

THE RAMBUS® SYSTEMS

TEST AND MEASUREMENT GUIDE

VERIFYING, CHARACTERIZING, AND DEBUGGING
YOUR RAMBUS DESIGN



Tektronix®



THE RAMBUS SYSTEMS TEST AND MEASUREMENT GUIDE

FOR VERIFYING, CHARACTERIZING, AND DEBUGGING YOUR RAMBUS DESIGNS

Table of Contents

A Message from Rambus® Inc.	ii
Section 1: Introduction	1
Test and Measurement Implications	2
Section 2: Impedance Measurements	3
Overview	3
Rambus Impedance Tolerances	3
Circuit Board Impedance Measurements	4
Rambus Impedance Measurements with a TDR	4
Differential Measurements Make a Difference	6
Thoughts on Troubleshooting with a TDR	6
Summary: Rambus Impedance Measurements	6
Section 3: Signal Measurements	7
Overview	7
A Closer Look at RSL Signals and Their Environment	7
Rambus Signal Acquisition	8
Some Rambus RSL Measurement Examples	10
Stimulus Sources Reinforce Rambus Design and Troubleshooting Efforts	11
Summary: Rambus Signal Measurements	12
Section 4: Protocol Measurements	13
Overview	13
Logic Analyzer Views Fast, Complex Rambus Protocol Details as They Occur	13
Unfolding Cryptic Packets of Rambus Data	13
The Case of the Disappearing Read DATA Signal	15
Connections: So Many Pins, So Little Time	16
Sometimes Two Instruments are Better Than One	16
Summary: Rambus Protocol Measurements	17
Acknowledgements	18
About the Authors	18
The Tektronix ER Team	18
Rambus Inc	18

A MESSAGE FROM RAMBUS® INC.

System memory bandwidth is more important now than ever before. With the increase in processor performance, multimedia, and 3D functions, high-bandwidth memory is essential to sustain system performance. The transition to Rambus DRAMs (RDRAM™) allows a performance gain of up to ten-fold over the current SDRAM technology. This big step in performance requires a system-level approach to building a robust memory subsystem. Rambus is working with companies such as Tektronix to provide the system characterization and debug tools that enable OEMs to build RDRAM-based systems successfully and quickly.

Rambus' approach to providing increased memory system performance is based on a high-speed, chip-to-chip interface that transfers data over a relatively narrow bus, referred to as the Rambus Channel. The interface, which uses Rambus Signaling Levels (RSL), provides the highest data transfer rate per pin of any memory technology available.

Relative to conventional DRAM and SDRAM technology, Rambus memory offers major advantages for system designers. The first is raw speed. By driving data twice each clock cycle at a burst frequency of 400 MHz, Rambus DRAM (RDRAM) devices support a transfer rate of 800 Mbps on each of the data pins. Coupled with the pipelined, multi-bank architecture of the RDRAM, this results in an actual sustainable throughput that is close to the maximum peak level.

Equal length, matched-impedance data and clock signals allow a simpler, more efficient physical design. Clock skew and variable signal-to-signal capacitive loads are reduced as the address, control, and data bits travel together synchronously along the Rambus Channel with virtually no pin-to-pin timing skew. Switching noise is minimized in the Rambus system as it uses low-voltage (0.8 volt) signal swings and fewer data lines.

The Rambus Channel is specified as a system. All critical, high-speed system engineering issues which arise while developing a memory subsystem (memory controller ASICs, RDRAM, and the Rambus Channel) have been resolved for the designer by Rambus. Designers can easily implement a Rambus system using validated Rambus components and reference board layouts. System bring-up consists of simply verifying the proper operation of the Rambus Channel and components using test equipment available from Tektronix.



SECTION 1: INTRODUCTION

One of the constants in computer technology is the continuing advancement in operational (clock) speed. A few short years ago, a 66 MHz PC was considered “lightning fast.” Today’s common desktop machine operates at many times that frequency. The much-repeated axiom about microprocessors and their periodic doubling in clock rates still holds true.

All this speed is the foundation of a trend toward visual computing, in which the PC display becomes ever more graphical, animated, and three-dimensional. Visual computing enables exciting new games, video applications, and user interface possibilities, ensuring that the market’s appetite for increasingly fast PCs will continue to grow.

In this quest for speed, most of the attention is focused on the microprocessor. But a PC’s memory system is equally important in supporting the new capabilities of visual computing. And commodity dynamic RAMs (DRAMs), the mainstay of PC memory architecture, have fallen behind the microprocessor in their ability to handle data in the volume necessary to support complex, lifelike graphics.

A solution for this deficiency is the innovative Rambus® Channel memory architecture created by Rambus Inc. At its heart, this memory scheme uses ordinary DRAM cells to store information. But the access to those cells, and the physical, electrical, and logical construction of a Rambus memory system is entirely new – and much, much faster than conventional DRAMs. Direct Rambus devices deliver data rates up to 1.6 Gb/s, up to 8 times faster per pin than other currently available technologies.

The new technology presents some challenges for the engineer designing a memory system incorporating the Rambus Channel. There are the usual complications that come with higher frequencies, of course. But more importantly, there are critical issues that cannot be ignored:

1. Impedance environment
2. Physical circuit layout
3. The “analog” quality of digital signals
4. Timing and synchronization
5. Distributed effects
6. Logical protocols
7. Any combination of the above!

Figure 2 summarizes some of the “test points” that must be considered in any Rambus design. Rambus Inc. promulgates design rules and reference implementations for most of these elements. Compliance with these guidelines must be verified from the very beginning of any Rambus system project.




Figure 1. Rambus data rate vs. typical SDRAM data rate.

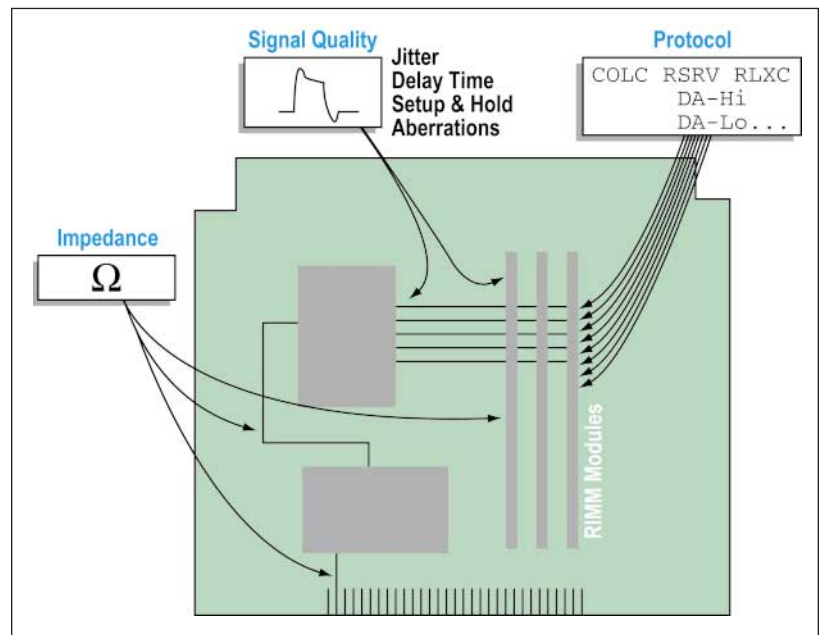


Figure 2. Rambus signal paths and measurement areas.

Test and Measurement Implications

Since Rambus is a system specification containing both physical and logical layers, a test strategy is required in order to verify new designs as they evolve. To meet today's aggressive time-to-market schedules, it's essential to choose measurement methods, procedures, and tools that will provide repeatable performance validation, and do it quickly. Beginning with the initial design, each step becomes the foundation for the next. With impedances verified, signal quality measurements can begin; with those tests complete, debugging at the protocol level can proceed. Finally, system-level tests confirm the design's ability to handle real-world applications.

What's important here is not only the exactness of these tests, but also the time that can be saved by planning and executing a systematic measurement

regime – and by using the right tools. Certain instruments can save literally weeks in the design of a Rambus memory system, simply by revealing circuit performance details that elude lesser tools... or by making measurement results easier to comprehend and analyze.

The essential Rambus measurement toolkit for designers consists of a high-resolution Time Domain Reflectometer (TDR), a multi-GHz, multi-channel oscilloscope, and a logic analyzer that offers high channel count and high acquisition speed, as well as decoding features to disassemble and interpret complex data transmissions.

This document will look at the most pressing Rambus design issues, and will discuss measurement tools, techniques, and strategies to ensure a smooth, fast development process.

SECTION 2: IMPEDANCE MEASUREMENTS

Overview

Because the Rambus Channel operates at data rates up to 800 million transfers/second, it exhibits all the properties of an RF signal. Phenomena like reflections and crosstalk take on unprecedented importance in the Rambus environment. The key to a successful design implementation is step-by-step adherence to the Rambus design rules, starting with the all-important circuit board impedance specification. This section discusses the measurement tools and procedures that underpin any good Rambus system design.

It's no secret that high-speed signals travelling along a transmission line tend to reflect energy backward (toward their source) when they encounter a change in impedance. The amount of reflected energy depends on both the energy of the original transmission and the magnitude of the impedance change.

Due to the innately high speed of Rambus circuits, reflected energy can make the difference between a circuit that works and one that doesn't. Consider this startling fact: at a given instant in time, approximately three bits of information are in transit between a source and a destination (for example, between the memory controller IC and an RDRAM). In conventional dynamic RAM circuits, a single "bit" (that is, the energy representing a pulse) travels down the transmission line alone, reaches its destination, and dissipates any reflections before the next bit is launched. Not so in the Rambus world.

At Rambus speeds, the second and third bits in transit encounter any reflections that occur when that first bit hits the impedance mismatch. A common mismatch point is the node at which the signal enters a connector or IC pin. The reflection can degrade signal margins and cause timing errors, potentially causing failures. Figure 3 depicts a pulse that exhibits the degrading effects of reflections*¹.

Thus it is absolutely essential to minimize the impedance variations across the entire Rambus circuit. Fortunately, the Rambus design guidelines in this area are explicit and easy to interpret. But the specifications are stringent, and proving that your circuit complies with them may be more of a challenge.

*¹ TDR waveform plots in this section are adapted from measurements taken with a Tektronix 11801C TDR and processed with Tektronix Wavestar waveform publishing software. Vertical scale increments are in rho, the standard TDR measurement unit.

Rambus Impedance Tolerances

When designing conventional PC motherboards and memory systems, the accepted impedance environment is 65 Ω , with a $\pm 15\%$ tolerance at any point. With almost 10 Ω tolerance on either side of the design center, there's a fair amount of margin for error in the manufacture of the circuit board. Impedance control on the end item is often ignored – relegated to a process or batch qualification on line width only.

Designing a PC motherboard using high-speed Rambus components is a different story. Rambus devices are designed to operate in a 28 Ω impedance environment. Due to the concerns explained above, Rambus has defined an impedance tolerance of $\pm 10\%$ across the entire bus, including circuit traces, connectors, and components. The acceptable bus impedance range, therefore, is 25.2 Ω to 30.8 Ω .

What's important to note here is that the net measured impedance must be within the $\pm 10\%$ tolerance, including the additive and subtractive effect of all the individual deviations at points along the bus.

Some manufacturers who fabricate individual Rambus sub-components have elected to measure their products (such as motherboards, RIMM™ modules, or Continuity RIMM modules) to an even stricter tolerance to ensure that their products aren't the source of impedance deviations that can drive the bus out of compliance. Many of these vendors also choose to implement 100% test strategies, even though this is not a specific Rambus requirement.




Figure 3. Reflections tend to degrade otherwise clean pulses.

Figure 4 shows a typical Rambus motherboard and RIMMs (much simplified) and the impedance tolerance over the entire bus.

Circuit Board Impedance Measurements

Meeting the Rambus $28\ \Omega \pm 2.8\ \Omega$ impedance specification calls for rigorous policing of design guidelines and fabrication procedures. Dielectric thickness and material properties, trace width, plating thickness... all these variables and more contribute to the board's ultimate impedance environment. The details of the

layout and etching processes are beyond the scope of this document, but several broad procedural tips apply here:

- ▶ 3D Field Solvers are the most accurate means of calculating board and trace geometries to meet the $28\ \Omega$ guideline. Field solver results are generally good predictors of actual board characteristics. Alternatively, reference circuit board designs are available from Rambus. Either approach is a satisfactory starting point for your board design.
- ▶ Cross-sectional measurements of prototype boards are necessary to confirm adherence to the material and physical tolerances.
- ▶ One or more test builds likely will be required; it's desirable to incorporate test "coupons" (sample test areas) to facilitate connection to impedance measurement tools.

Much more design information is available at the Rambus website:

www.rambus.com

and on the Intel® Developer's website:

<http://developer.intel.com/design/chipsets/-memory/rdram/>

Rambus Impedance Measurements with a TDR

The industry-standard tool for circuit board impedance measurements is, of course, the Time Domain Reflectometer (TDR). TDRs are offered with diverse ranges and capabilities, but special care must be taken when choosing a solution for Rambus applications. Following are some general TDR performance guidelines for Rambus impedance measurements:

- ▶ Ability to resolve short trace lengths (<4 mm), connector and IC pad features, etc.
- ▶ Fast system risetime (<35 ps) and high bandwidth
- ▶ Differential measurement capability

A TDR makes a measurement by sending a known pulse down the transmission medium (in this case a circuit board trace or differential pair of traces) and capturing the resulting reflections. The instrument then calculates the Rho (the ratio of reflected to incident signal energy) and uses that information to compute and display the impedance in ohms. The heart of the TDR is an extremely high-bandwidth (20 GHz) sampling oscilloscope and fast-risetime sampling head with an integral step generator.

The TDR measurement circuit as a whole consists of more than just the circuit board trace itself. The instrument also "sees" the effects of its own internal resistances, as well as the effects of the cable and probe connecting the TDR to the unit-under-test. Each of these media has a "signature" that appears in the reflected pulse. Moreover, most TDRs are designed to address a $50\ \Omega$ environment.




Figure 4. Rambus impedance tolerance is $\pm 10\%$ over the length of the bus.



Figure 5. Typical High-Resolution TDR (Tektronix 11801C with SD-24 TDR Sampling Head).

To counteract these error contributions and to account for the 28 Ω nature of the Rambus traces, it's necessary to pre-calibrate the TDR using a known standard impedance. The most accurate standard is an air line having a characteristic impedance close to 28 Ω and at least 15 cm length. The result of a TDR measurement on such a device is shown in Figure 6.

A less accurate, but still acceptable alternative, is to use a pair of characterized 50 Ω airlines (again, at least 15 cm in length) connected in parallel, yielding 25 Ω impedance.

Using the automated capability that advanced TDRs offer, the TDR pre-calibration measurement itself is straightforward. The results can be used as an offset value. If, for example, the test on the 28 Ω standard produces a reading of, say, 26.5 Ω then the 1.5 Ω “discrepancy” is factored into subsequent measurements.

In setting up for actual circuit board impedance measurements, the probing approach is another big concern. Particularly in dense high-speed Rambus circuit areas, the space available to firmly attach a handheld TDR probe is very limited. Yet probe connection integrity is essential. The probe's ground pad should be close to the signal pad for best results, and of course the probe tip connection must be positive. Any compromises in these areas will inevitably appear as inaccuracies in the TDR reading. Figures 7A and 7B contrast the results obtained with “compromised” probing techniques and good probing practices.

One solution is to omit conventional probes altogether and mount output connectors on a “dummy” RIMM module. The RIMM becomes the probe, offering a repeatable, low-loss path directly to the TDR sampling head. The beauty of this approach is that it makes it equally easy to connect to single-ended Rambus Data and Control lines or differential Clock lines, and provides positive ground connections in both instances.

Specially-designed microprobes are also available for TDR measurement applications. This architecture typically provides excellent results, although it's the most costly of the available probing alternatives.

Whatever the means of calibration and connection, the TDR itself must be capable of delivering uncompromised results. The best modern TDRs provide a 35 ps system risetime and 20 GHz bandwidth. These parameters ensure a clean, sharp outgoing (stimulus)

pulse and minimal degradation of the returning reflection. When applied to a well-controlled test methodology, a TDR of this caliber will handle Rambus demands with performance to spare.




Figure 6. TDR measurement on a 28 Ω standard.




Figure 7. (A) Poor probing technique results; (B) Proper probing technique results.

Differential Measurements Make a Difference

As mentioned above, Rambus Clock signals travel through balanced differential lines. Unlike the Data and Control lines, which emerge from the Rambus ASIC Cell (RAC) and go to the RIMM modules, the 400 MHz bus clock signal originates at the far end of the Rambus Channel. Its paired traces loop through the RAC and return to a termination at the far end again. Figure 8 depicts this scheme.

From a procedural standpoint, differential measurements are just like single-ended – if the TDR permits it. The TDR must be able to execute differential probing and measurements, and calculate the result.

Attempting to measure the two halves of the differential pair separately can produce misleading results. Two traces in close proximity, as the differential traces are, tend to read a lower impedance than their characteristic impedance as a pair. If this information is used to guide the board fabrication process, the result may be a Clock signal line with reduced voltage and timing margins – clearly unacceptable.

Therefore the best TDR solution is one that's innately differential. Differential probing produces valid, meaningful measurements that correlate closely with the Rambus device's real-world operational context. Figure 9 is an example of a differential TDR measurement.

Thoughts on Troubleshooting with a TDR

The TDR is the most effective tool for making impedance measurements. It's also a powerful troubleshooting aid during both design and later production steps.

A subtle benefit of the TDR is that it's a passive test environment, since the TDR provides the stimulus. The TDR can test a board before significant assembly costs accrue. The board is "blank," not loaded with components and not powered up. Moreover, the

board doesn't need to "work" in the functional sense. Yet TDR measurements help stabilize and control board fabrication variables and even some assembly parameters.

The TDR's performance level and capabilities pay off during the troubleshooting phase of the design process. An instrument with high bandwidth and a system risetime on the order of 35 ps makes it possible to distinguish board characteristics as little as 3.5 mm apart, and to view individual aberrations resulting from trace imperfections, connectors, pads, and other features.

A proven TDR troubleshooting technique helps to pinpoint areas of process deviation along a trace while probing at the board's edge connector (or any point at the end of a trace). An aberration in the circuit board trace shows up as a "bump" or dip in the TDR display, even if the board is covered with solder mask or some other insulator. The bump's location on the TDR display correlates with the actual physical location of the problem on the board. If you put your finger (or a smaller object if necessary; some users employ a disconnected scope probe) on the trace, it too will cause a bump on the TDR display. As you move your finger closer to the problem area, the two bumps will get closer together. Using this technique, it's possible to "home in" on the problem very quickly.

Summary: Rambus Impedance Measurements

Impedance measurements are a necessary initial step in any Rambus design project. The Rambus Channel is designed to operate at 28 Ω impedance; anything other than that can cause signal reflections, in turn reducing voltage and timing margins.

The Time Domain Reflectometer is the cornerstone of practical, accurate impedance measurement technology. When properly calibrated for the 28 Ω environment, a fast, capable TDR can minimize the time and cost of designing and troubleshooting new Rambus circuit designs.

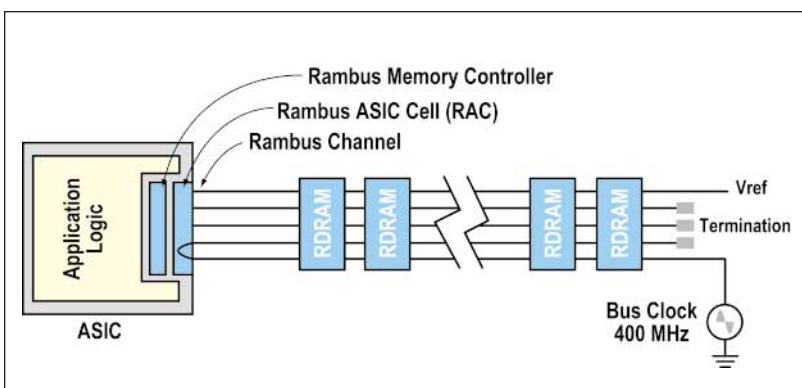


Figure 8. Rambus signal paths.




Figure 9. A 28 Ω differential Rambus impedance measurement.

SECTION 3: SIGNAL MEASUREMENTS

Overview

It should come as no surprise that clean, well-controlled digital signals are at the heart of every successful Rambus system design. The Rambus Channel is based on a very high-speed, chip-to-chip interface that transfers data on each edge – rising or falling – of a 400 MHz differential clock. The data, clock, and control lines have 800 mV logic levels that must meet stringent timing requirements. These are some of the “analog” characteristics that make up the Rambus Signaling Level (RSL) environment. In addition, the unique Rambus logical protocol can make it difficult to capture meaningful pulses and data combinations with conventional measurement tools and triggering approaches.

A Closer Look at RSL Signals and Their Environment

A typical Rambus Channel is made up of 30 controlled-impedance, matched transmission lines, including:

- ▶ Two 9-bit wide data buses (DQA and DQB)
- ▶ A 3-bit wide ROW bus
- ▶ A 5-bit wide COLumn bus
- ▶ ClockToMaster and ClockFromMaster differential clock buses

In designing the Rambus Channel architecture, Rambus engineers went to great lengths to maintain excellent signal quality and transmission characteristics. The architecture makes use of carefully controlled terminations, low-voltage signaling, active current control, high packaging density, and other innovative techniques throughout. These requirements are spelled out in detailed documentation from Rambus, Inc.

www.rambus.com

and on the Intel Developer’s website:

<http://developer.intel.com/ial/home/sp/-index.htm>

Looking again at the simplified schematic in Figure 8, the physical architecture consists of a memory controller at one end, termination resistors at the other, and RDRAMs in between. The Rambus ASIC Cell (RAC) interfaces the low-voltage Rambus signal environment with the standard CMOS logic levels common in most logic device types.

The standard Rambus Channel high-speed signal has an 800 mV excursion centered about a reference voltage of +1.4 V, as shown in Figure 10.

Of course, this “textbook” representation is very different from the real-world RSL signals on an actual Rambus Channel. Even with all the layout and design precautions outlined in the Rambus specifications, active signals are bound to show some effects of crosstalk, jitter, reflections, overshoot, and any of a dozen other artifacts. In Figure 11 a pulse shape traced from a valid Rambus signal acquisition is superimposed on the same ideal pulse template.

It is these signal imperfections that are the starting point of many design and troubleshooting measurements. The tools you select for capturing and analyzing RSL signals can make a difference in your ability to see waveform aberrations, and ultimately in your ability to complete a reliable Rambus-based design.




Figure 10. Ideal Rambus signal characteristics.




Figure 11. Observed Rambus signal characteristics.

Rambus Signal Acquisition

The cornerstone of basic signal acquisition is of course the oscilloscope.

Several factors come into play when choosing an oscilloscope for the demanding task of Rambus measurements.

The oscilloscope must have the fundamental performance to handle fast Rambus signals.

Clearly, the scope needs the bandwidth to handle 400 MHz clock signals. While many economical oscilloscope models boast 500 MHz bandwidth, remember that this specification actually refers to the 3 dB rolloff point. Viewing a 400 MHz RSL signal with an instrument in this class (while trying to see edge transition details with any degree of clarity) may lead to significant errors in displayed signal amplitude. Instruments specified with 2 GHz or 3 GHz bandwidth are recommended for Rambus applications.

Oversampling is a necessity if the digitizing oscilloscope is to capture fast edges, transients, and one-time events faithfully. Some instruments deliver ample bandwidth but may lack the sample rate, especially if more than one input is in use, to capture signal details or infrequent glitches. The best solutions are those with 3X to 5X oversampling ratios on all channels simultaneously.

Equally important, the instrument's performance must be available at the probe tip, not just at the input to the scope mainframe.

Figure 13 depicts a Rambus clock waveform from a circuit experiencing intermittent failures. Note that the pulse is split almost in half from an aberration that dips down almost to the V_{ref} level. To the circuit, this might look like two pulses rather than one, which would explain the intermittency. An oscilloscope's bandwidth and sample rate can make the difference between capturing this detail on the first try, and spending the next five weeks looking for it.

The oscilloscope should be equipped to interface to external triggers from logic analysis instruments.

This is a valuable asset in making measurements on complex digital devices and circuits.

The scope must have an exceptionally accurate time base system and a low-jitter triggering system.

State-of-the-art instruments have a time base accuracy of 10 ppm and a trigger jitter specification of 8 ps. Several accuracy-related specifications are of interest here:

The **Trigger Jitter** specification defines the amount of jitter the scope's triggering system adds to the existing jitter caused by the device-under-test. It's applicable only to repetitive acquisition jitter measurements, not the single-shot jitter measurements explained later in this document. It is, however, important when making jitter measurements using traditional histogram/statistics techniques.

The **Delta Time Accuracy** specification is critical when making single-shot timing measurements because it determines how closely the measurements will match the real values. Delta Time takes into account both the repeatability and the resolution specifications discussed below. It defines the worst-case peak deviation from the expected value. In an oscilloscope, Delta Time is based on a number of factors, including sample interval, timebase accuracy, quantization error, interpolation error, vertical amplifier noise, and sample clock jitter. Each of these factors contributes to the timing error implicit in the Delta Time measurement accuracy specification.

Today's best-in-class oscilloscopes feature a Delta Time accuracy specification of 15 ps. This is a worst-case specification determined under many input conditions, and, depending upon the input applied, these instruments can actually measure much smaller errors. Figure 18 shows a clock waveform measured over 10,000 cycles; it was found to have <1.5 ps RMS jitter and <11 ps peak error.




Figure 12. 3 GHz Digitizing Oscilloscope (Tektronix TDS 694C DSO).




Figure 13. Rambus clock waveform, showing amplitude aberrations (shown using a Tektronix TDS 694C DSO).

The **Repeatability** specification determines how closely repeated measurements match one another. Repeatability is valuable but cannot stand alone. If, for example, the instrument always returns an 11 ns measurement for a clock pulse whose duration is actually 10 ns, repeatability is excellent – but accuracy is not! It's very important, therefore, that good repeatability be combined with excellent Delta Time accuracy.

The **Resolution** specification defines the smallest measurement increment the instrument can “perceive.” If, for example, the measurement system returns 10.001 ns for a 10 ns period, its resolution is 1 ps. Like repeatability, this specification is contained within the Delta Time accuracy specification.

The oscilloscope's probing solutions must extend the instrument's performance all the way to the device-under-test.

Probing is a challenge in the Rambus environment. RSL signals operate at low signal levels with DC offsets. Unlike other common logic families, RSL Logic “0” isn't at ground potential; instead it's the more positive of two DC values.

The oscilloscope probe must present a very low degree of loading to the signal, and it must be able to eliminate the DC component while working at maximum frequency. Conventional passive probes simply won't do the job. The preferred solution for day-in, day-out probing of Rambus signals is an active (FET buffered) probe that delivers uncompromised performance at the test point.

Rambus clock signals are fast, low-voltage, true differential signals. The only acceptable solution for probing in this environment is a high-speed differential probe with high CMRR. While there are measurement “shortcuts” involving the use of two scope channels and standard probes for measuring differential signals, this approach simply isn't adequate at Rambus speeds. A differential probe with 60 dB CMRR and appropriate bandwidth is a satisfactory solution for Rambus work.

The probe's physical size is also a concern. Rambus circuits are made up of densely-packed SMD devices with very closely-spaced board traces. For manual probing, the probe tip must be able to fit into the tiny

spaces between components and between traces. Probe-to-PCB adapters are also valuable in making reliable connections to devices mounted on prototype boards and buses.

For productivity reasons, the oscilloscope should simplify everyday Rambus measurements.

In the course of almost any Rambus design project, it will be necessary to measure signal rise and fall times, setup and hold times, jitter, and more. Clearly, it is desirable to automate these routine tests. Many digitizing oscilloscopes have pre-programmed procedures for these tests and others.

The oscilloscope should offer a variety of triggering options and conditions.

In many cases, a specialized trigger (such as a State, Runt, or Pulse Width trigger) is the only way to capture a troublesome event that is occurring erratically. For example, a narrow, infrequent transient might be difficult to distinguish using normal edge triggering, but with the Pulse Width trigger setting, the oscilloscope can be set to trigger when a shorter-than-normal pulse occurs.



Figure 14. A compact differential probe (Tektronix P6248 shown) meets high-speed Rambus requirements.

Some Rambus RSL Measurement Examples

The explanation above discusses some Rambus RSL signal measurements requirements in general terms. The next step is to examine some specific measurements and see how instrument performance can affect their outcomes.

Setup and hold time specifications of Rambus are very narrow. Figure 15 is adapted from the Rambus RDRAM Data sheet (Rambus Document DL0059), showing the edge relationship between the CFM Clock and the valid data pulse associated with it. At the minimum specified cycle time (2.5 ns), both t_s (setup time) and t_H (hold time) are only 200 ps each! Although this is an extreme example, Rambus setup and hold times ranging from low nanoseconds to a

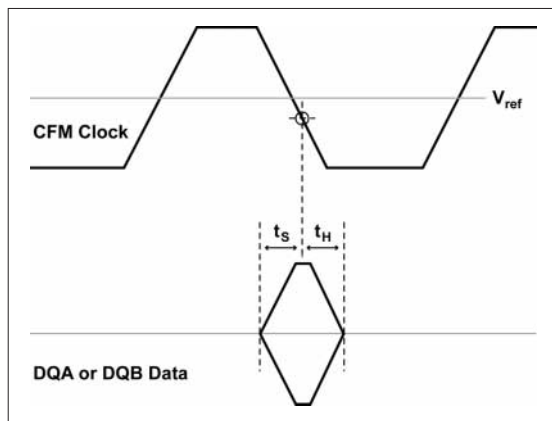


Figure 15. Rambus RSL Clock-to-Data Setup Time definition.




Figure 16. Cursor-based Setup Time measurement (shown using a Tektronix TDS 694C DSO).

few hundred picoseconds are not uncommon. Due to the pipeline clocking nature of Rambus signals, valid data can be captured only through single-shot acquisitions – a compelling reason to use the fastest oscilloscope available.

There are two ways to approach the Setup Time measurement. The first, and most common, is the cursor measurement method. The measurement requires two oscilloscope inputs, of course. To work at Rambus speeds, it's important to choose an oscilloscope that doesn't trade off sample rate as additional inputs are connected. And to measure setup time values in the low nanosecond range and below, it's also prudent to deskew the probes that will be connected to the clock and data lines. Even a minor delay difference between these probes can invalidate the measurement result. A deskew fixture is the best way to equalize the delay differences.

With the deskewing completed, the measurement process starts with a stable, positive probe connection. It may be necessary to use a soldered-on probe adapter to capture the high-frequency, low-voltage RSL signals without adding unwanted aberrations.

For capturing and viewing fast Rambus clock and data signals, the scope's highest real-time sample rate is normally used. In Figure 16, the sample rate is (automatically) set to 10 GS/s and the data signal has triggered an acquisition.

NOTE: It may be tempting to use the "equivalent time" (ET) acquisition techniques if the oscilloscope lacks adequate real-time sample rate.

However, this method is sure to add error to the time measurement due to the scope's own trigger jitter. Moreover, ET sampling is not usable for non-repetitive signals.

The final step is to position the cursors at the points of interest, usually the 50% point of the relevant edge transitions. The oscilloscope should automatically display the time reading. In this illustrated example, the setup time is 1.42 ns. That completes the setup time measurement using the "conventional" procedure.

A second setup time measurement alternative has appeared recently. This software-based method, like the procedure above, uses the information from a single acquisition. However, it has the ability to measure the timing on every valid edge on a record made up of thousands of cycles' worth of clock and data transitions. This has the advantage of providing a meaningful statistical analysis of the measurement's behavior over time, and also allows you to correlate specific clock edges to their associated data transitions. Figure 17 shows the display resulting

from a setup time measurement using the Tektronix TDS 694C oscilloscope. Note that the analysis occurred over 119 timing intervals of interest.

Excessive clock jitter can reduce the already narrow timing margins in a Rambus circuit. Essentially, jitter is the unintended variation of an edge's timing placement over many cycles of operation. Measuring this quantity in the Rambus environment challenges an oscilloscope's repeatability, resolution, sample interval accuracy, and its own trigger system jitter.

As with the setup and hold time measurements, there are several ways to proceed with jitter testing. Jitter is not an absolute figure; simply looking at it on the oscilloscope screen is not a useful exercise. Because jitter is inherently statistical in nature, all of the measurement methods deliver numerical statistics as the result – peak, mean, standard deviation, and so on.

The first, simplest method is the least likely to produce comprehensive results on a high-speed Rambus circuit. But it may be useful in determining whether circuit jitter exceeds gross tolerances. Many oscilloscope models can perform automatic Period measurements. When such routine is carried out, the Mean reading and Standard Deviation provide statistical information about the jitter in the circuit. Similarly, performing a Period measurement and applying the Min-Max function yields a reading of the jitter peak values.

Many advanced oscilloscopes provide built-in histogram features. The histogram technique works with both equivalent-time and real-time acquisitions, although with some tradeoffs in accuracy when equivalent time is used. The histogram is activated by a front-panel button or menu, and automatically produces a readout of mean, peak-to-peak, and standard deviation results from the waveform currently displayed.

The third and most effective of the jitter measurement techniques is, like the setup tests above, a software-based analysis of a single acquisition containing many waveform cycles. Figure 18, produced by the Tektronix TDS 694C oscilloscope, describes a 300 MHz Rambus CFM clock. The analysis encompasses almost 11,000 clock cycles. Here, the peak-to-peak period jitter is less than 71 ps, easily within the 100 ps peak-to-peak requirements of the Rambus Channel.

An additional reminder about Rambus jitter measurements: the solution, whether it's an automated front-panel scope function or a more sophisticated software routine, must be capable of making many types of jitter acquisitions: Cycle-to-cycle, input duty cycle, frequency of modulation, jitter over 1 to 6 cycles, and more.

Aside from jitter, there are other circuit problems that are statistical in nature and therefore difficult to find. To statistically characterize signal behavior, many data points are required. The Digital Phosphor Oscilloscope (DPO), with its real-time intensity-graded display, is especially powerful in these respects. Some DPOs can acquire up to 200,000 waveforms per second, greatly increasing statistical confidence and making it much easier to spot intermittent errors.

Stimulus Sources Reinforce Rambus Design and Troubleshooting Efforts

Up to this point, most of our discussion has focused on Rambus measurements. But as a new system design evolves, problems due to interactions among the growing number of components and buses can arise. Tolerances can build up, ground bounce phenomena can emerge, and signals can deteriorate. It may be difficult to pinpoint the blame on any one part of the circuit.

We have talked about measuring setup and hold times and jitter, for example. But what is the recourse when there is too much jitter? Jitter is a symptom – what's causing the symptom? Other potential problems range from improper clock and data waveform symmetry to excessive noise on the transmission lines.

External data and clock generators – signal sources – can augment measurement tools such as oscilloscopes and logic analyzers in resolving design problems. They can help find the root cause of problems that seem to have no clear origin. The process is one of signal substitution and fault isolation.




Figure 17. Automated Setup Time measurement results (shown using a Tektronix TDS 694C DSO).




Figure 18. Rambus clock jitter measurement using software-based analysis (shown using a Tektronix TDS 694C DSO).

As always in the Rambus environment, the data generator must deliver exceptional performance at high frequencies. A data rate of at least 1 Gb/s is needed to match the Rambus Channel's data rates. In addition, it must provide an ultra-low-jitter signal (preferably <30 ps) with edge risetimes of 150 ps or less. To carry out certain stress tests, the data generator must deliver a wide range of signal amplitudes – approximately 0.25 V_{p-p} to 2.5 V_{p-p} . Programmable channel-to-channel delay and complementary outputs are also useful for the differential signal lines common in Rambus architecture. Following are some applications that make good use of a high-performance signal source.

Design margin testing

Even a design that “works” has its limits. What are those limits? It's mandatory to confirm that the design offers sufficient margin to manage the vagaries of the manufacturing process. Tolerance buildup in clock signal path terminations and device AC parameters is another common problem in high-speed circuits.

With an appropriate signal source, system elements can be individually stressed with signals that are too early, too late, too large, too noisy, and so on.

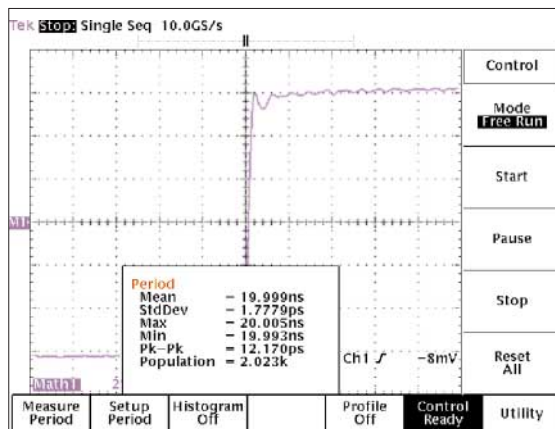


Figure 19. A low-jitter signal from a Tektronix DG2040 Data Generator can be used for Rambus clock substitution during troubleshooting (shown using a Tektronix TDS 694C DSO).

Clock substitution

Memory access failures can have many causes. One such cause is setup and hold timing violations due to excessive clock jitter. By substituting an external low-jitter source for the clock that feeds the RDRAMs, it is possible to “divide and conquer” the circuit problem.

Figure 19 is a jitter characterization of a Tektronix DG2040 Data Generator output, measuring only 12.170 ps jitter. Comparing this reading to that of Figure 18 (70.444 ps), it's obvious that this data generator is capable of handling Rambus clock substitution chores.

An external clock with programmable complementary outputs can provide paired signals to evaluate the effects of amplitude variations, skew between the ClockToMaster (CTM) and ClockFromMaster (CFM), and symmetry referenced to V_{ref} .

Dispel the effects of ground bounce and transients

Simultaneous bus access across a multi-channel Rambus system draws far more current than usual, causing ground bounce and possibly voltage transients that are echoed in the CTM signal line.

An external clock source can usually provide much greater current source and sink capacity than the on-board clock generator. This extra output capacity offers greater resistance to waveform deformation permitting the designer to evaluate the effects of a more perfect clock in the circuit.

Summary: Rambus Signal Measurements

As a Rambus system design evolves, engineering efforts must focus on maintaining the signal quality and timing relationships set forth in the Rambus-published specifications. Given Rambus data rates, signal levels, and transmission line environments, there's an absolute requirement for measurement tools with matching bandwidth, timing accuracy, and probing capabilities. Nothing less will suffice.

Fortunately, there's a new generation of measurement solutions that delivers unprecedented detail for signal acquisition, analysis, and stress testing applications.

SECTION 4: PROTOCOL MEASUREMENTS

Overview

As the functional debug phase of a Rambus circuit design project commences, your preparatory efforts in making thorough impedance, signal quality, and timing measurements begin to pay off. These earlier steps provide the best possible environment for high-speed Rambus devices to send and receive information. But these transactions are themselves complex. The Rambus protocol presents new challenges to the designer who is accustomed to working with conventional DRAM circuitry. To stay on track with demanding time-to-market schedules, it's necessary to use a suitable logic analyzer with a Rambus probe adapter to examine Rambus activity at the protocol level.

A logic analyzer's protocol trace brings out the internal details of memory system functionality in human-readable alphanumeric formats. On the display, Rambus operations such as REFRESH or WRITE appear in their actual mnemonic form, followed by the hexadecimal content of the data transfer. As soon as a new Rambus circuit design is capable of sending, receiving, and responding to signals, the logic analyzer becomes the cornerstone of protocol analysis for:

- ▶ Data throughput
- ▶ Refresh operations
- ▶ Bank management
- ▶ Power management
- ▶ Bus protocol latency
- ▶ Memory controller latency
- ▶ Rambus Signal Level (RSL) calibration
- ... and more.

In addition, the logic analyzer is valuable for general troubleshooting and resolving system integration issues. Its state display provides a useful waveform-based overview of signal behavior in a binary state context. This view is the stepping stone for even more detailed signal capture procedures to be discussed later in this document.

Logic Analyzer Views Fast, Complex Rambus Protocol Details as They Occur

With its multiplicity of probes and inputs, its ability to trigger on complex logic conditions, and its capacity to store volumes of acquired data, the logic analyzer captures detailed "snapshots" of bus operations. Of course, Rambus data rates and signal characteristics place exacting demands on the logic analyzer.

To handle Rambus measurement requirements, the logic analyzer must be fast – but how fast? At first glance, the 200 MHz maximum synchronous clock limit of today's best logic analyzers might seem to fall short of the Rambus Channel's 800 MHz data transfer demands. But the distinction between data transfer rate (frequency) and data bandwidth (in GB/s) is the real issue here. As long as the logic

analyzer has sufficient acquisition data bandwidth to acquire all the data and protocol of a particular bus, the clock speed challenge can easily be handled by a bus-specific probe adapter.

As a case in point, the logic analyzer mentioned above with its 200 MHz synchronous acquisition capability actually has an acquisition data bandwidth of 3.4 Gb/s – easily capable of handling the highest Rambus data rates of 1.6 Gb/s. In fact, that same instrument has capacity enough to also acquire the protocol signals (another 800 MB/s) and parity (another 200 MB/s) signals at the same time as the data. With modular expandability, some logic analyzers can be configured to achieve even higher data bandwidths.

In reality, because all data and protocol information is transferred on the Rambus Channel in bursts of eight cycles, the data acquired by the logic analyzer is much more useful when the probe adapter performs a physical-to-logical conversion in real-time. This delivers the data to the logic analyzer in a form that is eight times as wide at 100 MHz. To summarize the concept, fast synchronous buses need the logic analyzer's data bandwidth, not just its synchronous clock rate in MHz.

Unfolding Cryptic Packets of Rambus Data

Acquisition data bandwidth is important, but it's by no means the only attribute needed for proper analysis of Rambus operation.

Protocol signals on the Rambus Channel have their content structured for maximum data throughput, not for direct human comprehension. The data stream is made up of packets which occur in eight transfer bursts. ROW packets consist of 24 bits (8 bursts of 3 bits); COLUMN packets are made up of 40 bits of information (8 bursts of 5 bits). DATA packets consist of 144 bits (8 bursts of 18 bits, including parity) that follow COLUMN packets for certain operations. ROW, COLUMN, and DATA



Figure 20. The Tektronix TLA 700 Series Logic Analyzer's 3.4 Gb/s acquisition data bandwidth surpasses Rambus data rate requirements of 1.6 Gb/s.

packets are largely independent; individual bus operations employ different combinations of these packet types. Even packets of the same type are not all aligned on similar clock boundaries; instead, alignment timing is adjusted dynamically to maximize performance. Subsequent packets can begin an arbitrary number of clock cycles after a prior packet completes. Figure 21 shows the Rambus pipeline architecture.

The ROW and COLUMN packets control a variety of operations: device, row, and column selection, row activation, read/write control, refresh, power management, and more. Due to the “bursted” nature of these signals, simply probing the bus with a logic analyzer and viewing the result would be almost meaningless.

To be useful, the logic analyzer and Rambus probe adapter must reconstruct logical packets from the bursted signals and display them in logical groupings with understandable mnemonics. Figure 22 illustrates this concept.

One way to do this reconstruction is to post-process the data with disassembly software after the acquisition. Unfortunately, this approach does nothing to help the designer trigger on the protocol itself, which is the only way to ensure that data is captured in the context of the logical bus activity occurring at the time.

A more efficient way to do the reconstruction is to implement it in real-time within the probe adapter

hardware. With the packets de-serialized and dynamically aligned in real-time, the user can define triggering scenarios using a syntax that represents the logical operation itself. This is much simpler than attempting to mentally interpret the cryptic burst format of the physical bus.

The upper diagram depicts the native Rambus data layout of a ROWR packet (ROWR is one specific class of ROW packets). Clearly, the de-serialized packet beneath it is much easier to read and comprehend. The de-serialized elements are spread out in a common-sense numeric format with the LSB on the right.

De-serializing and dynamically aligning the packets in real-time is essential to meaningful triggering of the logic analyzer. Without these two operations, the instrument would be unable to trigger effectively because it has no way to evaluate the data every 1250 ps, to say nothing of evaluating every burst transfer and every possible packet alignment combination.

For example, when troubleshooting it’s often useful to capture the data that’s being written into a particular RDRAM location. To do so, it’s necessary to identify – that is, trigger on – a specific combination of Device, Bank, and Operation fields, and the Data itself. The user may need to trigger on a combination of ROW and COLUMN packets that access the same specific DRAM device. By reconstructing packets in real-time, the logic analyzer can trigger on a unique ROW followed by a specific COLUMN. The DATA that appears when the trigger fires is that which was stored at the event of interest.

The de-serialized logical packets are very useful for storage qualification as well. The user may want to only store packets that access a specific Device, Row, or Column, or to store only packets that initiate specific operations. The easily-understood de-serialized packets hasten analysis of specific classes of operations.




Figure 21. Rambus pipeline architecture.

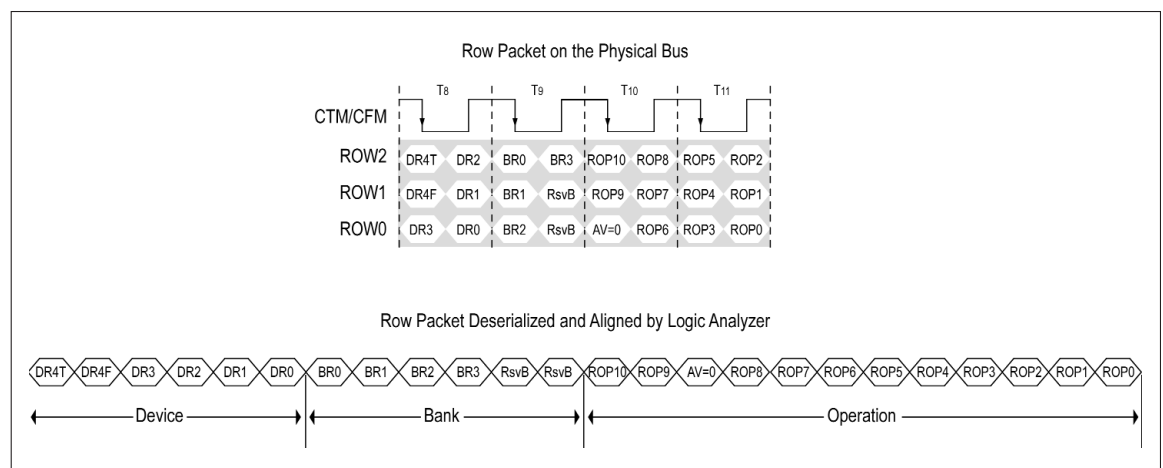


Figure 22. Physical vs. de-serialized Rambus ROWR packet.

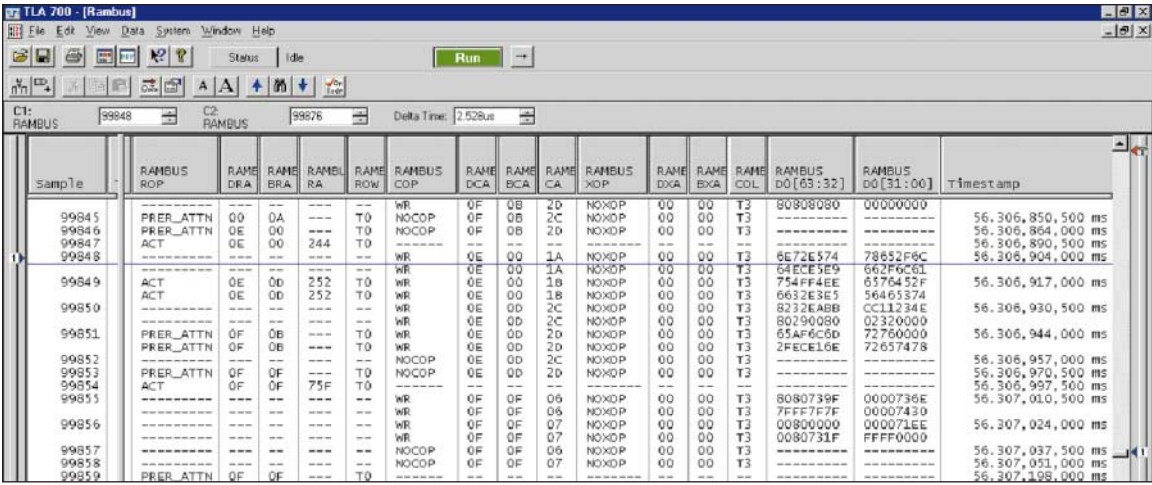


Figure 23. De-serialized Rambus transactions (using a Tektronix TLA 700 Logic Analyzer with a TMS810 Rambus adapter).

The information acquired by the logic analyzer during any given transaction is a complete logical packet. De-serialization simplifies the data layout while still capturing every single detail of logical activity. Figure 23 is a logic analyzer screen display of a Rambus acquisition, showing the result of the de-serialization process.

The Case of the Disappearing Read DATA Signal

The RSL (Rambus Signaling Levels) logic used on all high-speed Rambus signals uses a low-voltage-swing over a carefully controlled low-impedance transmission line. The signal path design is carefully tuned for optimum signal quality and power dissipation through the use of single-ended terminations. One of the by-products of this bus topology is the unique behavior of the DATA signals during a data Read operation (when READ data is driven on the bus by an RDRAM). In normal operation, the data simply cannot be viewed with conventional acquisition techniques.

Look at Figure 24 to understand why. The DATA and control (ROW and COLUMN) signals are parallel

terminated (to V_{term}) at only one end of the transmission line (opposite the memory controller/chipset). The other end of the line is effectively only terminated by the series output impedance of the Rambus ASIC Cell (RAC) in the chipset, a form of termination that disappears when the bus is reversed for READ data. Therefore, when an RDRAM delivers a DATA signal during the Read operation, the signal on the bus encounters a perfect termination at one end (to the right in Figure 24) and a very high impedance on the other (to the left in Figure 24).

It's important to note here that the RDRAM “drives both ways” down the bus, which in effect means it drives two parallel $28\ \Omega$ impedances. The result is that it encounters a $14\ \Omega$ environment – not the $28\ \Omega$ impedance characteristic of other Rambus lines. The DATA signal traverses the bus at half amplitude. But because this is the expected behavior for the DATA line under these circumstances, the circuit has been designed to accommodate it.

When the signal encounters the high impedance at the RAC, reflections occur. At the RAC connection point itself, the reflections are virtually in phase with the incident waves, and actually double the

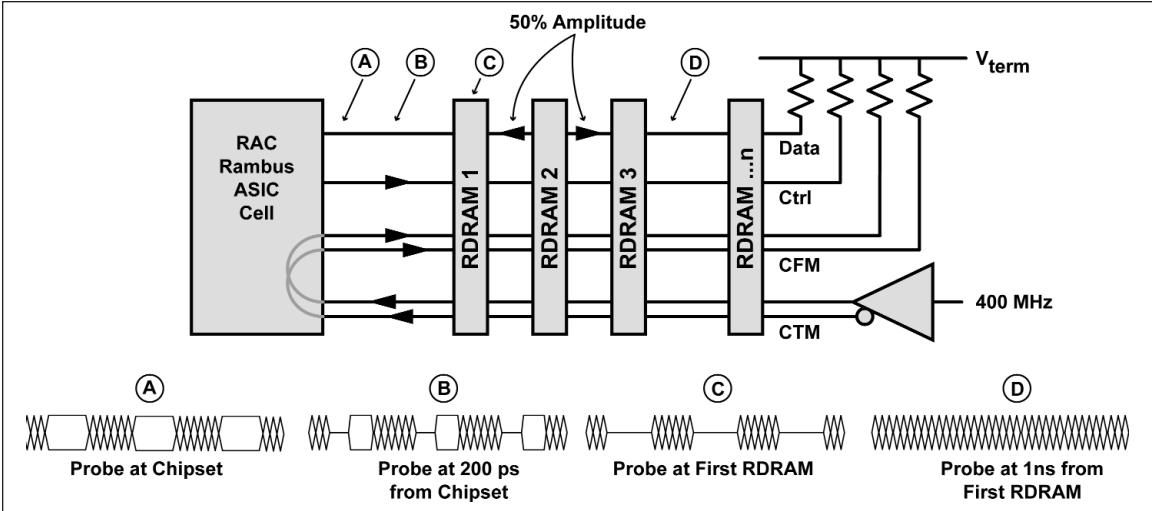


Figure 24. Rambus Read DATA signal behavior.

signal's amplitude at that point. The RAC, therefore, receives full-amplitude DATA pulses.

However, as you probe down the length of the DATA line (away from the RAC), the relative phase difference of the two signals increases. The reflections begin to cancel out more and more of the DATA signal. The DATA pulses steadily erode and ultimately become unintelligible. Given that the guaranteed data valid window is only about 400 ps at the input of the RAC, it erodes to nothing within a very short distance from the end of the bus (a fraction of an inch).

Because DATA is never shared directly from RDRAM to RDRAM, this characteristic is not a problem for normal Rambus operations. Read DATA, even when it's destined to go only from one RDRAM to its nearest neighbor, always goes to the RAC first, at which point the signal is still in usable form. Having passed through the RAC, the information goes out through the RAC driver on a separate Write DATA operation. In effect, this buffers the whole transaction and sends the Write DATA down a properly terminated path.

Note that the control signals, which are terminated the same way, never exhibit the reflection problem because they are always driven by the RAC. Like the Write data signal, these other RAC-driven signals appear about the same regardless of where the probe is placed. The clocks, which are differential and terminated differently, do not exhibit the reflection problem.




Figure 25. Rambus Probe Adapter (Tektronix TMS810).

Currently there is no solution available for directly viewing Rambus Read DATA signals using normal, voltage-threshold acquisition techniques.

Connections: So Many Pins, So Little Time

With all the emphasis on Rambus “speeds and feeds,” it's easy to overlook the mechanical aspects of connecting multi-channel probes to dense logic devices and boards. In reality, connectivity is one of the biggest problems in protocol analysis. With Rambus circuits, the challenge is compounded by critical termination, impedance, and signal quality requirements. Failure to heed these requirements can produce erratic or invalid results.

Some logic analyzer architectures require the user to individually connect scores of tiny probe tips to specially-prepared test pins on a large external pre-processor. Aside from the time-consuming cumbersome nature of this method, the likelihood of making a misconnection is high. With, say, only two probes reversed, the instrument might still trigger, but on entirely different information than intended. The misconnection could easily go unnoticed, because the triggering would appear to be operating normally. Yet the data from the acquisition would be worse than invalid – it would be misleading.

A much cleaner Rambus probing solution is the RIMM Probe Adapter architecture such as the Tektronix TMS810 Rambus Probe Adapter shown in Figure 25. A special interface board substitutes for a standard RIMM module on the Rambus Channel and plugs directly into the RIMM socket, making the same reliable bus connection that an actual RIMM would provide. Bus terminations, impedances, and transmission paths are undisturbed, and the probe adapter “acts” like an operating RIMM. The result is data monitoring and acquisition under conditions that closely approximate actual Rambus operation.

On the logic analyzer side of the probe adapter, the instrument connects via latching high-density connects and ribbon cables, protecting against both misconnections and accidental disconnections.

Importantly, the probe adapter offers local buffering and also de-serializes the Rambus protocol packets as described earlier. It samples Rambus protocol signals at speed and effectively divides the frequency (not the data bandwidth!) by eight. The resulting signal sent to the logic analyzer is one-eighth as fast but eight times as wide as the sampled signal.

Sometimes Two Instruments are Better Than One

Often, it's necessary to view minute signal details that are outside the realm of the logic analyzer alone. The logic analyzer may detect the symptom – a read or write error, for example, but may be unable to home in on the timing violation that's causing it. The fastest edge transitions and transients are better captured by a high-speed oscilloscope. The 3 GHz real-time acquisition performance of a state-of-the-art oscilloscope is easily capable of identifying the violation. But the oscilloscope's triggering features, usually limited to four inputs and simple binary combinations thereof, may not be sufficient to respond to complex data combinations in the Rambus Channel.

The solution is to use a logic analyzer together with a Rambus Probe Adapter to detect the occurrence of a specific complex logical event based on protocol and Write data. Today's top logic analyzers can acquire up to 680 logic signal inputs simultaneously – certainly ample for any Rambus Channel measurement! Using a specially designed internal trigger state machine, the logic analyzer evaluates multiple simultaneous Boolean or “If-Then-Else” conditions or even specific events such as a memory read operation with up to 16 states of possible interactions. The trigger goes to a suitably-equipped oscilloscope, which in turn captures the signal on the bus at the time the trigger occurred. This process is known as Time-Correlated Cross Triggering. Some high-speed oscilloscopes, such as the Tektronix TDS 694C, have a special logic analyzer cross trigger mode that simplifies the correlation of protocol and data

acquired by the logic analyzer with the physical signals acquired simultaneously by the external oscilloscope.

For certain timing or characterization measurements, statistical analysis is needed. This too is the province of the oscilloscope. Even when triggered by the logic analyzer, the oscilloscope can produce histograms on pulse edges and widths relating to specific logical events on the Rambus Channel.

The pairing of a high-speed logic analyzer with an equally capable oscilloscope is the best toolset for identifying glitches and “analog” problems that can cause intermittent failures on the Rambus Channel. The combination also supports stress, reliability, and QA testing, in which signal quality is verified over a range of functional conditions.

Summary: Rambus Protocol Measurements

When performing functional debug and validation procedures on Rambus circuits, the engineer must confront challenges of complexity, clock rates, and connectivity. Protocol analysis tools (logic analyzers) designed to penetrate complex, packetized Rambus bus operations are now available. These instruments offer both the bandwidth and the signal de-serialization needed to capture Rambus data and display it meaningfully. When paired with a high-speed oscilloscope, these tools can reveal the smallest details of signal activity. Advanced probing and connectivity solutions complete the package. A modern logic analyzer is a cornerstone of Rambus design measurements.

ACKNOWLEDGEMENTS

About the Authors

As a leading technology manufacturer for verifying, characterizing, and debugging of Rambus implementations, Tektronix is proud to offer this informative test and measurement guide. Developed by Tektronix' own Engineering Rambus (ER) Team of experts, this guide was produced to help high-speed digital design engineers swiftly integrate their Rambus system designs.

The Tektronix ER Team

The Tektronix ER Team is made up of five technical specialists with core competencies in electrical signal quality and protocol testing. Each primary ER Team contributor to this guide has the hands-on expertise as well as the broad test and measurement experience necessary to deliver this authoritative reference tool on Rambus system verification, characterization, and debug.

Rambus Inc

Rambus, RDRAM, and the Rambus logo are registered trademarks of Rambus Inc. Direct Rambus, Direct RDRAM, RIMM, and SO-RIMM are trademarks of Rambus Inc.

Rambus assumes no responsibility or liability for any use of the information contained herein. RAMBUS SHALL HAVE NO LIABILITY, DIRECT OR INDIRECT, OR OTHERWISE, ARISING OUT OF THIS DESIGN GUIDE OR THE USE THEREOF.

Components containing Rambus technology, such as Rambus DRAMs and RIMM modules, are manufactured and sold by Rambus licensees. Rambus licensees provide data sheets specific to their Rambus components' products. For a list of Rambus licensees who are provided Rambus memory components, refer to the Rambus web site:


<http://www.rambus.com>

For further information, contact Tektronix:



Worldwide Web: for the most up-to-date product information visit our web site at: www.tektronix.com/rambus

ASEAN Countries (65) 356-3900; Australia & New Zealand 61 (2) 9888-0100; Austria, Central Eastern Europe, Greece, Turkey, Malta, & Cyprus +43 2236 8092 0; Belgium +32 (2) 715 89 70; Brazil and South America 55 (11) 3741-8360; Canada 1 (800) 661-5625; Denmark +45 (44) 850 700; Finland +358 (9) 4783 400; France & North Africa +33 1 69 86 81 81; Germany + 49 (221) 94 77 400; Hong Kong (852) 2585-6688; India (91) 80-2275577; Italy +39 (2) 25086 501; Japan (Sony/Tektronix Corporation) 81 (3) 3448-3111; Mexico, Central America, & Caribbean 52 (5) 666-6333; The Netherlands +31 23 56 95555; Norway +47 22 07 07 00; People's Republic of China 86 (10) 6235 1230; Republic of Korea 82 (2) 528-5299; South Africa (27 11)651-5222; Spain & Portugal +34 91 372 6000; Sweden +46 8 477 65 00; Switzerland +41 (41) 729 36 40; Taiwan 886 (2) 2722-9622; United Kingdom & Ireland +44 (0)1628 403300; USA 1 (800) 426-2200.



From other areas, contact: Tektronix, Inc. Export Sales, P.O. Box 500, M/S 50-255, Beaverton, Oregon 97077-0001, USA 1 (503) 627-6877.

Copyright © 1999, Tektronix, Inc. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.