Secrecy Capacities and Multiterminal Source Coding

Prakash Narayan

Joint work with Imre Csiszár and Chunxuan Ye

Multiterminal Source Coding

The Model

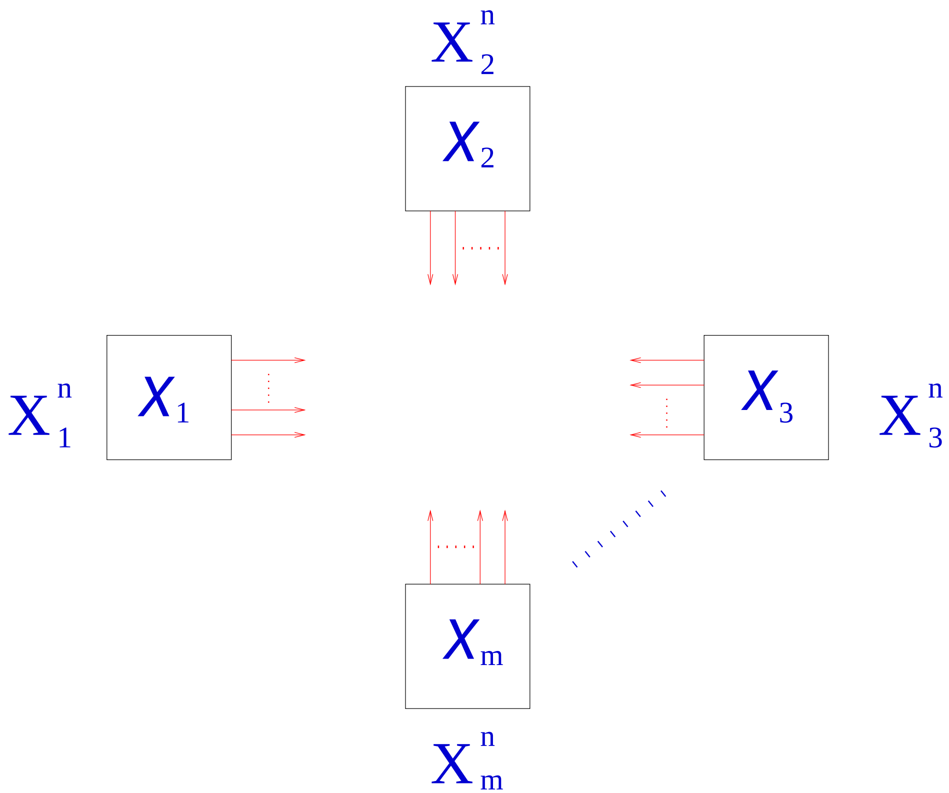

- $m \geq 2$ terminals.
- $X_1, \ldots, X_m$, $m \geq 2$, are rvs with finite alphabets $\mathcal{X}_1, \ldots, \mathcal{X}_m$.
- Consider a discrete memoryless multiple source with components $X_1^n = (X_{11}, \ldots, X_{1n}), \ldots, X_m^n = (X_{m1}, \ldots, X_{mn})$.
- Terminal $\mathcal{X}_i$ observes the component $X_i^n = (X_{i1}, \ldots, X_{in})$.

The Model

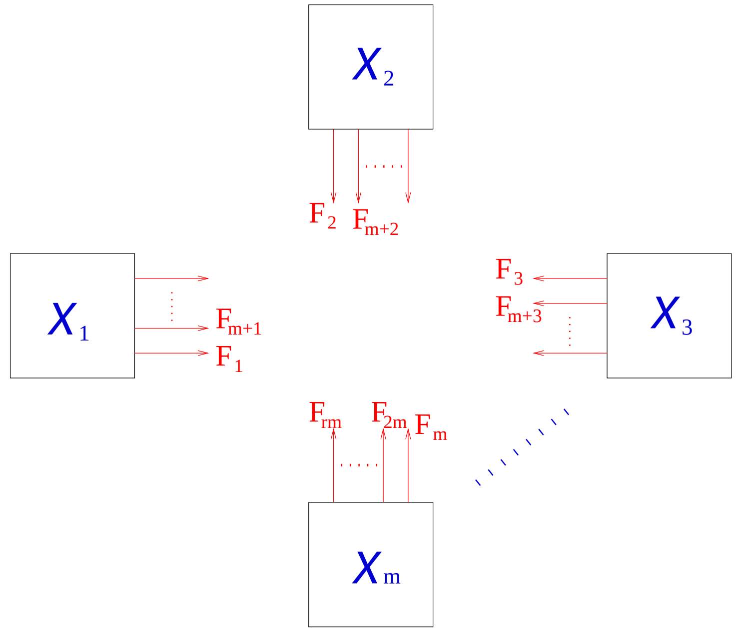

- The terminals are allowed to communicate over a *noiseless* channel, possibly interactively in several rounds.
- All the transmissions are observed by all the terminals.
- No rate constraints on the communication.
- Assume w.l.o.g that transmissions occur in consecutive time slots in r rounds.
- Communication depicted by rvs $\mathbf{F} \triangleq F_1, \ldots F_{rm}$, where
  - ∗ F_ν = transmission in time slot ν by terminal $i \equiv \nu \mod m$.
  - ∗ F_ν is a function of X_i^n and $(F_1, \ldots, F_{\nu-1})$.

Communication for Omniscience

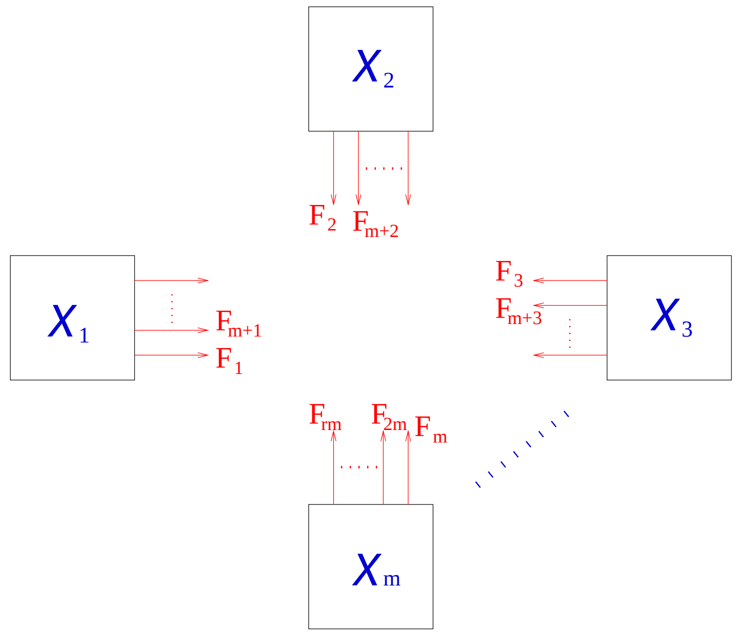

- Each terminal wishes to become "omniscient," i.e., recover $(X_1^n, \ldots, X_m^n)$ with probability $\geq 1 - \varepsilon$.

- What is the smallest achievable rate of communication for omniscience (CO-rate), $\lim_n \frac{1}{n} H(F_1, \ldots, F_{rm})$?

Minimum Communication for Omniscience

Proposition [I. Csiszár - P. N., '02]: The smallest achievable CO-rate, $\lim_n \frac{1}{n} H(F_1^{(n)}, \ldots, F_{rm}^{(n)})$, which enables $(X_1^n, \ldots, X_m^n)$ to be ε_n-recoverable at all the terminals with communication $(F_1^{(n)}, \ldots, F_{rm}^{(n)})$ (with the number of rounds possibly depending on n), with $\varepsilon_n \to 0$, is

$$R_{min} = \min_{(R_1, \ldots, R_m) \in \mathcal{R}_{SW}} \sum_{i=1}^{m} R_i,$$

where $\mathcal{R}_{SW} = \left\{ (R_1', \cdots, R_m') : \sum_{i \in B} R_i' \geq H(X_B | X_{B^c}), \quad B \subset \{1, \ldots, m\} \right\}$.

Remark: The region $\mathcal{R}_{SW}$, if stated for *all* $B \subseteq \{1, \ldots, m\}$, gives the achievable rate region for the multiterminal version of the Slepian-Wolf source coding theorem.

Case: $m = 2$; $R_{min} = H(X_1|X_2) + H(X_2|X_1)$.

Communication for Omniscience

Proof of Proposition: The proposition is a source coding theorem of the "Slepian-Wolf" type, with the additional element that interactive communication is not a priori excluded.

Achievability: *Straightforward extension of the multiterminal Slepian-Wolf source coding theorem; the CO-rates can be achieved with noninteractive communication.*

Converse: *Nontrivial; consequence of the following "Main Lemma."*

Common Randomness

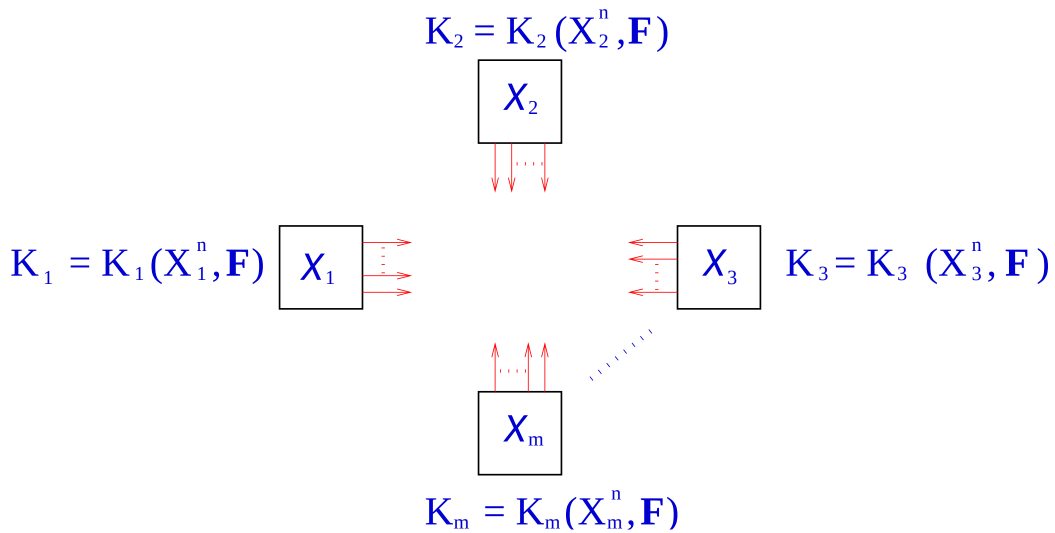

Common Randomness (CR): A function K of $(X_1^n, \cdots, X_m^n)$ is ε-CR, achievable with communication **F**, if

$$Pr\{K = K_1 = \cdots = K_m\} \geq 1 - \varepsilon.$$

Thus, CR consists of random variables generated by different terminals, based on

– local measurements or observations

– transmissions or exchanges of information

such that the random variables agree with probability $\cong 1$.

Main Lemma

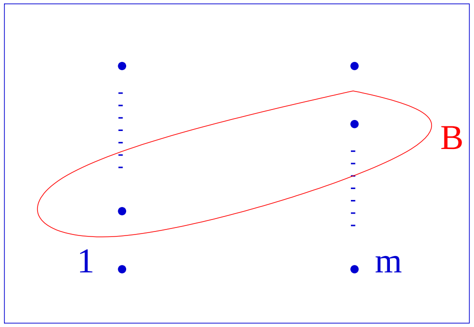

Lemma [I. Csiszár - P. N., '02]: If K is ε-CR for the terminals $\mathcal{X}_1, \cdots, \mathcal{X}_m$, achievable with communication $\mathbf{F} = (F_1, \cdots, F_{rm})$, then

$$\frac{1}{n} H(K|\mathbf{F}) = H(X_1, \cdots, X_m) - \sum_{i=1}^{m} R_i + \frac{m(\varepsilon \log |\mathcal{K}| + 1)}{n}$$

for some numbers $(R_1, \cdots, R_m) \in \mathcal{R}_{SW}$ where

$$\mathcal{R}_{SW} = \left\{ (R_1^{'}, \cdots, R_m^{'}) : \sum_{i \in B} R_i^{'} \geq H(X_B | X_{B^c}), \quad B \subset \{1, \ldots, m\} \right\}.$$

Remark: Decomposition of total joint entropy $H(X_1, \ldots, X_m)$ into the normalized conditional entropy of any achievable ε-CR conditioned on the communication with which it is achieved, and a sum of rates which satisfy the SW conditions.

Secrecy Capacities

The General Model

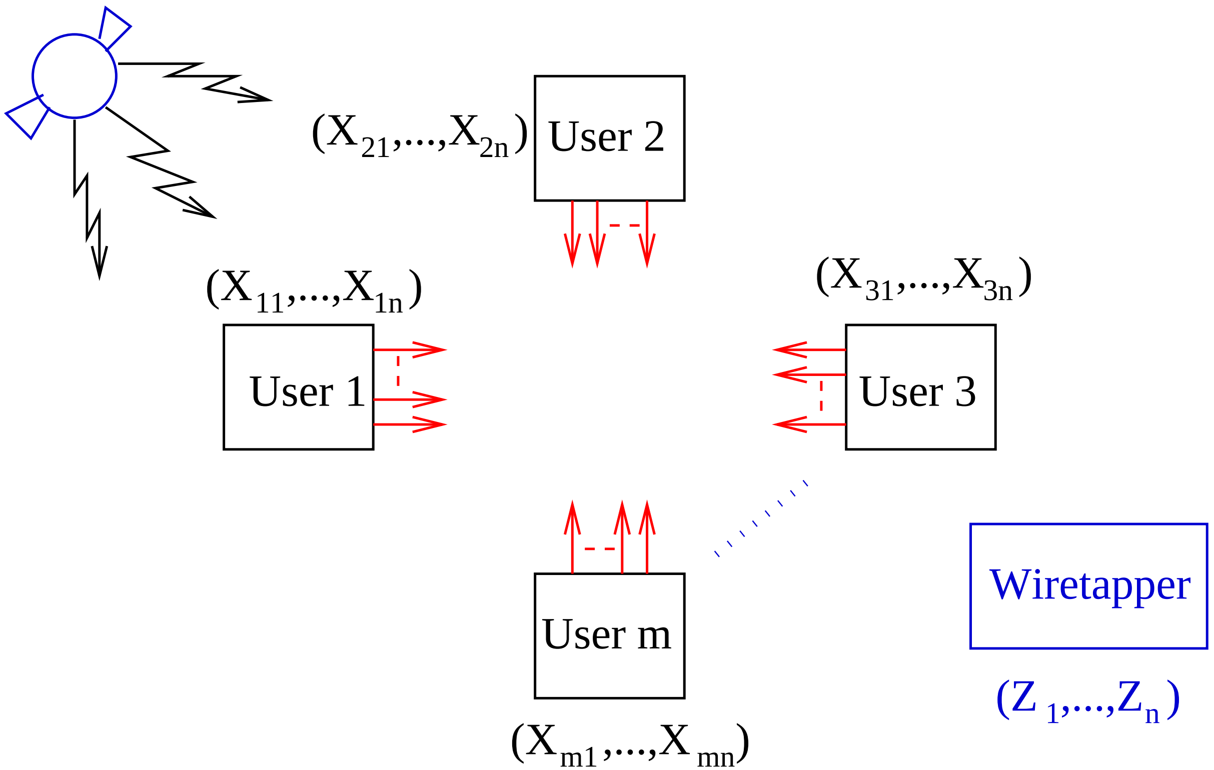

The user terminals wish to generate CR which is effectively concealed from an eavesdropper with access to the public interterminal communication or from a wiretapper.

Secret Key

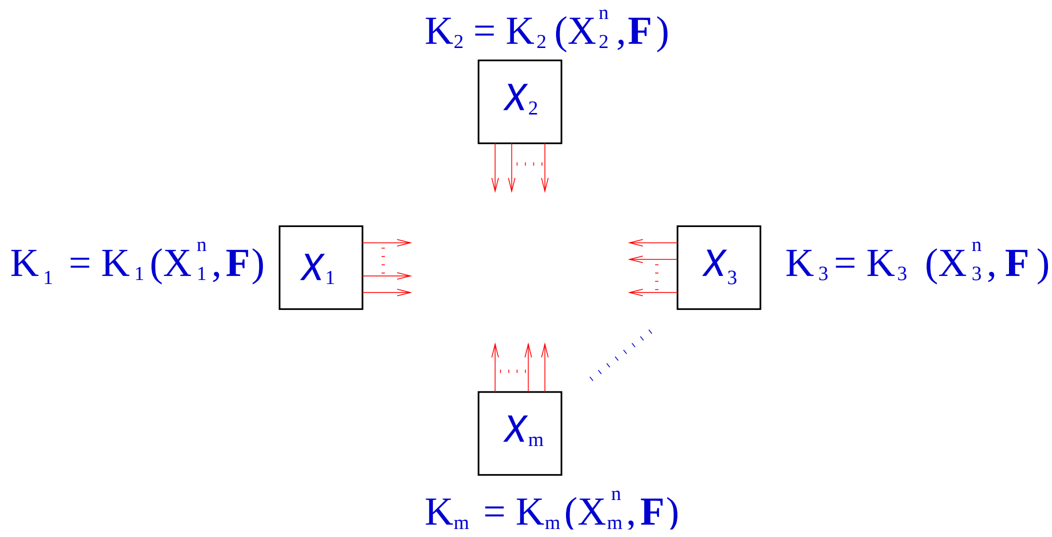

Secret Key (SK): A function K of $(X_1^n, \cdots, X_m^n)$ is an ε-SK, achievable with communication $\mathbf{F}$, if

- $Pr\{K = K_1 = \cdots = K_m\} \geq 1 - \varepsilon$ (“ε-common randomness”)
- $\frac{1}{n} I(K \wedge \mathbf{F}) \leq \varepsilon$ (“secrecy”)
- $\frac{1}{n} H(K) \geq \frac{1}{n} \log |\mathcal{K}| - \varepsilon$ (“uniformity”)

where $\mathcal{K}$ = set of all possible values of K.

Thus, a secret key is effectively concealed from an eavesdropper with access to $\mathbf{F}$, and is nearly uniformly distributed.

Secret Key Capacity

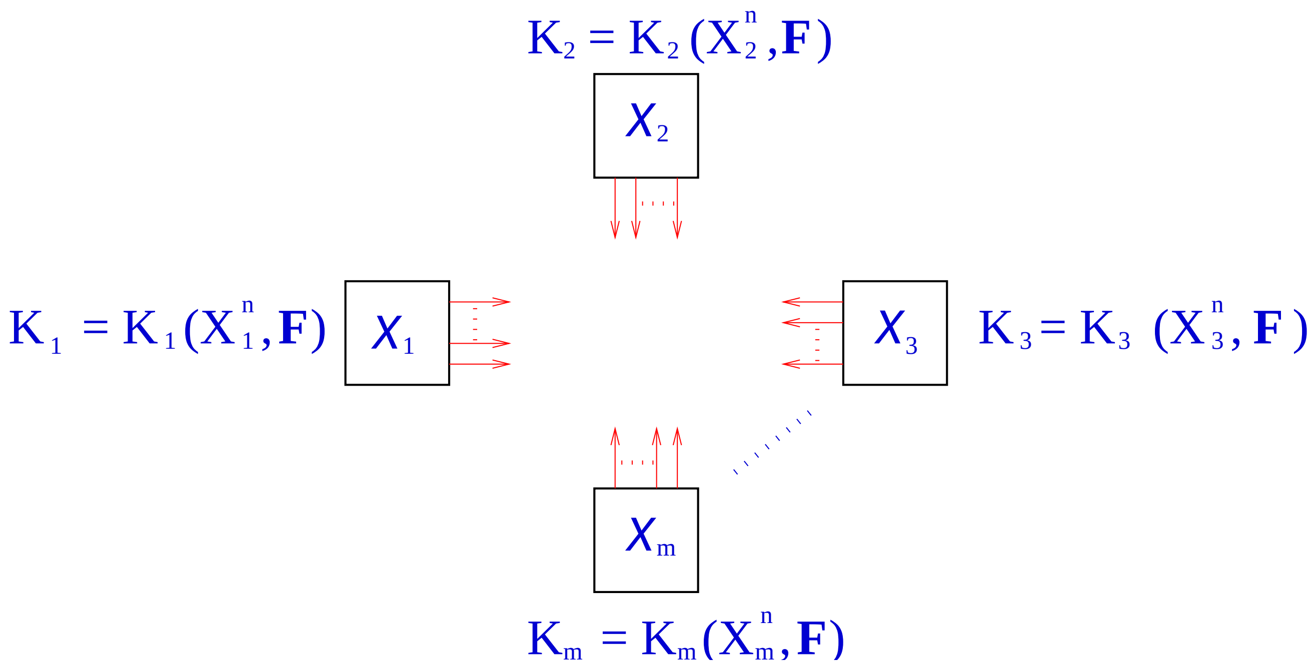

- Achievable SK-rate: The (entropy) rate of such a SK, achievable with suitable communication (with the number of rounds possibly depending on n).
- SK-capacity C_{SK} = largest achievable SK-rate.

Some Recent Related Work

- Maurer 1990, 1991, 1993, 1994, $\cdots$
- Ahlswede-Csiszár 1993, 1994, 1998, $\cdots$
- Bennett, Brassard, Crépeau, Maurer 1995.
- Csiszár 1996.
- Maurer - Wolf 1997, 2003, $\cdots$
- Venkatesan - Anantharam 1995, 1997, 1998, 2000, $\cdots$
- Csiszár - Narayan 2000.
- Renner-Wolf 2003.

$\vdots$

$\vdots$

The Connection

Special Case: Two Users

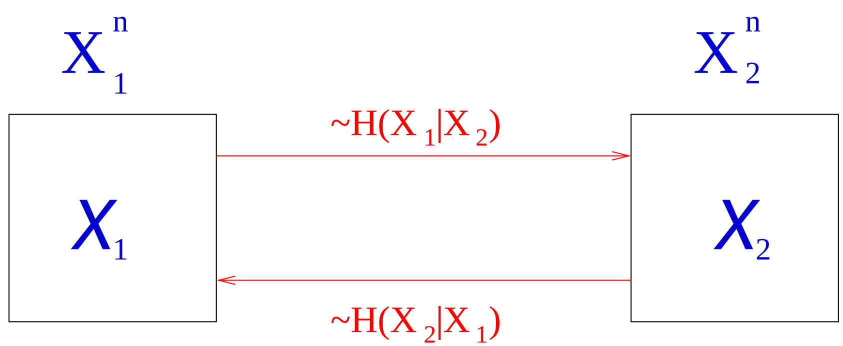

Observation

$$
\begin{aligned}
C_{SK} &= I(X_1 \wedge X_2) \qquad \text{[Maurer 1993, Ahlswede - Csiszár 1993]} \\
&= H(X_1, X_2) - [H(X_1|X_2) + H(X_2|X_1)] \\
&= \text{Total rate of shared } CR - \text{ Smallest achievable} \\
&\quad \text{CO-rate } (R_{min}).
\end{aligned}
$$

The Main Result

- SK-capacity [I. Csiszár - P. N., '02]:

$$C_{SK} = H(X_1, \ldots, X_m) - \text{Smallest achievable CO-rate, } R_{min}\text{, i.e., smallest rate of communication which enables each terminal to reconstruct all the } m \text{ components of the multiple source.}$$

- A single-letter characterization of R_{min}, thus, leads to the same for C_{SK}.

Remark: The source coding problem of determining the smallest achievable CO-rate R_{min} *does not involve any secrecy constraints.*

Secret Key Capacity

Theorem [I. Csiszár - P. N., '02]: The SK-capacity C_{SK} for a set of terminals $\{1, \ldots, m\}$ equals

$$C_{SK} = H(X_1, \ldots, X_m) - R_{min},$$

and can be achieved with noninteractive communication.

Proof: Converse: *From Main Lemma.*

Idea of achievability proof: *If L represents ε-CR for the set of terminals, achievable with communication $\mathbf{F}$ for some block length n, then $\frac{1}{n}H(L|\mathbf{F})$ is an achievable SK-rate if ε is small. With $L \cong (X_1^n, \ldots, X_m^n)$, we have*

$$\frac{1}{n}H(L|\mathbf{F}) \cong H(X_1, \ldots, X_m) - \frac{1}{n}H(\mathbf{F}).$$

Remark: The SK-capacity is not increased by randomization at the terminals.

Case: $m = 2$; $C_{SK} = I(X_1 \wedge X_2)$.

Example

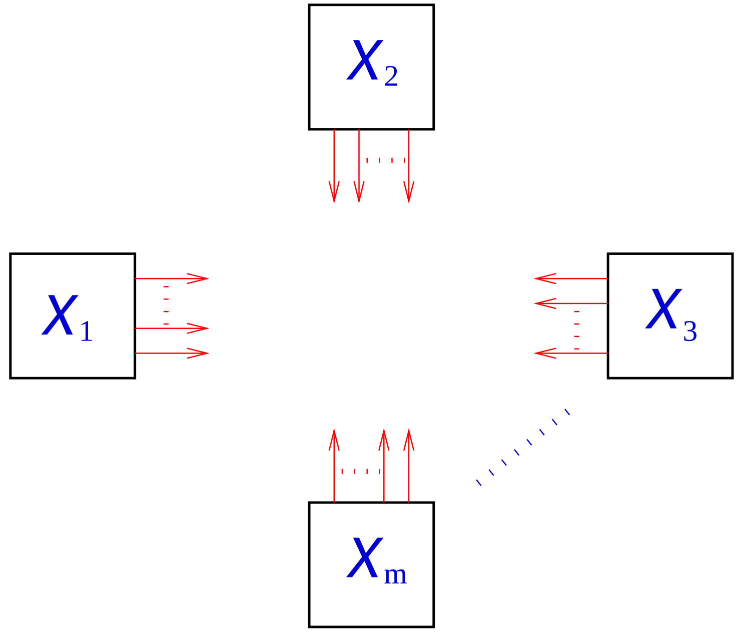

[I. Csiszár - P. N.,'03]:

- $X_1, \cdots, X_{m-1}$ are $\{0,1\}$-valued, mutually independent, $(\frac{1}{2}, \frac{1}{2})$ rvs, and

$$X_{mt} = X_{1t} + \cdots + X_{(m-1)t} \; mod \; 2, \qquad t \geq 1.$$

- Total rate of shared CR$=H(X_1, \ldots, X_m) = H(X_1, \ldots, X_{m-1}) = m-1$ bits.
- $R_{min} = \ldots = \frac{m(m-2)}{m-1}$ bits
- $C_{SK} = (m-1) - \frac{m(m-2)}{m-1} = \frac{1}{m-1}$ bit.

Example – Scheme for Achievability

- **Claim**: 1 bit of perfect SK (i.e., with $\varepsilon = 0$) is achievable with observation length $n = m - 1$.
- *Scheme with noninteractive communication:*

  - Let $n = m - 1$.

  - For $i = 1, \cdots, m-1$, $\mathcal{X}_i$ transmits $F_i = f_i(X_i^n) =$ block X_i^n excluding X_{ii}.

  - $\mathcal{X}_m$ transmits $F_m = f_m(X_m^n) = (X_{m1} + X_{m2} \; mod \; 2, \; X_{m1} + X_{m3} \; mod \; 2, \cdots, X_{m1} + X_{mn} \; mod \; 2)$.
- $\mathcal{X}_1, \cdots, \mathcal{X}_m$ all recover $(X_1^n, \cdots, X_m^n)$. (Omniscience)
- In particular, X_{11} is independent of $\mathbf{F} = (F_1, \cdots, F_m)$.
- X_{11} is an achievable perfect SK, so $C_{SK} \geq \frac{1}{m-1} H(X_{11}) = \frac{1}{m-1}$ bit.

Eavesdropper with Wiretapped Side Information

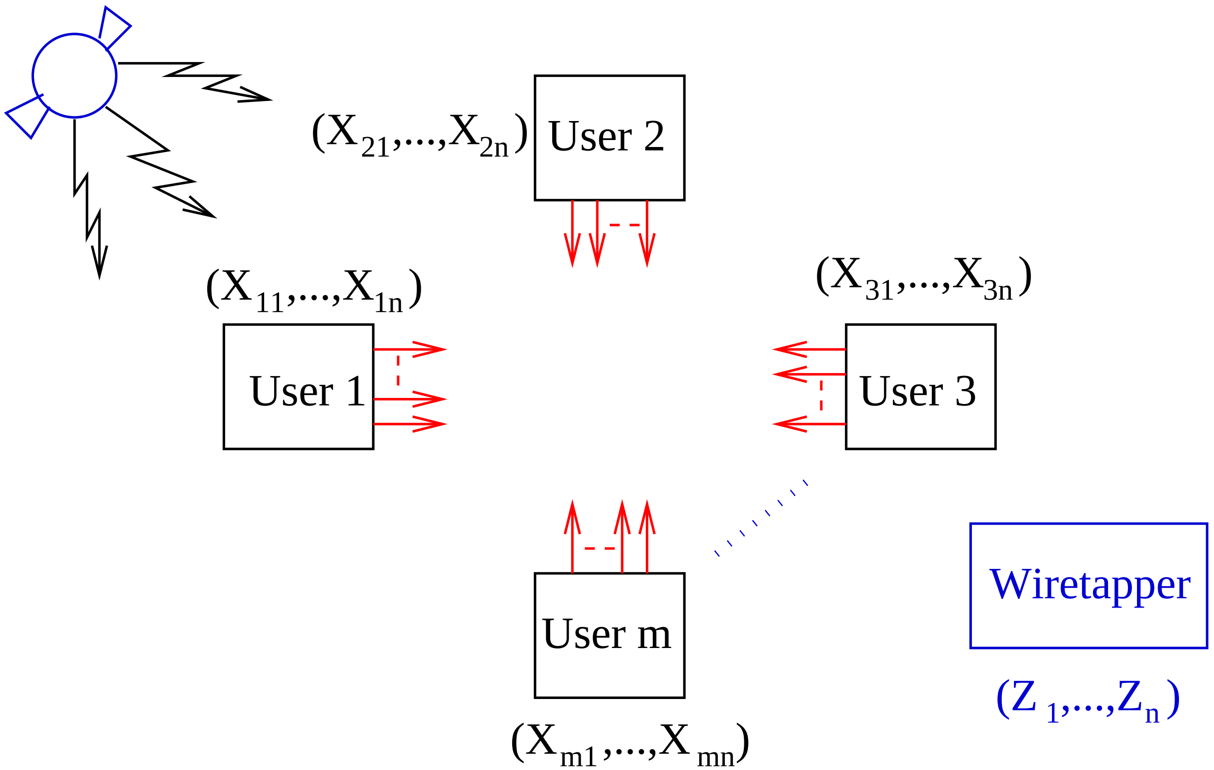

- The secrecy requirement now becomes

$$\frac{1}{n} I(K \wedge \mathbf{F}, Z^n) \leq \varepsilon.$$

- General problem of determining the "Wiretap Secret Key" capacity, C_{WSK}, remains unsolved.

Wiretapping of Noisy User Sources

The eavesdropper can wiretap noisy versions of some or all of the components of the underlying multiple source. Formally,

$$\Pr\{Z_1 = z_1, \ldots, Z_m = z_m | X_1 = x_1, \ldots, X_m = x_m\} = \prod_{i=1}^{m} \Pr\{Z_i = z_i | X_i = x_i\}.$$

Theorem [I. Csiszár - P. N., '03]: The WSK-capacity for a set of terminals $\{1, \ldots, m\}$ equals

$$\begin{aligned} C_{WSK} &= H(X_1, \ldots, X_m, Z_1, \ldots, Z_m) - \text{“Revealed” entropy } H(Z_1, \ldots, Z_m) \\ &\quad - \text{Smallest achievable CO-rate for user terminals} \\ &\qquad \text{when they additionally know } (Z_1, \ldots, Z_m) \\ &= H(X_1, \ldots, X_m | Z_1, \ldots, Z_m) - R_{min}(Z_1, \ldots, Z_m), \end{aligned}$$

provided that randomization is permitted at the user terminals.

Case: $m = 2$; $C_{WSK} = I(X_1 \wedge X_2 | Z_1, Z_2)$.

A Few Variants

Secret Key Capacity with Helpers

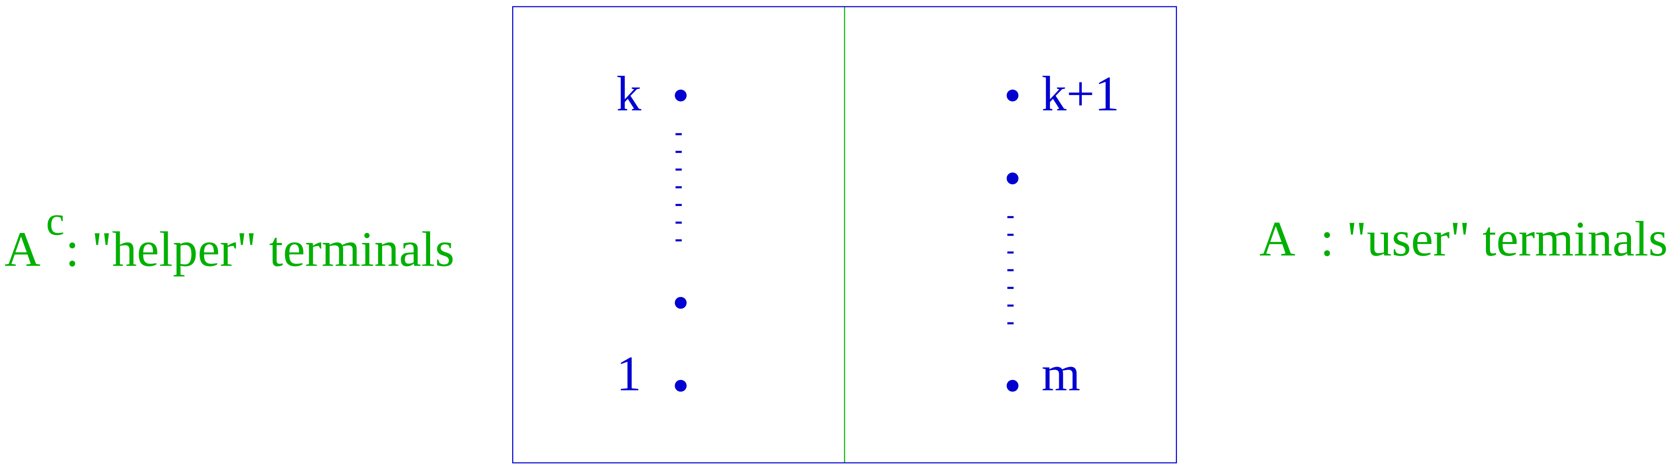

Theorem [I. Csiszár - P. N.,'02]: The SK-capacity for the terminals in A, with the terminals in A^c as helpers, is

$$C_{SK}(A) = H(X_1, \ldots X_m) - \text{Smallest achievable CO-rate for user terminals in } A$$
$$= H(X_1, \ldots X_m) - R_{min}(A).$$

Case: $m = 3$, $A = \{2, 3\}$, $A^c = \{1\}$; $C_{SK}(A) = \min\{I(X_1, X_2 \wedge X_3),\ I(X_1, X_3 \wedge X_2)\}$.

Private Key Capacity

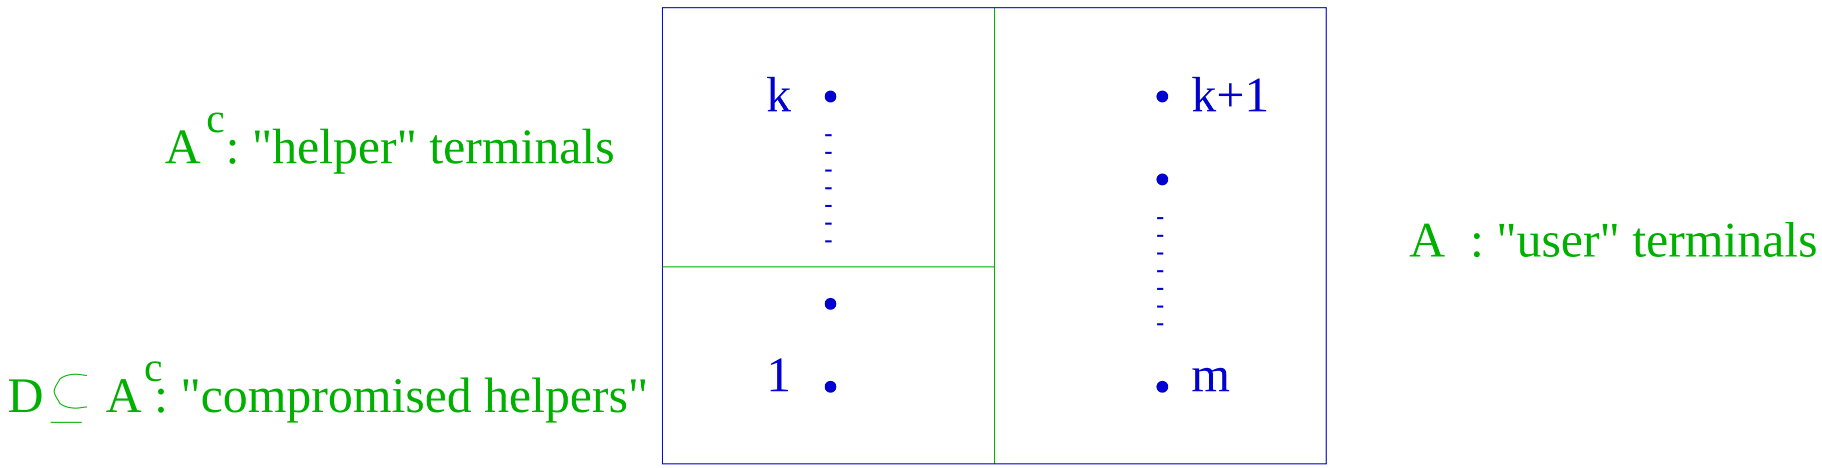

Theorem [I. Csiszár - P. N.,'02]: The PK-capacity for the terminals in A, with privacy from the set of wiretapped helper terminals $D \subseteq A^c$, is

$$C_{PK}(A|D) = H(X_1, \ldots, X_m) - \text{"Revealed" entropy } H(\{X_i,\ i \in D\})$$

$$- \text{ Smallest achievable CO-rate for user terminals in } A \text{ when they additionally know } \{X_i,\ i \in D\}$$

$$= H(X_1, \ldots, X_m | \{X_i,\ i \in D\}) - R_{min}(A|D).$$

Case: $m = 3,\ A = \{2, 3\},\ A^c = D = \{1\};\ C_{PK}(A|D) = I(X_2 \wedge X_3 | X_1)$.

Example

Markov Chain on a Tree [I. Csiszár - P. N.,'03]

- A tree with vertex set $\{1, \cdots, m\}$, i.e., a connected graph G containing no circuits.

- For $(i, j) \in$ edge set $E(G)$ of G, let

$$B(i \leftarrow j) \triangleq \text{set of all vertices connected with } j \text{ by a path containing the edge } (i, j).$$

- The random variables $X_1, \cdots, X_m$ form a *Markov chain on the tree* G if for each $(i, j) \in E(G)$, the conditional pmf of X_j given $\{X_l, l \in B(i \leftarrow j)\}$ depends only on X_i.

- If G is a chain, then $X_1, \cdots, X_m$ form a (standard) Markov chain.

Markov Chain on a Tree

- $C_{SK} = \min_{(i,j)\in E(G)} I(X_i \wedge X_j)$.

- When an eavesdropper wiretaps $Z_1, \cdots, Z_m$ which are noisy versions of $X_1, \cdots, X_m$,

$$C_{WSK} = \min_{(i,j)\in E(G)} I(X_i \wedge X_j | Z_1, \cdots, Z_m).$$

- $C_{SK}(A) = \min_{(i,j)\in E(G(A))} I(X_i \wedge X_j)$,
  where $G(A)$ is the smallest subtree of G whose vertex set contains A.

- $C_{PK}(A|D) = \min_{(i,j)\in E(G(A))} I(X_i \wedge X_j | \{X_l, l \in D\})$.

Multiple Levels of Secrecy

Simultaneous Generation of Multiple Keys

- Simultaneous generation of *multiple* keys
  - by different groups of terminals (with possible overlaps),
  - with protection from prespecified terminals as also from an eavesdropper;
  - at the outset of operations.
- Useful, for instance, when some terminals are disabled or cease to be authorized, and their keys are compromised.

Two Private Keys for Three Terminals

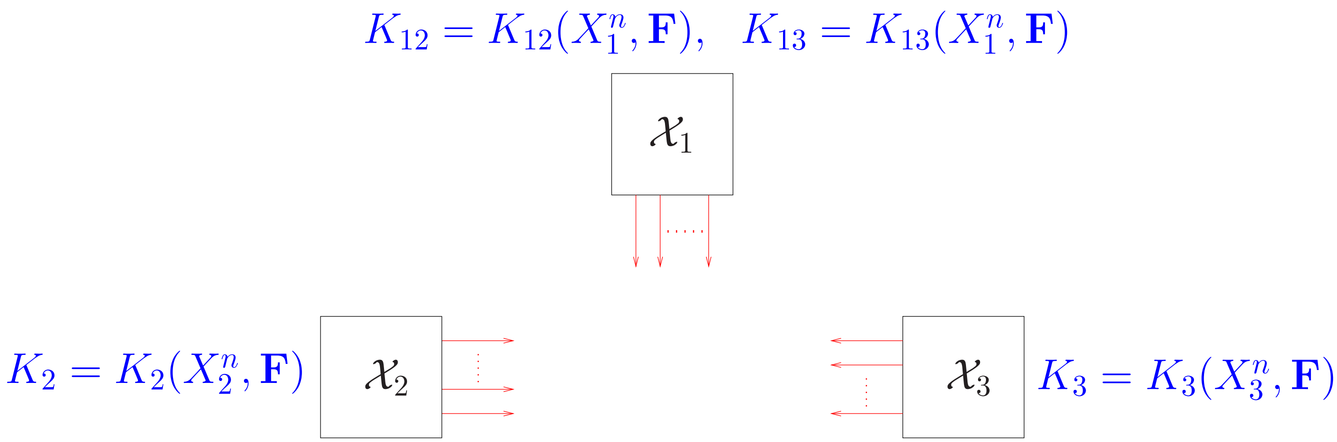

Private Keys for $(\mathcal{X}_1, \mathcal{X}_2)$ and $(\mathcal{X}_1, \mathcal{X}_3)$

- $Pr\{K_{12} = K_2\} \geq 1 - \varepsilon, \quad Pr\{K_{13} = K_3\} \geq 1 - \varepsilon$ ("ε-common randomness")
- $\frac{1}{n} I(K_{12} \wedge \mathbf{F}, X_3^n) \leq \varepsilon, \quad \frac{1}{n} I(K_{13} \wedge \mathbf{F}, X_2^n) \leq \varepsilon$ ("secrecy")
- $\frac{1}{n} H(K_{12}) \geq \frac{1}{n} \log |\mathcal{K}_{12}| - \varepsilon, \quad \frac{1}{n} H(K_{13}) \geq \frac{1}{n} \log |\mathcal{K}_{13}| - \varepsilon.$ ("uniformity")

Thus, a "central" terminal $\mathcal{X}_1$ establishes a separate key with each terminal $\mathcal{X}_2$ (resp. $\mathcal{X}_3$) which is concealed from the remaining *helper* terminal $\mathcal{X}_3$ (resp. $\mathcal{X}_2$), as also from an eavesdropper with access to $\mathbf{F}$; and the keys are nearly uniformly distributed.

Private Key Capacity Region

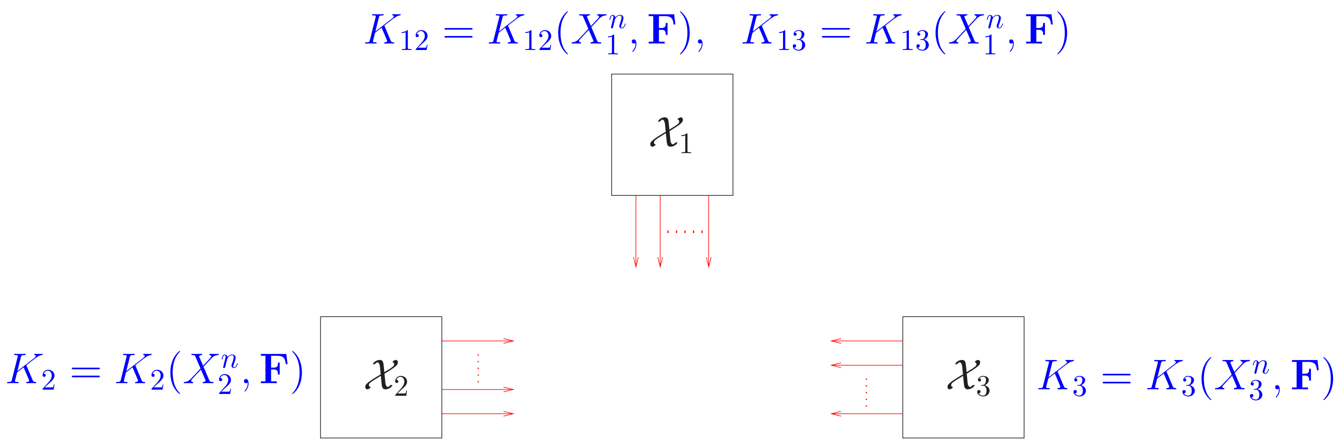

Theorem [C. Ye, '03]: If X_2 and X_3 are *deterministically correlated*, the PK-capacity region equals the set of pairs (R_{12}, R_{13}) which satisfy

$$R_{12} \leq I(X_1 \wedge X_2 | X_3), \qquad R_{13} \leq I(X_1 \wedge X_3 | X_2),$$
$$R_{12} + R_{13} \leq I(X_1 \wedge X_2, X_3) - I(X_1 \wedge X_{mcf}),$$

where X_{mcf} is the *maximal common function* of X_2 and X_3.