TELEVISION: A COMPACT TUTORIAL

APPENDIX I

This appendix evolved from a tutorial, originally written for a nontechnical audience.[1] It is organized as follows:

I.1 Television: video plus audio

¶ **1.** Television involves the remote delivery of a moving picture plus sound. It is accurate to think of the *sound* as continuous; however, the *picture* is captured, and then delivered, as a succession of still images, at a rate fast enough that the viewer perceives a scene of continuous motion.[2]

¶ **2.** Television is distinguished further, of course, by the transmission of this movie-like content to the remote viewer. Originally this was carried out exclusively by terrestrial transmission, via radio waves, to the viewer's antenna and television set. Over time, other methods of transmission have been added – electrical cable,[3] optical fiber, direct satellite transmission via microwaves, and, of course, the Internet – along with recording methods such as magnetic videotape (Betamax, VHS, D-VHS), and optical discs (Laserdisc, VideoCD, DVD,[4] HD-DVD, Blu-Ray, and others).

I.1.1 The audio

¶ **3.** The *audio* portion of television is perhaps more easily understood, as it differs little from ordinary sound-recording techniques. A microphone converts the instantaneous sound pressure variations into an electrical signal; that is, it creates as its output an electrical voltage that at each moment is proportional to the pressure of the sound wave to which it is exposed. Contemporary audio recording and delivery usually employs two or more microphones, creating "stereo" sound (i.e., two channels), or multichannel sound (e.g., "5.1 channel sound").

¶ **4.** Traditionally these signals were processed, stored, and delivered by "analog" methods, which means simply that the signals were treated as smoothly varying voltages as they passed through the electronic innards of the amplifiers, recorders, modulators, and so on.[5] Contemporary "digital" technology does it differently: almost as soon as possible, the microphone's signal (the varying voltage that represents the sound) is converted to a succession of numbers (it is *digitized*), and everything that follows is some form of arithmetic on this torrent of numbers that come tumbling out. Only at the final stage – recreating the recorded sound for the listener – is the digital representation converted back to an analog voltage, and then, in the loudspeaker, to a reproduction of the original sound pressure wave.

¶ **5.** Just to give a sense of the quantity of numbers involved, in the standardized recording technology of the

[1] Following the stylistic conventions of that audience, the paragraphs are numbered sequentially.

[2] For conventional cinema-style movies, the rate is 24 frames/second; television in the United States uses a rate of approximately 30 frames/second.

[3] Known technically as *coaxial transmission line*.

[4] "Digital Versatile Disc."

[5] Common analog recording technologies, now nearly obsolete, include the vinyl record (where the audio signal waveform is embossed as small displacements of a fine groove), and the audio cassette tape (where the audio signal waveform is recorded as patterns of magnetization on a thin layer of magnetic oxide coating on a flexible plastic tape).

compact disc (CD), the instantaneous sound is *sampled* at a rate of 44,100 times per second (in both stereo channels simultaneously), and each such sample pair is converted ("digitized") to a 16-bit binary number.[6] So the bits are tumbling out at a rate of $2\times44{,}100\times16 = 1{,}411{,}200$ bits per second, or nearly 100 million bits per minute.[7]

¶ **6.** One might ask why any sane person would want to deal with such a quantity of numbers, when the original analog representation of the sound was so much simpler – just a pair of voltages that were varying at most 20,000 times per second.[8] The reasons are several, but they boil down to the contemporary ease and economy of digital processing, combined with the higher efficiency and quality of storage and transmission of audio (and video) that has been properly digitized. To get a sense of those advantages, one need only marvel at the gorgeous images transmitted daily from planetary probes visiting Mars and Saturn – images that are free of "snow" and other artifacts irreparably added to analog transmission by the effects of unavoidable electrical interference – to appreciate the benefits of error-free digital transmission. And, to get a sense of the density of digital storage, we note that a contemporary 5″ optical disc (dual-layer Blu-ray disc) holds 80 hours of CD-quality audio, or ten times that amount if modestly "compressed," compared with a mere hour's storage of analog audio on the 12″ vinyl LPs of yesteryear.[9]

I.1.2 The video

¶ **7.** The *video* is by far the more complicated part of television. The challenge is to reproduce a scene with motion, in color, while preserving adequate fidelity and absence of artifacts. And this must be done within the resources of the storage and delivery channels – that is, with finite disk storage and speed, and with finite transmission (via broadcast tower, cable, Internet, or satellite) *bandwidth*.[10]

¶ **8.** Video systems begin with a camera that has an electronic sensor (analogous to a digital camera) and that converts the two-dimensional color scene on that sensor into a succession of *frames*, each of which represents the image at those successive times (for US TV, the rate is approximately 30 frames per second). In traditional *analog* television, the two-dimensional image is converted into an electrical signal by the following method: imagine a single frame, that is, a still picture. To keep it simple, imagine further that it is monochrome; that is, "black and white."[11] We begin at the upper left, and move horizontally across the picture, generating an electrical voltage proportional to the brightness at each point as we pass by. When we reach the right-hand border, we jump back to the left edge, continuing with another horizontal path, slightly below the first. See Figure I.1. We continue in this way until we reach the bottom right-hand corner, at which time we have scanned the entire frame once, in what is known as a *raster* ("grid" in German) pattern.[12]

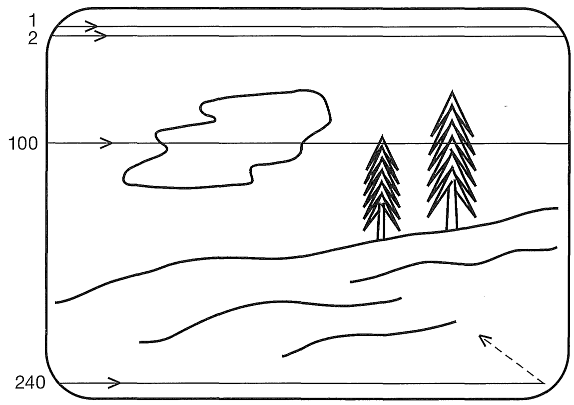

Figure I.1. A static two-dimensional image is "raster scanned" to create a video waveform (Figure I.2) representing intensity along the scan lines.

¶ **9.** What we have done, then, is to generate an elec-

[6] That is, a number ranging from −32768 to +32767, those bracketing the "full-scale" range of the recorded sound.

[7] The *recorded* bitrate is roughly triple this figure because of coding, error-correction redundancy, and the like.

[8] Or 20 kHz, the upper limit of human hearing; and that only for one of relative youth, such youth further possessed of sufficient wisdom to avoid deafening rock concerts.

[9] And a contemporary 3 terabyte 3″ magnetic hard drive that you can hold in your hand holds yet another factor of 60, or 50,000 hours of excellent quality (128 kbps AAC) compressed stereo audio; that's 15 years of 40-hour per week music!

[10] *Bandwidth* refers to the range of frequencies that can be carried on the cable or other transmission medium. It is technically accurate to think of this, for example, as the range of stations on the radio dial that could be carried with fidelity by a single electrical cable (or other medium). The term is sometimes used loosely to refer to rate of data transfer.

[11] Or, more accurately, *grayscale*.

[12] Traditional standard definition television (SDTV, usually called "NTSC," for National Television System Committee, and going back to the 1940s) in the United States divides the whole picture into 480 horizontal lines, along each of which roughly 640 features (picture elements, or *pixels*) could be resolved; a computer user would not be terribly impressed – he or she would say that standard NTSC television has only "VGA" resolution (i.e., 640×480).

trical representation, in time (a varying voltage proportional to brightness at each point in the image) of a single two-dimensional image; that is, we've converted a two-dimensional image into a one-dimensional output voltage. See Figure I.2.

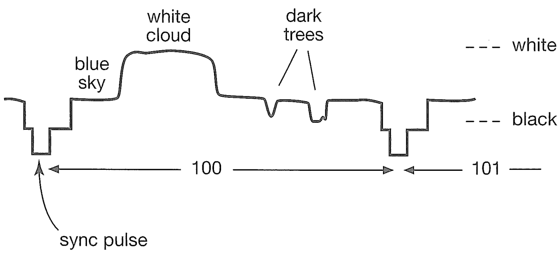

Figure I.2. A portion of the video waveform from Figure I.1, representing one of the 240 horizontal scan lines.

¶ **10.** This time-varying voltage is called the *video* signal, and it is the first step in creating a television image. In traditional NTSC analog television, this signal was transmitted by analog methods, after a process called *modulation* (more later), and was recovered and used by the television set to paint the picture on the screen, performing the same raster scan (left to right, top to bottom). Each frame follows in sequence, presenting a succession of 30 pictures per second on the television set's viewing screen.[13]

¶ **11.** To complete the video signal, it is necessary to add some synchronizing information, so the television set knows when to begin painting a frame, and also when each horizontal line begins. In traditional NTSC television this is done by adding a horizontal *sync pulse* at the beginning of each horizontal line, which is just a short[14] voltage pulse that, if it were in the middle of a picture, would be interpreted as "blacker than black." See Figure I.3 The television set detects these pulses and uses them to synchronize its scanning across each line. Likewise, a unique *vertical sync pulse* is transmitted for each field which informs the

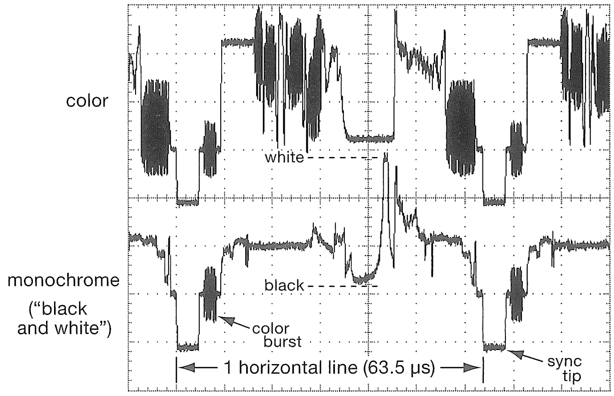

Figure I.3. Composite analog video signal of one horizontal line, framed by horizontal sync pulses. The brightness ("luminance") is represented by its amplitude. Color is accommodated by adding a modulated 3.58 MHz "chrominance" subcarrier, whose amplitude represents degree of saturation and whose phase encodes the colors.

television set when to return to the top to begin painting the next field or frame (see Figure I.4. The complete video picture signal, with its added sync pulses, is called *composite video*.

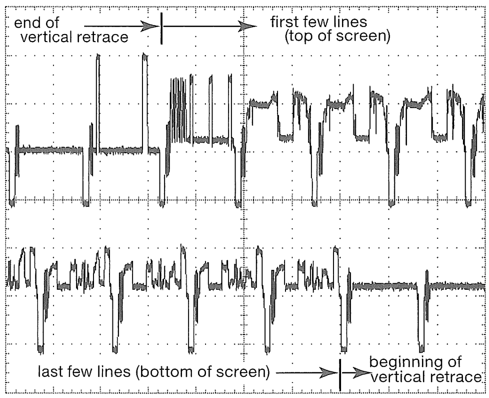

Figure I.4. The vertical retrace (beginning of a new field) is signaled by a set of tailored sync pulses, the first and last of which are shown here.

I.2 Combining and sending the audio + video: modulation

¶ **12.** Continuing for the moment with traditional NTSC television (as opposed to *digital* television, whose standards are known as *ATSC*, set by the Advanced Television Systems Committee, and which will be explained later), the composite video, along with the audio, must now be

[13] To complicate things, NTSC uses a method known as *interlacing*, in which a coarser raster – omitting every other horizontal line – is performed at twice the rate. Thus, in standard NTSC US television, the viewer sees 60 pictures ("fields") per second, each of which has only 240 horizontal lines; two such fields, with their interlaced lines, form one complete 480-line frame. This is sometimes called a "480i" format, to distinguish it from formats with higher resolution (e.g., high-definition TV, HDTV, with 720 lines or 1080 lines, or "4K Ultra HDTV," with 2160 lines), or from those without interlacing (which are known as *progressive*; e.g., *720p*).

[14] About 4.5 millionths of a second.

sent, via transmitting tower or cable, to the home viewer. Naively one might think of simply transmitting these signals "as is." This is not done, however, for at least two reasons: first, if the composite video signals were transmitted directly, then any two television signals would overlap and jam each other (because they would all share the same frequency band, namely that of the raw video signal itself); second, some wavelengths are more conveniently generated and propagated than others. For these reasons, the audio and video content of television signals (and, indeed, all communications and broadcast signals) are instead used to vary some aspect of a "carrier" wave, chosen at some specified wavelength. That carrier wavelength (or, equivalently, frequency) defines the "channel"; and the process of impressing the information (video and audio) onto the carrier wave is known as *modulation*.

¶ 13. Radio stations use the same technique: AM stations vary the strength (amplitude) of the carrier (hence "amplitude modulation"), whereas FM stations vary the frequency ("frequency modulation"). The carrier frequency itself defines the channel: in the US, AM stations are assigned to carrier frequencies between 520 and 1710 kilohertz (kHz, thousands of cycles per second), while FM is assigned to the band of carriers from 88 to 108 megahertz (MHz, millions of cycles per second). In the US, broadcast television begins at 54 MHz (Channel 2), and ends at 698 MHz (Channel 51), with gaps for FM, aeronautical, and other services.[15]

¶ 14. When information (video, for example) is modulated upon a carrier, the resultant signal spreads out and occupies a small band of frequencies. For example, when an FM station varies the frequency of its assigned carrier to carry its audio signal, the resulting signal occupies about 150 kHz. So FM stations are assigned channels separated by 200 kHz (to allow a "guard band" of 50 kHz in addition to their 150 kHz signal) – and that is why the frequencies of FM stations always end in an odd number after the decimal point (for example New York City's WNYC is at 93.9 MHz), ensuring a minimum spacing of 0.2 MHz (= 200 kHz).

¶ 15. Traditional analog NTSC broadcast television used a variant of AM for the picture signal (composite video) and, separately, FM for the sound signal.[16] The assigned TV channels are spaced apart by 6 MHz, each station being permitted to occupy nearly that amount, after allowing for a small guard band. Television sets "know" the frequencies allocated for each channel and tune to the correct frequency when the user chooses the channel number. For example, if (during television's analog era) one tuned to Channel 13 in the New York City area, the television set's electronics selected the station transmitting on 210 MHz (assigned by the FCC as Channel 13), namely WNET. The electronics in the set *demodulates* the received signal, recovering composite video and audio. The video, with its embedded sync signals, is used to paint the picture, frame after frame, while the audio is sent to the loudspeakers.[17]

¶ 16. Broadcast television (and radio) takes place on what is often called the "public airwaves." One needs only a television set (or radio) and an antenna to receive these over-the-air (OTA) public transmissions. Although some countries require licensing of receiving devices (radios and television sets), in the United States the broadcast services are freely available to anyone within range of a transmitting tower.

¶ 17. Depending on the distance and path from the broadcast station to the viewer, the "antenna" can be as simple as an indoor "bowtie" or pair of "rabbit ears," or as elaborate as a roof-mounted multi-element structure. Whatever its form, the antenna's function is to create an electrical signal on the feedline, induced by the speed-of-light broadcast signal passing by the antenna's site. Receiving antennas intended for broadcast television are designed to work over the range of frequencies used by broadcasters (see ¶13); thus the electrical signal delivered to the television set includes multiple stations, and it is the job of the TV tuner to select and process the channel to be viewed.

¶ 18. *Cable* television sends traditional analog TV channel signals in almost exactly the same way as broadcast. An evident difference, however, is that the channelized signals are received at the viewer's end from a coaxial cable (rather than being received by the viewer's television antenna), and then connected to the television set directly (i.e., to its normal antenna connector on the rear). Alternatively, for additional cable services (such as premium channels) the incoming cable connects to a "set-top box" (STB) provided by the cable company, the output of which is connected to the viewer's television set (or flat-screen monitor,

[15] You can download a gorgeous multicolor wall-sized spectrum allocation figure from `http://www.ntia.doc.gov/files/ntia/publications/spectrum_wall_chart_aug2011.pdf`.

[16] That is, the picture and sound signals are carried simultaneously on a pair of designated carrier frequencies within the single assigned television channel.

[17] In this primer we have ignored details associated with reproduction of *color* (versus black and white).

projector, etc.; see ¶56ff).[18] The channel *frequencies* are also somewhat different, with Channels 2–13 chosen the same as for terrestrial broadcast, but with the remaining channels reassigned to eliminate gaps; this can be done because the cable is a private world of its own, isolated so as not to interfere with, or be interfered by, other broadcast services.

¶ 19. A third difference is that some of the cable content is delivered as a subscription, for which the viewer pays additional fees; examples are premium services such as Home Box Office (HBO). These require some method for permitting or denying viewable delivery of selected channels or programs. Continuing for the moment with analog cable (whose days are numbered!), this can be done in several ways: the simplest is by installing filters (to block unsubscribed channel frequencies) at the utility pole, where the subscriber's cable splits off from the trunk running along the street;[19] A more sophisticated method involves scrambling the cable-borne analog signals of subscription programs,[20] and then instructing the set-top box (via digital communication from the cable provider to the individual STB) as to which programs may be unscrambled.

¶ 20. It is worth noting that cable companies have been required to carry the broadcast stations in their area, normally as analog cable channels.[21] Each such program occupies a cable channel (frequency). However, they may distribute additional services via digital methods ("digital cable") on additional channels (frequencies), which they much prefer; that is because, with digital methods, it is possible to carry up to *ten* NTSC-quality programs (i.e., SDTV, for Standard-definition TV) on a single channel. This is called *multicast*: the ability to carry multiple *programs* on a single channel (i.e., frequency). And, note that a cable can carry more than 100 such carriers – permitting more than 1000 simultaneous programs.

[18] For better picture quality, the latter connection is usually made not to the set's antenna input (called "RF," for Radio frequency, meaning the modulated channels discussed above, in ¶¶12–15), but to special audio–video inputs, with names like *s-video*, *component video*, *composite video*, *DVI*, or *HDMI*; see ¶64ff. The latter pair are *digital* connections, discussed below in connection with digital TV.

[19] Vintage cable subscribers will remember calling the cable company to add a movie channel, whereupon a cable truck appeared, the cable guy went up the pole to fiddle with something (changing the filter), and, voilà, movies on your television!

[20] For example, by suppressing the horizontal sync pulses or inverting the video (interchanging black and white).

[21] Unless all subscribers are provided with STBs that can receive digital delivery.

¶ 21. Previewing some additional characteristic of digital television: digital cable permits flexible subscriptions, with a program being authorized on-the-fly (e.g., pay-per-view, or video-on-demand). It also provides a natural mechanism for content protection via encryption. It allows for interactive participation, via a reverse channel back to the cable provider. It permits the delivery of high-definition content, with more than the 480 lines of NTSC (up to 1080 lines, at the highest quality currently supported). Finally, it provides a natural way to time-shift, pause, or replay live programs, via computer-type hard-disk storage.

¶ 22. Analog broadcast was sent into retirement in the United States in June of 2009, and all television broadcast delivery is now done by digital methods (more to follow). This conversion-to-digital process is taking place worldwide and will likely be complete by 2020 or earlier.

I.3 Recording analog-format broadcast or cable television

¶ 23. Video recording was complex and costly (and therefore confined to the broadcast studios) until 1975, when home video-recording devices were introduced in the US by Sony ("Betamax") and its competitors ("VHS," for video home system). These devices replicate the "front end" of a television set, to recover video and audio from the incoming signal (broadcast or cable), and use a clever spinning tape head arrangement[22] to capture on magnetic tape a reasonable replica of an NTSC analog television program. The technique is entirely analog (no digitization, no numbers) and records only onto special video-tape media (no computer media, no hard disks, etc.), as a magnetic recording (analogous to an analog audio tape recording; see the footnote at ¶4).

¶ 24. Videotape technology has been upstaged by digital alternatives such as optical disc recording (most famously in the form of DVDs and Blu-Ray discs – whether sold with prerecorded content, or recorded with a disk recorder), which creates a permanent copy of the video material; or by recording to a computer-type hard-disk drive (hdd), where the video copy is stored as a computer file. These digital methods require that the program material be converted from analog to digital form, if it is not already. (This is done internally and automatically in devices like TiVo®

[22] This is known as a "helical" tape head, which creates successive narrow slanted tracks across the slowly moving tape, each one holding one field of video. The use of a rapidly moving tape *head* eliminates the need for rapidly moving *tape*.

and other personal video recorders.) Digital television and digital video are discussed next.

I.4 Digital television: what is it?

¶ **25.** Just as an *audio* signal can be digitized (i.e., its instantaneous amplitude is measured, at rapid intervals, and converted to a succession of numbers), and subsequently transmitted, stored, or processed (¶¶4–6, above), so it is possible to digitize the *video* signal that represents successive frames of picture. Although one could imagine simply sending the digitized version of traditional NTSC as "digital TV," in practice one takes advantage of the enormous processing finesse of contemporary digital electronics to economize by *compressing* the raw video signal to a small fraction of its native size before it is delivered. The use of compression, along with the fact that a digital signal is "just numbers," permits the delivery of multiple programs on what would otherwise carry just a single video signal (program), typically by a factor of five to ten.

¶ **26.** There are several reasons for this improvement. One is the ability to detect and correct transmission errors by numerical techniques, allowing one to operate with received signal levels that are close to the "noise" (from interference, or signal loss owing to range or obstructions); with purely analog transmission a large received signal/noise ratio is necessary to reduce the visible effects of noise ("snow") to acceptable levels.

¶ **27.** A second reason is the spectral efficiency of digital transmission – or, more accurately, its improvement compared with the *inefficiency* of analog signaling. This can be seen in Figures I.5 and I.6, a pair of spectra taken directly off the home antenna of one of the authors in March of 2009, a time during the switchover to digital when both analog and digital broadcasts were taking place (see also Figure I.7).

¶ **28.** Compression aims to reduce by a large factor (tenfold or more) the quantity of numbers needed to describe the succession of picture frames, without significantly degrading the image quality. This seemingly impossible task takes advantage of redundancies in a moving picture, and of limitations in human visual perception.

¶ **29.** Contemporary digital-video compression is a rich and mathematically complex subject, the result of enormous effort in the applied mathematics and electrical engineering communities over the last several decades. But the basic tricks are easy enough to understand. The process begins by exploiting the fact that portions of an image near each other tend to be similar; so one can encode and send the (smaller) *differences* of brightness and color from a set of reference points, rather than the full description of brightness and color at each point. Likewise, successive frames tend to be similar, so one can define a sparse collection of index frames and send only the differences for intervening frames.[23] A further trick exploits the fact that the image often contains moving objects or a panning camera; so it is efficient to calculate "motion vectors" predicting

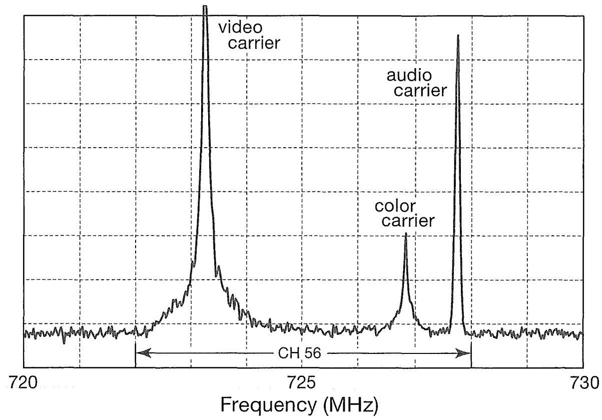

Figure I.5. The spectrum of 6 MHz-wide analog Channel 56 in Boston, as seen in May 2009. The video information resides in the sideband tails, while most of the transmitted power is wasted in the non-informational video carriers.

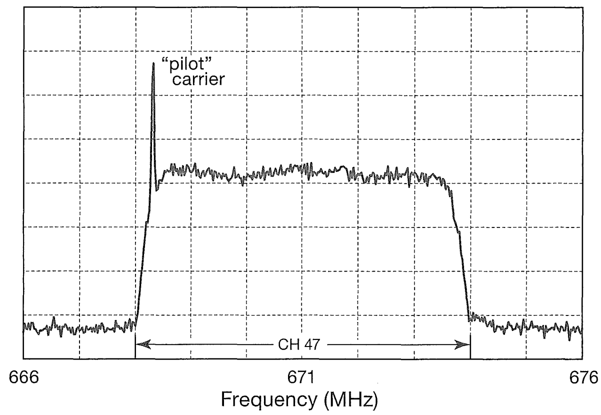

Figure I.6. Digital Channel 47, seen also in May 2009, fills its 6 MHz spectrum allocation with digitized video. It carries five times as many programs, with comparable (or better) picture quality.

[23] More precisely, it is the corrections from an *interpolated guess* between index frames (or reference points within a frame) that are sent.

the approximate motions, and then send only the (smaller) corrections from the predicted values.

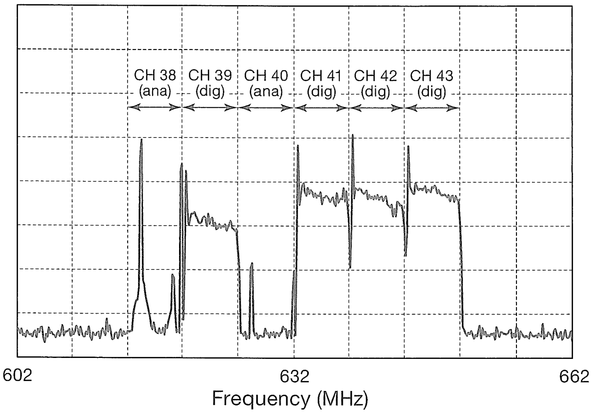

Figure I.7. A country in transition: RF spectrum of Channels 36–45 (each permitted 6 MHz of spectrum) as seen on our antenna feedline in Cambridge, Massachusetts, on May 6, 2009. Analog Channels 38 and 40 each carry one standard-definition (SDTV) program, whereas Digital Channels 39 and 41–43 can carry up to five SDTV programs each (though it's more common to see one HDTV and one SDTV program).

¶ **30.** These methods greatly reduce the needed bitrate (number of numbers per second), and they do so *without any loss of picture quality whatsoever* – they are "lossless." That is because the original digital image can be fully and exactly recovered by applying the numerical differences in the reverse order. However, further bitrate reduction is desirable (and often necessary), and this is accomplished by *lossy* compression. This consists essentially of discarding the less-important image information (from a psychovisual standpoint); the tradeoff is a somewhat degraded image (the degree of degradation depending on the degree of compression), which, however, can differ from the pristine original in ways that are hardly perceptible to the viewer.[24] The mathematics involves methods with names like discrete cosine transform, variable quantization, and Huffman coding; but the bottom line is that these methods permit a large reduction in bitrate with a relatively small reduction in perceived image quality.[25]

¶ **31.** The tradeoff of image quality with bitrate is gradual, and somewhere in the process a decision is made as to the desired final bitrate.[26] A major constraint is imposed by the fact that both digital broadcast and digital cable television in the US is sent on channels that conform to the same 6 MHz channel spacing that has been used for television since the 1940s. In practice (see ¶35, below) it is possible to send about 20 million bits per second (Mbps) on an over-the-air digital-broadcast channel, and nearly double that on a digital cable or satellite channel. A typical compressed bitrate for over-the-air NTSC-quality (SDTV) digital video is about 4 Mbps; thus digital-broadcast television stations are able to combine up to 5 or so NTSC-quality programs on a channel. (Recall that a frequency "channel" is no longer a single "program," because of multiplexing. More on this beginning at ¶33, below.) High-definition (HDTV) content requires nearly the full broadcast bitrate, so only one HDTV program can be broadcast on a channel. By contrast, cable or satellite systems, which are not constrained to MPEG-2 compression, are able to combine as many as eight HDTV programs onto one channel when using efficient H.264/MPEG-4 encoding. The second revision of the broadcast standard (ATSC 2.0) incorporates these more efficient codecs, as well as a host of transport and delivery enhancements that are aimed at mobile and interactive viewing, thus allowing over-the-air broadcasting to compete with services available on the internet.

¶ **32.** It is worth admiring the impressive bitrate reductions that these methods are achieving: a simple calculation[27] shows that digitizing an HDTV program without any compression would produce a bitrate of roughly 1000 Mbps, whereas contemporary compression methods reduce this to a modest (and deliverable) 20 Mbps, a 50-fold reduction! And comparable reductions are routinely achieved with SDTV.

[24] If such effects are noticeable, they are called *compression artifacts*; these are sometimes seen in over-compressed "jpeg" still photographs, as the patchy blocks or the "mosquito noise" around edges. Similar considerations apply to lossy audio compression, for example highly compressed "MP3" music files.

[25] The video-compression recipe currently used for all digital TV broadcasting in the US is named "MPEG-2" and described in the Advanced Television Systems Committee documents A/53 and A/54 (see www.atsc.org). An improved set of compression methods is incorporated in the set of standards known as MPEG-4; these are widely used by the cable and direct broadcast satellite services, as well as for video streaming over the internet.

[26] Which is permitted to vary, as program content changes. This is known as *variable bitrate*, or VBR, as distinguished from *constant bitrate*, or CBR.

[27] Bitrate $\approx$ 1080 lines $\times$ 1920 pixels/line $\times$ 30 frames/second $\times$ 16 bits/pixel = 995,328,000 bits/second.

I.5 Digital television: broadcast and cable delivery

¶ **33.** Over-the-air digital television broadcast and digital cable television both use traditional frequency "channels," upon which they put a stream of numbers (the compressed video described in ¶¶25–32, plus the associated digitized audio[28]), instead of the continuous analog waveform that was used in traditional NTSC television. Because of the economical bitrates produced by compression, there is adequate capacity on a single cable (or broadcast) channel frequency to accommodate several simultaneous programs. This is called multicasting, and permits up to four or five SDTV programs (a "multiplex") to be carried on a single broadcast channel frequency.[29] One can think of these as subchannels.[30]

¶ **34.** For either OTA digital broadcast or digital cable, the set-top box (STB) or equivalent hardware within the television set receives the multiple channel frequencies, each with its multiplex of programs. The STB or television knows the program assignments within each channel and is able to pull out the subchannel that the viewer selects, which it identifies by assigning a "virtual channel number." That is what the viewer chooses – it is displayed on the STB and on the screen during selection. For example, the viewer might select HBO, which is assigned a virtual channel (e.g. 82) and which might actually be just one of ten subchannel programs carried on one digital cable channel frequency. The STB then captures the HBO stream, decrypts and decodes its MPEG-2 encoding, and converts it to displayable video for a television monitor (or flat screen, etc.).

¶ **35.** In more detail, and in the language of digital television engineering, the delivery channel (digital broadcast or digital cable) is called the "transport stream," which can be thought of as a data pipe carrying some 20 Mbps (broadcast) or 38 Mbps (cable) in each frequency channel.[31] The ATSC specification dictates that the data put onto the transport stream must be broken into little *packets* of data, each of length 188 bytes and each belonging to an individual program. When multiple programs are sent on one transport stream, it is called a "multiprogram transport stream," or MPTS; if a single program, it's a "single-program transport stream" (SPTS). Repeating what was said earlier: a broadcaster can can put five standard-definition programs (or one HD and one SD program) onto a single broadcast channel's MPTS. The individual packets are identified by program, and they are interleaved in time (see Figure I.8).

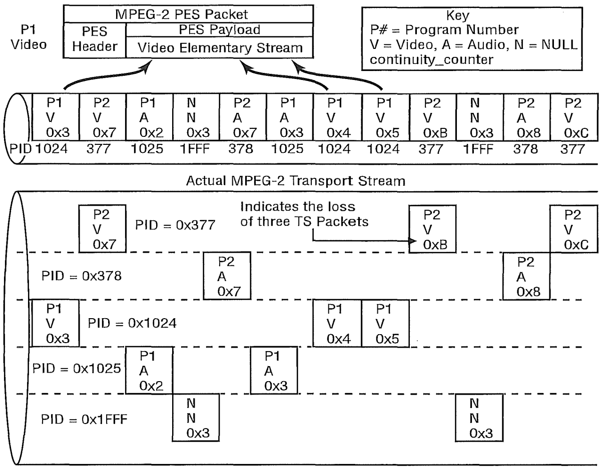

Figure I.8. Multiple programs can be interleaved into one digital transport stream, as a "multiprogram transport stream." Their individual video and audio packets are tagged with program identifiers (PIDs), by which they can be selected and reassembled (adapted from Figure 7.1 of ATSC Doc. A/54A, courtesy ATSC; readers are recommended to reference the current version of the standard or recommended practice available on the ATSC website).

¶ **36.** To put it another way, the several programs that will be put into a multiprogram transport stream (a "program multiplex") are cut into short pieces (packets, about 40 μs in length for digital cable), tagged with unique identifiers (called PIDs, for "program identifiers"), and then interleaved with the pieces (packets) from the other programs

[28] And also some "conditional access" (CA) information that enables legitimate subscribers to decrypt and view protected content; see ¶47.

[29] *Cable* delivery is more efficient, and permits as many as ten SDTV programs on a single channel.

[30] Because the cable (or broadcast channel) has a fixed total bitrate, the bitrates of the individual programs that are being multiplexed must be adjusted such that their total combined bitrate matches the channel capacity. This is called *bitrate grooming* and involves *null padding* (adding nulls, to increase a program's bitrate), on-the-fly compression (to reduce a program's bitrate), or even time shifting of program content (to prevent unlucky alignment of peak bitrates of the various programs). Digital television packets include "presentation time stamps," so it's OK to move things around a bit as they flow through the various digital pipes on their way to the television screen.

[31] The disparity has to do with the particular modulation schemes used: for broadcast it is called "8-VSB," whereas cable uses the more efficient "256-QAM" (pronounced "quahm"), thereby exploiting cable's better signal-carrying properties to carry roughly twice the information content.

that share the same transport stream. At the STB or television set the packets belonging to the selected program are identified ("filtered") by looking for their PIDs, and then reassembled into a single program transport stream to be decoded and displayed. In Figure I.8 there are two programs (P1 and P2), each with video (V) and audio (A), with their corresponding PIDs (1024, 1025, 377, 378); they are shown as the interleaved multiprogram stream at the top, and as filtered into their respective single-program streams at bottom.

I.6 Direct satellite television

¶ **37.** Satellites provide an alternative to over-the-air or cable/fiber delivery of television programming, and satellite delivery is particularly welcome in areas not served by wired broadband connections. This is known variously as direct-to-home (DTH), direct broadcast satellite (DBS), or broadcast satellite services (BSS), and exploits (usually) satellites in the geostationary constellation, i.e., in the equatorial "Clarke" orbit[32] of radius 42,200 km, where a satellite's period matches Earth's rotational period of 23h56m4s.[33] As surprising as it may seem, a single satellite with a ~100 W transmitter can deliver a half-dozen high-definition programs (or 30 "standard-definition" programs) simultaneously to small dish antennas on houses everywhere in the continental US. Typical direct-broadcast satellites are equipped with a dozen or more such "transponders," and the receiving dish antennas use multiple "feeds" (as many as four, for DirecTV or DISH Network) to point at several satellites, making available many hundreds of television programs.

¶ **38.** Early DBS systems used a 4 GHz (C-band) downlink, and required large dishes (>3m diameter) and expensive RF electronics. Contemporary systems operate around 12 GHz (Ku-band), with mass-produced oval receiving dishes (typically 0.6m×0.8m) that incorporate several low-noise RF amplifiers, each with local oscillator, downconverting mixer, and IF amplifier in an integrated LNBF (low-noise block-downconverter plus feed) unit that sits at the focus of the parabolic dish. That's an impressive amount of hardware for $100.

[32] See "Extra-terrestrial relays – can rocket stations give world-wide radio coverage?" [Sir] Arthur C. Clarke, *Wireless World*, October 1945.

[33] Ha! You thought it was 24 hours. Not so – we're in orbit around the sun, which adds an extra 4 minutes (1/365th part of a day) to Earth's true rotational period (the *sidereal* day) to arrive at the 24-hour *solar* day (the average time from solar noon to solar noon).

¶ **39.** One may be puzzled by the peculiar orientation of home dish antennas – why are they sometimes pointed so low that they seem to be aimed at or below the horizon? There are two parts to the answer: first, the constellation of geostationary satellites spans an arc across the southern sky, populated worldwide with some 200 satellites; over the longitudes of the US alone there are some 35 satellites parked[34] in geostationary orbit. The line of satellites across the southern sky dips down to the horizon at its eastern and western ends.[35] Second (and technologically more interesting), the geometrical arrangement of the receiving dish is what's called an "offset-feed paraboloid." That is, the little conical feeds are offset below so they do not block the incoming signal. This makes the dish appear to be pointing some 25° lower than it is, thus the explanation for the apparently "subterranean" satellites. This peculiar arrangement is used to eliminate blockage of incoming signal by the feedhorns, and also to reduce the encroachment of thermal radio noise that is emitted by the surrounding environment – another bit of thoughtful design that enables successful direct satellite broadcasting.

¶ **40.** The individual transponder RF channels are 27 MHz wide, with as many as 32 such transponders on a satellite, which (with guard bands) adds up to a total downlink bandwidth of ~1000 MHz. Satellite transponders typically use phase-shift keying (QPSK, 8PSK), with downlink bitrates of 40 Mbps per channel, adequate to deliver a half-dozen or more 1080i HD programs with efficient H.264/MPEG-4 encoding. The transponder channels are divided into two sets, one of each circular polarization at each transponder frequency, so that the total downlink spectrum occupied by a given satellite is about 500 MHz. This is downconverted in the LNBF to a pair of 500 MHz-wide IF bands, centered at 1.2 GHz and 1.9 GHz.

¶ **41.** Rather more hardware resources are expended on the *up*link, with steerable dishes of ~10 m-diameter class (Figure I.9) illuminating the uplink receiving dishes at the satellite, with transmitted power of the order of hundreds of watts per transponder; dishes of that size produce diffraction-limited beam diameters of about 0.15°

[34] Hardly "parked," of course – they are in equatorial orbits, whizzing around the Earth at nearly 7000 miles per hour, to keep up with Earth's rotation.

[35] You can get a good sense of the satellite arc with with mobile apps like DishPointer or Satellite Locator: when you point your smartphone into the sky it shows the satellites as bright red circles (with their locations) on a red arc, superimposed on the camera view (with obstructing trees, etc.) seen by the mobile device.

($\theta \approx \lambda / D$), small enough to prevent illumination of a satellite in an adjacent orbital slot, but requiring some active steering to maintain alignment on the desired satellite. Direct-broadcast providers such as DirecTV and DISH Network like to position their satellites close enough along the Clarke belt so that a single dish with several feeds can capture the downlinked signals from several of them; for example, both DirecTV and DISH Network offer an oval dish with three feeds, targeting satellites separated by 10° (currently at 110°W, 119°W, and 129°W longitudes).

Figure I.9. A few of the several dozen transmitting antennas at EchoStar's Network's uplink facility in Cheyenne, Wyoming.

¶ 42. You hook the coax cables from the receiving dish to a set-top box (usually with DVR), similar to what's used for cable TV, but which is designed to power the LNBF electronics up there on the roof, to select from the dish's several feeds and polarizations, and to receive the intermediate-frequency (IF) bands coming down from the dish on standard 75 Ω video coax (quad-shielded RG-6, usually).[36] As with cable TV, the content providers control your available programming, via subscription-enabled decryption. Satellite systems are not bound by over-the-air terrestrial standards (e.g., MPEG-2 encoding), and they generally use more efficient schemes such as MPEG-4. When compared with cable or fiber television service, the inability to target millions of subscribers individually (along with the absence of a reverse uplink channel to the satellite) limits the possibilities for interactive services such as video-on-demand.

[36] There may be an intervening module or two, with names like "multi-switch" or "node," to deal with selection of polarizations and IF-band stacking.

I.7 Digital video streaming over internet

¶ 43. With steady improvements in internet bandwidth (i.e., speed), it has become practical to deliver video (and associated audio) through the same internet infrastructure that serves personal computers and mobile devices (cellphones, tablets) with their email, web browsing, and so on. Some familiar examples of internet video streaming are news services such as CNN, government services such as NASA TV, movie streaming services such as Netflix and Hulu, and peer-to-peer services such as Skype. Just as with broadcast or cable delivery, the ultimate payload is a stream of numbers that constitute the video and audio content, in some efficient compressed format that takes advantage of sophisticated encoding schemes with names like "H.264" (also known as AVC, for advanced video coding), one of the current favorites. At the viewer's end the digital content is decoded[37] to recover the video and audio. For some services (e.g., Skype) a dedicated client program must be installed, whereas for others (e.g., NASA TV or CNN) a standard internet browser (such as Internet Explorer, Safari, or Firefox) suffices.

¶ 44. Compared with delivery via broadcast or cable, however, internet delivery of the time-critical video data presents some unique challenges. That is because data sent over the internet (via "Internet Protocol," IP) travels as independent packets of data, each typically some thousand bytes in length, and each including headers that specify its destination (IP address). Packets make their way through the multiply-connected nodes of the internet, and (usually) reach their destination. But there is no reserved pathway for a stream of packets (as there is for the "circuit-switched" architecture of the telephone system), and no guarantee of speedy delivery, sequential delivery, error-free delivery, or, indeed, of delivery of any sort. Various schemes are used to circumvent these evident deficiencies, for example by requesting retransmission of missing packets (they are numbered) or of corrupted packets (they include error-revealing "checksums"). These work well, and a data file that is downloaded via "TCP/IP"[38] is effectively guaranteed to be bit perfect.

[37] Hence "codec," a contraction for coder–decoder, usually appended to the name of the compression scheme, e.g., "the H.264 codec."

[38] Transmission Control Protocol/Internet Protocol, a universally used standard for transfer of data that needs to be transported and reconstituted without error.

¶ **45.** Because of this disorganized "packet switching" of internet data, streaming video may suffer interruptions or intervals during which the average delivery rate is reduced. For this reason it is common for the receiving end to store ("buffer") a few seconds of video beyond what is currently being displayed; and if the internet delivery speed is inadequate (as evidenced by buffer underruns), the sender will reduce the data rate (thus delivering lower-quality video). Contemporary broadcast-quality high-definition TV (known as 1080i, meaning that the picture consists of 1080×1920 pixels, delivered 30 frames/sec) requires upward of 2 megabits/sec of download speed, available now in most homes with a broadband internet connection (cable, optical fiber, or telephone-line DSL[39]). That speed is well within the capability of wireless (Wi-Fi) connections, so high-definition video can be delivered to mobile devices such as laptop computers and tablets.

I.8 Digital cable: premium services and conditional access

¶ **46.** Cable television providers offer premium services, such as Showtime, HBO, and pay-per-view, for which the subscriber pays additional monthly fees. A subscriber whose subscription includes Showtime, for example, is able to view (and record, if the STB includes a DVR) programming delivered on Showtime channels. The cable provider needs a method to control each subscriber's access to the full channel lineup. Although each customer has a cable coming into the home (and therefore one could imagine that different content is sent to each home), in fact the same signal is sent to all the homes within a neighborhood group, known as a "service group."[40]

¶ **47.** To limit access within the full suite of distributed programming, the cable provider includes *conditional access* information along with the video and audio. This is done by including "CA" packets, along with the usual V (video) and A (audio) packets that comprise one program within the multiprogram transport stream. The STB includes decryption hardware, and uses the CA packets to provide the key information to unlock the encryption that is imposed on the video and audio packets by the cable company.

[39] Digital subscriber line, technology for bidirectional digital transport over analog telephone lines.

[40] See ¶49, below, for more details.

I.8.1 Digital cable: video-on-demand

¶ **48.** How is it possible to provide program material specifically for an individual subscriber, for example "on-demand" delivery of a previously broadcast program, or of a movie, on a cable network that serves an entire city?[41] Such services go by names like "video-on-demand" (VOD), and are made possible by the fact that the cable provider is able to tailor the actual signals carried on its cables that go to different groups of cable subscribers in an area.

¶ **49.** A cable network in a metropolitan area is more complex than one might at first imagine: rather than a city-wide distribution of common program material, the network is organized into smaller groupings of *nodes* and *service groups*. A service group consists of segments of cable, typically running past no more than 500 homes, carrying identical material; those signals are fed into the cable at a "node" in the neighborhood, which in turn is fed via optical fiber from a more distant "hub" at which the provider inserts the group of channels that is to go to that particular service group.

¶ **50.** The trick to providing individual on-demand material, then, is first to ensure that there are at least a few extra channels (in addition to the standard lineup) available to carry such content; and, second, to divide up the city into many service groups, so that those extra channels can carry a different suite of on-demand content to the different service groups. For example, suppose there are five channels available for on-demand material and that each can carry ten programs (as described earlier); if a service group includes 500 houses, of which 200 are cable subscribers, then the cable provider can satisfy 25% of those subscribers with simultaneous on-demand programming (because its five extra channels, each sending ten custom programs, deliver 50 simultaneous programs).

¶ **51.** On-demand programming requires also a reverse channel for each subscriber, so that the subscriber can select programming, and also pause (or fast-forward or rewind) the material. However, these reverse channels need carry only a few simple commands up to the provider (as opposed to the high-bitrate video coming down in response), and are easily accommodated in an "interactive" uplink signaling band of the cable network.

[41] The term *unicast* is sometimes used to distinguish such individual delivery from *broadcast*.

I.8.2 Digital cable: switched broadcast

¶ **52.** Some cable providers use another service that exploits the flexibility inherent in the cable network's use of separate service groups, namely "switched broadcast." Switched broadcast delivers programs only when requested by viewers, on a service-group basis, as compared with delivering all programming to all customers. This allows the cable provider to offer more programming choices than could be carried simultaneously on the available number of channels.[42]

¶ **53.** Although switched broadcast and VOD use similar methods to deliver their content to a subscriber, it is worth noting a difference: switched broadcast content is delivered whenever that program material is normally scheduled, and not at the whim of the subscriber; what distinguishes it from normal program material is that it is not put onto the service group's cable at all, if no one in that group has tuned to the program. Once a subscriber has tuned that program (causing it to be sent to that service group, on a particular subchannel), that same program is present on the cable serving any additional subscribers in that service group, on the same subchannel. By contrast, VOD is delivered on-demand, at the time requested by the subscriber; likewise, it can be paused, or fast-forwarded, etc. (tasks that are performed on-the-fly by the cable provider, not by the set-top box). That is possible because VOD content is being sent on a particular subchannel on the subscriber's service group only, and is available for the requesting subscriber only (via encryption and user-specific "entitlement control," in the case of protected content). The terms *narrowcast* and *unicast* are sometimes used to refer to service-group-specific cable delivery of these two varieties of content, namely scheduled ("linear") material, and interactive user-specific on-demand material, respectively.

I.9 Recording digital television

¶ **54.** The conversion of analog material (audio and video of the real world) into digital form is hard work – but it makes the task of *recording* straightforward. That is because a single program received at the STB is, in essence, just a stream of numbers, which can be filtered from the multiprogram stream, assembled in a temporary memory "buffer," and written to a hard-disk file just like any computer file. In that sense, a set-top box with DVR is simply a special-purpose computer, with the usual processor, hard drive, etc., and having additional hardware to do the special STB tasks – receive the cable signal, generate the displayable output, take control commands from the infrared remote "clicker," and so on. A typical contemporary set-top box with DVR contains a dual-processor chip, 64 MB program memory, a 160 GB (minimum) hard drive, and various video-related additional hardware (input tuner, video memory, display and audio drivers, etc.). Along with buffering and storing the video and audio content, the STB controls the access, decryption, and re-encryption of the (protected) video content.

¶ **55.** It is worth noting that any digital storage medium of adequate speed and capacity can be used to store digital video content; at the consumer level there are many "personal video recorders" (PVRs) that store programs on recordable DVDs or onto solid-state "flash" memory chips. There are also digital video tape recorders that can store both SDTV and HDTV onto a digital variant of VHS tape.

I.10 Display technology

¶ **56.** For over half a century television images were displayed with a cathode-ray tube (CRT), in which electrons emitted from a hot cathode and accelerated to potentials of kilovolts were deflected (usually magnetically) to paint a raster, at the video frame rate, on a phosphor surface coated on the interior of the viewing face of the evacuated tube. The intensity was modulated with a grid electrode near the cathode. Early CRTs were monochrome (black and white); color tubes used arrays (triads or stripes) of red, green, and blue phosphors, aligned with a metallic mask so that electrons from each of three electron-emitting cathodes (or, in Sony's Trinitron, steered from a single electron gun) struck only one color of phosphor.

¶ **57.** Cathode-ray tube displays worked; but they were heavy (over 100 lbs for a TV with 32″ display), bulky, and required fussy "convergence" adjustment to get the colors to track. Shortening the tube to reduce cabinet depth exacerbated the convergence and geometry problems. CRTs are now obsolete, replaced by several technologies, among them liquid-crystal displays (LCDs), plasma displays, and organic light-emitting diodes (OLEDs).

¶ **58.** In an LCD a liquid-crystal layer is sandwiched between a pair of crossed optical polarizers; an applied electric field alters the polarizing properties of the layer, thus varying the optical transmission. In the classic display there's a uniform white rear illumination (from white LEDs, or from one or more cold-cathode fluorescent lamps – CCFLs – combined with diffusers and light pipe

[42] This is sometimes called *narrowcast*, versus broadcast or unicast.

material); an array of electrodes applies local electric fields to the image pixels, which are overlaid with red, green, and blue color filters. The array of liquid-crystal pixels act as video-rate dimmable shutters; all the light originates in the rear illuminator.

¶ 59. LCDs are dominant in computer displays and popular in televisions. They are thin (a centimeter or so) and bright. But they have somewhat limited dynamic range (or contrast ratio: ratio of maximum brightness to "maximum darkness") and speed, and their color balance and black level degrades when viewed off-axis. There have been great improvements in speed and in off-axis performance, owing to methods with names like in-plane switching (IPS), fringe-field switching (FFS), and the like. And the dynamic range can be improved by using LED-array backlighting, which can be dimmed locally and rapidly, adapting to the light and dark areas of the displayed image.[43]

¶ 60. For the most realistic rendering of cinematic material, however, the plasma display is superior to the LCD. It consists of an array of tiny cells, in each of which a switchable gas discharge generates ultraviolet light that causes a spot of phosphor to glow. A high-definition display of 1080×1920 pixels has three such cells for each pixel (one each with red, green, and blue phosphors) to generate the pixel's overall emitted color. Unlike an LCD (which filters an underlying light source), the cells of a plasma display generate the emitted light directly. You can think of it as an array of 6 million little CRTs ($1080\times1920\times3$), each one time-switched to achieve the required light intensity.

¶ 61. Plasma displays retain their color fidelity and contrast ratio regardless of viewing angle, and they have fast response. In larger sizes they are currently somewhat less expensive than LCDs. They are not as bright as LCDs, however, and a static pattern that is displayed for a long time can cause some image retention, or (in extreme cases) phosphor "burn." Their contrast ratio is very good, but not infinite, because the gas discharge in every pixel must be sustained at a low level (i.e., it cannot be switched off entirely) so that it can be rapidly modulated.

¶ 62. LCD and Plasma have dominated display technology, but the future is likely to be ruled by OLED (organic LED), a direct-emitting array of tiny LEDs (either in three colors, or white LEDs with filters). These emerged first in small displays (e.g., cellphones and camera viewfinders), but by 2014 they had made the big time, with full 4K "Ultra HD" resolution (3840×2160), and screen sizes to 65″ and beyond. Unlike earlier display technologies, OLEDs can be made flexible, and the current fad is *curved* TV screens (capable of 3D, if you're interested in that). OLED is likely to be the ultimate winner, because of their very high dynamic range (1,000,000:1), wide viewing angle, low power consumption, elegant form factor (3 mm thick, lightweight, and nearly borderless), and potential manufacture by an inkjet-like process. These have been winning Best of Show awards, and they deserve it.[44]

¶ 63. Two other technologies that looked good, but have fallen onto hard times, are field-emission display (FED) and surface-conduction electron-emitter display (SED). Both involve an array of phosphor cells (like the plasma display), but with electron (rather than ultraviolet) excitation of the phosphor. An SED prototype from Canon generated great enthusiasm[45] at the 2006 CES, but subsequent patent disputes and economic realities took their toll. FED and SED may rise again – but don't hold your breath.

I.11 Video connections: analog (composite, component) and digital (HDMI/DVI, DisplayPort)

¶ 64. What are all those incompatible cables, anyway? To add to the analog–digital confusion, the *consumer* market (of large-screen TVs and flat panels, etc.) and the *computer* market (of LCD monitors) have gone their (mostly) separate ways. Here's a quick rundown of the most-used connections; their connectors are pictured in Figure I.10.

¶ 65. In the *consumer television* world there are four types of connections (and connectors), the first two of which are nearly extinct.

Composite video. Low quality standard-definition (SDTV: 480i) analog video, recognizable by the yellow RCA-type connector ("phono jack"), usually bundled with an audio pair (red = right, white = left). The single yellow line carries bandwidth-limited luminance

[43] For marketing purposes these LCDs are sometimes called "LED TVs." Don't be fooled: it's an LCD, but with LED rear illumination replacing the CCFL. And it may or may not have local dimming – read the fine print.

[44] The reviewers are *gushing*: "The best picture I've seen on any TV, ever." (CNET), "The best-looking TV I've ever seen. Ever." (Digital Trends), and "The best direct-view display – of any size, at any price – we've ever laid eyes on." (HDTVtext, UK).

[45] A breathless review in SlashGear (19Oct2006) proclaimed "SED-TV is something that no amount of words can describe. It is something that must be SEEN to be believed; literally." and "SED-TV is the future of digital image displays; it's as simple as that."

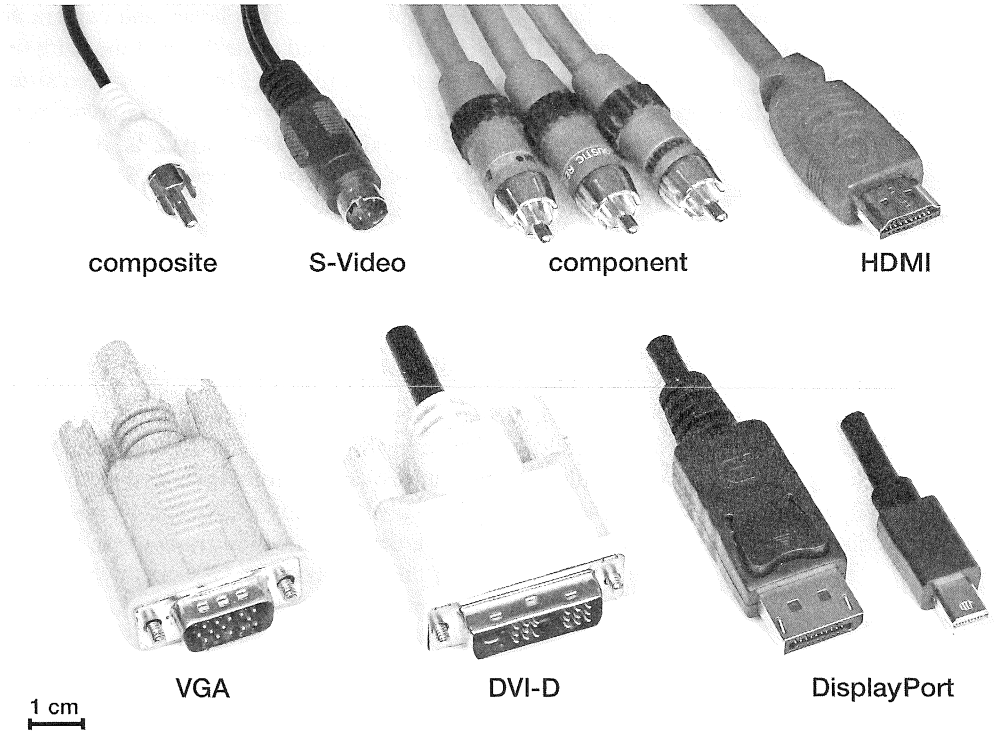

Figure I.10. Cable connectors for computer monitor and video signals. Top row: TV video connectors, from the legacy analog composite to contemporary HDMI high-definition digital video. Bottom row: computer monitor connectors, from the legacy analog VGA to the popular digital DVI and newer DisplayPort (standard and mini connectors).

("luma," abbreviated Y, essentially the grayscale image) and chrominance ("chroma," abbreviated C, the color difference signal pair that is modulated on a 3.58 MHz color subcarrier), along with line and frame synchronization pulses; it is sometimes called CVBS (composite video, blanking, and sync). Don't use such a connection unless there's nothing else available!

S-video. Somewhat better SDTV analog video, recognizable by the 4-pin miniDIN connector with fragile pins. It separates the luma (+ sync) and chroma signals, retaining more bandwidth. Avoid this one, too, unless you like fuzzy pictures.

Component video. Now we're talking! This analog format uses three 75 Ω coax lines, with (usually) RCA-type connectors (or, occasionally, BNCs) that are colored red, green, and blue, and that can carry full-bandwidth HDTV. The colors are misleading: the green line carries luminance (+ sync), while blue and red carry color *difference* signals. That is, the GBR-colored connectors denote "YPbPr," where Y is luminance (red+blue+green), Pb is blue-minus-Y, and Pr is red-minus-Y. As with composite and S-video, component video carries only video; the audio needs it own cables. Component video does not know, or care, about things like content protection; for that reason it is not embraced by content providers, who may prevent full-resolution HDTV (1080×1920) output on the component jacks (e.g., on a Blu-Ray player).

HDMI. A purely digital format, whose initials stand for high-definition multimedia interface. HDMI is the digital alternative to component video. It is recognizable by the flat 19-pin USB-like connector (sadly without any required latching mechanism), and electrically it is equivalent to the DVI-I computer-monitor format (see below). It carries both audio (up to eight channels, digitized to 24 bits, up to 192 ksps) and video (digitized at 8–16 bits per component, at rates adequate for full HDTV ("1080p," 1080×1920 progressive at 60 Hz; HDMI versions 1.4 and later support full "4K," 4096×2160, with 60 fps progressive from version 2.0). The video data is a digitized version of analog video: uncompressed numerical data representing the amplitude of the color components, sent over four twisted pairs (R, G, B, clock). HDMI supports digital content protection (HDCP, high-bandwidth digital content protection), a protocol by which an HD video source authenticates a display device before sending (encrypted) data, so you are allowed to view the full-resolution video. HDCP seems to work, most of the time anyway (though you may get annoying messages and glitches).

¶ **66.** In the *computer monitor* world there are three types of connections (and connectors) in wide use.

Analog VGA. Legacy analog link, recognizable by the 15-pin D-type connector with locking screws. VGA (for video graphics array) carries separate RGB analog signals, plus Hsync and Vsync (thus "RGBHV"). In contemporary implementations it also has an I^2C channel for monitor identification and control. VGA will work up to resolutions of 1600×1200 or so (there's no specified limit, but you'll see smearing going rightward of sharp features when pushing the resolution), but many monitors have abandoned VGA altogether in favor of the digital formats: DVI and DisplayPort.

DVI. Currently the standard digital interface, recognizable by the 29-pin (maximum) connector with locking screws, looking somewhat like a "VGA on steroids." It is electrically similar to the more compact and inexpensive HDMI, above, that evolved from it for the consumer television market, and that includes audio (thus a single connection from a cable box or Blu-Ray player to the television monitor). It comes in several variants, all using the same connector (in which some pins may not be loaded): DVI-D is video-only, and comes in single-link and dual-link varieties (the latter needed for resolutions greater than 1920×1200 at 60 Hz, for example the 2560×1600 at 60 Hz used in 30″ displays); DVI-A is analog video only (for compatibility with analog monitors); and DVI-I ("integrated") has both digital (single- or dual-link) and analog video. DVI, like VGA, carries no audio.

DisplayPort. Newer standard, intended to supersede DVI; it uses a 20-pin USB-like connector with latching mechanism, and supports very high data rates (up to 4.3 Gbps on each of four differential pairs, thus 17.3 Gbps). It departs from the "digitized-raster" scheme of DVI/HDMI, using instead a packetized data protocol; but it's got enough bandwidth to handle the full video bandwidth of dual-link DVI (which can be converted to DisplayPort protocol). It supports up to 16 bits per color, and 8-channel audio at 24 bits and 192 ksps. It includes provision for fiber optics (instead of copper) for long cable runs (to 50 m or more), and it supports both existing 56-bit HDCP and its own DPCP (DisplayPort content protection, with stronger 128-bit AES encryption). The current revision of DisplayPort can handle 4K 60 fps progressive with ease.