

Evaluation of Static and Dynamic Scheduling for Media Processors

Jason Fritts and Wayne Wolf

Abstract

This paper presents the results of an architecture style evaluation that compares the performance of static scheduling and dynamic scheduling for media processors. Existing programmable media processors have predominantly used statically-scheduled architectures. As future media processors progress to higher frequencies and higher degrees of parallelism, the dynamic aspects of processing become more pronounced and dynamic hardware support may be needed to achieve high performance. This paper explores many of the dynamic aspects of media processing by evaluating various fundamental architecture styles and the frequency effects on those architecture models. The results indicate that dynamic out-of-order scheduling methods enable significantly higher degrees of parallelism and are less susceptible to high frequency effects. Consequently, dynamic scheduling support may be necessary for maximizing performance in future media processors.

Conclusions

This paper examined the range of static and dynamic scheduling and explored the effects of increasing frequency on those architecture models in order to determine the most appropriate scheduling model for media processing. It was found that statically-scheduled VLIW architectures and in-order superscalar processors perform comparably, while dynamic out-of-order scheduling performs significantly better, with a 64% average improvement over VLIW performance. An evaluation of compiler techniques indicated that hyperblock optimization has the best performance, but its minimal performance improvement over superscalar optimization probably does not warrant the cost of additional hardware for conditional execution. Finally, a variety of processor issue widths were evaluated for the three architecture models. Only minimal performance gain was found when using more than 4 issue slots. Furthermore, the 2-issue out-of-order superscalar demonstrated better average performance than both the 8-issue VLIW and 8-issue in-order superscalar.

This study also examined the impact of increasing frequency on all architecture models to

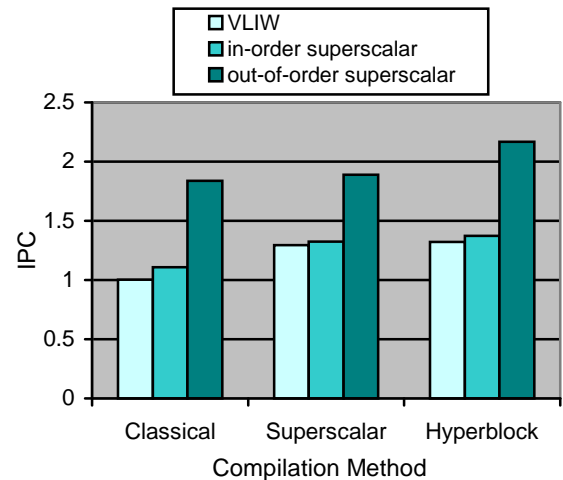


Figure 1 – Comparison of three processor models, simulated with real caches.

determine the most suitable scheduling model for future high frequency media processors. The effect of increasing frequency was analyzed via two experiments, the first of which examined scheduling performance on three processor frequency models with varying instruction and memory latencies, while the second examined the performance variation between immediate and delayed bypassing. The experiments showed doubling processor frequency degrades parallelism by an average of 16% (for a total execution speedup of 72%), while adding a one-cycle delay to bypassing causes a significant IPC degradation of 32%. Among the various architecture and compiler optimization alternatives, dynamic out-of-order scheduling and superscalar optimization are about 30% less susceptible to IPC degradation than the other architectures and compiler methods.

Overall we have shown that some dynamic hardware support may be desirable in future media processors. Dynamic out-of-order scheduling demonstrated significantly better performance than other architectures, plus it is less susceptible to high frequency effects of longer instruction and memory latencies, and delayed bypassing. The cost of out-of-order scheduling and high frequency design in media processors is currently prohibitive, but as these costs continue to diminish, we expect high frequency and out-of-order scheduling will be used to maximize performance in future media processors.