

Understanding Headphone Power Requirements

- **Headphone Sensitivity**
- **Listing of Headphones**
- **HC 6 Power vs. Loudness**
- **HC 4 Power vs. Loudness**
- **Table of Common Sound Pressure Levels**

**Dennis Bohn
Rane Corporation**

**RaneNote 100
© 1983 Rane Corporation**

INTRODUCTION

Much confusion abounds regarding headphone power requirements. This Rane Note is intended to disperse some of the mist surrounding headphone specifications and hopefully give you a clearer understanding of how much power is really needed for your application.

HEADPHONE SENSITIVITY

Headphone manufacturers specify a “sensitivity” rating for their products that is very similar to loudspeaker sensitivity ratings. For loudspeakers, the standard is to apply 1 watt and then measure the sound pressure level (SPL) at a distance of 1 meter. For headphones, the standard is to apply 1 milliwatt ($1 \text{ mW} = 1/1000$ of a watt) and then measure the sound pressure level at the earpiece (using a dummy head with built-in microphones). Sensitivity is then stated as the number of dB of actual sound level (SPL) produced by the headphones with 1 mW of input; headphone specifications commonly refer to this by the misleading term “dB/mW.” What they really mean is dB SPL for 1 mW input.

Think about these sensitivity definitions a moment: headphone sensitivity is rated using $1/1000$ of a watt; loudspeaker sensitivity is rated using 1 watt. So a quick rule-of-thumb is that you are going to need about $1/1000$ as much power to drive your headphones as to drive your loudspeakers since both of their sensitivity ratings are similar (around 90-110 dB SPL). For example, if your hi-fi amp is rated at 65 watts, then you would need only 65 mW to drive comparable headphones. (Actually you need less than 65 mW since most people don’t listen to their loudspeakers at 1 meter.) And this is exactly what you find in hi-fi receivers—their headphone jacks typically provide only 10-20 mW of output power.

Take another moment and think about all those portable tape players. They sound great, and loud. Why, you can even hear them ten feet away as the teenage skateboarder that ran over your foot escapes.

Power output? About 12 mW.

THE LIST

As an aid in finding out how much power is available from either the HC 6 or HC 4 Headphone Consoles, we have compiled a listing of popular headphones. Included is a column giving the maximum SPL obtainable using the HC 6 or HC 4 and any particular headphone—ultimately, it all gets down to actual SPL. The power rating really doesn’t matter at all—either it’s loud enough or it isn’t (of course it has to be clean power, not clipped and distorted). The SPL numbers shown are for maximum *continuous* SPL; for momentary peak SPL add 3 dB.

Note that the maximum achievable SPL varies widely for different models and manufacturers, ranging from a low of 107 dB to a harmful 146 dB! The table also shows there is very little relationship between headphone impedance and sensitivity, and that power output *alone* means nothing, since in one case 80 mW produces a maximum SPL of 107 dB, yet in another case the same 80 mW yields an SPL of 124 dB!

Sensitivity (dB) is the measured sound pressure level with 1 mW of power. The **Max Power (mW)** columns are typical continuous average (rms) power, 20 Hz-20 kHz, with THD less than 0.1%.

If headphones are not yet owned, or replacements are desired, use this listing as a guide for selecting headphones with sufficient sensitivity for the maximum desired SPL.

Note: headphones with an impedance of less than 32 ohms are not recommended for use with the HC 6 or the HC 4.

Manufacturer	Model	Impedance (ohms)	Sensitivity (dB)	HC 4 Max Power (mW)	HC 4 Max SPL (dB)	HC 6 Max Power (mW)	HC 6 Max SPL (dB)
AKG	K141M	600	98	89	117	80	117
	K240M, K240DF	600	88	89	107	80	107
	K270S	75	92	239	115	380	118
	K301	100	94	225	118	285	119
	K401, K501	120	94	220	117	290	119
Audio-Technica	ATH-COM1, ATH-COM2, ATH-908	40	90	220	113	440	116
	ATH-910	40	92	220	115	440	118
	ATH-P5	40	100	220	123	440	126
	ATH-M40	60	100	238	123	400	126
	ATH-D40	66	102	235	126	295	127
Beyerdynamic	ATH-M2X, ATH-M3X	45	100	230	123	435	126
	DT150	250	97	160	119	175	119
	DT211, DT311	40	98	220	121	440	124
	DT250	80	98	240	121	360	123
	DT411	250	102	160	124	175	124
Fostex	DT 531	250	95	160	116	175	116
	DT431, DT331	40	86	220	109	440	112
	DT770PRO, DT990PRO	600	96	89	115	80	115
	DT801, DT811, DT511	250	94	160	116	175	116
	DT901, DT911	250	98	160	120	175	120
Grado	T-5	44	96	225	119	435	122
	T-7	70	98	240	121	385	124
	T-20	50	96	233	120	425	122
	T-40	50	98	233	122	425	124
	SR 325	40	96	220	119	440	122
Hosa	HDS-701	40	91	220	114	440	117
Koss	A/250, A/200, A/130, TD/80	60	98	238	123	320	125
	R/200	60	84	238	108	400	110
	R/100, R/45	60	85	238	109	400	111
	R/90, HD/2, SB/15	60	100	238	123	400	126
	R/80, R/35S, R/20, Porta Pros	60	101	238	124	400	127
MB Quart	R/70B, R/55B, SB/50, SB/35	60	101	238	124	400	127
	R/40	60	90	238	114	400	116
	R/30S	60	106	238	130	400	132
	R/10	60	103	238	127	400	129
	TD/75	60	95	238	119	400	121
Sennheiser	TD/65	90	101	235	124	340	126
	TD/61	38	93	212	116	440	119
	QP 805	300	98	145	120	80	117
	HD 400, 433, 435, 470	32	94	200	117	450	121
	HD25	70	120	240	144	380	146
Sony	HD445	52	97	235	121	390	123
	HD25SP	85	100	235	123	350	125
	HD265, 525, 535, 545, 565	150	94	207	117	190	117
	HD455, 475	60	94	238	118	400	120
	HD465	100	94	225	118	285	119
Stanton	HD 570	120	95	220	110	290	120
	HD580, 600	300	97	145	118	80	116
	MDR-V100MK2	32	98	200	121	450	125
	MDR-85	40	102	220	125	440	128
	MDR-V600, MDR-D77	45	106	230	129	435	132
Telex	MDR-CD10	32	96	200	119	450	123
	MDR-CD550, CD750	45	100	230	123	435	126
	MDR-CD6	45	110	230	133	435	136
	MDR-CD850, CD950	32	102	200	125	450	129
	MDR-CD1000, CD3000	32	104	200	127	450	131
Yamaha	MDR-D33, MDR-D55, MDR-7504	45	104	230	127	435	130
	MDR-7506	63	106	240	129	400	132
	MDR-7502	45	102	230	125	435	128
	ST PRO, DJ PRO 1000	32	100	200	123	450	127
	PH-6	600	105	89	124	80	124
Yamaha	RH5MA	32	98	200	121	450	125
	RH1	32	90	200	113	450	116
	RH2	32	95	200	118	450	122
	RH3	60	95	238	119	400	121
	RH10M	40	102	220	125	440	128
	RH40M	32	103	200	126	450	130

Sound Pressure Level Equivalents

SPL-dB	Common Example	SPL-dB	Common Example
140	Irreparable damage	60	Normal conversation
130	Jet aircraft taking off	50	Elevator music
120	Threshold of pain / Thunder	40	Normal home background (kids asleep)
110	Threshold of discomfort	30	Studio background
100	Dirt bike / Riveter	20	Rustling of leaves / Quiet whisper
90	Start of unsafe levels	10	Butterfly swoop
80	Average factory	0	Threshold of hearing
70	Kids at play		

Permissible Noise Exposures

Extracted from the U.S. Department of Labor Noise Regulations

Duration Per Day, Hour	Sound Level (dB), A-Weighting
8	90
6	92
4	95
3	97
2	100
1½	102
1	105
½	110
¼ or less	115

NOTE 102

Professional Audio Products



ANALOG I/O STANDARDS

- 3-PIN CONNECTOR WIRING
- BALANCED INPUTS & OUTPUTS
- FLOATING UNBALANCED OUTPUTS
- RFI & EMI PROTECTION
- GROUNDING
- PASSIVE BYPASS

INTRODUCTION

Quality engineering necessitates standards. Recognizing this, Rane has adopted an input/output (I/O) convention for all of its products which parallels international standards where applicable and accepted practices elsewhere. This Note describes and discusses each of these so the user of our equipment will never be in the dark regarding polarities, grounding or connector wiring. In general, we are a left-to-right, front-to-back, top-to-bottom society—which has absolutely nothing to do with this Note. It's just interesting.

Dennis Bohn
V.P. Research & Development

Although written in 1982, Rane customers continue to request this Note. They refer to it often for its basic information and background on connectors, circuit and wiring conventions. Indeed, the derivation given for the differential-mode and common-mode input impedance of the everyday difference amplifier continues to be referenced by other authors. While some of it may seem dated (and *is*) we continue to keep it among our current Rane Notes for its historical perspective and general usefulness. Please read it with its origin date in mind. Seeded throughout the text are a few parenthetical *Update Notes*. These aid in making certain important concepts current. For exact detailed specifics on any current Rane product, contact our Technical Service Department.

CONNECTOR CONVENTION

The single biggest source of wiring confusion in the pro audio industry is 3-pin connectors. Even their name is a source of confusion. Most commonly called “XLR” jacks, they are also known as “Cannon plugs”, “3-pin connectors”, “XLR-type” jacks, and “circular connectors”. The name confusion makes sense since this type of jack has never been given a generic title. Other audio connectors are commonly referred to as “phono” (originally called RCA jacks), or “phone” (¼" jacks, after headphones and telephones—their original usage), or “DIN” plugs (Deutsch Industry Norm) a German standards organization and European standard.

Nothing comparable has caught on for 3-pin connectors, so they have become known by the original manufacturer’s model number. The original manufacturer was ITT-Cannon and the model number was “XLR”. Today, “XLR” is a registered trademark of ITT-Cannon and cannot legally be used to describe any other manufacturer’s version of this connector. And since very few use ITT-Cannon XLR connectors, violation of their trademarked name is commonplace. Until something better comes along, Rane will refer to this type of jack as a “3-pin connector”.

3-PIN CONNECTORS

Much to many people’s surprise there is a standard for wiring 3-pin connectors—not two, as commonly believed. The conflicting so-called “European” and “American” standards are myths. There is one standard. It is *IEC 268*¹ issued in 1975 and signed by 17 countries, including the UK and USA. In addition, the same convention was adopted separately in the UK as *BS 5428*². And finally, has been made an American National Standard. ANSI PH7.102-1983. There *should* be no conflict in connectors. There obviously is, and that is a sorry reflection of how badly informed are many manufacturers of pro audio equipment. (*Update Note: In 1992, the Audio Engineering Society finally issued it’s Standard AES14.1992 agreeing with the above standards.*)

Rane wires its 3 and 5-pin connectors per IEC 268 as follows:

3-Pin Connectors

- Pin 1** Ground (Shield, screen, etc.)
- Pin 2** Positive (Signal, Hot, etc.)
- Pin 3** Negative (Return, Common, etc.)

5-Pin Connectors

- Pin 1** Ground
- Pin 2** Left Positive
- Pin 3** Left Negative
- Pin 4** Right Positive
- Pin 5** Right Negative

¼" PHONE JACKS

On Rane equipment outfitted with input/output phone jacks, the wiring convention for standard tip-ring-sleeve (TRS) connectors is:

TRS ¼" Jacks

- Tip = Positive**
- Ring = Negative**
- Sleeve = Ground**

TS ¼" Jacks

- Tip = Positive**
- Sleeve = Ground**

For ¼" TRS output jacks designed exclusively for headphone use, the standard is:

Headphone Jacks

- Tip = Left Positive**
- Ring = Right Positive**
- Sleeve = Common Ground**

TRANSFORMERLESS BALANCED INPUTS


Rane incorporates automatic balanced-unbalanced inputs on all of its pro audio products. Balanced inputs are a thing of beauty. They take two signal lines and amplify only the difference between them while rejecting anything that is common. The common stuff is all the garbage induced in the lines as they dutifully connect two pieces of equipment. Hum, noise, AM radio stations, police, fire and ambulance radio signals, radar, fluorescent lights—even life itself—all competing to violate the pure signal as it tries heroically to traverse from one piece of gear to another.

Many solutions exist to accommodate balanced inputs. Transformers are the most common. And the most expensive. And the most troublesome.

So much for transformers.

Rane uses the best active balanced input design that fits the application and budget for each particular unit. These range from the true instrumentation amplifier designs used in Flex (and other) units (see the *Flex Users Guide* for details), to the very popular and successful difference amplifier found throughout the pro-audio industry.

A difference amplifier, like its name implies, is one that amplifies only the *difference* between two input signals. For instance, two inputs of +1 volt and -1 volt respectively result in an output of +2 volts for unity gain designs, i.e., a *difference* of +2 volts. Or, two signals equal in value, say, +1 volt each, result in 0 volts output (their difference). Two equal signals are called *common*, and the amplifier is said to be operating *common mode* (as opposed to *differential mode*). For real-world audio input signals the amplifier operates in both modes simultaneously—it amplifies anything that is different (signal) and rejects anything that is common (garbage).



**Figure 1. TRANSFORMERLESS
BALANCED INPUT STAGE**

With a single op-amp difference amplifier, you can have *either* equal common mode input impedance or equal differential mode input impedance, but not both (see sidebar box on back page for details). Of the two the most important is common mode impedance; it must be equal for the minimization of hum and RFI pick up.

With equal common mode impedance, you get unequal differential impedance. For some designs (where all resistors are equal), the mismatch can be as much as three to one, with the positive input impedance being three times the negative input impedance. This sounds like a big problem, when in fact, it is not. All it means is that the load impedance seen by the positive and negative line drivers are not equal, so the negative driver has to deliver more current. But, as long as the impedance are within reason (at least 5k ohms or greater), you really do not have a problem.

AUTOMATIC BALANCED-UNBALANCED OPERATION

Figure 1 shows a balanced stage with a TRS 1/4" phone jack input. There is a bit of cleverness going on here that comes for free. Using a stereo phone plug into the stereo jack gives you the balanced line input. Not so obvious is that by using a mono phone plug the circuit automatically switches to unbalanced single line operation. Nice, huh? What happens is that the ring of the jack gets shorted to ground by the mono plug and turns the difference amp into a non-inverting stage with a gain of two (assuming all resistors equal), but there is a 6dB pad (a loss of 1/2) hung on its positive input, so the net result is a unity gain unbalanced input. The stage becomes universal now, being either balanced or unbalanced, depending on whether the input plug is stereo or mono.

TRANSFORMERLESS BALANCED OUTPUTS

Balanced outputs have evolved in similar fashion to balanced inputs, starting out with everything using transformers and gradually moving toward transformerless solutions. The most commonly seen solid-state circuit involves two op amps as shown in Figure 2 (or some variation of this configuration). There is one problem and one myth that must be overcome

before this circuit can be used. The problem is one op-amp gets shorted to ground if an unbalanced (mono) plug is used for interconnect. If allowed to happen this can cause possible damage to the op amp, along with possible distortion interaction with the other op-amp if they are part of the same IC package. Rane solves both of these problems very simply by adding resistors in series with each output (isolation of reactive output loads require these anyway), and by ensuring that separate op-amp packages are used. These things, along with using IC's capable of driving lines, guarantee that shorting either of the outputs does not damage the op-amp or induce any distortion products into the other side.

The myth involves a mistaken belief that the cross-coupled output stage developed by Thomas Hay at MCI (AES preprint no. 1723) retains the 6dB headroom advantage of balanced lines when used in an unbalanced fashion. Hay's circuit does two nice things when either side is shorted: [1] The shorted op amp's output is forced to zero volts. It does not drive the output resistor at all, so no stress is put on the output stage. [2] It doubles the *gain* of the unshorted op amp. This way, the output still delivers the same level for the same input either balanced or unbalanced. For example, a 1 volt input produces either ± 1 volt output balanced, or 2 volts unbalanced. On the surface, this suggests the 6dB headroom improvement offered by all balanced output stages is preserved when used unbalanced. *This is not true at all.* Since the *gain* of the output stage is doubled then it will run out of headroom 6dB earlier than when operated balanced. The op amp can only swing so many volts. When only one op amp is driving the line there will be a 6dB loss of headroom—period. It does not matter how cleverly the op amp is configured.

Since a discretely implemented cross-coupled output design requires many carefully trimmed precise parts, without producing any real advantages. Rane does not use it. Where appropriate. Rane does use a special laser-trimmed integrated circuit developed for this purpose. This design exhibits the same 6dB loss of headroom when used unbalanced, but otherwise is an excellent balanced line driver in a small package. (See note 124 for additional details.)




Figure 2. BALANCED OUTPUT STAGE

FLOATING UNBALANCED 1/4" OUTPUTS

On many Rane products fitted with 1/4" outputs, a floating unbalanced drive circuit is used. What is required is an output scheme compatible with balanced or unbalanced inputs that allows signal ground and shield (chassis) grounds to be isolated.




Figure 3. FLOATING UNBALANCED 1/4" TRS OUTPUT

The best solution, as is so often true, is the simplest. It is no more complex than Figure 3. By taking a regular single-ended (unbalanced) line driver and floating its ground you create a compatible system for driving differential (balanced) inputs that is trouble free. Figure 4 shows the interface between the two systems. Note that while the chassis of each unit may be at the same potential, their signal grounds are allowed to be at different potentials. This is very important in keeping hum common mode, and not differential. Any difference of potential existing between the two grounds is seen as a common mode signal and is rejected. This problem is particularly acute where units of various manufacturers

are bolted together in the same rack. All of Rane's products may be intermixed with other brands with minimal hum problems.

Like the balanced input stage of Figure 1, the floating output stage of Figure 3 automatically switches to a conventional unbalanced output stage if a mono phone jack is used for interconnection. The ring of the jack gets shorted to the sleeve, thus grounding the sleeve as you would in a normal unbalanced system.

CHASSIS GROUND

Signal ground is tied to chassis ground by a network consisting of a 10 ohm resistor and a .01 microfarad capacitor in parallel. (Update Note: the resistor has been replaced by a Ground Lift switch on most products.) The small resistor allows the signal ground and the chassis ground to be at different potentials for all audio frequencies; while the capacitor shorts the two grounds together for high radio frequencies. The actual chassis is carefully tied to the chassis ground path at only one point. This is critical if proper chassis shielding is to be realized.

SIGNAL POLARITY THROUGH RANE PRODUCTS

While not wishing to engage in the controversy that rages on in the audio world regarding whether or not signal inversion is audible, it is Rane's policy to not invert signals through its products. Our position is very simple—overall signal inversion is not necessary, so why do it?




Figure 4. INTERCONNECTION




Figure 5. PASSIVE BYPASS WIRING

PASSIVE BYPASS

It is Rane's philosophy that Bypass switches should not require power to operate. For this reason all bypass circuitry is totally passive as shown in Figure 5. *(Update Note: Some digitally controlled products violate this rule, while others use relays to accomplish the same function. Technically not passive, but they work in a fail-safe mode so that bypassing is automatic with power failure.)* With the Bypass switch engaged, all signal connections are rerouted directly to the output (including ground) thus turning the unit into a rather expensive 3-wire patch cord. (Note that the floating output ground reverts to a straight through balanced ground connection in bypass condition.) If at any time, power is lost to a Rane product, simply pressing the Bypass switch completely restores the signal path. Some Rane products use relays to automatically restore signal path upon power loss. Check data sheets.

RFI & EMI PROTECTION

Radio frequency interference (RFI) is one of the most annoying problems that plagues any piece of equipment featuring high gain (e.g. mic preamp stages in mixers). It normally enters as a common mode signal induced into the mic input cables, although it also gets in via the line cord, or in extremely bad areas directly onto the output lines.

The use of balanced inputs and outputs on Rane products greatly reduces their susceptibility to RFI pick up. In addition, RFI filter traps are featured on all high gain inputs. The single point chassis ground concept discussed earlier further helps in reducing RFI problems. Praying doesn't hurt either.

Electro-magnetic interference (EMI) can also be very troublesome if proper steps are not taken during product design. EMI is primarily due to one piece of equipment inducing 50Hz or 60Hz hum into another piece of equipment. For example, power amplifiers with their large transformers and associated magnetic fields may easily induce hum into low level signal processors having high gain. Rule: Thou shalt not mount thy amplifier upon thine mixer.

Proper shielding is essential in reducing any piece of equipment's susceptibility to EMI. All Rane products utilize steel chassis and front panels to help shield against stray field pick up. Balanced inputs further help in rejecting EMI since it is nearly always a common mode signal. And, again, the single point chassis ground is another important safeguard.

INPUT IMPEDANCE OF DIFFERENCE AMPLIFIERS

Differential Or Common Mode Positive Input Impedance

Figure 1 shows a standard difference amplifier. On the surface it looks simple enough, but close examination of the available literature and text books covering difference amplifiers reveals a subtle complexity that causes great confusion. It all revolves around the fact that the negative input impedance is *dynamic* while the positive input impedance is *static*. That is, the positive impedance is unrelated to the type of input signal or feedback resistor. It is always equal to equation (1), whether the input signal is differential or common mode.

$$R_{in} (+) = R_3 + R_4 \quad (1)$$

This is not true for the negative impedance: the negative impedance is *different* depending on the nature of the input signal. This is due to the dynamic servo loop nature of the feedback circuitry, which is always working to keep the voltage at the negative input to the op amp equal to the voltage at the positive input (the virtual short concept).

Differential Mode Negative Input Impedance

The negative differential input impedance is not equal to resistor R_2 as popularly believed. It is equal to equation (2):

$$R_{in} (-) = \frac{R_3 + R_4}{2R_4 + R_3} R_2 \quad (2)$$

For the most common design where all four resistors are equal, R_{in} equals $2/3 R_2$.

An example may be the best way to understand why the impedance is $2/3 R_2$. Assume there is +1 volt at the tip and -1 volt at the ring of the input jack and that all resistors are equal. There is therefore, $+1/2$ volt at the

plus input port of the op amp—and the minus port (virtual short). The ring input current equals the voltage across resistor R_2 , divided by R_2 . The voltage equals $+1/2V$ —(-1V), or $3/2$ volts. The current is then $3/2 R_2$. The ring input impedance is, by definition, the input voltage (-1V) divided by the input current ($-3/2R_2$), (the minus sign indicates current is flowing out of the circuit) giving an answer of $2/3 R_2$.

Common Mode Negative Input Impedance

The negative common mode input impedance is found from equation (3):

$$R_{in} (-) = \frac{R_3 + R_4}{R_3} R_2 \quad (3)$$

For the case where all four resistors are equal, R_{in} equals $2 R_2$.

Again, an example may be the best way to illustrate equation (3). For common mode input signals, the op amp works to reject them such that the output will be zero. If the output is zero and the input currents to the op amp may be ignored (a reasonable assumption with modern IC's) then the negative input impedance is simply $R_1 + R_2$. Equation (3) is the general equation found by writing the input current equation as before, and dividing it into the negative input voltage. It serves to illustrate the negative input bootstrap effect caused by the feedback.

Comment

For further confusion (clarification?) it is necessary to reiterate that both cases occur simultaneously. That is, the negative input impedance is lower than R_2 for all differential input signals and higher than R_2 for all common mode signals—at the same time. Audio input signals are almost always a mixture of signal (differential input) and interference (common mode input).

REFERENCES

1. International Electrotechnical Commission Standard No. 268. Part 12, titled. "Circular Connectors for Broadcast and Similar Use", 1975. Signers were Australia, Belgium, Canada, Denmark, Egypt, Hungary, Israel, Japan, Netherlands, Norway, Romania, South Africa, Sweden, Switzerland, Turkey, the UK, and USA.

2. British Standards Institution No. BS 5428. Part 5. Sec. 3, 1980/81.

RaneNote 104**Microphones and
the RA 27**

- **RA 27 Microphone Characteristics**
- **Powering**
- **Line Driving Capability**
- **Using Other Microphones with the RA 27**

**Dennis Bohn
Rane Corporation**

© 1983 Rane Corporation

INTRODUCTION

It has become obvious to many users of the RA 27 Realtime Analyzer that the RA 27 microphone is an excellent omnidirectional with uses other than room equalization. It has become equally obvious to many others that they have a need to use different microphones with the RA 27. Such is the world. This Rane Note attempts to describe in gritty, solder-smelly detail how to accomplish both tasks.

RA 27 MICROPHONE CHARACTERISTICS

The RA 27 microphone is designed around a back electret condenser cartridge with true omnidirectional pick up characteristics. It has a built-in N-channel JFET amplifier and must be externally powered from a 1.5-10 VDC source to operate. The nominal sensitivity is -64 dB (0 dB=V/microbar) with an output impedance of 1800 ohms and a maximum current consumption of 500 microamps when powered from a 2 VDC source. Its frequency response is essentially flat out to 16 kHz where it typically displays a rise of 2-4dB before rolloff at 20 kHz. And it handles a maximum SPL of 140 dB. (See the MIC 1 Data Sheet for complete details.)




Figure 1. RA 27 Mic Battery Power Supply

POWERING

Figure 1 shows a simple power supply that operates the RA 27 microphone. It can be built into any small plastic or metal enclosure in just a few minutes once all the parts are on hand (that's always the rub, isn't it?). By using a standard 1/4" TRS jack as shown, you get a battery on-off switch for free. Soldering the negative side of the battery to the ring instead of directly to ground guarantees that the battery is only *on* when the mic is plugged *in*. (Plugging the TS plug of the RA 27 mic into the TRS input jack causes the ring to be shorted to the sleeve which completes the battery circuit to ground.

Resistor R1 is the JFET load resistor and must be 2.2k ohms (either 5% or 10% tolerance is fine). Capacitor C1 blocks the DC voltage so that the output feeds any mic level input. Its value may be anything larger than 1 microfarad, with a voltage rating equal to or greater than the battery used. Be sure to observe the proper polarity when soldering in place. Resistor R2 prevents reverse charge from building up on C1 which would cause snapping and popping when connecting to other equipment and may be any value greater than 20k ohms.

Using another 1/4" TRS jack on the output and wiring as shown, allows a floating output with all of its advantages. A standard 1/4" TS jack may be substituted if desired.

The batteries used may be any size from AA to D. The larger the size the longer they will last. Also since the RA 27 may be powered from 1.5-10 VDC, a single battery within this range may be substituted; for example, the ubiquitous 9V transistor battery.

LINE DRIVING CAPABILITY

Since the output impedance of the RA 27 mic is fairly high at 1800 ohms, exercise some precaution before attempting to drive long lines. The best way to define a "long line" is by total capacitance rather than total length. Since lines will exhibit so many picofarads per foot, they are really one and the same thing, but it is the total capacitance that is the problem—not the length. With an output impedance of 1800 ohms, the RA 27 mic exhibits a 3 dB loss of 16 kHz when driving 5500 pF. So something around 5000 pF becomes a good rule-of-thumb. Anything greater starts showing up as a loss of high frequencies. Therefore, if your cable has 50 pF per foot of capacitance (a common figure), then you can safely drive 100 feet with the RA 27 microphone.




Figure 2. RA 27 Dynamic Mic Adaptor

USING OTHER MICROPHONES WITH THE RA 27

Since the 1/4" TRS mic INPUT jack on the RA 27 has DC voltage (+2.2 V) internally wired to the tip contact, you must use mucho discretion before plugging any other manufacturers microphone directly into the RA 27. Most microphones take a great deal of offense of having DC voltage applied to their delicate little bodies; this being true, then what to do? Block the DC, that's what!

This can be done in a couple of easy ways. The first is to simply use one of several line matching transformer adapters available for adapting 3-pin mic connectors to 1/4" TS jacks. These are the "in-line" type built into the connector assembly. Examples would be the Shure A95FP, Radio Shack 274-16, or Cal Switch PS-42.

The transformer blocks the DC voltage from the microphone and makes for a very convenient adapter. Since the voltage applied to the transformer secondary is very small and the DC current is limited to about 1mA by the internal 2200 ohm JFET load resistor, the transformer still operates just fine.

A second method for do-it-yourselfers who would like to save the cost of the transformer adapter is to wire up a blocking capacitor as shown in Figure 2. You will need a 3-pin female connector, a 1/4" TS jack, a 10k ohm or larger resistor, and a capacitor rated 4.7 microfarads at 3 V or larger. Wire them together as shown, observing the proper capacitor polarity. If available, you can use a Switchcraft 3-pin to 1/4" TRS adapter model number 386A and put the capacitor and resistor inside. Wire the red and black wires as shown in Figure 2. Otherwise wire the circuit into a small plastic or metal enclosure.

Happy soldering.

NOTE 106



USING THE HC 6 HEADPHONE CONSOLE AS A BALANCED LINE DISTRIBUTION AMPLIFIER


INTRODUCTION

The Model HC 6 Headphone Console is a six channel headphone amplifier, consisting of six stereo amplifiers, two master stereo inputs and six individual stereo inputs. The two master stereo inputs feed a separate left and right bus to each of the six headphone amps (see Figure 1). In addition, there are individual stereo inputs to each headphone amplifier that disconnect that particular amp from the master stereo bus and allows a separate mix to be used.

From time to time, we are asked if the HC 6 may be configured for use as a balanced line distribution amplifier. The answer is, yes, with a special input jack arrangement.

Dennis Bohn
V.P. Research & Development


Professional Audio Products



HC 6 DISTRIBUTION AMPLIFIER

The requirements for using the HC 6 as a balanced line distribution amplifier are that *the input drive must be a balanced (or floating) mono signal. It will not work with a stereo balanced pair or an unbalanced mono signal.*

The trick is to wire a special input jack assembly per Figure 2. What is required is for the tip of the input jack to feed the tip of one of the input males and the RING of the other. And for the ring of the input jack to feed the ring of one and TIP of the other. What this does is drive the two inputs with opposite polarities, such that the left input is driven positive while the right input is driven negative, and vice-versa. This causes all of the left headphone amps to drive the output lines positive while all the right headphone amps drive the other side of the line negative, and vice-versa. Which creates six balanced line drivers, all fed from the one common input line. So, with this special input jack, the HC 6 is now configured as a one-to-six distribution amplifier.



SPECIFICATIONS

The new specifications are listed below and are quite good, with perhaps one exception. The signal-to-noise ratio is nothing to write home about, but this is not unexpected. The HC 6 is a power amp. In fact, it is 12 mini power amps capable of driving headphones to ear splitting levels. This tends to degrade the signal-to-noise ratio. But 74dB is quite respectable and will not be a problem in the majority of distribution amplifier applications.

For applications requiring a distribution amplifier driven from stereo lines (balanced or unbalanced), or several individual inputs (balanced or unbalanced), or requiring exceptional signal-to-noise and THD performance, then the user is directed to Rane's Model SM 26 Splitter Mixer. The SM 26 boasts S/N ratios of 90dB, with THD levels below .01%.

SPECIFICATIONS

Nominal Input Level: -10dBV(316mV)
 Maximum Input Level: +20dBm (7.75V)
 Input Impedance: 20k Ω
 Maximum Available Gain: 34dB
 Output Impedance: 15 Ω
 Maximum Output Level: +25dBm (14V)
 Minimum Load Impedance: 32 Ω
 Typical Load Impedance: 10k Ω
 THD: less than .05%, 20-20kHz (+20dBm, 10k Ω load)
 Frequency Response: 20-20kHz +0/-3dB
 S/N: 74dB re +4dBm (flat, 20kHz BW, 12dB gain)

Linkwitz-Riley Crossovers

- **ZERO LOBING ERROR**
- **24 dB/OCTAVE SLOPES**
- **STATE VARIABLE SOLUTION**
- **TIME CORRECTION**

The assumption is made that the reader is familiar with active crossovers and how they are used in professional sound systems. For those who are not and want to review the basics, one of the best references will be found in an article entitled. "Crossover Basics". by Richard Chinn, in the September 1986 issue of STEREO REVIEW.

Dennis Bohn
Rane Corporation

RaneNote 107
© 1983 Rane Corporation

INTRODUCTION

What's a Linkwitz-Riley, and why do I care? First off, its not "what's", but "who": Siegfried Linkwitz and Russ Riley are two Hewlett-Packard R&D engineers who wrote a paper¹ in 1976 describing a better mousetrap in crossover design. Largely ignored (or unread) for the past several years, it is now receiving the attention it deserved in 1976. Typical of most truly useful technical papers, it is very straight-forward and unassuming. A product of careful analytical attention to details, with a wonderfully simple solution.

The, "why do I care?", part is easily answered by stating that a Linkwitz-Riley crossover will give you a clearer and more accurate sound system. Period. It will automatically clean up the messiness that mars most systems at their crucial crossover points. (It is at the crossover points that most systems lose it.)

It is seldom whether to cross over, but rather, how to cross over. Active crossovers have proliferated over the past few years at a rate equal to the proverbial lucky charm. The potential crossover buyer must choose from among a dozen different manufacturers and designs. Some are adequate; some are even good; but none seem to offer just the right mix of features, technology and cost. Until now.

An attempt will be made within this Note to present the essence of a Linkwitz-Riley design, and to introduce Rane's answer to a truly affordable crossover that features the very best technology, with exactly the right features.

A 4th-order state variable active filter² has been developed by Rane Corporation to implement the Linkwitz-Riley alignment for crossover coefficients. In addition to the active crossover, the unit features a variable time delay circuit so the user may effectively "move" the drivers into front-to-back alignment. With both these tools, the professional sound person now has the means to smooth out and perfect the crucial crossover region, resulting in a sound system that exhibits unsurpassed clarity and accuracy.

A PERFECT CROSSOVER

Mother nature gets the blame. Another universe, another system of physics, and the quest for a perfect crossover might not be so difficult. But we exist here and must make the best of what we have. And what we have is the physics of sound, and of electromagnetic transformation systems that obey these physics.

A perfect crossover, in essence, is no crossover at all. It would be one driver that could reproduce all frequencies equally well. Since we cannot have that, then second best would be multiple speakers, along the same axis, with sound being emitted from the same point, i.e., a coaxial speaker that has no time shift between drivers. This gets closer to being possible, but still is elusive. Third best, and this is where we really begin, is multiple drivers mounted one above the other with no time shift, i.e., non-coincident drivers adjusted front-to-rear to compensate for their different points of sound propagation. Each driver would be fed only the frequencies it is capable of reproducing. The frequency dividing network would be, in reality, a frequency gate. It would have no phase shift or time delay. It's amplitude response would be absolutely flat and its roll-off characteristics would be the proverbial brick wall. (Brings a tear to your eye, doesn't it?)

Using digital technology, such a crossover is possible, but not at a price that is acceptable to most working musicians. What is possible at an affordable price is a very good compromise known as the Linkwitz-Riley crossover.

LINKWITZ-RILEY CROSSOVER

What distinguishes the Linkwitz-Riley crossover design from others is its perfect combined radiation pattern of the two drivers at the crossover point. Stanley P. Lipshitz³ has coined the term "lobing error" to describe this crossover characteristic. It's a good term and should spread through the industry as the standard. It derives from the examination of the acoustic output plots (at crossover) of the combined radiation pattern of the two drivers (see Figures 1 & 2). If it is not perfect, the pattern forms a lobe that exhibits an off-axis frequency dependent tilt with severe amplitude peaking.

Interpretation of Figure 1 is not particularly obvious. Let's back up a minute and add some more details. For simplicity, only a two way system is being modeled. The two drivers are mounted along the vertical center of the enclosure (there is no side-to-side displacement, i.e., one driver is mounted on top of the other.) Any front-to-back time delay between drivers has been corrected. The figure shown is a polar plot of the side-view, i.e., the angles are vertical angles.

It is only the vertical displacement sound field that is at issue here. All of the popular crossover types (constant voltage⁴, Butterworth all-pass⁵, etc.) are well behaved along the horizontal on-axis plane. To illustrate the geometry involved here, imagine attaching a string to the speaker at the mid-point between the drivers. Position the speaker such that the mid-point is exactly at ear level. Now pull the string taut and hold it up to your nose (go on, no one's looking). The string should be parallel to the floor. Holding the string tight, move to the left and right. This is the horizontal on-axis plane. Along this listening plane, all of the classic crossover designs exhibit no problems. It is when you lower or raise your head below or above this plane that the problems arise. This is the crux of Siegfried Linkwitz's contribution to crossover design. After all these years and as hard as it is to believe, he was the first person to publish an analysis of what happens off-axis with non-coincident drivers (not-coaxial). (Others may have done it before, but it was never made public record.)

Figure 1A represents a side view of the combined acoustic radiation pattern of the two drivers emitting the same single frequency. That is, a plot of what is going on at the single crossover frequency all along the vertical plane. The pattern shown is for the popular 18 dB/octave Butterworth all-pass design with a crossover frequency of 1700 Hz and drivers mounted 7 inches apart¹.

What is seen is a series of peaking and cancellation nodes. Back to the string. Holding it taut again and parallel to the floor puts you on-axis. Figure 1A tells us that the magnitude of the emitted 1700 Hz tone will be 0 dB (a nominal reference point). As you lower your head, the tone will increase in loudness until a 3 dB peak is reached at 15 degrees below parallel. Raising your head above the on-axis line will cause a reduction in magnitude until 15 degrees is reached where there will be a complete cancellation of the tone. There is another cancellation axis located 49 degrees below the on-axis. Figure 1B depicts the frequency response of the three axes for reference.




Figure 1A.




Figure 2. Linkwitz-Riley Radiation Response at Crossover




Figure 1B.

Figure 1. Butterworth All-Pass Design Radiation Pattern at Crossover.

For a constant voltage design, the response looks worse, having a 6 dB peaking axis located at -20 degrees and the cancellation axes at +10 and -56 degrees, respectively. The peaking axis tilts toward the lagging driver in both cases, due to phase shift between the two crossover outputs.

The cancellation nodes are not due to the crossover design, they are due to the vertically displaced drivers. (The crossover design controls *where* cancellation nodes occur, not *that* they occur.) The fact that the drivers are not coaxial means that any vertical deviation from the on-axis line will result in a slight, but very significant differences in path lengths to the listener. This difference in distance traveled is effectively a phase shift between drivers. And this causes cancellation nodes — the greater the distance between drivers, the more nodes.

In distinct contrast to these examples is Figure 2, where the combined response of a Linkwitz-Riley crossover design is shown. There is no tilt and no peaking. Just a perfect response whose only limitation is the dispersion characteristics of the drivers used. The main contributor to this ideal response is the in-phase relationship between the crossover outputs.

Two of the cancellation nodes are still present but are well defined and always symmetrical about the on-axis plane. Their location changes with crossover frequency and driver mounting geometry (distance between drivers). With the other designs, the peaking and cancellation axes change with frequency and driver spacing.

Let's drop the string and move out into the audience to see how these cancellation and peaking nodes affect things. Figure 3 shows a terribly simplified, but not too inaccurate stage-audience relationship with the characteristics of Figure 1 added.

The band is cooking and then comes to a musical break. All eyes are on the flautist, who immediately goes into her world-famous 1700 Hz solo. So what happens? The people in the middle hear it sweet, while those up front are blown out of their seats, and those in the back are wondering what the hell's all the fuss!

Figure 4 shows the identical situation but with the Linkwitz-Riley characteristics of Figure 2 added. Now the people in the middle still hear everything sweet, but those up front are not blown away, and those in the back understand the fuss!

I think you get the point.

Now let's get real. I mean really real. The system isn't two way, it's four way. There isn't one enclosure, there are sixteen. No way are the drivers 7 inches apart — try 27. And time corrected? Forget it.

Can you even begin to imagine what the vertical off-axis response will look like with classic crossover designs? The further apart the drivers are, the greater the number of peaks and cancellations, resulting in a multi-lobe radiation pattern. Each crossover frequency will have its own set of patterns, complicated by each enclosure contributing even more patterns. And so on.

(For large driver spacing the Linkwitz-Riley design will have as many lobes as other designs, except that the peaks are always 0 dB, and the main lobe is always on-axis.)

Note that all this is dealing with the direct sound field, no multiple secondary arrivals or room interference or reverberation times are being considered. Is it any wonder that when you move your realtime analyzer microphone 3 feet you get a totally different response?

Now let me state clearly that using a Linkwitz-Riley crossover will not solve all these problems. But it will go a long, long way toward that goal.

The other outstanding characteristic of the Linkwitz-Riley alignment is the rolloff rate of 24 dB/octave (Figure 5). With such a sharp drop-off, drivers can be operated closer to their theoretical crossover points without the induced distortion

normally caused by frequencies lying outside their capabilities. Frequencies just one octave away from the crossover point are already attenuated by 24 dB (a factor of about 1/16). The importance of sharp cutoff rate and in-phase frequency response of the crossover circuitry cannot be over-stressed in contributing to smooth overall system response.

A summary of the characteristics of a Linkwitz-Riley crossover reads:

1. Absolutely flat amplitude response through out the pass-band with a steep 24 dB/octave rolloff rate after the crossover point.
2. The acoustic sum of the two driver responses is unity at crossover. (Amplitude response of each is -6 dB at crossover, i.e., there is no peaking in the summed acoustic output)
3. Zero phase difference between drivers at crossover. (Lobing error equals zero, i.e., no tilt to the polar radiation pattern.) In addition, the phase difference of zero degrees through crossover places the lobe of the summed acoustic output on axis at all frequencies.
4. The low pass and high pass outputs are everywhere in phase. (This guarantees symmetry of the polar response about the crossover point)
5. All drivers are always wired the same (in phase).

A casual reading of the above list may suggest that this is, indeed, the perfect crossover. But such is not so. The wrinkle involves what is known as "linear phase". A Linkwitz-Riley crossover alignment is not linear phase: meaning that the amount of phase shift is a function of frequency. Or, put into time domain terms, the amount of time delay through the filter is not constant for all frequencies. Which means that some frequencies are delayed more than others. (In technical terms, the network has a frequency-dependent group delay, but with a very gradually changing characteristic.)

Is this a problem? Specifically, is this an audible "problem"? In a word, no.

Much research has been done on this question⁶⁻⁹, with approximately the same conclusions: given a slowly changing non-linear phase system, the audible results are so minimal as to be non-existent; especially in the face of all of the other system non-linearities. And with real-world music sources (remember music?), it is not audible at all.




Figure 3. Butterworth All-Pass Crossover Stage-Audience Relationship Crossovers-4




Figure 4. Linkwitz-Riley Crossover Stage Audience Relationship

STATE VARIABLE SOLUTION

One of the many attractions of the Linkwitz-Riley design is its utter simplicity, requiring only two standard 2nd-order Butterworth filters in series. The complexities occur when adjustable crossover frequencies are required.

After examining and rejecting all of the standard approaches to accomplish this task, Rane developed a 4th-order state-variable filter specifically for implementing the Linkwitz-Riley crossover. The state-variable topology was chosen over other designs mainly for the following reasons:

1. It provides simultaneous high-pass and low-pass outputs that are always at exactly the same frequency.
2. Changing frequencies can be done simultaneously on the high-pass and low-pass outputs without any changes in amplitude or Q (quality factor).
3. The sensitivities of the filter are very low. (Sensitivity is a measure of the effects of non-ideal components on an otherwise, ideal response.)
4. It offers the most cost-effective way to implement two 4th-order responses with continuously variable crossover frequencies.

TIME OR PHASE CORRECTION

Implicit in the development of the theory of a Linkwitz-Riley crossover design is the key assumption that the sound from each driver radiates from the same exact vertical plane, i.e., that the drivers have no time delay with respect to each other. The crossover then prohibits any lobing errors as the sound advances forward simultaneously from the two drivers. Figure 6 illustrates such a front-to-back displacement, which causes the lobing error shown in Figure 7a.

A Linkwitz-Riley crossover applied to drivers that are not time-corrected loses most of its magic. The lobing error is no longer zero; it exhibits a frequency dependent tilt with magnitude errors as shown in Figure 7b.

This being the case, Rane incorporates either adjustable time delay or phase shift circuits into its Linkwitz-Riley crossovers.




Figure 5. Frequency Response of 4th-order Linkwitz-Riley Active Crossover




Figure 6. Driver Displacement




Figure 7. Adding Delay to the Forward Driver Time-Aligns the Phase of Both Drivers, Reducing Lobing Error.

Figure 7a Without Time Alignment.

Figure 7b. With Time Alignment.

REFERENCES

1. S. H. Linkwitz "Active Crossover Networks for Non-coincident Drivers," *J. Audio Eng. Soc.*, vol. 24, pp. 2-8 (Jan/Feb 1976).
2. D. Bohn. "A Fourth Order State Variable Filter for Linkwitz-Riley Active Crossover Designs," presented at the 74th Convention of the Audio Engineering Society, New York, Oct. 9-12, 1983, preprint no. 2011.
3. S. P. Lipshitz and J. Vanderkooy, "A Family of Linear-Phase Crossover Networks of High Slope Derived by Time Delay," *J. Audio Eng. Soc.*, vol. 31, pp. 2-20 (Jan/Feb 1983).
4. R. H. Small, "Constant-Voltage Crossover Network Design," *J. Audio Eng. Soc.*, vol. 19, pp. 12-19 (Jan 1971).
5. J.R. Ashley and A. L. Kaminsky. "Active and Passive Filters as Loudspeaker Crossover Networks," *J. Audio Eng. Soc.*, vol. 19, pp. 494-502 (June 1971).
6. B. B. Bauer, "Audibility of Phase Shift," *Wireless World*, (Apr. 1974).
7. S. P. Lipshitz, M. Pocock, and J. Vanderkooy. "On the Audibility of Midrange Phase Distortion in Audio Systems," *J. Audio Eng. Soc.*, vol. 30, pp. 580-595 (Sep 1982).
8. R. Lee. "Is Linear Phase Worthwhile," presented at the 68th Convention of the Audio Engineering Society, Hamburg, Mar 17-20, 1981, preprint no. 1732.
9. H. Suzuke. S. Morita. and T. Shindo. "On the Perception of Phase Distortion," *J. Audio Eng. Soc.*, vol. 28, no. 9, pp. 570-574 (Sep 1980).
10. RaneNote, "Linkwitz-Riley Crossovers Up to 8th-Order," 1989.

NOTE 111

Professional Audio Products



FEEDBACK FINDER FOR MATRIX MIXERS

INTRODUCTION

Feedback during a performance is every sound person's nightmare. We've all heard it; we've all had it happen; and we all hate it. This Rane Note describes a clever technique suggested by Spence Burton (Audio Engineer from Richmond, VA) for feedback monitoring of all outputs of matrix mixers with an RE 27 Realtime Equalizer or RA 27 Realtime Analyzer.

APPLICATION

A common application for a matrix mixer uses it to produce six different mixes for on-stage live music monitoring. Often, each of the output channels will incorporate an equalizer in the EQ loop to facilitate the stage speaker. This allows a dedicated one-third octave equalizer for feedback control of each speaker.

Most matrix mixers have separate cue pushbuttons, or a selector switch for each output, arranged such that with an equalizer in the loop the effect of that equalizer can be monitored through headphones plugged into the mixer. Pushing each cue pushbutton allows you to hear what is going to each speaker. Spence came up with an idea that also allows you to see what is going to each speaker.

Spence Burton
Audio Engineer

Dennis Bohn
Rane Corporation

FEEDBACK FINDER

The trick is to use a wye adapter from the mixer headphone jack. One side goes to your headphones and the other side is transformer coupled into the mic input of your RE 27 or RA 27 Realtime Equalizer/Analyzer. Brilliant.

With this arrangement, anytime a cue pushbutton is pressed, you hear that channel through your headphones and simultaneously you see the response on the realtime display. If a feedback situation occurs you can identify the problem area by watching the display go into the red. You turn the display down (not the sound) and watch the lights drop out of their red condition. **The last band out of the red zone is the feedback frequency band.** Don't worry about time here, the musician will back out of the hot zone immediately, which gives you enough time to identify the frequency.

A typical arrangement might use a mixer with six matrix outputs, one RA 27, three ME 60's and an MA 6S. The equalizers are patched into the mixer EQ loops using a standard 1/4" TRS plug with one end split out into two TS 1/4" plugs. The tip from the mixer goes into the input of the equalizer and the ring from the mixer goes into the output of the equalizer. This follows the normal convention of "tip send/ring return" wiring of mixing board EQ loops. Each matrix output of the mixer then feeds the MA 6S Power Amplifier.

DETAILS

Figure 1 shows the fun stuff. You need a regular 1/4" TRS wye adapter having one male plug and two female jacks. Plug the adapter into the mixer Headphone Output jack and your headphones into one side of the wye. Now for the soldering.

Refer to Figure 1 for the detail drawing of the special box. This is a line matching transformer wired to a mono input jack and a mono output jack. Note that the **ring** connection of the input jack is left open—**do not solder it to anything or you will short out one side of the mixer headphone amplifier.** Also leave the center tap connection on the primary open.

This transformer does two things: it drops the level down to something compatible with the RA 27 mic input stage, and it isolates the DC power on the mic input from the mixer headphone amp. (The DC voltage is small enough that it doesn't hurt the transformer.)

Common Radio Shack part numbers are given for reference, but any 1/4" TS jacks will work as well as just about any kind of line matching transformer. (You are matching line level to mic input level, i.e., high-Z to low-Z.)

One final tip is to mark the equalizers "1" through "6" in agreement with the mixer matrix output channels, so that during the performance you know which equalizer corresponds to each cue pushbutton.




Figure 1. Feedback Finder for Matrix Mixers

Linkwitz-Riley Active Crossovers up to 8th-Order: An Overview

- Linkwitz-Riley Alignment
- Butterworth Alignment
- 1st to 8th Order Filters
- Vector Diagrams
- Transient Response
- Phase Response

**Dennis Bohn
Rane Corporation**

**RaneNote 119
© 1989 Rane Corporation (rev. 12/01)**

INTRODUCTION

In 1976, Siegfried Linkwitz published his famous paper [1] on active crossovers for non-coincident drivers. In it, he credited Russ Riley (a co-worker and friend) with contributing the idea that cascaded Butterworth filters met all Linkwitz's crossover requirements. Their efforts became known as the Linkwitz-Riley crossover alignment. In 1983, the first commercially available Linkwitz-Riley active crossovers appeared from Sundholm and Rane [2].

Today, the de facto standard for professional audio active crossovers is the 4th-order Linkwitz-Riley (LR-4) design. Offering in-phase outputs and steep 24 dB/octave slopes, the LR-4 alignment gave users the tool necessary to scale the next step toward the elusive goal of perfect sound.

Now a new tool is available: the 8th-order Linkwitz-Riley (LR-8) active crossover [3]. With incredibly steep slopes of 48 dB/octave, the LR-8 stands at the door waiting for its turn at further sound improvements. Using a LR-8 cuts the already narrow LR-4 crossover region in half. Just one octave away from the crossover frequency the response is down 48 dB. The LR-8 represents a major step closer to the proverbial brick wall, with its straight line crossover region.

Before exploring the advantages of LR-8 designs, it is instructional to review just what Linkwitz-Riley alignments are, and how they differ from traditional Butterworth designs of old.




Figure 1. 1st order crossover network

1st-Order Network

It begins with a resistor and a capacitor. It never gets more complicated than that—just resistors and capacitors: lots and lots of resistors and capacitors. Resistors are the great emancipators of electronics; they are free of frequency dependence. They dissipate energy without frequency prejudice. All frequencies treated equally. Capacitors, on the other hand, selectively absorb energy; they store it, to be released at a later time. While resistors react instantly to any voltage changes within a circuit, capacitors take time to charge and discharge.

Capacitors are so frequency dependent, they **only** pass signals with frequency associated with them. Direct-current (what we call zero frequency) will not pass at all; while, at the other end of the spectrum, very high frequencies will not absorb. Capacitors act like a piece of wire to high frequencies; hardly there at all.

We use these facts to create a crossover network. Figure 1 shows such a circuit. By interchanging the positions of the resistor and capacitor, low-pass (low frequencies = LF) and high-pass (high frequencies = HF) filters result. For the low-pass case (LF), the capacitor ignores low frequencies and shunts all high frequencies to ground. For the high-pass case (HF), the opposite occurs. All low frequencies are blocked and only high frequencies are passed.




Figure 2. 1st-order amplitude response

1st-Order Amplitude Response

Using 1kHz as an example and plotting the amplitude versus frequency response (Figure 2) reveals the expected low-pass and high-pass shapes. Figure 2 shows that the 1st-order circuit exhibits 6 dB/octave slopes. Also, that 6 dB/octave equals 20 dB/decade. Both ways of expressing steepness are useful and should be memorized. The rule is: **each order, or degree, of a filter increases the slopes by 6 dB/octave or 20 dB/decade**. So, for example, a 4th-order (or 4th-degree—interchangeable terms) circuit has 24 dB/octave (4x6 dB/octave) or 80 dB/decade (4x20 dB/decade) slopes.

Using equal valued resistors and capacitors in each of the circuits causes the amplitude responses to ‘cross over’ at one particular frequency where their respective -3 dB points intersect. This point represents the attenuation effect resulting when the impedance of the capacitor equals the resistance of the resistor.

The equivalent multiplying factor for -3 dB is .707, i.e., a signal attenuated by 3 dB will be .707 times the original in level. Ohms law tells us that if the voltage is multiplied by .707, then the current will also be multiplied by .707. Power is calculated by multiplying voltage times current. Therefore, a voltage multiplied by .707, and a current multiplied by .707, equals 0.5 power. So the -3 dB points represent the one-half power point—a useful reference.

Lastly, Figure 2 shows the flat amplitude response resulting from summing the LF and HF outputs together. This is called **constant voltage**, since the result of adding the two output voltages together equals a constant. The 1st-order case is ideal in that **constant power** also results. Constant-power refers to the summed power response for each loudspeaker driver operating at the crossover frequency. This, too, results in a constant. Since each driver operates at ½ power at the crossover frequency, their sum equals one—or unity, a constant.

1st-Order Phase Response

Much is learned by examining the phase shift behavior (Figure 3) of the 1st-order circuit. The upper curve is the HF output and the lower curve is the LF output. The HF curve starts at +90° phase shift at DC, reduces to +45° at the crossover frequency and then levels out at 0° for high fre-




Figure 3. 1st-order response

quencies. The LF curve starts with 0° phase shift at DC, has -45° at the crossover frequency and levels out at -90° for high frequencies.

Because of its reactive (energy storing) nature each capacitor in a circuit contributes 90° of phase shift, either positive or negative depending upon its application. Since the HF section places the capacitor directly in the signal path, this circuit starts out with $+90^\circ$ phase shift. This is called **phase lead**. The LF section, which starts out with 0° and eventually becomes -90° is called **phase lag**.

Examination of Figure 3 allows us to formulate a new rule: **each order, or degree, of a crossover network contributes $\pm 45^\circ$ of phase shift at the crossover frequency** (positive for the HF output and negative for the LF output).

Once again, Figure 3 shows the idealized nature of the 1st-order case. Here the result of summing the outputs together produces 0° phase shift. Which is to say that the summed amplitude and phase shift of a 1st-order crossover equals that of a piece of wire.

1st-Order Group Delay Response

We shall return to our rules shortly, but first the concept of **group delay** needs to be introduced. Group delay is the term given to the ratio of an incremental change in phase shift divided by the associated incremental change in frequency (from calculus, this is the first-derivative). The units for group delay are seconds. If the phase shift is **linear**, i.e., a constant rate of change per frequency step, then the incremental ratio (first-derivative) will be constant. We therefore refer to a circuit with linear phase shift as having **constant group delay**.

Group delay is a useful figure of merit for identifying linear phase circuits. Figure 4 shows the group delay response for the Figure 1 1st-order crossover circuit. Constant group delay extends out to the crossover region where it gradually rolls off (both outputs are identical). The summed response is, again, that of a piece of wire.

The importance of constant group delay is its ability to predict the behavior of the LF output **step response**. A circuit with constant group delay (linear phase shift) shows no overshoot or associated damping time to a sudden change (step) in input level (Figure 5). The circuit reacts smoothly to




Figure 4. 1st-order group delay response

the sudden change by rising steadily to meet the new level. It does not go beyond the new level and require time to settle back. We also refer to the step response as the **transient** response of the circuit. The transient response of the summed outputs is perfect since their sum is perfectly equal to one.

For clarity purposes normally only the step response of the LF network is shown. Nothing is learned by examining the step response of the HF network. A step response represents a transition from one DC level to another DC level, in this case, from -1 volt to $+1$ volt. A HF network, by definition, does not pass DC (neither does a loudspeaker), so nothing particularly relevant is learned by examining its step response. To illustrate this, Figure 5 shows the HF step response. It begins and ends with zero output since it cannot pass DC. The sharp edge of the input step, however, contains much high frequency material, which the HF network passes. So, it begins at zero, passes the high frequencies as a pulse, and returns to zero.

The HF and LF outputs are the exact complement of each other. Their sum equals the input step exactly as seen in Figure 5. Still, we learn everything we need to know by examining only the LF step response; looking for overshoot and ringing. From now on, just the LF output will be shown.

Vector Diagrams

A vector is a graphical thing (now we're getting technical) with magnitude and direction. We can use vectors to produce diagrams representing the instantaneous phase shift and amplitude behavior of electrical circuits. In essence, we freeze the circuit for a moment of time to examine complex relationships.

We shall now apply our two rules to produce a vector diagram showing the relative phase shift and amplitude performance for the 1st-order crossover network at the single crossover frequency (Figure 6a). By convention, 0° points right, $+90^\circ$ points up, -90° points down, and $\pm 180^\circ$ points left. From Figures 2 & 3 we know the HF output amplitude is -3 dB with $+45^\circ$ of phase shift at 1 kHz, and the LF output is -3 dB with -45° phase shift. Figure 6a represents the vectors as being $.707$ long (relative to a normalized unity vector) and rotated up and down 45° . This shows us the relative phase difference between the two outputs equals 90° .




Figure 5. 1st-order transient response




Figure 6. 1st-order Vector and 2nd through 4th-order Butterworth Vector Diagrams

Next we do vector addition to show the summed results. Vector addition involves nothing more complex than mentally moving one of the vectors to the end of the other and connecting the center to this new end point (it is like constructing a parallelogram). Doing this, results in a new vector with a length equal to 1 and an angle of 0° . This tells us the recombined outputs of the HF and LF networks produce constant voltage (i.e., a vector equal to 1), and is in phase with the original input of the circuit (i.e., a vector with 0° phase rotation).

The 1st-order case is ideal when summed. It yields a piece of wire. Since the responses are the exact mirror images of each other, they cancel when summed, thus behaving as if neither was there in the first place. Unfortunately, all optimized higher order versions yield flat voltage/power response, group delay or phase shift, **but not all at once**. Hence, the existence of different alignments and resultant compromises.

2nd, 3rd & 4th-Order Butterworth Filters

There are many types of filters (most named after mathematicians). Each displays a unique amplitude characteristic throughout the passband. Of these, only **Butterworth** filters have an absolutely flat amplitude response. For this reason, Butterworth filters are the most popular for crossover use. Butterworth filters obey our two rules, so we can diagram them for the 2nd, 3rd and 4th-order cases (Figures 6b-6d). The 2nd-order case has $\pm 90^\circ$ phase shift as shown. This results in the outputs being 180° out of phase. Vector addition for this case produces a zero length vector, or complete cancellation. The popular way around this is to reverse the wiring on one of the drivers (or, if available, electronically inverting the phase at the crossover). This produces a resultant vector 90° out of phase with the input and 3 dB (1.414 equals +3 dB) longer. This means there will be a 3 dB amplitude bump at the crossover region for the combined signals.

The 3rd-order Butterworth case (Figure 6c) mimics the 1st-order case at the crossover frequency, except rotated 180° . Hence, we see the HF vector rotated up 135° ($3 \times 45^\circ$) and the LF vector rotated down the same amount. The phase shift between outputs is still 90° . The resultant is constant voltage (unity) but 180° out-of-phase with the input.

Linkwitz-Riley Crossovers-4

The 4th-order Butterworth diagram (Figure 6d) shows the HF vector rotated up 180° and the LF vector rotated down the same amount. The phase difference between outputs is now zero, but the resultant is +3 dB and 180° out-of-phase with the input. So, the 4th-order and the inverted phase 2nd-order produce 3 dB bumps at the crossover frequency.

Linkwitz-Riley Alignment

Two things characterize a Linkwitz-Riley alignment: 1) In-phase outputs (0° between outputs) at *all* frequencies (not just at the crossover frequency as popularly believed by some) and 2) Constant voltage (the outputs sum to unity at all frequencies).

Linkwitz-Riley in-phase outputs solve one troublesome aspect of crossover design. The acoustic lobe resulting from both loudspeakers reproducing the same frequency (the crossover frequency) is always on-axis (not tilted up or down) and has no peaking. This is called zero lobing error. In order for this to be true, however, both drivers must be in correct time alignment, i.e., their acoustic centers must lie in the same plane (or electrically put into equivalent alignment by adding time delay to one loudspeaker). Failure to time align the loudspeakers defeats this zero lobing error aspect. (The lobe tilts toward the lagging loudspeaker.)

Examination of Figure 6 shows that the 2nd-order (inverted) and 4th-order Butterworth examples satisfy condition 1), but fail condition 2) since they exhibit a 3 dB peak. So, if a way can be found to make the amplitudes at the crossover point -6 dB instead of -3 dB, then the vector lengths would equal 0.5 (-6 dB) instead of .707 (-3 dB) and sum to unity—and we would have a Linkwitz-Riley crossover.

Russ Riley suggested cascading (putting in series) two Butterworth filters to create the desired -6 dB crossover points (since each contributes -3 dB). Voila! Linkwitz-Riley alignments were born.

Taken to its most general extremes, cascading any order Butterworth filter produces 2x that order Linkwitz-Riley. Hence, cascading (2) 1st-order circuits produces a 2nd-order Linkwitz-Riley (LR-2); cascading (2) 2nd-order Butterworth filters creates a LR-4 design; cascading (2) 3rd-order Butterworth filters gives a LR-6, and so on. (Starting with LR-2, every other solution requires inverting one output. That is, LR-2 and LR-6 need inverting, while LR-4 and LR-8 don't.)




Figure 7. Linkwitz-Riley vector diagrams for 2nd to 8th-order cases.

LR-2, A Transient Perfect 2nd-Order Crossover

As an example of this process, let's examine a LR-2 design. Referring to Figure 1, all that is required is to add a buffer amplifier (to avoid interaction between cascaded filter components) to each of these two outputs and then add another resistor/capacitor network identical to the first. We now have a 2nd-order Linkwitz-Riley crossover.

The new vector diagram looks like Figure 7a. Each vector is .5 long (from the fact that each 1st-order reduces by 0.707, and $.707 \times .707 = .5$) with phase angles of $\pm 90^\circ$. Since the phase difference equals 180° , we invert one before adding and wind up with a unity vector 90° out of phase with the original.

Figure 8 shows the amplitude response. The crossover point is located at -6 dB and the slopes are 12 dB/octave (40 dB/decade). The summed response is perfectly flat. Figure 9 shows the phase response. At the crossover frequency we see the HF output (upper trace) has $+90^\circ$ phase shift, while the LF output (lower trace) has -90° phase shift, for a total phase difference of 180° . So, we invert one before summing and the result is identical to the LF output.

These results differ from the 1st-order case in that the summed results do not yield unity (a piece of wire), but instead create an all-pass network. (An all-pass network is characterized by having a flat amplitude response combined with a smoothly changing phase response.) This illustrates

Garde's [4] famous work.

Cascading two linear phase circuits results in linear phase, as shown by the constant group delay plots (all three identical) of Figure 10. And constant group delay gives the transient perfect LF step response shown in Figure 11.

LR-4 and LR-8 Alignments

Looking back to Figure 7b., we see the vector diagrams for 4th and 8th-order Linkwitz-Riley designs. The LR-4 design shows the resultant vector is unity but 180° out of phase with the input at the crossover frequency.

Cascading (2) 4th-order Butterworth filters results in an 8th-order Linkwitz-Riley design. Figure 7c. shows the vector diagram for the LR-8 case. Here, we see the phase shift for each output undergoes 360° rotation returning to where it began. The resultant vector is back in phase with the original input signal. So, not only, are the outputs in phase with each other (for all frequencies), they are also in phase with the input (at the crossover frequency).

8th-Order Comparison

A LR-8 design exhibits slopes of 48 dB/octave, or 160 dB/decade. Figure 12 shows this performance characteristic compared with the LR-4, 4th-order case for reference. As expected, the LR-4 is 80 dB down one decade away from the




Figure 8. LR-2 amplitude response




Figure 9. LR-2 phase response




Figure 10. LR-2 group delay

corner frequency, while the LR-8 is twice that, or 160 dB down. Of interest here, are the potential benefits of narrowing the crossover region by using a LR-8 alignment.

Figure 13 magnifies the responses shown in Figure 12 to reveal a clearer picture of the narrower crossover region, as well as showing the flat summed responses. (The slight difference in summed amplitudes at the crossover frequency is due to a slight gain difference between the two circuits.) The critical crossover region for the LR-8 case is one-half of what it is for the LR-4 case. The exact definition of where the crossover region begins and ends is ambiguous, but, by whatever definition, the region has been halved.

As an example of this, a very conservative definition might be where the responses are 1 dB down from their respective passbands. We would then refer to the crossover region as extending from the -1 dB point on the low-pass response to the -1 dB point on the high-pass response. For LR-8, these points are 769 Hz and 1301 Hz respectively, yielding a crossover region only $\frac{3}{4}$ -octave wide. As a comparative reference, the LR-4 case yields -1 dB points at 591 Hz and 1691 Hz, for a 1.5-octave wide region.

For the LR-8 case, it is interesting to note that the -1 dB point on the low-pass curve corresponds almost exactly to the




Figure 12. LR-4 and LR-8 slopes

-20 dB point on the high-pass curve (the exact points occur at 760 Hz and 1316 Hz). So if you want to define the region as where the response is down 20 dB, you get the same answer. The entire region for the LR-8 case is $\frac{3}{4}$ -octave wide, or it is one-half this number for each driver. That is, the loudspeaker driver (referred to as 'driver' from now on) has to be well behaved for only about 0.4 octave beyond the crossover point. This compares with the 4th-order case where the same driver must behave for 0.8 octave.

The above is quite conservative. If other reference points are used, say, the -3 dB points (895 Hz & 1117 Hz), then the LR-8 crossover region is just $\frac{1}{3}$ -octave wide, and drivers only have to stay linear for $\frac{1}{6}$ -octave. ($\frac{1}{6}$ -octave away from the crossover frequency the drive signal is attenuated by 12 dB, so the output driver is operating at about $\frac{1}{16}$ power.)

The extremely steep slopes offer greater driver protection and linear operation. Beyond the driver's linear limits all frequencies attenuate so quickly that most nonlinearities and interaction ceases being significant. Because of this, the driver need not be as well behaved outside the crossover frequency. It is not required to reproduce frequencies it was not designed for. For similar reasons, power handling capability can be improved for HF drivers as well. And this




Figure 11. LR-2 transient response




Figure 13. Figure 12 magnified

narrower crossover region lessens the need for precise driver time alignment since the affected spectrum is so small.

The caveat, though, is an increased difficulty in designing good systems with sharp slopes. The loudspeakers involved have differing transient responses, polar patterns and power responses. This means the system designer must know the driver characteristics thoroughly. Ironically, sometimes loudspeaker overlap helps the system blend better even when on-axis amplitude response is flat.

LR-8 Phase Response

Figure 14 shows the respective phase response for LR-4 (upper trace) and LR-8 (lower trace) designs. As predicted by the vector diagram in Figure 7b, the LR-4 case has 180° ($4 \times 45^\circ$) of phase shift at the crossover frequency. Thus, the output signal is out-of-phase with the input signal at the crossover frequency for the LR-4 case. Both outputs are in-phase with each other, but out-of-phase with the input.

The LR-8 design eliminates this out-of-phase condition by bringing the outputs back in sync with the input signal at the crossover frequency. The lower trace shows the 360° phase shift for the LR-8 alignment.

LR-8 Transient Response

Butterworth functions do not have linear phase shift and consequently do not exhibit constant group delay. (First-order networks are not classified as Butterworth.) Since Linkwitz-Riley designs (higher than LR-2) are cascaded Butterworth, they also do not have constant group delay.

Group delay is just a measure of the non-linearity of phase shift. A direct function of non-linear phase behavior is overshoot and damping time for a step response. The transient behavior of all Linkwitz-Riley designs (greater than 2nd-order) is classic Butterworth in nature. That is, the filters exhibit slight overshoot when responding to a step response, and take time to damp down.

Figure 15 compares LR-8 and LR-4 designs and shows the greater overshoot and damping time for the 8th-order case. The overshoot is 15% for the LR-4 case and twice that, or about 30%, for the LR-8 case. As expected, the LR-8 design takes about twice as long to damp down. The initial rise-time differences are due to the group delay value differences.




Figure 14. LR-4 and LR-8 phase response

Is It Audible?

The conservative answer says it is not audible to the overwhelming majority of audio professionals. Under laboratory conditions, some people hear a difference on non-musical tones (clicks and square waves).

The practical answer says it is not audible to anyone for real sound systems reproducing real audio signals.

Linkwitz-Riley Power Response

Linkwitz-Riley alignments produce constant voltage response (voltage vectors sum to unity) at the crossover frequency, but they may produce constant power. At the crossover frequency, each voltage output is $\frac{1}{2}$ of normal. This produces $\frac{1}{2}$ the normal current into the loudspeakers. Since power is the product of voltage times current, the power is $\frac{1}{4}$ of normal. Considering a simple two-way system, the combined total power at the crossover frequency will be $\frac{1}{2}$ of normal ($\frac{1}{4}$ from each driver), producing a dip of 3 dB at the crossover frequency in the overall power response, *provided there is no additional phase shift contributed by the drivers themselves — such is never the case.*

The power response of loudspeakers with noncoincident drivers is a complex problem. See the Vanderkooy and Lipshitz⁵ study for complete details.

These responses are just for the electronic filters. For real world results including driver response, see Siegfried Linkwitz's website at www.linkwitzlab.com/frontiers.htm#F.




Figure 15. LR-4 and LR-8 transient response

References

1. S. Linkwitz, 'Active Crossover Networks for Non-Coincident Drivers,' *J. Audio Eng. Soc.*, vol. 24, pp. 2-8 (Jan./Feb. 1976).
2. D. Bohn, 'A Fourth-Order State-Variable Filter for Linkwitz-Riley Active Crossover Designs,' presented at the 74th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 31, p. 960 (Dec. 1983), preprint 2011.
3. D. Bohn, 'An 8th-Order State-Variable Filter for Linkwitz-Riley Active Crossover Designs,' presented at the 85th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 36, p. 1024 (Dec. 1988), preprint 2697.
4. P. Garde, 'All-Pass Crossover Systems,' *J. Audio Eng. Soc.*, vol. 28, pp. 575-584 (Sep. 1980).
5. J. Vanderkooy & S.P. Lipshitz, "Power Response of Loudspeakers with Noncoincident Drivers — The Influence of Crossover Design," *J. Audio Eng. Soc.*, Vol. 34, No. 4, pp. 236-244 (Apr. 1986).

RAP – Remote AC Power: An Idea Long Overdue

- **Hum Suppression**
- **Safety Agencies**
- **Performance Comparison**

**Dennis Bohn
Rane Corporation**

**RaneNote 121
© 1989 Rane Corporation**

INTRODUCTION

RAP is an acronym standing for **Remote AC Power**. RAP means remote power supplies. It means two pieces, where before you had one. RAP simply takes the transformer from inside the box and puts it outside the box. Call it a remote power supply. Call it an AC adapter. Call it a desktop power supply. Call it an external power supply. Call it a three-letter word: Call it RAP.

The idea for RAP grew out of the recognition that power supplies for truly professional audio signal processing equipment should be remote. They should not be in the same box with audio. Fifty years ago we got off on the wrong foot and have been stumbling about ever since. Everyday we pay the price for beginning wrong and not fixing it. The aim of professional audio products should be maximum performance, not maximum convenience. Maximum performance should be what separates professional audio products from consumer audio products. Yet the very best consumer products offer remote power supplies for that last percent of perfection, while most mainstream professional audio products do not. Only the top-of-the-line mixing consoles recognize power supplies must be remote to get maximum performance.

To make remote power supplies more palatable to the end-user, there should be an international standard specifying the voltages, connectors and wiring. All signal processing equipment would then have the same power jack and run off the same voltages. And instead of several remote power supplies, one central unit power supply would run everything through flexible cabling. Eventually, power supplies would become accessories, just like cables are today. In 1988, Rane took the first steps to make this a reality.

Due to Rane's efforts, members of the pro audio industry has been meeting and discussing such a standard since April, 1988. Later that year, the Audio Engineering Society (AES) created a Working Group to draft a standard. This Note discusses progress of the Working Group and presents Rane's decision to proceed with remote power supplies for most of its product line.*

BACKGROUND

In 1986 Rane began incorporating provisions for optional remote powering of its products¹. Next, Rane called upon the professional audio industry to recognize the merits of remote power supplies, and to seek ways to standardize them^{2,3}. RAP represents what we learned these past years.

Before getting into the details of RAP, we should discuss the need for remote power. Two quite different factors: performance and safety, dominate the case for remote power supplies.

PERFORMANCE

That audio products improve without hum sources in the box is self-evident; yet, a little review of the basics cannot hurt.

Hum Basics

The main ingredient of internal power supplies is the transformer used to step the AC line voltage down to usable levels for audio. This transformer operates by converting electrical energy into a magnetic field used to regenerate electrical energy at lower levels. The magnetic field is a strong source of either 50 Hz or 60 Hz hum. (50 Hz hum fields are much worse, which explains why the same audio products are quieter in 60 Hz countries.) Rich in odd-order harmonics, this field easily induces hum into sensitive audio lines (discrete or printed circuit) unless extreme (read, expensive) measures are taken to reduce this phenomenon.

Hum Suppression

Common techniques include using special copper bands around the transformer, installing special mu-metal shielding cans over the transformer, or going to expensive toroidal designs to reduce the spread of the magnetic field. All adding significantly to product costs.

If not contained, these hum field components leak into otherwise pristine audio to reappear as audible hum within the noise floor. Since these components are periodic, the ear is excellent at picking out and focusing on their repetitive nature within the noise. All high-gain microphone preamp stages, all summing nodes, and all high impedance points along the audio path are vulnerable to magnetic field induced hum.

Hidden Hum

Experienced audio designers learn many techniques to reduce these effects and get very impressive performance specifications. And experienced audio marketing people learn many equally valuable techniques to *hide* whatever small residual effects remain after being released to production.

Such as specifying the noise of graphic equalizers with all the sliders set flat. For it is when the sliders are in use that the equalizer is vulnerable to picking up radiated hum components from the transformer, not when set flat. (Which is why Rane also specifies its equalizers with the sliders positioned at their extremes.)

Tricks like using A-weighting to roll-off the hum components before measuring noise. Things like specifying mic

stages with input referred noise (Equivalent Input Noise), instead of total output noise. Things like specifying compressors, limiters, and gates with all dynamics turned off.

These things account for why one product *sounds* quieter than another, yet they have identical noise specs. The difference is one has hum components mixed with the noise and the other does not. The ear syncs into the repetitive nature of the hum and finds it much more objectionable than random noise of equal magnitude.

Ho Hum

You must understand that this discussion concerns very *small* amounts of hum. What my old boss used to describe as “picking fly droppings out of pepper.”

Making good better, and making excellent superlative — that is what remote power supplies are all about. Improving already excellent products is the goal, not fixing unacceptable designs. Nothing short of absolute perfection is acceptable. Power supply transformers are a major detraction from that goal. So, you get rid of them. You move them outside the box.

SAFETY

“Protecting” public safety has become very big business, with hundreds of millions of dollars at stake. (In 1994, Underwriters Laboratories grossed \$281 million in revenue and posted \$17.1 million in profits⁵.) Safety agencies have a job to do, but they also have a job to protect. As a result, we see more mandatory safety agency compliance, not less. In the US, the number of audio contracting jobs *requiring* UL listing grows at an alarming rate. Outside the US, it’s become a weapon.

Safety compliance, in its essence, is not really very difficult. Any safety agency’s main concern in electrical products is shock hazard. Shock hazard is confined to safety concerns of the line (mains) connected AC primary circuits. Remove the AC primary circuits from the box and you remove the shock hazard. Remove the shock hazard and you remove the need for compliance. That is the essence of RAP.

UL, NRTL and JAIL

In the United States, failure to have products listed by a “Nationally Recognized Testing Laboratory” (NRTL) is a crime in some jurisdictions — a **crime**, not a civil matter, but a **criminal** matter. On the federal level it is OSHA (Occupational Safety and Health Administration) that most directly affects audio makers. Congress gave OSHA exclusive jurisdiction to ensure the health and safety of workers *wherever* they may work — private industry, private schools, private churches, private recording studios, etc. In 1981, the agency adopted regulations requiring all equipment that uses electricity to be listed. OSHA Regulations, Section 1910.399, state “Electronic equipment is acceptable ... if it is ... listed ... by a Nationally Recognized Testing Laboratory ...” Use of equipment in the work place that is not “acceptable” is a violation of federal law and subjects the user to a broad range of penalties⁴. Very powerful lobbying.

*See Note on last page.

Canadian Standards Association (CSA)

In Canada, nearly all Provinces require equipment to be CSA certified. The Saskatchewan restrictions are typical. They totally ban the sale or use of equipment not CSA certified. Penalties include fines, recalls, or imprisonment. And in Europe, things are really beginning to boil.

Harmonization's Ugly Side

Now that the European Communities (EC, i.e., Common Market) harmonization is in effect, many people forget that the whole reason for harmonization is to create one large European market. A market with few *internal* walls, but with very strong *external* walls. Where possible, the aim of harmonization is to make it easier for EC members, and harder for outsiders, namely the United States and Japan, to sell products *within* the EC. The lesson here, is don't confuse a "common market" with a free market.

Low Voltage Directive

Sometimes making it easier for one must make it easier for the other. An example is the 1973 (this is not a misprint; it's an indicator of just how long harmonization efforts have been going on) "Low Voltage Directive" that unified the product safety standards for electronic products throughout the 12 member EC (Great Britain, France, West Germany, Belgium, Netherlands, Luxembourg, Spain, Portugal, Italy, Greece, Denmark and Ireland). The Low Voltage Directive, designated 73/23/EEC, acknowledged that the product safety requirements of some governments of the EC had become "repressive" and had to be eliminated⁵.

Among the repressive requirements was the need for equipment to receive the national mark of each country to be sold. Also, the standards would no longer be drafted by the individual countries, but would be chosen by the Common Market itself. This resulted in one common safety document (EN60065 for audio products), and all members of the EC must accept each other's safety marks. So, for example, VDE works in Great Britain, and DEMKO works in Luxembourg, etc.

To U.S. manufacturers, this means that compliance in one country, means compliance in all countries of the EC, plus Switzerland and Norway. (Expect other countries to agree to the common product safety code as harmonization spreads.) That is one side of the coin — the good side. The other side of the coin is that compliance with the requirements established by IEC 65 (now referred to as EN60065) is mandatory for all products entering into EU after 1/1/97.

Punitive Safety, or No Pay, No Play

Mandatory safety agency compliance too often becomes a de facto trade barrier preventing otherwise high quality products from being used or available to customers or contractors. For example, in the United States if UL is required, only Japanese and a few US and foreign products qualify. They are the only ones who can afford the UL process. Small companies simply do not have the resources to get compliance with multiple agencies (UL, CSA, VDE, etc.) for multiple products. For example, in Rane's case, for, say 30 products, this would entail over 90 separate filings and compliance records. The cost would amount to more than \$350,000 and the total time required would approach 10 years!

And that is not the bad news. The really bad news is that, too often, these safety agencies have conflicting requirements. What satisfies UL will not satisfy CSA, and so on, so you must build *separate* products for each safety agency. The limited resources of most audio companies, and the total available sales volume in most foreign countries simply does not allow building separate products.

Only the largest audio companies (read, mostly Japanese) can afford to do business in this manner. Many smaller American audio companies will be forced to withdraw from international markets, and try to survive on what non-UL business exists in the US.

Bleak? You bet it is.

Compliance Exemption

Enter remote power supplies, to save the day.

Remote power supplies may represent the survival of the small entrepreneurial audio company in America and elsewhere. *By using remote power supplies, and sizing them correctly, all products powered by them become as exempt as possible from safety agency compliance.* Only the remote power supply itself must comply with each safety agency. With proper design, the number of products requiring compliance reduces to 2 or 3. This is manageable by any audio company.

The hedging ("...as exempt as possible...") recognizes the reality that any government agency can require safety agency markings on anything they so choose. Even if something runs off a 9 V battery, they can argue it could be a fire hazard and demand compliance. There are never any guarantees when dealing with governments. But the use of approved remote power supplies makes compliance simple, fast and economical.

REMOTE POWER SUPPLIES

Beginning in 1989, improving performance and providing instant safety agency compliance caused Rane to adopt remote power supplies for all new products, and to convert many existing products. Instead of waiting for the AES to agree on a standard, Rane chose to move ahead and provide their customers with all the benefits of remote power.

AC Power

Today Rane uses RAP for remote power. This differs from what was described in the now discontinued *Rane Note 118*, "Remote Power System," and sold as the RS 10, Remote Power Supply. RAP is remote **AC** power; the RS 10 is remote **DC** power — big difference.

Rane's early efforts in remote power supplies centered on DC systems. Most of Note 118 is still valid, and RS 10 based systems work perfectly; but, we have learned DC's limitations. We have learned there are easier ways. Ways that allow more freedom in mixing manufacturers and different wiring techniques for complex professional audio systems. Our work in chairing the AES Working Group and our own experience taught us the superiority of distributed AC systems over DC systems regarding ground loops.

Adopting AC is similar to simply removing the existing power transformers from within the units and collecting them

together remotely in a common box. This box may contain one transformer, like our RS 1, or it may contain 5 transformers with 10 separate secondaries, like our RAP 10. If each power transformer has isolated secondary windings, then each unit powered has its own isolated ground system.

RAP supplies also allow us to use simple AC doubling techniques for generating higher voltage levels when required, without having to resort to expensive DC-to-DC converters. And adopting AC makes for simpler, more cost-effective external supplies.

Voltage Level

The AES Working Group's choice of voltage level occurred after a study of international safety regulations for professional audio products. A cross section of countries investigated found the maximum voltage permitted before requiring safety compliance. Table 1 shows the results of this study.

At least half the countries studied feel that a risk of shock occurs with exposure to greater than 42.4 volts peak (30 volts RMS). The other half defines shock hazards at lower voltages. And there are obvious contradictions. For instance, Table 1 shows that while Italy and the U.K. are members of the EC and legally bound to the tenets of Low Voltage Directive 73/

Table 1. Maximum voltage levels allowing greatest exemption from safety codes.

Country	Vpeak	VACrms	Comments
U.S.A. (UL)	42.4	30	Class 2
Canada (CSA)	42.4	30	Class 2
Japan (JET)	42.4	30	
Swiss (SVE)	42.4	30	
W. Germany (VDE)	42.4	30	VDE SELV Definition EC Member
Australia (SECV)	70.7	50	
Italy (IMQ)	100	34	EC Member
Denmark (DEMKO)	50	35.4	EC Member
Finland (FEMKO)	42.2	29.8	
Norway (NEMKO)	42	29.7	
Sweden (SEMKO)	35	25	
U.K. (BSI)	34	24	BS415 (IEC 65) EC Member
Dir 73/23/EEC	75	50	Low Voltage Directive
IEC 65*	34	24	CENELEC** Safety Standard Covering Audio

* IEC 65, entitled "Safety Requirements for Mains Operated Electronic and Related Apparatus for Household and Similar General Use, is the approved Harmonization Document (HD 195 S3).

** CENELEC is the European Committee for Electrotechnical Standardization.

(1989 Data)

23/EEC, they differ widely on the required voltage for compliance. And look at the differences between the Directive and IEC 65. The Working Group decided upon two remote power supply voltages. One, aimed at low voltage products, would be 9 VAC RMS. The other would be the best compromise for worldwide acceptance. Therefore, since Europe was more restrictive than the United States, it would decide the voltage level. Further investigation revealed that IEC 65 was the controlling harmonization document, and that any disagreements between IEC 65 and any EC member country must be resolved. All of which boiled down to 24 VAC RMS being the maximum voltage allowable in any product for the greatest exemption from worldwide safety codes.

Based on this input, the final voltage level picked by Rane was 18 VAC RMS. This allows worst case high AC line conditions, combined with light load situations, never to let the voltage level exceed 24 VAC. This voltage level is low by professional audio standards and required a complementary decision that all Rane RAP products incorporate regulated voltage doublers to create the necessary headroom demanded by our products.

All Rane RAP products qualify worldwide as Safety Extra Low Voltage (SELV) units. They pose no shock hazard at all.

Connector

Once you accept the need for RAP, the hardest and most controversial aspect becomes connector choice. The AES Working Group favors, and Rane uses, telephone style modular connectors.

Three problems head the list of what is wrong with the connectors in use today (mostly barrel and mini-phone jacks):

1. They do not lock.
2. They are prone to shorting (either during the act of plugging them in, or just by having them dangle and touch something).
3. There is no interchangeability between manufacturers adapters

After a painstaking search, Rane concluded the modular series developed by AT&T was the best choice. These connectors offer the following advantages:

1. Keyed and locking design.
2. Pins fully protected against accidental short circuits.
3. Small size.
4. Inherently strain relieved.
5. Very economical.
6. Multi-sourced from many manufacturers.
7. Extremely reliable.
8. Field repairable.

All currently used connectors fail the first two critical criteria. Alternative proposals tended to fail 1, 3 and 5. Only the modular connector satisfied *all* these requirements. On reliability, AMP Corporation applications engineers agree

that the telephone modular connector is one of the most tested and approved connectors ever developed. When AT&T develops and approves something for consumer use, it is reliable.

Addressing the obvious question of potential telephone/audio equipment wiring mix-ups was a major concern of RAP. In the end, simply using red mod jacks, with a special RAP logo (Fig. 1) silkscreened around them, reduced all telephone confusion. Further precautions dictated a wiring convention (Fig. 2) that prevents potential damage resulting from inadvertent telephone equipment hook-up.

Selecting the 6-pin connector for the 18 VAC designs (and the 4-pin handset connector for 9 VAC use), further improved the power handling, performance and reliability. Using the 6-pin version allowed parallel pins and wiring for each side of the transformer and center-tap lead. Figure 2 shows this arrangement.

Addressing the Negatives

Remote power supplies have been around for a long time; long enough to have earned a bad reputation. Most complaints boil down to these:

The remote power supply unit must plug directly into the wall or an outlet strip. Doing so, covers up other AC outlets, preventing their use.

The connecting cable is too short to reach the unit from where you want to put the outlet strip.

There is no way to mount the power supply securely so it will not become disconnected or damaged.

Rane's RAP supplies address each of these complaints in a simple and straightforward way: Rane's units are a desktop design, with a 6 foot AC mains line cord out one side, and a 6 foot low voltage cable out the other. And each power supply comes complete with integral mounting ears.

Now you are free to mount a RAP supply almost anywhere you want to, without covering any AC outlets. Typically, RAP supplies securely mount to the side or bottom of the equipment rack. This makes for a neat, out of the way location with all wiring bundled and tied. Larger installations get even easier by using the RAP 10 Remote Power Supply to eliminate several individual RAP units altogether.

Terry Pennington, Rane's MIS Director, refers to the RAP supply as "the fat spot in the line cord." Looked at from this view point, it's no big deal. You mount it once, and forget it.

Hum Avoidance

Mounting RAP supplies requires common sense. Since they are nothing but a hum producing transformer, mount them far away from sensitive high gain inputs. Locate them near AC outlet strips, or power amplifiers (another source of large hum fields), or together off in a corner. Avoid putting them near microphone inputs, mixers, and other small signal processing devices. Mount them as far away from audio lines as possible. No need to get paranoid about it, just use common sense.



(a) RAP-4 Logo




(b) RAP-6 Logo


Figure 1. RAP Logos

PERFORMANCE COMPARISON

Converting an existing product to a RAP product results in varying performance improvements. Some products benefit more than others. Performance improvement stems from the lack of hum fields within the unit. Removal of these hum fields improves the noise and distortion performance of RAP products. The improvement in distortion results from the reduction of the noise floor. Since distortion is measured as total harmonic distortion plus noise (THD+N), reducing noise also reduces distortion. Most other parameters of RAP



(a) RAP-4 9VAC wiring.



(b) RAP-6 20VAC W/CT wiring.

Figure 2. RAP Wiring Convention

products remain as before.

RAP graphic equalizers benefit most in their cut noise performance. Flat and boost figures remain about the same. It is always annoying to experience an increase in noise when using sliders to cut the signal. Intuition says if you are reducing levels through the equalizer, the noise should reduce, not increase. Well, most equalizers exhibit a degrading noise performance with increased cutting of their sliders. Partially, this is due to the antenna effect produced by having all sliders tied together along the front panel. This forms a loop-antenna of sorts, and moving the sliders from their center-detent positions allows them to pick-up magnetic hum produced by the transformer. Hum picked-up from boosting usually gets masked by the larger signal level. Hum picked-up when cutting stands out from the reduced signal level. RAP graphic equalizers do not suffer from the malady.

As an example of this phenomenon, let's examine the GE 30 before and after RAP conversion (see Table 2). For proper perspective, remember the GE 30 is a leader the world over among active graphic equalizers for noise. It is among the quietest active equalizers ever designed. Against this background, the performance advantages provided by RAP are startling. Performance of this magnitude was previously only possible from passive equalizers.

As expected, we see a couple of dB improvement in the flat settings. The 4 dB increase in signal-to-noise for full boost and maximum gain is good, but nothing compared to the full cut improvement. The RAP GE 30 is over 15 dB quieter when used in the Cut Mode. We think that's worth a little inconvenience in mounting an outboard power supply.

Another dramatic improvement happened with the RAPping of the HC 6 Headphone Console. This box demanded the extra current of a model RS 2 remote supply. The signal-to-noise ratio improved by 16 dB in Channel 1! The rest of the Channels improved by 6 dB or better. Simply because the transformer was located next to Channel 1. Now all Channels are equally quiet.

Some dramatic improvements do not show up on the data sheet. An example is the SM 26B. The Mono Outputs of the SM 26B have always been quiet (96 dB re +4 dBu), but the Mix Outputs have always had a tendency to pick up transformer and AC line hum. Over the years we have improved this by adding a metal shield can to the transformer and routing the signal path differently, but it has always been a fight. And when used in 50 Hz countries, it is always worse. Still very quiet, but not as quiet as we would have liked. Then we converted it to RAP. Now the Mix Outputs exhibit at least 95 dB signal-to-noise ratio for unity gain settings, and over 82 dB with all controls set for maximum gain. These figures are at least 6 dB better than without the outboard power supply.

Some products show no measurable improvement at all. This is the case with the DC 24. It was so extraordinarily quiet before, that no further improvement is practical.

The AC 22 and AC 23 Active Crossovers are similar, with one exception. The exception is the 230 VAC, 50 Hz units. In these units the overall noise is down about 10 dB for the Ch. 1 Low Output. Not surprisingly, this is the output closest to the power transformer. All other outputs show little improvement. Again, over the years we had to add transformer shielding

with other separate shields to hold performance. Now, all cans and shields are gone and the low noise performance is repeatable, predictable and the lowest we have ever measured.

RAP, RAP, RAPPING UPON MY DOOR

So there you have it: RAP, an idea long overdue. The only price is having to mount a small power supply using the supplied hardware. The benefits range from being essential to the recording studio producing digital recordings; to being absolute to the commercial sound installer, who is no longer prevented from using Rane because of National or local safety codes.

Here's to the "fat spot in the line cord" — a small inconvenience, a large benefit.

REFERENCES

1. D. Bohn, "Remote Power System," *Rane Note 118*, Rane Corp., (discontinued, 1987).
2. D. Bohn, "Remote Power Supply Standards: A Proposal," *S&VC*, vol. 5, pp. 70-78, (Nov. 15, 1987).
3. L. Winter, "Why Bother with Industry Standards?," *Sound & Communications*, vol. 34, p.6, (Feb., 1988).
4. G. Dash, "Product Safety in the United States and Canada," *Compliance Engineering*, vol. VI, pp. 219-230, (Jan., 1989).
5. G. Dash, "Product Safety in the European Communities," *Compliance Engineering*, vol. VI, pp. 232-238, (Jan., 1989).
6. Edward T. Pound, Gary Cohen, & Penny Loeb, "Tax Exempt," *U.S. News & World Report*, (Oct.2, 1995).

DISSOLUTION OF AES WORKING GROUP ON POWER SUPPLY INTERFACING*

The AES Working Group grew out of an ad hoc committee started by Bob French (Ashly) and Jim Furman (Furman) at the April 1988 NSCA Expo in Reno, with representatives from ART, Ashly, Furman, Rane and Symetrix attending. During the summer of 1988, research was carried out to determine the maximum voltage level allowable to power commercial audio products (country-by-country) before requiring compliance with safety agency codes.

AES Working Group Formed

The results of this study made it clear that all audio products would benefit from remote power supplies – both in performance and safety. In August 1988, the ad hoc committee agreed to form an AES Working Group on remote power supply standardization. The first meeting was held during the November 1988 AES Conference in Los Angeles. By that time, over 20 companies were on the mailing list, expressing interest in standardizing remote power supplies. Representatives attending the November 1988 meeting were from ART, Ashly, Crown, dbx, Furman, Korg, Lexicon, Rane, and Shure.

At this meeting, a formal proposal was handed out for study and discussion. Over the next several months, revisions were made to this report, loosely referred to as the RAP (Remote AC Power) proposal.

A revised version circulated in February 1989 to official representatives of 25 companies for comment, solicited only one response. A draft proposal was then drawn up and submitted to the AESSC (AES Standards Committee) in July 1989. Along with the proposal, went the results of balloting all interested parties. Of 23 companies responding, 16 companies (70%) favored the proposal, while seven companies (30%) rejected the proposal.

Draft Proposal Summary

Remote power voltage would be alternating current (AC). There would be two levels: level 1 would be 9 VAC single-ended, and level 2 would be 20 VAC with center-tap connections.

The connectors would be telephone style modular plugs. Level 1 would use the 4-pin handset modular plug (AMP no. 641334, or equal) with parallel pins and wire for each side of the 9 VAC supply. Level 2 would use the 6-pin RJ-12 modular plug (AMP no. 641337, or equal) with parallel pins and wire for all three connections of the 20 VAC with center-tap supply. An alternate proposal would use the offset version of the 6-pin plug to prevent telephone miswiring. All connector plugs and matching jacks would be color-coded red.

All plug and jack wiring was specified (as much as possible) to prevent any damage to or from any telephone equipment inadvertently miswired into these connectors.

A unique identifying logo was specified to be silk-screened around the red modular jack as a further preventive measure against any possible telephone confusion.

Disagreements With Draft Proposal

Of the seven companies disagreeing with the proposal, two were against the voltage and five were against the connector.

The two companies voting against the voltage wanted DC instead of AC. They believed a ground loop free system of distributed DC was possible and superior to the proposed distributed AC system.

The five companies voting against the modular connector proposal all voiced reservations that many companies voting for it also shared. The most obvious one was confusion with telephone equipment. They did not feel that the offset latch plug, red color, and unique logo were sufficient to prevent confusion.

The second concern was with the reliability of the plastic latching tab. Many felt it could be broken off too easily, and that the entire plug was too vulnerable to breakage.

Further concern was expressed by the lack of readily available 6-pin crimping tools and connectors, particularly if the offset latch version were adopted. It was pointed out this issue becomes even more acute outside the United States.

And last, reservations were stated regarding the current handling capability of modular plugs, especially over long distances. Even though the modular pin carries a 2 amps/pin UL rating, it restricts wire size to 26 AWG. And two 26 AWG conductors in parallel only allows 1.5 amps maximum.

Yet after 18 months of discussion, no one presented an alternative connector proposal better suited than the modular plug.

AES Working Group Dissolved

The AESSC rejected the proposal for *Journal* draft publication because four of the seven companies disagreeing were large companies. After a poorly attended meeting during the October 1989 AES Conference in New York, the AESSC Executive Committee dissolved the Working Group on the grounds of insufficient interest.

* Condensed from the original article appearing in the *J. Audio Eng. Soc.*, Vol. 39, No. 4, April 1991, pp. 275-276.

Limiters Unlimited

- **Auto-Slave**
- **Protecting Loudspeakers**
- **Protecting Systems**
- **Signal Overload Protection**
- **Limiting Volume**

Monty Ross
Rane Corporation

RaneNote 127
© 1991 Rane Corporation

INTRODUCTION

Limiters are in the protection business, limiting audio systems to safe levels. These limits protect loudspeakers, protect the audio signal from clipping, protect the neighbors, and protect ears.

A limiter continuously monitors the audio signal, looking for levels exceeding its adjustable threshold. A limiter normally operates at unity gain and has no effect on the signal. If excessive levels are detected, the Voltage Controlled Attenuator (VCA) automatically reduces the gain. If the level never exceeds the threshold, the signal remains unaffected.

AUTO-SLAVE, EXTENDING YOUR LIMITS

In the applications and examples that follow, you will read about multi-channel limiters with *auto-s/ave*. Auto-slaving provides the ability to tie together the circuitry that controls the VCA's of two or more limiters. Each limiter channel maintains its independent threshold adjustment. The amount of gain reduction is shared with all slaved channels when any one has exceeded its threshold. Figure 1 shows a simplified block diagram of a quad limiter with this feature.

PROTECTING LOUDSPEAKERS WITH LIMITERS

Loudspeakers manufacturers provide recommended power ratings for their drivers. These ratings are not absolute, and should be used only as a guide. Sound contractors' diverse experiences using loudspeakers sometimes lead them to formulate their own power ratings. Limiters allow for real world applications. With a limiter you can customize your speaker protection scheme for maximum reliability.




Figure 1. Quad Limiter with Auto-Slave

It is important to note that midrange and tweeter (high-frequency) drivers cannot handle as much power as woofers (low-frequency drivers). Therefore it is especially important that tweeters and midranges are protected by different limiter thresholds than the woofer.

Multi-amped systems often use limiters for speaker protection. Limiters are typically placed at the input of each power amplifier for each driver. In Figure 2 we show a tri-amped system protected by a limiter. Set each limiter's threshold to match the power handling of each driver. Whenever the high, mid or low limiters reach their set limit threshold, all three bands limit by the same amount. This is due to the three slaved channels in the limiter, assuring the system's spectral balance stays constant. (If only one band limited and the other two increased in volume, the result would be an abrupt change in the tonal balance of the system.) This clever trick provides the least audible and most effective form of multi-amped system power limiting.

TAKING YOUR SYSTEM TO THE LIMIT

A properly designed system has the ability to operate as loudly as needed for the application. When limiters are introduced into a system their thresholds are set to protect the systems' components.

Sometimes limiting occurs before a satisfactory volume level is reached. For example, a touring band could perform in a venue that is much bigger than its equipment can handle. Higher power speakers or more powerful amplifiers may be called for. This permits higher limit levels to be set, thus allowing a louder system.

ALL SYSTEMS UNDER CONTROL

Signal overload problems may occur when cascading multiple audio signal processors. Take for example a parametric equalizer with a single band boosted by 10 dB to flatten a room. If a loud note in this same frequency range is present, the overload light on the equalizer illuminates. By then it is too late, the signal is clipped. Additionally, overloading could be overlooked if the overload indicators are




Figure 2. Active Tri-Amped System with Protection




Figure 3. Signal Overload Protection

obscured by security covers or a closed rack. This problem is corrected by adding a 4-channel limiter with all four channels slaved together as shown in Figure 3. The first channel is the primary input limiter and up to 3 points are simultaneously monitored by the limiter's other 3 channels. The *threat* of overload automatically activates the limiter rather than lighting an LED, thus preventing overload altogether. Simply add a wye connector in the audio path anywhere you need protection and connect it to one of the 3 other inputs. When any single processor nears an unwanted level, the first channel limits the gain at the input of the whole system, preventing overload.

REIGN OVER YOUR VOLUME

A benefit of limiters is their automatic gain control ability. No one needs to be present to watch and change the levels of a system once you properly set the thresholds. A limiter prevents the need to gain ride a system. In this sense, the limiter acts as a compressor. Increase the volume so that the quiet parts are loud enough to be heard and the limiter automatically prevents the loud passages from clipping the amplifiers or overpowering the speakers.

Professional audio systems are often operated by untrained people. Take a music/paging system in a restaurant for instance. The owner of a restaurant wants specific levels of music in certain dining rooms so his customers or neighboring businesses will not be annoyed. Placing a limiter in the

audio chain will prevent everyday users from turning the volume up too loud. The limiter's circuit tends the ship.

Another application involves hearing. We sometimes take our hearing for granted. Hearing is non-linear and constantly changing. When the sounds our ears pick up get loud (sensitivity to loudness varies with different people), some automatic systems come into play. Some of these changes are aural (involuntary muscular responses in the ear) and some are neural (the way the brain responds to impulses from the ear). The mechanism between the ear and brain begins to shut down in an attempt to protect itself. This action manifests itself by dulling your hearing senses. The natural response to this is to boost the volume or EQ controls to compensate for the perceived drop in level. Extended exposure to high sound levels can result in permanent hearing damage. Limiters help control these levels. When using a limiter for this purpose we recommend security covers to make the settings tamper proof.

SUMMARY

Put limiters in the signal path—anywhere there is a place where signals can't be monitored by someone or where there might be too much signal. At the output of the signal source, at the input of an equalizer, at the input of a mixer, at the input of a power amplifier or at the input to an active cross-over. You can be very creative with limiters to achieve your particular goal. Especially if your limiter has an auto-slave option.

REFERENCES

1. M. Ross, "An Investigation into How Amplifier Clipping is Said to Burn-Out Loudspeakers, and How Limiters Can Save Them," presented at the 89th Convention of the Audio Engineering Society, *J. Audio Eng. Soc. (Abstracts)*, vol. 38, p.870 (Nov 1990), Preprint 2956.

Power Amplifier Clipping and its Effects on Loudspeaker Reliability

- **Harmonic Theory**
- **Power Ratings**
- **Blown Tweeters**
- **Clipping and its Indicators**
- **Limiter Protection**

**Monty Ross
Rane Corporation**

**RaneNote 128
© 1991 Rane Corporation**

INTRODUCTION

Power amplifier clipping is quite common. This note examines the clipping phenomenon which *allegedly* damages loudspeakers. We suggest that this form of distortion is not the cause. Rather, we show that amplitude compression of the audio spectrum is the culprit. Rane limiters provide a solution to amplitude compression, thus preventing loudspeaker failure.

WHY DO LOUDSPEAKERS NEED PROTECTING?

All loudspeaker drivers have power handling limits. Once exceeded, damage occurs. There are several ways a loudspeaker suffers *power* damage. A couple of these warrant explanation.

The first is over-excursion of the diaphragm. The diaphragm of a loudspeaker is the radiating surface that moves in response to an electrical signal. This surface may be conical, domed or flat in shape, and it creates sound by physically pushing and pulling the air in the room. The laws of physics say that in order to play louder or to reproduce lower frequencies, the diaphragm must move further toward its mechanical limits. If it is asked to move still farther, it experiences over-excursion. This most often occurs in woofers but can affect midranges or tweeters if low frequencies are not limited. If the loudspeaker cannot handle over-excursion, mechanical destruction of the driver is likely the result.

Another enemy of loudspeakers is heat, generated by power losses in the voice coils. No device is 100% efficient. For loudspeakers, 1 watt of input power does not produce 1 watt of acoustic energy. In fact most loudspeakers are typically well under 10% efficient [1] [3]. These losses convert to heat that builds up in the voice coils, causing mechanical deformation, like melting of the voice coil former. It causes weakening of the structure by charring the voice coil former, which later shakes apart. The heat causes the glues to bubble up, fill the air gap and glue the voice coil solidly in the gap. Often the voice coil wire melts like a fuse link, resulting in an open driver. Obviously we wish to prevent this.

Music power-handling capability for multi-way loudspeakers always presents a problem to the loudspeaker user and designer. Users who must replace blown tweeters often feel they didn't do anything wrong, because their amplifier only put out 50 watts and their speaker had a 200 watt rating.

Yet, the tweeter blew up. This recurring problem motivated engineers to find out why this happens. Many opinions developed. Some of these have been scientifically verified—others remain theory.

CONFLICTING “FACTS”

Studies show the typical spectral energy for different types of music have high frequency energy considerably lower in level than low frequency energy [2]. This knowledge has further complicated the studies of how tweeters get destroyed. It seems that woofers should blow rather than tweeters if the high frequencies are lower in amplitude.

Loudspeaker manufacturers use this knowledge about the energy distribution of music when they design their products. This knowledge allows them to make better sounding tweeters because they can use lighter moving structures. Smaller wire in the voice coils can be used because there is significantly less power in the high frequency ranges. Since smaller wire is lighter, it takes less energy to move. For a speaker system rated to handle a given number of watts, the tweeter by itself can probably handle less than one-tenth that amount.

From all this came a theory that spread quickly through the industry. Since there is more musical energy at low frequencies than high frequencies, there is not enough high frequency power to blow out tweeters. Therefore, high frequencies loud enough to burn out tweeters must come from somewhere else. Where do they come from?

Well, it was reasoned, if there is enough low frequency energy to clip the amplifier, then it perhaps would produce enough high frequency distortion products (as a result of clipping) to blow up the tweeter.

This theory convinced many in the early 70's and slowly evolved into “fact”. While doing research into the reliability and protection of power amplifiers, I had to study how the

Table 1. Harmonic Amplitudes of a 100 Hz Square Wave, 0 dB = 100 Watts (Instantaneous Power)

Harmonic	Amplitude	in dB	Watts	Frequency
1	1	0	100	100 Hz
2	0	-∞	0	200 Hz
3	1/3	-9.54	11.12	300 Hz
4	0	-∞	0	400 Hz
5	1/5	-13.98	4	500 Hz
6	0	-∞	0	600 Hz
7	1/7	-16.9	2.04	700 Hz
8	0	-∞	0	800 Hz
9	1/9	-19.1	1.23	900 Hz
10	0	-∞	0	1 kHz
11	1/11	-20.8	0.83	1.1 kHz
12	0	-∞	0	1.2 kHz
13	1/13	-22.3	0.589	1.3 kHz
:	:	:	:	:




Figure 1. Harmonic Components of a 100 Hz Squarewave vs. a 100 Hz Sinewave.

typical consumer used amplifiers and speakers. I found that clipping is a common occurrence and is not as audible as most people think. I also found that the operation of many clipping indicators is very slow and does not always show actual clipping. (Many manufacturers slow them down, using their own rule of thumb for how much clipping can occur until it lights the indicator.)

Newer and better sounding amplifiers, including amplifiers with soft clipping circuits, still blew tweeters. But amplifiers with higher power were having fewer incidences of blown tweeters. This appeared to reinforce the theory that clipping caused tweeter blowouts. One thing was clear, when clipping occurred, tweeters blew.

If you're getting the idea I don't believe in the clipping/harmonic theory, you're right. So let's investigate the phenomena further.

WHEN SINE WAVES CLIP

When sine waves clip severely they resemble square waves in shape, introducing massive distortion. In the extreme case, a perfect square wave has the highest level of harmonic components (See Figure 1). A less clipped sine wave has components at the same frequencies but at lower levels.

Let's look at the square wave example shown in Table 1 (at left). Fourier analysis shows the harmonic structure.

As you can see, the total amount of instantaneous power left to make it through an ideal 1kHz crossover (and on to the tweeter) is less than two watts ($0.83 + 0.589 = 1.419W$). Hardly a problem. And remember, this simulates severe overdrive of a 100 watt amplifier with a sine wave to make an ideal square wave. Driving it harder will not increase the harmonics.

This analysis shows if a small tweeter that only handles 5 or 10 watts is used in a 100 watt speaker system it would not blow out, even under square wave conditions. Yet it does.

It takes a lot more than this to cause major failure. So what's happening?

Compression is what's happening [3].

Today's newer higher quality amplifiers have greater dynamic range and sound better when clipped with musical transients than older amplifier designs. So it is more likely for a user to overdrive and clip newer amplifiers on low frequency dynamic peaks because of lower audible distortion. This results in compression of the dynamics of the music. The




Figure 2. Low Level, High Frequency Sinewave Mixed with a High Level, Low Frequency Sinewave Burst.




Figure 3. 100 W Amp with 3 dB Overdrive




Figure 4. 100 W Amp with 10 dB Overdrive

high frequencies get louder but the low frequencies can't. This may be heard as an increase in brightness of the sound. Some may simply interpret it as louder with no change in tonal balance.

For example, in a 100 watt amplifier, as you turn up the level, the low frequency components will limit (clip) at 100 watts. Meanwhile the high frequency components continue to increase until they (the high frequencies) approach the 100 watt clipping point.

The graphs in Figures 2, 3 & 4 are scaled in volts. With an 8 ohm load the 100 watt level corresponds to 40 volts peak. Below clipping, the low frequencies reach 100 watts (40 volts peak) but the high frequencies are only 5 or 10 watts (9 to 13 volts peak).

Let's assume a musical signal with low and high frequency components driving a 100 watt (8 ohms) amplifier. We use a low level/high frequency sinewave mixed with a high level/low frequency sinewave burst. (See Figure 2). The high frequencies reproduced by the tweeter are at least 10 dB lower in level than the low frequencies. Now as we turn up the amplifier to clip the signal (3 dB overdrive—See Figure 3). Notice that only the low frequency burst portion of the waveform clips but the high frequency portion increases in level. The clipping, of course, produces harmonics but not nearly as much as the square waves discussed earlier. The amplitude of the high frequencies went up by 3 dB in relation to the low frequency fundamental. (3 dB compression).

If you overdrive the amplifier by 10 dB, the high frequency amplitude goes up by 10 dB. This goes on dB for dB as you turn up the volume, until the high frequency reaches the 100 watt level. Meanwhile the peak level of the low frequency portion can not increase above 100 watts (See Figure 4). This now represents nearly 100% compression (no difference between HF amplitude and LF amplitude).

Now it is easy to see how the high frequency portion exceeds the 5 or 10 watts tweeter rating. Sure, clipping is producing extra harmonics but *they never approach the levels of the amplified high frequency source signals.*

It may be argued that the signal's distortion would be intolerable. Don't fool yourself. It really surprises people how much clipping they tolerate before they cannot listen anymore. Just disconnect the clipping indicator on a power amplifier and see how loud someone drives it. Watch the amplifier output with an oscilloscope. There will be a surprising level of clipping. 10 dB clipped off the top of low frequency transients is not an uncommon occurrence when the purpose is to impress your neighbors.

WHAT CAN WE DO ABOUT IT?

If we can prevent an amplifier from clipping, we could better utilize our loudspeakers. Limiters play an important role in preventing clipping and the resulting amplitude compression. The Rane MA 6S power amplifier, DC 24 dynamic controller, CP 64 and CP 52 commercial processors, VP 12 voice processor, and the RPM series processors are all products that limit. These limiters prevent the compression mentioned earlier because when *any frequency reaches threshold all frequencies are turned down by the same amount.*

The MA 6S six channel power amplifier is specifically designed not to clip by the use of its internal voltage controlled attenuators on each channel.

The DC 24 limiters have user adjustable threshold controls. This allows you to customize your total system for maximum reliability with no compromise on sound quality.

REFERENCES

1. K. Blair Benson, *Audio Engineering Handbook*, (McGraw-Hill, New York, 1988)
2. R. A. Greiner and Jeff Eggars, "The Spectral Amplitude Distribution of Selected Compact Discs," *Journal of the Audio Engineering Society*, vol. 37, pp.346-275 (April, 1989)
3. Carlo Zuccatti, "Thermal Parameters and Power Ratings of Loudspeakers" *Journal of the Audio Engineering Society*, vol. 38, pp.34-39 (Jan-Feb, 1990)

Squeeze Me, Stretch Me: The DC 24 Users Guide

- **DYNAMICS 101**
- **THRESHOLD & RATIO CONTROLS**
- **COMPRESSORS & LIMITERS**
- **GATES & EXPANDERS**
- **SPLIT BAND PROCESSING**
- **GUITAR, BASS & RECORDING**

Jeff Davies
Dennis Bohn
Rane Corporation

RaneNote 130
© 1993 Rane Corporation

INTRODUCTION

Compressors, expanders, and their cohorts – limiters and gates, are all in the business of automatically controlling the volume, or dynamics of sound. Lumped together they can be called dynamic controllers, which would also have to include your hand on the fader and the fat man dancing in front of the midrange cabinet.

Used wisely, often in conjunction with each other or with equalization or filtering, dynamic controllers can improve the intelligibility of voice and the subjective effect of music. But in the wrong hands they can sound terrible, and compressors are the worst offenders.

Our goal is to de-mystify dynamic controllers as best we can within the limitations of printed media. By understanding a given tool's strengths and weaknesses, you can put it to its best use.

Roger Nichols - "I have used the DC 24 on every album project I have done since I've had it". He has had a DC 24 since 1988. Projects include mixdown on Riki Lee Jones *Flying Cowboys*, recording and mixdown on Donald Fagens *Kamakiriad*, and numerous others.

Walter Becker - "The DC 24 is great for bass and guitar. I suggest you check it out". Walter is a member of the popular group, *Steely Dan*.

DYNAMICS 101: A PRIMER

Let's start with what a dynamic controller actually does. No matter how you cut it, these are electronic volume controls. It is a hand on a control, turning the volume down and turning it up again. The hand is really quick and really accurate, but it's just turning a volume control.

SIGNAL CHAIN

Conceptually, dynamic controllers have two internal paths, the signal and the side chains. The signal chain is the path the main signal takes through the unit: through the input circuits, the gain control device and then through the output circuits. The signal chain goes through the "volume control" in the "hand on a control" analogy.

SIDE CHAIN

The side chain is the hand which turns the control. Side chain circuitry examines the input signal and issues a control voltage to adjust the amplification of the signal. There are a number of parameters governing side chain activity, but the four most commonly discussed are threshold, ratio (or slope), attack time and release time. Some dynamic controllers offer adjustment of each of these parameters, while others have one or more preset at an optimum setting for the application.

THRESHOLD

The threshold, like crossing through a doorway, is the point at which gain adjustment begins. When the input signal is below the threshold, a dynamic controller should be like a straight wire. Above, the side chain asserts itself and turns the volume down.

RATIO

Once the threshold is exceeded, just how far the volume goes down depends on the ratio (or slope) setting. An ordinary preamp or a straight wire has a ratio of 1:1, that is, the output level tracks the input level perfectly. A 2dB change at the input produces a 2dB change at the output. A severe ratio is perhaps 8:1 or 10:1. For a 10:1 ratio, a 10 dB blast at the input would rise only 1 dB at the output – *heavy* compression. Kinder, gentler ratios are in the 2:1 to 3:1 range.

ATTACK TIME

Attack time is the time which passes between the moment the input signal exceeds the threshold and the moment that the gain is actually reduced. Attack times generally range between 1ms and 30ms.

RELEASE TIME

Release time is the time which passes between the moment the input signal drops below the threshold and the moment that the gain is restored. Typical release times are between .1 seconds and 4 seconds.

Some of the oldest compressors were called levelers, which are becoming popular again. They had very slow attack times and very long release times to provide volume adjustment of overall program level for broadcast. If you shouted repeatedly, the level would slowly fall off for about 30 seconds, then it would take another minute or so to recover.

COMPRESSORS

A compressor, when the input signal reaches the level set by the Threshold control, begins turning down the signal by an amount set by the Ratio control. Most modern compressors make the loud signals quieter, but do not make the quiet parts louder. (However, by keeping the loud signals under control, you can turn up the output level which will make the quiet parts louder along with the rest of the signal.) Some compressor designs actually do raise quiet signals below the threshold. These designs might be called "upward expanders".

LIMITERS

A limiter is a special form of compressor set up especially to reduce peaks for overload protection. In other words, it is a compressor with a maximum ratio. A compressor is usually set up to change the dynamics for purposes of aesthetics, intelligibility, or recording or broadcast limitations. Once the threshold of a limiter is reached, no more signal is allowed through. A limiter has a relatively high threshold, very fast attack and release times and a very high ratio, approaching infinity:1.

EXPANDERS

An expander is a compressor running in reverse. *Above* the threshold, a compressor reduces the gain; *below* the threshold an expander reduces the gain. A compressor keeps the loud parts from getting too loud, an expander makes the quiet parts quieter.

GATES

A gate is an expander with the ratio turned up. With the proper settings (low threshold and a high ratio), a gate can be applied to remove noise between louder sounds, and is often called a noise gate for the way it can lock out background noise.

GATE / EXPANDERS

A low ratio acts as an expander that turns quieter signals down, while a high ratio acts as a gate that shuts signals off.

SIDE CHAIN EXTRA #1: SEND/RETURN

The gain control voltage is derived from the side chain audio. If you were to put a signal with treble boost into the side chain audio, it would not effect the treble in the main signal path, but it would cause the high frequencies to cross the threshold sooner or more often. Large peaks of treble could be set to cause heavy compression with virtually no compression at other times. What we've just designed here is the basic de-esser, a circuit to remove excess sibilance. With a bass boost you can make a de-thumper and with a midrange boost a de-nasaler. Most compressors have a send and return available in a side chain loop to patch in an equalizer for these purposes.

SIDE CHAIN EXTRA #2: SLAVE

Many compressors and expanders make the side chain control voltage available to connect to a neighboring unit, or to tie internal channels together. This is called slaving or linking the compressors, and it causes the units to compress

simultaneously when only one has an input over the threshold. This feature is normally used to preserve stable stereo imaging, or to preserve spectral balance when the compressors are used in the high and low frequency ranges of a mono signal.

THIS IS ALL VERY INTERESTING. SO WHAT'S THE PROBLEM?

The problem is that heavy compression (low threshold and a high ratio), almost always has nasty side effects. In the first place, the timbre of the sound itself changes; it becomes “hard” and “closed” and not nearly as sweet and open as the sounds you envisioned when you got into this business. Second, attack times optimized for pleasant compression will not track initial transients quickly enough, and many instruments audibly suffer. Third, heavy compression will usually be accompanied by “breathing,” i.e., the background noise rises way out of proportion to the foreground sound as the compressor releases. Bottom line: it just doesn't sound good. Take anybody's compressor, run just about any sound through it, compress it severely and run the results on Family Feud: survey says, 89% of the audience won't like it.

SO WHAT'S THE SOLUTION?

Many designs have appeared throughout the years to produce gentle, smooth, natural-sounding compression. They include tubes, FETs, VCAs, soft-knee compressors, electro-optical attenuators, and self-adjusting attack and release times. Today, some digital workstations compress without snipping transients, by looking ahead into the digital future. Is that cheating or what? So what has Rane done to make its compressors sound so great?

An independent panel of judges has studied Rane's compressor designs and unanimously decided there has been no cheating. Rane has combined a number of perfectly even handed, meat-and-potatoes ideas to make its compressors so capable and transparent that we just *seem* like we're not being fair.

IDEA NUMBER ONE

Use self-adjusting attack and release times. The compressor and expander sections in the DC 24 change attack and release times automatically to suit the program material by using dedicated RMS-sensing ICs in the side chain. If the input is predominantly low-frequency, the times are made more gradual and slowed. If a quick transient comes flashing down the wires, the times are tightened to deal with it. Our experience has shown that attack and release controls, when present, are confusing and easy to misalign.

IDEA NUMBER TWO

Combine an expander/gate function with the compressor. The expander/gate, the compressor (and the limiter: see Idea Number Three) in the DC 24 can be used independently, but a big reason they are together is to share the work of clean compression. An expanded or gated source of sound exhibits less “breathing” when compressed. Instead of looking for another patch cord when you realize you need a bit of gating, you just turn a control.

IDEA NUMBER THREE

Combine a peak limiter function with the compressor. Tracking with this idea of burden-sharing, Rane has put a peak limiter in the same path as the expander and the compressor in the DC 24. With a limiter right there, you won't be asking the compressor to clamp the wild excursions. The limiter, with auto-attack, auto-release and adjustable threshold optimized, will play level police while the compressor persuades more gently.

Rane designed a patented servo-locked limiting circuit, which places the limiter within a servo loop and effectively stops peaks from exceeding the threshold. The attack time varies with the source material, but is never allowed to produce diode-like hard clipping.

REALLY GREAT IDEA NUMBER FOUR, WHICH DESERVES ITS OWN SECTION

Here's the special twist in the DC 24: the two sections, fabulous as they are, can also be assigned to *different frequency ranges* of the same channel of sound. This is not a new idea, but it's a great idea. In the past, the difficulty has been that split-band compression has required a lot of equipment: at least two compressors and a set of bandpass filters per channel, or a very expensive difficult-to-set large unit. What the DC 24 offers is not just innovative engineering but a lot of powerful, interactive functions crammed into one rack space.

SPLIT-BAND DYNAMIC PROCESSING

We haven't talked much about split-band processing, but it's one of the easiest ways to compress transparently. Broadcast stations have used split-band compression for years, often dividing the spectrum into four or five bands. When it's done right, the radio station sounds great: loud, present, with no squashing or pumping at all.

The great Dolby noise reduction systems, from Dolby A all the way through B, C, S and SR, all use some variation on compression, expansion and band-splitting. Dolby's goal has always been maintenance of the purity of sound, with no artifacts of the processing. It works.

Split-band compression works well for several reasons: You can optimize each set of dynamic processors (the compressor, expander and limiter) to a particular range of audio. That is, the ratio and threshold controls can be suited to each part of the spectrum.

You can decide to process different ranges of an instrument differently. You could use no compression at all on the low end of a bass, with heavy compression on the top end to put the string slaps in balance with the bottom. Or you could tighten the boomy bottom up with compression but leave the top less controlled for that open feeling.

Any massive anomaly like a low frequency breath noise for example, only triggers gain reduction within its range, leaving the desired vocal unaltered. And the decidedly unmusical phenomenon of a popped ‘P’ sucking the overall level back 10 dB is a thing of the past.





Figure 1. Results of the Separate / Combine and Dual switches in Crossover Mode.

VISUALIZING ALL THIS

Figure 1 shows a couple of different configurations, depending on the positions of the Crossover / Dual / Combine switches that truly make the DC 24 a multi-function unit. It can be a two way crossover, with independent processing on the low and high outputs. The outputs can be summed with the Separate/Combine switch, and the Crossover can be switched in (see rear panel above), so that processing the low and highs separately can take place in a mono or send/return application. Even though the outputs are summed at Channel 2's output, Channel 1 is still outputting the lows which might be valuable to a bass player running a full range along with a bass bin (see Figure 1).

Figure 2 shows how the gate/expander, compressor, and limiter all can work together on the same program material in a single channel. The vertical axis is the output level, and the horizontal axis is the input level. When all Ratios are set at 1:1, the input and output of the circuit are the same as illustrated by the straight diagonal line running at 45° across the graph. Each of the Threshold controls acts like a "hinge point", activating gain reduction only when the input signal reaches the level set by this control. The Ratio controls how much of an "angle" the hinge will bend, or more realistically how much gain reduction will occur once the threshold is reached. Graphically, the ratio can swing this hinge from 45° (no processing) to almost 90° (full ratio). It is also possible, by adjusting the Thresholds, to have each of these circuits overlap and interact

with each other to develop a dynamic curve. The solid black line shows the curve produced when the controls are set as shown.

In this example, the gate/expander circuit works on the quiet parts, and the compressor and limiter work on the louder parts. The gate/expander can range from just turning down the quiet parts a little to a lot. The compressor and limiter are a lot more flexible when used separately at different thresholds, even though they have the same job of keeping the loud stuff under control. Got it?




Figure 2. Compressor / Limiter / Gate Input to Output Graph

DC 24 APPLICATIONS

STARTING OUT

Sometimes it's necessary to start from scratch. The panel above shows where the controls should be for *no* processing. Then you can adjust each section one at a time.

TWO CHANNEL COMPRESSOR/LIMITER

In this case, the audio path on channel 1 is completely separate from channel 2, allowing you to use it as a stereo unit or for doing two completely different processes to two completely different signals. For stereo use, the front panel has a "Dual/Slave" mode switch that allows you to slave channel 2 to channel 1. This assures that both signals are affected identically. In this application, the crossover is disengaged (this button is located in the middle of the rear panel.) The "Separate/Combine" switch on the rear panel should be in the "Separate" mode. Set the rear panel "-10/+4" switches accordingly, depending on whether you are running your system at -10 dBV or +4 dBu levels.

CROSSOVER WITH BUILT-IN LIMITING

Let's say that you want to run a bi-amped system and process the low end a little differently than the high end. This is a handy way of saving your woofers from over-exursion. In this application the crossover is engaged and set at the recommended crossover point for your speakers, let's just say 1.2k for the sake of example. In this instance the Separate/Combine switch on the far left of the rear panel is set on Separate. Your input would connect to the Channel 1 input jack. This gives you a separate output on Channel 1 (signals below 1.2k, the "lows") and everything from 1.2k and above on Channel 2 (the "highs"). This setup allows you to better contain the low end without unnecessarily limiting the high end. A crossover *and* a processor all in one rack space!

GUITAR

John Albani (Canadian Musician Magazine)- "By now, I'm sure that you have heard that a low stage volume is essential to your sound man getting a better house mix. Well, here are a few suggestions on how to achieve a lower volume without sounding like you're playing out of a transistor radio.

"Marshall's and other 4 x 12 cabinets give a great 'chunky' sound, but it is also accompanied by an annoying 'woofing' on the lower end. This stereo compressor has the unique feature of becoming a two-way crossover with independent low end and high end compressors. With this I was able to achieve what was previously only possible with the dynamics section of the SSL console that was used for my guitar sounds on the Lee Aaron "Bodyrock" album. Take the preamp output

of the loop into the DC 24 Channel 1 Input. The Channel 2 Output should return to the main amp input of the effects loop or the power amp (via your effects). Set the switches to Dual / Crossover / Combine. Now you can set a crossover point on the front panel (try around 400 Hz) and compress the bottom end at a 10:1 ratio. While chugging on a chord where you notice a lot of woofing, set the gain reduction with the Threshold control to read 6 dB. When you hit an open chord, there should be no gain reduction. If there is, back off on the Threshold, not on the Ratio. Now compress the top end between a 1.5:1 to 2:1 ratio, with 3 dB gain reduction when an open chord is hit, to give your sound a lot more attack. Also, no matter where you play on the neck, the bottom end of the sound will be even, without woofing, giving your overall tone punch and clarity.

"Warning: Do not over compress the top end or the pick attack will be slurred. If you want to hear more attack, turn up the top end Level of the DC 24 after setting the above-mentioned compression for the top end.


"Right now you are 99% on your way to retaining your sound or bettering it, without blasting everyone to Palookaville, or deafening your sound man.

"On stage you must work within the tonal range of your instrument. I hear guitarists with huge sounds that are great until the bass player fires up. He can't hear because the bottom from the 4 x 12s is blurring out his bottom end. So you end up in a volume war, which puts you out of the front mix. Try this: Once you have the sound you like, back off on the bottom Level control of the DC 24. Your bass player is already operating in that tonal range and you won't miss the sub lows when he's playing with you anyway."


BASS

Now for you bass guitar players out there...How many times have you been yanked out of the mix by your soundman because you're overdriving the system? You'd love to be able to keep the high-end attack without booming on the low end. Well, try this. Set the switches to Crossover *in*, Separate/Combine to Combine, and Dual Mode on the front panel. Now, plug into Channel 1 from your preamp output, and come out of Channel 2 into your amp. What you have done is split your mono signal, with a crossover point, then you've run it through separate processors and combined the signal back together on the Output of Channel 2.

Where does the unit go in the signal chain? Well, that depends on how you want it to function. If it's a comp/limiter for the input signal, it would go after the bass (if the bass has a line-level output) and before the preamp. If it's to function as a limiter to protect the speakers in the bass rig, it would go after the preamp and before the power amp. Another method



"Straight Wire" Setup



Bass Guitar Setup

is to insert the unit in the effect loop of the preamp. This allows the bass signal to be affected by the pre-amp first, then the comp/limiter, and then sent to the power amp. This can be desirable with tube pre-amps.

This unit can also be used for biamp rigs. For this, it is placed in the signal chain after the preamp and before the power amps. The output from the preamp is the signal that is processed and split at the selected crossover point. For biamp purposes, the Combine/Separate switch should be in the Separate position. Channel 1 processes the lows and channel 2 processes the highs. The low and high outputs are independent and correspond to Channels 1 and 2.

The DC 24 has two great advantages over other compressors—the crossover and the dual channels. It gives you complete control of the signal and processing of it. This is something that wasn't available before in a single unit. One stereo or two mono comp/limiters and one crossover would be required to do what the DC 24 does in a single rack space. This unit solves many compressor blues. For more attack, you can turn up the Level on the top end. Notice that when you stop playing, that amp buzz and hiss goes away. Nice, huh?

Dave Freeman (Bassics Magazine)- "I tested this unit in the combine mode with the crossover set at 200 Hz. I used my 4 string Music Man bass, famous for its ear splitting high end, as test in different channel settings. I set Channel 1 (low end) for mild compression at 2:1 with the Threshold at -10 dB. I set Channel 2 (high end) for heavy compression at 6:1, and the Threshold at -20 dB. I turned the volume and the treble controls on the bass on full, and slapped and popped like a madman. So what happened? Well, the high end was compressed down to the low end level. The sound was balanced and didn't have a compressed tone. I could slap away 'til my fingers went numb without having the comp/limiter clamp down on the entire signal. Impressive results!

"I then tested the unit with my 5 string Ken Smith bass. I set the lows for mild compression at 2:1 at -20 dB Threshold. I set the high end for the same compression but with the Threshold at -10 dB. I wanted just a bit on the bottom for the low B string and less processing on the highs. I slapped and popped on all the strings including the low B. The result was slight processing on the lows which tightened the bottom, but didn't make it *sound* controlled or processed. The highs had subtle compression that sounded natural, unlike others that 'breathe' when compressing."

RECORDING

Use it on bass guitar, piano, drums, vocals—anywhere you've used a compressor/limiter before. The DC 24 gives you more control and a less tortured sound. In fact, split-band processing works so well that a DC 24 sounds good compressing an entire mix (two required for stereo in split-band mode).

Of special interest are instruments which have large level differences in their different tonal ranges. String pops on a bass are one, but flute is another. The higher tones require more breath and are much louder than the lower. Another good application would be a drum mix or submix. A split-band compressor does a better job of smoothing the performance out.

Roger Nichols (Engineer)- He uses the DC 24 primarily on bass and guitar. He sets the Crossover at 100 Hz, the Gates and Compressors to 1:1, and engages the Combine and Dual Mode switches. This gives him separate Limiters to control the high and low peaks separately on a mono signal.

Brent Hurtig (EQ Magazine)- "In the studio, the crossover has some different applications. With Combine selected and the Crossover engaged, a signal entering Channel 1 is split into two bands. These two bands again may receive separate processing. What's different here, though, is that the two bands' signals are merged at the Channel 2 Output. This little exercise allows you to apply different amounts of compression and limiting to the low and high ends of a piano. Or let's say the saxophonist sounds great, but every time she hits the high C she pins the meters: Just the high end of the sax could be limited. Very clever.

"You also can use the Separate mode in the studio. With this setting, the crossover acts like a low pass filter to signals in Channel 1, and like a high pass filter to signals in Channel 2. We found some great sounding guitar, vocal, and keyboard tones using the DC 24 in this equalizer-like manner."

Digital Recording: Use it to compress an extremely wide dynamic range into a signal that won't go into digital overload, i.e. clipping. The limiter is the primary circuit here to keep things under control, but a little compression with its threshold set just under the limiter threshold setting will help keep the limiting even more subtle. Also, the gate can be set just above the noise floor with a low threshold and high ratio to remove mixer or tape hiss between cuts. To control a stereo mix, set the switches to Normal / Separate / Slave.

Microphones for the AP 13 Acoustic Preamp

- **AKG**
- **Audio Technica**
- **Beyerdynamic**
- **Countryman**
- **Crown**
- **Donnell**
- **K&K**
- **Microvox**
- **Shure**
- **Sony**
- **Telex**

**Jeff Davies
Chris Duncan
Rane Corporation**

**RaneNote 131
© 1993, 1999 Rane Corporation**

General application notes and disclaimers:

All of the information is gathered here for the curious regarding “What mics can I use with the AP 13?”. This miniature mic collection was assembled from the respective manufacturer’s published specs. Rane does not guarantee specifications will be accurate when modifications, as required, are made. Rane also assumes no liability for altering another manufacturer’s microphone or warranties which may be voided due to modifications described herein. Each manufacturer has the right to change any specification or wiring without notice. Some information was furnished by actual users of these microphones. If you have additional information on any of these or know of any good models we missed, we urge you to contact the factory so we can update our records.

AKG C 411/B-lock

1/8" locking mini plug

Electret pickup 20 Hz-20 kHz

Phantom power: 6v

Includes: reusable adhesive compound for mounting directly to the instrument.

The 1/8" connector can be amputated then the hot lead connects to the tip (or ring) and shield to the sleeve of a TRS jack in the guitar body, or a mini jack to 1/4" adapter can be used outside the instrument. *See AKG MicroMic Series Note.

AKG C 416/B-lock

1/8" locking mini plug

Hypercardioid condenser 20 Hz-20 kHz

Phantom power: 6v

Includes: built on gooseneck, mounting plate, and windscreen.

Uses the same element as the C 419, but with a 9" gooseneck intended for low profile mounting inside guitar, piano, accordion, autoharp, dulcimer, etc. The 1/8" mini can be amputated and the 2 leads connected to the tip (or ring) and sleeve of a 1/4" TRS jack in the guitar body, or a mini jack to 1/4" adapter can be used outside the instrument. *See AKG MicroMic Series Note.

AKG C 417/B-lock

1/8" locking mini plug

Omni condenser 20 Hz-20 kHz

Phantom power: 6v

Includes: clip tie pin and windscreen.

The 1/8" connector can be amputated and the 2 leads connected to the tip (or ring) and sleeve of a TRS jack in the guitar body, or a mini jack to 1/4" adapter can be used outside the instrument. *See AKG MicroMic Series Note.

AKG C 418/B-lock

1/8" locking mini plug

Hypercardioid condenser 50 Hz-20 kHz

Phantom power: 6v

Includes: built on gooseneck, mounting plate, and windscreen.

This microphone is designed for drum and percussion devices. It is similar in design to the 419 but has shorter gooseneck and is suited for higher SPL applications. The 1/8" connector can be amputated and the 2 leads connected to the tip (or ring) and sleeve of a TRS jack in the guitar body, or a mini jack to 1/4" adapter can be used outside the instrument. *See AKG MicroMic Series Note.

AKG C 419/B-lock

1/8" mini plug

Hypercardioid condenser 20 Hz-20 kHz

Phantom power: 6v

Includes: built on 6" gooseneck and clip, windscreen.

An excellent choice for horn players, the rubberized finish clip attaches to the bell. The 1/8" connector can be mated with a mini jack to 1/4" adapter. *See AKG MicroMic Series Note.

Audio Technica ATM 35cW

Lemo connector

Cardioid condenser 35 Hz-20 kHz

Phantom power: 6v

Includes: built in clip, gooseneck and windscreen.

The lemo connector can be amputated, and the red (pin 3) wire going to the tip (or ring), and the yellow (pin 2) and shield (pin 1) tied together at the sleeve of a TRS jack in the guitar.

Audio Technica AT831c

Unterminated

Cardioid condenser 40 Hz-20 kHz


Phantom power: 6v

Includes: AT8411 tie clip, guitar adapter, AT8116 windscreen.

The red wire goes to the tip (or ring), and the yellow or white and shield ties together at the sleeve of a TRS jack in the guitar body.

Beyerdynamic MCE 5.9

***AKG Micro Mic Series Note:** These mics are incompatible with the AP 13 as it is shipped. A minor modification must be performed to be able to use these microphones. Remove the top and bottom covers and replace R220 (near the INPUT jack) with a 15k ohm resistor. Once modified with the 15k resistor, the input won't be compatible with other mics.



Unterminated
 Omni condenser 20 Hz-20 kHz
 Phantom power: 6v
 Includes: tie clip, windscreen.
 Recommended: MAG 5 or MAG 5.1 guitar clip.
 Mounts for a violin, viola, cello, flute, trumpet, sax, trombone are available. Connect the green wire to the tip (or ring), and the outer shield to the sleeve. Do not connect the inner shield to anything. One of the best sounding mics we have actually tested.

Beyerdynamic MCE 52.16

Lemo connector
 Omni pressure mic 35 Hz-20 kHz
 Phantom power: 6v
 Includes: built on clip.
 Connect the green wire to the tip (or ring), and the outer shield to the sleeve. Do not connect the inner shield to anything.

Beyerdynamic MCE 53.16

Lemo connector
 Omni pressure mic 35 Hz-20 kHz
 Phantom power: 6v
 Includes: built on 85mm gooseneck and clip.
 Connect the green wire to the tip (or ring), and the outer shield to the sleeve. Do not connect the inner shield to anything.

Countryman Isomax B3*

XLR
 Omnidirectional condenser 20 Hz-20 kHz
 (*Use Red Band version, which has lower sensitivity than standard model.)

Countryman Isomax EMW

XLR
 Omnidirectional condenser 20 Hz-20 kHz

Countryman Isomax II C

XLR
 Cardioid condenser 50 Hz-20 kHz
 Countryman Isomax II H
 XLR
 Hypercardioid condenser 50 Hz-20 kHz

Countryman Isomax II O

XLR
 Omni condenser 50 Hz-20 kHz

Countryman Isomax II B

XLR
 Bidirectional condenser 50 Hz-20 kHz
 Phantom power all models: 6v
 Includes: clip and windscreen. Sax or flute clip also available.
 Each microphone comes with electronics in the XLR connector. The connectors have an internal EQ circuit that increases gain. The EQ section in the AP 13 would replace

the EQ circuit in the connector, but mic frequency response is unknown. To install one of these mics without the XLR, amputate it and connect the red wire to the tip (or ring), and tie the green and shield together at the sleeve. When using the XLR, not connecting the green wire defeats the EQ circuit.

Crown GLM 100/E

Unterminated
 Omni condenser 20 Hz-20 kHz
 Phantom power: 6v
 Includes: GLM-UM universal mount, windscreens, tie bar mount and belt clip.
 The red goes to the tip (or ring) and the white and shield get tied together at the sleeve.

Crown GLM 200

XLR preamp
 Hypercardioid condenser 60 Hz-20 kHz
 Phantom power without inline preamp: 6v
 Includes: GLM-UM universal mount, windscreens, tie bar mount and belt clip.
 The GLM 200 comes with a preamplifier module. When using their preamplifier, switch phantom OFF, and use a power supply battery. Connect pin 2 to the tip (or ring) and pins 1 and 3 go to the sleeve. Their preamplifier can be removed and lead connections are same as the GLM 100, with phantom set to 6v.

Donnell Mini-Flex 135

1/4" TRS jack
 Cardioid condenser 80 Hz-16 kHz
 Phantom power (battery): off
 This is an efficient one piece device that installs through the endblock and exits the body via a 1/4" TRS jack in place of the strap button. A pickup may be wired to the ring of the TRS. An 'AA' battery supplies the current for the condenser, so switch phantom OFF.

K & K Silver Bullet

XLR Plug
 Softcardioid condenser 20 Hz-20 kHz
 Phantom power: 6v
 The white wire goes to the tip (or ring) and shield connects to the sleeve. Amputate the XLR and connect white wire according to your needs.

Microvox M400

Gold Phono Plug
 Omni condenser 30 Hz-20 kHz
 Phantom power: 6v
 The center wire goes to the tip (or ring) and shield connects to the sleeve.
 Sennheiser MKE 2 -5
 Unterminated
 Omni condenser 40 Hz-20 kHz
 Phantom power: 6v
 The red wire goes to the tip (or ring) and the blue and shield tie together at the sleeve.

Shure SM11

XLR connector
 Omni dynamic 50 Hz-15 kHz
 Phantom power: none
 Includes: tie tack and tie bar mount
 Recommended: RK279 instrument mounting kit
 Amputate the XLR connector. Connect the red wire to the tip (or ring). Black and shield tie at the sleeve.

Shure SM98A

Includes XLR preamp
 Cardioid condenser 40 Hz-20 kHz
 Phantom power: 6v*
 Includes: windscreen, swivel adapter, supercardioid polar modifier, preamplifier.

recommended: A98KCS Keen Clamp.

This model comes with a barrel preamplifier. In order to guarantee frequency response, the preamplifier works with an adapter that takes pin 2 to the tip (or ring), and pins 1 and 3 go to the shield. *Phantom power should be off when used in this configuration, and two 9v batteries must be used in the preamplifier. It is possible to disconnect the preamplifier and just use the capsule, but response is unknown. Connect the red wire (pin two of the QG connector) to the tip (or ring), and the black or blue and shield (pins 1 and 3) tie together at the sleeve. Set phantom to 6v.

Sony ECM-44BBT

XLR
 Sony capsule part number #A4510056A
 Underminated
 Omni condenser 40 Hz-15 kHz

Sony ECM-55BBT

XLR
 Sony capsule part number #A4510051A
 Underminated
 Omni condenser 30 Hz-18 kHz

Sony ECM-77BBT

XLR
 Sony capsule part number #A4510050A
 Underminated
 Omni condenser 40 Hz-20 kHz
 Phantom power all units: 6v*
 Includes: windscreen and two mic clips for the ECM units.

The regular ECM versions of these mics come with an attached preamplifier capsule. *When using this capsule, a battery must be used and phantom switched OFF. Use an adapter from pin 2 of the XLR to the tip (or ring) of a TRS, and tie pins 1 and 3 to the sleeve. A mic capsule can be ordered separately from Sony without the preamplifier with the part number. Connect the red to the tip (or ring) and the white and shield get tied to the sleeve.

Telex ELM-22PT

Underminated
 omni condenser 20 Hz-20 kHz
 Telex ELM-33PT
 Underminated
 cardioid condenser 100 Hz-10 kHz
 Phantom power all models: 6v
 Connect the red wire to the tip (or ring) and the black and white wires together at the shield.


WIRING

Most condenser microphones come with a balanced output, i.e., three wires. The input of the AP 13 is unbalanced stereo, i.e. two wires, per input. Most of the time, the red wire is the hot or + wire, and should be wired to the tip or ring of the plug. The other two wires, regardless of color (blue, black, white) and the shield will get tied together at the sleeve or common of the plug. This is a very general rule. When in doubt, always check with the microphone manufacturer before making modifications or connections.

DEMOS

To most effectively demonstrate the power of Rane's AP13 acoustic preamp, use a combination of a piezo pickup and miniature microphone—pickup for gain, mic for clarity. You may wish to try out several mics before actually committing a purchase.

Detailed below is a wiring diagram, of a custom made cable, to take the guitar's piezo pickup output, add the microphone, and end up with a single 1/4" TRS connector ready to plug into the AP 13. We suggest retailers make one of these and keep it handy for doing demonstrations. The necessary adapters may be available off the shelf (PROCO, Hosa, etc.) to build this. A 'wye' cable with a TRS female to two 1/4" TS males will work like the diagram below. One TS goes to the guitar and the other goes to the mic. A 1/4" TS female to (XLR, 1/8" , or 1/4") female adapter will get you from the mic to the wye cable.



Demo Mic Adaptor Wye Cable

Home Cinema Systems

- **Dolby Pro-Logic & THX**
- **Home Theater Equalization**
- **THX 44 Equalizer**
- **THX 22 Equalizer**
- **SSE 35 Equalizer**
- **Interfacing Consumer & Pro Gear**
- **Troubleshooting & Alignment**

Jeff Davies
Rane Corporation

RaneNote 132
© 1994-1999 Rane Corporation

Dolby is a trademark of Dolby Laboratories.

THX is a trademark of THX Ltd.

DTS is a trademark of Digital Theater Systems.


Radio Shack is a trademark of Tandy Corporation.

INTRODUCTION

A good home theater pulls you into the movie, letting you forget about flashing video light and amplified sound. A better playback system equals a better experience. Video quality at home can't match 35 mm projection, but the sound quality at home *can* surpass a commercial theater. Subjective opinion about audio reproduction is nothing new—we all know “experts” in the business who will tell you the “right” way to do it. There are several “right” ways of doing it, depending on your philosophy and taste. Realism, hearing through the directors ears, being blown away, and still having a system that sounds good when you play a CD is a multi-faceted goal. System component selection is important. Correct calibration is equally important. Rane equalizers are used in the finest recording studios in the world, including Lucasfilm's Skywalker Ranch. This is why Rane was originally approached to build the first equalizer for Home THX systems. This same technology is available for recreating the original studio mix in your living room.

DOLBY PRO-LOGIC®

In a stereo system, the left and right speakers produce all the imaging, including center mono and ambient signals. In movies, the added center speaker is the *most* important channel in the system, while the surrounds produce the ambience. Dolby Pro-Logic derives four channels with full backward compatibility to stereo and mono, with minor compromises. Rather than re-explain how Dolby Pro-Logic works (*see References*), this note covers a few popular confusing topics.



Surround frequency

response. The Dolby Pro-Logic surround output is mono. Decoder manufacturers may add their own stereo simulators to the mono surround output, delivering simulated left and right surround outputs. During Pro-Logic encoding, surround information gets a 100 Hz high pass filter and a 7 kHz low pass filter. Surround speakers only need to reproduce 100 Hz-8 kHz for Pro-Logic, but require full range for Dolby Digital and DTS systems (more on those later).

Center Channel Modes. Normal mode: Also called **Small** on some controllers. 20 Hz-20 kHz is delivered to the left and right outputs. 100 Hz-20 kHz is delivered to the center output. All material below 100 Hz is delivered to the left and right outputs. This is the mode preferred when using subwoofers (through a crossover, internal or external), and/or with a smaller center speaker that has less range than the left and right. **Wide mode:** Also called **Large** on some controllers. 20 Hz-20 kHz is delivered to left, center and right outputs. This mode is preferred when the left, center and right speakers are all full range and a subwoofer is either not used, or with a improves the “wrap-around” effect of the surround track.

Surround Timbre Matching Equalization: An EQ correction curve matches front and surround response to our ears, to alleviate timbre shifts when sound pans from front to the rear.


Crossover: A subwoofer output, crossed over at 80 Hz, derived from the LCR channels with a response of 20-80 Hz. All other channels reproduce 80 Hz-20 kHz. To be a Home

THX System, other criteria in system response must be met.



Home Cinema-2

Front THX speakers have wide horizontal dispersion and focused vertical dispersion. Within this listening window they must have a flat response and produce 105 dB without artifacts. Ceiling reflections can cause unpredictable imaging, acceptable for music but not for film reproduction. Living rooms are all shapes, sizes and textures, and controlling dispersion guarantees an accurate sound focus toward the screen. Smaller front speakers enable closer speaker placement and localization to the action. Subwoofer(s) take physical size away from the front speakers (more on that later).



Surround THX speakers are dipolar radiation in design, with opposite phase drivers aiming toward and away from the screen and the listener sitting in the “null” area, improving the diffuse effect, imitating a row of speakers as in a theater. Power must be flat to 102 dB without artifacts.

THX subwoofers produce an in-room response from 20-80 Hz at 105 dB without artifacts.

Power amplifiers must meet noise, phase, load and distortion specs, with power required to drive the speakers.

It is certainly alright to have a mixed system of THX and non-THX components. But to be a bona-fide *Home Ultra THX System*, it must include a THX approved controller, amps and speakers at a minimum and be aligned by a certified THX installer who has completed the course at Lucasfilm.

THX media. Lucasfilm’s THX does *not* do any different encoding/decoding than Dolby Pro-Logic—they are the *same*. THX is just a *playback* system for material recorded in Dolby Pro-Logic. The THX Studio media program is designed to produce the highest visual and audio quality discs, but these do *not* require a THX system for playback. These discs look and sound fantastic on any Pro-Logic system, whether it is THX or not. The program came into existence because there was a lack of consistent quality control during the many steps of mastering. With a THX DVD, you can be assured you are getting the sharpest picture and best sound possible with the present state-of-the-art.

DOLBY DIGITAL®

Previously known as AC-3 on laserdiscs, Dolby Digital has been chosen as the standard for DVD and HDTV. Each channel delivers 20 Hz-20 kHz bandwidth to all five channels, along with a sixth “effects” channel used for additional subwoofer information. Crosstalk is much lower than Pro-Logic, providing a wide soundscape. Dolby Digital is not a replacement for THX, for some controllers add the benefits of timbre-matching and *dynamic* decorrelation to the surrounds. Even with discrete surround channels, dynamic decorrelation widens and prevents off-axis speaker localization of mono surround information.



DTS®

Digital Theater Systems' audio compression scheme applies to laserdisc and CD technology for home theater use. Though it has a smaller title library and controller support than Dolby Digital, it has a core of enthusiastic followers such as directors Steven Spielberg and James Cameron. It employs less compression than Dolby Digital, and proponents claim it has a much cleaner sound. A small but growing number of music CDs and DVDs are available in the DTS format.

EQUALIZATION

Equalizers somehow got a bad reputation in the hi-fi stereo world. Sure, there's the argument that the less electronics in the signal path, the fewer artifacts (noise, phase shifting, etc.). Check out the RaneNote, *Exposing Equalizer Mythology*. However, music recordings are rare that did not employ an equalizer at least once during production, and in the case of a movie soundtrack you can be assured there were several. Those bass and treble controls you've been using for years are crude broad-band equalizers. In a perfect listening environment, with flat speakers, you won't need an equalizer. The "flat" speaker has not been invented yet, and even if it existed the room would color its sound. Speakers *sound* different in different rooms, for the room is an extension of the speaker enclosure. Most recording and film studios use equalizers to correct speaker and room response, even in carefully designed studios. Living rooms have even more random acoustics, and reproduction of the same audio as heard in a recording studio requires a similar equalizer.

A quality equalizer when used correctly, makes a system more accurate and transparent to its environment. Don't use an equalizer to solve a problem that can be solved with speaker placement—place the speakers, *then* EQ the room. Good speaker placement tips are usually found in the speaker's manual. An equalizer should be used minimally, with preferably no more than ± 6 dB boost/cut. When setting the equalizer, try to use more cutting than boosting. Use a $\frac{1}{3}$ -octave spectrum analyzer, with multiple readings around the listening area for each channel, and averaged to set the EQ curves. Equalizers in a home theater fine-tune the room, and help timbre-match non-matching speakers. A lot of sounds pan between the LCR channels, so the same speaker model should be used whenever possible. Employing equalization is even *more* important when adding a different model center channel speaker to a stereo pair.

To decide whether or not an equalizer is really necessary, try running a $\frac{1}{3}$ -octave spectrum analyzer in the room to see the differences between channels. If all channels average flat within ± 3 dB, you're just plain lucky.

Left, Center, and Right.

Most room equalization curves have one thing in common.

Higher frequencies require less EQ, while the room affects frequencies below 1 kHz and gets worse as they get lower. Looking at a $\frac{1}{3}$ -octave analyzer with good speakers, the lower octaves will have greater dips and peaks, while the upper octaves may have few if any adjustments. The Rane **THX 44** equalizer was specified by Lucasfilm with this in mind. Rather than use $\frac{1}{3}$ -octave sliders across the spectrum, the LCR

channels get the $\frac{1}{3}$ -octave treatment from 80 to 800 Hz.


Each channel has two tunable

parametric bands that smooth the higher frequencies more accurately with fewer controls and circuitry. Two tunable parametric bands are also provided for the subwoofer. Since the THX 44 is intended as a *system-flattening* EQ, it is intended to be set *once* in the final stages of system alignment. A security cover is provided to prevent temptation to play with all those little sliders. Though THX 44 meets the Lucasfilm standards, it is an ideal system equalizer for all Pro-Logic and 5.1 digital systems. And speaking of 5.1, add a **THX 22** to the surround channels for the same equalization control all around the room.

For systems that need equalization but budget or space considerations rule out the THX 44, the **SSE 35** is a single rack space three channel graphic equalizer intended for home theater. $\frac{2}{3}$ -octave sections are provided for the LCR channels above 160 Hz, and $\frac{1}{3}$ -octave sections are provided for Left and Right (or subwoofer channels) below 100 Hz. Crossovers provide mono or stereo subwoofer outputs if desired, adding a crossover to controllers or subwoofers without one (see *Subwoofers below*). The SSE 35 is a sweetener for systems that allow connections between the processor and power amplifiers, and a must-have when a different model center channel speaker is used. The SSE 35 was designed for Dolby Pro-Logic systems, but in Dolby Digital systems use the SSE 35 without the subwoofer crossover. Additionally, the SSE 35 works well in systems that are primarily used for stereo music listening with occasional home theater use—see the Subwoofer section for more on this subject.



Home Cinema-3



Rane THX 44 Home Cinema Equalizer

SURROUNDS

Equalizing the surround speakers has less of an impact than equalizing the front channels in Dolby Pro-Logic systems, but more important for 5.1 digital systems. Proper surround placement should have the effect of enveloping the listener in a subtle, non-localized sound. Depending on the room, interior decorator, and speaker design intentions, surround speakers may be mounted on the side walls, back walls, aiming down, up, or any way to *maximize* room reflections to give a multiple speaker effect. This is the *opposite* goal of the front speakers that must keep attention toward the screen. A matter of preference exists for some 5.1 digital users that like a point/source surround, but this is less like an actual movie theater and may not translate with how the director intended the surround sound field. Any localized sound to the back or side that makes you want to look over your shoulder is not good, until they come up with 360° screens. Surrounds “fit” in the system better when they have the same relative timbre to the front speakers. In-wall surrounds benefit from equalization when chosen for their appearance over sound. Even though Dolby Pro-Logic surround information rolls off below 100 Hz and above 7 kHz, there will still be some signal present outside of this range. This is unintended “leakage” from the Dolby Pro-Logic decoder, and doesn’t need to be boosted by the equalizer—leave it flat outside of this range; the front speakers will reproduce it. Even with 5.1 digital systems, equalizing the surrounds is not as important as equalizing the front channels, but the purist will use the same $\frac{1}{3}$ -octave equalization as the front channels.

The best choice for surrounds is the Rane **THX 22** Stereo Equalizer. It has the same $\frac{1}{3}$ -octave graphic/parametric combo and look as the THX 44, in a single rack height chassis. It may also be used in any stereo application, such as auxiliary speakers in a home where the main system serves another room or rooms for music. For those who would rather not deal with a parametric, use the Rane ME 60. It is the easiest equalizer to use with a $\frac{1}{3}$ -octave realtime analyzer.


SUBWOOFER

Fact: low frequencies require a larger speaker enclosure to provide the larger wavelengths. Since a small center channel speaker needs to be placed close to the screen, and speakers and drivers need to be matched for left, center and right, large full range enclosures are impractical for the front channels—hence the need for subwoofers. Two subwoofers provide greater room coverage, though one can be sufficient if placed correctly. A second subwoofer will activate different room nodes and provide more even coverage throughout the room, but if incorrectly placed can cause more problems by canceling frequencies at the listening position. Frequencies below 100 Hz are hard to localize, and movie soundtracks are mixed with this in mind. *Tip: to place a subwoofer, put it in the primary listening chair, and turn on the pink noise. Crawl around places in the room where the subwoofer might go, reading your SPL meter (try corners). The place that gives you the loudest reading is the best place to put the subwoofer.*

The reasons for using stereo left/right subwoofers can use the same points. Two subs can smooth out different resonances from different parts of the room. Music CDs often carry stereo information in the sub range, especially in some classical recordings. Turning off the main speakers and just listening to stereo subs proves the point. Symphonic and pipe organ works reveal a wealth of out-of-phase sub information that is not directly localizable, but perceivably stereo. If a stereo recording has out-of-phase sub material, it may be cancelled out in a summed mono system.


So—if the system is primarily a home theater, mono sub(s) can suffice; if watching movies is secondary to playing music, stereo subwoofers or large full range speakers are called for. Movie soundtracks sound just as good with stereo subwoofers, and music sounds even better. The **SSE 35** is one of the few components available that delivers stereo subwoofer outputs.

Equalizing a subwoofer is admittedly not an easy task. Using pink noise and a spectrum analyzer is tricky, several readings must be taken and averaged, and a $\frac{1}{3}$ -octave analyzer won’t tell you the narrow problem frequencies that get even




Rane THX 22 Stereo Equalizer


THX 44 or THX 22 GRAPHIC EQUALIZATION CHART



LEFT



CENTER



RIGHT

more specific in the sub range. A frequency generator and an SPL meter might show general problems, but use your ears as the final judge. Just run some pink noise through the subwoofer, and listen for an even, low roar without any coloration. You might find a strong resonance at one particular frequency. Be sure you're not hearing rattles coming from the environment, these need to be dealt with mechanically.

A parametric equalizer is called for once the best subwoofer position is found and rattles have been dealt with. The **THX 44** has a two band version, which should be enough if you are using a THX approved subwoofer. For more bands, see the Rane PE 17 in the professional products division.

Setting a parametric by ear is a bit of an art—someone with good tuning ears can use the following method: Begin with all LEVEL controls set at “0”. One band at a time, turn up the Boost control about 6 dB. The BANDWIDTH should be pretty tight, about a tenth of an octave. Sweep the FREQUENCY control slowly until you hear an obvious volume increase at a certain frequency (you can use an SPL meter or your ears). There's your problem point, so sweep a little and listen to how wide of a “bump” it is, and adjust the BANDWIDTH control to fit the bump. Now pull the BOOST control down to a negative setting where things sound smooth. Repeat this process for the other bands as necessary.

SUBWOOFER CROSSOVER

A crossover takes the audio spectrum and divides it, so that frequencies below the crossover point are sent to the sub, and frequencies above are sent to the main speakers. This lets the whole system have a wider dynamic range with less distortion. Bass requires a lot of power to reproduce correctly, and the subwoofer needs it's own amplifier. A crossover circuit for the subwoofer is found in all THX controllers, frequency set at 80 Hz. Dolby Pro-Logic components may have a sub output added, but it is just a summed mono signal with maybe a low pass filter that removes nothing from the LCR channels. Some powered subwoofers also may claim to have a crossover, but in reality it's just a 6 or 12 dB per octave low pass filter that does nothing to reduce the drain and increase the headroom of the main speakers. In all Dolby Pro-Logic center modes (except Wide or Large), a bass-splitting circuit only sends information over 100 Hz to the center and surround outputs and keeps the bass in the left and right channels. Sure, some low frequency material exists at the center and surround outputs, but is leakage, rolled off, hard to localize, and nothing that isn't already in the left and right channels, so don't fret about trying to reproduce it in the surrounds. All are good arguments to the benefits of a 5.1 digital system, which has the dedicated subwoofer channel with no crossover required. However, some smart controllers will include a subwoofer crossover anyway to route bass information from non-5.1 sources to the subwoofer — check your manual.

INTERFACING CONSUMER AND PRO GEAR

The THX 44, THX 22, ME 60 and SSE 35 all sport RCA connectors to make connection easy, following the consumer audio standard. Questions arise with non-RCA balanced equipment, such as a Rane RPE 228d Programmable Equalizer or PE 17 Parametric Equalizer. Here's a can of worms that books are written about, but here are a few basics, starting with definitions and ending with a cure.

Balanced: One audio signal on three wires: positive (+) signal, negative (–) return signal, and a shield ground. The audio signal voltage is on the (+) and (–) wires, and the shield protects the two from outside interference as well as providing a ground reference. Use balanced connections in professional installations where cable lengths exceed 10 feet (3 meters). Balanced is also sometimes referred as +4 dBu because of it's high typical signal. Connectors can be XLR (Cannon or 3-pin), screw terminal, or ¼" TRS (tip-ring-sleeve) phone plugs.

Unbalanced: One audio signal on two wires: positive (+) signal, and a grounded shield. Since the shield also contains the (–) signal, it is more susceptible to outside interference. This is usually not a problem with short cable runs (under 10 feet or 3 meters). Unbalanced is sometimes referred as –10 dBV because of its lower typical signal. Connectors are usually RCA or ¼" TS (tip-sleeve) phone plugs.


Connecting these two systems cause level mismatches, ground loops and hums, and headaches for anyone attempting to get the two to talk to each other. For more in-depth information on grounding solutions and custom cables, refer to the RaneNote, *Sound System Interconnection*.

Rule #1. Any signal that is running on an unbalanced (2-conductor) cable longer than 10 feet (3 meters) is likely to hum. The longer the cable, the worse the hum. Balanced equipment, isolation transformers or interface boxes can help, but any line over 10 feet must use two conductors with shield.

In consumer (unbalanced) systems:

1. Locate the equipment as close as possible to each other.
2. Keep cables as short as possible, using good quality shielded cables.
3. Keep AC cables away from signal cables. Never run them parallel to each other.

Rule #2. Audio equipment designers don't read the same textbooks regarding signals and grounding schemes. Some will have great ideas that are compatible with their own equipment, but not with the rest of the world. The worst offenders are grounding rules. Some ground their chassis—others don't. Some connect sleeve to chassis ground—others to signal ground. Some have grounding AC line cords—some don't. Some include RF suppression, some don't. Rack mounting sometimes solves ground problems, other times causes ground problems. Even though the *Audio Engineering Society* has decreed that pin 2 is 'hot' on balanced XLR connectors, some still use pin 3 out of habit. Often the answer lies in trial and error, one unit at a time along the signal chain, from the amp back to the source. When you get to a piece that does appear to cause hum, it may be just in the interconnection — you'll get a “no problem found” reply on the bill if you send it in for repair. Good sources to help avoid this, along with *Sound System Interconnection*, are in the references. Troubleshoot along the signal chain from the speaker




Rane BB 44X Balance Buddy, Front and Rear Panels

back to the audio source, one component, one wire at a time, following rules and checking grounds. There are unfortunately no easy steps to follow, just several items to try.

1. Before changing any grounds, turn off the system or at least turn down the volume to prevent pops or nasty surprises.
2. Try a different or better quality signal cable. Different cable companies use different shielding practices.
3. Try connecting a heavy gauge wire from one chassis to another, possibly by a rear panel screw and a star washer.
4. Try reversing the polarity on non-polarized AC plugs.
5. For RF suppression, ferrite beads or ferrite cable clamps can be installed at the input connector.

Rane makes a couple of solutions to get out of all this, the Rane **BB 44X** and the **BB 22**. The BB 44X is a 19" W x 1.75" H box with no controls, only XLR and RCA connectors. There is no power supply or active electronics to introduce any noise into the system. It uses high grade transformers to isolate, balance or unbalance, and level match between +4 dBu and -10 dBV signals (depending on how you connect it). There are two female and two male XLR connectors, each going through a transformer to an RCA connector, equating to stereo in and stereo out of a balanced device. The **OPTION 88** accessory expands the 4 channels into 8 within the same enclosure, providing 4-in and 4-out. The BB 44X solves all grounding problems when used with properly wired cables. The BB 22 is half of a BB 44x, simply converting two unbalanced RCA inputs to two balanced XLR outputs in a small 5" box. No more overloaded inputs. No more ground loop. No more hum.



Rane BB 22 Balance Buddy, Top and Rear Panels
(jacks are only on the rear).

ALIGNING A HOME CINEMA

Once everything is hooked up correctly, situated, rattle and hum removed, the tuning stage can finally begin. Basic tools required are a $\frac{1}{3}$ -octave realtime analyzer and a hand held SPL meter.

Most installers use a Radio Shack SPL meter—it works, it's inexpensive, and easy to find. Set it to *C Weighting, Slow*. It's primary job is balancing loudness of all channels during the Dolby test signal found on all decoders. Place the meter at arms length in the central listening position. Seated ear height is the best, but do try keeping it away from reflective surfaces or seat cushions. Setting the meter to read 75 dB on all channels during the Dolby test sequence check with the preamp volume at "0" will bring the system to Dolby Reference Level. In other words, the living room is calibrated to the same volume of a Dolby Stereo movie theater, and is the "proper" level to watch a movie. If your preamp does not have a 0 dB indicator of this level tied to the volume control, you could mark or create one, say at "noon" or "2 o'clock" and set channel levels at the power amps.

Choose an analyzer that is $\frac{1}{3}$ -octave, with a flat response mic and a pink noise source, and preferably with a memory averaging function. Setting EQ correctly requires mic readings from different listening area locations and then averaging them. Goldline and Audio Control are good analyzer manufacturers. Minimize any background noises like ventilation systems, traffic, or wild animals because these may throw off analyzer readings. It is not necessary to endure pink noise at high volume. 75 dB is recommended, but can 65 work as well. Ear plugs are allowed.

If the amplifiers have level controls, turn them all down. Connect the pink noise output of the analyzer directly to each input channel of the equalizer. With the equalizer set flat (all sliders at 0), turn up the amplifier level so the SPL meter reads 75 dB. Beginning with the center channel, analyze, average, and flatten each channel one at a time per the analyzer instructions. Getting the pink noise to be read flat on the analyzer within ± 3 dB is the goal.

Don't be picky with irregularities above 1 kHz. The purpose is to correct the power response in the room, not redesign the speaker. The lower the frequency, the more the room affects it.

If the amplifiers don't have front volume controls, connect the pink noise with an RCA "wye" cable to a left and right aux input. With Dolby Pro-Logic on, the center channel produces the noise. Start with center and then do left, then right. Disconnecting the right channel moves the noise to the

left, and vice versa. This method only works for the front channels.

If you need to equalize your surrounds and the power amplifier doesn't have level controls, run the pink noise into the left input on your processor temporarily, set the mode to Stereo or Bypass (so that no processing will occur) and take the left preamp output to each surround equalizer channel in turn, thereby using the volume control on your processor.

If you are aligning a THX or 5.1 digital system or are using an R-2 Analyzer, refer to the Home THX Audio System Room Equalization Manual available from Lucasfilm or the Rane website. The equalizer is not set any differently for a THX system, it's just a more comprehensive manual.

ANALYZING WITH NO ANALYZER

Sometimes an analyzer just isn't available without great expense or wait, but with a less than \$60 investment, equalizers can be aligned and systems greatly improved. Tools: the Radio Shack SPL meter, and an audio test CD. This method doesn't average readings throughout the room, but works great for creating the "sweet spot" for one listener in the center of the room.

There are several audio test CDs available at better record and hi-fi stores, but look for one with $\frac{1}{3}$ -octave test signals on individual tracks. This method can be useful before installing an equalizer to see how flat the system is. Get a pencil and plot the band readings with the supplied chart to see the response of the room before equalization.

Start by aligning the system to Dolby level as described above. Adjust the output from all the speakers in the system to 75 dB, with the subwoofer disconnected or its amplifier off.

Begin with the processor in Dolby Pro Logic Mode. The test signals appear primarily in the center channel. Switch your processor to Center Only mode if available—if not, disconnect the other speakers or channels or turn their amplifiers off. Select the CD track that corresponds to the first slider on the EQ. On the THX 44 or THX 22 it would be 80 Hz. On the SSE 35 it would be 160 Hz. Slide the filter on the EQ until the SPL reading hovers around 75 dB.

Advance the CD track to adjust the next slider and adjust the equalizer. In most rooms you will find more variance in the 80 to 800 Hz range. To double check your adjustments, go back to the track you started with and advance tracks, watching the meter for any bands that deviate.

In Dolby Pro-Logic, mono signals go to the center. The $\frac{1}{3}$ -octave frequencies on the test disc are mono, and go to the center when left and right are connected. To get signal to the left channel, disconnect the CD player's right output. Reverse this to get right channel signal. After reconnecting left and right amplifiers and speakers, analyze and equalize each in turn.

Do the sub channel last. Reconnect the subs, and connect

both CD channels, with center in OFF mode or disconnected. You will find an usually large amount of level variations from 40 through 100 Hz. These variations change in different locations in the room. Look at the meter reading range as the needle is waving, and use the center of that range to go by. Just try to correct major peaks and valleys, and trust your ears for an uncolored sound.

Make sure you re-test all channel levels with the Dolby TEST signal and the SPL meter to 75 dB when you are done with the equalization process.

LISTEN

You may now trust your ears to adjust the subwoofer. If there are differences between the front channels, use the center channel as a reference and get the left and right to sound similar. This is where higher frequency controls above 1 kHz are better adjusted by ear than trusting the meter or analyzer. Most people find that a completely flat equalizer setting can sound overly bright. Not equalizing much above 1 kHz helps, and allowing a natural rolloff of high frequencies in the front channels, but not much more than -3 dB at 16 kHz.

The above approach is a shortcut to doing the room with an averaging analyzer, but it allows demonstrable results. Just listen with the equalizer *in* the system, and *bypassed*. Use the enclosed security cover to guard your valuable settings, and keep a record with the supplied equalization charts.

REFERENCES

1. J. Sehring, "The Advantages of Stereo Subwoofers", *Audio*, February 1994.
2. D. Bohn, "Exposing Equalizer Mythology", *Rane Note*, 1986
3. T. Holman "New Factors in Sound for Cinema and Television", presented at the 89th Convention of the Audio Engineering Society, Los Angeles, CA, 1990 September 21-25.
4. Lucasfilm LTD, "Home THX System Room Equalization Manual", 1993
5. D. Bohn, "Sound System Interconnection", *Rane Note*, 1985
6. G. Davis and R. Jones, "Sound Reinforcement Handbook", Yamaha Corporation, Hal Leonard Publishing.
7. P. Giddings, "Audio System Design and Installation", Howard Sams Books #22672
8. R. Dressler, "Dolby Pro Logic Surround Decoder Principles of Operation", Dolby Laboratories 1988, 1993
10. H. Ott, "Noise Reduction Techniques in Electronic Systems", Wiley Interscience

DC Powering Rane RAP (Remote AC Power) Units

- Voltage & Current Requirements
- Connector Wiring
- Battery Operation
- DC-DC Converters

Dennis Bohn
Rane Corporation

Mark Wentling
E&E Exports

RaneNote 133
© 1991 Rane Corporation

Introduction

This note explains how to operate Rane units designed for remote AC operation from DC power—whether batteries, DC-DC converters or DC power supplies. The techniques shown *do not require any unit modification*. All wiring is external to the unit. There are techniques that *do* require unit modification, but these are better left explained by our Technical Service personnel.

This note applies only to these listed products:

AC 22/22B
AC 23/23B
AD 22d
AP 13
AVA 22d
CP 31/52/64
DA 26/216a
DC 22/24
DMS 22
GE 30/130/215
HC 4/6
MLM 42/82a
MM 8z
MP 22z/24z/2016
MS 1b
PE 15/17
PS 1
RPE 228d
SM 26B/82
SRM 66
SSE 35
THX 44/22
TTM 52i/54i/56
VP 12

Background

First, a little history helps unravel our twisted DC-AC-DC path.

In 1987, Rane established a precedence among pro audio signal processing manufacturers by adding DC jacks to all its units. These jacks allowed users the option of powering several units from one remote power supply. Advantages cited included lower noise, system power sequencing, universal safety approvals, etceteras¹. That same year saw Rane call for the pro sound industry to adopt remote power supplies as standard².

A year later in 1988, Rane joined an industry-wide committee attempting to see if sufficient agreement could be reached to get the Audio Engineering Society (AES) to issue a Recommended Practice document concerning remote power supplies. It was to include voltage type, levels and connectors. Two years later in 1990, the committee disbanded without reaching sufficient agreement for AES action; however, a summary of progress did appear in the AES journal³.

Much was learned, and Rane decided to continue the pursuit of a separate remote power supply standard of its own. Participation in the AES committee demonstrated that using AC remote power was, overall, superior to DC power. Rane took this to heart and nicknamed their revised line of products *RAP*, an acronym for Remote AC Power. All units equipped with black DC jacks were redesigned with red AC jacks, and a remote power supply was shipped with each unit. Gone was DC, and gone was the option—now remote power was required⁴.

But Murphy likes to play her little jokes. Now we find ourselves once again explaining how to run Rane products from DC. Life is never dull.

DC Powering RAP Units

It came at us from several directions: Disney wanted to run Rane units from the 24 volt batteries used to power their magnificent floats. Consultants and contractors needed ways to comply with emergency paging requirements dictating that certain audio equipment run from batteries upon AC power failure. Car installers started calling about modifying units for competition auto installations. Requests came in for retrofitting Rane units for the upgrade RV market and even executive jets. A Southwest police department wanted to install FPE 13 parametric equalizers into police cars to improve intelligibility of reception. There was even a sightseeing outfit installing Rane products into tour buses for use by their guides.

Mark Wentling from E&E Exports (Los Angeles) gets the credit for elegantly solving the puzzle of how to DC power units designed for remote AC use. Mark worked for years as an audio design engineer (MXR, Music Man, Fender, to name a few) before getting into the export business. One day he looked at our voltage doubler circuit and mused that maybe if we were to add just two diodes (soon to become known as *the DC diodes*) it should be possible to power all Rane RAP units from DC. Amazing. He was right.

It took over a year to add the DC diodes to all the units⁵. This is a long time, but Rane knew that once the re-engineering was complete the actual cost of the diodes was small

enough to warrant adding them to every RAP unit. As suspected, a grateful few benefited immensely.

Minimum DC Voltage

All Rane RAP units operate from an 18 VAC remote transformer equipped with a center tap. This powers a dual output DC voltage doubler found in each unit. The voltage doubler steps up the low AC voltage (necessary for world-wide safety agency compliance and exemption) to high enough levels to allow each unit's positive and negative voltage regulators to operate. These voltage regulators put out the industry standard levels of ± 15 VDC (*see the DC-DC Converter section for an exception*). This voltage runs all the internal circuitry and indicators.

These regulators require an absolute minimum of ± 17 VDC to operate. In a typical RAP product the voltage doubler circuits create at least ± 20 VDC. Even with a 10% low AC mains line (normal pro audio design-limits) this guarantees regulator operation. Taking the ± 17 VDC number and allowing an additional one volt drop across the added DC diodes brings us to a *minimum required DC input of ± 18 VDC* to guarantee normal operation.

Maximum DC Current

Every RAP unit has the maximum required AC current silk-screened under the RJ12 modular POWER jack and also listed on its Data Sheet. The DC current equals *one-third (1/3) the AC current*. (See the VC 18 Voltage Converter Data Sheet for a listing of these DC currents.) The DC current is much less than the AC current for two reasons: DC eliminates the diode AC-to-DC conversion factor and the inefficient voltage doublers. These two combine for a factor of about three.

Connector Wiring

Rane uses RJ12 modular jacks wired per Figure 1. The details and history of this jack appear in *Rane Note 121*⁴. There is a distinguishing factor to note about this wiring. Even though it is not obvious why this configuration was adopted (it has to do with accidental telephone hook-up), the wiring is *symmetrical*. Although unintentional, this produced the very favorable byproduct of the plug *not being polarized when used for DC voltages*. When wiring-up or buying cable, it is not necessary to worry about pin numbers, or beginning from right or left, or whether it uses twisted or straight flat cable. Just remember: negative voltage goes to the center two pins, positive to the next outer pair, and the common goes to the two outside pins. Two parallel conductors for each connection are standard. This reduces wiring impedance and doubles current capacity. Rane uses and recommends 26 AWG size wire (two parallel 26 AWG wires are rated at 1.5 amps). Also observe the following:

Use only 6-wire cable and modular connectors for all interconnection. Common 4-wire cable and connectors will not work. Do not substitute 4-wire for 6-wire cable and connectors.




Figure 1. RJ 12 Modular Jack & Plug Wiring




Figure 2. Battery Wiring for RAP units




Figure 3. DC-DC Converter Wiring

Battery Operation

Figure 2 shows the correct wiring for battery operation. Note this requires two batteries wired in series to create the necessary bipolar input voltage. One is not enough. If only one battery is available, see the next section on DC-DC Converters. Also note the batteries must be rated 24 V. The more common 12 V batteries do not provide enough input voltage to operate the regulators. If the regulators do not operate, the unit's noise performance is apt to be unacceptable. The unit *will work*, but it *will not work properly*. Sometimes you are lucky and can get away with it, but not too often. We have seen enough noisy 12 VDC installations to make us recommend 24 V batteries *only*.

DC-DC Converter Operation

Powering RAP units from only one battery requires a DC-DC converter. Figure 3 shows the typical arrangement. What is not common are the output voltage levels of ± 18 VDC. These satisfy our previously discussed minimum for regulator operation. Unfortunately the more commonly found ± 15 VDC converters should not be used. Again, maybe you will be lucky and not have noise problems, but don't bet on it. Higher output voltage converters may also be used, but should not exceed ± 24 VDC. Greater than this begins to tax the regulator's maximum power limit.

Caveat: The PE 17 is an exception. The PE 17 uses internal ± 17 VDC regulators, and requires at least a ± 20 V DC-DC converter. Beware.

Remember to observe the maximum DC current requirements when buying the converter.

DC Power Supply Operation

A DC power supply is an alternative to using batteries or DC-DC converters. Figure 4 diagrams the hook-up required. As shown, wire a dual power supply (or two single supplies) in series to produce a bipolar output voltage. Setting each supply for 18 volts results in a single ± 18 VDC source. This arrangement is popular for emergency use where many RAP units run from one large DC power supply operated by a master AC mains UPS source.




Figure 4. DC Power Supply Wiring

References

1. D. Bohn, *Rane Note 118: Remote Power System*, (Rane Corporation, 1987). Although out of print, interested parties may obtain photocopies upon request.
2. D. Bohn, "Remote Power Supply Standards: A Proposal," *S&VC*, Vol. 5, pp. 70-78, Nov.15, 1987).
3. "Summary of Progress Before Dissolution of AES Working Group on Power Supply Interfacing," *J. Audio Eng. Soc. (News of the Standards Committee)*, Vol. 39, pp. 275-276 (Apr 1991).
4. D. Bohn, *Rane Note 121: RAP—Remote AC Power: An Idea Long Overdue*, (Rane Corporation, 1989).
5. The PS 1, RPS 4 and SM 82 RAP units do not have DC diodes. On the PS 1 and the RPS 4 it was felt they would never be needed, and on the SM 82 the addition proved too difficult.
6. Richard Clark & David Navone, "Isolation," *AUTOSOUND 2000 TECH BRIEFS*, June/July 1992, pp. 149-153. Subscriptions and individual copies available from Autosound 2000 Tech Briefs, 2563 Eric Lane, Suite D, Burlington, NC 27215, (910) 570-0341. Very highly recommended reference series for anyone involved in automotive sound installation. These guys are the gurus—no one knows more.

Good Dynamics Processing

- **A COMPRESSOR**

- **Threshold**
- **Ratio**
- **Attack time**
- **Release time**
- **Gain control**
- **Hold time**
- **Compression caused distortion**
- **Signal measurement**
- **What to look for when buying**

- **A LIMITER**

- **A DOWNWARD EXPANDER**

- **What to look for in an expander**

- **A GATE**

- **HOW TO “KICK THE TIRES” OF A GATE**

A COMPRESSOR

Producing a good, effective compressor is not a trivial task. It has to “rein in” excessive program dynamics without being unduly intrusive. It does this by reducing the gain in the audio path when the signal level exceeds a pre-determined threshold. Four important parameters need to be controlled. Refer to Figure 1 for the following definitions.

Ray Bennett
Rane Corporation

RaneNote 141
© 1998 Rane Corporation




Figure 1. Short Segment of Signal Dynamics

Note that the signal increase is too fast to keep under control with the illustrated attack time. In this case, either the attack time must be decreased, or the high frequency content must be limited. Although not totally obvious, limiting the high frequency content also limits the rate-of-change of the signal level—that is, the sound can only get louder at some maximum rate.

Threshold

The signal level above which the compressor reduces system gain. This parameter is almost always adjustable. Note that a threshold that exceeds the system clip level is essentially a bypass function. A workable range for this control is -40 dBu to $+20$ dBu. Notice that if the threshold is set to a low level of -40 dBu, the compressor begins to look very much like a Leveler or Automatic Gain Control (AGC).

Ratio

A measure of how much the gain will be reduced under given conditions. A ratio of 2:1 means that once the signal level exceeds the threshold, the output signal level is allowed to increase by only 1 dB for every 2 dB of input increase. The desired setting during use should be the minimum that will provide the required overload protection. Normal ratio for a compressor ranges from 1.5:1 to 10:1. This parameter is almost always adjustable. Note that a ratio of 1:1 is analogous to a functional bypass.

Attack time

How long the compressor takes to control the signal after the actual overload occurs. In a good compressor with adjustable attack time, expect a range of adjustment from 500 microseconds (us) to 100 milliseconds (ms). Some manufacturers specify this time from the beginning of overload to some arbitrary gain reduction – perhaps -3 dB. Others specify it from the start of overload to the time the gain stops changing. Perhaps the most meaningful is to specify from start of overload to when the system gain is within, say, 3 dB of the final control point for 10 dB of gain reduction. Unfortunately, in the war of specsmanship, this measurement suffers by

comparison. In a system where this parameter is adjustable, ridiculously short times are often spec'd just to look good. Some manufacturers state the attack time as from onset of a very large, very fast overload to the *beginning* of gain control. This is an artificially short time since it really doesn't reflect how long it takes to get the overload under control. The important consideration is "how long does it take to control an overload?" Equally important as speed, is the *shaping* of the attack function. If badly done, even a slow attack will sound abrupt and "clicky." Unfortunately, the buyer is not usually informed as to this critical part of the design. This is where a reviewer's article could pay off, as well as a carefully done listening test by the potential buyer. If during testing, the compressor sounds intrusive when the attack is reduced below 1 ms, try a different compressor.

Release time

How long the compressor takes to relinquish control once the overload passes. Same problems of specification as the attack time, but of less consequence. Normal adjustment range is from 100 ms to 3 sec or more. A short release time of 100 to 500 ms is a good starting range for spoken voice, while the longer times are better for instrumental music. This time inversely affects the distortion added by the compression process, as will be shown. Release time is usually adjustable. Here again, the buyer is at the mercy of the designer. Much has been done with release circuitry over the years to produce good compressors. Terms like "program dependant" and "dynamics dependant" abound. Some have genuine meaning, while some are hype to get the buyer's attention. Bottom line? How does it sound when compressing?




Figure 2. Signal Dynamics with Limited Frequency Response

In this example, note that the signal attack time is slow enough to be kept under control. The hold time, if adjustable, could have been shorter.

Gain control

This is sometimes called “make up gain.” It is used to adjust the output level to the desired optimum. This control usually has a range of plus and minus 15 to 20 dB. A center detent at 0 dB of gain is often provided. Beware of too much control, as it makes it difficult to adjust to the desired value. Some feel that ± 10 dB is plenty. The author believes ± 15 dB is a good range. If the user is employing lots of compression, then more gain range might be appropriate. However, if subtlety is the buyer’s style, 10 dB may be better.

In addition, the following terms may be encountered:

Hold time

How long the compressor takes to *begin* the release process once the overload passes. Not normally spec’d in a compressor, but a valid parameter nonetheless.

If provided, the hold time allows the operator to maintain the important “tail” of cymbals, triangles, etc., without distorting the dynamics. In most compressors, this is accomplished by making the release time quite long. A hold function allows the operator to more accurately tailor the compressor function to the dynamics of the program material. If the adjustment is provided, it should be from 50 ms to 1 sec or more. Adjusting the hold time entails using the minimum setting that preserves the program dynamics. Speech would require minimum hold, while solo instruments such as vibraphone would probably sound better with a hold time of 1-2 seconds.

Compression caused distortion

This is largely ignored but is very significant. In one U.S. made compressor, which is generally considered to be quite good, operating with 10 dB of compression, minimum release time, ratio of 10:1, at a frequency of around 100 Hz, the distortion exceeds 5% THD (Total Harmonic Distortion)! This is not at all abnormal. As the release time is increased, this distortion decreases but never goes away completely. There is a technique that will eliminate it completely, but that is the subject of another paper. Incidentally, the author believes this distortion often causes the operator to adjust the release time longer than should be used. The assumption is that the sound improved because of the longer release time, but perhaps it actually sounded better because of the reduced distortion (to a point). The condition that saves the day is that this distortion is relatively difficult to hear, since it tends to be masked by the program dynamics that brought it about in the first place. The same effect masks occasional clipping. For example the cannons in the *1812 Overture* are almost totally clipped but still manage to sound reasonably good. Try treating a vocal that way and the singer will probably be unrecognizable.

The source of this compressor caused distortion is ripple remaining superimposed on the control voltage presented to the voltage-controlled amplifier (VCA). Because the rectification and log conversion is full-wave, the ripple is at twice the signal frequency. The release time acts as a single pole low-pass filter, which reduces this ripple at the rate of 6 dB per octave. That is, as the signal frequency goes up, so does the ripple frequency. As the ripple frequency increases one octave, it is reduced in amplitude by 6 dB (by one half) by the action of the release filter. Since the release filter is set to

around 0.5 seconds or more, its implied “rolloff frequency” (-3 dB point) is 0.32 Hz or less. (The formula is $\text{freq} = 1 / (2 \cdot \pi \cdot \text{TC})$ where TC is the time constant – 0.5 sec in this example.) One would think that 0.32 Hz is so far below any signal of interest that there just can’t be any ripple remaining. If the VCA didn’t deal with it the way it does, that supposition would be true; however, the VCA functions by multiplying the control voltage times the signal voltage. This means that the ripple (at twice the signal frequency) multiplied times the input signal causes even-ordered harmonics. Often the designer has worked hard in the design process to eliminate the generation of these harmonics elsewhere in the audio path, only to generate them in the VCA.

Signal measurement

Much has been written concerning the proper rectification/log conversion process used to measure the incoming signal level. Some advocate peak logging, while others say that RMS log conversion is the only way to go. It must be kept in mind that, by definition, RMS log conversion requires averaging over several cycles of the signal. At most frequencies this makes the attack time unacceptably long. Often a compromise is used where the averaging time for the log conversion process is essentially a peak function at low frequencies and RMS at high signal frequencies. In-between it’s anyone’s guess. This sounds sloppy, but in practice, it works quite well. Ultimately, the release time filter converts the peak log to an averaged peak log. Since the logging process reduces the apparent extremes in the signal dynamics, this averaging is a reasonably good first-order approximation. One way to improve the ripple removal is to use a two-pole release filter. This two-pole approach reduces the ripple twice as fast – 12 dB per octave at the cost of more circuitry. Another advantage to using two-pole filtering (both attack and release) is the shaping of the transitions (the “knees”) tends to be much smoother and less abrupt. The attack and release can be adjusted to be faster while still allowing the compression process to sound good. Unfortunately, unless the manufacturer “brags” about the method used, there really isn’t any easy way to tell how the circuitry was designed. The bottom line remains the same – How does it sound? Keep in mind, how it sounds is the *only thing that will remain in your recording*. All the fancy lights, expensive extruded front panel, stainless steel cabinet, etc won’t be in the credits.

What to look for when buying a compressor

A compressor is enough of an investment that, like a new car, it should be “test driven.”. Look at the front panel. The function of each knob and switch should be clear and understandable. If the salesman has trouble explaining it so that the function makes sense, consider another brand or another salesman. Turn the knobs, feel the “action.” It should be smooth, not “gritty.” While this may seem silly for some, it is a good indication of how much effort the builder puts into buying quality parts. Listen to it carefully with program material similar to what will be used in the buyer’s situation. Adjust for minimum attack, maximum ratio, mid-range release (perhaps 0.5 sec). Adjust the threshold until the gain reduction meter shows 10 dB or so of gain reduction with no

clipping during the loud music. No gain reduction meter or clip indicator? This is a good time to look at another brand. There should be at least a threshold light to help set the control levels. Once the unit is set up, listen very carefully to the transition from quiet material to the loud transients, such as drumbeats. There should be nothing added by the compressor except decreased dynamic range. If it sounds worse when compressing, look elsewhere. Understand that the rather extreme compression that was just set up will sound somewhat “squashed”, but it shouldn’t sound “clicky” or “nasty.” The click sound usually comes from the attack not being shaped properly, while the nasty, distorted sound is simply the result of a poor compressor. This is the area that the buyer should try to get the best possible performance for the money being spent. Once this part is satisfied, the buyer can look for attractive graphics, lots of lights and extra functions. Be careful about paying extra for a function that isn’t needed. Although, one extra function that is almost imperative is a limiter with separate controls. Often, a downward expander is also included. Depending on the intended use, this may be worthwhile.

A LIMITER

What makes a good limiter? One that is used infrequently but effectively. As will be shown, an effective limiter has constraints that preclude it sounding good. That may seem like a contradiction, but it really isn’t. That’s not to say that there aren’t good limiters – there certainly are. There just aren’t *good* limiters that make good *compressors*.

A Rane DC 24 is an example of a good feature set. It has two channels of compression, limiting, and downward expansion.

If we make it a condition that the limiter cannot pass any signal that exceeds some threshold, then we automatically impose certain restrictions as follows:

To prevent transient overload, it must have an attack time faster than the fastest rise time expected. See Figures 1 & 2. The fastest rise time may be defined as one divided by the highest system frequency. This frequency is estimated as four times the -3 dB rolloff of the system for a single processing device. “Four times” is an arbitrary and conservative constraint. It assumes that the highest frequency will be 12 dB lower in amplitude than the program material due to the 6 dB per octave rolloff attributable to the high frequency limitation. The 6 dB figure is correct only if the system rolloff is defined by a single filter. If several devices are in sequence in the audio path and they have similar rolloffs, the maximum frequency can be downgraded to two times the rolloff of the system. For the sake of discussion, let’s assume there are two devices in the path and they each rolloff around 50 kHz at 6 dB per octave each. We estimate the highest frequency of interest at 100 kHz. Therefore, the attack time required is $1 / 100 \text{ kHz}$, or 10 μs . An attack time this short represents a possible, but challenging goal.

A more common approach is to allow very short transients to pass through and be clipped by the system. Generally speaking this isn’t too objectionable. In fact, it probably sounds better than allowing the attack time to be as fast as 10 μs , for reasons to be explained. A common compromise

seems to be 100 μ s to 1ms for the attack time. Slower than 1 ms and the clipped peaks become wide enough to be quite audible. To avoid second-order effects, the fast attack transfer function must have “gentle transitions” known commonly as “soft knees.” All this means is that as the change is made from no limiting to full limiting, the transition be performed smoothly rather than abruptly. If this isn’t done, the engaging of the limiter introduces an audible click. This artifact is one of the main reasons operators “dial in” a relatively slow attack time in limiters with this parameter adjustable. *The operator finds the occasional clipping preferable to the click.* Consequently, if the click can be minimized, a faster attack time can be used. As the attack time becomes faster, this click becomes more difficult to avoid. Discussion abounds on the ideal shaping for this transfer function. Most common is an RC (exponential) charge with some diode action to “soften” the start of the limiting.

The limiter must have enough control ratio to guarantee the signal can’t clip once it has begun to decrease system gain. Normally, the ratio of a limiter ranges from around 20:1 to “infinity.” With an infinite ratio, once the input signal exceeds the limiter’s threshold, the output is not allowed to increase at all. A 20:1 limiter is easily attainable in a “feedforward limiter” while a “feedback limiter” is required for an infinite ratio. A *feedforward* limiter is one, which measures the input level, then calculates the gain applied by other functions to obtain the assumed output level. It sounds complicated, but it isn’t that bad. A *feedback* limiter simply measures the output level then generates whatever control voltage is necessary to prevent that level from exceeding the threshold. Why use one over the other? In a system where other dynamic gain control is used, such as a compressor, there is already a voltage available that indicates the level in dB of the input signal. In a feedback limiter, the rectification and logging (if any) must be added at the output. Also, all limiters add distortion to the audio. Some of the reasons have already been discussed. Another source of distortion is the ripple remaining on the control voltage. This ripple is inversely proportional to the release time and the signal frequency. If a half-wave rectifier is used (common if applied at the output), the ripple is at the signal frequency and double the amplitude compared to that left by full-wave rectification. With full-wave rectification, the ripple is twice the signal frequency. Also, if the rectified voltage is converted to a log function, it leaves less ripple than a simple, rectified voltage. This ripple causes distortion in the VCA as it did in the compressor, only worse. A serious problem with the feedback limiter is that the distortion caused by the limiter is already present in the signal being measured. In other words, the limiter reprocesses the distortion along with the signal and adds increasingly complex distortion. Fortunately, this “second-pass” distortion is much reduced in level. The ripple voltage may be minimized by use of a multiple-order release filter. However, as the order of the filter increases, it becomes more difficult to obtain reasonably short attack times. A second-order filter is a workable compromise. All this means is that the first release filter is followed by a second nearly identical filter. This reduces the ripple by 6 dB (by one half) at a given signal frequency compared to the standard first-

order release filter at the expense of additional circuitry. Nearly all limiters use a single-order filter for this function.

Another parameter is the release time. If too short, the action of the limiter sounds choppy; if too long, the limiter “breaths”. That is, after the offending overload, the limiter holds the system gain down too long. As the gain is brought back to normal, the increase in level is audible – not very pleasant. Also, as mentioned above, as the release time is shortened, the distortion increases. For a limiter operating with a release time of perhaps 500 ms, at 500 Hz this distortion easily exceeds several percent THD. It is quite audible. This may be improved by using a second-order release filter as mentioned, but ultimately the best fix is to use the limiter as little as possible.

A serious limiter *must* use full-wave rectification to measure the signal level. Let’s assume in a half-wave system that the positive peaks are measured. It is entirely possible to have a major overload on the negative half of the signal only. In such a case, the overload wouldn’t be detected and there wouldn’t be any limiting. Full wave rectification adds so little to the cost that the choice is trivial. Unfortunately, the buyer has no easy means to determine the method used. If the ratio is mentioned and it’s greater than about 20:1, it is *probably* a feedback limiter. A feedback limiter “tacked on” to another function (usually a compressor) is *not likely* to sound good if used heavily.

Normally the limiter is part of some other function, such as a compressor, as in the Rane DC 24. A limiter by itself is difficult to evaluate. It should be listened to with a compressor ahead of the limiter in the sound processing. With the compressor set up and functioning normally, have the clip light come on *occasionally*. Adjust the threshold of the limiter to just light on the loudest music peaks, and then only briefly. The clip light *should not come on except very rarely*. The limiter should be considered “disaster prevention.” With it operating as described, its action should be quite unobtrusive. If it can be heard working, it probably is either not set up properly or is a bad limiter. It is the author’s opinion that there are a lot more bad limiters than good ones. Fortunately, if used carefully, even a bad one won’t be a total disaster. However, it *must be fast enough* to avoid clipping on all but the most difficult program material. If it can’t do this properly, it’s no better than a really bad compressor.

A DOWNWARD EXPANDER

The function of a good downward expander is to increase the *apparent* dynamic range of the system by decreasing the gain during the relatively quiet times thereby moving the apparent noise floor downward. It does this by comparing the signal level to a threshold. When the signal level drops below this threshold, the expander decreases the system gain by some ratio. This ratio is the same as defined in the compressor, only opposite in sense. The difference between this ratio and that of the compressor is that the gain of the expander decreases when the input signal is *below* the threshold. Expanders usually operate with a rather gentle ratio – perhaps 1.5:1 to 2.5:1. Occasionally higher ratios are used, but the audio tends to take on an odd breathing sound as the background noise comes and goes due to the expander’s action.

To be effective, the dynamics of the expander must have moderately fast attack and fairly slow release times. The attack is the time spent restoring the system gain to normal when the signal level increases. Times of 0.5 ms to 1.5ms work well. The release is the time spent reducing system gain when the signal decreases. For the release, a time of 0.5 sec to 1.5 sec is common. The waveform of the control voltage must avoid sudden changes. Transitions must happen smoothly with a decidedly “soft knee.” This is especially important with the downward expander since the audio will be, by definition, rather quiet at the threshold. Without audio dynamics to mask the gain adjustment, it must be done smoothly and gently or the action adds objectionable artifacts to the sound. To recap, the five determinates of the dynamics of a good downward expander are:

Threshold. The signal level below which the expander begins to reduce the system gain. This parameter is almost always adjustable and typically ranges from -50dBu to 0 dBu.

Ratio. The rate at which the gain reduces expressed as a ratio of output level change to input level change, such as 2:1. The normal range of adjustment is 1:1 (bypass) to 5:1 or so.

Release Time. The time it takes the expander to reduce system gain after the signal has dropped below the threshold level. If this parameter is adjustable, it is set to fit the dynamics of the particular audio source. Speech typically requires a fairly short release time while music sounds better with a longer release time. Adjustment range should be from 100 ms to 3 sec or more.

Attack Time. The time it takes the expander to restore the system gain after the signal level rises above threshold. A variety of methods exist to define this time, but it seems to be less of a problem than with a Gate, Limiter, or Compressor. The most descriptive is to declare the time it takes the expander to increase the gain to within some percentage of the final level after the signal is suddenly restored. This time may depend on how much the input level exceeds the threshold. Attack time is usually fixed. If it is adjustable, it should range from 0.5 ms to 10 ms. If it is fixed, an attack time from 1 to 5 ms would be appropriate.

Depth of gain reduction. This refers to the limit of the gain reduction caused by the expander. As a practical matter, it's usually limited to perhaps 60 dB of gain reduction. Sometimes the amount of gain reduction is adjustable to keep the system from sounding “dead.” If the depth is fairly modest, this becomes a “ducking” function. If the system gain is decreased to a very low level, the constraints placed on the control dynamics are relaxed. For example, if the system is not passing audio due to heavy downward expansion, the transition at the beginning of the attack time is not so critical. However, the apparent attack time appears longer because of the time it takes the control to transit through the relatively inaudible region.

What to look for in an expander

Normally an expander is packaged as part of some other processing function, such as a compressor or a gate. A good listening test for an expander is to set the threshold *higher* than the threshold of the companion compressor. This would *not* be a normal setting, but the overall control function

should still be smooth and without “odd anomalies.” All this means is that it shouldn't “go wacky” even though the controls are set in a non-standard way. When operating by itself, it should be possible to set the controls so that the expander doesn't breath excessively nor sound choppy. An expander is particularly effective with spoken material. Subtle control should normally be the goal – avoid extremes. If the ratio is fixed, make sure it is gentle, perhaps 1.5:1 or so. Make sure there is a way to defeat the expander if the program material doesn't require it – a real bypass function is better than a very low threshold. If a mic is available for the listening test, set the threshold to substantially decrease the gain with no input to the mic. While listening with earphones, say something like “putt, putt, putt” into the mic. Make sure the expander “comes alive” quickly enough to include all of the “p” sound. If the release is adjustable, it should be possible to set it such that the gain decrease after the input ends is not overly noticeable, yet the softer “t” sound isn't chopped off. If used properly, the expander should do a good job of cleaning up a sound track during the quiet times without making it sound like the background noise is “doing pushups.”

A GATE

An excellent gate is a very challenging balance of constraints. A gate is a device used to pass or not pass audio, based on the signal level. Refer to Figure 3. If the level exceeds an established threshold, the audio is passed (the gate is “open”). If the level is below the threshold, the audio is blocked (the gate is “closed”). It's never “partially open” except during transitional times. A deceptively simple definition. The easy approach is to use the equivalent of a simple switch – either on or off. Indeed, some are done that way. Unfortunately, they sound terrible.

By definition, sound is present when the switch is “thrown.” The rapid change in level causes a serious “click” to be added to the audio. This happens both when the gate is opened and when closed. Various tricks are added to avoid these artifacts and that is where the design challenge happens. One method, which has some acceptance is to only open or close the gate when the actual signal waveform crosses zero volts (remember, it's alternating voltage). If the VCA is a reasonably good one, no output signal is produced at that instant. The problem is that most waveforms resulting from a percussion instrument have a very large initial transition followed by smaller “tremors.” By the time the threshold is exceeded and the next zero crossing happens, the main event is over. The loud drum “rim shot” is reduced to a dull thud, for example.

Another approach is to use a downward expander with a very high ratio and a very fast attack. The problem with this method is that almost no downward expander is capable of the dramatic attack times required. The attack time is defined as the time required for the gate to open, or pass signal, after the signal level exceeds the established threshold. To be effective, this time must be less than 100 us. One gate measured, advertises 5 us, although in actual tests it didn't do anything before 100 us, and it took an additional 50 us to actually control the signal! Another gate has been measured at 20 us




Figure 3. Typical Gate Dynamics

In this example, what we want isn't going to sound like what we may get! First, part of the drum hit is lost. Second, the end of the tail is attenuated too soon. Third, unwanted sound is picked up. In a really good gate, the drum will sound exactly the same as the original, only the unwanted sound will be gone. Incidentally, part of the solution to this problem may be to reposition the microphone to remove some of the unwanted sound. No gate can perform magic, though with the right settings it can seem like it.

(start of +20 dBu burst of 20 kHz to *beginning* change of control voltage). In addition, this gate uses an analog delay to “hold back” the audio path to give the gate time to open. It appears to open before the increase in signal level happens. These are rather expensive but effective tricks just to make a “switch” sound good.

The following parameters are required in a superior gate:

Threshold. The signal level above which the gate opens to pass audio. The operator almost always adjusts this parameter. A normal control range is from -50 dBu to +20 dBu. A threshold setting *below* the noise floor becomes a bypass –gate always open.

Attack Time. The time from when the signal exceeds the threshold to when the gate opens. Unfortunately, it is often defined as when the gate just begins to open in order to “improve the spec’s.” This is probably the first specification a prospective buyer should look at when evaluating a gate. When an attack time of less than 100 μ s is specified, be suspicious of how it was measured. This parameter is sometimes adjustable, often by a 2-position switch (fast and faster). If adjusted by potentiometer, it needs to be set fast enough to include all of the percussive rise time of the audio. If it won’t open that fast, go buy a better gate. The attack time should be adjustable, or fixed, to 500 μ s or *less*. A reasonable maximum is 1.5 ms.

Release Time. The time from when the signal drops below the threshold to when the gate closes. This parameter is usually adjustable. It is set to include the “tail” of the sound before the switch completely closes. For a kick drum, this time is fairly short, while for a triangle it may be as long as 4 seconds. The range of adjustment should be from 100 ms to 4 sec or more.

The following parameters are optional:

Hold Time. Companion to the release control. The time from when the signal drops below the threshold to when the gate just *begins* to close. If present, this parameter is normally adjustable by the operator and ranges from 10 ms to 4 seconds. This parameter is used to preserve the dynamics of the “tail” of a sound, such as a cymbal or bell. If used thoughtfully with the release time the over all quality of the transition is much improved.

Depth of Cut. The amount of gain reduction when the gate is “switched off.” Since a VCA is usually used for the actual gain control element, a total switch function is not practical. Also, in some situations, total gain reduction is not wanted. If total silence isn’t necessary while the gate is closed, the attack time is improved by limiting the gain reduction. Therefore, this parameter is sometimes adjustable. It should range from 90 dB reduction to 0 dB (effective bypass). A depth of 60 dB to 40 dB would be typical.




Figure 4. Anatomy of a Good Gate

The details may vary, features may be added or left out, but this is a reasonable starting point. The Side Chain switch (Ext./Int.) is actually a function of the Side Chain jack—that is, plug into the jack and it automatically switches to External.

High-Pass and Low-Pass Filters. These are applied in the control signal path (referred to as the “side chain”) **not** the actual audio signal path. Usually there is a switch called a “key listen switch” that allows listening to the control signal to judge the affect of the filtration. Often there is a back panel jack called the “side chain jack” that allows this control signal path to be interrupted and modified externally then returned to control the gating. Any delay in this signal path degrades the attack time of the gate, so filters should be used sparingly since **all** filters introduce delay (we’ll save the semantics of “lead” and “lag” for another discussion).

An analog delay in the audio path. If used, it is essential that the curve of the delay-versus-frequency be smooth and free of sudden changes. Ideally, it would seem that a constant delay would be best, however, a delay that varies **smoothly** with frequency works well. In any case, its function is to delay the audio path slightly to give the gate time to react before too much of the signal transition has occurred. This delay should range from 10 μ s to 100 μ s or more. If it is too long it will cause other problems. This delay is difficult to evaluate in testing. The only practical approach is to measure phase shift at a rather low frequency. Incidentally, this delay thoroughly “trashes” a square wave, or any other complex wave when viewed on an oscilloscope, due to phase dispersal. Fortunately, the signal sounds unchanged, but it certainly doesn’t look unchanged.

Notice that there is no reference to “ratio.” For those who insist on thinking in terms of a downward expander, the ratio of a gate is infinite. That is, above the threshold it’s on and below the threshold it’s off. Nothing in between other than transitional shaping.

Note that, because of the very fast attack time required, shaping of the attack function is very important. This alone will make or ruin an otherwise good gate. The real secret is a very fast turn on done carefully and smoothly – not a trivial task. Most, but not all, gates that are fast enough to be effective have a serious “click” at the turn on-transition.

HOW TO “KICK THE TIRES” OF A GATE

The best way is to play a drum track recording which has picked up low levels of some other instrument. The drum should be a snappy snare, preferably with lots of “rim shots.” The sound to look for is a sharp percussive attack with some trailing sound. With the audio signal passing through the gate, adjust the threshold to trigger on the drum but not quite trigger on the background sound. Set the attack time for the minimum. There should be no “click” added even at minimum attack time. Listen very carefully to the drum sound while enabling and disabling the gate action. The sound of the drum *should not change* while the background sound between drum hits should be gone. Adjust the release and hold controls for the best result. This is the essential purpose of a gate. If it can’t do it well, find a gate that will. Take heart – good gates do exist.

In summation, let it be said – buyer, be careful. Don’t expect top performance if you’re shopping for the lowest price. By the same token, a high price doesn’t ensure the best sound. If you don’t need “gimmicks”, don’t be talked into them. The beautifully done, extruded front panel may look really sharp, but it doesn’t make a recording sound one bit better. If the cabinetry is tinny and cheap looking, the circuit design *may* also be minimal and of poor quality, but this isn’t necessarily true. Review articles are a good starting point, but be careful of the reviewer’s personal prejudices – they may not coincide with those of the buyer. Listen to the opinions of colleagues, again with a view to personal prejudice. When all is said and done – listen to the equipment. Make the test as representative of the buyer’s actual situation as possible. Listen carefully, be critical. If something sounded funny, don’t necessarily believe the salesman’s claim that there’s “something wrong with the recording.” If necessary, the buyer should provide a recording for the listening. Make sure the equipment can be returned if it doesn’t work out. After all, the ultimate test is how the potential hit recording sounds. Will it go platinum or just be another “not bad” recording?

Product	Rack Spaces	Inputs	Outputs	Hardware	Power Switch	Power Supply	Phantom Power
---------	-------------	--------	---------	----------	--------------	--------------	---------------



Headphone Amplifiers

HC 6

1U

(2) Master Line
(6) Stereo Line

(6) Stereo Phone



-

**HC 4
MH 4**

1U

(2) Master Line

(4) Stereo Phone



-

Compressors

DC 24

1U

(2) Line

(2) Line



-

VP 12

1U

(1) Mic
(1) Line

(2) Line



-



DC 22

1U

(2) Line

(2) Line



-

Crossovers

AC 22

1U

(2) Line

Stereo 2-Way
Mono 3-Way



-

AC 22B

1U

(2) Line

Stereo 2-Way
Mono 3-Way



-

AC 23

1U

(2) Line

Stereo 3-Way
Mono 4/5-Way



-

AC 23B

1U

(2) Line

Stereo 3-Way
Mono 4/5-Way



-

DC 24

1U

(1) Line

Mono 2-Way



-

**SAC 22
MX 22**

1U

(2) Line

Stereo 2-Way



-

**SAC 23
MX 23**

1U

(2) Line

Stereo 3-Way



-

**RPM 26v
RPM 26i**

1U

RPM 26v = (2) Line
RPM 26i = AES3
digital input

Stereo 2-Way
Stereo 3-Way
Mono up to 6-Way



-



-


All Rane products have full agency compliance (UL, cUL, & CE).


All Graphic Equalizers are Constant-Q design.

Most Outputs are balanced with the exception of RCA jacks and ¼" mixer sends.

Refer to individual Data Sheets for specific and additional product details.

Product	Rack Spaces	Inputs	Outputs	Hardware	Power Switch	Power Supply	Phantom Power
DJ Mixers							
TTM 52i	-	(1) Mic, (2) Line/Phono	Main Outs	 	-		-
TTM 54i	-	(1) Mic, (2) Line/Phono, (1) Stereo Loop	Loop and Main Outs	 	-		-
TTM 56	-	(1) Mic, (2) Line, (2) Phono, (1) Stereo Loop	Aux, Loop & Main Outs	 	-		-
MP 2	1U	(1) Mic, (4) Stereo Line (2 Phono)	Tape, Aux & Main Outs	 	-		-
MM 8z	4U	(2) Mic, (4) Stereo Line/Phono, (1) Stereo Loop	Tape, Loop & Main Outs	  	-		-
MP 22z	4U	(2) Mic, (7) Stereo Line (2 Phono), (1) Stereo Loop	Tape, Loop, (2) Zones & Main Outs	  	-		-
MP 24z	4U	(2) Mic, (9) Stereo Line (3 Ph), (2) Stereo Loops	Tape, (2) Loops, Zone, Booth, Light & Main Outs	  	-		
MP 44	5U	(2) Mic, (8) Stereo Line (4 Ph), (2) Stereo Loops	Tape, (2) Loops, Zone, Booth, Light & Main Outs	  	-		
MP 2016	3U	(2) Mic, (9) Stereo Line (4 Ph), (1) Stereo Loop	Tape, Loop, Booth & (2) Main Outs	  	-		-
Mic/Line Mixers and Mic Preamplifiers					also see Paging Processors		
DA 216a	1U	(2) Mic / Line	(16) Line	         			

Product	Rack Spaces	Channels	Sliders/ Control	Hardware	Power Switch	Power Supply	Special Features
RANE Equalizers: 1/3-Octave Graphic							
ME 30B	1U	1 Ch	20 mm				
ME 60	2U	2 Ch	20 mm				Interpolating
GE 130	2U	1 Ch	45 mm		-		External supply mounts rear panel
GE 60	3U	2 Ch	45 mm				Interpolating
GE 30	2U	1 Ch	60 mm				Line-out transformer; Cut-only mode; Interpolating
RPE 228d	1U	2 Ch			-		Interpolating
SEQ 30S MQ 302S	1U	2 Ch	20 mm				Ganged Stereo Sliders
SEQ 30L MQ 302L	2U	2 Ch	45 mm				Ganged Stereo Sliders
Equalizers: 2/3-Octave Graphic							
ME 15B	1U	2 Ch	20 mm				
GE 215	2U	2 Ch	45 mm		-		External supply mounts rear panel
Equalizers: Parametric							
PE 15	1U	1 Ch					5 bands
PE 17	1U	1 Ch			-		5 bands +2 cut filters
Equalizers: Combination							
THX 44	2U	4 Ch	20 mm		-		(3) 11-band 1/3-octave (3) 2-band parametric (1) 2-band sub parametric
THX 22	1U	2 Ch	20 mm		-		(2) 11-band 1/3-octave (2) 2-band parametric
SSE 35	1U	3 In 3-5 Out	20 mm		-		(3) 11-band 2/3-octave (2) 7-band 1/3-octave

Product	Rack Spaces	Inputs	Outputs	Hardware	Power Switch	Power Supply	Phantom Power
RANE Distribution Amplifiers							
DA 216a	1U	(2) Mic / Line	(16) Line		-		 15V
SM 26B	1U	(2) Master Line (6) Line	(2) Master Line (6) Line				-
SRM 66	1U	(6) Line	(6) Line		-		-
Paging Processors							
CP 31	1U	(1) Mic (1) Stereo Line	(2) Main Line (1) Expand Line		-		-
CP 52	1U	(1) Mic (4) Stereo Line	(2) Main Line (1) Expand Line		-		 15V
CP 64	1U	(2) Mic (4) Stereo Line	(2) Main Line (2) Expand Line (1) Zone Line		-		 15V
DA 26	1U	(1) Mic (priority) (1) Mono Line	(6) Line		-		-
Network Audio							
NM 84	2U	(8) Mic / Line	(4) Line (8) Mic/Line (8) Direct 100Base-T	 100Base-T  RS-232	-		 48V
NM 48	2U	(4) Mic / Line 100Base-T	(4) Line (8) Mic/Line (8) Direct	 100Base-T  RS-232	-		-
RaneWare Products							
RPE 228d	1U	(2) Line	(2) Line	 RW 232	-		-
RPM 26v RPM 26i	1U	RPM 26v = (2) Line RPM 26i = AES3 digital input	(6) Line	 AES3 26i only  RW 232	-		-
RPD 1	1U	-	-	 modem  RW 232	-		-
ECM 64e	1U	(4) Mic / Line (2) Line	(3) Line	 RW 232	-		 12V
ECB 62e	1U	(6) Port Line (6 stereo w/ECS 62) (6) Mixer Line	(6) Port Line	 RW 232	Expand with ECM 82e		-
ECM 82e	1U	(8) Mic / Line	(2) Line (8) Post (8) Direct	 RW 232	Requires ECB 62e		 12V

Emerging Standards for Networked Audio System Control

- **Hardware**
- **Network Planning & Design**
- **Software**
- **ActiveX**
- **Necessary Skills for Tomorrow's Installations**

Devin Cook
Rane Corporation

Ethernet Links

John's Closet

<http://www.digitalmx.com/wires/index.html>

A down to earth set of networking how-to's

Ethernet Tutorial

<http://www.lantronix.com/htmlfiles/mrktg/cataloget.htm>

A wonderful tutorial in plain English from Lantronix

Macmillan's Personal Bookshelf

<http://www.mcp.com/personal/>

Free online books about programming and the internet

Network Design Tutorials & Other Resources

<http://www.alaska.net/~research/Net/nwpages.htm>

An industrial strength link list of networking topics

RaneNote 144
© 1998 Syn-Aud-Con

Digital devices must be able to co-exist at the hardware, protocol and application level in order to create a seamless network of digitally controlled audio devices. Slowly these issues are being resolved so that we can soon look forward to larger, better integrated installations.

Hardware

Up to now, most audio manufactures have relied on either RS-232 or RS-485 to communicate with their units. This will continue for some time because these interfaces are easy to implement, with support built into most common low-cost microcontrollers designed into systems today. While adhering to electrical standards, both have their drawbacks.

RS-232 is common on computers, but can't be easily connected to multiple devices without some major kludging. RS-485 can more readily talk to multiple devices, but some sort of box must be used to convert the RS-232 to RS-485. Though you may be able to electrically interface to more than one device, there is no generally accepted protocol to talk to devices from different manufacturers. Even if you could overcome these limitations, speed starts becoming a problem as control for these devices becomes more complicated.

Ethernet has emerged as a viable electrical interface to solve these problems (along with all those I neglected to mention). In the past, hardware cost and software complication has kept this technology out of the hands of most small audio manufacturers. Now, microcontrollers are emerging with Ethernet controllers built-in, and the cost of external Ethernet chipsets has fallen to the point where it is practical for smaller companies to justify including the interface in their devices.

One pitfall of Ethernet is that the technology is a bit of a moving target. Just as 10Base-T (10 Megabits per second) Ethernet is becoming cheap enough for practical control applications, the computer industry is converting over to 100Base-T (100 Megabits per second), with work on a Gigabit standard currently being finalized. 10Base-T is more than adequate for control and monitoring, but digital audio distribution requires 100Base-T to be practical, and that's still a bit expensive.

Network Planning & Design

Surprising to many is the fact that not only can you connect devices together and control them from your computer via its NIC (Network Interface Card), you can control devices from *any* computer on the network. This opens up the possibilities for remote monitoring and control (and I haven't even mentioned control over the Internet yet!). If simple control and monitoring is all that is required, these devices may easily be incorporated into an existing data network, lowering the hardware and wiring costs. If you decide to take the big plunge and distribute audio on the network, along with control, a separate network is probably a must.

Software

Now that an Ethernet hardware standard appears to be emerging, software remains a barrier to creating a truly integrated network. There already exists many excellent applications available from different manufactures to control a single device, or family of devices, but running a separate program for each device is nobody's idea of a "system", let alone providing the end-user with a simple network interface.

Though efforts to create an interoperable environment in the past have been disappointing, MediaLink being one notable disaster, they have pointed us towards the control environment of the future. To take full advantage of a network, control software must provide the following services:

- Control of ALL devices on the network
- Customizable interface to simplify user operation
- Complex device interactions to allow the above

ActiveX

The most promising solution to the application interface problem appears to be ActiveX, Microsoft's evolution of OLE, into an Internet friendly programming model. ActiveX components can be used in web pages, visual programming languages such as Delphi, Power Builder and Visual Basic, and even tools such as Lab View. Each ActiveX control is made up of "Properties", values associated with the control which might include things as level settings and meter readings, and "Events", which tell the computer something significant has happened, such as a switch closer or clip detection.

ActiveX allows the manufacturer to create an object which fully encapsulates (describes) a device, while hiding the implementation details such as protocol from the programmer. For example, no longer would you need to know that "byte 5" of a 24 byte status message meant that the unit was limiting. With an ActiveX control, you might simply refer to the limiting status as "Device1.LimitStatus".

By hiding all the nitty-gritty communication details, there is no need for different manufacturers devices to *agree* on protocol. The lack of a protocol standard means that no cooperation between manufacturers is required. This is a good thing in light of the progress, or lack there of, on a common protocol such as AES-24. It allows manufactures to choose the best protocol for their device, just as you might choose JPEG versus GIF images for your web page, trading off size and speed versus quality. As a manufacturer I may choose a custom lightweight protocol versus a larger more standard protocol depending on the power of my device.

You might also add control of devices from outside the industry, such as lighting, as long as an ActiveX module is provided for you. Recently, the ESTA (Entertainment Services & Technology Association) has just published a standard title "Recommended Practice for Ethernet Cabling Systems in Entertainment Lighting Applications" and there are quite a few Ethernet to DMX512 (an ESTA standard control protocol for lighting equipment) bridges on the

market. Also available are Ethernet bridges for RS-232 and RS-485, although I haven't spotted any specific ActiveX controls available yet—but many bridges do use standard internet protocols for which there are controls available, such as telnet (the standard terminal protocol).

By linking various controls yourself and adding a bit of programming you can do things you can't imagine with just simple control software. For example, let's say you want your system to automatically turn down the output level of your speaker processor box whenever the amplifier clips. This might mean trying to get two devices from two different software interfaces to cooperate. This is not an easy task, but if each device was represented by an ActiveX control in a web page perhaps, a VBScript (Visual Basic Script) to control it might look something like:

```
Sub AMP1_OnClip()  
if DSP1.OutLevel >0 then DSP1.OutLevel = DSP1.OutLevel - 1  
end sub
```

This script is executed every time the AMP1 control generates a "Clipping" event. The script checks to see if the DSP box (DSP1) has its output level turned up (above 0) and if so, turns it down one step.

The example is simple, but illustrates staggering possibilities. Add a couple of buttons, graphics and text and you've got a customized installation interface tying together devices from separate manufacturers.

Necessary Skills for Tomorrow's Installations

Building a network is not nearly as complicated as some would have you think. Start by getting some experience building a small home or office network. The cabling is simple, and as long as you don't need the equipment to talk to the Internet or a corporate network, you won't need switches, routers or any of the heavy-duty network equipment (which must be designed in carefully). A small hub will do, and this is not even necessary if you're simply trying to connect two Ethernet interfaces together. Getting two computers to talk together via TCP/IP (the de-facto king of network protocols) should give you enough exposure to figure out just what is involved with networking. The supplied links are good launching points for more details.

To make the most of a network, programming skills will be required. These skills may range from simply being able to use web publishing software to create web pages and a few simple scripts to tie controls together as in the earlier example, to creating custom applications in Visual Basic, Delphi or C++. Again, there are easy ways to get your feet wet. Programs such as Microsoft's *FrontPage*, Adobe's *GoLive* or Macromedia's *Dreamweaver* are more than adequate for building web pages with ActiveX controls and VBScript.

Whatever level you wish to start, it is apparent that computers are not going away and that networks will be finding their way into new installations more and more often. It's time to get some exposure to the technology so your design options will grow *with* the technology.

Controlling Audio Systems with ActiveX

- **ActiveX**
- **Network Hardware**
- **CobraNet**
- **Ethernet**

**Stephen Macatee
Devin Cook
Rane Corporation**

**RaneNote 152
© 1999 Rane Corporation**

INTRODUCTION

Control of audio systems and transport of audio over computer networks are newsworthy topics in the professional audio industry. Several schemes are available today. Using ActiveX is not new to the audio industry, but controlling CobraNet devices with ActiveX is. The secondary problem of incorporating real world relays, switches, indicators and existing non-networked, serially-controlled audio and interface products needs addressing to transition into our future. Also, the ability to link on-screen controls to such relays, switches, indicators and equipment must be supported.

To implement such control over Ethernet networks, Microsoft ActiveX controls have been employed as a control implementation. A brief primer on ActiveX as well as ActiveX advantages and disadvantages are discussed. With years of work on the old topic of true cross-manufacturer interoperability still not generating the critical mass it once was, a view toward these topics is also revisited. This paper attempts to show that the nirvana of true manufacturer interoperability — regardless of protocol — is achievable over networks without the need for manufacturers to agree — except on Microsoft ActiveX.

BACKGROUND

Much has been written on the varying needs, implementations, benefits and pitfalls of system control. The Audio Engineering Society (AES) has held conferences and paper sessions on computer-controlled audio systems[1]. With the seemingly constant drop in price and availability of computers, it is easy to anticipate further and even ubiquitous incorporation of computers as the control platform of choice for professional audio systems. Thousands of systems already incorporate personal computers, Macintosh and other computers.

The natural evolution of audio and control over standard Ethernet computer networks already taking place furthers the need to enable disparate computer-controllable devices to share both a common user interface and the transport mechanism over Ethernet. Peak Audio's CobraNet technology offers one such solution which uses Ethernet for transport and SNMP (Single Network Management Protocol) for control. Others to watch for increased applicability in the pro audio world include the currently consumer and professional audio-deployed IEEE 1394 (FireWire) and USB. ATM and other technologies from the computer industry will continue to infiltrate the conservative pro audio world.

The complexity of audio systems continues to grow, driven by users' requirements and expectations and the constant march of technology. The need for untrained users to easily control such complex systems is perhaps the hardest thing to provide when equipment from multiple manufacturers is incorporated into a single system.

Control interfaces began with device front panel controls and evolved through the following: wired and wireless remotes, MIDI, RS-232, RS-422, RS-485, DMX512, etc. The list continues through Ethernet and many other proprietary and non-proprietary schemes. Yet, with all these control and several transport schemes to choose from, it is still difficult to control disparate products from a common user interface and have the results work cohesively.

The dream of a common hardware and software approach permitting cross-manufacturer interoperability has lost a lot of steam since the Lone Wolf days of the pro audio industry. Much is learned from examining where we've been; where we're going; how we got here; why we were going there in the first place and what mistakes were made along the way.

The Lone Wolf dream of linking all manufacturers' products together with a common user interface and transport mechanism created great excitement and anticipation. The idea was simple, but the implementation proved too difficult. The Lone Wolf dream failed due to sale of the technology before it worked well enough for the scaleable needs of the industry, as well as the sudden licensing of too many manufacturers before the resources to support them were in place. One thing is certain, Lone Wolf and SC-10 succeeded by providing the audio industry with a vision and the facility for cooperation toward that vision. It is now taking the technology of the computer industry to get us there though.

The need for a common communications and transport protocol applicable across disparate components from many manufacturers was achieved by the computer industry many

years ago. It is often discussed that many standards come about through the necessity of one or more entities in an industry having the initiative to solve the problem at hand. Then others recognize its significance and jump on the bandwagon thus creating the momentum for a future standard, and/or competitive alternate solutions are developed. As is happening in many industries, adopting techniques and technology from the long-standing and significantly larger computer industry with its enticing economies of scale and other advantages seems a worthwhile venture. This is where the ActiveX rubber meets the road.

ACTIVEX BACKGROUND

Microsoft ActiveX controls are of concern to the pro audio community. This technology allows designers of computer-controlled sound systems to create common front-end software control panels that operate different manufacturers' units, without having to know anything about their internal code or protocols. This is powerful. If manufacturers incorporate ActiveX controls, systems designers will not be limited to products offered by a single, platform-specific manufacturer.

Each ActiveX control is made up of Properties and Events. ActiveX control Properties are values associated with the control which include things such as level settings, mute condition and meter readings. ActiveX control Events tell the computer something has happened, such as a switch closure, button press or clip detection. ActiveX allows the manufacturer to create an object (a piece of software code) which fully describes a device, while hiding the implementation details such as protocol from the programmer. By hiding the communication details, there is no need for manufacturers to agree on protocol. This lack of a protocol standard means that cooperation between manufacturers is not required. It allows each manufacturer to choose the best protocol for their devices.

For example, no longer would you need to know that the 17th byte of a 32-byte status message meant that the unit's second output channel was muted. With an ActiveX control, you might simply refer to the device's output 2 mute status as "Device1.Out2Mute."

WHAT IS ACTIVEX ANYWAY?

ActiveX is a Microsoft-developed software technology released in 1996. ActiveX, formerly called OLE (Object Linking and Embedding), is based on the Component Object Model (COM), but provides substantially different services to developers. An ActiveX control is a unit of executable code, such as an .exe file, that follows the ActiveX specification for providing software objects. This technology allows programmers to assemble reusable software controls into applications and services. However, software development using ActiveX technology should not be confused with Object-Oriented Programming (OOP). OOP is concerned with creating objects, while ActiveX is concerned with making objects work together. Simply stated, ActiveX is a technology that lets a program – the ActiveX control – interact with other programs over a network, regardless of the language in which

they were written. ActiveX controls can do similar things as Java, but they are quite different. Java is a programming language, while ActiveX controls can be written in any language (e.g., Visual Basic, C, C++, even Java). Also, ActiveX runs in a variety of applications, while Java and Javascript usually run only in Web browsers. ActiveX controls can be used in web pages and within visual programming languages such as Borland's Delphi, Sybase's PowerBuilder, Microsoft's Visual Basic and even in tools such as Adobe's GoLive and National Instrument's LabVIEW.

For pro audio applications, objects are the sliders, buttons, indicators and other graphical screen entities. The objects have properties: slider position, slider range, on or off for buttons and indicators. Once the screen objects are chosen and placed, ActiveX controls can link the objects' properties to other ActiveX controls such as the device parameters inside an audio device. For example, linking a slider to the ActiveX control for a device's level control. Then, moving the level control graphic slider varies the audio level and vice-versa.

Here's an example: A computer is used to control an audio system over an Ethernet network and something on the computer's screen controls some function of the system. The idea is to place controls on the computer screen and link them, using ActiveX, to a parameter in the system. What's important is that only the controls required by the computer's end-user need be displayed. Additionally, more detailed interfaces including hidden or password-protected web pages can then be created to provide any level of system parameter access desirable — from complete system control, to a lone system power button, or anything in-between. No longer are systems limited to the number of security levels provided by a vendor's software, nor are you limited to controlling a single system parameter per screen control. For example, you can link multiple ActiveX controls to a single screen object, thus adjusting EQ level simultaneously with master level control and limiter threshold. You can also program actions when certain events occur, such as triggering audio playback or turning a system off at a certain time or adjusting delay time as the temperature changes.

You can control different parameters inside the same device from different computers on the network as well as control the same parameter from multiple computers. This is one of the major advantages of networks — multiple control locations automatically update when changes are made by any control location.

However, ActiveX controls are not limited to just Ethernet implementations. You can create RS-232 or MIDI capable ActiveX controls. For audio networking purposes, Ethernet-enabled ActiveX controls are discussed here. It is wise as a development architecture to separate the communications code from the other software development pieces to allow portability of ActiveX controls onto other transport mechanisms.

One of many popular software packages used to create user interface web pages for computer-controlled systems is Microsoft's FrontPage. Web pages may or may not be

accessible over the Internet. Using FrontPage or any of the many ActiveX-ready software packages, ActiveX provides, literally, an infinite number of programming possibilities.

The procedure for using ActiveX with FrontPage is simple: insert a manufacturer-supplied ActiveX control in your web page. Set the control's Properties — such as the IP address and name for the manufacturer's device. As needed, insert a button (or scroll bar, check box...) and set its Properties like the button's name, value and control range. Use Microsoft's Script Wizard that is included with FrontPage to link the button's Events to the device's internal ActiveX parameter(s) or link other Events and Actions as needed for your application.

NETWORKING RELAYS, SWITCHES AND INDICATORS

Many systems must incorporate relays, switches and luminous indicators for control, event or status indication. The Rane Via 10 Ethernet Bridge ships with ActiveX controls which allow its logic port states to be linked to a web page ActiveX control. The device's logic ports are capable of driving relays or LEDs, or reading switches, relays or any zero to 5 volt source. Thus, on-screen objects can be linked with these real world hardware tools to incorporate their use in networked systems. Applications for the zero to five volt input port include switches, but temperature or humidity sensors, potentiometers or any variable 5 volt source are suitable. Since this is a networked system, multiple locations can control and monitor the logic ports, and the on-screen status always remains linked to the hardware tools — change either and their states always correspond. This product is one way of implementing relays, et cetera into network controlled systems.

CONTROLLING EXISTING EQUIPMENT

Using existing RS-232-based and other serially-controlled equipment on a network is a likely requirement as the audio industry transitions into the network world. It may be some time before Ethernet control jacks are found on a wide variety of pro audio equipment. Until then, carrying RS-232 data over an Ethernet network can be achieved in several ways. Several vendors sell inexpensive RS-232-to-Ethernet adapters. The previously mentioned Rane product simultaneously supports both the logic ports and the proprietary RW 232 protocol of Rane products. Also, many CobraNet devices offer asynchronous transport of RS-232 data. You are still faced with RS-232's one-to-one Master-Slave issues with these devices. Some CobraNet devices now incorporate memories that are recallable from contact closures. These memory closures can also be transported over the network allowing a single switch to recall memories in multiple, similarly-equipped CobraNet devices.

Logic ports on existing devices can drive or be driven from relays that are monitored by the network. Or, combining existing ports with those on the previously mentioned Rane device's logic ports, one can easily incorporate non-networked products into networks. Applications for these input and output logic ports include memory changes and/or system monitoring.

ACTIVEX EXPERIENCES

ActiveX control of system parameters provides several advantages:

World Unity

ActiveX unifies divergent protocols from multiple manufacturers without the need for cooperation. ActiveX can significantly widen product choices for system designers and end-users, providing economies of scale, cost, availability, flexibility and serviceability.

Manufacturers need only develop ActiveX implementations for parameters requiring network control. Though, easier said than done, it's a small price given the advantages and direction of the pro audio industry. For some applications and products, ActiveX may eliminate the need for a manufacturer to implement a customized software user interface for a product or products.

Customization and Unlimited Security Levels

ActiveX permits easy customization of user interfaces and access levels, thereby avoiding constraints to the security levels offered by individual products or manufacturers. Complete interface customization, specific for the system end-user or the installer can easily be incorporated. The use of full color graphics with photos of the system controls or its components can be placed and scaled on a computer screen. The common password-protected web pages and firewalls of the computer world offer the security necessary for any system.

Third Party Software Support & Availability

ActiveX is supported by many software vendors who offer third-party education and support avenues. Sources of information abound for those who want to dive into this technology. Development of ActiveX controls is possible using any of the previously mentioned tools. With a variety of ActiveX capable packages available, needs from the simple to the advanced are satisfied.

Microsoft's Visual C++ includes a library called Active Template Library (ATL) which is designed to allow programmers to avoid the inefficiencies of other less efficient approaches to ActiveX. While ATL does allow development of ActiveX files that can be more size and time efficient, the drawback to this approach is the added difficulty in debugging the code.

Same Business Model as AMX & Crestron

ActiveX offers the same business model to the audio industry as room controller software providers. Many sound contractors either have in-house programmers for such software development or this service is contracted out. Additionally, these same sound contractors may already have a web master on staff that is capable of ActiveX implementations for sound system control.

Those who have followed the trends in room controller technology already recognize the industry's acceptance of

Ethernet and IP-based implementations. Thus, it is a simple step for room control developers to support ActiveX controls. Plus, their development tools are already web-publishing based, so current room control developers should have a small learning curve in the ActiveX world.

Therefore, for network systems with one or more PCs as the main controller, instead of assigning AMX or Crestron programmers the system control programming tasks, you find a web site provider who can implement the system control front-ends.

The success or failure of this new relationship remains the same: the system's control implementation relies on the system designer's ability to communicate the needs to the software provider. This is no different than the current room controller business model.

Non-proprietary and Ubiquitous

ActiveX controls are widespread and non-proprietary. Web page creators have been utilizing ActiveX controls for years, thus providing a large and knowledgeable group of ActiveX-familiar providers. ActiveX also provides the ability for these providers to create their own customized ActiveX controls. The color, apparent texture, size, shape, shading, look and feel of each control can be created completely from scratch and designed to match any on-screen décor imaginable.

Fits well into existing PCs

ActiveX is easily incorporated into existing PCs running Windows and Internet Explorer avoiding the need to create a new technology, protocol or hardware-based solution.

Multiple, Simultaneous Control Locations

While not an ActiveX advantage, the combination of network technologies and protocols permits incorporation of multiple simultaneous control locations. This was a previously difficult task in the control industry because serial streams such as RS-232 and RS-485 are unable to share multiple masters. With web-based architectures like ActiveX, controls on multiple web pages offer simultaneous control of the same system parameter(s) from many locations. For the CobraNet ActiveX controls implemented, Simple Network Management Protocol (SNMP) was used to provide the control capabilities found in computer network management, control and diagnostics tools such as Hewlett-Packard's OpenView. The case where changes made from a device's front panel are reflected in multiple network control screens is also satisfied. This is one of the major advantages of networks – multiple control locations can automatically be updated when changes are made by any control location.

Another advantage to network technology is its inherent redundancy ability. With careful network design, both the control and audio transport paths offer automatic fault redundancy.

ActiveX controls also have limitations.

Signal Indication

Perhaps the most important need within controlled systems is the need for accurate and trustworthy on-screen signal indicators. ActiveX control signal indicators have yet to be implemented within a CobraNet system.

Timing

ActiveX has timing limitations, particularly as the number of controls on a single web page are increased. As controls are added to a page, more RAM is used and the multiplexed processing time allotted for each control is diminished thus creating practical timing limits. The response times are not easily measured or calculated. Timing is a function of the computer's speed, the network's speed which can be constantly changing, the ActiveX control's size and code efficiencies and – in the case of CobraNet-based systems – how often the CobraNet interface and the CobraNet device itself scans for network and/or hardware changes. In the CobraNet case, the ActiveX control data is being transported over SNMP.

While it is impractical to suggest precise delay times for system changes, a unicast CobraNet system with 120 ActiveX controls on a single page, utilizing four, 24-port managed switches, with 64 audio channels and two simultaneously running web pages appeared to have a response time around one second. This is not heart warming, but the intent is to communicate performance levels to allow applicability decisions.

Time-critical applications such as show control and synchronized events may not be deterministic enough for ActiveX.

Uncommon and New to the Audio Industry

Few pro audio manufacturers offer ActiveX implementations of their software-controlled parameters which creates a Herculean obstacle for significant implementations across multiple manufacturers' products. Also, there is no current support for ActiveX in Netscape or the Macintosh world.

There are certainly other ways to implement cross-manufacturer control. For CobraNet-based systems, CobraCAD software from Peak Audio offers significant advantages through its use of SNMP – particularly as the CobraCAD package matures and the use of SNMP for control may offer fewer obstacles as the future unfolds. No other currently available tools seem to offer the advantages of ActiveX for the pro audio industry.

CONCLUSION

ActiveX provides a viable solution to the problem of user interface across disparate equipment in computer-controlled systems – particularly CobraNet and other Ethernet-based systems. While no solution – including ActiveX – can solve all problems, ActiveX appears to be a viable one at this juncture for the above issues.

REFERENCES

1. Proceedings of the AES 13th International Conference on Computer-Controlled Sound Systems (Audio Engineering Society, NY, 1994).
2. Tanenbaum, Andrew S. “*Computer Networks*” 3rd edition (Prentice Hall, ISBN 0-13-349945-6, New Jersey, 1996).
3. Perkins, David & Evan McGinnis, “*Understanding SNMP MIBs*” (Prentice Hall, ISBN 0-13-437708-7, Upper Saddle River, 1997).
4. Bates, Jonathan, “*Creating Lightweight Components with ATL*” (Sams, ISBN 0-672-31535-1, Indianapolis, 1999).
5. Cook, Devin, “*Emerging Standards for Networked Audio System Control*,” RaneNote 144, 1998.

Fixed-Point vs. Floating-Point DSP for Superior Audio

- **Dynamic Range**
- **Proper Gain Setting**
- **Right Tool for the Right Job**

**Greg Duckett
Terry Pennington
Rane Corporation**

**RaneNote 153
Revised 2005
© 2002, 2005 Rane Corporation**

Introduction

A popular belief is that for pro audio applications, floating-point DSPs are better than fixed-point DSPs – it depends, as you will learn. Read on to see why you may be paying too much and receiving too little when it comes to buying DSP signal processors.

A lot of confusion abounds in the audio industry over the issue of DSP internals. Things like: “Our box uses 32-bit floating point DSP chips, while their box uses only 24-bit fixed-point DSP chips. Obviously 32 is better than 24.” Obviously – if you don’t know the facts. What they fail to point out is the “rest of the story,” which is that the 24-bit fixed-point DSP box can operate in “double-precision” mode, making it a 48-bit box. And in this case, 48 really is better than 32, only it has little to do with size. With today’s technology, both fixed-point and floating-point can be equal if chosen correctly.

Since fixed-point is the most picked on, let’s begin with how it can be superior to floating-point. Here is the executive summary of why fixed-point DSPs can make for superior audio:

1. Less dynamic range (yes, in DSPs used for audio, this can be a feature).
2. Double-precision capable 48-bit processing.
3. Programming flexibility that can guarantee proper behavior under the adverse conditions presented by audio signals — truly one of nature’s oddest phenomena.
4. Lower-power consumption (floating point hardware is more complicated than fixed point; more transistors require more watts).

With that said, let’s back up and review a few things.

A truly objective comparison of DSP architectures is not easy; in fact, it may not be possible. *In the end, it is the application and the skill of the software programmer implementing the application that determines superior audio performance.* But people don’t want to hear that. They want the easy answer. Everyone is looking for the secret word, the single number that defines the difference and makes the duck drop down and deliver a \$100 bill (*apologies made to all readers who have not seen the original You Bet Your Life, NBC 1950-1961, TV shows hosted by Groucho Marx*). Yet the truth is that there is no single number that quantifies the differences. Not the number of bits, the MIPS or FLOPS rate, the clock rate, the architecture, or any one thing.

Two distinct types of DSP implementations dominate pro audio applications: one utilizes *fixed-point* processing while the other features a *floating-point* solution. Both produce the same results under most conditions; however, it is the word “most” that creates the difference.

Looking under the hood of an IEEE 32-bit floating-point processor and a 24-bit fixed-point processor reveals that each DSP design offers the same 24-bit processing precision — *precision* is not the issue. The defining difference is that the fixed-point implementation offers double precision, while the floating-point device features increased dynamic range. In floating-point processors scaling the data increases dynamic range, but scaling does not improve precision, and in fact degrades performance for audio applications (more on this later). And it turns out that the strength of the fixed-point approach is the weakness of the floating-point, giving fixed-point a double advantage.

The benefit is most obvious in low frequency audio processing. This is important since most of the energy in audio lies in the low-frequency bands (*music and speech have an approximate 1/f spectrum characteristic, i.e., each doubling of frequency results in a halving of amplitude*). The simple truth is that the floating-point technique struggles with large amplitude, low-frequency computations. In fact, building a high-Q, low frequency digital filter is difficult no matter what method you use, but all things considered fixed-point double-precision is superior to floating-point single-precision.

To thoroughly explain, in a scientific engineering manner, the advantages and disadvantages of the two techniques as they relate to broadband audio applications is a vastly complex subject and lies beyond the scope of this article. An objective direct comparison involves a steep slippery slope of complexities, qualifications and definitions, all necessary in order to avoid apples to oranges error. A task as daunting as this has already been done by Dr. Andy Moorer[1] (cofounder of the Stanford University Center for Computer Research in Music and Acoustics, then CTO at the Lucasfilm Droid Works, then cofounder of Sonic Solutions, and now Senior Computer Scientist at Adobe Systems) and is recommended for the detail-curious and mathematically courageous. The goal here is to draw a simplified illustration of the audio challenges faced by each DSP solution. (See the *Let's Be Precise About This...* section for a more in-depth mathematically oriented example — sorry, but mathematics is all we are dealing with here.)

Dynamic Range

Higher dynamic range is better, yes? Just as lower distortion is better and lower noise is better. We're reminded of a sales guide published by a hi-fi manufacturer in the mid-seventies: This guide had a "lower is better" and "higher is better" approach to equipment specifications. The author of that promotional material would be shocked to hear an audio manufacturer claim that higher dynamic range can be a problem. Nonetheless, it is a fact when examined in relationship to the ultra-high dynamic range capabilities of the 32-bit floating-point processors found in some of the current DSP audio signal processors.

As mentioned earlier, both DSP designs have a 24-bit processor for the mainstream functions. The fixed-point technique adds double precision giving it 48-bit processing power, while the floating-point design adds an 8-bit exponent. The 8-bit exponent gives the floating-point architecture an astonishing dynamic range spec of 1500 dB (8-bits = 256, and 2^{256} equals approximately 1500 dB) which is used to manipulate an operating window, within which its 24-bit brain operates. Floating-point

processors automatically scale the data to keep it within optimum range. This is where the trouble lies and this is why fixed-point is better than floating-point for audio. It is not that the dynamic range is the problem so much as the automatic scaling over a 1500 dB range that is the problem. Fixed-point, with its 48-bits, gives you 288 dB of dynamic range — enough for superior audio — but the programmer has to scale the data carefully. Floating-point programmers leave it up to the chip, but, unless they are careful, that creates serious errors and noise artifacts. All the jumping about done by the continuous signal boosting and attenuating can produce annoying noise pumping.

Floating-point evangelists flaunt their DSP dynamic range as a plus, but it only shows their weakness. *The dynamic range specification of a DSP chip has little to do with the overall dynamic range of the finished product.* The dynamic range of the "box" is bounded by the A/D converter on its input, to some extent on the processing in the center of the device, of which the DSP chip is a part, and on the D/A converter on the output (*even without converters the output of both DSP types is a 24-bit fixed-point word*). The dynamic range of a DSP chip is the ratio between the largest and smallest numbers it can handle. If a DSP device sporting an advertised dynamic range of 1500 dB resides between the input converters and the output converters its contribution to the overall dynamic range of the product is limited to the dynamic range of the converters. Is this a bad thing? No, not in itself.

What's bad about a floating-point processor with a dynamic range of 1500 dB is that it scales its processing range based on the amplitude of the signal its dealing with, but when dealing with signals of differing amplitudes (i.e., real audio), the scaling may not be optimized for the mixed result. When dealing with audio signals the installer cannot simply ignore the subtleties of system setup because they have a floating-point processor in their box.

Consider the typical audio mixer scenario: At any given moment a mixer can have multiple levels present at its many input ports. Input-1 might have a high-level sample to deal with while Input-2 has a very low level, Input-3 somewhere in the middle of its upper and lower limits and so on. A 32-bit floating-point DSP chip makes a determination about the appropriate window within which to work on a sample-by-sample basis but finally represents its calculations in the same 24-bit manner as its fixed-point counterpart. Even in a simple two-channel stereo processor signal levels between channels, while similar in average level, can be vastly different instantaneously due to phase differences.

Nothing is gained by using a floating-point device in an audio application but much may be lost. It does not have the 48-bit double precision capability of a fixed-point solution, and noisy artifacts may be added.

Importance of Proper Gain Setting

The reality here is that so long as we have finite precision/dynamic range in the converters and DSPs, the installer plays the final and most important role in maintaining the proper processing window alignment for a given installation. Improperly set gain structure can overload fixed-point processors. While floating-point DSPs give the flexibility to misadjust the system (too much internal gain) without noticeable internal clipping, they

still suffer the unintended consequences of the misalignment (say, in trying to mix two channels of very different audio levels) that floating-point processors cannot fix. They merely mask the problem from the installer's view. Or, worse, produce audible and annoying rise in quantization noise when filters are used below 100 Hz. In this sense the fixed-point processors force the installer to maintain the 144 dB processing window by avoiding internal clipping through proper gain structure/setup and so make maintaining overall quality easier than floating-proper processor based boxes.

Double Precision

The double precision 48-bit processing is used when long time constants are required. This occurs when low frequency filters are on the job and when compressors, expanders and limiters are used with their relatively slow attack and release times. If 24 bits are all that are available when more precision is required, the results are a problem. The function misbehaves and the least damaging result is poor sound quality. The worst result is amplifier or loudspeaker damage due to a misbehaving DSP crossover, making double precision a must-have for superior audio.

Examples and Counterexamples

Floating-point evangelists like to use an example where the processor is set up for 60 dB attenuation on the input and 60 dB make-up gain on the output. Leaving aside the absurdity of this fabricated example, let's use it to make our fixed-point-is-better point: add a second input to this example, with the gain set for unity, a 0 dBu signal coming in, and configure the processor to sum both these channels into the output and listen to the results — you will not like what you hear.

Another revealing example is how you never hear floating-point advocates talk about low-frequency/high-Q filter behavior. The next time you get the opportunity, set up a floating-point box parametric filter for use as a notch filter with a center frequency of 50 Hz and a Q of 20. First listen to the increase in output noise. Now run an input sweep from 20 Hz to 100 Hz and listen to all the unappetizing sounds that result. Audio filters below about 100 Hz require simultaneous processing of large numbers and small numbers — something fixed-point DSPs do much better than their floating-point cousins.

Free the Developers

The real determinant of quality in audio DSP is the skill of the programmers. They must devise accurate and efficient algorithms; the better their understanding of the (sometimes-arcane) math, the better the algorithm; the better the algorithm, the better the results. Fixed-point processing delivers a load of responsibility to the hands of the developer. It also delivers an equal amount of flexibility. A talented engineer with a good grasp of exactly what is required of a DSP product can fashion every detail of a given function down to the last bit. This is not so with floating-point designs. They offer an ease of programming that is seductive, making them popular when engineering talent is limited, but not the best choice. On one hand it is easier to program but on the other hand it is less controlled as to the final results — and, as we all know, that is what is important.

The Right Tool for the Right Job

If fixed-point DSP devices are so good, then why do floating-point DSPs exist? Fair enough question. They exist because DSP applications differ widely. Some of the more popular floating-point applications are found in physics, chemistry, meteorology, fluid dynamics, image recognition, earthquake modeling, number theory, crash simulation, weather modeling, and 3-D graphics. If you are designing an image processor, a radar processor, anything to do with astronomy, or a mathematics matrix inverter, the choice is clearly a floating-point solution. As always, the application dictates the solution.

What is required in a floating-point DSP to achieve superior audio? Here are some pretty nasty "ifs" necessary for floating-point to overtake fixed-point: *if* it is a 56-bit floating-point processor (i.e., 48-bit mantissa plus 8-bit exponent) or 32-bit with double-precision (requiring a large accumulator), *if* the parts run at the same speed as the equivalent fixed-point part, *if* they use the same power, and *if* they cost the same, then the choice is made.

Another possibility is if the floating point DSPs evolve to offer significantly more processing power for the same price (enough to overcome the low-frequency, high-Q issues in firmware) and offer a compatible peripheral chip set, then this could tip the scales even if they still offer only a 32-bit fixed numerical format.

Let's Be Precise About This ...

An example is the best way to explain how you lose precision when floating-point processors scale data. Assume you have two mythical 3-digit radix-10 (i.e., decimal) processors. One is "fixed-point", and one is "floating-point." For simplicity, this example uses only positive whole numbers. (On real fixed- or floating-point processors, the numbers are usually scaled to be between 0 and 1.)

The largest number represented in single precision on the fixed-point processor is 999. Calculations that produce numbers larger than 999 require double precision. This allows numbers up to 999999.

Let the floating-point processor use 2 digits for the exponent, making it a 5-digit processor. This means it has a dynamic range of 0 to 999×10^{99} = HUGE number. To see how this sometimes is a problem, begin with the exponent = 0. This allows the floating-point processor only to represent numbers up to 999 — same as the fixed-point single-precision design. Calculations that produce numbers larger than 999 require increasing the exponent from 0 to 1. This allows numbers up to 9990. *However*, notice that the smallest number (greater than zero) that can be represented is $1 \times 10^1 = 10$, *meaning numbers between 1-9 cannot be represented (nor 11-19, 21-29, 31-39, etc.)*. Increasing the exponent to 3 only makes matters worse, but you can cover (almost) the same range as the fixed point processor (up to 999000); however the smallest number now represented is $1 \times 10^3 = 1000$, *meaning numbers between 1 and 999 cannot be represented*. And the next increment is $2 \times 10^3 = 2000$, meaning the represented number jumps from 1000 to 2000. So that now numbers between 1001 to 1999 cannot be represented. With exponent = 3, each increment in the mantissa of 1 results in an increase in the number of 1000, and another 999 values that cannot be represented.

Is this as big a problem as it first appears – well, yes and no. At first it looks like the floating-point processor has lost the ability to represent small numbers for the entire calculation's time, but the scaling happens on a *per-sample* basis. The loss of precision only occurs for the individual samples with magnitude greater than 999. Now you might think that everything is fine, because the number is big and it does not need the values around zero. But a few wrinkles cause trouble. When calculations involve large and small numbers *at the same time*, the loss of precision affects the small number and the result. This is especially important in low-frequency filters or other calculations with long time constants. Another wrinkle is that this happens automatically and beyond the control of the programmer. If the programmer does not employ the right amount of foresight, it could happen at a bad time with audible results.

In the fixed-point case, the programmer must explicitly change to double precision – there is nothing automatic about it. The programmer changes to double precision at the start of the program section requiring it and stays there till the work is done.

The Big and the Small

Over and over in audio DSP processing you run into the same simple arithmetic repeated over and over: multiply one number by another number and add the result to a third number. Often the result of this multiply-and-add is the starting point for the next calculation, so it forms a running total, or an accumulation, of all the results over time. Naturally enough, adding the next sample to the previous result is called an “accumulate” and it follows that a multiply followed by an accumulate is called a MAC. MAC's are the most common of all operations performed in audio DSP, and DSP processors typically have special hardware that performs a MAC very, very quickly.

As results accumulate, errors also accumulate. As well, the total can get large compared to the next sample. To show this in action return to the mythical 3-digit processors. Say we have the series of numbers shown in the row labeled “Samples” in Table 1; a strange looking set of numbers, perhaps, but it represents the first part of a simple sine wave. Multiply the first number by a small constant (say, 0.9) and add the result to the second number: $0 \times 0.9 + 799 = 799$. Multiply this result by 0.9 and add it to the third number: $799 \times 0.9 + 1589 = 2308$. And again: $2308 \times 0.9 + 2364 = 4441$. Continue this pattern and it forms a simple digital filter. The results using double precision fixed-point are shown in the row labeled “Fixed-Point Results” in Table 1.

Sample #	Samples	Fixed-Point Results	Floating-Point Results
1	0	0	0
2	799	799	799
3	1589	2308	2290
4	2364	4441	4420
5	3115	7112	7080
6	3835	10236	10200
7	4517	13729	13600
8	5154	17510	17300
9	5739	21498	21200

Table 1: Results Between Floating- and Fixed-Point Accumulation

What about the floating-point processor? Start with exponent = 0. The results are: 0, 799, ... the next number is too big, so increase the exponent to 1 ... 2290, 4420, etc. Notice that the floating-point values are smaller than they should be because the limited precision forces the last one or two digits to be 0. It's easy to see that each result has an error, and the errors are carried forward and accumulate in the results. Algorithms with long time constants, such as low frequency filters, are especially prone to these errors.

You'll also notice that the accumulated values are getting larger than the input samples. The long time constant in low frequency filters means that the accumulation happens over a longer time and the accumulated value stays large for a longer time. Whenever the input signal is near zero (at least once every cycle in a typical audio signal) the samples can be small enough that they are lost; because the accumulated value is large, the samples fall entirely outside the precision range of the floating point processor and are interpreted as zero. The double precision available in the fixed point processor helps the programmer to avoid these problems.

Further Study

Digital audio is a vast and complex subject with many subtleties when it comes to superior signal processing. An article this short touches only some of the important issues. Selecting just one book from all the possibilities for recommendation for further study is easier than it may seem. That book is John Watkinson's *The Art of Digital Audio, 3rd ed.* (Focal Press ISBN 0-240-51587-0, Oxford, England, 2001). The title says it all. One of the best digital audio references you can own.

References

1. Moorer, James A. (www.jamminpower.com) “48-Bit Integer Processing Beats 32-Bit Floating Point For Professional Audio Applications,” presented at the 107th Convention of the Audio Engineering Society, New York, September 24-27, 1999, preprint no. 5038.