
IgH Ether**CAT**[®] Master 1.4.0 Preliminary Documentation

Dipl.-Ing. (FH) Florian Pose, fp@igh-essen.com
Ingenieurgesellschaft **IGH**

Essen, February 4, 2009
Revision 1656

Contents

Conventions	viii
1 The IgH EtherCAT Master	1
1.1 Feature Summary	1
1.2 License	3
2 Architecture	5
2.1 Master Module	7
2.2 Master Phases	8
2.3 Process Data	8
3 Application Interface	11
3.1 Master Configuration	11
3.1.1 Slave Configuration	11
3.2 Cyclic Operation	13
3.3 Concurrent Master Access	13
4 Ethernet Devices	15
4.1 Network Driver Basics	15
4.2 EtherCAT Device Drivers	17
4.3 Device Selection	19
4.4 EtherCAT Device Interface	19
4.5 Patching Network Drivers	19
5 State Machines	21
5.1 State Machine Theory	22
5.2 The Master's State Model	24
5.3 The Master State Machine	27
5.4 The Slave Scan State Machine	27
5.5 The Slave Configuration State Machine	30
5.6 The State Change State Machine	30
5.7 The SII State Machine	33
5.8 The PDO State Machines	34
6 Mailbox Protocol Implementations	39
6.1 Ethernet over EtherCAT (EoE)	39
6.2 CANopen over EtherCAT (CoE)	41

7	Userspace Interfaces	43
7.1	Command-line Tool	43
7.1.1	Character Devices	43
7.1.2	Setting Alias Addresses	43
7.1.3	Displaying the Bus Configuration	44
7.1.4	Displaying Process Data	45
7.1.5	Setting a Master's Debug Level	45
7.1.6	Configured Domains	45
7.1.7	Master and Ethernet Devices	46
7.1.8	Sync Managers, PDOs and PDO Entries	46
7.1.9	SDO Dictionary	47
7.1.10	SDO Access	48
7.1.11	Slaves on the Bus	49
7.1.12	SII Access	50
7.1.13	Requesting Application-Layer States	52
7.1.14	Generating Slave Description XML	52
7.2	System Integration	53
7.2.1	Init Script	53
7.2.2	Sysconfig File	53
7.2.3	Starting the Master as a Service	54
7.3	Debug Interfaces	55
8	Timing Aspects	57
8.1	Application Interface Profiling	57
8.2	Bus Cycle Measurement	58
9	Installation	61
9.1	Building the Software	61
9.2	Building the Interface Documentation	61
9.3	Installing the Software	61
9.4	Automatic Device Node Creation	64
	Bibliography	65
	Glossary	67
	Index	69

List of Tables

3.1	Specifying a Slave Position	12
5.1	A typical state transition table	23
8.1	Application Cycle on a 2.0 GHz Processor	57
9.1	Configuration options	62

List of Figures

2.1	Master Architecture	6
2.2	Multiple masters in one module	7
2.3	Master phases and transitions	8
2.4	FMMU Configuration	10
3.1	Master Configuration	12
3.2	Slave Configuration Attachment	13
3.3	Concurrent Master Access	14
4.1	Interrupt Operation versus Interrupt-less Operation	18
5.1	A typical state transition diagram	23
5.2	Transition diagram of the master state machine	28
5.3	Transition diagram of the slave scan state machine	29
5.4	Transition diagram of the slave configuration state machine	31
5.5	Transition Diagram of the State Change State Machine	32
5.6	Transition Diagram of the SII State Machine	33
5.7	Transition Diagram of the PDO Reading State Machine	35
5.8	Transition Diagram of the PDO Entry Reading State Machine	35
5.9	Transition Diagram of the PDO Configuration State Machine	36
5.10	Transition Diagram of the PDO Entry Configuration State Machine	37
6.1	Transition Diagram of the EoE State Machine	40
6.2	Transition diagram of the CoE download state machine	42

Conventions

The following typographic conventions are used:

- *Italic face* is used for newly introduced terms and file names.
- **Typewriter face** is used for code examples and command line output.
- **Bold typewriter face** is used for user input in command lines.

Data values and addresses are usually specified as hexadecimal values. These are marked in the *C* programming language style with the prefix `0x` (example: `0x88A4`). Unless otherwise noted, address values are specified as byte addresses.

Function names are always printed with parentheses, but without parameters. So, if a function `ecrt_request_master()` has empty parentheses, this shall not imply that it has no parameters.

If shell commands have to be entered, this is marked by a dollar prompt:

\$

Further, if a shell command has to be entered as the superuser, the prompt is a mesh:

#

1 The IgH EtherCAT Master

This chapter covers some general information about the EtherCAT master.

1.1 Feature Summary

The list below gives a short summary of the master features.

- EtherCAT master implementation conforming to IEC/PAS 62407 [2] [3].
 - Runs as kernel module for Linux 2.6.
 - Multiple masters possible on one machine.
- EtherCAT-capable versions of standard Linux drivers for wide-spread Ethernet devices.
 - The Ethernet hardware is operated without interrupts.
 - Drivers for additional Ethernet hardware can easily be implemented using the common device interface (see sec. 4.4) provided by the master module.
 - Operation possible with any device supported by the standard drivers, including PCMCIA devices.
- The master code supports any Linux realtime extension through its independent architecture.
 - RTAI, Xenomai, etc.
 - Operation possible without any realtime extension at all.
- Common “Application Interface” for kernel-space realtime applications (see chap. 3).
 - Requesting and releasing masters.
 - Dynamic slave configuration, even for slaves that are offline.
 - Detailed configuration of the slaves’ PDOs and SDOs.
 - Creation of process data domains (see below). Registration of PDO entries for exchange within a domain.
 - Monitoring the states of masters, slave configurations and domains.
 - SDO handlers for application-triggered CoE transfers (see below).
 - Avoidance of unnecessary copy operations for process data.

- *Domains* are introduced, to allow grouping of process data transfers with different slave groups and task periods.
 - Management of PDO groups with different sample rates.
 - Automatic calculation of process data mapping, FMMU and sync manager configuration within the domains.
 - Process data exchange can be monitored via a per-domain mechanism.
- Master finite state machine (FSM).
 - The same state machine runs both in idle mode and in realtime operation.
 - Bus monitoring: Slave states are read cyclically. Automatic scanning of the bus after a topology change.
 - Automatic configuration of slaves, if a application-layer state change is requested.
- Implementation of the CANopen over EtherCAT (CoE) mailbox protocol.
 - Configuration of CoE-capable slaves.
 - SDO information service (dictionary listing).
 - SDO transfers both via the application interface and the command-line tool.
- Implementation of the Ethernet over EtherCAT (EoE) mailbox protocol.
 - Virtual network interface for any EoE-capable slave.
 - Both a switched and a routed EoE network architecture is natively supported and configurable with standard tools.
- Userspace command-line-tool “ethercat” (see sec. 7.1)
 - Detailed information about master, slaves, domains and bus configuration.
 - Reading/Writing alias addresses.
 - Listing slave configurations.
 - Viewing process data.
 - SDO download/upload; listing SDO dictionaries.
 - Slave SII (EEPROM) access.
 - Controlling application-layer states.
 - Generation of slave description XML from existing slaves.
- Seamless integration in any GNU/Linux distribution.
 - “Linux Standard Base”-compatible init script for master control.
 - Master and Ethernet device configuration via sysconfig file.
- Virtual read-only network interface for debugging and traffic monitoring purposes (using Wireshark [6], etc.). No additional hardware necessary.

1.2 License

The master code is released under the terms and conditions of the GNU General Public License (GPL [\[4\]](#)), version 2. Other developers, that want to use EtherCAT with Linux systems, are invited to use the master code or even participate on development.

2 Architecture

The EtherCAT master is integrated into the Linux 2.6 kernel. This was an early design decision, which has been made for several reasons:

- Kernel code has significantly better realtime characteristics, i. e. less latency than userspace code. It was foreseeable, that a fieldbus master has a lot of cyclic work to do. Cyclic work is usually triggered by timer interrupts inside the kernel. The execution delay of a function that processes timer interrupts is less, when it resides in kernelspace, because there is no need of time-consuming context switches to a userspace process.
- It was also foreseeable, that the master code has to directly communicate with the Ethernet hardware. This has to be done in the kernel anyway (through network device drivers), which is one more reason for the master code being in kernelspace.

Figure 2.1 gives a general overview of the master architecture.

The components of the master environment are described below:

Master Module Kernel module containing one or more EtherCAT master instances (see sec. 2.1), the “Device Interface” (see sec. 4.4) and the “Application Interface” (see chap. 3).

Device Modules EtherCAT-capable Ethernet device driver modules, that offer their devices to the EtherCAT master via the device interface (see sec. 4.4). These modified network drivers can handle network devices used for EtherCAT operation and “normal” Ethernet devices in parallel. A master can accept a certain device and then is able to send and receive EtherCAT frames. Ethernet devices declined by the master module are connected to the kernel’s network stack as usual.

Application Modules A kernel module that uses the EtherCAT master (usually for cyclic exchange of process data with EtherCAT slaves). These modules are not part of the EtherCAT master code¹, but have to be generated or written by the user. An application module can “request” a master through the application interface (see chap. 3). If this succeeds, it has the control over the master: It can provide a bus configuration and exchange process data.

¹Although there are some examples provided in the *examples/* directory.

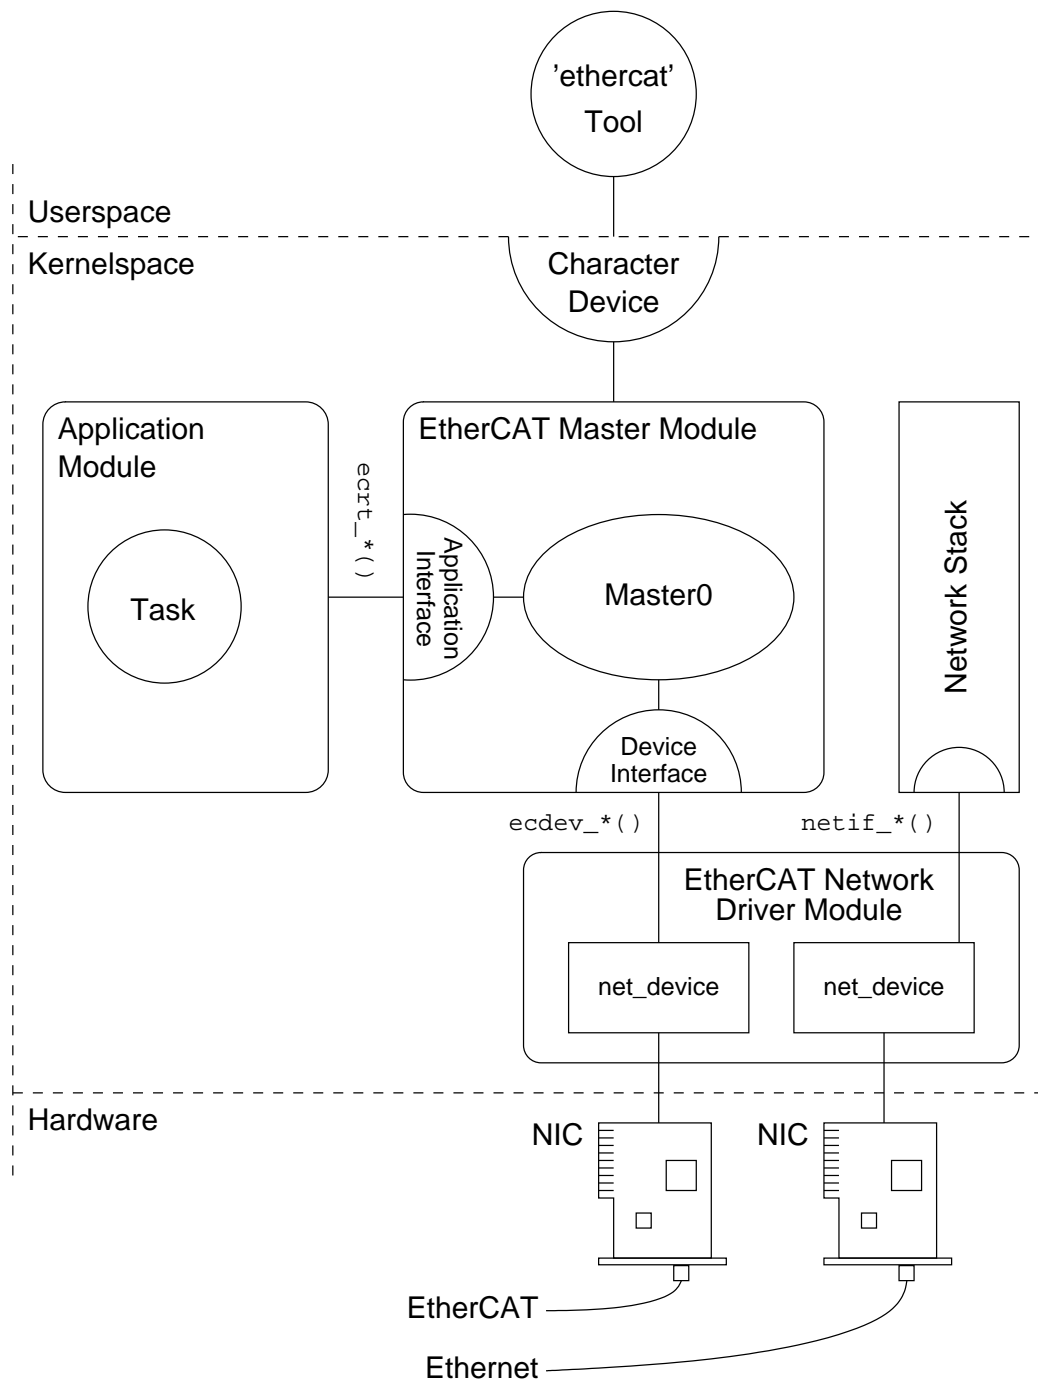


Figure 2.1: Master Architecture

2.1 Master Module

The EtherCAT master kernel module *ec_master* can contain multiple master instances. Each master waits for a certain Ethernet device identified by its MAC address. These addresses have to be specified on module loading via the *main_devices* module parameter. The number of master instances to initialize is taken from the number of MAC addresses given.

The below command loads the master module with a single master instance that waits for the Ethernet device with the MAC address 00:0E:0C:DA:A2:20. The master will be accessible via index 0.

```
# modprobe ec_master main_devices=00:0E:0C:DA:A2:20
```

MAC addresses for multiple masters have to be separated by commas:

```
# modprobe ec_master main_devices=00:0E:0C:DA:A2:20,00:e0:81:71:d5:1c
```

The two masters can be addressed by their indices 0 and 1 respectively (see figure 2.2). The master index is needed for the `ecrt_master_request()` function of the application interface (see chap. 3) and the `--master` option of the *ethercat* command-line tool (see sec. 7.1), which defaults to 0.

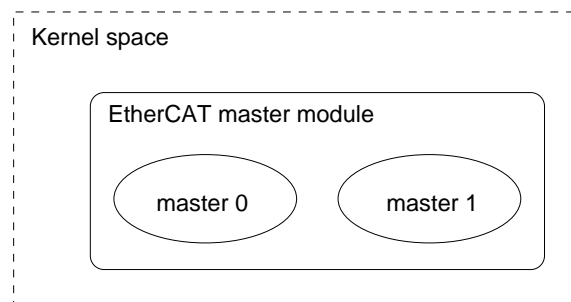


Figure 2.2: Multiple masters in one module

Init Script In most cases it is not necessary to load the master module and the Ethernet driver modules manually. There is an init script available, so the master can be started as a service (see sec. 7.2).

Syslog The master module outputs information about its state and events to the kernel ring buffer. These also end up in the system logs. The above module loading command should result in the messages below:

```
# dmesg | tail -2
EtherCAT: Master driver 1.4.0
EtherCAT: 2 masters waiting for devices.
```

```
# tail -2 /var/log/messages
Jul  4 10:22:45 ethercat kernel: EtherCAT: Master driver 1.4.0
Jul  4 10:22:45 ethercat kernel: EtherCAT: 2 masters waiting
                                for devices.
```

All EtherCAT master output is prefixed with `EtherCAT` which makes searching the logs easier.

2.2 Master Phases

Every EtherCAT master provided by the master module (see sec. 2.1) runs through several phases (see fig. 2.3):




Figure 2.3: Master phases and transitions

Orphaned phase This mode takes effect, when the master still waits for its Ethernet device to connect. No bus communication is possible until then.

Idle phase takes effect when the master has accepted an Ethernet device, but is not requested by any application yet. The master runs its state machine (see sec. 5.3), that automatically scans the bus for slaves and executes pending operations from the userspace interface (for example SDO access). The command-line tool can be used to access the bus, but there is no process data exchange because of the missing bus configuration.

Operation phase The master is requested by an application that can provide a bus configuration and exchange process data.

2.3 Process Data

This section shall introduce a few terms and ideas how the master handles process data.

Process Data Image Slaves offer their inputs and outputs by presenting the master so-called “Process Data Objects” (PDOs). The available PDOs can be either determined by reading out the slave’s TXPDO and RXPDO SII categories from the E²PROM (in case of fixed PDOs) or by reading out the appropriate CoE objects (see sec. 6.2), if available. The application can register the PDOs’ entries for exchange during cyclic operation. The sum of all registered PDO entries defines the “process data image”, which is exchanged via datagrams with “logical” memory access (like LWR, LRD or LRW) introduced in [2, sec. 5.4].

Process Data Domains The process data image can be easily managed by creating so-called “domains”, which allow grouped PDO exchange. They also take care of managing the datagram structures needed to exchange the PDOs. Domains are mandatory for process data exchange, so there has to be at least one. They were introduced for the following reasons:

- The maximum size of a datagram is limited due to the limited size of an Ethernet frame: The maximum data size is the Ethernet data field size minus the EtherCAT frame header, EtherCAT datagram header and EtherCAT datagram footer: $1500 - 2 - 12 - 2 = 1484$ octets. If the size of the process data image exceeds this limit, multiple frames have to be sent, and the image has to be partitioned for the use of multiple datagrams. A domain manages this automatically.
- Not every PDO has to be exchanged with the same frequency: The values of PDOs can vary slowly over time (for example temperature values), so exchanging them with a high frequency would just waste bus bandwidth. For this reason, multiple domains can be created, to group different PDOs and so allow separate exchange.

There is no upper limit for the number of domains, but each domain occupies one FMMU in each slave involved, so the maximum number of domains is de facto limited by the slaves.

FMMU Configuration An application can register PDO entries for exchange. Every PDO entry and its parent PDO is part of a memory area in the slave’s physical memory, that is protected by a sync manager [2, sec. 6.7] for synchronized access. In order to make a sync manager react on a datagram accessing its memory, it is necessary to access the last byte covered by the sync manager. Otherwise the sync manager will not react on the datagram and no data will be exchanged. That is why the whole synchronized memory area has to be included into the process data image: For example, if a certain PDO entry of a slave is registered for exchange with a certain domain, one FMMU will be configured to map the complete sync-manager-protected memory, the PDO entry resides in. If a second PDO entry of the same slave is registered for process data exchange within the same domain, and it resides in the same sync-manager-protected memory as the first one, the FMMU configuration is

not altered, because the desired memory is already part of the domain's process data image. If the second PDO entry would belong to another sync-manager-protected area, this complete area would also be included into the domains process data image.

Figure 2.4 gives an overview, how FMMUs are configured to map physical memory to logical process data images.

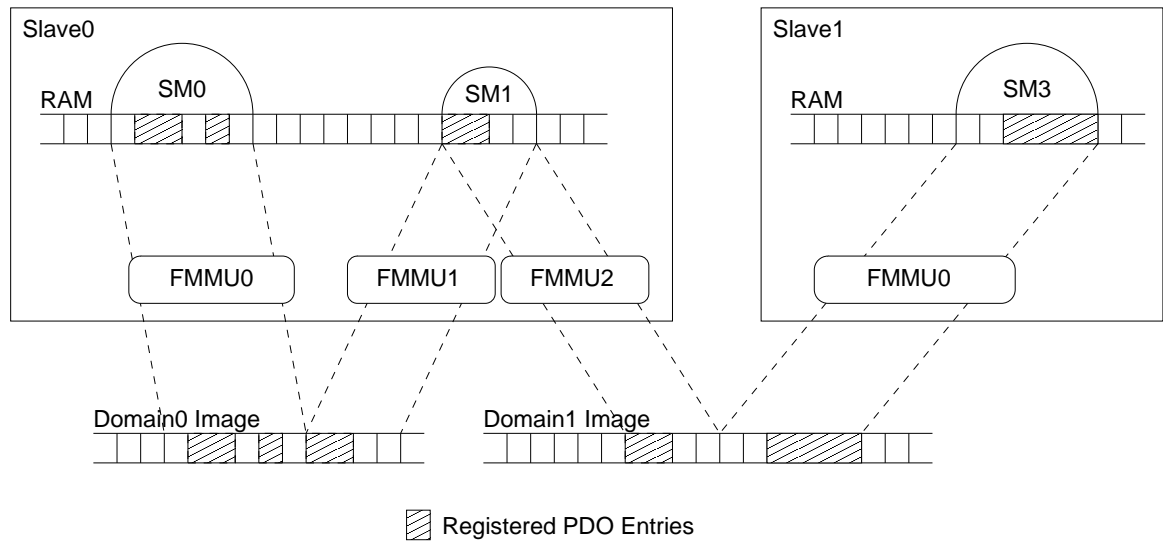


Figure 2.4: FMMU Configuration

3 Application Interface

The application interface provides functions and data structures for applications to access an EtherCAT master. The complete documentation of the interface is included as Doxygen [10] comments in the header file *include/ecrt.h*. It can either be read directly from the file comments, or as a more comfortable HTML documentation. The HTML generation is described in sec. 9.2.

The following sections cover a general description of the application interface.

Every application should use the master in two steps:

Configuration The master is requested and the configuration is applied. For example, domains are created, slaves are configured and PDO entries are registered (see sec. 3.1).

Operation Cyclic code is run and process data are exchanged (see sec. 3.2).

Example Applications There are a few example applications in the *examples/* sub-directory of the master code. They are documented in the source code.

3.1 Master Configuration

The bus configuration is supplied via the application interface. Figure 3.1 gives an overview of the objects, that can be configured by the application.

3.1.1 Slave Configuration

The application has to tell the master about the expected bus topology. This can be done by creating “slave configurations”. A slave configuration can be seen as an expected slave. When a slave configuration is created, the application provides the bus position (see below), vendor id and product code.

When the bus configuration is applied, the master checks, if there is a slave with the given vendor id and product code at the given position. If this is the case, the slave configuration is “attached” to the real slave on the bus and the slave is configured according to the settings provided by the application. The state of a slave configuration can either be queried via the application interface or via the command-line tool (see sec. 7.1.3).




Figure 3.1: Master Configuration

Slave Position The slave position has to be specified as a tuple of “alias” and “position”. This allows addressing slaves either via an absolute bus position, or a stored identifier called “alias”, or a mixture of both. The alias is a 16-bit value stored in the slave’s E²PROM. It can be modified via the command-line tool (see sec. 7.1.2). Table 3.1 shows, how the values are interpreted.

Table 3.1: Specifying a Slave Position

Alias	Position	Interpretation
0	0 – 65535	Position addressing. The position parameter is interpreted as the absolute ring position in the bus.
1 – 65535	0 – 65535	Alias addressing. The position parameter is interpreted as relative position after the first slave with the given alias address.

Figure 3.2 shows an example of how slave configurations are attached. Some of the configurations were attached, while others remain detached. The below lists gives the reasons beginning with the top slave configuration.

1. A zero alias means to use simple position addressing. Slave 1 exists and vendor id and product code match the expected values.
2. Although the slave with position 0 is found, the product code does not match, so the configuration is not attached.

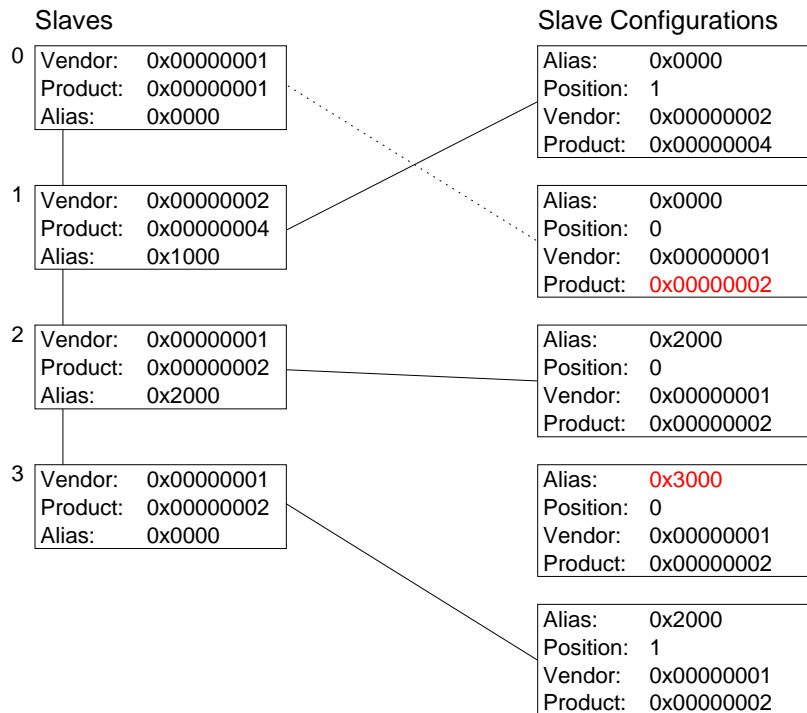


Figure 3.2: Slave Configuration Attachment

3. The alias is non-zero, so alias addressing is used. Slave 2 is the first slave with alias 0x2000. Because the position value is zero, the same slave is used.
4. There is no slave with the given alias, so the configuration can not be attached.
5. Slave 2 is again the first slave with the alias 0x2000, but position is now 1, so slave 3 is attached.

3.2 Cyclic Operation

To enter cyclic operation mode, the master has to be “activated” to calculate the process data image and apply the bus configuration for the first time. After activation, the application is in charge to send and receive frames.

3.3 Concurrent Master Access

In some cases, one master is used by several instances, for example when an application does cyclic process data exchange, and there are EoE-capable slaves that require to exchange Ethernet data with the kernel (see sec. 6.1). For this reason, the master is a shared resource, and access to it has to be sequentialized. This is usually done by locking with semaphores, or other methods to protect critical sections.

The master itself can not provide locking mechanisms, because it has no chance to know the appropriate kind of lock. For example if the application module uses RTAI functionality, ordinary kernel semaphores would not be sufficient. For that, an important design decision was made: The application that reserved a master must have the total control, therefore it has to take responsibility for providing the appropriate locking mechanisms. If another instance wants to access the master, it has to request the master lock by callbacks, that have to be set by the application. Moreover the application can deny access to the master if it considers it to be awkward at the moment.




Figure 3.3: Concurrent Master Access

Figure 3.3 exemplary shows, how two processes share one master: The application's cyclic task uses the master for process data exchange, while the master-internal EoE process uses it to communicate with EoE-capable slaves. Both have to acquire the master lock before access: The application task can access the lock natively, while the EoE process has to use the callbacks. See the application interface documentation (chap. 3) for how to use the locking callbacks.

4 Ethernet Devices

The EtherCAT protocol is based on the Ethernet standard, so a master relies on standard Ethernet hardware to communicate with the bus.

The term *device* is used as a synonym for Ethernet network interface hardware. There are device driver modules that handle Ethernet hardware, which a master can use to connect to an EtherCAT bus.

4.1 Network Driver Basics

EtherCAT relies on Ethernet hardware and the master needs a physical Ethernet device to communicate with the bus. Therefore it is necessary to understand how Linux handles network devices and their drivers, respectively.

Tasks of a Network Driver Network device drivers usually handle the lower two layers of the OSI model, that is the physical layer and the data-link layer. A network device itself natively handles the physical layer issues: It represents the hardware to connect to the medium and to send and receive data in the way, the physical layer protocol describes. The network device driver is responsible for getting data from the kernel's networking stack and forwarding it to the hardware, that does the physical transmission. If data is received by the hardware respectively, the driver is notified (usually by means of an interrupt) and has to read the data from the hardware memory and forward it to the network stack. There are a few more tasks, a network device driver has to handle, including queue control, statistics and device dependent features.

Driver Startup Usually, a driver searches for compatible devices on module loading. For PCI drivers, this is done by scanning the PCI bus and checking for known device IDs. If a device is found, data structures are allocated and the device is taken into operation.

Interrupt Operation A network device usually provides a hardware interrupt that is used to notify the driver of received frames and success of transmission, or errors, respectively. The driver has to register an interrupt service routine (ISR), that is executed each time, the hardware signals such an event. If the interrupt was thrown by the own device (multiple devices can share one hardware interrupt), the reason

for the interrupt has to be determined by reading the device's interrupt register. For example, if the flag for received frames is set, frame data has to be copied from hardware to kernel memory and passed to the network stack.

The `net_device` Structure The driver registers a `net_device` structure for each device to communicate with the network stack and to create a “network interface”. In case of an Ethernet driver, this interface appears as `ethX`, where X is a number assigned by the kernel on registration. The `net_device` structure receives events (either from userspace or from the network stack) via several callbacks, which have to be set before registration. Not every callback is mandatory, but for reasonable operation the ones below are needed in any case:

`open()` This function is called when network communication has to be started, for example after a command `ip link set ethX up` from userspace. Frame reception has to be enabled by the driver.

`stop()` The purpose of this callback is to “close” the device, i. e. make the hardware stop receiving frames.

`hard_start_xmit()` This function is called for each frame that has to be transmitted. The network stack passes the frame as a pointer to an `sk_buff` structure (“socket buffer”, see below), which has to be freed after sending.

`get_stats()` This call has to return a pointer to the device's `net_device_stats` structure, which permanently has to be filled with frame statistics. This means, that every time a frame is received, sent, or an error happened, the appropriate counter in this structure has to be increased.

The actual registration is done with the `register_netdev()` call, unregistering is done with `unregister_netdev()`.

The `netif` Interface All other communication in the direction interface → network stack is done via the `netif_*()` calls. For example, on successful device opening, the network stack has to be notified, that it can now pass frames to the interface. This is done by calling `netif_start_queue()`. After this call, the `hard_start_xmit()` callback can be called by the network stack. Furthermore a network driver usually manages a frame transmission queue. If this gets filled up, the network stack has to be told to stop passing further frames for a while. This happens with a call to `netif_stop_queue()`. If some frames have been sent, and there is enough space again to queue new frames, this can be notified with `netif_wake_queue()`. Another important call is `netif_receive_skb()`¹: It passes a frame to the network stack, that was just received by the device. Frame data has to be included in a so-called “socket buffer” for that (see below).

¹This function is part of the NAPI (“New API”), that replaces the kernel 2.4 technique for interfacing to the network stack (with `netif_rx()`). NAPI is a technique to improve network performance on Linux. Read more in <http://www.cyberus.ca/~hadi/usenix-paper.tgz>.

Socket Buffers Socket buffers are the basic data type for the whole network stack. They serve as containers for network data and are able to quickly add data headers and footers, or strip them off again. Therefore a socket buffer consists of an allocated buffer and several pointers that mark beginning of the buffer (`head`), beginning of data (`data`), end of data (`tail`) and end of buffer (`end`). In addition, a socket buffer holds network header information and (in case of received data) a pointer to the `net_device`, it was received on. There exist functions that create a socket buffer (`dev_alloc_skb()`), add data either from front (`skb_push()`) or back (`skb_put()`), remove data from front (`skb_pull()`) or back (`skb_trim()`), or delete the buffer (`kfree_skb()`). A socket buffer is passed from layer to layer, and is freed by the layer that uses it the last time. In case of sending, freeing has to be done by the network driver.

4.2 EtherCAT Device Drivers

There are a few requirements for Ethernet network devices to function as EtherCAT devices, when connected to an EtherCAT bus.

Dedicated Interfaces For performance and realtime purposes, the EtherCAT master needs direct and exclusive access to the Ethernet hardware. This implies that the network device must not be connected to the kernel's network stack as usual, because the kernel would try to use it as an ordinary Ethernet device.

Interrupt-less Operation EtherCAT frames travel through the logical EtherCAT ring and are then sent back to the master. Communication is highly deterministic: A frame is sent and will be received again after a constant time, so there is no need to notify the driver about frame reception: The master can instead query the hardware for received frames, if it expects them to be already received.

Figure 4.1 shows two workflows for cyclic frame transmission and reception with and without interrupts.

In the left workflow "Interrupt Operation", the data from the last cycle is first processed and a new frame is assembled with new datagrams, which is then sent. The cyclic work is done for now. Later, when the frame is received again by the hardware, an interrupt is triggered and the ISR is executed. The ISR will fetch the frame data from the hardware and initiate the frame dissection: The datagrams will be processed, so that the data is ready for processing in the next cycle.

In the right workflow "Interrupt-less Operation", there is no hardware interrupt enabled. Instead, the hardware will be polled by the master by executing the ISR. If the frame has been received in the meantime, it will be dissected. The situation is now the same as at the beginning of the left workflow: The received data is processed and a new frame is assembled and sent. There is nothing to do for the rest of the cycle.

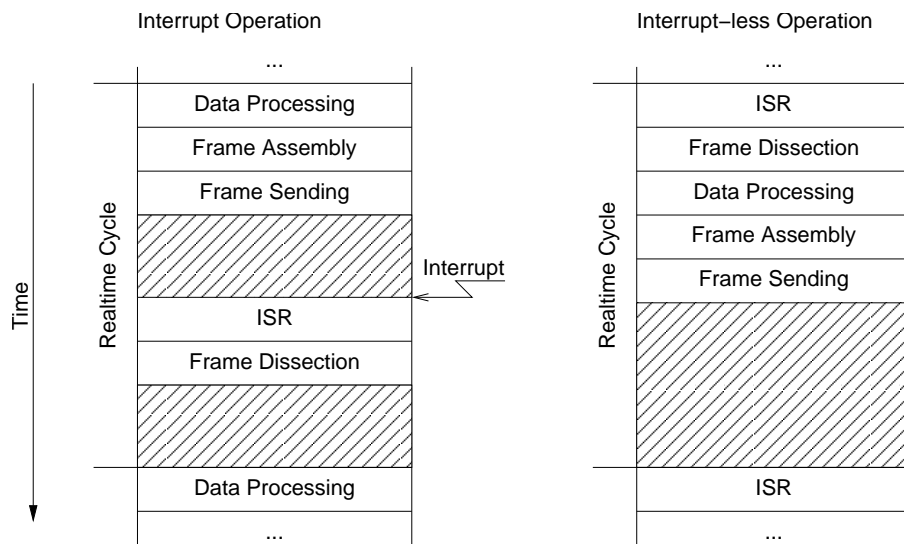


Figure 4.1: Interrupt Operation versus Interrupt-less Operation

The interrupt-less operation is desirable, because hardware interrupts are not conducive in improving the driver's realtime behaviour: Their indeterministic incidences contribute to increasing the jitter. Besides, if a realtime extension (like RTAI) is used, some additional effort would have to be made to prioritize interrupts.

Ethernet and EtherCAT Devices Another issue lies in the way Linux handles devices of the same type. For example, a PCI driver scans the PCI bus for devices it can handle. Then it registers itself as the responsible driver for all of the devices found. The problem is, that an unmodified driver can not be told to ignore a device because it will be used for EtherCAT later. There must be a way to handle multiple devices of the same type, where one is reserved for EtherCAT, while the other is treated as an ordinary Ethernet device.

For all this reasons, the author decided that the only acceptable solution is to modify standard Ethernet drivers in a way that they keep their normal functionality, but gain the ability to treat one or more of the devices as EtherCAT-capable.

Below are the advantages of this solution:

- No need to tell the standard drivers to ignore certain devices.
- One networking driver for EtherCAT and non-EtherCAT devices.
- No need to implement a network driver from scratch and running into issues, the former developers already solved.

The chosen approach has the following disadvantages:

- The modified driver gets more complicated, as it must handle EtherCAT and non-EtherCAT devices.

- Many additional case differentiations in the driver code.
- Changes and bug fixes on the standard drivers have to be ported to the EtherCAT-capable versions from time to time.

4.3 Device Selection

After loading the master module, at least one EtherCAT-capable network driver module has to be loaded, that offers its devices to the master (see sec. 4.4). The master module knows the devices to choose from the module parameters (see sec. 2.1). If the init script is used to start the master, the drivers and devices to use can be specified in the sysconfig file (see sec. 7.2.2).

4.4 EtherCAT Device Interface

An anticipation to the section about the master module (sec. 2.1) has to be made in order to understand the way, a network device driver module can connect a device to a specific EtherCAT master.

The master module provides a “device interface” for network device drivers. To use this interface, a network device driver module must include the header *devices/ecdev.h*, coming with the EtherCAT master code. This header offers a function interface for EtherCAT devices. All functions of the device interface are named with the prefix *ecdev*.

The documentation of the device interface can be found in the header file or in the appropriate module of the interface documentation (see sec. 9.2 for generation instructions).

4.5 Patching Network Drivers

This section will describe, how to make a standard Ethernet driver EtherCAT-capable. Unfortunately, there is no standard procedure to enable an Ethernet driver for use with the EtherCAT master, but there are a few common techniques.

1. A first simple rule is, that `netif_*()` calls must be avoided for all EtherCAT devices. As mentioned before, EtherCAT devices have no connection to the network stack, and therefore must not call its interface functions.
2. Another important thing is, that EtherCAT devices should be operated without interrupts. So any calls of registering interrupt handlers and enabling interrupts at hardware level must be avoided, too.

3. The master does not use a new socket buffer for each send operation: Instead there is a fix one allocated on master initialization. This socket buffer is filled with an EtherCAT frame with every send operation and passed to the `hard_start_xmit()` callback. For that it is necessary, that the socket buffer is not be freed by the network driver as usual.

An Ethernet driver usually handles several Ethernet devices, each described by a `net_device` structure with a `priv_data` field to attach driver-dependent data to the structure. To distinguish between normal Ethernet devices and the ones used by EtherCAT masters, the private data structure used by the driver could be extended by a pointer, that points to an `ec_device_t` object returned by `ecdev_offer()` (see sec. 4.4) if the device is used by a master and otherwise is zero.

The RealTek RTL-8139 Fast Ethernet driver is a “simple” Ethernet driver and can be taken as an example to patch new drivers. The interesting sections can be found by searching the string “ecdev” in the file `devices/8139too-2.6.24-ethercat.c`.

5 State Machines

Many parts of the EtherCAT master are implemented as *finite state machines* (FSMs). Though this leads to a higher grade of complexity in some aspects, it opens many new possibilities.

The below short code example exemplary shows how to read all slave states and moreover illustrates the restrictions of “sequential” coding:

```
1 ec_datagram_brd(datagram, 0x0130, 2); // prepare datagram
2 if (ec_master_simple_io(master, datagram)) return -1;
3 slave_states = EC_READ_U8(datagram->data); // process datagram
```

The `ec_master_simple_io()` function provides a simple interface for synchronously sending a single datagram and receiving the result¹. Internally, it queues the specified datagram, invokes the `ec_master_send_datagrams()` function to send a frame with the queued datagram and then waits actively for its reception.

This sequential approach is very simple, reflecting in only three lines of code. The disadvantage is, that the master is blocked for the time it waits for datagram reception. There is no difficulty when only one instance is using the master, but if more instances want to (synchronously²) use the master, it is inevitable to think about an alternative to the sequential model.

Master access has to be sequentialized for more than one instance wanting to send and receive datagrams synchronously. With the present approach, this would result in having one phase of active waiting for each instance, which would be non-acceptable especially in realtime circumstances, because of the huge time overhead.

A possible solution is, that all instances would be executed sequentially to queue their datagrams, then give the control to the next instance instead of waiting for the datagram reception. Finally, bus IO is done by a higher instance, which means that all queued datagrams are sent and received. The next step is to execute all instances again, which then process their received datagrams and issue new ones.

This approach results in all instances having to retain their state, when giving the control back to the higher instance. It is quite obvious to use a *finite state machine* model in this case. Section 5.1 will introduce some of the theory used, while the

¹For all communication issues have been meanwhile sourced out into state machines, the function is deprecated and stopped existing. Nevertheless it is adequate for showing its own restrictions.

²At this time, synchronous master access will be adequate to show the advantages of an FSM. The asynchronous approach will be discussed in sec. 6.1

listings below show the basic approach by coding the example from above as a state machine:

```
1 // state 1
2 ec_datagram_brd(datagram, 0x0130, 2); // prepare datagram
3 ec_master_queue(master, datagram); // queue datagram
4 next_state = state_2;
5 // state processing finished
```

After all instances executed their current state and queued their datagrams, these are sent and received. Then the respective next states are executed:

```
1 // state 2
2 if (datagram->state != EC_DGRAM_STATE_RECEIVED) {
3     next_state = state_error;
4     return; // state processing finished
5 }
6 slave_states = EC_READ_U8(datagram->data); // process datagram
7 // state processing finished.
```

See sec. 5.2 for an introduction to the state machine programming concept used in the master code.

5.1 State Machine Theory

A finite state machine [7] is a model of behavior with inputs and outputs, where the outputs not only depend on the inputs, but the history of inputs. The mathematical definition of a finite state machine (or finite automaton) is a six-tuple $(\Sigma, \Gamma, S, s_0, \delta, \omega)$, with

- the input alphabet Σ , with $\Sigma \neq \emptyset$, containing all input symbols,
- the output alphabet Γ , with $\Gamma \neq \emptyset$, containing all output symbols,
- the set of states S , with $S \neq \emptyset$,
- the set of initial states s_0 with $s_0 \subseteq S, s_0 \neq \emptyset$
- the transition function $\delta : S \times \Sigma \rightarrow S \times \Gamma$
- the output function ω .

The state transition function δ is often specified by a *state transition table*, or by a *state transition diagram*. The transition table offers a matrix view of the state machine behavior (see table 5.1). The matrix rows correspond to the states ($S = \{s_0, s_1, s_2\}$) and the columns correspond to the input symbols ($\Gamma = \{a, b, \varepsilon\}$). The table contents in a certain row i and column j then represent the next state (and possibly the output) for the case, that a certain input symbol σ_j is read in the state s_i .

Table 5.1: A typical state transition table

	a	b	ε
s_0	s_1	s_1	s_2
s_1	s_2	s_1	s_0
s_2	s_0	s_0	s_0

The state diagram for the same example looks like the one in figure 5.1. The states are represented as circles or ellipses and the transitions are drawn as arrows between them. Close to a transition arrow can be the condition that must be fulfilled to allow the transition. The initial state is marked by a filled black circle with an arrow pointing to the respective state.




Figure 5.1: A typical state transition diagram

Deterministic and non-deterministic state machines A state machine can be deterministic, meaning that for one state and input, there is one (and only one) following state. In this case, the state machine has exactly one starting state. Non-deterministic state machines can have more than one transitions for a single state-input combination. There is a set of starting states in the latter case.

Moore and Mealy machines There is a distinction between so-called *Moore machines*, and *Mealy machines*. Mathematically spoken, the distinction lies in the output function ω : If it only depends on the current state ($\omega : S \rightarrow \Gamma$), the machine corresponds to the “Moore Model”. Otherwise, if ω is a function of a state and the input alphabet ($\omega : S \times \Sigma \rightarrow \Gamma$) the state machine corresponds to the “Mealy model”. Mealy machines are the more practical solution in most cases, because their design allows machines with a minimum number of states. In practice, a mixture of both models is often used.

Misunderstandings about state machines There is a phenomenon called “state explosion”, that is often taken as a counter-argument against general use of state

machines in complex environments. It has to be mentioned, that this point is misleading [8]. State explosions happen usually as a result of a bad state machine design: Common mistakes are storing the present values of all inputs in a state, or not dividing a complex state machine into simpler sub state machines. The EtherCAT master uses several state machines, that are executed hierarchically and so serve as sub state machines. These are also described below.

5.2 The Master's State Model

This section will introduce the techniques used in the master to implement state machines.

State Machine Programming There are certain ways to implement a state machine in *C* code. An obvious way is to implement the different states and actions by one big case differentiation:

```
1 enum {STATE_1, STATE_2, STATE_3};
2 int state = STATE_1;
3
4 void state_machine_run(void *priv_data) {
5     switch (state) {
6         case STATE_1:
7             action_1();
8             state = STATE_2;
9             break;
10        case STATE_2:
11            action_2()
12            if (some_condition) state = STATE_1;
13            else state = STATE_3;
14            break;
15        case STATE_3:
16            action_3();
17            state = STATE_1;
18            break;
19    }
20 }
```

For small state machines, this is an option. The disadvantage is, that with an increasing number of states the code soon gets complex and an additional case differentiation is executed each run. Besides, lots of indentation is wasted.

The method used in the master is to implement every state in an own function and to store the current state function with a function pointer:

```
1 void (*state)(void *) = state1;
```



```

2
3 void state_machine_run(void *priv_data) {
4     state(priv_data);
5 }
6
7 void state1(void *priv_data) {
8     action_1();
9     state = state2;
10 }
11
12 void state2(void *priv_data) {
13     action_2();
14     if (some_condition) state = state1;
15     else state = state2;
16 }
17
18 void state3(void *priv_data) {
19     action_3();
20     state = state1;
21 }

```

In the master code, state pointers of all state machines³ are gathered in a single object of the `ec_fsm_master_t` class. This is advantageous, because there is always one instance of every state machine available and can be started on demand.

Mealy and Moore If a closer look is taken to the above listing, it can be seen that the actions executed (the “outputs” of the state machine) only depend on the current state. This accords to the “Moore” model introduced in sec. 5.1. As mentioned, the “Mealy” model offers a higher flexibility, which can be seen in the listing below:

```

1 void state7(void *priv_data) {
2     if (some_condition) {
3         action_7a();
4         state = state1;
5     }
6     else {
7         action_7b();
8         state = state8;
9     }
10 }

```

③ + ⑦ The state function executes the actions depending on the state transition, that is about to be done.

³All except for the EoE state machine, because multiple EoE slaves have to be handled in parallel. For this reason each EoE handler object has its own state pointer.

The most flexible alternative is to execute certain actions depending on the state, followed by some actions dependent on the state transition:

```
1 void state9(void *priv_data) {
2     action_9();
3     if (some_condition) {
4         action_9a();
5         state = state7;
6     }
7     else {
8         action_9b();
9         state = state10;
10    }
11 }
```

This model is often used in the master. It combines the best aspects of both approaches.

Using Sub State Machines To avoid having too much states, certain functions of the EtherCAT master state machine have been sourced out into sub state machines. This helps to encapsulate the related workflows and moreover avoids the “state explosion” phenomenon described in sec. 5.1. If the master would instead use one big state machine, the number of states would be a multiple of the actual number. This would increase the level of complexity to a non-manageable grade.

Executing Sub State Machines If a state machine starts to execute a sub state machine, it usually remains in one state until the sub state machine terminates. This is usually done like in the listing below, which is taken out of the slave configuration state machine code:

```
1 void ec_fsm_slaveconf_safeop(ec_fsm_t *fsm)
2 {
3     fsm->change_state(fsm); // execute state change
4                             // sub state machine
5
6     if (fsm->change_state == ec_fsm_error) {
7         fsm->slave_state = ec_fsm_end;
8         return;
9     }
10
11    if (fsm->change_state != ec_fsm_end) return;
12
13    // continue state processing
14    ...
```

- ③ `change_state` is the state pointer of the state change state machine. The state function, the pointer points on, is executed. . .
- ⑥ . . . either until the state machine terminates with the error state . . .
- ⑪ . . . or until the state machine terminates in the end state. Until then, the “higher” state machine remains in the current state and executes the sub state machine again in the next cycle.

State Machine Descriptions The below sections describe every state machine used in the EtherCAT master. The textual descriptions of the state machines contain references to the transitions in the corresponding state transition diagrams, that are marked with an arrow followed by the name of the successive state. Transitions caused by trivial error cases (i. e. no response from slave) are not described explicitly. These transitions are drawn as dashed arrows in the diagrams.

5.3 The Master State Machine

The master state machine is executed in the context of the master thread. Figure 5.2 shows its transition diagram. Its purposes are:

Bus monitoring The bus topology is monitored. If it changes, the bus is (re-)scanned.

Slave configuration The application-layer states of the slaves are monitored. If a slave is not in the state it supposed to be, the slave is (re-)configured.

Request handling Requests (either originating from the application or from external sources) are handled. A request is a job that the master shall process asynchronously, for example an SII access, SDO access, or similar.

5.4 The Slave Scan State Machine

The slave scan state machine, which can be seen in figure 5.3, leads through the process of reading desired slave information.

The scan process includes the following steps:

Node Address The node address is set for the slave, so that it can be node-addressed for all following operations.

AL State The initial application-layer state is read.

Base Information Base information (like the number of supported FMMUs) is read from the lower physical memory.

Data Link Information about the physical ports is read.

SII Size The size of the SII contents is determined to allocate SII image memory.

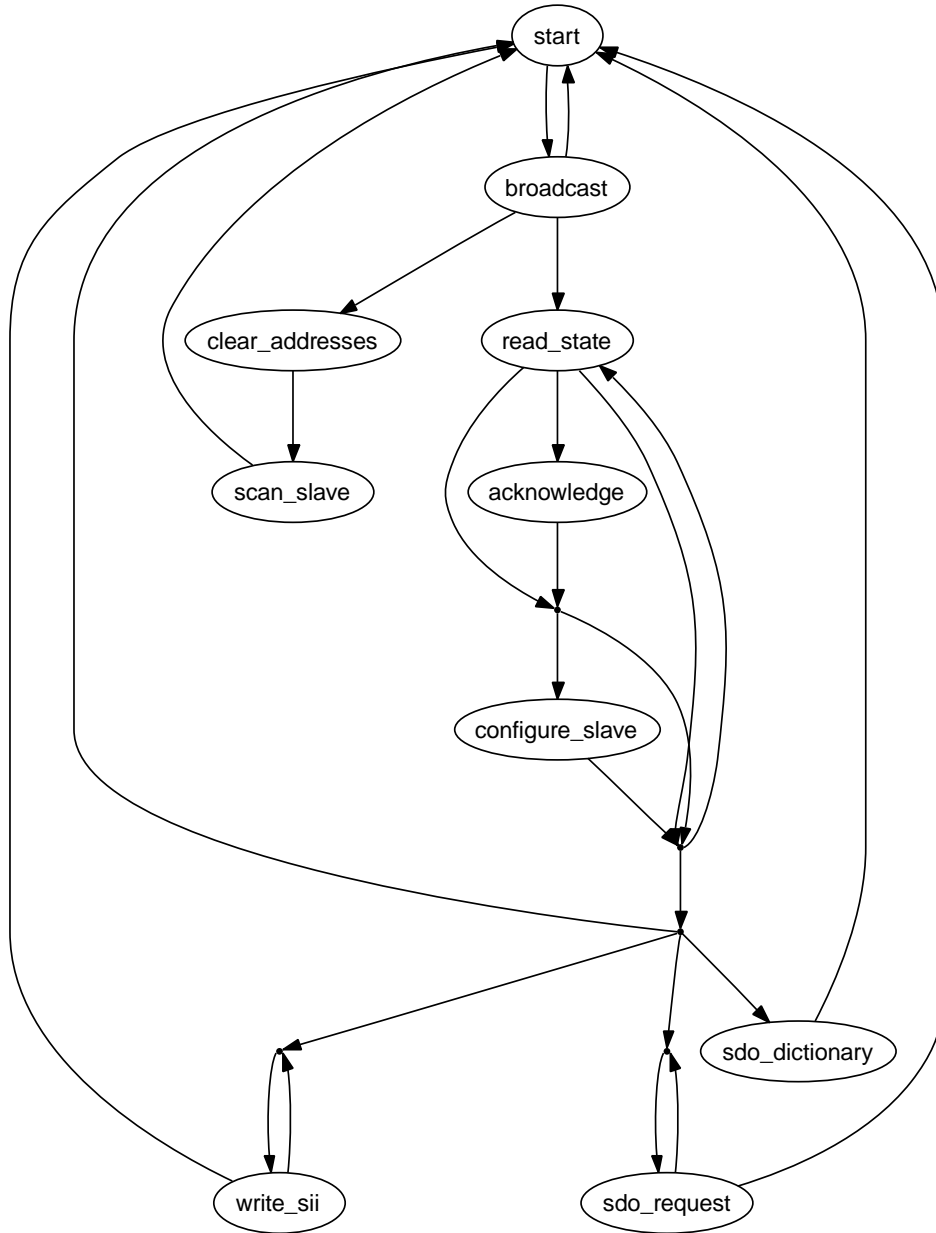


Figure 5.2: Transition diagram of the master state machine




Figure 5.3: Transition diagram of the slave scan state machine

SII Data The SII contents are read into the master's image.

PREOP If the slave supports CoE, it is set to PREOP state using the State change FSM (see sec. 5.6) to enable mailbox communication and read the PDO configuration via CoE.

PDOs The PDOs are read via CoE (if supported) using the PDO Reading FSM (see sec. 5.8). If this is successful, the PDO information from the SII (if any) is overwritten.

5.5 The Slave Configuration State Machine

The slave configuration state machine, which can be seen in figure 5.4, leads through the process of configuring a slave and bringing it to a certain application-layer state.

INIT The state change FSM is used to bring the slave to the INIT state.

FMMU Clearing To avoid that the slave reacts on any process data, the FMMU configuration are cleared. If the slave does not support FMMUs, this state is skipped. If INIT is the requested state, the state machine is finished.

Mailbox Sync Manager Configuration If the slaves support mailbox communication, the mailbox sync managers are configured. Otherwise this state is skipped.

PREOP The state change FSM is used to bring the slave to PREOP state. If this is the requested state, the state machine is finished.

SDO Configuration If there is a slave configuration attached (see sec. 3.1), and there are any SDO configurations are provided by the application, these are sent to the slave.

PDO Configuration The PDO configuration state machine is executed to apply all necessary PDO configurations.

PDO Sync Manager Configuration If any PDO sync managers exist, they are configured.

FMMU Configuration If there are FMMUs configurations supplied by the application (i. e. if the application registered PDO entries), they are applied.

SAFEOP The state change FSM is used to bring the slave to SAFEOP state. If this is the requested state, the state machine is finished.

OP The state change FSM is used to bring the slave to OP state. If this is the requested state, the state machine is finished.

5.6 The State Change State Machine

The state change state machine, which can be seen in figure 5.5, leads through the process of changing a slave's application-layer state. This implements the states and transitions described in [3, sec. 6.4.1].

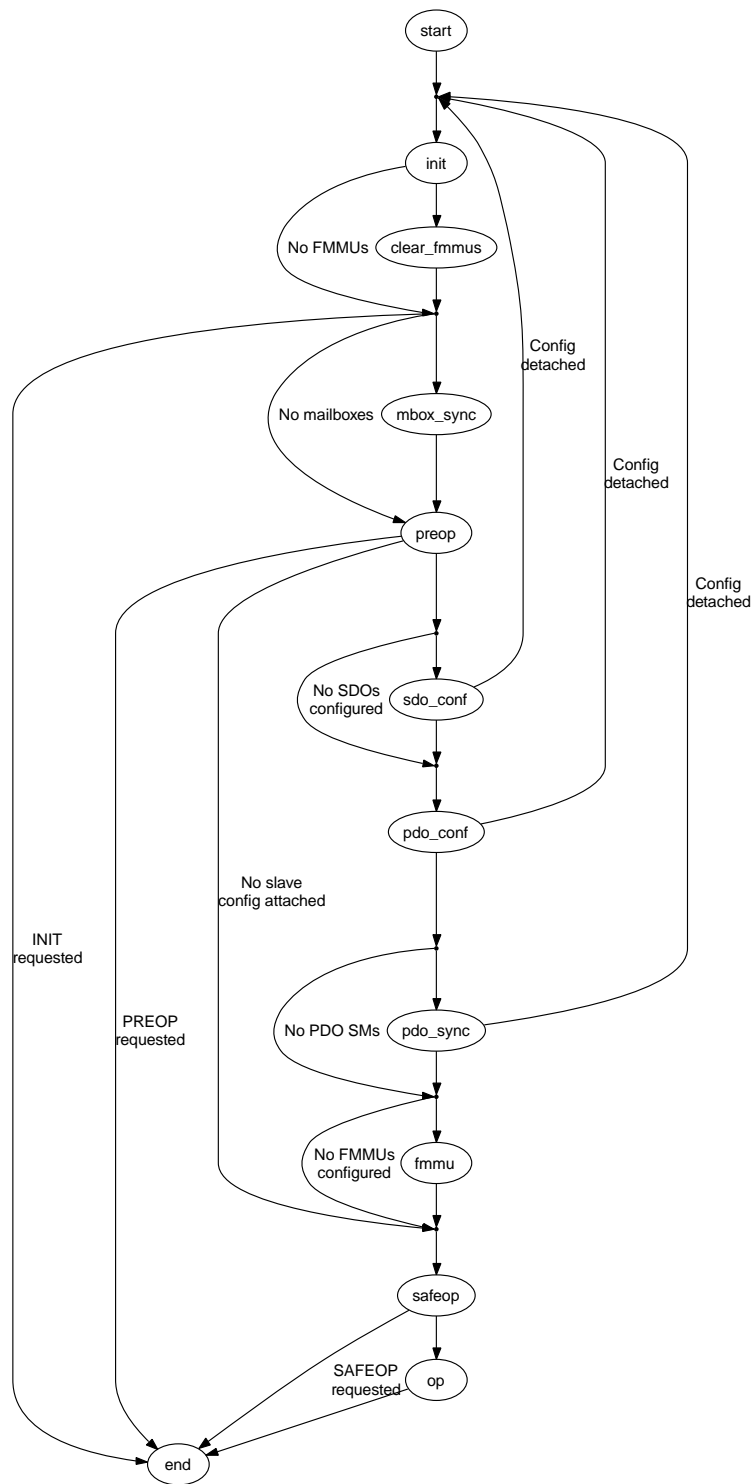


Figure 5.4: Transition diagram of the slave configuration state machine




Figure 5.5: Transition Diagram of the State Change State Machine

Start The new application-layer state is requested via the “AL Control Request” register (see [3, sec. 5.3.1]).

Check for Response Some slave need some time to respond to an AL state change command, and do not respond for some time. For this case, the command is issued again, until it is acknowledged.

Check AL Status If the AL State change datagram was acknowledged, the “AL Control Response” register (see [3, sec. 5.3.2]) must be read out until the slave changes the AL state.

AL Status Code If the slave refused the state change command, the reason can be read from the “AL Status Code” field in the “AL State Changed” registers (see [3, sec. 5.3.3]).

Acknowledge State If the state change was not successful, the master has to acknowledge the old state by writing to the “AL Control request” register again.

Check Acknowledge After sending the acknowledge command, it has to read out the “AL Control Response” register again.

The “start_ack” state is a shortcut in the state machine for the case, that the master wants to acknowledge a spontaneous AL state change, that was not requested.

5.7 The SII State Machine

The SII state machine (shown in figure 5.6) implements the process of reading or writing SII data via the Slave Information Interface described in [2, sec. 6.4].

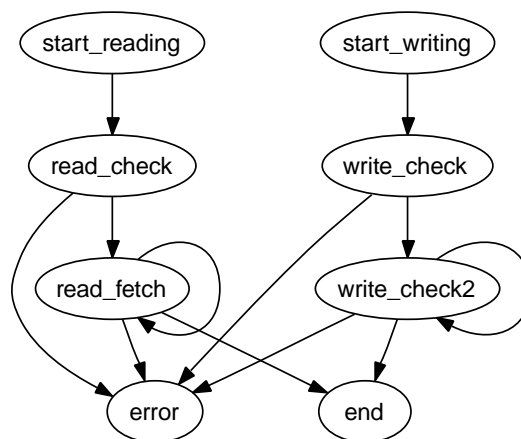


Figure 5.6: Transition Diagram of the SII State Machine

This is how the reading part of the state machine works:

Start Reading The read request and the requested word address are written to the SII attribute.

Check Read Command If the SII read request command has been acknowledged, a timer is started. A datagram is issued, that reads out the SII attribute for state and data.

Fetch Data If the read operation is still busy (the SII is usually implemented as an E²PROM), the state is read again. Otherwise the data are copied from the datagram.

The writing part works nearly similar:

Start Writing A write request, the target address and the data word are written to the SII attribute.

Check Write Command If the SII write request command has been acknowledged, a timer is started. A datagram is issued, that reads out the SII attribute for the state of the write operation.

Wait while Busy If the write operation is still busy (determined by a minimum wait time and the state of the busy flag), the state machine remains in this state to avoid that another write operation is issued too early.

5.8 The PDO State Machines

The PDO state machines are a set of state machines that read or write the PDO assignment and the PDO mapping via the “CoE Communication Area” described in [3, sec. 5.6.7.4]. For the object access, the CANopen over EtherCAT access primitives are used (see sec. 6.2), so the slave must support the CoE mailbox protocol.

PDO Reading FSM This state machine (fig. 5.7) has the purpose to read the complete PDO configuration of a slave. It reads the PDO assignment for each Sync Manager and uses the PDO Entry Reading FSM (fig. 5.8) to read the mapping for each assigned PDO.

Basically it reads the every Sync manager’s PDO assignment SDO’s (0x1C1x) number of elements to determine the number of assigned PDOs for this sync manager and then reads out the subindices of the SDO to get the assigned PDO’s indices. When a PDO index is read, the PDO Entry Reading FSM is executed to read the PDO’s mapped PDO entries.

PDO Entry Reading FSM This state machine (fig. 5.8) reads the PDO mapping (the PDO entries) of a PDO. It reads the respective mapping SDO (0x1600 – 0x17ff, or 0x1a00 – 0x1bff) for the given PDO by reading first the subindex zero (number of elements) to determine the number of mapped PDO entries. After that, each subindex is read to get the mapped PDO entry index, subindex and bit size.




Figure 5.7: Transition Diagram of the PDO Reading State Machine

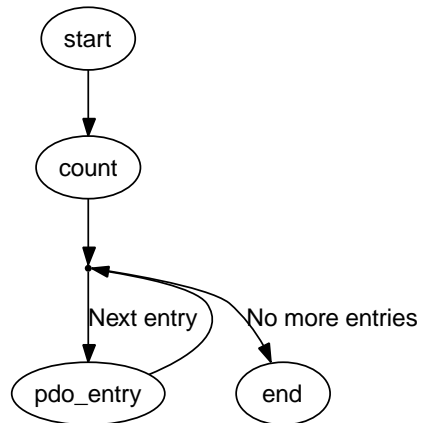


Figure 5.8: Transition Diagram of the PDO Entry Reading State Machine

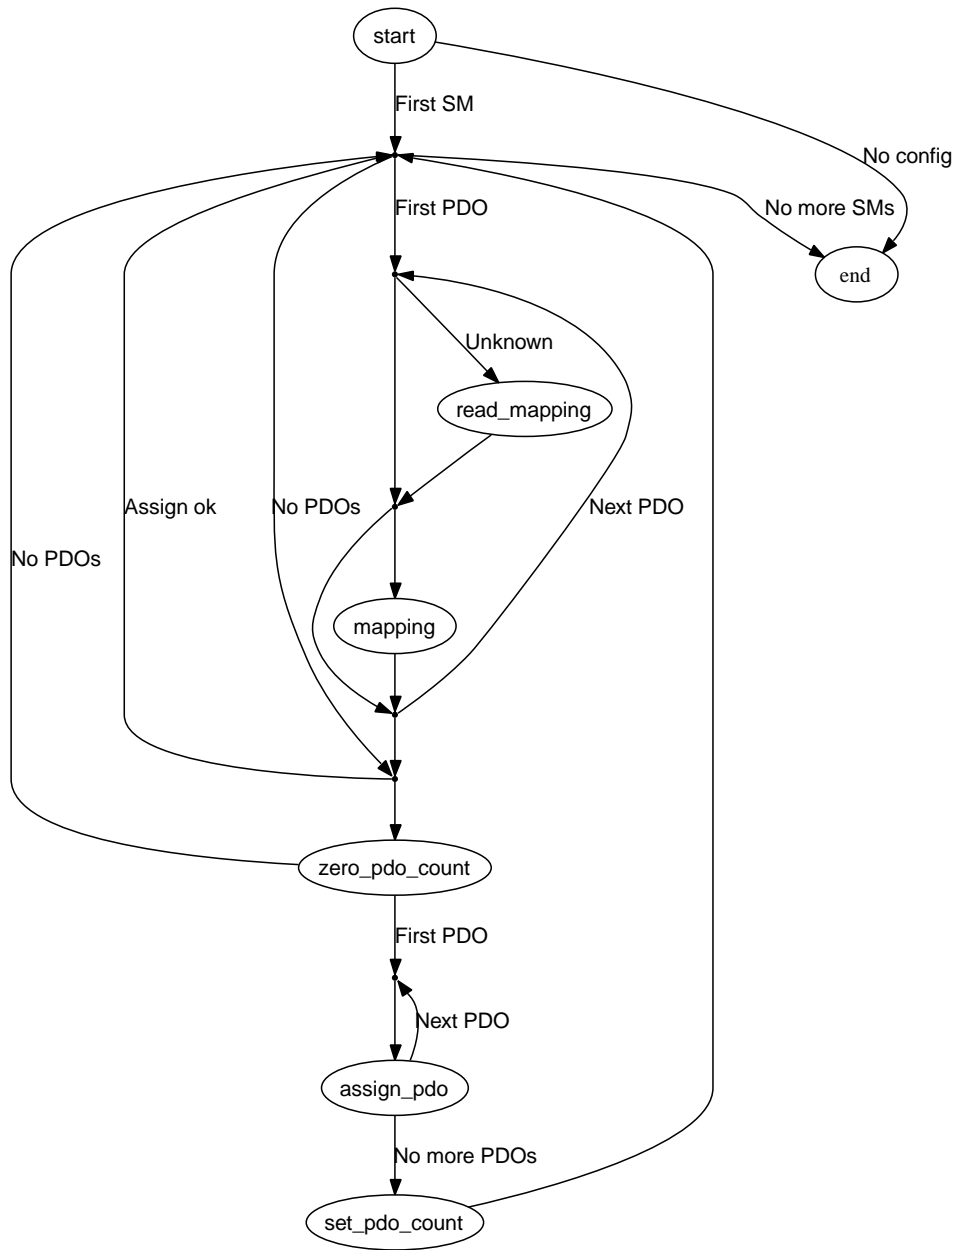


Figure 5.9: Transition Diagram of the PDO Configuration State Machine

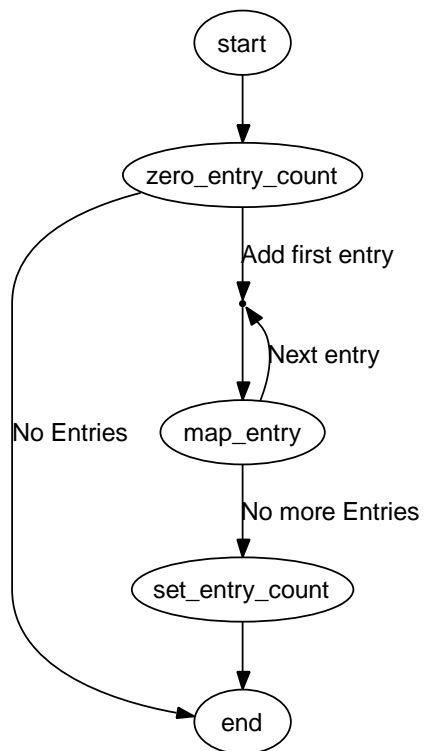


Figure 5.10: Transition Diagram of the PDO Entry Configuration State Machine

6 Mailbox Protocol Implementations

The EtherCAT master implements the EoE and the CoE mailbox protocols. See the below section for details.

6.1 Ethernet over EtherCAT (EoE)

The EtherCAT master implements the Ethernet over EtherCAT mailbox protocol [3, sec. 5.7] to enable the tunneling of Ethernet frames to special slaves, that can either have physical Ethernet ports to forward the frames to, or have an own IP stack to receive the frames.

Virtual Network Interfaces The master creates a virtual EoE network interface for every EoE-capable slave. These interfaces are called either

eoEXsY for a slave without an alias address (see sec. 7.1.2), where X is the master index and Y is the slave's ring position, or

eoEXaY for a slave with a non-zero alias address, where X is the master index and Y is the decimal alias address.

Frames sent to these interfaces are forwarded to the associated slaves by the master. Frames, that are received by the slaves, are fetched by the master and forwarded to the virtual interfaces.

This bears the following advantages:

- **Flexibility:** The user can decide, how the EoE-capable slaves are interconnected with the rest of the world.
- Standard tools can be used to monitor the EoE activity and to configure the EoE interfaces.
- The Linux kernel's layer-2-bridging implementation (according to the IEEE 802.1D MAC Bridging standard) can be used natively to bridge Ethernet traffic between EoE-capable slaves.
- The Linux kernel's network stack can be used to route packets between EoE-capable slaves and to track security issues, just like having physical network interfaces.

EoE Handlers The virtual EoE interfaces and the related functionality is encapsulated in the `ec_eoe_t` class. An object of this class is called “EoE handler”. For example the master does not create the network interfaces directly: This is done inside the constructor of an EoE handler. An EoE handler additionally contains a frame queue. Each time, the kernel passes a new socket buffer for sending via the interface’s `hard_start_xmit()` callback, the socket buffer is queued for transmission by the EoE state machine (see below). If the queue gets filled up, the passing of new socket buffers is suspended with a call to `netif_stop_queue()`.

Creation of EoE Handlers During bus scanning (see sec. 5.4), the master determines the supported mailbox protocols for each slave. This is done by examining the “Supported Mailbox Protocols” mask field at word address 0x001C of the SII. If bit 1 is set, the slave supports the EoE protocol. In this case, an EoE handler is created for that slave.

EoE State Machine Every EoE handler owns an EoE state machine, that is used to send frames to the corresponding slave and receive frames from it via the EoE communication primitives. This state machine is shown in figure 6.1.




Figure 6.1: Transition Diagram of the EoE State Machine

RX_START The beginning state of the EoE state machine. A mailbox check datagram is sent, to query the slave’s mailbox for new frames. → RX_CHECK

RX_CHECK The mailbox check datagram is received. If the slave’s mailbox did not contain data, a transmit cycle is started. → TX_START

If there are new data in the mailbox, a datagram is sent to fetch the new data.
→ RX_FETCH

RX_FETCH The fetch datagram is received. If the mailbox data do not contain a “EoE Fragment request” command, the data are dropped and a transmit sequence is started. → TX_START

If the received Ethernet frame fragment is the first fragment, a new socket buffer is allocated. In either case, the data are copied into the correct position of the socket buffer.

If the fragment is the last fragment, the socket buffer is forwarded to the network stack and a transmit sequence is started. → TX_START

Otherwise, a new receive sequence is started to fetch the next fragment. → RX_START

TX_START The beginning state of a transmit sequence. It is checked, if the transmission queue contains a frame to send. If not, a receive sequence is started. → RX_START

If there is a frame to send, it is dequeued. If the queue was inactive before (because it was full), the queue is woken up with a call to *netif_wake_queue()*. The first fragment of the frame is sent. → TX_SENT

TX_SENT It is checked, if the first fragment was sent successfully. If the current frame consists of further fragments, the next one is sent. → TX_SENT

If the last fragment was sent, a new receive sequence is started. → RX_START

EoE Processing To execute the EoE state machine of every active EoE handler, there must be a cyclic process. The easiest solution would be to execute the EoE state machines synchronously with the master state machine (see sec. 5.3). This approach has the following disadvantage:

Only one EoE fragment could be sent or received every few cycles. This causes the data rate to be very low, because the EoE state machines are not executed in the time between the application cycles. Moreover, the data rate would be dependent on the period of the application task.

To overcome this problem, an own cyclic process is needed to asynchronously execute the EoE state machines. For that, the master owns a kernel timer, that is executed each timer interrupt. This guarantees a constant bandwidth, but poses the new problem of concurrent access to the master. The locking mechanisms needed for this are introduced in sec. 3.3.

Automatic Configuration By default, slaves are left in PREOP state, if no configuration is applied. If an EoE interface link is set to “up”, the requested slave’s application-layer state is automatically set to OP.

6.2 CANopen over EtherCAT (CoE)

The CANopen over EtherCAT protocol [3, sec. 5.6] is used to configure slaves and exchange data objects on application level.

...

SDO Download State Machine The best time to apply SDO configurations is during the slave’s PREOP state, because mailbox communication is already possible and slave’s application will start with updating input data in the succeeding SAFEOP state. Therefore the SDO configuration has to be part of the slave configuration state machine (see sec. 5.5): It is implemented via an SDO download state machine, that is executed just before entering the slave’s SAFEOP state. In this way, it is guaranteed that the SDO configurations are applied each time, the slave is reconfigured.

The transition diagram of the SDO Download state machine can be seen in figure 6.2.

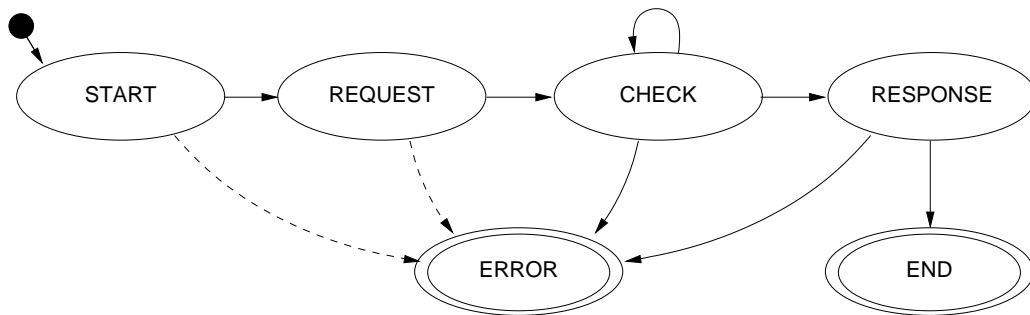


Figure 6.2: Transition diagram of the CoE download state machine

START The beginning state of the CoE download state machine. The “SDO Download Normal Request” mailbox command is sent. → REQUEST

REQUEST It is checked, if the CoE download request has been received by the slave. After that, a mailbox check command is issued and a timer is started. → CHECK

CHECK If no mailbox data is available, the timer is checked.

- If it timed out, the SDO download is aborted. → ERROR
- Otherwise, the mailbox is queried again. → CHECK

If the mailbox contains new data, the response is fetched. → RESPONSE

RESPONSE If the mailbox response could not be fetched, the data is invalid, the wrong protocol was received, or a “Abort SDO Transfer Request” was received, the SDO download is aborted. → ERROR

If a “SDO Download Normal Response” acknowledgement was received, the SDO download was successful. → END

END The SDO download was successful.

ERROR The SDO download was aborted due to an error.

7 Userspace Interfaces

For the master runs as a kernel module, accessing it is natively limited to analyzing Syslog messages and controlling using *modutils*.

It was necessary to implement further interfaces, that make it easier to access the master from userspace and allow a finer influence. It should be possible to view and to change special parameters at runtime.

Bus visualization is another point: For development and debugging purposes it is necessary to show the connected slaves with a single command, for instance (see sec. 7.1).

Another aspect is automatic startup and configuration. The master must be able to automatically start up with a persistent configuration (see sec. 7.2).

A last thing is monitoring EtherCAT communication. For debugging purposes, there had to be a way to analyze EtherCAT datagrams. The best way would be with a popular network analyzer, like Wireshark [6] (the former Ethereal) or others (see sec. 7.3).

This chapter covers all these points and introduces the interfaces and tools to make all that possible.

7.1 Command-line Tool

7.1.1 Character Devices

Each master instance will get a character device as a userspace interface. The devices are named */dev/EtherCATx*, where $x \in \{0 \dots n\}$ is the index of the master.

Device Node Creation The character device nodes are automatically created, if the *udev* Package is installed. See sec. 9.4 for how to install and configure it.

7.1.2 Setting Alias Addresses

```
ethercat alias [OPTIONS] <ALIAS>
```

```
Write alias addresses.
```

Arguments:

ALIAS must be an unsigned 16 bit number. Zero means removing an alias address.

If multiple slaves are selected, the `--force` option is required.

Command-specific options:

```
--alias    -a <alias>
--position -p <pos>    Slave selection. See the help of
                        the 'slaves' command.
--force    -f          Acknowledge writing aliases of
                        multiple slaves.
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.3 Displaying the Bus Configuration

```
ethercat config [OPTIONS]
```

Show slave configurations.

Without the `--verbose` option, slave configurations are output one-per-line. Example:

```
1001:0  0x0000003b/0x02010000  3  OP
|      |                      |  |
|      |                      |  \- Application-layer
|      |                      |    state of the attached
|      |                      |    slave, or '-', if no
|      |                      |    slave is attached.
|      |                      \- Absolute decimal ring
|      |                      position of the attached
|      |                      slave, or '-' if none
|      |                      attached.
|      \- Expected vendor ID and product code (both
|          hexadecimal).
\ - Alias address and relative position (both decimal).
```

With the `--verbose` option given, the configured PDOs and SDOs are output in addition.

Configuration selection:

Slave configurations can be selected with the `--alias` and `--position` parameters as follows:

- 1) If neither the `--alias` nor the `--position` option is given, all slave configurations are displayed.
- 2) If only the `--position` option is given, an alias of zero is assumed (see 4)).
- 3) If only the `--alias` option is given, all slave

configurations with the given alias address are displayed.

- 4) If both the `--alias` and the `--position` option are given, the selection can match a single configuration, that is displayed, if it exists.

Command-specific options:

```
--alias      -a <alias>  Configuration alias (see above).
--position   -p <pos>    Relative position (see above).
--verbose    -v          Show detailed configurations.
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.4 Displaying Process Data

```
ethercat data [OPTIONS]
```

Output binary domain process data.

Data of multiple domains are concatenated.

Command-specific options:

```
--domain -d <index>  Positive numerical domain index.
                    If omitted, data of all domains
                    are output.
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.5 Setting a Master's Debug Level

```
ethercat debug <LEVEL>
```

Set the master's debug level.

Debug messages are printed to syslog.

Arguments:

```
LEVEL can have one of the following values:
  0 for no debugging output,
  1 for some debug messages, or
  2 for printing all frame contents (use with caution!).
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.6 Configured Domains

```
ethercat domains [OPTIONS]
```

Show configured domains.

Without the `--verbose` option, the domains are displayed one-per-line. Example:

```
Domain0: LogBaseAddr 0x00000000, Size 6, WorkingCounter 0/1
```

The domain's base address for the logical datagram (LRD/LWR/LRW) is displayed followed by the domain's process data size in byte. The last values are the current datagram working counter sum and the expected working counter sum. If the values are equal, all PDOs were exchanged during the last cycle.

If the `--verbose` option is given, the participating slave configurations/FMMUs and the current process data are additionally displayed:

```
Domain1: LogBaseAddr 0x00000006, Size 6, WorkingCounter 0/1
  SlaveConfig 1001:0, SM3 ( Input), LogAddr 0x00000006, Size 6
  0x00 0x00 0x00 0x00 0x00 0x00
```

The process data are displayed as hexadecimal bytes.

Command-specific options:

```
--domain -d <index> Positive numerical domain index.
                    If omitted, all domains are
                    displayed.

--verbose -v        Show FMMUs and process data
                    in addition.
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.7 Master and Ethernet Devices

```
ethercat master [OPTIONS]
```

Show master and Ethernet device information.

Command-specific options:

```
--master -m <index> Index of the master to use. Default: 0.
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.8 Sync Managers, PDOs and PDO Entries

```
ethercat pdos [OPTIONS]
```

List Sync managers, PDO assignment and mapping.

The information is displayed in three layers, which are indented accordingly:

- 1) Sync managers - Contains the sync manager information from the SII: Index, physical start address, default size, control register and enable word. Example:

```
SM3: PhysAddr 0x1100, DefaultSize 0, ControlRegister 0x20, Enable 1
```

- 2) Assigned PDOs - PDO direction, hexadecimal index and the PDO name, if available. Note that a 'Tx' and 'Rx' are seen from the slave's point of view. Example:

```
TxPDO 0x1a00 "Channel1"
```

- 3) Mapped PDO entries - PDO entry index and subindex (both hexadecimal), the length in bit and the description, if available. Example:

```
PDO entry 0x3101:01, 8 bit, "Status"
```

Note, that the displayed PDO assignment and PDO mapping information can either originate from the SII or from the CoE communication area.

Command-specific options:

```
--alias      -a <alias>
--position   -p <pos>      Slave selection. See the help of
                           the 'slaves' command.
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.9 SDO Dictionary

```
ethercat sdos [OPTIONS]
```

List SDO dictionaries.

SDO dictionary information is displayed in two layers, which are indented accordingly:

- 1) SDOs - Hexadecimal SDO index and the name. Example:

```
SDO 0x1018, "Identity object"
```

- 2) SDO entries - SDO index and SDO entry subindex (both hexadecimal) followed by the data type, the length in bit, and the description. Example:

```
0x1018:01, uint32, 32 bit, "Vendor id"
```

If the `--quiet` option is given, only the SDOs are output.

Command-specific options:

```
--alias    -a <alias>
--position -p <pos>    Slave selection. See the help of
                       the 'slaves' command.
--quiet    -q          Only output SDOs (without the
                       SDO entries).
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.10 SDO Access

```
ethercat download [OPTIONS] <INDEX> <SUBINDEX> <VALUE>
```

Write an SDO entry to a slave.

This command requires a single slave to be selected.

The data type of the SDO entry is taken from the SDO dictionary by default. It can be overridden with the `--type` option. If the slave does not support the SDO information service or the SDO is not in the dictionary, the `--type` option is mandatory.

These are the valid SDO entry data types:

```
int8, int16, int32, uint8, uint16, uint32, string,
octet_string.
```

Arguments:

```
INDEX      is the SDO index and must be an unsigned
            16 bit number.
SUBINDEX   is the SDO entry subindex and must be an
            unsigned 8 bit number.
VALUE      is the value to download and must correspond
            to the SDO entry datatype (see above).
```

Command-specific options:

```
--alias    -a <alias>
--position -p <pos>    Slave selection. See the help of
                       the 'slaves' command.
--type     -t <type>   SDO entry data type (see above).
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

```
ethercat upload [OPTIONS] <INDEX> <SUBINDEX>
```


Read an SDO entry from a slave.

This command requires a single slave to be selected.

The data type of the SDO entry is taken from the SDO dictionary by default. It can be overridden with the `--type` option. If the slave does not support the SDO information service or the SDO is not in the dictionary, the `--type` option is mandatory.

These are the valid SDO entry data types:

```
int8, int16, int32, uint8, uint16, uint32, string,
octet_string.
```

Arguments:

```
INDEX      is the SDO index and must be an unsigned
            16 bit number.
SUBINDEX   is the SDO entry subindex and must be an
            unsigned 8 bit number.
```

Command-specific options:

```
--alias     -a <alias>
--position  -p <pos>      Slave selection. See the help of
                           the 'slaves' command.
--type      -t <type>    SDO entry data type (see above).
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.11 Slaves on the Bus

Slave information can be gathered with the subcommand `slaves`:

```
ethercat slaves [OPTIONS]
```

Display slaves on the bus.

If the `--verbose` option is not given, the slaves are displayed one-per-line. Example:

```
1 5555:0 PREOP + EL3162 2C. Ana. Input 0-10V
| | | | | |
| | | | | \- Name from the SII if available,
| | | | | otherwise vendor ID and product
| | | | | code (both hexadecimal).
| | | | | \- Error flag. '+' means no error,
| | | | | 'E' means that scan or
| | | | | configuration failed.
| | | | | \- Current application-layer state.
| | | | | \- Decimal relative position to the last
| | | | | slave with an alias address set.
```

```
| \- Decimal alias address of this slave (if set),  
|     otherwise of the last slave with an alias set,  
|     or zero, if no alias was encountered up to this  
|     position.  
\- Absolute ring position in the bus.
```

If the `--verbose` option is given, a detailed (multi-line) description is output for each slave.

Slave selection:

Slaves for this and other commands can be selected with the `--alias` and `--position` parameters as follows:

- 1) If neither the `--alias` nor the `--position` option is given, all slaves are selected.
- 2) If only the `--position` option is given, it is interpreted as an absolute ring position and a slave with this position is matched.
- 3) If only the `--alias` option is given, all slaves with the given alias address and subsequent slaves before a slave with a different alias address match (use `-p0` if only the slaves with the given alias are desired, see 4)).
- 4) If both the `--alias` and the `--position` option are given, the latter is interpreted as relative position behind any slave with the given alias.

Command-specific options:

```
--alias    -a <alias>  Slave alias (see above).  
--position -p <pos>    Slave position (see above).  
--verbose  -v          Show detailed slave information.
```

Numerical values can be specified either with decimal (no prefix), octal (prefix `'0'`) or hexadecimal (prefix `'0x'`) base.

Below is a typical output:

```
$ ethercat slaves  
0      0:0  PREOP  +  EK1100 Ethernet Kopplerklemme (2A E-Bus)  
1  5555:0  PREOP  +  EL3162 2K. Ana. Eingang 0-10V  
2  5555:1  PREOP  +  EL4102 2K. Ana. Ausgang 0-10V  
3  5555:2  PREOP  +  EL2004 4K. Dig. Ausgang 24V, 0,5A
```

7.1.12 SII Access

It is possible to directly read or write the complete SII contents of the slaves. This was introduced for the reasons below:

- The format of the SII data is still in development and categories can be added in the future. With read and write access, the complete memory contents can be easily backed up and restored.

- Some SII data fields have to be altered (like the alias address). A quick writing must be possible for that.
- Through reading access, analyzing category data is possible from userspace.

```
ethercat sii_read [OPTIONS]
```

Output a slave's SII contents.

This command requires a single slave to be selected.

Without the `--verbose` option, binary SII contents are output.

With the `--verbose` option given, a textual representation of the data is output, that is separated by SII category names.

Command-specific options:

```
--alias      -a <alias>
--position   -p <pos>      Slave selection. See the help of
                           the 'slaves' command.
--verbose    -v           Output textual data with
                           category names.
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

Reading out SII data is as easy as other commands. Though the data are in binary format, analysis is easier with a tool like *hexdump*:

```
$ ethercat sii_read --position 3 | hexdump
00000000 0103 0000 0000 0000 0000 0000 0000 008c
00000010 0002 0000 3052 07f0 0000 0000 0000 0000
00000020 0000 0000 0000 0000 0000 0000 0000 0000
...
```

Backing up SII contents can easily done with a redirection:

```
$ ethercat sii_read --position 3 > sii-of-slave3.bin
```

To download SII contents to a slave, writing access to the master's character device is necessary (see sec. 7.1.1).

```
ethercat sii_write [OPTIONS] <FILENAME>
```

Write SII contents to a slave.

This command requires a single slave to be selected.

The file contents are checked for validity and integrity. These checks can be overridden with the `--force` option.

Arguments:

FILENAME must be a path to a file that contains a positive number of words. If it is '-', data are read from stdin.

Command-specific options:

--alias -a <alias>
--position -p <pos> Slave selection. See the help of the 'slaves' command.
--force -f Override validity checks.

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

```
# ethercat sii_write --position 3 sii-of-slave3.bin
```

The SII contents will be checked for validity and then sent to the slave. The write operation may take a few seconds.

7.1.13 Requesting Application-Layer States

```
ethercat states [OPTIONS] <STATE>
```

Request application-layer states.

Arguments:

STATE can be 'INIT', 'PREOP', 'SAFEOP', or 'OP'.

Command-specific options:

--alias -a <alias>
--position -p <pos> Slave selection. See the help of the 'slaves' command.

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.1.14 Generating Slave Description XML

```
ethercat xml [OPTIONS]
```

Generate slave information XML.

Note that the PDO information can either originate from the SII or from the CoE communication area. For slaves, that support configuring PDO assignment and mapping, the output depends on the last configuration.

Command-specific options:

```
--alias      -a <alias>
--position  -p <pos>      Slave selection. See the help of
                          the 'slaves' command.
```

Numerical values can be specified either with decimal (no prefix), octal (prefix '0') or hexadecimal (prefix '0x') base.

7.2 System Integration

To integrate the EtherCAT master as a service into a running system, it comes with an init script and a sysconfig file, that are described below.

7.2.1 Init Script

The EtherCAT master init script conforms to the requirements of the “Linux Standard Base” (LSB, [5]). The script is installed to *etc/init.d/ethercat* below the installation prefix and has to be copied (or better: linked) to the appropriate location (see sec. 9), before the master can be inserted as a service. Please note, that the init script depends on the sysconfig file described below.

To provide service dependencies (i. e. which services have to be started before others) inside the init script code, LSB defines a special comment block. System tools can extract this information to insert the EtherCAT init script at the correct place in the startup sequence:

```
# Default-Stop:      0 1 2 6
# Short-Description: EtherCAT master
# Description:       EtherCAT master 1.4.0
### END INIT INFO
```

```
#-----
```

```
MODPROBE=/sbin/modprobe
RMMOD=/sbin/rmmod
MODINFO=/sbin/modinfo
ETHERCAT=/opt/etherlab/bin/ethercat
```

7.2.2 Sysconfig File

For persistent configuration, the init script uses a sysconfig file installed to *etc/sysconfig/ethercat* (below the installation prefix), that is mandatory for the init script. The sysconfig file contains all configuration variables needed to operate one or more masters. The documentation is inside the file and included below:

```
9 #
10 # Master devices.
11 #
12 # The MASTER<X>_DEVICE variable specifies the Ethernet device for a master
13 # with index 'X'.
14 #
15 # Specify the MAC address (hexadecimal with colons) of the Ethernet device to
16 # use. Example: "00:00:08:44:ab:66"
17 #
18 # The broadcast address "ff:ff:ff:ff:ff:ff" has a special meaning: It tells
19 # the master to accept the first device offered by any Ethernet driver.
20 #
21 # The MASTER<X>_DEVICE variables also determine, how many masters will be
22 # created: A non-empty variable MASTER0_DEVICE will create one master, adding
23 # a non-empty variable MASTER1_DEVICE will create a second master, and so on.
24 #
25 MASTER0_DEVICE=""
26 #MASTER1_DEVICE=""
27 #
28 #
29 # Ethernet driver modules to replace with EtherCAT-capable ones.
30 #
31 # The init script will try to unload the Ethernet driver modules in the list
32 # and replace them with the EtherCAT-capable ones, respectively. If a certain
33 # (EtherCAT-capable) driver is not found, a warning will appear.
34 #
35 # Possible values: 8139too, e1000.
36 # Separate multiple drivers with spaces.
37 #
38 # Note: The e1000 driver is not built by default. Enable it with
39 # the --enable-e1000 configure switch.
40 #
41 DEVICE_MODULES=""
42 #
43 #
44 # Flags for loading kernel modules.
45 #
46 # This can usually be left empty. Adjust this variable, if you have problems
47 # with module loading.
48 #
49 #MODPROBE_FLAGS="-b"
50 #
51 #-----
```

7.2.3 Starting the Master as a Service

After the init script and the sysconfig file are placed into the right location, the EtherCAT master can be inserted as a service. The different Linux distributions offer different ways to mark a service for starting and stopping in certain runlevels. For example, SUSE Linux provides the *insserv* command:

```
# insserv ethercat
```

The init script can also be used for manually starting and stopping the EtherCAT master. It has to be executed with one of the parameters `start`, `stop`, `restart` or `status`.

```
# /etc/init.d/ethercat restart
```

```
Shutting down EtherCAT master      done
Starting EtherCAT master            done
```

7.3 Debug Interfaces

EtherCAT buses can always be monitored by inserting a switch between master and slaves. This allows to connect another PC with a network monitor like Wireshark [6], for example.

For convenience, so-called “debug interfaces” are supported. Debug interfaces are virtual network interfaces allowing to capture EtherCAT traffic with a network monitor (like Wireshark or tcpdump) running on the master machine without using external hardware. To use this functionality, the master sources have to be configured with the `--enable-debug-if` switch (see sec. 9).

Every EtherCAT master registers two read-only network interfaces, corresponding to the physical Ethernet devices. These are named *ecdbgmX* (main device) and *ecdbgbX* (backup device, for future use), where X is the master index. The below listing shows debug interfaces among some standard network interfaces:

```
# ip link
1: lo: <LOOPBACK,UP> mtu 16436 qdisc noqueue
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
4: eth0: <BROADCAST,MULTICAST> mtu 1500 qdisc noop qlen 1000
    link/ether 00:04:61:03:d1:01 brd ff:ff:ff:ff:ff:ff
8: ecdbgm0: <BROADCAST,MULTICAST> mtu 1500 qdisc pfifo_fast
    qlen 1000
    link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
9: ecdbgb0: <BROADCAST,MULTICAST> mtu 1500 qdisc noop qlen 1000
    link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
```

While a debug interface is enabled, all frames sent or received to or from the physical device are additionally forwarded to the debug interface by the corresponding master. Network interfaces can be enabled with the below command:

```
# ip link set dev ecdbgm0 up
```

Please note, that the frame rate can be very high. With an application connected, the debug interface can produce thousands of frames per second.

Attention The socket buffers needed for the operation of debug interfaces have to be allocated dynamically. Some Linux realtime extensions do not allow this in realtime context!

8 Timing Aspects

Although EtherCAT's timing is highly deterministic and therefore timing issues are rare, there are a few aspects that can (and should be) dealt with.

8.1 Application Interface Profiling

One of the most important timing aspects are the execution times of the application interface functions, that are called in cyclic context. These functions make up an important part of the overall timing of the application. To measure the timing of the functions, the below cyclic code was used:

```
c0 = get_cycles();
ecrt_master_receive(master);
c1 = get_cycles();
ecrt_domain_process(domain1);
c2 = get_cycles();
ecrt_domain_queue(domain1);
c3 = get_cycles();
ecrt_master_send(master);
c4 = get_cycles();
```

Between each call of an interface function, the CPU timestamp counter is read. The counter differences are converted to μs via the `cpu_khz` variable, that contains the number of counts per ms for the IA32 architecture's timestamp counter.

For the actual measurement, a system with a 2.0 GHz CPU was used, that ran the above code in an RTAI thread with a period of 1 ms. The measurement was repeated $n = 10000$ times and the results were averaged. These can be seen in table 8.1.

Table 8.1: Application Cycle on a 2.0 GHz Processor

Function	$\mu(\Delta t)$ [μs]	$\sigma(\Delta t)$ [μs]
<code>ecrt_master_receive()</code>	6.13	1.11
<code>ecrt_domain_process()</code>	< 0.01	0.07
<code>ecrt_domain_queue()</code>	< 0.01	0.17
<code>ecrt_master_send()</code>	1.15	0.65
Complete Cycle	7.28	1.31

It is obvious, that the functions accessing hardware make up the lion's share. The `ec_master_receive()` executes the ISR of the Ethernet device driver, dissects the received frame and copies the datagram contents into the memory of the corresponding datagram objects. The `ec_master_send()` function assembles a frame from different datagrams and copies it to the hardware buffers. The functions that only operate on the masters internal data structures are very fast ($\Delta t < 1 \mu\text{s}$).

For a realtime cycle makes up about $10 \mu\text{s}$, the resulting theoretical frequency could be up to $1/10 \mu\text{s} = 100 \text{ kHz}$. For two reasons, this frequency keeps being theoretical:

1. The processor must still be able to run the operating system between the real-time cycles.
2. The EtherCAT frame must be sent and received, before the next realtime cycle begins. The determination of the bus cycle time is difficult and covered in sec. 8.2.

8.2 Bus Cycle Measurement

For measurement the time, a frame is “on the wire”, two timestamps must be taken:

1. The time, the Ethernet hardware begins with physically sending the frame.
2. The time, the frame is completely received by the Ethernet hardware.

Both times are difficult to determine. The first reason is, that the interrupts are disabled and the master is not notified, when a frame is sent or received (polling would distort the results). The second reason is, that even with interrupts enabled, the interrupt latency (i.e. the time from the event to the notification) is unknown. Therefore the only way to confidently determine the bus cycle time is an electrical measurement.

Anyway, the bus cycle time is an important factor when designing realtime applications, because it limits the maximum frequency for the cyclic task. In practice, these timing parameters are highly dependent on the hardware and often a trial and error method must be used to determine the limits of the system.

An essential question is: What happens, if the cycle frequency is too high? The EtherCAT frames that have been sent at the end of the cycle could have been not yet received when the next cycle starts. First this is noticed by the domain, because the working counters of the datagrams are zero. This can be queried in realtime context via the application interface and is output via Syslog¹. In this case, the process data keeps being the same as in the last cycle, because it is not erased by the domain. When the domain datagrams are queued again, the master notices, that they are already queued (and marked as sent). The master will mark them as unsent again and output a warning, that datagrams were “skipped”.

¹To limit Syslog output, a mechanism has been implemented, that outputs a summarized notification at maximum once a second.

On the mentioned 2.0 GHz system, the possible cycle frequency can be up to 25 kHz without skipped frames. This value is highly dependant on the chosen hardware.

9 Installation

9.1 Building the Software

The current EtherCAT master code is available at [1] or can be obtained from the EtherLab CD. The *tar.bz2* file has to be unpacked with the commands below (or similar):

```
$ tar xjf ethercat-1.4.0.tar.bz2
$ cd ethercat-1.4.0/
```

The tarball was created with GNU Autotools, so the build process follows the below commands:

```
$ ./configure
$ make all modules
```

Table 9.1 lists important configuration switches and options.

9.2 Building the Interface Documentation

The source code is documented using Doxygen [10]. To build the HTML documentation, the Doxygen software has to be installed. The below command will generate the documents in the subdirectory *doxygen-output*:

```
$ make doc
```

The interface documentation can be viewed by pointing a browser to the file *doxygen-output/html/index.html*. The functions and data structures of the application interface are covered by an own module “Application Interface”.

9.3 Installing the Software

The below commands have to be entered as *root*: The `install` target will install the EtherCAT header, init script, sysconfig file and the userspace tool to the prefix path. The `modules_install` target will install the kernel modules to the kernel’s modules directory. The final `depmod` call is necessary to include the kernel modules into the *modules.dep* file to make it available to the `modprobe` command, that is used in the init script.

Table 9.1: Configuration options

Option/Switch	Description	Default
<code>--prefix</code>	Installation prefix	<i>/opt/etherlab</i>
<code>--with-linux-dir</code>	Linux kernel sources	Use running kernel
<code>--with-rtai-dir</code>	RTAI path (only for RTAI example)	
<code>--enable-oe</code>	Enable EoE support	yes
<code>--enable-cycles</code>	Use CPU timestamp counter. Enable this on Intel architecture to get finer timing calculation.	no
<code>--enable-debug-if</code>	Create a debug interface for each master	no
<code>--enable-debug-ring</code>	Create a debug ring to record frames	no
<code>--enable-8139too</code>	Build the 8139too driver	yes
<code>--with-8139too-kernel</code>	8139too kernel	†
<code>--enable-e1000</code>	Enable e1000 driver	no
<code>--with-e1000-kernel</code>	e1000 kernel	†

† If this option is not specified, the kernel version to use is extracted from the Linux kernel sources.

```
# make install modules_install
# depmod
```

If the target kernel's modules directory is not under */lib/modules*, a different destination directory can be specified with the `DESTDIR` make variable. For example:

```
# make DESTDIR=/vol/nfs/root modules_install
```

This command will install the compiled kernel modules to */vol/nfs/root/lib/modules*, prepended by the kernel release.

If the EtherCAT master shall be run as a service¹ (see sec. 7.2), the init script and the sysconfig file have to be copied (or linked) to the appropriate locations. The below example is suitable for SUSE Linux. It may vary for other distributions.

```
# cd /opt/etherlab
# cp etc/sysconfig/ethercat /etc/sysconfig/
# ln -s etc/init.d/ethercat /etc/init.d/
# inserv ethercat
```

Now the sysconfig file */etc/sysconfig/ethercat* (see sec. 7.2.2) has to be customized. The minimal customization is to set the `MASTER0_DEVICE` variable to the MAC address of the Ethernet device to use (or `ff:ff:ff:ff:ff:ff` to use the first device offered) and selecting the driver(s) to load via the `DEVICE_MODULES` variable.

After the basic configuration is done, the master can be started with the below command:

```
# /etc/init.d/ethercat start
```

At this time, the operation of the master can be observed by viewing the Syslog messages, which should look like the ones below. If EtherCAT slaves are connected to the master's EtherCAT device, the activity indicators should begin to flash.

```
1 EtherCAT: Master driver 1.4.0
2 EtherCAT: 1 master waiting for devices.
3 EtherCAT Intel(R) PRO/1000 Network Driver - version 6.0.60-k2
4 Copyright (c) 1999-2005 Intel Corporation.
5 PCI: Found IRQ 12 for device 0000:01:01.0
6 PCI: Sharing IRQ 12 with 0000:00:1d.2
7 PCI: Sharing IRQ 12 with 0000:00:1f.1
8 EtherCAT: Accepting device 00:0E:0C:DA:A2:20 for master 0.
9 EtherCAT: Starting master thread.
10 ec_e1000: ec0: e1000_probe: Intel(R) PRO/1000 Network
11         Connection
12 ec_e1000: ec0: e1000_watchdog_task: NIC Link is Up 100 Mbps
```

¹Even if the EtherCAT master shall not be loaded on system startup, the use of the init script is recommended for manual (un-)loading.

```
13             Full Duplex
14 EtherCAT: Link state changed to UP.
15 EtherCAT: 7 slave(s) responding.
16 EtherCAT: Slave states: PREOP.
17 EtherCAT: Scanning bus.
18 EtherCAT: Bus scanning completed in 431 ms.
```

- ① – ② The master module is loading, and one master is initialized.
- ③ – ⑧ The EtherCAT-capable e1000 driver is loading. The master accepts the device with the address 00:0E:0C:DA:A2:20.
- ⑨ – ⑯ The master goes to idle phase, starts its state machine and begins scanning the bus.

9.4 Automatic Device Node Creation

The `ethercat` command-line tool (see sec. 7.1) communicates with the master via a character device. The corresponding device nodes are created automatically, if the `udev` daemon is running. Note, that on some distributions, the `udev` package is not installed by default.

The device nodes will be created with mode `0660` and group `root` by default. If “normal” users shall have reading access, a `udev` rule file (for example `/etc/udev/rules.d/99-EtherCAT.rules`) has to be created with the following contents:

```
KERNEL=="EtherCAT[0-9]*", MODE="0664"
```

After the `udev` rule file is created and the EtherCAT master is restarted with `/etc/init.d/ethercat restart`, the device node will be automatically created with the desired rights:

```
# ls -l /dev/EtherCAT0
crw-rw-r-- 1 root root 252, 0 2008-09-03 16:19 /dev/EtherCAT0
```

Now, the `ethercat` tool can be used (see sec. 7.1) even as a non-root user.

If non-root users shall have writing access, the following `udev` rule can be used instead:

```
KERNEL=="EtherCAT[0-9]*", MODE="0664", GROUP="users"
```


Bibliography

- [1] Ingenieurgesellschaft IgH: EtherLab – Open Source Toolkit for rapid realtime code generation under Linux with Simulink/RTW and EtherCAT technology. <http://etherlab.org/en>, 2008.
- [2] IEC 61158-4-12: Data-link Protocol Specification. International Electrotechnical Commission (IEC), 2005.
- [3] IEC 61158-6-12: Application Layer Protocol Specification. International Electrotechnical Commission (IEC), 2005.
- [4] GNU General Public License, Version 2. <http://www.gnu.org/licenses/gpl-2.0.html>. October 15, 2008.
- [5] Linux Standard Base. <http://www.linuxfoundation.org/en/LSB>. August 9, 2006.
- [6] Wireshark. <http://www.wireshark.org>. 2008.
- [7] *Hopcroft, J. E. / Ullman, J. D.*: Introduction to Automata Theory, Languages and Computation. Adison-Wesley, Reading, Mass. 1979.
- [8] *Wagner, F. / Wolstenholme, P.*: State machine misunderstandings. In: IEE journal “Computing and Control Engineering”, 2004.
- [9] RTAI. The RealTime Application Interface for Linux from DIAPM. <http://www.rtai.org>, 2006.
- [10] Doxygen. Source code documentation generator tool. <http://www.stack.nl/~dimitri/doxygen>, 2008.

Glossary

CoE CANopen over EtherCAT, Mailbox Protocol, page 41

ecdev EtherCAT Device, page 19

EoE Ethernet over EtherCAT, Mailbox Protocol, page 39

FSM Finite State Machine, page 21

ISR Interrupt Service Routine, page 15

PCI Peripheral Component Interconnect, Computer Bus, page 18

RTAI Realtime Application Interface, page 1

Index

- Application, 5
- Application interface, 11
- Bus cycle, 58
- CoE, 41
- Concurrency, 13
- Debug Interfaces, 55
- Device interface, 19
- Device modules, 5
- Domain, 9
- EoE, 39
- Example Applications, 11
- FMMU
 - Configuration, 9
- FSM, 21
 - EoE, 40
 - Master, 27
 - PDO, 34
 - SII, 33
 - Slave Configuration, 30
 - Slave Scan, 27
 - State Change, 30
 - Theory, 22
- GPL, 3
- Idle phase, 8
- Init script, 7, 53
- Interrupt, 15, 17
- ISR, 15
- LSB, 53
- MAC address, 7
- Mailbox, 39
- Master
 - Architecture, 5
 - Features, 1
 - Installation, 61
- Master Module, 5
- Master module, 7
- Master phases, 8
- net_device, 16
- netif, 16
- Network drivers, 15, 19
- Operation phase, 8
- Orphaned phase, 8
- PDO, 9
- Process data, 9
- Profiling, 57
- Service, 54
- SII, 33, 40
 - Access, 50
- Socket buffer, 16, 17
- Sysconfig file, 53
- Syslog, 63
- Userspace, 43