

Article

Controller Area Network (CAN) Bus Transceiver with Authentication Support and Enhanced Rail Converters

Can Hong ¹, Weizhong Chen ¹, Xianshan Wen ¹, Theodore W. Manikas ² , Ping Gui ¹
and Mitchell A. Thornton ^{1,3,*} 

¹ Department of Electrical and Computer Engineering, Southern Methodist University, Dallas, TX 75205, USA; canh@smu.edu (C.H.); weizhongc@smu.edu (W.C.); xianshanw@smu.edu (X.W.); pgui@lyle.smu.edu (P.G.)

² Department of Computer Science, Southern Methodist University, Dallas, TX 75205, USA; manikas@lyle.smu.edu

³ Darwin Deason Institute for Cyber Security, Southern Methodist University, Dallas, TX 75205, USA

* Correspondence: mitch@smu.edu

Abstract: This paper presents an advanced Controller Area Network (CAN) bus transceiver designed to enhance security using frame-level authentication with the concept of a nonphysical virtual auxiliary data channel. We describe the newly conceived transceiver security features and provide results concerning the design, implementation, fabrication and test of the transceiver to validate its functionality and robust operation in the presence of systemic error sources including Process, Voltage, and Temperature (PVT) variations. The virtual auxiliary channel integrates CAN frame authentication signatures into the primary data payload via phase modulation while also providing compatibility with existing CAN protocols, interoperability with non-enhanced systems and requiring no network or software modifications. Enhanced rail converters are designed to facilitate single-rail to dual-rail data conversion and vice versa, preserving phase information and minimizing phase errors across various nonideal effects such as frequency drift, Process, Voltage, and Temperature (PVT) variations, and cable phase mismatch. This ensures reliable data transmission and robust authentication in the presence of adversarial cyberattacks such as packet injection. The receiver recovers both the CAN frame data and the security signature, comparing the latter with an authorized signature to provide a real-time “GO/NO_GO” signal for verifying packet authenticity and without exceeding the CAN clock jitter specifications.



Citation: Hong, C.; Chen, W.; Wen, X.; Manikas, T.W.; Gui, P.; Thornton, M.A. Controller Area Network (CAN) Bus Transceiver with Authentication Support and Enhanced Rail Converters. *Chips* **2024**, *3*, 361–378. <https://doi.org/10.3390/chips3040018>

Received: 15 September 2024

Revised: 20 October 2024

Accepted: 30 October 2024

Published: 4 November 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: automotive security; CAN bus transceiver; packet authentication; dual-rail communication; phase mismatch

1. Introduction

The Controller Area Network (CAN) protocol, developed by Bosch in the 1980s, has become a fundamental communication standard in automotive and industrial control systems. Its widespread adoption is attributed to its reliability, real-time capabilities, and cost-effectiveness, making it indispensable for in-vehicle networking and various industrial applications. However, the protocol’s original design did not anticipate the current landscape of cybersecurity threats, leaving CAN-based systems vulnerable to a range of attacks. Recent research has exposed significant security vulnerabilities in the CAN protocol, most notably demonstrated by the remote exploitation of a 2014 Jeep Cherokee, which underscored the potential risks associated with unsecured in-vehicle networks [1]. The lack of native support for message authentication and encryption in CAN makes it susceptible to various forms of cyberattacks, including message spoofing, replay attacks, and denial of service (DoS), all of which could compromise vehicle safety and data integrity.

To address these vulnerabilities, numerous solutions have been proposed, ranging from the use of cryptographic techniques such as Hash-based Message Authentication Codes (HMACs) and Advanced Encryption Standard (AES) encryption to other hardware

modifications and software-based approaches [2–6]. While these methods enhance security, they often face challenges in terms of compatibility, implementation complexity, and cost. For example, cryptographic methods generally require significant modifications to both hardware and software, impacting overall system performance and increasing the cost. Furthermore, adding these forms of cryptographic approaches results in a CAN system that is not compliant with or backward-compatible with non-equipped systems. As described in more detail in a later section of this paper, our approach is backward-compatible with unequipped systems, and this fact was validated in our laboratory testing of a set of fabricated Integrated Circuits (IC). Similarly, approaches like dynamic identification (ID) virtualization and Electronic Control Unit (ECU) fingerprinting, though innovative, introduce additional layers of complexity and incompatibility with the CAN specifications that may not be feasible for all applications [6]. There have also been recent applications of machine learning (ML) to CAN bus security. For example, Zhang et al. [7] uses a convolutional encoder network (CEN) for intrusion detection, while Shi et al. [8] combines a convolutional neural network (CNN) with entropy-based clustering. However, these approaches require the addition of hardware and software to the CAN bus to implement the ML methods. Furthermore, even when the ML methods achieve high accuracy, they still result in non-zero false negative and false positive classification errors, which are unacceptable in safety-critical systems such as automotive CAN-based communications. For example, the method in [7] reports non-zero False Negative Rates (FNRs). Although these rates are small, even a rarely occurring false negative with respect to, for example, a CAN frame commanding the automobile's brakes to engage could result in a fatal collision to occur. Given that a modern automobile processes a huge number of CAN frames in a typical single usage of the vehicle, even very small FNRs can result in the probability of a fatal accident to be unacceptably high. In general, the fact that even the best ML methods have non-zero FNR values is a primary reason that such approaches are not often implemented in safety-critical systems such as automobiles. In contrast, our approach of providing deterministically computed authentication signatures does not suffer from non-zero FNR values since we do not employ a ML approach.

This paper presents a novel CAN bus transceiver design that integrates phase-preserving dual-rail converters, offering a secure authentication mechanism without altering the existing CAN protocol or hardware. The proposed transceiver design is “backward compatible”, which means that it can work with other existing (non-secure) CAN transceivers in the CAN bus system. The only required change to an existing CAN system is the replacement of the transceiver circuits with our new modified security-enhanced transceiver. Since our transceiver is backward-compatible, it is not required to replace the transceivers in ECUs where safety is not deemed to be critical. By embedding an auxiliary data channel within the primary data stream using phase modulation, as shown in Figure 1, this approach ensures compatibility with existing CAN systems while providing an additional layer of security. This method draws inspiration from earlier research on phase modulation for data authentication in high-speed communication protocols [9], adapting it to the unique constraints of CAN networks. The added security is accomplished by incorporating an authentication signature that can be unique for each CAN frame. Furthermore, this signature is selectable and is in the form of either an eight-bit or a sixteen-bit signature.

The proposed solution aims to enhance the security of CAN-based systems by providing a backward-compatible cost-effective method for secure communication. Our proposed approach authenticates each CAN frame, which can mitigate a significant number of attacks such as message spoofing, replay, DoS, various man-in-the-middle (MitM) attacks including MitM insertion attacks, or frame flooding attacks. Although no security-enhancing improvements are guaranteed to cover all possible attacks, particularly since it is impossible to predict future zero-day attacks that have not yet been invented or observed, the list of attacks provided above indicates that our approach in providing frame-level authentication is effective in mitigating the effects of a large class of threats. The following sections of this paper detail the design, implementation, and evaluation of the proposed transceiver,

demonstrating its efficacy in mitigating a large class of security risks while preserving the functional integrity of existing CAN infrastructure.




Figure 1. Proposed CAN asynchronous serial link with a virtual auxiliary channel.

2. Overall Architecture

A block diagram of the CAN transceiver with enhanced rail converters is shown in Figure 2. During phase modulation within the transmitter (TX), the phase of the primary data is modulated with the authentication data that is transmitted via the virtual auxiliary channel. The TX rail converter converts the modulated single-rail signal (TX_{IN}) to dual-rail signals before asynchronously driving the modulated CAN frame onto the CANH/CANL cable's dual-rail transmission lines. The receiver (RX) recovers both the primary data and authentication signature by extracting the phase information embedded in the primary data. The dual-rail signals present on the CANH/CANL cable are converted back into a single-rail CMOS signal (RX_{OUT}) in the front end of the RX before the CAN frame authentication data is extracted using a phase demodulator.




Figure 2. CAN transceiver block diagram.

2.1. Phase Modulation

The serial CAN bus transmission speed in our design can achieve up to a 1 Mb/s data rate and supports both internal synchronization and authentication functions that are unique to this proposed solution and not inherent to standard CAN transceivers. A block diagram of the TX phase modulation circuitry and an associated timing diagram is shown in Figure 3. Each data bit can incorporate up to 25 CAN standard time quanta (TQ), providing for a fine-grained time resolution accuracy of 40 ns within the RX for internal clocking. We limit the allowable portions of the CAN frame in which to incorporate a modulation signature to the data payload portions that are guaranteed to be received after

arbitration has occurred. Specifically, we do not modulate the arbitration fields nor the dominant acknowledgement bit in the CAN frame.




Figure 3. Phase modulation and timing diagram.

The concept of dividing CAN frame timing into small units referred to as “time quanta” is explained in detail in the CAN specification [10]. Specifically, a “time quantum” is a fixed unit of time derived from the oscillator period defined by a programmable prescaler that is typically implemented within the ECU. With respect to the TX phase modulator, a constant time value of three TQs is implemented to selectively delay the serial primary data stream when the authentication signature bit is “1”, or no time shift is induced into the frame when the authentication signature bit is “0”. The system clock is derived from a local clock with a period of $1/TQ$ or 25 MHz that is divided by 25 to generate the 1 MHz clock CLK0 that sets the CAN frame serial data rate that internally synchronizes the 1 Mb/s primary CAN frame data. CLK0_D is generated by delaying CLK0 by 3 TQs, or 120 ns.

The 40 ns TQ design choice is based on the worst-case scenario of a 25 MHz CAN bus clock, which is the maximum data rate of 1 Mbps as per the specifications [10]. For reduced data rates, the 3 TQ modulation index would remain, but the value of 40 ns would accordingly increase to accommodate slower data rates. The 3 TQ modulation constant falls within the CAN specification in view of allowable jitter and clock drift values. Also, the use of 25-stage Time-to-Digital Converter circuits (TDCs) in our receiver circuit would remain the same and would not be affected or need to be changed for slower clock rates as validated by our testing and simulation studies.

Depending on whether the auxiliary signature bit is ‘1’ or ‘0’, modulation is performed when the D flip-flop selects either CLK0 or CLK0_D to re-sample the primary data, thus producing the modulated CAN frame as illustrated in the timing diagram of Figure 3. As specified in the CAN communication protocol, the number of consecutive ‘1’s or ‘0’s in the CAN frame payload cannot exceed five; thus, the frequency of the auxiliary or authentication signature is set to be one-fifth of the CAN frame data rate to ensure at least one data transition is present for modulation for every five primary data bits. This modulation schema enables the authentication signature to be recovered by the receiver (RX) demodulator by detecting the edge transition times, or phase, in the single-rail signal

that represents the primary CAN frame data with Non-return to Zero (NRZ) signaling as specified in the CAN physical layer protocol standard [10].

This modulation scheme can be considered as rising edge or “phase modulation” of an NRZ signal, which is a variation of Pulse Position Modulation (PPM) and Pulse Width Modulation (PWM) since modulating the timing position of the rising edge affects both the position of a pulse and the width of a pulse in the single-rail signal representing CAN frame data. It is not strictly a PWM method, since the pulse widths do vary in an NRZ signal, nor is it strictly a PPM method, since only the rising edge of the NRZ pulse is modified and not the falling edge; however, it does have similarities to both approaches. Specifically, the modulation equation for the timing position of the rising pulse edge, t_{rise} , in reference to the rising edge of the local 25 MHz clock, CLK0, t_{R_CLK0} , indicates the value of a bit as modulated into the signature.

$$t_{R_CLK0} - t_{rise} = \begin{cases} 0, & \text{Signature bit is "0"} \\ 120 \text{ ns}, & \text{Signature bit is "1"} \end{cases}$$

The delay value of 3 TQ = 120 ns combined with the resynchronization that realigns subsequent NRZ pulse edges means that the smallest possible reduction in a single pulse in an NRZ stream is 80 ns, since a 1 TQ reduction due to jitter is allowed in the CAN specifications, (i.e., 3 TQ – 1 TQ = 120 ns – 40 ns = 80 ns). Therefore, the increase in required bandwidth of the modulated versus unmodulated NRZ signal is very small and widens the frequency-domain *sinc* function corresponding to the smallest possible pulse in the NRZ stream by an amount due to the difference of the unmodulated smallest pulse size of 24 TQ in comparison to the smallest possible modulated pulse width of 22 TQ. This occurs since the smallest possible NRZ pulse at the highest data rate has a width of 96 ns (assuming an unfavorable maximum jitter value of 1 TQ or 40 ns) and the corresponding smallest possible pulse in our modulated scheme has a value of 88 ns, assuming that an unfavorable jitter value of 40 ns is present and that a signature bit with a “1” value is present). Therefore, the overall worst-case signal bandwidth increase changes from 10.417 MHz to 11.364 MHz or an overall increase in bandwidth requirements of 947 kHz. This increase in signal bandwidth is easily supported by the components used in typical fabrication processes, including more current fabrication processes with smaller feature sizes than that used in this work.

The RX includes a both primary data recovery path for extraction of the CAN frame data and an auxiliary data recovery path for recovering each frame’s authentication signature, as illustrated in Figure 4. To generate a clock signal that is synchronized with the modulated primary data, 25 different clock signals at 1 MHz with 40 ns time offset spacings are produced by a clock divider and a delay line. Synchronization is accomplished by the RX by selecting 1 of the 25 time-shifted clock signals that most closely aligns with edge transitions in the received data stream. The clock selection block (CLK_SEL1) chooses 1 of the 25 clock signals to designate as CLK1 that samples the modulated primary data. CLK_SEL1 operates as a hard synchronization mechanism, ensuring precise selection of the clock signal that most closely aligns with the edge transition in the received data stream. A time-to-digital converter (TDC1) with one TQ or 40 ns resolution detects the time difference between CLK1 and the modulated primary data edge transitions, adjusting CLK_SEL1 through Decoder1 to maintain CLK1 at an optimal sampling point, within a 1 TQ time interval. The RX also contains an auxiliary data recovery path that recovers the modulated authentication signature that implements a second TDC (referred to as TDC2) that detects the time difference between CLK2 rising edges and edge transitions within the modulated primary data. Based on TDC2’s output, edge transition timing, or phase information, of the modulated primary data is detected, allowing for recovery of the authentication signature or auxiliary data.




Figure 4. Primary data and auxiliary data recovery path.

2.2. Phase Extraction

The RX phase extraction circuit block diagram is shown in Figure 4 within the Auxiliary Data Recovery Path, and an associated timing diagram for the circuitry depicted in the Figure 4 block diagram is shown in Figure 5. When the RX detects the CAN start of frame (SOF), indicated by the first '1' to '0' edge transition in the received CAN frame, TDC1 begins updating CLK_SEL1 to ensure the falling edges of CLK1 are aligned with the edges of the modulated primary data whenever a misalignment is detected, otherwise known as 'soft synchronization'. The TDC1 updates CLK_SEL1 to keep the falling edges of CLK1 aligned with the edges of the modulated primary data when it detects a misalignment between these signals. The CLK_SEL1 operates as a hard synchronization mechanism, ensuring precise selection of the clock signal that most closely aligns with the edge transitions in the received data stream. We have designed CLK_SEL1 to be fast enough to follow and allow for edge timing variances due to modulation, jitter, and drift as has been validated in our laboratory tests. This alignment ensures the rising edges of CLK1 are optimal, occurring with a resolution of 1 TQ time unit for sampling the primary data. In the authentication signature, or auxiliary data, recovery path circuitry, TDC2 uses CLK2 as the reference clock to extract phase modulation from the primary data. CLK2 is generated by the CLK_SEL2 circuit that is used to align with the SOF, thus implementing a 'hard synchronization' function. CLK_SEL1 serves as hard synchronization at the beginning of the frame, ensuring that the receiver's clock aligns precisely with the transmitter's clock. This guarantees accurate tracking of phase modulation and proper sampling of incoming data. CLK_SEL2, in contrast, is used for resynchronization and received frame signature recovery during the frame transmission. It helps to maintain synchronization by adjusting for any potential clock drift that may occur over time during the inter-frame processing. By incorporating the Auxiliary Data Recovery Logic into the resynchronization portion of the receiver, the received frame signature is demodulated simultaneously with updates to the CLK2 signal that compensate for any potential clock drift. The phase-modulated auxiliary data, serving as the authentication signature, is extracted and recovered by TDC2 and the Auxiliary Data Recovery Logic that implements the phase demodulator within the RX. The auxiliary data recovery logic outputs a '0' as the recovered authentication data when the extracted phase of the primary data bit is less than 2 TQ, and outputs a '1' if the extracted phase is 2 TQ or greater. After the entire CAN frame has been received, a comparator (COMP) compares either the 8 or 16 bits of authentication signature to indicate whether the authentication

is verified since the transceiver provides for the authentication signature to be present as either a one- or two-byte word as specified at system reset.




Figure 5. Phase extraction timing diagram.

To successfully retrieve the auxiliary data, the RX must be capable of recovering the phase-modulated signature despite the systemic errors that may be present due to jitter and frequency drift between TX and RX clock signals, phase mismatch between CANH and CANL cable lines, or other errors due to changing environmental characteristics. Since the proposed receiver is fully digital with a 1 TQ or 40 ns time resolution, the jitter-induced error in the TDC output will not exceed the allowable 1 TQ value as per the specifications [10], provided that overall peak-to-peak jitter of the entire transmission chain remains less than 40 ns at the highest allowable data rate of 1 Mbps with a 25 MHz clock, regardless of the initial phase of the modulated primary data, as illustrated in Figure 6. The phase modulation constant delay value is set at 3 TQ, ensuring that even with a 1 TQ TDC detection error, the Auxiliary Data Recovery Logic can accurately extract the phase-modulated authentication signature word. Setting the phase modulation constant to a value greater than 3 TQ could cause the modulated edge delays to exceed allowable CAN system specifications when also considering the possibilities of clock drift, jitter, and variations due to the environmental extremes as provided in the CAN physical layer specifications.

2.3. Circuit Design Considerations

The effect of frequency drift varies depending on the CAN frame data length since phase errors induced by frequency discrepancies accumulate over time during transmission of the serial asynchronous signal. Thus, longer CAN frames become proportionately more susceptible to systemic errors. In the CAN bus transceiver described here, when the authentication signature word size is set to 16 bits, the minimum length CAN frame payload must correspond to a length of at least 80 bits (2000 TQ) to ensure that enough edge transitions are present within the frame to allow for the required number of transitions to occur for modulation purposes. This 80-bit value is determined in consideration that an edge transition must be present in the single-rail NRZ CAN frame signal for a maximum of every 5 consecutive bits of the same value as specified in the CAN protocol standard. To ensure accurate 16-bit auxiliary data recovery, the maximum phase error present within the final modulated primary data bitstream must remain within 1 TQ. Thus, the transceiver can handle a frequency error of up to 0.05% (1 TQ/2000 TQ) for any combination of primary

CAN frame and associated authentication data. However, a smaller 8-bit signature can optionally be used and was also implemented and tested in our transceiver IC.




Figure 6. Input time difference versus TDC detected time difference.

Phase mismatch can occur due to differences in the lengths of CANH and CANL lines within the cables that serve as the communications medium for dual-rail physical CAN signals that interconnect ECU nodes within a system, resulting in a load, or parasitic RC, mismatch. This mismatch can lead to pulse width errors that translate into phase errors at the RX output. Since accurate phase information is crucial for the phase extraction circuit to correctly recover a CAN frame authentication signature, enhanced single-rail to dual-rail, and dual-rail to single-rail, converter circuits are designed and implemented that minimize the effects of phase mismatch between the CANH and CANL transmission lines.

2.4. Rail Converter

CAN rail converter circuits are included within the TX and RX blocks of the CAN frame transceiver. Single-rail CAN data (TX_{IN}) is converted into dual-rail signals V_{CANH} and V_{CANL} by the TX. Likewise, the RX detects the voltage difference (V_{DIF}) between V_{CANH} and V_{CANL} and generates a corresponding NRZ single-rail output signal, RX_{OUT} , as shown in Figure 7.




Figure 7. Transmission of CAN bus transceiver.

A schematic of a representative and conventional non-phase-preserving CAN bus transceiver TX is shown in Figure 8. To generate a dominant bit as defined in the CAN specification, the driver activates transistors M_1 and M_2 by quickly pulling V_{CANH} to V_{DD} and V_{CANL} to ground (GND). In contrast, a recessive bit, as defined by the CAN standard, is pulled to a common-mode voltage (V_{CM}) by the termination resistors denoted as R_T . In this configuration, the speed of the transition from a dominant to a recessive bit is dictated by the value of the pull-up/pull-down resistor R_T . For faster transitions, R_T needs to be minimized resulting in a relatively large constant current flow from V_{CANH} and V_{CANL} to V_{CM} during the dominant bit transition, thereby limiting the circuit's power efficiency.




Figure 8. Conventional CAN dual-rail transmitter.

To address this limitation, the proposed TX rail converter includes transistors M_1 and M_6 that drive the dominant bit and transistors M_2 and M_5 that drive the recessive bit, as illustrated in Figure 9. Transistors M_3 , M_4 , M_7 , and M_8 act as switches controlled by the internal enable signal (EN) that disconnect the TX from the CANH and CANL cable lines when no data is being transmitted. The enable signal, EN, allows the driver amplifiers in the enhanced CAN transmitter of Figure 9 to send data through the switching transistors M_1 to M_8 into the dual-rail CAN cable lines. Since phase is critical in our design, the timing of the rising and falling edges of each bit is important. In contrast, for past CAN transceiver circuits that are not equipped with our security mechanism, the dominant-to-recessive bit transition is driven by the global resistors on the CANH and CANL lines, not by the CAN transmitters resulting in a CAN cable output driver circuit as shown in Figure 8. However, in this design, to ensure that the rising and falling edge times are equal, the dominant-to-recessive bit transition is also controlled by the transmitter. To prevent a short-to-ground situation, the EN signal is used to disable the transmitter when it is not transmitting data. Also, since both the dominant-to-recessive bit and recessive-to-dominant bit are driven by the transmitter, the terminal resistor is no longer the main driving source, which means that the single-resistor termination can be used. The enhanced transceiver operates with the EN signal normally unasserted and it is only asserted internally when the transceiver is actively driving CAN data onto the bus. This design allows for rapid transitions between recessive and dominant bits without consuming significant static power, as verified through measurements of the fabricated enhanced CAN transceiver circuit in our laboratory testing [11].

The RX implements a hysteretic comparison of voltage difference between V_{CANH} and V_{CANL} with positive trigger points defined by V_{THH} and V_{THL} , as shown in Figure 10. This approach guarantees that the dominant and recessive bits are not activated by V_{CANH} or V_{CANL} alone since the signal should be in a dual-rail form; thus, a sufficient voltage margin to avoid improper triggering is ensured.




Figure 9. The CAN dual-rail transmitter with enhanced driving capability.




Figure 10. (a) Hysteretic comparison with V_{THH} , V_{THL} ; (b) realization with V_{CANH} and V_{CANL} .

The rail converter circuit within the RX is shown at the transistor level in Figure 11. The comparator comprises two input pairs in a differential configuration. These are the NMOS pair M1 and M2 for V_{CANH} (0.9–1.8 V) and the PMOS pair M3 and M4 V_{CANL} (0–0.9 V). Two Direct Current (DC) common-mode voltages V_{REFH} (1.35 V) and V_{REFL} (0.45 V) are provided for voltage comparison. The two input differential pairs compare V_{CANH} and V_{CANL} with V_{REFH} and V_{REFL} , respectively, and then combine the resulting differential currents for hysteretic triggering purposes. The two trigger points are defined as $V_{THH} = V_{REFH} - V_{REFL} + V_{OS}$ and $V_{THL} = V_{REFH} - V_{REFL} - V_{OS}$, where $(V_{REFH} - V_{REFL})$ is equal to 0.9 V, and V_{OS} , determined by the sizes of M₆, M₇, M₈, and M₉, has a nominal value of 0.3 V [4]. Note that a mismatch in input pairs can cause phase errors. However, since this design has tolerance for phase extraction, a phase error of less than 1 TQ is acceptable as per the CAN specifications and as we have verified in our laboratory testing of the fabricated enhanced-security CAN transceiver circuit described here.

As the authentication signature data is extracted from RX_{OUT} , it is essential to ensure that the phase of RX_{OUT} , which can vary due to pulse width variation, matches with that of TX_{IN} even in the presence of non-ideal circuit conditions such as Process, Voltage, and Temperature (PVT) variations. PVT variations can influence V_{OS} , causing the triggering points V_{THH} and V_{THL} to shift, potentially inducing timing errors with respect to the rising edge (t_{th}) and falling edge (t_{fl}) of RX_{OUT} . However, since V_{THH} and V_{THL} are symmetrically positioned around 0.9 V, the voltage variation on V_{THH} (ΔV_{THH}) is the opposite of that on V_{THL} (ΔV_{THL}). Therefore, with a uniform slope on both the rising and falling edges, due to the balanced driving capability of the TX, t_{th} and t_{fl} remain identical and ensure that no errors are introduced in the pulse width of RX_{OUT} as shown in Figure 12a. Allowable PVT extremes are provided in automotive specifications and we used a special-

ized PVT chamber to validate our fabricated CAN transceiver chip during our laboratory testing exercises.




Figure 11. Schematic of hysteretic dual-to-single receiver.




Figure 12. (a) RX_{OUT} 's pulse width relationship with V_{THH} and V_{THL} ; (b) Phase mismatch compensation under different situations.

Another nonideal factor that can contribute to phase error on RX_{OUT} is phase mismatch occurring due to unequal latency on the dual-rail signaling lines of the CAN cable denoted as CANH and CANL. When there is an unequal transmission latency among the CANH or CANL lines, timing errors t_{dr} and t_{df} may arise on RX_{OUT} 's rising/falling edge. These timing errors are due to the signal arriving at the RX input with a greater delay, causing later triggering to occur. Because the triggering depends on the combined current from both differential input pairs, the bit transition of RX_{OUT} is primarily determined by the later-arriving signal among the two CANH/CANL lines, resulting in t_{dr} and t_{df} to be equal in a first-order manner. If the ideal (i.e., without phase mismatch) pulse width of RX_{OUT} is T , then the pulse width with phase mismatch becomes $T' = T - t_{dr} + t_{df} \approx T$. As a result, the impact of phase mismatch on RX_{OUT} 's pulse width is minimized as shown in Figure 12b.

2.5. Signature Generation and GO/NO_GO Signal

Our transceiver is designed such that both the transmitter and the receiver portions of the transceiver are equipped with signature generators. In this way, the receiver can compute the expected signature concurrently while demodulating the signature of a CAN frame that is actually received. The expected signature and the received, demodulated signature are then applied to a bitwise comparator circuit (COMP) to produce a single-rail signal, GO/NO_GO, that indicates if the signatures match or not. If the signatures do not match, the received frame can be dropped and an error frame can optionally be issued. The signature generation circuits are assumed to be provided by individual Hardware Security Modules (HSM). There are many different HSM cores available that can perform this function, and our intent is to design the transceiver in a manner that is agnostic to the choice of which particular HSM is employed for signature generation.

Our proposed transceiver design allows signature generation within the transceiver circuit to be accomplished in a manner such that signature synchronization is maintained among all nodes in the CAN subsystem. As a simple example, a possible signature generation approach is to implement a Linear Feedback Shift Register (LFSR), perhaps augmented with an additional hash function or transceiver-dependent Physically Unclonable Function (PUF) fingerprint. Signature synchronization is implemented through issuing broadcast messages to all ECUs and their associated transceivers, either during CAN bus initialization or at any later time. The transceivers then periodically increment the LFSR to generate new signatures as commanded by the Microcontroller Command Unit (MCU) within each CAN system node. With this approach, each transceiver on the bus can use the generated signature from their HSMs for each frame to modify and compute the signatures and GO/NO_GO signals. This approach enables the MCUs in each CAN system node to maintain synchronization among signatures in the CAN system. Furthermore, this approach allows CAN system implementers to choose the frequency with which updated signature values are generated within the system. At the finest-grain level, each CAN frame can have its own unique signature, or systems can employ a reduced granularity where groups of subsequent messages share a common signature that is updated less frequently depending on how often their HSMs are commanded to generate a new signature root value. For our testing, we focused on generating new signatures within this different range of granularities and validated that synchronization was maintained.

Our specific transceiver chip, as described here, is implemented with an exemplary HSM for signature generation that comprises an LFSR with its output further scrambled by a small hash function designed to preserve security properties including adherence to the avalanche effect, and was tested using the approach described in Section 3 to ensure synchronization is maintained. This simple and exemplary HSM is specified as a Register Transfer Level (RTL) Verilog module that was then synthesized into a small block of digital circuitry and incorporated into the overall transceiver circuit layout. It would be expected that users of the transceiver would use a more secure HSM, but we were focused on incorporating authentication in the transceiver rather than investigating the well-studied area of HSM architectures.

As the choice of HSM is likely to vary from manufacturer to manufacturer, our design goal is to produce an enhanced-security CAN transceiver circuit that is agnostic to the choice of a particular HSM since the use of a specific HSM will be accomplished by automotive manufacturers. Typically, the specific details of the HSM are closely guarded proprietary information to prevent adversaries from attempting to exploit the signature generation circuitry. Therefore, the security-enhanced CAN transceiver offers up to frame-level authentication in a manner such that users of the IC have the freedom to choose the actual signature generator to be deployed; whether it is implemented as a subcircuit within the transceiver or as a module, as is the case for our transceiver chip described here, or within the MCU/ECU as software or firmware, although this latter case is optional and not strictly required to use our security-enhanced transceivers. This approach allows a large degree of freedom in choosing the actual HSM and its means of implementation; however,

should users of the enhanced CAN transceiver decide to use the simple signature generator we implemented, it is also available for use as an integrated digital subcircuit within the transceiver itself.

3. Measurement Results

The CAN transceiver chip was fabricated using a 180 nm process and a chip die photo is presented in Figure 13. The transceiver communication test setup is shown in Figure 14, where two noncoherent signal generators are used to provide a 25 MHz clock to each of Chip#1 and Chip#2 individually. Figure 14a shows a block diagram of the testing setup, while Figure 14b shows a photograph of the testing setup with the addition of a temperature chamber. The use of these two noncoherent signal generators mimics the practical situation among the transceivers present within two automotive ECUs that comprise a CAN implementation. A Field-Programmable Gate Array (FPGA) development kit is programmed and used to emulate an MCU/ECU CAN node that supplies the primary data (i.e., CAN frame) and auxiliary data (i.e., authentication signature) to one of the transceiver chips. A logic analyzer is employed to capture and observe the CAN transceiver outputs. Chip#1 functions as the transceiver TX in a first ECU that sends CAN frames to Chip#2 through CANH/CANL cables where Chip#2 functions as the transceiver in a second ECU, and vice versa.




Figure 13. Chip die photo for CAN transceiver.




Figure 14. (a) CAN bus chip testing setup. (b) CAN bus chip testing setup with temperature chamber.

3.1. Functionality and Measurement against Frequency Drift between TX and RX

Figure 15 shows the measured waveforms captured by the logic analyzer that verifies the functionality of the CAN transceiver chips. Figure 15a shows the results for a +0.05% frequency drift, while Figure 15b shows the results for a −0.05% frequency drift. The measured RX_{OUT} is demonstrated to show that the demodulated authentication signal recovered by the Chip#2 RX preserves the same phase information as was introduced by modulated primary data in the Chip#1 TX_{IN} and to further show that the 16-bit extracted auxiliary data (Auth_Rec) matches with the auxiliary data on TX side (Auth_TX), despite a $\pm 0.05\%$ frequency drift between TX and RX, where the frequency drift is defined as $(f_{TX}/f_{RX} - 1)$. With −0.05% frequency drift, the Auth_Rec signal begins to diverge from Auth_TX after the extraction of the 16-bit authentication signature since the accumulated phase error due to the clock frequency mismatch between Chip#1 and Chip#2 starts to exceed 1 TQ.




Figure 15. Measurement results with frequency drift between TX and RX: (a) +0.05% frequency drift, (b) -0.05% frequency drift.

3.2. Transceiver Evaluation with PVT Variations

The transceiver's functionality is further verified by making measurements with V_{DD} varying by $\pm 10\%$, ranging from 1.62 V to 1.98 V, and temperatures ranging from $-28.5\text{ }^{\circ}\text{C}$ to $120\text{ }^{\circ}\text{C}$. Figure 16 shows the pulse width error of RX_{OUT} compared to 1 TQ (40 ns). Despite variations in voltage and temperature, the pulse width error remains below 25% of 1 TQ.




Figure 16. RX_{OUT} pulse width error with temperature ($-28.5\text{ }^{\circ}\text{C}$ to $120\text{ }^{\circ}\text{C}$) and voltage ($V_{DD} = 1.8\text{ V}/1.98\text{ V}/1.62\text{ V}$) variation.

3.3. Measurement with Cable Lengths Mismatch to Induce Phase Shift Errors

Figure 17 shows the zoomed in waveforms of the rising and falling edge of the transceiver's RX output and V_{CANH}/V_{CANL} in the presence of phase mismatch due to unequal CANH/CANL line lengths. The phase mismatch between V_{CANH} and V_{CANL} is set to 50 ns. The measured t_{dr} is 75 ns, t_{df} is 68 ns, and the phase error in RX_{OUT} is suppressed to 8 ns.




Figure 17. Timing errors with phase mismatch between V_{CANH} and V_{CANL} (a) t_{dr} , (b) t_{df} .

The measured pulse width of RX_{OUT} in the presence of induced phase mismatch between the CANH and CANL lines with unequal latencies is shown in Figure 18. Rather than using unequal line lengths, the phase mismatch is introduced by adding additional Resistor/Capacitor (RC) components to add a delay to either V_{CANH} or V_{CANL} . The RC time delay circuit is a more convenient method for purposely inducing line mismatch delay than using an additional length of cabling, since a nontrivial length of additional cabling would be required at these relatively low data throughput speeds (i.e., 1 Mb/s) and the induced delay values are easier to vary and more accurately controlled. The ideal pulse width of RX_{OUT} is $1 \mu s + 3 TQ$ (1.12 μs) due to the chosen phase modulation parameters previously discussed. With a 55 ns/100 ns phase mismatch purposely induced on the CANH/CANL cable lines, the resulting error in RX_{OUT} 's pulse width is only 10 ns/5 ns indicating that the effect of phase mismatch on RX_{OUT} 's pulse width is significantly suppressed. Figure 19 shows the measurement results of the three-pulse width of RX_{OUT} as 3 μs without phase modulation and with phase modulation RX_{OUT} 3 $\mu s + 3 TQ$ (3.12 μs). The phase mismatch of 110 ns between CANH and CANL has no impact on RX_{OUT} 's pulse width.




Figure 18. Measured pulse width of RX_{OUT} with phase mismatch between CANH and CANL. (a) Extra RC on V_{CANH} , (b) extra RC on V_{CANL} .

A total of ten CAN transceiver Integrated Circuits (ICs) were tested, all demonstrating consistent results, thereby validating the proposed CAN bus transceiver architecture and circuits. Additionally, the rail converter has been verified under conditions of delay mismatch as well as voltage and temperature variations.




Figure 19. Measured pulse width of RX_OUT (a) without phase modulation, (b) with phase modulation.

3.4. Transceiver Test with Signature Mechanism Disabled

We designed our transceiver to function as a normal off-the-shelf (unequipped) transceiver by disabling selectively the signature mechanism. This allowed us to test a multiple transceiver CAN bus with some transceivers functioning like off-the-shelf devices and some to function with our security modifications. Our tests revealed that system functionality was maintained although, as expected, the signature comparison failed in the transceivers that had no added security functionality; however, the CAN still functioned normally as would occur in a non-equipped system.

4. Conclusions

A secure CAN bus transceiver with enhanced security via the incorporation of individual and independent CAN frame authentications is devised and developed from initial concept through fabrication using a 0.18 μm CMOS process and capable of operating with a 1 Mb/s bandwidth. The fabricated CAN transceiver chips are physically validated to demonstrate functionality and robust operation in the presence of emulated systemic and environmental extremes consistent with those encountered in automobiles. Unlike previous approaches, the transceiver described herein maintains compatibility with current and past CAN physical protocol specifications, is interoperable with CAN systems comprising conventional non-enhanced ECU transceivers, and does not require extensive hardware or software modifications to incorporate into existing systems.

Enhanced CAN system security is achieved by phase-modulating the CAN frame payloads to embed frame-specific authentication signatures allowing for additional security data to be concurrently transmitted over a nonphysical virtual channel that avoids the need for inclusion of an additional data communications channel. The modulation subsystem parameters are carefully chosen to adhere to the phase jitter tolerances of the CAN standard, ensuring that the approach remains compatible with existing systems and presenting as noise or jitter that falls below allowable maximum thresholds. The TX modulation circuitry is very simple, comprising a D flip-flop to mix the authentication signature with the CAN frame payload data and a secondary TDC circuit in the RX to demodulate and recover the received authentication signature. Real-time CAN frame authentication is accomplished through the inclusion of a bit-parallel comparator circuit that detects the presence of any mismatches between the recovered and the expected CAN frame signatures producing a GO/NO_GO signature at each ECU that can be used to discard frames that are in error or that have been maliciously injected into the CAN system. The transceiver also includes enhanced rail converter circuits that preserve edge transition timings during conversion from single- to dual-rail signaling and vice versa.

The enhanced rail converter circuitry in combination with the other security-enhancing circuitry enable the CAN transceiver to benefit from improved rising and falling edge performance, and the receiver effectively mitigates phase errors caused by PVT environmental

variation, phase mismatches due to unequal latencies in the CAN cable CANH and CANL lines, and frequency variations among the noncoherent local ECU clock generator circuits due to jitter and drift. Comprehensive measurements under different PVT conditions and other systemic error sources validate the robustness of the CAN bus transceiver architecture, confirming its ability to support robust, enhanced, and secure communications within a CAN bus system even in the presence of malicious adversarial cyber activity such as a control frame injection attack.

Author Contributions: Conceptualization, resources, project administration and funding acquisition, M.A.T. and P.G.; supervision, P.G., M.A.T. and T.W.M.; data curation and visualization, C.H.; methodology, software, validation, formal analysis, investigation, writing—original draft preparation and writing—reviewing and editing, C.H., W.C., X.W., T.W.M., P.G. and M.A.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Toyota Motor North America, Inc. (TMNA).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in this article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Miller, C.; Valasek, C. Remote Exploitation of an Unaltered Passenger Vehicle. In Proceedings of the Black Hat USA, Las Vegas, NV, USA, 1–4 August 2015; pp. 1–91.
2. Maruaisap, A.; Kumhom, P. A hardware-based security scheme for in-vehicle CAN. In Proceedings of the International Computer Science and Engineering Conference (ICSEC), Chiang Mai, Thailand, 14–17 December 2016; pp. 1–5.
3. Wu, Y.; Kim, Y.-J.; Piao, Z.; Chung, J.; Kim, Y.-E. Security protocol for controller area network using ECANDC compression algorithm. In Proceedings of the IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC), Hong Kong, China, 5–8 August 2016; pp. 1–4.
4. Mun, H.; Han, K.; Lee, D.H. Ensuring Safety and Security in CAN Based Automotive Embedded Systems: A Combination of Design Optimization and Secure Communication. *IEEE Trans. Veh. Technol.* **2020**, *69*, 7078–7091. [[CrossRef](#)]
5. Avatefipour, O.; Hafeez, A.; Tayyab, M.; Malik, H. Linking received packet to the transmitter through physical-finger printing of controller area network. In Proceedings of the IEEE Workshop on Information Forensics and Security (WIFS), Rennes, France, 4–7 December 2017; pp. 1–6.
6. Song, H.M.; Kim, H.R.; Kim, H.K. Intrusion detection system based on the analysis of time intervals of CAN messages for in-vehicle network. In Proceedings of the International Conference on Information Networking (ICOIN), Kota Kinabalu, Malaysia, 13–15 January 2016; pp. 63–68.
7. Zhang, X.; Cui, X.; Cheng, K.; Zhang, L. A Convolutional Encoder Network for Intrusion Detection in Controller Area Networks. In Proceedings of the 2020 16th International Conference on Computational Intelligence and Security (CIS), Guangxi, China, 27–30 November 2020; pp. 366–369.
8. Shi, J.; Xie, Z.; Dong, L.; Jiang, X.; Jin, X. IDS-DEC: A novel intrusion detection for CAN bus traffic based on deep embedded clustering. *Veh. Commun.* **2024**, *49*, 100830. [[CrossRef](#)]
9. Wang, X.; Liu, T.; Guo, S.; Thornton, M.A.; Gui, P. A 2.56-Gb/s Serial Wireline Transceiver That Supports an Auxiliary Channel in 65-nm CMOS. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* **2020**, *28*, 12–22. [[CrossRef](#)]
10. ISO 11898-1:20039(E); Part 1: Data Link Layer and Physical Signalling. International Standards Organization (ISO): Geneva, Switzerland, 2003.
11. Chen, W.; Hong, C.; Wen, X.; Thornton, M.A.; Gui, P. Controller Area Network (CAN) Bus Transceiver with Enhanced Rail Converter. In Proceedings of the IEEE Midwest Symposium on Circuits and Systems (MWSCAS), Springfield, MA, USA, 11–14 August 2024.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.