

Leveraging High Density Packaging for High Performance DSP Systems

Pronita Mehrotra, Vikram Rao, Tom Conte, Paul Franzon
Department of Electrical and Computer Engineering,
North Carolina State University, Raleigh, NC 27695
{pmehrot,vsrao,conte,paulf}@eos.ncsu.edu
Ph: (919) 515 7351 Fax: (919) 515 2285

Abstract

The high connectivity of the SHOCC technology can be exploited to increase the number of memory channels in a DSP system. This paper describes the physical and logical architecture of a high performance FFT system enabled by a combination of high density packaging and good memory management schemes, with an emphasis on signal integrity issues.

Introduction

The aim in this paper is to demonstrate how high density packaging can lead to large gains in system performance if the design team is willing to re-engineer the chip set. The example used in this paper is that of an FFT engine for use in Synthetic Aperture Radar (SAR) applications. This application is characterized by large total data volumes (1GB or more), necessitating the heavy use of DRAM, high performance requirements 10,000,000 point FFTs at TFLOP rates), and tight power and system size budgets.

In this paper we describe how we used high density packaging, and developed new architectures to design an FFT engine that can compute a million-point FFT in 1.31ms. This is at least an order of magnitude better than that achievable with conventional off-the-shelf systems.

Physical Architecture

The physical architecture is shown in Figure 1. The chip set contains 1 GByte of memory distributed amongst 128 64Mbit DRAM chips, and four custom 1 sq.cm. micro-accelerator chips. Each micro-accelerator chip contains an array of 64 32-bit multiply-accumulate units and 32 memory interface channels. Each of the 60-pin DRAMs are edge-mounted onto a high density 8 cm x 8 cm SHOCC interposer substrate [1] using a previously developed solder-bump edge mounting technique [7]. Each memory chip is wired directly to one (and only one) memory port on a micro-accelerator chip. The high density substrate thus contains 128 independent 16-bit memory buses, which together with the control and inter-accelerator bus, make up approximately 8000 total nets to be accommodated. Each micro-accelerator has 2500 signal pins.

One significant challenge in the physical design is the tight wiring requirements under the micro-accelerator chips. We conducted a routability and signal integrity analysis to make sure that this level of wiring was achievable.

Routability and Signal Integrity Analysis

The SHOCC technology is capable of permitting very high levels of off-chip connectivity [1]. The SHOCC substrate, in our design, consists of a Si base with alternating dielectric and metal layers. The SHOCC substrate stack up is shown in Fig.2 The stack-up consists of three signal layers (S1, S2, S3), power and ground. A microstrip stackup is used to maximize the amount of decoupling capacitance present between the power and ground planes, which is 19.2pF/cm² in our case. We also evaluated a two-signal layer stackup but determined that the crosstalk associated with the high density breakout wiring was too high.

Two of the micro-accelerator chips read in 32 complex numbers each, in 4 cycles, and put the results on the 256-bit high speed bus for the other two chips. For 64-bit complex numbers, the number of signals to be routed out of a chip, including control signals, is around 2000. We have taken a stricter estimate of 2500 signal I/Os that need to be routed out for our analysis. Also, the ratio of power and ground bumps to I/O bumps has been assumed to be 1:1. This makes the total number of bumps on a die to be 5000. For a 1cmx1cm chip, the bump pitch required to support these many bumps is 140 μ . This is easily achievable in existing technology. A two-stage breakout routing approach has been taken where the pitch in the final stage is 36 μ as shown in Fig.3. The first pitch corresponds to the initial breakout, i.e. the pitch required for the 2500 I/Os to break out of the 1cmx1cm perimeter in 2 layers. The second pitch is required to move to an XY routing from the initial breakout and we have assumed that the perimeter in this case does not exceed 2cmx2cm. The final pitch corresponds to the pitch that can be supported under the DRAMs.

To carry out a noise and timing analysis for this substrate, we first determined the R,L and C parameters for the SHOCC lines using Ansoft's Maxwell Q-3D. To do the crosstalk noise analysis, we extended the model used in [2] to include couplings from the four nearest neighbors. We then simulated a SHOCC line, driven by a 5 stage CMOS driver (with a stage ratio of 3), in SPICE. The on-chip and the solder bump parameters were taken from [3]. The crosstalk noise for a 5 stage driver in various routing stages is shown in Fig.4. The crosstalk noise for a signal in the top layer is more than double that for a signal in the bottom layer in the breakout regions (Fig 4(a) and (b)). The presence of a local ground in these regions makes the bottom signal layer behave like a stripline instead of a microstrip. The

difference in the final case is not as much where the signal line in both the top and the bottom layer is a microstrip. The bottom layer, being closer to the power/ground planes shows lesser noise. The reflection noise is shown in Fig. 5. Delays for the SHOCC line in the three stages is shown in Fig.6. Since the capacitance of the SHOCC line in the top layer is much less than that on the bottom layer, the propagation delay for a SHOCC line on the top layer is lesser than that for a bottom layer. The total noise budget can be approximated by the root sum square of the individual components of the noise. To compute the crosstalk noise, we add the crosstalk noise components from the different regions. The escape lengths in the two breakout stages are 1cm and 0.8cm. The total crosstalk noise, for a total routing length of 4cm, is 0.56V. With a reflection noise component of 0.03V and an SSN of 0.2V [4], the total noise comes to 0.6V. This is within the noise budget of 0.7V (for the 0.25 μ technology). The worst-case off-chip skew on an 8cmx8cm substrate is around 0.2ns. After adding factors for on-chip skew and jitter, we can have a cycle time of at least 2ns. This gives us an I/O bandwidth > 100GB/sec.

Logical Architecture

The high bandwidth of the SHOCC (Seamless High Off-Chip Connectivity) technology can be used in increasing memory channels in the system. For most memory-starved applications, this would directly lead to an improvement in performance. For the FFT engine, we used the extra channels to make a Radix-64 engine as opposed to standard Radix-4 or Radix-8 systems. For a 0.25 μ technology, a 32 bit multiplier and adder takes up an area <1mm². Taking into account the area required for memory interface and some control unit, we can accommodate around 64 of these units. Using 4 such chips, we can build a radix-64 FFT engine. Each chip holds 4 radix-8 units and 4 such chips are required to build the 64-point FFT. For the 0.25 μ technology, a 32-bit multiplication followed by an addition can be done in \approx 2ns. The FFT computation is pipelined in two stages. In each stage the 8-point FFTs are computed and the outputs are twiddled before passing the results to the next stage. The sequence of operations is shown in Fig. 7.

However, the overall key to this design is the ability to use all the DRAMs to exploit the full memory bandwidth available. We achieve this by rotating the outputs from the chips for each stage of the FFT using the relation given below:

$$\text{DRAM \#} = (\text{FFT \# \% 64}) + (\text{index \%64})$$

Where FFT # = $\lfloor (\text{index} / 64) \rfloor$ and index refers to the index number of the data ($0 < \text{index} < 1,048,575$)

This introduces the stagger of 64 in the data, which is required for the next stage. After the data is read in the next stage, it needs to go through only an on-chip shift register stage to be sorted in the right order before it goes to the multiplier/accumulator stage. In our analysis we have used DDR SDRAM (MT46V4M16) from Micron Semiconductor Products Inc[8]. In the chosen DDR SDRAMs, a random access request takes 60ns to serve (compared with less than 5ns for modern SRAMs). Thus if the data accesses were all truly random the total memory bandwidth of our system would be 15Gb/s and the FFT performance unacceptably low. We worked out a memory mapping scheme that generates bank, row and column addresses in such a way such that a row hop in the same bank is done only after a minimum of four 64-point FFT cycles. This hides the row precharge time of the DRAM and reads and writes can be done every 5ns i.e. SRAM performance with DRAM density. The twiddle factors required for the million-point FFT are generated on chip using a base vector set, stored in an on-chip SRAM bank. Twiddle factor generation, which requires floating point multipliers, is done in the two stages of the 8-point FFT which don't involve any multiplications.

Performance of the FFT system

The FFT engine computes a 64 point FFT in 20ns. A million-point FFT, can be done in 4 stages, with 16,384 64-point FFTs in each stage. The total time to compute the million-point FFT is therefore, 1.31ms.

Comparison with Conventional Off-the-shelf Systems

We compared the performance of our FFT engine with two other systems. The first one uses four BOPS Inc. DSP chips [5]. The system has 4 32-bit memory channels. Each chip has 4 PEs and each of the PE has 5 FP units. This system would take up approximately 80 sq.cm of PCB and would perform a million-point FFT in 21.5 ms, more than an order of magnitude slower than ours. We also compared our engine with a G4 Velocity Engine implementation of the FFT, using the Motorola's AltiVec technology. Computing a million-point FFT on their system takes 511ms, almost 400 times slower than ours[6].

Conclusions

We have shown a high-density packaging technology like SHOCC, can improve the performance of memory intensive systems by more than an order of magnitude. Our FFT system, implemented using this technology can compute a million-point FFT in 1.31 ms while remaining within the noise and timing budgets.

References:

- [1] L. Schaper, M. Dibbs, P. Garrou, C.C. Chow, Y. So, D. Frye, J. Wagner, J. Ousley, G. Baugher, R. Picard, G. Connor, D. Winn, P. Deane, "Seamless High Off-Chip Connectivity", IC/Package Design Integration, 1998, Proceedings, Pages: 39-44
- [2] S. Afonso, W.D. Brown, L.W. Schaper, J.P. Parkerson, "Signal Propagation on Seamless High Off-Chip Connectivity (SHOCC) Interconnects", EPEP, 1998, Pages:31-34
- [3] A. Deutsch, G.V. Kopsay, C.W. Surovic, B.J. Rubin, L.M. Terman, R.P. Dunne Jr., T.A. Gallo, R.H. Dennard, "Modeling and Characterization of long on chip interconnects for high performance microprocessors", IBM J. Res. Develop., Vol. 39, No. 5, September 1995, Pages: 547-56
- [4] J.T. Schaffer, A.W. Glaser, P.D. Franzon, "A High Throughput Implementation of the Data Encryption Standard and Triple-DES", submitted for publication to IEEE Journal of Solid State Circuits
- [5] <http://bopsnet.com/cores/>
- [6] R. Crandall, J. Klivington, "Supercomputer-style FFT library for Apple G4", 1999
- [7] M. Nakkar, A.W. Glaser, P. Franzon, K. Williams, M. Roberson, G. Rinne, "Three Dimensional MCM Package Assembly and Analysis," in Proc. IEEE/IMAPS Conference on High Density Packaging and MCMS, Denver, CO., May 1999, pp. 188-192.
- [8] DDR SDRAM (MT46V4M16), www.micron.com




Fig.1 Layout of memory and micro-accelerator chips




Fig.2 Substrate Stackup







Fig 3. Two stage breakout strategy



4 (a)



4 (b)



4(c)

Fig. 4. Crosstalk Noise for (a): initial breakout (b): XY routing and (c): final routing







Fig.5 Reflection Noise for final case



6(a)



6(b)



6(c)

Fig. 6. Worst case delays for (a): initial breakout (b): XY routing and (c): final routing




Fig.7 Sequence of operations for the FFT engine