

An 11.8-in Flat Panel Display Monitor

The HP S1010A flat panel display is designed to be a plug-compatible replacement for CRTs used with HP workstations. This compatibility is provided by an interface board that uses the same analog signals that drive the CRTs to create digital signals to drive a high-resolution, high-performance LCD color display.

by David J. Hodge, Bradly J. Foster, Steven J. Kommrusch, and Tom J. Searby

The HP S1010A is a color flat panel liquid crystal display (LCD) monitor. It is designed to be a plug-compatible replacement for conventional CRT monitors on HP's workstation platforms in applications requiring lighter weight, lower power, or a smaller footprint than CRT displays. The HP S1010A uses an 11.8-in diagonal active matrix TFT (thin-film transistor) color LCD module that has a resolution of 1024 by 768 pixels (see Fig. 1).

To maintain compatibility with CRTs, the HP S1010A monitor is designed to accept analog red, green, and blue input signals with composite sync on the green line* and the same video display timing used for the CRTs on standard HP workstations. During implementation we found that the available LCD flat panel displays all had digital inputs and required a clock signal. Since CRT monitors don't use a

* The red and blue signals do not have sync pulses.

clock, the pixel clock for the flat panel display had to be generated internally.

Another feature of the HP S1010A monitor is its backlight saver and replaceable backlight assembly. Since their transmissivities are relatively low, color LCDs require a very bright backlight. Small, bright backlights tend to have relatively short operating lifetimes compared to CRTs. To minimize the problems associated with these short lifetimes we chose a panel with a replaceable backlight assembly. To minimize the need for replacement, we designed a circuit that can use the screen saver feature of windowing software to detect when the display is not in use and extinguish the backlight. The screen saver should be set up to display an all-black image when there is no user input for some period of time. To avoid cycling the backlight unnecessarily (for example, when the user taps the mouse as soon as the



Fig. 1. The HP S1010A flat panel display monitor weighs 12 lb including the stand, and measures 13.0 in (330 mm) wide, 14.8 in (377 mm) high, and 6.4 in (163 mm) deep.

screen saver turns on), the backlight only turns off after the black image has been present for about 15 seconds.

We incorporated a number of diagnostic features in the HP S1010A monitor. If no video is present, or the timing is incorrect and the display cannot lock onto the video signal, the display shows six horizontal color bars. Many computer system problems are hard to distinguish from monitor problems, since in either case, there is nothing on the display. The presence of the color bars tells users or service personnel that the display is functional, but that there is most likely a problem with the input video from the workstation.

The HP S1010A monitor has three LEDs that are visible from the rear of the display. One indicates whether power is present, another indicates that the backlight is turned on, and the third indicates that the display is properly locked on the input video signal. The LEDs can be used to determine quickly whether there is a problem with the power supply, the backlight, or the video.

Maintaining Plug Compatibility

Providing the capability to plug the HP S1010A into a typical HP workstation without any modification to the workstation meant that we had to design the circuits in the HP S1010A's display interface board to handle the same analog RGB composite signals sent to a CRT display. Essentially, we had to design circuits that are able to take analog signals and their associated timing and convert them to digital signals for driving counters and cells in an LCD matrix.

A conventional CRT builds an image by sweeping one or more modulated electron beams across the phosphors on the face of the tube from left to right and top to bottom. The signals that control the sweep of the beam are the horizontal and vertical sync signals, HSYNC and VSYNC respectively (see Fig. 2). As the beam sweeps across the face of the tube, it creates the image for a single horizontal line. At the end of the line, the beam is blanked, so that the retrace won't be visible, and the HSYNC signal causes the beam to sweep back to the start of the next line. The front porch provides some time for the beam to shut off before the HSYNC signal causes the retrace to begin. The back porch provides some time for the sweep circuits to stabilize before the beam is unblanked




Fig. 2. Composite analog video signals.

to start the next line. At the bottom of the frame the same thing happens during vertical retrace to get the beam back to the top.

Table I shows the timing for the composite video signals coming from an HP workstation and going to the HP S1010A. To keep everything properly synchronized, the vertical timing is all in exact integer multiples of the horizontal line time for the noninterlaced monitors used in workstations. Horizontal timing is in pixel periods, which is the time required to set up and display one pixel.

Table I
Timing for the Composite Video Signals
Input to the HP S1010A

	Horizontal (Pixels)	Vertical (Lines)
Period (Active + Blanking)	1344	840
Active	1024	768
Blanking (Porches + Sync)	320	72
Front Porch	64	4
Sync	128	4
Back Porch	128	64

If we think of an active matrix LCD as a big RAM, the display update process consists of writing to all of the cells. If we always update the cells the same way, all of the addressing functions can be done by horizontal and vertical address registers inside the display. Most LCDs accept multiple pixels on each clock to keep the clock rates down. Some actually update different parts of the display at the same time. The HP S1010A's LCD takes two consecutive pixels on each clock as it scans across the horizontal lines and down the frame. After each clock, the internal horizontal address counter increments to point to the next pair of pixels. At the end of the line, there are two things that need to happen: the horizontal counter needs to be reset and the vertical counter needs to be incremented. The signal that performs this function is very similar to the horizontal sync signal in a CRT, so we call it HSYNC. After the last line at the bottom of the frame, we need to reset the vertical address counter to start over at the top of the next frame. Since the signal to do this behaves very much like the vertical sync in a CRT, we call it VSYNC.

One important thing to note about the timing of CRTs and LCDs is that the LCD has much more flexibility. The CRT sweep circuits must drive the large inductive load of the deflection coils, so it is not practical to adjust the timing dynamically. Within limits, the LCD's timing can be varied. We took advantage of this characteristic in the design of the HP S1010A.

Hardware Architecture

Fig. 3 shows a block diagram of the main components of the HP S1010A flat panel display monitor. The analog video from the workstation comes into the termination network which matches the impedance of the video cable. The video then goes to the analog-to-digital converters (ADCs). The digital output is sampled and used to control the level adjust

Liquid Crystal Display Technology

Liquid crystal displays, or LCDs, are divided into two main classes: active and passive matrix LCDs. Passive matrix displays scan each of the cells, or pixels, sequentially. They are less complex and less expensive than active matrix devices, but the addressing technique they use means that each cell is only driven for a small fraction of the time. To maintain the image, the cells must hold their state for a long time (analogous to long decay phosphors in a slow-scan CRT). The disadvantage is that the response time of the display is slowed, which leads to ghosting on rapidly changing images, such as when the cursor is moved.


Active matrix displays have circuitry associated with each cell. The usual technique for building active matrix LCD circuits is to use a thin film of silicon grown on the display glass. This technique is known as thin-film transistor, or TFT. This type of LCD can be thought of as a big dynamic RAM, with one cell for each pixel. The RAM drives each liquid crystal cell continuously, and the RAM cells are refreshed by scanning the display. This enables the liquid crystal cell to be faster, improving response time dramatically. It also allows the display to have much higher resolution, since the drive time of each cell is not reduced by adding more cells. The increased resolution makes it more practical to build a color display since a color display has at least three times as many effective pixels as a monochrome display (one red, one green, and one blue subpixel for each pixel).

LCDs offer a number of advantages over CRTs. Because they are fabricated with a lithographic process, they offer excellent linearity, convergence, and purity. LCDs have no electron beam, making them unsusceptible to magnetic fields. They have lower power requirements (about 55W versus 100W for a comparable-size CRT). Because conventional CRTs need to deflect an electron beam, they must have a greater depth, and thus a larger footprint, than an LCD monitor. CRTs need to accelerate the electron beam, which requires high voltages that are not necessary for LCD monitors. Side effects of the high voltages include x-ray emissions and potential electrostatic problems in some environments. Finally, LCD monitors don't need a heavy glass bottle to maintain a vacuum, so they weigh much less than CRTs.

Accurate Positioning

To create a pleasing visual image without pixels jittering around, it is critical to position the sampling clock edges precisely and repeatably. If a pixel is sampled at horizontal position x in one frame, but at position $x+1$ in the next frame, the user will easily see the difference as noise on the monitor. Furthermore, the sample must be taken when the signal is stable and not transitioning between pixels. This requires the sampling edge precision to be significantly less than a single pixel time. The HP S1010A is able to recover the digital pixel data very cleanly.

For the HP S1010A monitor, positioning the sampling edge requires that it regenerate a dot clock with exactly the same frequency as the clock used in the workstation to generate the analog video data. The phase-locked loop circuit is used to synchronize an internally generated sync pulse with the horizontal sync pulse on the green video line. The horizontal sync signal provides only one synchronization event every 15.888 μ s. At 84 Mpixels/s, pixel time is only 11.82 ns, allowing 1344 pixels (one horizontal line of data) between each synchronization event. The transition time on the analog input from one level to the next is 4 ns worst case, leaving only 7.82 ns for setup and uncertainties over all temperature, voltage, and component variations.



*There is one ADC each for red, green, and blue analog signals.

Fig. 3. A block diagram of the components in the interface board for the HP S1010A flat panel display monitor.

Circuits which set the dc level of the ADC inputs and the ADC reference voltage so that the entire dynamic range of the ADCs is used. The video sync signal is extracted from the input video and sent to the dot clock regeneration circuit which generates the clock signals used by the ADCs to sample the video signal. The digital outputs of the ADCs are sent to the frame buffer which synchronizes the input signal to the flat panel timing and sends it to the panel.

Analog-to-Digital Conversion

The ADC circuitry converts the incoming video analog stream into the digital data that the flat panel display can accept. In the case of the HP S1010A monitor, 84 Mpixels/s are converted so that the four most-significant bits of each color's digital representation do not vary from one frame to the next because of digital noise. One of these four bits is used for control, and the other three bits define the three-bit color that the flat panel uses to generate eight intensity levels for red, green, and blue, providing a total of 512 displayable colors. The one-volt swing of the input analog video is divided into three main regions: the sync level, the blank level (typically 286 mV above the sync voltage), and the video level (ranges from 54 mV above blank signifying black to 714 mV above blank signifying white). These voltage levels are shown in Fig. 2. The period of the sync signal on the analog input is the time required to provide data for one horizontal line of pixels. On the HP S1010A monitor this time is 15.888 μ s.

Fig. 4 shows that the ADC circuitry is cleanly divided between the dot clock regeneration circuit and the support circuitry for the ADC. The dot clock regeneration circuit uses the horizontal sync signal from the green analog input line to create a clock which is used by the ADC to sample the analog video signal. The ADC support circuit is primarily concerned with controlling the offset of the input analog signal and the reference voltage to ensure that the ADC outputs have the desired digital range.

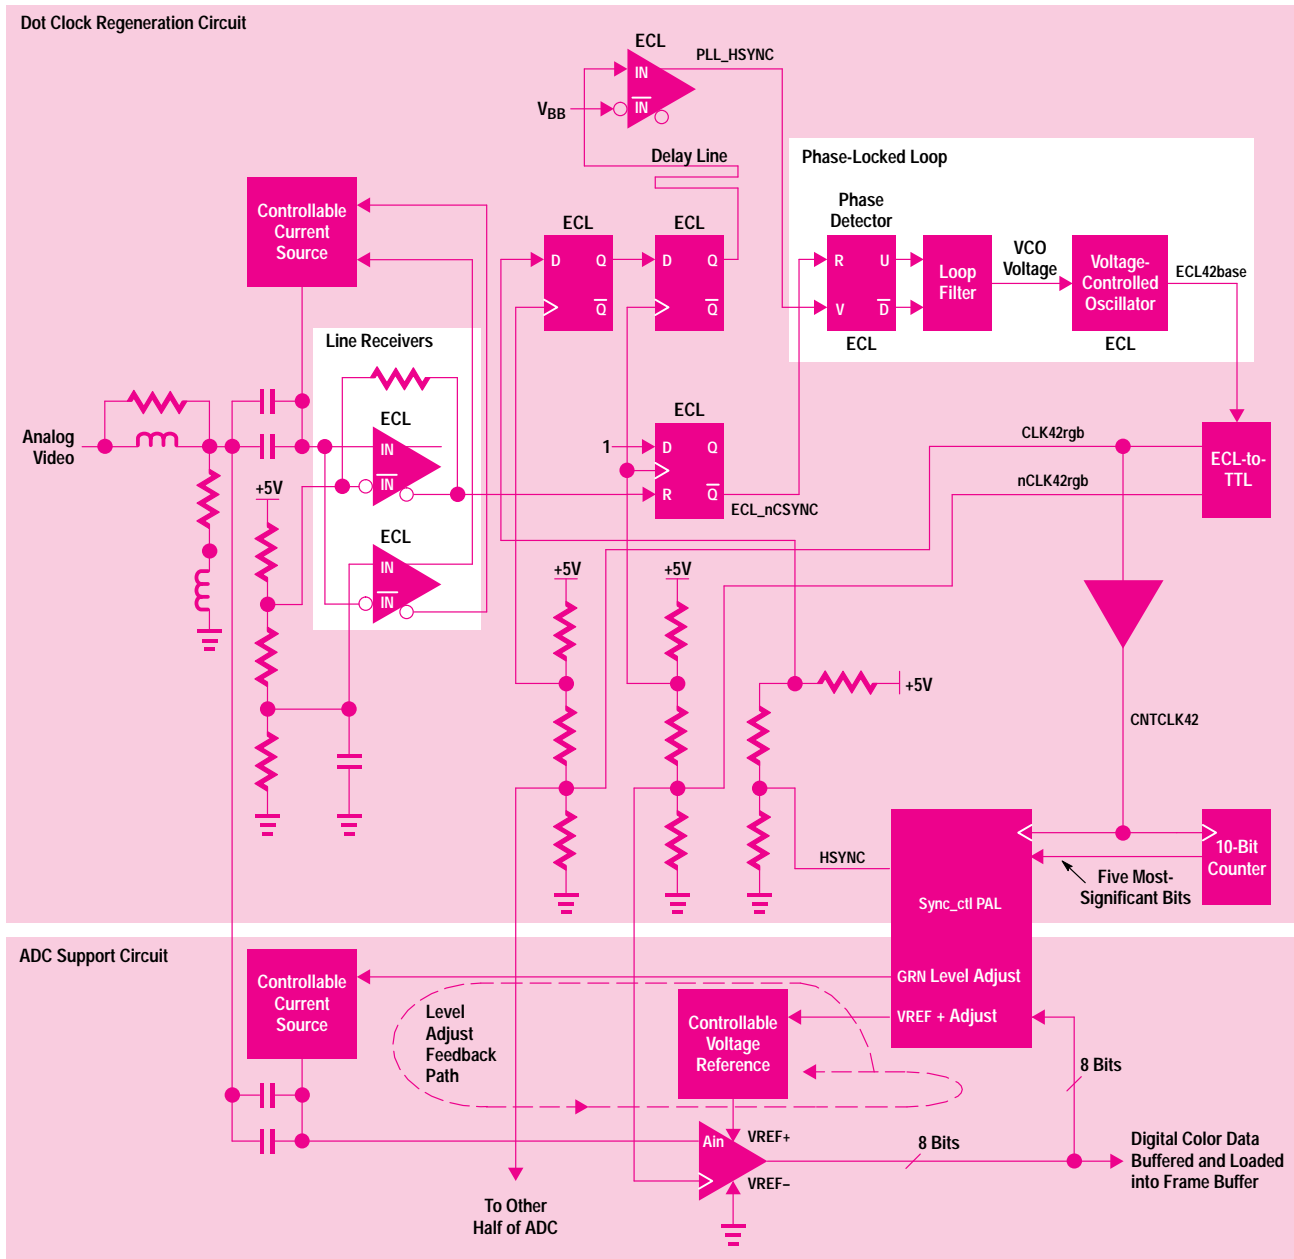


Fig. 4. The analog-to-digital conversion circuitry in the HP S1010A interface board.

Skew can be introduced in the three major areas: sync separation (extracting the sync signal from the analog input to be used by the phase-locked loop), internal clock skew (skew control between the edge of the internal sync signal and the clock sent to the ADC for sampling the data), and phase locking (degree to which the internal sync signal can be matched with the analog input sync signal).

Fig. 5 shows how the signals used in the HP S1010A monitor line up. Ultimately, we are trying to get the rising edges of CLK42rgb and nCLK42rgb to be precisely positioned with respect to the analog input. This is because these signals are used to clock the ADC that produces the digital pixel data.

The ECL_nCSYNC signal is generated from the analog video signal after the analog video signal has been terminated and voltage shifted through bypass capacitors. Separating the sync signal from the analog video signal is a task that all

display monitors must do and there are many parts commercially available for this purpose. However, none of these parts were acceptable for use in the HP S1010A monitor because of the great uncertainty about the propagation delay of these parts. Therefore, we used a carefully biased ECL differential line receiver (MC10E416) arranged in a Schmitt-trigger configuration. The MC10E416 is sensitive enough to distinguish between the sync voltage on the video input and the blank voltage which is only 286 mV higher. The MC10E416 has a small propagation uncertainty; its speed ranges between 200 ps and 550 ps, giving a net uncertainty of only 350 ps in propagation delay. Another differential line receiver is used to control the offset of the analog input to the first receiver. This biases the sync-to-blank transition voltage into the range where the MC10E416 is most sensitive.

Generating the internal sync signal (PLL_HSYNC) begins with the signal ECL42base, the output of the voltage-controlled oscillator in the phase-locked loop circuit. ECL42base is converted into TTL (CLK42rgb) and then buffered to the rest of the system as CNTCLK42. A set of counters with the appropriate control logic generates the TTL HSYNC signal which has a period equal to 1344 84-MHz clock periods (one full input line time). This signal is delayed by one and a half 42-MHz clock periods to produce one input to the phase detector in the phase-locked loop circuit. The final flip-flop, which generates PLL_HSYNC, is clocked by a simple TTL-to-ECL resistor ladder from nCLK42rgb, one of the clocks used for the ADC circuit.

Because of its high speed, ECL technology was used in time-critical functions in the dot clock regeneration circuitry, and TTL was used elsewhere because of its low cost.

After being generated by the flip flop, PLL_HSYNC goes through a delay line and then through one of the line receivers on the same MC10E416 that does the sync separation for the signal ECL_nCSYNC. The delay line compensates for all propagation delay shifts and results in correctly positioning the CLK42rgb positive edge within the analog video pixel period. The pass through the line receiver is done primarily to minimize skew uncertainty contributed by the line receiver itself. Passing HSYNC through the same part eliminates the 350-ps skew uncertainty mentioned above and allows us to use the specification for “within-device skew” which is only 50 ps! The fact that the MC10E416 also buffers the delay line from the sensitive phase detector of the phase-locked loop is an added bonus. Similarly, ECL_nCSYNC passes asynchronously through the same part used to generate PLL_HSYNC so as to minimize the same uncertainties mentioned above.

Throughout the synchronization process, the goal is to minimize the skew (actual propagation times are irrelevant). The commercial sync separators with propagation delays varying from 5 ns to 25 ns were unacceptable, but a part with a minimum delay of 24 ns and a maximum delay of 25 ns would have been acceptable, since another clock cycle and a delay line could have been accommodated. The end result is a 6-sigma skew budget of 11.44 ns, which is barely within the 11.82 ns pixel time.

One final issue in dot clock regeneration is the behavior of the phase-locked loop which locks the edges of PLL_HSYNC and ECL_nCSYNC by modifying the frequency of ECL42base. The phase detector controls the VCO voltage such that a 1-ns difference at one HSYNC edge will cause almost a 1-ns phase shift by the next HSYNC edge. However, because of the loop filter used, the average period of HSYNC changes little. If the HSYNC period is correct but out of phase by a few nanoseconds, on the next cycle the HSYNC period will still be correct, and the phase will no longer be shifted. Mathematically, this creates a response for the entire phase-locked loop system that is almost exactly critically damped. Lock is achieved quickly (less than 2 ms after connection to the analog input signal) and HSYNC drift is minimized. The phase detector we chose is one of the best on the market and can detect edge differences between PLL_HSYNC and ECL_nCSYNC of as little as 300 ps, which is referred to in phase-locked loop literature as the “dead zone.”

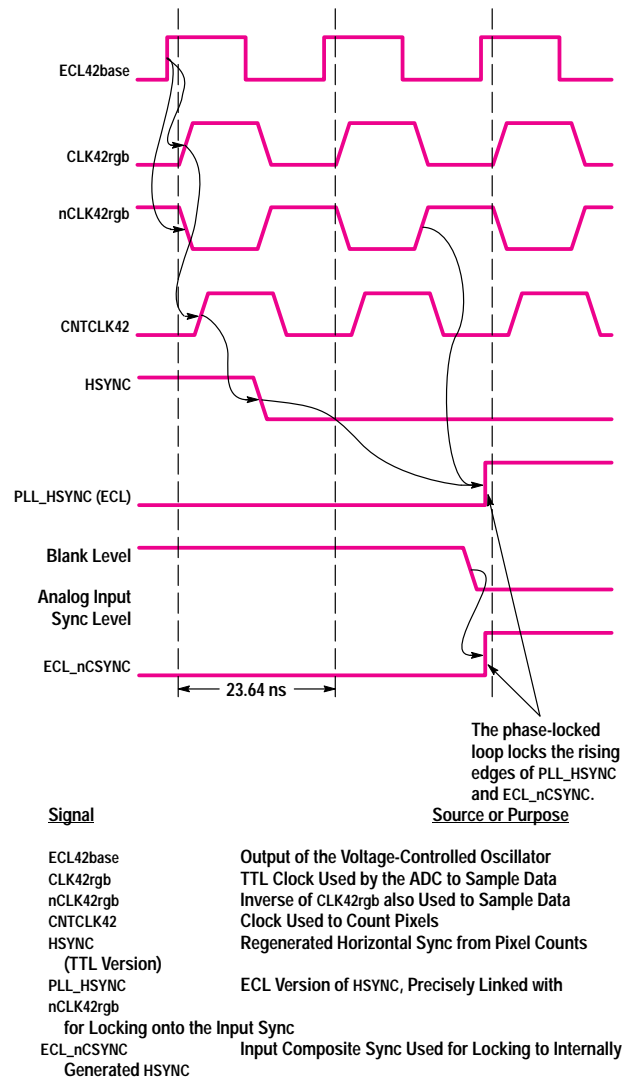


Fig. 5. The timing of the phase-locked loop input signals.

Automatic Scaling

To make the best use of the eight shades per color provided by the flat panel display, the HP S1010A uses feedback to control the ADC so that the full digital range is available. The control logic uses timing information from the dot clock counters to determine when the video is exhibiting a blank voltage and when it is exhibiting a sync voltage. The GRN (green) level adjust signal in the sync_ctl PAL in Fig. 4 is used to raise or lower the analog voltage to set the sync level at a digitized pixel value of 74.5. If the digitized data collected during the sync period is 74 or less, the GRN level adjust causes the input to the ADC to rise, whereas if the digital data during sync is 75 or more, it causes the ADC input to go down. In a similar fashion, VREF adjust is used to set the blank level to an eight-bit value of 123.5.

Setting sync (0 mV on the analog input signal) to 74.5 and blank (286 mV) to 123.5 gives a full white (1000 mV) level of 246 (Fig. 2). The most-significant bit of the ADC output is low during sync and blank and high during active video. The next four bits from the ADC represent the pixel value sent to the frame buffer. Since the flat panel uses only three bits, any ADC output between 240 and 255 will cause the

pixel to be displayed at full brightness. If the display is connected to another system with an analog output that is consistently 5% higher, sync will still be at 74.5 (still a 0-mV input), blank will be adjusted to 123.5 (now a 300-mV input), and full white will still be 246 even though the analog signal is at 1050 mV.

With a consistent digital translation being made on the voltage waveform, the input can be optimized for noise immunity. Typically, the input analog video will have 256 shades for each red, green, and blue signal, and with 660 mV for video, this is less than 3 mV per color. If all colors are allowed on the input, 3 mV of noise will be noticeable in the color sampling. On a CRT monitor, 3 mV of noise from frame to frame will be unnoticed by the user since this is such a small intensity difference. However, on the HP S1010A monitor only eight shades of each color are available, so any color change will be noticed. If noise appears on a color that lies near the digital sampling transition from one of the eight color levels to the next, this 3 mV of noise may cause the color to be 1/8 brighter or dimmer from frame to frame, which will clearly be noticed by the user. Our solution to this problem is called color centering and involves a software modification to host systems capable of creating an analog input to the HP S1010A monitor. This modification simply maps the 256 normal output shades into one of the eight levels that the HP S1010A is able to recognize with large noise margins. The host system performs this mapping by controlling what is written into the color map of the output RAMDAC (random access memory digital-to-analog converter) that generates the video.

Frame Rate Matching

One obstacle we had to overcome in the design of the display interface board was matching the slower frame rate of the LCD panel to the faster frame rate of HP's typical workstation video timing. To solve this problem, we included a full-frame VRAM frame buffer. Even with a frame buffer the two data streams have to be properly synchronized or the faster video input side will eventually catch up and pass the flat panel side sometime during the visible part of the frame. This can cause a tearing phenomenon in which part of the displayed frame on the flat panel comes from one input frame and the rest of the displayed frame comes from an

other input frame. When the image is changing rapidly, this phenomenon is visible to the user. To avoid this artifact, we have to store the digitized video information and shift it out to the flat panel display in such a way that only whole input frames are displayed, while still satisfying all of the flat panel display requirements.

One aspect of this frame buffer is that neither data stream (digitized video into the frame buffer or display data to the flat panel display) can be interrupted. This is not usually the case in typical graphics systems because the designer can interrupt the flow of data into the frame buffer when necessary. For the HP S1010A flat panel design an algorithm had to be devised to write and read the frame buffer in such a way that neither data stream is interrupted, while keeping things synchronized to prevent video tearing.

The digitized video feeding into the frame buffer is running at a 75-Hz frame rate, but the LCD monitor is not capable of running at more than about a 66-Hz frame rate. The ratio of the two frame rates is about 9:8, so we chose to match them by discarding every ninth frame of input video. If this is done synchronously, so that only whole frames are discarded, the user will not notice the skipped frame. In fact, the only effect will be a delay of no more than 28 ms in updating the display if the image changes during the skipped frame.

VRAM Frame Buffer Architecture

A four-pixel architecture is used for the HP S1010A monitor. Thus, every frame buffer write cycle stores four pixels of information. Each pixel has three colors, with four bits for each color (only three of these bits are used for color resolution so the least-significant bit is not used). Therefore, each RAM write cycle will write 48 bits (four pixels \times three colors \times four bits/color) into the frame buffer.

The frame buffer consists of six 256K \times 8-bit video RAMs (VRAMs), arranged in two banks of three. Each VRAM stores the data for one color (red, green, or blue) in two adjacent pixels in a single word. The digitized input video is written through the random port of the VRAMs, and the flat panel video is read from the serial port (Fig. 6).

Two things should be noted about this architecture. First, all six VRAMs are controlled by the same control signals (RAS, CAS, DSF, WBWE, DT/OE, and so on). This simplifies the control circuitry because separate signals are not required for each bank. Second, an entire horizontal line of information (1024 pixels) can be stored in half a row in the VRAM array. This also simplifies the control circuitry because each line fits in a single half of the split serial port register.

To simplify the control logic, the frame buffer begins to store data from the ADCs at the end of a vertical sync period. Since the vertical back porch (the delay between the end of sync and the start of displayed data of the digitized video) is 64 lines, the first 32 rows of data in the VRAMs are meaningless. The counters in the serial port control block are pre-loaded with the correct value so that the first 64 lines of data from the frame buffer are never displayed. This scheme eliminates the extra counter that would be needed to count scan lines to get through the vertical back porch.




Fig. 6. The frame buffer arrangement in the display interface board for the flat panel display.

(continued on page 58)

Product Design of the HP S1010A Flat Panel Display

Simplicity and elegance were the two main underlying objectives for the product design of HP's first standalone flat panel display monitor. Because of the high cost and resolution of the display technology, the product design needed to radiate innovation and quality. The use of many subtle curves gave the product a very soft and sophisticated look and feel.

Other mechanical objectives were to design a small-footprint, yet stable package with a wide tilt range and swivel, require no fan, be desktop or wall mountable, and have built-in security and cable management features.

Simplicity

The design is made up of two assemblies: the chassis/display assembly which houses the display module and control electronics, and the stand assembly which provides structure and dynamic movements (see Fig. 1). The stand has no electronics and is detachable. The overall structure is C-shaped which helped to reduce the footprint by balancing the display over the stand and provided a wider tilt range (Fig. 2). It also gave it an elegant, floating display look. An added benefit to the C shape is that a keyboard can fit under the display portion to free up even more desk space.

Because of the schedule and available engineering resources, simplicity was taken seriously. A human factors study was completed giving the desired height, tilt, and swivel ranges. However, designing individual height, tilt, and swivel

adjustment mechanisms would have taken more time and resources and potentially resulted in a bulkier design. For simplicity, a fixed height with a wider tilt range was decided upon that would meet most users' needs. As for swivel, the simple answer was: just slide it around. With the appropriate feet material, the monitor is light enough to be easily swiveled and slid anywhere on the desk. For thermal, size, and simplicity reasons, it was decided not to incorporate the power supply and instead use an external power module (off the desk, out of sight).

Chassis/Display Assembly

This chassis/display assembly shown in Fig. 1a consists of the LCD module, the interface printed circuit board, the power and brightness switch board, an aluminum chassis, a protective and conductive glass over the display, and cosmetic plastic covers. An aluminum chassis (as opposed to steel) was chosen to reduce weight and for EMI containment. The chassis contains a stainless-steel gasket to provide EMI contacts around the video and power connectors. A steel bracket is attached to the rear to provide a more rigid mounting location for the hinge and stand assembly. The plastic middle and back covers are heatstaked to the metal chassis. The printed circuit boards are snapped and then screwed into place. A protective glass, which is conductive and provides EMI containment, is taped to the display module metal housing. The display module is connected via cables to the control board and then screwed to the chassis. The plastic front cover hooks at the top on the middle cover and is then screwed underneath into the chassis. The

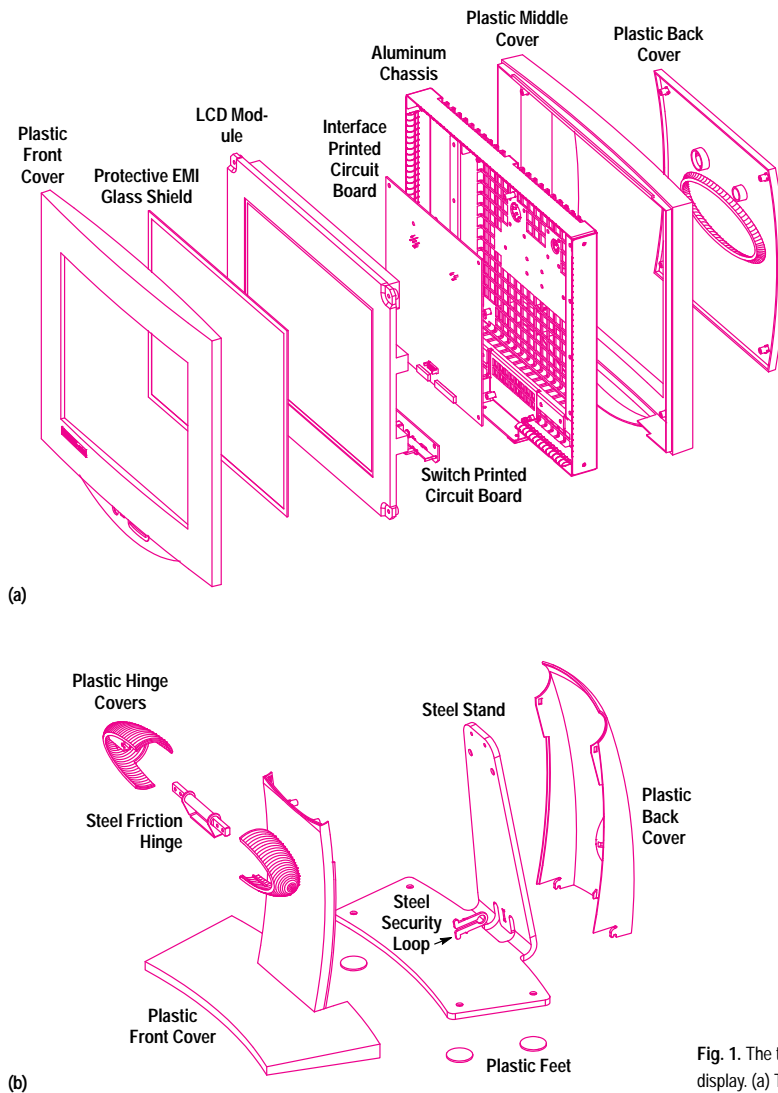


Fig. 1. The two assemblies that make up the product design for the HP S1010A flat panel display. (a) The chassis/display assembly. (b) The stand assembly.




Fig. 2. The tilt range of the flat panel display.

assembly weighs approximately 6 lb (2.7 kg) with the LCD module weighing approximately 3 lb (1.4 kg).

The Stand Assembly

To provide a stable base for the display assembly, the stand was designed out of heavy sheet steel with a counterbalancing shape (Fig. 1b). The stand assembly includes a custom steel friction hinge, a stainless-steel security loop, and cosmetic plastic covers. The security loop snaps into the metal stand. The plastic front cover is heatstaked to the metal stand. The plastic hinge covers are screwed to the hinge shaft. The hinge assembly is screwed to the stand. The plastic back cover, which incorporates a cable management recess, is hooked at the bottom into the front cover and rotates and snaps at the top into the front cover.

Final Assembly

The stand assembly is mounted to the display assembly via two screws. To expose the mounting holes, the back stand cover is snapped off and the hinge shaft is aligned with the chassis mounting bracket and secured with two screws. Fig. 3 shows different views of the final assembly of the display.

Conclusion

As a testimony to our adherence to the original design goals of simplicity and elegance, the product has won two major design awards: Design Zentrum Red Dot for High Design Quality (Germany 1994) and The Industrial Design Excellence Award-Gold 1994 (United States), featured in the June 6, 1994 issue of *Business Week*.




Fig. 3. Different views of the HP S1010A flat panel display.

VRAM Frame Buffer Control

The frame buffer control is responsible for, among other things, deciding when to do split data transfers from the VRAM array to the serial port shift registers and controlling which incoming frames get written into the frame buffer. (See a "A Note About VRAMs," on page 59 for some definitions of the terms used in this section.)

To ensure that the flat panel display always has the correct information to display, a split data transfer must occur on the random port side of the VRAMs for every line on the display. Table II shows the differences in the horizontal and vertical rates for the digitized video coming from the workstation and going to the LCD monitor.

Since the horizontal rate of the incoming video is faster than the horizontal rate of the LCD monitor, we are guaranteed that if we choose a point in the incoming horizontal period (say when we transition to horizontal blank) and do a single split data transfer every time we reach that point, we will always do at least one split data transfer for every horizontal line on the monitor. Although this seems to be straightforward, the scheme is complicated by a timing constraint of the VRAMs which prohibits split data transfers too close to the time when the monitor ends its current line. Fig. 7 shows the times (A and B) where split data transfers are not allowed.

Table II
Horizontal and Vertical Rates

	Horizontal		Vertical	
	Displayed μs	Total μs	Displayed ms	Total ms
Digitized Video	12.105	15.888	12.202	13.346
LCD Monitor	18.954	19.250	14.784	15.015

This requirement can be met if only one data transfer is allowed per flat panel horizontal line (one period of flat panel horizontal sync). This is achieved by looking at the sense of the qsf signal at the present time and comparing it to the sense it had the last time a data transfer was performed. If they are the same, this data is already in the shift register, so no data transfer is done. If they are different, a data transfer is done. In this way, no data transfers are performed at times A or B.

Every 15.888 μs there is a potential data transfer window. If a data transfer window falls in the middle of time B, that transfer will be inhibited because there will already have




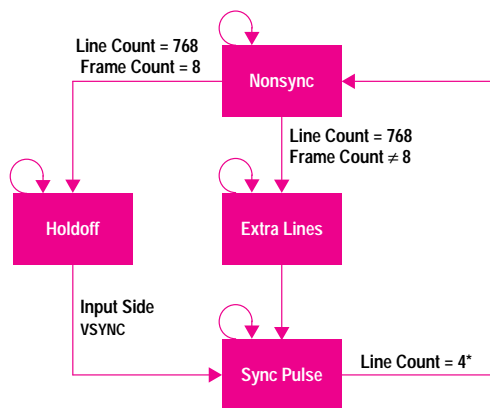
Fig. 7. Restricted split data transfer times.

been a transfer earlier in the same flat panel line. The qsf signal comes from one time domain and must be synchronized to another time domain before it is used to enable data transfers. By the time a change in qsf propagates through the synchronization and setup logic, time A is past, so no data transfer will occur there either.

Frame Rate Synchronization Algorithm

Data transfer in and out of the frame buffer must be synchronized to prevent the tearing phenomenon described above. Our synchronization technique does not try to synchronize the front-end clock (the output of the phase-locked loop) with the back-end clock (from an on-board oscillator). Instead, our technique uses events on the faster video input side to trigger events on the flat panel side. Specifically, for every eighth flat panel frame, the back end holds off asserting the vertical sync to the panel until it receives a vertical sync from the input video side. By doing this every eight frames and choosing the frequency of the back-end oscillator carefully and adding extra vertical front porch lines during the previous seven frames, a robust “on the fly” synchronization algorithm can be implemented.

Fig. 8 shows that there are only four states in the monitor’s vertical state machine. The Nonsync state represents the time when the flat panel display is receiving active video (i.e., the 768 lines of digitized input video). The Sync Pulse state represents the time when the vertical sync signal is sent to the flat panel display. The other two states (Extra Lines and Holdoff) represent the times when the flat panel display is in its vertical front porch. The flat panel doesn’t require any vertical front porch lines. However, to synchronize the incoming and outgoing frames in the frame buffer to avoid video tearing, we needed to add a few extra lines to the vertical front porch portion of the flat panel display to



*The LCD requires its VSYNC signal to last for four horizontal line times.

Fig. 8. The vertical state machine for the HP S1010A monitor.

A Note About VRAMs

Video RAMs, or VRAMs, are a variety of two-port dynamic RAM. They are designed to work well in graphics and video applications. The main port allows random access to any cell of the RAM. The other port consists of shift registers that are controlled by an independent clock. In the HP S1010A, the random port runs in the video input clock domain, and the serial port runs in the flat panel clock domain.

A data transfer operation loads the shift registers with data from the RAM array. The shift registers can be treated as two semi-independent halves, so that one half can be loaded without interfering with the data being shifted out of the other half. This provides more flexibility, since a data transfer operation (called a split data transfer in this case) can happen at any time while the other half is active, and transfers can be arranged so that there will be no interruption in the data flow out of the shift registers. The VRAM provides a signal called qsf to indicate which half of the shift register is active. When the data in the active half of the shift register is exhausted, qsf toggles, and the other half becomes active. This signals the HP S1010A’s control logic that it’s time to get ready for another split data transfer.

get an even ratio of eight flat panel frames for every nine input video frames.

After transfer of the last active line of video to the flat panel display, the state machine goes to the Extra Lines state where it will stay for eight horizontal flat panel lines (the eight-line vertical front porch for the current frame). The state machine then goes to the Sync Pulse state for four lines where it drives the flat panel vertical sync signal. It then goes back to the Nonsync state where it begins a new active line. This cycle repeats for seven flat panel display frames. On the eighth frame, the transition out of the Nonsync state goes to the Holdoff state. The state machine stays in Holdoff until VSYNC arrives from the input side. This is the signal to start the cycle again with another flat panel vertical sync. At this point, the input side and the flat panel display side of the frame buffer are synchronized.

Remember that one frame gets discarded. This frame is the first incoming frame after the synchronization event. Since the video input side has a 64-line vertical back porch, and the flat panel display side has none, the flat panel side will require a new frame of data immediately following the synchronization event, but the input side will not even start to write any valid data into the frame buffer until 64 horizontal lines later. In addition, since the input side is faster, it will catch up to the flat panel side sometime during this first frame causing a video tear. By not writing this first frame of data into the frame buffer, we can avoid the video tear. This is illustrated in Fig. 9. The numbers in the figure indicate




Fig. 9. Frame rate synchronization timing.

frames. Note that frame 0, which is the first frame in a sequence of nine frames, is not written into the frame buffer. The write enable signal controls writing a frame into the frame buffer. When low, the frame is skipped by not writing it into the frame buffer.

Acknowledgments

Many people made significant contributions to the design of this product. Bob Myers and Monish Shah gave valuable technical advice, John Metzner wrote the color centering software, Stuart Yoshida worked on the mechanical design, Howell Felsenthal managed the mechanical design team, and

Paul Cacciola designed the power supply. Osamu Suzuki and Sunny Hattori were our communication channel to the flat panel vendors. Nancy Venturato provided marketing support, and Steve Grotheer and Tony Barton managed the unusual manufacturing requirements. The people who helped with field support, regulatory compliance, printed circuit board layout, and environmental test were critical, and unfortunately, too many to name. Last, but by no means least, we would like to thank Steve Becker, project manager for the electrical design team, and Mike Myshatyn, section manager for the electrical design team, for their support.