

11. Local Memory Interface

The TMS34010 local memory interface consists of a triple-multiplexed address/data bus and associated control signals. Several types of memory cycles, including read, write, screen-refresh, and DRAM-refresh cycles are supported. During a memory cycle, the row address, column address, and data are transmitted over the same physical bus lines. The row and column addresses necessary to address DRAMs and VRAMs are available directly at the address/ data pins, eliminating the need for external multiplexing hardware.

The TMS34010 interfaces directly to DRAMs and VRAMs, and can be programmed to perform DRAM-refresh cycles at regular intervals. $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ or $\overline{\text{RAS}}$ -only refresh cycles may be selected. The GSP can also be programmed to perform screen refresh by scheduling VRAM shift-register transfer cycles to occur at regular intervals.

The local memory interface provides a hold/hold acknowledge capability that allows external devices to request control of the bus. After acknowledging a hold request, the GSP releases the bus by driving its address/data bus and control outputs into high impedance.

Section	Page
11.1 Local Memory Interface Pins	11-2
11.2 Local Memory Interface Registers	11-3
11.3 Memory Bus Request Priorities	11-4
11.4 Local Memory Interface Timing	11-5
11.5 Addressing Mechanisms	11-23

11.1 Local Memory Interface Pins

TMS34010 pin functions are described in detail in Section 2. This section briefly summarizes the local memory interface pins.

LAD0–LAD15

These pins form the local multiplexed address/data bus.

$\overline{\text{DEN}}$ The local data enable signal is driven active low to allow data to be written to or read from LAD0–LAD15. (Connects to the $\overline{\text{G}}$ pins of a pair of optional '245-type octal bus transceivers.)

DDOUT The local data direction out signal is driven high to enable data to be output on LAD0–LAD15. It is driven low to enable data to be input on LAD0–LAD15. (Connects to the DIR pins of a pair of optional '245-type octal bus transceivers.)

$\overline{\text{LAL}}$ The high-to-low transition of the local address latched signal is used by an external '373-type latch to capture the column address from LAD0–LAD15.

$\overline{\text{RAS}}$ The local row address strobe signal drives the $\overline{\text{RAS}}$ inputs of DRAMs and VRAMs.

$\overline{\text{CAS}}$ The local column address strobe signal drives the $\overline{\text{CAS}}$ inputs of DRAMs and VRAMs.

$\overline{\text{W}}$ The local write strobe signal drives the $\overline{\text{W}}$ inputs of DRAMs and VRAMs.

$\overline{\text{TR}}/\overline{\text{OE}}$ The local shift register transfer/output enable signal connects to the $\overline{\text{TR}}/\overline{\text{OE}}$ (or $\overline{\text{DT}}/\overline{\text{OE}}$) pins of a VRAM.

LRDY The local ready signal is driven low by external circuitry to inhibit the TMS34010 from completing a local memory cycle.

INCLK TMS34010 processor functions are synchronous to this input clock signal. (Video timing is controlled by VCLK.)

**LCLK1,
LCLK2** These output clocks are available to the board designer for synchronous control of external circuitry.

**$\overline{\text{LINT}}\ 1,$
 $\overline{\text{LINT}}\ 2$** Interrupt requests are transmitted to the GSP on these pins.

11.2 Local Memory Interface Registers

The local memory interface registers are summarized below. These registers are a subset of the I/O registers which are detailed in Section 6.

The memory **CONTROL** register contains several programmable parameters that provide control of the local memory interface:

RM. DRAM refresh mode, bit 2. Selects $\overline{\text{RAS}}$ -only or $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ refresh cycles.

RR. DRAM refresh rate, bits 3 and 4. Controls the frequency of DRAM refresh cycles.

T. Transparency enable, bit 5. Enables or disables the pixel attribute of transparency.

W. Window violation detection mode, bits 6 and 7. Selects the course of action the GSP will follow when it detects a window violation.

PBH. PixBlt horizontal direction, bit 8. Determines the horizontal direction (increasing X or decreasing X) for pixel operations.

PBV. PixBlt vertical direction, bit 8. Determines the vertical direction (increasing Y or decreasing Y) for pixel operations.

PPOP. Pixel processing operation select, bits 10–14. Selects among several Boolean and arithmetic pixel processing options.

CD. Instruction cache disable, bit 15. Enables or disables the instruction cache.

The **CONVDP** register contains the destination pitch conversion factor that is used during XY-to-linear conversion of a destination pixel address.

The **CONVSP** register contains the source pitch conversion factor that is used during XY-to-linear conversion of a source pixel address.

The **PMASK** (plane mask) register selectively disables or enables various planes in a multiple-bit-per-pixel bit map display.

The **PSIZE** (pixel size) register specifies the number of bits per pixel.

The **REFCNT** (refresh count) register generates the addresses output during DRAM-refresh cycles and counts the intervals between successive DRAM-refresh cycles.

11.3 Memory Bus Request Priorities

The GSP's local memory interface controller assigns priorities to requests from various sources, both on and off chip, for local memory cycles. Table 11-1 lists these priorities (priority 1 is highest).

Table 11-1. Priorities for Memory Cycle Requests

Priority	Memory Cycle Requested
1	Hold request from external bus master device
2	Screen-refresh cycle
3	DRAM-refresh cycle
4	Host-initiated indirect read or write cycle
5	GSP CPU memory cycle

A GSP CPU memory cycle is a read or write performed by the GSP's on-chip 32-bit processor. Insertion of a field (or a portion of a field spanning multiple words) into a word requires *two* CPU memory cycles when the field does not begin and end on word boundaries. The two cycles are a read followed by a write. This sequence is called a read-modify-write operation. The read and write are performed as separate memory cycles, but are treated as indivisible; that is, the memory controller will not permit another memory request to be serviced between the read and its accompanying write. The only exception to this statement is the hold request.

While a read-modify-write operation on an individual memory word is indivisible, the accesses necessary to extract or insert a field spanning multiple memory words are not. For example, if a field spans portions of two memory words, a higher priority access such as a host-indirect cycle can occur between the two read-modify-write operations required to insert the field.

The hold request has the highest priority. An external device requests control of the bus by signalling a hold request to the GSP. The external device may perform multiple memory cycles following acknowledgment from the GSP. However, the device should not control the bus for so long that necessary screen-refresh and DRAM-refresh cycles are prevented from occurring. Indirect accesses initiated by a host processor will be blocked as long as the external device continues to control the bus. If the host processor attempts to initiate another indirect access during this time, the host will be forced to wait at the host interface (the GSP sends it a not-ready signal) until the external device releases the local bus.

A memory cycle already in progress will always be permitted to complete, even if a higher priority request is received while the cycle is still in progress.

11.4 Local Memory Interface Timing

The TMS34010 memory interface contains a triple-multiplexed address/data bus on which row addresses, column addresses and data are transmitted. Figure 11-1 illustrates multiplexing of addresses and data.




Figure 11-1. Triple Multiplexing of Addresses and Data

The TMS34010 LAD pins directly provide the multiplexed row and column addresses needed to drive dynamic RAMs and video RAMs. Any eight adjacent pins in the range LAD0-LAD10 provide 16 contiguous logical address bits; the eight MSBs are output as part of the row address, and the eight LSBs are output as part of the column address. For example, Figure 11-1 shows that logical address bits 5-20 are output at LAD1-LAD8.

The control signals output to memory support direct interfacing to DRAMs and VRAMs. At the beginning of a memory cycle, the address is output in multiplexed fashion as a row address followed by a column address. The remainder of the cycle is used to transfer data between the TMS34010 and memory. Figure 11-2 (page 11-6) illustrates general timing (the local address/data pins are identified as the LAD Bus)




Figure 11-2. Row and Column Address Phases of Memory Cycle

Figure 11-3 through Figure 11-8 show functional timing of the local memory interface. Several timing features are common to the memory read and write cycles in Figure 11-3 and Figure 11-4, and to the shift-register-transfer cycles in Figure 11-6 and Figure 11-7. A row address is output on LAD0-LAD15 at the start of the cycle, and is valid before and after \overline{RAS} falls. A column address is then output on LAD0-LAD15. The column address is valid briefly before and after the falling edge of \overline{LAL} , but is not valid at the falling edge of \overline{CAS} . The column address is clocked into an external transparent latch (such as a 74AS373 octal latch) on the falling edge of \overline{LAL} to provide the hold time on the column address required for DRAMs and VRAMs. A transparent latch is required so that the row address is available at the outputs of the latch during the start of the cycle.

11.4.1 Local Memory Write Cycle Timing

Figure 11-3 illustrates a memory write cycle. Data are output on LAD0-LAD15 following the latching of the column address. $\overline{\text{DEN}}$ goes active low at the same time the data become valid, and remains low as long as the data remain valid. In a large system that requires buffering of the data bus to memory, $\overline{\text{DEN}}$ is typically used as the enable signal to an external bidirectional buffer (such as a 74AS245 octal buffer). DDOUT is used as the direction control signal to the buffer. The write strobe, $\overline{\text{W}}$, goes active low after the data have become valid and $\overline{\text{CAS}}$ is low. This is interpreted as a "late write" cycle by the DRAMs and VRAMs, which are prevented by the inactive-high $\overline{\text{TR}}/\overline{\text{OE}}$ signal from enabling their read drivers. Because the data are valid on both sides of the $\overline{\text{W}}$ write strobe, external devices can latch the data on either the high-to-low or low-to-high edge of $\overline{\text{W}}$.




Figure 11-3. Local Bus Write Cycle Timing

11.4.2 Local Memory Read Cycle Timing

Figure 11-4 illustrates a memory read cycle. LAD0-LAD15 are forced to high impedance following the latching of the column address. $\overline{\text{DEN}}$ and $\overline{\text{TR}}/\overline{\text{OE}}$ both go active low after $\overline{\text{CAS}}$ becomes low in order to enable read data from the memory to the LAD pins. $\overline{\text{TR}}/\overline{\text{OE}}$ enables the output drivers of the DRAMs and VRAMs. $\overline{\text{DEN}}$ enables the external bidirectional buffers needed with memories so large that external buffering (using a device such as a 74AS245 octal buffer) of the data bus is required. The DDOUT signal serves as the direction control for the external bidirectional buffers, and is low well in advance of the high-to-low transition of $\overline{\text{DEN}}$, and remains low well after the low-to-high transition of $\overline{\text{DEN}}$. The data that is read from memory must be valid during the middle of the Q4 clock phase, as indicated in Figure 11-4. The low-to-high transitions of $\overline{\text{TR}}/\overline{\text{OE}}$ and $\overline{\text{DEN}}$ occur well in advance of the time at which the LAD drivers turn on to output the row address of the next cycle. This prevents bus conflicts.




Figure 11-4. Local Bus Read Cycle Timing

11.4.3 Local Shift-Register-to-Memory Cycle Timing

A shift-register-to-memory cycle is a special type of cycle used in systems with VRAMs. The cycle transfers the contents of the VRAM's internal shift register to a selected row of its internal memory array. The cycle typically affects all VRAMs in the system.

During the shift-register-to-memory cycle shown in Figure 11-5, both $\overline{TR}/\overline{OE}$ and \overline{W} are low during the fall of \overline{RAS} . VRAMs will recognize this timing as the beginning of a shift-register-to-memory cycle. Conventional DRAMs may need to be de-selected (by withholding the row or column address strobe, for example) to prevent them from interpreting the cycle as a conventional read cycle. Alternately, the output enable signal required by a DRAM such as the TMS4464 can be synthesized by connecting \overline{DEN} and \overline{DDOUT} to the inputs of a two-input OR gate. (In fact, any pair of the signals \overline{DEN} , \overline{DDOUT} , and $\overline{TR}/\overline{OE}$ will work.) The low-to-high transition of $\overline{TR}/\overline{OE}$ occurs after the fall of \overline{CAS} but prior to the rising edge of \overline{RAS} . This timing provides compatibility with a variety of VRAMs.

The GSP performs a shift-register-to-memory cycle when writing to a pixel while the DPYCTL register's SRT bit is set to 1. For example, the instruction `PIXT A0,*A1` writes the pixel in A0 to the address pointed to by A1. The PSIZE register should contain the value 16 so that the write cycle is not preceded by a read cycle. When SRT is set to 1, this write is converted to the shift-register-to-memory cycle shown in Figure 11-5. The row address is selected from bits 12-26 of A1, which are output on LAD0-LAD14 during the cycle.




Figure 11-5. Local Bus Shift Register to Memory Cycle Timing

11.4.4 Local Memory-to-Shift-Register Cycle Timing

A memory-to-shift-register cycle is a special type of cycle used in systems with VRAMs. The cycle transfers the contents of a selected row of a video RAM's memory array to its internal shift register.

VRAM memory-to-shift-register cycles are primarily used to refresh the screen of a CRT monitor. The cycles referred to elsewhere in this document as *screen-refresh cycles* are in fact memory-to-shift-register cycles. The GSP also performs a memory-to-shift-register cycle when reading a pixel (for example, by executing a PIXT *A0,A1 instruction) while the SRT bit of the DPYCTL register is set to 1.

During the memory-to-shift-register cycle shown in Figure 11-6, $\overline{TR}/\overline{QE}$ is low during the fall of \overline{RAS} , but \overline{W} remains high. VRAMs will recognize this timing as the beginning of a memory-to-shift-register cycle, and their data outputs will remain in high impedance. Conventional DRAMs may need to be deselected to prevent them from interpreting the cycle as a memory read cycle. Alternately, the output enable signal required by a DRAM such as the TMS4464 can be synthesized by connecting \overline{DEN} and \overline{DDOUT} to the inputs of a two-input OR gate. The low-to-high transition of $\overline{TR}/\overline{QE}$ occurs after the fall of \overline{CAS} but prior to the rising edge of \overline{RAS} . This timing provides compatibility with a variety of VRAMs.




Figure 11-6. Local Bus Memory to Shift Register Cycle Timing

11.4.5 Local Memory $\overline{\text{RAS}}$ -Only DRAM Refresh Cycle Timing

During the $\overline{\text{RAS}}$ -only DRAM refresh cycle shown in Figure 11-7, $\overline{\text{RAS}}$ and $\overline{\text{LAL}}$ are the only active control signals. All other signals, including $\overline{\text{CAS}}$, $\overline{\text{W}}$, and $\overline{\text{TR}}/\overline{\text{OE}}$, remain inactive high through the cycle. The row address, output on the LAD pins during the high-to-low transition of $\overline{\text{RAS}}$, is generated by the REFCNT (DRAM-refresh counter) register.




Figure 11-7. Local Bus $\overline{\text{RAS}}$ -Only DRAM-Refresh Cycle Timing

11.4.6 Local Memory $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ DRAM Refresh Cycle Timing

During the $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ DRAM-refresh cycle shown in Figure 11-8, $\overline{\text{CAS}}$ goes low before $\overline{\text{RAS}}$ goes low. Certain types of DRAMs and VRAMs recognize this as the beginning of a DRAM-refresh cycle in which the address of the row to be refreshed is generated by a counter on the RAM chip itself, rather than by an external device such as the GSP. The row address output by the GSP during the cycle is the same as would be output if the GSP were configured to perform a $\overline{\text{RAS}}$ -only refresh cycle rather than a $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ cycle. The address bits output on LAD0-LAD13 remain stable from the start of the row address time (start of Q2) to the end of the column address time (end of Q4). During row address time, LAD14 will be the same value as LAD6. LAD15, on which the RF bus status bit is output, will be low during the row address times. LAD14 and LAD15 are both high during column address time. In contrast to other types of cycles in which $\overline{\text{RAS}}$ goes low, the LAL output goes low at the start of Q3, after the fall of $\overline{\text{CAS}}$ and before the fall of $\overline{\text{RAS}}$. The timing of $\overline{\text{LAL}}$ is designed to support the use of decode circuitry which latches the state of selected address/data pins during the fall of $\overline{\text{LAL}}$, and which recognizes a $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ cycle by detecting a high level at the $\overline{\text{RAS}}$ output during the fall of $\overline{\text{LAL}}$.




Figure 11-8. Local Bus $\overline{\text{CAS}}$ -Before- $\overline{\text{RAS}}$ DRAM-Refresh Cycle Timing

11.4.7 Local Memory Internal Cycles

When the GSP is not performing one of the memory operations shown in Figure 11-3 through Figure 11-8, its memory interface control signals remain inactive, as shown in Figure 11-9. This is called an internal cycle. Figure 11-9 shows two sequential internal cycles. During internal cycles, the LRDY input is ignored.




Figure 11-9. Local Bus Internal Cycles Back to Back

11.4.8 I/O Register Access Cycles

A special memory read or write cycle is performed when the GSP addresses an on-chip I/O register. During this cycle, the external RAS signal falls, but the external CAS remains inactive high for the duration of the cycle. I/O register locations begin at address >C000 0000, and all 32 bits of the I/O register address are decoded internally. The two MSBs of the 32-bit logical address are not output at the LAD0-LAD15 pins.

Figure 11-10 shows an I/O register read cycle and Figure 11-11 shows an I/O register write cycle. These cycles occur when one of the TMS34010's on-chip I/O registers is accessed by the on-chip processor or by the host processor via a host-indirect access. An address in the range >C000 0000 to >C000 01FF is interpreted as an I/O register access by on-chip decode logic, and the read or write cycle is modified as shown in Figure 11-10 or Figure 11-11. The two MSBs of the internal address (bits 30 and 31) are available internally and are included in the internal decoding operation.

An I/O register read or write cycle is always two clock periods in duration, and LRDY is ignored. The control outputs that are active low during the cycle are RAS and LAL. The W, TR/OE, and DDOUT outputs all remain inactive high. The row and column addresses output at the LAD0-LAD15 pins are all 0s. All three bus status bits are inactive (RF is high, IAQ is low, and TR is high). During the read cycle shown in Figure 11-10, the LAD0-LAD15 pins are driven to high impedance during the data phase of the cycle. During the write cycle shown in Figure 11-11, the LAD0-LAD15 pins contain the valid data being written to the I/O register.




Figure 11-10. I/O Register Read Cycle Timing




Figure 11-11. I/O Register Write Cycle Timing

11.4.9 Read-Modify-Write Operations

The GSP's read-modify-write operation, which consists of separate read and write cycles, is not the same as the read-modify-write cycle specified for some DRAMs. As explained in Section 5, when inserting a field into memory that is not aligned to 16-bit word boundaries, the GSP memory interface logic may be required to perform read-modify-write (RMW) operations on one or more words in memory. A RMW operation consists of the following sequence of steps:

- 1) A word is read from memory.
- 2) The portion of that word corresponding to the field being inserted is loaded with the new value.
- 3) The modified word is written back to memory.

The read cycle is as shown in Figure 11-4 (page 11-8), and the write cycle is as shown in Figure 11-3 (page 11-7).

11.4.10 Local Memory Wait States

The timing shown in Figure 11-3 through Figure 11-8 assumes that the LRDY input remains high during the cycle. The LRDY pin is pulled low by slower memories requiring a longer cycle time. The GSP samples the LRDY input at the end of Q1, as indicated in the figures. If LRDY is low, the GSP inserts an additional state, called a *wait state*, into the cycle. Wait states continue to be inserted until LRDY is sampled at a high level. The cycle then completes in the manner indicated in Figure 11-3 through Figure 11-8.

The LRDY input is ignored by the GSP during internal cycles, as indicated in Figure 11-9.

Figure 11-12 shows an example of a read cycle extended by one wait state. The first time LRDY signal is sampled, a low level is detected by the GSP, causing the cycle to be delayed by a wait state. When LRDY is sampled again one local clock period later, a high level is detected, permitting the cycle to complete. The time during which $\overline{\text{RAS}}$, $\overline{\text{CAS}}$, $\overline{\text{LAL}}$, $\overline{\text{TR/OE}}$, $\overline{\text{DEN}}$, and $\overline{\text{DDOUT}}$ remain low is extended by one state (one local bus clock period).




Figure 11-12. Local Bus Read Cycle with One Wait State

Local Memory Interface Bus - Local Memory Interface Timing

Figure 11-13 is an example of a write cycle extended by one wait state. The first time LRDY signal is sampled, a low level is detected by the GSP, causing the cycle to be delayed by a wait state. When LRDY is sampled again one local clock period later, a high level is detected, permitting the cycle to complete. The time during which $\overline{\text{RAS}}$, $\overline{\text{CAS}}$, $\overline{\text{LAL}}$, $\overline{\text{W}}$ and $\overline{\text{DEN}}$ remain low is extended by one state.




Figure 11-13. Local Bus Write Cycle with One Wait State

Figure 11-14 (page 11-18) is an example of a shift-register-to-memory cycle extended by one wait state. The first time the LRDY signal is sampled, a low level is detected by the GSP, causing the cycle to be delayed by a wait state. When LRDY is sampled again one local clock period later, a high level is detected, permitting the cycle to complete. The time during which $\overline{\text{RAS}}$, $\overline{\text{CAS}}$, and $\overline{\text{LAL}}$ remain low is extended by one state. The $\overline{\text{W}}$ and $\overline{\text{TR/OE}}$ low times are not extended, however. Similarly, during a memory-to-shift-register cycle, $\overline{\text{TR/OE}}$ is not extended.




Figure 11-14. Local Bus Shift-Register-to-Memory Cycle with One Wait State

11.4.11 Hold Interface Timing

The TMS34010 contains a hold interface through which external bus-master devices can request control of the local memory bus. When the GSP grants a hold request, it drives its local memory address/data bus and control outputs to high impedance, and the requesting device becomes the new bus master. When the requesting device no longer requires the bus, it removes its hold request, and the GSP again assumes control of the local bus.

Figure 11-15 shows the GSP releasing control of the local bus in response to a hold request. The GSP samples the state of the HOLD input at each LCLK2 rising edge (at the end of the Q1 phase of the clock). The state of the hold acknowledge signal (active or inactive) is output on the HLDA/EMUA pin during the Q3 and Q4 clock phases (LCLK1 low). During the first (or leftmost) LCLK2 rising edge, the hold request is inactive. Consequently, the hold acknowledge signal remains inactive during the first LCLK1 low phase. By the second LCLK2 rising edge, the hold request has been activated, and the GSP responds by acknowledging the hold request during the next LCLK1 low phase. The address/data lines and majority of the control lines are driven to high impedance at the start of the next Q2 phase (LCLK2 rising edge). The DDOUT and DEN pins are driven to high impedance a quarter clock later.

Figure 11-16 shows the GSP resuming control of the local bus after deactivation of the hold request. Again, the GSP samples the state of the $\overline{\text{HOLD}}$ input at each LCLK2 rising edge. During the first LCLK2 rising edge, the hold request is still active, and the GSP responds during the next LCLK1 low phase with an active hold acknowledge signal. By the second LCLK2 rising edge, the hold request has been removed. The GSP responds by outputting an inactive hold acknowledge signal during the next LCLK1 low phase. At the next LCLK2 rising edge, the GSP begins to drive its address/data pins and the majority of its control pins to logic-high or logic-low levels. The $\overline{\text{DEN}}$ and DDOUT signals remain in high impedance for one additional quarter clock before they too begin to be driven.




Figure 11-15. TMS34010 Releases Control of Local Bus

In Figure 11-15, the first active-low pulse of the $\overline{\text{HLDA}}/\overline{\text{EMUA}}$ output is an early acknowledgment, and the bus will not be released for another three quarters of a clock. The early acknowledgment gives advance warning to the device requesting the hold that the bus is about to be released by the GSP, allowing the device time to prepare to become the new bus master. The GSP outputs the active hold acknowledge signal only when it is prepared to release the bus within the next clock period. If the GSP must wait longer than this to release the bus, its hold acknowledgment will be withheld until it can release the bus.

For instance, if the LRDY signal in Figure 11-15 were low instead of high at the second rising edge of LCLK2, the GSP would be forced to wait, and would therefore not acknowledge the hold request until later, when the not-ready condition was removed. Also, if the hold request in Figure 11-15 was asserted initially during the first LCLK2 rising edge rather than the second, the GSP would delay its hold acknowledgment until the second LCLK1 low clock phase, knowing that the cycle in progress would not be completed until the third Q2 phase in the diagram.

A hold request has a higher priority than any internally generated memory cycle requests, including:

- Screen refresh
- DRAM refresh
- Indirect access by the host processor
- GSP instruction fetch or data access

A hold request will be delayed only to allow a memory cycle already in progress to complete.

External devices can activate or deactivate the $\overline{\text{HOLD}}$ input at any time, as long $\overline{\text{HOLD}}$ is at a valid logic level during each rising edge of LCLK2, and meets the required setup and hold times with respect to this edge. After the GSP grants the bus to an external device (via an active-low level on the $\overline{\text{HLDA}}/\overline{\text{EMUA}}$ output during the Q3 clock phase), it will continue to acknowledge the hold request during the Q3 phases of subsequent clock cycles. The external device will retain control of the bus until it deactivates its hold request.

External devices should avoid placing the GSP in hold for long periods. While the GSP is in hold, it can perform neither screen-refresh nor DRAM-refresh cycles. Furthermore, a host processor attempting to access the GSP's local memory through the host interface registers while the GSP is in hold may receive a not-ready signal. When this occurs, the host will be forced to wait to complete its access until the GSP leaves the hold state.

If a request for a DRAM-refresh or screen-refresh cycle is generated within the GSP while an external device controls the bus, the GSP will retain the request and perform the DRAM-refresh or screen-refresh cycle after the external device has returned control of the bus to the GSP. However, if a requested DRAM-refresh cycle is prevented from occurring for so long that a second DRAM-refresh cycle is requested before the first DRAM-refresh cycle can occur, the first DRAM-refresh request will be lost. Similarly, if a screen-refresh request is prevented from occurring for so long that a second screen-refresh cycle is requested before the first screen-refresh cycle can occur, the first screen-refresh request will be lost.

The $\overline{\text{HLDA}}/\overline{\text{EMUA}}$ output is multiplexed between the hold acknowledge ($\overline{\text{HLDA}}$) and emulate acknowledge ($\overline{\text{EMUA}}$) signals. The $\overline{\text{HLDA}}$ signal is output during the LCLK1 low phase, and the $\overline{\text{EMUA}}$ signal is output during the LCLK1 high phase.

Local Memory Interface Bus - Local Memory Interface Timing




Figure 11-16. TMS34010 Resumes Control of Local Bus

11.4.12 Local Bus Timing Following Reset

Figure 11-17 shows the timing of the local bus signals following reset. At the end of reset, the TMS34010 automatically performs a series of eight $\overline{\text{RAS}}$ -only refresh cycles, as required to initialize certain DRAMs (such as the TMS4256 and TMS4464) and VRAMs (such as the TMS4461) following power-up. The asynchronous low-to-high transition of $\overline{\text{RESET}}$ is sampled at the second high-to-low LCLK1 transition in Figure 11-17. In less than two local clock periods following this LCLK1 transition, the first of the eight $\overline{\text{RAS}}$ -only cycles begins, as shown at the right side of Figure 11-17.

Each of the eight $\overline{\text{RAS}}$ cycles following reset is two clock periods in duration, but can be extended by a not-ready signal (LRDY low). The timing for each cycle is identical to that of a $\overline{\text{RAS}}$ -only DRAM-refresh cycle, including the bus status codes output during the row and column address times. The row address for each of the eight $\overline{\text{RAS}}$ -only cycles is all 0s.




Figure 11-17. Local Bus Timing Following Reset

11.5 Addressing Mechanisms

The GSP addresses memory by means of a 32-bit logical address. As explained in Section 3, each 32-bit logical address points to a bit in memory.

Logical address bits are numbered from 0 to 31, where bit 0 is the LSB and bit 31 is the MSB. Figure 11-18 illustrates the manner in which address bits 4-29 are output to physical memory. Each column in the figure indicates an address/data bus pin, LAD0-LAD15, and below it is the corresponding bit of the logical address output at the LAD pin during the fall of $\overline{\text{RAS}}$ and during the fall of $\overline{\text{CAS}}$. Bus status bits $\overline{\text{RF}}$, $\overline{\text{TR}}$ and IAQ are output on LAD14-LAD15.

		LAD Pin Numbers															
		15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
GSP Logical Address Bits†	At Fall of $\overline{\text{RAS}}$	$\overline{\text{RF}}$	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12
	At Fall of $\overline{\text{CAS}}$	IAQ	$\overline{\text{TR}}$	29	28	27	14	13	12	11	10	9	8	7	6	5	4

† Bus status signals:
 $\overline{\text{RF}}$ - DRAM refresh cycle
 IAQ - Instruction acquisition cycle
 $\overline{\text{TR}}$ - Shift-register-transfer cycle

Figure 11-18. External Address Format

Key features of the local bus addressing mechanism include the following:

- The two MSBs of the 32-bit logical address (bits 30 and 31) are not output.
- The four LSBs of the 32-bit logical address (bits 0 to 3) are not output, but are used internally to designate a bit boundary within a 16-bit word accessed in the external memory.
- The address bits output on LAD0-LAD10 during the falling edges of $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ are aligned so that 16 consecutive bits from the logical address are available at any eight consecutive pins in the range LAD0 to LAD10. The address bits are output in this way in order that the 8-bit row address and 8-bit column address presented to the dynamic RAMs can always be taken from the same eight address/data pins. This eliminates the need for external address multiplexers.
- Logical address bits 12-14 are output twice during a memory cycle - during both the $\overline{\text{RAS}}$ and $\overline{\text{CAS}}$ falling edges - but at different pins. This allows a variety of memory organizations and decoding schemes to be used, as will be explained shortly.

Pins LAD0-LAD10 form an 11-bit zone in which logical address bits 12-14 are overlapped (that is, they are issued in both cycles, but on different pins). The row and column address bus is connected to any eight consecutive pins within this zone. The actual position is determined by the bank-decoding scheme selected for a particular memory organization.

Output along with the address are three bus status signals:

- The \overline{RF} (DRAM refresh) bit is output on LAD15 during the fall of \overline{RAS} . It is low if the cycle that is just beginning is a DRAM-refresh cycle (either \overline{RAS} -only or \overline{CAS} -before- \overline{RAS}); otherwise, \overline{RF} is high.
- The \overline{TR} (VRAM shift-register-transfer) bit is output on LAD14 during the fall of \overline{CAS} , and is low if the cycle in progress is a video RAM shift-register transfer. Otherwise, \overline{TR} is high. In either event, the state of the \overline{TR} bit reflects the state of the $\overline{TR}/\overline{OE}$ output during the falling edge of \overline{RAS} within the same cycle.
- The IAQ bit is output on LAD15 during the fall of \overline{CAS} , and is high if the cycle is an instruction fetch; otherwise, IAQ remains low. The term *instruction fetch* includes not only reads of opcodes, but also immediate data, immediate addresses, and so on. The instruction cache may or may not be disabled.

IAQ is active high when words are fetched from memory to load the instruction cache. A block (or subsegment) of words is fetched in a series of read cycles, during which IAQ is active high. The PC points to an instruction word within the block, but the block may contain data as well as instruction words (opcodes, immediate addresses, immediate data, and so on). Only during execution will the GSP distinguish instruction words from data words residing in the cache. Instruction words will be fetched from the cache as they are needed, but data inadvertently loaded into the cache will be ignored and all memory data reads or writes will result in accesses of the memory rather than the cache.

When the cache is disabled, IAQ is active high only during a cycle in which an instruction word (a word pointed to by the PC) is fetched.

11.5.1 Display Memory Hardware Requirements

The minimum number of bits of memory required to implement the display memory is the product of the total number of pixels (on-screen and off-screen areas combined) and the number of bits per pixel. The minimum number of VRAMs required to contain the display memory is calculated as follows:

$$\text{Number of VRAMs} = \frac{(\text{pixels per line}) \times (\text{lines per frame}) \times (\text{bits per pixel})}{\text{Number of bits per VRAM}}$$

This calculation yields the minimum number of VRAMs needed, but additional VRAMs may be required in some applications. For instance, XY addressing can be supported by making the number of pixels per line of the display memory a power of two, but this may require more than the minimum number of VRAMs needed to contain the display.

11.5.2 Memory Organization and Bank Selecting

During a single local memory cycle, one data word (16 bits) is transferred between the GSP and the selected bank of memory. The memory is partitioned into a number of banks, where each bank contains the number of memory devices that can be accessed in a single memory cycle. The number of devices per bank is therefore determined by dividing the width of the data bus by the number of data pins per device. The GSP data bus is 16 bits wide, and can access 16 memory data pins during a single cycle. This means, for example, that a bank composed of 64K-by-1 RAMs contains 16 RAM devices. A bank composed of 64K-by-4 RAMs contains 4 RAM devices.

In a typical system, the local memory is divided into two parts, one consisting of the display memory and the other consisting of additional DRAMs needed to store programs and data. This additional RAM can be called the *system* memory. A high-order address bit is typically used to select between the display memory and system memory. Within the display memory or system memory, some address bits are provided as the row and column addresses to the selected bank, while other address bits are used to select one of the banks.

The number of banks of VRAM needed for the display memory is calculated by dividing the total number of VRAMs by the number of VRAMs per bank. This in turn determines how many bank select bits must be decoded.

11.5.3 Dynamic RAM Refresh Addresses

DRAMs (and VRAMs) require periodic refreshing to retain their data. The GSP automatically generates DRAM-refresh cycles at regular intervals. The interval between refresh cycles is programmable, and DRAM refreshing can be disabled in systems that do not require it.

The GSP can be configured to generate one of two types of DRAM-refresh cycle timing: $\overline{\text{RAS}}$ -only or $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$. Figure 11-7 shows $\overline{\text{RAS}}$ -only timing, and Figure 11-8 shows $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ timing. During a $\overline{\text{RAS}}$ -only refresh cycle, the GSP provides the 8-bit row address needed to refresh a particular row within each of the DRAMs in the memory system. DRAMs that support $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ cycles each contain an on-chip counter which generates the row address needed during the cycle. In other words, these devices do not rely on the GSP to provide the row address during the $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ cycle.

The row address output by the GSP during a DRAM-refresh cycle is the same regardless of whether the GSP is configured for $\overline{\text{RAS}}$ -only or $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ refresh timing. The fact that the GSP outputs a valid row address during a $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ cycle makes possible a hybrid system in which some DRAMs use $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ refresh timing while others use $\overline{\text{RAS}}$ -only timing. This hybrid approach configures the GSP to perform $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ refresh, and relies on external decode logic to prevent the active-low column address strobe from reaching those DRAMs that require $\overline{\text{RAS}}$ -only refreshing. The decode logic detects the fact that $\overline{\text{CAS}}$ falls before $\overline{\text{RAS}}$ during a $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ cycle, and uses this to inhibit transmitting the $\overline{\text{CAS}}$ signal to the $\overline{\text{RAS}}$ -only DRAMs.

Local Memory Interface Bus - Addressing Mechanisms

Several bits in the CONTROL register determine the manner in which the GSP performs DRAM refreshing. The RM bit selects the type of DRAM-refresh cycle:

- RM=0 selects $\overline{\text{RAS}}$ -only cycles
- RM=1 selects $\overline{\text{CAS}}$ -before- $\overline{\text{RAS}}$ cycles

The RR bits determine the interval between DRAM-refresh cycles:

- RR=00 selects refreshing every 32 local clock periods
- RR=01 selects refreshing every 64 local clock periods
- RR=10 is a reserved code
- RR=11 inhibits DRAM refreshing

At reset, internal logic forces the RM bit to 0 and the RR field to 00. While the $\overline{\text{RESET}}$ signal to the GSP is active, no DRAM-refresh cycles are performed. Following reset, the GSP begins to automatically perform DRAM-refresh cycles at regular intervals.

Both the interval between DRAM-refresh cycles and the addresses output during the cycles are generated within the REFCNT (DRAM-refresh count) register. Bits 2-15 of REFCNT form a continuous binary counter. The RINTVL field occupies bits 2-7, and counts the length of the interval between successive internal requests for DRAM-refresh cycles. The eight MSBs of REFCNT form the ROWADR field, containing the row address output to memory during the DRAM-refresh cycle.




Figure 11-19. Row Address for DRAM-Refresh Cycle

During a DRAM-refresh cycle, the 8-bit row address in the ROWADR field of the REFCNT register is output on the LAD pins during the high-to-low transition of $\overline{\text{RAS}}$. As shown in Figure 11-19, the eight bits of ROWADR are output on LAD0-LAD7. The seven LSBs of ROWADR are also output on LAD8-LAD14. LAD15 transmits the $\overline{\text{RF}}$ bus status signal, low during the fall of $\overline{\text{RAS}}$.

Assume that LAD2-LAD9 are used as the 8-bit row address by a bank of DRAMs, as indicated in Figure 11-19. The address bits output on LAD2-LAD9 are the same eight bits output on LAD0-LAD7, but in a different order. During a series of 256 DRAM-refresh cycles, the row addresses output on LAD0-LAD7 and LAD2-LAD9 contain the same bits. Thus, if the addresses output on LAD0-LAD7 cycle through all 256 row addresses then the addresses output on LAD2-LAD9 will also cycle through all 256 row addresses, but in a different order.

11.5.4 An Example - Memory Organization and Decoding

As an example, consider a memory organization based on the address decoding scheme shown in Figure 11-20. Three logical address bits (4, 21, and 26) are used as bank-select bits. Logical address bits 5-12 are used as the 8-bit column address, and bits 13-20 are used as the 8-bit row address. Referring to Figure 11-18, the row and column addresses are multiplexed out on the same eight pins, LAD1-LAD8. The total number of address bits used to address external memory is 19, for a total address reach of one megabyte. The remaining address bits output by the GSP are not used for this example.




Figure 11-20. Address Decode for Example System

Bank select bit 2 (BS2) in Figure 11-20 selects between the display memory (BS2=0) and the system memory (BS2=1). System memory is a block of conventional DRAM used for program and data storage. BS2 becomes valid before $\overline{\text{RAS}}$ falls, and thus can be used to determine whether the row-address strobe is gated to the display memory or to the system memory. The average power dissipation is reduced because only one or the other (the display memory or the system memory) is enabled during a particular memory read or write cycle.

Figure 11-21 shows the structure of the display memory. Its dimensions are 1024 by 1024 at four bits per pixel. Bank select bit 1 (BS1) selects between the top (BS1=0) and bottom (BS1=1) halves of the display memory. Since BS1 becomes valid before the fall of $\overline{\text{RAS}}$, it can be used to gate $\overline{\text{RAS}}$ to either the upper or lower half of the display memory during a memory read or write

cycle. By transmitting the row address strobe to only half of the display memory, the power dissipation for the cycle is significantly reduced.

Bank select bit 0 (BS0) selects between the even word and odd word of each pair of adjacent words in the display memory. Each word contains four adjacent pixels. Odd and even words are stored in two separate banks of VRAMs, and the decode logic gates the column address strobe to the selected bank only. The row address strobe is gated to both banks (odd and even words). This increases the power dissipation over that required if only one bank were active. A compensating benefit of this organization, however, is that it reduces the rate at which each VRAM must supply serial data to refresh the screen. During screen refresh, the bank containing the even words and the bank containing the odd words alternately provide data to the video monitor. Alternating between the two banks in this fashion reduces the data bandwidth requirements of each bank to about 10 MHz, which is an eighth of the video bandwidth.




Figure 11-21. Display Memory Dimensions for the Example

The decode logic must be capable of more than just selecting a particular bank of the display memory or system memory during a memory read or write cycle. It must also be capable of enabling all DRAMs and VRAMs during a DRAM-refresh cycle, and enabling all VRAMs during a screen-refresh (memory-to-shift-register) cycle. This means that the decode logic must distinguish DRAM-refresh and screen-refresh cycles from memory access cycles, and during a refresh cycle broadcast the row and column address strobes to all devices that require them. The timing of the \overline{RF} and \overline{TR} bus status bits has been designed to make these signals convenient for the design of the decode logic.

During a read or write cycle, the value of BS2, output with the row address, determines whether \overline{RAS} is gated to the display memory or to system memory. During a DRAM-refresh cycle, the decode logic must broadcast the row-address strobe to all dynamic RAMs (including the VRAMs). The decode logic must be able to determine prior to the fall of the row address strobe whether the cycle that is beginning is a DRAM-refresh cycle, or a memory read or write cycle. This is the reason the GSP outputs the \overline{RF} bus status signal prior to the fall of \overline{RAS} .

The decode logic uses the value of BS1 to determine whether the top or bottom half of the display memory receives an active row-address strobe during a memory read or write cycle. The same logic must also be capable of broadcasting \overline{RAS} to all VRAMs during either a DRAM-refresh cycle or a shift-register-transfer cycle. The decode logic therefore monitors the state of the GSP's $\overline{TR}/\overline{OE}$ output prior to the fall of \overline{RAS} . A low level on $\overline{TR}/\overline{OE}$ indicates that the cycle just beginning is a shift-register-transfer cycle, and that \overline{RAS} should be broadcast.

While the decode logic uses the value of BS0 to determine whether the even or odd word receives a column-address strobe during a read or write cycle involving the display memory, the same logic must be capable of broadcasting \overline{CAS} to all VRAMs during a screen-refresh cycle. Rather than require an external latch to capture the state of the $\overline{TR}/\overline{OE}$ during the fall of \overline{RAS} , the GSP outputs the same information a second time in the form of the \overline{TR} bus status signal, which is valid prior to and during the fall of \overline{CAS} .