



Chipelets:

The path to IoT diversity



Contents

1	Executive Summary	02
2	Opportunities abound in the IoT market – if the price is right	03
3	The traditional approach to chip development limits the progress of IoT devices	05
3.1	Reducing node size is costly and time-consuming using the traditional monolithic approach	05
3.2	Analogue functions do not scale down well	05
3.3	Drawbacks of forcing all chip functionalities onto a single node size	06
4	Chiplets offer a way to remove the barriers to entry for IoT	07
4.1	Reduced investment costs	08
4.2	Lower cost of specialisation	08
4.3	Lower production costs	08
4.4	Shorter time to market	08
4.5	Lower supply risk for OEMs	08
4.6	Simpler architectural partitioning	08
5	Chiplets provide the path to IoT diversity	09
6	The chiplets approach is driving change in the industry, but more needs to be done	11
6.1	Packaging and interconnect	11
6.2	Implications for manufacturing processes	12
6.3	Business impacts	12
7	Chiplets offer a bright future	13
	Appendix	14
A.1	Performance gains at smaller node sizes are diminishing	14
A.2	Why are the costs to migrate a monolithic chip so high?	14
A.3	How do chiplets lower the production costs?	15
	References	16
	Authors	17
	Glossary	17

1 Executive Summary

With global spending on the Internet of Things (IoT) forecast to reach the trillion dollar mark by 2022, clearly the IoT is here to stay. In the past few years, we have seen increasingly intelligent products, from smart medical implants that monitor our health to smart speakers that allow us to shop online. But IoT technology could be used for so much more. As silicon chips become increasingly capable of artificial intelligence and machine learning, we can envision a world where even the smallest, most inexpensive everyday items incorporate intelligence for the benefit of the end user.

Unfortunately, this type of product has not been possible to date, due to the investment of time and money required to develop and manufacture the underlying, traditionally monolithic semiconductor chip technology. Each IoT product requires a unique suite of functionalities comprising sensor technology, memory, processing power and connectivity, which are simply too expensive and time-consuming to develop using the traditional semiconductor system on

a chip (SoC) development cycle. Thus, many innovative products never see the light of day. However, there is an alternative to monolithic chips – silicon chipelets, which allow customised functionality to be developed quickly using modular building blocks and manufactured at very low cost. Although chipelets bring their own technical and business challenges, the semiconductor industry is clearly embracing the concept and developing solutions to drive chipelet architectures forward. In this whitepaper, we explore how chipelets offer a path towards increased IoT diversity, bringing benefits to original equipment manufacturers (OEMs) and semiconductor companies alike.



2 Opportunities abound in the IoT market – if the price is right

Semiconductor investment has been focused on device miniaturisation to achieve cutting edge performance, but this is not the optimal approach for IoT devices

With the advent of 5G and cloud computing, the Internet of Things (IoT) has been predicted to transform businesses of all sizes as part of the global drive towards digital transformation.

Global spending on the IoT is forecast to grow at double digit growth rates for at least the next few years, reaching the \$1 trillion mark by 2022 according to some forecasts.¹ All aspects of society are touched by IoT, from consumers to businesses to government. Manufacturing, transportation and utilities are all investing heavily in this new technology, and consumers will also drive demand for intelligent devices to keep them informed about their health and homes.

However, IoT implementations continue to be costly and complicated to deploy in practice. Some estimate that to deploy an IoT solution in a business environment costs at least US \$1 million² with each sensor costing between \$10 and \$50. Consumer applications of IoT have also gained attention for audio, medical or fitness applications, but to date products like fitness bands or voice-activated smart speakers have been limited in variety and produced in very large volumes to achieve a realistic return on investment.

In theory, IoT could be applied to hundreds of different applications in the consumer market and even the fast-moving consumer goods (FMCG) market. As silicon chips become increasingly capable of artificial intelligence and machine learning, we can envision a world where even the smallest, most inexpensive everyday items can incorporate IoT technology for the benefit of the end user. Smart disposable syringes could provide tailored drug delivery for each patient, for example, or intelligent tags could be placed on food products to alert a shopper that the product is no longer fresh and safe to eat.

But these types of products tend to be very sensitive to cost and traditionally it has been prohibitively expensive and time consuming to develop the underlying silicon chip technology. Unlike semiconductor chips for, say, cloud or smart phone applications, which tend to be homogeneous in performance and highly commoditised, each IoT product potentially requires a unique suite of functionalities comprising sensor technology, memory, processing power and connectivity. This is simply too expensive and time-consuming to develop using the traditional semiconductor chip development approach for all companies except the largest and most established in the industry with substantial R&D budgets.

Thus numerous innovative applications never see the light of day.

The burgeoning number of IoT use cases presents a huge opportunity for OEMs that want to offer innovative products, as well as for the semiconductor industry that supplies the underlying silicon-based components, especially as more and more novel use cases present themselves. However, to enable these new use cases and further propel the market forward will require a reduction in cost of the IoT devices themselves and more specifically the silicon chips that provide the intelligence.

Over the last half century investment in semiconductors has been strongly focused on increasing integrated circuit performance of the single monolithic chips used for high-performance applications. This has led to a continual reduction in the size of the individual transistors within semiconductor chips (what is known as 'process node size' in the industry). Historically, smaller node sizes have generally provided benefits for semiconductor manufacturers across nearly all areas – performance, end product cost and power consumption – with few drawbacks. However, performance gains have begun to plateau in recent years (see [Section 3](#)).

“We can envision a world where even the smallest, most inexpensive everyday items incorporate IoT technology.”

With the inception of the IoT, semiconductor companies are reconsidering what makes a successful chip. Unlike the commodity laptop market, for example, which has a narrow focus on ever-increasing amounts of compute, the IoT market does not have a single 'killer application', but rather is characterised by a large number of diverse use cases in more modest volumes. IoT chips are also 'mixed signal', that typically contain a combination of digital and analogue functions. Together these factors place a very different set of demands on the underlying silicon technology than the quest for higher-performance compute. The functionality of digital components, like memory and processing, scales down well with smaller node sizes, with performance improving as the node size decreases. Conversely, the functionality of analogue functions, such as sensors and radios, does not scale down well, for a variety of reasons, not least the large effect that physical size has on performance level (see Section 3). This means that significant investment is required to port analogue functionality to smaller node sizes, which may not offer any performance benefits, and may even cause a deterioration in performance.

Companies across the semiconductor supply chain have re-invested in older, more established node sizes for IoT, but this approach is unlikely to be sustainable

To address the IoT market, semiconductor companies have made a second wave of investment to repurpose old, established node sizes for use in mixed-signal IoT applications using the same monolithic system on a chip (SoC) approach. This is evidenced by the fact that demand for wafers at larger node sizes is still considerable, even as newer, smaller process nodes are introduced (see Figure 1).

The migration of analogue functionality from 55 nm and 40 nm to 22 nm, which is currently underway, is a significant challenge for semiconductor companies, costing tens of millions of dollars (see Section 4.1). As geometries become smaller, the process of migration becomes increasingly challenging and costly.

This problem is further exacerbated by the drive towards digital functionality like machine learning and artificial intelligence in IoT devices. Keeping more functionality on end devices rather than in the cloud offers clear benefits regarding latency and bandwidth for some applications and privacy can be improved by keeping data closer to the end user. As a result, state-of-the-art node sizes in IoT devices will be critical to ensure fast, low power execution. This creates a problem for IoT chip vendors that wish to rely on older node sizes for analogue functionality, and it means that achieving optimal digital and analogue performance on a monolithic chip presents an even greater challenge.

Given that the IoT market is expected to be a significant area of growth for technology companies over the next decade – with predictions that by 2022 IoT modules will make up 51% of total global devices/connections, a total of 14.6 billion connections globally³ – semiconductor companies will want to create a portfolio of chips that allows them to address the range of possible applications across various industries. However, due to the breadth of IoT use cases, achieving this at an appropriate cost point is going to be very difficult with the current monolithic approach, which presents challenges across various dimensions.

“Creating monolithic chips to address the diverse range of IoT applications is time-consuming and cost prohibitive.”

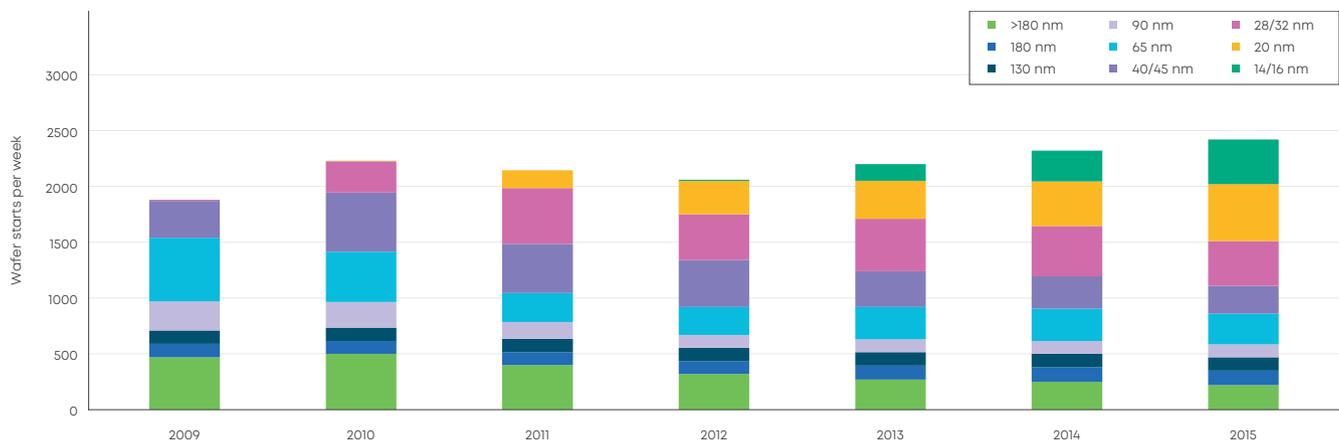


Figure 1 - Historical wafer production per week of different node sizes (Source – VLSI)⁴

3 The traditional approach to chip development limits the progress of IoT devices

As previously mentioned, the current focus on improving chip performance by reducing node size is not ideal for IoT applications for a number of reasons. First is a general issue with the continual drive towards smaller process node sizes, whereas the second and third issues relate to scaling analogue functionality and using a single node size. The latter two are particularly pertinent to creating a portfolio of IoT chips. The issues can be summarised as:

- Reducing node size generally is costly and digital gains are starting to diminish
- Analogue functionality does not scale well to new node sizes
- Using a single node size reduces flexibility

3.1 Reducing node size is costly and time-consuming using the traditional monolithic approach

Whilst we expect the semiconductor industry to continue to develop smaller and smaller node sizes, the performance gains that result are beginning to diminish with each additional node size reduction (for further details, see Appendix A.1). Reducing the node size brings other consequences too. Semiconductor companies are taking increasingly long times to produce chips at smaller node sizes, due to the difficulty of maintaining high yields, leading companies like Intel and AMD to move away from monolithic approaches in their latest CPUs.⁵

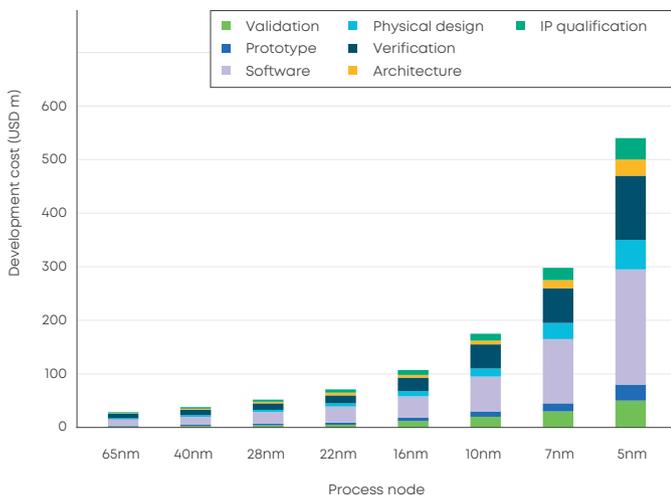


Figure 2 - Development cost for new process node of SoC devices (Source - IBS)⁶

Further, the cost to move to a new node size is also increasing. The small geometries demand more numerous and complex process steps to manufacture the chips, leading to expensive mask-sets for lithography. There is also an increasing probability that multiple design iterations ('re-spins') will be required with smaller process nodes to achieve the desired performance and high yields. It is not uncommon for two or three re-spins to be required for a new chip, which significantly increases both the cost and the time to market. Tooling costs at smaller process nodes conforming to the additional process design rules and the extra fabrication steps are also more expensive due to the increased complexity of the design and the extra fabrication steps. An increased amount of verification is also required for smaller nodes. Figure 2 shows the exponential increase in development cost as process node size is reduced.

Usually, once a new process node has become established, the costs of producing further chip designs for that process decline over time, due to reduced costs to create the lithographic masks, for example. However, it is unclear whether this depreciation in migration cost will occur to the same extent for smaller node sizes.

Since many novel IoT applications will require custom silicon chip functionality, these large development costs severely inhibit the entrance of new IoT-based products to the market and slows innovation and improvement of existing IoT-based products.

3.2 Analogue functions do not scale down well

As mentioned above, IoT products require both analogue and digital functionality. Unlike digital functionality, analogue components, such as sensors, do not necessarily scale down well for manufacture on the reduced process nodes typically used for digital components like memory, for both technical and business reasons. There are multiple benefits from staying at a larger node size (i.e. 28 nm and above) compared to digital components.

Since the technical performance of analogue devices tends to improve with increasing physical dimensions, certain performance characteristics are compromised when the device is transitioned to smaller node geometries.⁷

On the business side, there are significant time delays, costs and risk involved in porting analogue functionality between process nodes, although there are companies that are actively working to reduce these effects. Agile Analog,⁸ for example, claims to have developed an innovative design methodology to optimise analogue components. However, more typically, developing novel analogue functionality includes significant engineering effort, electronic design automation (EDA) tools, run time, mask costs, fabrication processing time, re-verification, re-validation and re-characterisation efforts as well as logistics management. It may also involve re-qualification costs and effort to ensure the product conforms to industry standards.

Based on both the performance degradation and increased cost and complexity to manufacture at smaller node sizes, it no longer makes sense to integrate analogue and digital functions on the same piece of silicon for many IoT products.

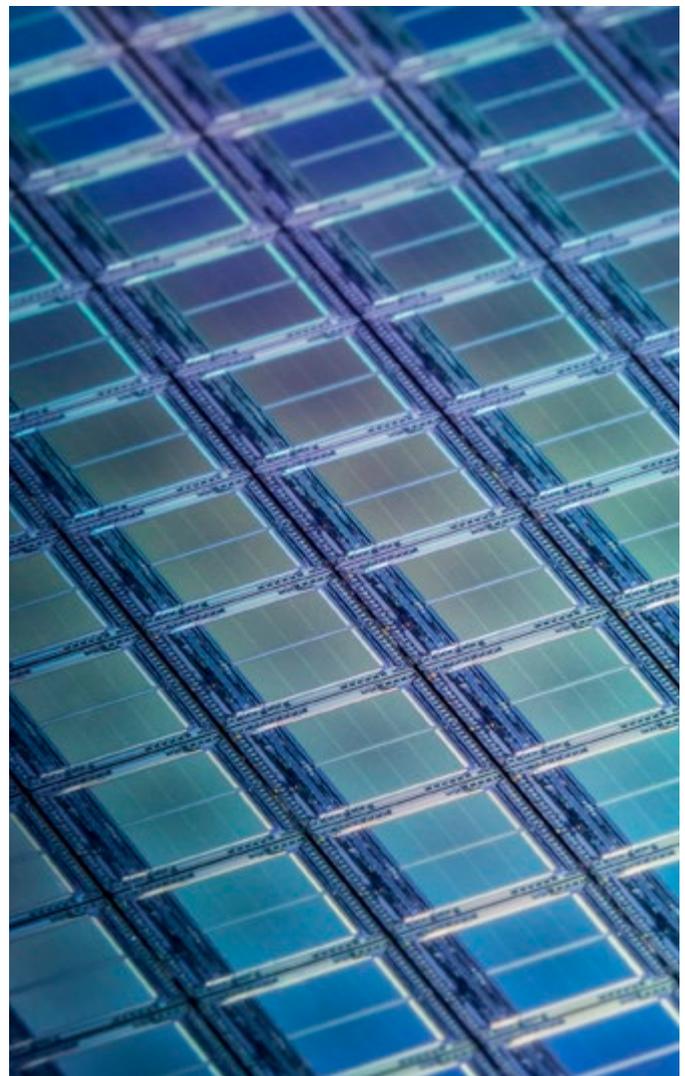
3.3 Drawbacks of forcing all chip functionalities onto a single node size

The differing requirements for analogue and digital functionality can also give rise to additional complications:

- Issues with combining technology from different chip designers. In the semiconductor marketplace, there is a considerable portfolio of existing chip layouts and functional blocks that are developed for repeatable use on a fee basis. By using this third-party intellectual property (IP), OEMs can reduce the cost of developing a chip for their specific application. However, it is not straightforward to combine third party IP onto monolithic chips because the designs and manufacturing steps may not be compatible.⁹ To overcome this sometimes requires the removal of features from the chip, the selection of a less optimal process node or the re-working of any underlying IP.
- Less flexibility in terms of the physical size of the chip. The more complex the monolithic chip, the more likely it is to include functionality that is redundant for highly-specific IoT applications. Complex designs are more expensive to fabricate, so the cost of the end product is driven up.

“It no longer makes sense to integrate analogue and digital functions on the same piece of silicon at the smallest nodes.”

- Complex fabrication processes required for monolithic chips can affect the performance of the final product. The processing techniques required for one functional block on a monolithic chip can affect the functionality of other blocks, and consequently the final integrated circuit (IC). For example, non-volatile memory requires extra processing steps which potentially affect the integrity of the base process device characterisations.¹⁰
- Even incremental functionality improvements incur large development costs. Adding new functionality to a monolithic chip, however incremental, can often require a substantial redesign and re-planning of the full chip. Consideration must be given to how the new functionality will affect other parts of the chip and potential for introducing effects such as crosstalk, magnetic coupling and RF interference, as well as the potential interaction with the package and external components.



4 Chipllets offer a way to remove the barriers to entry for IoT

Monolithic chips, then, are not the optimal approach for implementing mixed-signal IoT devices in all situations. An alternative approach is the use of chiplet architectures, which allow multiple chiplets, with different process node sizes, to be used in a single package or on a single substrate. These types of architectures bring a host of benefits for IoT product developers compared to the SoC approach, including:

- Reduced investment costs
- Lower cost of specialisation
- Lower production costs
- Shorter time to market
- Lower supply risks for OEMs
- Simpler architectural partitioning

“More choice, less cost, faster time to market, lower risk – chiplet-based architectures bring a host of benefits to IoT product developers.”

Figure 3 shows an example of how the functionality of a traditional monolithic IoT chip could be achieved using the modular chiplets approach. On the left-hand side of the figure, we see the traditional monolithic chip, with all functionalities designed on a single piece of silicon and hence a single process node size. The chiplets approach is shown on the right – the chip functionalities have been divided into physically discrete pieces interconnected via a substrate. This separation relaxes the requirement to use a single node size, allowing each functionality to be optimised separately. Furthermore, the ability to easily swap IP blocks means that an existing functional block can be evolved to suit a new purpose much more quickly and easily and at a lower overall cost. This is particularly important for OEMs looking to create a portfolio of chips to serve different IoT applications or to react quickly to new IoT market opportunities.

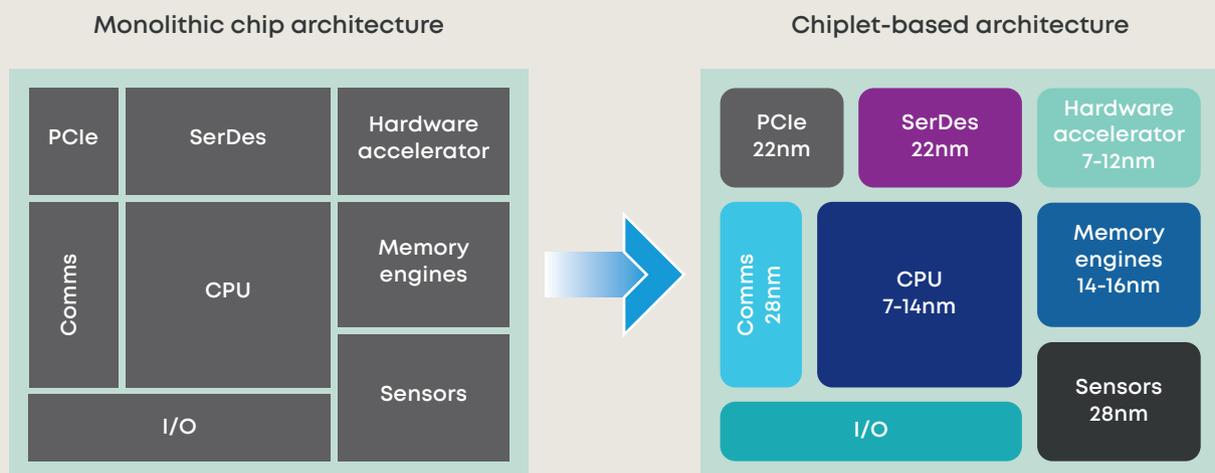


Figure 3 - Monolithic vs chiplet architecture (Source – Cadence, Intel).^{11 12} *Node sizes are for illustration purposes only.

4.1 Reduced investment costs

As discussed above, developing a new target process node for a family of chips often requires a significant investment and planning to secure a reasonable return on investment (ROI). The initial investment is predominantly defined by the size and complexity of the chip as well as the selected process node. Currently the semiconductor industry is transitioning monolithic IoT chips from higher node sizes, like 40 and 55 nm, to 22 nm. However, we estimate that this type of migration requires substantial development investment, as illustrated in Figure 4, where the cost of developing a 22 nm IoT chip could cost around US\$ 60million (discussed further in Appendix A.2). The costs to develop similar functionality would be significantly lower for chiplet-based architectures which use functional blocks at more established node sizes rather than designing them onto newer, smaller nodes.

4.2 Lower cost of specialisation

The chiplets approach also means the design cost is reduced because for any particular end product, OEMs need only to design the specialty chiplets required for their application and then combine them with commercially available chiplets for the standard functionalities.

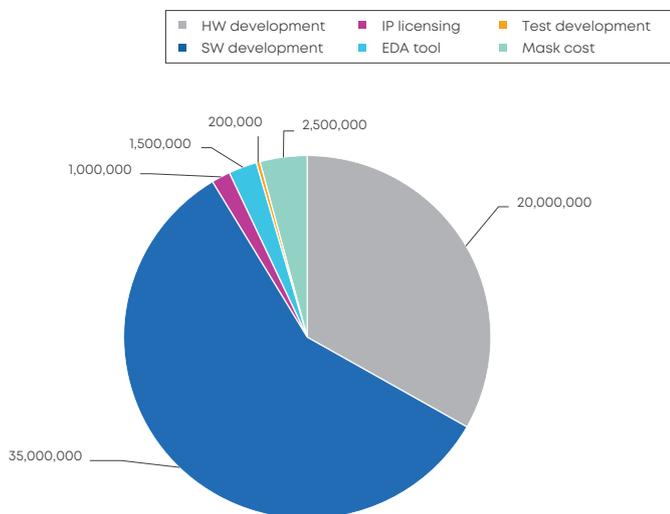


Figure 4 - Development cost (US\$ 60 million) of 22 nm node IoT SoC (Source – Cambridge Consultants estimate, GSA)¹⁶

4.3 Lower production costs

Once a chip has been designed and validated for performance, it is ready for transfer to a facility for volume manufacturing. The chiplets approach can reduce the cost and time of the transfer, as well as the unit manufacturing cost of the chip and, by extension, the final product. By using larger process nodes for the analogue components, chiplet architectures lead to (further details are provided in Appendix A.3):

- Reduced number of design revisions ('re-spins') increased yields
- Simpler fabrication processes

4.4 Shorter time to market

A key commercial benefit of chiplets is that they enable a faster time to market of new products in two ways:

- Faster prototyping – rather than requiring a large investment to create a new monolithic chip, ensuring each of the blocks function correctly on the selected node size, semiconductor companies and OEMs can mix and match different blocks of silicon, including custom chiplets, to create a trial product for rapid market testing.
- Easier upgrades – selecting functional blocks from a 'silicon library' of third-party IP could allow new functionality to be added to a chip without a complete redesign of the chip, thereby avoiding the time required for testing and validation. This is particularly useful if certain functionalities are developing at a faster rate than others, which is the case for ML chips currently.

4.5 Lower supply risk for OEMs

Chiplet designs allow OEMs to have more freedom in sourcing their functional blocks from multiple suppliers versus a sole supplier, thereby reducing the supply risk.

4.6 Simpler architectural partitioning

Chip designers frequently approach the design from a top down view, meaning that it is function driven. Once the data flow through the system is decided, it is the hardware architecture which is crucial for correct system operation. However, post-repartitioning is often required to change some of the physical implementations. Chiplets provide a more flexible means to position the functional elements within the chip, so that it can meet the power and regulatory requirements more easily and accommodate late changes in partitioning.¹³ Flexibility in re-design minimises the time and cost to accommodate changes, an important factor for OEM product release deadlines.

5 Chipselets provide the path to IoT diversity

Chipselets are well-suited to mixed signal IoT devices

As mentioned previously, IoT applications tend to be highly diverse in terms of their required functionality. In Table 1 we show the functionalities required for a variety of use cases, along with the state-of-the-art node size for optimal performance. IoT devices vary significantly in terms of complexity, ranging from inexpensive, disposable RFID tags used for tracking of goods – which comprise an RF chip, CPU and data storage – to complex sensor and compute units used, for example, in advanced driver assistance systems

(ADAS), which combine multiple sensors, advanced machine learning or artificial intelligence chips, processing capability, memory and connectivity.

Given the benefits in section 4, chipselets allow both rapid prototyping for new applications, and short time to market for final products, with a suitably low cost-point to make them viable in the market even at relatively low production volumes.

Application	Sensor	Actuator	Connectivity	CPU	GPU/ML	Persistent memory
Digital, analogue or mixed	Analogue	Analogue	Mixed signal	Digital	Digital	Digital
Optimal node size for performance	40 - 180 nm	130 - 180 nm	40 - 180 nm	5 - 14 nm	5 - 14 nm	22 - 40 nm
Active RFID tracking tag ¹⁵			●	●		●
ADAS image processor ¹⁶	●		●	●	●	●
Smart speaker	●	●	●	●		●
Smart medical implant	●		●	●		●
Fitness band	●	●	●	●		●
Medical inhalers	●	●	●	●		●

Table 1 - Example IoT applications and functionality (Source – Cambridge Consultants)

*A note of caution is required here: node size has traditionally been used to denote the size of transistors inside an IC, with smaller node sizes being more desirable as they correspond to more powerful CPUs. However, as node sizes have shrunk to the width of tens of atoms, feature geometries have changed to address the manufacturing and performance limitations that arise, so there is no longer a direct correlation between node size and chip performance. As a result, it is possible that a 10 nm chip from one manufacturer will be comparable to a so-called 7 nm chip from another vendor.

Example: a smart implant

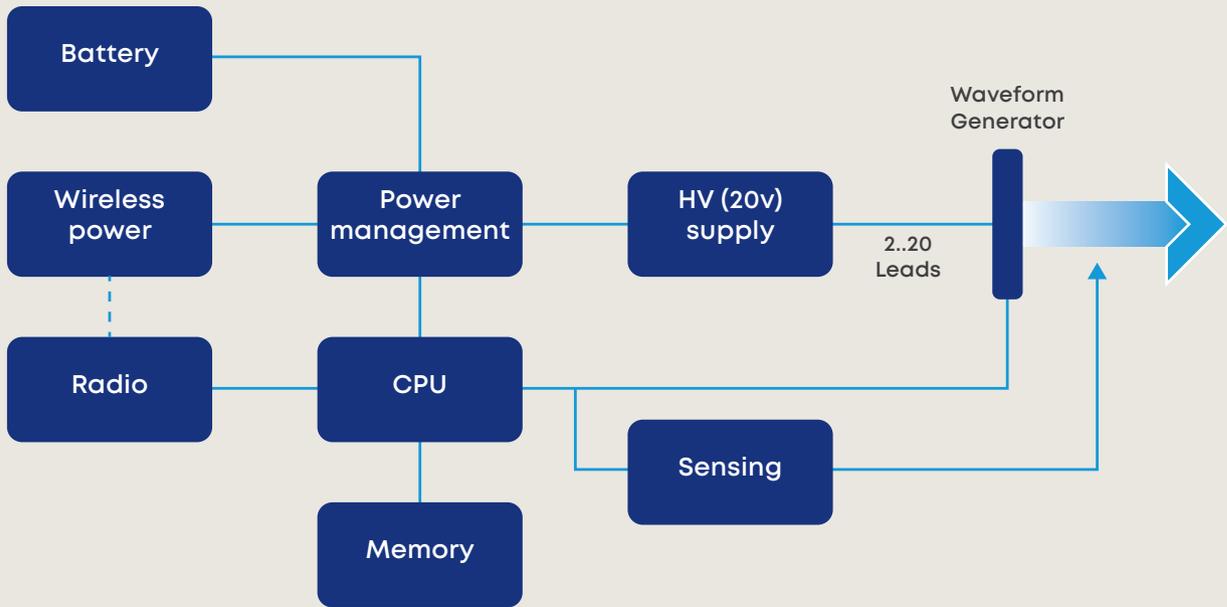


Figure 5 - Smart implant block diagram (source – Cambridge Consultants)

There are numerous potential applications of IoT in the medical field, from insulin monitoring and delivery systems to connected asthma inhalers. In this example, we discuss how the chiplets approach could be used to quickly develop a smart implant for continuous monitoring of a medical condition. A block diagram of the functionality required is illustrated in Figure 5. Smart implants are not usually driven by cost, because they must comply to strict medical regulatory standards. However, chiplet-based architectures provide other benefits for smart implants. First, medical implants can present tough performance challenges, such as low leakage currents and high voltages (up to 20V), which might necessitate a specific process node that is different from digital node sizes. Second, medical implants can pose micromechanical assembly challenges, which can be solved by the different form factors chiplets allow. Lastly, separating functionality onto different chiplets minimises the impact of regulatory changes, since only the affected chiplets need to be redesigned and verified. All these factors ultimately mean a lower cost to develop the implant, as well as faster time to market.



6 The chiplets approach is driving change in the industry, but more needs to be done

Due to the use of different geometries within a chiplet-based chip, there are design and business impacts which do not arise for monolithic chips, such as the packaging and interconnection technology; the testing techniques and even the issue of liability in the case of failures in the field. All these factors potentially affect the OEMs that use these chips in their products, either from a cost or a legal perspective, so it is important to understand their implications.

6.1 Packaging and interconnect

To be useful in the final product, the chiplets must be connected together on a suitable substrate and packaged for robustness and ease of handling by the OEM. This forms an unavoidable part of the cost of the final product. The selection of the packaging and interconnect technology for chiplets is highly dependent on the final end use of the device. Higher performance packaging is required for high-performance, long-life devices, such as autonomous driving systems, for example, whereas inexpensive, simple packaging is adequate for disposable fast moving consumer good (FMCG) applications.

There is a great deal of effort being expended on packaging and interconnect technology by a number of large players in the industry. For example, TSMC has recently announced its plans to start chiplet style production with its System-on-Integrated Circuit (SoIC) 3D packaging system. It can be used for chips that are made using 10 nm or finer manufacturing processes. Manufacture is expected to start in 2021.¹⁷ Intel has developed a technology called Omni-Directional Interconnect¹⁸ which provides low power, low latency connections across chips. It allows both horizontal and vertical communication within the package. Intel has also announced a new version of its AIB (Advanced Interface

Bus) standard called MDIO,¹⁹ which enables chiplet-to-chiplet interfacing. Intel has made some of its technologies available to the industry. A range of other vendors are also focused on developing their own interconnect technologies, including AMD's Infinity Fabric and Nvidia's NVLink.²⁰

However, these types of packaging for high-performance applications can be expensive, proprietary and unnecessarily high performance for many low-cost IoT applications, potentially inhibiting innovation in the lower-cost IoT product market, so lower cost alternatives need to be explored. Anisotropic conductive pastes, for example, are used extensively in printed circuit board (PCB) packaging. Given their significantly lower cost, they offer potential for low-cost IoT applications²¹ as does micro-transfer printing, which is being offered by companies like X-Celeprint for integration of heterogeneous components.²²

Clearly the semiconductor industry is taking the need for high-speed interconnect seriously and are driving improvements in this area. Industry standardisation of interconnection would allow chiplets from different suppliers to be easily integrated and help to realise some of the benefits of chiplets. We expect methods of interconnect to become standardised as the chiplet industry matures.

“Chiplets bring novel technical and business challenges, but these are not insurmountable.”



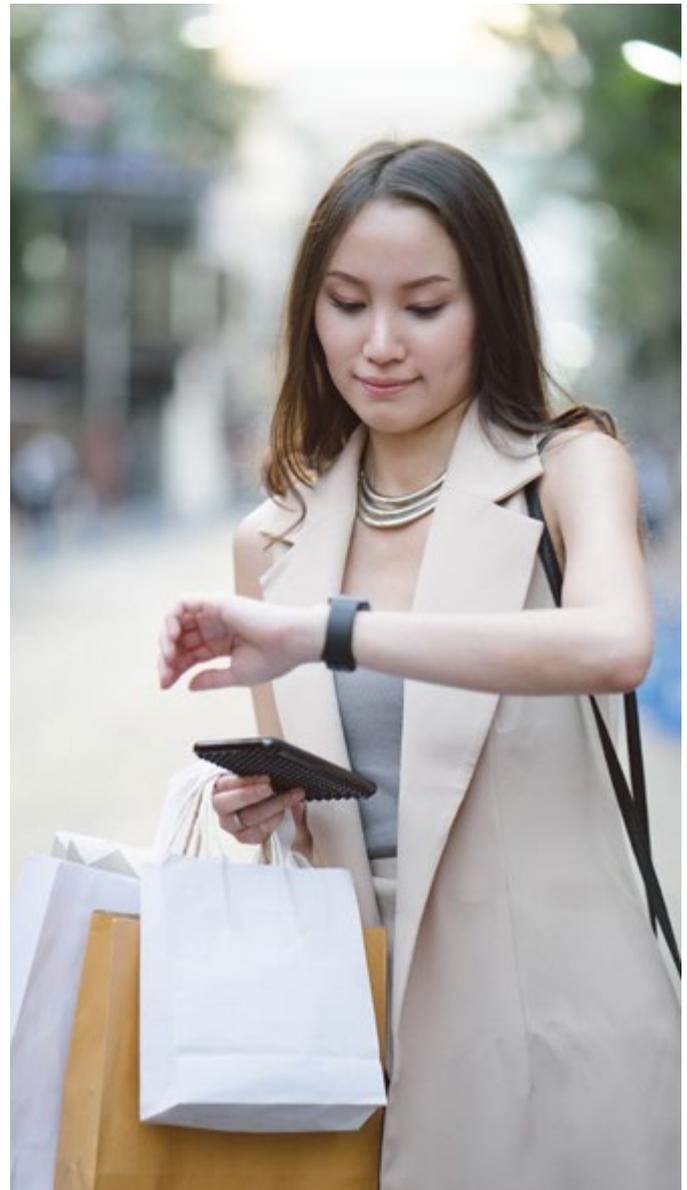
6.2 Implications for manufacturing processes

The chiplet approach might also require a new way of thinking about the manufacturing methodologies in use today, such as testing, for example. For monolithic chips made in high volumes, a standard testing process typically comprises fairly rudimentary wafer screening prior to packaging and then final test of the packaged chip.²³ Since only a single SoC is involved, the wafer sort testing can be minimal, and sometimes even omitted if the fabrication process is high yield, since the cost of proceeding to final package and test is incremental.²⁴

In contrast, chiplet-based architectures comprise a number of chiplets (our example in Figure 3 includes eight), potentially from different manufacturers, any of which could cause the packaged module to fail at final test, increasing waste and the final product cost for OEMs. To avoid this means introducing suitable testing early on in the process to ensure that only good chiplets are used in the final device. Further, it will be important that chiplets are designed to be tested in isolation so that they do not need to be placed with another chiplet for testing purposes. Finally, given the diverse products that could be coming off the production line, chiplet testing capabilities will need to be flexible enough to handle the variety and combinations of performance parameters that will necessarily arise. Other process changes will likely become evident as the chiplets approach evolves. However, any additional costs must be kept low, so as not to offset the benefits of chiplet architectures for OEMs.

6.3 Business impacts

Moving to chiplet architectures will also have an impact on the existing business models of the semiconductor industry. For example, the use of multiple components inside a package raises the question of who is ultimately responsible, and hence liable, if the final device fails in the field. The chip might contain chiplets from multiple suppliers and might also have been assembled and packaged by yet another supplier. This could be particularly challenging for applications in highly regulated industries such as automotive and healthcare which require high reliability. However, there are many other products on the market that comprise numerous components from multiple suppliers (a car is an obvious example), so this issue is not insurmountable, and we expect it to be resolved as chiplet architectures become more widespread.



7 Chiplets offer a bright future

The chiplet approach is gaining momentum and is fuelled by the growing diversity of potential low cost IoT applications. The traditional approaches to servicing this emerging market are not a fit economically. This is compounded by the consolidation in the semiconductor industry, which limits the variety of chips available to OEMs, resulting in devices that are not tailored from either a cost, feature or performance perspective. Chiplets clearly hold promise to enable the next generation of novel, low-cost IoT products due to their flexibility, fast time to market and reduced cost to develop and manufacture. These architectures allow OEMs to easily combine state-of-the-art analogue and digital functionalities into their end products, thereby providing not only the best performance, but also more flexibility in design. The mix-and-match approach can also mean reduced costs to develop these products, since the unique functionality of the product can be designed as a chiplet and then combined with commercial off the shelf chiplets. Time to market is also greatly improved due to the simpler development and transfer to manufacturing steps involved.

The approach also has benefits to semiconductor companies. Chiplets will allow semiconductor companies to provide more optimal offerings for a wider variety of applications whilst keeping the portfolio smaller and reducing development costs and timescales. This interest

“Chiplets benefit semiconductor companies and OEMs alike.”

across the semiconductor industry is demonstrated by the continuing advances in packaging and interconnects to enable these use cases. So what do we see looking forward?

Integration technologies will continue to evolve to enable the whole spectrum of IoT devices, from highly sophisticated ML-enabled autonomous vehicle sensors in which performance is paramount, to lower cost smart tags for consumer applications.

The business ecosystem will also evolve. To enable a truly dynamic marketplace in which chiplets from any manufacturer can be combined quickly and cost-effectively will require the evolution of open technical standards, and co-operative business processes. This work has already begun and will undoubtedly evolve over time.²⁵ An example might be the emergence of new businesses that focus on the integration of chiplets and supply of the resulting modules.

Chiplets offer a welcome alternative to traditional monolithic chip approaches and will enable a host of new, low-cost IoT devices which were simply not possible before. This can only be a good thing for semiconductor companies and OEMs alike.

To discuss semiconductor strategies for IoT products, contact:

Jez Stark, Head of Semiconductor Services,
Wireless and Digital Services
jez.stark@cambridgeconsultants.com



Appendix

A.1 Performance gains at smaller node sizes are diminishing

Figure 6 shows the evolution of (digital) CPU characteristics since 1970. It can be seen that the SpecINT value (a measure of CPU performance) has failed to scale with the number of transistors on a CPU. Additionally, node sizes reached a point around 15 years ago where transistors became so small that the clock speed can no longer be continually increased without increasing power. This necessitated a move to multi-core CPUs in around 2005, as the maximum clock frequency plateaued.

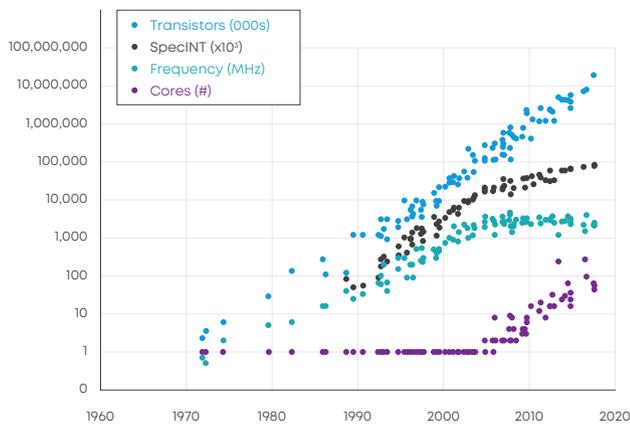


Figure 6 – Evolution of CPU technology (source – Karl Rupp)²⁷

A.2 Why are the costs to migrate a monolithic chip so high?

The development cost for a monolithic chip is made up of several different components, each of which will vary depending on the complexity of the chip. In our example, the software and hardware components, which mostly consist of engineering labour, contribute approximately 90% of the total development cost. In the case of IP licensing, many foundries offer sponsored (free) foundation IP, typically including I/O pads, static random-access memory (SRAM) and standard cell logic libraries. However, additional cost of IP integration and validation complexity can be added to the hardware/software engineering costs. The EDA tooling costs will depend on the number and complexity of the design rules for the process, so are lower for larger process nodes, and may also be affected by the amount and accuracy of the up-front verification required before tapeout. Due consideration also needs to be given to the risks of non-functioning or compromised silicon which impact the schedule and the costs of re-spinning against tool costs, verification efforts and methodologies.

Newly released products such as smartphone ICs on cutting-edge process nodes typically cannot fail in the marketplace without a significant negative commercial and reputational impact. Therefore, the simulation and validation costs for these types of products are high to ensure the best chance of success. Simpler products like IoT chips, on the other hand, will likely require less complex (and therefore typically cheaper) tools as they require less verification effort. Also, the financial and schedule delay costs of re-spins on larger nodes may have less of an impact than in the former case. Moreover, for companies with a portfolio of products the incremental tool cost is likely to be minimal.²⁶

A.3 How do chiplets lower the production costs?

Reduced number of design re-spins

New chip designs frequently require the design to be revisited and adjusted to achieve the required functionality ('re-spin'). Using older (larger) nodes for analogue functionality not only optimises the performance, but also improves the likelihood that the design will work first time, thereby minimising the number of re-spins especially since the analogue functions tend to be sensitive to process. This reduction in re-spins translates to lower cost and shorter time to market for OEMs.

Increased yield

While a single defect on one part of a monolithic chip will cause the loss of the whole chip, this is not the case for chiplets. If any of the chiplets are defective there is the

potential to replace them depending on the manufacturing process. Also, leaving the analogue chips on higher legacy process nodes produces lower failure rates and production line issues, hence, a better yield²⁸. Since yield has a direct impact on the unit cost of semiconductor components, higher yields help reduce the cost of the final product for OEMs and allows foundries to deliver more product to their OEM customers.

Simpler fabrication process

There are further advantages to using older process techniques at larger node sizes when it comes to actually manufacturing the chips. The smaller node sizes of cutting-edge chips require complex lithographic processes.²⁸ Figure 7 below illustrates the potential impact of decreasing node size on wafer (production) cost.²⁹ In addition, process cost for any given node size tends to decrease over time as the foundry gains experience with the process. Again, this translates to lower unit cost per chip, thereby improving product cost to the IoT OEM.

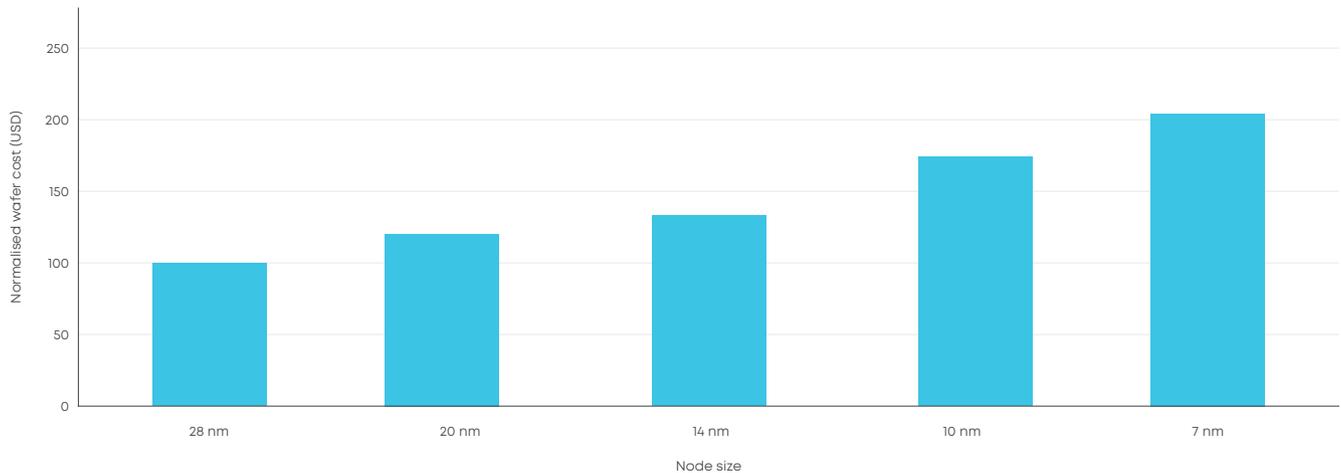


Figure 7 - The impact of reducing process node size on wafer cost (source – Tech Design Forum)²⁹

References

- 1 <https://www.idc.com/getdoc.jsp?containerId=prUS44596319>
- 2 <https://www.raconteur.net/technology/iot-challenges>
- 3 <https://blogs.cisco.com/sp/2018-complete-vni-forecast-update>
- 4 VLSI data from <https://www.electronicdesign.com/eda/reported-death-moore-s-law-premature>
- 5 <https://www.pcworld.com/article/3407802/intel-takes-the-chiplet-concept-to-the-next-level-with-co-emib-odi-connections.html>
- 6 <https://semiengineering.com/the-impact-of-moores-law-ending/>
- 7 Noise is introduced at low node sizes (flicker noise at lower frequencies and thermal noise at radio frequencies). Larger node sizes also allow higher voltages to be used, meaning better dynamic range. Lastly, there is better matching leading to similar threshold voltages across the device as well as lower leakage currents.
- 8 <https://www.agileanalog.com/home>
- 9 The IP may be designed for a particular manufacturing process (node size, metal stack etc) or performance characteristic (such as threshold voltage). The lowest common process options across all of the IP on the IC will restrict the process node that can be selected.
- 10 The addition of a process option may therefore require the complete re-characterisation and potential redesign of multiple IP blocks in order to have sufficient confidence that the final IC will operate in accordance with the correct process device models. Similarly, analogue components may require completely re-verifying and re-signing off; potentially involving long and complex simulations that may expose parasitic extraction differences, resulting in redesign, re-layout and re-signing off.
- 11 https://community.cadence.com/cadence_blogs_8/b/spi/posts/chiplets
- 12 <https://www.techspot.com/news/76079-intel-details-future-chip-design-hot-chips-2018.html>
- 13 <https://semiengineering.com/partitioning-drives-architectural-considerations/>
- 14 The IC Foundry Almanac 2017 edition, Global Semiconductor Alliance
- 15 <https://technology.ihc.com/596669/smart-speakers-teardowns>
- 16 <https://orthostreams.com/disruptive-trend-in-orthopedics-the-smart-implant-revolution/>
- 17 <https://www.eenewsanalog.com/news/tsmc-preps-chiplet-style-manufacturing-2021>
- 18 <https://www.extremetech.com/computing/294659-intels-new-omni-directional-interconnect-combines-emib-foveros>
- 19 <https://newsroom.intel.com/news/intel-unveils-new-tools-advanced-chip-packaging-toolbox/#gs.y9yl1x>
- 20 <https://www.nvidia.com/en-gb/data-center/nvlink/>
- 21 <https://internetofthingsagenda.techtarget.com/blog/IoT-Agenda/Get-linked-up-with-anisotropic-conductive-film-for-your-IoT-devices>
- 22 <https://x-celeprint.com/technology>
- 23 <https://anysilicon.com/ultimate-guide-wafer-sort/>
- 24 https://community.cadence.com/cadence_blogs_8/b/breakfast-bytes/posts/chips-program
- 25 See the work by DARPA for example, <https://www.darpa.mil/program/common-heterogeneous-integration-and-ip-reuse-strategies>
- 26 <http://www.adapteva.com/andreas-blog/semiconductor-economics-101/>
- 27 <https://www.karlsruhp.net/2015/06/40-years-of-microprocessor-trend-data/>
- 28 Multiple masking and metallization steps are involved, each one adding time to the production cycle, as well as potential for defects or contamination
- 29 <https://www.techdesignforums.com/practice/technique/assessing-the-true-cost-of-node-transitions/>

Authors

Oliver Matthews, Consultant, Technology Strategy
Robert McCubbin, Director, ASICs Wireless and Digital Services
Robert Milner, Head of Smart Systems, Wireless and Digital Services
Ivan Petrov, Consultant, Technology Strategy
Heidi Pinkney, Lead Consultant, Technology Strategy
Jez Stark, Head of Semiconductor Services, Wireless and Digital Services
Eugene van Wyk, Senior Consultant, Global Medical Technology

Glossary

Analogue chip	<i>A chip that usually consists of a transistor together with passive components such as resistors, capacitors and inductors</i>
ASIC	<i>An application-specific integrated circuit is a microchip designed for a special application</i>
Crosstalk	<i>An effect that occurs when a signal transmitted on a circuit causes unwanted effects in another circuit or channel</i>
CPU	<i>The central processing unit implements the instructions of a computer program</i>
Digital chip	<i>A chip that comprises components where the signal is interpreted as one of two discrete levels (0/1, false/true, on/off, etc.)</i>
EDA tooling	<i>Electronic design automation tools are software tools for designing electronic systems</i>
Flicker noise	<i>An unwanted effect found in all active electronic components and passive devices; the flicker noise magnitude decreases with frequency</i>
GPU	<i>The graphics processing unit consists of large number of cores which can handle a high volume of software threads simultaneously and is primarily used for rapid image rendering and memory altering</i>
IC	<i>Integrated circuits are small chips having a certain function (e.g. timer, oscillator, memory, etc.)</i>
Lithography	<i>A process used to form the pattern of an integrated circuit on a piece of silicon (or another semiconductor material). One step of lithography involves passing light through patterned glass plates called lithographic masks.</i>
Monolithic integrated circuit	<i>A set of electronic circuits placed on a single chip made of semiconductor material</i>
PCIe	<i>Peripheral component interconnect express is an interface standard used in connecting components of high speed</i>
Persistent memory	<i>A method of storing data structures so that they can still be accessed with the use of memory instructions or memory APIs</i>
Process node	<i>A specific semiconductor manufacturing process with certain design rules; smaller node sizes contain smaller condensed transistors to optimise for performance and efficiency</i>
SerDes	<i>A serialiser/deserialiser is used in high speed communications to compensate for limited input/output</i>
SpecINT	<i>A type of computer benchmarking for CPU integer processing power</i>
Tapeout	<i>The final stage of the design cycle of integrated circuits before being sent for manufacturing</i>
Thermocompression bonding	<i>A process where two metal surfaces are bonded together under controlled conditions (time, temperature and pressure cycle)</i>
Transistor	<i>A semiconductor device capable of switching or amplifying electrical power and electronic signals</i>

About Cambridge Consultants

Cambridge Consultants is a world-class supplier of innovative product development, engineering and technology consulting. We work with organisations globally to help them manage the business impact of the changing technology landscape.

With a team of around 800 staff in Cambridge (UK), Boston, San Francisco and Seattle (USA), Singapore and Tokyo, we have all the in-house skills needed to help you – from identifying the impact of emerging technologies on markets, to creating innovative concepts right through to taking your product into manufacturing. Most of our projects deliver prototype hardware or software and trial production batches. Our technology consultants can help you to develop strategies, policies and plans that take account of the step change often delivered by cutting-edge technologies.

We're not content to create 'me-too' solutions that make incremental change: we specialise in helping companies achieve the seemingly impossible. We work with some of the world's largest blue-chip companies as well as with some of the smallest, most innovative start-ups who want to change the status quo fast.

To discuss semiconductor strategies for IoT products, contact:

Jez Stark, Head of Semiconductor Services,
Wireless and Digital Services
jez.stark@cambridgeconsultants.com



UK — USA — SINGAPORE — JAPAN

www.cambridgeconsultants.com

Cambridge Consultants is part of Capgemini Invent, the innovation, consulting and transformation brand of the Capgemini Group. www.capgemini.com