

BITS&CHIPS

Besi holds the best cards in advanced packaging, but the game has only just begun – November 1, 2023

René Raaijmakers

1 November

Die-to-wafer hybrid bonding is about to penetrate advanced chip systems for servers, data centers and later possibly cell phones. Back-end equipment builder Besi and partner Applied Materials are in the front row with their machines.


Expectations for hybrid bonding are high, and Richard Blickman, CEO of BE Semiconductor Industries (Besi), likes to drive the enthusiasm. “Hybrid bonding, that’s what you all came for,” he said when welcoming investors and analysts at Besi’s Capital Markets Day last June. The new interconnection technology should provide a new growth spurt for the Netherlands-based company in the coming years.

In technical terms, die-to-wafer hybrid bonding is a fascinating process. It precisely layers a bare, clean and very flat chip onto a wafer that has undergone the same beauty treatment. The surfaces must be clean at virtually atomic level, and then there’s the crucial operation of precisely mirroring and bonding the insulating silicon oxide and copper contacts. In a super-clean environment and with the right mechanical-chemical treatments, this creates a sandwich of integrated circuits with thermal and mechanical properties very similar to a monolithic chip.

A similar process, but in a form where two wafers are bonded together, has been used in industry for some time. Almost every cell phone contains an image sensor made with wafer-on-wafer hybrid-bonding processes. Likewise, the CMOS image sensors that provide spectacular light sensitivity in Sony’s high-end cameras are made by bonding very clean wafers together.

The equipment for die-to-wafer hybrid bonding must be very clean and precise. In part because of those requirements, it falls into the price range of front-end semicon equipment. A hybrid-bonding line that fills a 300 mm wafer with dies starts with a machine for cleaning and applying plasma activation to prepare the chips and wafers for the bonding process. In the second step, the bonder accurately places the ICs on the wafer. Applied Materials is moving into this market with Besi, which supplies the bonders. ASM PT and EV Group provide the same equipment for placement and preparation, respectively, but they don’t operate in an exclusive partnership.





A wafer with individual chips inserted, most probably NAND memories. Credit: EV Group

Growth strategy

The total cost of such a cleaning and bonding line amounts to 5-6 million euros. Applied and Besi each grab half of that, depending on the application – memories or logic.

The Dutch backend specialist is currently well-positioned in die-to-wafer hybrid bonding because of its historic relationship with Intel and TSMC. According to CEO Blickman, the Taiwanese foundry asked his company eight years ago to start developing bonders for its technology. “They’ve helped us throughout the learning curve,” he said in an interview with Pierre Ferragu of New Street Research last year. “We’re in a unique situation, with the right customers. We’ve been picking winners from the beginning. Our partnership with Applied Materials is enormously helpful in understanding the requirements in cleanroom environments.”

TSMC placed a large order for hybrid-bonding lines with Applied and Besi this year. Blickman expects a similar order for multiple machines from Intel before the end of the year.

Besi’s hybrid flagships are a central part of its current growth strategy. They’re advanced machines with high price tags, and the company is preparing to achieve sales of hundreds of millions of euros with them in the coming years. In 2021, in the semicon hype during the COVID crisis, Besi announced that both Intel and TSMC had promised to purchase fifty hybrid bonders. Orders are only really taking off this year, so these plans seem a bit delayed, but Besi says it already has the capacity to produce 180 hybrid bonders per year. If that capacity is filled, it will mean additional sales of 400 million euros – with price tags per machine of 2-2.5 million euros.

Looking at the growth target with which Besi is currently delighting investors, this is where most of the additional sales will come from. By the way, the previous target – 800 million euros per year – hasn’t been met yet. Besi held out that carrot to investors several years ago, but the current recession threw a

spanner in the works. In 2021, Besi's sales stuck at 749 million euros; this year, the company expects to sell 600 million euros worth of machines.


Thermo-compression

The wafer-on-wafer variant of hybrid bonding has been proving itself for some time in camera sensors and 3D NAND memories. Now the largest manufacturers are preparing for die-to-wafer. Intel, Samsung and TSMC are expected to be the first to deploy the technology on a large scale in the heterogeneous manufacturing processes they use to assemble complex chip systems from dozens of chips from different nodes.

Meanwhile, however, advanced chip systems such as Nvidia's AI tiles and microprocessors from AMD and Intel have been assembled with thermo-compression bonding (TCB) for years. This interconnect technology is well established and, in addition, is still developing. The question is whether, and how much, thermo-compression can slow down the hybrid-bonding train.

Kulicke & Soffa (K&S), provider of TCB bonders, sees this as a more realistic scenario and attractive opportunity. Among other things, it supplies the machines that Intel currently uses to assemble its most advanced chip systems, like Meteor Lake, which is due out next December. If K&S succeeds in reducing the space between contact points (pitch) of thermo-compression, it may offer a cheaper alternative to chip stacking than hybrid bonding.

The Singapore-based outfit is an established player in the TCB market for advanced packaging. Discussing last quarter's results, CEO Fusen Chen said his company will earn 60-68 million dollars in TCB systems this year. He expects this to grow to 100 million dollars by 2025.



Nvidia's latest GH200 Grace Hopper superchip is basically a composite of graphics processors, microprocessors and around it more than a handful of high-bandwidth memories. No hybrid bonding comes into play yet for composing this tile. Credit: Nvidia

K&S recently announced a partnership with UCLA CHIPS (Center for Heterogeneous Integration and Performance Scaling at the University of Los Angeles) to get the TCB pitch below 5 micrometers. “I believe it will be very difficult to get this to high-volume manufacturing in the next ten years”, comments Stefan Chitoraga, technology and market analyst for packaging and assembly at Yole Intelligence. In practice, TCB sits at a pitch above 30 microns. Hybrid bonding achieves 10 microns in die-to-wafer processes and under 1 micron for wafer-to-wafer in the case of 3D NAND.

In addition, K&S is working on a process without solder, using formic acid to make the copper bobbles very clean and then connecting them directly together. Paul Lindner, executive technology director at EV Group, admits that this is a sensible strategy, with a chance of success. He agrees that cheaper solutions that work usually win in the back-end. “The industry will do everything possible to hold on to and expand the established technology.”

Yole’s Chitoraga is more outspoken about the drawbacks of thermo-compression. “It’s true that TCB is less expensive than hybrid bonding, but keep in mind that you’re talking about the current status of thermo-compression machines with a pitch of 40 microns. If we go to 10 micron, which by the way is still under development, then the cost will definitely increase compared to 40 micron.”

The Yole analyst has doubts about the reliability of thermo-compression for small pitches. “If you reduce the pitch, then you also have to reduce the size of the micro-bumps. There are quite a few challenges there, because if those micro-bumps are smaller, then they’re much more sensitive to thermal stress.”

Foveros

Besi CEO Blickman isn’t the only one pepping up the buzz around hybrid bonding. Last July, Intel spoke clearly in favor of the technique at Semicon West. It envisions a future where computer systems consist of complex packages made up of many separate chips – also called cores or chiplets in the case of heterogeneous integration. “We call it a sea of cores,” said Babak Sabi, senior vice president of assembly test technology development at Intel, during a forum discussion on hybrid bonding at Semicon West.

Composing chip systems from many chiplets or cores is one way to increase performance, Sabi argues. He sees it as the way for Intel to reduce interconnection signal delays and achieve the highest bandwidth. “In chip packaging, it’s all about interconnections. If you’re talking about chip stacking with vertical interconnects, then you want to use hybrid bonding. Also to drastically reduce energy consumption. To achieve that, we need innovation beyond what we have now.”

Chitoraga of Yole says Intel is even already preparing to use hybrid bonding in the silicon carriers of its Foveros interposers. That would be a formidable step because it would greatly expand hybrid-bonding applications. Interposers are like expensive printed circuit boards made of silicon, into which active parts can be incorporated if necessary. In back-end processes, interposer chips are placed on top with advanced techniques such as TCB, but according to Yole, Intel is now looking at hybrid bonding, although it won’t completely replace the Foveros with thermo-compression.

The TCB-assembled Meteor Lake package, for example, includes four advanced chips: a graphics processor, a system chip and an I/O chip, all manufactured by TSMC, and a general-purpose processor from Intel’s own fabs. The TCB connections that connect these four chips to the Foveros interposer are 36 micrometers apart. That’s already a substantial improvement over the 55-micron bump pitch used for Lakefield, the first Foveros application.

Foveros enables 770 connections per square millimeter in Meteor Lake. In the future, the pitch will go to 25 and 18 microns. Yole’s Chitoraga says Intel may switch to hybrid bonding between logic chips and interposer as early as next year. Foveros Direct will use die-to-wafer hybrid bonding with an expected pitch of around 9-10 microns.



For several years, Austria's EV Group has supplied a system for cleaning chips and wafers and preparing them for a placement system with plasma activation. EVG has a collaboration with ASM Pacific, which supplies the bonders. Credit: EV Group.

Bloody back-end


Last year's New Street interview demonstrates that even Blickman is capable of putting things into perspective. "Until now, the world hesitated over whether this hybrid bonding was really the answer. Of course, there's also thermo-compression bonding with its pluses and minuses." While the Besi CEO also sees ASM Pacific coming up with solutions in hybrid bonding, he finely points out his company's edge. "It's ultimately about accuracy and speed. That leads to efficiency and ultimately to cost of ownership."

The Besi top executive realizes that a technological edge is quickly squandered, especially in the bloody back-end equipment segment. "The main challenge for Besi is certainly to stay ahead of this development in the next two, three, four, five years," he told New Street's Ferragu. But, he said, "the economics and financial aspects of this technology make it very interesting for anyone in the back-end industry to move in that direction."

Blickman is assured of the support of Intel and TSMC for now, but in technological terms, Besi is vulnerable. Hybrid bonding may be a knowledge-intensive and advanced technology, but the entry threshold for competitors is low when compared to front-end machines. ASM Pacific and EVG, among others, have strong positions.

Austria's EVG has already proven itself in hybrid wafer-on-wafer bonding, a market in which it's leading with hundreds of machines in the market. Just about all the sensors in cell phones have gone through EVG equipment in wafer-on-wafer processes. In the high-end CMOS sensor market, EVG is competing with Japanese TEL.

Chinese manufacturer Yangtze Memory Technologies Corp (YMTC) is already using the Austrian equipment to stack 3D NAND memories. These are the memories Apple wanted to include in its iPhones, but it refrained from doing so after mounting geopolitical pressure and criticism from US policymakers.



Chip manufacturers also expect to use hybrid-bonding processes in the coming years to package high-bandwidth memory memories. Credit: SK Hynix

Cost issue

In discussing Besi's latest quarterly results, Blickman – master of half-truths – refers to the forum discussion at Semicon West. In it, top executives from Intel, AMD and Qualcomm, spoke out in favor of hybrid bonding, he claims. This is two-thirds true. AMD is using the technology for its high-end processors and Intel is poised to tackle hybrid bonding on a large scale. For Qualcomm, however, it's a bridge too far for now from a cost standpoint. AMD and Intel, however, are willing to invest. Then again, both are playing in segments, the computer and data center industry, where customers are shelling out hundreds to thousands of dollars for a system tile.

It remains to be seen when hybrid bonding will penetrate application processors such as Qualcomm's. That kind of chip system only costs tens of dollars. At least at the Semicon West forum, PR 'Chidi' Chidambaram was adamant: "We don't have a product on our roadmap yet, because it doesn't meet our cost requirements."

EV Group's Lindner, however, is convinced that hybrid bonding will eventually be used in cell phones. "I can't predict when we'll see a hybrid-bonded apps processor in cell phones, but I'm very sure it will come because it enables an interconnect density that you can't compare to a soldered copper connection."

"Hybrid bonding at the wafer level delivers proven efficiency," says Lindner. "We get back from customers that our machines ensure 100 percent bonding. If errors do occur, in all cases they can be traced to other factors, such as incoming contamination or things to do with wafer preparation, such as surface

roughness. 3D NAND wouldn't be produced with hybrid bonding if it didn't render. There, it's just a cost issue. Flash memories aren't expensive products. 3D NAND manufacturers are using hybrid bonding because they require a pitch of one micron."

The technology director of EV Group believes that with TCB, pitches of seven microns will be possible in the future, maybe even 5 or 3 microns. But Yole's Chitoraga thinks this is "extremely aggressive from a technology point of view." He assesses that TCB has no potential to replace hybrid bonding in 3D NAND and flash memory applications.

Yole Systemplus, a sister company of Yole Intelligence, did a teardown of YMTC's 3D NAND chips. "The pitch of YMTC's latest 3D NAND chips is 0.8 microns," says Chitoraga. The precision offered by die-to-wafer hybrid bonding currently lags behind that. "Using die-to-wafer bonding, TSMC puts the SRAM cache memories of AMD's Epyc and Ryzen microprocessors on the CPU with a pitch of just under 10 microns."

But even pitches of 10 microns with thermo-compression aren't ready for commercialization in high volumes, says Chitoraga, who refers to research at CEA-Leti in Grenoble for a glimpse into the TCB future. There, they achieve pitches below 3 micrometers with die-to-wafer hybrid bonding. "This isn't yet commercially available and there's always the question of what the yield will be in practice. We really believe that hybrid bonding will be adapted more and more. But it won't replace thermo-compression. Both will continue to coexist."



A look inside the cleanroom of SK Hynix, the largest producer of high-bandwidth memories, which consist of a stack of dynamic RAMs. Credit: SK Hynix

Optical

With Intel and TSMC as lead customers, Besi appears to have a formidable lead over competitors K&S and ASM PT when it comes to die-to-wafer. Along with Samsung, Intel and TSMC are the forerunners in

hybrid bonding. Blickman outlines a “unique transition” in the semicon market, with the major chip manufacturers going big for 5-nanometer chips and next-generation hybrid bonding. “It has become clear over the past 12 months that it’s no longer a question of whether hybrid bonding will reach the mainstream. It’s just a question of how soon and for what applications.”

Nvidia’s latest GH200 Grace Hopper superchip is basically a composite of graphics processors, microprocessors and around it more than a handful of high-bandwidth memories. No hybrid bonding is yet involved in putting together this tile.

Another question is whether future packages will need millions of connections per square millimeter. Perhaps serial communication will be a better and cheaper solution than parallel information exchange. In the computer and PC industry, parallel communication technology such as SCSI was already supplanted in the 1990s by serial standards such as USB.

Optical communications can make this possible. Last September, TSMC made it clear that it sees a future for this technology. Douglas Yu, VP of pathfinding for system integration at the Taiwanese foundry, says he wants to bet heavily on optical channels between chips. He sees in this both an opportunity to increase performance and a chance to reduce power consumption and heat problems in chip packaging. “This could well be the beginning of a new era,” he said during a discussion session at Semicon Taiwan last September.

During the Semicon West forum discussion on hybrid bonding, Intel’s Babak Sabi was also explicit about the promises of optical interconnects. “I know we all love copper here,” he said. “But eventually, even copper doesn’t hit enough for high communication speeds and we’ll have to go to optical.” Sabi estimates that by the end of this decade, we’ll start seeing optical interconnects in many varieties. His company has already developed an optical connector for PCBs. “If you look a little further into the future, electrical networks will be replaced by optical networking.”

According to Sabi, the world of artificial intelligence is crying out for more bandwidth. “Then we need to go beyond the current assembly on wafers,” he said at Semicon West. “We’re then talking about integrating an entire product at the wafer level.” To achieve the required system complexity, the Intel VP thinks communication on a system wafer via glass is going to become very important. “With that, you virtually eliminate the delay in signals.”

Main picture credit: EV Group

<https://bits-chips.nl/artikel/besi-holds-the-best-cards-in-advanced-packaging-but-the-game-has-only-just-begun/>