

Transport protocols and message processing inside Linux

Dr. Oliver Hartkopp

Author

Dr. Oliver Hartkopp
Volkswagen AG
P. O. Box 1777
DE-38436 Wolfsburg

Link

www.volkswagenag.com

Today CAN networks are used in a wide range of control applications. Usually the content-based addressing is practiced to transfer information by sending single CAN frames. In the automotive context periodical transmissions are used to detect the absence of the data source within a reasonable time. E.g. when the period status information of a CAN controlled media player is not received every 200 ms the media player can be disabled in the HMI of the vehicles Infotainment system.

Additionally so-called CAN transport protocols are used in the vehicles on-board diagnostics (OBD) to provide a virtual point-to-point connection between the diagnosis tester and the CAN-connected ECU. With help of these CAN transport protocols (like ISO 15765-2) PDU (process data unit) length up to 4095 byte can be transferred via CAN, which is used e.g. for firmware updates.

Character device drivers for CAN

Having these two use cases in mind the known CAN drivers for Linux were not able to fulfill any of them appropriately in 2002. Due to the fact that known CAN drivers based on the character device driver model the interface to the CAN controller was simple and di-

rect – as known from character based serial drivers. The various CAN character device drivers presented a more or less abstract programming interface that is specific to CAN controller capabilities or vendor requirements. Especially the

“For novice CAN application developers there is no break in the programming philosophy known from other networking technologies and the stable socket programming interface promises a protection of investment. Especially CAN drivers that found their way into the Linux mainline kernel are continuously maintained and fixed in a community process so that they stay operational for all upcoming Linux kernel versions.”

Dr. Oliver Hartkopp

functionality provided by the CAN driver is reduced to apply CAN identifier filters.

To implement content filters for the payload of cyclic CAN messages these messages had to be processed in user-space context. E.g. if only a state change in the payload was relevant for the application each received CAN message had to be transferred from kernel-space into user-space for a comparison with the former one.

The implementation of a CAN transport protocol has to fulfill restrictive timing requirements down to a

few milliseconds. This cannot be ensured in the user-space context where several processes in a multitasking system share the same CPU. Depending on the system load the processes get their CPU time regulated by the system scheduler and the resolution of the system timer. The common alternative to implement CAN transport protocols for multitasking operating systems is to have a separated embedded CAN CPU or to use a real-time variant of the selected multitasking operating system (OS).

CAN networking in Linux

The point-to-point network communication in multitasking operating systems is well known from the Internet protocol (IP) communication. Different connection oriented and connectionless protocols based on the IP like the TCP and UDP are state of the art in common operating systems. Adapting this internet networking technology to realize the similar requirements of CAN transport protocols leads to an implementation of the CAN access inside the OS network stack.

This approach has several advantages over the formerly described character device driver model:

- ◆ Standard programming interfaces for system calls ('network sockets') ▷

- ◆ Standard network driver model known from Ethernet drivers
- ◆ Established abstraction layers in the network stack
- ◆ Communication protocols implemented inside the operating system context
- ◆ Multiuser access to the network

The definition of a network device driver interface for CAN is the first step to unify the CAN network access. As known from Ethernet networking the driver abstraction allows the replacement of networking hardware without modifying the user applications like e-mail clients or web browsers. On the other hand the CAN network device definition allows the hardware vendor to focus on the driver development without being in charge to specify and implement user-space programming interfaces and tools.

The direct access to network devices can be performed with a privileged

access to the PF_PACKET socket interface. But using the PF_PACKET protocol family sockets for CAN network devices has some vital drawbacks:

- ◆ Privileged access rights are needed (administrator only)
- ◆ No full network transparency for local CAN applications
- ◆ No efficient traffic filtering based on CAN identifiers
- ◆ No integration of CAN specific (transport) protocols

Especially the missing network transparency (second bullet) turns out as a knock-out criterion in a multiuser environment, as different CAN applications would have a different view of the existing CAN traffic.

The PF_CAN network protocol family

To overcome with the limitations caused by accessing CAN network devices with PF_PACKET sockets a new protocol family for

the Controller Area Network has been created. The integration of this PF_CAN protocol family into the Linux network stack can be derived from existing network protocol families like DECnet, Appletalk or AX.25 that also use the Linux networking infrastructure with different networking hardware and protocols.

To establish a separate data path for CAN messages and CAN devices the data buffers and network devices are identified with new type definitions:

- ◆ ETH_P_CAN identifies network data buffers that contain CAN messages
- ◆ ARPHRD_CAN identifies CAN network devices that handles network data buffers marked as ETH_P_CAN type

Based in these definitions a new PF_CAN protocol family implementation can register itself to be responsible for ETH_P_CAN marked buffers containing CAN messages.

Together with the PF_CAN/AF_CAN value defini-

tions and a new CAN socket address structure the missing link for the socket-programming interface for user-space applications is established. From this point the protocol family PF_CAN provides a framework that manages the CAN data flow from the user to the CAN network driver and vice versa.

Inside the operating system PF_CAN offers programming interfaces for various CAN based protocols. Comparable to UDP and TCP being different protocols of the protocol family PF_INET for internet protocol networking new CAN relevant functionalities can be hosted inside the PF_CAN framework. The different CAN protocols can be accessed from user-space applications with different protocol numbers that are functional available when a CAN protocol implementation registers itself at the PF_CAN framework.

To provide a multiuser access to CAN traffic the PF_CAN core offers inter-

SPI CAN driver for Linux

On the 13th iCC, Sakari Junnila introduced in his presentation the Wapice Custom Can Driver (WCCD), which is targeted for embedded Linux and optimized for ARM-based platforms. He presented also the results of some measurements comparing the WCCD approach with SocketCAN and LinCAN drivers. They author working with Wapice, a Finnish CiA member, has tested his driver's performance with extremely high bus-loads and have been able to receive over 20000 CAN messages per second. To achieve this, they made significant amount of optimizations also to the platform and low-level CAN and SPI driver code for the MCP 2515 controller outside of the SocketCAN driver framework. But as the evaluation of the executed performance measurements showed, the usage of the WCCD driver architecture provides benefits over alternative CAN driver implementations.

The biggest downside for the usage of the WCCD is the lack of a standard interface. Standard software cannot be used with it without porting it to the WCCD driver API. However, embedded systems applications are often proprietary and developed to some specific purpose.

Reference: Sakari Junnila and others – Design of high-performance CAN driver architecture for embedded Linux. In: 13th iCC proceedings, Nuremberg 2012.

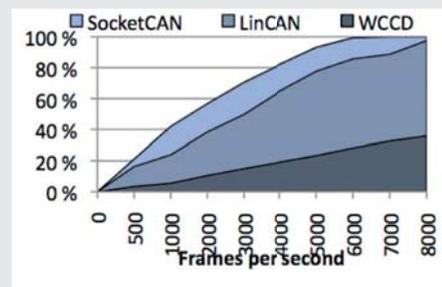


Figure I: CPU load for received messages

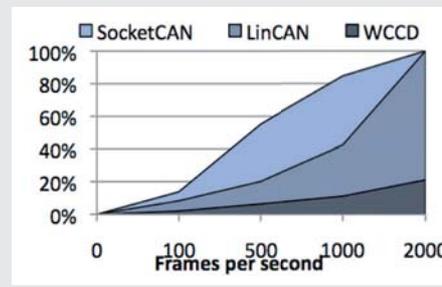


Figure II: CPU load for echoed messages

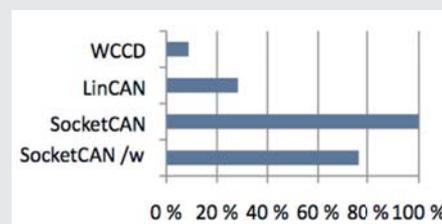


Figure III: CPU load when transmitting 1000 messages per second

Summary

Since version 2.6.25 (issued 2008-04-17) the Linux mainline kernel supports the network protocol family PF_CAN providing standardized programming interfaces for CAN users and CAN driver developers. This article provides an overview of the implemented technologies and challenges to integrate CAN interfaces into a non-real-time multiuser/multitasking operating system. Due to the standardized network driver model for CAN hardware a wide range of CAN controllers and CAN IP-cores are supported by Linux out-of-the-box. In opposite to usual embedded CAN ECUs the Linux networking system is designed to handle multiple CAN applications using multiple CAN networks at the same time. The integration of the CAN infrastructure into the networking stack allows implementing CAN transport protocols like ISO 15765-2 or high-performance CAN frame gateways inside the operating system context. The article provides solutions for expected prioritization issues when executing multiple CAN applications and summarizes requirements for Linux-preferred CAN controller concepts.

SocketCAN: The “official” CAN-API for Linux

Marc Kleine-Budde from Pengutronix discussed in his 13th iCC paper the tradeoffs of the SocketCAN approach. He also presented the CAN-BCM broadcast manager, which is used for example to support the J1939 and ISO 15765-2 transport protocols. The first ideas of a socket-based networking stack for CAN devices went back to 2002. There were several CAN implementations for Linux available back then, and some still are. Nowadays, SocketCAN is the official CAN application programmer interface (API) to the Linux kernel. It has been included in the kernel more than three years ago. Although, the kernel internal interfaces are stable, there is a constant evolution in the kernel going on. Re-evaluation of existing concepts, improvements, additional features and consolidation doesn't stop before the SocketCAN core or it's drivers. For example, the error handling in the CAN drivers is going to be consolidated and unified.

Reference: Marc Kleine-Budde – SocketCAN: The official CAN API of the Linux kernel. In: 13th iCC proceedings, Nuremberg 2012 (see < <http://www.can-cia.org/fileadmin/cia/files/icc/13/kleine-budde.pdf> >)

nal functions to CAN protocol implementations to send and receive CAN frames:

`can_send()`

Send a data buffer containing a CAN frame to a CAN device. Performs sanity checks for CAN frames (e.g. dlc) and ensures the local echo functionality.

`can_rx_register()`

Register a callback function that is executed on the reception of a CAN identifier that matches the given mask/value filter.

`can_rx_unregister()`

Remove a formerly registered subscription.

Depending on the given mask/value filter different filter-sets are used to reduce the CPU consumption when checking the registered subscriptions at runtime. The per-interface created filter-sets are checked in a software-interrupt on every CAN frame reception. Due to the described multiuser requirement the CAN drivers intentionally do not support any system wide CAN identifier filters. The performance of the software-interrupt based CAN filters has been evaluated in [4].

Higher-layer protocols in PF_CAN

The protocol CAN_RAW offers a similar programming interface as known from the CAN character device drivers. Analog to opening a character device file in Linux (e.g. `/dev/can0`) the CAN application programmer creates a network socket and reads/writes specific data structures representing CAN frames.

Additional to the fact that multiple instances of CAN_RAW sockets can be

created at the same time several different CAN identifier filters can be applied to each created socket separately. Due to the different filters for each CAN_RAW socket a specific view to the selected CAN network traffic can be archived and separately handled by different applications and users on the system.

The CAN_RAW sockets identifier filters reference directly to the per-interface filter-sets provided by the PF_CAN framework. The simple callback registration delivers the subscribed CAN frames directly to the requesting socket instance to be read by the user-space CAN application.

In automotive networks, the CAN messages are not only used to transfer signal values in the CAN frames payload. Sending CAN frames periodically e.g. every 200 ms allows detecting failures of the originating CAN node. This mode of operation has two drawbacks regarding the bus-load and the message processing. Even when the reception of CAN traffic is reduced by filters and/or message buffers the CAN frame payload needs to be checked whether the relevant signal has changed.

To reduce the effort for payload content filtering the CAN broadcast manager protocol (CAN_BCM) has been integrated into

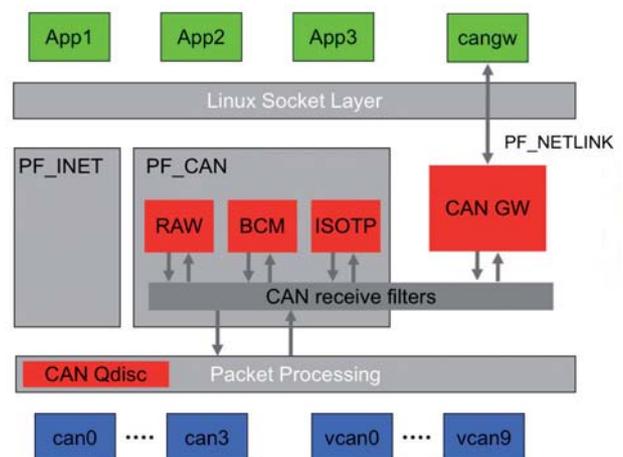


Figure 1: CAN network layer protocols and CAN frame processing



We offer CAN and CANopen training for development engineers and system integrators. Topics include application fields, physical layer, protocol, communication services, standardization, and certification.

CAN training

2012-10-17 Nuremberg (DE)
2012-12-11 Nuremberg (DE)

CANopen[®] training

2012-10-04 Zürich (CH)
2012-10-18 Nuremberg (DE)
2012-10-25 Essen (DE)
2012-11-06 Nuremberg (DE)
2012-11-07 Praha (CZ)
2012-11-15 Helsinki (FI)
2012-12-04 Gothenburg (SE)
2012-12-12 Nuremberg (DE)

Special CANopen[®] training

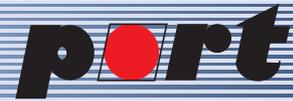
2012-09-27 Safety Nuremberg (DE)
2012-11-14 Special purpose cars Nuremberg (DE)

In-house seminars

CiA offers in-house seminars for companies with specific training requirements. The agenda of an in-house seminar is tailored to the customer's needs. CiA trainers are experienced in teaching CAN technology.

For more details please contact the CiA office at
events@can-cia.org

www.can-cia.org



www.port.de

Professionals for Industrial Communication



port provides to you a faster **Time-to-Market** with powerful DesignTools and solid and mature Protocol Stacks

port connects Industrial Worlds
reliable – fast – worldwide



Contact: +49 345 / 7 77 55 0
port GmbH +49 345 / 7 77 55 20
Regensburger Str. 7b service@port.de
GERMANY – 06132 Halle/Saale www.port.de

PF_CAN. The CAN_BCM is a programmable content filter and timeout handler for incoming CAN frames and can also manage the sending of cyclic messages in operating system context.

Due to the Linux high-resolution timers the precise sending of messages is independent from process scheduling. The content filter that checks changes in the 64 bit payload is processed in the software interrupt at CAN frame reception time to ensure that only relevant changes are passed to the CAN application in user-space.

CAN_BCM transmission path functions:

- ◆ Cyclic transmission of a CAN message using a given interval time
- ◆ Redefinition of CAN message content and interval timers at runtime
- ◆ Counting of performed intervals with automatic switching to a second interval time
- ◆ Immediate transmission of changed CAN message content with/without starting the timer interval
- ◆ Single transmission of CAN messages

BCM receive path functions:

- ◆ Receive filters to check changes in relevant CAN message elements (payload, data length code)
- ◆ Receive filters for single CAN identifiers (no content matching)
- ◆ Receive filters for multiplex CAN messages (e.g. with indices inside the CAN message payload)
- ◆ Receive filters for monitoring the data length code of the CAN message
- ◆ Respond to CAN remote frames
- ◆ Time-out-monitoring of CAN messages
- ◆ Reduction of the update rate for content filter changes messages (throttling)

Depending on the use case payload content filtering

and throttling can provide a significant reduction of the system load (see [5] p. 102-103). Additionally the comfortable CAN_BCM socket-programming interface reduces the complexity of the CAN application and supercedes a complex and imprecise timer handling in user-space.

Transport-protocols in OS context

Known solutions that implement CAN transport protocols use an embedded CAN processor in the CAN hardware interface to ensure the ambitious protocol timing requirements. The availability of Linux high resolution timers for the precise sending of CAN messages lead to the question if these timing requirements of CAN protocols can be handled inside the PF_CAN framework. A public available implementation [6] of the ISO 15765-2 CAN transport protocol has been evaluated and compared to a commercial diagnosis tool. It turned out that minimum response time of the Linux implementation was at least twice the time and in the worst case ten times longer than with the embedded CAN CPU solution. As the 4 ms delay of the worst-case scenario was still covered by the specified protocol timeout of 1000 ms the PF_CAN based implementation is conform to ISO 15765-2. Besides the measured overall timing ranges the open source implementation processed the CAN protocol messages two times faster than the embedded solution (average values).

CAN frame routing using PF_CAN

In Internet protocol networking the routing and forwarding of IP traffic through different network devices is common practice. The routing and modifying operations for IP packets are

based on IP addressing schemata. As the qualified CAN addressing is based on the CAN identifier and the CAN interface the existing routing implementations inside the operating system context are not suitable for the Controller Area Network.

Attempts to route CAN frames efficiently with user-space applications like candump failed through the existing process scheduling which lead to drops and non-deterministic delays depending in the size of the per-socket receive buffer size.

To integrate an efficient CAN message routing in Linux the created CAN gateway (CAN_GW) makes use of the PF_CAN filter subscription infrastructure. The registered filters trigger callback functions in the CAN gateway where the received CAN frames can optionally be modified before they are sent to the outgoing CAN interface. The one-hop CAN message forwarding supports the following on-the-fly modifications of the CAN frame elements CAN-ID, CAN-DLC, CAN-DATA:

- ◆ AND: logical 'and' element with value
- ◆ OR: logical 'or' element with value
- ◆ XOR: logical 'xor' element with value

- ◆ SET: set element to new value

After performing one or more of these optional CAN frame modifications in the described order a potentially corrupted payload data checksum can be built on demand. Additional to a simple XOR checksum calculation for a given data set, three CRC8 profiles can be selected that can calculate a CRC8 value based on an individual 256-byte CRC table. With this functionality the CAN gateway is able to create Autosar End-to-End-Library compliant CAN messages on-the-fly even after modification.

The implementation of the CAN gateway is optimized towards performance as the optional modification and checksum calculating operations are executed in the software interrupt context. Detailed performance measurements of CAN_GW have been made. The CAN gateway implements the Linux routing capability for the protocol family PF_CAN and supports the netlink configuration interface that is also used for the routing configuration of internet protocol data.

CAN frame traffic shaping

Due to the multiuser capabilities of the CAN network- ▶

deltaT	ID	data
0.200	123 _h	00 00 00 20
0.200	123 _h	00 00 00 21
0.200	123 _h	00 00 00 22
0.100	730 _h	02 03 04 05 06 07 08 09
0.001	730 _h	02 03 04 05 06 07 08 09
0.001	730 _h	02 03 04 05 06 07 08 09
0.001	730 _h	02 03 04 05 06 07 08 09
0.001	730 _h	02 03 04 05 06 07 08 09
(500 CAN frames later)		
0.001	730 _h	02 03 04 05 06 07 08 09
0.001	730 _h	02 03 04 05 06 07 08 09
0.001	123 _h	00 00 00 23
0.001	123 _h	00 00 00 24
0.001	123 _h	00 00 00 25
0.001	123 _h	00 00 00 26
0.200	123 _h	00 00 00 27
0.200	123 _h	00 00 00 28

Table 1: Data traffic without shaping



CiA

Product Guides

Comprehensive catalogs for CANopen, CAN & J1939 products and services

- ▶ Products listed are highly integrated into the new CAN Newsletter Online
- ▶ Find products by search function or by name, manufacturer, category, and profile implemented
- ▶ Advertise CANopen, CAN, and J1939 products, software and services
- ▶ Dedicated features, standardized and edited entries enhance readability
- ▶ Find worldwide contact data for each company

Linked to CiA's Product Guides.
Integrated into CiA's website.
Complementing CAN Newsletter.
Associated to iCC proceedings.



For more details please contact the CiA office at
publications@can-cia.org

www.cia-productguides.org



NEUER FREIRAUM

Energy 1500 W digital Primärschaltregler

EINSATZBEREICHE

- Automatisierungs- und Prüfabläufe
- Labortechnik

FÜR HOHEN LEISTUNGSBEDARF

- Spannungs-, Strom- und Leistungsregelung
- Alle Signale und Schnittstellen potentialgetrennt
- Standardmäßig mit RS 232, USB und CAN-Schnittstelle mit CANopen-Protokoll (CiA); LAN, IEEE optional
- Programmierbare Sicherheitsfunktionen
- Einfache Bedienung über Quick-Taste und Drehknopf
- Schnelle Umprogrammierung sowie Rückspeisefähigkeit durch 300 W E-Last
- Integriertes Mikroprozessorsystem zur zentralen Verknüpfung aller Daten

Kniel macht den Unterschied.

KNIEL System-Electronic GmbH
D-76187 Karlsruhe · Tel. 0721 95920 · www.kniel.de

Summary

The Linux subsystem for CAN standardizes the programming interfaces for CAN driver programmers and CAN application programmers. In comparison to a programming interface with a character device driver model, the chosen approach allows to implement various CAN specific communication protocols inside the operating system as well as the reuse of established and powerful networking techniques like network traffic control. The clear and simple network device driver interface known from Ethernet drivers leads to an easy exchangeability of CAN hardware without changing the existing applications. The open-source development process in the Linux kernel enables the contribution of new CAN drivers and CAN based protocols for everyone in an open community and is independent from single vendors.

Related articles

This article as iCC proceeding:
<http://www.can-cia.org/fileadmin/cia/files/icc/13/hartkopp.pdf>
 Dr. Sakari Junnila:
Design of high-performance CAN driver architecture for embedded Linux
 (in iCC proceedings 2012; see < <http://www.can-cia.org/fileadmin/cia/files/icc/13/junilla.pdf> >)

ing implementation in Linux a challenge to manage the host access to a single CAN medium arises. For example a CAN application A sends status information on CAN-ID 0x123 every 200 ms. Application B sends a 4095 byte PDU via ISO 15765-2 without a tx-delay requested from the communication partner. This leads to an expected transmission sequence shown in the Table.

Due to the fact that the ISO 15765-2 protocol pushed more than 500 CAN frames en-bloc into the CAN network device queue the relevant status information is not send on the CAN network for the ISO-TP frames transmission time. Depending on the used bitrate the block transfer time the cyclic status information can be inhibited which leads to inconsistencies and timeout errors in the data sink.

The described data flow effect is a known problem in Internet protocol networking too. To be able to browse the web in a smooth way while having peer-to-peer data communication the traffic can be controlled depending on its content. This so-called 'traffic shaping' for peer-to-peer connections is part of a traffic control framework inside the Linux networking. The various traffic control capabilities to prioritize, throttle or drop packets are designed to handle Internet protocol traffic and are not aware of Controller Area Network identifiers. Therefore the classifiers to sort and separate packets into network queues have been extended with a new CAN classifier implementation.

This CAN classifier support enables the usage of the entire traffic control (TC) capabilities of the Linux networking subsystem although not every queuing discipline is suitable for CAN use cases. The traffic control configuration requires administrative access rights, which

concentrate the host configuration to a single point. This allows setting up the local CAN node fulfilling the requirements of the different CAN applications and the other CAN devices.

This setup can be used to solve the initial example with the modified 'tc' tool (see listing).

```
# create a can0 prio
traffic control
handle
$ tc qdisc add dev
can0 root handle 1:
prio

# sort CAN-ID 0x123
to a separate queue
$ tc filter add dev
can0 parent 1:0 prio
1 \can sffid 0x123
flowid 1:1

# sort CAN-ID 0x124
to a separate queue
$ tc filter add dev
can0 parent 1:0 prio
2 \can sffid 0x124
flowid 1:2

# catch the rest
into the default
class
$ tc filter add dev
can0 parent 1:0 prio
3 \can sffid 0x0:0x0
flowid 1:3
```

When the CAN identifier 0x123 is to be sent, the CAN frame is sorted into the priority 1 queue. The frames in the priority 3 queue are sent to the CAN network when the higher priority queues are empty.

An alternative utilization for traffic control can be the throttling of outgoing traffic with a token bucket filter. This can be used to slow down the traffic on a local virtual CAN interface to have a realistic throughput like 125 kbit/s even without real CAN hardware. Usually the runtime created virtual CAN network devices are not limited in bandwidth like the loopback device for local Internet protocol traffic.

Multiuser requirements obviously need to be solved on the host differently than

in known embedded systems which has an impact on the CAN network driver implementation. Most CAN controllers provide a set of several transmit objects that can be accessed via different memory registers. Based e.g. on the CAN identifiers stored in the transmit objects an "intelligent" CAN controller is able to decide, which object to transmit first. To be able to specify the sequence of sent CAN messages with queuing disciplines the CAN controller has been configured in a simple FIFO mode. Therefore only one transmit object is used in the Linux CAN network drivers even when the CAN controller itself provides more than one transmit object. To ensure the maximum CAN throughput so-called shadow registers are used if available. Shadow registers allow writing the next CAN message into the CAN controller while the current transmission is in progress.

The CAN classifier for Linux queuing disciplines is currently evaluated with a prototypic implementation. Analogue to the CAN gateway this traffic control approach has the potential to become part of the Linux mainline kernel. ◀