

# Embedded Real-time HD Video Deblurring

Timothy J. Dysart and Jay B. Brockman  
Emu Solutions, Inc.

South Bend, IN 46617

Email: {tdysart,jbrockman}@emutechnology.com

Stephen Jones and Fred Bacon

Aerodyne Research, Inc.

Billerica, MA 01821

Email: {sjones,bacon}@aerodyne.com

**Abstract**—This paper explores a computational deblurring algorithm that will ultimately be implemented in an embedded system with a targeted form factor of 2”x2”x3”. The deblurring algorithm completes a Fourier filtering step followed by a wavelet transform denoising step on a 1080x1920 Bayer input 30 frame per second video feed. A major challenge in performing this processing in real time is that the wavelet denoising process utilizes the stationary wavelet transform, thus exploding the bandwidth requirements of the algorithm. To reach the desired form-factor and performance rate, a hardware accelerator is required. While both GPU and FPGA implementations have been pursued, this paper limits itself to describing our successful implementation using a desktop GPU card. Additionally, we briefly highlight methods, left for future work, for improving GPU performance based on our FPGA implementation efforts that should aid in scaling from our current desktop implementation to an embedded implementation.

## I. INTRODUCTION

The goal of this work is to implement a high definition (HD) video deblurring algorithm in real time at 30 frames per second; as we assume that the point spread function of the camera and lens is known *a priori*, we are implementing a non-blind image deconvolution. The overall system consists of an out-of-focus HD video camera capturing 1080x1920 Bayer color formatted frames at a rate of 30 frames per second and a computational component that process these frames to provide a reconstructed video feed to the end user. As our ultimate goal is an implementation that fits into a 2”x2”x3” volume, using a hardware accelerator is required.

Both GPU and FPGA based hardware accelerators have been explored with the GPU based implementation being the main target; as such, we do not discuss the FPGA implementation in depth here. We favored the GPU platform as it provides a more flexible structure for investigating changes to the underlying frame reconstruction algorithms. By being able to quickly investigate various algorithmic options, we could then target an FPGA implementation and minimize our need to make expensive design modifications to explore different algorithmic choices. However, as we learned during the FPGA implementation phase, the capabilities of GPUs can hide implementation details that result in a potentially sub-optimal solution. The key capability difference between the devices was that the GPU had a peak memory bandwidth of almost 10x than the FPGA framework we studied.

Memory bandwidth is a key consideration in any implementation as moving a color image of 1080x1920 24-bit pixels between memory and the device requires moving 5.93 MB;

thus just copying 30 frames/sec requires almost 356 MB/sec of bandwidth. The deblurring algorithm utilized, the ForWaRD (Fourier-Wavelet Regularized Deconvolution) algorithm [1], has a Fourier filtering component and a wavelet transform denoising component and, even after some algorithmic reductions, requires at least 55 GB/sec of memory bandwidth in our current GPU implementation. Clearly, minimizing this value will have a major performance impact.

Parallelizing the computation and minimizing memory bandwidth requirements were two major factors that enabled us to go from needing nearly 30 seconds to deblur a 1024x2048 grayscale frame with an unoptimized implementation on a CPU to easily meeting the 30 frames per second requirement to fully process 1080x1920 color frames on an Nvidia graphics card. While we have needed to use a rather powerful graphics card to meet the timing requirements of our computational deblurrer, we are continuing to develop our implementation to meet our performance goals on an embedded platform like the Nvidia Jetson Pro automotive development kit.

The scope of this paper is the implementation of our computational deblurring system using GPUs and how we have met the real-time processing requirement. As a result, many of the details regarding the choice of image processing algorithms and the remainder of the system are not contained here.

Section II sets up the deblurring problem and describes the algorithms used here. We describe our implementation of the algorithms in Section III. Performance results are in Section IV. Section V discusses the implementation modifications that we will explore to be able to meet our desired frame rate in an embedded system. Section VI concludes the paper.

## II. DEBLURRING PROBLEM AND ALGORITHMS

The image blurring process can be modeled as

$$y(x) = g(x) * f(x) + n(x) \quad (1)$$

where  $y$  is the blurred image,  $f$  is the clear image,  $g$  is a point spread function that characterizes the blur, and  $n$  is an additive noise process. The goal of the deblurring process is to recover the image from the noisy blurred measurements.

The naïve approach is to use the fact that convolution in the spatial domain is equivalent to multiplication in the frequency domain and attempt to recover  $f$  by applying  $1/G(\omega)$  to  $y$ . However, since  $G(\omega)$  invariably has regions of small amplitude, the effect is to cause a large amplification of the noise so this approach is rarely practical. The Wiener filter,

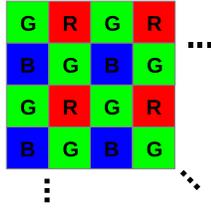


Fig. 1. Bayer filter format.

$$\frac{G^*(\omega)}{|G(\omega)|^2 + \frac{S_n(\omega)}{S_f(\omega)}} \quad (2)$$

minimizes the mean square error and performs much better than Eq. 1, but it requires knowledge of the power spectra of  $f$  and  $n$  which usually are not known. A useful, and frequently applied, technique is to assume that the spectra of  $f$  and  $w$  are white and that the signal to noise ratio is reasonably large. In that case, Eq. 2 becomes

$$\frac{G^*(\omega)}{|G(\omega)|^2 + K} \quad (3)$$

where  $K$  is a small number. This regularized Fourier domain filter can produce a reasonably sharp image, albeit noisy. The approach used by the ForWaRD algorithm, discussed next, is to follow this operation with wavelet denoising. The wavelet processing is very effective at reducing the noise while not degrading detail in the image [2].

### A. Deblurring Algorithms

There are two main algorithms required to convert the input frames, which are blurry and have a Bayer input format, into a reconstructed frame which has been demosaicked and deblurred. Demosaicking is the process of taking a color input frame in the Bayer format, which has only one hue for each pixel as shown in Fig. 1, and interpolating the missing pixels so that each pixel has all three hues.

We initially utilized the one-step alternating projections demosaicking approach in Ref. [3] which was based on the approach of Ref. [4]. The results of Ref. [4] demonstrated a clear improvement over the bilinear interpolation method of demosaicking while the method of [3] demonstrated a performance improvement over the earlier approach. Additionally, we later explored the use of bilinear interpolation for the demosaicking and found that it had some desirable properties that we will describe in more depth in the following section.

To complete the image deblurring, we utilized the Fourier-Wavelet Regularized Deconvolution (ForWaRD) approach of Ref. [1]. The main aspects of this approach, shown in Fig. 2 are the pre-compute section (light gray), the threshold update section (white), and the deblurring section (dark gray). The following notations are used in the figure:

- PSF: Point-spread function
- FFT/IFFT: 2D forward and inverse Fourier transforms
- WT/IWT: 2D forward and inverse wavelet transforms

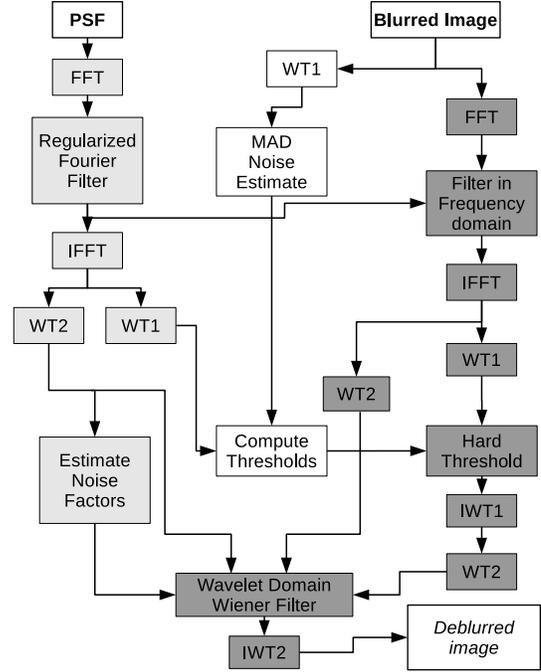


Fig. 2. ForWaRD Algorithm. See text for an explanation of the abbreviations used.

- MAD: Median Absolute Deviation

The three major sections of the algorithm do the following (the latter two are done separately for each color):

- Pre-compute (light gray boxes): This computes a Fourier filter, via Eq. 3, and initial threshold and noise factors based on the point-spread function of the camera lens which is defocused by a known amount. The initial threshold value is computed after transforming the Fourier filter back into the time domain and is based on the Daubechies-4 wavelet transform (WT1). The initial noise factors are based on the Haar wavelet transform (WT2). Once these values are computed, they do not change while the system is operating.
- Threshold Update (non-input/output white boxes): This section of the algorithm takes the blurred input frame and needs to compute just a single-subband of the first level Daubechies wavelet transform on the blurred frame after it has been demosaicked. The MAD Noise Estimate sorts the pixel values, finds the median, and uses the median to update the threshold values used in the deblurring section. While this section of the algorithm does not have to be computed with each frame, it is desirable to do so.
- Deblurring section (dark gray boxes): This is the main component of the algorithm and consists of a filtering step in the frequency domain followed by a denoising step, via wavelets, after returning to the spatial domain. As can be seen in Fig. 2, one wavelet transform goes directly into the wavelet domain Wiener filter while the other is thresholded and then after an inverse

and forward wavelet transform (required to match up the wavelet bases) also goes into the wavelet domain Wiener filter.

One of the most challenging components of implementing this algorithm is that we are utilizing the stationary wavelet transform (aka, *algorithme à trous*) which maintains all of the output rather than subsampling the output. As a result, the stationary wavelet transform produces four full-sized frame outputs at each level of 2D transform rather than being able to remain stored in the space of the original frame. Similarly, when doing the inverse stationary wavelet transform, we are collapsing four frames back into a single frame at each level.

Considering just the Daubechies-4 wavelet transforms that occur before and after the Hard Threshold step, a direct implementation of a three level forward transform has a total of 9 frame reads and 18 frame writes and results in 10 frames left for use in the Hard Threshold step. The total number of reads and writes is flipped for performing the inverse transform. Even with embedding the thresholding step into the inverse transform processing to reduce memory bandwidth, over 35 GB/sec of memory bandwidth is required for just this processing step (single precision floating point values).

### III. ALGORITHM IMPLEMENTATION

We have developed the DDGPU (Demosaic and Deblur on GPU) library based on Nvidia’s CUDA GPU programming model to implement the algorithms in the prior section. Throughout the development process, we frequently modified the demosaicking and deblurring algorithms in order to reach the desired 30 frames/sec processing rate. In this section, we highlight the major algorithmic changes and performance optimizations that enabled us to reach the desired frame rate.

One feature of the DDGPU library is that it preallocates and maintains all of the memory device space necessary to run the algorithm. This preallocation, approximately 1GB of device memory, prevents losing time to create and free memory spaces as each frame is being processed.

#### A. Demosaicking Modifications

While our original demosaicking approach was the alternating one-step projections of Ref. [3], we found that this approach was not particularly amenable to a GPU based implementation. With this algorithm, each color of an input frame is broken into four one-fourth sized subframes and then the majority of the demosaicking operation is performed on the subframes (total of 12 subframes). Once the subframe processing has completed, the twelve subframes are stitched back together into a final color image. The major challenge of this approach is that the initial breaking of the frame and the final stitching of the frame require processing that is not coalesced<sup>1</sup>. This lack of coalescing occurs because each 2x2 pixel block in the input frame has a single pixel to direct to each subframe. Similarly, a single row of the output frame has to pull data from 6 subframes so one can either optimize the reads or the writes to be coalesced, but the other will not be.

<sup>1</sup>Memory *coalescing* refers to a group of threads (a warp) accessing contiguous memory, preferably aligned to the L2 cache line width of the GPU.

While the memory coalescing issues can be managed, this demosaicking method is also computationally expensive as four forward and eight inverse 2D Fourier transforms are also required for each input frame. This, coupled with the issues outlined above, led us to investigate bilinear interpolation as a demosaicking approach. By changing our approach to using bilinear demosaicking, we were able to quickly separate an input frame into its respective color frames by applying a mask (precomputed and stored) for each color onto the entire input frame. This results in a separate image for each color which is then transformed into the frequency domain. The demosaicking is then done for each color by embedding it into the frequency domain filtering step by first multiplying the frequency domain version of the input by a precomputed demosaicking filter and then multiplying this result by the regularized Fourier filter. While some additional processing was required to write out and perform an inverse 2D Fourier transform on the blurry demosaicked frame (the result of the first multiply) for updating the thresholds, we greatly reduced the time required to demosaic a frame.

One of the unexpected benefits from switching the demosaicking algorithm was that some of the edge effects using the one-step approach, shown in Fig. 3(a), disappeared when using the bilinear interpolation approach as shown in Fig. 3(b). While not a major impact on the final frame output, it does provide a more visually appealing output.



Fig. 3. Edge differences after running our version of the deblurring algorithm after using (a) one-step alternating projections and (b) bilinear demosaicking.

#### B. Deblurring Implementation

Overall, we made two major changes to the ForWaRD algorithm to reach our performance goals. The first was the elimination of the wavelet domain Wiener filter and the second modification was using only a subset of the frame, rather than a full frame, to update the thresholds used in the wavelet denoising process (the MAD noise estimate and Compute Thresholds steps). After applying these changes, there was only a minimal impact to the PSNR ratios between the original frames and the frames output from applying the full ForWaRD algorithm and the reduced ForWaRD algorithm, thus the algorithm modifications were acceptable.

As the benefit of eliminating the wavelet domain Wiener filter is obvious based on the bandwidth requirements of a wavelet transform, we focus on the changes to the threshold update steps. One of the major components in performing the MAD noise estimate step is that the frame under consideration has to be sorted to identify the median pixel value in the frame. For a 1080x1920 input frame, this requires sorting over 2 million values. By using only a portion of the frame, e.g., 64x64 pixels taken from the middle of the frame, the sort is reduced to about 4,000 values; clearly, this is a major performance benefit.

In terms of the actual implementation of the ForWaRD algorithm, substantial performance improvements came from two major sources: the reduction of host/device memory transfers and reducing unnecessary data movements, particularly frame transposes, in the denoising process. Host/device memory transfers were generally eliminated by precomputing as much data as possible and placing it into device memory as well as keeping small sequential functions on the device rather than the host. The second improvement, reducing data movements, was extremely beneficial as frame transposes perform no practical work. In the full ForWaRD algorithm, when using a wavelet transform depth of 4 levels, the longest path (through the “Hard Thresholding” step) required 100 frame transposes. The majority of these transposes were in reorienting the frame to align the columns of the frame onto the rows to do the wavelet transform on the columns and have coalesced memory accesses and then reorienting the frame back to its original alignment.

To eliminate these unnecessary transposes, we utilized the block based approach of Ref. [5] to perform the transform operation. This approach does the horizontal (row-based) pass in a standard fashion. However, for the vertical (column-based) pass, they use “slabs” to perform the wavelet transform on a set of rows within a column. The slabs read in a block of data that is a power-of-2 number of columns wide to ensure data coalescing and the number of rows to be updated along with additional rows to handle boundaries. This is an extremely advantageous GPU operation as each slab is then computationally independent of the other slabs and can be a single operational block.

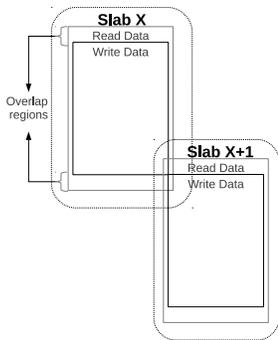


Fig. 4. Slabs for doing wavelet transform operations on a set of columns

Figure 4 shows an overview of the read and write regions for two slabs with slab X+1 shifted by a group of columns (e.g. 16 or 32) to improve the visual quality. As the individual wavelet transforms operate on “local” data, minimizing the overlap regions is a desirable option. For example, in our implementation, we utilize a slab that is 16 columns wide and 120 rows long that then writes out 80 rows of data. While this is a significant amount of read overhead at nearly 50%, it is a far superior approach to using frame transposes. Quantitatively, if we assume that two transposes are necessary to rotate the frame and then return it to its original orientation, nearly 16MB of data is transferred as compared to a read overhead of about 4MB (per color, single point floating point values). Stretching this out to all three colors and our desired frame rate, replacing 2 transposes with a slab-based operation saves over 1GB/sec of memory bandwidth.

TABLE I. PSNR VALUES (DB) FOR EACH VERSION OF FIG. 5 AS COMPARED TO THE ORIGINAL.

	Y	Cb	Cr
Blurry	25.64	40.73	38.25
Full	30.39	40.50	38.8
Reduced	32.74	41.32	40.64

TABLE II. GPU CARDS USED HERE. EACH CARD HAS A TOTAL OF 64 KB OF LOCAL MEMORY/CACHE PER SYMMETRIC PROCESSOR.

	GTX 580	GTX 670	GTX 780
Compute Cores	512	1344	2304
Cores per symmetric processor	32	192	192
Max GFLOPs (single-precision)	1030	2460	3977
Memory Bandwidth (GB/s)	192.4	192.2	288.4
Cuda Architecture and Capability	Fermi, 2.0	Kepler, 3.0	Kepler, 3.5

### C. Image Quality

Figure 5 shows four snips of an image from our test set to demonstrate the visual performance of our implementation of ForWaRD. Table I provides the PSNR results (Y=luminance, Cb, Cr = chrominance) when comparing the full-sized figures in (b)-(d) against the original image. As the figure and table results demonstrate, this is an effective deblurring approach and there is no degradation in quality using the reduced version of the algorithm (for this case, the reduced version is actually better).

## IV. PERFORMANCE

While implementing the DDGPU library, we ended up using three different Nvidia GPU cards. Table II shows the capabilities of each of these cards. Overall, the majority of our algorithm modifications and development occurred on the GTX 580 and GTX 670 cards. Since our initial tests of the GTX 780 card easily met our desired frame rate, we did not pursue any further optimizations. For reference, a 30 frames per second rate provides a window of just 33ms to process a frame. In general, we have written custom CUDA kernels except for the sort (CUDA Thrust library), the FFT/IFFT operations (CUFFT, CUDA FFT library), and a few instances where the CUBLAS (CUDA BLAS) library could be utilized.

### A. Demosaicking Performance

Our best performing version of the alternating one-step projections demosaicking method required around 7.5ms per frame on the GTX 670 card. After integrating the bilinear demosaicking in with the Fourier filtering step, we reduced the demosaicking time to approximately 1.5ms per frame. However, that demosaicking time is slightly misleading figure as some frames need to be demosaicked and output in their blurry state to update the noise thresholds. This results in having to write an intermediate data value from the Fourier filtering step to memory, performing an inverse 2D FFT on it, and then scaling it for use in the threshold update steps. While this extra processing is done for each color, the total processing “penalty” is about 5ms, thus bringing the total demosaicking time to 6.5ms which is still an improvement over the one-step projection approach. Overall, the major performance boost occurs when threshold calculations do not have to be done with each frame as the algorithm can skip three 2D inverse Fourier transforms in these cases.



Fig. 5. Various image snippets: (a) Original (b) blurred (c) full algorithm (d) our reduced implementation.

### B. Deblurring Performance

In the prior section we discussed our algorithmic changes prior to our discussion on the implementation of the deblurring algorithm. We reverse that discussion as the limits of the implementation improvements was the impetus for making the algorithmic changes. As mentioned previously, the major performance improvement was in eliminating frame transposes. In our initial attempts to reduce the number of transposes, we skipped realignment transposes when possible. This effort removed slightly more than half of the transposes and when coupled with some additional improvements in reducing other memory bandwidth and placing small “sequential” functions on the GPU rather than the host, we improved our time per frame on the GTX 580 card from about 165ms to 100ms (one-step demosaicking, 4 level deep wavelet transforms) per frame (no threshold updates; about 12ms to complete them).

After applying the approach of Ref. [5] to remove the remaining transposes in the wavelet transform processing step, our run time on the GTX 580 card was reduced to about 72ms per frame (no algorithmic changes). One of the more interesting results we had was when we switched from the GTX 580 to the GTX 670 card at this stage of development. This switch in card caused a performance reduction as DDGPU was optimized for the Fermi GPU architecture [6] rather than the Kepler GPU architecture of the GTX 670 [7].

Once we optimized for the GTX 670 and had not made significant progress on the run time per frame, we began to investigate the ForWaRD algorithmic changes above to go along with the bilinear interpolation demosaicking. The first optimization was to reduce the number of levels of wavelet transform to three as each level added several milliseconds of processing time. Further study then identified that the wavelet domain Wiener filter was having almost no impact on the resulting output image, so it was removed as well. With these changes our processing time per frame without updating the thresholds approached 30ms, but updating the thresholds took a much longer time in proportion – almost 20ms.

This 20ms difference consists of the 5ms penalty for having to compute the demosaicked blurry frame and nearly 5ms per color for updating the thresholds. Further analysis showed that the main time component of the threshold update step was in sorting the image to find the median. Although we utilized the CUDA Thrust library [8] to have a native CUDA sort function, there were few options for improving performance short of using the smaller window as discussed in the prior section. By using a 64x64 pixel subframe to update the thresholds, we reduce the time per color to about 0.7ms. This then leaves a roughly 7-7.5ms difference in per frame computation times when the thresholds are updated and when they are not.

While we are at the edge of performing the demosaic and

TABLE III. PERFORMANCE SUMMARY.

Time per frame (ms)	Color	GPU	Major Improvements
70	Gray	GTX 580	Initial GTX 580 implementation
55	Gray	GTX 580	Removed additional host/device transfers
165	Color	GTX 580	Initial color implementation; demosaicking nearly "free"
100	Color	GTX 580	Removed unnecessary computation, especially redundant frame transposes
72	Color	GTX 580	Removed remaining transposes except those that may be in CUDA libraries
37.5	Color	GTX 670	Full reduced algorithm (includes Threshold Update section)
23	Color	GTX 780	Upgraded hardware; safely met performance requirement

deblurring process in real time, these performance estimates do not consider the time required to move the frame out of the camera nor the time for the host system to display the output frame. In order to reach our goal of operating the system in real time, we acquired a GTX 780 card (also a Kepler based GPU architecture) and found it capable of processing frames, including updating the thresholds, in roughly 23ms which left plenty of time for the other required operations.

Table III summarizes the algorithm performance at various stages of development.

## V. MOVING TO AN EMBEDDED GPU SYSTEM

To date, we have implemented our deblurring solution on desktop based GPU cards. While it simplifies development and demonstrates the desired capability, a desktop card is not a feasible solution for the desired system volume of 2"x2"x3". However, should a slightly larger volume be acceptable, then a system similar to the Nvidia Jetson [9] may be a viable solution. If we can combine this board with its Tegra K1 chip with a single Kepler multiprocessor (7 Kepler multiprocessors are on the GTX 670) and have a second board with a discrete GPU, it is reasonable to think that 3-4 Kepler multiprocessors will be available for use.

Assuming this type of configuration, the DDGPU library will need to reduce its runtime by nearly 50% again (based on the GTX 670 performance). It is in attempting to reach this performance level that the lessons learned from trying to implement our deblurring algorithm in an FPGA can be applied. In particular, since our FPGA system did not have the same memory bandwidth as the GPUs, we identified two possible implementation modifications that can be done in the wavelet processing to enable more computation per CUDA kernel and reduce memory bandwidth requirements. As we have shown, reducing bandwidth is a highly productive approach for improving processing speed.

The first method divides the image into 16x16 or 32x32 pixel blocks and performs the full wavelet denoising on each block. We have validated the correctness of this approach in Matlab. This method would cause a massive reduction in memory bandwidth as only one frame read and one frame write would be required per pixel block. The main limitation, however, is that a large amount of local memory may be needed to complete the processing steps; in turn, this may limit the amount of computation done on the device at a given time and impact performance in a negative manner.

The second method would be to immediately expand and collapse frames that go through the hard threshold step but are not used in the next level of the wavelet transform. Doing this would remove several frame reads and writes from both the forward and inverse wavelet transform steps. Additionally, this method should be applicable to the prior method of doing the full wavelet transform on small pixel blocks to reduce the local memory required to process each block.

Another modification, separate from the wavelet processing, is that since the demosaicking process is quite expensive when needing to update the noise thresholds, it may be worthwhile to investigate if demosaicking should be applied to the blurry frame prior to performing the 2D forward FFT on it. While this may slow the processing time for frames that do not update the thresholds, it may reduce the processing time of frames where the thresholds are updated by enough to be a net positive.

## VI. CONCLUSION

We have shown the implementation of a HD video demosaicking and deblurring algorithm, a modified version of ForWaRD, that can run at greater than 30 frames per second on desktop graphics cards. We have also described a viable path for further development, applying lessons learned from developing an FPGA based implementation, that would allow embedded scale graphics hardware to run this algorithm in real time. In the near term, we will continue to advance our implementation on the desktop graphics card to reach a per frame processing time that would allow for an embedded system, such as the Nvidia Jetson Pro, to perform the final implementation in real-time.

## ACKNOWLEDGMENT

The authors would like to thank the United States Army Research Laboratory for their support of this work.

## REFERENCES

- [1] R. Neelamani, H. Choi, and R. Baraniuk, "Forward: Fourier-wavelet regularized deconvolution for ill-conditioned systems," *Signal Processing, IEEE Transactions on*, vol. 52, no. 2, pp. 418-433, Feb 2004.
- [2] S. Mallat, *A wavelet tour of signal processing*, 3rd ed. Academic Press, 2009.
- [3] Y. Lu, M. Karzand, and M. Vetterli, "Demosaicking by alternating projections: Theory and fast one-step implementation," *Image Processing, IEEE Transactions on*, vol. 19, no. 8, pp. 2085-2098, Aug 2010.
- [4] B. Gunturk, Y. Altunbasak, and R. Mersereau, "Color plane interpolation using alternating projections," *Image Processing, IEEE Transactions on*, vol. 11, no. 9, pp. 997-1013, Sep 2002.
- [5] W. van der Laan, A. Jalba, and J. B. T. M. Roerdink, "Accelerating wavelet lifting on graphics hardware using cuda," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 22, no. 1, pp. 132-146, Jan 2011.
- [6] N. Corporation, "Nvidia's next generation cuda compute architecture: Fermi," Nvidia Corporation, White paper, 2009.
- [7] —, "Nvidia's next generation cuda compute architecture: Kepler gk 110," Nvidia Corporation, White paper, 2012.
- [8] Nvidia, "Thrust::cuda toolkit documentation," Online, <http://docs.nvidia.com/cuda/thrust/>.
- [9] —, "Jetson automotive development platform," Online, <http://www.nvidia.com/object/jetson-automotive-development-platform.html>.