Security for the Internet of Things: standards, problems, open issues

Nicola Laurenti

UNIVERSITÀ
DEGLI STUDI
DI PADOVA

DIPARTIMENTO
DI INGEGNERIA
DELL'INFORMAZIONE

Ph.D. Summer School in Information Engineering
Bressanone/Brixen, 4 July 2016

Outline

Outline

Security goals, threats, services and mechanisms

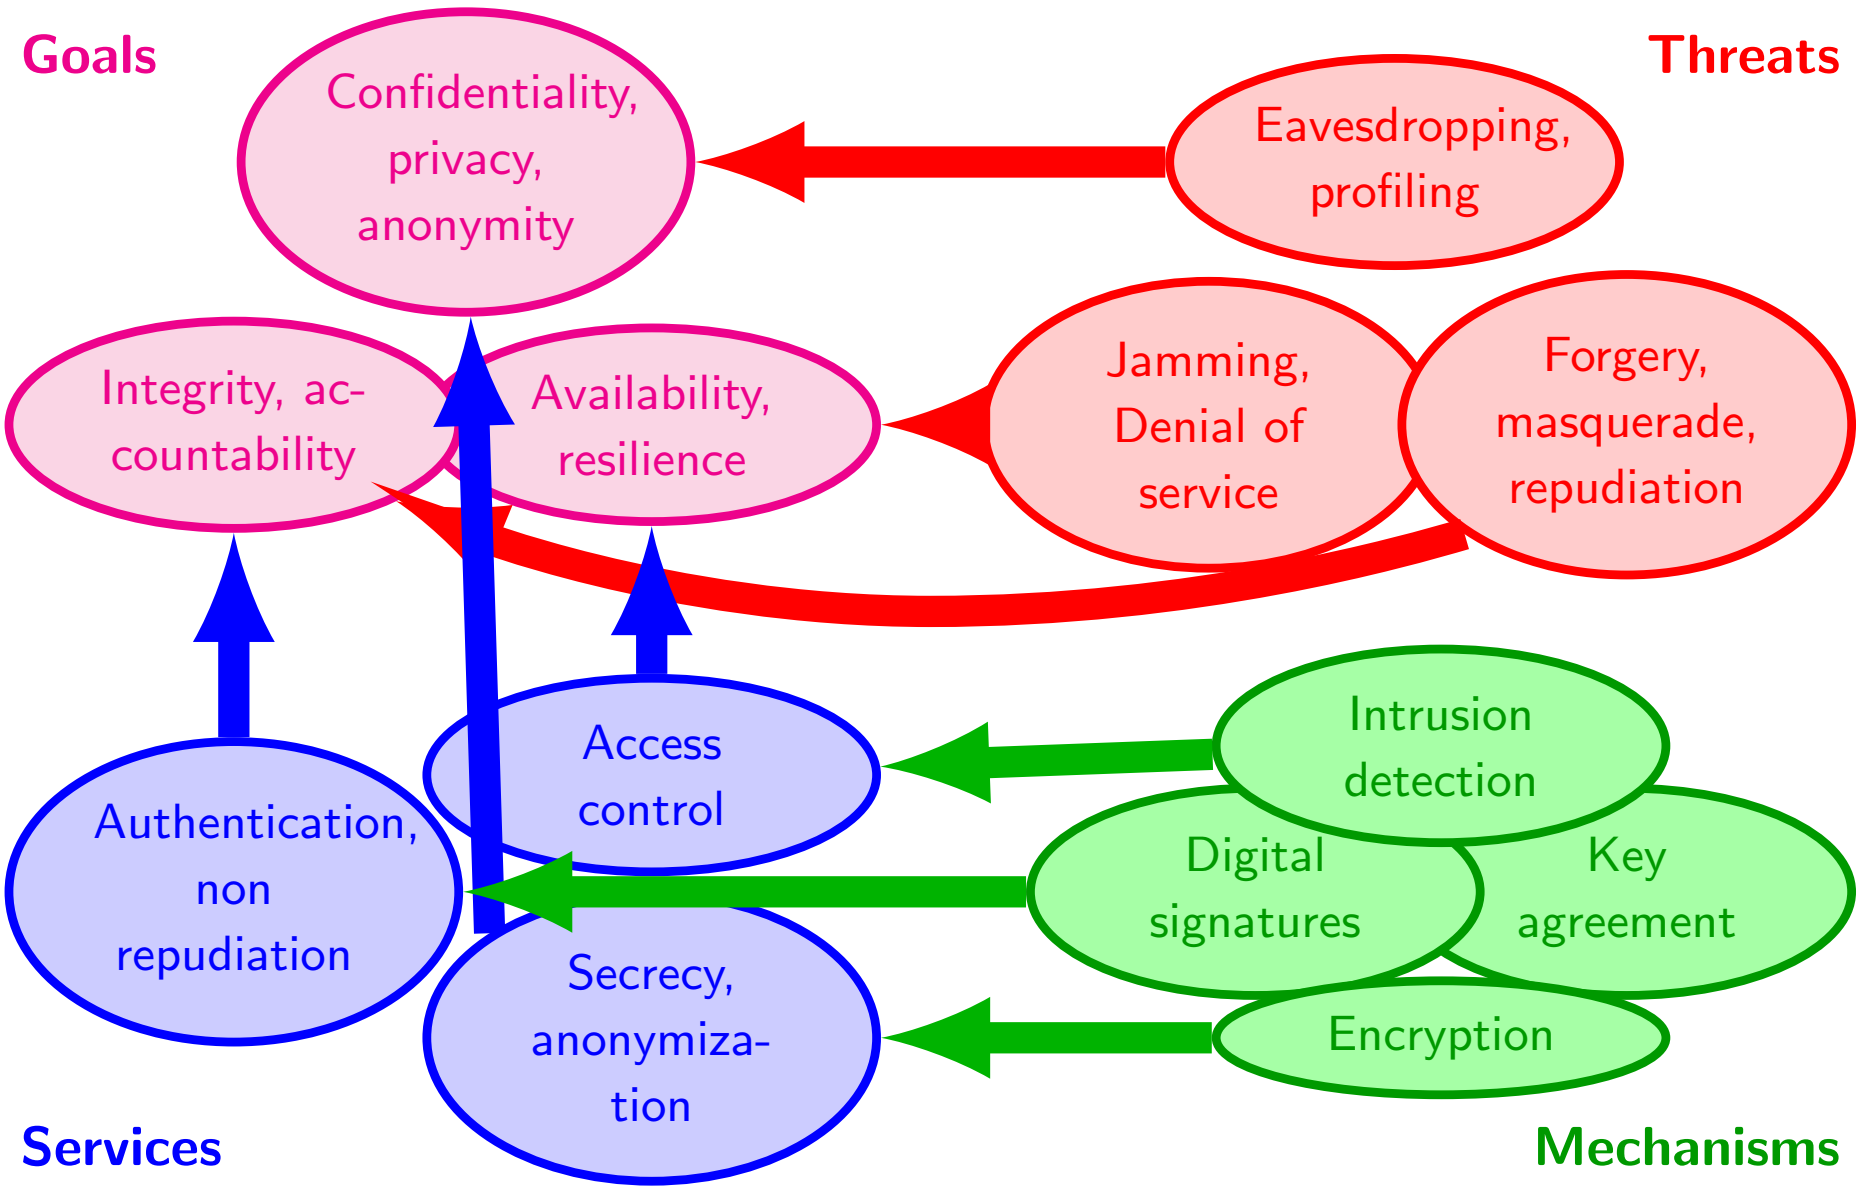

What security do we need?

FROM THE EDITORS

What Was Samsung Thinking?

Susan Landau
Associate Editor in Chief

In February 2015, the press discovered that if Samsung's Smart TV voice recognition system is activated, the television sends voice commands to Samsung and then to a third-party provider for processing. Any other conversations that are overheard are also sent. The company's user manual explains:[1]

trained to tackle. The real issue is how we employ devices. This is different from the concern about the voice channel being hacked.

Privacy Expectations and Social Context

People use smartphones to make a call, check their email, or look for directions. Although smartphones transmit users' location to the network, the assumption is that if a phone isn't actively being used, only location information is transmitted. (Yes, we've all seen the television show in which a cell phone is surreptitiously used to tape a conversation. That's ...e of the phone and is a violation of the ...otocol.) The same is true of ... fact, when users ask Glass to ...

SYN

Hacker shows off vulnerabilities of wireless insulin pumps

By ARUNDHATI PARMAR

Post a comment / 20 Shares / Mar 1, 2012 at 3:16 PM

Security and Privacy Vulnerabilities of In-Car Wireless Networks: A Tire Pressure Monitoring System Case Study

Ishtiaq Rouf[a], Rob Miller[b], Hossen Mustafa[a], Travis Taylor[a], Sangho Oh[b]
*Wenyuan Xu[a], Marco Gruteser[b], Wade Trappe[b], Ivan Seskar[b] **
[a] *Dept. of CSE, Univ. of South Carolina, Columbia, SC USA*
{*rouf, mustafah, taylort9, wyxu*}*@cse.sc.edu*
[b] *WINLAB, Rutgers Univ., Piscataway, NJ USA*
{*rdmiller, sangho, gruteser, trappe, seskar*}*@winlab.rutgers.edu*

Low energy, small storage security

Possible approaches

- provide "minimum" (i.e., "no") security
- expect electronic and battery performance increase
- harvest energy from environment

Outline

MAC layer: IEEE 802.15.4 security

Security modes

no security

authentication with AES-CBC-MAC

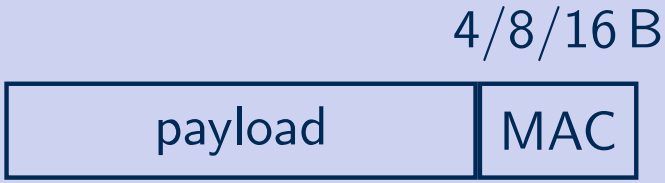

encryption with AES-CTR

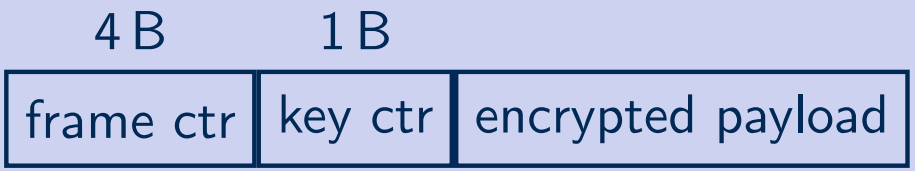

auth + encryption with AES-CCM

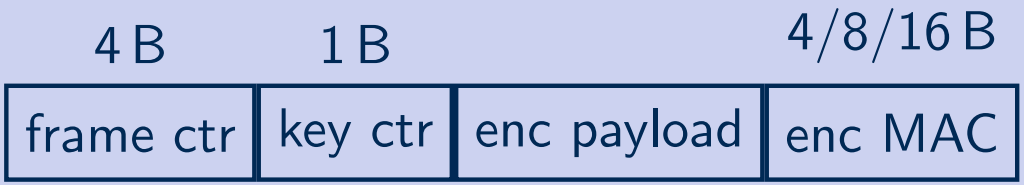

Network layer: 6LoWPAN security

No security mechanisms are defined in 6LoWPAN, but the following issues are identified:

secure routing: neighbour descovery and mesh routing may be vulnerable to threats, 802.15.4 AES-based security could provide mechanisms for securing routing

port number compression: the overload of ports may be exploited by applicatons not honoring the reserved sets, it should be integrity protected

RPL routing security

The RPL protocol supports security services for routing messages:

- Authentication and integrity protection through AES-CBC-MAC codes
- Secrecy through AES-CTR encryption
- Semantic security through counters and nonces
- Key management with pairwise symmetric keys, group keys, digital signatures

Application layer: CoAP and DTLS

CoAP adopts Datagram TLS to transparently apply security. This guarantees secrecy, authentication, integrity protection, non replay by adopting AES/CCM and nonces

Security modes

no security

preshared symmetric keys preprogrammed into the device

public key identity-based

certificate for public-key infrastructure

public key and certificate modes support elliptic curve cryptography

Outline

Broadcast authentication: the TESLA protocol

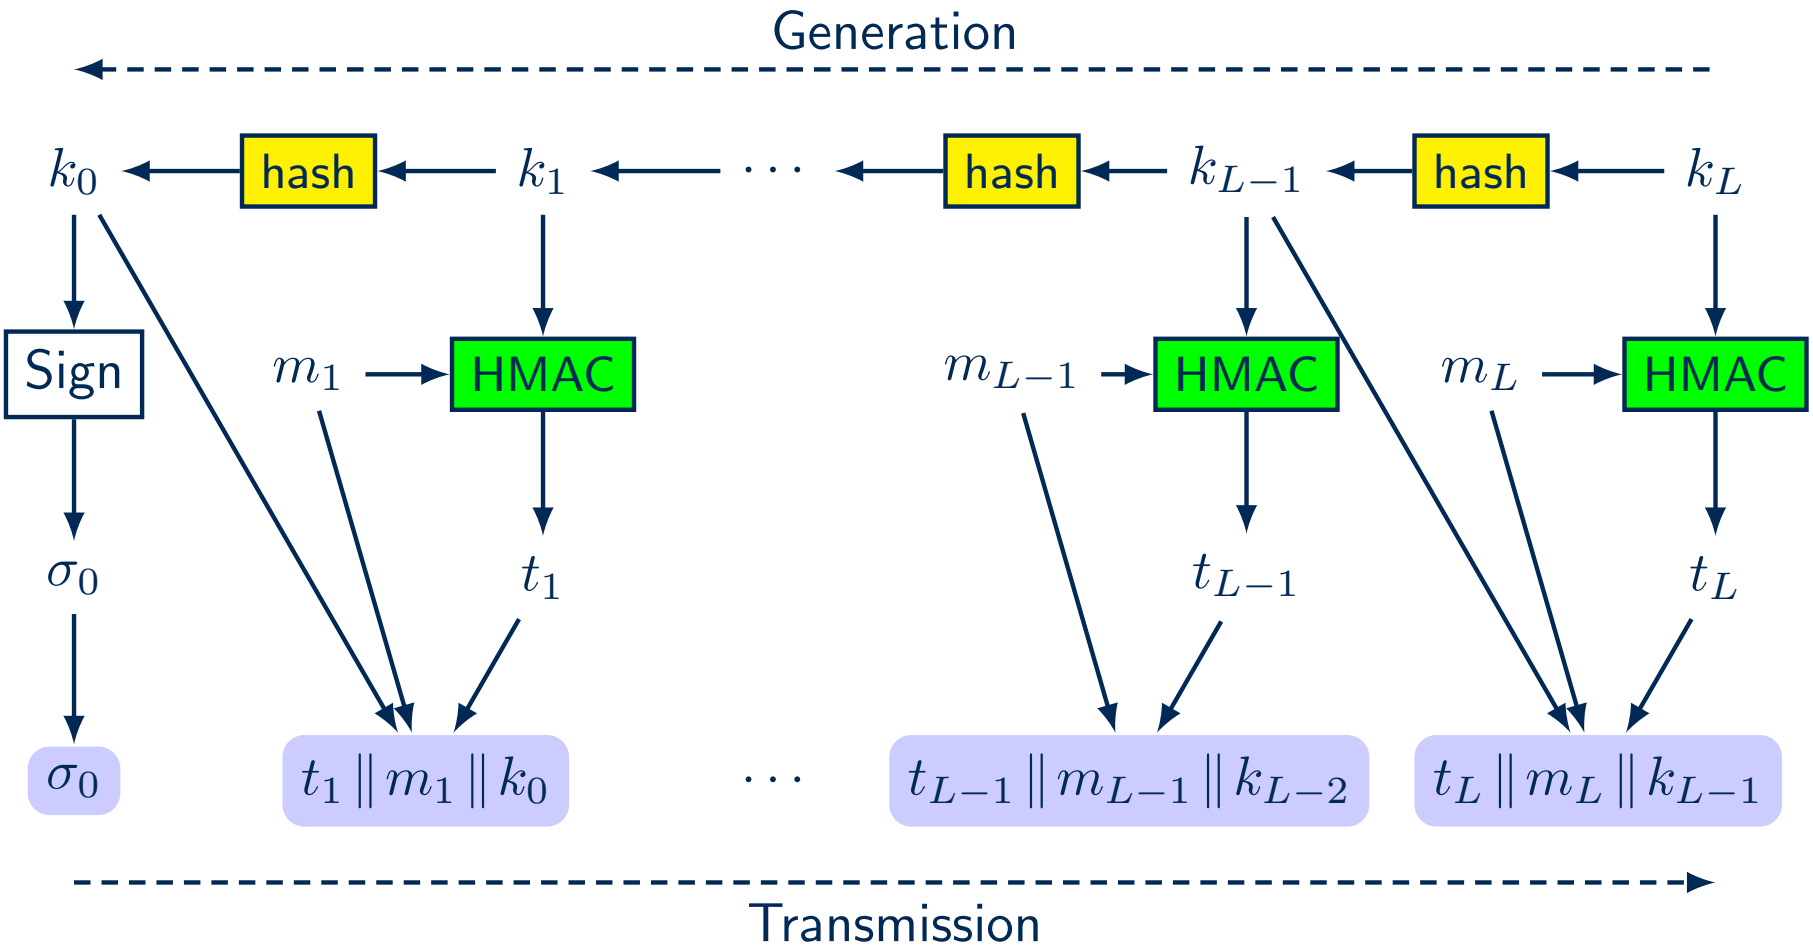

Security of TESLA

The security of TESLA is based on the one-wayness of the hash function in the chain, i.e. on the difficulty of finding possible future keys that match the so far disclosed keys

Previous work

The security of TESLA has been proved

- in the assumption that the whole hash chain is a pseudo random function [Perrig *et al.*, '00]
- when each key is separately generated and its committment obtained by a single hash [Archer, '02]

PHY authentication: System model

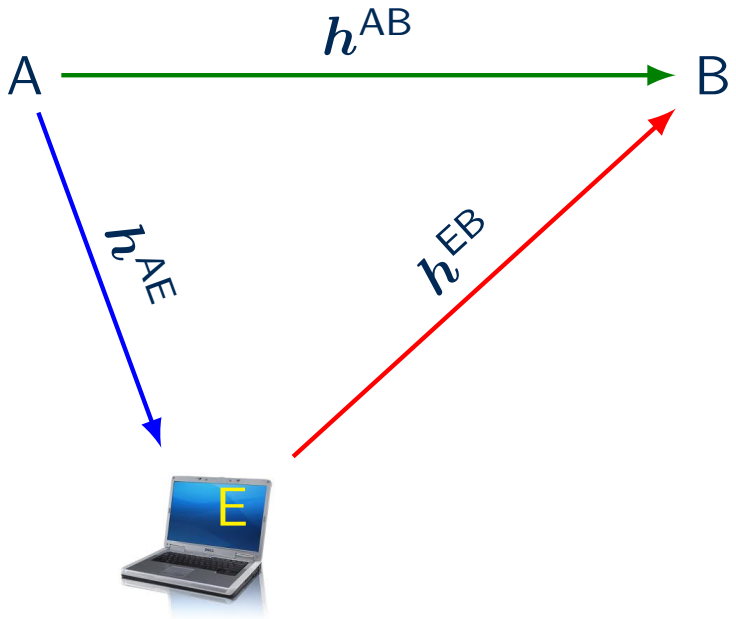

$\boldsymbol{h} = [h_0, \ldots, h_{N-1}]$:
channel fading coefficients (e.g., impulse response, frequency response, channel matrix entries)

channel statistics

complex, jointly Gaussian, circularly symmetric

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

channel reciprocity

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

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

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

Phase II: hypothesis testing

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

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

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

The wiretap channel [Wyner, '75]

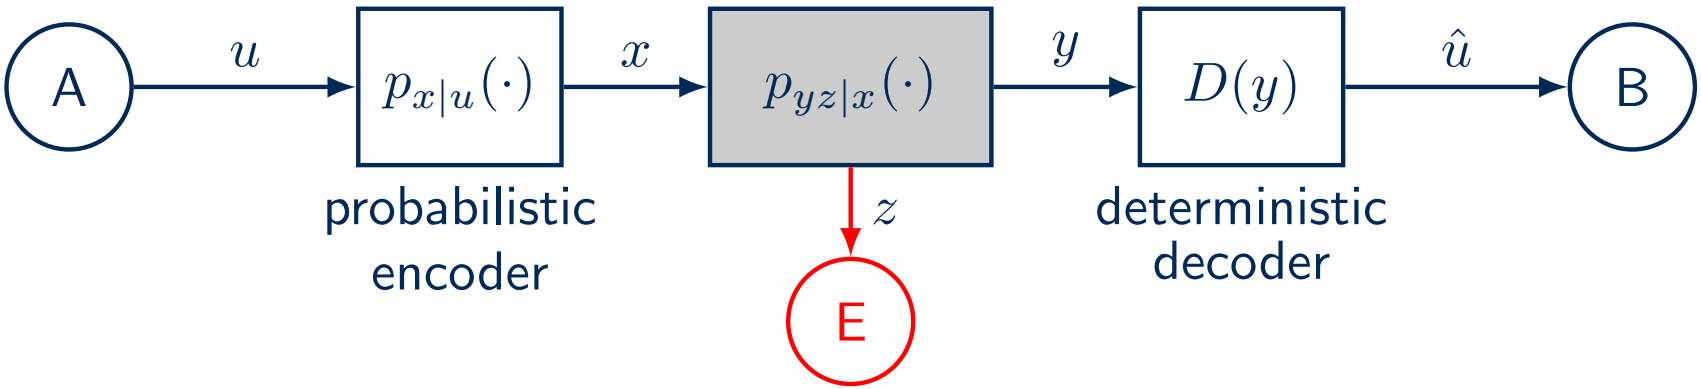

We aim for reliable transmissions to B, and secrecy with respect to E

Secrecy capacity

$C_{\mathsf{s}} = \lim_{n\to\infty} \max_{u, p_{x|u}, D} \left[\frac{1}{n} H(\boldsymbol{u})\right]$ subject to:

reliability: $\lim_{n\to\infty} P\left[\boldsymbol{u} \neq \hat{\boldsymbol{u}}\right] = 0$

(strong) secrecy: $\lim_{n\to\infty} I(\boldsymbol{u}; \boldsymbol{z}) = 0$

or (weak) secrecy: $\lim_{n\to\infty} \frac{1}{n} I(\boldsymbol{u}; \boldsymbol{z}) = 0$

Theorem [Csiszàr-Körner, '78]

For memoryless channels

$$\begin{aligned} C_{\mathsf{s}} &= \max_u \left[I(u;y) - I(u;z)\right]^+ \\ &\geq \max_x \left[I(x;y) - I(x;z)\right]^+ \\ &\geq \left[C_{\mathsf{AB}} - C_{\mathsf{AE}}\right]^+ \end{aligned}$$

Toy example: uniform channel

Consider a wiretap channel in which

$$p_{y|x}(b|a) = \begin{cases} 1/N_{y|x} & , b \in T_{y|x}(a) \\ 0 & , b \notin T_{y|x}(a) \end{cases} , \quad p_{z|x}(c|a) = \begin{cases} 1/N_{z|x} & , c \in T_{z|x}(a) \\ 0 & , c \notin T_{z|x}(a) \end{cases}$$

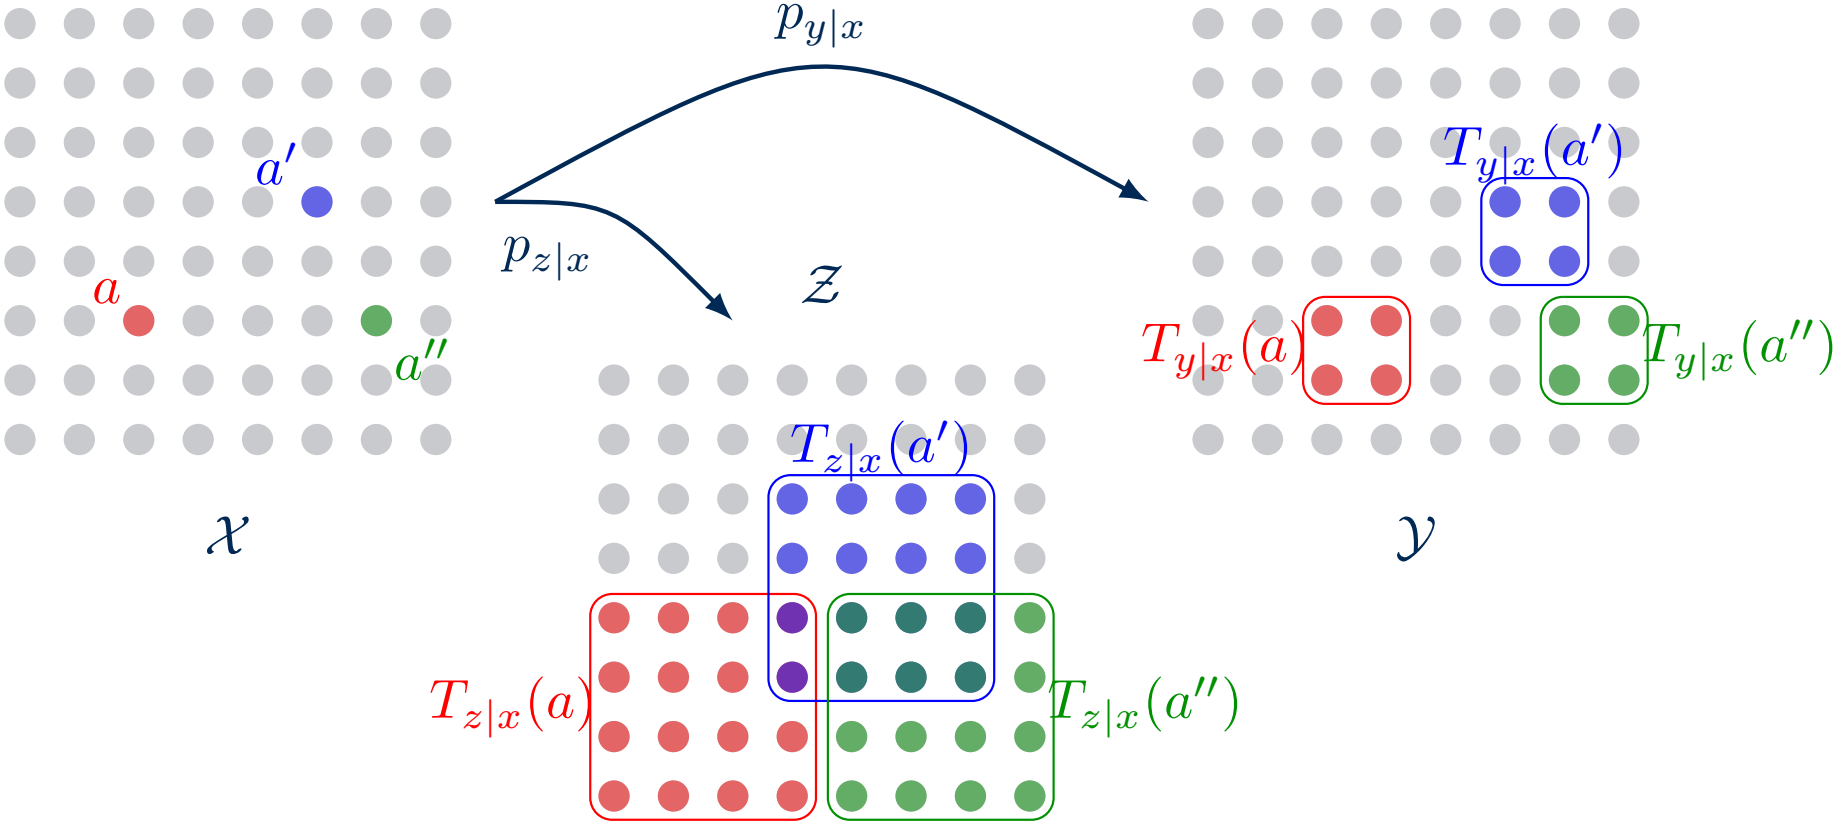

Random binning

If we can find:

- a subset $\mathcal{X}' \subset \mathcal{X}$ such that $\forall a \neq a' \in \mathcal{X}'$, $T_{y|x}(a) \cap T_{y|x}(a') = \emptyset$
- a set $\mathcal{M}$ and a partition of $\mathcal{X}'$ into $\left\{T_{x|u}(d)\right\}_{d \in \mathcal{M}}$ such that

$$\bigcup_{a \in T_{x|u}(d)} T_{z|x}(a) = \mathcal{Z} \quad , \quad \forall d \in \mathcal{M}$$

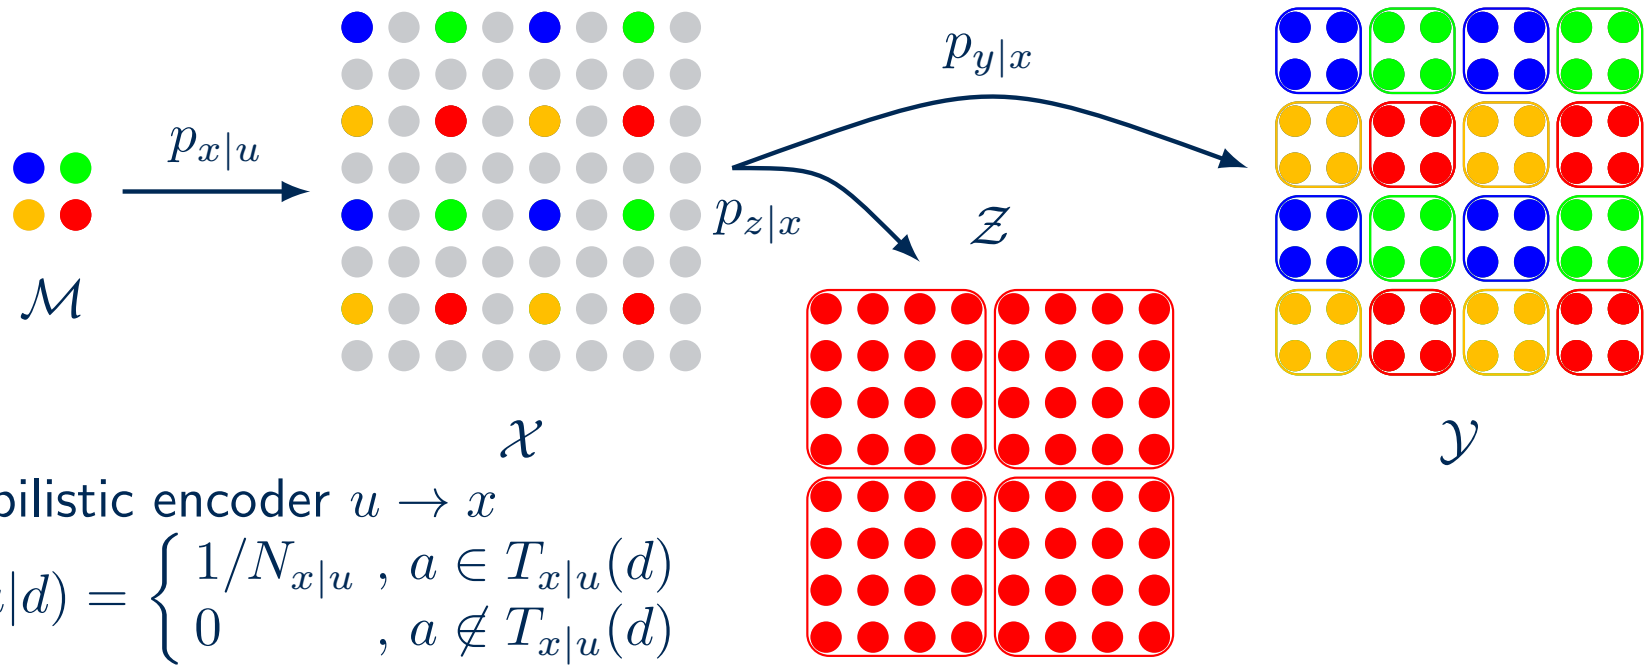

probabilistic encoder $u \to x$

$$p_{x|u}(a|d) = \begin{cases} 1/N_{x|u} & , a \in T_{x|u}(d) \\ 0 & , a \notin T_{x|u}(d) \end{cases}$$

How many secrecy bits?

For perfect reliability to B:

$$|\mathcal{X}'| \leq \frac{|\mathcal{Y}|}{N_{y|x}}$$

For perfect secrecy with respect to E:

$$N_{x|u} \geq \frac{|\mathcal{Z}|}{N_{z|x}}$$

For both

$$M = |\mathcal{M}| \leq \frac{|\mathcal{X}'|}{N_{x|u}} \leq \frac{|\mathcal{Y}|}{N_{y|x}} \frac{N_{z|x}}{|\mathcal{Z}|} = 2^{I(x;y)-I(x;z)}$$

Outline

Security of TESLA with very long hash chains

In assessing the security of TESLA it is assumed that each key k_i in the chain is unpredictable.
However, repeated hashing brings collisions and hence reduces entropy of keys.
Is it possible for an attacker to leverage the reduction in entropy and forge a valid key chain?

Aim of our work

To assess the security of the TESLA hash chain against attacks that aim at forging messages by leveraging a forged key chain

Hash chain example

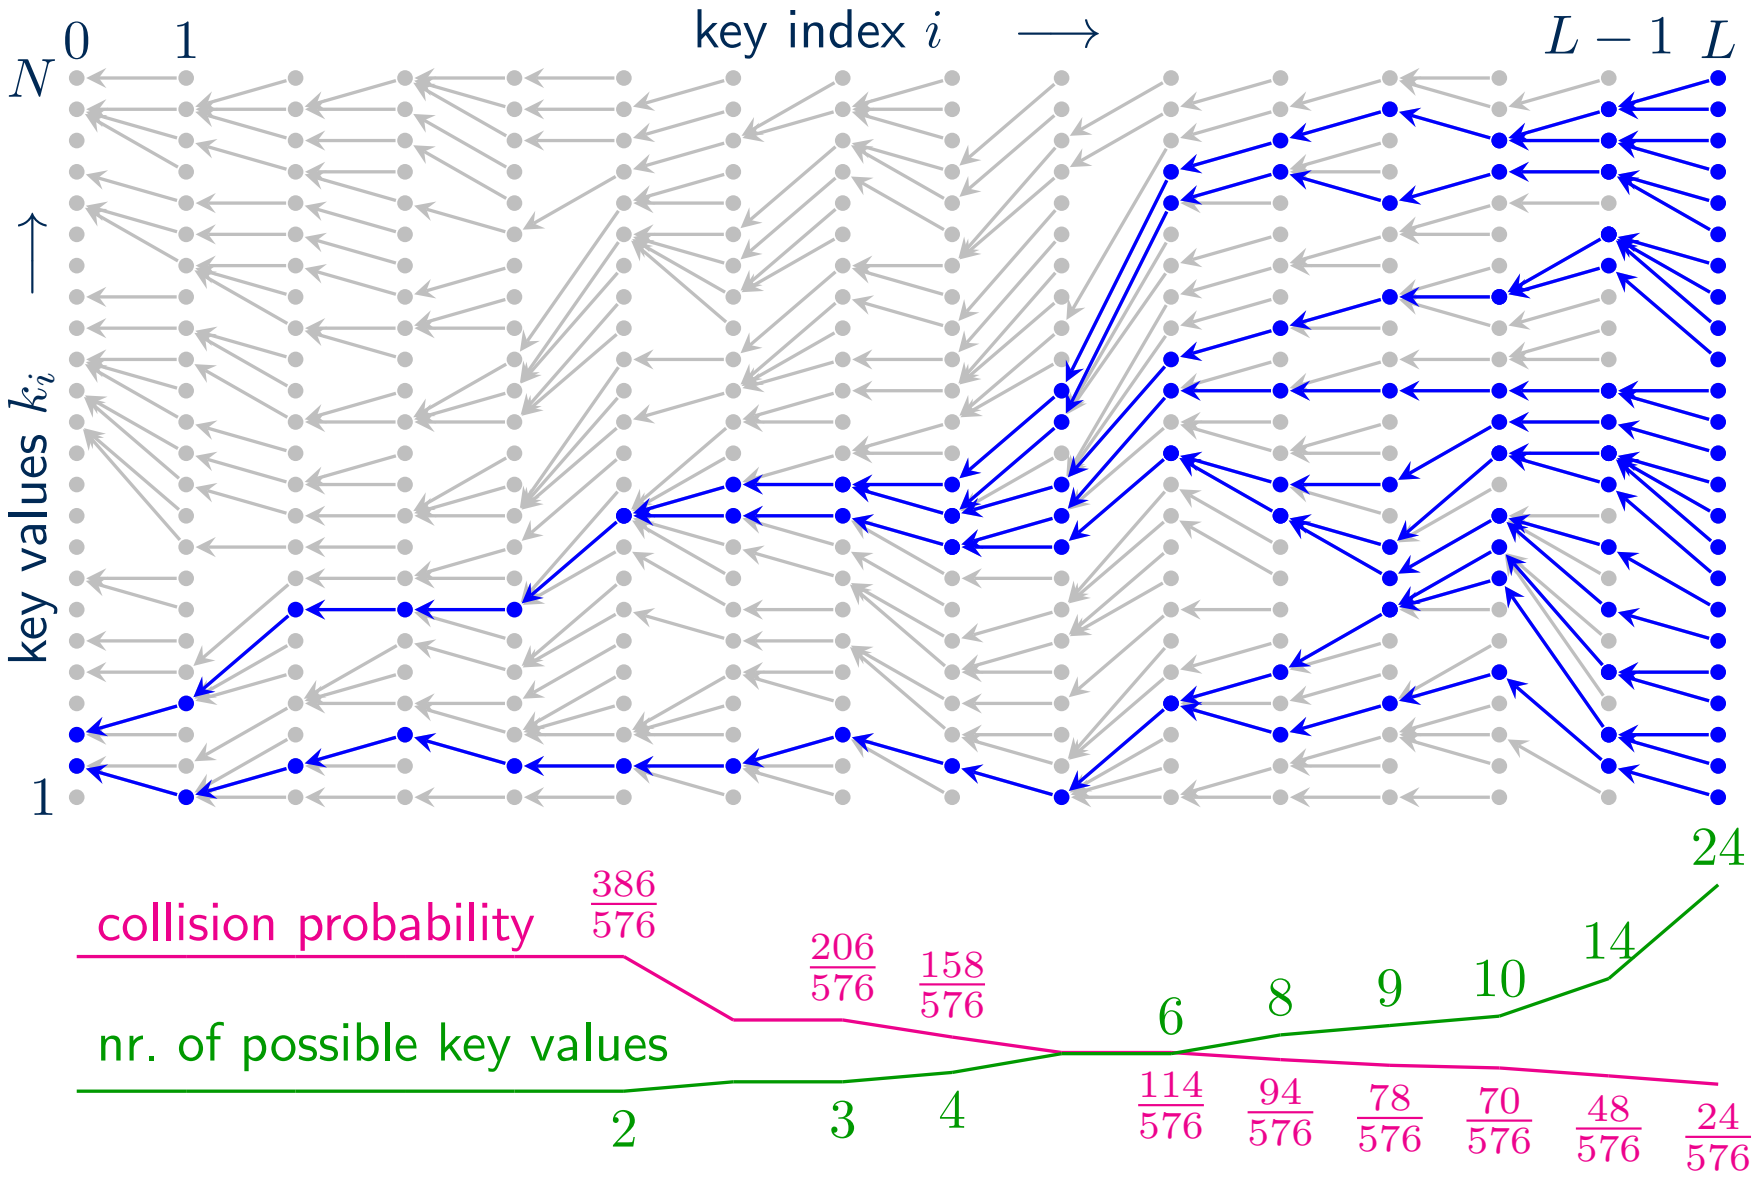

Attack model

Attack parameters

$T_A = \ell T_k$	length of the target forging interval
R_{h}	hashing rate available to the attacker
T	total time dedicated to attack computing
N_A	total number of guesses (attempts)

For an attack that starts at the disclosure of k_i and last for ℓ key intervals, the attacker tries to find $\hat{k}_{i+1}, \ldots, \hat{k}_{i+\ell}$ such that

$$f_i(\hat{k}_{i+1}) = k_i, f_{i+1}(\hat{k}_{i+2}) = \hat{k}_{i+1}, \ldots, f_{i+\ell-1}(\hat{k}_{i+\ell}) = \hat{k}_{i+\ell-1}$$

He guesses a random value for $\hat{k}_{i+\ell}$, computes ℓ hashes and checks if he gets k_i. If not, he tries another guess.
The success event for a single guess $j = 1, \ldots, N_A$ is

$$S_j(i, \ell) = \{\hat{k}_i^j = k_i\}$$

Attack example

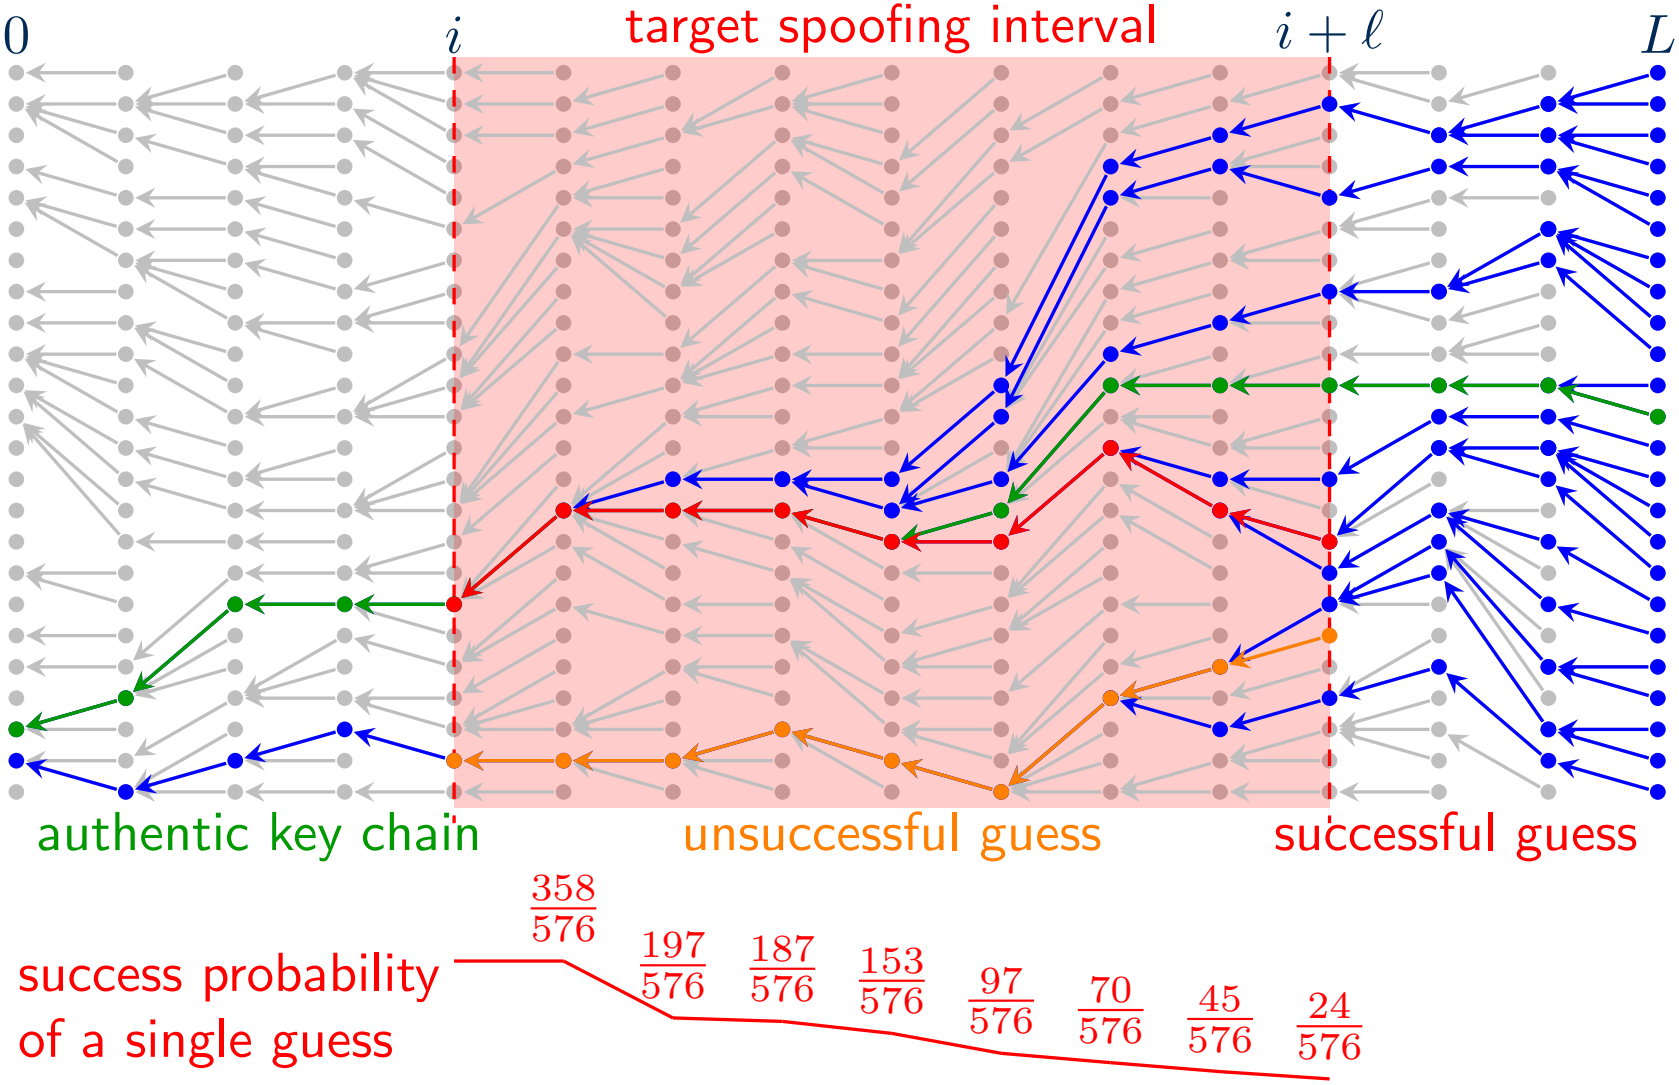

Attack success probability

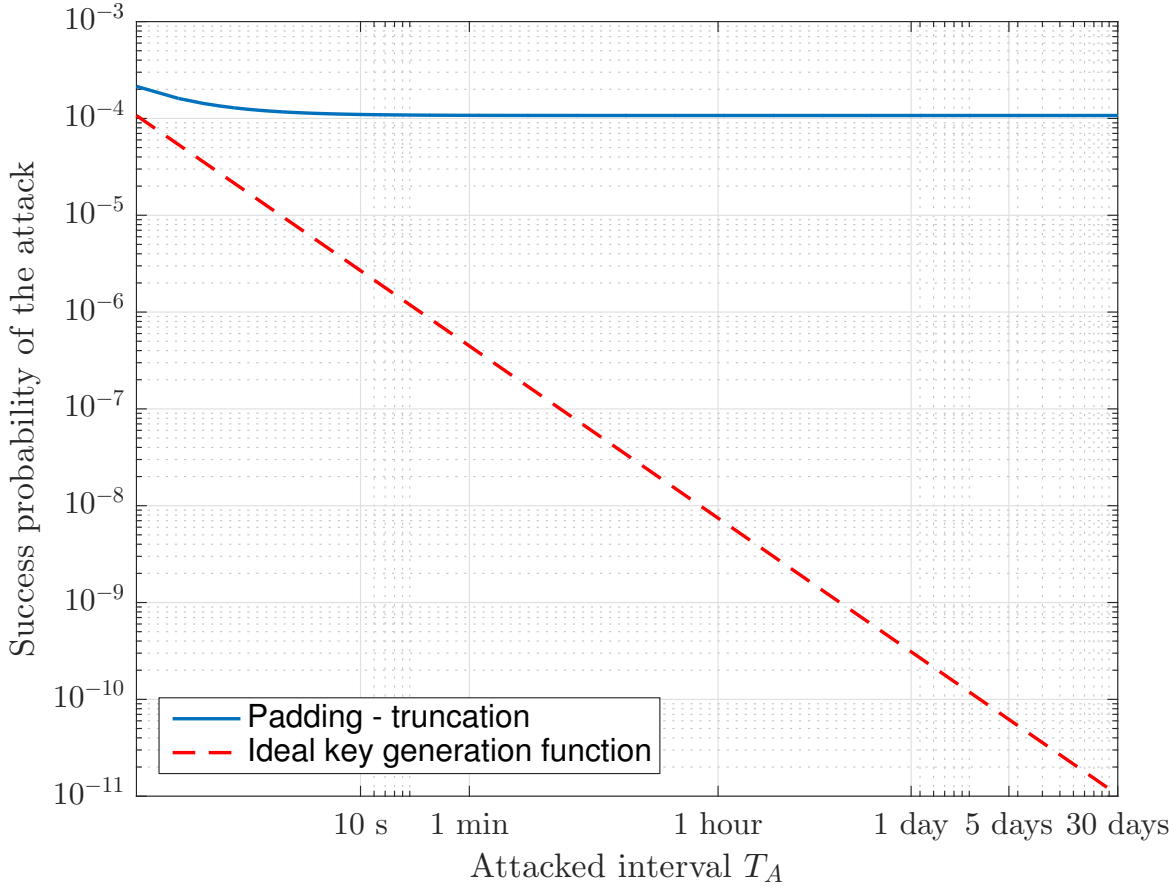

Success probability for an attack against a 80-bit key chain ($N = 2^{80}$) with a new key released every $T_k = 0.25\,\text{s}$, with hashing rate $R_\text{h} = 5 \cdot 10^{13}\,\text{hash/s}$, and attack duration $T = 30\,\text{days}$.

PHY Authentication for the IoT

Our assumptions

We consider:

- a CIoT network
- *anchor nodes*, that are trusted and securely connected with the concentrator

Our goals

1. provide a key-less message authentication scheme
2. assuming that anchor nodes have a limited energy availability, propose suitable scheduling policies for the activation of the anchor nodesto maximize the anchor network lifespan

Reference Scenario

Consider an CIoT scenario with M legitimate sources, N anchor nodes (with indices i), and one concentrator node C over a narrowband channel

$$\downarrow$$

The complex channel gains (including path loss λ_i and fading) from source node S to anchor node $i = 1, \ldots, N$ are collected into the vector

$$\boldsymbol{h}(\mathsf{S}) = [h_1(\mathsf{S}), \ldots, h_N(\mathsf{S})]$$

Assumptions:

- communication between the anchor nodes and the concentrator node C is secure
- when involved in the authentication of a node, an anchor node consumes a fixed amount of energy so that each anchor node is able to perform at most Q message authentications

Proposed Authentication Protocol I

Authentication procedure:

1. *initialization*: anchor nodes receive a message coming from source S that has been authenticated by some other methods (e.g., by a key-based authentication procedure) and the anchor nodes estimate the channel gain vectors $\hat{\boldsymbol{h}}^{(0)}(\mathrm{S})$. This estimate is reported to C.
2. *runtime*: upon reception of the k-th message reportedly coming from source S, a sub-set of anchor nodes (*configuration*) estimate the channel on the message and report the estimate to the concentrator node

Proposed Authentication Protocol II

The *anchor node configuration* for the k-th message from S is denoted by $\boldsymbol{c}(\mathsf{S}, k) \in \{0,1\}^{N\times 1}$ where

$$[\boldsymbol{c}(\mathsf{S}, k)]_i = \begin{cases} 1 & \text{if node } i \text{ is active in the authentication} \\ 0 & \text{otherwise} \end{cases}$$

$$\downarrow$$

C obtains from the active anchor nodes the estimated channel gain vector $\hat{\boldsymbol{h}}^{(k)}(\mathsf{S})$:

- $\mathcal{H}_0$: actual transmitter is S. Then

$$\hat{h}_i^{(k)}(\mathsf{S})[\boldsymbol{c}(\mathsf{S}, k)]_i \approx h_i(\mathsf{S})[\boldsymbol{c}(\mathsf{S}, k)]_i \quad i = 1, \ldots, N$$

- $\mathcal{H}_1$: A is transmitting. Then

$$\hat{h}_i^{(k)}(\mathsf{S})[\boldsymbol{c}(\mathsf{S}, k)]_i \approx g_i(\mathsf{S})[\boldsymbol{c}(\mathsf{S}, k)]_i \quad i = 1, \ldots, N$$

Admissible Configurations

Definitions:

- *admissible configuration* = configuration s.t. $P_{\mathrm{FA}} \leq P_{\mathrm{FA}}^{\mathrm{thr}}$ and $P_{\mathrm{MD}} \leq P_{\mathrm{MD}}^{\mathrm{thr}}$
- *efficient admissible configuration* = admissible configuration with a minimal set of active nodes

$$\downarrow$$

We collect all efficient admissible configurations into the $N \times A$ binary matrix

$$\boldsymbol{C} = [\ \boldsymbol{c}_1(1)\ \cdots\ \boldsymbol{c}_{a_1}(1)\ \cdots\ \boldsymbol{c}_1(M)\ \cdots\ \boldsymbol{c}_{a_M}(M)\ ]$$

where

- a_{S} = number of efficient admissible configurations for source node S
- $\boldsymbol{c}_\ell(\mathsf{S})$ = ℓ-th efficient admissible configuration for S
- $A = \sum_{m=1}^{M} a_m$ is the total number of efficient admissible configurations

$$\downarrow$$

Efficient admissible configurations are chosen *randomly* according to a predetermined probability distribution $p_\ell(\mathsf{S})$, that can be found in vector

Network Lifespan

Reference performance metric:
anchor network lifespan L = smallest number of authentication processes after which at least one anchor node runs out of power
Since the choice of the configuration is random, L is a random variable: in the paper, derivations of

- upper/lower bound on $F_L(x)$
- approximation of $F_L(x)$, neglecting correlation among anchor nodes

How to compute π?

Define the utilization vector $\boldsymbol{u} \in [0,1]^{N\times 1}$ as

$$\boldsymbol{u} = \frac{1}{M}\boldsymbol{C\pi}$$

where $1/M$ is the probability that each source node is transmitting

Optimization Problems

Objective functions:

- Least Squares Problem

$$\min_{\boldsymbol{\pi}} \sum_{i=1}^{N} u_i^2$$

- Minimum Variance Problem

$$\min_{\boldsymbol{\pi}} \sum_{i=1}^{N} \left(u_i - \frac{1}{N} \sum_{j=1}^{N} u_j \right)^2$$

- Min-Max Problem

$$\min_{\boldsymbol{\pi}} \max_{i=1}^{N} u_i$$

Constraints:

$$\boldsymbol{u} = \frac{1}{M} \boldsymbol{C} \boldsymbol{\pi}$$

$$0 \leq \pi_\ell(\mathsf{S}) \leq 1 \quad \ell = 1, \ldots, a_{\mathsf{S}}, \ \mathsf{S} = 1, \ldots, M$$

Observations

Least squares prob. (33) and minimum variance probl. (33) are *convex*

$\Downarrow$

Solved fast using well-known techniques (e.g., interior point method)

Min-max prob. (33) can be *linearized*

$\Downarrow$

Solved fast using the simplex algorithm

Anchor Network Lifespan

a) Least squares problem

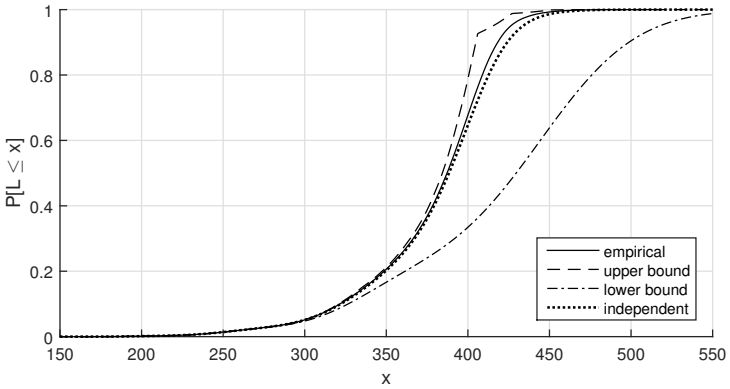

b) Minimum variance problem

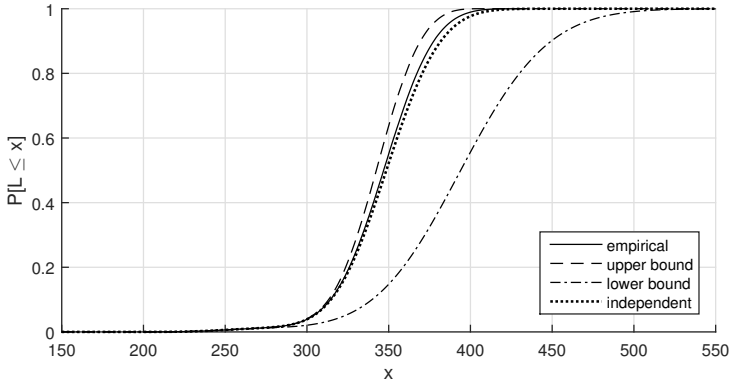

c) Min-max problem

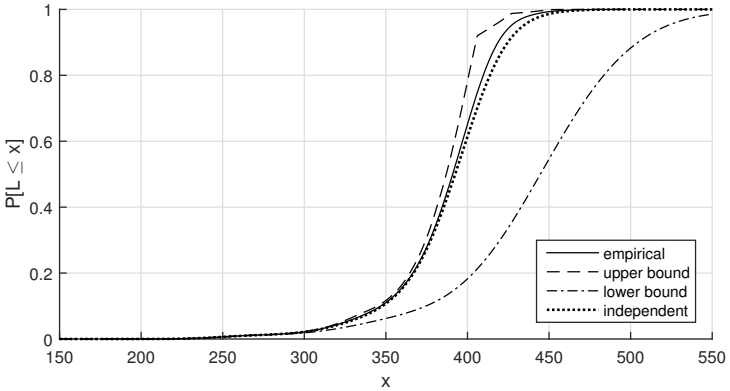

Log-scale comparison

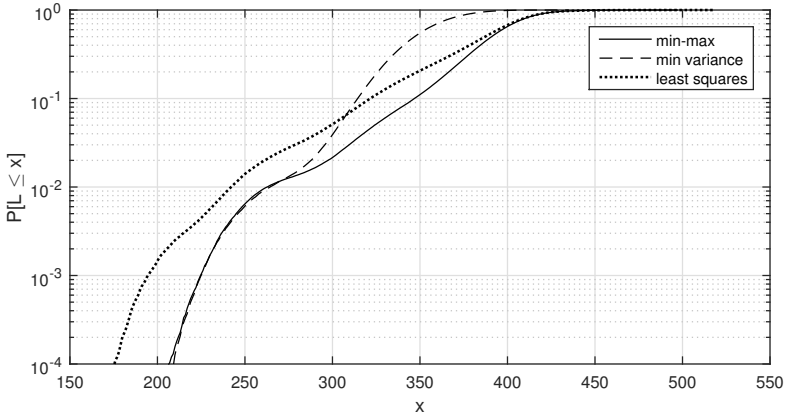

Parameters

$N = 9$, $\eta = 2$

SNR $= 30$ dB, $\rho = 0.1$

Results

a) and c) better than b)

Up to $4\times$ the default lifespan!

PHY secrecy and Energy Harvesting

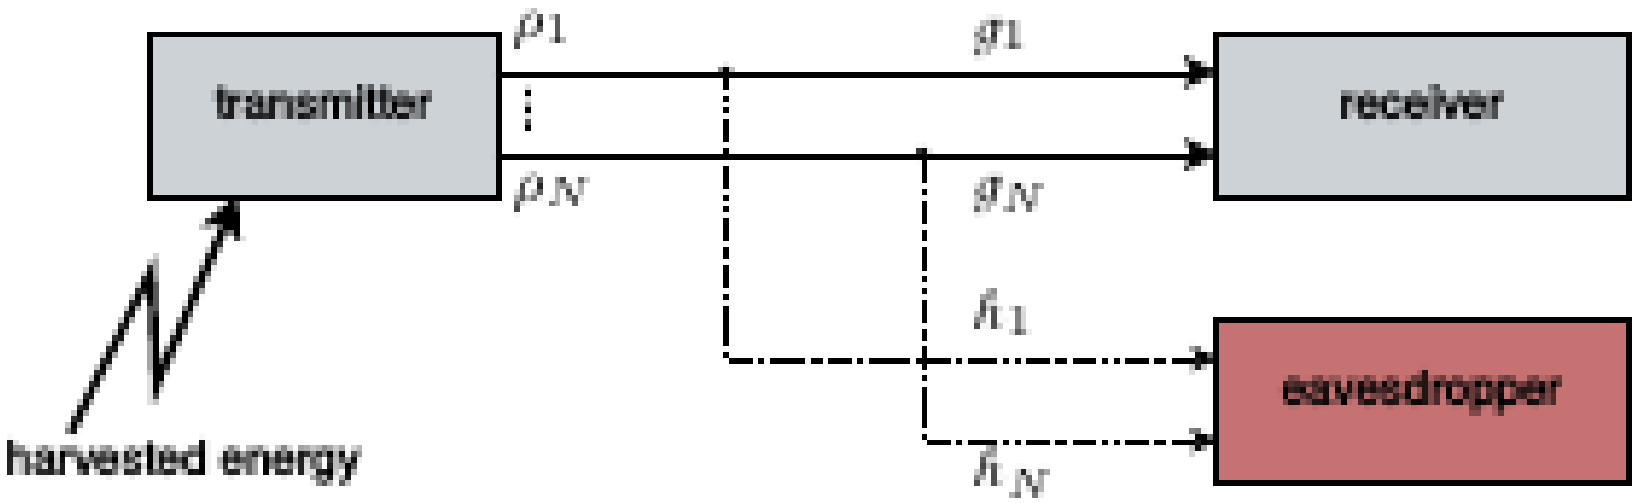

The transmitter harvests energy from an external, non-controllable and renewable energy source

Transmit power is split among N subchannels

Our goal is to maximize the achievable secrecy rate, by choosing the Optimal Secrecy Policy

Problem statement

Given We model our system with a Markov Chain (MC) with a finite number of states. For every MC state $(e, \boldsymbol{g}, \boldsymbol{h})$, a *power allocation policy* μ is the set of rules

$$\mu = \{\mu(\cdot; e, \boldsymbol{g}, \boldsymbol{h}),\ \forall e \in \mathcal{E},\ \forall \boldsymbol{g} \in g,\ \boldsymbol{h} \in h\},$$

where

$$\mu(\boldsymbol{\rho}; e, \boldsymbol{g}, \boldsymbol{h}) \triangleq \mathbb{P}\left(\begin{array}{c}\text{using a power}\\ \text{splitting vector } \boldsymbol{\rho}\end{array} | e, \boldsymbol{G} = \boldsymbol{g}, \boldsymbol{H} = \boldsymbol{h}\right),$$

Optimal policy is deterministic

Theorem

There exists a deterministic OSP, i.e., an optimal secrecy policy in which, for every MC state $(e, \boldsymbol{g}, \boldsymbol{h})$, $\exists \boldsymbol{\rho}^{\star}_{e,\boldsymbol{g},\boldsymbol{h}}$ *such that*

$$\mu^{\star}(\boldsymbol{\rho}; e, \boldsymbol{g}, \boldsymbol{h}) = \begin{cases} 1, & \text{if } \boldsymbol{\rho} = \boldsymbol{\rho}^{\star}_{e,\boldsymbol{g},\boldsymbol{h}}, \\ 0, & \text{otherwise,} \end{cases}$$

where $\boldsymbol{\rho}^{\star}_{e,\boldsymbol{g},\boldsymbol{h}}$ *depends upon the current MC state in general.*

Efficient optimization

Theorem

The maximization of C_μ can be decomposed into two steps:

1. *fix a value x and the channel gain vectors $\boldsymbol{g}$, $\boldsymbol{h}$ and find the optimal power splitting choice*

$$\begin{aligned} & \boldsymbol{\rho}^\star = \arg\max_{\boldsymbol{\rho}} c(\boldsymbol{\rho}, \boldsymbol{g}, \boldsymbol{h}), \\ \text{s.t.} \quad & \boldsymbol{\rho} \in \mathcal{P}_=(x) \triangleq \left\{ \boldsymbol{\rho} \ : \ \boldsymbol{\rho} \succeq 0, \ x = \mathbf{1}_N^T \boldsymbol{\rho} \right\}; \end{aligned}$$

2. *maximize C_μ by considering only μ^{tot}*

$$\begin{aligned} & {\mu^{\text{tot}}}^\star = \arg\max_{\mu^{\text{tot}}} C_{\boldsymbol{P}}, \\ \text{s.t.} \quad & \mu^{\text{tot}} \text{ and } \mu \text{ are consistent,} \\ & \boldsymbol{\rho}_{e,\boldsymbol{g},\boldsymbol{h}} \text{ solves the above with } x = \rho^{\text{tot}}_{e,\boldsymbol{g},\boldsymbol{h}}, \\ & \qquad \forall e \in \mathcal{E}, \forall \boldsymbol{g} \in g, \ \forall \boldsymbol{h} \in h, \end{aligned}$$

The optimal $\mu^\star$ can be found by fixing ${\rho^{\text{tot}}}^\star$ according to point 2) and choosing $\boldsymbol{\rho}$ with the optimal power splitting choice of point 1).

Monotonicity of optimal policy

Theorem

Consider $N = 1$. The transmission power of OSP is non-decreasing with g and non-increasing with h (we omit the "1" subscripts). Formally

- *if $g'' \geq g'$, then $\rho^{\text{tot}\star}_{e,g'',h} \geq \rho^{\text{tot}\star}_{e,g',h}$;*
- *if $h'' \geq h'$, then $\rho^{\text{tot}\star}_{e,g,h''} \leq \rho^{\text{tot}\star}_{e,g,h'}$.*

Theorem

Consider $N = 1$. With partial CSI, the transmission power of OSP is non-decreasing with g (we omit the "1" subscripts). Formally, if $g'' \geq g'$, then $\rho^{\text{tot}\star}_{e,g''} \geq \rho^{\text{tot}\star}_{e,g'}$.

Numerical results

Secrecy rate vs battery size

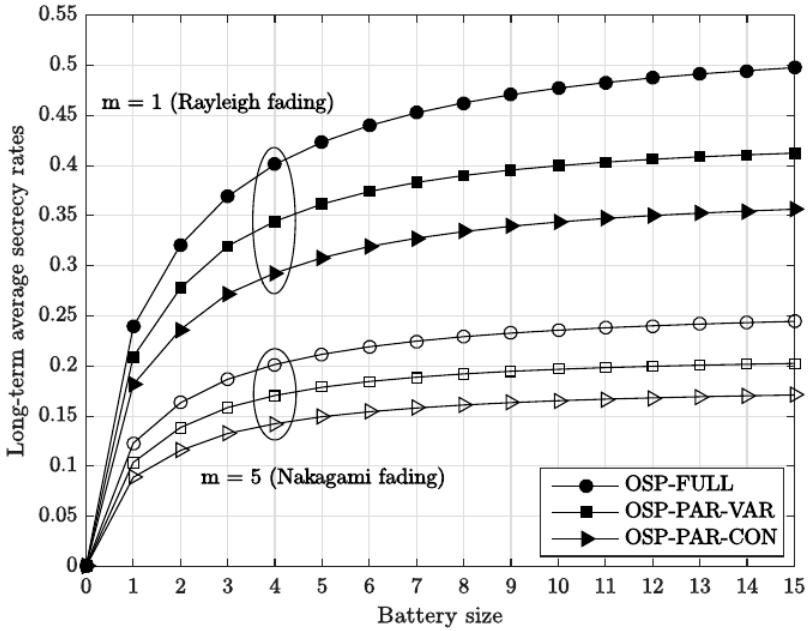

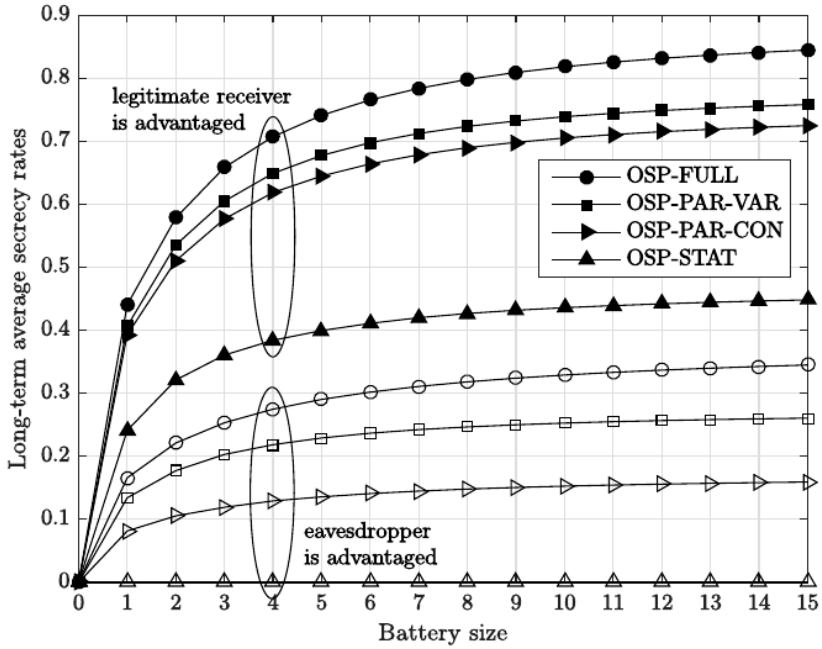

No need for large batteries

Essential references

J. Granjal, E. Monteiro, and J. Sa Silva
"Security for the Internet of Things: A Survey of Existing Protocols and Open Research Issues"
IEEE Communications Surveys & Tutorials, 17(3): 1294–1312, 2015.

W. Trappe, R. Howard, and R. S. Moore
"Low-Energy Security: Limits and Opportunities in the Internet of Things"
IEEE Security & Privacy, 13(1): 14–21, 2015.

A. Perrig, R. Canetti, J. D. Tygar, and D. Song
"Efficient authentication and signing of multicast streams over lossy channels"
2000 IEEE Symposium on Security and Privacy, 2000, pp. 56-73.

L. Xiao, L. J. Greenstein, N. B. Mandayam, and W. Trappe
"Using the physical layer for wireless authentication in time-variant channels"
IEEE Trans. on Wireless Communications, 7(7): 2571–2579, 2008.