

9. Screen Refresh and Video Timing

The TMS34010 generates the synchronization and blanking signals used to drive a video screen in a graphics system. The GSP can be programmed to support a variety of screen resolutions and interlaced or noninterlaced video. If desired, the GSP can be programmed to synchronize to externally generated video signals. The GSP also supports the use of video RAMs by generating the memory-to-shift-register cycles necessary to refresh a screen.

This section includes the following topics:

Section	Page
9.1 Video Timing Signals	9-2
9.2 Screen Sizes	9-3
9.3 Video Timing Registers	9-4
9.4 Horizontal Video Timing	9-6
9.5 Vertical Video Timing	9-8
9.6 Display Interrupt	9-14
9.7 Dot Rate	9-15
9.8 External Sync Mode	9-16
9.9 Video RAM Control	9-19

9.1 Video Timing Signals

The TMS34010 generates horizontal sync, vertical sync, and blanking signals ($\overline{\text{HSYNC}}$, $\overline{\text{VSYNC}}$, and $\overline{\text{BLANK}}$) on chip. The GSP's video timing logic is driven by the video input clock (VCLK). The sync and blanking signals control the horizontal and vertical sweep rates of the screen and synchronize the screen display to data output by the VRAMs.

$\overline{\text{HSYNC}}$ is the horizontal sync signal used to control external video circuitry. It may be configured as an input or an output via the DXV and HSD bits in the DPYCTL register. When DXV=0 and HSD=0, external video is selected and $\overline{\text{HSYNC}}$ is an input. Otherwise, $\overline{\text{HSYNC}}$ is an output.

$\overline{\text{VSYNC}}$ is the vertical sync signal used to control external video circuitry. It may be configured as an input or an output via the DXV bit in the DPYCTL register. If DXV=1, internal video is selected and $\overline{\text{VSYNC}}$ is an output. If DXV=0, external video is selected and $\overline{\text{VSYNC}}$ is an input.

$\overline{\text{BLANK}}$ is used to turn off a CRT's electron beam during horizontal and vertical retrace intervals. The signal output at the $\overline{\text{BLANK}}$ pin is a composite of the internally generated horizontal and vertical blanking signals. $\overline{\text{BLANK}}$ can also be used to control starting and stopping of the VRAM shift registers.

VCLK is derived from the dot clock of the external video system. VCLK drives the internal video timing logic.

9.2 Screen Sizes

The TMS34010's 26-bit word address provides direct addressing of up to 128 megabytes of external memory. This address reach supports very high-resolution displays. For example, the designer of a large TMS34010-based system could decide to use the lower half of the address space for display memory, and use the upper half for storing programs and data. Half of this memory space, for example, could be used as a display memory, and the remaining memory can be used for programs and data. The 64-megabyte display memory in this example could support the following display sizes:

- 8192 by 4096 pixels at 16 bits per pixel
- 8192 by 8192 pixels at 8 bits per pixel
- 16,384 by 8192 pixels at 4 bits per pixel
- 16,384 by 16,384 pixels at 2 bits per pixel
- 32,768 by 16,384 pixels at 1 bit per pixel

The video timing registers also support high-resolution displays. The 16-bit vertical counter register, VCOUNT, directly supports screen lengths of up to 65,536 lines. The 16-bit horizontal counter register, HCOUNT, does not directly limit the horizontal resolution. Each horizontal line can be programmed to be up to 65,536 VCLK (video clock) periods long. The VCLK period, however, is an arbitrary number of dot-clock periods in length, depending on the external divide-down logic used to produce the VCLK signal from the dot clock. Thus, the number of pixels per line supported by the GSP horizontal timing registers is limited only by the amount of video memory that is present.

9.3 Video Timing Registers

The video timing registers are a subset of the I/O registers described in Section 6. The values in the video timing registers control the video timing signals. These registers are divided into two groups:

- **Horizontal timing registers** control the timing of the $\overline{\text{HSYNC}}$ signal and the internal horizontal blanking signal.
 - HCOUNT** counts the number of VCLK periods per horizontal scan line.
 - HESYNC** specifies the point in a horizontal scan line at which the $\overline{\text{HSYNC}}$ signal ends.
 - HEBLNK** specifies the endpoint of the horizontal blanking interval.
 - HSBLNK** specifies the starting point of the horizontal blanking interval.
 - HTOTAL** defines the number of VCLK periods allowed per horizontal scan line.
- **Vertical timing registers** control the timing of the $\overline{\text{VSYNC}}$ signal and the internal vertical blanking signal.
 - VCOUNT** counts the horizontal scan lines in the screen display.
 - VESYNC** specifies the endpoint of the $\overline{\text{VSYNC}}$ signal.
 - VEBLNK** specifies the endpoint of the vertical blanking interval.
 - VSBLNK** specifies the starting point of the vertical blanking interval.
 - VTOTAL** specifies the value of VCOUNT at which $\overline{\text{VSYNC}}$ may begin.

Figure 9-1 illustrates the relationship between the horizontal and vertical timing signals in the construction of a two-dimensional raster display pattern. The vertical sync and blanking signals span an entire frame. The horizontal sync and blanking signals span a single horizontal scan line within the frame. $\overline{\text{HBLNK}}$ and $\overline{\text{VBLNK}}$ are the internal horizontal and vertical blanking signals that combine to form the $\overline{\text{BLANK}}$ signal output. The display is active (not blanked) only when $\overline{\text{HBLNK}}$ and $\overline{\text{VBLNK}}$ are both inactive high.




Figure 9-1. Horizontal and Vertical Timing Relationship

Horizontal front porch refers to the interval between the beginning of horizontal blanking and the beginning of the horizontal sync signal. Horizontal back porch is the interval between the end of the horizontal sync signal and the end of horizontal blanking.

Vertical front porch refers to the interval between the beginning of vertical blanking and the beginning of the vertical sync signal. Vertical back porch is the interval between the end of the vertical sync signal and the end of vertical blanking.

9.4 Horizontal Video Timing

The following discussion applies to internally generated video timing (the DXV and HSD bits in the DPYCTL register are set to 1 and 0, respectively). Horizontal timing signals are the same for interlaced and noninterlaced video.

The HESYNC, HEBLNK, HSBLNK, and HTOTAL registers control horizontal signal timing as shown in Figure 9-2. All horizontal timing parameters are specified as multiples of VCLK. The time between the start of two successive HSYNC pulses is specified by HTOTAL. HCOUNT counts from 0 to the value in HTOTAL and then repeats. The value in HTOTAL represents the number of VCLK periods, minus one, per horizontal scan line. The value in HESYNC represents the duration of the sync pulse, minus one. The values in HEBLNK and HSBLNK specify the beginning and end points of the horizontal blanking interval.




Figure 9-2. Horizontal Timing

Figure 9-3 shows the internal logic used to generate the horizontal timing signals. HCOUNT is incremented once each VCLK period (on the high-to-low transition) until it equals the value in HTOTAL. On the next VCLK period following HCOUNT=HTOTAL, HCOUNT is reset to 0, and begins counting again.

The limits of the horizontal sync pulse are defined by the values in HESYNC and HTOTAL. HSYNC is driven active low when HCOUNT=HTOTAL; it is driven inactive high when HCOUNT=HESYNC. After HCOUNT becomes equal to HTOTAL or HESYNC, there is a one-clock delay before the active/inactive transition takes place at the HSYNC pin.

The internal $\overline{\text{HBLNK}}$ signal is driven active low after HCOUNT=HSBLNK; it is driven inactive high after HCOUNT=HEBLNK. $\overline{\text{HBLNK}}$ is logically ORed (negative logic) with $\overline{\text{VBLNK}}$ to produce the $\overline{\text{BLANK}}$ signal; that is, $\overline{\text{BLANK}}$ goes low when either $\overline{\text{HBLNK}}$ or $\overline{\text{VBLNK}}$ is low. After HCOUNT becomes equal to HSBLNK or HEBLNK, there is a one-clock delay before the transition takes place at the $\overline{\text{BLANK}}$ pin.




Figure 9-3. Horizontal Timing Logic - Equivalent Circuit

Figure 9-4 illustrates horizontal signal generation. In this example, $HTOTAL=N$, $HSBLNK=N-2$, $HESYNC=2$, and $HEBLNK=4$. Signal transitions at the \overline{HSYNC} and \overline{BLANK} pins occur at high-to-low VCLK transitions. After HCOUNT becomes equal to HTOTAL, HSBLNK, HESYNC, or HEBLNK, there is a one-clock delay before the transition takes place at the \overline{HSYNC} or \overline{BLANK} pin.

When $HCOUNT=HSBLNK$ (shortly before the end of the horizontal scan), horizontal blanking begins. At this time, the DIP (display interrupt) bit in the INTPEND register will be set to 1 if $VCOUNT=DPYINT$. The next screen-refresh cycle may also occur at this time - the GSP can be programmed to refresh the screen after one, two, three, or four scan lines.




Figure 9-4. Example of Horizontal Signal Generation

9.5 Vertical Video Timing

The following discussion applies to internally generated video timing (the DXV bit in the DPYCTL register is set to 1).

The VESYNC, VEBLNK, VSBLNK, and VTOTAL registers control vertical signal timing as shown in Figure 9-5. All vertical timing parameters are specified as multiples of the horizontal sweep time H, where

$$H = (HTOTAL + 1) \times (VCLK \text{ period})$$

VTOTAL specifies the time interval between the start of two successive vertical sync pulses; this value is the number of H intervals, less one, in each vertical frame. VESYNC represents the duration of the VSYNC pulse, less one, in each vertical frame. VSYNC's high-to-low and low-to-high transitions coincide with high-to-low transitions at the HSYNC pin.

VSBLNK and VEBLNK specify the starting and ending points of vertical blanking. Blanking begins when VTOTAL=VSBLNK and ends when VTOTAL=VEBLNK. Assuming that horizontal blanking is active at the start of each HSYNC pulse, transitions of the internal vertical blanking signal, VBLNK, occur while horizontal blanking is active.




Figure 9-5. Vertical Timing for Noninterlaced Display

Figure 9-6 shows the internal logic that generates the vertical timing signals. VCOUNT increments at the beginning of each HSYNC pulse until it equals the value in VTOTAL. When VCOUNT=VTOTAL, VCOUNT is reset to 0 and begins counting again. VSYNC is driven active low after VCOUNT=VTOTAL; it is driven inactive high after VCOUNT=VESYNC. The internal VBLNK signal is driven active low after VCOUNT=VSBLNK; it is driven inactive high after VCOUNT=VEBLNK. VBLNK is logically ORed (negative logic) with HBLNK to produce the BLANK signal. This description applies to a noninterlaced display. The vertical timing changes slightly for an interlaced display.




Figure 9-6. Vertical Timing Logic - Equivalent Circuit

9.5.1 Noninterlaced Video Timing

Noninterlaced scan mode is selected by setting the NIL bit in the DPYCTL register to 1. In this mode, each video frame consists of a single vertical field. Figure 9-7 shows the path traced by the electron beam on the screen. Box A shows the vertical retrace, which is an integral number of horizontal scan lines in duration. Box B shows the active portion of the frame. Solid lines represent lines that are displayed; dashed lines are blanked.




Figure 9-7. Electron Beam Pattern for Noninterlaced Video

Screen Refresh and Video Timing - Vertical Video Timing

Figure 9-8 illustrates the video timing signals that generate the display. In this example, $VSBLNK=8$, $VTOTAL=9$, $VESYNC=1$, and $VEBLNK=2$. (In actual applications, much larger values are used; these values were chosen for illustration only.) Each horizontal scan line is preceded by a horizontal retrace. The horizontal scan pattern repeats until $VCOUNT=VTOTAL$; $VCOUNT$ is then reset to 0, and vertical retrace returns the beam to the top of the screen. $BLANK$ is active low during both horizontal and vertical retrace intervals.

$VCOUNT$ is incremented each time $HCOUNT$ is reset to 0 at the end of a scan line. The \overline{VSYNC} output begins when $VCOUNT=VTOTAL$, coinciding with the start of \overline{HSYNC} . The \overline{VSYNC} output ends when $VCOUNT=VESYNC$; this also coincides with the start of an \overline{HSYNC} pulse.

The starting screen-refresh address is loaded from $DPYSTRT$ into $DPYADR$ at the end of the last active horizontal scan line preceding vertical retrace. This load is triggered when $HCOUNT=HSBLNK$ and $VCOUNT=VSBLNK$.




Figure 9-8. Noninterlaced Video Timing Waveform Example

9.5.1.1 Interlaced Video Timing

Interlaced scan mode is selected when the NIL bit in the DPYCTL register is set to 0. In this mode, each display frame is composed of two fields of horizontal scan lines. The display consists of alternate lines from the two fields. This doubles the display resolution while only slightly increasing the frequency with which data is supplied to the screen.

Figure 9-9 illustrates the path traced by the electron beam on the screen. Figure 9-10 shows the timing waveforms used to generate the display in Figure 9-9. In this example, VSBLNK=6, VTOTAL=7, VESYNC=1, and VEBLNK=2. (In actual applications, much larger values are used; these values were chosen for illustration only.)

In interlaced mode, two separate vertical scans are performed for each frame – one for the even line numbers (even field) and one for the odd line numbers (odd field). The even field is scanned first, starting at the top left of the screen (see Figure 9-9 *b*). When VCOUNT=VTOTAL, the vertical retrace returns the beam to the top of the screen, and the odd field is scanned (Figure 9-9 *d*). The electron beam starts scanning the odd and even fields at different points. The reason for this is illustrated in Figure 9-10. The end of the VSYNC pulse that precedes the even field coincides with start of an HSYNC pulse; however, the VSYNC pulse that precedes the odd field ends exactly halfway between two HSYNC pulses




Figure 9-9. Electron Beam Pattern for Interlaced Video

In interlaced mode, video timing logic operation is altered so that the odd field begins when $HCOUNT = HTOTAL/2$. The beam is thus positioned so that horizontal scan lines in the odd field fall between horizontal scan lines in the even field. To place each line of the odd field precisely between two lines of the even field, load $HTOTAL$ with an odd number.

The transition from *d* to *a* in Figure 9-9 shows that the vertical retrace at the end of the odd field begins at the end of a horizontal scan line; that is, it coincides with the start of an \overline{HSYNC} pulse, which results from the condition $HCOUNT = HTOTAL$. The \overline{VSYNC} pulse duration is an integral number of horizontal scan retrace intervals. When vertical retrace ends and the active portion of the next even field begins, the beam is positioned at the beginning of a horizontal scan line.

Horizontal timing is similar for interlaced and noninterlaced displays. $HCOUNT$ is reset to 0 at the end of each horizontal scan line. A screen-refresh cycle begins before the end of the line, coinciding with the start of the horizontal blanking interval. Assuming that the starting corner of the display is the upper left corner, the $DUPDATE$ field of the $DPYCTL$ register is added to the screen-refresh address ($SRFADR$ in the $DPYADR$ register) to generate the row address for the next screen-refresh cycle. In interlaced mode, the $DUPDATE$ value must be twice that of the value needed to produce the same display in noninterlaced mode (that is, two times the difference in addresses between consecutive scan lines). This causes the screen refresh to skip alternate lines during the odd and even fields.

At the beginning of each vertical blanking interval, the screen-refresh address ($SRFADR$ in the $DPYADR$ register) is loaded with the starting value specified by the $DPYSTRT$ register. When vertical blanking precedes an even field, the new $DPYADR$ row address is incremented by half the value in the $DUPDATE$ field. This is in preparation to display line 2 (Figure 9-9 *b*). When vertical blanking precedes an odd field, the row address loaded into $DPYADR$ from $DPYSTRT$ is not incremented. In this case, the starting row address in $DPYSTRT$ points to the beginning of line 1 (Figure 9-9 *d*).

Screen Refresh and Video Timing - Vertical Video Timing




Figure 9-10. Interlaced Video Timing Waveform Example

9.6 Display Interrupt

The TMS34010 can be programmed to interrupt the display when a specified line is displayed on the screen. This is called a display interrupt. It is enabled by setting the DIE bit in the INTENB register to 1 and loading the DPYINT register with the desired horizontal scan line number. When VCOUNT = DPYINT, the interrupt request is generated to coincide with the start of horizontal blanking at the end of the specified line.

The display interrupt request can be polled by disabling the interrupt (setting DIE=0) and checking the value of the DIP bit in the INTPEND register. Writing a 0 to DIP clears the request.

The display interrupt has several applications. It can be used to coordinate modifications of the bit map with the display of the bit map's contents, for example. While the bottom half of the screen is displayed, the GSP can modify the bit map of the top half of the screen, and vice versa.

The display interrupt is also useful in split screen applications. By modifying the contents of the DPYADR register halfway through a frame, different parts of the bit map can be displayed on the top and bottom halves of the screen. No special steps are necessary to ensure that loading a new value to DPYADR will not interfere with an ongoing screen-refresh cycle. The display interrupt is requested at the beginning of the horizontal blanking interval. If a screen-refresh cycle occurs during the same horizontal blanking interval, the GSP cannot respond to the interrupt request until the refresh cycle and subsequent updating of DPYADR are complete. This is true whether the interrupt is taken or the GSP simply polls the DIP bit and detects a 0-to-1 transition. After DIP has been set to 1, DPYADR can be loaded with a new value to achieve the split screen anytime before the next screen-refresh cycle.

9.7 Dot Rate

A typical screen must be refreshed 60 times per second for a noninterlaced scan or 30 times per second for an interlaced scan. For a noninterlaced display, the dot period (time to refresh one pixel) is estimated as:

$$\text{Dot Period} = \frac{(0.8)(1/60 \text{ second})}{(\text{pixels/line}) \times (\text{lines/frame})}$$

For an interlaced display, the dot period is estimated as

$$\text{Dot Period} = \frac{(0.8)(1/30 \text{ second})}{(\text{pixels/line}) \times (\text{lines/frame})}$$

The 0.8 factor in the numerator accounts for the fact that the display is typically blanked for about 20% of the duration of each frame. This factor varies somewhat from monitor to monitor.

During each dot period, the complete information for one pixel must be obtained from the display memory (or frame buffer). Thus, the rate at which video data must be supplied from the display memory (which is usually the limiting factor for large systems) is a function of pixel size as well as screen dimensions.

9.8 External Sync Mode

External sync mode allows the TMS34010 to use horizontal and vertical sync signals from an external source. This permits graphics images generated by the GSP to be superimposed upon or mixed with images from external sources.

External sync mode is selected by setting the DXV and HSD bits in the DPYCTL register to 0. HSYNC and VSYNC are now configured as inputs. (Alternately, HSYNC can be configured as an output and VSYNC as an input by setting DXV=0 and HSD=1.) When an active-low sync pulse is input to one of these pins, the corresponding counter (HCOUNT or VCOUNT) is forced to all 0s. By forcing the counters to follow the external sync signals, the blanking intervals and screen-refresh cycles are also forced to follow the external video signals.

The HSYNC and VSYNC inputs are sampled on each VCLK rising edge. HCOUNT or VCOUNT will be cleared 2.5 clock periods (on a VCLK falling edge) following a high-to-low transition at the HSYNC or VSYNC pin, respectively. BLANK remains an output, but its timing is affected because the point at which HCOUNT and VCOUNT are cleared is controlled by the external sync signals. The 2.5-clock delay must be considered when selecting values for the HSBLNK and HEBLNK registers.

9.8.1 A Two-GSP System

One GSP can generate video timing for two GSPs. As Figure 9-11 shows, GSP #1 is configured for internal sync mode (DXV=1) and generates the sync timing. GSP #2 is configured for external sync mode (DXV=0 and HSD=0), and receives the HSYNC and VSYNC inputs from GSP #1. Assume that the video timing registers of the two devices are named as follows:

GSP #1	GSP#2
HCOUNT1	HCOUNT2
HESYNC1	HESYNC2
HSBLNK1	HSBLNK2
HEBLNK1	HEBLNK2
HTOTAL1	HTOTAL2
VCOUNT1	VCOUNT2
VESYNC1	VESYNC2
VSBLNK1	VSBLNK2
VEBLNK1	VEBLNK2
VTOTAL1	VTOTAL2

GSP #2's registers should be programmed in terms of the values in GSP #1's registers, as shown in Table 9-1. The BLANK signals from GSP #1 and GSP #2 are the same, and switch in unison on the same VCLK edges. When HCOUNT1 is cleared on a VCLK falling edge, HCOUNT2 is cleared three full VCLK periods later. For short horizontal blanking periods, HEBLNK2 may need to be loaded with a value that is less than zero. For example, assume that HSBLNK1=HTOTAL1-4 and HEBLNK1=1 (that is, the horizontal blanking interval is six VCLK periods). To ensure that GSP #2's horizontal blanking interval begins and ends at the same time as GSP #1's, GSP #2's registers must be loaded with values so that HSBLNK2=HTOTAL1-8 and HEBLNK2=HTOTAL1-2.




Figure 9-11. External Sync Timing - Two GSP Chips

The values in HTOTAL2 and VTOTAL2 must be large enough so that the conditions HCOUNT=HTOTAL and VCOUNT=VTOTAL do not cause HCOUNT and VCOUNT, respectively, to be cleared before the leading edges of the external horizontal and vertical sync pulses occur. In the example in Table 9-1, HTOTAL2 and VTOTAL2 are set to their maximum values. The value of HESYNC2 must be such that HCOUNT=HESYNC2 occurs between the end of an external HSYNC pulse and the beginning of the next external HSYNC pulse. The value of VESYNC2 must be such that VCOUNT=VESYNC2 occurs between the end of an external VSYNC pulse and the beginning of the next external VSYNC pulse.

Table 9-1. Programming GSP #2 For External Sync Mode

HEBLNK2	=	HEBLNK1 - 3
HSBLNK2	=	HSBLNK1 - 3
HTOTAL2	=	65535
HESYNC2	=	(HEBLNK2 + HSBLNK2)/2 †
VEBLNK2	=	VEBLNK1
VSBLNK2	=	VSBLNK1
VTOTAL2	=	65535
VESYNC2	=	(VEBLNK2 + VSBLNK2)/2 †

† Suggested value; see description in text.

Since the internal counter can only be resolved to the nearest VCLK edge, precise synchronization with an external video source can be achieved only when VCLK is harmonically related to the external horizontal sync signal. In general, however, the HSYNC and VSYNC inputs are allowed to change asynchronously with respect to VCLK, although the precise VCLK edge at which an external sync pulse is recognized can be guaranteed only if the setup and hold times specified for sync inputs are met.

9.8.2 External Interlaced Video

External sync mode can be used for both interlaced and noninterlaced displays. When locking onto external interlaced sync signals, the GSP discriminates between the odd and even fields of the external video signals based on whether its internal horizontal blanking is active at the time that the start of the external vertical sync pulse is detected. In Figure 9-10, for example, the even field begins at a point where $\overline{\text{HBLNK}}$ is active low, and the odd field begins while $\overline{\text{HBLNK}}$ is high.

In interlaced mode, the discrimination between the even and odd fields of an external video source is based on the value of HCOUNT at a point two VCLK periods past the rising VCLK edge at which the GSP detects the $\overline{\text{VSYNC}}$ input's high-to-low transition. If HCOUNT contains a value greater than the value in HEBLNK, but less than or equal to the value in HSBLNK, the GSP assumes that the vertical sync pulse precedes the start of an odd field. Otherwise, the next field is assumed to be even. Alternatively, the GSP can be placed in noninterlaced mode, even though the external sync signals it is locking onto are for an interlaced display. In this case, the GSP will simply cause identical display information to be output to the monitor during the odd and even fields.

9.9 Video RAM Control

The TMS34010 automatically schedules the VRAM (video RAM) memory-to-shift-register cycles needed to refresh a video monitor screen. These cycles are referred to as *screen-refresh* cycles.

In addition to automatic screen-refresh cycles, the GSP can be configured to perform memory-to-shift-register and shift-register-to-memory cycles under the explicit control of software executing on the GSP's internal processor. One of the primary uses for this capability is to facilitate nearly instantaneous clearing of the screen. The screen is cleared in 256 memory cycles or less by means of a technique referred to here as *bulk initialization* of the display memory.

9.9.1 Screen Refresh

A screen-refresh cycle loads the VRAM shift registers with a portion of the display memory corresponding to a scan line of the display. The internal requests for these cycles occur at regular intervals coinciding with the start of the horizontal blanking intervals defined by the video timing registers. When horizontal blanking ends, the contents of the shift registers are clocked out serially to drive the video inputs of a monitor. A screen-refresh cycle typically occurs prior to each active line of the display.

9.9.1.1 Display Memory

The *display memory* is the area of memory which holds the graphics image output to the video monitor. This memory is typically implemented with VRAMs. During a screen-refresh cycle, a portion of the display memory corresponding to one (or possibly more) scan lines of the display are loaded into the VRAM shift registers. Depending on the screen dimensions selected, not all portions of the display memory are necessarily output to the monitor.

The width of the display memory is referred to as the *screen pitch*, which is the difference in 32-bit memory addresses between two vertically-adjacent pixels on the screen. The screen pitch is also the difference in starting memory addresses of the video data for two consecutive scan lines. When XY addressing is used, the screen pitch must be a power of two to facilitate the conversion of XY addresses to memory addresses. The value loaded into the DUDATE field of the DPYCTL register represents the screen pitch, and is the amount by which the screen-refresh address is incremented (or decremented) following each screen-refresh cycle.

The portion of display memory that is actually output to the monitor is referred to as the *on-screen memory*. The starting location of the on-screen memory is specified by the SRFADR field in the DPYSTRT register.

The starting screen-refresh address is output during the screen-refresh cycle that occurs at the start of each frame. At the end of the screen-refresh cycle, the address is incremented to point to the area of memory containing the pixels for the second scan line. The process is repeated for each subsequent scan line of the frame.

A screen-refresh cycle typically affects all video RAMs in the system. A memory-to-shift-register cycle transfers data from a selected row of memory to the internal shift register of each VRAM. The data is then shifted out to refresh the display.

A screen-refresh cycle takes place during the horizontal blanking interval that precedes a scan line to be displayed. Typically, the shift registers containing the video data for the line are clocked only during the active portion of the scan line, that is, when the BLANK output is high. At higher dot rates, the pixel clock or dot clock used to shift video data to the monitor is run through a frequency divider to create the VCLK signal input to the GSP.

The 8-bit row address output during the screen-refresh cycle specifies the row in memory to be loaded into the shift register internal to the VRAM. The number of bits of video data transferred to the shift registers of all the VRAMs in the system during a single screen-refresh cycle is calculated by multiplying the number of VRAMs times the length of the shift register in each VRAM. For example, 64 TMS4161 (64K-by-1) VRAM devices are sufficient to contain the bit map for a 1024-by-1024-pixel display with four bits per pixel. The length of the shift register in each TMS4161 is 256 bits. Thus, in a single screen-refresh cycle, a total of 64 times 256, or 16,384, bits are loaded. This is enough data to refresh four complete scan lines of the display. In general, a single screen-refresh cycle performed during a horizontal blanking interval is sufficient to supply one or more complete scan lines worth of data to the video monitor screen.

9.9.1.2 Generation of Screen-Refresh Addresses

The DPYADR, DPYCTL, DPYSTRT, and DPYTAP registers are used to generate the addresses output during screen-refresh cycles. Figure 9-12 shows these four registers, and indicates the register fields which determine the way in which screen-refresh addresses are generated.




Figure 9-12. Screen-Refresh Address Registers

Screen Refresh and Video Timing - Video RAM Control

- DPYADR contains the SRFADR field, which is a counter that generates the addresses output during screen-refresh cycles.
- DPYSTRT contains the SRSTRT field, the starting address loaded into SRFADR at the beginning of each frame.
- DPYCTL contains several fields that affect screen-refresh addresses. The 8-bit DUDATE field is loaded with seven 0s and a single 1 that points to the bit position within SRFADR (bits 2–9 of DPYADR) at which the address is to be incremented (or decremented) at the end of each screen-refresh cycle. The ORG bit determines whether the screen-refresh address is incremented or decremented. If ORG=0, the screen origin is located at the top left corner of the screen and the address is incremented; otherwise, it is decremented. The NIL bit determines whether the GSP is configured to generate an interlaced (NIL=0) or noninterlaced (NIL=1) display. The generation of screen-refresh addresses can be modified to accommodate either type of display.
- The DPYTAP register is used to specify screen-refresh address bits to right of the position at which DUDATE increments the address. DPYTAP provides the additional control over screen-refresh address generation necessary to allow the screen to pan through the display memory.

Bits not directly involved in address generation are shaded in Figure 9-12.

The address output during a screen-refresh cycle identifies the starting pixel on the scan line about to be output to the monitor. Figure 9-13 (page 9-22) shows a 32-bit logical address of the first pixel on one of the scan lines appearing on the screen. The screen-refresh address consists of bits 4–23 of the logical address, which are generated by combining the values contained in SRFADR and DPYTAP. Where SRFADR and DPYTAP overlap (bits 10–17 of the logical address), the address bits are generated by logical ORing the corresponding bits of SRFADR and DPYTAP. The 8-bit DUDATE value contains seven 0s and a single 1 pointing to the position at which SRFADR is to be incremented (or decremented). The DPYTAP register should be loaded with the portion of the pixel address in Figure 9-13 lying to the right of the position indicated by the DUDATE pointer bit. SRFADR contains the portion of the pixel address that is incremented by the DUDATE pointer bit.

Following system power up, the software loads the starting screen-refresh address into the DPYSTRT register and the increment to the screen-refresh address into the DPYCTL register. For a typical CRT display, the starting address is the address in memory of the pixel that appears in the upper left corner of the display. If ORG bit in DPYCTL is 0, the *1's complement* of the starting address should be loaded into DPYSTRT. If ORG=1, the starting address loaded into DPYSTRT should *not* be complemented.

DPYADR is automatically loaded with the starting display address from DPYSTRT prior to the start of each frame. As shown in Figure 9-14 *a*, bits 2–15 of DPYSTRT (SRSTRT) are loaded into bits 2–15 of DPYADR (SRFADR). The load occurs coincident with the start of the horizontal blanking interval that occurs just at the end of the last active scan line of the preceding frame.





Figure 9-13. Logical Pixel Address


The address output during each screen-refresh cycle is contained in bits 2 through 15 of the DPYADR register (the 14-bit SRFADR field). As shown in Figure 9-14 *b*, DPYADR bits 4-15 are output at the LAD0-LAD11 pins during the row address time of the screen-refresh cycle. If ORG=0, the DPYADR bits are inverted before being output; otherwise, they are output unaltered. Zeros (logic-low level) are output on LAD12-LAD14, and a one (logic-high level) is output on LAD15; this is the \overline{RF} status bit.

During the column address time of the screen-refresh cycle, bits 2-6 of DPYADR are output at LAD6-LAD10. If ORG=0, the DPYADR bits are inverted before being output. DPYTAP bits 6-10 are ORed with DPYADR bits 2-6 and output at LAD6-LAD10. Bits 0-5 and 11-13 of DPYTAP are output at LAD0-LAD5 and LAD11-LAD13, respectively. Zeros are output at LAD14-LAD15 (the \overline{TR} and IAQ status bits).

After the row and column addresses have been output, the address in DPYADR bits 2-15 is decremented by the 8-bit value in DPYCTL bits 2-9 (the DUPDATE field). This is done in preparation for the next screen-refresh cycle. The 8-bit DUPDATE value is a binary number consisting of seven 0s and a single 1. This single 1 indicates the position at which DPYADR will be decremented. If ORG=0, the screen-refresh address in DPYADR is effectively incremented; the one's complement of the address contained in DPYADR is decremented by the DUPDATE amount, but is inverted before being output. This is equivalent to incrementing the address. If ORG=1, the address is decremented.




(a) Display-Address Initial Value




(b) Row-Address Time

Figure 9-14. Screen-Refresh Address Generation



(c) Column-Address Time

Figure 9-14. Screen-Refresh Address Generation (Continued)



(d) Display-Address Update

Figure 9-14. Screen-Refresh Address Generation (Concluded)

9.9.1.3 Screen Refresh for Interlaced Displays

The size of the DUDATE increment specified for an interlaced display should be twice that required for a noninterlaced display of the same dimensions. This allows every other line to be skipped during the even or odd field of an interlaced frame. Before the start of the even field, half the value of the DUDATE increment is added to the starting address loaded into DPYADR to obtain the necessary starting displacement. The SRSTRT field in DPYSTRT points to the area of memory containing the video data for scan line 1 in the example of Figure 9-9 on page 9-11.

9.9.1.4 Panning the Display

The DPYTAP register supports horizontal panning of the screen across a display memory that is larger than the screen. The value contained in the low-order bits of DPYTAP furnish the LSBs of the column address output during the screen-refresh cycle. Incrementing this value results in panning to the right; decrementing this value results in panning to the left.

9.9.1.5 Scheduling of Screen-Refresh Cycles

The internal request for a screen-refresh cycle is generated when horizontal blanking begins. This gives the GSP essentially the entire horizontal blanking interval in which to perform the screen-refresh cycle. The delay from the start of horizontal blanking to the start of the screen-refresh cycle is called the *screen-refresh latency*, and is determined by the internal memory controller.

The best and worst case screen-refresh latencies are given in Table 9-2. In the best case, the delay from the high-to-low transition of the BLANK output to the start of the screen-refresh cycle (the time the row address is output) is only 3.25 machine states (or local clock periods). In the worst case, the delay is $(7.25 + 2W)$ states, where W represents the number of wait states required per memory cycle. The worst case number is based on the fact that the start of the screen-refresh cycle can be delayed by up to three states if a read-modify-write operation began one state before the memory controller received the request for the screen-refresh cycle. A screen-refresh request is given higher priority than requests for DRAM-refresh, host-indirect or GSP CPU cycles; hence, no further delays will occur unless an external device generates a hold request.

Table 9-2. Screen-Refresh Latency

Min	Max
3.25 states	$(7.25 + 2W)$ states

Note: W is the number of wait states per memory cycle.

The horizontal blanking interval should be sufficiently long in duration for the screen-refresh cycle to be completed before blanking ends. The required minimum blanking interval is therefore about $(9.25 + 3W)$ machine states, depending on how soon after the end of blanking the external video logic begins clocking the VRAM shift registers. Of course, this time must be translated from machine states (local clock periods) to VCLK periods to program the HEBLNK register.

The horizontal sync pulse is permitted to be as small as a single VCLK period in duration.

No screen-refresh cycles are performed during vertical blanking until nearly the end of vertical blanking - at the start of the horizontal blanking interval that precedes the first active scan line of the new frame.

9.9.2 Video Memory Bulk Initialization

VRAMs may be rapidly loaded with an initial value using a special GSP feature that converts pixel accesses to shift register transfers. This rapid loading method is referred to as bulk initialization of the video memory, and can be used with VRAMs such as the TMS4161 and TMS4461. When the SRT (shift register transfer) bit in the DPYCTL register is set to a 1, all reads and writes of pixel data are converted at the memory interface of the GSP to shift-register-transfer cycles. When SRT=0, pixel accesses are performed in normal fashion.

When SRT=1, the processor can initiate shift-register-transfer cycles under explicit program control. By performing a series of such cycles, some or all of the display memory can be set to an initial background color or pattern very rapidly (in a small fraction of one frame time). First, the VRAM shift registers are loaded with the initial value. The video memory is then set to the initial color or pattern one row at a time by writing the shift register contents to the memory.

During a shift-register-transfer cycle (when SRT=1), the row and column addresses are output in unaltered form; that is, the address is not affected by the state of SRT. The 8-bit row address output during the cycle designates which row in memory is involved in the transfer. The direction of the transfer is determined by whether the cycle is a read or a write. A write cycle such as a PIXT transfer from a general-purpose register to memory is converted to a VRAM shift-register-to-memory cycle. Similarly, a read cycle such as a PIXT transfer from memory to a general-purpose register is converted to a VRAM memory-to-shift-register cycle.

Only pixel transfers are affected by the SRT bit. The manner in which all other data accesses and instruction fetches are performed is not affected.

Before bulk initialization of the display memory, the VRAM shift registers are loaded with the solid color or pattern with which the display memory will be loaded. This can be done in one of two ways, by either

- Serially shifting bits into the shift register

or

- First loading a row of display memory with the color or pattern using a series of "normal" pixel writes (when SRT=0), and then loading the contents of this row into the shift register by means of a PIXT memory-to-register instruction (executed while SRT=1).

To speed up the bulk initialization operation further, a series of transfers can be made more rapidly by using a single FILL instruction in place of a series of PIXT instructions. The fill region is selected so that each pixel write cycle generates a new row address. The fill region is specified to be precisely 16 bits wide, the width of the memory data bus. Also, plane masking is disabled, transparency is turned off, and the pixel processing *replace* operation is selected. This ensures that each row is addressed only once during the course of the fill operation.

The number of bits of the display memory that are altered by a single shift-register-to-memory transfer cycle is calculated by multiplying the number of VRAM devices by the number of shift register bits in each device. The entire frame buffer is loaded with the initial color or pattern in 256 memory cycles.