



WHITE PAPER

Vision-Critical Networked Video

The Exceptional Benefits of Real-Time Video Transmission over Gigabit Ethernet

This whitepaper examines the benefits of Gigabit Ethernet (GigE) networked video, illustrates some of its key characteristics in the context of a number of primary applications, and draws conclusions as to its ease of implementation and future development.

Pleora Technologies Inc.
www.pleora.com

Updated: September 25, 2017

Table of Contents

- Executive Summary..... 1
- Introduction..... 2
- The Importance of Networked Video..... 2
- The GigE Vision Standard..... 4
- Vital Statistics of Real-Time Communication..... 5
- Vision-Critical Applications..... 8
- The Future of Ethernet..... 11
- Conclusions..... 12
- References..... 13
- About Pleora..... 13
- Appendix A - Comparison of Vision System Interfaces..... 14
- Appendix B - Ethernet and GigE Vision..... 16
- Appendix C - Ultra-Low Latency and Jitter in GigE Vision Systems..... 19

Executive Summary

Whether a doctor performs a life-saving operation using image-guided surgery, or a pick-and-place robot precisely positions a component on a fully automated assembly line, each requires instant verification. This kind of real-time visual confirmation via video is now found in an ever-expanding range of demanding applications, all of which rely on decisions being made in fractions of a second. Fully networked vision systems that use Gigabit Ethernet (GigE) provide both high performance and flexibility, and GigE's fast growing industry adoption proves the point.



While various vision system architectures are available, one of the factors behind the popularity of GigE is its universal deployment in practically every consumer and commercial network around the globe. Its high level of maturity means that off-the-shelf equipment is widely available, very affordable, and exceptionally easy to install and maintain. For real-time video applications, the globally recognized GigE Vision@standard has further facilitated the use of Ethernet by providing a common frame of reference that is fully compatible and tightly integrated with numerous international standards.

Imaging device manufacturers and system integrators serving the industrial automation, medical imaging, military imaging, and intelligent transportation markets are able to take full advantage of GigE Vision compliant components for new designs and installations, system upgrades, or cost-effective conversion from other network topologies. With the 2.0 release of GigE Vision and its formal inclusion of 10 GigE in the standard, the development of GigE Vision over NBASE-T technologies, together with other important characteristics such as compression and precision triggering, the future of GigE Vision as a dominant communications medium for vision-critical applications is assured.

Introduction

Automated industrial imaging, using analog cameras in earlier machine vision systems, paved the way for the digital systems of today. High-speed digital cameras

The AIA's GigE Vision® standard has enabled a cost-effective and fast means of capturing, transmitting, analyzing, and displaying video

have exceeded the capabilities and performance of analog cameras, enabling the rapid move to an all-digital infrastructure. The point-to-point topologies largely in use until now were the only viable means of real-time, reliable transmission, but these suffered from a number of drawbacks, including limited reach and less flexible interfaces.

Ethernet's flexible architecture, however, with point-to-multipoint and multipoint-to-multipoint connectivity, which provides native video distribution and aggregation, coupled with the AIA's GigE Vision standard, has enabled a cost-effective and fast means of capturing, transmitting, analyzing, and displaying video.

This whitepaper examines the benefits of GigE networked video, illustrates some of its key characteristics in the context of a number of primary applications, and draws conclusions as to its ease of implementation and future development.

The Importance of Networked Video

Real-time functionality is often achieved using point-to-point connections between a vision sensor and an image capture board or frame grabber in a PC – see Figure 1. Complexities arise when the images need to be viewed on more than one display or processed on more than one PC, necessitating the configuration of additional point-




Figure 1: Common network topologies; distributed networks offer the greatest flexibility

to-point connections on extra PCs, display controllers, and other pieces of specialized hardware.

Consequently, point-to-point connections can be costly to install, difficult to manage, and expensive to scale. Moreover, as sensors continue to evolve to higher resolutions and faster frame rates, the high bandwidth needed for real-time image transfer becomes a limiting factor.

Ethernet, on the other hand, offers exceptional networking flexibility using standard switches, supporting almost every conceivable connectivity configuration, including point-to-point, star (point-to-multipoint), and distributed (multipoint-to-multipoint).

Ethernet is supported by a well-understood infrastructure based on cost-effective and non-proprietary chip sets, switches, and cabling.

As the primary standard deployed in most of the world's networks, including those for demanding military and industrial applications, Ethernet is supported by a well-understood infrastructure based on cost-effective and non-proprietary chip sets, switches, and cabling. (See Appendix A for a comparison of vision system interface technologies.)

Additionally, different traffic rates between 10 Mb/s up to 10 Gb/s can be handled by the same switch, ensuring backward compatibility and permitting system upgrades without sacrificing legacy cameras or imaging devices already in place

More importantly, Ethernet offers long reach—allowing spans of up to 100 meters between network nodes over standard, low-cost Cat 5/6 copper cabling—and even greater distances with switches or fiber. Its superior scalability, supporting distributed network configurations, easily accommodates the addition of new processing nodes, displays and switches.

Networked video means that all elements—such as image sensors, cameras, video receivers, video servers, control units, and displays—are fully interconnected, with each of them using the same standard framework to transmit or receive video and control data.

Therefore, the distinct advantages of networked transmission for high-resolution video are:

- **Reduced system costs:** using standard Ethernet components.
- **Greater design flexibility:** permitting more configuration and distribution options
- **Easy system maintenance:** expanding and upgrading with little or no interruptions.

The GigE Vision Standard

The benefits of Ethernet for transporting high-speed images were first recognized in the machine vision sector more than ten years ago, when GigE was introduced into mainstream networks. The popularity of GigE for industrial vision applications led to the introduction of GigE Vision in 2006, a global open standard governing the distribution of video and control data over Ethernet networks. The GigE Vision standard establishes a standardized environment for the delivery of networked video applications based on switched client/server Ethernet architectures. Its success is illustrated by industry data for the most common camera interface technologies in use in machine vision applications – see Figure 2. (See Appendix B for more details on GigE Vision and its evolution.)

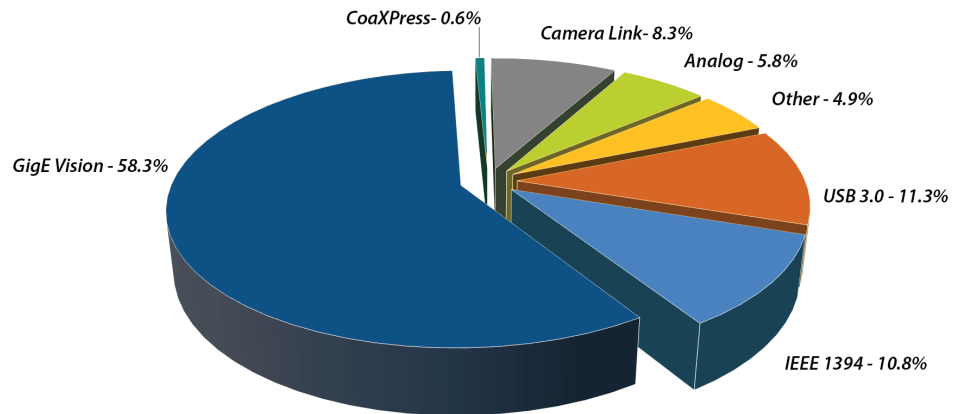


Figure 2: Camera Interface Technologies Sales by Units
(AIA 2016 Machine Vision Camera Study)

The GigE Vision standard defines four main areas that are specific to machine vision networked systems – see Figure 3. Importantly, GigE Vision leverages about 25 existing industry standards rather than introducing new proprietary schemes. These include IEEE 802.3 (Ethernet), IEEE 1588 (time synchronization), IETF RFC2026 (jumbo frames) and EMVA GenICam (XML device description file).

While the GigE Vision standard creates a robust framework for device interconnectivity, the quality of the vendor implementation can significantly impact its operation, because the standard itself neither defines nor recommends performance criteria. The levels of latency and jitter inherent in a system are a key area that can be influenced by the interface implementation – see Appendix C. In the final event, vision systems must deliver images practically in real-time, therefore an examination of the important factors that influence high-speed transmission are covered in the next section of this whitepaper.

Device Discovery	Defines how compliant devices obtain IP addresses and are identified on the network
GigE Vision Control Protocol (GVCP)	Defines how to specify stream channels and control and configure compliant devices
GigE Vision Stream Protocol (GVSP)	Defines how images are packetized and the mechanisms by which images can be transferred
XML Device Description File GenICam Compliant	Provides computer-readable device datasheet permitting access to controls and image stream

Figure 3: The four main elements of the GigE Vision standard

Vital Statistics of Real-Time Communication

Many performance characteristics that are imperative to real-time video—such as low and consistent latency, low jitter, high throughput, guaranteed data delivery, and low CPU usage—can vary greatly from system to system. Of all these, low latency and low jitter are critical factors in real-time systems, because delays or inaccurate timing over certain thresholds are generally not tolerable.

When Ethernet networking was initially conceived, delays in the order of seconds were considered acceptable. Even now, closed circuit security video networks experience typical delays of around two seconds because most of the video is stored for later viewing, so a delay is less of an issue. By contrast, today's sophisticated industrial vision systems have "glass-to-glass" (from lens to monitor) latency threshold tolerances in the order of 100 ms or less.

Closely associated with latency are the effects of signal jitter. While systems can compensate for the effects of jitter, they generally cause an increase in the delay, making latency compensation considerably more complex. Figure 4 shows the causes and effects of both.

Where dedicated electrically connected systems, such as those based on frame grabbers, exhibit very low and consistent latency, standard Ethernet packetized systems have inherent latency and jitter, introduced, in part, by overheads in the data packets and traffic handling protocols. High-performance GigE Vision implementations largely address these issues by optimizing the packet handling process on both the transmit and receive ends of the link. Tests show that the glass-to-glass latency — which takes into account exposure and readout time, as well as the time required for processing, display, and more — and is typically <100 milliseconds (μ s); well within the tolerances of the vast majority of vision systems.

	Definition	Causes	Effects
Latency	The minimum time between image capture and image reception/display.	Data processing, line speed, and the receive and retransmit delay.	Delay between the actual event and the availability of the image data for processing.
Jitter	The variation in the latency.	CPU processing and scheduling, buffering in network switches, crosstalk, and EMI (electromagnetic interference).	Synchronization errors, delayed images, late/useless data, and reduced processing time.

Figure 4: Causes and effects of latency and jitter in vision systems

Use of Jumbo Frames

Other measures that help increase performance even further include the use of jumbo Ethernet frames, resulting in a reduction in CPU usage. A jumbo frame carries a much larger payload, typically about 9,000 bytes compared with a standard frame of 1,500 bytes. At six times the size, its overhead is just 0.55%, compared with 4% for standard frames—a significant amount considering that approximately 13,800 jumbo packets per second can be sent at maximum throughput over a 1 GigE link.

Therefore, the use of jumbo frames not only has a positive effect on throughput, but also on CPU usage—see Figure 5—because the CPU processes and copies each frame on arrival. For example, in the tests in Figure 5, a data throughput of 700 Mb/s was the maximum achievable over a GigE connection. By switching to jumbo frames, the reduced overhead raised the data throughput to 960 Mb/s. Appendix C provides more details.

Reduction in CPU Usage

One of the key differences between GigE Vision and Camera Link is that GigE Vision has no independent incoming data management. A Camera Link frame grabber has the capability to assemble complete images via its FPGA, and generate a single OS interrupt. GigE Vision, on the other hand, relies on the CPU to process each packet, with associated multiple interrupts per received image, and then forward each image to the application. Therefore, the CPU is not only processing the data stream, but also assembling images from received network packets, which can add up to a significant burden if not managed appropriately. A high CPU load will inevitably result in processing delays in the host PC, as well as a reduced image processing capability. The use of jumbo frames clearly diminishes the number of interrupts, lowering the CPU burden accordingly, and increasing the processing bandwidth.





Figure 5: CPU utilization using standard frames ...



... and jumbo frames.

GigE Vision filter, stack and performance drivers can further reduce the burden. Such drivers, for example, transfer the larger video packets directly to the application software, freeing the CPU of a considerable processing load – see Appendix C.

Ethernet technology has proven to be a reliable and robust medium for vision interfaces and real-time image transmission. The broad range of mission-critical applications across multiple markets using GigE Vision compliant devices provides further evidence of its versatility.

Vision-Critical Applications

System designers for a wide range of medical, military, transportation and industrial video applications look for transmission capabilities with low latency, high reliability, and cost-effective implementation.

For example, achieving the performance required for image-guided surgery needs careful selection of the devices that fit the required criteria. Manufacturers of X-ray equipment, flat panel detectors (FPDs) and other medical imaging appliances have been embedding GigE Vision networking interfaces into their systems for some time. These products, such as small-footprint video transmitters and receivers, stream imaging data from sensors and flat panel detectors to PCs in real time for processing and display. The precise implementation reduces system costs by replacing image capture boards (frame grabbers) with performance-oriented, GigE Vision-compliant software.

GigE is a natural choice for video transmission within military ground vehicles, due to its lightweight cabling, networking capabilities, and support for a range of different computing platforms. In a local situation awareness (LSA) application, real-time video from cameras and sensors is transmitted to display panels for crew members to navigate the windowless vehicle and survey surroundings. (Figure 6).

External frame grabbers convert the video feed from analog and Camera Link cameras into a GigE Vision image stream. The video is streamed uncompressed with low, consistent “glass-to-glass” latency over the multicast Ethernet network to displays and processing equipment within the vehicle. End-to-end (or glass-to-glass) system latency is under 80 milliseconds (ms)—a small delay that equates to real-time viewing. A large part of this latency is associated with the native low frame rate of NTSC (30 fps or 33 ms latency per frame) and PAL (25 frames per second or 40 ms latency per frame) analog video. In systems where higher-frame rate digital cameras with a GigE Vision interface are used, the end-to-end latency will be lower.

Video, control data, and power are transmitted over the single cable; lowering component costs, simplifying installation and maintenance, and reducing “cable clutter” in the vehicle. The networking flexibility of GigE allows images from multiple cameras/sensors to be aggregated to a single port on a mission computer or processing unit, and/or imaging data to be multicast from one camera/sensor to multiple displays.

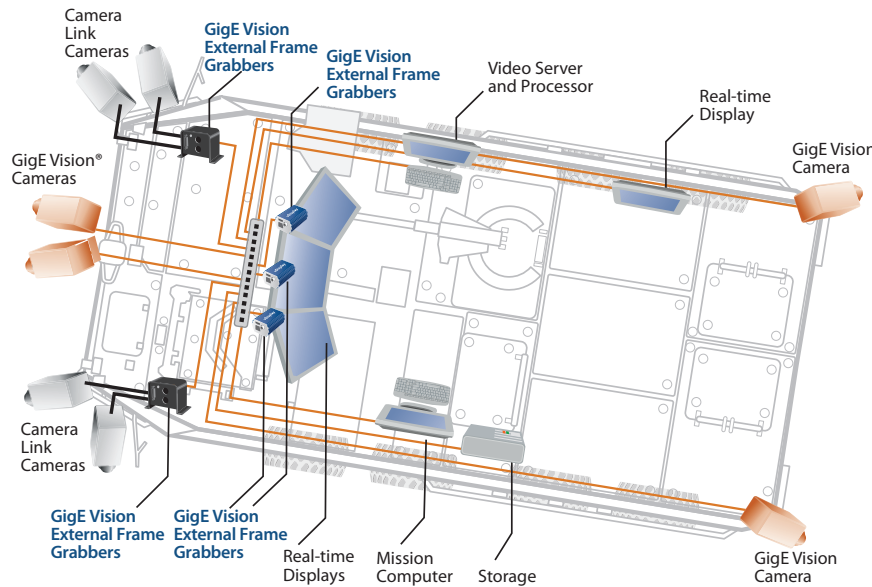


Figure 6: A vision-enabled local situational awareness system for military vehicles

In the transportation sector, distance and reach are some of the deployment issues that need to be accounted for in the design phase. Intelligent transportation systems, such as those used in automatic toll road lanes, will likely include fiber links for backhauling data, as well as local cable runs of up to 100 meters to reach across highways. Both are readily achievable with GigE networks using standard interfaces and switches.

Industrial automation systems were among the first to employ machine vision technology, which emphasized the need for ultra-low latency in order to facilitate the use of fast robots and rapid assembly lines in production facilities. Similarly, speed is essential in automated sorting facilities, such as those found in postal and parcel services. Additionally, GigE Vision employs a data resend capability, which can ensure data delivery even in the event of packet loss or corruption, whether due to EMI or any other reason. In the industrial sector, cable reach across factory floors is also vital, so existing Ethernet infrastructures can easily be leveraged for this purpose.

In Figure 7, an automatic inspection process uses the imagery from line scan cameras to analyze products for possible defects—each PC with a different defect-detection algorithm. In This example, external frame grabbers convert the video feed from Camera Link cameras into GigE Vision. Alternatively, manufacturers can develop 1G, 2.5G, 5G and 10G cameras using GigE Vision and GigE Vision over NBASE-T embedded hardware. The master controller PC aggregates the results and sends grading instructions to the GigE Vision stamping device. Trigger synchronization provided by the programmable logic controller (PLC) uses an encoder to provide a common time reference, ensuring that latency and jitter—whether generated by the network, by the host PC, or by the application itself—are removed from the equation




Figure 7: An automated inspection process using machine vision on a GigE network

With advances in higher resolution, higher sensitivity sensors, NBASE-T technology offers a natural evolution for advanced imaging thanks to its bandwidth support, low-cost cabling, and compatibility with the GigE Vision standard. The IEEE 802.3bz standard, which encompasses the NBASE-T specification, defines a new type of Ethernet signaling that boosts the speed of twisted-pair cabling well beyond its designed limit of 1 Gbps to support 2.5 and 5 Gbps speeds at distances up to 100 meters.

With GigE Vision over NBASE-T, designers can transmit uncompressed images at throughputs up to 5 Gbps over low-cost Cat 5e copper cabling. The extended-reach, flexible, and field-terminated cabling can be easily routed through systems to ease installation and maintenance. Designers can create NBASE-T imaging devices and vision systems that are natively compatible with GigE Vision compliant software. An NBASE-T network interface card (NIC) and Gigabit Ethernet NIC are treated the same by Windows, Linux, and other operating systems. This means existing GigE Vision-compliant software and software development kits (SDKs) are compatible with NBASE-T without any modifications.

GigE Vision over NBASE-T provides a straightforward upgrade path for manufacturers, enabling the design of next-generation imaging devices for advanced applications.

The Future of Ethernet

Ethernet technologies continue to be enhanced for consumer, enterprise and networking applications. The Ethernet Alliance, an industry consortium that works with the IEEE and other standards bodies, publishes an annual roadmap that outlines the expected evolution of Ethernet technologies.

In its most recent roadmap, published in 2016, the Alliance outlines a continuing evolution to higher bandwidth, with router and switch ports running at 200 Gbps and 400 Gbps and greater than terabit per second transmission for the near future. The Ethernet Alliance also highlights an expanding footprint for next-generation networked applications, with Internet of Things (IoT), Internet of Device (IoD) and cloud services pushing demand for 10 GbE, 40 GbE, 100 GbE, and 400 GbE technologies.

Similarly, the NBASE-T Alliance – a consortium of over 40 companies representing all major facets of networking infrastructure including the vision market – is supporting the widespread use and deployment of 2.5G and 5G Ethernet through promotion of the IEEE 802.3bz standard, and testing and compliance programs to facilitate the development and deployment of interoperable products.

The NBASE-T specification defines a new type of Ethernet signaling that boosts the speed of twisted-pair cabling well beyond its designed limit of 1 Gbps to support 2.5 and 5 Gbps speeds at distances up to 100 meters. The specification supports autonegotiation between the new NBASE-T rates, and slower 1 Gbps rates, or – if the network infrastructure supports it – 10 Gbps.

Conclusions

As they continue to evolve, GigE video networks will serve as important technology platforms for vision-critical applications enabling real-time display, processing and storage.

Their broad range of attributes appeal to many users due to:

- Versatile system design
- Universal adoption
- Cost-effective infrastructure
- Excellent interoperability
- International standardization

Above all, quality implementation will define the exceptional performance levels demanded by the industry. Characteristics imperative to real-time video—ultra-low and consistent latency, low jitter, high throughput, guaranteed data delivery, and low CPU usage—are required to meet the needs of systems integrators and designers, enabling future expansion into new markets.

References

Mark Claypool and Jonathan Tanner: The Effects of Jitter on the Perceptual Quality of Video. Worcester Polytechnic; 2000.

Wikipedia: Ethernet Frame. Source: http://en.wikipedia.org/wiki/Ethernet_frame; September 2011.

Ethernet Alliance: Ethernet Jumbo Frames. Whitepaper; November 2009.

Ethernet Alliance: The 2016 Ethernet Roadmap. Market Report, March, 2016.

NBASE-T Alliance: NBASE-T 2.5G/5G Ethernet Technology Basis for the IEEE 802.3bz Standard. Whitepaper, December 2016.

Neil Martin: Bridging the bandwidth gap in vision. MVPro; March, 2017

GigE Vision Standard: <http://www.visiononline.org/vision-standards-details.cfm?type=5>; Version 2.0; November 2011.

AIA: AIA 2016 Machine Vision Camera Study. Market Report; 2016.

About Pleora

Pleora Technologies Inc. pioneered the use of Gigabit Ethernet for real-time networked video connectivity and revolutionized industrial imaging. With this spirit of innovation, Pleora engineers networked video connectivity solutions for mission-critical applications. Working with its rich portfolio of video networking elements Pleora partners with customers to tailor solutions to their individual needs, from definition to deployment, with full integration support.

To find out why more military, medical, and manufacturing integrators and OEMs choose—and stay with—Pleora, visit <http://www.pleora.com>.

Pleora Technologies Inc.
340 Terry Fox Dr., Suite 300
Kanata, ON K2K 3A2
Canada

Appendix A - Comparison of Vision System Interfaces

The three key vision system interfaces in use today are Camera Link®, CoaXPress, and GigE Vision®. Figure 9 compares these three interfaces, as well as Camera Link HS®, USB3Vision®, and GigE Vision over 10 GigE, which are three emerging interfaces.

Attribute	Camera Link	Camera Link HS	CoaXPress	USB	GigE Vision	GigE Vision over 10 GigE
Cable Type	Camera Link	Copper or Optical	Coaxial	USB 3.0	Cat5/6 or Optical	Copper or Optical
Power over Cable	13 W	Possible	4 W	4.5 W	13 W	15.4 W
Industry Adoption	Broad	Low	Limited	Broad	Broad	Emerging
Relative System Cost	High	Medium	Medium	Low	Low	Medium
Max Payload Throughput on Single Cable	2.08 Gb/s	16.8 Gb/s	6.25 Gb/s	2.8 Gb/s	980 Mbps (GigE) 2.4 & 4.8 Gb/s (GigE over NBASE-T)	9.0 Gb/s
Requires Non-standard Hardware	Yes	Yes	Yes	No	No	No
Max Distance	10 m	15 m (copper, 2.4 Gb/s) 100 m (optical 16.8 Gb/s)	130 m @ 1.25 Gb/s 40 m @ 6.25 Gb/s	3 m	100 m (Cat 5) 40 km (optical) Unlimited with switches	7 m (Direct Attach Copper) 100 m (Cat 6a) 40 km (optical) Unlimited with switches
Network Topology	Point-to-point	Point-to-point	Point-to-point	Star	Distributed	Distributed
Native OS Support	No	No	No	Yes	Yes	Yes
Guaranteed Delivery	No	Yes	No	Yes	Yes	Yes
Standard Ratification	2000	2012	2009	2013	2006	2006

Figure 8: Key attributes of the major vision system interfaces

The background on each interface technology can be summarized as follows:

- **Camera Link** has been in use since 2000 and has a large installed base, mainly because of its high throughput and low latency. However, its 26-wire connection cables and its limited cable length add to its installation complexity and expense.
- **Camera Link HS** was released in 2012 to overcome the speed limitations of Camera Link for line scan cameras; however the short cable reach of up to 15 meters is a limitation.
- **CoaXPress** has some advantages in its speed, reach, and straightforward coaxial cabling, but like the other non-Ethernet based standards, has no networking capability. Introduced in 2010, it was designed to address some large existing coax installations, and as a consequence, has gained only narrow industry acceptance. Speeds comparable with 10 GigE are only attainable using more expensive multi-core and active cables, and over significantly shorter distances.
- GigE Vision was introduced in 2006 and is based upon the 802.3 Ethernet standard, which has been in use since 1980, and is therefore by far the most mature and best recognized standard. Cat 5/6 cables are widely available and have full-duplex 8-strand cores. The NBASE-T specification boosts the speed of twisted-pair cabling beyond its designed limit of 1 Gbps to support 2.5 and 5 Gbps speeds at distances up to 100 meters
- 10 GigE Vision Version 2.0, released in November 2011, introduced support for 10 GigE. Where GigE Vision over 10 GigE ensures backwards compatibility of existing GigE Vision devices, the same will not be true of next generation Camera Link devices.
- USB3 Vision is a global standard for transporting high-speed imaging and video data to computers over the widely available USB 3.0 bus at throughputs approaching 3 Gb/s. The USB3 Vision standard, ratified in February 2013, helps reduce the design, deployment, and maintenance costs of high-speed video applications by making it simpler to leverage the native performance attributes of the USB 3.0 platform, such as its high bandwidth, power over cable, and ease-of-use.

The three vision system interfaces of Camera Link, CoaXPress and GigE Vision have one important common factor—they all stem from higher level standards.

Of all the interfaces, only GigE Vision and USB3 Vision have multi-channel aggregation capability (not to be confused with link aggregation) in which a number of cameras can be received via a single port. Channel aggregation, achieved with standard Ethernet switches and USB 3.0 hubs, illustrates the versatility of these standards. Additionally, GigE Vision is the only standard which can operate using a distributed network configuration—the most flexible, cost-effective, and scalable topology.

In summary, GigE Vision and Ethernet have significant advantages over Camera Link and CoaXPress, while USB3 Vision is suited for higher bandwidth, shorter-reach applications. GigE Vision delivers a unique combination of networking, throughput, flexibility, distance, and scalability that makes it the optimal choice as a digital video interface

Appendix B - Ethernet and GigE Vision

The GigE Vision standard is open and globally accepted. Since its introduction in 2006, it has been embraced by leading hardware and software companies that develop systems for high-performance video applications in the military, aerospace, medical, and manufacturing sectors.

The 1000 Mb/s Gigabit Ethernet bus provides sufficient bandwidth to transmit uncompressed 5 Megapixel images at 15 frames/sec over distances of up to 100 meters using inexpensive copper cabling, and even further if using switches or fiber. The fourth generation of Ethernet, 10 GigE, now widely available in mainstream markets, delivers 10 Gb/s. Longer distances can be bridged using switches and routers. The IEEE 802.3bz standard, based on the NBASE-T specifications defines 2.5 and 5 Gigabit Ethernet speeds over twisted pair copper cabling (2.5GBASE-T and 5GBASE-T).

Where Ethernet excels is in its unsurpassed cost-effectiveness:

- All Ethernet generations employ the same frame format, ensuring backward compatibility and permitting system upgrades without sacrificing legacy cameras.
- System costs are reduced through the use of low-cost and widely available components and elimination of the need for cable repeaters.
- High scalability reduces future capital expenditure as GigE Vision continues to be developed for higher bandwidths.

GigE Vision Protocol

The GigE Vision protocols operate at OSI Layers 5-7 for the delivery of video and control data over Ethernet networks. Figure 10 illustrates how the GigE Vision standard fits into the OSI and TCP/IP models.

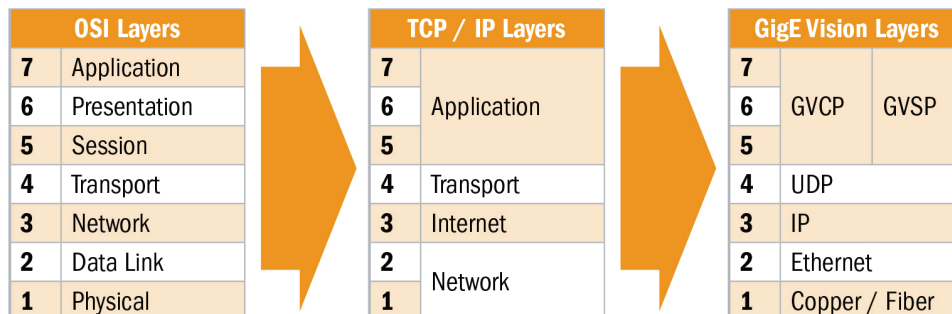


Figure 9: Comparison of the OSI, TCP/IP and GigE Vision networking models

One of the outstanding advantages of standardization through GigE Vision is guaranteed interoperability between a broad range of devices.

Of note is that User Datagram Protocol (UDP) is used to handle transport at layer 4, rather than Transport Control Protocol (TCP). UDP was selected for its simplicity, low overhead and multicast support. It is ideally suited for low-latency networked video, but does not guarantee data delivery. To address this limitation, the GigE

Vision standard includes an optional mechanism that allows video sources to resend undelivered data to video receivers. This mechanism is part of GigE Vision Control Protocol (GVCP) and GigE Vision Stream Protocol (GVSP), and together with other recommendations in the standard, allows performance-oriented implementations of the GigE Vision standard to guarantee video transport and achieve low, predictable latency—even during a resend.

One of the outstanding advantages of standardization through GigE Vision is guaranteed interoperability between a broad range of devices including image sensors, cameras, video transmitters and receivers, video servers, control units, and displays – see Figure 10. Automated test suites permit rapid verification of interoperability, and a number of international plugfests have confirmed its ease of use and clearly demonstrated common certification.

Similarly, designers can create NBASE-T imaging devices and vision systems that are natively compatible with GigE Vision compliant software. An NBASE-T network interface card (NIC) and Gigabit Ethernet NIC are treated the same by Windows, Linux, and other operating systems. This means existing GigE Vision-compliant software and software development kits (SDKs) are compatible with NBASE-T without any modifications.

The 2.0 Release

The GigE Vision standard continues to evolve. Version 2.0, released in early 2012, formally includes 10 GigE—key for high-throughput applications. Consequently, switching from Ethernet will be an attractive option for a wide range of applications. Aggregation of a number of cameras onto a single link is especially advantageous because of a quantifiable reduction in the cost of cabling and system components such as frame grabbers. Two more features included in the 2.0 release are precision timing for trigger synchronization, and the transmission of compressed video streams providing savings in bandwidth and storage.

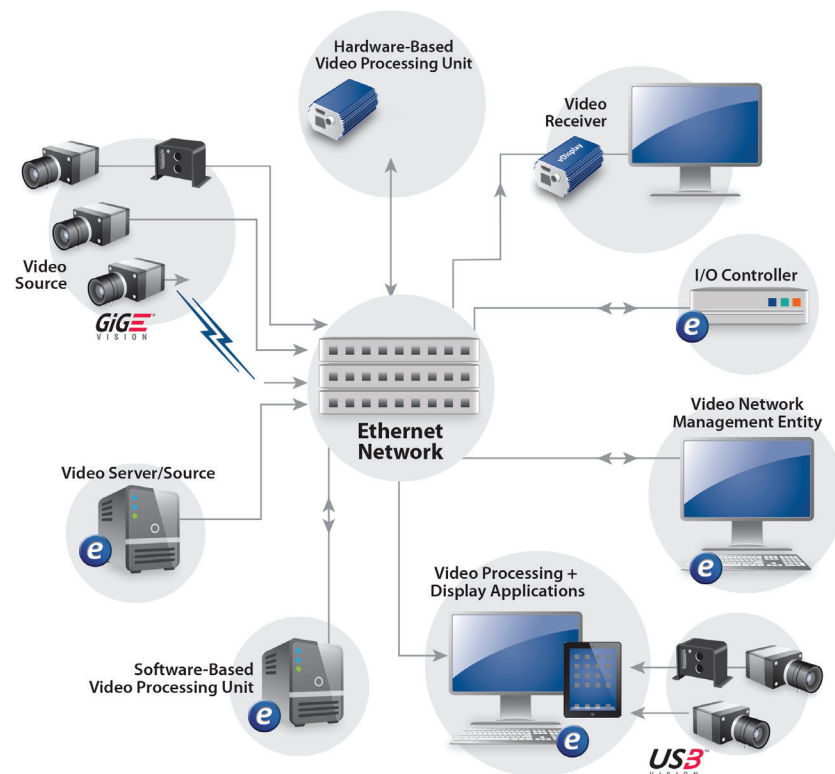


Figure 10: Elements of a networked distribution system (new one)

Appendix C - Ultra-Low Latency and Jitter in GigE Vision Systems

Ultra-low latency and jitter are important characteristics of any digital communications network, and are closely related. The levels of latency and jitter will determine the quality and timely delivery of the video, and minimizing them is highly desirable— particularly where real-time analysis is required.

In essence, latency is the total delay between the sensor and the receiver (display or image processing application), whereas jitter is the variation in the delay.

The link between latency and jitter is that a network with constant latency has no jitter, whereas a network with variable latency may have high jitter.

Interestingly, typical compensation techniques can actually add to the problem:

- Buffering can compensate for jitter but introduces latency.
- Interleaving, redundancy, or retransmission may compensate for lost packets but can add latency and jitter, degrading overall system performance.

Quantifying Latency and Jitter

The absolute levels of latency and jitter that are acceptable vary enormously with the application. In video security systems, latency of about two seconds is common, much of it attributable to encoding and decoding of the video such as in H.264 compression. In real-time applications, however, delays of more than 200 ms might be deemed unacceptable, especially where lives may be at risk as in image-guided surgery or in closed-hatch driving of a military vehicle.

It is worth noting that video in real-time vision-critical systems is generally uncompressed because of the additional processing required to decompress the image at the PC. Also, some compression engines can be “lossy” in nature (i.e. not all of the data is faithfully reproduced, and image integrity may be compromised). While compression was introduced into version 2.0 of the GigE Vision standard, its primary benefit is to reduce transmission bandwidth requirements.

Jumbo Frames

There are a number of methods of increasing system performance. In Gigabit Ethernet, a significant method is through the use of jumbo frames that have an expanded payload and thereby enable a much higher throughput. The table below shows the composition of a standard Ethernet frame of about 1,500 bytes. A typical jumbo frame has around 9,000 bytes, and therefore carries six times the payload of a standard frame as illustrated in Figure 11.

802.3 Ethernet Frame Structure									
Preamble	Start of frame delimiter	MAC destination	MAC source	802.1Q tag (optional)	Ethertype or length	IP, UDP, GVSP Headers	Payload (Raw Data)	Frame check sequence	Interframe gap
7 bytes	1 byte	6 bytes	6 bytes	(4 bytes)	2 bytes	36 bytes	1464 bytes	4 bytes	12 bytes
1542 bytes at full payload									




Figure 11: Composition of a standard Ethernet frame (top) and comparison with a jumbo frame (below)

The protocol efficiency for Ethernet frames may be calculated as: Protocol efficiency = payload size ÷ frame size

Maximum efficiency is achieved with the largest allowable payload sizes.

For a standard frame:

$$1464 \div 1542 = 94.9\% \text{ efficiency}$$

For a jumbo frame:

$$8964 \div 9042 = 99.14\% \text{ efficiency}$$

For a GigE connection transmitting at 960 Megabits of information per second, the use of jumbo frames is a significant method of increasing throughput. The marginal level of latency introduced through the processing of the longer jumbo packets versus short standard packets is considered negligible.

Minimizing Introduced Latency

When a device, such as an analog to GigE converter or a frame grabber, is built into a camera or inserted into the video link, latency is necessarily introduced, however incremental. Keeping it to the minimum is essential. Well-designed interfaces are fast and efficient at processing the video stream. For example, only 400 μ s of additional latency is introduced with a 2-Megapixel camera over a 1-Gb/s Ethernet link, using a filter driver on the PC and an IP engine at the camera.

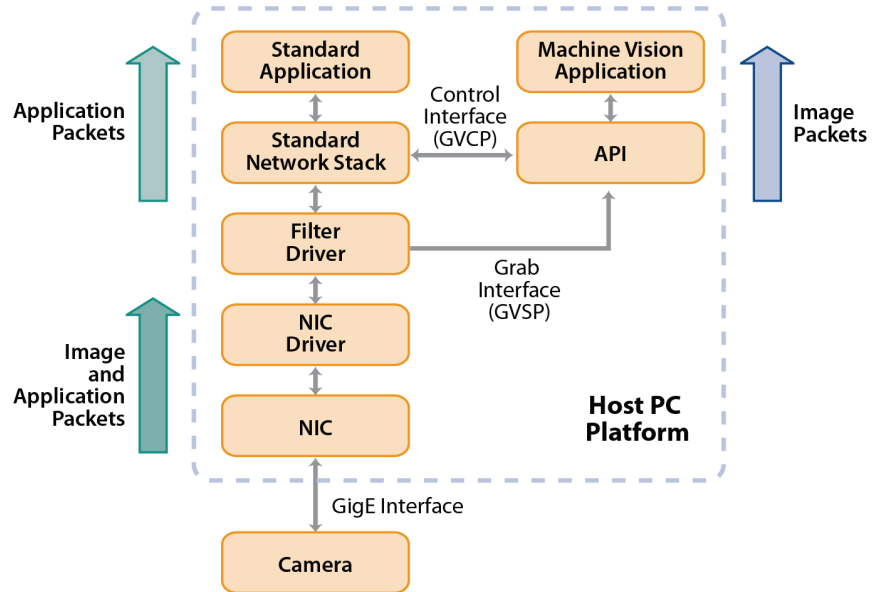


Figure 10: 12: Filter drivers reduce CPU usage and associated latency

The levels of latency and jitter will determine the quality and timely delivery of the video.

A related area that is often overlooked is the CPU usage introduced by software tasks. The host CPU receiving images from networked devices has to share its processing time between all of the services on that PC, and unnecessary latency could be introduced by devices and software that place a high processing overhead on the CPU.

Software drivers can reduce CPU usage in high-performance software development kits. The driver simply inspects the content of each packet, grabs all those with GigE Vision video content, and hands them directly off to the application programming interface (API)—see Figure 12—thus bypassing the native Windows or Linux standard network stack, enabling such operations as image analysis to be prioritized accordingly. Filter drivers have the major advantage of being device independent, and supporting all network interface cards (NICs) rather than being manufacturer specific. Performance drivers provide marginally faster processing times by replacing the NIC driver but are highly device dependent.

Windows Vista Operating System	Average Latency	Average Jitter	Average CPU Usage
Without driver	352.75 μ s	39.62 μ s	10.4%
With eBUS Universal Pro SDK filter driver	386.53 μ s	43.70 μ s	5.2%

Figure 13: Test results for latency and jitter with and without a filter driver. [Test parameters: 1600x1200 pixels, 40.4 fps, 620.8 Mbps; Pleora eBUS SDK and Universal Pro driver]

Measuring Introduced Latency

Laboratory tests on latency and CPU usage show a significant improvement in performance when filter drivers are employed – see Figure 13. Measurements are taken from the time of the transmission of the first Ethernet packet to the completion of the reassembled image by the SDK. The latency and jitter introduced are not only deterministic, but of negligible levels, while CPU usage is reduced by half from its already low level. It is worth noting that the level of excellence achieved in such results can be very much vendor specific, depending on the quality and method of the implementation.

The use of jumbo frames, high-performance filter drivers, well-designed interfaces, and a low overhead network protocol successfully achieves ultra-low latency and jitter in GigE Vision interfaces. Continued development in such areas as low-jitter triggering will further enhance performance.