Poster: Real ORNL Automotive Dynamometer (ROAD) CAN Intrusion Dataset

Samuel C. Hollifield, Miki E. Verma, Michael D. Iannacone, Robert A. Bridges, Bill Kay, and Frank L. Combs
Oak Ridge National Laboratory
{hollifieldsc, vermake, iannaconemd, bridgesra, kaybw, combsfl}@ornl.gov

This extended abstract and poster summarize our paper under review [6] describing the ROAD Dataset https://0xsam.com/road/. Modern vehicles are increasing in complexity and automation, relying on constant communication of small embedded devices called electronic control units (ECUs). This communication is typically accomplished via controller area networks (CANs). CAN furnishes a broadcast bus protocol, meaning each node that is connected to the network receives all the messages being transmitted. While CAN is a lightweight and dependable technology, it has many demonstrated security vulnerabilities. For instance, CAN does not have authentication nor sender/receiver information in messages, and it is difficult to functionally encrypt CAN messages.

CAN packets or frames are comprised of a few technical components. Most importantly for vehicle function are the Arbitration Field and Data Field.

- **Arbitration Field**: Contains an *Arbitration ID (AID)* which serves to label and assign priority to the message.
- **Data Field**: Contains the functional payload of the CAN message. The data field contains a maximum of 8 bytes.

Security research for CAN-based technologies has seen appreciable growth in the last few years, particularly targeting intrusion detection systems (IDSs), although many proof-of-concept attack works exist. The current taxonomy of attacks has been widely accepted from Cho & Shin's research into vulnerabilities [4] with three primary categories:

- **Fabrication Attacks** involve injected messages with malicious AIDs and data fields, the simplest attack. Examples of *fabrication attacks* include denial of service (e.g., sending AID 0x000 at a high frequency to overwhelm the bus), fuzzing attacks (sending messages with random AID and payloads), and targeted AID attacks (injecting messages with a specific AID and altered data field).
- **Suspension Attacks** involve techniques that silence or prevent a node from communicating. Examples of this attack include Cho & Shin's *Bus Off Attack*, which exploits the error-handling features of CAN to prevent a targeted ECU from communicating with the rest of the network [3].
- **Masquerade Attacks**, the most sophisticated of this trichotomy, silence the transmission of a legitimate message and inject malicious data field contents in its place. Miller & Valasek's infamous 'Jeep Hack' used a version of a masquerade attack to silence a safety-critical ECU and disable the braking mechanism [5]. The *bus off attack* mentioned earlier could also be the initial step of a *masquerade attack*.

Data Problems

Current CAN security research is limited by two major problems: data obfuscation and availability.

Data Obfuscation Problem: Manufacturers of passenger vehicles obfuscate the CAN message contents. This has created an asymmetric approach to IDS research, as many publications explore the relationship between message arrival times & physical layer attributes as opposed to functional signal values.

Data Availability Problem: It is costly and difficult to produce reliable attack datasets for three primary reasons. First, the difficulty of producing advanced attacks can inhibit researchers. Most current datasets provide a wealth of fabrication attacks but fewer more advanced attacks. It is relatively simple to connect to an automotive CAN and inject basic messages, but more advanced exploits often require a dedicated research vehicle, with its associated expenses. Second, the safety of equipment, occupants, and bystanders imposes an inherent risk when attacking a moving automobile. Thus, a dedicated facility with property safety precautions (such as a rolling dynamometer) are desired when launching attacks that have unknown consequences. Finally, disclosure of sensitive information is often an inhibitor to data release. Since manufacturers may consider their CAN data or their CAN message encodings as intellectual property, researchers must responsibly disclose and release data in a way that does not to expose themselves to potential litigation.

To our knowledge, there exists six datasets for CAN IDS research. Only three of these datasets include real, verified attacks on an automobile. The remaining three datasets use

This manuscript has been co-authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan). This research was funded by the Laboratory Directed Research and Development (LDRD) Program of the Oak Ridge National Laboratory (ORNL), managed by UT-Battelle, LLC, for the U.S. Department of Energy under Contract DE-AC0500OR22725.

Dataset	Org.	Year	Real Attacks	# Logs
CAN Intrusion (OTIDS)[1]	HCRL	2017	X	3
Survival Analysis for Automobile IDS[2]	HCRL	2018	X	3
Car Hacking for Intrusion Detection[3]	HCRL	2018	X	4
SynCAN[4]	Bosch	2019		5
Automotive CAN Bus Intrusion v2[5]	TU Eindhoven	2019		7
Can Log Infector[6]	CrySyS Lab	2020		7

TABLE I: Comparison of available CAN datasets

synthetic (not real) data entirely, or simulate injection attacks by adding messages in post-processing. See Table I.

Approach

To solve the shortcomings in existing CAN IDS datasets, we present the Real ORNL Automotive Dynamometer dataset[7] [6] consisting of 33 attack captures totalling about 30 minutes, and 12 ambient captures containing about 3 hours of ambient data. The logs are enumerated in Table II. We collected the CAN data using SocketCAN [1] software on a Linux computer with a Kvaser Leaf Light V2 connected to the OBD-II port. Data is saved in the standard SocketCAN ".log" format. All data is collected from a single, undisclosed vehicle. The data is obfuscated in such a way to not provide specific details about the vehicle while retaining important aspects of the data for an IDS. All attack data was generated on a four wheel rolling dynamometer with verified physical reaction by the vehicle. The ambient captures were collected both on the dynamometer and on-road.

A. Attacks

Attack captures are detailed below in order of complexity.

1) Fuzzing Attack: We mounted a less stealthy version of the fuzzing attack, which injects frames with random AIDs, with a maximum payload of `0xFFFFFFFFFFFFFFFF` every 0.005s. Many physical effects are observed as a result of this attack: accelerator pedal is unresponsive, dash and warning lights activate, fan blows at max speed, etc.

[1]http://ocslab.hksecurity.net/Dataset/CAN-intrusion-dataset. The attacks are not labeled and documentation on the injection intervals is unclear and possibly incorrect.

[2]http://ocslab.hksecurity.net/Datasets/survival-ids. The attacks are very noisy due to the exceptionally high frequency of injections, thus they could be detected with a very simple frequency-based detector.

[3]http://ocslab.hksecurity.net/Datasets/CAN-intrusion-dataset. At the conclusion of each attack, there is a large gap in the messages being transmitted. We suggest researchers using this dataset, particularly for intrusion detection, trim the attack captures to the time preceding the gap.

[4]https://github.com/etas/SynCAN. The data provided is entirely synthetic which leads to a 'cleaner' experience than real data. Additionally, authors claim that all ambient data should be used for training but do not provide additional ambient data for testing; thus, it is difficult to test an IDS's false positive rate.

[5]https://doi.org/10.4121/uuid:b74b4928-c377-4585-9432-2004dfa20a5d. Timestamps are altered in post-processing which alters data and fidelity. Further, many attacks are not realistic. For instance in the DoS example, 10s worth of frames are overwritten which is not how a real DoS attack appears. Further, the messaged are dispersed too greatly to affect vehicle functionality.

[6]https://www.crysys.hu/research/vehicle-security/. The attacks added are entirely in post processing and have a limited number of options. Notably, *Can Log Infector* allows users to change only whole bytes. Considering CAN signals are often dispersed across multiple (or partial) bytes, this means signals will be altered unrealistically.

[7]Dataset is available at https://0xsam.com/road/

TABLE II: Logs in ROAD CAN Intrusion Detection Dataset

Attack Description	Modified	# Logs
Accelerator Attack (In Drive)		2
Accelerator Attack (In Reverse)		2
Correlated Signal Fabrication Attack		3
Correlated Signal Masquerade Attack	X	3
Fuzzing Attack		3
Max Engine Coolant Temp Fabrication Attack		1
Max Engine Coolant Temp Masquerade Attack	X	1
Max Speedometer Fabrication Attack		3
Max Speedometer Masquerade Attack	X	3
Reverse Light Off Fabrication Attack		3
Reverse Light Off Masquerade Attack	X	3
Reverse Light On Fabrication Attack		3
Reverse Light On Masquerade Attack	X	3
Dynamometer Various Ambient		10
Road Various Ambient		2

2) Targeted ID & Masquerade Attacks: Fabrication attacks were mounted with a *flam* delivery, meaning a message is injected immediately after an ambient message with the target AID. This allows for a dynamic injection, as the legitimate message is effectively overwritten by the malicious injection.

3) Accelerator Attacks: The *accelerator attack* is an advanced attack that does not fit into the general taxonomy of injection attacks. This attack is specific to the make/model of the vehicle used for our experiments and has been responsibly disclosed to the manufacturer. We will not disclose details of how to implement this attack. Instead, we include the CAN data post-exploit. The effect afterward is that the driver has far less control of the vehicle, specifically in terms of acceleration when in gears D (drive) and R (reverse). The accelerator attack does not have maliciously injected message, but contains the CAN data from the vehicle in this altered state.

Discussion

Our dataset contains examples of attacks with varying levels of sophistication and deployment mechanisms. Notably, we were able to verify vehicle malfunction on each of the attack logs captured. Further, our dataset was verified as a key improvement in CAN datasets by Blevins & Moriano et al. in their research of time-based CAN IDSs [2].

References

[1] SocketCAN. https://python-can.readthedocs.io/en/master/interfaces/socketcan.html.

[2] Deborah H. Blevins, Pablo Moriano, Robert A. Bridges, Miki E. Verma, Michael D. Iannacone, and Samuel C Hollifield. Time-based CAN intrusion detection benchmark. preprint: https://arxiv.org/abs/2101.05781, under review, 2021.

[3] Kyong-Tak Cho and Kang G Shin. Error handling of in-vehicle networks makes them vulnerable. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1044–1055. ACM, 2016.

[4] Kyong-Tak Cho and Kang G Shin. Fingerprinting electronic control units for vehicle intrusion detection. In *25th USENIX Security Symposium USENIX Security 16)*, pages 911–927, 2016.

[5] Chris Valasek and Dr Charlie Miller. CAN message injection. http://illmatics.com/can%20message%20injection.pdf, Jun 2016.

[6] Miki E. Verma, Michael D. Iannacone, Robert A. Bridges, Samuel C. Hollifield, Bill Kay, and Frank L. Combs. ROAD: The Real ORNL Automotive Dynamometer Controller Area Network Intrusion Detection Dataset (with a comprehensive CAN IDS dataset survey & guide). preprint: https://arxiv.org/abs/2012.14600, under review, 2020.

Real ORNL Automotive Dynamometer (ROAD) CAN Intrusion Dataset

Samuel C. Hollifield, Miki E. Verma, Michael D. Iannacone, Robert A. Bridges, Bill Kay, Frank L. Combs

Oak Ridge National Laboratory

Controller Area Networks & Security

What is a Controller Area Network?

The Controller Area Network (CAN) protocol is a message-based standard which is implemented in all modern automobiles to control vehicle function and share information between electronic control units (ECUs).

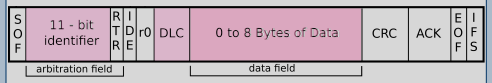

CAN messages exist with many components. Most important for vehicle operation are the arbitration field and data field:

- **Arbitration Field**: Labels and prioritizes messages, also referred to as an *AID*.
- **Data Field**: Payload binary which contains command or state information.

Although CAN is ubiquitous in modern vehicles, it is bereft of security measures. Previous research has demonstrated exploitations which could impact property and safety. Thus, CAN-based intrusion detection is a growing field of research.

CAN-Based Attacks

CAN lacks message encryption and authentication and is critically vulnerable to exploitation. Current research describes three primary categories of attacks[1]:

1. Fabrication Attacks: Injected messages with malicious AIDs and data fields. Examples of *fabrication attacks* include denial of service (such as sending AID 0x000 at a high frequency to prevent transmission of other messages), fuzzing attacks (messages with random AID and payloads), and targeted AID attacks (messages injected with a specific AID and manipulated data field)

2. Suspension Attacks: Preventing or silencing an ECU from sending one or more messages

3. Masquerade Attacks: The most sophisticated category, *masquerade attacks* involve suspending the transmission of a legitimate message and replacing the data field with a malicious payload.

Problem

Reliable CAN Data with Labeled Attacks is Scarce

CAN data with high-fidelity labeled attacks is unavailable for several reasons:

- Such data is costly to produce, except for simple-injection attacks.
- Producing realistic CAN attack data carries inherent risks to property, drivers, and passengers.
- OEMs consider their CAN encodings as intellectual property, thus the disclosure of sensitive information inhibits data release.

Intrusion Detection Dataset Comparison

Dataset	Org.	Year	Real Attacks	# Logs
CAN Intrusion (OTIDS)	HCRL	2017	✓	3
Survival Analysis for Automobile IDS	HCRL	2018	✓	3
Car Hacking for Intrusion Detection	HCRL	2018	✓	4
SynCAN	Bosch	2019		5
Automotive CAN Bus Intrusion v2	TU Eindhoven	2019		7
Can Log Infector	CrySyS Lab	2020		7
ROAD CAN Intrusion Dataset	**ORNL**	**2020**	**✓**	**10**

There are CAN IDS datasets which exist in the research community; however, the ROAD CAN Intrusion Dataset provides high quality attack samples with:

- Physically verified attacks on real vehicles with a stealth injection mechanism.
- Simulated masquerade attacks which are otherwise unproducible.
- Advanced attacks that entail no injected messages.

References

[1] K.-T. Cho and K. G. Shin, "Fingerprinting electronic control units for vehicle intrusion detection," in Proceedings of the 25thUSENIX Security Symposium, 2016, pp. 911–927

[2] SocketCAN, https://python-can.readthedocs.io/en/master/interfaces/socketcan.html.

[3] can-utils, "https://github.com/linux-can/can-utils"

[4] D. H. Blevins and P. Moriano and R. A. Bridges and M. E. Verma and M. D. Iannacone and S. C. Hollifield, "Time-Based CAN Intrusion Detection Benchmark", *https://arxiv.org/abs/2101.05781 (under review)*

[5] Miki E. Verma, Michael D. Iannacone, Robert A. Bridges, Samuel C.Hollifield, Bill Kay, and Frank L. Combs. ROAD: The Real ORNL Automotive Dynamometer Controller Area Network Intrusion Detection Dataset. preprint: https://arxiv.org/abs/2012.14600 under review, 2020.

Approach

Produce a Dataset With Numerous Verified Attacks

- Real ORNL Automotive Dynamometer (ROAD) Dataset (https://0xsam.com/road/, DOI: 10.13139/ORNLNCCS/1728694) [5]:
  - 33 attack captures totaling ~30 minutes of dynamometer driving
  - 12 ambient captures containing ~3 hours of both dynamometer and on-road driving
- Data was collected using SocketCAN[2] software on a Linux computer with a Kvaser Leaf Light V2 connected to the OBD-II port.
- All data is collected from a single vehicle and obfuscated to maintain the anonymity of the vehicle while preserving aspects important to an IDS.

Data Obfuscation

- Absolute timestamps may all be shifted by a scalar, but relative times are preserved.
- Messages from particular AIDs that were deemed unimportant were replaced with a "filler message" (ID#Data in hex) *FFF#0000000000000000.*
- Messages on reserved IDs (e.g., greater than 0x700, diagnostic messages) have been removed.
- Arbitration IDs have been anonymized in such a way that *arbitration order/priority is not preserved*. There is a one-to-one mapping between the original and anonymized AIDs.
- Data fields have been scrambled such that signals have been preserved, and the fields are scrambled consistently for each AID.

Example Message Distribution (Max Speedometer Attack)

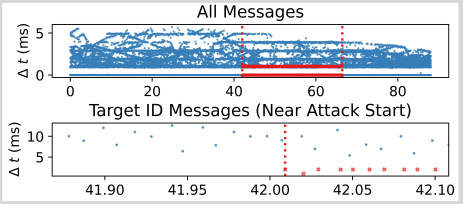

Example of timing/frequency of messages during the max speedometer attack. Notably, our injection method does not generate a large amount of 'noise' due to the stealth/flam delivery.

Results

To our knowledge, this represents the most diverse collection of verified attacks on an automotive CAN.

Our dataset was verified in research with Blevins & Moriano et al with a comparison study on time-based CAN IDS technologies[4].

Log Information

- All of the CAN data files are logged using the standard can-utils[3] candump format:

Unix Timestamp	Channel	AID (hex)	Data Field (hex)
(1569510697.667343)	can0	5E1#	893FE0070A000080

Logs in ROAD CAN Intrusion Detection Dataset

	Modified	# Logs
Accelerator Attack (In Drive)		2
Accelerator Attack (In Reverse)		2
Correlated Signal Fabrication Attack		3
Correlated Signal Masquerade Attack	✓	3
Fuzzing Attack		3
Max Engine Coolant Temp Fabrication Attack		1
Max Engine Coolant Temp Masquerade Attack	✓	1
Max Speedometer Fabrication Attack		3
Max Speedometer Masquerade Attack	✓	3
Reverse Light Off Fabrication Attack		3
Reverse Light Off Masquerade Attack	✓	3
Reverse Light On Fabrication Attack		3
Reverse Light On Masquerade Attack	✓	3
Dynamometer Various Ambient		10
Road Various Ambient		2

Attacks Available

1. Fuzzing Attack: Frames are injected with random IDs and maximum payloads (0xFFFFFFFFFFFFFFFF) every 0.005s.
2. Targeted ID Fabrication & Masquerade Attacks: These attacks are performed using a *flam* injection technique, meaning a message is injected immediately when a target ID is seen. In a *masquerade attack*, we remove the preceding message in post-processing. The attacks are as follows.
   - **Correlated Signal —** A message which contains each wheels' speed is injected with four false wheel speed values (each two bytes) that are different. The car rolls to a stop and inhibits acceleration.
   - **Max Speedometer —** The speedometer signal (one byte) is targeted by modifying the signal value to the maximum (0xFF). The car falsely displays a maximum value on the instrument cluster.
   - **Max Engine Coolant Temperature —** We target the engine coolant signal (one byte), modifying the signal value to the maximum (0xFF). The car falsely displays an "engine coolant too high" warning.
   - **Reverse Light —** A binary (one bit) signal which communicates the state of reverse lights (on/off). The car illuminates the reverse lights during the injection.
3. Accelerator Attack: An advanced attack that does not fit into the general framework. This attack exploits a vulnerability of the particular make/model that puts the ECU into a compromised state. This vulnerability has been responsibly disclosed and we do not provide vehicle data during the exploit. However, the driver experiences less control such as a lack of accelerator pedal input and fixed acceleration to a constant speed.

Contact Information

- Samuel C. Hollifield, hollifieldsc@ornl.gov

Cyber Security Research Group
Cyber Resilience and Intelligence Division
Oak Ridge National Laboratory

Acknowledgements

This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan(http://energy.gov/downloads/doe-public-access-plan)

This research was funded by the Laboratory Directed Research and Development (LDRD) Program of the Oak Ridge National Laboratory (ORNL), managed by UT-Battelle, LLC, for the U.S. Department of Energy under Contract DE-AC0500OR22725

Oak Ridge National Laboratory
MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY