Multi-terminal Secrecy in Linear Non-coherent Packetized Networks

Mahdi Jafari Siavoshani
Christina Fragouli

École Polytechnique Fédérale de Lausanne
June 2012

Outline

- Introduction and Motivation
- Problem Statement
- Secrecy Upper Bound
- Secrecy Lower Bound: Achievability Scheme
- Conclusion

Motivation

- Consider that m terminals communicate through a network performing randomized linear network coding
- Goal: Creating a common secret key <K> amongst them which is concealed from a passive eavesdropper Eve
  - This can be done using public-key cryptography:
    - Based on some unproven hardness problems
    - The computational power of Eve is limited
- Alternative approach: Propose a scheme that guarantees information theoretical secrecy

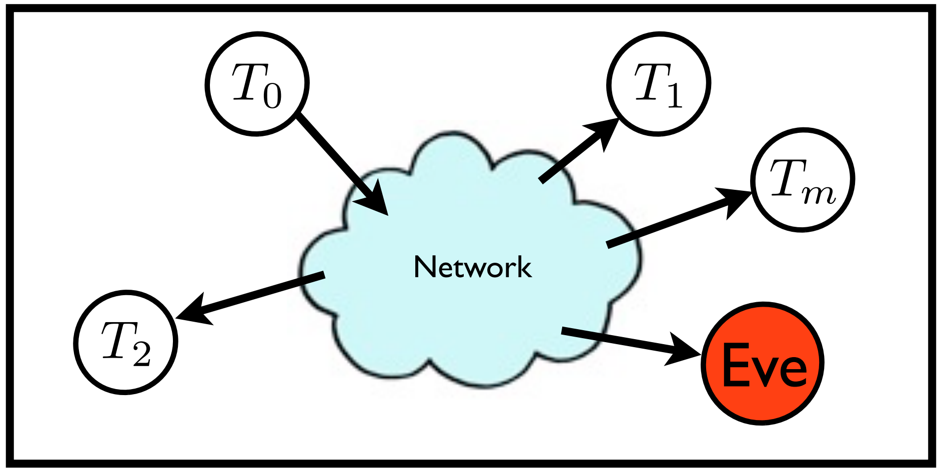

Problem Statement

- Goal: m trusted (authenticated) terminals aim to create a common secret key which is secret from a passive eavesdropper Eve
  - There is a broadcast channel from one of the terminals (Alice) to the others including Eve
  - Assume the availability of a costless public discussion channel
  - Terminals can interact in many rounds

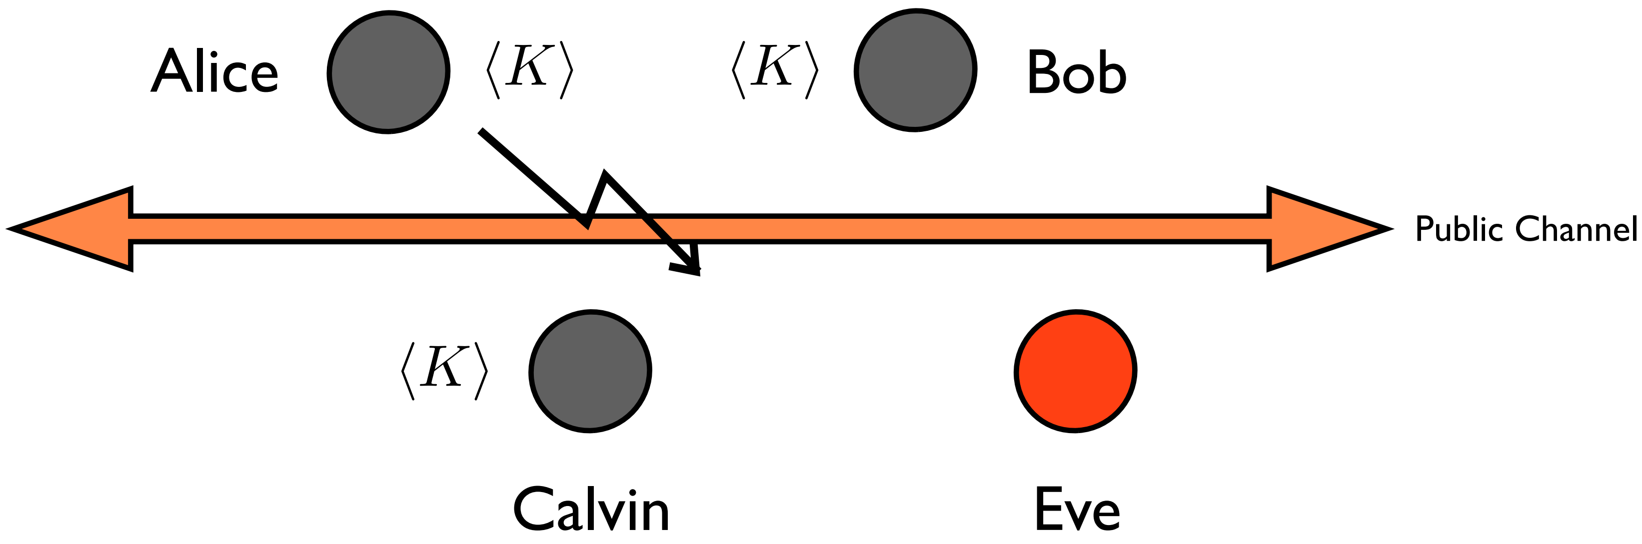

  - In general, the exact characterization of the secrecy rate is open

Problem Statement

- Assumptions:
  - Broadcast channel is a non-coherent network coding channel:
    1. The non-coherent NC is modeled by a matrix channel with uniform distribution over the transfer matrix:
       $$X_r[t] = H_r[t]X_{\mathsf{A}}[t], \qquad r \in \{1, \ldots, m, \mathsf{E}\}$$
    2. The input symbols are matrices of size $n_A \times L$ over $\mathbb{F}_q$
    3. The output symbols are matrices of size $n_r \times L$ over $\mathbb{F}_q$
  - The channels from Alice to the rest of terminal are independent, namely:
    $$P_{X_1 \cdots X_m X_E | X_A}(x_1, \ldots, x_m, x_E | x_A) = P_{X_E|X_A}(x_E|x_A) \prod_{i=1}^{m} P_{X_i|X_A}(x_i|x_A)$$
  - We study the asymptotic behavior of the secrecy capacity, by stating upper and lower bounds as the field size q increases

Related Work

- Multi-terminal secrecy:
  - Wiretap channel (Wyner 1975, Csiszar and Korner 1978)
  - Observation (Maurer 1993): Feedback can increase the secret key generation rate
  - Multi-terminal Secrecy Problem without Eve's side information (Csiszar and Narayan 2008), completely solved
  - Multi-terminal Secrecy Problem with Eve's side information (Gohari and Anantharam 2010), open even for two terminals!
- Secure Network Coding:
  - Cai and Yeung 2002, Feldman et. al. 2004, Rouayheb et. al. 2007
  - Jaggi et. al. 2008, Silva et. al. 2011

Upper Bound: Independent Broadcast Channel

- Theorem: By applying Csiszar and Narayan (2008) result (and by adding a dummy terminal) for the upper bound we can write:

$$C_s \le \max_{P_{X_0}} \min_{\lambda \in \Lambda([0:m])} \left[H(X_{[0:m]}|X_E) - \sum_{B \subsetneq [0:m]} \lambda_B H(X_B|X_{B^c}, X_E) \right]$$

where $\Lambda([0:m])$ is the set of all collections $\lambda = \{\lambda_B : \; B \subsetneq [0:m], \; B \neq \emptyset\}$ of weights $0 \le \lambda_B \le 1$ satisfying $\sum_{B \subsetneq [0:m], \; i \in B} \lambda_B = 1$

- Theorem: For independent broadcast channel, we can show that the above bound simplifies to:

$$\begin{aligned} C_s &\le \max_{P_{X_0}} \min_{i \in [1:m]} I(X_0; X_i|X_E) \\ &\le \min_{i \in [1:m]} \max_{P_{X_0}} I(X_0; X_i|X_E) \end{aligned}$$

Upper Bound

- Theorem: The secret key generation capacity is asymptotically upper bounded by:

$$C_s \le \min_{i \in [1:m]} \max_{P_{X_A}} I(X_A; X_i | X_E)$$
$$= \min_{i \in [1:m]} \Big[(\min[n_A, n_i + n_E] - n_E)(L - \min[n_A, n_i + n_E]) \Big] \log q$$

- Sketch of the proof:
  - Coding over subspace (row span of X_A) is a maximizer
  - Considering the input distribution which is uniform over subspaces of the same dimension is sufficient
  - Finally, we have to solve a convex optimization problem on $O(\min[n_A, L])$ variables, instead of $q^{n_A L}$

Lower Bound

- Theorem: The secret key sharing rate given by the solution of the following convex optimization problem can be asymptotically achieved:

$$\text{maximize} \quad \left[\min_{r\in[1:m]} \sum_{\mathcal{J}\ni r} \theta_{\mathcal{J}}\right] (L - n_A) \log q$$

$$\text{s.t.} \quad \theta_{\mathcal{J}} \geq 0, \quad \forall \mathcal{J} \subseteq [1:m],\ \mathcal{J} \neq \emptyset,$$

$$\theta_{\mathcal{J}_1} + \cdots + \theta_{\mathcal{J}_k} \leq \dim\left(U_{\mathcal{J}_1} + \cdots + U_{\mathcal{J}_k} + \Pi_E\right) - \dim(\Pi_E)$$
$$\forall k,\ \forall \mathcal{J}_1, \ldots, \mathcal{J}_k :\ \emptyset \neq \mathcal{J}_i \subseteq [1:m],\ \mathcal{J}_i \neq \mathcal{J}_j \text{ if } i \neq j$$

where for every non-empty $\mathcal{J} \subseteq [1:m]$, $U_{\mathcal{J}}$ is chosen uniformly at random from $\Pi_{\mathcal{J}}$ with dimension:

$$\dim(U_{\mathcal{J}}) = \dim(\Pi_{\mathcal{J}}) - \dim\left(\sum_{i\in\mathcal{J}^c} \Pi_{i\mathcal{J}} + \Pi_{E\mathcal{J}}\right)$$

Lower Bound: Sketch of the Proof

- Suppose that Alice broadcast $X_A[t]$ at time t of the following form:

$$X_A[t] = \left[\begin{array}{cc} I_{n_A \times n_A} & M[t] \end{array} \right]$$

- $M[t] \in \mathbb{F}_q^{n_A \times (L - n_A)}$ is a uniformly at random distributed matrix
- Legitimate terminals learn the channel and reveal $H_r[t]$ publicly
- => Alice can reconstruct subspaces $\Pi_r \triangleq \langle X_r \rangle$ for all of the legitimate terminals
- Subspaces Π_r are chosen independently and uniformly at random from Π_A => $\dim(\Pi_r) = n_r$ w.h.p.

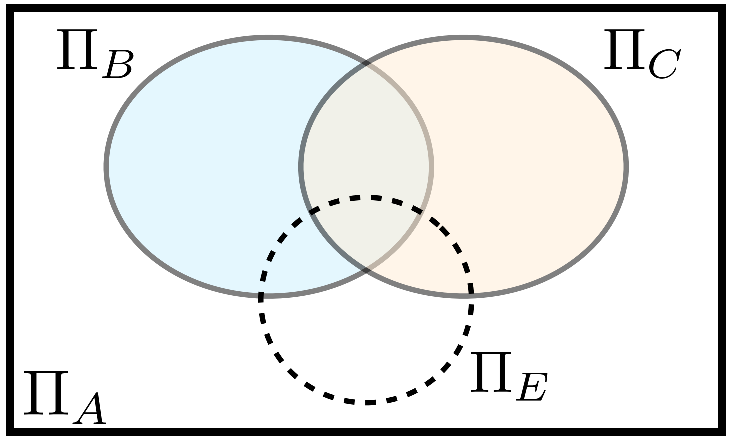

Lower Bound: Sketch of the Proof

- Suppose that Alice broadcast $X_A[t]$ at time t of the following form:

$$X_A[t] = \left[\begin{array}{cc} I_{n_A \times n_A} & M[t] \end{array} \right]$$

- $M[t] \in \mathbb{F}_q^{n_A \times (L - n_A)}$ is a uniformly at random distributed matrix
- Legitimate terminals learn the channel and reveal $H_r[t]$ publicly
- => Alice can reconstruct subspaces $\Pi_r \triangleq \langle X_r \rangle$ for all of the legitimate terminals
- Subspaces Π_r are chosen independently and uniformly at random from Π_A => $\dim(\Pi_r) = n_r$ w.h.p.
- For each non-empty $\mathcal{J} \subseteq [1:m]$ define:

$$U_{\mathcal{J}} \triangleq \Pi_{\mathcal{J}} \backslash_s \left(\sum_{i \in \mathcal{J}^c} \Pi_{i\mathcal{J}} + \Pi_{\mathsf{E}\mathcal{J}} \right)$$

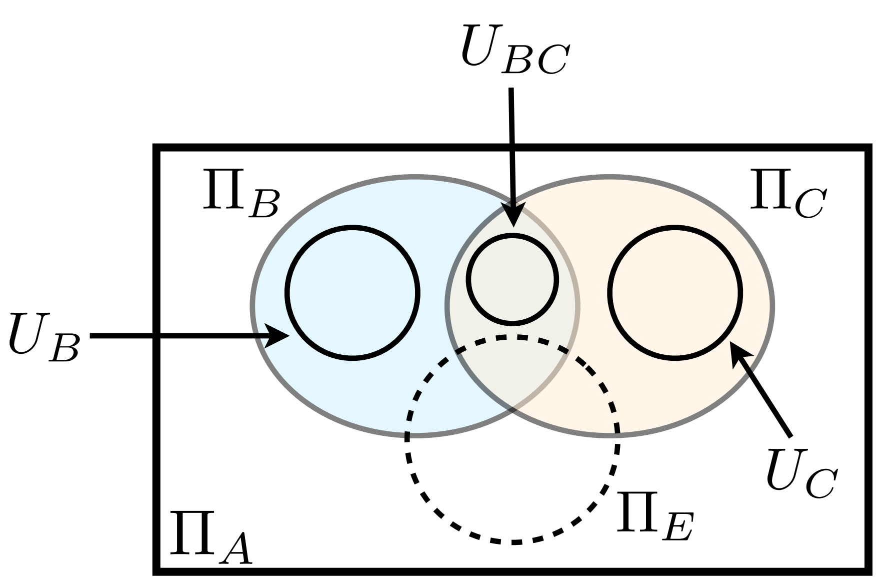

Lower Bound: Sketch of the Proof

- From definition of "$\backslash_s$" => dimension of $U_{\mathcal{J}}$ is equal to:

$$\dim(U_{\mathcal{J}}) = \dim(\Pi_{\mathcal{J}}) - \dim\left(\sum_{i\in\mathcal{J}^c} \Pi_{i\mathcal{J}} + \Pi_{\mathsf{E}\mathcal{J}}\right)$$

- Assuming q is large, Alice can calculate $\dim(U_{\mathcal{J}})$ w.h.p. even without knowing Π_E
- Observation: If Alice randomly chooses a subspace of dimension $\dim(U_{\mathcal{J}})$ from $\Pi_{\mathcal{J}}$ it satisfies w.h.p.:

$$U_{\mathcal{J}} \triangleq \Pi_{\mathcal{J}} \backslash_s \left(\sum_{i\in\mathcal{J}^c} \Pi_{i\mathcal{J}} + \Pi_{\mathsf{E}\mathcal{J}}\right)$$

- To each subset $\emptyset \neq \mathcal{J} \subseteq [1:m]$ we assign a parameter $\theta_{\mathcal{J}} \geq 0$ s.t.

$$\theta_{\mathcal{J}_1} + \cdots + \theta_{\mathcal{J}_k} \leq \dim(U_{\mathcal{J}_1} + \cdots + U_{\mathcal{J}_k} + \Pi_E) - \dim(\Pi_E)$$

for every k and any different selection of subsets: $\mathcal{J}_1, \ldots, \mathcal{J}_k$

Lower Bound: Sketch of the Proof

- Lemma*: There exist subspaces $U'_{\mathcal{J}} \sqsubseteq U_{\mathcal{J}}$ such that $\dim(U'_{\mathcal{J}}) = \theta_{\mathcal{J}}$ and all $U'_{\mathcal{J}}$ and Π_E are orthogonal subspaces w.h.p., namely:

$$\dim(\Pi_E + \sum_i U'_{\mathcal{J}_i}) = \dim(\Pi_E) + \sum_i \theta_{\mathcal{J}_i}$$

- Lemma: Alice can use a basis of $U'_{\mathcal{J}}$ to share a secret key $\mathcal{K}_{\mathcal{J}}$ with all terminals in $\mathcal{J}$. This key is secure from Eve and all terminals in $\mathcal{J}^c$
- To this end: Alice sends publicly a set of coefficients for each terminal in $\mathcal{J}$ => each of them reconstruct $U'_{\mathcal{J}}$
- Even having access to the coefficients, Eve cannot recover any info about $\mathcal{K}_{\mathcal{J}}$

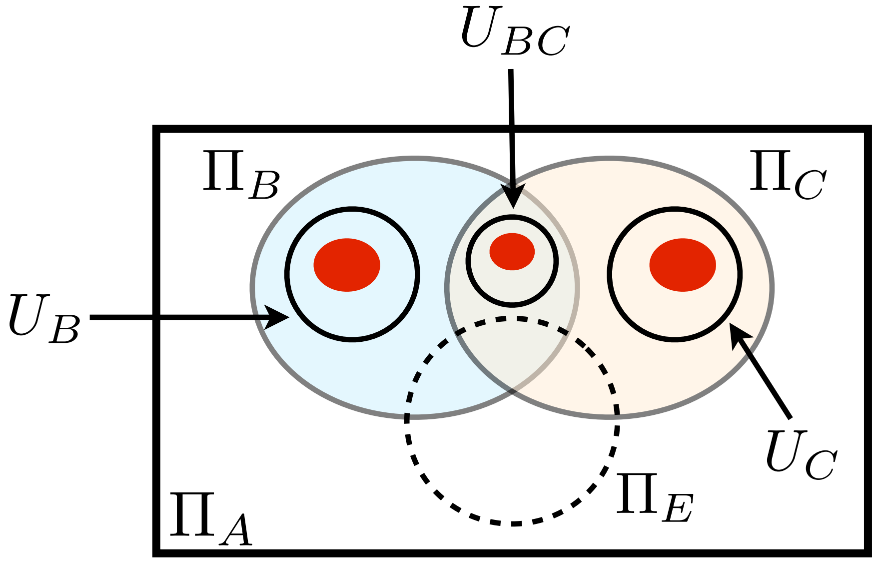

[*] Khojastepour et. al., Multicast achievable rate region of deterministic broadcast channel, 2011.

Lower Bound: Sketch of the Proof

Reconciliation Phase

- Using $\mathcal{K}_{\mathcal{J}}$ Alice can send a message of size $\theta_{\mathcal{J}}(L - n_A)\log q$ secretly to terminals in $\mathcal{J}$ over the public channel
- Now, Alice can use an MDS code to achieve the secrecy rate:

$$\left[\min_{r\in[1:m]} \sum_{\mathcal{J}\ni r} \theta_{\mathcal{J}}\right](L - n_A)\log q$$

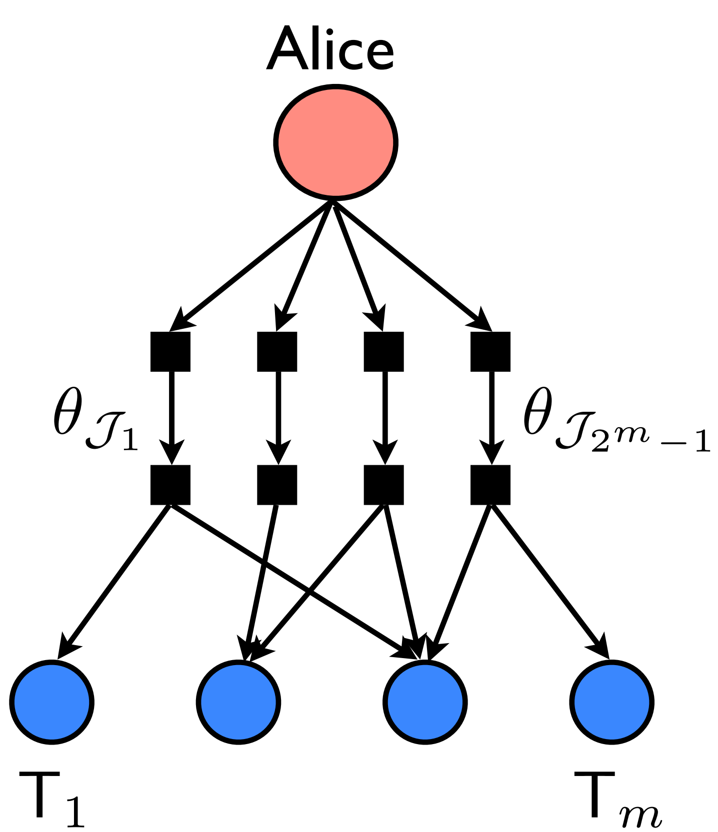

Lower Bound

- Theorem: The secret key sharing rate given by the solution of the following convex optimization problem can be asymptotically achieved:

$$\text{maximize} \quad \left[\min_{r\in[1:m]} \sum_{\mathcal{J}\ni r} \theta_{\mathcal{J}}\right](L-n_A)\log q$$

$$\text{s.t.} \quad \theta_{\mathcal{J}} \geq 0, \quad \forall \mathcal{J} \subseteq [1:m],\ \mathcal{J} \neq \emptyset,$$

$$\theta_{\mathcal{J}_1} + \cdots + \theta_{\mathcal{J}_k} \leq \dim\left(U_{\mathcal{J}_1} + \cdots + U_{\mathcal{J}_k} + \Pi_E\right) - \dim(\Pi_E)$$
$$\forall k,\ \forall \mathcal{J}_1, \ldots, \mathcal{J}_k:\ \emptyset \neq \mathcal{J}_i \subseteq [1:m],\ \mathcal{J}_i \neq \mathcal{J}_j \text{ if } i \neq j$$

where for every non-empty $\mathcal{J} \subseteq [1:m]$, $U_{\mathcal{J}}$ is chosen uniformly at random from $\Pi_{\mathcal{J}}$ with dimension:

$$\dim(U_{\mathcal{J}}) = \dim(\Pi_{\mathcal{J}}) - \dim\left(\sum_{i\in\mathcal{J}^c} \Pi_{i\mathcal{J}} + \Pi_{E\mathcal{J}}\right)$$

Example: 3 Terminals Problem

- Three terminals problem, $n_A = 60$ and $n_B = n_C = 15$

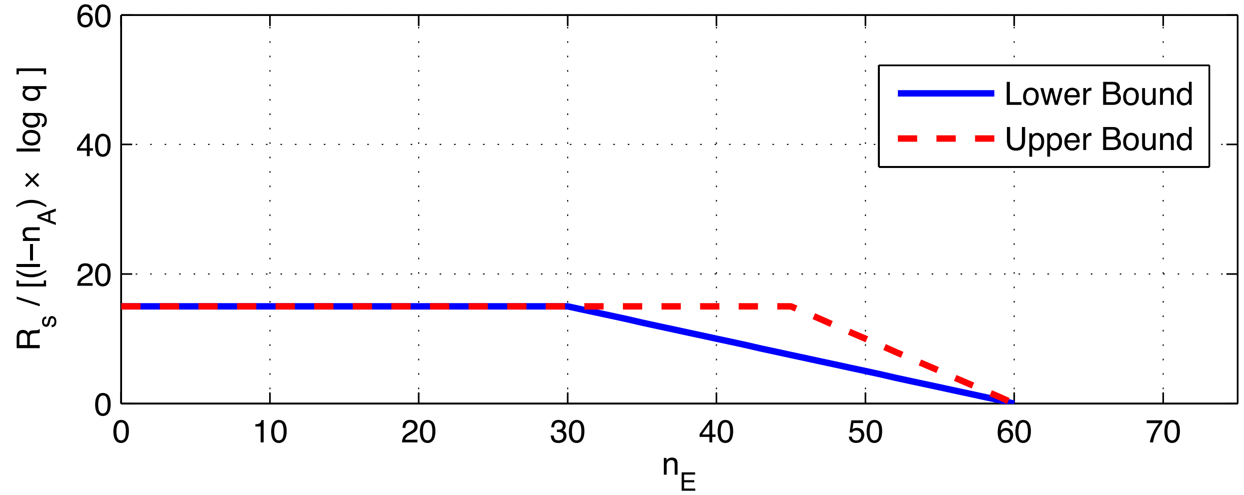

- Three terminals problem, $n_A = 60$ and $n_B = n_C = 45$

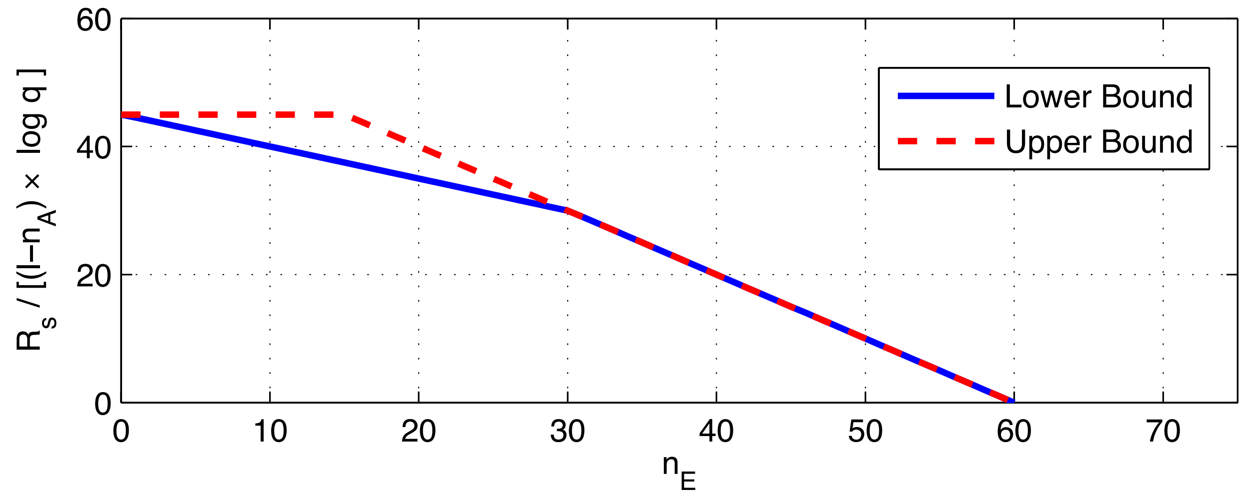

Conclusion

- We have considered the problem of secret key sharing among m terminals in the presence of a passive eavesdropper
  - Terminals communicate through a network performing randomize network coding => a non-coherent scenario
  - Terminals can discuss over a public channel
- We provide asymptotic upper and lower bounds for large field size
- For some channel parameters: the upper and lower bounds match

Thank You!