Chapter 1 How and Why I got involved in measuring flutter

Learning the Basics from Jack Mullin at 3M

When I arrived at 3M as a new college graduate in 1965, my bosses, Jack Mullin and Ken Clunis, were the unquestioned masters of low-noise, low-flutter audio recorders. The pilot run of 6 of the original 3M Dynatrack machines was already complete, and tape recording would never be the same. Those machines extended the dynamic range of audio recording by over 10 dB and reduced the flutter content by about an order of magnitude. This was also, I believe, the first silicon planar transistor audio recorder. In addition to low flutter, the transport featured full dynamic braking with motion sensing.

A key instrument in the lab was the D&R FL-14 flutter meter that featured not only mechanical flutter measurements out to 300 Hz, but also scrape flutter capability out to 5 kHz. With this meter we could look at and tame most of the flutter components due to mechanical and scrape excitations.

The D&R meter included a very handy accessory - small built-in CRT display. In addition to the normal time based sweeps, the 'scope also accepted an external frequency to create a Lissajous or x-y pattern. The display would show a partially static pattern of an ellipse or line when the frequency of the external oscillator was adjusted to match one of the frequency components of the flutter signal. This technique, using the x-y mode of any 'scope, works well to analyze the composition of the flutter signal.

Around 1968 we also acquired a new flutter meter - the Micom Model 8100W. This unit couldn't measure anything beyond 250 Hz, but it did include a tunable wave analyzer that could be used to isolate the frequency components of the flutter signal. This technique was more convenient that the x-y method, but the finite bandwidth of the tunable filter could allow more than one frequency component through to the meter. Overall, the Micom meter was quite a step forward in flutter measurement.

The Micom line had two other meters for measuring instrumentation machines to the standards set by IRIG - the Intermediate Range Instrumentation Group that set the telemetry standards for rocket scientists. These meters had a wide range of test frequencies and measurement bandwidths. Unfortunately, the nicer unit, the 8300, had a mechanical interlock that locked out the combination of a 12.5 kHz carrier and 5 kHz analysis bandwidth. All of the Micom gear was expensive, and the IRIG models were way beyond the budget of normal recording studios.

We used our flutter meters quite a bit at 3M to solve design and manufacturing problems. For example, the new Model 23 recorders utilized a rubber-tired flywheel rather than the belt drive used in the Dynatrack machines. We had gotten Walter White, the expert at reconditioning Ampex 300 flywheels, to build us a rubber-tired flywheel. The performance was excellent - the best ever seen on a 3M Isoloop. The rubber added some damping to the drive train, thereby reducing both the flutter components and the acoustic noise coming from the motor. The resilient coupling also allowed a small range of speed adjustment.

The new flywheel was designed into the new product, but we were in for a big surprise when we tried to buy production quantities of the flywheel. Walter White wasn't an affordable source for large quantities of flywheels, so we tried the rubber vendors in L.A. None of them could make a tire that didn't have a once-around thump.

The tires for the flywheels were cut from a long cylinder of rubber. The cylinder started with a metal core or mandrel of the desired inner diameter, and then a thin sheet of rubber was wrapped around the mandrel until the multiple layers reached the desired thickness. The wrapped mandrel was then vulcanized and slit into rings for attachment to the flywheel. Even though the tire on the flywheel went through a final grinding step to achieve the exact desired outer diameter, the discontinuity in the rubber that was formed by the edge of the first layer would show up in our flutter measurements. We evaluated dozens and dozens of flywheels, but never saw anything that was as good as Walter White's original prototype.

I believe that someone must have known how to make what we needed, but we never got hooked up with the proper source. The flywheels continued to be a major problem for 3M as long as that design was around.

Reverting to the old Dynatrack configuration utilizing a belt drive eventually skirted the problem. The original design had two problems - the seam in the belt and dirt buildup. The early belts were made from the ribbon material used for brassiere straps. The seam was a butt splice stitched across the belt. The seam always produced a measurable, but not necessarily unacceptable, flutter thump.

(Jack Mullin had a humorous story about trying to find someone in the New York City garment district to make an emergency replacement for a belt that failed while Jack was demonstrating the Dynatrack prototype at Columbia or RCA in New York.)

The other problem with the belt was dirt pickup. A seamless Mylar belt had been tried as a replacement for the brassiere strap with excellent flutter results, but dirt clumps would build up on the drive pulley on the motor. Dirt hadn't been a big problem with the fabric belt, maybe because the belt was porous and the dirt could migrate through the belt. It doesn't take much dirt to introduce a once-around thump at the 30-Hertz rate of the motor. To make matter worse, the plastic belt made a good Van de Graff static electricity generator that would attract dirt particles.

Ken Clunis conducted extensive studies on how to stop the dirt from building up. His solution included several aspects. He utilized a chrome-plated steel flywheel rather than the brass flywheel utilized in the Dynatrack machines. For the motor drive spud, he chose a hard-anodized aluminum pulley that was ground to final size after installation on the motor shaft. And then, for good measure, he enclosed the entire belt assembly in a shroud to keep dirt away. The combination completely eliminated the dirt clumps!

We didn't change the Model 23's to replace the rubber tire with the Mylar belt, but we did include the new technology in our newer prototypes, including the new two-inch transport design.

We had another flutter crisis on the Model 23 that involved the reversing idler. The Isoloop (a name coined from `Isolated Loop') tape path presses the tape against both sides of a fairly large capstan. A ridge and groove pattern on the Isoloop capstan and complementary patterns on the ingoing and outgoing pinch rollers meter the tape into the Isoloop area at a rate slightly slower than the tape exiting the loop. This small difference in entry and exit speeds creates the tape tension across the heads. A `reversing idler' with tape wrapped around 180 degrees of its surface reverses the path of the tape and sends it back to the capstan and outgoing pinch roller for the exiting pass.

The reversing idler is a critical flutter component since it is inside the tight loop. Any flutter due to the idler is transmitted directly to the record and playback heads on either side of the idler. The idler, which is about 2" in diameter, was designed to minimize this flutter.

Ball bearings must be properly `preloaded' to give long bearing life and quiet operation. (Preloading is applying a force along the axis of the bearing to guarantee that the rolling balls stay in contact with the inner and outer races at all times. Wavy spring washers and springs are commonly used for preloading.) The 3M reversing idler design utilized a rubber O-ring to provide both the preload and some vibration damping. The ring was placed under the outer bearing flange of the upper bearing. The rubber acted as a damper to absorb some of the mechanical resonances of the idler, reducing the overall flutter.

Our problems started when one of the bearing suppliers offered to make the entire idler assembly for us rather than just supplying us with bearings. `Who should know more about precision rotating assemblies than a bearing manufacturer?' we thought. Well, we learned a very expensive lesson!!

When the first idlers were received from the vendor, we tested them in the lab. Flutter was terrible, and they howled like a banshee at fast winding speeds. You could even feel `bumps' when you rotated the idler by hand. To make a long and painful story short, the vendor had change the preload technique. In fact, he had completely eliminated the compliant preload altogether! His technique was to screw down the retaining screw on the bearing `just the right amount'. Well, first of all, how much was just enough? If you went just a little too far, the balls would flatten or dent the races - a problem called Brinneling. And if, by some stroke of good luck, you got a good preload, what would happen as the temperature changed? The idler was aluminum, but the support shaft was stainless steel. The difference in thermal expansion coefficients will never hold the preload constant as the temperature varies. And, of course, there are severe problems if the rigidly locked idler is bumped or dropped.

It took a long time to educate the vendor and solve this problem. We tried hand-selecting idlers, but even that was not `kosher' since the design had so many flaws. We eventually changed the design slightly in the 2" transport to avoid problems.

While on the subject of the reversing idler, let me mention a few experiments to further improve flutter. The underside of the idler was hollowed out to reduce the inertia of the idler, allowing it to accelerate and decelerate with the tape without slipping. Unfortunately, the idler would ring like a bell when it was tapped. If the demodulated output of the flutter meter is fed to a monitor speaker, you can hear this ringing if you tap the idler with the rubber eraser of a pencil.

We conducted a series of experiments to try to control this ringing. The hollowed out ring was moved to the top of the idler to form a `bathtub' into which we could place damping material. We tried Brillo pads and oil, pennies and oil, and brass disks cut to simulate pennies. (The real pennies out performed the brass copies!) We got some significant improvements in flutter, but the practical considerations such as vertical mounting and leaking oil doomed the project.

I didn't get heavily involved in the transport design until we started working on the 2" versions. I guess I was the champion of the transition to 2", so Jack Mullin let me have a free reign. Our prototype for 2" investigations was made from many of the standard Model 23 parts. Our first 2" reversing idler, which I still have in my collection, is solid aluminum with the standard anodized coating, but it has no bleed grooves to reduce the entrapment of air between the idler and the tape at high speeds. You could bring the deck up to full rewind speed and then stop the reversing idler with your finger and spin it in the opposite direction! What a bizarre sight to see the tape moving 600 ips in one direction and the idler coasting in the other direction!

We needed stronger reel motors for eight-pound reels of 2" tape. The small Bodine NSH-14 motors were barely capable of 1" operation, so they was not an option for the wider 2" tape. Fortunately, the Mincom plant also manufactured large instrumentation recorders that handled 14" reels of 1" tape. Those transports utilized a similar motor, but in the bigger NSH-34 frame. The only way to make the bigger motors fit between the webs of our audio transport casting was to shave the sides off the motor. The motor had a heavy cast iron shell that held everything together. I determined that we could cut completely through the sides of this shell if we just didn't cut into the four screws that held the end caps onto the barrel. The barrel would become just two separate arcs, but the motor was still one solid piece.

The machinist assigned to shave the motor didn't believe the scheme would work. He was convinced that the whole motor would fall apart when he made the cut that broke through the barrel. He dragged me out to the machine shop to watch as he made the final cut so that I could observe the catastrophe first-hand. He cut through the side and - nothing happened. He vigorously poked at the motor to prove his point and - nothing happened. I just turned and walked away.

We used those cut-down motors for quite a while. When we took that prototype deck to the Spring AES Convention in 1968, we covered up the gaping holes in the sides of the motors with masking tape, and then sprayed the tape with black Krylon paint to hide the sin.

The 2" deck required reinforcement to the pinch rollers. Don Kahn, our mechanical designer, devised a support yoke for the top of the idler that included a fine adjustment of pressure at the top. We also changed to a simplified reel hub that had no lock-down mechanism. In retrospect, I think we should have provided something to hold the reel centered on that hub.

Off to Graduate School

When I left 3M 4 years later in 1969 to attend graduate school at New Mexico State University, the one tool that I couldn't find in the university lab was a good flutter meter. I did take a course in active filter design, so I undertook a special-projects class to build a flutter meter that would include a 5-pole active filter. I chose the Signetics 560 PLL (phase locked loop) chip as the FM demodulator.

Since I wanted to measure all the FM components out to 5 kHz, I needed a carrier that was more than twice that frequency. I chose a 12.5 kHz test tone as a reasonable frequency for machines running at professional speeds. (The signal must contain not only the carrier, but also the sidebands produced by the flutter. In this case the first order sidebands are at 12.5 kHz +/- 5 kHz, requiring good response to 17.5 kHz. This is reasonable at 15 and 30 ips, and marginal at 7.5 ips.)

I built a prototype that had the basic flutter measurement capabilities for rotational and scrape flutter. The report was entitled ` A one transistor flutter meter' since I had only one transistor in the box - along with a dozen IC's!!.

I was so proud of the unit that I offered to share my schematics with 3M's Mincom Division, the new owners of the Bahrs/Micom line of flutter meters, and allow 3M to test the unit for a couple hundred dollars. They took the unit and sent me a check, along with a document that claimed exclusive rights to my product in return for the token amount I had negotiated for a 2-week evaluation!! Fortunately, I still had some friends at 3M who cut through this bit of subterfuge and got me my money without compromising my rights.

Adventures at Ampex

After grad school I went to Ampex to be the product manager of the MM1100, AG440C, and the start of the ATR100 project. While there, Alastair Heaslett and I used my prototype flutter meter to measure just how terrible the MM1100 was after we were badly trounced by a 3M M79 in a fly off at Wally Heider's studio. The scrape flutter of the MM1100 without any scrape flutter idlers (the original configuration of the machine) was 10 times worse than the M79. Nobody at Ampex at that time believed that scrape flutter was important enough to give it a second thought - until we got trounced.

The Beginnings of Altair Electronics

After 3 years with Ampex and another year bouncing around in Hollywood, I wound up teaching EE at Cal State University at Northridge. I spent a lot of time on campus in my Audio Lab, so I decided to pull out the flutter meter and turn it into a real product. In the ensuing years I had learned that although flutter was important, there were other sidebands produced by magnetic recorders that were also important to the cleanness of the sounds that are recorded and reproduced. The missing piece was the AM sidebands produced by variations in the magnetic tape coating and variations in the contact between the heads and the tape.

Fortunately, the Signetics 560 PLL that I had used in New Mexico had a companion product, the 561, that could simultaneously decode AM and FM sidebands. The 561 featured two phase comparators to accomplish this job. (The 561 mysteriously disappeared from the Signetics product line after it was found to be an ideal chip for making bootleg TV descramblers. Many of these schemes used suppressed carrier techniques that were easily decoded by the 561.) I redesigned the meter to offer both FM and AM measurements, the first and only instrument ever offered to do both jobs.

(It is very important to keep the FM and AM sidebands properly separated, since they are caused by very separate and distinct phenomena on the tape recorder. A spectrum analyzer typically does not track the phase of these individual sidebands, and therefore isn't able to distinguish bad tape from flutter.)

My company's name was Altair Electronics, and I called the box the Altair Tape and Transport Diagnostic System (T squared DS). The unit offers FM (flutter) measurements with filters for Weighted, Rotational, Scrape and Wideband measurements, and AM (modulation noise and dropouts) with the same band limiting filters. The unit has a crystal-referenced source with both FM and AM modulation capability for calibration purposes. The metering can function as either NAB (average responding/RMS calibrated) or DIN peak. Residual FM noise is typically .007% with the widest filter, and .0005% for the Weighted filter, well below any recorder I have every met. The highest line in the noise spectrum is 70 dB below the highest full-scale meter range of 1.0%. AM noise is similarly well below the typical readings of tape machines.

I actively sold the unit for a number of years, but Altair Electronics was flushed down the toilet in a divorce settlement. Yes, Fred, I do still have a few units available at $985. They are identical to the original units. The instruction manual you asked about is a cassette tape that walks you through the theory, setup and actual testing of some tape machines. I recorded parts of the tape while I was conducting tests on a typical transport (not an Ampex!!) Some of the original information is available on my website at www.manquen.com. (Update: This website is not currently available.)

When Audio Precision began to consider adding flutter to their System One, they kept running into fanatics like Neil Muncy who were using my box. My customers swore by the need to measure scrape flutter so strongly that, after a couple of years, AP decided to include my `highband' flutter measurement techniques in their flutter option. They did not, however, understand that the AM portion was also very valuable, and chose to ignore the AM function.

Why were these customers so outspoken? Because this instrument allows you to `see' what is really going on with both the FM and the AM sidebands that degrade the purity of the program material. The wideband nature removes the artificial and misleading measurement standards such as `Weighted Flutter' that lead people to ignore some of the most useful information because the machine's specifications are very restricted. And if anyone wants to doubt that this information is critical to how clean a tape machine sounds, we can conduct a poll here on the list server to see what people who have dealt with these problems have to say.

Chapter 2 High Resolution Flutter Analysis with MARCAP

The demodulated output from a flutter meter is a very complex set of (mostly) audible signals. All of the speed anomalies are transformed into frequency components. The frequency of an individual component represents the frequency of the vibration or speed variation that created the flutter component. Since most irregularities are not clean sinusoidal variations, the flutter sources generate an array of harmonics just as square waves, sawtooth waves, etc. The first great crime of flutter measurement (Sorry, Jay!) is to use a meter to read flutter. This is like saying that you can pick out the harmonies in a specific hit tune by watching only a VU meter. The meter is useful only as a go/no go gauge. It is almost useless as a diagnostic tool. (By diagnostic tool, I mean something that tells you how well the parts of a device is performing. When you go to a mechanic to have your car diagnosed, you don't want him to say "Well the engine runs, so it must be OK. You want to know that the ignition, valves, alternator, etc. are each performing properly.)

If your mechanic is an old timer, he might pull on the throttle linkage a few times to gun the engine and then tell you "Sounds like you have one plug missing, the timing belt is slapping, and one of the valve lifters is sticking." He can do this because he is using a very complex spectral analysis device called the human hearing system. As Neil Muncy keeps stressing, listening to the demodulated output of a flutter meter takes advantage of the incredible power of our hearing. Our hearing can distinguish the individual major components and even include the complex harmonic structure of each by distinguishing between thumps, clicks and other combinations.

I used the phrase `mostly audible signals' above because flutter extends to very low frequencies. Most transports can generate components as low as .5 Hz. The rotation of the reels is in this speed range, so a dragging brake or eccentricity of the reel holddown (which I have measured) will show up as subsonic frequencies. A common trick to extend your frequency response is to place your fingers on the cone of the speaker hooked to the demodulated output. This isn't great, but better than nothing.

After building a number of Altair boxes, I embarked upon the next step of my journey into flutterdom. My ears allowed me to qualitatively evaluate a tape transport, but I wanted quantitative signatures. In fact, I wanted repeatable signatures that could be saved and compared with prior measurements of the same machine to detect gradual degradation. With a high-resolution picture, I would be able to pick out the signature of each rotating component on a deck and tell if that component was OK. My effort grew into a suite of recorder analysis techniques that I named MARCAP - Magnetic Analog Recorder Comprehensive Analysis Program.

The key element of MARCAP was an interface with my flutter meter that provided high-resolution FFT analysis and display of the flutter components. This interface included a programmable 12-bit A/D with antialias filter to encode the flutter meter's demodulated output, a control system to vary the sampling rate so that I could measure and display bandwidths from 5 kHz down to 2 Hz with zoom capability, and all the software to control, capture, convert, graph, save, print/plot, and compare signatures.

At the time, I was using a Hewlett Packard 9836 computer. I started with the A/D interface. My control program was written in HP Rocky Mountain BASIC with high-speed subroutines written in PASCAL. Some of the FFT code came from my work with Deane Jensen's COMTRAN program. (I wrote interface modules for Deane to acquire data from some of the HP waveform recorders and spectrum analyzers. Deane also used my machine for a while to duplicate his release copies on the 5 ¼" disks used on the 9845.)

Once I got the program completely implemented, I went looking for a market. (This was very dumb `reverse marketing'!!) Since everyone in the U.S. was busy abandoning the manufacture of analog recorders, I didn't find many takers. So I spent time trying to convince operations people that they could diagnose their recorders periodically as preventive maintenance to avoid catastrophic failures.

The best `bite' I got was from Modern Video. They had me conduct periodic analysis of their Sony 1" helical video machines to look for evidence of progressive wear. We were able to see gradual changes to things like the pinch roller. They chose not to continue because I couldn't offer them an affordable package of their own. (My HP 9836 with an 8 MHz 68000 and 1 MB of RAM was about $15,000 and then there was my hardware and software. Where was today's PC when I needed it!!

MARCAP contained additional modules that would model the wavelength and equalization effects of recorders and control/equalize a flux loop.

I eventually gave up on converting the world regarding flutter and went on to other activities that lead to joining the design team for the Flying Faders fader automation system.

Flutter sources

Disclaimer: I am a designer of tape recorders; therefore, I approach the problem of flutter from the engineering side. What components are making flutter and why? How can each flutter component be minimized? A user might take an entirely different approach, wondering only if something is broken and needs fixing.

Rule #1

The best time to check the flutter of a machine (and the AM content) is BEFORE you buy it. This rule is normally applied to new equipment. The reasoning is thus: Many of the sources of flutter are inherent in the design of the machine. Scrape flutter, for example, is primarily a geometry problem. You can't easily change the geometry of the tape path if one day you decide you want lower scrape flutter.

Corollary #1

If you are designing a tape transport, most of the flutter characteristics will be set by your early conception of the machine's basic layout. If you want low scrape flutter, choose a design that generates low scrape flutter. Don't try to add a bunch of Band-Aids to correct the problem later.

Flutter isn't a static phenomenon like harmonic distortion in an amplifier. We have not only time-varying flutter effects, but we also have compounding of the playback flutter on top of the record flutter. Let's take that a bit slower&

We find several types of speed/timing irregularities on a tape deck. The most obvious is eccentricities of round rotating things in the tape path. Our biggest fear is that the speed-determining components, usually the capstan and pinch roller, will have eccentricities that will try to pull/meter the tape at fluctuating velocities:

1. We could have a capstan that is not concentric. For example, the 3M Isoloop capstan assembly had a large steel shaft with a 2" diameter anodized aluminum capstan. The concentricity of the bearings, steel shaft, and the inner and outer diameters of the capstan all can introduce concentricity errors. Add to this the required tolerances between the shaft and bearings and the shaft and capstan so that the parts can slide together. A common technique is to `true up' the capstan surface after all the bearings and parts are assembled to minimize eccentricity or `run-out'.

2. The shaft could be bent. This wasn't likely on the 1" shaft of the Isoloop, but it is quite possible on the slender capstan motor shafts of many of the Ampex designs.

3. The bearing supporting the capstan can be loose. The Ampex direct drive hysteresis synchronous motors use a bushing at the top, nearest the tape, because a bushing usually has less radial play than a ball bearing. Unfortunately, that bushing also has a very high side load due to the pinch roller pressing against the capstan. When the bushing wears, the shaft may hunt around in the enlarged bearing hole. The shaft may repeatedly `climb' to a different position in the hole and then fall back, like trying to climb a wall. This erratic motion will not show up as a simple harmonic of the shaft rotation speed. Don't neglect the lubrication for the capstan and pinch roller bushings!!

4. The traction surface that pulls the tape may be spotty. If one side of the capstan face has better traction than the other side, the bite of the capstan will vary as the capstan turns.

5. We may have a dirt buildup on the capstan that changes the effective diameter as the capstan rotates.

Eccentricities usually show up as a flutter spectrum line at the rotational rate of the eccentric device. (I wonder what my frequency is? I am certainly eccentric!) We may also see harmonics of the primary rate.

Next, we may have `rattlings' that are due to parts rattling and banging around. The worn bushing in #3 above is an example. The special case of ball bearings will be covered below. We also have brushes clicking inside DC motors, vibration from fans, and sticky spots on pinch rollers due to splicing tape adhesive. We might also have an oscillation in the capstan or reeling servos. In the flutter spectrum, these rattlings produce many frequency components, usually with a dense set of harmonic since the motion isn't nicely sinusoidal.

Last we have pure random phenomena. Scrape flutter is a prime example of a random excitation. The scraping produces a wide range of frequencies, just like the sound of sandpaper on wood. The band of random frequencies may show a peaked shape due to the tuning effect of the span of tape that is vibrating.

Now we will consider the special case of ball bearings. A ball bearing consists of a number of `round' steel balls that roll between an inner and outer ground-steel race. The bearings are kept evenly spaced by a ball retainer that moves with the balls. How many irregularities can we have? We can have

1. Eccentricities of both the race and the outer surface of the outer race,

2. Eccentricities of both the race and inner bore of the inner race,
3. The roundness of each and every ball,
4. The smoothness of the inner and outer race surfaces where the balls roll,
5. The retainer or balls striking a deformed bearing shields that covers the side of the bearing,
6. Irregular distribution of the oil or grease used to internally lubricate the bearing, and
7. Erratic motion of the ball cage as is moves around.

We can get many frequencies from a ball bearing. We have the rotational speed of the shaft, the rotational rate of the balls, and the rate at which the balls progress around the bearing.

One particular hazard for ball bearings is a shock or tap on the bearing that can cause a dent in the race or a flat spot on a ball. This is called Brinnelling. This will cause thumping as balls hit the dent or the flat on the ball hits the race. This acts like repeated shocks or spikes that produce flutter pulses with lots of harmonics.

We can calculate the various rates for a given ball bearing by measuring the diameters of the inner and outer races, counting the number of balls, and then applying the following formulae:

(Do is the outer race diameter, Di is the inner race diameter, and n is the number of balls)

If we have a flatten ball, the disturbance rate will be 2*Do*Di/(Do^2+Di^2) times per revolution.

If we have a dent in the outer race, balls will strike it n*Di/(Do+Di) times per revolution.

If we have a dent in the inner race, balls will strike it n*Do/(Do+Di) times per revolution.

The ball retainer migrates Do/(Do+Di) of a revolution for each revolution of the bearing.

As an example, the bearing used in the tension sensor arms of the ATR100 has the following values:

Flattened ball rate = 1.8 thumps/rev   Typically 1.5-4

Outer race dent rate = 2.96 thumps/rev   Typically 2.5-5.5

Inner race dent rate = 5.04 thumps/rev   Typically 4-8

Retainer migration rate = 5/8 turn/rev   Typically .5-.75

Another problem with ball bearings is that they require a preload. For a ball bearing to run smoothly and last a long time the balls must be in firm contact with the inner and outer races at all times. To assure this rolling contact, a preload force is applied between the inner and outer races to clamp the balls into place. Various kinds of wavy washers and spring devices are used to create the force. The preload mechanism must have some compliance or give so that temperature variations and normal operation and handling will not overload or unload the bearing. I believe that at least one of the Ampex audio capstan servo motors contains an internal assembly with some springs to maintain preload. If we don't have preload, the bearing will be loose and cause rattling (and premature failure.)

Frequencies of flutter components

Rotating components generate flutter frequencies that are determined by the rotational speed of the component. This rotational speed is calculated by measuring the diameter of the part and then dividing the circumference of the part (pi times diameter) into the tape velocity. For example, at 15 in/s (AKA ips) a 3/8" scrape flutter idler will turn 12.73 times a second. (Actually, the neutral axis of the tape runs at the designated linear speed. Add one tape thickness to the diameter for a rough estimate of the neutral axis at the middle of the tape thickness. How thick is the tape? A typical 1.5 mil backing usually gives a tape that is about 2.0 to 2.1 mils thick. (Am I correct, Bill Lund?) So our idler would probably run at 12.665 revs/sec. (The neutral axis calculation is much more involved that what I have shown here. It depends upon the detailed composition of the `belt', the amount of wrap and other factors. In our case we are dealing with a layered `belt' that has very different strengths and stiffnesses for the polyester backing and the magnetic coating.)

Measuring Flutter

So now we are ready to measure and characterize the variations in speed produced by all these defects. Flutter can be defined as time base instabilities that create frequency modulation byproducts that corrupt our audio signal. The frequency modulation of the flutter process produces new sideband tones that were not in the original music. The frequencies of the sidebands are combinations of the music frequencies and the flutter frequencies. For example, a 1 kHz tone fed through a tape recorder that has capstan flutter at 20 Hz (a 1200 RPM motor) will have new sideband components at 1 kHz plus and minus 20Hz or 980 Hz and 1020 Hz. (Since tape machine flutter is typically less than .1%, our system is referred to as a very narrow-band FM system, leading to simplifications in some of the describing equations. In our case, it means that we need worry about only the first-order sidebands described above when the flutter content is low.)

We want a method to pick out the disturbing sideband components while ignoring the intentional test tone. The tool for this job is an FM discriminator, a device that locks to a fixed frequency test tone. The discriminator eliminates the `carrier' test tone and recovers only the sidebands. In the demodulation process, all of the sidebands are folded around the carrier and shifted down to the `baseband' spectrum starting at zero Hertz instead of the carrier's frequency. The folding process combines the upper and lower sidebands created by a specific modulating frequency. In our example case, the two sidebands - each caused by the same 20 Hz modulation - are combined into one baseband component at 20 Hz.

Our preferred test technique is to reproduce a constant frequency tone that was recorded on a `flutter free' recorder and measure any frequency modulation due to the playback machine. The frequency components of the demodulated signal will be an `image' of the speed disturbances due to flutter.

Life would be quite simple if we could pull our handy dandy MRL flutter-free flutter test tape from the shelf and play it on our machine. Unfortunately, recording the flutter tape leaves the flutter signature of the recording machine. If you really want to measure only your machine, the flutter of the test tape recorder must be somewhere between 1/3rd to 1/10th the values of your machine. This is not possible with the high quality of our better analog audio recorders. There are other cases, however, where the flutter on the test tape can be adequately low to permit basic measurements.

(Once we invoke the high-resolution analysis of the FFT technique, a pre-recorded `standard' tape would introduce all of the spectral components of the original recorder. Some day we might be able to have a tape that has a floppy disk attached that contains the FFT spectrum of the test tape so that we can subtract the test tape's components from the measurement made on our machine.)

If we have a high-quality recorder, our solution is to use our machine to record and play back the tape. But this really opens a can of worms. When you record the tape, the eccentricities produce speed variations. If we play the tape back, we might wind up with the same eccentricities in exactly the same phase as when we recorded. This would give us not net speed error. Or on the other hand all of the eccentricities may start out on the opposite side of the bump, causing twice the speed error. So each time we rewind the tape, we play Russian roulette with how the record and playback components add up.

"Well why not make the test while recording?" I heard someone in the back row ask. This introduces a time delay equal to the transit time from the record to the playback head. If the circumference of an eccentric roller evenly divides into the head-to-head distance, the eccentricity will cancel out completely. So Rule #1 is: Don't use record mode for flutter testing if you want to be assured you measure all the flutter components.

In a simple flutter meter, all of the demodulated flutter components are lumped together and fed to the analog meter. We can determine the total flutter, but this doesn't give us much insight into the true sources of the flutter.

Fortunately, we have various methods of segregating the flutter components. The simplest method is to use bandpass filters to carve out chunks of the flutter spectrum. Typical filters may pass all of the primary rotational components (.5 - 250 Hz), all of the components due to vibrations in the tape (250 - 5k Hz), or the infamous `weighting' curve (a narrowband filter with the peak at 3 Hz) which supposedly judges the objectionability of flutter.

A more sophisticated technique is to use a tunable narrowband filter to look at specific frequencies. This allows fine surgery, but we loose the overall view while in the narrowband mode. Several flutter meters include a tunable wave analyzer.

The ultimate wave analyzer would allow us to look at all of the individual frequency components at the same time. A real time analyzer is a meter of this type, but the resolution of the display doesn't give us the very tight groupings that we require. A Fourier Transform (more correctly, a DFT or Discrete Fourier Transform since we have sampled data) can give us this fine detail. The DFT spectrum consists of frequency `bins' or subdivisions of the total frequency range. As we take more time samples, we can have more subdivisions and finer detail. I worked with DFT's with 1024 to 16,384 samples. The tradeoff is that more samples take more time to acquire and more processing power/time to crunch the numbers. We also control the sampling rate, which determines how long it takes to acquire our string of samples. The upper frequency limit of a DFT is the sampling rate/2.

Let's say we want to cover the full output range of an Altair flutter meter - 5 kHz. We need to sample at more than twice this rate or more than 10 ksamples/sec. In fact, to make the antialias filtering not too bad, we could choose 15 ksamples/sec. This gives us a transition band of 2.5 kHz between the Nyquist limit of 7.5 kHz for our sampling and the 5 kHz of our desired data. The antialias filter must drop any and all flutter and noise components below the floor of our sampling resolution before we hit 7.5 kHz. This is much looser than the requirement for a CD antialias filter with 20 kHz audio and 22.05 kHz Nyquist limit.

If we choose to use a sample string with 1024 data points, we will achieve a frequency resolution of (sampling frequency/number of samples) or 15 ksps/1024samples = 14.65 Hz/bin. Doubling the number of points to 2048 will improve our frequency resolution to 7.32 Hz/bin.

The number of samples we take is a multi-edged sword. Lots of samples improve our resolution, but how do we display things on a VGA computer screen with 640 x 480 pixels? One answer is to include a Zoom function that allows us to expand specific parts of the spectrum.

But life is not quite that simple. We must remember that the DFT assumes that this waveform repeats the picture captured by our string of samples forever. But we are dealing with flutter that looks more like random data than a nicely repeating waveform. The solution is to `window' our sampled data to make it look like it is part of a repeating string. In the audio world this would be like starting with the fader down at the beginning of the sample string. Then we `slowly' push the fader up and pull it back down, arriving at the fader bottom again at exactly the end of the sample string.

The process of windowing distorts the data in a way that changes our resolution. After all, we calculated our resolution based upon a full string of samples, but then we badly attenuated the beginning and end of the string. Not fair!! As a result, the DFT looses some of its accuracy for the size and exact frequency of the components. We get some `leakage' between adjacent bins. The shape of the windowing - how fast we push the fader up and down - determines the amount of distortion. There are many, many possible windowing shapes with different tradeoffs of amplitude and frequency accuracy.

I used a `Hanning' or `raised cosine window' window. Each input data point is multiplied by a number between 0 and 1 that rise from 0 to 1 and fall back to 0 (like our fader) with values ½[1-cos(2*pi*n/(N-1))] where n is the number of the sample between 0 and N, the number of samples. The math isn't important here. Just think `Push up fader, pull down fader.'

(The reason I am giving all the gory details is because some of you will be trying this with your sound cards and FFT software. Done properly, the results will be very useful. Done improperly, the results will be a mess. For the rest of you, just bear with me&)

So we are ready to take data. We set up a sampling rate determined by the span of our desired spectrum. We select the number of samples based upon our computer resources and computation time. We set an antialias filter to avoid false images. We capture the data, multiply each point by the appropriate windowing coefficient, calculate the FFT and get a magnitude for each frequency bin. We do not care about phase information, just magnitude info. We then plot a graph (I used amplitude in dB vertically since we hear the flutter with our logarithmic ears. FFT data comes out on a linear frequency scale with equal sized bins for each frequency step.

Now we look at our spectrum around 12.7 Hz to find the primary frequency of our flutter idler. Usually we will find a signature at that frequency, and possibly at the second, third& harmonics. If we are foolish and try to resolve low-frequency components on a wideband plot (finding 12.7 Hz on a 5 kHz span), we will be disappointed. In this case we should use a span of about 200 Hz so that we will have bins at 400/1024 or .391 Hz spacing. We would have bins at 11.72, 12.11, 12.5, 12.89, 13.28 Hz, etc. Our data may still fall on the boundary of two bins and be split between them. If we see this happening, we can double the number of samples or reduce our span to narrow the bins.

If you don't have all this fancy hardware and software, you stick with those fabulous ears glued to your head. They are a great spectrum analyzer, but the `Graph' function doesn't work well.

In the MARCAP program I made several playback runs from the same recording, each time re-orienting the rollers to assure random phase of the flutter components. I then accumulated the peak value for each of the FFT frequency bins to get the worst-case spectrum. I figured a producer who was objecting to bad flutter wouldn't be calmed by a tech telling him that, on the average, the flutter was so and so if he was hearing the worst case!

Based upon the way you stated the distance numbers, I think you may not quite have the picture correct, so just a quick review. You state that the distance from the record to play gaps is 2.59 inches if we add the two increments together - the distance `as the tape flies.' At 15 ips, this is one 5.8th (1/5.8) of a second. Anything rotating near 5.8 times per second or multiples thereof will be `synchronized' to the time delay between the two heads in record/play and the flutter of that component will `disappear.' The standard idler is turning about 12.7 rps, which is a ratio of 2.2 with the intergap timing. At 2.0 we would expect full cancellation of any once-around, and at 2.5 we would expect a maximum increase. That roller is near the middle of this range, so we should expect a fairly normal `average' contribution to flutter during record/play. The optional idler has a rate of about 10.2 Hz, giving a ratio of about 1.76 - again about midway, but this time between 1.5 and 2.0. The capstan, pinch roller and reel idler can be checked with the same technique.

The distance from the head gap to the roller isn't of much use by itself for either rotational or scrape analysis. The scrape flutter modes are determined by the spans of tape between rollers. All sliding surfaces are exciters. Worn heads probably change the way the tape is excited, but I don't recall any test that describe how. I suppose we would need the exact position information if we had a complete mathematical model of the tape and machine and we solved the problem analytically as a boundary value problem. It would be fun to modify a deck so that the heads could be slid sideways while measuring flutter. Or maybe a series of tests of various tape types, including 1 and 1.5 mil backings, to compare the effects of roughness and elasticity on scrape flutter.

Please accept these comments in the `mentoring' spirit in which they are offered. I am certainly not criticizing anything you stated. This is as much for my benefit as yours! I have not had any discussions on flutter for a long time, so please bear with me while I enjoy the exchanges.

-----Original Message-----

n: **Steve Fuller [SMTP:sbfuller@compuserve.com]**

:: Monday, August 28, 2000 6:02 AM

dmanquen@email.msn.com

ect: RE: Fritz Chambers

Dale,

> No, that name doesn't ring a bell. Who is/was he? I just
> pulled out the old file to refresh my memory. I dealt with
> Transcom down here in Orange County. By the way, thanks

I knew Fritz when he was with Global Wulfsburg which I believe was a Sundstrand holding at the time and then Sundstrand was bought by Allied Signal which was then bought by Honeywell. As a result of this mess he moved on to in-flight entertainment and is now with Air Show.

> Thanks for taking the time supply us with flutter
> measurements. Remember, though, you made your tests in
> record/play, so you don't know if you missed a critical
> component that has a circumference that is an even multiple
> of the record-to-repro head spacing. I think both the
> scrape flutter idler and the capstan have circumferences
> just over an inch. How does that compare with the
> gap-to-gap distance?

The linear gap to gap distance is about 2.55 inches but this is broken into two pieces by the 3/8 inch diameter scrape idler. The record to idler gap is about 1.14 inch and the idler to play head gap is about 1.45 inch. The scrape idler preceding the record head is 15/32 inch in diameter. The flat surfaces at the gaps are now about 3/32 inch so that will effectively decrease the associated distances to the idlers. I seem to recall an Ampex paper wherein they supported the diameter choices for the flutter idlers and the manual insists that the 3/8 inch idler be between the record and playback heads. The manual also states that the use of both scrape idlers places the scrape frequency at about 10,000 Hz where the use of one places it at about 7,000Hz.

We were talking about actually running flutter tests. We record the pure test tone once, then play it back several times, re-orienting the rollers and capstan each time. This will require disarming the reel servos on some machines.

A note about the test tone source is appropriate. The tone must be very pure in frequency and free of noise. Function generators are generally not good enough to be a flutter tone source. The residual ripple on the power supply rails will cause FM in the Voltage Controlled Oscillator. When in doubt, feed the source directly into the flutter meter and measure the residual noise of the combination.

I chose to use a crystal oscillator as the primary source for two reasons. First, dividing a 4 MHz crystal down to 12.5 kHz averages the FM content of the crystal oscillator to produce a very clean source. Using a filtered digital square wave also gave very good amplitude stability for the AM measurements.

Second, a crystal-accurate tone can be used for speed/drift testing. A digital counter can be used to count the absolute frequency of the tone for speed checks. Flipping the recorded test tape head to tail will give the worst case speed change from head to tail. (For the absolute purist, use the VCO output in the Altair PLL as the counter source. This frequency tracks the tape frequency exactly, but it has none of the level fluctuations due to dropouts that are on the signal directly from the tape.)

Here is something to consider: If you used the flutter meter's stable 12.5 kHz test tone as the high frequency alignment tone on the front of each reel of tape, you would create a Rosetta Stone that would give you amplitude, speed/frequency, flutter, and dropouts. In fact, if you had your flutter meter patched into the console, you could quickly check the tape machine AND THE ROLL OF TAPE in about 20 seconds. More about tape testing after we finish flutter.

I guess I skipped over meter readings for multiple passes. This is where the difference between spectral analysis and a meter really shows up. With the meter we are stuck with a hodge podge of random components that come and go. There is no `good' way to really understand what is going on based upon a single number.

But don't let that reality stop anyone! I believe the flutter standards call for averaging the readings from 5 to 10 passes. Lowell probably remembers how the standards deals with splitting up the record and play to get a single number.

Maybe we should spend a bit of time discussing the `readout' problem. The NAB chose to standardize on a conventional meter with average responding, RMS calibrated specs. The meter damping was in the neighborhood of a VU meter so that fast peaks were sliced off a bit, but the meter wasn't `hyperactive'. The European approach was to follow their own ideas of level indicators and go for a quasi-peak reading meter. Both are reasonable choices based upon the `culture' of each group. Instrumentation flutter meters use statistical nomenclature with `sigma' notation. One sigma, two sigma, three sigma indicate how often a given level will be exceeded. In my opinion, all of these techniques are equally valid, which isn't saying anything good about any of them since a meter just isn't adequate for the job!

But according to my understanding, the FFT isn't perfect either. I believe I mentioned earlier that the FFT can have very high dynamic range if we take enough samples. The random noise of the system averages out over the length of a long time sample. Unfortunately, the same thing happens to the random flutter components! If we have a nice, sustained flutter component due to a rotating element, the signal level is read very accurately. If, on the other hand, a component (let's say an oscillation burst in a servo amp) is present for only part of the sampling interval, the FFT will average that component over the entire sampling interval. So we can be misled by the spectral plots.

I chose to use the peak values. The real test was to see if the resulting plots were `fingerprints' of the machine. When I was measuring the multiple BVH-2000's at Modern Video, I had a chance to check out this concept. First, the peak method with 5 passes gave me plots that were repeatable on a given machine to about +/- 1 dB for each of the major flutter components. And each machine had its own characteristic plot!

During a MARCAP demo I measured an Ampex 1" C format machine at Complete Post that had come from Ampex via their training department or the Olympics or some other route. When I looked at the plot, it didn't seem correct. I wasn't overly familiar with the machine, but the engineer and I poked around to see what was going on. One of the guides seemed a bit strange, so I investigated it further. Turned out that guide was really the flutter idler, equivalent to the reel idler on an audio deck. This idler had a clutch or brake that changed modes for high-speed winding. Well, the mechanism was frozen and the roller had not turned for many moons. The tape wore such a groove that I thought it was a guide! You can imagine what happened when we finally freed things up and that flatted shaft started rotating! Flutter was off the map!

The point here is that we can establish baseline spectra of what to expect, and quickly detect deviations. With just a meter, this problem might have slipped through, even though it was degrading the operation of the machine.

Micha wrote:

In the seventies and eighties I was doing service work on tape recorders (Many more than I knew existed) including cassettes of those days.

One of the confusing issues was that units with the same specs (Frequency response, S/N Wow and Flutter etc) sounded unexpectedly different. This until I used a W/F analyzer, plugged it into a scope and a frequency analyzer and checked the components of the frequency deviation. From there on things became clear:

1)any unit that had any repeat pattern (most of them) had a sound signature of its own.

It did not sound always like the typical W/F but it had a sound. 2)once you found the major frequencies, remedy became much easier, you looked for anything That rotated at that frequency, or its multiple, and bingo!

This is a testament to the value of what we are discussing. Unfortunately, however, Micha was dealing with only part of the available information. I would guess that the flutter meter didn't go beyond 200 Hz. He was able to clearly diagnose the rotational nature of the transport. But he had no quantitative information regarding the scrape flutter or the AM products. These higher-frequency artifacts determine the openness or cleanness - transparency, if you will - of the recorder.

Micha also commented about injecting flutter into digital programs. I had a few good chuckles when people said "Now that I have a digital recorder, I don't need to worry about flutter any more!" When `clock jitter' reared its ugly head, someone finally got around to noting that jitter follows the same Bessel functions that describe scrape flutter! Again we had something that degraded the transparency of the sound, but in a way that was hard to describe. More about this when I (maybe with Lowell's help) rail on the concept of `weighted flutter', one of the big sins of the 20th Century.

I also laughed to think that people thought that their digital transport wouldn't have flutter that could affect the machine's ability to record a very high-density digital signal with the necessary timebase accuracy. Flutter on a digital machine becomes data errors. The source of the problem is harder to detect since the digital process transforms the errors from the obvious pitch changes of an analog machine to more complicated error phenomena.

Chapter 5 Interpreting Flutter Graphs

The website contains 5 flutter spectra obtained on an ATR-102. Four of these are a series of plots that progressively widen the span or frequency range of the graph, starting at 20 Hz, then moving to 50, 500 and finally 5kHz. The fifth graph is a repeat of the 5 kHz sweep, but with the scrape flutter idler stalled. The vertical scales are logarithmic and are calibrated in RMS values.

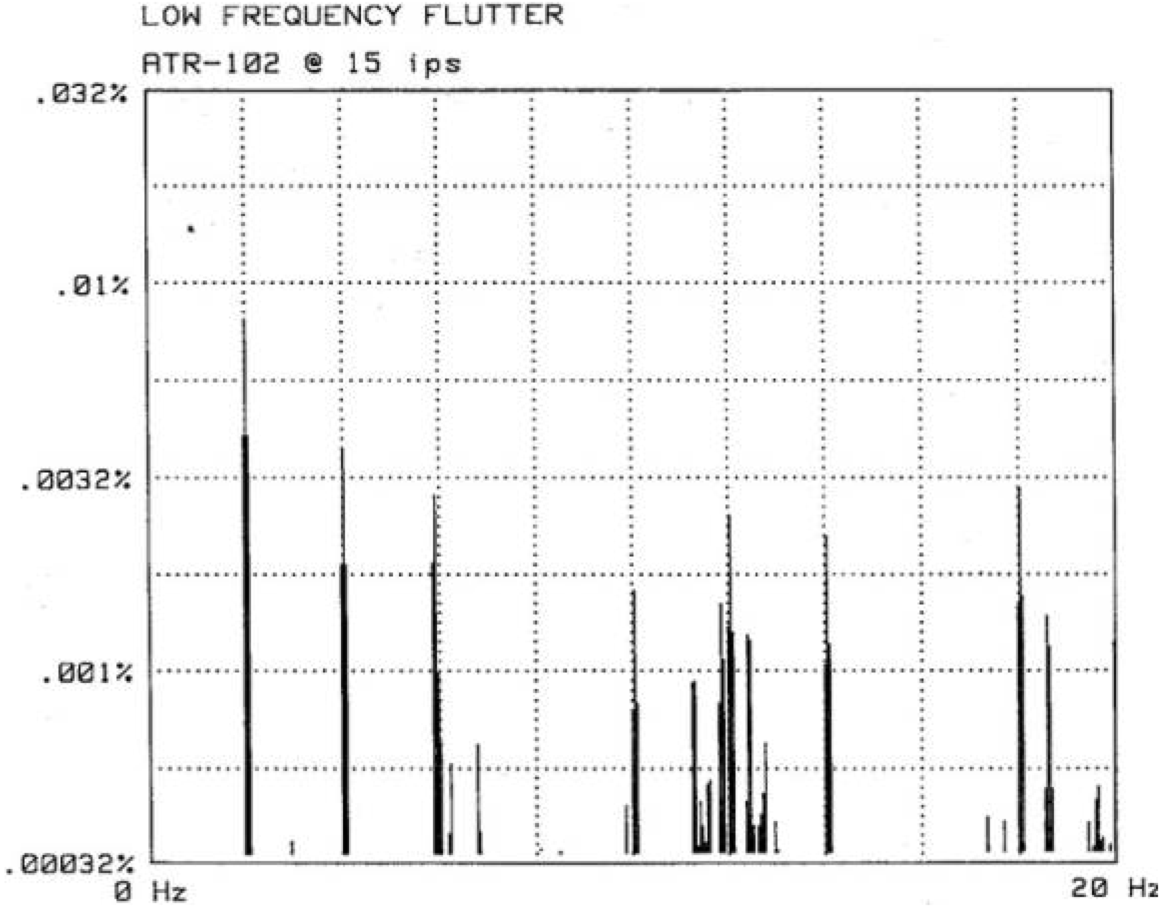

Figure 1 Low frequency ATR flutter

The graph shows the high resolution that can be attained with the equipment and technique that we have been discussing. First we note a spectrum line at 2 Hz, the rotation rate of the capstan (and, I presume, the left-side timer idler. This line is all by itself, 28 dB above the bottom edge of the graph. The noise floor appears to be well below the .00032% level at the bottom of the graph.

We are applying a Hanning window to the data to make it `pseudo repeating'. This window causes some leakage into adjacent bins. For example, notice that the 2 Hz line has small spurs about 6 dB down from the peak. (Each dot on the vertical grid is .5 dB.) Both spurs are exactly symmetric, indicating that our data point is exactly in the middle of the 2 Hz frequency bin. The spurs of the big line near 6 Hz is skewed toward the left side, indicating that the signal frequency fell slightly below the bin frequency represented by the spike.

We also note a second harmonic component at 2 Hz. We call it a second harmonic of the capstan since there are no elements of the size that would give a 2 Hz rotation rate.

We might also think that the line near 6 Hz is the third harmonic, but it isn't. The line is slightly below 6 Hz. The source of this line is an unknown something rotating at this rate. The 6 Hz line that would be the capstan's third harmonic appears to be at least 10 dB below this mystery line.

(We are at a disadvantage without the program running on a computer. The built-in cursor would allow us to read out the exact frequency and amplitude of each point on the graph and zoom in for a more detailed examination.)

There is a hole at 8 and 16 Hz, but we have lines at 10, 12, 14 and 18 Hz. We could examine this data and compare it with anything and everything for which we can measure a diameter. For example, the scrape flutter idler will show up at ? Hz.

We also have ball bearings that have internal ball rates. In some cases we can find data that will give the ball diameter and the inner and outer race diameters so that we can calculate the ball rate. We have ball bearings on the capstan assembly, the tension sensing rollers, the timer idler and the spooling motors.

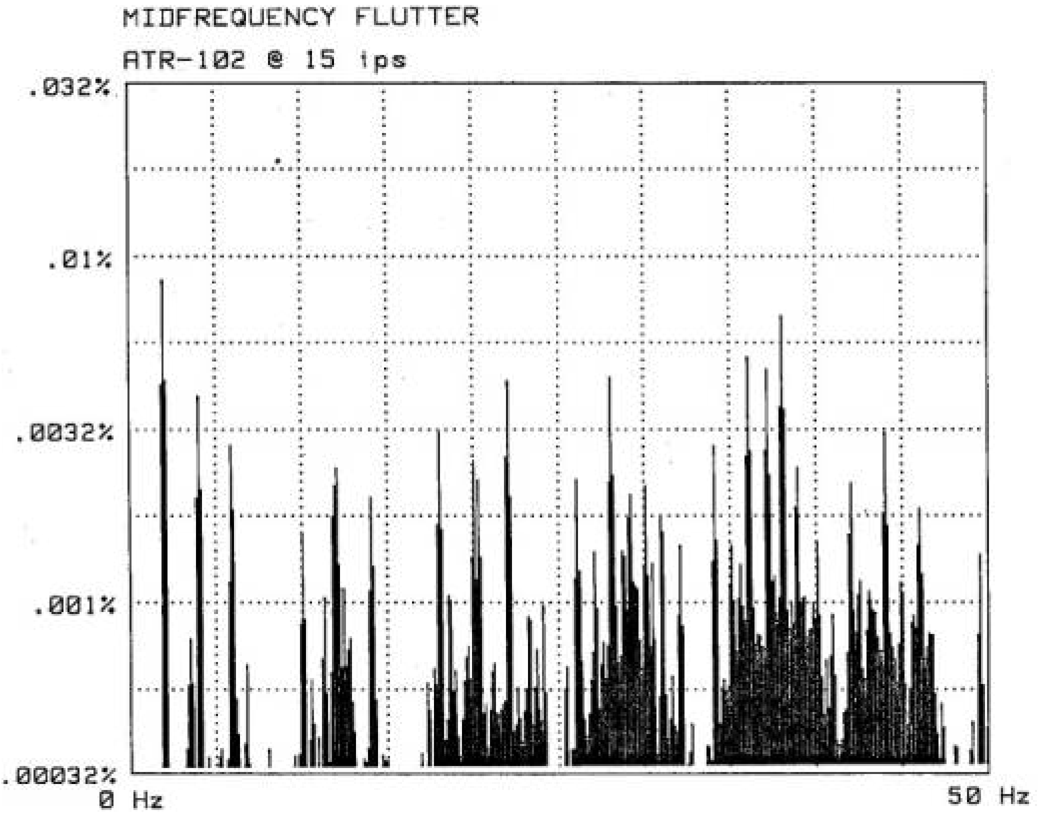

Figure 2 Low Midfrequency Flutter out to 50 Hz

This graph is from a different set of data. For example, note the line 3.5 Hz on the new graph that was absent on Figure 1. But also compare the amplitudes of the first three major lines. They are virtually identical, indicating good repeatability for non-random components.

On this graph we begin to detect `clumps' of data that are about 8 Hz wide. I don't know the source of these clumps, but this structure continues out to beyond 500 Hz!

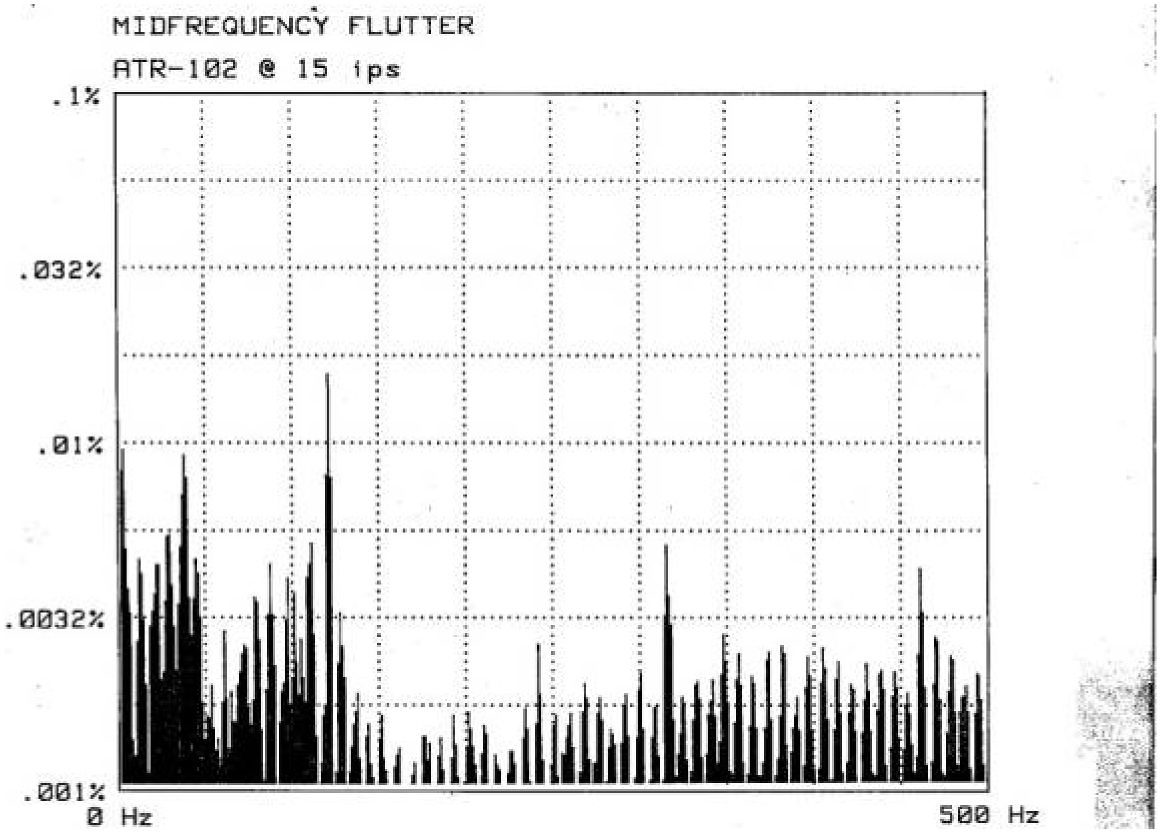

Figure 3 High Midfrequency Flutter out to 500 Hz

This view shows us that many of the mechanical rates are dying off by the time we reach 50 Hz. We then go through a not-too-active zone until we reach 120 Hz. Here we encounter the largest single line in any of the normal spectra. Surprisingly, we don't see anything at 240, 360, or 480 Hz! The 120 Hz component might be due to the 120 Hz ripple on the capacitors that drive the various servos, or they might be due to a simple ground loop that permits the hefty charging pulses for the big capacitors to couple into the ground bus of the servo speed-control amps. My guess is a ground loop - a very small ground loop that would probably go unnoticed without our high-resolution analysis.

We also see lines at 315 Hz and 461 Hz - sources unknown.

By the way, I pulled out a printed circuit motor that I believe is similar to the ATR motor to see what the torque cogging rate should be. The armature on my motor has 141 commutator conductors, which would give us 282 torque pulses per second at 15 ips. There is nothing unusual in the spectrum at this frequency. The capstan servo is a current-mode arrangement that does a great job of smoothing out any pulsing.

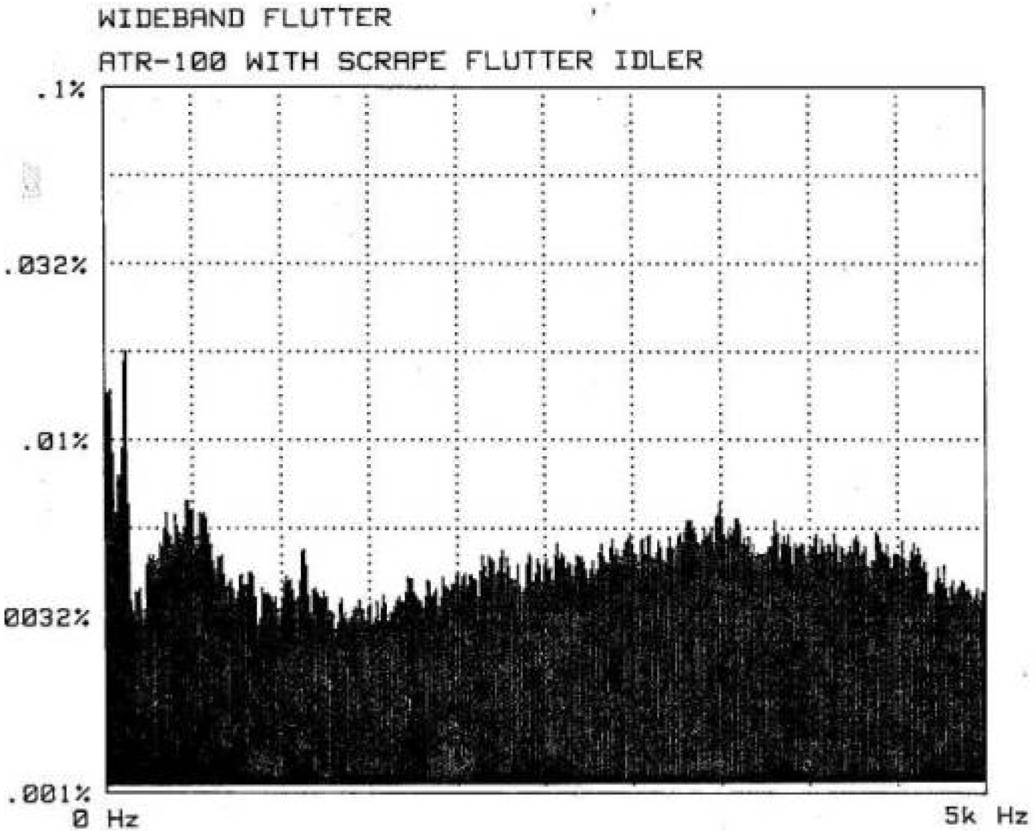

Figure 4 Wideband Flutter out to 5 kHz

The wideband view leaves behind all of the mechanical rattles and clanks. We enter the realm of vibrating tape sliding over stationary things. So what is stationary on an ATR. The heads, of course, but also the ceramic guide points at the edges and center of the head assembly. We also have ceramic flanges on the tension dancer arms. Our graph has a low point at about 1.5 kHz. Beyond that point the flutter gradually rises about 5 dB to a maximum at about 3.5 kHz. We see some rolloff beyond this point, but the flutter meter is also running out of bandwidth. We would need a frequency response sweep of our instrumentation to get more detailed here. I would rate the ATR's scrape flutter performance as good, but not great.

I compared this plot to an MCI JH-114. In general, the MCI plot is about 5 dB higher at most frequencies. The scrape flutter curve also peaks around 3.5 kHz, but falls off a bit faster. It is down 9 dB by the edge of the graph compared to the ATR's 4 dB.

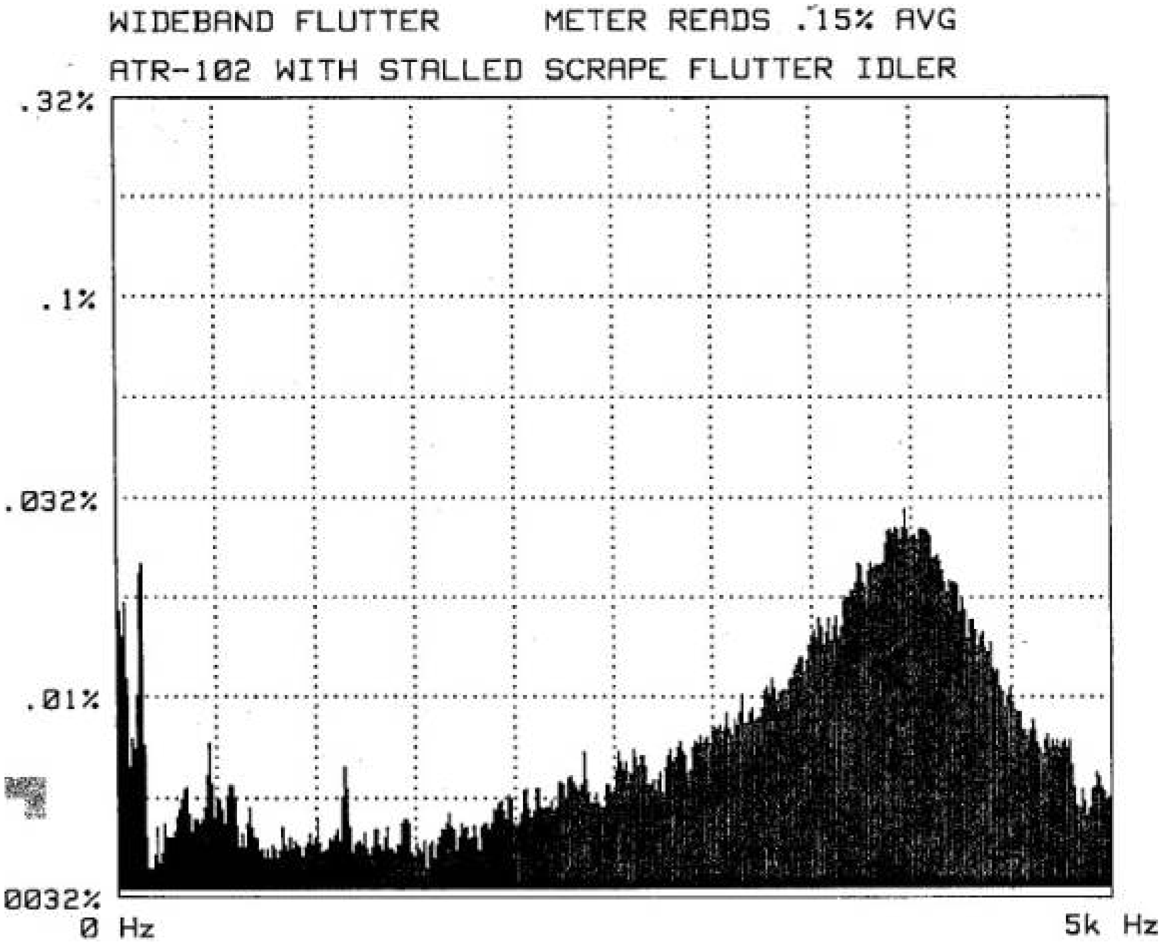

Figure 5 Wideband Flutter with the scrape idler stalled

In this case we reach in with a pencil and stall the scrape flutter idler. The scrape components shoot through the roof, rising at least 13 dB to a value higher than any other line on the spectrum. (Note that the vertical axis of the graph was shifted by 10 dB.) This curve would be very similar to any pre-440 Ampex machine without a scrape flutter idler. The total energy in the scrape flutter - the amplitude times the frequency - is huge, totally dwarfing the energy on the left half of the plot.

But did I hear someone ask `Is it fair to stall the scrape flutter idler? Doesn't that cause extra friction?" Yes, it does. But the curves for a stalled scrape idler, a missing scrape idler and a scrape idler that is slipping are all very similar. The amplitude may change a bit, but the peaked shape is still there.

The narrow-band analysis that we have just finished really doesn't give us a clue regarding what the meter would read. We could mathematically integrate all of the components to get a value, but I just used the meter on the Altair box.

Is this much care and detail necessary? Take, for example, the 2 Hz capstan signature at .008%. Is this too small to care about? Not if you believe in preventive maintenance. This is a reference against which you can periodically test the machine. You won't be able to make absolute judgements right away, but after you accumulate a stack of graphs an inch and a half thick like I have, you begin to see trends. You don't expect a hysteresis synchronous capstan motor on an AG-350 to match the ATR (unless you are Kurt). But you do know the characteristic signature of that motor with all its harmonics.

Tuning the capacitor on a sync motor is a very good use for this technique. The value of the capacitor and sometimes the addition of some series resistance can make a significant difference in the flutter generated by a sync motor. Tuning for the proper value is easy when you can watch the specific spectrum lines generated by the motor.

Chapter 6 Amplitude Modulation

Amplitude modulation comes from a number of sources, and, like flutter, identifying the probable source is the key to diagnosis/evaluation.

When the tape is manufactured, the first step is to mix all of the ingredients is a milling process. The goals of the process are to separate and coat each magnetic particle with a coating of binder glue. Too much milling will fracture the particles, producing low-coercivity particles that raise the printthrough level of the tape. Not enough milling leaves clumps of particles together that have coercivities much higher than the desired mean coercivity. If you listen to the demodulated AM output, you will hear popcorn noise due to the large coercivity particles and 'white' noise related to the uniformity of dispersion of the single particles.

In addition, coating abnormalities and surface damage will show up as dropouts.

And any distortion of the tape backing will produce variations in the contact with the record and play heads.

On the other side of the coin we have the contributions due to the transport. If the tension is inadequate, the tape will not be held snugly against the heads. The amount of 'snugness' required varies with the transport design and the amount of wear on the face of the head. Heads that are worn flat have a lower contact pressure at the gap.

Another factor is any grooving of the heads due to wear. This can cause liftoff of the tape due to pinching in the groove. Most professional heads are undercut or grooved where the top and bottom edges of the tape run to eliminate the formation of a groove.

Sometimes the AM analyzer can be used to diagnose problems that repeat cyclically. For example, knowing that tape-slitter wheels are about 4" in diameter, calendering rolls are 10-12" in diameter, and jumbo rolls are perhaps 18-24" in diameter can help pinpoint the cause of erratic output due to physical distortion of the base film. Just measure the period of the disturbance and convert that into tape distance and then equivalent diameter.

There are also problems due to the heads and electronics. The bias frequency and waveform and the record head gap length (width to us old farts) determine how the recording flux is distributed within the oxide layer.

The contour of the face of the head can also be a factor. At 3M we had a problem on the early machines with `rocks' in bass guitar notes. When we compared our new machine with our old Ampex 300 `reference machine', the Ampex was much better. Jack Mulling noticed that the heads on the 300 were quite a bit flatter in contour, so we tried flattening the face of the 3M heads. It did improve the noise.

We also had a problem with an external magnetic field causing noise problems at the erase and record heads. The magnetic field of the left-hand pinch roller solenoid was being `conducted' along the steel linkage to a point quite close to the erase and record heads. The resulting `DC' flux at the heads was upsetting the erasing and biasing fields generated at the gaps of the heads, leading to noisy recordings. We replaced the steel linkage with non-magnetic stainless steel to eliminate the problem.

The bottom line is that AM is affected by the type of tape, the 'goodness' of the specific lot, the design of the recorder, and the 'wellness' of the recorder. Thus a roll of tape will have different AM readings on different machines. If the machines are both good designs that are well maintained, the readings will be close.

I have one case history of a major studio that used my meter to test incoming lots of tape for AM. After about a year of baseline information, they got a batch that read 2 to 3 times higher than the baseline. Same transport, same adjustment - but new tape vs. a roll of their reference tape. They rejected the lot before ever using any of it in the studio. About 9 months later they had a similar experience. When they checked the lot numbers, guess what? Yep, it was the same bad batch shipped to them a second time!

We are not talking 10% difference here. This was day and night difference. (Ampex Tape and 3M Tape both had my Altair boxes.)

Another case history occurred at Martinsound Studios. At the beginning of a session, the first playbacks on the MCI 2" recorder revealed a burbly popping noise. The first stab was to replace a microphone and its power supply. Then the alignment of the machine was re-checked. Finally, in desperation, someone suggested changing the tape. The problem immediately disappeared. Over a half-hour of session time was lost.

Below is the AM spectrum of a normal roll of tape on an MCI machine compared with the spectrum of the bad roll of tape on a similar machine. Note that the low frequency AM content is much higher, again by a factor of about 3.

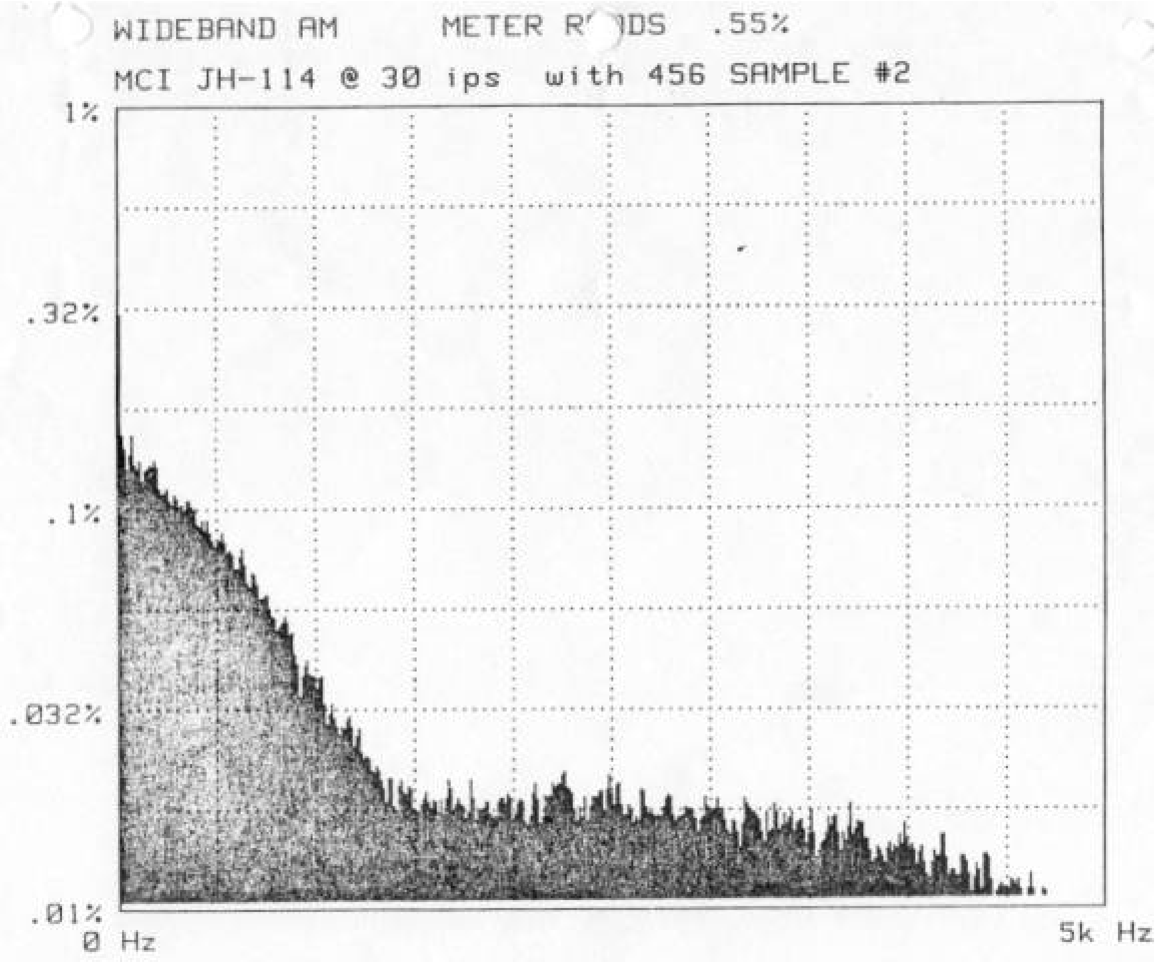

Figure 6. AM spectrum for normal roll of tape

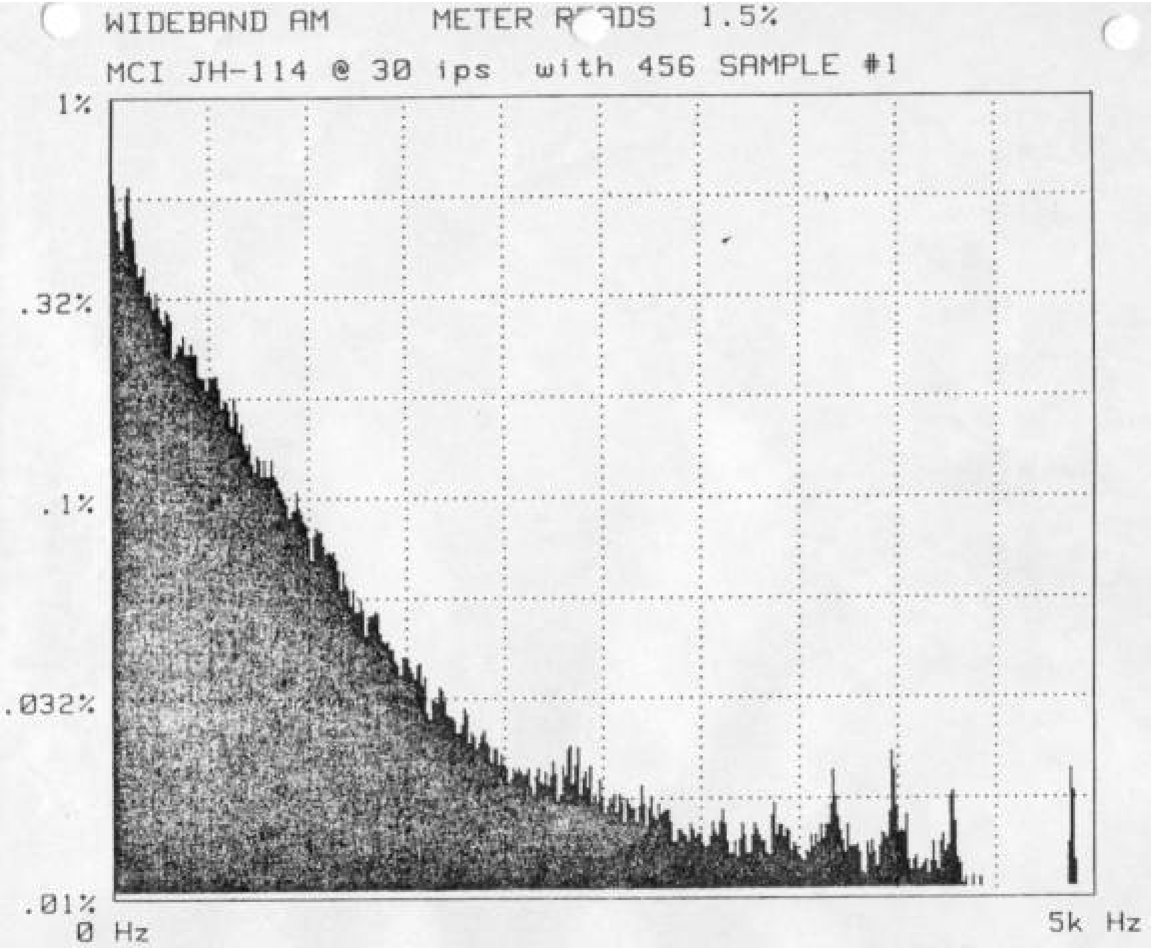

Figure 7. AM spectrum for bad roll of tape

An interesting demonstration is to add the demodulated AM from the Altair box to the 12.5 kHz test tone with dropout/AM using the 'Add' function on a 'scope. If the gains are set properly, one edge of the signal from the tape will be a flat line. Unfortunately, the other edge is twice as bad since we are not un-AMing, just adding the amplitude of the AM envelope! But this demo did convince me that I was really reading the variations that create the ragged edge.

* About 6 years ago I had an opportunity to teach the designers at Alesis about current-mode repro circuits. We set up the standard sync/audio head on an ADAT - a standard VHS-format head with a sync track at the bottom and two audio tracks at the top - to play back cueing or longitudinal timecode. I shoved the head signal into the summing junction of an op amp as Jay discussed. This simple circuit would play back timecode at 40 times normal speed with a good waveform, just as Stan noted. Shunting out the coil and lead wire capacitance makes a ***big*** difference, and the automatic eddy current compensation does the rest.

The integrator effect (6 dB/octave rolloff) continues throughout the usable passband, thereby treating all harmonics of the timecode similarly to avoid waveform distortion at high speeds.

But note that reading timecode at high speed is not too difficult since the wavelength of the recorded timecode doesn't change as you speed up the machine. Only the temporal frequency changes. Playing back a 100 kHz timecode signal at 600 ips isn't like trying to record and play back 100 kHz at 15 ips.

I don't know if this circuit was used in the cueing feature of the new version of the ADAT with the Panasonic transport. There was a lapse of a few years between our early development work with the Panasonic and the introduction of the final product. We had other interesting moment, like when Keith Barr, the chief honcho, declared that we should forget balanced outputs at +4 dBu because the newer 5 volt analog IC's were the way to go. I believe the machine still has +4 outputs.

By the way, I also used that same sync/record head to test flutter on the transport. I still have a little adapter box that switches the head between `Record' (connected to the flutter meter's oscillator through a booster amp) and `Play' (connected to the flutter meter's measurement input through a preamp.) I couldn't convince the Alesis production people that looking at the raw tape motion could help them diagnose transport problems. They relied solely upon the RF waveform and the recovery of the digital information to test for 'goodness'.

Back to top

PORTAL >> AUDIO - PENNY & GILES FADERS - TRAINS - FORUMS - SWAP STOP << PORTAL