The Mysteries of
Shannon's Channel and Capacity:
Then and Now

IIT Kanpur Shannon Centennial, October 2016

Michelle Effros

California Institute of Technology

TO UNDERSTAND SHANNON’S WORK,
IT IS USEFUL TO KNOW SOMETHING ABOUT
HIS TIME.

Claude Shannon

was born in Michigan in 1916
to Claude Elwood and Mabel Wolf Shannon.

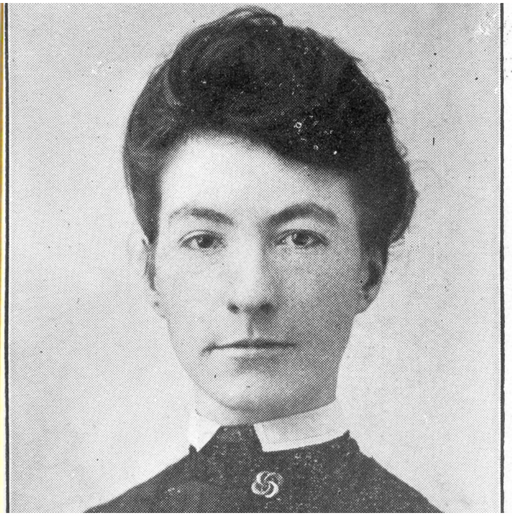

In 1916, telephony was new ...

1892: Bell placing the first New York to Chicago phone call

... but it was catching on quickly.

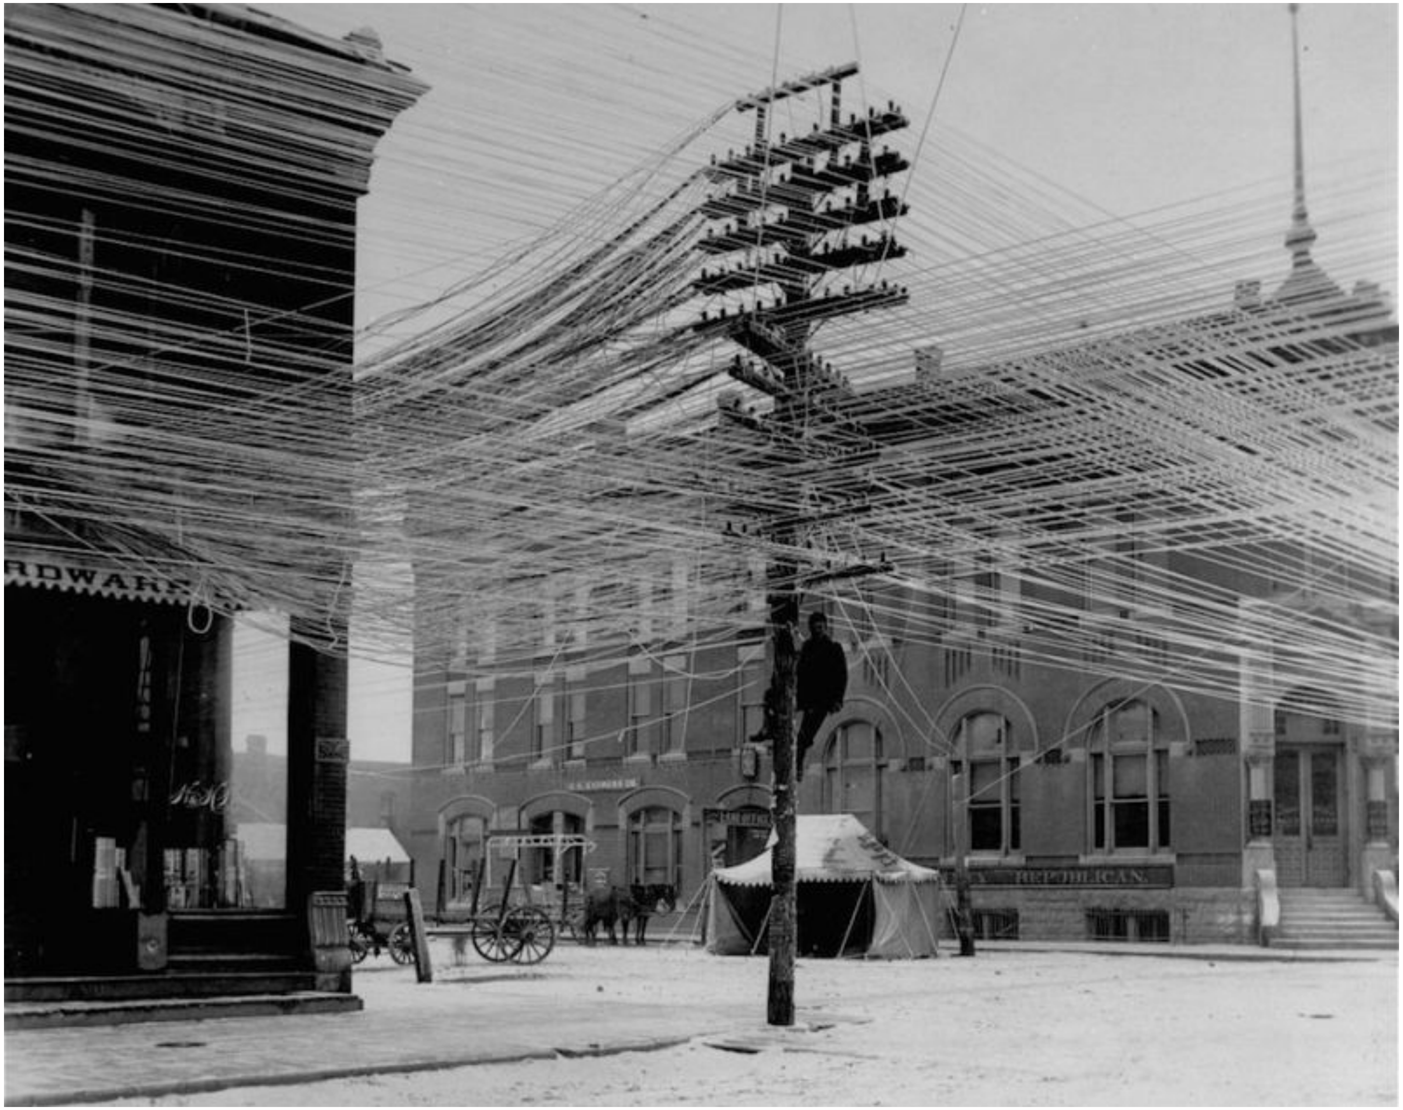

Pratt, Kansas 1911 (pop 11,156)

IN THE LATE 1930s, SHANNON BEGAN WORK ON A NEW THEORY OF "TRANSMISSION OF INTELLIGENCE."

He was interested in information representation ...

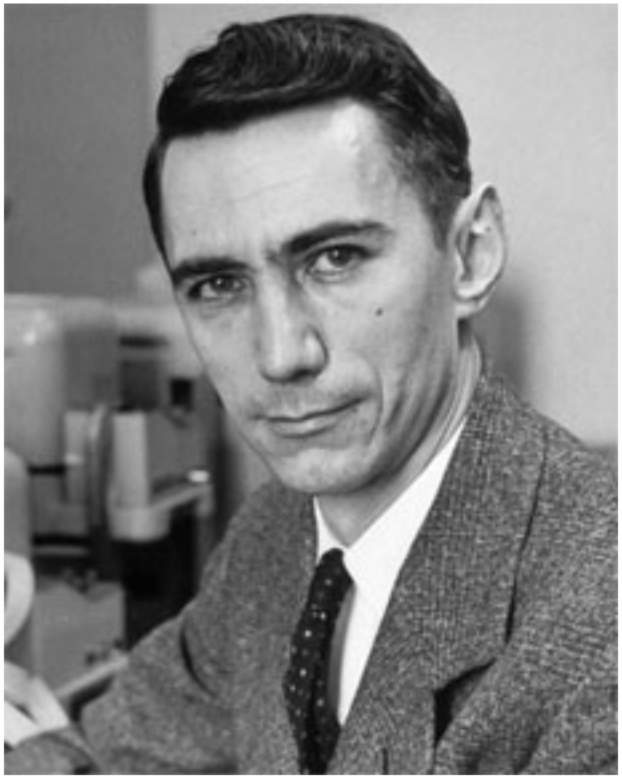

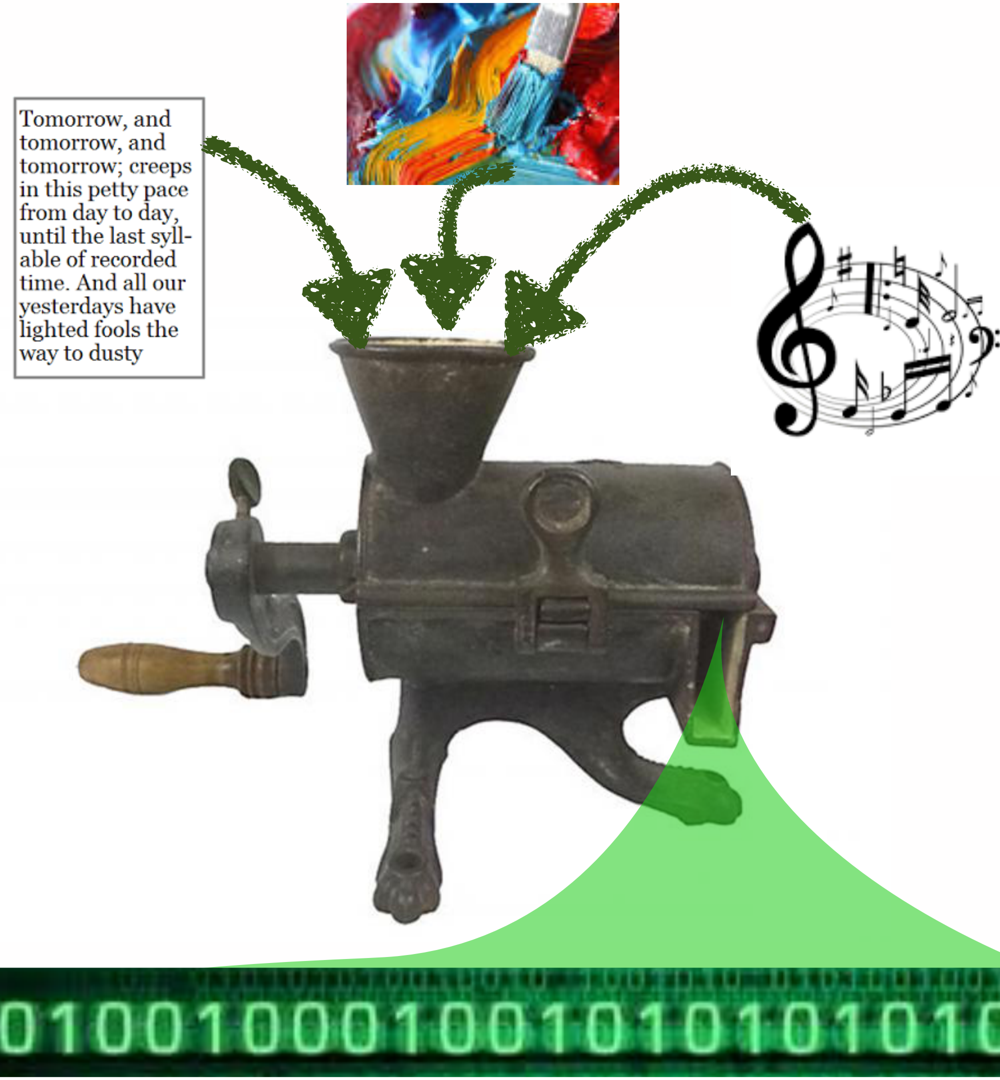

... as well as reliable communication.

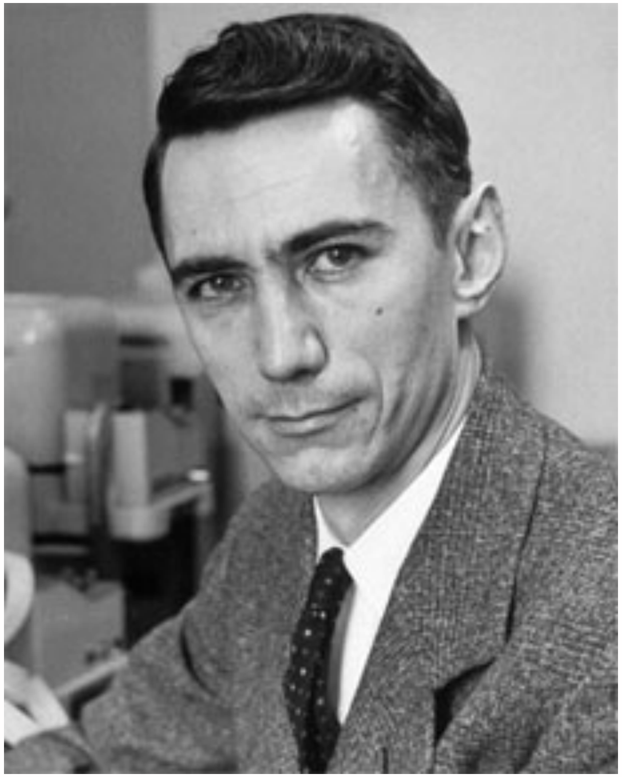

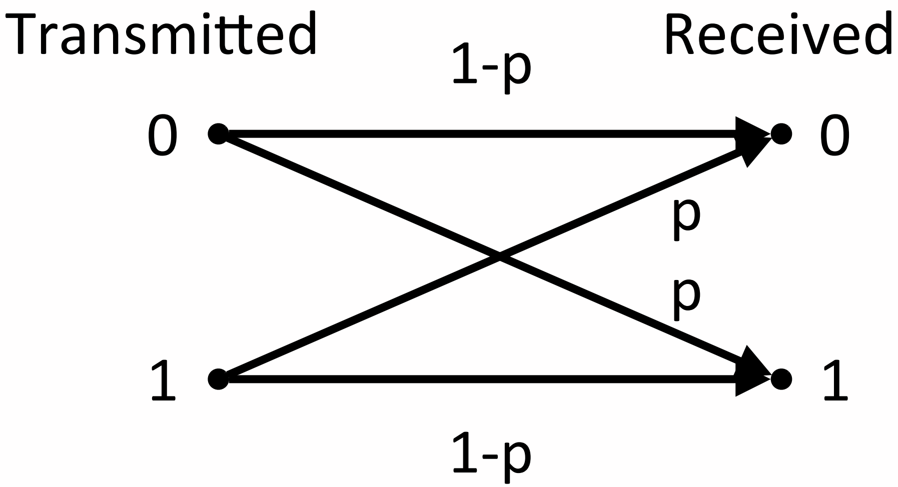

To make communication reliable, add redundancy.

Message	0	0	1	0	1
Transmitted	000	000	111	000	111
Received	010	000	110	100	010
Decoding	0	0	1	0	0

Rate = 1/3

But the more you repeat, the less you can say.

Message	0	0	1	0	1
Transmitted	000000000	000000000	111111111	000000000	111111111
Received	010000010	000010000	110110011	100010001	010110111
Decoding	0	0	1	0	1

Rate = 1/9

Can we increase reliability without forcing rate to 0?

Perhaps we can be more clever than repetition ...

Message	0	0	1	0			
Transmitted	0	0	1	0	1	1	1

Can we increase reliability without forcing rate to 0?

Perhaps we can be more clever than repetition ...

Message	0	0	1	0			
Transmitted	0	0	1	0	1	1	1

Can we increase reliability without forcing rate to 0?

Perhaps we can be more clever than repetition ...

Message	0	0	1	0			
Transmitted	0	0	1	0	1	1	1

Can we increase reliability without forcing rate to 0?

Perhaps we can be more clever than repetition ...

Message	0	0	1	0			
Transmitted	0	0	1	0	1	1	1

Can we increase reliability without forcing rate to 0?

Shannon showed that increasing reliability does not necessarily force rate to zero.

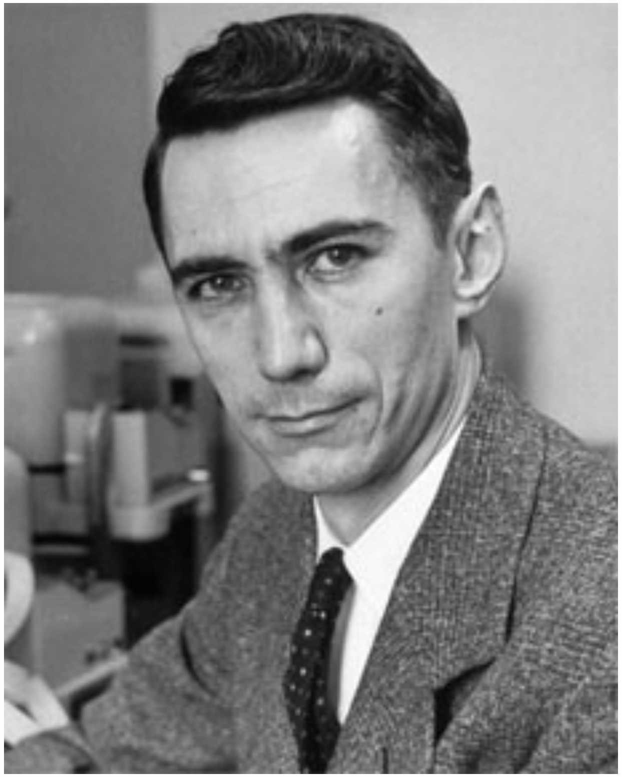

For each channel there is a range of rates achievable with arbitrary reliability.

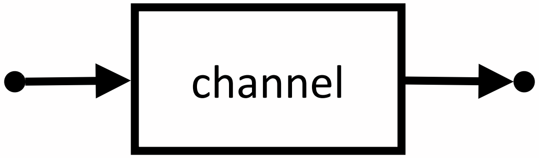

$$C = \max_{p(x)} I(X;Y)$$

The maximal such rate is called the capacity.

The capacity of a point-to-point channel is the number of bits per channel use that the link can reliably deliver.

SHANNON'S CHANNEL CAPACITY IS SOLVED.

MANY MORE MYSTERIES REMAIN.

Shannon's channel model

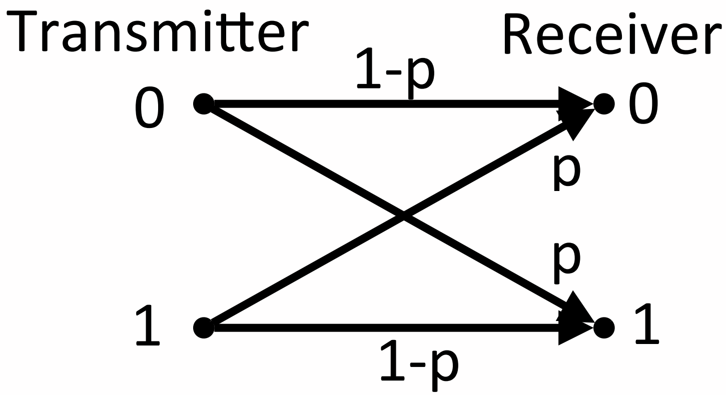

captures a world that looks like this.

Shannon’s channel model

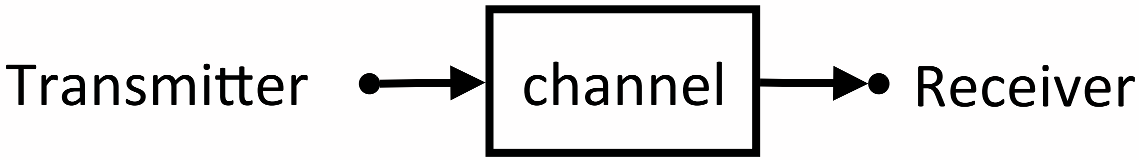

captures a world that looks like this.

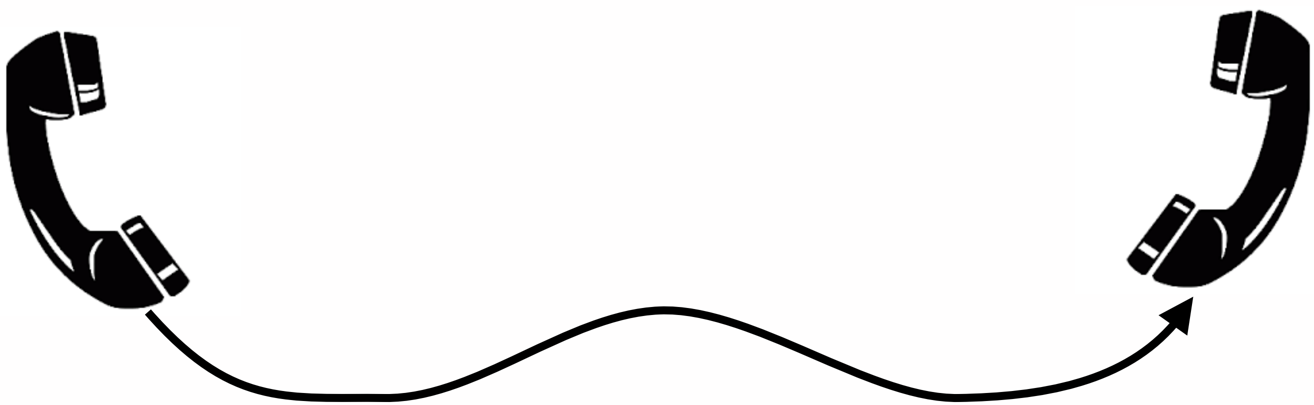

But even then, the network was far more complex.

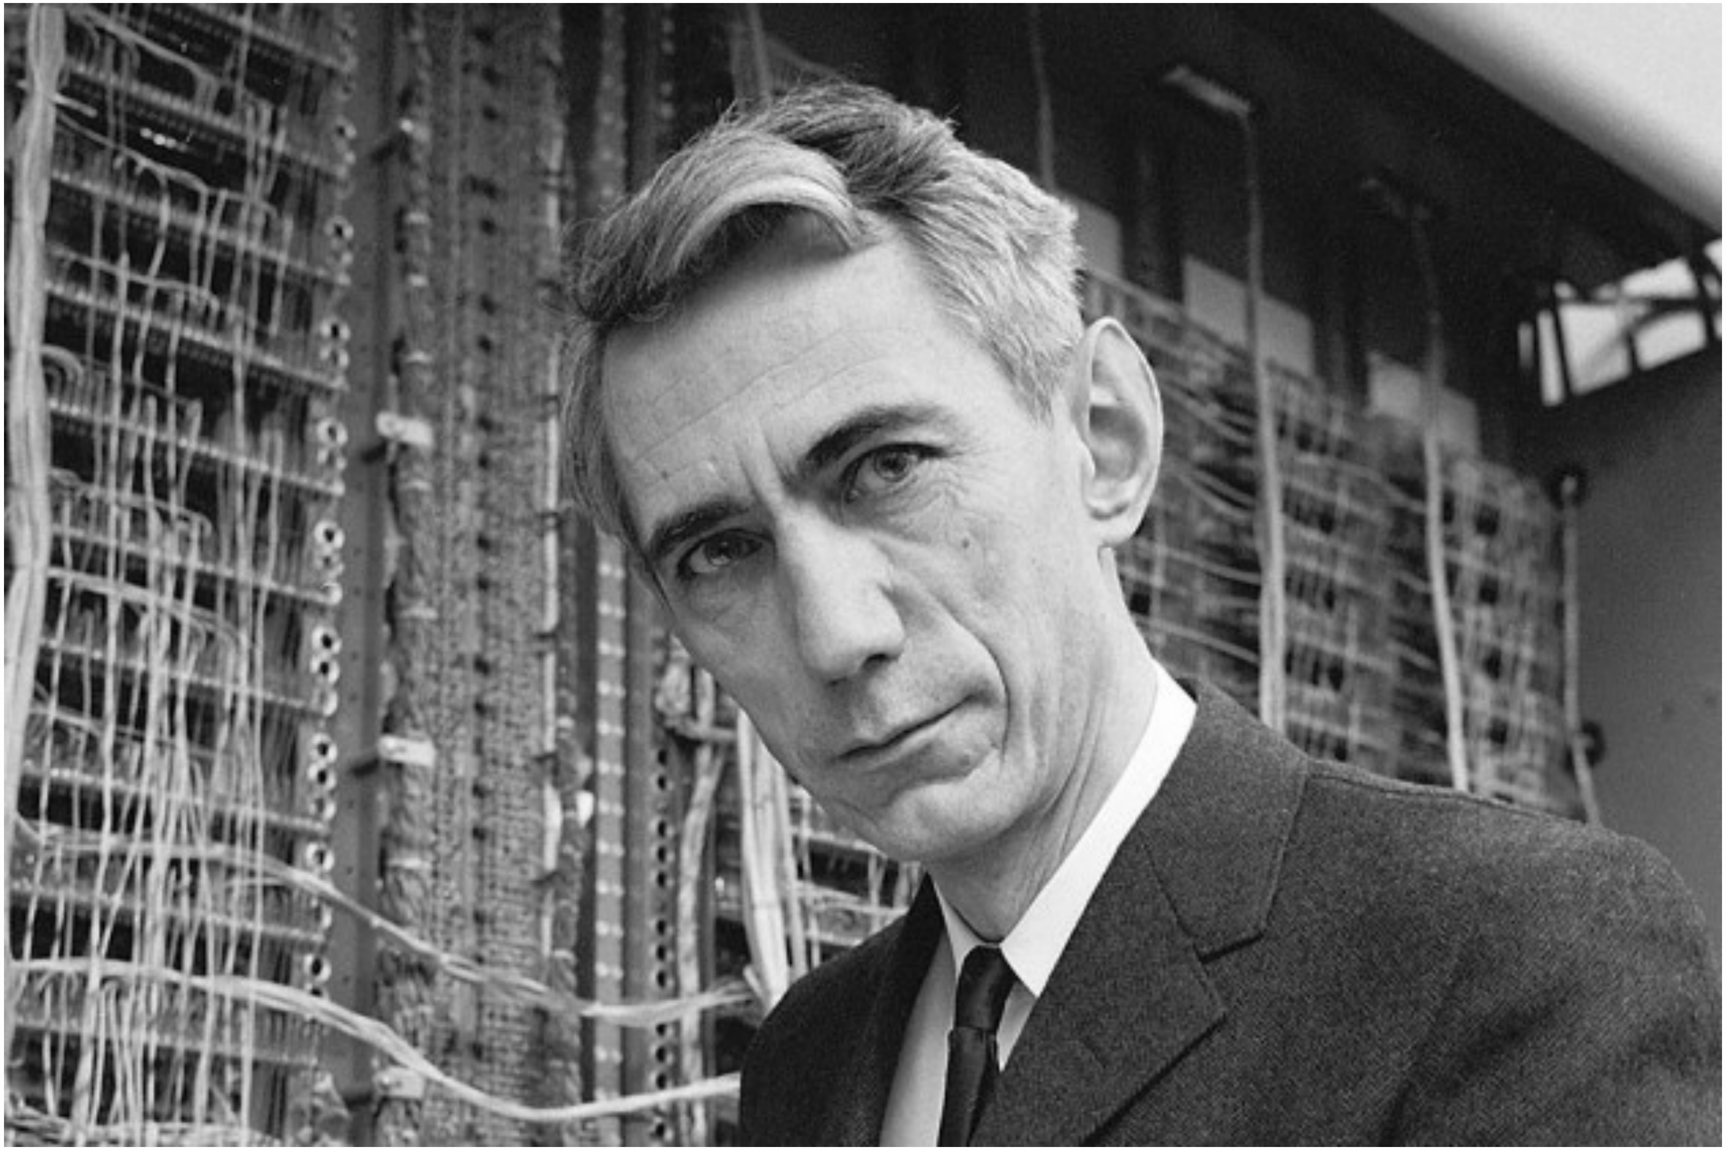

IS SHANNON'S CHANNEL'S CAPACITY RELEVANT TO THE NETWORK'S CAPACITY?

The capacity of a network is the set of rate vectors at which all source & receiver pairs can be simultaneously satisfied.

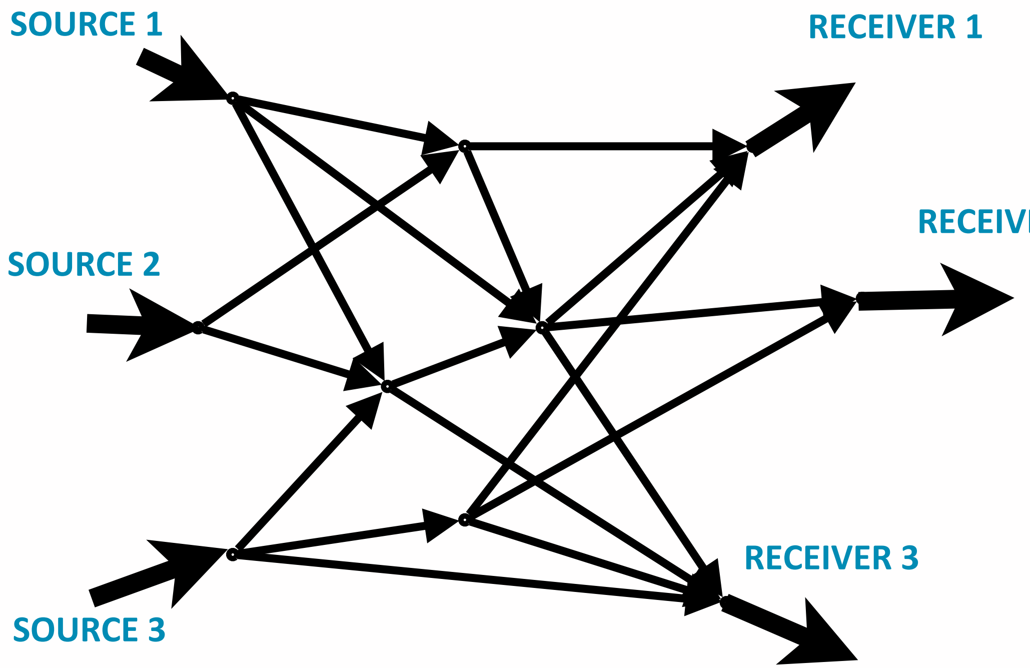

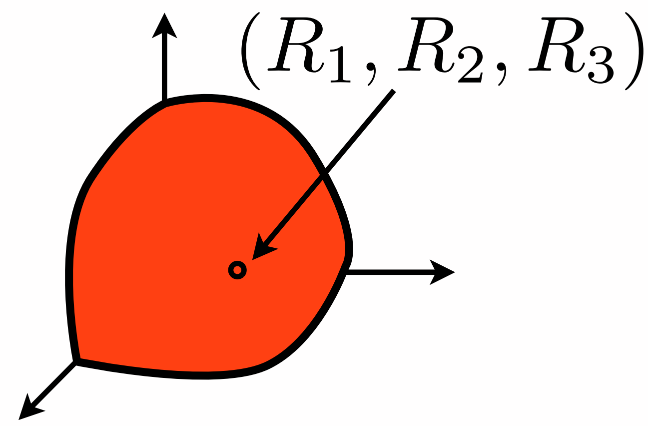

Does the capacity capture the essence of a channel's behavior?

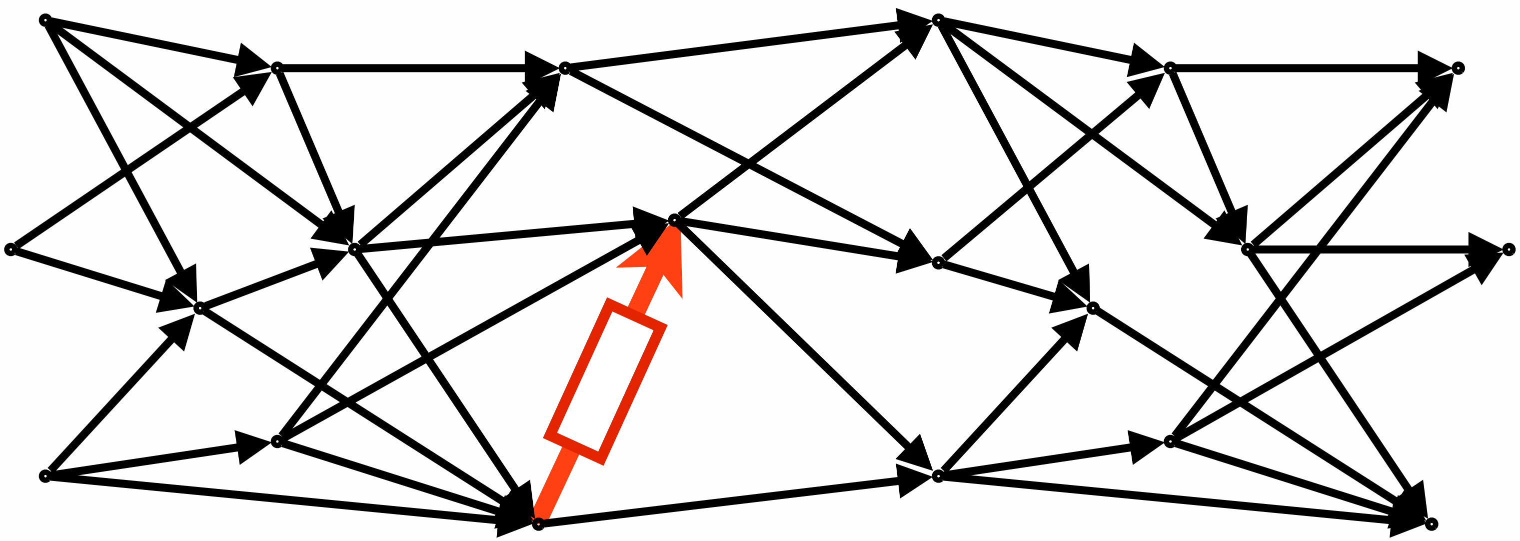

Does the capacity capture the essence of a channel's behavior?

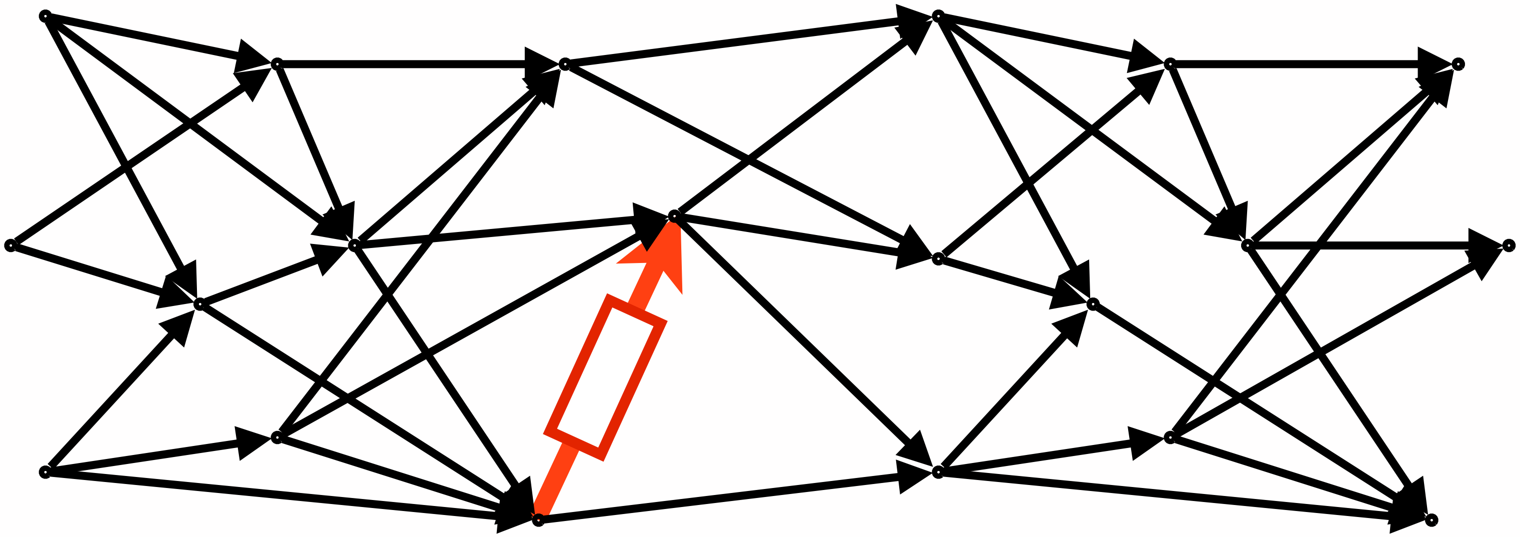

A noisy channel has the same impact on network capacity as a lossless link of the same capacity.

[Koetter, Effros, Medard 2009, 2011]

Does the capacity capture the essence of a channel's behavior?

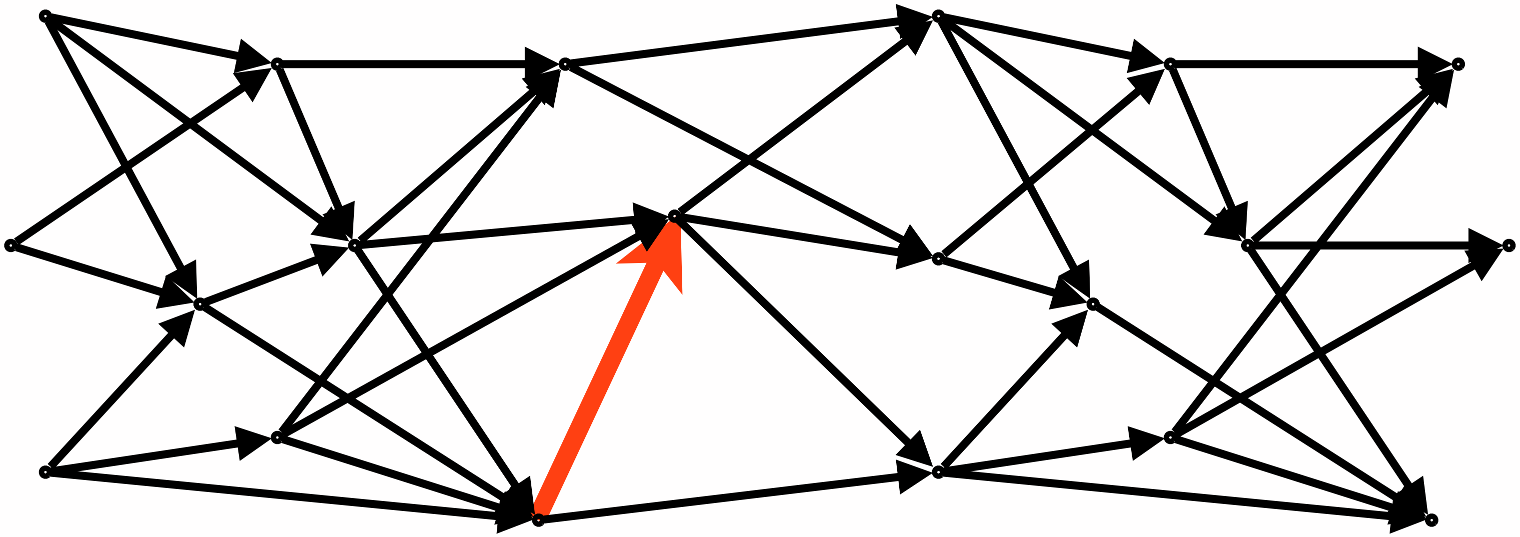

A noisy channel has the same impact on network capacity as a lossless link of the same capacity.

[Koetter, Effros, Medard 2009, 2011]

Does the capacity capture the essence of a channel's behavior?

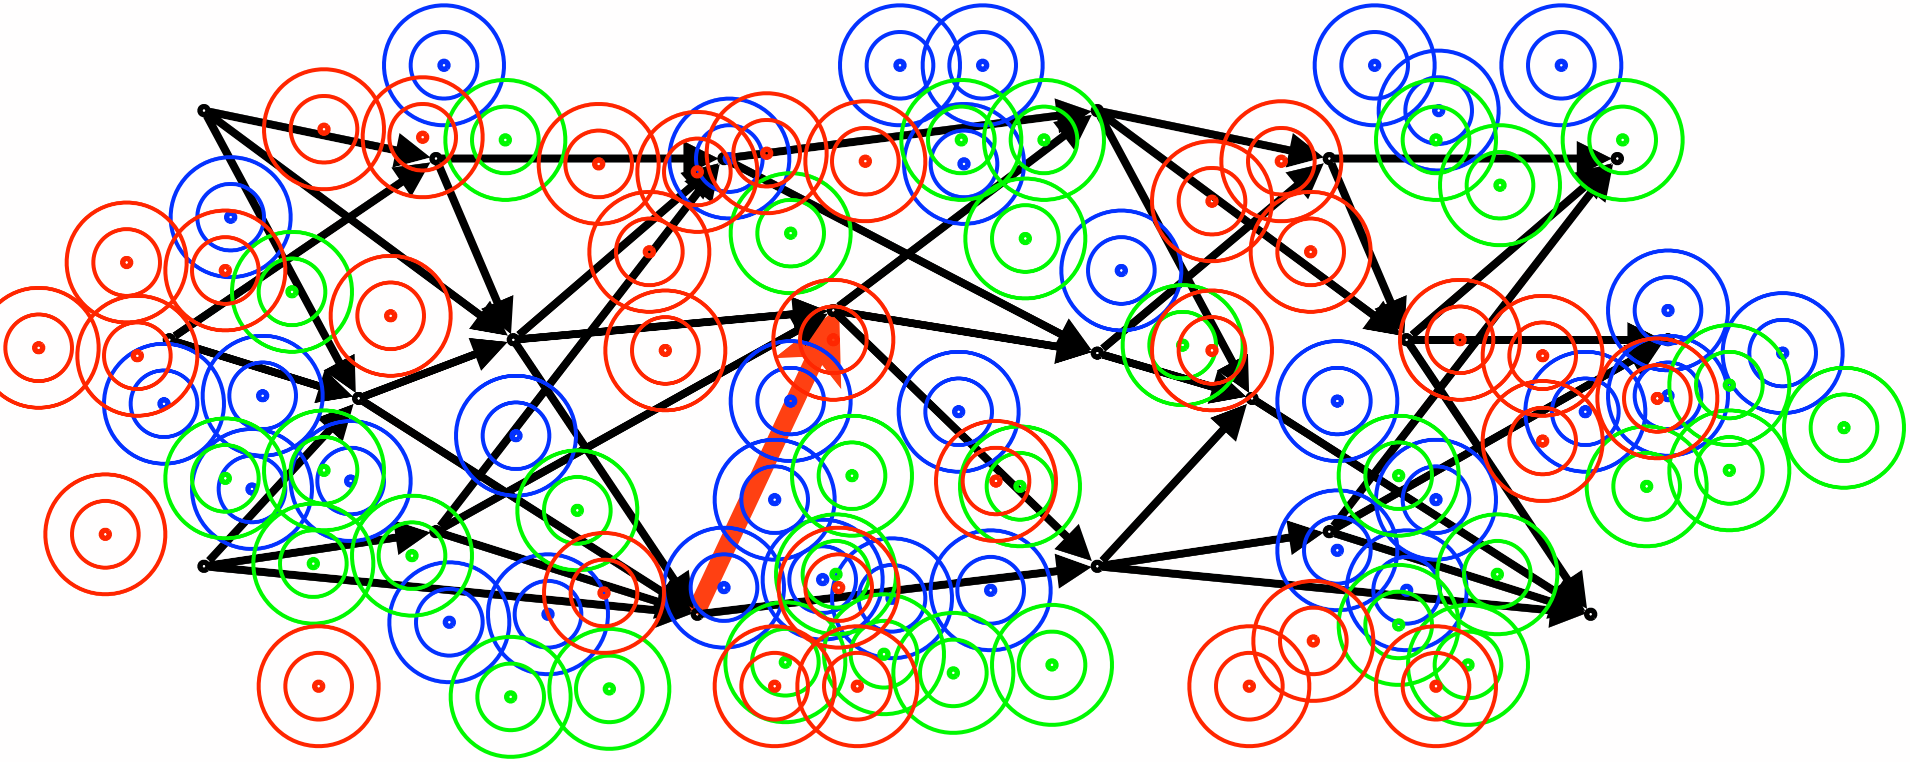

A noisy channel has the same impact on network capacity as a lossless link of the same capacity.

[Koetter, Effros, Medard 2009, 2011]

WHAT IS THE IMPACT
OF A SINGLE ONE OF SHANNON’S CHANNELS
ON A NETWORK’S CAPACITY?

Edge Removal in Wireline Networks

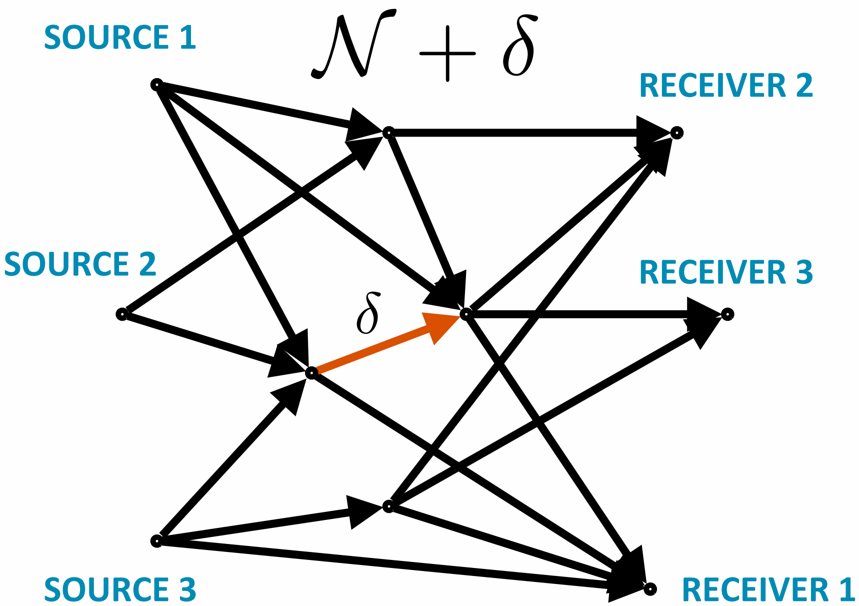

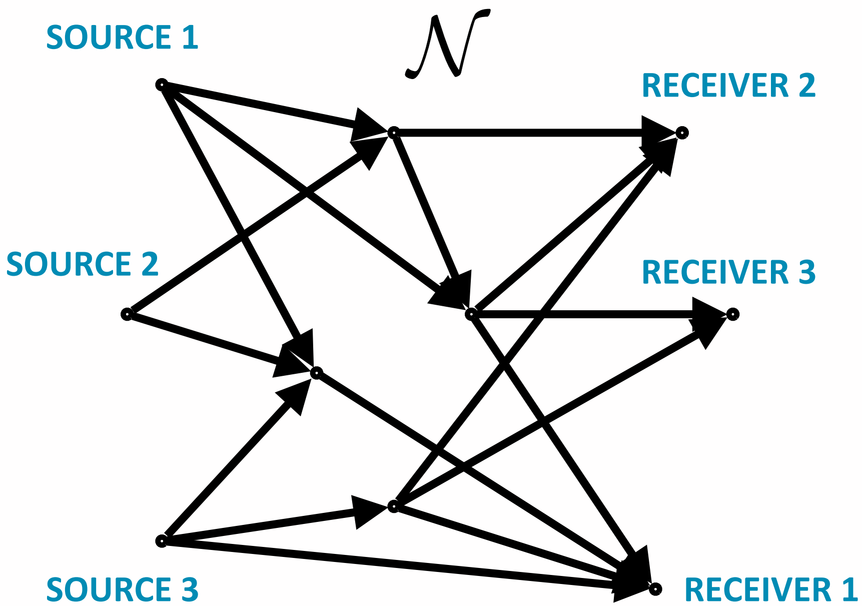

If I remove a "Shannon's channel" of capacity δ, how much can the network capacity change?

Is it true that the benefit of a capacity-δ edge can never exceed δ?

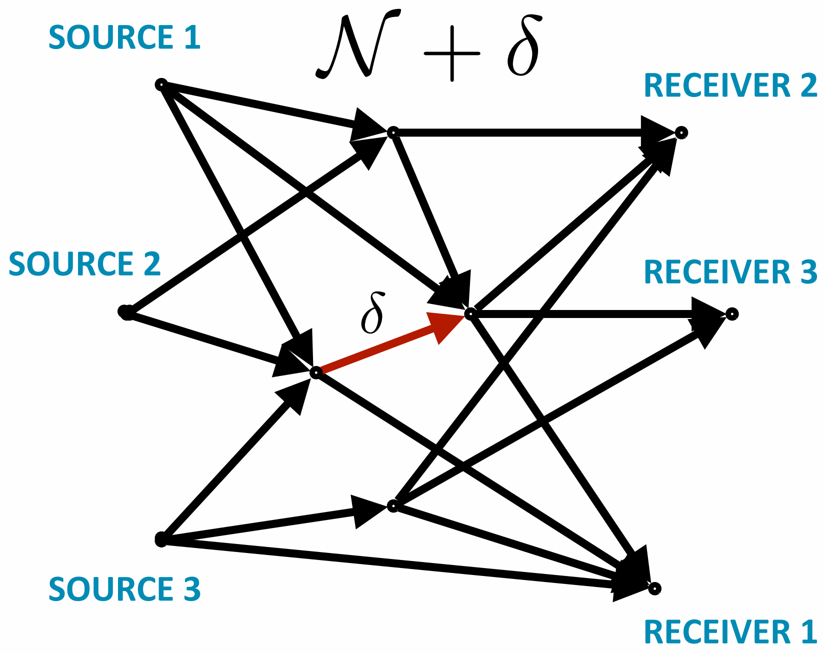

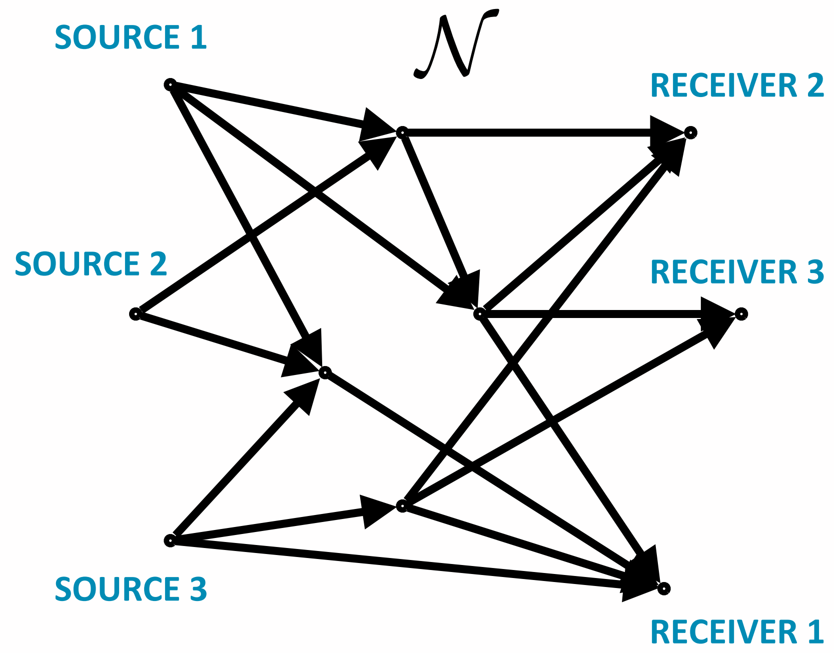

$$\text{Does } (R_1, R_2, R_3) \in \text{Capacity}(\mathcal{N} + \delta),$$

$$\text{imply}$$

$$(R_1 - \delta, R_2 - \delta, R_3 - \delta) \in \text{Capacity}(\mathcal{N})?$$

Is it true that the benefit of a capacity-δ edge can never exceed δ?

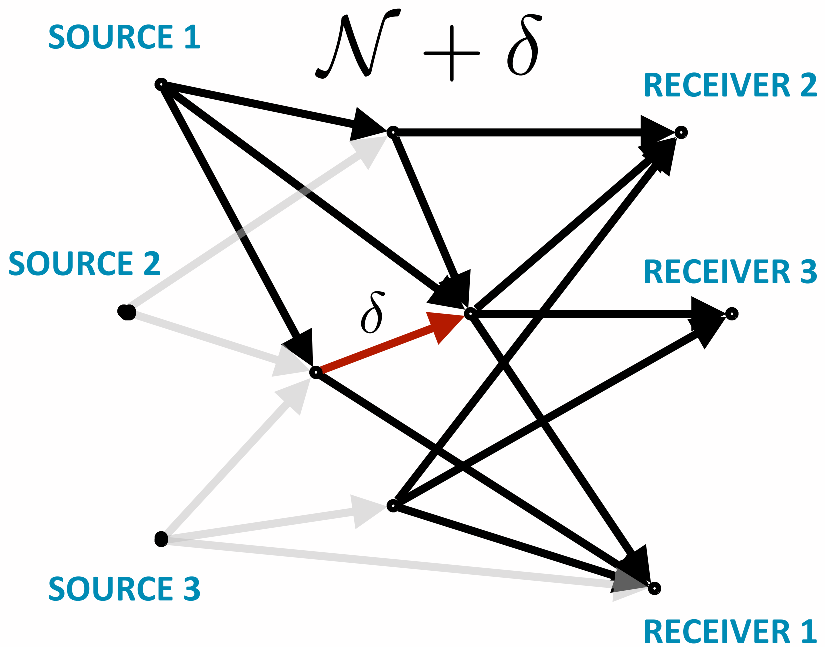

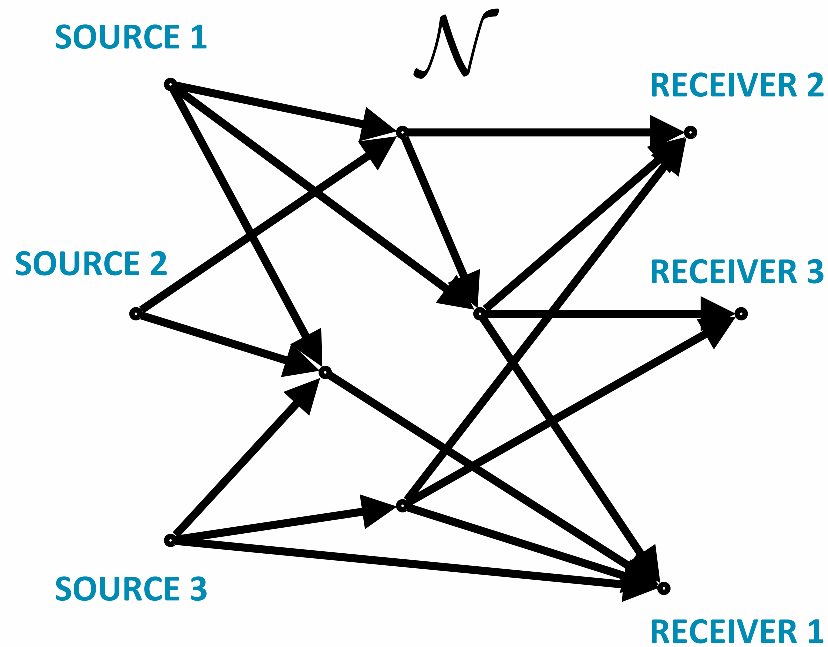

$$\text{Does } (R_1, R_2, R_3) \in \text{Capacity}(\mathcal{N} + \delta),$$

$$\text{imply}$$

$$(R_1 - \delta, R_2 - \delta, R_3 - \delta) \in \text{Capacity}(\mathcal{N})?$$

Is it true that the benefit of a capacity-δ edge can never exceed δ?

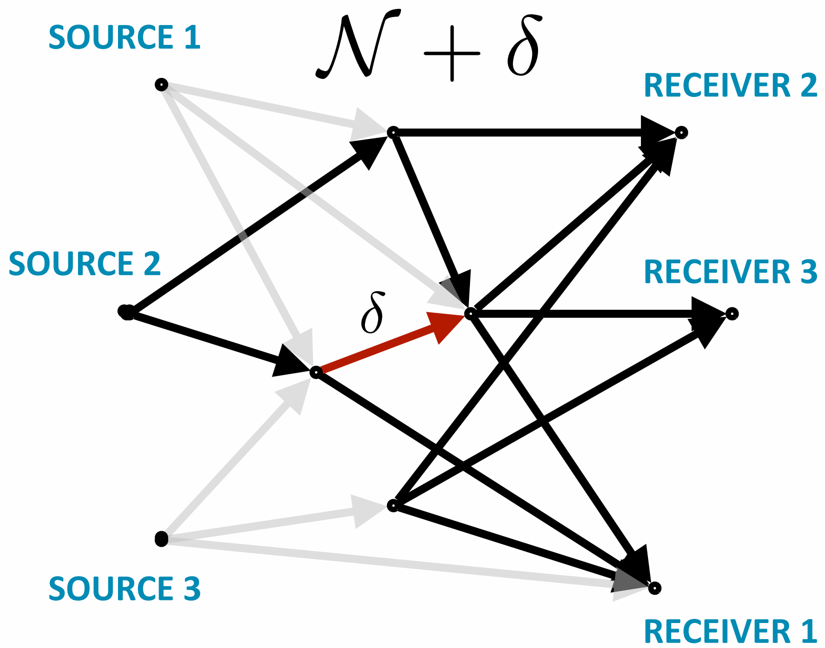

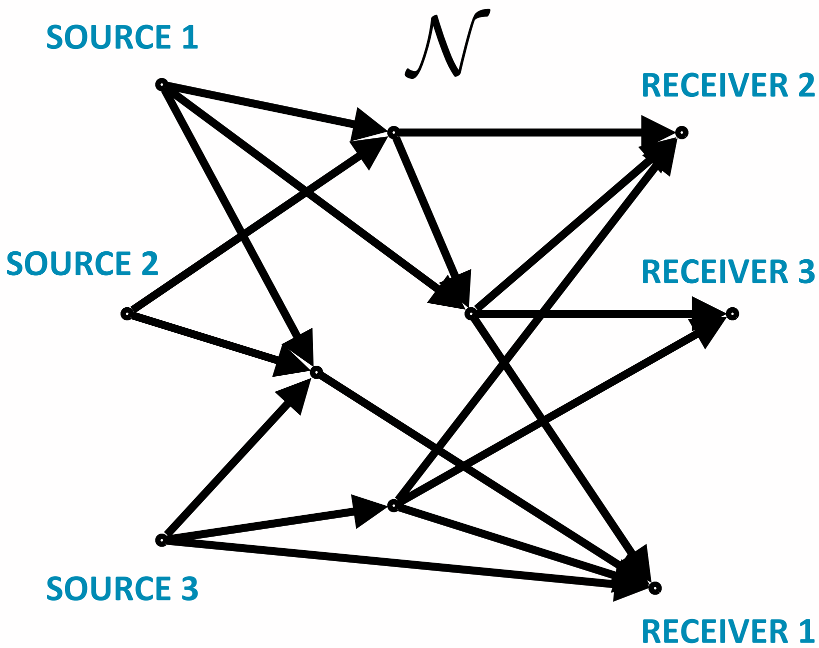

$$\text{Does } (R_1, R_2, R_3) \in \text{Capacity}(\mathcal{N} + \delta),$$

$$\text{imply}$$

$$(R_1 - \delta, R_2 - \delta, R_3 - \delta) \in \text{Capacity}(\mathcal{N})?$$

Is it true that the benefit of a capacity-δ edge can never exceed δ?

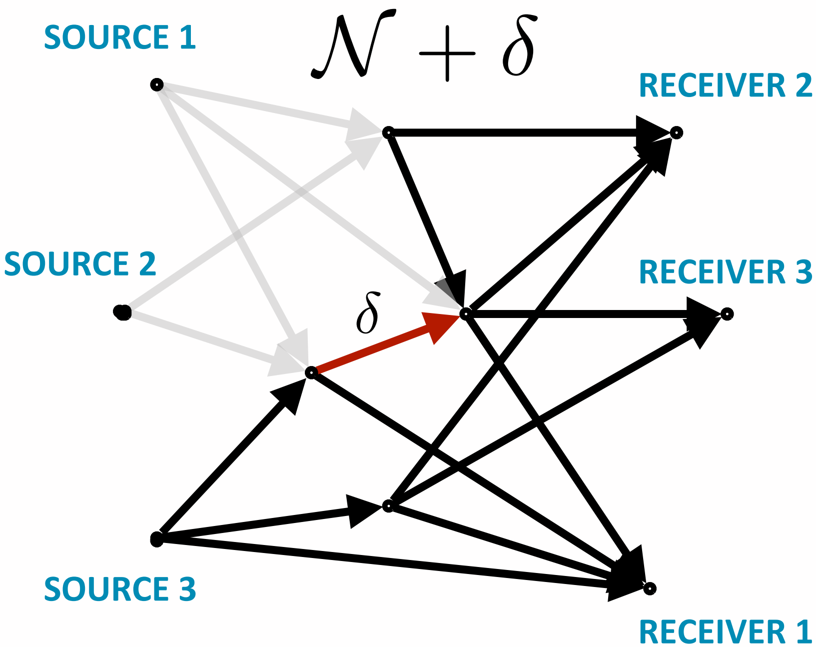

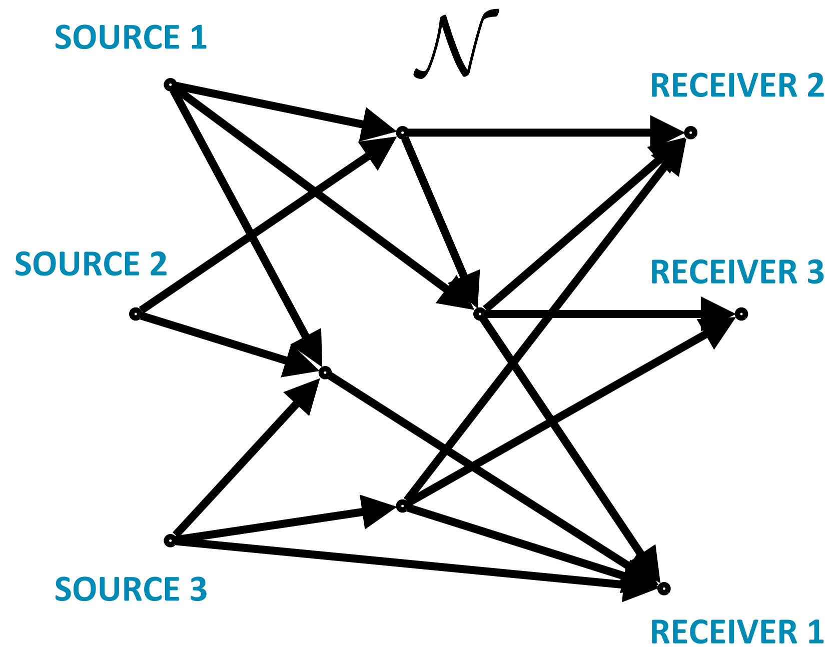

$$\text{Does } (R_1, R_2, R_3) \in \text{Capacity}(\mathcal{N} + \delta),$$

$$\text{imply}$$

$$(R_1 - \delta, R_2 - \delta, R_3 - \delta) \in \text{Capacity}(\mathcal{N})?$$

Is it true that the benefit of a capacity-δ edge can never exceed δ?

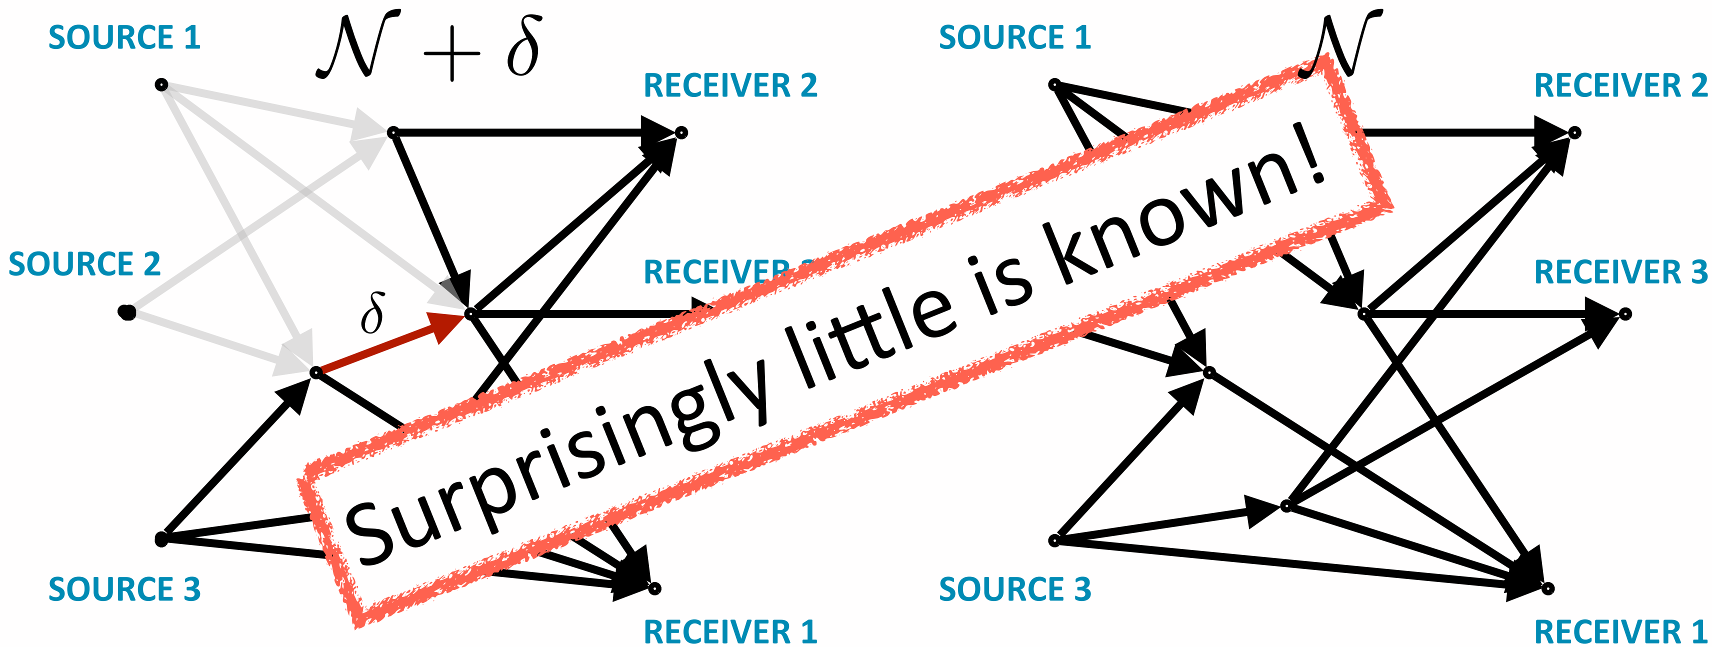

$$\text{Does } (R_1, R_2, R_3) \in \text{Capacity}(\mathcal{N}+\delta),$$

$$\text{imply}$$

$$(R_1-\delta, R_2-\delta, R_3-\delta) \in \text{Capacity}(\mathcal{N})?$$

The question remains unsolved for network coding.

[Jalali, Effros, Ho 2011, 2012, Langberg, Effros 2012, Lee, Langberg, Effros 2013]

- The edge removal property holds (='yes') for some networks.
  - cut-set bounds are tight (e.g., single- & multi-source multicast)
  - co-located sources, super-source networks, terminal edges
  - linear codes, "separable" codes
  - index coding
- No proof that the property always holds.
- No examples where property fails.
- The edge removal property holds for outer bounds.
  - Cut-set bound
  - Generalized network sharing bounds [Kamath, Tse, Anantharam 2011]
  - Linear Programming (LP) bound [Yeung 1997, Song, Yeung 2003]
- Equivalence to other problems (0- vs. ϵ-error, dep srcs, NC vs. IC, …)

Wireline networks: Intuition

[Jalali, Effros, Ho 2011, 2012, Langberg, Effros 2012, Lee, Langberg, Effros 2013]

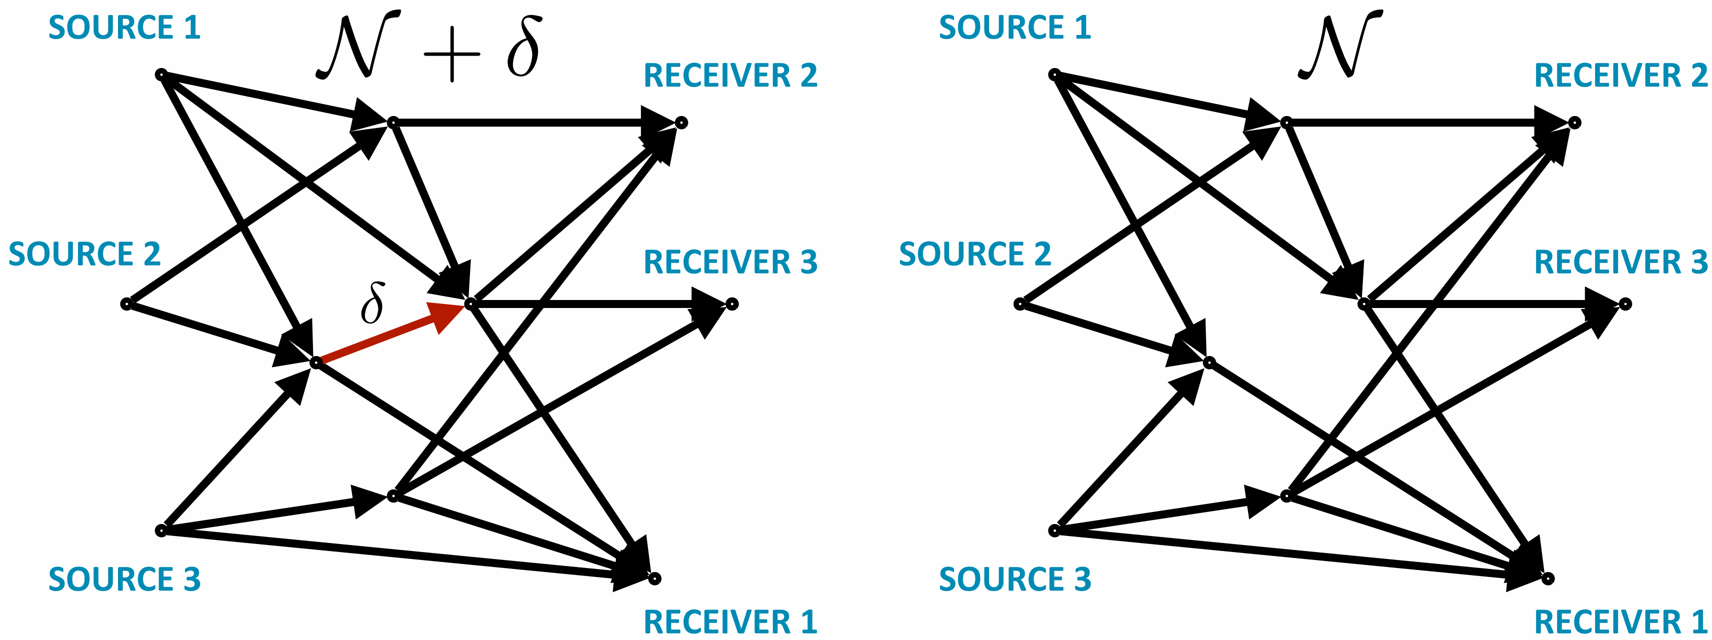

Only send source values that give the most common transmission across our connection.

The number of such transmissions supports rate $(R_1 - \delta, R_2 - \delta, R_3 -. \delta)$

Challenge: This strategy may not always be possible.

OUR WORLD IS INCREASINGLY WIRELESS.
DOES THE ANSWER CHANGE?

What happens in wireless networks?

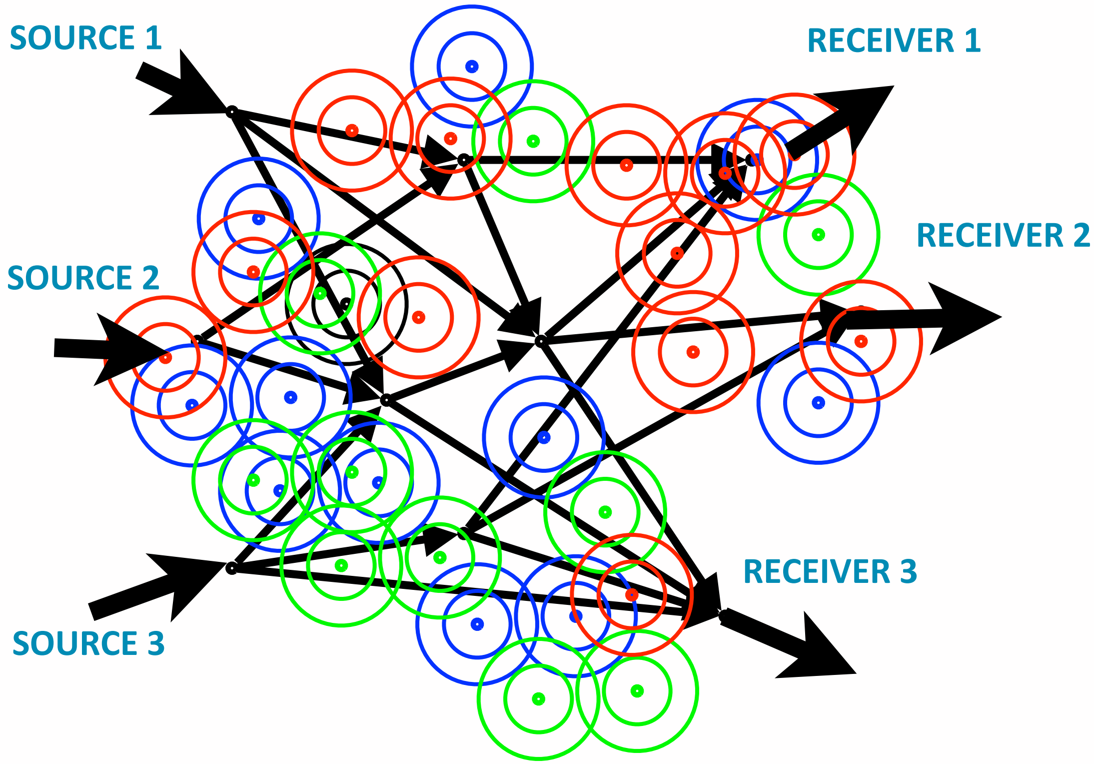

$$\text{Does } (R_1, R_2, R_3) \in \text{Capacity}(\mathcal{N} + \delta),$$

$$\text{imply}$$

$$(R_1 - \delta, R_2 - \delta, R_3 - \delta) \in \text{Capacity}(\mathcal{N})?$$

In prior literature, the impact of any edge was bounded by the capacity of that edge.

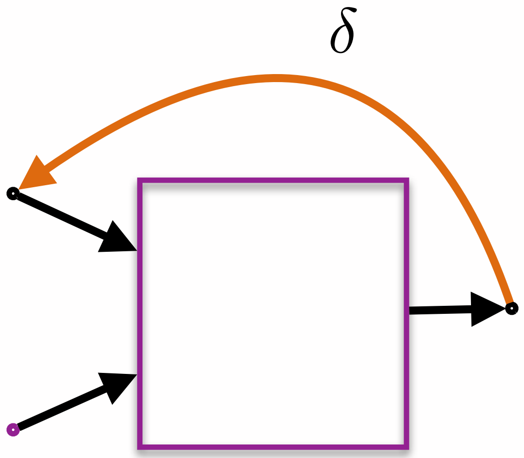

[Sarwate & Gastpar 2009]

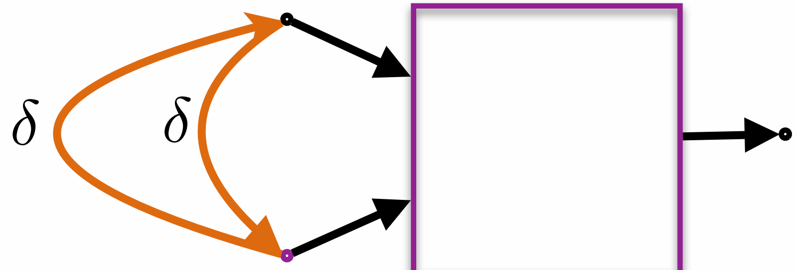

[Willems 1983]

$$\text{Does } (R_1, R_2, R_3) \in \text{Capacity}(\mathcal{N} + \delta),$$

$$\text{imply}$$

$$(R_1 - \delta, R_2 - \delta, R_3 - \delta) \in \text{Capacity}(\mathcal{N})?$$

YES.

For general memoryless networks, the edge removal property sometimes fails.

[Noorzad, Effros, Langberg, Ho 2014]

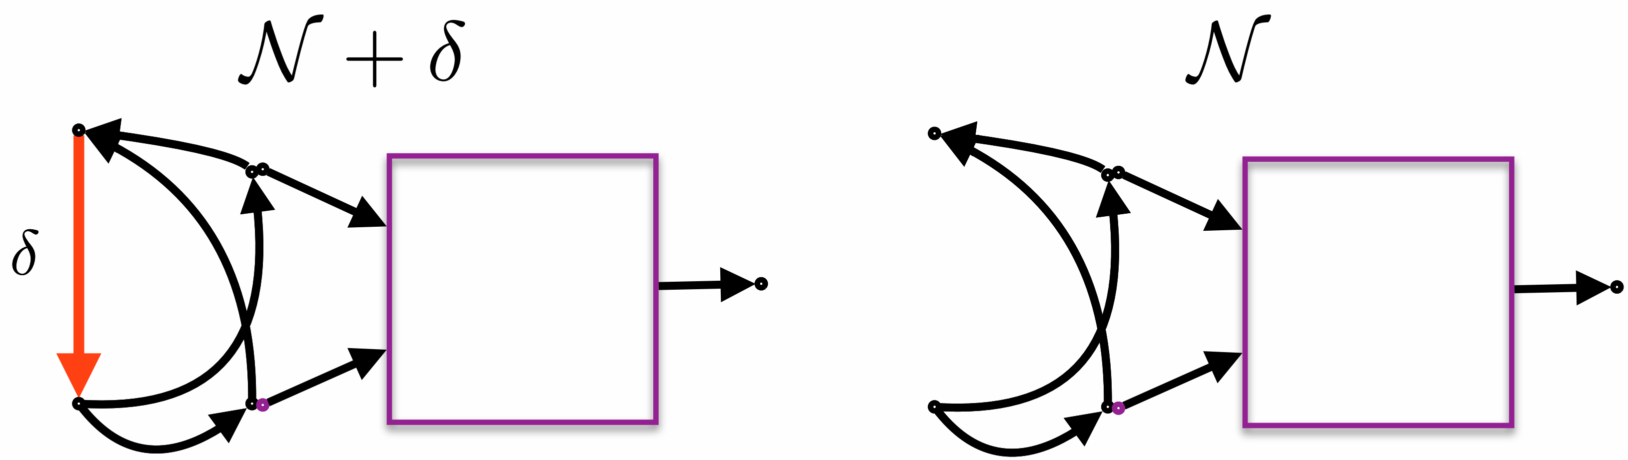

$$\text{Does } (R_1, R_2, R_3) \in \text{Capacity}(\mathcal{N} + \delta),$$

$$\text{imply}$$

$$(R_1 - \delta, R_2 - \delta, R_3 - \delta) \in \text{Capacity}(\mathcal{N})?$$

NO.

In fact, the property fails even if we loosen the constraint.

[Noorzad, Effros, Langberg, Ho 2014]

$\mathcal{N} + \delta$

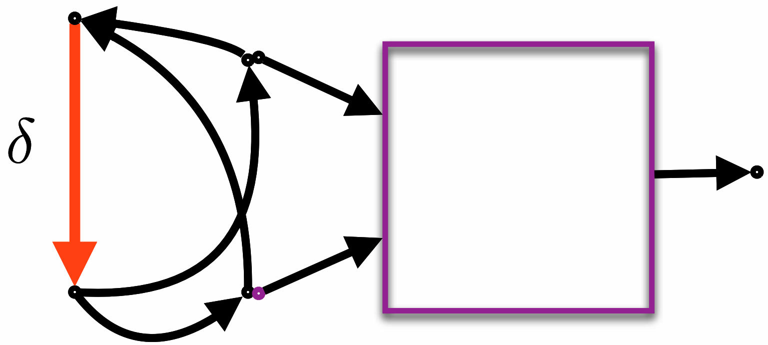

$\mathcal{N}$

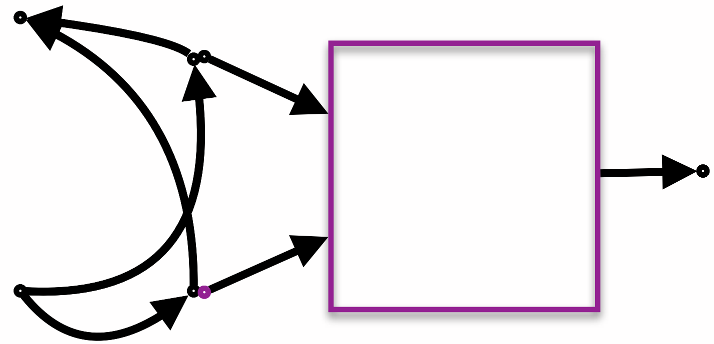

Does

$$(R_1, R_2, R_3) \in \text{Capacity}(\mathcal{N} + \delta) \text{ imply}$$

$$(R_1 - f(\delta), R_2 - f(\delta), R_3 - f(\delta)) \in \text{Capacity}(\mathcal{N})$$

NO!!! (for ANY polynomial f)

The power of a connection can FAR exceed its capacity!

[Noorzad, Effros, Langberg, Ho 2014]

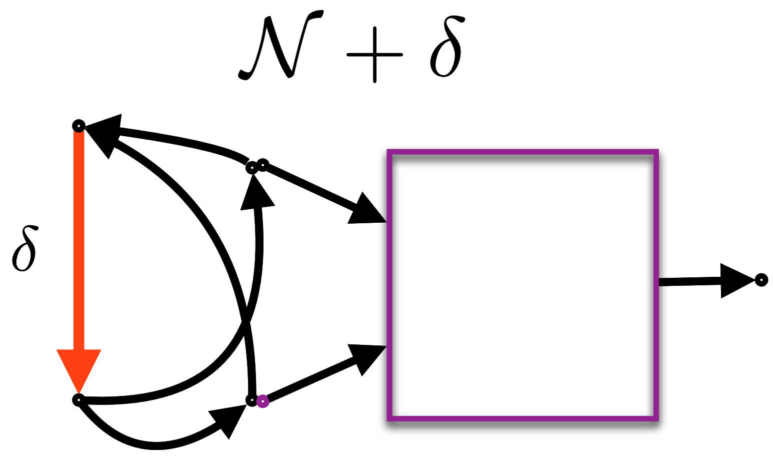

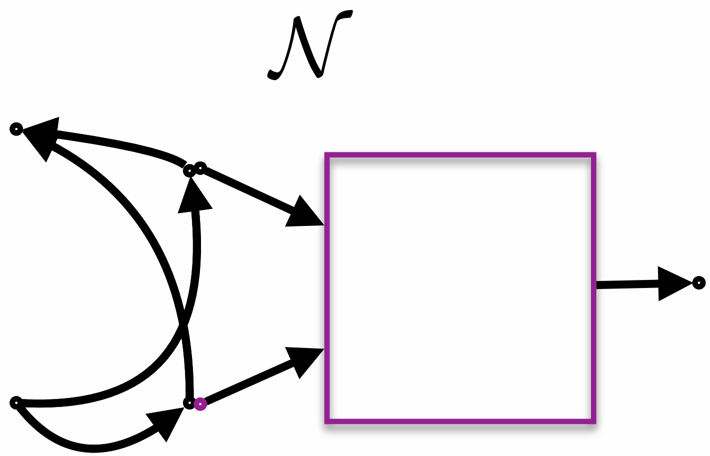

Adding a δ-capacity link can increase the network capacity ALMOST EXPONENTIALLY in δ.

The power of a connection can FAR exceed its capacity!

[Noorzad, Effros, Langberg, Ho 2014]

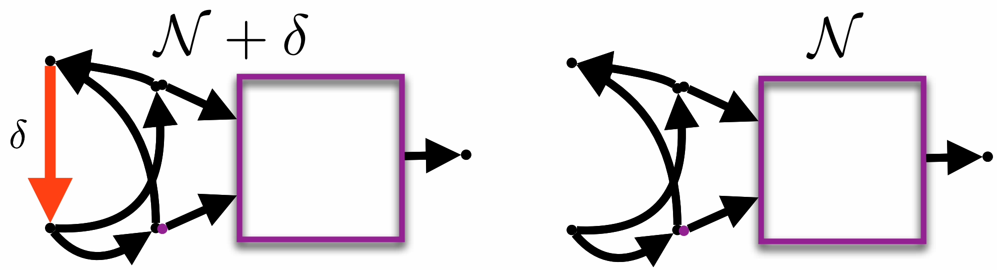

$$\mathcal{X}_1 = \mathcal{X}_2 = \{1, \ldots, 2^m\}$$

$$\mathcal{Y} = (\mathcal{X}_1 \times \mathcal{X}_2) \cup \{E\} \text{ (E denotes "erasure")}$$

$$B = \begin{bmatrix} b(1,1) & b(1,2) & \ldots & b(1,2^m) \\ b(2,1) & b(2,2) & \ldots & b(2,2^m) \\ \vdots & \vdots & \ddots & \vdots \\ b(2^m,1) & b(2^m,2) & \ldots & b(2^m,2^m) \end{bmatrix}$$

$$p(y|x_1,x_2) = \begin{cases} 1(y=(x_1,x_2)) & \text{if } b(x_1,x_2)=0 \\ 1(y=E) & \text{if } b(x_1,x_2)=1 \end{cases}$$

Proof (counter-example)

[Noorzad, Effros, Langberg, Ho 2014]

$\mathcal{X}_1 = \mathcal{X}_2 = \{1, \ldots, 2^m\}$

$\mathcal{Y} = (\mathcal{X}_1 \times \mathcal{X}_2) \cup \{(E, E)\}$ (E denotes "erasure")

$$B = \begin{bmatrix} b(1,1) & b(1,2) & \ldots & b(1,2^m) \\ b(2,1) & b(2,2) & \ldots & b(2,2^m) \\ \vdots & \vdots & \ddots & \vdots \\ b(2^m,1) & b(2^m,2) & \ldots & b(2^m,2^m) \end{bmatrix}$$

$$p(y|x_1,x_2) = \begin{cases} 1(y = (x_1,x_2)) & \text{if } b(x_1,x_2) = 0 \\ 1(y = E) & \text{if } b(x_1,x_2) = 1 \end{cases}$$

$\underline{\exists\ B \text{ such that:}}$

$\exists\ \frac{2^m}{2^{\log(m \log m)}}$-partition of $\mathcal{X}_1$ ($\mathcal{X}_2$) s.t. each "cell" contains ≥ 1 "0"

Ensures $\mathcal{C}(\mathcal{N} + \delta)$ large ($R_1 + R_2 = 2m - 2\log(m \log m)$ ach)

Every sufficiently large sub-matrix has fraction $\geq 1 - \epsilon$ "1"s

Ensures $\mathcal{C}(\mathcal{N})$ small ($R_1 + R_2 < 1.25m$)

The benefit of an edge can far exceed its capacity...

[Noorzad, Effros, Langberg, Ho 2014]

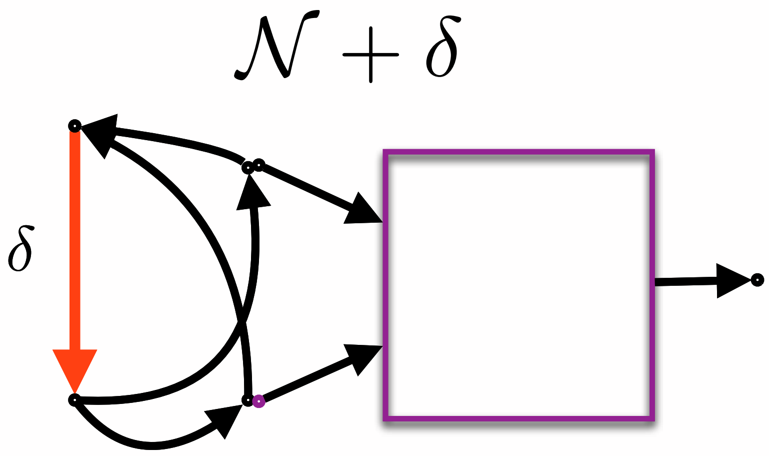

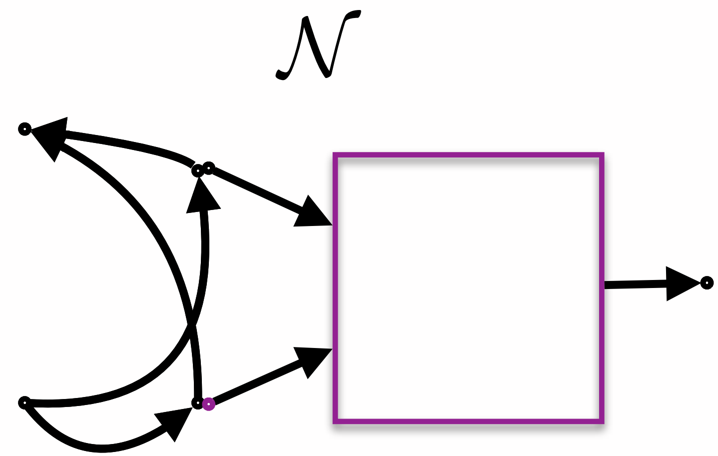

But this is an artificial example...

What happens in more realistic channels?

[Noorzad, Effros, Langberg 2015]

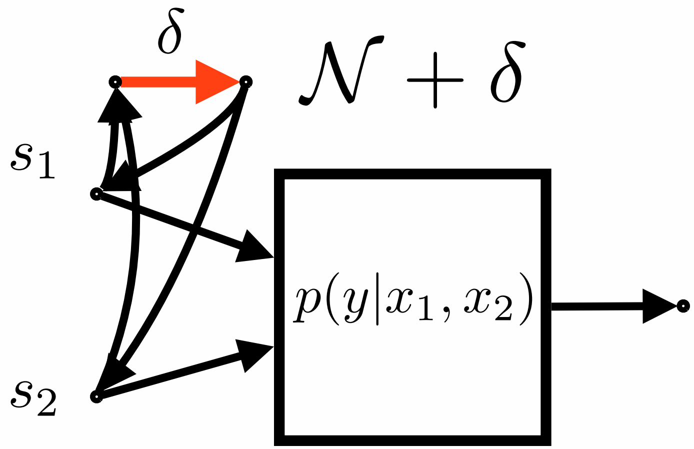

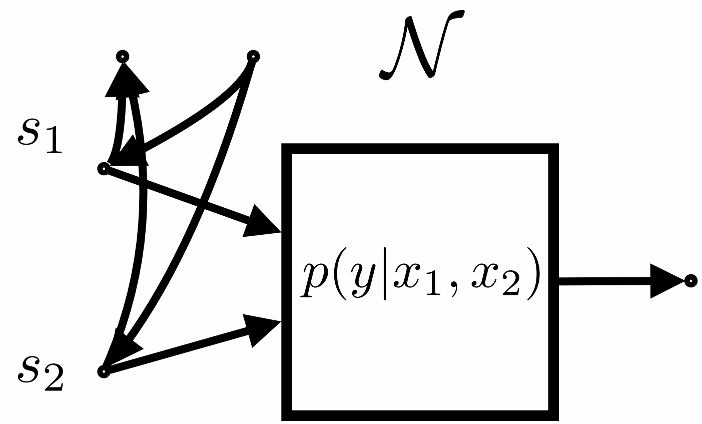

If cooperation helps at all, then a little cooperation helps a LOT!

Can rate-0 cooperation ever help???

[Noorzad, Effros, Langberg 2016]

Surprisingly, at least in
the case of zero-error capacity,
the answer is YES!

[Langberg & Effros 2016]

In this case, even
a single bit can change capacity!

Summary

- Shannon started a communication revolution by characterizing the capacity of a single channel.
- Shannon's work is the *first* (not *last*) word on the impact of a channel.
- For wireline networks, it is unknown whether the benefit of a single edge can ever exceed its capacity.
  - In some cases, it provably cannot.
  - Current outer bounds likewise suggest that it cannot.
  - The question is related to other interesting unsolved questions.
- For networks with wireless connections, the benefit of a a single edge can FAR exceed its capacity.
  - The gap can be large.
  - The slope can be infinite.
  - The benefit can be discontinuous.
- The question of a channel's impact on network capacity is, perhaps, the most fundamental open question in information theory.