Quantitative Evaluation of Pairs and RS Steganalysis

Andrew Ker

Oxford University Computing Laboratory
adk@comlab.ox.ac.uk

Royal Society University Research Fellow / Junior Research Fellow
at University College, Oxford

SPIE EI'04 19 January 2004

Simple Classification

The primary question an Information Security Officer (Warden) wants to ask is

Does this image contain hidden data?

(as opposed to estimating any hidden message length or trying to decode any hidden data).

This work focuses solely on evaluating the reliability of hypothesis tests for this question.

Reliability is a two-dimensional measure, showing how false positive and missed detections trade off against each other. Traditionally this is displayed as a Region of Confidence curve.

Distributed Steganalysis Project

- A number of large libraries of “natural images” (many JPEG compressed)

  Currently over 30,000 images in total, with more to come

- Optimised and portable program to simulate steganography and compute detection statistics

  Includes over 100 variants of steganalysis statistics

- Heterogeneous cluster of computing machines to spread the work

  Has been 7-50 machines at any one time

- Calculations queued and results stored in a relational database

  Presently over 13 million rows of data, expected to grow to over 100 million

Scope of Investigations

Covers

Grayscale bitmaps (which quite likely were previously subject to JPEG compression)

Embedding method

LSB steganography using a set proportion of evenly-spread pixels

Steganalysis statistics

"Pairs" [Fridrich *et al*, SPIE'03]

"RS" a.k.a. "dual statistics" [Fridrich *et al*, ACM Workshop '01]

Will focus on "interesting" cases, in this case embedding rates of 0.01-0.2 secret bits per cover pixel.

Sample Output

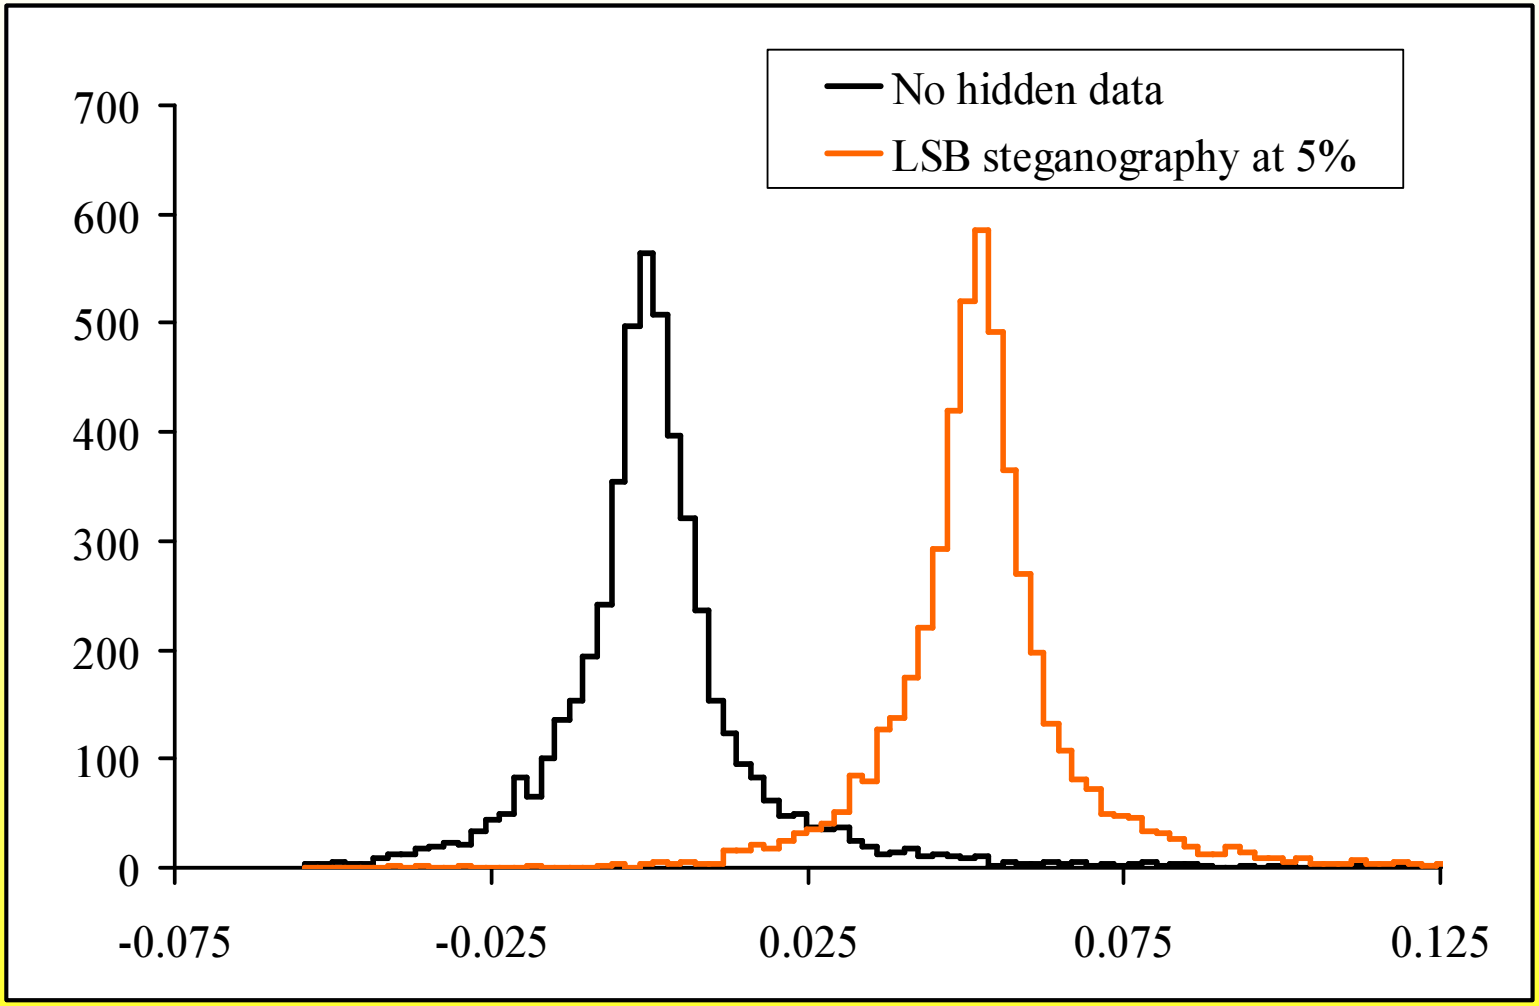

Histograms of the standard RS statistic, generated from 5000 JPEG images.

Sample Output

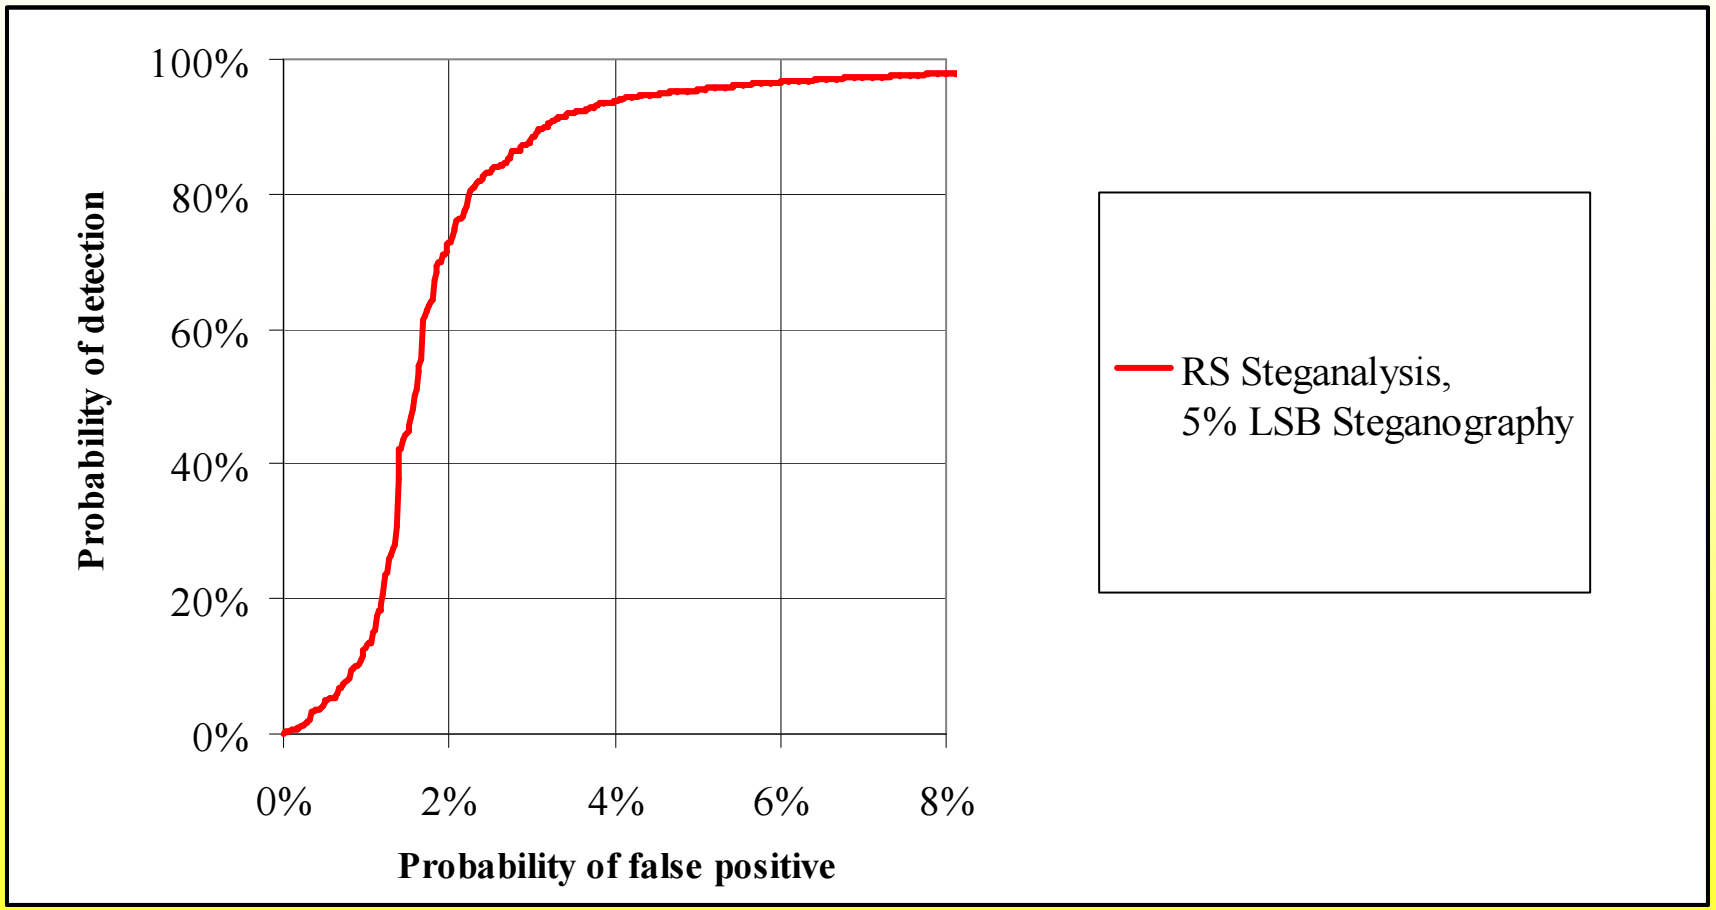

ROC curves generated from 5,000 JPEG images

Sample Output

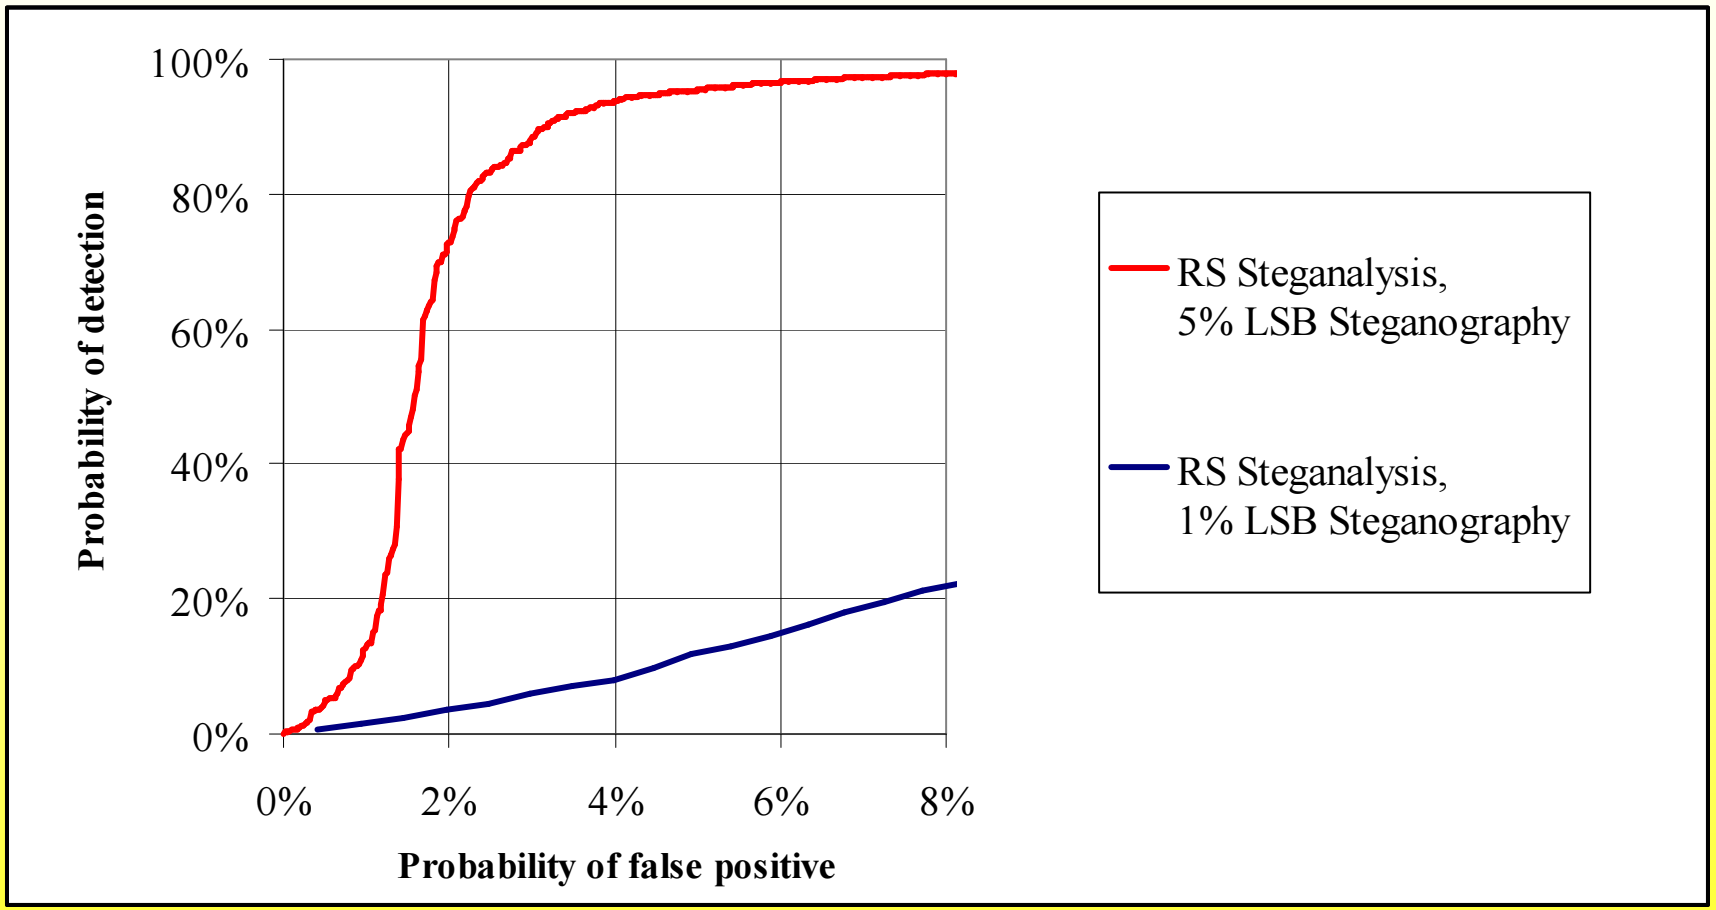

ROC curves generated from 5,000 JPEG images

Choosing the RS "Mask"

The "mask" in RS Steganalysis determines how the pixels are grouped and which pixels of each group are LSB-flipped.

In [Fridrich *et al*, ACM Workshop '01] the masks $[0 \quad 1 \quad 1 \quad 0]$ and $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ were used.

We experimented with a number of alternative masks including:

$[0 \quad 1 \quad 0]$ $[0 \quad 1 \quad 1 \quad 1 \quad 0]$ $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ $\begin{bmatrix} 0 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 0 \end{bmatrix}$ $\begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$

Uniformly, $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ was the best performer.

Choosing the RS "Mask"

There is a small but useful improvement:

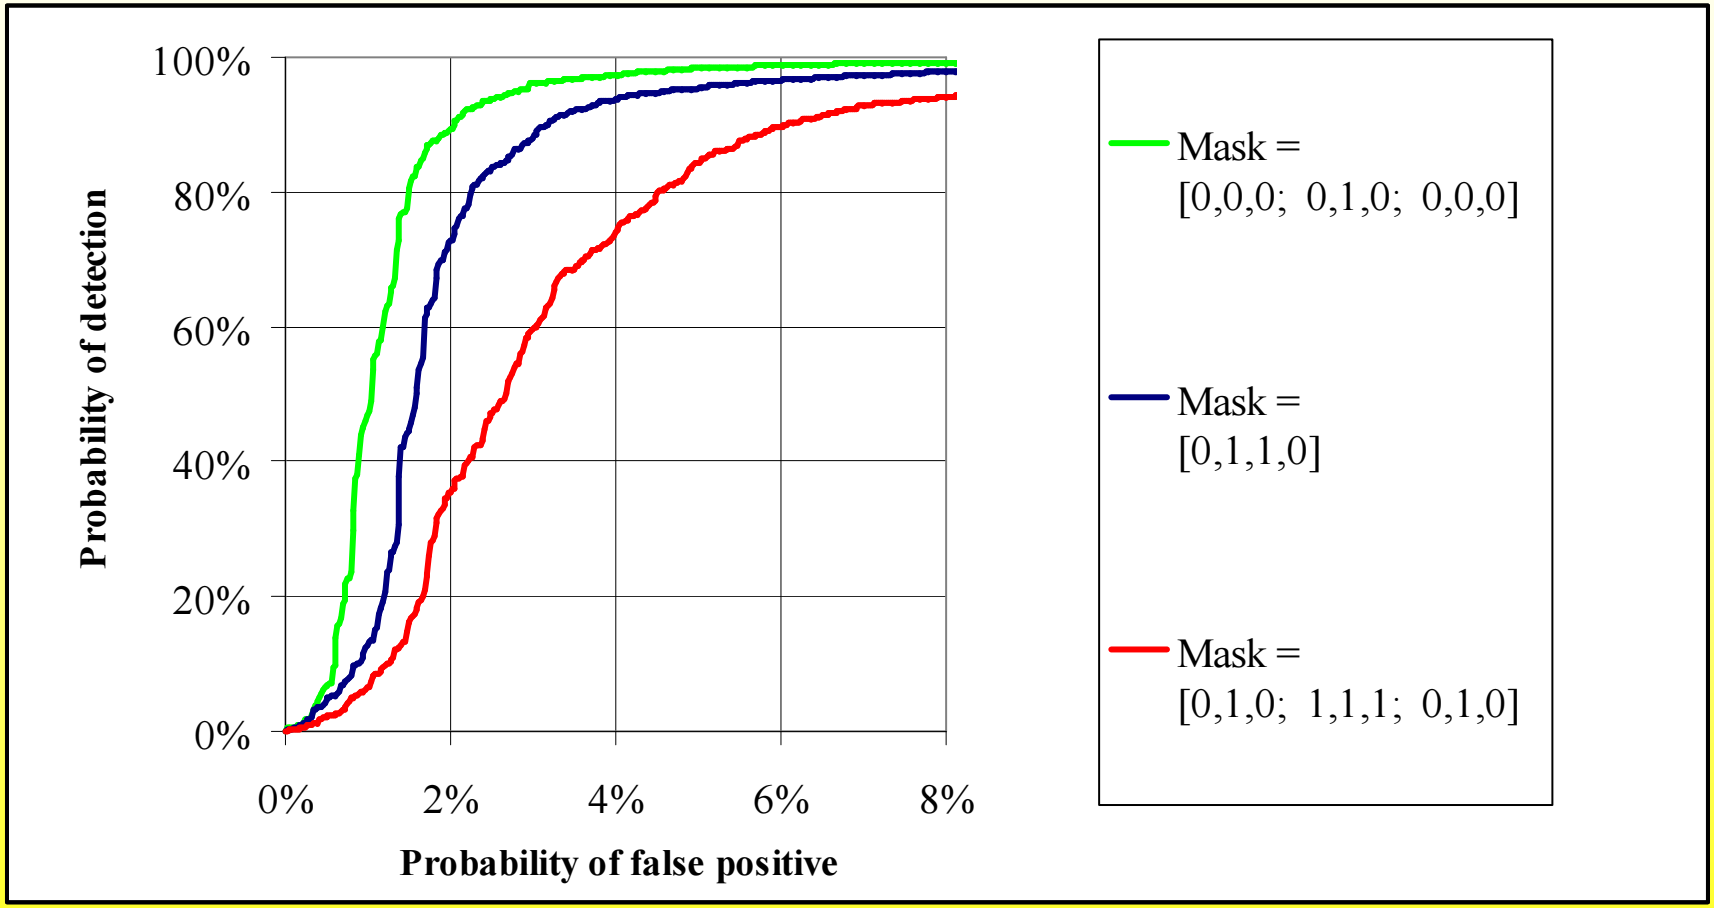

ROC curves generated from 5,000 JPEG images; 5% LSB Steganography was used

Improved Pairs Analysis

Pairs Analysis works by forming the *colour cuts* and then measuring relative homogeneity:

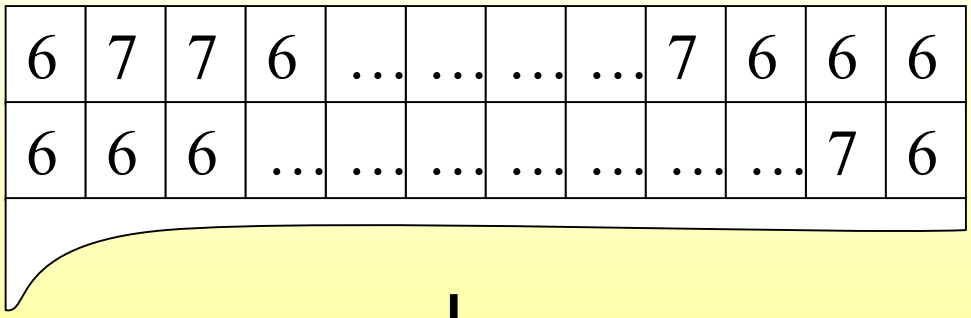

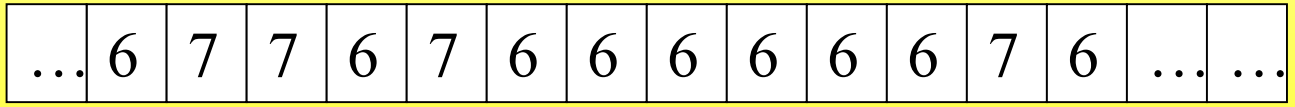

Improved Pairs Analysis

Pairs Analysis works by forming the *colour cuts* and then measuring relative homogeneity:

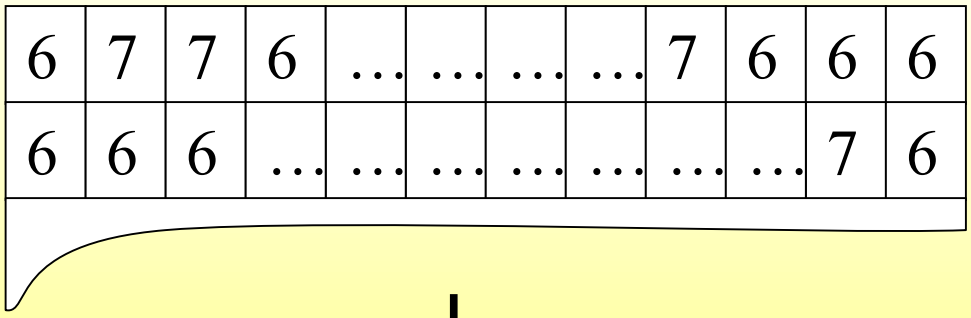

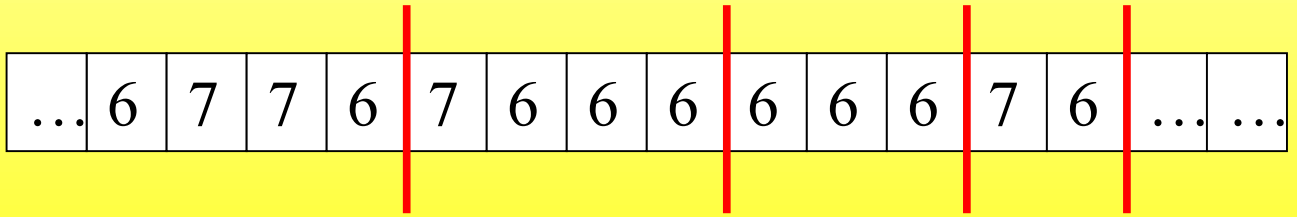

Improved Pairs Analysis

Count

E = # adjacent pixels of equal value

F = # adjacent pixels which differ by being "LSB flipped" (e.g. (6,7))

C = # adjacent pixels which differ by being "LSB contraflipped" (e.g. (7,8))

Let

$$Q=E/(E+F) - E/(E+C)$$

Then Q is quadratic in the length of LSB-embedded message, which can be solved for in the usual way [Fridrich *et al*, SPIE'02].

… c.f. [Dumitrescu *et al*, IHW'02]

Improved Pairs Analysis

Results in (very roughly) reduction of false positives by approximately half:

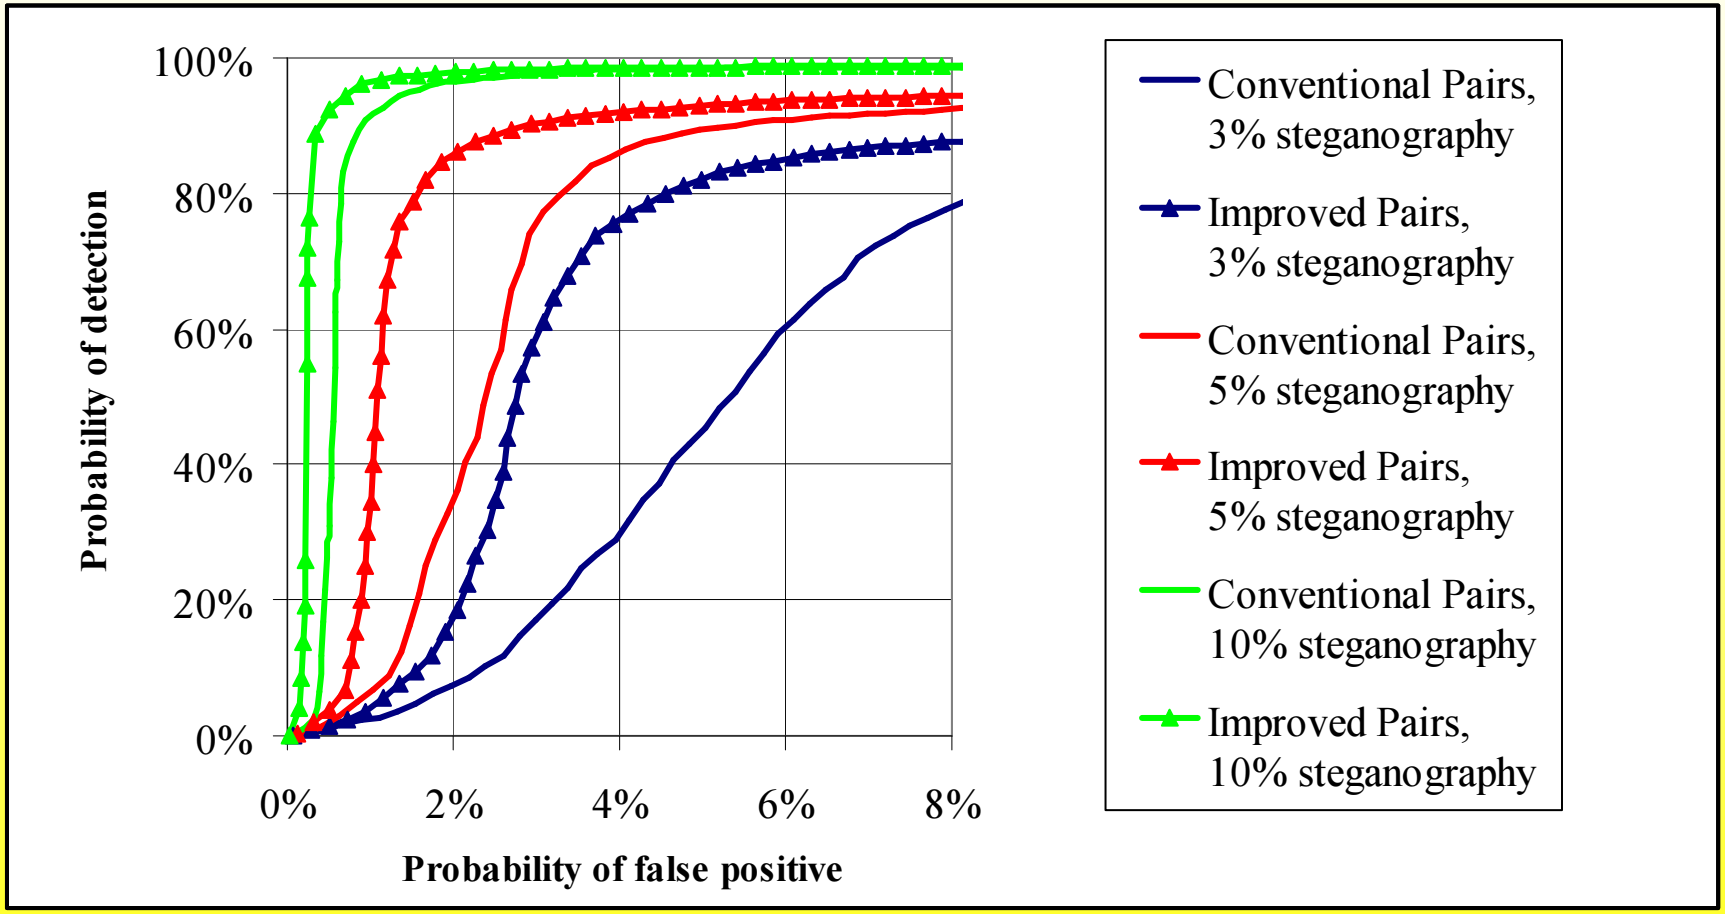

ROC curves generated from 15,000 JPEG images

Conclusions

The first results from the distributed steganalysis project focus only on LSB steganography in grayscale bitmaps. So far we have:

- Determined the best-performing “mask” for RS steganalysis,
- Substantially improved the performance of Pairs steganalysis

Conclusions

The first results from the distributed steganalysis project focus only on LSB steganography in grayscale bitmaps. So far we have:

- Determined the best-performing "mask" for RS steganalysis,
- Substantially improved the performance of Pairs steganalysis,
- Showed the null distribution of RS statistic is leptokurtic,
- Illustrated that 2-dimensional variants of Pairs and RS are no more useful than the standard versions,
- Exposed some pitfalls in the selection of a representative set of "natural images".

Further Work

Have only examined the tip of the iceberg! Still to do:

- Consider other steganalysis algorithms for LSB steganography,
- Look at LSB steganography in RGB, palette, JPEG images,
- Plenty of other methods of embedding besides LSB…

Further Work

Have only examined the tip of the iceberg! Still to do:

- Consider other steganalysis algorithms for LSB steganography,
- Look at LSB steganography in RGB, palette, JPEG images,
- Plenty of other methods of embedding besides LSB…

- What is it about JPEG compressed images that causes variation in reliability results?
- Can we classify images as a first stage, and then apply the best-performing steganalysis statistic for each class?
- Do some statistical analysis of accuracy of ROC curves generated by simulation.

End