Solutions of CAN and CAN FD in a mixed network topology

Scott Monroe, David Stout, and John Griffith, Texas Instruments

While CAN with Flexible Data Rate (CAN FD) promises to revolutionize the data rates and data frame lengths in CAN networks, the lack of interoperability of CAN nodes with CAN FD nodes limits the mixing of these types of devices on the same network. There are several possible solutions allowing the mixture of CAN FD nodes into existing CAN networks. These solutions have varying levels of system impact and trade-offs. Some of these solutions are described and explained.

Controller Area Network (CAN)

Controller Area Network (CAN) is very well established across many industries with hundreds of millions and possibly billions of control units already deployed to the field. These units are based on ISO11898-1 CAN data link layer (original CAN protocol, 2003) and ISO 11898-2 and ISO 11898-5 CAN standards for CAN high-speed medium access unit or physical layer (transceiver). Various application and protocol level abstractions including CANopen, DeviceNet, NMEA 2000, ARINC 825, MilCAN, SAEJ 2284 and ISO 11783 have been built on these foundations. For this paper a control unit based on the original ISO 11898-1 CAN protocol will be called a CAN node. A control unit that uses the proposed CAN with Flexible Data Rate protocol will be called a CAN FD node.

CAN with flexible data rate challenge

CAN with Flexible Data Rate (CAN FD) was introduced in 2011 and 2012 by Robert Bosch, GmbH. This introduction sparked renewed interest and energy within the CAN world and created some controversy as the new protocol is not directly interoperable with the original CAN protocol. Since CAN FD designed systems may operate using the original CAN protocol, successful integration into original CAN networks is possible. However, if CAN FD communication is exposed to a CAN node, the CAN node will see the CAN FD frame as an error and block the communication, thereby rendering the network useless. By using original CAN protocol on the CAN FD nodes all advantages of CAN FD are lost. Other than starting to “future proof” for a full CAN FD network there are no tangible benefits with this approach.

Thus the big challenge is how to mix CAN FD nodes into existing systems and effectively use their advanced capability. This requires added layers of complexity and CAN nodes capable of “ignoring” CAN FD in an efficient way until all microprocessors and nodes may be converted to CAN FD.

Mixed CAN and CAN FD proposals

There is a lot of work on going to standardize the CAN FD protocol and transceiver performance for FD data rates, but little discussion in how to prepare current CAN networks for mixed use. Depending on the application, different solutions and trade-offs may be required. Four options providing this capability are described below:

1. Software method utilizing silent mode with current transceivers.
2. Partial Networking with FD Passive Transceiver
3. CAN FD Blanking Transceiver
4. Managed Hub

These methods have varying impact to current software and hardware designs. System designers need to assess the trade-offs of mixed CAN and CAN FD versus implementation complexity, cost, latency, and bandwidth improvement.

Software with silent mode method

Many CAN transceivers today offer a silent mode that disables the driver but leaves the receiver active. This allows the CAN node to listen to or monitor the bus. However, it cannot transmit or disrupt the communication flow on the bus since the driver is disabled making the unit silent or mute.

A simple method of mixing CAN FD into CAN nodes is built primarily in software utilizing this very simple feature of existing CAN transceivers. Due to the simplicity, the benefits are limited to systems where long periods of time, such as a software update, the bus may be dedicated to CAN FD message traffic or where a very short burst of CAN FD traffic is needed. The basic application hardware is shown in Figure 1.

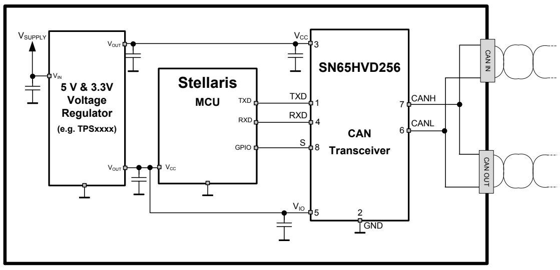

Figure 1: Software & silent application

The application software utilizes the silent mode of the transceiver to prevent CAN nodes from corrupting bus traffic during CAN FD communication. A CAN FD node will broadcast a special CAN message directing all CAN nodes to switch to silent mode on their transceivers and wait for a second CAN message to re-enable the transceivers into normal mode. This blocks error frames generated in CAN nodes from reaching the bus since the drivers in their transceivers are turned off.

The CAN FD nodes now have the freedom on the bus to transmit CAN FD frames without corruption. If enough messages are sent, the CAN nodes will go bus off due to an overflow of their error counter. The CAN FD node has the responsibility in the system to return the bus to normal CAN communication. If the number of CAN FD messages were below the threshold of sending the CAN nodes “bus off”, the CAN FD node may send the second special CAN message telling all nodes to return their transceivers to normal mode. However, if the number of CAN FD messages broadcast transitioned the CAN nodes into “bus off”, additional actions will be necessary to return the CAN nodes to error active state allowing CAN communication.

Depending on the CAN application implementation for bus off, it may be possible via a special CAN message to reset the error counters. If the application does not allow for a reset, the CAN FD node must send enough valid CAN messages and interframe spaces to clear the error counters and return the CAN nodes to error active or bus on condition.

Once the nodes are bus on, the final special CAN message is sent. This message commands the CAN nodes to put the transceivers in normal mode.

Partial networking with FD passive

Another method of mixing CAN FD and CAN on the same network utilizes the other emerging CAN transceiver standard, CAN with selective wake, or partial networking functionality. This method is somewhat similar to the FD Blanking transceiver that will be described later, but uses a significantly more complex transceiver and software to go back to CAN communication after CAN FD communication.

The selective wake functionality described in ISO11898-6 is the base with an FD passive feature added. This feature sets the transceiver's frame decoder to ignore CAN FD frames. A special CAN message will be sent to put the CAN nodes in selective wake mode with FD passive. Upon completion of the CAN FD communication a CAN FD node sends a CAN message encoded as a wake up frame (WUF) for the CAN nodes. Upon reception of this wake up frame the CAN nodes will return to normal CAN communication.

CAN FD blanking transceiver

A more straight forward solution requires simple additions to the current physical layer. The concepts build easily on the current CAN transceivers based on ISO11898-2 and -5 standards. This new hardware in conjunction with application software will provide a mechanism to blank the FD frames from the CAN nodes and then provide a simple Return to CAN request once the bus is ready for CAN messages again. This method prevents the error counters from increasing in the CAN nodes and allows for much faster bus transition between CAN and CAN FD than the software method relying on silent mode of the CAN transceiver if the CAN nodes went bus off.

In order to implement this solution a method is needed to signal the CAN nodes when CAN FD communication is complete. This method cannot overlap or interfere with the current transceiver standards. For cost and ease of implementation known concepts are used much as possible.

An example of such an implementation is shown in Figure 2. This transceiver equivalent uses a mechanism to disable the standard CAN transceiver while enabling a second receiver which is looking for a new Return to CAN request while in the FD Blanking mode.

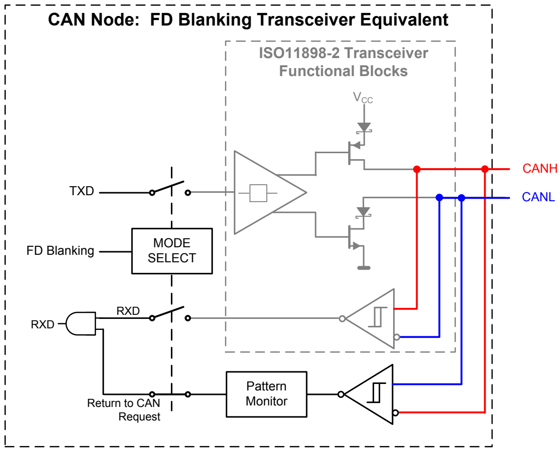

Figure 2: CAN node FD blanking transceiver

The system software will utilize this new hardware and signaling in the following way. The CAN FD node will send a special CAN message telling the CAN nodes to set their transceivers for CAN FD Blanking mode. This mode disables the normal transceiver by turning off the TXD to driver path and receiver to RXD path thus blanking all CAN FD communication.

To re-use known concepts as much as possible the application software is set to look for a signal on RXD. In ISO11898-5 the wake request is signaled on RXD corresponding to a valid wake up pattern on the bus. For a CAN FD Blanking transceiver the Return to CAN request is signaled on RXD corresponding to a valid Return to CAN signal on the bus.

The CAN FD Blanking mode uses a second receiver which is wired in reverse to CANH and CANL from a normal transceiver. This receiver along with a pattern monitor looks for the Return to CAN request which matches the wake pattern from ISO11898-5 except with an inverted differential signal.

Once the CAN FD communication is completed, the CAN FD node uses a new driver circuit that generates the Return to CAN request. Figure 3 shows a simple implementation for the CAN FD Node transceiver with the Return to CAN transceiver function.

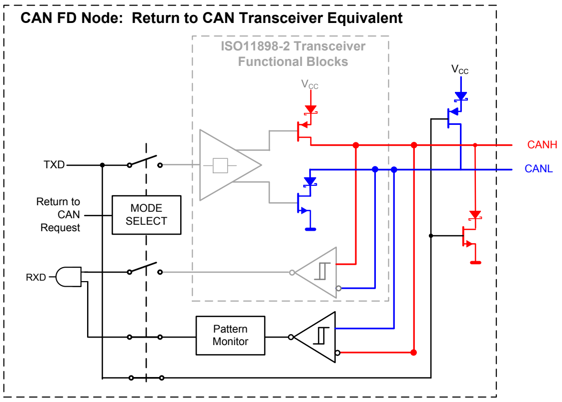

Figure 3: Return to CAN transceiver equivalent

The bus signaling with this new third state, an inverted dominant, for the FD Blanking with Return to CAN signal is shown in Figure 4.

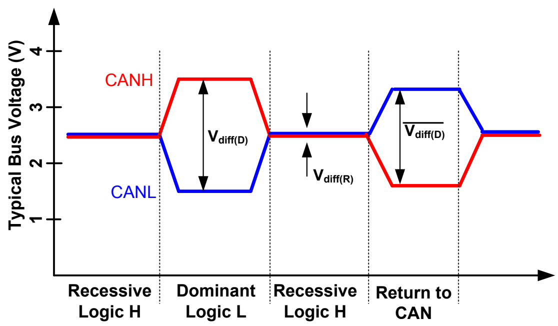

Figure 4. Bus signaling for FD blanking and return to CAN

An alternative and simpler implementation is shown in Figure 5. The bus signaling for this implementation is shown in Figure 6.

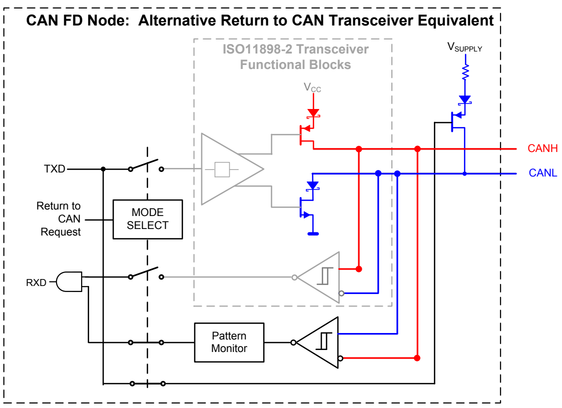

Figure 5: Alternative return to CAN transceiver equivalent

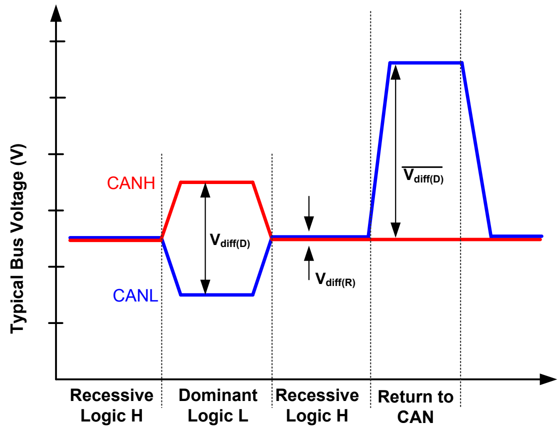

Figure 6: Alternative bus signaling for FD blanking and return to CAN

Once the CAN nodes have received the Return to CAN request they set their CAN FD Blanking transceivers back to normal mode and CAN communication resumes.
This method is scalable in implementation since the requirements on the Return to CAN portion of the transceiver only needs to receive the pattern and does not need to meet the full requirements of CAN or CAN FD communication. The Return to CAN request receiver and driver may be optimized to minimize cost and parasitic impact to the bus similar to the simplified receiver requirements defined for low power wake up in ISO11898-5.
The communication overhead cost for this method is two or three CAN messages. This is potentially far less than the software and silent mode method where enough CAN messages may have to be sent to reduce the error counter to under 128. This will return the CAN node to error active and bus on condition. This method also does not require specialized CAN wake up frame messaging and software control needed for the Partial Networking with FD Passive approach.
Bus waveform measurements using a discrete implementation based on existing CAN transceivers are shown in Figure 8, Figure 9, and Figure 10. These oscilloscope pictures were taken using the schematics shown in Figure 7. Simplified wave forms to represent CAN FD and Return to CAN request were used for clarity of the bus and transceiver behavior in the oscilloscope pictures.

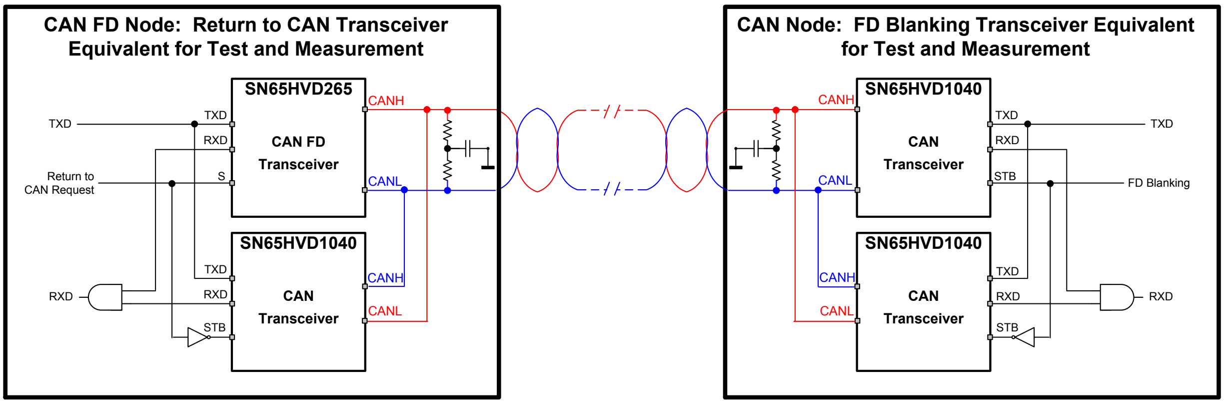

Figure 7: Test and measurement set up

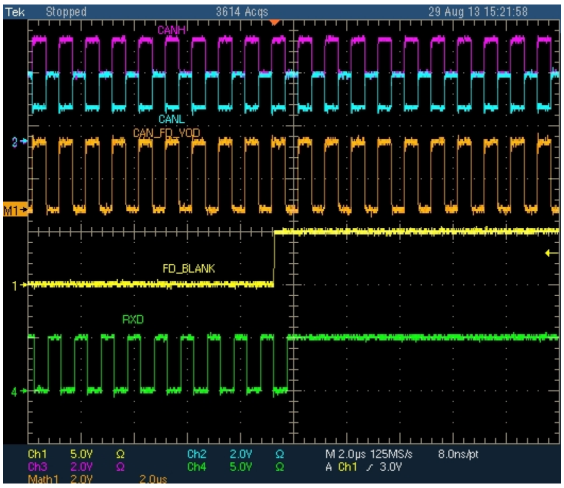

Figure 8: CAN node with transceiver blanking measurement

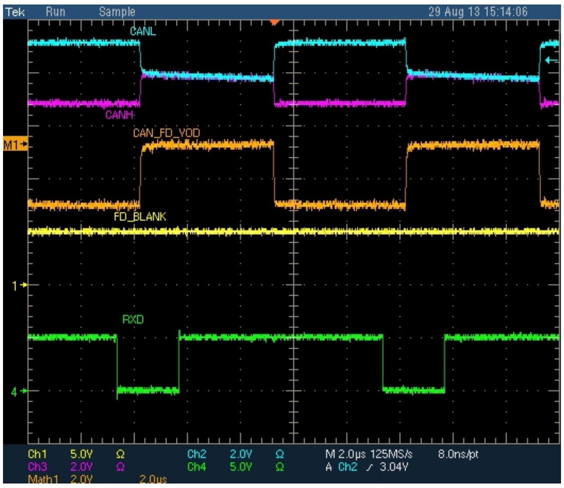

Figure 10: CAN node with blanking transceiver, receiving the return to CAN request

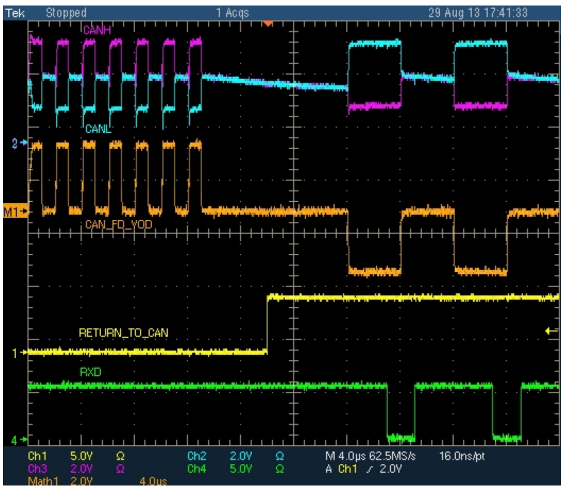

Figure 9: CAN FD node with return to CAN request measurement

Figure 8 illustrates the effect of using the blanking function, screening the 2Mbps CAN FD signal present on the bus from the CAN node.

Figure 9 illustrates the end of a CAN FD transmission followed by the Return to CAN request. The inverted differential signal generating the Return to CAN request is seen on the bus highlighted by the polarity shift shown in the differential voltage signal (orange). The RXD being monitored is the CAN node with the FD Blanking signal turned on. This illustrates the CAN FD signal is blanked from the node but the Return to CAN request gets through.

Figure 10 illustrates the Return to CAN request in more detail as seen at the CAN node.

Managed hub

The fourth method of mixing CAN FD nodes and CAN nodes into a network offers the highest performance but also at the most expense. This method involves the use of a managed hub, which may also be called a repeater, gateway, bridge, or even router. Essentially this managed hub breaks the network up into sub-networks of homogeneous nature and manages the communication flow between the sub-networks. An example of a managed hub system is shown in Figure 11.

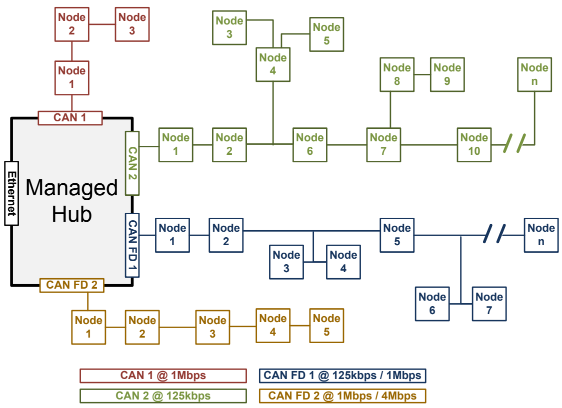

Figure 11: Managed hub

Adding this active component into the mix provides network segmentation ensuring the highest level of performance in the sub-networks where it is needed. Thus CAN FD nodes may be added on their own sub-network which may be wired with an ideal bus topology to maintain the highest network signal integrity and data rates. The CAN nodes may remain on their sub-network which may not have been architected for high data rates and highest signal integrity.

This approach is extremely scalable and flexible but due to the active nature of the managed hub it is also the most expensive. There may also be latency issues between different CAN and CAN FD sub-networks as the messages will have to be actively re-broadcast across sub-network domains.

Summary and conclusion

While the future looks bright for higher data rates and larger payloads in CAN systems via CAN FD, in the short term these solutions offer options to start gaining the CAN FD benefits sooner with mixed networks.

The cost versus benefit and system trade-offs of having CAN FD messages mixed with CAN messages must be evaluated to determine the best approach for the specific application.

Finally, these solutions provide various options that trade-off CAN messaging overhead, latency, software complexity, hardware, and cost, which all must be taken into account for effective system designs using CAN and CAN FD networks.

Scott Monroe
Texas Instruments, Inc.
12500 TI Boulevard, MS 8701
US-75243 Dallas, TX
Tel.: +1 214-567-5432
s-monroe1@ti.com
www.ti.com

David Stout
Texas Instruments, Inc.
12500 TI Boulevard, MS 8701
US-75243 Dallas, TX
Tel.: +1 214-480-6847
david.stout@ti.com
www.ti.com

John Griffith
Texas Instruments, Inc.
12500 TI Boulevard, MS 8701
US-75243 Dallas, TX
Tel.: + 214-567-0337
johngriffith@ti.com
Website: www.ti.com

References

[1] ISO11989-1, Controller area network (CAN) – Part 1: Data link layer and physical signaling
[2] ISO11989-2, Controller area network (CAN) – Part 2: High-speed medium access unit
[3] ISO11989-5, Controller area network (CAN) – Part 5: High-speed medium access unit with low-power mode
[4] ISO11989-6, Controller area network (CAN) – Part 6: High-speed medium access unit with selective wake-up functionality
[5] A Comprehensible Guide to Controller Area Network, Wilfried Voss, Copperhill Media Corporation
[6] CAN with Flexible Data-Rate, White Paper Version 1.1, Robert Bosch, GmbH.
[7] SN65HVD256, SN65HVD265, and SN65HVD1040 Datasheets, Texas Instruments, Inc.