Introduction

The Direct Rambus™ DRAM is a high speed dynamic random access memory developed by Rambus, Inc. which provides more bandwidth from a single chip than any other DRAM memory technology available today. Direct Rambus DRAMs achieve this bandwidth because they are designed to be an integral part of a complete memory subsystem solution. All of the components that make up the Direct Rambus systems have been engineered to work together in order to achieve the following goals:

- High bandwidth (1.6GB/s)
- Minimum system latency
- Maximum data bus utilization
- Minimum system cost
- Pluggable upgrades
- Same form factor as SDRAM solutions
- Scalable with device density and power supply

The operation of the Direct RDRAM™ is dictated by the requirements of the Direct Rambus subsystem. In order to understand the details of the application specification of the Direct Rambus DRAM, it is important to understand the Rambus™ memory subsystem as a whole.

Direct Rambus Memory System

An example of a Direct Rambus memory subsystem is shown in Figure 1. The subsystem is comprised of the following components:

- Direct Rambus Controller
- Direct Rambus Channel
- Direct Rambus Connector
- Direct Rambus RIMM™
- Direct Rambus DRAMs.

The physical, electrical, and logical portions of all these components have been predetermined and specified by Rambus, Inc. This detail is required to enable the compatibility needed for high speed operation of the Direct Rambus subsystem.

Figure 1. Direct Rambus Memory Subsystem

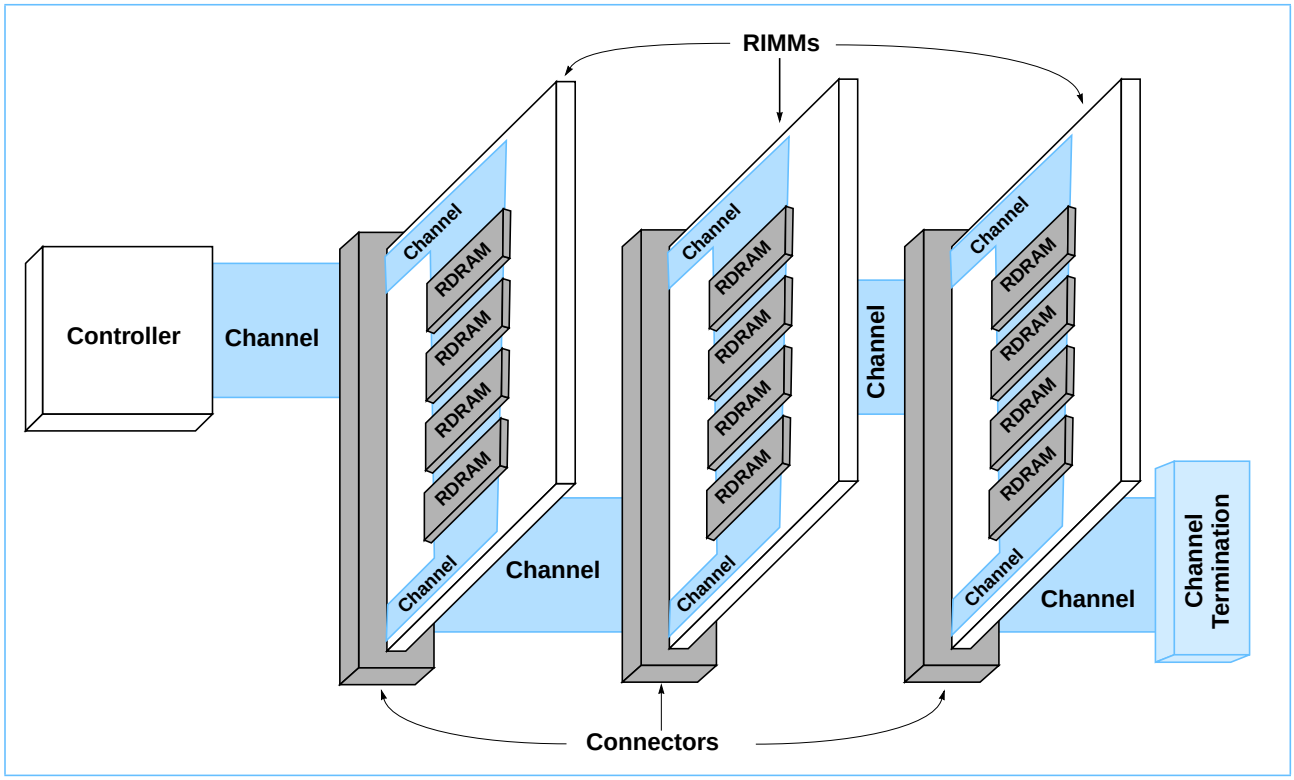

Direct Rambus Controller

The Direct Rambus memory controller is the bus master of the memory subsystem. This controller is placed on a logic chip such as a PC chipset, microprocessor, graphics controller, or an ASIC. Up to four Direct Rambus controllers can be physically placed on a single logic chip. The controller is the interface between the logic chip and the Direct Rambus Channel. Responsibilities of the controller include generating requests, controlling the flow of data and keeping track of the Direct RDRAM states and refresh.

The Direct Rambus controller consists of two separate circuit blocks: the Rambus ASIC Cell (RAC) and the Rambus Memory Controller (RMC). The RAC is clocked at the channel frequency which can be as high as 400MHz. It also functions as the physical and electrical interface to the channel. It consists of the I/O buffers and multiplexers as well as the clocking circuitry for the channel interface. The RAC is designed for a ball grid array (BGA) footprint and requires the use of 76 pins on the logic chip.

The RMC is the control logic for the channel. It acts as the interface between the logic chip and the RAC. The RMC consists of a block of synthesizeable logic which is used to take memory access requests from the logic chip and translate them into the Direct Rambus protocol. Other duties of the RAC include tracking the RDRAM states and optimizing channel performance.

Both the RAC and RMC designs are provided by Rambus, Inc. and are currently available in a wide number of ASIC processes from several ASIC vendors.

Direct Rambus Channel

The Direct Rambus Channel creates an electrical connection between the Rambus controller and the Direct RDRAMs. Channel operation is based on 30 high speed signals which use both edges of a 400MHz clock, creating data transfer rates of 800Mb/s/pin. These high signaling rates are achieved by using the following techniques:

- High quality terminated transmission lines
- Low voltage swing signaling
- Well-defined channel topology
- Precise differential clocking
- Current mode drivers
- Dense device packing with short stub lengths

The channel can be implemented on today's standard FR4 system boards and it fits into the same form factor as current SDRAM memory subsystems. Rambus, Inc. will provide board manufacturers with step-by-step channel layout documentation to ensure the compatibility of different designs.

High Speed Signals

The high speed signals of the channel use Rambus Signaling Levels (RSL) and can be divided into three independent busses: a 16/18 bit data bus, a 3 bit row control bus, and a 5 bit column control bus. These busses operate independently of each other. The data bus is bidirectional and transfers read and write data packets to and from the RDRAMs. The column and row busses are unidirectional and transmit column and row commands to the RDRAMs from the controller. Using separate independent busses for column and row commands, in addition to the pipelined nature of the RDRAMs, allows the data bus to achieve utilizations of up to 95%.

Channel Termination

All of the high speed RSL signals in the Direct Rambus channel are terminated at the end of the channel opposite the controller. This termination consists of resistors which are connected to a termination voltage of 1.8V. The resistors are given values which match the impedance of the channel, enabling them to absorb a signal reaching the end of the channel without generating any reflections.

Channel Clocking

To obtain data transfer rates of 800Mb/s/pin, the Direct Rambus clocking scheme differs significantly from current SDRAM memory systems. The Direct Rambus system uses two differential clocks: clock-to-master (CTM and CTMN) and clock-from-master (CFM and CFMN). Using differential clocks improves noise immunity and reduces jitter in the RAC and Direct RDRAM clocking circuitry.

Figure 2. Direct Rambus System Clocking

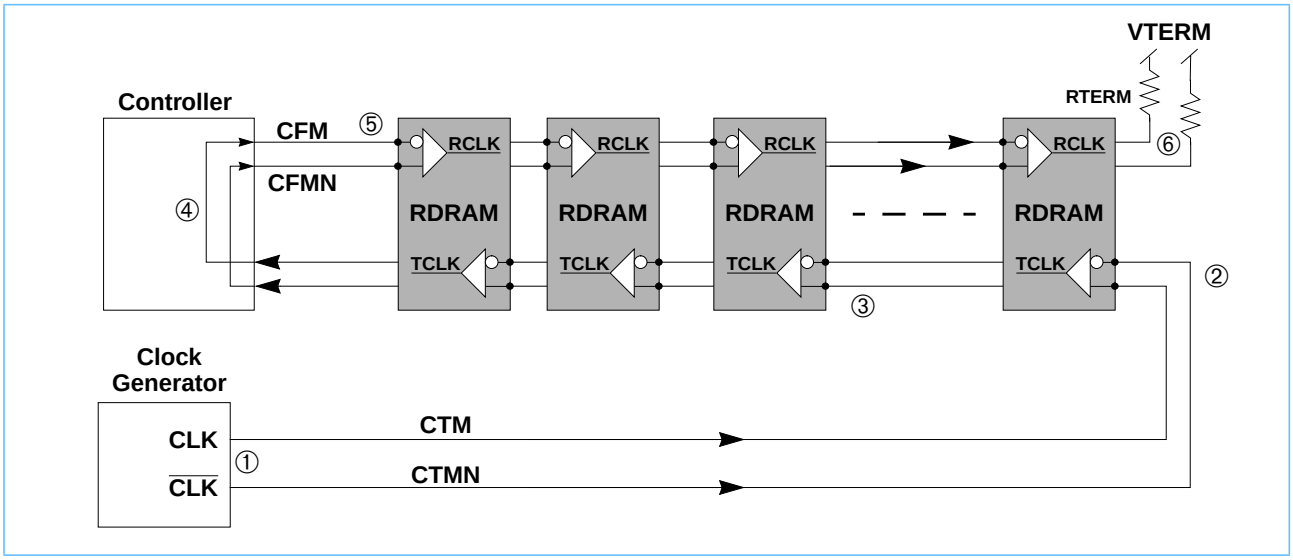

For the following discussion on the Direct Rambus clocking scheme, refer to Figure 2. The differential 400MHz clock can be generated either on the application's logic chip or on a separate clock chip (1). This clock is then routed into the Rambus channel starting from the termination end (2). It travels down the channel toward the controller contacting all of the CTM and CTMN clock inputs of the Direct RDRAMs (3). When it reaches the controller it loops back out (4) into the channel as the CFM and CFMN clocks (5). After contacting all of the CFM and CFMN clock inputs of the RDRAMs, the clock reaches the termination end of the channel where it is resistively terminated to the RSL termination voltage (6).

The CTM differential clock pair operates as the Direct RDRAM transmit clock (TCLK) and the controller receive clock. All read data is aligned to the CTM clock.

The CFM differential pair operates as the Direct RDRAM receive clock (RCLK) and the controller send clock. All write data as well as column and row packet data are aligned to the CFM clock.

Data and its corresponding clock are aligned and always flow in the same direction through the Rambus channel. This allows the controller and Direct RDRAMs to be able to reliably latch all inputs at the 800Mb/s data rate independent of the RDRAM's location on the channel.

Low Speed Signals

The channel also has several low speed LVTTL signals which make up a serial control bus. The primary functions of this low speed bus are initialization of the devices on the channel and performing RDRAM control register reads and writes. These low speed signals are synchronous and use their own low speed clock (SCK).

Packets

Data and control information is transferred across the channel in a packet-oriented protocol. Each of the packets consists of a burst of eight bits over the corresponding signal lines of the channel. For example, a row packet is a burst of eight over the three signal row bus, giving the row command a total of 24 bits. Correspondingly, a column packet consists of 40 bits and a data packet is 128 bits. Using packets for the command and data has two advantages. The first is a reduction in the number of pins or channel signals needed to communicate to or from the RDRAM. This allows the channel to be physically placed in a small area which reduces overall cost and minimizing skew between the channel signals. Keeping the signal skew low is crucial in allowing the Rambus channel to transfer data at 800Mb/s/pin. The second advantage of packetizing is that reserved bits have been placed in the command packets such that Direct Rambus devices from 32Mb to 1Gb use the same number of control signals. This eliminates the need to redesign the mem-

ory subsystem in order to add another address signal every time an advance in device density occurs.

Direct Rambus Connector

The Direct Rambus connector offers a low inductance interface between the Direct Rambus channel on the RIMM module and the Direct Rambus channel on the motherboard. The connector is a through hole connector with 168 contacts that are spaced 1mm (40mil) apart. Contacts are clustered in two groups of 84, at either end of the RIMM socket.

Direct Rambus RIMM

The Direct Rambus RIMM is the expansion memory module which contains one or more Direct RDRAMs and the channel wiring. An eight chip RIMM is shown in Figure 3. The channel enters the RIMM on one end, crosses all of the Direct RDRAMs, and exits out the other end of the RIMM. This channel path is shown in Figure 1. The RIMM essentially connects the channel from one end of the connector to the other. Empty RIMM connectors are not allowed because they would disconnect the channel from the termination and clock sources located on the motherboard at the end of the channel. To solve this problem, RIMMs with only the channel and no RDRAMs were developed. These RIMMs are called “continuity modules” and are required to fill any unpopulated RIMM connectors on the Direct Rambus channel.

Direct Rambus RIMMs are designed to have dimensions similar to SDRAM DIMMs. This allows RIMMs to fit into all current motherboard form factors. Direct Rambus RIMMs have 168 pins spaced on a 1mm (40mil) pitch. Additionally, Direct Rambus RIMMs have been designed to support Serial Presence Detect (SPD) which is used on current SDRAM DIMMs. Unlike SDRAM DIMMs, Direct Rambus RIMMs can be designed with any integer number of Direct RDRAM chips up to the maximum number allowed.

Direct Rambus RIMMs are available as either single-sided or double-sided. The single-sided RIMM uses a six-layer board and can hold one to eight Direct RDRAMs. The double-sided RIMM uses an eight-layer board and can hold up to sixteen Direct RDRAMs. In order to ensure compatibility between RIMMs from different vendors, reference designs and layout rules are provided by Rambus, Inc.

Memory Expansion

A single Direct Rambus channel can support a maximum of 32 Direct RDRAMS on up to three Direct Rambus RIMMs. Using 64Mbit, 128Mbit, and 256Mbit devices, the maximum memory capacity per channel is 256MBytes, 512MBytes, and 1GByte, respectively. To maintain the channel integrity, all empty RIMM slots must be filled with “continuity” RIMMs (specially wired RIMMs with no devices on them).

In order to expand a channel beyond 32 devices, up to two repeater chips can be used. With one repeater chip, a channel can support 64 devices on 6 RIMMs; with two repeater chips, 128 devices can be supported on 12 RIMMs.

Figure 3. Direct Rambus RIMM

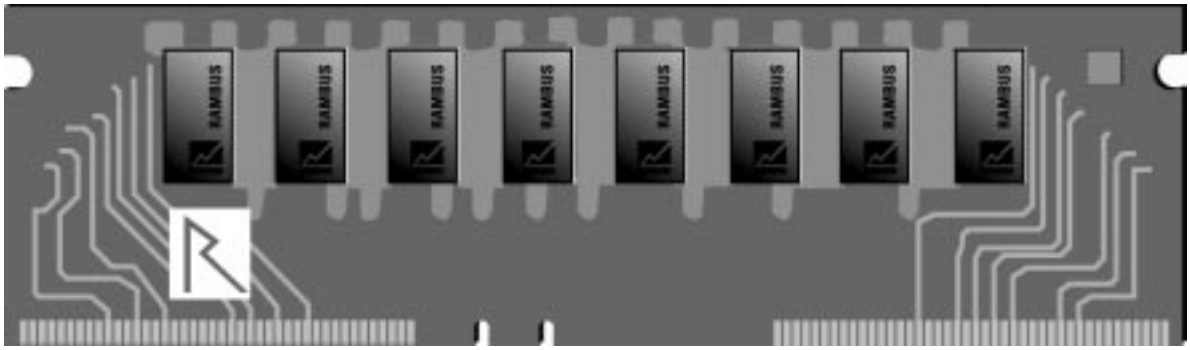

Direct Rambus DRAMs

Direct Rambus DRAMs make up the data storage portion of a Rambus memory subsystem. All devices in a system are electrically located on the channel between the controller and the termination. Direct Rambus Devices can only respond to requests from the controller which makes them bus slaves or responders. A Direct Rambus Device can be divided into two parts, the Rambus interface logic and the DRAM core.

Direct RDRAM Interface Logic

The interface logic translates requests and data between the narrow, packetized, 400MHz RSL inputs and outputs to the wide 100MHz, CMOS, DRAM core which uses standard RAS/CAS-like controls. For example, the interface logic converts a 128 bit data packet, which is transferred across the 16 bit data channel on 8 clock edges of a 400MHz clock, to or from the 100Mb/s/pin 128 bit DRAM core interface. The interface logic consists of the OCDs, receivers, clocking circuitry, control registers, multiplexors and demultiplexors. Rambus, Inc. has designed the interface logic for the Direct RDRAM and supplies this logic to all of the Direct Rambus DRAM licensees. This ensures compatibility of all Direct Rambus DRAMs independent of which DRAM vendor builds the device.

Direct RDRAM Core

The Direct RDRAM core consists of the DRAM array, sense amplifiers, decoders, redundancy, and the row and column control logic. Unlike the interface logic, the DRAM core of the Direct RDRAM is designed by the DRAM vendor. This allows the DRAM vendors to customize their cores for power, cost or performance.

The 64Mbit DRAM core of the Direct RDRAM has 16 dependent banks with each bank organized as a 32K x 128bit DRAM. The wide (x128) DRAM core allows it to operate on the single edge of a 100MHz clock while the narrow (x16) Direct RDRAM interface operates on both edges of a 400MHz clock. The DRAM core uses standard RAS/CAS-like controls which are issued by the Rambus interface logic.

Direct RDRAMs introduce the concept of dependent banks. Dependent banks operate differently than the independent banks used in current SDRAMs. The advantage of dependent banks is the ability to share sense amplifiers across adjacent banks. This makes room for more banks in a smaller die area, leading to a lower device cost. The drawback of using dependent banks is that once a bank is activated, it needs to be precharged before either of the two adjacent banks can be activated. As a result, the maximum number of banks which can be open at the same time is one half of the total number.

For more information contact your IBM Microelectronics sales representative or
visit us on World Wide Web at http://www.chips.ibm.com