Understanding HPEC Computing: The Ten Axioms

Read About

HPEC

HPEC and HPC market

Embedded System market

Sensor market

The 10 Axioms of HPEC

Introduction

High Performance Embedded Computing (HPEC) is bringing previously unobtainable levels of processing power to compute-intensive defense and aerospace applications, through its unique convergence of high performance computing (HPC), embedded systems, and sensor systems. The requirements of HPEC systems, used to satisfy demanding signal, radar and image processing, comprise a distinct market in which select best-of-breed technology is adopted from large commercial markets and then modified for use in military environments.

In HPEC system designs, commercial HPC hardware and software standards are leveraged and innovatively packaged to ensure that they can meet the demands of rugged embedded deployment. Various industries provide the open standards and building blocks that comprise advanced HPEC systems. HPEC embraces adjacent market technologies from a wide variety of markets including military, information technology (office IT), telecommunications, industrial automation, scientific discovery/ research, test and measurement, and finance.

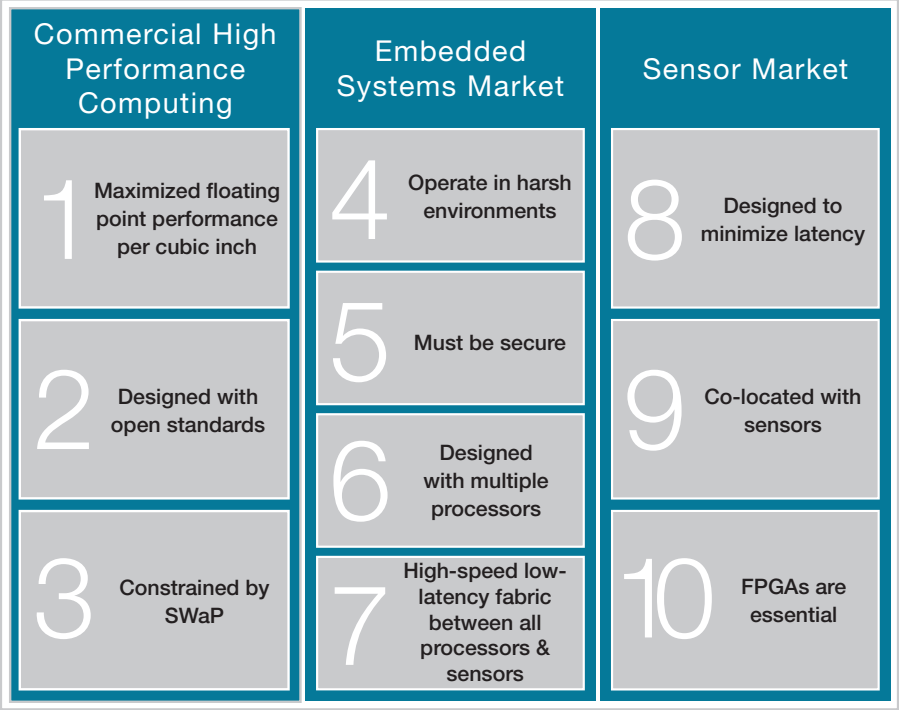

Figure 1: The 10 Axioms of HPEC systems

Info

curtisswrightds.com

Email

ds@curtisswright.com

The technologies that comprise an HPEC system are primarily derived from three distinct markets – the Commercial HPC, Embedded Systems and the Sensor markets. HPEC tracks and benefits from advances in the technologies that service these markets. As their technologies evolve, the HPEC market adapts by selecting the appropriate technologies that meet the Ten Axioms of HPEC that define the essential elements of an HPEC system.

An HPEC system must meet the criteria outlined in the following Ten Axioms of HPEC to be able to satisfy the requirements of today's most demanding defense and aerospace rugged embedded applications. To help system designers to better understand what to look for in an HPEC system, and what to avoid in less capable system alternatives, this paper will discuss the Ten Axioms, organized by the industry (HPC, Embedded Systems and Sensor) from which each is derived. The paper will explain in detail how each of these essential elements relates to and meets the unique requirements of these high performance-processing solutions.

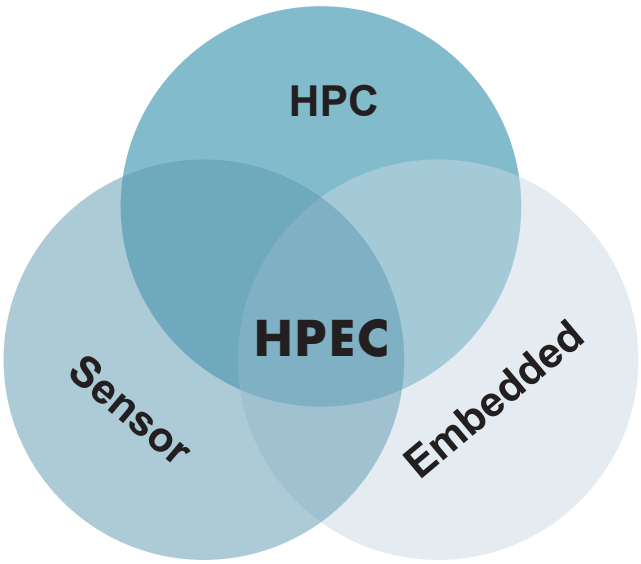

Figure 2: HPEC contributing markets

HPEC and the HPC Market

Axiom 1: Maximized Floating Point Performance per cubic inch (FLOPS)

Today's commercial HPC market produces processing systems that consist of clusters of Intel® or Intel-compatible personal computers. These x86 architecture processors may be augmented with powerful math accelerators, such as general purpose graphic processor units (GPGPUs), and are typically inter-connected via Ethernet or InfiniBand®-based networks. The performance of these commercial HPC systems is measured in floating-point operations per second (FLOPS). The www.top500.com web site lists the largest of these systems.

A characteristic of these HPC systems is that their application program interface (API) sets, such as VSIPL, MPI, and BLAS for core-to-core or processor-to-processor communications, are typically converged with a math library to enable measurement of their FLOPS performance using a benchmark such as LINPACK. As the HPC market is called on to service an increasing set of applications that require cluster computing solutions, the number of HPC middleware packages is also expanding. This trend is evidenced on http://en.wikipedia.org/wiki/Computer_cluster.

One of the newest and most significant HPC applications to emerge recently is "Big Data". Big Data applications work on extremely large datasets - a classic example is Google's search engine software for the Internet. As these datasets grow, current hardware and software architectures are increasingly less able to keep pace. The result is that processing cannot be performed in the required, reasonable time span. Designers of HPC systems seek the highest available performing FLOPS processor for use as the basis, or smallest node, for these large processing systems. This is also true for HPEC systems.

Axiom 2-3: Open standards and SWaP

Much like the HPEC industry, the commercial HPC industry got its initial start with system solutions based on proprietary hardware, software and infrastructures. After HPC became widely deployed, the cost of hardware, software and maintenance limited the continued adoption of proprietary processing technologies. Today, HPC systems are typically based on open architecture designs using Intel processors, which are sometimes augmented with GPGPUs. These designs rarely use fabrics other than Ethernet or InfiniBand.

In recent years, as HPC systems have continued to increase in size, system designers have responded by turning to power-optimized processors that make it possible to install more powerful systems in the limited existing space envelope used by legacy systems, without having to increase their existing power and cooling infrastructure. These emerging SWaP constraints could disrupt today's use of standard off-the-shelf processors in HPC system designs.

HPEC system designers look to the HPC market, leveraging and adopting as much as possible of the proven hardware technologies it has developed. As HPEC designs increasingly adopts the compute hardware used in HPC systems it becomes easier for the HPEC market to embrace the large amount of tested and proven open standard software developed for use in the HPC realm. While SWaP constraints are starting to emerge as an issue for HPC system designers, they are a familiar challenge for designers of HPEC platforms deployed on military air and ground vehicles. As SWaP becomes a design hurdle for HPC system integrators it brings that aspect of the two markets even closer together. On the software front, leveraging popular HPC APIs provides HPEC system developers with access to Open Source HPC application code and tools. This approach can provide significant benefits. For example, reprogramming legacy code to these API sub-sets simplifies and speeds the porting of software from one generation of processor technology to the next.

The defense and aerospace customers supported by HPEC system vendors increasingly demand open standard solutions to lower costs, which fosters competition and ensures long life program support. As the HPC market has moved away from proprietary architectures and embraced open standards, the HPEC market has reaped the benefit. The open architecture hardware and software standards used in the HPC world provide a proven path for HPEC system designers, enabling them to meet their goal of providing open standards and addressing SWaP requirements. HPEC systems tend to comprise a blend of HPC and SWaP-critical applications.

On the hardware side, there have been some attempts, both with commercial telco and military system designers, to use rugged personal computers (PCs) or rackable blade processors such as BladeCenters to satisfy the compute requirements of HPC and HPEC applications in non-office environments. Very few of these systems do provide high-speed, low-latency interconnects, and can be co-located to the sensor array, such as a ballistic missile defense system on a Navy ship. Unlike the cramped, hot environment of combat ground vehicles, fighter jets and helicopters, the relatively spacious and less rugged environments found in some locations on naval ships provide large enough SWaP to make the use of these commercial appliances practical.

However, some shipboard locations require more solutions that exceed the survivability of rugged PCs and BladeCenters. For example, rugged HPEC technology, rather than commercial HPC processors, is more appropriate for use with the beam-former located in the ship's conning tower where the sensor array is hosted. Ultimately, the use of PCs or BladeCenters in the belly of the Navy ships might be limited as application processing and power requirements continue to grow.

To meet the low latency requirements and SWaP constraints of compute hungry applications, HPEC designers have created systems that feature a mixture of processor elements such as FPGAs, proprietary ASICs, game industry ASICs, and GP-GPUs, all supported with a variety of high-speed, low-latency fabrics such as RapidIO®, Serial Front Panel Data Port (sFPDP), Ethernet and MIL-STD-1553.

On the software side, system designers can leverage proprietary or open source real-time operating systems designed to minimize latency. HPEC system designers increasingly look for software vendors that abstract the subtleties of the system's hardware with middleware. This approach speeds development and lowers costs by maximize the ability to re-use existing code on new systems.

Today, the U.S. military is making large strides in the effort to define middleware baselines. For example, the various branches are supporting initiatives, including the Navy's

OSA (Open Software Architecture), The Air Force's FACE, the Army's VICTORY and the SIG-INT industry's Iron Symphony , that promise to simplify the interconnection between the various electronic functions. HPEC systems for military applications will need to comply with the new standards resulting from these initiatives. While there are hundreds, if not thousands, of acceleration libraries and middleware frameworks currently available for use in the HPC market, the HPEC market will identify and select for its use the software that provides the best performance, price and portability (ease of use). The result will be a widely supported catalog of middleware, operating systems and fabrics.

Embedded Systems Market

Axiom 4: Harsh environment operation

In general, embedded systems are designed to operate outside of the benign environments provided by offices and laboratories. Embedded systems typically function in relatively harsh environments and must satisfy a range of SWaP constraints. A key metric for an embedded system is how economically it performs or operates within those SWaP constraints. Embedded systems vary greatly in how much of a burden they place on available resources (such as battery life and fuel usage) provided by the platform on which they reside. The greater the SWaP of a given embedded system, the more of a burden that system will have on potential mission duration. Also, the greater the SWaP, the more critical the cost of energy to deploy the embedded system.

Axiom 5: Security

The latest and newest axiom to the embedded system market that also applies to HPEC is the requirement for security. The newest major threat to global commercial industries and the military defense industries is the war of information. Driven by the rise of international industrial espionage as well as by global military conflicts, data security and the effective detection of incursions and hostile system intrusions has become a mandate. Data security must be provided for military applications, such as to protect electronic systems that fall into enemy hands, or defend against the theft and duplication of intellectual property. Embedded system designers can apply security to various parts of the embedded system at the box (outer enclosure), PCB board or the module/chip level, but the need for companies to add security has stopped being an option. HPEC systems by their very nature require security.

Axiom 6-7: More than one processor connected by fabrics

All HPEC systems are embedded, but not all embedded systems are HPEC. While numerous companies market ruggedized laptops, PCs and cell phones to operate in harsh environments, only a sub-set of these embedded systems require scalable fabrics to connect more than a single processor. In general, HPEC systems are embedded systems that need more GFLOPS or raw processing capability than one, or even two, commercial processors can easily provide. As a result, HPEC system designs require a communications fabric to interconnect the system

TABLE 1 Embedded vs. non-embedded equivalents

	CONTROL I/O	DSP	HIGH-PERFORMANCE DSP
Rugged	Modules 	Computing density per any SWaP 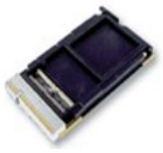	HPEC 
Commercial	Blades 	High-end Blades 	

I/O to the processing units, and connect the processing units to each other. The fabrics used in HPEC systems can be heterogeneous. Depending on the sensor(s) used on the deployed platform, the fabrics supported on an HPEC system are often a mix of analog, or data interfaces such as 1553, sFPDP, Camera-link, Ethernet, RapidIO, and ATM.

However, one factor common to all HPEC system data fabrics is that latency must be carefully optimized to meet the application requirement. Small embedded processor systems that are tightly coupled with the platform's sensors must use a network fabric that leverages the large middleware base of HPC. Otherwise the system will have difficulty scaling to larger systems. Often, smaller HPEC systems will remove middleware to eliminate the related processing overhead in order to save on SWaP. But, this limits scalability.

Moore's Law works against HPEC systems. As technology advances and silicon becomes ever more capable, the amount of work measured in FLOPS increases per square inch and Watt. Applications such as SONAR, that required HPEC-type system architectures in the past, have migrated to small embedded computers with no real need for high-speed, low-latency networks.

Conversely, applications such as RADAR processing, which have over time become less analog-based and more digital, have steadily increased their need for processing power and will most likely require HPEC systems for RADAR deployed in non-office or lab environments.

Sensor Market

Axiom 8: Latency

As one military RADAR/HPEC designer has explained, "the RADAR is badly designed if the RADAR system detects the missile after it destroys you." Thus, the main axiom of HPEC system designers is to minimize latency. The most important component of an HPEC system architecture is the sensor used in the application. Sensors are the elements, arrays, antennas or cameras that start the process of converting the analog world data into digital information. Simply stated, sensor processing drives the need for HPEC systems.

Axiom 9: Sensor-to-processing co-location

Typically, the sensors supported by an HPEC system require a vast amount of processing power directly coupled with a high-speed, low-latency fabric interconnection. A military sensor application cannot accept the amount of time it would take to for the large amount of sensor data to be transported over a slow, non-deterministic network to a large HPC cluster. Therefore, HPEC systems require co-located sensor systems. In the commercial arena MRI and baggage scanners identify threats very quickly and for a time produced sensor data at rates that exceeded the capacity of state-of-the-art networks. These commercial systems have taken advantage of Moore's Law as applied to the networks and processing silicon available. These applications do not typically need more than one or two general-purpose processors.

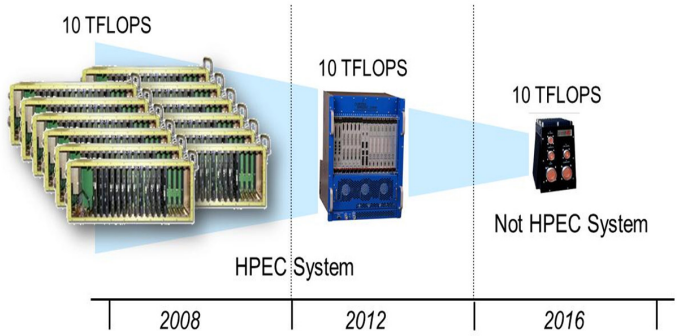

Figure 3: Moore's Law: SWaP – more processing, less space, power and weight

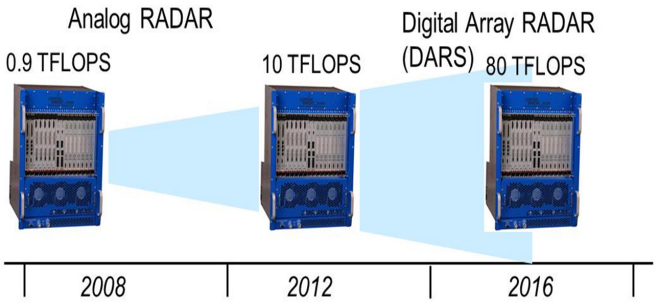

Figure 4: Moore's Law: fixed SWaP

Also, the long distance data networks have improved their bandwidth sufficiently that they are adequately served by small HPC systems in centrally located IT rooms. In these cases Moore's Law has ended the need for HPEC systems for MRI in hospitals and baggage scanners in airports.

For military systems such as aircraft and ground vehicle RADAR, SIG-INT and EW, the high level of SWaP and network speed performance requirements demand that the sensors and processing be co-located. For these applications, HPEC is a must. To explicate the difference between commercial HPC and HPEC requirements, we will examine three systems that have similar functions. In our example, we will consider the air space RADAR application for Ground Mobile RADAR, Airborne Radar and Airport RADAR. All three systems monitor the air space and provide traffic control for the areas where the systems are located. HPEC is required for are Ground Mobile RADAR and Airborne Radar because they need harsh environment operation and the sensor is the same location as the processing system. The HPC center, used in Airport RADAR cannot replace the HPEC system in some environmentally controlled office because the network to access to this system would not be sufficient to meet the latency demands for air-traffic control.

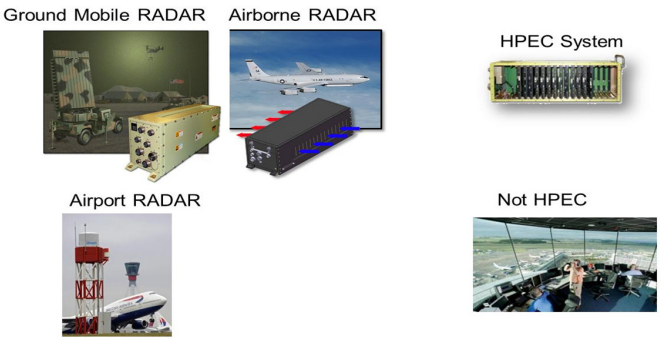

Figure 5: Air traffic control RADAR (HPEC and not HPEC)

Axiom 10: FPGAs are essential

A typical HPEC system can be broken down into five basic functions: HMI (Human-Machine Interface), store, DSP, acquire, and digitize. When we consider these five functions the importance of FPGAs (Field Programmable Gate Arrays) becomes clear. The sub-system in the HPEC system that requires the least intensive bandwidth is the HMI. Arguably, the video, keyboard or mouse interfaces require data rates less than 100 MBytes per second (MB/s) or .1 gigabytes per second (GB/s) since human beings would have difficulty making decisions on datasets that exceeded these interfaces.

The second most demanding function is storage, and is used to store the HPEC system's algorithms, status logs or post-processed sensor data for later review. Storage systems may be part of the HPEC system, as well as part of the test, maintenance, or debug protocols used to verify the functionality of the system. The third function, DSP, is the component leveraged from the commercial HPC arena. It processes the data from the sensor(s) to identify the signal of interest, and produce the resulting data in a size that can either be stored or acted upon by the HMI.

The DSP function can comprise many types of processors, such as those discussed previously, and often include the FPGAs. The acquire function is essentially provided by the application's sensor(s). The function that links the sensors to the other functions is the digitization function, which is the primary domain of the FPGAs.

FPGAs are the silicon of choice for interfacing analog-to-digital converters (ADCs), digital-to-analog converters (DACs), synchronizing or decimating video data streams, or for converting digital sensor data from a proprietary network to a standard data networks supported by the DSP or storage functions. FPGAs are adaptable silicon that provide the first step of filtering for the sensor data to identify the signals of interest. This is because FPGAs provide massive bandwidth capability and are virtually infinitely re-programmable to well-defined tasks in a very efficient manner. The trade-off with FPGAs is ease of use, as these parts are a challenge to program and do not use the readily available tools, math libraries or program tools available from the HPC industry. However, FPGAs are the most flexible programmable silicon interface in existence today.

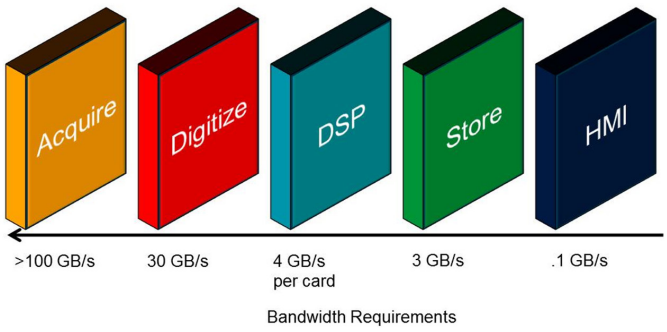

Figure 6: Five functions of HPEC

Faster Fabrics are Needed (and Coming)!

As Each successive generation of CPUs, GPGPUs and FPGAs brings improvements in floating point performance per cubic inch. To keep up, the I/O bandwidth needed to supply data to/from these processing elements must also improve. In 2012, the state of the art processing technology in deployed systems used 5 Gbaud technologies such as Serial RapidIO® (SRIO) Gen2, DDR InfiniBand and PCI Express® (PCIe) Gen2. The good news is that the next generation of faster 8-10 Gbaud fabrics such as 40 GbE, QDR InfiniBand and PCIe Gen3 are being designed into products today and will be available to the market by the end of 2013. Even better, these technologies all employ open standards.

One of the biggest concerns within the embedded industry is whether an OpenVPX™ system can support 10 Gbaud signaling. Curtiss-Wright has done extensive simulations and testing on hardware over the OpenVPX backplane to ensure that OpenVPX systems built on 10 Gbuad signaling are robust.

Summary

The convergence of HPC, Embedded and Sensor technologies is the foundation of HPEC. When these three distinct markets are combined, distinct technological trade-offs result. If any of the HPEC Axioms discussed above is not present for a given application, the system designer may satisfy their requirements with an HPC, embedded, or sensor system rather than using an HPEC architecture. For example, if the application demands the latest fastest, high power silicon, HPEC might not be the right solution because of resulting SWaP constraints. An HPEC customer has limited available space and must address challenging operational environments. For HPEC the packaging capability to cool and power the latest high-powered, 100+ Watt silicon will not work. Therefore, applications that do not have SWaP constraints will naturally migrate to HPC system designs. HPC will suffice if the network capacity is large enough or fast enough to meet application latency, if the customer can remove the processing from the sensor, or can store the data and process that data later.

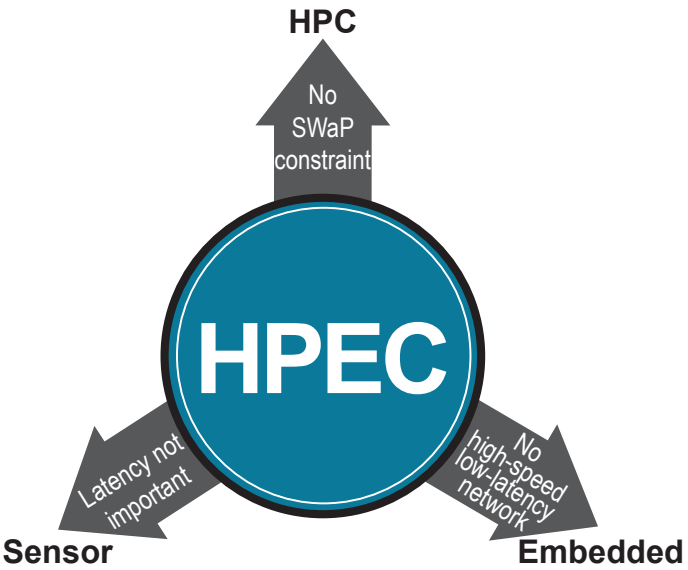

Figure 7: HPEC vs non-HPEC

Lastly, if the embedded system designer can replace an HPEC system with a small SBC combined with a GPGPU to meet the application requirements, then all the fabric and middleware complexity, except for outside-the-box communications, can also be removed. These are the classic quandaries for the system designer: should the designer create a system that can scale to larger capability, and how does that system designer scale the system at the module, chassis or rack level, within a given volume? Does the designer make an HPEC system or remove the complexity to optimize for SWaP? These questions must be addressed because the answers will each present unique trade-offs that may lead to radically different system approaches.

The future of HPEC

With sensor designers making faster, higher bandwidth sensors, such as Tektronix's 12 Gigasample 8-bit ADCs, the amount of data created by these devices will only keep increasing. This causes a technology gap when compared to network technology such as the newer 40/100 Gigabit Ethernet networks that are just now coming to market. This gap is an opportunity for HPEC systems, if the sensor in use must be deployed outside the lab or office environment.

Author

Steve Edwards
Director, Product Management - ISR Solutions & Technical Fellow
Curtiss-Wright Defense Solutions

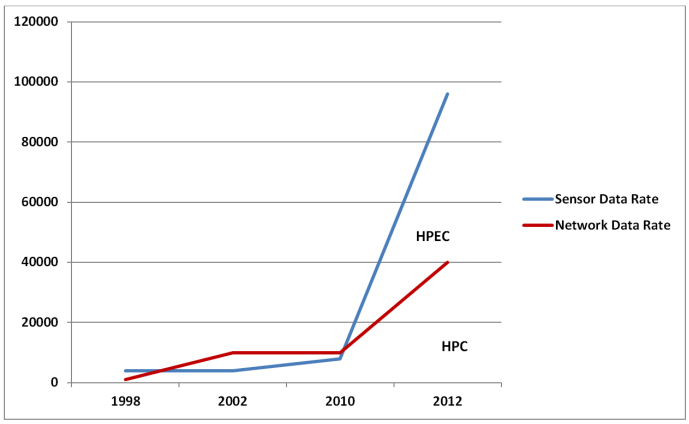

Figure 8: ADC to network technology gap

Realistically, the data acquired by a high bandwidth sensor will overwhelm an HPC network's ability to transfer the data without first performing some decimation or compression. If the processing requirements of that data exceed the capabilities of one or two processors, the system designer should consider deploying an HPEC system. All indications are that sensor data rates will continue to grow faster than available network resources. As this technology gap grows, so too will the demand for HPEC systems.

Learn More

High Performance Embedded Computing (HPEC)

Curtiss-Wright HPEC Development Platform

Press Release: Curtiss-Wright Launches New OpenHPEC Initiative to Bring Supercomputing Software Tools to Embedded COTS Systems

White Paper: 5 Traps to Avoid When Designing COTS HPEC Sensor Processing Systems