Signal processing for unconditional security

Nicola Laurenti

UNIVERSITÀ DEGLI STUDI DI PADOVA

DIPARTIMENTO DI INGEGNERIA DELL'INFORMAZIONE

Ph.D. Summer School in Information Engineering
Bressanone/Brixen, 7–11 July 2014

Outline

Outline

Outline

1. **What is unconditional security?**
   - What is security?
   - Computational vs unconditional security
   - Why do we need unconditional security?
2. **Signal processing for unconditional secrecy**
3. **Signal processing for unconditionally secure key agreement**
4. **Unconditionally secure authentication**

Security services and mechanisms

Security services

- Secrecy
- Data integrity
- Accountability
- Controlled access
- Availability

Security services and mechanisms

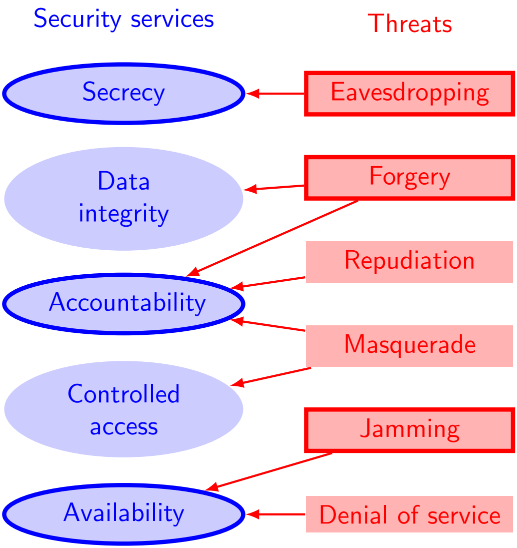

Security services and mechanisms

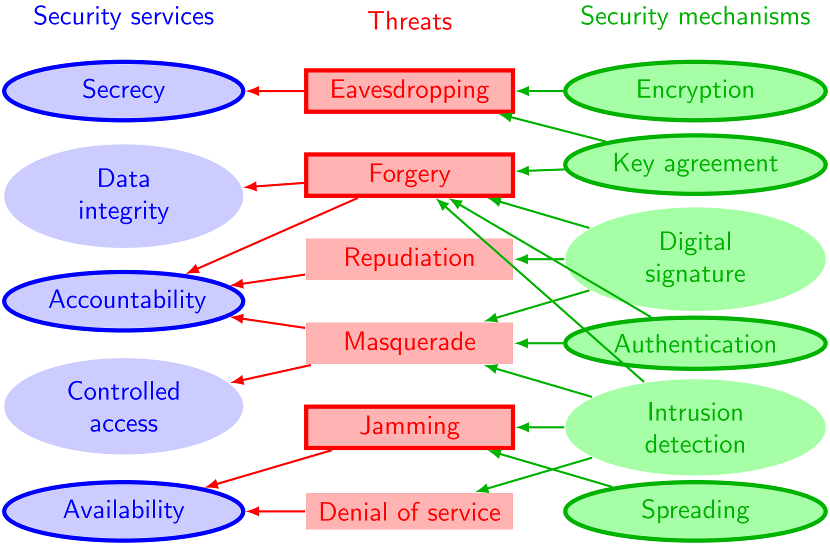

Outline

Computational security

The complexity vs. success probability tradeoff

For a (probabilistic) attack

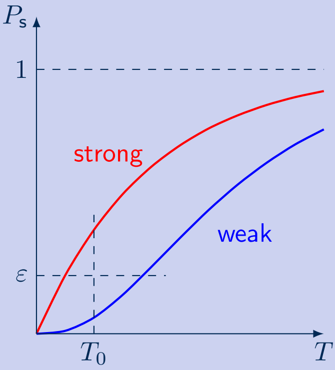

Computational security

The complexity vs. success probability tradeoff

For a (probabilistic) attack

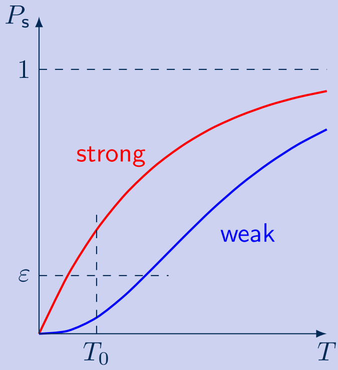

Concrete security (T_0, ε)

For any probabilistic attack with complexity T and success event S, it must be $\mathrm{P}[S, T < T_0] < \varepsilon$

Computational security

The complexity vs. success probability tradeoff

For a (probabilistic) attack

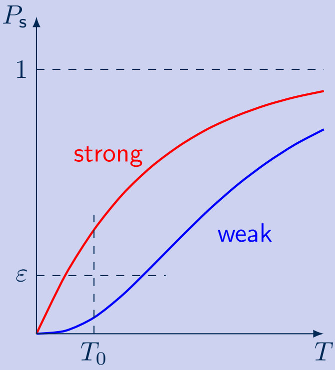

Concrete security (T_0, ε)

For any probabilistic attack with complexity T and success event S, it must be $\mathrm{P}\left[S, T < T_0\right] < \varepsilon$

Asymptotic security in key length n

For any probabilistic attack with complexity T and success event S, it must be $\mathrm{P}\left[S, T < P(n)\right] < \varepsilon(n)$ with vanishing

$$\varepsilon(n) = o(1/Q(n))$$

for any polynomials $P(n), Q(n)$.

Computational security

The complexity vs. success probability tradeoff

For a (probabilistic) attack

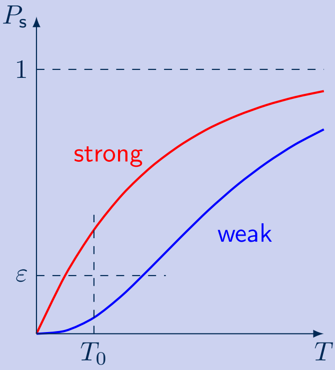

Concrete security (T_0, ε)

For any probabilistic attack with complexity T and success event S, it must be $\mathrm{P}[S, T < T_0] < \varepsilon$

Asymptotic security in key length n

For any probabilistic attack with complexity T and success event S, it must be $\mathrm{P}[S, T < P(n)] < \varepsilon(n)$ with vanishing

$$\varepsilon(n) = o(1/Q(n))$$

for any polynomials $P(n), Q(n)$.

Ex.: "brute force" attack with N trials: $T \propto N$, $P_s = N/2^n$

Physical layer security - Motivation

- Wireless communications are inherently vulnerable to various attacks
- Any device is a potential eavesdropper/jammer
- Cryptographic mechanisms (e.g., WPA) require costly key renewal
- Little is done to protect transmissions at the physical layer directly
- Diversity and randomness of the channels can be leveraged to provide security

Computational security systems can be broken by an attacker with enough computational power

Computational security systems can be broken by an attacker with enough computational power

In unconditional security, the attacker is not better off at guessing by observing the protocol communications. However, in designing the system, (statistical) knowledge of the attacker channel is often required

Computational security systems can be broken by an attacker with enough computational power

Post-quantum security systems have not been shown breakable by quantum computers in polynomial time

In unconditional security, the attacker is not better off at guessing by observing the protocol communications. However, in designing the system, (statistical) knowledge of the attacker channel is often required

Unconditional vs computational security

computational security

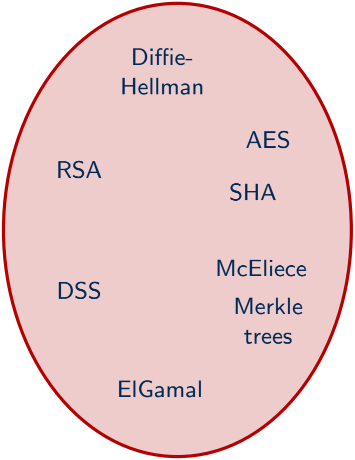

Unconditional vs computational security

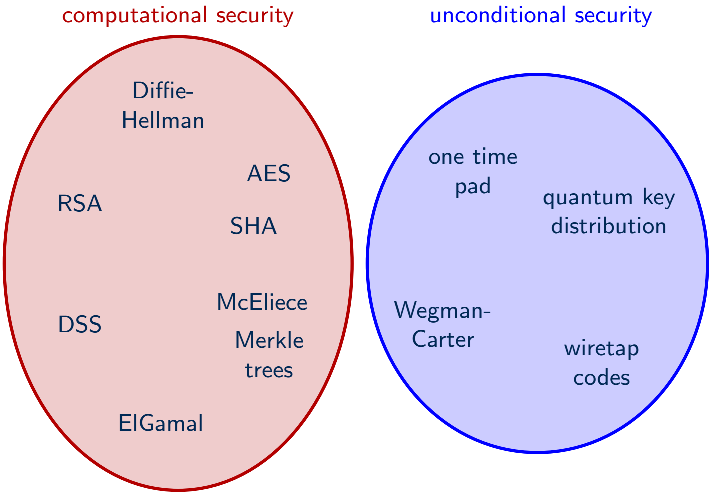

Unconditional vs computational security

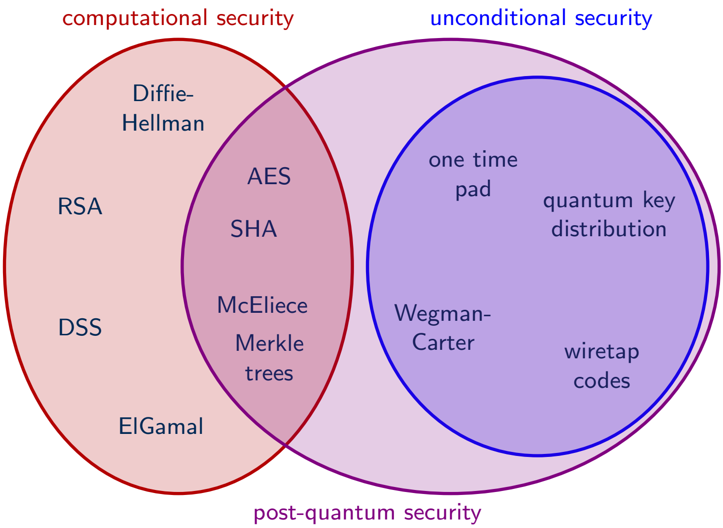

Unconditional vs computational security

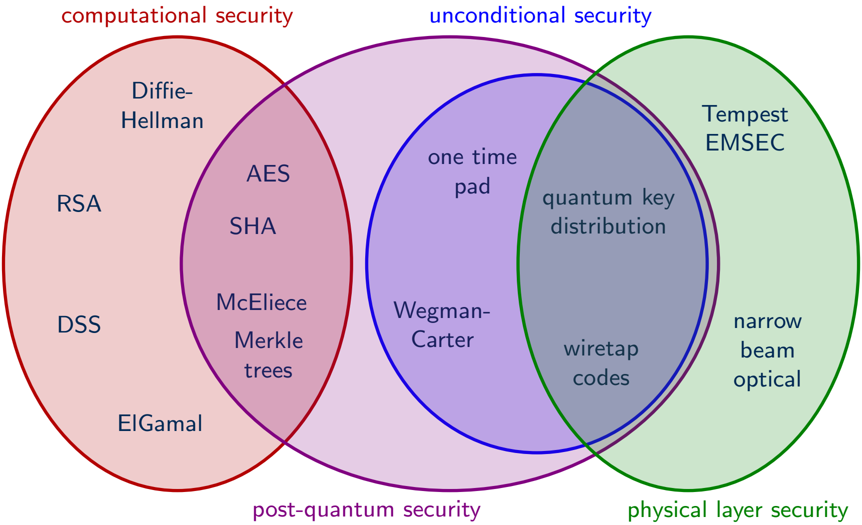

Outline

Do we really need unconditional security?

Bruce Schneier on Quantum Cryptography

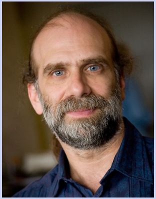

"Quantum cryptography doesn't address the weak points of the system.

Mathematical cryptography is the strongest link in most security chains. The real problems are elsewhere: computer security, network security, user interface and so on."

It's like defending yourself by putting a stake in the ground. Whether the stake is 50 feet tall or 100 feet tall, the attacker will go around it.

It's not that quantum cryptography might be insecure; it's that cryptography is already sufficiently secure."

WIRED, 16 Oct 2008

Do we really need unconditional security?

A more suitable simile, in my opinion. . .

It is true that computational security is still the strong point of security, and we should defend the weaker points. . .

Do we really need unconditional security?

A more suitable simile, in my opinion. . .

It
sti

. . . until a computational (technological or algorithmic) breakdown comes along. . .

Do we really need unconditional security?

A more suitable simile, in my opinion. . .

It
sti

. .

. . . and unless we have unconditional security. . .

Do we really need unconditional security?

Ross J. Anderson on Quantum Computing and Cryptography

"Why quantum computing is hard — and quantum cryptography is not provably secure
we still cannot perform [quantum] computation with more than about three qubits and are no closer to solving problems of real interest than a decade ago.
In consequence we dispute the claim that a quantum cryptosystem based on EPR pairs must be secure."

ArXiv, 30 Jan 2013

Scott Aaronson' response

"quantum mechanics might someday be superseded by an even deeper theory
but the fact that quantum computing still hasn't progressed beyond a few qubits does not [...] overthrow quantum mechanics."

Shtetl-Optimized, 4 Feb 2013

Outline

Outline

The wiretap channel [Wyner, '75]

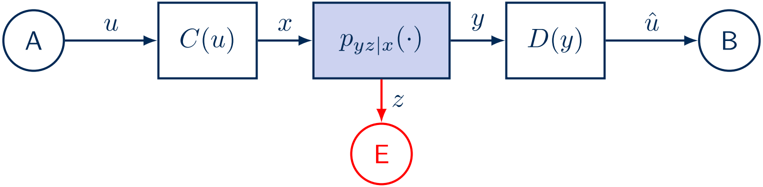

We aim for reliable transmissions to B, i.e. $\lim_{n\to\infty} P[\boldsymbol{u} \neq \hat{\boldsymbol{u}}] = 0$, under the constraint of secrecy with respect to E

Secrecy constraints

- Perfect secrecy, [Shannon, '49]: $I(\boldsymbol{u}, \boldsymbol{z}) = 0$
- Asymptotic perfect secrecy: $\lim\limits_{n\to\infty} I(\boldsymbol{u}, \boldsymbol{z}) = 0$
- Vanishing information rate, [Wyner, '75]: $\lim\limits_{n\to\infty} \frac{1}{n} I(\boldsymbol{u}, \boldsymbol{z}) = 0$

Random binning: a toy example

Random binning encoding & channel resolvability

- The basic idea is to use a probabilistic encoder $u \to x$

Random binning encoding & channel resolvability

- The basic idea is to use a probabilistic encoder $u \to x$
- Consider a subset $\mathcal{X}'_n \subset \mathcal{X}^n$ that allows to simulate the channel, that is $p_{z|x\in\mathcal{X}'_n}(\cdot) = p_{z|x\in\mathcal{X}^n}(\cdot) = p_z(\cdot)$

Random binning encoding & channel resolvability

- The basic idea is to use a probabilistic encoder $u \to x$
- Consider a subset $\mathcal{X}'_n \subset \mathcal{X}^n$ that allows to simulate the channel, that is $p_{z|x\in\mathcal{X}'_n}(\cdot) = p_{z|x\in\mathcal{X}^n}(\cdot) = p_z(\cdot)$
- Map each possible message u to a disjoint $\mathcal{X}'_n(u)$

Random binning encoding & channel resolvability

- The basic idea is to use a probabilistic encoder $u \rightarrow x$
- Consider a subset $\mathcal{X}'_n \subset \mathcal{X}^n$ that allows to simulate the channel, that is $p_{z|x \in \mathcal{X}'_n}(\cdot) = p_{z|x \in \mathcal{X}^n}(\cdot) = p_z(\cdot)$
- Map each possible message u to a disjoint $\mathcal{X}'_n(u)$
- Choose the codeword x randomly from $\mathcal{X}'_n(u)$

Random binning encoding & channel resolvability

- The basic idea is to use a probabilistic encoder $u \to x$
- Consider a subset $\mathcal{X}'_n \subset \mathcal{X}^n$ that allows to simulate the channel, that is $p_{z|x\in\mathcal{X}'_n}(\cdot) = p_{z|x\in\mathcal{X}^n}(\cdot) = p_z(\cdot)$
- Map each possible message u to a disjoint $\mathcal{X}'_n(u)$
- Choose the codeword x randomly from $\mathcal{X}'_n(u)$

Channel resolvability [Han-Verdù, '93]

The minimum number of typical codewords in $\mathcal{X}'_n$ is $|\mathcal{X}'_n| \geq 2^{nI(x;z)}$

Random binning encoding & channel resolvability

- The basic idea is to use a probabilistic encoder $u \to x$
- Consider a subset $\mathcal{X}'_n \subset \mathcal{X}^n$ that allows to simulate the channel, that is $p_{z|x\in\mathcal{X}'_n}(\cdot) = p_{z|x\in\mathcal{X}^n}(\cdot) = p_z(\cdot)$
- Map each possible message u to a disjoint $\mathcal{X}'_n(u)$
- Choose the codeword x randomly from $\mathcal{X}'_n(u)$

Channel resolvability [Han-Verdù, '93]

The minimum number of typical codewords in $\mathcal{X}'_n$ is $|\mathcal{X}'_n| \geq 2^{nI(x;z)}$

Secrecy rates and secrecy capacity

Transmission rates for which we can satisfy the secrecy constraint and guarantee reliability are called achievable secrecy rates.

Random binning encoding & channel resolvability

- The basic idea is to use a probabilistic encoder $u \to x$
- Consider a subset $\mathcal{X}'_n \subset \mathcal{X}^n$ that allows to simulate the channel, that is $p_{z|x\in\mathcal{X}'_n}(\cdot) = p_{z|x\in\mathcal{X}^n}(\cdot) = p_z(\cdot)$
- Map each possible message u to a disjoint $\mathcal{X}'_n(u)$
- Choose the codeword x randomly from $\mathcal{X}'_n(u)$

Channel resolvability [Han-Verdù, '93]

The minimum number of typical codewords in $\mathcal{X}'_n$ is $|\mathcal{X}'_n| \geq 2^{nI(x;z)}$

Secrecy rates and secrecy capacity

Transmission rates for which we can satisfy the secrecy constraint and guarantee reliability are called achievable secrecy rates. The secrecy capacity is the supremum of all achievable secrecy rates.

Secrecy capacity

Theorem

The secrecy capacity of the wiretap channel in bit/channel use is

$$C_{\mathsf{s}} = \max_u [I(u;y) - I(u;z)]^+ \geq \max_x [I(x;y) - I(x;z)]^+$$

Secrecy capacity

Theorem

The secrecy capacity of the wiretap channel in bit/channel use is

$$C_{\mathsf{s}} = \max_u [I(u;y) - I(u;z)]^+ \geq \max_x [I(x;y) - I(x;z)]^+$$

Visualization of the proof

Secrecy capacity

Theorem

The secrecy capacity of the wiretap channel in bit/channel use is

$$C_{\mathsf{s}} = \max_{u}[I(u;y) - I(u;z)]^{+} \geq \max_{x}[I(x;y) - I(x;z)]^{+}$$

Visualization of the proof

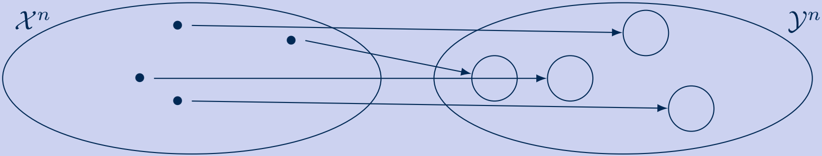

Secrecy capacity

Theorem

The secrecy capacity of the wiretap channel in bit/channel use is

$$C_{\mathsf{s}} = \max_u [I(u;y) - I(u;z)]^+ \geq \max_x [I(x;y) - I(x;z)]^+$$

Visualization of the proof

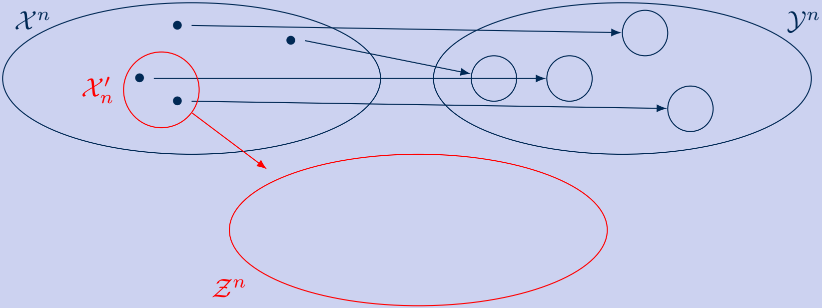

Secrecy capacity

Theorem

The secrecy capacity of the wiretap channel in bit/channel use is

$$C_\mathsf{s} = \max_u [I(u;y) - I(u;z)]^+ \geq \max_x [I(x;y) - I(x;z)]^+$$

Visualization of the proof

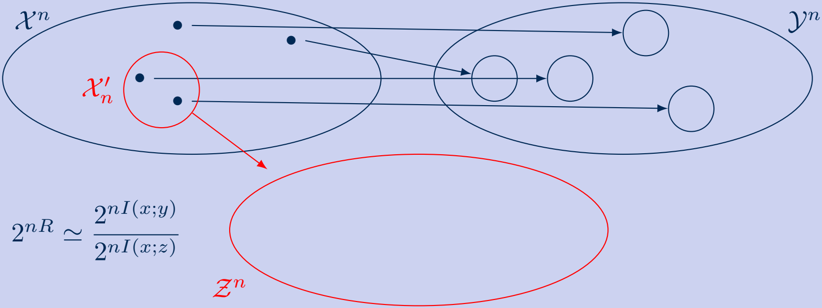

$$2^{nR} \simeq \frac{2^{nI(x;y)}}{2^{nI(x;z)}}$$

Outline

1. What is unconditional security?
2. **Signal processing for unconditional secrecy**
   - Random binning
   - Precoding and beamforming for MIMO and OFDM
3. Signal processing for unconditionally secure key agreement
4. Unconditionally secure authentication

Orthogonal frequency division multiplexing (OFDM)

Assume the legitimate nodes are communicating via OFDM modulation in presence of an eavesdropper.

Motivation for choosing OFDM:

- widely adopted as the physical layer for wireless, high-rate links
- efficient use of channel frequency diversity (high spectral efficiency)
- low complexity transceivers (FFT-based devices)

Fundamental performance limits for wiretap OFDM

- achievable secrecy-rates with OFDM transmission (and robustness wrt system parameters)
- is an OFDM receiver the best for Eve?

System block diagram

- symbol-by-symbol analysis

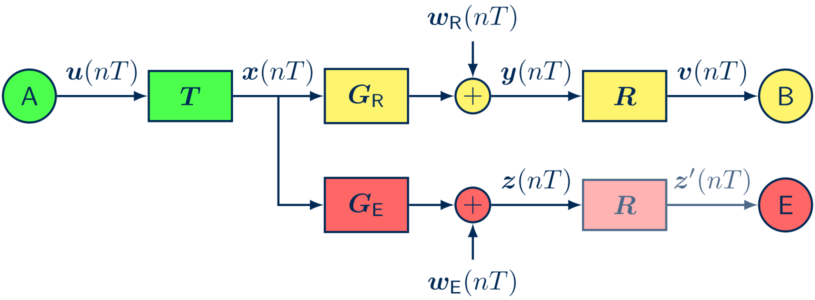

System block diagram

- symbol-by-symbol analysis
- stationarity

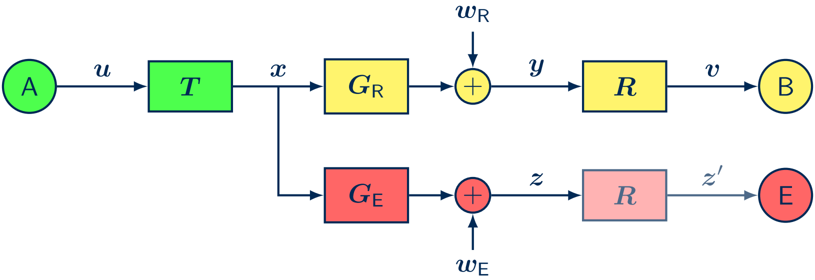

System block diagram

- symbol-by-symbol analysis
- stationarity

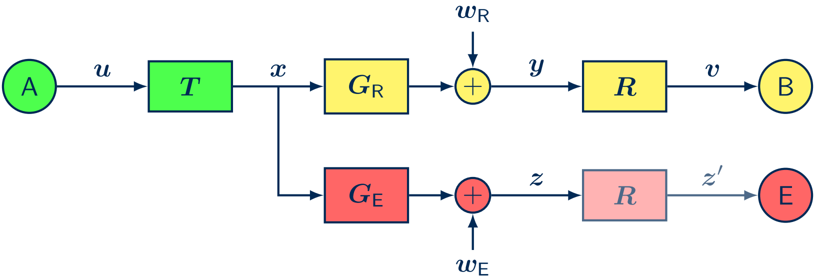

- instance of MIMO Gaussian wiretap channel (MIMOME) with $\boldsymbol{H}_\mathsf{R} = \boldsymbol{R}\boldsymbol{G}_\mathsf{R}\boldsymbol{T}$ diagonal, and $\boldsymbol{H}_\mathsf{E} = \boldsymbol{R}\boldsymbol{G}_\mathsf{E}\boldsymbol{T}$

System block diagram

- symbol-by-symbol analysis
- stationarity

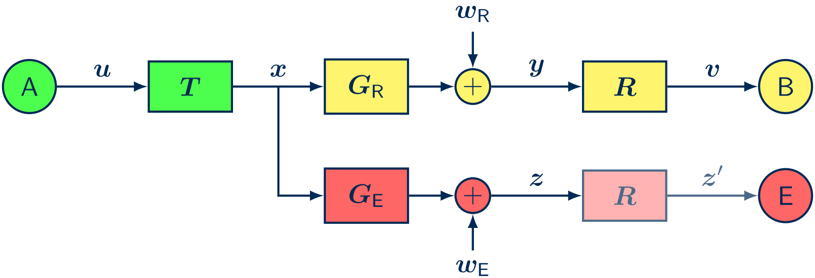

- instance of MIMO Gaussian wiretap channel (MIMOME) with $\boldsymbol{H}_{\mathsf{R}} = \boldsymbol{R}\boldsymbol{G}_{\mathsf{R}}\boldsymbol{T}$ diagonal, and $\boldsymbol{H}_{\mathsf{E}} = \boldsymbol{R}\boldsymbol{G}_{\mathsf{E}}\boldsymbol{T}$
- complete CSI on both the main and eavesdropper channel

System block diagram

- symbol-by-symbol analysis
- stationarity

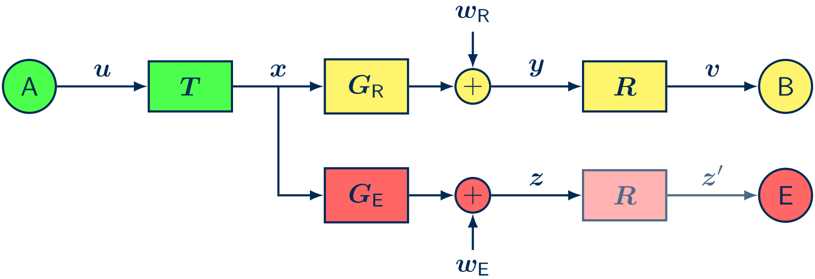

- instance of MIMO Gaussian wiretap channel (MIMOME) with $\boldsymbol{H}_{\mathsf{R}} = \boldsymbol{R}\boldsymbol{G}_{\mathsf{R}}\boldsymbol{T}$ diagonal, and $\boldsymbol{H}_{\mathsf{E}} = \boldsymbol{R}\boldsymbol{G}_{\mathsf{E}}\boldsymbol{T}$
- complete CSI on both the main and eavesdropper channel
- transmitter power constraint $\mathrm{tr}(\boldsymbol{K}_{\boldsymbol{x}}) = \mathrm{tr}(\boldsymbol{T}\boldsymbol{K}_{\boldsymbol{u}}\boldsymbol{T}^*) \leq P$

OFDM secrecy capacity (I)

Definition

$$C_s(P) = \max_{\boldsymbol{u}:\mathrm{tr}(\boldsymbol{K_x})\leq P} [I(\boldsymbol{u};\boldsymbol{v}) - I(\boldsymbol{u};\boldsymbol{z})]$$

Lemma

The secrecy capacity is achieved by a Gaussian $\boldsymbol{u}$

Proof.

Use the analogous result for a matrix covariance constraint $\boldsymbol{K_u} \preceq \boldsymbol{P}$

Let $\mathcal{K}_P = \{\boldsymbol{K} \succeq \boldsymbol{0} : \mathrm{tr}(\boldsymbol{TKT}^*) \leq P\}$, so $\bigcup_{\boldsymbol{P}\in\mathcal{K}_P} \{\boldsymbol{K} : \boldsymbol{K} \preceq \boldsymbol{P}\} = \mathcal{K}_P$

$$\begin{aligned} C_s &= \max_{\boldsymbol{P}\in\mathcal{K}_P} \max_{\boldsymbol{u}:\boldsymbol{K_u}\preceq\boldsymbol{P}} [I(\boldsymbol{u};\boldsymbol{v}) - I(\boldsymbol{u};\boldsymbol{z})] \\ &= \max_{\boldsymbol{P}\in\mathcal{K}_P} \max_{\substack{\boldsymbol{u}\sim\mathcal{CN}(\boldsymbol{0},\boldsymbol{K_u}) \\ \boldsymbol{K_u}\preceq\boldsymbol{P}}} [I(\boldsymbol{u};\boldsymbol{v}) - I(\boldsymbol{u};\boldsymbol{z})] \\ &= \max_{\substack{\boldsymbol{u}\sim\mathcal{CN}(\boldsymbol{0},\boldsymbol{K_u}) \\ \boldsymbol{K_u}\in\mathcal{K}_P}} [I(\boldsymbol{u};\boldsymbol{v}) - I(\boldsymbol{u};\boldsymbol{z})] \end{aligned}$$

$\square$

OFDM secrecy capacity (II)

Theorem

The secrecy capacity of the OFDM wiretap channel is given by

$$C_{\mathsf{s}} = \max_{\mathrm{tr}(\boldsymbol{K})\leq P} \left[\log|\boldsymbol{I} + \tilde{\boldsymbol{H}}_{\mathsf{R}}\boldsymbol{K}\tilde{\boldsymbol{H}}_{\mathsf{R}}^{*}| - \log|\boldsymbol{I} + \tilde{\boldsymbol{H}}_{\mathsf{E}}\boldsymbol{K}\tilde{\boldsymbol{H}}_{\mathsf{E}}^{*}|\right]$$

(non convex problem) where

$$\tilde{\boldsymbol{H}}_{\mathsf{R}} = \begin{cases} \boldsymbol{H}_{\mathsf{R}}\boldsymbol{D}_{\mathsf{CP}}\boldsymbol{F} & \text{for CP} \\ \boldsymbol{F}\boldsymbol{D}_{\mathsf{ZS}}\boldsymbol{H}_{\mathsf{R}} & \text{for ZS} \end{cases} \quad , \quad \tilde{\boldsymbol{H}}_{\mathsf{E}} = \begin{cases} \boldsymbol{H}_{\mathsf{E}}\boldsymbol{D}_{\mathsf{CP}}\boldsymbol{F} & \text{for CP} \\ \boldsymbol{H}_{\mathsf{E}} & \text{for ZS} \end{cases}$$

$$\boldsymbol{D}_{\mathsf{CP}} = \left[\begin{array}{c|c} \boldsymbol{I}_{M-\mu} & \boldsymbol{0} \\ \hline \boldsymbol{0} & \frac{1}{\sqrt{2}}\boldsymbol{I}_{\mu} \end{array}\right] \quad , \quad \boldsymbol{D}_{\mathsf{ZS}} = \left[\begin{array}{c|c} \frac{1}{\sqrt{2}}\boldsymbol{I}_{\mu} & \boldsymbol{0} \\ \hline \boldsymbol{0} & \boldsymbol{I}_{M-\mu} \end{array}\right]$$

The corresponding input covariance is given by

$$\boldsymbol{K}_{\boldsymbol{u}} = \begin{cases} \boldsymbol{F}\boldsymbol{D}_{\mathsf{CP}}\boldsymbol{K}^{\star}\boldsymbol{D}_{\mathsf{CP}}\boldsymbol{F} & \text{for CP} \\ \boldsymbol{K}^{\star} & \text{for ZS} \end{cases}$$

where $\boldsymbol{K}^{\star}$ *maximizes* C_{s} *above.*

Asymptotic values of secrecy capacity

High SNR limit

If $\tilde{\boldsymbol{H}}_{\mathsf{E}}$ has full column rank, then

$$\lim_{P\to\infty} C_{\mathsf{s}}(P) = \sum_{i=1}^{M} \left[\log_2 \sigma_i^2(\tilde{\boldsymbol{H}}_{\mathsf{R}}\tilde{\boldsymbol{H}}_{\mathsf{E}}^{\dagger})\right]^+$$

Low SNR limit

As $P \to 0$

$$C_{\mathsf{s}}(P) = \frac{P}{(1+\rho)\ln 2}\left[\lambda_{\max}(\tilde{\boldsymbol{H}}_{\mathsf{R}}^*\tilde{\boldsymbol{H}}_{\mathsf{R}} - \tilde{\boldsymbol{H}}_{\mathsf{E}}^*\tilde{\boldsymbol{H}}_{\mathsf{E}})\right]^+ + o(P)$$

- At high SNR, transmit on all the SVD directions of $\tilde{\boldsymbol{H}}_{\mathsf{R}}\tilde{\boldsymbol{H}}_{\mathsf{E}}^{\dagger}$ in which the legitimate receiver has higher gain than the eavesdropper.
- At low SNR, transmit only on the direction that gives the best advantage

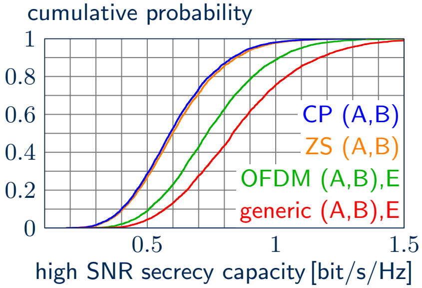

Achievable rates with Gaussian inputs

Generalized SVD

- Choose the Gaussian parallel inputs with uniform power

Achievable rates with Gaussian inputs

Generalized SVD

- Choose the Gaussian parallel inputs with uniform power

Water Filling

- Pretend the eavesdropper channel is diagonal too, with $\boldsymbol{H}_{\mathsf{E}} = \mathrm{diag}(G_{\mathsf{E}}(f_1), \ldots, G_{\mathsf{E}}(f_M))$
- Choose the optimal distribution [Li *et al.*, '06]

Achievable rates with Gaussian inputs

Generalized SVD

- Choose the Gaussian parallel inputs with uniform power

Water Filling

- Pretend the eavesdropper channel is diagonal too, with $\boldsymbol{H}_{\mathsf{E}} = \mathrm{diag}(G_{\mathsf{E}}(f_1), \dots, G_{\mathsf{E}}(f_M))$
- Choose the optimal distribution [Li *et al.*, '06]

Power allocation

- Restrict to diagonal $\boldsymbol{K_u}$
- Choose the optimal power allocation by optimization in $\mathbb{R}^M$

Achievable secrecy rates with finite inputs

Use 2^{n_i}-QAM on subchannel i

Lemma

$$\lim_{\boldsymbol{n}\to\boldsymbol{\infty}} R(\boldsymbol{n},\boldsymbol{P}) = R_{\mathsf{U}}(\boldsymbol{P})$$

Achievable secrecy rates with finite inputs

Use 2^{n_i}-QAM on subchannel i

Lemma

$$\lim_{\boldsymbol{n}\to\infty} R(\boldsymbol{n}, \boldsymbol{P}) = R_{\mathsf{U}}(\boldsymbol{P})$$

Lemma

$$\lim_{\boldsymbol{P}\to\infty} R_{\mathsf{U}}(\boldsymbol{P}) = \lim_{\boldsymbol{P}\to\infty} R_{\mathsf{G}}(\boldsymbol{P})$$

Achievable secrecy rates with finite inputs

Use 2^{n_i}-QAM on subchannel i

Lemma

$$\lim_{\boldsymbol{n}\to\infty} R(\boldsymbol{n},\boldsymbol{P}) = R_{\mathsf{U}}(\boldsymbol{P})$$

Lemma

$$\lim_{\boldsymbol{P}\to\infty} R_{\mathsf{U}}(\boldsymbol{P}) = \lim_{\boldsymbol{P}\to\infty} R_{\mathsf{G}}(\boldsymbol{P})$$

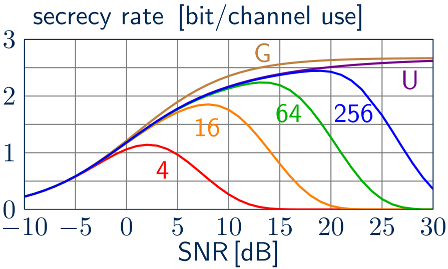

Proposition

$$\lim_{\boldsymbol{P}\to\infty}\lim_{\boldsymbol{n}\to\infty} R(\boldsymbol{n},\boldsymbol{P}) = \lim_{\boldsymbol{P}\to\infty} R_{\mathsf{G}}(\boldsymbol{P})$$

In the high SNR limit, any rate that is achievable by independent Gaussian inputs is also achievable by uniform QAM inputs with sufficient cardinality

Conclusions

- We have proved the single letter characterization of the secrecy capacity for an OFDM system with a general eavesdropper
- We have expressed in closed form the secrecy capacity at high SNR, and its derivative at low SNR, showing the loss with respect to the OFDM eavesdropper case over the statistics of a fading channel model.
- We have numerically evaluated efficient optimal power allocation schemes for generic eavesdropper, and compared them with other methods.
- We have shown that even with uniform QAM and bit loading on the main channel the high SNR secrecy capacity can be achieved.

Outline

Outline

Cryptographic key agreement [Diffie-Hellman, '76]

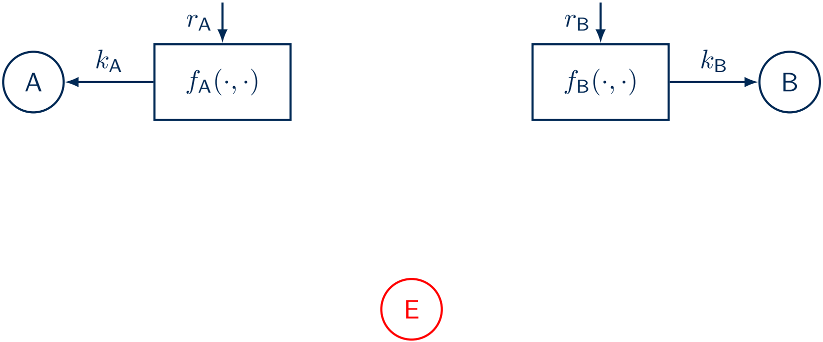

Cryptographic key agreement [Diffie-Hellman, '76]

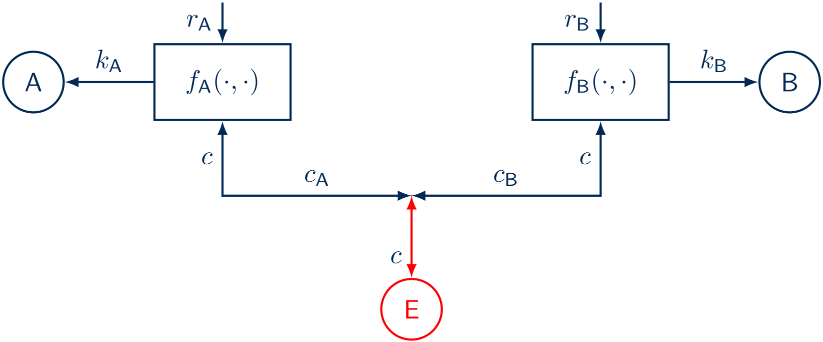

Cryptographic key agreement [Diffie-Hellman, '76]

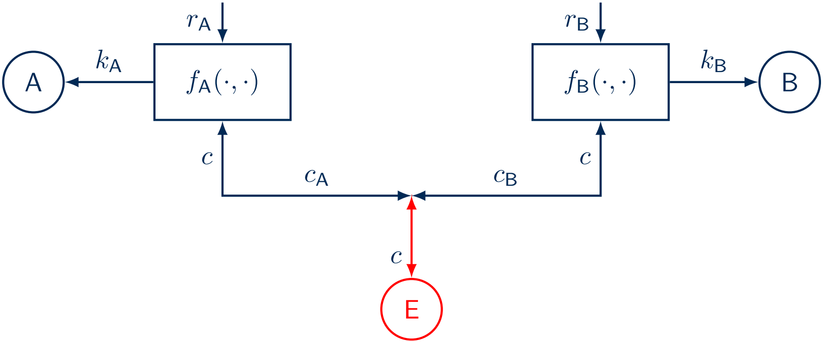

Objective

$$\max L(k_A) \quad \text{subject to:}$$

correctness: $k_A = k_B$

secrecy: infeasible to derive k from c

uniformity: $p_{k_A}(a) \approx 1/2^{L(k_A)}$

Unconditional key agreement [Ahlswede-Csiszar, '93]

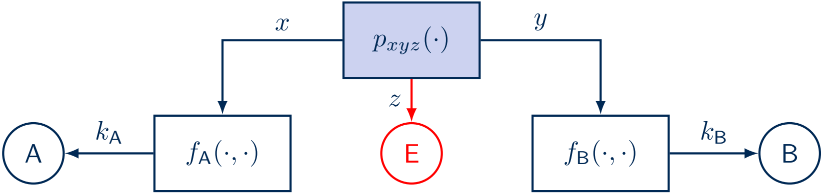

Unconditional key agreement [Ahlswede-Csiszar, '93]

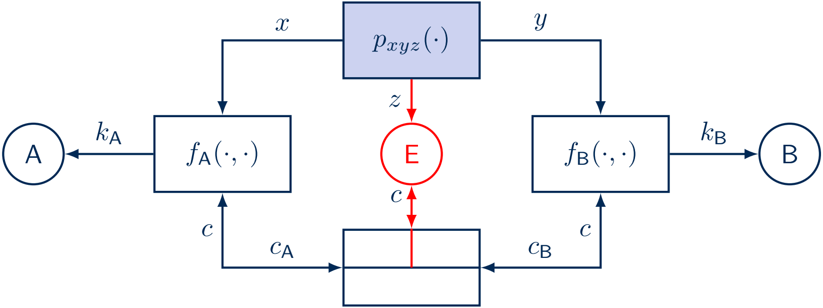

Unconditional key agreement [Ahlswede-Csiszar, '93]

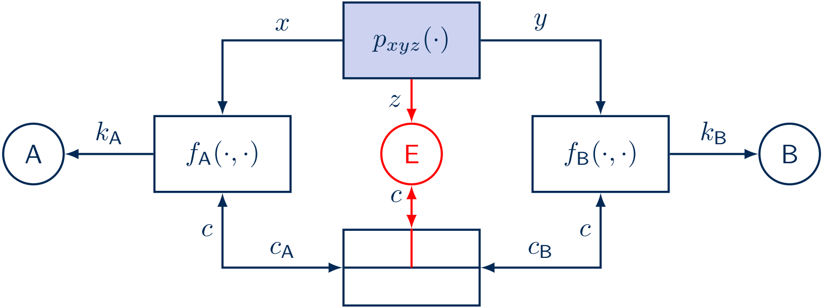

Objective

$$\max_{f_A, f_B} H(k_A) \quad \text{subject to:}$$

correctness: $\mathrm{P}\left[k_A \neq k_B\right] < \varepsilon$

secrecy: $I(k_A, k_B; z, c) < \varepsilon'$

uniformity: $L - H(k_A) < \varepsilon''$

Unconditional key agreement [Ahlswede-Csiszar, '93]

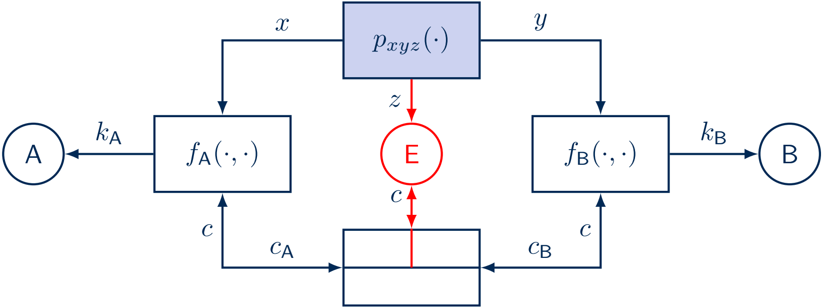

Objective

$$\max_{f_{\mathsf{A}},f_{\mathsf{B}}} H(k_{\mathsf{A}}) \quad \text{subject to:}$$

correctness: $\mathrm{P}\left[k_{\mathsf{A}} \neq k_{\mathsf{B}}\right] < \varepsilon$

secrecy: $I(k_{\mathsf{A}}, k_{\mathsf{B}}; z, c) < \varepsilon'$

uniformity: $L - H(k_{\mathsf{A}}) < \varepsilon''$

Secret-key capacity

$$S = \lim_{n\to\infty} \max_{f_{\mathsf{A}},f_{\mathsf{B}}} \left[\frac{1}{n} H(k_{\mathsf{A}})\right]$$

and $\varepsilon, \varepsilon', \varepsilon'' \to 0$

upper bound: $S \leq I(x;y|z)$

Unconditional key agreement [Maurer, '93]

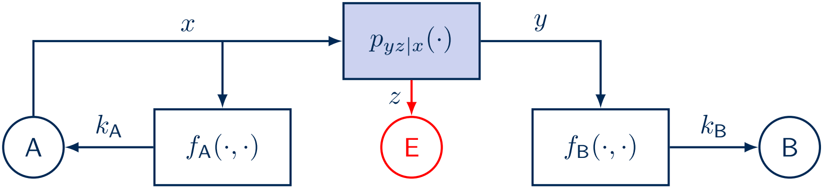

Unconditional key agreement [Maurer, '93]

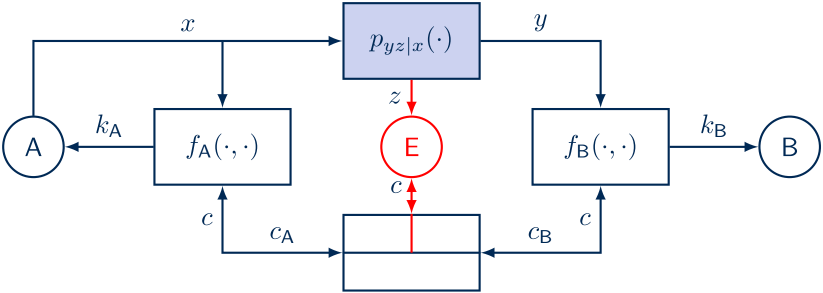

Unconditional key agreement [Maurer, '93]

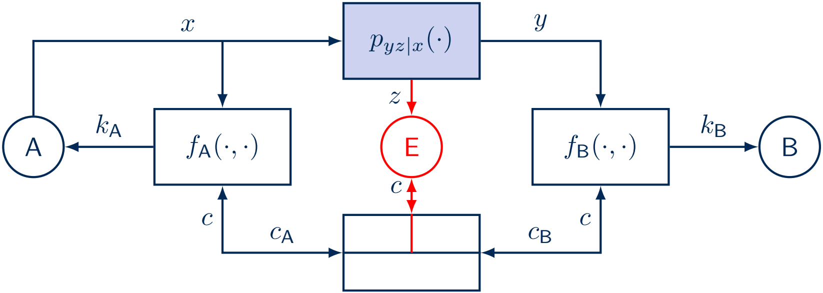

Upper and lower bounds

$$S \leq \max_x I(x;y|z)$$

$$S \geq \max_x \left[I(x;y) - I(x;z)\right]$$

Unconditional key agreement [Maurer, '93]

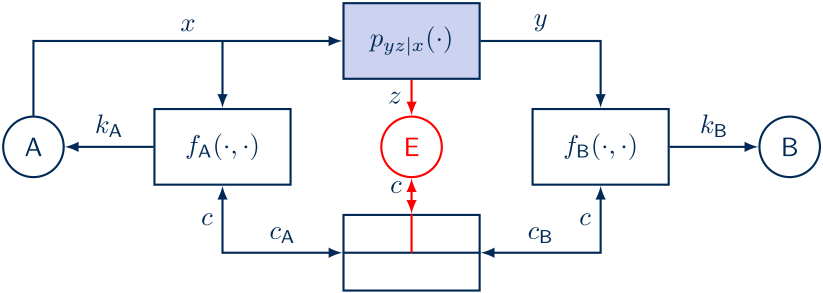

Upper and lower bounds

$$S \le \max_x I(x;y|z)$$

$$S \ge \max_x [I(x;y) - I(x;z)]$$

Optimal 3-step protocol

1. randomness sharing
   to maximize $I(x;y|z)$
2. information reconciliation
   for correctness
3. privacy amplification
   for secrecy and uniformity

Divide et impera

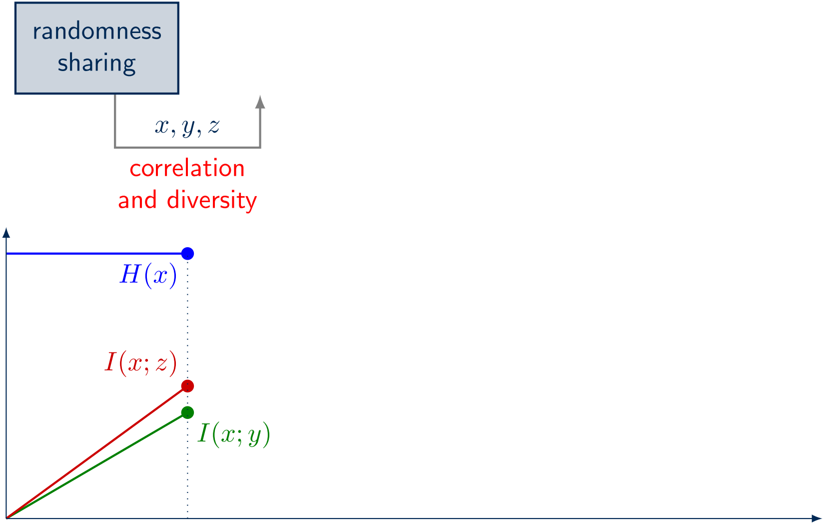

Divide et impera

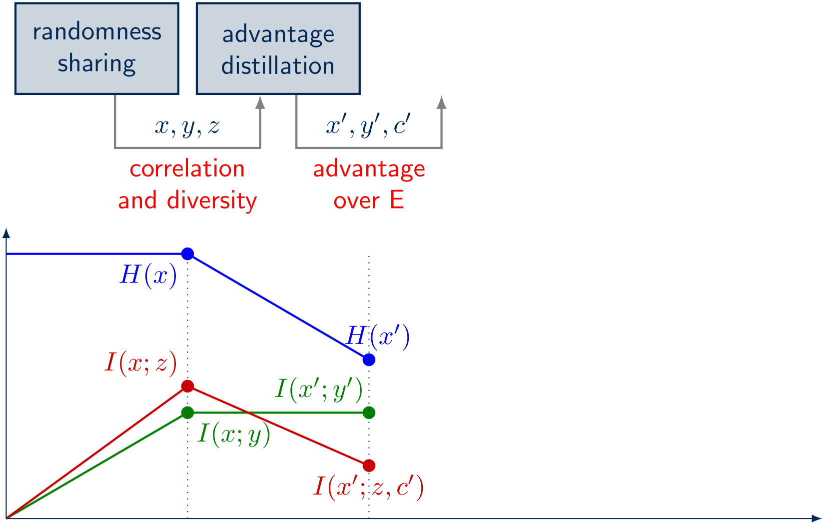

Divide et impera

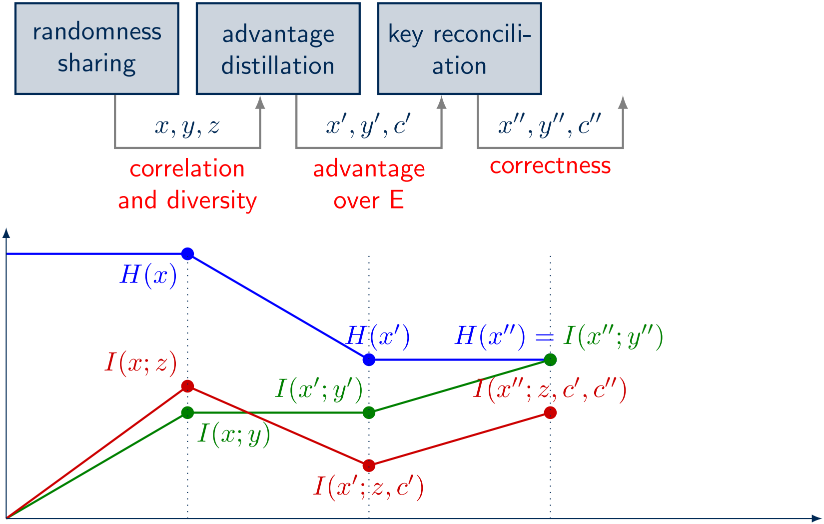

Divide et impera

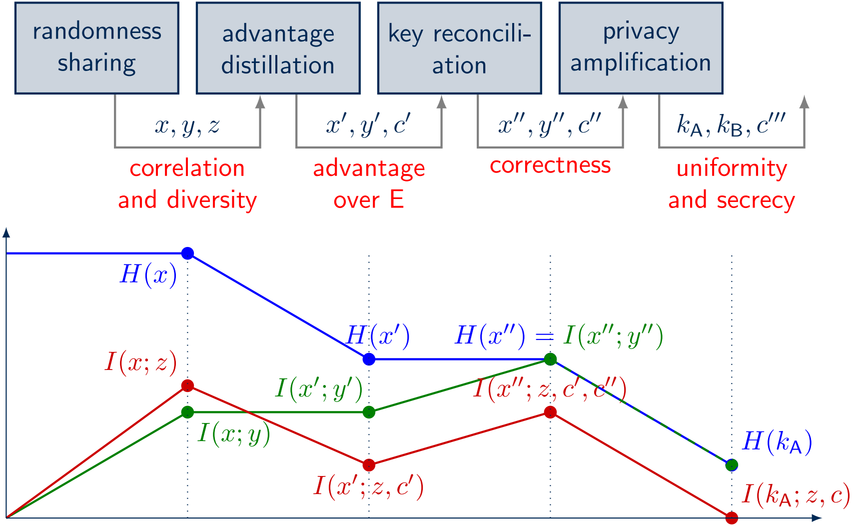

Outline

1. What is unconditional security?
2. Signal processing for unconditional secrecy
3. Signal processing for unconditionally secure key agreement
   - unconditionally secure key agreement
   - Information reconciliation
   - Privacy amplification
   - Precoding and beamforming for MIMO randomness sharing
4. Unconditionally secure authentication

Reconciliation of sifted keys

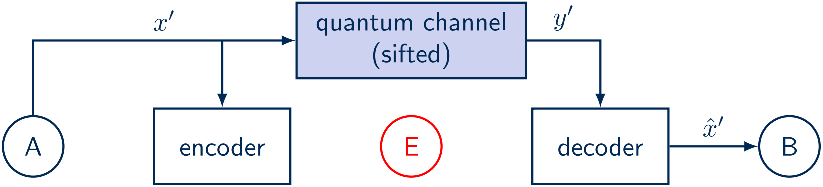

Reconciliation of sifted keys

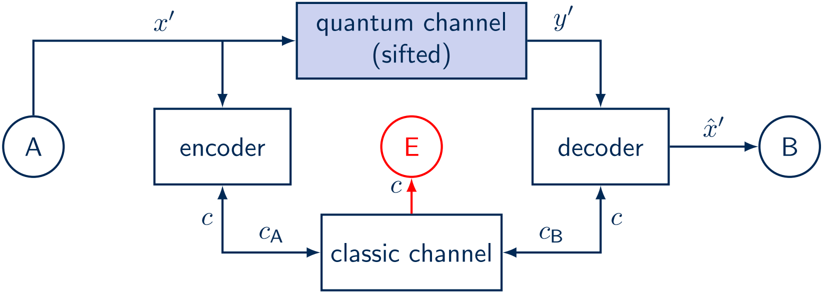

Reconciliation of sifted keys

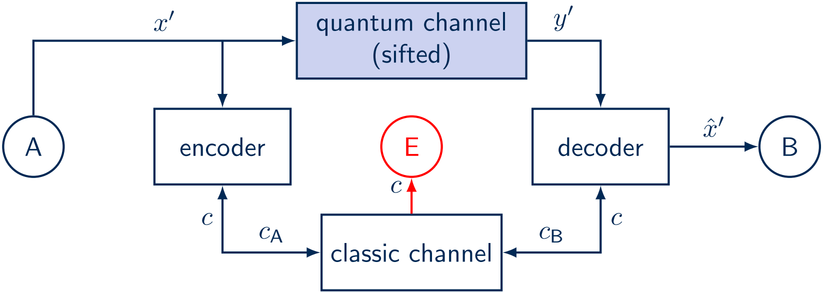

quantum channel	classic channel
private low rate unreliable	public high rate reliable

Reconciliation of sifted keys

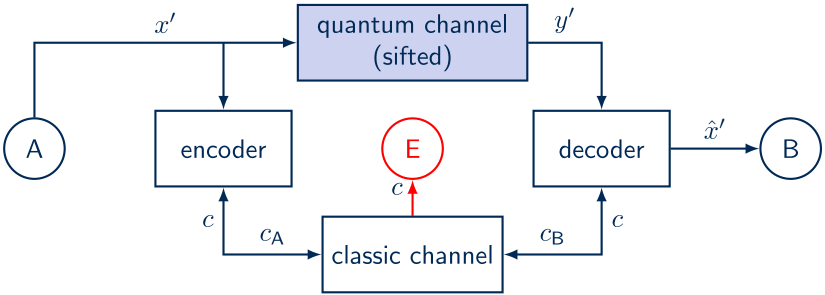

quantum channel	classic channel
private low rate unreliable	public high rate reliable

Aim

To allow B to reliably reconstruct $\hat{x}' = x'$, by transmitting $c = (c_\mathsf{A}, c_\mathsf{B})$ publicly, with the minimum leakage of information $I(x'; c)$ to E.

Existing models and solutions

Coding techniques for reconciliation fall into 1 of 3 categories:

Existing models and solutions

Coding techniques for reconciliation fall into 1 of 3 categories:

cascade iteratively (and interactively) split the keys to locate single errors and correct them [Brassard-Salvail, ‘93]

Existing models and solutions

Coding techniques for reconciliation fall into 1 of 3 categories:

cascade iteratively (and interactively) split the keys to locate single errors and correct them [Brassard-Salvail, '93]

hashing given a $(n, n-r)$ parity check matrix $\boldsymbol{H}$
Alice transmits $\boldsymbol{c} = \boldsymbol{H}\boldsymbol{x}'$.
Bob chooses $\hat{\boldsymbol{x}}' = \arg\min_{\boldsymbol{a}:\boldsymbol{H}\boldsymbol{a}=\boldsymbol{c}} d(\boldsymbol{a}, \boldsymbol{y})$
Examples: Winnow [Buttler *et al.*, '03]
LDPC [Elkouss *et al.*, '09]

Existing models and solutions

Coding techniques for reconciliation fall into 1 of 3 categories:

cascade iteratively (and interactively) split the keys to locate single errors and correct them [Brassard-Salvail, '93]

hashing given a $(n, n-r)$ parity check matrix $\boldsymbol{H}$
Alice transmits $\boldsymbol{c} = \boldsymbol{H}\boldsymbol{x}'$.
Bob chooses $\hat{\boldsymbol{x}}' = \arg\min_{\boldsymbol{a}:\boldsymbol{H}\boldsymbol{a}=\boldsymbol{c}} d(\boldsymbol{a}, \boldsymbol{y})$
Examples: Winnow [Buttler *et al.*, '03]
LDPC [Elkouss *et al.*, '09]

systematic pick a $(n+r, n)$ generating matrix $\boldsymbol{G} = \begin{bmatrix} \boldsymbol{I}_n \\ \boldsymbol{A} \end{bmatrix}$
Alice transmits $\boldsymbol{c} = \boldsymbol{A}\boldsymbol{x}'$.
Bob chooses $\hat{\boldsymbol{x}}' = \arg\min_{\boldsymbol{a}\in\mathcal{C}} d(\boldsymbol{a}, \boldsymbol{y})$
Examples: LDPC [Mondin *et al.*, '10]
BCH [Traisilanun *et al.*, '07]

Existing models and solutions

The choice of the coding technique for reconciliation depends on the model for the classical channel

layer	ch. type	condition	delays	codes used
Physical	AWGN	high SNR	none	systematic (soft)
Data link	binary	low BER	low	systematic (hard)
Net & up	packet	error free	long	cascade, hashing

Cascade and Winnow: common structure

$i = 0, \varepsilon_i = \varepsilon_{\mathsf{q}}$

segment x' and y' into blocks of length L_i, with $\varepsilon_i L_i \ll 1$

check if parity of each block is the same in x' and y'

correct error in blocks of y' with different parities

$i \leftarrow i+1$, estimate ε_i, equally permute x' and y'

for $i = I$, let $\hat{x}' = y'$

- the condition $\varepsilon_i L_i \ll 1$ ensures that multiple errors in a block are unlikely
- the block parities need to be exchanged $(c_{\mathsf{A}}, c_{\mathsf{B}})$
- both algorithms can correct a single error per block

Outline

Privacy amplification

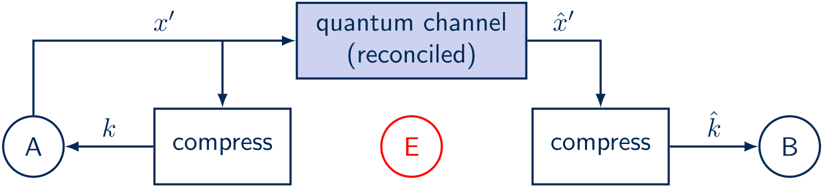

Privacy amplification

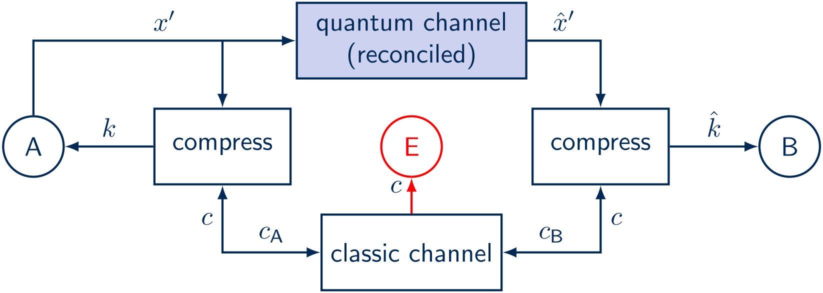

Privacy amplification

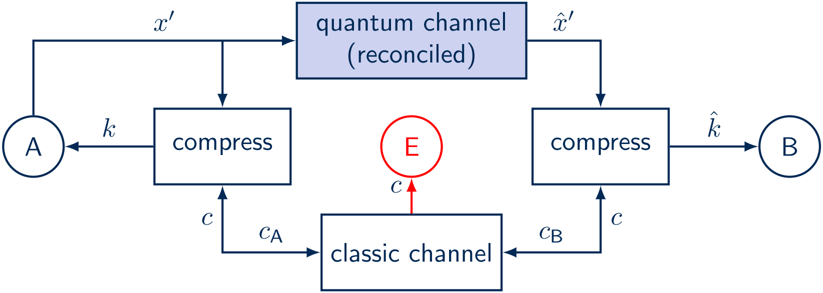

quantum channel	classic channel
private low rate	public high rate

Privacy amplification

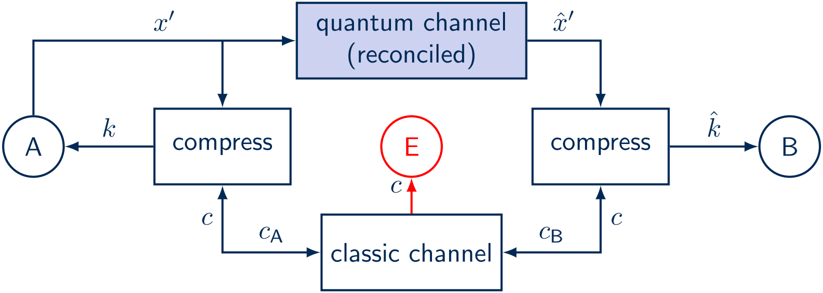

quantum channel	classic channel
private low rate	public high rate

Aim

To allow A and B to remove any information E might have from $\hat{k} = k$, by publicly agreeing on the compressing function, and with the minimum amount of compression.

Choosing a compression function

- Must be chosen randomly, after transmission
- Must be compactly representable

Assume we know that Eve has observed some t-bit linear function of the reconciled key

$$\boldsymbol{z} = \boldsymbol{M}\boldsymbol{x}' \quad , \qquad \text{with } \boldsymbol{M} \in \{0,1\}^{t \times n}$$

(include $\boldsymbol{c}$ observed during reconciliation)

Theorem (Universal hashing functions [Bennett *et al.*, '95]**)**

If the compressing function $\boldsymbol{A}$ is chosen uniformly from a class of universal hashing $s \times n$ matrices, then on average (over $\boldsymbol{M}$ and $\boldsymbol{A}$)

$$I(\boldsymbol{k}; \boldsymbol{z}, \boldsymbol{A}) \leq \frac{1}{\ln 2} 2^{s+t-n}$$

Choosing a compression function

Once we choose a hashing matrix $\boldsymbol{A}$, we would like to obtain

1. $H(\boldsymbol{k}) = s$ (perfect uniformity)
2. $I(\boldsymbol{k}; \boldsymbol{z}) = 0$ (perfect secrecy)

Choosing a compression function

Once we choose a hashing matrix $\boldsymbol{A}$, we would like to obtain

1. $H(\boldsymbol{k}) = s$ (perfect uniformity)
2. $I(\boldsymbol{k}; \boldsymbol{z}) = 0$ (perfect secrecy)

Lemma 1

If $\mathrm{rank}(\boldsymbol{A}) = s$ and $\boldsymbol{x}'$ is uniform over $\{0,1\}^n$, then $\boldsymbol{k}$ is uniform over $\{0,1\}^s$

Choosing a compression function

Once we choose a hashing matrix $\boldsymbol{A}$, we would like to obtain

1. $H(\boldsymbol{k}) = s$ (perfect uniformity)
2. $I(\boldsymbol{k};\boldsymbol{z}) = 0$ (perfect secrecy)

Lemma 1

If $\mathrm{rank}(\boldsymbol{A}) = s$ and $\boldsymbol{x}'$ is uniform over $\{0,1\}^n$, then $\boldsymbol{k}$ is uniform over $\{0,1\}^s$

Example: binary Toeplitz matrices

- $\boldsymbol{A}$ is uniquely specified by $n+s-1$ bits $\boldsymbol{a} = [a_{-r+1}, \ldots, a_{n-1}]$
- If $\boldsymbol{a}$ is uniform in $\{0,1\}^{n+s-1}$, $\mathrm{P}\left[\text{rank}(\boldsymbol{A}) < s\right] = 1/2^{n-s+1}$

Choosing a compression function

Once we choose a hashing matrix $\boldsymbol{A}$, we would like to obtain

1. $H(\boldsymbol{k}) = s$ (perfect uniformity)
2. $I(\boldsymbol{k};\boldsymbol{z}) = 0$ (perfect secrecy)

Lemma 1

If $\mathrm{rank}(\boldsymbol{A}) = s$ and $\boldsymbol{x}'$ is uniform over $\{0,1\}^n$, then $\boldsymbol{k}$ is uniform over $\{0,1\}^s$

Example: binary Toeplitz matrices

- $\boldsymbol{A}$ is uniquely specified by $n+s-1$ bits $\boldsymbol{a} = [a_{-r+1}, \ldots, a_{n-1}]$
- If $\boldsymbol{a}$ is uniform in $\{0,1\}^{n+s-1}$, $\mathrm{P}\left[\text{rank}(\boldsymbol{A}) < s\right] = 1/2^{n-s+1}$

Lemma 2

If $\dim \mathcal{N}(\boldsymbol{M}) - \dim\left(\mathcal{N}(\boldsymbol{M}) \cap \mathcal{N}(\boldsymbol{A})\right) = \mathrm{rank}(\boldsymbol{A})$ and $\boldsymbol{x}'$ is uniform over $\{0,1\}^n$, then $I(\boldsymbol{k};\boldsymbol{z}) = 0$

Choosing a compression function

Theorem

If $\dim \mathcal{N}(\boldsymbol{M}) - \dim(\mathcal{N}(\boldsymbol{M}) \cap \mathcal{N}(\boldsymbol{A})) = s$ *and* $\boldsymbol{x}'$ *is uniform over* $\{0,1\}^n$*, then* $\boldsymbol{k}$ *is uniform and perfectly secret.*

Choosing a compression function

Theorem

If $\dim \mathcal{N}(\boldsymbol{M}) - \dim(\mathcal{N}(\boldsymbol{M}) \cap \mathcal{N}(\boldsymbol{A})) = s$ *and* $\boldsymbol{x}'$ *is uniform over* $\{0,1\}^n$, *then* $\boldsymbol{k}$ *is uniform and perfectly secret.*

Illustration

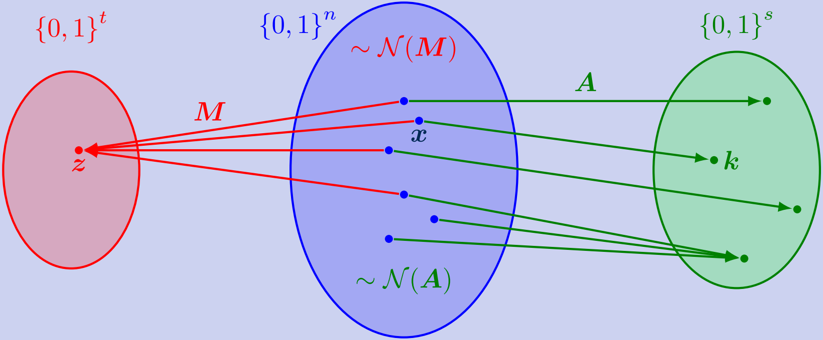

Choosing a compression function

Theorem

If $\dim \mathcal{N}(\boldsymbol{M}) - \dim(\mathcal{N}(\boldsymbol{M}) \cap \mathcal{N}(\boldsymbol{A})) = s$ *and* $\boldsymbol{x}'$ *is uniform over* $\{0,1\}^n$*, then* $\boldsymbol{k}$ *is uniform and perfectly secret.*

Illustration

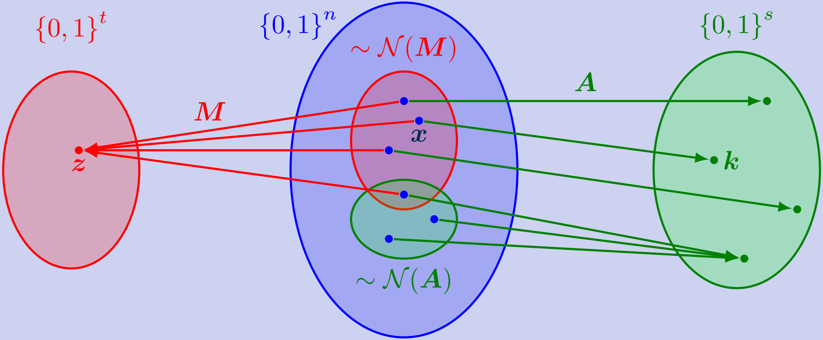

Outline

Secret-key rate

ℓ: length of $\boldsymbol{k}$

n: number of noisy channel uses

$R = \ell/n$: key rate

Definition

A secret-key rate R is achievable if

- $\lim_{n\to\infty} \mathrm{P}\left[\boldsymbol{k} \neq \hat{\boldsymbol{k}}\right] = 0$ (reliability)
- $\lim_{n\to\infty} I(\boldsymbol{k}; \boldsymbol{z}, \boldsymbol{r}_{\mathsf{A}}, \boldsymbol{r}_{\mathsf{B}}) = 0$ (secrecy)
- $\lim_{n\to\infty} H(\boldsymbol{k}) - nR = 0$ (uniformity)

Secret-key capacity

$$S = \sup\{R : \ R \text{ is achievable}\}$$

Secret-key agreement over MIMO channels

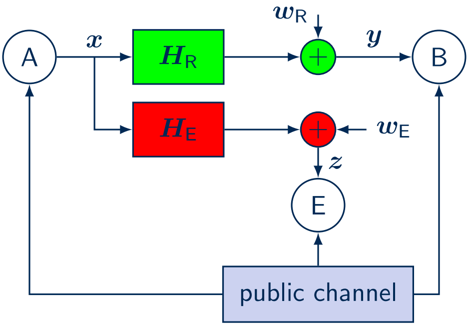

- quasi-static MIMO channels (OFDM as a particular case)
- assume $\boldsymbol{H}_\mathsf{E}$ full column rank (otherwise d.o.f.)
- average power constraint, $\mathrm{tr}(\boldsymbol{K}_{\boldsymbol{x}}) \leq P$
- complete CSI on both the main and eavesdropper channel

Secret-key capacity

Lemma

The secret-key capacity is achieved with a Gaussian $\boldsymbol{x}$ and is given by

$$S(P) = \max_{\mathrm{tr}(\boldsymbol{K_x}) \leq P} \log \left| \boldsymbol{I} + \boldsymbol{K}_{\boldsymbol{x}}^{\frac{1}{2}} \boldsymbol{H}_{\mathsf{R}}^{*} \boldsymbol{H}_{\mathsf{R}} \boldsymbol{K}_{\boldsymbol{x}}^{\frac{1}{2}} (\boldsymbol{I} + \boldsymbol{K}_{\boldsymbol{x}}^{\frac{1}{2}} \boldsymbol{H}_{\mathsf{E}}^{*} \boldsymbol{H}_{\mathsf{E}} \boldsymbol{K}_{\boldsymbol{x}}^{\frac{1}{2}})^{-1} \right|$$

Proof.

- $I(\boldsymbol{x};\boldsymbol{y}|\boldsymbol{z}) = h(\boldsymbol{y},\boldsymbol{z}) - h(\boldsymbol{z}) - h(\boldsymbol{w}_{\mathsf{R}})$
- optimality of Gaussian $\boldsymbol{x}$ analogous to MIMO secrecy capacity [Khisti-Wornell, '10]
- $h(\boldsymbol{y},\boldsymbol{z}) = \log(2\pi e)^{n_{\mathsf{R}}+n_{\mathsf{E}}} \det \begin{bmatrix} \boldsymbol{H}_{\mathsf{R}}\boldsymbol{K_x}\boldsymbol{H}_{\mathsf{R}}^{*} & \boldsymbol{H}_{\mathsf{R}}\boldsymbol{K_x}\boldsymbol{H}_{\mathsf{E}}^{*} \\ \boldsymbol{H}_{\mathsf{E}}\boldsymbol{K_x}\boldsymbol{H}_{\mathsf{R}}^{*} & \boldsymbol{H}_{\mathsf{E}}\boldsymbol{K_x}\boldsymbol{H}_{\mathsf{E}}^{*} \end{bmatrix}$
- use block determinant and matrix manipulation

$\square$

High-SNR secret-key capacity (I)

Proposition

The high-power secret-key capacity when $\boldsymbol{H}_\mathsf{E}$ has full column rank is

$$S(\infty) = \lim_{P\to\infty} S(P) = \sum_{i=1}^{s} \log(1+\sigma_i^2),$$

where $\sigma_1, \dots, \sigma_s$ are the generalized singular values of $(\boldsymbol{H}_\mathsf{R}, \boldsymbol{H}_\mathsf{E})$.

Proof of achievability.

We build $\{\boldsymbol{K}_{\boldsymbol{x}}(P)\}_{P\geq 0}$ such that $\lim_{P\to\infty} I(\boldsymbol{x};\boldsymbol{y}|\boldsymbol{z}) = \sum_{i=1}^{s} \log(1+\sigma_i^2)$

From the GSVD: $\boldsymbol{\Psi}_\mathsf{R}^* \boldsymbol{H}_\mathsf{R} \boldsymbol{V} = \begin{bmatrix} \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{D}_\mathsf{R} \end{bmatrix}$, $\boldsymbol{\Psi}_\mathsf{E}^* \boldsymbol{H}_\mathsf{E} \boldsymbol{V} = \begin{bmatrix} \boldsymbol{I} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{D}_\mathsf{E} \end{bmatrix}$

choose $\boldsymbol{x} = \boldsymbol{V} \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{t} \end{bmatrix}$ with $\lim_{P\to\infty} \lambda_{\mathsf{min}}(\boldsymbol{K}_{\boldsymbol{t}}) = \infty$ and $\mathrm{tr}(\boldsymbol{K}_{\boldsymbol{x}}(P)) \leq P$

$$I(\boldsymbol{x};\boldsymbol{y}|\boldsymbol{z}) = \log \frac{|\boldsymbol{I} + (\boldsymbol{D}_\mathsf{R}^*\boldsymbol{D}_\mathsf{R} + \boldsymbol{D}_\mathsf{E}^*\boldsymbol{D}_\mathsf{E})^{-1}\boldsymbol{K}_{\boldsymbol{t}}^{-1}|}{|\boldsymbol{I} + (\boldsymbol{D}_\mathsf{E}^*\boldsymbol{D}_\mathsf{E})^{-1}\boldsymbol{K}_{\boldsymbol{t}}^{-1}|} + \log \frac{|\boldsymbol{D}_\mathsf{R}^*\boldsymbol{D}_\mathsf{R} + \boldsymbol{D}_\mathsf{E}^*\boldsymbol{D}_\mathsf{E}|}{|\boldsymbol{D}_\mathsf{E}^*\boldsymbol{D}_\mathsf{E}|}$$

$\square$

High-SNR secret-key capacity (I)

Proposition

The high-power secret-key capacity when $\boldsymbol{H}_{\mathsf{E}}$ has full column rank is

$$S(\infty) = \lim_{P\to\infty} S(P) = \sum_{i=1}^{s} \log(1+\sigma_i^2),$$

where $\sigma_1, \ldots, \sigma_s$ are the generalized singular values of $(\boldsymbol{H}_{\mathsf{R}}, \boldsymbol{H}_{\mathsf{E}})$.

Proof of the converse.

We prove that $\forall \boldsymbol{x}$, it is $I(\boldsymbol{x};\boldsymbol{y}|\boldsymbol{z}) \leq \sum_{i=1}^{s} \log(1+\sigma_i^2)$.

$$\begin{aligned} I(\boldsymbol{x};\boldsymbol{y}|\boldsymbol{z}) &= h(\boldsymbol{y}|\boldsymbol{z}) - h(\boldsymbol{y}|\boldsymbol{x}) = \min_{\boldsymbol{\Theta}} h(\boldsymbol{y}-\boldsymbol{\Theta}\boldsymbol{z}) - h(\boldsymbol{w}_{\mathsf{R}}) \quad \text{(LMMSE)} \\ &\leq h(\boldsymbol{y}-\boldsymbol{H}_{\mathsf{R}}\boldsymbol{H}_{\mathsf{E}}^{\dagger}\boldsymbol{z}) - h(\boldsymbol{w}_{\mathsf{R}}) = h(\boldsymbol{w}_{\mathsf{R}}-\boldsymbol{H}_{\mathsf{R}}\boldsymbol{H}_{\mathsf{E}}^{\dagger}\boldsymbol{w}_{\mathsf{E}}) - h(\boldsymbol{w}_{\mathsf{R}}) \\ &= \log|\boldsymbol{I}+\boldsymbol{H}_{\mathsf{R}}(\boldsymbol{H}_{\mathsf{E}}^{*}\boldsymbol{H}_{\mathsf{E}})^{-1}\boldsymbol{H}_{\mathsf{R}}^{*}| = \textstyle\sum_{i=1}^{s}\log(1+\sigma_i^2) \end{aligned}$$

Hence $S(P) \leq \sum_{i=1}^{s} \log(1+\sigma_i^2)$, for all P. $\square$

High-SNR secret-key capacity (II)

Corollary

If $\boldsymbol{H}_{\mathsf{R}}$ has full column rank, $S(\infty)$ is achieved by any Gaussian $\boldsymbol{x}$ such that $\lim_{P\to\infty} \lambda_{\min}(\boldsymbol{K}_{\boldsymbol{x}}) = \infty$.

Remark 1 If $\operatorname{rank}(\boldsymbol{H}_{\mathsf{E}}) < n_{\mathsf{T}}$, Alice can transmit information in $\mathcal{N}(\boldsymbol{H}_{\mathsf{E}})$

$$S(P) = \sum_{i=1}^{s} \log(1+\sigma_i^2) + \log\left|\boldsymbol{I} + \frac{P}{p}(\boldsymbol{H}_{\mathsf{R}}^*\boldsymbol{H}_{\mathsf{R}} + \boldsymbol{H}_{\mathsf{E}}^*\boldsymbol{H}_{\mathsf{E}})\boldsymbol{H}_{\mathsf{E}}^\sharp\right| - o(1),$$

$\boldsymbol{H}_{\mathsf{E}}^\sharp$ is the projection onto $\mathcal{N}(\boldsymbol{H}_{\mathsf{E}})$ and $p = \dim \mathcal{N}(\boldsymbol{H}_{\mathsf{R}})^\perp \cap \mathcal{N}(\boldsymbol{H}_{\mathsf{E}})$.

Remark 2 In contrast with secrecy capacity, the high SNR secret-key capacity is achieved by transmitting along all the directions obtained with the GSVD, including those with $\sigma_i \leq 1$.

Low-SNR secret-key capacity

Proposition

$$\dot{S}(0) = \frac{1}{\ln 2} \lambda_{\mathsf{max}}(\boldsymbol{H}_{\mathsf{R}}^* \boldsymbol{H}_{\mathsf{R}})$$

and it is achieved by beamforming along the corresponding eigenspace.

$$\ddot{S}(0) = -\min_{\{\alpha_i\}} \frac{1}{\ln 2} \sum_{i=1}^{\ell} \alpha_i^2 \left(\lambda_{\mathsf{max}}(\boldsymbol{H}_{\mathsf{R}}^* \boldsymbol{H}_{\mathsf{R}})^2 + 2\lambda_{\mathsf{max}}(\boldsymbol{H}_{\mathsf{R}}^* \boldsymbol{H}_{\mathsf{R}}) \|\boldsymbol{H}_{\mathsf{E}} \boldsymbol{u}_i\|^2 \right),$$

where $\boldsymbol{u}_i$ form an orthonormal basis of the $\lambda_{\mathsf{max}}(\boldsymbol{H}_{\mathsf{R}}^* \boldsymbol{H}_{\mathsf{R}})$ eigenspace and $\sum \alpha_i = 1$. It is achieved by $\boldsymbol{K}_{\mathsf{x}} = P \sum_{i=1}^{\ell} \alpha_i \boldsymbol{u}_i \boldsymbol{u}_i^*$

Second-order Taylor expansion as $P \to 0$:

$$S(P) = \dot{S}(0)P + \frac{\ddot{S}(0)}{2} P^2 + o(P^2)$$

Observe that the optimal signaling does not depend on the eavesdropper's channel and also achieves low-power, main channel capacity.

Numerical results for finite SNR

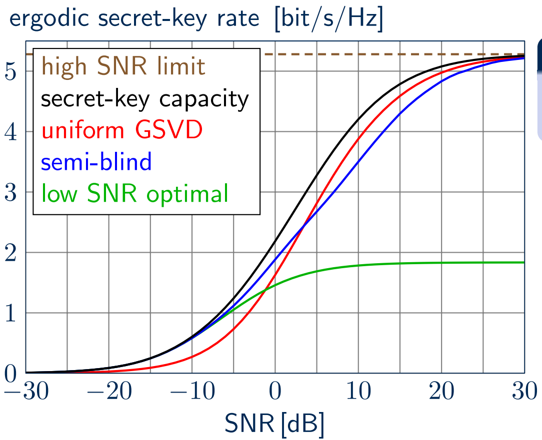

Parameters

$n_\mathsf{T} = n_\mathsf{R} = n_\mathsf{E} = 3$

1000 channel realizations

- Secret-key capacity: computed numerically via KKT conditions
- Semi-blind: input that achieves capacity of $\boldsymbol{H}_\mathsf{R}$, regardless of $\boldsymbol{H}_\mathsf{E}$

The semi-blind solution is optimal at low and high SNR and nearly optimal in intermediate power regimes.

Blind key agreement: outage analysis

Assume the transmitter:

- has perfect CSI on $\boldsymbol{H}_\mathsf{R}$
- has statistical CSI on $\boldsymbol{H}_\mathsf{E}$ (Rayleigh fading)
- uses low-power optimal input

$$R = \log\left(1 + \frac{P\lambda_{\max}(\boldsymbol{H}_\mathsf{R}^*\boldsymbol{H}_\mathsf{R})}{1 + P\|\boldsymbol{H}_\mathsf{E}\boldsymbol{u}_1\|^2}\right)$$

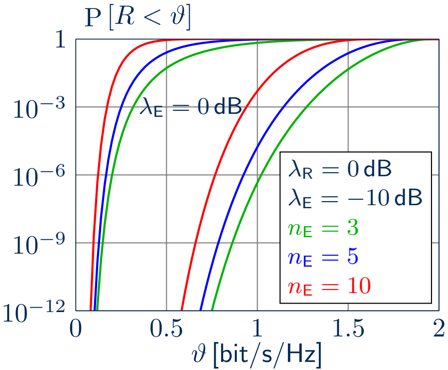

Outage probability

$$\mathrm{P}[R < \vartheta] = 1 - \frac{1}{(n_\mathsf{E}-1)!}\,\gamma\left(n_\mathsf{E}\,,\ \frac{\lambda_{\max}(\boldsymbol{H}_\mathsf{R}^*\boldsymbol{H}_\mathsf{R})}{2^\vartheta - 1} - \frac{1}{P}\right)$$

Conclusions

- We have derived closed-form expressions of the secret-key capacity in the high and low-power regimes.
- The low-power optimal signaling is independent from the eavesdropper's channel.
- We propose a semi-blind approach: the (unconstrained) capacity achieving input is optimal in the asymptotic regimes, and performs well in the intermediate regimes.
- We evaluate the secret-key rate outage probability to perform strictly blind key-sharing with statistical CSI about the eavesdropper's channel.

Outline

1. What is unconditional security?
2. Signal processing for unconditional secrecy
3. Signal processing for unconditionally secure key agreement
4. **Unconditionally secure authentication**
   - Universal hashing
   - Physical layer authentication for MIMO systems
   - Authentication based on channel estimation
   - Effective attack strategies

Outline

1. What is unconditional security?
2. Signal processing for unconditional secrecy
3. Signal processing for unconditionally secure key agreement
4. **Unconditionally secure authentication**
   - Universal hashing
   - Physical layer authentication for MIMO systems
   - Authentication based on channel estimation
   - Effective attack strategies

Unconditional authentication / integrity protection

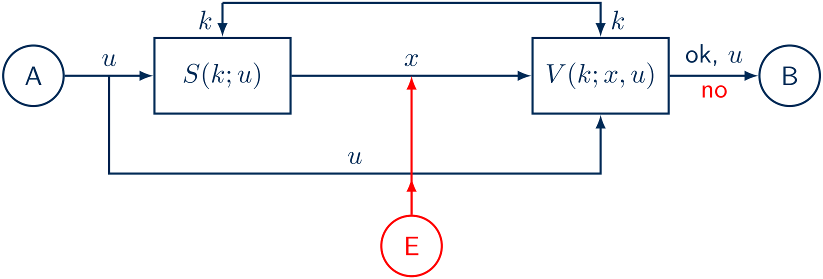

Unconditional authentication / integrity protection

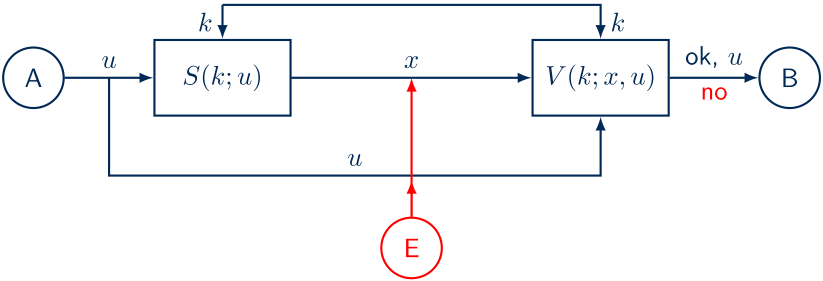

Kerchoff's-like Assumption

E knows:

- the functions $S(\cdot;\cdot)$, $V(\cdot;\cdot)$
- the distributions $p_u(\cdot)$, $p_k(\cdot)$

Non forgeability of x is only based on hiding the key k

Unconditional authentication / integrity protection

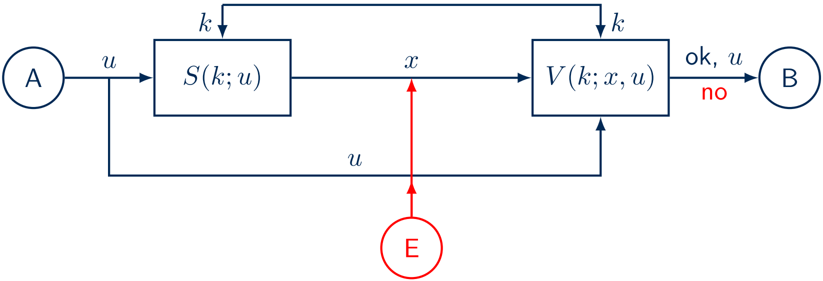

Kerchoff's-like Assumption

E knows:

- the functions $S(\cdot;\cdot)$, $V(\cdot;\cdot)$
- the distributions $p_u(\cdot)$, $p_k(\cdot)$

Non forgeability of x is only based on hiding the key k

Unconditionally secure authentication

Ask for $p_{\mathsf{MD}} < \varepsilon$, while $p_{\mathsf{FA}} \to 0$
$I(k; x|u) \geq -\log\varepsilon$,
$H(k|u, x) \geq -\log\varepsilon$

It requires $H(k) \geq -2\log\varepsilon$

Unconditionally secure authentication

Unconditionally secure authentication can be obtained with a One Time Pad

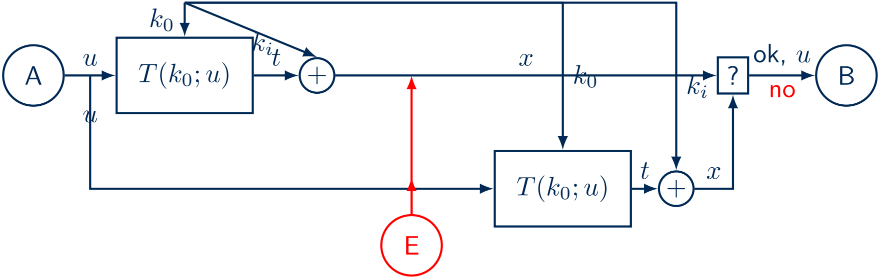

Unconditionally secure integrity protection

Need $\{T_k(u)\}_{k \in \mathcal{K}}$ to be a class of universal$_2$ hashing functions for some parameter ε, that is

1. $T_k \; : \mathcal{U} \to \mathcal{T}, \forall k \in \mathcal{K}$

The lowest possible (ideal) value of ε is $1/\left|\mathcal{T}\right|$.

Unconditionally secure integrity protection

Need $\{T_k(u)\}_{k\in\mathcal{K}}$ to be a class of universal$_2$ hashing functions for some parameter ε, that is

1. $T_k \ : \ \mathcal{U} \to \mathcal{T}, \forall k \in \mathcal{K}$
2. (uniform mapping) $\forall u \in \mathcal{U}, t \in \mathcal{T}$, it must be $|\mathcal{K}_{u\to t}| \leq \varepsilon\, |K|$, where

$$\mathcal{K}_{u\to t} = \{k \in \mathcal{K} \ : \ T_k(u) = t\}$$

The lowest possible (ideal) value of ε is $1/\,|\mathcal{T}|$.

Unconditionally secure integrity protection

Need $\{T_k(u)\}_{k\in\mathcal{K}}$ to be a class of universal$_2$ hashing functions for some parameter ε, that is

1. $T_k\ :\ \mathcal{U}\to\mathcal{T},\ \forall k\in\mathcal{K}$
2. (uniform mapping) $\forall u\in\mathcal{U}, t\in\mathcal{T}$, it must be $|\mathcal{K}_{u\to t}|\leq\varepsilon\,|K|$, where

$$\mathcal{K}_{u\to t}=\{k\in\mathcal{K}\ :\ T_k(u)=t\}$$

3. (uniform collisions) $\forall u_1,u_2\in\mathcal{U},$, it must be $|\mathcal{K}_{u_1u_2}|\leq\varepsilon\,|K|$, where

$$\mathcal{K}_{u_1u_2}=\{k\in\mathcal{K}\ :\ T_k(u_1)=T_k(u_2)\}$$

The lowest possible (ideal) value of ε is $1/\,|\mathcal{T}|$.

Classes of universal hashing functions

Example

All the functions The class of all the functions mapping $\mathcal{U}$ to $\mathcal{T}$ is universal with $\varepsilon = 1/|\mathcal{T}|$. Its cardinality is $|\mathcal{K}| = |\mathcal{T}|^{|\mathcal{U}|}$

Classes of universal hashing functions

Example

All the functions The class of all the functions mapping $\mathcal{U}$ to $\mathcal{T}$ is universal with $\varepsilon = 1/|\mathcal{T}|$. Its cardinality is $|\mathcal{K}| = |\mathcal{T}|^{|\mathcal{U}|}$

Example

All the linear functions (matrices) If $\mathcal{U} = \mathbb{F}^{\ell_u}$, $\mathcal{T} = \mathbb{F}^{\ell_t}$, with $\mathbb{F}$ a finite field, the class of all the matrices $\mathbb{F}^{\ell_t \times \ell_u}$ is universal with $\varepsilon = 1/|\mathcal{T}|$. Its cardinality is $|\mathcal{K}| = |\mathcal{T}| \cdot |\mathcal{U}|$

Classes of universal hashing functions

Example

All the functions The class of all the functions mapping $\mathcal{U}$ to $\mathcal{T}$ is universal with $\varepsilon = 1/|\mathcal{T}|$. Its cardinality is $|\mathcal{K}| = |\mathcal{T}|^{|\mathcal{U}|}$

Example

All the linear functions (matrices) If $\mathcal{U} = \mathbb{F}^{\ell_u}$, $\mathcal{T} = \mathbb{F}^{\ell_t}$, with $\mathbb{F}$ a finite field, the class of all the matrices $\mathbb{F}^{\ell_t \times \ell_u}$ is universal with $\varepsilon = 1/|\mathcal{T}|$. Its cardinality is $|\mathcal{K}| = |\mathcal{T}| \cdot |\mathcal{U}|$

Example

All the Toeplitz matrices If $\mathcal{U} = \mathbb{F}^{\ell_u}$, $\mathcal{T} = \mathbb{F}^{\ell_t}$, with $\mathbb{F}$ a finite field, the class of all the Toeplitz matrices in $\mathbb{F}^{\ell_t \times \ell_u}$ is universal with $\varepsilon = 1/|\mathcal{T}|$. Its cardinality is $|\mathcal{K}| = |\mathbb{F}|^{\ell_t + \ell_u - 1}$

Outline

1. What is unconditional security?
2. Signal processing for unconditional secrecy
3. Signal processing for unconditionally secure key agreement
4. **Unconditionally secure authentication**
   - Universal hashing
   - **Physical layer authentication for MIMO systems**
   - Authentication based on channel estimation
   - Effective attack strategies

Motivations

The problem of message authentication is certainly, together with that of message confidentiality, one of the most common tasks in information security.

Classical solution is cryptographic: hash and sign protocols.

Physical Layer Secrecy already enjoys a rich literature. It is not so for authentication.

What could be the purpose of PHY authentication?

Provide an outer defense, to reduce the amount of attacks that higher layers must repel?

Previous work

Information theory results

- With secret key and noiseless transmission [Maurer, '00]
- Allowing for distortion of the message [Martinian *et al.*, '05]
- Introducing noisy channel for key and message [Lai *et al.*, '09]

Device identification schemes

- Pre-shared key used in modulation
- Wireless fingerprinting

Channel-based schemes

- With spatial diversity from cooperating receivers [Chen *et al.*, '07]
- Diversity from estimation of a wide band channel
  [Xiao *et al.*, '06–'10], but no attack at the physical layer...

System model

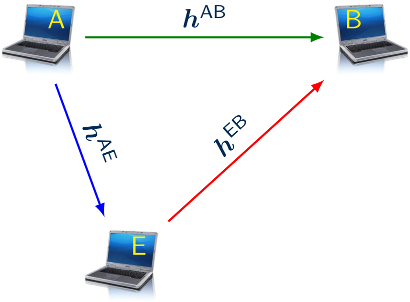

$\boldsymbol{h} = [h_0, \dots, h_{N-1}]$:
channel fading coefficients (e.g., impulse response, frequency response, channel matrix entries)

System model

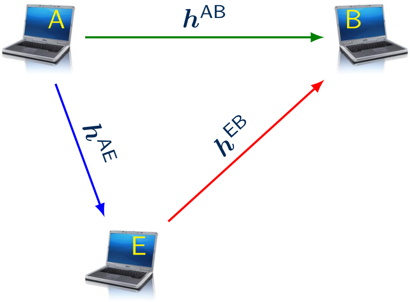

$\boldsymbol{h} = [h_0, \ldots, h_{N-1}]$:
channel fading coefficients (e.g., impulse response, frequency response, channel matrix entries)

channel statistics

complex, jointly Gaussian, circularly symmetric

$$\boldsymbol{h}^{(AB)} \sim \mathcal{CN}(\boldsymbol{0}_{\nu\times 1}, \boldsymbol{R}^{(AB)})$$
$$\boldsymbol{h}^{(AE)} \sim \mathcal{CN}(\boldsymbol{0}_{\mu\times 1}, \boldsymbol{R}^{(AE)})$$
$$\boldsymbol{h}^{(EB)} \sim \mathcal{CN}(\boldsymbol{0}_{\varphi\times 1}, \boldsymbol{R}^{(EB)})$$

channel reciprocity

$$\mathrm{E}\left[\boldsymbol{h}^{(AB)}\boldsymbol{h}^{(AE)*}\right] = \boldsymbol{R}^{(AB,AE)}$$
$$\mathrm{E}\left[\boldsymbol{h}^{(AB)}\boldsymbol{h}^{(EB)*}\right] = \boldsymbol{R}^{(AB,EB)}$$
$$\mathrm{E}\left[\boldsymbol{h}^{(AE)}\boldsymbol{h}^{(EB)*}\right] = \boldsymbol{R}^{(AE,EB)}$$

Outline

Authentication scheme [Xiao *et al.*, '08]

Phase I: training

- A (securely) sends a training sequence to B
- B obtains a (reliable) ML estimate $\hat{\boldsymbol{h}}^{\mathsf{AB}}$ of the channel

$$\hat{\boldsymbol{h}}^{\mathsf{AB}} = \boldsymbol{h}^{\mathsf{AB}} + \boldsymbol{w}^{\mathsf{I}} \quad , \qquad \boldsymbol{w}^{\mathsf{I}} \sim \mathcal{CN}(\mathbf{0}, \sigma_{\mathsf{I}}^2 \boldsymbol{I})$$

Phase II: hypothesis testing

For every received packet, B estimates the channel response $\hat{\boldsymbol{h}}(t)$ and checks it against the hypotheses

$$\begin{aligned} \text{(authentic) } \mathcal{H}_0 &: \hat{\boldsymbol{h}}(t) = \boldsymbol{h}^{\mathsf{AB}} + \boldsymbol{w}^{\mathsf{I}}(t) \quad , \qquad \boldsymbol{w}^{\mathsf{II}}(t) \sim \mathcal{CN}(\mathbf{0}, \sigma_{\mathsf{II}}^2 \boldsymbol{I}) \\ \text{(forged) } \mathcal{H}_1 &: \hat{\boldsymbol{h}}(t) = \boldsymbol{g}(t) + \boldsymbol{w}^{\mathsf{I}}(t) \quad , \qquad \boldsymbol{g}(t) \text{ arbitrary} \end{aligned}$$

Generalized likelihood ratio test (GLRT)

Formulation

- log likelihood ratio: $\Psi = \log \dfrac{f_{\hat{\boldsymbol{h}}|\mathcal{H}_1,\boldsymbol{g}}(\hat{\boldsymbol{h}}|\hat{\boldsymbol{h}})}{f_{\hat{\boldsymbol{h}}|\mathcal{H}_0}(\hat{\boldsymbol{h}})} \propto \dfrac{2}{\sigma^2}\sum_{n=0}^{\nu-1}\left|\hat{h}_n - \hat{h}_n^{(\mathrm{AB})}\right|^2$
- compare with a threshold : $\begin{cases} \Psi \leq \vartheta : & \text{decide for } \mathcal{H}_0\,, \\ \Psi > \vartheta : & \text{decide for } \mathcal{H}_1\,. \end{cases}$

Probability of False Alarm and Missed Detection

Ψ is a chi-square variable

$$P_{\mathsf{FA}} = \mathrm{P}\left[\Psi > \vartheta \,|\, \mathcal{H}_0\right] = 1 - F_{\chi^2,0}(\vartheta) P_{\mathsf{MD}} = \mathrm{P}\left[\Psi < \vartheta \,|\, \mathcal{H}_1\right] = F_{\chi^2,\beta}(\vartheta)$$

If we fix a target P_{FA}, we get $P_{\mathrm{MD}}(\beta) = F_{\chi^2,\beta}\left(F^{-1}_{\chi^2,0}\left(1 - P_{\mathrm{FA}}\right)\right)$

Outline

1. What is unconditional security?
2. Signal processing for unconditional secrecy
3. Signal processing for unconditionally secure key agreement
4. **Unconditionally secure authentication**
   - Universal hashing
   - Physical layer authentication for MIMO systems
   - Authentication based on channel estimation
   - **Effective attack strategies**

Effective attack strategies

Knowledge assumptions

We assume that E has estimated her channels to A and B

$$\hat{\boldsymbol{h}}^{\mathsf{AE}} = \boldsymbol{h}^{\mathsf{AE}} + \boldsymbol{w}^{\mathsf{AE}} \quad , \quad \hat{\boldsymbol{h}}^{\mathsf{EB}} = \boldsymbol{h}^{\mathsf{EB}} + \boldsymbol{w}^{\mathsf{EB}}$$

with $\boldsymbol{w}^{\mathsf{AE}} \sim \mathcal{CN}(\mathbf{0}, \sigma^2_{\mathsf{AE}}\boldsymbol{I}), \boldsymbol{w}^{\mathsf{EB}} \sim \mathcal{CN}(\mathbf{0}, \sigma^2_{\mathsf{EB}}\boldsymbol{I})$

Optimal strategy for a single attack

If the horizon of E is a single attack, her optimal strategy is the ML estimate of $\hat{\boldsymbol{h}}^{\mathsf{AB}}$

$$\bar{\boldsymbol{g}} = -\left([\boldsymbol{R}^{-1}]_{11}\right)^{-1}\left([\boldsymbol{R}^{-1}]_{12}\hat{\boldsymbol{h}}^{(\mathrm{AE})} + [\boldsymbol{R}^{-1}]_{13}\hat{\boldsymbol{h}}^{(\mathrm{EB})}\right)$$

with $\boldsymbol{R}$ the covariance matrix of $[\hat{\boldsymbol{h}}^{\mathsf{AB}}, \hat{\boldsymbol{h}}^{\mathsf{AE}}, \hat{\boldsymbol{h}}^{\mathsf{EB}}]$

A repeated attack strategy

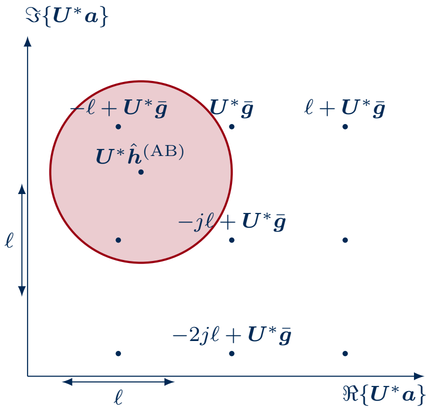

Sequential guessing problem. . .

. . . with distortion and lies [Arikan-Merhav, '98], on a continuous space.

For the ease of tractability

- consider a discrete set $\mathcal{Z}$ of regularly spaced points
- at any attempt τ, choose the next best guess among them, given the previous failed attempts

$$\bar{\boldsymbol{g}}(t) = \arg\max_{\boldsymbol{z}\in\mathcal{Z}} \mathrm{P}\left[\boldsymbol{z} + \boldsymbol{w}^{\mathsf{II}}(t) \in \mathcal{S} \,|\, \cap_{t'=0}^{t-1} \left\{\bar{\boldsymbol{g}}(t') + \boldsymbol{w}^{\mathsf{II}}(t') \notin \mathcal{S}\right\}\right]$$

A repeated attack strategy

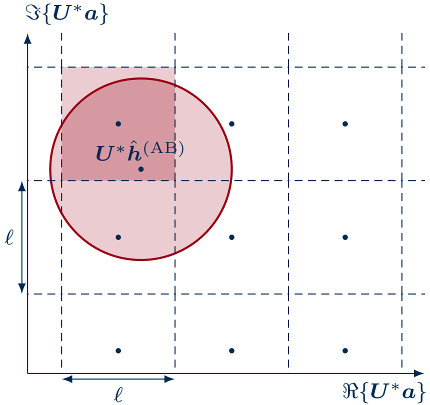

Evaluation of probabilities

As a further simplification

- partition $\mathbb{C}^\nu$ into ν-dimensional cubes centered in Z
- replace $\mathbb{S}$ with the cube in which $\hat{h}^{\mathrm{AB}}$ lies

It becomes a discrete guessing problem without distortion.

$$q(z|a) = \mathrm{P}\left[\hat{h}(t) \in \mathcal{R}(z) |\, \bar{g}(t) = a\right]$$

$$p(z) = \mathrm{P}\left[\hat{h}^{\mathrm{AB}} \in \mathcal{R}(z) | \hat{h}^{\mathrm{AE}}, \hat{h}^{\mathrm{EB}}\right]$$

$$\bar{g}(t) = \arg\max_{a} \sum_{z \in \mathcal{Z}} p(z) q(z|a) \prod_{t'=1}^{t-1} (1 - q(z|\bar{g}(t')))$$

Bound on the achievable $(\mathrm{E}[p_{\mathsf{FA}}], \mathrm{E}[p_{\mathsf{MD}}])$ region

Let $d(x,y) = x \log \frac{x}{1-y} + (1-x) \log \frac{1-x}{y}$
Then, for any authentication procedure that makes use of $\hat{\boldsymbol{h}}^{\mathsf{AB}}, \hat{\boldsymbol{h}}$,

$$d(\mathrm{E}[p_{\mathsf{FA}}], \mathrm{E}[p_{\mathsf{MD}}]) \leq D\left(p_{\hat{\boldsymbol{h}},\hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_0} \,||\, p_{\hat{\boldsymbol{h}},\hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_1}\right)$$

$$d(\mathrm{E}[p_{\mathsf{MD}}], \mathrm{E}[p_{\mathsf{FA}}]) \leq D\left(p_{\hat{\boldsymbol{h}},\hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_1} \,||\, p_{\hat{\boldsymbol{h}},\hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_0}\right)$$

The above outer bounds depend on the attack strategy $f_{\hat{\boldsymbol{h}}|\mathcal{H}_1,\hat{\boldsymbol{h}}^{(\mathrm{AE})},\hat{\boldsymbol{h}}^{(\mathrm{EB})}}$ as under $\mathcal{H}_1$, $\hat{\boldsymbol{h}}$ is independent of $\hat{\boldsymbol{h}}^{(\mathrm{AB})}$, when conditioned on $\hat{\boldsymbol{h}}^{(\mathrm{AE})}, \hat{\boldsymbol{h}}^{(\mathrm{EB})}$.
We consider
$f_{\hat{\boldsymbol{h}}|\mathcal{H}_1,\hat{\boldsymbol{h}}^{(\mathrm{AE})},\hat{\boldsymbol{h}}^{(\mathrm{EB})}} = f_{\hat{\boldsymbol{h}}|\mathcal{H}_0,\hat{\boldsymbol{h}}^{(\mathrm{AE})},\hat{\boldsymbol{h}}^{(\mathrm{EB})}}$

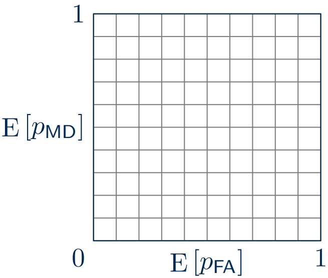

Bound on the achievable $\left(\mathrm{E}\left[p_{\mathsf{FA}}\right], \mathrm{E}\left[p_{\mathsf{MD}}\right]\right)$ region

Let $d(x,y) = x \log \frac{x}{1-y} + (1-x) \log \frac{1-x}{y}$
Then, for any authentication procedure that makes use of $\hat{\boldsymbol{h}}^{\mathsf{AB}}, \hat{\boldsymbol{h}}$,

$$d\left(\mathrm{E}\left[p_{\mathsf{FA}}\right], \mathrm{E}\left[p_{\mathsf{MD}}\right]\right) \leq D\left(p_{\hat{\boldsymbol{h}},\hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_0} \,||\, p_{\hat{\boldsymbol{h}},\hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_1}\right)$$

$$d\left(\mathrm{E}\left[p_{\mathsf{MD}}\right], \mathrm{E}\left[p_{\mathsf{FA}}\right]\right) \leq D\left(p_{\hat{\boldsymbol{h}},\hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_1} \,||\, p_{\hat{\boldsymbol{h}},\hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_0}\right)$$

The above outer bounds depend on the attack strategy $f_{\hat{\boldsymbol{h}}|\mathcal{H}_1,\hat{\boldsymbol{h}}^{(\mathrm{AE})},\hat{\boldsymbol{h}}^{(\mathrm{EB})}}$ as under $\mathcal{H}_1$, $\hat{\boldsymbol{h}}$ is independent of $\hat{\boldsymbol{h}}^{(\mathrm{AB})}$, when conditioned on $\hat{\boldsymbol{h}}^{(\mathrm{AE})}, \hat{\boldsymbol{h}}^{(\mathrm{EB})}$.
We consider
$f_{\hat{\boldsymbol{h}}|\mathcal{H}_1,\hat{\boldsymbol{h}}^{(\mathrm{AE})},\hat{\boldsymbol{h}}^{(\mathrm{EB})}} = f_{\hat{\boldsymbol{h}}|\mathcal{H}_0,\hat{\boldsymbol{h}}^{(\mathrm{AE})},\hat{\boldsymbol{h}}^{(\mathrm{EB})}}$

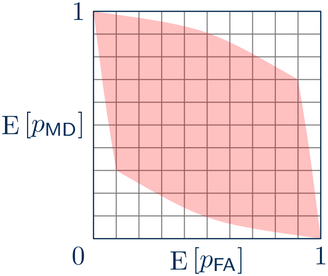

Bound on the achievable $(\mathrm{E}[p_{\mathsf{FA}}], \mathrm{E}[p_{\mathsf{MD}}])$ region

Let $d(x,y) = x \log \frac{x}{1-y} + (1-x) \log \frac{1-x}{y}$
Then, for any authentication procedure that makes use of $\hat{\boldsymbol{h}}^{\mathsf{AB}}, \hat{\boldsymbol{h}}$,

$$d(\mathrm{E}[p_{\mathsf{FA}}], \mathrm{E}[p_{\mathsf{MD}}]) \leq D\left(p_{\hat{\boldsymbol{h}}, \hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_0} \,||\, p_{\hat{\boldsymbol{h}}, \hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_1}\right)$$
$$d(\mathrm{E}[p_{\mathsf{MD}}], \mathrm{E}[p_{\mathsf{FA}}]) \leq D\left(p_{\hat{\boldsymbol{h}}, \hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_1} \,||\, p_{\hat{\boldsymbol{h}}, \hat{\boldsymbol{h}}^{(\mathrm{AB})}|\mathcal{H}_0}\right)$$

The above outer bounds depend on the attack strategy $f_{\hat{\boldsymbol{h}}|\mathcal{H}_1, \hat{\boldsymbol{h}}^{(\mathrm{AE})}, \hat{\boldsymbol{h}}^{(\mathrm{EB})}}$ as under $\mathcal{H}_1$, $\hat{\boldsymbol{h}}$ is independent of $\hat{\boldsymbol{h}}^{(\mathrm{AB})}$, when conditioned on $\hat{\boldsymbol{h}}^{(\mathrm{AE})}, \hat{\boldsymbol{h}}^{(\mathrm{EB})}$.
We consider
$f_{\hat{\boldsymbol{h}}|\mathcal{H}_1, \hat{\boldsymbol{h}}^{(\mathrm{AE})}, \hat{\boldsymbol{h}}^{(\mathrm{EB})}} = f_{\hat{\boldsymbol{h}}|\mathcal{H}_0, \hat{\boldsymbol{h}}^{(\mathrm{AE})}, \hat{\boldsymbol{h}}^{(\mathrm{EB})}}$

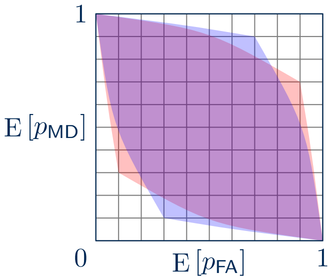

Average P_{MD} vs channels correlation

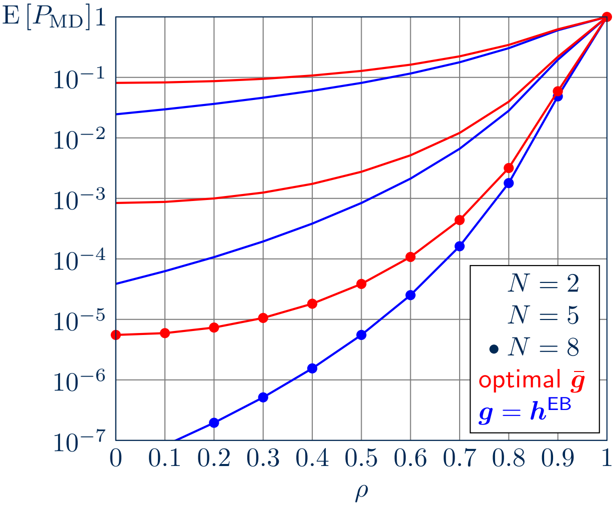

Parameters

OFDM scenario

N iid subcarriers

Rayleigh fading

$$\mathsf{SNR}^{\mathsf{I}} = 15\,\mathrm{dB}$$
$$\mathsf{SNR}^{\mathsf{II}} = 20\,\mathrm{dB}$$
$$P_{\mathsf{FA}} = 10^{-4}$$

CDF of first success for multiple attack strategy

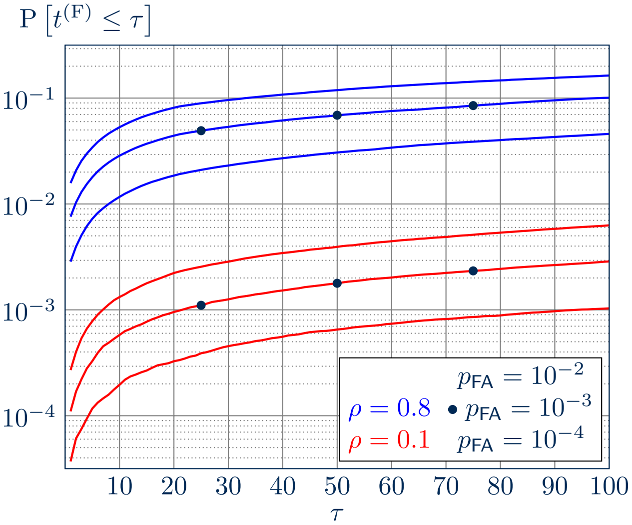

Parameters

OFDM scenario
$N = 5$ iid subcarriers
Rayleigh fading

$\mathsf{SNR}^{\mathsf{I}} = 15\,\mathrm{dB}$
$\mathsf{SNR}^{\mathsf{II}} \to \infty$

Conclusions

We have generalized the physical-layer technique of [Xiao *et al.*, '06–'10] to provide authentication between Alice and Bob, also assuming a more general model for the attack employed by Eve.

We provide the optimal strategy for Eve in the case of single attack and we perform an analytical computation of $\mathrm{E}[P_{\mathrm{MD}}]$ with respect to channel distribution.

Moreover, we formulate a suboptimal multiple attacks strategy for Eve consisting in a sequence of messages and channel guesses aiming to break authentication.

Numerical results confirm the merits of the considered method when diversity is sufficiently high and when correlation among channels is low.