UMassAmherst

Everlasting Security and Undetectability in Wireless Communications

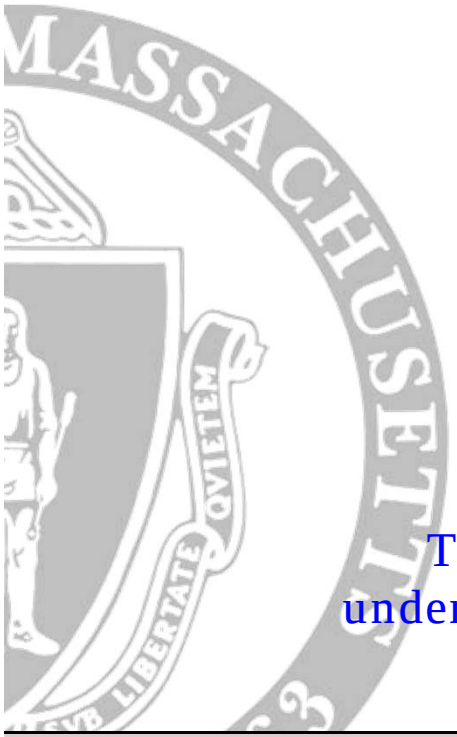

ICNC Lecture
February 6, 2014

Dennis Goeckel

University of Massachusetts Amherst

This work is supported by the National Science Foundation under Grants CNS-1019464, CCF-1249275, and ECCS-1309573.

Motivation

Everlasting Secrecy:

We are interested in keeping something secret forever. A challenge of cryptography (e.g. the VENONA project) is that recorded messages can be deciphered later.

From: *The Guardian*

Undetectability:

1. A stronger form of security than any encryption: computational or information-theoretic.
2. Often more important than encryption: whom is talking to whom (so called "metadata")

The Alice-Bob-Eve Scenario in Wireless

It might be Eve in the parking lot listening, or….

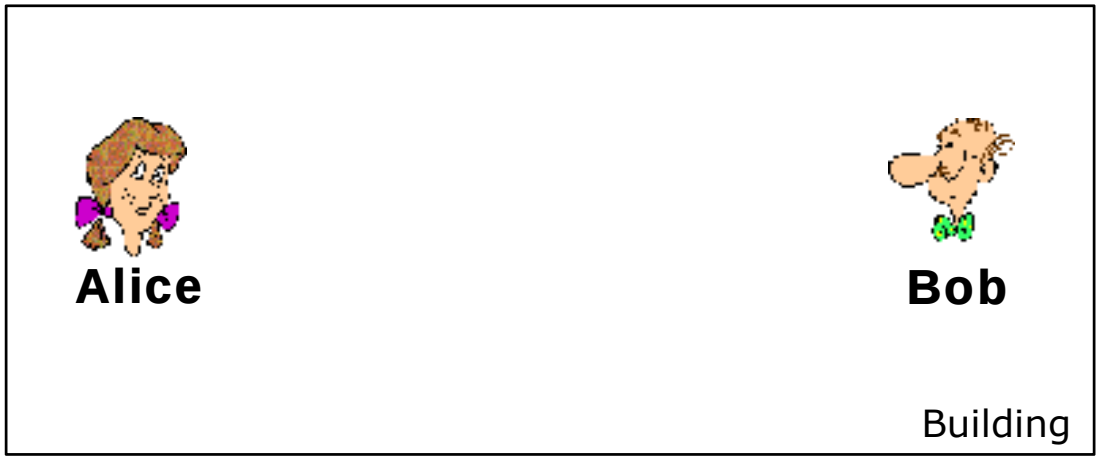

…it might be Eve in the building!

Important Challenge: the "near Eve" problem…

…and you very likely will not know where she is.

Computational Security (Cryptography)

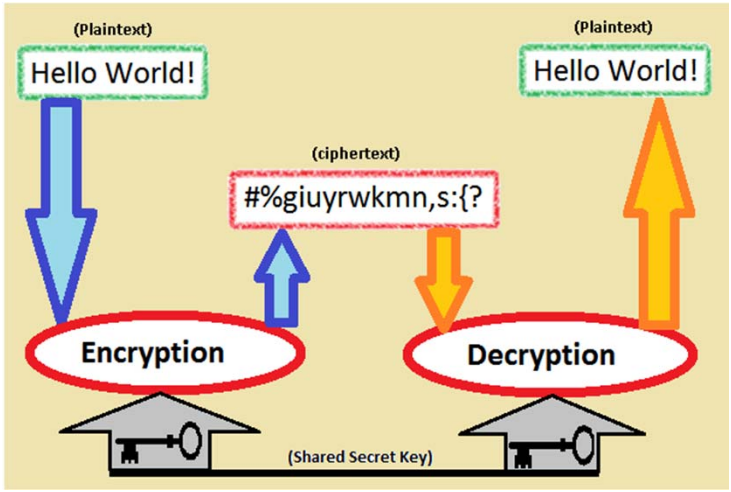

Eve can see the transmitted bits perfectly, but cannot solve the "hard" problem presented to her.

Advantages:

1. Well-studied and efficient algorithms
2. Does not suffer from the "Near Eve" problem

Disadvantages:

1. Implementations often broken (although the primitive is fine)
2. Computational assumptions on Eve
3. Message can be stored and decrypted later

Information-theoretic secrecy

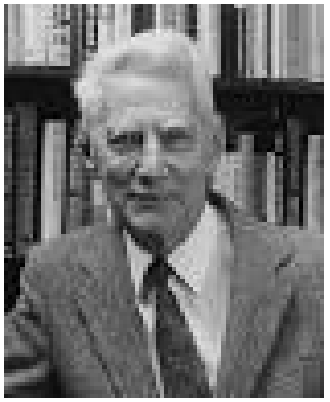

Information is encoded in such a way that Eve gets no information about the message…if the scenario is right

Advantages:

1. No computational assumptions on Eve.
2. *If* the transmission is securely made, it is secure forever.

Disadvantages (key part of this talk):

Information-theoretic secrecy generally relies on a (known) advantage for Bob over Eve (e.g. less noisy). If that is not true, Eve gets the message *today*.

Many would argue that we have traded a long-term computational risk for a short-term scenario risk…**no thank you**!

Outline

1. **Computational and Information Theoretic security basics**
   a. **Computation Security: Diffie-Hellman**
   b. **IT Security: The wiretap channel (Wyner et al)**
   c. **Application to wireless…and challenges**
2. Potential solutions
   a. Exploiting fading
   b. Two-way communications
   c. Attacking the receiver's hardware
   d. Cooperative jamming
3. Asymptotically-large networks
   a. Cooperative jamming
   b. Network coding
4. Undetectable communications (LPD)
   a. Steganography
   b. Emerging approaches for wireless channels
5. Current and Future Challenges

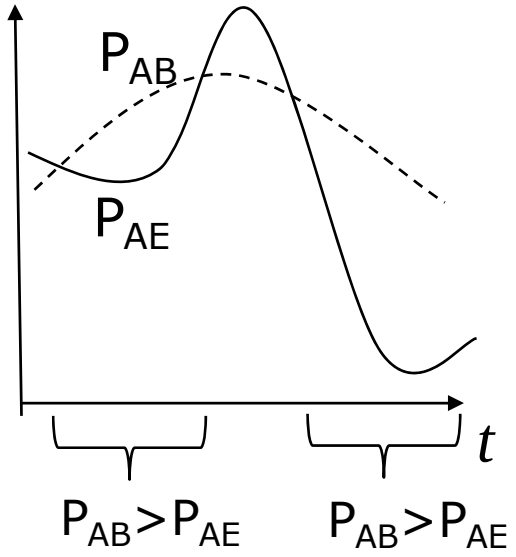

Outline

1. Computational and Information Theoretic security basics
   a. Computation Security: Diffie-Hellman
   b. Information-Theoretic Security: The wiretap channel (Wyner et al)
   c. Application to wireless…and challenges
2. **Potential solutions**
   a. Exploiting fading
   b. Two-way communications
   c. Attacking the receiver's hardware
   d. Cooperative jamming
3. Asymptotically-large networks
   a. Cooperative jamming
   b. Network coding
4. Undetectable communications (LPD)
   a. Steganography
   b. Emerging approaches for wireless channels
5. Current and Future Challenges

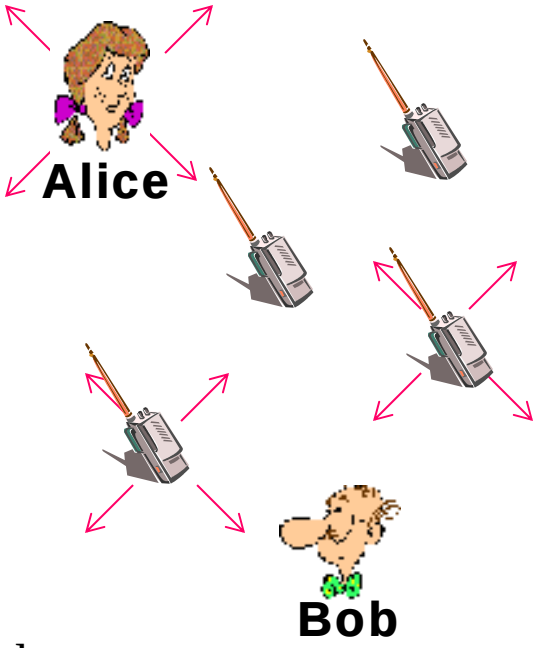

Outline

1. Computational and Information Theoretic security basics
   a. Computation Security: Diffie-Hellman
   b. Information-Theoretic Security: The wiretap channel (Wyner et al)
   c. Application to wireless...and challenges
2. Potential solutions
   a. Exploiting fading
   b. Two-way communications
   c. Attacking the receiver's hardware
   d. Cooperative jamming
3. **Asymptotically-large networks**
   a. Cooperative jamming
   b. Network coding
4. Undetectable communications (LPD)
   a. Steganography
   b. Emerging approaches for wireless channels
5. Current and Future Challenges

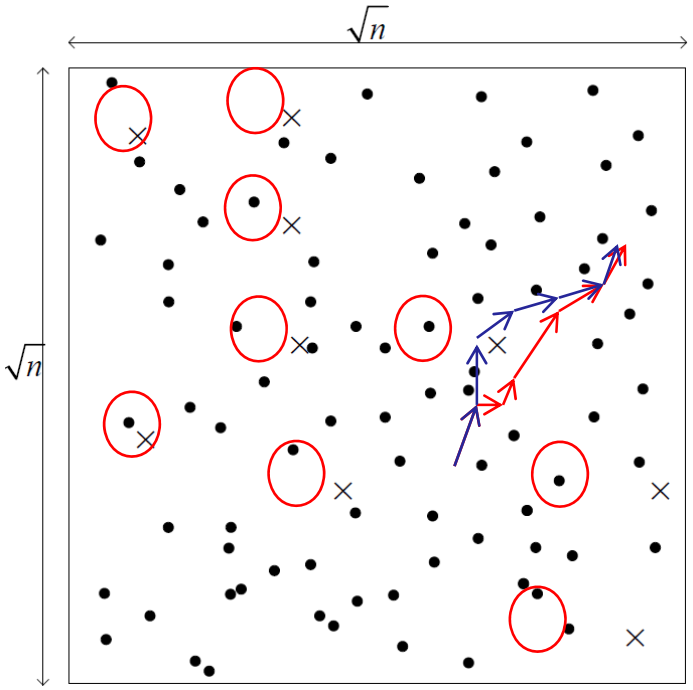

Outline

1. Computational and Information Theoretic security basics
   a. Computation Security: Diffie-Hellman
   b. Information-Theoretic Security: The wiretap channel (Wyner et al)
   c. Application to wireless...and challenges
2. Potential solutions
   a. Exploiting fading
   b. Two-way communications
   c. Attacking the receiver's hardware
   d. Cooperative jamming
3. Asymptotically-large networks
   a. Cooperative jamming
   b. Network coding
4. **Undetectable communications (LPD)**
   a. Emerging approaches for wireless
   b. Experiments
5. Current and Future Challenges

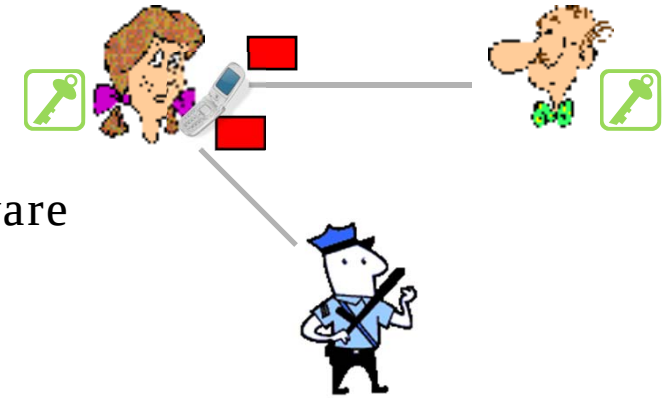

Outline

1. Computational and Information Theoretic security basics
   a. Computation Security: Diffie-Hellman
   b. Information-Theoretic Security: The wiretap channel (Wyner et al)
   c. Application to wireless...and challenges
2. Potential solutions
   a. Exploiting fading
   b. Two-way communications
   c. Attacking the receiver's hardware
   d. Cooperative jamming
3. Asymptotically-large networks
   a. Cooperative jamming
   b. Network coding
4. Undetectable communications (LPD)
   a. Emerging approaches for wireless channels
   b. Experiments
5. **Current and Future Challenges**

Outline

1. **Computational and Information Theoretic security basics**
   a. **Computation Security: Diffie-Hellman**
   b. **Information-Theoretic Security: The wiretap channel (Wyner et al)**
   c. **Application to wireless…and challenges**
2. Potential solutions
3. Asymptotically-large networks
4. Undetectable communications (LPD)
5. Current and Future Challenges

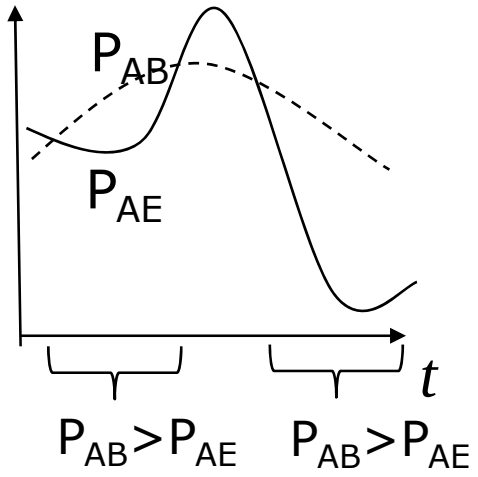

Diffie-Hellman: Establishing a key "on the fly"

Alice

y_a

Set-up: Alice announces a large prime p and a primitive root g mod p.

y_b

Bob

Eve

1. Alice chooses a secret random integer a, $1 < a < p-1$, and broadcasts $y_a = g^a \bmod p$.

2. Bob chooses a secret random integer b, $1 < b < p-1$, and broadcasts $y_b = g^b \bmod p$.

3. Alice forms the key $K = y_b^a \bmod p = g^{ab} \bmod p$

3. Bob forms the key $K = y_a^b \bmod p = g^{ab} \bmod p$

Eve is left trying to solve the *discrete logarithm* problem, which is believed to be "hard".

Diffie-Hellman: An Example

Alice

y_a

Set-up: Alice announces a large prime p=19 and a primitive root g=2.

y_b

Bob

Eve

1. Alice chooses a secret random integer a=5, $1 < a < 18$, and broadcasts $y_a = 2^5 \bmod 19 = 13$.

2. Bob chooses a secret random integer 6, $1 < b < 18$, and broadcasts $y_b = 2^6 \bmod 19 = 7$.

3. Alice forms the key: $K = 7^5 \bmod 19 = 11$

3. Bob forms the key: $K = 13^6 \bmod 19 = 11$

Eve is left trying to solve the *discrete logarithm* problem, which is believed to be "hard".

From: J. Talbot and D. Welsh, *Complexity and Cryptography*

Diffie-Hellman: How could it be broken?

1. The discrete logarithm is not hard (unlikely?)

2. Somebody obtains the key in some other manner (e.g. side-channel analysis on power utilization of a processor).

[Courtesy: C. Paar]

[from: "wired.com"]

3. Advances in computing

This motivates approaches of "keyless security", where what the eavesdropper receives does not contain enough information to (ever) decode the message…**information-theoretic secrecy**.

Shannon and the one-time pad [C. Shannon, 1948]

- Consider *perfect* secrecy over a ***noiseless*** wireline channel:

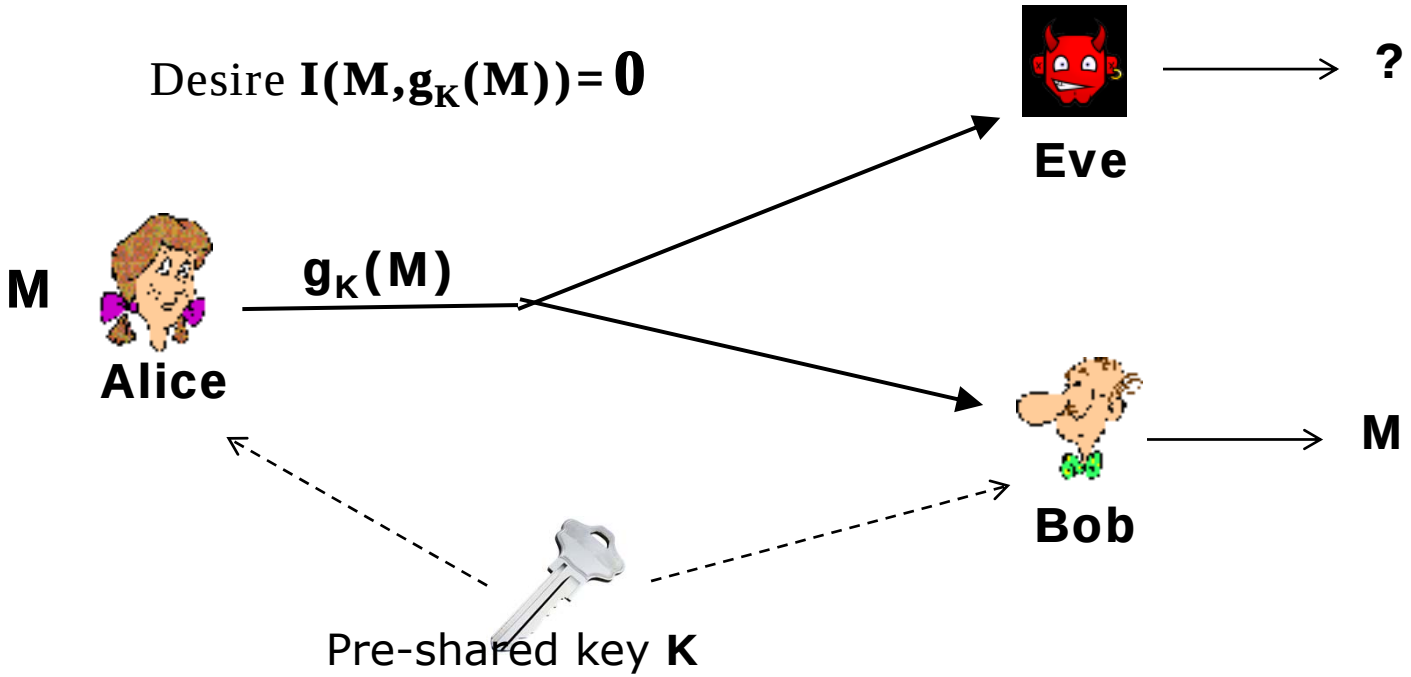

Questions:
1. How long must **K** be for an N-bit message **M**?
2. How do you choose $g_K(\mathbf{M})$?

Shannon and the one-time pad [C. Shannon, 1948]

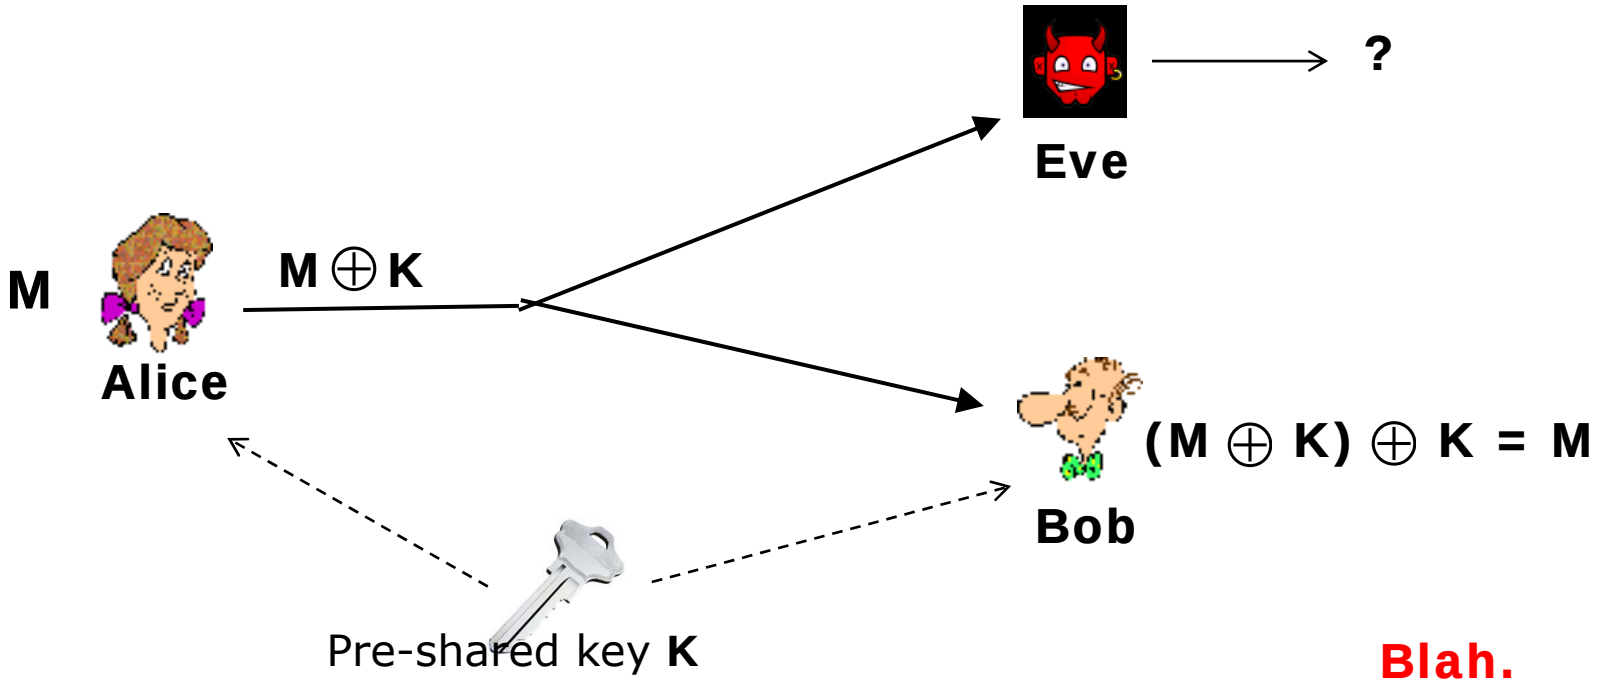

Blah.

Answers:

1. You need an N-bit key **K** for an N-bit message **M**.
2. $g_K(\mathbf{M}) = \mathbf{M} \oplus \mathbf{K}$

Example: Why not a "two-time pad"?

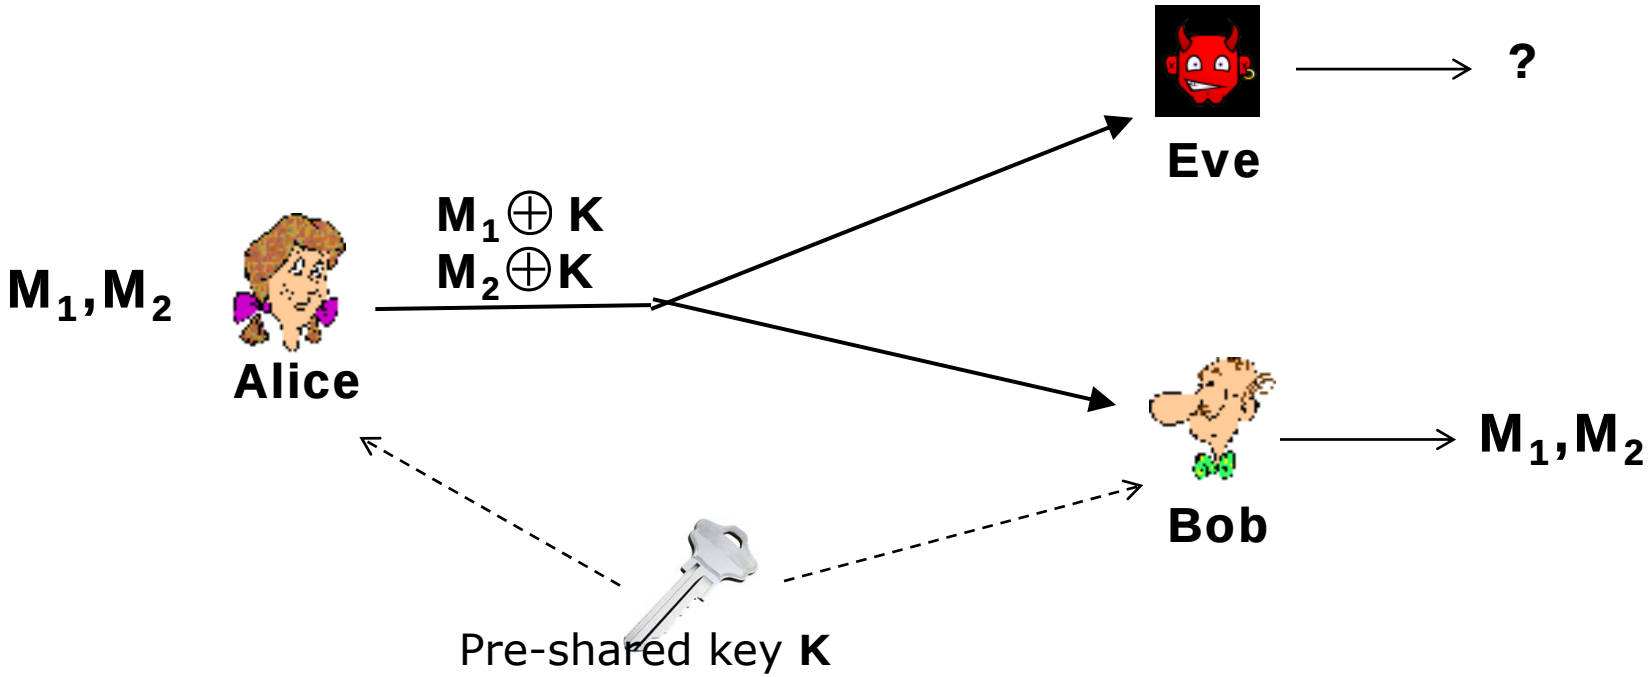

What does Eve do? $(M_1 \oplus K) \oplus (M_2 \oplus K) = M_1 \oplus M_2$

This is an information leak of K bits! Not information-theoretic secure. (VENONA exploited this).

The Wiretap Channel [Wyner, 1975; Cheong and Hellman, 1978]

But suppose:
1. There is noise in the system.
2. The eavesdropper has a worse view of the transmitted signal than Bob.

R_{AB}: Capacity of channel from Alice to Bob
R_{AE}: Capacity of channel from Alice to Eve

$R_{AB} > R_{AE}$

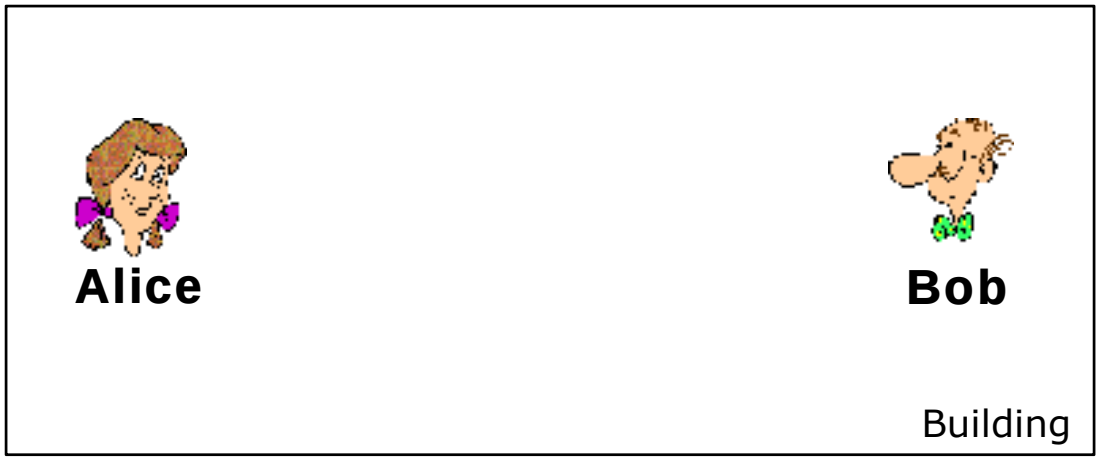

Gaussian channels: $R = \log_2(1 + SNR_{AB}) - \log_2(1 + SNR_{AE})$

Positive rate "if Bob's channel is better", and Eve gets nothing.

Wiretap Code Construction [Wyner, 1975]

$$R_o = R_{AB} - R_{AE} = 2 \qquad R_{AB} = \log(1+SNR_{AB}) = 4 \qquad R_{AE} = \log(1+SNR_{AE}) = 2$$

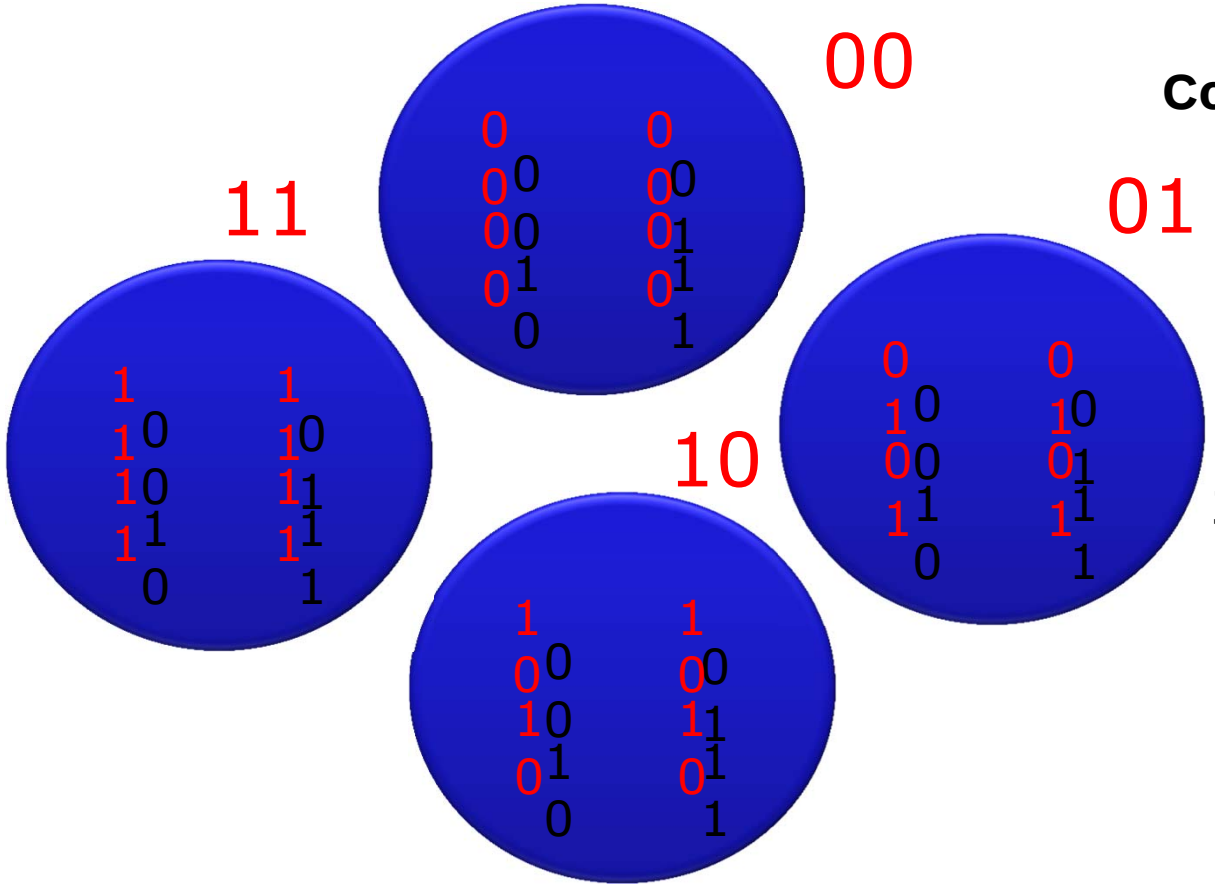

Code Book Construction:

1. Alice generates $2^{NR_{AB}}$ random codewords.
2. She splits them randomly into $2^{NR_{AE}}$ bins.
3. Codebook is broadcast to everybody, including Eve

Wiretap Code Encoding

$$R_o = R_{AB} - R_{AE} = 2 \qquad R_{AB} = \log(1+SNR_{AB}) = 4 \qquad R_{AE} = \log(1+SNR_{AE}) = 2$$

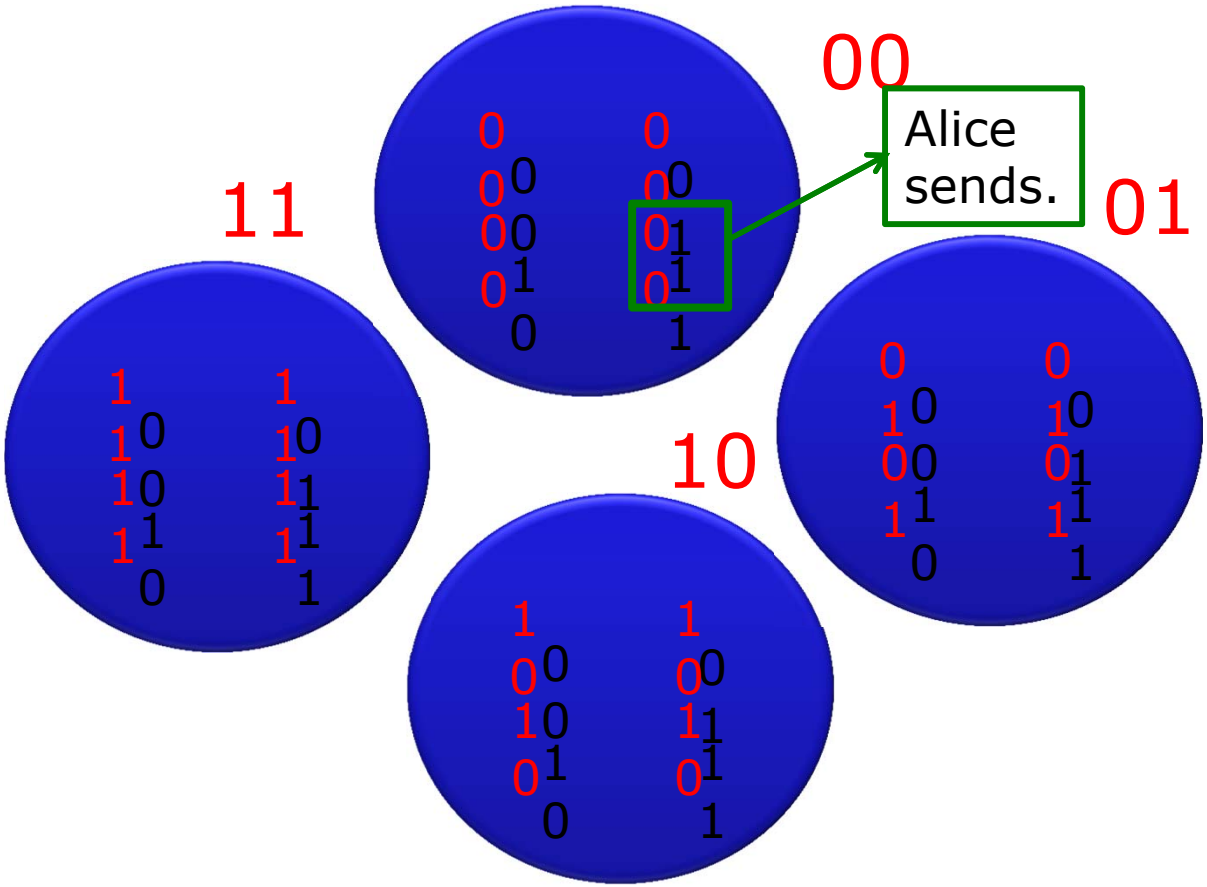

Alice picks a circle at random and uses R_o information bits to pick a codeword within that circle.

Wiretap Code Decoding

$$R_{AB} = \log(1+SNR_{AB}) = 4$$

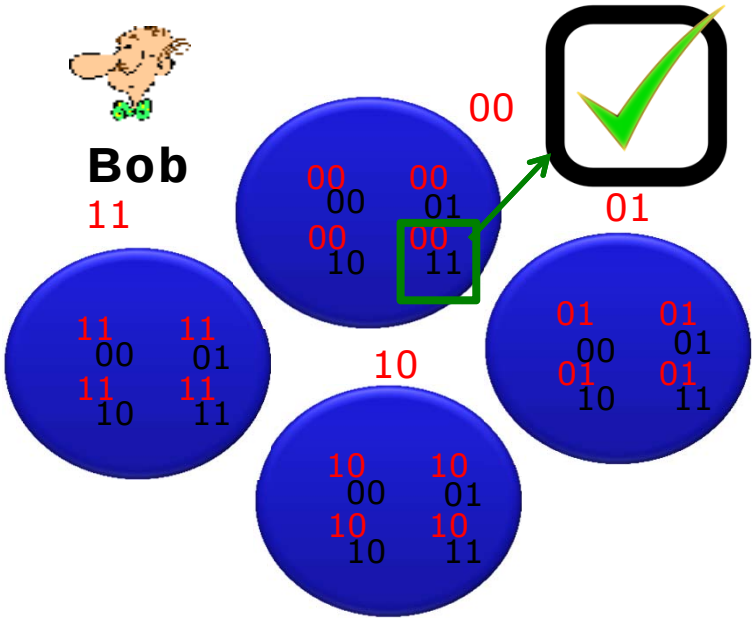

Bob's able to decode information bits 11 corresponding to codeword 0011

$$R_{AE} = \log(1+SNR_{AE}) = 2$$

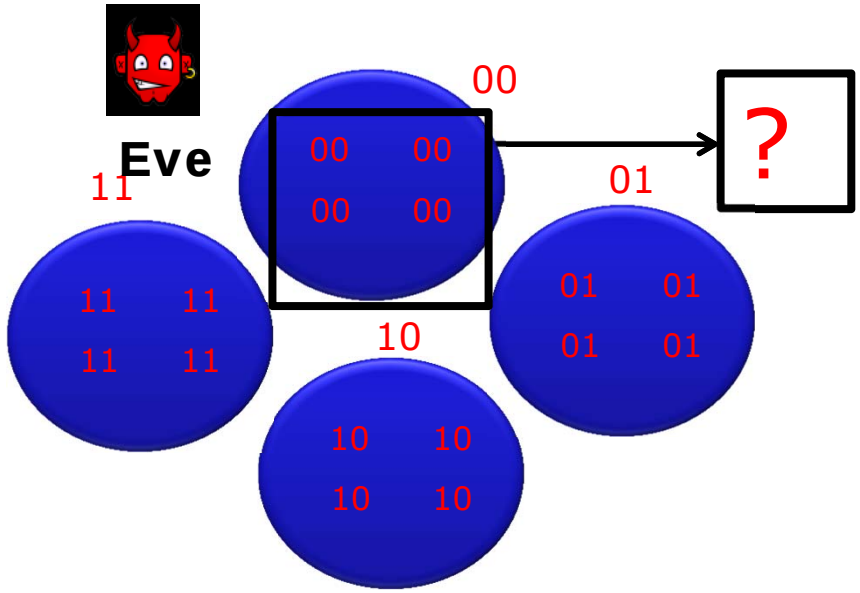

Secrecy capacity is given by the difference in the capacities between the main channel and the eavesdropper channel.

Secrecy Outage Event

Outage
$R_S < 2$

$R_S = 2$

Secrecy
$R_S > 2$

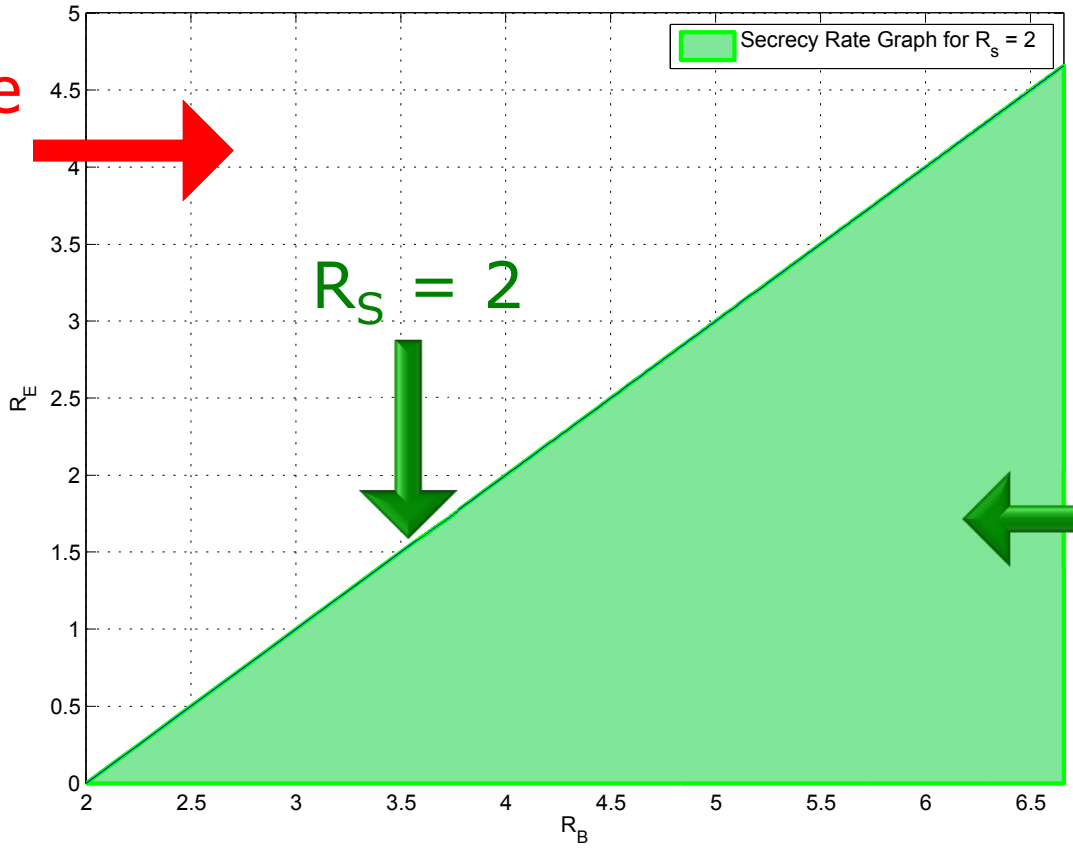

Eve Channel Capacity

Bob Channel Capacity

Critical Issue in Trying to Plan Secrecy: What do we Know?

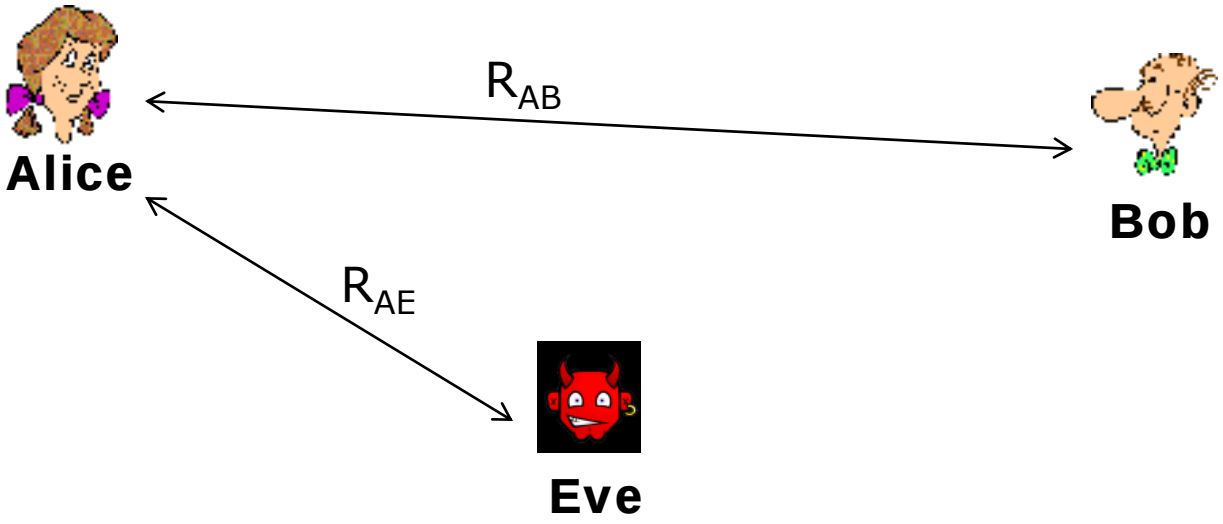

Do we know R_{AB}? Maybe.

Do we know R_{AE}? Probably not.

Known CSI to Bob

(known R_b puts on this line)

Outage
$R_S < 2$

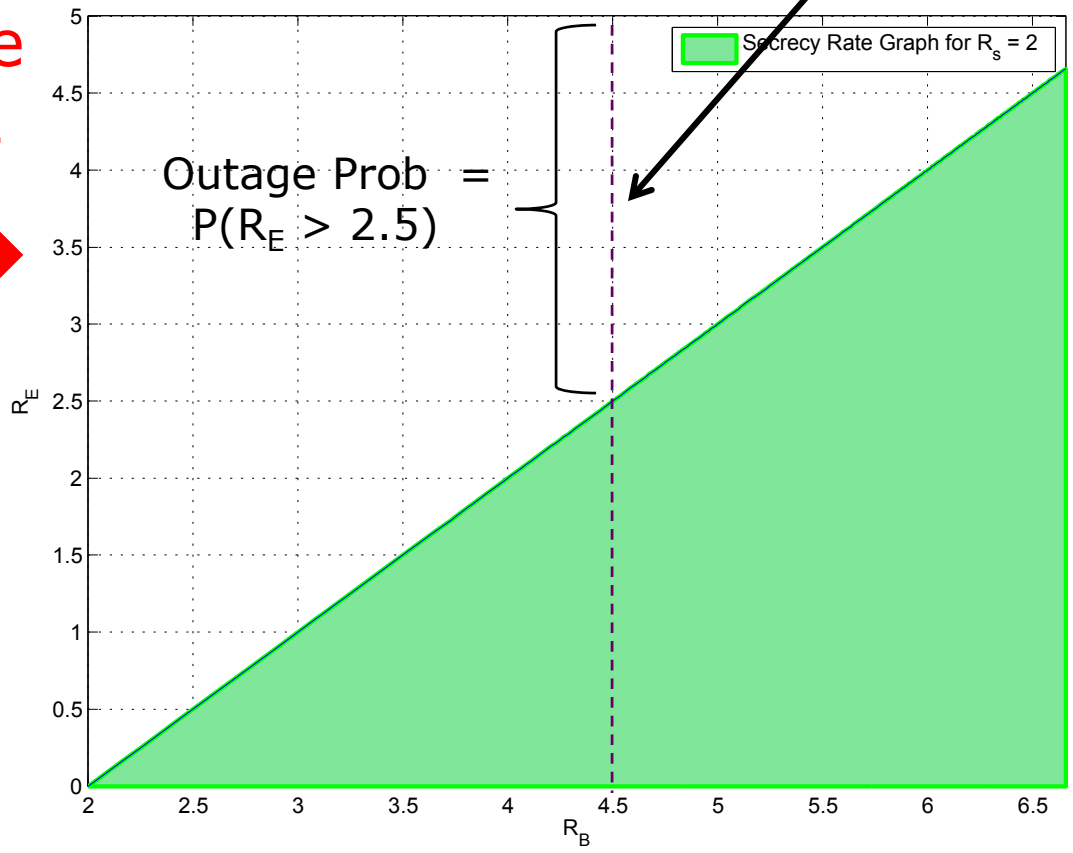

Eve
Channel
Capacity

Secrecy
$R_S > 2$

Bob Channel Capacity

Wireless Channels: Path Loss

What happens to the transmitted wave on the way from the cell phone to the tower?

Goes in all directions (**broadcast**) and the signal strength **weakens**. Let's model it.

$$P_r \propto \frac{P_t}{d^n}$$

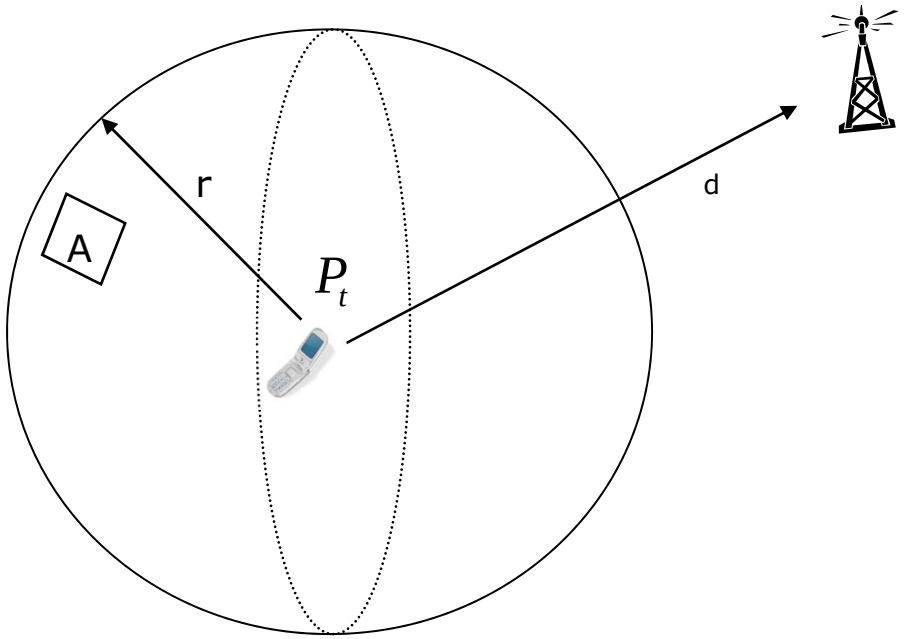

Table 4.2 Path Loss Exponents for Different Environments

Environment	Path Loss Exponent, *n*
Free space	2
Urban area cellular radio	2.7 to 3.5
Shadowed urban cellular radio	3 to 5
In building line-of-sight	1.6 to 1.8
Obstructed in building	4 to 6
Obstructed in factories	2 to 3

T.S. Rappaport, Wireless Communications

Note that the differences in received powers can be **huge**, for example, in a cell phone system:

$$(100m/2000m)^3 = 39dB!$$

Small Scale: Multi-path Fading

What happens to the signal on the way from the cell phone to the tower?

It gets reflected by many objects and the reflections add up at the receiver.

Example:

- Two paths of different lengths, signals arrive at slightly different times.
- One path is 100ft longer: 100ns difference (*big deal?).*
- Carrier might be at 1 GHz -> period 1ns (So, *yes, big deal)*
- Walk around the room while you speak on the phone -> the two signals keep adding up or canceling at the receiver as you move.

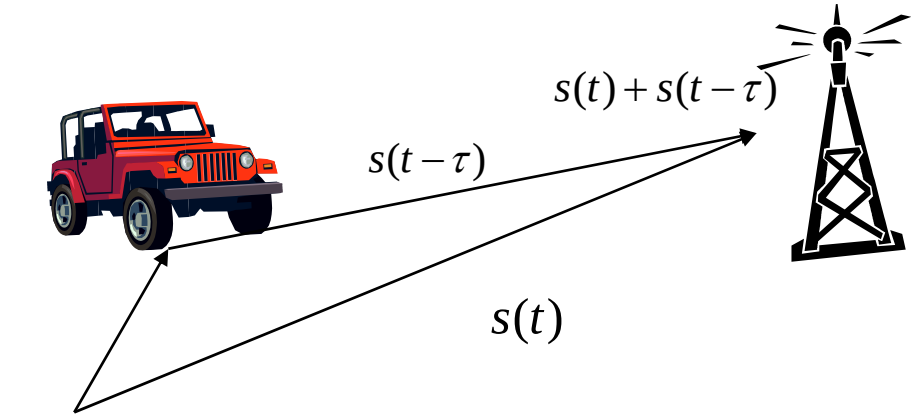

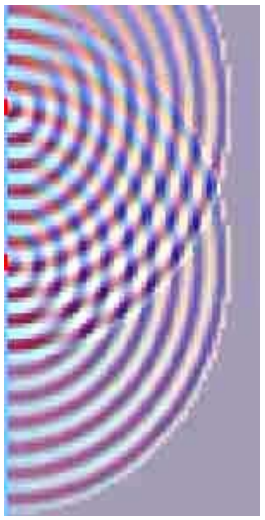

Key: Varies 3 ways:

1. **Spatially**
2. **Temporally**
3. With frequency

Alice

Bob

Building

Eve

Computational Security:

Drawback: "Only" computational security.

Plus: No problem with the "Near Eve".

Information-Theoretic Security:

$$R = \log_2(1 + SNR_{AB}) - \log_2(1 + SNR_{AE})$$

Plus: Positive rate "if Bob's channel is better", and Eve gets nothing.

Drawback: Zero rate "if Eve's channel is better", and Bob gets nothing.

Question: Can we get the best of both worlds.

Outline

1. Computational and Information Theoretic security basics
2. **Potential solutions**
   a. **Exploiting fading**
   b. **Two-way communications**
   c. **Attacking the receiver's hardware**
   d. **Cooperative jamming**
3. Asymptotically-large networks
4. Undetectable communications (LPD)
5. Current and Future Challenges

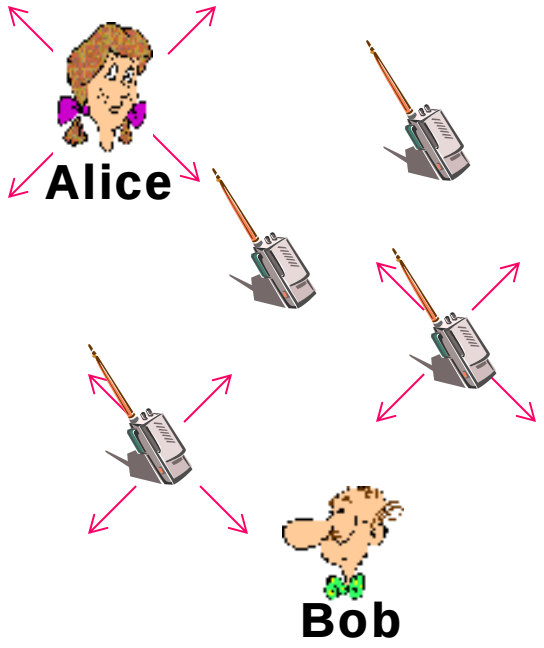

Exploiting Fading I: signal when we have the advantage

$$R = \log_2(1 + P_{AB}) - \log_2(1 + P_{AE})$$

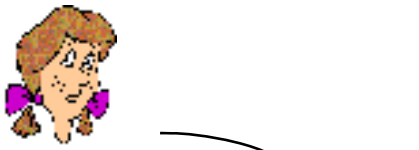

Alice

Eve

Bob

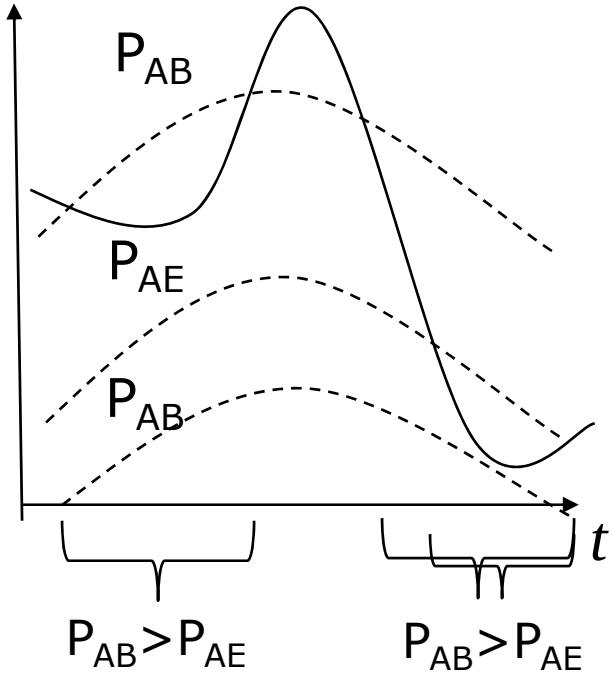

CSI = Path-loss and fading

Situations: (1) Known CSI, (2) Partial CSI (all Bob, path-loss to Eve), **(3) No CSI**

Problem: If I do not know where Eve is, how do I choose a strategy/rate?

Exploiting Fading II: derive a key from it

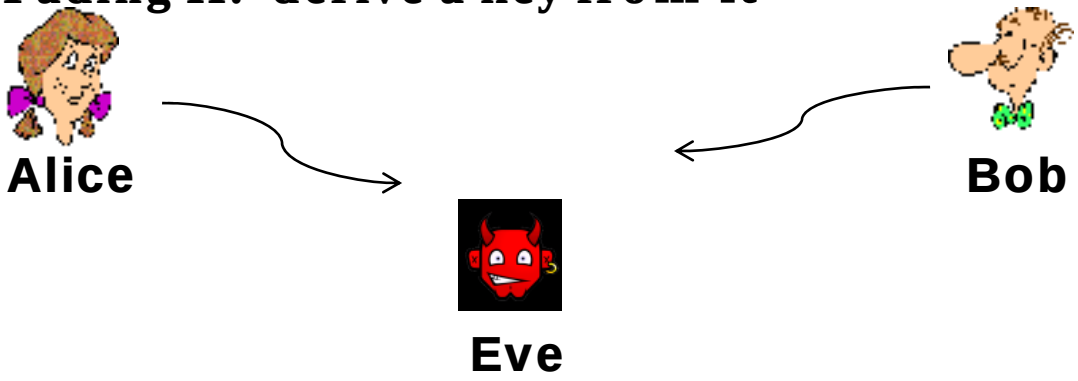

1. Alice broadcasts a pilot signal. Bob measures the channel H_{AB}

2. Bob broadcasts a pilot signal. Alice measures the channel H_{BA}

3. Assuming reciprocity, $H_{AB}=H_{BA}$, and we have a source of common randomness. Alice and Bob reconcile their channel estimates to form a common key K.

4. Alice broadcasts with a one-time pad: X=M XOR K

Drawback: Limited number of bits can be extracted

Public Discussion [Maurer, 1993][Ahlswede and Csiszar, 1993]

Alice

Bob

Eve is closer than Bob.

Consider binary symmetric channels (0 or 1 in, 0 or 1 out):

Eve

Bob's Channel (p_B:prob of error)

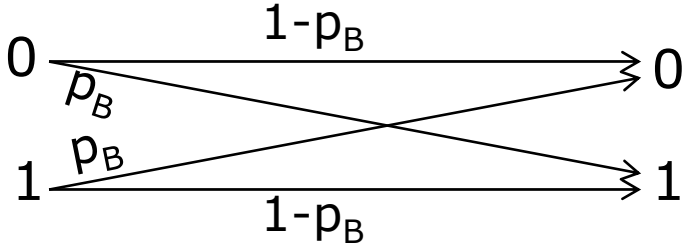

Eve's Channel (p_E:prob of error)

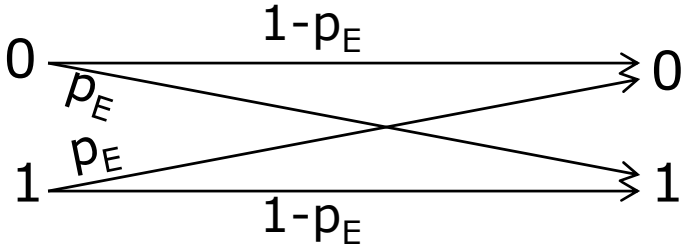

$p_B > p_E$ (since Eve is closer); hence, the secrecy capacity is zero. What to do?

Public Discussion II

M

Alice

Bob

Eve

…and we have a channel of positive capacity.

But it could be a really small capacity… If Eve is close to Bob…and how do you choose the rate?

1. ***Bob*** transmits X:

   Alice Receives: $Y = X \oplus D$

   Eve Receives: $Z = X \oplus E$

2. Alice transmits (on noiseless, public channel):

   $M \oplus Y = M \oplus X \oplus D$

   Bob Receives: $M \oplus X \oplus D$

   Eve Receives: $M \oplus X \oplus D$

3. Bob forms:

   $(M \oplus X \oplus D) \oplus X = M \oplus D$

   Eve forms:

   $(M \oplus X \oplus D) \oplus X = M \oplus D \oplus E$

Public Discussion III

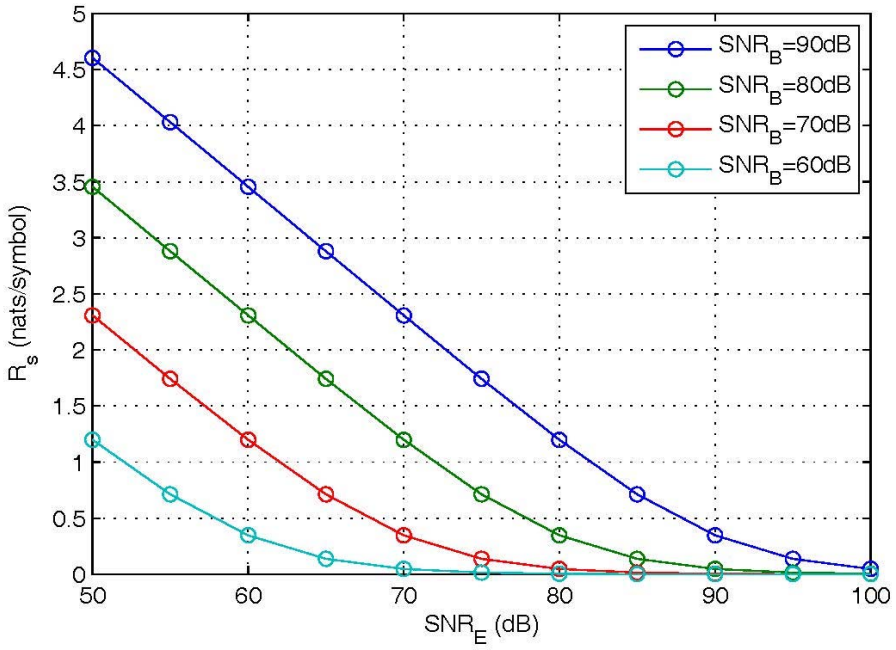

Problem: Rate becomes limited for a very near Eve.

What if Eve picks up the transmitter?

Challenges

1. Exploiting when Alice -> Bob **channel is better** than Alice -> Eve
   Challenge: unknown Eve location
2. Exploiting common randomness of **channel reciprocity**
   Challenge: limited number of key bits
3. Exploiting "**public discussion**"
   Challenge: two-way communication and unknown Eve

Attacking the Hardware I: Bounded Memory Model

Cachin and Maurer introduced the "bounded memory model" to achieve everlasting secrecy [Cauchin and Maurer, 1997]. An eavesdropper with memory < M cannot store enough to eventually break the cipher.

However, it is hard to pick a memory size that Eve cannot use beyond:

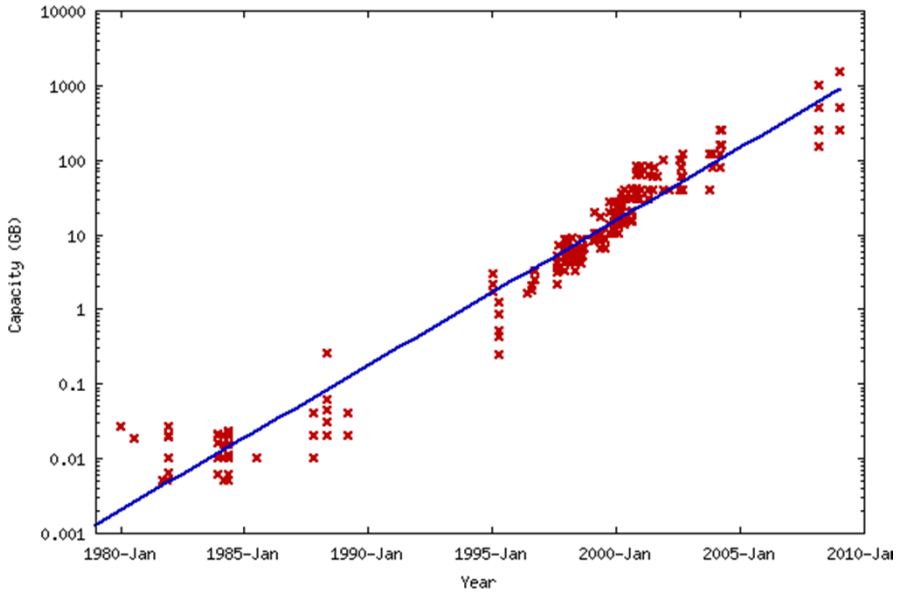

[From "blog.dshr.org"]

1. The density of memories grows quickly (Moore's Law)

2. Memories can be stacked arbitrarily subject only to (very large) space limitations.

Potential Solutions: (a) Exploiting fading (b) 2-way comms (c) Rcvr hardware (d) Jamming

Bounded Conversion Model

Perhaps Cachin and Maurer attacked the wrong part of the receiver:

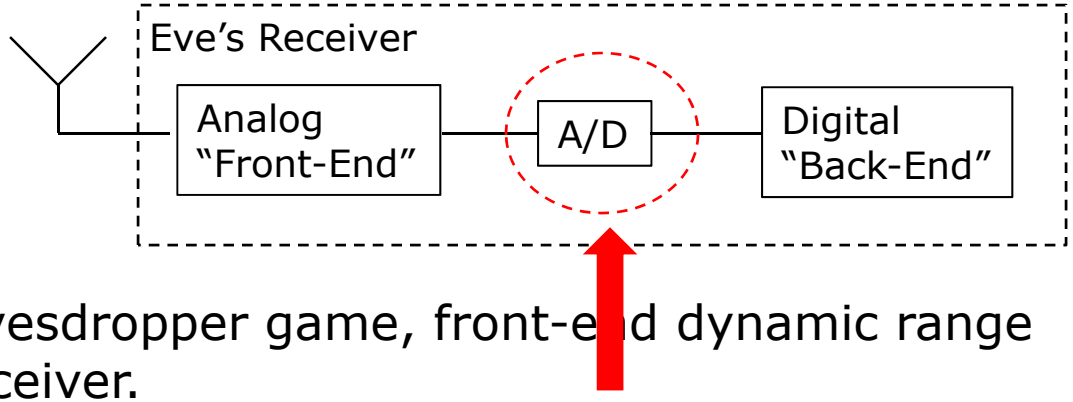

1. In the combative sender-eavesdropper game, front-end dynamic range is a critical aspect of the receiver.
2. A/D Technology progresses very slowly.
3. High-end A/D's are already stacked to the limit of the jitter.

Idea:

Alice

Eve

Bob

1. IT security requires a channel advantage.
2. Establish cryptographic security (e.g. Diffie-Hellman)
3. Use a *short-term* cryptographic to establish the channel advantage.

System model and approach:

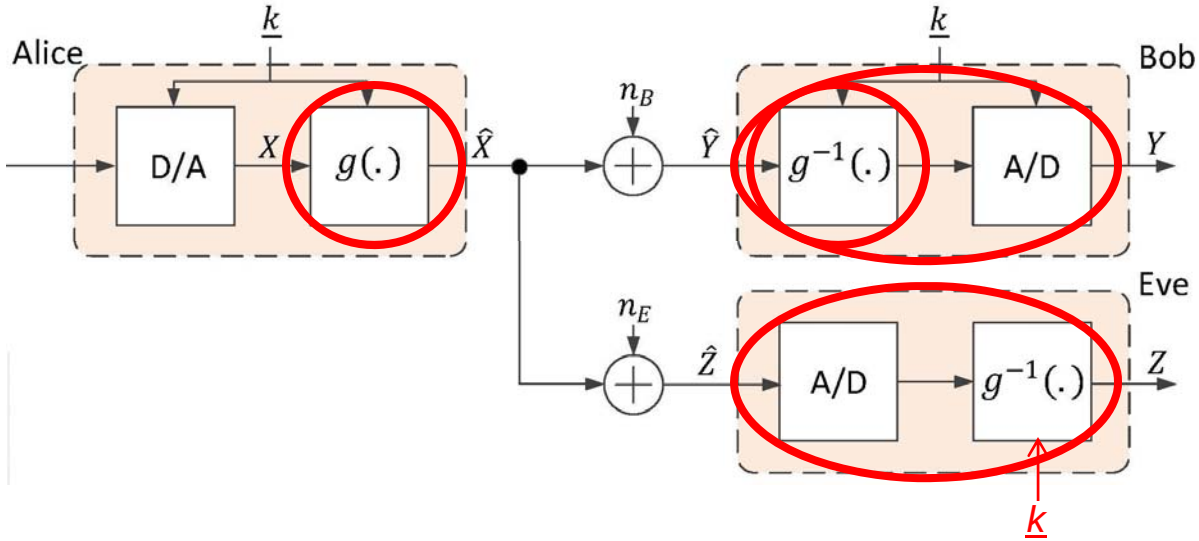

1. Alice and Bob pre-share an "emphemeral" cryptographic key $\underline{k}$ to choose $g(.)$. **Note:** Key will be handed to Eve after transmission.
2. A/D is a non-linear element. Non-commutativity of non-linear elements: potential information-theoretic security.
3. Secrecy rate is a shaping gain: $R_s = E_g[h(X) - h(g(X))]$
   h(X): differential entropy

…but, unlike traditional "shaping gains", gain can be huge.

Rapid power modulation for secrecy:

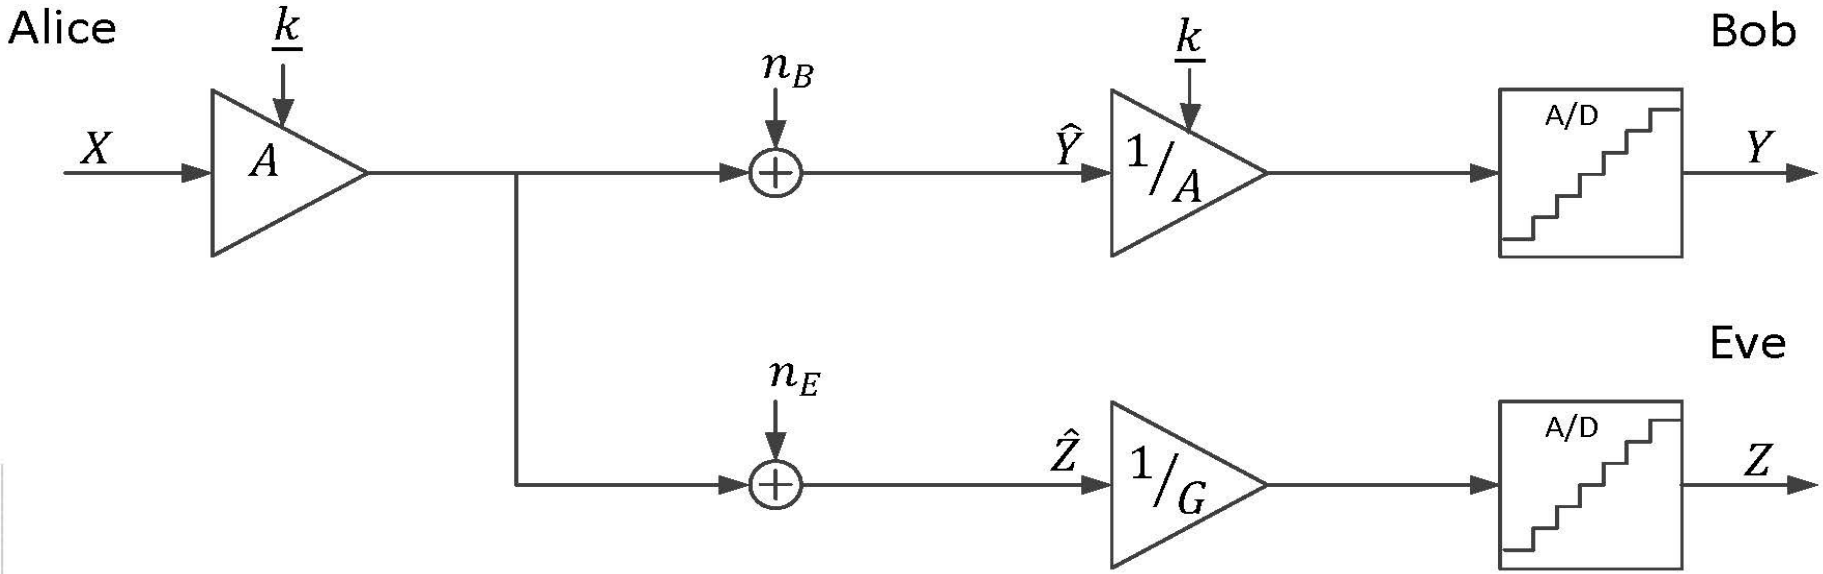

Idea: Key used to rapidly power modulate transmitter. Bob's receiver gain control can follow, while Eve's struggles.

Rapid power modulation for secrecy:

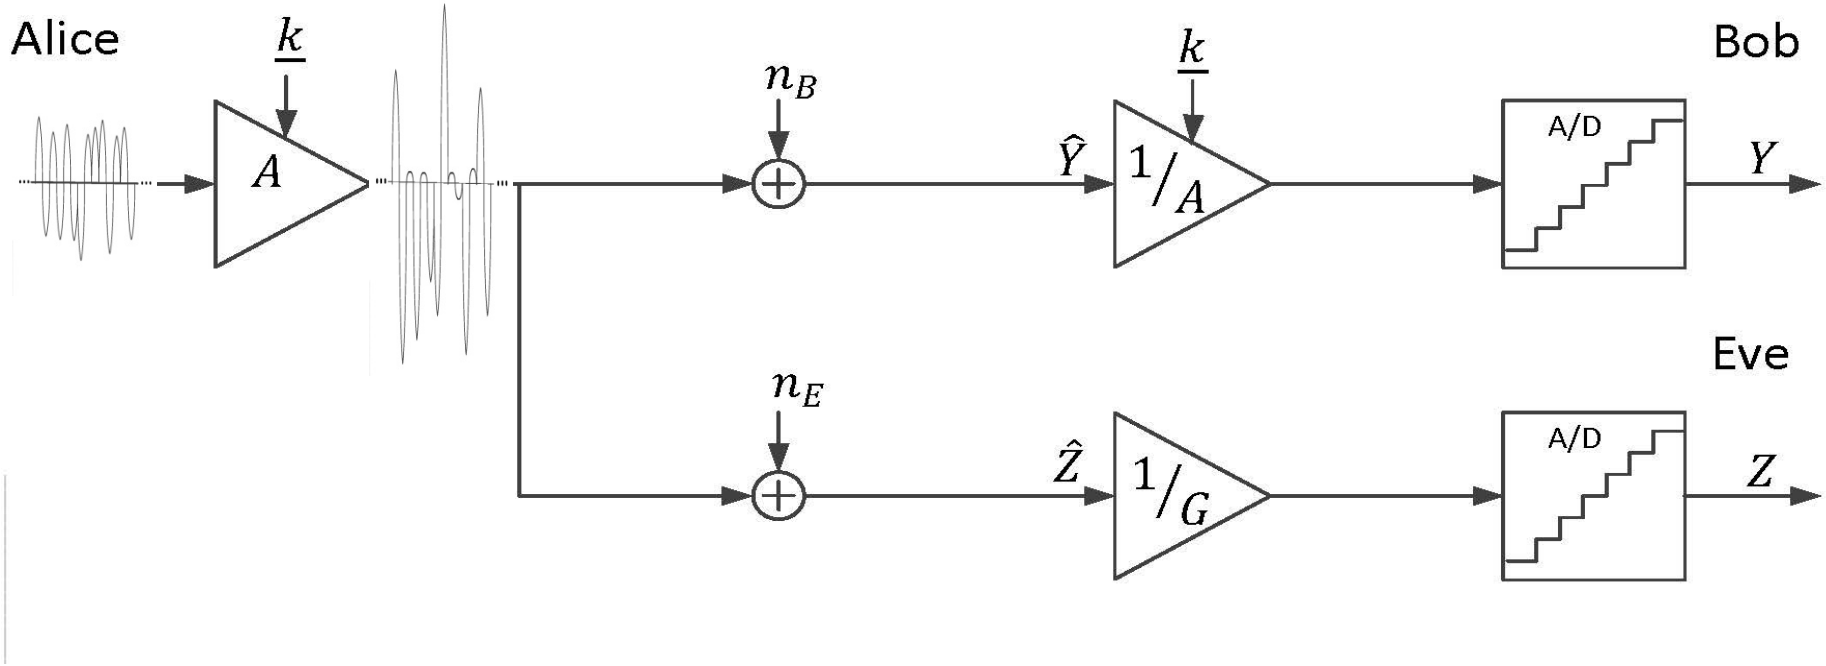

Rapid power modulation for secrecy:

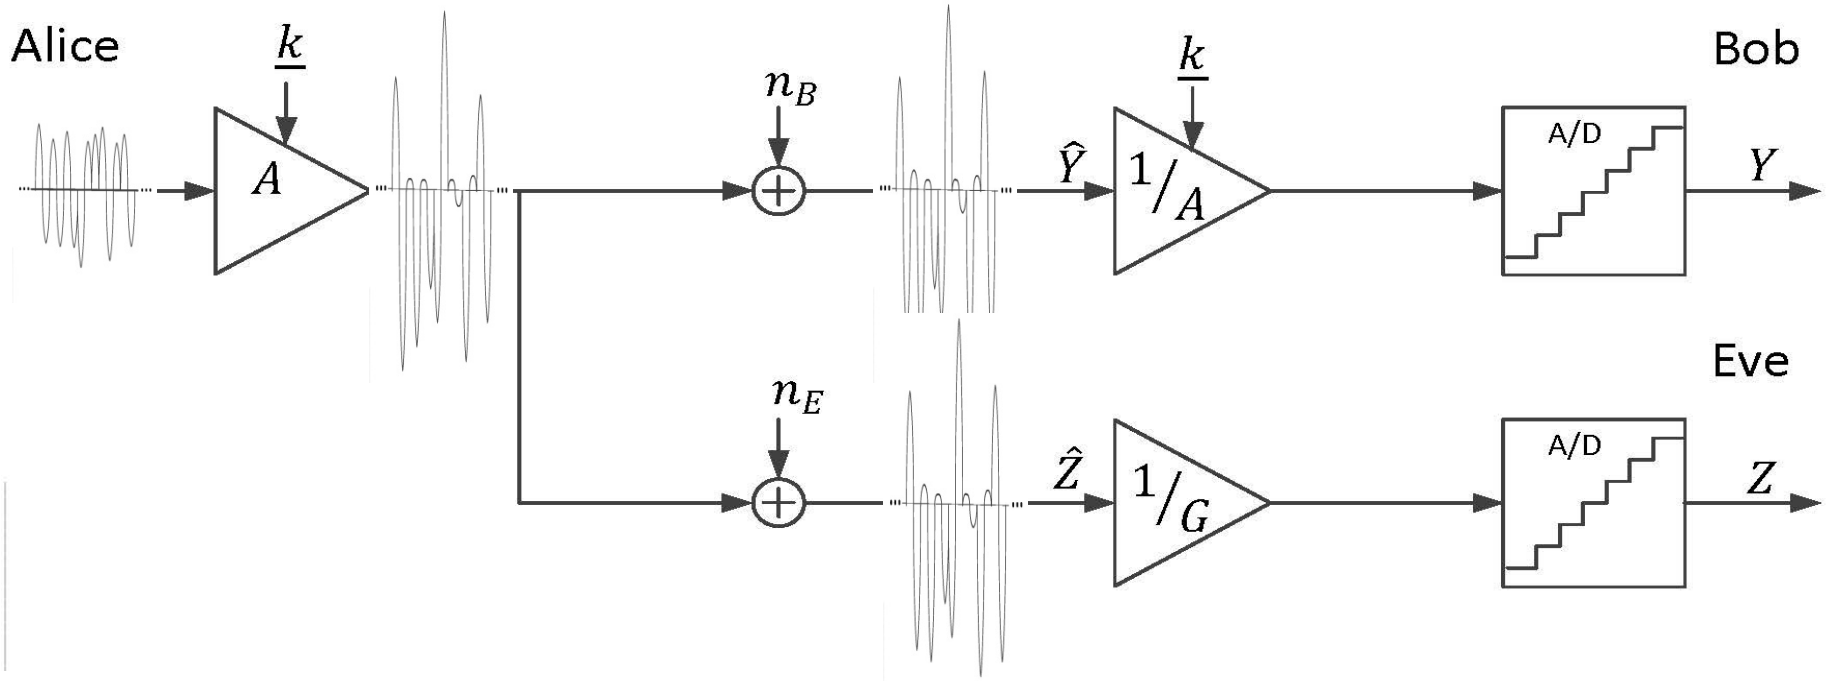

Rapid power modulation for secrecy:

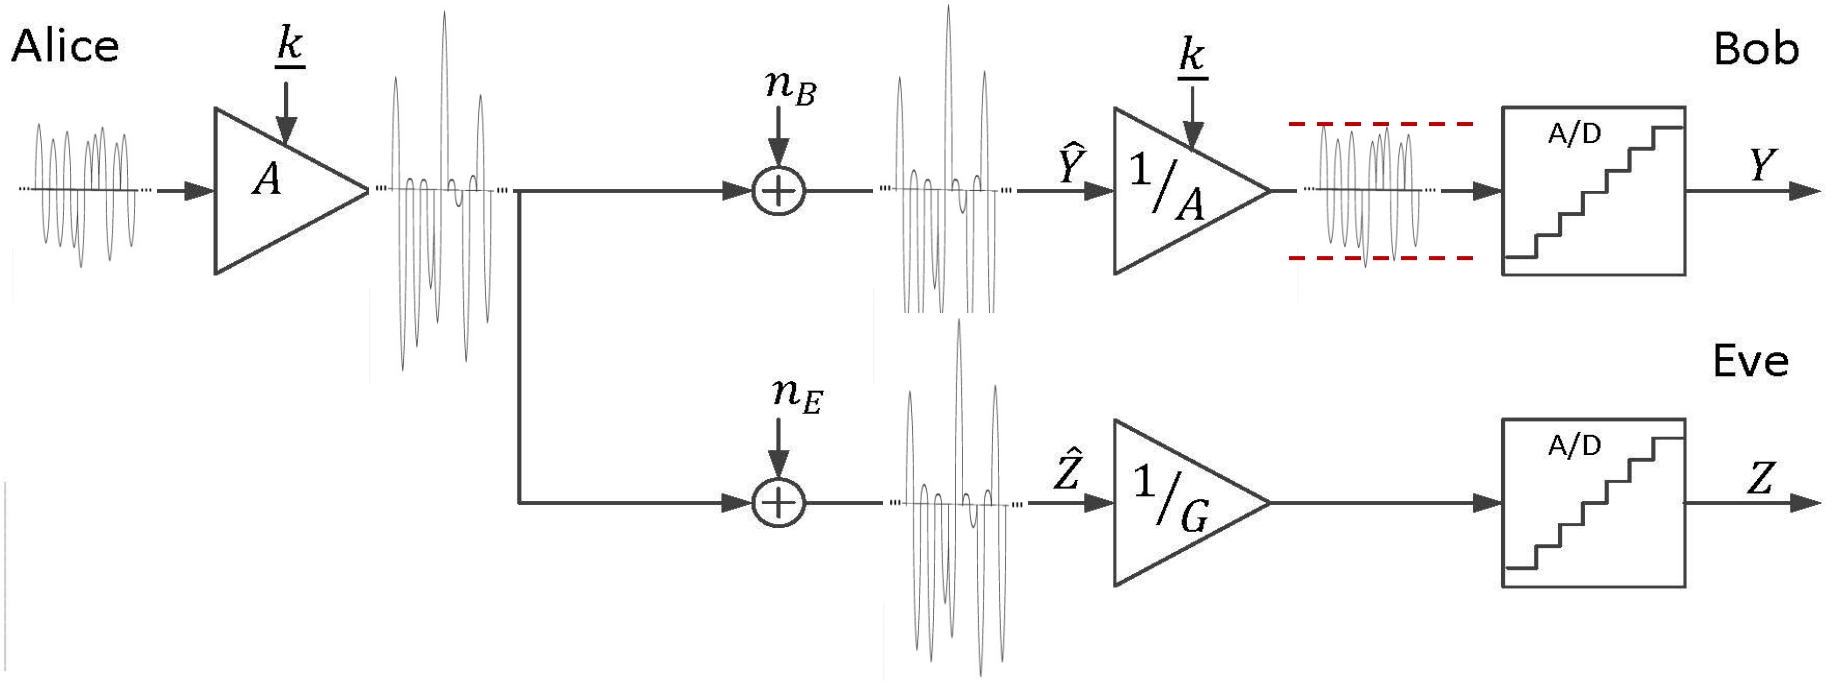

Bob's gain control is correct: input well-matched to A/D span.

Rapid power modulation for secrecy:

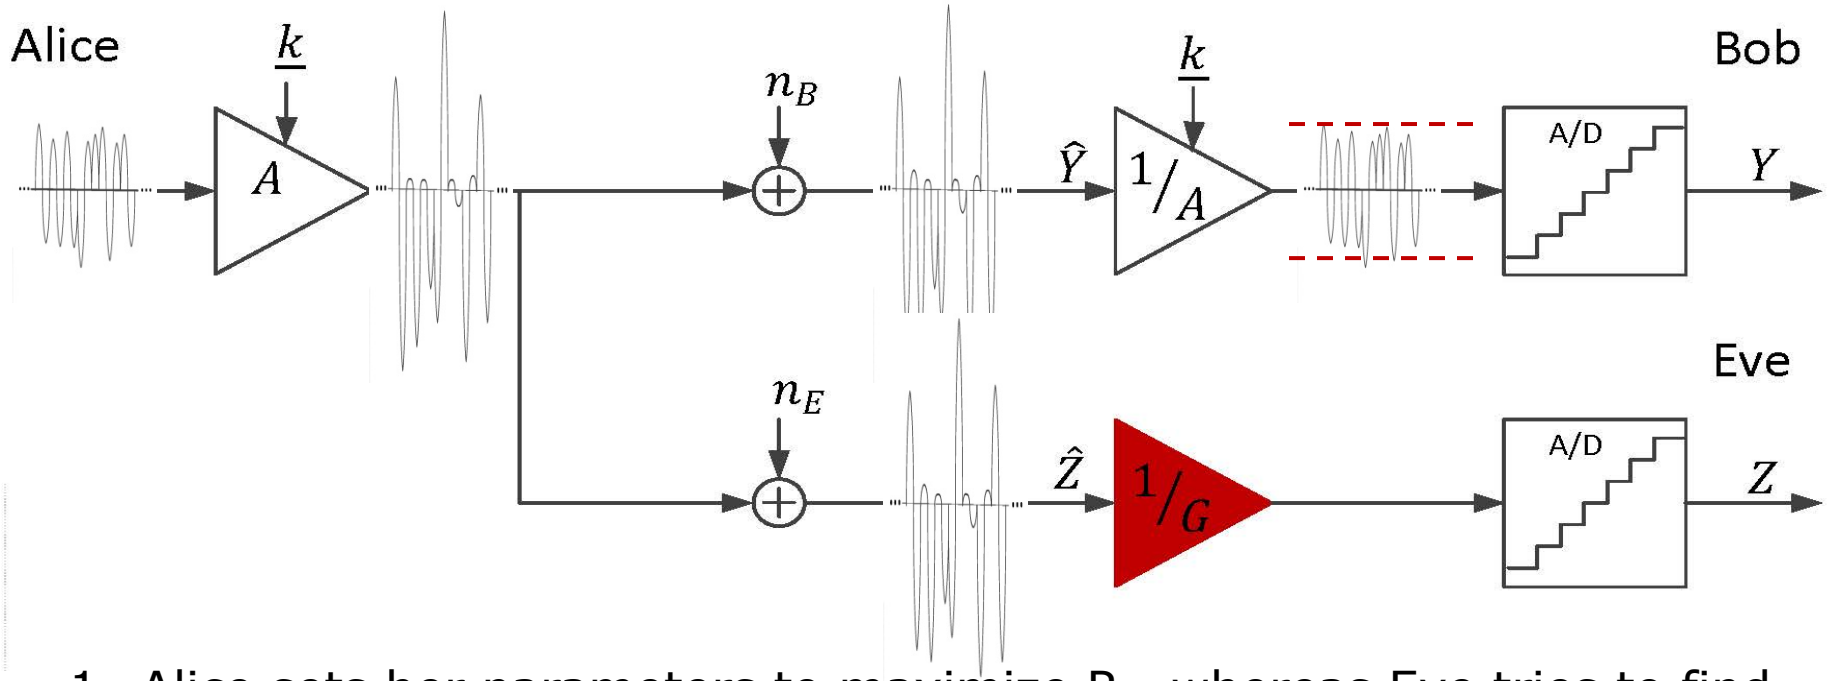

1. Alice sets her parameters to maximize R_s, whereas Eve tries to find a gain G that minimizes the secrecy rate R_s given Alice's choice:

$$R_s = \max_S \min_G R_s(S, G)$$

2. It is easy to show that the optimal strategy (for Eve) is to pick a single gain G.

Rapid power modulation for secrecy:

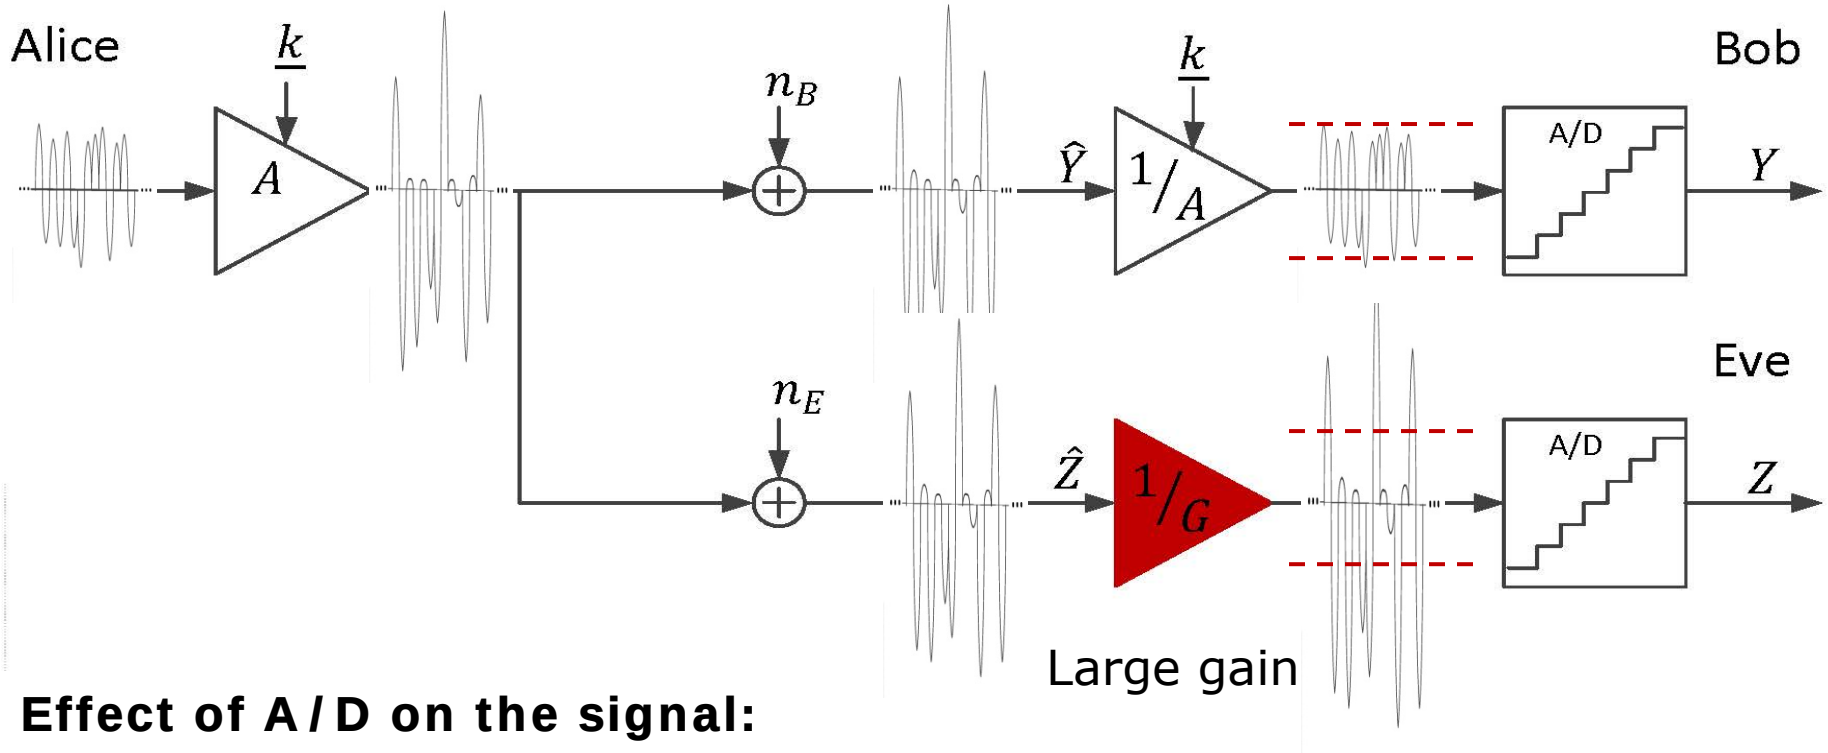

Effect of A/D on the signal:

- Clipping (due to overflow)

Rapid power modulation for secrecy:

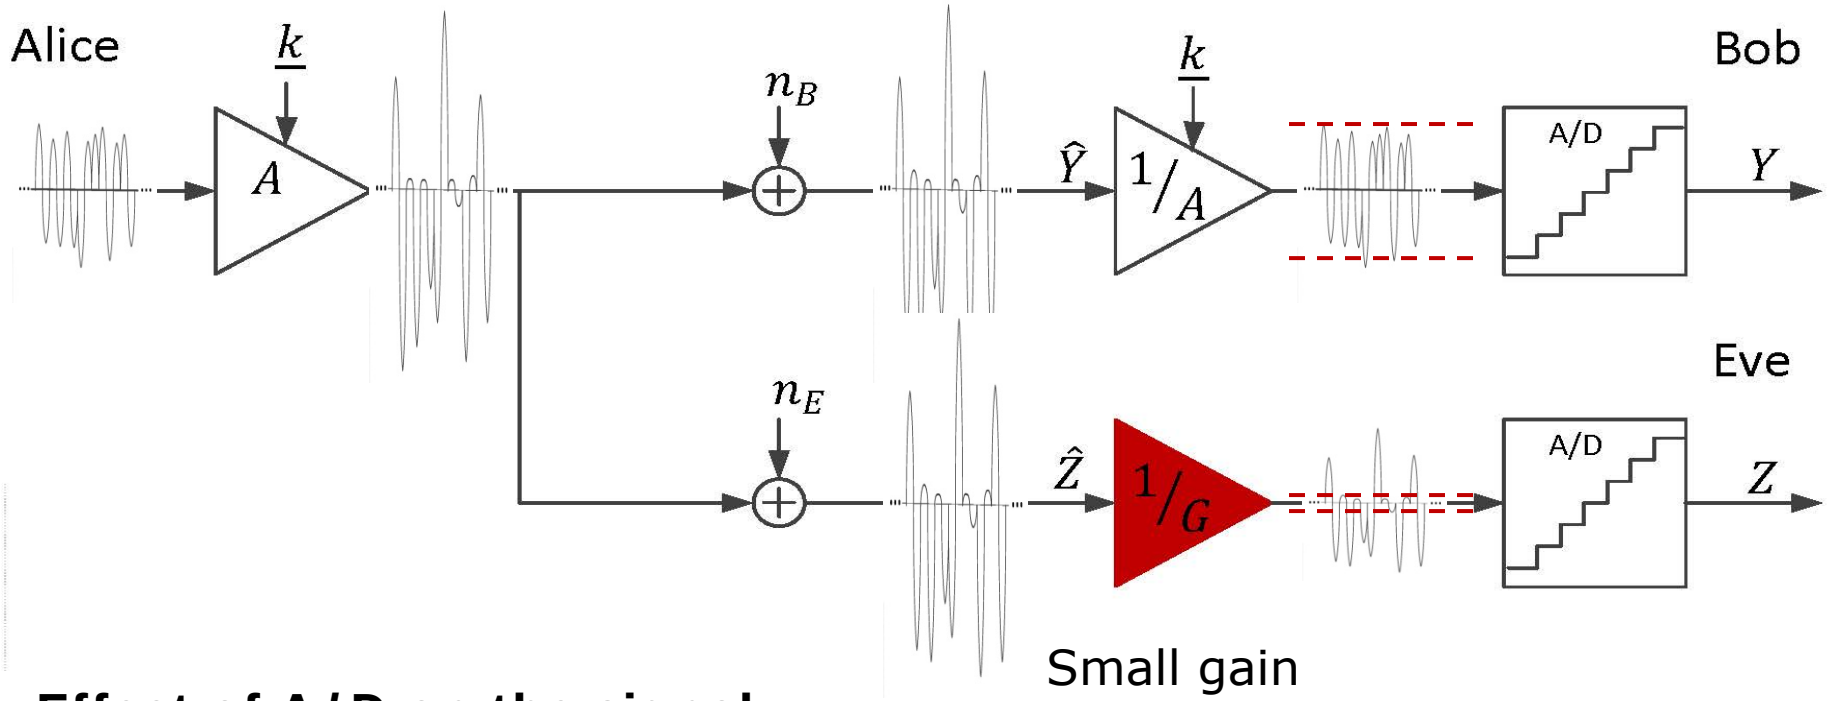

Effect of A/D on the signal:

- Clipping (due to overflow)
- Quantization noise (uniformly distributed)

Rapid power modulation for secrecy:

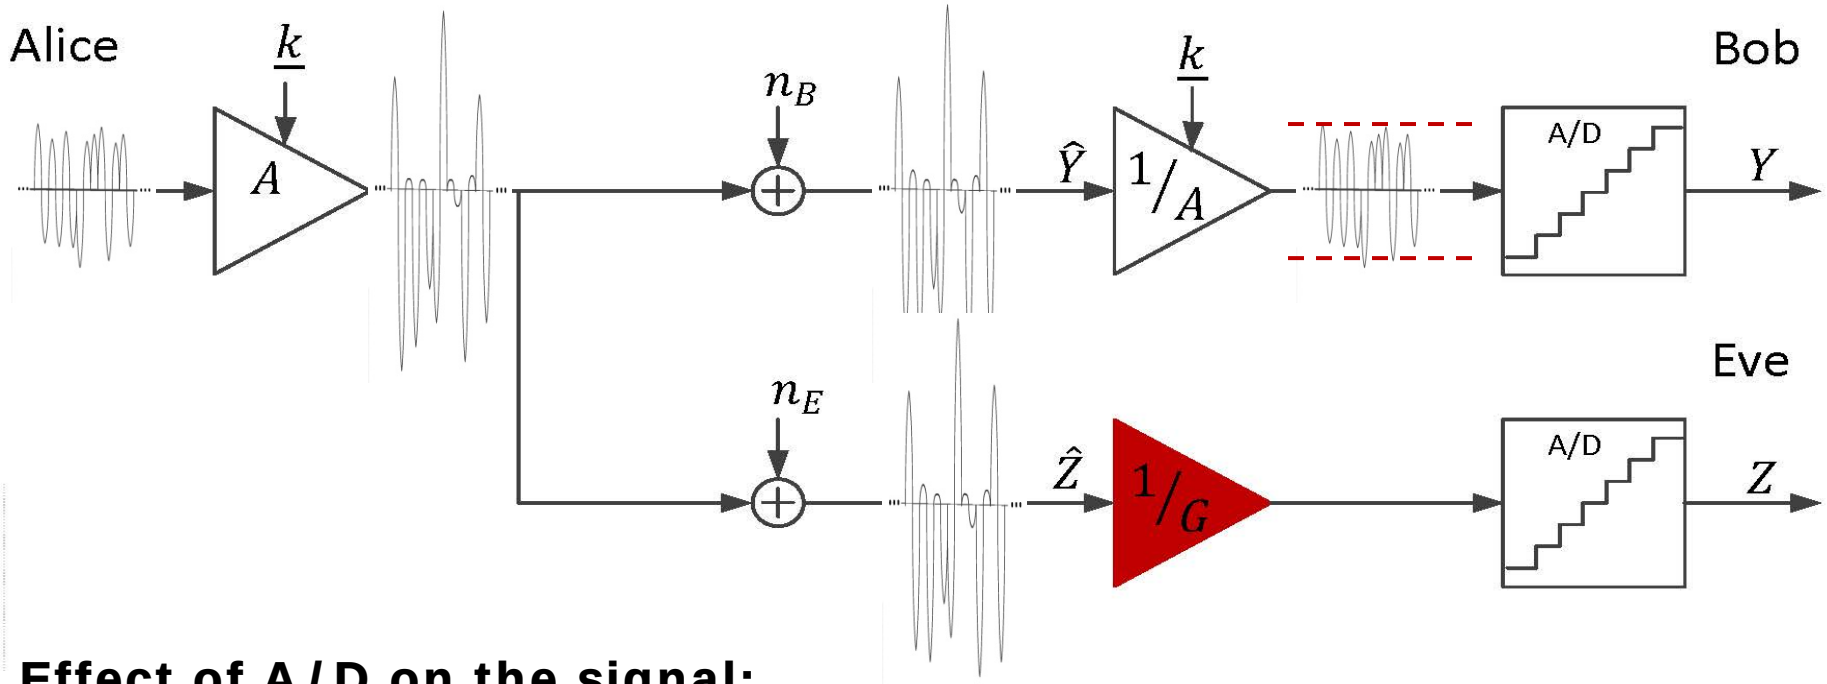

Effect of A/D on the signal:

- Clipping (due to overflow)
- Quantization noise (uniformly distributed)

Trade-off between choosing a large gain and a small gain:

- Eve needs to compromise between more A/D overflows or less resolution.

Potential Solutions: (a) Exploiting fading (b) 2-way comms (c) Rcvr hardware (d) Jamming

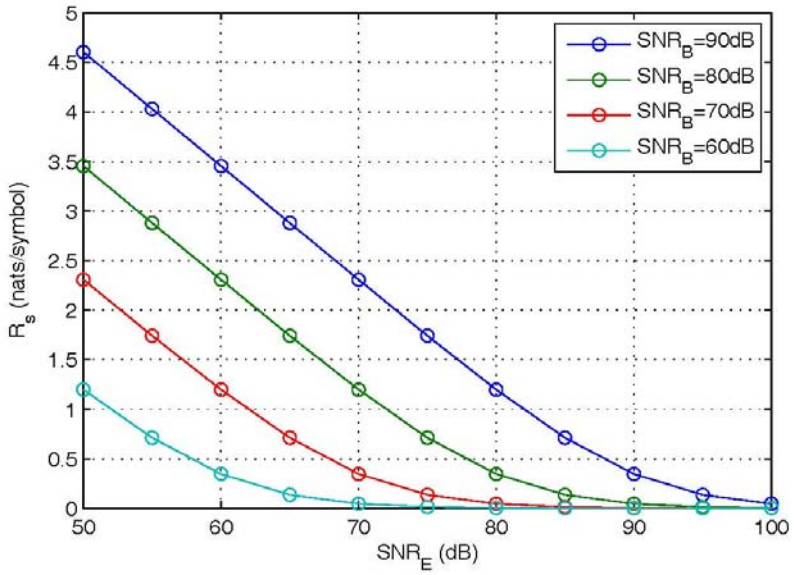

(a) Public Discussion

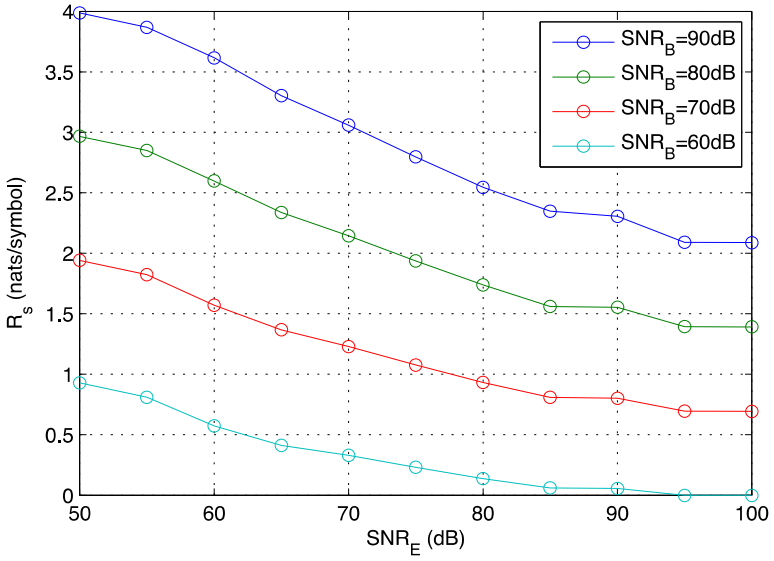

(b) Power modulation

(Although they are not really competing techniques. Power modulation approach could be used under public discussion.)

What if Eve picks up the transmitter?

Secrecy rate vs. SNR at Bob, Eve has perfect access to the signal

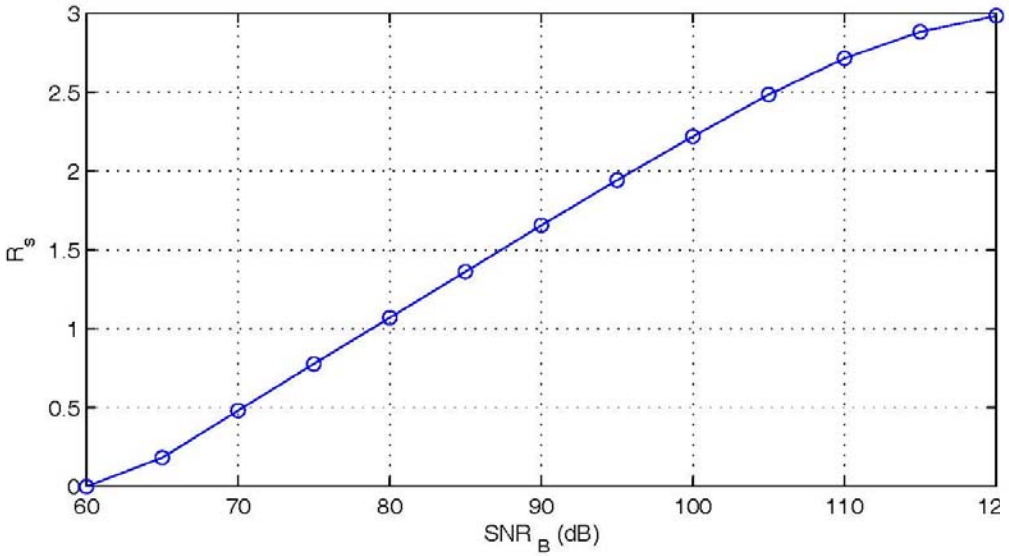

Noisy channel to Bob, **noiseless** channel to Eve.

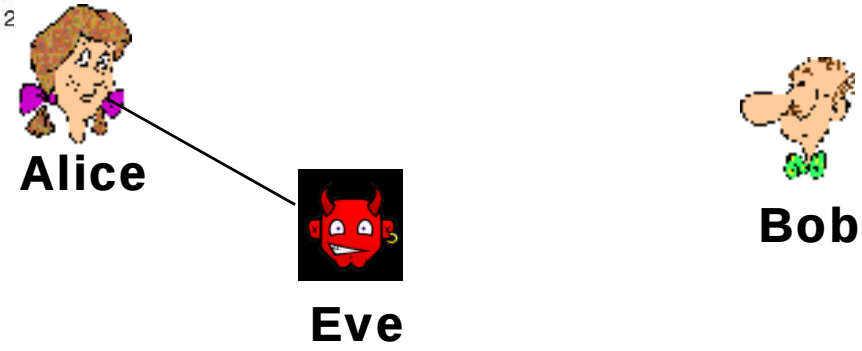

Challenges: (1) Only effective (so far) in short-range environments.
(2) Risk Eve has a better receiver than you thought.

Challenges

1. Exploiting when Alice -> Bob **channel is better** than Alice -> Eve
   Challenge: unknown Eve location
2. Exploiting common randomness of **channel reciprocity**
   Challenge: limited number of key bits
3. Exploiting "**public discussion**"
   Challenge: two-way communication and unknown Eve
4. Attacking Eve's **receiver hardware**
   Challenge: short range, assumptions on Eve's hardware

Cooperative Jamming for Secrecy [Negi/Goel, 2005]

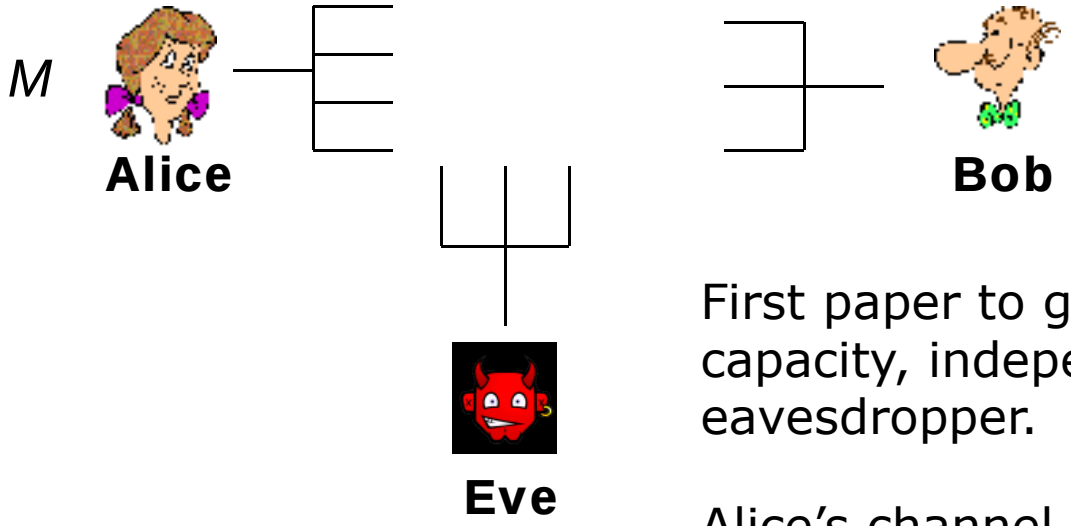

First paper to guarantee a minimum secrecy capacity, independent of the location of the eavesdropper.

Alice's channel state information knowledge:

1. She knows the channel to Bob
2. She does not know the channel to Eve

Idea: Jam in the null space of Bob's receiver.

Problem: Asymmetric capabilities are backward! (more powerful Alice than Eve)!

Cooperative Jamming for Secrecy II [Goel/Negi, 2008]

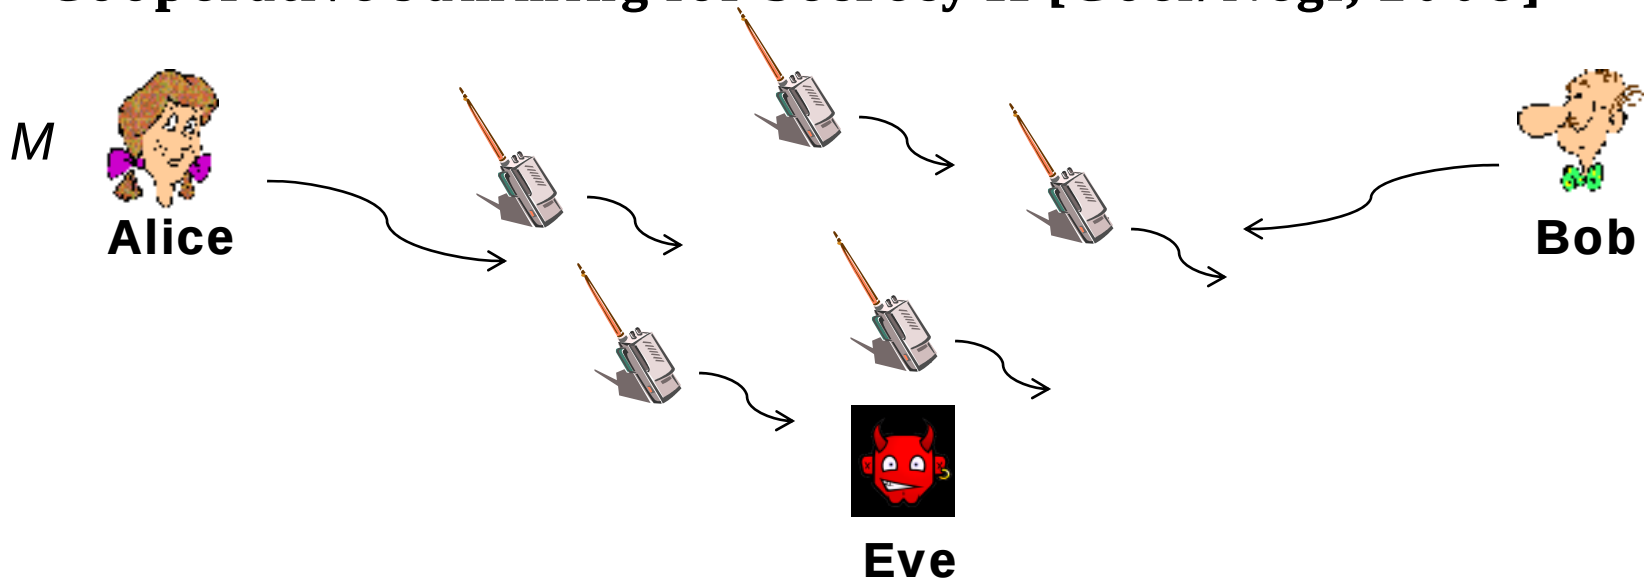

1. Stage 1: Alice and Bob send "noise messages"
2. Stage 2:
   - Alice sends the message plus a signal to cancel relay chatter.
   - Relays "chatter"

Relay chatter only affects Eve due to interference pre-cancellation by Alice.

Challenge: Interference cancellation challenging in a near-far environment.

Challenges

1. Exploiting when Alice -> Bob **channel is better** than Alice -> Eve
   Challenge: unknown Eve location
2. Exploiting common randomness of **channel reciprocity**
   Challenge: limited number of key bits
3. Exploiting "**public discussion**"
   Challenge: two-way communication and unknown Eve
4. Attacking Eve's **receiver hardware**
   Challenge: short range, assumptions on Eve's hardware
5. **Interference Cancellation**
   Challenge: near-far environment

Recall: Small Scale: Multi-path Fading

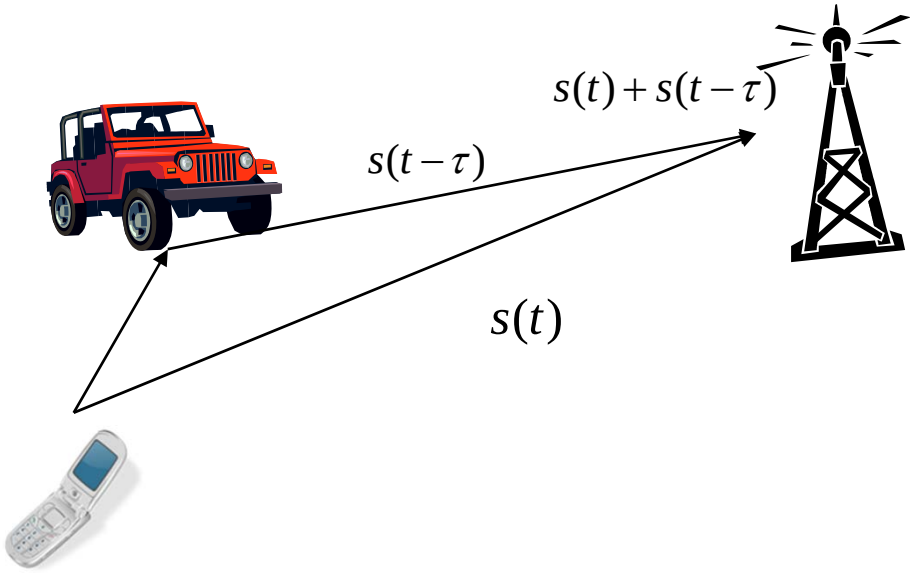

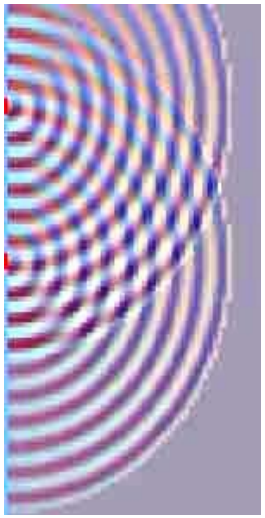

Key: Varies three ways:

1. **Spatially**
2. **Temporally**
3. With frequency

Exploiting spatial fading:

Could wait until channel until channel from Alice to Bob is really good…

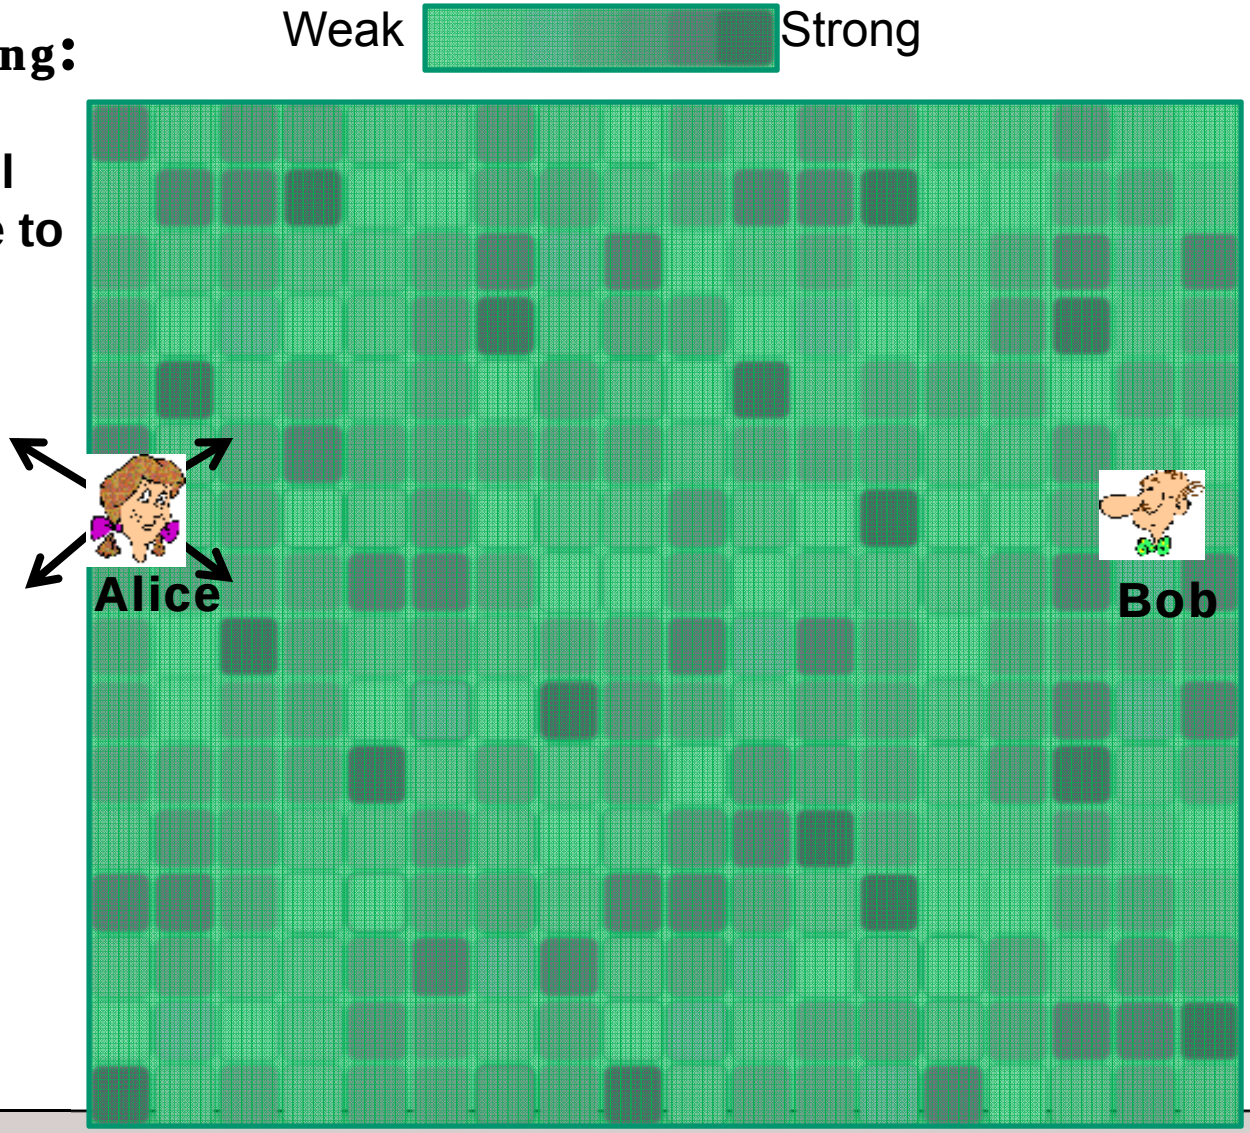

Potential Solutions: (a) Exploiting fading (b) 2-way comms (c) Rcvr hardware (d) Jamming

Could wait until channel until channel from Alice to Bob is really good…

…or could look at relays which yield a random spatial sampling of the field.

For Rayleigh fading, power at "best" relay of N relays goes as $P_{S,R}^{*} \sim \log N$.

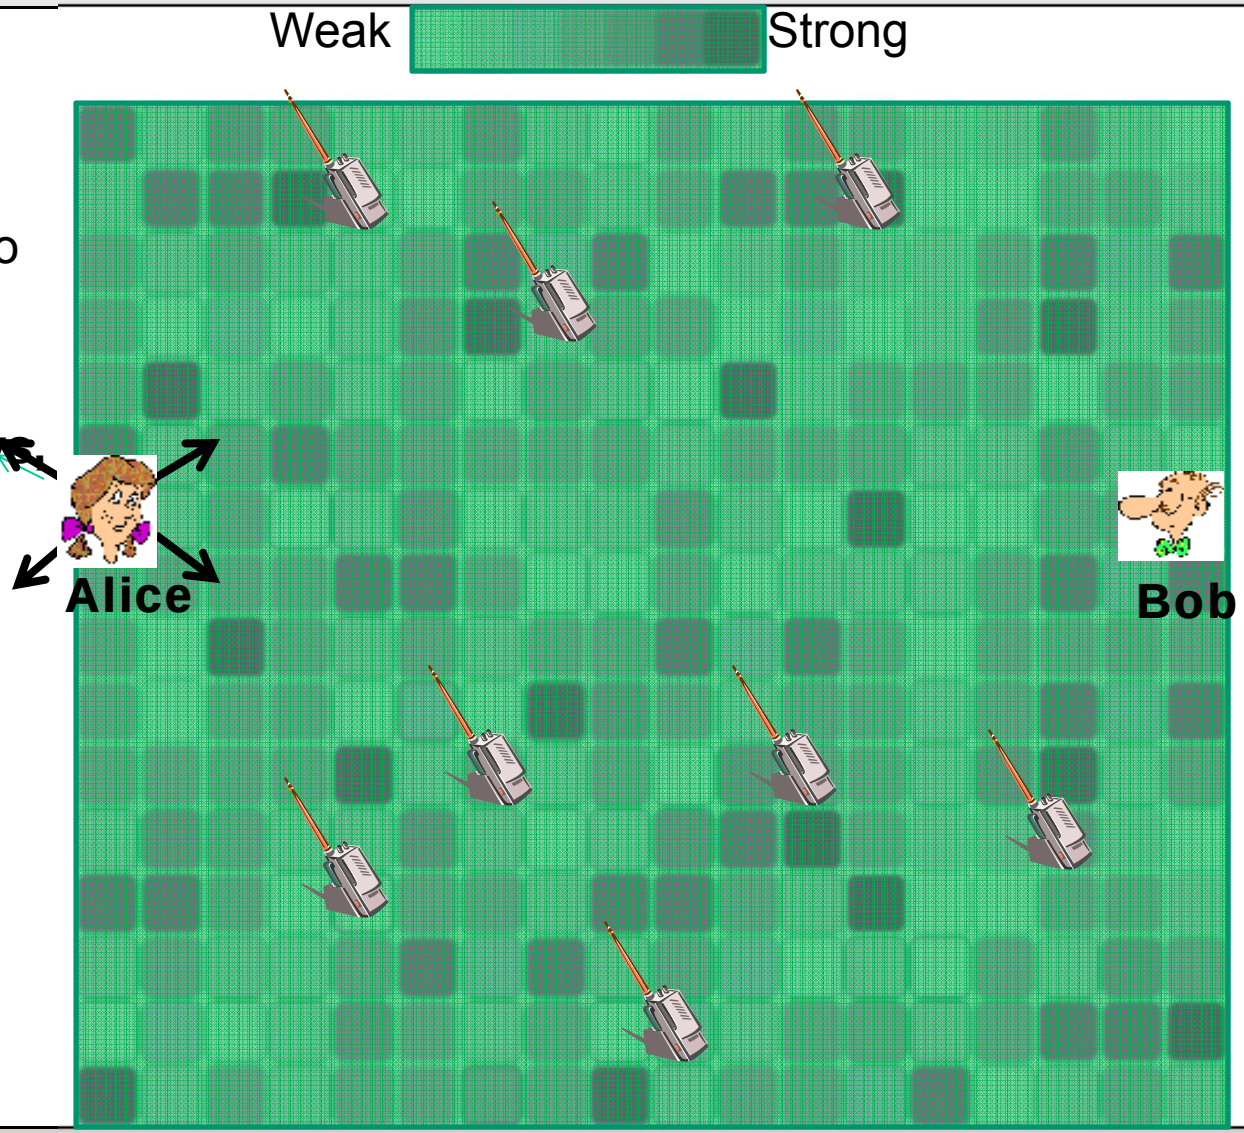

The "Two-Hop" Case

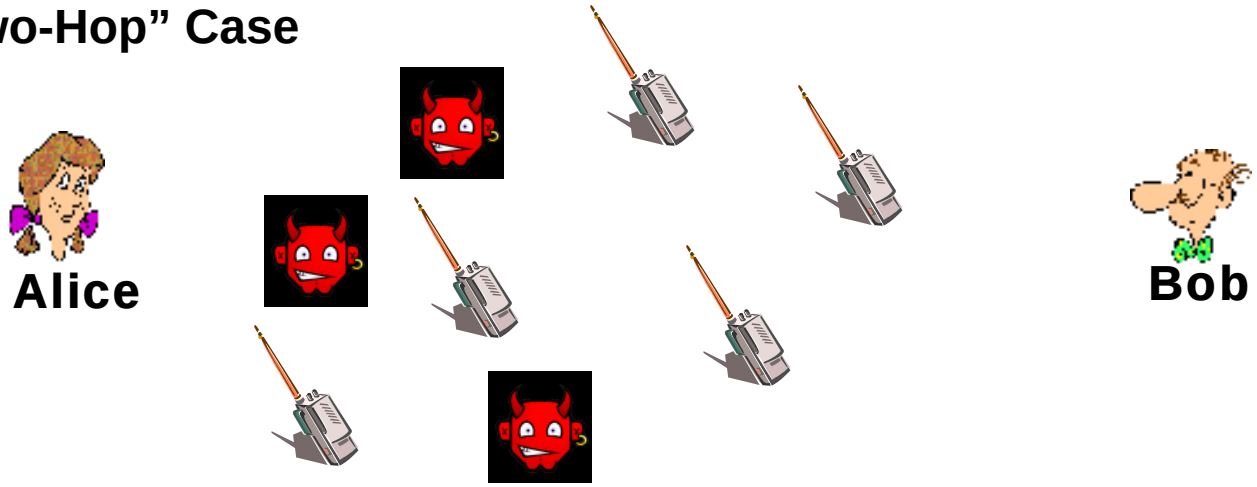

Scenario: Source "S" wants to communicate securely with the destination "D" in the presence of M eavesdroppers (M=3 above) with the help of N relays (N=5 above).

Goal: Deliver a packet from S to D of rate R, such that D "always" decodes the packet, and none of the eavesdroppers decodes the packet.

$$\lim_{N\to\infty} P(\{S \to D\}) = c > 0$$

$$\lim_{N\to\infty} P(\{S \to E_0\} \bigcup \{S \to E_1\} \bigcup \ldots \{S \to E_{M-1}\}) = 0$$

Question: How many eavesdroppers can be tolerated?

Approach: (Ignore path-loss for now) Find the relay R* that has the best fading gains on S -> R* and R* -> D. More precisely,

$$R^* = \arg\max_{R_i : i=0,1,\ldots,N-1} \min(|h_{S,R_i}|^2, |h_{R_i,D}|^2)$$

Assuming an exponential tail on the magnitude squared of the fading distribution of any link,

$$\max \min(|h_{S,R^*}|^2, |h_{R^*,D}|^2) \sim \frac{1}{2}\log N$$

Then, can transmit with power 2/log N, and **tolerate sqrt(N) eavesdroppers**.

Basic Idea: If you cannot hear somebody "talk", then, by reciprocity, they will not be able to hear you. That allows you to be a jammer with little pain to the system when that person receives.

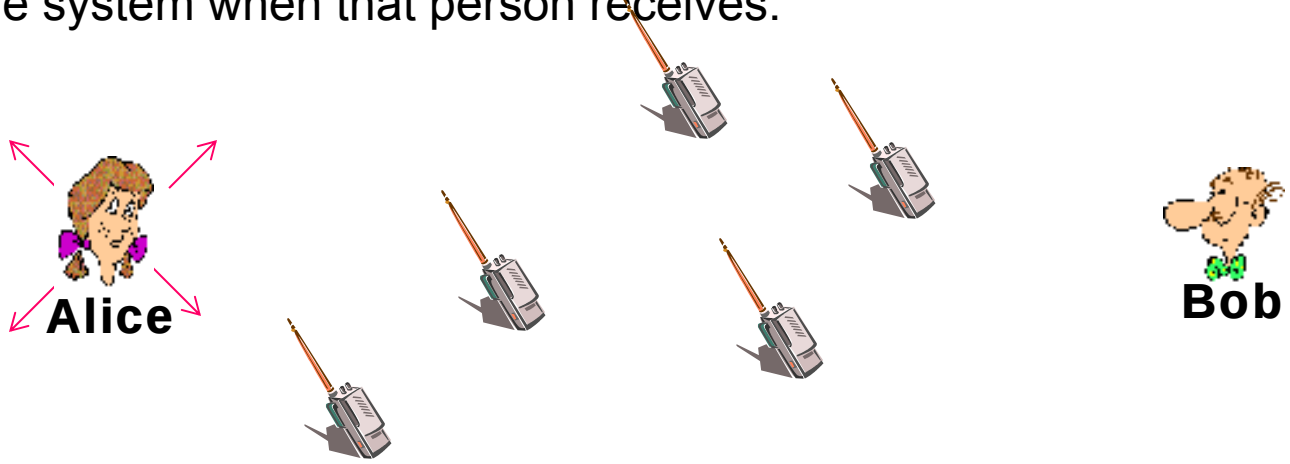

Scenario: Source "S" wants to communicate securely with the destination "D" in the presence of M eavesdroppers (M=3 above) with the help of N relays (N=5) above.

Step 1: Source broadcasts pilot. Relay i measures S -> R_i channel, i=0,1,…N-1.

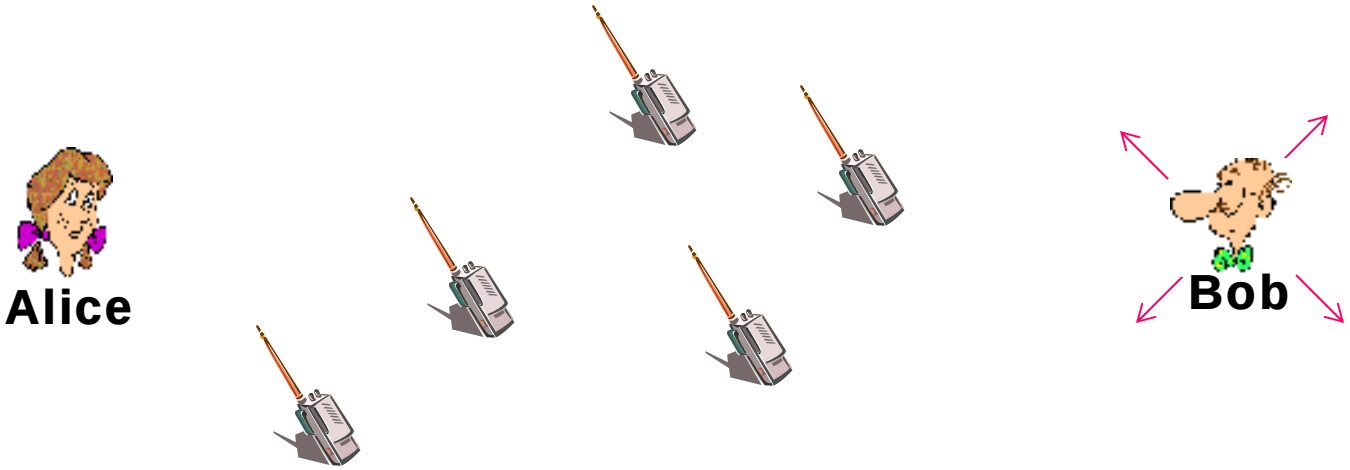

Scenario: Source "S" wants to communicate securely with the destination "D" in the presence of M eavesdroppers (M=3 above) with the help of N relays (N=5) above.

Step 1: Source broadcasts pilot. Relay i measures S -> R_i channel, i=0,1,…N-1.

Step 2: Destination broadcasts pilot. Relay i measures D -> R_i channel, i=0,1,…N-1.

Potential Solutions: (a) Exploiting fading (b) 2-way comms (c) Rcvr hardware (d) Jamming

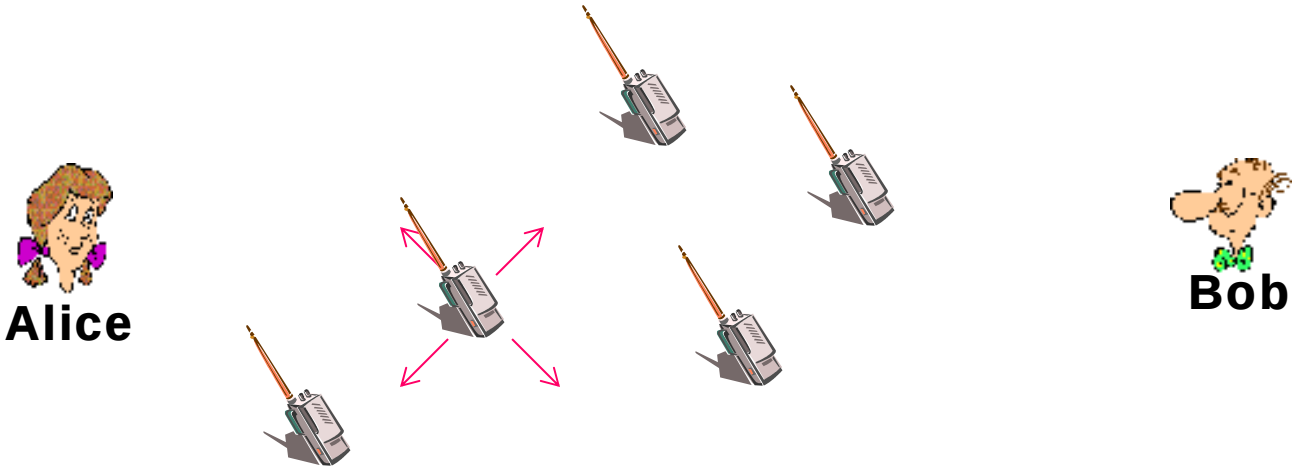

Scenario: Source "S" wants to communicate securely with the destination "D" in the presence of M eavesdroppers (M=3 above) with the help of N relays (N=5) above.

Step 1: Source broadcasts pilot. Relay i measures S -> R_i channel, i=0,1,…N-1.

Step 2: Destination broadcasts pilot. Relay i measures D -> R_i channel, i=0,1,…N-1.

Step 3: "Best" relay broadcasts pilot. Relay i measures R^* -> R_i channel, i=0,1,…N-1. (Note that this can be done in a distributed manner, with success prob c > 0).

Potential Solutions: (a) Exploiting fading (b) 2-way comms (c) Rcvr hardware (d) Jamming

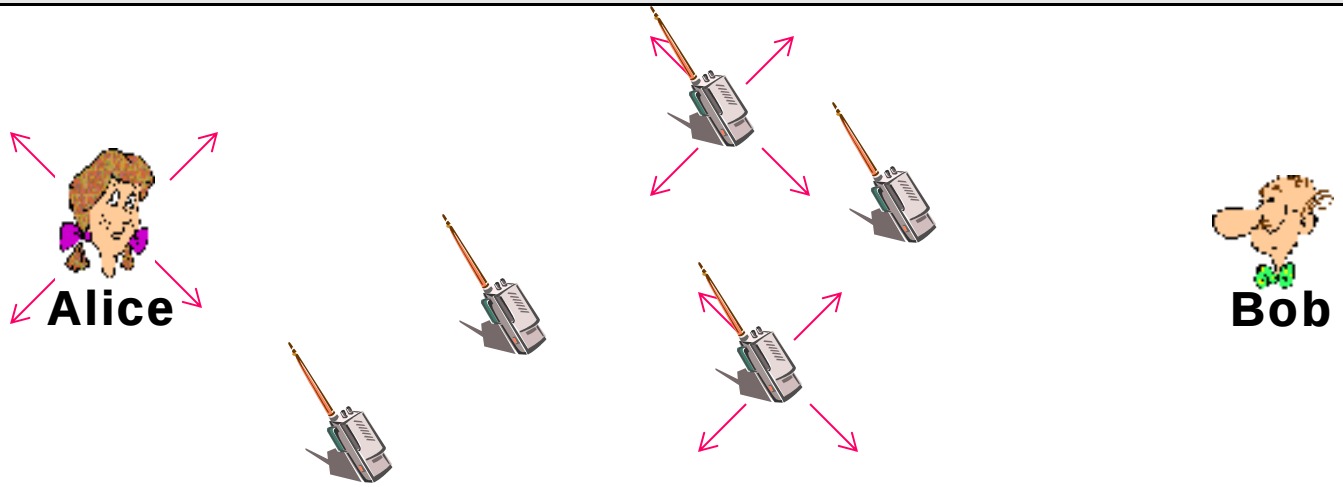

Scenario: Source "S" wants to communicate securely with the destination "D" in the presence of M eavesdroppers (M=3 above) with the help of N relays (N=5) above.

Step 1: Source broadcasts pilot. Relay i measures S -> R_i channel, i=0,1,…N-1.

Step 2: Destination broadcasts pilot. Relay i measures D -> R_i channel, i=0,1,…N-1.

Step 3: "Best" relay broadcasts pilot. Relay i measures R^* -> R_i channel, i=0,1,…N-1.

Step 4: Source transmits message. Relays with "bad" (to be defined later) R_i -> R^* channels generate random noise (to confuse eavesdroppers).

Potential Solutions: (a) Exploiting fading (b) 2-way comms (c) Rcvr hardware (d) Jamming

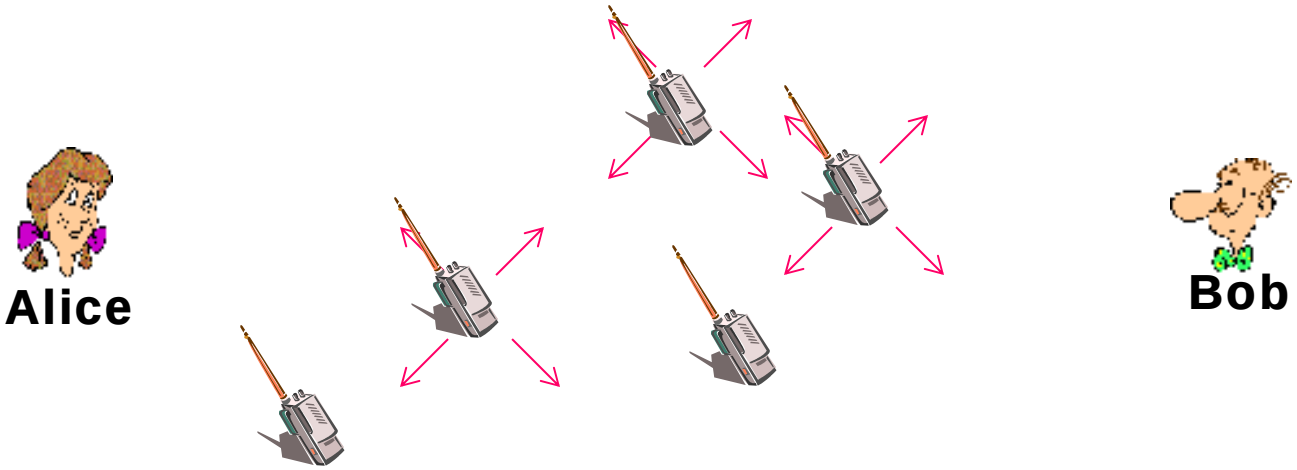

Scenario: Source "S" wants to communicate securely with the destination "D" in the presence of M eavesdroppers (M=3 above) with the help of N relays (N=5) above.

Step 1: Source broadcasts pilot. Relay i measures S -> R_i channel, i=0,1,…N-1.

Step 2: Destination broadcasts pilot. Relay i measures D -> R_i channel, i=0,1,…N-1.

Step 3: "Best" relay broadcasts pilot. Relay i measures R^* -> R_i channel, i=0,1,…N-1.

Step 4: Source transmits message. Relays with "bad" R_i -> R^* channels generate random noise (to confuse eavesdroppers).

Step 5: Relay R^* transmits message. Relays with "bad" (to be defined later) R_i -> D channels generate random noise (to confuse eavesdroppers).

Consider how the (possible) chatter of N intermediate nodes might help us to communication secretly in the presence of Eve. Can achieve a secrecy rate R with probability 1 (over the locations of nodes) for asymptotically large N if the number of eavesdroppers satisfies:

	Standard MU diversity: best relay/power control	**Proposed Scheme: intelligent "chatter"**
Lower bound on Eve distance from S	$o(\sqrt{N})$	$o(c_1^{\,c_2\sqrt{N\log N}}), \quad c_1 > 1$
Uniformly distributed eavesdroppers	$o(\log N)$	$o(N)$

Problem: Convergence (in N) is *slow*. Not practical (we tried!).

Solution (?): Use the (multi-hop) network.

Challenges

1. Exploiting when Alice -> Bob **channel is better** than Alice -> Eve
   Challenge: unknown Eve location
2. Exploiting common randomness of **channel reciprocity**
   Challenge: limited number of key bits
3. Exploiting "**public discussion**"
   Challenge: two-way communication and unknown Eve
4. Attacking Eve's **receiver hardware**
   Challenge: short range, assumptions on Eve's hardware
5. **Interference Cancellation**
   Challenge: near-far environment
6. **Relay Chattering**
   Challenge: great in theory, not so much in practice (density of nodes).

Outline

1. Computational and Information Theoretic security basics
2. Potential solutions
3. **Asymptotically-large networks**
   a. **Cooperative jamming**
   b. **Network coding**
4. Undetectable communications (LPD)
5. Current and Future Challenges

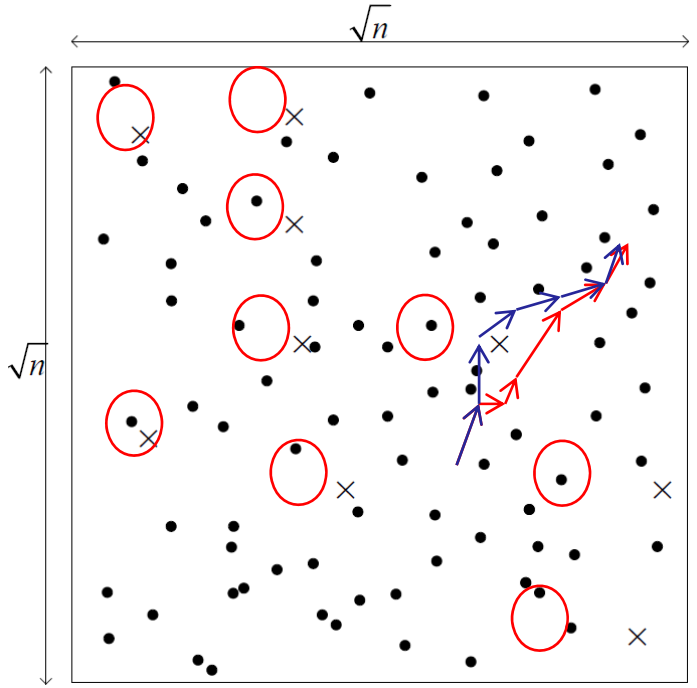

The Network Scale:

So far we have considered:

Alice

Eve

Bob

But now we want to consider:

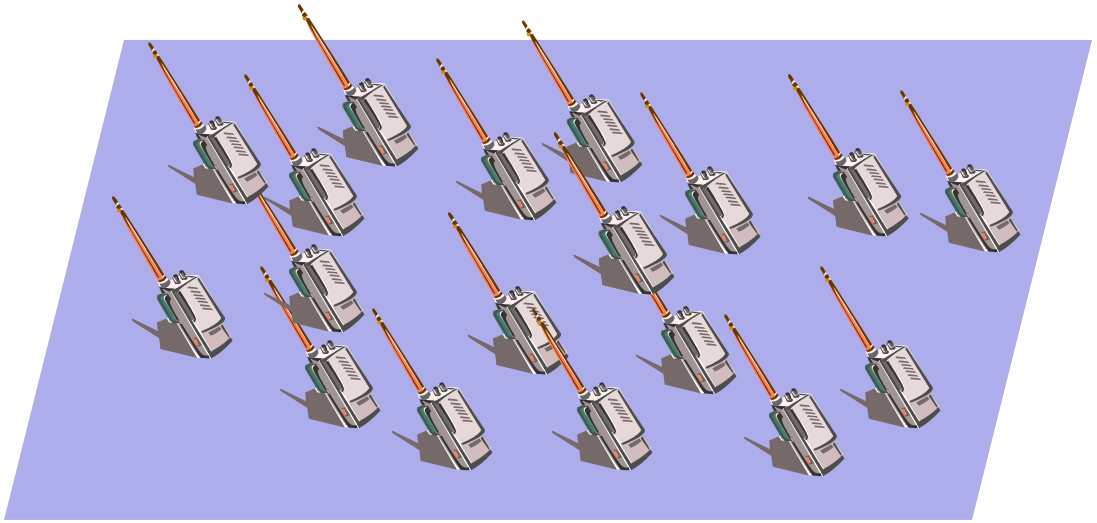

Questions:

1. How much secret information can be shared by a network of wireless nodes in the presence of eavesdropper nodes? [Gupta/Kumar et al]
2. … and how many eavesdroppers can the network tolerate?

Problem: Secrecy Scaling

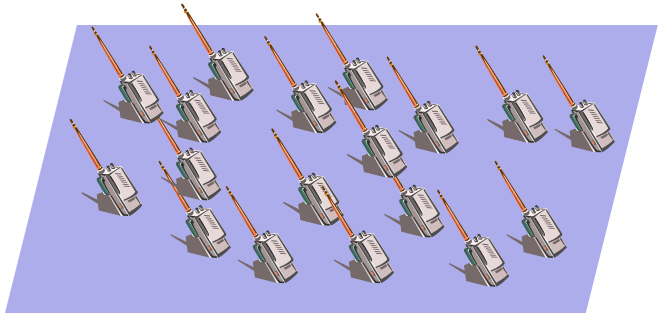

- *n* good guys (matched into pairs)
- *m* bad guys.

Secure throughput per pair?
For what *m*?

Gupta-Kumar: *n* nodes each can share $\Theta(1/\sqrt{n \log n})$ bits per second.

Want to achieve this throughput *securely*, in the presence of *m* eavesdroppers of **unknown** location.

Focus (as always): Avoiding the known eavesdropper location assumption.

1-D Networks (worst-case topology)

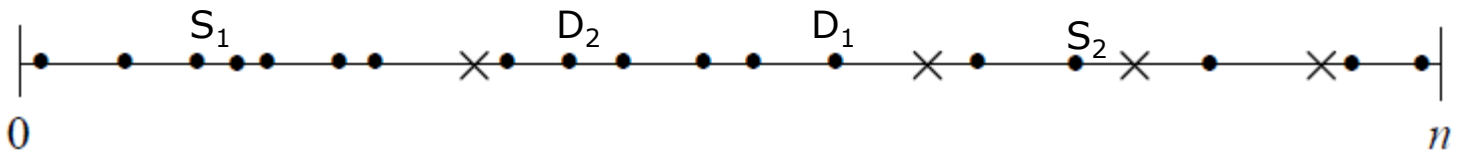

Network (random extended network):

- Nodes are placed in the interval [0, n].
- Legitimate nodes randomly placed: Poisson with unit density (*n* good guys on average)
- Eavesdroppers Poisson with $\lambda_e(n)$ ($m(n)$ = $\lambda_e(n)\,n$ bad guys on average.)
- n nodes matched into source-destination pairs uniformly at random.

n randomly located nodes -> how much secret information in the presence of $m(n)$ eavesdroppers ?

1-D Networks (worst-case topology)

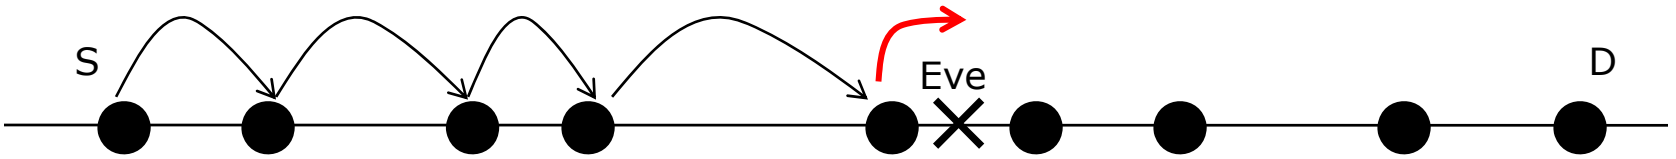

- A **single** eavesdropper of **known** location -> 1-D **disconnected (zero secret bits)**!
- Why? For any node to Eve's left, Eve is closer (has larger SINR) than the good nodes located to Eve's right.
- What to do? Answer: Nodes help each other to achieve secrecy -> **Cooperation.**

Cooperative Jamming:

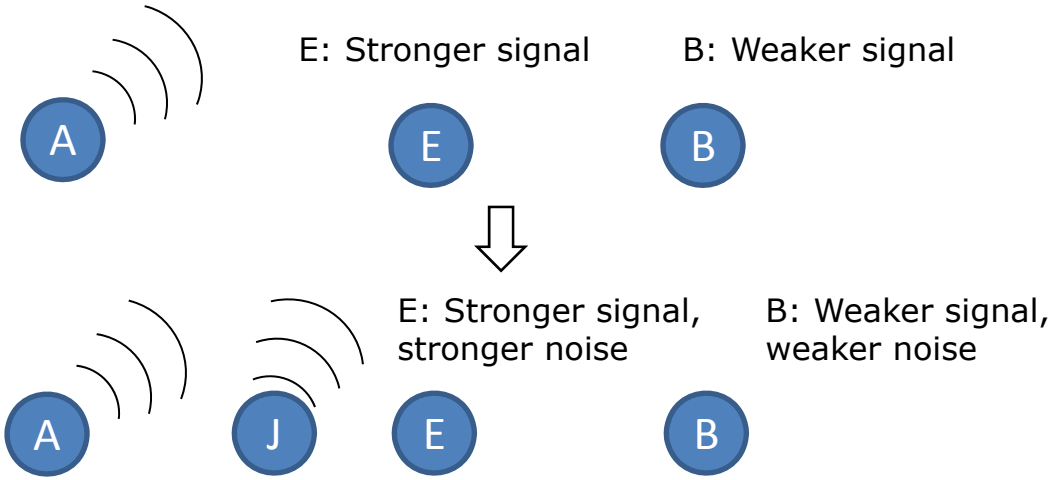

Routing Algorithm:

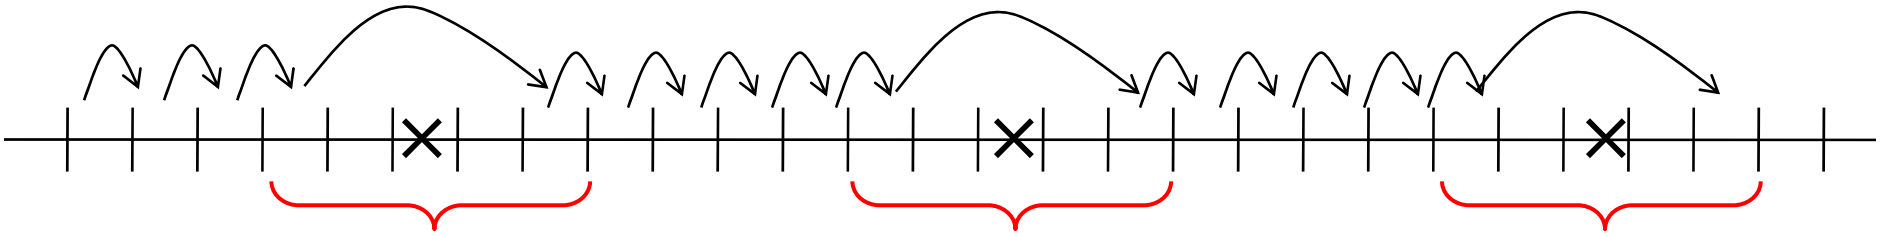

- So, the idea to connect each S-D pair through a sequence of many single-cell hops + one multi-hop **jump** until reaching D.
- Jamming works if you know where the eavesdroppers are.

But what if you **don't know** where the eavesdroppers are?

Unknown Eavesdropper Locations:

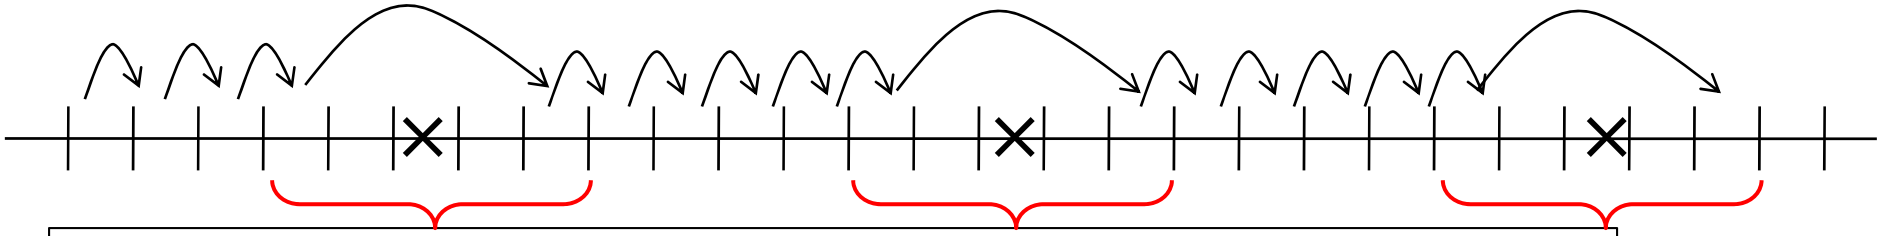

- I know how to protect the message if eavesdroppers are spaced far apart.
- But this time eavesdroppers can be anywhere.

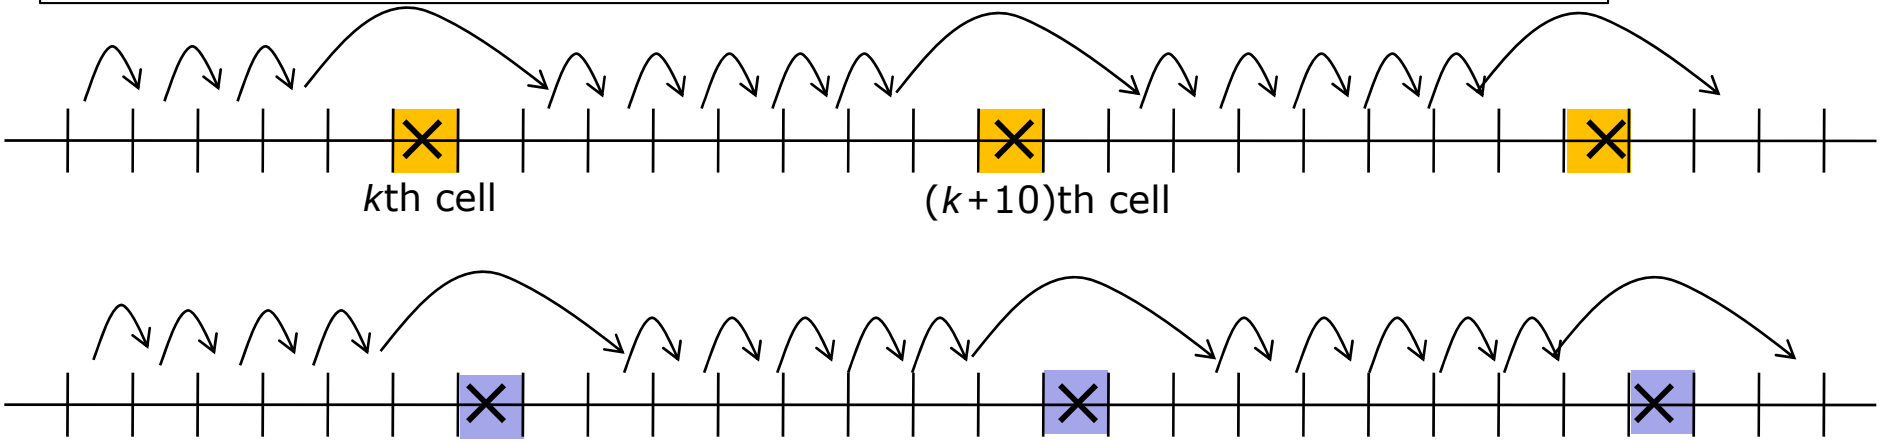

Solution comes with secret sharing at the source

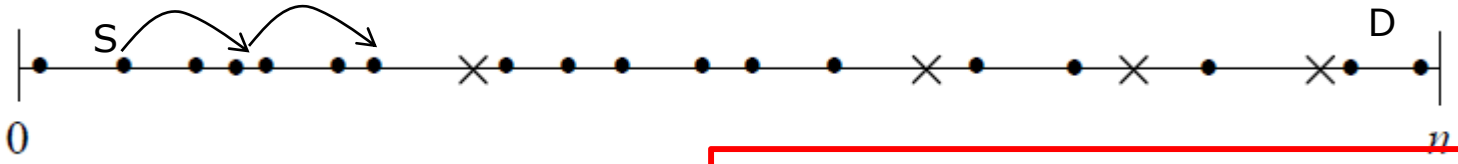

- x: the b-bit secret message.
- S generates t-1 b-bit **packets** w_1, w_2, …, w_{t-1} randomly, *sets* w_t *to be such that*

$$x = w_1 \oplus w_2 \oplus w_3 \oplus \cdots \oplus w_t$$

	x	= 1 0 0 1 1
random	w_1	= 0 1 0 1 1
random	w_2	= 1 0 1 1 1
random	w_3	= 0 0 0 1 0
	w_4	= 0 1 1 0 1

- Anyone who has all t packets has the message.
- Anyone who misses **at least one packet** has no information about the message.

- For one message, S sends t packets.
- The packets are sent in **separate** transmissions.
- **Idea**: Ensure an eavesdropper anywhere in the network **misses** at least one packet.

Asymptotically-large networks: (a) Cooperative jamming (b) Network coding

Unknown Eavesdropper Locations:

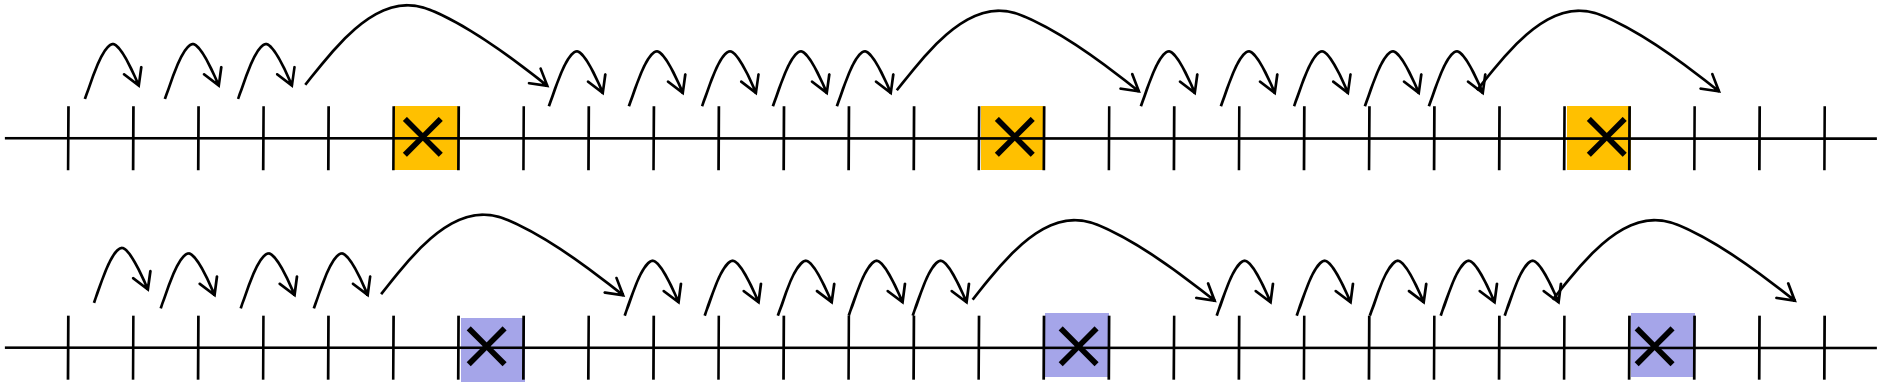

Divide the network into regions: Coloring the network!

$$\Theta(1/n)$$

Secrecy Analysis:

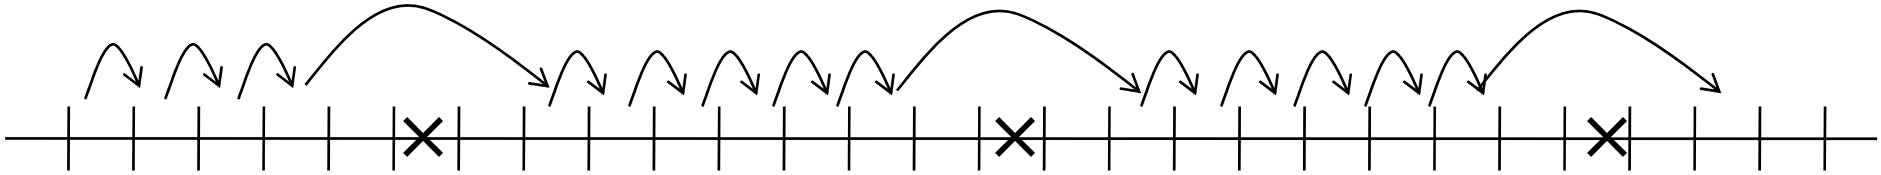

The only potentially **unsecure** places: *start* or the *end* of a route.

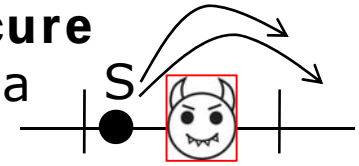

Near-eavesdropper ⟹ **n/logn** eavesdroppers

Throughput Analysis:

Standard Scheduling applied: cells take turns in transmissions: $\Theta(1/n)$ throughput per node pair (standard).

Two-Dimensional Networks (Review)

(Insecure) Throughput:

- $\Theta\left(1/\sqrt{n \log n}\right)$ bits per second (per-node throughput) [Gupta-Kumar, 2001]
- $\Theta(1/\sqrt{n})$ bits per second (per-node throughput [Franceschetti et al, 2007]
- Multi-hop route connecting source-destination pairs.

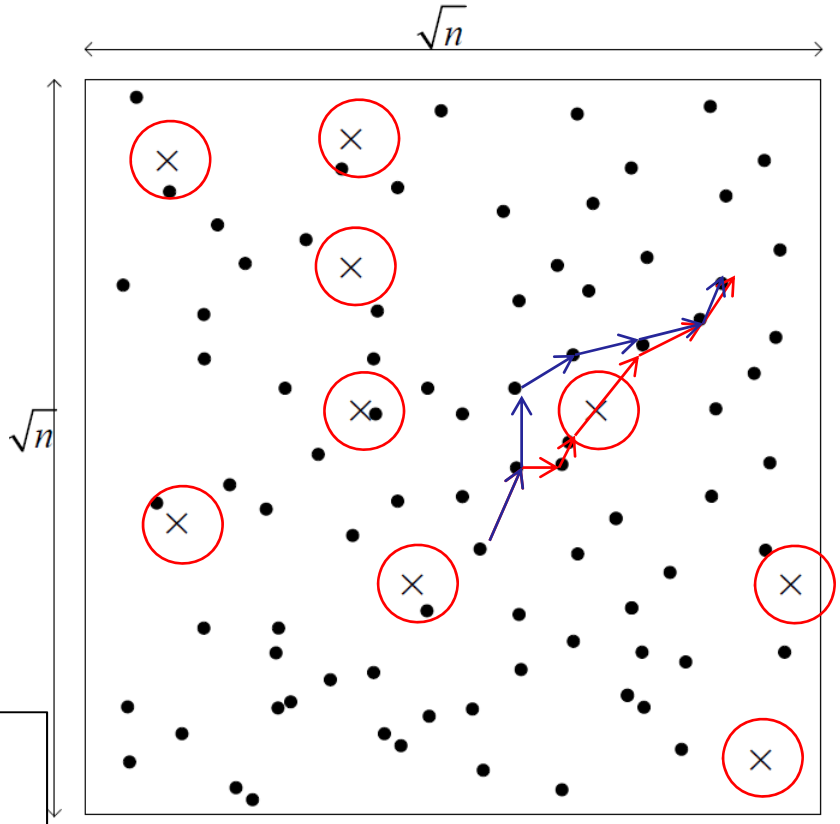

If eavesdropper locations are known:

- Can route around the "holes" as long as ***m* = *o*(*n*/log*n*)** [Koyluoglu et al, 2011]
- Also: Securing each hop individually is sufficient to secure the end-to-end route [Koyluoglu et al, 2011]

Two-Dimensional Networks

Unknown eavesdropper location. What to do?

First answer: Try **Cooperative Jamming**.

At each hop, some nodes transmit artificial noise to protect the message from eavesdroppers around.

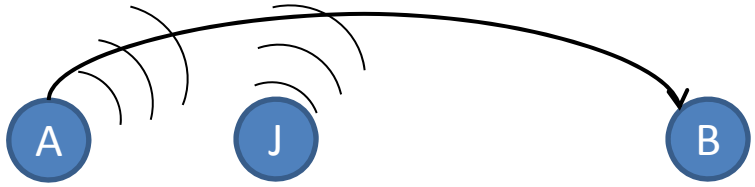

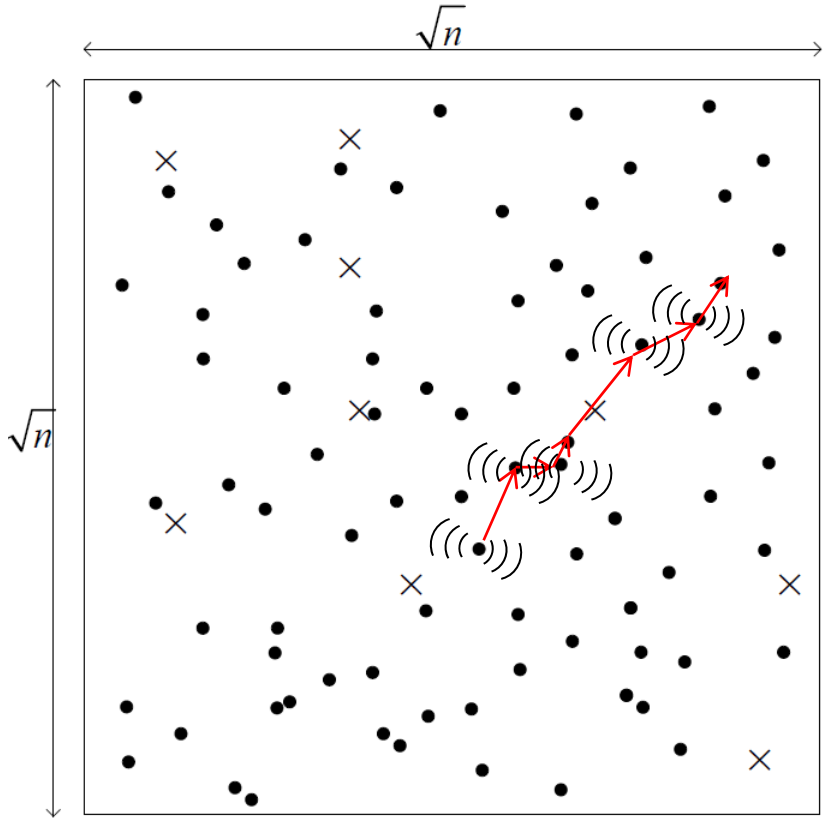

Can tolerate $m(n) = \log n$ eavesdroppers (only). Is this the cost of unknown eavesdropper positions?

[Vasudevan et al, 2011]

Two-Dimensional Networks (Secret Sharing)

$$x = w_1 \oplus w_2 \oplus w_3 \oplus \cdots \oplus w_t$$

	x	= 1 0 0 1 1
random	w_1	= 0 1 0 1 1
random	w_2	= 1 0 1 1 1
random	w_3	= 0 0 0 1 0
	w_4	= 0 1 1 0 1

- An eavesdropper cannot be close to many paths at once…
- Except when close to the source or the destination.

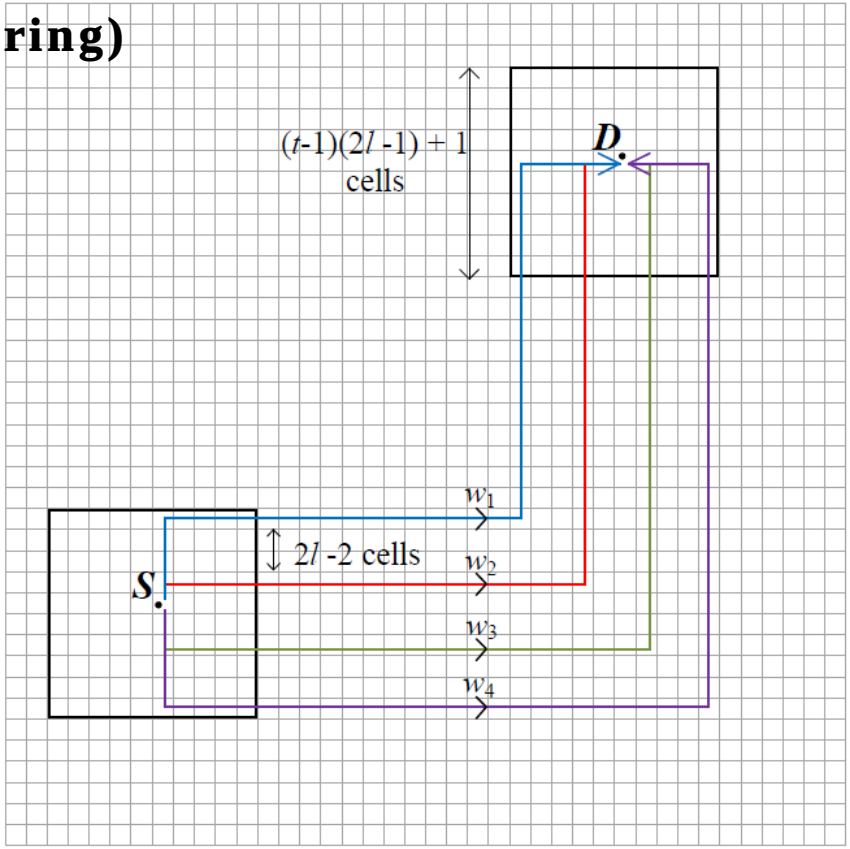

[Capar, Goeckel, Towsley, 2012]

Can tolerate $n/\log n$ eavesdroppers.

What about those near eavedroppers?

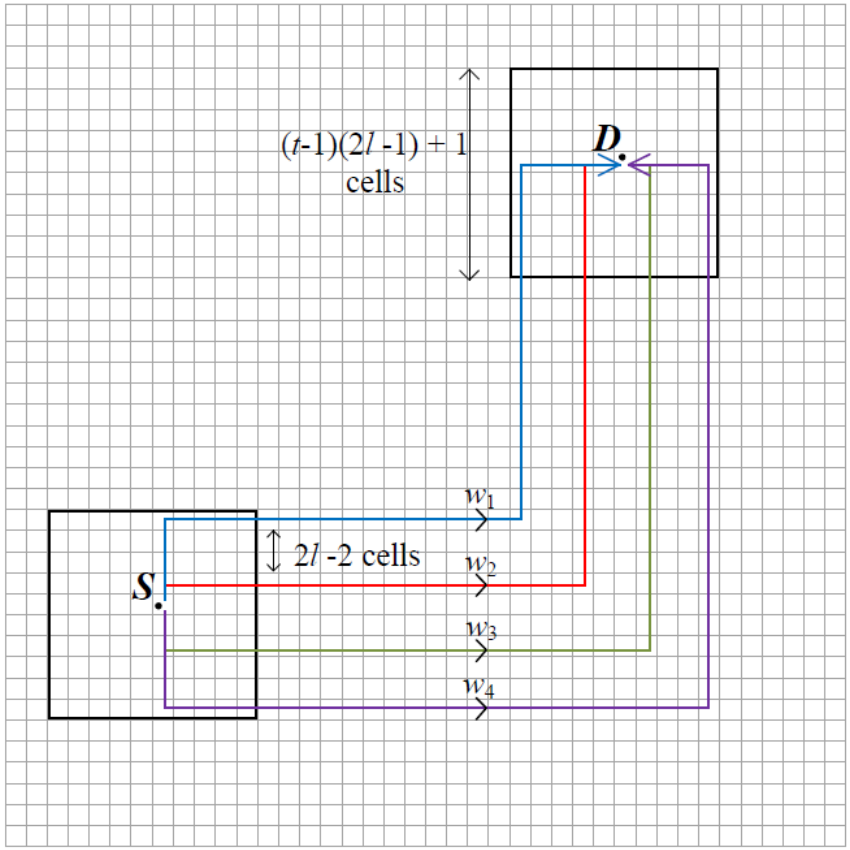

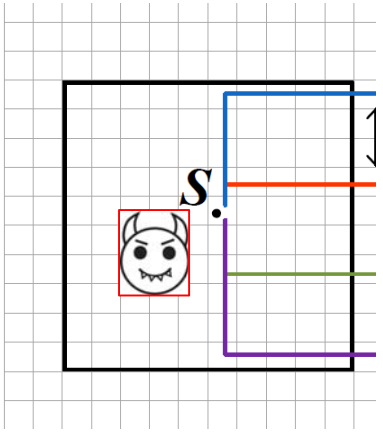

Remember, the problem here was: **Near eavesdropper of unknown location**.

Secure Network Coding

- A formal way to study a **wiretap network**.
- Start with a given **graph** representing the network.
- Some of the edges are **tapped**.

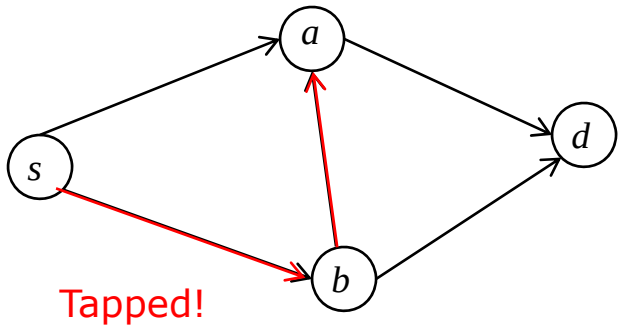

- Gives necessary and sufficient condition to go securely from *s* to *d*,
- and (if possible) tells you how to do it.
- Nice formal way to check secrecy capability, but wireless secrecy: **we don't have a graph**.

[Cai and Yeung, 2002], [Jain 2004]

Secrecy Graph: Wireless Network -> Wiretap Graph

Takes your wireless network, puts out a graph.

Start with given locations of good guys $\{x_i\}$, and bad guys $\{y_i\}$. Then, a simple rule.

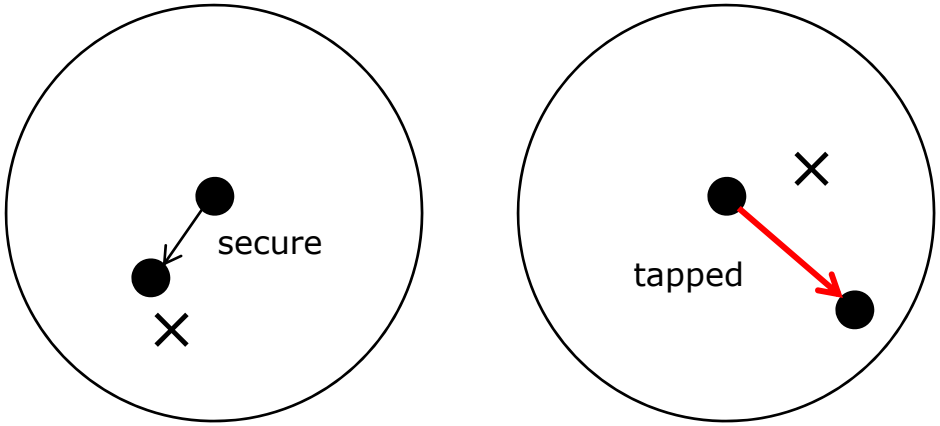

Baseline graph

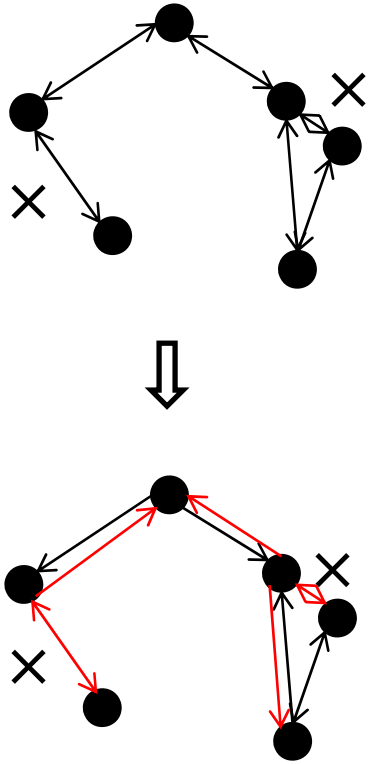

- Draw an edge from x_i to x_j if within radius.
- Edge (x_i, x_j) is tapped if $d(x_i , y_k) \leq d(x_i , x_j)$

[Haenggi, 2008]

Toy Example 1 (Basic Two-way Scheme).

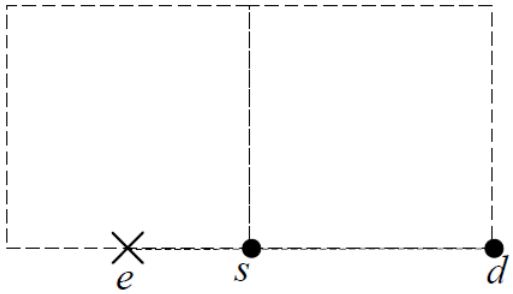

Wireless Network

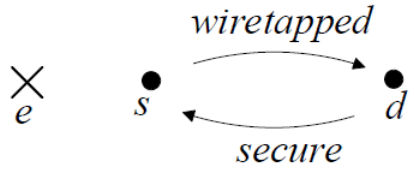

- Two nodes + one eavesdropper
- *e* catches whatever *s* says.
- *An incoming secure edge is sufficient for secrecy.*

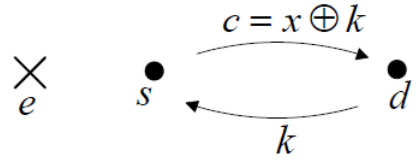

Secrecy protocol for connecting *s* to *d*

1) *d* generates a random message *k* and sends it to *s*.

2) *s* replies with $c = x \oplus k$ (*k* is used as a one-time pad.)

3) *d* extracts *x* from *c*, *k*.

e misses k, cannot decode x.

- An incoming connection can be very useful.
- This simple trick has an important implication for wireless secrecy.
- Two-way helps address the **near eavesdropper** problem

Toy Example 2:

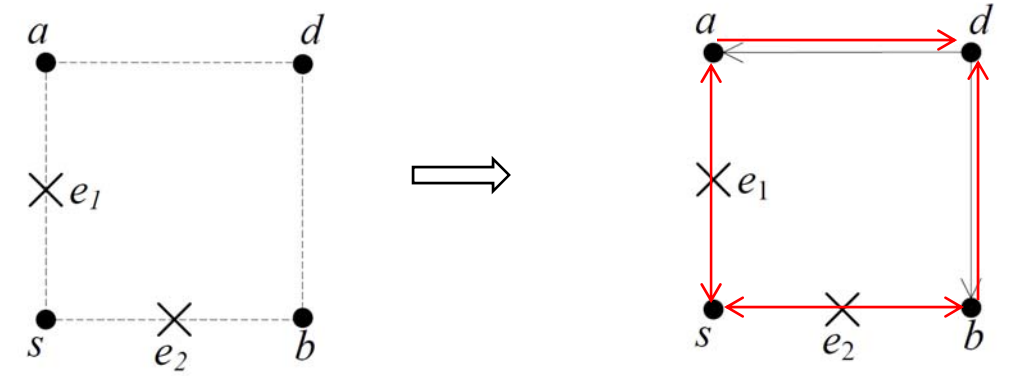

Wireless Network

Disconnected!

- Four nodes on the corners of a square.
- Two eavesdroppers in the middle of two edges.

s **disconnected** from **both** neighbors in both directions!

Still hope?

Toy Example 2 (continued):

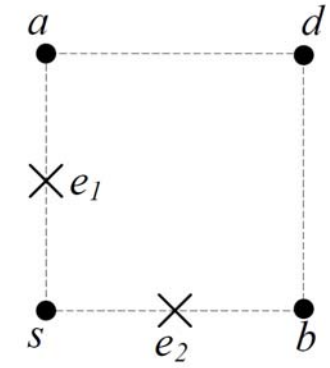

Wireless Network

s **disconnected** from **both** neighbors in both directions!

Still hope?

*e*1 catches whatever *a* or *s* says.

*e*2 catches whatever *b* or *s* says.

- *s* blocked from both sides.
- Can still achieve secrecy if these blockages are due to **separate non-collaborating** eavesdroppers.

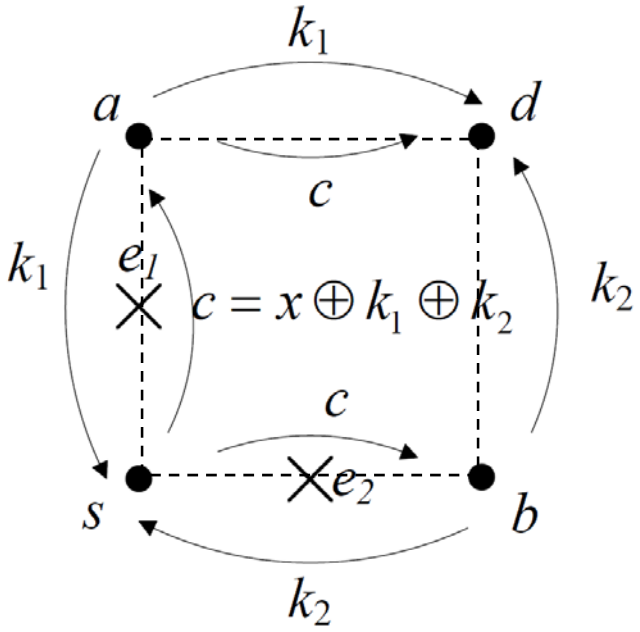

Secrecy protocol for connecting *s* to *d*

1) *a* generates, sends *k*1, *b* generates, sends *k*2.

2) s replies with $c = x \oplus k_1 \oplus k_2$

3) *d* gets *c*, *k*1, *k*2 -> extracts *x*

*e*1 misses *k*2, *e*2 misses *k*1

Back to Secrecy Capacity for the Main Result.

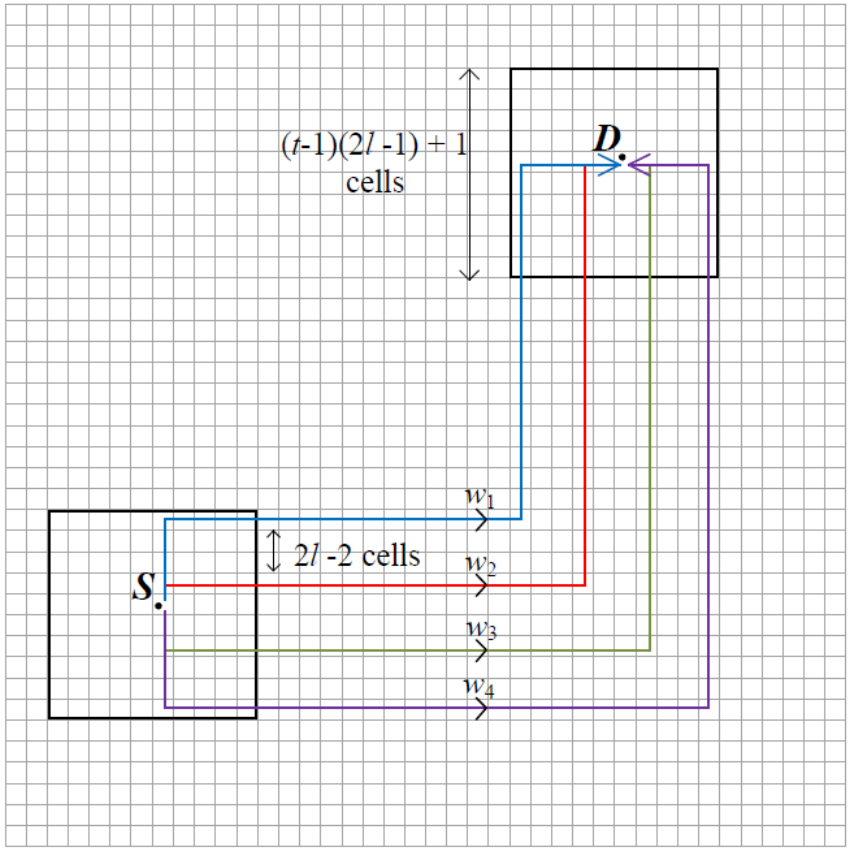

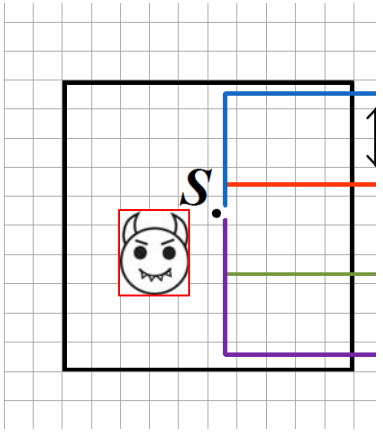

Remember, the problem here was:
Near eavesdropper of unknown location.

Two-Dimensional Networks

Draining Phase: How we initiate at the source

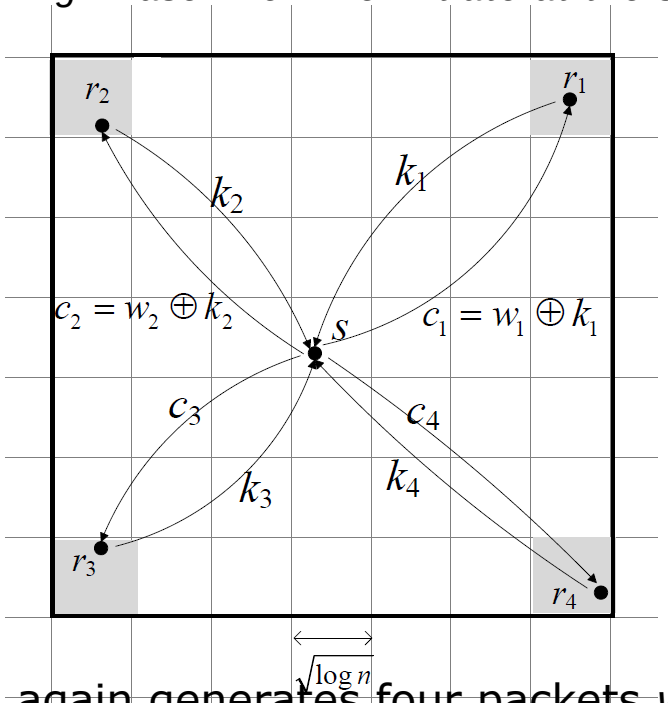

- S again generates four packets *w*1, *w*2, *w*3, *w*4.
- But this time, does the two-way scheme with four relays to deliver these packets.

Routing Algorithm: Draining, routing, delivery

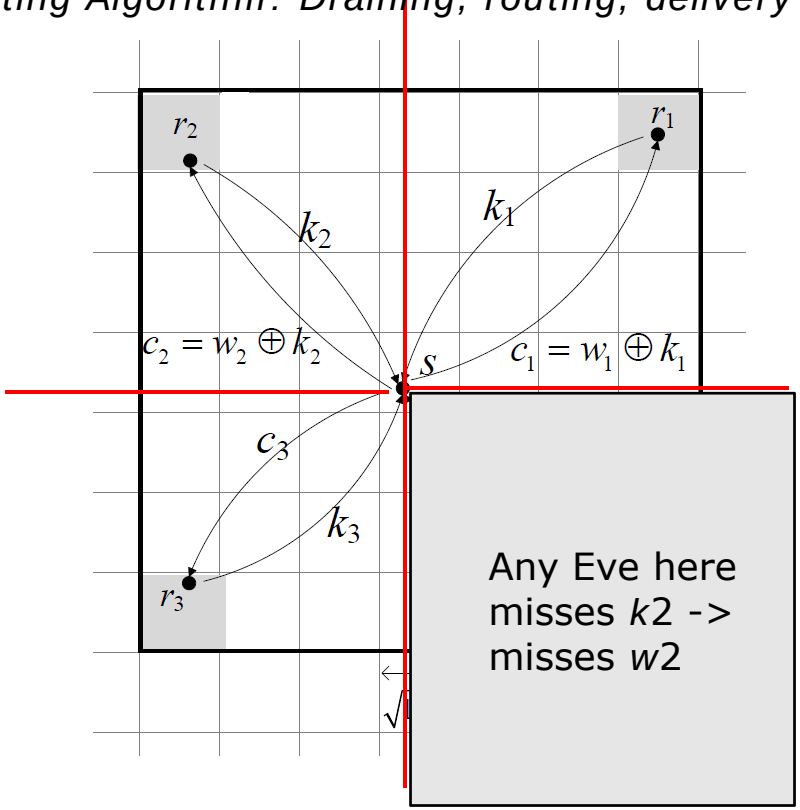

- Come from four directions.
- No Eve can be in between for all four *r*-*s* pairs.

Two-Dimensional Networks

Result: Network can tolerate *any* number of eavesdroppers of *arbitrary location* at the Gupta-Kumar per-pair throughput.

Note: Importantly, the same technique can be used in practical networks that are:
1. sufficiently dense
2. add infrastructure

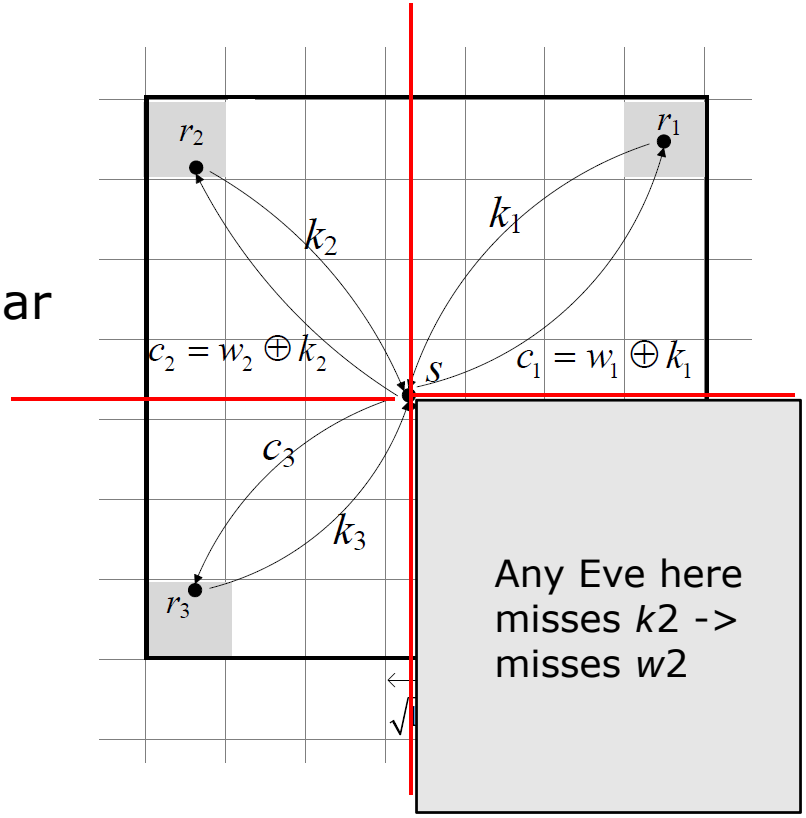

- Come from four directions.

- No Eve can be in between for all four *r*-*s* pairs.

Outline

1. Computational and Information Theoretic security basics
2. Potential solutions
3. Asymptotically-large networks
4. **Undetectable communications (LPD)**
   a. **Emerging approaches for wireless channels**
   b. **Experiments**
5. Current and Future Challenges

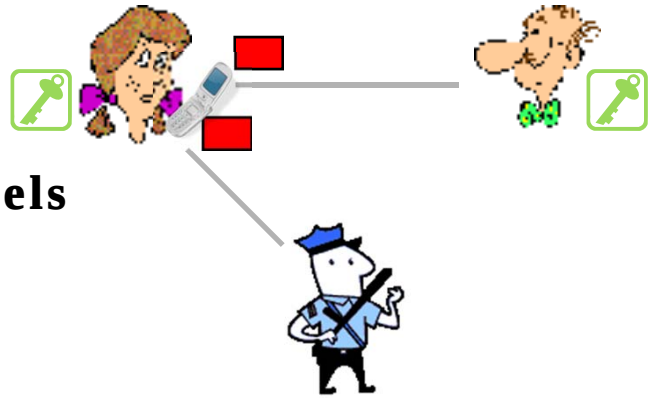

LPD Communications – of our interest

From: *The Guardian*

- Problem: conceal **presence** of the message: "metadata" as opposed to concealing message *content* (encryption)
- Why? Lots of applications…
  - Encrypted data looks suspicious
  - "Camouflage" military operations
  - etc…
- LPD="low probability of detection"
  - Limit adversary's detection capability to tolerable level
- **Fundamental limits of LPD communication**

Scenario

- Alice uses radio to covertly communicate with Bob
  - They share a secret key

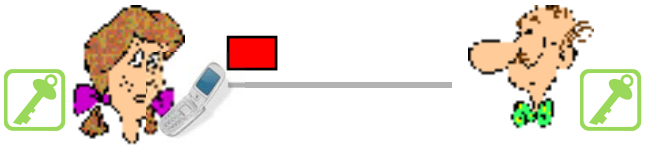

Scenario

- Alice uses radio to covertly communicate with Bob
  - They share a secret key
- Willie attempts to detect if Alice is talking to Bob
  - Willie is passive, doesn't actively jam Alice's channel

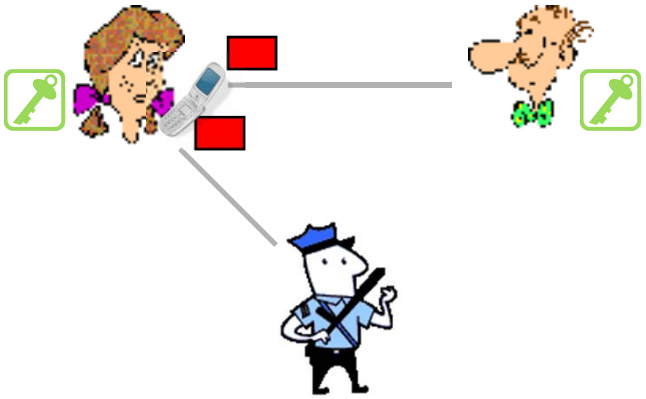

Scenario

- Alice uses radio to covertly communicate with Bob
  - They share a secret key
- Willie attempts to detect if Alice is talking to Bob
  - Willie is passive, doesn't actively jam Alice's channel

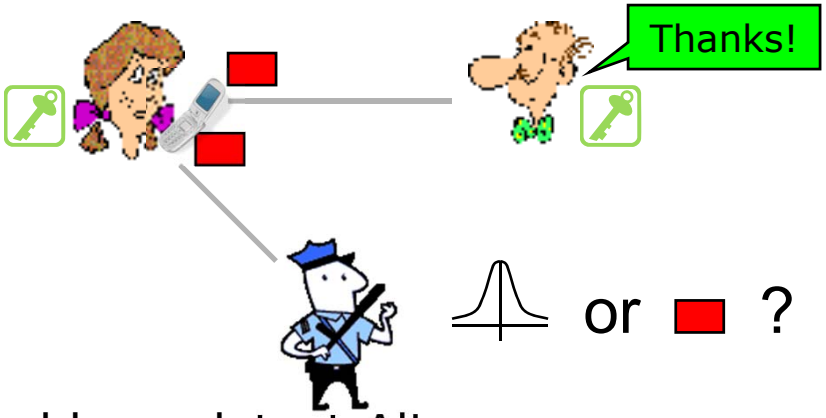

- Willie's problem: detect Alice
- Alice's problem: limit Willie's detection schemes
- Bob's problem: decode Alice's message

Results from Steganography:

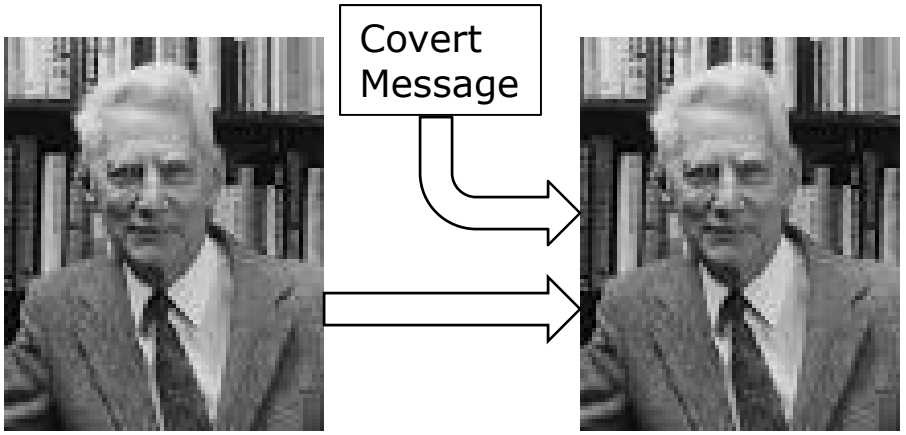

Problem: Modify characters in a cover text (message, picture, etc.) to convey secret message without detection.

Results:

1. $\sqrt{n}$ symbols can be modified in a cover text of length n symbols without detection.
2. $\sqrt{n}\log n$ bits of information can be encoded in those n symbols without detection

But this is on a finite alphabet channel. What about a physical (e.g. wireless) channel?

Spread Spectrum

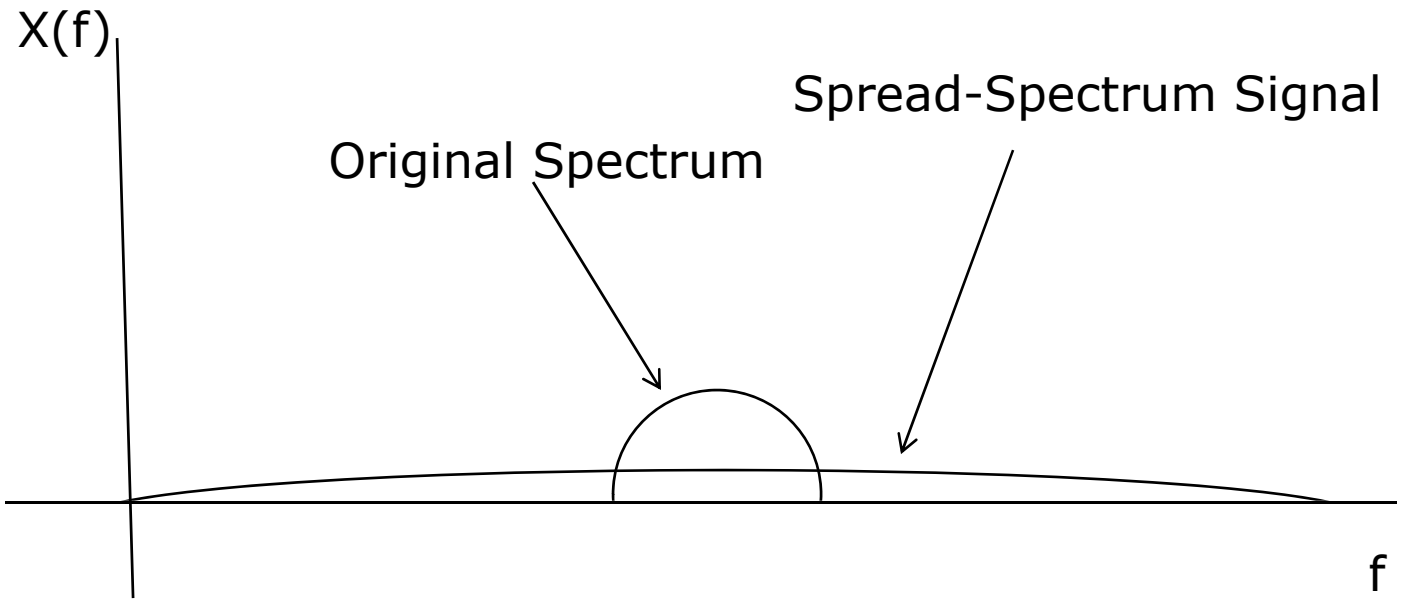

"Hide signal in the noise" at a spectrum loss of 1/N (the bandwidth expansion or processing gain). But what is the fundamental tradeoff?

Channel Model

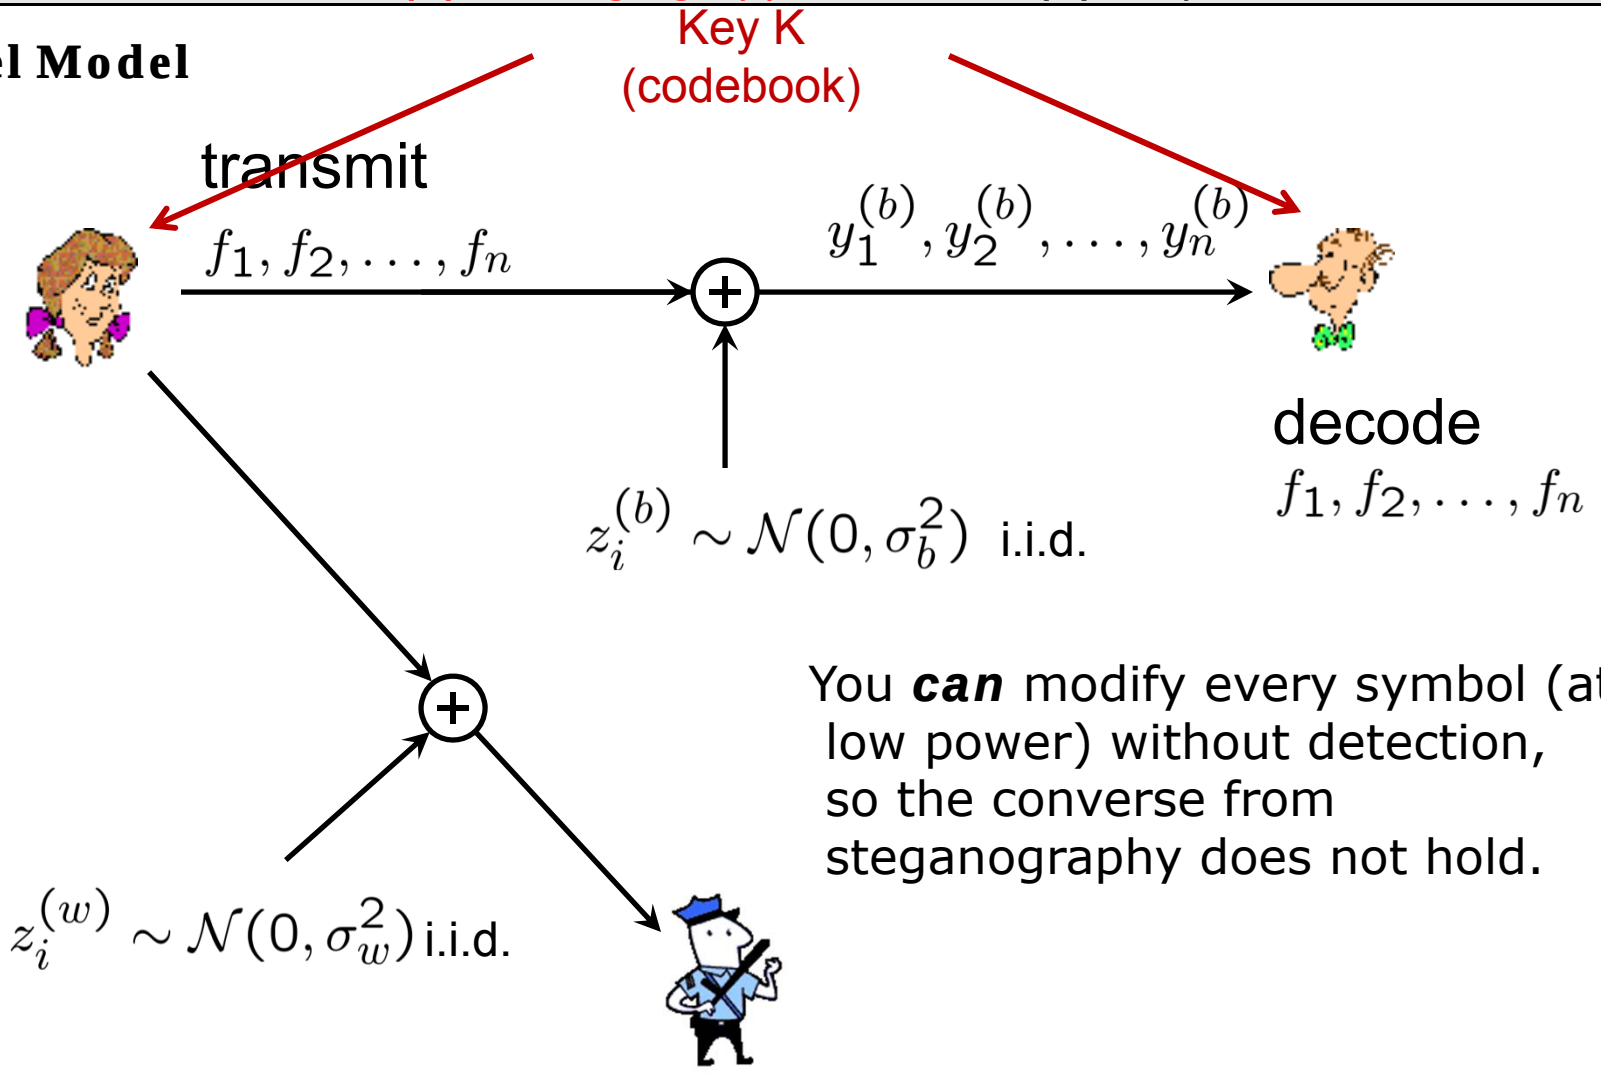

You ***can*** modify every symbol (at low power) without detection, so the converse from steganography does not hold.

The Detection Side (Willie): ROC Curves

Characterization of Willie's detector:

1. Low threshold: always say "Yes".
2. High threshold: always say "No".
3. Tradeoff curve for an "average" detector.

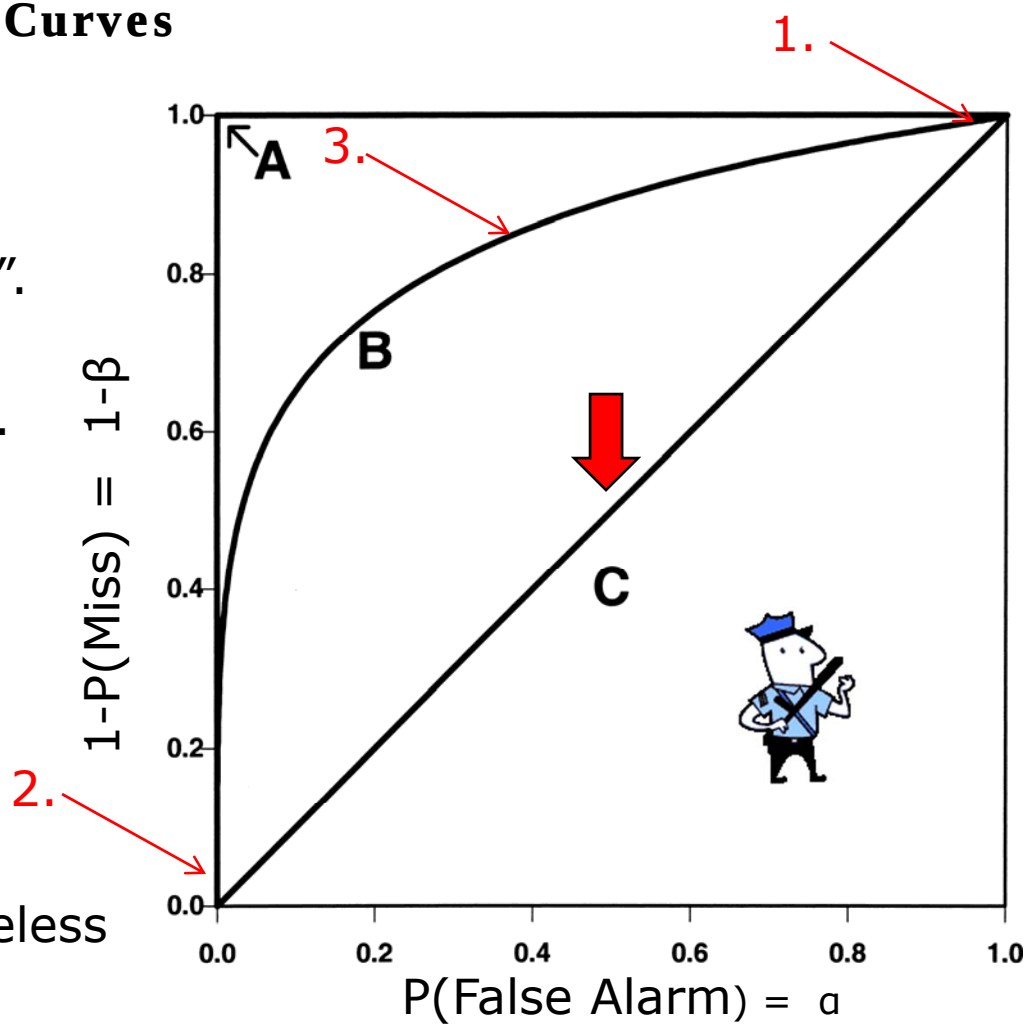

[from "circ.ahajournals.org"]

Goal: Drive Willie to Curve "C", useless detector.

Main Result: The Square Root Law

Consider a party Alice trying to communicate to Bob in the presence of a Warden Willie, with all channels AWGN:

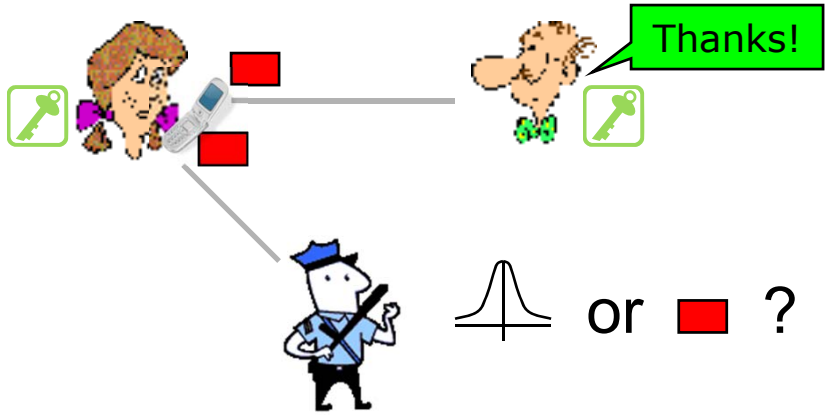

1. Alice can send $O(\sqrt{n})$ bits reliably to Bob with probability of detection at Willie < ε for any ε > 0
2. Conversely, if Alice tries to send $\omega(\sqrt{n})$ bits to Bob, one of the following occurs:
   - Bob's decoding error is bounded away from zero, or
   - Alice's transmission is detected with probability 1.

[Bash, Goeckel, Towsley, 2013]

Structure of the proof

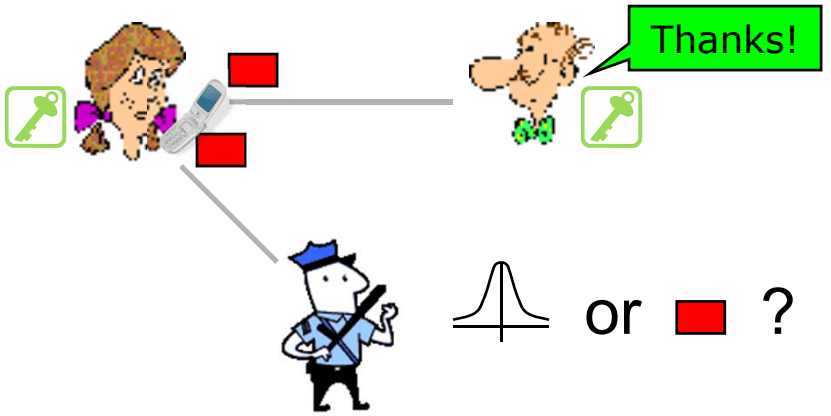

Achievability

+ Willie doesn't detect transmission, despite a detector, *and*
+ Bob decodes the message reliably with a (possibly) sub-optimal scheme

Converse

+ Bob cannot decode the message with an optimal detector, *or*
+ Willie can detect the transmission with a sub-optimal scheme

Achievability: Construction

- Random codebook with average symbol power P_f

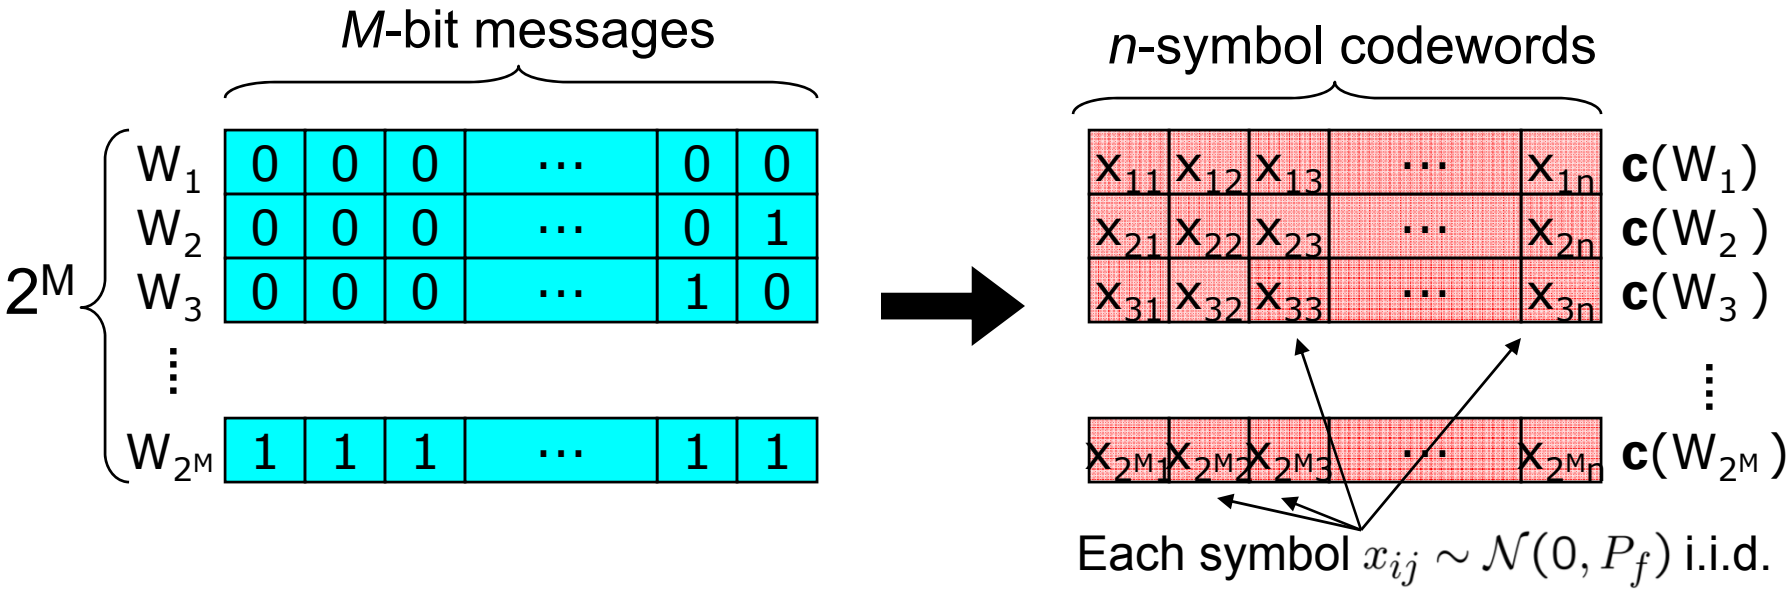

- Codebook revealed to Bob, but not to Willie
- Willie knows how codebook is constructed, as well as n and P_f
  - System obeys Kerckhoffs's Law: all security is in the key used to construct codebook

Achievability: Analysis of Willie's Detector

- Willie collects *n* observations
  - $\mathbf{P}_0^n$ -- distribution when Alice quiet, $\mathbf{P}_0 = \mathcal{N}(0, \sigma_w^2)$
  - $\mathbf{P}_1^n$ -- distribution when Alice transmitting, $\mathbf{P}_1 = \mathcal{N}(0, P_f + \sigma_w^2)$

- For any hypothesis test:

$$\alpha + \beta \geq 1 - TV(\mathbf{P}_0^n, \mathbf{P}_1^n)$$

Total Variation: 1-norm distance on the distributions

- Bounding Willie's detection:

$$TV(\mathbf{P}_0^n, \mathbf{P}_1^n) \leq \sqrt{\frac{n}{2} D(\mathbf{P}_0 \| \mathbf{P}_1)} \leq \frac{P_f}{2\sigma_w^2}\sqrt{\frac{n}{2}} \Rightarrow \boxed{P_f = \frac{c}{\sqrt{n}}}$$

Relative entropy

Back of the Envelope

Hence, the number of bits conveyed in *n* symbols is:

$$n\log(1+\frac{c}{\sqrt{n}}) \rightarrow O(\sqrt{n})$$

(That is not quite rigorous, as Shannon capacity is for a fixed *R* as *n* goes to infinity.)

But there is a simple workaround to finish the proof.

Converse

- When Alice tries to transmit $\omega(\sqrt{n})$ bits in *n* channel uses, using arbitrary codebook, either
  - Detected by Willie with arbitrarily low error probability
  - Bob's decoding error probability bounded away from zero

- Two step proof:
  1. Willie detects arbitrary codewords with average symbol power $\omega(1/\sqrt{n})$ using a simple power detector
  2. Bob cannot decode codewords that carry $\omega(\sqrt{n})$ bits with average symbol power $\mathcal{O}(1/\sqrt{n})$ with arbitrary low error

Done at Raytheon/BBN Technologies with collaborator Saikat Guha

Experiment: given n possible slots for PPM symbols, Alice and Bob secretly agree on a random subset of expected size $c_\tau \sqrt{n}$ to use for message transmission

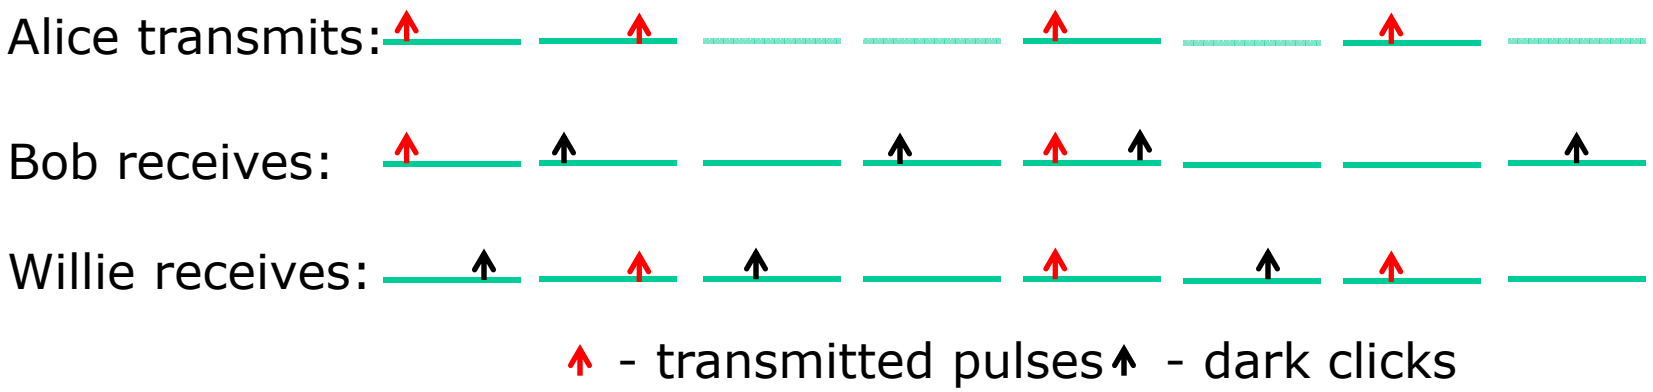

- Bob ignores the "empty" symbols, but Willie cannot since he doesn't know where they are

Experimental Set-Up

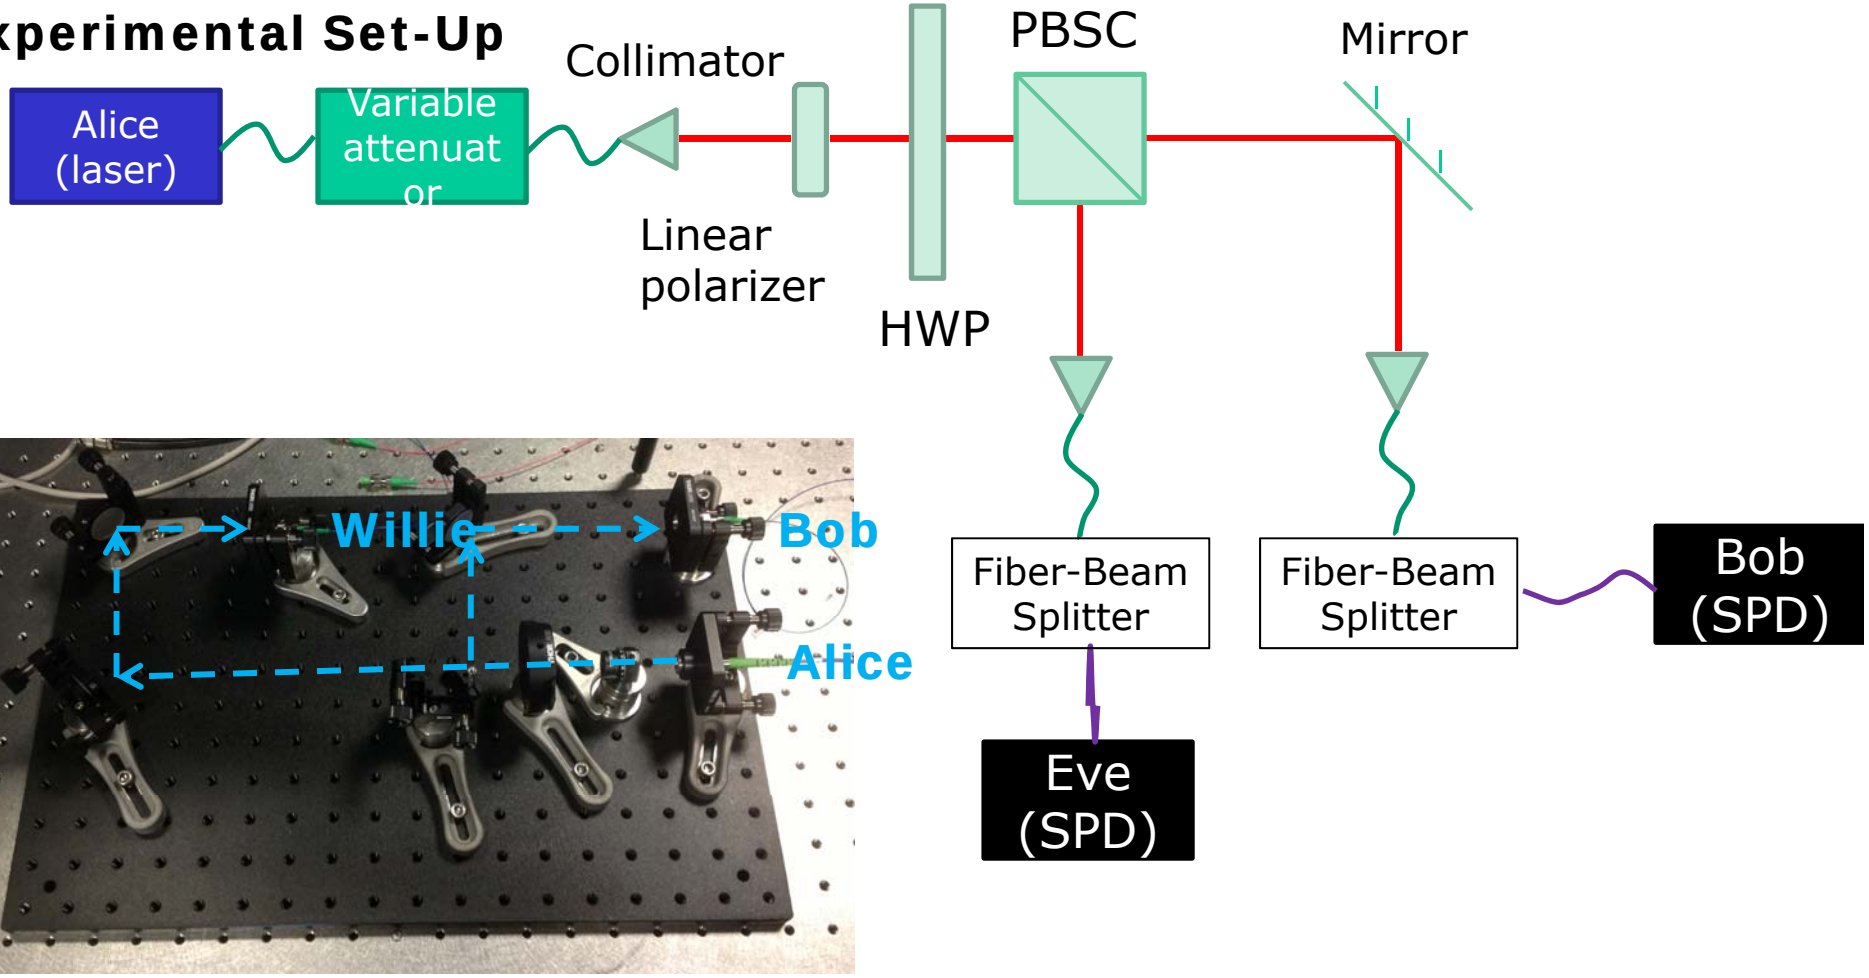

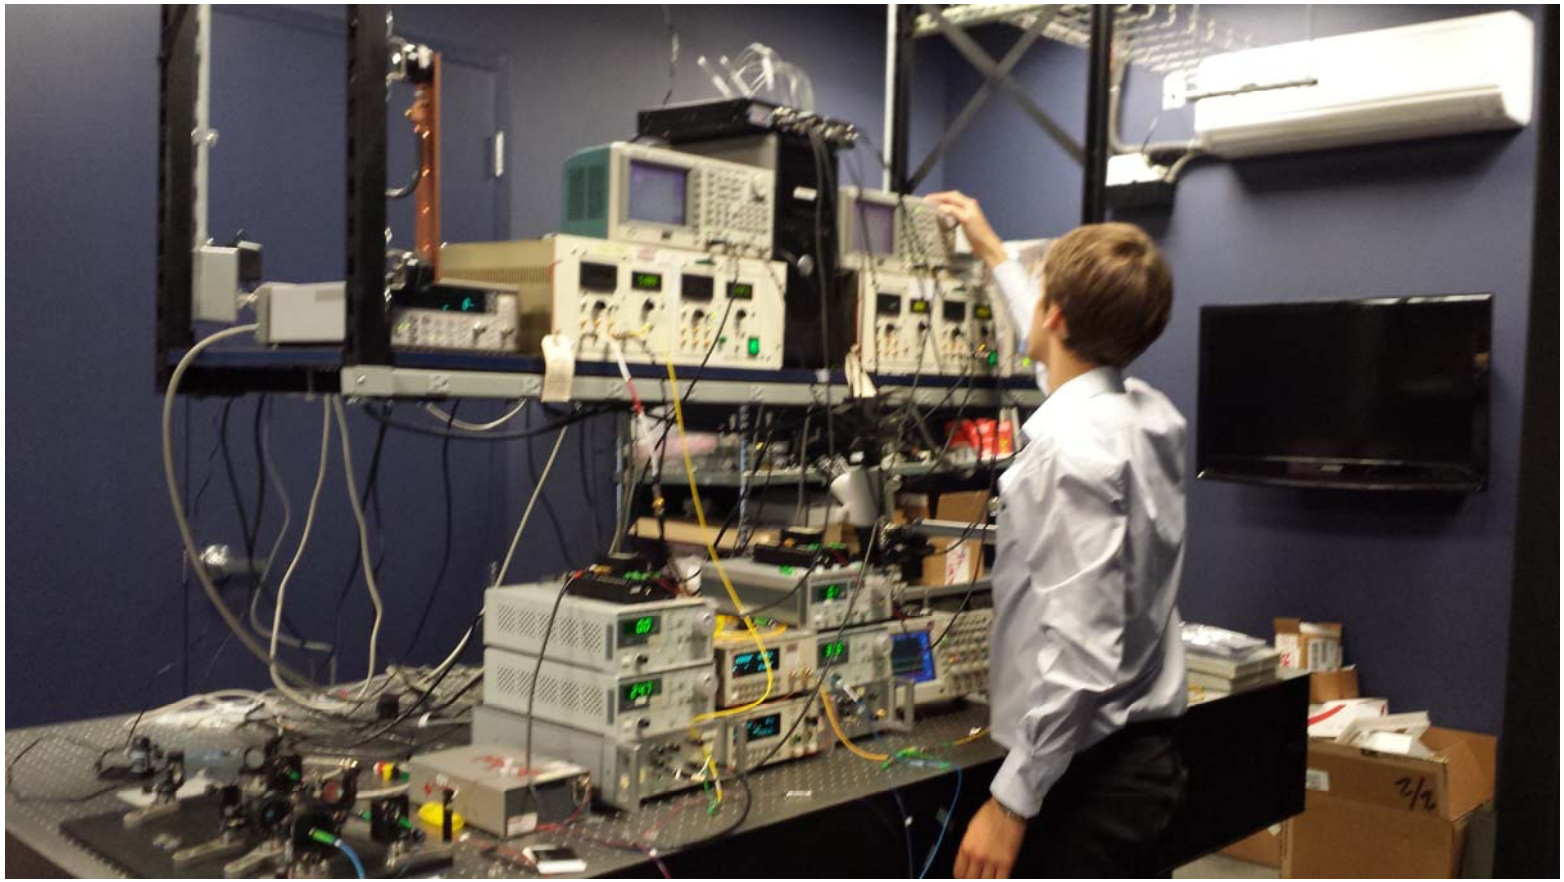

Credit for the setup: Andrei Gheorghe (Amherst College), Jon Habif (BBN) and Monika Patel (BBN)

Preliminary Data - Bob

	Bob	Willie
P(dark click)	≈4.6e-6	≈8.2e-5
P(miss pulse)	≈0.25	≈0.97
c_τ	0.25	
PPM order	32	
R/S code rate	1/2	

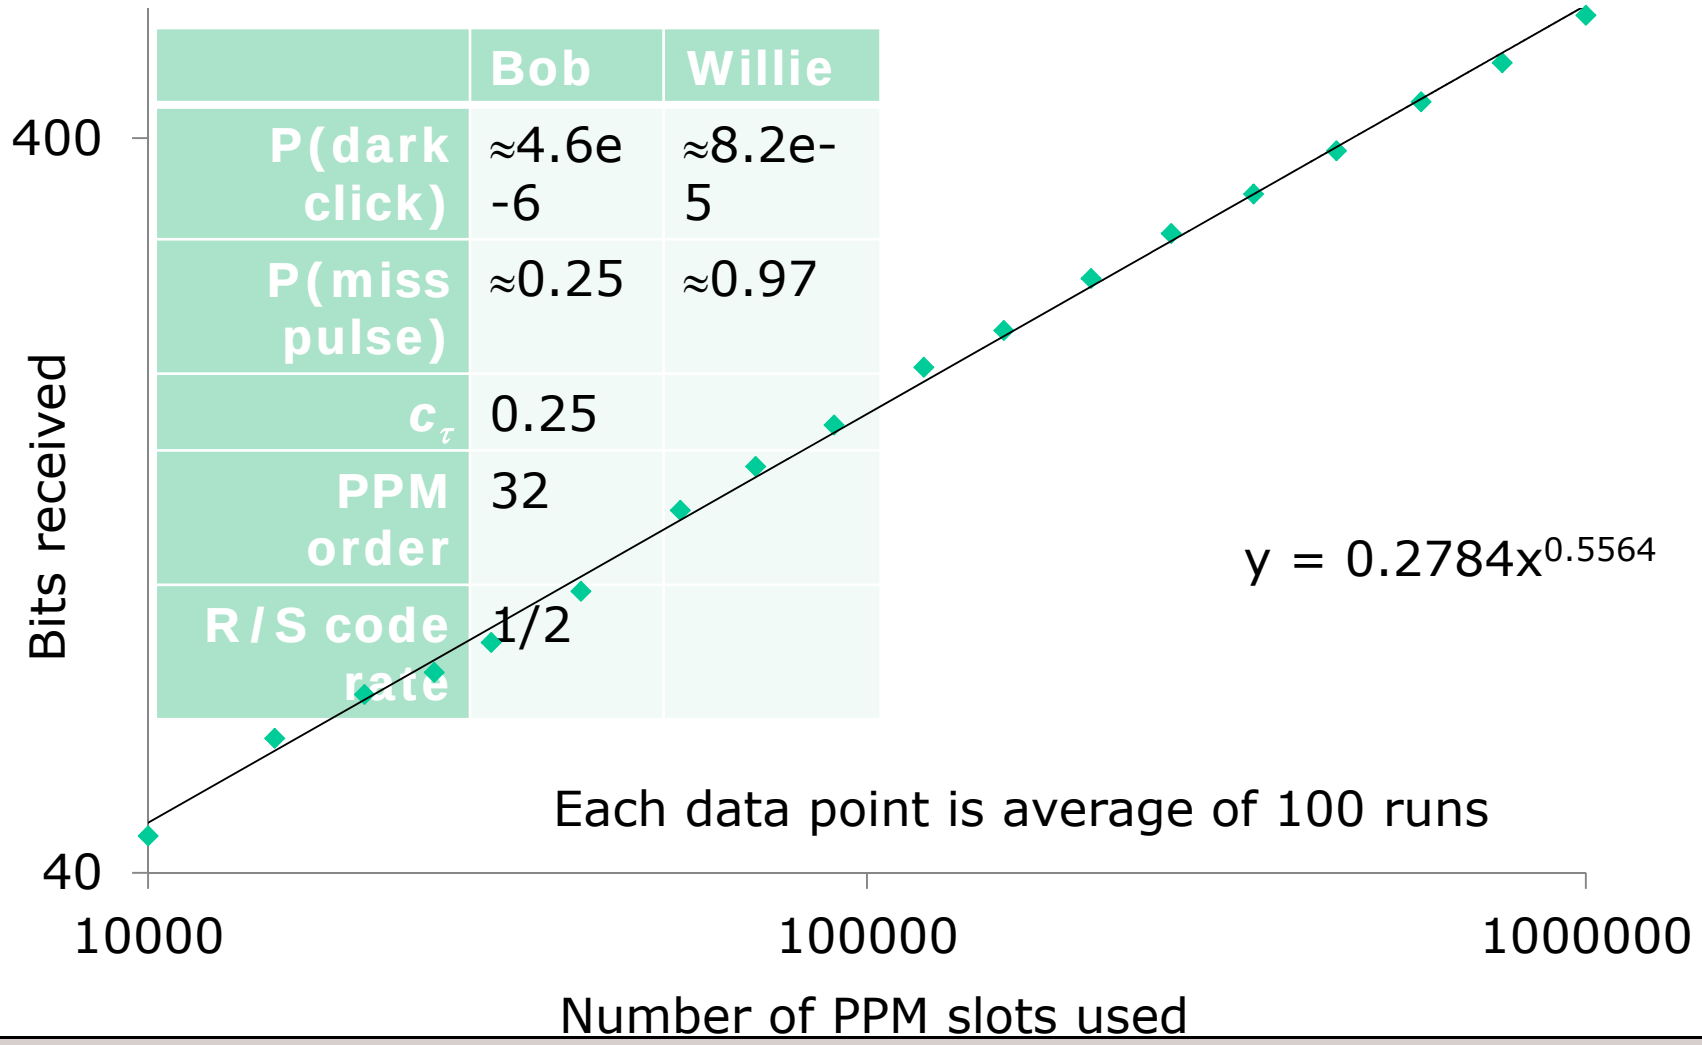

Preliminary Data - Willie

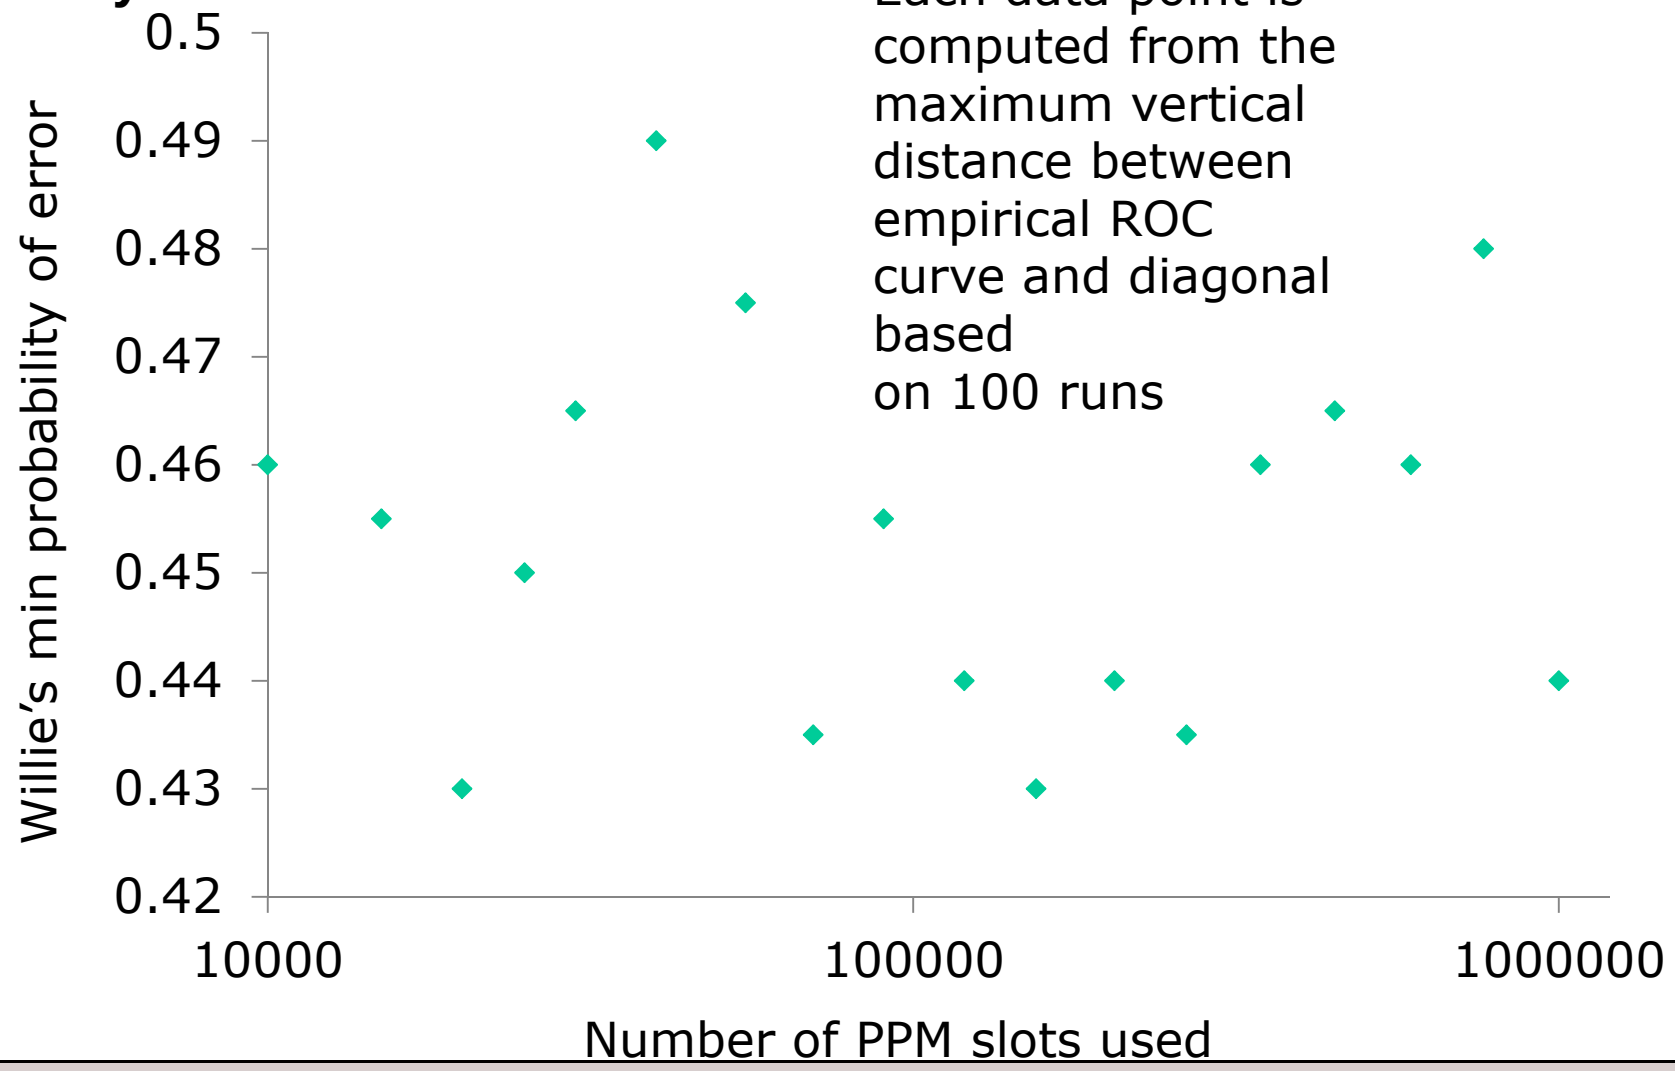

Other Recent Advances in LPD Communications

1. **No Shared Codebook**

For Binary Symmetric Channels (BSCs):

Bob's Channel (p_B:prob of error)

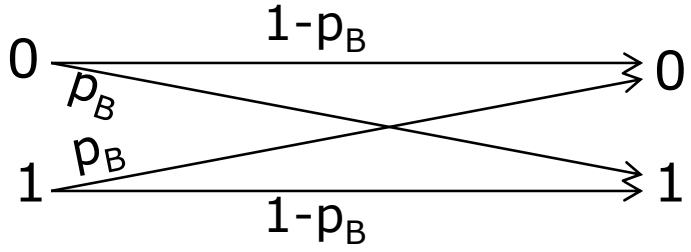

Willie's Channel (p_W:prob of error)

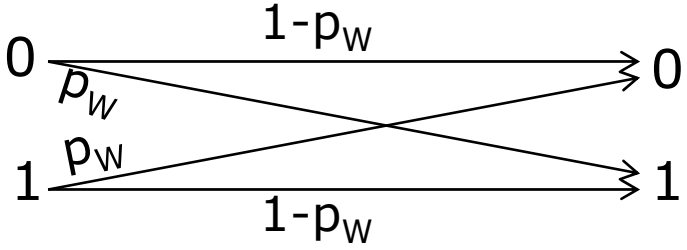

If $p_B > p_W$ (Bob is closer): can get $O(\sqrt{n})$ bits ***without shared codebook.***

[Che, Bakshi, Jaggi, 2013]

Other Recent Advances in LPD Communications

2. **If Willie does not know the time of the message:**

(For example, Alice-to-Bob secret:
"I will send the message at 4:23pm today.")

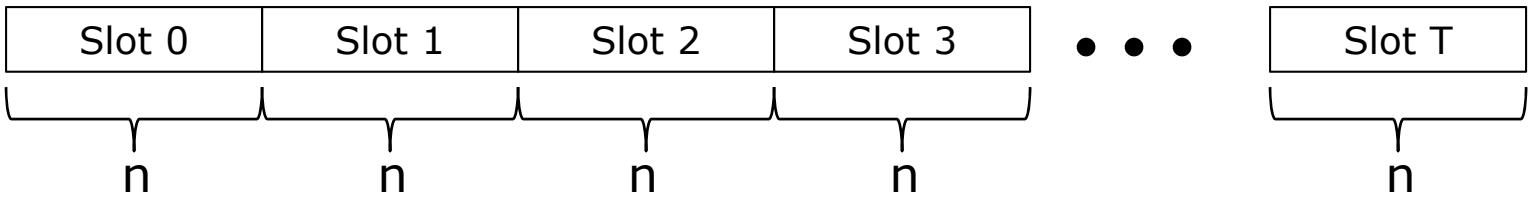

Willie has to watch a much larger time interval.

Can transmit $O(\sqrt{n \log T})$ bits in n channel uses.

[Bash, Goeckel, Towsley, 2014]

Outline

1. Computational and Information Theoretic security basics
2. Potential solutions
3. Asymptotically-large networks
4. Undetectable communications (LPD)
5. **Current and Future Challenges**

Challenges

1. Exploiting when Alice -> Bob **channel is better** than Alice -> Eve
   Challenge: unknown Eve location

2. Exploiting common randomness of **channel reciprocity**
   Challenge: limited number of key bits

3. Exploiting "**public discussion**"
   Challenge: two-way communication and unknown Eve

4. Attacking Eve's **receiver hardware**
   Challenge: short range, assumptions on Eve's hardware

5. **Interference Cancellation**
   Challenge: near-far environment

Is the Network the Solution?

6. **Relay Chattering**
   Challenge: great in theory, not so much in practice (density of nodes).

Challenges for the Future

1. Biggest question: Is information-theoretic security in wireless just a waste of (mostly academic) resources?

There are certainly lots of doubters:

1. "My problem with IT security is you can't guarantee it." -Andrew Worthen, MIT-LL
2. "If cryptographic security primitives are broken, the world collapses with or without IT security." -Dakshi Agrawal, IBM-Watson
3. …

Is it only used in a defense-in-depth approach under cryptographic stuff? But then is there really value-added?

2. **Undetectable Communications**: Can we build "shadow networks"?