

Tracing the Source of Printed Documents with Edge-Refined Stegatones

Robert Ulichney, Hewlett-Packard Co., Marlboro, MA
Matthew Gaubatz, Hewlett-Packard Co., Seattle, WA

Abstract

A system for tracking and tracing secure documents is described by embedding print job data on each page in a steganographic halftone, preserving the aesthetics of the page. The system demonstrates the ability to transmit job tracking data with perfect accuracy using data hidden in a high-quality rendering of a graphics object. This work addresses two key issues: (1) control of halftoning edge effects for graphic objects like logos, and (2) compensation for non-affine warping effects induced by the print-scan process. The solution eliminates noisy edges without sacrificing data-carrying capacity in natural images, while providing automated alignment and robust recovery with the use of error correction codes.

The Need for Tracing Documents

In environments where sensitive or confidential documents are handled, such as in a financial institution, it is often desirable to have the ability to trace who printed a document as well as when and where it was printed. The desired workflow of a sensitive document is as follows. A user, who is authenticated to gain access to both the digital document and a printer on a secure network, initiates a print job. The print driver generates a log of the print job that includes the user who initiated the job, the date and time of that job, and what printer was used. This step is already common practice in many enterprise networks. The next step in the desired workflow is embedding this print job data inconspicuously on every page as it is printed. When the sensitive document is left unattended in an insecure location, or in the possession of someone not authorized to have it and is then discovered, the means to trace that document is needed. With that means, operations audit staff can recover the print job log directly from the hard copy to determine who printed the document.

There are many solutions for embedding data in hard copy documents that are in use for labels and packing that involve placing overt marks on the page. The most common are 1D, 2D, and 3D (color) barcodes. A review of halftoning-based approaches is presented in [1]. While the associated data densities can be high, these markings can significantly disturb the aesthetics of office documents and thus would not be feasible in some enterprise applications. Another candidate solution could involve the use of special inks, but this method would add considerable cost to every printer. In our solution, data embedding is achieved through the use of steganographic halftones, or steganones. Stegatones [1] are clustered-dot halftones that encode a data payload by single-pixel shifts of selected clusters. Unlike barcodes, a steganone does not disturb the aesthetics of the page composition, and yet can represent a large number of bits.

Embedding Job Data with Stegatones

The means for achieving this automatic data embedding on every page can be enabled by enhancing the printer driver as illustrated in Figure 1. A user initiates a Print Job, which is processed by a Print Driver that normally delivers the page content to the printer and generates a Print Job Log. In the enhanced driver, this Job Data is processed by the Payload Generator which prepares data for embedding in the steganone by arranging it into a more compact form.

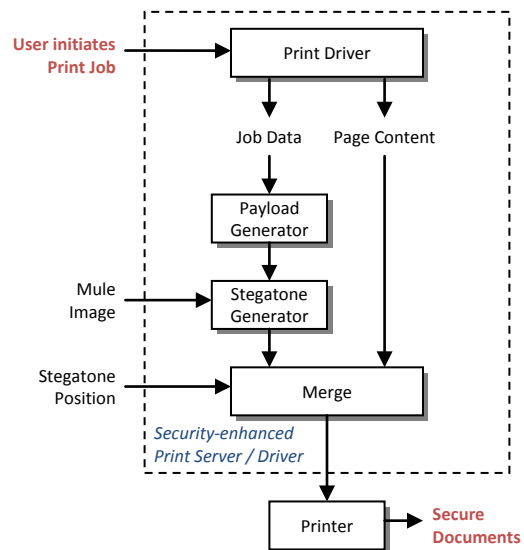


Figure 1. Print Job Embedding System.

This Payload is then delivered to the Stegatone Generator, the operation of which is described in detail in [1]. The data is carried in a pre-selected input image or “Mule”; this image appears in every page, ideally in a pre-specified location. One candidate Mule could be a company logo such as the example seen in Figure 2. The resulting steganone is merged with the page content and sent to the printer. The Merge operation requires the coordinates for placement of the steganone on the page. The resulting output pages appear normal but are embedded with data from the Print Job Log.

Bank of Stegatonica

Figure 2. Example input logo at 600 dpi, with gray text on a white background.

The Print Job Data Recovery System is shown in Figure 3. The secure documents are first scanned using information that specifies the position of the stegatone on the page so that only that region is captured. The scanned data is delivered to the Stegatone Recovery System, which also takes the Mule image used to generate the stegatone as input. The output from the Stegatone Recovery System is the embedded Payload. This information is sent to the Payload-to-Job-Data Generator to produce the desired Job Data.

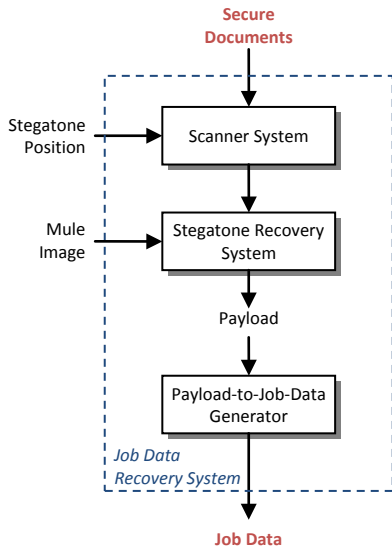


Figure 3. Print Job Data Recovery System.

Edge Refinement for Steganographic Halftones

While stegatones were designed to be applied with any input image as a mule for carrying data, a problem surfaced when using them with graphic elements such as logos for this application. An enlarged portion of the reference halftone constructed from the logo in Figure 2 (displayed at actual size) is illustrated in Figure 4 and resulting stegatone carrying the job data is shown in Figure 5. Note the satellite clusters around the edges of the logo in the stegatone contribute to a fuzzy or noisy appearance. The reason these artifacts appear is due to the fact that halftone cells identified as data carriers are flattened via averaging to produce a predictable shape in the resulting clustered dot. Carrier cells must thus sacrifice some rendition detail in order to store recoverable data. For photo-acquired image mules, this loss of detail is negligible, but can be quite pronounced for synthetic graphic content if the percentage of halftone cells along potentially problematic edge regions is high enough, as in this case.

The encoding system [1] uses the input image values associated with the 4x4 cell to classify it as either a reference cell or a carrier cell. The solution to the noisy edge problem is achieved by adding a condition to this classification procedure to exclude sharp edges. If a cell that would be classified as a carrier contains an edge then it is re-classified as a reference cell and is not flattened. The edge information in that that cell is rendered with a “partial-dot” (a term used to describe an edge-bearing cluster in halftoning) to properly reveal the edge detail.



Figure 4. Portion of a clustered-dot halftone using the classical screen of the logo from Figure 2 shown at 5.6x real size.



Figure 5. Portion of a Stegatone of the logo from Figure 2 revealing noisy edges.

A number of edge detecting criteria, such as the Laplacian operator, are candidates techniques for perform this classification step. In this case, a more complex classifier does not necessarily create a higher quality result, and in fact, a simple range test performs very well. Therefore, cells are allowed to be carriers only if the dynamic range associated with the input cell pixels is below a threshold. Figure 6 shows the result of suppressing carriers with a minimum range of 0.25, where the difference between white and black is 1.0.



Figure 6. Stegatone created using an additional edge condition based on range, with the threshold range set to 0.25.

Notice that almost all of the satellite clusters seen in Figure 5 are gone with the exception of a few small dots. The cells responsible for these cases are largely white pixels with only a few anti-aliasing transition pixels from the edge of the gray text. Even these degradations are removed by lowering the range to 0.06 as is shown in Figure 7. The elimination of carriers around the edges has a cost in terms of reduced data carrying capacity. Another concern is for hybrid mule images that contain both photographically-acquired and graphics pixels. By way of

example, consider the example logo in Figure 2, and the 600x600 photo shown in Figure 8.



Figure 7. Stegatone created using an additional edge condition based on range, with the threshold range set to 0.06. All unwanted satellite clusters are gone.



Figure 8. Example 600x600 pixel image.

Table 1 shows the effect that the range-based edge criterion has on the number of carrier bits for these two mule images. It expresses data carrying capacity as a function of the Range Threshold in terms of bit count, as well as a percentage of the capacity achieved when edge refinement is not used.

Table 1. Relative loss of data capacity, expressed as total bits and as a percent of the “No edge refinement” capacity for two image types, achieved as a function of Range Threshold, T, using the rule “range ≥ T”.

Range Threshold	Edge Noise	Data Carrying Capacity (bits)			
		logo		photo	
No edge refinement	Severe	4401		18961	
0.25	Improved	3717	84%	17078	90%
0.13	Improved	3655	83%	14166	75%
0.06	Absent	3584	81%	11419	60%

This particular solution improves the aesthetics of the results for graphic content at the expense of a reduction in capacity that is less than 20%. One problem with this solution, however, is that data carrying capacity loss becomes more significant for photo-acquired images, which do not benefit from any noticeable improvement in visual quality due to edge preservation. The reduction in data carrying capacity is most severe when a range threshold of 0.06 is used; the capacity is reduced to 60% of its original value.

This issue is resolved by exploiting the fact that graphics images, such as text, have foreground (or background) pixels that are all the same value. Even for flat regions in a photo, sensor noise usually guarantees some small variation. Thus, a second

criterion that can be cascaded with the range check is a test for solidity that compares the frequency of the mode pixel value against a minimum threshold. This step will eliminate almost all photo-based cells that pass the range test. To continue our example with the test images of Figure 2 and Figure 8, with a frequency threshold of 5 pixels, there is only a small decrease in data carrying capacity (see Table 2). There is no difference in capacity for the logo image, as expected, but a vast improvement for the example photo.

Table 2. Relative loss of data capacity, expressed as total bits and as a percent of the “No edge refinement” capacity for two image types as a function of Range Threshold, with the additional refinement condition “frequency of mode ≥ 5”.

Range Threshold	Edge Noise	Data Carrying Capacity (bits)			
		logo		photo	
No edge refinement	Severe	4401		18961	
0.25	Improved	3717	84%	18961	100%
0.13	Improved	3655	83%	18931	99%
0.06	Absent	3584	81%	18826	99%

Adding edge conditions to the Reference Map Generator to check for range and mode frequency thresholds thus vastly improves the quality of graphics with minimal impact on data carrying capacity.

Improving Robustness for Recovery

Recovery of the payload data involves two main phases: an alignment procedure, which registers a scanned stegatone with respect to a reference halftone, and a decoding procedure, which maps regions of pixels in the aligned stegatone to shifted versions of the corresponding reference halftone cell. The alignment process, at the very least, requires estimation of an affine transform (linear perspective change, including asymmetric scale, rotation, and keystone) that best matches the scanned stegatone to the reference halftone. Various geometric distortions induced into a scanned stegatone via a print-scan cycle can cause alignment to be a rather difficult problem. This issue is complicated by the fact that sub-pixel alignment errors can render stegatoness unreadable.

One of the issues that can defeat certain global alignment strategies is the variety in pixel pitch that occurs across a horizontal swath of the page in the print-scan process. Because of this effect, even if scale and rotation parameters are estimated correctly, the pixels within the scanned image may be up to several pixels off from those in the reference halftone. Figure 9 illustrates this behavior for one sample printer and scanner. A magenta stegatone is globally aligned with a cyan reference halftone.

Bank of Stegatonion

Figure 9. Sample non-affine warping induced by the Print-Scan Process. The figure contains a globally-aligned reference halftone (cyan) and stegatone (magenta). Even though the end letters are properly aligned, the letters in the middle of the image are not.

Note that though both ends of the stegatone are properly aligned, the letters in the middle of the logo are offset from those in the reference halftone by several pixels.

As a result, a tiered approach to alignment is used that determines approximate scale and rotation parameters via a global estimation phase, then performs a local alignment procedure on vertical strips of the stegatone. The local alignment scheme is an iterative process that assesses the quality of different candidate alignment parameters, and chooses the best one by optimizing for some functional measurement. A number of measurements can be used for this purpose, such as per-pixel correlation or squared error between the reference halftone and the scanned image. It is important to note that using this approach requires that results from the second recovery phase be fed back into the alignment system (see Figure 10) prior to invoking the second phase for a final decode operation.

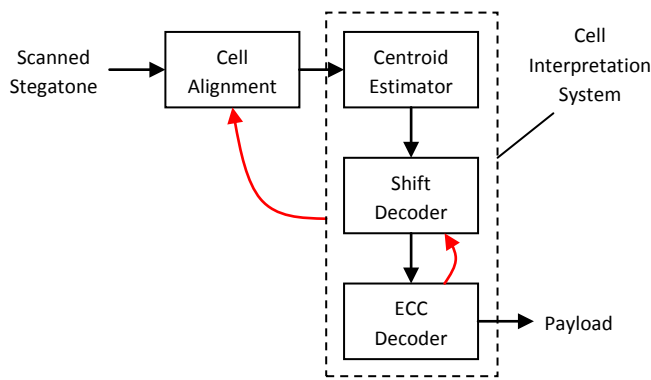


Figure 10. Simplified summary of Stegatone Recovery System. Feedback loops are depicted in red.

Because perfect alignment is difficult to achieve in practice, reading markings on printed media is not an error-free process. There are a number of types of error correction codes (ECC) available that can help make the process more robust. Reed-Solomon [2] and convolutional and repeat accumulate codes [3] are two examples that have been applied for conveying data through print scan channels. ECC codes, however, are most efficient when the block length becomes large. In the application at hand, block lengths may not be very long, simply because the payload capacities might be limited by size. Therefore, it is important to qualify how different effects induced by the print-scan process [4] affect data integrity. The information-theoretic channel capacity of the stegatone cells is thus of interest. In the ideal case, each stegatone cell represents two bits of data, but in practice, the (average) number of bits that can be robustly transmitted by each cell is lowered by print and scan distortions.

Performance Testing

The efficacy of different error coding schemes was analyzed for the aforementioned choice of printer and scanner using the following procedure. Data associated with a typical XML record generated by the enterprise print driver, 112 bytes in size, was transmitted in a logo-based Stegatone with a data carrying capacity of 448 bytes using the image in Figure 2. It is well-known that

turbo codes offer some of the most competitive error coding schemes when used to protect large amounts of data. Since the number of bits used for transmission is relatively low, however, and asymptotic trends may not apply, a number of different alternatives are investigated. The tested methods used to improve robustness of transmission include (1) repetition codes, (2) Hadamard codes, (3) a combination of turbo and Hadamard codes, (4) a combination of turbo and Reed-Solomon codes and (5) the placement of regularly spaced cells with pre-defined shifts to create patterns that are easier to locate during the alignment process. To isolate the effects of the current alignment strategy, two versions of each Stegatone were decoded: one version aligned by the proposed approach, and another aligned using fiducials printed around the logo to help locate each cell stegatone accurately.

Table 3. Decoding performance achieved with a sample compressed Print Log data in conjunction with five different redundancy schemes. The best results in each column are highlighted in bold.

Protection scheme	Fiducial-assisted alignment			
	Channel capacity (bits/cell)	Raw BER	Final BER	# byte errors
Repetition code	1.78	0.0170	0.0045	3
Hadamard code	1.84	0.0140	0	0
Turbo+Hadamard	1.74	0.0229	0	0
Turbo+Reed-Solomon	1.73	0.0243	0	0
Alignment markers	1.86	0.0092	0.0092	10

Protection scheme	Blind alignment (without using fiducials)			
	Channel capacity (bits/cell)	Raw BER	Final BER	# byte errors
Repetition code	1.51	0.0550	0.0156	8
Hadamard code	1.40	0.0843	0.0134	7
Turbo+Hadamard	1.38	0.0893	0.0190	9
Turbo+Reed-Solomon	1.40	0.0776	0	0
Alignment markers	1.60	0.0419	0.0419	49

Results are compiled in Table 3, and are reported for images printed at 600 dpi on an HP LJ1200, and scanned at 1200dpi on an HP G4050. The data include the information-theoretic capacity of the Stegatone print-scan channel given the alignment scheme, the raw and final bit error rates (BERs), i.e., the fraction of incorrectly decoded bits before and after the application of any error

correction scheme, and the final number of byte errors in the decoded payload. In several cases, with fiducial-assisted alignment, perfect transmission of the payload is achieved. Since a small percentage of bit errors can lead to a larger percentage of byte errors (if the bits are evenly distributed), the combination of turbo and Reed-Solomon coding offers a good trade-off between the ability to recover from both randomly distributed bit errors as well as bursts of bit errors. This particular protection scheme results in perfect transmission even when fiducial markings are not used for alignment.

The amount of redundant information required will depend on the target application. In instances where borders or other structured data are present near a stegatone signal, fiducial-assisted alignment is a definite possibility. With respect to robustness, the advantage of using fiducials can be significant: in this test, the information-theoretic capacity of the stegatone print-scan channel decreases by almost twenty percent on average when fiducial markings are not used. It is worth noting that while using extra bits to create cells with pre-specified shifts results in the least robust overall protection scheme, doing so does create a channel with the highest capacity. These observations highlight the mutual importance of a good alignment strategy and the application of error correction coding.

We have demonstrated the ability to achieve 100% data recovery from logos via the application of ECC. Stegatone generation now creates noise-free graphics with no loss of data capacity or quality for photos. The system is well suited for our application of tracing the source of printed documents by covertly embedding data.

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] V. Guruswami and A. Rudra, "Limits to List Decoding Reed-Solomon Codes," IEEE Transactions on Information Theory, vol. 52, pp. 3642-3649, Aug. 2006.
- [3] O. Bulan, G. Sharma, and V. Monga, "Orientation Modulation for Data Hiding in Clustered-Dot Halftone Prints," IEEE Transactions on Image Processing, vol. 19, no. 8, pp. 2070-2084, Aug. 2010.
- [4] K. Solanki, U. Madhow, B. S. Manjanuath, S. Chandrasekaran and I. El-Khalil, "'Print and Scan' Resilient Data Hiding in Images," IEEE Transaction on Information Forensics and Security, vol. 1, pp. 464-478, Dec. 2006

Author Biography

Robert Ulichney received a Ph.D. from MIT in electrical engineering and computer science. He was with Digital Equipment Corp. for several years where he focused on image and video implementations for both hard copy and display products. From 1998-2002 he was with Compaq's Cambridge Research Lab where he led a number of research efforts in video and image processing. Bob is now a Distinguished Technologist with HP Labs.