IMPLEMENTING FAN-OUT WAFER-LEVEL PACKAGING (FOWLP) WITH THE MENTOR HDAP FLOW

JOHN FERGUSON AND KEITH FELTON, MENTOR, A SIEMENS BUSINESS

WHITEPAPER

Mentor®
A Siemens Business

DESIGN TO SILICON

www.mentor.com

Fan-out wafer-level packaging (FOWLP) is an emerging type of high-density advanced packaging (HDAP) technology in the semiconductor industry that is rapidly gaining popularity in the market. But what exactly is FOWLP? Why do we need it? How do we take advantage of it? What limitations still need to be overcome? And… how can Mentor Graphics help?

FOWLP BACKGROUND

FOWLP brings single and multi-die designs together, combining multiple die from heterogeneous processes into a compact package (Figure 1). The die are embedded in an epoxy mold compound (EMC), then die-to-die and die-to-ball-grid-array (BGA) connectivity are established directly through the package redistribution layer (RDL), without the need for through-silicon vias (TSVs) or a laminate package substrate. While FOWLP typically requires a silicon wafer as a carrier, this wafer does not remain in the package. Eliminating TSVs and the substrate not only reduces production costs, but also enables companies to make very thin packages, which provides electrical advantages.

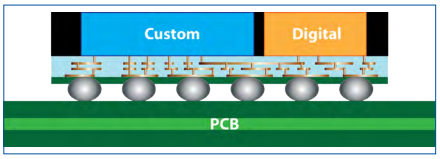

Figure 1: FOWLP can significantly reduce the package footprint.

FOWLP designs also provide the option of placing even more components vertically, using through-package vias (TPVs) that allow for package-on-package (POP) design. Unlike TSVs, TPVs act much more like the traditional package vias, reducing concerns about yields and reliability. TPVs are also less sensitive to inductive and heating effects than TSVs, as well as being less expensive to produce. The PoP design provides a smaller footprint, which enables designers to stay within vertical requirements, and allows a wide choice of 3rd party DRAM in the package. With their promise of smaller form factors providing low power with high performance [1], it's easy to understand the growing popularity of FOWLP. Smartphones and wireless multimedia are the first large commercial segment to use FOWLP designs, as seen in the newest iPhones from Apple.

FOWLP BENEFITS

FOWLP offers multiple advantages over conventional packaging technologies:

- Higher performance
- Shorter interconnect paths lead to fewer parasitics and less delay.
- Shorter paths to heat sinks reduce concerns about thermal impacts and resistance.
- Lower power consumption.
- Improved form factor
- Thinner and lighter than traditional 2.5D and 3D packaging
- Option to place even more components vertically, reducing footprint

While silicon interposer (SI) 2.5 and 3D designs do compete against FOWLP, many in the industry see them as two technologies that are targeted for different markets and different applications. FOWLP is frequently considered a better fit for automotive, radio frequency (RF), and mobile products, whereas SI and 3D designs leveraging TSVs are best suited for memory, FPGA, and GPU/CPU applications.

FOWLP MANUFACTURING (IN A NUTSHELL)

The basic FOWLP manufacturing process starts with a carrier material, typically a silicon wafer, attached to "double-sided tape." Individual die are placed on the other side of the tape, and covered with the EMC, with space allocated between each die for additional input/output (I/O) connection points. The carrier and tape are then removed, leaving the die embedded in an artificial EMC wafer. This wafer is then flipped upside down, exposing the die. RDLs are formed on top using physical vapor deposition (PVD) metallization and subsequent electroplating/patterning to re-route I/O connections on the die to the mold compound regions in the periphery. Solder balls are attached, and the die are singulated.

There are several ways to implement FOWLP, such as chips-first vs. chips-last and die face-up vs face-down. In the chips-first process, the die is attached to a temporary or permanent material structure prior to creating the RDL that extends from the die to BGA or land grid array (LGA) interface. With chips-first, the yield loss associated with creating the RDL occurs after the die are attached, where even known good die can face yield limitations imposed by the RDL packaging. In the chips-last process, the RDL is created first, and then the die are attached, which allows the use of only known good die and known good RDL. The RDL structures can either be electrically tested or visually inspected for yield loss before die placement, avoiding the placement of good die on bad sites. The difficulty is aligning the chips correctly to the already-generated RDL to form connections, which can introduce new yield concerns. For low I/O die, where RDL is minimal and yields are very high (>99%), a chips-first flow is preferred. However, for high-value die (large I/O), a chips-last process is preferred.

FOWLP METHODOLOGIES

Once you've decided to use FOWLP, the next big challenge is to decide which fan-out technology to adopt.

The first commercially viable FOWLP methodology was the embedded wafer-level ball grid array (eWLB) approach invented by Infineon. Assembled directly on a silicon wafer, the eWLB approach is unconstrained by die size, providing designers with the flexibility to incorporate an unlimited number of interconnects between the package and the application board for maximum connection density. Infineon chose to license their eWLB flow, so it is now readily available from most of the leading outsourced assembly and test (OSAT) providers. However, it is very difficult for OSATs to compete when they all use the exact same recipe and are constrained in pricing due to the licensing costs. As a result, while all major outsourced OSAT companies support eWLB packages, many also provide alternative proprietary versions.

As such, eWLB is not the only FOWLP technology in the market. A group of companies has formed the High-Density Fan-Out Wafer Level Packaging consortium, led by A*Star (a Singapore R&D organization), and including Amkor, Nanium, STATS ChipPAC, NXP, GLOBALFOUNDRIES, Kulicke & Soffa, Applied Materials, Dipsol, JSR Micro, KLA-Tencor, Kingyoup, Orbotech and TOK [2]. This consortium is looking at both the technical and cost issues of high-density FOWLP, and driving towards a high volume manufacturing solution to hasten its adoption in the industry. Currently, the consortium is developing two types of high-density fan-out packages—mold-first and RDL-first/chips-last.

For example, Amkor's silicon wafer integrated fan-out technology (SWIFT) incorporates some unique features not associated with conventional integrated circuit (IC) packages, such as the inclusion of polymer-based dielectrics, multi-die and large die capability, interconnect density down to 2 µm line/space (critical for SoC partitioning applications), copper pillar die interconnect down to 30 µm pitch, and the use of through-mold vias (TMVs) or tall Cu pillars. One drawback of the SWIFT technology is that it requires an extra silicon component (like a silicon interposer, but without TSVs). Because the OSATs can't manufacture these components, involvement of an additional 3rd party foundry process is necessary. Table 1 lists some of the OSAT FOWLP options currently available [3].

Table 1 - OSAT FOWLP options

Company/Package	Minimum Line/Space (µm)	Process Features	Application Targets
Amkor SLIM	<1 (0.4 possible)	BEOL Cu damascene	Partitioned die
Amkor SWIFT	2	RDL	Dual die, application and baseband processors
SPIL NTI	2	RDL	Dual die, application processors
Xilinx SLIT	0.4/0.4 for 3 Cu layers	Selective silicon removal	High to mid-range performance devices

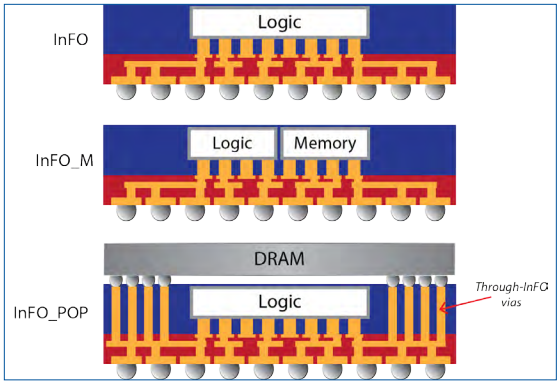

Figure 2: Cross-section of TSMC InFO design variations. TSMC also offers all of these options with an additional substrate included.

Foundries have taken note of the growing popularity of FOWLP processes as well. TSMC recently entered the FOWLP market with their integrated fan-out (InFO) wafer-level packaging [4], a silicon-validated technology that comes in different package sizes: 8x8mm2 (which allows mono-die or multi-die, and supports up to 600 I/O count), 15x15mm2 (allowing up to 2000 I/O count), and 25x25mm2 (allowing up to 3600 I/O count). As shown in Figure 2, one version of InFO, known as the integrated fan-out package-on-package (InFO-PoP), places DRAM on top of the FOWLP, using a TPV (which TSMC calls a through-InFO via, or TIV). Another variation, called InFO-M, supports multiple die placements.

Intel Custom Foundry recently introduced their embedded multi-die interconnect bridge (EMIB) as a cost-effective approach to in-package high density interconnect of heterogeneous chips (Figure 3) [5]. While not a true FOWLP design, it is a package option that eliminates the need for TSVs by replacing the large silicon interposers with small die-to-die silicon bridges, providing access to the fine pitch interconnects of a silicon process for the die-to-die interfaces. They are also rumored to be working on a true fan-out wafer-level packaging process, but no details have yet been made public.

With the inclusion of foundries in the mix, package designers have a new factor to consider: OSAT vs. foundry. The primary benefits touted by the foundries are 1) faster turn times due to the ability to share manufacturing lines, 2) single point of contact and services, 3) qualified tools, and 4) decades' worth of experience in delivering full design kits. For example, TSMC is currently moving the InFO technology into production [6], meaning they are building and validating InFO design flows, complete with required sign-off decks, providing designers confidence that the final taped-out package will work with reasonable yields. It is likely that the OSATs will ultimately be required to follow suit in this area, putting more resources into quantifying and validating their manufacturing tolerances.

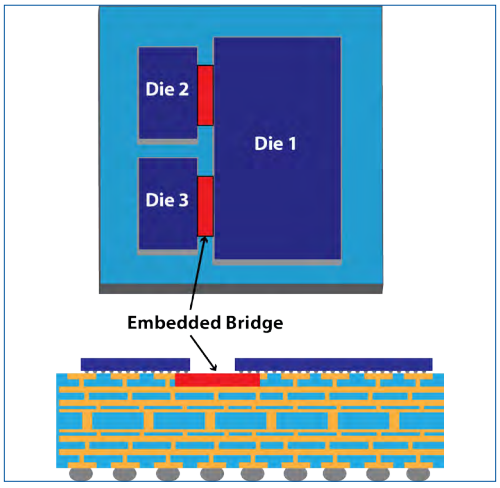

Figure 3: Top and cross-section view of Intel EMIB package.

Comparing a foundry FOWLP process to that available from an OSAT company illustrates some of the advantages and disadvantages of each approach:

Foundry

- Single-source process
  - All die tied to single foundry
  - Faster throughput time
  - Single point of communication
- Dedicated process design kit (PDK) and design methodology
  - Certified EDA tools and flows
  - Strict PDK compliance required
- Controls manufacturing capacity and cost

OSAT

- Multi-party chain
  - Can combine die from multiple foundries and processes
  - May incur transition delays
  - Multiple points of communication
  - Who owns yield issues?
- Historically, no rigorous PDK support or enforcement
- Typically compete on lowest price

The most obvious advantage for design companies is the option to work with a single point of contact and a single process flow. Foundries have extensive expertise in generating PDKs for advanced IC processes. Now they are bringing that knowledge and expertise into the package co-design world. Not only do they have the advantage of significant experience and long, historic relationships with the EDA companies, they also control both the package and the die process details. This control enables them to more readily access information across the IC and package domains than might be possible when working with multiple parties.

For example, it would be difficult for an OSAT to put together a generic solution to solve parasitic extraction issues where die signals and package signals affect each other, because the OSAT only has detailed knowledge about the package. Not only does it have little insight as to what might be in the dies, it may also have various legal restrictions as to what it is allowed to divulge to other IC suppliers.

What is not typically available through TSMC or another pure-play foundry is the ability to build an assembly based on dies produced by multiple foundries. It's not likely, for example, that one would send a die manufactured at Samsung to TSMC to incorporate into an InFO package. On the one hand, using a single foundry eliminates the need to transfer die from a foundry to an OSAT, saving time and ensuring that all the information stays together from start to finish. Alternatively, choosing an OSAT offers greater design freedom, but places the burden of ensuring all the components work together (and debugging in the case of errors) on the package designer, to work out through multiple 3rd party partners. This choice may materially impact pricing of a full FOWLP package, depending on a foundry's die manufacturing offerings.

This variety and customization means package design companies now have an increasing array of FOWLP options to choose from. Determining which offering is best for a given design requirement can be complex. In general, beyond price and manufacturer, product offerings are typically measured by what they can provide in terms of package footprint. Footprint is determined by the minimum vertical height that can be delivered, the minimum horizontal coverage as dictated by the minimum spacing allowed between die, and the minimum line width and space of the interconnecting RDL and BGA pitch.

What we can say is that it appears whether you target a foundry or OSAT offering, FOWLP is here to stay. While it does not replace all cases where a silicon interposer is required, it enables a new design methodology with several key advantages that package designers will likely want to take advantage of. You can expect FOWLP to be a continued area of growth over the next several years, as 3D IC technologies continue to expand.

FOWLP WITH MENTOR GRAPHICS

FOWLP design presents its own set of unique design and verification challenges. Historically, package design and IC design have been isolated processes, with little in common. Bringing them together requires melding the two into a single cohesive flow.

ICs are designed primarily in a Linux operating environment, using electronic design automation (EDA) tools certified by a target foundry or fab, and associated with a specific PDK. System-on-chip (SoC) designs are typically built using primarily Manhattan geometries, and represented in gridded formats like GDSII or OASIS. When the IC design is sign-off ready, the final approved design files (tapeout) are sent to the foundry for manufacture. When finalized, a die abstract (indicating the die's size and the individual pin locations) is passed to a package design team using any one of several formats (LEF, AIF, etc.).

The package design community predominantly works with EDA tools designed to run on a Microsoft Windows operating system. Packaging design uses non-Manhattan geometries extensively, and does not often map nicely to the gridded formats of the IC world. In fact, the two worlds share very few standards in terms of data representation. Package designs, along with the physical die, are historically transferred to a package house or OSAT facility using a number of formats, such as AIF, ODB++, and even Gerber. There is typically very little in the form of formal sign-off requirements that accompany the package design, other than a textual document describing the intended design rules.

Surprisingly, even though there are very few common standards between these two communities, no new standards are needed to bring the two together. It turns out that data format conversion in combination with tool-to-tool interfaces and communication protocols can be used to meet the needs of all the parties involved,. While communicating between Microsoft Windows and Linux can be tricky, experience has shown it can be readily and pragmatically achieved using a virtual network connection (VNC).

FOWLP DESIGN

Implementing a unified FOWLP design approach requires a significant expansion in the communication between the IC design world and the package design world. For instance, if you want to optimize the package design for size and/or performance, then you must optimize the entire system, not just the individual elements. An IC designer might be able to design a really small IC, but in doing so, it might be more difficult to connect that die into the package, expanding the package footprint. Similarly, a package designer might be able to design a clean and tight package, while making it impossible for the IC designers to get their die I/Os to match specific locations. To optimize the entire package design, IC designers must know more about the intended package, and package designers must know more about the ICs included in the package.

There are also a number of technical details that must be decided, such as the decision to fabricate the fan-out layers using IC or packaging technologies. Some of these issues include:

- Silicon RDL layers, organic RDL layers, or mold/plastic layers? Each type has its own unique design rules to increase yield
- Stress relief. Minimizing straight lines helps, but makes verification more challenging.
- Copper areas require complex outgassing perforations to prevent substrate warping, and copper density balancing across the design
- The connection path from die pad to copper conductor path requires filleting, the addition of metal to smooth the transition between geometry dimensions.
- Potential die shift during assembly requires flexible, adaptive patterns.

PACKAGE DESIGN

Mentor's Xpedition® Substrate Integrator technology helps IC, packaging, and printed circuit board (PCB) co-design teams visualize and optimize complex single or multi-chip packages integrating silicon on board platforms. The Xpedition Substrate Integrator co-design methodology automates the planning, optimization and connectivity from a chip through multiple packaging variables. Xpedition Package Integrator users can drive rule-based I/O-level optimization and perform pin and ball-out studies from their respective domains, visualizing the impact across the complete system, and generating an automated central data library in the process.

For complete package implementation, the Xpedition Package Designer tool provides the path from initial prototype to fabrication-ready GDS for HDAP that includes FOWLP, multi-chip module (MCM), system-in-package (SiP), and interposer-based 2.5D ICs, as well as traditional laminate plastic BGA (PBGA) designs.

Together, the Xpedition HDAP solution includes:

- Connectivity management
- Design domain connectivity models
  - Hardware description language (HDL), table(spreadsheet)-based and schematic editing modes
- Cross-domain signal shorting and splitting with automatic pin-mapping
- System-level connectivity tracking (layout vs. schematic (LVS) comparison) and verification (source netlist)
- Design tool aggregation
- Interconnect optimization
  - Visualize and optimize interconnect across system, including ICs, interposers, packages, and PCBs
  - Smart unraveling of interconnect paths
  - Layer reduction, improved signal quality

To address the rapidly changing design rules of HDAP/FOWLP, the Xpedition Package Designer tool combines a powerful geometry processing engine with a field solver using HyperLynx® DRC technology. The HyperLynx DRC technology supports creation of custom rules checks, and provides independent, in-design verification above and beyond standard design rule checking (DRC). These capabilities allow package designers to produce a cleaner design without the iterations of the GDS verification engineering change order (ECO) flow.

The Xpedition Package Designer tool also includes the proven HyperLynx Fast 3D Solver field solver for full-package model extraction (Figure 7). It's ideally suited for power integrity, low-frequency simultaneous switching noise (SSN) and simultaneous switching output (SSO), and complete-system SPICE model creation, while accounting for skin effect impact on resistance and inductance.

Figure 4 illustrates the use of the Xpedition Package Designer and HyperLynx tools in the FOWLP process.

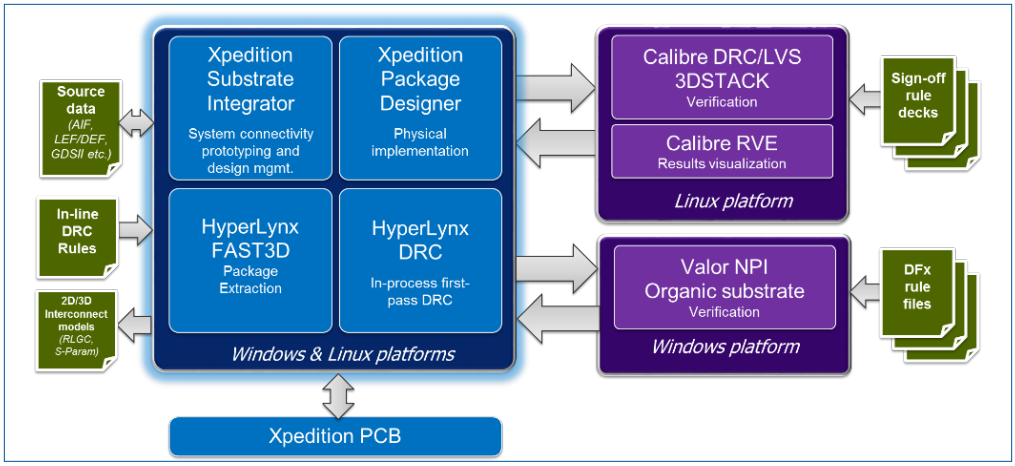

Figure 4: The integration of Xpedition package design and HyperLynx tools enables a FOWLP co-design flow.

Traditional SoC design processes contain fully-qualified verification methods for IC designs in the form of process design kits (PDKs) from the foundries. IC designers use these PDKs to not only reduce their risk, but also to improve overall productivity, by implementing repeatable, proven verification techniques.

Until now, chip design companies and OSATs have had no corresponding signoff verification process to ensure that an IC package meets manufacturability and performance requirements. This lack of qualified verification processes put package designers at a significant disadvantage compared to their IC design counterparts. Package die are often produced using multiple processes and multiple foundries, which not only raises the level of complexity, but also exposes the need for a process that ensures these disparate products can be manufactured within a single package. Package designers must often work with poorly-documented and/or inconsistent processes each time they submit a package design. In addition, the new class of packages now coming into the market enhances the interactions between the layers, so there is no clear separation between the traditional die and package.

The challenges of verifying IC packages are numerous. Die design teams often have very different goals from the package design team, which creates unforeseen integration issues. Chips often come from different foundries, and are verified using different processes, making package failures hard to identify and fix. There hasn't even been a formal verification process to ensure the connectivity from a single die to the BGA of the package is correct. We simply haven't had good characterization of package processes and requirements, leaving chip designers and OSATs to muddle through.

Not that the industry isn't trying. OSATs have built and manufactured packages for years, and developed a variety of tools to help in the process. However, each design team must write its own rules for each assembly, with no reference signoff deck that ensures the manufacturability and performance of the package. In addition, OSATs rarely hold their customers to a hard and fast rule deck like a foundry does.

Mentor has pioneered an entirely new approach for sign-off quality physical verification of packages, what we call an assembly design kit (ADK). The purpose of an ADK is similar to that of the PDK—ensure manufacturability and performance. What makes that happen, in both PDKs and ADKs, are things like standardized rules that ensure consistency across a process, qualified tool flows, interface formats, input/output formats—in short, everything a designer needs for successful design, tested and qualified and proven to produce working products. For example, one component of an ADK would be sign-off checks for the package layers and interacting die layers, without reference to any kind of layer mapping. To jumpstart the supply chain into producing ADKs, Mentor offers an OSAT Alliance program that can ease the adoption of emerging HDAP technologies.

EDA vendors who support an ADK ensure their tools can implement these checks and enable debugging of any errors. An ADK provides many benefits to the industry: reduced risk of package failure, increased packaging business, and increased use of 2.5/3D packages.

Why Can't We Use Existing Verification Flows?

There are multiple reasons why IC-centric verification solutions do not work for package verification. Physical verification (PV) tools rely on GDSII or similar layout formats for their inputs. Because these formats contain no data about a geometry's vertical placement, that information is inferred through layer mapping and the use of typical layer-naming conventions. For example, PV tools infer that metal2 is vertically higher than metal1 and lower than metal3, and establish the layers' electrical connectivity through the appropriate via layers. All geometries mapped to the same layer are considered to be co-planar, and any geometries that overlap or abut are treated as if they are a single polygon, as they will be merged during the mask generation process.

However, there is generally little to no commonality in layer mapping from one foundry to another, and in some cases, even from one node to another at the same foundry. The design rule checking (DRC) tool would interpret the geometries located on the same layer as the geometries from the package RDL as co-planar, and considers those geometries as if merged, when in reality there is a significant vertical displacement between them. For this reason, traditional IC PV tools do not easily lend themselves to this kind of packaging task.

FOWLP VERIFICATION

To sum it up, an entirely new approach is needed for sign-off quality physical verification of FOWLP packages. Let's start with the requirements of FOWLP verification. Fortunately, physical verification is significantly simplified when you realize that you do not need to check every geometry in every die. Each individual die in a FOWLP package has already been taped out for its target foundry with respect to DRC and layout vs. schematic (LVS) comparison. What is required in the FOWLP verification is to check the interactions between the dies, although that is not to say this is a trivial process. In some cases, this checking can require extracting several layers within each die to see what their impacts are each other and the full package. To ensure the results of these checks are accurate, the tool used must understand the layering per die and per placement.

Physical verification of FOWLP designs must address the unique physical characteristics of these packages:

- Multiple components
  - Multiple die from different sources/processes
  - Same layer may be used in different die placements (not co-planar)
  - Package
  - BGA
  - Substrate
- Non-Manhattan shapes
  - Vias, tapers, BGA, etc.
- Disparate file standards
  - GDSII vs MCM, ODB++
  - SPICE vs comma-separated values (CSV) spreadsheet

FOWLP physical verification consists of two primary steps (Figure 5):

- Verify each unique component per process technology
- Verify interfaces to the package
  - Package geometry DRCs
  - Package geometry to die geometry
  - Connectivity checking from die through package

Package DRCs, package LVS, and interface checks can all be combined into a single Calibre 3DSTACK deck and checked in one run. The only individual runs required are for the die-specific DRCs and LVS.

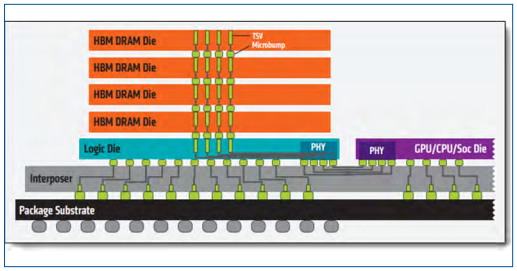

Figure 5: FOWLP physical verification consists of component and interface verification checks.

As mentioned earlier, a verification tool must be able to implement sign-off checks for the package layers and interacting die layers, without reference to any kind of process specific layer mapping. Such checks include ensuring that connections are intact within the package (die to die) and from the package to the outside (die to BGA).

Next is validation of the electrical connectivity across the assembly. In the IC space, designers rely heavily on standards such as SPICE, which represents connectivity all the way down to each individual device component. By associating device components with their corresponding device models, and tracing their connections to all other devices in the circuit, they can perform electrical simulations. Unfortunately, despite being called a "netlist," SPICE does not actually contain direct information on connections. Instead, it is a list of devices and device pins from which connections are inferred. This inference does not work in the package assembly space, because package design tools do not deal with data down to the transistor level. With no true devices, a traditional LVS tool relying on a SPICE-style netlist has nothing to check in a package. Other file formats, such as AIF, are used to represent the connections from the pins of one die to its surrounding.

Fortunately, the Calibre nmPlatform can handle all of these verification needs. The Calibre 3DSTACK functionality embedded in the Calibre nmPlatform works by separating process and assembly integration rule requirements from the design assembly definition. The package house or OSAT is responsible for creating, validating, qualifying and delivering the package and interface rules. The package designer is responsible for describing how the assembly is put together, so the tool can understand how to differentiate the layers per die. This step can be automated by extracting the assembly design information from the design tools used to build the flow.

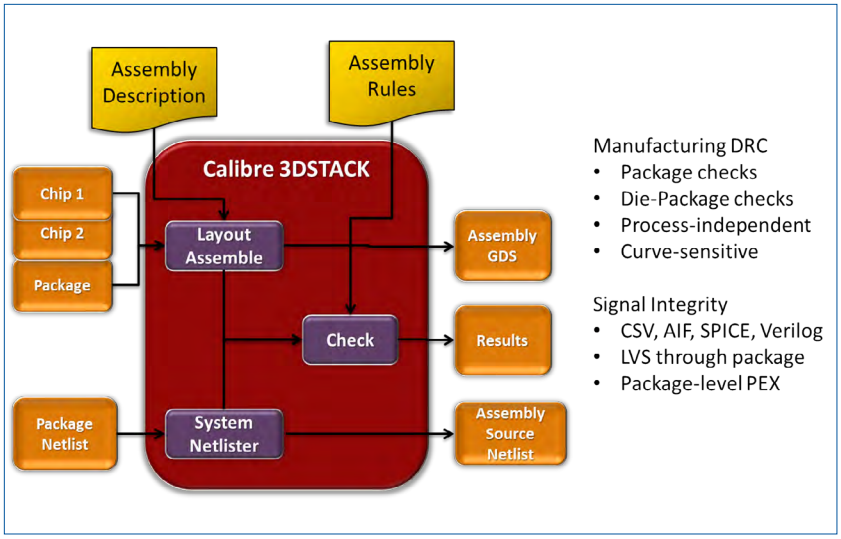

Figure 6: The Calibre 3DSTACK functionality enables and supports all of the unique verification needs of FOWLP designs.

Syntax in the Calibre 3DSTACK tool can combine the two types of rules and provide rule checking capabilities. The Calibre 3DSTACK method works by separating out the process and assembly integration rule requirements from the design assembly definition. With this information, the Calibre 3DSTACK functionality builds its own view of the assembly, and checks to determine that it satisfies all the requirements (Figure 6). This checking includes the comparison of the physical package assembly against a source netlist, which can come in traditional IC forms like SPICE or Verilog, or traditional package tools like AIF or CSV spreadsheet formats. On output, the Calibre 3DSTACK process provides errors that can be highlighted back into the package design using the traditional Calibre RVE™ display capabilities, as well as a generated extracted netlist representing the assembly for cross-probing of connectivity results.

The OSAT or foundry determines the connectivity stack for LVS comparisons on FOWLP designs, and creates design rules addressing package-specific requirements as well as die-to-package geometric interface requirements. For example, package wire widths vary depending on whether they are over die, between die, or completely outside the die, so rules are needed to govern the design of wires in the package, including specifics such as size, distance, etc. Design rule checks can then be implemented to ensure all requirements are satisfied in the package. Typical checks (Figure 7) include:

- Verify DRC and LVS across interface boundaries between package and dies
- Perform DRC checks on RDL layers
  - Avoid orthogonal/90° angles on RDL geometries
  - Area-dependent line width and spacing (L/S) checks
- Perform metal density checks
  - Fill check for floating metal fill pattern and/or biased metal plane with slotting
  - Local-based density checking

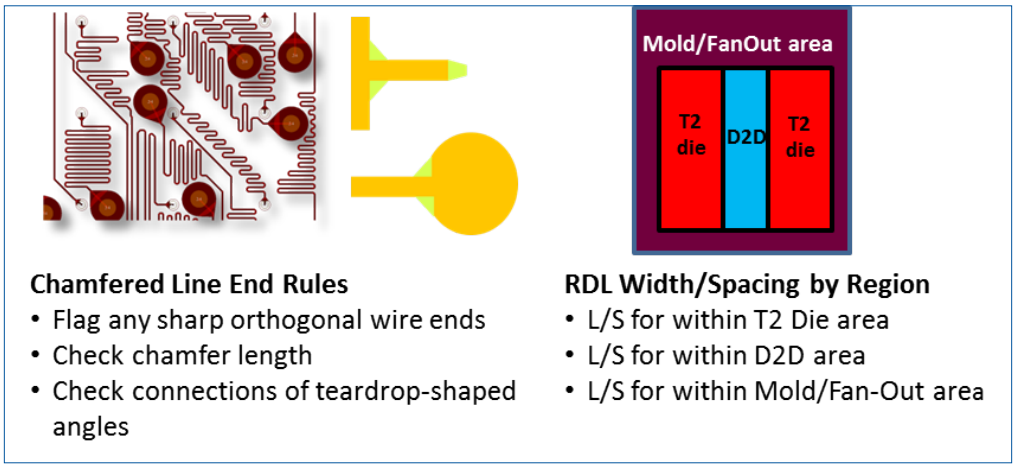

Figure 7: Design rule checks in an ADK target specific geometries found in package designs.

The OSAT or foundry also writes the rules for the types of elements and configurations permitted in the package.

- Assembly rule checks
  - Die-to-die edge
  - Die-to-package edge
  - Die-to-package alignment
  - Corner rules
  - Landside Caps to BGA spacing
- Mechanical stress checks
  - Die-Package Fan-out ratio

To eliminate layer mapping issues, the Calibre 3DSTACK functionality automatically re-assigns all layers for each component (die, interposer, packaging layers, etc.) to separate layers. This reassignment must be done per placement to avoid problems caused by multiple placements of the same die. However, reassignment alone is insufficient. DRC is typically coded for the process layers available in a chip die, but because these layers are now all

reassigned to unique names, the implication is that assembly checking will always require a unique rule file specific to the die layers in the specific configuration. Instead, the Calibre 3DSTACK tool provides the rule writer a level of abstraction to indicate types of layers (pads, bumps, RDL, etc.), so rules can be associated with these types.

When the Calibre 3DSTACK process is run, the Calibre platform can associate the specific layers of each die or package component to the appropriate layer type. Knowing where each component interacts to others allows the Calibre 3DSTACK tool to automatically regenerate a detailed set of checks between the appropriate layers from each component, independent of how the design is configured, how many dies are placed, or for what processes those dies are targeted.

For LVS, the first challenge is the absence of a traditional source netlist for a package. The de facto industry substitute, a CSV spreadsheet format, is typically used. LVS connectivity through the package RDL is confirmed by verifying the physical connections between external pins to the die. To enable designers to check the connectivity of a standalone package design (no devices), as well as improve the existing solution for capturing netlist (connectivity intention) for a fully-stacked system (IC + package), virtual die are generated from the spreadsheet netlist pin coordinates. The Calibre 3DSTACK tool can perform connection checking through the package validating die-to-die and die-to-BGA connections, and identifying any unintended shorts and opens.

A first-level extracted netlist is generated by automatically combining the extracted package netlist from the Calibre 3DSTACK tool and filling in the die details with the already extracted parasitic netlists for each die. This task can be done using the System Netlist Generator (SNG), a part of the Calibre 3DSTACK offering that can be used to create source netlists for input to the Calibre 3DSTACK tool, or to stitch multiple netlists together into a single one to create the parasitic netlist. SNG can be run in batch or GUI mode, and can also output a reduced netlist for specific paths cross-die through the assembly. Once created, these netlists can also be used for detailed SPICE analysis or other electrical analyses, such as reliability checking for electrostatic discharge (ESD) or electromigration (EM) with the Calibre PERC™ tool.

But what about the potential parasitics between die? There may be capacitive or inductive coupling from routing in a die to an external bump, or even from routing in one die to routing in its neighbor. The solution is to use the Calibre 3DSTACK tool to re-extract each die, but include one or two layers from the neighboring die. Of course, these new layers must be accounted for in the calibration. This is achieved by performing a much simpler 'incremental calibration' that extends to the existing MiPT files generated by the target foundry for the stand-alone die processes. With this new parasitics per die information, designers can again combine the data into the assembly netlist to generate a more accurate extracted netlist.

Using the Calibre 3DSTACK tool, the OSAT or foundry creates a rule file for the FOWLP process that can be used by any designer targeting this package technology at that OSAT or foundry, regardless of what processes the dies are, or how many dies are in the package. The rule file checks the manufacturing constraints of the package RDL and the die-to-die constraints, and verifies the connectivity through the package from die-to-die and die-to-BGA. It is entirely independent of any specific design tool used to generate the package. Designers applying this rule file can leverage traditional Calibre functionalities for debugging and review.

The Calibre PERC reliability verification platform provides verification beyond the traditional DRC, LVS and electrical rule checking (ERC) available with the Calibre 3DSTACK functionality. It delivers a unique reliability solution that can consider the design content and provide a system view of how design elements interact (Figure 8). Typical application areas include:

- ESD, electrical overstress (EOS), and latch-up detection
- Voltage-aware DRC
- Low/multi-power domain designs
- General reliability issues

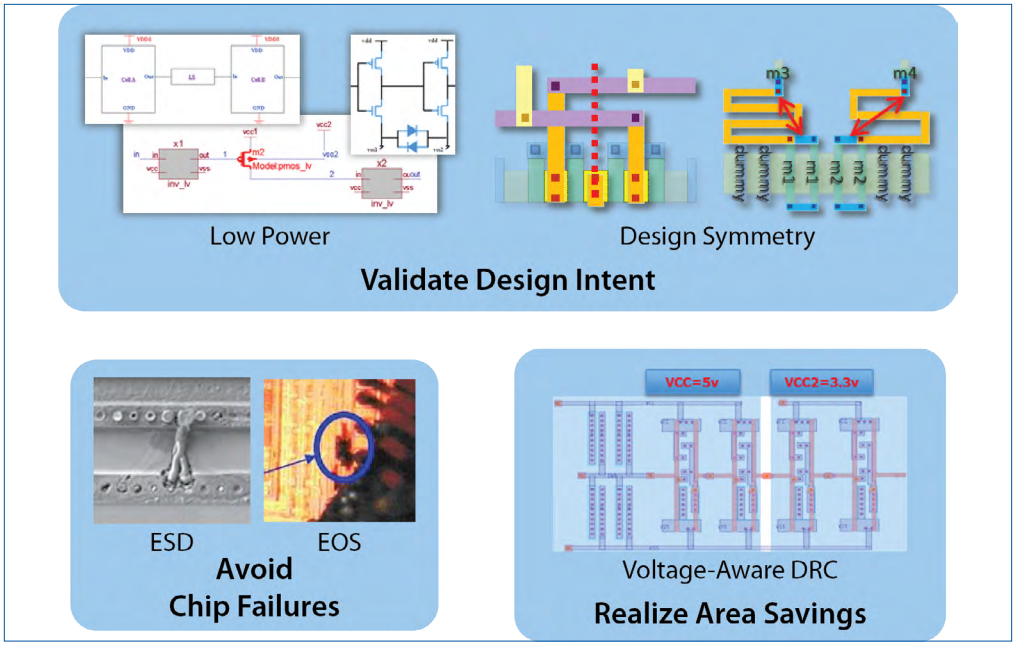

Figure 8: The Calibre PERC reliability platform can assess a wide range of potential performance and reliability issues in FOWLP packages.

PRODUCTION-PROVEN RESULTS

Mentor Graphics successfully developed and proved the viability of an assembly design kit (ADK) using the building blocks of the Calibre 3DSTACK solution. A Calibre 3DSTACK deck was created and qualified by STATS ChipPAC, a leading OSAT, for their eWLP wafer level packaging process. This rule file was adopted and proven in production at Qualcomm on an assembly comprising two die connected through the package RDL. The rules from STATS ChipPAC included the DRC rules for the package layers themselves, plus checks to ensure that the dies interacted and connected within the package properly.

SUMMARY

FOWLP design and verification introduces new requirements and challenges, which require new EDA tools, functionalities, and flows. Because FOWLP design and verification requires cooperation and collaboration between design houses, OSATs, foundries, and EDA vendors, using common tools that have the integration and functionality needed to operate in both the IC and packaging domains can reduce risk of package failure, while also reducing turnaround time for the component providers, OSATs and foundries. Using Mentor Graphics EDA tools ensures that FOWLP package designers have all the design and verification capabilities they need to take advantage of this exciting new packaging technology and the markets it serves.

REFERENCES

[1] C. C. Liu et al., "High-performance integrated fan-out wafer level packaging (InFO-WLP): Technology and system integration," Electron Devices Meeting (IEDM), 2012 IEEE International, San Francisco, CA, 2012, pp. 14.1.1-14.1.4.
doi: 10.1109/IEDM.2012.6479039

[2] Fan-Out Wafer-Level Packaging (FOWLP) Consortium, https://www.a-star.edu.sg/ime/INDUSTRY/INDUSTRY-CONSORTIA/Fan-Out-Wafer-Level-Packaging-FOWLP-Consortium.aspx

[3] TechSearch International, Inc., "The Future of Packaging and Assembly Technology," presentation at Central Texas Electronics Association, 2016. http://www.smta.org/chapters/files/Central-Texas_CTEAVardaman2016.pdf

[4] Letter to Shareholders, Taiwan Semiconductor Manufacturing Co Ltd, 2015. http://www.tsmc.com/download/ir/annualReports/2015/english/pdf/e_1_1.pdf

[5] Intel Custom Foundry. "Embedded Multi-die Interconnect Bridge," Intel Corporation, http://www.intel.com/content/www/us/en/foundry/emib.html

[6] Taiwan Semiconductor Manufacturing Co Ltd Earnings Call, October 15, 2015. http://www.tsmc.com/uploadfile/ir/quarterly/2015/3C2bO/E/TSMC 3Q15 transcript.pdf

For more information on the Mentor Graphics tools and functionalities discussed in this paper:

Xpedition HDAP co-design

HyperLynx simulation solutions

Calibre 3DSTACK multi-die verification

or

contact a Calibre representative

For the latest product information, call us or visit: www.mentor.com

Corporate Headquarters
Mentor Graphics Corporation
8005 SW Boeckman Road
Wilsonville, OR 97070-7777
Phone: 503.685.7000
Fax: 503.685.1204

Sales and Product Information
Phone: 800.547.3000
sales_info@mentor.com

Silicon Valley
Mentor Graphics Corporation
46871 Bayside Parkway
Fremont, CA 94538 USA
Phone: 510.354.7400
Fax: 510.354.7467

North American Support Center
Phone: 800.547.4303

Europe
Mentor Graphics
Deutschland GmbH
Arnulfstrasse 201
80634 Munich
Germany
Phone: +49.89.57096.0
Fax: +49.89.57096.400

Pacific Rim
Mentor Graphics (Taiwan)
11F, No. 120, Section 2,
Gongdao 5th Road
HsinChu City 300,
Taiwan, ROC
Phone: 886.3.513.1000
Fax: 886.3.573.4734

Japan
Mentor Graphics Japan Co., Ltd.
Gotenyama Trust Tower
7-35, Kita-Shinagawa 4-chome
Shinagawa-Ku, Tokyo 140-0001
Japan
Phone: +81.3.5488.3033
Fax: +81.3.5488.3004

Mentor® A Siemens Business

MASB 06-17 SSWP-00004