University of Groningen

Accelerating Wavelet Lifting on Graphics Hardware Using CUDA

Laan, Wladimir J. van der; Jalba, Andrei C.; Roerdink, Jos B.T.M.

Published in:
IEEE transactions on parallel and distributed systems

DOI:
[10.1109/TPDS.2010.143](https://doi.org/10.1109/TPDS.2010.143)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2011

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Laan, W. J. V. D., Jalba, A. C., & Roerdink, J. B. T. M. (2011). Accelerating Wavelet Lifting on Graphics Hardware Using CUDA. *IEEE transactions on parallel and distributed systems*, 22(1), 132-146.
<https://doi.org/10.1109/TPDS.2010.143>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Accelerating Wavelet Lifting on Graphics Hardware Using CUDA

Wladimir J. van der Laan, Andrei C. Jalba, and Jos B.T.M. Roerdink, *Senior Member, IEEE*

Abstract—The Discrete Wavelet Transform (DWT) has a wide range of applications from signal processing to video and image compression. We show that this transform, by means of the lifting scheme, can be performed in a memory and computation-efficient way on modern, programmable GPUs, which can be regarded as massively parallel coprocessors through NVidia's CUDA compute paradigm. The three main hardware architectures for the 2D DWT (row-column, line-based, block-based) are shown to be unsuitable for a CUDA implementation. Our CUDA-specific design can be regarded as a hybrid method between the row-column and block-based methods. We achieve considerable speedups compared to an *optimized* CPU implementation and earlier non-CUDA-based GPU DWT methods, both for 2D images and 3D volume data. Additionally, memory usage can be reduced significantly compared to previous GPU DWT methods. The method is scalable and the fastest GPU implementation among the methods considered. A performance analysis shows that the results of our CUDA-specific design are in close agreement with our theoretical complexity analysis.

Index Terms—Discrete wavelet transform, wavelet lifting, graphics hardware, CUDA.

1 INTRODUCTION

THE wavelet transform, originally developed as a tool for the analysis of seismic data, has been applied in areas as diverse as signal processing, video and image coding, compression, data mining, and seismic analysis. The theory of wavelets bears a large similarity to Fourier analysis, where a signal is approximated by superposition of sinusoidal functions. A problem, however, is that the sinusoids have an infinite support, which makes Fourier analysis less suitable to approximate sharp transitions in the function or signal. Wavelet analysis overcomes this problem by using small waves, called *wavelets*, which have a compact support. One starts with a wavelet prototype function, called a *basic wavelet* or *mother wavelet*. Then, a wavelet basis is constructed by translated and dilated (i.e., rescaled) versions of the basic wavelet. The fundamental idea is to decompose a signal into components with respect to this wavelet basis, and to reconstruct the original signal as a superposition of wavelet basis functions; therefore, we speak a *multiresolution analysis*. If the shape of the wavelets resembles that of the data, the wavelet analysis results in a sparse representation of the signal, making wavelets an interesting tool for data compression. This also allows a client-server model of data exchange, where data are first decomposed into different levels of resolution on the server, then progressively transmitted to the client, where the data can be *incrementally* restored as it arrives ("progressive refinement"). This is

especially useful when the data sets are very large, as in the case of 3D data visualization [1]. For some general background on wavelets, the reader is referred to the books by Daubechies [2] or Mallat [3].

In the theory of wavelet analysis, both continuous and discrete wavelet transforms are defined. If discrete and finite data are used, it is appropriate to consider the *Discrete Wavelet Transform* (DWT). Like the discrete Fourier transform (DFT), the DWT is a linear and invertible transform that operates on a data vector whose length is (usually) an integer power of 2. The elements of the transformed vector are called *wavelet coefficients*, in analogy of Fourier coefficients in case of the DFT. The DWT and its inverse can be computed by an efficient filter bank algorithm, called Mallat's pyramid algorithm [3]. This algorithm involves repeated downsampling (forward transform) or upsampling (inverse transform) and convolution filtering by the application of high and low-pass filters. Its complexity is linear in the number of data elements.

In the construction of so-called *first-generation* wavelet bases, which are translates and dilates of a single basic function, Fourier transform techniques played a major role [2]. To deal with situations where the Fourier transform is not applicable, such as wavelets on curves or surfaces, or wavelets for irregularly sampled data, *second-generation* wavelets were proposed by Sweldens, based on the so-called *lifting scheme* [4]. This provides a flexible and efficient framework for building wavelets. It works entirely in the original time/space domain and does not involve Fourier transforms.

The basic idea behind the lifting scheme is as follows: It starts with a simple wavelet, and then, gradually builds a new wavelet, with improved properties, by adding new basis functions. So, the simple wavelet is *lifted* to a new wavelet, and this can be done repeatedly. Alternatively, one can say that a complex wavelet transform is factored into a sequence of simple lifting steps [5]. More details on lifting are provided in Section 3.

• W.J. van der Laan and J.B.T.M. Roerdink are with the Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen, PO Box 407, 9700 AK Groningen, The Netherlands. E-mail: w.j.van.der.laan@rug.nl, j.b.t.m.roerdink@rug.nl.

• A.C. Jalba is with the Institute for Mathematics and Computing Science, Eindhoven University of Technology, PO Box 513, 5600 MB Eindhoven, The Netherlands. E-mail: a.c.jalba@tue.nl.

Manuscript received 4 Sept. 2009; revised 23 Feb. 2010; accepted 2 June 2010; published online 28 July 2010.

Recommended for acceptance by D.A. Bader, D. Kaeli, and V. Kindratenko. For information on obtaining reprints of this article, please send e-mail to: tpsds@computer.org, and reference IEEECS Log Number TPDS-2009-09-0403. Digital Object Identifier no. 10.1109/TPDS.2010.143.

Also, for first-generation wavelets, constructing them by the lifting scheme has a number of advantages [4]. First, it results in a faster implementation of the wavelet transform than the straightforward convolution-based approach by reducing the number of arithmetic operations. Asymptotically, for long filters, lifting is twice as fast as the standard algorithm. Second, given the forward transform, the inverse transform can be found in a trivial way. Third, no Fourier transforms are needed. Lastly, it allows a fully in-place calculation of the wavelet transform, so no auxiliary memory is needed. With the generally limited amount of high-speed memory available, and the large quantities of data that have to be processed in multimedia or visualization applications, this is a great advantage. Finally, the lifting scheme represents a universal discrete wavelet transform which involves only *integer* coefficients instead of the usual floating point coefficients [6]. Therefore, we based our DWT implementation on the lifting scheme.

Custom hardware implementations of the DWT have been developed to meet the computational demands for systems that handle the enormous throughputs in, for example, real-time multimedia processing. However, cost and availability concerns and the inherent inflexibility of this kind of solutions make it preferable to use a more widespread and general platform. NVidia's G80 architecture [7], introduced in 2006 with the GeForce 8800 GPU, provides such a platform. It is a highly parallel computing architecture available for systems ranging from laptops or desktop computers to high-end compute servers. In this paper, we will present a hardware-accelerated DWT algorithm that makes use of the Compute Unified Device Architecture (CUDA) parallel programming model to fully exploit the new features offered by the G80 architecture when compared to traditional GPU programming.

The three main hardware architectures for the 2D DWT, i.e., row-column, line-based, or block-based, turn out to be unsuitable for a CUDA implementation (see Section 2). The biggest challenge of fitting wavelet lifting in the SIMD model is that data sharing is, in principle, needed after every lifting step. This makes the division into independent computational blocks difficult, and means that a compromise has to be made between minimizing the amount of data shared with neighboring blocks (implying more synchronization overhead) and allowing larger data overlap in the computation at the borders (more computation overhead). This challenge is specifically difficult with CUDA, as blocks cannot exchange data at all without returning execution flow to the CPU. Our solution is a sliding window approach which enables us (in the case of separable wavelets) to keep intermediate results longer in shared memory, instead of being written to global memory. Our CUDA-specific design can be regarded as a hybrid method between the row-column and block-based methods. We implemented our methods both for 2D and 3D data, and obtained considerable speedups compared to an *optimized* CPU implementation and earlier non-CUDA-based GPU DWT methods. Additionally, memory usage can be reduced significantly compared to previous GPU DWT methods. The method is scalable and the fastest GPU implementation among the methods considered. A performance analysis shows that the results of our CUDA-specific design are in close agreement with our theoretical complexity analysis.

The paper is organized as follows: Section 2 gives a brief overview of GPU wavelet lifting methods, and previous work on GPU wavelet transforms. In Section 3, we present the basic theory of wavelet lifting. Section 4 first presents an overview of the CUDA programming environment and execution model, introduces some performance considerations for parallel CUDA programs, and gives the details of our wavelet lifting implementation on GPU hardware. Section 5 presents benchmark results and analyzes the performance of our method. Finally, in Section 6, we draw conclusions and discuss future avenues of research.

2 PREVIOUS AND RELATED WORK

In [8], a method was first proposed that makes use of OpenGL extensions on early nonprogrammable graphics hardware to perform the convolution and downsampling/upsampling for a 2D DWT. Later, in [9], this was generalized to 3D using a technique called tile boarding.

Wong et al. [10] implemented the DWT on programmable graphics hardware with the goal of speeding up JPEG2000 compression. They made the decision not to use wavelet lifting, based on the rationale that, although lifting requires less memory and less computations, it imposes an order of execution which is not fully parallelizable. They assumed that lifting would require more rendering passes, and therefore, in the end be slower than the standard approach based on convolution.

However, Tenllado et al. [11] performed wavelet lifting on conventional graphics hardware by splitting the computation into four passes using fragment shaders. They concluded that a gain of 10-20 percent could be obtained by using lifting instead of the standard approach based on convolution. Similar to [10], Tenllado et al. [12] also found that the lifting scheme implemented using shaders requires more rendering steps, due to increased data dependencies. They showed that for shorter wavelets, the convolution-based approach yields a speedup of 50-100 percent compared to lifting. However, for larger wavelets, on large images, the lifting scheme becomes 10-20 percent faster. A limitation of both [11] and [12] is that the methods are strictly focused on 2D. It is uncertain whether, and if so, how they extend to three or more dimensions.

All previous methods are limited by the need to map the algorithms to graphics operations, constraining the kind of computations and memory accesses they could make use of. As we will show below, new advances in GPU programming allow us to do in-place transforms in a *single pass*, using intermediate fast shared memory.

Wavelet lifting on general parallel architectures was studied extensively in [13] for processor networks with large communications latencies. A technique called *boundary postprocessing* was introduced that limits the amount of data sharing between processors working on individual blocks of data. This is similar to the technique we will use. More than in previous generations of graphics cards, general parallel programming paradigms can now be applied when designing GPU algorithms.

The three main hardware architectures for the 2D DWT are *row-column* (RC), *line-based* (LB), and *block-based* (BB), see, for example, [14], [15], [16], [17], and all three schemes are based on wavelet lifting. The simplest one is RC, which applies a separate 1D DWT in both the horizontal and

Fig. 1. Structure of the forward wavelet transform with J stages: recursively split a signal c^0 into approximation bands c^J and detail bands d^J .

vertical directions for a given number of lifting levels. Although this architecture provides the simplest control path (thus being the cheapest for a hardware realization), its major disadvantage is the lack of locality due to the use of large off-chip memory (i.e., the image memory), thus decreasing the performance. Contrary to RC, both LB and BB involve a local memory that operates as a cache, thus increasing bandwidth utilization (throughput). On FPGA architectures, it was found [14] that the best instruction throughput is obtained by the LB method, followed by the RC and BB schemes which show comparable performances. As expected, both the LB and BB schemes have similar bandwidth requirements, which are at least two times smaller than that of RC. Theoretical results [15], [16] show that this holds as well for ASIC architectures. Thus, LB is the best choice with respect to overall performance, for a hardware implementation.

Unfortunately, a CUDA realization of LB is impossible for all but the shortest wavelets (e.g., the Haar wavelet), due to the relatively large cache memory required. For example, the cache memory for the Deslauriers-Dubuc (13, 7) wavelet should accommodate six rows of the original image (i.e., 22.5 KB for 2-byte word data and HD resolutions), well in excess of the maximum amount of 16 KB of shared memory available per multiprocessor, see Section 4.3. As an efficient implementation of BB requires similar amounts of cache memory, this choice is again not possible. Thus, the only feasible strategy remains RC. However, we show in Section 5 that even an improved (using cache memory) RC strategy is not optimal for a CUDA implementation. Nevertheless, our CUDA-specific design can be regarded as a hybrid method between RC and BB, which also has an optimal access pattern to the slow global memory (see Section 4.1.2).

3 WAVELET LIFTING

As explained in Section 1, lifting is a very flexible framework to construct wavelets with desired properties. When applied to first-generation wavelets, lifting can be considered as a reorganization of the computations leading to increased speed and more efficient memory usage. In this section, we explain in more detail how this process works. First, we discuss the traditional wavelet transform computation by subband filtering, and then, outline the idea of wavelet lifting.

3.1 Wavelet Transform by Subband Filtering

The main idea of (first generation) wavelet decomposition for finite 1D signals is to start from a signal $c^0 = (c_0^0, c_1^0, \dots, c_{N-1}^0)$, with N samples (we assume that N is a power of 2). Then, we apply convolution filtering of c^0 by a low-pass analysis filter H and downsample the result by a factor of 2 to get an “approximation” signal (or “band”) c^1 of length $N/2$, i.e., half the initial length. Similarly, we apply convolution filtering of

Fig. 2. Structure of the inverse wavelet transform with J stages: recursively upsample, filter, and add approximation signals c^j and detail signals d^j .

c^0 by a high-pass analysis filter G , followed by down-sampling, to get a detail signal (or “band”) d_1 . Then, we continue with c^1 and repeat the same steps, to get further approximation and detail signals c^2 and d^2 of length $N/4$. This process is continued a number of times, say J . Here, J is called the number of *levels* or *stages* of the decomposition. The explicit decomposition equations for the individual signal coefficients are

$$c_k^{j+1} = \sum_n h_{n-2k} c_n^j, \quad d_k^{j+1} = \sum_n g_{n-2k} c_n^j,$$

where $\{h_n\}$ and $\{g_n\}$ are the coefficients of the filters H and G . Note that only the approximation bands are successively filtered, the detail bands are left “as is.”

This process is presented graphically in Fig. 1, where the symbol $\downarrow_2$ (enclosed by a circle) indicates downsampling by a factor of 2. This means that after the decomposition, the initial data vector c^0 is represented by one approximation band c^J and J detail bands $d^1, d^2, \dots, d^J$. The total length of these approximation and detail bands is equal to the length of the input signal c^0 .

Signal reconstruction is performed by the *inverse* wavelet transform: first upsample the approximation and detail bands at the coarsest level J , then apply synthesis filters $\tilde{H}$ and $\tilde{G}$ to these, and add the resulting bands. (In the case of orthonormal filters, such as the Haar basis, the synthesis filters are essentially equal to the analysis filters.) Again, this is done recursively. This process is presented graphically in Fig. 2, where the symbol $\uparrow_2$ indicates upsampling by a factor of 2.

3.2 Wavelet Transform by Lifting

Lifting consists of four steps: split, predict, update, and scale, see Fig. 3 (left).

1. **Split:** This step splits a signal (of even length) into two sets of coefficients, those with even and those with odd index, indicated by even^{j+1} and odd^{j+1} . This is called the *lazy wavelet transform*.
2. **Predict lifting step:** As the even and odd coefficients are correlated, we can predict one from the other. More specifically, a prediction operator P is applied to the even coefficients and the result is subtracted from the odd coefficients to get the detail signal d^{j+1} :

$$d^{j+1} = \text{odd}^{j+1} - P(\text{even}^{j+1}). \quad (1)$$

3. **Update lifting step:** Similarly, an update operator U is applied to the odd coefficients and added to the even coefficients to define c^{j+1} :

$$c^{j+1} = \text{even}^{j+1} + U(d^{j+1}). \quad (2)$$

Fig. 3. Classical lifting scheme (one stage only). Left part: forward lifting. Right part: inverse lifting. Here, “split” is the trivial wavelet transform, “merge” is the opposite operation, P is the prediction step, U the update step, and K the scaling factor.

4. **Scale:** To ensure normalization, the approximation band c^{j+1} is scaled by a factor of K and the detail band d^{j+1} by a factor of $1/K$.

Sometimes, the scaling step is omitted; in that case, we speak of an *unnormalized* transform.

A remarkable feature of the lifting technique is that the inverse transform can be found trivially. This is done by “inverting” the wiring diagram, see Fig. 3 (right): undo the scaling, undo the update step ($\text{even}^{j+1} = c^{j+1} - U(d^{j+1})$), undo the predict step ($\text{odd}^{j+1} = d^{j+1} + P(\text{even}^{j+1})$), and merge the even and odd samples. Note that this scheme does not require the operators P and U to be invertible: nowhere does the inverse of P or U occur, only the roles of addition and subtraction are interchanged. For a multistage transform, the process is repeatedly applied to the approximation bands, until a desired number of decomposition levels are reached. In the same way as discussed in Section 3.1, the total length of the decomposition bands equals that of the initial signal. As an illustration, we give in Table 1 the explicit equations for one stage of the forward wavelet transform by the (unnormalized) Le Gall (5, 3) filter, both by subband filtering and lifting (*in-place* computation). It is easily verified that both schemes give identical results for the computed approximation and detail coefficients.

The process above can be extended by including more predict and/or update steps in the wiring diagram [4]. In fact, any wavelet transform with finite filters can be decomposed into a sequence of lifting steps [5]. In practice, lifting steps are chosen to improve the decomposition, for example, by producing a lifted transform with better decorrelation properties or higher smoothness of the resulting wavelet basis functions.

Wavelet lifting has two properties which are very important for a GPU implementation. First, it allows a fully *in-place* calculation of the wavelet transform, so no auxiliary

memory is needed. Second, the lifting scheme can be modified to a transform that maps *integers to integers* [6]. This is achieved by rounding the result of the P and U functions. This makes the predict and update operations nonlinear, but this does not affect the invertibility of the lifting transform. Integer-to-integer wavelet transforms are especially useful when the input data consist of integer samples. These schemes can avoid quantization, which is an attractive property for lossless data compression.

For many wavelets of interest, the coefficients of the predict and update steps (before truncation) are of the form $z/2^n$, with z integer and n a positive integer. In that case, one can implement all lifting steps (apart from normalization) by integer operations: integer addition and multiplication, and integer division by powers of 2 (bit-shifting).

4 WAVELET LIFTING ON GPUS USING CUDA

4.1 CUDA Overview

In recent years, GPUs have become increasingly powerful and more programmable. This combination has led to the use of the GPU as the main computation device for diverse applications, such as physics simulations, neural networks, image compression, and even database sorting. The GPU has moved from being used solely for graphical tasks to a fully fledged parallel coprocessor. Until recently, General Purpose GPU (GPGPU) applications, even though not concerned with graphics rendering, did use the rendering paradigm. In the most common scenario, textured quadrilaterals were rendered to a texture, with a fragment shader performing the computation for each fragment.

With their G80 series of graphics processors, NVidia introduced a programming environment called CUDA [7]. It is an API that allows the GPU to be programmed through more traditional means: a C-like language (with some C++ features such as templates) and compiler. The GPU programs, now called *kernels* instead of shaders, are invoked through procedure calls instead of rendering commands. This allows the programmer to focus on the main program structure, instead of details like color clamping, vertex coordinates, and pixel offsets.

In addition to this generalization, CUDA also adds some features that are missing in shader languages: random access to memory, fast integer arithmetic, bitwise operations, and shared memory. The usage of CUDA does not add any overhead, as it is a native interface to the hardware, and not an abstraction layer.

TABLE 1
Forward Wavelet Transform (One Stage Only) by the
(Unnormalized) Le Gall (5, 3) Filter

Subband filtering	$c_k^{j+1} = \frac{1}{8} (-c_{2k-2}^j + 2c_{2k-1}^j + 6c_{2k}^j + 2c_{2k+1}^j - c_{2k+2}^j)$ $d_k^{j+1} = \frac{1}{2} (-c_{2k}^j + 2c_{2k+1}^j - c_{2k+2}^j)$
Lifting	Split: $c_k^{j+1} \leftarrow c_{2k}^j, \quad d_k^{j+1} \leftarrow c_{2k+1}^j$ Predict: $d_k^{j+1} \leftarrow d_k^{j+1} - \frac{1}{2}(c_k^{j+1} + c_{k+1}^{j+1})$ Update: $c_k^{j+1} \leftarrow c_k^{j+1} + \frac{1}{4}(d_{k-1}^{j+1} + d_k^{j+1})$

4.1.1 Execution Model

The CUDA execution model is quite different from that of CPUs, and also different from that of older GPUs. CUDA broadly follows: the data-parallel model of computation [7]. The CPU invokes the GPU by calling a *kernel*, which is a special C-function.

The lowest level of parallelism is formed by *threads*. A thread is a single scalar execution unit, and a large number of threads can run in parallel. The thread can be compared to a fragment in traditional GPU programming. These threads are organized in *blocks*, and the threads of each block can cooperate efficiently by sharing data through fast shared memory. It is also possible to place synchronization points (barriers) to coordinate operations closely, as these will synchronize the control flow between all threads within a block. The Single Instruction Multiple Data (SIMD) aspect of CUDA is that the highest performance is realized if all threads within a *warp* of 32 consecutive threads take the same execution path. If flow control is used within such a *warp* and the threads take different paths, they have to wait for each other. This is called *divergence*.

The highest level, which encompasses the entire kernel invocation, is called the *grid*. The grid consists of blocks that execute in parallel, if multiprocessors are available, or sequentially if this condition is not met. A limitation of CUDA is that blocks within a grid cannot communicate with each other, and this is unlikely to change as independent blocks are a means to scalability.

4.1.2 Memory Layout

The CUDA architecture gives access to several kinds of memory, each tuned for a specific purpose. The largest chunk of memory consists of the *global* memory, also known as device memory. This memory is linearly addressable, and can be read and written at any position in any order (random access) from the device. No caching is done in G80; however, there is limited caching in the newest generation (GT200) as part of the shared memory can be configured as automatic cache. This means that optimizing access patterns is up to the programmer. Global memory is also used for communication with the CPU, which can read and write using API calls. *Registers* are limited per-thread memory locations with very fast access, which are used for local storage. *Shared memory* is a limited per-block chunk of memory which is used for communication between threads in a block. Variables are marked to be in shared memory using a specifier. Shared memory can be almost as fast as registers, provided that bank conflicts are avoided. *Texture memory* is a special case of device memory which is cached for locality. Textures in CUDA work the same as in traditional rendering, and support several addressing modes and filtering methods. *Constant memory* is cached memory that can be written by the CPU and read by the GPU. Once a constant is in the constant cache, subsequent reads are as fast as register access.

The device is capable of reading 32, 64, or 128-bit words from global memory into registers in a single instruction. When access to device memory is properly distributed over threads, it is compiled into 128-bit load instructions instead of 32-bit load instructions. The consecutive memory locations

must be simultaneously accessed by the threads. This is called *memory access coalescing* [7], and it represents one of the most important optimizations in CUDA. We will confirm the huge difference in memory throughput between coalesced and noncoalesced access in our results.

4.2 Performance Considerations for Parallel CUDA Programs (Kernels)

Let us first define some metrics which we use later to analyze our results in Section 5.3 below.

4.2.1 Total Execution Time

Assume that a CUDA kernel performs computations on N data values and organizes the CUDA “execution model” as follows: Let T denote the number of threads in a block, W the number of threads in a warp, i.e., $W = 32$ for G80 GPUs, and B denote the number of thread blocks. Further, assume that the number of multiprocessors (device specific) is M , and NVidia’s occupancy calculator [18] indicates that k blocks can be assigned to one multiprocessor (MP); k is program-specific and represents the total number of threads for which (re)scheduling costs are zero, i.e., context switching is done with no extra overhead. Given that the amount of resources per MP is fixed (and small), k simply indicates the occupancy of the resources for the given kernel. With this notation, the number of blocks assigned to one MP is given by $b = B/M$. Since, in general, k is smaller than b , it follows that the number α of times k blocks are rescheduled is $\alpha = \lceil \frac{B}{Mk} \rceil$.

Since each MP has eight stream processors, a warp has 32 threads and there is no overhead when switching among the warp threads, it follows that each warp thread can execute one (arithmetic) instruction in four clock cycles. Thus, an estimate of the asymptotic time required by a CUDA kernel to execute n instructions over all available resources of a GPU, which also includes scheduling overhead, is given by

$$T_e = \frac{4n}{K} \frac{T}{W} \alpha k l_s, \quad (3)$$

where K is the clock frequency and l_s is the latency introduced by the scheduler of each MP.

The second component of the total execution time is given by the time T_m required to transfer N bytes from global memory to fast registers and shared memory. If thread transfers of m bytes can be coalesced, given that a memory transaction is done per half-warp, it follows that the transfer time T_m is

$$T_m = \frac{2N}{W m M} l_m, \quad (4)$$

where l_m is the latency (in clock cycles) of a memory access. As indicated by NVidia [19], reported by others [20] and confirmed by us, the latency of a *noncached* access can be as large as 400-600 clock cycles. Compared to 24 cycle latency for accessing the shared memory, it means that transfers from global memory should be minimized. Note that for *cached* accesses, the latency becomes about 250-350 cycles.

One way to effectively address the relatively expensive memory transfer operations is by using fine-grained *thread parallelism*. For instance, 24 cycle latency can be hidden by

running six warps (192 threads) per MP. To hide even larger latencies, the number of threads should be raised (thus increasing the degree of parallelism) up to a maximum of 768 threads per MP supported by the G80 architecture. However, increasing the number of threads while maintaining the size N of the problem fixed implies that each thread has to perform less work. In doing so, one should still recall 1) the paramount importance of coalescing memory transactions and 2) the Flops/word ratio, i.e., peak Gflop/s rate divided by global memory bandwidth in words [20], for a specific GPU. Thus, threads should not execute too few operations nor transfer too little data such that memory transfers cannot be coalesced. To summarize, a trade-off should be found between increased thread parallelism, suggesting more threads to hide memory transfer latencies on the one hand, and on the other hand, memory coalescing and maintaining a specific Flops/word ratio, indicating fewer threads.

Let us assume that for a given kernel, one MP has an occupancy of kT threads. Further, if the kernel has a ratio $r \in (0, 1)$ of arithmetic to arithmetic and memory transfer instructions, and assuming a *round-robin* scheduling policy, then the *reduction* of memory transfer latency due to latency hiding is

$$l_h = \sum_{i=0}^{\infty} \left\lfloor \frac{kT}{8} r^i \right\rfloor. \quad (5)$$

For example, assume that $r = 0.5$, i.e., there are as many arithmetic instructions (flops) as memory transfers, and assume that $kT = 768$, i.e., each MP is fully occupied. The scheduler starts by assigning eight threads to an MP. Since $r = 0.5$, chances are that four threads execute each a memory transfer instruction, while the others execute one arithmetic operation. After one cycle, those four threads executing memory transfers are still asleep for at least 350 cycles, while the others just finished executing the flop and are put to sleep too. The scheduler assigns now another eight threads, which again can execute either a memory transfer or a flop, with the same probability, and repeats the whole process. Counting the number of cycles in which four threads executed flops reveals a number of 190 cycles so that the latency is decreased in this way to just $l_m = 350 - 190 = 160$ cycles. In the general case, for a given r and occupancy, we postulate that (5) applies.

The remaining component of the total GPU time for a kernel is given by the synchronization time. To estimate this component, we proceed as follows: Assume that all active threads (i.e., $kT \leq 768$) are about to execute a flop, after which they have to wait on a synchronization point (i.e., on a barrier). Then, assuming again a round-robin scheduling policy and reasoning similar as for (3), the idle time spent waiting on the barrier is

$$T_s = \frac{T_e}{n} = \frac{4}{K} \frac{T}{W} \propto k l_s. \quad (6)$$

This agrees with NVidia's remark that if within a warp thread, divergence is minimal, then waiting on a synchronization barrier requires only four cycles [19]. Note that the expression for T_s from (6) represents the minimum synchronization time, as threads were assumed to execute

(fast) flops. In the worst case scenario—at least one active thread has just started before the synchronization point, a slow global memory transaction—this estimate has to be multiplied by a factor of about $500/4 = 125$ (the latency of a noncached access divided by four threads).

To summarize, we estimate the total execution time T_t as $T_t = T_e + T_m + T_s$.

4.2.2 Instruction Throughput

Assuming that a CUDA kernel performs n flops in a number c of cycles, then the *estimate* of the asymptotic Gflop/s rate is $G_e = \frac{32}{c} \frac{M}{n} \frac{K}{T_t}$, whereas the *measured* Gflop/s rate is $G_m = \frac{nN}{T_t}$; here, K is the clock rate and T_t the (measured) total execution time. For the 8800 GTX GPU, the peak instruction throughput using register-to-register MAD instructions is about 338 Gflop/s and drops to 230 Gflop/s when using transfers in/from shared memory [20].

4.2.3 Memory Bandwidth

Another factor which should be taken into account when developing CUDA kernels is the memory bandwidth $M_b = \frac{N}{T_t}$. For example, parallel *reduction* has very low arithmetic intensity, i.e., one flop per loaded element, which makes it bandwidth-optimal. Thus, when implementing a parallel reduction in CUDA, one should strive for attaining peak bandwidth. On the contrary, if the problem at hand is matrix multiplication (a trivial parallel computation, with little synchronization overhead), one should optimize for peak throughput. For the 8800 GTX GPU, the pin bandwidth is 86 GB/s.

4.2.4 Complexity

With coalesced accesses, the number of bytes retrieved with one memory request (and thus, one latency) is maximized. In particular, coalescing reduces l_m (through l_h from (5)) by a factor of about 2. Hence, one can safely assume that $l_m/(2W) \rightarrow 0$. It follows that the total execution time satisfies

$$T_t \sim 4n \frac{N}{WM D}, \quad (7)$$

where n is the number of instructions of a given CUDA kernel, N is the problem size, D is the problem size per thread, and $\sim$ means that both left- and right-hand side quantities have the same order of magnitude.

The *efficiency* of a parallel algorithm is defined as

$$E = \frac{T_s}{C} = \frac{T_s}{M T_t}, \quad (8)$$

where T_s is the execution time of the (fastest) sequential algorithm, and $C = M T_t$ is the *cost* of the parallel algorithm. A parallel algorithm is called *cost efficient* (or *cost optimal*) if its cost is proportional to T_s . Let us assume that $T_s \sim n_s N$, where n_s is the number of instructions for computing one data element and N denotes the problem size. Then, the efficiency becomes

$$E \sim \frac{n_s W D}{4n}. \quad (9)$$

Thus, according to our metric above, for a given problem, *any* CUDA kernel which

1. uses *coalesced* memory transfers (i.e., $l_m/(2W) \rightarrow 0$ is enforced),
2. avoids thread divergence (so that our T_s estimate from (6) applies),
3. minimizes transfers from global memory, and
4. has an instruction count n proportional to $(n_s W D)$

is *cost-efficient*. Of course, the smaller n is, the more efficient the kernel becomes.

4.3 Parallel Wavelet Lifting

Earlier parallel methods for wavelet lifting [13] assumed an MPI architecture with processors that have their own memory space. However, the CUDA architecture is different. Each processor has its own shared memory area of 16 KB, which is not enough to store a significant part of the data set. As explained above, each processor is allocated a number of threads that run in parallel and can synchronize. The processors have no way to synchronize with each other, beyond their invocation by the host.

This means that data parallelism has to be used, and moreover, the data set has to be split into parts that can be processed as independently as possible, so that each chunk of data can be allocated to a processor. For wavelet lifting, except for the Haar [4] transform, this task is not trivial, as the implied data reuse in lifting also requires the coefficients just outside the delimited block to be updated. This could be solved by duplicating part of the data in each processor. Wavelet bases with a large support will, however, need more data duplication. If we want to do a multilevel transform, each level of lifting doubles the amount of duplicated work and data. With the limited amount of shared memory available in CUDA, this is not a feasible solution.

As kernel invocations introduce some overhead each time, we should also try to do as much work within one kernel as possible so that the occupancy of the GPU is maximized. The sliding window approach enables us (in the case of separable wavelets) to keep intermediate results longer in shared memory, instead of being written to global memory.

4.4 Separable Wavelets

For separable wavelet bases in 2D, it is possible to split the operation into a horizontal and a vertical filtering step. For each filter level, a horizontal pass performs a 1D transform on each row, while a vertical pass computes a 1D transform on each column. This lends itself to easy parallelization: each row can be handled in parallel during the horizontal pass, and then, each column can be handled in parallel during the vertical pass. In CUDA, this implies the use of two kernels, one for each pass. The simple solution would be to have each block process a row with the horizontal kernel, while in the vertical step, each block processes a column. Each thread within these blocks can then filter an element. We will discuss better specific algorithms for both passes in the upcoming sections.

4.5 Horizontal Pass

The simple approach mentioned in the previous section works very well for the horizontal pass. Each block starts by reading a line into shared memory using so-called *coalesced reads* from device memory, executes the lifting steps in place in fast shared memory, and writes back the result using *coalesced writes*. This amounts to the following steps:

Fig. 4. Horizontal lifting step. h thread blocks are created, each containing T threads; each thread performs computations on $N/(hT) = w/T$ data. Black quads illustrate input for the thread with id 0. Here, w and h are the dimensions of the input and $N = w \cdot h$.

1. Read a row from device memory into shared memory.
2. Duplicate border elements (implement boundary condition).
3. Do a 1D lifting step on the elements in shared memory.
4. Repeat steps 2 and 3 for each lifting step of the transform.
5. Write back the row to device memory.

As each step is dependent on the output in shared memory of the previous step, the threads within the block have to be synchronized every time before the next step can start. This ensures that the previous step did finish and wrote back its work.

Fig. 4 shows the configuration of the CUDA execution model for the horizontal step. Without loss of generality, assume that $N = w \cdot h$ integers are lifted at level i . Note that if the lifting level $i = 0$, then w and h are the dimensions of the input image. For this step, a number $B = h$ of thread blocks are used, with T threads per block. Thus, each thread performs computations on w/T integers. In the figure, black quads illustrate locations which are processed by the thread with id 0. Neither the number nor the positions of these quads need to correspond to the actual number and positions of locations where computations are performed, i.e., they are solely used for illustration purposes.

By reorganizing the coefficients [21], we can achieve higher efficiency for successive levels after the first transformation. If the approximation and detail coefficients are written back in interleaved form, as is usually the case with wavelet lifting, the reading step for the next level will have to read the approximation coefficients of the previous level in interleaved form. These reads cannot be coalesced, resulting in low-memory performance. To still be able to coalesce, one writes the approximation and detail coefficients back to separate halves of the memory. This will result in a somewhat different memory layout for subbands (Fig. 5), but this could be reorganized if needed. Many compression algorithms require the coefficients stored per subband anyhow, in which case this organization is advantageous.

4.6 Vertical Pass

The vertical pass is more involved. Of course, it is possible to use the same strategy as for the horizontal pass, substituting rows for columns. But this is far from efficient. Reading a column from the data would amount to reading one value per row. As only consecutive reads can be coalesced into one read, these are all performed individually. The processing

Fig. 5. Wavelet lifting for a row of data, representing the result in interleaved (top) and deinterleaved (bottom) form. Here, x_i and y_i are the approximation and detail bands at level i .

steps would be the same as for the horizontal pass, after which writing back is again very inefficient.

We can gain a 10 times speedup by using coalesced memory access. Instead of having each block process a column, we make each block process multiple columns by dividing the image into vertical “slabs,” see Fig. 6. Within a block, threads are organized into a 2D grid of size $V_x \times V_y$, instead of a 1D one, as in the horizontal step. The number S of columns in each slab is a multiple of V_x such that the resulting number of slab rows can still be coalesced, and has the height of the image. Each thread block processes one of the slabs, i.e., $S/V_x \times h/V_y$ data. Using this organization, a thread can do a coalesced read from each row within a slab, do filtering in shared memory, and do a coalesced write to each slab row.

Another problem arises here, namely, that the shared memory in CUDA is not large enough to store all columns for any sizable data set. This means that we cannot read and process the entire slab at once. The solution that we found is to use a sliding window within each slab, see Fig. 7a. This window needs to have dimensions so that each thread in the block can transform a signal element, and additional space to make sure that the support of the wavelet does not exceed the top or bottom of the window. To determine the size of the window needed, how much to advance, and at which offset to start, we need to look at the support of each of the lifting steps.

In Fig. 7a, height is the height of the working area. As each step updates either odd or even rows within a slab, each row of threads updates one row in each lifting step. Therefore, a good choice is to set it to two times the number of threads in the vertical direction. Similarly, width should be a multiple of the number of threads in the horizontal direction, and the size of a row should be a multiple of the coalescable size. In the figure, rows in the top area have been fully computed, while rows in the overlap area still need to go through at least one lifting step. The rows in the working area need to go through all lifting steps, while rows in the bottom area are untouched except as border extension. The sizes of overlap, top, and bottom depend on the chosen wavelet. We will elaborate on this later.

Fig. 6. Vertical lifting step. w/S blocks are created, each containing $T = V_x \cdot V_y$ threads; each thread performs computations on $S/V_x \times h/V_y$ data. Black quads illustrate input for the thread with id (0, 0), whereas vertical lines depict boundaries between image-high slabs.

Fig. 7. (a) The sliding window used during the vertical pass for separable wavelets. (b) Advancing the sliding window: the next window is aligned at the bottom of the previous one, taking the overlap area into account.

4.6.1 The Algorithm

Algorithm 1 shows the steps for the vertical lifting pass. Three basic operations are used: **read** copies rows from device memory into shared memory, **write** copies rows from shared memory back to device memory, and **copy** transfers rows from shared memory to another place in shared memory. The shared memory window is used as a cache, and to manage this, we keep a read and a write pointer. The read pointer *inrow* indicates where to read from, the write pointer *outrow* indicates where to write back. After reading, we advance the read pointer; after writing, we advance the write pointer. Both are initialized to the top of the slab at the beginning of the kernel (lines 1 and 2 of Algorithm 1).

Algorithm 1. The sliding window algorithm for the vertical wavelet lifting transform (see Section 4.6). Here top, overlap, height, bottom are the length parameters of the sliding window (see Fig. 7), and h is the number of rows of the data set. The pointer *inrow* indicates where to read from, the pointer *outrow* indicates where to write back.

- 1: *inrow* $\leftarrow$ 0 {initialize read pointer}
- 2: *outrow* $\leftarrow$ 0 {initialize write pointer}
- 3: *windows* $\leftarrow$ $(h - \text{height} - \text{bottom}) / \text{height}$ {number of times window fits in slab}
- 4: *leftover* $\leftarrow$ $(h - \text{height} - \text{bottom}) \% \text{height}$ {remainder}
- 5: **read**(height + bottom from row *inrow* to row top + overlap) {copy from global to shared memory}
- 6: *inrow* $\leftarrow$ *inrow* + height + bottom {advance read pointer}
- 7: **transformTop**() {apply vertical wavelet lifting to rows in shared memory}
- 8: **write**(height - overlap from row top + overlap to row *outrow*) {write transformed rows back to global memory}
- 9: *outrow* $\leftarrow$ *outrow* + height - overlap {advance write pointer}
- 10: **for** $i = 1$ to *windows* **do** {advance sliding window through slab and repeat above steps}
- 11:   **copy**(top + overlap + bottom from row height to row 0)
- 12:   **read** (height from row *inrow* to row top + overlap + bottom)
- 13:   *inrow* $\leftarrow$ *inrow* + height
- 14:   **transformBlock**() {vertical wavelet lifting}

TABLE 2
Filter Weights of the Two Lifting Steps for the
Deslauriers-Dubuc (13, 7) [22] Wavelet

Offset	Prediction	Update
-3		$-\frac{1}{16}$
-2	$\frac{1}{16}$	
-1		$\frac{9}{16}$
0	$-\frac{9}{16}$	$\bullet$
1	$\bullet$	$\frac{9}{16}$
2	$-\frac{9}{16}$	
3		$-\frac{1}{16}$
4	$\frac{1}{16}$	

The current element being updated is marked with $\bullet$.

```

15:  write(height from row top to row outrow)
16:  outrow ← outrow + height
17:  end for
18:  copy(top + overlap + bottom from row height to
    row 0)
19:  read(leftover from row inrow to row
    top + overlap + bottom)
20:  transformBottom() {satisfy bottom boundary condition}
21:  write(leftover + overlap + bottom from row top to
    row outrow)

```

The first block has to be handled differently because we need to take the boundary conditions into account. So, initially, rows are copied from the beginning of the slab to shared memory, filling it from a certain offset to the end (line 5). Next, we apply a vertical wavelet lifting transform (**transformTop**, line 7) to the rows in shared memory (it may be required to leave some rows at the end untouched for some of the lifting steps, depending on their support; we will elaborate on this in the next section). After this, we write back the fully transformed rows from shared memory to device memory (line 8). Then, for each block, the bottom part of the shared memory is copied to the top part (Fig. 7b), in order to align the next window at the bottom of the previous one, taking the overlap area into account (line 11). The rest of the shared memory is filled again by copying rows from the current read pointer of the slab (line 12).

Further, we apply a vertical wavelet lifting transform (**transformBlock**, line 14) to the rows in the working area. This does not need to take boundary conditions into account as the top and bottom are handled specifically with **transformTop** and **transformBottom**. Then, height rows are copied from shared memory row *top* to the current write pointer (line 15). This process is repeated until we have written back the entire slab, except for the last leftover part. When finishing up (line 20), we have to be careful to satisfy the bottom boundary condition.

4.6.2 Example

We will discuss the Deslauriers-Dubuc (13, 7) wavelet as an example [22]. This example was chosen because it represents a nontrivial, but still compact enough case of the algorithm, that we can go through step by step. The filter weights for the two lifting steps of this transform are shown in Table 2. Both the prediction and update steps depend on two coefficients before and after the signal element to be computed. Fig. 8 shows an example of the performed computations. For this example, we choose *top* = 3,

Fig. 8. The vertical pass for the Deslauriers-Dubuc (13, 7) [22] wavelet. Lifted rows in each step are marked with a cross, and dependent rows are marked with a bullet.

overlap = 2, *height* = 8, and *bottom* = 3. This is a toy example, as, in practice, *height* will be much larger when compared to the other parameters.

Starting with the first window at the start of the data set, step 1 (first column), the odd rows of the working area (offset 1, 3, 5, 7) are lifted. The lifted rows are marked with a cross, and the rows they depend on are marked with a bullet. In step 2 (second column), the even rows are lifted. Again, the lifted rows are marked with a cross, and the dependencies are marked with a bullet. As the second step is dependent on the first, we cannot lift any rows that are dependent on values that were not yet calculated in the last step. In Fig. 8, this would be the case for row 6: this row requires data in rows 3, 5, 7, and 9, but row 9 is not yet available.

Here, the *overlap* region of rows comes in. As row 6 of the window is not yet fully transformed, we cannot write it back to device memory yet. So, we write everything up to this row back, copy the overlapping area to the top, and proceed with the second window. In the second window, we again start with step 1. The odd rows are lifted, except for the first one (offset 7) which was already computed, i.e., rows 9, 11, 13, and 15 are lifted. Then, in step 2, we start at row 6, i.e., three rows before the first step (row 9), but we do lift four rows.

After this, we can write the top eight rows back to device memory and begin with the next window in exactly the same way. We repeat this until the entire data set is transformed. By letting the second lifting step lag behind the first, one can do the same number of operations in each, making optimal use of the thread matrix (which should have a height of 4 in this case).

All separable wavelet lifting transforms, even those with more than two lifting steps, or with differently sized supports, can be computed in the same way. The transform can be inverted by running the steps in reverse order and flipping the signs of the filter weights.

4.7 3D and Higher Dimensions

The reason that the horizontal and vertical passes are asymmetric is because of the coalescing requirement for reads and writes. In the horizontal case, an entire line of the data set could be processed at a time. In the vertical case, the data set was horizontally split into image-high slabs. This allowed the slabs to be treated independently and processed using a sliding window algorithm that uses

coalesced reads and writes to access lines of the slab. A consecutive, horizontal span of values is stored at consecutive addresses in memory. This does not extend similarly to vertical spans of values, these will be separated by an offset at least the width of the image, known as the *row pitch*. As a slab is a rectangular region of the image of a certain width that spans the height of the image, it will be represented in memory by an array of consecutive spans of values, each separated by the *row pitch*.

When adding an extra dimension, let us say z , the volume is stored as an array of slices. In a span of values oriented along this dimension, each value is separated in memory by an offset that we call the *slice pitch*. By orienting the slabs in the xz -plane instead of the xy -plane, and thus, using the *slice pitch* instead of the *row pitch* as offset between consecutive spans of values, the same algorithm as in the vertical case can be used to do a lifting transform along this dimension. To verify our claim, we implemented the method just described, and report results in Section 5.2.7. More than three dimensions can be handled similarly, by orienting the slabs in the $D_i x$ plane (where D_i is the dimension i) and using the pitch in that dimension instead of the row pitch.

5 RESULTS

We first present a broad collection of experimental results. This is followed by a performance analysis which provides insight in the results obtained, and also shows that the design choices we made closely match our theoretical predictions.

The benchmarks in this section were run on a machine with a AMD Athlon 64 X2 Dual Core Processor 5200+ and a NVidia GeForce 8800 GTX 768MB graphics card, using CUDA version 2.1 for the CUDA programs. All reported timings exclude the time needed for reading and writing images or volumes from and to disc (both for the CPU and GPU versions).

5.1 Wavelet Filters Used for Benchmarking

The wavelet filters that we used in our benchmarks are integer-to-integer versions (unnormalized) of the Haar [4], Deslauriers-Dubuc (9, 7) [22], Deslauriers-Dubuc (13, 7) [22], Le Gall (5, 3) [23] (integer approximation of) Daubechies (9, 7) [2], and the Fidelity wavelet—a custom wavelet with a large support [24]. In the filter naming convention (m, n) , m refers to the length of the analysis low-pass and n to the analysis high-pass filters in the conventional wavelet subband filtering model, in which a convolution is applied before subsampling. They do not reflect the length of the filters used in the lifting steps, which operate in the subsampled domain. The implementation only involves integer addition and multiplication, and integer division by powers of 2 (bit-shifting) (cf. Section 3.2). The coefficients of the lifting filters can be found in [24].

5.2 Experimental Results and Comparison to Other Methods

5.2.1 Comparison of 2D Wavelet Lifting, GPU versus CPU

First, we emphasize that the accuracies of the GPU and CPU implementations are the same. Because only integer operations are used (cf. Section 5.1), the results are identical.

We compared the speed of performing various wavelet transforms using our optimized GPU implementation, to an optimized wavelet lifting implementation on the CPU, called Schrödinger [24]. The latter implementation makes use of vectorization using the MMX and SSE instruction set extensions, thus can be considered close to the maximum that can be achieved on the CPU with one core.

Table 3 shows the timings of both our GPU-accelerated implementation and the Schrödinger implementation when computing a three-level transform with various wavelets of a $1,920 \times 1,080$ image consisting of 16-bit samples. As it is better from an optimization point of view to have a tailored kernel for each wavelet type than to have a single kernel that handles everything, we used a code generation approach to create specific kernels for the horizontal and vertical pass for each of the wavelets. Both the analysis (forward) and synthesis (inverse) transforms are benchmarked. We observe that speedups by a factor of 10-14 are reached, depending on the type of wavelet and the direction of the transform. The speedup factor appears to be roughly proportional to the length of the filters. The Haar wavelet is an exception, since the overlap problem does not arise in this case (the filter length being just 2), which explains the larger speedup factor.

To demonstrate the importance of coalesced memory access in CUDA, we also performed timings using a trivial CUDA implementation of the Haar wavelet that uses the same algorithm for the vertical step as for the horizontal step, instead of our sliding window algorithm. Note that this method can be considered an improved (using cache) row-column, hardware-based strategy, see Section 2. While our algorithm processes an image in 0.80 milliseconds, the trivial algorithm takes 15.23, which is almost 20 times slower. This is even slower than performing the transformation on the CPU.

Note that the timings in Table 3 do not include the time required to copy the data from (2.4 ms) or to (1.6 ms) the GPU.

5.2.2 Vertical Step via Transpose Method

Another method that we have benchmarked consists in reusing the horizontal step as vertical step by using a “transpose” method. Here, the matrix of wavelet coefficients is transposed after the horizontal pass, the algorithm for the horizontal step is applied, and the results are transposed back. The results are shown in columns 3 and 4 of Table 3. Even though the transpose operation in CUDA is efficient and coalescable and this approach is much easier to implement, the additional passes over the data reduce performance quite severely. Another drawback of this method is that transposition cannot be done in-place efficiently (in the general case), which doubles the required memory, so that the advantage of using the lifting strategy is lost.

5.2.3 Comparison of Horizontal and Vertical Steps

Table 4 shows separate benchmarks for the horizontal and vertical steps, using various wavelet filters. From these results, one can conclude that the vertical pass is not significantly slower (and in some cases even faster) than the horizontal pass, even though it performs more elaborate cache management, see Algorithm 1.

TABLE 3

Performance of Our CUDA GPU Implementation of 2D Wavelet Lifting (Column 5) Compared to an Optimized CPU Implementation (Column 2) and a CUDA GPU Transpose Method (Column 3, See Text), Computing a Three-Level Decomposition of a $1,920 \times 1,080$ Image for Both Analysis and Synthesis Steps

Wavelet (analysis)	CPU (ms)	GPU transpose (ms)	Speed-up	GPU our method (ms)	Speed-up
Haar	10.31	5.58	1.9	0.80	12.9
Deslauriers-Dubuc (9, 7)	16.84	6.01	2.8	1.50	11.2
Le Gall (5, 3)	14.03	5.89	2.4	1.34	10.5
Deslauriers-Dubuc (13, 7)	19.52	6.08	3.2	1.62	12.0
Daubechies (9, 7)	22.66	6.54	3.5	2.05	11.1
Fidelity	28.82	6.45	4.5	2.11	13.7
Wavelet (synthesis)	CPU (ms)	GPU transpose (ms)	Speed-up	GPU our method (ms)	Speed-up
Haar	9.11	6.33	1.4	0.83	11.0
Deslauriers-Dubuc (9, 7)	15.93	6.40	2.5	1.45	11.0
Le Gall (5, 3)	13.02	6.29	2.1	1.28	10.2
Deslauriers-Dubuc (13, 7)	18.22	6.48	2.8	1.55	11.8
Daubechies (9, 7)	21.73	7.03	3.1	2.04	10.7
Fidelity	27.21	6.86	4.0	2.18	12.5

5.2.4 Timings for 16-Bit versus 32-Bit Integers

We also benchmarked an implementation that uses 32-bit integers, see Table 5. For small wavelets like Haar, the timings for 16- and 32-bit differ by a factor of around 1.5, whereas for large wavelets, the two are quite close. This is probably because the smaller wavelet transforms are more memory-bound and the larger wavelets are more compute-bound; hence, the increased memory bandwidth does not affect the performance significantly.

5.2.5 Comparison of 2D Wavelet Lifting on GPU, CUDA versus Fragment Shaders

We also implemented the algorithm of Tenllado et al. [12] for wavelet lifting using conventional fragment shaders and performed timings on the same hardware. A three-level Daubechies (9, 7) forward wavelet transform was applied to a $1,920 \times 1,080$ image, which took 5.99 milliseconds. In comparison, our CUDA-based implementation (see Table 3) does the same in 2.05 milliseconds, which is about 2.9 times faster. This speedup probably occurs because our method effectively makes use of CUDA shared memory to compute

intermediate lifting steps, conserving GPU memory bandwidth, which is the bottleneck in the Tenllado method. Another drawback that we noticed while implementing the method is that an important advantage of wavelet lifting, i.e., that it can be done in place, appears to have been ignored. This is possibly due to an OpenGL restriction by which it is not allowed to use the source buffer as destination, the same result is achieved by alternating between two halves of a buffer, resulting in a doubling of memory usage.

Fig. 9 further compares the performance of the Schrödinger CPU implementation [12] and our CUDA-accelerated method. A three-level Daubechies (9, 7) forward wavelet decomposition was applied to images of different sizes, and the computation time was plotted versus image size in a log-log graph. This shows that our method is faster by a constant factor, regardless of the image size. Even for smaller images, our CUDA-accelerated implementation is faster than the CPU implementation, whereas the shader-based method of Tenllado is slower for 256×256 images, due to OpenGL rendering and state setup overhead. CUDA kernel calls are relatively lightweight, so this problem does not arise in our approach. For larger images, the overhead

TABLE 4

Performance of Our GPU implementation on 16-Bit Integers, Separate Timings of Horizontal and Vertical Steps on a One-Level Decomposition of a $1,920 \times 1,080$ Image

Wavelet (analysis)	Horizontal (ms)	Vertical (ms)
Haar	0.26	0.19
Deslauriers-Dubuc (9, 7)	0.44	0.42
Le Gall (5, 3)	0.39	0.34
Deslauriers-Dubuc (13, 7)	0.47	0.47
Daubechies (9, 7)	0.62	0.62
Fidelity	0.63	0.76
Wavelet (synthesis)	Horizontal (ms)	Vertical (ms)
Haar	0.29	0.19
Deslauriers-Dubuc (9, 7)	0.39	0.44
Le Gall (5, 3)	0.35	0.36
Deslauriers-Dubuc (13, 7)	0.42	0.48
Daubechies (9, 7)	0.58	0.79
Fidelity	0.59	0.64

TABLE 5

Performance of Our GPU Implementation on 16 versus 32-Bit Integers (Three-Level Transform, $1,920 \times 1,080$ Image)

Wavelet (analysis)	16-bit (ms)	32-bit (ms)
Haar	0.80	1.09
Deslauriers-Dubuc (9, 7)	1.50	1.64
Le Gall (5, 3)	1.34	1.45
Deslauriers-Dubuc (13, 7)	1.62	1.75
Daubechies (9, 7)	2.05	2.13
Fidelity	2.11	2.72
Wavelet (synthesis)	16-bit (ms)	32-bit (ms)
Haar	0.83	1.15
Deslauriers-Dubuc (9, 7)	1.45	1.81
Le Gall (5, 3)	1.28	1.66
Deslauriers-Dubuc (13, 7)	1.55	1.90
Daubechies (9, 7)	2.04	2.35
Fidelity	2.18	2.80

Fig. 9. Computation time versus image size for various lifting implementations; three-level Daubechies (9,7) forward transform. (a) The Schrödinger CPU implementation [12] and our CUDA-accelerated method in a log-log plot. (b) Just the two GPU methods in a linear plot.

averages out, but as the method is less bandwidth efficient, it remains behind by a significant factor.

5.2.6 Comparison of Lifting versus Convolution in CUDA

Additionally, we compared our method to a convolution-based wavelet transform implemented in CUDA, one that uses shared memory to perform the convolution plus downsampling (analysis), or upsampling plus convolution (synthesis) efficiently. On a $1,920 \times 1,080$ image, for a three-level transform with the Daubechies (9,7) wavelet, the following timings are observed: 3.4 ms for analysis and 5.0 ms for synthesis. The analysis is faster than the synthesis because it requires less computations—only half of the coefficients have to be computed, while the other half are discarded in the downsampling step. Compared to the 2.0 ms of our own method for both transforms, this is significantly slower. This matches the expectation that a speedup factor of 1.5-2 can be achieved when using lifting [4].

5.2.7 Timings for 3D Wavelet Lifting in CUDA

Timings for the 3D approach outlined in Section 4.7 are given in Table 6. A three-level transform was applied to a 512^3 volume, using various wavelets. The timings are compared to the same CPU implementation as before,

TABLE 6
Performance of Our GPU Lifting Implementation in 3D, Compared to an Optimized CPU Implementation; a Three-Level Decomposition for Both Analysis and Synthesis Is Performed on a 512^3 Volume

Wavelet (analysis)	CPU (ms)	GPU (ms)	Speed-up
Haar	1037.4	147.2	7.0
Deslauriers-Dubuc (9, 7)	2333.1	192.0	12.2
Le Gall (5, 3)	1636.2	179.3	9.1
Deslauriers-Dubuc (13, 7)	3056.1	200.5	15.2
Daubechies (9, 7)	3041.3	234.6	13.0
Fidelity	5918.4	239.2	24.7
Wavelet (synthesis)	CPU (ms)	GPU (ms)	Speed-up
Haar	926.5	150.9	6.1
Deslauriers-Dubuc (9, 7)	2289.9	184.7	12.4
Le Gall (5, 3)	1631.1	173.7	9.4
Deslauriers-Dubuc (13, 7)	2983.9	192.1	15.5
Daubechies (9, 7)	2943.5	232.0	12.7
Fidelity	5830.7	230.9	25.3

extended to 3D. The numbers show that the speedups that can be achieved for higher dimensional transforms are considerable, especially for the larger wavelets such as Deslauriers-Dubuc (13, 7) or Fidelity.

5.2.8 Summary of Experimental Results

Compared to an optimized CPU implementation, we have seen performance gains of up to nearly 14 times for 2D and up to 25 times for 3D images by using our CUDA-based wavelet lifting method. Especially for the larger wavelets, the gains are substantial. When compared to the trivial transpose-based method, our method came out about two times faster over the entire spectrum of wavelets. When regarding computation time versus image size, our GPU-based wavelet lifting method was measured to be the fastest of three methods for all image sizes, with the factor mostly independent of the image size.

5.3 Performance Analysis

We analyze the performance of our GPU implementation, according to the metrics from Section 4.2, for performing one lifting (analysis) step. Without loss of generality, we discuss the Deslauriers-Dubuc (13, 7) wavelet (cf. Section 4.6). Our systematic approach consists first in explaining the total execution time, throughput, and bandwidth of our method, and then, in discussing the design decisions we made. The overhead of data transfer between CPU and GPU was excluded, since the wavelet transform is usually part of a larger processing pipeline (such as a video codec), of which multiple steps can be carried out on the GPU.

5.3.1 Horizontal Step

The size of the input data set is $N = w \cdot h = 1,920 \cdot 1,080$ two-byte words. We set $T = 256$ threads per block, and given the number of registers and the size of the shared memory used by our kernel, NVidia's occupancy calculator indicates that $k = 3$ blocks are active per MP, such that each MP is fully occupied (i.e., $k \cdot T = 768$ threads will be scheduled); the number of thread blocks for the horizontal step is $B = 1,080$. Given that the 8800 GTX

GPU has $M = 16$ MPs, it follows that $\alpha = 23$, see Section 4.2. Further, we used *decuda* (a disassembler of GPU binaries; see <http://wiki.github.com/laanwj/decuda>) to count the number and type of instructions performed. After unrolling the loops, we found that the kernel has 309 instructions, 182 of which are arithmetic operations in local memory and registers, 15 instructions are half-width (i.e., instruction code is 32-bit wide), 82 are memory transfers, and 30 are other instructions (mostly type conversions). Assuming that half-width instructions have a throughput of two cycles, and others take four cycles per warp, and since the clock rate of this device is $K = 1.35$ GHz, the asymptotic execution time is $T_e = 0.48$ ms. Here, we assumed that the extra overhead due to rescheduling is negligible, as was confirmed by our experiments.

For the transfer time, we first computed the ratio of arithmetic to arithmetic-and-transfers instructions, which is $r = 0.67$. Thus, from (5), it follows that as many as 301 cycles can be spared due to latency hiding. As the amount of shared memory used by the kernel is relatively small (i.e., 3×3.75 KB used out of 16 KB per MP) and the size of the L2 cache is about 12 KB per MP [20], we can safely assume that the latency of a global memory access is about 350 cycles so that $l_m = 49$ cycles. Since $m = 4$ (i.e., two 2-byte words are coalesced), the transfer time is $T_m = 0.15$ ms. Note that as two MPs also share a small but faster L1 cache of 1.5 KB, the real transfer time could be even smaller than our estimate. Moreover, as we also included in our counting shared memory transfers (whose latency is at least 10 times smaller than that of global memory), the real transfer time should be much smaller than its estimate.

According to our discussion in Section 4.5, five synchronization points are needed to ensure data consistency between individual steps. For one barrier, in the ideal case, the estimated waiting time is $T_s = 1.65 \mu\text{s}$; thus, the total time is about $8.25 \mu\text{s}$. In the worst case, $T_s = 0.2$ ms so that the total time can be as large as 1 ms.

To summarize, the estimated execution time for the horizontal step is about $T_t = 0.63$ ms, neglecting the synchronization time. Comparing this result with the measured one from Table 4, one sees that the estimated total time is 0.16 ms larger than the measured one. Probably, this is due to L1 caching contributing to a further decrease of T_m . However, It is essential that the total time is dominated by the execution time, indicating a compute-bound kernel. As the timing remains essentially the same (cf. Tables 3 and 5) when switching from 2-byte words to 4-byte words data, this further strengthens our finding.

The measured throughput is $G_m = 98$ Gflop/s, whereas the estimated one is $G_e = 104$ Gflop/s, indicating, on average, an instruction throughput of about 100 Gflop/s. Note that with some abuse of terminology, we refer to flops, when, in fact, we mean arithmetic instructions on integers. The measured bandwidth is $M_b = 8.8$ GB/s, i.e., we are quite far from the pin-bandwidth (86 GB/s) of the GPU; thus, one can conclude again that our kernel is indeed compute-bound. This conclusion is further supported by the fact that the flop-to-byte ratio of the GPU is 5, while in our case, this ratio is about 11. The fact that the kernel does not achieve the maximum throughput (using shared

memory) of about 230 Gflop/s is most likely due to the fact that the synchronization time cannot simply be neglected and seems to play an important role in the overall performance.

Let us now focus on the design choices we have made. Using $T = 256$ threads per block amounts to optimal time slicing (latency hiding), see discussion above and in Section 4.2, while we are still able to coalesce memory transfers. To decrease the synchronization time, lighter threads are suggested implying that their number should increase, while maintaining a fixed size of the problem. NVidia's performance guidelines [19] suggest that the optimal number of threads per block should be a multiple of 64. The next higher than 256 multiple of 64 is 320. Unfortunately, using 320 threads per block means that at most two blocks can be allocated to one MP, and thus, the MP will not be fully occupied. This, in turn, implies that an important amount of idle cycles spent on memory transfers cannot be saved, rendering the method less optimal with respect to time slicing. Accordingly, our choice of $T = 256$ threads per block is optimal. Further, our choice on the number of blocks also fulfills NVidia's guidelines with respect to current and future GPUs, see [19].

5.3.2 Vertical Step

While conceptually more involved than the horizontal step, the overall performance figure for the vertical step is rather similar to the horizontal one. The CUDA configuration for this kernel is as follows: Each 2D thread block contains a number of $16 \times 8 = 128$ threads, while the number of columns within each slab is $S = 32$, see Fig. 6. Thus, since the input consists of 2-byte words, each thread performs coalesced memory transfers of $m = 4$ bytes, similar to the horizontal step. As the number of blocks is $w/S = 60$, $k = 4$ (i.e., four blocks are scheduled per MP) and the kernel takes 39,240 cycles per warp to execute, the execution time for the vertical step is $T_e = 0.46$.

Unlike the horizontal step, now $r = 0.83$ so that no less than 352 cycles can be spared in global memory transaction. Note that when computing r , we only counted global memory transfers, as in this case, more, much faster shared memory transfers take place, see Algorithm 1. As the shared memory usage is only 4×1.8 KB, this suggests that the overhead due to slow accesses to global memory can be neglected so that the transfer time T_m can be neglected. The waiting time is $T_s = 0.047 \mu\text{s}$, and there are 344 synchronization points for the vertical step kernel so that the total time is about $15.6 \mu\text{s}$. In the worst case, this time can be as large as 1.9 ms. Thus, as $T_t = 0.46$ (without waiting time), our estimate is very close to the measured execution time from Table 4—this being, in turn, the same as that of the horizontal step. Finally, both the measured and estimated throughputs are comparable to their counterparts of the horizontal step.

Note that compared to the manually tuned, optimally designed matrix multiplication algorithm of [20] which is able to achieve a maximum throughput of 186 Gflop/s, the performance of 100 Gflop/s of our lifting algorithms may not seem impressive. However, one should keep in mind that matrix multiplication is much easier to parallelize efficiently, as it requires little synchronization. Unlike

matrix multiplication, the lifting algorithm requires a lot more synchronization points to ensure data consistency between steps, as the transformation is done in-place.

The configuration we chose for this kernel is $16 \times 8 = 128$ threads per block and $w/S = 60$ thread blocks. This results in an occupancy of 512 threads per MP, which may seem less optimal. However, to increase the number of threads per block to 192 (next larger multiple of 64, see above) would mean that either we cannot perform essential, coalesced memory accesses, or that extra overhead due to the requirements of the moving-window algorithm would have to be accommodated. Note that we verified this possibility, but the results were unsatisfactory.

5.3.3 Complexity

Based on the formulas in Section 4.2, we can analyze the complexity of our problem. For any of the lifting steps using the Deslauriers-Dubuc (13, 7) wavelet, considering that the number of flops per data element is $n_s = 22$ (20 multiply or additions and two register shifts to increase accuracy), the numerator of (9) becomes about 700 D . For the horizontal step, $D = w/T = 7.5$ so that the numerator becomes about 5,000. In this case, the number of cycles is about 1,250 so that one can conclude that the horizontal step is indeed *cost-efficient*. For the vertical step, $D = (S h)/T = 270$ so that the numerator in (9) becomes about 190,000, while the denominator is 39,240. Thus, the vertical step is also cost-efficient, and actually, its performance is similar to that of the horizontal step (because $5,000/1,250 \approx 190,000/39,240 \approx 5$). Of course, this result was already obtained experimentally, see Table 4. Note that using vectorized MMX and SSE instructions, the *optimized* CPU implementation (see Table 3) can be up to four times faster than our T_s estimate above. However, even in this case, both our CUDA kernels are still cost-efficient. Obviously, both steps are also *work-efficient*, as their CUDA realizations do not perform asymptotically more operations than the sequential algorithm.

6 CONCLUSION

We presented a novel, fast wavelet lifting implementation on graphics hardware using CUDA, which extends to any number of dimensions. The method tries to maximize coalesced memory access. We compared our method to an optimized CPU implementation of the lifting scheme, to another (non-CUDA-based) GPU wavelet lifting method, and also to an implementation of the wavelet transform in CUDA via convolution. We implemented our method both for 2D and 3D data. The method is scalable and was shown to be the fastest GPU implementation among the methods considered. Our theoretical performance estimates turned out to be in fairly close agreement with the experimental observations. The complexity analysis revealed that our CUDA kernels are cost- and work-efficient.

Our proposed GPU algorithm can be applied in all cases where the Discrete Wavelet Transform based on the lifting scheme is part of a pipeline for processing large amounts of data. Examples are the encoding of static images, such as the wavelet-based successor to JPEG, JPEG2000 [25], or video coding schemes [24], which we already considered in [26].

ACKNOWLEDGMENTS

The authors thank the anonymous reviewers for several helpful suggestions to improve their paper. This research is part of the "VIEW: Visual Interactive Effective Worlds" program, funded by the Dutch National Science Foundation (NWO), project no. 643.100.501.

REFERENCES

- [1] M.A. Westenberg and J.B.T.M. Roerdink, "Frequency Domain Volume Rendering by the Wavelet X-Ray Transform," *IEEE Trans. Image Processing*, vol. 9, no. 7, pp. 1249-1261, July 2000.
- [2] I. Daubechies, *Ten Lectures on Wavelets*, vol. 61. Soc. for Industrial and Applied Math., 1992.
- [3] S. Mallat, *A Wavelet Tour of Signal Processing*. Academic Press, 1998.
- [4] W. Sweldens, "The Lifting Scheme: A Construction of Second Generation Wavelets," *SIAM J. Math. Analysis*, vol. 29, no. 2, pp. 511-546, 1998.
- [5] I. Daubechies and W. Sweldens, "Factoring Wavelet Transforms into Lifting Steps," *J. Fourier Analysis and Applications*, vol. 4, no. 3, pp. 247-269, 1998.
- [6] A.R. Calderbank, I. Daubechies, W. Sweldens, and B.-L. Yeo, "Wavelet Transforms That Map Integers to Integers," *Applied and Computational Harmonic Analysis*, vol. 5, no. 3, pp. 332-369, 1998.
- [7] E. Lindholm, J. Nickolls, S. Oberman, and J. Montrym, "NVIDIA Tesla: A Unified Graphics and Computing Architecture," *IEEE Micro*, vol. 28, no. 2, pp. 39-55, Mar./Apr. 2008.
- [8] M. Hopf and T. Ertl, "Hardware Accelerated Wavelet Transformations," *Proc. EG/IEEE TCVG Symp. Visualization (VisSym '00)*, pp. 93-103, May 2000.
- [9] A. Garcia and H.W. Shen, "GPU-Based 3D Wavelet Reconstruction with Tileboarding," *The Visual Computer*, vol. 21, pp. 755-763, 2005.
- [10] T.T. Wong, C.S. Leung, P.A. Heng, and J. Wang, "Discrete Wavelet Transform on Consumer-Level Graphics Hardware," *IEEE Trans. Multimedia*, vol. 9, no. 3, pp. 668-673, Apr. 2007.
- [11] C. Tenllado, R. Lario, M. Prieto, and F. Tirado, "The 2D Discrete Wavelet Transform on Programmable Graphics Hardware," *Proc. Fourth IASTED Int'l Conf. Visualization, Imaging, and Image Processing*, 2004.
- [12] C. Tenllado, J. Setoain, M. Prieto, L. Piñuel, and F. Tirado, "Parallel Implementation of the 2D Discrete Wavelet Transform on Graphics Processing Units: Filter Bank versus Lifting," *IEEE Trans. Parallel and Distributed Systems*, vol. 19, no. 3, pp. 299-310, Mar. 2008.
- [13] W. Jiang and A. Ortega, "Parallel Architecture for the Discrete Wavelet Transform Based on the Lifting Factorization," *J. Parallel and Distributed Computing*, vol. 57, no. 2, pp. 257-269, 1999.
- [14] M. Angelopoulou, K. Masselos, P. Cheung, and Y. Andreopoulos, "Implementation and Comparison of the 5/3 Lifting 2d Discrete Wavelet Transform Computation Schedules on FPGAs," *J. Signal Processing Systems*, vol. 51, no. 1, pp. 3-21, 2008.
- [15] N.D. Zervas, G.P. Anagnostopoulos, V. Spiliotopoulos, and Y. Andreopoulos, "Evaluation of Design Alternatives for the 2D Discrete Wavelet Transform," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 11, no. 12, pp. 1246-1262, Dec. 2001.
- [16] M.Y. Chiu, K.-B. Lee, and C.-W. Jen, "Optimal Data Transfer and Buffering Schemes for JPEG 2000 Encoder," *Proc. IEEE Workshop Design and Implementation of Signal Processing Systems*, pp. 177-182, 2003.
- [17] T. Acharya and C. Chakrabarti, "A Survey on Lifting-Based Discrete Wavelet Transform Architectures," *J. VLSI Signal Processing Systems*, vol. 42, no. 3, pp. 321-339, 2006.
- [18] Nvidia, *CUDA Occupancy Calculator*, http://developer.download.nvidia.com/compute/cuda/CUDA_Occupancy_calculator.xls, 2009.
- [19] NVIDIA Corporation, *Compute Unified Device Architecture Programming Guide*, <http://developer.nvidia.com/cuda>, 2010.
- [20] V. Volkov and J. Demmel, "LU, QR and Cholesky Factorizations Using Vector Capabilities of GPUs," technical report, Univ. of California at Berkeley, 2008.
- [21] S. Chatterjee and C.D. Brooks, "Cache-Efficient Wavelet Lifting in JPEG 2000," *Proc. IEEE Int'l Conf. Multimedia and Expo*, pp. 797-800, 2002.

- [22] G. Deslauriers and S. Dubuc, "Symmetric Iterative Interpolation Processes," *Constructive Approximation*, vol. 5, no. 1, pp. 49-68, Dec. 1989.
- [23] D. LeGall and A. Tabatabai, "Sub-Band Coding of Digital Images Using Symmetric Short Kernel Filters and Arithmetic Coding Techniques," *Proc. IEEE Int'l Conf. Acoustics, Speech and Signal Processing*, vol. 2, pp. 761-764, 1988.
- [24] BBC Research, *Dirac Specification 1.0.0pre7*, <http://dirac.sourceforge.net/specification.html>, 2010.
- [25] A.N. Skodras, C.A. Christopoulos, and T. Ebrahimi, "JPEG2000: The Upcoming Still Image Compression Standard," *Pattern Recognition Letters*, vol. 22, no. 12, pp. 1337-1345, 2001.
- [26] W.J. van der Laan, A.C. Jalba, and J.B.T.M. Roerdink, "Accelerating Wavelet-Based Video Coding on Graphics Hardware Using CUDA," *Proc. Sixth Int'l Symp. Image and Signal Processing and Analysis (ISPA '09)*, pp. 614-619, Sept. 2009.

Wladimir J. van der Laan received the MSc degree in computing science in 2005 from the University of Groningen, where he is currently working toward the PhD degree at the Johann Bernoulli Institute for Mathematics and Computer Science. His research interests include computer graphics, visualization, and GPU computing.

processing, and parallel computing.

Andrei C. Jalba received the BSc and MSc degrees in applied electronics and information engineering from "Politehnica" University of Bucharest, Romania, in 1998 and 1999, respectively, and the PhD degree from the Institute for Mathematics and Computing Science, University of Groningen, in 2004. Currently, he is an assistant professor at the Eindhoven University of Technology. His research interests include computer vision, pattern recognition, image

struction from 1986 to 1992. He was appointed an associate professor in 1992 and a full professor in 2003, respectively, in the Institute for Mathematics and Computing Science at the University of Groningen, where he currently holds a chair in Scientific Visualization and Computer Graphics. His current research interests include biomedical visualization, neuroimaging, and bioinformatics. He is a senior member of the IEEE.

Jos B.T.M. Roerdink received the MSc degree in theoretical physics from the University of Nijmegen, The Netherlands, in 1979, and the PhD degree from the University of Utrecht in 1983. He was a postdoctoral fellow in the area of stochastic processes at the University of California, San Diego, from 1983 to 1985. He joined the Centre for Mathematics and Computer Science in Amsterdam. There he worked on image processing and tomographic recon-

► **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**