THE ROAD AHEAD

The changing automotive landscape



Imagination

Foreword

It's a hugely exciting time for the automotive industry. The move from fossil-fuel consuming combustion engines to the electric drive train is a step-change for both manufacturers and consumers, and reflecting this, the in-car experience is changing too. Internally, digital is replacing analogue, both in terms of the human-machine interface (HMI) and under the hood, with electronics playing an ever-increasing role in the functions of the car, from tyre pressure monitoring to sophisticated driver assistance systems, all the way up to fully self-driving cars. Meanwhile, to cater for these changes, the underlying digital infrastructure of the car is evolving, moving from multiple, disparate electronic control units (ECUs) to a more centralised computing architecture.

While these advancements are positive, they are presenting new and unprecedented challenges to car manufacturers. New players, with new operating models and culture are taking on the industry incumbents and in the face of this new competition, both car manufacturers and suppliers are having to rapidly learn and adapt, as the technology and processes with which they have been so familiar evolve at a rapid pace.

The automotive electronics is today a $500 billion industry, and this is only set to grow. Imagination has been supplying its technology into the industry for over 15 years and enjoys a 45% market share for in-car HMI: over 300 million cars on the road today use our graphics technology.

In the coming pages, we will talk more about the changes facing the architecture of automotive electronic platforms. We will discuss advances in HMI, the rise of AI and autonomous driving, the importance of functional safety and the changes coming to the car's connectivity backbone: and explain how we are here to help you solve those problems.

The car industry is on a journey and it is one we look forward to joining you on. Let's get started.

Jamie Broome,
Vice President of Product Management,
Automotive, Imagination





Remake, remodel:
How the architecture of the car is changing

The electronics architecture of the vehicle is undergoing a dramatic evolution. As we move into an electrified automotive world without the distinguishing sound and feel of a combustion engine, software and semiconductor technology will increasingly be the basis for competitive differentiation.

As such, car manufacturers are seeking to advance the human-machine interface (HMI), providing exciting new experiences for the driver and passengers. The latest top-of-the-range cars boast huge multi-screen displays, up to 8K in resolution, cutting-edge head-up displays (HUDs) provide AR-type functions inside the windscreen, while the use of voice and gesture control is increasing. Additionally, the advent of ADAS and self-driving capabilities means that the processing power required, and therefore the number of semiconductors in a vehicle, is set to more than double by 2025. Indeed, according to one major German OEM soon "over 80% of a car's value will be based on the underlying electronics available".

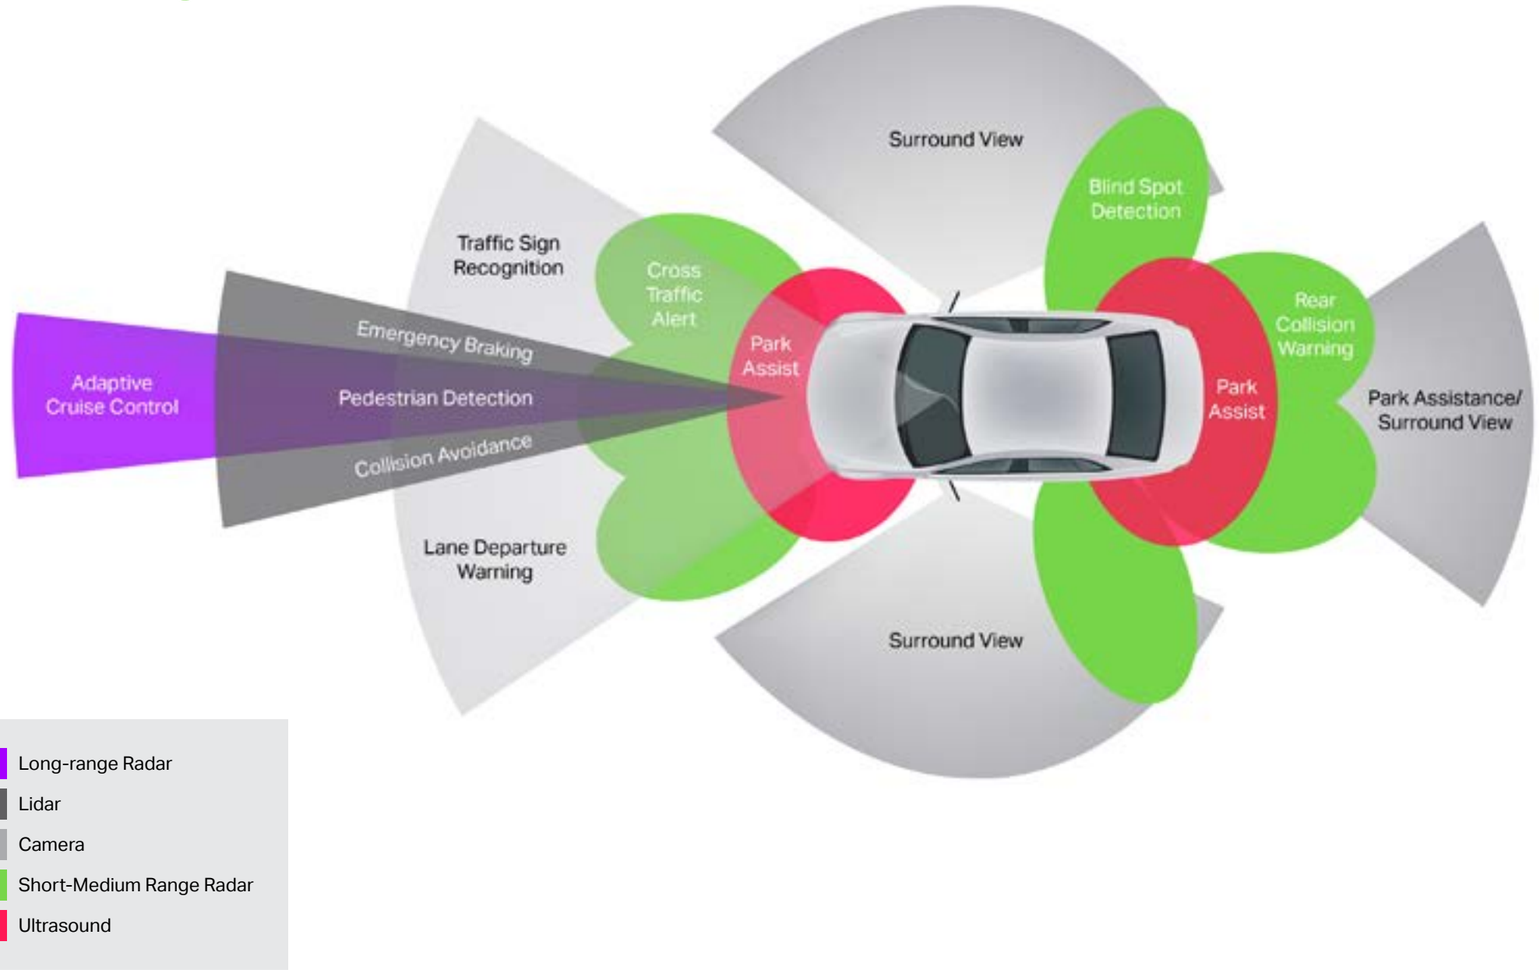


The massively increasing amount of data is driving an increase in the underlying processing requirements of silicon and even today, up to four terabytes must be moved around the car internally. The rise of ADAS and self-driving cars is bringing a proliferation in the number of integrated sensors as manufacturers turn to several technologies to help the car understand the world, from multiple image-based cameras to lidar, radar, infra-red and thermal.

As such, the amount of data being generated and processed will increase exponentially, and to cope with this, gigabit-capable networks will be required.

The architecture of a typical car today comprises anything from 30-100 electronic control units (ECUs), depending on the value and feature level of the car, performing varying amounts of processing from 8-bit micro to multiple 32/64-bit core application-specific SoCs.

Self-Driving Car



Remake, remodel: How the architecture of the car is changing

With the increasing demands of processing and data, this high number of ECUs is viewed as impractical so manufacturers are looking to consolidate these into a smaller number of domain controllers, each controlling a function such as the engine, chassis, or cockpit, in combination with one or multiple, larger central processors. This will bring several benefits. Fewer ECUs will lower costs, increase reliability, and make over- the-air updates easier to manage and reduce weight, helping to offset EV powertrains, which are heavier than the combustion engine equivalent.

To reduce the data traffic load and lower latency, processing is also occurring at the “edge” of the vehicle, as in, directly on the sensors themselves. This central ECU vs edge sensor processing debate is a key question facing the industry, but we believe there will be a hybrid approach for the foreseeable future, depending on the use case and functionality.



The changing face of the automotive value chain: How OEMs are taking control of their destiny.

As the roar of the combustion engine is replaced by the hum of an electric motor, manufacturers are facing a new challenge of how to create the best in-car user experience. In this new automotive landscape, the car is becoming a mobile data centre; a software-defined platform delivering entertainment, navigation, advanced driver assistance systems (ADAS) and autonomy.

In this new world, silicon will be the difference and software will define the future. It is not surprising, therefore, that many OEMs are looking at vertical integration, based around fundamental IP building blocks placed at the heart of the vehicle. Here, custom silicon is designed specifically to handle the precise demands of next-generation autonomous vehicles; from robotaxis to platooning convoys of heavy vehicles, to self-driving personal cars.

On its path to verticalise and gain control of the automotive semiconductor revolution, Silicon Valley has so far stolen a march over the traditional motor heartlands of Detroit and Stuttgart – with China closing in the rear-view mirror. Realising this, many OEMs and Tier 1s around the world are looking to control their destinies and, with an eye on their cutting-edge neural network skills have invested in autonomous vehicle start-ups, while seeking the most appropriate hardware platform to achieve their silicon dreams.



"The automotive electronics is a $500 billion industry."

The Human and the Machine

It is an exciting time for in-car interfaces. Today's embedded, low-power GPUs can deliver extremely high-performance 3D graphics, and even affordable entry-level smartphones offer smooth, fluid interfaces. As such, when consumers get into their vehicles, their expectations are sky high – they will not tolerate slow, clunky, and unappealing interfaces.

Furthermore, as consumers start to move to electric cars, Tier 1s and OEMs will look to the HMI to differentiate themselves from their rivals – it already is a key point of competitive advantage, and the emphasis on this will only grow.

As such, innovation in this space is increasing. Not only are we seeing crisper imagery on larger, higher-resolution screens, but multiple displays will be the norm: the driver's main display, a centre console on the dashboard and even a third for the front passenger. Soon wraparound digital experiences will be possible and there will even be displays for the passengers in the rear. HMI isn't just about the screen: for hands-off use, voice interaction is also becoming increasingly popular.

With the display so dominant, HMIs cannot rely on conventional 2D and primitive designs. Complex, visually arresting 3D graphics will be the order of the day. 3D is already used for features such as sat-nav or surround view, but these will get more detailed and sophisticated.

Ray tracing, a graphics technique that mimics how light behaves in the real world to create astonishing 3D realism will elevate in-car graphics to new heights. We are world leaders in this and are readying IP to bring this revolution to embedded GPUs.

With 3D at their heart, HMI creators are now turning to game engines to create their designs. This is because these software environments are inherently designed for 3D, they enable the designers to review, iterate and collaborate in real-time.



Game graphics is a familiar area for Imagination GPU IP. Our embedded GPU IP effectively launched the mobile gaming revolution and we have been leaders in this space for over 15 years. Our GPUs offer wide API support with many cutting-edge features such as physically-based rendering (PBR), tessellation and high dynamic range (HDR) all supported by the PowerVR SDK.

Remarkably, some car HMIs are now being used to let front occupants play games (though naturally only when the vehicle is parked) as well as rear seat gaming. In this way, the vehicle is becoming an entertainment centre on wheels, but our GPUs have the performance and scalability to meet this challenge.

Our BXS GPU family are a range of automotive-focused cores derived from our IMG B-Series which can provide solutions from entry-level or mid-range HMI experiences (BXS-4-64) up to multiple high-resolution, 4K+ displays (BXS-32-1024 MC2). Our four-core BXS-32-1024 MC4 GPU can deliver an incredible 6 TFLOPS of performance for full blown console-quality gaming – and we have even more powerful solutions waiting in the wings. This is all done with power efficiency in mind: B-Series is up to 30% more power-efficient than our previous generation designs.



IMG B-Series: Learn More

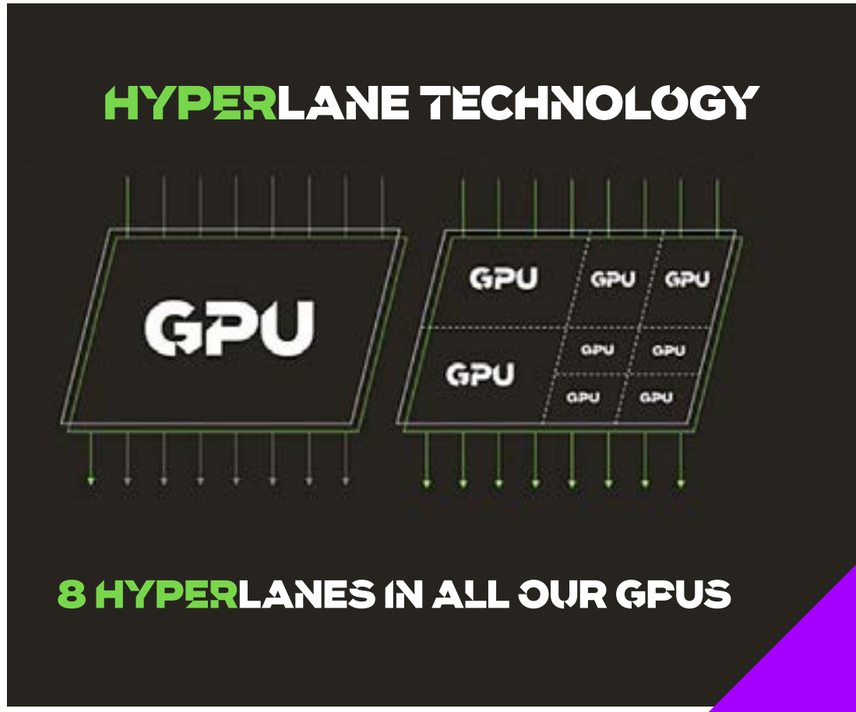
“Imagination has a 45% market share for in-car graphics: over 300 million cars on the road today use our technology.”

The multi-tasking master

Our software driver, combined with our HyperLane technology enables this performance to be used in different ways. HyperLane offers up to eight individual hardware control lanes, each isolated in memory, so different tasks can be submitted to the GPU simultaneously, for fully secure multi-tasking – with prioritisation techniques used to ensure that the necessary performance is delivered to the required applications without any waste of GPU utilisation. For example, a sat-nav could be guaranteed 60fps so there was no possibility of slowdown to impact the driver, with the rest of the available performance given to running the infotainment. Alternatively, a single application, such as gaming, , could use all the cores and lanes at once for maximum performance if desired.

HMI is not just about graphics. The car will employ AI in many places, such as in-cabin monitoring to analyse driver attentiveness, potentially triggering active safety options for autonomous driving modes if the driver is compromised in any way, all the way to natural language processing for voice commands making it easier to control functions in the car without the driver having to take their hands off the wheel and AI algorithms for route planning. AI-powered sensors will even recognise the passenger as they approach, open the door for them and set up the interior exactly the way they like it, be it driving position, interior lighting, or infotainment preferences.

The hardware on which these compute tasks run will depend on the nature of the algorithm. For non-machine learning (ML) algorithms and simple AI tasks our GPUs offer excellent performance, but pure machine learning tasks will run even more efficiently on dedicated silicon such as our IMG Series4 neural network accelerator.



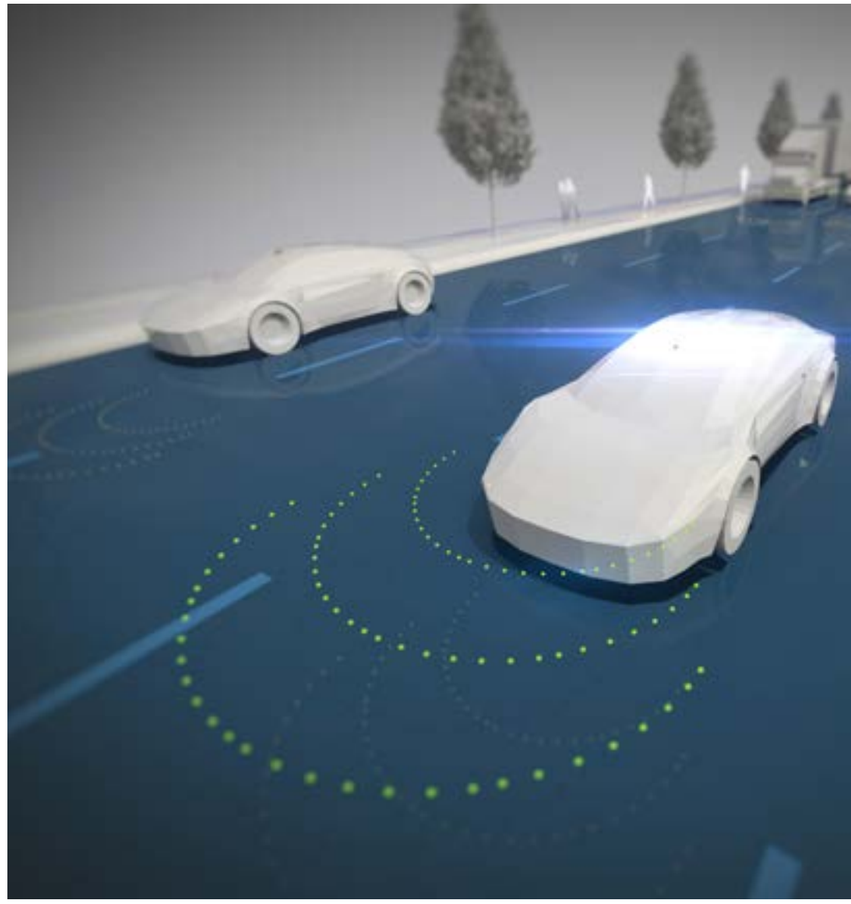
IMG B-Series: Learn More

AI Everywhere

The car concept as we know it is currently gaining significant upgrades in terms of its brawn and its brain. Efficient electric motors are replacing the combustion engine, and eventually, a central processor will remove the unpredictability of the human element from the driving experience. The Society of Automotive Engineers defines six levels of automation – from Level 0 that has no automation systems at all, to Level 5, where the vehicle performs all driving tasks in all conditions and a steering wheel is not even required. The path to the latter will be a long one, but will certainly require a combination of advanced AI, powerful software, and high-performance computing capabilities in order to extract and process data from multiple sensors, such as vision cameras, lidar, radar, and thermal imaging.

Autonomy aside, AI will play an increasingly important role in every aspect of the car – from in-cabin monitoring that analyses driver attentiveness (potentially activating autonomous driving modes if the driver is compromised in any way), to natural language processing for voice commands. AI-powered systems will even scan the road ahead and adjust the vehicle suspension to accommodate any imperfections in the road surface, improving ride comfort. AI will even play a role even before anyone gets into the car: sensors will recognise the passenger as they approach, open the door for them and set up the interior exactly the way they like it, be it driving position, interior lighting, or infotainment preferences.

The most challenging task for AI will be to power advanced driver assistance systems and full autonomous platforms. To do this will require hardware to absorb and process data from multiple sensors, such as vision cameras, lidar, radar, and thermal imaging.



These systems will depend on neural networks and while these networks can be run on CPUs and GPUs there are now better solutions for this type of compute task.

The compute performance required for Level 3 autonomy will require hundreds of TOPS, while 1000 TOPS will be required for Level 4 and 5. To deliver this in an area and power-efficient way can only be done using dedicated silicon, such as Imagination's neural network accelerator (NNA).

Just three years ago our PowerVR Series2NX NNA led the market with performance offering up to 4.1 tera operations per second (TOPS) – but today a single IMG Series4 core delivers up to 12.5 TOPS. Thanks to our groundbreaking, highly scalable, multi-core, multi-cluster architecture, Series4 can now deliver 500 TOPS of performance and beyond – providing the industry with the definitive platform on which to deploy autonomous driving systems in real-world vehicles.

For Imagination, this is not just theoretical! Our award-winning IMG Series4 NNA IP has already been licensed and is currently being deployed in real-world silicon to enable a truly next-generation autonomous system.

While the focus on any conversation about AI is always on machine learning (ML) as mentioned earlier, there is a need for different flavours of compute to support and complement these ML algorithms. These compute tasks are best served by the parallelism inherent in a GPU. Therefore, a GPU and NNA, running in collaboration, designed with an emphasis on low-power and safety, will address the multiple types of compute needed to power next-generation automotive systems.



IMG Series4 NNA: Learn More

Under the hood

The power of the IMG Series4 lies in its architecture. It uses on-chip memory and shared SRAM along with advanced scheduling techniques to enable multiple neural networks to be run concurrently, delivering incredible performance with low latency, which will be vital to ensure quick response times for autonomous cars. And thanks to ongoing research and optimisation we have new techniques that bring the flexibility to run the latest, updated neural network models on our accelerator.

Power efficiency is also key. Today, Tier 1 manufacturers are training and testing their autonomous vehicles using systems that are derived from desktop and server architectures, but these are simply too power-hungry and heat-generating to be suitable for mass-market production vehicles. In an industry where every joule of energy counts and every millimetre of silicon eats into the bottom line, our heritage as a leading player in fanless, power-efficient embedded IP will be a game-changer.

Some new automotive players are using a vertical model to create their own IP for ADAS and full self-driving, but by working with Imagination, silicon providers can provide the incumbents with the means to not just compete with the start-ups on AI – but to lead.



The data backbone

Cars are becoming more intelligent. With a rise in ADAS systems, a push towards an autonomous driving future, and a raft of in-cabin features designed to make the driver safer and more comfortable, the need to push more data around the vehicle is increasing.

Lidar, radar and high-resolution cameras are among the many sensors that create large volumes of data every second. This data must be moved to the vehicle's compute centres, processed, and then moved on to affect action elsewhere in the vehicle.

To take just one sensor, a single camera can require more than 1Gbps of data transfer speed, and many modern ADAS-equipped vehicles feature up to 16 cameras in the front of the vehicle and bumper alone. In combination with all the other sensor data that might be in the vehicle, this means that many terabytes of data must be moved quickly around the system.

Moving data around a vehicle is not a new concept; CANbus has been in use since 1987, but with this exponential rise in vehicular data transfer, its bespoke but low-bandwidth design is no longer passing muster. To get terabytes of data around a vehicle, quickly, safely, and efficiently, we need to look to another well-established protocol – Ethernet.

Ethernet is a flexible and high-throughput standard; a single twisted pair of modern IEEE standard Ethernet cable can move up to 100Gbps around the vehicle – that caters for a lot of cameras. It can also be routed more effectively with well-defined and easy-to-implement options for cabling redundancy, increasing reliability, and importantly, safety.

However, does all of this come at a high cost? Not necessarily. Thanks to its open standard and universal availability, Ethernet can reduce the cost of in-vehicle connectivity by up to 80%.



So, Ethernet is faster, more cost-effective, and safer than the market alternative. But cabling is only a part of the solution. Creating a reliable networking infrastructure to deliver the data to a central processing unit is critical for the safety and integrity of the system. To route and prioritise traffic you need a switch and router which can handle the data throughput and provide effective distribution.

Our Ethernet Packet Processor (EPP) is exactly that. The IP offers a switch and router that offers low-power and low-area high-performance Ethernet packet processing with minimum CPU load.

The EPP is not just fast and efficient. Our EPP supports all the time-sensitive network standards that the industry demands to provide a safe system and all of our automotive IP is developed following ISO 26262 processes, which have been externally validated by industry-recognised experts HORIBA MIRA. Our EPP has over 12 years of history, having successfully been deployed in ASIL-D silicon in the automotive market to a high-profile autonomous driving customer and as such, has already been proven where it counts – on the road.



Ethernet Packet Processors: Learn More

Built with safety at heart

With automotive systems relying so heavily on electronics, car manufacturers need to ensure that their suppliers offer IP that meets the requirements of ISO 26262, the safety standard for automotive electronics. However, meeting this standard is far from trivial. It requires a holistic approach to creating IP, encapsulating all stages of the design workflow.

Imagination has many years of experience working to meet this standard. We have innovative ways of looking at the system-level safety problem, employing innovative techniques such as System Theoretical Process Analysis (STPA), originally developed at MIT.

As a reflection of our rigorous systematic approach to processes we have been awarded a statement of conformance to ISO 26262 ASIL B by HORIBA MIRA, a provider of independent testing and consultancy to the automotive sector. This means our GPU and NNA IP production lines are aligned to the safety needs of our customers and the ISO 26262 standard.

Imagination was the first in the industry to offer a safety-critical software driver for functionally safe GPU IP. This driver, based on the OpenGL® SC™ 2.0 API ensures that HMI designers can unleash their 3D design creativity, confident that they can leverage the power of a GPU to create their dream interfaces while knowing that the underlying hardware and software drivers are functionally safe. Our safety-critical driver has also been fully designed in-house, so you can be sure that it will extract the maximum performance from the IP.

Safety-Critical Driver: Learn More



"Through our collaboration with leading automotive suppliers such as Renesas and TI, our licensed IP technology features in many of the bestselling cars in the world today."

Built with safety at heart

Our BXS family of IMG B-Series cores are the first functionally safe GPUs in the industry and offer specific features for the automotive market. First, they use error-checking memory so that in the event of any randomly flipped bit, the GPU will detect this and rerun the workload.

Second, our BXS-family employs a feature called **Tile Region Protection** (TRP). This enables the 3D pipeline to be used to display safety-critical HMI information, and it does so without significantly compromising performance. To handle safety-critical information on a screen, data must be handed off to a functionally safe component, which, today, will likely be a CPU. However, CPUs are not designed for graphics, placing a significant limitation on what can be achieved in a functionally safe manner. The result is that anything safety-critical can only be displayed in a visually basic way.

Our functionally safe BXS GPUs remove this limitation. The result is richer, more enticing displays with more advanced functionally.



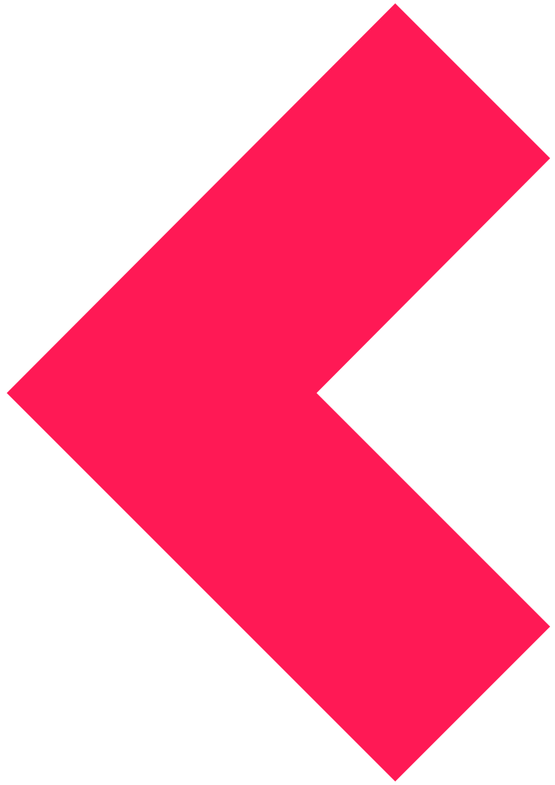
Watch Tile Region Protection Demo

For example, a surround-view system could use camera data to overlay information about potential hazards to the driver directly onto the surround view 3D image of the car, making it easier to spot hazards. As this information is implicitly safety-critical it must be treated as such – so a functionally safe GPU is essential. It benefits HMI designers too. As everything is handled by a GPU, without having to hand over to another component it simplifies the workflow, saving time and cost.

Finally, the multi-core architecture of our GPUs provides car manufactures with a cost-effective method of delivering ISO 26262 compliance by providing built-in redundancy. Each core can be set to check the other's operation so there is no need for an entirely separate GPU. Furthermore, when this functional safety operation is not required, it can be set to deliver full performance for tasks such as in-car entertainment – even gaming!

Equally, our IMG Series4 NNA is built with these same functionally safe processes, so our customers can build ISO 26262 compliant platforms to power next-generation autonomous driving systems.

At Imagination, our many years' experience in the automotive industry has given us the understanding of how to implement the needs of the automotive industry such as robustness, reliability, safety and security; not just at the IP level, but at the integration level. System integrators that work with us are therefore able to benefit from our experience, enabling them to implement their designs at minimal cost and without compromise.



XS Family GPUs: Watch Webinar

Our three-pronged automotive solution

On this journey through the challenges facing the new automotive world we have examined: how the underlying architecture of the car is changing; how innovative new HMIs are revolutionising the in-car experience, discussed the importance of functional safety and shown how AI will impact every facet of the car – with Ethernet underpinning all the required data transfer. And while vehicles will need more GPU and AI horsepower and faster connectivity, this can't come at the expense of power budget, silicon area or safety.

We have decades of experience in producing and bringing up automotive semiconductors. Our IP can be found in 45% of current vehicles on the road today, powering graphics and system intelligence. As a leading automotive supplier, we are tackling these challenges with a three-pronged solution:

- **IMG BXS GPU** – class-leading low-power, multi-core GPUs offering unprecedented performance and designed using ISO 26262 certifiable processes.
- **IMG Series4 NNA** – Award-winning low-power, low-area neural network accelerators that delivers to the automotive industry a platform to enable ADAS and full self-driving.
- **Ethernet Packet Processor (EPP)** – Silicon-proven, very high-bandwidth, low-power, high-performance networking to support next-gen automotive solutions.

With our IP inside, and in conjunction with the wider ecosystem and partners, car manufacturers can be sure that by using suppliers that work with Imagination they will gain access to truly best-of-breed solutions so they can deliver on the promise of a visually rich, power efficient and safe automotive future.



GPU



NNA



EPP

[Contact Us Today To Discuss Your Automotive Challenges]

[Surround View: Watch Demo]

Our three-pronged automotive solution: GPU

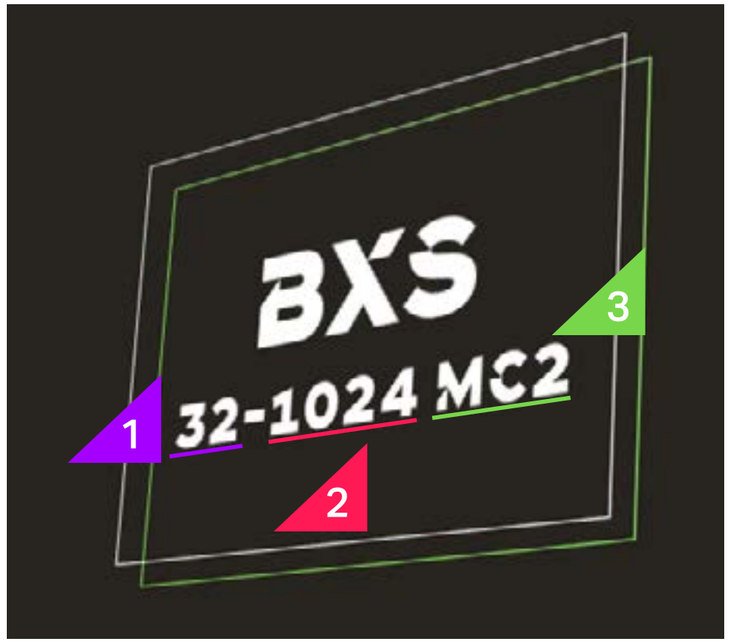
IMG BXS

- Multi-core – for flexible performance scaling and error checking
- ISO 26262 ASIL-B ready design
- Tile Region Protection – functionally safe 3D with maximum performance
- IMGIC – the most advanced image compression technology in the market for ultimate bandwidth saving

IMG B-Series: Learn More

Our cores offer a range of options for range of core options for entry-level UI up to high-resolution, multi-screen complex UI.

1. **Refers to the pixel throughput e.g., 32 pixels per clock, per core – the higher the number the higher the screen resolution capability.**
2. **Refers to the amount of compute performance available: so 1024 = one teraflop per core for greater image complexity.**
3. **MC2 = indicates core count: ranges from 1 to 4.**



Our three-pronged automotive solution: NNA

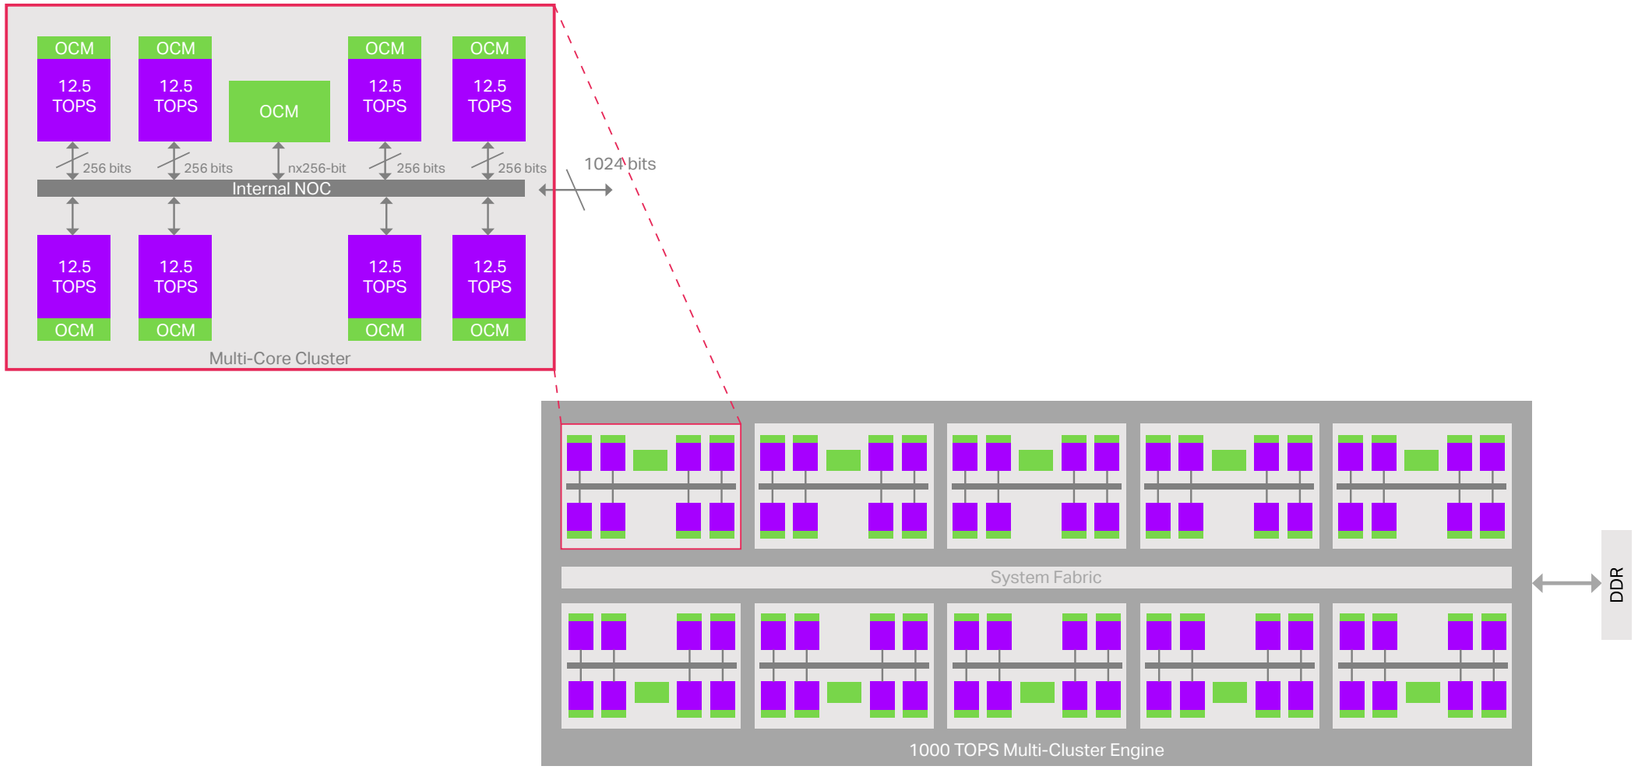
IMG Series4 NNA

The IMG Series4 scalability delivers many hundreds of TOPS to provide a powerful platform for ADAS and autonomous driving.

- Ultra-low area
- Ultra-low latency
- Ultra-low bandwidth
- Ultra-low power consumption
- ISO 26262 ASIL-B ready design

IMG Series4 NNA: Learn More

1000 TOPS performance (illustrative architecture)



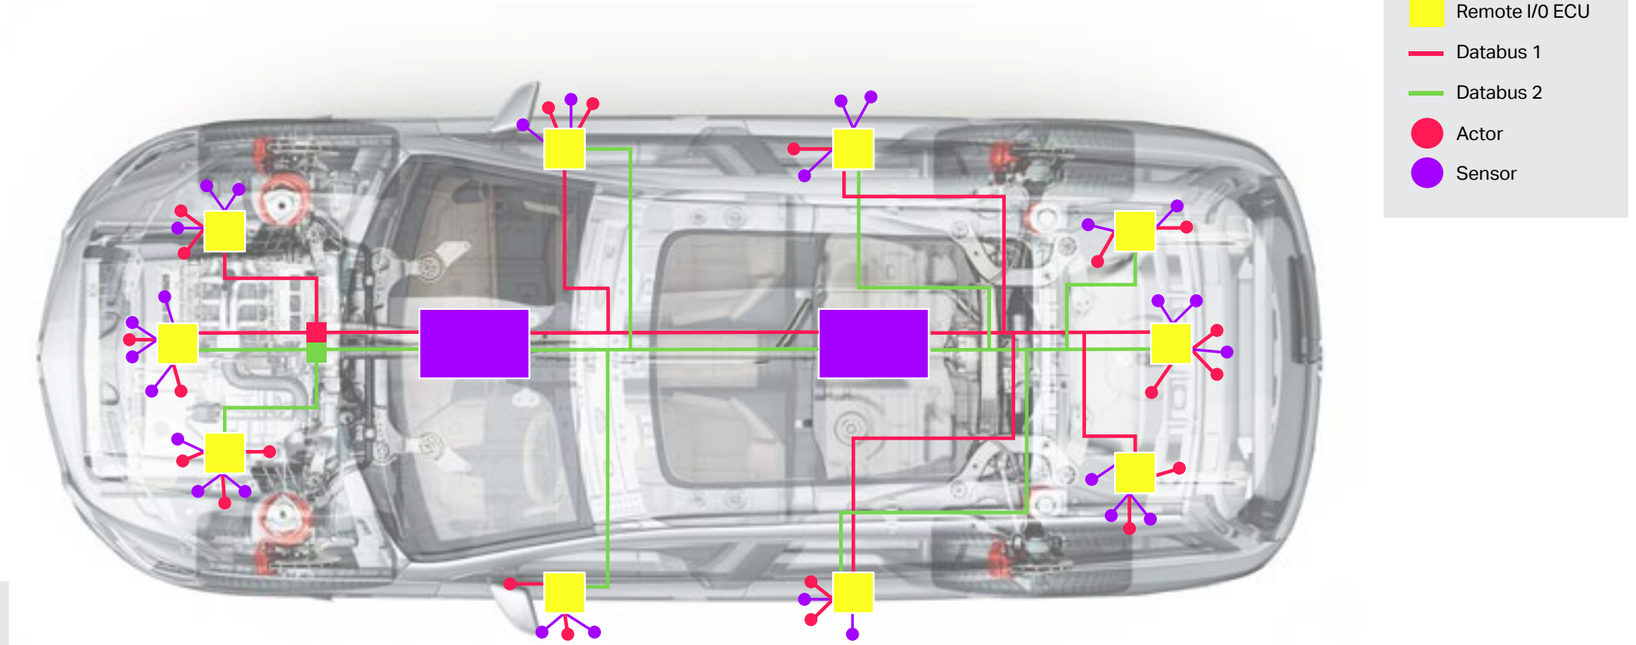
Our three-pronged automotive solution: Ethernet

EPP

Ethernet is the gold standard of in-vehicle networking. Imagination's Ethernet Packet Processor (EPP) is a silicon-proven, fast and efficient combined switch and router designed specifically for automotive.

- Over 400Gbit data transfer across 32 ports for high-bandwidth data transfer and management.
- Flexible configuration to provide the right level of performance at the right price.
- Safety focused – developed to ISO 26262 standards and capable of use in ASIL-B systems.
- Offers crossover into multiple protocols, including Wi-Fi, 5G, CANbus and more, for seamless integration into all systems technology in the market for ultimate bandwidth saving.

Inside the Connected Car



Ethernet Packet Processor: Learn More

Imagination

www.imaginationtech.com