

TTEthernet – A Powerful Network Solution for All Purposes

Ethernet is still gaining ground although it has been around for more than 30 years. With Ethernet being used as a universal network solution in office and web applications, production facilities, safety-critical systems, airplanes and automobiles, it has a big potential for cost savings in all areas. Engineering, maintenance and training costs are considerably lower than those for many proprietary bus systems. Additionally, Ethernet offers far higher bandwidths. But when Ethernet was developed in the first place, tasks with time-critical, deterministic or safety-relevant conditions were not taken into account. TTEthernet expands classical Ethernet with powerful services to meet all new requirements.

From Ethernet to TTEthernet

Over the last few years there has been a lively discussion about how to adapt Ethernet for new application domains. Its focus has been on the use of Ethernet for real-time tasks in industrial environments (Industrial Ethernet). More than 20 different approaches are now struggling for recognition in industrial automation alone. Competition in this area has grown with the design of additional solutions such as Safety Ethernet that are intended to meet the requirements of engine and plant construction. There also has been made an effort to adapt Ethernet for special requirements in other areas, e.g. LXI in measurement technology or AFDX in the aerospace industry.

Today's Ethernet systems have limits when it comes to combining them with classical Ethernet networks, devices and services. The scalability of these systems is also limited and the network solution is tailored for a specific application area. The use of Real-Time and Safe Ethernet systems outside engine and plant construction is no option. The industry is striving to reduce the number of networks and to cut costs for effort and resources. TTEthernet combines the proven determinism, fault-tolerance and real-time properties of the time-triggered technology with the flexibility, dynamics and legacy of "best effort" of Ethernet and is therefore suited for all types of applications.

Considering networking technologies in general, one can distinguish between closed, statically configured embedded networks and open, dynamic networks allowing free-form communication. Whereas statically configured communication networks enable reliable data transmission in real time (to different extents), free-form communication networks perform only on a best-effort basis, i.e., there is no guarantee if and when data messages are transmitted. As statically configured systems usually need to comply with strict safety are based on dedicated standards, the number of communication nodes is known and not very flexible. Open standards, like the TCP/IP/Ethernet stack that drives the Internet, form the basis of free-form communication and the number of nodes in systems is arbitrary. TTEthernet brings together the high flexibility of free-form systems and the reliability and speed of statically configured systems (see Figure 1).

Scalable Real-Time Ethernet Platform

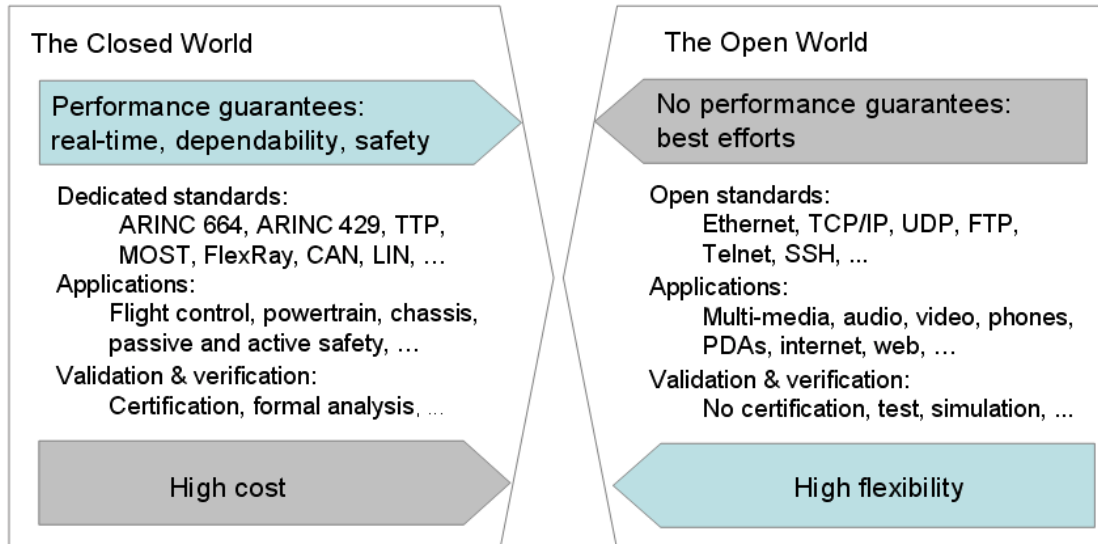


Fig. 1: TTEthernet combines closed-world and open-world systems on the basis of IEEE 802.3 Ethernet standards

TTEthernet Design Objectives

TTEthernet enables the seamless communication of all applications by way of Ethernet. Conventional PCs, web and office devices, multimedia systems, real-time systems and safety-critical systems are to use the same network. A single network that is completely compatible with the IEEE Ethernet 802.3 standards is suited for data transmission among different applications with various requirements. A single network solution could thus be used for all applications in airplanes, ranging from the entertainment program, to board supply, electronic navigation and guidance system, and internet access in passenger seats. Critical areas are accordingly made fail-safe or fail-operational. Fault tolerance mechanisms avoid the fault propagation in the system and prevent potential hackers from unauthorized access to resources.

TTEthernet is scalable. Networks that now connect uncritical applications shall be able to transmit real-time data in distributed controls and shall be suited for safety-critical applications in the future. Existing applications need not be changed when the network is extended in terms of functionality. Time-critical messages always take precedence over less important messages in TTEthernet. This does not affect conventional applications. The temporal behavior of the time-critical messages is predictable (deterministic) and can be characterized depending on the required quality.

TTEthernet is used for safety-critical fail-operational applications. This means that the system remains fully functional even if a failure occurs (supporting a single or double fault hypothesis). No matter if a node, a switch or a network branch is faulty, the network continues safe communication. This fact accounts for the essential difference between TTEthernet and other Safe Ethernet systems. The latter systems, which are

Scalable Real-Time Ethernet Platform

Page 3

sufficient for industrial applications (fail-safe), detect faults in the network and switch the system to a safe state, e.g. stopping the engine. In order to secure the availability of the system even if a failure occurs, TTEthernet provides a variety of network services such as a clock synchronization service, a startup service and clique detection and recovery services. The behavior of TTEthernet is precisely predictable and thus formally verifiable.

TTEthernet System Properties

TTEthernet has time-triggered services that enable time triggered communication over Ethernet. These time-triggered services establish and maintain a global time, which is realized by the close synchronization of local clocks of the devices. The global time forms the basis for system properties such as temporal partitioning, precise diagnosis, efficient resource utilization, or composability.

Temporal Partitioning: The global time can be used as a powerful isolation mechanism when devices become faulty; we say that the global time operates as a “temporal firewall”. In case of failure it is not possible for a faulty application to untimely access the network. Depending on the location of the failure, either the communication controller itself or the switch will block faulty transmission attempts. Failures of the switch can be masked by powerful end-to-end arguments such as CRCs or by high-integrity designs.

Efficient Resource Utilization: The global time contributes to efficient resource utilization in several ways. Time-triggered communication allows minimizing the memory buffers in network devices as the time-triggered communication schedule is free of conflicts. Hence, switches do not have to be prepared for bursts of messages that have to be delivered over the same physical link. A minimal time-triggered switch design could even multiplex media access logic such as reception or transmission logic. A second way of effective resource utilization is buffer memory in the nodes, which can be minimized as the sensor values can be acquired according to the global time, immediately before sending the message. Finally, a third way of effective resource utilization is power management in which energy can be seen, and saved, analogously to memory.

Precise Diagnosis: A global time stamping service simplifies the process of reconstruction of a chain of distributed events. On the other hand, the synchronous capturing of sensor values allows building snapshots of the state of the overall systems.

Composability: The global time allows the specification of devices not only in the value domain, but also in the temporal domain. This means that already during the design process of devices, the access pattern to the communication network can be defined. The devices can then be developed in parallel activities. Upon integration of the individual devices, it is guaranteed that prior services are stable and that the individual devices operate as a coordinated whole.

Scalable Real-Time Ethernet Platform

Dataflow Options in TTEthernet

TTEthernet specifies services that enable time-triggered communication on top of Ethernet. The time-triggered services can be viewed parallel to the usual OSI layers: a communication controller that implements these services is able to synchronize with other communication controllers and switches in the system. The communication controller can then send messages at points in time derived from this system-wide synchronization. These messages are then called time-triggered messages.

As TTEthernet supports communication among applications with various real-time and safety requirements over a network, three different traffic types are provided: time-triggered (TT) traffic, rate-constrained (RC) traffic, and best-effort (BE) traffic. If required, the corresponding traffic type of a message can be identified based on a message's Ethernet Destination address. The relation of the TTEthernet traffic types to existing standards is depicted in Figure 2.

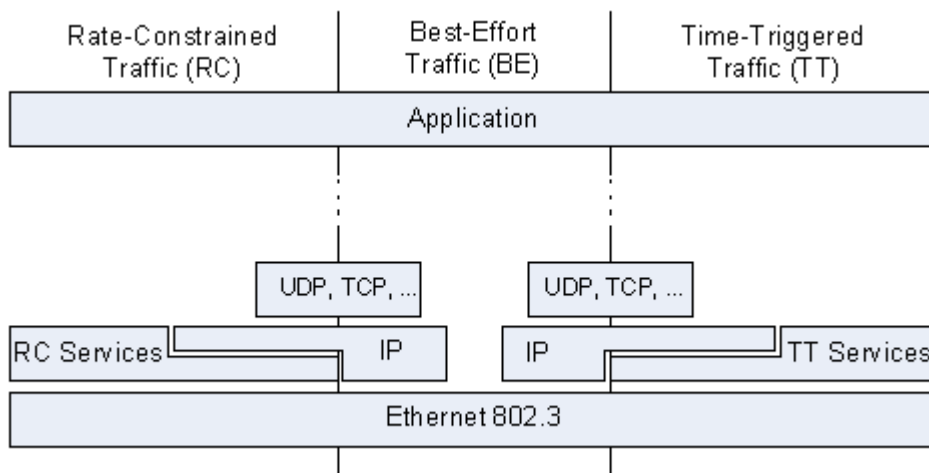


Fig. 2: Relation of TTEthernet to existing communication standards

Messages from higher layer protocols, like IP or UDP, can be “made” time-triggered without modifications of the messages’ contents itself. The TTEthernet protocol overhead is transmitted in dedicated messages termed protocol control frames, which are used to establish system-wide synchronization. In short, TTEthernet is only concerned with “when” a data message is sent, not with specific contents within in a message.

TT messages are used for time-triggered applications. All TT messages are sent over the network at predefined times and take precedence over all other traffic types. TT messages are optimally suited for communication in distributed real-time systems. TT messages are typically used for brake-by-wire and steer-by-wire systems that close rapid control loops over the network. TT messages allow designing and testing strictly deterministic distributed systems, where the behavior of all system components can be specified, analyzed and tested with sub-micro second precision.

Scalable Real-Time Ethernet Platform

RC messages are used for applications with less stringent determinism and real-time requirements than strictly time-triggered applications. RC messages guarantee that bandwidth is predefined for each application, and delays and temporal deviations have defined limits. RC messages are used for safety-critical automotive and aerospace applications that depend on highly reliable communication and have moderate temporal quality requirements. Typically, RC messages are also used for multimedia systems.

In contrast to TT messages, RC messages are not sent with respect to a system-wide synchronized time base. Hence, different communication controllers may send RC messages at the same point in time to the same receiver. As a consequence, the RC messages may queue up in the network switches, leading to increased transmission jitter. As the transmission rate of the RC messages is bound a priori and controlled in the network switches, an upper bound on the transmission jitter can be calculated off-line and message loss is prevented.

BE messages follow a method that is well-known in classical Ethernet networks. There is no guarantee whether and when these messages can be transmitted, what delays occur and if BE messages arrive at the recipient. BE messages use the remaining bandwidth of the network and have less priority than TT and RC messages. Typical user of BE messages are web services. All legacy Ethernet traffic (e.g. internet protocols) without any QoS requirement can be mapped to this service class. TTEthernet implements strong partitioning between non-critical BE traffic and all other service classes (see Figure 3).




Fig. 3: TTEthernet includes TT, RC and BE messages

TTEthernet as Transparent Synchronization Protocol

TTEthernet is a transparent synchronization protocol, i.e., it is able to co-exist with other traffic, potentially legacy traffic, on the same physical communication network. For reasons of fault tolerance a multitude of devices can be configured to generate synchronization messages. The devices generating the synchronization messages may be distributed with a high number of intermediate devices in between each other.

TTEthernet defines basic building blocks that allow the transparent integration of the time-triggered services on top of message-based communication infrastructures such as standard Ethernet. For this, TTEthernet defines a novel application of the transparent clock mechanism that enables the concept of the permanence point in time, which allows re-establishing the send order of messages in a receiver:

Scalable Real-Time Ethernet Platform

Page 6

- Application of transparent clock mechanism: all devices in the distributed computer network that impose a dynamic delay on the transmission, reception, or relay of a synchronization message add this dynamic delay into a dedicated field in the synchronization messages used for the synchronization protocol.
- Novel precise calculation of the permanence point in time: the application of transparent clock mechanism allows a precise re-establishment of the temporal order of synchronization messages. In a first step the worst case delay is calculated off-line. In a second step, each synchronization message is delayed for "worst case delay minus dynamic delay" upon reception of the synchronization message, where the dynamic delay is the delay added to the synchronization message, as the synchronization message flows through the communication channel. This point after the reception point in time will be called the permanence point in time.

For fault-tolerant algorithms in general, and fault-tolerant synchronization algorithms in particular, the message send order is of highest importance. The re-establishment of the send order of synchronization messages is required for any fault-masking synchronization protocol that ensures synchronization of local clocks in a distributed computer network.

Safety and Fault Tolerance

A high level of safety is provided by the time-triggered method of TTEthernet, which detects failures and irregularities in the network and certain systems. Additional measures need to be taken to achieve maximum safety, availability and fault tolerance.

TTEthernet networks can be set up with multiple redundant end systems, switches and segments. Thus the system will remain in operation even if faults occur. Redundant network paths are always used in fault-tolerant TTEthernet systems so that the failure of a single system or messages can be tolerated without affecting the application. If multiple redundancy is implemented, multiple faults can be tolerated. It is important that the entire system remains in operation without interrupts under the same temporal conditions as defined before.

TTEthernet allows the integration of guardians in switches and end systems. Guardians check if the communication on the network works in compliance with the predefined parameters. If faulty systems block network segments, the guardian disconnects the network segment or port. Multiple redundant guardians can be implemented to meet the highest safety requirements (see Figure 4).

Scalable Real-Time Ethernet Platform

Page 7




Fig. 4: TTEthernet provides implicit fault tolerance mechanisms

Fault-Tolerant Capabilities

TTEthernet is designed to scale over a multitude of cross-industry applications. As such, TTEthernet comprises demanding fault-tolerant capabilities.

- TTEthernet is scalable: TTEthernet can be configured to operate as a simple master-slave synchronization protocol for industrial control or a multi-master synchronization protocol for civil avionics. This scalability gives a vast economic benefit because the cost of the realization of TTEthernet throughout different application domains can be decreased significantly. Likewise, the cross-domain usage of TTEthernet increases the probability of latent failure detection in the realization of TTEthernet and contributes to the “service history” of TTEthernet, when deployed in systems with a comparable level of criticality.
- TTEthernet tolerates multiple inconsistent faults: When configured to a multi-master mode, TTEthernet tolerates a fully inconsistent-omission faulty communication path and even an inconsistent-omission faulty end system at the same point in time. This failure mode means that each faulty device can arbitrarily drop messages on any of its incoming communication links and on any of its outgoing communication links with potential inconsistent dropping behavior for each message. TTEthernet therefore allows a more cost-efficient realization of system architectures that require tolerance of multiple concurrent failures in the system.

Scalable Real-Time Ethernet Platform

Page 8

- TTEthernet tolerates arbitrary end system failures: The switches in TTEthernet can be configured to execute a central bus guardian function. The central bus guardian function guarantees that even if a set of end systems becomes arbitrarily faulty, the system-wide impact of these faulty end systems is masked. The arbitrarily faulty failure mode also includes babbling idiot behavior and similar failure modes. TTEthernet switches establish fault containment boundaries.
- TTEthernet tolerates arbitrary transient disturbances even in presence of permanent failures: In addition to fault tolerance, TTEthernet also provides self-stabilization properties, i.e., the synchronization will be re-established even after transient upsets in a multitude of devices in the distributed computer system. TTEthernet stabilizes from an arbitrary system state to a synchronized system state. This self-stabilizing property becomes more and more important with decreasing feature sizes of computer chips and, therefore, resulting increase in transient upsets. The design of future reliable distributed computer networks depends on an effective and sound tolerance of multiple transient upsets.

Network Structure

TTEthernet supports all physical layers specified in IEEE 802.3 for switch-based networks. Even sub-networks with different bandwidths (100 Mbit/s, 1 Gbit/s, ...) are supported.

Switches in TTEthernet have the central role of organizing the data communication. TT messages are routed in the switch according to a predefined schedule with as little delay as possible. Precise planning at the time of system design precludes resource conflicts at runtime. TT messages have the highest priority level. If the planned transmission time of one of these messages arrives, this message is immediately transmitted. Due to the predefined transmission of the message the switch ensures that the medium is free at the time of transmission and delays are precluded.

RC messages are routed with little delay. If TT messages are to be transmitted via the same outgoing port at the same time, the TT messages take priority over the RC messages. TT messages can delay RC messages. RC messages are transmitted if no planned transmission of TT messages is pending and the sender observes the minimal transmission distance. The switch is responsible for arranging several RC messages at an outgoing port.

BE messages always have little priority. RC and TT messages can delay or discard BE messages at the same outgoing port. The switch uses the remaining bandwidth for BE messages if no TT or RC messages are to be transmitted. BE messages are transmitted after all pending RC messages. This method exploits the bandwidth of the network in an optimal way. Tools are used to design and verify a TTEthernet system in advance. This ensures that the bandwidth for TT, RC and BE messages is always sufficient according to the requirements of the application and interrupts are reduced to a minimum. Later incremental changes of the system configuration are possible.

Scalable Real-Time Ethernet Platform

Page 9

TTEthernet switches allow the simultaneous distribution of TT messages to groups of end systems or the connection of unsynchronized TTEthernet networks. This is how TTEthernet networks can be divided into smaller application-specific sub-networks and the design can be facilitated.

Supported Topologies

TTEthernet allows synchronizing local clocks in a distributed computer network. Of particular interest are computer networks that exchange information via messages that are sent on communication links between devices in the network. In standard Ethernet end systems are connected via network switches via bi-directional communication links. An end system will communicate with a second end system or a group of end systems via sending a message to the switch, which will then relay the message to the receiving end system or end systems. Also, switches can be connected to each other via bi-directional communication links. In this case the resulting architecture is referred to as a multi-hop architecture and the links between any two switches as the multi-hop link.

Communication links and switches are said to form a communication channel between end systems. End systems can be connected directly to each other via bi-directional communication links, which makes a clear differentiation between end systems and switches in certain configurations difficult. Hence, generally we use the term device to refer to a physical device that can be either end system or switch. Whether a device is regarded as an end system or a switch is determined by its usage rather than its physical appearance.

Synchrony

Events in time-triggered systems occur at predefined times with a precision at the single microsecond level. This also includes the communication of TT messages. The system design specifies when the TT messages are transmitted by which participants and who shall receive them. This ensures that the network processes TT messages without collisions (i.e. without data congestion in the switches) and the recipient can continuously check the quality of the deterministic system if, for example, a message fails to arrive at the predefined time or does not arrive at all. This makes TTEthernet suited for applications of the highest safety integrity level.

Synchrony among all participants is crucial for the transmission of TT messages. TTEthernet always transmits clock synchronization messages to keep the clocks of the end systems and switches in synchrony. For this purpose TTEthernet relies on a redundant hierarchical master-slave method that has a distributed fault-tolerant majority of master nodes and master switches to provide the time in the system. This guarantees both the fail-safe operation and the high quality in synchronization. This method is unique for TTEthernet and can be combined with other mechanisms such IEEE 1588 (see Figure 5).

Scalable Real-Time Ethernet Platform

Page 10




Fig. 5: TTEthernet allows for flexible, even redundant network topologies and synchronization with other systems such as IEEE 1588

IEEE 1588 specifies a synchronization protocol for Ethernet. The global time base of TTEthernet can be leveraged to synchronize native IEEE 1588 synchronization clients, too. For this purpose, additional functionality can be realized on top of a TTEthernet device that generates IEEE 1588 clock synchronization frames. TTEthernet provides means to compensate for delays through the TTEthernet network. Outside the TTEthernet network, in a native IEEE 1588 network, the clock synchronization messages can be handled as native IEEE 1588 clock synchronization messages.

Synchronization Approach

TTEthernet takes a two-step approach to synchronization. In the first step synchronization masters send protocol control frames to the compression masters. The compression masters then calculate an averaging value from the relative arrival times of these protocol control frames and send out a new protocol control frame in a second step. This new protocol control frame is then also sent to synchronization clients.

The decision on which devices are configured as synchronization masters, synchronization clients, and compression masters arises from the requirements on the system architecture. End systems can be configured as synchronization masters and switches as compression masters. But system configurations with end systems configured as compression masters and switches as synchronization masters are also possible. Switches and end systems not configured either as synchronization or compression masters will be configured as synchronization clients (see Figure 6).

Scalable Real-Time Ethernet Platform

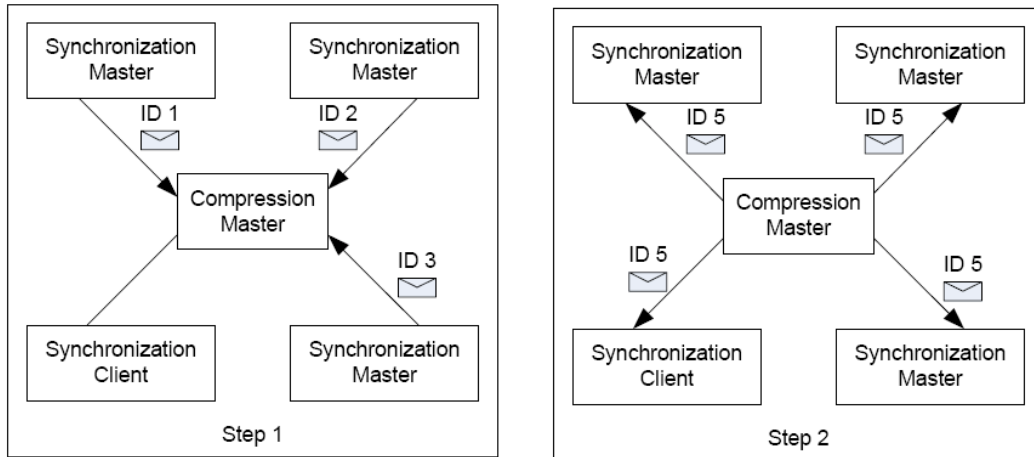


Fig. 6: TTEthernet follows a two-step synchronization approach

Synchronization Topology

TTEthernet distinguishes four different levels in synchronization topology. On the lowest level, TTEthernet defines the device level that comprises synchronization masters, synchronization clients, and compression masters. The cluster level groups devices with the same synchronization priority and the same synchronization domain to a single cluster. On the multi-cluster level, several clusters with different synchronization priorities but same synchronization domain are grouped together. Finally, the network level groups different clusters (potentially multi-clusters) with different synchronization priorities and different synchronization domains (see Figure 7).

Network Level	y Synchronization Domains, (y,z) Synchronization Priorities
Multi-Cluster Level	One Synchronization Domain, x Synchronization Priorities
Cluster Level	One Synchronization Domain, One Synchronization Priority
Device Level	Synchronization Masters, Synchronization Clients, Compression Masters

Fig. 7: The TTEthernet synchronization topology has four levels

TTEthernet specifies the concept of a cluster. A TTEthernet cluster is a group of end systems and switches that have the same synchronization priority and synchronization domain. TTEthernet clusters could be used in large TTEthernet networks, where different clusters shall be able to run in isolation, but shall be able to operate in a master-slave mode, once a high priority cluster joins the network or is powered on.

Scalable Real-Time Ethernet Platform

A TTEthernet simple cluster consists of a set of end systems that are connected to each other via an optionally redundant set of communication channels, where each communication channel consists of one switch only (see Figure 8). In a TTEthernet cascaded cluster configuration each communication channel consists of more than one switch (see Figure 9).

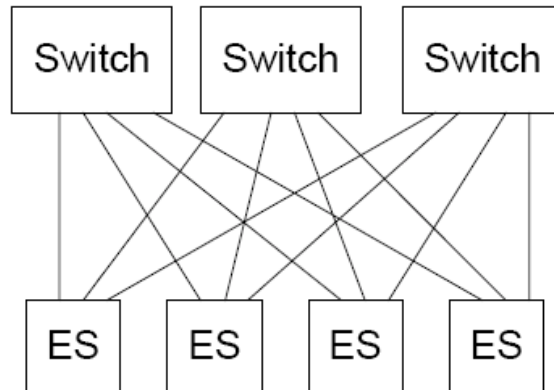


Fig. 8: TTEthernet simple cluster with three redundant channels




Fig. 9: TTEthernet cascaded cluster with three redundant channels

TTEthernet specifies different synchronization priorities. Synchronization in a multi-cluster system is usually done according to a master-slave paradigm, where the devices will synchronize towards the highest synchronization priority. TTEthernet also specifies different synchronization domains. A synchronization domain is a group of TTEthernet clusters that will not synchronize to each other. However, dataflow between two TTEthernet clusters in different synchronization domains can be done using RC or BE traffic.

Variable Implementation

A TTEthernet integration for end systems can be implemented in hardware or software, depending on such requirements as temporal quality, safety and fault tolerance. A TTEthernet system can always be connected to conventional Ethernet systems without affecting the predefined behavior. But there might be a lack of bandwidth for this additional system.

Scalable Real-Time Ethernet Platform

Even standard PCs can participate in a TTEthernet system. These scenarios are possible:

- A PC with a conventional Network Interface Card (NIC) can send and receive BE messages. Equipped with dedicated software the PC can also receive and analyze TT and RC messages.
- A PC with conventional NIC and a TTEthernet stack is a software-based end system (SES) that allows the reception and transmission of TTC, RC and BE messages. But the PC software is the limiting factor to the temporal precision.
- A PC with specific TTEthernet NIC can send and receive TT, RC and BE messages with the highest temporal precision.

Non-PC-based embedded systems can implement TTEthernet. This can happen in stack-based software on standard Ethernet hardware or in dedicated hardware controllers.

This broad implementation freedom is a result of TTEthernet's compatibility to the Ethernet standard as only Ethernet messages are used in TTEthernet. However, there is a natural trade-off between the implementation options and the temporal quality of TTEthernet. A TTEthernet protocol stack that is implemented on a standard PC with standard NICs may, for example, achieve a precision in the order of hundred microseconds, while a dedicated hardware implementation will come down to a one digit microsecond precision and below. Still, the lower temporal quality arising from standard Ethernet controllers is sufficient for a multitude of real-time control processes. In both cases the deterministic properties of time-triggered systems can be maintained.

Available Products

^{TT}Evaluation Systems are available as 100 Mbit/s and 1 Gbit/s solutions. They allow the investigation of TTEthernet features and the development of distributed real-time applications. In order to create and modify the time-triggered schedule, ^{TT}E-Tools are also part of the evaluation system packages.

These tools cover the entire life cycle of the network. Automatic and manual modeling tools allow the intuitive system design in terms of temporal behavior, network and topology. ^{TT}EPlan and ^{TT}EBuild generate configuration data that comply with the communication schedule. Later ^{TT}ELoad loads the data into the involved systems. The tools can be integrated as plug-ins into the open Eclipse framework. ^{TT}EMonitoring Switches und ^{TT}EView display the network traffic on-line and off-line. ^{TT}EVerify checks the accuracy and consistency of a designed system including the temporal behavior of TT and RC messages. It also generates detailed reports for approving a system in compliance with application regulations such as DO-178B in the aerospace industry (see Figure 10). Open XML data exchange formats allow the simple and seamless integration with third-party tools.

Scalable Real-Time Ethernet Platform




Fig. 10: The TTEthernet tool chain covers the entire life cycle of a network

To define and develop a TTEthernet system there are several switches available: ^{TTE}Development Switch supporting 1 Gbit/s and 2 or more channels and another version supporting 100 Mbit/s and 1 channel. The ^{TTE}Monitoring Switch enables monitoring of data streams. The ^{TTE}Monitoring System serves as lab test equipment. It is designed to support recording of up to 30 minutes with 1 Gbit/s traffic.

Network interface cards are available for the 1 Gbit/s variant of TTEthernet. Supported form factors are PMC and PCI Express. Software-based clients are supported for the 100 Mbit/s variant of TTEthernet.

Standardization Activities

At present, SAE AS-2D plans to standardize TTEthernet, and to work closely with other standardization bodies in their respective industry (e.g. ISO, SAE J, IEEE and others) with target date 2012.

Conclusion

TTEthernet enables time-triggered communication over Ethernet networks in all application areas. The network provides all necessary mechanisms for applications as diverse as classical web services and time-critical and safety-critical control system in airplanes. Existing networks can be extended step by step using TTEthernet-capable switches and end systems without the need to change existing applications and end systems. Reducing network solutions to established and recognized Ethernet standards opens up saving potentials that secure major advantages in competitive markets. Honeywell is the first company to use

Scalable Real-Time Ethernet Platform

Page 15

TTEthernet for production programs in the aerospace and automation industries. TTEthernet will thus be used not only in extremely demanding aerospace applications but also in completely new application areas.

Contact

TTTech Computertechnik AG
Schoenbrunner Strasse 7
A-1040 Vienna, Austria
Tel.: +43 1 585 34 34-0
Fax: +43 1 585 34 34-90
E-mail: office@tttech.com
Web: www.tttech.com