The Square Root Law Does Not Require a Linear Key

Andrew Ker

adk@comlab.ox.ac.uk

Royal Society University Research Fellow
Oxford University Computing Laboratory

12th ACM Multimedia & Security Workshop
Rome, Italy, 10 September 2010

The Square Root Law Does Not Require a Linear Key

Outline

- Imperfect embedding
- Square root law & a linear key
- Asymptotically perfect security with no stego key
  - definition of security in the absence of a key
- Asymptotically perfect security with Hamming syndrome codes

Imperfect embedding

Perfect embedding preserves **all** statistics of the cover source.

- It is undetectable.
- It has a linear capacity law.
- It is not practically realisable.

We contend that all practical steganography is imperfect.

Imperfect embedding makes **changes** to elements of the cover, in a way which does not preserve their statistics.

- Capacity follows a ‘Square Root Law’.

Notation:	cover size n	(‘pixels’)
	payload size m	(bits)

Classic square root law

Cover consists of n 'pixels', some are used to carry payload,
of which some are changed.

Model:

- Cover pixels: i.i.d. random variables with p.m.f. $p(x)$,
- Changed pixels: i.i.d. random variables with p.m.f. $q(x)$,
- Embedding: m used pixels selected uniformly at random, each changed with probability ½,
- $p(x)$ known to the detector, $\forall x . p(x) \neq 0, 1,\ \ \exists y . p(y) \neq q(y)$.

perfect information

covers not deterministic

imperfect embedding

Classic square root law

Cover consists of n 'pixels', some are used to carry payload,
of which some are changed.

Model:

- **Cover pixels:** **i.i.d. random variables with p.m.f.** $p(x)$,
- **Changed pixels: i.i.d. random variables with p.m.f.** $q(x)$,
- **Embedding:** m **used pixels selected uniformly at random, each changed with probability ½,**
- $p(x)$ **known to the detector,** $\forall x . p(x) \neq 0, 1, \quad \exists y . p(y) \neq q(y)$.

As cover size $n \to \infty$,

1. If $m/\sqrt{n} \to \infty$ then an asymptotically perfect detector exists.
2. If $m/\sqrt{n} \to 0$ then we have asymptotically perfect security.

So $\sqrt{n}$ ***is the critical payload 'rate'.***

Classic square root law

Cover consists of n 'pixels', some are <u>used</u> to carry payload,
of which some are <u>changed</u>.

Model:

- **Cover pixels:** **i.i.d. random variables with p.m.f.** $p(x)$,
- **Changed pixels:** **i.i.d. random variables with p.m.f.** $q(x)$,
- **Embedding:** m **used pixels selected uniformly at random,** **each changed with probability ½,**
- $p(x)$ **known to the detector,** $\forall x . p(x) \neq 0, 1, \quad \exists y . p(y) \neq q(y)$.

To tell the recipient which pixels are used requires $O(m \log m)$ stego key.

<u>Theorem (MM&Sec 09)</u>
If the stego key length is not at least $O(m)$ then an asymptotically perfect detector exists, regardless of payload rate.

Non-shared selection channel

Avoid telling the recipient the location of the changes (but still have the message extractable).

Well-solved by wet paper codes [Fridrich et al, 2004]:

- Reduce everything to binary (e.g. pixel LSBs).
- Create an $m \times n$ matrix $\mathbf{D}$ (possibly public). ***How generated?***
- Change the cover $\boldsymbol{c}$ into a stego object $\boldsymbol{s}$ such that $\mathbf{D}\boldsymbol{s} = \boldsymbol{p}$, where $\boldsymbol{p}$ is the desired payload.

How many changes? ***How solved?***

Difficult to analyse the predictability of the changes.

Possible flaws already highlighted [Böhme, 2005].

Simplest example

[Anderson & Petitcolas, 1998]

- Reduce everything to binary (e.g. pixel LSBs).
- Divide into m groups:

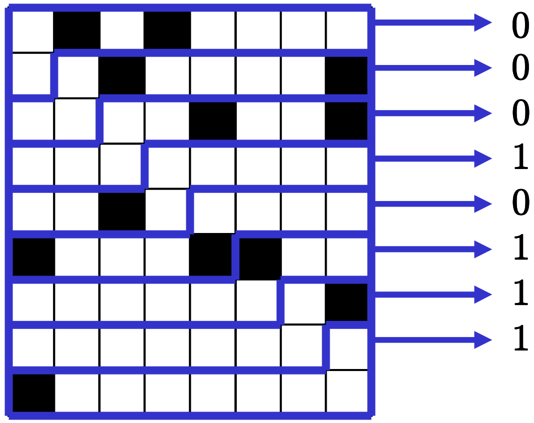

(The groups can be made public.)

- Carry payload bit i as the parity of the sum of the pixels in group i.
- When a group in the cover needs its parity flipping, pick one of its pixels to change uniformly at random.

— ***We know exactly how predictable the changes are.***

Theorem

Cover consists of m publicly known groups of pixels each of size $\lfloor n/m \rfloor$.

Model:

- **Cover pixels:** **i.i.d. random variables with p.m.f.** $p(x)$,
- **Changed pixels:** **i.i.d. random variables with p.m.f.** $q(x)$,
- **Embedding:** **each group unchanged with probability ½, otherwise one randomly selected pixel changed,**
- $p(x)$ **known to the detector,** $\forall x . p(x) \neq 0, 1, \quad \exists y . p(y) \neq q(y)$.

As cover size $n \to \infty$,

1. If $m/\sqrt{n} \to \infty$ then an asymptotically perfect detector exists.
2. If $m/\sqrt{n} \to 0$ then we have asymptotically perfect 'security'.

Up to $\sqrt{n}$ groups, each at least $\sqrt{n}$ big, spreads the payload thinly enough.

Proof idea

Consider one group of pixels $(X_1, \ldots, X_k),\ k = \lfloor n/m \rfloor$.

Let $\mathcal{P}$ and $\mathcal{Q}$ be the probability laws for cover and stego pixel groups.

$$
\begin{aligned}
&D_{\mathrm{KL}}(\text{cover} \parallel \text{stego}) \\
&\quad = \quad m\, D_{\mathrm{KL}}(\mathcal{P} \parallel \mathcal{Q}) \\
&\quad = \quad -m\, \mathrm{E}\left[\log\left(\frac{\frac{1}{2}\prod_i p(X_i) + \frac{1}{2k}\sum_j q(X_j)\prod_{i\neq j} p(X_i)}{\prod_i p(X_i)}\right)\right] \\
&\quad = \quad -m\, \mathrm{E}\left[\log\left(\boxed{\frac{1}{2} + \frac{1}{2k}\sum_{j=1}^{k}\frac{q(X_j)}{p(X_j)}}\right)\right]
\end{aligned}
$$

Random variable, mean 1, satisfying some analytic conditions.

Proof idea

Consider one group of pixels $(X_1, \ldots, X_k),\ k = \lfloor n/m \rfloor$.

Let $\mathcal{P}$ and $\mathcal{Q}$ be the probability laws for cover and stego pixel groups.

$$
\begin{aligned}
&D_{\mathrm{KL}}(\text{cover} \parallel \text{stego}) \\
&\quad= \quad m\, D_{\mathrm{KL}}(\mathcal{P} \parallel \mathcal{Q}) \\
&\quad= \quad -m\,\mathrm{E}\left[\log\left(\frac{\frac{1}{2}\prod_i p(X_i) + \frac{1}{2k}\sum_j q(X_j)\prod_{i\neq j} p(X_i)}{\prod_i p(X_i)}\right)\right] \\
&\quad= \quad -m\,\mathrm{E}\left[\log\left(\frac{1}{2} + \frac{1}{2k}\sum_{j=1}^{k}\frac{q(X_j)}{p(X_j)}\right)\right] \\
&\quad\sim \quad \frac{m}{2}\,\mathrm{Var}\left[\frac{1}{2k}\sum_{j=1}^{k}\frac{q(X_j)}{p(X_j)}\right] \\
&\quad\sim \quad \frac{m^2}{2n}\boxed{\tfrac{1}{4}\mathrm{Var}\left[\tfrac{q(X_1)}{p(X_1)}\right]}
\end{aligned}
$$

This coefficient is known as Steganographic Fisher Information (SFI).

It turns out that the SFI of uniformly-spread embedding is also

$$\tfrac{1}{4}\mathrm{Var}\left[\tfrac{q(X_1)}{p(X_1)}\right].$$

Steganographic security

With the prior scheme, there is no steganographic key at all.

- Everyone knows the pixel groups and so can read the message.
- The content is not confidential.

Steganographic security is distinct from cryptographic security and the latter still requires a shared crypto key.

NB: if the hidden payload is encrypted, the encryption must have the property that cyphertexts cannot easily be recognised.

Syndrome codes

Making one change to a group of pixels can carry more than one bit.

Well-studied topic called matrix embedding [Crandall, 1998].

- Divide pixels into groups.
- Use syndromes of some code with low covering radius
  (like solving $\mathbf{D}\boldsymbol{s} = \boldsymbol{p}$ in each group).

Again, we should be concerned that the locations of the changes might be predictable.

Theorem

Payload of size m, embedded using largest possible binary Hamming code.

Model:

- **Cover pixels:** **i.i.d. random variables with p.m.f.** $p(x)$**,**
- **Changed pixels:** **i.i.d. random variables with p.m.f.** $q(x)$**,**
- **Embedding:** **make minimum changes and move to uniformly random coset.**
- $p(x)$ **known to the detector,** $\forall x . p(x) \neq 0, 1, \quad \exists y . p(y) \neq q(y)$**.**

As cover size $n \to \infty$,

[1. If $m/\sqrt{n} \log n \to \infty$ then an asymptotically perfect detector exists.]

2. If $m/\sqrt{n} \log n \to 0$ then we have asymptotically perfect security.

Conclusions

- The old 'parity of a block' idea is asymptotically perfectly secure, below the square root bound.
  - *The opponent gains nothing by knowing the groups.*
  - *No stego key is required: 'public key steganography'.*
  - *A crypto key is still required, for confidentiality.*

- The old 'syndrome of a Hamming code' idea is asymptotically perfectly secure, with number of changes below the square root bound.
  - *This means the payload capacity is* $O(\sqrt{n}\log n)$.

- We should consider the finer asymptotics of matrix embedding and related schemes.
  - *Steganographic Fisher Information:*

$$2\lim D_{\mathrm{KL}}(\text{cover} \parallel \text{stego})\frac{n}{m^2}$$

Conclusions

- The old 'parity of a block' idea is asymptotically perfectly secure, below the square root bound.
  - *The opponent gains nothing by knowing the groups.*
  - *No stego key is required: 'public key steganography'.*
  - *A crypto key is still required, for confidentiality.*

- The old 'syndrome of a Hamming code' idea is asymptotically perfectly secure, with number of changes below the square root bound.
  - *This means the payload capacity is* $O(\sqrt{n}\log n)$.

- We should consider the finer asymptotics of matrix embedding and related schemes.
  - *'Equivalent Steganographic Fisher Information':*

$$2\lim D_{\mathrm{KL}}(\text{cover} \parallel \text{stego})\frac{n\log n}{m^2} \quad ?$$