

Building 3D-ICs: Tool Flow and Design Software Part 2

Robert Patti

11/21/2011 11:18 AM EST

The industry's current enthusiasm for 3D-ICs is widespread and well warranted, but designing those 3D devices presents a challenge. Normal 2D tool flows, thoroughly honed and refined over many years, nonetheless fail to address some of the critical issues of 3D design. A new 3D design process is evolving gradually from that 2D heritage. When Tezzaron designed its first 3D circuits in 2003, the designers used standard 2D CAD tools and cobbled together a 3D DRC and LVS flow based on scripts. Today there are tools to handle a complete backend flow and strides are being made to enable true 3D design partitioning, synthesis, placement, and routing (see Figure 1).

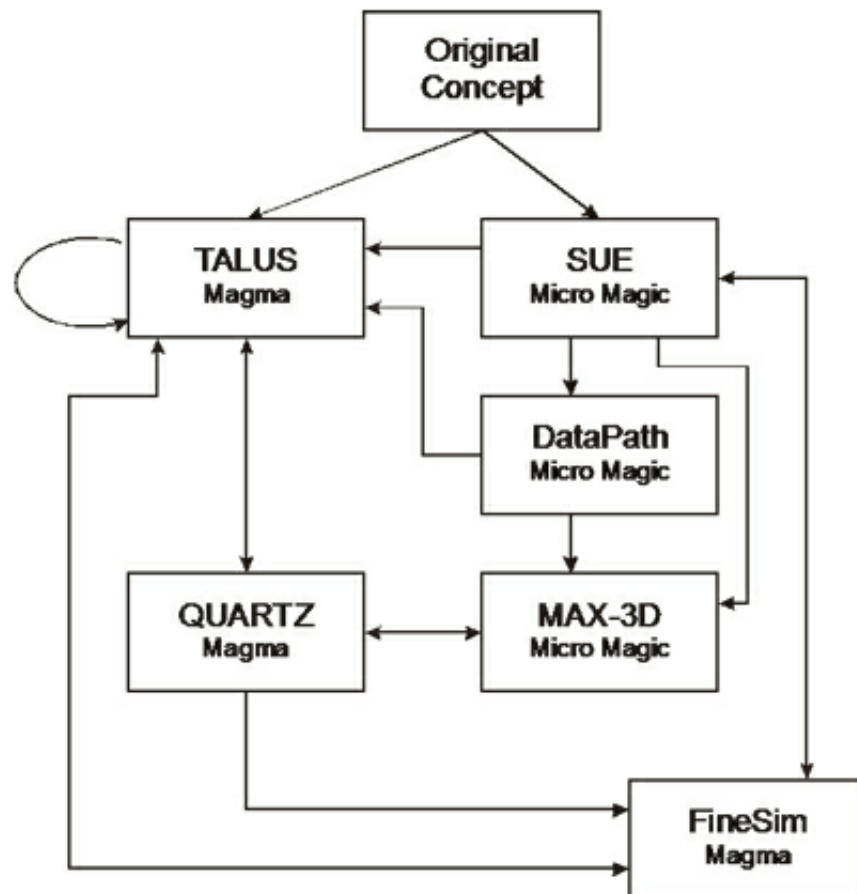


Figure 1 – Current 3D design flow

This article discusses the current state of 3D tools and software, describes a working flow, and identifies the areas where more progress is needed. We base the discussion on a specific next-generation demonstration device taken from a design that Tezzaron is prototyping with several partners. The demo design contains an

advanced ARM® processor stack, an “off the shelf” FPGA die, and a DRAM memory stack, all assembled onto an active silicon circuit board acting as an interposer, as shown in Figure 2 below.

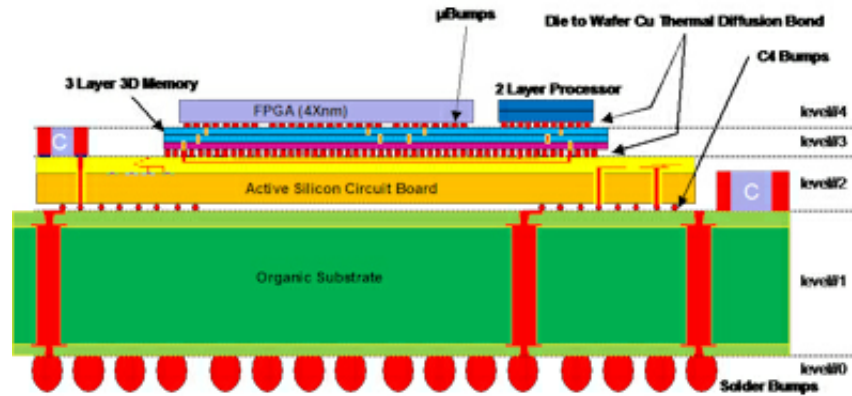


Figure 2 – The demo device

[Part 1 of this article](#) talked about the various pieces of this system and looked in depth at the process. Part 2 will look at the other pieces of the system.

The Memory

The memory design follows a similar flow to the processor, but instead of the primary design entry being Verilog, it is schematic captured. The schematic entry tool we use here is the Micro Magic SUE editor. It is closely tied to the MAX-3D physical editor and has data path compilation capabilities that are handy for the highly bit slice oriented design of a memory. The Micro Magic tools, like the Magma tools, use TCL extensively. In fact, much of Micro Magic’s tool set is written in TCL, and thus is very open to adaptation. The designer can use the DPC tool, which extends from the schematic editor, to do the placement and get early estimates of the delays and physical size. Much of memory design consists of fitting together puzzle pieces in an optimal way. The memory cell itself sets wordline and bitline pitch; these in turn dictate sense amplifier and wordline driver pitches, and from here the pitch of the data path elements is set as well. Timing generators and state machines are often complicated, and need to fit into very specific areas to satisfy time of flight constraints and circuit delay matching. It can be difficult to express physical constraints to a layout engineer, so the designer is deprived of the layout engineer’s insight into the issues and knowledge of alternate implementations that might work better. In our experience, all too often a block is designed in such a way that the layout engineer must struggle to make things fit. Good interaction between designer and the layout engineer is absolutely essential, otherwise the design ends up much less than optimal. The tight coupling of the SUE editor to the MAX-3D physical editor allows information to pass effortlessly between the designer and the layout engineer. This saves a significant amount of time.

About 3D Physical Editing

In 3D design, physical editing is a very big deal. The MAX-3D physical editor is a true 3D editor that allows differing technology files to coexist in the editing environment. We can edit 2D layers or cut and paste across different layers in 3D. The power of this tool becomes obvious when doing real 3D design editing. With 2D tools, simply moving a 3D connection requires separate editing in different tool sessions and data mapping and transfers and external DRC/LVS checking. The MAX-3D tool loads all layers at once and incorporates basic DRC and LVS functions; moreover, the 3D stackup itself is displayed in a direct manner.

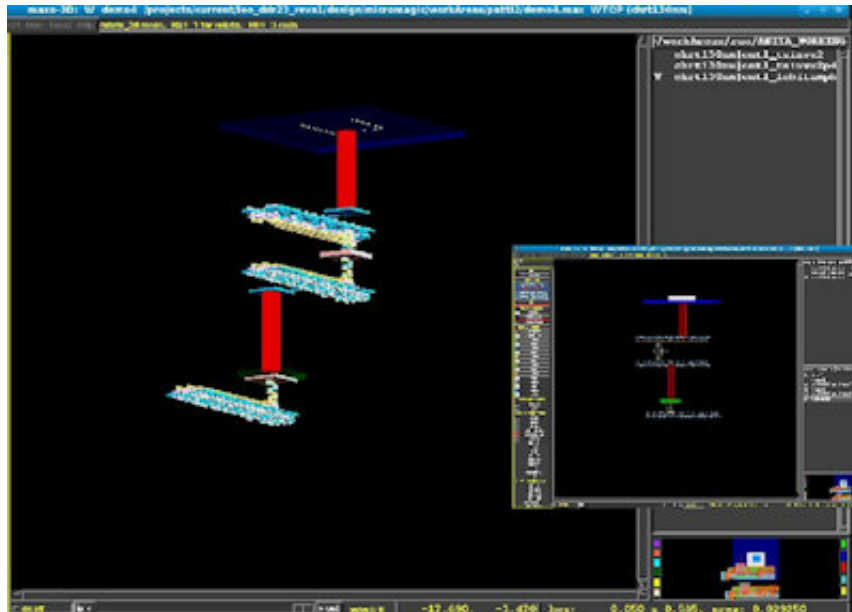


Figure 3 – MAX-3D orthogonal view screen shot

Another vital element in 3D design is layer directions. Sometimes connections are face to face, other times face to back, or even back to back. In a stack with several different layers there can be dozens if not hundreds of possibilities. Using 2D tools makes 3D designs difficult and scary; good 3D tools keep track of the layers and understand the 3D connections.

Process Separation

Getting back to our 3D memory device, we need to design and lay out two different types of layers for our 3D DRAM. Tezzaron 3D memories use two different types of wafers, constructed with two different process technologies and actually built by two different manufacturers. This process separation enables extraordinary levels of optimization, yielding DRAM with low power, high density, and near SRAM like performance. At the same scale of manufacturing it also allows lower cost, due to higher array utilization and better yields, but that discussion is for another article. The bottom layer of the memory stack is the controller, a logic process technology wafer. This contains the sense amplifiers, I/O interface circuits, test and repair circuits – basically, everything in the memory other than the bit cells and the sub-wordline drivers. Using the logic process to build this layer produces smaller, more compact layouts, much higher performance, and perhaps lower operating voltage.

The memory cell layer(s) must be built in a DRAM process to achieve high bitcell density and data retention. DRAM processes typically have large peripheral design rules, high device thresholds, and thick oxide devices. DRAM bitcell selects require internally boosted voltages. Also, virtually all DRAM devices run with negative substrates and wordline bias to further reduce the bitcell select transistor leakage. For practical reasons, we actually split the voltage generators to use both the controller layer and the cell layers. The logic process provides better analog function and faster feedback for the charge pump. The DRAM process is designed for generation and operation at high voltages; therefore, the high voltage circuit portions of the charge pump power supplies are better placed on the cell layers.

Design Capture

The memory device is designed mostly by schematic capture. This is due to the very rigid layout restrictions imposed by memory designs. Memories are true bottom up designs, where the memory cell is the basic building block. Days and weeks are spent optimizing the layout of the memory cell to minimize its size and maximize its yield and function. After the memory cell is fixed, the drivers and sensing circuits are fitted to the pitch dictated by the memory cell. Higher levels of circuitry are fitted around the memory arrays, and so on. Each higher level in the design has less impact on the overall size. This bottom up approach can be managed only by schematic driven design.

Micro Magic's SUE editor is much like other schematic editors except that it provides the capability to pass information to the MAX-3D tool. This is critically important in order to reduce layout iterations. SUE also has the same TCL interface, which again comes in handy for add-ons not envisioned by the tool designers. At Tezzaron we have added some functions to deal with extremely large bus widths, as much as 1 million bits, which memory devices can have internally. Intelligent auto-labeling becomes a necessity. Another obscure requirement of the schematic is generating SPICE netlists for simulation and LVS, which have millions of pins in a subcircuit. The Tezzaron 3D DRAM device has more than 2 million interconnects from the logic/controller layer to the cell layer. Not all schematic editors can write netlists with this many pins, but SUE can.

Some areas of circuitry have fewer restrictions in absolute timing and physical layout. These can be created with Verilog or the Micro Magic Data Path Compiler (DPC). The Verilog flow follows the standard Magma Talus 2D flow for synthesis, placement, and routing of blocks. The TCL interface allows control to be imposed over size, shape, timing, routing layers, directions, etc., so that the block can be easily meshed with the substantial hand layouts that the memory contains. Very few commercial EDA tools are as open and flexible as Talus.

The DPC flow is an extension of the SUE schematic editor. Within SUE the schematic is entered as 1x logic gates from a standard cell library. There is a built-in timing analyzer that, much like a Verilog synthesis tool, adjusts driver sizes based on wire lengths and gate loads. The gates are fitted and placed using the schematic placement as a template. Toggling to the DPC screen displays the actual placement and rats nest wiring. This allows the designer to iterate, improving the design as needed based on what he actually sees as the timing and size and shape. DPC works hierarchically so that the designer can choose to optimize various levels individually or to address them globally. DPC also accepts non standard cell blocks in its placement. Custom analog or pre-routed blocks can be added to the DPC library, allowing these objects to be incorporated as part of the plan rather than somehow planned around. The final DPC product is a .DEF file that is taken to the Talus router where the Magma tool can do straight routing or additional optimizations, or use the .DEF information as part of a larger place and route task.

Simulation

A memory device is an enormous analog circuit that must be SPICE verified. A unique aspect of the Tezzaron DRAM devices is the process separation that is employed. Simulation of a composite device can require pulling devices from numerous libraries. Care must be taken to have unique device names across the various libraries. The simulator should also have a large device capacity so that it can verify sufficiently large sections of a 3D design. A complete DRAM has billions of devices. Simulation, even of small sections, must handle millions or tens of millions of transistors. The Magma FineSim Pro simulator is an ideal tool for the task of simulating mega transistor count devices. Memories and 3D designs in general require transistor level SPICE simulations. SPICE data is the least common denominator when the designs are complete. If voltage variation or the effects

of temperature rise in the 3D assembly are to be modeled, only the SPICE environment offers the fine grain and accuracy needed. FineSim Pro is one of only a few simulator products that can distribute the simulation across numerous processors while keeping the simulation memory footprint under control.

Figure 4 below shows simulation of a block with more than 20 million transistors. The results of a few hundred memory clock cycles were generated in 20 minutes. Quick and accurate results are key to debugging and ensuring process corner operation of large transistor level designs like memories.



Figure 4 – Magma FineSim Pro screen shot showing results of memory block simulation.

The FPGA

This is an off the shelf device in die form. It is treated as a black box to the overall design and handled just as a packaged part might be handled for the PCB design. Die level integration has other challenges, such as getting KGD and the required physical interface data, but employing this information in a low power 3D design has few new challenges over a stand MCM. Thermal issues and power distribution would indeed ratchet up the issues and risks if the FPGA were to dissipate tens of watts, but this is not the case in the configuration of the demo device.

The Silicon Circuit Board.

In this device, the SiCB is more than just a passive interposer; it is an active element of the design. Many companies see that the future potential for the interposer or SiCB lies in extending its usage to adding value, like power regulators. If the interposer serves merely as a fanout element, the market is small and will remain small. The manufacturers of organic substrates continue to improve their offerings and will come at least close to what glass or silicon can do today; but if there are to be transistors, built-in decoupling capacitors, termination resistors, interface buffers, and perhaps optical devices, silicon is the only solution. Adding features such as these also provides real value to the silicon employed in the interposer and eliminates the need for other integrated circuit devices.

To optimize the power of our multicore CPU we operate each core and cache at an optimized voltage and frequency. Under a very light workload perhaps only a single core may be operating, and at just a few tens of

megahertz. This active CPU can operate at just above threshold with the remaining CPUs stop clocked at an even lower voltage, just enough to hold a static state. Under heavy loads all the cores may be operating, but due to process variations, an optimal voltage for one CPU is likely not optimal for another. To maximize power efficiency we use a near point of load power regulation technique. The proximity and the number of power supplies, perhaps as many as 20 or more in this example, virtually dictate the use of a local multichannel smart power device. A silicon circuit board (SiCB) is ideal for power regulation circuitry. The transistor process technology employed is large geometry and thus can be very high yielding, which is important for large pieces of silicon. The large geometry technology can also accept higher voltage inputs for switching regulators, thus reducing the current flow to the assembly and simplifying the printed circuit boards and other pieces of the overall system design.

The SiCB also contains bulk capacitors built with a high k material and trenching that provides several tens of micro-farads of decoupling for the attached ICs. The switching power supplies use attached discrete capacitors and inductors to optimize the efficiency and form factor.

Lastly, the SiCB includes polysilicon resistors that terminate the high speed clock lines from the system. A dynamic impedance matching circuit tunes the termination network for maximum clock swing without overshoot or undershoot.

Design Capture

The design parameters and rules for the SiCB put the tool requirements firmly between those used for a printed circuit board and those for a chip. The need for transistors used by our local power regulators makes the decision clear: chip tools are the correct path. We capture the design using the Micro Magic SUE editor as we are also targeting MAX-3D for pad placements and routing of the critical nets. We can tag net information in SUE so that the layout engineers know whether nets are signals or power, analog or digital, etc. Just as was illustrated before with the memory design, this reduces the iterations as well as avoiding surprises in design reviews.

Physical Design

Again we employ MAX-3D as the physical editor. This tool has the necessary hooks to support 3D elements such as TSVs and backside metal. It also has an indispensable feature for looking at high speed signals and signal integrity. MAX-3D permits a 3D cut extracting the GDS of a circuit trace. It can also use a definable halo and grab neighboring structures for signal integrity analysis. The entire 3D silicon stack (levels 2, 3, and 4, as shown in Figure 2) can be manipulated, analyzed, and edited simultaneously within MAX-3D.

Verification and Simulation

The verification and simulation flows mirror those already described above. The SiCB has its own DRC and LVS decks as well as transistor, resistor, and capacitor models. It really is just another piece of silicon circuitry. There are some important new rules regarding maximum wire lengths and other items that address mechanical stress management, but these are ultimately just additional design rules. Magma's Quartz DRC and LVS are the workhorse tools for all forms of 3D verification. FineSim Pro, because of its enormous device capacity and superior performance, is again the tool of choice for simulation.

Final Top Level Design

The Schematic

To perform the schematic capture of the top level 3D assembly, we again use Micro Magic SUE because of its message passing capabilities to the MAX-3D editor. There are no other strong requirements that make SUE a better choice than any other schematic entry tool. But because other elements have been entered with SUE, a netlist can be generated down to the transistor level for the much of the design. The netlist can be attached to the Talus synthesized layers and a vendor supplied Verilog model for the FPGA.

The Physical Design

MAX-3D is the tool used here for the final 3D design. The additional backside layers that are implemented on the memory can be added at this top level if they weren't added at the memory device level. At the very least, the complete design can be viewed and first level verified using built-in DRC and LVS capabilities. The 3D visualization capabilities give a real view of the entire 3D device assembly, including the SiCB. It should be noted that the database for the entire assembly is enormous, easily approaching 100Gbytes. Very few physical editors can manage databases of this size. Again, MAX-3D can sliver off signals of interest for in-depth signal integrity analysis and generate all of the required GDS files and views to perform the final signoff DRC and LVS confirmations.

Verification

The final 3D assembly must adhere to a global set of design rules and obviously must be checked for manufacturability, using technology files provided by the 3D assembly house. The most realistic way to check the design is to penetrate the individual devices only to the physical and electrical interface. This is basically the same as a printed circuit board approach. If necessary, the verification can drill all the way down to the transistor, but there is no obvious benefit to this.

With the appropriate model statements used in SUE, our 3D netlist provides the required verification path to confirm the 3D interconnect path. Next, MAX-3D generates the required GDS files for each silicon layer as a black box containing only the top and/or bottom metals as necessary. MAX-3D also generates the GDS with extra data to check the 3D positional information, e.g., do the backside memory pads physically line up with the frontside CPU pads? File generation is guided by a tech file provided by the 3D assembly house. Today's 3D tech files are generated as one-off hand assembled files, but in the future we expect to see additional tools to help gather the correct pieces into design control files. The 3D tech file embodies the various capabilities of a 3D assembly house. Die to die, die to wafer, wafer to wafer, flipchip, and wire bonding techniques all can come into play in these new 3D systems. There is no one single standard format today that encompasses this information.

With the GDS files generated and the netlist in hand, Quartz is once again called to the task of 3D verification. Much as it did with the 3D chips themselves, it checks the new complete 3D silicon assembly piece by piece against a schematic black box, ensuring that all pins are accounted for and that the footprint locations are correct. The entire 3D assembly is then checked against the top level schematic via the generated SUE netlist. Quartz is then used to DRC the physical 3D connections in the stack. Quartz can also perform ERC, DFM, signal integrity checks, and IR drop analysis. Other standalone tools can be employed using either the GDS

information or the Quartz extracted data as part of the final verifications.

Another 3D Tool

R3Logic's R3Integrator is a new tool to the Tezzaron 3D environment and it shows a lot of promise. It brings in the additional features of pathfinding and up-front design partitioning. The tool helps to evaluate the tradeoffs and aid in placement of TSVs. It is already set up to cover the organic substrates and packaging reflected in the level 0 and 1 of Figure 2. R3Integrator can accept and manipulate data from both the IC world of GDS formats and microns and the PCB world of Gerbers and mils. Like the Micro Magic physical tools, R3Integrator is a tool designed from the ground up for 3D design and analysis. It understands the nature and use of TSVs, e.g., it knows that TSVs through transistors is a bad thing. We expect to add this tool to the baseline flow that Tezzaron already supports.

Missing Pieces

Any discussion of 3D-IC design tools must address the missing pieces. The 3D flow described here absolutely works. Tezzaron has used it to produce many of its own devices as well as dozens of proprietary designs for its customers. Our PDK is in its eighth generation, based on the GlobalFoundries 130nm process that supports our flow as well as numerous other tools and flows. But there are still some gaps in the flow, and their impact can be large or small depending on the application.

Thermal Modeling

A common concern in 3D designs is thermal modeling. We have reviewed a few tools, but none of them quite fill the bill. The ultimate thermal modeling tool must understand the likely power usage of a design and the structure's ability to conduct heat, all while modeling millions if not billions of devices in 3D. Efforts are under way to create this tool, but so far the issue is still addressed only by 2D approximations and very careful design.

To date, Tezzaron's devices have not required elaborate thermal modeling. Tezzaron employs very thin layers, about 12 microns each. This reduces the overall thermal issue in our devices enough that the limiting factor is typically the thermal interface material between the die stack and the heat sink. In the demo design presented here the processor and FPGA are both at the top of the 3D stack, allowing direct contact to a heatsink. This is the case with most of Tezzaron's 3D designs; high power devices are placed on top and their signals are routed vertically through other layers, typically memory. Tezzaron's memories have hundreds of thousands of extra TSVs to accommodate routing the power and signals from the substrate to the high powered logic stacked above it. Putting the memory on top would be unrealistic, as it would present a thermal barrier between other device layers and their heatsink.

TSV Stress

Another 3D concern is the effect of TSV stress on local transistors. The root cause of this stress is the mismatch in thermal coefficient of expansion (TCE) between silicon and the TSVs (see Figure 5).




Figure 5 – Copper TSVs after thermal cycling (early results); TCE mismatch caused oxide failure.

Copper TSVs require large keepout zones to avoid breaking the SPICE models or at least altering inherit matching. A few tool companies are working on modeling and bringing stress effects into the simulation environment. If the TSV keepouts are to be minimized, the effects must be accurately modeled rather than designed around.

Tezzaron avoids much of the stress problem by using tungsten TSV material in its current products and supported processes. The TCE mismatch between silicon and tungsten is a fraction of that between silicon and copper. Tungsten alleviates almost all of the TSV induced stress issues, but there are limits on its application. The deepest tungsten TSVs are perhaps 20 μ m, no more, restricted by a firm process limitation.

3D Synthesis

The ultimate 3D synthesizer is not even on the horizon. This tool must take a ball of RTL code and make speed, power, and cost tradeoffs to partition the design into the most cost effective, lowest power, and highest performance device. This is many years away. Magma's current tool architecture provides some of the features and can simultaneously handle multiple technology files and design rule sets. Tezzaron has done some very early work making this flow operate through the existing TCL interface, but it only scratches the surface. It also requires a pathfinding tool such as the R3Integrator to do high-level partitioning first. Complete partitioning is still very much a manual job. For at least the next few years, 3D partitioning is going to be hard work.

3D Test

Another whole article can be written on the requirements and issues related to 3D test. A few new standards are appearing, such as 3D enhanced IEEE 1500. There is little magical about testing 3D structures, but a lot of fundamental work must be done in order to extend the historic methodologies on test. In addition, there are likely to be new methodologies and standards to address device repair and redundancy.

Summary

This article has walked through a very elaborate 3D design, touching on the tools involved in turning a concept into reality. While a 3D device this complicated will not be in your phone or tablet next year, it is destined to be there within the next several years. This past year has seen dozens of different 3D integrated circuit designs and thousands of devices delivered to customers for first looks at the ultimate in integration. Passive interposers are already shipping today, albeit in small volumes, and true 3D integrated circuit designs will ship in volume next year. A tool chain for 3D integration is available today, as illustrated in Figure 1, but it is not as complete or

mainstream as most designers would desire. Nonetheless, the number of 3D designs is starting to exponentiate, and the available 3D tools are maturing and moving forward. Existing 2D tools will rapidly adjust to 3D, adopting the 3D paradigm as a normal flow element. The line between packaging and integrated circuit will virtually disappear as 3D technologies drive toward common goals of smaller size, lower power, lower cost and higher performance. We expect that the next few years will see 3D design tools expanded, honed, and adopted into the mainstream of IC design.



Robert Patti, CTO and VP of Design Engineering; Director, Tezzaron Semiconductor (Singapore) Pte Ltd Previous experience: Founder and President of the predecessor company, ASIC Designs, Inc.; Member of Technical Staff for Tellabs, Inc. Member of IEEE. Former Vice-Chairman of JEDEC's DDRIII / Future Memories Task Group. Holder of 16 US patents, numerous foreign patents and many more pending patent applications in deep sub-micron semiconductor chip technologies; BS (EE/CS) and BS (Physics), Rose-Hulman Institute of Technology.

If you found this article to be of interest, visit [EDA Designline](#) where you will find the latest and greatest design, technology, product, and news articles with regard to all aspects of Electronic Design Automation (EDA).

Also, you can obtain a highlights update delivered directly to your inbox by signing up for the EDA Designline weekly newsletter – just [Click Here](#) to request this newsletter using the Manage Newsletters tab (if you aren't already a member you'll be asked to register, but it's free and painless so don't let that stop you [grin]). This is part 1 of a two part article. The second part will appear next Monday 21st November

The industry's current enthusiasm for 3D-ICs is widespread and well warranted, but designing those 3D devices presents a challenge. Normal 2D tool flows, thoroughly honed and refined over many years, nonetheless fail to address some of the critical issues of 3D design. A new 3D design process is evolving gradually from that 2D heritage. When Tezzaron designed its first 3D circuits in 2003, the designers used standard 2D CAD tools and cobbled together a 3D DRC and LVS flow based on scripts. Today there are tools to handle a complete backend flow and strides are being made to enable true 3D design partitioning, synthesis, placement, and routing (see Figure 1).

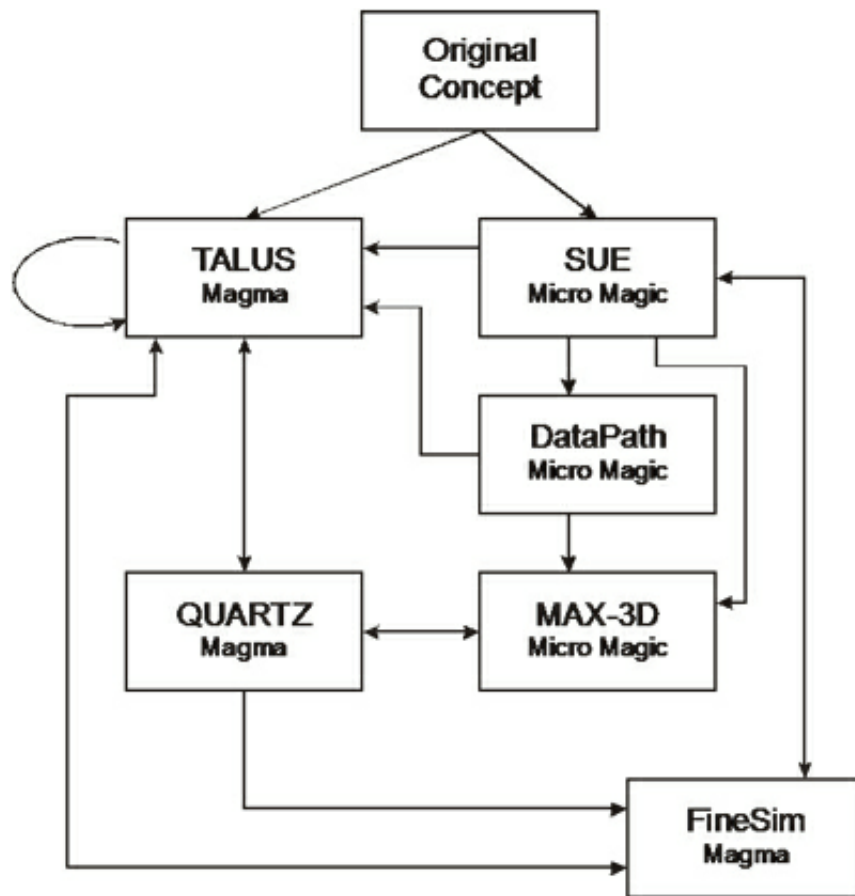


Figure 1 – Current 3D design flow

This article discusses the current state of 3D tools and software, describes a working flow, and identifies the areas where more progress is needed. We base the discussion on a specific next-generation demonstration device taken from a design that Tezzaron is prototyping with several partners. The demo design contains an advanced ARM® processor stack, an “off the shelf” FPGA die, and a DRAM memory stack, all assembled onto an active silicon circuit board acting as an interposer, as shown in Figure 2 below.

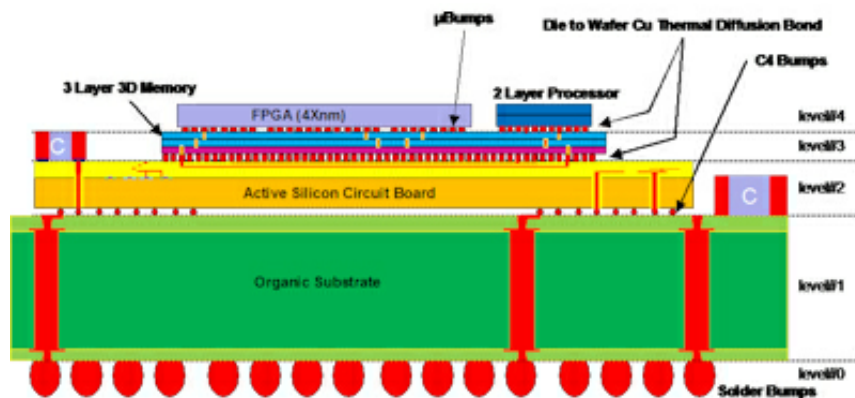


Figure 2 – The demo device

The Parts

The CPU is an implementation of an ARM® Cortex™-A processor, part of a new class of highly scalable ultra-efficient processors. An out-of-order execution unit with superscalar pipelining is tightly coupled to a low-latency level-2 cache. We take advantage of the multicore capability to support 1 to 4 layers of processing for a

scalable design implementation. Each layer of the 3D processor stack contains two CPU cores, their respective L1 caches, and a shared 2MByte L2 cache. The processor layers communicate vertically via a 1024 bit ultra wide low latency bus interface block. The interface block isolates powered-down cores and copes with bus masters of differing frequencies. The gate count per core is about 2 million gates, excluding the cache memories. The FPGA is a standard 2D device. This could just as well have been an ASIC or other logic element. It provides the glue in our demo system.

The DRAM, a three layer 8Gbit memory, is representative of a next generation Tezzaron memory. It acts as the main memory in this system.

The final element in our demonstration system is an active interposer or silicon circuit board (SiCB). Beyond signal redistribution, our demo system takes advantage of the SiCB to provide power regulation, decoupling, and signal termination.

The Processor

The 3D processor is assembled using a wafer to wafer process to produce a frontside to frontside bonded wafer pair. Wafer to wafer bonding allows the smallest, densest vertical interconnect structures and opens many avenues for design optimization.




Figure 3 – Details of a two layer logic device




Figure 4 – Three layers of silicon, top two layers thinned to 5.5 μm

On the down side, wafer to wafer bonding does not provide the yield benefit that die to die techniques can offer with a Known Good Die (KGD) approach. Overall, a wafer to wafer 3D design has about the same yield as a 2D equivalent when looking at the total area of all layers.

We scale the performance of the processing element in this demo device by adding a pair of CPU cores in each layer of the 3D stack. This reduces our partitioning effort and allows us to treat the design effort more like a 2D design with extensions. Our early 3D designs were done this way because of the primitive state of 3D tools eight years ago. In this case, the 3D enhancement of the 2D design can be relatively brute force; TSVs and vertical interconnect planning can be inserted manually. The synthesis and placement are done in 2D in this case. However, to provide insight into a more sophisticated and powerful flow, let us explore some alternatives.

Partitioning and Placement

The CPU begins as a Verilog module. It can be partitioned in a few different ways. For instance, the CPU logic could be put on one level and the cache on a separate level. In this case the manufacturing process for the cache could be done in a different node or technology than that of the CPU core itself. Process separation is one of the most beneficial aspects of 3D integration. Alternatively, we could use recent tool advances to create a 3D optimized placement of the CPU internals themselves. Using Magma's Talus tool, with its unified data model and highly flexible TCL interface, we can easily optimize the Verilog into partitioned regions. While the tool itself is not presently 3D aware, we can use it to create identical regions in a flat 2D model and define interconnect elements that mirror between regions. The mirroring in fact forces the 3D interconnect to align vertically. Another approach is to use the Talus tool to do full block physical synthesis and early partitioning and then import the .DEF files to Micro Magic's 3D-FloorPlanner (see Figure 5 and Figure 6 below).

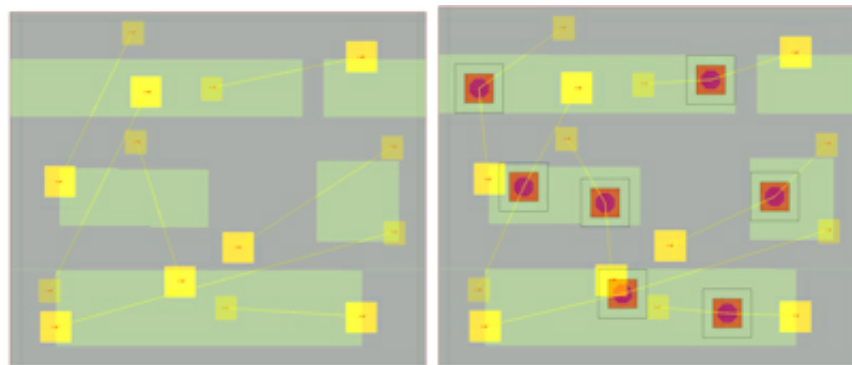


Figure 5 – TSV insertion example using Micro Magic 3D MAX auto placement

Left: Top down view of two layers to be vertically connected

Right: TSVs inserted into areas without obstructing transistors, placed to minimize wire length




Figure 6 – Orthogonal view of 2 layer circuit after TSV insertion

This tool auto inserts TSVs into the floor plan, optimizes the TSV placement, and further optimizes the cell floor plan. One of the benefits of the Micro Magic environment is that its MAX-3D tool allows true 3D physical editing. The design can be tweaked as necessary to improve the partitioning and pre-route critical nets. After the Micro Magic adjustments, the database can once again be transferred to the Magma Talus environment for routing. The advantages of partitioning the entire design in 3D is increased performance and lower power. In theory, 3D can reduce the length of wires by a factor of N where N is the number of layers used. This is overly optimistic, but as long as the 3D interconnect structures (face to face or TSV) are reasonably fine grained, the wire shortening is significant. The longest nets tend to be clocks, data busses, and global controls – all of which have the most repeaters and consume the most power. Applying 3D wiring to these nets can easily improve the system speed and power by 30 to 50%.

Someday, the ultimate 3D tool will take into account both separate processes and true 3D placement optimization. Today, however, this tool is somewhere over the horizon.

Routing

We have now successfully partitioned our processor and done the initial placement. The vertical interconnects in the placed design are modeled as standard cells or blocks and these are fixed to prevent the router from breaking the inherent 3D wiring that the placement created. Additional cells and blocks can either float or be fixed, as desired. We finish the routing with the Magma Talus router. The Talus router now sees just a simple series of fixed regions that have “wormhole” connections – that is, the 3D interconnect points provide connectivity between the fixed regions with no obvious wire length. The 3D interconnect is modeled as a lumped fixed element in the wiring network. As mentioned previously, the vertical interconnect locations are locked down, but the Talus tool is free to make final adjustments to the clustering and balancing of the cells and blocks as necessary and complete the routing process. A key to making this flow work is the extensive flexibility that Talus offers. We have glossed over the obvious strengths of the Talus tool at handling a multi-gigahertz, high performance core processor through physical synthesis, placement, and routing. At 28nm these tasks are no small issue. The need for a clock guard ring and ultra tight signal matching is paramount. Needless to say, if these basic requirements of state of the art tools were not already present in the Talus tool, the 3D extensions and use would be substantially diminished.

Verification

The design verification involves design rule checking (DRC) and comparing the layout versus schematic (LVS). Magma's Quartz, which already supports 3D DRC and LVS decks, can check the 2D and 3D design requirements. The key to 3D verification tools is that they must support the simultaneous use of multiple databases and technology files. Quartz, like Talus, has a built in multi-database capability that fulfills this basic requirement. To simplify and allow reuse of the existing 2D rule decks provided by the foundries, the design is treated as a set of 2D designs, each of which is separately checked and verified. The whole is then checked as a 3D assembly of the 2D designs. The 3D rule sets should be provided by the 3D assembly house – in this case, TezzaronDRC

As mentioned previously, 3D DRC is performed by first checking each of the 2D layers with the DRC deck supplied by its respective wafer foundry. The use of unadulterated foundry decks is important. If we were to

significantly customize a deck for use with 3D, the continuing maintenance and revisions from the foundries would make updating that deck a never-ending chore. Error results are thus generated separately for each of the silicon layers in the 3D stack.

After the 2D DRC processing is complete, 3D DRC is run. The TCL interface to Quartz allows simple scripts, supplied by the 3D assembly house, to check all of the 2D layers and then process the 3D rules as well. 3D DRC must look at the top and bottom of each of the 2D layers and the interconnect between the layers, such as bonding or backside metals. If necessary, 3D DRC also performs checks for layer specific pieces of the 3D interconnect – for example, TSVs might not be part of the 2D rules from the foundry. An important nuance is that there are multiple sets of silicon layers that may or may not have identical GDS layers. Micro Magic's MAX-3D tool can specify unique GDS numbering for the 3D layer stackup as part of its 3D tech file. In this way metall on the first layer of silicon can be unique from metall on the second and third layers, etc. The 3D GDS file written from MAX-3D contains all of the layers that are required to check the design to assure that, mechanically, the design can be correctly fabricated. The 3D tech file and the 3D DRC files should be written and supported by the 3D assembly house.

LVS

Talus provides a post place and route Verilog netlist. This netlist can be used for simulation and physical verifications such as LVS. Quartz LVS operates in 3D much as Quartz DRC does. A script interface checks each of the 2D designs against a set of 2D netlists and then an overall 3D check is done, comparing a top level netlist against the extracted 3D interconnect. Again, the original LVS decks as supplied by the foundries can be used for the 2D checks, eliminating the need for a cumbersome and error prone rewrite of the extraction and comparison rules. The 3D LVS is very simple in that there are no devices – just blocks with pins and wiring between those blocks. This greatly simplifies the 3D assembly house's task of creating 3D LVS decks. The 3D LVS can be generalized to easily accommodate 2D layers with differing technology.

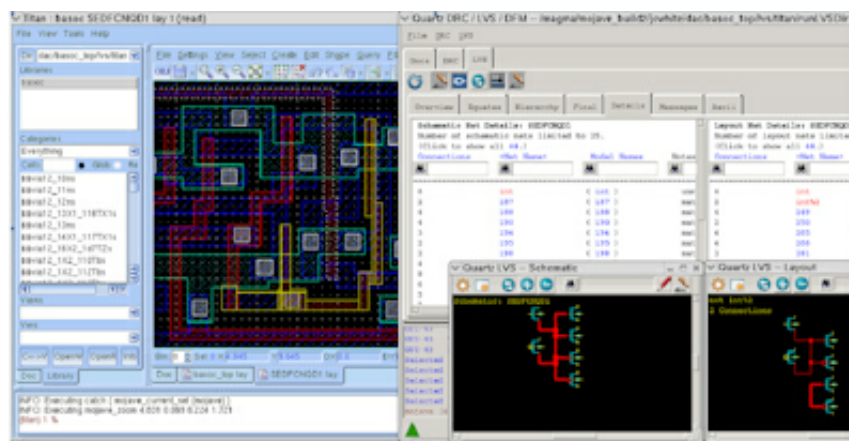


Figure 7 – Quartz LVS screen shot

The second part of this article will appear next Monday 21st November and cover the other pieces of the design.



Robert Patti, CTO and VP of Design Engineering; Director, Tezzaron Semiconductor (Singapore) Pte Ltd Previous experience: Founder and President of the predecessor company, ASIC Designs, Inc.; Member of Technical Staff for Tellabs, Inc. Member of IEEE. Former Vice-Chairman of JEDEC's DDRIII / Future Memories Task Group. Holder of 16 US patents, numerous foreign patents and many more pending patent applications in deep sub-micron semiconductor chip technologies; BS (EE/CS) and BS (Physics), Rose-Hulman Institute of Technology.

If you found this article to be of interest, visit [EDA Designline](#) where you will find the latest and greatest design, technology, product, and news articles with regard to all aspects of Electronic Design Automation (EDA).

Also, you can obtain a highlights update delivered directly to your inbox by signing up for the EDA Designline weekly newsletter – just [Click Here](#) to request this newsletter using the Manage Newsletters tab (if you aren't already a member you'll be asked to register, but it's free and painless so don't let that stop you [grin]).