



Technical Introduction and Overview

December 2004



EtherCAT — Ethernet Control Automation Technology

This paper introduces EtherCAT, an Ethernet-based fieldbus system. EtherCAT sets new performance standards. Handling is straightforward and similar to a fieldbus, thanks to flexible topology and simple configuration. Moreover, since EtherCAT can be implemented very cost-effectively, the system enables fieldbusses to be used in applications where fieldbus networking was not an option in the past.

EtherCAT is an open technology for which IEC standardization is in progress. Support is provided by the EtherCAT Technology Group, an international association of users and manufacturers with more than 150 member companies.

1. Introduction

Fieldbusses have become an integrated component of automation technology. They have been tried and tested and are now widely established. It was fieldbus technology that enabled the wide-scale application of PC-based control systems. While the performance of controller CPUs - particularly for IPCs - is increasing rapidly, conventional fieldbus systems tend to represent "bottlenecks" that limit the performance control systems can achieve. An additional factor is the layered control architecture consisting of several subordinate (usually cyclic) systems: the actual control task, the fieldbus system and perhaps local expansion busses within the I/O system or simply the local firmware cycle in the peripheral device. Response times are typically 3-5 times higher than the controller cycle time - an unsatisfactory solution (see Fig. 1).

Above the fieldbus system level, i.e. for networking controllers, Ethernet has already been the state of the art for some time. What is new is its application at the drive or I/O level, i.e. in areas that were dominated by fieldbus systems in the past. The main requirements

for this type of application are high real-time capability, suitability for small data quantities, and naturally cost-effectiveness. EtherCAT meets these requirements and at the same time makes internet technologies available at the I/O level.

1.1 Ethernet and real-time capability

There are many different approaches that try and provide real-time capability for Ethernet: for example, the CSMA/CD media access procedure is disabled via higher level protocol layers and replaced by the time slice procedure or polling; other propositions use special switches that distribute Ethernet packets in a precisely controlled timely manner. Whilst these solutions may be able to transport data packets more or less quickly and accurately to the connected Ethernet nodes, the times required for the redirection to the outputs or drive controllers and for reading the input data strongly depend on the implementation.

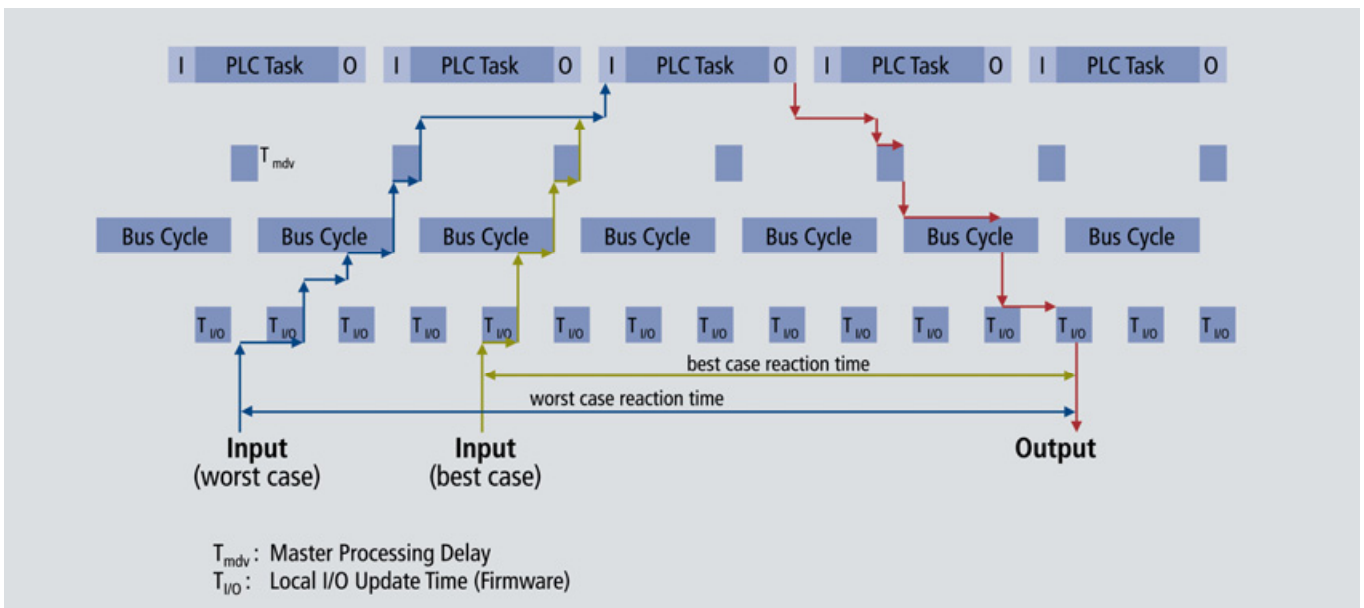


Figure 1: Response times of conventional fieldbus systems

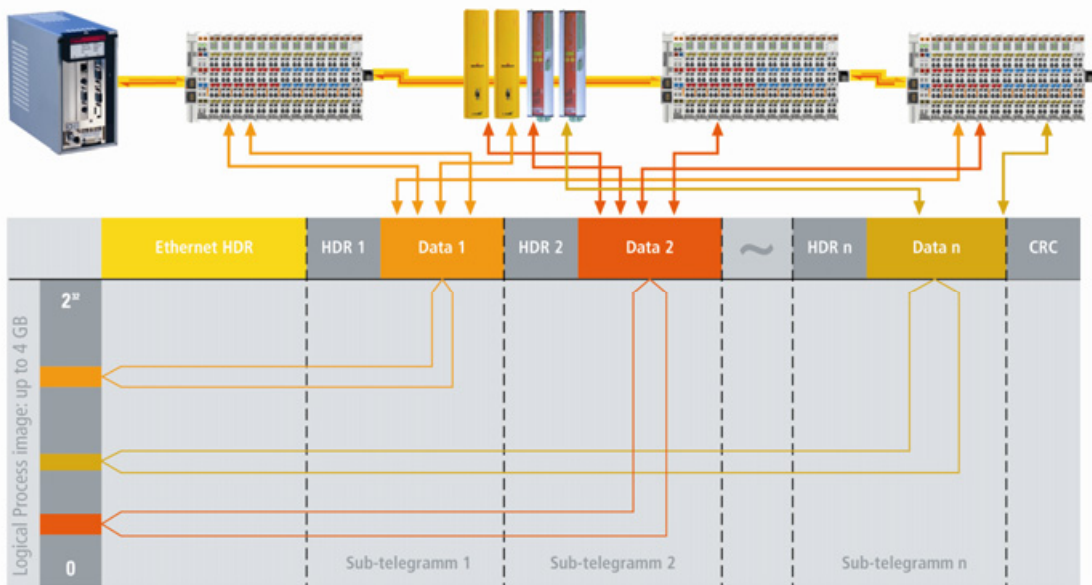


Figure 2: Process data is inserted in telegrams

If individual Ethernet frames are used for each device, the usable data rate is very low in principle: The shortest Ethernet frame is 84 bytes long (incl. inter-packet gap IPG). If, for example, a drive cyclically sends 4 bytes of actual value and status information and accordingly receives 4 bytes of command value and control word information, at 100% bus load (i.e. with infinitely short response time of the drive) a usable data rate of only $4/84 = 4.7\%$ is achieved. At an average response time of $10 \mu\text{s}$, the rate drops to 1.9%. These limitations apply to all real-time Ethernet approaches that send an Ethernet frame to each device (or expect a frame from each device), irrespective of the protocols used within the Ethernet frame.

2. EtherCAT operating principle

EtherCAT technology overcomes these inherent limitations of other Ethernet solutions: the Ethernet packet is no longer received, then interpreted and process data then copied at every device. The EtherCAT slave devices read the data addressed to them while the frame passes through the

node. Similarly, input data is inserted while the telegram passes through (see **Fig. 2**). The frames are only delayed by a few nanoseconds.

Since an Ethernet frame comprises the data of many devices both in send and receive direction, the usable data rate increases to over 90%. The full-duplex features of 100BaseTX are fully utilised, so that effective data rates of $> 100 \text{ Mbps}$ ($>90\%$ of $2 \times 100 \text{ Mbps}$) can be achieved (see **Fig. 3**).

The Ethernet protocol according to IEEE 802.3 remains intact right up to the individual terminals; no sub-bus is required. In order to meet the requirements of the electronic terminal block, only the physical layer in the coupler is converted from twisted pair or optical fibre to E-bus (alternative Ethernet physical layer: LVDS according to [4,5]). The signal type within the terminal block (E-bus) is also suitable for transfer via a twisted pair line over short distances (up to 10 m). The terminal block can thus be extended very cost-efficiently. Subsequent conversion to Ethernet 100BASE-TX physical layer is possible at any time.

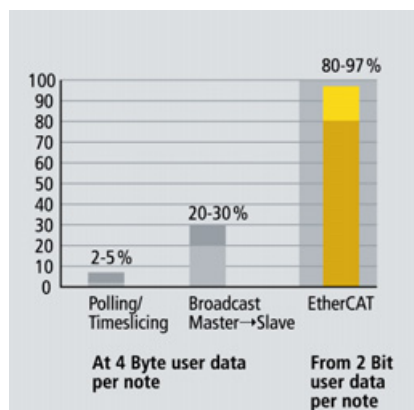


Figure 3: Comparison of Bandwidth Utilisation

3. EtherCAT features

3.1 Protocol

The EtherCAT protocol is optimised for process data and is transported directly within the Ethernet frame thanks to a special Ethertype. It may consist of several sub-telegrams, each serving a particular memory area of the logical process images that can be up to 4 gigabytes in size. The data sequence is independent of the physical order of the Ethernet




Figure 4: EtherCAT: Standard Frames according to IEEE 802.3 [3]

terminals in the network; addressing can be in any order. Broadcast, Multicast and communication between slaves are possible. Direct Ethernet frame transfer is used in cases where maximum performance is required and the EtherCAT components are operated in the same subnet as the controller.

However, EtherCAT applications are not limited to a subnet: EtherCAT UDP packages the EtherCAT protocol into UDP/IP datagrams (see Fig. 4). This enables any control with Ethernet protocol stack to address EtherCAT systems. Even communication across routers into other subnets is possible. In this variant, system performance obviously depends on the real-time characteristics of the control and its Ethernet protocol implementation. The response times of the EtherCAT network itself are hardly restricted at all: the UDP datagram only has to be unpacked in the first station.

EtherCAT only uses standard frames according to [3] - the frames are not shortened. EtherCAT frames can thus be sent from any Ethernet controller (master), and standard tools (e.g. monitor) can be used.

3.2 Topology

Line, tree or star: EtherCAT supports almost any topology (see Fig. 5). The bus or line structure known from the field-busses thus also becomes available for Ethernet.

Particularly useful for system wiring is the combination of line and branches or stubs: the required interfaces exist on

the couplers; no additional switches are required. Naturally, the classic switch-based Ethernet star topology can also be used.

Wiring flexibility is further maximised through the choice of different cables. Flexible and inexpensive standard Ethernet patch cables transfer the signals optionally in Ethernet mode (100BASE-TX) or in E-bus signal representation. Plastic optical fibres (POF) complement the system for special applications. The complete choice of Ethernet wiring - such as different optical fibres and copper cables - can be used in combination with switches or media converters.




Figure 5: Flexible Topology: Line, Tree or Star

The Fast Ethernet physics (100BASE-TX) enables a cable length of 100 m between two devices, the E-bus line is intended for distances of up to 10m. For each connection, the signal variant can be selected individually. Since up to 65535 devices can be connected, the size of the network is almost unlimited.

3.3 Distributed clocks

Accurate synchronisation is particularly important in cases where spatially distributed processes require simultaneous actions. This may be the case, for example, in applications where several servo axes carry out coordinated movements simultaneously.

The most powerful approach for synchronisation is the accurate alignment of distributed clocks, as described in the new IEEE 1588 standard [6]. In contrast to fully synchronous communication, where synchronisation quality suffers immediately in the event of a communication fault, distributed aligned clocks have a high degree of tolerance vis-à-vis possible fault-related delays within the communication system.

With EtherCAT, the data exchange is fully based on a pure hardware machine. Since the communication utilises a logical (and thanks to full-duplex Fast Ethernet also physical) ring structure, the master clock can determine the propagation delay offset to the individual slave clocks simply and accurately - and vice versa. The distributed clocks are adjusted based on this value, which means that a very precise network-wide timebase with a jitter of significantly less than 1 microsecond is available (See Fig. 6). External synchronisation, e.g. across the plant, is then based on IEEE 1588.

However, high-resolution distributed clocks are not only used for synchronisation, but can also provide accurate information about the local timing of the data acquisition. For example, motion controllers typically calculate velocities from sequentially measured positions. Particularly with very short sampling times, even a small temporal jitter in the position measurement leads to large step changes in the computed velocity. With EtherCAT, timestamp data types are introduced as a logical extension. The high-resolution system time is linked to the measured value, which is made possible by the large bandwidth offered by Ethernet. The accuracy of a velocity calculation then no longer depends on the jitter of the communication system. It is orders of magnitude better than that of measuring techniques based on jitter-free communication.

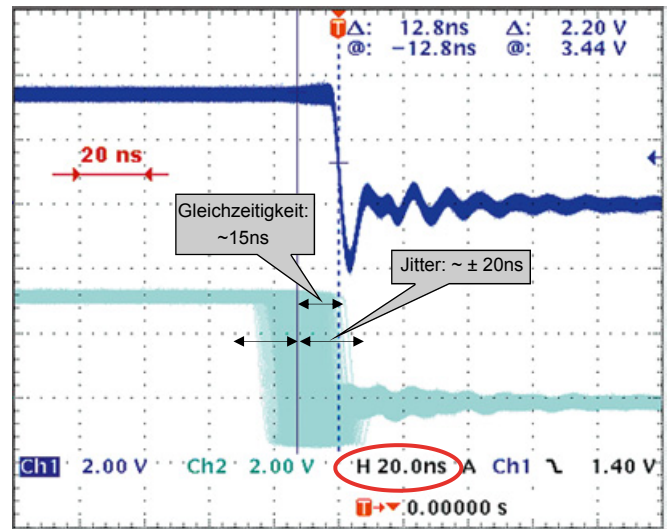


Figure 6: Synchronicity and Simultaneity: Scope view of two distributed devices with 300 nodes and 120m of cable between them

3.4 Performance

EtherCAT reaches new dimensions in network performance. Thanks to hardware integration in the slave and direct memory access to the network card in the master, the complete protocol processing takes place within hardware and is thus fully independent of the run-time of protocol stacks, CPU performance or software implementation. The update time for 1000 I/Os is only 30 µs - including I/O cycle time (see Table 1). Up to 1486 bytes of process data can be exchanged with a single Ethernet frame - this is equivalent to almost 12000 digital inputs and outputs. The transfer of this data quantity only takes 300 µs.

The communication with 100 servo axes only takes 100 µs. During this time, all axes are provided with command values and control data and report their actual position and status. The distributed clock technique enables the axes to be synchronised with a deviation of significantly less than 1 microsecond.

Prozessdaten	Update-Zeit
256 verteilte digitale E/A	11 µs = 0,01 ms
1000 verteilte digitale E/A	30 µs
200 analoge E/A (16 bit)	50 µs ↔ 20 kHz
100 Servo Achsen, je 8 Byte Ein- und Ausgangsdaten	100 µs
1 Feldbus Master-Gateway (1486 Bytes Eingangs- und 1486 Bytes Ausgangsdaten)	150 µs

Table 1: EtherCAT Performance Overview

The extremely high performance of the EtherCAT technology enables control concepts that could not be realised with classic fieldbus systems. With EtherCAT, a communication technology is available that matches the superior computing capacity of modern Industrial PCs. The bus system is no longer the bottleneck of the control concept. Distributed I/Os are recorded faster than is possible with most local I/O interfaces. The EtherCAT technology principle is scalable and not bound to the baud rate of 100 MBaud – extension to GBit Ethernet is possible.

3.5 Diagnosis

Experience with fieldbus systems shows that availability and commissioning times crucially depend on the diagnostic capability. Only faults that are detected quickly and accurately and located unambiguously can be rectified quickly. Therefore, special attention was paid to exemplary diagnostic features during the development of EtherCAT.

During commissioning, the actual configuration of the nodes (e.g. drives or I/O terminals) should be checked for consistency with the specified configuration. The topology should also match the configuration. Due to the built-in topology recognition down to the individual terminals, this verification can not only take place during system start-up, automatic reading in of the network is also possible (configuration upload).

Bit faults during the transfer are reliably detected through evaluation of the CRC checksum: The 32 bit CRC polynomial has a minimum hamming distance of 4. Apart from broken wire detection and localisation, the protocol, physical layer and topology of the EtherCAT system enable individual quality monitoring of each individual transmission segment. The automatic evaluation of the associated error counters enables precise localisation of critical network sections. Gradual or changing sources of error such as EMI influences, defective connectors or cable damage are detected and located, even if they do not yet overstrain the self-healing capacity of the network.

3.6 EtherCAT instead of PCI

With increasing miniaturisation of the PC-components, the physical size of Industrial PCs is increasingly determined by the number of required slots. The bandwidth of Fast Ethernet, together with the process data width of the EtherCAT communication hardware enables new directions: clas-




Figure 7: Decentralized Fieldbus Interfaces

sic interfaces that are conventionally located in the IPC are transferred to intelligent EtherCAT interface terminals (see Fig. 7). Apart from the decentralised I/Os, drives and control units, complex systems such as fieldbus masters, fast serial interfaces, gateways and other communication interfaces can be addressed.

Even further Ethernet devices without restriction on protocol variants can be connected via decentralised hub terminals. The central IPC becomes smaller and therefore more cost-effective. One Ethernet interface is sufficient for the complete communication with the periphery (see Fig. 8).

3.7 Device profiles

The device profiles describe the application parameters and the functional behaviour of the devices including the device class-specific state machines. For many device classes, fieldbus technology already offers reliable device profiles, for example for I/O devices, drives or valves. Users are familiar with these profiles and the associated parameters and tools. No EtherCAT-specific device profiles have therefore been developed for these device classes. Instead, simple interfaces for existing device profiles are being offered. This greatly assists users and device manufacturers alike during the migration from the existing fieldbus to EtherCAT.




Figure 8: EtherCAT leads to smaller Controllers

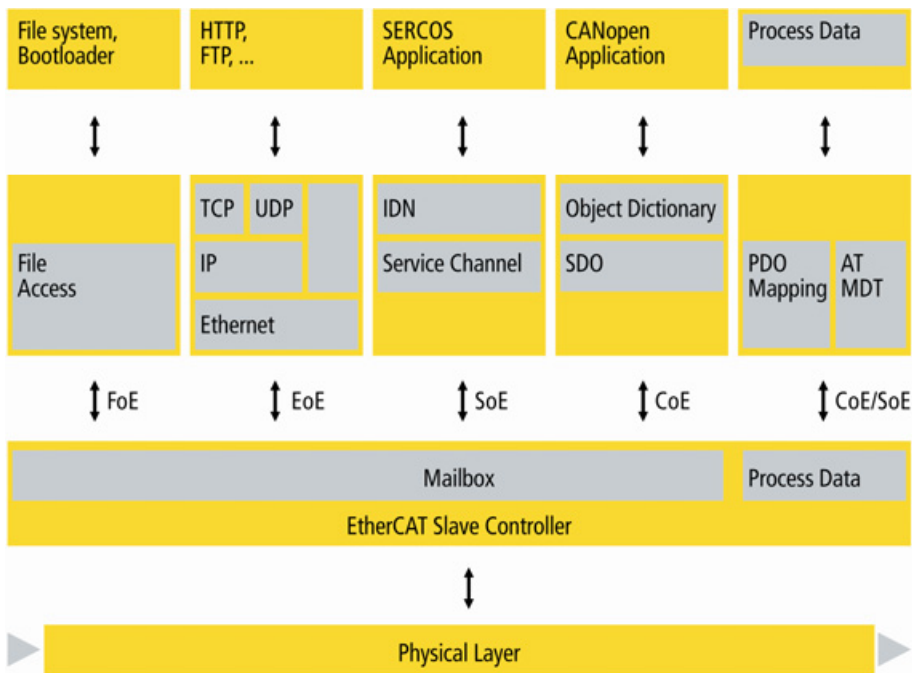


Figure 9: Several Device Profiles and Protocols can co-exist side by side

3.7.1 CANopen over EtherCAT (CoE)

CANopen device and application profiles are available for a wide range of device classes and applications, ranging from I/O components, drives, encoders, proportional valves and hydraulic controllers to application profiles for plastic processing or textile machines, for example. EtherCAT can provide the same communication mechanisms as the familiar CANopen [9] mechanisms: object dictionary, PDO (process data objects) and SDO (service data objects) - even the network management is comparable. EtherCAT can thus be implemented with minimum effort on devices equipped with CANopen. Large parts of the CANopen firmware can be reused. Objects can optionally be expanded in order to account for the larger bandwidth offered by EtherCAT.

3.7.2 Servodrive-Profil according to IEC 61491 over EtherCAT (SoE)

SERCOS interface™* is acknowledged and appreciated worldwide as a high-performance real-time communication interface, particularly for motion control applications. The SERCOS profile for servo drives and the communication technology are covered by the IEC 61491 standard [10]. This servo drive profile can very easily be mapped to EtherCAT. The service channel, and therefore access to all parameters and functions residing in the drive, is based on the EtherCAT mailbox (see Fig. 9).

Here too, the focus is on compatibility with the existing protocol (access to value, attribute, name, units etc. of the

IDNs) and expandability with regard to data length limitation. The process data, with SERCOS in the form of AT and MDT data, are transferred using EtherCAT slave controller mechanisms. The mapping is similar to the SERCOS mapping. The EtherCAT slave state machine can also be mapped easily to the phases of the SERCOS protocol. Real-time Ethernet technology is therefore already available for this device profile, which is particularly widespread in CNC applications. The benefits of the device profile are combined with the benefits offered by EtherCAT. Distributed clocks ensure precise network-wide synchronisation. Optionally, the command position, speed or torque can be transferred. Depending on the implementation, it is even possible to continue using the same configuration tools for the drives.

3.8 Ethernet over EtherCAT (EoE)

The EtherCAT technology is not only fully Ethernet-compatible, but also characterised by particular openness "by design": the protocol tolerates other Ethernet-based services and protocols on the same physical network - usually even with minimum loss of performance. There is no restriction on the type of Ethernet device that can be connected within the EtherCAT segment via a switch port. The Ethernet frames are tunneled via the EtherCAT protocol, which is the standard approach for internet applications (e.g. VPN, PPPoE (DSL) etc.). The EtherCAT network is fully transparent for the Ethernet device, and the real-time characteristics are not impaired (see Fig. 10).

*SERCOS interface is a trademark of the Interest Group SERCOS e.V.




Figure 10: Transparent for all Ethernet Protocols

EtherCAT devices can additionally feature other Ethernet protocols and thus act like a standard Ethernet device. The master acts like a layer 2 switch that redirects the frames to the respective devices according to the address information. All internet technologies can therefore also be used in the EtherCAT environment: integrated web server, e-mail, FTP transfer etc.

3.9 File Access over EtherCAT (FoE)

This very simple protocol similar to TFTP enables access to any data structure in the device. Standardized firmware upload to devices is therefore possible, irrespective of whether or not they support TCP/IP.

4. Implementation

4.1 Master

EtherCAT uses standard Ethernet controllers where real cost savings can be achieved: in the master. No communication coprocessors are required, since usually only one Ethernet frame has to be sent per cycle. EtherCAT therefore is the only Ethernet solution for demanding real-time requirements that does not require special master plug-in cards. The on-board Ethernet controller or a cost-effective standard NIC card is sufficient. The master is usually implemented as a pure software solution. Implementation of an EtherCAT master is very easy, particularly for small and medium-sized control systems and for clearly defined applications.

For example a PLC with a single process image: if it does not exceed 1488 bytes, cyclic sending of a single Ethernet frame with the cycle time of the PLC is sufficient. Since the header does not change at run time, all that is required is a constant header to be added to the process image and the result to be transferred to the Ethernet controller (see Fig. 11).

Ethernet Header			ECAT	EtherCAT Telegram			Ethernet		
DA	SA	Type	Frame HDR	EtherCAT HDR	Data	CTR	Pad.	FCS	
(6)	(6)	(2)	(2)	(10)	(0...1486)	(2)	(0...32)	(1)	
constant Header			completely sorted (mapped) process data			Working Counter: constant	Padding Bytes and CRC generated by Ethernet Controller (MAC)		

Figure 11: Master-Implementation with one Process Image

The process image is already sorted, since with EtherCAT mapping does not occur in the master, but in the slaves - the peripheral devices insert their data at the respective points in the passing frame.

4.1.1 Master Sample Code

Master sample code for supporting a master implementation is available for a nominal fee. The software is supplied as source code and comprises all EtherCAT master functions, including Ethernet over EtherCAT. All the user has to do is adapt the code, which was created for Windows environments, to the target hardware and the RTOS used (see Fig. 12).

4.1.2 Monitor Tools

Since EtherCAT uses standard Ethernet frames according to IEEE 802.3, any commercially available Ethernet monitoring tool can be used for monitoring EtherCAT communication. In addition, free parser software for Ethereal (an open source monitoring tool) and the Microsoft network monitor are available for processing and displaying recorded EtherCAT data traffic.

4.2 Slave

A cost-effective EtherCAT slave controller (ASIC or FPGA) is used in the slave device. For simple devices, no additional microcontroller is required. For more complex devices, EtherCAT communication performance is almost independent of the performance capability of the controller used, making the device cost-effective.

EtherCAT slave controllers are available or are in preparation from several manufacturers. The slave controllers developed by Beckhoff (available via semiconductor distribution) feature an internal DPRAM and offer a range of interfaces for accessing this application memory:

- The serial SPI (serial peripheral interface) is intended particularly for devices with small process data quantity, such as analog I/O modules, sensors, encoders or simple drives.
- The 32-bit parallel I/O interface is suitable for the connection of up to 32 digital inputs/outputs, but also for simple sensors or actuators operating with 32 data bits.
- The parallel 8/16-bit microcontroller interface corresponds to conventional interfaces for fieldbus controllers with DPRAM interface. It is particularly suitable for more complex devices with larger data volume.

For supporting a slave implementation, a cost-effective evaluation kit is available, including slave application software in source code and a test master.

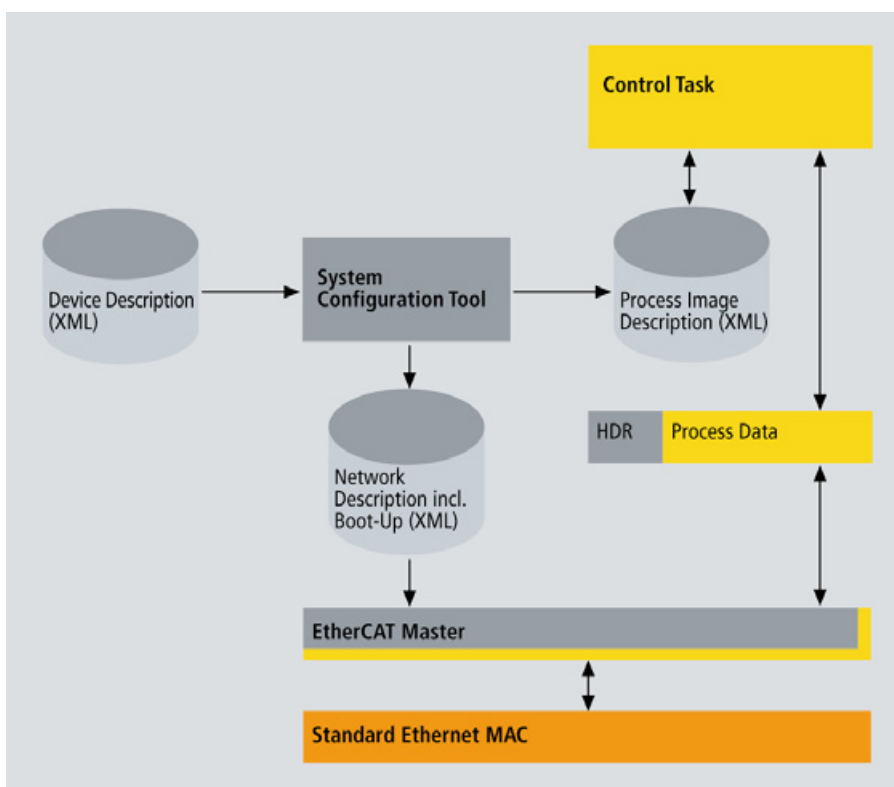


Figure 12: Structure of Master Sample Code

4.3 Infrastructure costs

Since no hubs and switches are required for EtherCAT, costs associated with these devices including power supply, installation etc. are avoided. Standard CAT5 cables and standard connectors are used, if the environmental conditions permit this.

5. EtherCAT Technology Group

Everyone should be able to use and implement EtherCAT. The EtherCAT Technology Group promotes this philosophy. The ETG is a forum for end users from different sectors, and for machine manufacturers and suppliers of powerful control technology with the aim of supporting and promoting EtherCAT technology. The wide range of industry sectors that are represented ensures that EtherCAT is optimally prepared for a large number of applications. With their qualified feedback, the system partners ensure simple integration of the hardware and software components in all required device classes.

ETG was established in November 2003 and is the fastest growing fieldbus organisation, with the number of member companies currently standing at 140. ETG has member companies from Austria, Belgium, Canada, China, Germany,

Finland, France, Great Britain, Israel, Italy, Korea, Liechtenstein, the Netherlands, , Singapore, Sweden, Switzerland, Taiwan, and the USA (see Fig. 13).

6. International Standardisation

Disclosure is not only driven from within the EtherCAT Technology Group - the international standardisation of EtherCAT has also been initiated already. The IEC management board has accepted the EtherCAT Technology Group as an official liaison partner of the IEC working groups for digital communication. The EtherCAT specification was submitted to IEC in November 2004 and has already advanced to the voting process. ISO has also accepted an accelerated standardisation procedure for EtherCAT, so that EtherCAT is expected to obtain the status of an official IEC and ISO specification quite soon.

7. EtherCAT is proven

EtherCAT is already used commercially: following positive experience with a pilot system, Schuler Pressen AG decided to approve EtherCAT for their Proffline press range. A repre-



Figure 13: ETG Member Companies (as of Nov 04)

sentative from the company Schuler said: "This system enables us to realise fast drive and hydraulic controls for all applications currently used in the Schuler Group. Another crucial factor is that, due to EtherCAT's performance, we still have enough potential for solving complex control tasks in future without performance problems." Meanwhile, more than 3500 nodes have been supplied to Schuler and other customers.

8. Summary

EtherCAT is characterised by outstanding performance, very simple wiring and openness for other protocols. EtherCAT sets new standards where conventional fieldbus systems reach their limits: 1000 I/Os in 30 μ s, optionally twisted pair cable or optical fibre and, thanks to Ethernet and Internet technologies, optimum vertical integration. With EtherCAT, the costly Ethernet star topology can be replaced with a simple line structure - no expensive infrastructure components are required. Optionally, EtherCAT may also be wired in the classic way using switches, in order to integrate other Ethernet devices. Where other real-time-Ethernet approaches require special connections in the controller, for EtherCAT very cost-effective standard Ethernet cards (NIC) are sufficient.

EtherCAT makes Ethernet down to the I/O level technically feasible and economically sensible. Full Ethernet compatibility, internet technologies even in very simple devices, maximum utilisation of the large bandwidth offered by Ethernet, outstanding real-time characteristics at low costs are outstanding features of this network. As a high-speed drive and I/O bus for Industrial PCs or in combination with smaller control units, EtherCAT is expected to be widely used in a wide range of applications.

9. Literature

- [1] EtherCAT Technology Group, <http://www.ethercat.org>
- [2] Beckhoff GmbH, <http://www.beckhoff.com>
- [3] IEEE 802.3-2000: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications.
- [4] IEEE 802.3ae-2002: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications; Amendment: Media Access Control (MAC) Parameters, Physical Layers, and Management Parameters for 10 Gb/s Operation.
- [5] ANSI/TIA/EIA-644-A, Electrical Characteristics of Low Voltage Differential Signaling (LVDS) Interface Circuits
- [6] IEEE 1588-2002: IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems
- [7] Janssen, Dr. Dirk, Büttner, Holger, "EtherCAT – the Ethernet Fieldbus". PC Control Magazine 3/2003.
- [8] EtherCAT Communication Specification, EtherCAT Technology Group 2004
- [9] EN 50325-4: Industrial communications subsystem based on ISO 11898 (CAN) for controller-device interfaces. Part 4: CANopen.
- [10] IEC 61491-2002: Electrical equipment of industrial machines - Serial data link for real-time communication between controls and drives