
IDT[®]
Bridging the PCI Express
Interconnect Divide
Application Note

8000000_AN008_02

September 22, 2009

6024 Silver Creek Valley Road San Jose, California 95138

Telephone: (408) 284-8200 • FAX: (408) 284-3572

Printed in U.S.A.

©2009 Integrated Device Technology, Inc.

GENERAL DISCLAIMER

Integrated Device Technology, Inc. ("IDT") reserves the right to make changes to its products or specifications at any time, without notice, in order to improve design or performance. IDT does not assume responsibility for use of any circuitry described herein other than the circuitry embodied in an IDT product. Disclosure of the information herein does not convey a license or any other right, by implication or otherwise, in any patent, trademark, or other intellectual property right of IDT. IDT products may contain errata which can affect product performance to a minor or immaterial degree. Current characterized errata will be made available upon request. Items identified herein as "reserved" or "undefined" are reserved for future definition. IDT does not assume responsibility for conflicts or incompatibilities arising from the future definition of such items. IDT products have not been designed, tested, or manufactured for use in, and thus are not warranted for, applications where the failure, malfunction, or any inaccuracy in the application carries a risk of death, serious bodily injury, or damage to tangible property. Code examples provided herein by IDT are for illustrative purposes only and should not be relied upon for developing applications. Any use of such code examples shall be at the user's sole risk.

Copyright © 2009 Integrated Device Technology, Inc.
All Rights Reserved.

The IDT logo is registered to Integrated Device Technology, Inc. IDT is a trademark of Integrated Device Technology, Inc.

1. Bridging the PCI Express Interconnect Divide

While the transition from PCI and PCI-X to PCI Express® took some time, few that will argue that PCI Express is now the ubiquitous interconnect available on the latest chipsets and processors. Initial processor offerings provided both PCI Express and PCI/X (i.e. PCI or PCI-X), but PCI/X is quickly disappearing from the processor landscape due to the inherent cost in terms of pin count. However, while processors may have made the transition to PCI Express, the same doesn't apply for many peripherals. This is especially true for legacy I/O peripherals as well as the numerous ASICs and FPGAs that were designed around the PCI and PCI-X standards.

From a block diagram perspective, solving this new interconnect mismatch between PCI Express-based processors and legacy peripherals is straightforward. Simply use a block entitled "PCI Express to PCI Bridge" and the job is done. But putting a bridge in a block diagram is much different than solving the interconnect challenge in a way that doesn't compromise the resultant system in terms of power, performance, cost and reliability.

This paper will discuss key areas of concern for designers selecting a PCI Express Bridge. Specifically, the challenges of successfully bridging between PCI Express and PCI or PCI-X will be addressed. IDT's high quality PCI Express Bridges; the Tsi381, Tsi382 and Tsi384; will be examined along with the features and attributes that make IDT's PCI Express Bridges the best choice for fulfilling unique application requirements.

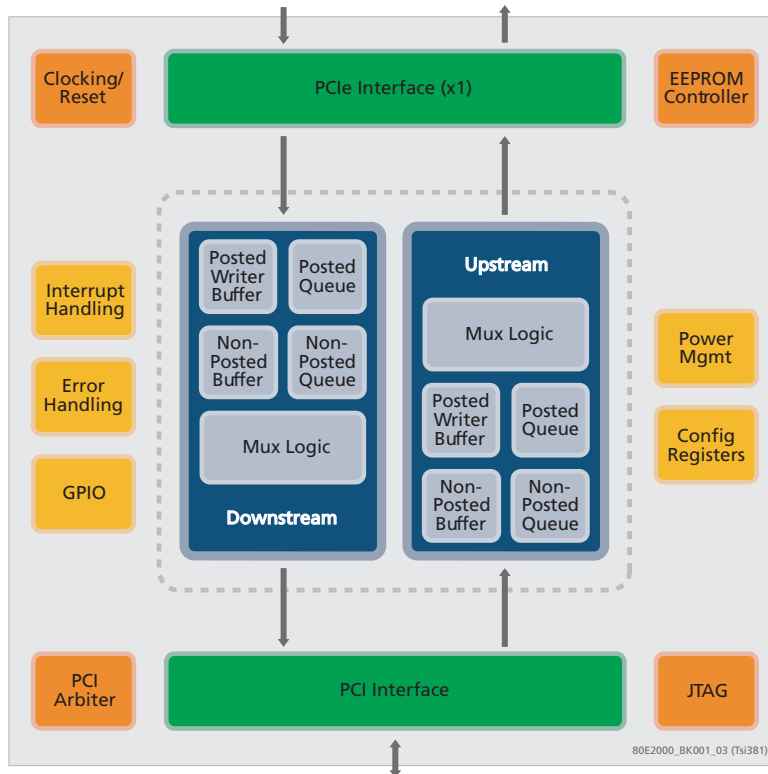
1.1 Tsi381, Tsi382, and Tsi384 PCI Express Bridges

Designing an interconnect bridge is very difficult to do successfully. Not only must different interconnect protocols be properly mapped to one another, mapping must be done in such a way that system level concerns aren't compromised in the process.

IDT has expertise in bridging several standards to PCI and PCI-X (e.g. RapidIO®, HyperTransport™, VME, Power Architecture™, XScale, etc.) and knows what is required to overcome the difficulties of designing a bridge. With a focus on system interconnects, IDT takes innovative approaches to improving bandwidth and lowering latency through its bridging and switching solutions, while also focusing on system level concerns such as cost and power savings.

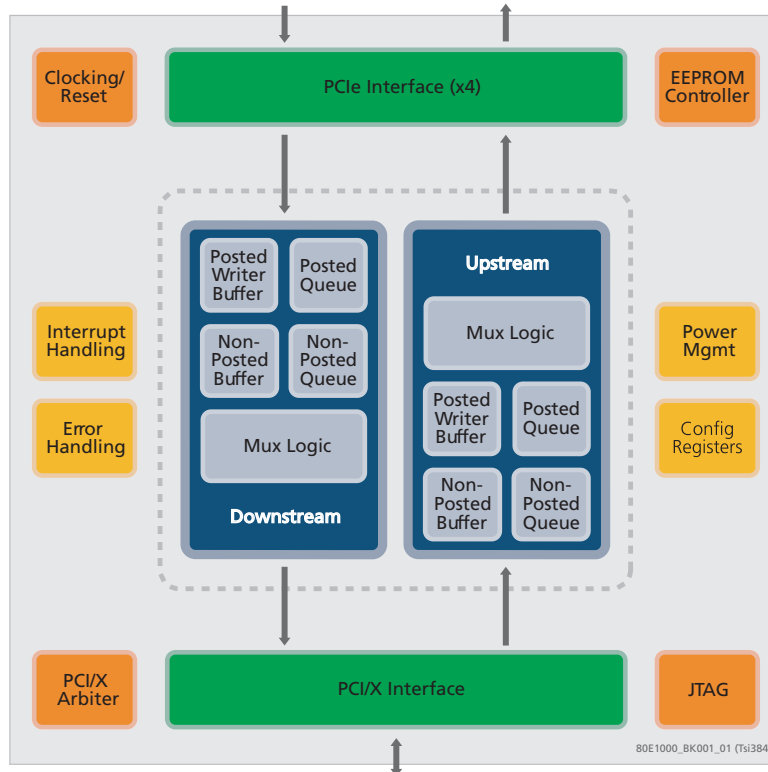
The Tsi381 and Tsi382 were designed to solve the PCI Express to PCI interconnect challenge. Both devices bridge a single lane PCI Express port to a 32-bit, 66 MHz PCI port. The Tsi381 is offered in a 13 x 13 mm PBGA package with a 1 mm ball pitch while the Tsi382 is offered in the smallest available package for a x1 PCIe bridge (10 x 10 mm PBGA, 0.8 mm ball pitch) to assist designers in solving the bridging challenge in a constrained space.

Figure 1: Block Diagram of Tsi381 and Tsi382



The Tsi384 is a high performance bridge that connects up to four lanes of PCI Express to a 64-bit, 133 MHz PCI-X port. With a focus on power and performance, the Tsi384 has proven successful in solving the interconnect challenge in a variety of high performance applications.

Figure 2: Block Diagram of Tsi384



All of these bridges are designed for systems in which the processor or root complex is PCI Express-based. This type of bridging is known as Forward Mode Bridging and is the most commonly required type of bridge in use today. In contrast to Forward Bridging, Reverse Bridging requires that the processor or root complex be resident on the PCI/X bus. By design, the Tsi38x family does not implement Reverse Bridging. While this type of bridging was required during the infancy of PCI Express when processors were still PCI/X-based, the need for this type of bridge has quickly diminished. Focusing on Forward Bridging allowed IDT to reduce device complexity and optimize internal buffer structures, resulting in several benefits that are discussed later in this document.

1.2 The Challenges of Bridging

A PCI Express bridge can solve an interconnect mismatch between PCI Express and PCI/X, but the use of a bridge can introduce unforeseen issues into a design. These issues include areas such as:

- Power Requirements
- Performance
- System Cost
- Reliability
- Sourcing Limitations
- Testability

Choosing the bridge that best addresses these issues is critical to avoiding these pitfalls.

1.3 The Power Challenge

As processing and I/O performance increase, so does the power of the associated processors and I/O peripherals. The overall power consumption of these devices can present a challenge in meeting the power budget of a system. Add a PCI Express bridge to the architecture and in many cases the power consumption of the bridge can push the power of a design outside of the maximum power budget. At the very least, a bridge adds to the overall power consumption of a design and in many cases power is a major competitive feature for a system.

Circuitry associated with the use of a bridge must also be taken into consideration when addressing the power challenge. Clock buffers and power sequencing circuitry all consume power, as does the circuitry required with each power rail used by a bridge.

All of IDT's PCI Express Bridges have been designed with a focus on both device power as well as system power. With multiple PCI/X clock outputs on each of the devices, designers can typically implement a design without the use of an external clock buffer. And with no power sequencing constraints and only two power rails, power supply circuitry can be greatly simplified resulting in lower associated power and cost.

The power consumption of the devices themselves is also quite low. For example, operating with four lanes of PCI Express and a 133 MHz, 64-bit PCI-X bus, the Tsi384 has a typical power consumption of only 1.3W. This is the lowest power consumption of any PCI Express to PCI-X bridge available.

Finally, Active State Power Management (ASPM) allows power to be incrementally reduced on the PCI Express link as the link becomes less active. ASPM has been incorporated into PCI Express as a method to better manage system power consumption. Although ASPM was defined in the PCI Express Base Specification, Microsoft selected PCI Express Base Specification 1.1 as the required level of support to enable ASPM. Both the Tsi381 and Tsi382 fully support Active State Power Management.

1.4 The Performance Challenge

The addition of a bridge into a data path presents an inherent performance challenge. Not only does a bridge represent additional latency that must be taken into consideration, but the internal buffering structure of a bridge can also compromise performance as data packets must be buffered internal to the bridge as part of the transfer from one interconnect to another.

1.4.1 Forward Mode Buffer Optimization

Data movement in a system is typically driven by the peripherals themselves. The processor may orchestrate data movement by setting up DMAs in the peripherals, but it removes itself from data movement at that point in order to focus on other tasks. Processor based data movement is the exception in systems and is typically not performance critical.

Bridges that support both Forward and Reverse modes must structure their internal buffers to account for data movement driven from both PCI/X peripherals (in Forward mode systems) as well as PCI Express peripherals (in Reverse mode systems). This complicates the overall design of the bridge and often results in sub-par performance, increased errata resulting from the complexities of designing a dual mode bridge, or both.

By designing Forward mode only bridges, IDT was able to optimize the buffer structure of the Tsi381, Tsi382, and Tsi384. The results are bridges with superior latency, higher throughput, and much less complexity than many of the other bridge offerings.

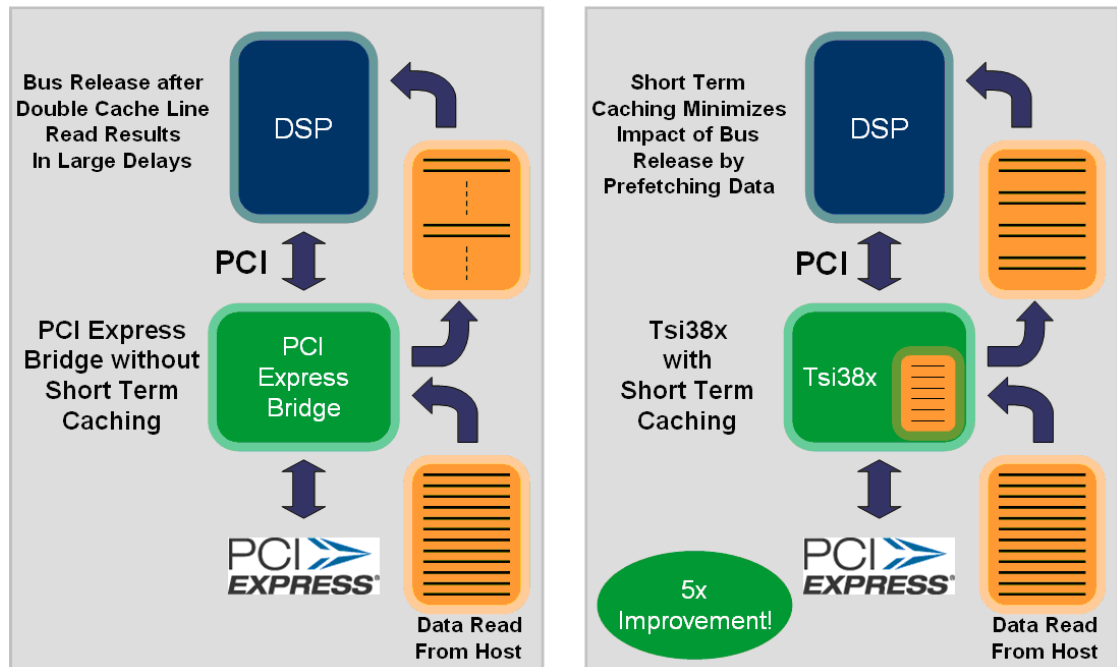
For the performance critical world of PCI-X based peripherals the Tsi384 was designed for a maximum payload size of 512 bytes. This is double the maximum payload size of some of the other PCI Express bridge offerings and can result in a significant performance increase in cases where a system can take advantage of the larger payload size.

As an example of the impact of these performance based buffer optimizations, the Tsi384 offers up to a 48% latency improvement and up to 30% to 45% throughput improvement on Reads over the PEX 8114 from PLX. While this is only one example, it helps to illustrate the impact of a performance oriented design.

1.4.2 Short Term Caching

Aggregate throughput performance is often contingent on the read performance of a bridge. Writes from PCI peripherals to a bridge are typically posted in an internal write buffer to overcome the inherent performance penalty a bridge imposes. However, PCI Reads are almost always problematic. While the PCI-X protocol avoids this problem through the use of Split Transactions, the older PCI protocol does not implement this protocol feature. Add to this the limited PCI Express packet sizes and the fact that some PCI devices were designed to automatically release the bus after receiving one or two cache lines of data and it quickly becomes apparent that there could be a significant performance challenge with PCI Reads.

IDT's PCI Express Bridges offer an innovative feature known as Short Term Caching to assist in overcoming this performance challenge. This optional feature allows data to be prefetched from the attached PCI Express device during an initial PCI Read cycle and temporarily stored in the bridge for quick turnaround during subsequent Read cycles. Figure 3 illustrates how Short Term Caching can offer a significant performance improvement.

Figure 3: Performance Impact of Short Term Caching

Customers who have been able to take advantage of the Short Term Caching feature have experienced up to a 500% improvement in PCI Read performance. While Short Term Caching is not appropriate for all applications, it can make a dramatic impact in many systems with PCI based peripherals.

1.5 The System Cost Challenge

System cost must take into account everything from device cost to power supply circuitry to decoupling capacitors. While the PCI Express bridges available today may have comparable device costs, the impact these devices have on the overall system cost can vary.

1.5.1 Power Supply Costs

The number of power rails required for a typical design has risen in recent years as process geometries have continued to shrink. While 3.3V is common to almost all designs, other voltage requirements demand the use of voltage regulators and associated circuitry.

IDT's PCI Express Bridges require only two power rails – 3.3V and 1.2V. While 1.2V may need to be derived locally, it is a fairly common power rail for 130 nm devices and is typically available in most designs. Also, the requirement for only two power rails is in contrast to the three power rails required by some of the PCI Express bridges offered today.

While the cost of circuitry associated with a power rail may be small, it adds into the overall system cost where every cent counts.

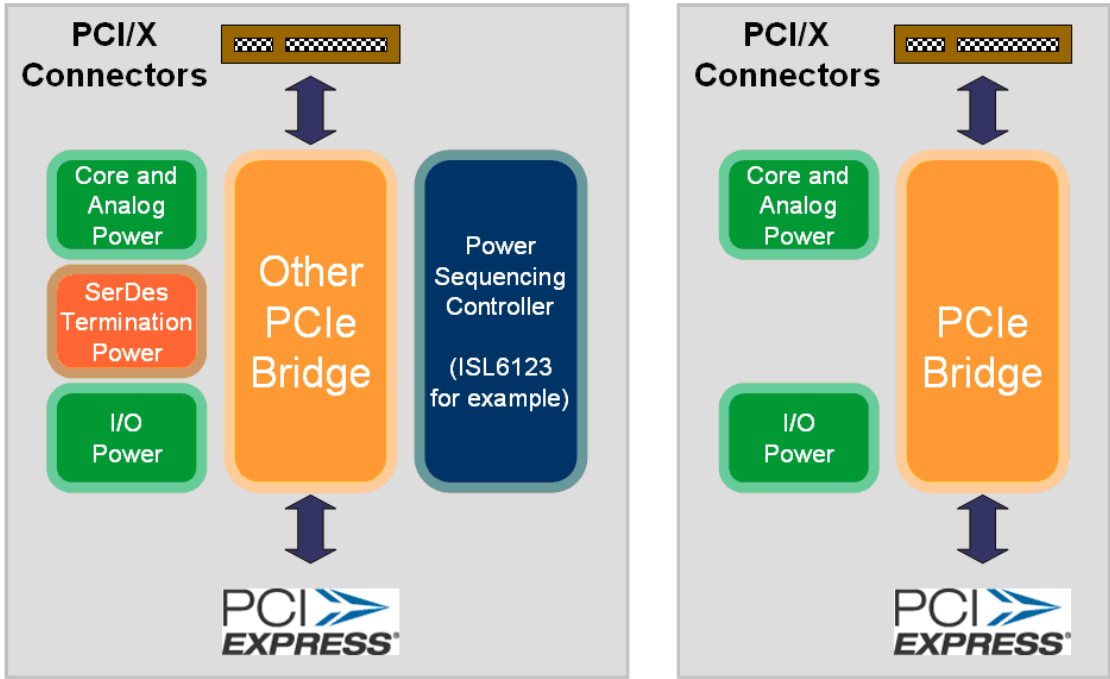
1.5.2 Power Sequencing Costs

Power sequencing requirements can often result in an engineering nightmare. Properly sequencing the voltages for one device can be difficult by itself, but when multiple devices have sequencing requirements the management of the sequencing becomes very difficult and can add both cost and complexity to a design.

The Tsi38x family of PCI Express Bridges was designed with no power sequencing constraints. This provides flexibility in a design that not only simplifies the design process, but can save cost and power associated with sequencing circuitry.

Figure 4 shows a comparison of two implementations of a PCI Express to PCI/X Expansion Card. These implementations are based on an actual customer design. The solution on the left requires three power rails as opposed to the two rails required by the IDT solution and also has power sequencing constraints that result in the need for a Power Sequencing Controller in the design. The IDT solution is not only a simpler design, but a design with lower overall system cost and power.

Figure 4: Case Study of PCI Express to PCI/X Expansion Card



1.6 The Reliability Challenge

The introduction of an additional device such as a bridge into a design adds to the overall risk of a design. Designs based around a single component are subject only to the errata and reliability concerns of the single component. As more components are added into the design, however, designers must concern themselves with an ever increasing list of reliability challenges. At the top of this list is typically the errata imposed by the additional devices, although other device attributes such as operating temperature can impact overall reliability as well.

1.6.1 Errata

A Forward mode bridge versus a Forward/Reverse mode bridge allows for performance enhancements due to buffer optimization as well as reduction in overall bridge complexity, that can result in a reduction in overall device errata.

The Tsi38x family has benefited greatly from this reduction in complexity. Extreme stress testing both within IDT and by some of its key customers has allowed IDT to remove all but the most minor of errata from the devices. IDT's PCI Express Bridges are the cleanest bridges available. At the time of this writing, the production version of the Tsi384 contained only two minor errata and no known errata existed in the production versions of the Tsi381 and Tsi382.

The benefits of minimizing the bridge errata are numerous, but perhaps the most significant are an increase in overall reliability of the PCI Express to PCI/X link and the reduction in debug effort associated with the use of a bridge.

1.6.2 Operating Temperature

The Tsi381, Tsi382, and Tsi384 operate under ambient temperature conditions of -40 to +85 degrees C without the cost penalty associated with many "Industrial Temp" components. This provides a much wider operating environment than "Commercial Temp" devices and could result in increased flexibility with respect to system cooling requirements. The wider range also results in a higher tolerance to temperature variations outside of the commercial range without compromising long term device reliability.

1.7 The Sourcing Challenge

1.7.1 Pinout and Footprint Compatibility with PLX Devices

The Tsi381 was designed to be pinout and footprint compatible with the PEX 8111 and PEX 8112 from PLX. Likewise, the Tsi384 is pinout and footprint compatible with the PEX 8114. While there are device differences that don't allow a direct drop-in replacement of one device for the other, the similarities in the device designs are such that a board can be easily designed to accommodate either device with simple build options. No software changes are necessary. While the advantages of the IDT Bridges over the comparable PLX devices are such that the IDT devices are typically chosen as the primary source, the pinout and footprint compatibility can give customers a build option that provides a second source for the bridge component. IDT provides documentation to assist customers in designing their board to accommodate either device.

1.7.2 The Testability Challenge

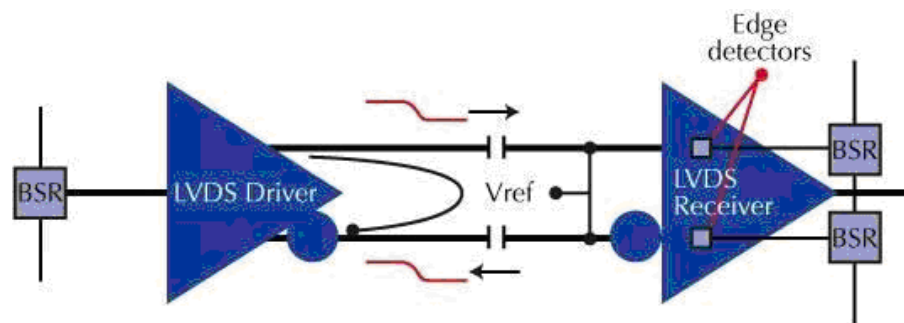
Much like the challenges with reliability, the challenges with testability increase with the addition of new devices in a design. Functional testing must take into account the additional devices and in many cases the testing must take various operating modes into account as well. Boundary scan testing can alleviate some of the issues associated with functional testing by helping to identify manufacturing issues prior to the expense involved with a functional test. However, the introduction of high speed serial links such as PCI Express has introduced challenges to testing via boundary scan.

1.7.3 Advanced JTAG Controller

While standard JTAG (IEEE 1149.1) is sufficient for the boundary scan testing of digital interfaces, it is insufficient for coverage of AC-coupled signals such as PCI Express.

All members of the IDT PCI Express Bridge family have a JTAG Interface that is compliant with both IEEE 1149.1 Standard Test Access Port and Boundary Scan Architecture as well as IEEE 1149.6 Boundary Scan Testing of Advanced Digital Networks standards. In fact, the Tsi381 and Tsi382 are the only x1 PCI Express bridges on the market that are compliant with IEEE 1149.6. The IEEE 1149.6 standard extends the capability of the IEEE 1149.1 standard to include AC-coupled and/or differential nets. This allows for JTAG testing of the PCI Express portion of the design.

Figure 5: IEEE 1149.6 Test Circuit



1.8 Summary

While a PCI Express bridge can solve the interconnect mismatch between PCI Express based processors and legacy peripherals, it also presents a new set of challenges in terms of power, performance, system cost, reliability, sourcing, and testability.

With industry-leading expertise in System Interconnect, IDT has designed a family of PCI Express Bridges to overcome these challenges. Examining the problems that need to be solved and focusing on solutions that address these problems has resulted in devices that offer a breadth of features and attributes that no other bridge in the market can offer. Designers are increasingly taking advantage of IDT's solutions to solve the PCI Express interconnect challenges in a number of applications, including server and storage, security DVR cards and systems, and many other embedded computing and PC applications.



CORPORATE HEADQUARTERS
6024 Silver Creek Valley Road
San Jose, CA 95138

for SALES:
800-345-7015 or 408-284-8200
fax: 408-284-2775
www.idt.com

for Tech Support:
email: ssdhelp@idt.com
phone: 408-284-8208
document: 8000000_AN008_02