Effective Design Exploration through Multi-Objective High-Level DSP Synthesis

M. S. Bright
Cardiff University
Circuits and Systems Research Group
Electronic Engineering Division
Newport Road
Cardiff CF24 0YF
UK
msbright@iee.org

T. Arslan
University of Edinburgh
SLI Institute
The Kings Buildings
Mayfield Road
Edinburgh EH9 3JL
UK

ABSTRACT - This paper presents a CAD framework that allows exploration of high-level designs within a multi-objective framework. The designs are manipulated with high-level transformations in a sequence determined by an external stochastic process. The effectiveness of the framework is evaluated with DSP systems, illustrating the ability for pointing the designer to a set of designs with 'good' power-area tradeoffs.

Index terms—high-level synthesis, low-power, VLSI

I. INTRODUCTION

Fast time to market is increasingly becoming important in today's electronic industry [1,2]. This is especially true in the portable electronics market where complex System-On-Chip (SOC) based systems are designed under multiple design constraints that specify power as well as throughput and area constraints. Complex SOC based design is centered around high level synthesis tools [3,4,5], where strategical decisions effecting cost and performance are made. In addition, such high-level synthesis tools provide the user with the ability to analyze high-level design alternatives before time consuming synthesis steps are performed, thus significantly compressing the design time. For optimal and rapid design there is a continuos demand for effective CAD tools which allow manipulation of the high level designs in order to map onto physical designs (post logic synthesis) that best satisfy design constraints [6].

A technique for high-level manipulation of power, speed and area was presented in [7]. This technique uses high-level transformations to manipulate Digital Signal Processing (DSP) algorithms represented as behavioral Data Flow Graphs (DFGs). The transformations are applied to the DFG specification prior to synthesis, manipulating the DFG to optimize the results of the synthesis process. As illustrated in figure 1, a key disadvantage of this technique is that pursuit of minimum power consumption leads to considerable area increase.

The technique increases the speed of the DSP algorithms (enabling reduced voltage operation) by increasing the amount of operations that can be processed in parallel, hence increasing the area requirements for the processing units. Therefore, power and area are conflicting objective parameters for VLSI devices. Improvement in one objective is achieved at the expense of degradation in the other. The design process must trade-off between these conflicting parameters to determine the best solution that meets specified requirements. Such a design process requires the consideration of multiple design objectives; it requires a multi-objective optimization framework.

A typical technique for multi-objective optimization is to explore the solution space and present a range of alternative non-dominated solutions (NDS) that are each optimal for a single parameter [8,9]. The alternative solutions can then be analysed by the designer to select the solution that satisfies the specified requirements. This removes the need to prioritise parameters during the optimisation process, which can lead to sub-optimal designs. The presentation of a suite of alternative designs enables the designer to use his expert knowledge in selecting the best solution.

This paper describes a new approach to high level synthesis of complex DSP systems under multiple design constraints. The technique allows manipulation of a high level design, represented as a DFG, through the application of various high level transformations which effect performance parameters such as area, power, speed. Due to the complexity of the high-level design space, the sequence of transformations applied by a stochastic process [10], which uses a class of search algorithms called Genetic Algorithms (GAs)

Previous CAD tools based on high-level transformations and stochastic search techniques have shown success in solving the low-power synthesis problem. Examples of such stochastic CAD tools are the HYPER [7] and GALOPS [10] systems. GALOPS, which uses a GA to solve the complex problem of high-level low power design, has been used as the basis of the multi-objective search tool described in this paper.

0-7803-5632-2/99/$10.00 © 1999 IEEE

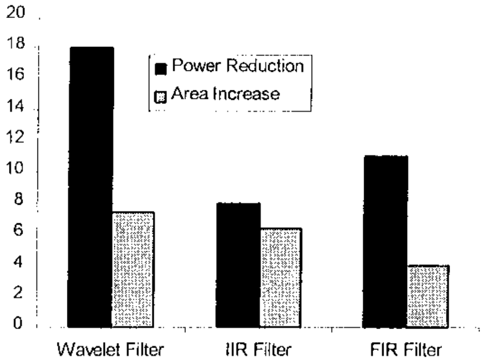

Figure 1. Power Reduction and Area Increase for VLSI-Based DSP Algorithms. Results collated from HYPER[7] and GALOPS[10]

A GA performs a parallel search of the solution space to determine the optimum design. During the optimisation process the GA evaluates many alternative solutions. The information obtained from these evaluations can be used to illustrate the trade-offs between different parameters in the optimisation problem. The trade-off between competing parameters is typically presented as a Pareto-surface [11].

The paper describes the design approach and its embodiment within the stochastic search technique for multi-objective design synthesis/exploration. The effectiveness of the framework is evaluated with a range of DSP systems, illustrating the ability for guiding the designer, through a set of Pareto-surfaces, to an effective set of designs with 'good' area-power tradeoffs.

II. IMPLEMENTATION

The design procedure commences with a DFG representation of the DSP algorithm. This is obtained after translation of design specification to the appropriate algorithmic description. The DFG forms a design seed that initializes the GALOPS stochastic process. GALOPS uses the seed to create a number of designs, with different performance specifications, through the application of one or more transformations from a *transformation set* which includes the following [7]:

1. Retiming
2. Pipelining
3. Loop Unfolding
4. Automatic Pipelining

The above transformation set forms the input for an evolutionary process in which designs are selected based on satisfying area, speed, and power constraints. These designs are manipulated through the application of transformations with high probability of producing better designs.

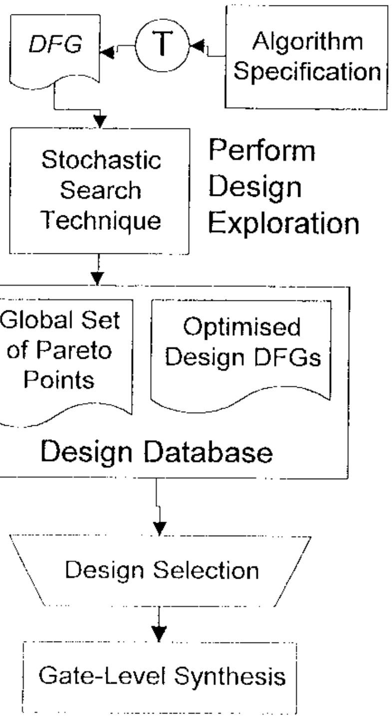

Figure 2. Identification of Pareto-Points in GALOPS

The GALOPS tool, in addition to containing power-optimizing transformations, also contains a high-level power estimation strategy discussed in [10]. This power estimation strategy extrapolates area, capacitance and speed data from the high-level DFG to estimate the VLSI implementation parameters. This enables the stochastic process to select designs based on their performance.

A key feature of GALOPS is its parallel search technique, evaluating many designs throughout the search process. Modification of the standard GA search technique enables collection of a set of NDS or Pareto-points from this parallel search, rather than the production of a single low-power solution as those illustrated in figure 1.

The search process is illustrated in figure 2. The original specification is translated (T) into a DFG representation before submission to the stochastic search technique. The stochastic search applies the high-level exploration techniques to generate a set of Pareto-optimal designs. These designs are generated through a 2-stage process:

- The first stage uses the GA framework established in GALOPS to produce a set of minimal-power designs for a range of solutions over the area axis.
- The second stage then processes each low-power design in turn, pruning all designs that have a larger area than that found for previous designs.

Executing both of these steps produces a *design database*. This database comprises the NDS (the area-power Pareto-optimal points) and the set of related high-level designs (one for each NDS).

The design data is presented to the VLSI designer as a set of NDS on a Pareto-optimal surface, showing the power-area tradeoffs for the original high-level design. Through analysis of the Pareto-surface information the designer selects the design that best meets the implementation requirements. The associated design from the *design database* can then be fed through the subsequent stages of the VLSI synthesis process.

As this technique contains attributes of the GALOPS tool it embodies the main advantages of that tool, specifically the ability to perform this optimisation process on a wide range of DSP algorithms of varying function and complexity.

III. RESULTS

This section illustrates the area and power trade-offs for a range of common DSP benchmark designs often used in evaluation of synthesis tools. The designs cover a range of recursive and non-recursive signal processing algorithms. AVEN8DI is an 8th order direct implementation of an Avenhaus filter [7]. BIQAUD3 is the 3rd Order Biqaudratic filter presented in [12]. GMLAT4 is a 4-stage Gray-Markel filter [13] and LMS5 is a 5th order least mean square algorithm [13].

Figure 3 depicts the Pareto-surface for each of the benchmark designs on a single chart. The power consumption for each design is expressed as a percentage of the power consumption of the original design, to illustrate the power savings that can be obtained with the corresponding area. A solid line denotes the pareto-surface.

The pareto-surface is presented as a straight line rather than a curve joining the pareto-points. This is due to the nature of the VLSI synthesis problem. Solutions exist at discrete points in the solution space; a curve may imply that there is a range of solutions between two points when no feasible solutions actually exist between those points.

Each Pareto-surface follows the trend of increased area enabling larger power reductions (due to increased speed allowing a reduction in supply voltage). The advantages of the Pareto-surface information can be demonstrated with the example of the AVEN8DI benchmark in selecting a solution that balances power and area reductions. The single-point, lowest-power solution that would be generated by a greedy low-power optimizer specifies a design with a power consumption of 19.64% and an area of 23mm^2 (an increase of approx. 450% compared to the initial solution). The pareto-surface data illustrates that a power consumption of 20% is achievable with an area of 13mm^2 (an area increase of approx. 260%). Using the pareto-chart the VLSI designer can decide whether the further increase in area is worth the small further reduction in power. The graph presents information for a family of solutions to the VLSI designer; the required solution can be selected based upon the particular implementation criteria. A similar analysis could be applied to the other benchmark design Pareto-surfaces.

IV. CONCLUSIONS

The paper has presented an effective multi-objective framework within which different design topologies could be selected for fast synthesis turnaround time. The effectiveness of the approach has been demonstrated in a low-power design environment demonstrating a high degree of flexibility. The results obtained with a number of DSP systems indicate that overall power savings of up to 80% are achievable while maintaining area increases significantly below that indicated with single-point solutions.

V. REFERENCES

1. V.K. Madisetti, *VLSI Digital Signal Processors : An Introduction to Rapid Prototyping and Design Synthesis*, Newton, MA: Butterworth-Heinman, 1995
2. V.K. Madisetti, "Rapid digital system prototyping: current practice, future challenges," in *IEEE Design and Test of Computers*, Fall 1996, pp. 12-22
3. G. De Micheli, "High-level synthesis of digital circuits," in *IEEE Design and Test of Computers*, Oct. 1990, pp. 6-7
4. D.D. Gajski, L. Ramachandran, "Introduction to high-level synthesis," in *IEEE Design and Test of Computers*, Winter 1994, pp. 45-54
5. R.A. Walker, "The status of high-level synthesis," in *IEEE Design and Test of Computers*, Winter 1994, pp. 42-43
6. D. Singh, J.M. Rabaey, M. Pedram, F. Catthoor, S. Rajgopal, N. Sehgal, T.J. Mozdzen, "Power conscious CAD tools and methodologies: a perspective," in *Proc. of the IEEE*, vol. 83, no. 4, April 1995, pp. 570-594
7. A.P. Chandrakasan, M. Potkonjak, R. Mehra, J. Rabaey and R.W. Broderson, "Optimizing power using transformations," in *IEEE Trans. On CAD Of Integrated Circuits and Systems*, vol. 14, no. 1, Jan 1995, pp. 12-31
8. C.M. Fonseca, P.J. Fleming, "An overview of evolutionary algorithms in multiobjective optimization," in *Evolutionary Computation 3(1)*, Massachusetts, MA: MIT-Press, 1995, pp. 1-16
9. H. Esbensen and E.S. Kuh, "Design space exploration using the genetic algorithm," in *Proc. IEEE Int.*

Symposium on Circuits and Systems, ISCAS 96, Atlanta, USA, May 1996, pp. 500-503

10. M. S. Bright and T. Arslan, "A genetic algorithm for the high-level synthesis of DSP systems for low power," in *Proc. IEE/IEEE Conf. on Genetic Algorithms in Engineering Systems, Innovations and Applications (GALESIA '97)*, Glasgow, UK, 2-4 Sept. 1997, pp. 174-179
11. D.E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*, Reading, Mass: Addison-Wesley Publishing Co. Inc., 1989
12. J.M. Rabaey, C. Chu, P. Hoang, M. Potkonjak, "Fast prototyping of datapath-intensive architectures," in *IEEE Design and Test Of Computers*, June 1991, pp. 40-51
13. P.R. Gelabert and T.P. Barnwell III, "Optimal automatic periodic multiprocessor scheduler for fully specified flow graphs," in *IEEE Trans. on Signal Processing*, vol. 41, no. 2, Feb 1993, pp. 858-888

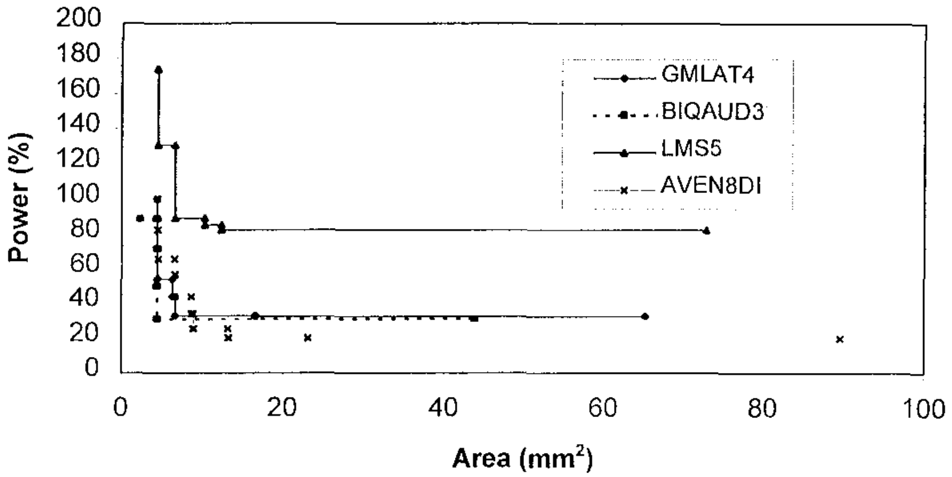

Figure 3. Pareto-surface charts for benchmark DFG algorithms