BY SAM GOLDWASSER
SERVICE EDITOR

CD Information Storage and Playback

BEFORE WE CAN GET INTO THE NITTY-GRITTY OF REPAIRING A CD PLAYER OR CD-ROM DRIVE, WE NEED A LITTLE MORE INFORMATION ON HOW CD INFORMATION STORAGE AND PLAYBACK WORKS. LET'S TACKLE THAT NEXT.

The actual information to be recorded on a CD undergoes a rather remarkable transformation as it goes from raw audio (or digital data) to microscopic pits on the disc's surface. For commercial or professional audio recording, the process starts with pre-filtering to remove frequencies above 20 kHz. It continues with analog-to-digital conversion, usually at a sampling rate of 48K samples/second for each stereo channel. The resulting data stream is then recorded on multi-track digital magnetic tape. All mixing and pre-mastering operations are done at the same sampling rate. The final step is conversion through re-sampling (sample-rate conversion including sophisticated interpolation) to the 44.1K samples/second rate actually used on the CD (88.2K total for both channels). (In some cases, all steps may be performed at the 44.1 K rate.) That is followed by sophisticated coding of the resulting 16-bit "two's-complement" samples (alternating between L and R channels) for the purpose of error detection and correction. Finally, the data is converted to a form suitable for the recording medium by Eight-to-Fourteen modulation (EFM) and then written on a master disc using a precision laser cutting lathe. A series of electroplating, stripping, and reproduction steps then produce multiple "stampers," which are used to actually press the discs you put in your player. Of course, it is possible to create your own CDs with a modest priced CD-R recorder (which does not allow erasing or re-recording), and now with re-writable CD technology with fully reusable discs that allow editing similar to what can be done using cassette tape.

Like a phonograph record, the information is recorded in a continuous spiral. However, with a CD, that track (groove or row of pits—not to be confused with the selections on a music CD) starts near the center of the CD and spirals counter-clockwise (when viewed from the label side) toward the outer edge. The readout is through the 1.2 mm polycarbonate disc substrate to the aluminized information layer just beneath the label. The total length of the spiral track for a 74-minute disc is over 5000 meters—which is more than 3 miles in something like 20,000 revolutions of the disc!

The digital encoding for error detection and correction is called the Cross Interleave Reed Soloman Code, or CIRC. To describe that as simply as possible, the CIRC consists of two parts: interleaving of data so that a dropout or damage will be spread over enough physical area (hopefully) to be reconstructed and a CRC- (Cyclic Redundancy Check) like error-correcting code. Taken together, those two techniques are capable of some remarkable error correction. The assumption here is that most errors will occur in bursts as a result of dust specs, scratches, or imperfections such as pinholes in the aluminum coating, etc. For example, the codes are powerful enough to totally recover a burst error of greater than 4000 consecutive bits—about 2.5 mm on the disc. With full error correction implemented (that is not the case with every CD player), it is possible to put a piece of 2-mm tape radially on the disc or to drill a 2-mm hole in the disc, and have no audio degradation. Some test CDs have just this type of defect introduced deliberately.

Two approaches are taken with uncorrectable errors: interpolation and muting. If good samples surround bad ones, then linear or higher-order interpolation might be used to reconstruct the bad samples. If too much data has been lost, the audio is smoothly muted for a fraction of a second. Depending on where those errors occur in relation to the musical context, even such drastic measures might be undetectable to the human ear.

Note that the error correction for CD-ROM formats is even more involved than for CD audio as any bit error is unacceptable. That is one of many reasons why it is generally impossible to convert an audio CD player into a CD-ROM drive. However, since nearly all CD-ROM drives are capable of playing music CDs, much can be determined about the nature of a problem by first testing a CD-ROM drive with a music CD.

Compact Disc Construction

As the following discussion proceeds, we will be expanding on some of the concepts introduced above.

The information layer uses "pits" as the storage mechanism. Pits are depressions less than 0.2 µm (1 µm = 0.001 mm = 0.000001 meter = 1/25,400 of an inch) in depth ($^1/_4$-wavelength of the 780-nm laser light, taking into consideration the actual wavelength inside the polycarbonate plastic). Thus, the reflected beam is 180 degrees out of phase with incident beam making for high-contrast edges and good signal-to-noise ratio. Everything that is not a pit is a "land". Pits are about 0.5-µm wide; their length varies with the information content—with each bit being represented by a 0.278-µm increment.

Each byte of the processed information is converted into a 14-bit run-length-limited code taken from a code book (lookup table) such that there are no fewer than two or more than ten consecutive 0s between 1s. By then making the 1s transitions from pit to land or land to pit, the minimum length of any feature on the disc is no less than 3P and no more than 11P, where P is 0.278 µm. This is called Eight-to-Fourteen Modulation—EFM. Thus the length of a pit ranges from 0.833 to 3.054 µm.

Each 14-bit code word has 3 additional sync and low-frequency-suppression bits added, for a total of 17 bits representing each 8-bit byte. Since a single bit is 0.278 µm, a byte is then represented in a linear space of 4.72 µm. EFM in conjunction with the sync bits assures that the average signal has no DC component and that there are enough edges to reliably reconstruct the clock for data readout. These words are combined into 588-bit frames (see Table 1). Each frame contains 24 bytes of audio data (6 samples of L+R at 16 bits) and 8 bits of information used to encode (across multiple frames) information like the time, track, index, etc.

A block, which is made up of 98 consecutive frames, is the smallest unit that can be addressed on an audio CD and corresponds to a time of $^1/_{75}$ of a second. Two bits in the information byte of each frame are currently defined. These are called P and Q. P serves a kind of global sync function, indicating start and end of selections, time in between selections, and so forth. The 98 Q bits of each block encode the time, track and index number, as well as many other possible functions depending where on the disc it is located, what kind of disc it is, and so forth.

Information on a CD is recorded at a Constant Linear Velocity—CLV. That is both good and bad. For CD audio at 1× speed, this CLV is about 1.2 meters per second. (It really isn't quite constant—due to non-constant coding packing density and data buffering—but instead varies between about 1.2 and 1.4 meters-per-second). CLV permits packing the maximum possible information on a disc since it is recorded at the highest density regardless of location.

However, for high-speed access, particularly for CD-ROM drives, it means there is a need to rapidly change the speed of rotation of the disc when seeking between inner and outer tracks. Of course, there is no inherent reason why, for CD-ROMs, the speed could not be kept constant meaning that data transfer rate would be higher for the outer tracks than the inner ones. Modern CD-ROM drives with specs that sound too good to be true (and are), may run at constant angular speed achieving their claimed transfer rate only for data near the outer edge of the disc.

Note that unlike a turntable, the instantaneous speed of the spindle is not what determines the pitch of the audio signal. For one thing, there is extensive buffering in RAM inside the player. That buffering is used as a FIFO to smooth out data read off of the disc to ease the burden on the spindle servo, as well as to provide temporary storage for intermediate results during decoding and error correction. Pitch (in the music sense) is instead determined by the data readout clock—usually a crystal oscillator—which controls the D/A and LSI chip-set timing. The only way to adjust the player's pitch is to vary that clock. Some high-end players include a pitch adjustment that does just that.

Since the precision of the playback of a CD player is determined by a high-quality quartz oscillator, wow and flutter—key measures of the quality of phonograph turntables—are so small as to be undetectable. Ultimately, the sampling frequency of 44.1K samples-per-second determines the audio output. For this, the average bit rate from the disc is 4.321M bits-per-second.

Tracks are spaced 1.6-µm apart—a track pitch of 1.6. Thus, a 12-cm disc has over 20,000 tracks for its 74 minutes of music. Of course, unlike a hard disk and like a phonograph record, it is really one spiral track over 3 miles long! (However, as noted above, the starting point is near the center of the disc.) Compare that to an LP record: A long long-playing LP might have a bit over 72 minutes of music on two sides, or 36 minutes per side. (Most do not achieve anywhere near this much music since the groove spacing needs to vary depending on how much bass content the music has; bass requires wider grooves, and wide grooves occupy more space.) At 33-$^1/_3$ rpm, this is just over 1200 grooves in about 4 inches compared to 20,000 tracks on a CD in a space of just over 1.25 inches! The readout styles for an LP has a tip radius of perhaps 2 to 3 mils (50 to 75 µm)

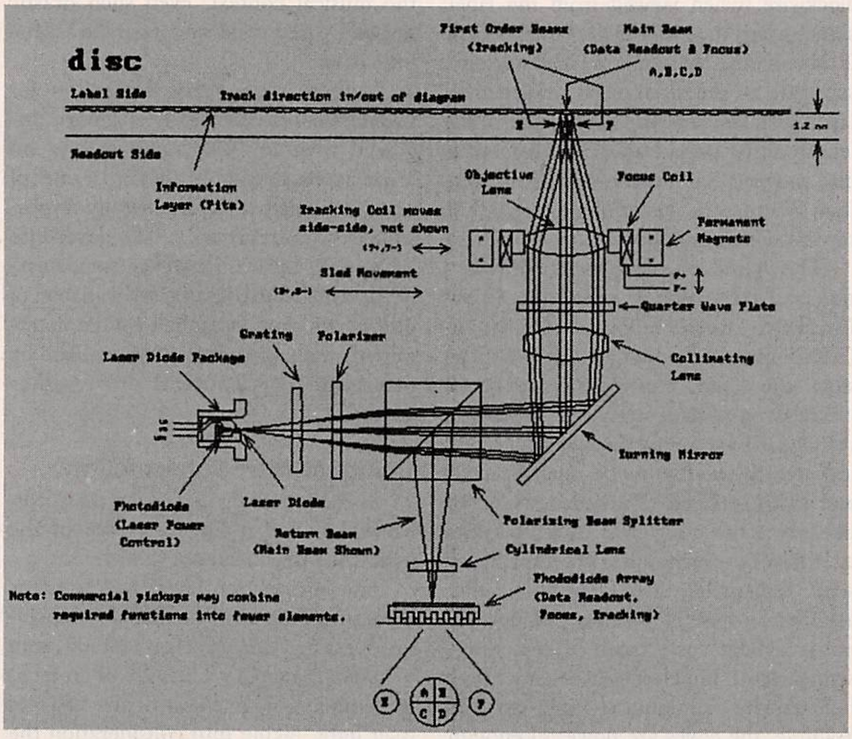

FIG. 1—HOW A TYPICAL THREE-BEAM PICKUP WORKS. Though there are many variations on this design, all pickups work in a similar fashion.

compared to 1 μm for the focused laser beam of a CD player!

At a constant linear velocity of about 1.2 meters per second, the required tracking precision is astounding. To put the required CD player servo-system performance into perspective, here is an analogy: Proper tracking of a CD is equivalent to driving down a 10-foot-wide highway (assuming an acceptable tracking error of less than +/-0.35 μm) for more than 3200 miles for one second of play, or over 14,400,000 miles for the entire disc without accidentally crossing lanes! And you thought that driving on a narrow winding country road was pretty tough!

Actually, it is worse than that: Focus must be maintained all this time to better than 1 μm as well (say, +/-0.5 μm). So, it is more like piloting a aircraft down a 10-foot-wide flight path at an altitude of about 12 miles (4 mm (typical) focal-length objective lens) with an altitude error of less than +/-7 feet while the target track below you is moving both 1 mile horizontally (CD and spindle run-out of 0.35 mm) and 3 miles vertically (disc warp and spindle wobble of up to 1 mm) per revolution! In addition, you are trying to ignore various types of garbage (smudges, fingerprints, fibers, dust, etc.) below you, which on this scale have mountain-sized dimensions. (Sorry for the mixed units, and my apologies to the rest of the world where the proper units are used for everything.)

The required precision seems unbelievable, but is just another day in the entertainment center for the CD player's servo systems. Even more surprising, this level of precision is achieved using mass-produced technology that dates to the late 1970s. And, don't forget that a properly functioning CD player is remarkably immune to small bumps and vibration—more so than an old style turntable.

Of course, we better hope that our technological skills are never lost—a phonograph record can be played using the thorn from a rosebush and a potter's wheel for a turntable. As you can see, there's just a bit more technology needed to read and interpret the contents of a CD!

Optical-Pickup Principles

The purpose of the optical pickup in a CD player, CD-ROM drive, or optical-disk drive, is to recover digital data from the encoded pits at the information layer of the optical medium. For CD players, the resulting data stream is converted into high-fidelity sound. For CD-ROMs or other optical-storage devices, it might be interpreted as program code, text, audio or video multimedia, color photographs, or other types of digital data.

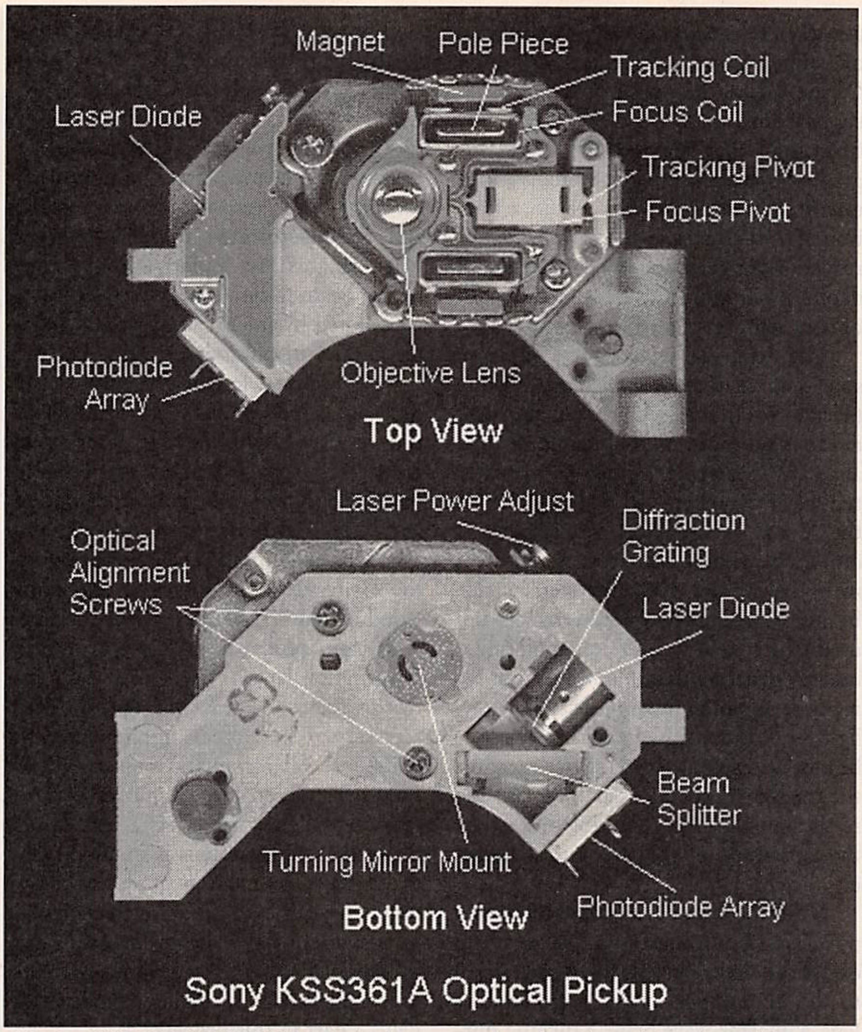

FIG. 2—THIS SONY KSS361 is a typical mass-produced modern pickup. It combines multiple functions into fewer distinct parts, resulting in improved robustness and lower cost.

The following (simplified) description of an optical pickup assumes a three-beam device—the most common type. A diagram of the pickup we will be discussing is shown in Fig. 1. To accomplish the same objectives, many variations on this design—such as a single-beam pickup—are possible. In addition, modern mass-produced pickups like the very common Sony KSS361A model shown in Fig. 2 combine multiple functions into fewer distinct parts. The result is improved robustness and lower cost. However, most of the basic operating principles are similar.

TABLE 1

Sync	(24 + 3)
Control and Display	(14 + 3)
Data	(12 × 2 × (14 + 3))
Error Correction	(4 × 2 × (14 + 3))
Total Bits/Frame	588

It is often stated that the laser beam in a CD player is like the stylus of a phonograph turntable. While this is a true statement, the actual magnitude of this achievement is usually overlooked. Consider that the phonograph stylus is electromechanical. Stylus positioning—analogous to tracking and focus in an optical pickup—is based on the stylus riding in the record's grooves and controlled by the suspension of the pickup-cartridge and tone arm. The analog audio is sensed most often by electromagnetic induction produced by the stylus's minute movements wiggling a magnet within a pair of sense coils.

The optical pickup, on the other hand, must perform all of those functions without any mechanical assistance from the CD. It is guided only by a fraction-of-a-mW of laser light and a few milligrams of silicon-based electronic circuitry. Furthermore, the precision involved is easily more than two orders of magnitude finer than that required for a phonograph. Sophisticated servo systems maintain focus and tracking to within a fraction of a micrometer of optimal. Data is read out by detecting the difference in depth of pits and lands of $^1/_4$-wavelength of laser light (about 0.15 µm in the CD)!

The laser beam is generated by a solid-state laser diode emitting at 780 nm (near IR). Optical power from the laser diode is no more than a couple of mW and exits in a wedge-shaped beam with a typical divergence of 10 × 30 degrees in the X and Y directions, respectively. A diffraction grating splits the beam into a main beam and two (first-order) side beams. (The higher-order beams are not used.) Note that the diffraction grating is used to generate multiple beams, not for its more common function of splitting up light into its constituent colors. The side beams are used for tracking and straddle the track that is being read. The tracking servo maintains this centering by keeping the amplitude of the two return beams equalized.

Next, the laser beam passes through a polarizing beam splitter (a type of prism or mirror that redirects the return beam to the photodiode array), a collimating lens, a quarter-wave plate, a turning mirror, and the objective lens before finally reaching the disc.

The collimating lens converts the diverging beam from the laser into a parallel beam. A turning mirror (optional, depending on the specific optical path used) then reflects the laser light up to the objective lens and focus/tracking actuators.

The objective lens is similar in many ways to a high-quality microscope objective lens. It is mounted on a platform that provides for movement in two directions. The actuators operate similarly to the voice coils in loudspeakers. Fixed permanent magnets provide the magnetic fields that the coils act upon. The focus actuator moves the lens up and down. The tracking actuator moves the coil in and out with respect to the disc center. The collimated laser beams (including the 2 side beams) pass through the objective lens and are focused to diffraction-limited spots on the information-pits layer of the disc (after passing through the 1.2 millimeters of clear polycarbonate plastic that forms the bulk of the disc).

The reflected beam retraces the original path up until it passes through the polarizing beam splitter, at which point it is diverted toward the photodiode array. (The polarizing beam splitter passes the horizontally polarized laser beam straight through. However, two passes—source and return—through the quarter-wave plate rotates the polarization of the return beam to be vertical instead, and it is reflected by the polarizing beam splitter toward the photodiode array.)

A cylindrical lens slightly alters the horizontal and vertical focal distances of the resulting spot on the photodiode array. The spot will then be perfectly circular only when the lens is positioned correctly. Too close or too far and the spot will be elliptical (*e.g.*, elongated on the 45-degree axis if too close, and elongated on the 135-degree axis if too far). The main return beam from the disc's information layer is used for servo control of focus and tracking, and for data recovery. The actual implementation could use an astigmatic objective lens rather than a separate cylindrical lens to reduce cost, but the effect is the same. Since the objective lens is molded plastic, it costs no more to mold an astigmatic lens (though grinding the original molds might have been a treat!). It is even possible that in some cases, the natural astigmatism of the laser diode itself plays a part in this process.

In essence, the optical pickup is an electronically steered and stabilized microscope that is extracting information from tracks $^1/_{20}$ the width of a human red blood cell while flying along at a linear velocity of 1.2 meters per second!

Now that you know everything (almost) there is to know about how CDs are made and work, we will wind up our theoretical discussion and go on to the good stuff, how to fix a player or a drive when it is broken. Tune in next time for our first CD troubleshooting segment. In the meantime if you have any problems or questions that just can't wait, go to my Web site at www.repairfaq.org. For questions to me, address them to me via e-mail at sam@stdavids.picker.com. (Note: While I would love to answer all your questions, regrettably, the finite number of microseconds in a day prevents me from being able to reply to letters sent via the postal service. However, I will respond to all e-mail requests in a timely manner—usually within 24 hours. Thanks for your understanding in this.)

See you next time. EN