

Coarse grain parallelization of H.264 video decoder and memory bottleneck in multicore architectures

Ahmet Gurhanli^{a,1,*}, Charlie Chung-Ping Chen², Shih-Hao Hung³

^aAddress: 3F-1, No.24, Ln. 225, Sec. 1, Wenhua 3rd Rd., Linkou Township, Taipei County 244, Taiwan (R.O.C.), Tel: +886-927-679-623, Fax: +886-2-2362-6806

Abstract

Fine grain methods for parallelization of the H.264 decoder have good latency performance and less memory usage. However, they could not reach the scalability of Group of Pictures (GOP) level approaches although assuming a well-designed entropy decoder which can feed the increasing number of parallel working cores. We would like to introduce a GOP level approach due to its high scalability, mentioning solution approaches for the well-known latency and memory issues. Our design revokes the need to a scanner for GOP start-codes which was used in the earlier methods. This approach lets all the cores work on the decoding task. Our experiments showed that the memory operations may degrade the scalability of parallel applications substantially. The multicore cache architecture appeared to be a critical point for getting the desired speedup.

Keywords: video compression, H.264 decoder, parallel processing, high-performance computing, image processing

1. Introduction

Video encoding and decoding is among the tasks demanding very high performance computing. This demand keeps increasing when we consider the recent developments in video technologies like 3D TV and Ultra High Definition Video. The computation power of a home appliance may become insufficient to decode the huge amount of data needed by those applications.

After Instruction Level Parallelism approached to saturation, Thread Level Parallelism gained more and more importance day by day. Today chip multi-processors (CMPs) are wide-spread, and the number of cores in a CMP is expected to be doubled in every 3 year [1]. This phenomenon increases the importance of data-parallel algorithms for applications needing high computing power.

Recent works on video decoder parallelization preferred fine-grain methods, but they could not reach the scalability of coarse-grain schemes. In this paper, we introduce a GOP-level approach for H.264 video decoding. The main point of our design is that; it does not need a GOP start-code scanner. We accept that GOP-level parallelism needs more memory resources and has long latency problem. We will mention the possible solutions to these problems.

Our work showed that memory bandwidth in multicore architectures is the main bottleneck for highly scalable applications. Programs processing large amount of data has to deal with a huge amount of memory load-store operations. When each core has not a direct path to the memory for write and read operations the bandwidth becomes insufficient for parallel working cores. We will see the importance of efficient memory usage in the results section. In a program processing big amount of data, choosing a function inefficiently may cause the program run 2 times slower.

2. Structure of the H.264 decoder and parallelization opportunities

The system first entropy decodes the coming stream. There are two choices for entropy coding: CAVLC and CABAC. CABAC is available only in the main profile

*Corresponding author

Email addresses: ahmet@ntu.edu.tw (Ahmet Gurhanli),
cchen@cc.ee.ntu.edu.tw (Charlie Chung-Ping Chen),
hungsh@csie.ntu.edu.tw (Shih-Hao Hung)

¹National Taiwan University, Collage of Electrical Engineering and Computer Science, Graduate Institute of Electronics Engineering

²National Taiwan University, Collage of Electrical Engineering and Computer Science, Graduate Institute of Electronics Engineering

³National Taiwan University, Collage of Electrical Engineering and Computer Science, Graduate Institute of Networking and Multimedia

and uses more computations for achieving a better compression. Entropy decoding is mostly a sequential computation and hard to parallelize.

After entropy decoding the pictures are reordered and passed through inverse quantization. This is followed by inverse integer transform. Then, intra prediction or motion compensation is performed according to the picture type. Finally the picture is passed through a deblocking filter.

There are several alternatives for data level parallelism offered by the structure. Figure 1 shows the data structure in a video sequence. The sequence is first divided into GOPs, each having a certain number of frames. Frames are divided into slices of variable sizes. Slices consist of macroblocks each containing 4 luma and 2 chroma blocks. Each level is a candidate for partitioning the data to be decoded.

Ref. [7] gives a good summary of the parallelization methods. The highest level is the Group of Pictures level parallelism. The video stream is divided into groups of pictures for enabling video control functions and synchronization. GOP level parallelism provides high scalability but requires more memory resources. Since there are no dependencies between GOPs, the threads do not need to wait for synchronization operations in this approach.

In frame-level parallelism, each frame is handled by one thread. The main problem of this approach is the frame referencing. In H.264 all I, P, and B frames might be referenced by others. This flexibility makes frame level parallelism hard to implement.

Another choice is slice level parallelism. In order to ensure error resilience, frames are divided into slices, which are independent from each other. So slices might be processed in parallel. However, scalability of this approach is limited, because the number of slices is decided by the encoder. Besides increasing slice number results in increase in the bit rate.

Recently, most popular approach is the macroblock level parallelism. In this approach macroblocks are processed in parallel after dependencies have been resolved. Major problem of this approach is its dependence on a high performance entropy decoder.

Finally, the finest approach is block level parallelism. Block level parallelism might be used with SIMD instructions, for operations like deblocking, IDCT and interpolations which are done at block level.

3. Parallelization strategy

In a highly scalable model, dependencies between parallel threads must be as low as possible and sequen-

tial section of the algorithm must be minimized.

Another important issue is efficient resource usage. Manager-worker architectures allocates one core in the system for task scheduling, and only P-1 of P processors can work on the parallel decoding task. So, in a 4 core CMP we may only expect a speed up of 3.

Most of the parallel decoders we mentioned in related work part, uses a scheduling component. Besides, allocating one core for scheduling issue, this adds a sequential section to the overall decoding task. The scheduler may easily become insufficient for feeding the increasing number of parallel working cores.

Scalability of a fine grain scheme is limited due to the dependencies among the macroblocks. Before starting to decode a macroblock the system needs to wait for all other referenced macroblocks be decoded. Another limit for the fine-grain methods is the speed of the entropy decoder. A separate core should be employed in entropy decoding, and this sequential decoding must be very fast for feeding lots of parallel macroblock decoders.

Here we present a GOP-level scheme that doesn't need a start-code scanner: Scanless-GOP. In GOP-level parallelism each group of pictures is handled in a processor as shown in figure 2. We employ closed GOPs in our evaluation for the sake of simplicity, but methods for open GOP structures are also introduced in previous works like [1] and [2]. In a closed GOP structure there aren't any references between two GOPs. So each GOP can be decoded in a separate core without any dependency.

However there is another issue limiting the scalability of the system: Before starting to decode a GOP, a process must know the start point of it. In previous approaches a separate process scans the input stream for GOP start-codes, cuts the stream into segments and places these in a task queue.

Actually, we do not need to search for these start codes every time we decode the video. We know their positions in the stream during the encoding process. So these start points can be written in the header of the video stream or in a separate file. For our evaluation, we produced the start points file by means of the decoder, but the same job may easily be done by the encoder as well. Before starting to decode, each process can read these start points into an array, and decode its own portion without waiting for anything.

After this observation we may suggest the following parallel decoding algorithm:

1. Read the start points into an array.
2. i = process id number

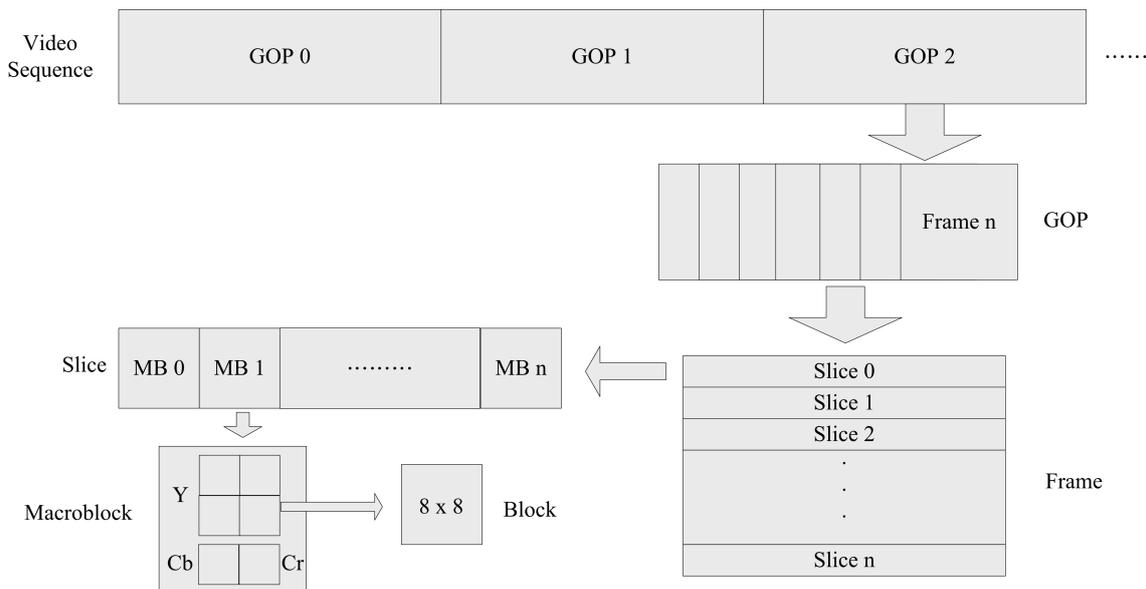


Figure 1: Data structure of H.264 video sequence and partitioning alternatives

3. While (not end of the stream)
4. Decode (i)th GOP
5. $i = i + (\text{total process number})$

The Decode (i)th GOP part of the algorithm includes writing the decoded frames into a disk or a display device or a collector process, but we will exclude this time in our performance calculations.

4. Experimental methodology

In our performance evaluations, we used the H.264/AVC reference software *JM* as the starting point. After parallelizing the program according to our method, we tested our parallel algorithm using a cluster of 5 IBM x3550 machines. Each machine has two 4-core Xeon 5500 series processors. We have tested 3 different situations using 3 different GOP sizes. GOP size affects the performance since it has a direct effect on the memory usage. In order to see the performance on a bounty of computing resources, first we ran just one process in each machine. Then we performed a test on one machine only, to see the performance on a shared-memory platform. And then, we checked the performance using all of the 40 cores. We used Intel's VTune to profile the processor events during the program execution. Profiling results has been very helpful for finding out the bottlenecks of the parallel program. Finally, we tested the run-time and speedup in a different machine having a different multicore cache architecture to see the effect of cache structure on scalability.

5. Results and discussion

5.1. Speedup improvement when we remove the start-code scanner

The speedup gain in coarse grain approaches may be limited in CMP architectures due to memory issues. If we also add some sequential scheduling part to the overall parallel program we might encounter with very disappointing results like in figure 3. Here we see that in a 2 x 4 core machine, the run-time of the algorithm using a start-code scanner hardly decreases from 13.286 sec to 11.057 sec in the minimum case when using 6 cores and than starts to increase again.

In figure 3 we see that when we remove the scanner from the algorithm by using ready start codes, we can get the run time decrease from 13.233 sec to 6.105 sec in the minimum case when using 7 cores.

Removing the GOP start-code scanner makes our algorithm purely parallel. Now each process is decoding its own part of data without sharing data with other processes. Besides there is not any scheduler or manager processes assigning tasks to others. So ideally, this algorithm should let each process work without any synchronizations.

We can see in figure 4 that maximum speedup jumps from 1.2 to 2.16 when we remove the scanner.

5.2. Performance in multi machines

We have seen that the speedup in 1 machine is bounded at 2.16. That is not compatible with the ex-

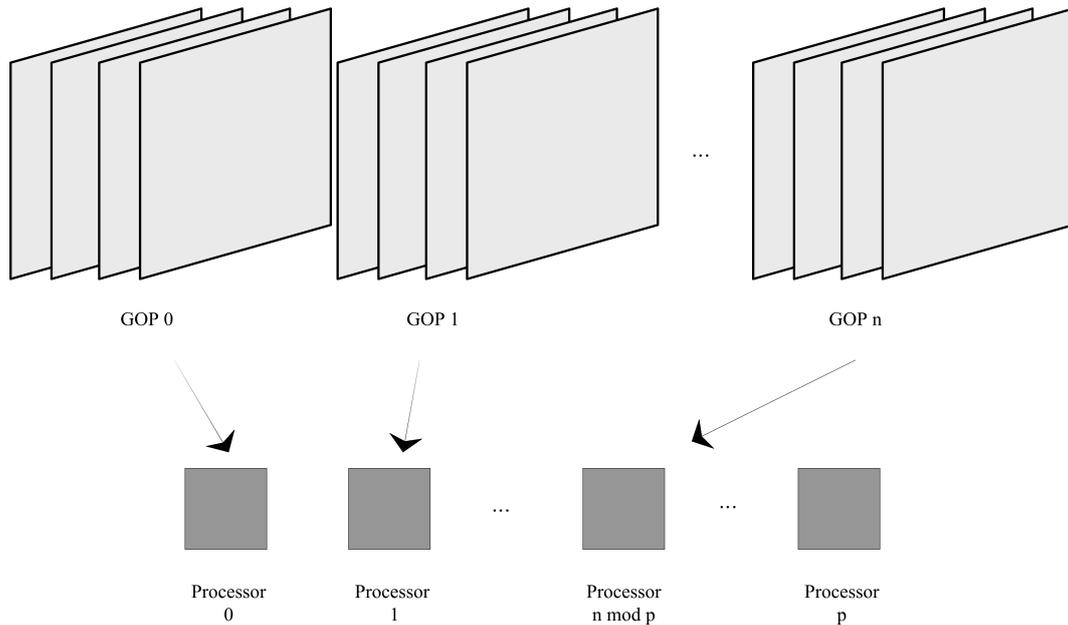


Figure 2: GOP-level parallelization

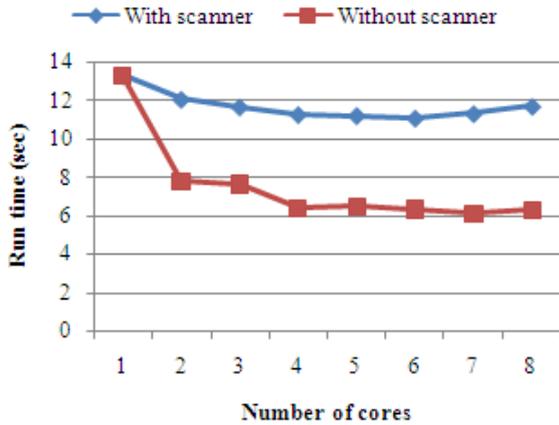


Figure 3: Run-time comparison of the algorithm with start-code scanner vs. the algorithm without a start-code scanner

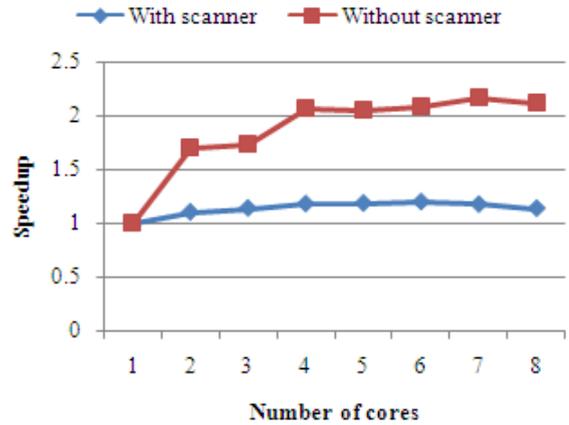


Figure 4: Speedup comparison of the algorithm with start-code scanner vs. the algorithm without a start-code scanner

pected ideal performance of this algorithm. In order to be sure we can check the results in figure 5. Here we see the speedups for different GOP sizes when we run only 1 process in each machine. This figure is very close to what we expect: we see a very beautiful, linear, one-to-one speedup!

To be sure, we may also check what would happen if we used all the available 40 cores in the 5 machines. Figure 6 shows that we can get a maximum of 11 times speedup when we keep the GOP size very small. GOP

size affects the speedup, because it affects the memory usage. A speedup of 11 is reasonable. It approximately equals the speedup in 1 machine, 2.16, times the number of machines, 5.

But, what is the bottleneck in shared-memory platforms?

5.3. The bottleneck in shared-memory platforms

We have seen that our algorithm does not have a very good scalability in single machines when using a shared

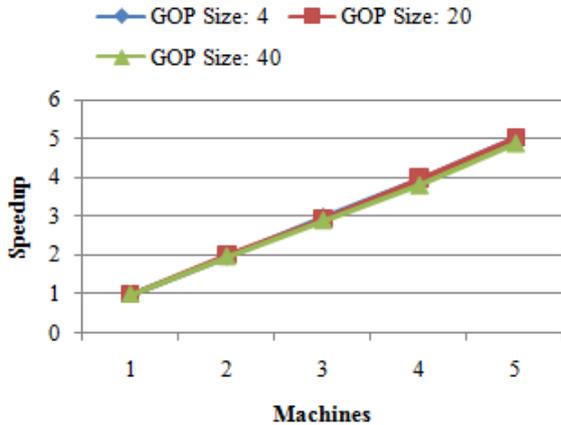


Figure 5: Speedup when only one process is run in each machine

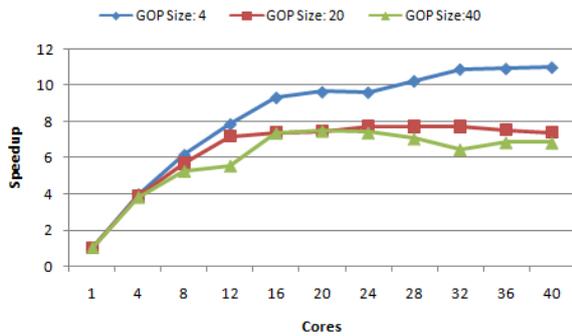


Figure 6: Speedup when we use all the 40 cores in 5 machines

memory, even though it produces a very good plot in multiple machines. As we mentioned in our conference paper [8] cache pollution is thought to be the main factor affecting the performance in chip multiprocessors. In order to be sure let's check out the profiling results obtained by Intel's vTune.

We see the impact of parallelization on branch mispredictions in figure 7. There is a slight increase in mispredicted branches. This must be due to the increase in instruction count. The maximum number of branch mispredictions is just about 420, which is a low number.

A reason of performance degradation might be register allocation table stalls. We can see how RAT stall number changes as we increase the number of parallel working processors. There is a general slight increase as we see in figure 8. The increase in RAT stall number may be counted as normal, since we get all the idle cores in the machine work.

We see a surprising result in figure 9 and figure 10.

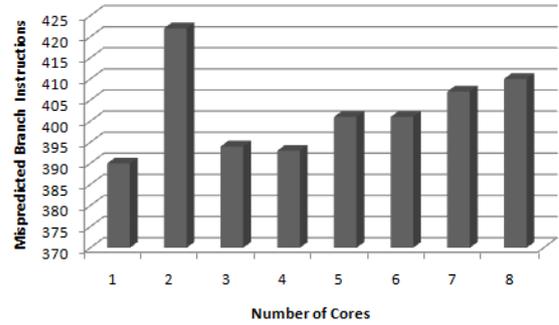


Figure 7: Parallelization impact on branch mispredictions

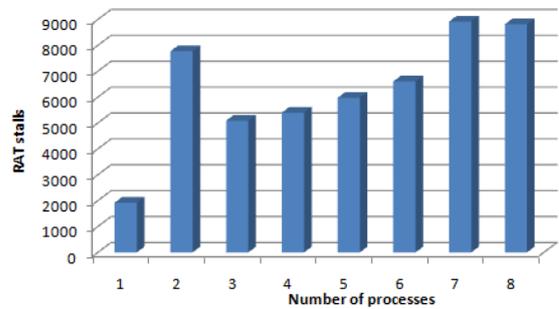


Figure 8: Register allocation table stalls

Major result of cache pollution is that it increases the cache misses. The data written into the cache newly, causes replacement of the data already in the cache. So when the program needs the replaced data it has to reload it from the memory system. The amount of data processed by a video decoder is very large. It is expected that when each core processes 1 GOP its cache be filled quickly. So we expect a big number of cache misses. However, we see a pretty good L1 cache performance and an almost ideal L2 cache performance: 17 misses at most.

So, what causes the speedup degradation? Figure 11 takes us close to the answer. 189072 resource stalls occur when 1 core is work, this number increases as we increase the parallel working core number and reaches 736559 when we employ 8 cores. This may explain the matter. Now, we know that there is a resource shortage, but which resource is not sufficient?

Figure 12 gives the answer. We see that almost all of the resource stalls are because of load-store operations. We cannot even see the plot for other stalls since their number is too small when compared with the resource stalls due to the load-store operations. When 1 core is

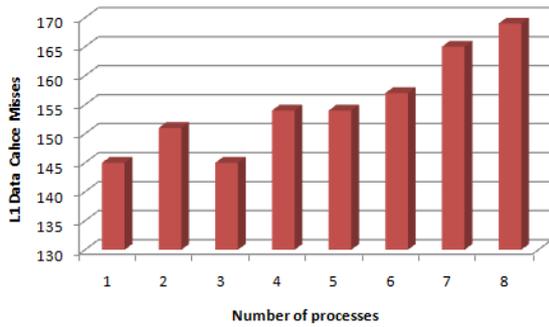


Figure 9: Level 1 data cahce misses

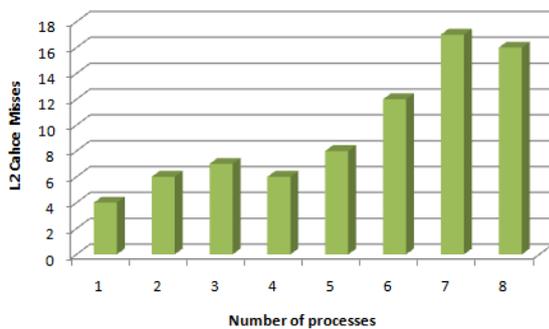


Figure 10: L2 cahce misses

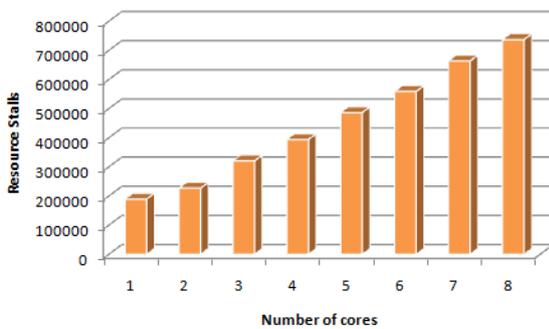


Figure 11: Effect of parallel processing on resource stalls

working there are 177,117 stalls due to load-store operations. The stall number increases to 732,792 when all the 8 cores are working in parallel. We have seen that cache performance is quite good. So we should suspect the performance of store operations.

5.4. Most active functions

It is the time now to check the most active functions in the system. Figure 13 is very interesting. Function

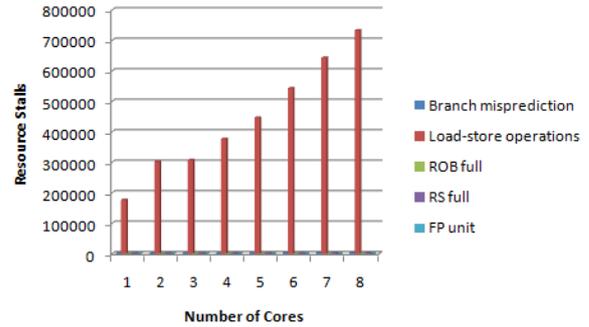


Figure 12: Effect of parallel processing on different kinds of resource stalls

memset takes the 70.47% of the whole system for serial execution. This means that most of the time is spent for setting memory locations to a specific value. Basically, we expect store operations to be faster than the load operations. Because during a store the processor can continue to execute other instructions after sending it to the memory system, it does not need to wait for the store operation to finish. But store operations are also done via cache levels. They affect the cache bandwidth and take longer than calculation instructions. When too many stores are executed one after another, the store buffer of the processor fills up, and this causes the load-store resource stalls.

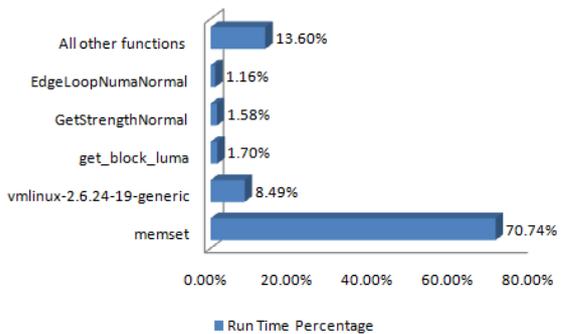


Figure 13: Most active functions in the system during serial run

5.5. Impact of memory initializations

We noticed that most of the calls to *memset* function is done by the *calloc* function which allocates a location in the memory and initializes these locations to zero. But are all these initializations needed? The answer is system dependent. For portability reasons, in the reference software, all the memory allocations are done with

an initialization. Not affecting the correct program execution for linux, we have changed most of the *calloc* functions to *malloc*, which allocates memory, but does not initialize the allocated addresses. We can see the great difference in run time and speedup in figures 14 and 15.

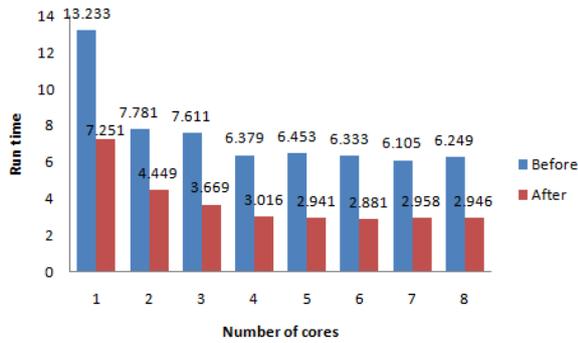


Figure 14: Run time improvement after removing unnecessary memsets

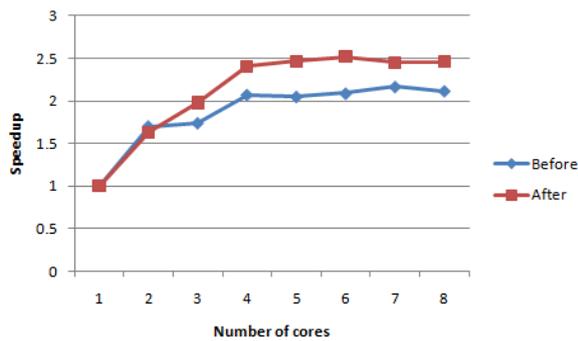


Figure 15: Speedup improvement after removing unnecessary memsets

We see that the runtime changed from 6.105 to 2.881, that is we got a speedup of 2.119 just by changing the function *calloc* to *malloc* in suitable places. This shows the impact of memory operations on the performance. The scalability of our algorithm also improved. The speedup due to parallel processing was 2.167 at most before, now we can see a speedup of 2.516.

Figure 16 shows the most active functions in the system after the modification. We see that the %age of the *memset* function decreased to 60.03%.

We have seen that resource stalls due to the store operations is the main bottleneck of this application. Most of the execution time is used by *memset* function which sets a memory area to a specific value. This is done when resetting a data-structure before starting to use it

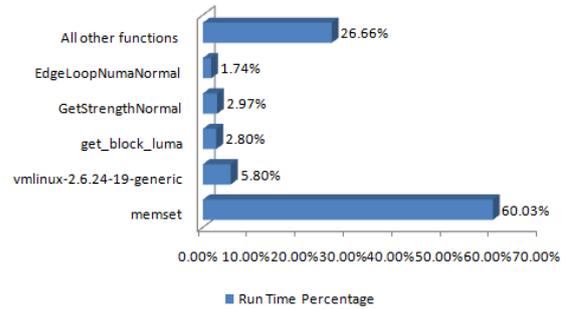


Figure 16: Most active functions in the system during serial run

for some different calculation. But this resetting gets very expensive when it is done again and again. It consumes 60% of the runtime even after we removed the unnecessary ones for linux operating system.

The processors, repeat the store operation of the same value for the number of items to be initialized. When we have some big amount of data to be initialized this results in a long latency. A possible way to make these memory initializations faster might be using the non-temporal stores which is offered by SSE instruction set. These instructions stores the data directly to memory without cache allocation. So in the first glance, we may think that they might have a better scalability, because the cores will not wait for each other for writing to the same cache. However, they will bypass the cache system, only if the related address is not cached. For the case of memory initializations, almost all the data we deal with is already in the cache, because we are resetting some data that we have already used. So we should not be so hopeful about non-temporal store instructions for our problem. Indeed, when we changed the *memset* function with a function using non-temporal stores, we did not observe any speedup.

5.6. Impact of cache architecture

Finally, we should check the effect of cache architecture on scalability. The processor we used till now is Intel Xeon 5500 Series 4-core architecture. It has a *shared* L2 cache of 8MB. We may suspect that a shared cache may degrade the performance of parallel processing. As we increase the parallel working core number the cache size available for each core decreases!

The results we obtained up to this point reveal that shared cache architectures are not suitable for highly scalable data-intensive parallel applications. Even though we haven't encounter a misprediction problem due to enough cache size, the store operations caused the cores wait for each other to write to the same cache.

The performance of shared cache structures might be enhanced by providing a separate read-write port for each core in the environment. In our case, we saw that even though the cores are working on totally different data, they had to wait for using the shared resources. We measured 177,117 stalls due to load-store operations when working with 1 core. The stall number increased to 732,792 when all the 8 cores joined the calculation.

Therefore, we tested the performance of our algorithm in a different platform. We used AMD Turion 64 X2 dual-core mobile technology processor for comparison. The architecture of this processor is given in figure 17. We see that in this processor there are two separate level 2 caches for each core. Each cache is 512KB. We should not be surprised that the machine with AMD dual core processor is slower because it is a notebook having 1.8GHz core speed and a total of 1 MB L2 cache, while the machine with the Intel processor is a server having 2.5GHz core speed and a total of 8 MB L2 cache.

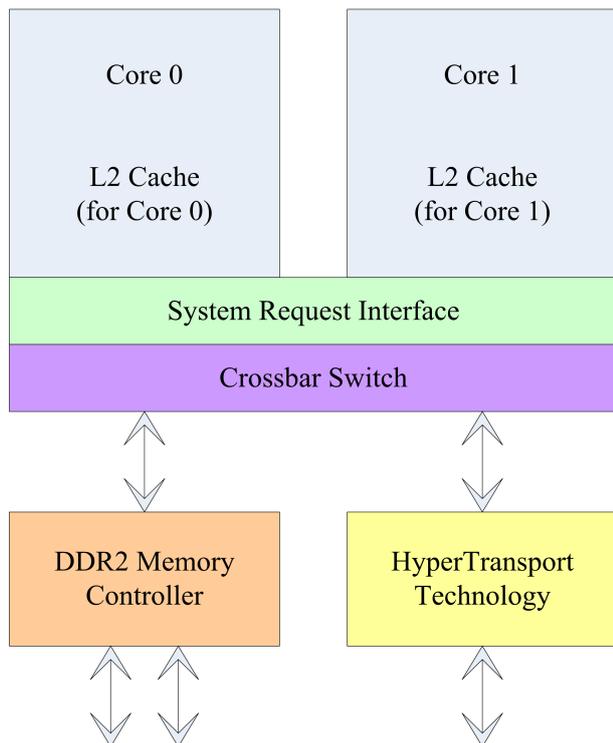


Figure 17: The architecture of AMD Turion 64 X2 mobile technology processor

Table 1 shows the impact of cache architecture on the scalability. We see that we can get a speedup of 1.962 with 2 cores in a shared memory platform if we use a separate cache for each core. This is almost an ideal

result! This result proves that separated caches are more suitable for multicore machines in order to get a better scalability. The result also proves the efficiency of GOP-level partitioning without a start-code scanner.

	Shared Cache		Separated Cache	
	1 core	2 cores	1 core	2 cores
Run Time (s)	13.233	7.781	21.438	10.925
Speedup	1.701		1.962	

Table 1: Effect of cache architecture on speedup

6. Related work

First GOP-level parallel video decoder was introduced a long time ago in [2]. It is a real-time MPEG-1 decoder consisting of parallel processing 16 nodes having distributor and collector components. The task of the distributor is to cut the video sequence into segments. An other GOP level parallelization of MPEG-1 encoder for MIMD multiprocessors was presented by Shen [9].

We see a real-time parallel MPEG-2 decoder in [3]. Both GOP-level and slice-level approaches are evaluated. This system also has scan and display processes. Scan process is responsible for reading the encoded video from the disk and placing encoded GOPs into a task queue. For GOP approach they observed almost linear speedup in all cases. The bottleneck of the design is that, the memory requirement increases with the size of the GOP, size of the picture and number of parallel processors used.

Various slice level parallelism methods have been suggested. Lee introduced a slice level parallel MPEG-2 decoder for HDTV [10]. A manager-worker style parallel H.263 decoder was implemented by Lehtoraanta [11] using 1 manger and 3 worker DSP cores.

A hierarchical parallelization approach for H.264 encoder is introduced in [4]. In this paper authors suggest that a GOP-level scheme and a slice-level scheme might be used together for overcoming the latency problem. Another hierarchical method was proposed by Chen [14]. In this study frame level and slice level parallelism employed together. When the frame level saturates, slice level parallelization is used for further partitioning. They declare a 4.5X speedup in a machine having 8 cores.

Roitzsch [15] suggests modifying the encoder for obtaining slices that are equivalent in decoding time, not in macroblock number. This approach needs encoder

modification, and increases slice number and bit rate to increase scalability.

A task level decomposition method has been introduced by Gulati [12]. His system both encodes and decodes H.264 video sequences in real time by means of a control processor and 3 DSPs. Schoffmann [13] also suggests a pipeline model, but at macroblock level.

Among fine grain methods, 2D-wave approach [5], and 3D-wave technique [6] declares pretty high scalabilities. However, fine-grain approaches depend on a well designed CABAC accelerator, since entropy decoding of a single slice or frame is mostly sequential.

7. Conclusions

We have introduced a GOP level parallelization method for the H.264 video decoder. Our method revokes the need for a start-code scanner, thus lets all the processors in the environment contribute to the decoding task. This technique also lets the processors work without waiting for a new task assignment. So in the ideal computation environment it has perfect scalability.

We have observed a one-to-one linear speedup in parallel working machines. This is because the memory resources does not change when we increase the parallel working processor number. So we can observe a speedup close to ideal.

The speed of memory store operation degrades the speedup in shared memory platforms. As the parallel working processors increase, the number of simultaneous store operations also increase. This causes a lot of resource stalls due to fullness of the store buffer. We saw a maximum speedup of 2.516 when working with 6 processors, and the speedup get close to saturation after 4 processors.

We saw that memory load-store operations are very expensive, and they should be carefully utilized. When we replaced the *calloc* function with *textitmalloc*, both of which are doing memory space allocation but *calloc* is also doing initialization of allocated memory, we get a run-time decrease from 6.105 sec to 2.881 sec, and maximum parallel processing speedup increased from 2.167 to 2.516.

The data-structure initializations gets very expansive when we deal with big amount of data like in video encoding and decoding. So a special hardware and additional instruction may be designed to make these initializations less time-consuming.

Finally, we saw the effect of multicore cache architecture on scalability. The speedup changes from 1.701

to 1.962 for 2 cores when we use separated L2 caches in state of shared ones. An almost ideal speedup of 1.962 with 2 cores shows the efficiency of GOP level parallelization without a start-code scanner.

References

- [1] P. Stenstrom, Chip-multiprocessing and Beyond, Proc. Twelfth Int.Symp. on High-Performance Computer Architecture, 2006, pp. 109 - 109.
- [2] M. K. Kwong, P. T. P. Tang, and B. Lin. A Real-Time MPEG Software Decoder Using a Portable Message-Passing Library. Mathematics and Computer Science Division, ANL, Argonne, IL 60439-4844, 1995.
- [3] A. Bilas, J. Fritts, and J. Singh, Real-time parallel mpeg-2 decoding in software, Parallel Processing Symposium, 1997. Proceedings., 11th International, pp. 197203, 1-5 Apr 1997.
- [4] A. Rodriguez, A. Gonzalez, and M. P. Malumbres, Hierarchical parallelization of an h.264/avc video encoder, Proc. Int. Symp. on Parallel Computing in Electrical Engineering, 2006, pp. 363368.
- [5] Mauricio Alvarez Mesa, Alex Ramirez, Arnaldo Azevedo, Cor Meenderinck, Ben Juurlink, Mateo Valero, Scalability of Macroblock-level Parallelism for H.264 Decoding, icpads, pp.236-243, 2009 15th International Conference on Parallel and Distributed Systems, 2009
- [6] A. Azevedo, B.H.H. Juurlink, C.H. Meenderinck, A. Terechko, J. Hoogerbrugge, M. Alvarez, A. Ramirez, M. Valero, A Highly Scalable Parallel Implementation of H.264, Transactions on High-Performance Embedded Architectures and Compilers (HiPEAC), September 2009.
- [7] Cor Meenderinck, Arnaldo Azevedo, Ben Juurlink, Mauricio Alvarez Mesa, Alex Ramirez, Parallel Scalability of Video Decoders, Journal of Signal Processing Systems, v.57 n.2, p.173-194, November 2009.
- [8] Ahmet Gurhanli, Charlie Chung-Ping Chen, Shih-Hao Hung, GOP-level Parallelization of the H.264 Video Decoder without a Start-code Scanner, International Conference on Signal Processing Systems, July 2010, in press
- [9] Shen, K., Rowe, L. A., Delp, E. J., Parallel implementation of an MPEG-1 encoder: Faster than real time. In Proc. SPIE, digital video compression: Algorithms and technologies 1995 (Vol. 2419, pp. 407418).
- [10] Lee, C., Ho, C. S., Tsai, S.-F., Wu, C.-F., Cheng, J.-Y., Wang, L.-W., Implementation of digital hdtv video decoder by multiple multimedia video processors., In International conference on consumer electronics, 1996 (pp. 98), 57 June.
- [11] Lehtoranta, O., Hamalainen, T., Saarinen, J., Parallel implementation of h.263 encoder for cif-sized images on quad dsp system. In The 2001 IEEE international symposium on circuits and systems, ISCAS 2001 (Vol. 2, pp. 209 212), 69 May.
- [12] Gulati, A., Campbell, G. Efficient mapping of the H.264 encoding algorithm onto multiprocessor DSPs. In Proc. embedded processors for multimedia and communications II, 5683(1), 94103, March 2005.
- [13] Klaus Schffmann, O. L., Fauster, M., Bszrmenyi, L., An evaluation of parallelization concepts for baseline-profile compliant H.264/AVC decoders. In Lecture notes in computer science. Euro-Par 2007 parallel processing, August.
- [14] Chen, Y., Tian, X., Ge, S., Girkar, M., Towards efficient multi-level threading of h.264 encoder on intel hyper-threading architectures. In Proc. 18th int. parallel and distributed processing symposium, 2004.
- [15] Roitzsch, M., Slice-balancing H.264 video encoding for im-

proved scalability of multicore decoding. In Work-in-progress proc. 27th IEEE real-time systems symposium, 2006.



Ahmet Gurhanli is a Ph.D. candidate in National Taiwan University, Graduate Institute of Electronics. He graduated from Hacettepe University, Electronics Engineering Department in 2003. He received his MS degree from National Taiwan University in 2006. His research interests are Computer Architecture, Parallel Computing, Electronic Design Automation and Video Processing



Charlie Chung-Ping Chen is a professor at National Taiwan University, Graduate Institute of Electronics Engineering. He received his Ph.D. from the University of Texas at Austin, USA. His research interests are VLSI CAD, Microprocessor Design, and RF Mix/Signal Circuit Design.



Shih-Hao Hung is an assistant professor at National Taiwan University, Graduate Institute of Networking and Multimedia. His research interests are Co-Optimization of Applications and Architectures, Computer Architecture and Parallel Computing, Computer Performance Characterization and Optimization, Commercial Servers and Applications, and Design of Embedded Systems.