Ahh…the power of Ethernet

Overview to recent advances in Ethernet used in industrial IO communication networks

Ethernet will be the mainstream communication technology for years to come; computer and telecom industries, corporate IT, streaming video and audio distribution, and now even machine control is moving towards this technology of packet-oriented communciation on all levels. Many vendors offer Ethernet based communication technology; and from Coax to Twisted Pair cable to Fiber optic there is barely a media, distance or installation scenario that can not be covered by Ethernet. Since about a year, commercially available IO products based on Ethernet or higher protocol layers promise openness and usability in control applications. This article gives an overview to existing technology implementations. Recent product announcements allow to predict that Ethernet is ready to move into the control Input / Output communication world and has the potential to take the fieldbus world to the next level.

Several architectural approaches are in discussion between users, manufacturers and supporting user- and marketing groups (ODVA; IAONA; Profibus Organization) or are simply introduced by single manufacturers. Despite the commonality of using 100 Mbit clock rate and Ethernet communciation chips; incompatibility of selected protocol implementations is reality: Ethernet IO does not mean that all of the products can be used by the same protocol stack, and declared "openness" obiously does not mean interoperability . A control would have to provide stacks for all different products to use them at the same time – and the user has the difficulty of choice of the lasting product.

Ethernet is a packet oriented protocol; and certainly was not intended to cover deterministic communications: based on CSMA/CD (IEE 802.3), multiple network nodes share a media and may even access at the same time (Carrier Sense Multiple Access); while detecting a collision (Collision Detection): the physical layer and lower protocol levels manage an arbitration of a collision situation that forces nodes into a back-off period with a new attempt to communciate. Inside a collision domain, realtime achievements are relative; only acceptable if the IO signal bandwidth is significantly lower than the network bandwidth; in other words, for very slow processes where a few seconds of delay have no impact. Other protocol layers add to these collission delays; which limit the bandwidth of the network.

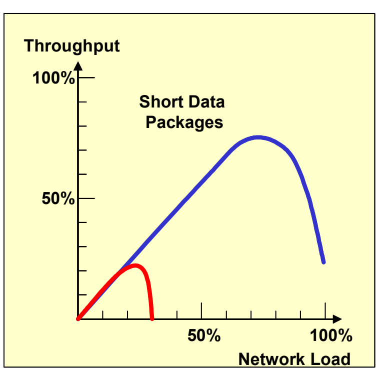

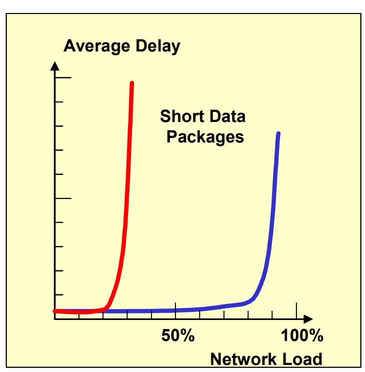

Throughput and average delay in IEE Network technology:
IEEE 802.4 Token Ring „deterministic Method"
IEEE 802.3 CSMA/CD „non deterministic Method"

It is possible to avoid collisions at all by use of switches or by using Token Ring technology (IEE 802.4), both approaches separates traffic of the full-duplex system by buffering messages and directing telegrams to the correct ports only. A comparison of the usable network bandwidth for the typically small IO telegrams in collision domains versus the collision free networks is shown in picture 1 “Throughput and average delay in IEE 802 Network technology”.

Therefore; some vendors started to promote switched IO nodes as a base for sufficiently fast and collision free IO traffic on Ethernet networks; e.g. Schneider, Siemens, (Beckhoff as an intermediate step). For a while; the industry gave the impression that, if everything was connected collision-free through switches, real-time deterministic commiunication behaviour was achievable. The obvious downside of this approach one can find on the cabling side: every IO node requires a switch outlet; the cabling topology is a star scheme; thus increasing the cost per node for Ethernet as fieldbus replacement significantly and making it commercially unattractive for many midsized applications: Profibus; DeviceNet and all the other fieldbus technologies offer best direct wire replacement practice by utilizing more efficient bus, tree or ring structures.

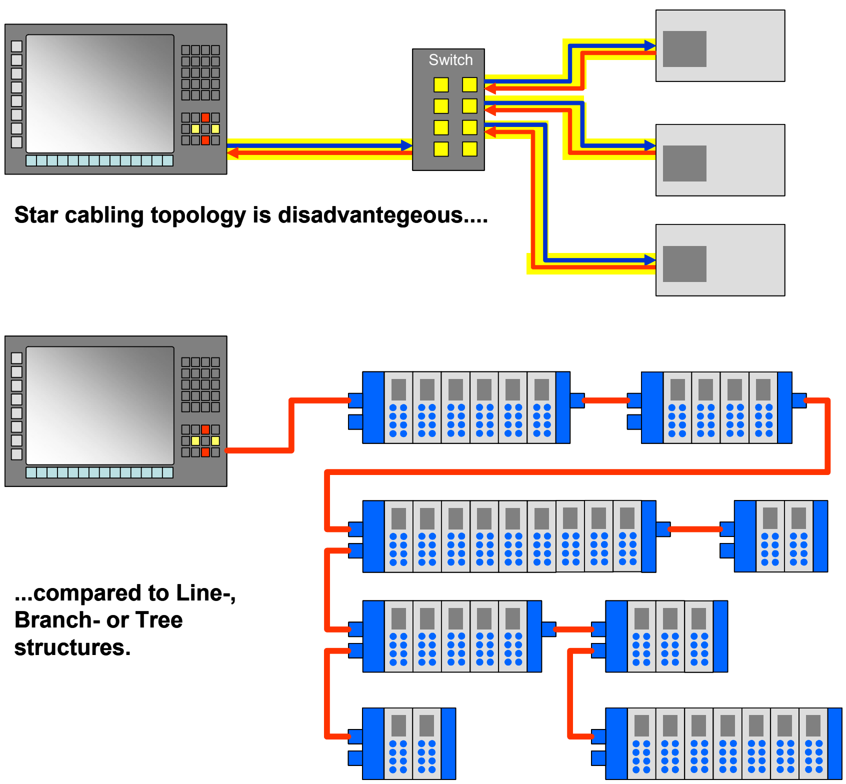

Picture: Cabling through hubs or switches is not really advantegeous compared to line cabling

Secondly, and not so obvious, closer analysis proves that the crux lies in the significant delays which are not in the collisions, but in the switch software itself, and the higher layer protocol stacks which add significantly more delay ot the Ethernet network technology: many vendors offer products based on higher layers of Ethernet protocols: IP, TCP, UDP to mention the main ones: picture 2 “Nesting of protocols in Ethernet frames”.

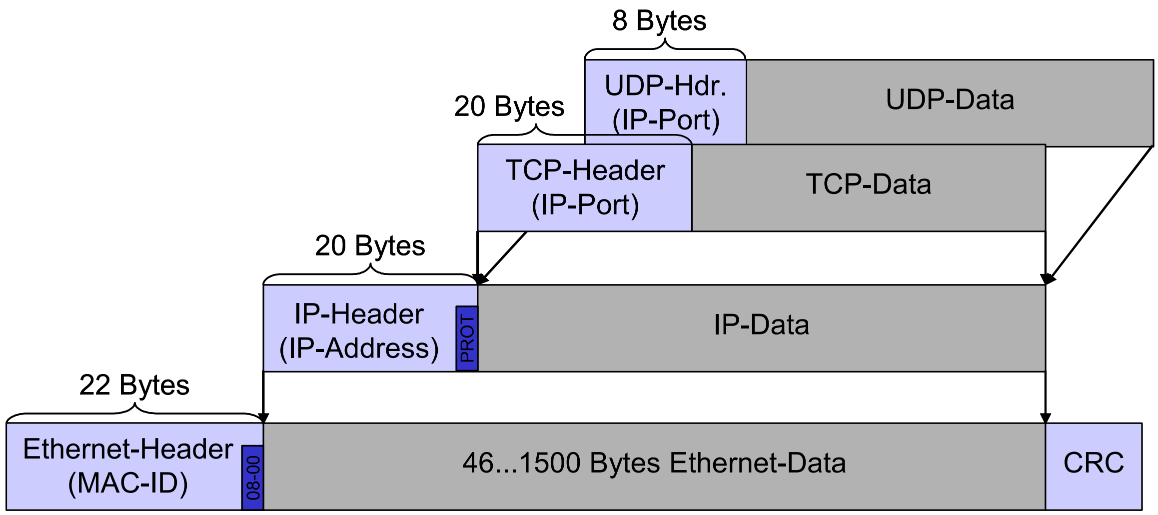

Nesting of protocols in Ethernet frames

Ethernet may contain various other internal packets, one of which is IP (Internet Protocol). IP is responsible for the address resolution for IP addresses to MAC ID's of computers globally; inside the IP frame, either TCP, UDP or something else may travel along. TPC is a confirmed connection between computers, UDP is an unconfirmed connection protocol:

- IP will resolve telegram destination IP address to MAC address for the Ethernet telegram It caches the cross reference in the own computer, or queries the network server or, with it's help, the next higher computer in the hierarchy to resolve the MAC address of the destination computer. In order to do this efficiently, a protocol called ARP (address resolution protocol) is used for that service: it takes some time (sometimes even long times of many seconds) to refresh and dynamically build ARP caches; and this interferes by principle with any deterministic way of communciation. Approaches to use IP deterministically all result in static ARP configurations, and that is no longer "plain vanilla IP"; but needs careful engineering in every single case. It takes the "plug and play" character from Ethernet and creates proprietary mainstream-incompatible situations.
- TCP is a confirmed network protocol; which requires a confirmation telegram for every packet sent. The establishing and termination of such connections require many communciation cycles; and a fairly deterministic behaviour is achievable only by building a proprietary TCP stack with non-standard behaviour. As TCP depends on acknowledge-messages, its behaviour is principally dependent on the slave device and its response time. TCP stacks are quite complex and arbitration of the protocol consumes typically many milliseconds per event. The delays in regular TCP stacks are significant and a multitude larger than regular IO cycle times. TCP does not allow for broad- and multicast messages, these are handled by UDP:
- UDP is an unconfirmed message protocol: the sender delivers its message but does not bother for any confirmation. This makes this protocol quite unsuited for critical IO communication with the need to deliver IO status in 100% of all machine cycles. However, it's simpler structure comes with much faster arbitration times; therefore it is used by some manufacturers despite it's unconfirmed communication status. Using UDP in IO communication seems like driving a car blindfoldedly, stating that "experience shows that there never was a problem with this before".

Conclusion: Ethernet protocol alone may allow for fast and deterministic communciation in collision free domains. IP, TCP and UDP protocols are neccesary to allow for routing and worldwide connection, but bring overhead and unpredictable delays along. With IP, TCP and

UDP, used by COTS network cards and communication stacks, cycle times for information (from the control to the IO-node AND back) are more likely in the area of 20 to 50 ms, even if the IO node is advertised as “fast response in 1 ms”: this is acceptable in process control, but unsuitable in machine control. Cabling installation in collision-free star architecture is not efficient and expensive. Therefore, the utilization of Ethernet in it’s standard form and by use of standard networking equipment needs an entirely different approach.

Exisiting market developments for Ethernet for IO communication:

Rockwell promotes Ethernet IP through ODVA and ControlNet International: Ethernet IP represents the “higher layer” communciation protocol. A complex layer of abstraction allows for powerful device profiling (CIP; Control and Information Protocol); and it’s common Application Object Library is quite far reaching approach to describe automation objects (pump motor, drive, etc.) completely. CIP uses regular TCP/IP or UDP/IP and therefore is limited to these protocol’s abilities to perform on the IO level, e.g. in a 20 to 50 ms cycle. ODVA recommends ControlNet or DeviceNet, if real-time communciation in the single millisecond domain is required – so, no serious intention to convert this business to the Ethernet domain. Ethernet IP uses a CAN similar publisher-subscriber model to distribute information and distributes (by default) broadcast messages to publish information. In case an Ethernet IP installation should be connected somehow to a corporate network, it is necessary to use and configure routers to filter that broadcast traffic; because hubs, switches and unconfigured routers flood the Ethernet IP messages throughout all networks connected – the nature of broadcast. As ARP address resolution relies on broadcasting as well, this router filter configuration can prove tricky and somewhat defeats the idea of “easy to use Ethernet” with Ethernet IP.

So far, only a few products are available; Rockwell, HMS, Applicom, Hilscher, and Mettler Toledo have announced or shown products using Ethernet IP, many of them rather tools than IO or controls.

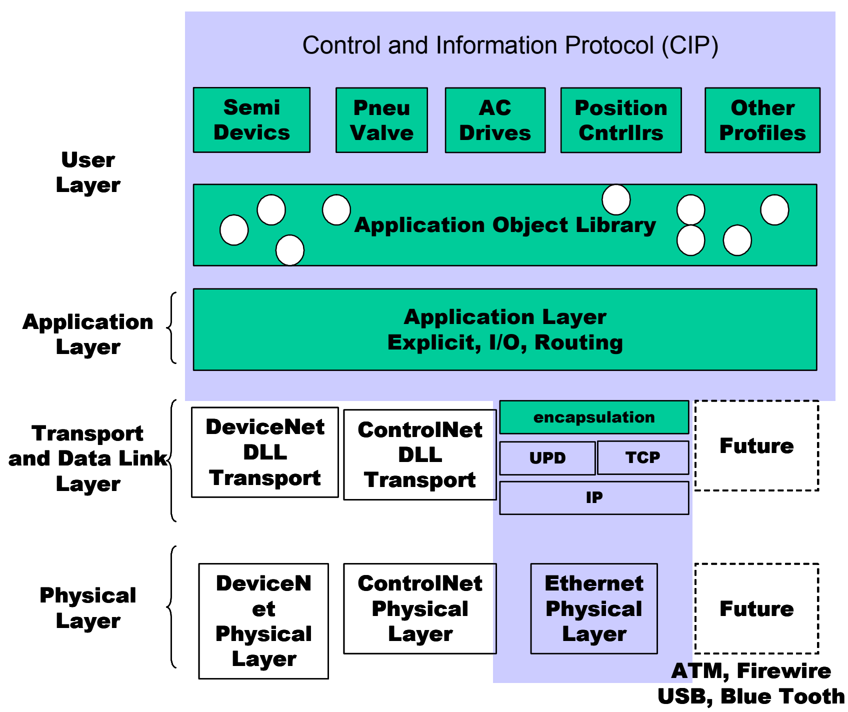

Rockwell: Ethernet IP is using TCP/IP and UDP/IP as protocol basis.

Siemens presents Profinet in various steps as the Ethernet solution of PTO (Profibus Trade Organization): starting from a gateway to Profibus control and IO systems through a proxy system for Ethernet communciating controls, to Profibus on Ethernet to “soft-realtime” ProfiNet to a hard realtime ProfiNet solution.

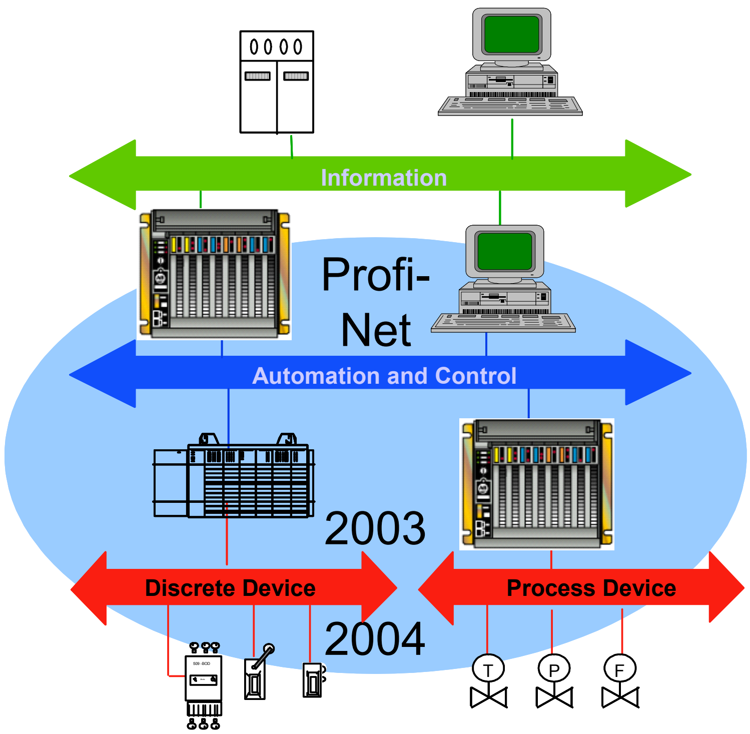

Picture: Siemens Profinet Roadmap: 2003 control to control; 2004 device level; 2005 hard real time

The gateway functionality is based on RPC (remote procedure calls) via TCP or UDP and used DCOM – meanwhile, seen as an intermediate integration tool for Profibus DP installations into Profinet.

Since autumn 2002, Profinet Ver. 2.0 specification is published, covering it’s object and component model, runtime communication, proxy concept, and engineering. So far, Profinet covers the control level, only. Recently, April 2003, Siemens announced “soft realtime” capabilites to come; and for end of 2003, an API for Profibus DP direct IO access services is announced.

Siemens uses it’s own registered Ethertype telegram to communciate directly between controls and for direct IO access. The performance of Profinet in Ethernet environments depends on the delays generated by the overall network setup. Siemens will manufacture Ethernet chips with reduced functionality and isochrone communication features to facilitate IO node integration. The Profinet version 1.2 is based on TCP and does not improve TCP’s performance; the “soft-realtime (SRT)” approach, accounced for 2004, utilizes a modified communciation stack to bypass TPC and IP and uses an “Ethernet for IO” protocol, it includes network management; and depending on stack performance, communication times of 5 to 10 ms and a network-dependent jitter of 15% are expected. In 2005, an Ethernet based hard realtime communciation solution for isochrone traffic is announced (IRT), with 1 µs Jitter and cycle times of 250µs to 1 ms, however requiring special Profinet master cards by Siemens to time-slice between IRT traffic and a remaining TCP portion on the Ethernet wire. Regular TCP traffic for computers or printers may coexist once connected through special manufactured switches which include a gateway to the timeslicing method used to separate Profinet isochrone traffic from regular network traffic. As a side effect; the announced timeslicing switches will allow for a line-structure in installation which is certainly an improvement to star-cabling.

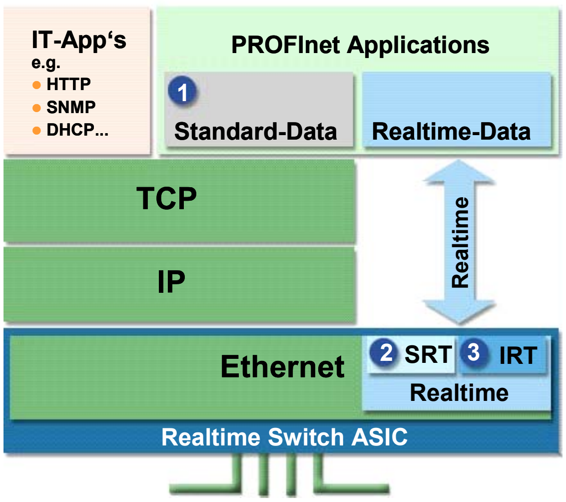

Picture: Siemens Profinet, 1: regular Ethernet traffic; 2: Soft Realtime; modified stack to bypass TCP and IP, 3: Isochronous Realtime; hardware-scheduled Ethernet communication.

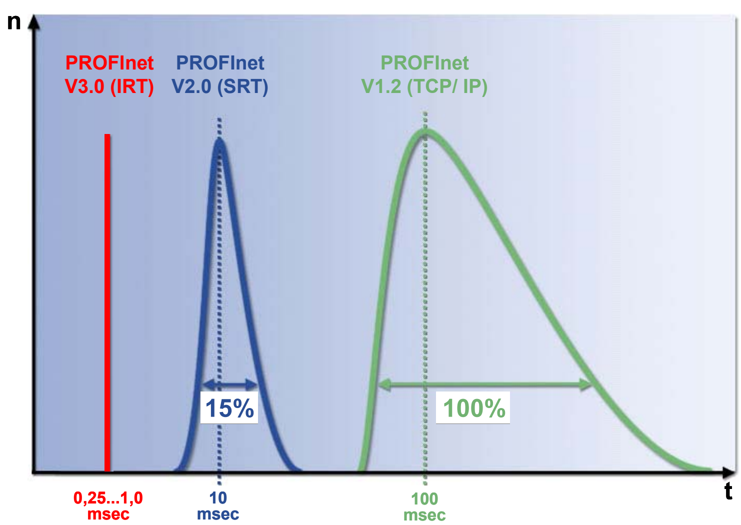

Picture: Siemens performance projections for SRT Profinet and IRT Profinet

Schneider, Phoenix, Lenze, Sick, and others formed the IDA (interface description interest group; using RTI's middleware NDDS using a Publisher – Subscriber model called RTPS on Ethernet. A spec for IDA is in distribution since 11/2001. In short: the implementation was too ressource-hungry and therefore NDDS was dropped; RTI and Lenze left IDA; and Schneider is now focusing back on Modbus/TCP. Modbus is an older protocol for serial IO communciation, and embedded in TCP it is used for the data representation layer. It uses a simple master – slave polling mechanism, lacks complex services and, as it is based on TCP, not intended for hard real-time communication. IDA Member Jetter stated in Hannover 2003 that a 10 ms communication cycle time is targeted as sufficient for control applications by his company. However, the Modbus part is easy to understand, and a variety of devices from various vendors understand Modbus or Modbus/TCP.

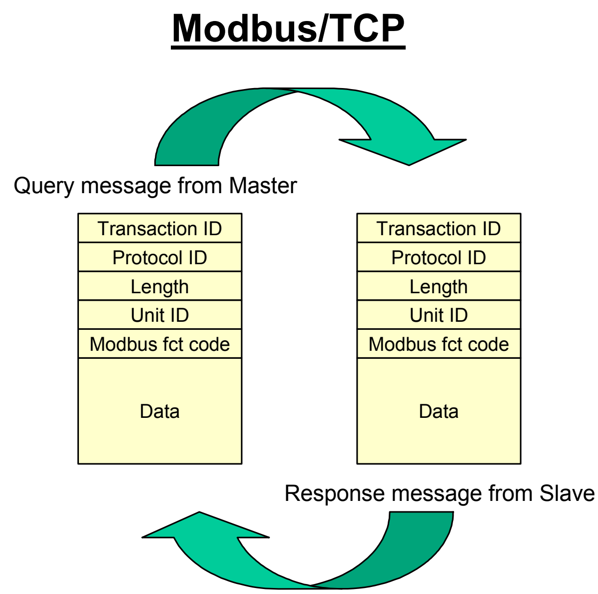

IAONA is an international organization attempting to standardize the various protocol layers for IO communication on Ethernet as an umbrella organization in Europa as well as North America. IAONA cooperates with IDA and ODVA on Ethernet IP; and focuses on areas where these two organizations do not work. IAONA has recently come up with a quite helpful Ethernet installation guideline; however, standardization work shows very slow progress.

B&R offers an Ethernet based IO technology named Ethernet Powerlink with a main purpose to communciate with drives and IO in a realtime environment. Ethernet Powerlink is is meantime supported by a group of companies including end user Kuka (robotics), and drive vendors Lenze and connector manufacturer Hirschmann. Ethernet Powerlink uses Ethernet physical layer and requires hubs, it applies standard fieldbus polling mechanisms similar to Profibus to communicate with slaves, using a fixed time slice pattern for isochronous communciation and an additional time slot for all asynchronous communication (compare to Profinet). The isochronous (deterministic) communication mode works on a private network; a second operation mode on regular networks is available but lacks determinism. Acyclic communication is transferred via TCP or UDP, regular TPC traffic requires the use of gateways or bridges to tunnel this traffic through a Powerlink private network used for determinism. Powerlink uses special scanner cards and does not work with regular NIC cards. Recent publications reference a cycle time of 2.4 ms for 100 nodes (50 drives and 50 IO nodes with each 40 IO signals).

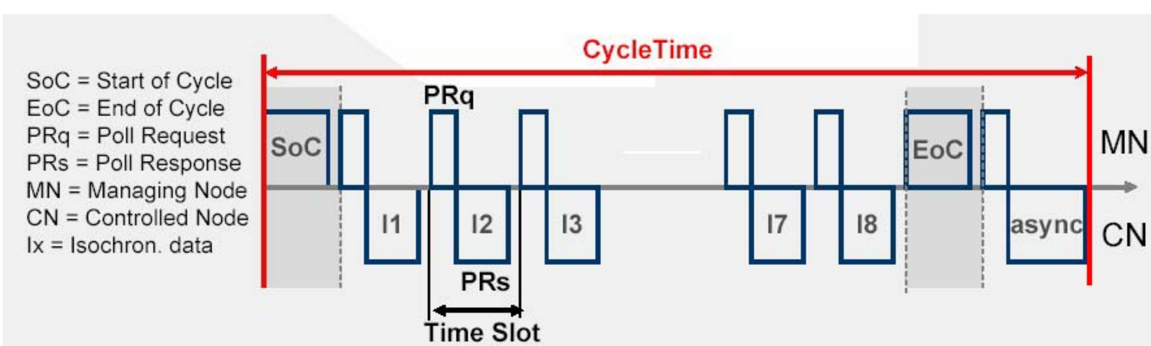

B&R's Powerlink polls slave stations comparable to Profibus and ProfiNet

Beckhoff has made a technology announcements and demonstrated first products in April 2003 for a Ethernet based ultra-highspeed realtime IO communication technology named EtherCAT (Ethernet Control and Automation Technology). Beckhoff designed a new IO system with hardcoded protocol chips in every single IO terminal that read and write IO information directly from and to an EtherCAT telegram – at full speed of 100 Mbit and without buffering the EtherCAT telegram, which is a registered Ethertype telegram. In EtherCAT IO, Ethernet builds the communication technology for the backbone bus of the IO block. Decoding the telegram, output information extraction and insertion of input information is achieved serially (including necessary modifications of the Ethernet CRC block) at full speed, therefore full bandwidth can be utilized for IO information, no time is spent on query - answer protocols. No microprocessors are used for the ultra-highspeed EtherCAT IO telegram decoding and processing. This technology is completely transparent to regular Ethernet traffic as regular Ethernet traffic is ignored by the EtherCAT decoder chips and passed through. Most important, regular NIC cards are utilized, as Beckhoff achieves real-time with extremely low jitter (short term: < 1 µs, long term: +/- 15 µs) with intelligent driver technology for Windows OS on regular NIC cards.

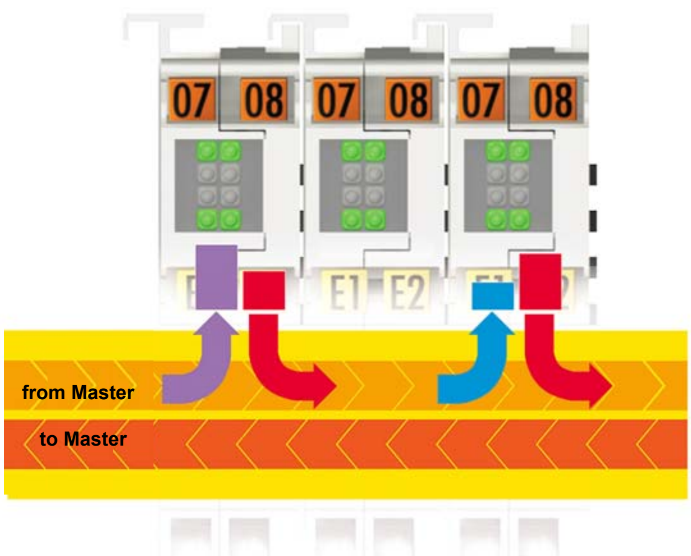

Beckhoff: Ultra Highspeed IO reads and inserts data into Ethernet (EtherCAT) Telegrams at full 100 Mbit speed

The performance of EtherCAT is according to the full utilization of Ethernet 100 Mbit speed by hardcoded protocol chips: cycle times for 1000 IO data exchange are in 15 µs; 100 drives in 100 µs; 200 analog IO (16 Bit) in 50µs; 12.000 IO in 350 µs. most important, Ethernet including 6000 ft. optical transmitters, may be used at any point in the installation without influence to the named performance: the 200 analog IO may be separated each by 6000 ft... No address switches are needed as the system powers up automatically, instead of expensive “switched or hubed” star cabling, simple line cabling schemes, or tree and branch schemes may be used. Redundant cabling is on the roadmap. Beckhoff transforms the physical layer to E-Bus, a noise immune low voltage differential signal on ther terminal level and up to 30 ft. distance, thus improving the noise immunity. At any time, reconversion to Ethernet physical media is possible.

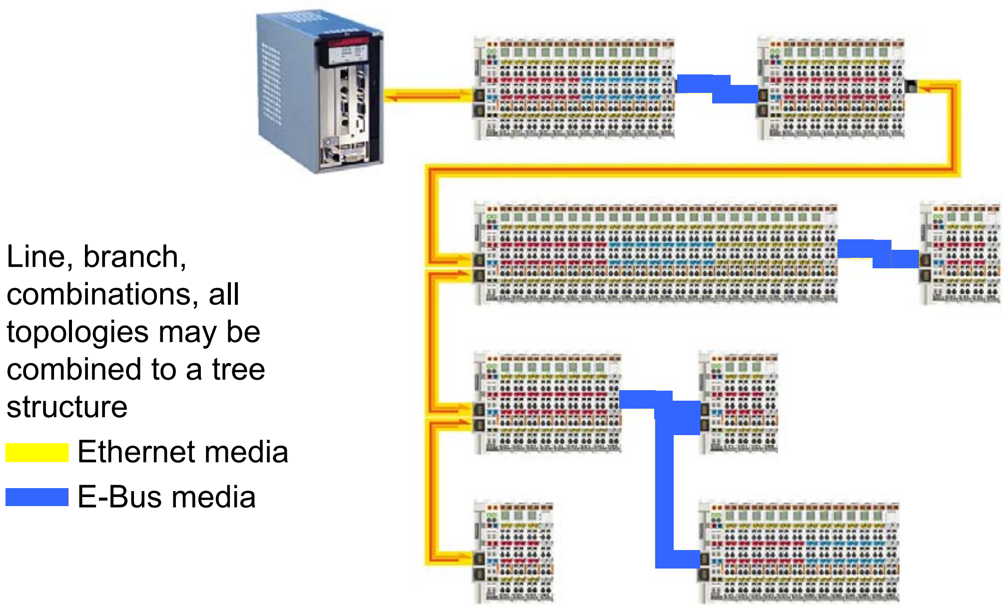

Beckhoff EtherCAT allows to cable in any of line, branch and tree structures.

The address size for a fieldbus terminal starts at 2 Bit up to 64 kBye per terminal (!) and up to 4 Gbyte per EtherCAT system (32 bit address identifier). In order to avoid any overhead communciation for mapping and addressing schemes during the process data exchange, the system loads mapping information into the terminals at startup: as a result, the controller exchanges it's process images, and the terminals identify the correct position of information in the logical process image by comparing the address like a Memory Management Unit. As a result, the location of information in the telegram is of no importance; an EtherCAT telegram with its 46 to 1500 Byte user information may include several subtelegrams if desired. With this method, individual scan times per terminal may be achieved. Also, the effective data rate for read or write events can be above 95% of the Ethernet maximum bandwidth (!). As inputs and outputs may utilize the same location in the telegram, in one cycle a 1500 byte fullsized telegram may transport 1500 bytes output and return 1500 bytes input: the effective data rate is above 100%!

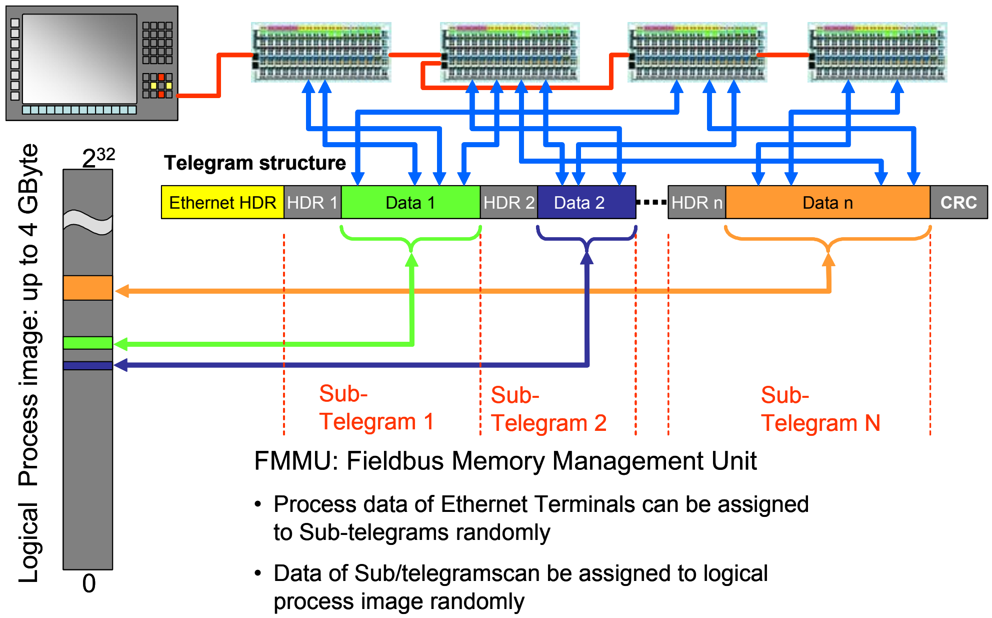

EtherCAT can address 4 Gbyte IO information. The terminals store mapping information and need no complex telegram to insert and extract data.

Determinism is achieved by usign advanced driver technology, controling the moment of telegram transmission to the µs on regular Ethernet NIC cards. During the real-time slot, regular traffic is buffered and may use the EtherCAT media in periods with no deterministic traffic.

EtherCAT uses Beckhoff RTEthernet software technology, providing real-time communciation on standard NIC cards. The driver executes the EtherCAT traffic deterministically, holding other traffic in a buffer until there is time to pass it on. To improve the timing on the EtherCAT system and to hold the jitter of the system always below 1% of the faster and faster cycle time, synchronized distributed clocks with 10 nanosecond resolution are provided. Also, oversampling data types allow for fastest data acquisition rates even beyond the 100 Mbit transmission bandwidth. For future systems, Gigabit compatibility is inherent…

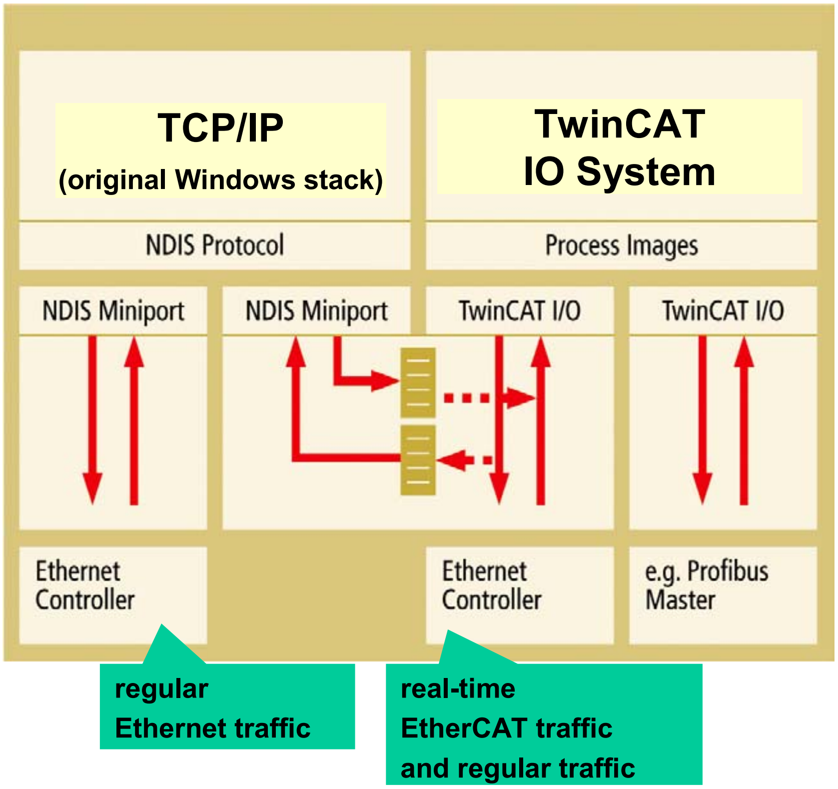

Compatibility to existing Fieldbus systems is achieved by gateway terminals to Profibus, DeviceNet, etc., that build an entire scanner card with an EtherCAT interface. Buscouplers for EtherCAT guarantee compatibility for existing BusTerminals from Beckhoff, which can be used along in the new system.

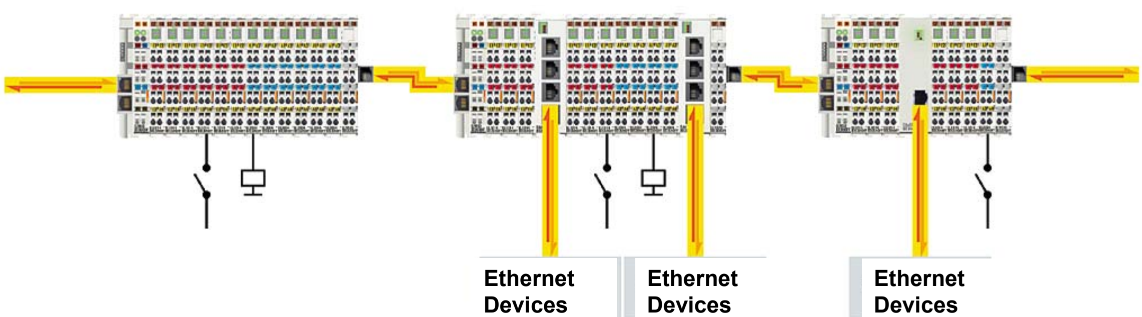

EtherCAT allows for integration of regular TCP/IP devices at any point

If the system needs to be remotely located in a regular network, EtherCAT allows for a less performant UDP/IP based mode; in which only one single MAC ID for the first coupler is needed: the protocol frames are routable and can travel worldwide.

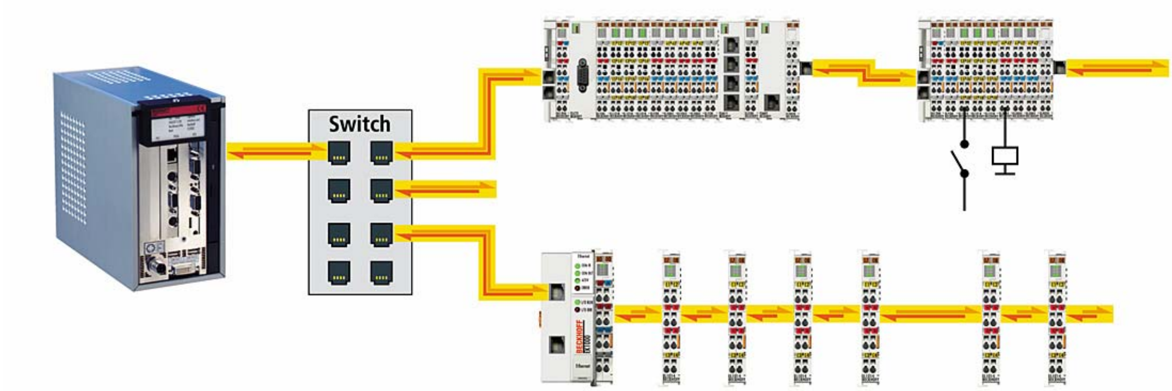

EtherCAT is routable with help of couplers with MAC Id.

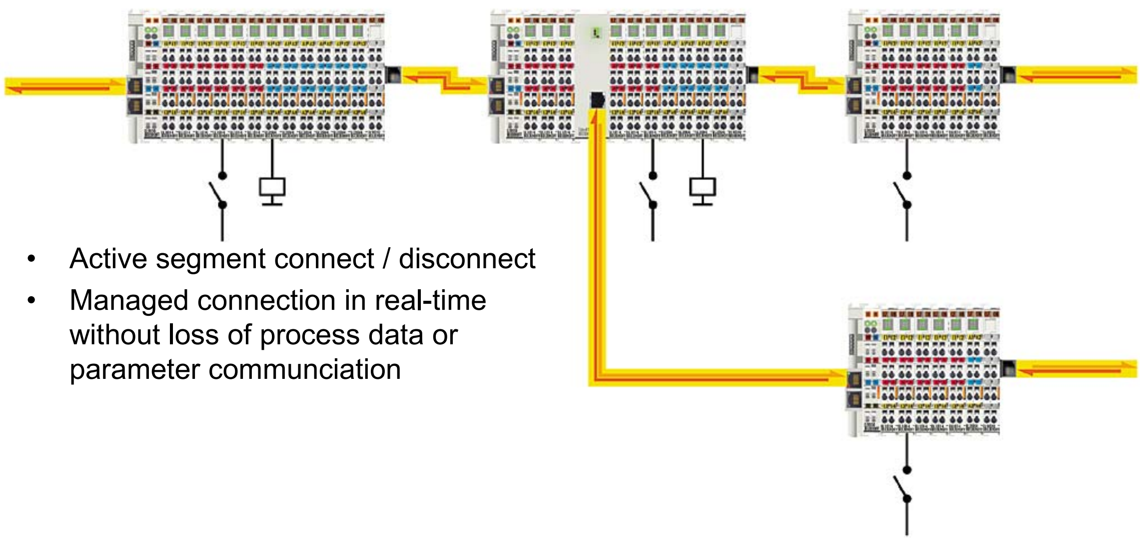

TwinCAT allows for hot connect of dedicated segments…

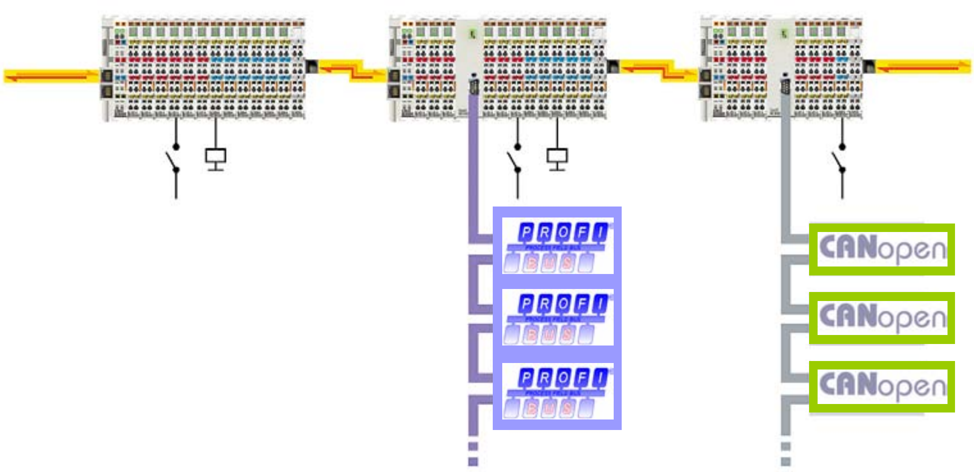

and integrates existing fieldbus systems seamlessly by high performance scanner terminals

All of that comes at the lower cost of COTS Ethernet: scanner cards, cabling, etc. all come at market price. IO will be as competitive as the worldwide accepted BusTerminal system.

With the presented approach, Beckhoff provides full utilization of Ethernet by bringing Ethernet into the backbone of IO terminals at full speed of 100 Mbit/s. All othern system use Ethernet as media but re-use the strucutures of the 80's fieldbus world. Decades of experience on more than 18 networks, starting with Lightbus in 1989; resulted in mciroprocessor-free, storage-free data exchange by the powerful Memory Management capabilties of the terminal electronics. Therefore, Beckhoff's approach leapfrogs the performance of industrial IO and provides a system with so enormous gains in throughput and reduction of communciation delays, that for most applications, EtherCAT IO can be considered to be equal to direct boardlevel IO.