

Reasons to use CAN XL

Kent Lennartsson has explored engineering-driven reasons, why CAN XL (extended data-field length) is gaining traction across automotive and industrial applications. He explains these reasons in detail, giving background information on architectural and practical advantages that address real-world design challenges, from cabling cost and robustness to deterministic latency and functional safety.

(Source: Adobe Stock)

1st reason: CAN XL can transfer unfragmented Ethernet frames

CAN XL features data fields up to 2048 byte, enabling an unfragmented Ethernet frame to be placed in the CAN XL data frame. This facilitates the use of TCP/IP (Transmission Control Protocol/Internet Protocol) and UDP (User Data Protocol) services. Background: CAN CC (classic) was designed for fast, predictable, and real-time control. This is achieved by using short data frames that ensure high-priority frames can get fast access with low latency and predictable bandwidth. This is perfect for real-time control, but inefficient for service tasks such as ECU (electronic control unit) reprogramming or diagnostics. It is necessary to use a CAN-mappable transport layer protocol to cut the long TCP segment into 8-byte pieces to map them into a CAN CC data frame, and at the receiver assemble the data back into a TCP segment. This is a relatively long process, and if it fails, it is necessary to have a method to recover from the fault. Even with CAN FD (flexible data rate), protocol overhead can approach 50 percent, meaning half of the available bandwidth is consumed by framing rather than payload.

2nd reason: CAN SIC XL transceivers are an integrated solution

CAN SIC (signal improvement capability) XL transceivers are a robust technology that can be connected directly to CAN wires without any external components. The 10BASE-T1S transceivers are designed with complex 1,8-V logic, that typically require a bus interface network to protect the physical layer components. They also require a 25-MHz clock with ± 100 -ppm accuracy to run the internal logic. This adds cost and PCB (printed-circuit board) area and increases the design and validation effort.

Table 1: Comparing physical layer space requirements and costs (Source: Kvaser)

Parameter	CAN FD/XL	CAN CC	T1S driver (1)	T1S + EMC (2)
Comment area	9 mm ²	9 mm ²	16 mm ²	16 mm ²
Oscillator	n.a.	n.a.	8 mm ²	8 mm ²
Bus IF network	n.a.	n.a.	9 mm ²	9 mm ²
EMC choke	n.a.	n.a.	n.a.	9 mm ²
EMC varistors	n.a.	n.a.	n.a.	12 mm ²
Total area	9 mm ²	9 mm ²	60 mm ²	81 mm ²
Cost (10 nodes)	US-\$ 7,3	US-\$ 3,3	US-\$ 27,5	US-\$ 33

(1) Requires at least three capacitors, three resistors, and an oscillator
 (2) Includes (1) plus two varistors and a common-mode choke

As the network wires directly connect to the CAN transceiver, circuit complexity between the CAN transceiver and the CAN wires is significantly reduced.

3rd reason: CAN SIC XL transceivers are intrinsically robust

CAN SIC XL transceivers are robust, meeting most EMC (electromagnetic compatibility) requirements. There are CAN CC transceivers available for a more extreme environment – ± 80 V on the CAN wires. It is only in exceptional cases that you need to add external components to protect the CAN transceivers.

Table 2: Comparing EMC requirements (Source: Kvaser)

Test type	CAN FD/CC (1)	CAN XL (2)	T1S driver (3)	T1S + EMC (4)
Max. voltage	± 65 V	± 42 V	-27 V to +42 V	± 65 V
ESD (human-body model)	± 45 kV	± 10 kV	± 8 kV	± 45 kV
ESD (IEC air-gap model)	± 30 kV	± 15 kV	unknown	± 30 kV
ESD (IEC contact model)	± 12 kV	± 8 kV	± 1 kV	± 12 kV
Common-mode range	± 25 V	± 25 V	± 32 V	± 32 V
Cost (10 nodes)	US-\$ 7,3	US-\$ 3,3	US-\$ 27,5	US-\$ 33

(1) MAX3301E transceiver by Analog Devices
 (2) TCAN6062 by Texas Instruments
 (3) Requires at least three capacitors and three resistors (36 mm²)
 (4) Includes (1) plus two varistors and a common-mode choke

In practice, EMC filters are often added to the CAN XL driver to compensate for poor wiring-harness design, while a good cable layout will save ECU cost and increase robustness and vice versa.



CAN XL is gaining traction across automotive and industrial applications. There are many reasons to use CAN XL.

Kent Lennartsson (Research Manager at Kvaser)

4th reason: CAN SIC XL transceivers are cost effective

The CAN transceiver is an efficient power amplifier that provides good signal integrity without any complexity. All necessary complexity is integrated in the CAN XL controller.

In contrast, 10BASE-T1S integrates significant real-time logic directly into the PHY, resulting in a transceiver comparable in complexity to a CAN XL controller.

- ◆ CAN XL needs 2 pins, not 14 pins for an xMII interface.
- ◆ CAN XL uses robust logic design compatible with legacy production facilities.
- ◆ CAN SIC XL transceiver does not need a 25-MHz, $\pm 0,01\%$ oscillator.
- ◆ CAN XL controller operates with a $\pm 1,5\%$ oscillator tolerance.

Table 3: Comparison of physical layer component features (Source: Kvaser)

Component	Type	Price [US-\$]	Interface	Package	Space [mm ²]
LAN8671 (1)	10BASE-T1S	2,6	xMII, 14 pin	VQFN-24	16
TJA1462 (2)	CAN FD	0,73	TX/RX, 2 pin	HVSON-8	9
NCV7357 (3)	CAN CC	0,33	TX/RX, 2 pin	DFNW-8	9
Oscillator	n.a.	0,1	4 pin	n.a.	8
Choke	n.a.	0,8	4 pin	n.a.	9

(1) Microchip
(2) NXP
(3) Onsemi

CAN transceivers are based on mature, robust technology well suited for the electrically harsh conditions found in automotive and industrial environments. While 10BASE-T1S PHYs benefit from high-density semiconductor processes, these processes are inherently more expensive. In addition, the wide xMII interface, external oscillator, power pins, and EMC components significantly increase PCB area, package size, and bill of materials cost. Although higher production volumes can reduce unit price, a larger pin count and increased system complexity will always carry a cost penalty.

5th reason: CAN XL supports low-cost cable layouts

With CAN XL, it is possible to match cable cost to the required performance. CAN XL supports bit rates from 50 kbit/s up to 20 Mbit/s, allowing the bit rate to be selected according to the cable layout. For example, CAN XL communication at 125 kbit/s will work reliably in a star topology with stub length of several meter. CAN XL can operate at up to 20 Mbit/s when the bus layout is designed to support higher bit rates.

In contrast, 10BASE-T1S or 100BASE-T1L operate at fixed bit rates of 10 Mbit/s resp. 100 Mbit/s and require a cable layout that can support such high bit rates. This cost is incurred even when bandwidth utilization is very low.

Higher bit rates also demand better control of impedance variations along the bus. Most twisted wires in a cable have continuous impedance by design.

The main sources for impedance variations are:

- ◆ Star points, multiple drop lines close together
- ◆ Relatively long drop lines
- ◆ Connectors
- ◆ EMC filters
- ◆ Poor termination

To optimize cost, it is necessary to optimize the ECU design and bit rate according to the cable design.

How the cable affects the signal: When the transmitter changes its output voltage, it also changes the current flowing into the twisted pair. The resistance of the cable is very low, so one might expect a high current. If you have a twisted pair that is 3 m long with AWG24 wires, there is a resistance of 85 m Ω /m. The total resistance in one wire over 3 m is 250 m Ω . If you make a shortcut at the other end, the total resistance will be 0,5 Ω . A 1-V_{DC} step would therefore suggest a current of 2 A. This is only true for DC; with fast voltage transitions, the cable behaves as a transmission line with a characteristic impedance of about 100 Ω , which limits the initial current. The voltage step propagates along the cable and is reflected at impedance discontinuities, inverting at a short circuit and doubling at an open end. To avoid repeated reflections, the transmitter's output impedance is matched to the cable, causing most signal energy to be dissipated in the driver rather than the cable. As a result, a stable communication link requires point-to-point connections with matched impedances at both transmitter and receiver.

6th reason: CAN XL supports flexible drop line lengths

Communication is all about transporting energy over a cable. To switch your light bulb on and off is straightforward, but switching a communication signal one billion times per second is challenging for both the switch and the cable. At 1 Gbit/s, each bit is 1 ns long and covers about 20 cm of cable. For a 10-m cable, the sender will transmit 50 bits before the first bit reaches the receiver. If the communication link has impedance variations, the energy will be reflected into neighboring bits, a phenomenon known as Inter Symbol Interference (ISI). CAN uses arbitration, which requires a bit-length to be at least four times the signal propagation delay of the cable. With such long bits relative to the length of the communication wire, ISI becomes negligible because all energy from previous bits has dissipated before the sample point is reached. For this reason, CAN CC was inherently protected against ISI. With CAN FD and CAN XL, higher data-phase bit rates are used after arbitration. In these cases, the bit length will be close to the delay or even shorter than the propagation delay of the cable. As a result, the bus layout must be designed to cope with shorter bit durations.

As described earlier, signal degradation is caused by reflections due to impedance variations. If the physical length of an impedance discontinuity is less than one-tenth of the signal wavelength, its influence on signal quality is minimal. ▶

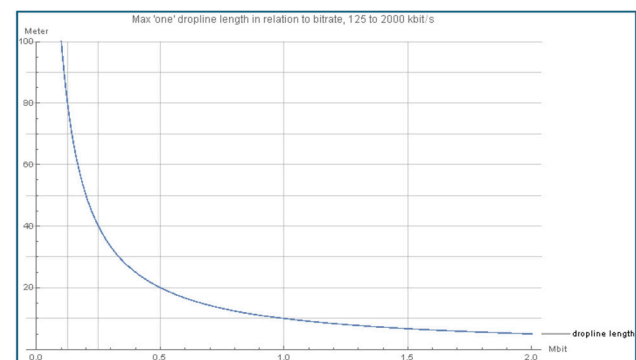


Figure 1: Maximal drop-length in relation to the bit rate (Source: Kvaser)

Figure 1 represents a conservative, safe total drop-line length, and you can read:

- ◆ At 125 kbit/s, drop lines of up to 80 m are acceptable.
- ◆ At 1000 kbit/s, it is possible to have a 10-m drop line.
- ◆ Doubling the bit rate requires the allowable drop line length to be halved.
- ◆ The drop lines can be longer if they are of equal length and evenly distributed along the main bus line.
- ◆ Shorter limits apply if drop lines are concentrated in a star configuration.

Following these guidelines will ensure you remain on the safe side and you will have a very nice eye-diagram. It is possible to violate these limits, but each violation will make the eye smaller and increases the risk of bit errors. With the introduction of CAN FD and CAN XL with a bit rate above 1 Mbit/s, drop line design has suddenly become a critical concern. At 20 Mbit/s, the drop line should be less than 0,5 m. For 10 nodes, that would be a 5-cm drop line each. It can be longer, if the nodes are evenly spread, or by analyzing and testing the effect of the cable layout. The 10BASE-T1S communication is due to the 8b/10b coding, which is actually 12,5 Mbit/s, but the analog signal is optimized for this specific frequency. This makes it possible to have longer drop lines compared to CAN XL at 12,5 Mbit/s. It could be possible to make a CAN transceiver optimized to one single bit rate, but that will prevent such component from being used at any other bit rate.

7th reason: CAN XL provides inherent real-time control

Arbitration is a simple, built-in, and efficient mechanism to secure predictable real-time behavior:

- ◆ The highest priority data frame has guaranteed latency.
- ◆ By limiting repetition interval on each priority level, all priorities can have guaranteed latency.
- ◆ Work the same from 2 nodes to 200 nodes.
- ◆ Additional nodes with lower priority can be added or removed without modification of existing nodes.

Achieving the same predictability with Ethernet requires TSN (Time Sensitive Networking), which is a logic that ensures that Ethernet frames are sent in a controlled order without collision, guaranteeing delivery within a certain latency.

However, TSN (Ethernet) switches with six ports are complex and costly. One TSN switch is equivalent to more than 20 CAN XL controllers and typically relies on a microcontroller with a 500-kbyte code. This cost is acceptable for the central computer, with or without AI, to help a driver control the car. For distributed I/O-modules, a less sophisticated solution is a multi-drop PLCA (Physical Layer Collision Avoidance).

As indicated, 10BASE-T1S is similar to CAN: multiple nodes share a single communication medium. This has clear advantages for synchronized broadcasting communication, as every receiver has the exact same information at the same time. The challenge is that there must be a set of rules to ensure that package transfer is not delayed by collision. ▶

CAST



CAN Controller IP Core

COMPLETE

CAN CC, CAN FD, CAN XL plus TTCAN
AUTOSAR & SAE optimization

SECURE

Optional
CANsec

SAFE

Designed for FuSa:
ASIL-D Ready Certified

RELIABLE

Hundreds of Successful
Customer Deployments
Industry-Leading Support

FLEXIBLE

ASICs or FPGAs; Works with any Transceiver

www.cast-inc.com

When using a common media, there are two approaches:

1. Orthogonal channels (frequency, code, or physical separation).
 - a. Easy to use, but expensive.
2. Common channel with base band signaling, and time-division between different users. The division of time can be done in several ways.
 - a. Fixed timeslots for different users;
 - b. Token passing.

(Both a and b demand some kind of organizer that controls all members of the media.)

Another approach is to use CSMA (Carrier Sense Multiple Access), of which there are two variants:

- ◆ CSMA/CA (Collision Avoidance) – Used in classic Ethernet, nodes wait a random time after collisions. This works well at low utilization but becomes unpredictable at high load due to increasing collision rates. This is the fallback solution if the PLCA scheduling breaks down.
- ◆ CSMA/CR (Collision Resolution) – Used in CAN, arbitration occurs bit by bit. Nodes with lower priority drop out, allowing the highest priority message to continue without data loss. The only requirement is a careful CAN-ID assignment, which is comparable in complexity to assigning IP addresses.

Modern Ethernet (>10 Mbit/s) typically uses point-to-point full duplex links, eliminating collisions on the physical link. However, this effectively moves collisions to the switch. When Ethernet frames arrive simultaneously and must exit through the same port, a resource conflict occurs, whereby one frame must be sent before the other frame. In a standard office switch, this is solved with a simple FIFO, where the first frame that arrived gets through first. In a real-time system, it is necessary to control the latency for all critical information. This makes it necessary to add some rules that prevent low priority information from delaying important information. This is solved by placing a TSN logic and configuring this logic so that high priority Ethernet frame reaches the destination member with lowest possible latency.

If you fail to assign CAN-IDs to get a working real-time control system, you will find that configuring TSN is not simpler.

TSN is relatively expensive and needs complex configuration tools. To remove TSN complexity and cost, real-time logic can be included in the 10BASE-T1S transceiver. Node-ID “0” acts as the coordinator and sends a beacon. After this beacon, each node has an opportunity to send Ethernet frames, as determined by the node-ID order. If a node has nothing to send, it passes its opportunity to the next node.

To make the system work, it is necessary to set up the PLCA in every installed node. Optimal performance is achieved when all nodes from node-ID 0 to N are present, forming a continuous round-robin sequence. Missing nodes are tolerated, but a full sequence is preferred.

Following configurations must be done in each 10BASE-T1S transceiver to have a working PLCA:

- ◆ Set the unique PLCA local ID (ID = 0 will send beacon).
- ◆ Set the PLCA node count, maximum number of nodes.
- ◆ Set the PLCA transmit opportunity timer register. All units must have the same value.

- ◆ Set the PLCA maximum burst; this sets the maximum number of frames to be sent.
- ◆ Set the PLCA burst timer, limiting the time for the burst.

There are several other registers to control the device, status information, handling events, and to get diagnostic information. If node-ID = 0 is lost, the beacon will be lost, and the communication will fall back to classic CSMA/CA.

8th reason: CAN XL supports adding or removing modules without real-time interference

Once the real-time system has been tested and all latency requirements have been met, new CAN data frames can be added, if the priority level (CAN-ID) is lower than the CAN-ID of any CAN-data frame with strict latency constraints. It is possible to design the system for future high-priority CAN data frame, but this requires those data frames to be simulated during all tests to ensure that high-priority traffic continues to function correctly under the additional high-prior CAN traffic.

9th reason: CAN XL provides inherent consistency and time-synchronized data distribution

Every CAN data frame is broadcasted to all nodes connected to the common CAN bus. A valid CAN data frame will ensure that all installed nodes receive the same information at the same instance in time, thus providing natural time alignment across nodes. Ethernet demands time stamping and acknowledgement to achieve this. This time stamping induces a cost in terms of latency and overhead (additional data in the package) and to achieve acknowledgment, it requires a second frame in return, which also takes bandwidth and time to return this information to the data producer.

CAN provides a simple solution for free that is good for most control systems. The same function can be achieved with 10BASE-T1S by using broadcast, ensuring that all connected units receive the same Ethernet frame at the same time. If a switch exists between the sender and receivers, synchronization becomes more complex as it is necessary to configure TSN to ensure that sending to all receivers starts in the same instance in time. A simpler solution would be to synchronize the data utilization to a certain time event when you know that all receivers have received the information. This requires that the system has a global common time, which is useful in many other cases, for example, to compensate for latency in the control algorithms.

10th reason: CAN XL enables fast and secure functional safety

CAN has been designed to deliver real-time performance and high reliability with minimal logic. The effort required to analyze and to verify communication logic grows at least linearly with protocol complexity. As shown in Table 4, only very simple interfaces such as SPI (Serial ▷

Peripheral Interface), I²C (Inter Integrated Circuit), UART (Universal Asynchronous Receiver/Transmitter), and LIN (Local Interconnect Network) require less logic than CAN.

For UART and LIN, this reduced hardware complexity comes at the cost of real-time software involvement: protocol timing, arbitration, and error handling must be managed in software.

Table 4: Comparison of communication node implementations (Source: Kvaser)

Interface type	LUT (a)	BRAM (b)	RAM (c)
UART	600	1	100
LIN	800	1	100
CAN CC	1800	1	100
CAN FD	2300	1	300
CAN XL	3500	8	8000
CAN CC safe	7000	1	100
CAN FD safe	8300	1,5	500
CAN XL safe	12000	13	12000
PLCA (1)	1200	n.a.	100
Ethernet MAC	6000	9	9000
CANsec (2)	9000	n.a.	n.a.
MACsec (3)	18000	n.a.	n.a.
TSN 1 simple	2000	12	12000
TSN switch (4)	40000	355	355000

(1) plus Ethernet MAC
 (2) designed for CAN
 (3) can be used also for CAN
 (4) plus microcontroller
 (a) look-up table
 (b) block RAM
 (c) random access memory [byte]

11th reason: CAN XL provides inherent detection of alien attacks

There are several CAN protocol features allowing to detect security issues:

- ◆ Every CAN data frame has a priority level (CAN-ID) that is owned by a single node in the system.
- ◆ All CAN data frames are sent by broadcast on the network.
- ◆ The owner of a CAN-ID will detect, if this CAN-ID is misused by an alien node or alien software in some other legal node. This detection can be reported and used to prevent the effect of this misuse of communication.
- ◆ In real-time systems, CAN data frames are sent at a certain time interval.
 - ◇ A receiver can easily detect if the time interval is too short or too long.
 - ◇ This is always a functional problem, even if this is caused by some internal problem. An attack on signals is a signal quality problem, independent of the source (internal or from an alien).
- ◆ All advertised automotive hacks worked because the system ignores this inherent protection in the CAN protocol as described above.



Same Specs. New Design.

#1 Choice for Engineers has a new color.

As part of the HMS Networks group, we have changed our visual appearance. What stays are our core values continuity, reliability and transparency for our customers.

We keep on developing Hardware and Software for Automotive, Medical and Industrial Communication—for Engineers by Engineers.

www.peak-system.com



Hardware
Meets
Software™