The KIT COSMOS Processor: Some Ideas on Realizing Complexity-Effective Superscalar Processors

Toshinori Sato[1,2] Toshiyuki Yamamoto[3] Itsujiro Arita[1]
[1] Department of Artificial Intelligence, Kyushu Institute of Technology
[2] Center for Microelectronic Systems, Kyushu Institute of Technology
[3] Alpha Systems Inc.

Abstract

This paper introduces a new microarchitecture, which we call COSMOS, for realizing large scale superscalar processors with high clock frequencies. In order to achieve the goal, several design techniques on instruction supply mechanism, instruction window, register files, and operand bypass logic are proposed. Based on simulation results of an 8-way dynamically scheduled superscalar processor, we estimate that the contribution of COSMOS microarchitecture on instruction level parallelism is only 1.3% degradation from the conventional superscalar processor with an equivalent scale. Due to its low complexity, a COSMOS superscalar processor works at higher clock frequencies than the conventional one. Therefore, COSMOS microarchitecture is one of the promising candidates for future superscalar processors.

Keywords: superscalar processors, instruction supply mechanism, instruction window design, variable latency pipeline, clustered architecture

1 Introduction

Future microprocessors will rely on higher clock frequencies, wider instruction issue width, deeper pipeline stages, and larger instruction windows in order to improve performance. As the width and size increase, the scaling of the clock frequencies becomes difficult to attain. Obstacles for the scaling are register renaming logic, large register files, and operand bypass logic as well as instruction window wakeup and select logic [11,14]. There are many proposals to alleviate these constraints on the renaming logic [1], the registers [7, 24], and the instruction window [4, 14]. This paper proposes a new microarchitecture, which realizes large scale superscalar processors with high frequencies emerging in the near future. We introduce an implementation example, which we call KIT COSMOS processor [20], and describe its unique features. This paper focuses on its characteristics based on the superscalar paradigm, while it is a simultaneous multithreading processor [20,25].

The features of the KIT COSMOS processor, which improve clock frequencies are summarized as follows.

- Clustered architecture [20]
- High bandwidth instruction cache [21]
- Decoupled instruction window [19]
- Simplified instruction issue logic [22]
- Variable latency pipeline [23]

It is important to consider impact of hardware complexity on processor cycle time for large scale superscalar processors. In order to solve the problem, it is widely investigated to split the large processor into a number of small processing elements as the first feature. As instruction issue width is increase, requirement of instruction fetch bandwidth also increases. The conventional multiple-port cache combined with the instruction alignment circuit becomes the bottleneck deciding processor cycle time. The second feature of COSMOS attacks this problem. The next two deal with large scope instruction scheduling. The large instruction window is required for efficient instruction scheduling, while it is the most critical in the future [14]. The hierarchical and simplified design helps minimize performance degradation without increasing cycle time. The result bypass logic is an another critical problem, which is considered by the last one.

The organization of the rest of this paper is as follows. Next section introduces the KIT COSMOS processor. Section 3 explains the high bandwidth instruction supply mechanism. The large instruction window and the operand bypass logic are dealt with in Sections 4 and 5 respectively. We present some evaluation results for these techniques based on timing simulations using SimpleScalar tool set [2]. The current processor model evaluated is an 8-way dynamically scheduled superscalar processor. For measuring performance, we use in this paper the committed instructions per cycle

(IPC) as a metric. Only useful instructions are considered for counting the IPC. We do not count nop instructions. Section 6 surveys related works. Finally, our conclusions are presented in Section 7.

2 KIT COSMOS Processor

Figure 1 shows the block diagram of a COSMOS processor, which consists of one instruction supply mechanism and two instruction execution mechanisms (clusters). The number of the clusters is not necessary two. The instruction supply mechanism will be explained in Section 3. The instruction buffer and the scheduling window will be described in Section 4. The detail of a functional unit will be found in Section 5. These components contribute to high clock frequencies. On the other hand, this section focuses on the clustered architecture, which is combined with the value predictor.

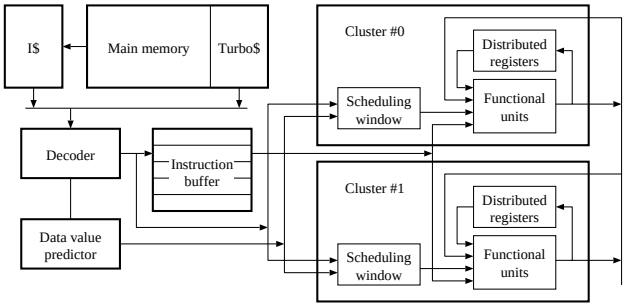

Figure 1: A dual-clustered COSMOS processor

The delay of both the large register file and the long bypass busses is severe and is degrades processor clock frequencies [14]. The clustered architecture attacks this problem by distributing the register files and the functional units across multiple clusters. Each register file becomes small and the bypass busses can be short. Therefore, the clustered processor can operate at higher frequencies than the superscalar processor with an equivalent scale.

In COSMOS processor microarchitecture, a register file is divided into multiple local register files and distributed across the clusters. There are not any copies of each register file nor any global register files. When an instruction is decoded, the steering logic decides in which cluster it is to be issued. If any source operand of the instruction is not held in the local register file, it is required to move the operand from the remote cluster. This incurs penalties at least one cycle. The proposed microarchitecture solves this problem by exploiting value prediction [9,13]. The instruction which requires any source operand located in the remote cluster uses its predicted value and executes speculatively. If the prediction is correct, the processor does not suffer the cycle penalty from moving the operand between clusters. Only when the prediction is incorrect, it suffers the penalty. One of the advantage of this value prediction strategy is that instructions held in the value prediction table are limited. Thus, the prediction table is efficiently utilized and its capacity can be reduced without performance loss.

In the remaining of this paper, we assume a single-clustered COSMOS processor as a representative for large scale superscalar processors.

3 Instruction Supply Mechanism

We propose a simpler trace cache [15, 16], named non-consecutive basic block buffer (NCB) [21]. The followings describe the NCB fetch mechanism and explains the difference between two mechanisms.

The NCB is an extension of branch target buffer (BTB), and caches source basic blocks as well as target basic blocks, i.e. the NCB stores multiple non-consecutive basic blocks. Figure 2 indicates the NCB fetch mechanism. The instruction cache and the NCB are indexed by the program counter. The branch predictor is used for selecting an instruction sequence from the ones supplied by the instruction cache and the NCB. If the branch outcome is predicted taken, the instruction sequence (a trace) from the NCB is selected. Otherwise, the one from the instruction cache is selected. In the case that the trace is not cached in the NCB, instructions are provided from the instruction cache.

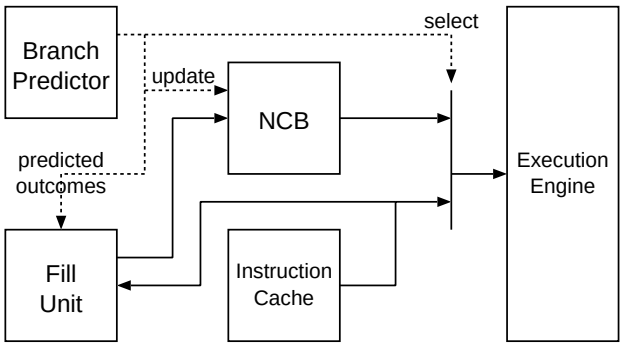

Figure 2: NCB fetch mechanism

Trace constructing process executed in the fill unit is as follows. `Step 1:` The fill unit temporary buffers instructions from the instruction cache in a line buffer. `Step 2:` If there is a branch instruction (both conditional and unconditional) in the instructions and the branch outcome is predicted taken, the instruction se-

quence reaching the branch instruction is stored in the line buffer. `Step 3:` After next instructions from the instruction cache is fetched, the former instructions stored in the line buffer and the latter instructions are combined into a trace. The trace constructing process is finished when

1. the line buffer becomes full, or
2. the trace contains two branches predicted taken (including indirect jump).

And then, the constructed trace is stored in the NCB.

This process resembles the one performed by the collapsing buffer [6]. In the case of the collapsing buffer, `Steps 1-3` have to be processed in one cycle. On the other hand, in the case of the proposed scheme, these steps can be pipelined because this process goes out from the instruction supplying process. Therefore, the complicated fill unit structure does not have any impact on the processor cycle time.

The differences between the trace cache and the NCB are summarized as follows.

1. The NCB stores only those traces containing taken branches, eliminating unnecessary replication of trace and basic block.
2. In the case of the NCB, a basic block can be split across two traces just like the trace packing [15]. And hence, starting addresses of traces do not always match with basic block boundary.
3. A single indirect jump, return or trap instruction does not terminate trace construction, reducing unused instruction slots.
4. The trace construction is based on predicted branch outcomes rather than actual ones. This decouples the NCB from the execution engine, reducing hardware complexity. In addition, this enables each trace to be constructed early.

According to these characteristics, the NCB fetch mechanism becomes simpler than the trace cache.

Figure 3 compares processor performance. We simulate three models: the baseline model (64K I-cache), the baseline model whose instruction cache is doubled in capacity (128K I-cache), and the evaluated model (64K I-cache + 32K NCB). Results of the last two models are normalized by the first one. It can be easily found that the large cache model (denoted as `128`) can not contribute to processor performance. On the other hand, it is observed that the IPC of the NCB model (denoted as `N`) is increased by 9.4% on average and a maximum of 13.7%. This confirms that the NCB can more efficiently supply instructions than the conventional instruction cache.

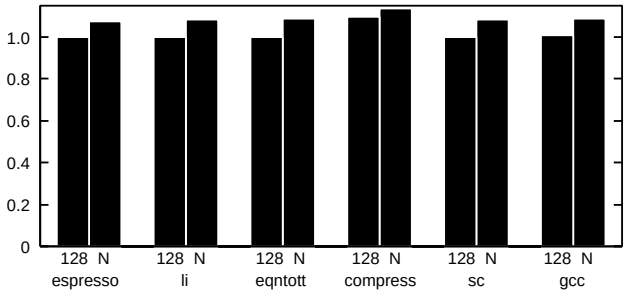

Figure 3: Performance contribution of NCB

4 Simple Instruction Window

This section considers how to realize large instruction window. First, we decouple the instruction window into relatively small instruction scheduling windows and a large instruction buffer which is a backup for the first one. Second, a simplified instruction wakeup logic utilized in the small scheduling window is proposed. This removes associative lookup and thus the scalability is improved.

4.1 Decoupled instruction window

COSMOS processor architecture relies on data speculation technique as described in Section 2. The data speculation is a new technique removing serialization on a program execution caused by data dependences based on data value prediction [9, 13]. When a predicted value is correct, it becomes possible to execute the predicted instruction and its dependent instructions simultaneously, thereby more instruction level parallelism (ILP) is extracted. Otherwise, it is necessary to revert processor state to a safe point where the speculation is initiated. Instruction reissue is such a technique recovering the processor state when a misspeculation occurs. It invalidates instructions dependent upon the misspeculated instruction selectively and then reissues them in the instruction window. It is possible to realize the instruction reissue with a simple hardware structure [18]. However, the instruction reissue technique has a problem. In order to realize the instruction reissue, each instruction remains in the instruction window until it is not speculative. In other words, it can not release its entry in the instruction window until it is committed. Thus, effective capacity of the instruction window is reduced, diminishing the freedom of the instruction scheduling [5]. In order to maintain the freedom, it is necessary to increase the instruction window size.

In order to realize a large instruction window, we propose to decouple the recovery mechanism for data speculation from dynamic instruction scheduling structure [19]. As shown in Figure 1, the decoupled

window consists of relatively small instruction windows for the scheduling and a large instruction buffer for the instruction reissue. The small scheduling windows are distributed across clusters. On the other hand, the large buffer is centralized. After an instruction is fetched and decoded, it enters both the scheduling window and the instruction buffer. When the instruction is dispatched to a functional unit, it leaves the scheduling window and release its entry but remains in the instruction buffer. When it is committed, it leaves the instruction buffer and releases its entry. In the case that either the window or the buffer is full, instruction issuing stalls.

The small instruction window works for dynamic instruction scheduling. Thus, most of the times each instruction is dispatched from a small window. Since instructions are aggressively deallocated from the scheduling window when they are dispatched, the problem reducing its effective capacity is solved. In addition, its relatively small size does not have serious impact on processor cycle time. However, it is impossible to reissue misspeculated instructions inside the scheduling window. The large instruction buffer works as the backup for the small window and performs the instruction reissue. Each instruction remains in the buffer until it is committed. When a misspeculation occurs, reissued instructions are obtained from the instruction buffer. In order to aggressively speculate instructions, the instruction buffer should be very large. Since it is difficult to access such a large buffer in one cycle, the wakeup and selection logic of the buffer is pipelined to maintain high clock frequencies. It is expected that the pipelining does not degrade processor performance, since the instruction buffer is active only when misspeculations are detected.

It is straightforwardly decided which structure dispatches an instruction to a functional unit. When an instruction is misspeculated, its dependent instructions which have already dispatched and thus which should be reissued are obtained from the instruction buffer only. They have already left the scheduling window. The dependent instructions which have not been dispatched remain in both the scheduling window and the buffer. However, the instructions are obtained only from the scheduling window, because they are not candidates for reissue.

Figure 4 presents the performance contribution of the decoupled instruction window. We evaluate three models. The first (denoted as `B`) is a conventional processor model which has a 512-entry instruction window. The second (denoted as `P`) is the one whose instruction window is pipelined by 2 cycles. And the last (denoted as `D`) is the decoupled instruction window model. The instruction buffer and the scheduling window has 512 and 256 entries respectively. The buffer is pipelined by 2 cycles but the scheduling window is not pipelined. The 256-entry instruction window is still too large to achieve high clock frequencies. This is resolved in the next subsection. Each model utilizes 4096-entry stride value predictor [9], and its performance is normalized by the baseline model without the value predictor. In this evaluation we use 128K conventional instruction cache. It is observed that the pipelining of the instruction window is seriously degrades processor performance. However, it can be found that the performance degradation due to the pipelining is compensated by its decoupling. For all cases, processor performance is improved over the baseline model, 3.7% improvement on average. Compared with the conventional model with the value predictor, performance of the decoupled model is lower. However, it is expected that the clock frequencies of the decoupling model is faster than that of the conventional one. Therefore, the processor performance of the decoupling model will be comparable to that of the conventional one.

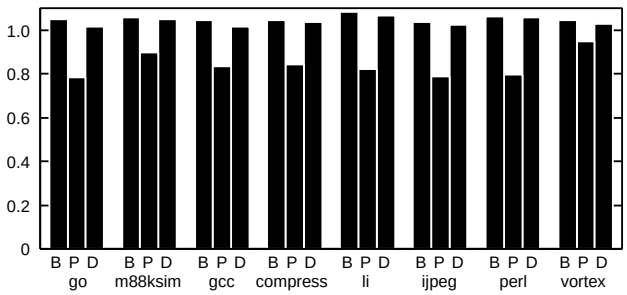

Figure 4: Performance contribution of decoupled instruction window

4.2 EDF Instruction Window

In order to eliminate anti- and output-dependences, modern dynamically scheduled processors perform register renaming. There are two common ways to implement the register renaming. One is using a separated renaming registers which are usually constructed by reorder buffer. The other combines the renaming registers with architected registers in a single register file. We focus on the latter case, especially based on MIPS R10000 [26].

In order to improve the scalability of the instruction window by reducing the delay of the instruction wakeup logic, we propose the explicit data forwarding (EDF) instruction window [22]. The main purpose of the EDF instruction window is using RAM which

has more scalability than CAM. The EDF instruction window consists of the RAM instruction window and a table named dataflow management table (DMT). In order to replacing CAM by RAM, the dependences between instructions are definitely explained by any means. The DMT keeps the dependences. Figure 5 depicts the DMT with attached to the register mapping hardware. It is indexed by physical register number and each entry holds IDs indicating specified instruction window slots. In Figure 5, the number of IDs which are registered in each entry is one. However, it can be increased by analyzing the tradeoff between performance and its hardware cost. The dependences between instructions are registered when every instruction is issued, and the DMT is referred when instructions complete. The registration process is as follows. As shown in Figure 5, the DMT is indexed by the physical operand register numbers and the ID associated with the instruction which requires the operands are registered. Since the instruction in the figure has two operands, the ID corresponding to the instruction is registered in 2 entries. The identifiers (denoted as `A` and `B`), that tell which operand for the instruction it is, are also held in the DMT.

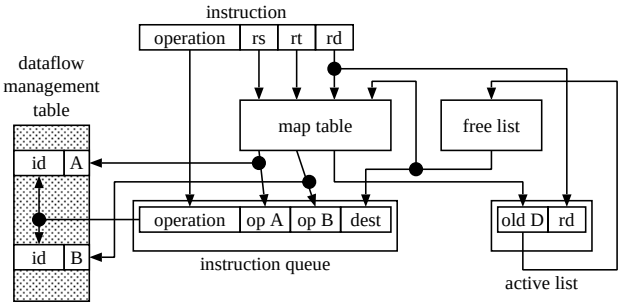

Figure 5: Dataflow management table

The reference process and its following instruction wakeup process are explained in Figure 6. When an instruction completes, the DMT is indexed by the result tag of the instruction, that is physical destination register number. From the table, instruction ID which requires the execution result is obtained. Using the ID, the ready bit of an entry associated with the instruction is set. If all ready bits in the entry are set, the instruction is ready for execution (wakeup). As can be seen, there are no associative lookup in the instruction wakeup logic. And thus, the instruction window is implemented using RAM.

It is important to mention how to handle branch mispredictions. When a branch is mispredicted, dependences held in the DMT are incorrect. Therefore,

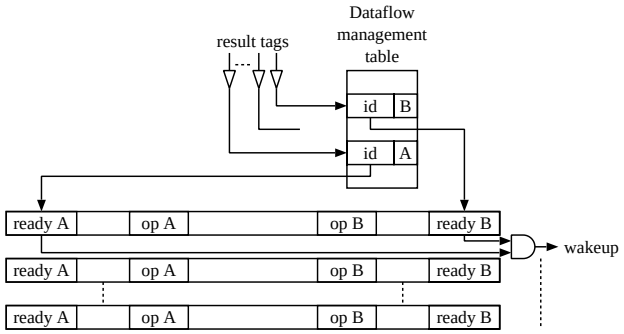

Figure 6: EDF instruction window

it is necessary to revert the table to a safe point where the speculation is initiated. This is easily handled. Every time a branch is predicted, a checkpoint of the DMT is made, just like the map table [26].

Next, we consider the delay of the instruction windows. In the case of the conventional instruction window, it is realized using CAM in order to the process of the wakeup logic. Since each entry of the window requires $2*IW$ comparators, where IW is the instruction issue width, the CAM instruction window size is in the order of $O(IW * WS)$, where WS is the instruction window size. From the detailed analysis by Palacharla et al. [14], the delay of the CAM window is in the order of $O(IW^2 * WS^2)$. On the other hand, the delay of the proposed instruction window is estimated as follows. The sizes of the DMT implemented using RAM and the RAM instruction window are both in the order of $O(WS)$. Based on the analysis on the register renaming logic by Palacharla et al. [14], the delay of them are both in the order of $O(IW^2)$. Therefore, the EDF instruction window is more scalable than the conventional instruction window.

Figure 7 compares processor performance. Performance for the proposed RAM instruction window is normalized by that of the conventional CAM instruction window. Every instruction window has 256 entries. For each group of three bars in the figure, the left, middle, and right bars indicate the results in the cases where the number of the ID field in the DMT is 1, 2, and 3, respectively.

It is found that the DMT with three ID slots is required for the EDF instruction window to achieve processor performance comparable to that of the conventional CAM instruction window, except for the case of `124.m88ksim`. On average, it attains 95.4% of that of the conventional case. It is interesting that in the case of `130.li`, the processor performance of the RAM instruction window exceeds that of the CAM window. This is due to the indeterminate characteristics of out-

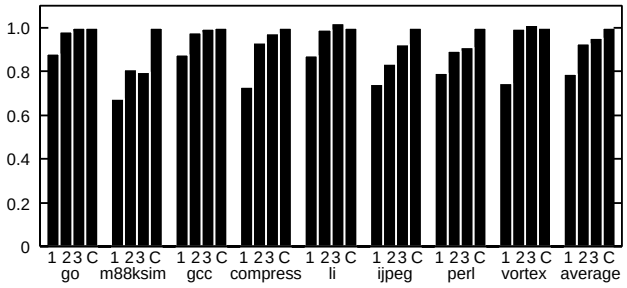

Figure 7: Performance contribution of EDF instruction window

of-order execution. For example, if the instruction issue stall reduces useless instructions which are executed on branch misprediction path, the performance is improved. When ID decreases to two, 92.7% of the conventional performance is achieved. If it still decreases to one, only 78.9% of the performance is achieved. Therefore, it is a good tradeoff point that the DMT has two ID slots.

5 Variable Latency Pipeline

The execution stage consists of execution latency and result drive time for the operand bypassing [11]. It is almost impossible to move the bypass logic from the execution stage to the issue stage, since the issue stage is already critical [26]. The delay of the bypass logic imposes the execution stage to be divided into several stages, increasing the execution latencies. Pipelining is one of the techniques realizing the high speed circuits and can improve the throughput of a function. However, it also increases the latency of the function and thus processor performance sometimes can not be improved by the technique. Therefore, careful considerations are required for applying the pipelining to the bypass logic. On the other hand, asynchronous and pseudo-asynchronous circuits are the other techniques realizing high speed circuits. However, they have the following problems. The asynchronous and pseudo-asynchronous circuits need a completion detector of each operation. The detector becomes the critical path of the circuits and increases processor cycle time. Furthermore, the throughput can be diminished for specific operands. From these considerations, techniques to reduce execution latencies including bypassing are required.

In order to mitigate the constraints on the bypass logic, we exploits variable latency pipeline (VLP) structure [23]. Figure 8 shows the concept of the VLP [12]. A function can be implemented by several kinds of circuits whose design policy are different with each other. In Figure 8, two circuits are used for implementing the function. Circuit A is designed so that most of the longest path of each operation is shorter than a processor cycle time. Circuit B is designed so that the critical path of the circuit is shorter than the cycle time, and is pipelined. Combining these two kinds of circuits reduces the effective latency to execute the function and also maintains the throughput of the function even for the operations which are executed in two cycles.

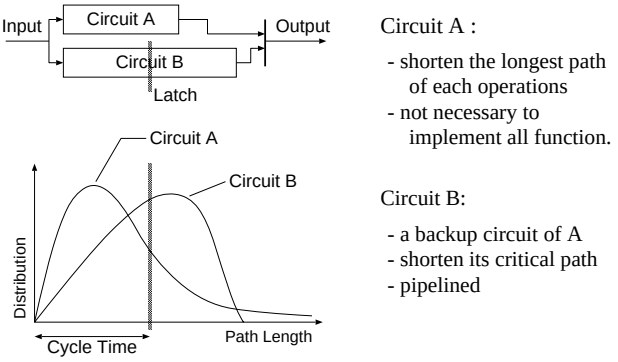

Figure 8: Variable latency pipeline structure

In order to select one from two results, a completion detector for circuit A is required. The detector can be pipelined since the pipelined circuit B works as the backup of the circuit A. Thus, the detector does not increase the critical path. From the above considerations, it can be seen that high frequencies and short latency function is implemented. It is true that the total transistor count increases in order to implement two circuits. However, it is possible to reduce the count if the circuits A and B share circuitry.

Figure 9 depicts the impact of the VLP on processor performance. We evaluate integer ALUs which utilize VLP. The latency of the ALU utilizing the VLP is 2 cycle when a carry propagation over 16 bits occurs. Otherwise it is 1 cycle. Candidate functions for the VLP include not only addition and subtract instructions but also load, store, and branch instructions. Multipliers and dividers do not utilize the VLP technique. For each group of four bars, the left bar is for the processor consisting of two latency pipelined ALUs and the right bar is for that consisting of the variable latency pipelined ALUs. Every simulation result is normalized by the processor model whose ALUs have 1 cycle latency. We can find that, utilizing the VLP, processor performance for most of the programs is comparable to performance of the baseline processor. It achieves 96.4% of the baseline.

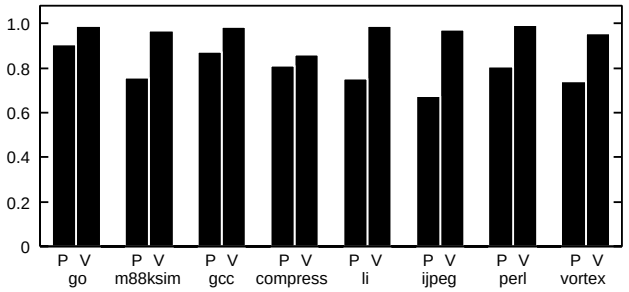

Figure 9: Performance contribution of VLP

6 Related Work

There are several studies on the clustered architecture [3, 8, 14], in which the register files, the instruction queue, and functional units are distributed across multiple clusters. Its motivation is to increase the clock frequencies by reducing the size and complexity of each cluster. The motivation of COSMOS clustered microarchitecture is the same, while there are some differences from the previous proposals. First, there are not any global register file nor copies of local register files in COSMOS microarchitecture. Second, COSMOS exploits the value prediction technique for the instruction assignment mechanism, which decides in which cluster an instruction is issued. Third, due to the introduction of the value prediction, the instruction window in a COSMOS processor has the hierarchical characteristics. And last, each cluster in COSMOS is a relatively large scale superscalar processor. In this paper, we have evaluated 256-entry local instruction queue.

Trace processor architecture [17] is an another related study. The NCB resembles the trace cache [15, 16], while some differences have been explained in Section 3. The main difference from Trace processor is that Trace processor exploits both instruction and trace (thread) level parallelism, while a COSMOS processor explained in this paper only relies on ILP. And thus, the purpose of the value prediction in Trace processor is initiating thread level speculation, while that in COSMOS is supporting instruction assignment. In addition, Trace processor also has global registers and its instruction window is not hierarchical.

The EDF instruction window is strongly influenced by Dualflow architecture [10]. It owes the basic idea of explicitly explaining data communication for Dualflow, which hybridizes control- and data-driven architectures. Instruction sequence is control-driven, while result forwarding between instructions is data-driven. Destinations of a result is explicitly explained in each instruction, removing associative lookup in instruction wakeup logic while Dualflow performs out-of-order execution. One of the disadvantages of Dualflow is explosion of program code. It is reported [10] that the code size is increased by more than 100% and approximately 50% of dynamic instructions are useless. It also requires considerable changes in the instruction set architecture. On the other hand, the EDF instruction window maintains binary compatibility by translating implicit result forwarding into explicit ones dynamically.

7 Conclusion

This paper introduced COSMOS microarchitecture for realizing large scale superscalar processors with high clock frequencies. The key features are the partitioned execution engine (clustered architecture), the high bandwidth instruction supply mechanism (non-consecutive basic block buffer), the hierarchical (decoupled instruction window) and simplified (EDF instruction window) instruction scheduling mechanism, and the mitigated bypass logic (variable latency pipeline). The contributions on ILP of the components are 9.4% and 3.7% improvements of the NCB and the decoupled instruction window and 4.6% and 3.6% degradations of the EDF instruction window and the VLP. The total ILP is improved by 4.3%. Since a processor based on COSMOS microarchitecture can operate at higher frequencies than the superscalar processor with an equivalent scale, the total processor performance can further improved. This confirms that COSMOS is one of the promising candidates for future superscalar processors.

One of the primary future studies dealing with COSMOS microarchitecture is evaluating the cluster assignment mechanism based on the value prediction. This can affects the total processor performance. Currently, we expect that this is not a severe issue for COSMOS due to the following consideration. The degradation of ILP which is occurred by cluster partitioning is at most 6% [14]. The degradation in COSMOS can not be worse, since we can remove the proposed cluster assignment strategy. Therefore, the total ILP including the clustering effect is only 1.3% degradation. This is easily compensated by COSMOS's high clock frequencies.

Acknowledgment

This work is supported in part by the Grant-in-Aid for Scientific Research from Japan Society for the Promotion of Science (No.12780273 and No.13558030). Toshinori Sato was supported in part by the grant from Fukuoka Industry, Science & Technology Foundation (No.H12-1).

References

[1] B.Black, J.P.Shen, "Scalable register renaming via the quack register file," Tech. Rep. CMuART-2000-01, Department of Electrical Computer Engineering, Carnegie Mellon University, 2000.

[2] D.Burger, T.M.Austin, "The SimpleScalar tool set, version 2.0," ACM SIGARCH Computer Architecture News, 25(3), 1997.

[3] R.Canal, J-M.Parcerisa, A.Gonzalez, "Dynamic cluster assignment mechanisms," 6th Int. Symp. on High Performance Computer Architecture, 2000.

[4] R.Canal, A.Gonzalez, "A low-complexity issue logic," 14th Int. Conf. on Supercomputing, 2000.

[5] G.Z.Chrysos, J.S.Emer, "Memory dependence prediction using store sets," 25th Int. Symp. on Computer Architecture, 1998.

[6] T.M. Conte, K.N. Menezes, P.M. Mills, B.A. Patel, "Optimization of instruction fetch mechanisms for high issue rates," 22nd Int. Symp. on Computer Architecture, 1995.

[7] J-L.Cruz, A.Gonzlez, M.Valero, N.Topham, "Multiple-banked register file architectures," 27th Int. Symp. on Computer Architecture, 2000.

[8] K.I.Farkas, P.Chow, N.P.Jouppi, Z.Vranesic, "The Multicluster architecture: reducing cycle time through partitioning," 30th Int. Symp. on Microarchitecture, 1997.

[9] F.Gabbay, "Speculative execution based on value prediction," Tech. Rep. #1080, Department of Electrical Engineering, Technion, 1996.

[10] M. Goshima, N. H. Ha, A. Agata, H. Mori, S. Tomita, "Proposal of the Dualflow architecture," 12th Joint Symp. on Parallel Processing, 2000.

[11] T.Hara, H.Ando, C.Nakanishi, M.Nakata, "Performance comparison of ILP machines with cycle time evaluation," 23rd Int. Symp. on Computer Architecture, 1996.

[12] Y.Kondo, N.Ikumi, K.Ueno, J.Mori, M.Hirano, "An early-completion-detecting ALU for a 1GHz 64b datapath," Int. Solid State Circuit Conf., 1997.

[13] M.H.Lipasti, C.B.Wilkerson, J.P.Shen, "Value locality and load value prediction," Int. Conf. on Architectural Support for Programming Languages and Operation Systems VII, 1996.

[14] S.Palacharla, N.Jouppi, J.Smith, "Complexity-effective superscalar processors," 24th Int. Symp. on Computer Architecture, 1997.

[15] S.J.Patel, D.H.Friendly, Y.N.Patt, "Critical issue regarding the trace cache fetch mechanism," Tech. Rep. CSE-TR-335-97, Department of Electronics Engineering and Computer Science, University of Michigan, 1997.

[16] E.Rotenberg, S.Bennet, J.Smith, "Trace cache: a low latency approach to high bandwidth instruction fetching," 29th Int. Symp. on Microarchitecture, 1996.

[17] E.Rotenberg, Q.Jacobson, Y.Sazeidas, J.Smith, "Trace processors," 30th Int. Symp. on Microarchitecture, 1997.

[18] T.Sato, "Data dependence speculation using data address prediction and its enhancement with instruction reissue," 24th Euromicro Conf., 1998.

[19] T.Sato, "Decoupling recovery mechanism for data speculation from dynamic instruction scheduling structure," 5th Int. Euro-Par Conf., 1999.

[20] T.Sato, I.Arita, "The KIT COSMOS processor: introducing CONDOR," Int. Conf. on Parallel and Distributed Processing Techniques and Applications, 2000.

[21] T.Sato, "Evaluating trace cache on moderate scale processors," IEE Proc. on Computers and Digital Techniques, 147(6), 2000.

[22] T.Sato, Y.Nakamura, I.Arita, "Revisiting direct tag search algorithm on superscalar processors," Workshop on Complexity-Effective Design held in conjunction with 28th Int. Symp. on Computer Architecture, 2001.

[23] T.Sato, I.Arita, "Execution latency reduction via variable latency pipeline and instruction reuse," 7th Int. Euro-Par Conf., 2001.

[24] S.Wallace, N.Bagherzadeh, "A scalable register file architecture for dynamically scheduled processors," Int. Conf. on Parallel Architecture and Compilation Techniques, 1996.

[25] T.Yamamoto, T.Sato, I,Arita, "The KIT COSMOS processor: eliminating ineffectual branch instructions via concurrent dynamic optimization," COOL Chips IV, 2001.

[26] K.C.Yeager, "The MIPS R10000 superscalar microprocessor," IEEE Micro, 16(4), 1996.