

# An integrated chipset for embedded communications applications

By Lewis Zook

## The problem

Historically, great strides have been made in the architecture and performance of Intel-architecture chipsets designed for and targeted at the mainstream server, desktop, and laptop markets. While the market for embedded chipsets shares some of the same requirements, it also has notable, unique requirements that are not cleanly addressed by chipsets targeted at these applications.

Some of the applications being targeted and served by embedded PCs in the communications market are:

1. Network storage such as network access server (NAS) and storage area network (SAN). In this area we are still expecting CAGRs (composite annual growth rates) of greater than 25 percent.
2. Security applications such as virtual private networks (VPNs), firewalls, intrusion detection, and virus detection. Viruses and worms continue to impact networks and translate into billions of lost dollars per event.
3. Network management such as SSL acceleration and load balancing. These are driven by the need for higher density in existing rack space along with maximized performance for existing networks.

For the embedded market, the requirements for these applications translate into higher-performance processors and the ability to address large physical memory for fast access to large databases. For these applications a server or desktop type of chipset would seem to fit.

Embedded PCs, especially in communications applications, must generally fit into very tight form factors to achieve the density required in real-world deployments. This means that the chipsets must occupy a minimum of board space, consume very little power, and require little cooling. To

meet these requirements, a chipset targeted for laptop-type applications would seem to be appropriate.

In most communications applications a large portion of the development time is spent in validating and certifying system design. This effort consists of lab testing, customer lab and field trials, and independent lab certification. This effort can constitute a considerable portion of the development expense and directly increases time-to-market.

When a system has been validated and certified, the manufacturer's primary requirement is to maintain a constant bill of material for the hardware portion of the system. Because these systems may have a deployed life of three to five years, the manufacturer needs chipsets that will be available for this time. Since the mainstream PC market is making technology leaps approximately every 18 months in accordance with Moore's Law, very few

mainstream chipsets have lifetimes longer than that.

The PCI bus is the most widely accepted data-transfer standard in embedded PCs, especially in the CompactPCI form factor for carrier-class deployments. When implementing an embedded PC using a conventional chipset for an embedded communications application, there is generally a requirement for a second PCI bus – which may be required for additional I/O devices or for communication to a system-wide bus in an architecture such as CompactPCI. Unfortunately, most mainstream chipset architectures require a PCI-to-PCI bridge on the primary PCI bus for communication to the second PCI bus. This architecture limits the performance of the second PCI bus by requiring all traffic on the second PCI bus to share bandwidth on the primary PCI bus with devices on the primary PCI bus as well as other I/O such as hard drives, USB, keyboard, mouse, etc. (see Figure 1).

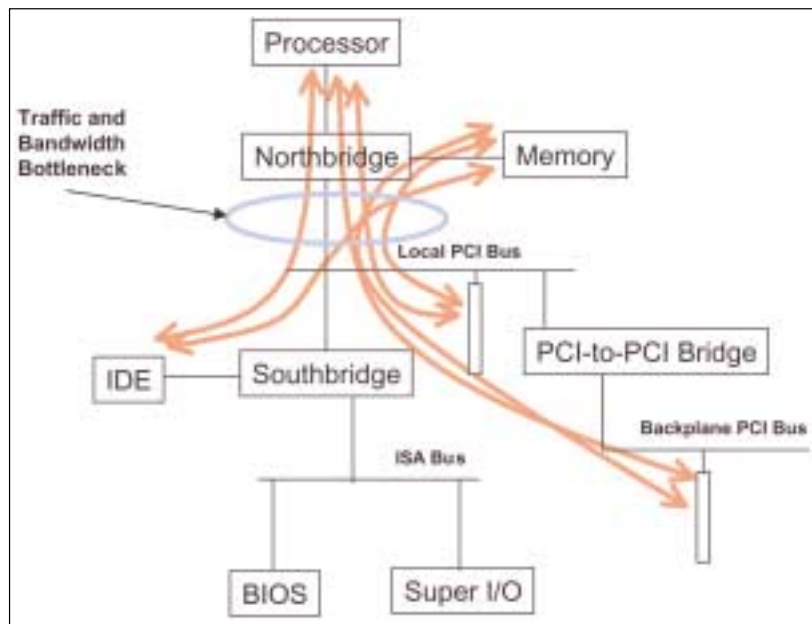


Figure 1. Typical chipset implementation with I/O bottleneck

Clearly a chipset designed to meet the large memory capacity, board space, power dissipation, and PCI bus performance requirements of the embedded communications market is needed. It is also clear that existing mainstream chipsets do not meet all the requirements and are generally not offered as long life products.

### Conceptual solution

A typical mainstream chip solution requires the following four components for a system that supports a second PCI bus such as a CompactPCI CPU card (see Figure 2):

1. **Northbridge:** This is the component that interfaces to the processor, the main system memory, and the primary, or local, PCI bus. The Northbridge also contains the arbiter for the local PCI bus.
2. **Southbridge:** This is the component that interfaces to IDE disk drive devices, an ISA bus, and is a device on the local PCI bus.
3. **Super I/O:** This component is a device on the ISA bus and provides the interfaces to other PC peripherals such as keyboard, mouse, COM ports, parallel port, and USB.
4. **PCI-to-PCI bridge:** This component is required to produce the second PCI bus, commonly referred to as the “backplane PCI bus” in CompactPCI applications. From the viewpoint of the Northbridge, this is also a component on the local PCI bus.

This typical implementation consumes a large amount of board space in an embedded application. With the availability of 0.25micron CMOS technology, a better approach would be to integrate the four major components into a single package that allows for a 50-60 percent reduction in board space. This would make valuable board space available for other application specific I/O or storage features and allow system designers to more easily differentiate their products.

A single-chip implementation of an embedded chipset produces a second tangible benefit. Implementing the four components described above requires layout, board routing, and debug time during the system design. But a single-chip implementation would greatly reduce both the design and debug time of a PC-based

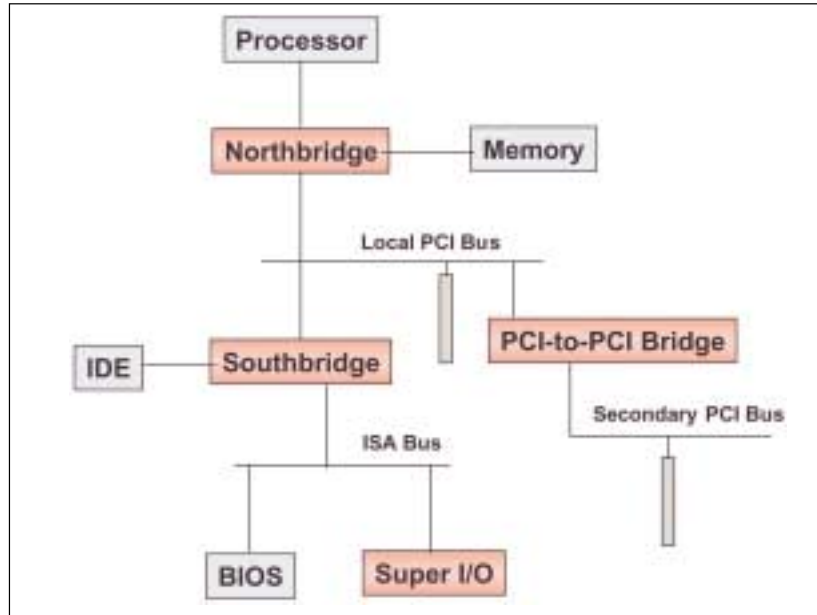


Figure 2. Typical four-component PC implementation

embedded system by simplifying the layout and drastically reducing the number of interconnections that need to be checked for proper routing and signal integrity.

As shown in Figure 1, it is very simple for the local PCI bus to become saturated in the conventional chipset architecture. This bus is responsible for carrying all traffic from devices installed on it, any traffic from the secondary PCI bus along with any traffic from IDE disk drive devices and other I/O such as USB. If you integrate the four components required in the mainstream chipset architecture into a single chip, you can alleviate this traffic bottleneck.

The single-chip implementation offers direct, on-chip access to the both the processor and memory interfaces. With this access, it would be preferable to dedicate an independent I/O path for all the major consumers of local PCI bus bandwidth in the I/O architecture. An independent I/O path could be allocated for the local PCI bus, the secondary PCI bus, IDE disk drive devices, USB, and an independent path for the set of lower bandwidth I/O devices (see Figure 3). These lower bandwidth devices include keyboard, mouse, COM ports, parallel port, and the system management bus (SMBus). Such an independent I/O path chip architecture would improve system performance by eliminating PCI bus traffic bottlenecks. The only limitations to system data transfer performance would be the performance

of the individual buses and the bandwidth of the memory and processor interfaces.

This architecture would also alleviate a serious shortcoming in mainstream chipset implementations – speed limitations on the secondary bus. Because this secondary PCI bus is generated by a PCI-to-PCI bridge that is a device on the local PCI bus, its operation speed can be no greater than the speed of the local PCI bus. If the local bus is supporting lower speed PCI devices – such as video or 10/100Base-T Ethernet adapters – then the speed of the local PCI bus is generally limited to 33 MHz. This directly limits the maximum operation speed of the secondary bus. If the secondary bus is intended to support higher speed PCI peripherals such as SCSI controllers or a CompactPCI backplane bus in the 66 MHz mode, the higher speed peripherals are forced to run at the slower 33 MHz speed.

While many conventional mainstream chipsets support large physical memory there is an issue in the level of memory chip integration supported by these chipsets. Most conventional mainstream chipsets support SDRAM memory densities up to 256 Mbits. For a memory system of 1 Gbyte, this would require a total of 32 memory components – which eats up a large amount of board space in an embedded application. A more desirable implementation in a single-chip solution is to design the SDRAM controller to accept the

latest 512 Mbit memory technology. This will allow a 1 Gbyte memory system to be built with a only 16 memory components.

A second requirement for the memory system is the ability to add ECC – a memory controller technology that allows the addition of eight bits to the memory bus bandwidth. These bits store a code that allows the detection of single- and multiple-bit memory errors and on-the-fly correction of single-bit memory errors. ECC is a stringent requirement for telecommunications applications, especially in the carrier-class market. Very few mainstream chipsets support the capability to add ECC memory capability.

### Implementation

Based on our experience designing highly integrated embedded chipsets for the 80386 and 80486 processors, RadiSys has developed a Pentium III-class integrated chipset specifically targeted at the embedded market. This chipset – the 82600 – is a single-chip integrated implementation of the Northbridge, Southbridge, PCI-to-PCI Bridge, and a selected subset of the conventional PC Super I/O (see Figure 4).

The Northbridge functionality of the 82600 includes a processor interface to the Celeron, Pentium II, and Pentium III processors, including the Tualatin version of the Pentium III. This flexibility in processor selection allows the performance, power, and price point of a system based on the 82600 to scale with the needs and limitations of the application. The 82600 supports 66 MHz, 100 MHz, and 133 MHz processor side bus (PSB) and memory bus frequencies to support this wide range of processors. The 82600 Northbridge section also contains an advanced SDRAM controller. To support high-availability telecom applications, the SDRAM controller supports ECC memory protection. It also supports up to 2 Gbytes of physical memory and 512 Mbit SDRAMs to allow large physical memory to be implemented with a minimum component count.

The Southbridge portion of the 82600 includes a DMA/66-capable EIDE interface capable of interfacing to conventional hard drives, CD-ROMs, and Flash drives. The Southbridge portion also includes the standard PC engine logic including USB controller, real-time clock, standard PC DMA engines, SMBus controller, and timer/counters.

The 82600 also contains a PCI-to-PCI bridge for applications requiring a secondary PCI bus and several additional

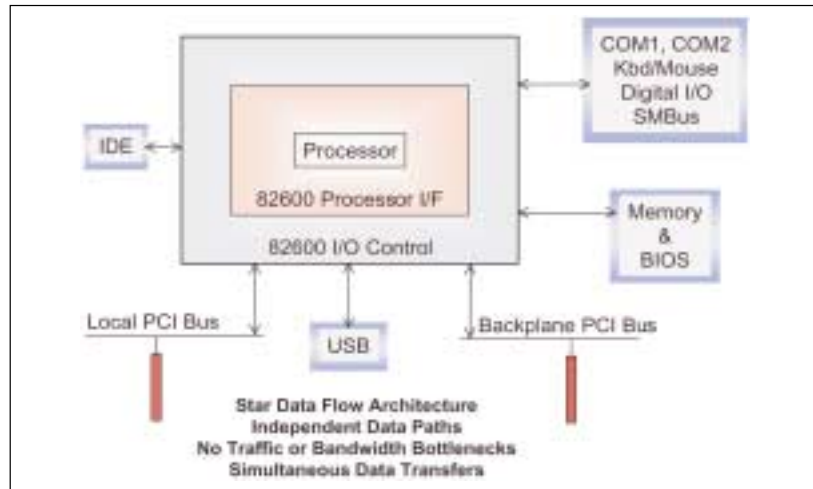


Figure 3. Independent I/O path architecture

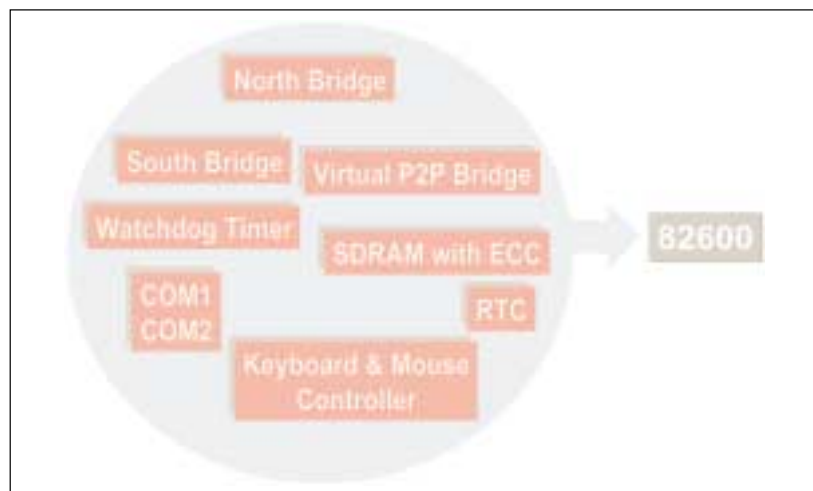


Figure 4. 82600 integrated functional blocks

enhancements over a conventional chipset bridge approach have been incorporated in the 82600. To support applications requiring either a transparent or non-transparent PCI-to-PCI bridge, such as CompactPCI or PrPMC, the 82600 PCI-to-PCI bridge offers two operating modes. Based on the state of a control pin at power-up, the PCI-to-PCI bridge in the 82600 can become either a conventional transparent PCI-to-PCI bridge or become a non-transparent bridge. This capability allows a single-board design based on the 82600 to be used in either the system controller or peripheral bridge slot of a CompactPCI system (see Figure 5 and Figure 6).

With the integration of the PCI-to-PCI bridge on the same chip as the processor and memory interface, one major enhancement was added to improve data transfer performance over the secondary, or backplane PCI bus. In a Pentium III-based system, the data for host-initiated DMA trans-

fers to and from PCI targets cannot be cached. This limitation leads to a maximum data transfer rate for these DMA operations of about 15 Mbits/sec. To overcome this limitation, the 82600 includes four high-speed DMA engines that accelerate the transfer rate for DMA transfers between main memory and PCI targets on the secondary or backplane PCI bus (see Figure 7). With these DMA engines, transfer rates of up to 160 Mbits/sec for host-initiated DMA transfers can be achieved.

All the data paths between the peripherals and the processor and main memory are integrated onto a single chip, thus alleviating the performance limitation that results from all the peripheral data traffic passing through the local PCI bus. All high-bandwidth peripherals have their own I/O path to and from the processor and main memory in the 82600 internal architecture. Traffic to and from the backplane PCI bus, USB, and EIDE

devices does not cause any traffic on the local PCI bus.

A second performance-enhancing benefit arises from this internal data transfer architecture. If lower speed 33 MHz devices such as video or 10/100Base-T Ethernet controllers are installed on the local PCI bus, the performance of this bus is limited to 33 MHz. Since the secondary PCI bus is completely independent of the

local PCI bus, it can be configured to operate at 66 MHz with higher-speed devices, such as SCSI controllers, without being constrained by the local PCI bus performance.

The 82600 is fabricated in a mainstream 0.25micron process. This process was chosen because it is an economically available wafer fabrication process and it lends itself well to producing 5V tolerant

buffers for the appropriate signals on the 82600. As a result of the process used and the power savings from only driving the required signals off-chip, the typical power dissipation of the 82600 is 1.2 watts when operating at 133 MHz PSB and memory bus with both PCI buses operating at 66 MHz. At lower frequencies the power is even less. A typical mainstream chipset would require 4 to 5 watts to operate all four components.

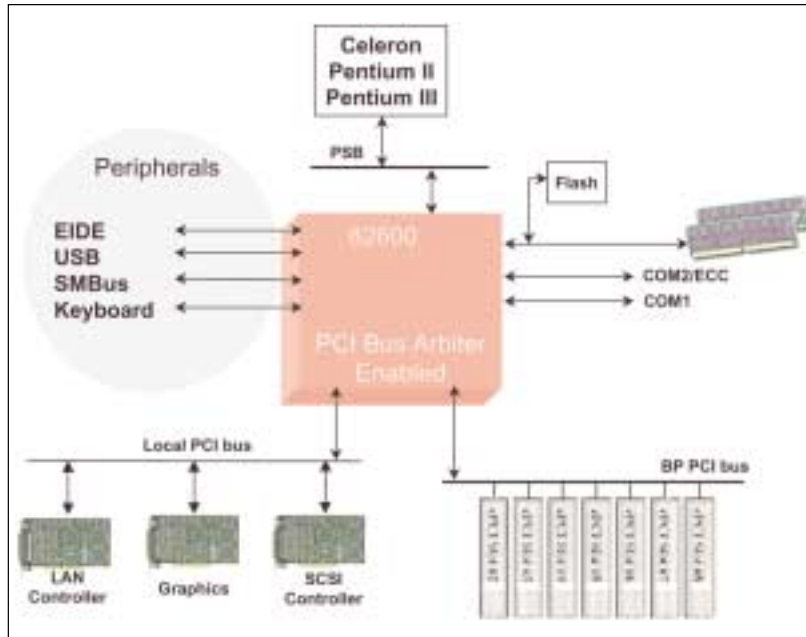


Figure 5. 82600 in the central resource mode

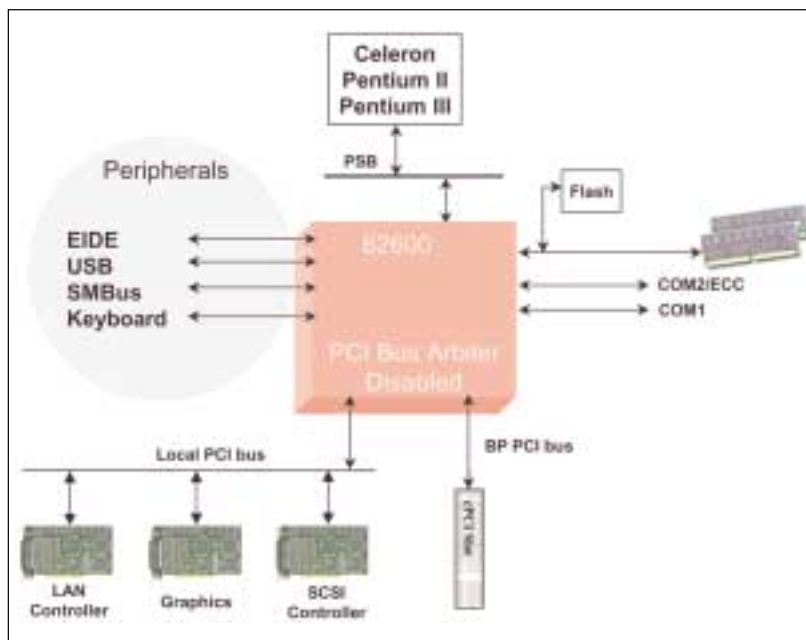


Figure 6. 82600 in the peripheral bridge mode

To conserve on device pin count and cost, the following design tradeoffs were made for the 82600:

1. In general, embedded applications need one or two EIDE devices – to conserve pin count, the 82600 supports one DMA/66 EIDE channel with two EIDE devices.
2. Support for a parallel port was removed from the 82600 because that port is rarely used in embedded applications.
3. Most embedded applications require a COM port to run in headless (or no VGA) mode and few require two COM ports. The 82600 supports one COM port directly and a second COM port through redefinition of device pins with a power-up option strap – which offers a choice between support of a second COM port or ECC on eight of the 82600 pins.
4. The 82600 does not include an ISA bus, which has not been generally used in Pentium III designs for several years. For applications requiring an ISA bus for legacy peripherals, there are commercially available PCI-to-ISA bridge devices. The absence of the ISA bus required that we find another location for the boot Flash. In the 82600, the SDRAM controller has been enhanced to allow X16 Flash devices to be gluelessly attached directly to the SDRAM bus. This results in a simpler overall system design along with the ability to add up to 64 Mbytes of Flash memory area to the 82600.

**Deployment example**

Although there are several designs using the 82600 in process, the commercially available deployment example at this time is the EPC-3331 reference design. This reference design demonstrates the deployment of the 82600 in a CompactPCI layout. It also features a socket processor to allow development and benchmarking of a large variety of processor performance points. To maximize flexibility, the EPC-

3331 has the main memory in conventional DIMMs and has standard PCI connectors on the local PCI bus. With the inclusion of a large variety of test points, the EPC-3331 is an ideal platform for hardware and software development and serves as a proven starting point for system designs using the 82600.



*Lewis Zook is a product marketing manager for RadiSys Corporation in Hillsboro, Oregon.*

Lewis can be contacted at:

Lewis Zook

**RadiSys Corporation**

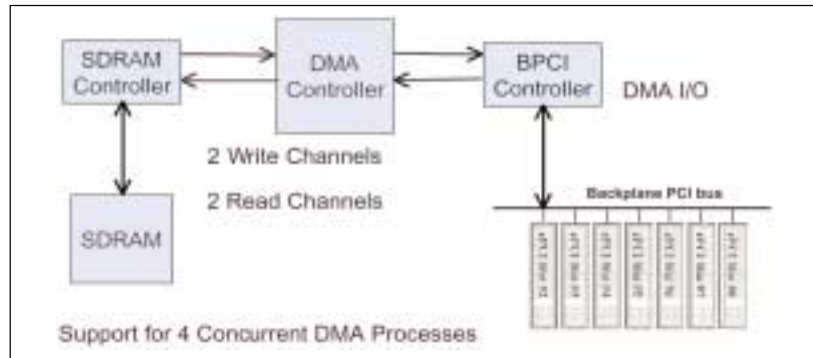
5445 NE Dawson Creek Drive

Hillsboro, OR 97124

Tel: 503-615-1100

E-mail: [lewis.zook@radisys.com](mailto:lewis.zook@radisys.com)

Web site: [www.radisys.com](http://www.radisys.com)



**Figure 7. 82600 high-speed DMA engines**