Designing and Validating High-Speed Memory Buses

Application Note 1382-2

by Gregg Buzard, FuturePlus Systems
and Perry Keller, Agilent Technologies

Who Should Read This Application Note?

Digital design engineers working on high-speed buses. This application note covers new tools and measurement techniques for validating general high-speed buses and specifically covers DDR (double data rate synchronous dynamic random access memory).

High-Speed Design Requires New Techniques

The obvious technology trend today is higher clock speeds. But the many related changes have an equal or greater impact on designs. Faster clock speeds require smaller voltage swings and shorter setup and hold times. Suddenly, data-valid windows are orders of magnitude smaller. The decreasing size of the data-valid window means that jitter-induced noise, crosstalk, and intersymbol interference further reduce its size, creating errors. Because the noise margins are so small, noise and timing budgets can no longer tolerate phenomena that were previously ignored. LVDS (low voltage differential signaling), double-pumped clocks, and point-to-point designs can also create new challenges.

Other improvements have resulted in yet more design challenges. Clock speeds are now reaching frequencies formerly used only by RF and microwave engineers.

Higher frequencies result in shorter wavelengths, so board layout becomes critical. Designers have to watch for breaks in the ground plane, match impedances to reduce reflections, and worry about trace separation, trace length, and even discontinuities in the FR4 material.

Today's high-speed bus designer needs to carefully consider proper termination of signals and correct impedance matching. Buses are no longer clocked with one or two signals; source-synchronous bus designs create situations where there can be a dozen or even more clocks or strobes moving the data. Rambus® designers have to deal with training or calibration algorithms to compensate for thermal drift. With InfiniBand and RIO (rapid IO), buses are moving from traditional multidrop buses to switched fabrics. While this may be common to ATM (asynchronous transfer mode) switch designers, for many this is a major shift in technology.

Table of Contents

Agilent Technologies

The New Challenge of DDR

DDR SDRAM (double data rate synchronous dynamic random access memory) is quickly becoming an accepted technology in the PC (personal computer) industry. Its low cost, high performance, and increasingly wide availability make it very desirable for PC memory buses and embedded designs such as high-end graphics systems. As more key industry players announce their intention to support DDR, you may find the need to implement such a memory system in your future.

DDR SDRAM dramatically increases DRAM bus speeds to 266 MHz by activating data output on both the rising and falling edges of the system clock rather than on just the rising edge. This potentially doubles the data output rate. However, the design and validation of DDR buses present new challenges to design teams and tool developers; the techniques used for slower buses, such as PCI-X, will not work for DDR.

Unlike PC100 and PC133, which have a common clock, DDR is a source-synchronous bus. Instead of one clock, the bus now has 19—one for the control signals and 18 strobes (or clocks) for the data groups. Each data group has its own strobe that is most likely skewed from the other data groups. To make things even more interesting, DDR write data is shifted 90 degrees from the read data. The only similarity between PC133 and DDR is that they are both multipoint buses (there are multiple devices on the bus that have to arbitrate for bus access).

	PC133	DDR
Speed	100/133 MHz	200/266 MHz+
# Clocks	1	Up to 19
Eye Size	2V x 4 ns	700 mv x 1.5 ns
Clock Edges	Single	Double
Clock Method	Common	Source Synchronous
Interconnect	Multipoint	Multipoint
Timing	Centered	Centered / Straddle

Table 1. PC133 versus DDR SDRAM.

Table 1 shows the similarities and differences between PC133 and DDR. The key difference is the size of the “eye”, or data-valid window. The DDR eye is 700 mv by 1.2 ns, resulting in a data-valid window that is only one fourth the size of PC133’s data-valid window, so it takes significantly less energy to distort a signal and create a glitch on the DDR bus.

Figure 1 is an eye diagram of a single data channel. The first data bit has a data-valid window with an amplitude of 700 mV and a data-valid time of 1.5 ns, which is within the specification. But a few bits later, a small amount of jitter has reduced the data-valid window to 1 ns.

Figure 1. This eye diagram of a single data channel shows how the data-valid time can shrink from 1.5 ns to 1.0 ns with just a small amount of jitter.

Causes of Jitter

With the reduced size of the data-valid window, events that create jitter need to be dealt with. One cause of jitter is intersymbol interference (ISI). At higher frequencies the voltage may not have time to reach its maximum amplitude, as when a 0 changes to a 1 and then back to 0. But when a series of 1s occurs, for example 01110, the voltage does have time to reach the maximum voltage (figure 2).

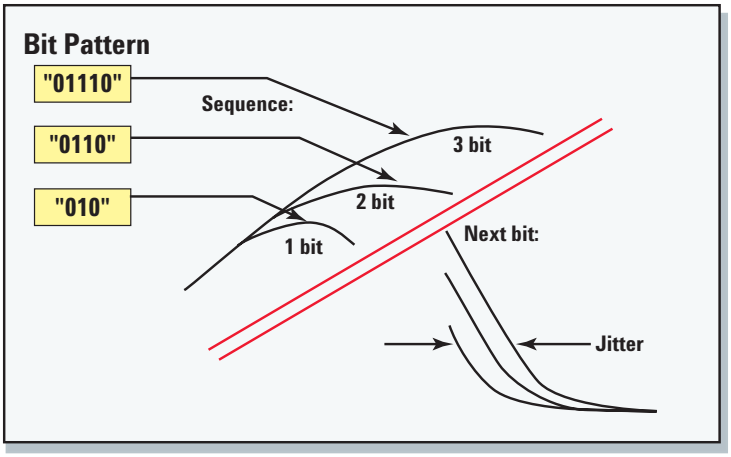

Figure 2. Intersymbol interference occurs when the signal may not have time to reach its maximum amplitude. A bit pattern with more consecutive 1s results in a higher maximum voltage.

ISI also causes jitter on the falling edge. A data pattern of 01110 is going to take longer to go from the maximum voltage to the threshold for a 0 than a pattern of 010, which never achieved the maximum voltage and will reach the threshold faster. For example, pattern 010 takes 1 ns for the signal to go from a high (1) voltage level and cross the voltage threshold, whereas a pattern of 01110 takes 1.5 ns to go from the maximum voltage to the threshold for a 0. This example results in 0.5 ns of jitter. Since the data-valid window for DDR can be as little as 1.2 ns, a jitter of 0.5 ns is significant.

Another source of jitter is crosstalk. Higher frequencies cause even small amounts of inductive and capacitive coupling to transfer a significant amount of energy between traces. The energy coupled into the victim line acts as a voltage offset and results in jitter.

Stubs

At DDR speeds the effects of stubs cannot be ignored. Figure 3a shows a classical load-terminated bus. At frequencies below 100 MHz, the stubs are short relative to the edge speed of the signals on the bus. This design works fine because any ringing has time to settle out before reaching the device.

Even at frequencies of 133 MHz, damping resistors must be provided, as shown in figure 3b. The value of 22 ohms is listed in the DDR specification. The 10-ohm termination on the backplane is not a set value and is best determined by SPICE modeling.

Clock signals should look like the one on the left of figure 4, which was measured at the DRAM. The signal on the right is obviously not desirable, but the reality is that even a small amount of reflection can cause distortion like this. This distortion could be caused by reflections from improperly terminated stubs.

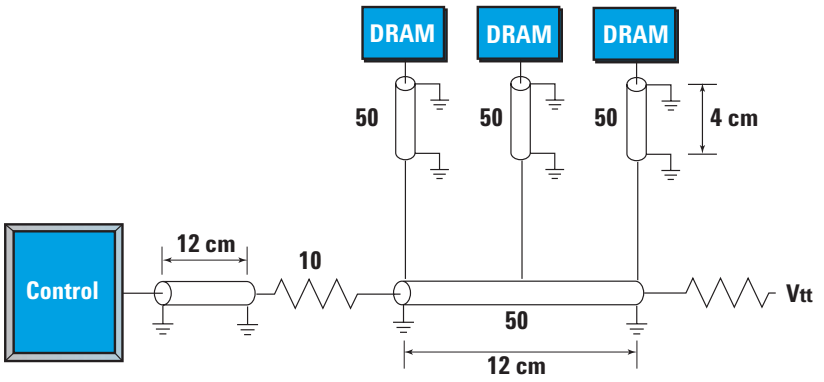

Figure 3a. At frequencies below 100 MHz, the stubs are short relative to the edge speed of the signal.

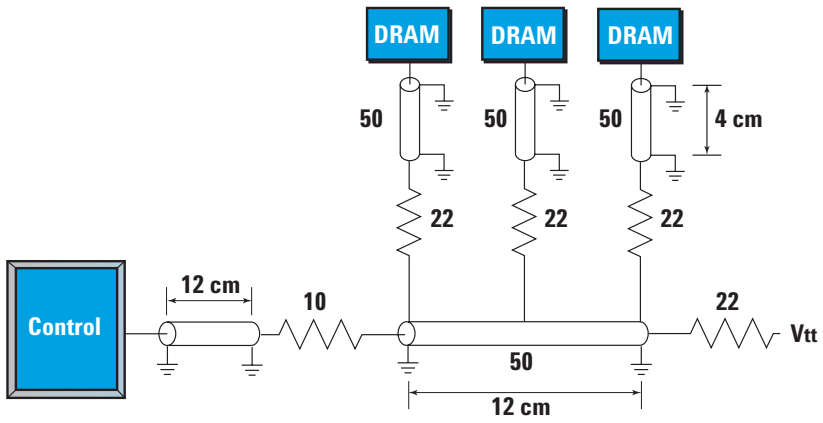

Figure 3b. At DDR speeds, you cannot ignore the effects of stubs.

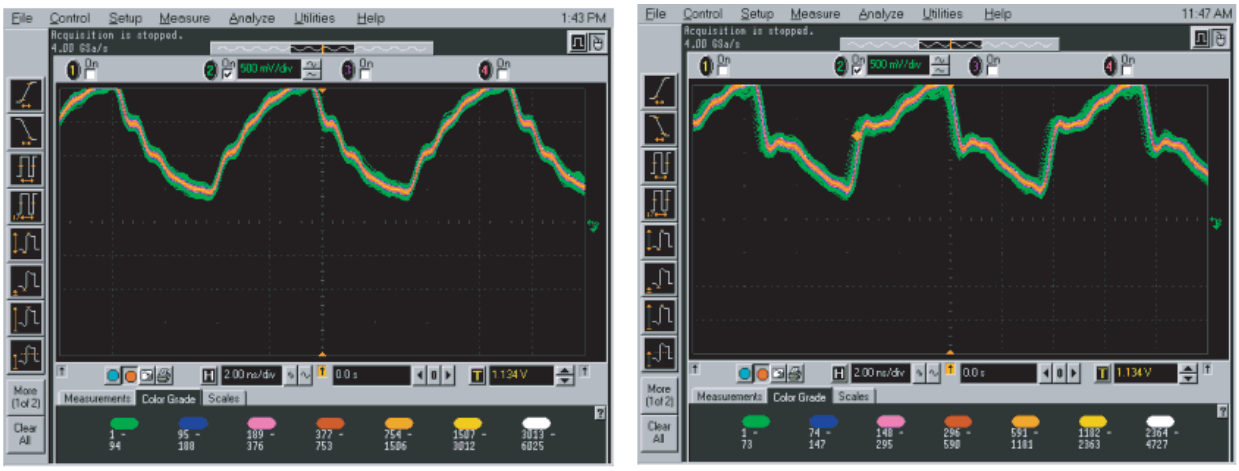

Figure 4. A properly designed signal path results in a clean signal as shown on the left. The distortion of the signal on the right is caused by a small amount of reflection.

Beyond 266 MHz

As bus speeds move beyond 266 MHz, there will be even more challenges to overcome. For example, the specification for Rambus running at 800 MT/s (megatransfers per second) requires that traces must be routed like Via 5 and Via 6 in figure 5. If a signal is routed as shown between Via 2 and Via 3, a stub is created, violating the Rambus specification and causing unwanted reflections. At InfiniBand and RIO frequencies, test points as illustrated with Via 1 turn into antennas, creating more signal integrity problems.

To optimize a design with a high-speed memory bus, it is necessary to run full SPICE simulations of the bus with 3D modeling software. There is never one perfect value; all choices are tradeoffs between noise margin and timing.

Higher-performance buses are driving changes in debugging, verification, and validation methodologies. With DDR implementations, signal integrity is critical because of the small size of the data-valid window. As frequencies increase, designers need to be aware of how jitter further reduces its size. Phenomena such as crosstalk and intersymbol interference need to be taken into account. The higher frequencies mean shorter wavelengths, so traces need to be treated as transmission lines and even package leads need to be considered as potential antennas.

So an increase in bus speeds requires many changes to design methods. These same changes affect the way buses are tested and measured.

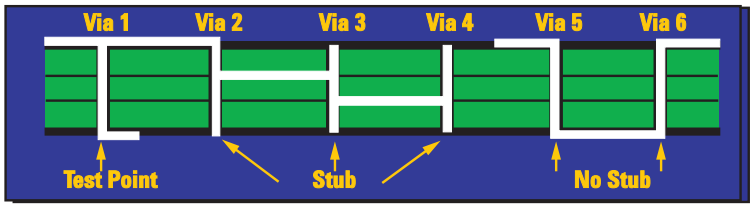

Figure 5. As bus speeds increase beyond 1 GHz, vias can become stubs and test points can become antennas.

Minimizing the Effects of Probing

One serious issue affecting testing is probing. Everyone is familiar with how the Heisenberg uncertainty principle applies to measurements: whenever a measurement is made, the connection of the test equipment affects the measurement to some extent. The effect of probing on a circuit becomes more pronounced as the frequency increases. A good test design must minimize the impact of the test probing on the measurement.

Figure 6 shows circuit models of passive probes and how each impacts the edge of a signal. As the signal frequency (or the frequency components of the signal edge) increases above 50 MHz, the stray inductance of the probe lead has a greater effect on the circuit. Combined with the capacitance of the probe, this causes the signal to start to ring. At higher frequencies such as 266 MHz, you can't just connect a typical probe; you need to probe with a properly terminated transmission line. Probing and proper termination are critical to accurate measurements.

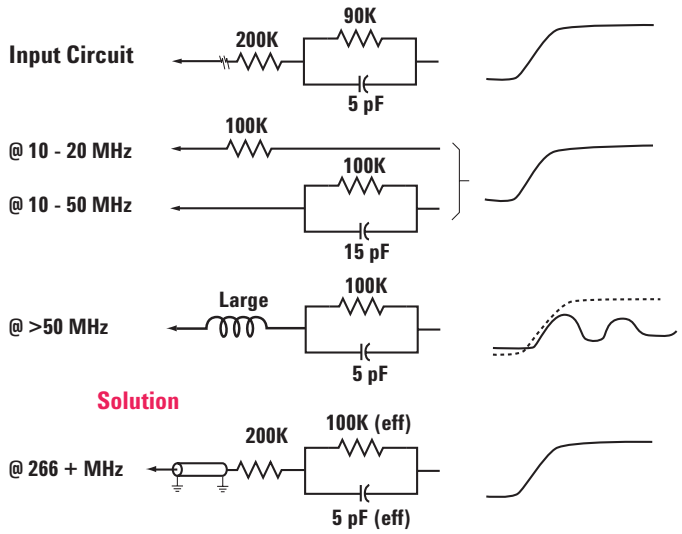

Figure 6. As frequencies increase, you eventually have to resort to transmission lines to probe your circuits.

Test tools need to be able to accurately clock the data even when the signal is distorted. The data strobe in figure 7 is used to clock the data into the logic analyzer. The Vref line shows where the standard transition between 0 and 1 is detected. If the probe uses only Vref to determine the transition point on the rising and falling edge, it will clock improperly. This is due to the wide shoulders on the signal causing a false clock because the output of a comparator glitches. To solve this problem and capture accurate data, the probe should use a dual-threshold trigger (a Schmitt trigger). This probe has special circuitry that causes it to wait, ensuring proper clocking.

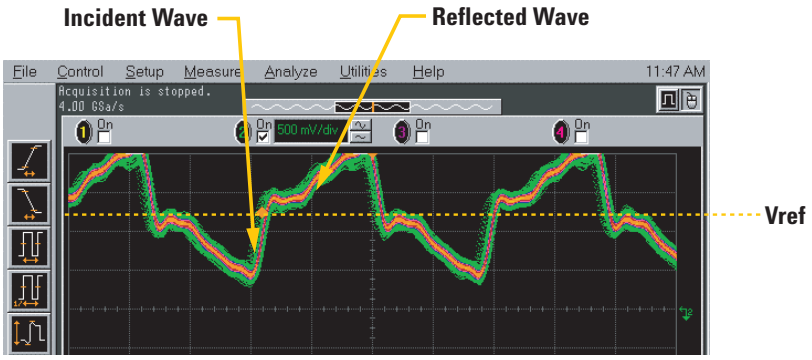

Figure 7. Test equipment must be able to accurately clock data even when the signal is distorted.

Advanced Tools Provide a Clear View

To accurately measure a high-speed bus, engineers need tools that can help them quickly see the behavior and performance of the system under test without affecting system behavior or creating problems that aren't there. Advanced analysis probes are an important component of the tool set.

Probes

Analysis probes can be designed as either active or passive. As with everything else in engineering, this means making tradeoffs. The ideal analysis probe would be easy to set up, have powerful triggering, and be able to properly clock on small data-valid windows (advantages of an active probe). At the same time it would be reasonably priced and, most importantly, would be available as soon as silicon was available (advantages of a passive probe). It must also provide accurate timing information and accurate state data by dealing with the clock issues, and should not create reflections on the bus. Moreover, it must combine the power of an active probe and the flexibility of a passive one.

Eye Finder Measurement Technology

Source-synchronous data buses create an entirely new challenge. Unlike the good old days of the Motorola 6800, valid data cannot simply be clocked on the rising edge of Address Strobe (AS). For example, the DDR bus uses 18 source-synchronous clocks or strobes. Each strobe has 4 data lines assigned. However, a logic analyzer cannot sample using 19 clocks—only one can be chosen. Furthermore, skews between sections may be large, and the data-valid windows are tiny. Realigning the data bus and clocking the data in at the right time are challenges for both the high-speed bus designer and for the test equipment measuring the bus (figure 8a). It is possible to do this manually but it would take many hours, even days, to accomplish. Technology to automate this process would greatly speed up testing.

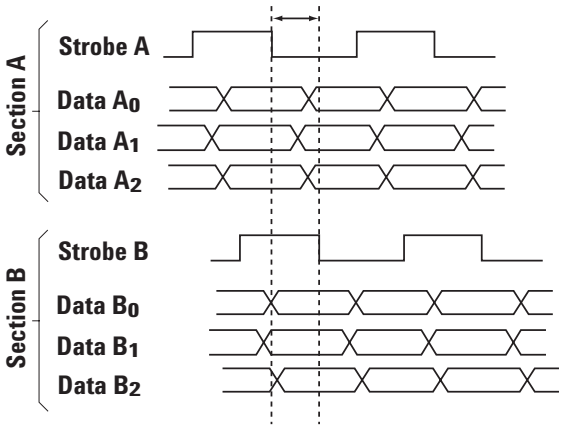

Figure 8a. It is not possible to sample with all 19 clocks of the DDR bus to compensate for the skew between sections. Using a single clock to sample all the skewed bus data can be a challenge.

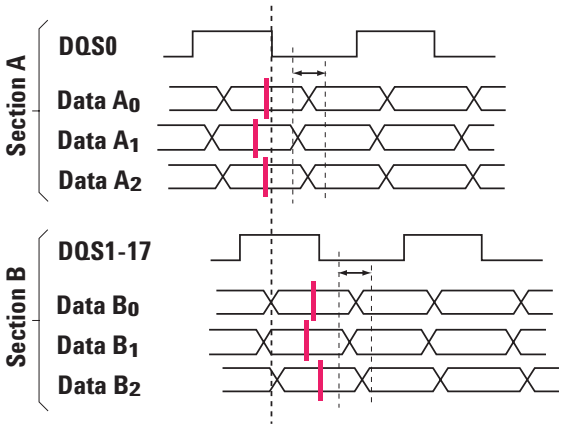

Figure 8b. Agilent's eye finder technology automatically solves the channel-to-channel skew problem in a matter of minutes.

Figure 8b shows how Agilent Technologies' automated "eye finder" technology can solve in a matter of minutes, with 100-ps resolution, the channel-to-channel skew problem created by a source-synchronous bus. The sample and hold setting for each channel is aligned with DQS0.

There is an added benefit to using the eye finder feature. Because the user can see the size of the eye and can visually compare it to all the other channels, detection of potential problems is easy. If one eye looks much smaller than other eyes it might be a problem area. The user can then take a scope and make an exact measurement of the size of the eye to obtain more detailed and quantitative information about the possible problem.

Summary

Probing is no longer just clipping a probe to a test point. The probing system must be specifically designed for the bus under test. New tools such as eye finder technology are needed to deal with source-synchronous buses.

An additional challenge is that a bus does not operate in a vacuum, but is always interacting with other buses. The CPU front-side bus (FSB) may request data from a SCSI disk drive and place that data in memory. If a problem occurs in getting that data into the memory, discovering the problem's source can become a quagmire. Obviously it is not enough to look at just the DDR memory bus or just the SCSI bus. The ideal solution would be to look at all of the buses involved, determine if they are working, and even look for performance issues. This capability is referred to as cross-bus analysis.

So system and tool developers must cope with a myriad of changes in design and testing techniques when they begin working with advanced high-speed buses. Leading-edge engineers need to seek out the most advanced probing and logic analysis tools to solve today's and tomorrow's complex design challenges.

Probe and Logic Analyzer Team Up to Tame DDR

Analyzing a bus such as DDR is simplified with a combination of the Agilent 16700 Series logic analysis system and the FuturePlus FS2330 DDR DIMM (dual in-line memory module) probe. Since DDR is a complex source-synchronous bus, it is important that the analysis probe and logic analyzer manage the complexity for the user. The FS2330 provides all of the setup files for the clocking protocols. It also provides full inverse assembly of the data and command lines. In addition, Agilent's patented eye finder technology provides automatic alignment of the clock and data with 100-ps resolution.

Because the probe provides termination as close to the DIMM slot as possible, the impedance mismatches between the probe and the DUT (device under test) are minimized, and the need for active probing is eliminated. Because there are no active circuits in the path, accurate timing can now be measured. The logic analyzer cable connections and terminations are directly connected to the board. This simplifies the connection to the logic analyzer, provides proper impedance matching, and creates a reliable mechanical connection.

The FS2330 provides full-speed, 266-MHz state capture and timing information at 2 GHz on a single data capture. Simultaneous capture of state and timing data provides quick and easy insight into the behavior of the DDR bus. In addition to passively probing all of the signals, this analysis probe has the flexibility required to deal with many of the different DDR configurations.

The most important feature of the FS2330 analysis probe is its size, both physical and electrical. Physically it is the same size as a standard DIMM. Because the flexible cable comes out of the top, it can be placed in any slot. More importantly, its effect on the bus is the same as a registered DIMM, so it does not impact the behavior of the bus. Not only does the bus run at full speed; it behaves the same with or without the FS2330 analysis probe. If your design uses unregistered DIMMs, a switch is provided on the probe to accommodate the change in clock latency.

Agilent Technologies' Test and Measurement Support, Services, and Assistance
Agilent Technologies aims to maximize the value you receive, while minimizing your risk and problems. We strive to ensure that you get the test and measurement capabilities you paid for and obtain the support you need. Our extensive support resources and services can help you choose the right Agilent products for your applications and apply them successfully. Every instrument and system we sell has a global warranty. Support is available for at least five years beyond the production life of the product. Two concepts underlie Agilent's overall support policy: "Our Promise" and "Your Advantage."

Our Promise
Our Promise means your Agilent test and measurement equipment will meet its advertised performance and functionality. When you are choosing new equipment, we will help you with product information, including realistic performance specifications and practical recommendations from experienced test engineers. When you use Agilent equipment, we can verify that it works properly, help with product operation, and provide basic measurement assistance for the use of specified capabilities, at no extra cost upon request. Many self-help tools are available.

Your Advantage
Your Advantage means that Agilent offers a wide range of additional expert test and measurement services, which you can purchase according to your unique technical and business needs. Solve problems efficiently and gain a competitive edge by contracting with us for calibration, extra-cost upgrades, out-of-warranty repairs, and on-site education and training, as well as design, system integration, project management, and other professional engineering services. Experienced Agilent engineers and technicians worldwide can help you maximize your productivity, optimize the return on investment of your Agilent instruments and systems, and obtain dependable measurement accuracy for the life of those products.

Related Literature

Publication Title	Publication Type	Publication Number
Agilent 16700 Series Logic Analysis System	Product overview	5968-9661E
Agilent Infiniium 54800 Series Oscilloscopes	Color brochure	5980-2388EN/EUS

By internet, phone, or fax, get assistance with all your test & measurement needs

Online assistance:
www.agilent.com/find/assist

www.agilent.com/find/LAsystems.com

www.futureplus.com

Phone or Fax
United States:
(tel) 1 800 452 4844

Canada:
(tel) 1 877 894 4414
(fax) (905) 282 6495

China:
(tel) 800 810 0189
(fax) 1 0800 650 0121

Europe:
(tel) (31 20) 547 2323
(fax) (31 20) 547 2390

Japan:
(tel) (81) 426 56 7832
(fax) (81) 426 56 7840

Korea:
(tel) (82 2) 2004 5004
(fax) (82 2) 2004 5115

Latin America:
(tel) (305) 269 7500
(fax) (305) 269 7599

Taiwan:
(tel) 080 004 7866
(fax) (886 2) 2545 6723

Other Asia Pacific Countries:
(tel) (65) 375 8100
(fax) (65) 836 0252
Email: tm_asia@agilent.com

Product specifications and descriptions in this document subject to change without notice.

Printed in USA December 20, 2001
5988-4497EN

Agilent Technologies