

An Integrated Real-Time Platform Can Deliver Improved DSP Performance at Lower Costs

Compared with systems based on DSP coprocessors, an approach that integrates processing cores on a single real-time platform yields economic and operational benefits that are catalysts for the increasing migration from DSPs.

by Andy The, IntervalZero

Digital Signal Processors (DSPs) have specialized architectures that are optimized for heavy math computations and they have dominated the real-time digital signal processing market as a cost-effective solution for many complex designs.

Although DSPs remain a viable choice for many systems, developers are finding that proprietary DSPs are no longer required to perform real-time digital signal processing. Instead, developers are migrating to software-centered Real-time Platforms, which comprise multicore x86 general purpose processors (GPPs), the Windows operating system and an SMP-enabled real-time software extension to Windows. They are leveraging this innovative architecture to outperform DSPs, to significantly reduce costs, and to streamline development cycles.

DSPs rely on specialized architectures designed for heavy math computations, but not for general processing. Because of DSPs' computational focus,

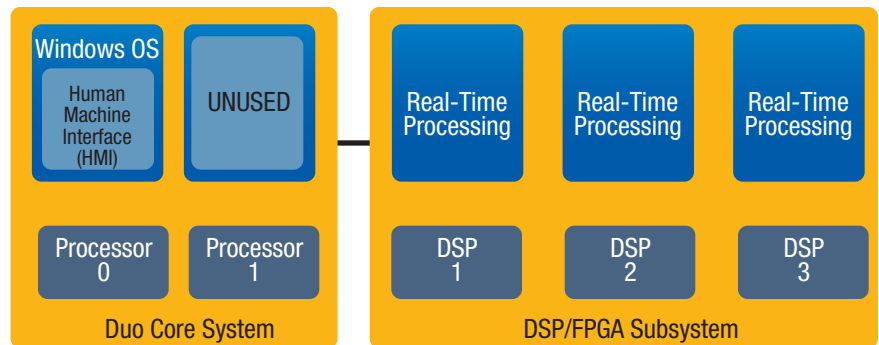


FIGURE 1

A DSP subsystem on a separate device connected to a Windows-based system involves a variety of development tools and issues with the interface and latencies.

GPPs are used for human machine interface (HMI) and general purpose functions such as input and output (Figure 1). Additionally, it should be noted that DSPs require separate memory devices for each device and separate buses, such as PCIe or serial, are needed for all inter-processor communications.

The Real-time Platform takes a different approach, with the multiple cores divided up for dedicated real-time processing and general purpose processing.

Figure 2 illustrates a multicore x86 device with cores dedicated for Windows and cores dedicated to Windows real-time extension for real-time signal processing functions. Unlike the design with a dedicated DSP, memory is shared using the same high-speed bus, allowing for data sharing and inter-processor communications. Table 1 shows a comparison of the two architectures, and in the sections that follow, this paper will discuss the merits of the architectures in detail.



Hardware Design

The x86/x64 multicore advances from Intel and AMD are changing the way system developers meet real-time signal processing needs. New GPPs not only have multiple cores with extremely high clock speeds, but also perform complex math efficiently. They deliver several times the performance of standard DSPs, as can be seen from the benchmarks in the comparison chart. Additionally, multicore x86 devices are delivered with commercial off-the-shelf (COTS) hardware, which is not usually an option for DSP users. COTS hardware rather than custom hardware translates into reduced costs, shorter time-to-market and less risk.

When replacing DSPs with x86 multicore, attention should be focused on the number of cores needed, as well as which x86 processor family best fits the system requirements. As a rough guideline, there can be a 1:1 relationship between DSPs and x86 cores. Because of the extremely high clock speeds with the x86 cores, the design will likely require fewer cores, but the 1:1 ratio is a conservative start, with room for optimization.

Both Intel and AMD have different processor families based on cost and power. For example, Intel's Atom and AMD's Fusion devices are focused on embedded systems requiring low power and cost minimization.

In DSP-based systems, engineers spend a great deal of time selecting the right processor so that the board can be designed and built in time for the software. Having to make hardware commitments early in the design often creates challenges, and any miscalculation can lead to a hardware redesign, which will have a negative impact on the delivery schedule.

When using the real-time platform there are no strict deadlines for choosing and designing the hardware. Because the real-time software extension for Windows is x86 based, standard COTS hardware is available for both development and production. With no custom boards or drivers required for development, the final hardware selection is performed later in the design cycle. This reduces risk and greatly improves the chances for releasing products on time.

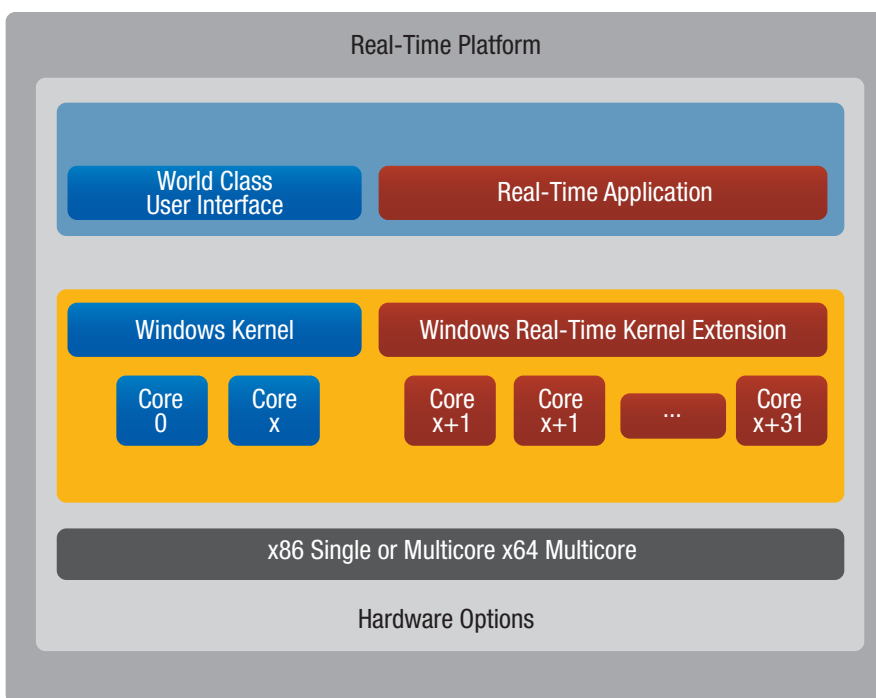


FIGURE 2

Real-Time Platform

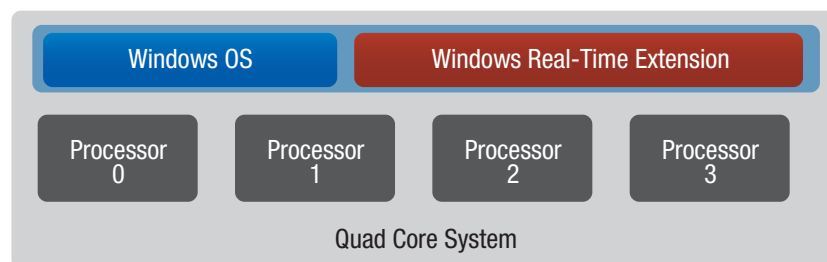


FIGURE 3

An SMP-enabled Scheduler

Also, because the real-time software is based on an x86 Windows architecture, engineers can use PCs as both development and target machines. No separate target system with specialized in-circuit emulators is needed to develop and debug code. As the comparison chart shows, the Real-time Platform's use of COTS hardware reduces both engineering costs and risks.

DSP-based systems rely on custom hardware and software interfaces when communicating between devices. Serial lines and sometimes PCI buses are used for inter-processor communications. These custom interfaces are troublesome to design on custom hardware and can be difficult to upgrade as requirements change.

The Real-time Platform uses shared memory and formalized APIs to communicate between the processes. Because everything is running on a single device, it is easy to share data and messages between threads running on different cores. This communication architecture makes programming easy and scalable.

While DSPs require separate memory devices, with the Real-time Platform memory is consolidated into a single memory device. Also, the x86-based platform can use standard off-the-shelf memory modules, such as 4 Gbyte SDRAM DIMM, and memory modules can be easily changed as system requirements dictate. DSPs require onboard memory chips because of strict timing and routing requirements. Switching out memory chips

Category	DSP Platform	Real-time Platform	RT Platform Benefit
Hardware Design	<ul style="list-style-type: none"> • Custom board design • Long development cycle • Difficult to modify hardware • Requires in Circuit Emulators (ICE) 	<ul style="list-style-type: none"> • (COTS) • Short development cycle • Easy to change hardware • Local Debug – no ICE needed 	<ul style="list-style-type: none"> • Reduced Cost • Reduced Risk • Faster Time To Market
Drivers/BSP	<ul style="list-style-type: none"> • Custom Drivers-device specific 	<ul style="list-style-type: none"> • Standard Drivers 	<ul style="list-style-type: none"> • Cost • Productivity
System Communication	<ul style="list-style-type: none"> • Custom protocol • Custom buses 	<ul style="list-style-type: none"> • Shared memory • Standard APIs 	<ul style="list-style-type: none"> • Ease of use • Future proof
Memory	<ul style="list-style-type: none"> • Chip Down – hard to change • Limited selection • Limited memory reach (1GB) • No MMU 	<ul style="list-style-type: none"> • PnP/COTS Modules – easy to change • Large selection • Large memory reach (>4GB) • MMU – virtual memory 	<ul style="list-style-type: none"> • Reduced system size, cost and complexity
RTOS	<ul style="list-style-type: none"> • Proprietary 	<ul style="list-style-type: none"> • winAPI 	<ul style="list-style-type: none"> • Familiar API
Development Environment	<ul style="list-style-type: none"> • Proprietary Tools • Manufacture specific • Proprietary Compilers 	<ul style="list-style-type: none"> • Standard x86 Tools • Microsoft Visual Studio • Intel and Microsoft Compilers 	<ul style="list-style-type: none"> • Common tool chain throughout the product
Code Base	<ul style="list-style-type: none"> • Assembly and C • Architecture specific 	<ul style="list-style-type: none"> • C,C++, C# • Universal/Portable 	<ul style="list-style-type: none"> • Uses other higher level languages for non real-time
DSP Libraries	<ul style="list-style-type: none"> • Processing libraries from OEM 	<ul style="list-style-type: none"> • Intel supplied IPP library • Other 3rd party libraries 	<ul style="list-style-type: none"> • No requirement to use OEM proprietary libraries.
Benchmarks	<ul style="list-style-type: none"> • 1.25 GHz – Max core speed • 20 GFLOPS/Device Max 	<ul style="list-style-type: none"> • 3.0 GHz – Max core speed • >50 GFLOPS/Device Max 	<ul style="list-style-type: none"> • Faster • Higher performing

TABLE 1

DSP vs. Real-time Platform

is difficult and requires board changes to add or remove memory. The DSP-based approach is more costly, complex and less flexible than the Real-time Platform.

Software Design

Software advances have also had significant impact on DSP-based systems. The demand for standardized tooling in lieu of the proprietary tool sets is a recent change for many DSP developers. Programming DSPs requires proprietary tools and lower-level languages such as C and Assembly. Many engineering teams find that proprietary DSP tools and low-level languages require specialized expertise, which is often very difficult to find and expensive to acquire. The use of standardized development tools and higher-level languages, such as Visual Studio and C++, not only increases productivity but also greatly reduces engineering costs.

The demand for complex graphical user interfaces (GUIs) has also impacted DSP-based systems. Customers are in-

creasingly seeking elaborate touch-and-gesture-based user interfaces on top of a real-time subsystem. DSPs have never had the GUI and I/O support to satisfy most complex system requirements. That is why GPPs running general purpose operating systems are typically found next to DSPs to handle all of the I/O and the complex user interfaces.

As a result of the strong demand for powerful GUIs, Windows is becoming the preferred operating system because of its abundant tools support and its standardization. The use of standardized tools like Visual Studio and the large support structure of Microsoft’s Developers Network make transitioning from DSPs to the Real-time Platform very straightforward.

DSPs use proprietary RTOSs from device manufacturers such as DSP/BIOS from Texas Instruments (TI) or VDK from Analog Devices (ADI). These DSP-based RTOSs are quite capable but have significant limitations because of the lack of communication among the different

DSP cores. Each DSP runs its own application and very little inter-processor communication is used.

The Real-time Platform implements a single symmetric multiprocessing scheduler, such as IntervalZero’s SMP-enabled RTX, across all of the cores in the real-time subsystem (Figure 3). Instead of a separate RTOS and application running on each real-time core, the real-time platform uses a single real-time scheduler to schedule threads across a number of different cores. The flexible SMP scheduling model makes synchronization and communication among the different cores/threads simple and easy to implement.

Load balancing can be easily implemented from the APIs provided by the real-time subsystem. The subsystem’s APIs enable threads to be moved between processors during runtime. This scalability translates well when moving the application to systems with different core configurations. When cores are either added or removed from a system, the user can easily use logic within their application to load balance threads as needed.

DSP manufactures provide proprietary tools such as Code Composer Studio from TI and VisualDSP++ from ADI. While these tools are very useful for developing DSP-based applications, market demand and cost pressures are providing impetus for more standardized tooling. Additionally, DSPs’ low-level tooling and proprietary coding practices do not promote code reuse and portability. This translates into longer development times and higher risks. Many developers are finding that working with proprietary DSP tools is not only challenging, but also costly to maintain.

The real-time Windows extension uses Microsoft’s Visual Studio, which is the industry standard tool for x86/x64-based programming. Finding engineering talent for Microsoft tools is easier and more cost-effective than is the case with DSPs. Leveraging the support and robustness of Visual Studio and the Microsoft Developers Network (MSDN) makes for increased productivity and reduced costs. Most engineering teams prefer the use of standardized tools for their robustness and their support.

Migrating code base to a Real-time

Platform such as IntervalZero's RTX Platform is straightforward. When considering a port, the developer needs to be aware that there are two parts to the application. First, there is the general purpose, or non-real-time processing, and then there is the DSP/real-time processing. The non-real-time/general purpose code will directly port over using a standard Visual Studio project and will run on the Windows operating system. The DSP/real-time code will also be built using Visual Studio, but it will be running on the RTX controlled cores (real-time subsystem).

DSPs are usually programmed in both C and Assembly. While the Assembly code is not portable and will need to be written in C or C++, the C code can be ported over using Visual Studio. The platform enables real-time users to code efficiently in C++ rather than having to start with C and Assembly programming. While DSPs do have some C++ support, object-oriented languages lose too much efficiency to be useful when compiled for DSP architectures.

The RTX Platform's strong support for C++ is a distinct advantage. Through the use of powerful x86 devices, programmers can use C++ to increase productivity. By keeping the code base in higher level languages like C/C++, portability and code reuse can be leveraged to reduce costs and risks.

	DSP	Intel Multicore
Benchmarks	1.25 GHz - Max core speed 20 GFLOPS / Core	3.0 GHz - Max core speed 32 GFLOPS / Core

TABLE 2

The comparison is between aC66xx multi-core DSP from Texas Instruments and an Intel i7 Sandy Bridge device.

Texas Instruments and Analog Devices both provide optimized DSP libraries to help developers with performance and time-to-market. To take the place of these DSP libraries, Intel created the Integrated Performance Primitives (IPP). The IPP library is essentially a collection of optimized functions, such as DSP, imaging, video, etc., for the multicore x86 architectures. The IPP library helps to increase productivity and performance by providing optimized routines for the most common DSP-based functions.

There are a number of ways to measure performance between processors. Because the focus is on DSPs, measuring the number of floating point operations per second (GFLOPS) is used in this comparison (Table 2).

The TI C66xxx @ 1.25 GHz outputs 20 Single Precision GFLOPS per core. The Intel i7 Sandy Bridge processor outputs 32 Double Precision GFLOPS per core. Although the GFLOPS performance is very strong with the TI DSP, it does not match the raw performance and speed of

the Intel processors.

A real-time platform marries the powerful Windows 7 interface with the real-time signal processing capabilities of Intel's and AMD's multicore architectures. Using standardized COTS hardware and standardized tooling greatly simplifies the engineering effort and reduces cost. The SMP-enabled platform increases innovation, portability and scalability while also reducing costs.

There will always be a need for digital signal processing, but because of the many hardware and software advances, dedicated DSPs are an option, not a requirement. For complex systems requiring powerful graphical user interfaces along with real-time signal processing, there are higher-performing, more scalable, less costly options. ▲

IntervalZero
Waltham, MA.
(781) 996-4481.
[www.intervalzero.com].