Parallel Numerics '02, 204-214 R. Trobec, P. Zinterhof, M. Vaiteršic, A. Uhl (Eds.)
Chapter 6: Video and Data Coding ISBN 961-6303-39-2

Optimization of Bitstream Assembly in Parallel Multimedia Compression

Rade Kutil *

Dept. of Scientific Computing
University of Salzburg
Jakob Haringer-Str. 2, A-5020 Salzburg, Austria

Multimedia compression algorithms usually take DCT or wavelet transformed input data and produce a stream of bits as output. It has been shown that compression algorithms can be parallelized in a data driven way so that a sequence of bit-stream parts is produced by each processing element (PE) corresponding to the PE-local part of the partitioned and distributed input data. Unfortunately, the collection and assembly of these bit-stream parts by a single PE turns out to be a major bottleneck because it is sequential. This work addresses this problem by sequential optimizations and parallelization of the assembly process itself.

1 Introduction

The most successful method for image and video compression is transform coding. The input data is usually given as two- or three-dimensional arrays of pixels. This data is transformed by a blocked discrete cosine transform (DCT) (as in JPEG [1] or MPEG [2, 3]) or by a discrete wavelet transform (as in JPEG2000 [4] or SPIHT [5]). The transformed data is then quantized and the resulting integer values are encoded bit by bit.

Although these algorithms are efficient in terms of rate-distortion performance as well as with respect to computational complexity, the amount of data that has to be processed causes prohibitive computational and memory demands (especially for real-time applications). Therefore, MIMD architectures seem to be an interesting choice for such algorithms.

*Corresponding author. E-mail: rkutil@cosy.sbg.ac.at

The parallelization of a multimedia compression scheme is divided into at least two phases: Parallel transform and parallel compression. Parallel transform is easy for blocked DCT, since DCT blocks do not interfere. A significant amount of work has been done on parallel wavelet transform [6, 7, 8]. The parallelization of the compression algorithm did not catch so much attention, although there exists parallel implementations of MPEG [9], JPEG [10, 11], JPEG2000 [12], EZW [13, 14] and SPIHT [15, 16].

The author's own investigations showed that the SPIHT algorithm can be fully parallelized except for the bit-stream assembly as the only sequential part. Although bit-stream assembly seems to be a simple and non-complex process, it turned out to be a major bottleneck in parallel SPIHT compression. This work is dedicated to the comparison of several approaches to efficient bit-stream assembly.

For obtaining experimental results, MPI implementations on a Cray T3E were employed which is situated at the High-Performance Computing-Center in Stuttgart and consists of 512 DEC Alpha EV5 processors (900 MFLOPS each) interconnected by a 3-D torus with 3 GB/s bandwidth in each network node. Results were also conducted on a shared memory machine (SGI Powerchallenge GR at RIST++, Salzburg Univ.) with 20 MIPS R10000 processors. However, these results are very similar to the Cray T3E results. Therefore, they are not shown in this paper. The parallel SPIHT algorithm is applied to 3-D monochrome video data with $864 \times 88 \times 72$ pixels. In all examples, a bit-rate of 0.09 bpp is chosen.

2 Separators

Most multimedia compression algorithms can be viewed as an abstract procedure which processes a set of primitive tasks in a certain order. For each task, a certain amount of bits is written into a bit-stream. Moreover, each task is associated with a certain approximate spatial position (i.e. a DCT-block or a wavelet). Hence, there exists a natural approach to data-driven parallelization, that is to distribute the source data among the PE's and to distribute the primitive tasks corresponding to the data distribution.

The problem is now to assemble the bit-stream parts produced by each PE both correctly and efficiently. A good way to do this is to use separators. Separators are dummy entries in the list of tasks which indicate a change of the PE the tasks belong to. So, each PE only has local tasks in its local part of the task list and inserts a separator wherever the global list continues with other PE's tasks. While tasks are processed, new tasks can be created and appended at the end of the task list. If the creation of these tasks depends only

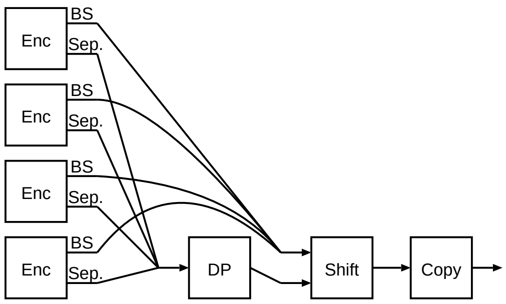

Figure 1: Method A. Enc = Encoding process (produces bit-stream parts and separators), BS = bit-stream, Sep. = separators, DP = calculation of destination positions, Shift = bit-alignment of bit-stream parts, Copy = merge bit-stream parts to a single bit-stream.

on local information (i.e. other tasks that belong to the same PE), then it is possible that each PE handles its separators correctly without any additional communication. When separators are processed as part of the task list, two things have to be done: First, a separator has to be inserted into the bit-stream. This is necessary to identify the bit-stream parts that have to be assembled in an alternating way at the end of the encoding process. Second, a separator has to be inserted at the end of the task list to separate newly created tasks correctly from newly created tasks of other PEs.

Experiments show that, very often, there are no actual tasks between separators. This is the case when many tasks do not initiate new tasks, so only separators are appended at the end of the task list. To reduce memory demands, it is necessary to keep such groups of separators together in a single entry – associated with a counter. For the bit-stream, separators are implemented as arrays of bit-pointers. Accordingly, equal pointers are kept together and are associated with a counter.

In this work, a parallel implementation [15, 16] of the well-known SPIHT algorithm [5, 17] is used. In this algorithm, the task list is split into three lists. Each task is associated with a single wavelet coefficient. Additionally, each list is reread when its end is reached so as to encode another bit-plane for higher precision. Entries that may not be processed anymore have to be removed.

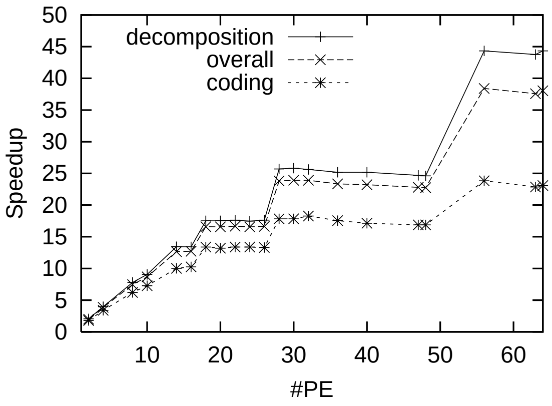

Figure 2: Speedups for method A

3 Method A: Sequential Assembly

The first and most easy-to-implement way of bit-stream assembly is to collect all bit-stream parts together with their separators and to put the parts together sequentially on a single PE. This process is visualized in Figure 1. The assembly itself is split into three blocks: The first one determines the destination position of the bit-stream parts in the output bit-stream. The second part shifts the bit-stream parts to the correct bit-position, since the bit-stream parts are not byte-aligned. The last block simply copies the parts together.

The first block - the calculation of the destination positions - can be implemented as follows:

$p \leftarrow 0$
while still separators left
  for k in $1 \ldots \#\mathrm{PE}$
    select new separator from PE_k
    $l \leftarrow$ difference to old separator from PE_k
    if $l > 0$
      output $(p, p + l)$ as destination positions
      $p \leftarrow p + l$

A first improvement of this algorithm is to reduce the number of iterations of the "while still separators left" loop. This can be done because, very often, subsequent separators are equal (and kept together in a single entry). If this is the case for all PE's separator lists at the beginning of the loop body, the minimum equal-separator count can be removed from all lists at once without corrupting the output. This modification reduces the time for bit-

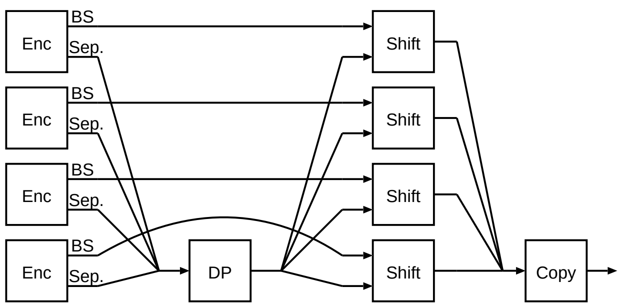

Figure 3: Method B

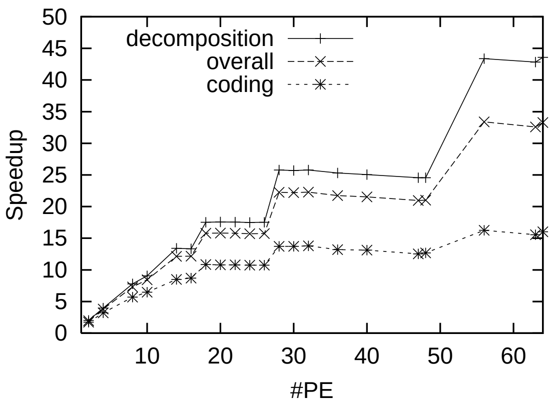

Figure 4: Speedups for method B

stream assembly (without collection) by about 40%.

Speedup results for the wavelet transform, the SPIHT encoding and the overall encoding process are shown in Figure 2. Although the speedups for the coding part are not as good as those of the parallel wavelet decomposition, the overall speedups are quite reasonable.

4 Method B: Parallel Shift

The next approach is to parallelize the shift operation because it is the most time consuming operation in the process of bit-stream assembly. To do this, the destination positions have to be distributed to enable the PEs to determine the correct bit-alignments of the bit-stream parts (see Figure 3). This additional communication step is the main disadvantage of this approach. Thus, the speedups are worse than for method A (see Figure 4).

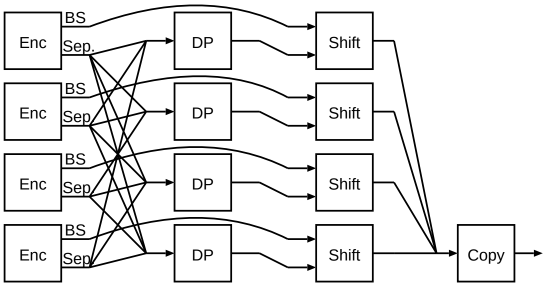

Figure 5: Method C

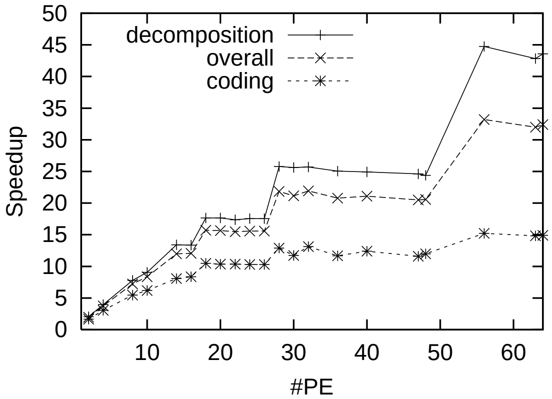

Figure 6: Speedups for method C

5 Method C: Redundant Calculation of Destination Positions

To avoid the additional communication step of distributing the destination positions, this method calculates the destination positions redundantly on every PE. Unfortunately, this involves n-to-n communication because each PE needs the full separator lists of each other PE. Therefore, the speedups of this approach are also degraded because of communication overhead as can be seen in Figure 6.

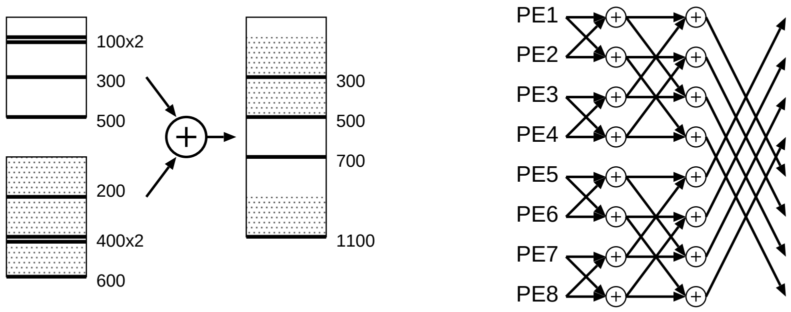

Figure 7: Calculation of destination positions in method D by an all-reduce operation

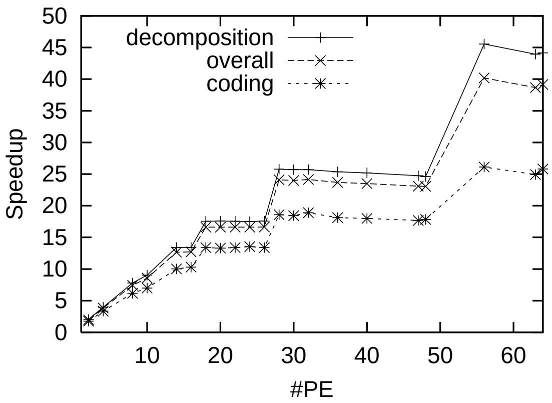

Figure 8: Speedups for method D

6 Method D: Parallel Calculation of Destination Positions

To calculate the destination positions for its own bit-stream parts, a PE does not really have to know all other's separator lists in detail. It would be sufficient to have a separator list of a fictive bit-stream in which all other PE's bit-stream parts are assembled. Such a separator list can easily be constructed by adding corresponding separators as in Figure 7.

With the help of this construction, it is possible to gather enough information on each PE by an **all-reduce** operation as shown in Figure 7. If the PE's are hypercube-connected, the communication can be performed in parallel as depicted in Figure 7, which makes this operation have logarithmic complexity.

The speedup results are shown in Figure 8. Comparing the results to those of method A (Figure 2), we see that they are approximately equal. This means that although we have parallelized the bit-stream assembly almost completely, the increased communication overhead makes the additional speedup potential vanish. However, given a better communication bandwidth and/or lower computational power of processor elements, method D could outperform method A.

7 Conclusions

Four different approaches to efficient bit-stream assembly in parallel multimedia were presented. While method B and C showed major drawbacks in performance due to significant communication overhead, methods A and D show similar performance results. The advantage of method A is that it is easy to implement. However, method D could outperform method A given better communication bandwidth and/or lower computational power of processor elements. This is because method D manages to parallelize the bit-stream assembly almost completely (speedups are primarily degraded by communication overheads) while method A is a purely sequential approach.

Acknowledgments

The author was supported by the Austrian Science Fund FWF, project no. P13903. The author also wants to thank the High-Performance Computing-Center (HLRS) in Stuttgart for granting access to their computing facilities.

References

[1] G.K. Wallace. The JPEG still picture compression standard. *Communications of the ACM*, 34(4):30–44, 1991.

[2] L. Chiariglione-Convenor. MPEG-2: Generic coding of moving pictures and associated audio information. ISO/IEC JTC1/SC29/WG11, July 1996.

[3] R. Koenen. Overview of the MPEG-4 standard. ISO/IEC JTC1/SC29/WG11, July 1996.

[4] ISO/IEC JPEG committee. JPEG 2000 image coding system — ISO/IEC 15444-1:2000, December 2000.

[5] Amir Said and William A. Pearlman. A new, fast, and efficient image codec based on set partitioning in hierarchical trees. *IEEE Transactions on Circuits and Systems for Video Technology*, 6(3):243–249, June 1996.

[6] R. Kutil and A. Uhl. Optimization of 3-d wavelet decomposition on multiprocessors. *Journal of Computing and Information Technology (Special Issue on Parallel Numerics and Parallel Computing in Image Processing, Video Processing, and Multimedia)*, 8(1):31–40, 2000.

[7] M-L. Woo. Parallel discrete wavelet transform on the Paragon MIMD machine. In R.S. Schreiber et al., editor, *Proceedings of the seventh SIAM conference on parallel processing for scientific computing*, pages 3–8, 1995.

[8] J. Fridman and E.S. Manolakos. On the scalability of 2D discrete wavelet transform algorithms. *Multidimensional Systems and Signal Processing*, 8(1–2):185–217, 1997.

[9] S.M. Akramullah, I. Ahmad, and M.L. Liou. A data-parallel approach for real-time MPEG-2 video encoding. *Journal of Parallel and Distributed Computing*, 30:129–146, 1995.

[10] G. W. Cook and Edward J. Delp. An investigation of JPEG image and video compression using parallel processing. In *Proceedings of the IEEE International Conference on Accoustics, Speech and Signal Processing, ICASSP '94*, pages 437–440, Adelaide, South Australia, Australia, April 1994.

[11] J.J. Falkemeier and G. Joubert. Parallel image compression with JPEG for multimedisa applications. In J.J. Dongarra et al., editors, *High Performance Computing: Technologies, Methods & Applications*, number 10 in Advances in Parallel Computing, pages 379–394. North Holland, 1995.

[12] Peter Meerwald, Roland Norcen, and Andreas Uhl. Parallel JPEG2000 image coding on multiprocessors. In *Proceedings of the International Parallel & Distributed Processing Symposium 2002 (IPDPS'02)*, page 2, Fort Lauderdale, FL, USA, April 2002. IEEE Computer Society Press.

[13] C.D. Creusere. Image coding using parallel implementations of the embedded zerotree wavelet algorithm. In B. Vasudev, S. Frans, and S. Panchanathan, editors, *Digital Video Compression: Algorithms and Technologies 1996*, volume 2668 of *SPIE Proceedings*, pages 82–92, 1996.

[14] F.W. Wheeler and W.A. Pearlman. Low-memory packetized SPIHT image compression. In *Proceedings of the Asilomar Conference on Signals, Systems, and Computers*, October 1999.

[15] R. Kutil. Zerotree based video coding on MIMD architectures. In S. Panchanathan, V. Bove, and S.I. Sudharsanan, editors, *Media Processors 2001*, volume 4313 of *SPIE Proceedings*, pages 61–68, January 2001.

[16] R. Kutil. Approaches to zerotree image and video coding on MIMD architectures. *Parallel Computing*, 28(7–8):1095–1109, 2002.

[17] B.J. Kim and W.A. Pearlman. An embedded wavelet video coder using three-dimensional set partitioning in hierarchical trees (SPIHT). In *Proceedings Data Compression Conference (DCC'97)*, pages 251–259. IEEE Computer Society Press, March 1997.