Perturbation Hiding and the Batch Steganography Problem

Andrew Ker

adk@comlab.ox.ac.uk

Royal Society University Research Fellow

Oxford University Computing Laboratory

10th Information Hiding Workshop, Santa Barbara, CA

19 May 2008

Perturbation Hiding and the Batch Steganography Problem

Outline

- *The batch steganography problem*
- *Paradox!*
- *Kerckhoffs' principle*
- *Perturbation hiding*
- *Two theorems*
- *Conclusions*

Batch Steganography

- *Spreading a payload amongst multiple covers*

A. Ker, *Batch Steganography & Pooled Steganalysis,* Proc. 8th Information Hiding Workshop, 2006.

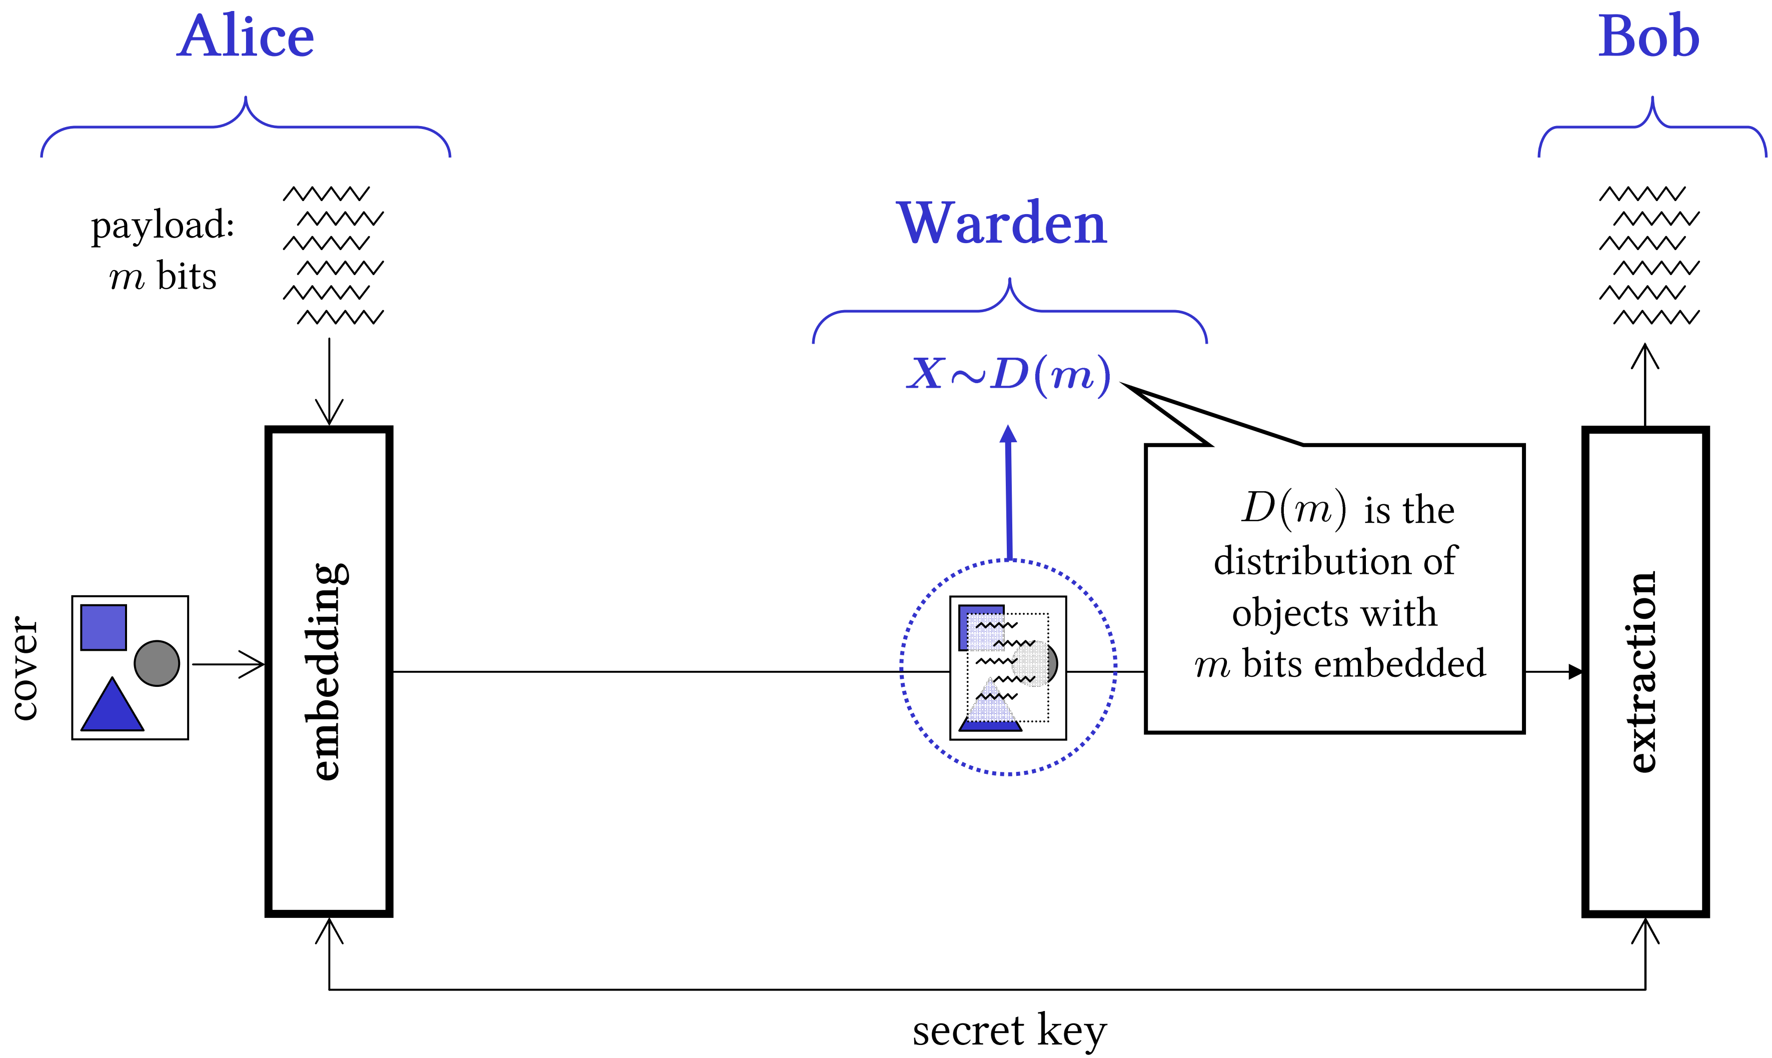

Alice
Warden
Bob
payload:
m bits
$X \sim D(m)$
$D(m)$ is the distribution of objects with m bits embedded
cover
embedding
extraction
secret key

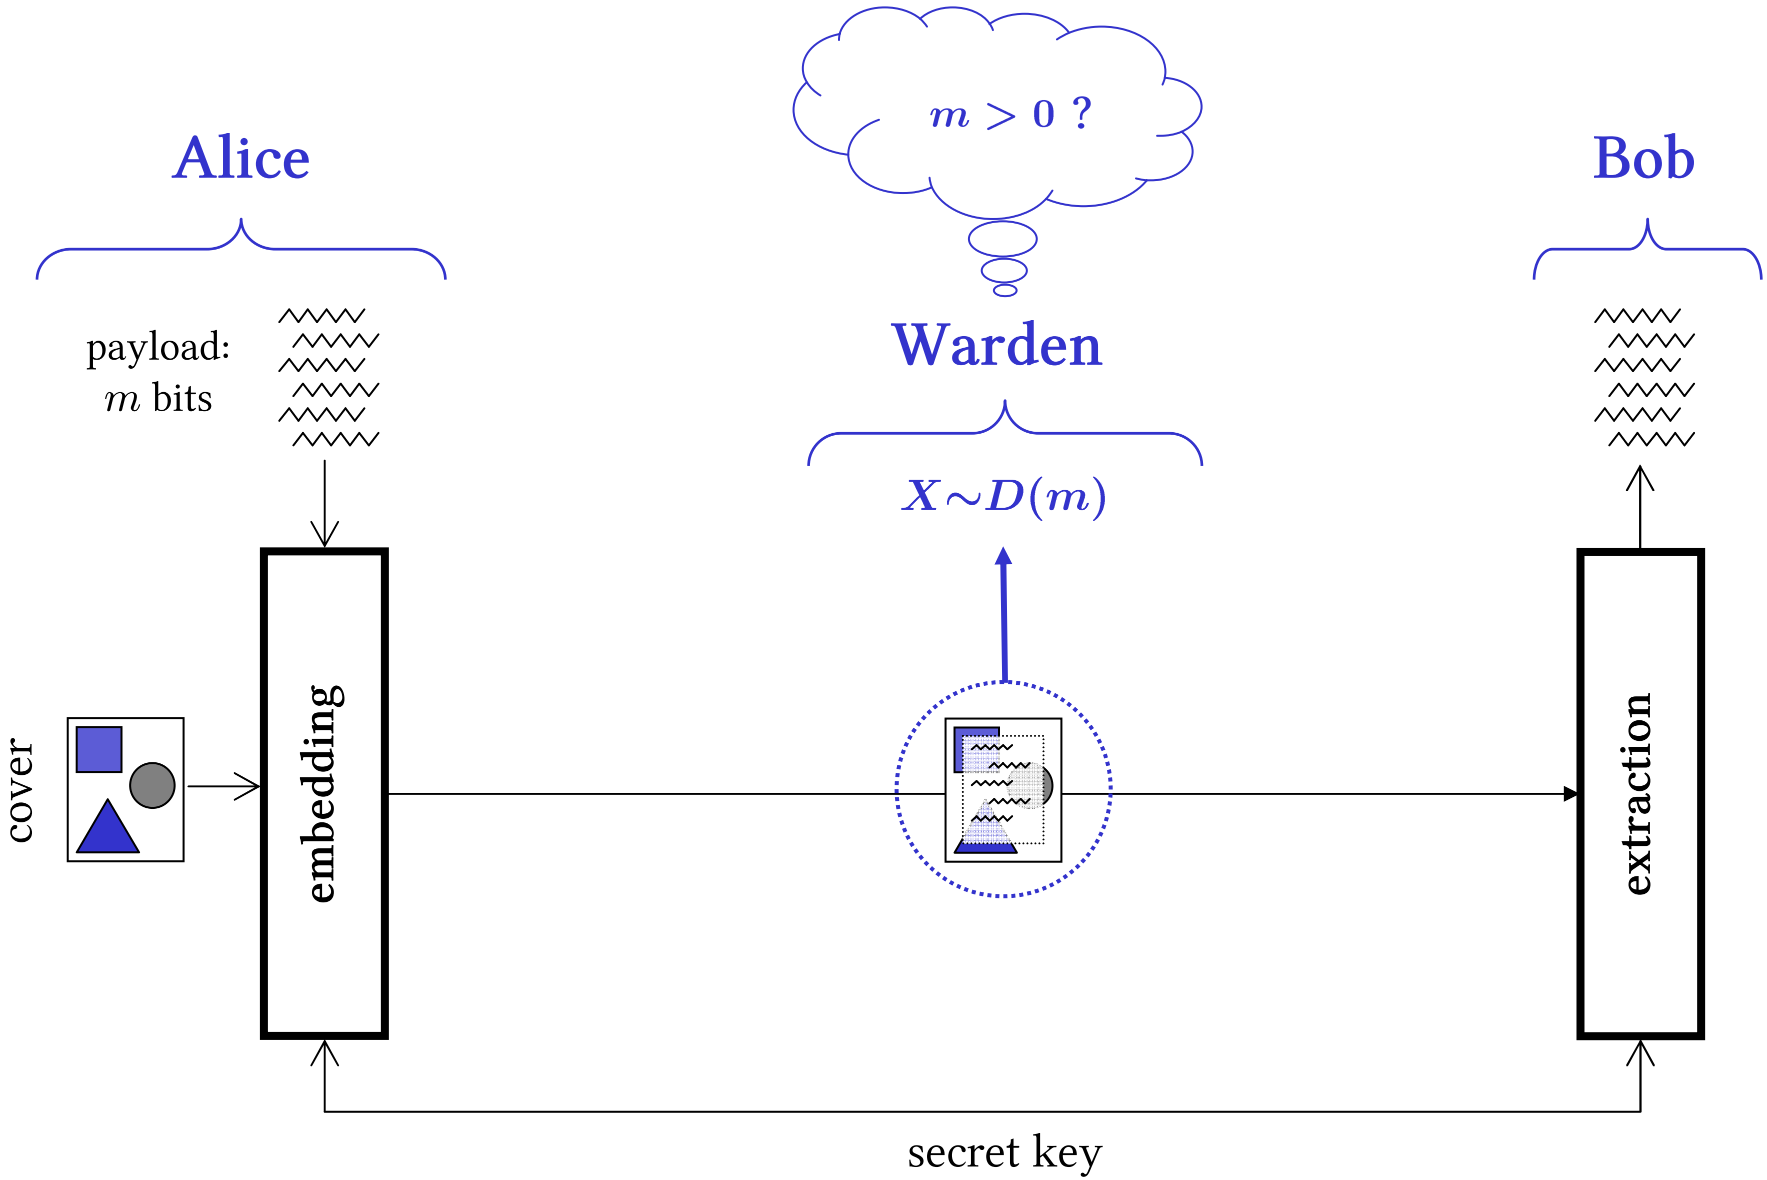

$m > 0$?
Alice
Bob
Warden
payload:
m bits
$X \sim D(m)$
cover
embedding
extraction
secret key

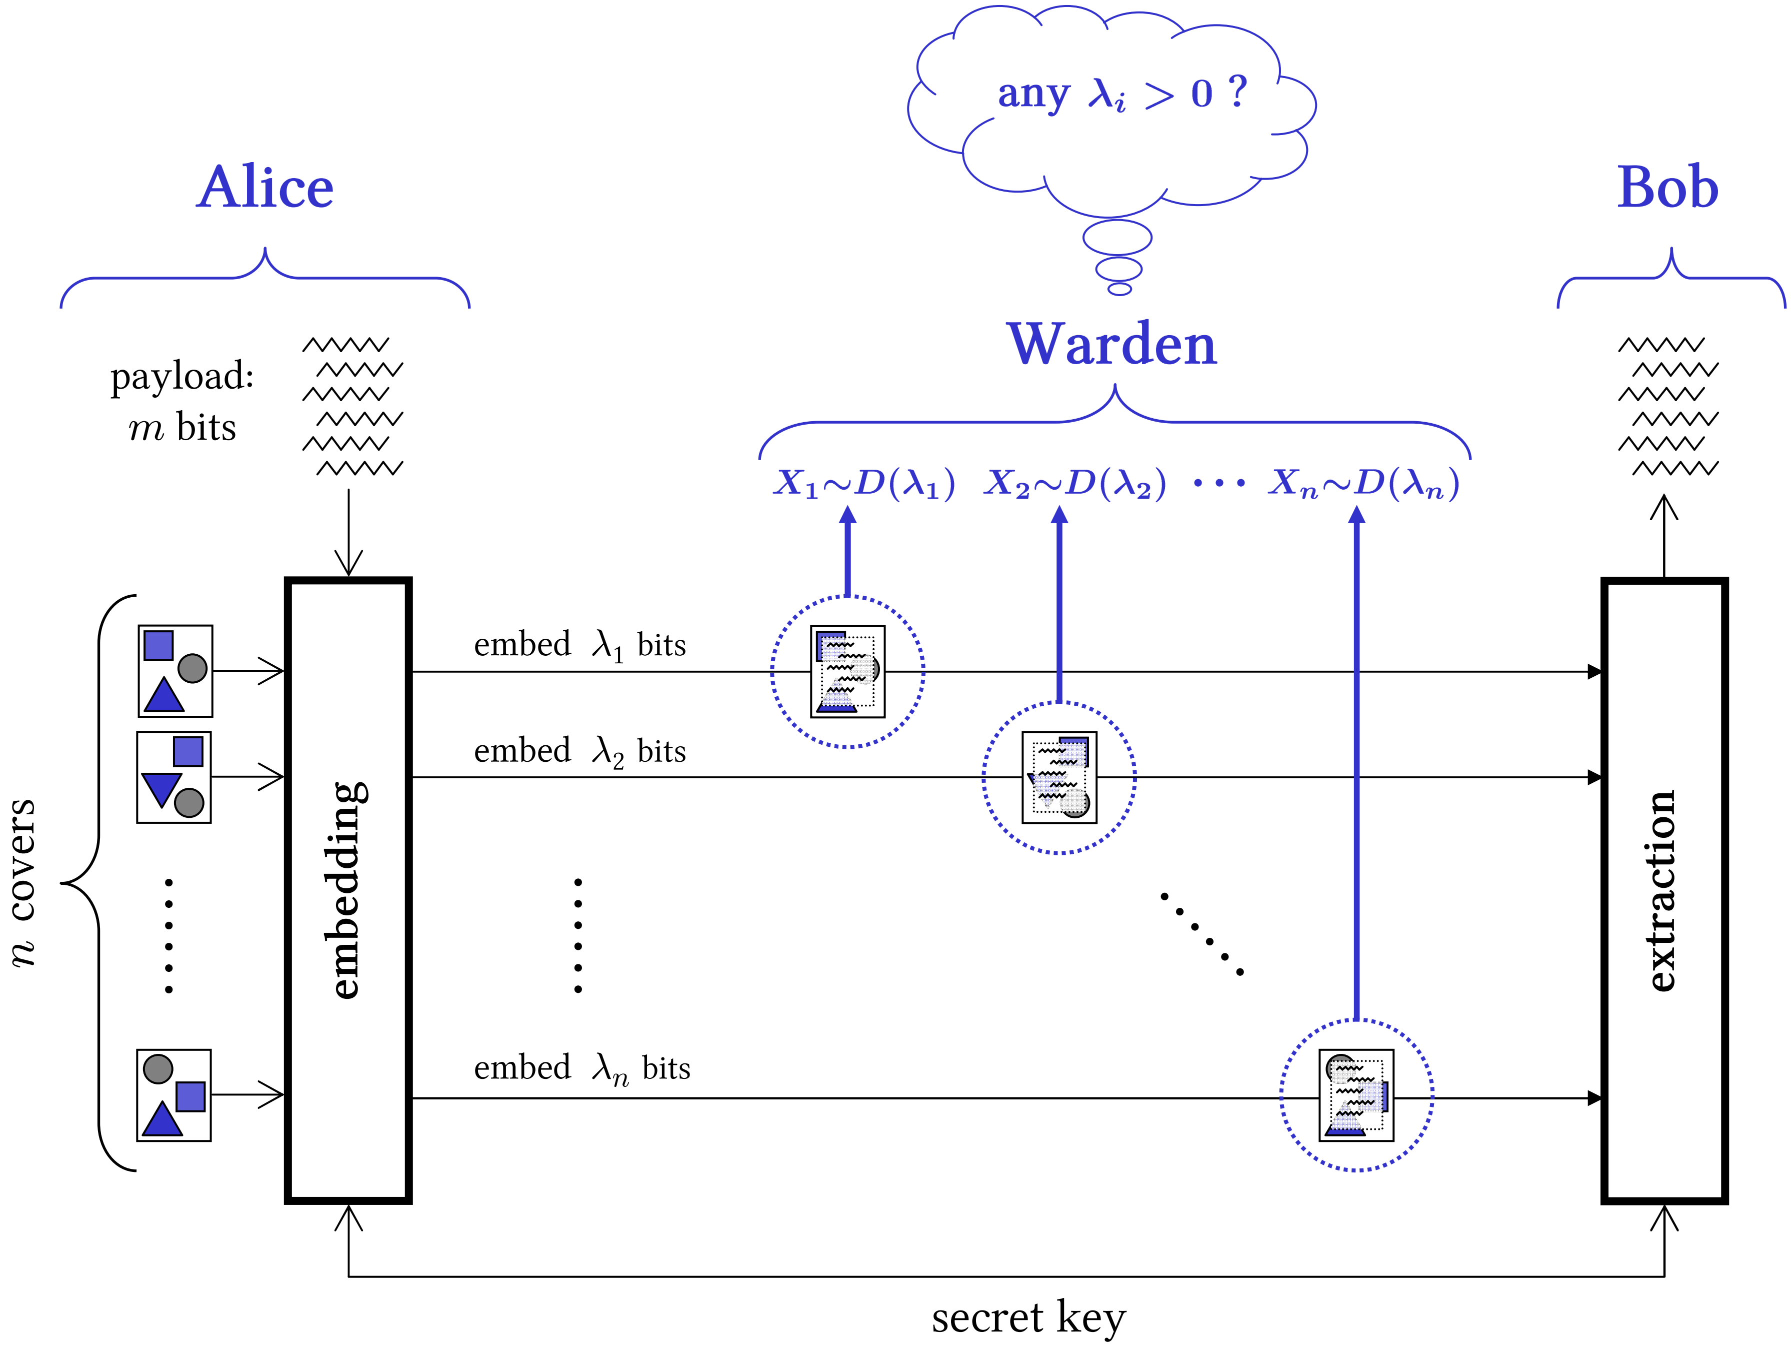

any $\lambda_i > 0$?
Alice
Bob
Warden
payload:
m bits
$X_1 \sim D(\lambda_1)$ $X_2 \sim D(\lambda_2)$ $\cdots$ $X_n \sim D(\lambda_n)$
n covers
embedding
embed λ_1 bits
embed λ_2 bits
embed λ_n bits
extraction
secret key

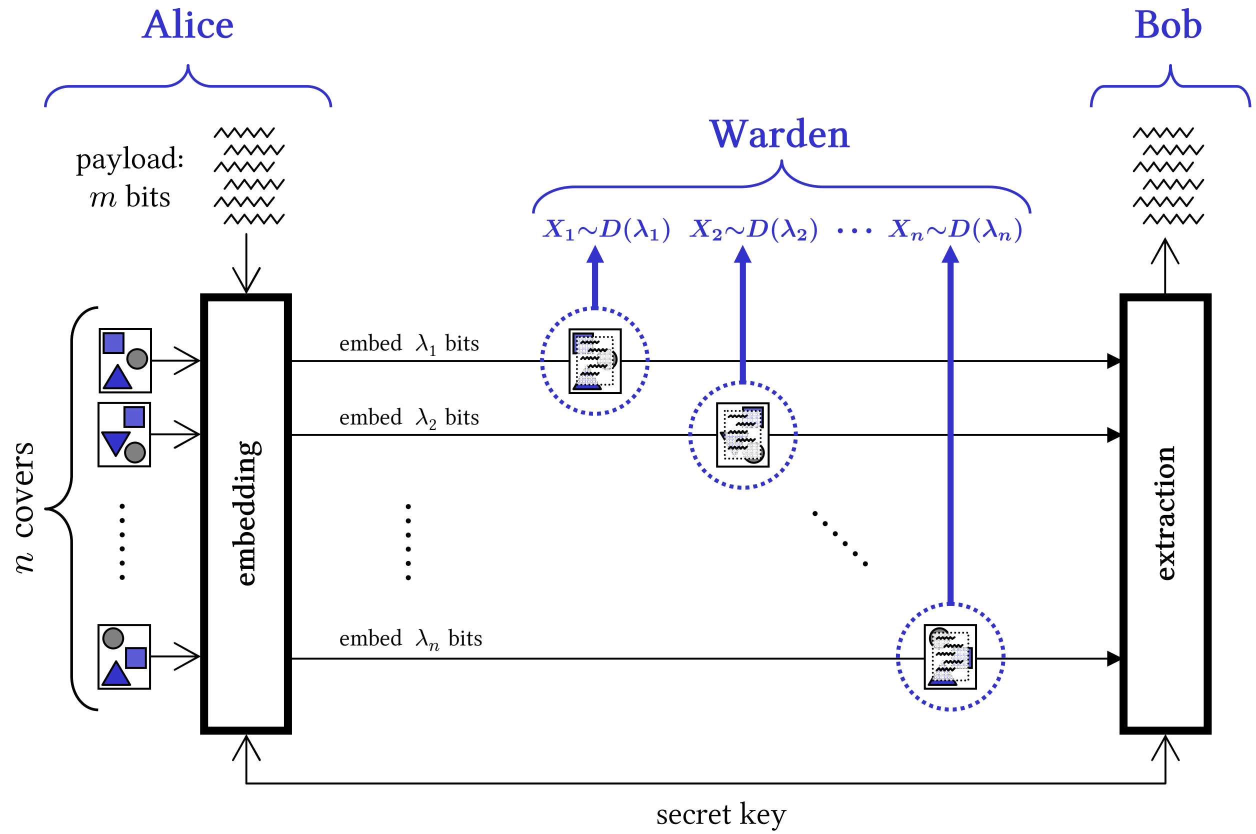

The Batch Steganography Problem

How should Alice distribute her payload between the covers, to make detection as difficult as possible?

- *This question is almost unavoidable in covert communication.*
- *The answer is not as obvious as you might think!*

Analysing Batch Steganography

Can fix a particular behaviour for Warden and optimize with respect to that, e.g. [1], [2].

- *These results have limited applicability.*

Alternatively, consider $D_{\mathrm{KL}}(\boldsymbol{X} \parallel \boldsymbol{Y})$, where $X_i \sim D(0),$
$Y_i \sim D(\lambda_i).$

We can seek to minimize the KL divergence, e.g. [3].

- *These results have a problem.*

[1] A. Ker, *Batch Steganography & Pooled Steganalysis,* Proc. 8th Information Hiding Workshop, 2006.
[2] A. Ker, *Batch Steganography & the Threshold Game,* Proc. SPIE/IS&T Electronic Imaging, 2007.
[3] A. Ker, *Steganographic Strategies for a Square Distortion Function,* Proc. SPIE/IS&T Electronic Imaging, 2008.

Paradox!

Hide m bits in one object out of n, in independent covers:

$$\text{one } \lambda_j = m, \quad \text{all other } \lambda_i = 0.$$

$$D_{\mathrm{KL}}(\boldsymbol{X} \,\|\, \boldsymbol{Y}) = \sum_{i=1}^{n} D_{\mathrm{KL}}\big(D(0) \,\|\, D(\lambda_i)\big) = D_{\mathrm{KL}}\big(D(0) \,\|\, D(m)\big)$$

which is independent of n!

- *No harder to detect 1 stego object in 10 than 1 in 1000 !?*

The problem is that KL divergence bounds the performance of **simple** hypothesis tests. For batch steganography, we have

H_0: all $\lambda_i = 0$ $\qquad$ H_1: some $\lambda_i > 0$

Not a simple hypothesis

but we measured the security of

H_0: all $\lambda_i = 0$ $\qquad$ H_1: $\boldsymbol{\lambda} = \boldsymbol{\lambda}'$

A specific, known, alternative

Kerckhoffs' Principle

"It must not be necessary to keep the system secret:
it should not cause trouble if it falls into enemy hands."

(But we do not assume that the enemy knows the secret crypto key!)

Often coupled with ***chosen plaintext model*** or the ***Dolev-Yao model*** for protocol analysis.

What motivates these pessimistic assumptions?

Conservatism, and in particular the possibility of traitors.

In steganography, what should we grant the Warden?

- ✓ Complete knowledge of embedding algorithm.
- ✓ Complete knowledge of cover source.
- ✗ Complete knowledge of the payload.

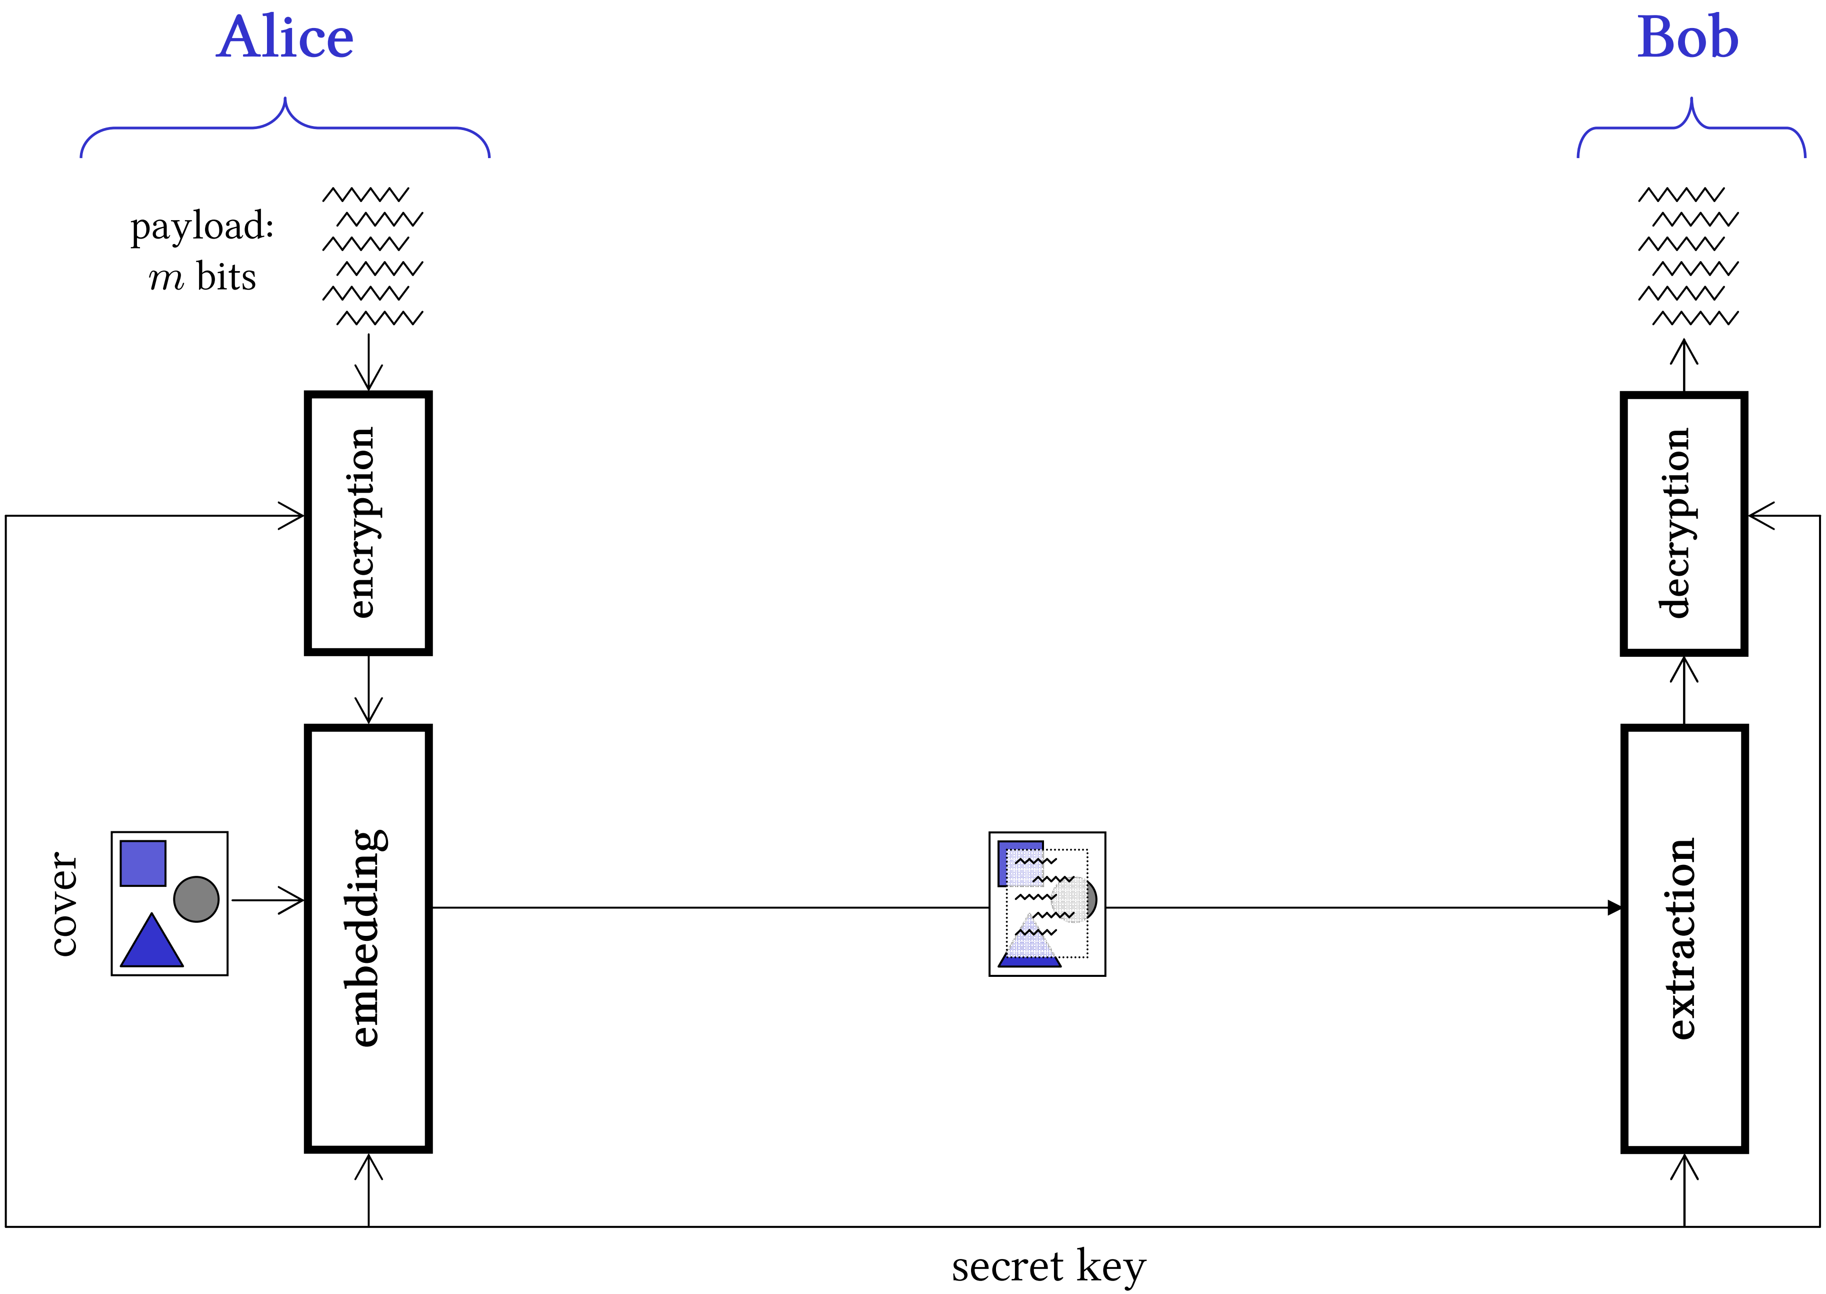

Alice
Bob
payload:
m bits
encryption
decryption
cover
embedding
extraction
secret key

Kerckhoffs' Principle

"It must not be necessary to keep the system secret:
it should not cause trouble if it falls into enemy hands."

(But we do not assume that the enemy knows the secret crypto key!)

Often coupled with ***chosen plaintext model*** or the ***Dolev-Yao model*** for protocol analysis.

What motivates these pessimistic assumptions?

Conservatism, and in particular the possibility of traitors.

In steganography, what should we grant the Warden?

- ✓ Complete knowledge of embedding algorithm.
- ✓ Complete knowledge of cover source.
- ✗ Complete knowledge of the payload.
- ✓ Knowledge of **size** of payload.

Options for Warden's Knowledge

In batch steganography, what should we grant the opponent?

- ✓ Complete knowledge of embedding algorithm for individual objects.
- ✓ Complete knowledge of cover source.
- ✗ Complete knowledge of the payload.

  Knowledge of size of payload...

  - ✗ 1. The sizes of the individual payload chunks in each cover.
  - ✓ 2. The sizes of the individual payload chunks, but not their correspondence with covers.
  - ? 3. The total size of payload, but not its division into chunks.

The Perturbation Hiding Problem

Suppose a fixed one-parameter family of probability distributions $D(\lambda)$ defined for $\lambda \geq 0$, an integer $n \geq 2$, and a constant $m > 0$.

We must choose a nonnegative vector of parameters $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_n)$ subject to $\sum \lambda_i = m$ to minimize $D_{\mathrm{KL}}(\boldsymbol{X} \parallel \boldsymbol{Z})$,

where $X_1, \ldots, X_n$ are iidrv with distribution $D(0)$,

$Y_1, \ldots, Y_n$ are independent with distributions $D(\lambda_1), \ldots, D(\lambda_n)$,

$(Z_1, \ldots, Z_n) = \Pi(Y_1, \ldots, Y_n)$ with Π drawn uniformly at random from S_n.

Places a uniform prior on the correspondence between payload sizes and covers.

Theorem 1

If $D(\lambda)$ is an **exponential family** with a natural reparameterization, the natural parameter is convex nondecreasing, and the variance nondecreasing, in λ, then the solution to the perturbation hiding problem is

$$\lambda_i = m/n$$

i.e. spread payload equally amongst all covers.

It is more important to minimize the overall distortion than to keep the Warden guessing as to the location of the payload.

One such case is $D(\lambda_i) \sim \mathrm{N}\big(\phi(\lambda_i), \sigma^2\big)$ when ϕ is convex and monotonic.

Numerical Results

For distributions which are not an exponential family, we would like to explore the problem numerically.

But

$$D_{\mathrm{KL}}(\boldsymbol{X} \parallel \boldsymbol{Z}) = \mathrm{E}\left[-\log\left(\frac{1}{n!}\sum_{\pi\in S_n}\prod_{i=1}^{n}\frac{f(X_i;\lambda_{\pi(i)})}{f(X_i;0)}\right)\right]_{X_i\sim D(0)}$$

can only be estimated for very small n.

Numerical Results

Let $D(\lambda)$ be a t-distribution with df parameter ν and location parameter λ.

The ν parameter controls the weight of the tails: small $\nu \to$ heavy tails
large $\nu \to$ light tails

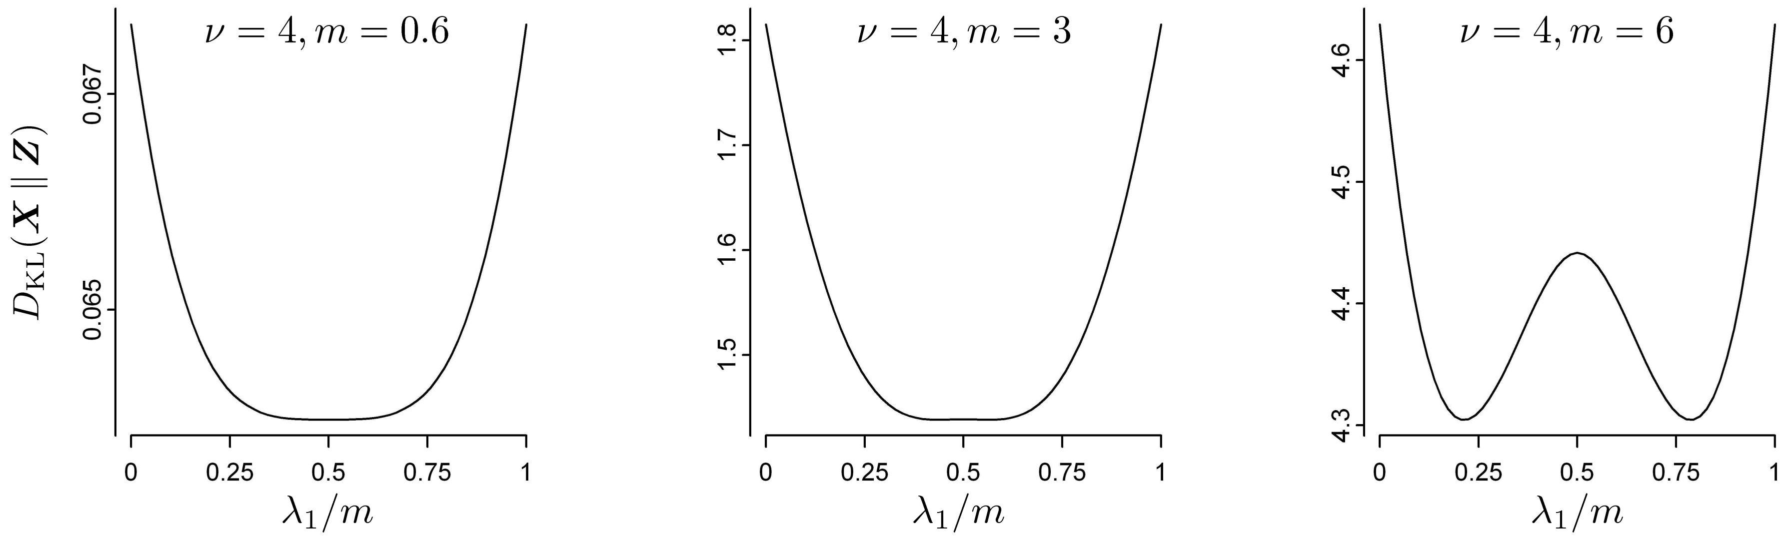

- *Appears that equal distribution is optimal for small enough payload*

Numerical Results

Let $D(\lambda)$ be a t-distribution with df parameter ν and location parameter λ.

The ν parameter controls the weight of the tails: small $\nu \rightarrow$ heavy tails
large $\nu \rightarrow$ light tails

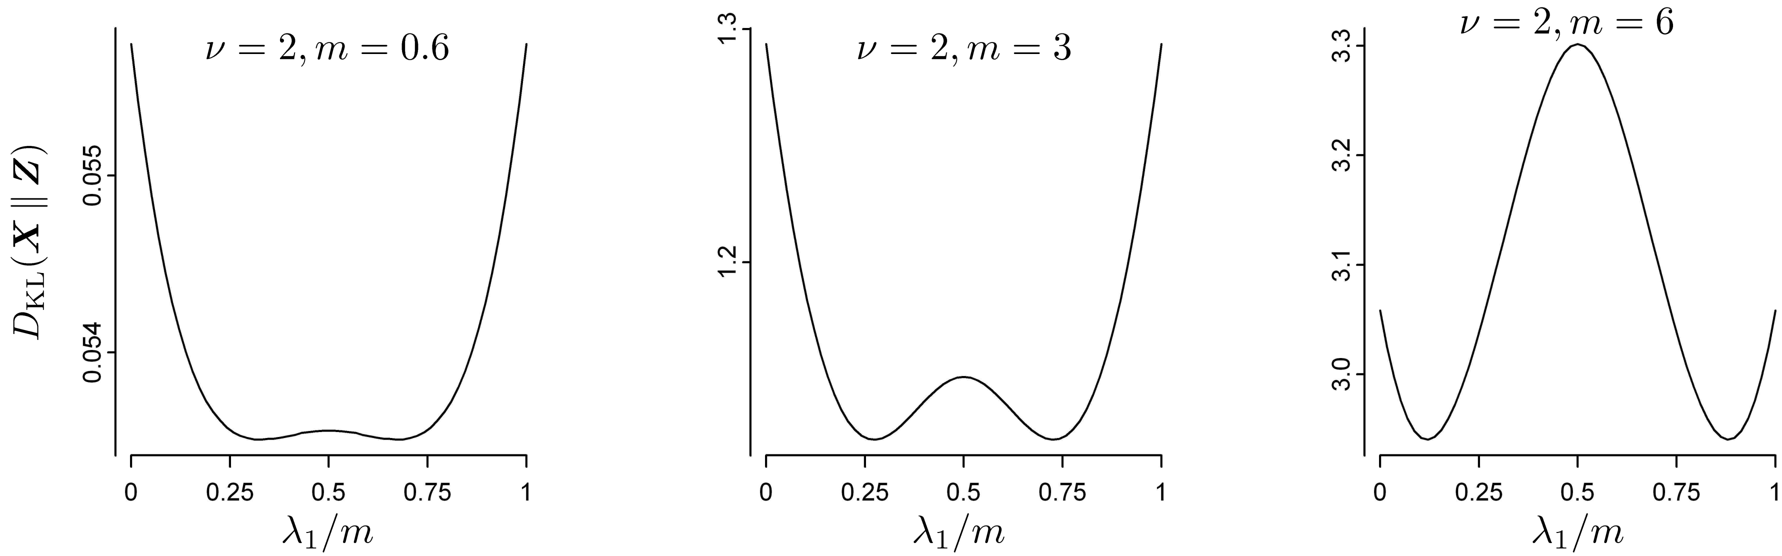

- *Appears that equal distribution is optimal for small enough payload*
  ... for distributions not having heavy tails?

Theorem 2

Assuming sufficient regularity, as $m \to 0$ we have

$$D_{\mathrm{KL}}(\boldsymbol{X} \,\|\, \boldsymbol{Z}) \;\sim\; \underbrace{c_1\left(\sum \lambda_i\right)^2 + c_2\left(\sum \lambda_i\right)^3}_{\textit{fixed}} + \underbrace{c_3\left(\sum \lambda_i^2\right)\left(\sum \lambda_i\right)} + O(m^4)$$

$$c_3 = \frac{1}{2n}\mathrm{E}\left[\ell_\lambda(X)^3 + \ell_\lambda(X)\ell_{\lambda\lambda}(X)\right]$$

$$\ell_\lambda(x) = \frac{\partial}{\partial\lambda}\log f(x;\lambda)|_{\lambda=0}$$

$$\ell_{\lambda\lambda}(x) = \frac{\partial^2}{\partial\lambda^2}\log f(x;\lambda)|_{\lambda=0}$$

$c_3 > 0 \longrightarrow$ ***spread payload equally***

$c_3 < 0 \longrightarrow$ ***concentrate as much as possible***

$c_3 = 0 \longrightarrow$ ***need to go to fourth order***

Conclusions

- The Perturbation Hiding problem is a mathematical abstraction of batch steganography.

  Its level of abstraction is different to traditional information-theoretic analyses of security.

- It has been solved for the case of convex exponential families, and we have explored its asymptotics for small payloads.

  Future work:
  - ***nonuniform covers,***
  - ***asymptotics for large n.***

- We must be careful about the information asymmetry in the batch steganography problem.

  Kerckhoffs' Principle has to be interpreted carefully in steganography, and particularly in batch steganography.

End

adk@comlab.ox.ac.uk