advantage of low latency, read and write operations with auto-precharge are recommended.

The MB81E161622 chip is targeted for small-scale systems. For that reason, the output buffer capacity has been reduced to approximately half that of a PC100 memory. There is a version for operation up to a maximum frequency of 100 MHz with *CL* = 2, and a version for operation up to 83 MHz with *CL* = 2. It is estimated that the FCRAM chip requires approximately 25% less current in normal operation and 12% less current in burst operation than a standard SDRAM chip.

Based on the FCRAM technology, Fujitsu also developed several memory devices for cellular telephone applications. Usually, two types of memory are used in cellular phones: the flash memory and low-power SRAM. The low-power SRAM is available at a reasonable price, but it is limited to capacities up to 8 Mbits. As a new alternative to the low-power SRAM, at the middle of the year 2000 Fujitsu introduced a 16-Mbits Mobile FCRAM (MB82D01160). This device uses an asynchronous interface (the standard RAM in cellular phones), so that the cellular phone designers can upgrade easily the memory without making major architectural changes.

A newer version of the Mobile FCRAM, MB82D01171A, reduces the standby current to 70 μA, which is 30% lower than the standby current of the previous version. This version supports the power-down function proposed by Fujitsu. This function uses the FCRAM chip as temporary data storage for a CPU in a cellular phone, thereby eliminating the need for data retention in the working memory while in standby mode. Therefore, RAM current consumption in standby mode is reduced (the power-down current is 10 μA). This contributes to longer battery life, the major factor in cellular phones.

Toshiba is co-developing other FCRAM products with Fujitsu and also manufactures FCRAM chips.

4.4.6.10. DDR SDRAM

Overview

DDR (*Double Data Rate*) SDRAM derives its name from the technique of transmitting data on both the rising and falling edges of the clock signal. For example, with a bus clock of 133 MHz, the DDR memory performs the data transfers at a frequency of 266 MHz. This technique highly increases the effectiveness of the memory bus for data transfer. A DDR SDRAM module with a clock of 133 MHz provides a peak bandwidth of 2100 MB/s, compared to a peak bandwidth of 1066 MB/s for a standard SDRAM module running at the same frequency. However, this does not mean that the average bandwidth is also doubled.

DDR is not a completely new technology. Proprietary DDR SDRAM variants began appearing in 1997, and later on the DDR technique has been used for several high-performance video cards. This technique has also been used for memory devices other than SDRAM: MDRAM (*Multi-bank* DRAM or MoSys DRAM), SLDRAM (*SyncLink* DRAM), and DRDRAM (*Direct Rambus* DRAM). From proprie-

tary solutions, open specifications have emerged, which are not protected by licenses; they have been standardized by the JEDEC committee in 2000. The first DDR SDRAM main-memory modules based on these specifications, operating with a 133-MHz clock, were available in limited quantities at the end of 2000, but the wider distribution began in 2001.

DDR SDRAM is essentially a 2*n* prefetch architecture with an interface designed to transfer two data words per clock cycle at the I/O pins. A single read or write access for the DDR SDRAM consists of a single 2*n*–bit wide data transfer at the internal DRAM array during one clock cycle, and two corresponding *n*–bit wide data transfers at the I/O pins, each during one-half of the clock cycle. Therefore, the internal data bus is twice as wide as the external interface would indicate.

A bidirectional data strobe (*DQS*) is transmitted externally, along with the data, to be used at the receiver. *DQS* is a strobe transmitted by the DDR SDRAM during read operations and by the memory controller during write operations. The data strobe adjusts for variations in clock skew, capacitance, and interconnect length in systems with multiple memory modules.

DDR SDRAM operates with two differential clock signals: *CK* and $\overline{CK}$. The crossing of the positive edge of the *CK* signal with the negative edge of the $\overline{CK}$ signal is referred to as the positive edge of the clock (*CK*). Commands (address and control signals) are registered on every positive edge of the clock signal. Input data is registered on both edges of the *DQS* signal, and output data is referenced on both edges of the *DQS* signal, as well as on both edges of the clock signal.

Read and write accesses to the DDR SDRAM are burst oriented. Accesses start at a selected location and continue for a programmed number of locations in a programmed sequence. The burst length can be programmed to 2, 4, or 8 locations. Accesses begin with the registration of an *Active* command, which is then followed by a *Read* or *Write* command. The address bits registered with an *Active* command are used to select the bank and row to be accessed. The address bits registered with a *Read* or *Write* command are used to select the bank and the starting column location for the burst access. An *Auto Precharge* function may be enabled to provide a row precharge that is initiated at the end of the burst access.

The *CAS* latency of DDR SDRAM can be set to 2 or 2.5 clock cycles. A *CAS* latency of 2.5 clock cycles means 5 clock edges (rising or falling edges). The JEDEC standard also specifies optional *CAS* latencies of 1.5 and 3, which might be supported by some manufacturers.

Because of the high data-transition speeds, on a heavily loaded bus the low-voltage TTL interface cannot be used for the I/O buffers. For that reason, DDR SDRAM uses an I/O interface called SSTL_2 (*Stub-Series Terminated Logic*). This interface is developed for devices operating at 2.5 V; there is a previous version developed for devices operating at 3.3 V, called SSTL_3. The SSTL interfaces do not require off-chip resistors.

DDR SDRAM also uses DIMM modules, but they are not compatible with the DIMM modules used by conventional SDRAM. Therefore, DDR SDRAM can only be used in motherboards with adequate sockets. The main difference is the in-

crease of the number of pins from 168 to 184. However, the DDR DIMM modules have the same dimensions as the conventional SDRAM DIMM modules. The DDR SDRAM memory chips are operating at 2.5 V, instead of 3.3 V. Therefore the power consumption is reduced by about 25%.

Read Operation

Burst read operations are initiated with a *Read* command. The starting column and bank addresses are provided with this command. The *Auto Precharge* operation is either enabled or disabled for that burst access. If *Auto Precharge* is enabled, the row that is accessed will start precharge at the completion of the burst operation.

During burst read operations, the valid data element from the starting column address will be available following the *CAS* latency after the *Read* command. Each subsequent data element will be valid at the next positive or negative clock edge (i.e., at the next crossing of *CK* and $\overline{CK}$). The *DQS* signal is transmitted by the DDR SDRAM along with the output data.

Data from any *Read* command may be concatenated with or truncated with the data from a subsequent *Read* command. In either case, a continuous flow of data can be maintained. The first data element from the new burst operation follows either the last element of a completed burst operation, or the last desired data element of a longer burst operation which is being truncated. The new *Read* command should be issued *x* cycles after the first *Read* command, where *x* equals the number of desired data element pairs (pairs are required by the 2*n* prefetch architecture). This is illustrated in Figure 4.25. If the burst length is 4, the data from the two *Read* commands are concatenated. If the burst length is 8, the second burst operation interrupts the first. The *Read* commands shown must be to the same memory device. For the generic *Read* commands shown in this figure, the *Auto Precharge* operation is disabled.

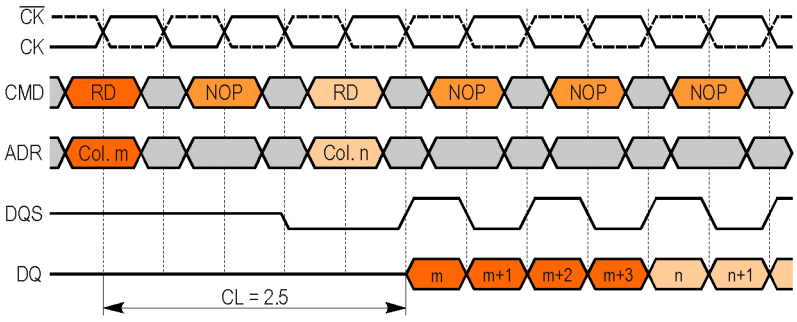

Figure 4.25. Consecutive burst *Read* operations for the DDR SDRAM memory.

Write Operation

Burst write operations are initiated with a *Write* command. The starting column and bank addresses are provided with the *Write* command, and the *Auto Precharge*

operation is either enabled or disabled for that access. If *Auto Precharge* is enabled, the row being accessed is precharged at the completion of the burst operation.

During burst write operations, the first valid input data element will be registered on the first rising edge of the *DQS* signal following the *Write* command, and subsequent data elements will be registered on successive edges of the *DQS* signal. The time between the *Write* command and the first corresponding rising edge of the *DQS* signal is specified with a relatively wide range (from 75% to 125% of one clock cycle).

DDR SDRAM Types

Usually, the various DDR SDRAM types are named after their peak transfer rate. For example, the DDR SDRAM which uses a 100-MHz clock, with a maximum transfer rate of 1600 MB/s, is called PC1600. Similarly, the 133-MHz version is called PC2100. Technically, these DDR SDRAM memories should be called PC200 and PC266, which was the original plan. The name was changed because Rambus uses names in the same style, for example, PC800 for a Rambus DRAM which operates at 800 MHz. In the JEDEC standard, other names are used for DDR SDRAM, based on the operating frequency. For example, DDR-200 indicates a DDR SDRAM performing data transfers at 200 MHz (with a 100-MHz clock), and it is equivalent with PC1600 DDR SDRAM.

Table 4.6 shows the main types of DDR SDRAM memories, their clock frequency and maximum theoretical bandwidth.

Table 4.6. Types of DDR SDRAM and their maximum theoretical bandwidth.

Memory type	Clock frequency	Maximum bandwidth
DDR-200 (PC1600)	100 MHz	2 x 100 x 8 = 1600 MB/s
DDR-266 (PC2100)	133 MHz	2 x 133 x 8 = 2128 MB/s
DDR-300 (PC2400)	150 MHz	2 x 150 x 8 = 2400 MB/s
DDR-333 (PC2700)	166 MHz	2 x 166 x 8 = 2656 MB/s

As in the case of other memory types, the maximum theoretical bandwidth does not indicate the overall performance, because the average bandwidth must also be considered. For example, a DDR-266 memory should have an average bandwidth of about 1200 MB/s (or lower), with an efficiency of 56.4%. In comparison, a *Direct Rambus* DRAM operating at 800 MHz has an average bandwidth of about 1500 MB/s, with an efficiency of 94.7%. However, the initial latency of the DRDRAM is much higher.

Usually, the DDR-200 (PC1600) memories have latencies of 2:2:2. The other versions have higher latencies, expressed in number of clock cycles. For instance, the DDR-266 (PC2100) may have a latency designation of 2.5:3:3, which means a *CAS* latency of 5 clock edges (rising or falling edges), or 2.5 clock cycles. There are higher-

quality DDR-266 chips (called DDR-266A or PC2100A) with latencies of 2:2:2, but they are more expensive.

The DDR memories are used for two different types of applications: module-based and component-based (point-to-point) applications. Module-based applications generally refer to main-memory systems; the modules contain several memory chips, depending on the memory organization (usually, ×4 or ×8). For example, in the case of the ×8 organization, 8 memory chips are needed for a 64-bit bus. Module-based systems are limited in speed by bus loading and line lengths. In the case of point-to-point applications, such as graphics cards, higher speeds are possible. For these applications, usually memory chips with ×16 or ×32 organizations are used.

In order to use the DDR SDRAM memory, a supporting chipset is needed. Several chipset manufacturers are supporting DDR SDRAM, such as Via Technologies (Apollo Pro 266) or AMD (AMD 760). They expect DDR SDRAM to be used in all segments of the PC market. Until recently, Intel does not supported the DDR SDRAM memory for desktop PCs because of his agreement with Rambus Inc. This agreement, revealed in the summer of 2000, does not allow Intel to implement chipset support for other memories than Rambus, except if the bandwidth is lower than 1 GB/s. The agreement only refers to the desktop PCs and does not include the server market. It seems that Intel demanded Rambus to remove from the licensing contract the clause which does not allow it to develop a DDR chipset until 2003. As a result, at the beginning of 2002 Intel activated the support for DDR SDRAM memory in some of their chipsets.

Intel does implement support for the DDR memory in the chipsets targeted to servers and high-end workstations based on the 64-bit Itanium processor and its successors Foster and McKinley. Intel licensed Via Technologies to develop DDR-supporting chipsets for the Pentium 4 processor.

Only a fast processor, such as the Pentium 4, can take advantage of the DDR memory. AMD's Athlon processors (Thunderbird, XP) can also use the DDR memory, because they have a fast, 200/266 MHz front side bus (FSB). Intel's Pentium III Coppermine processors with a 133-MHz FSB cannot make an efficient use of the DDR memory.

Currently, at the beginning of 2002, DDR SDRAM chips with capacities of 64 Mbits, 128 Mbits, and 256 Mbits are available. It is expected that 512-Mbit chips will be available in the first half of 2002. The development of 1-Gbit DDR SDRAM chips is underway.

A low-latency variant of the DDR memory, called RLDRAM (*Reduced Latency* DRAM), was proposed by Micron Technology and Infineon Technologies AG. This variant offers a higher sustained bandwidth, of up to 2.4 GB/s. Other improvements include: reduced row cycle time; reduced initial latency, which is half that of DDR SDRAM; higher bus utilization; more banks (8 vs. 4). No address multiplexing is used, which allows a fast access, similar to that of SRAM memories. The production of RLDRAM devices began at the beginning of 2002.