

Summary

The Cost of Flexibility in Systems on a Chip Design for Signal Processing Applications

Kiran Yellajosula

April 30, 2005

1 To be completed before class

What are the problems solved by this paper? (50 words)

The cost of flexibility for computing in terms of energy, chip area is not well understood. The general approach of providing unlimited flexibility for applications is inefficient due to the lack of proper metrics which considers flexibility as a dimension in design. The new design considerations like low power and low cost computation requires a careful consideration of flexibility.

What are the approaches attempted by this paper? (50 words)

This paper proposes to use flexibility as part of the algorithm/architecture codesign space. The authors defines efficiency metrics for energy (MOPS/mW) and area (MOPS/mm²) and use them to compare a variety of designs and architectures for signal processing applications. The authors propose to study and understand situations where it is better to build dedicated hardware or hardware with least flexibility.

What are the main conclusions of this paper? (50 words)

The authors study Digital signal processing applications and conclude that flexibility leads to upto 3 orders of inefficiency in energy and cost. . The architectural factors that contribute to the inefficiency were identified and analyzed and the amount of parallelism was also identified as a critical parameter. The authors conclude that dedicated solutions with high efficiency can be built once flexibility is reduced.

2 To be completed after class

Did this paper address an important issue? Explain. (100 words)

In the last few years, computing solutions of low cost and low energy consumption are being investigated. The authors look at unlimited flexibility as a possible reason for the inefficiency in the design space of architecture and algorithms. The authors point out the technological assumptions of the past are no longer valid and that hardware complexity is not a major concern in design. The authors investigate the design space to understand the key trade-offs and the costs of important architecture parameters and further apply them to DSP architectures for high-performance and energy-sensitive portable applications.

Are the proposed approaches valid? Describe its strength and weakness. (100 words)

Strengths:

1. They propose ideal solutions for the ideal environments, i.e., reduce the cost of the processor and energy consumed.

2. Since the amount of heat produced decreases automatically, the cost of cooling the applications decreases, if such a cost is involved.
3. The amount of chip area consumed decreases.
4. The number of errors to deal with reduces (AVF and other error rates depending on the area of the chip decreases)
5. The problem of dividing the problem space into hardware and software is hard and though this method can be used as hints to solve the problem, it is still hard.

Weakness:

1. The design and testing time for every product increases due to the fact that every product is to be designed separately.
2. The authors propose metrics for studying cost and energy of a design but it is still unclear if a proposed design is how close to the optimal design. A designer might have to spend considerable time reiterating the design as there is no way to determine efficiency.
3. The design may not be intergrable with existing software, there by leading to new cost in designing software.

Do the results support the conclusions? Explain. (100 words)

The authors study apply and study their metrics to compare many general purpose processors, the popular embedded ARM processor and DSP Comm. processor. Figure 2 shows that dedicated solutions have better power and area efficiency when compared with Microprocessors and General purpose DSPs. The authors relate flexibility with the number of supported operations and study the area and capacitance per operation. They explain that parallelism of operations is hard to obtain in C programs. Figure 6 compares the power and area used by a DSP processor, DSP extension and dedicated hardware and this supports the conclusion that with reduced flexibility, as the pipeline stages and huge circuitry decreases, thereby leading to the observations. The results thus support the conclusions that it might be efficient to design dedicated components to reduce energy consumption and decrease cost.

Describe the potential future works? (100 words)

The authors can apply this method to most of the embedded applications and propose a solution closer to the optimal. It would be interesting if this approach is applied to sub-parts of the current general purpose processors and the components are redesigned to optimize existing general purpose applications which have unlimited flexibility. The authors study DSP and FPGA applications alone but exclude co-processors. This approach can be applied to the multi-core processors and exclude unnecessary instructions and hardware components.