SIEMENS DIGITAL INDUSTRIES SOFTWARE

Real-world considerations for advanced vehicle networks design

Executive summary
This paper describes some of the challenges and considerations which go into developing the electrical and electronic (E/E) architectures of today's sophisticated and highly connected on and off-highway vehicles. The interdependent design objectives that have to be considered are discussed, as well as the technologies that can be used to reduce the risk and time taken in resolving these complexities. Once the E/E architecture is defined, detailed design of the networks takes place ensuring that all signal data is available to functions across the vehicle. This paper examines several of the challenges faced by the network architects and designers. Without assistance, managing these multiple system considerations could lead to a very complex design process. Today, assistance for architects and designers comes in the form of advanced design tools.

siemens.com/networks

SIEMENS

Contents

Introduction and technological context 3

Network load, gateway load 6

Ethernet and switches 7

Functional safety 10

Cyber security 11

Power modes 13

Complexity 14

Summary 15

Introduction and technological context

The E/E architectures used in the automotive industry today are usually highly complex, with many vehicle features delivered by functionality distributed across multiple discrete ECUs. Commonly, the ECUs, sensors and actuators are not all directly connected and much of the communication of data takes place across networks, often through gateways over several networks. Modern E/E architectures have become more formally organized around functional domains, increasingly with domain computers or controllers acting as a centralized compute and absorbing much of the higher level functions for that domain.

The next trend is the increasing use of service oriented architectures (SOA), enabled by Ethernet networking, which allow principals from the information technology (IT) domain to be re-used in automotive applications. The main evolution with such an architecture is the move from thinking about discrete signals to services, which often provide multiple related signals. Functions now provide and subscribe to services appropriate to their functional needs.

Moving towards a SOA is happening in parallel with changes to the physical architecture. Computing power is increasingly centralized, with domain ECUs being reorganized into a zonal layout. Some OEMs and integrators opt for high computing power in the zonal ECUs that reside near the sensors and actuators, while others aim to keep them to relatively simple gateways. Furthermore, moving away from functional domains cascades high integrity requirements into more ECUs across the architecture.

The forces driving OEMs towards centralized or zonal architectures are broader than only reducing the ECU count in vehicles. These architectures can streamline the scalability of functionality, with processing and memory headroom only needing to be provisioned in the central compute unit. Centralized or zonal architectures can also help reduce harness mass and lower bill-of-material (BoM) costs to the OEM. It is worth noting that the speed of this transition is variable across organizations, regions and vehicle market sectors. Well-funded OEMs known for selling premium passenger cars are usually seen as the first movers, but there

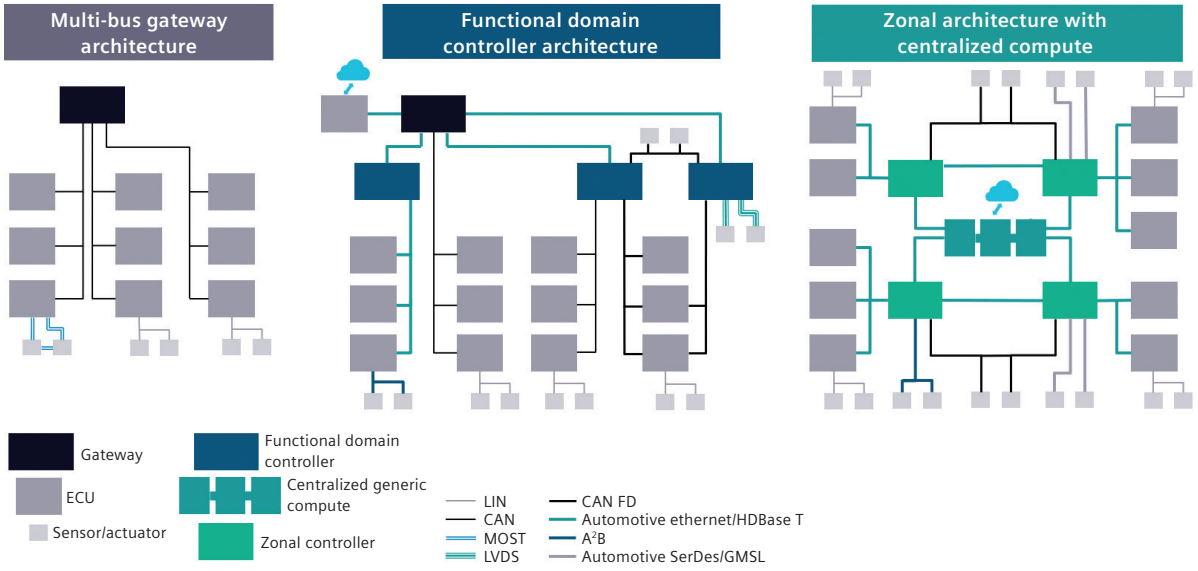

Figure 1. Centralized versus zonal architectures.

are also regional variations and new entrants that do not conform to the traditional timelines of automotive companies, and some very advanced vehicles in more specialist market segments.

Through the rest of this paper we will refer to an example to help illustrate the theory and challenges in each topic, while also highlighting the importance and benefits of full system design. Usually, these design challenges can be considered by the network designer in relative independence. But, each design decision has an impact on the full system. This impact should be considered at the time of design, allowing system testing later in the process to confirm correct behavior rather than uncover issues requiring design iterations.

The instrument cluster, central display(s) and heads-up display (HUD) are increasingly part of one integrated system, an extension of the driver information and infotainment systems. However, there is often more information in the instrument cluster with functional safety considerations. Therefore, several partitions are still required on the compute platforms hosting these functions. Such partitions may be separate processors or merely separate cores. These systems may also have discrete LEDs, switches and other peripherals connected to fully meet their requirements, though traditional automotive vehicle control switchgear is often connected to a body controller or gateway.

Besides the media functions that may be hosted in the infotainment portion of this system, there are many other pieces of critical information to convey to the driver, ranging from vehicle speed and faults, through to driving modes, ice warnings, navigation directions, estimated range and more. Some OEMs refer to the cluster as a 'combination meter', reflecting the history of bringing together what were once multiple instrumentation gauges,

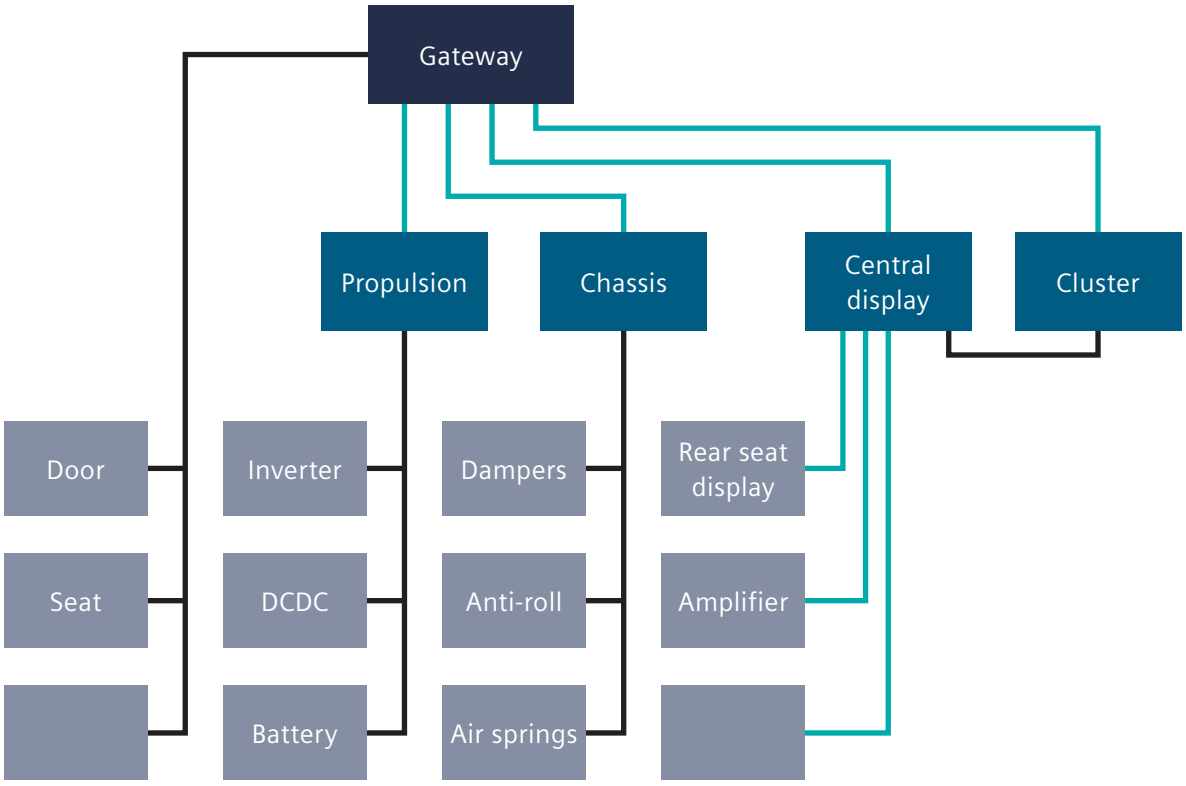

Figure 2. An example subset of ECUs in a vehicle architecture.

warning lights, and trip computer elements into one system. The domain consolidation we hear so much about today is in line with such historical examples, which also illustrate the pattern of functional growth after each consolidation. In other words, every time functions are consolidated, new functions tend to be added in new components or ECUs - from trip computers through to clocks, temperature gauges and ice warnings.

Taking the cluster example, it is common for displayed data coming from a specific control ECU or sensor to be relatively raw. Meanwhile, data from a domain controller or mode function, such as a chassis mode for off-road or sport driving, to be more heavily processed to resolve a status or adjust the context/highlighting of information.

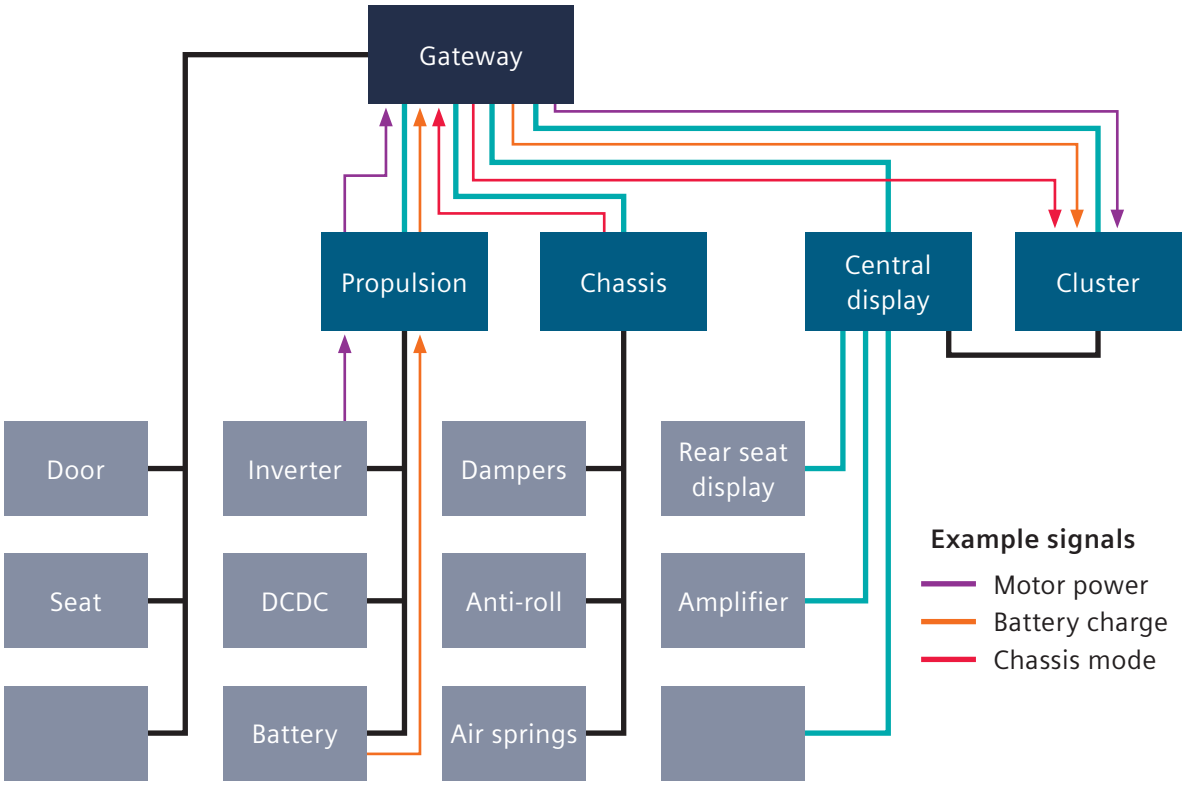

Figure 3. An example subset of signal flows.

Network load, gateway load

In the days of CAN-to-CAN gateways, the network designer had a choice between gatewaying whole frames of signals, or each signal individually, re-packed into a new frame. This provided a simple trade-off between more processing at the gateway or less-efficient use of network bandwidth. That compromise is still made today: AUTOSAR provided protocol data units (PDUs), as a design element to support gatewaying between different network technologies. This streamlines the gatewaying of PDUs, while preserving the option of reading out individual signals and repacking them into a new PDU.

This basic compromise is much more complicated in detail. The designer may need to consider how the gateway will trigger the sending of the gatewayed data it's forwarding, which in turn influences the overall latency across the system. In practice, this is often constrained by re-use of existing ECUs or, in some cases, frame and PDU designs that support an integration with a supplier across multiple vehicle programs.

Some OEMs, however, have a preference towards using an existing library of pre-designed network messages with pre-defined frame packing to support ECU re-use without changes. The use of a standardized protocol, such as SAE J1939 for heavy-duty and off-highway vehicles, provides the ability to connect vehicles and equipment of different brands together reliably. Both of these approaches reduce the scope for optimization by the designer, but don't reduce the need to consider the performance and behavior of the design.

Networks designers follow design rules, defined by the respective OEM, that consider many technical details of the system. For example, the prioritization or scheduling of each network technology has to be considered. LIN and FlexRay are time triggered networks with a schedule, though FlexRay's dynamic segments give some flexibility for event-driven data and variations in payload size. CAN uses an arbitration mechanism based on the frame ID, meaning that in most designed networks there are frame ID ranges reserved for different types of payloads, including both functional and network management needed to run most networks, also occasional data such as service and diagnostics. Higher priority is often assigned to data that would affect vehicle functions with variable jitter.

The use of a model-based design tool permits a left-shift of network validation by understanding the behavior of the network and the consequences of design decisions in worst case scenarios. Capital Networks has inbuilt models allowing the worst case consequences of design decisions to be predicted, confirming the validity of the design. Rules for frame IDs, consistency of frame packing and the correct configuration of gateways can all be checked and confirmed to be correct-by-design. Further generative automation accelerates the completion of often laborious tasks, saving time and reducing manual errors.

Ethernet and switches

Today, there is an increasing need to consider more network technologies across the architecture. Ethernet may have first appeared in the infotainment system or diagnostics over IP (DoIP) systems, now frequently expanding across domains forming a back-bone between the functional domain controllers. At this point, Ethernet design, including switch configuration, may expand from being a topic specific to the network traffic of a single domain to include more general vehicle data, e.g. data passing between traditional networks and Ethernet, benefitting from full system considerations.

Initially, using Ethernet adds another set of network behaviors to consider and a more complex set of standards and protocols. However, these are more scalable than specialized automotive networks, with the same communication software in use regardless of Ethernet physical layer type, making updates over time easier. And, as Ethernet networks can interoperate at multiple baud rates, it can be used across a large section of the vehicle, reducing the technological complexity over time. MOST is already fading from use, and was never incorporated into AUTOSAR. FlexRay and high-baud-rate CAN have a seemingly ever-reducing set of use-cases for which they are the ideal solution. And only time will show if 10Base-T1S or CAN-XL can fully take over for networks where a 10 Mbps network is utilized.

Ethernet networks introduce additional configuration options for the network designer. Protocols, methods, and elements of different levels can be used to ensure that priority data, signals, and services are available in a timely manner while allowing multiple types of data on the same physical network.

Meanwhile, virtual local area networks (VLANs) are used to segregate different types of data, and allow these various data types to be prioritized, limited (in terms of bandwidth utilization), and even disabled. A specific VLAN may, for example, be used to implement software updates allowing regulation of the bandwidth utilized for specific functions, more or less depending on the vehicle status or mode.

Network	Max baud rate	Max frame payload	AUTOSAR support	Priority/ timing	Segregation	Topology
LIN	20 kbps	8 Bytes	Y	Schedule	Physical network	Linear
CAN	1 Mbps	8 Bytes	Y	Priority arbitration	Physical network	Linear/Star
CAN-FD	8 Mbps	64 Bytes	Y	Priority arbitration	Physical network	Linear/Star
CAN-XL	10 Mbps	2048 Bytes	In development	Priority arbitration	Physical network	Linear/Star
FlexRay	10 Mbps	254 Bytes	Y	Schedule	Physical network	Linear/Star/ Hybrid
MOST25	25 Mbps	64 Bytes	N	Schedule	Physical network	Ring
10Base-T1S	10 Mbps	1500 kB	Y	AVB/TSN	VLAN	Linear
100Base-T1	100 Mbps	1500 kB	Y	AVB/TSN	VLAN	Switched flexible
1000Base-T1	1 Gbps	1500 kB	Y	AVB/TSN	VLAN	Switched flexible

Table 1. A quick reference guide to common automotive network technologies.

AVB		IEEE standard
Time Sync	gPTP	802.1AS-2011
Reservation	Stream reservation protocol	802.1Qat
Quality of service	Credit based shaper	802.1Qav
Transport protocol	AVTP	1722
Transport protocol	RTP over AVB	1733

Table 2. Audio Visual Bridging (AVB) was originally created for audio/visual media streams.

Audio visual bridging (AVB) was initially created to add specific shaping, or prioritization of audio and visual data flows on Ethernet networks, to ensure that audio and visual data could be sent across the network without pops, crackles or other distortions due to variable data rates. AVB was adopted by early automotive Ethernet users, usually in conjunction with scalable service oriented middleware over IP (SOME/IP) and service discovery (SD), as a method to enable SOA communication. Time sensitive networking (TSN) is a development of AVB specifically for functions and use-cases that have high integrity requirements, such as automotive. TSN extends some of the elements of AVB, while also adding others that were not available before.

TSN		IEEE standard(s)
Redundant Time Sync	gPTP	802.1AS-2020
Reservation	Stream reservation enhancement	802.1Qcc
Reservation	Path control and reservation	802.1Qca
Quality of service	Time aware shaper	802.1Qbv
Quality of service	Frame preemption	802.3br and 802.1Qbu
Quality of service	Cyclic queue forwarding	802.1Qch
Quality of service	Asynchronous shaping	802.1Qcr
Redundancy	Frame replication and elimination	802.1CB
Transport protocol	AVTP	1722-2016

Table 3. Time Sensitive Networking (TSN) was developed from AVB as an expansion to include necessary elements to support high integrity use cases of automotive.

AUTOSAR has either directly included or supported the above technologies and standards as they have been needed. This has avoided the need to reinvent technology that was originally developed for specific purposes. Standards and functionality needed by both Classic and Adaptive versions of AUTOSAR are standardized in the Foundation standard, ensuring compatibility and consistency.

Figure 4. The Foundation standard of AUTOSAR provides functionality needed in both the Classic and Adaptive platform versions.

In the instrument cluster example, any information displayed to the driver must be current ('current' meaning of an adequate maximum age from the original measurement or calculation), to represent the actual state of the vehicle. Examples include a fuel gauge or battery state of charge reading, or trip computer functions. The charge or fuel level reading shouldn't change quickly, and can likely be updated once a second, potentially with some damping. An energy usage reading, such as a power gauge, instantaneous fuel consumption or similar graphic, on the other hand, will likely need much more frequent updates. These elements of similar data will come from different ECUs: one from a battery management ECU or fuel tank sender and the other from an inverter or powertrain control ECU, although these are sometimes consolidated.

Network design tools are vital to ensuring the correct implementation of these protocols and standards, both traditional automotive networks such as CAN, and those utilized in, but not originating from the automotive industry, like Ethernet. Model based design tools support consistent implementations, predicting configuration problems and mismatches across the full architecture. Capital Networks facilitates the consistent design of data signals, and implementation of protocols across full

architectures, considering vehicle variants, options, and the different technologies used. Sophisticated design tools can model full E/E architecture behavior, including timing, bandwidth utilization and other performance measures. These models can also consider the characteristics of each technology, and the prioritizations and allocations made for the data types in use. Configuration of VLANs, the assignment of the relevant sub-set of ECUs to each, the AVB/TSN priorities and more can all be managed centrally in the system design, ensuring consistent and correct implementations.

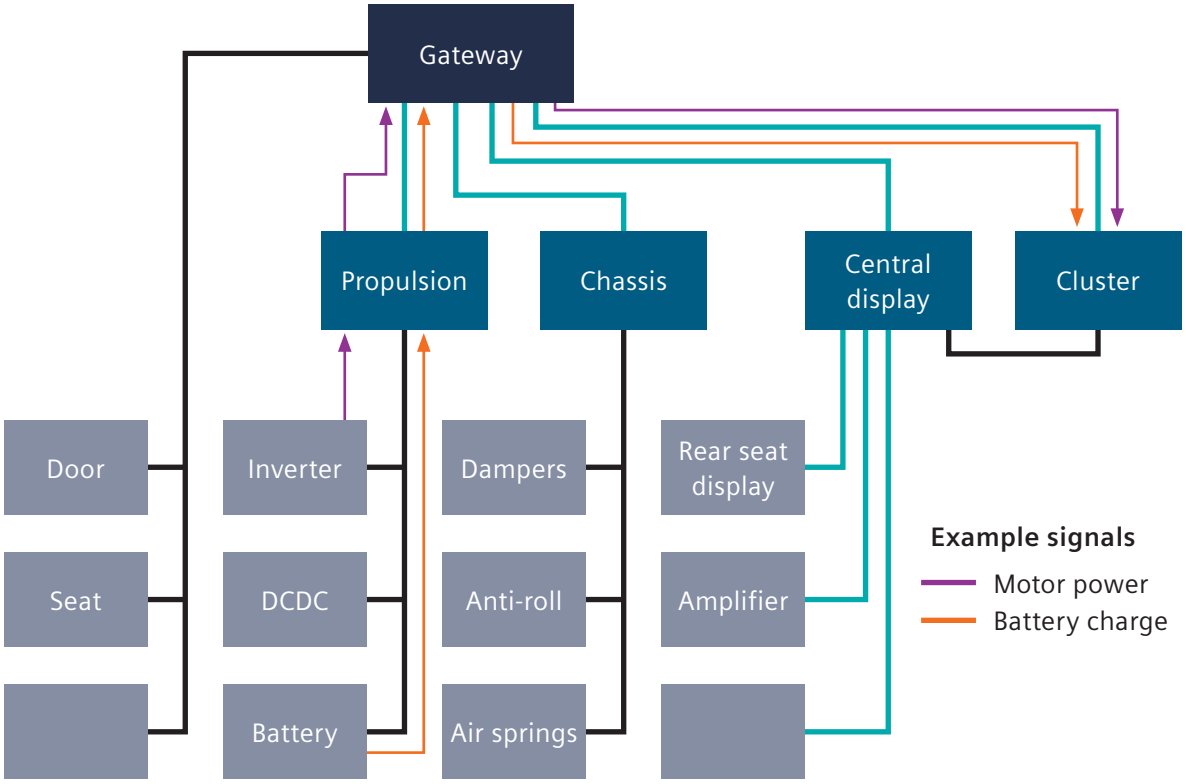

Figure 5. Apparently similar data may come from different sources.

Functional safety

Network designers have been designing with functional safety in mind for many years now, and in most cases the mechanisms used are well understood. The larger data elements and objects used with higher levels of driver assistance and automation have added updated mechanisms or schemas in recent AUTOSAR releases.

The conventional approach with networks is to treat them as quality measures (QM), as defined in ISO 26262, as a mechanism, and thus add elements to the design to validate that the data is being received regularly and accurately. Increasing system integrity requirements now demand redundant routings for some data, but this is often a system-level design consideration, and thus will only be encountered by the network designer as additional design rules. For instance, a design rule may prevent the routing signals A1 and A2 on the same networks from sender to receiver.

Data that carries a potential safety consequence if incorrect or missing primarily receives end to end (E2E) protection, where by a group of signals are packaged in a common message or PDU, and treated as a single entity in terms of the network bus, gateways and COM stacks. These grouped signals have a cyclic redundancy check (CRC) calculated for them, some form of counter (alive, frame or other depending on the scheme chosen), and (often) a data ID, although other methods can be used to overcheck the identity of the group. These protection methods are identified as schemas by AUTOSAR, which has included common mechanisms used to provide the protections, including the CRC calculations. This allows OEMs and systems integrators to shape their own design rules according to the risks identified in their system design methodology.

The network designer groups signals based on the functional requirements identified in the systems design phase and structures these groups in the network design, or re-includes the groups in the case of carry-over from existing projects. Documentation demonstrating that the design fulfills requirements, rules and standards in place supports auditing of the application of E2E protections. These mechanisms are set by the sender and used by the receiver to confirm that the data is fresh, valid, and from the correct sender. The integrated protections, however, are only designed to defend against errors and faults, and not targeted malicious actions. Thus, the system design has to be robust enough to cope with potentially correct data occasionally being rejected, or, invalid data being accepted, both infrequently and usually as single occurrences, but over thousands of hours of usage of millions of vehicles, these infrequent events occur.

For E2E protection to be meaningful, both the sending and receiving ECU need to be designed with appropriate functional safety considerations. Protecting a signal which may have been sent incorrectly due to software of inappropriate integrity is meaningless and a waste of bandwidth. Further, the frame headers of most network types contain some protections that will cause the data in the frame to be rejected should a fault be detected. But, these usually are not considered to be robust enough for application level data checks, and are only suitable for bus level errors.

The information presented to the driver is, by its nature, of mixed criticality. Some elements are safety related or legislated, requiring continuous availability with sufficient accuracy. Other elements of information are presented simply for the comfort or convenience of the driver. The simultaneous need to present both safety-critical and non-safety-critical information introduces requirements not just into the processing and presenting of such data, but also the networks design. Sometimes this includes redundant information sources, though, more commonly, the requirements call for E2E signal

protection, or even both. Examples may include airbag warnings, which, if the cluster or warning lamp system misses the data from the airbag controller, will put the warning lamp on assuming there to be a fault. In contrast, some of the data in a trip computer function is for convenience only, and can safely not be shown when not available. Likewise, other than damaging the impression of vehicle quality, erroneous values may be accepted by design.

As with general network design, OEMs have design rules and standards instructing the network design engineers which protections to use in which system design scenarios. Capital Networks includes editors that assist with the set up these safety mechanisms. The design model enables consistency checks to guide the designer towards problems and ensure that the networks are correct-by-design and consistently described in the software configuration outputs for each ECU.

Cyber security

There are similarities in the approaches taken with functional safety and cyber security, in that the base technology is built upon, and in terms of network design, signal protection elements are added. However, due to the need to mitigate against malicious attacks, the mechanisms added need to be more sophisticated. Even still, these mechanisms may not always meet the needs of the functional safety on their own and OEMs must consider the actual risk when selecting the mitigations to be used across a system or platform for potential threats and/or faults. Aligning with functional safety, there is a standard developed (ISO/SAE 21434) to set out best practice principals and process when designing vehicle systems in consideration of cyber security.

To satisfy functional safety, received data is checked to be consistent and correct with what was sent, with a limited check that the signal group is correct. Cyber security includes additional checks to authenticate that the data is from the correct sender, and sometimes (depending on the systems design) includes encryption of the data itself, though both generally are not needed together.

Introducing new challenges for the network designer, modern vehicle systems can exchange data such as phone numbers, addresses, payment details and more. These types of data are, or contain, personally identifiable information (PII). This data needs to be encrypted both during transmission and in storage, and thus encryption keys are also needed to write and read the data.

Meanwhile, data used to determine control decisions with safety relevance needs to be trustworthy. In some cases, the overall system design may contain sufficient redundancy in the sourcing or sensing of this data that full protection is not needed on every element, and a fusion algorithm may be used to resolve conflicts. It is also possible that this part of the system design is constrained by the sourced system components and network technology available (bandwidth, maximum PDU size, etc). Eliminating or reducing these constraints is a primary driver to higher baud rate networks with larger payloads per frame.

The control data coming from the decision algorithm, which may be instructions on control inputs for steering, acceleration, braking and more, has a direct impact on the vehicle behavior. The system design has to assure that this data is correct, and, thus, authenticating control data at the target motor or actuator is highly desirable. Possible authentication mechanisms include a hashed (#) version of the signal group that enables the receiver to perform an additional keyed check of the data. Control data may also include some indicator for

time, or a step counter. Counters used in automotive networks for functional safety are often 4 or 8 bits, while cyber security counters can be extended to 16, 32 or more bits to boost the mitigation against a replay attack. It's common to use multiple protections, as per functional safety, to mitigate different risks. Redundant copies or paths for the data can help, however, determining which data to trust when a conflict occurs is an important design consideration.

These protections, and any others, have a bigger impact to both the network designer and the architect than the functional safety protections. With cyber security, data sizes tend to be much larger, consuming more bandwidth. The larger frame sizes of higher baud rate networks such as Ethernet, FlexRay, CAN FD, or CAN XL can better support the greater demands of cyber security, as opposed to traditional CAN or LIN networks. An additional firewall before a CAN or LIN network, with some additional health monitoring of the sub-systems from the secure ECU, can provide additional protection to vulnerable parts of the system.

VLANs can be used to segregate traffic types on Ethernet and IP networks allowing bandwidth utilization limits to be set. This can help prevent denial-of-service-type attacks from affecting multiple systems, and enables certain traffic types to be turned off in various modes of operation. For example, software updates can be prevented while the vehicle is in motion. Additionally, firewalls are increasingly deployed at entry points to the vehicle, such as the telematics, and between the entry points and ECUs hosting higher risk functions.

A further consideration is that, while functional safety is a systems design challenge that requires appropriate software, hardware and point-to-point design of the data flows, it can usually be considered in relative isolation. In contrast, cyber security protections manifest as built-up layers of defense across the platform, its cloud connections and more. Special care must be taken to ensure all appropriate layers are in place for systems determined to be at risk.

If we consider the cluster example, functions that relay phone or navigation information onto the driver display might contain PII. This data typically should be encrypted so that it is not readable from a data log without the correct key. This data, however, usually does not warrant protection from corruption such that duplicate information is displayed to the driver. The encryption of data containing PII may be more important in markets with privacy legislations such as Europe's GDPR protections, but is likely good practice overall.

As described previously, the consequence of this is additional design rules and standards that must be followed by the network designer. Capital Networks models the full vehicle E/E architecture and thus can perform consistency checks across the full system ensuring that the design is consistent and that, where added, the cyber security protections are complete, ensuring correct-by-design outputs.

Power modes

Traditionally, vehicle networks have been designed to remain all awake to ensure functionality is available when needed. Under this approach, special attention is paid to designing robust shutdown procedures that occur when the vehicle is in an appropriate state. This method sustains safety-related functions and backup functionality to, for example, enable parking brakes or maintain limited powertrain operation in the event of a faulted network. For maximum energy efficiency, however, it is desirable to shut down what isn't currently needed, and what won't be needed on shorter notice than the wake time of the components which are asleep or powered down.

Partial networks allow some networks to be shut down when not needed. Pretended networking is also used occasionally, where some ECUs go into a low-power mode, but continue to be active on the networks. Power modes can, and do, get more complicated. It is very important that needed signals and data can be generated by awake ECUs, using awake sensors, and sent over networks that are awake. Power modes can therefore quickly constrain the routing of signals.

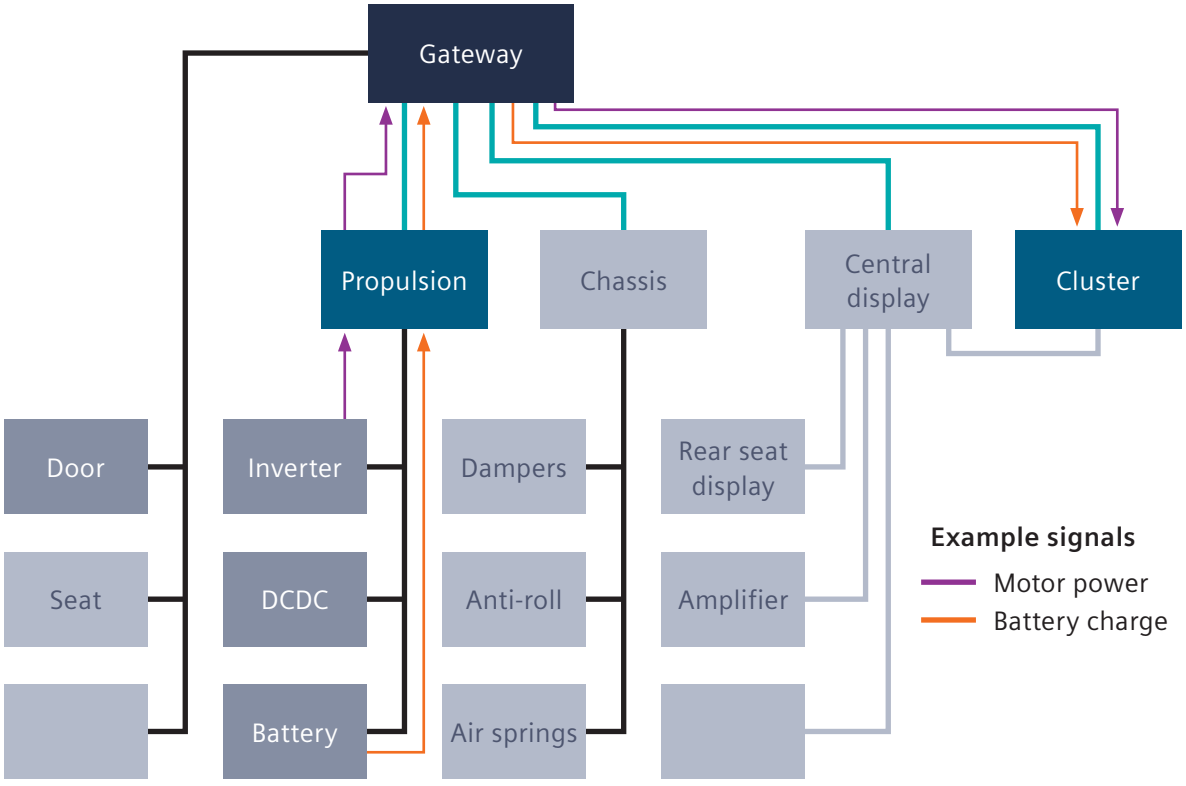

Figure 6. During some modes, such as vehicle charging, only a sub-set of the full display information is needed and available.

Returning to the cluster example, the user has an expectation that certain data is visible in given vehicle modes, such as charging a battery electric vehicle. This data might be offered for the convenience of the driver, but could also have legal or safety considerations, such as the current charging status, vehicle odometer or parking brake status. This means if this data is not calculated or stored in the cluster, it needs to be available to be displayed. While this usually will have been considered in the E/E architecture definition, it still creates network routing and signal packing requirements for the network designer.

Capital Networks includes editors that assist and automate setup of these power saving modes, with model-based consistency checks ensuring correct-by-design setup of mechanisms which can be very complicated to implement.

Complexity

Finally, complexity must be considered at a full system level as it is impacted by everything discussed thus far. In this paper, complexity means options and variants. Most vehicle programs share a common underlying E/E architecture across a range of vehicles of different sizes, body types, for different markets and more. A two door vehicle may use fewer door modules than a four door car, and a low-specification car will likely use fewer ECUs overall than a high specification car, driven by features included in the vehicle. All of this is to be considered in the design. Some signals must be available on all variations, while others may change source due to being calculated or measured differently for different vehicle types. A vehicle speed algorithm, for instance, will consider different wheel slip behavior on two and four–wheel-drive vehicles. The mechanisms covered for functional safety and cyber security also need to consider the relevant vehicle variations.

Capital Networks allows full consideration of vehicle and ECU variants, using consistency checks to ensure correct-by-design implementations. Implementation files, AUTOSAR configurations or other, can be shared with software teams and suppliers appropriate to the ECU versions and variants needed.

Summary

In this paper, we have discussed a wide range of day-to-day challenges and considerations that occur during network design phase of E/E systems development. Each of these challenges and decisions can have widespread, cross-domain effects that are difficult to predict or even fully understand. Connecting the many disciplines enables designers to understand the downstream impacts of their decisions during development, and is critical to accelerating the vehicle development process. Networks design considers and implements many elements vital to both ensuring correct vehicle functionality and protecting the entire system from incorrect sub-system behavior. It is important to select a solution which consistently and correctly generates the configurations and documentation used in the development and validation of each ECU making up the full system.

Capital Networks has been developed to address the specific needs of the design of vehicle networks. Bringing together learnings from its predecessors, which were used by multiple OEMs across the world, and the AUTOSAR flow and framework, which also includes many years of industry learning, to offer the most robust networks design solution. Capital Networks is a model-based design solution, offering generative design capabilities, that ensures efficient design of performant networks across the multiple interdependent complexities of modern E/E architectures, used across multiple vehicle platforms. Built-in consistency checks use design rules and models to ensure design correctness, guiding the user to the any areas needing attention. Correct-by-design outputs in AUTOSAR or other formats are generated to configure each ECU, or validate whole networks, including the needed elements for functional safety and cyber security.

Siemens Digital Industries Software

Americas: 1 800 498 5351

EMEA: 00 800 70002222

Asia-Pacific: 001 800 03061910

For additional numbers, click here.

About Siemens Digital Industries Software

Siemens Digital Industries Software is driving transformation to enable a digital enterprise where engineering, manufacturing and electronics design meet tomorrow. Xcelerator, the comprehensive and integrated portfolio of software and services from Siemens Digital Industries Software, helps companies of all sizes create and leverage a comprehensive digital twin that provides organizations with new insights, opportunities and levels of automation to drive innovation. For more information on Siemens Digital Industries Software products and services, visit siemens.com/software or follow us on LinkedIn, Twitter, Facebook and Instagram. Siemens Digital Industries Software – Where today meets tomorrow.

siemens.com/software

83770-D3 6/21 H