New Fast Wavelet Transform Methods

V3C

The discrete wavelet transform is a widely applied mathematical tool used among others fields in computer graphics and multimedia. With the start of its widespread use, the need for its efficient implementation is becoming increasingly more important. Our research is focused on novel implementations on traditional processing units (CPU), modern graphics cards (GPU), and other platforms (FPGA). We have created several vectorized and parallelized algorithms of discrete wavelet transform, mainly using a lifting scheme..

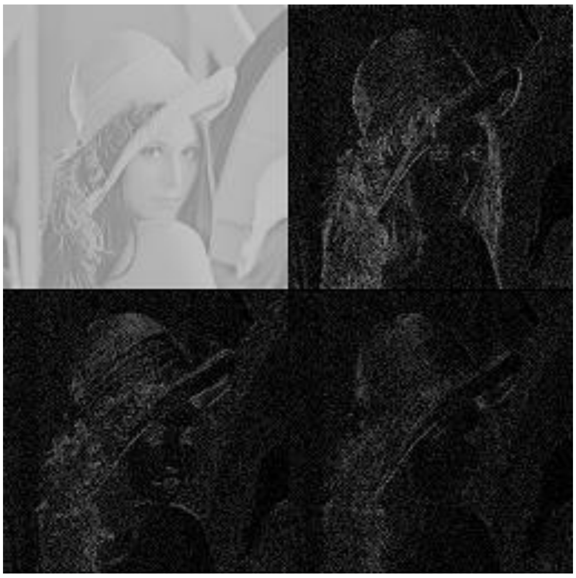

Figure 1

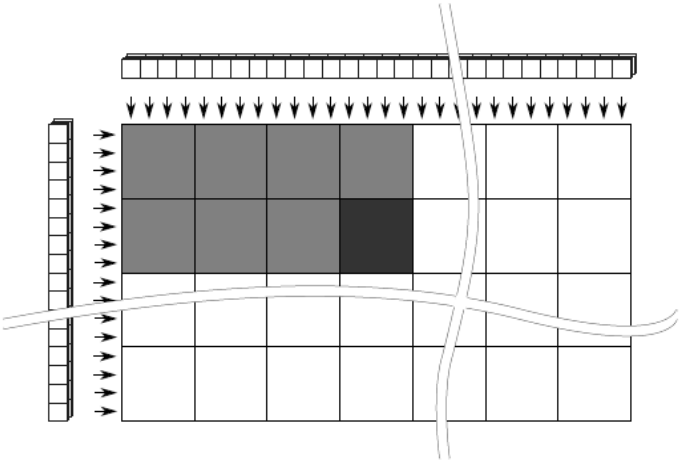

Figure 2

The image processing applications of the wavelet transform include tone mapping, compression, contrast enhancement, watermarking, edge detection, or denoising. Figure 1 demonstrates a single-level analysis of a grayscale test image (Lena image). You can easily see four resulting sub-bands of the wavelet decomposition.

Our research

Our research started with a simple 1-D signal processing. A naïve implementation of this transform suffers from the necessity of using several reads and writes of intermediate results, which slows down the computation. Of course, more efficient methods exist. Above these methods, we have designed a diagonal vectorization of lifting scheme. This new method processes pairs of coefficients one by one immediately when available. This strategy can be especially useful on systems equipped with a small CPU cache.

Further, we have developed several novel SIMD-vectorized and parallelized algorithms of 2-D discrete wavelet transform. At the beginning, a stand-alone core of an already known single-loop approach was extracted. This core was further simplified, reorganized, and finally vectorized. The method is friendly to CPU caches and the best of the proposed algorithms scale almost linearly with the number of threads. Out new method achieved a speedup of nearly 40 times using 4 threads compared to a naïve single-threaded implementation. Figure 2 shows the complete image processing using one of our methods. This method reads an input image per 4 times 4 blocks and instantly produces an output. Figure 3 shows the final comparison to the naïve single-threaded method.

Graphics cards

Considering the modern programmable GPUs, we have presented an improved version of an algorithm suitable to compute the 2-D discrete wavelet. Our approach is focused on utilization of parallel capabilities of modern GPUs. We reached an average speed-up at least 30 percents on tested graphics cards. We have also compared it with the tuned CPU implementation. One such a comparison can be seen in Figure 4.

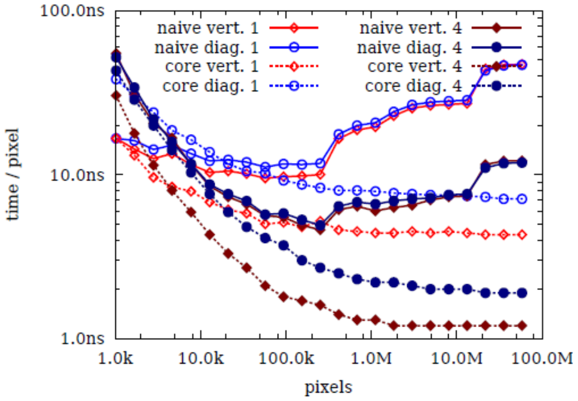

Figure 3

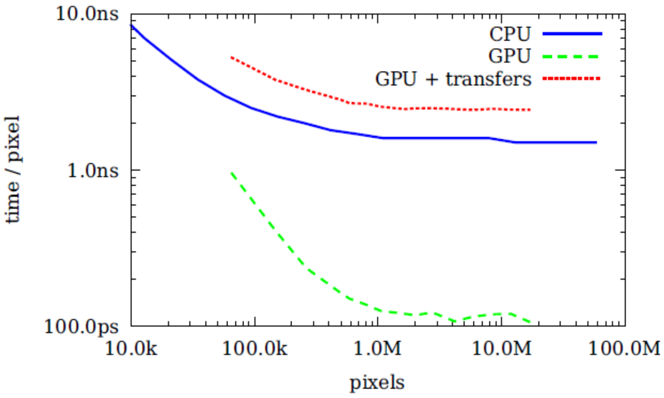

Figure 4

Title: New Fast Wavelet Transform Methods
Author: David Bařina (ibarina@fit.vutbr.cz)

Faculty of Information Technology
Brno University of Technology
Božetěchova 1/2
612 66 Brno, Czech Republic
Tel.: +420 54114-1144
E-Mail: info@fit.vutbr.cz

Research Project: V3C - Visual Computing Competence Center
No. TE01020415
www.v3c.cz
Project Manager: Pavel Zemčík (zemcik@fit.vutbr.cz)
Tel.: +420 603 487 824

Technologická agentura
České republiky

Centra
kompetence

Acknowledgment

This work was supported by the Technology Agency of the Czech Republic.