Fundamental rate delay tradeoffs in multipath routed and network coded networks

Steven Weber
(joint work with John M. Walsh)
Drexel University, Dept. of ECE
Philadelphia, PA 19104
sweber@ece.drexel.edu

September 11, 2008
DARPA ITMANET Project Meeting
Arlington, VA

MANET (S,D) pairs employ multiple paths with uncertain end to end delays

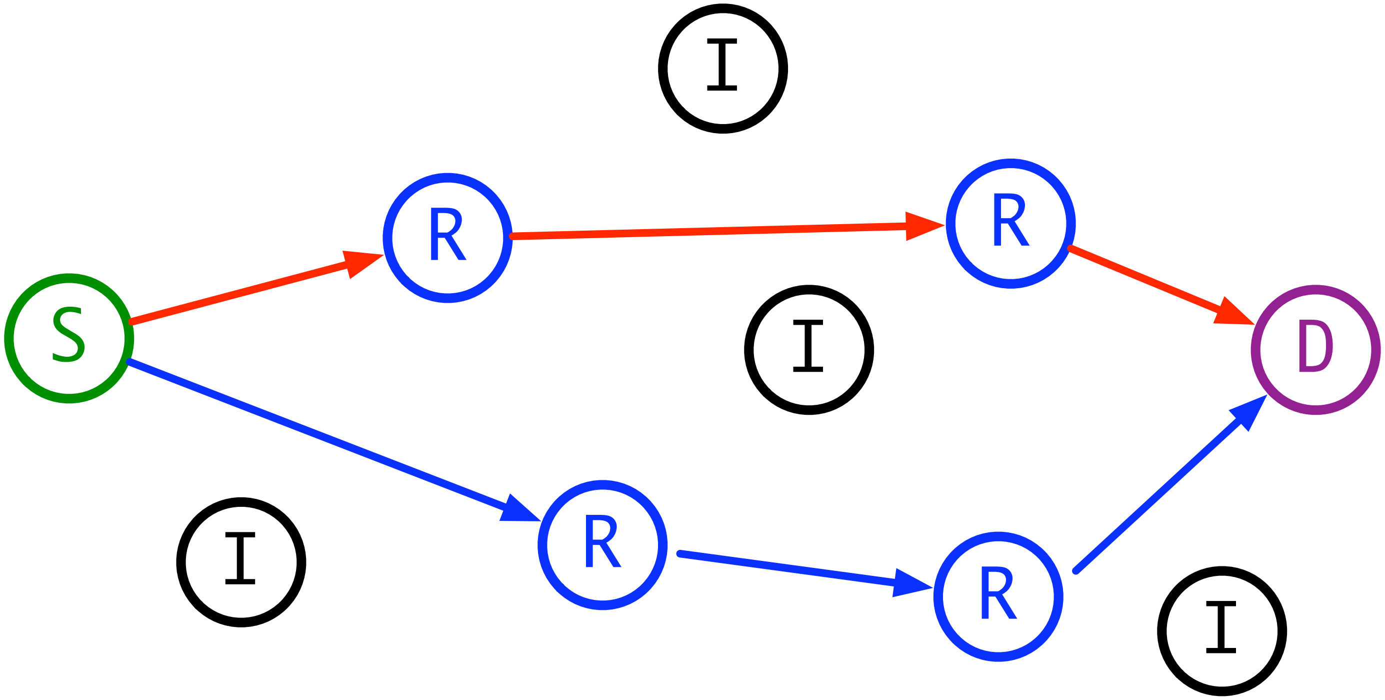

- "Opportunistic" routing necessary due to movement, channel conditions, buffer congestion...
- Multiple paths with uncertain end to end delays means packets may arrive out of order at destination

Sending an ordered source over a packet permuting network

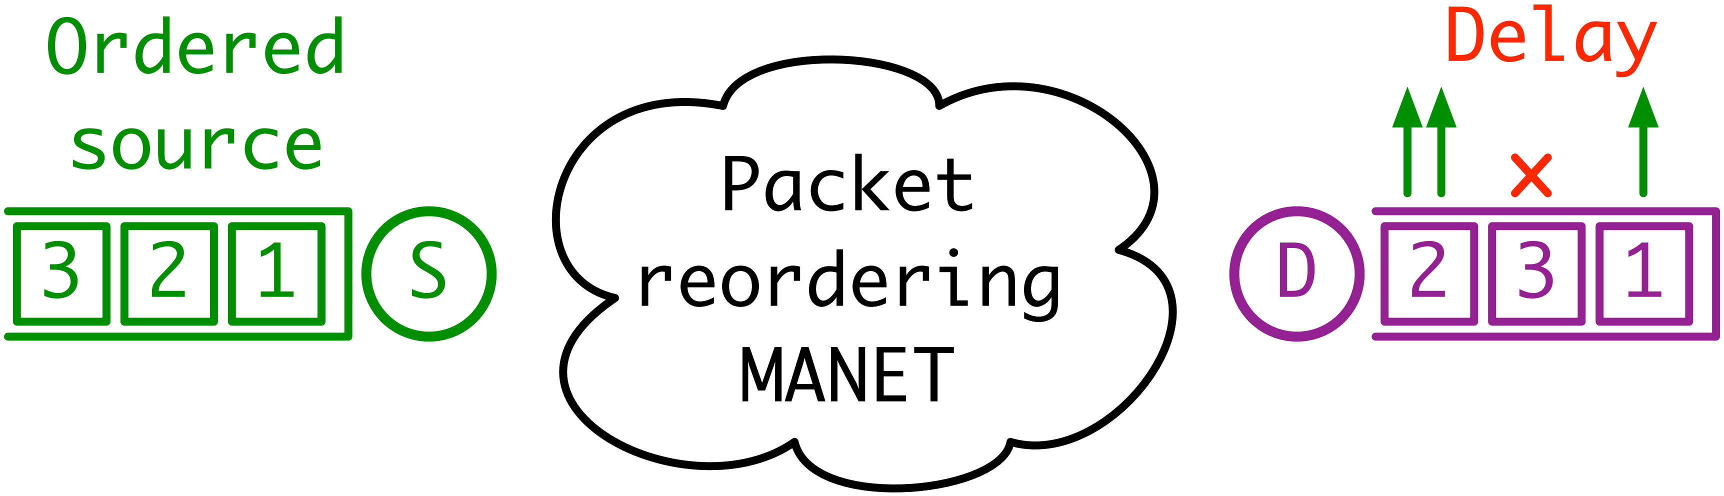

- Network will reorder and delay but not lose packets
- Playback of (uncoded) ordered source delayed by reordering

How to use two paths to a destination?

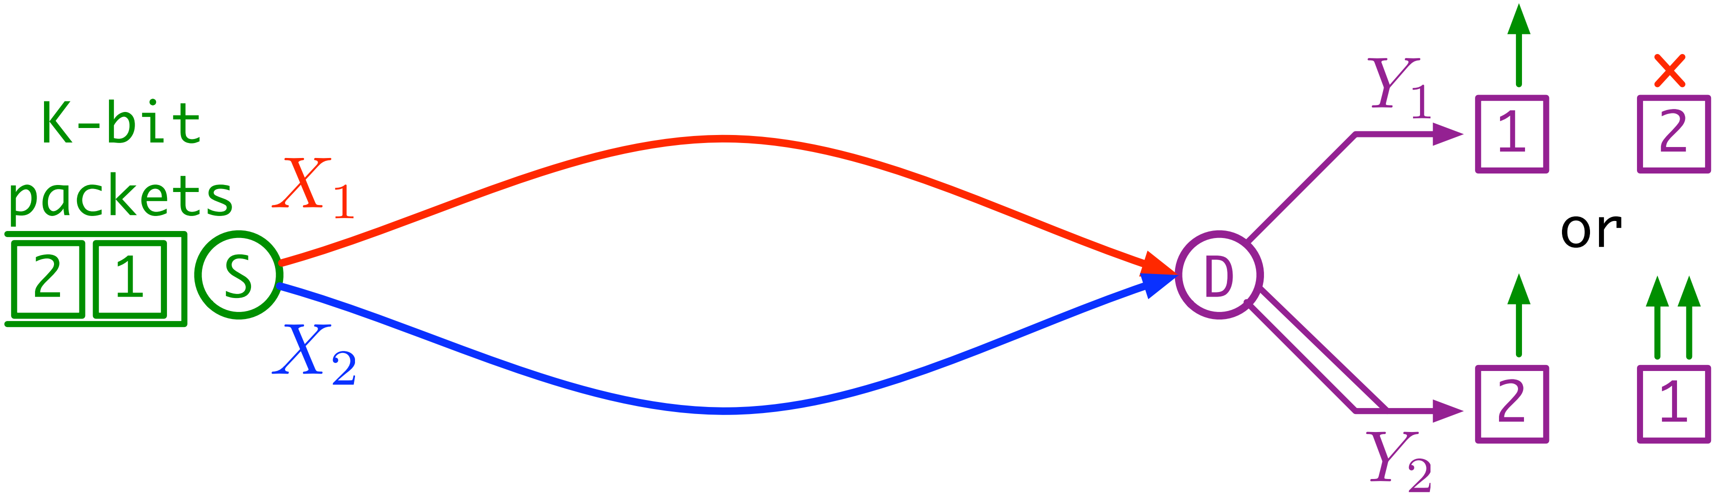

$$\begin{aligned} Y_1 &= X_{\pi(1)} \\ Y_2 &= (X_{\pi(1)}, X_{\pi(2)}) \end{aligned} \qquad \begin{aligned} p(Y_1 = X_1) &= 1 - p \\ p(Y_1 = X_2) &= p \end{aligned}$$

- Packets arrive in order w.p. $1 - p$, in reverse order w.p. p
- Y_1 (Y_2) denotes information available upon first (second) arrival
- Decoding delay modeled as a low/zero rate to receiver Y_1

How to use two paths to a destination?

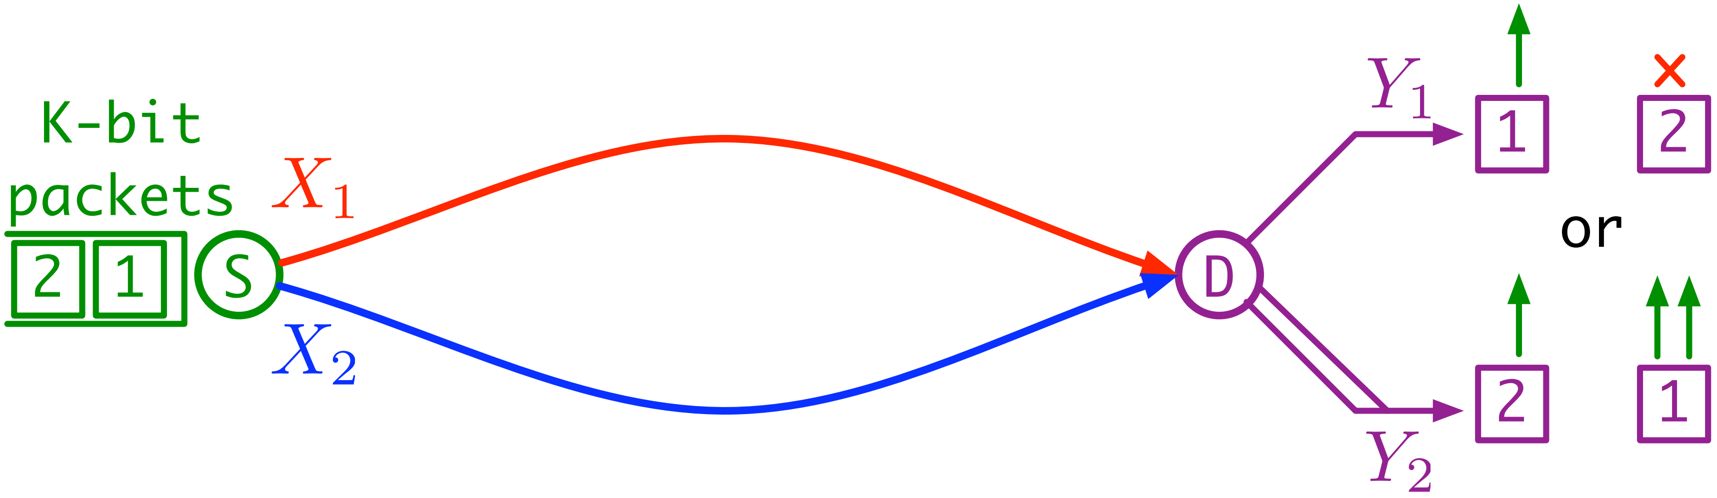

- Minimize "delay": ensure content arrives in order by sending same content on both paths ⇒ low throughput
- Maximize "throughput": send independent packets on the two paths ⇒ higher delay
- Each packet reception instant is a receiver, rate tradeoff among receivers captures T-D tradeoff

T-D tradeoff in the NequIT architecture

Non-equilibrium Information Theory

Intellectual Architecture

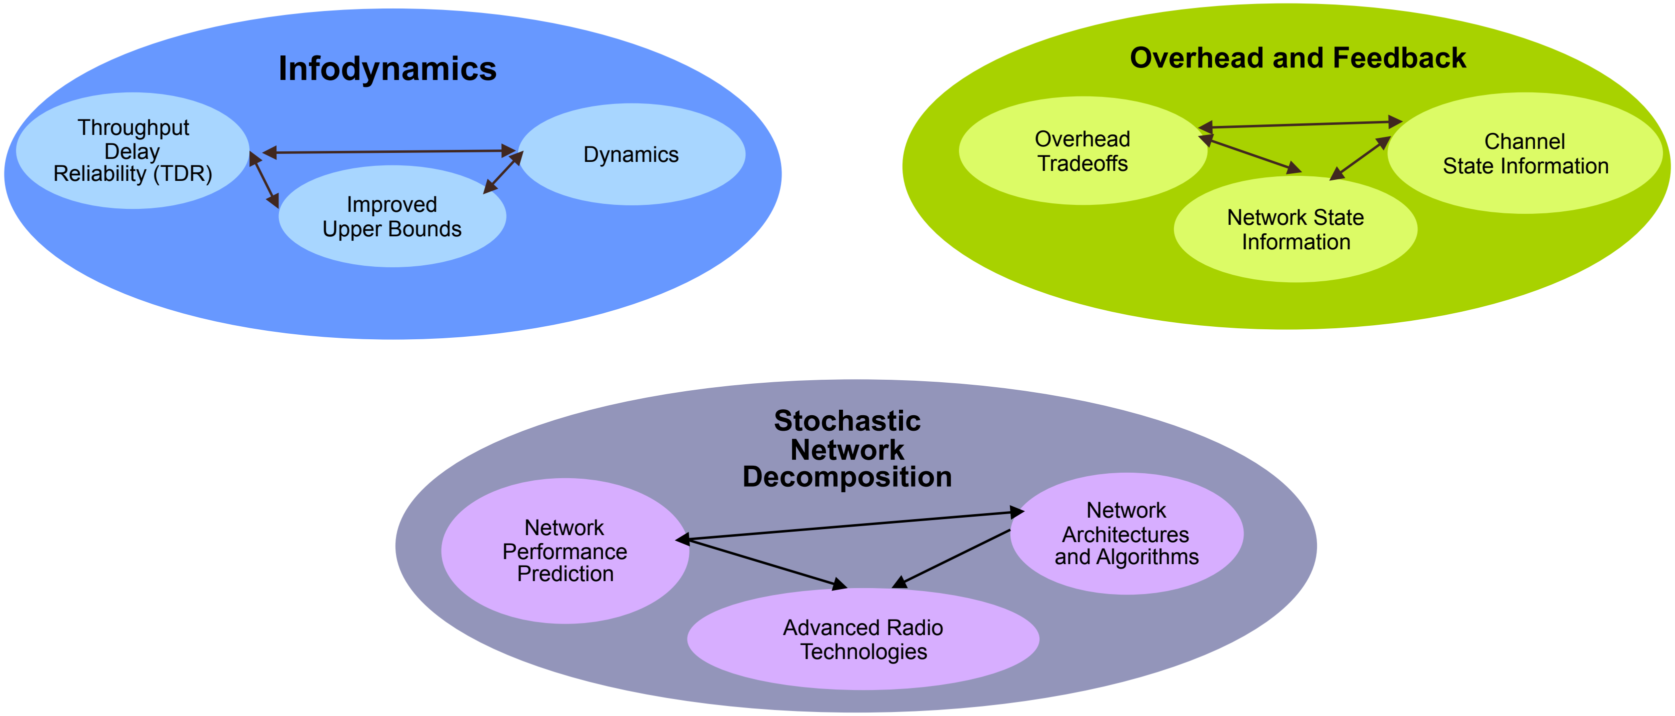

Infodynamics

Throughput Delay Reliability (TDR)

Dynamics

Improved Upper Bounds

Overhead and Feedback

Overhead Tradeoffs

Channel State Information

Network State Information

Stochastic Network Decomposition

Network Performance Prediction

Network Architectures and Algorithms

Advanced Radio Technologies

Same problem, depicted as a broadcast channel

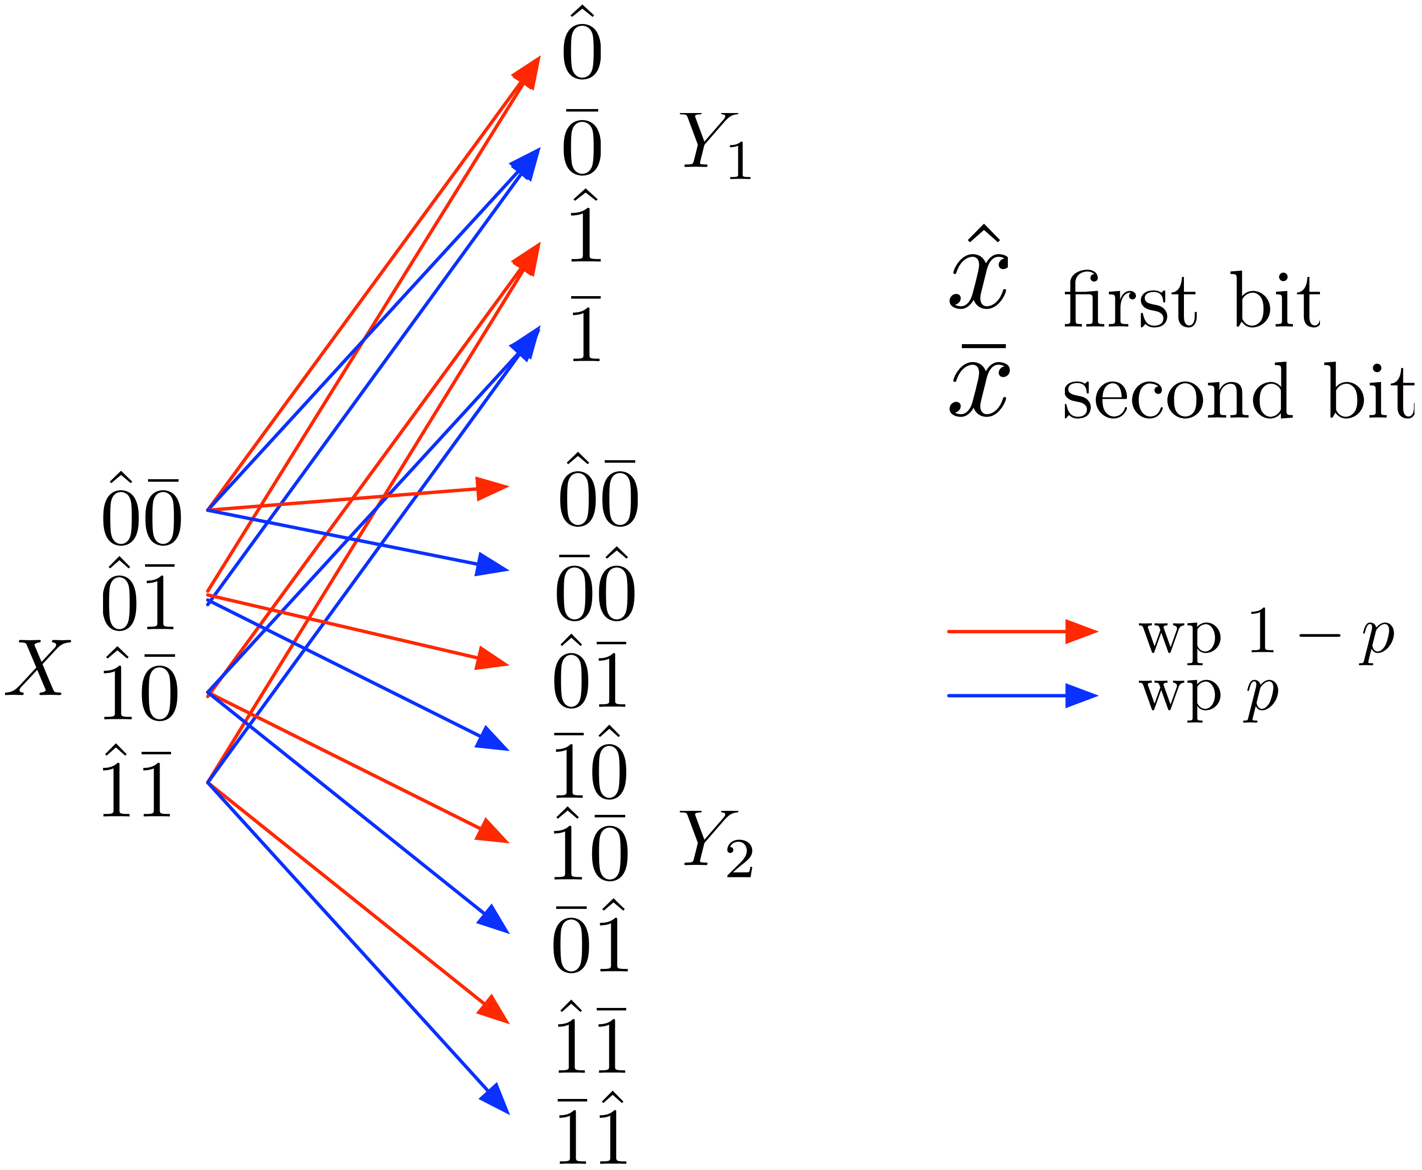

In fact, it is a degraded broadcast channel

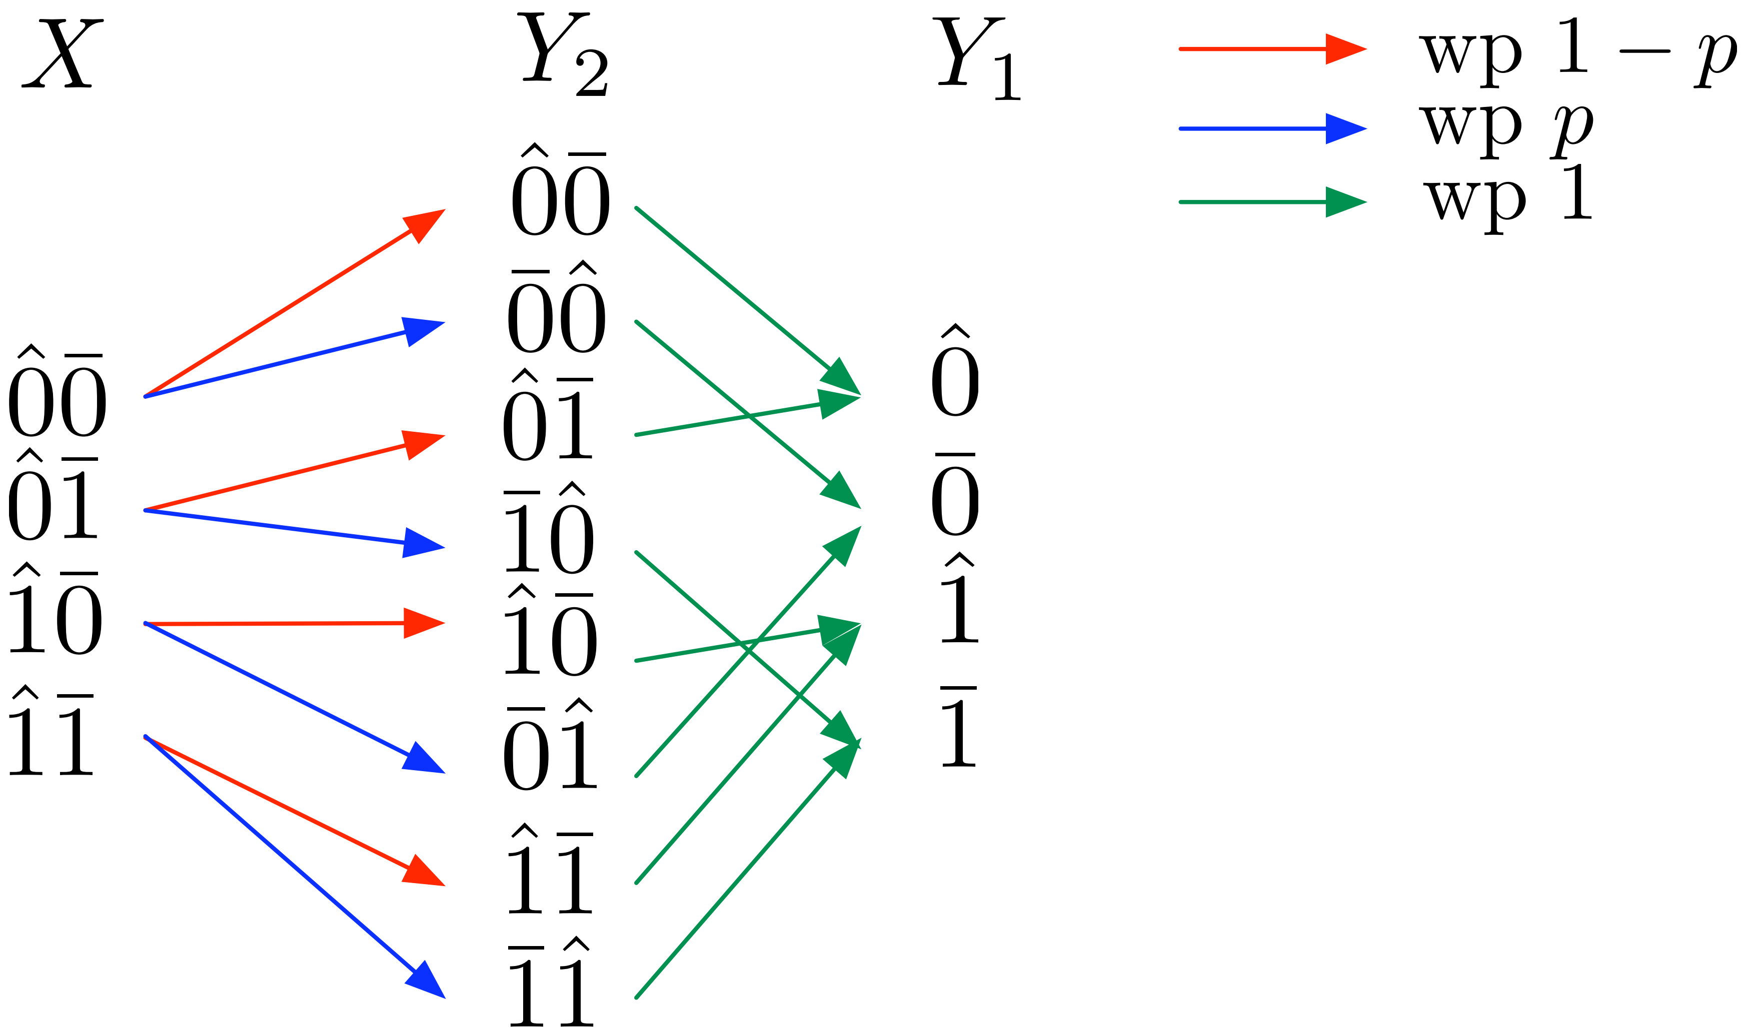

X

$\hat{0}\bar{0}$

$\hat{0}\bar{1}$

$\hat{1}\bar{0}$

$\hat{1}\bar{1}$

Y_2

$\hat{0}\bar{0}$

$\bar{0}\hat{0}$

$\hat{0}\bar{1}$

$\bar{1}\hat{0}$

$\hat{1}\bar{0}$

$\bar{0}\hat{1}$

$\hat{1}\bar{1}$

$\bar{1}\hat{1}$

Y_1

$\hat{0}$

$\bar{0}$

$\hat{1}$

$\bar{1}$

🟥 wp $1-p$

🟦 wp p

🟩 wp 1

The capacity region of the generic DBC is known

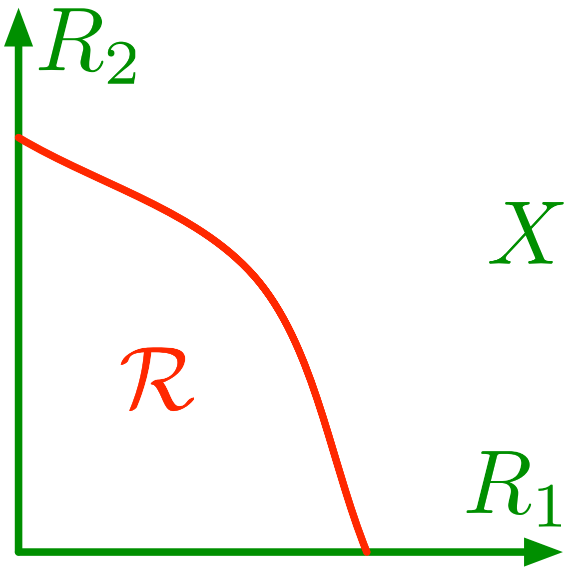

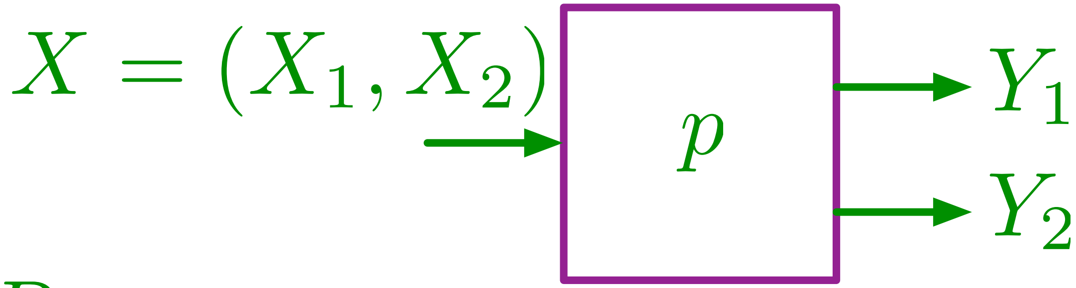

The capacity region of the discrete memoryless DBC is the closure of the convex hull of the region $\mathcal{R}$ of rates satisfying

$$\begin{aligned} R_1 &\leq I(Y_1; X_1) \\ R_2 &\leq I(Y_2; X|X_1) \end{aligned}$$

for some discrete rv X (with bounded support) where $X_1 \to X \to Y_2 \to Y_1$ form a Markov chain. (Cover 1972, El Gamal 1978, Cover 1998)

Extreme points: duplication and independence

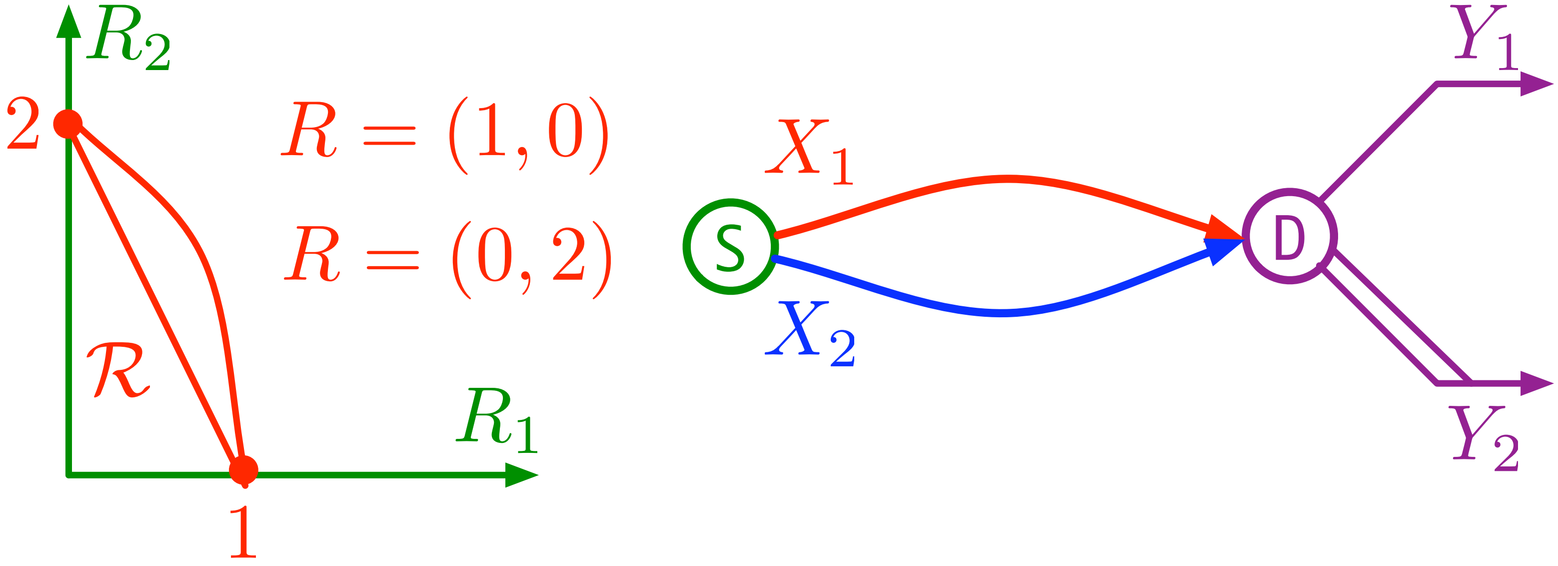

- Rate point $(1, 0)$ requires first receiver be able to decode regardless of which packet arrives: achieved by duplication $(X_1 = X_2)$
- Rate point $(0, 2)$ requires no decoding of first receiver: achieved by sending independent packets $(X_1 \perp X_2)$
- Time-sharing gives easy inner bound

Extreme cases: uniform and degenerate distributions

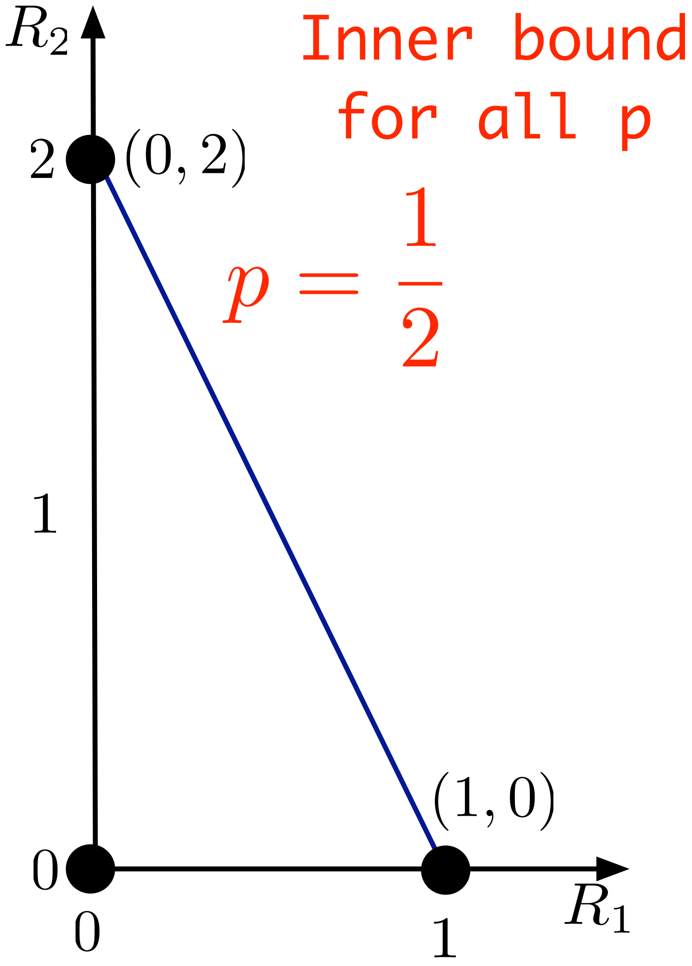

Inner bound for all p

$p = \frac{1}{2}$

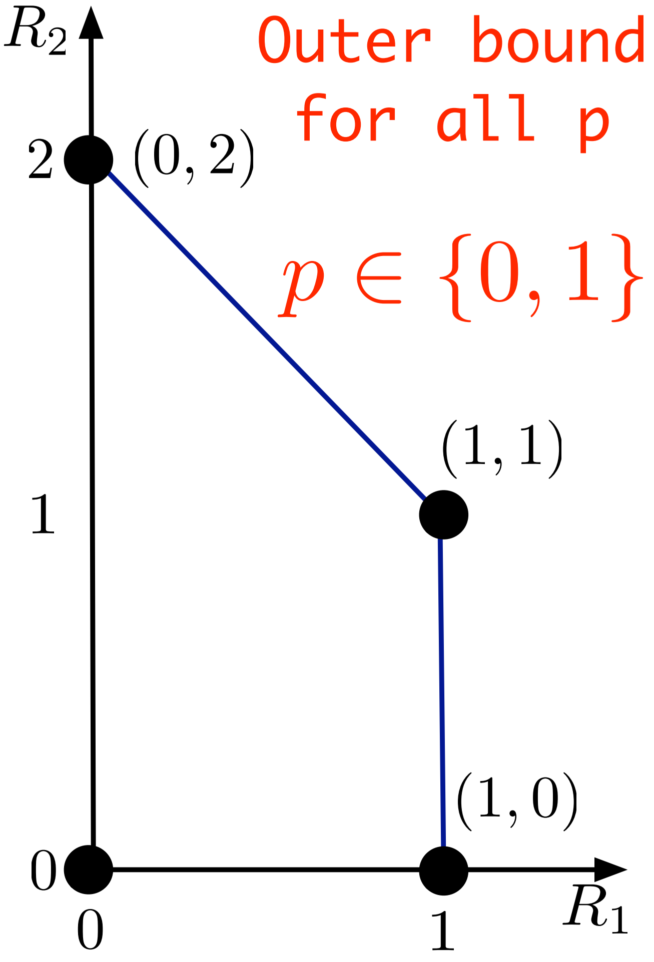

Outer bound for all p

$p \in \{0, 1\}$

Capacity region for two packets

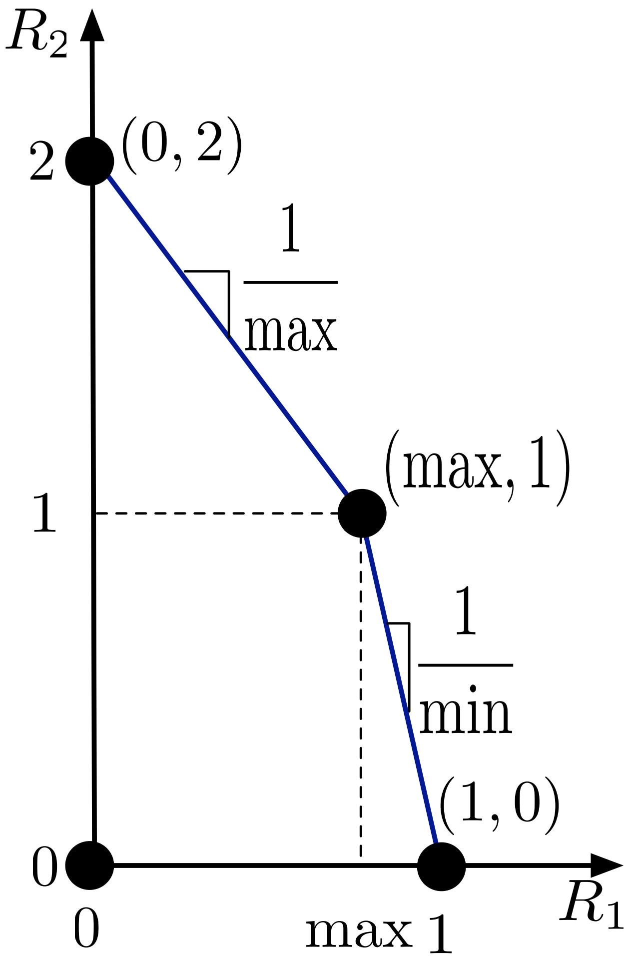

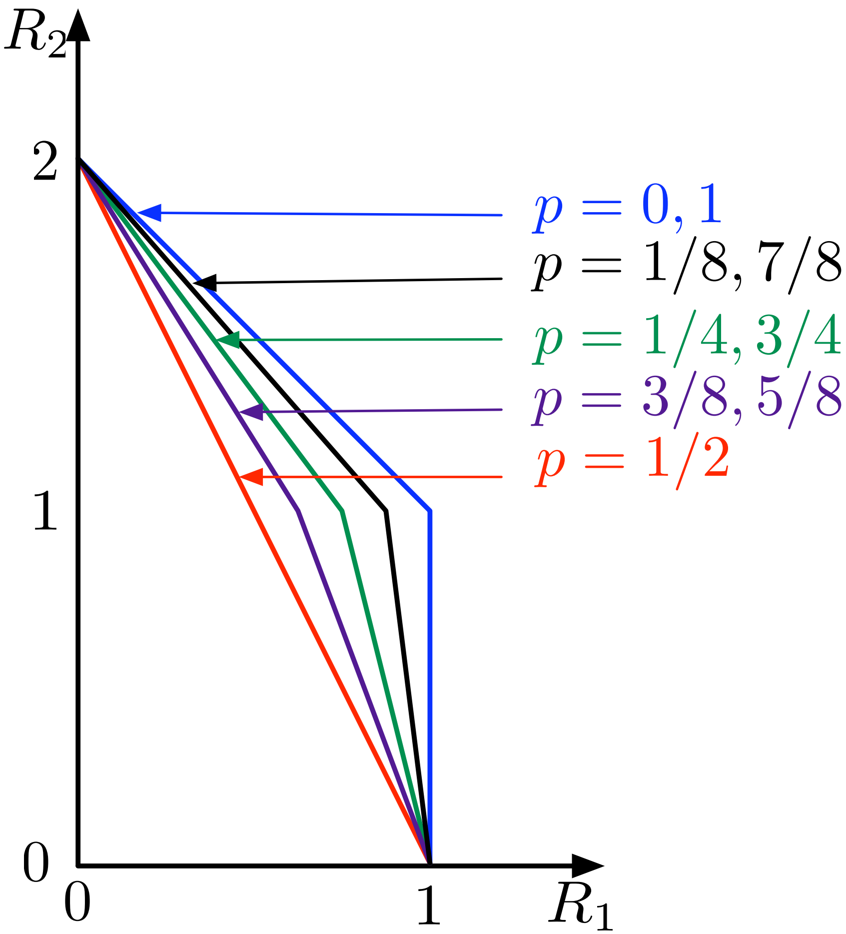

$$\begin{aligned} \max &\equiv \max\{p, 1-p\} \\ \min &\equiv \min\{p, 1-p\} \end{aligned}$$

Rate-delay and rate-reliability tradeoffs

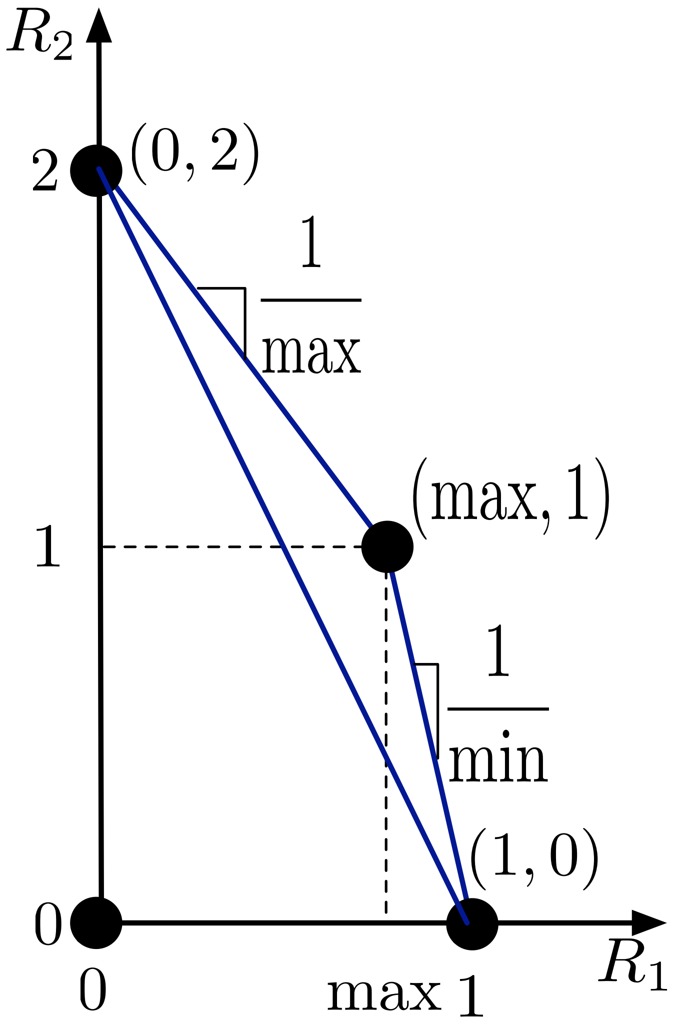

- Increasing rate R_1 (decreasing playback delay) requires decrease in rate R_2 at slope $1/\max > 1$ or $1/\min > 1$
- Time-sharing inner bound achievable by single channel use
- Outside this region requires coding across multiple channel uses (moving away from capacity boundary increases reliability)

How to use three paths to a destination?

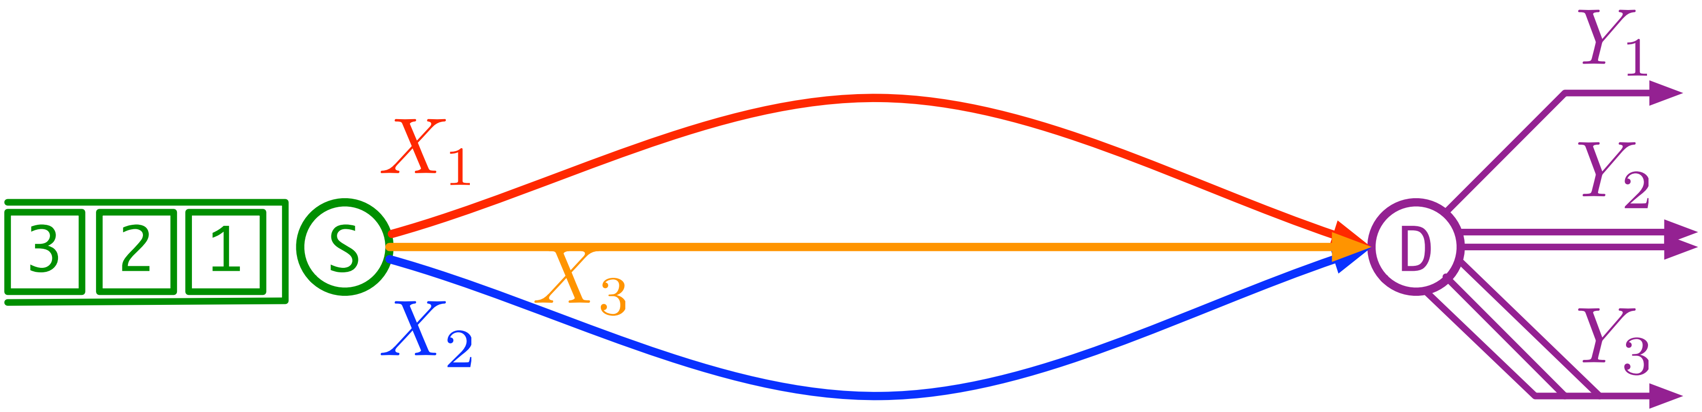

$$
\begin{aligned}
\Pi &= \{123, 132, 213, 231, 312, 321\} \\
p(\pi) &= (p_{123}, p_{132}, p_{213}, p_{231}, p_{312}, p_{321}) \\
Y_1 &= X_{\pi(1)} \\
Y_2 &= (X_{\pi(1)}, X_{\pi(2)}) \\
Y_3 &= (X_{\pi(1)}, X_{\pi(2)}, X_{\pi(3)})
\end{aligned}
$$

Distributions on first packet and first pair of packets

- Six possible permutations for three packets with a specified distribution:

$$\begin{aligned}\Pi &= \{(123),(132),(213),(231),(312),(321)\} \\ p(\pi) &= (p_{123},p_{132},p_{213},p_{231},p_{312},p_{321})\end{aligned}$$

- p_i^1 is the probability that packet i is the first to arrive

$$\mathbf{p}^1 = \begin{pmatrix} p_1^1 \\ p_2^1 \\ p_3^1 \end{pmatrix} = \begin{pmatrix} p(123)+p(132) \\ p(213)+p(231) \\ p(312)+p(321) \end{pmatrix}$$

- p_{ij}^2 is the probability that the *unordered* pair of packets $\{i,j\}$ is the first pair of packets to arrive

$$\mathbf{p}^2 = \begin{pmatrix} p_{12}^2 \\ p_{13}^2 \\ p_{23}^2 \end{pmatrix} = \begin{pmatrix} p(123)+p(213) \\ p(132)+p(312) \\ p(231)+p(321) \end{pmatrix}$$

10 ways of assigning 3 packets to 3 "users", and the associated rates

(a_1, a_2, a_3)	(r_1, r_2, r_3)
$(3, 0, 0)$	$(1, 0, 0)$
$(0, 3, 0)$	$(0, 2, 0)$
$(0, 0, 3)$	$(0, 0, 3)$
$(0, 1, 2)$	$(0, \max\{p_{12}^2 + p_{13}^2, p_{12}^2 + p_{23}^2, p_{13}^2 + p_{23}^2\}, 2)$
$(0, 2, 1)$	$(0, \max\{1 + p_{12}^2, 1 + p_{13}^2, 1 + p_{23}^2\}, 1)$
$(1, 0, 2)$	$(\max\{p_1^1, p_2^1, p_3^1\}, 0, 2)$
$(1, 2, 0)$	$(p_1^1, 1 + p_{23}^2, 0)$ $(p_2^1, 1 + p_{13}^2, 0)$ $(p_3^1, 1 + p_{12}^2, 0)$
$(2, 0, 1)$	$(\max\{p_1^1 + p_2^1, p_1^1 + p_3^1, p_2^1 + p_3^1\}, 0, 1)$
$(2, 1, 0)$	$(p_1^1 + p_2^1, p_{13}^2 + p_{23}^2, 0)$ $(p_1^1 + p_3^1, p_{12}^2 + p_{23}^2, 0)$ $(p_2^1 + p_3^1, p_{12}^2 + p_{13}^2, 0)$
$(1, 1, 1)$	$(p_1^1, \max\{p_{12}^2 + p_{23}^2, p_{13}^2 + p_{23}^2\}, 1)$ $(p_2^1, \max\{p_{12}^2 + p_{13}^2, p_{13}^2 + p_{23}^2\}, 1)$ $(p_3^1, \max\{p_{12}^2 + p_{13}^2, p_{12}^2 + p_{23}^2\}, 1)$

Capacity region is convex hull of 10 rate points

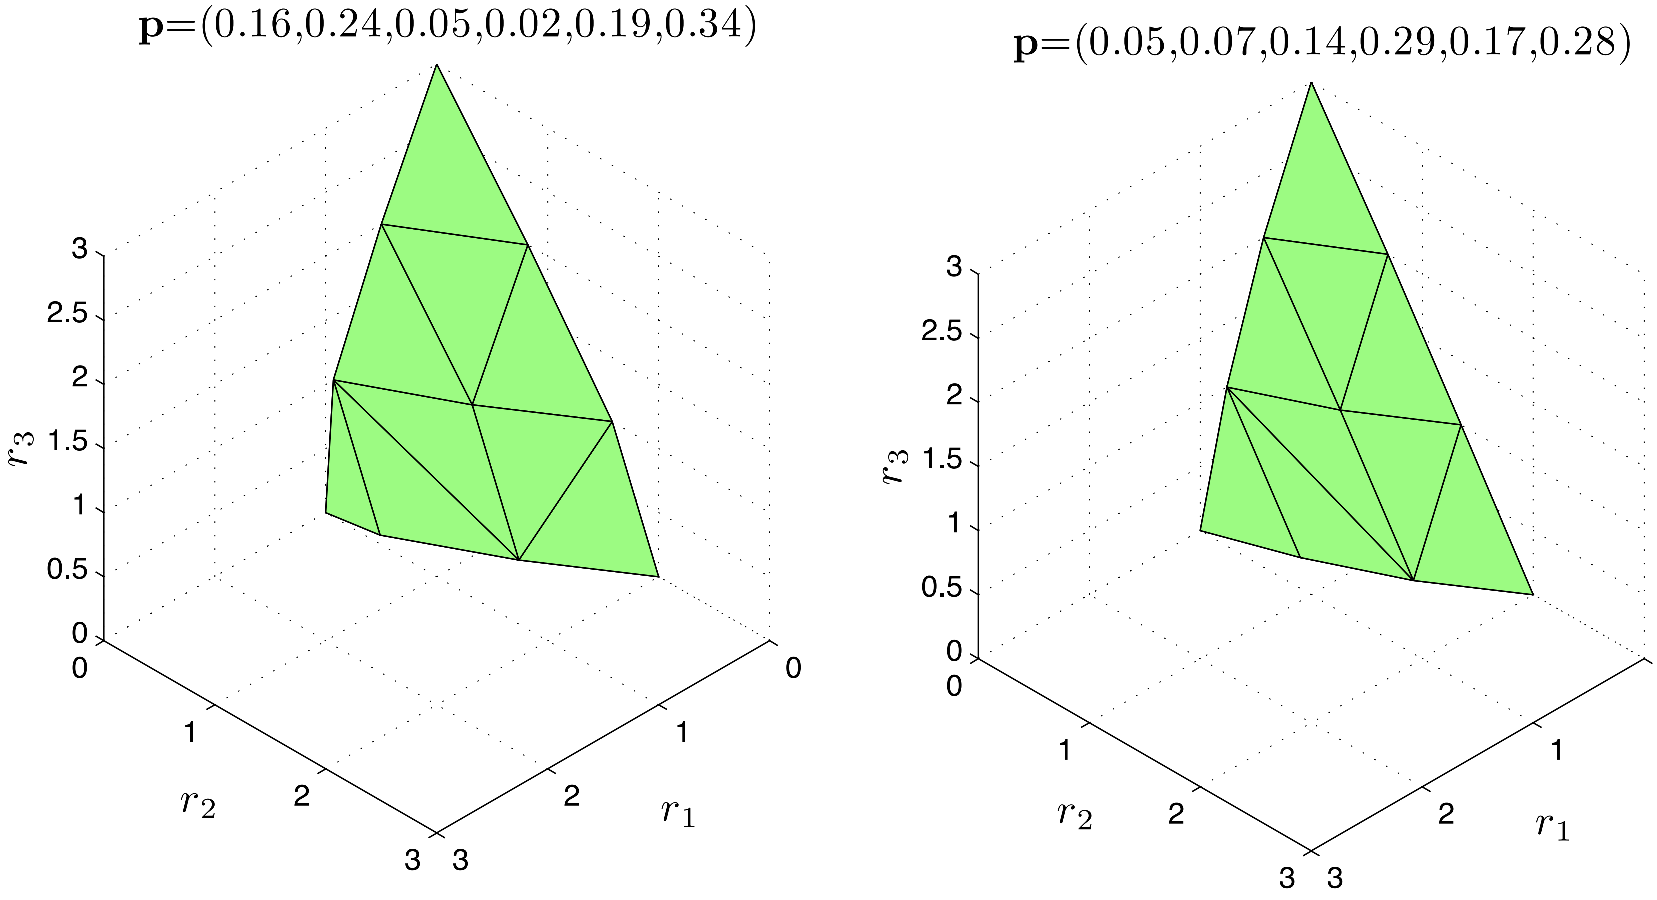

$\mathbf{p}=(0.16,0.24,0.05,0.02,0.19,0.34)$

$\mathbf{p}=(0.05,0.07,0.14,0.29,0.17,0.28)$

Problem statement: sending M packets over a permutation DBC

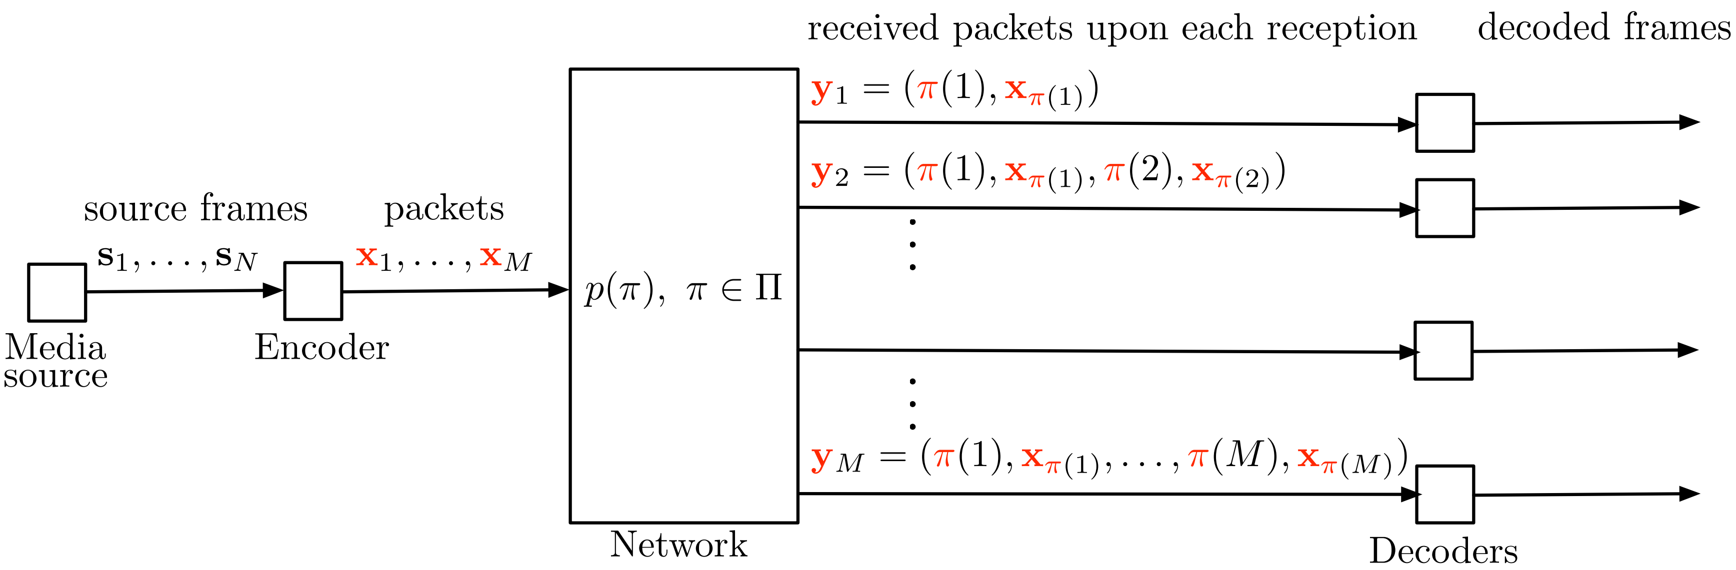

Goal: characterize the region of achievable rate points, $\mathcal{R}$, in terms of the specified permutation distribution $\mathbf{p} = (p(\pi), \pi \in \Pi)$.

Special case: uniform disbn (inner bound for all π)

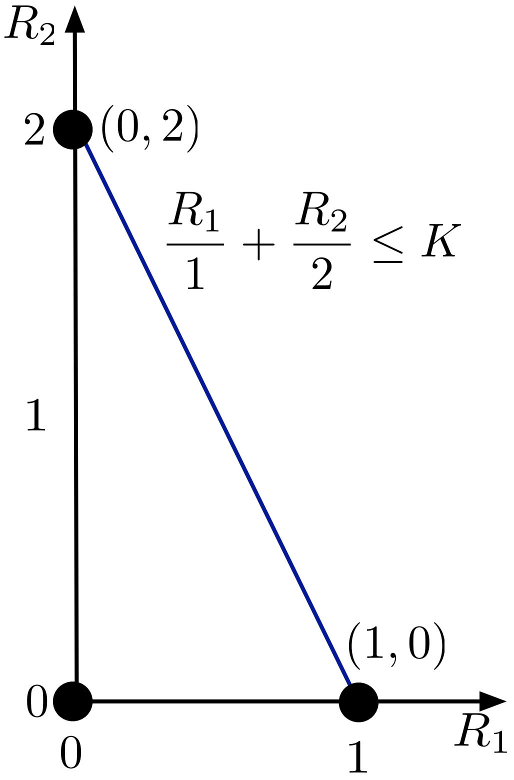

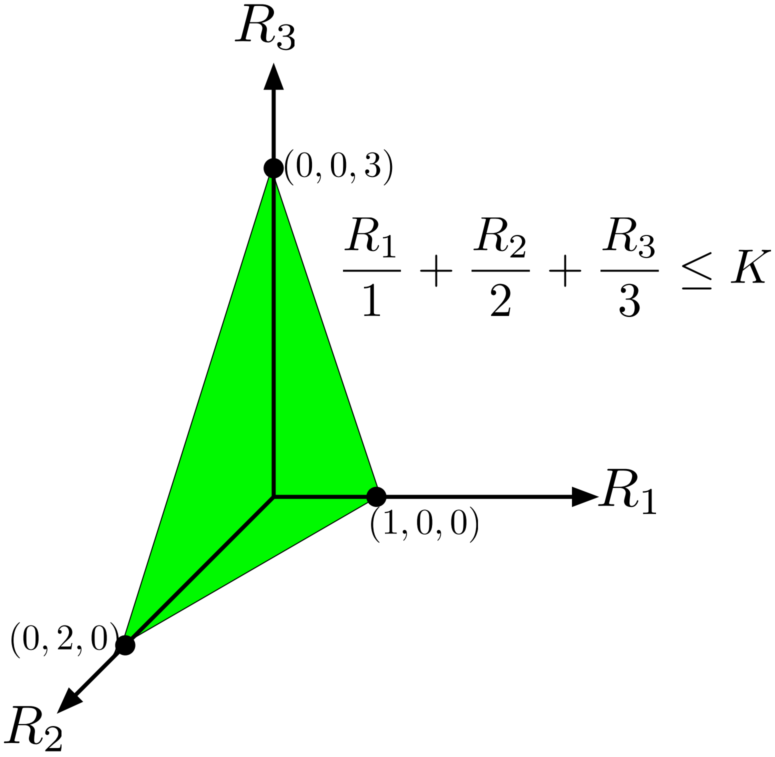

If all permutations are equally likely, the capacity region is specified by the inequality:

$$\sum_{m=1}^{M} \frac{R_m}{m} \le K.$$

Special case: degenerate disbn (outer bound for all π)

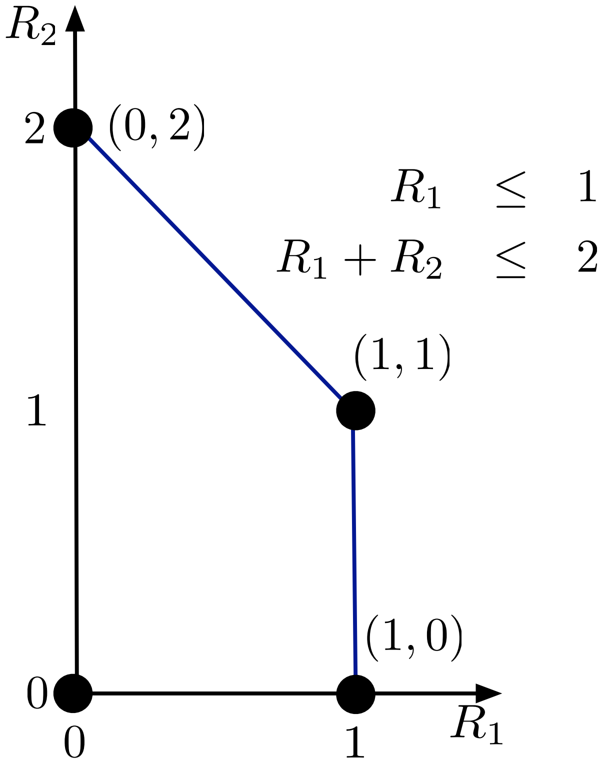

$$\begin{aligned} R_1 &\leq 1 \\ R_1 + R_2 &\leq 2 \end{aligned}$$

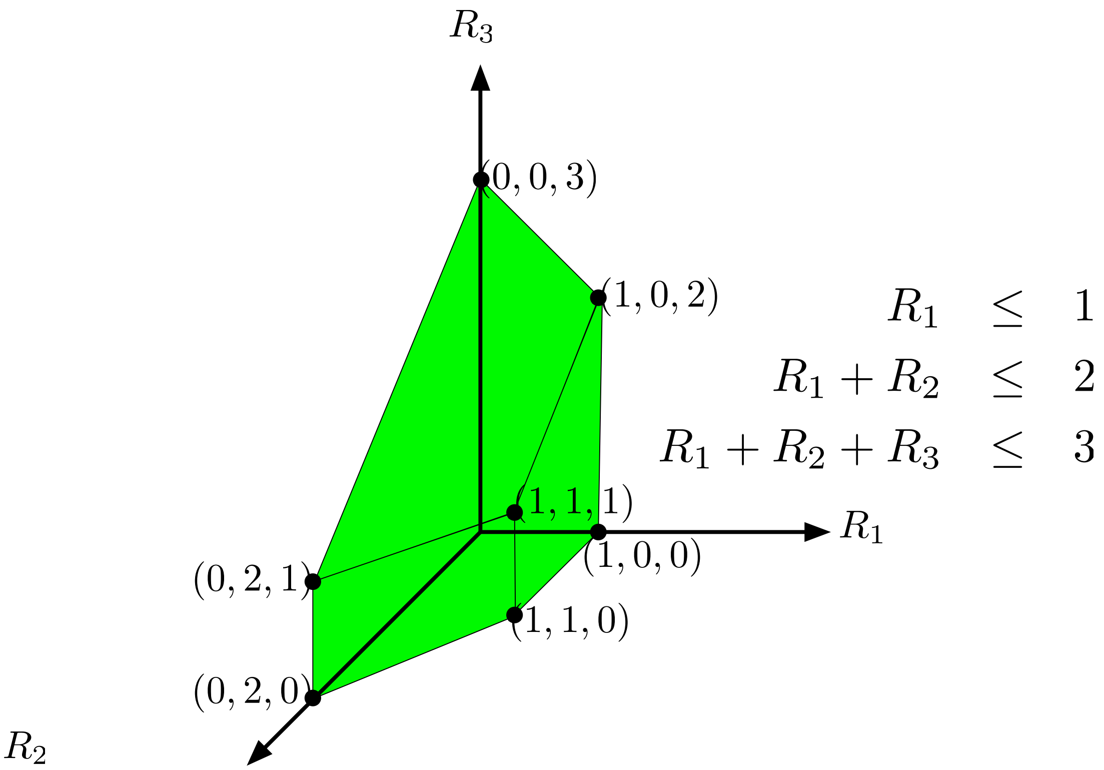

$$\begin{aligned} R_1 &\leq 1 \\ R_1 + R_2 &\leq 2 \\ R_1 + R_2 + R_3 &\leq 3 \end{aligned}$$

If only a single permutation $\pi \in \Pi$ is possible, the capacity region is specified by the M inequalities

$$\sum_{i=1}^{m} R_i \leq mK, \quad m \in [M].$$

Achievability (inner bound): packet assignments

The set of packet assignments, $\mathcal{A}_M$, is the set of all partitions of $[M]$:

$$\mathcal{A}_M \equiv \left\{ \mathbf{A} = (A_1, \ldots, A_M) \;:\; A_m \subseteq [M], \; \bigcup_{m=1}^{M} A_m = [M], \; A_m \cap A_{m'} = \emptyset \right\}.$$

An assignment of packets $A_m \subseteq [M]$ to user m means all of the K bits in each of those packets are used to encode the source bits for user m.

Achievability (inner bound): point to point channels

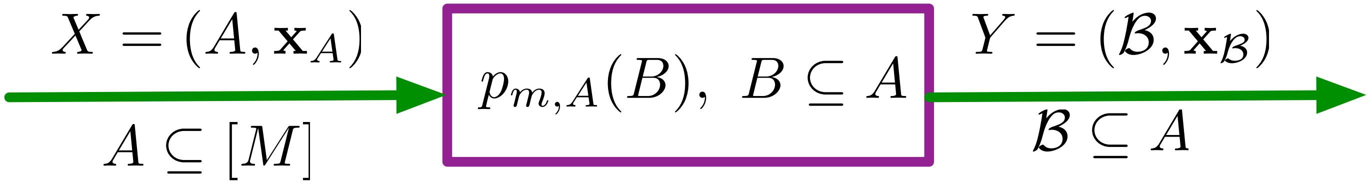

$X = (A, \mathbf{x}_A)$

$A \subseteq [M]$

$p_{m,A}(B),\ B \subseteq A$

$Y = (\mathcal{B}, \mathbf{x}_{\mathcal{B}})$

$\mathcal{B} \subseteq A$

Proposition: $C(A) = K \sum_{B \subseteq A} p_A(B)|B| = K\mathbb{E}_A[|\mathcal{B}|].$

Proof. Let $\mathcal{Q}$ be the collection of disbns $\mathbf{q}$ on the contents of packets $\mathbf{x}_A$.

$$\begin{aligned} C(A) &\equiv \max_{\mathbf{q}\in\mathcal{Q}} H(Y) - H(Y|X) \\ &= \max_{\mathbf{q}\in\mathcal{Q}} (H(\mathcal{B}) + H(\mathbf{x}_{\mathcal{B}}|\mathcal{B})) - (H(\mathcal{B})) \;=\; \max_{\mathbf{q}\in\mathcal{Q}} H(\mathbf{x}_{\mathcal{B}}|\mathcal{B}) \\ &= \max_{\mathbf{q}\in\mathcal{Q}} \sum_{B\subseteq A} p_A(B) H(\mathbf{x}_B|\mathcal{B}=B) \le \sum_{B\subseteq A} p_A(B) K|B| \end{aligned}$$

Equality is achieved by the uniform distribution $\mathbf{q}$.

Achievability theorem (inner bound)

- Fix assignment $\mathbf{A} = (A_1, \ldots, A_M) \in \mathcal{A}_M$: consider M P2P channels:

$$X_m \equiv (A_m, \mathbf{x}_{A_m}) \quad \rightarrow \quad Y_m \equiv (\mathcal{B}_m, \mathbf{x}_{\mathcal{B}_m}), \quad m \in [M]$$

- Achievable rate point for $\mathbf{A}$: $\mathbf{R}(\mathbf{A}) = (R_1(A_1), \ldots, R_M(A_M))$

$$R_m(A_m) = K \sum_{B \in \mathcal{S}(A_m, m)} p_{A_m}(B)|B|, \; m \in [M]$$

- A region of achievable rates is defined by the polytope formed by the convex hull of the rate points for each possible packet assignment:

$$\mathcal{R}_M^{\text{in}} = \text{conv}\left(\{\mathbf{R}(\mathbf{A}), \; \mathbf{A} \in \mathcal{A}_M\}\right)$$

- Non-vertex points in this set are achievable via time-sharing across channel uses

General DBC capacity theorem

- Define $u_m = (x_1, \ldots, x_m)$ and $y_m = (x_{\pi(1)}, \ldots, x_{\pi(m)})$
- General theorem for DBC with M receivers:

$$R_m \leq I(y_m; u_m | u_{m-1}), \; m \in [M].$$

- Specialized to permutation channel with M receivers:

$$R_m \leq \sum_{G \in \mathcal{G}_m} p_m(G) I(\mathbf{x}_G; u_m | u_{m-1}), \; m \in [M]$$

for appropriate support $\mathcal{G}_m$ and distribution $p_m(G)$, defined in terms of $(p(\pi), \pi \in \Pi)$.

DBC capacity theorem as a linear map from mutual information vectors to rate points

Achievable rates (R_1, R_2) are a linear map from a mutual information vector to the rate points:

$$\begin{aligned} R_1 &\leq pI(x_1; u_1) + (1-p)I(x_2; u_1) \\ R_2 &\leq H(x_1, x_2 | u_1) \end{aligned}$$

or

$$\begin{bmatrix} R_1 \\ R_2 \end{bmatrix} = \begin{bmatrix} p & 1-p & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} I(x_1; u_1) \\ I(x_2; u_1) \\ H(x_1, x_2 | u_1) \end{bmatrix}$$

Feasible mutual information vectors don’t form a polytope

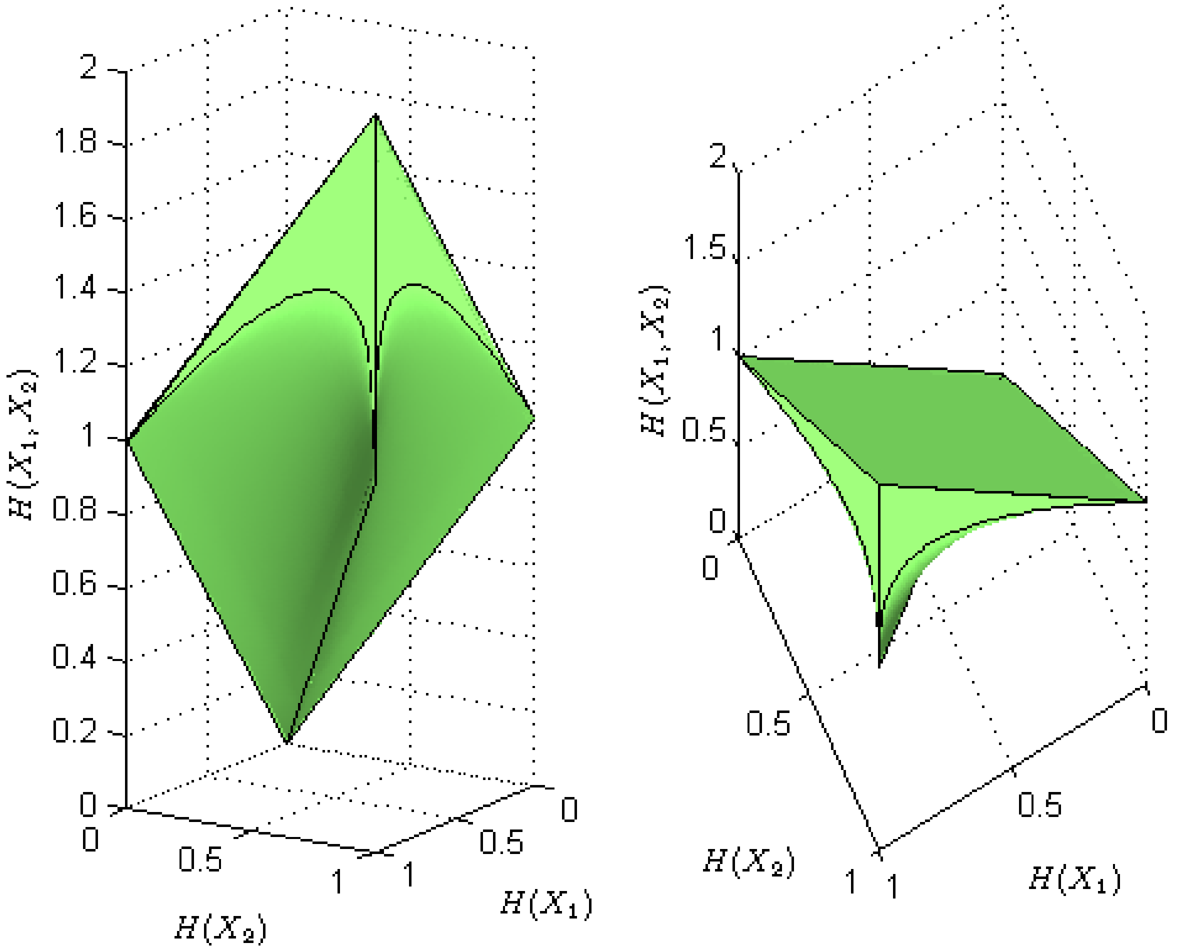

Take the convex hull of the feasible set of mutual information vectors

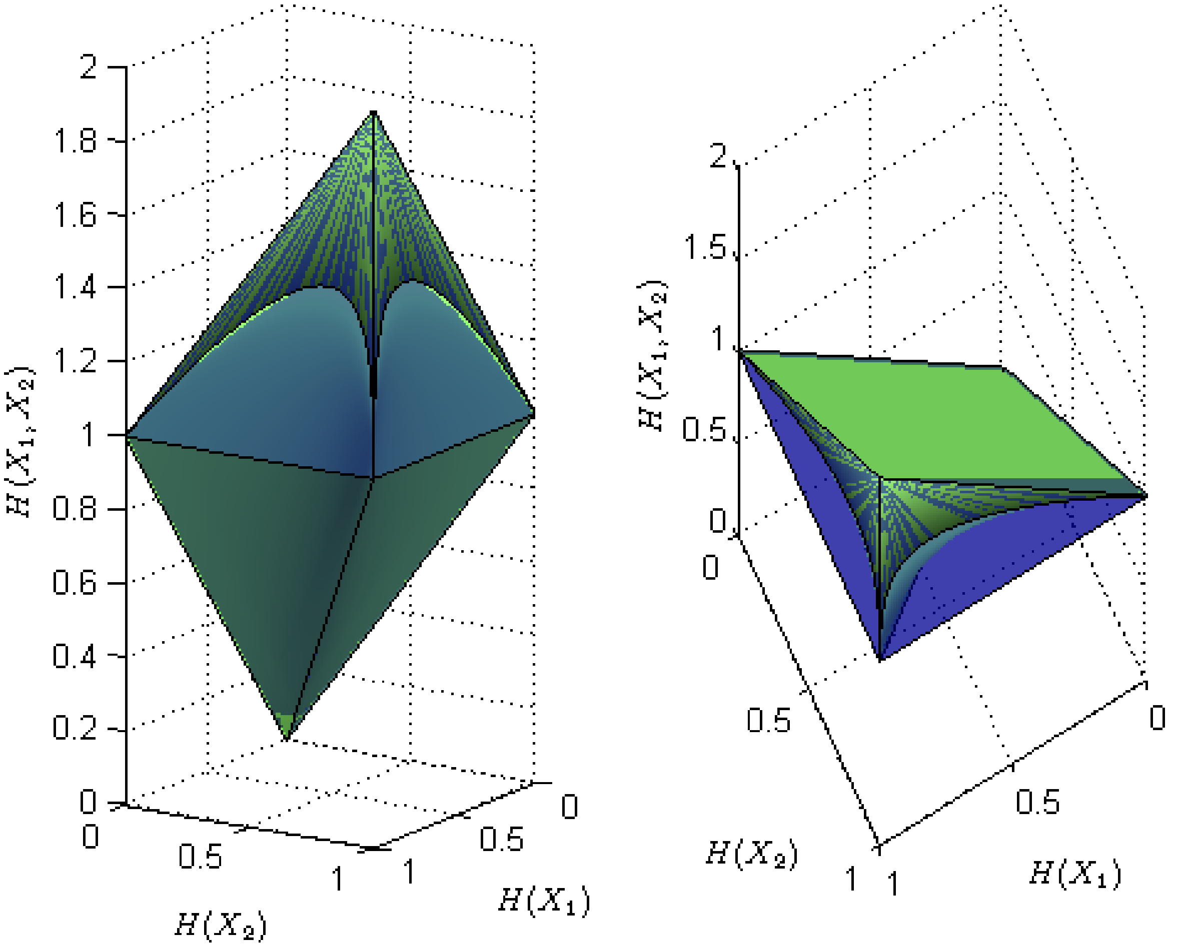

The convex hull of feasible mutual information vectors maps to a lower dimensional convex set of feasible rates under the linear map

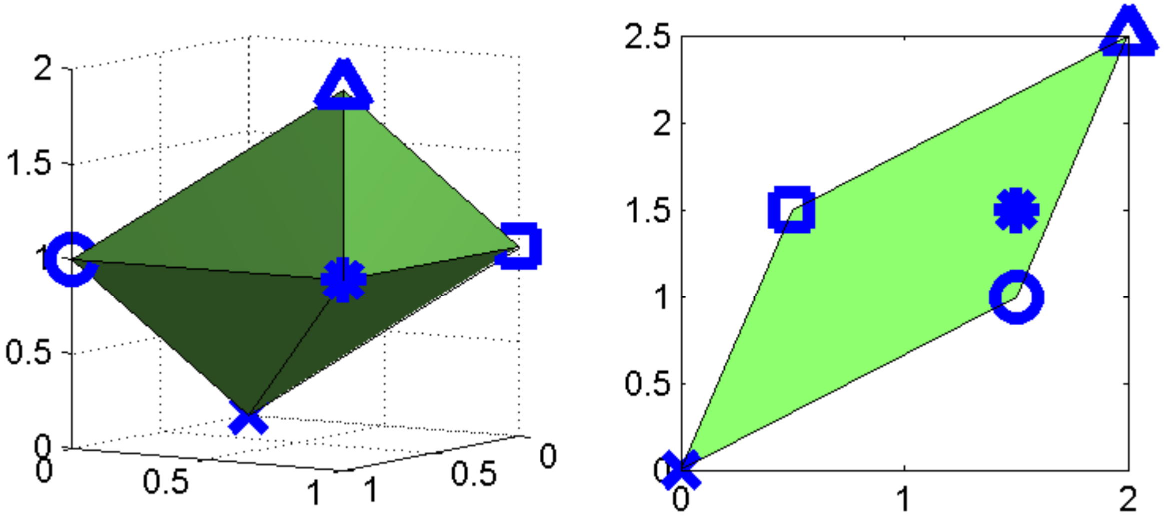

Entropy vectors through the linear map yields achievable rate points

- Convex hull of set of feasible mutual information vectors generated by a set of points $\{\mathbf{I}\}$
- Map each M.I. vector to a rate point: $\mathbf{R} = \mathbf{T}\mathbf{I}$
- Form convex hull of these rate points $\mathcal{R}_M^{\mathrm{out}} = \mathrm{conv}(\{\mathbf{R}\})$

This procedure can be automated for any specified permutation distribution $(p(\pi), \pi \in \Pi)$.

Summary

- **Theorem (inner bound):** $\mathcal{R}_M^{\text{in}} = \text{conv}(\{\mathbf{R}(\mathbf{A})\})$ for $\mathbf{A} \in \mathcal{A}_M$.
- **Theorem (outer bound):** $\mathcal{R}_M^{\text{out}} = \text{conv}(\{\mathbf{R}\})$ for $\mathbf{R} = \mathbf{TI}$.
- **Conjecture:** $\mathcal{R}_M^{\text{in}} = \mathcal{R}_M^{\text{out}}$. Verified for $M = 2, 3$.

Related work

- PET codes (Albanese, Blömer, Edmonds, Luby, Sudan, Trans. IT 1996)
- Fountain codes (Shokrollahi, Trans. IT 2006)
- Rateless codes (Sanghavi, ITW 2007)

Extensions

- More sophisticated measures of throughput and delay are possible, yield explicit T-D tradeoffs for large M by solving associated CoV problem
- Incorporating lost packets
- Unlabeled packets

References

1. John M. Walsh and Steven Weber. “Capacity region of the permutation channel”. *Proceedings of the 46th Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, September 2008.

2. John M. Walsh, Steven Weber, and Ciira wa Maina. “Optimal rate delay tradeoffs and delay mitigating codes for multipath routed and network coded networks”. Submitted in April, 2008 to *IEEE Transactions on Information Theory*.