

Generation of Custom DSP Transform IP Cores: Case Study Walsh-Hadamard Transform

Fang Fang James C. Hoe Markus Püschel Smarahara Misra
*Department of Electrical and Computer Engineering
Carnegie Mellon University, Pittsburgh, PA*

Abstract

Hardware designers are increasingly relying on pre-designed DSP (digital signal processing) cores from IP libraries to improve their productivity and reduce design time. Unfortunately, static DSP cores also introduce an inefficiency because a designer cannot make application-specific trade-offs. We are proposing to automatically generate *customized* DSP cores that can be tailored for specific design requirements. This parameterized core generation technology enables a designer, with no specific background in DSP transform mathematics, to quickly create and evaluate different design choices to determine what is most suitable for his/her application. In this talk, we present a generator for the Walsh-Hadamard Transform (WHT). The generator accepts as input the WHT's sample size and two implementation parameters that control the degree of hardware reuse. The output is an RTL-level Verilog description of the desired implementation. We present results from exhaustive generation of WHT₆₄ and compare the different designs in terms of their resource requirements, computation latency and throughput.

DSP Transform IPs

Creating an optimized hardware implementation of DSP (digital signal processing) transforms requires combined expertise in both transform mathematics and hardware design to take advantage of the full design space available through algorithmic and architectural choices. Although the field of DSP hardware implementation has been well studied, it remains a highly specialized field with a high entry threshold. Instead of hand-crafting custom implementation of DSP transforms, typically, hardware designers make use of static, pre-designed IP blocks from ASIC/FPGA vendors and third-party IP libraries. These IP libraries provide implementations of common DSP transforms of specific sample sizes and data formats. For example, Xilinx's LogiCore library provides implementations of FFT for $N=16, 64, 256$ and 1024 on 16-bit 2's complement complex numbers. An obvious shortcoming of these pre-designed IP blocks is the loss of customization. A static IP makes no allowance for achieving higher performance (more concurrent computations) when more hardware resources or power budget is available, nor can a static IP trade excess performance for reduced size or power consumption.

We have been working on the technology to systematically and automatically generate DSP cores that can be optimized with respect to user specified constraints on performance, size, power, numerical accuracy, and I/O bandwidth. Our approach is based on machine-manipulatable formula representations of fast transform algorithms. By encapsulating the mathematics of DSP transforms in a few well-designed transform factorization rules and algebraic formula manipulations rules, we can develop a systematic approach to explore the algorithmic space of linear DSP transforms. Here we present our experience in developing a parameterized generator for the Walsh-Hadamard Transform (WHT).




Figure 1: Pease WHT algorithm. From left to right: completely flattened, horizontally folded, fully horizontally and vertically folded. Two lines meeting at a node means addition, and a dashed line means scaling by -1 .




Figure 2: From left to right: area, latency, and throughput variations over the 24 implementations of WHT_{64}

Generation of Walsh-Hadamard Transform

We chose the Walsh-Hadamard transform (WHT) as our initial case study because this very simple transform has a dataflow that is typical for many DSP transforms (e.g. discrete Fourier transform). Our generator is based on the Pease algorithm for the WHT, which we express in a formula notation as

$$\text{WHT}_{2^n} = ((I_{2^{n-1}} \otimes F_2) L_2^{2^n})^n, \quad F_2 = \begin{pmatrix} 1 & -1 \\ 1 & -1 \end{pmatrix} \quad (1)$$

where ‘ I ’ denotes an identity matrix, ‘ \otimes ’ the Kronecker product of matrices and ‘ L ’ is a stride permutation. Figure 1 shows a dataflow representation of (1) for $n = 3$. In general, (1) is a grid of $(2^{n-1} \times n)$ F_2 blocks, which is reflected by the constructs ‘ $I_{2^{n-1}} \otimes$ ’ and ‘ $(\cdot)^n$ ’. The rectangular grid structure exhibits the possibility of *horizontal folding* and *vertical folding* (i.e. reusing the same F_2 hardware instances over multiple columns or rows of the conceptual F_2 grid). Figure 1 shows, for $n = 3$, a block diagram of a horizontally folded WHT (middle) and a horizontally and vertically folded WHT (right). Our WHT generator currently accepts h , v , and n as degrees of freedom in core generation where h and v represent the degree of folding in the horizontal and vertical dimensions. For WHT of size 2^n , there are $H \times V$ folding strategies, where H and V are the number of divisors of n and 2^{n-1} , respectively.

For $n = 6$, the WHT core generator produces $4 \times 6 = 24$ implementations. Figure ?? shows how resource usage, latency and throughput varies over these 24 design choices. These results are estimates produced by Synplify after compiling the generated Verilog descriptions for Xilinx XCV????-??. The figures indicate that resource usage reduces almost linearly with increase in either h or v . The latency is not affected by horizontal folding, but increases with the degree of vertical folding. In our implementation, columns of F_2 ’s are pipelined so throughput decreases with horizontal folding. Vertical folding also reduces throughput by increasing the number of times an F_2 must be reused. In summary, folding reduces the resource requirement, but at the same time, reduces the computation performance. Our core generation methodology can help hardware designer choose a correct trade-off points over the space of possibilities.