

Combined Covert Data Embedding and Forensic Markings for Graphic Objects

Robert Ulichney, Hewlett-Packard Co., Andover, MA, USA

Stephen Pollard, Hewlett-Packard Co., Bristol, UK

Matthew Gaubatz, Hewlett-Packard Co., Seattle, WA, USA

Steven Simske, Hewlett-Packard Co., Ft. Collins, CO, USA

Abstract

We introduce a method for combining the covert encoding of data in hardcopy using a steganographic halftone with the extraction of a microscopic print signature from around the outside of the halftone suitable for forensic levels of authentication. Outline models for the extracting a unique forensic signature include the input image, the edge-refined reference halftone, and scans of sample stegatones. The results for recovery of several printed samples are reported. An example with a 3mm character can reliably carry 32 bits of data and carry a forensic signature with a false positive rate of 7×10^{-21} .

Introduction

Authentication of hard copy documents is important for a number of applications to protect against counterfeiting, product security, warrantee fraud and others. There are also a number of needs for embedding data in hard copy where overt marks such as bar codes would damage the aesthetics of a document. The novel method outlined in this paper simultaneously addresses both of these needs by combining forensics and steganographic halftoning on the same printed object, and describes a system for both encoding and decoding such objects.

Encoding System

The function of the Encoding System is the creation of a secure hardcopy document with an embedded payload along with filing its forensic fingerprint in a Registry. An overview of the Encoding System is illustrated in Figure 1.

Stegatone Generation

To covertly embed data in a printed object the method for creating a steganographic halftone or “stegatone” was reported earlier [1]. A stegatone generator takes a data payload and input image we call a “mule” because it is the vehicle that transports the payload when printed. In this paper, we are using dark grayscale images surrounded by white space to carry the payload. Glyphs are members of this class and an example that we use throughout for illustrative purposes is shown in Figure 2(a). At 600 dpi, this lower case 20-point “a” would appear 3 mm tall but our illustration shows it at 18x real size.

After some preprocessing, a reference halftone (Figure 2(b)), which is a standard clustered-dot halftone, is generated from the mule image. All halftone cells are classified in a reference map as either 0-bit, 1-bit, 2-bit, or 3-bit data carriers. These cells are depicted in Figure 2(c). 0-bit carriers, colored black, are called “Reference Cells” because they are unchanged and can be used for

alignment. The green cells for this example represent 2-bit carriers (as they can shift in one of four positions). Cells can be reference cells because they are too large to be shifted or too small to be detected; but more interestingly we can force cells to be reference cells if their unaltered shape is important to retain edge detail [2] or if they are needed to assist alignment. The payload is encoded by means of single pixel shift of the halftone clusters. The data carrying capacity of this example image is 234 bits. A 234-bit payload is encoded in the stegatone as shown in Figure 2(d). Along with producing this image, the stegatone generator outputs stegatone decode support data needed for decoding the stegatone, possible regeneration of the stegatone, and aid in finding print signatures. Stegatone decode support data includes the mule image, reference halftone, reference map, along with the shift rule. The stegatone is then printed to create secure hardcopy.

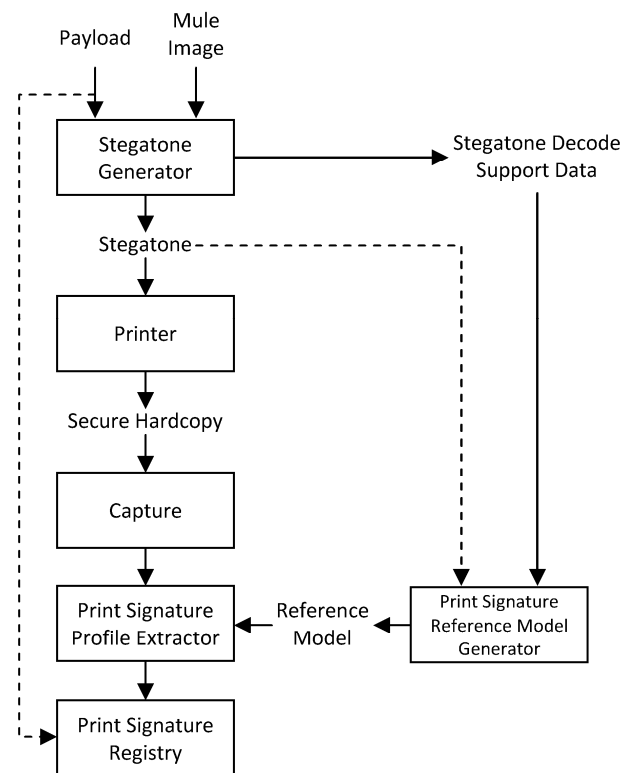




Figure 1. Forensic Stegatone Encode System.




(a) Mule image.




(b) Reference halftone.




(c) Reference map.




(d) Stegatone.



(e) Captured hardcopy.



(f) Standard outline model.



(g) Halftone outline model.

Figure 2. Example steps for a forensic stegatone (18x actual size).

Print Signature Profile Extraction

A unique forensic signature of the edges of this printed object can be recorded if the capture hardware is of sufficiently high resolution. The Secure Hardcopy is digitized in the Capture system shown in Figure 1. One device capable of this is the Dyson Relay CMOS Imaging Device (DrCID) [3] which was used to capture the image shown in Figure 2(e). A reference model of the outline surrounding the graphic image is the ideal against which the actual edge is measured. Deviations from this model will separate the truly random part of the outline in each individually printed glyph. It uses as input the stegatone decode support data mentioned above. When the mule image is used, a “standard” outline model is generated as described in [4]. This is depicted as the green outline in Figure 2(f). Alternatively forensic signatures can be obtained using more accurate outline models that more closely follows that of the printed stegatone. It is important to point out that the stegatone generated in Figure 2(d) employs an edge refinement process. This process retains a cleaner edge surrounding the graphic object and is consistent for all stegatones that it produces. Halftone cells near the edge that would be carrier cells instead retain their original edge detail and become fixed reference cells. So in this case the edges of the Reference Halftone (Figure 2(b)) can be used as the reference model for all resulting stegatones because those edges do not change with the payload. This halftone outline model is shown as the green line in Figure 2(g).

In both reference models shown in Figure 2(f) and Figure 2(g) the span of the Print Signature is depicted by the red and blue lines. The process for quantifying the print signature profile is detailed in [5]. The print signature is normalized in part by defining it with respect to the reference model, and then stored in a print signature registry. In practice the print signature can be reduced to a few hundred bits of data using a variance coding process in which the signature (which is typically 2000 to 4000 elements long) is broken into N segments (N of about 200 is optimal) whose variances are measured. Each segment is then coded with respect to the mean variance and rounded (see again [5] for details). The print signature is stored in a registry (possibly using the stegatone data as a key) so that it can subsequently be compared against a second signature recovered from the same document to prove its authenticity.

Recovery and Verification

Given the secure hardcopy document, hardware similar to what was used in the capture stage of the encode system is used to create a digitized version of the document. First the data is recovered from the stegatone [1]. Key to recovering the payload is precise alignment of the captured image. In this case we again use the outline model(s) used to extract the Print Signature to achieve the alignment. This in turn allows the boundaries of each halftone cell to be determined. For the purpose of illustration the halftone cell boundaries are shown overlaying the digital images in Figure 2(b), (d) and (e). The cell shifts are found by comparing the captured and registered stegatone to the reference halftone (Figure 2 (b)) and the payload is recovered. Because of imperfection in the print-capture process errors inevitably occur. For this reason error correction codes are used to increase recovery rates.

The generation of a print signature profile for authentication is the same as that performed during the encoding process. Finding the reference model only needs the mule input image if the standard model is used, or the reference halftone if the halftone model is used, both of which are part of the stegatone decode support data set. Once the reference model is established the print signature can be extracted and coded using the same variance coding method as previously. The difference between the newly extracted code and that stored in the registry is given by the sum of the absolute difference of their respective variance codes (essentially a modified hamming distance) which is called the shape distortion encoding distance or SDED [5].

Sample Results

To illustrate the operation of our forensic stegatone system 80 prints of the example stegatone from Figure 2 were printed on a HP M4345 laser printer and captured at high resolution. On average 86.3% of the 234 raw encoded stegatone bits were recovered. With error correction coding we can expect 100% recovery of 32 bits, a payload equal to many alternative large image watermark solutions.

To test and compare the performance of forensic matching we conducted a total of *four* separate experiments, each with a different outline models for generating the print signatures. In addition to the standard model based on the input image shape shown in Figure 2(f) and the model based on the reference halftone shown in Figure 2(g), two additional image models were generated from an average outline fitted to scans of the actual prints of the glyph.

The first image model used scans of the original set of 80 prints from the same printer. Again, since the quality of the forensic print signature increases when the truly random part of the print is separated from the repeatable part of the outline, the goal of the outline model is to be as accurate as possible. As expected the sampled printed image model performed better than the standard model and the halftone model. However, in practice all prints form a forensic stegatone system would not come from the same printer, so the performance of the same-printer-image model is unnaturally high. A more realistic test is to generate the model from sample prints made on some other printer or printers. The “other printer” model still performed better than the reference halftone model but not as good as the image model from the same printer.

In each of our four experiments 80 print signatures were generated as the deviation of the signal around our glyph from the outline model, and stored in the print signature registry. Each of the prints were then captured a second time and the print signature form that capture compared with the 80 pre-stored signatures arriving at 80 SDEDs (shape distortion encoding distances). Figure 3 shows a plot of these 6400 (80x80) SDEDs for the experiment that used the other printer image model. For each of the 80 samples the red points indicate the SDED for comparison to that same print’s stored signature. The blue prints represent the SDEDs for comparison with the 79 other print signatures.

Figure 4 is a histogram of the SDED values for all 80*79 comparisons with other print signatures for the data in Figure 3. This approximate Gaussian shape is typical for all models used. Assuming a Gaussian distribution we can measure a average z-

score as the number of standard deviations the average matching SDED is from this distribution. For this example using the image model from another printer, that z-score is 9.3. This means that the probability of a false positive will effectively never happen – the probability associated with that z-score is 7×10^{-21} .




Figure 3. SDED values for 80 prints using an Image-based model produced on different printer. Comparisons of the same print in red, comparisons with all different prints in blue.




Figure 4. Histogram of SDED values for the 80*79 non-matches in Figure 3

However, false negatives can occur due to noisy or blurred captures. Three such false negatives are evident in Figure 3 where the red points are found in the non-match distribution. Upon examination we found the captures associated with those print signatures were indeed of low quality and subsequent repeat captures resulted in the SDED values indicated by the circles on that plot

The false positive z-scores for all four of our outline models are plotted in Figure 5. The span values indicated describe the

range on either side of the outline model over which the print signature is sampled. (The distance between the red and blue lines in Figure 2(f) and Figure 2(g).) They are expressed as a percent of the average linear size of the glyph. The case detailed in Figure 3 and Figure 4 is represented as the “Image Model Other Printer” with a span of 6%.




Figure 5. False positive z-scores for the four outline models tested.

In most forensic authentication applications it is acceptable to have false negatives that can be corrected with a better capture, but it is not acceptable to allow false positives. Even the worst case shown in Figure 5 using the Standard Model with a span of 10% has a z-score of 6.7 indicating the probability of a false positive is 1 out of 100 billion. Our example demonstrates high forensic reliability along with covert embedding of a 32 bit payload in a single 3 mm character.

References

- [1] R. Ulichney, M. Gaubatz, and S. Simske, “Encoding Information in Clustered-Dot Halftones”, IS&T NIP26 (26th Int. Conf. on Digital Printing Technologies), Austin, TX, 602-605, Sep 2010.
- [2] R. Ulichney, M. Gaubatz, “Tracing the Source of Printed Documents with Edge-Refined Stegatores”, IS&T NIP27 (27th Int. Conf. on Digital Printing Technologies), Oct 2011.
- [3] G. Adams, “Hand held Dyson Relay lens for anti-counterfeiting”, IEEE IST, 2010.
- [4] S. Pollard, S. Simske, G. Adams, “Model based print signature profile extraction for forensic analysis of individual text glyphs”, IEEE WIFS, 2010.
- [5] G. Adams, S. Pollard, S. Simske, “An imaging system for simultaneous inspection, authentication and forensics”, IEEE IST, 2010.

Author Biography

Robert Ulichney is a Distinguished Technologist with HP Labs. He received a Ph.D. from MIT in electrical engineering and computer science. Before joining HP he was with Digital Equipment Corp for several years then with Compaq’s Cambridge Research Lab where he led a number of research projects on image and video implementations for both hard copy and display products.