# Interview about EtherCAT, the real-time Ethernet solution

➤ Kai Binder, chief editor of SPS-MAGAZIN, talked to Dr. Dirk Janssen, Holger Büttner and Martin Rostan from Beckhoff about the concept, technology and costs of the industrial Ethernet solution with real-time capability.

### What is the main difference between EtherCAT and conventional fieldbus systems?

| One of the main differences is performance: EtherCAT is several orders of magnitude faster than conventional fieldbus systems like PROFIBUS, CANopen, and DeviceNet, and also significantly faster than SERCOS.

### What is the purpose of this kind of performance? Aren't current systems fast enough?

| Today's fieldbus systems offer cycle times with approximately the same order of magnitude as the control tasks themselves, i.e. about 5 to 10 ms. This is often regarded as adequate, since the orders of magnitude are compatible. However, what is often neglected is the fact that what we are dealing with here are cascaded cycles, i.e. a multitude of cycles at different levels that are usually not synchronized with each other. At the lowest level, a cyclic firmware operates with or without a local I/O expansion bus, above it we have a cyclic fieldbus, and at the end of the chain we have a master card with its own firmware cycle that has to be completed at the end of the fieldbus cycle. The data is then copied into the shared memory and released, so that the PLC can access it. With this methodology, it takes several cycles for the data to travel through the different stages. This can lead to rather long delays, which is particularly significant in applications requiring rapid responses to changes in inputs. It doesn't take a rocket scientist to understand that, compared with directly wired I/Os, the additional delay caused by the fieldbus architecture with its multiple subordinate cycles slows down the applications. Thanks to EtherCAT and its performance, you can simply forget these subordinate cycles. As a matter of fact, they don't even exist any more, since

EtherCAT operates directly down to the electronic terminal block. This means that the controller can use the current process input image for starting the next PLC task, and no longer an image that is already two cycles old by the time it has risen through the cascading levels.

### But isn't this only really relevant for high-end applications, where every microsecond counts?

| This rigorous synchronization across all levels and the significant shortening of the response time also has benefits for conventional controllers, not just for ultra-high" performance applications. One feature that particularly distinguishes us from other Ethernet applications is that we have consolidated the fieldbus, the backplane bus of the controller, and the sub-bus, which is usually manufacturer-specific. Other Ethernet concepts continue to assume that modular devices have a manufacturer-specific sub-bus. Nothing has changed in this respect. We, on the other hand, have consolidated these, which provides a speed advantage and also means that the system no longer requires certain gateway components. This means that the performance gained through EtherCAT also benefits "normal" applications without using a faster CPU and without modifying the application.

Another difference compared with fieldbus systems is much simpler wiring. Users are no longer bound to lines with very short or non-existing drop lines and can use a flexible tree topology. Another issue is network size. This often comes up in discussions with suppliers offering logistics, material flow or transport systems, i.e. applications where 500 m is quite a normal distance. In this situation, the classic fieldbus quickly reaches its limits. Another advantage is that EtherCAT can also replace the backplane bus, for example in

Martin Rostan, EtherCAT
product manager, Beckhoff

Dr. Dirk Janssen, manager
software development, Beckhoff

Holger Büttner, manager development in the
Berlin branch, Beckhoff

applications using classic fieldbuses, whose process data are traditionally transferred from the fieldbus master to the control via a backplane bus. A system with a certain spatial extent may feature one or several fieldbus systems that control certain system components. All signals have to be made available to the master control. This can now be done with EtherCAT: The fieldbus interface is no longer located in the controller, but in the field.

### Is Ethernet not relatively complicated compared with conventional fieldbus systems and requires a lot of IT know-how?

| Simple configuration is extremely important. With EtherCAT, no node addresses have to be set. The node addresses can be retained even if the system is subsequently expanded – the process data mapping therefore remains intact. Transfer rates do not have to be set either. Compared with conventional fieldbus systems, configuration therefore even becomes simpler. But you are right, the opposite is usually true when it comes to Ethernet. With EtherCAT, Internet protocols, i.e. things like FTP, http or web servers are now also available at a very low level in the field. The fine new world of Internet technologies is suddenly also available to users at the fieldbus level. This is another important argument for using Ethernet here.

### However, for users it is not only functionality that is important, but also costs. How does EtherCAT compare in this respect?

| Another advantage of EtherCAT is that it is cost less than conventional fieldbus systems. This has mainly two reasons. On the one hand the infrastructure is unbeatably cost-effective. In an IP 20 environment you can use a patch cable that costs 1 Euro complete with two connectors – ready-made! In con-trast, PROFIBUS, for example, involves significant material costs for the connectors, not to mention labor time for assembly. The cable material itself is not the issue. The question his how quickly and elaborately the connectors have to be mounted. On the other hand our infrastructure requires no switches, which is different from other Ethernet solutions. And, we no longer use a special master card in the IPC or in the PLC, which would normally cost several hundred Euros. This means that we can move at an unprecedented cost level.

### The EtherCAT system has no master?

| Of course we use a master. However, it is a pure software implementation. All the master needs is a conventional Ethernet controller. By the way, we consider this to be a significant distinguishing feature compared with all other real-time Ethernet systems, which invariably require a plug-in card with co-processor and local intelligence, i.e. a special component requiring a powerful 32-bit CPU and plenty of memory. EtherCAT doesn't need this. The only requirement is an Ethernet controller installed on a conventional motherboard or an inexpensive standard LAN card.

### How is this realized technically? Where does the user become involved?

| A distinction has to be made between master and slave. As already mentioned, each standard Ethernet controller in the EtherCAT system can be used as a master. Normal "off-the-shelf" Ethernet infrastructures can be used. On the slave side a special slave controller is used. This system can be used to realize small devices that can respond very quickly and accumulate minimum

Kai Binder is chief editor of SPS-Magazin,
TeDo-Verlag, www.sps-magazin.de

delay during communication. At this point a special component is therefore required that is specially optimized for fieldbus technology. This slave controller is used to realize the EtherCAT principle of operation: A telegram sent by the master is not only sent to one device and returned by the device and so on. Instead, all devices – in some cases several thousand – are addressed at once, and the process data is exchanged with a single Ethernet frame. The telegram therefore passes from one device to the next and finally – having been processed as required – returns to the master. The operating principle is transparent for the user: The process data is visible, and parameterization can also be dealt with via the network.

### Does ASIC not make the system unnecessarily expensive?

| On the contrary! Ethernet is popular since it is widely available and, relative to their performance, the controllers are therefore comparatively cost-effective. But for particularly simple slave devices they are still too expensive. The ASIC solution we chose for EtherCAT means that the connection costs for slave devices are lower than the costs that can be achieved with an Ethernet controller. Furthermore, Ethernet controllers weren't designed for fieldbus technology. Their application interface, for example, requires a lot of processor power. The EtherCAT ASIC was developed specially for picking up or delivering process data directly. This means that simple devices even no longer require a microprocessor. The digital interface offered by EtherCAT ASIC can be connected directly to digital I/Os via associated optocouplers.

### What does this approach mean for users?

| On the slave side the user is not aware whether a special EtherCAT slave controller or an off-the-shelf Ethernet controller is used. After all he doesn't go to the supermarket to buy an I/O component. But he could control an EtherCAT network using a supermarket computer – even though we would recommend a decent industrial computer. It therefore makes sense to do

away with the special master module and to use standard components. That is exactly what we do. On the slave side, however, the aim is to reduce the costs per connection and to achieve a performance level that cannot be achieved with standard components. The hardware implementation also offers performance that is independent from the firmware implementation or processing power of the device. This means that with EtherCAT the response time does not depend on how well a company has implemented its slave. We achieve all this by using the slave controller as FPGA or ASIC.

### What is ASIC likely to cost?

| The ASIC is expected to cost less than five Euros.

### Where will the slave controller be available from?

| The FPGA variant is available from Beckhoff, and from the end of November also from a familiar semiconductor distributor. The ASIC will be handled in a similar way.

### When will the ASIC be ready?

| Fortunately, our main goal does not have to be to make the ASIC available on the market as soon as possible, because with the FPGA we already have an operational and cost-effective solution. It was therefore initially more important for us to ensure availability of the FPGA. We managed to achieve this some time ago. The suggestions that are currently still flowing back from the development projects of the device manufacturers will benefit the ASIC. ETG helps to ensure that EtherCAT is positioned more widely and meets the requirements of more manufacturers and applications. If we had only implemented our own ideas, the project would be further advanced by now. It means that our original target date, i.e. the end of the year, is unlikely to be met. We hope that the ASIC will be available for the Hanover Fair 2005.