Design of a 600-Pixel-per-Inch, 30-Bit Color Scanner

Simply sampling an image at higher resolution will not give the results a customer expects. Other optical parameters such as image sharpness, signal-to-noise ratio, and dark voltage correction must improve to see the benefits of 600 pixels per inch.

by Steven L. Webb, Kevin J. Youngers, Michael J. Steinle, and Joe A. Eccher

The objective of a scanner is to digitize exactly what is on the document that is being scanned. To do this perfectly would require a CCD (charge coupled device) detector with an infinite number of pixels and a lens with a modulation transfer function of 1.0, which does not exist. Modulation transfer function, or MTF, is a measure of the resolving power or image sharpness of the optical system. It is analogous to a visual test that an optometrist would use to measure a human eye's resolving power.

In the real world, the scanner user does not require a perfect reproduction of the original because the human eye does not have infinite resolving power. However, as originals are enlarged and as printers are able to print finer detail, the imaging requirements of the scanner are increased.

The HP ScanJet 3c/4c scanner, Fig. 1, is designed to obtain very finely detailed images for a variety of color and black and white documents and three-dimensional objects that are typically scanned. Its optical resolution is 600 pixels per inch, compared to 400 pixels per inch for the earlier HP ScanJet IIc. It produces 30-bit color scans compared to the ScanJet IIc's 24-bit scans, and its scanning speed is faster. The ScanJet 3c and 4c differ only in the software supplied with them.

Fig. 1. *HP ScanJet 4c 600-dpi, 30-bit color scanner.*

Optical Design

The HP ScanJet 3c/4c optical system is similar to that of the HP ScanJet IIc scanner,[1] with improvements to increase the optical resolution to 600 pixels per inch. Just sampling an image at higher resolution will not give the results a customer expects. Other optical parameters, such as MTF (i.e., image sharpness), signal-to-noise ratio, and dark voltage correction must improve to see the benefits of 600 pixels per inch.

The major optical components are:

- Two laminated dichroic composite assemblies used for color separation
- A fluorescent lamp with a custom mixture of phosphors
- A six-element double Gauss lens
- A three-row CCD sensor that has 5400 pixels per row
- Four front-surface mirrors.

The color separator composites, double Gauss lens, and CCD are shown in Fig. 2.

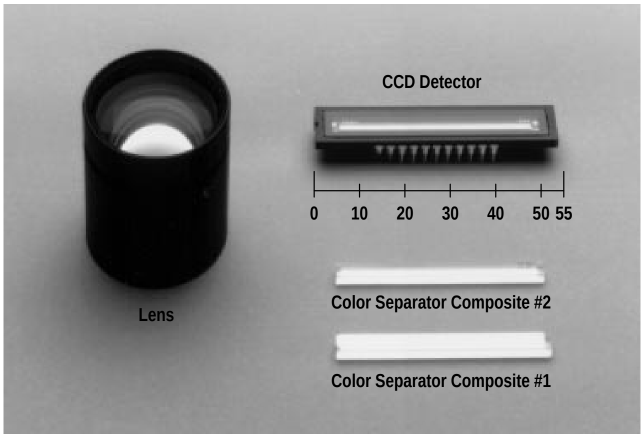

Fig. 2. *Lens, CCD (charge-coupled device) detector, and color separator composites.*

The color separation system (Fig. 3) consists of the two dichroic assemblies and the three-sensor-row CCD. With this method, red, green, and blue are scanned simultaneously, so only one pass is needed to scan all three colors. Each dichroic assembly is constructed of three glass plates that are bonded to each other with a thin layer of optical adhesive. Red, green, and blue reflective dichroic coatings are deposited onto the glass before lamination. The order of the coatings is reversed for the second dichroic assembly. The thickness of the glass plates between the color coatings and the flatness, tilt, and alignment are precisely controlled to ensure accurate color separation and image sharpness.

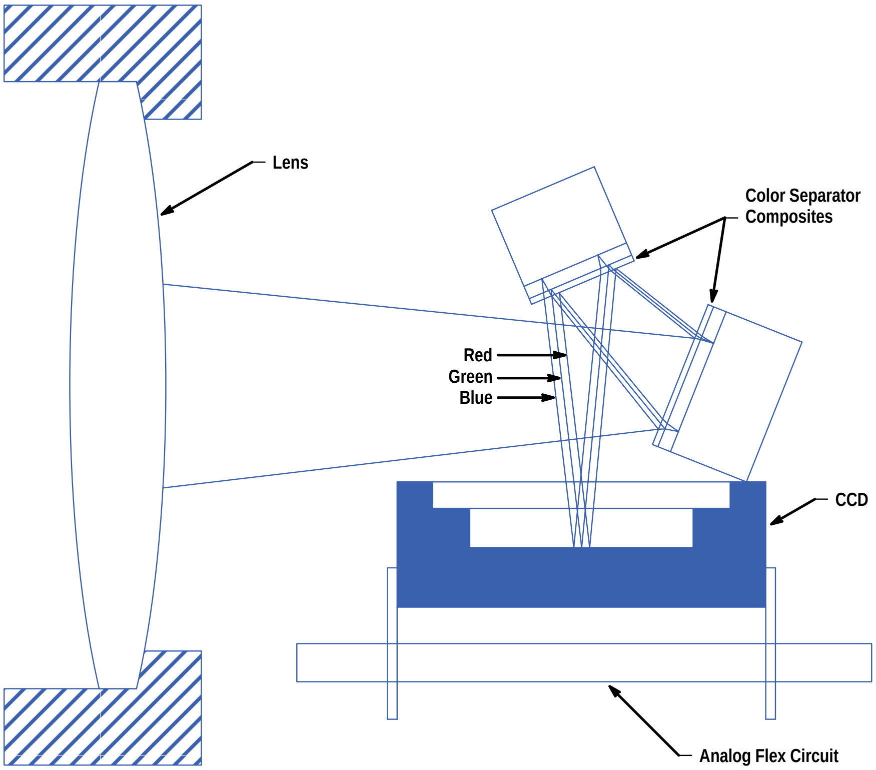

Fig. 3. *The color separation method uses two dichroic assemblies (composites) and a three-row CCD.*

Each color component is focused onto a CCD sensor row consisting of 5100 imaging pixels. Additional pixels are used for closed-loop dynamic light control, dark voltage correction, and reference mark location. By having all three rows integrated onto a single silicon chip, precise distances between the three rows are obtained. Production consistency is guaranteed by the integrated circuit process. Each CCD pixel generates a voltage signal that is proportional to the amount of light focused onto each pixel. The signal for each pixel is then processed and digitized. This data is sent to a computer or a printer.

Focus Optimization for Each Color

Two dichroic assemblies are used to equalize the path lengths of the three colors. A six-element double Gauss lens is used to focus the light onto the CCD sensors. However, the variation of the index of refraction of glass as a function of wavelength causes two of the three colors to obtain optimum focus at different locations. This phenomenon of differential refraction caused by wavelength dependence is best demonstrated by holding a prism up to a white light source and observing the colors. The light spectrum is separated because the shorter wavelengths (blue) are refracted or bent more than the longer wavelengths (red). Since lenses are made of glass that refract light of varying wavelengths at different angles, it is difficult to have all three colors focus at the same location.

To achieve simultaneous focus for all three colors there are several possible solutions. One is to design the focusing optics with curved front-surface mirrors only. However, these systems can be expensive, and it can be hard to correct other optical aberrations and difficult to image enough light onto the CCD. Another possible solution is to use an achromatic doublet. However, this type of lens can minimize chromatic aberration for only two of the three colors.

The ScanJet 3c/4c scanner optical design minimizes the chromatic aberration caused by the lens. An uncorrected optical system is shown in Fig. 4a, and a corrected optical system is shown in Fig. 4b. Lens chromatic aberration is corrected by adjusting the thickness of the dichroic coated plates. The path length of each color is adjusted to obtain optimum focus.

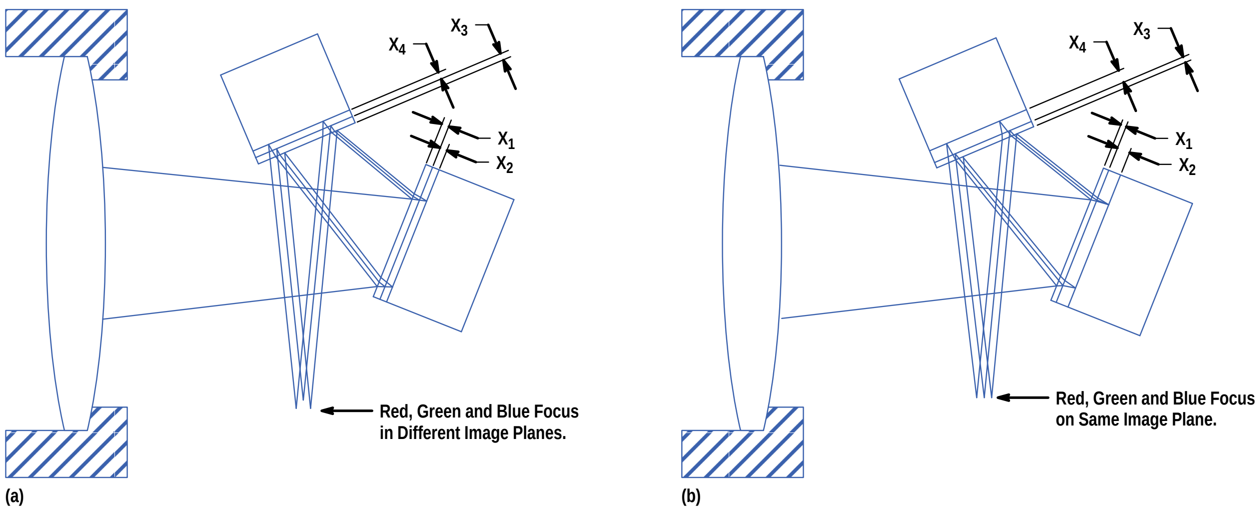

Fig. 4. *(a) Chromatic aberration of an uncorrected system.* $X_1 = X_2 = X_3 = X_4$. *(b) To ensure that simultaneous focus for red, green, and blue is achieved in the HP ScanJet 3c/4c scanner, unequal path lengths are used to compensate for the chromatic aberration of the lens.* $X1 = X3$ *and* $X2 = X4$, *but* $X2 \neq X1$.

Unequal path lengths for red, green, and blue would cause color registration error across the scan region. To prevent this, the CCD sensor row lengths are adjusted as shown in Fig. 5. Each row has the same number of pixels. However, the center-to-center spacing (pixel pitch) is slightly larger for a small number of pixels in rows 1 and 3. The pixels with slightly larger pitch are strategically placed to correct for the lateral chromatic aberration of the lens. This eliminates any color registration error that would have been caused by the lens.

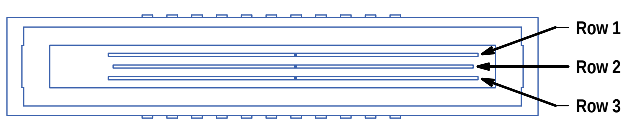

Fig. 5. *CCD row lengths are adjusted to compensate for color separator plate thicknesses.*

Optical System Layout

The lamp, lens, mirrors, color separators, and CCD are mounted into an aluminum carriage that is translated or scanned along the length of the document. The carriage is pulled underneath the glass platen by a belt connected to a stepper motor. The optical layout is shown in Figs. 6 and 7. Fig. 6 shows the mechanical design model of the carriage and light path. Fig. 7 shows part of the light path in more detail.

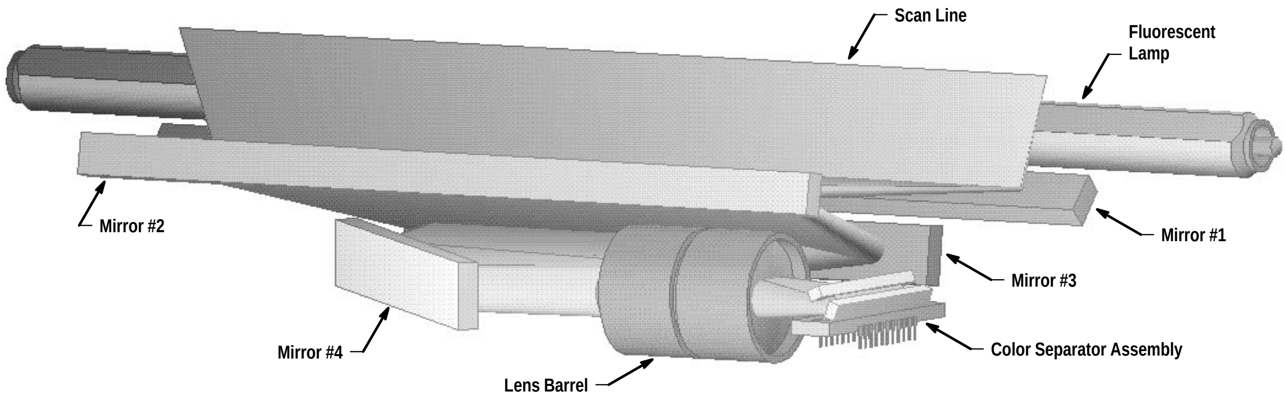

Fig. 6. *Mechanical design layout of the HP ScanJet 3c/4c optical path. The light path is from the scan line to mirror #1 to mirror #2 to mirror #3 to mirror #4 to the lens to the color separator to the CCD detector.*

The optical system was designed and evaluated using a commercially available optical design program. The sensitivity of optical tolerances such as lens centering, radii, thickness, and index of refraction were evaluated to determine the effects on image quality. The manufacturing assembly and mounting tolerances of key optical components in the carriage assembly were also evaluated. Image quality parameters such as MTF, color registration error, illumination uniformity, and distortion were emphasized.

To achieve precise optical alignment, custom assembly tooling was designed and implemented to meet production goals.

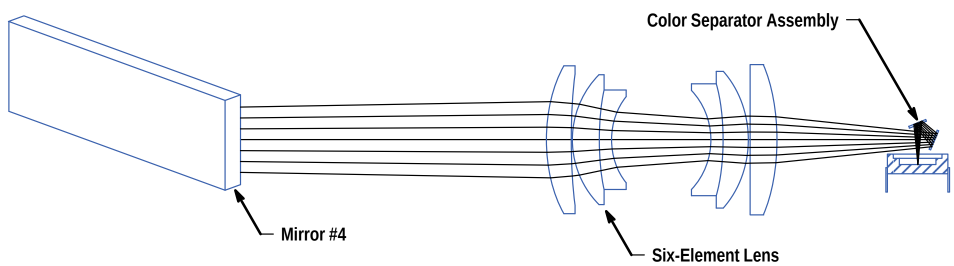

Fig. 7. *Ray trace of the optical path from mirror #4 to the color separator assembly (one color only).*

Fluorescent Lamp Driver

The fluorescent lamp is driven by a circuit that allows the lamp current to be varied over a range of 90 to 425 milliamperes. Since the lamp output is proportional to current, the lamp intensity is also varied.

A block diagram of the lamp driver circuit is shown in Fig. 8. The control inputs to the circuit provide the following functions:

- PREHEAT_L allows the filaments to be heated before the lamp is ignited.
- LAMP_PWM provides a pulse width modulated signal to set the desired current level.
- LAMPON_L turns the lamp on.

The filaments of the lamp are preheated for one second before lamp turn-on to reduce the amount of filament material that gets deposited on the insides of the glass. The deposits reduce light output, causing the light level to drop off near the ends of the lamp. This could create a lamp profile problem if preheating were not implemented.

The LAMP_PWM signal provides the desired current level plus a sync signal to the oscillator. The switching of the lamp driver power transistors occurs while the CCD (charged coupled device) is being reset. This helps keep switching noise from contaminating the CCD measurements. The lamp current command is derived from LAMP_PWM via the low-pass filter. The output of the low-pass filter is a voltage proportional to the amount of current desired.

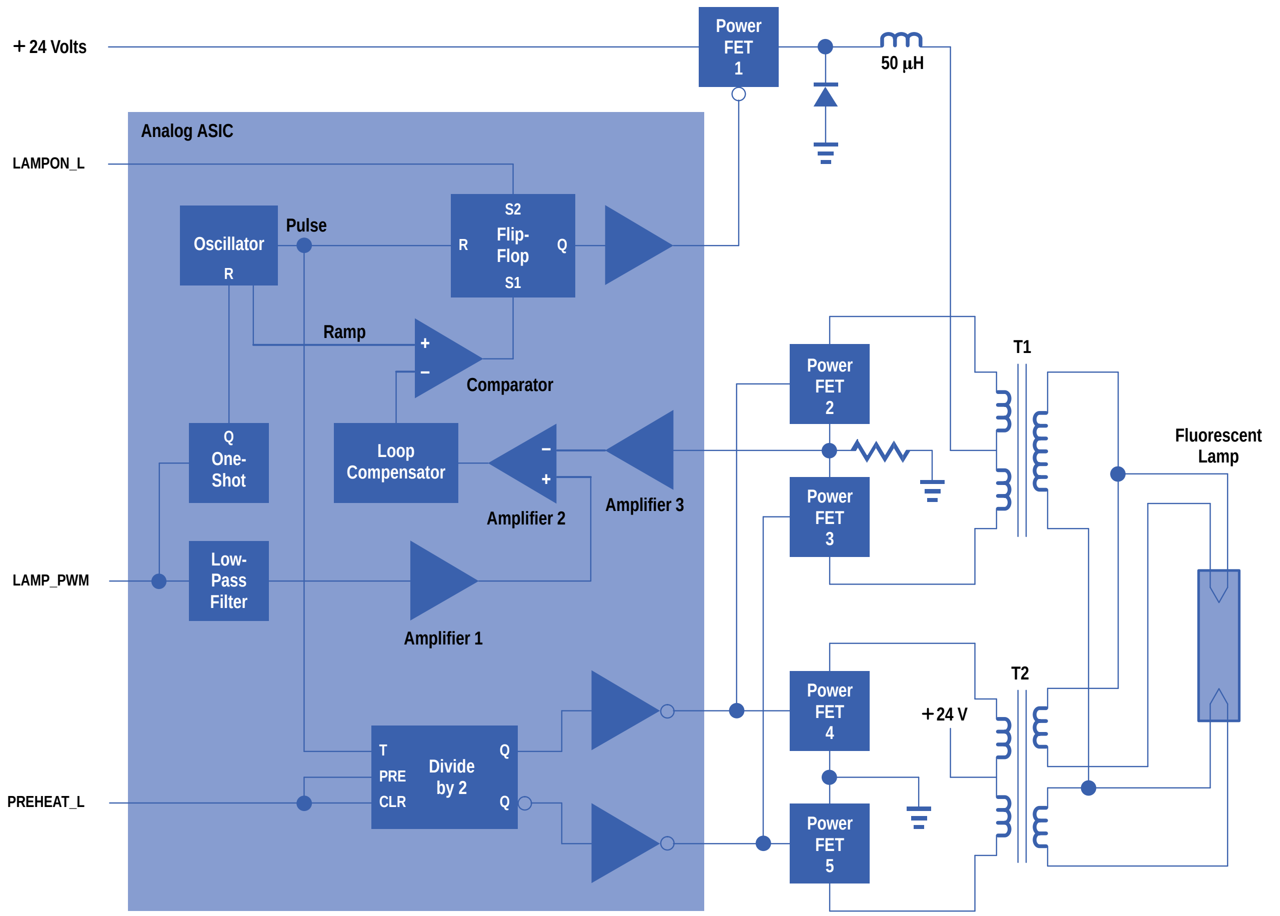

Fig. 8. *Block diagram of the fluorescent lamp driver.*

The LAMPON_L signal holds the flip-flop in the set mode until it is time to turn the lamp on. When the flip-flop is set, power FET 1 is held off via the buffer.

Operation of the lamp driver begins by taking PREHEAT_L to a logic zero. This allows the divide-by-2 circuit to begin toggling. When PREHEAT_L is high, both Q and $\overline{Q}$ are high, which turns off power FETs 4 and 5 via the inverting buffers. The toggling of the divide-by-2 circuit drives power FETs 4 and 5 out of phase. This provides a 24-volt square wave on the primary of T1 which is stepped down to 3.6V to drive the filaments. When LAMPON_L is activated, the flip-flop is reset on the next LAMP_PWM pulse, turning on power FET 1. The lamp appears as a high impedance in the off state, which results in power FETs 2 and 3 avalanching as a result of collapsing magnetic fields. The avalanche voltage of the power FETs is approximately 120 volts, half of which, or 60V, appears at the center tap of T1. This voltage is multiplied by the 1:6 turns ratio of T1 to produce 360V across the lamp. This voltage starts the lamp and the voltage drops to the low forty volt range. Current now flowing in the lamp is reflected back to the primary, where it is sensed. Amplifier 3 amplifies the voltage across the sense resistor and amplifier 2 subtracts it from the current command (output of amplifier 1).

The output of amplifier 2 is passed through the loop compensator (proportional plus integral) and applied to the comparator. The oscillator output is applied to the other input to the comparator. In the steady state, the loop compensator will stabilize at a voltage that produces the proper duty cycle on power FET 1 to maintain the commanded current. At this time the voltage across the 50-µH inductor will be in volt-second balance.*

All of the low-power analog and digital circuits are contained in an analog ASIC.

* The voltage across the inductor switches from positive to negative as the FET turns on and off. When the product of the positive voltage and its duration equals the product of the negative voltage and its duration, the inductor voltage is in volt-second balance.

Firmware Design

The firmware inside the ScanJet 3c/4c has many tasks. Two of the most critical (and most interesting to work on) were the start-stop algorithm and the light control algorithms.

Start-Stop. During some scans the host computer's I/O rate may not be able to keep up with the scanner's data generation rate. This will cause the internal buffer in the scanner to fill. When this occurs the scanner may need to stop and wait for the host to catch up (empty the internal buffer) before restarting the scan. This is called a start-stop. The scanner must restart the scan in the same place that it stopped or the user will see artifacts in the final image. If the scanner's drive system can start and stop within a fraction of the y-direction sampling size then no repositioning is needed. If the scanner's drive system cannot stop or start fast enough then it must back up and reposition the scan bar to be able to restart at the correct location (see Fig. 9).

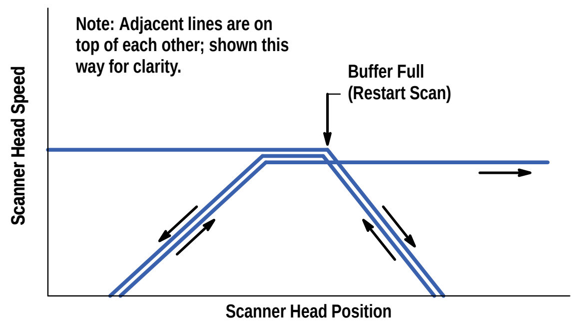

Fig. 9. *Start-stop profile.*

The ScanJet 3c/4c uses variable-speed scanning in the y-direction (along the length of the scan bed). Variable-speed scanning has two main advantages: better y-direction scaling and fast scan speeds at low resolution. The ScanJet3c/4c has a wide range of scan speeds (20 to 1), so the drive system needs some acceleration steps (of the stepper motor) to reach most of the final scanning speeds. This also means that the drive system cannot start or stop in one step. This dictated the need for a reposition movement for each start-stop.

There are three parts to a start-stop. First, when the internal buffer becomes full, the firmware marks the position and time of the last scan line and stops the drive system. Second, the firmware calculates how far to back up and then backs up and stops. Third, when there is enough space in the internal buffer the firmware accelerates the drive system up to the correct scanning speed and then restarts the scan line at the correct scan position.

The scanner firmware controls the step rate of the drive system. It uses its internal timer with a hardware interrupt to control the time between steps precisely. During acceleration, the firmware gets the next time interval from the acceleration table. Once at the proper scanning speed, the time interval is constant and the firmware just reloads the timer with the same interval. Deceleration uses the same table as acceleration in the reverse order. The firmware also keeps track of how many motor steps have occurred. Each motor step represents 1/1200 inch of travel for the scan head. This allows the firmware to keep track of the location of the scan head.

The scanner firmware also keeps track of when each scan line occurs (relative to a motor step). The scan lines are spaced 4.45 ms apart (for normal speed). A scan line may coincide with a motor step or may be between two motor steps, depending on the y-direction scan resolution). For example, for a 600-dpi scan there are exactly two motor steps for each scan line ($2 \times 1/1200 = 1/600$, so the scan head moves 1/600 inch in 4.45 ms). For a 500-dpi scan there would be 2.4 motor steps for each scan line.

When restarting the scan, the firmware must restart the CCD at least seven scan lines before putting scan data into the buffer. This is to allow the CCD to flush any extra charge in the system caused by restarting the CCD. The number of motor steps for seven scan lines depends on the y-direction scanning resolution. The number of steps to accelerate also depends on the y-direction scanning resolution. There is also a minimum number of steps that the drive system must be backed up to remove any mechanical backlash. These requirements determine the number of steps the scan head must be backed up (see Fig. 10). Once this is determined the firmware backs up the scan head and waits for the host to remove enough data from the internal buffer.

The internal buffer capacity inside the 3c/4c scanner is 256K bytes. Under the DOS operating system a typical receive block is 32K bytes (it can be larger). The ScanJet 3c/4c will restart a scan when the buffer is half full or holds less than twice the current receive block size, whichever is less.

Once there is enough space in the buffer the firmware restarts the scan. First, the scan head is accelerated up to the final scanning speed. A hardware interrupt is programmed to restart the CCD exactly seven scan lines before the position at which the last scan line was put into the buffer. Then, half a scan line away from the restart position, the buffer is reenabled such that the next line is put into the buffer. At this point the scan has been restarted and the start-stop is completed.

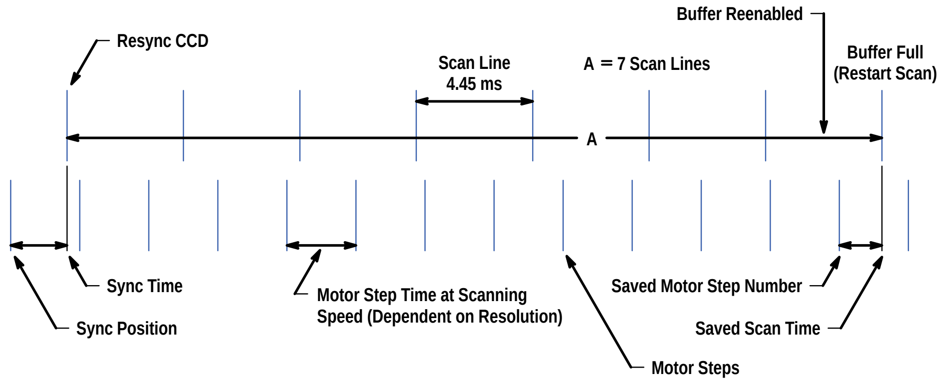

Fig. 10. *Start-stop timing.*

The start-stop accuracy of the ScanJet 3c/4c scanner is specified at half the y-direction scanning resolution. The typical resolution is between one-eighth and one-quarter pixel at the normal speed.

Light Control. The lamp in the ScanJet 3c/4c scanner is a special triphosphor fluorescent bulb. Using a fluorescent bulb has a number of trade-offs. The good news is that fluorescent bulbs have a range of phosphors to choose from. This allows the designer to balance the light spectrum with filters to give good colorimetric performance. The three phosphors in the ScanJet 3c/4c scanner give off red, green, and blue light. Florescent bulbs are also efficient, and give a reasonable amount of light for the energy used.

The bad news is that the intensity of the light is dependent on the bulb temperature. This means that as the bulb heats up the light gets brighter. If the bulb gets too hot, then the light gets dimmer again. What is worse, the bulb does not heat evenly across its length. The ends heat first and fastest and then the center of the bulb slowly heats up. The phosphors also have different efficiency-versus-temperature characteristics. This means that as the bulb heats up, it shifts color. At some nominal temperature, and only at that temperature, the phosphors are at their design efficiency, and the light is balanced with the filters. What makes this really bad is that the time it takes to complete a scan can vary between 15 seconds and 5 minutes. Fluorescent bulbs also have a long-term aging effect—a decrease in efficiency that affects performance—and the phosphors we have chosen age at different rates.

One solution to some of these problems is to leave the light on all the time. Then the bulb is at one stable temperature for the full scan. This solution has its own set of problems. For example, the bulb needs to be customer replaceable and the power consumption of the unit is high during idle time.

The ScanJet 3c/4c solves some of these problems with a real-time control system that controls the output of the light by modifying the power into the bulb during a scan. It also has separate red, green and blue system gains that are adjusted each time the light is turned on to help balance the overall color of the system. The light control system in the ScanJet 3c/4c uses the same CCD that is used for scanning. The CCD is wide enough so that it can look beyond the document being scanned at a white strip that runs along the length of the scan bed underneath the scanner top cover. This area of the CCD is called the *light monitor window.*

The light control algorithm for the ScanJet 3c/4c scanner has three parts. Part one turns on the power to the lamp and waits until some minimum level of light is detected. Part two tries to balance the output of the red, green, and blue channels by adjusting the independent system gains. Part three adjusts the power to the lamp to keep the green output at a fixed value during the scan. The purpose of part one of the lamp control is to turn the lamp on and make sure it is fluorescing at some minimum level. The goal for the startup algorithm (part two) is to have the lamp bright enough to scan with low system gains, which helps maximize the signal-to-noise ratio. The purpose of part three is to maintain the lamp at a given level for the entire scan.

Part one first sets the red, green, and blue gains to a low level. Then it turns on the preheaters (the coils at each end of the lamp) for about one second. It then turns on the lamp power, which is controlled by a pulse width modulation signal, to 20% for 4.5 ms and then to 80%. The first step at 20% is to help prevent the power supply from ringing. Once the lamp power is at 80% the control loop monitors the lamp output using the light monitor window. When the output of the lamp reaches or exceeds the minimum threshold, part two of the control algorithm starts. If the threshold is never reached the control loop will time out with an error (after about 5 minutes).

Part two of the algorithm waits about one second for the lamp to warm up (at 80% power). After the warmup delay the lamp power is lowered to 50% and the red, green, and blue system gains are adjusted. In the ScanJet 3c/4c there are two light monitor windows. One always reads the green channel's output, and the other reads either the red channel or the blue channel. The gain control loop adjusts the level of each system gain and tries to make the output of the light monitor window match a set value called the *desired value.* The window output is checked against the desired value on each end-of-scan-line interrupt, or every 4.45 ms. When the output of the green light monitor window matches its desired value (within some margin) 200 times in a row, the gains are considered stable and the green gain is fixed at its current value. If the control loop is unable to match the desired values by adjusting the gains, that is, the gains are at maximum or minimum values, it times out. The green gain is then fixed at slightly above the minimum value or slightly below the maximum value (to give the red and blue gains some margin).

Once the green gain has been fixed, the control loop switches from controlling the gains to controlling the power to the lamp. This is part three of the light control algorithm. The lamp power control loop uses only the green channel. It uses an eight-line running average to damp the control loop. If the control loop sees a difference of one count for eight lines or eight counts for one line between the light monitor window and the desired value, it changes the lamp power by one count. When the control switches from the gains to the lamp power, there is a short delay to load the eight-line average used in the lamp power control loop. After the short delay, the output of the green light monitor window is compared to its desired value, and if they match (within some margin) 200 times in a row, the light is considered stable and the scan is allowed to start. During this stabilization period the red and blue gains are being controlled. Once the light is considered stable the red and blue gains are fixed. The control loop for the lamp power using the green channel continues to operate during the scan. If the light fails to match the desired output 200 times in a row, the scanner will time out with a lamp error. Once the scan has started, if the control loop is unable to keep the output of the green light monitor window within some tolerance of its desired value, a lamp error is issued.

RFI and ESD Design

The ScanJet 3c/4c color scanner was a challenging design with respect to RFI (radio frequency interference) and ESD (electrostatic discharge). To begin with, the mechanical design didn't lend itself to stellar RFI and ESD performance. In an attempt to lower cost and weight, the design specified a plastic chassis instead of a sheet-metal chassis. Secondly, the design spread key electrical systems throughout the scanner. For example, the controller board was positioned in the lower rear of the product. The controller board clock is derived from a 36-MHz crystal oscillator. It generates the CCD clocks, motor control signals, and lamp control signals, processes all of the image data, and controls the SCSI interface. It also controls the optional automatic document feeder or the optional transparency adapter. Not only is the controller board a source of a lot of RF energy, it also has multiple interconnections that increase the difficulty of containing that RF energy. The controller board connects to the power supply, to the carriage, to the SCSI interface, and to any optional accessory.

Another key electrical system is the power supply assembly. Besides generating +5V, +24V, +12V, and −12V, the power supply assembly also contains the lamp and motor drivers. It has a total of five cable connections including the ac power cord, the dc power cable to the controller board, the lamp cable, the motor cable, and the LED power-on indicator cable (see Fig. 11).

The third key electrical system is the carriage, which has characteristics that dominate the scanner's basic EMC (electromagnetic compatibility) performance. The carriage is a metal casting that rides on two steel guide rods. The steel guide rods are held in place by a sheet-metal plate in the rear and by the plastic chassis in the front. A fluorescent lamp is mounted on the carriage and is connected through its own dedicated cable, the lamp cable, to the lamp driver in the power supply. The lamp cable is about 15 inches long and travels along the right side of the scanner as the carriage moves under the glass window. The imaging flex circuit is a two-layer circuit that is wrapped around the outside of the CCD and is connected through the carriage cable to the controller. It is located in the left rear of the carriage. The carriage cable is a single-layer unshielded flexible cable that carries CCD clocks, which can run at speeds over 1 MHz, to the imaging circuit from the controller board. This cable also returns the resulting analog image data. The carriage cable, which is about 25 inches long, travels along the left side of the scanner as the carriage is in motion (see Fig. 11).

The carriage is a source of energy from the imaging circuit. It is also an antenna whose electrical length changes with the position of the carriage. At least three different electrical structures change as the carriage moves from the back of the scanner to the front. These include the carriage cable, the lamp cable, and the current path through the steel guide rods and the carriage. Because of this dynamic antenna structure, the radiating efficiency for any specific frequency will be optimized at one corresponding specific position of the carriage over its range of travel. One can think of it as a "self-tuning" antenna. Typical RFI control approaches that merely retune energy from one frequency to another simply do not work because the new frequency to which the RF energy is shifted will just correspond to a different carriage position at which the antenna efficiency is optimized for that frequency.

A number of RFI suppression techniques were considered. Putting a Faraday cage around the whole scanner was, of course, impossible because the top needed to be glass. Trying to enclose all the electronics and shield all of the cables also proved futile. Enclosing the controller board only seemed to make things worse. Using power and ground islands didn't help. Ferrites didn't seem to have a lot of impact, and extrapolating their performance, we estimated that RFI might only decrease by 5 dB if the box were completely filled with ferrite. Using capacitors to roll off clock or clock-like signals only seemed to increase emissions below 300 MHz.

We decided that the best approach to keeping RFI emissions down was to reduce all possible sources as much as possible. We needed to minimize the energy that got onto the carriage structure, because any energy that got there would be radiated efficiently at some point in the carriage's travel. We began to work on some new approaches that were guided by theory and that we later confirmed with experiment. First of all, we revisited the equation that describes the radiation from a current loop. Because this radiation is proportional to the product of the frequency squared, the current, and the loop area, we tried to minimize the areas of current loops and to minimize the current in those circuits with series impedance. Because we did

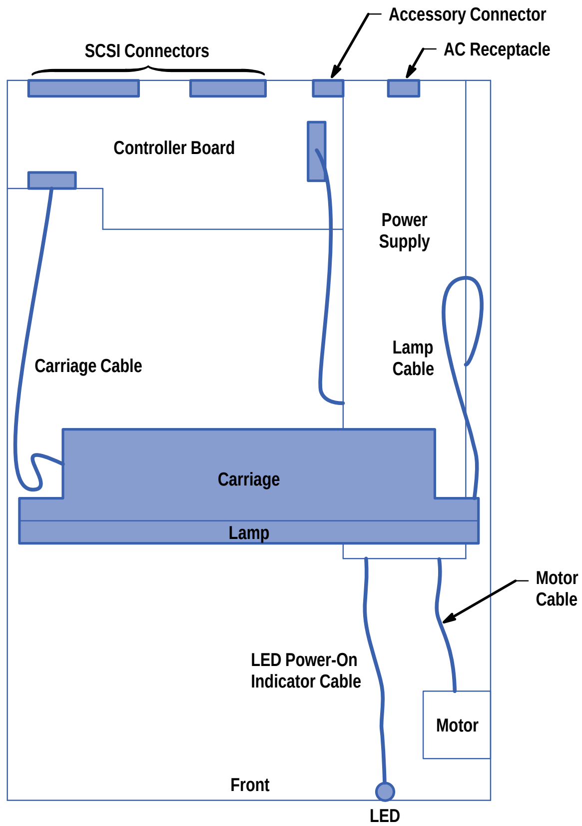

Fig. 11. *Scanner internal layout showing key components for RFI design.*

not want energy traveling onto the self-tuning antenna, we purposely tried to mismatch cable impedances so that most of the energy would be reflected back onto the controller board rather than traveling out onto the carriage cable. To do this, grounding and shielding needed to be minimized. This meant doing things that were just the opposite of what would normally be done. Instead of routing the carriage cable close to metal, it was raised away from any metal to increase its effective transmission line impedance. Although the carriage cable became a better antenna, far less high-frequency energy was able to get onto that antenna because of the impedance mismatch.

ESD also required an unusual approach. Initially, the scanner was highly susceptible to static discharges. An air discharge of only 1 kV would usually cause the SCSI bus to hang even if there was no data transfer in progress. This problem was ultimately improved by over an order of magnitude by the inclusion of a part affectionately known as the BMP or *big metal plate*. The BMP is simply the flat metal plate that is affixed to the bottom of the scanner. Its exact physical dimensions turn out to be relatively unimportant because it doesn't perform its function through any shielding or plane imaging phenomenon. It is attached to the SCSI cable shield and merely serves as a huge charge sink. The BMP could be connected to the SCSI shield without regard to three-dimensional position and it would always improve the ESD air discharge performance to over 10 kV, even while data was being transferred over the SCSI interface.

The ScanJet 3c/4c also inspired an interesting solution to a common ESD/RFI problem. Often, different methods of connecting the chassis to dc ground will have different effects on RFI and ESD. In the ScanJet 3c, if the chassis was connected directly to dc ground at the SCSI connectors, ESD performance was improved. However, if chassis ground wasn't connected at all to dc ground except in the power supply, RFI was improved. In the end, by connecting chassis ground to dc ground through parallel diodes oriented in opposite directions (see Fig. 12), good performance for both RFI and ESD was achieved.

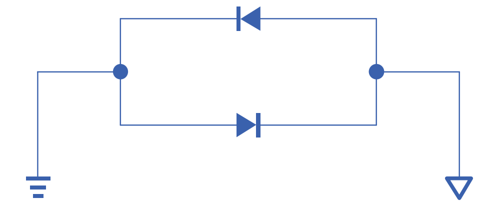

Fig. 12. *Diode connection for ESD and RFI suppression.*

Acknowledgments

The authors would like to acknowledge the contributions of the following individuals to the design of the HP ScanJet 3c/4c scanner: Bob Emmerich, Ray Kloess, Greg Degi, Irene Stein, Nancy Mundelius, Dave Boyd, Kent Vincent and Hans Neumann of HP Laboratories, project managers Gerry Meyer and Gordon Nuttall, section manager Jerry Bybee, and R&D manager Dean Buck.

Reference

1. K.D. Gennetten and M.J. Steinle, "Designing a Scanner with Color Vision," *Hewlett-Packard Journal*, Vol. 44, no. 4, August 1993, pp. 52-58.

▶ Go to Subarticle 8a
▶ Go to Next Article
▶ Go to Journal Home Page