IMPLEMENTING THE NEXT GENERATION OF TELECOMMUNICATION SATELLITE ON-BOARD DIGITAL PROCESSORS: A FEASIBILITY STUDY OF MODULE HARDWARE DESIGN

Mark Walker (mark.walker@astrium.eads.net)
Processor Products Group, Astrium-Space, Stevenage, Hertfordshire, England.

ABSTRACT: This paper describes a novel hardware scheme intended to realise future telecommunications satellite on-board digital processor performance targets, whilst meeting stringent power and mass constraints by using large ceramic substrates. The analysis takes a bottom-up approach, starting with trace geometry and substrate layer stack design, followed by architecture partitioning, trace routing capacity calculation, power distribution network analysis and surface budget estimation. The feasibility study, from which the work described in this paper is drawn, was completed in Spring 2006. Furthermore, the current substrate test piece design activity is briefly described and further work suggested. A mixture of Imperial and Metric units is used, as is common practice within the industry.

I. INTRODUCTION

The Astrium Processor Product Group's (PPG) [1] current generation of telecommunications satellite digital processors reached its zenith in the form of the Inmarsat 4 (I4) mobile payload [2], characterized as follows:

Inmarsat 4 (I4) processor key metrics:

- Over 200 digitally formed spot beams per satellite.
- Total routed capacity: 2 x 126 MHz.
- Mass: 2 x 80 kg (2 large units), power ≈ 2 kW.
- Two operational craft, launched 11-03-05 & 09-11-05.

Although I4 was the largest commercial digital telecommunications payload when launched, future missions will require a significantly increased processed bandwidth and power density, with reduced mass.

Figure 1: A deployed Inmarsat 4 satellite.

Each I4 payload consists of two units (forward and return), each containing about 25 modules that attach to a baseplate, with integral heatpipes below, and interconnect via an active backplane above. The modules consist of a printed circuit board (PCB), clamped within an alloy frame, with four high temperature co-fired ceramic (HTCC) multi-chip modules (MCMs) mounted on either side. Each MCM accommodates several wire bonded ASIC die. The main thermal path is provided by the MCM substrate, module frame and baseplate, while the module and backplane PCBs provide the majority of the electrical connectivity and also accommodate lower power packaged devices, including the analogue to digital converters (ADCs) and digital to analogue converters (DACs). The active devices interconnect using synchronous parallel busses operating at 16 to 50 MHz.

The Next Generation Processor (NGP) aims to achieve a large increase in total processed bandwidth for a range of possible missions including mobile, flexible C/Ku and broadband Ka, a wider range of available beams per processor and a significant reduction of mass and power. These objectives [3] are impossible to achieve by scaling the I4 architecture because power and mass both become impractically large. Several enabling technologies must therefore be simultaneously developed to realise NGP, including:

- Innovative algorithm and architectural design.
- Multi-mega gate, deep sub-micron ASICs.
- High-speed, low power ADCs and DACs.
- High-speed serial interconnect.
- High efficiency, low voltage power supplies.
- Advanced embedded test techniques.

The NGP hardware uses novel techniques to maintain the necessary low mass and volume, whilst achieving good thermal performance and reliability. The digital processor is implemented as a 3 stage Clos network [4], with the switching nodes replaced by ASICs, implemented in an inherently radiation hardened gate array technology, and the Clos' mesh interconnection implemented with multi-Gbit/s high-speed serial links.

The following feasibility analysis focuses primarily on module design. The thermal design margin, substrate trace routing capacity, power distribution network (PDN) design and surface budget allocation are scrutinized.

II. Analysis

The use of a ceramic substrate elegantly solves the thermal and mechanical problems associated with the ASICs but dictates material choice. This lack of choice makes the electrical design challenging because the substrate's dielectric and conductor electrical properties are poor in comparison to organic PCBs, especially when designing for multi-Gbit/s high-speed serial links with their wide signal bandwidths.

A. Trace Geometry

High-speed serial links typically use a point-to-point 100 Ohm differential controlled impedance channel that may be realized with edge-coupled symmetrical striplines in plane-trace-plane (PTP) layer stacks. A simple 2D electromagnetic (EM) field solver [5] was used to calculate the trace geometry for the target impedance, whilst varying the dielectric thickness, trace width and intra-pair separation. The relative permittivity was held constant. A sample of the results was also verified against another 2D [6] and two 3D [7, 8] EM field solvers.

Figure 2 shows the resulting trace geometries as a plot of trace width *vs.* separation for various standard dielectric thicknesses. As trace width increases, trace impedance is eventually dominated by trace to plane coupling and trace separation quickly becomes impractically large. Even a 4 thou (minimum width) trace requires a large (15.4 thou) intra-pair separation and the use of thick (15 thou) dielectric layers to achieve the target impedance, due to the dielectric's high relative permittivity. For comparison, a 100 Ohm differential stripline implemented in FR4, with a relative permittivity of 4.0 [9] (at 1 GHz) could have a trace width, separation and height between planes of: 4.0, 4.28 and 16.0 thou, respectively. The wide intra-pair trace separation and relatively thick dielectric lead to low horizontal and vertical routing density, respectively, while the narrow trace width, thin internal conductor thickness, dielectric loss tangent Debye [10] peak and high conductor resistivity, lead to high signal loss per unit length. A commercial FPGA vendor [11] recommends an 8 thou trace width for high-speed serial channels to ensure sufficiently low loss, although this is aimed at a link of 30” over FR4 / copper, plus two backplane connectors, where-as the NGP inter-ASIC links will be shorter.

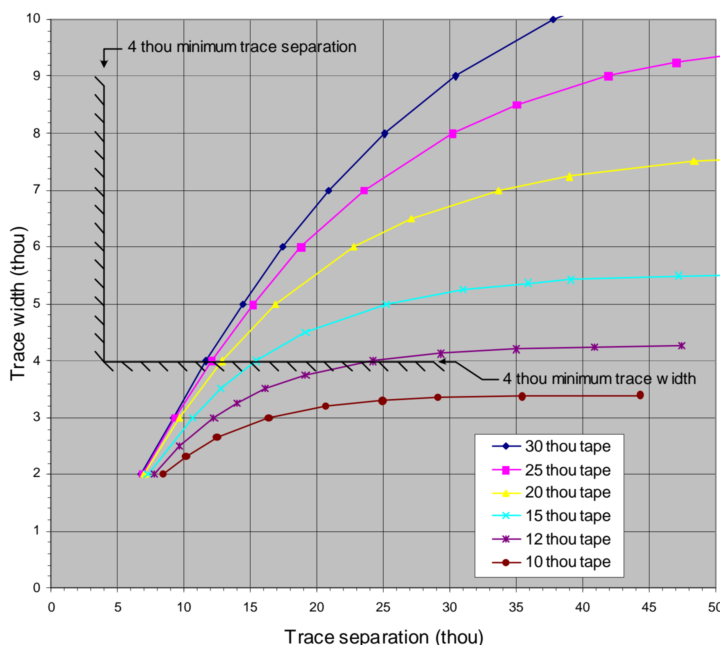

Figure 2: Stripline trace geometry, for Z_0 = 100 Ohms

The trace separations were calculated for standard dielectric thicknesses and integer trace widths ranging from 4 to 10 thou, with only 22 of the 90 combinations able to achieve the target impedance. The trace density was then calculated, assuming an inter-pair trace separation of twice the pitch, with a maximum value of 445 pairs per square inch, in comparison to 1,566 for the example FR4 trace geometry. Similar analysis was carried out for plane-trace-trace-plane (PTTP) offset edge-coupled and symmetrical broadside-coupled stripline geometries, yielding similarly low routing densities. The PTP layering with 15 thou dielectric, 4 thou trace width and 15.4 thou trace separation was baselined because it yields the highest routing density, previously quoted.

B. Substrate Layer Stack Design

The substrate manufacturing process restricts the maximum substrate thickness by both layer count and absolute thickness. The generic layer stack has a pair of closely spaced power planes just below the pad layers, with the core consisting of multiple PTP or PTTP layerings. PTP layer stacks with feasible layer counts were analysed for total thickness and routing density. The baseline stack achieves maximum manufacturable thickness and offers the greatest routing density of 108 pairs per inch of substrate, with the baseline trace geometry. The PTTP stacks offer slightly greater parallel routing densities, of up to 117 pairs per inch, or routing capacities of up to 352 pairs per horizontal square inch, but with the added constraints of either maintaining inter-pair trace separation or routing orthogonality on adjacent trace layers to maintain cross-talk isolation.

C. Architecture Partitioning

This analysis establishes the feasibility of the proposed architectures in terms of the substrate's ability to dissipate the ASICs thermal load and accommodate the trace routing. Vertical Clos partitioning was assumed, as shown in figure 3 so that the ASICs on a particular substrate do not interconnect with each other. Of the fourteen proposed processor architectures [12], the three variants of the baseline architecture were analysed for module power dissipation and horizontal (*i.e.* 'East-West') routing capacity, for hardware implementations using various numbers of modules. The baseline trace geometry, along with a conservative assumed 60% usable fraction of substrate width, due to breakout via arrays under the ASICs, was used to calculate an effective number of pairs per PTP routing layer of 47. The number of ASICs per module was varied and the thermal and routing margins calculated. Optimum partitioning has no spare ASICs per module and uses all the available modules. For a typical architecture and module design, this corresponds to 8 ASICs, with a 33% thermal margin and 15% routing margin. All 3 variants of the baseline architecture are feasible when partitioned across all available modules, but only the most conservative and therefore least attractive architecture is feasible with an alternative implementation using a reduced number of modules.

D. Vertical Routing Capacity

The module's ASIC and connector placement was considered to determine whether the substrate's vertical routing capacity (*i.e.* North-South, as opposed to East-West) was sufficient. The conceptual routing flow for four ASICs, each with 32 ingress and 32 egress pairs, is shown in Figure 3:

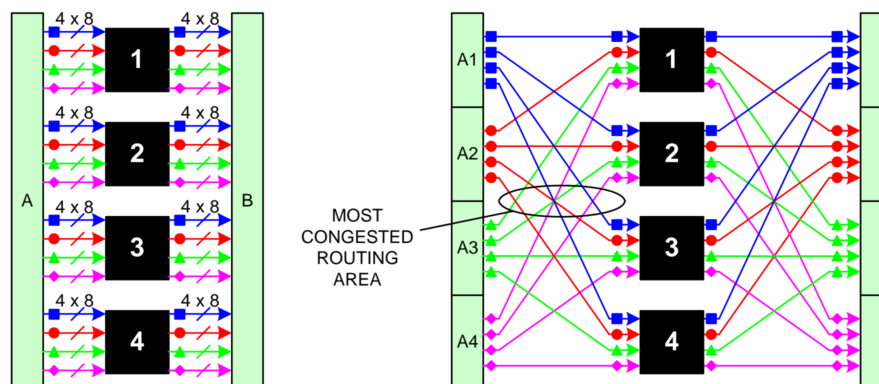

Figure 3: Logical and physical routing flow for 4 ASICs with mesh interconnection between 2 connector banks

The logical connection between connectors and ASICs on a processor module is shown on the left and the physical routing flow on the right, with the links effectively routing in groups of 8. The most congested vertical routing area is the centre of each mesh, where all 128 pairs must pass. The ASIC and connector placement allowed the available cross-sectional substrate area to be calculated, along with its theoretical routing capacity. With the baseline stack, single-sided ASIC placement is infeasible, with a -5% margin, whereas double-sided ASIC placement is probably feasible, with a +39% theoretical trace routing margin for the baseline architecture.

E. Power Distribution Network (PDN) design

PDN design is critical to the reliable operation of contemporary integrated circuits (ICs) with high current transients. The PDN must supply sufficiently clean power to the ICs to enable them to operate without error and provide a controlled, low impedance reference path for signals to maintain signal quality. A 2 V IC supply and 10% maximum voltage deviation were assumed. The IC switching current was then estimated and used to determine the bulk supply decoupling requirements. The high-frequency decoupling is designed to maintain a target PDN impedance, up to a critical frequency, F_{max}, which is dependant on the edge rates of the IC's output drivers. The NGP module PDN circuitry cannot reach F_{max} due to decoupling capacitor and IC mounting inductances, so the shortfall must be provided by on-package and on-die sources. A more reasonable but still challenging F_{max} of 500 MHz was therefore used. To achieve this, the capacitor mounting inductance must be minimized by using reverse geometry components (*e.g.* 0508, instead of 0805), using multiple large diameter vias per pad (closely spaced to reduce loop area), locating the power planes close to the pad layers (to minimize via length) and minimizing plane separation (to minimize spreading inductance). A simple spreadsheet model [13] was derived for the NGP module and shows that the baseline is feasible.

F. Module Surface Budget

The surface budget was estimated by accounting for the connectors, ASICs, ASIC PDN passive components, using typical commercial FPGA transceiver data [11] and various mechanical features. The budget is estimated as 132% (*i.e.* 32% oversubscribed) for single-sided substrate ASIC placement and 70.8% / 106.5% when using both the top and bottom substrate surfaces respectively.

III. CONCLUSIONS

The innovative thermal and mechanical design dictates substrate choice, which has poor electrical properties in comparison to organic PCB materials [9, 14]. The 100 Ohm differential trace impedance requirement and high relative permittivity result in narrow traces, thick dielectrics and large intra-pair gaps. The resulting low maximum routing density of 445 pairs per square inch, for the baseline edge-coupled symmetrical stripline trace geometry, assumes a separation factor of twice the pitch. For PTP layering, the highest achievable routing density is 108 pairs, per horizontal inch of substrate. PTTP stacks offer potentially significantly greater routing density, but only if orthogonal routing is used, which appears unlikely for the component placements and Clos topology considered. The three variants of the baseline architecture are feasible for thermal and horizontal routing capacity, assuming the architecture is partitioned across the maximum number of available modules. Analysis has shown the PDN circuitry is feasible. Double-sided ASIC placement is required for feasible vertical routing and is most likely to achieve a surface budget with positive margin. Time domain link simulation [13] has also shown that a realistic worst case channel is feasible, using commercial FPGA transceiver models [11].

IV. FURTHER WORK

Work on the design of a substrate test piece is currently (June 2006) under way. Various structures, such as ring resonators, striplines, vias, passive PDN components, dummy ASICs and representative connectors will allow materials properties characterization, the correlation of measurements with 3D EM simulation results, end-to-end link emulation and surface mount assembly trials. Calculating the signal loss per unit length for the possible trace geometries will allow this to be included in the trace geometry trade-off and calculating trace geometries, densities and stack-ups for typical commercial PCB materials will provide a useful benchmark. Cross-talk analysis will allow an accurate setting of the inter-pair trace separation factor and return current cross-talk analysis, for PTP layering, will determine if the signal layers really are isolated by the thin planes, or whether the relatively large skin depth will present a problem. Alternative Clos partitionings and the optimal substrate size, ignoring manufacturing constraints, should also be investigated.

ACKNOWLEDGEMENTS

The author would like to thank his colleagues, Rajan Bedi, Peter Dunford, Lewis Farrugia and John Phillips and academic supervisor, Dr. Izzat Darwazeh, for their assistance and stimulating discussions during this project, the LCS organizing committee for providing a forum for this paper and ESA for their collaboration and co-funding.

REFERENCES

[1] Astrium PPG: http://www.astrium.eads.net/corp/prod/ (select payload equipment -> digital processors).

[2] I4: http://about.inmarsat.com/satellites.aspx?top_level_id=3&language=EN&textonly=False.

[3] NGP: http://telecom.esa.int/telecom/www/object/index.cfm?fobjectid=18809.

[4] Charles Clos. "A study of non-blocking switching networks," the bell system technical journal, 1953, vol. 32, No. 2, pages 406 to 424.

[5] Polar instruments Si6000B 2D field solver V3.10. See http://www.polarinstruments.com .

[6] Agilent ADS Linecalc. See http://eesof.tm.agilent.com/docs/adsdoc2005A/pdf/linecalc.pdf.

[7] Ansoft HFSS. http://www.ansoft.com/products/hf/hfss.

[8] CST microwave studio. See: http://www.cst.de/Content/Products/MWS/Overview.aspx.

[9] Nelco N4000-7 EF (FR4) see http://www.parkelectro.com.

[10] P. Debye, "Polar Molecules," Chemical Catalog Company, New York, 1929, Chapter V.

[11] The Xilinx RocketIO transceiver user guide, UG035. See http://www.xilinx.com/support/library.htm

[12] Various contributors, "Proposal for development of the next generation processor," SP-DG0025606.

[13] Peter Dunford, John Phillips, Mark Walker. "NGP high-speed serial links phase 1 final report." December 2005, RP-DG0043397-ASTR.

[14] Rogers 4003, see http://www.rogerscorporation.com.

AUTHOR BIOGRAPHY

Mark's Engineering experience began in 1993, initially in professional audio equipment manufacture at Soundcraft and Citronic and more recently developing Ethernet switches at the Premises Distribution Division of 3Com, in Hemel Hempstead, (1997-2001) and designing prototype satellite processor hardware at Astrium-Space, in Stevenage, (2001 to present). He obtained his first degree in Electronic Engineering from the University of East Anglia, Norwich, in 1992, an MSc in telecommunications from UCL in 2003 and has been studying for an EngD in communications at UCL, since March 2005.