

Entropy-Based Dark Frame Subtraction

Michael Goesele¹, Wolfgang Heidrich², Hans-Peter Seidel¹

¹Max-Planck-Institut für Informatik, Saarbrücken, Germany

²The University of British Columbia, Vancouver, Canada

Abstract

Noise due to dark current is a serious limitation for taking long exposure time images with a CCD digital camera. Current solutions have serious drawbacks: interpolation of pixels with high dark current leads to smoothing effects or other artifacts – especially if a large number of pixels are corrupted. Due to the exponential temperature dependence of the dark current, dark frame subtraction works best for temperature controlled high end CCD imaging systems.

On the physical level, two independent signals (charge generated by photons hitting the CCD and by the dark current) are added. Due to its random distribution, adding (or subtracting) the dark current noise signal increases the entropy of the resulting image. The entropy is minimal if the dark current signal is not present at all.

A dark frame is a good representation of the dark current noise. As the generated dark current depends on the temperature equally for all pixels, a noisy image can be cleaned by the subtraction of a scaled dark frame. The scaling factor can be determined in an optimization step which tries to minimize the entropy of the cleaned image.

We implemented a software system that effectively removes dark current noise even from highly corrupted images. The resulting images contain almost no visible artifacts since only the noise signal is removed. This extends the range of usable exposure times of digital cameras without temperature control systems by about one to two orders of magnitude.

1. Introduction

The ability of modern CCD digital cameras to take images with long exposure times is seriously limited by dark current noise. Images with exposure times of several seconds are often highly corrupted. Only a small portion of current cameras is equipped with a cooling system to reduce the amount of dark current noise.

Given no additional information about the properties of this noise, the only possibility to remove it is to apply general noise detection and removal techniques such as median filtering [4]. However these techniques only detect that a pixel is noisy and try to infer the true pixel value from the values of neighboring pixels.

As each sensor element on a CCD chip generates a characteristic amount of dark current [7], it is possible to capture this information in a separate calibration step. It can be used afterwards to subtract the dark current noise from noisy images (dark frame subtraction). This technique has the advantage of reconstructing the true pixel value of a noisy pixel instead of inferring it from neighboring values.

The amount of dark current generated by a particular sensor element is approximately constant under fixed external conditions. Unfortunately, it depends exponentially on the temperature of the CCD chip and doubles for a temperature increase of about 5 Kelvin [1]. This exponential behavior makes dark frame subtraction hard to use unless the temperature can be controlled very strictly. If the temperature can only be measured but not controlled, scaling the entire dark frame with a correction factor k before subtraction works well in principle. The correct k has to be derived from the temperature difference between the two images. But due to the exponential behavior it is difficult to set k correctly.

In this paper we describe a technique to find the correct factor k based on the entropy of the resulting image. In contrast to techniques that try to maximize the entropy e.g. to remove blur [7] we try to minimize the entropy via dark frame subtraction.

In Section 3 we describe the theoretical foundation for this technique, Section 4 shows its relation to compression algorithms, and Section 5 discusses some implementation issues. In Section 6 we show some results before we finish with our conclusions in Section 7.

2. Physical Foundation

When light hits a sensor element of a CCD chip an amount of charge proportional to the incoming light is generated and stored. The most important other source of charge is the so called dark current – the amount of charge that is generated even in absence of light [6]. Apart from random fluctuations, each sensor element generates a characteristic amount of dark current which depends exponentially on the temperature. In addition there are other noise sources in a CCD imaging system which are not discussed in this

paper [6].

For the sensors on a CCD chip we use the following model that was derived from our experiments: for a sensor element j the total amount of charge Q_j collected during an exposure with exposure time T can be written as

$$Q_j = Q_{light,j} + kQ_{noise,j} + Q_{other,j} \quad (1)$$

$$= \int_T I_{light,j}(t) dt + \int_T kI_{noise,j} dt + Q_{other,j} \quad (2)$$

$$= \int_T I_{light,j}(t) dt + kTI_{noise,j} + Q_{other,j} \quad (3)$$

with a single temperature-dependent constant k for all sensor elements. This corresponds to a physical model of a sensor element where all charge is generated by current sources and the current does not depend on the amount of charge already stored in the sensor element. The dark current $I_{noise,j}$ is therefore assumed to be constant over time. Furthermore we assume that $Q_{other,j}$ is either much smaller than $kQ_{noise,j}$ or compensated by other techniques and can therefore be neglected.

The constant $Q_{noise,j}$ scaled by an unknown temperature constant k_u can then be determined for a given exposure time T by taking an image with no light hitting the sensor. The accuracy can be improved by averaging several images which were taken under the same conditions to remove random fluctuations. Such an image containing only the dark current noise is called a *noise image*, an image of a scene that also contains dark current noise is called a *target image*. A *cleaned image* is a target image for which some dark frame subtraction has been performed.

As the amount of charge due to dark current depends on two variables – the exposure time and the temperature – a database of noise images containing the appropriate image for each combination of exposure time and temperature would be very large and impractical even if the temperature could be controlled or measured exactly. The goal of our algorithm is therefore to use a single noise image generated under roughly the same conditions as the target image and than to find a suitable k that removes the contribution of $Q_{noise,j}$ as accurately as possible.

2.1. Analog-Digital Conversion

When an image is taken the amount of charge collected on each sensor element is converted into a digital value leading to an image P with pixel values P_j . In the following we assume that the pixel values P_j are proportional to the amount of charge Q_j collected during the exposure. Depending on the properties of the actual camera system used this may require additional processing steps e.g. to correct for the internal gamma factor setting of the camera.

3. Entropy

According to Shannon’s information theory [8], the entropy H of some data taken from an alphabet with n characters with probabilities p_0, \dots, p_{n-1} can be expressed as

$$H = - \sum_{i=0}^{n-1} p_i \log p_i. \quad (4)$$

The entropy is a measure for the information content of the data. It becomes maximal if all symbols occur with the same probability. It is minimal if only one symbol occurs in the data.

A digital image P is a (two-dimensional) array of pixel values P_j . The set of possible pixel values (or alternatively the set of values with a probability larger than zero, i.e. the set of all values that occur at least once in the image) forms an alphabet which can be used to compute the entropy of the image as described above.

The entropy of an image is not an ideal measure for its information content. It depends only on the probability of the elements of the alphabet and totally disregards the spatial distribution of the pixel values in the corresponding image. Therefore the image of a gray ramp can have the same entropy as random noise as long as the values have the same probability distribution. Furthermore the entropy is only a meaningful measure for the information content of a real image if the number of used elements of the alphabet is much smaller than the number of pixels. As images are usually digitized and stored with a very limited precision (e.g. 8 or 12 bits per value) this is normally true. In spite of these facts we use the entropy because it helps to understand the principle of our technique and only leads to problems for special cases. In Section 4 we describe other implementations that use different measures for the information content of an image.

3.1. Image Properties

The entropy of an image depends heavily on its contents. A noiseless or almost noiseless image (e.g. an image captured with a digital camera at a short exposure time) has – except for some special cases – a rather low entropy. In contrast to that the entropy of a noise image is very high due to the random properties of the dark current noise.

For the same reason a CCD image with long exposure time containing a considerable amount of noise has a much higher entropy than the same image after dark frame subtraction with a correct factor k . Furthermore, the entropy increases again if the fixed pattern noise is overcompensated (k is too large). This leads to the addition of a “negative dark current noise” which has the same random properties and is therefore again a high entropy signal.




Figure 1: Curves for test image A. The image shows the compressed file size in bits as a function of k with $k \in [0, 1]$ using Huffman encoding. The values were clamped to 0 for the lower curve. Negative values were allowed for the upper curve. The optimum is located at $k = 0.46$.

Dark frame subtraction with an optimal correction factor k therefore minimizes the entropy of the corrected image.

3.2. Optimization

This property can be used in an optimization step to find a good correction factor k : given a target image P with dark current noise and a noise image P_{noise} , a series of images P_k is computed where the value of a pixel j is given as

$$P_{k,j} = P_j - kP_{noise,j}. \quad (5)$$

For each image, the entropy H_k is computed as a quality measure. The correction factor k , that minimizes the entropy of the corrected image, is selected as the optimal solution. Pixels for which clipping occurred due to saturated sensor elements during analog-digital conversion must be excluded from the entropy computation.

The optimization can return a wrong result if negative values, which can occur due to overcompensation or random noise, are not included as well. Figure 1 shows the result of two optimization runs where negative values were once clipped and once allowed. The same optimal correction factor $k = 0.46$ was nevertheless found in both cases. The size of the Huffman encoding of the images was used as quality measure, which is described in Section 4.

3.3. Practical Considerations

Due to the limited precision of image data, it is not recommended to use an arbitrary correction factor. Instead, an additional noise image should be acquired under conditions that match those of the target image more closely if k is too small or too large.

The entropy of a cleaned CCD image depends on other factors as well – especially on the characteristics of the analog-digital conversion. In some cases, the actual precision (bits per pixel value) of the conversion process may be lower than the requested resolution of the image data format. The optimization process then introduces new pixel values. This extension of the alphabet may lead to a drastic increase of the entropy. In this case, the image values $P_{k,j}$ should be rounded to the same precision as the original data.

Depending on the amount of noise in an image it is often sufficient to use only a small region of the image for the computation of k which can reduce the cost of this approach significantly.

4. Compression Algorithms

The ability of a compression algorithm to compress data is closely related with the information content of the data. The Huffman compression algorithm [5], which is directly based on the entropy of the data, assigns short code words to values with high probability and long code words to values with low probability. The compressed data has therefore a very high entropy. So instead of computing the entropy, it is possible to use the size of the Huffman compressed images as a quality measure. It is minimal for an optimal k .

The compression efficiency of other generic compression algorithms like gzip [3] generally also depends on the information content of the compressed data, which makes them also suitable for this technique. Details depend on the specific algorithm. One advantage of some of these algorithms is that they consider not only the probability of single values but also that of repeating patterns or constant data regions.

4.1. Image Compression Algorithms

This is especially true for image compression algorithms which exploit the two-dimensional nature of the images. Some of them even take into account specific properties of natural images such as constant regions. They are often implemented in specific file formats.

Figure 2 shows the resulting file sizes for an example image. The curves of the gzip-compressed raw data and the lossless image compression techniques show an identical behavior and lead to the same correction factor as the entropy based Huffman code (see Figure 1).

The two examples of lossy JPEG compressions show also a significant change in the behaviour of the curves for an optimal k , which is more noticeable for the one with high quality setting. Although the file size is not minimal for an optimal k , it is marked by a change of the gradi-




Figure 2: Curves for test image A. The image shows the compressed file size in byte for a series of k with $k \in [0, 2]$ and for several compression methods (left border top to bottom: RAW image with gzip, JPEG quality 100%, TIFF with LZW compression, PNG with quality 100%, JPEG quality 75%). The optimum is located at $k = 0.46$.

ent. Using the size of a JPEG compressed image as quality measure is especially interesting as many cameras support JPEG compression in hardware.

5. Implementation Issues

In theory our technique requires only a single noise image. Due to the limited precision of real image data, other noise sources, and nonlinearities, a small number of additional noise images should be used to arrive at a reasonable k .

In some of our test cases with different exposure times of the target image and the dark image, the minimum criteria was not sufficient to select an optimal k . However, other characteristics of the curve (e.g. a change in the gradient as in the JPEG examples in Figure 2) can be used to determine the optimal k . Alternatively, a target image can always be corrected using a noise image with the same exposure time.

The dark current properties of a CCD chip change only slowly over time, allowing the use of a single set of noise images for the calibration during an extended period of time.

High resolution CCD cameras often require a host computer to generate the final images. Our dark frame subtraction technique can be easily implemented on the host computer. Alternatively the noise images can be stored on the camera itself. The camera's signal processing hardware can be used for histogram calculation (to compute the probability of the values in the image) or for image compression.

5.1. Additional Image Processing

Once an optimal k is found a corrected version of the original image can be computed. Here again, pixels that were overexposed in the original image or in the noise image have to be treated separately. If the goal is to get a faithful representation (i.e. for scientific purposes) they can be marked as invalid. Otherwise linear interpolation or another technique can be used to reconstruct the missing values.

6. Results

We applied this technique to various images with different exposure times of up to 25 s. The images were taken with a single chip 6 million pixel professional digital camera (Kodak DCS 560). The original sensor values were reconstructed for each pixel. The optimal correction factor k was robustly found using the quality measures described above and even highly corrupted images could be cleaned.

Figure 3 shows an example of the cleaning process. Note that although the test image and the noise image were taken with the same exposure time the optimal correction factor is $k = 0.65$ and not 1.0.

Some noise pixels were usually left untouched by our technique. They seem to be due to other noise sources or random processes beyond our control. For visually pleasing images these pixels can be treated manually or with the use of generic noise removal techniques.

Some of our implementations consumed a considerable amount of resources as a huge amount of data was produced and analyzed. Currently the computation time for an optimization consisting of 200 steps for a 6 million pixel image on an SGI Octane is 88 s if we use the Huffman code length as quality measure. Smaller image regions lead to a linear speedup. In an ideal case where the exposure times of the target and the noise image were identical we could find the optimal correction factor using only 2000 pixels.

7. Conclusion and Future Work

Dark frame subtraction is an effective technique to reduce dark current noise if the dark frame is scaled appropriately. Using entropy based techniques the scale can be robustly determined without additional knowledge about the conditions during the exposure. The true pixel values can be reconstructed even for highly corrupted images without introducing smoothing artifacts.

We developed this approach to extend the usable range of exposure times of our camera, which we use a measurement tool e.g. to determine the surface characteristics of different test objects. As there is still a tradeoff between

very accurate results and computation cost if a large number of images is processed, we hope to further improve the efficiency. Other measures for the information content could allow the use of the simple minimum entropy criteria even for very different exposure times. Finally, it could be interesting to incorporate this noise removal technique directly into a camera.

For this technique a patent application has been filed.

8. Acknowledgments

We would like to thank Thomas Neumann and Jan Uschok for their help in acquiring images for this project. Thanks to Hartmut Schirmacher, Philipp Slusallek, and Philippe Bekaert for proofreading and many helpful suggestions.

9. References

- [1] Richard L. Baer. CCD Requirements for Digital Photography. In *Proceedings of IS&T's PICS Conference*, pages 26–30, March 2000.
- [2] Ed Chang, Shiufun Cheung, and Davis Pan. Color Filter Array Recovery Using a Threshold-based Variable Number of Gradients. In *SPIE Proceedings Volume 3650. Sensors, Cameras, and Applications for Digital Photography*, pages 36–43, 1999.
- [3] P. Deutsch. Gzip file format specification version 4.3. RFC 1952, May 1996.
- [4] T. Huang, G. Yang, and G. Tang. A fast two-dimensional median filtering algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 27(1):13–18, February 1979.
- [5] David A. Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, September 1952.
- [6] Arch C. Luther. *Video camera technology*. The Artech House audiovisual library. Artech House Publishers, 1998.
- [7] John C. Russ. *The image processing handbook*. CRC Press, 3. ed. edition, 1999.
- [8] C. E. Shannon. A mathematical theory of communication. *Bell System Tech. J.*, 27:379–423, 623–656, 1948.

10. Biographies

Michael Goesele studied computer science at the University of Ulm, Germany, and at the University of North Carolina at Chapel Hill, USA. He received his Master in computer science from the University of Ulm in 1999 and is now working as a PhD student in the computer graphics group of the Max-Planck-Institut für Informatik in Saarbrücken, Germany. His interests include color science, acquisition techniques for computer graphics, and image based rendering.

Wolfgang Heidrich is an Assistant Professor at the University of British Columbia, Canada. Previously he was a Research Associate at the computer graphics group of the Max-Planck-Institut für Informatik in Saarbrücken, Germany, where he chaired the activities on image-based and hardware-accelerated rendering. He received a PhD in computer science from the University of Erlangen, Germany, in 1999, a Master in mathematics from the University of Waterloo, Canada, in 1996, and a Diploma in computer science from the University of Erlangen, Germany, in 1995. His research interests include hardware-accelerated and image-based rendering, global illumination, and interactive computer graphics.

Hans-Peter Seidel is the scientific director and chair of the computer graphics group at the Max-Planck-Institut für Informatik and a professor of computer science at the University of Saarbrücken, Germany. He is an adjunct professor of computer science at the University of Erlangen, Germany, and at the University of Waterloo, Canada. Seidel's research interests include computer graphics, geometric modeling, high-quality rendering, and 3D image analysis and synthesis. He has published some 150 technical papers in the field and has lectured widely on these topics.






Figure 3: Excerpt of test image B taken with 25 s exposure time. Top: original version of the image. Middle: captured noise image with 25 s exposure time. Bottom: cleaned version for optimal $k = 0.65$. The checkerboard pattern is due to a color filter array on the CCD chip.





Figure 4: Color images taken from test image B. The upper version was not corrected, the lower image was corrected using our method. The images were reconstructed using the threshold-based variable number of gradients method [2].