A Construction of Universal Secure Network Coding

Ryutaroh Matsumoto[1]

[1]Tokyo Institute of Technology, Japan

Sept. 22, 2010
at Institute of Network Coding
Chinese University of Hong Kong

Secure network coding

- Secure network coding (NC) is NC that prevents the eavesdropper (Eve) from gaining any information by wire-tapping the specified number of intermediate links [Cai and Yeung 02].
- Traditionally, secure NC has to be reconstructed when the network topology or coding at intermediate nodes are changed, which is inconvenient with the random network coding.
- Silva and Kschischang [2008] constructed the secure NC that provides secrecy with every coding at intermediate nodes, provided that the number of tapped links is below specified value, based on the MRD (maximum rank distance) codes.

But there is difference in the eavesdroppers' strength between [Cai and Yeung 02] and [Silva and Kschishang 08]...

Different eavesdropping models

- Silva et al.' method encodes the source information over several time slots,
- while Cai et al.' method encodes the source information in single time slot.

$\Downarrow$ this leads to . . .

- The set of tapped links has to be the same during multiple time slots in Silva et al.' method, while
- the set of tapped links is allowed to change at each time slot in Cai et al.'s method.

Is Silva et al.' method secure against Cai et al.'s eavesdropper?

We considered this question because ...

- The default behavior of the Internet (IPv4) allows intermediate routers to split an IP packet into multiple fragments and carry those fragments over different routes.

  $\Downarrow$ this means

  the set of tapped links can change at any time, even if the set of tapped links are physically the same.
- and just theoretical curiosity.

More on the set of tapped links

The argument (in the last slide) assumed the network coding is implemented as an application layer overlay network [Zhu et al. IEEE JSAC 2004].

- Nodes capable of network coding is rare.
- Those capable nodes communicate with each other by UDP or TCP.
- There are multiple physical links between two capable nodes.
- The selection of physical links between two capable nodes varies with time. This is a feature of the Internet protocol, and cannot be avoided.

Vulnerability of Silva et al's secure network coding

The MRD-code-based secure NC cannot be universally secure if the set of tapped links changes during transmission of single codeword. It remains insecure even if the number of tapped links per time slot is 1.
See our ISIT 2010 paper for details of the results.

It is convenient if we have a universal secure network coding that allows the set of tapped links can change at each time slot. We will make one.

Features of our construction

- Use the (unpopular) privacy amplification theorem
- Arbitrary small but nonzero mutual information
- Universal
- Can be combined with any kind of network error correcting codes.
- The same idea can be used for the wiretap channel and the broadcast channel with confidential messages (talk in the next week).

Family of two-universal hash functions

We will propose a construction attaching a hash function at the source node to an existing network code.

Definition

Let $\mathcal{F}$ be a set of functions from a set $\mathcal{S}_1$ to $\mathcal{S}_2$, F a (uniform) RV on $\mathcal{F}$. If for all $x_1 \neq x_2 \in \mathcal{S}_1$ we have

$$\Pr[F(x_1) = F(x_2)] \leq \frac{1}{|\mathcal{S}_2|},$$

then $\mathcal{F}$ is said to be a family of two-universal hash functions.

Privacy amplification theorem

(X, Z): a pair of discrete RVs
$\mathcal{F}$: a family two-universal hash functions from $\mathcal{X}$ to $\mathcal{S}$
F: an RV on $\mathcal{F}$ statistically independent of (X, Z).

$$I(F(X); Z|F) \leq \frac{|\mathcal{S}|^\rho \mathbf{E}[P_{X|Z}(X|Z)^\rho]}{\rho}$$

for all $0 < \rho \leq 1$.

All the logarithms, including ones in H and I, have to be the natural ones.

The form of the theorem depends on the choice of the base in the logarithm, which is unusual.

$\rho = 1$: Bennett et al. (1995).
$0 < \rho \leq 1$: Hayashi (2009).
$\rho = 1$ gives the best result for our application.

Inverse hashing construction (Csiszár 1996?)

Assume further that each $f \in \mathcal{F}$ is

- surjective, and
- for all $s \in \mathcal{S}$, $f^{-1}(s)$ has constant numbers of elements.

The set of all surjective linear maps satisfies the required condition.

S: secret information, uniform RV on $\mathcal{S}$
The uniformity assumption will be removed later.
F: RV on $\mathcal{F}$, independent of S.
$F^{-1}(S)$: the conditionally uniform RV on $\{x \in \mathcal{X} \mid F(x) = S\}$ (notation abuse)

When the source node wants to send S secretly, it sends $F^{-1}(S)$ to the outgoing edges, by using existing network coding.

We will apply the privacy amplification theorem to $F^{-1}(S)$ and Z (Eve's information) to show

$$I(S; Z, F) = I(S; Z|F)$$

can become arbitrary small.

Eavesdropping model and result summary

An edge carries one $\mathbf{F}_q$ symbol per time slot
Eve can tap up to μ links per time slot
The set of tapped links can change at each time slot.
The secret information S belongs to $\mathbf{F}_q^k$.
The source node sends n $\mathbf{F}_q$ symbols per time slot, and uses m time slots to send symbols.
Z: Eve's information, belonging to $\mathbf{F}_q^{m\mu}$.

$$I(S; Z, F) = I(S, Z|F) \leq q^{-m(n-\mu-k/m)}.$$

If k/m is kept constant below $(n-\mu)$ then $I(S; Z, F) = I(S, Z|F) \to 0$ as $m \to \infty$.

Proof will be given later.

On decoding

A sink node receives an element x in $F^{-1}(S) = \{x \in \mathcal{X} \mid F(x) = S\}$. Since $I(S, Z|F) = \mathbf{E}_f I(S, Z|F = f)$ is small, almost every choice of hash function f makes $I(S, Z|F = f)$ small. The source node and sink nodes can agree on the choice of f in advance, and a sink node can compute S by applying f to received $x \in f^{-1}(S)$.

On the asymptotic optimality

Since

$$I(S;Z,F) = I(S,Z|F) \leq q^{-m(n-\mu-k/m)},$$

we can send up to $(n-\mu)$ $\mathbf{F}_q$ symbols per time slot securely from Eve. This is the optimum. The precise statement and the proof (of the converse part) will be given later (if there is enough time).

Necessity of the uniformity assumption

In order to apply the privacy amplification theorem, we must ensure $F^{-1}(S)$ and F is statistically independent. Generally this is not the case.

However, we assumed for each $f \in \mathcal{F}$,

- $f^{-1}(S)$ has the constant number of elements,
- f is surjective, and
- S is uniform,

$F^{-1}(S)$ is uniform for every realization of F, which means that $F^{-1}(S)$ and F are independent.

Without additional assumption, the uniformity of S does not seem to be removed.

Validity of the uniformity assumption

In almost all (or most?) the research in information theoretic security, the uniform distribution of secret message is assumed. But is it valid? Realistic?

Mahdavifar and Vardy questioned the uniformity assumption and emphasized that their construction did not rely on that assumption (ISIT2010).

I will show a technique to remove the uniformity assumption from any inverse hashing construction.

Removing the uniformity assumption 1

$\mathcal{F}$: two-universal hash family from $\mathcal{X}$ to $\mathcal{S}$
For each $f \in \mathcal{F}$ and $x \in \mathcal{X}$, f is surjective and $|f^{-1}(x)|$ is constant.
σ: cyclic shift of indices of $\mathcal{S}$.
S': any RV on $\mathcal{S}$
$Z()$: a stochastic map from $\mathcal{X}$ to $\mathcal{Z}$, modeling eavesdropper's activity.
Σ: the uniform RV on $\{\sigma^i \mid i = 1, \ldots, |\mathcal{S}|\}$.
$G = \Sigma \circ F$

$$\begin{aligned} I(S', Z(G^{-1}(S'))|G) &= I(S', Z(G^{-1}(\sigma^i S'))|\sigma^{-i} \circ G) \\ &= I(\sigma^i S', Z(G^{-1}(\sigma^i S'))|G) \end{aligned}$$

Removing the uniformity assumption 2

$$
\begin{aligned}
I(S', Z(G^{-1}(S'))|G) &= \frac{1}{|S|}\sum_{i=1}^{|S|} I(\sigma^i S', Z(G^{-1}(\sigma^i S'))|G) \\
&\leq I(S, Z(G^{-1}(S))|G),
\end{aligned}
$$

where S is uniformly distributed. The last inequality follows from

- $\frac{1}{|S|}\sum_{i=1}^{|S|} P_{\sigma^i S'} = P_S$, and
- the concavity of the mutual information with respect to input distribution given fixed channel transition probability, here $Z \circ G^{-1}$ is regarded as a channel.

Proof of the result 1

An edge carries one $\mathbf{F}_q$ symbol per time slot
Eve can tap up to μ links per time slot
The set of tapped links can change at each time slot.
The secret information S belongs to $\mathbf{F}_q^k$.
The source node sends n $\mathbf{F}_q$ symbols per time slot, and uses m time slots to send symbols.
Z: Eve's information, belonging to $\mathbf{F}_q^{m\mu}$.

$$\begin{aligned} I(S;Z,F) &= I(S,Z|F) + I(S;F) \\ &= I(S,Z|F) \\ &\leq |\mathcal{S}|\mathbf{E}[P_{F^{-1}(S)|Z}(F^{-1}(S)|Z)] \end{aligned}$$

Proof of the result 2

Since

- $F^{-1}(S)$ is uniformly distributed on $\mathbf{F}_q^{mn}$,
- Z is a linear image of $F^{-1}(S)$, and
- Z belongs to $\mathbf{F}_q^{m\mu}$,

$P_{F^{-1}(S)|Z}(x|z) \leq q^{-m(n-\mu)}$. Thus

$$\begin{aligned} & I(S,Z|F) \\ \leq\ & |S|\mathbf{E}[P_{F^{-1}(S)|Z}(F^{-1}(S)|Z)] \\ \leq\ & q^{-m(n-\mu-k/m)}. \end{aligned}$$

A capacity definition

Let e_m be a stochastic encoder from $\mathbf{F}_q^{m\kappa_m}$ to $\mathbf{F}_q^{mn}$, d_m either stochastic or deterministic decoder from $\mathbf{F}_q^{mn}$ to $\mathbf{F}_q^{m\kappa_m}$, and S_m the uniform random variable on $\mathbf{F}_q^{m\kappa_m}$, for $m = 1, 2, \ldots$. If

$$\begin{aligned}
\lim_{m\to\infty} \Pr[S_m \neq d_m(e_m(S_m))] &= 0, \\
\lim_{m\to\infty} \max_{B\in\mathbf{F}_q^{\mu m\times mn}} I(S_m; Be_m(S_m))/m &= 0, \text{ (weak security)} \\
\liminf_{m\to\infty} \kappa_m &\geq R,
\end{aligned}$$

then the rate R is said to be *achievable* for the universal secure network coding with up to μ eavesdropped links per time slots. We also define the $\mathbf{F}_q$*-linear universal secure network coding capacity* as the supremum of such R over all the sequences of encoders and decoders.

The capacity is $n - \mu$. We have already proved the direct part of the capacity result.

Proof of the converse part 1

$$\text{Assume } \lim_{m\to\infty} I(S_m; Be_m(S_m))/m = 0.$$

For every m, take B as the horizontal concatenation of the $\mu m \times \mu m$ identity matrix and the zero matrix. Define $X_m^{(1)}$ as the first μm components in the random vector $e_m(S_m)$, and $X_m^{(2)}$ as the remaining components in $e_m(S_m)$. We have

$$\begin{aligned}
& H(S_m|e_m(S_m)) \\
= \; & H(S_m|X_m^{(1)}, X_m^{(2)}) \\
= \; & H(S_m) - I(S_m; X_m^{(1)}, X_m^{(2)}) \\
= \; & H(S_m) - I(S_m; X_m^{(1)}) - I(S_m; X_m^{(2)}|X_m^{(1)}) \text{ (by the chain rule)} \\
= \; & H(S_m) - I(S_m; Be_m(S_m)) - I(S_m; X_m^{(2)}|X_m^{(1)}) \\
\geq \; & H(S_m) - I(S_m; Be_m(S_m)) - H(X_m^{(2)}) \\
\geq \; & m\kappa_m \log q - I(S_m; Be_m(S_m)) - m(n-\mu)\log q \\
\geq \; & m((\kappa_m - n + \mu)\log q - I(S_m; Be_m(S_m))/m),
\end{aligned}$$

Proof of the converse part 2

Fano's inequality gives

$$\begin{aligned} & \Pr[S_m \neq d_m(e_m(S_m))] \\ \geq & \frac{H(S_m|e_m(S_m)) - 1}{\log |\mathbf{F}_q^{m\kappa_m}|} \\ \geq & \frac{m((\kappa_m - n + \mu)\log q - 1/m - I(S_m; Be_m(S_m))/m)}{m\kappa_m \log q}. \end{aligned}$$

We see $\limsup_{m\to\infty} \Pr[S_m \neq d_m(e_m(S_m))] \geq \delta/(n - \mu + \delta) > 0$, where

$$\liminf_{m\to\infty} \kappa_m \geq n - \mu + \delta.$$

Concluding discussion

The constructed universal secure network coding

- allows the set of tapped links to change,
- attaches the inverse of hash function to an existing network coding, the same idea already appeared in Csiszár (1996),
- uses Bennett et al's privacy amplification theorem.

The inverse hashing method does not depend on details of network coding, which means

- it does not depend on the network topology nor coding at intermediate nodes,
- it can be used with any network error-correcting codes,
- it can construct the wiretap codes (Hayashi 2009), and
- it can construct codes for the broadcast channels with confidential messages (the topic of next week).

The paper was registered to arXiv.org under the same title.