3D Integrated Packaging Approach for High Performance Processor-Memory Module

S. C. Polzer, W. L. Wilkins, J. L. Fasig, M. J. Degerstrom, B. K. Gilbert, E. S. Daniel
200 First Street SW
Rochester, MN 55905
(507) 284-5506, (507) 284-9171, polzer.stephen@mayo.edu

Abstract

As high performance computing (HPC) system performance requirements increase, it is necessary to investigate new methods for integrating system components. Of interest is the applicability of 3D packaging approaches to HPC systems. Using thermal test chips, we designed and assembled a 3D processor-memory module with an integrated power delivery network to investigate interconnect density, integration, testability, and rework issues with 3D integrated packaging in an HPC environment. The design was based on interconnection and power delivery requirements for a processor-memory module capable of supporting 64 full-duplex 30G SerDes, routing for 800 processor-to-memory pins, an integrated multi-tiered power delivery network, and a thermal management solution capable of dissipating a nominal processor heat flux of 100 W/cm2. The technologies selected—semi-rigid flex, power connectors, land grid array (LGA) attach with an anisotropic film, and cold plate-based cooling—are all commercially available technologies, which we adapted for this HPC module. As more advanced 3D packaging and integrated circuits become available, these assemblies and components can be incorporated into our approach to increase integration and performance. This design approach also accommodates substitution of thermal test chips in place of functional components, allowing validation of thermal management solutions ahead of the final module design. We will present the electrical-to-mechanical design strategy used to build this module and results of the thermal and electrical analyses, and point to several areas where further development work would be beneficial in the areas of interconnect, power delivery, and mechanical design.

Key words

3D Packaging; integrated thermal management; cold-plate cooling; high performance computing (HPC); processor-memory module

I. INTRODUCTION

As system performance requirements for high performance computing (HPC) systems continue to increase, the need to increase the packaging density to shorten interconnect distances also increases. One technique for accomplishing this is to move to 3-dimensional heterogeneous integration of system components. Of particular interest are the all-to-all processor-memory nodes common to a class of special purpose HPC systems. A customized study conducted with TechSearch International in 2009 [1] revealed several technology deficiencies with respect to the HPC application space: interconnect density, thermal management, reworkability, reliability, field repairability, and testability. A review of literature confirmed that current 3D and integration packaging technologies are being driven by handheld devices with cost being a primary concern [2]. Although there is ongoing research and development work in 3D silicon integration, such as [3], [4], the currently available 3D package solutions [5] do not completely address the deficiencies listed above. As a result of this study, we undertook an effort to design a test vehicle using currently available technologies to address these deficiencies based on the power delivery, interconnect density and thermal management requirements of a processor-memory module that is representative of what we have seen in some of the current and upcoming special purpose HPC systems. We used thermal test chips in place of functional processor and memory parts in the design to allow for future thermal validation of the design. Other design considerations included the ability to test the various component parts prior to full assembly, the use of material sets and assembly methods that allow the 3D module to be reworked, and the ability to replace modules in the field. This paper describes the development of this HPC 3D Integrated Thermal Test Vehicle (3DI TTV). It starts with a description of the final design followed by a discussion of design considerations, looking at requirements, design methods and analyses. It concludes with an assessment of where additional development in the areas of interconnect, power delivery, and mechanical design would be beneficial.

II. Overview of 3DI TTV Design

The basis for the 3DI TTV design is the use of currently available technologies. We selected semi-rigid flex as the base interconnecting substrate. Although semi-rigid flex has been used for high-density, low-cost packaging in consumer products such as cellular phones and laptops, there is increasing capacity for high density interconnect rigid flex circuits [6]. The rigid islands allow for mounting of fine-pitch BGA devices and the flexible regions allow for the assembly to be folded into compact shapes. One approach for using the semi-rigid technology in an HPC is to place a processor on a central rigid area and attach memory die on adjacent rigid islands, which can be folded above the central rigid area thereby creating a module with a small footprint. The 3DI TTV module is based on this packaging approach, using a left and right semi-rigid memory islands that flank a central island accommodating the ASIC processor.

The 3DI TTV is composed of thermal test chips (TTC), a semi-rigid flex base substrate, and an integrated power delivery network. The TTC used to emulate the processor and the memories for this design were purchased from Thermal Engineering Associates (TEA). Figure 1 shows the various configurations of the TTC. The emulated processor consists of a 5 x 5 array (12.5 mm x 12.5 mm x 0.3 mm) of TTCs that are flip-chip attached to a 27 mm x 27 mm x 2.0 mm substrate. TEA provided this to us in the form of a ball-grid array (BGA) package with lead-free (SAC alloy) bumps on a 1 mm pitch. The emulated memories consist of 4 x 4 arrays (10.5 mm x 10.5 mm x 0.3 mm) that have been thinned, backside metallized with 1000 Å Titanium, 3,000 Å Nickel, 3000 Å Silver, and bumped with lead-tin eutectic solder on a 250 µm pitch.

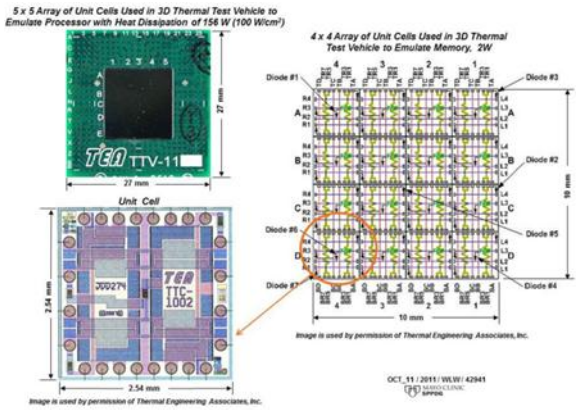

Fig 1: Thermal Engineering Associates thermal test chips and package used in 3DI TTV module

The processor and memory TTC components are designed to attach to a semi-rigid flex circuit fabricated by HEI, Inc. Figure 2 depicts an unfolded view of the semi-rigid flex. The red square is the BGA package containing thermal chips that emulate the processor. The green squares are thermal test chips that represent memory. The power regulators are mounted on separate cards, one for each memory island and one for the processor island. The flex circuit carries power and signals associated with the thermal diodes and the power cards. On the bottom of each memory island, there are pads for a connector to the power card, which contains the power regulators for the memory.

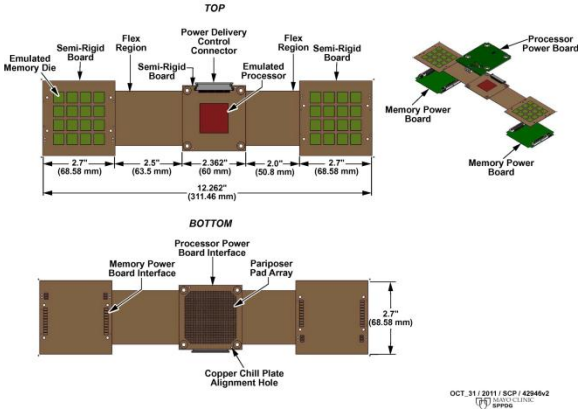

Fig 2: Unfolded view of semi-rigid flex substrate

On the top of the processor island there is a connector for plugging in the card containing the processor power components. Because the power components are on separate cards that attach to the semi-rigid flex with a connector, failed power components can easily be replaced by swapping out the power component cards. Figure 3 depicts the folded view of the semi-rigid flex assembly.

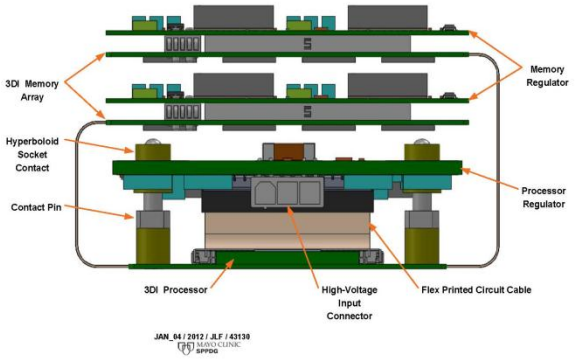

Fig 3: Folded view of semi-rigid flex assembly with copper heat spreaders removed

Copper heat spreader substrates are placed in between each rigid area of the folded assembly for mechanical support and to transfer heat into a base cold plate located at the bottom of the module. Water is pumped through a base cold plate to complete the module heat transfer process.

The 3DI TTV module is designed to connect to a printed circuit board (PCB) through an anisotropic conductive film, PariPoser™, made by Paricon. The PariPoser™ material

consists of a silicon membrane that contains columns of silver plated nickel balls. When the material is compressed, an interconnection is made in the vertical direction. This type of interconnect was selected because it supports high-speed signaling, which we were able to verify [7], and the land grid array (LGA) attach provides for relatively simple assembly and disassembly for field replacement. When used in an actual processor-memory module configuration, the high-speed inter-processor links would be routed through the PariPoser into the PCB to allow interconnection with other processor-memory modules. A detailed view of the final 3DI TTV is depicted in Figure 4, including a photograph of an assembled module.

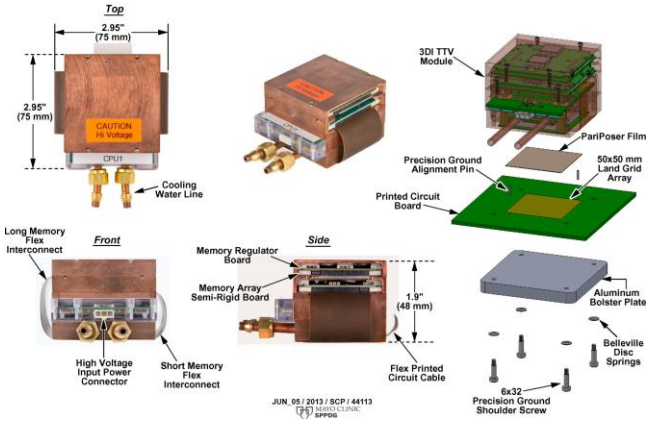

Fig 4: Detailed view of 3DI TTV module

III. DESIGN CONSIDERATIONS

The 3DI TTV design required a number of trade-offs to be considered in order to accommodate the signals, power, thermal and mechanical requirements. The process of creating this design allowed us to more fully investigate the issues associated with applying 3D integrated packaging to HPC systems. We have divided the design considerations into three areas: interconnect, integrated power delivery, and mechanical. For each area, we describe the requirements, and design and analysis methodology used.

A. Interconnect

The 3DI TTV design did not impose any special requirements on the signal interconnects beyond the number of inputs and outputs (I/O), as there are no high-speed interconnections. To assess the capability of the interconnect technology to support this application, we conducted a routing study and preliminary crosstalk analysis. To be viable for use in the specialized HPC systems we have worked with, the semi-rigid flex technology needed to maximize the memory bandwidth as well as support 64 full-duplex SerDes pins. The first step was to determine the I/O needed for memory. To support the equivalent of six DDR3 DIMMs, 800 pins are more than adequate (a typical DDR3 DIMM has 120 to 130 signal pins). For maximum bandwidth, four memory chips and one controller are required per DIMM, and the 30(=5x6) chips required per module should be easily accommodated between the two memory islands. At DDR3 2133 Mb/s speed grades, each DIMM can support up to 136.5 Gb/s of memory bandwidth yielding a module capable of 0.819 Tb/s of memory bandwidth. Based on emerging DDR4 standards, replacement of the DDR3 memory with DDR4 could push the bandwidth over 1 Tb/s.

Figure 5 shows the results of the routing study for the lower-left quadrant of the 50 x 50 pin array using a 1 mm pin pitch. The remaining quadrants are assumed to have similar pin counts and assignments. The inner-most 24 x 24 array of pins is dedicated to power delivery for the processor core logic. Thirty-two SerDes pairs are required for 16 SerDes full-duplex channels per quadrant. The SerDes pin pairs are designed with a generous void to enable the low pin capacitance required for 20 Gb/s to 30 Gb/s data rates. Moreover, each SerDes pair is isolated from other signal pins by a ground or power pin to limit crosstalk. Our primary interest in this routing study was to determine whether 207 signal pins (200 pins for memory with 7 spares) could "break out" between the voids in the planes required around the SerDes pins. This study showed that using line width and spacing of 1.5 mils and 2.0 mils, respectively, the breakout of 207 pins was possible.

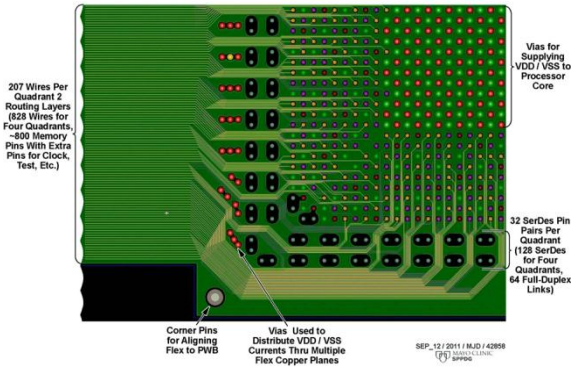

Fig 5: Semi-rigid flex study showing memory I/O routing through 64 full-duplex SERDES pins within a 50x50 1mm pitch pin-field quadrant

Although the semi-rigid technology offers tremendous wiring density per layer, to meet the minimum allowable bend radius for the flex portion, only 2 or 3 wiring layers and associated ground plane layers is considered feasible. In general, memory traces are routed using minimum spacing for very short distances – only a fraction of the crosstalk saturation length. However, the simulation results in Figure 6 show that crosstalk can be managed for very tight wiring

spacing – in this case a 0.0017" spacing was used (equal to the line width).

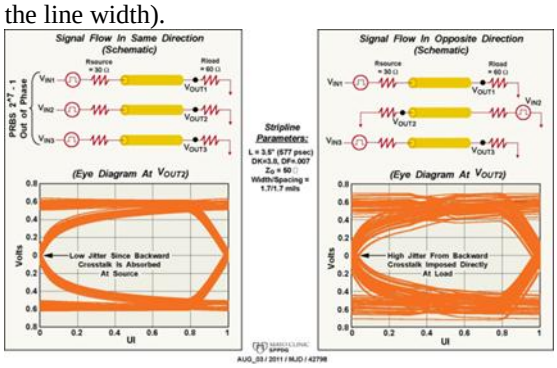

Fig 6: Simulation results for semi-rigid flex crosstalk analysis

The key to managing crosstalk with tight wiring spacing is to group all wires having the same signaling direction together. By adhering to this rule, low crosstalk can result in a very clean eye opening, as shown in the lower-left eye diagram of Figure 6. However, if signal flow is allowed in opposite directions between tightly coupled striplines, then considerable cross talk can degrade the eye opening significantly, as is shown in the lower-right eye diagram. This crosstalk mitigation approach requires good impedance matching, such as that specified for DDR3 memory, to the interconnect impedance at both source and destination points.

B. Integrated Power Delivery

The power delivery network design was primarily influenced by two factors: 1) A load dissipation representative of future HPC systems; 2) A physical form factor that was compatible with the 3DI TTV module architecture. To be what we think will be representative of future HPC systems, we chose a nominal power dissipation of 156 W total for the process load to provide to provide a power density approaching 100 W/cm2. The memory array dissipation was set at 2 W per memory package for a total allocation of 32 W. Based on these constraints, we settled on the power architecture summarized in Figure 7.

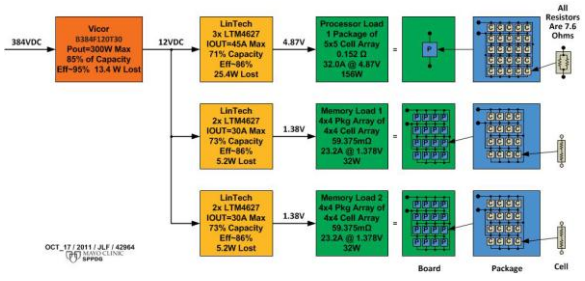

Fig 7: 3DI TTV module power delivery network architecture

Since the TTCs used in the processor and memory arrays are fixed resistive loads, we had no performance requirement regarding transient response beyond maintenance of the output voltage and thus avoided the typically substantial effort of a decoupling-network design.

In addition to meeting specific electrical requirements, the power system needed to accommodate the assembly of the 3DI semi-rigid flex and the subsequent folding of that assembly into the final shape of the finished 3DI TTV itself. In the course of exploring this design constraint, we considered several variations of high current interconnect including new products from Molex and other manufacturers intended to serve the commercial server market, but we were repeatedly frustrated by the sheer size and limited height-compliance of server components. These difficulties were compounded by the fact that the metal plate that cools the processor and the processor voltage-regulator needed to be sandwiched vertically between the processor and processor-regulator while maintaining a low-resistance electrical path as the 3DI TTV expanded and contracted due to thermal effects. Ultimately, we settled on a board-mountable pin-and-socket configuration manufactured by Hypertronics Corporation. Their hyperboloid socket design is inherently insensitive to engagement variations while offering low insertion and extraction forces and low contact resistance. The resulting 3DI TTV power regulator boards are shown in Figure 8.

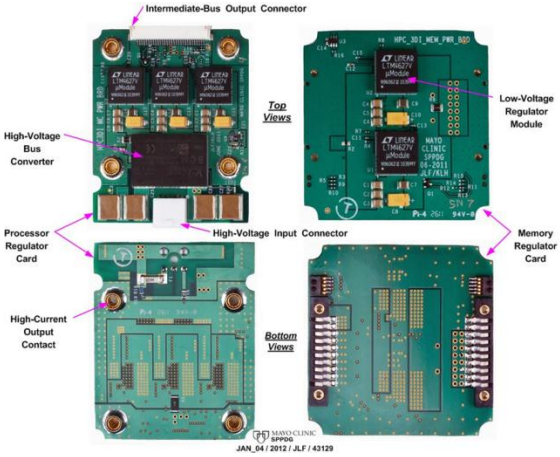

Fig 8: Pictures of 3DI TTV module power regulator boards

The high voltage is delivered to the module via a cable attached to a three-pin connector located at the front edge of the processor regulator card, and the low-voltage output is delivered to the processor by pins and hyperboloid sockets near the corners of the rigid section of the processor card and regulator, respectively. The 12 V intermediate bus is routed down to the processor by a separate length of Flexible Printed Circuit (FPC) cable through mating connectors at the rear edge of the processor and processor

regulator cards, and subsequently routed to the memory arrays and memory regulators through the integral flexible interconnects.

C. Mechanical and Thermal Design

The mechanical design of the 3DI TTV module was an iterative process based on interconnection, power delivery and thermal requirements. The design needed to accommodate the interconnect (both within the module as well as to a printed circuit board), while supporting the power delivery network and providing an adequate thermal solution within the 3DI TTV. In addition, it needed to meet the overall goals of maximizing packaging density, while allowing for assembly rework, test, and field replacement. SolidWorks computer aided design (CAD) software was used at the beginning of the 3DI TTV design cycle to create high level concepts of the module based on the TTC component form factors, which were turned into electrical designs using Cadence Allegro CAD layout software. Mechanical 3D CAD models were then generated from Allegro designs using CircuitWorks, bi-directional intermediate data format (IDF) file translation software. The minimum footprint area of the 3DI TTV was dictated by the processor power regulator card design. An example of this translation is shown in Figure 9.

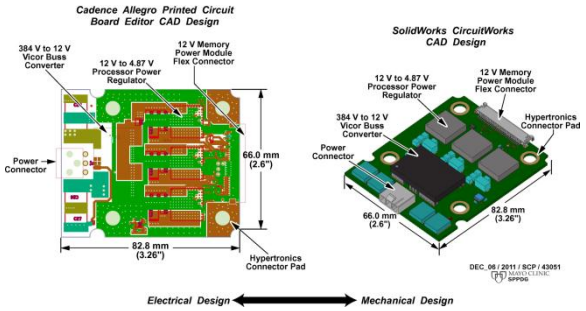

Fig 9: Example of electrical to mechanical 3D CAD translation used in 3DI TTV design

To predict the structural integrity and thermal performance of the 3DI module, we used the finite element analysis (FEA), computational flow dynamics (CFD) and tolerance stack-up analysis applications that are integrated within the SolidWorks 3D CAD environment. The PariPoser film loading and alignment requirements were analyzed with SolidWorks Simulation FEA and Tolanalyst software. From these analyses, we selected bolster plate, base chill plate geometries and materials that met the 0.05 mm (0.002") PariPoser parallelism requirement. SolidWorks CFD Flow Simulation was used to optimize the thermal performance of the water cooling channels in the base chill plate and predict the maximum operating temperatures of components within the 3DI TTV. The following parameters were used in the CFD analysis: processor power = 156 W, memory die = 2 W, Vicor Bus converter = 13.4 W, three LTM4627 voltage regulators (processor) = 8.5 W each, four LTM4627 voltage regulators (memory) = 2.6 W each, incoming water temperature = 25 ºC, incoming water flow rate = 0.1 L/minute. Based on these assumptions, the CFD analysis predicted the temperature of the components were well below the generally accepted 85 ºC maximum operating limit. Specific temperature results from the analysis are summarized in Table I.

Table I: Results of 3DI TTV computational fluid dynamics analysis

Component	Max. Temp. (ºC)
TEA Processor	50.9
TEA Memory Die (level 2)	45.7
TEA Memory Die (level 1)	37.3
Vicor Bus Converter	44.9
LTM4627 Processor Regulator	55.0
LTM4627 Memory Regulator (level 2)	45.1
LTM4627 Memory Regulator (level 1)	49.4

The HPC 3DI TTV was designed to be assembled to a PCB for testing. The final assembly involves connecting each 3DI TTV to a motherboard by aligning the PariPoser film and then compressing the film by tightening the four 6-32 precision shoulder screws using a torque wrench to apply a uniform load. Because we selected an LGA attach methodology using anisotropic conductive film as interconnect, the entire 3DI TTV module can be removed from the motherboard and either reworked or replaced.

IV. DISCUSSION

Mayo designed the 3DI TTV using thermal test chips to investigate thermal management, interconnect density, integration, testability, and reworkability issues with using 3D integrated packaging in an HPC environment. The design was based on the interconnection and power delivery requirements we have seen in special purpose HPC systems. In addition, we designed the module to allow for rework at

assembly, test, and field replacement while maximizing the packaging density. The technologies selected—semi-rigid flex, the power connectors, LGA attach with an anisotropic film, and cold plate-based cooling—are all commercially available technologies, which we adapted for an HPC module. Although our design is not something we have seen offered by industry, it has the advantage of leveraging existing fabrication infrastructure as well as having the flexibility to accommodate advances in component technology. When compared to standard 2D layout practices, our approach of integrating a power delivery network with a folded semi-rigid processor-to-memory interconnect flex design resulted in a 4X improvement in packaging density. The flexibility of this modular design should also scale in density as advances are made in 3D packaging and integrated circuits. Finally, this design approach can be used as a validation tool for thermal management designs as TTC devices can be substituted for functional components ahead of the final module design.

The 3DI TTV design does point to several areas where further development work would be beneficial in the areas of interconnect, power delivery, and mechanical design. For the interconnection, the use of advanced low-K dielectrics to improve wiring density and reduce flex circuit thickness, and the incorporation of multiple thicker metal layers to help distribute high-current, low-voltage could be investigated.

The primary area for further development is in power delivery. This design effort highlighted power delivery as a major technology limitation for using 3D integrated packaging in HPCs. Both the signals and power contend for common resources, trade-offs impact both camps. If this design were to be used with an ASIC or FPGA, multiple voltages and voltage-decoupling would need to be accommodated. The space taken by the power delivery network would likely double the module area. Alternatively, to minimize module size, miscellaneous voltages could be supplied from the motherboard, through LGA pins into the module. The development of custom power connectors and integrated power components would substantially reduce the size of the power delivery network.

A major portion of this design effort ended up being the mechanical design. The intent was not to optimize the mechanical design, but to find a solution. For example, it may be possible to redesign the power pin connectors to either reduce their size or to integrate them with the mechanical alignment and loading. We acknowledge that the use of copper for the structure and water cooling adds complexity, weight and cost. Other custom chill plate options exist [8] that may improve the weight, performance, and reduce cost when compared to the support structure designed in this study. In addition, it might be possible to incorporate the alignment mechanism with custom high current connectors, reducing the mechanical overhead.

V. CONCLUSION

In summary, we demonstrated that this type of design approach to 3D-integrated packaging has the potential for providing denser, high performance packaging for processor-memory modules in HPC systems. Although the design is not optimized, it achieved the base requirements for power, interconnect, and thermal management, as well as accommodating our objectives for assembly, test, rework, and field replacement. In addition, as part of the design effort, we were able to identify several areas for further design work and technology development that show potential for reducing the module weight and size and increasing performance.

References

[1] E. J. Vardaman, "3D integration technology analysis for large-scale, high-performance computing applications," TechSearch International, Inc. April, 2009.

[2] C. E. Uzoh, "Technology drivers for 3DIC," presented at SMTA Upper Midwest Chapter Meeting, March 28, 2012.

[3] Y. Kurita, et al, "A 3D stacked memory integrated on a logic device using SMATFTI technology," in *Proc. Electronic Components and Technology*, 2007, pp 821- 829.

[4] V. Solberg, "Stacked Multi-Chip μZ™-F Packaging for the Next Generation Electronics," *Journal of SMT*, vol. 15, issue 2, 2002, pp. 18 – 24.

[5] B. E. Swiggett, Prismark Partners, "Advanced packaging outlook: trends and implementation challenges," presented at SMTA Upper Midwest Chapter Meeting, March 28, 2012.

[6] M. Buetow., "AT&S Sees Q1 Profits Drop," Printed Circuit Design and Fab [Online]. Available: http://pcdandf.com/cms/fabnews/9199-atas-sees-q1-profits-drop. [Accessed May 2013].

[7] M. J. Degerstrom, D. J. Post, B. K. Gilbert, E. S. Daniel, "PCB pin-field considerations for 40 Gb/s SerDes channels," in *Proc. Electronic Components and Technology*, Las Vegas, NV, May 30, 2013, pp 1392-1400.

[8] Lytron Total Thermal Solutions, "Cold plate custom design & manufacturing," 2011. [Online]. Available: http://www.lytron.com/Cold-Plates/Custom. [Accessed January 2012].