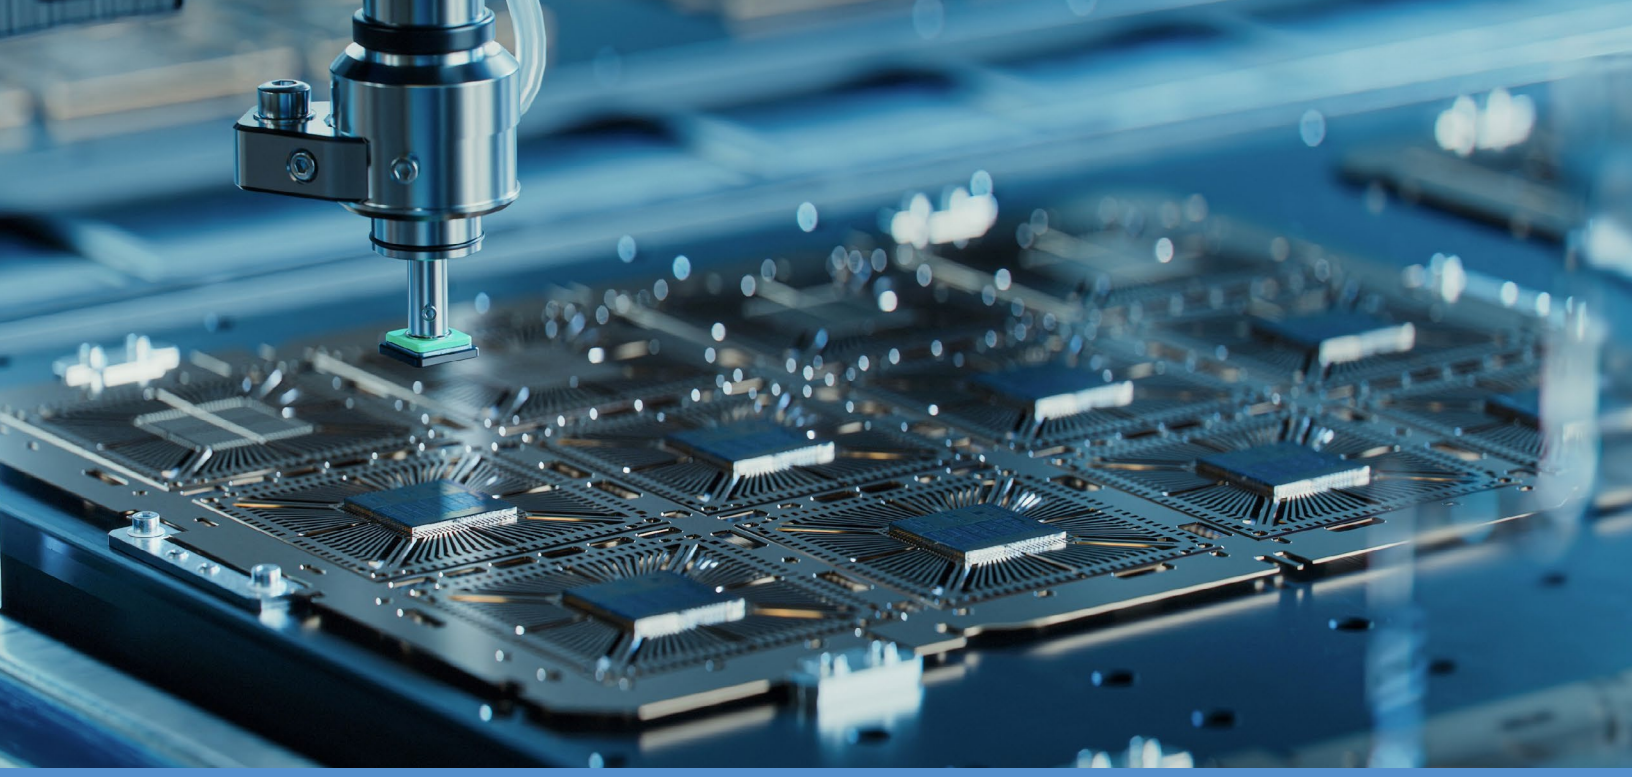

# Advanced Packaging Fundamentals

## for Semiconductor Engineers

SEMICONDUCTOR ENGINEERING
DEEP INSIGHTS FOR THE TECH INDUSTRY

# Contents

# Introduction

**By Bryon Moyer**
With the SE Staff
Designed by Jesse Allen

**Technical advisors:**

- John Park, product management group director for advanced IC packaging at Cadence Design Systems
- Lee Harrison, director of Tessent automotive IC solutions at Siemens EDA
- Robert Patti, president and CEO of NHanced Semiconductors
- Gretchen Patti, senior member of technical staff at NHanced Semiconductors

Integrated circuits have been housed in packages since the beginning of the semiconductor industry. The original idea was largely to protect the delicate piece of silicon inside from the elements outside, but the nature and role of packaging have evolved dramatically in the last decade. While chip protection is still important, it has become the least interesting of the roles that packaging provides.

This eBook addresses the biggest changes in packaging, typically referred to as *advanced packaging*. There's no clear definition of what *advanced* means. Instead, the term broadly covers a number of possible packaging scenarios, all of which are far more complex than traditional one-chip packaging. Advanced packages typically encapsulate multiple components, but the means of assembling the unit can vary widely.



Common in such a discussion are mentions of 2.5D or 3D packaging, with those descriptions referring to the ways in which internal components may be arranged.

This starts with a discussion of package types as viewed from the outside, moving inward from there to the basic components that advanced packages integrate. After that it explores each of those components in more detail. The bulk of the discussion will deal with the varying processes by which an advanced package can be assembled.

The report wraps up by examining four topics that must accompany any discussions of technology — how engineers design advanced packages, how they test them, the reliability implications of advanced packages in general, and any security implications.

Two related broad topics are addressed briefly, as well. The first is bonding. Although a necessary part of packaging, it is a big topic on its own and one that will not be discussed in detail here. The second includes the various types of components that aren't integrated circuits, but which may be included in a package. Optical components and MEMS (microelectromechanical systems) are two prominent examples, and each has its own set of considerations that can push outside the scope of this eBook. So again, they are discussed at a high level.

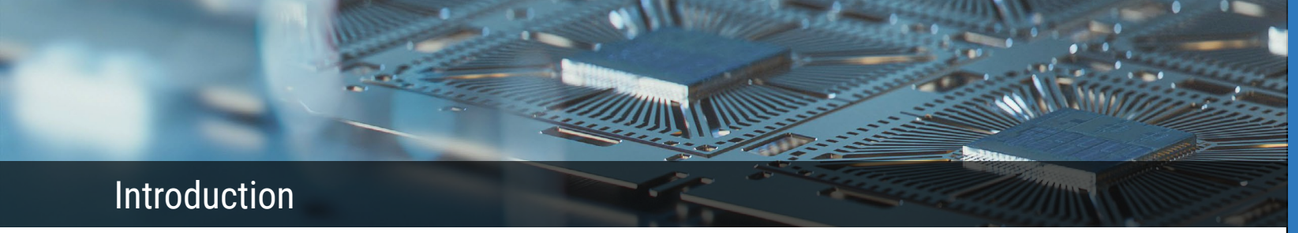

# Introduction

Coverage focuses on techniques largely in production today, since they're relatively stable. Developments since mid-2024 on are not included because they typically deal with new ideas that are still in flux and haven't settled out yet. Semiconductor Engineering will continue to refresh the content in future reports, adding any topics that have made it into commercial production since the prior edition.

Our expectation is that readers are electrical/electronic engineers — or at least that they have a good grasp of basic engineering and semiconductor concepts. Electrical notions in packaging largely relate to connections and signal integrity. Circuits or transistors, or any of the functions that may reside on the chips in the package, are not part of this report. Readers are assumed to know what processors, memories, and other fundamental chips are.

Packaging also involves a wide range of materials, and those may be less familiar to someone trained to work with digital logic or analog circuits. So, where helpful, this eBook will describe any such materials in more detail. Similarly, processing steps may use a number of arcane terms and brands, which will be defined and described, because they're typically less familiar to an electrical engineer. The goal is to keep this accessible and relevant to anyone with a working knowledge of semiconductors.

This topic also raises a question about the distinction between a chip and a die. In the past, there was no real distinction, because a package housed a single chip. The question is, however, does the chip refer to the packaged product, or is it simply another name for a die? Some people use the latter definition, while others now say a chip is the packaged unit. By that definition, an advanced package is likely to be a single chip containing more than one die. We have not tried to adjudicate this debate and have tried to use the words in ways that should be obvious from the context.

Finally, as this is a new area, there are very many configurations and approaches to advanced packaging. Most of them are in research; a few of them are in commercial production, but without any standards. The focus will be on those shipping in volume, adding more in the future as they gain traction.

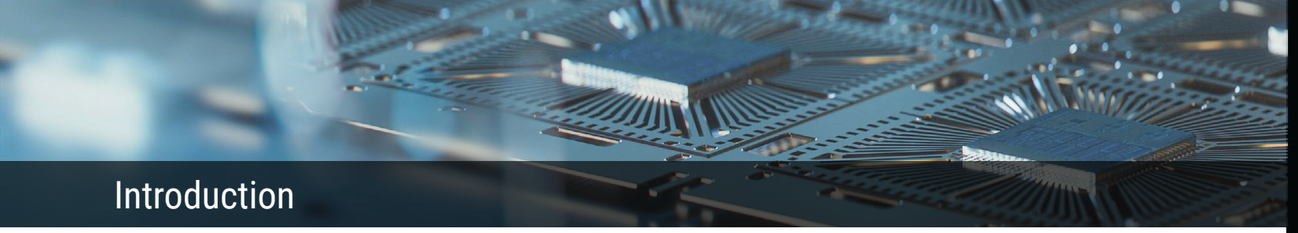



## The overloaded notion of Heterogeneous Integration

Advanced packaging implements what is often called *heterogeneous integration*. Although liberally bandied about, the term has no strict definition, and various entities use it in slightly different ways.

Some use it to mean integration of more than one element in a package, although this is infrequent because it doesn't really meet the definition of heterogeneous if the components are merely multiples of the same die. Others use it to indicate that different chips or chiplets are co-packaged (where, loosely speaking, a chiplet is a die that has some I/Os designed to communicate with another die within a package). Yet others refer to the inclusion of non-electrical components, such as optical.

We largely avoid the term here as a buzzword with imprecise connotations, other than the fact that there are multiple things in a single package. All of what follows could be considered heterogeneous integration, depending on one's definition.

# Why advanced packages?

Advanced packaging follows the age-old move toward increasing integration. In this case, rather than integrating onto a single chip, it's integration of multiple components into one package. The motivation for doing so ties directly into a few different trends, although those trends are often intertwined. One is the increasing use of chiplets, while another is the halting move toward co-packaged optics. These two examples typify the two primary issues driving this whole movement — bandwidth and power. Cost also may be a factor motivating chiplets, but that's specific to chiplets, and a fully costed advanced package may still be expensive, despite the use of any chiplets.

## Higher bandwidth

The bandwidth in this case refers to the speed with which components can exchange data — particularly when accessing memory. Faster physical communication mechanisms, such as PAM4 and PAM8, are part of the solution, but at this point increasing gains are hard won.

Given a set speed for a signal, another way to increase bandwidth is simply to make more signals available. *Printed circuit boards* (*PCBs*, also known as *printed wiring boards*, or *PWBs*) provide a given set of line and spacing rules that limit the number of signals that can run next to each other. Additional layers on the PCB can help with the routing, but at some point costs mount, and longer, more tortuous routes also hinder speed.

Achievable dimensions inside a package are much more aggressive than those available for PCBs, allowing more signals. This is specifically what gives *high-bandwidth memory* (*HBM*) its value — a much wider bus than is possible on a PCB. Distances are also much shorter, permitting more aggressive signaling techniques.

This notion of how many signals to run ties into the concept of *beachfront*. Beachfront is measured in terms of I/Os per square millimeter and relates to the line/space rules inside the package, as well as the ball dimensions of the dies. Although older chips had peripheral pads for I/Os, modern chips have arrays of balls. The first few rows of that array may be usable as related I/Os in a bus, but the number of usable rows is limited by how easy it is to route inner rows out.

In theory, any ball can be routed out, but when it comes to buses, the signals need to be as well matched as possible to limit skews across the bus. Those considerations limit the number of ball rows that can be part of a single bus, which along with the ball pitch determines how many I/Os

are available for a bus per mm2 of the edge of the die. Routing may be further limited if noise considerations necessitate ground lines between signals.

## Lower power

Reducing power is the other big motivator, and that relates directly to the distances signals must travel. A longer trace requires more signaling energy to ensure that data travels to its endpoint in good enough shape to be accurately received and interpreted.

Many more signals will travel inside a package than would be possible on a PCB, but because distances are measured in millimeters rather than centimeters, drive strength can be eased, saving energy. It's possible that net energy could still be higher given the greater number of signals. In that case, the bandwidth motivation (if not simply the space savings) would be the more important one, sacrificing total power for higher performance. But even then, the power per signal will be lower.



# Package types

The semiconductor industry has developed countless types of packages for integrated circuits (ICs). Most of them contain one chip and serve both to protect that chip from the ambient environment and to achieve a means of attaching the chip to a PCB. We will not be addressing most of those package types.

"Advanced packaging" is a vague and somewhat unhelpful term. At any given point in time, the latest packaging technologies, whatever they are, qualify as advanced. So expectations of what that means today may change in the future. This report will narrowly define which types of packages are being discussed.

## Through-hole vs. surface-mount

Older packages have pins intended to go through a hole drilled in a PCB. These are simpler packages for simpler boards. They're assembled by placing all the components on the top side of the board and then running it through a wave-soldering process, where a "wave" of molten solder gently wipes the bottom of the board, sticking to the appropriate pads and being wicked up into the holes around the package leads to form reliable connections.

This is a well-established technology, and it's relatively inexpensive. The downside is that only the top side of the board is available for components. The through-holes and solder wave make assembly on the back side impossible. So this approach results in a low density of components.

Surface-mount technology addresses that challenge by eliminating pins that go through a board. Instead of pins, solder balls are attached to the outside of the package. All such components are placed on a board and run through a heat cycle that slightly melts (or *reflows*) the solder balls so that they make a clean

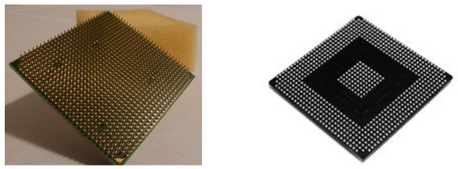

Pin-grid array Ball-grid array

*Fig. 1: Pin-grid array vs. ball-grid array. The left image shows the bottom of a package with pins intended to go through holes in a PCB. The right image shows the surface-mount equivalent, with solder balls that mount to the surface of the PCB. Note that the arrays don't have to be fully populated with pins or balls. (Sources: Left: Liam McSherry, CC BY-SA 3.0, via Wikimedia Commons; Right: ASE Group)*

connection to a pad on the surface of the PCB. That leaves the back side of the PCB available for other components.

The types of packages considered here typically have a huge number of connections, and *pin-grid arrays* (PGAs) and *ball-grid arrays* (BGAs) provide two examples of high-lead-count package. The former is a through-hole package; the latter is surface mount.

This eBook focuses only on surface-mount technology, and the BGA is the most widely used style of package that meets the remaining criteria.

## Edge leads vs. lead arrays

Older packaging technology works by using wires to attach chip bonding pads to a leadframe that routes the signals from the die's bond pads to the package pins. Those bond pads are all located on the edge of the die, and the resulting pins are at the edge of the package.

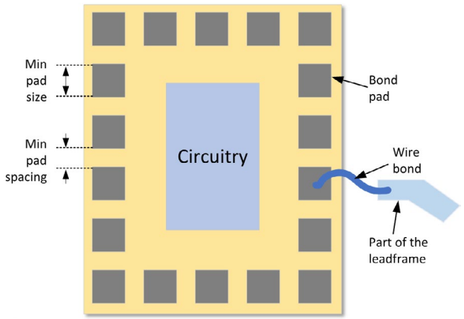


*Fig. 2: Top view of a pad-limited die. The size of the die is set by the pads on the periphery. Shrinking the circuitry in the middle would not result in a smaller die unless the pads were removed.*

Such an arrangement places a limit on how small a die can be since bond pads on a die have a minimum required size and spacing governed by the size of the wires, not by the silicon process. A very small circuit that needs lots of connections may need so much room for the bond pads that the die size is determined only by the pads, not the circuitry on the chip. Such a die is said to be *pad-limited*.

For hundreds or thousands of connections, leads on the edge would result in enormous packages and horrible performance due to the length of the connections and leadframes. Instead, modern large packages have arrays of leads — balls in the case of a BGA. Array leads can derive from edge pads on a die if the die is extremely large, or the die itself can have an array of balls with circuitry skirting around them.

## Single component vs. multi-component

There are many possible reasons for wanting to integrate multiple chips into a single package. The resulting size will use less space on a PCB than if multiple packages were used. Performance can be higher and more efficient because of shorter connections. And in many cases, fewer connections to the PCB may be necessary.

The latter effect is related to a relationship established years ago between the number of gates on a chip and the resulting number of I/Os. Called Rent's Rule, it finds that as the number of gates on a chip increases, so does the number of I/Os — but not as fast. That's because many connections remain internal to the die.

The same effect happens with the package. If two chips that might otherwise be in separate packages have connections between them, those connections will disappear from the PCB because they're made inside the package instead.

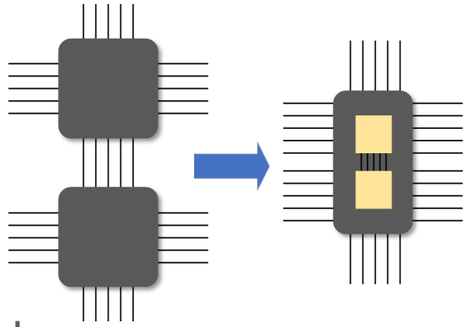

*Fig. 3: The left image shows two chips that share five connections. If those two chips are co-packaged together, then those five shared connections are made inside the package and disappear from the package's leads.*

## The confusing notion of an RDL

Many packages include what's called a *redistribution layer*, or *RDL*. The original concept was that of a few layers of interconnect for routing signals from one pattern to another — typically from the lead or ball pattern of a chip package to a landing pattern on a PCB. This becomes particularly important when the connection pitch on the package is too tight for line and space rules on a PCB. An RDL is necessary for taking those signals and spreading them out.

That's the general idea, and it applies in most advanced packages. But many elements can serve to reroute signals, including interposers and package substrates. Technically, these act as RDLs. But the term *RDL* seems to have a more specific usage that refers to the addition of routing layers above the metal layers of a die, added after die passivation, or on the backside of the die, rerouting *through-silicon vias* (*TSVs*).

In general, they have a few layers built up out of organic resin that provide nothing more than a rerouting of signals. This more limited definition of an RDL doesn't permit component embedding passives or other components in a way that might be possible with interposers.

## Fan-in vs. fan-out

Because older technology runs wires from die pads to a leadframe, the signals have nowhere to go but away from the die. The package footprint on the PCB is larger than the die it contains. In current parlance, the signals fan out from the die.

Advanced packaging techniques include the use of RDLs that can route the signals anywhere. If the die has few connections, they can be routed under the die, with the resulting package being only slightly larger than the die itself. Called *chip-scale packaging* (*CSP*), it's the smallest possible practical packaging technique, given there's no way to go smaller than the die. Anything up to about 1.2X the size of the chip is called chip-scale. And because in this situation the signals route from the edges inward under the die, the technique is called *fan-in*.

The opposite is *fan-out*, where some leads are routed away from the die, making the package larger than the die (even if some signals fan in as well). This may be necessary even for a single chip because PCB design rules require connections to be much farther apart than bond pads on a chip.

Finally, the advanced assembly processes include manufacturing of *panels*. These are larger rectangular boards that effectively act like PCBs, but with more aggressive dimensions. Panels haven't yet achieved mainstream status.

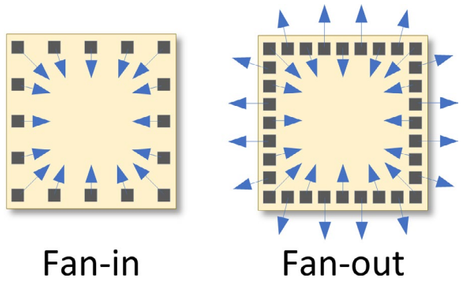


*Fig. 4: Fan-in vs. fan-out. The left image shows all signals being routed towards the center of the package, which is possible because of the limited number of signals. The right shows fan-out, with some signals routing out beyond the die. With fan-out, some signals can still route inward.*

In summary, then, advanced packages for the purposes of this eBook will be characterized as:

- Using surface-mount technology (probably BGA or related);
- Having bump arrays rather than edge connections;
- Encapsulating multiple components (with the exception of CSP); and
- Having a fan-out RDL or some other element that routes signals away from a die.



# Advanced-package components

Older package styles tend to have few components: a substrate, a leadframe, a chip, and then molding compound (for a plastic package) or some other enclosure. Advanced packages are based on this architecture, but they add a few elements:

- Substrates provide the connection between the package contents and the PCB to which the package will eventually be attached.
- Balls made of solder form the PCB connections.
- Interposers act similarly to substrates, but they typically can accommodate tighter spacing between metal lines, and bumps can form many more connections than balls.
- Microbumps are even smaller connections between a die and another die or an interposer.
- Depending on the interposer material (discussed below), bridges may provide connections between the components mounted to the interposer.

In addition to these basic elements, additional components for managing heat, for example, may be added.

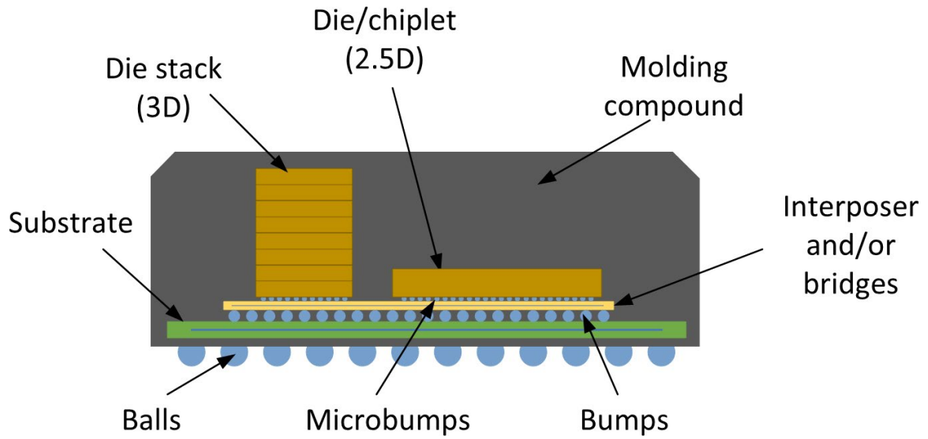


*Fig. 5: Basic advanced-package components. Like all packages, it contains a substrate. It also includes an interposer on which components are mounted and routed to each other. Microbumps connect a die to an interposer, bumps connect the interposer to the substrate, and balls connect the substrate to the PCB.*

## 2D, 2.5D, and 3D

When integrating multiple components in an advanced package, the industry has developed a notion of dimensionality that isn't strictly accurate but helps to describe how things are arranged in the package. It's unusual to see 2D, but that would refer to the default arrangement in a standard package, where the chip is placed directly on the substrate. 2.5D involves multiple chips laid out on an interposer. That interposer lies above the main substrate, so there's a bit of vertical going on — a half-dimension's worth.

3D refers to the literal stacking of components atop each other. The best example of this today is HBM, which consists of several memory chips stacked to look like one large memory. But AMD, for example, also uses the approach to place its V-Cache above a compute die, and more heterogeneous examples like this are expected in the future.

Real-world designs tend to feature a combination of 2.5D and 3D, often referred to as 3.5D. For example, one might have one or more compute chiplets and one or more I/O chiplets arranged next to each other along with HBM stacks. Figure 5 and Figure 6 illustrate such a combination.

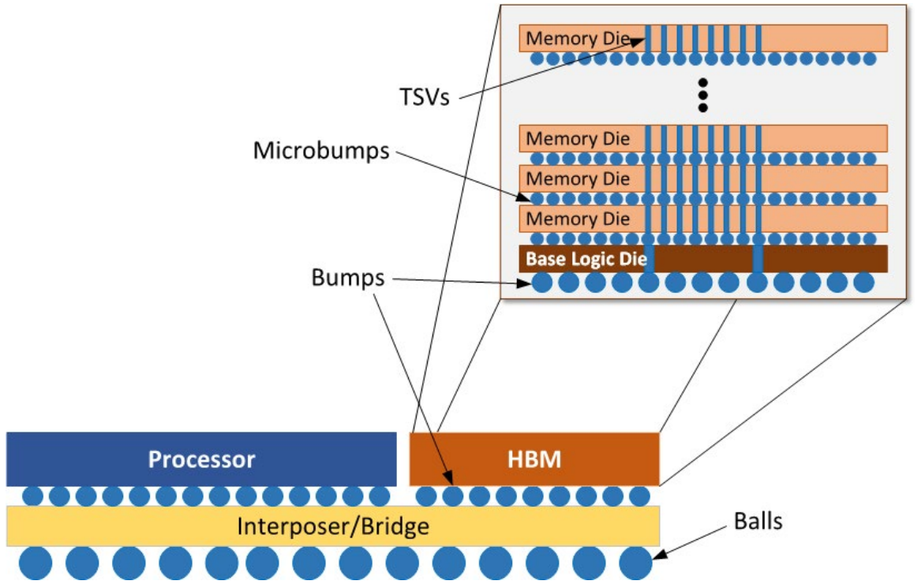


*Fig. 6: HBM as a 3D stack in an advanced package. A single HBM unit contains multiple individual thinned-down memory dies that communicate through small-pitch microbumps. TSVs carry signals from the top die down to the base die. As shown, the processor is next to the HBM in a 2.5D configuration.*



# Package substrates

The substrates used in packages resemble PCBs, although they're most similar to high-density-interconnect (HDI) PCBs. Like PCBs, they consist of alternating layers of a dielectric and metal. In this manner, they also resemble the back-end-of-line (BEOL) of a semiconductor, which also has alternating layers of dielectric and metal. But substrates feature an organic dielectric rather than an oxide. The differences between PCBs and substrates are more about dimensions than materials.

Substrates typically start with a core, which is a rigid layer of organic dielectric with copper on both sides. Construction is *additive*, meaning that the substrate is built by adding materials — in this case, building up additional layers of dielectric and metal. This leads to the other descriptive name for the process: a *build-up* process.

The metal layers act in two capacities. The obvious one is to route signals from a die connection inside the package to a connection for soldering onto a PCB. Depending on how those signals route, one or more layers may be necessary. Vias provide connections between metal layers, with three types:

- *Through* (or *through-hole*) *vias* are accessible on both sides of the finished substrate.
- *Blind vias* are accessible only on one side, terminating on an interior layer.
- *Buried vias* start and end on internal layers and aren't accessible from outside the substrate.

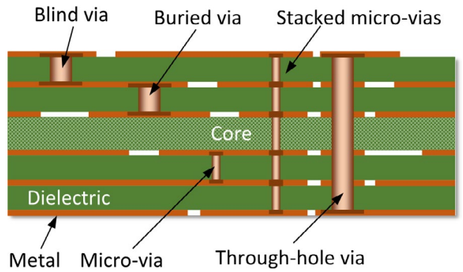


*Fig. 7: A cross-section of a package substrate. Much like a PCB, it consists of layers of metal separated by layers of organic dielectric. Vias of different types allow connections between metal layers.*

Micro-vias are simply vias with smaller diameters — below 150 µm. These allow greater via density, but they're harder to build, requiring laser drilling and higher precision. With a narrower "barrel," the aspect ratio (ratio of height to width) must also be considered since plating into high-aspect-ratio holes is more difficult.

The other role that metal layers provide is as power and ground planes. They primarily contribute to power stability for the component(s) in the package. But for high-frequency signals or high-performance circuitry, where noise is the enemy, the

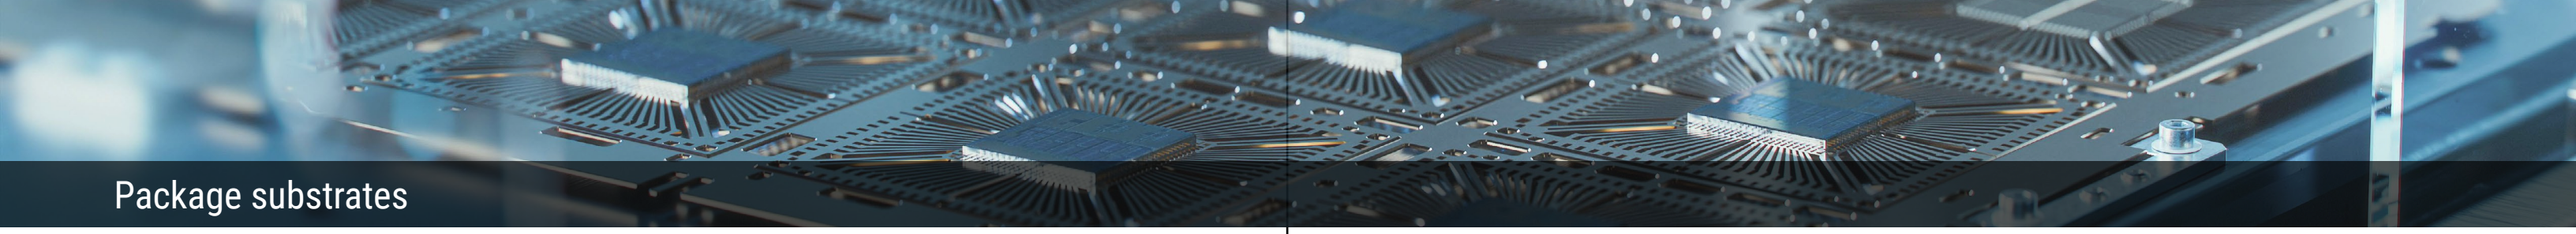



planes act as shields that keep metal layers from interacting through the dielectric. Ground lines also may be routed on a single layer between signals to mitigate any crosstalk between the signals.

On PCBs, such layers are also essential for creating strip- or microstrip-lines with controlled impedances. Substrates' smaller dimensions make such structures less common because there will be fewer signals acting as transmission lines. A 6 GHz signal, for example, has a wavelength of around 50 mm, so traces longer than about 25 mm (a half-wavelength) bear consideration as transmission lines when carrying signals of that frequency. Only the largest packages have dimensions in that range, and few signals do. Careful routing can keep the traces low enough. But if controlled impedances are necessary, the ground planes can assist with that.

## Material choices

Substrates have two main materials that remain after processing, the dielectric and the metal. The metal is overwhelmingly copper, with solder making connections. In a lead-free environment, SAC (tin/aluminum/copper) solder dominates.

In contrast with the metal, the dielectric provides more choice. The two most common ones are different forms of epoxy resin that are curable by heat (also known as *thermoset*).

FR-4 (also written FR4) is by far the best-known resin for PCBs. The "FR" stands for flame-retardant; the  4 was assigned by the National Electrical Manufacturers Association (NEMA). It's a composite, consisting of a fiberglass cloth impregnated with epoxy.

For higher performance, BT epoxy (which stands for bismaleimide-triazine) is available. It's more tolerant of high heat as measured by the *glass temperature* ($T_g$), the temperature at which the resin starts to reflow and lose structural integrity. It also provides a lower dielectric constant, which helps protect against inter-layer signal crosstalk.

Both of these materials are available in a form called *prepreg*, which stands for *pre-impregnated*. A fabric matrix is impregnated with resin and partially cured to stabilize it. Prepreg thus makes a convenient layer that can be laid down and then fully cured after it's in place. Once all the layers are in place, both heat and pressure provide thorough curing and bonding between layers.

A more recent material from a company called Ajinomoto is *ABF* (*Ajinomoto build-up film*). It provides improved dielectric and thermal properties for high-performance signals. It comes in a roll, encapsulated between an ortho-phenylphenol (OPP) film on one side, which is removed prior to application, and a polyethelene (PET) film on the other side, which is removed after application. The sheet of dielectric may come with a layer of copper.

In addition to its lower dielectric constant, it also has a coefficient of thermal expansion (CTE) closer to that of copper and other materials on the substrate. As a result, repeated thermal cycles are less likely to lead to cracks and other defects. It's helpful for packaging circuits that must demonstrate high reliability. The cost is higher, however, than simpler materials and processes.

Note that an entire substrate need not consist of the same dielectric. Different layers can feature different resins, depending on the needs of the signals they support.

Vias in PCBs are traditionally made by mechanical drilling, but the smaller dimensions of substrate vias make laser-drilling more common. Unlike holes drilled for through-hole pins, where solder wicks down into the hole, or for mechanical attachment, which needs no material in the hole, vias must conduct between layers. That's most commonly handled with electroplating, where a small amount of copper in the holes acts as a seed layer, and then copper is deposited in a bath with the board electrically connected as a cathode to attract the copper.

When drilling holes mechanically or by laser, nearby resin tends to melt, causing a "smear." For substrates with four or more layers, a *desmearing* process is necessary to clean up the surface. That process can be done either chemically or using plasma. The latter produces a cleaner, more uniform result but costs more.

## Building a substrate

The steps for building a substrate are straightforward conceptually — start with a core and add layers, patterning and drilling as you go. Buried vias and micro-vias can go on any layer, blind vias go on an outside layer, and through-hole vias are drilled once all layers are in place. A more detailed recipe would look like:

1. Start with a core metallized on both sides.
2. Drill and plate vias:
   a. Drill mechanically or with a laser.
   b. Desmear and clean.
   c. Apply copper seed.
   d. Electroplate.
3. Pattern the metal:
   a. Apply photoresist.
   b. Expose pattern.
   c. Remove developed photoresist.
   d. Etch copper.
   e. Remove all remaining photoresist and clean.
4. If more layers are needed, add another layer of resin and copper.
5. Repeat steps 2 and 3 until all layers are in place.
6. Bond layers together using pressure and heat.

The steps and materials shown are for the most common types of substrate. Other materials may be used for the core, such as ceramic or metal. Other special-purpose resins may be used, as well. When making choices on the materials to use, cost, reliability, thermal management, signal integrity, and power integrity must be balanced according to the needs of the application.

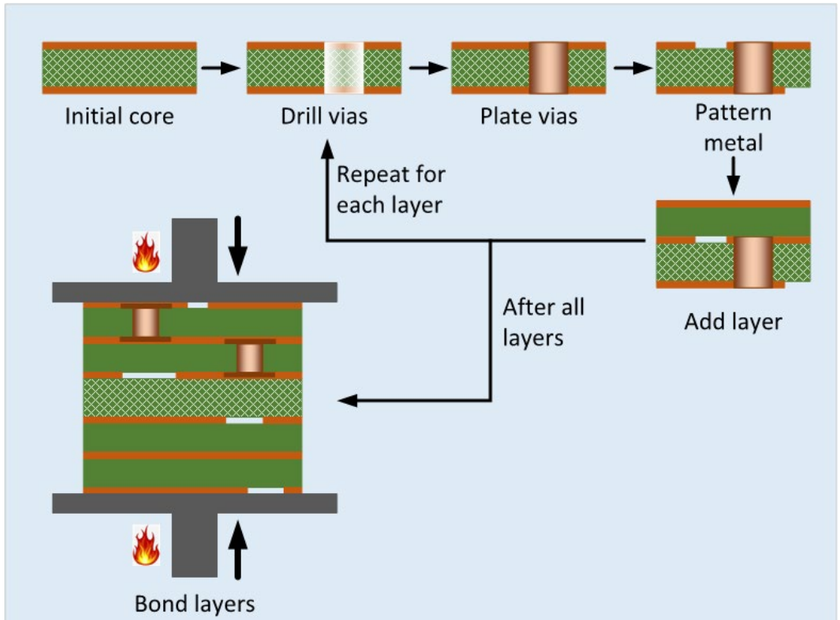


*Fig. 8: Process for manufacturing a typical substrate. Starting with a core, layers are successively added, drilled, and patterned until the layers are all in place. The final substrate layers are bonded together using heat and pressure.*



# Interposers

Substrates have been standard for packages for years, but they primarily provided a surface on which to mount a die. Today, substrates that also reroute signals are well established. In theory, one could place more than one chip on a substrate, but in practice, the resulting substrate would be too large if the number of die-to-die connections were too high. And some of today's chips have thousands of connections. It's impractical to route that many signals on a standard organic substrate in a reasonable (or manufacturable) size. In addition, signal routes are likely to be long and circuitous, increasing the power needed to communicate with high reliability.

This has been a primary motivator of the use of interposers. Technically, an interposer is any kind of intermediary or shim that makes or redirects connections between something on one side and something else on the other. In this case, silicon chips, passives, and other components lie on one side and the substrate lies below. Connections from the chips are through microbumps; connections to the substrate are through C4 bumps (both discussed below). This type of interposer is called passive. Silicon interposers also make possible active interposers that contain transistors.

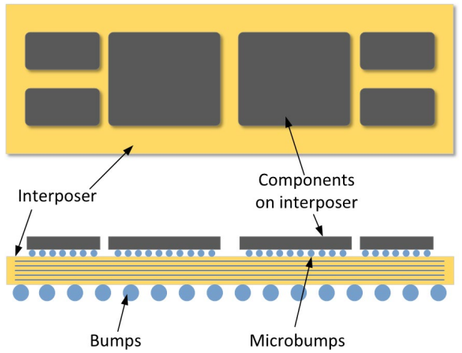


*Fig. 9: An interposer, top view and edge view. Components lie atop the interposer connected through microbumps. Layers within the interposer route signals between the components and to bumps for connection to a package substrate below.*

## Interposers vs. chips, packages, and PCBs

Interposers create an additional layer of connectivity hierarchy. At the top of that hierarchy is the die itself with its on-die connections. Below that is the interposer, which interconnects the components inside the package. Below that is the substrate, which makes connections between signals that need to be visible outside the package and the package pads. The final level is the PCB itself.

Each of those levels has different line/space and pad densities, with chips having the highest density and PCBs the lowest.

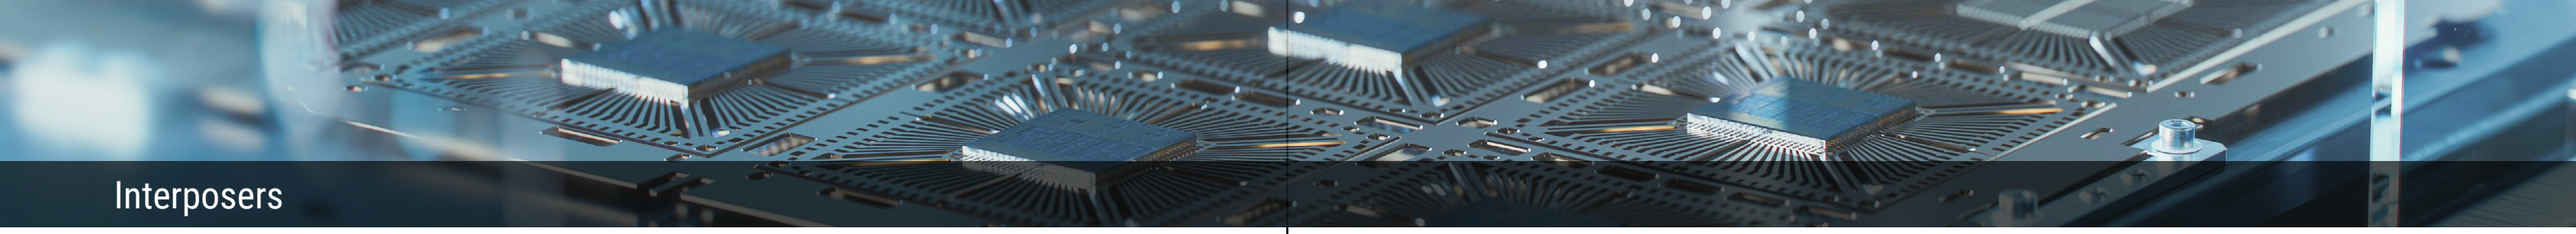

# Interposers

Interposers bring four fundamental benefits over package substrates and PCBs.

- The tighter metal and pad dimensions allow for many more signals to be routed either between in-package components or to package balls.
- More inter-component connections mean fewer signals leave the package.
- Die-to-die connections travel shorter distances and therefore have less signal-quality degradation.
- The shorter distances mean signal drivers can reduce the energy needed to drive the signal and the voltage swing, both of which reduce system power.

The downsides of using interposers are their cost, thermal considerations, and design complexity. Cost depends on the material used, but obviously having an interposer costs more than not having one. Still, the cost comparison is likely more favorable if taking into account what the cost would be using multiple packages instead of integrating into a single package.

Interposers don't introduce new thermal issues themselves, but packages using them are putting more silicon into a single package, which always has the potential to introduce heat challenges. Those issues can be addressed by ensuring, for example, that two high-power silicon chips aren't stacked atop each other or even placed side-by-side.

One classic challenge is placing HBM memory, which like all DRAM is highly sensitive to heat, as close as possible to the processor chip using it. The connections need to be as short as possible, and yet that brings the memory chips closer to the hot processor, potentially compromising memory performance.

Both cost and thermal concerns relate in general to a third concern — complexity. Cost, thermal, and any other issues typically can be addressed through careful design. But that design brings package, interposers, and chips together into one big co-design effort with many moving parts.

| | Die | Interposer/ Bridge | Substrate | PCB |
|---|---|---|---|---|
| Min. metal pitch (µm) | 0.02 – 0.04 | 4 | 50 | 250 |
| Typical number of metal layers | 9 | 4 | 8 | 15 |

*Table 1: Connectivity comparison between chips, interposers, package substrates, and PCBs. Dies are the densest, PCBs are the least dense. Adding layers raises routability but adds cost and can reduce signal integrity due to additional vias.*

## Different interposer materials

Interposers are primarily a means to route signals. The material used to build them therefore depends less on their electrical characteristics and more on their physical nature. Critical parameters include isolation between signals, thermal conductivity, and CTE as compared to that of the silicon above and the substrate below.

### Silicon interposers

The most common interposer material is silicon. The idea is that silicon manufacturing, depending on which process node is being used, allows far greater wire density than the organic materials typically used for PCBs and package substrates. Silicon interposers are thus made in semiconductor fabs, with TSMC currently being the biggest manufacturer of them.

Leading-edge silicon processes aren't necessary for interposers; they tend to remain on nodes such as 65 nm or 45 nm. That keeps their cost below what would be necessary for a leading node, but it's still more expensive per area than would be necessary for organic-material build-up. Adding to the cost, a single interposer must be large enough for all of the silicon it will carry, making it larger than the sum of all of the silicon sitting atop it. So cost per square micron is lower than that of an advanced die, but an interposer will have more square microns than a typical die.

Building a passive silicon interposer is like building a die, but using only the metal layers. As such, many layers can be created. But each layer adds cost, so minimizing layers while ensuring adequate signal routability with good signal quality is a design optimization challenge.

One frequent component on a silicon interposer is the through-silicon via (TSV), which routes a signal from one side straight through to the other. This is more typical for power and ground pins, but it can be used for signals, as well. TSVs have "keep-out" regions around them, where TSV creation may affect the neighboring silicon. A passive interposer never uses silicon's semiconducting properties, however, acting only as a medium for connections. So the preponderance of TSVs doesn't practically limit passive-interposer layout. It does, however, greatly affect the cost. Using thinner silicon can lower that cost because the TSVs can be shallower, but carrier wafers — wafers whose sole purpose is to act as a more substantial holder for a thin wafer — are required beyond a certain thinness to help maintain structural integrity

throughout the build process. The carriers are removed when finished.

A typical piece of silicon being manufactured is limited in size. For most dies, the limit is set by the size of the mask-holding apparatus, called a *reticle*. Most dies are far smaller than a reticle, and masks can include multiple dies within one reticle to improve wafer throughput. Other high-performance chips push to the limit of the reticle size.

A very few active chips exceed the reticle size, the most obvious case being Cerebras with an entire wafer acting as a single "chip." But silicon interposers can also exceed the reticle limit, although TSMC currently limits them to three reticles in size.

During photolithography steps, where patterns are exposed through masks onto the wafer, each reticle is a single exposure. The machine holding the reticle, called a *scanner*, repeatedly exposes the wafer, with the platen holding the wafer moving one reticle distance after each exposure. Over time, the entire wafer moves past the reticle.

In most cases, each exposure creates an independent chip. But for interposers (or very large chips), multiple exposures are required for a single chip. That means that the boundaries between the exposures must somehow be *stitched* together. That's easier to do at the relaxed dimensions used for interposers, but it's still a critical part of the manufacturing process that a silicon fab must perfect. Means of building larger interposers are under development. If they prove successful, they would eliminate the need for stitching.

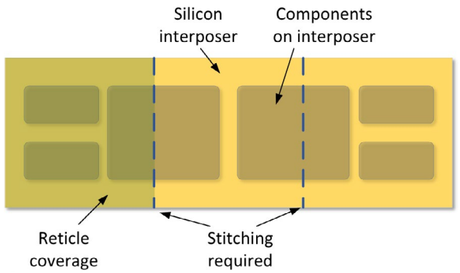


*Fig. 10: Multi-reticle interposer. When patterning the interposer, three different exposures are required in this example. Where the exposures meet, special care must be taken to ensure that any signals crossing the boundary are stitched together.*

## Glass interposers

The cost of silicon interposers has motivated the use of glass interposers. The manufacturing process for glass is very different from that of silicon, and it introduces some limitations. But it brings benefits — including better signal isolation — for some designs.

"Glass" is a broad term, and it includes many variations with different properties based on additives in the glass. Many of these are kept as trade secrets. Corning's process for producing strong glass used in cell phones and other mobile devices serves well for interposers, too. Its large-scale manufacturing in large sizes means that both wafers and panels are available for use as interposers.

The two primary features to build are the through-vias, here called *through-glass vias*, or *TGVs*, and the metal connections. Creating and filling vias is well established, and copper can be plated onto the glass. Glass interposers are still largely being researched. None is in high-volume production today.

## Organic interposers

The high cost of silicon interposers has developers moving in yet another direction, this time to organic interposers. These are largely the same as PCBs and package substrates, except that the features are much smaller. Metallization leverages equipment used for silicon rather than for PCBs, since the latter aren't capable of the necessary dimensions.

The fabrication of organic substrates is still in early days, and it's seen some production. But it hasn't yet taken over from silicon. Ultimately, if glass and organic interposers take off, the demand for silicon interposers should fall to only those designs that require the tightest dimensions or ones requiring active interposers.

## Active interposers

All three materials discussed are vying for roles as passive interposers, serving only to make connections. But silicon is a semiconductor, and it's possible to build circuits directly into the interposer itself, making it an active interposer.

No such approach is yet in production, but discussions include putting power management and I/O circuits near their respective signals in the interposer. This will increase the cost of the interposer, because it now needs front-end-of-line (FEOL) processing as well as BEOL.

Given the older process nodes used for interposers, these won't be leading-edge high-performance circuits, but rather ones that can either remove some circuitry from the chips atop it, or remove an entire chip outright. Depending on routing density, it may be that the circuits add no area to the interposer, so the increased material cost should be limited to the FEOL processing. But full costing will also rise somewhat due to, for example, the need for more extensive testing to ensure a known-good interposer.

# Silicon bridges

Silicon interposers have applications limited to those that can recover the advanced packaging cost. But the cost of a silicon component is famously related to the area of the component. And silicon interposers are very large compared with typical chips. Silicon bridges take the interposer concept and cut it down so that it uses a few small pieces of silicon rather than one big one.

Instead of using a silicon interposer, bridges are embedded into an organic interposer or substrate. The manufacturing flow has bridge manufacturers sending bridges to the interposer or substrate makers, who do the embedding. The completed interposer or substrate then goes to the packaging houses that perform the assembly.

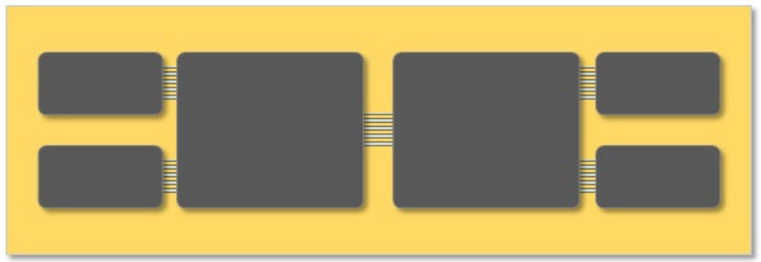

Silicon Interposer

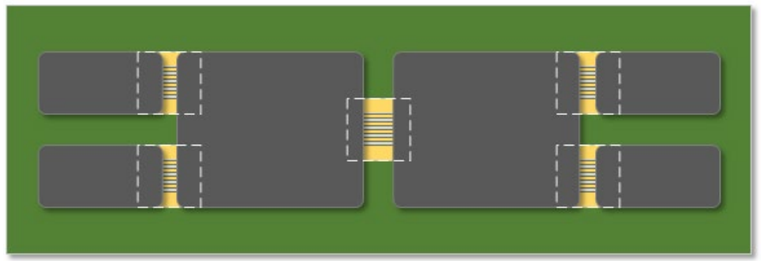

Organic Interposer with Silicon Bridges

*Fig. 11: Silicon interposer vs. silicon bridge. A silicon interposer uses a large silicon area, whereas a bridge puts silicon only where the interconnecting signals are.*



A silicon bridge is a very simple silicon die that requires only BEOL processing. That said, they're proprietary, and the manufacturing details aren't public. Intel's version is probably best known, called an Embedded Multi-Die Interconnect Bridge (EMIB). Amkor, ASE Group, Samsung, and imec also have been working on bridges.

Embedding the bridge in the interposer requires:

1. Building up the layers of the interposer up to the final layer.
2. Before encapsulation, creating cavities in the substrate where the bridge will go. Intel has some patents on cavity creation to improve cost and turnaround as compared to typical laser ablation.
3. The bridge is placed in the cavity and held there by adhesive. Alignment is critical.
4. The final substrate layers are built and other typical subsequent operations such as drilling are performed.

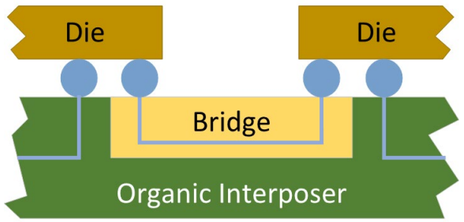


*Fig. 12: Cross-section of a silicon bridge. The bridge is embedded into the package substrate.*

The silicon technology used for building the bridge is capable of very fine lines. The limits on precision often are not determined by the bridge itself, but by the alignment of the bridge within the cavity. The tolerances on machines placing such components tends to be much looser than the line spacing on the bridge. Planarity also can limit the bridge size.

# Bonding

Bonding refers here to the attachment of a die to a substrate or attaching one substrate to another (including PCBs, package substrates, and interposers), as well as the signal connections. Many techniques are available for creating these bonds, and their details exceed the scope of this discussion. The goal here is to provide an overview of the different bonding techniques, highlighting those used more often for advanced packages.

For traditional packaging, die attach and making signal connections are two separate steps. With more recent packaging techniques, such as flip-chip, the signal connections become the die bond, although underfill materials improve mechanical and thermal stability.




*Fig. 13: The primary bonding techniques. With wire bonding, the die and signals are bonded separately. For the remaining techniques, the signal connections form the die connection as well.*

## Wire bonding

By far the most prevalent bonding techniques involve some type of adhesive between the die and the substrate. A eutectic bond may be formed to improve both stability and thermal conductivity. A eutectic bond is formed when an intermediate layer of a metal alloy is placed between the die and substrate and heated. The "eutectic" nature of the material means that the two metals in the alloy individually have higher melting points than the combined alloy, and once it melts, it does so all together rather than having some intermediate phase when there are solid and melted parts mixed together.

The wire bonds themselves can be attached using several techniques. *Wedge bonding* uses pressure to force the wire down onto the pad, squishing it into a wedge shape. It's directional in that the wedge must be aligned with the direction of the wire to its other end. *Ball bonding* lacks that requirement, making it faster and easier. In that case, a wire protrudes from the bonding equipment and is heated briefly at the end, which causes the end of the wire to form a ball that can then be placed on the pad. Any of these bonding techniques may involve a combination of pressure, heat, and ultrasonic vibration to soften the wire, scrub the pad, and form a solid, reliable connection.



*Fig. 14: 3D die stacking using wire bonds. This is a less expensive way to stack chips, but it requires the upper chips to be smaller than the lower ones.*

Although there's no formal definition of "advanced" packaging, wire bonding isn't generally considered an advanced technique. Early, less-expensive versions of die stacking still may use it as long as each die is smaller than the one it rests on so that the underlying die's pads are exposed.

Although wire bonding is less expensive, it limits the number of I/Os available based on clearances necessary for creating the wire bonds, where a machine attaches each wire to a pad on either end of the wire. It therefore can't be used in applications requiring high communication bandwidth.

## C4 balls and bumps

For higher connection density, particularly in BGA packages, *flip-chip* assembly has become the norm. It's so-named because, unlike with wire bonding, a die is flipped over with the active layer close to the substrate. Instead of wires forming the connections, solder balls serve that function.

After a die has completed processing, solder balls are formed on the die pads. The substrate pads may receive some flux, and then the inverted die is aligned and placed so that the balls land on the pads. The temperature is briefly raised in a *reflow* step, which causes the balls to partially melt and bond the die to the substrate. Because of the carefully engineered way this process proceeds, such connections are called *controlled-collapse chip connections*, or *C4*.

This technique can apply at several levels and with connections of varying size and density. At the bottom of the BGA package, balls form the connection to a PCB. Inside the package, bumps form connections between the die and substrate. Those bumps are smaller than the balls on the outside of the package. Finally, for 3D stacking, where one die stacks atop another, even smaller bumps called microbumps are employed, made possible by the finer lines and spaces that silicon technology permits.

Immediately after a die is bonded, only the metal connections provide mechanical adhesion, and this can result in reliability issues as temperature and other challenges may cause cracks or outright breaks in the solder. To stabilize things, an underfill material is applied after bonding to fill the gap. The material wicks under the die, and it helps with CTE mismatches and the dissipation of heat from the die onto the substrate.

## Thermocompression bonding

Standard flip-chip bonding is inexpensive and quick, but it has some drawbacks. Because reflow is performed in an oven, the entire board heats up, and thermal mismatch issues may weaken bonds or cause warping as everything cools back down. If the die or board isn't perfectly flat, then some bonds may be weak. In addition, metals such as aluminum form oxides, which must be breached to get a good connection.

One solution to this is thermocompression bonding (TCB), which applies heat and pressure chip-by-chip from the top. It can be used to bond multiple dies in a stack or to bond a package to a board. In the latter case, instead of reflowing by heating the entire board, only the chip and its balls are heated, which eliminates the warpage issue. The applied pressure helps to ensure a reliable bond, breaking through any oxides and forcing surface compliance between chip and board to combat any warpage. It's typically done with copper and aluminum but also can be performed using gold.

HBM uses thermocompression bonding extensively to bond the die stack. In addition to addressing the above issues, it also reduces the gap between the dies in the stack, resulting in a shorter stack. It also helps dissipate heat better than standard microbumps can.

The tradeoff is that this isn't a batch procedure in the way a reflow oven is. Instead of bonding a tray-full of chips all at once, the bonding tool — which is also more expensive than the tools used for microbumps — bonds each chip individually. The reduction in throughput makes this a more expensive process that's better tolerated with high-margin devices.

## Pillars

Microbumps cannot be made arbitrarily small. One issue is that, despite the controlled nature of the solder collapse during reflow, the exact profile of the finished connection isn't well controlled, limiting how close together they can be without interfering with each other. The other challenge is the fact that the bump size also determines the gap between the die and the substrate, sometimes called the *stand-off*. If the bumps are too small, then that gap will be too narrow to admit the underfill material.

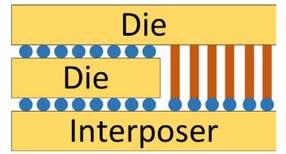


*Fig. 15: A die partially overlapping another die to which its attached can theoretically use bumps and pillars to manage the two different standoffs. This would challenge the achievable height of pillars with good yield. Die, pillar, and bump sizes not shown to scale.*

Pillars have stepped in to provide better-controlled pitch and stand-off control. Unlike a sphere, a cylinder can have independent heights and diameters, providing two degrees of freedom. Pillars and balls can even be used together on one die if it partially overlaps another die, requiring the short reach of a bump to the underlying die and then longer pillars (sometimes called columns) to reach the interposer where there is no underlying die.

The process of building pillars resembles that for building bumps. The difference is the step where the copper pillar is added, as Figure 16 shows.

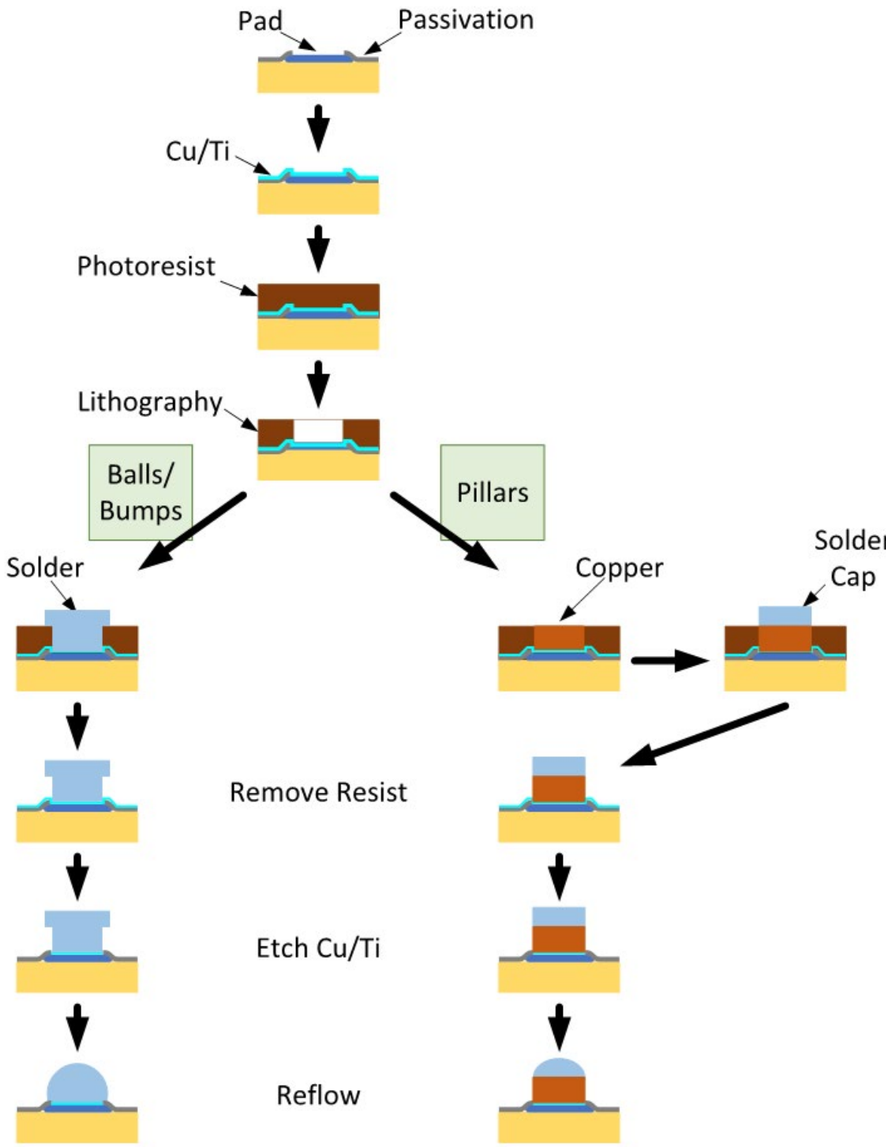


*Fig. 16: Steps to create solder balls or bumps and copper pillars. The steps are mostly the same. It's primarily the materials that change, with balls/bumps simply reflowing solder into a ball while pillars do the same with solder atop the copper pillar.*

## Hybrid bonding

The latest hot topic in advanced packaging is hybrid bonding, used primarily for die-to-die connections, including those made prior to wafer dicing when bonding wafers to wafers or dies to wafers. Instead of adding material like solder to form the connection, pads and the surrounding oxide are brought into intimate contact, making a connection with no intervening materials. The "hybrid" name refers to the fact that it's both oxide and metal forming the bond.

The metal pads are recessed slightly, causing the oxide to bond first, with the metal pads following suit. It's intended to improve the connection quality and electrical performance by eliminating the solder. Only the pad materials from the two chips make the connection. The bonded oxide provides mechanical strength.

It's a challenging process in practice, however, because all pads must be co-planar, and surface preparation is key to a reliable connection. It's been used in a limited number of applications such as newer flash memory and some image sensors, but it hasn't achieved widespread use and is subject to much ongoing research and development.

Each of these interconnect technologies permits a different size connection (e.g., ball diameter) and pitch. Most have a range between what's in widespread production and the leading edge.

| | Wirebond pad | Ball | Bump | Microbump | Hybrid |
|---|---|---|---|---|---|
| Diameter (µm) | 25 – 50 | 250 – 800 | 50 – 100 | 30 | 0.3 |
| Pitch (µm) | 50 – 100 | 300 – 500 | 50 – 150 | 10 – 50 | 0.4 – 1 |

*Table 2: Comparison of interconnect sizes and pitches. The low end tends to reflect advanced processes that may or may not be in high-volume production.*

# Assembly processes

Compared with silicon manufacturing processes, packaging is much less regimented. A given silicon node available from a foundry (or the manufacturing division of an integrated device manufacturer, or IDM, such as Intel or Samsung) is typically a fixed formula. For the most part, everything made on that process will go through the same sequence of steps.

Packaging is, at least at present, more flexible. Some manufacturers have a few well-known processes, but each company capable of such packaging is likely to have its own version of those same processes. For example, Amkor's process named HDFO is roughly equivalent to TSMC's CoWoS-R. Just as a given silicon node will differ in its details from foundry to foundry, the assembly steps will vary at different outsourced assembly and test (OSAT) houses.

It's also a time of rapid change, and the industry has not yet settled into neat, tidy processes. Each customer may want something slightly different, and manufacturers are accommodating them as much as possible. This report will review some of the better-known branded processes from TSMC and Intel, but they don't represent all of what's available or possible.

For a given process, several key parameters will vary. They include the number of layers available in a substrate, interposer, or RDL, the maximum size interposer (sometimes expressed in multiples of a reticle), and the bonding pitch. The latter is a function of the type of bond used as well as the capabilities of the manufacturer.

## Basic flip-chip packaging

Although single-chip packaging isn't the focus of this eBook, advanced techniques largely derive from basic flip-chip techniques, so understanding that process will make the other ones easier to grasp.



Illustrated in Figure 17, a die with balls attached is placed face down on the substrate following the deposition of solder flux onto the substrate. A reflow step melts the solder to form a tight connection, after which the flux is removed. Underfill then closes up any gaps between the die and the package for better mechanical stability. A final curing step completes the process.

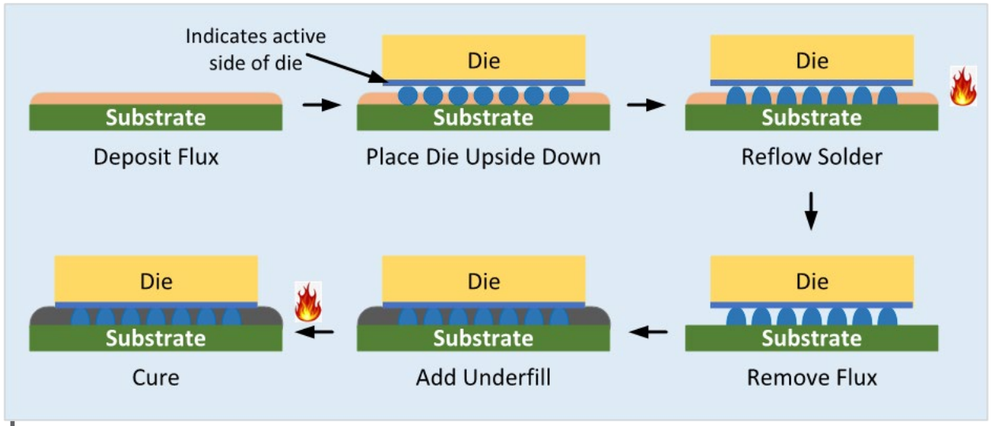


*Fig. 17: Flip-chip packaging. A bumped die is placed face-down on the package substrate. The solder is reflowed, underfill is added for mechanical stability, and then the whole unit is cured.*

## Package-on-package assembly

One of the simpler approaches to 3D assembly is to stack chips that are already in packages. Such approaches are often named with *PoP* (for package-on-package) in the brand. One specific application of this is for placing a DRAM chip above a logic chip. This is one version of what TSMC brands InFO.

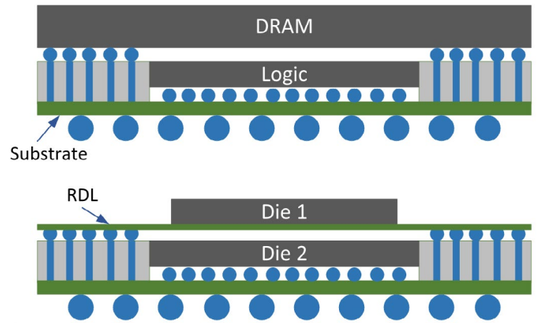


*Fig. 18: Package-over-package assembly. If the top die isn't larger than the bottom, then an RDL may be necessary. The top-die connections reach the board or bottom die through vias, being routed to the appropriate place using an RDL if necessary.*

## Chip on wafer

One of the earliest means of connecting dies in a package is to use a wafer as a carrier on which an RDL can be built, a technique branded CoW by TSMC. Two examples are shown below, illustrating two possible approaches to the assembly process.

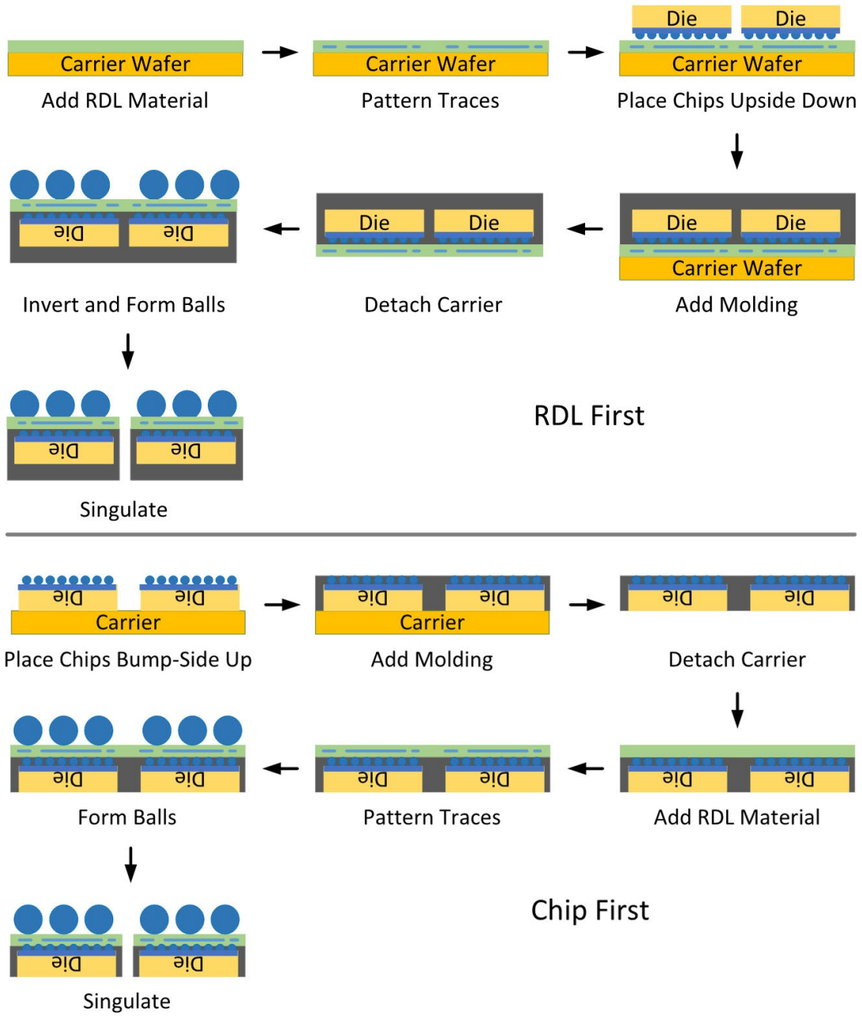


*Fig. 19: Two ways of implementing a chip-on-wafer process. In the top example, the RDL is built before the dies are placed; in the other, the RDL is built after the dies are in place.*

The first builds the RDL on the carrier wafer before placing the pre-singulated dies, bumps-down, onto the carrier. There they can be over-molded into what's effectively a reconstituted wafer. At that point, the carrier wafer can be removed, balls created, and the wafer singulated.

The other approach places the dies upside down on the carrier wafer before the over-molding. Once the carrier is removed, the RDL is built, balls are formed, and then finally the reconstituted wafer is singulated.

Intel's Foveros process is another variant intended for bonding two dies (or one die and an active interposer) face-to-face. The bottom die will be facing up, so it connects to the substrate using TSVs.

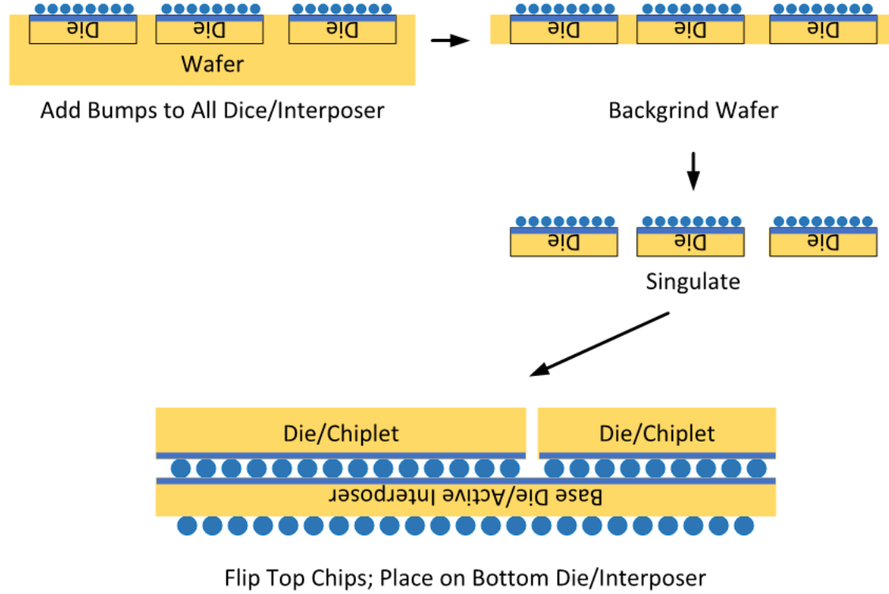


*Fig. 20: Intel's Foveros process. It connects dies or chiplets face-to-face.*

## Adding interposers

The prior approaches added only RDLs to route signals to balls. Adding interposers to the mix creates yet more routing flexibility. One well-known example from TSMC, CoWoS, has three variants, depending on the nature of the interposer. CoWoS-S is for silicon interposers; CoWoS-R implements an organic RDL; and CoWoS-L employs small chips whose function is to provide routing. The latter resembles a silicon bridge except that it also can include vias down to the substrate.

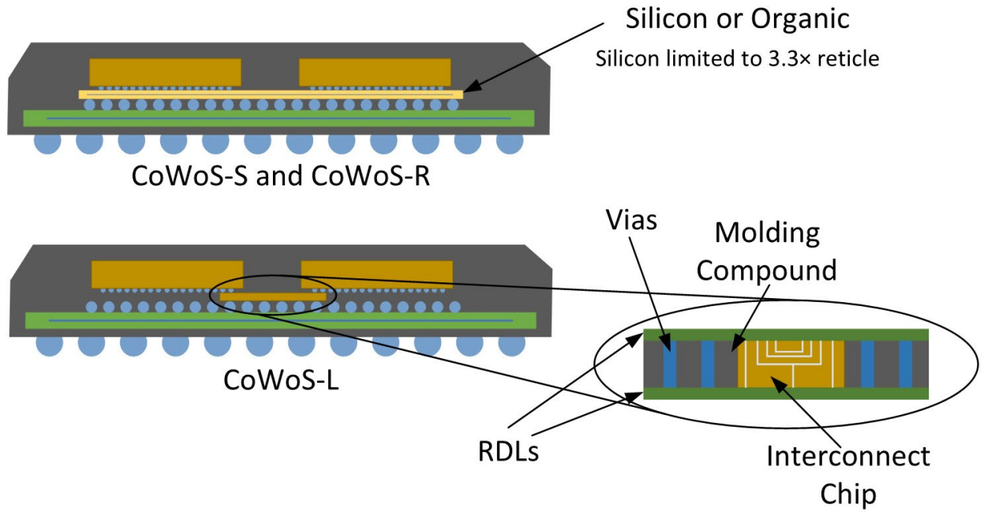


*Fig. 21: Packages with substrates. Processes such as CoWoS-S use silicon for the interposer; those similar to CoWoS-R employ organic interposers. The CoWoS-L approach includes an interconnect chip that resembles a silicon bridge.*



# Passives, optical, MEMS, and other components

The focus of this report so far has been on integrating multiple silicon dies into a package. But other components also can be included, with passives being the most common.

Passives include capacitors (the most common, used for decoupling to reduce noise), resistors, and inductors. The latter are likely only in packages that include radio-frequency (RF) capabilities. Resistors are less common, and can serve for signal termination.

Modern resistors and capacitors are exceedingly small, making it possible to embed them into organic interposers and substrates. Capacitor modules from companies such as Saras make possible the integration of a network or collection of capacitors with a single component, rather than dozens or hundreds of individual capacitors.

Optical and MEMS components are more typically mounted atop an interposer or substrate. It's possible to do so in the same manner what's done with another die, but alignment may be more critical.

Optical components typically involve ports for fiber in the package. That transition between the fiber and the optical receiver or transmitter is important for minimizing light loss, so the angle at which the fiber enters is important. The assembly process, if done manually, fiber-by-fiber, can be slow and expensive. The use of connectors forming an array of fibers that sit into so-called *V-grooves* can simplify the process.

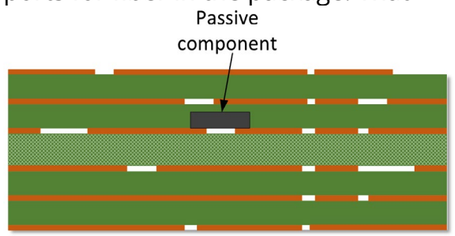


*Fig. 22: An embedded passive component in a package substrate. This would typically be a resistor or capacitor added during the build-up process.*

Some MEMS components also have important alignment considerations. Earlier accelerometers, for instance, typically handled only one dimension, meaning that three would be required for covering all three degrees of freedom (*x*, *y*, and *z*). Ideally, those three would need to be carefully aligned to be oriented precisely 90° from each other. Depending on the unit, some may have the capability of calibrating out small orientation errors.

Modern accelerometers (as well as other navigational units such as gyroscopes and magnetometers) integrate all three dimensions into a single die, where orientation is guaranteed by design. That makes the orientation of the die itself less of an issue.

# Thermal considerations

One of the functions of a package is to dissipate the heat generate by the die inside. Given the prevalence of inexpensive plastic packaging, which isn't a great thermal conductor, this duty hasn't put a strain on package design. But as more components are added, and as some of them increase in power, wicking heat away becomes critical. This is an issue with HBM today, and the challenges facing efforts to grow HBM's capacity include dealing with yet more heat to dissipate.

Package design therefore necessarily includes thermal analysis to identify whether the package will adequately remove heat without leaving any hot spots. Thermal analysis must now be performed for the package as a whole, including the components, to ensure the chip's ability to remain within its target power so that it can perform as specified.

If the leads, interposers, bridges, substrates, and molding compound are insufficient for maintaining proper temperatures across operating conditions, then the package may need to include components whose function lies only in moving heat around.

Examples of such components are heat sinks, heat spreaders, and thermal pipes. The former are pieces of metal (or any thermally conductive material) affixed externally to the top of the package. The second is similar, but it's embedded in the package. Spreaders smooth out internal hot spots, moving heat from where it's in abundance to areas with less heat, helping the package dissipate it.

Whereas heat sinks and spreaders are attached to the package on the side opposite the leads, thermal pipes use connection balls to dissipate heat. Although all signals and balls carry some heat out of the package along with their electrical functions, thermal pipes have no electrical function. Their only role is to form a connection between hot portions of the package contents and the PCB.

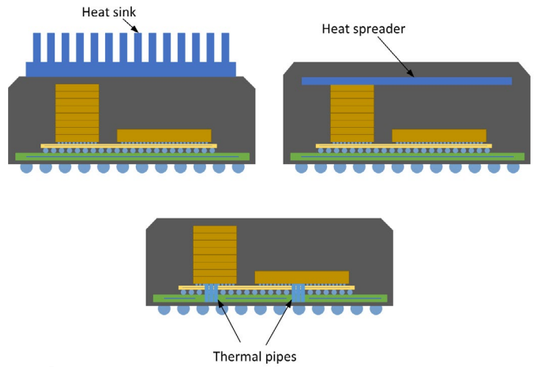


*Fig. 23: Thermal mitigation options. Extra inert metal structures may be necessary to provide adequate heat dissipation. Choices include, but are not limited to, heat sinks, heat spreaders, and thermal pipes.*



# Design implications

Advanced packaging challenges the traditional ways chips and their packages are designed. Those processes historically involved two separate groups, the chip designers and the packages designers. The former did the electronic design, while the latter focused more on the mechanical aspects of the housing. Being separate groups, the chip plans were largely tossed over to the packaging folks, and in the end, the chip went into the package.

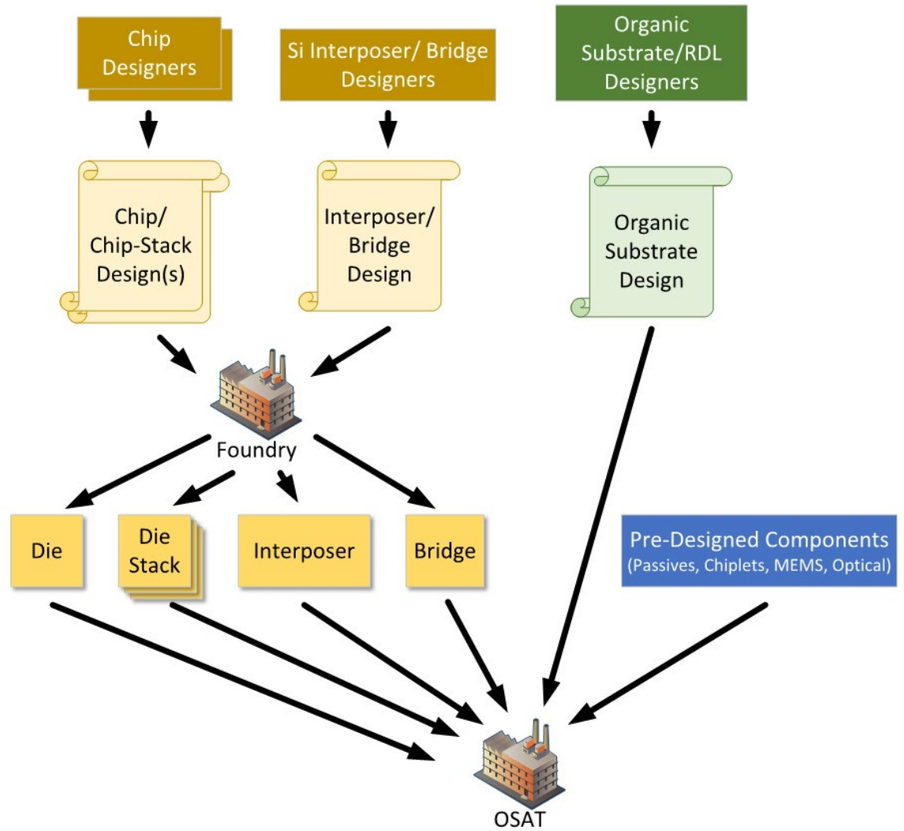


*Fig. 24: A simplified view of the supply chain for advanced packaging. One or more dies are designed and fabricated in a foundry. Silicon and glass interposers also typically require foundries. Organic elements typically come from the packaging houses. Additional components also may be included at assembly. With advanced packaging, the lines between foundries and OSATs are blurring.*

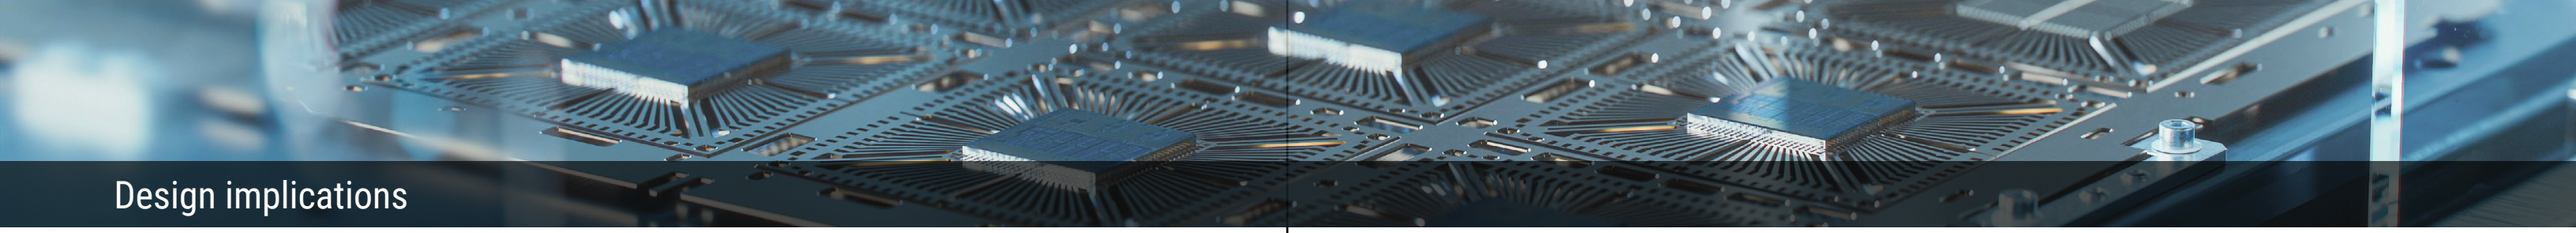

## Design implications

With advanced packages, such an arm's-length collaboration is insufficient. All stakeholders must participate early in the planning and design optimization process. Components that will co-reside in the package will come from a wide variety of sources. They include chip designers, designers of interposers or other substrates, package designers, and even makers of off-the-shelf items such as passives, MEMS, optical, or other electronic chips that are included in a given package.

Each of those roles has a specific set of tasks to perform. The items that silicon designers must attend to include:

- Meeting performance targets
- Meeting power targets
- Determining any chiplet partitioning and whether the pieces should lay out side-by-side or be stacked
- Place-and-route
- TSV placement
- Bump/microbump/pillar placement
- Power integrity
- Signal integrity
- Reliability
- Mechanical integrity, including thermal, stress, and co-planarity

Glass and silicon interposers require silicon-like design and fabrication, while organic interposers involve teams similar to those creating PCBs.

Regardless of the style of interposer or bridge, designers must pay attention to:

- Die and passive placement
- Routing to bumps
- Interposer TSVs (or more generally, through-interposer vias, or TIVs)
- Mechanical integrity, including thermal, stress, and co-planarity
- Reliability — particularly electromigration and IR drops (EMIR)

Package designers must include the following in their design efforts:

- Accurate stack-up definitions
- Physical and electrical constraint-driven signal routing (die-to-die and die-to-substrate)
- Surface-mount and embedded passive placement
- Power and ground plane generation and management
- Design for assembly
- Design for manufacturability, including stress
- Design for test
- Thermal analysis and management
- System-level power
- Die-to-die signal integrity (for interface compliance)
- Extraction of package parasitics
- Reliability



Each of the components being housed — from the highest-performance die to the lowliest capacitor — will have an impact on performance, power, and/or cost. Optimizing dies, interposers, bridges, substrates, and packaging requires the early collaboration of all designers, starting from the planning stages and continuing on from there.

Silicon design teams will create the principal chips or chiplets. That process may result in the partitioning of a single die into multiples. Those partitioned dies may lay next to each other as chiplets communicating laterally or they may be stacked, with signals flowing through TSVs.

Stacked chips can be simulated together in a chip-only environment, but chiplets arranged side-by-side must communicate through a substrate. That substrate could be the package substrate, but it will more likely be an interposer. In either case, the substrate or interposer will have an impact on performance and power. Simulating performance must include the effect of that interconnect. Passives will impact signal and power integrity. The placement and routing of signals to these passives will also impact performance.

Perhaps most critically, the arrangement of components in the package must be able to dissipate the heat that the components will generate. The operating junction temperature will impact the allowable performance — such as the maximum clock speed — and thus must be considered during silicon design.

A traditional flow could be viewed as serial, with package design following die design, or as a parallel effort. But in the latter case, the two designs — chip and package — come together only at the end. By contrast, the flow for advanced packaging requires

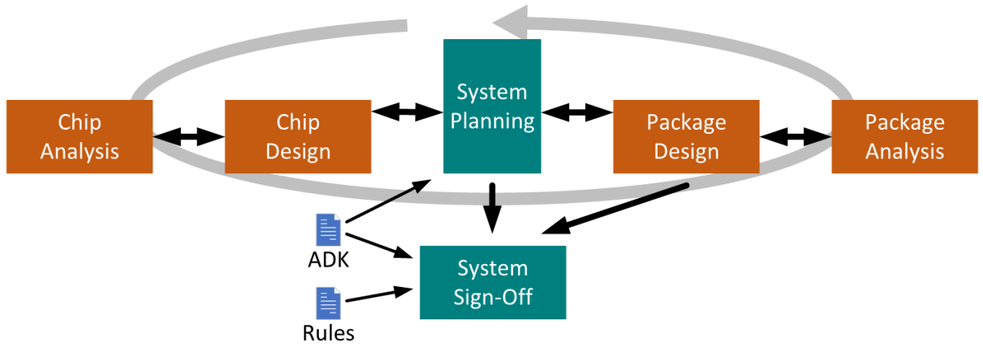


*Fig. 25: Advanced-packaging design flow. All components must be verified in parallel, communicating data between tools so that the entire system can be designed and optimized together.*

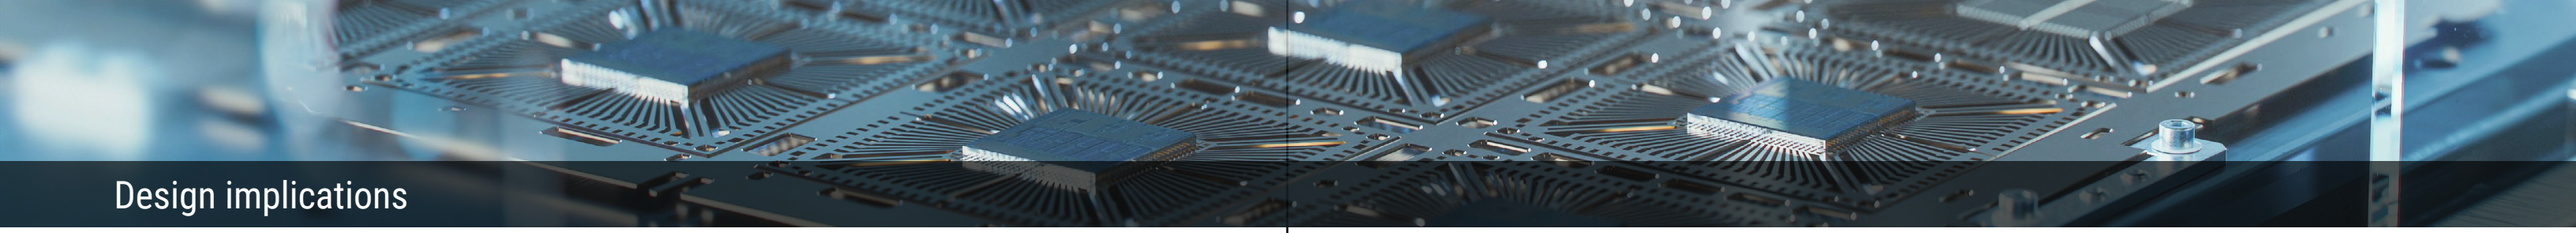



not only parallel design efforts, but constant communication between tools so that the effects of decisions on one team can be communicated to the other teams. Over time, initial estimates are replaced by simulated values as the design converges.

## Two very different scales

It would be ideal if system sign-off directly involved all the design inputs from chips, interposers, and packages. But the dimensions involved in silicon and packages are different by three orders of magnitude, with silicon working in nanometers and packages working in microns (or larger). Simply using one tool to verify at both scales would be extremely time-consuming.

What's more typical is that the chip-design data will be abstracted and fed to the sign-off engine by the system-planning tool. That's why the package design block above directly feeds the sign-off block, whereas the chip design block doesn't.

Meanwhile, the chip design goes through its own independent sign-off process that leads to tape-out. The system sign-off tools don't have the resolution necessary to bless the chip-design data.

## From Wild West to standardization

Advanced packaging creates an incredible number of options for designers — so many that each project is likely to evolve in ways different from prior projects. Variables include the number of dies, if and how they're divided and interconnected, where they're placed, other components, the interposer material, interposer vs. bridge vs. a combination of both, and materials to help with thermal issues just to name the obvious ones.

Silicon processes are also quite complex, and process design kits, or PDKs, have long been available as a way of informing electronic-design automation (EDA) tools of the many details associated with a given process. No such standard format exists for advanced packaging, although work is underway to establish assembly design kits (ADKs). Those will be more complex than PDKs owing to the number of considerations they must encompass, including:

- Technology files that specify details such as how layers stack up, the materials used and their properties and thickness, any physical or electrical layout constraints including line and space dimensions, special signals such as differential pairs, and any custom design-rule checks (DRCs) necessary to validate the design.
- Libraries specifying physical footprints and models for power and thermal behavior for all components including chiplets, passives, interposers, vias, routing between dies, and mechanical characteristics.
- Assembly rules that respect the constraints that a given set of pick-and-place tools will require, including spacing between devices, distances between devices and other elements or the edge of the package, and the maximum allowed stack height.
- Any electrical specifications with which signals must comply, including libraries of interconnects and I/Os, eye masks, jitter tolerance, and insertion or return losses.
- Manufacturing rules laying out checks for substrates, solder masks, soldering, and silkscreen patterns.

As the number of process variations evolve from numerous ones that cater to specific projects into a few well-accepted standard processes, ADKs will be critical enablers of further automation to ensure that the vast number of constraints have been met and that the design will function as intended.

# Test considerations

Testing chips in advanced packaging is fundamentally the same as testing a single die in a package, but the logistics are more difficult. Test circuitry and standards are all about one thing — making potential defects within a circuit controllable and observable. If you can't control certain nodes, then you can't test them thoroughly. If you can't observe the results, then even if you did manage to apply a test, you can't see the result, so it does no good. Combining multiple dies in a package makes controllability and observability more difficult.

For many years, two complementary test approaches have dominated the semiconductor business. The first is IEEE 1149.1, also known as JTAG (which stands for Joint Test Action Group — the committee that defined the standard originally). The second is so-called *design-for-test*.

JTAG enables *scan test*, which is the ability to serially scan data into test infrastructure, apply a test, and then scan the results out. The registers into which the data is scanned are dedicated to test, and the serial sequence of registers is known as the *scan chain*. The serial approach is important because, especially when the standard emerged, few pins were available for use in testing. As it is, a JTAG *test access port* (or *TAP*) consists of only four pins (with the option for a fifth reset pin).

JTAG originally was intended to test PC board connections. By loading data into each pin on one chip, the result could be detected on the connected chip, verifying the integrity of the PCB connection. The same approach allows testing chiplets on a substrate in a package.

But given the desire for testing the internals of a chip after packaging, companies ran scan chains internal to the chips, as well. In fact, at the time, internal test may have been more common than external test.

Driving scan chains through package pins for internal testing worked in the early days of testing, when the

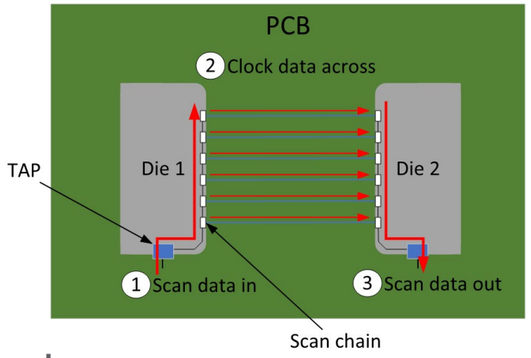


*Fig. 26: A basic JTAG scan chain testing board connections. Registers on the left side are loaded with data and then clocked across to the right side where the data is captured and scanned back out. On the test access port (TAP), only one signal is shown for simplicity: the data-in signal on the left and the data-out signal on the right.*



targeted faults were simple stuck-at faults. But as integration levels increased and new fault models were introduced, more efficient ways of testing became necessary. This was the era of design-for-test, or *DFT*, which involves automatic test-pattern generation (ATPG) and compression.

EDA companies developed technology that, when generating test patterns at design time, takes a mass of test input data and compresses it to speed up the time necessary to load it when testing. An on-chip circuit decompresses the test inputs, sending them across a dedicated test network. Rather than scanning out individual bits, the results are compressed into a much smaller signature, which then can be scanned out and compared to the expected result. The increased uptake of such testing technologies created a need for a universal mechanism to set up, configure, and control the test circuits in a manner similar to JTAG. This resulted in a new standard, IEEE 1687, informally known as internal JTAG, or IJTAG.

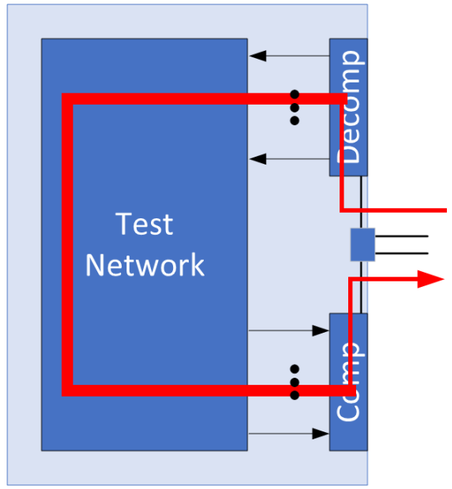


*Fig. 27: Internal chip test. Compressed test stimulus data is scanned in and decompressed into the test network. The results of the tests are then compressed into a small signature and scanned out for verification.*

## Standard modifications for special cases

Two special cases required revisions to these two standards. JTAG is a static DC test. It therefore cannot test signals that are AC coupled. AC coupling allows high-speed signals to be impedance matched, but capacitors lie between driver and line and between line and receiver. Whereas DC-coupled lines communicate via voltage levels, AC-coupled lines communicate via transitions that can pass through the capacitors. The benefit is the lack of a DC component to the current as well as the ability to cross voltage domains.

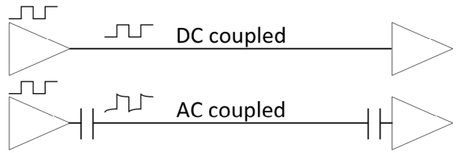


*Fig. 28: DC- vs. AC-coupled lines. Capacitors on AC-coupled lines eliminate DC current from the connection, but require signal transitions in order to pass through the capacitors.*

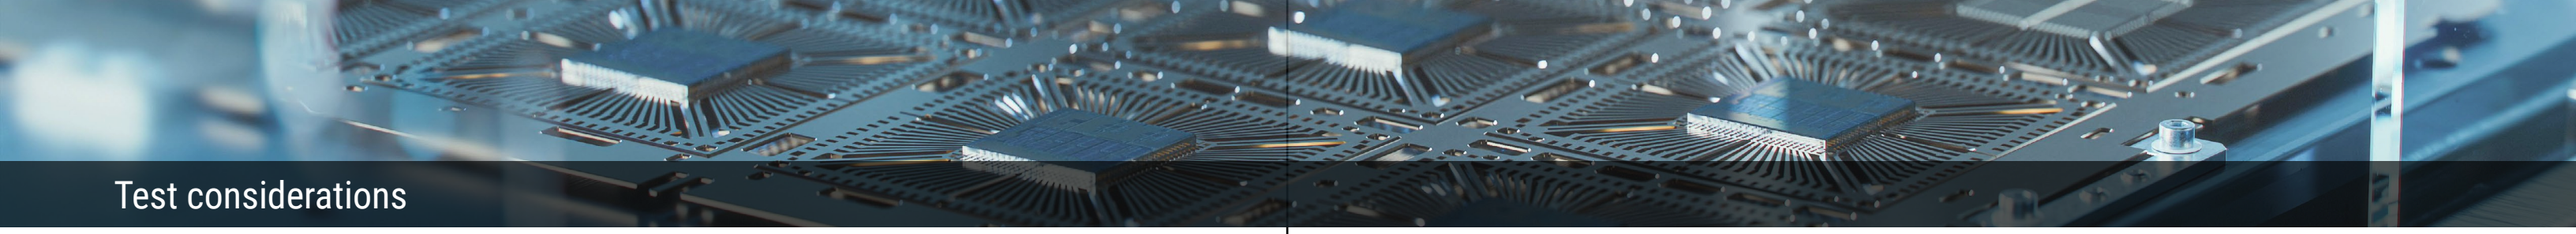

## Test considerations

IEEE 1149.6 provides a means of testing AC-coupled signals. It's complementary to 1149.1 and can reside on the same scan chain.

The internal test standard, meanwhile, works for digital logic, but not for analog blocks. The standard is in the process of being augmented to handle analog circuits. Currently called IEEE P1687.2 (The P indicates a work in progress), it will complement IEEE 1687. It allows the digitizing of the result of key parameters when compared to reference values. Each analog sub-block (essentially, some analog function) can have its own associated test block, or one test block can handle multiple sub-blocks, multiplexing analog signals, and reference values.

Some regular circuits, such as memories, can be outfitted with circuits that run tests internally, without the need for external test stimulus. Called *built-in self-test*, or *BIST*, such circuits can simplify the remaining test circuitry. Such BIST circuits can still be controlled by JTAG externally, and that's how manufacturing test proceeds. But they are particularly useful for systems requiring occasional in-field tests, such as in vehicles, where they're run by an internal JTAG controller rather than external JTAG pins.

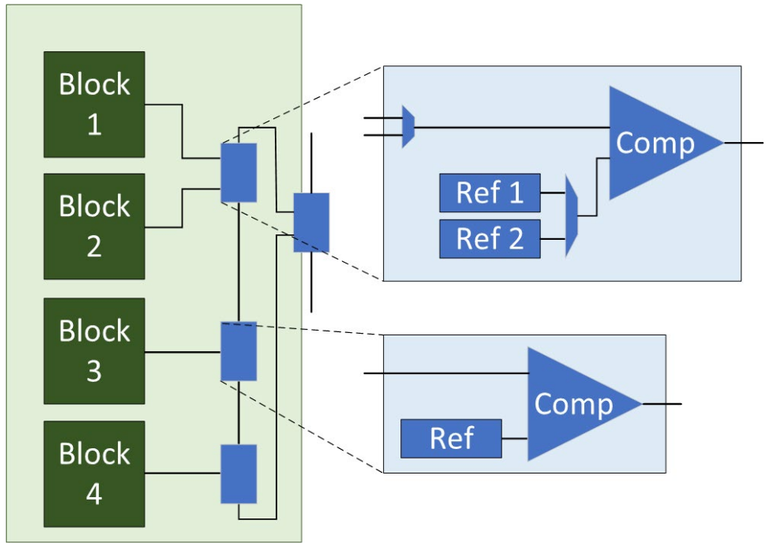


*Fig. 29: An example set of analog tests. The test blocks fit onto the scan chain, but they consist of comparisons of signal values to references. The specific tests are highly dependent on the analog function being performed. A test block can multiplex tests for multiple functions or each function can have its own test block.*

The final capability necessary for generating tests for an advanced package is software that can combine individual chip and other component tests into a single unified scan chain.

Figure 30 illustrates an example of a two-die-plus-HBM combination in a single package. The HBM stack can use memory BIST (MBIST) plus scan

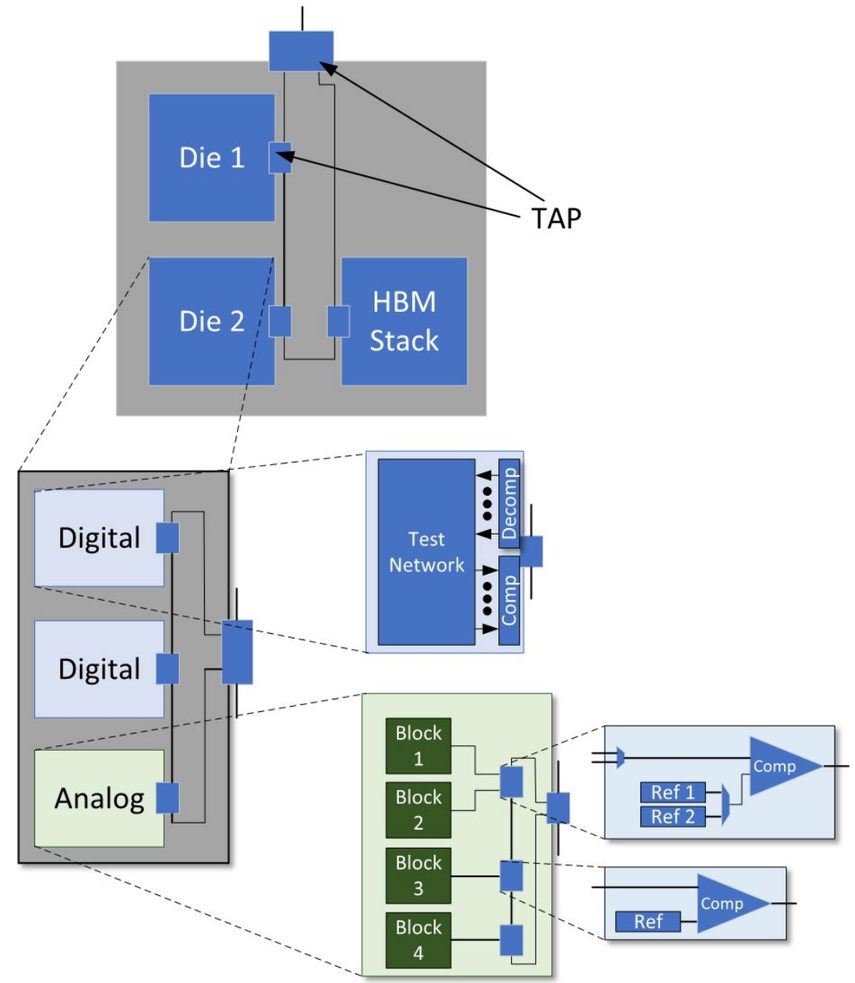


*Fig. 30: An example package with two dies and an HBM stack. One die has two digital blocks and one analog one. The digital block is shown as tested via IEEE 1687; the analog block is tested via a future IEEE 1687.2. The HBM stack can have its logic tested via JTAG, with MBIST testing the memory cells.*

testing. The other dies can use IEEE 1687 (or 1687.2) to test their internals.

In theory, this entire sub-system can be tested through a single TAP, although additional TAPs are possible for parallel testing. For the latter case, another standard, IEEE 1838, specifies how the multiple controllers are configured and interconnected, identifying a primary TAP (PTAP) and secondary TAPs (STAPs). IEEE 1838 specifically targets stacked dies, each of which has its own controller, but which are accessible only through the bottom die, with TSVs (typically) providing access to the upper dies.

A number of other test-related standards have been inactivated because IEEE 1149.1 has been broadly taken up and handles the situations the other standards covered. These include 1149.4 for mixed signals, IEEE 1532 for in-system programming of programmable chips, and IEEE 1581, which targeted memory chips that lacked TAPs.



# Reliability

Advanced packages share the same basic reliability concerns as standard packages, but they're amplified by the new materials and number of components being co-packaged. The largest concerns deal with three aspects — co-planarity, electromigration, and thermomechanical effects.

Co-planarity is always important for any die that has many connections, as is typical for a BGA package, for example. If not co-planar with the substrate or interposer it's being mounted on, then some of the balls may not make contact. Such a situation would be a test failure and the unit wouldn't be shipped to customers. But if the co-planarity weren't that far off, then solder balls might make poor connections on some pads — connections that could come loose after mechanical shocks or too many thermal cycles.

Warping is a concern particularly for a component with built-up layers, such as a substrate or an interposer. The different layers of materials can create internal stresses that cause bowing, so materials for these applications are carefully chosen for dimensional stability.

Electromigration has long been a concern, especially on silicon. It's a function of current density, with high currents literally pushing metal atoms around. Because silicon chips have very fine metal lines as compared with PCBs, they tend to have higher current density and therefore are more subject to migration.

Electromigration can happen anywhere that current density is too high, however, not just on a die. The whole point of using interposers and reducing the sizes of bumps is to allow higher interconnect density than is possible on a PCB. Those finer lines mean that electromigration will be more of an issue than it would be for

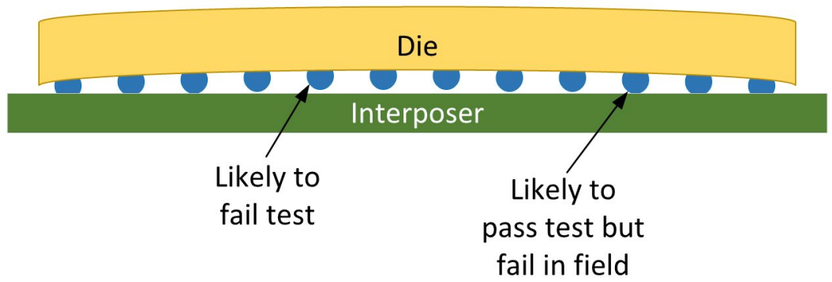


*Fig. 31: Co-planarity issues caused when one surface is warped more than the one to which it's bonded. If too far off, connections will fail outright and should be caught at test. But if a poor solder joint forms instead, it may not fail until in the field.*

a standard PCB. Analysis tools are important for identifying traces with high current density so they can be fixed prior to production.

Thermal considerations include two important aspects. The first is simply the ability to remove the heat being generated inside the chips. With multiple components, higher heat is possible than if those components were packaged separately. If that heat can't escape sufficiently, then junction temperatures will rise too high and chips will cease operating properly.

Of longer-term concern are the effects of repeated heating and cooling cycles on the assembly. As the unit heats up, different elements will expand at different rates according to their CTEs. For example, if not managed properly, the balls attached to a die may expand a different amount than the substrate and pad it's connected to, potentially breaking the connection — especially after many cycles.

The risk of such issues depends entirely on the materials used. If a silicon die is mounted onto a silicon interposer, the risk is lower because both elements are silicon. But that same die bonded onto an organic interposer may fare differently. Materials and physical arrangements are chosen to minimize the effects of CTE mismatches such as these, and some compliant materials are used to help dissipate the stresses such mismatch causes.

Although the assembly flows in commercial production today have addressed these issues to the extent possible, it's still early days for this type of assembly. So designers can't assume that all materials will be perfectly flat, that in-package metal lines can handle the current, and that the overall assembly will survive a lifetime of temperature cycles. Die and package analyses are important prior to taping out or committing to a packaging configuration to avoid possible future rework.



# Security

No discussion of electronic systems of any kind can be complete without considering security. In the semiconductor realm, that has primarily dealt with monolithic chips — especially systems-on-a-chip (SoCs), because so much valuable activity takes place on a single piece of silicon. Much effort has gone into protecting chips and the traffic between chips on boards. Their security involves not just protecting against hackers, but also against supply-chain threats that could enhance vulnerability to hacking or simply act as an economic leak as system builders buy fake components, perhaps unwittingly.

Advanced packages contain such chips plus other components — silicon or otherwise. The chips may be protected, but without extra thought, there is no unifying security that protects the entire package contents. Die protections already are well documented, but the additional considerations for advanced packages are not.

One of the important concepts for assessing vulnerability has to do with how much an attacker knows about the technology they're probing. A random hacker who has access only to the physical package will have no idea what's inside and so must play guessing games to break in. Such a hacking target acts as a black box.

At the other end of the spectrum would be an attacker who is part of the supply chain, and hence has access to design information, be it RTL (the hardware design specification) or GDSII (the physical mask data). Although it requires extensive work and sophisticated tools, much can be learned from this information, making hacking not so much a guess as a targeted strike. To this hacker, the chip is a white box.

Given the number of components in an advanced package, it's possible that a hacker may have more information about some components than others, giving a mixed black-box/white-box situation. But the package also contains more than simply the chips. The substrate, interposers, bridges, and passives all must be considered in addition to any active elements.

## Advanced-package vulnerabilities

Advanced packages suffer the same vulnerabilities that SoCs do, but the specific points of vulnerability and their impact differ. Two important considerations help characterize the nature of a specific type of attack. First, is it destructive or can (or must) it occur while the system is running? And second, does the attack occur during some stage of manufacturing and distribution, or does it occur in the field after deployment?

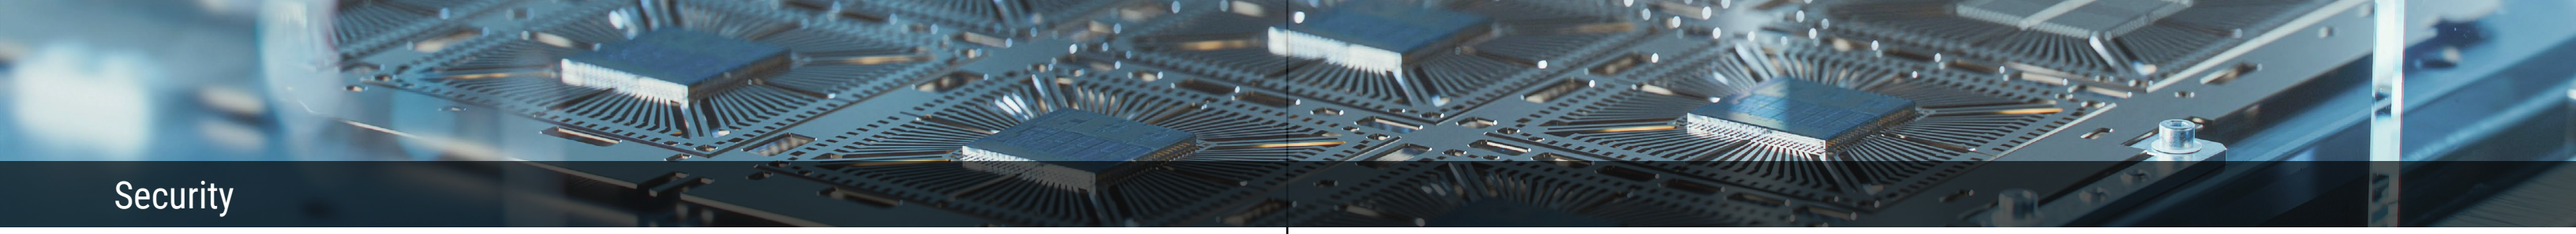



Considerations for 2.5D and 3D configurations are different. In general, multi-die stacks will be harder to probe and reverse engineer — especially in cases such as HBM, where the stack consists of chiplets all of the same size. As 3D connections evolve to hybrid bonds, which literally result in the knitting together of both the oxide and the copper on both chiplets, reverse engineering will become even more difficult due both to the greater difficulty in prying the chiplets apart and to the finer pitch and smaller pads that such technology allows. 2.5D arrangements expose more signals, and so most of the issues below deal with 2.5D.

The following are the different categories of threat and how they apply to advanced packages. Attack characteristics are denoted as D (destructive), N (non-destructive), S (supply chain), F (field), W (white box) or B (black box).

## Information leaks (N, F, W/B)

Although individual chips may be robustly protected, they communicate with each other and the outside world through interposers, bridges, redistribution layers, and the substrate. Anyone who can open the package delicately enough to not damage it can probe the inter-die connections to get information. An attacker in the supply chain might be able to do so prior to package encapsulation if there's a stage where the product can be powered up. The latter situation is less likely and more easily addressed by controlling the assembly and test flow to eliminate such opportunities. Attacks are obviously easier if the system is a white box.

## Control usurpation (N, F, W)

The prior attack merely releases information. This one allows an attacker to take control by accessing internal resources, such as registers and memory, and contaminating them to repurpose the system. This is most likely a white-box attack, although necessary information such as processor architectures can be found in trade publications, meaning that it doesn't necessarily have to be an insider performing the attack. It requires both access to signals and knowledge of how to apply them, although some amount of guessing is likely to determine finer details that haven't been published.

## Fault injection (N, F, B)

This sort of attack typically works by glitching the power in an attempt to put one or more chips into an illegal state that might release information or allow a control change. The latter works only if, after the change, the system can be returned to a legal state without power cycling (which might undo the control change). If the active chips in the package are well protected against fault-injection attacks, then the package will probably also be protected since the additional components are likely to be passive.

## Side-channel attacks (N, F, B)

The two most typical types of side-channel attack involve analysis either of power noise or electromagnetic emissions (EMI, where the I stands for interference). Both can be used to extract information (and so are a form of information leak), with the most common target being encryption keys during encryption and decryption. If such keys are unique to each device (as they should be), then analysis must be non-destructive since the keys will work only on that one unit. These attacks require a huge number of individual assaults in order to amass the amount of data necessary to statistically deduce key values, likely with the aid of AI.

## Reverse engineering (D, F, W/B)

Although some amount of reverse engineering can be done non-destructively, a thorough analysis will require careful deconstruction of the package and its components. Aside from the chips, the interconnect is the most likely target. That means interposers, bridges, and the substrate. The more advanced the technology being targeted, the more expensive the equipment necessary both to deconstruct the package and to analyze what it reveals. Delaminating built-up structures can reveal how the package components are interconnected.

## Trojan horses (N, S)

Supply-chain attacks include the possible insertion of Trojan horses at any of various design stages. A given die could have such circuitry either surreptitiously designed into the die by an attacker on the design team or it could inherit such a weakness in IP purchased for use in the die. Package-level interconnect infrastructure, especially if built out of silicon, could theoretically house active components, but a typical fabrication process wouldn't include the kinds of lithography and depositions required. More likely would be adding signals to the outside world that ought to remain inside, or re-routing the signals between components.

## Counterfeiting (N, S)

There are different opportunities in the supply chain for counterfeiting. In one case, legitimate units may be diverted through techniques such as overbuilding. Those devices will operate correctly. The impact is economic, with revenues going to the counterfeiter. In other cases, failed or marginal units might be diverted and sold, in which case the purchaser may get bad material. Finally, attempts to build counterfeit units based on reverse engineering could turn out units that work properly, with only an economic impact, or they may

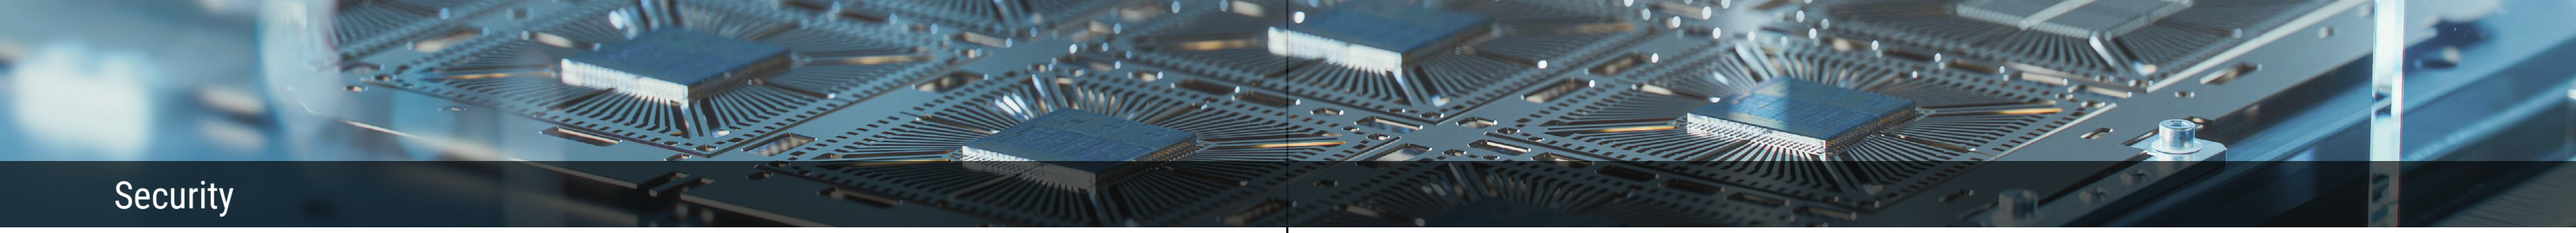



be unreliable if manufacturing and testing are sloppy or if the reverse engineering effort was only partially successful.

## Attack mitigations

The three main aspects to protecting the packaged components beyond mitigations already in place for the chiplets are the interposer, traffic in general, and side-channel vulnerability.

The interposer is vulnerable to probing and reverse engineering. Probing non-destructively is possible only if the accessible layer on top exposes critical signals. If critical signals are buried on inner layers, probing becomes more difficult. Turning probing into a successful attack also requires that the attacker know which signals are which. Reverse engineering may be necessary to make that determination unless the attacker can identify signals based on signaling patterns.

Complete reverse engineering requires delaminating the interposer layer-by-layer to trace out where the metal leads run. It may be possible to obfuscate the layout by including false routing, but ultimately, if an attacker can stack all the decoded layers graphically, they should be able to determine the intended routing from the pads connecting the chiplets. So obfuscation is likely to be only minimally helpful. Thus, keeping critical signals on inner layers may be the best defense against probing.

Silicon bridges should be harder to probe. The bridges expose pads on both sides that chips attach to. Those pads are under the chips, with the signals within the bridge being buried in the unit. Reverse engineering is possible by removing the chiplets and figuring out which signals the bridge connects. That may be an issue for a chiplet assembly of the type that's typical now, where a single company creates all of the chiplets. But in a future open chiplet market, those interfaces would be public, anyway. So bridges would appear to resist probing and, to a lesser degree, reverse engineering more effectively than interposers.

## Encrypting intra-package traffic

If probing can't be avoided, the best defense is the same one used for any network connection — encrypt the traffic. Although that's standard for some communication protocols, it may meet with resistance from chiplet designers. The idea of breaking a monolithic chip apart into different chiplets works best if chiplet-to-chiplet connections cause as little speed loss as possible. The fastest connection, as long as the topology permits, would simply be wires on the interposer.

But in the interests of standardization, connections typically use protocols that lay atop one of two physical arrangements, UCIe and Bunch of Wires (BoW). The very use of a protocol adds latency to signals crossing between chiplets. Adding encryption to that would further burden performance, potentially dramatically.

Although no standard specifies encryption for inter-chiplet communication, both UCIe and BoW are looking into security. Neither has established a standardized approach yet. Although added security always means some loss of performance, the cost for chiplets would be particularly significant. This is an area in active development, and changes in this space are likely over the next couple years.

Side-channel attacks, meanwhile, are most effectively protected at the die level, especially for power analysis. If every die has side-channel protections and the only components outside the chiplets are passive, then power analysis should be difficult.

If chiplets are protected against EMI radiation, then the package itself will have a smaller, but non-zero, EMI signature. If active countermeasures are used on all the chiplets, then the package EMI signature would be meaningless since it would consist of combined obfuscated signals from all chiplets. But if the chiplets rely only on shielding for protection, then the interposer and other connections may still radiate the signals.

Given complete chiplet designs, it's impractical to go back and tell the chiplet designers to add active EMI obfuscation. That means that signals outside the chiplets but inside the package may still radiate. Encryption would make such radiation far less meaningful, but it's unclear if encryption will become an accepted mitigation.

Adding obfuscation in the package would entail adding a component with the sole purpose of emitting confusing EMI signals, contributing to noise and consuming energy. It may work for some applications, but is unlikely to be appealing across the board.

If sensitive signals are radiating, further shielding at the package level may be necessary. That could mean an additional metal layer above the components and embedded within the substrate.

As with any discussion of security, likely attack surfaces and mitigations must be considered early in the design phase. After-the-fact stop-gap efforts are likely to be less effective — or even ineffective. Part of that calculus deals with the value of a given chip or chiplet in a system, but one can't lose sight of the fact that attacks on one component may not occur for the sake of probing that component. It may just be a way to get onto the network, and then get to some other component. So designers must look more broadly at the impact of security (or its lack) in the systems for which the chips are designed.

# A fast-moving technology

Advanced packaging remains an incredibly dynamic area. Although employed primarily on high-value chips for now, its use is growing at a rapid rate. This report presents a number of processes and options, some of them branded. Those should not be considered processes that are standardized across the industry. New ideas and techniques are the rule, not the exception, during this formative phase of the technology.

The high-level impact continues an integration trend started long ago. We've moved from systems on boards to systems on a chip, and although systems in a package aren't entirely new, they're now being embraced into the mainstream. Advanced packages are making that possible, shrinking systems and allowing more to be done in smaller spaces.

Semiconductor Engineering expects that new approaches will make appearances every year. Some will stay, others will fade as better ideas succeed them. Future updates to this eBook will look to sorting the lasting developments from the fleeting ones. The one guarantee is that from year to year, leading-edge packaging will look different.



# Recent developments

## Heterogeneous integration

### Challenges Of Heterogeneous Integration

**July 13, 2023**

Heterogeneous integration opens the door to an almost unlimited number of features in a single package, but it also adds system-level challenges into a small space filled with a whole spectrum of possible interactions. Mike Kelly, vice president of chiplets/FCBGA integration at Amkor Technology, talks about a variety of issues ranging from uneven aging, warpage, and different mechanical stresses, as well as some possible benefits.

### Smarter Systems Through Heterogeneous Integration: Highlights From 3D & Systems Summit

SEMI

**July 25, 2023**

It has taken decades of research and development and strong commitment to various industry programs, but the stars are finally aligning for 3D semiconductor systems. No one could have left the 3D & Systems Summit 2023 – held in late June in Dresden – with any doubt that heterogeneous integration, enabled by increasingly mature 3D packaging technologies, is becoming a key driver of the semiconductor industry's success.

### Experts at the Table: Heterogeneous Integration

**October 19, 2023** (Part 1)

**November 15, 2023** (Part 2)

**November 29, 2023** (Part 3)

Experts at the Table: Semiconductor Engineering sat down to discuss problems and potential solutions in heterogeneous integration with Dick Otte, president and CEO of Promex Industries; Mike Kelly, vice president of chiplets/FCBGA integration at Amkor Technology; Shekhar Kapoor, senior director of product management at Synopsys; John Park, product management group director in Cadence's Custom IC & PCB Group; and Tony Mastroianni, advanced packaging solutions director at Siemens Digital Industries Software. What follows are excerpts of that conversation. (Part 1)

### Many More Hurdles In Heterogeneous Integration

**January 18, 2024**

Advanced packaging options continue to stack up in the pursuit of "More than Moore" and higher levels of integration. It has become a place where many high-density interconnects converge, and where many new and familiar problems need to be addressed.

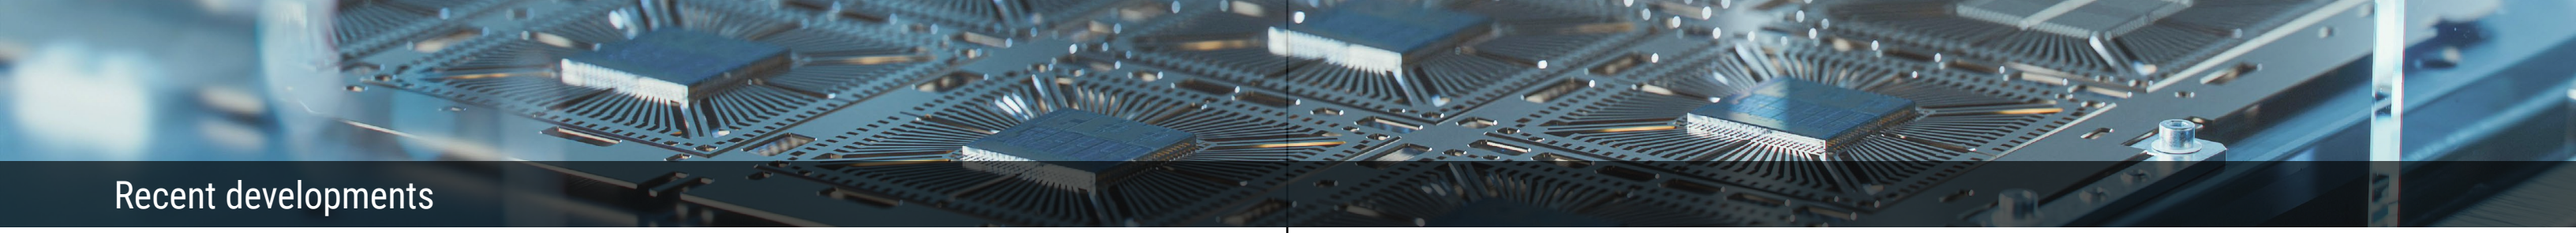



### Integration Hurdles For Analog And RF In Next-Gen Packages

**March 23, 2024**

A rapid increase in wireless connectivity and more sensors, coupled with a shift away from monolithic SoCs toward heterogeneous integration, is driving up the amount of analog/RF content in systems and changing the dynamics within a package.

### Powering The Automotive Revolution: Advanced Packaging For Next-Generation Vehicle Computing

Amkor

**April 18, 2024**

Automotive processors are rapidly adopting advanced process nodes. NXP announced the development of 5-nm automotive processors in 2020, Mobileye announced EyeQ Ultra using 5-nm technology during CES 2022, and TSMC announced its "Auto Early" 3-nm processes in 2023. In the past, the automotive industry was slow to adopt the latest semiconductor technologies due to reliability concerns and lack of a compelling need. Not anymore.

## Leadframes

### High Performance, Multi-Chip Leadframe Package With Internal Connections

Amkor

**December 19, 2023**

For high performance applications, demand for highly integrated packages has increased. This is due to the highly integrated package's electrical performance advantages of reduction of interchip distance (delay), high density I/O counts for multi-function and small form factor. With the increasing importance of highly integrated packages, the need for improved thermal management is also increasing. When the high-density I/O signals operate for the highest performance, heat generation increases on the die. The high heat generation without effective heat dissipation has adverse effects on reliability and electrical performance of electronic products.

## Bonding

### Reverse Laser Assisted Bonding (R-LAB) Technology For Chiplet Module Bonding On Substrate

Amkor

**September 21, 2023**

In the recent semiconductor market, there are many applications including smartphone, tablets, central processing units (CPUs), artificial intelligence (AI), data cloud and more that are expecting and experiencing rapid growth. As most of these applications require high performance, single-die Flip Chip packages may no longer be appropriate and multi-dies in a chiplet module packages could be the new solution. From the viewpoint of the bump interconnection methodology, mass reflow (MR), thermocompression bond (TCB) and laser assisted bonding (LAB) are widely used in the industry.

### Wirebonding Is Here To Stay

**October 19, 2023**

Few technologies in semiconductor manufacturing have stood the test of time as steadfastly as wirebonding. This process, which involves electrically connecting semiconductor devices to their packages, has been a cornerstone of the electronics industry since the beginning of the electronics industry.

### Gearing Up For Hybrid Bonding

**October 23, 2023**

Hybrid bonding is becoming the preferred approach to making heterogeneous integration work, as the semiconductor industry shifts its focus from 2D scaling to 3D scaling.

### LAB Flip Chip Reflow Process Robustness Prediction By Thermal Simulation

Amkor

**November 16, 2023**

Nowadays, there are many interconnects in IC chips. One of the packaging goals is to connect an IC to the next level of subsystem circuitry (package substrates/print circuit boards). Mass reflow (MR) of solder joints is a widely adapted and stable process in the industry. The applications of MR include flip chip, ball mounting, surface mount technology (SMT), and even reliability tests. SMT reflow ovens contain 8 to 20 heating zones and can be used for inline production. However, not only the silicon solder joints but the entire packages inside reflow ovens are heating up during the MR process. This thermal budget will increase package warpage and affect solder joint quality in fine pitch (i.e., less than 60 µm) conditions. Thermal compression with non-conductive paste (TCNCP) and laser assisted bonding (LAB) were introduced to address fine pitch devices with localized heating to reduce package warpage.

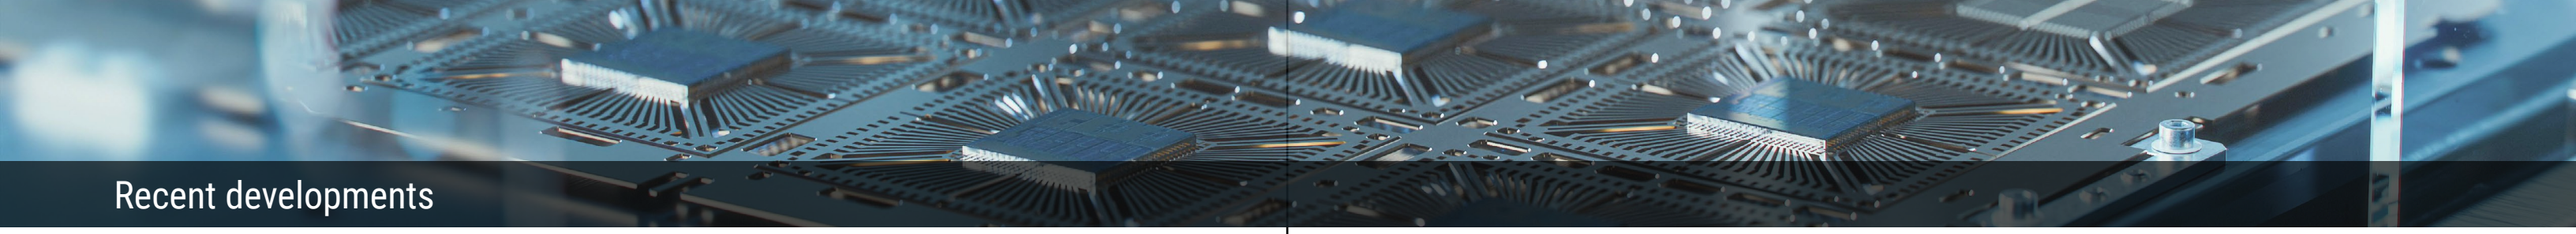



## Interposers and bridges

### Building Better Bridges In Advanced Packaging

**September 21, 2023**

The increasing challenges and rising cost of logic scaling, along with demands for an increasing number of features, are pushing more companies into advanced packaging. And while that opens up a slew of new options, it also is causing widespread confusion over what works best for different processes and technologies.

### The Race To Glass Substrates

**May 29, 2024**

The chip industry is racing to develop glass for advanced packaging, setting the stage for one of the biggest shifts in chip materials in decades — and one that will introduce a broad new set of challenges that will take years to fully resolve.

## Panels

### A Hybrid PLP Technology Based On A 650mm x 650mm Platform

Amkor

**July 25, 2023**

A panel-level (PL) approach to fan-out (FO) packaging has been discussed for several years to reduce the cost of chip-first FO packaging based on redistribution layer (RDL) technology. More recently, multilayer high-density chip-last packages have been introduced for more advanced applications. This technology would also benefit from PL processing for cost reduction. Due to the large package dimensions, applications such as an application processor (AP) or multichip module (MCM) will have greater benefits than the smaller power management integrated circuit (PMIC), transceiver or audio codec applications typical of chip-first FO packaging.

## Assembly

### Elimination Of Die-Pop Defect By Vacuum Reflow For Ultrathin Die With Warpage In Semiconductor Packaging Assembly

Siang Miang Yeo et al, Amkor, Universiti Tunku Abdul Rahman, Kajang, Malaysia

**January 18, 2024**

Semiconductor die thickness is getting thinner over time due to improvement of power efficiency in advance power electronic packages. Ultrathin die with convex warpage can easily deteriorate the solder void removal process during solder reflow, leading to various packaging reliability issues. In particular, a new type of packaging defect phenomenon—die-pop—is observed. Vacuum reflow process has been able to reduce the solder void size to minimum value consistently but there is an observable number of die-pop occurrence during the single-step pressure-profile reflow process for ultrathin dies that come with convex warpage. Nevertheless, optimization study has revealed that the application of the two-step pressure-profile and the medium reflow temperature profile within the vacuum reflow process has successfully eliminated the die-pop occurrence when packaging ultrathin dies. Together with the appropriate selection of the solder paste volume and the bond line thickness (BLT) of solder layer, majority of the samples attained the solder void size over die size of below 2% with zero die pop detected, without compromising the manufacturing productivity.

### Fan-Out Panel-Level Packaging Hurdles

**January 24, 2024**

Fan-out panel-level packaging (FOPLP) promises to significantly lower assembly costs over fan-out wafer-level packaging, providing the relevant processes for die placement, molding and redistribution layers (RDLs) formation can be scaled up with equivalent yield.

### Package Assembly Design Kits: The Future Of Advanced Package Design

Amkor

**May 23, 2024**

Why should there be an interest in Package Assembly Design Kits (PADK) today? For the most part, it is due to the advancement in the accumulation of files forming the PADK now offering a customized heterogeneous design experience that optimizes the device's intended package performance with complete connectivity verification, DRC, and assembly validation. Another primary reason is the ongoing focus on reducing the design cycle time. This is possible by beginning with clear and proven manufacturing and assembly checks that can be implemented during the initial layout phase, thereby reducing the need for rework of the design layout during post-process checking. In a blog discussing Package Assembly Kits, Paul McLellan mentions, "New challenges face both IC designers and package designers and new approaches are required." Finally, and possibly one of the most impactful reasons, is the need to maintain design integrity amid the growing demands and complexity of the high-density fan-out (HDFO) package design environment.

### Controlling Warpage In Advanced Packages

**June 24, 2024**

Warpage is becoming a serious concern in advanced packaging, where a heterogeneous mix of materials can cause uneven stress points during assembly and packaging, and under real workloads in the field.

## Thermal

### Navigating Heat In Advanced Packaging

**January 18, 2024**

The integration of multiple heterogeneous dies in a package is

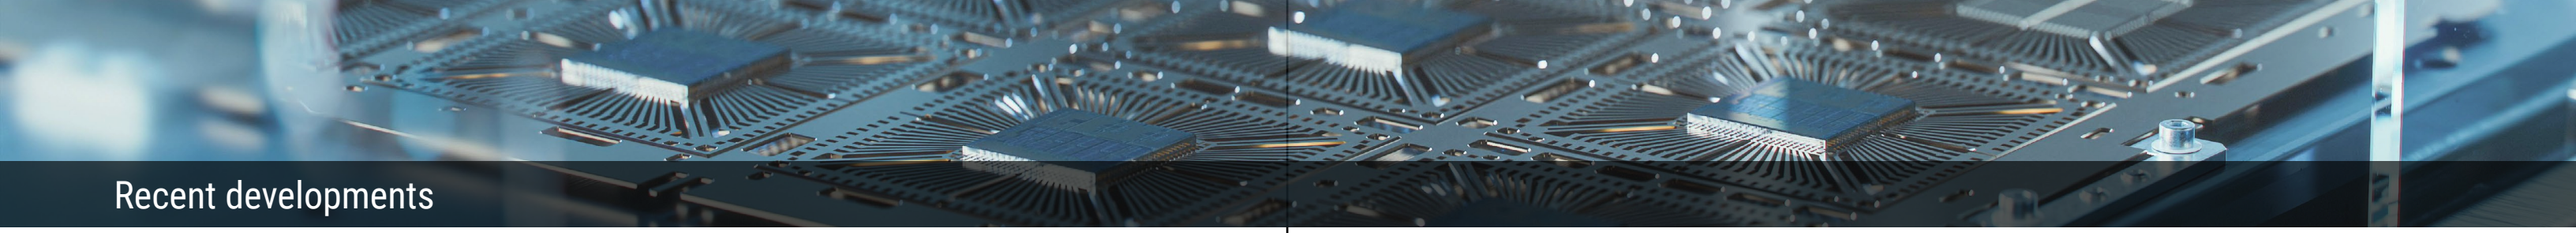



pivotal for extending Moore's Law and enhancing performance, power efficiency, and functionality, but it also is raising significant issues over how to manage the thermal load.

### Package Integrated Vapor Chamber Heat Spreaders

Amkor

**March 21, 2024**

With continuous increases in computational demand in nearly all electronics market segments, even historically lower power packaging is being driven into challenging thermal management situations. Node shrink alone is reaching a limit in maintaining track with Moore's law. The economics and yield challenges of large monolithic system on chip (SoC) designs are driving the development of silicon disaggregation or chiplet adoption. Trends in total power dissipation are driving the development of extremely low thermal resistance cooling systems. High-performance systems may implement a combination of heat pipes, vapor chambers, or liquid cooling, driving sink-to-ambient thermal resistance less than 0.5°C/W [1]. These low resistance cooling solutions translate to the package now comprising a much larger percentage of the total junction-to-ambient thermal resistance. As a result, these trends in power density and extremely low resistance cooling solutions are forcing the need for thermal enhancement closer to the silicon than ever before.

## Reliability

### Electromigration Performance Of Fine-Line Cu Redistribution Layer (RDL) For HDFO Packaging

Amkor

**January 18, 2024**

The downsizing trend of devices gives rise to continuous demands of increasing input/output (I/O) and circuit density, and these needs encourage the development of a High-Density Fan-Out (HDFO) package with fine copper (Cu) redistribution layer (RDL). For mobile and networking application with high performance, HDFO is an emerging solution because aggressive design rules can be applied to HDFO compared to the other package types such as Wafer Level Fan-Out (WLFO). HDFO allows assembly of more than one chip in one package and mostly fine Cu RDL is used to interconnect the chips. In addition, HDFO can be made in wafer and substrate level depending on the application, which has better scalability in terms of package size.

### Electromigration Concerns Grow In Advanced Packages

**April 18, 2024**

The incessant demand for more speed in chips requires forcing more energy through ever-smaller devices, increasing current density and threatening long-term chip reliability. While this problem is well understood, it's becoming more difficult to contain in leading-edge designs.

## Other

### Big Shifts In Power Electronics Packaging

**November 16, 2023**

The power semiconductor market is poised for remarkable growth in the next several years, fueled by the adoption of electric vehicles and renewable energy, but it also driving big changes in the packaging needed to protect and connect these devices.

# Standards

| Standards body | Standard | Coverage |
|---|---|---|
| IEEE | 1149.1 | JTAG |
| IEEE | 1149.4 | Mixed signals; now inactive since 1149.1 can handle |
| IEEE | 1149.6 | AC coupling |
| IEEE | 1149.7 | Enhanced JTAG |
| IEEE | 1532 | In-system programming; now inactive since 1149.1 can handle |
| IEEE | 1581 | Testing for memory ICs that lacked a JTAG TAP; now inactive due to prevalence of TAPs |
| IEEE | 1687 | Internal JTAG, or IJTAG |
| IEEE | 1838 | Multiple TAP controllers for stacked dies |