Integrated Device Technology, Inc.

IDT79R3051™ ADDRESS/DATA BUS TURN AROUND BEHAVIOR

APPLICATION NOTE AN-97

by Andrew Ng

INTRODUCTION

This application note describes the behavior of the R3051's multiplexed Address/Data, "A/D" bus and presents the issues of a particular topic called "Bus Turn Around." Bus Turn Around will be defined, design issues will be presented, and design solutions will be given for conventional R3051 systems, as well as a "DMA $\overline{\text{BusReq}}$" design solution for very low speed and very high speed systems.

Definition of the R3051

The IDT79R3051™ RISController™ is a highly integrated MIPS™ R3000™ instruction set compatible microprocessor that minimizes system cost and power consumption. The R3051 includes 4kB to 8kB of instruction cache, 2kB of data cache, an optional on-chip TLB memory management unit, 4-deep read and write buffers, on-chip DMA arbitration, a simple external bus interface, as well as the R3000A CPU execution engine — all in a single compact plastic 84-pin package.

Definition of the A/D Bus

One of the key features of the R3051 is its low pin count. The low pin count is largely a result of its simple control interface and its use of a multiplexed Address and Data bus, called A/D(31:0). As shown in Figures 1 and 2, the multiplexed A/D bus drives its address during the first phase of a read or write memory cycle. In the 2nd phase of a read memory cycle, the CPU expects the external memory system to drive the bus and return the data. In the 2nd phase of a write memory cycle, the CPU drives the data out to the memory system. Thus in a typical R3051 system, the address can be latched using a bank of transparent latches such as with the 54/74FCT373T or 54/74FCT841T as shown in Figures 4 and 5 so that the address is de-multiplexed from the data lines.

In systems using an ASIC, such as for a DRAM or DMA Controller or as an Integrated I/O Subsystem/Controller with on-chip programmable registers, the multiplexed A/D bus has an advantage over separate Address and Data busses in that the ASIC requires substantially fewer pins. The ASIC can latch the 32 Address bits internally, using the Address Latch Enable output from the CPU called "ALE", and then use the same input pins to provide data. In addition, the CPU has less noise from simultaneous switching of the 32 A/D lines than if it had to switch 64 separate Address and Data lines. Thus R3051 systems can often save cost and space by using inexpensive and low pin count ASICs.

Although a multiplexed bus may be thought of as a disadvantage in terms of system performance, this is rarely the case in R3051 systems. An analysis of memory behavior and the bus shows that in conventional memory systems (those that do not use exclusively high-speed, single cycle SRAMs for the entire memory system), the R3051 bus structure causes no real performance loss.

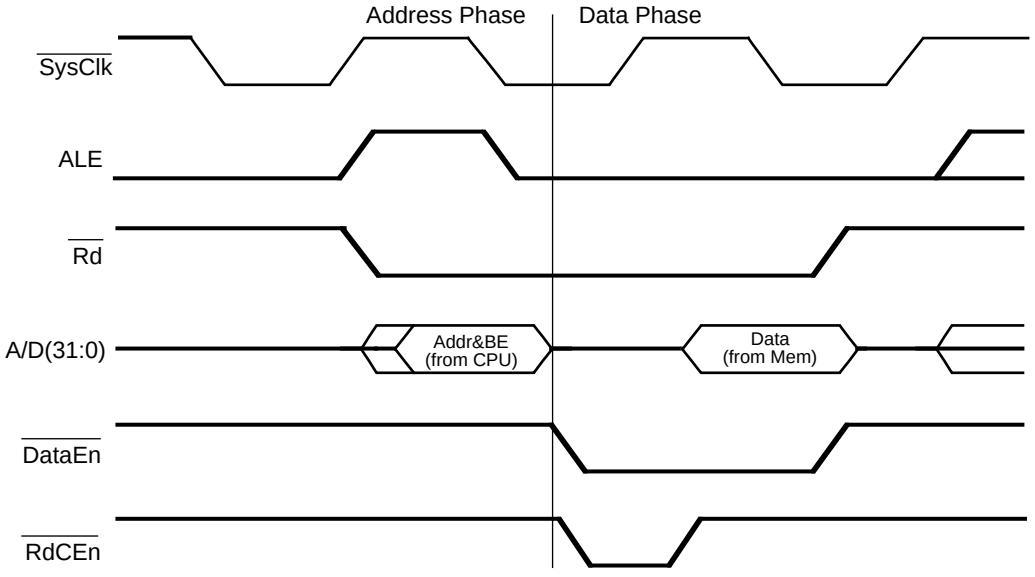

Figure 1. R3051 Read Cycle

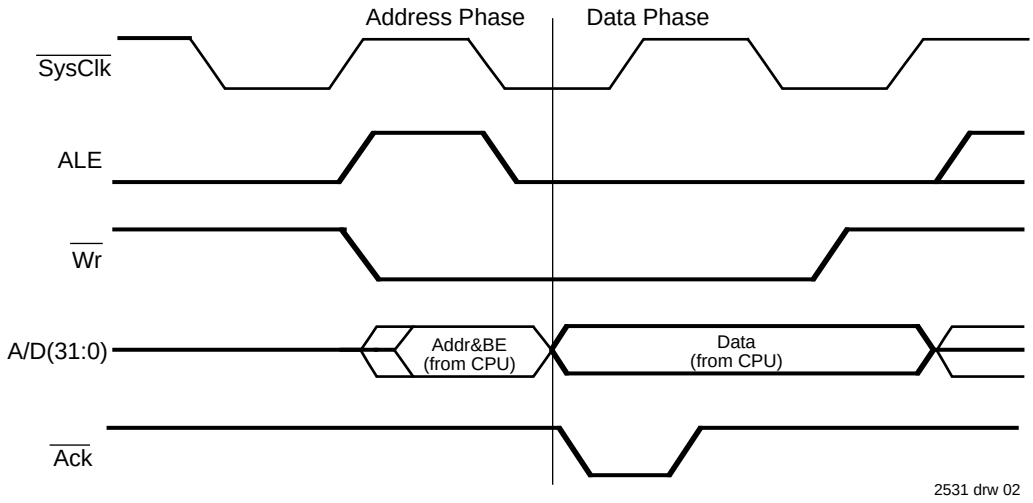

Figure 2. R3051 Write Cycle

For example, conventional memory systems use the address before the data is generated on read cycles or needed by write cycles. On read cycles, the address is always needed before the data array can be accessed. The multiplexed R3051 bus provides the address as early as a non-multiplexed bus would; thus, the read access is not delayed. Since memory read performance is described as "Address and Chip-Select valid to Data Available", the multiplexed bus causes no performance loss on reads.

Similarly, on write cycles, most memories (except for self-timed memories) require the address before the data in order to properly coordinate the write strobe with the correct internal row and column address decode/selects. The R3051 bus provides the write target address for one-half cycle, and then immediately presents the write data. That half cycle is required to perform address decoding, and to provide a Chip-Select to the memory device. Thus, once the address and Chip-Select are available to the memory, the data is also available.

Further, the R3051 decouples the system bus performance from processor performance based on the integration of on-chip resources. Specifically, the large on-chip caches minimize the number of main memory reads, thus making system read performance less criticial. The on-chip 4-deep write buffer isolates the processor from the memory system write speed, allowing it to continue execution while store operations are actually updated into the memory. Thus, R3051 performance, while somewhat dependant on memory system performance, is largely isolated from the memory system. Thus, high-performance systems using relatively slow EPROM and DRAM devices can be easily realized.

Definition of Bus Turn Around

The other consequence of a multiplexed bus arises from the fact that during a particular transaction, as well as from one transaction to the next, transitions between sources of the bus can occur. For example, a read transaction begins with the processor driving the address on the bus, and ends with the memory driving the data on the bus. Similarly, at the end of a read, the next transaction on the bus will begin again with the CPU driving an address on the bus.

Note that similar concerns are present even for non-multiplexed busses. For example, a read followed by a write results in the data bus first being driven by the memory, and then being driven by the CPU. Thus, bus turn-around is also a consideration in non-multiplexed bus systems.

Bus Turn Around behavior is the action that the CPU takes when its address/data bus transitions between the CPU and the memory, particularly when it changes direction from being a driver to being a non-driver or vice-versa. The actions that the CPU can take are:

1. Drive the address.
2. Drive the data.
3. Tri-state.

There are two basic times when the A/D bus will transition:

1. Intra-Cycle — Within a memory cycle as the address phase transitions into the data phase.
2. Inter-Cycle — Between two memory cycles when the data phase transitions into the address phase of the next memory cycle.

Intra-Cycle Bus Turn Around

A typical case of an address to data transition happens during a read cycle. As shown in Figure 1, when the Address Latch Enable (ALE) is negated, the address is externally latched and the CPU turns the bus around by tri-stating the A/D bus, so that the external memory system can begin to drive the expected data back to the CPU. The second case occurs during write cycles when the CPU finishes driving the address, it begins driving the data to the memory system. Since the CPU drives both the address and data during write cycles, bus turn around is not a significant issue during write cycles. The two intra-cycle transition cases are listed in Table 1, which shows the state of the CPU A/D output buffers during the address and data phases of the transaction.

Note that the processor provides an output, $\overline{\text{DataEn}}$, to indicate that this transaction has occurred. During the addressing phase, $\overline{\text{DataEn}}$ is negated, indicating the CPU is driving the A/D bus. During the Data Phase, $\overline{\text{DataEn}}$ is asserted, indicating that the bus is to be driven by the external memory system. During write cycles, and during idle cycles, $\overline{\text{DataEn}}$ is guaranteed to be negated, indicating that the external memory system should not be driving the A/D bus.

READ	A,Z
WRITE	A,D

Note: A — Address, D — Data, Z — Tri-State

Table 1. R3051 Address to Data Bus Transitional Behavior Within Memory Cycles

Inter-Cycle Bus Turn Around

A typical case of the transition between two memory cycles occurs on a read cycle that is immediately followed by a write cycle as shown in Figure 3. In this case, the memory system is required to turn the bus around by tri-stating the bus before the next write cycle begins to drive its address onto the A/D lines. Table 2 lists the R3051's behavior on each of the cases of inter-cycle memory transitions. The table lists the state of the CPU output buffers at the end of the first transaction, followed by the state of the buffers at the beginning of the next transaction. Note that if a read or write cycle occurs while the CPU is executing instructions from its internal cache, the next external memory cycle might not occur until many clocks later, in which case the A/D bus is tri-stated since it is idle. Also, many of the cases, such as the transitions after writes have both the data and address driven by the CPU. Thus bus turn around is not a significant issue after write cycles. Other transitions may not actually be possible. For example, it is impossible to have a read followed by a read. At least one idle cycle is required, to accomodate the internal fix-up cycle required by the processor (see the R3051 Hardware User's Manual for more detail).

From To	READ	WRITE	DMA	IDLE
READ	Z,A	Z,A	Z,Z	Z,Z
WRITE	D,A	D,A	D,Z	D,Z
DMA	Z,A	Z,A	Z,Z	Z,Z
IDLE	Z,A	Z,A	Z,Z	Z,Z

Note: A — Address, D — Data, Z — Tri-State

Table 2. R3051 Data to Address Bus Transitional Behavior Between Memory Transactions

TYPICAL SYSTEMS AND BUS TURN AROUND

To handle the timing associated with the bus turn around within a memory cycle, the Data Enable output, $\overline{\text{DataEn}}$ is provided by the R3051. As shown in Figure 1, on read cycles, $\overline{\text{DataEn}}$ gives an indication when the CPU has tri-stated the A/D bus. Thus after $\overline{\text{DataEn}}$ asserts, the memory system can begin driving data onto the A/D bus. The system designer can also look for the rising clock edge of $\overline{\text{SysClk}}$ after $\overline{\text{Rd}}$ asserts before allowing the memory system to drive data.

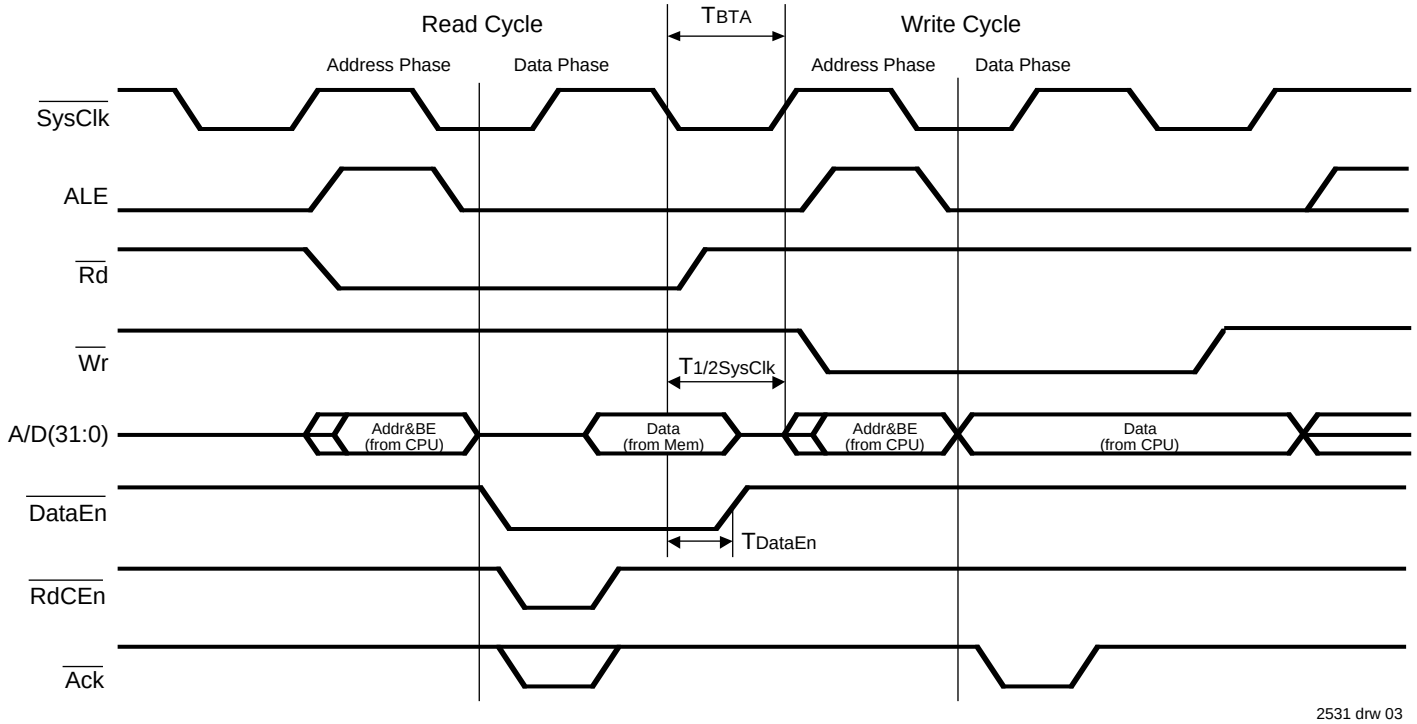

Figure 3. R3051 Read Cycle Followed by a Write Cycle

To handle the timing associated with the bus turn around between two memory transactions, consider the case of a read cycle immediately followed by a write cycle. The read cycle output enable control of the memory system must be such that the output drivers of the memory system turn off within 1/2 clock before the next address is driven by the write cycle. If the memory devices have an output disable to tri-state time (T_{OEZ}) of more than 1/2 clock, then they can be isolated from the A/D bus with a bank of data transceivers such as the 54/74FCT245T, 54/74FCT861, or 54/74FCT623T or with latched data transceivers such as the 54/74FCT543T or 54/74FCT646T as shown in Figure 4. All of these transceivers have very fast output disable times.

VERY FAST $\overline{\text{SysClk}}$ OR VERY SLOW T_{OEZ} AND BUS TURN AROUND

The majority of systems will use evenly matched memories relative to the system clock speed or use transceivers. However, two exceptions may occur:

1. Very Fast $\overline{\text{SysClk}}$ — Even with the highest speed transceivers, their output disable times (T_{OEZ}) are around 5-8 nsec. Thus at 40 MHz, if $\overline{\text{DataEn}}$ is used, it has a clock to de-assert time of 4 nsec. (Assume that the transceiver has two internally And'ed output enable inputs. For example, as shown in Figure 4, the FCT543T transceiver bank can use $\overline{\text{DataEn}}$ and the bank select for inputs to the output enables). If 1 nsec is allowed for clock skew, this just meets the worst case timing criterion of:

   $T_{1/2SysClk}$ (12.5) ≥ T_{DataEn} + T_{OEZ} + $T_{ClkSkew}$ + T_{Cap} (4+6.5+1+0)

   Some choices of transceiver and PLA-based output enable control combinations may need more time than is allowed by the above equation. Solutions to this problem will be given in the section below, "Using DMA $\overline{\text{BusReq}}$ to Match CPU and Memory Speeds."

2. Very Slow Memories — The second case occurs when relatively slow T_{OEZ} memories are attached directly to the A/D bus as shown in Figure 5. Such systems require these memories to turn off within 1/2 clock. A 20 MHz R3051 has a T_{DataEn} for the de-asserting edge of $\overline{\text{DataEn}}$ of 7 nsec. Assume that additional output enable control circuitry adds an additional delay of 10 nsec. 1 nsec is allowed for clock skew. For an inexpensive, slow 120 nsec EPROM, the output disable time is about 50 nsec, which seems to limit the clock speed to about 7 MHz:

   $T_{1/2SysClk}$ (71.4) ≥ T_{DataEn} + $T_{OutputEnableControl}$ + T_{OEZ} + $T_{ClkSkew}$ + T_{Cap} (7+10+50+1+0)

   However, as will be explained below in the section called, "Using DMA $\overline{\text{BusReq}}$ to Match CPU and Memory Speeds," the overall CPU speed does not have to be slowed down just because a slow T_{OEZ} memory is attached directly to the A/D bus.

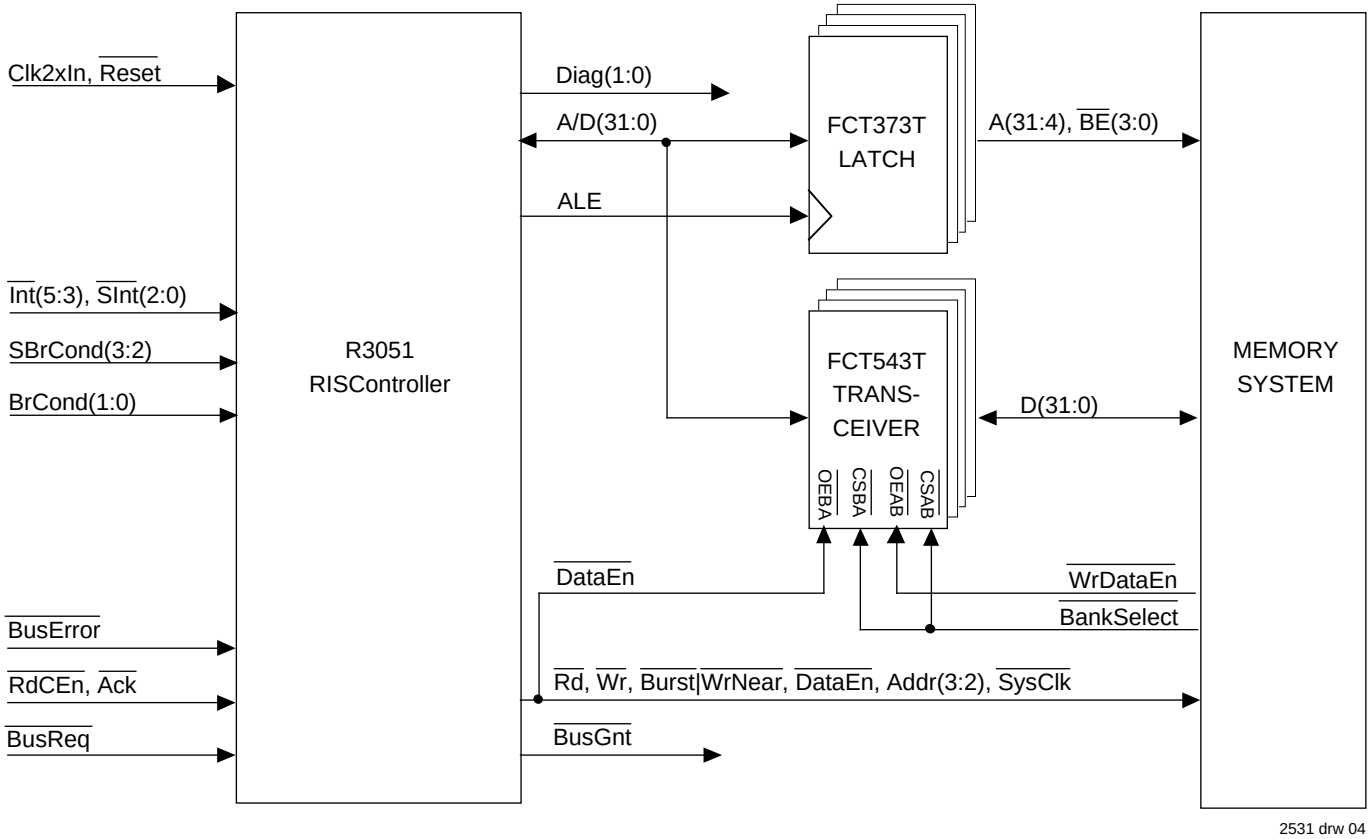

Figure 4. R3051 Memory System Isolated with Transceivers

USING DMA BusReq TO MATCH CPU AND MEMORY SPEEDS

For systems with very fast SysClk or very slow memories, a solution exists to the bus turn around timing constraints by using the Direct Memory Access (DMA) interface on the R3051. The R3051 DMA interface consists of two pins called BusReq and BusGnt as shown in Figure 6. Normally these pins are used for giving an external device control of the CPU bus instead of giving control of the bus to the R3051. In the R3051, when BusReq is asserted, DMA always has the highest priority immediately after the current memory cycle completes. The BusReq input is always sampled on the rising edge of SysClk. After the BusGnt is given, all of the CPU control line outputs, except SysClk and BusGnt are tri-stated. When the DMA device is finished with the bus, it de-asserts BusReq which then causes the CPU to de-assert BusGnt. The BusGnt output is always asserted on the rising edge of SysClk and de-asserted on the falling edge of SysClk.

Because a BusReq always has the highest priority, in a very fast SysClk system or a very slow memory system, asserting BusReq during the read cycle insures that the DMA request will always be granted at the end of the read cycle. After this happens, the BusReq pin can be de-asserted after the desired number of inter-cycle wait-states have been inserted. For example, as shown in Figure 7, by attaching the buffered read line, Rd to BusReq, the R3051 will grant the BusReq and immediately release it. Note that Rd needs to be buffered to meet the hold time of the BusReq input. Examine Figure 3, where a write cycle normally can follow a read cycle after 0.5 clocks and then compare it with Figure 7. In Figure 7, by using BusReq, it can be seen that a minimum of 1.5 clocks is guaranteed before the next memory cycle is started by the CPU.

Note that when using DMA, the system may choose to resistively pull-up or down its control signals since the DMA when granted will tri-state the CPU control output signals. Thus ALE could use a pull-down, while Rd, Wr, DataEn, and Burst|WrNear could use pull-ups. The resistor value of the pull-ups and pull-down is not that critical since the R3051 always drives the control signals to their de-asserted states before tri-stating them. Also, if the BusReq is needed for conventional DMA, a fixed-priority based arbiter can be used to allow bus turn around wait-state injection the highest priority and to allow conventional DMA the next priority.

Various improvements can be made to using the Rd line for BusReq. For example, instead of using the buffered Rd line, use the decoded chip select of the particular memory (e.g., the EPROM) that has the relatively slow TOEZ. Thus the extra wait-states are only asserted as needed (that is, after the slow memory is accessed).

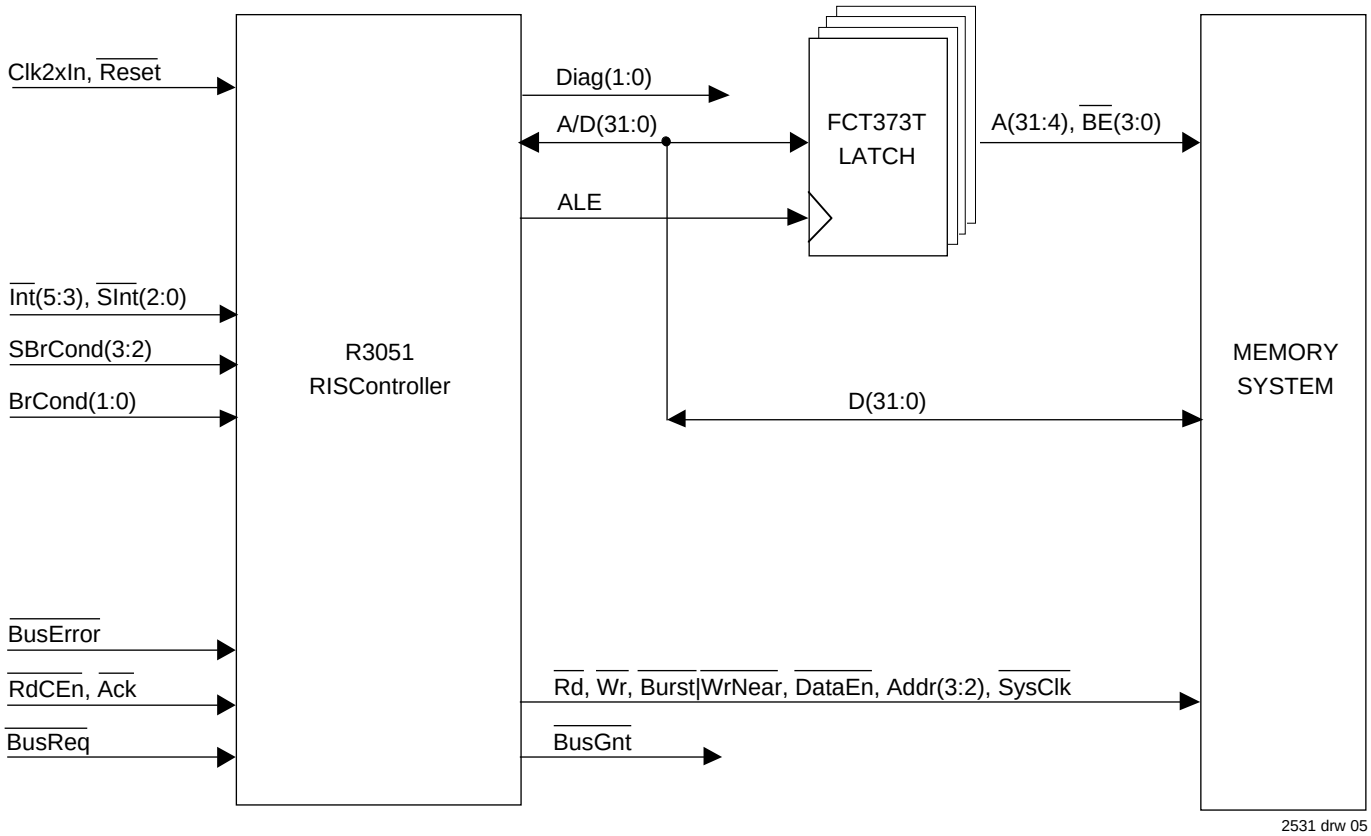

Figure 5. R3051 Memory System Connected Directly to the A/D Bus

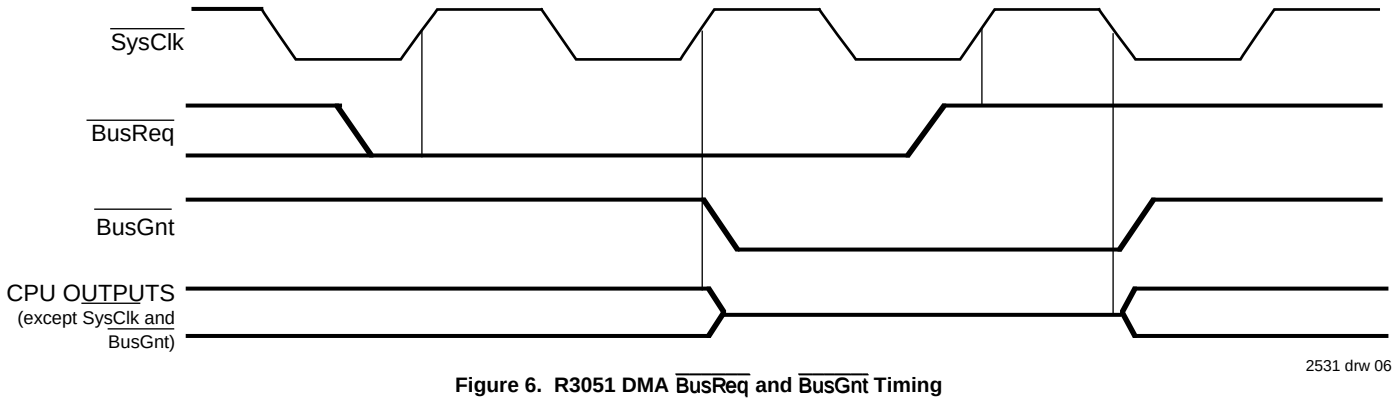

Figure 6. R3051 DMA $\overline{\text{BusReq}}$ and $\overline{\text{BusGnt}}$ Timing

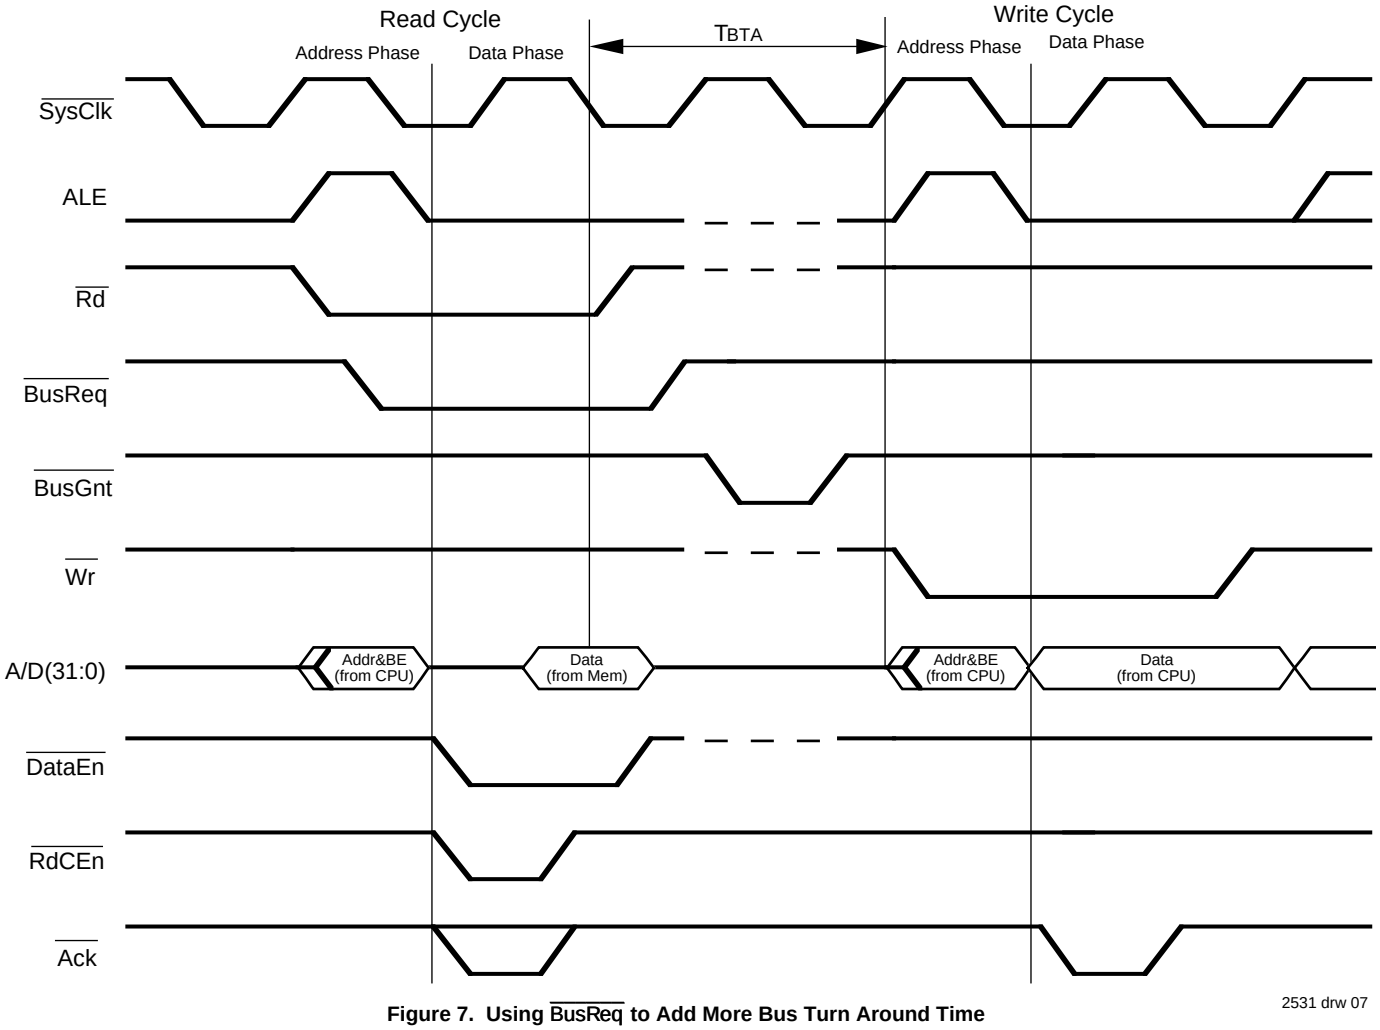

Figure 7. Using $\overline{\text{BusReq}}$ to Add More Bus Turn Around Time

SUMMARY

The R3051 allows inexpensive systems to be designed with the high throughput R3000 RISC instruction set architecture. The small 84-pin count is achieved with a multiplexed address and data bus, called “A/D”. The use of the multiplexed A/D bus allows ASICs and Memory Controllers such as the R3721 DRAM Controller to have fewer interface pins, with no real loss of system performance or real added complexity. However, as for any high-speed bus (either multiplexed or not) care has to be taken to avoid bus clashes as the bus transitions from one device to another. This applications note describes these considerations.

As shown in the text, the use of the A/D bus does not inherently limit the overall clock speed of the system, since either transceivers, or the described method of using the DMA $\overline{\text{BusReq}}$ input gives a solution for memory/CPU mismatches. Thus any memory or I/O system can use the multiplexed A/D bus and be designed to run at the full CPU clock frequency.

FOR FURTHER INFORMATION:

1. IDT79R3051 Family Hardware User's Manual, Integrated Device Technology, Inc., MAN-RISC-00051, Santa Clara, CA, 1991. — Describes the H/W features and functionality of the device as well the bus interface.
2. IDT 1991 RISC Data Book, Integrated Device Technology, DBK-RISC-00021, Integrated Device Technology, Inc., Santa Clara, CA, 1991. — Contains the data sheet with packaging, pinout, AC/DC electrical and thermal parameters.
3. G. Kane, *MIPS RISC Architecture*, Prentice Hall, Englewood Cliffs, NJ, 1988. — Describes the R3000/ R3051 instruction set architecture from a systems and assembly level programming perspective.
4. IDT 1991 Logic Data Book, Integrated Device TechnologyInc., Santa Clara, CA, 1991. — Contains the data sheets of many different high-speed FCT transceivers, latches, and buffers.