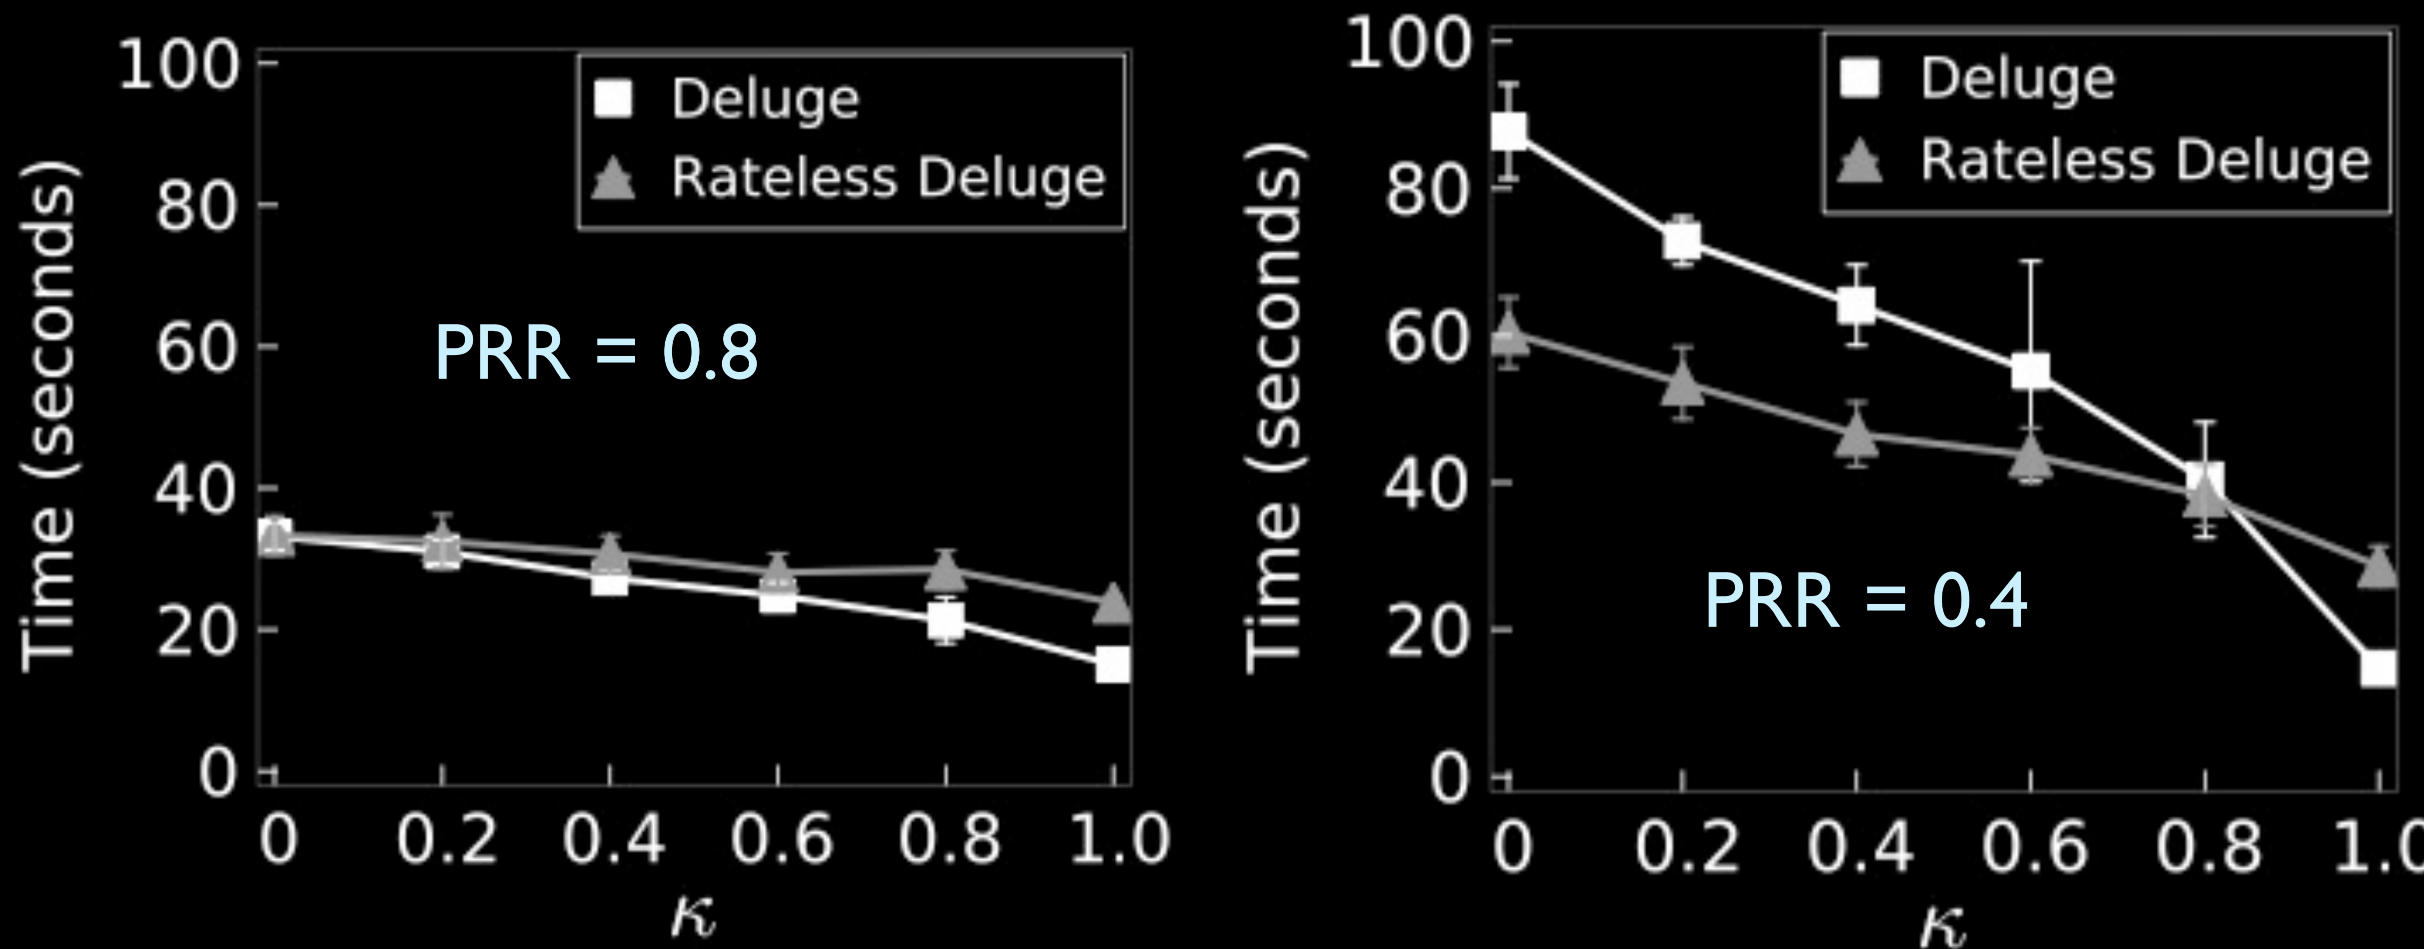

When PRR is low, Rateless Deluge is almost always better!

Dissemination Time Performance

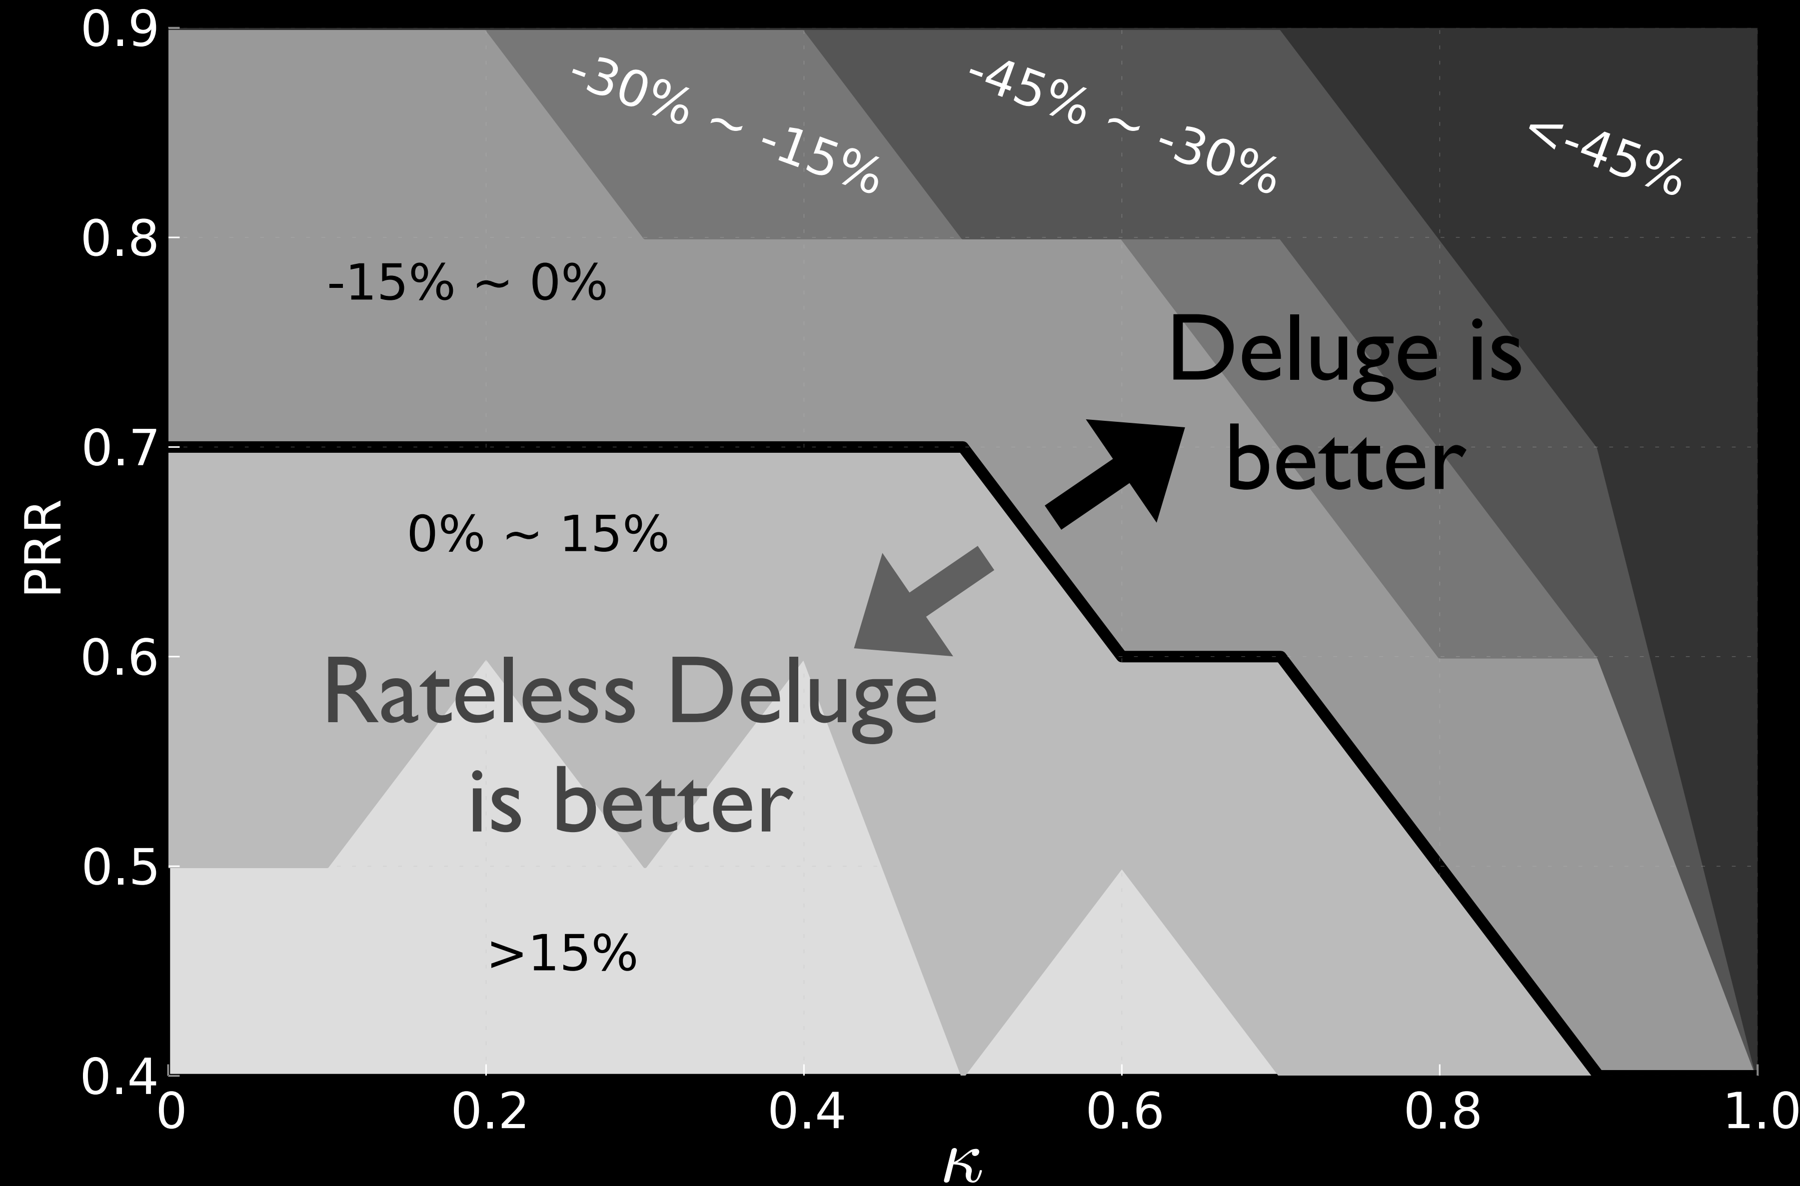

- Shows how much faster Rateless Deluge is over Deluge

Uncontrolled Experiment

- Measure **K** and then run the experiment
- 1 transmitter (injection point) and 8 receivers
- 3 setups
  - Ch 16: high correlation
  - Ch 26: medium correlation
  - Movement: low correlation

Deluge vs Rateless Deluge

Scenario	Avg κ	Avg PRR
Ch16 (High)	0.85	0.85

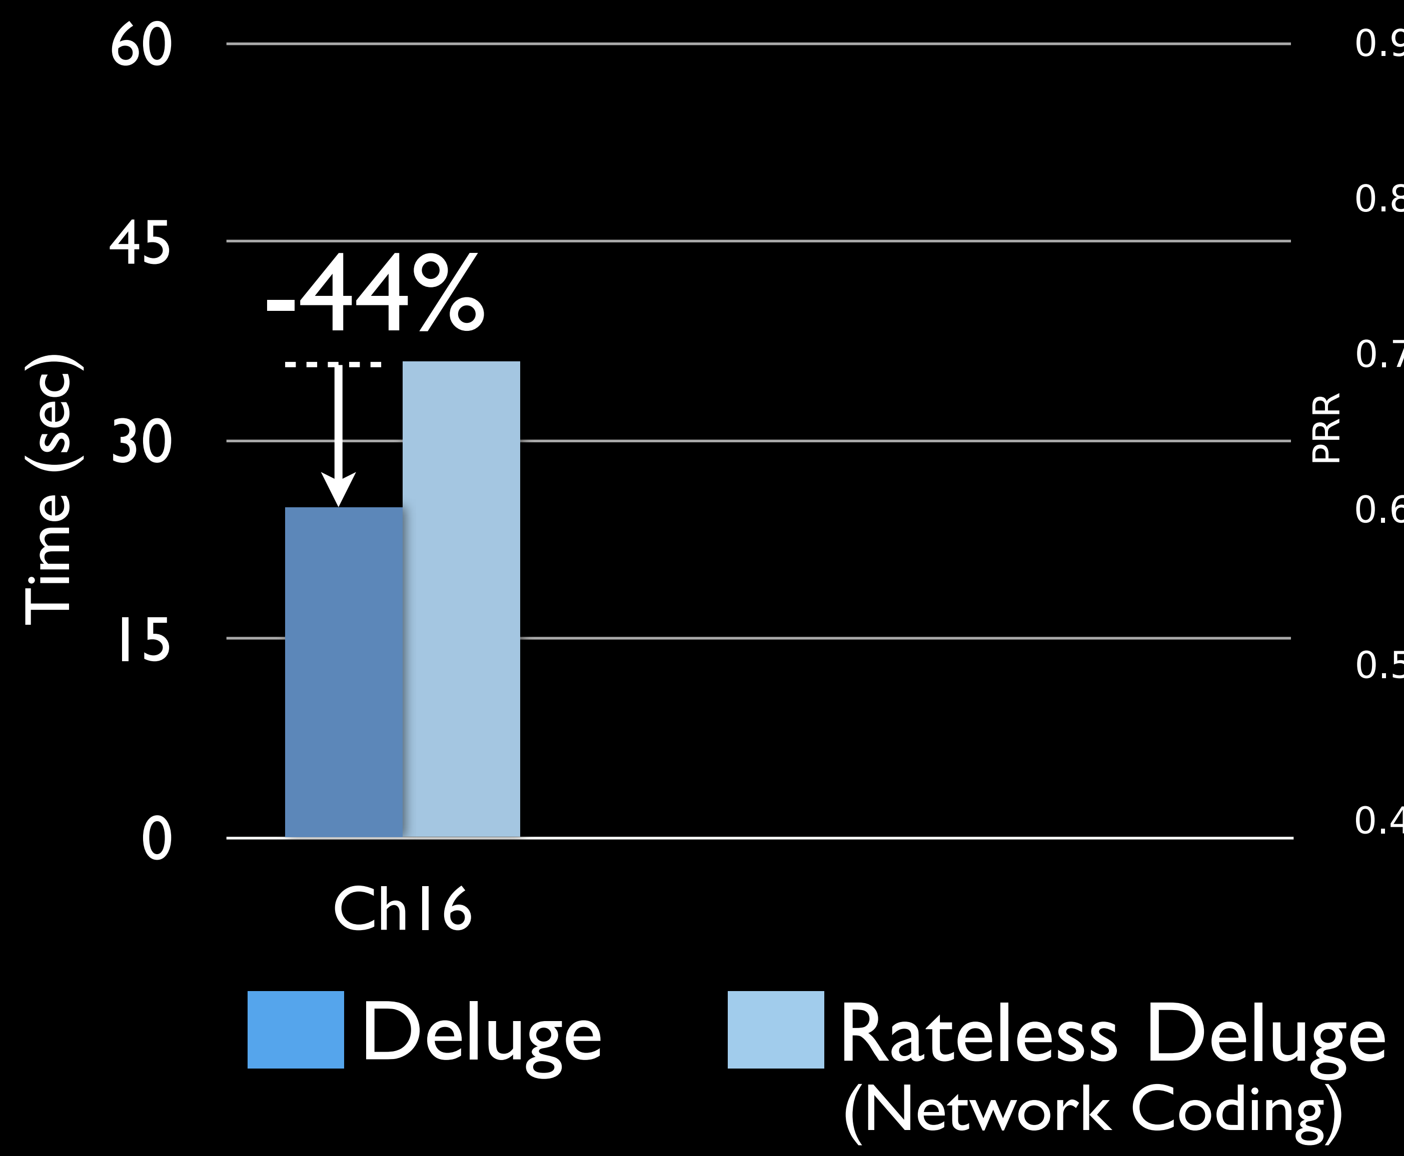

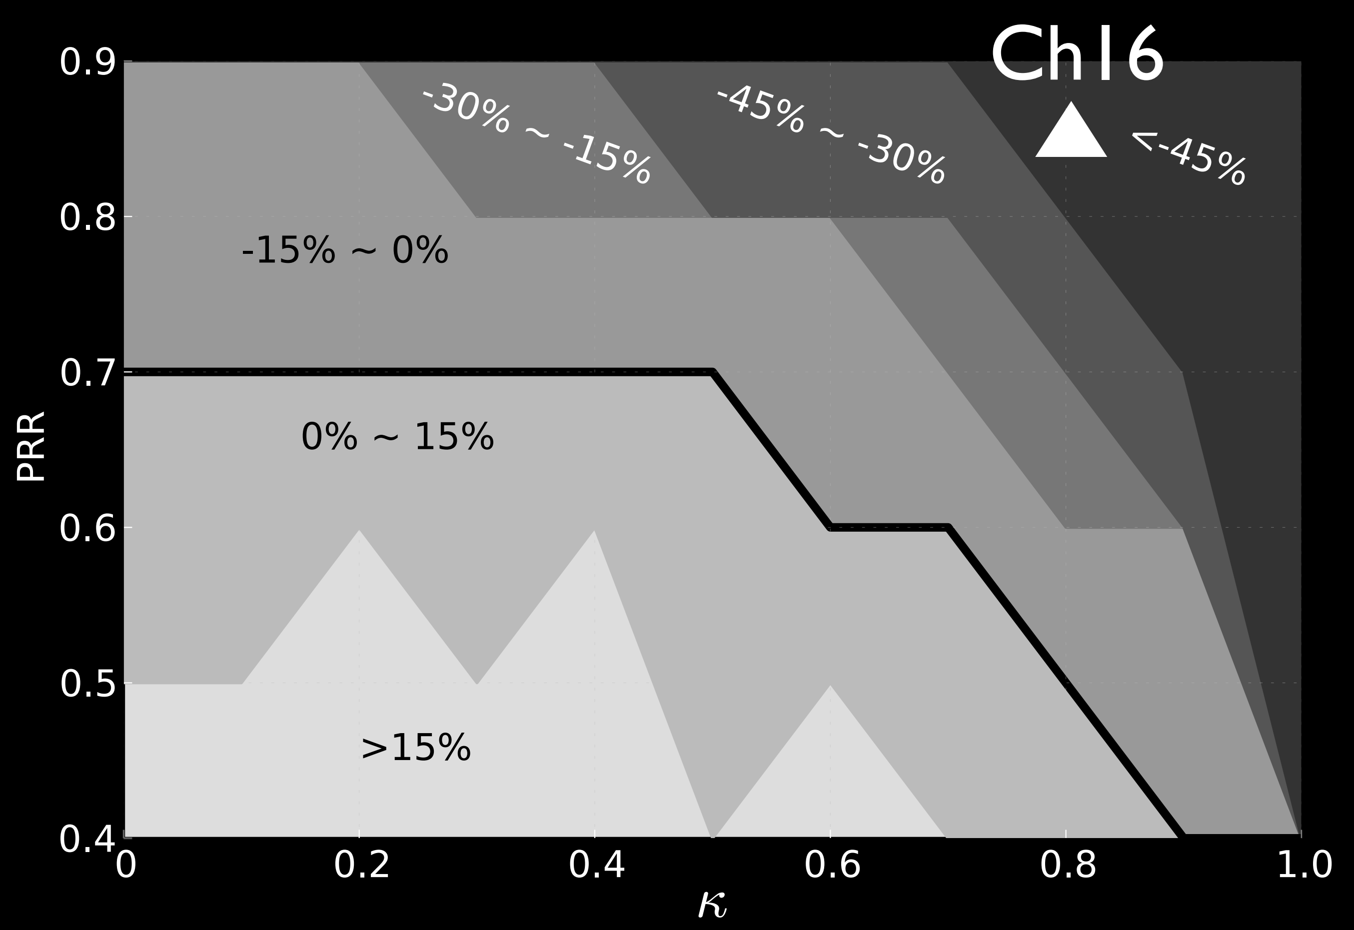

Deluge vs Rateless Deluge

Scenario	Avg κ	Avg PRR
Ch16 (High)	0.85	0.85
Ch26 (Medium)	0.55	0.91

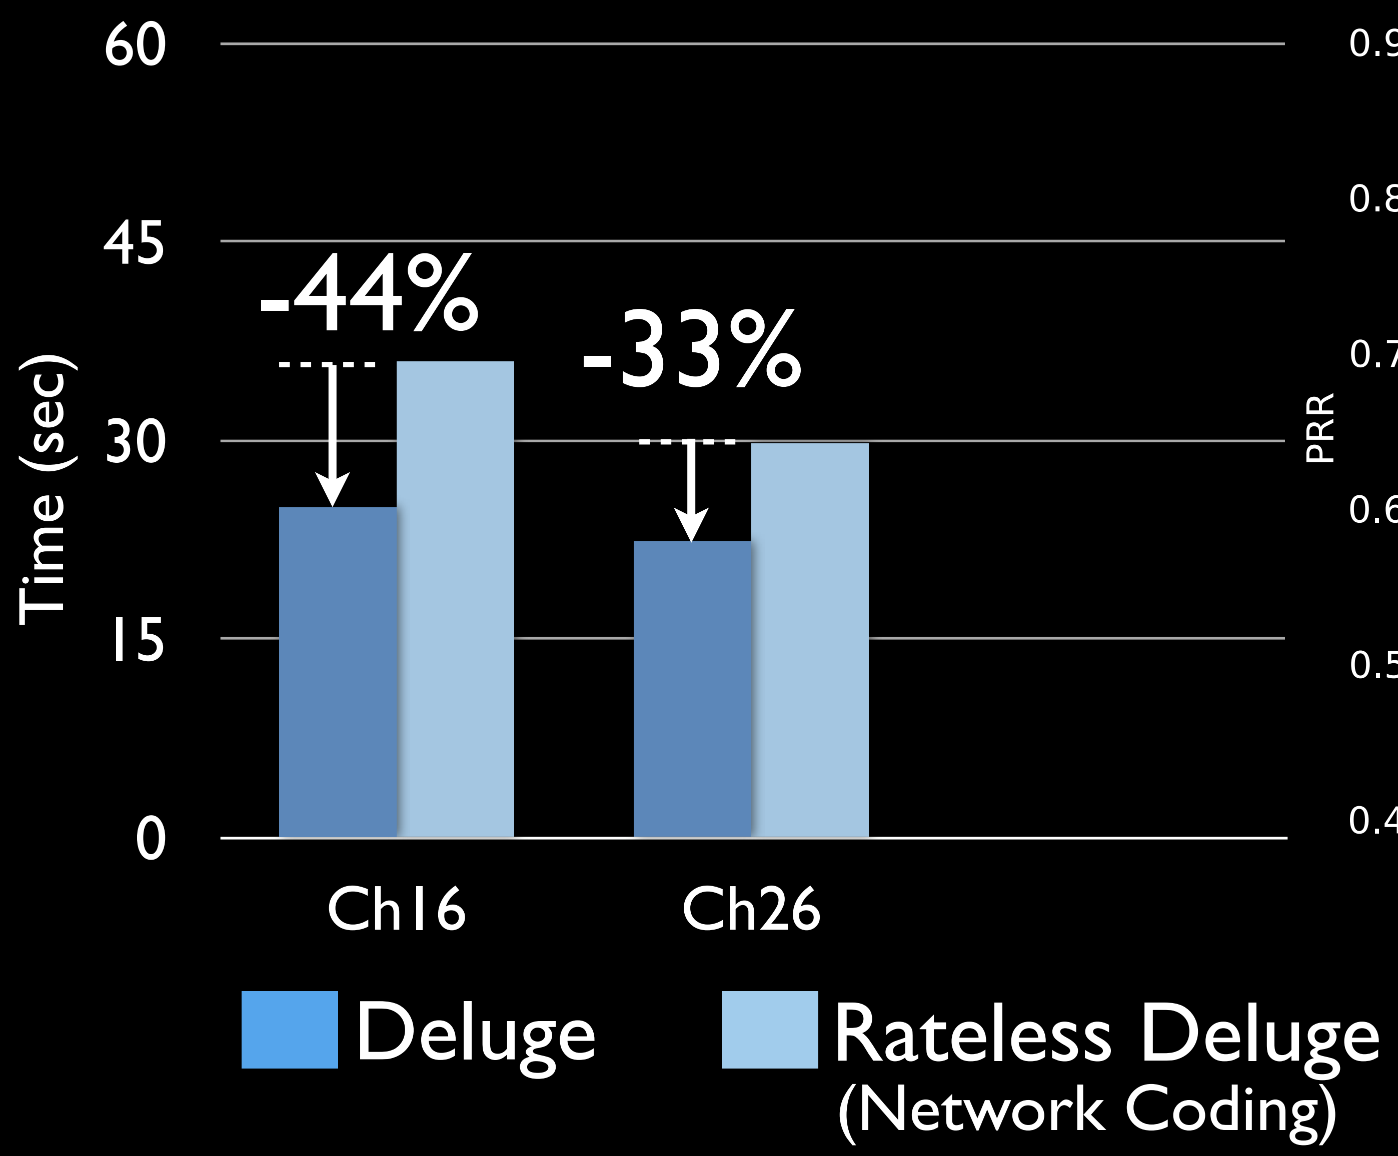

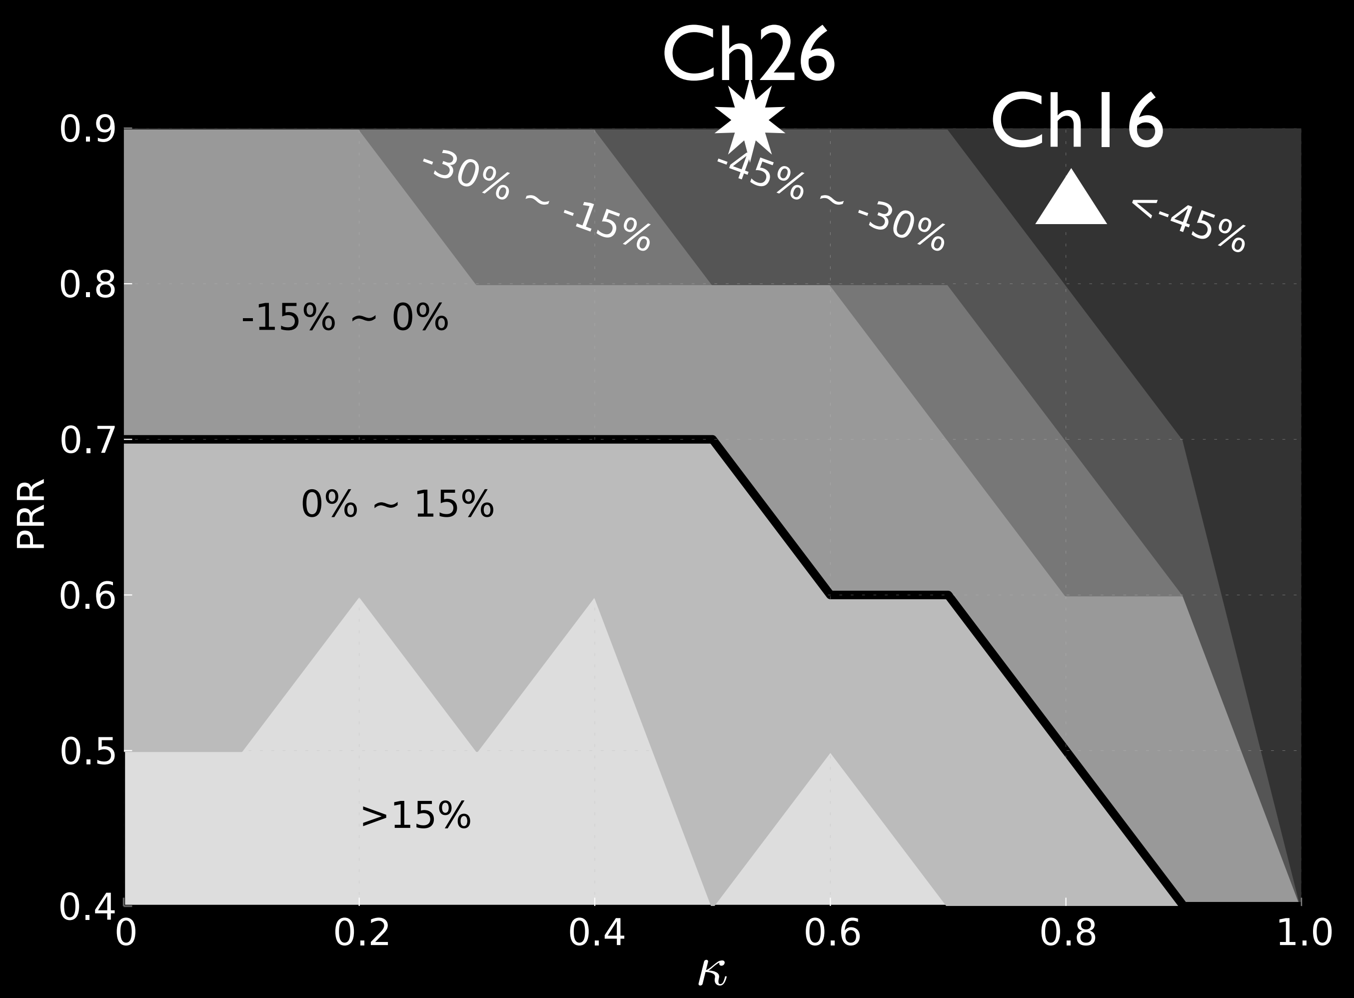

Deluge vs Rateless Deluge

Scenario	Avg κ	Avg PRR
Ch16 (High)	0.85	0.85
Ch26 (Medium)	0.55	0.91
Movement (Low)	0.04	0.57

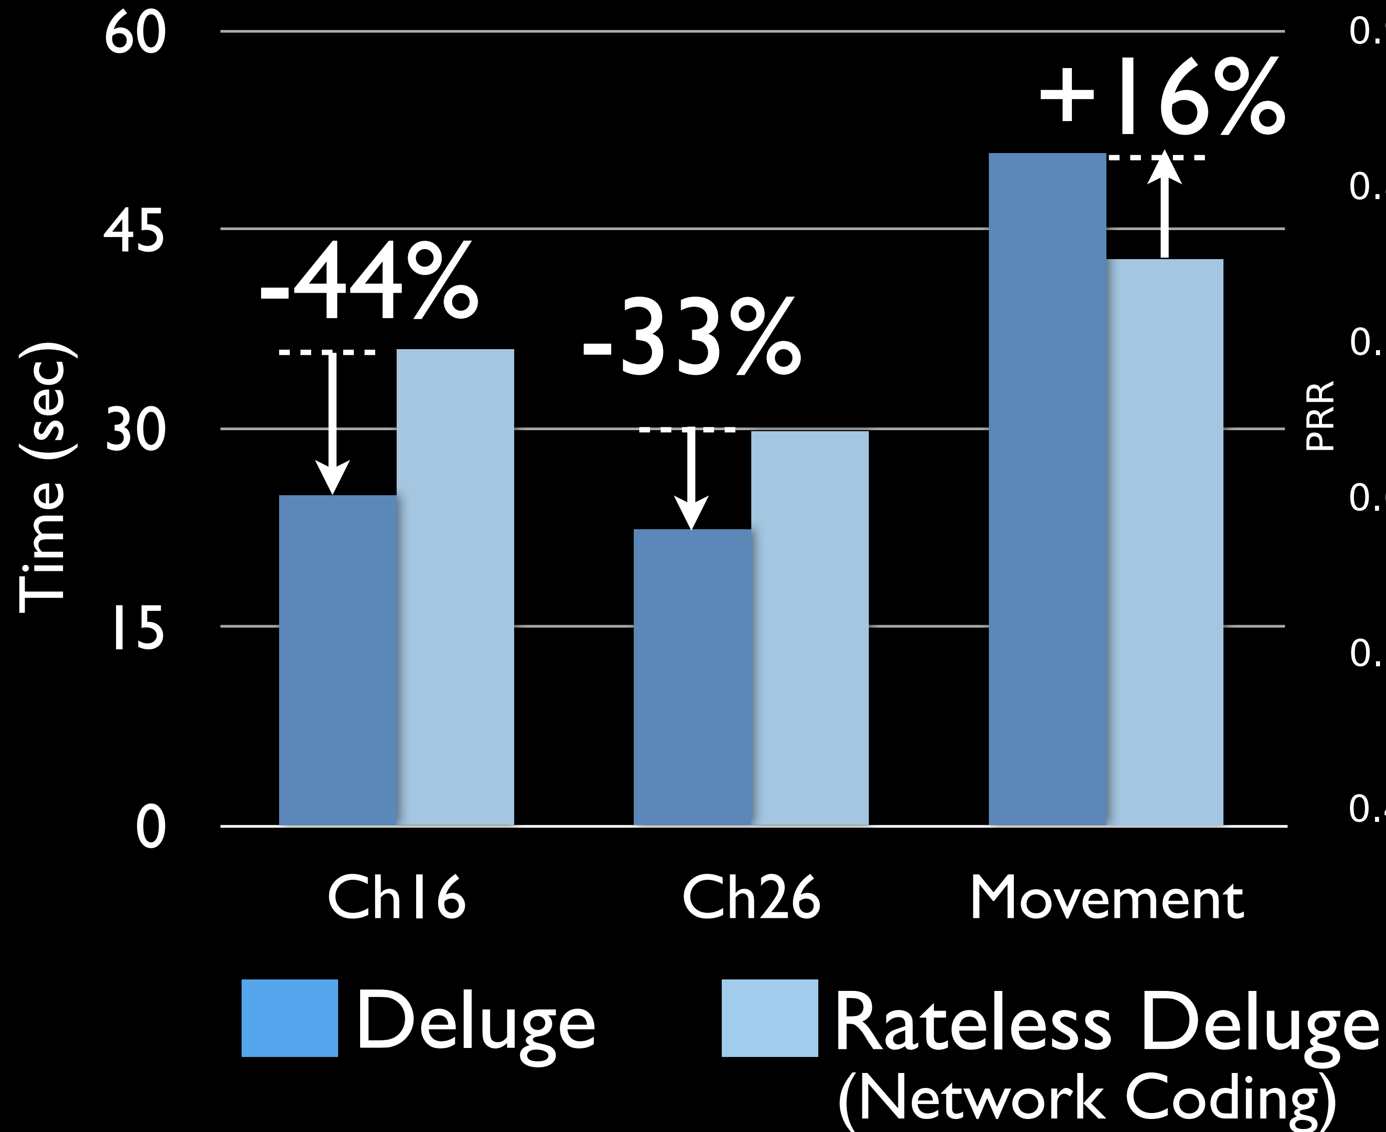

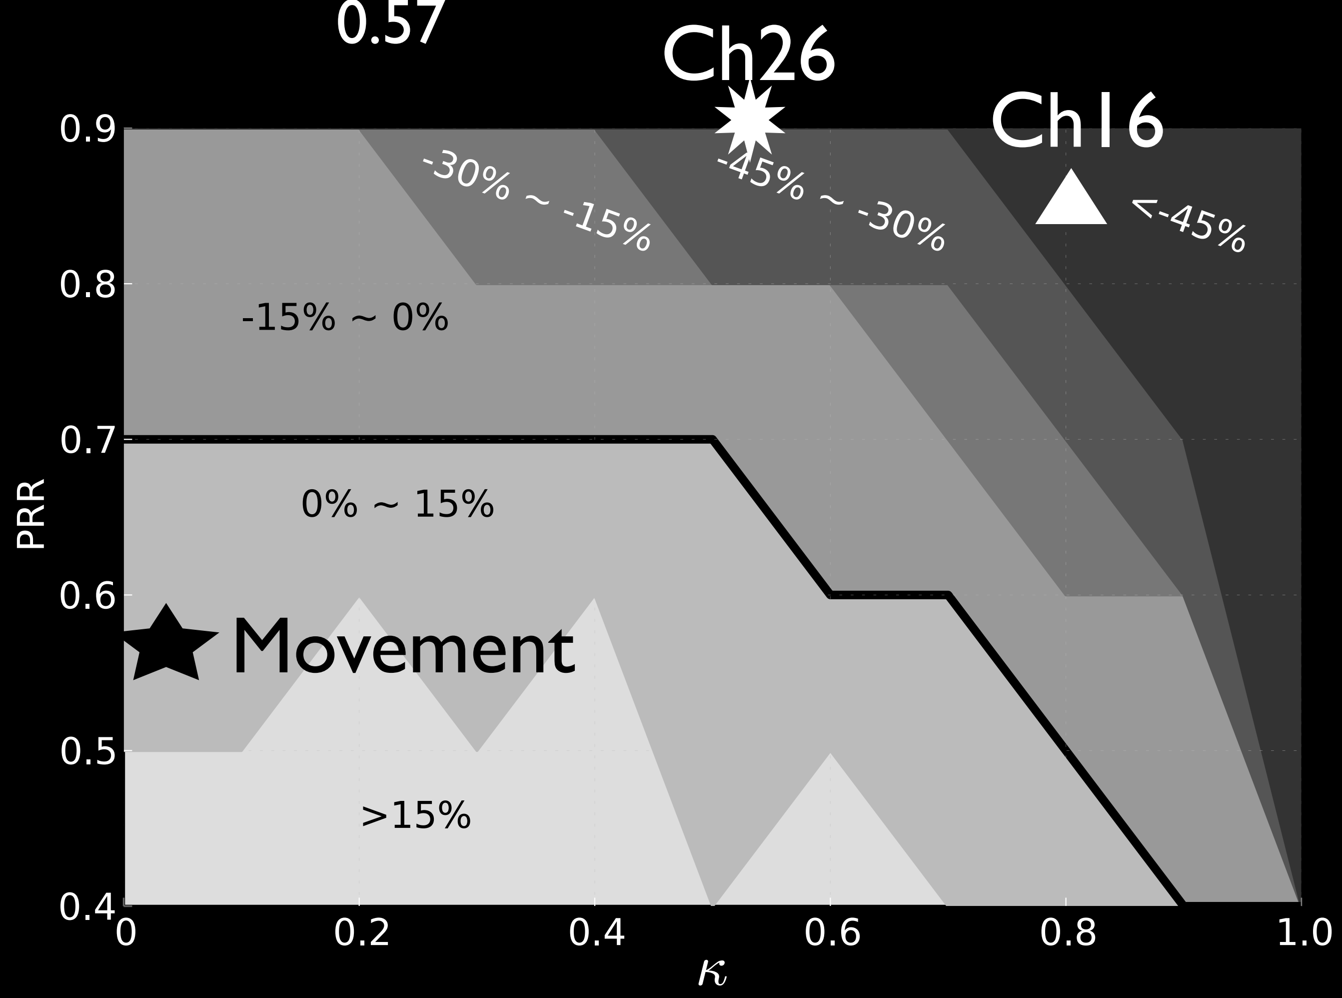

Outline

- Desired Metric Properties
- The κ Metric
- κ 's Usefulness
- Open Questions

Open Questions

κ can change over time

- how to measure it online?
  - useful for adaptive protocol design

Is κ useful with adaptive protocols?

- adaptive rate
- adaptive packet size
- adaptive channel and bandwidth

Summary

- Presented a spatial correlation metric, κ
  - κ does not conflate correlation with PRRs
- κ has great predictive qualities
  - predicts network coding protocol performance
  - κ shows how well opportunistic routing protocols perform

A Shameless Advertisement

- I'm looking for a faculty/research position

  contact: srikank@stanford.edu

- Mayank and Jung Il are looking for industrial research positions

  contact for Mayank: mayjain@stanford.edu
  contact for Jung Il: jungilchoi@stanford.edu

Backup Slides

κ on 802.11 networks

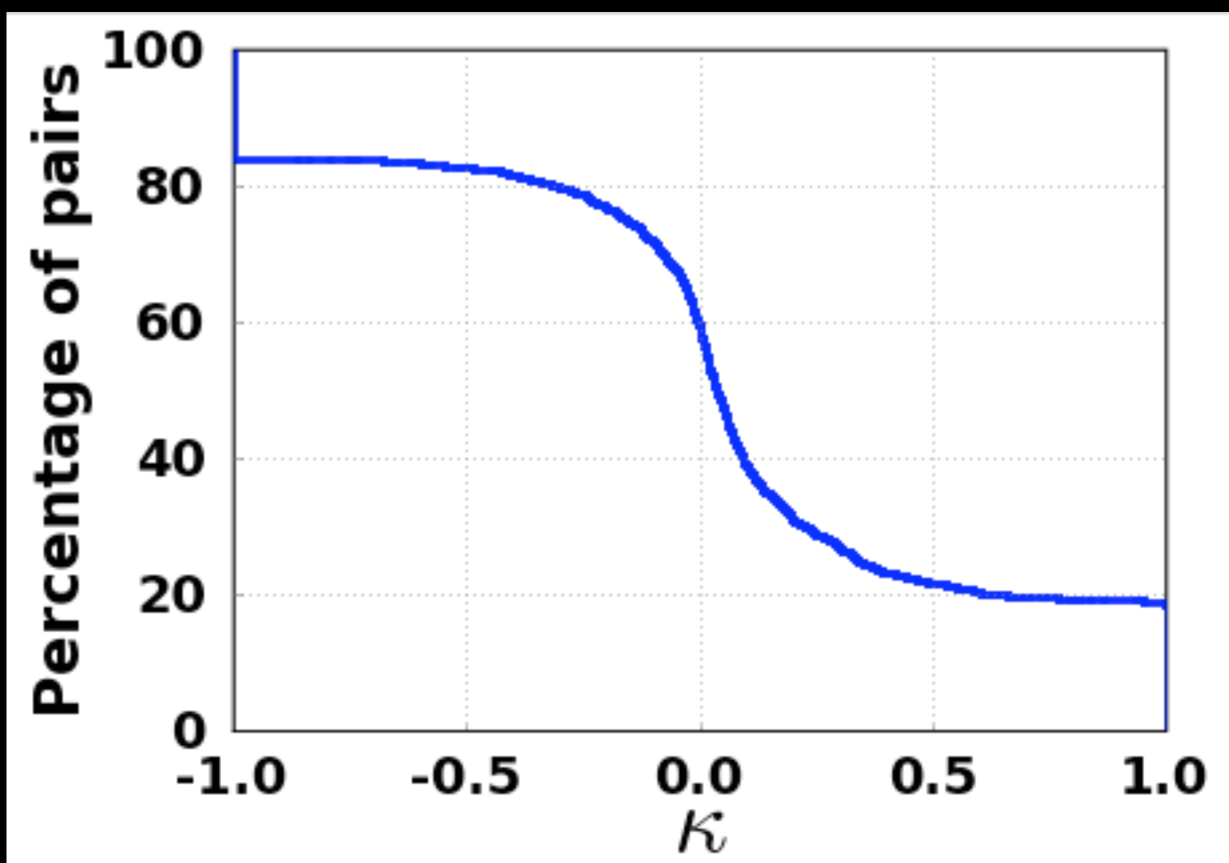

Roofnet (Outdoor)

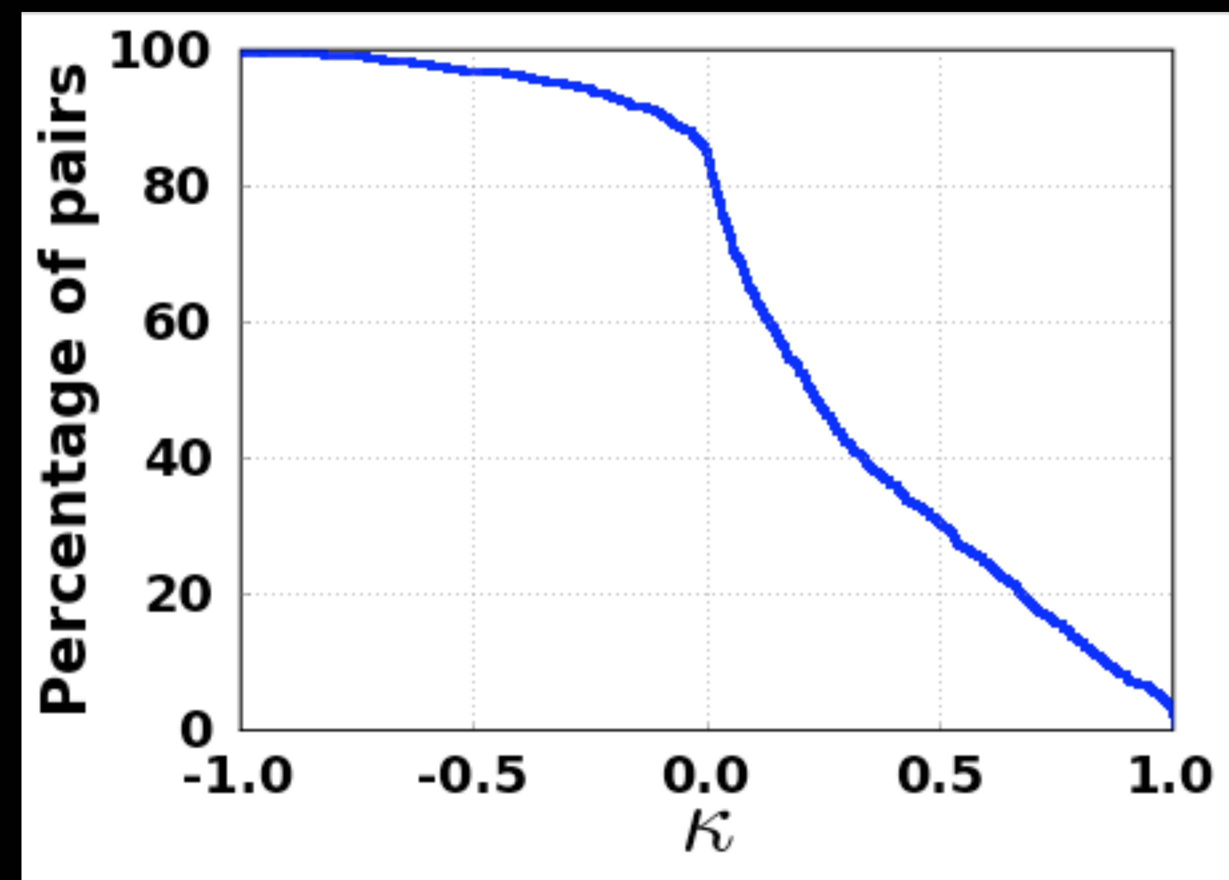

SWAN (Indoor)

Anypath ETX Ratio

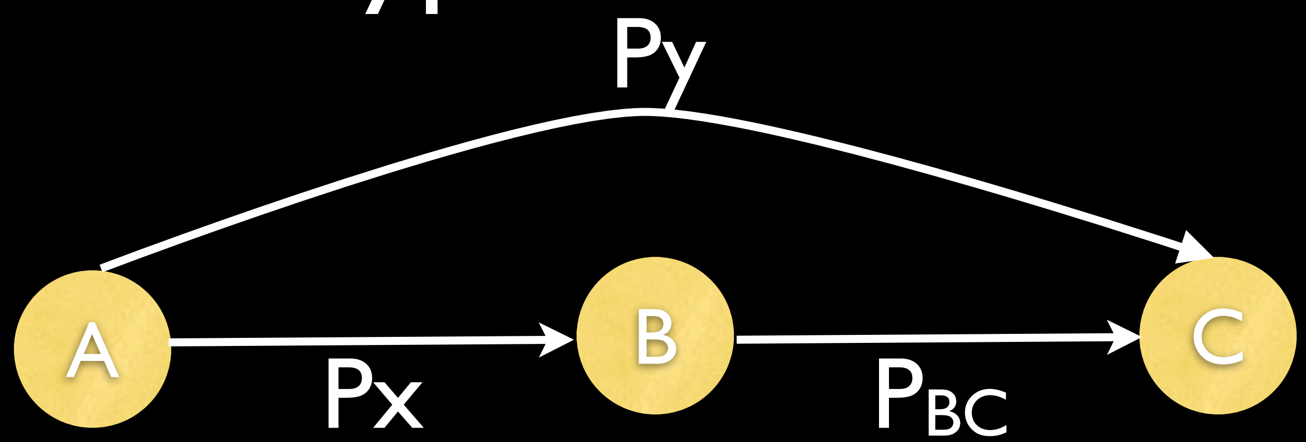

Anypath ETX Ratio =

$$\begin{cases} \frac{E[A]-E[A]_{inde}}{E[A]_{max}-E[A]_{inde}}, & E[A] \geq E[A]_{inde} \\ \frac{E[A]-E[A]_{inde}}{E[A]_{inde}-E[A]_{min}}, & \text{otherwise.} \end{cases}$$

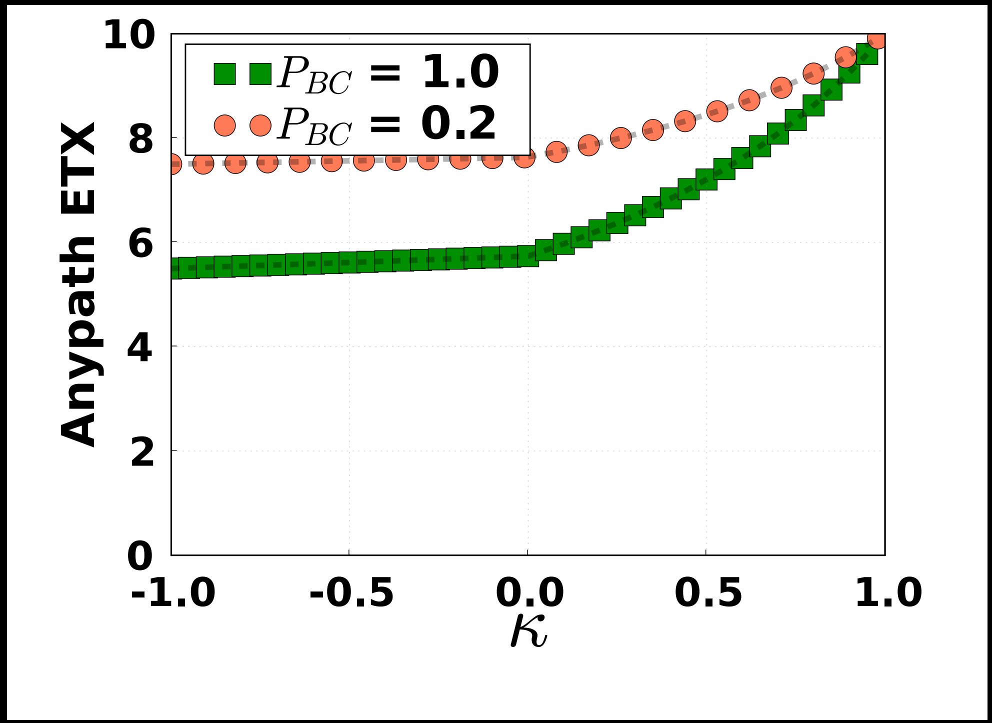

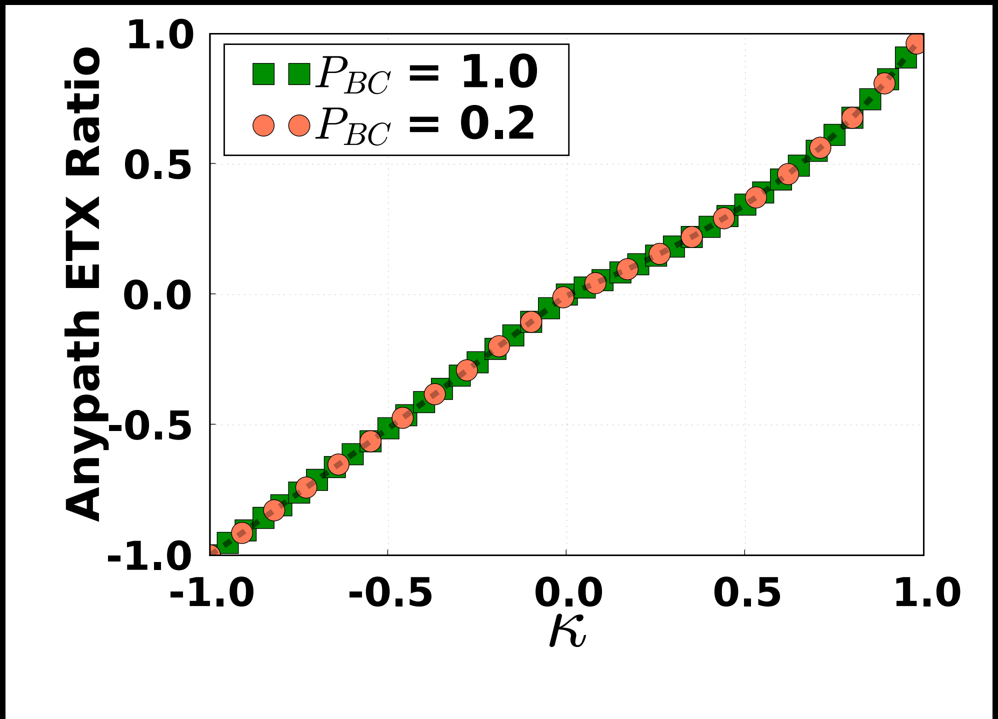

Opportunistic Routing: χ and κ

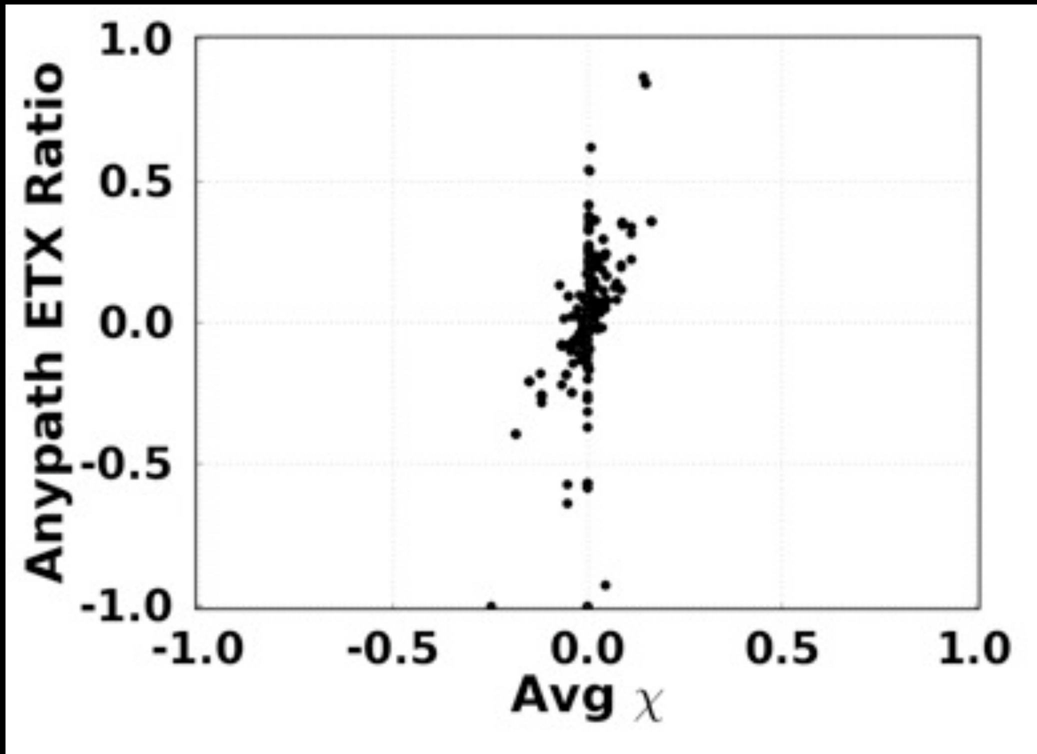

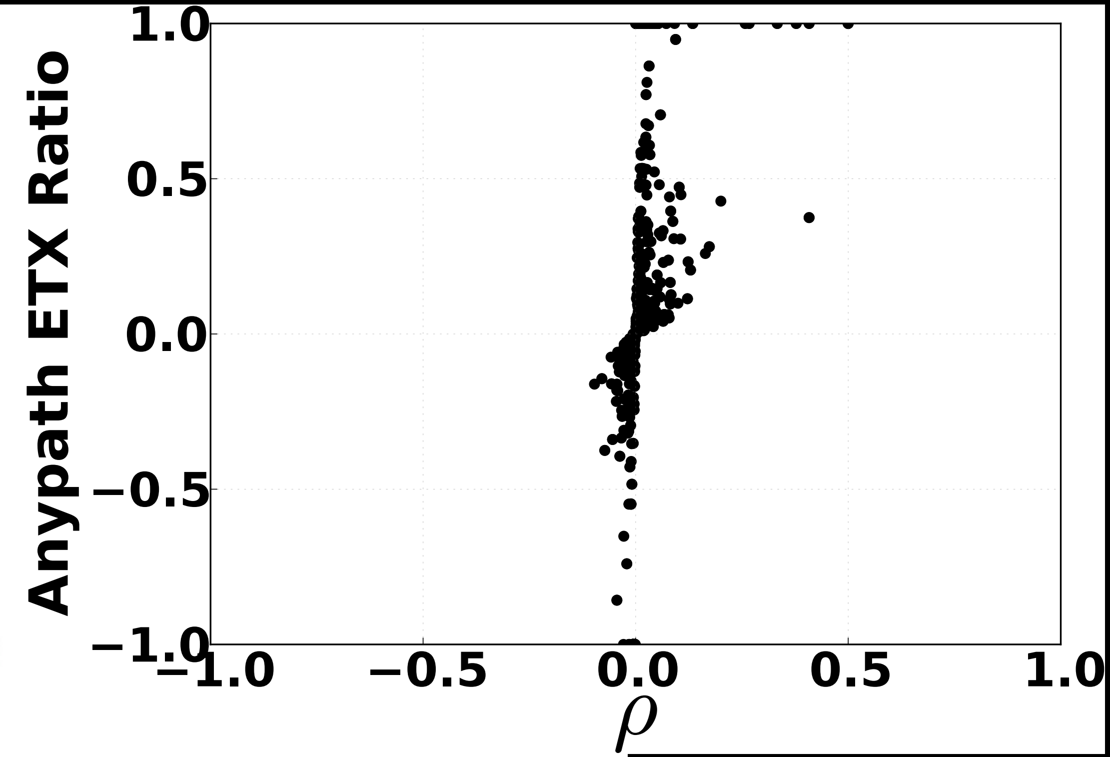

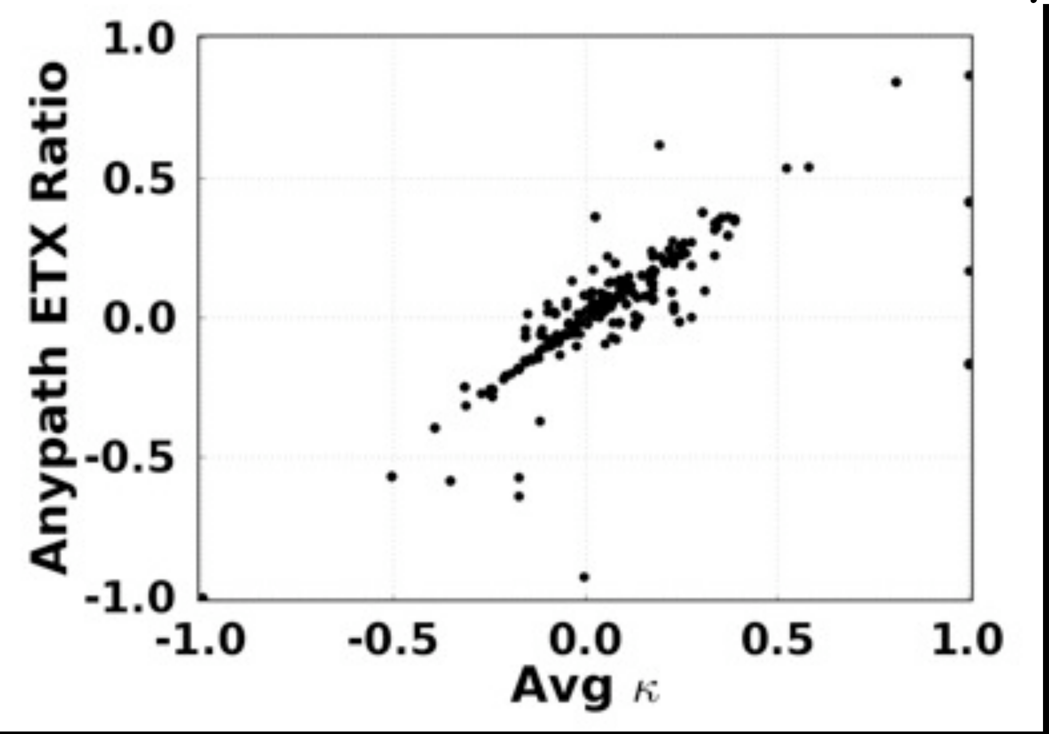

Roofnet 11Mbps

Causes

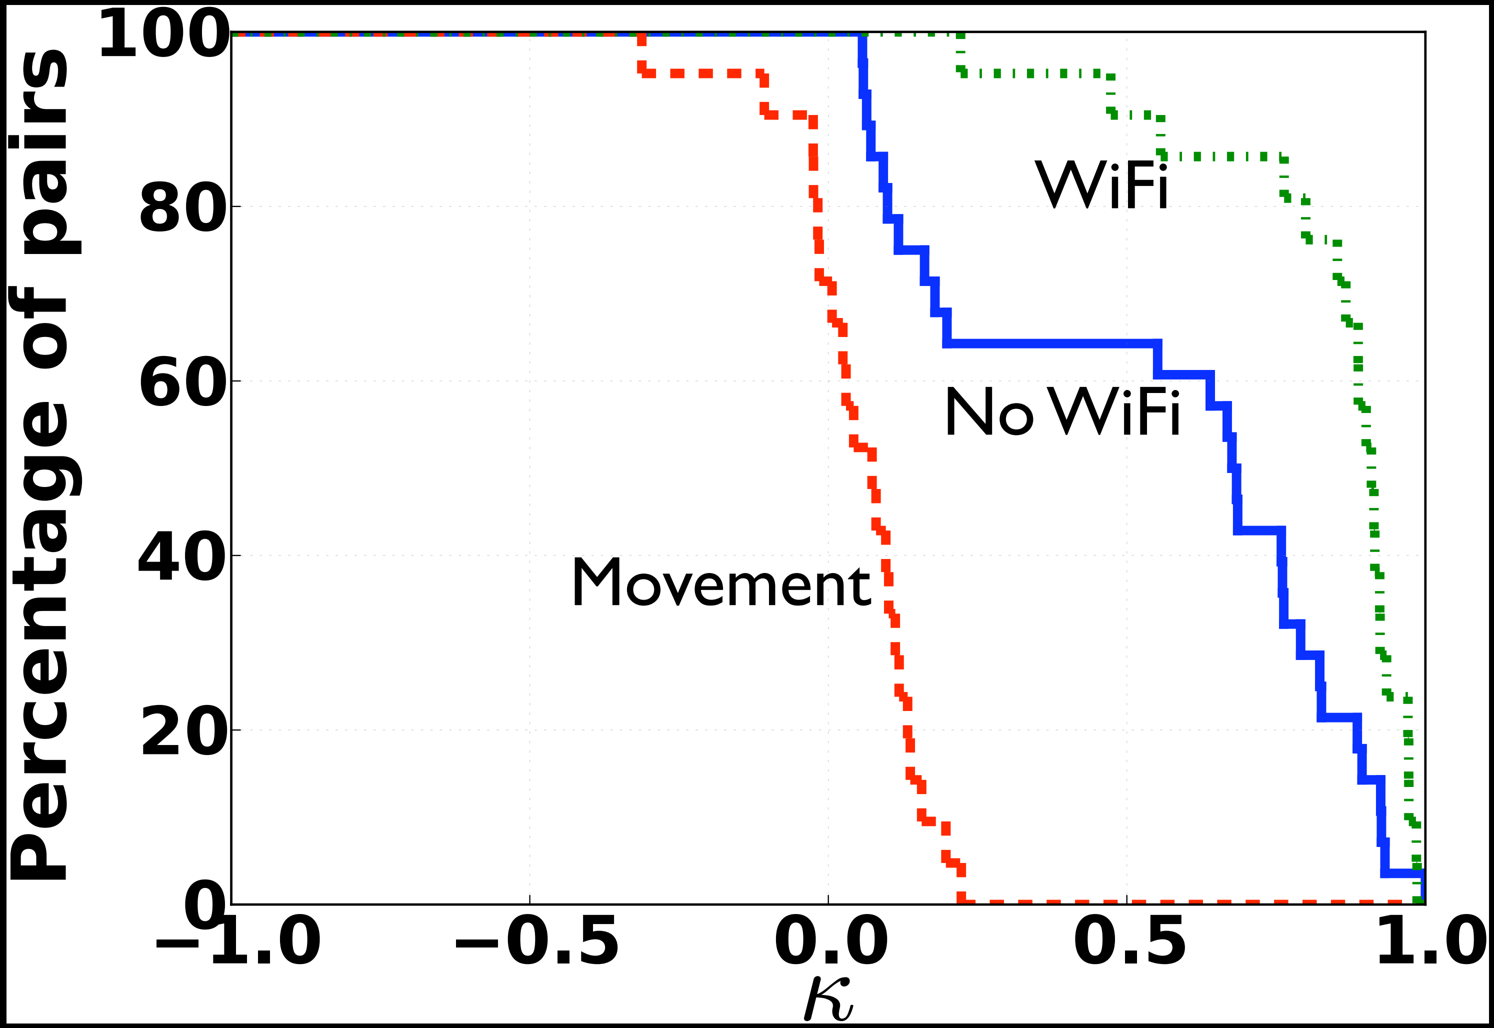

Percentage of pairs
100
80
60
40
20
0
−1.0
−0.5
0.0
0.5
1.0
κ
WiFi
No WiFi
Movement