CHALMERS

Proof of concept for Ethernet in Steer-by-wire

Master's Thesis in System, Control and Mechatronics
and Secure and Dependable Computer system

MATTIAS NILSSON
ROBERTH FRIBERG

Department of Signals & Systems
Communication Systems
Chalmers University of Technology
Gothenburg, Sweden 2012
Master's Thesis 2012:78

Proof of concept for Ethernet in Steer-by-wire

R. Friberg
M. Nilsson

Examiner: E. Ström

Chalmers University of Technology
University of Gothenburg
Department of Signals and Systems
SE-412 96 Göteborg
Sweden
Telephone + 46 (0)31-772 1000

Department of Signals and Systems
Göteborg, Sweden June 2012

Abstract

As the complexity of electrical systems on vehicles increases, so does the amount of communication. As of today communication often relies on controller area network (CAN) technology. The limitations of CAN in terms of bandwidth have begun to constitute a reason for concern within vehicle and boat industry and a change in technology is becoming more and more desirable, as proven by some major initiatives (from FlexRay to OPEN consortium). Earlier work at CPAC has shown standard Ethernet to be a possible candidate. This thesis aims to take that work further and to build a proof-of-concept, a steer-by-wire system that can be tested on a boat.

This study, about how a complete Ethernet based system could be structured, revealed several challenges but showed that the construction of a functional proof of concept system was possible.

The system was built using existing CAN to Ethernet converters, where the software was further developed to increase reliability and robustness. An evaluation of the use of non-standard cables with standard Ethernet components showed that the communication was sensitive to electromagnetic disturbance, a subject which requires further study to properly determine properties such as the cost, weight and resilience of different types of cables suitable for use in Ethernet communication. Several types of connectors, from different vendors, were found to be possible alternatives when designing a system for the marine environment. A major part in the design of an Ethernet system is the choice of network topology. Different topologies give the network different properties, which must be matched with the desired properties of individual systems. The possibility to easily use off the shelf products was seen as an obvious advantage, allowing for the assembly of a working demonstrative system with cheap wireless routers, an Ethernet camera and standard Android devices.

Keywords: Ethernet, CAN, Steer-by-wire, EMC, 100BASE-TX, Topology, Cables

Acknowledgements

First of all we would like to thank Marco Monzani at CPAC Systems, Erik Ström and Sun Wanlu at Chalmers University of Technology that made this thesis possible. Further we would like to thank all CPAC employees that has helped us with practical aspects as well as filling gaps in our knowledge. Last but not least we want to show our appreciation to other thesis workers at CPAC Systems that helped to make this work a very enjoyable journey.

Roberth Friberg and Mattias Nilsson, Gothenburg 2012-06-19

List of abbreviations

ACK Acknowledge

ARP Address Resolution Protocol

AUI Attachment Unit Interface

BPDU Bridge Protocol Data Unit

CAN Controller Area Network

CFI Canonical Format Indicator

CLT Central Limit Theorem

COTS Commercially available Off-The-Shelf

CRC Cyclic Redundancy Check

DLC Data Length Code

DTE Data Terminal Equipment

DUT Device Under Test

ECU Electronic Control Unit

IEEE Institute of Electrical and Electronics Engineers

EMC Electromagnetic Compatibility

EOF End Of Frame

FRNT Fast Reconfiguration Network Topology

ICMP Internet Control Message Protocol

IDE Identifier Extension

IEC International Engineering Commission

IGMP Internet Group Management Protocol

IP Internet Protocol

LSFR Linear Feedback Shift Register

MAC Media Access Control

MAU Medium Attached Unit

MII Media Independent Interface

NIC Network Interface Card

NMEA National Marine Electronics Association

PCS Physical Coding Sublayer

PHY Physical Layer of the OSI model

PLS Physical Layer Signalling

PMA Physical Medium Attachment

PMD Physical Medium Dependent

RFC Request For Comments

RSTP Rapid Spanning Tree Protocol

RTR Remote Transmit Request

SFD Start of Frame Delimiter

SNMP Simple Network Management Protocol

SOF Start Of Frame

SRR Substitute Remote Request

STP Spanning Tree Protocol

TP Twisted Pair

UDP User Datagram Protocol

VLAN Virtual Local Area Network

Contents

1

Introduction

CPAC Systems is a Volvo Group company that develops integrated safety-critical electronic control systems for every type of commercial vehicle. To keep ahead of competitors, the company needs to constantly improve current products and develop new functionality.

The last decades have brought an enormous leap in electronics and computer science. This change has not left the automotive industry untouched. The amount of electronics installed in an average vehicle has been steadily increasing. With the increasing number of electronic controllers comes an increasing demand for communication. Current systems often use one or several controller area network (CAN) buses for communication. This is widely used technology that has been relatively unchanged in the last two decades. Properties of this bus, such as its bandwidth has however become more and more limiting and different possible substitutes such as Ethernet and FlexRay have come into the spotlight.

Standard Ethernet has long been dismissed for use in time critical networks due to its non deterministic properties. This view has however started to change and prior work at CPAC Systems has shown that it is currently possible to model a network to decide maximum latency and jitter.

The aim of this thesis is to continue this work and create a proof-of-concept demonstrating the use of a standard Ethernet backbone for a steer-by-wire system, developed for Volvo Penta by CPAC Systems. Instead of just substituting the existing link, non-critical traffic from a camera is routed over the same network to show the possibility of separated and prioritised data streams. This proof-of-concept was seen to be working when it had no effect upon the steer-by-wire system that is noticeable to an operator. The complete proof-of-concept system is also modelled using software from Time Critical Networks to estimate latency and jitter.

Apart from the proof-of-concept this report addresses some of the aspects that arise from a system implemented using only Ethernet. This includes a possible choice of cable type (with focus on fault tolerance and other electromagnetic properties), network topology and type of connector. A lot of different aspects and technologies that might be interesting in the design of a full system have not been considered since this only was meant to produce a working proof-of-concept. One such technology is Ethernet over fibre cable (e.g. 100BASE-FX) which due to time constraints and lack of commercially available shelf products only has been briefly discussed.

The report will initially give the reader the needed background and knowledge to fully understand the subsequent chapters. It then continues to discuss how the work was carried out, before actual results from these investigations are presented. Finally there is a discussion regarding the result and possible directions of future work.

2

Theory

This chapter is meant to give an introduction to CAN and Ethernet communication and form the knowledge base needed to fully grasp the upcoming chapters. This includes brief explanations of CAN and Ethernet standards and statistics.

2.1 CAN

Controller Area Network (CAN) is a multi-master, communication bus capable of operation at speeds up to 1 Mb/s. Data transmission is initiated upon events within a node and the data is packaged in frames (Figure 2.1) containing up to eight bytes of data. Since transmissions are initiated on uncorrelated events within separate network nodes there is a risk of transmissions occurring at the same time. Thus there must exist mechanisms to handle such collisions. Nodes will only send frames if they judge that the bus is available but two nodes could decide to start sending at the exact same time. In case of such event the protocol falls back on non destructive bit-wise arbitration. Bit-wise arbitration means that each node is listening to the bus while sending and back down if they see a dominant bit when they send a recessive bit [1].

2.1.1 CAN frame

Information on a CAN network is broadcasted using frames on a multi-master bus. The frame identifier specifies the contents of the frame and nodes can therefore choose to process only the data relevant to them. Two frame formats are defined, both of which are described below in Figure 2.1 [1].

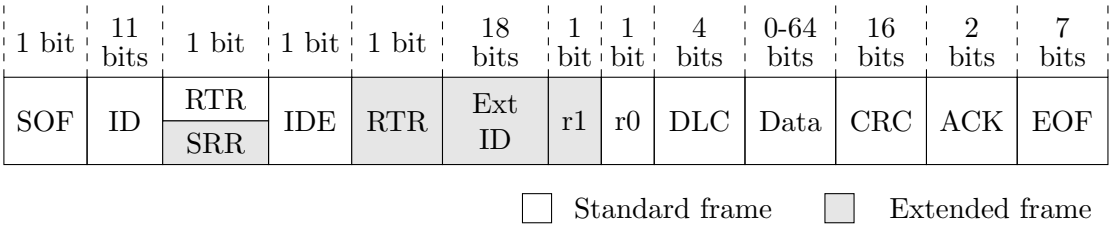

Figure 2.1: CAN frame.

SOF Start of frame, single bit indicating the start of a frame.

ID Identifier, eleven bits containing the identification number of a standard frame.

RTR Remote Transmit Request, used to request data from a node, which node is given by the identifier.

SRR Substitute Remote Request, works as a placeholder for the RTR of a standard frame when sending a extended frame.

IDE Identifier Extension, indicates the usage of an extended format.

Ext ID Extended Identification, extra bits for identification when using the extended format.

r1 and r0 Reserved bits for future use.

DLC Data Length Code, contains the data length.

CRC Cyclic Redundancy Check, checksum used for error detection.

ACK Acknowledge, the receiving nodes write a dominant bit in this slot to indicate that they received an error free frame.

EOF End Of Frame, indicates the end of the frame.

2.1.2 Arbitration

As previously stated, two nodes might initiate data transmission on the same bus at the same time. If this happens there is non-destructive bit-wise arbitration which means that the message that wins the arbitration is successfully transmitted and the concurrent message is retransmitted. In Figure 2.1 it is shown that the first field that is sent after a fixed preamble and SOF is the identifier which then correlates to the priority of the frame [1].

2.2 Ethernet

Communication may be established through a range of different media, common choices include copper cables and fibre optics. Except for the physical connection a variety of standards are used to enable an operational communication channel. These standards are needed for hardware configurations such as bit rate and signal amplitudes, as well as to define the meaning of the signal array that is to be passed over the channel.

Ethernet is a collection of such standards used to describe a well known and used framework. This collection defines a range of standards about different ways of signalling over different mediums. The standards are used to encode an array of bits to transmittable signals. A sequence of bits or information is called a frame which is illustrated in Figure 2.2. The Length/Type field defines the amount of data or if the data should be interpreted according to a different protocol such as IP [2].

7 byte	1 byte	6 byte	6 bytes	2 byte	46-1500 byte	4 byte
Preamble	SFD	Destination Address	Source Address	Length/ Type	Data	CRC

Figure 2.2: Ethernet frame.

This fundamental structure is very useful and is used to implement basic functionality such as frame filtering, self-identifying frames through the frame type field and contains information to verify that the message has been received intact in the cyclic redundancy check (CRC) field [3].

2.2.1 Layers

The Ethernet specification IEEE 802.3 specifies layered models of the communication. In Figure 2.3 the different layers used in 100BASE-T (100BASE-T connects 100 Mbit/s physical layers such as 100BASE-TX to IEEE 802.3 CSMA/CD MAC) is shown.

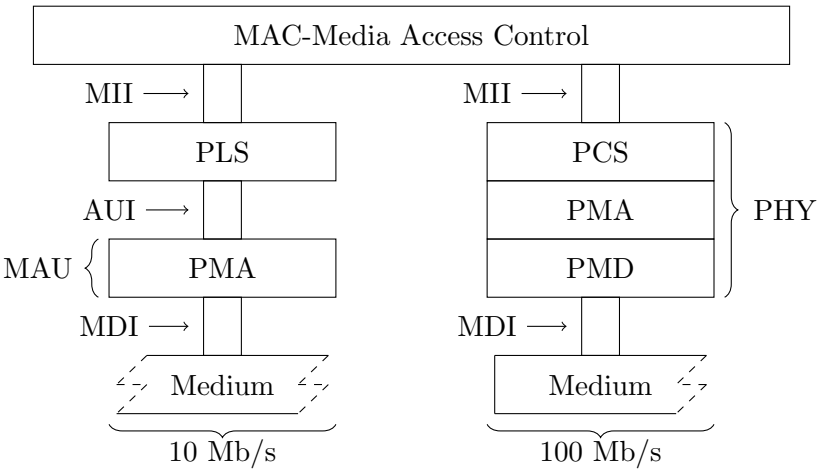

Figure 2.3: Communication layers in 100BASE-T.

Media Access Control (MAC)

The Media Access Control layer handles data and errors on a media independent level. It manages collision handling and avoidance as well as detection of physical medium transmission errors. Here the data is placed in frames (frame boundary delimitation and frame synchronisation) in which the MAC also places source and destination addresses (Section 2.2) [4].

Media Independent Interface (MII)

The Media Independent Interface is meant to provide a simple inexpensive and easy to implement interconnection between the MAC and the PHY for data transfers at 10 or 100 Mb/s. Data transmission is fully duplex over nibble (4 bits) wide data paths and synchronous to a clock reference. It is specified to use voltage levels compatible with a large amount of processors [5].

Physical Coding Sublayer (PCS)

The Physical Coding Layer is in charge of encoding and decoding data nibbles sent over the MII to or from five bit code groups (4B/5B). It will also generate Carrier Sense and Collision Detect indications as well as forwarding transmit and receive signals according to the MII specification. Serialisation and deserialisation to and from the underlying Physical Medium Attachment (PMA) is also performed on this layer [5].

Physical Medium Attachment (PMA)

The Physical Medium Attachment maps the transmit and receive code bits between the PMA's client and the underlying Physical Medium Dependent layer. It generates a control signal indicating the availability of the PMD to a PCS or other client and recover the clock from the Non-Return-to-Zero Inverted (NRZI, encoding in which a binary one is represented by a state transition and a binary zero as the lack of such transition) data supplied by the PMD. Also the Auto-Negotiation (when implemented) is done in this layer. Other optional features are sensing of receive channel failures and transmitting Far-End Fault indication [5].

Physical Medium Dependent (PMD)

The PMD layer in 100BASE-TX which handles the conversion to physical signals on the medium is specified by using the FDDI TP-PMD standard (ANSI X3.263:1995) with some modifications. These modifications can be found in section 25.4 in [5]. For further information on how this is done see Section 2.2.2.

Physical Layer Signalling (PLS), Attachment Unit Interface (AUI)

The Attachment unit interface is only present when using 10 Mb/s and is used to interconnect the Data Terminal Equipment (DTE) to a MAU that is not integrated as a physical component of the DTE. Using this, the DTE is provided with a media independent interface using an optional cable up to 50 meters [4].

Medium Attached Unit (MAU)

In the case of 100BASE-TX used at 10Mb/s the PMA is also called Medium Attached Unit. This name was before used (for example in 10BASE2) to specify both the PMA and the MDI [4].

2.2.2 Twisted Pair Physical Layer Medium Dependent (TP-PMD)

ANSI X3.263:1995 specifies the operation of the TP-PMD for 100BASE-TX and consists of three main parts; Scrambler/Descrambler, Encoder/Decoder and Driver/Receiver.

Scrambler/Descrambler

The scrambler is introduced to spread the transmission spectrum and by doing that minimising electromagnetic compatibility problems. It is also designed to be fully transparent to other parts of the protocol to decrease coupling [6].

The scrambling is performed by modulo 2 addition of the plaintext and a pseudorandom sequence. This sequence is generated using a Linear Feedback Shift Register (LFSR) with a length of 11 bits (see Figure 2.4) [6].

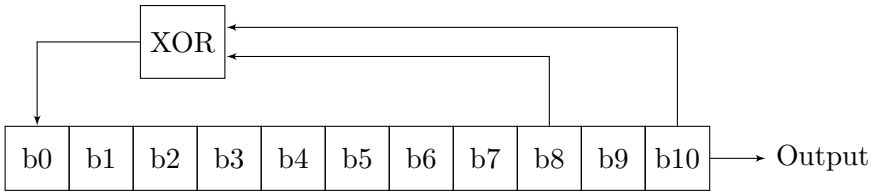

Figure 2.4: Linear Feedback Shift Register design.

The same operation is performed to descramble the sent text at the receiver side. To enable this, the receiver's LFSR must be synchronised with the transmitter's LFSR so they generate the same sequence. This synchronisation is achieved using the line states in which known sequences can be used to calculate the correct seed for the descrambler's LFSR [6].

Encoder/Decoder

The encoder converts the data stream from the scrambler into MLT-3 before presenting the data to the driver. The MLT-3 encoding consists of three levels and changes between them in a predefined way when the input is logical one. These changes can be seen in Figure 2.5 [6].

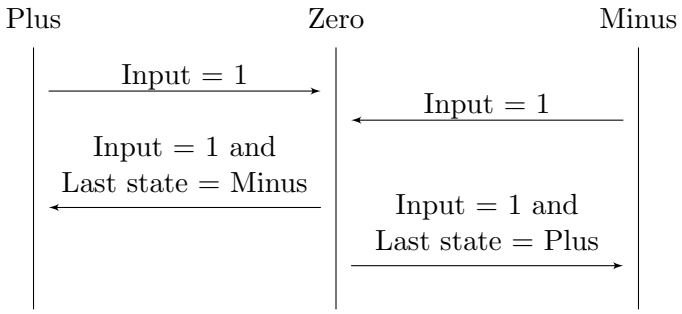

Figure 2.5: Encoder state diagram (MLT-3).

Driver/Receiver

The driver converts the MLT-3 data stream from the encoder to physical signals over a twisted pair. The different states are encoded using a differential voltage over the pair and can by that take the values plus, zero and minus.

2.2.3 Topologies

Ethernet can be used in different connection schemes. These have different advantages and disadvantages which need to be taken into consideration when an Ethernet based system is designed. The upcoming sections will illustrate different network topologies using circles for nodes and lines between them to show connections.

Point to point/Daisy chain

A point to point topology is the simplest of the different network topologies. It consists of two nodes directly connected by a single medium as shown in Figure 2.6 [7].

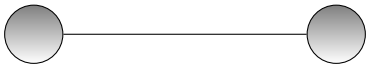

Figure 2.6: Point to point network.

This can also be extended by Daisy chaining where more devices are connected on a line (Figure 2.7) [7].

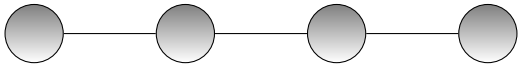

Figure 2.7: Daisy chained network.

Bus

The bus topology consists of a single cable that links all the nodes together (Figure 2.8). This comes with the advantage that it is simple and cheap to wire and it is also very easy to add nodes since they only need to tap into the existing cable. The single cable configuration does however also make fault isolation difficult and the network hard to diagnose. Since all communication is done over the same wire there is also a risk for collisions [8].

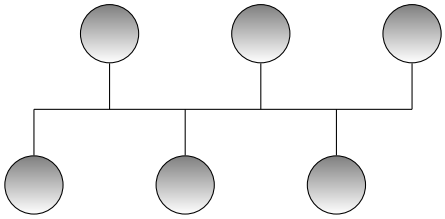

Figure 2.8: Bus network.

Star

In a star network the nodes are not directly connected to each other but all communication goes via a central unit such as a switch (Figure 2.9). The point to point configuration of the star makes the network more fault tolerant since a faulty link only affects a single node. But if the central unit fails the whole network breaks down. Using one cable between the central node and each other node can however result in a large amount of cabling for certain physical layouts which increases installation cost [8].

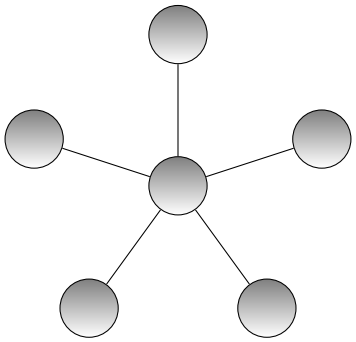

Figure 2.9: Star network.

Ring

In a ring network (also called circular network) each node is connected to two others in a way that they all form a ring (Figure 2.10). When a message is sent it is passed from node to node in a given direction until it reaches its destination or end up back at the sender node. Here a shorter length of transmission medium is often needed in comparison with the star topology. A faulty node can cut of the traffic going through it but it is less sensitive than the bus topology since the data can be transferred in principle both ways [8].

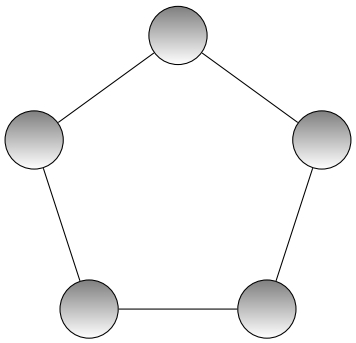

Figure 2.10: Ring network.

Mesh

In a mesh topology every node is connected to all (complete mesh, Figure 2.11) or some (partial mesh, Figure 2.12) of the other nodes. In the complete mesh topology there is no collisions or routing since all nodes are directly connected to all other nodes. This makes the network robust since traffic can be routed around a broken link but it also demands a large amount of cable. The partially connected mesh network need less cable and less interfaces per node but is not as resilient to link failure [8].

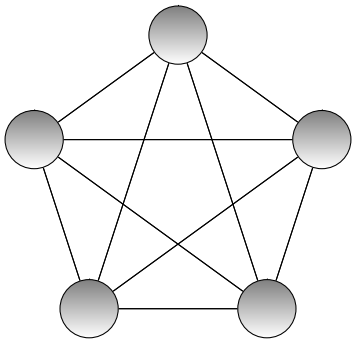

Figure 2.11: Fully connected mesh network.

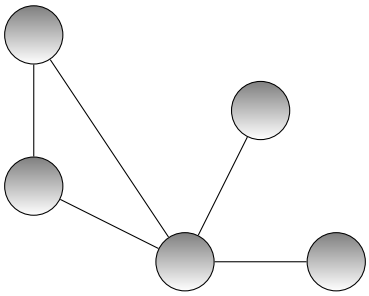

Figure 2.12: Partly connected mesh network.

Tree

A tree or hierarchical topology can be seen as a number of connected star networks. This makes the network easy to group and extend. This grouping can limit the extent of which a faulty node or link affects the network (see Figure 2.13). Apart from this possibility of grouping the properties of a tree topology is similar to that of a star topology.

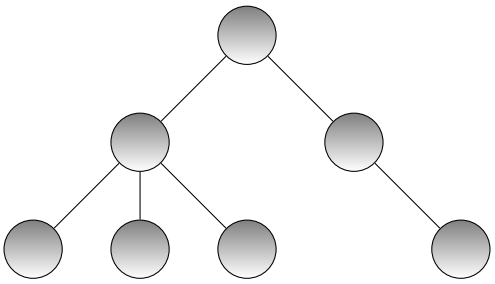

Figure 2.13: Tree network.

2.3 Structure protocols

On top of the Ethernet protocol there exist several different protocols to define and standardise communication on a level with higher abstraction.

2.3.1 Internet Protocol (IP)

The internet protocol implemented on top of Ethernet further specifies the data in the frame. Except the added logical addresses (called IP address) the protocol also specifies several control parameters which gives basic support to handle fragmented information. It also contains the time-to-live parameter that is decremented for every hop a frame makes and works as a guarantee that the frame cannot traverse the network for eternity. Due to that both the header and the data itself have non-fixed sizes, two length parameters are also embedded in the control part of the IP datagram. Figure 2.14 shows a non-detailed description of an IP datagram [3].

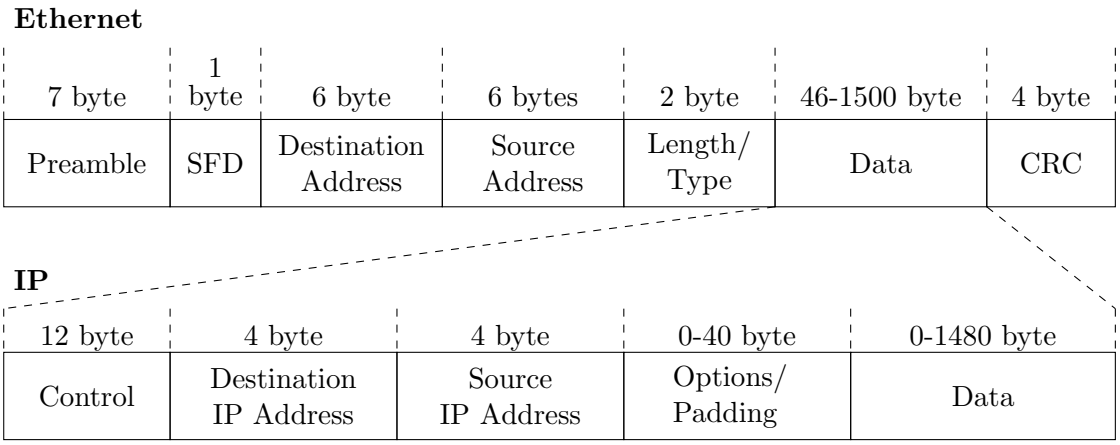

Figure 2.14: Illustrate how a IP frame is contained in an Ethernet frame.

The logical addresses are to uniquely define every device on the network, meaning that an address is only allowed to occur once on a network. They are controllable in that if a node is replaced the new node can be configured to have the same address as the replaced node. Furthermore the double addressing with both a physical and logical address is critical when a package is to be routed. The logical addresses are static while the package traverses the network whilst the physical addresses change at every hop. A node's physical addresses can be resolved given the logical address through the address resolution protocol (ARP) [2].

Today there exists a newer version of the IP, IP version 6 (IPv6). In this report all references to IP are referring to IP version 4 (IPv4).

2.3.2 User Datagram Protocol (UDP)

Another important piece of information that can be found in the control part of the IP header is the protocol parameter. This value defines how the data in the IP datagram is to be interpreted. The value corresponds to protocols such as UDP, TCP, ICMP and IGMP [2].

The UDP defines a connectionless and unreliable way for applications to exchange information over a network. There is in other words no built-in functionality to allow the sender to get a confirmation that the data is received correctly. This does however also mean that the protocol is not adding a large header, see Figure 2.15, nor needs it a procedure to establish a connection with the receiver to send data [2].

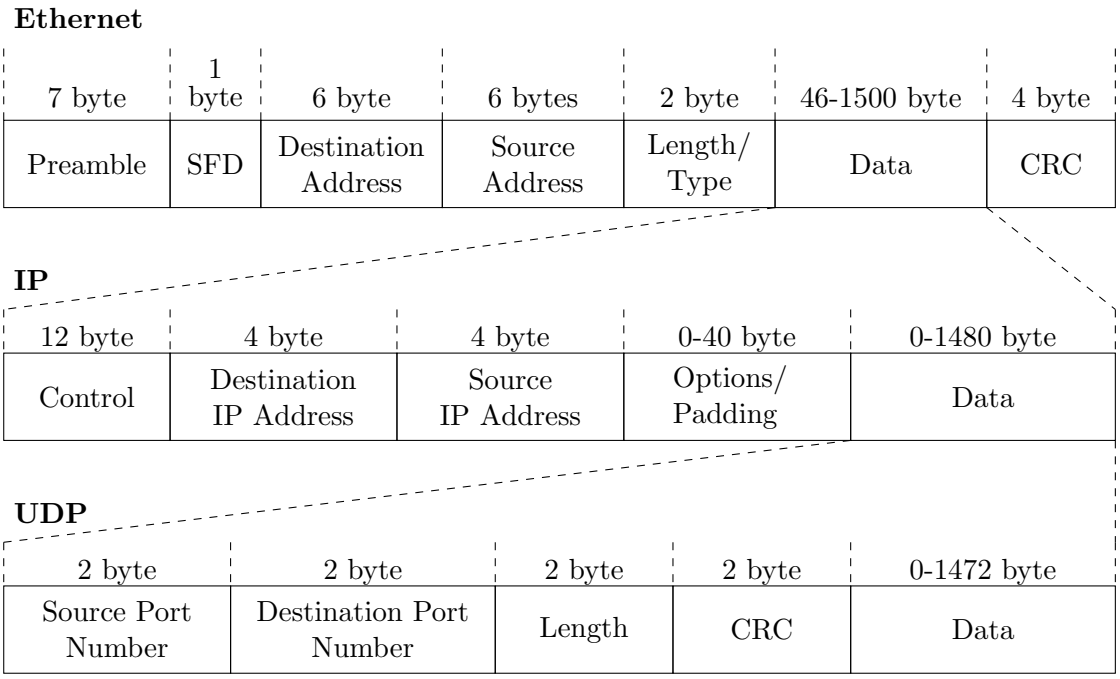

Figure 2.15: Illustrate how a UDP frame is contained in an IP frame which in turn is contained in an Ethernet frame.

The port numbers defined in the header of the UDP frame is used to direct the data to the right application at the host, which can have multiple applications using UDP to communicate with other nodes on the network [2].

So to conclude, when using UDP three different addresses are used for application to application communication where the third is used to address the right application within the receiver device.

2.4 Service mechanisms

Except for application communication protocols such as UDP and TCP, there is a collection of service protocols. These service mechanisms differ in purpose as well in the communication hierarchy level they are implemented on.

2.4.1 Internet Control Message Protocol (ICMP)

The Internet Protocol has no error control and lack assistance mechanisms. The ICMP utilise the IP itself to increase the possibility to diagnose the IP, and higher communication layers. This service is based on two different types of messages, error reporting and query messages. The error reporting can be done for a variety of reasons but are

mainly used to report that a frame has not been delivered, cannot be correctly received or could be forwarded more efficiently. An error report can be sent from routers as well as from the destination node. Note that unreported frame losses can still occur, but the protocol enables a framework for handling detected problems and errors [2].

The query messages were created to enable network diagnostic functionality. The two main queries currently in use are the echo and time stamp messages. The first is used to verify a working connection on the IP layer. The second is used to analyse the time needed for a frame to traverse the route between the sender and another node as well as to retrieve the destination node's universal time in milliseconds [2].

2.4.2 Simple Network Management Protocol (SNMP)

SNMP defines a framework that enables monitoring and managing of network devices such as routers and printers. The defined management protocol makes it possible to monitor devices from different brands and with different sets of features and performance. Further this open framework makes it possible for third party developers to develop software to manage the network.

The protocol is implemented on the UDP level and two specific ports are used for communication with the manager. One for the device to receive queries and another one to send traps, i.e. information about a present abnormal state of the device [2].

2.4.3 Internet Group Management Protocol (IGMP)

The protocol plays an important part in the enabling of multicast functionality in a network. The word multicast in computer communication context describes addressing of a group of nodes from a singe source, unicast one to one communication and broadcast one to all communication. Multicast addressing enables a more efficient bandwidth usage due to only the nodes concerned receiving the message, and the sender only has to send it once. This cannot be done with either unicast nor broadcast addressing.

IGMP defines how multicast routers can create and maintain a list linking the multicast internet protocol addresses to the prescribing nodes and their specific addresses. In a third version of this protocol the nodes cannot only specify multicast addresses of interest but also give specifications of "only allowed" or "not allowed" source IP addresses. The method of achieving a correct and updated list is based on a timer driven query system performed by the routers/switches, where each node is to report back [2].

The query interval and other timers and variables that affect the bandwidth and other hardware usage can be controlled. Details for the default and controllable parameters can be found in RFC 3376 [9].

2.4.4 Redundancy protocols

There exist several protocols to allow redundant wired links in an Ethernet network. The spanning tree protocol (STP) is a redundancy protocol based on the IEEE 802.1D standard which was published year 1990 and revised several times [10]. The main functionality of this type of protocol is to allow network devices (Bridge Protocol Data Units (BPDUs)) to have links between them without creating data link layer (LLC and MAC) (Layer 2) loops, causing broadcast storms [11] [12]. This is accomplished by detecting redundant links and deactivating them. If a active link becomes unavailable or broken the network will be re-evaluated and a deactivated link will be reactivated, if one exist.

As the name spanning tree indicates, the functionality is based on a graph data structure algorithm to find the optimal spanning tree. The links between are weighted differently depending on that specific link speed to create a root path cost. A high level description of the STP functionality can be described as: [12]

a) Elect root device based on given priority and device MAC address.

b) Elect root port by finding the port with lowest path cost to the root bridge, each non-root switch along the path adds a local cost.

c) Elect a designated port to handle traffic for each network segment. The port announcing the lowest root path cost becomes the designated port for a specific segment.

d) Remove bridge loops by deactivating ports that are neither set as root ports nor designated ports. Resulting in the removal of all data link layer loops in the network.

An improved version of the STP is the Rapid spanning tree protocol (RSTP) introduced in IEEE 802.1w (incorporated into IEEE 802.1D-2004) where a more complex algorithm poses a more rapid convergence during topology changes [12].

There also exist proprietary protocols for handling link redundancy. One example is the by Westermo developed Fast Reconfiguration of Network Topology (FRNT) protocol. The protocol can only be used in a ring topology in comparison to STP and RTSP which support mesh networks (topology independent). The fundamental limitation of FRNT enables a fast reconfiguration time when a redundant link is broken [11].

2.5 Virtual Local Area Network (VLAN)

The general concept of VLAN is to separate the logical connectivity from the physical connectivity using software configurations. The restriction of data streams are hereby not only set by physical cabling but also limited by filters in network devices that define subsets of the connected nodes. The VLAN is constructed by switch and/or end

station configurations. The precise possibilities of what can be achieved depends on the functionality provided by these devices [13].

The overall aim with this feature is to increase the security and the logic view of data routing through the network, often is a by-product of the implementation of VLAN a lower load on the network. If the only aim is to reach this by-product then multicast addressing is simpler to implement and sufficient for this purpose [13]. However, multicast addressing can still be a tool to decrease the network load within a VLAN.

One implementation of VLAN is the masked IP addressing, which allows subnets to be defined. These subnets are not able to communicate directly to each other, with the exception of when the limited broadcast address (255.255.255.255) is used as destination IP address[1]. To establish inter-subnet communication with unicast addressing, a router (layer 3 (network layer) switch) is needed [13].

The above described VLAN implementation is called implicit VLAN, which is a VLAN defined by rules based on subnets. Another implementation of VLAN is done by adding explicit tags (VLAN tag) to the frame. A common standard for tagged VLAN is the IEEE 802.1Q standard. The functionality is then embedded directly in the Ethernet frame. Figure 2.16 illustrates the VLAN-tagged frame defined by IEEE 802.1Q. The tag consists of two fields, each of two bytes. The first field, VLAN Protocol ID, is used to declare that the Ethernet frame is extended with the VLAN tag, this field has the same index as the type/length field of an non tagged Ethernet frame. The following field, Tag Control Info, carries the VLAN ID number. It also facilitates a three bit frame priority description and the Canonical Format Indicator (CFI). The priority feature is often referred to as 802.1P, even if such a standard has never been published. The CFI is interpreted differently depending on the technology of the network the frame traverses [13].

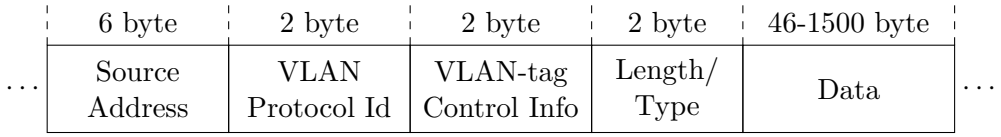

Figure 2.16: Illustrates how the VLAN tag is embedded in an Ethernet frame.

The VLAN-tagging provides both benefits and disadvantages. Tagged-VLAN nodes can belong to more than one VLAN, note that implicit VLANs can have this ability as well. With the help of edge switches (adds and removes the tags for incoming and outgoing frames respectively) can end nodes that do not support VLAN tagging communicate over a tagged VLAN. Exactly how this edge switch can be utilised depends on the implemented functionality on the network devices, as stated in The All-New Switch

[1]Note that a directed broadcast (network address followed by the broadcast host address) results in a frame being forwarded by layer two switches to all attached nodes. The nodes belonging to another subnet will discard the message when received. The directed broadcast may pose unnecessary traffic on the network.

book [13] *"The flexibility obtained from a commercial set of VLAN-aware switches will be a function of the features implemented in the devices, the limitations of those features (numbers of VLANs, permitted configurations, and so on), and the capabilities of the management software provided with the product."*. Other end nodes may support VLAN tagging (VLAN-aware) and therefore send and receive tagged frames, which can create additional possibilities compared to nodes that do not support VLAN tagging.

The functionality of explicitly tagged VLAN comes with the cost of increased complexity. Consider the mechanism that is implemented in switches to find and disable redundant connections to prevent bridge loops, it now depends on which VLAN the analysis is done for. Another disadvantage is the slight increase in data overhead due to that the four byte tag invalidates the Frame Check Sequence (implemented by a Cyclic Redundancy Check, CRC) of the Ethernet frame when inserted, which leads to additional computations [13].

2.6 Fail measurement statistics

The binomial distribution describes a collection of n independent samples from a stationary process where the outcome is either true with probability p or false (probability $1-p$) without regarding the order of the samples. The distribution's true mean value, μ and variance, σ^2 are described in equation 2.2 and equation 2.3 respectively. These can be derived from the distribution's probability density function, described by equation 2.1.

$$f(x) = \binom{n}{x} p^x (1-p)^{n-x} \tag{2.1}$$

$$\mu = np \tag{2.2}$$

$$\sigma^2 = np(1-p) \tag{2.3}$$

The central limit theory (CLT) states that an independent sample series of a stationary process will yield a normal distributed sample average of the measured process. The criteria for the CLT is that the measured process has finite expected value and variance and that the sample size is large enough. The sample average, $\bar{x}$ will approach normal distribution as the number of samples increase. The distribution of the sample average ($\bar{x}$) can be expressed as

$$\bar{x} \sim N\left(\mu, \frac{\sigma^2}{i}\right) \tag{2.4}$$

where μ is the mean and σ^2 is the variance of the stationary process, and i is the number of samples taken.

$\bar{x}$ and i is known when a collection of samples is retrieved from the process and in combination with a known variance (σ^2) for the process these parameters can be used to determine μ, with a certain percentile and confidence interval. A first step to show this is to bias and scale the normal distribution (equation 2.4) to obtain the standard normal distribution as done in equation 2.5.

$$\frac{\bar{x}-\mu}{\frac{\sigma}{\sqrt{i}}} \sim N(0,1) \tag{2.5}$$

Consider the special case where $n = 1$, for the binomial distribution to obtain the Bernoulli distribution with the expected value $\mu = p$ and variance $\sigma^2 = p(1-p)$. Together with equation 2.5 that yields equation 2.6.

$$\frac{\bar{x}-p}{\sqrt{\frac{p(1-p)}{i}}} \sim N(0,1) \tag{2.6}$$

A standard normal distribution and a chosen percentile of $1-\alpha/2$ yields an confidence level of $1-\alpha$. The confidence interval for the mean value of a measurement series can then be determined using equation 2.8. The confidence level is defined in equation 2.7 where $z_{1-\alpha/2}$ represents the value of the x-axis corresponding to the standard normal distribution and the used percentile.

$$1-\alpha = P(-z_{1-\alpha/2} < \frac{\bar{x}-p}{\sqrt{\frac{p(1-p)}{i}}} < z_{1-\alpha/2}) \tag{2.7}$$

$$I_p = \bar{x} \pm z_{1-\alpha/2}\sqrt{\frac{p(1-p)}{i}} \tag{2.8}$$

For a certain percentile and confidence interval the number of needed samples can be calculated as

$$i = z^2_{1-\alpha/2}\frac{p(1-p)}{d^2} \tag{2.9}$$

where d is the one sided length of the confidence interval. So given this information and the value of p, the needed number of measurements can be determined. But p is an unknown parameter that is to be estimated. An iterative procedure with the current estimated probability can be used to verify that the current number of measurements is enough.

Another approach to solve this problem is to calculate the worst case scenario, this will yield the smallest number of tests needed to establish the chosen confidence interval and percentile, regardless of the value of p. The function described by equation 2.10 and its derivative revealing the maximum point for the function which describes the worst case.

$$f(p) = p(1-p) \tag{2.10}$$
$$\Rightarrow \dot{f}(p) = 1 - 2p \tag{2.11}$$
$$0 = 1 - 2p \tag{2.12}$$
$$\Rightarrow p = \frac{1}{2} \tag{2.13}$$

Inserting equation 2.13 into equation 2.9 yields the number of needed measurements for the worst case as described by equation 2.14 .

$$i = z^2_{1-\alpha/2}\frac{1}{4d^2} \tag{2.14}$$

It should be clear that this reasoning is linked to the standard normal distribution and does not concern the uncertainties introduced by the approximation to this distribution motivated by the CLT. When p is one or zero the sample average will not be normal distributed. Close to the ends of the interval, a collection larger than shown above may be needed to not get a skewed distribution.

3

Method

This chapter describes strategies and the reasoning used in this thesis. It includes test set up for electromagnetic compatibility test, reasoning behind hardware selection and demonstration layout.

3.1 CAN to Ethernet converters

To minimise the impact on the current steer-by-wire system CAN to Ethernet converters were used for the proof-of-concept. This allowed a quick implementation of a complete working system using existing hardware.

3.1.1 Prior work

Construction of a CAN to Ethernet converter was the aim of a thesis work earlier carried out at CPAC [14]. The converter allowed the execution of specific timing measurements which in turn are required to identify the network model. In this thesis the initial work on these where to set up a build chain for the development boards that were used (IAR STM32F107VC-eval, Figure 3.1). A simple test in which the converters were used to bridge CAN communication was carried out to ensure proper functionality (Figure 3.2).

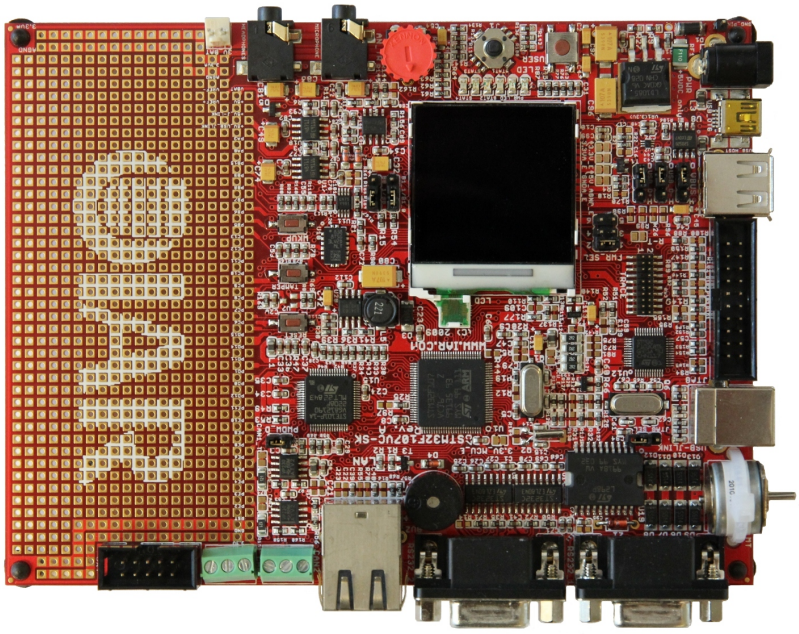

Figure 3.1: IAR STM32F107VC-eval evaluation board used for the CAN to Ethernet converter.

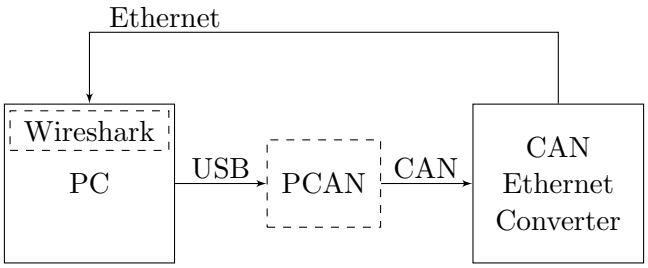

Figure 3.2: First CAN to Ethernet converter test setup.

3.1.2 Software robustness improvements

From a functional standpoint, the converter [14] essentially ready, as it could provide connectivity between CAN and Ethernet domains. However, a certain amount of work was required to make the device sufficiently robust for this application. The converters were now to be used in a real demonstration and needed therefore to be able to withstand harsher treatment without making the system fail or when system failure was unavoidable (i.e. when the communication channel is broken or saturated), re-establish the link as soon as possible without the need for any interaction.

To measure this resilience, tests were performed where first the CAN (using the same set up as the first converter test, Figure 3.2) and then Ethernet (Figure 3.3) inputs were subjected to a large amount of ingoing traffic.

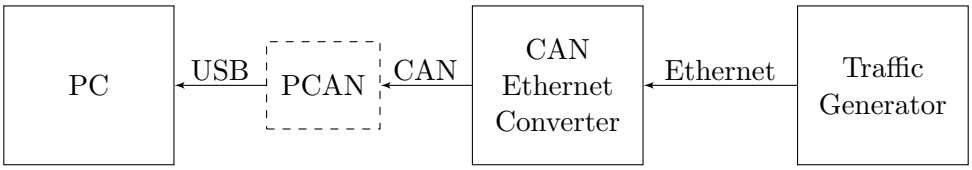

Figure 3.3: Ethernet overload test on converters.

3.2 Simple twisted pair cable

The cable that is standard for 100BASE-TX (Cat 5e) is not specifically developed for 100BASE-TX and has therefore properties that does not match the exact needs. A Cat 5e cable does for example contain four twisted pairs but only two are used. Due to this a cable consisting of two separate twisted pairs, much like two CAN cables where tested to see if it was compatible with 100BASE-TX.

3.2.1 UTP two pair cable construction

The cable was constructed using wires following the HD 21.7 standard ($0.75\,\mathrm{mm}^2$). 100 meter lengths of this cable was then twisted into two pairs using a drilling machine. After one of the pairs were twisted a communication test was carried out to decide if twisting the second pair was crucial. The pairs where then terminated with a male Deutch connector in each end (see Figure 3.4).

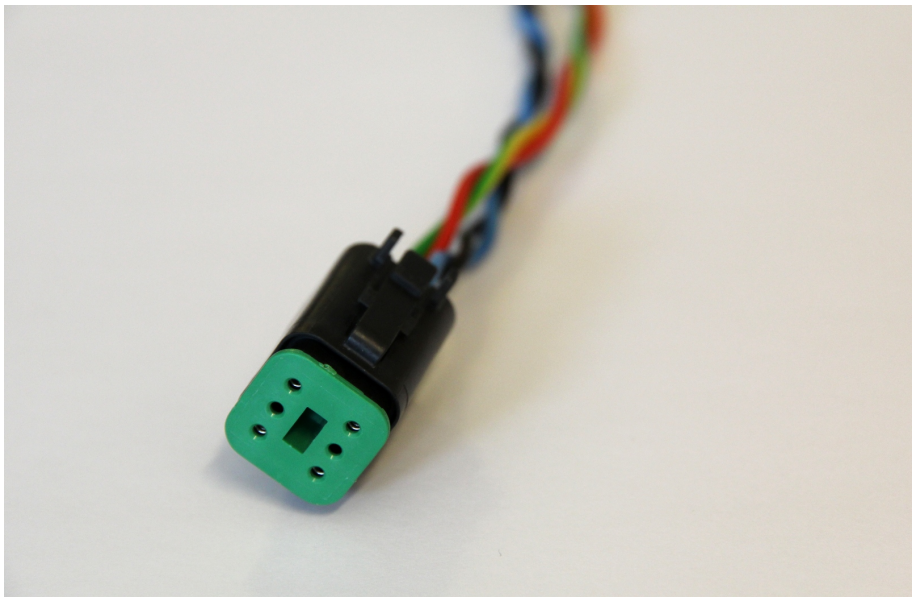

Figure 3.4: Deutch connector in the end of the constructed twisted pair cable.

3.2.2 RJ45 to Deutch converter

Since thicker non-conventional cables does not fit in a standard RJ45 contact a different connector had to be used for the cables that where tested. The CAN to Ethernet converters and a normal laptop has however no possibility to use a connector other than a RJ45. Therefore, short conversion cables (see Figure 3.5) where created to enable the use of a Deutch connector.

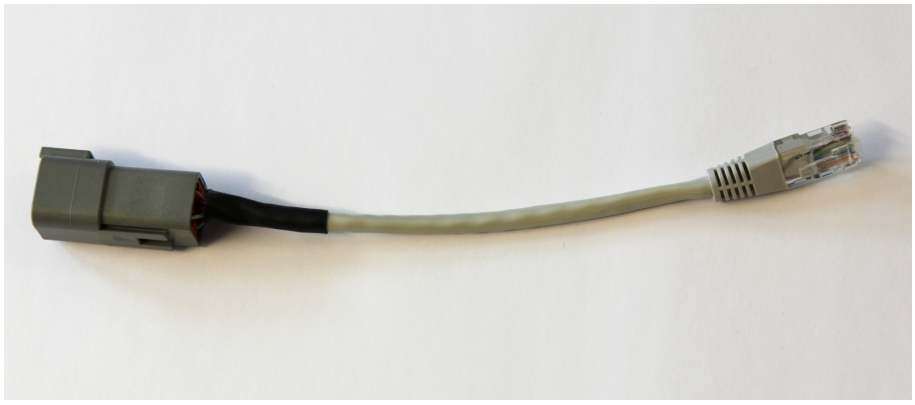

Figure 3.5: Converter between RJ45 to the right and Deutch connector to the left.

3.2.3 Testing

The finished cable was laid out along the office floor without any loops or bends except two 90 degree turns (roughly 10 cm in radius) approximately 10 and 20 meters from one end. It was then connected to two laptops using converters (Section 3.2.2) and the program used for EMC testing (described in Section 3.3.4) was used to send and verify 10 000 UDP frames.

A successful outcome of this test would indicate that the cable was valid for communication and therefore worth to test in the EMC test facility.

3.3 Electromagnetic Compatibility

Electromagnetic Compatibility tests where performed to get an approximation of the robustness of the communication link, regarding electromagnetic disturbance. This test also form the basis for recommendations of specific cable properties for a full installation. The cables tested are presented below together with the names they will be referenced by in quotations.

a) UTP Cat 5e 26AWG, "Cat 5e"

b) UTP Cat 5e 26AWG, using all four pairs, "parallel Cat 5e"

c) SF/UTP Cat 5e 26AWG, aluminium foil cable screen and braid, "shielded Cat 5e"

d) UTP Cat 6 23AWG, "Cat 6"

e) UTP two pairs 0.75 $mm^2 \sim [18 - 19]AWG$, "UTP two pair"

f) F/UTP two pairs 22AWG, cable screen of aluminium foil, "LonWorks"

Two different test types have been carried out, one for radiated emission and a second test for radiated immunity. The immunity testing was done with a functionality approach, the number of successful delivered frames is observed rather than the physical link behaviour.

3.3.1 Laboratory

The EMC tests were carried out in IVF's facilities in Mölndal, Sweden. The equipment allowed both radiated emission and immunity tests to be performed. The background noise for the emission test was measured to see how it compared to the measurements on the cables.

Another important aspect is the repeatability of the equipment used to measure the emissions. To test this two consecutive tests were run without any change in hardware.

The test rig does not contain any form of restraints for the cables and the placement will therefore differ between measurements of different cables. An extra test was run on the UTP two pair cable where it had been slightly moved to try and assess how small variations in placement affect the measured values.

3.3.2 Cable selection

A number of cables with different properties were tested. They had all been tested to work for a 10 meter cable in an office environment prior to the EMC tests and UTP two pair had been tested for a full length of 100 meters (see Section 3.2).

One aim of this test was to find out in what extent shield improves the resistance against radiated disturbance and attenuate the radiated emission. Another was to compare different cable types against each other where the cross-section area is one changing parameter. Even in the case where all these cables perform in a very similar manner this information is highly useful. Other aspects could then be evaluated to possibly reduce the number of cable candidates before further EMC testing.

The Cat 5e cable is the most common choice when using 100BASE-TX and has therefore served as a base case.

The fact that the Cat 5e cable contains four twisted pairs whilst 100BASE-TX only needs two opens up the possibility of connecting two parallel pairs. These can be used as a single pair to achieve redundancy in the cable. If the connection between the pairs is made between the transceiver and connector the same level of redundancy is achieved in the connector. But since this also changes the properties of the cable it needs to be tested separately.

The third cable to be tested is a shielded Cat 5e to try and determine a shields impact on radiated disturbance and emission.

The major difference between Cat 5e and Cat 6 is the attenuation on higher frequencies. Cat 5e is specified up to 100 MHz whilst Cat 6 is specified to be tested up to 250 MHz. This results in Cat 6 supporting more of the newer high speed protocols, which in turn makes it more probable that it will be available longer. Even though 100BASE-TX is not specified to need a bandwidth higher than 100 MHz, it is interesting to see if this has any impact on performance.

The last two cables to be tested are chosen to match the number of pairs needed. One of them consists of two parallel cables currently used for CAN (UTP two pair) and the other is a shielded cable with two pairs used in LonWorks systems.

It is not at this moment clear how these cables compare in price since listed prices only apply when acquiring small quantities. The result of this test was meant to narrow down the selection and make it easier to compare price and properties at a later stage.

3.3.3 Radiated emission

The emission test is to identify the energy amplitude of the radiated emission for different frequency bands. Important bands such as the VHF band are usually tested with higher resolution, i.e. a smaller band is gauged for each measurement. The measurements taken were however meant to give a rough indication on performance, and measurements using a narrower band for these frequencies would not have yielded enough additional information to justify the increase in the time needed for the tests.

The aim for the emission test was to see if the choice of cable had a noticeable impact on the electromagnetic emission. To test this, the cable under test was wound nine times around a block of Styrofoam and placed at a distance of one meter from the antenna (Figure 3.6 and Figure 3.7). During the test, a continuous stream of data was transferred between two evaluation boards in one direction. The data consists of UDP frames containing 19 bytes of data (for low level Ethernet data transmission see Section 2.2.2). The cards were controlled via one of the CAN interfaces using computers and CAN to Ethernet converters.

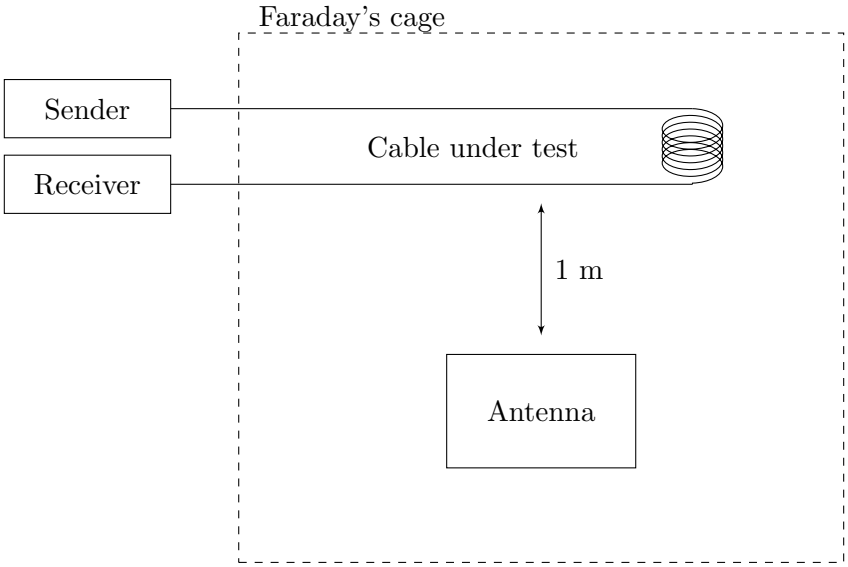

Figure 3.6: EMC emission test rig.

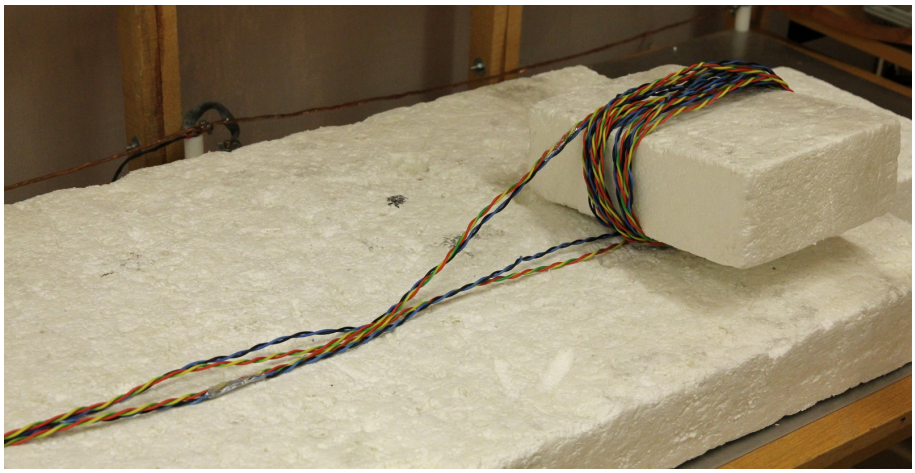

Figure 3.7: Cable placement for the EMC emission tests (UTP two pair).

Measurements where then carried out on each of the specified cable types between 30 MHz and 200 MHz using a Electro Metrics BIA-30S antenna and between 200 MHz and 1 GHz using a Electro Metrics EM-6950 antenna.

3.3.4 Radiated immunity

The main focus on the EMC test was to approximate the resilience of the Ethernet link in regard to electromagnetic disturbance. In Figure 3.8 the rig in which the test has been performed can be seen. Connections made for this test were very similar to the one used for radiated emission. The test focuses, as mentioned in Section 3.3 on application to application information loss. Therefore, only the probability for a dropped frame was approximated. With the simplification that there is no inter-symbol interference follows that the error probability for each bit is constant. This will in turn result in that the probability for a frame to include errors will increase with increasing frame size. The frame size used will due to this be kept constant during all the tests.

Two similar programs have been developed to simplify the EMC testing. One to do a sweep test where the drop rate is continuously updated, and one single point test where the number of packets can be specified and sent in one run.

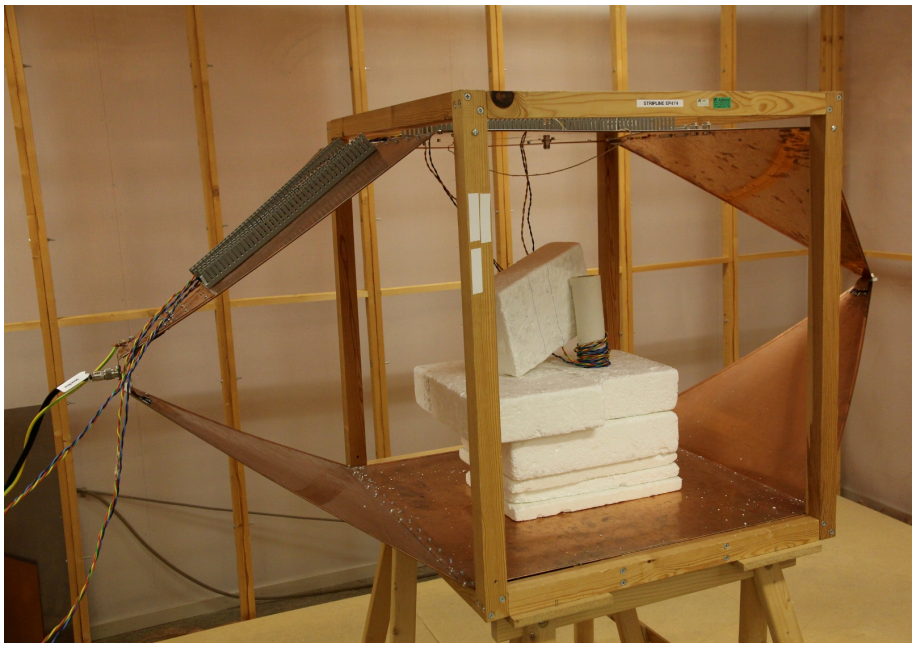

Figure 3.8: Measurement rig for the EMC immunity tests.

Note that there exist underlying mechanisms in the Ethernet standard IEEE 802.3 that define a link upstart procedure and ongoing channel monitoring. The mechanisms report link failure and if so, no data packages will be sent until a working link has been established. This enforces limitations on the immunity test in that a large ratio of frame drops can not be measured with the sweep test. Instead, a single point test would be able to make a good approximate of the link throughput even for high drop ratios. This due to that it analyses the link for a longer period of time which averages the effect of a temporarily dropped link.

Radiated immunity equipment limitations

The rig was defined to be able to generate electric fields up to 10 V/m for the amplitude modulated disturbance signal. For certain regions of the frequency spectrum the equipment was not able to deliver this field strength. The measurements carrier signal span the frequency range 25 MHz to 1 GHz. The amplitude modulated disturbance signal was configured to carry a 1000 Hz signal with 80 % modulation.

Radiated immunity sweep test

The radiated immunity sweep test is designed to quickly gauge the performance of a cable for a fixed disturbance amplitude over all frequencies of interest. This test is done

in an automated fashion, so that each disturbance frequency is tested for approximately two seconds.

The main concept for the software used in this test is to have a continuous gauge of dropped frames. This is achieved through one node sending numbered frames at a constant interval. The receiving node observes the sequence of numbered frames. If the received frame does not match the continuous sequence the frames in between are evaluated as dropped. The basis of this assumption is that the system used is not able to affect the order of the frames received.

The demand for the program to give a fast response is contradictory to being accurate. A compromise has to be made on the number of frames used to calculate the frame drop rate. A large number of observed frames will yield a slow change and that the test point is needed to be kept constant for a longer period. A small number of frames will result in faster change but the measurement loses precision and may not have time to stabilise. Every tested frequency is kept for approximately two seconds, the maximum time for getting valid data is therefore set to one second. Further constraints are set by the hardware which limits how fast frames can be sent to a two millisecond interval. These constraints result in that the last 500 frames were used to decide current frame drop.

The advantage of this test is the high test speed, which makes it a reasonable approach for a first evaluation of a collection of cable candidates. It was used to note at which frequencies the link was affected at maximal disturbance amplitude. The same program together with manual disturbance changes were then used to find the level of the disturbance for which errors started to appear on the link.

Radiated immunity single point test

This test is not as automated and fast as the sweep test. The main idea is to estimate a frame error rate with a higher precision motivated by statistical enquiries. This test is used to verify the results of the sweep test for combinations of noise amplitudes and frequencies of higher interest with certain cables. Further is the frame drop characteristics of concern and show the distribution of lost frames in time.

Here a large number of numbered frames are sent so that the receiver node can report the quantity and frame numbers of lost frames. The number of frames where set to 10 000, a number based on the reasoning in Section 2.6 [1]. One could argue for that the frame interval should be the same as for the sweep test to ensure that the estimated drop rate is unaffected by the test method. But the frame interval is not relevant as long as it is not too small since the bit error rate is estimated to be constant with respect to the usage level of the link.

[1] The use of the 97.5:th percentile (confidence level of 0.95) and the confidence interval ± 0.01 results in ≈ 9604 samples being needed

3.4 System modelling

To ensure that the designed Ethernet network are not causing problems in the time critical communication, Ethernet modelling software called TCN TimeAnalyzer is used. The software are to determine the performance of a specified system in terms of worst and best case latency and jitter for individual flows. The flows are build from one or several frames with specified inter frame time intervals. The software is managed from a graphical user interface where nodes and switches are added and linked together graphically. Frames can be defined and then used in one or several flows. A flow represent a periodic sequence of frames and have a defined source node and single node or a broadcast destination [15].

The user design a network and specify all flows in the network. So it is critical to model the network and add every flow correctly to enable the software to produce relevant results. If unknown flows are injected in the network the results obviously becomes invalid. Even in the case when all unknown flows are set to be low prioritised will they still pose a jitter problem due to that a started frame transmission always has to finish. The maximal introduced jitter from this type of scenario is therefore based on the maximal frame size that can be injected.

3.5 Frame size

When using Ethernet, the size of the frame can be substantially larger than when using CAN. On a single perfect link the largest possible size should be used to maximise throughput but there are other aspects such as packet drop and blocking that should be taken into consideration. A theoretical discussion is therefore made regarding the optimal frame size given some parameters.

3.6 Network topologies

The topology of the network form the foundation of the communication layer and has an impact on a number of different system properties. Different network topologies were evaluated with a focus on fault tolerance, wiring demands, and the need for additional hardware. This evaluation was purely theoretical since more a practical approach would have taken too much time. The three topologies that where evaluated are different combinations of the basic topologies presented in Section 2.2.3. There exist however no single solution that is best for every application. This due to that different systems have different properties and prioritise differently.

The three topologies discussed are selected from the properties of the system described below. The nodes in this system are distributed over two different areas, at the controls and the motors. To get a system with a feasible amount of cables, a solution with two

subsystems are used. The discussed network topologies can therefore be seen as two units connected by two Ethernet cables (for redundancy reasons). Since a bus layout only works for 10BASE-T [16] it is discarded and so is daisy chaining due to its high latency and lack of redundancy. That leaves tree/star, mesh and ring topologies to be discussed (see Section 2.2.3 for descriptions of the different topologies).

3.7 Hardware evaluation

A part of the work was to study possible choices of hardware for a full Ethernet installation. Measurements such as EMC have been made to form a basis for this evaluation.

3.7.1 Cable

Several properties of a cable are important in a production ready system. These properties include cost, weight, durability and resistance to electromagnetic disturbance. The EMC test described in Section 3.3 was used to see how the choice of cable affected the communication link's ability to withstand electromagnetic disturbance. The cables chosen for evaluation are the same cables chosen for the EMC test. The reasons for this choice of cable types can be found in Section 3.3.2.

Communication over a fibre cable would be very resilient to disturbance but will despite this not be further evaluated. It might however be a valid option and is here only discarded due to the thesis constraints.

3.7.2 Connector

When using Ethernet cables it is customary to also use RJ45 connectors. These do not withstand any harsher treatment and thus a different type of connector should be used for Ethernet in harsh environments. Apart from physical factors such as IP class and signal attenuation it is also crucial that it is readily available at a low cost. This means that it is desirable that it is produced by several different manufacturers with a long roadmap for the connector.

Discussion with CPAC System's employees lead to a selection of four different connector types to evaluate further. All these are classified for use in a marine environment and are in wide use today. This selection consists of:

a) NMEA 2000

b) IEC 61076-2-101

c) Deutch DT^{TM}

d) Molex MX150^{TM}

3.8 Android applications

Since this work is geared toward a proof-of-concept it is vital to show what possible advantages there are to using standard Ethernet. One such advantage could be the possibility to use any of the wide variety of products already built for Ethernet.

Android was chosen as development platform because of the ease in which a basic application can be developed. It was thought to have greater impact in a demonstration than a computer since it runs on a cheaper and more mobile device. Android is open source and run on a large set of inexpensive hardware platforms. This makes it a possible part of future prototypes as well as products. Especially if Ethernet is used since Android is built around the Linux kernel with a complete TCP/IP stack.

Two Android applications were developed to show the extendibility. One application intended for end user and another to show that Ethernet could simplify testing, debugging and maintenance. As of today the Ethernet infrastructure is highly available and could be leveraged to decrease the need to be physically present at the object of interest.

The end user do not need any low level information, but it can be interesting to show already present information on a different medium. It could also be interesting to show that a potentially wireless display could be constructed using off the shelf components.

The application that is intended for debugging and testing should on the other hand show all information. Since it is only meant to show the concept and possibilities such as wireless debugging, raw CAN messages are shown (with count and period).

3.9 Demonstration

The practical part of the thesis culminated in a demonstration on a boat (Nord West 1100). This was done to show that it is feasible to use Ethernet as a medium for time critical data. The time critical data stream was also mixed with the video stream from a TP-Link TL-SC3130G camera to show the possibility of using a single backbone system for data of different importance. In the demonstration the link between helm and motor was replaced with an Ethernet bridge as can be seen in Figure 3.9. In the figure CAN to Ethernet converters are marked CX, routers RX and switches SX. Each of the converters are connected to a CAN unit with a low speed CAN bus (125 kb/s). It was also tested to use the two routers to route all traffic between the two switches over a wireless bridge.

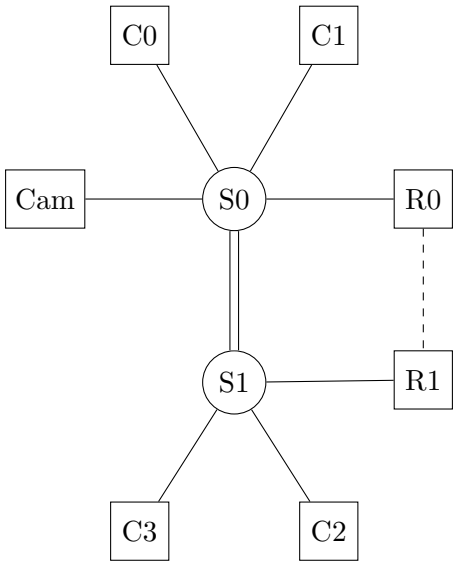

Figure 3.9: Layout of the system used in demonstration. (CX = CAN to Ethernet converter X, RX = Router X, SX = Switch X, Cam = Camera)

3.9.1 DC power supply

One problem that is outside the scope of the thesis work is the power supply compatibility of the hardware used in the demonstrator. The main issue being that the on board voltage vary greatly during normal usage.

The proof-of-concept hardware needs three different voltage levels 5, 12 and 19-48. CPAC Systems had DC-DC converters that deliver the needed voltage levels and are supposed to work under the described conditions. These converters were preferred to solutions such as using a separate battery due to convenience and availability.

3.9.2 Switches and routers

Industrial grade switches from Westermo have been used in the demonstration of the system. The choice to use Westermo switches for this thesis was motivated by that TCN TimeAnalyzer includes a model of these switches. They are furthermore robust and powerful switches that should not limit the functionality of the proof-of-concept.

Wireless routers are used for auxiliary functions such as to provide wireless CAN information and camera feed that can be displayed in a wireless unit. In addition to these functions a test where the Ethernet cables were substituted with two of these routers set up to work as a wireless bridge was also performed. This was done to give an idea of what is possible even if it currently is not a feasible solution. An approach very different from the selection of switches was taken regarding the selection of the routers. In this case it would not impact the basic functionality and they where instead chosen to show

possibility to use inexpensive Commercially available Off The Shelf (COTS) products. This resulted in the use of inexpensive Netgear routers.

Westermo switches

The used switches are made for industrial usage, where robustness is implemented in several ways. They are made to handle physical and electronic disturbances as well as a high communication load. The switches can be configured in multiple ways, such as via SNMP or a command line interface (e.g SSH or Telnet) but also through a more intuitive web interface. For the current system it is not needed to utilise anything more advanced than the web interface.

In the configuration used in the demonstration are two different CAN buses transmitted over a single 100BASE-TX connection. There exist different ways such as using different ports, subnets or VLANs to achieve a separation between the different streams. The most secure and easy solution, meaning that the same configurations of the converters could have been used, would have been to use 802.3Q VLAN. The Westermo switches support acting as both an edge and non-edge switch and could be connected to both the evaluation boards and routers. The routers that where meant to be able to work as a bridge between the switches dropped all incoming VLAN packages which meant that this approach could not be used. Subnets were instead used to logically separate the network streams.

3.9.3 Verification

To ensure that the priority mechanisms in the Westermo switches worked as expected, a verification test was performed. The test was based on a Volvo Penta steer-by-wire system mounted in a rig at CPAC systems. The CAN bus between the controls and motor was disconnected and each end of the bus was connected to a CAN to Ethernet converter. These were in turn linked to each other over Ethernet via a Westermo switch. In the same switch a third breakout board was connected and programmed to broadcast a fixed UDP frame as fast as possible. The frame was configured to have a different UDP port than that used to forward CAN frames by the converters. This results in that the frames where discarded when received by the converters.

A computer was also attached to the switch with the purpose to observe the communication via the software Wireshark. The switch was configured to prioritise the evaluation boards that were connected to the CAN steering bus, first using explicit VLAN tagging and second via the simpler port priority configuration.

4

Result

This chapter presents the results from measurements and theoretical evaluations. In the case of the EMC immunity measurements, the full set of results is shown in appendix A.

4.1 CAN to Ethernet converters

Since CPAC's current system uses CAN for communication and that the whole system should not be exchanged a converter between CAN and Ethernet was needed.

4.1.1 Prior work

Setting up a build chain and getting the converters [14] to work offered some minor difficulties due to that a newer version of the tool was used. Also the initial test showed that they only worked a small percentage of the time. But the converters seemed to work as intended after a minor bug fix.

4.1.2 Software robustness improvements

When running the converters as a bridge for a test rig they showed some unwanted behaviour, such as getting stuck in an infinite loop sending Ethernet frames when the CAN buffer was overrun. This and some other properties where adjusted with only minor modification to the code.

Routines to print statistics such as messages received and buffer overruns on the LCD display where added to simplify future debugging.

With these modifications the converters now lose their function when a data rate is too high since the buffers are overrun but continue their normal operation as soon as the data rate decreases. The same result is achieved if the network cable is unplugged during operation. If the converter is rebooted when the system is running it will restore the link as soon as it is fully configured.

It shall however be noted that if there is a lack in the power supply the card seems to be fully functioning but they will exhibit the exact same behaviour as if the network cable is unplugged.

4.2 Simple twisted pair cable

100BASE-TX does only use 2 out of 4 pairs in a standard Cat 5e network cable. Due to this a cable was constructed as described in Section 3.2.1.

4.2.1 Testing

The first test executed when only one of the two pairs were twisted was unsuccessful, since a stable connection could not be achieved and no packets could therefore be sent.

After the second pair was twisted as well a connection was successfully established and 10 000 UDP frames where sent with an interval of 5 ms. The test showed zero dropped frames and the cable was forwarded to the EMC test.

4.3 Electromagnetic Compatibility

An Electromagnetic Compatibility (EMC) test is normally made to test how a unit responds to electromagnetic disturbance and to check that emissions are within specified limits. In this case, an EMC test was used to aid in evaluation of cable types and to investigate how Ethernet behaves when used in an harsh electromagnetic environment. EMC consists of different types of disciplines which are based on their coupling, such as conducted or radiated and a combination of both [17].

This test focuses only on the actual cable performance which motivates that radiated tests was done. The results are still expected to be hardware dependent but the relative differences are expected to indicate the performance of the different cables described in Section 3.3.

The standard cable for a 100BASE-TX link is today the Cat 5e cable and hereby it is naturally the baseline for comparison with the other cable candidates.

4.3.1 Radiated emission

The radiated emissions were measured at two different frequency intervals, 20 to 200 MHz and 200 to 1000 MHz. The discontinuity of the test is due to that different antennas are used to test at these intervals.

The results can be grouped according to the attributes shielded/unshielded and low/high frequency interval.

During the radiated emission test on the parallel Cat 5e cable, it turned out not to work for the test hardware configuration (see Section 3.3.3). All cables had been verified to work as a link between two PC's and was again verified to work between two PC's when the problem occurred. The cable was excluded from the EMC test since the desired set up could not be used. This result implies that different network interface cards (NICs) have different performance and properties. Emphasising the fact that the performed measurements are limited and the relative differences should be in focus.

The difference between the cable candidate measurements and the background measurement is obviously relevant and a desirable aspect to illustrate. The digital background measurement for the low frequency radiated emission was unfortunately lost. A printed plot of the measurement was however obtained and a scanned version of it can be found in appendix B.

The quality and uncertainties of the measurements were examined by different types of measurements. In Figure 4.1 are the characteristics from subsequent repeatability are shown, the main sources of changes are expected to be the test equipment and environment noise.

For all measurements where a specific cable configuration of the device under test (DUT) is not presented, the configuration described in Section 3.3.3 is used.

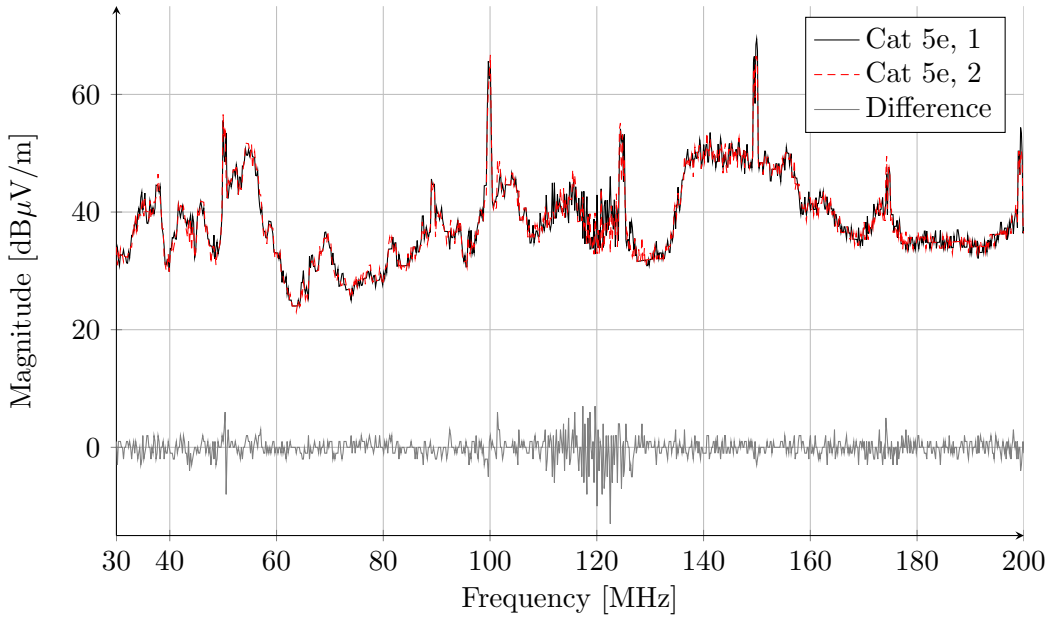

Figure 4.1: Plot showing emission measurements for two subsequent test runs of a Cat 5e cable.

To measure the effect of the low repeatability in the physical set up, minor changes reflecting these uncertainties of the physical layout was introduced. This change was then followed by a second measurement. The test were performed on the UTP two pair cable and the results are presented in Figure 4.2.

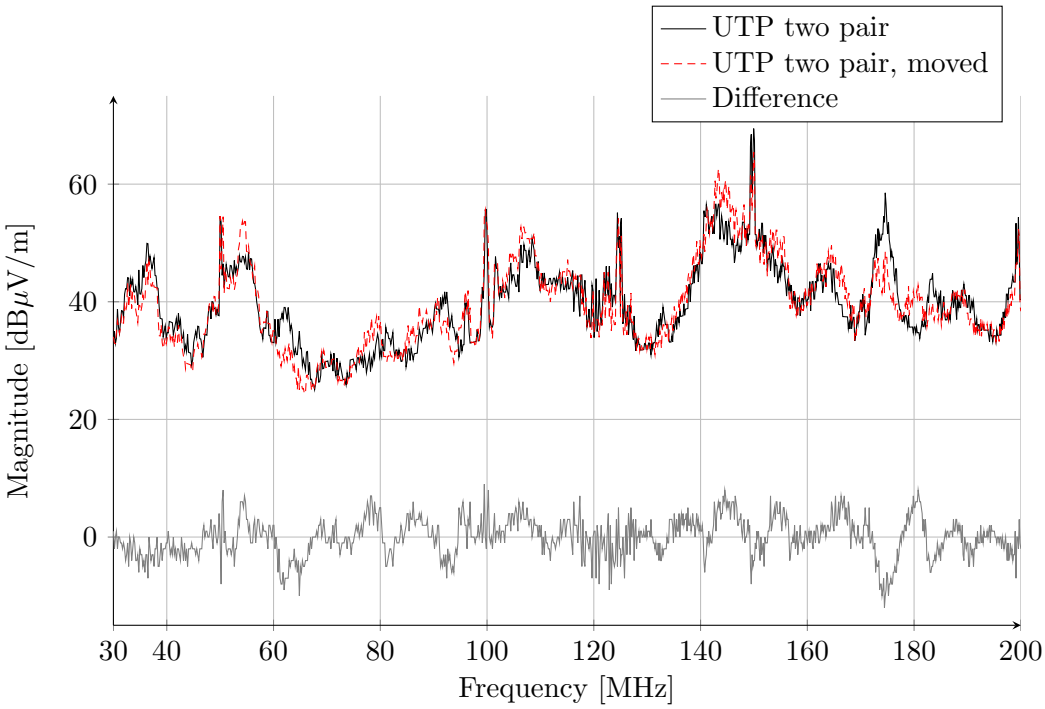

Figure 4.2: Plot showing emission measurements for two subsequent test runs of the UTP two pair cable which was slightly moved between the measurements.

Another investigated noise source was the AC adaptor connected to the computer which then supplied the development boards with 5 volts via a USB connection. Figure 4.3 shows measured results with and without AC adaptors connected.

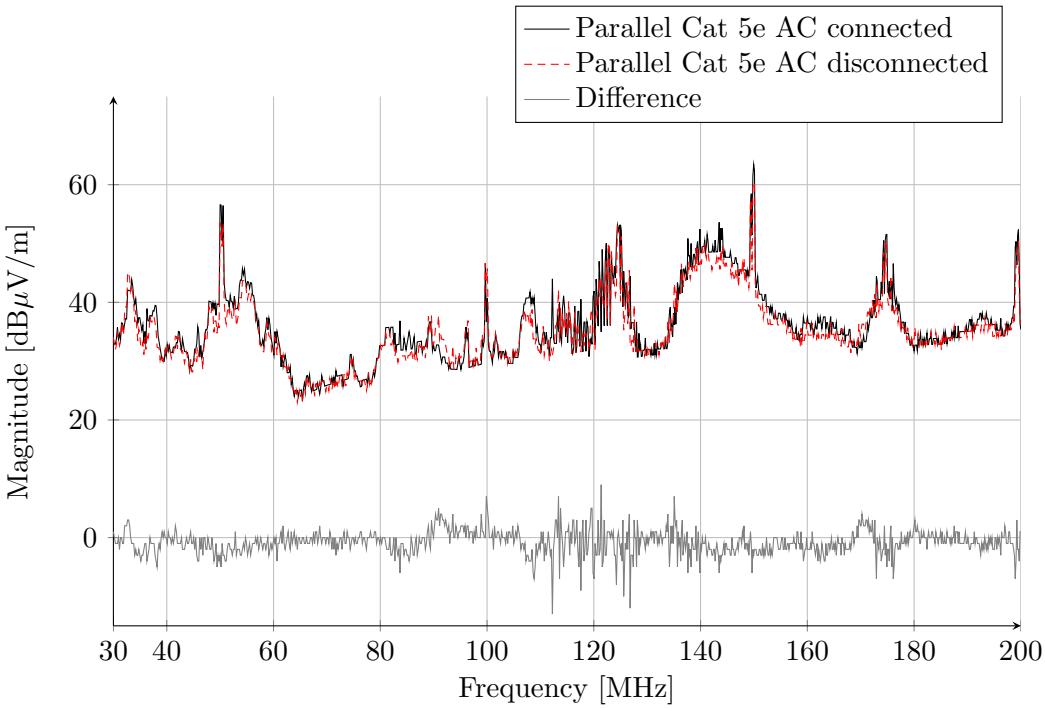

Figure 4.3: Plot showing emission measurements when the computers that supplied the experiment cards with power are plugged to an AC adapter and when supplied by battery.

The DUT configuration described in Section 3.3.3 was compared to a configuration where the cable was not looped, but only drawn straight in to the end of the Styrofoam where a 180 degree turn was made and then straight out from the measurement cage. The differences between these emissions for the UTP two pair cable are presented in Figure 4.4.

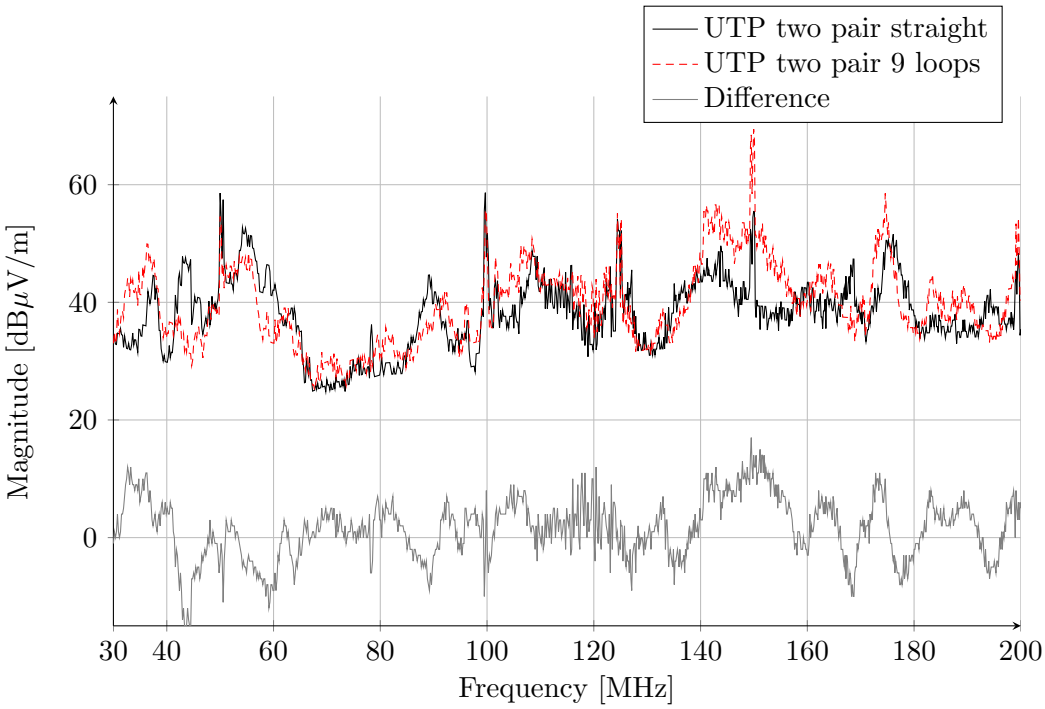

Figure 4.4: Plot showing the emissions of the UTP two pair cable configured as described in Section 3.3.3 compared to a different configuration.

Unshielded cables 30 to 200 MHz frequency emission

The emissions from the three cable candidates without shield are presented in Figure 4.5. Only larger characteristics and differences can be considered relevant due to uncertainties described in Section 3.3.3 and presented in Section 4.3.1. This motivates that all three measurements can be presented in a single figure.

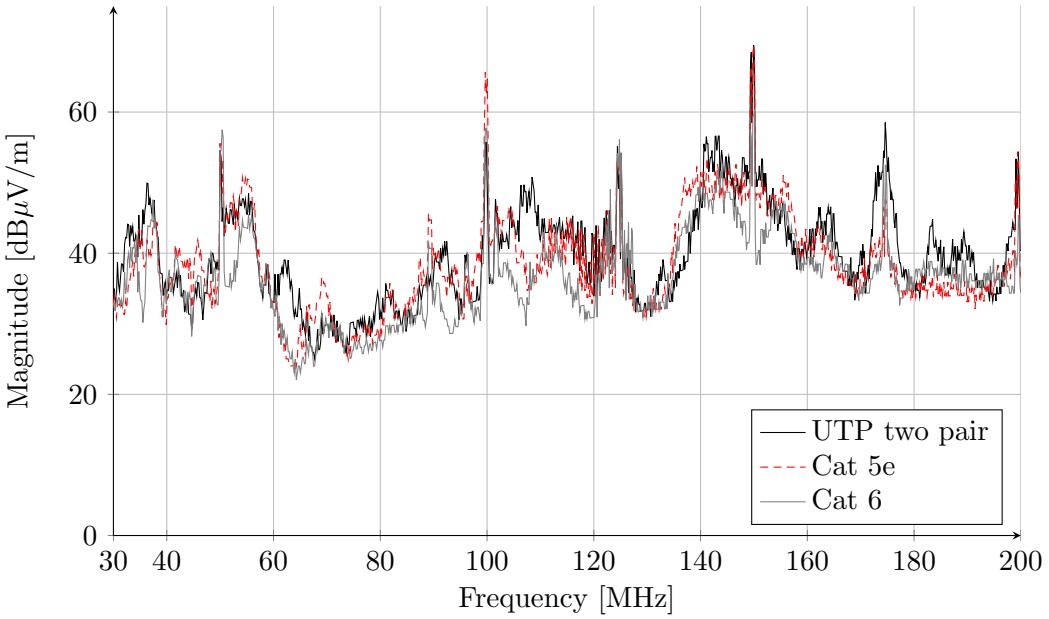

Figure 4.5: Plot showing emission measurements for the three different unshielded cables tested.

Shielded cables 30 to 200 MHz frequency emission

Figure 4.6 shows the results from the emission measurements from the two cable candidates with shield. Shields on measured cables were connected to ground only in one end except when a different configuration is explicitly stated.

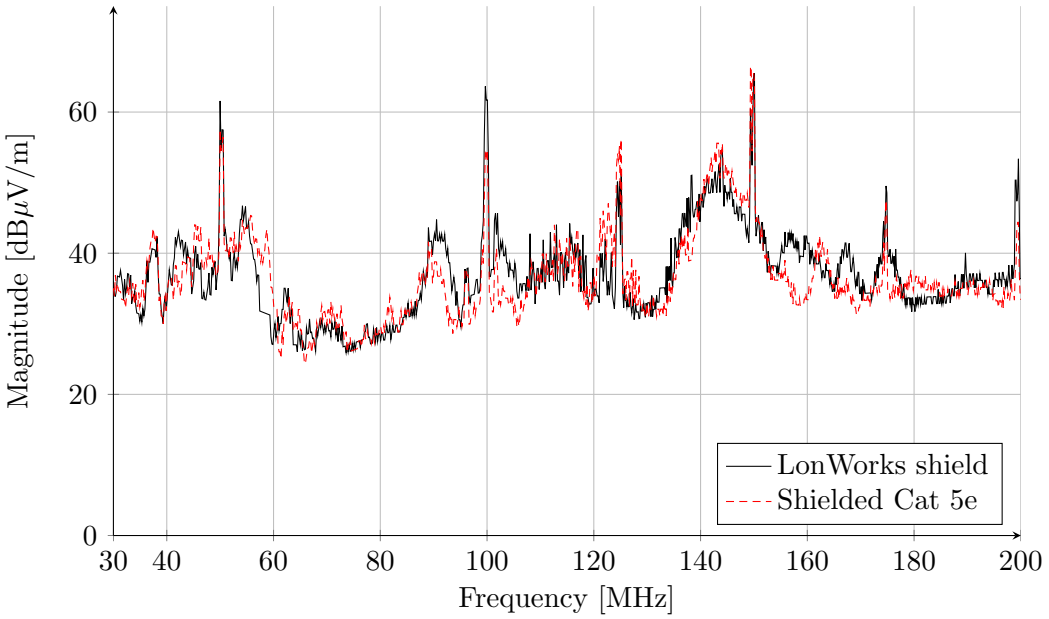

Figure 4.6: Plot showing emission measurements for the two different shielded cables tested.

The impact of how the shield is connected can be observed in Figure 4.7 and Figure 4.8.

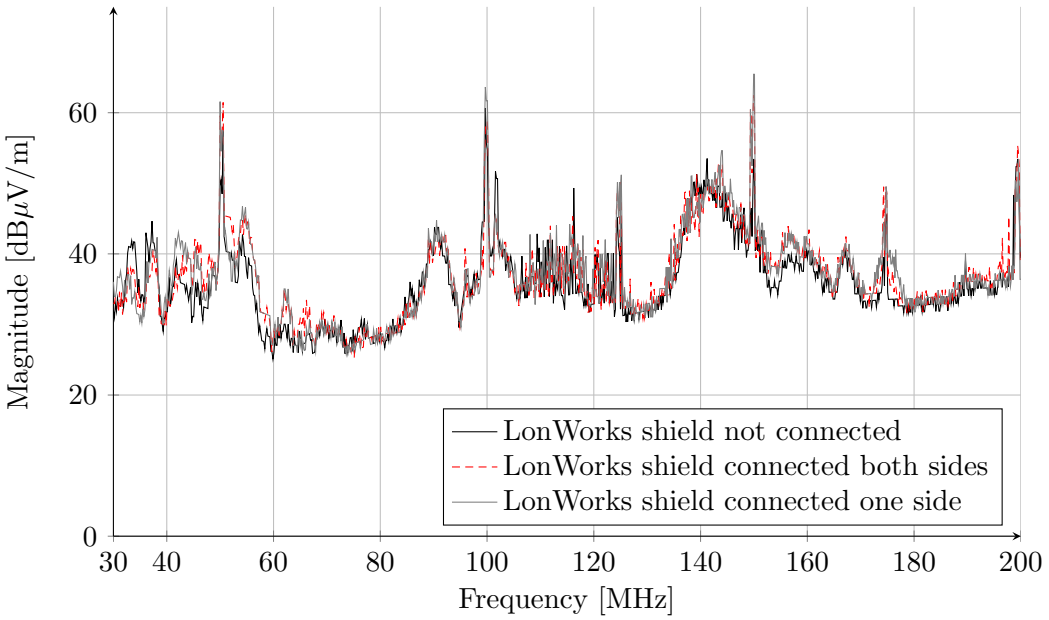

Figure 4.7: Plot showing emission measurements for the LonWorks cable with different shield configurations.

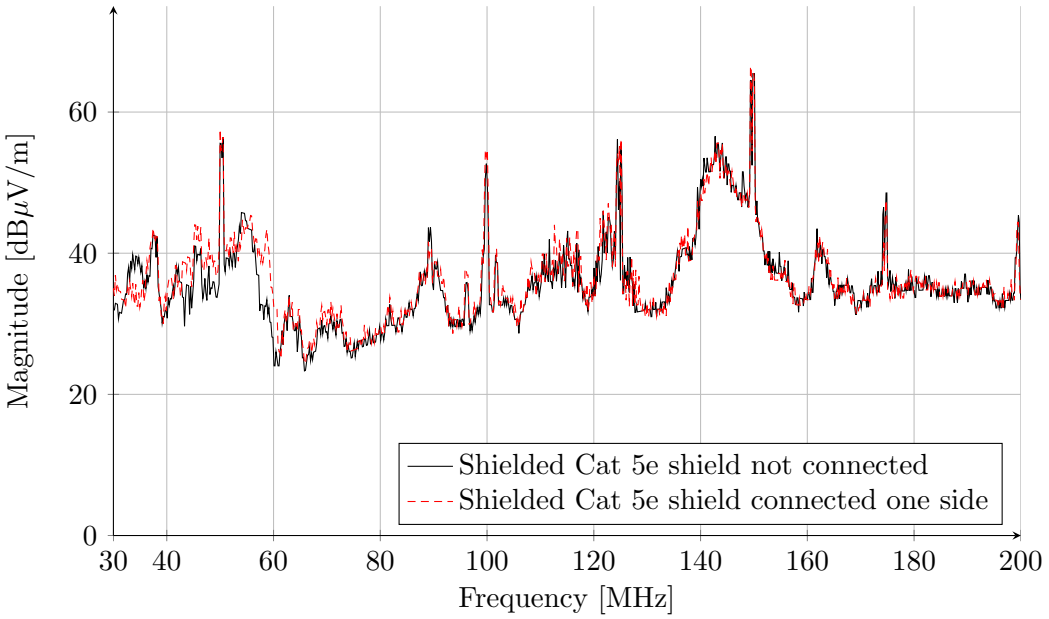

Figure 4.8: Plot showing emission measurements for the Cat 5e shielded cable with different shield configurations.

To compare the emission effect of a shield, a shielded cable it is compared to an unshielded cable and the result is illustrated in Figure 4.9.

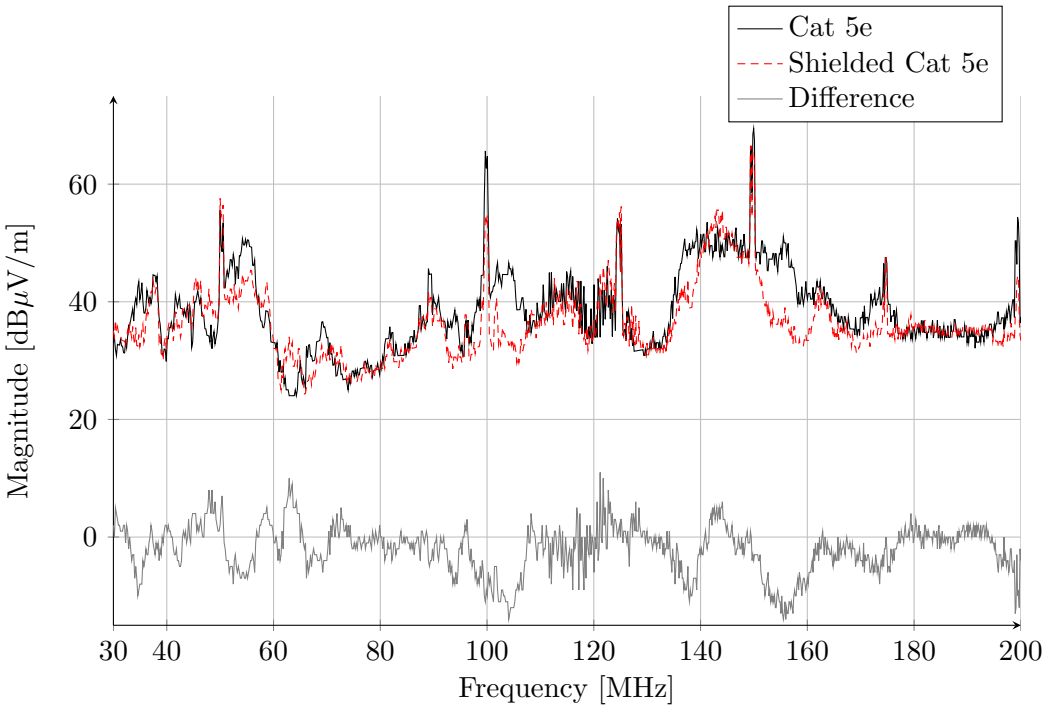

Figure 4.9: Plot showing emission measurements for the Cat 5e shielded and unshielded cable.

200 to 1000 MHz frequency emission

The higher frequency interval was not tested for all cables due to that this interval was not judged equally important and the lab time was limited. The chosen cables were the UTP two pair and the standard cable for 100BASE-TX, Cat 5e. Motivated by a rather poor performance in the low frequency test the UTP two pair cable was selected to be tested, with intention to investigate how the emission differs from the Cat 5e cable.

A repeatability test for emission where the physical set up remained unchanged were performed and the result is presented in Figure 4.10.

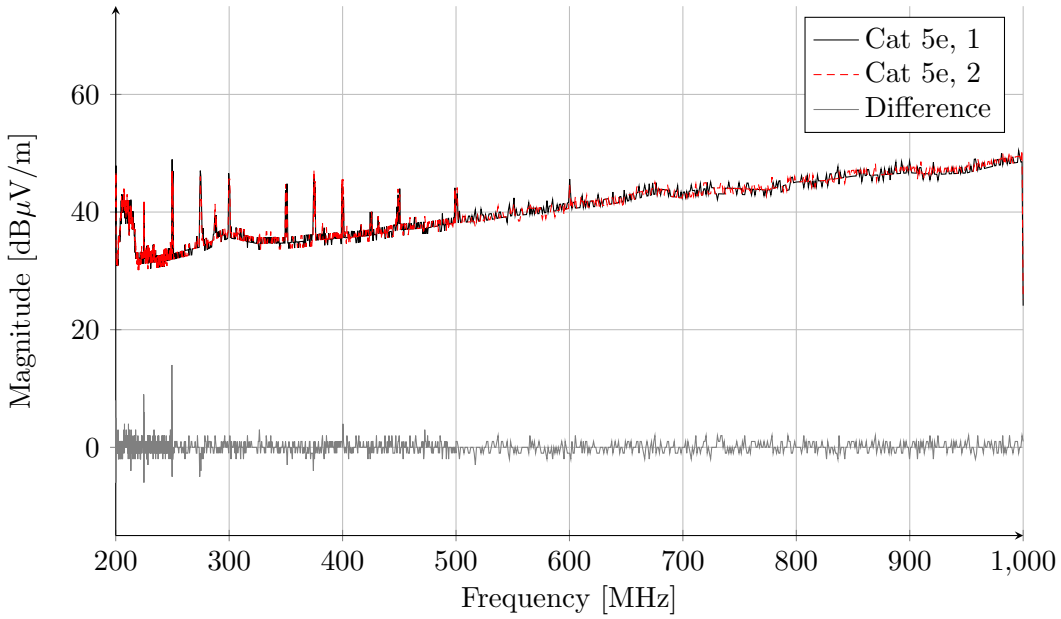

Figure 4.10: Plot showing the repeatability for emission measurements for the frequency interval 200 to 1000 MHz on the Cat 5e cable.

The measurements are affected by background and measurement noise. A background measurement was performed to to show the magnitude of the noise at the different frequencies and is presented together with the UTP two pair cable in Figure 4.11.

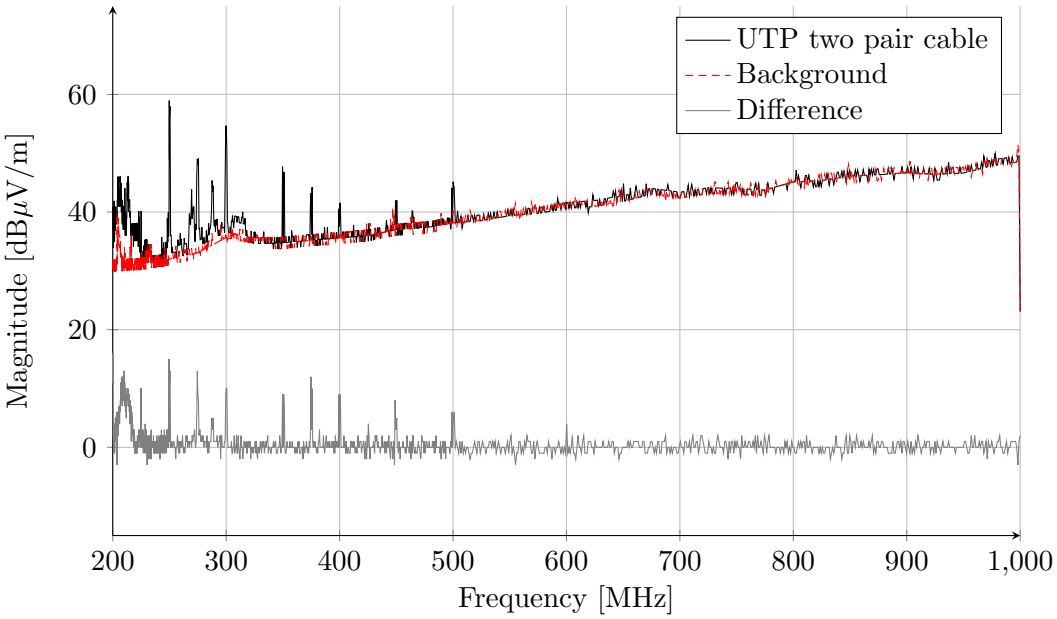

Figure 4.11: Plot showing the emission from the UTP double pair cable compared to the background noise.

The UTP two pair cable and the Cat 5e cable was tested in the high frequency interval emission measurements. The results together with the difference are presented in Figure 4.12.

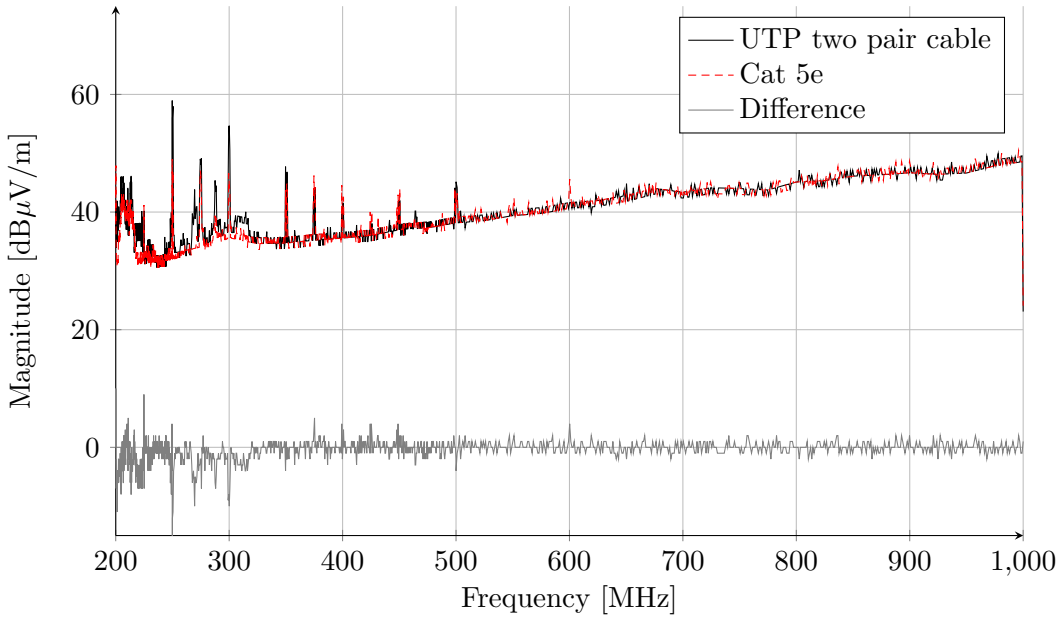

Figure 4.12: Plot showing the emission measurements comparing UTP two pair and Cat 5e cable.

4.3.2 Immunity

The electromagnetic radiated immunity was first and foremost tested for maintaining a lossless, fully functional link. Secondary tests were performed to evaluate the link degradation rate, i.e. how the link capacity and frame loss probability change with respect to disturbance amplitude on a certain frequency.

Full functionality tests

This immunity test was designed to find the maximum amount of electromagnetic disturbance that still allowed the link to maintain full functionality. The test consisted of a frequency sweep while the functionality was continuously monitored. When the link faltered the frequency was noted for closer investigation. The spectrum was then manually tested to find the point for when faults started to get introduced, with focus on earlier noted frequencies.

Due to limitations in the lab equipment the electric field strength of $10\,V/m$ could not be sustained for all frequencies. Figure 4.13 shows the maximum possible electric field strength of the lab equipment used.

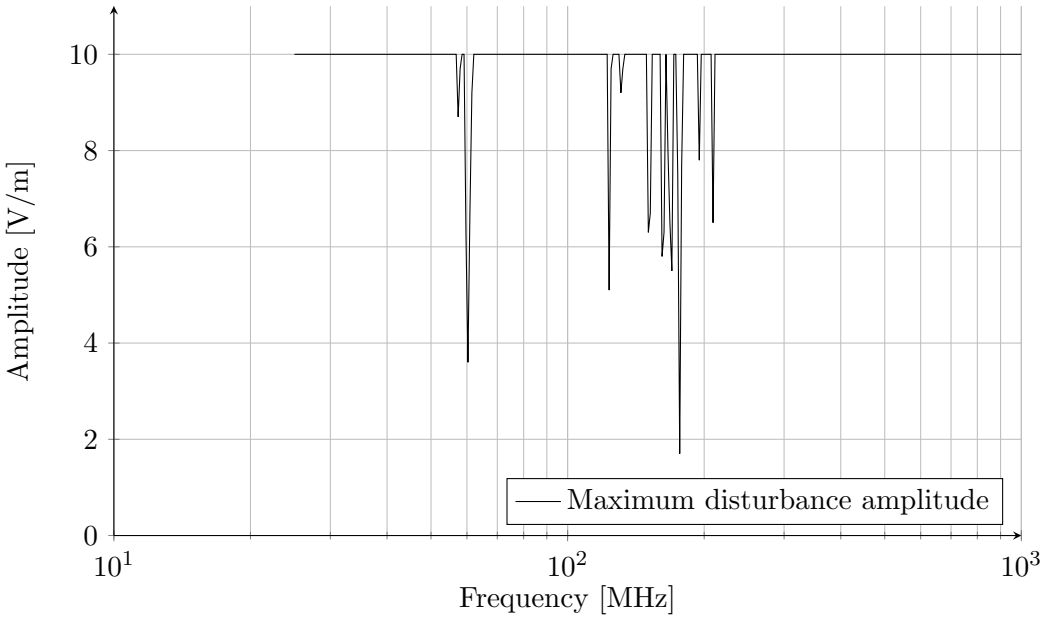

Figure 4.13: Plot showing the maximum amplitude of the disturbance signal of the radiated immunity equipment.

The frequency in Figure 4.13 refers to the disturbance signal's carrier frequency. The signal was amplitude modulated and carried a 1 kHz sinusoidal with a modulation index of 80%.

All cable candidates except the parallel Cat 5e cable were tested against immunity. The maximum disturbance amplitude were no frames where lost was logged. From this information no conclusions concerning the decay of the functionality can be estimated. Figure 4.14 indicates at which amplitudes and frequencies the link is not fully operational for the Cat 5e candidate. The complete results from the test can be found in appendix A.

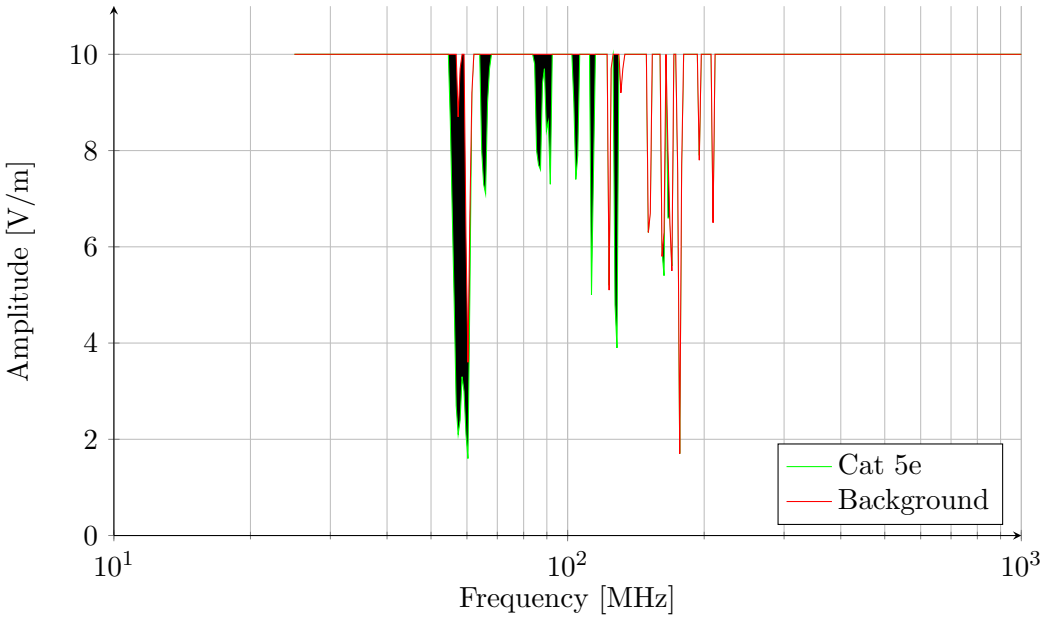

Figure 4.14: Plot showing link performance for Cat 5e cable where dark areas mark zones where full functionality is lost.

The immunity performance of the Cat 5e cable is shown together with the measured immunity of the UTP two pair cable in Figure 4.15.

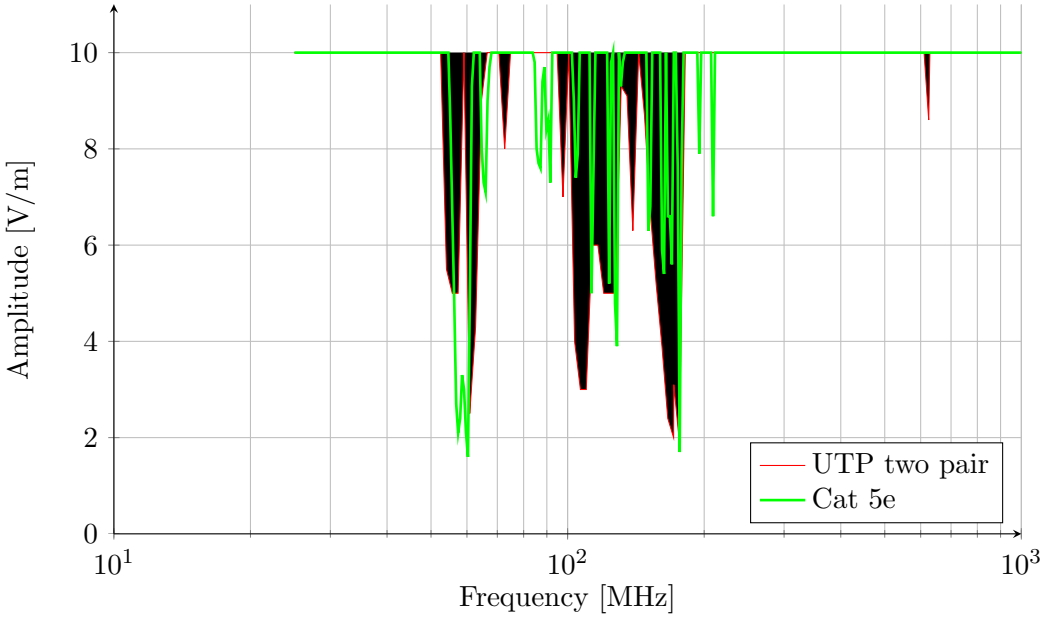

Figure 4.15: Plot showing link performance comparison for immunity between a Cat 5e cable and the UTP two pair cable.

The same test was made for a CAN link over the UTP two pair cable with 125 kbit/s link speed, between two of the development boards. For this case no functionality decrease could be detected. Two things should however be noted; first there exist acknowledgement mechanisms to resend broken frames for the CAN protocol (which was not observable with the test equipment) and secondly the link speed of 100BASE-TX is up to 800 times faster.

The immunity test results indicate that the UTP two pair cable performance is lower than the other candidates, but none of them manage to maintain full functionality throughout the test. No cable has distinctly higher performance than the others and the characteristics vary between candidates (see appendix A).

A behaviour that was observed during the tests was that the degradation seemed to have some hysteresis since the same point of measurement gave different results depending on the prior state of the link. The result for a single point was in other words better if the disturbance was increased from where the link was fully functional than if it was decreased from where it was heavily degraded. A possible explanation for this could be that the data is processed by an adaptive filter that makes an erroneous system identification when the link is downgraded.

Link degradation tests

To complement the full functionality measurements, link degradation was investigated. The same hardware for disturbance generation was used, but the test software was changed. For each disturbance amplitude and frequency, 10 000 Ethernet frames were sent over the link to the receiving computer which logged the number of dropped packages.

The s-shape graph in Figure 4.16 is a good example of the general link degradation characteristics. Even if the transition interval differs with respect to disturbance amplitude the general shape is reasonable and somewhat expected. Disturbances above a certain level will result in that the link will be diagnosed as faulty, and it has to be re-established as described in Section 3.3.4. This is thought to be the reason to the sudden increase in dropped frames.

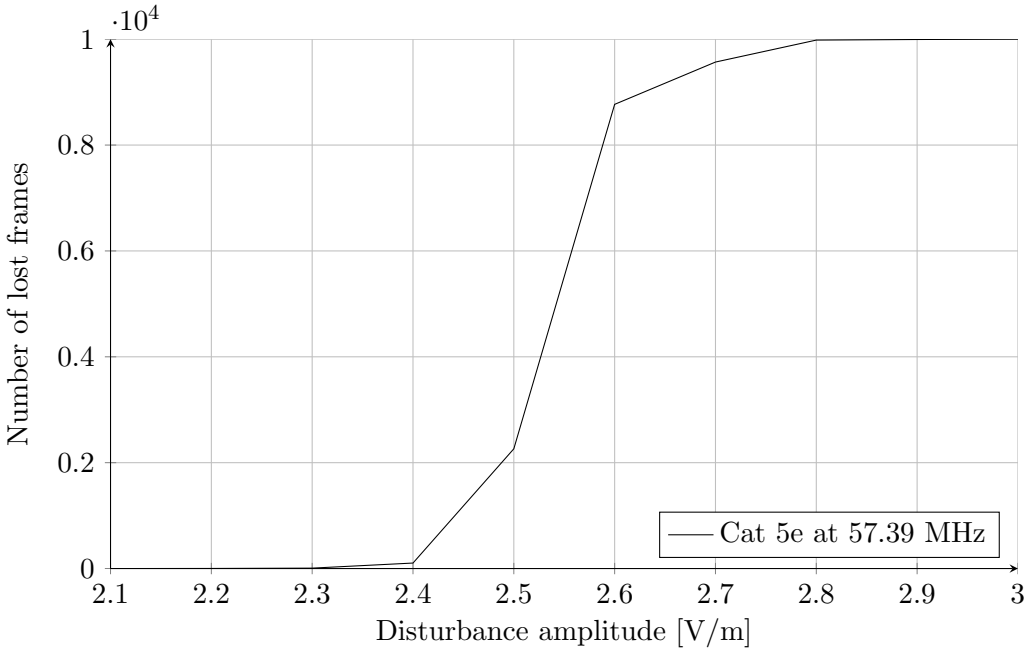

Figure 4.16: Plot showing measured link degradation for the Cat 5e cable at a fixed frequency.

The result obtained when the link functionality is measured at 91.60 MHz is illustrated in Figure 4.17. It shows that a large uncertainty is connected to each of the measured points, but still exhibit the same basic shape as Figure 4.16.

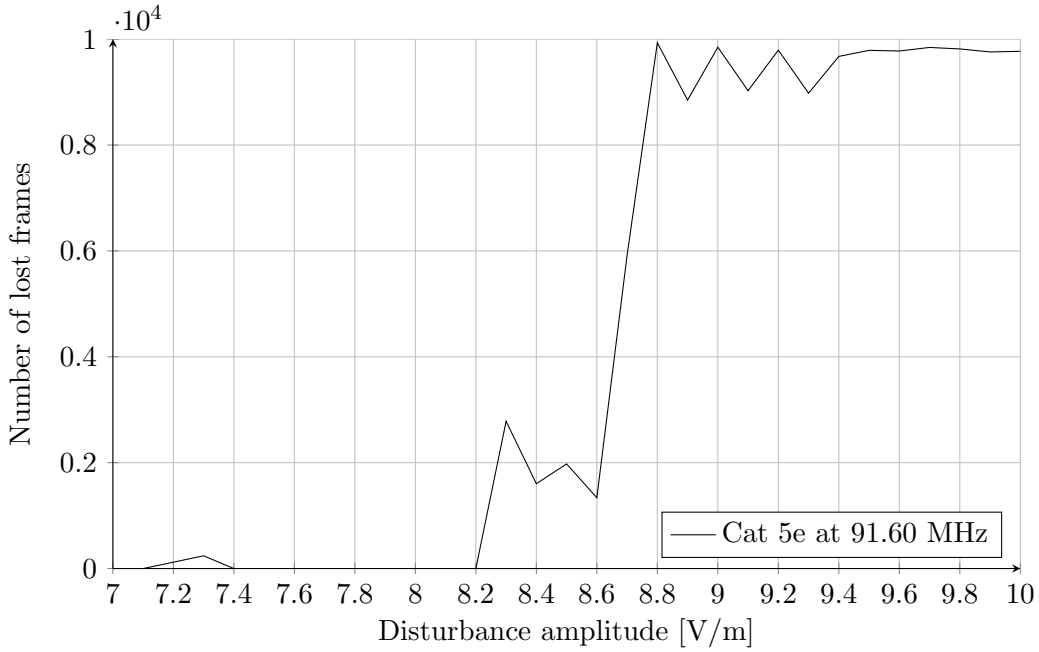

Figure 4.17: Plot showing measured link degradation for the Cat 5e cable at a fixed frequency.

Figure 4.18 is an example of the uncertainties and non repeatable measurement results. The peak of high frame losses between 4.8 and 5.8 V/m, is reduced at a higher disturbance amplitudes. Later measurements in this region give a much lower frame loss. This behaviour shows that the measurement environment was not constant and the results could for example be explained by a voltage increase on the evaluation board's power input and thus increased power to the transceivers. This type of increase could be an effect of a computer switching from charging, to having a fully charged battery.

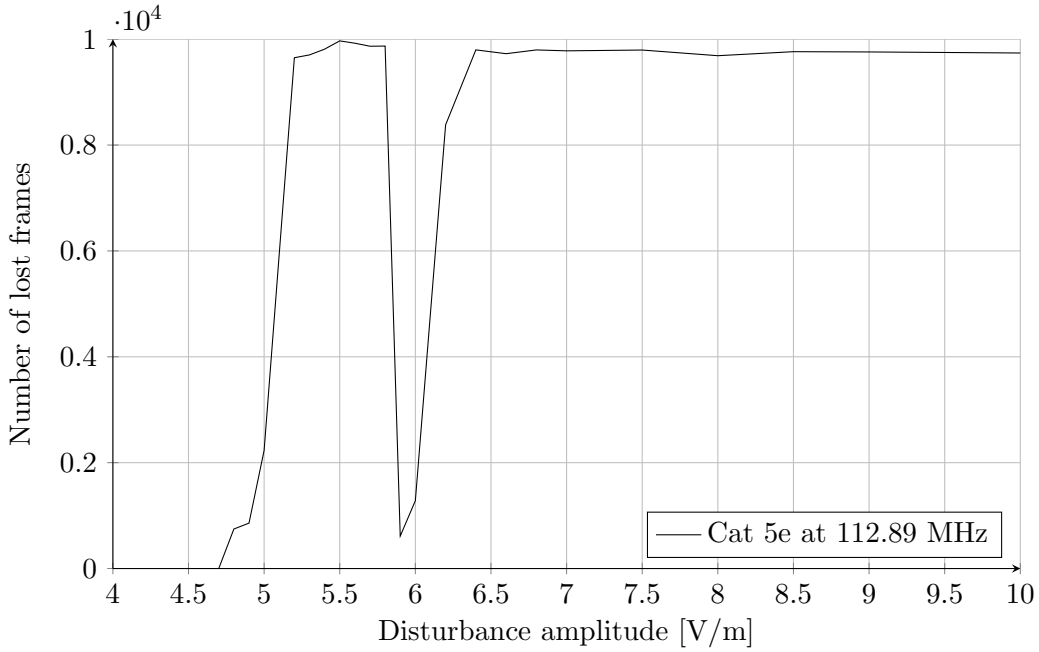

Figure 4.18: Plot showing measured link degradation for the Cat 5e cable at a fixed frequency.

4.4 System modelling

The demonstration network shown in Figure 3.9 was modelled using the TCN TimeAnalyzer software. Observe that the solution with subnet broadcast described in Section 3.9.2 causes unnecessary traffic since all frames are broadcasted to all nodes. Nodes in a different subnet will then discard the frame upon reception.

The model is based on specifications of the CAN bus communication. These specifications where supplied by CPAC Systems, the developers of this specific communication in the steer-by-wire system. An Ethernet communication system's best case latency is linked to the number of "hops" a frame has to do before it reaches its destination (transmission and switch forwarding times). Also the frame length will affect the resulting best case forward time (through transmission time), which for all time critical frames in this system is 12 or 18 μs. The maximal jitter for each flow is approximately 300 to 540 μs. A distinctly higher level of maximal jitter is obtained on the hardware port where the camera stream is forwarded. The amount of extra jitter on this port relates well to the expected blocking time of a frame with corresponding size and the fact that the camera stream has a lower priority than the time critical steer-by-wire flows.

The thesis work done by I.Muhammad, pp.98-101 [14] estimates the total conversion time of conversions from CAN to Ethernet and back to CAN to 35.816 μs. He further explains that the obtained results from TCN TimeAnalyzer should not only be added to

this conversation time but also an Ethernet serialisation time and an CAN serialisation time.

During this thesis the TimeAnalyzer software was in a beta stage and a dialogue with the developers was crucial to enable problems that arose to be solved. The results from the modelling was planned to be verified through Ethernet measurement with hardware especially designed for Ethernet timing measurements on the actual system as well as on a fake system to maximise the probability of that the worst case has been measured. The modelling has due to time spent waiting for a new beta release and bugs fixes not been able to be kept within the planed schedule. Therefore these verifications have not been possible to do within the thesis time constraints.

4.5 Frame size

In comparison to the CAN frame, as presented in Section 2.3.2 Ethernet, IP and UDP can hold a large amount of data (up to 1472 bytes). To decrease the overhead as much as possible, the maximum data length should be used. However there are reasons to claim that data should be packed into smaller frames.

4.5.1 System design

The usage of small frames could simplify system design in that unrelated data would not be packaged in a single frame. This would result in some decoupling between processes. This section focuses however on the throughput of the system and this aspect will therefore not be further discussed.

4.5.2 Real time aspect

When time critical frames are to be sent, frame tagging described by 802.3Q could implement to enable prioritisation of Ethernet frames. If only smaller frames are used in the system the upper bound for forwarding time of a high priority frame is lowered. If a node has data with different intervals a shorter high frequency frame can be designed, and even though the relative overhead is larger, the total bus load is lowered since less unnecessary bits are transmitted.

4.5.3 Frame drop

In the presence of disturbance the optimal frame size to maximise data throughput is not as easy to determine as in the ideal case. Consider the case of a constant bit error probability then as the frame size increases the probability of a frame to be successfully transmitted decreases as described by equation 4.1, where P_{bit} is the probability that a

bit is successfully transmitted and L_{header} and L_{data} are the frame header length and frame data length in bytes respectively. The header length refers to the sum of the Ethernet, IP and UDP non data bytes see Figure 2.15 for details.

$$P_{frame} = P_{bit}^{8(L_{header}+L_{data})} \tag{4.1}$$

It can be shown mathematically that for a value of P_{bit} less then one the probability for a successful frame transmission decreases when the frame size increases. This can also be seen in Figure 4.19.

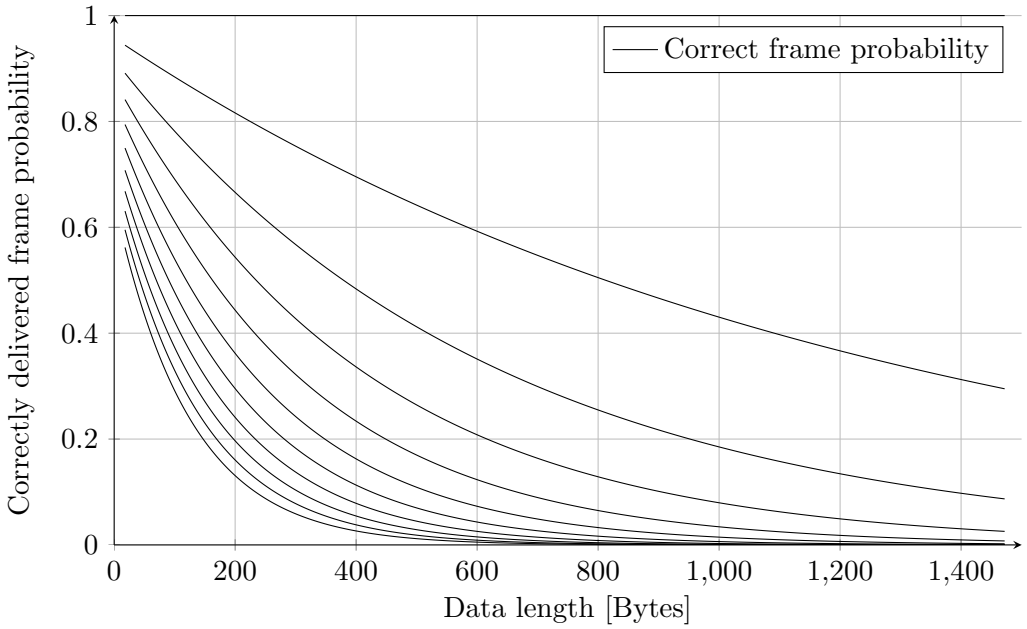

Figure 4.19: Plot showing how the probability for a successfully sent frame decreases with the data length in bytes and the decreases of P_{bit} from 1 to 0.999 with a step size of 0.0001.

Further the data overhead can be viewed with respect to utilisation. Utilisation is calculated as the number of useful bytes over the total number of bytes, as described by equation 4.2.

$$U = \frac{L_{data}}{L_{header} + L_{data}} \tag{4.2}$$

Both Figure 4.20 and mathematical analysis indicate that the utilisation factor approaches one as the frame data length approach infinity, note the UDP data length limit of 1472 bytes.

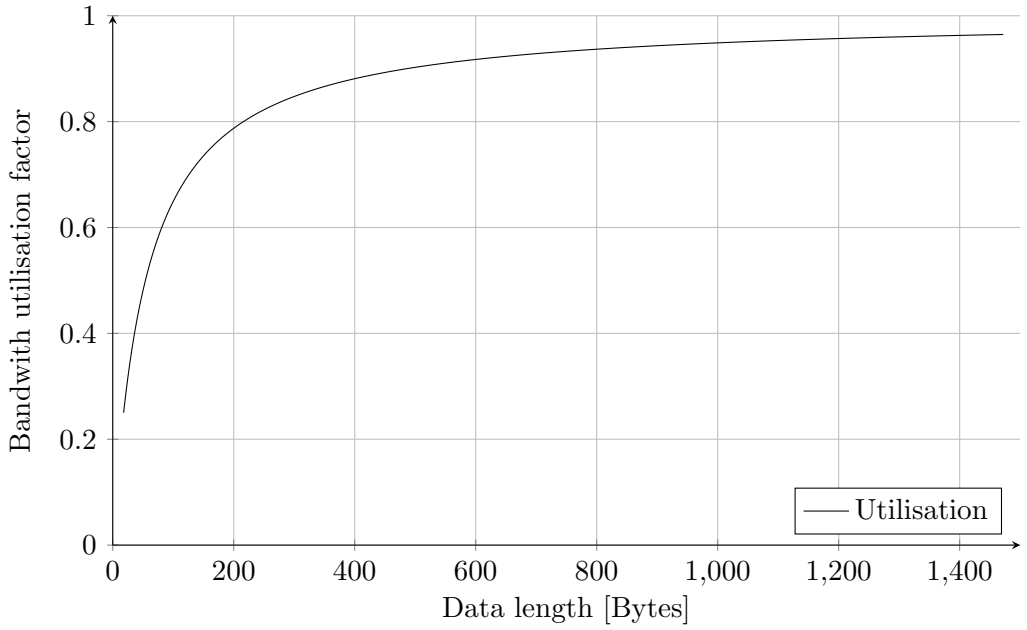

Figure 4.20: Plot showing how the bit probability affect the probability of frame loss for UDP packets sent over Ethernet.

Given a value of P_{bit} and a data length in a frame the maximum expected data throughput can be calculated using equation 4.3. Figure 4.21 illustrates this for selected values of P_{bit} and a constant value of one for the Ethernet link bandwidth B.

$$\begin{aligned} \mu_{troughput} &= B \cdot U \cdot P_{frame} \\ &= B \cdot \frac{L_{data}}{L_{header} + L_{data}} \cdot P_{bit}^{8(L_{header}+L_{data})} \end{aligned} \tag{4.3}$$

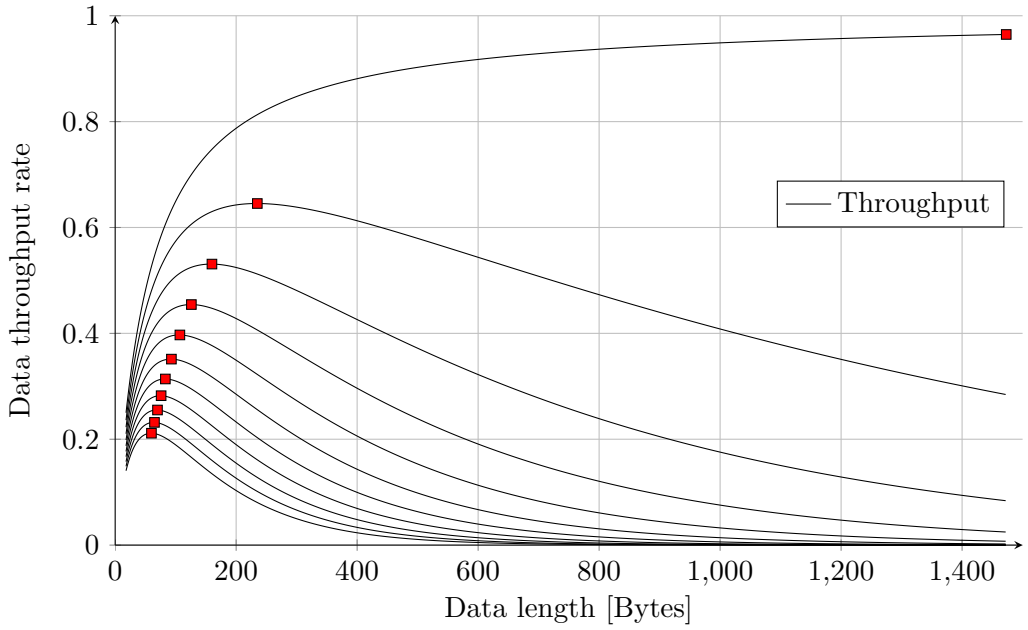

Figure 4.21: Plot showing the expected data throughput rate for a link bandwidth of one, given frame size and for P_{bit} spanning from 1 to 0.999 with a step size of 0.0001.

4.6 Network topologies

Three topologies that where relevant to the current system were chosen for evaluation in regard to hardware demand and fault tolerance. The topologies in this section have been drawn with a circle for a switch and a rectangle for a node/ECU.

A property that all of these network topologies have in common is that they consist of two groups connected by two cables. This type of connection where more cables than necessary are used to increase the fault tolerance is made possible by protocols such as RSTP and FRNT. If one link is broken the other will be used after a period of reconfiguration (see Section 2.4.4). In the event that both of these cables are broken, all discussed topologies will fail in that the two groups of nodes are completely separated. In a boat this would result in total loss of communication with the motors.

4.6.1 Tree topology

The tree topology consists of a number of star networks, in this case two (Section 2.2.3). This ensures that all nodes are directly connected to a switch and the number of hops is therefore minimised (Figure 4.22). In this specific network, every node can reach every other node in a maximum of three hops. A package between nodes N0 and N5 would for example take the route N0-S0-S1-N5. If a link between a node and a switch is broken,

the node is completely disconnected from the entire network. But no other nodes are affected by this lost link. One drawback to this topology is that if the switch is placed far away from the nodes connected to it, it could lead to a large amount of cabling due to that a cable is drawn from the switch to each of the nodes.

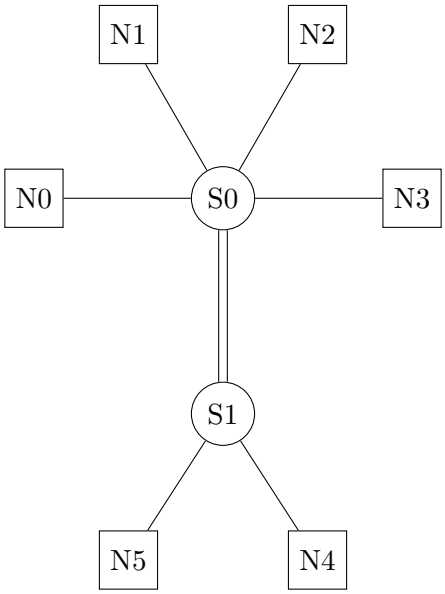

Figure 4.22: Suggestion of a tree topology.

The routing of packages is centralised to a few dedicated switches and the ECUs can therefore be constructed using a single network interface and a standard stack. The drawback is that all nodes has a single point of failure in that single cable between the node and the closest switch.

4.6.2 Ring topology

A ring topology is as the name suggests a ring of nodes where each node is connected to the nodes on either side. This will for some physical configurations result in that a smaller amount of cables is needed when compared to a star network. Another positive aspect is that the whole network stays connected when any single cable fail. Packages can be sent in either direction in such a constellation and a protocol such as RSTP or FRNT (Section 2.4.4) is needed to successfully route the packages. The multi route possibility creates a fault tolerance in the system but it also means that latency can be affected when routing changes. If for example the link between node N0 and switch S0 plus the link between node N5 and switch S1 are unused (see Figure 4.23). The transfer of a packet between the nodes N0 and N5 would consist of seven hops (N0-N1-N2-N3-S0-S1-N4-N5).

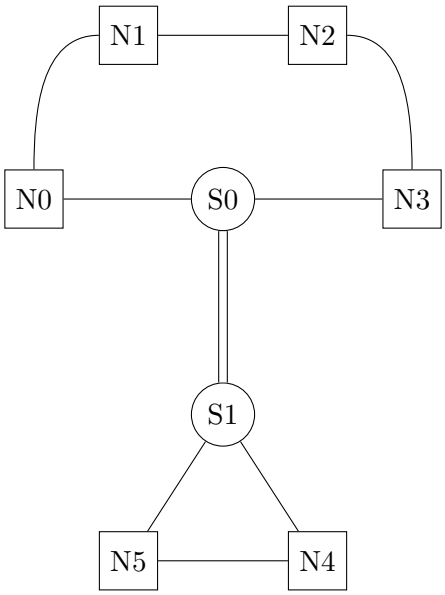

Figure 4.23: Suggestion of a ring topology.

The ring topology removes the need for switches with a large number of ports. Each of the nodes must instead be able to route packages coming from other nodes according to the current routing paths. So the routing becomes decentralised which means that each node becomes more complex but the fault tolerance is increased.

4.6.3 Mesh topology

A mesh topology is the least strict topology in that it consists of an arbitrary number of connections between nodes i.e. each node is connected to one or more other network units (see Section 2.2.3). This allows for a layout where some nodes are kept simple (such as N0 in Figure 4.24) and other more crucial nodes can have multiple connections to increase fault tolerance (such as N3 in Figure 4.24). Increased complexity does however mean that the system is harder to analyse, and need more complex routing algorithms. The maximum number of hops needed for a package transfer in Figure 4.24 can be seen when the link between N2 and S0 as well as the link between N5 and S1 are broken, and a package needs to be sent between N2 and N5. The route N2-N3-S0-S1-N4-N5 would be taken and five hops would therefore be needed

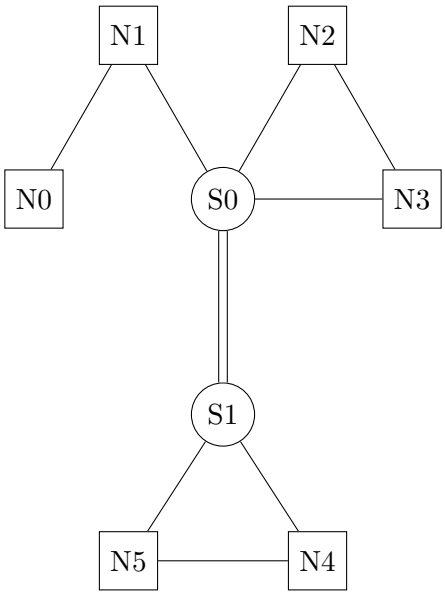

Figure 4.24: Suggestion of a mesh topology.

How a mesh network should be constructed depends not only on the nodes' placement but also between which nodes communication exists. If N0 only communicates with N1 the load on the rest of the network can be lightened if they have a direct link. The same reasoning should be applied to subgroups if larger and more complex networks are constructed. The mesh network can benefit from the advantages as well as suffer from the drawbacks of both ring and star topologies depending on how it is constructed.

4.7 Cable evaluation

A lot of different factors are important when choosing a cable for a critical communication system. These factors include attenuation, weight and return loss as well as non technical factors such as price and availability.

4.7.1 Category 5e and 6

The 5e and 6 category cables (often called just Cat 5e or Cat 6) are the most commonly used cables for 100BASE-TX Ethernet. These do however contain four twisted pairs whilst 100BASE-TX only makes use of two. They are widely available at relatively low cost since they are widely used in computer networks. The standardisation of these cables also entails that their performance is well specified. Some of the properties of Cat 5e and Cat 6 cables measured according to the TIA-EIA-568-A standard can be seen summarised in table 4.1.

	Cat 5e	Cat 6
Frequency Range	1-100MHz	1-250MHz
Attenuation	24 dB	21.3 dB at 100 MHz
		36 dB at 250 MHz
Near End Crosstalk	30.1 dB	39.9 dB at 100 MHz
		33.1 dB at 250 MHz
Return Loss	10 dB	18.6 dB at 100 MHz
		8 dB at 250 MHz

Table 4.1: Summary of Cat 5e and Cat 6 Ethernet cable properties [18].

4.7.2 LonWorks

LonWorks is a network standard for communication in control applications such as heating and home automation. Several different media can be used, one of which is a twisted pair copper cable. The cable needed share properties of those needed for 100BASE-TX (e.g. number of pairs). This is however a low speed network operating at 78 kbits/s [19].

4.7.3 UTP two pair

Section 4.2 shows that 100BASE-TX communication over a simple unshielded twisted pair cable is possible. But Section 4.3 raises the question of how much the transmission is affected. Only link availability and packet drop were measured which means that basic properties such as rise and fall times were not checked against specifications.

This shows that a simpler cable than Cat 5e might be a viable option but further studies need to be made to ensure a reliable link.

4.8 Connector evaluation

Since RJ45 connectors are not of a relevant IP class a different connector must be found if Ethernet is to be used in a marine environment. The following five connectors were as described in Section 3.7.2 chosen for evaluation.

4.8.1 NMEA 2000

National Marine Electronics Association (NMEA) has developed a standard for communication over a CAN bus at 250 kbit/s in a marine environment [20]. The connector for

this system is a five pin screw on variant [20] (see Figure 4.25). The testing for certification include areas such as impedance, propagation, attenuation, shielding, water/oil resistance and fire protection [21].

Figure 4.25: NMEA 2000 connector [22].

Five pins is sufficient since 100BASE-TX need 4 pins for communication and the fifth pin is used for the shield. Since the standard was developed by a association are there certified connectors available from several different producers which eliminates the dependency on a single company. Adoption by companies such as Garmin also suggests that production will continue to at least provide spare parts for a foreseeable future [23].

4.8.2 IEC 61076-2-101

The usage of a four pin M12 contact has been defined as an Industrial Ethernet standard according to IEC 61076-2-101 [24]. Male connectors have an encoding with four different pin radii to prevent mating with a different interface. The IEC 61076-2-101 specifies free and fixed connectors in both rewireable and non-rewireable variants (see Figure 4.26) [25]. The connectors are classified as waterproof since they are available with IP68 certification [24].

Figure 4.26: IEC 61076-2-101 connector [26].

The specifications of connectors are available in an open standard and the connectors are therefore produced by several manufacturers. A solution with data and power in the same cable is possible since there exist connector configurations with up to 17 pins [25]. Widespread adoption on the industrial side also suggests continued availability.

4.8.3 Deutsch DT™Series

The DT series from Deutsch are environmentally-sealed connectors specifically developed for communication in motor and transmission applications. Latches keep the connectors from disconnecting once mated (see Figure 4.27) and silicone seals are used to reach IP67 certification for AWG 18 through 14 wire [27].

Figure 4.27: Deutsch DT connector.

The Deutsch DT series is designed and produced by Deutsch and there is because of that no alternative supplier. Connectors with a pin count of two, three, four, six, eight and twelve is available [27].

4.8.4 Molex MX150™

Molex MX150 is a sealed connector system designed and produced by Molex (Figure 4.28). It is designed to be able to be used in marine environment with a IP6K9K classification. There exist both fixed and free connectors where the free are suitable for cables between 0.35 and 0.50 mm^2 [28].

Figure 4.28: Molex MX150 connector [29].

MX150 is just as Deutsch DT only available from a single supplier which increases the dependency on a single company. The connector is however available in 15 different pin configurations with a pin count between 2 and 20 [28].

4.9 Android applications

Two Android application were created to show the versatility and ease of which an Ethernet solution can be extended.

4.9.1 End user application

End users do not need any additional information from the system so rounds per minute, fuel level, trim angle and motor oil temperature were displayed. The result can be seen in Figure 4.29.

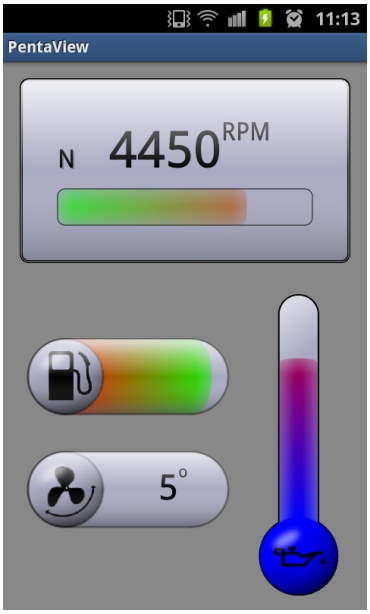

Figure 4.29: Screenshot from a Samsung Galaxy S2 running an application showing motor data.

4.9.2 Debugging application

In the application for debugging all available information is shown since any existing information could be relevant. A simple, yet useful CAN logging application was created where messages are shown in their raw form. But additional information such as number of frames with the same identifier and the time between them are also shown. For details of the interface, see Figure 4.30.

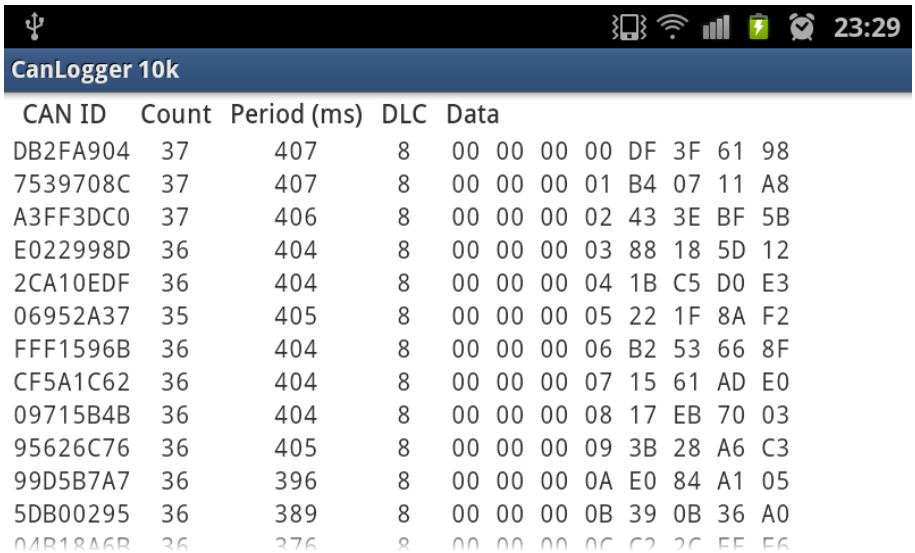

Figure 4.30: Screenshot from a Samsung Galaxy S2 running a CAN logger application.

4.10 Demonstration

In the end of the thesis, a demonstration on the boat that can be seen in Figure 4.31 (Nord West 1100) was planned and performed.

Before tests were run on the boat a verification of the demonstration hardware was carried out as described in Section 3.9.3. The rig was started and test functioned without exhibiting any abnormal behaviour in both test cases. The massive Ethernet load that was captured in Wireshark caused the otherwise well behaving laptop to freeze within seconds after the test was started.

This result strongly indicates that the Westermo switches can handle priority classification in combination with a large traffic load.

On the actual boat there were initially problems with a faulty extension cable. But the wired system worked well while connected between one driving station and the motors after these had been resolved.

Figure 4.31: Nord West 1100 used for demonstration of the proof-of-concept.

When the wires were disconnected and the wireless bridge was to be used, the system lost the steering wheel upon starting the starboard motor. The connection was however working and the motors where fully controllable. The problem with losing the steering wheel was however also observed when standard CAN connections later was used and might be due to a connector that was found not to be fully mated.

The demonstration was judged to be successful since the default system was working well after initial errors in cabling had been fixed.

5

Discussion and future work

The discussion of the results has been divided into the three categories Network and protocols, Hardware and Extensions. This is done to give a clear overview of results in the different areas and work that remains to be done.

5.1 Network and protocols

This thesis has discussed the possibility to use Ethernet as a future substitute for CAN. There exist a few alternatives to CAN in addition to Ethernet. One of these is FlexRay that is already implemented in number of solutions. Another possibility is BroadR-Reach which aimed at the automotive industry, provides standard Ethernet with a new physical layer that includes transmission over a single unshielded twisted pair at 100 Mbit/s.

CAN was designed with reliability and safety in mind. This has resulted in a protocol that is resilient to disturbance and includes mechanisms to ensure correct functionality. Ethernet is on the other hand originally not designed for use in safety and time critical applications. This has resulted in that most of the security mechanisms are implemented on higher layers and it shows sensitivity to electromagnetic disturbance. It shall however not be forgotten that it provides a large increase in bandwidth that could be used to overcome some of the limitations of the protocol.

The network topology of a system has a large impact on the characteristics it possesses. Every system has different priorities and there can therefore not exist a single optimal solution. Instead the topology must be carefully designed according to the desired properties of the system.

Once such a topology has been decided it can also be extended in a very dynamic way. Since a physical separation no longer is needed with a larger bandwidth, the system also

becomes more flexible when it comes to data distribution between nodes. If mechanisms such as DHCP and ARP are used the configuration of the system is also simplified.

An aid in network management and routing are available protocols such as RSTP and SNMP. These allow redundant network topologies to be used and analysed.

VLAN has been identified as a flexible and secure way to logically separate streams and groups. This was unfortunately not possible to use during the demonstration due to limitations in the routers. In a future system it would be recommended to use hardware with VLAN support because of its advantages such as being able to reduce the traffic generated by broadcast messages.

The possibility to select the frame size over a broad range of values contributes to the flexibility of an Ethernet based solution. The choice of frame size does affect several different aspects of the system. A small frame could be preferred in that it decreases jitter, risk for frame drop and process coupling. A large frame does on the other hand provide better utilisation of the link. This makes it one of the parameters that must be decided for each individual system. Some of these problems such as jitter can be modelled with software such as the TCN NetAnalyzer.

In the demonstration, a futuristic set up with wireless transmission was tested. This is with the used technology not a reasonable solution due to latency and security issues. A more specialised point to point wireless technology might however be used for everything but the most safety critical data.

5.2 Hardware

The proof-of-concept has been created with a development board from IAR. This board is very versatile and easy to work with. The drawback to use a development board is that the hardware is fixed. It would be interesting to see what effect more suitable hardware such as an Ethernet transceiver designed for automotive would have on its properties. This would also enable more relevant EMC tests where the converter can be tested to see how resilient it can become.

The connector is an integral part of a communications system. Aspects such as number of pins, signal attenuation and ingress protection play an important role in connector selection. Non-technical aspects such as the level of adoption and price might be equally important. This makes the selection of a specific connector not a purely technical choice but deeply dependant on other aspects as well.

Selection of a cable poses similar questions and issues to that of the connector selection. Section 4.2 shows that standard Ethernet communication can be conducted over a non standard cable even if the link's properties might be affected. That opens up for the possibility to use a large set of cables, but tests must be made to ensure reliability and performance.

A selection of standard physical media for Ethernet networking is available. Fibre optics have the advantage of not being sensitive to electromagnetic disturbance but have drawbacks such as being fragile. This option does however deserve a thorough investigation to see if the advantages outweigh the disadvantages.

It has become clear that there exist no single best alternative, selection of cable and connector must instead be made for a specific application or system. This means that this selection must be done when currently unclear aspects has been further investigated.

5.3 Extensions

One of the major advantages with Ethernet is the ease of which it can be extended. In this thesis this has been shown using COTS Android devices and simple application. There is no reason to only use this extensibility with phones and tablets. Android is an open system and could be used for building displays, controls or just speed up prototyping.

Consider a mobile application that allows on site personnel to stream system information live over mobile networks to a service center. This functionality could decrease time and cost for identifying unwanted system behaviour and create a higher service value for the customer. This is but one possibility that becomes easily implemented when using Ethernet as a medium of communication.

5.4 Final conclusions

This thesis has shown that the major drawbacks with Ethernet is the sensitivity to electromagnetic disturbance and its non deterministic properties. The non-determinism can be handled by modelling the system using appropriate software, but the sensitivity to electromagnetic disturbance needs further investigation. Advantages such as high bandwidth and extendibility could very well make Ethernet commonly used in vehicles within the next few years.

Bibliography

[1] S. Corrigan, Introduction to the Controller Area Network (CAN), Tech. rep., Texas Instruments (2008).

[2] B. A. Forouzan, TCP/IP Protocol Suite, Fourth Edition, McGraw-Hill, 2010.

[3] D. E. Comer, Internetworking With TCP/IP Volume I Principles, Protocols, and Architecture, Fifth Edition, PEARSON Prentice Hall, 2006.

[4] IEEE, IEEE Std 802.3 - 2005 Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications, IEEE Std 802.3-2005 (Revision of IEEE Std 802.3-2002 including all approved amendments) Section1.

[5] IEEE, IEEE Std 802.3 - 2005 Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications, IEEE Std 802.3-2005 (Revision of IEEE Std 802.3-2002 including all approved amendments) Section2.

[6] InterNational Committee for Information Technology Standards, Fibre Distributed Data Interface (FDDI) - Token Ring Twisted Pair Physical Layer Medium Dependent (TP-PMD), American National Standards Institute.

[7] InetDaemon, Local Area Network Topologies, accessed: 2012-04-11 (Mar. 2012).
URL `http://www.inetdaemon.com/tutorials/networking/lan/topology.shtml`

[8] A. K. Singh, Computer Networks, Laxmi Publications, 2006.

[9] B. Cain, S. Deering, I. Kouvelas, B. Fenner, A. Thyagarajan, Ericsson, Internet Group Management Protocol, Version 3, RFC 3376 , accessed: 2012-03-08 (Jul. 2002).
URL `http://www.ietf.org/rfc/rfc3376.txt`

[10] IEEE, IEEE 802.1, accessed: 2012-05-30 (Mar. 2012).
URL http://www.ieee802.org/1/

[11] Westermo, Westermo OS Management Guide Version 4.8.0-0, accessed: 2012-05-30 (2012).
URL http://www.westermo.net/dman/Document.phx/Manuals/Ethernet/Ethernet+Switches/WeOS+Management+Guide.pdf?folderId=%2FManuals%2FEthernet%2FEthernet+Switches&cmd=download

[12] S. McQuerry, D. Jansen, D. Hucaby, Cisco LAN switching configuration handbook, Second Edition, Cisco Press, 800 East 96th Street, Indianapolis, IN 46240 USA, 2009.

[13] R. Seifert, J. Edwards, The All-New Switch Book: The Complete Guide to LAN Switching Technology, Second Edition, Wiley Publishing, Inc, 2008.

[14] I. Muhammad, Ethernet in Steer-by-wire Applications, Master's thesis, KTH Royal Institute of Technology, Stockholm, Sweden (2011).

[15] Time Criticals Networks AB, TCN StreamAnalyzer (02 2012).

[16] DEW Associates Corporation, Introduction to Fast Ethernet, accessed: 2012-06-02 (Jan. 2000).
URL http://www.dewassoc.com/support/networking/fastethernet.htm

[17] T. Rybak, M. Steffka, Automotive Electromagnetic Compatibility (EMC), Kluwer Academic Publishers, 2004.

[18] Panduit, The Evolution of Copper Cabling Systems from Cat5 to Cat5e to Cat6, Tech. rep., Panduit (2003).

[19] Ab-Tech-Solutions, LonWorks, accessed: 2012-05-30 (Mar. 2012).
URL http://abtechteam.com/fieldbus_LonWorks.html

[20] S. Spitzer, Why NMEA 2000, accessed: 2012-05-20 (Mar. 2009).
URL http://www.nmea.org/Assets/nmea2000pdf.pdf

[21] D. Gratton, What does NMEA2000 compatible mean, accessed: 2012-05-20 (Jun. 2009).
URL http://www.nmea.org/Assets/what%20does%20nmea2000%20certified%20mean%2020090628.pdf

[22] Blue Heron Marine, Airmar WS2-C10 NMEA 2000 Connector Cable 10M, accessed: 2012-06-03.
URL http://www.blueheronmarine.com/Airmar-WS2-C10-NMEA-2000-Connector-Cable-10M-6603

[23] The Garmin blog team, Award winning NMEA 2000 adapters give your boat new life, accessed: 2012-05-20 (Feb. 2010).
URL http://garmin.blogs.com/my_weblog/2010/02/award-winning-nmea-2000-adaptors-give-your-boat-new-life.html

[24] Amphenol LTW Technology Co. Ltd., Industrial Ethernet M12 to RJ45 Adaptor, accessed: 2012-05-21 (May 2012).
URL http://www.ltw-tech.com/news/news.php?la=en-us&n=175

[25] International Electrotechnical Commission, Connectors for electronic equipment – Product requirements – Part 2-101: Circular connectors – Detail specification for M12 connectors with screw-locking.

[26] ShieldConnectors, M12 round female / male connector, panel receptacle version. IEC 61076-2-101 norms., accessed: 2012-06-03.
URL http://farm5.staticflickr.com/4041/4250300623_e099037aed_b.jpg

[27] Deutch, DT, accessed: 2012-05-21.
URL http://www.ltw-tech.com/news/news.php?la=en-us&n=175

[28] molex, MX150TM Sealed Connector System 3.50mm (.138") Pitch, accessed: 2012-05-21 (Aug. 2011).
URL http://www.molex.com/elqNow/elqRedir.htm?ref=http://rhu004.sma-promail.com/SQLImages/kelmscott/Molex/PDF_Images/987650-2411.PDF

[29] Molex, MX150 Sealed Connector System, accessed: 2012-06-03.
URL http://www.molex.com/elqNow/elqRedir.htm?ref=http://rhu004.sma-promail.com/SQLImages/kelmscott/Molex/PDF_Images/987650-2842.PDF

Appendix A

Radiated immunity EMC

Appendix showing complete results of the performed EMC immunity test performed at IVF's facilities in Mölndal. Observe that no results, for the Cat 5e cable with parallel pairs (parallel Cat 5e) are collected. This is due to that the cable performance could be proven degraded to a point where test equipment could not establish a connection, see Section 4.3.1 for details.

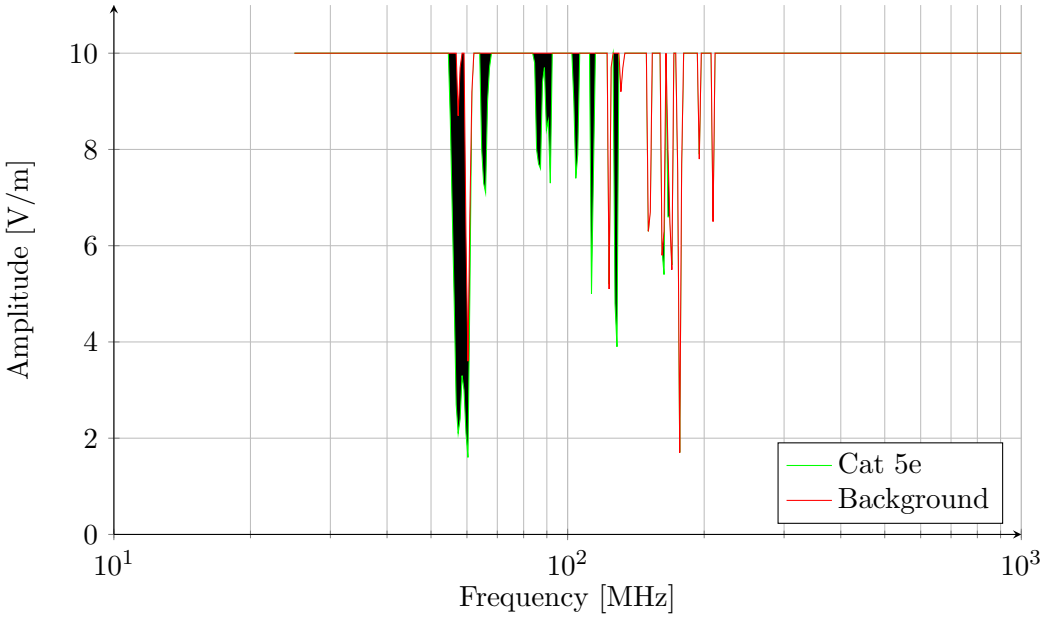

Figure A.1: Plot showing link performance for Cat 5e cable where dark areas mark zones where full functionality is lost.

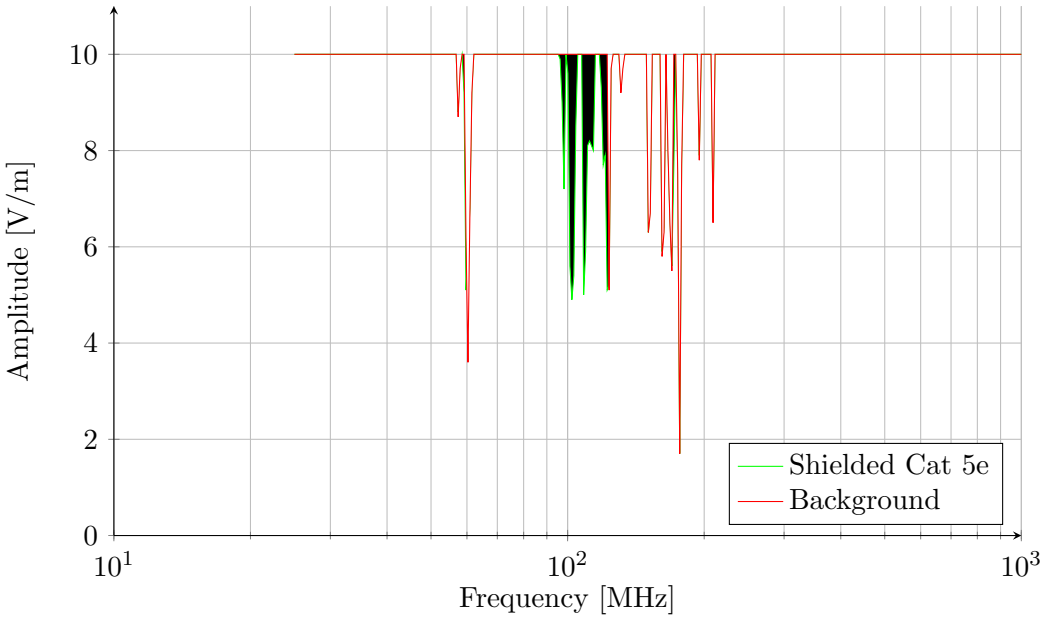

Figure A.2: Plot showing link performance for the shielded Cat 5e cable where dark areas mark zones where full functionality is lost.

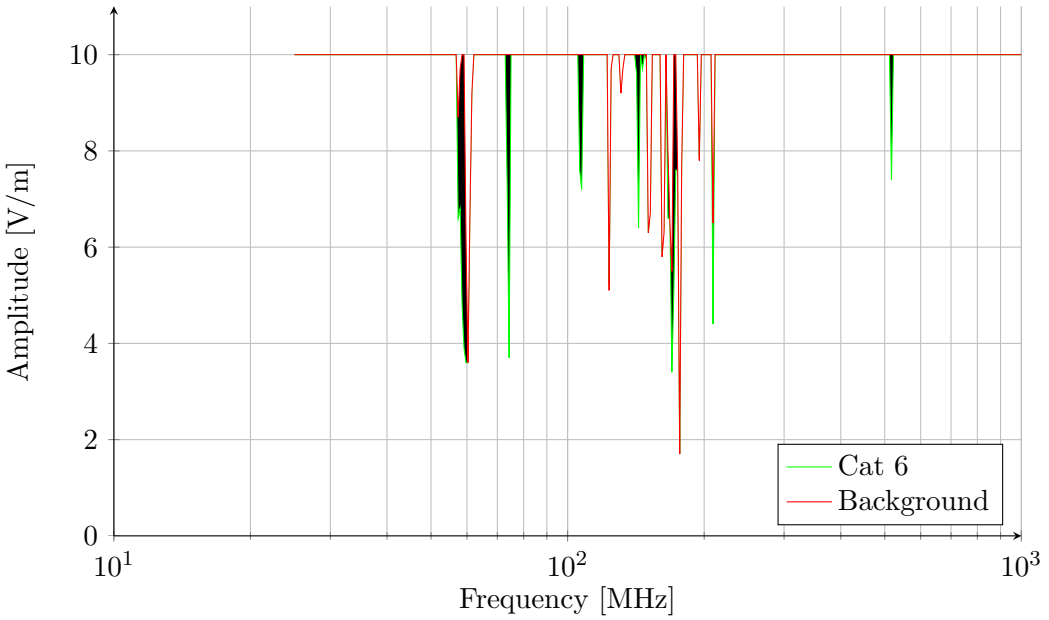

Figure A.3: Plot showing link performance for the UTP Cat 6 cable where dark areas mark zones where full functionality is lost.

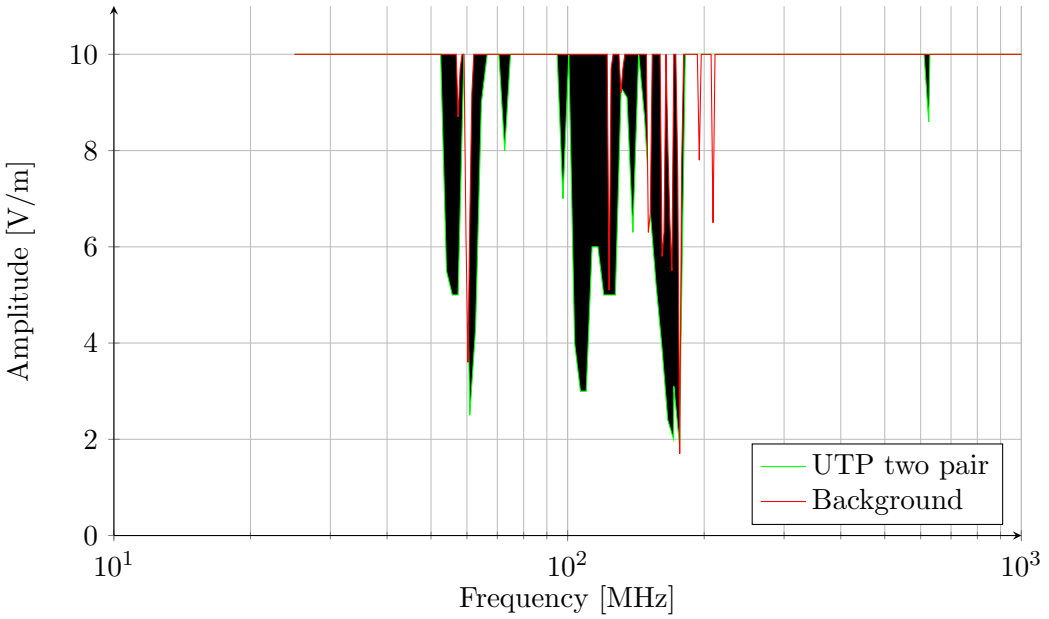

Figure A.4: Plot showing link performance for the UTP double pair cable where dark areas mark zones where full functionality is lost.

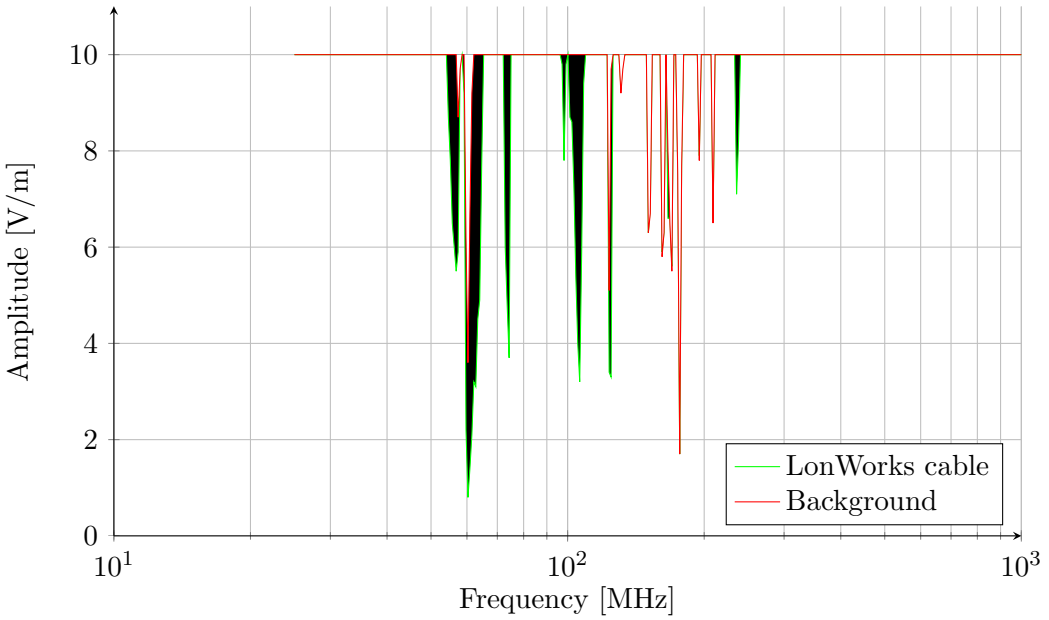

Figure A.5: Plot showing link performance for the LonWorks cable where dark areas mark zones where full functionality is lost.

Appendix B

Radiated emission background

Appendix containing a scanned copy (Figure B.1) of the printed results from the low frequency radiated emission background. The digital results from this measurement was unfortunately lost after the EMC measurements were performed.

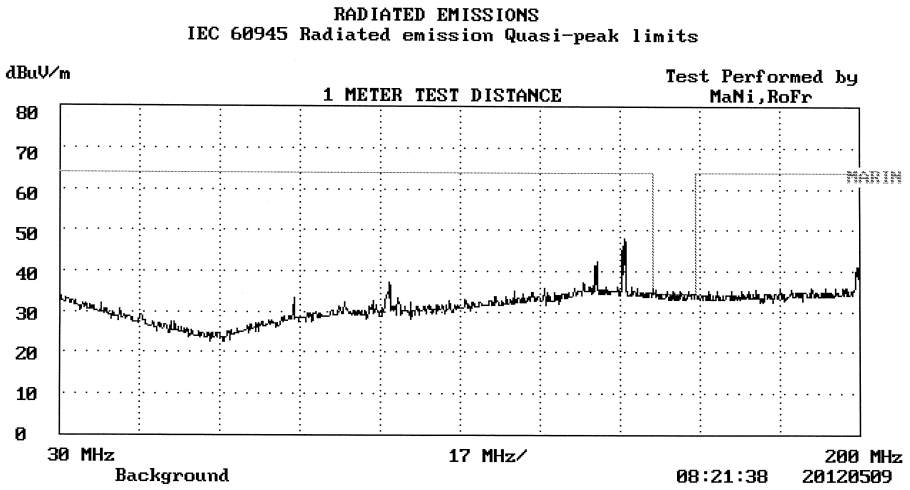

Figure B.1: Radiated emission background measurements.