Bonus chapter

PIC-A-STAR a Software Transmitter And Receiver

Unusually, this chapter is devoted to a single project. It is, however, a project that can be dipped into by those wishing to learn about modern design and construction techniques. It is by no means necessary to build the entire project as many of the techniques and modules can be used elsewhere. Pic-a-STAR may simply be read as a tutorial for anyone embarking on electronics construction in the 21st century.

PIC-A-STAR was originally published as a 20-part series in RadCom *[1] and was written by Peter Rhodes, BSc, G3XJP [2]. It is reproduced here almost in its entirity.*

This is a detailed construction project aimed at those of modest experience who would like to enhance both their craft and technology skills.

At the outset -like me - it may well be that you don't have the skills or knowledge to build this project. By the end, you will have. That is, as I see it, the whole idea.

By design, this is a project without end. From my perspective, it is the basis for years of happy building to come -and is my first investment in a new core transceiver platform in some 25 years. A glance at **Fig 8.1** tells you why I needed a new one.

From your perspective, it is a source of ideas for improving an existing transceiver - not least, upgrading the back-end with a powerful Digital Signal Processing (DSP) capability. There are also some craft techniques for handling small-size high-function components. So, there is something in this for all, with an eye on the self-education requirement of their licence.

SUMMARY

The heart of PIC-A-STAR is the DSP module. This provides both the back-end receiver functionality, as well as SSB/CW generation on transmit. The bottom line is absolutely superb audio quality on both transmit and receive. If you want to test the former, come on the home-brew net frequency (3.727MHz) any day around lunchtime where you will find at least one STAR in operation most days. If you want to test the latter then you will just have to make one.

Being implemented by software, it provides the opportunity to address both absolute performance as well as the delights of operational convenience - at zero incremental cost. This is precisely the basis for future developments, but the fundamental functionality together with some bells and even the odd whistle has been in daily use here for about nine months at the time of writing. This is the project on offer - but by the time you get there it will have moved on.

PIC-A-STAR is explicitly designed to be upgraded over the web, so there will no incremental DSP enhancement costs.

An outline of the main features of the project is shown in **Table 1**.

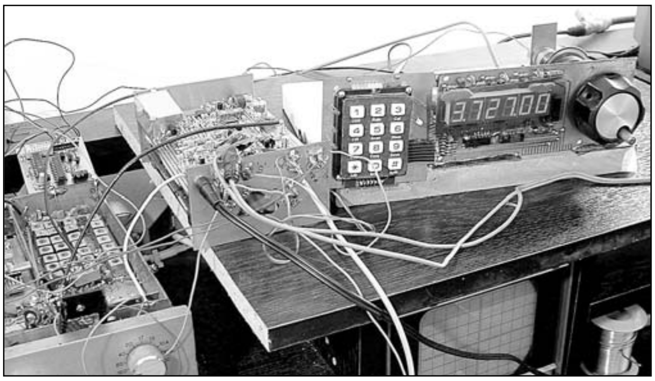

Fig 8.1: Early integration testing. Bottom left is the author's Third Method transceiver (borrowed front-end and PA), top right is Pic 'N' Mix DDS still on its original breadboard (injection and controls) - and in the middle is the new DSP module. Note that Pic 'N' Mix provides all the transceiver controls, leading to a clean and compact front panel

- SSB and CW detection and generation
- a bank of high-performance Rx filters
- impulse noise blanking
- non-coherent noise reduction
- auto-notch heterodyne removal
- variable AGC decay time
- synthetic stereo reception
- adjustable RF clipping on transmit
- very fast VOX and QSK operation
- the flexibility to change!

Table 8.1: A brief outline of the features of PIC-A-STAR

POSITIONING DSP

You might reasonably expect the author of a DSP project to have some serious knowledge in the field. So would I! Actually, in many ways, it is important to get this published before I acquire more than enough to be merely dangerous.

If, like me, you are at least in your late 50s, it is unlikely that DSP theory featured even in a formal engineering education. And if, equally like me, you have never worked in the engineering profession then you could reasonably start from the position that DSP is some kind of black magic which you could never understand in a life-time of trying. You might well be correct in this assumption because some of the theory is indeed very heavy.

But my personal discovery was that you don't need to understand DSP at other than a superficial level to be able to build it at home and to use it.

From a position of not being able to spell DSP, it took me two weeks to get my first DSP receiver working. The attraction is that everything since then has been incremental and I have not been off-air for a single day.

Design mistakes - and there have been many - have cost me my time but never any money - which is about perfect for a hobby. So this lends itself to a learn-as-you-go approach. In other words, unlike conversational French, you don't have to learn a lot before you can even get started.

SKILLS AND FACILITIES

A requirement of all my projects is that they can be built on the kitchen table with no access to professional facilities. Otherwise, it would not be amateur radio.

This one is no exception - though I have had to acquire new skills and hone them to the point of repeatability in order to build some of the hardware. This is all part of the adventure, part of the fun.

A simple (and in-expensive) technique for making precision PCBs will be covered - which includes mounting a 48-pin chip with a mere 0.5mm interval between pins. And you get to practice on a really easy one of 128 pins by 0.8mm first. If the prospect of this puts you off, I really can't help. If it sparks a 'can-do' spirit of adventure then we are in business.

INSPIRATION

THREE THINGS made this project possible. In the order in which I found them:-.

- *The Scientist and Engineer's Guide to Digital Signal Processing*, by Steven W Smith. This book is a little gem. If you flick through quickly, you will see copious examples and illustrations. What you do not see are lots of equations and impenetrable notation. I need just one quote:- "[this book] . . . is written for those who want to use DSP as a tool, not a new career." My kind of book!
- The Analog Devices website [3]. This contains a wealth of both theoretical and practical information - and specifically the electronic version of the above book. Most valuable to me were lots of DSP code examples for the ADSP-218x processors. The first incarnation of STAR was built by six of us on the ADSP-2181EZLITE evaluation board which had become somewhat of a standard over the years. Then over one fateful weekend when this project was 'finished', its price went from $90 to $275 - which spurred the challenge to home-brew a compatible and reproducible DSP board.
- DSP-10, a 2m DSP transceiver project published by *QST* in September - November 1999 - and reviewed in *RadCom*, Feb 2000. Although featured for VHF/UHF applications, the DSP core is totally universal. This project was designed by Bob Larkin, W7PUA [4], and I am indebted to Bob not only for the inspiration for this project, but for a significant amount of advice and help - including some code written specifically for PIC-A-STAR. Above all, Bob showed it can be done and whenever I get into problems, his material is the first place I look for clarification and understanding.

INTEGRATING PIC-A-STAR

The DSP module - designed to combine with the Tx/Rx RF stages of your choice - operates at a final IF of 15kHz as shown in **Fig 8.2**. This is a high enough frequency to make it immune from image responses, yet low enough to be affordable. And it is not a DSP audio add-on - which, coming after the product detector, will always struggle.

RF STAGES

Your HF IF can be derived from any reasonable transceiver front-end. My third method transceiver [5] and G3TSO's modular transceiver [6] have both been tested as representative - and there are lots of them out there. CDG2000 looks like a powerful approach

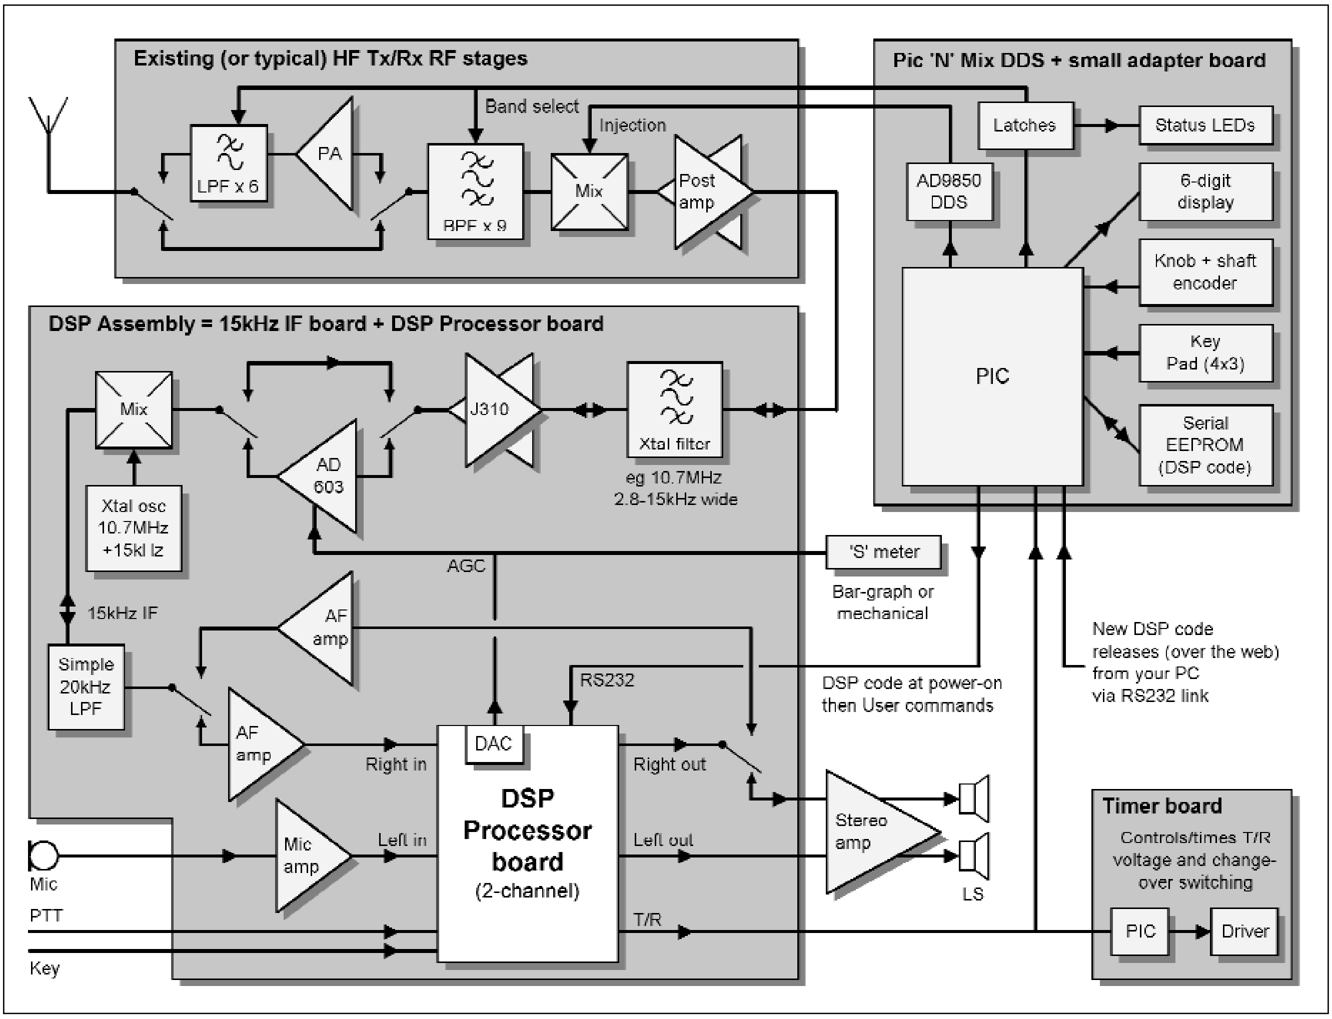

Fig 8.2: A typical transceiver incorporating PIC-A-STAR at a final IF of 15kHz. See text for a discussion of the major hardware elements

and its front-end could well be my next increment. The choice will substantially determine the overall receiver strong-signal handling capability - but not the intrinsic benefits of downstream DSP.

IF STAGES

In principle (and possibly in practice), you could modify an existing IF board so that its product detector produced a 15kHz output instead of straight audio but I don‘t recommend it in the long run. To cut a long story short, this design includes an IF board built for the job.

Details of this follow later, but it needs only a modest roofing filter (at any HF IF of your choosing) since all the serious filtering is implemented in DSP.

CONVERSION INJECTION

You won‘t be surprised to see Pic ‘N‘ Mix [7] used as the injection source to mix from RF to your chosen HF IF. This is not mandatory, but my records show 281 of them out there, so it is a non-trivial population.

COMMAND AND CONTROL

You need the ability to command the DSP for all the functions normally associated with front-panel controls. You may be some what surprised to see Pic ‘N‘ Mix used for this purpose as well.

A small adapter board is used to fit a more versatile and powerful PIC - which not only controls all the original DDS capability, but extends the existing keypad, display and tuning knob to control all the transceiver features.

Although highly recommended, use of Pic ‘N‘Mix is not mandatory. As an alternative, you can use your PC to load and control the transceiver - and a BASIC utility is provided to achieve this.

PIC-A-STAR DSP

Illustrated in **Fig 8.3** is the functionality implemented in software. You will have seen not dissimilar block diagrams implemented in analogue hardware - but not at this price and not inside a 28mm² chip!

Actually, there is an intrinsic overhead, namely that the analogue signals need converting to digital form before processing and back to analogue after. This is the purpose of the CODEC (encoder/decoder) referenced on several inputs and outputs - and is implemented on a separate chip.

The greatest appeal of the software approach, not least to the amateur, is the flexibility to change the line-up at a touch of the keyboard, so to speak. This allows easy experimentation (or overt tinkering, if you prefer), since at any time you can abandon the change and go back to the previous version. There are other subtleties.

For example, you will find five 15kHz oscillators scattered around the diagram. In fact their frequencies change depending on mode ie USB/LSB/CW. In DSP software terms the sinusoidal oscillator is simply a subroutine. To invoke it, all you need do is tell it what frequency/phase you want on any given occasion - and it is done.

Another example is the ‘delay’ in the receiver front-end image-cancelling I/Q mixer. In one path, there is a 90° phase shift, in

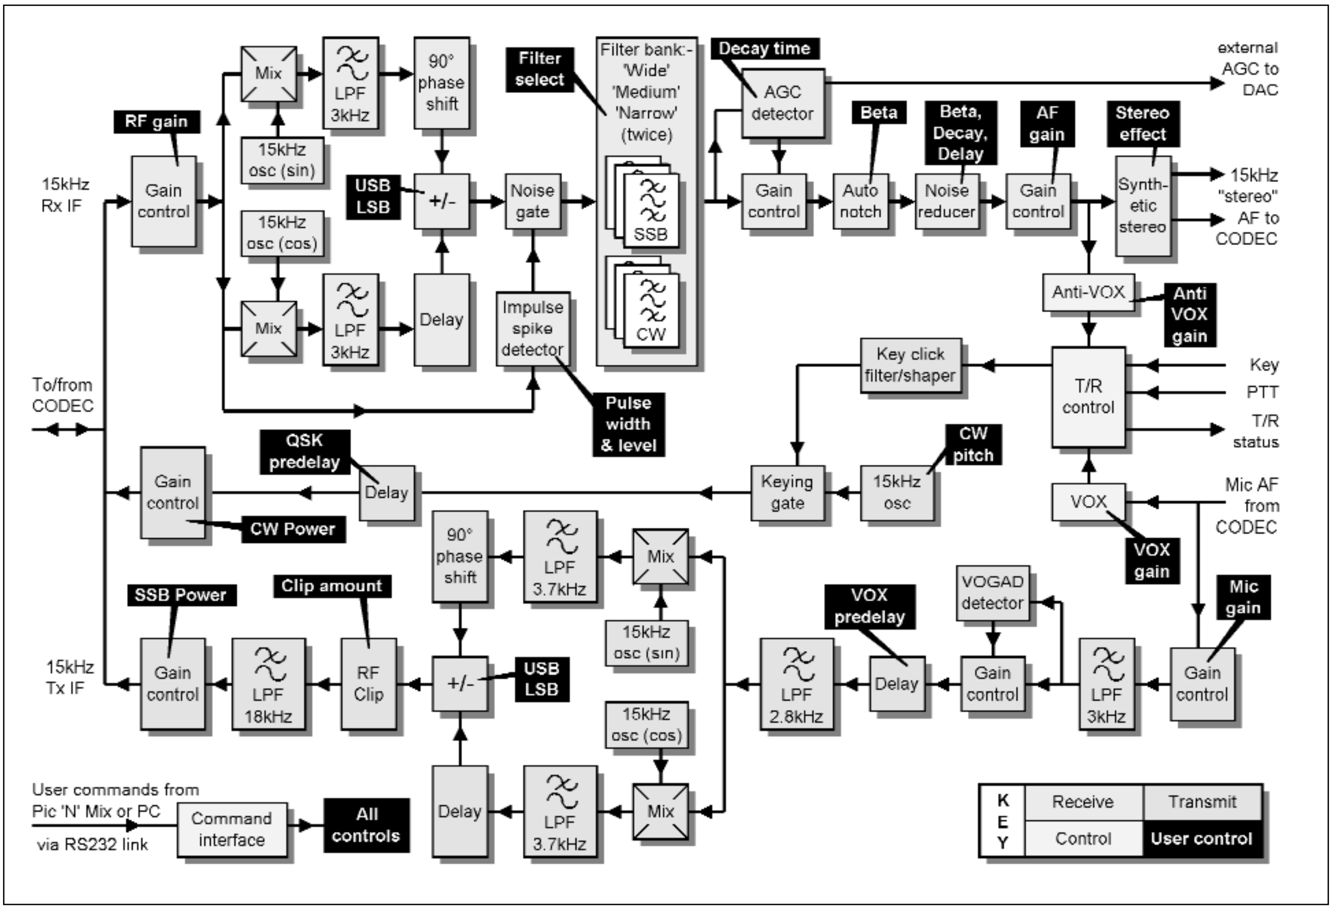

Fig 8.3: Software block diagram of PIC-A-STAR DSP functionality. Not shown are simple on / off switches associated with VOX, noise blanker, auto-notch, noise reducer and the RF clipper. The filter bank also has an off (ie bypass) switch to give a net maximum bandwidth of some 3kHz

the other a delay. The latter arises because it takes real elapsed time to produce the phase shift, so an equal amount of time has to be 'wasted' in the other channel to maintain that phase relationship.

Time is of the Essence

The basic understanding you need in order to grasp how DSP works is to note that time is the critical commodity. Every functional box in **Fig 8.3** takes time to execute. So does every individual instruction that goes to make up that functionality.

This would be of little concern were it not for our old friend Nyquist. He stated that, in order to process a signal faithfully, you must sample it at (at least) twice the rate of the highest frequency present.

For example, the incoming receive signal is around 15kHz, and so needs to be sampled at 30kHz or more. In fact, 48kHz is used to provide a useful margin.

The consequence of this is that, having grabbed one sample, you have no more than 20.83μs (by simple arithmetic) to do all the processing required before you have to get back to handle the next one.

So just how much processing can be achieved in 20 millionths of a second? The ADSP-2181 processor in this design executes an instruction in 30 nanoseconds. The simplistic answer is therefore 666 instructions-worth. But this is far from the whole story. During one processor cycle it can, for example, fetch two 16-bit numbers, multiply them to give a 32-bit product and add the result to a 40-bit accumulator. This MAC (Multiply & Accumulate) instruction is the essence of filter implementation and is critical because you need to loop around it many times. Meanwhile, in the background, the processor is also organising data samples in and out of the CODEC as well as handling any serial communications port activity.

Fig 8.4 shows a snatch of PIC-A-STAR code, so you can visualise just how much radio you get from each line of code.

Multi-rate Processing

There is a more structural solution to the issue of buying some time - which, equally, derives from Nyquist. Namely, once the receive signal has been mixed down to audio, you no longer need to process it at the 48kHz rate. Twice the audio frequency is fast enough.

PIC-A-STAR runs audio processing at 8kHz - by grouping the audio functions into six blocks and running one of them - but each in turn - during six successive 20.83μs time-slots. At the end of each slot the data is again processed at 48kHz because that is the sample rate used by the CODEC for outbound signals also.

```
{ Fetch Rx sample via CODEC and place in register mx0 ...}
        mx0=dm(Rx_in_buffer);
{ ... and fetch current RF gain value and place in register my0. }
        my0 = dm(RF_gain);
{ Multiply the two together to give a gain-controlled value ...}
        mr = mx0 * my0 (SU);
{ ... and keep the gain-controlled signal in register my0. }
        my0 = mr1;
{ Fetch the phase incremented value of LO and place in register ax0 ...}
        ax0 = dm(LO_ phase);
{ Pass the phase value to sin to get instantaneous sinusoid ampli-
tude ...}
        call sin;
{ ... and mix (ie multiply) it with the signal in register my0 }
        mr=ar*my0(SS);
```

Fig 8.4: Some early lines of code for the receiver. Yes, the last line truly is a mixer (otherwise known as a product detector)

So you can see that, all the way along the line, Nyquist is satisfied -and so am I because there is plenty of time for some exotic as well as the more mundane processing.

SOFTWARE PACKAGING

The desire to provide choice and flexibility, but above all upgradability (if that is an English word), leads to some complication in describing the various modules. The context will become clearer once the hardware functionality has been covered. Suffice it to say at this stage that, from an operator's perspective, the system is totally transparent, ie you just switch it on, wait about 20 seconds (as if for the valves to warm up) and then use it. The software comes in the following modules:

DSP Boot Utility

This code resides in PROM on the DSP board. At power-on time, besides running some basic hardware checks, it manages the on-board serial port to load the target DSP code. This utility was written by Bob Larkin, W7PUA, for PIC-A-STAR based on the original AD code.

DSP Transmit/Receive Code

This runs on the DSP board and provides the core functionality as in Fig 8.3. It needs to be loaded at power-on time, a process which takes some 20 seconds. Subsequent to loading it, you also need to be able to command it.

DSP Loader

This is a QBASIC utility which runs on your PC. It is written in very basic BASIC to enable you to adapt it or port it if you wish. It has two distinct alternative functions:

- to load and subsequently command the DSP code directly to the DSP board, via a COM port and a 9.6kB serial link.
- to load a new (or, of course, first) release of the DSP code to the PicAdapter board (see next) in Pic 'N' Mix. Subsequently, Pic 'N' Mix automatically loads the code at power-on time -and provides the command user interface.

These alternatives are not mutually exclusive. For early testing and use, the former gets you going quickly. The latter frees up your PC and, in my view, gives a much cleaner user interface - albeit with a little practice. The choice is yours. (There is a further option here. You could build a dedicated controller using any programmable device with an RS232 capability. The command syntax is simple and also provided - and is in any event self-evident from the QBASIC code. With some loss of maintainability, you could also burn the entire transmit/receive code into the boot PROM / EPROM.)

PIC 'N' MIX PicAdapter

Written in MicroChip Assembler, this code runs on a 16F870 (which replaces the present 16x84) to provide all the original DDS control functionality of Pic 'N' Mix and, in addition, it now integrates the ability to:

- download new release DSP code from your PC (via the web);
- subsequently upload that same code to the DSP board at power-on time;
- command the DSP using the self-same keypad, tuning knob and display as already fitted to Pic 'N' Mix.

Timer Board

Also in MicroChip Assembler, this code runs on a 16F627. It provides the sequencing and timing of receive/transmit transitions - both ways - to make them as clean and fast as possible. This board is designed to be general and will find uses on other transceiver projects.

Bargraph S-Meter

This is both optional for PIC-A-STAR and equally of general application. Also running on a 16F627, it controls a 12-LED bargraph on the Status board. It was built at all because 10 LEDs, as provided by most control chips, are not enough - and in any event, the PIC provides a lower-cost solution.

To Summarise

The programmed chips are a PROM, a 16F870 and two 16F627s. These provide the base infrastructure. The target DSP code - where most of the future enhancements will occur - is loaded from your PC using the QBASIC utility. No further hardware (eg a programmer) is needed.

SOFTWARE DISTRIBUTION

All the software itemised above will be available for your personal use at no charge. However, this does not mean it comes entirely free. The 'price' is that you need to send me an e-mail note requesting the software - giving an explicit undertaking that it is for your personal use and for the purpose of self-education.

Not least, this allows me to maintain a list of 'customers' to advise when updates become available; as mentioned previously, this is, by intent, a project without end. By software, I mean at this stage the loadable object code.

If, however, you want me to use my resources to program chips for you then I can supply them ready-programmed at a few pounds per chip - plus return postage (see [2] for author's address). It is worth pointing out in this context that you could build a programmer yourself for about £10.

Resistors 1/8-1/4W, 5-10%

R1	47R
R2, R48, R49, R76, R77	100R
R3	220R
R4	470R
R5, R6	560R
R7, R50, R64-R71	1k
R10, R43, R52, R72	2k2
R11	2k7
R12, R34, R37	22k
R13-R15	3k3
R16-R19	4k7
R20	5k6
R21-R32, R45-R47, R53-R55, R63, R73, R78	10k
R33, R51, R59-R61	47k
R35, R36	100k
R38	120k
R56-R58	10R
R62	56k
R74	1k x 4 SIL network (5 pin)
R75	1k x 8 SIL network (9 pin)
RV1, RV2, RV6	100k horizontal preset
RV3, RV4	10k horizontal preset
RV5	47k horizontal preset

The following resistors are all SMD 1206 size, 5%

R8, R9, R93	1k
R39, R40	22k
R41, R42, R97	220k
R44	3k3
R79-R89	10k
R90	330
R91, R92	2k2
R94-R96, R100-R103	4k7
R98, R99	100k
R104,R105	47k

Capacitors, 16V rating

C1-C191n solder-in feedthrough (C11 and C12 may be replaced with a stereo jack socket)

C21	10p ceramic plate
C20, C22, C23, C26, C69	220μ axial
C24	0μ radial
C25, C75, C76, C84	33μ radial
C27	47μ axial
C28	2μ2 radial
C29	4μ7 radial
C30	33μ radial
C31-C34	220n ceramic
C35, C36	1n ceramic
C37-C44, C59, C70, C71	10n ceramic
C45	22n ceramic
C46-C58, C60-C65, C67, C68 C74, C77-C83, C85-C89	100n ceramic
C72, C73	15p ceramic plate
C90, C91	22p ceramic plate
C92, C93	270p NPO ceramic
C94	47n ceramic
C95, C96	220n ceramic
C97	220p ceramic plate
C120, C121	18p ceramic plate
C122, C131	4μ7 radial
C123-C125, C128, C132	10μ radial
C126, C127, C129, C130, C133	1μ radial
C66, C98-C119	100n SMD 1206 size
VC1-3	100p polyethylene trimmer (1)

Inductors

RFC1-RFC3, RFC6, RFC7	1mH axial choke
RFC4, 5, RFC8-RFC10	100μH axial choke
RFC12	330μH axial choke
L1	4.85μH 38t 30SWG on T50-2 (2), (3)
T1	6t:3t 30SWG on FT37-43 (2)
T2	10t:3t 30SWG on FT37-43
T3	20t:4t 24SWG on T50-2 (3)
FB	1 turn through small ferrite bead

Semiconductors

IC1, IC13, IC15, IC17-IC19	78L05 regulator, 5V
IC2	AD603AQ (DIL)
IC3, IC4	FST3125M (SMD)
IC5	SBL-1 (7)
IC6, IC7, IC22	TL072 (DIL)
IC8	TLC7524CD (SMD)
IC9	PIC 16F870-ISP (DIL)
IC10	24LC256-IP (DIL)
IC11, IC28	PIC 16F627-04P(DIL)
IC12	ULN2803A (DIL)
IC14	7810 regulator, 10V
IC16	4094 (DIL)
IC20	74LVX125 (SMD)
IC21	ST232N (DIL)
IC23	74HC14 (SMD)
IC24	LE33CZ regulator, 3V3
IC25	AD1885JST
IC26	27C512 512KB PROM, programmed
IC27	ADSP-2181 KS130 or KS160
TR1, TR2	J310
TR3-TR5, TR12	2N3904
TR6	2N3906
TR7, TR8	BC517
TR9-TR11	VP0300LS P-ch MOSFET
D1-D6, D11, D15	1N4148
D7-D10	BA244 (8)
D12-D14, D34	1N4007 or similar
D16-D20, D35	3mm LED, red
D21	3mm LED, red/green tricolour
D22-D33	1.8mm LED, colours for S-meter
ZD1	4V7 200mW Zener diode

Miscellaneous

FL1	Crystal roofing filter (4)
X1	wire-ended BFO crystal (5)
X2	4MHz wire-ended crystal
X3	16.67MHz, probably custom
X4	24.576MHz low profile, Rapid Electronics
S-meter	1mA, optional
S1	8-pole on/off PCB DIL switch
S2, S3	skeleton push-to-make switch
All screened RF leads	RG174
Other screened leads	thin microphone lead
Spacers and 3mm nuts/bolts	4 sets (6)
PL1	2 strips of 9-way SIL header ie male/male
SK2-SK8	SIL socket strip
PL2-PL8	mating SIL plug strip
Turned pin socket IC6, 7, 10, 22	8 pin
Turned pin socket IC16, 21	16 pin
Turned pin socket IC11, 12, 28	18 pin
Turned pin socket used as spacer	18 pin
Turned pin socket IC9	28 pin by 0.3in
Turned pin socket IC26	8 pin by 0.6in

NOTES

(1) These three trimmers may all be replaced with fixed capacitors after adjustment.

(2) Assumes a 2.2k-ohm impedance.

(3) Assumes FL1 and X1 around 10.7MHz.

(4) Any HF centre frequency. Filter width between 2.7kHz and 15kHz. Around 3 - 6kHz is ideal.

(5) Centre frequency of FL1 plus 15kHz.

(6) Spacing specified to give clearance between FL1 and the DSP board.

(7) The SBL-1 mixer may be replaced by a stronger mixer and / or one using discrete components at your discretion. It is specified here as a well-established datum and works well.

(8) These diodes can be replaced with 1N4148 if you can't obtain RF switching diodes.

SUPPLIERS

Most of the components were procured from Farnell or Rapid Electronics. The exceptions are the toroids which I bought from Mainline - and the crystal filter, the specification of which is loose enough that it should be possible to find a surplus one. The two critical chips, namely IC25 and IC27, are manufactured by Analog Devices.

Table 8.2: Components list for all options, ie Timer board, DSP boards and assembly, IF board, PicAdapter board, Status board

PCB MANUFACTURE

IT HAS BECOME a tradition with my projects that each has been produced by a different one-off PCB production technique - in an attempt to dispel any unwarranted mystique and even some phobia. Frankly, I want to advance my own craft skills with every project. I used iron-on laser film to speed-up the development cycle and as the only realistic approach to making the DSP board - and found the results to be excellent.

Process Overview

In outline, the process is to photocopy the published artwork onto the film; then transfer the toner from the film to the board using heat and pressure from a clothes iron. This toner then acts as a superb resist while etching the copper in the normal way.

This is not entirely a precision engineering process and there are some experiential skills. Cleanliness is everything! Although incredibly fine lines and small spacings reproduce well in my experience, ironically there is sometimes difficulty with large black areas. But equally, these are the easiest to touch up with an indelible pen before etching - and, at times, you absolutely will need to.

The finished and populated board is illustrated in Fig 8.5.

Resources

You need access to a black-and-white laser photocopier. Almost all modern machines use this technology. It may well be that the copier in the corner shop will not be up to the job and you should probably consider the small cost of taking it to a professional copy shop as money well spent.

Some sheets of laser film (sold nowadays by several suppliers for this purpose) are required; a block of flat scrap wood larger than the PCB; and a domestic clothes iron constitute the tools.

The latter should preferably not be a steam-iron. If it is, ensure it is fully-drained, since water/steam and this process do not mix. Also, the steam holes in the sole-plate are unhelpful. Avoid if possible the more modern easy-iron technology which has fine ridges on the sole-plate. Check that the sole-plate is flat. Some have a slight curvature. If faced with these problems, it is best to use several (say, three) intervening layers of clean paper to provide a more evenly distributed heat source. Experiment!

The PCB material (all double-sided) can start out badly discoloured, but must not be mechanically damaged, ie no scratches. 1oz (or more) copper is better, but I used merely 0.5oz on GRP for all my boards.

The Process

1. Firstly, test-copy the artwork onto plain paper in order to check for copier quality and acceptable scaling error. Use the maximum contrast consistent with retaining a clean white background.
2. Copy the artwork onto the film. When viewed with the toner (matt) side down, you want to end up looking at the tracking with the correct orientation ie as if viewing the finished board. For most published artwork this requires the extra step of firstly copying it to a transparency, flipping it over and then copying that to the iron on film. In order to avoid this extra step - with some inevitable degradation - the PCB artwork in this project will be printed pre-flipped so to speak -and therefore should be copied directly to the film. The film itself is not "sided'.
3. Cut the PCB to size (or, preferably, somewhat over-size for now).

Remove all burrs and sharp edges. With cold water, wet a soap-impregnated wire-wool pad and use it to polish the copper - with increasingly light strokes - until immaculate; do not touch

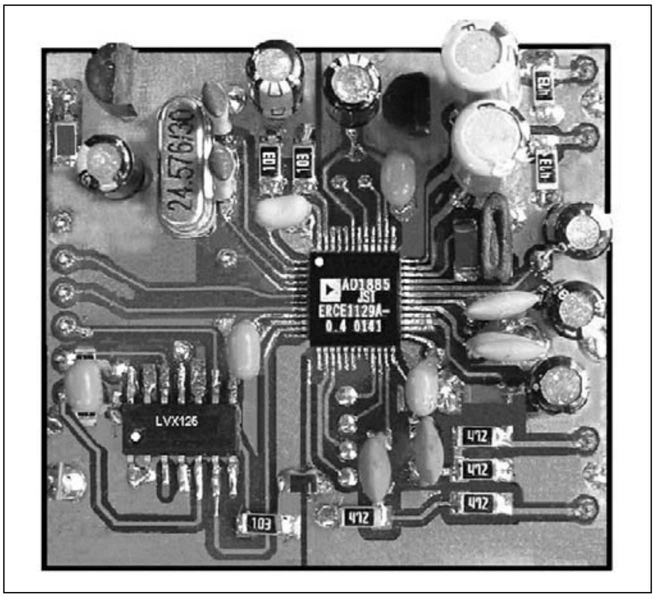

Fig 8.5: The completed CODEC board, illustrating the quality of PCB production available (0.2mm wide tracks at 0.5mm intervals) - using domestic kitchen resources. The CODEC chip is 7mm square

the surface thereafter. Polish both sides and then wash off all traces of soap residue with a clean paint-brush and cold water - and dry with kitchen paper.

4. Place the PCB on the scrap wood and clean it with some kitchen paper (uncoloured) moistened in acetone, isopropyl alcohol or cellulose thinners.
5. Heat the iron to about 140°C (cotton setting) and leave it for a few minutes to attain an even temperature across the sole-plate. At this temperature it should just scorch plain 80gsm copier paper.
6. Cut out the artwork to no larger than the PCB and register it toner side to the board.
7. With at least one sheet of clean paper interposed, lower the iron vertically onto the middle of the board and let it rest there for some five seconds. This will establish the registration of the artwork to the copper.
8. If the board is bigger than the iron, lift the iron off and relocate it every few seconds. Under no circumstances use an 'ironing' motion. Simply raise and lower it vertically - and frequently - until all the board has seen the iron and some applied pressure for about 20 seconds. For pressure, the weight of the iron plus about as much again is near enough and is not critical. Too little pressure and the toner will not transfer. Too much and the toner will migrate to widen the lines (ie smear) and reduce its depth. The former is correctable, the latter is absolutely not.
9. Inspect the result. You may see any areas which have not transferred as still retaining a somewhat glossy appearance. Repeat selectively as necessary. Pay particular attention to the edges.
10. Allow the board to cool naturally back to room temperature.
11. Carefully peel back the film from each corner and note that the toner has transferred. If any critical areas have not taken, the artwork will still be registered and you can selectively repeat.
12. Touch up any blemishes or areas of visibly thin toner with an indelible pen.
13. Now spray-mask the opposite side of the board and etch as normal.
14. Before removing the etch resist, centre-pop and/or drill the holes. They are easier to see at this stage. Clean off the etch resist with cellulose thinners and gently re-polish the board.

Fig 8.6: Timer board circuit diagram. This provides timed transitions between transmit and receive - in both directions. S1 allows you to set up the timing for your installation - and covers the range from slow relay-based RF T / R switching through to solid-state QSK

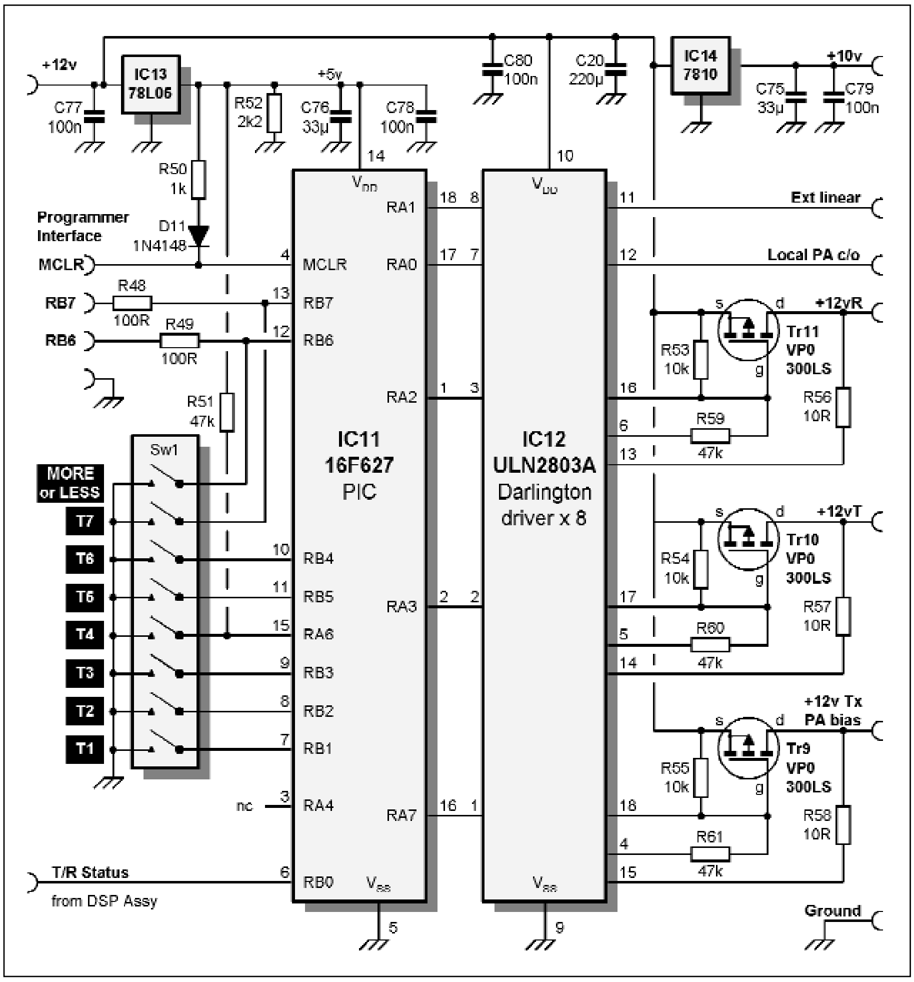

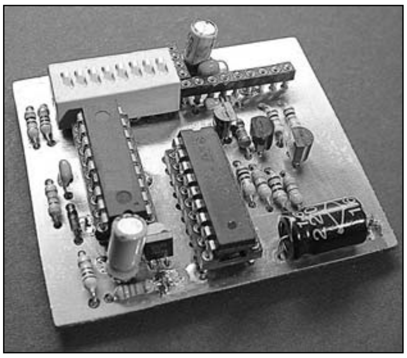

Fig 8.7: The Timer board, which will fit in a small corner in most transceivers and get rid of those DC switching relays

15. At this stage I lightly spray the board with SK10, which is both a protective lacquer and a flux. Available from Rapid Electronics, this makes for clean soldering and prevents contamination of the copper.

SUPERVOX

I have always operated VOX and, indeed, for many years did not bother to fit any PTT capability at all. Perhaps it would help the psychology if that switch on the microphone were known as "Lift To Listen'. The system has two elements, namely a timed solid-state switch for the various DC T/R lines and any relays, and an intelligent VOX system, implemented within DSP. But note that the Timer Board has been designed as a flexible and stand-alone general solution to manage the T/R switching in any transceiver.

Prerequisites

To be effective, any VOX system needs both T/R transitions to be free from clicks and thumps - electrical and mechanical. The first step to achieving this is to leave the maximum amount of circuitry powered up on both transmit and receive. Certainly all DC switching should be solid-state (hence the Timer Board), but the RF changeover can be more of an issue, especially at higher power levels. I use the circuit published in the 1988 *ARRL Handbook*, but hope to do better before the end of this series. Even if you use relays there are still significant benefits, though personally I just hate those acoustic rattling noises.

Timer Board

The T/R timer board circuit is shown in **Fig 8.6** and manages both the R to T and the T to R transitions. A completed board is in **Fig 8.7**.

The benefit of this approach is that the two transition sequences and timing can be - as indeed they should be - different. This cannot be achieved with the typical window comparator approach.

This board has one significant input, namely T/R Status from the DSP Assembly. This is a +5V logic signal (or floating) when on receive - and grounded to 0V to switch to transmit.

The PIC is a 16F627 which has the benefit of not needing an external crystal if timing accuracy requirements are modest; as a result the two crystal pins can be used for digital I/O purposes.

There are five timed outputs - which have been arbitrarily named for their most obvious general use. The ‘External Linear’ and ‘Local PA c/o’ lines are grounded on transmit, can sink 500mA -and are thus suitable for relay or solid-state switch control. The other three lines are at +12V when active and are explicitly grounded otherwise. They can each source/sink 500mA. They behave as follows.

Receive to Transmit (R/T)

The sequence for this transition follows, each step being followed by a timed delay:

External linear to Tx	T1
Local PA c/o to Tx	T2
12V Rx off	T3
12V Tx on	T4
Tx PA bias on	

There then follows a re-triggerable hang time, T5. This whole transition is not interruptible, see below.

Transmit to Receive (T/R)

For this transition the sequence is:

Tx PA bias off
External linear to Rx
Local PA c/o to Rx T6 12V
Tx off T7
12V Rx on

This transition is interruptible after step 1. That is, if you are part way through dropping back to receive when a transmit demand occurs, the T/R sequence will be aborted and the R/T sequence executed immediately. This interrupt logic is based on the view that it is always better to risk losing a moment of reception than to risk 'hot' switching.

Note that the 12V Tx and 12V Rx lines can never be energised at the same time.

Following a T/R transition, the PIC goes to SLEEP; that is, all dynamic activity ceases including its internal clock. Thus it can never act as a noise source to your receiver.

The process for adjusting the times for your installation will be covered later.

Construction Notes

Fig 8.8 (in the Appendix B) shows the PCB artwork ready for the iron-on process described earlier. The 10V regulator chip IC14 provides power to the STAR DSP board and may be omitted (as in the photograph) if you don't require this unswitched rail. Mounting holes for IC14 and the board (optional) have not been specified.

Start by fitting IC13, C76, C78 and C80, soldering one lead to the top ground plane. Then fit the socket for IC11 and solder pin 5 to the ground plane, followed by the socket for IC12 with pin 9 grounded. The remaining construction sequence is not critical. Mount the otherwise symmetrical switch so that, with the switches set away from the PIC, they are open circuit.

SUPER VOX IN DSP

Conventional VOX

When VOX detects the beginning of your speech, it initiates the R to T transition - which is going to take at least 3ms to complete. Further, if you have a T/R relay, the design must ensure the relay has settled in the transmit position before letting the RF through, typically adding a further 20 - 30ms. The result is that the leading edge of your speech is clipped off. Not by much in a good design, but often noticeable.

To disguise this effect, a VOX hang time is incorporated which is set to drop back to receive if you pause for breath, which at least minimises the number of truncated words. The other workaround you often hear from VOX operators is that they do not answer a direct question with - Yes". They tend to say "um, yes", probably subconsciously in order to avoid it coming over as merely "-esss".

How much better it would be if your transceiver started the R to T transition in anticipation, ie just before you started to speak! Sounds fanciful? In effect, this is what Super VOX does. And by the way, it applies equally to QSK CW operation.

Super VOX

The idea is to trigger the R to T transition immediately on detection of your voice, but then delay the 'voice' in DSP for the time it takes for the transition to complete. Thus the leading edge can never be clipped off.

Critically, this means in turn that you need no hang-time, since there is now no desire to minimise the number of transitions. Of course your delayed voice is still coming 'out of the antenna' for a few milliseconds after you stopped talking so you need to stay on transmit for that time - but absolutely no longer.

The net effect is that at a normal conversational speaking speed, you drop back onto receive not between breaths and sentences, but between every word - and often enough, between syllables, and if your T/R transitions are fast enough, you can listen through. Equally, someone listening to your transmission would be totally unaware that you were spending a significant percentage of your over on receive - in short, but very frequent, bursts.

The overall effect is very close in sensation to full duplex as in a normal (and therefore interruptible) conversation and, if widely practised, would do much to turn many a contact into a conversation rather than a series of speeches.

Anti-VOX

This normally works by comparing the microphone input with the speaker output - and if the same, concludes that it is not you speaking. STAR incorporates a further refinement in that the microphone input is compared with the output that did come from the speaker 4ms earlier. Why 4ms? Because this is the time it takes sound to travel 4ft in air, an assumed reasonable distance between the speaker and microphone. The improvement is noticeable and is worth having because the few extra lines of code don't cost anything.

AGC Implications

Normally, the AGC voltage decays to nothing shortly after you go to transmit. The result is that the receiver comes back on full gain in VOX gaps - which is not very comfortable in an 'S9' contact.

The approach adopted by STAR is to retain the AGC level established by the last 2s period of continuous receive - and apply that level during the gaps. It is important to ignore AGC levels established during the gaps for this purpose, so 2s was chosen as an arbitrary interval which is clearly longer than a casual pause. If, at any time, somebody other than you starts speaking, the normal AGC attack takes care of any adjustment in a few milliseconds.

So this is, if you like, extended-hang AGC where the 'hang' is extended over periods of transmission.

DSP BOARD

The board is based on (and is not incompatible with) the Analog Devices 2181 EZLITE board. That is, aspects of that board which are not used either by STAR or by W7PUA's DSP-10 have been omitted; the physical construction is completely different, and a current production and superior CODEC chip has been used. But conversely, the EZLITE board can be (and has been) used in this application and, if you already have your hands on one, e-mail me [2] for further details. Signal names as defined by Analog Devices are used throughout.

Mother Board

The mother board with her two daughters is shown in **Fig 8.9**. This form of construction was adopted to spread the risk during board manufacture and to allow upgrade of either the CODEC or Processor chips later.

Each board has its own regulator chips to spread the heat dissipation and to maintain modularity.

The CODEC daughter converts analogue signals to/from digital/analogue form for the benefit of the Processor. The digital signals are passed back and forth using a 12.288MHz industry

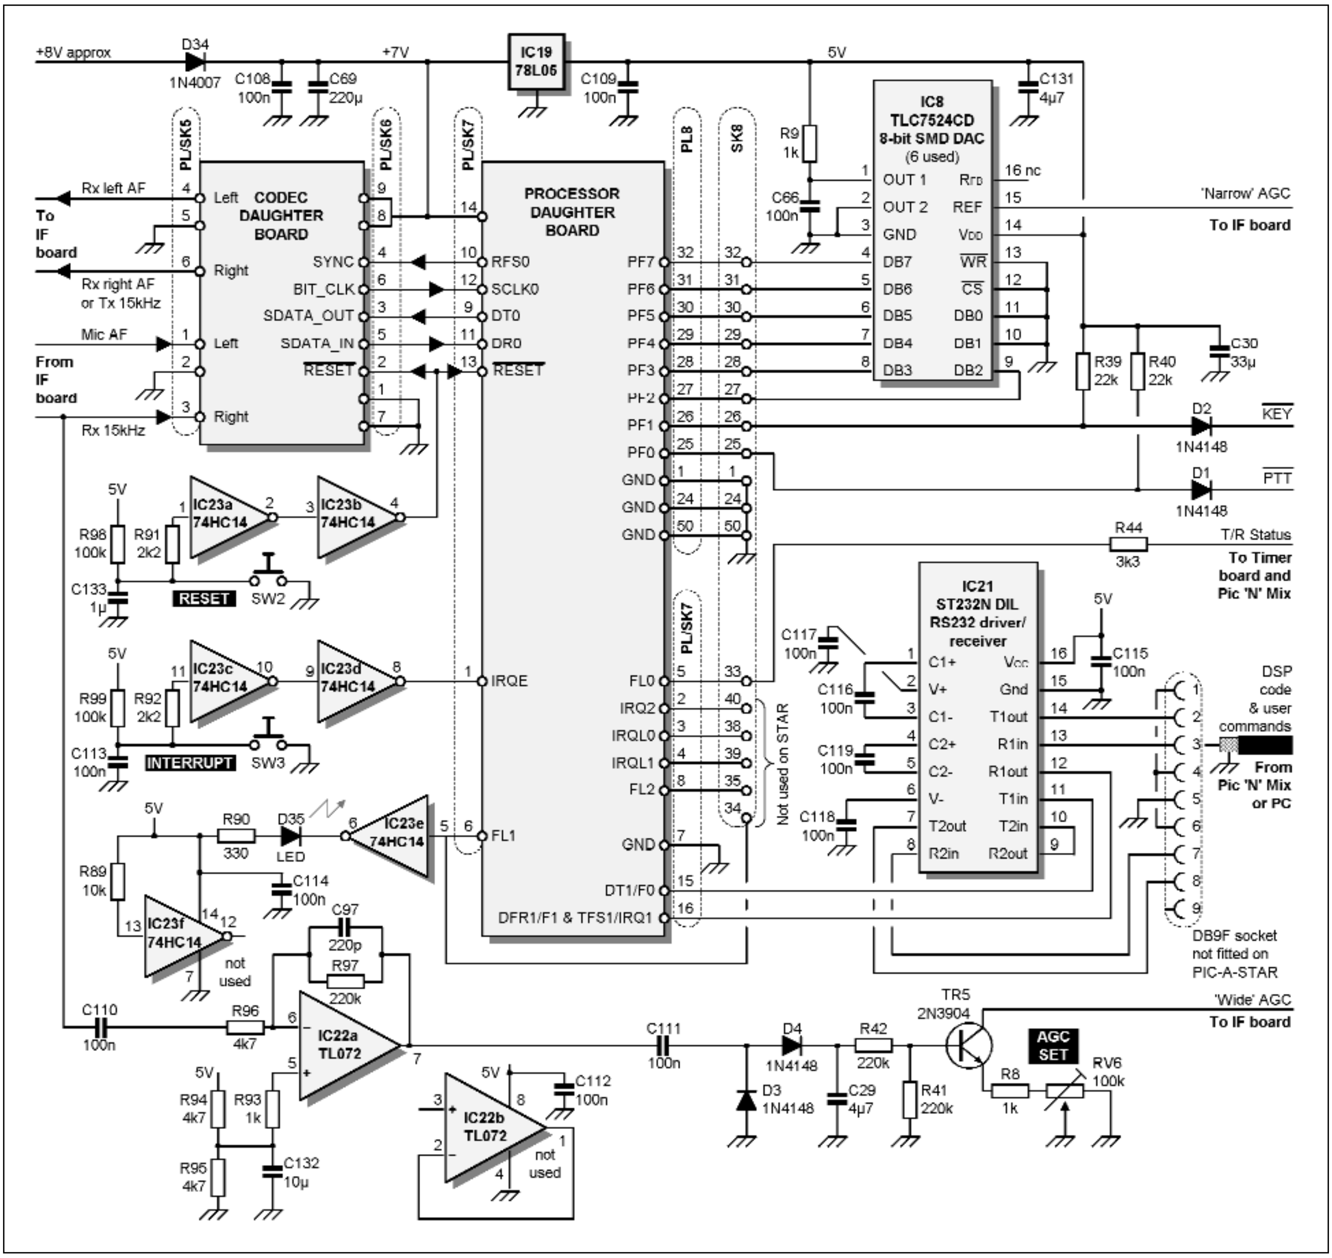

Fig 8.9: DSP mother / daughters relationship - and mother board circuit diagram

standard AC '97 serial bus -which multiplexes data in, data out and commands.

The Processor daughter does the DSP processing (no surprises there) -but has other control inputs / outputs as well.

Unlike the EZLITE board, the Mother Board carries IC8 and IC22/TR5 for generating AGC voltages for use on the IF Board accepts inputs from KEY and PTT lines -and generates the controlling system T/R line as a function of mode and control parameters eg VOX/QSK operation -thus customising it from the general to this particular transceiver application.

IC21 controls RS232 communications from a host - either your PC or the PIC in Pic 'N' Mix - and is used to upload the operational DSP code. It also accepts user commands to control the entire transceiver. IC23 buffers manual resets and interrupts , and drives an LED to show status.

Processor Board

This comprises the processor chip -and some memory used only at power-on (or Reset) time to boot load the real operational code. See **Fig 8.10**. For further detail, see the ADSP-2181 data sheet [3]. Being mostly track, the board is very quick and easy to build.

CODEC Board

This is a standard (albeit minimal) implementation of the AD1885JST CODEC chip. See **Fig 8.11**. For further detail, consult the data sheet [3]. The CODEC uses a 3V3 digital rail, but 5V on the analogue side; IC20 translates the 5V logic signals from the Processor to this 3V3 level. Outbound 3V3 lines to the Processor are already within its logic 1/0 definition range.

COMPONENT SPECIFICATION

SMD components have been specified here where space, cost, or performance considerations requires them - but not otherwise. 1206-size devices are used and these are no more difficult to handle than conventional leaded components.

Specifically, SMD electrolytic capacitors are not used because these are expensive - and the small space savings

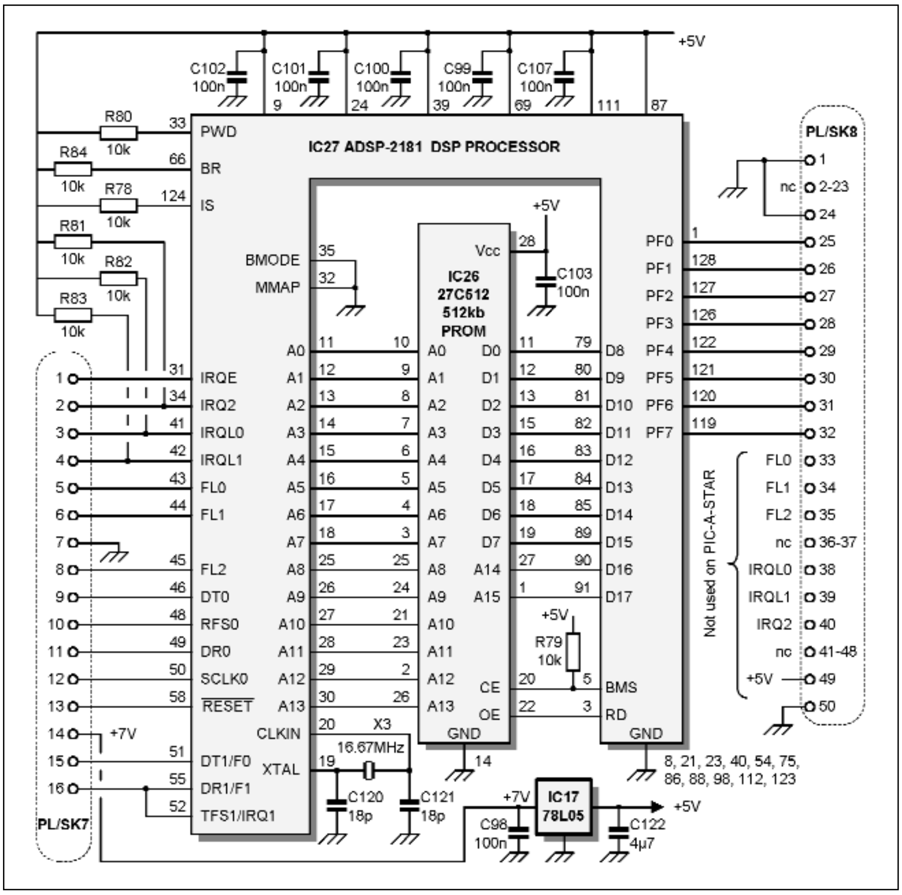

Fig 8.10: Processor daughter board circuit diagram. IC27 may be a KS-130 or KS-160 processor but, in any event, the slower KS-130 device is assumed. To retain compatibility with EZLITE, the unused or unconnected lines on PL / SK8 are available on the mother board for non-STAR applications

which are achievable are not necessary.

All the small coupling and decoupling capacitors are 100nF and, in general, they are SMD. However, on the CODEC board, C85-C89 are specified as wire ended disc ceramic units because their leads are used to couple power and ground between the two sides of the board.

SMD resistors are used throughout, since these save a great deal of space. The single exception is R78 - where a larger component was positively needed to span an otherwise unbridgeable gap.

In any event, all components are mounted on the track surface, but in some cases, leads are also soldered underneath. You need not take this to extremes, but every reasonable opportunity should be taken to interconnect the top and bottom grounds.

'EZLITE' COMPATIBILITY

IT IS ANTICIPATED that this DSP board will find applications in other DSP projects. If you are contemplating this, contact the author of that project in the first place for the current status. This hardware is a functionally compatible subset. It has the same overall dimensions albeit with different connector locations. The address, data and emulation expansion sockets have not been provided on this board - and nor, realistically, could they be. PIC-A-STAR uses a different CODEC chip, which requires a different DSP code module to handle it. A source code shell for this is available on request.

Fig 8.11: CODEC daughter board circuit diagram. Note that the ground plane is split between analogue and digital to minimise noise.

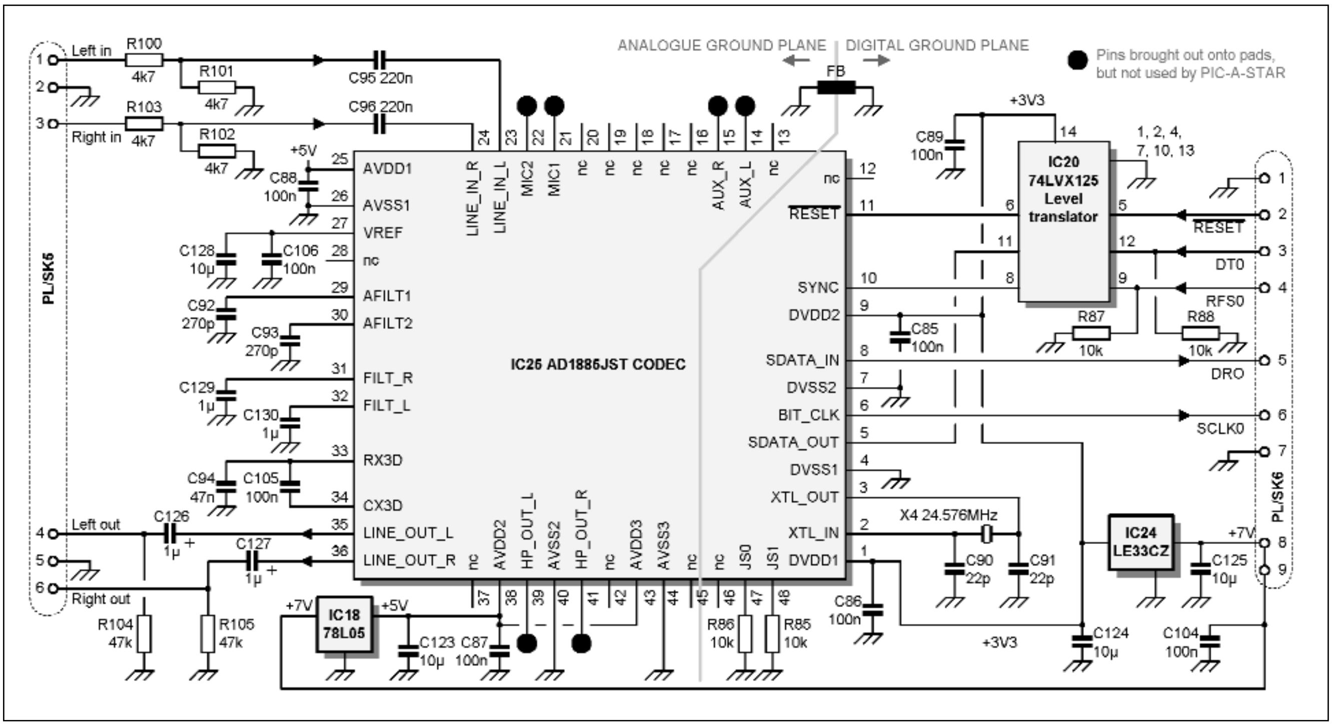

HARDWARE TEST

As the DSP board is progressively completed, it is highly desirable to test it in stand-alone mode before moving on. This process also proves the interface to your PC - which will be needed operationally later.

Prerequisites

The first requirement is that you are running *QBASIC* under Windows on your PC.

On older machines it is a standard application; later it was provided on the archive disc and on Windows ME it is not provided at all - but does run. In any event, it is an absolute prerequisite. The PC itself is totally uncritical.

You need to make up a lead from your computer's serial port - but only two of the lines are used. These are pin 3, the signal - and pin 5, the ground.

These connect (temporarily) to the mother board at 'DSP code and user commands' as per **Fig 8.12**. Ensure the ground lead is indeed grounded.

Set Up

On your PC, establish a new directory. The software assumes C:\STAR but you can edit the software for any other location.

In that directory, place the files testxx.xjp (where xx is the current version number of the test program) and XJPload.bas which is the utility used to load all STAR DSP software, not least this test program.

Open *QBASIC* and from there, open XJPload.bas. To run the test program, just follow the on-screen instructions!

Processor Test

This requires the Mother board with Processor daughter - but not necessarily the CODEC. The test process starts with D35 flashing at 1Hz. Once you start to load the test program, the LED will be permanently lit.

Once the program has loaded, the LED will be off if the CODEC was successfully initialised, on if it was not (particularly if it is not yet fitted). In either event, if you press the Interrupt switch, S3, the LED will toggle on and off. This verifies that the

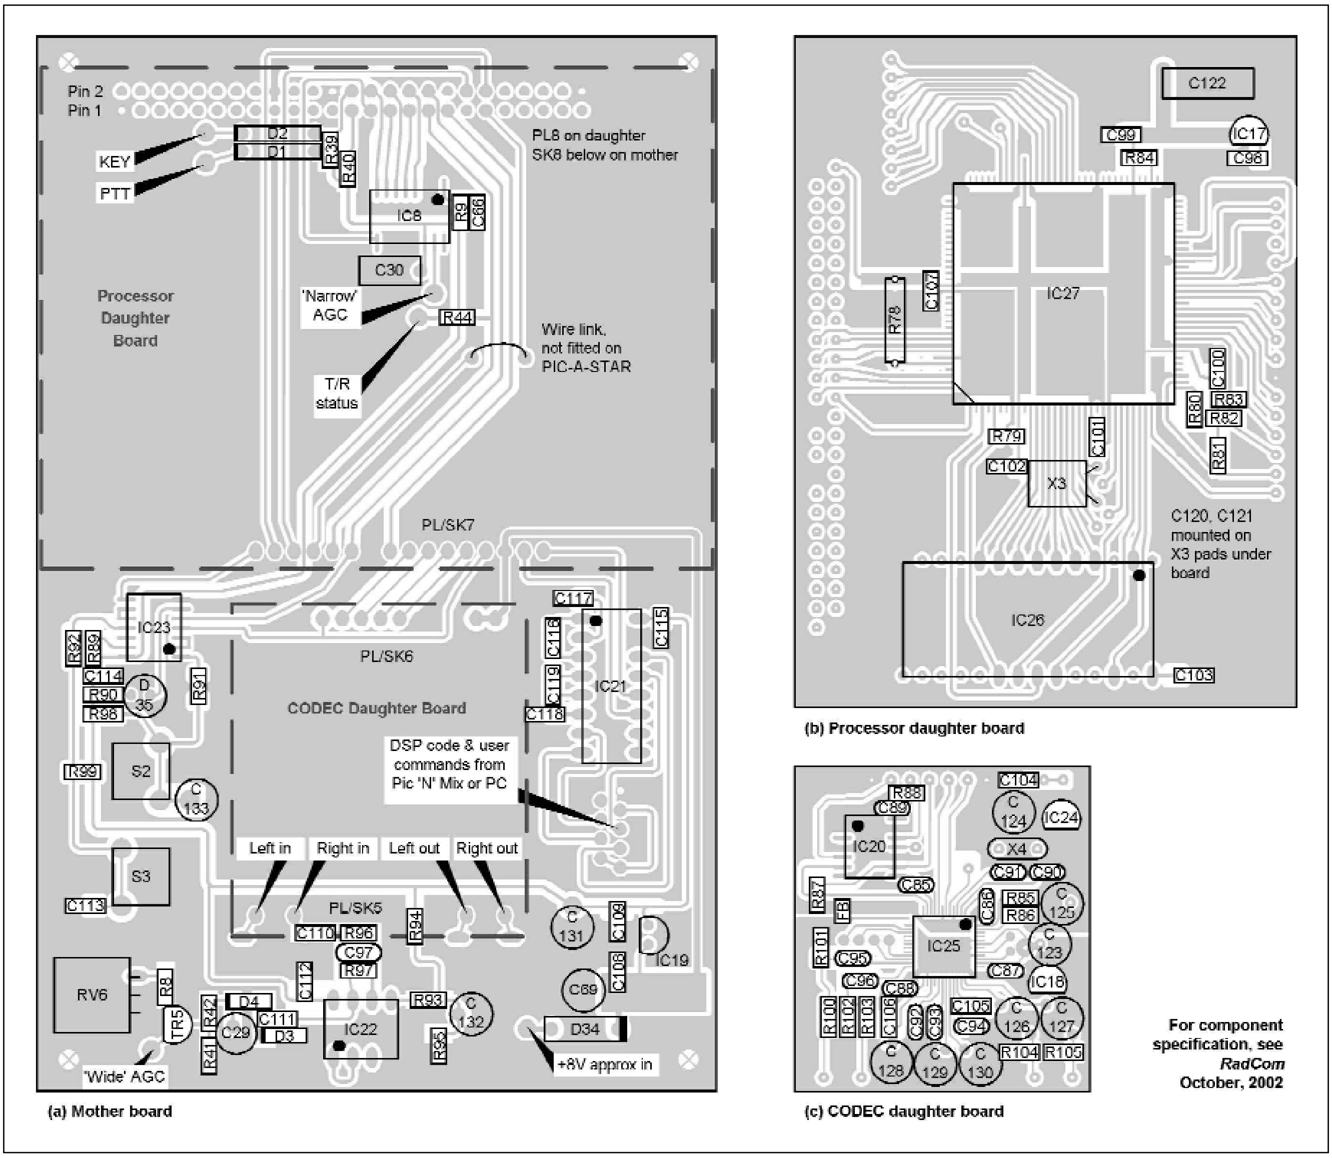

Fig 8.12: DSP boards component location diagram. SMD capacitors are shown as rectangles, disc ceramics with 'rounded' ends. The Processor daughter should be rotated clockwise through a right-angle to visualise the fit on the mother board. A significant number of pins on PL8 and SK8 are not used by PIC-A-STAR, but were included for compatibility with Analog Devices EZLITE board. These locations need not be populated for STAR

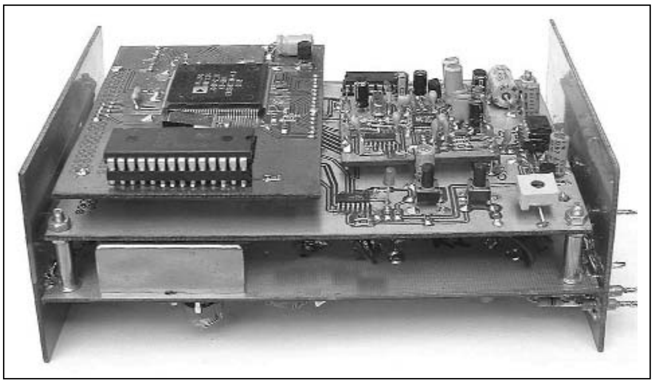

Fig 8.13: The DSP assembly. That is, the DSP mother board with CODEC and Processor daughter boards - mounted back-to-back with the IF board in its enclosure. The top, bottom and side screening panels are not fitted until after final test

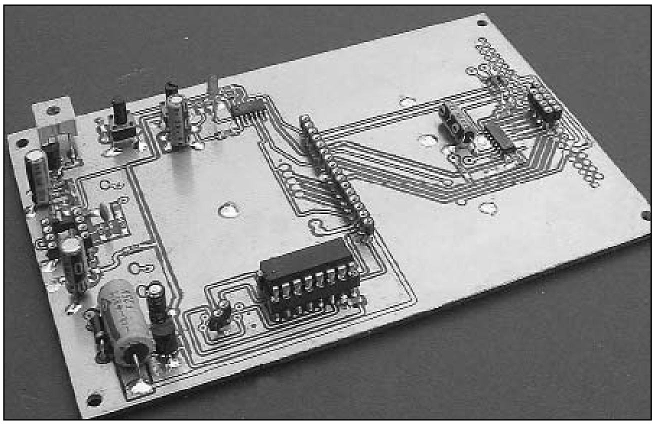

Fig 8.14: The DSP mother board, ready for daughter board fitting and test. Note IC8 and its associated components are located under the Processor daughter board. In fact, neither they nor IC22 need be fitted for stand-alone testing

code has loaded and that the processor is running and is in (or indeed, under) control. This also establishes your capability of loading any code over the serial link and unless and until you can achieve this, no further progress can be made.

CODEC Test

Once the CODEC daughter has been fitted and the previous test successfully repeated, power down and connect a patch lead from the CODEC left and right outputs to a stereo amplifier.

Power on again and reload the test program. A damp finger placed on the CODEC left or right inputs should now produce a corresponding hum on the respective output. Should you prefer something more exciting, you could connect up a microphone or any standard line-level stereo input. This is a test of a full loop-back on both channels. That is, the input is being digitised, sent to the processor where a minimal operation occurs in the digital domain before it comes back to the CODEC, where it is converted back to analogue form and thence to your ears.

Thus, when this test works, you have completely proved the CODEC and the vast majority of the processor functionality - and the interface between them.

If, however, it should fail, yet the processor successfully loaded the test program in the first place, the problem almost certainly lies on the CODEC board itself - or the link between it and the processor.

A 12.288MHz clock train is generated by the CODEC on the Bit Clock line. In response, the processor provides a 48kHz clock on the Sync line. If these are both present and you can see data pulses on the Data in / out lines, then the problem is probably on the analogue side of the CODEC. But, if you rigorously checked the board in the first place, there can't be a problem, can there?

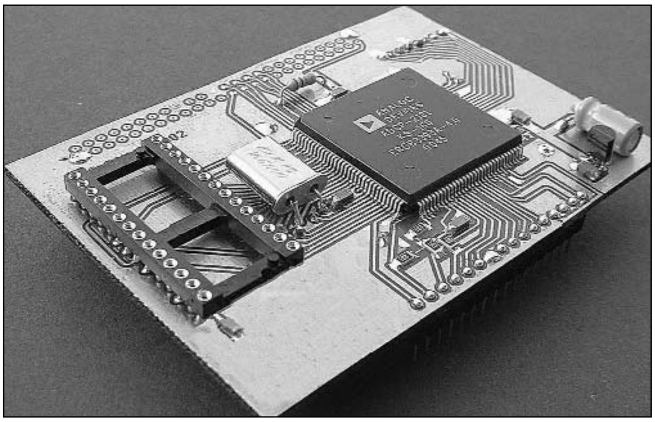

Fig 8.15: The finished Processor daughter board. Note that the crystal X3 is fitted after bending its leads - to reduce height. C120 and C121 are fitted under the board. IC26, when fitted in its socket defines the overall height of the complete DSP board

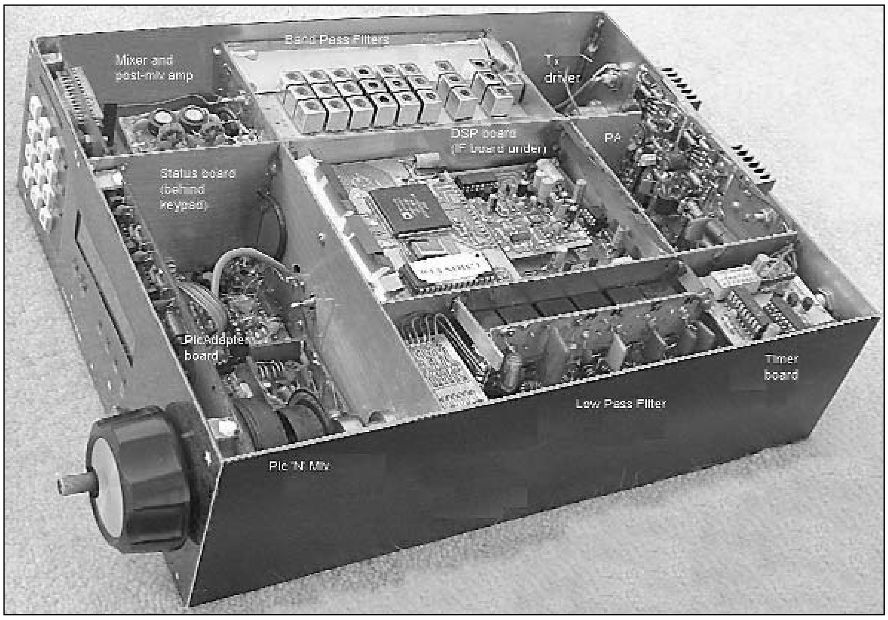

Fig 8.16: G3XJP's STAR built in a PCB enclosure - shown with all compartment cover-plates removed. The overall dimensions of the case are 310mm deep by 240mm wide by 85mm high. This generous size allows good in situ access to all the boards

DSP CODE DEVELOPMENT

If you want to develop your own code, you will need the tools. The author's code was developed using Analog Devices' older DOS-based development tools - which used to be supplied with its EZLITE board. These are available from its FTP site [8]. Nowadays it supplies its V*isualDSP++* environment which has the merit of a 'C'

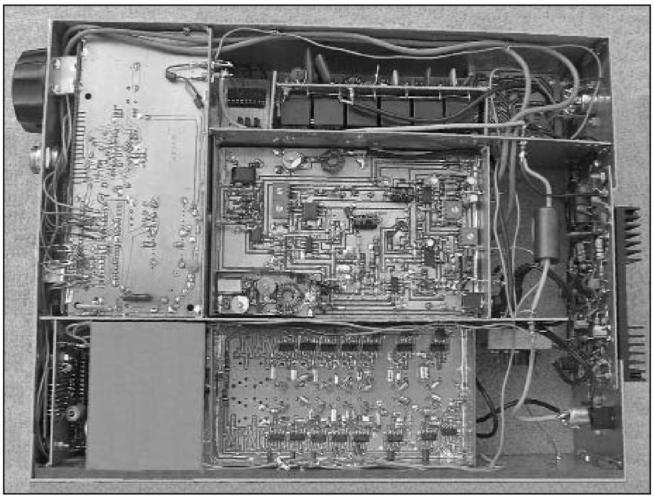

Fig 8.17: The view from underneath, traditionally somewhat less beautiful - so shown smaller

compiler. As an evaluation package, it also has a program memory limit but, at the time of writing, STAR would only use about half this limit. This world can change very quickly, so visit the AD site for the latest information.

DSP BOARD CONSTRUCTION

Although targeted specifically at the STAR DSP board, the technique for mounting the chips is totally general. There are no special tools required to mount these .difficult. chips - except a positive attitude. The author has heard much moaning about how these chips spell the end of home-brew construction - but it turns out the opposite is true. You can lay these chips down with a minimum of histrionics, and the following process - which is completely repeatable - came from AA7QU.

Tools

Firstly, the soldering iron. I used an Antex CS series iron (17W) with a 0.1mm tip, filed back from a mere point to a small chisel. Any bit about 1-2mm is fine. The other ingredients are:

- laser film, Farnell 895-945;
- some common solder;
- desolder braid, 2.7mm or less;
- a flux pen, Farnell 891-186;
- jam (home-brew, of course) or toothpaste.

The latter is for holding the CODEC chip in place long enough to tack its legs down. In fact, any water soluble non-setting stick is fine.

Construction Sequence

Make all three PCBs first as per Fig 8.20 (in Appendix B) using the iron-on process previously described.

The daughter boards are double-sided but, by design, only just. Under all circumstances, treat these as two-pass single sided boards. Any attempt to etch both sides in one pass is simply taking unnecessary risks. Do the complex topside first. If you want to use the artwork for the second side, drill all the holes, register the artwork with pins through those holes and then iron it on. But much easier, just sketch the trivial track and ground-plane in with an indelible pen, joining up the dots. When etching either side, merely spray mask the other.

When you have fully etched a board, absolutely check every track for continuity or shorts, either inter-track or to ground. If you get an open-circuit track the likelihood is that it will merely not work till you find the problem. If you have shorted tracks, however, the likelihood is that you will cook a chip and never find the problem.

If you have not used SMD Rs and Cs before, just tack one end down crudely, while holding it in position with a vertical screwdriver. Then solder the other end properly - and then revisit the first end.

Mother Board

Build this first, less the daughter board sockets. This board is completely unetched on the reverse (ground-plane) side. The only points to watch are the sockets for IC21 and IC22. Cut all their pins back to the shoulder except the grounded ones, which are soldered both sides. Check that all the obviously grounded areas on the board are indeed continuous and if not, add links through to the ground-plane side.

For the external connections, I simply countersunk the holes on the ground side, soldered stub wires to the pads on the track side - and then applied epoxy resin on the ground side to fabricate instant feed-through insulators.

Processor Board

Fit the inter-side links first. Then check the integrity of the tracking.

IC26 socket comes next. Cut back the pins which solder only to the top track; note that, exceptionally, pins 14 and 28 are soldered both sides.

Next the processor chip. Although it has more pins than the CODEC, it is somewhat easier to mount, since the pin spacing is greater and the chip is quite heavy so it is less inclined to skid around. The target time to mount this 128-pin PQFP chip is no more than 15 minutes - or you are doing something wrong!

Line the chip to the pads. Please check the orientation as you only have a 25% chance if you leave it to luck. The good news is

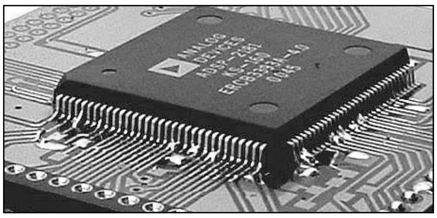

Fig 8.18: Processor chip before . . .

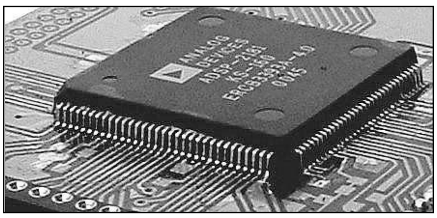

. . . and after removing excess solder. The target time to mount this 128-pin chip and clean up is 15 minutes

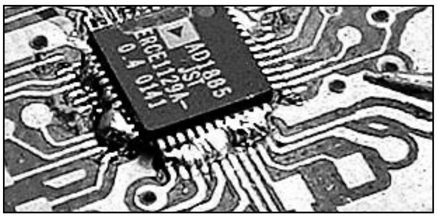

Fig 8.19: The CODEC chip before . . .

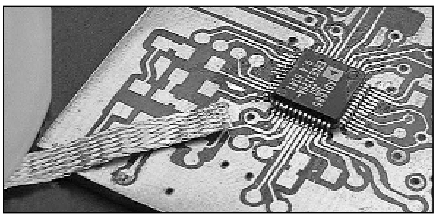

. . . and after desoldering. Don't panic, it works!

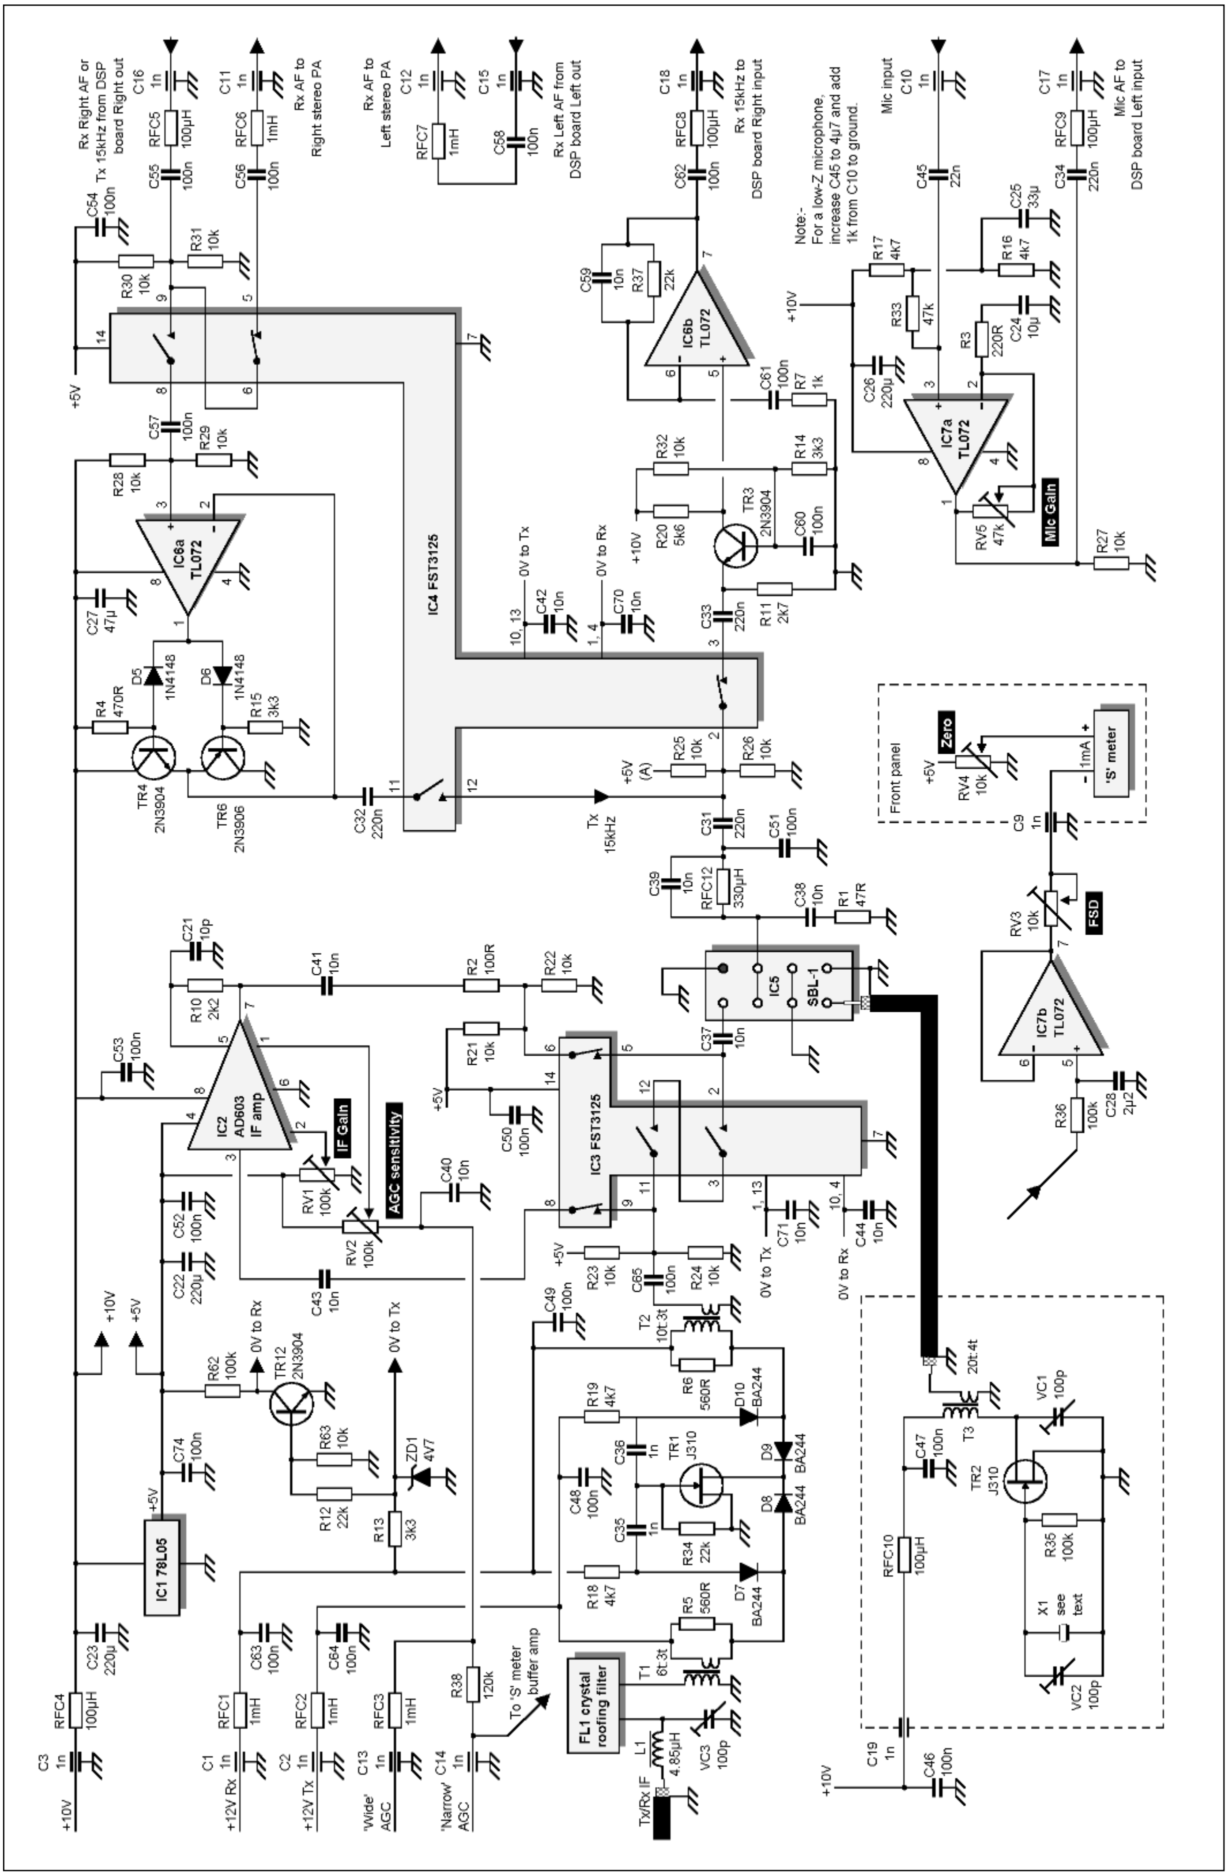

Fig 8.21: IF board circuit diagram. Takes an HF IF feed from a typical bi-directional mixer and post-mix amplifier, and translates it to/from 15kHz. This board mounts back to back with the DSP board. All components are mounted on the track side except for the crystal filter, FL1, and the SBL-1. The roofing filter is your choice and X1 must correspond. The switches in the receive path are shown as closed for illustration purposes only

that the correct quad-pack chip location on the board is totally unambiguous. Get someone else to hold it down while you roughly tack down a few legs in the middle of each side. It sounds cruel, but trust me, it feels no pain.

Running the iron and solder along each side at the point where the pins meet the track, run in a fillet of solder paying (almost) no attention to bridging the pins or the tracks. The only requirement at this stage is that every pin is indeed soldered to its track.

Three minutes elapsed.

Saturate some desolder braid with flux. Rest some fresh braid - over the top of the chip - on the bridged pins. Lightly apply the iron to the braid and, when you see the solder appear on the braid, withdraw. Then repeat as needed. Lay the braid on any bridged tracks - and repeat until all surplus solder has been removed. Do not draw the braid across the tracks, only along them.

Eight minutes elapsed.

Using a continuity meter, preferably with a 'beep' - and fabricating some probes from sewing needles - check that all bridges have indeed been removed. Finally, wash off any surplus flux under tepid water, and air dry.

Job done, seven seconds per pin.

Note that C120 and C121 mount on the pads of X3 on the underside of the board. Use SIL plug strip for both PL7 and PL8 - but use only the minimum population needed for the latter. Ensure the smaller diameter end of the plugs mates with the sockets.

For the sockets on the mother board, cut back the pins - except the grounded ones which solder both sides.

Fit the connectors dry to both the mother and daughter to ensure alignment - and then solder them to their respective boards.

The partial assembly may now be tested. Apply 8-10V power to the mother board and check the voltage rails before and then after fitting IC21 and IC22. Then plug in the Processor daughter and, after power up, D35 should flash at about 1Hz. Pressing the Reset button, S2 should cause a momentary hesitation before the flash resumes. Now run the test program, as described earlier.

CODEC Board

Having established that the digital and analogue ground-planes are mutually isolated, fit a wire link via a ferrite bead (FB) to join them. Then mount the CODEC chip as described for the processor, but in this case, use a *very* small amount of jam to hold the chip in register at first.

Then fit IC20 and C85-C89 and check integrity of ground and power. The other components should be mounted working outward from the chip, leaving the electrolytics till last. Finally, after rigorous checking and probing of every pin and every track (it must be right first time), I mounted the daughter to the mother using short lengths of component lead. In the case of the left and right inputs and outputs, their leads pass right through the mother board.

With this approach, should you ever want to remove the daughter subsequently, cut each wire first and then desolder both ends.

The complete board may now be tested by again running the test program, details of which were given earlier. The pleasure and pride of success at this stage is indescribable! .

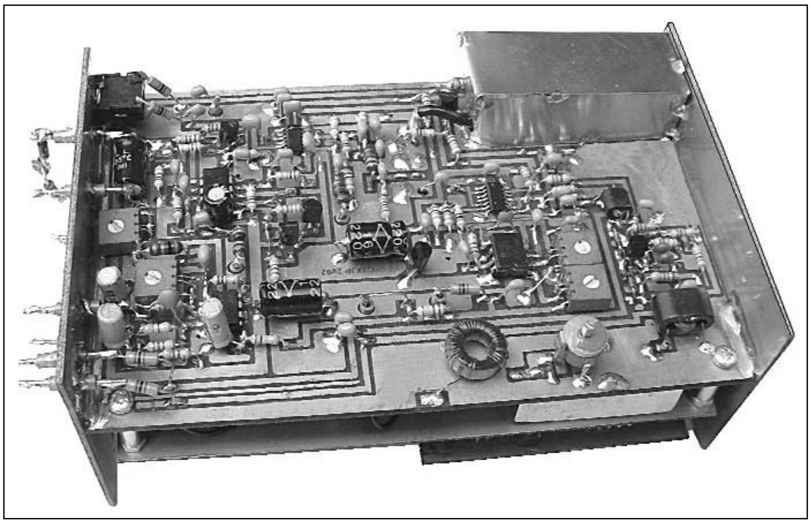

Fig 8.22: The IF board

IF BOARD OVERVIEW

The block diagram was shown in Fig 8.2 at the start of this chapter. The board has a bi-directional IF port - which is then translated to / from 15kHz where the DSP takes over both on transmit and receive.

The IF frequency can be at any HF frequency of your choosing, typically in the range 5 to 12MHz. The determinant is the availability of the crystal filter FL1.

Pic 'N' Mix [7] allows you to change the IF frequency injection offset in a matter of seconds, so there are no issues there.

IF CIRCUIT DESCRIPTION

Referring to **Fig 8.21**, the 50-ohm IF is matched to FL1 by L1/VC3. This is a standard L-match and should be modified - applying the textbook L-match equations - for your filter's frequency and impedance. VC3 is adjusted for maximum output in the first place, but thereafter for minimum passband ripple. The turns ratio of T1 also needs to be established for your filter impedance. The values given assume a 10.7MHz filter with 2200Ω impedance.

TR1 is the ubiquitous bi-directional J310 IF amplifier. It offers modest and quiet gain, and stable load and much convenience.

IC34 and IC4 provide fast (and silent) T/R signal switching - with high isolation and only a few ohms on-resistance. Resistive divider networks are used throughout to bias the signal paths to mid-rail.

On Receive

IC3 routes the signal to IC2, the AD603 IF amplifier. This is a quiet device with good AGC characteristics. As used here, it has a gain range from 0 to 40dB - with a linear dB response to a linear control voltage (an *increase* in control voltage produces an *increase* in gain). It is not the most inexpensive device available, but if you have been brought up on IF amplifiers that emulate snakes, you will appreciate the difference.

This 40dB AGC range is combined with a further 45dB in DSP to give 85dB in total - more than plenty by most standards.

The process for setting RV1 and RV2 follows later. The two AGC control signals ('wide' and 'narrow') are generated on the DSP board and summed at the junction of R38 and RFC3. The 'wide' AGC voltage is generated by detection over the full FL1 bandwidth - and is there only for emergency gain back-off in the presence of a very strong signal outside the DSP filter bandwidth. Normally the 'narrow' control voltage dominates - and it is

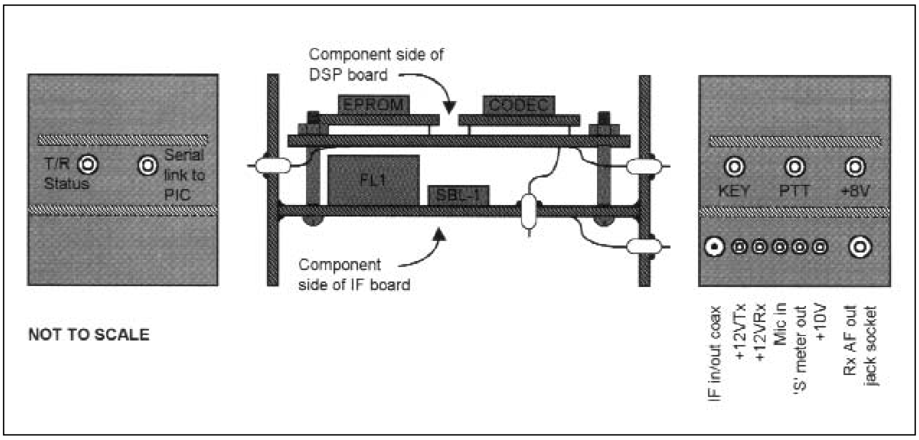

Fig 8.23: DSP assembly illustration and recognition drawing, not to scale. The IF board and two end-plates are seam-soldered to form an H-section. The height of the endplates is typically 6cm as a minimum, but can be up to the full height of the Tx / Rx enclosure. The DSP board is bolted to the back of the IF board. Note that critically, the external connections are brought out at different 'levels' depending on which side of which board they connect to - and at different ends depending on the destination. Two further sides and a top and bottom (not shown) complete the screening but are not added until after final commissioning

also routed via the buffer, IC7b, to drive the S-meter. This latter is shown as a 1mA movement - but later a bar-graph alternative will be offered - in which case RV3 sets the zero point and RV4 is not fitted at all.

The output from IC2 is routed via IC3 to the SBL-1 mixer, IC5. You may ultimately wish to fit a stronger device here - depending on the width of your roofing filter and your operating needs. The mixer injection port is fed from a basic crystal oscillator - and this could also be ‘beefed-up’ if required.

C38 and R1 terminate the sum (HF) mixer product - whereas RFC12, C39 and C51 pass the wanted 15kHz difference component.

TR3 is a low-noise, modest-gain amplifier which feeds IC6b. This latter has modest gain at low frequencies with the response rolled off rapidly by heavy negative feedback provided by C59.

On Transmit

IC7a provides modest shaping of the microphone audio and significant gain to get the level up to that required by the CODEC on the DSP board.

You should alter the input arrangements of IC7a to suit your microphone impedance - and C45 in particular for a good mid-range audio response with your voice. Some tailoring options may later be added in DSP as well.

The output of IC7a is routed unconditionally to one input of the DSP since it needs to monitor the mic input continuously for VOX purposes.

The transmit signal next appear as a 15kHz SSB or CW signal from the DSP which is routed via IC4 to the buffer, IC6a. This, in turn, drives the complementary pair, TR4 and TR6, which are there to deliver power into the low-impedance load presented by the SBL-1.

On transmit, the AD603 is out of circuit and IC3 routes the signal directly to the J310, TR1, and from there to the filter FL1, and thence out to your transmit IF strip.

T/R Switching Control

The J310 is switched by the 12V Rx and 12V Tx lines, the inactive one being taken to near ground. All other T/R switching is managed by IC3 and IC4. The switching voltages ae derived from the 12V Rx line only, with TR12 acting as a simple inverter. This approach is designed to prevent you from being on transmit and receive at the same time in the event of the loss of either the 12V Tx or 12V Rx supplies.

BUILDING THE IF BOARD

The IF board comprises a traditional PCB with two end-plates soldered on to form an H-section, as shown in **Fig 8.23**. The DSP board is subsequently mounted on the IF board as illustrated.

This form of construction is not strictly necessary. You could build the IF board and DSP board into two separate enclosures, but this approach was chosen because these two boards are highly interconnected.

The IF PCB dimensions are determined by the size of (and are just larger than) the DSP board - resulting in generous spacing between the functional blocks. The surplus board area has been allocated around the crystal filter and the crystal oscillator; the former so that any reasonably-sized filter may be fitted, the latter to give room for a more sophisticated oscillator if desired.

The PCB is assembled by soldering most of the components to the track side. This approach makes signal tracing easier and minimises the amount of hole-drilling. SMD components were not specified here because they are not needed, but most of the components are in fact mounted SMD-style.

Mask, etch and drill the PCB using the iron-on laser film technique covered earlier. On the ground-plane side, countersink the ungrounded holes associated with FL1 and the SBL-1. Both the coax lead to the SBL-1 and C37 are soldered directly to the SBL-1 pins - as opposed to PCB track - so drill generous clearance holes for these pins. All other holes are grounded both sides of the board and are not countersunk.

End-plate Dimensions

The width of the end-plates is that of the IF board. The task now is to determine their height - which is principally (but not entirely) determined by that of your crystal filter.

Fit FL1 and then, using spacers somewhat longer than the height of this filter, crudely trial-mount the DSP board as in Fig 8.23.

The height of the end-plates is now that of this assembly plus at least 20mm for the IF board components. The approximate sum is 24mm for the DSP board, plus 20mm for the IF board components, plus 2mm for the PCB thickness, plus the height of your chosen crystal filter plus 3mm margin. The latter two measurements also sum to give you the length of the four mounting spacers. Be generous.

End-plate Fitting

The end-plates are fitted before mounting the components, because this makes the board easier to build and handle without contaminating it with finger marks.

Mark the target position of both boards on the inside of the end-plates and then, looking at **Fig 8.24**, lay off the position of the feedthrough capacitors from the IF board. Drill holes for these now, but leave fitting the capacitors until later.

Clean both sides of the IF board and end-plates immaculately, and apply a light coat of spray flux / lacquer to both sides.

Now seam-solder the end-plates to the IF board with a large

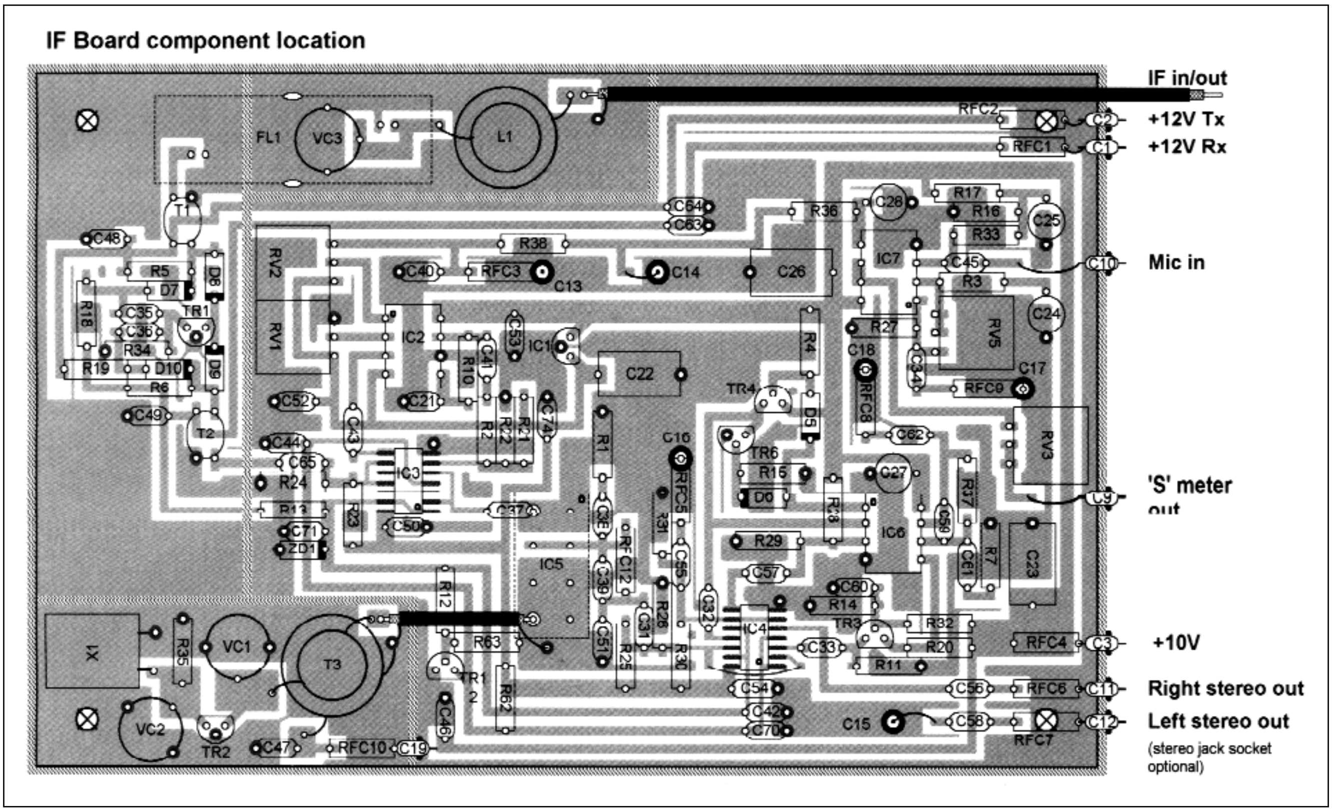

Fig 8.24: IF board PCB layout on double-sided board (see Appendix B for the PCB artwork). The reverse side is completely unetched to form a continuous ground plane and screen. All components with the exception of the crystal filter, FL1, and the SBL-1 mixer, IC5, are mounted on the track side. You may need to customise the tracking to suit your crystal filter. The 'holes' are shown on the component layout only to define the tracks should you be producing the PCB by some manual method. Only components which feed through to the back of the board require actual drilled holes and these are as defined on the tracking template. The tracking template image is mirrored (ie flipped left-to-right) for direct copying to iron-on laser film. The basic drilling size is 0.7mm - with holes for mounting, feedthroughs etc drilled larger to suit. Some internal screening partitions are made from PCB material or brass shim stock. Those

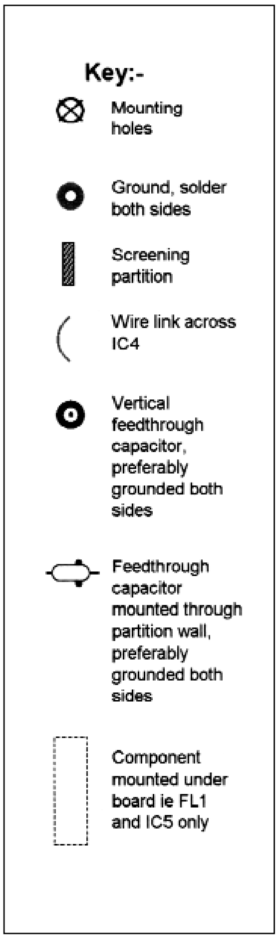

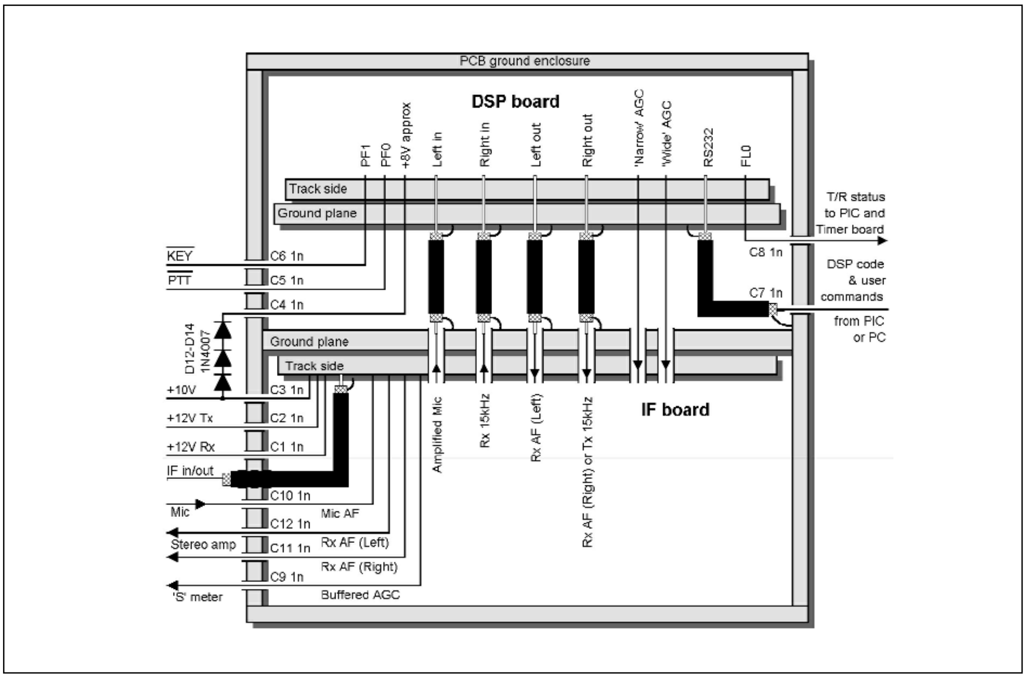

Fig 8.25: DSP sub-assembly. The IF board is bolted back-to-back with the DSP board using nuts, bolts and spacers. Both their ground planes form a screen to isolate the two halves of the box. Feedthrough capacitors are used to route between the two halves of the box - and to the rest of the transceiver. C11 and C12 may be replaced with a stereo jack socket

iron. If you mount both at the same time, you will be able to check on the geometry by eye. Progressively checking that all remains true and working both sides of the IF board, use small single tacks first of all, then multiple tacks and finally form neat fillets.

Component Mounting

Refer to Fig 8.24 and **Fig 8.25**. Tin all the pads except those under the FST3125s. Mount both FST3125 chips. Align the chip and tack down two opposite corners to the larger pads provided. For the remaining pins, offer the iron and solder to the track just short of the pin - and the solder will spread along the board and wet the pins by capillary action.

Fit the wire link across IC4. Cutting their ungrounded leads so that they sit just above the board, mount all the other components - except the preset capacitors. A pair of tweezers is useful for handling the smaller components.

To surface-mount the DIL ICs, cut off all the pins back to the shoulder except any grounded pins which pass through the board. Do not use sockets.

Fit the feedthrough capacitors ,which are typically made off to the IF board, by using a series RFC as a flying lead. Mount RFC1 and RFC2 at right-angles to each other in the vertical plane to minimise mutual coupling.

Trim and solder all the grounded leads on the back of the board. Check with a continuity meter that all grounded track is in fact grounded. Also perform all the usual basic tests such as checking isolation and integrity of the power rails.

Mask off the preset resistors and give both sides of the board a final and generous coat of spray lacquer. Finally, fit the preset capacitors, definitely unlacquered.

DSP ASSEMBLY – ASSEMBLY

This process starts when the IF board and DSP board are fully built, and the latter has been tested using the test program. The required DC supplies come from the Timer Board (or some equivalent arrangement).

Make off all the leads between the boards as shown in Fig 8.25. With the two boards at right-angles (but preferably less), trim their lengths and make off the other ends to their respective feedthroughs. Ground the braids to the adjacent groundplane.

Mate the two boards, and in the process, perhaps trim some excess lead lengths. Fit diodes D12 to D14 – outside the housing – to drop the 10V rail to a nominal 8V.

COMMISSIONING

The DSP Assembly is first proved in isolation and then crudely integrated with some existing transceiver for verification. The idea at this stage is to demonstrate hardware functionality, not system performance.

Basic DC Testing

As a preliminary, set RV1, RV2, RV5 to mid-travel and RV3, RV4 fully clockwise.

On the end-plates, connect up 10V, +12V Tx (grounded on receive), +12V Rx (+12V on receive) – and the stereo outputs, typically to some domestic amplifier.

For the first few seconds after power-on, a voltmeter on the S-meter feedthrough should show definite activity on a 5V range. Check that the T/R Status line is near +5V.

On the IF board, check all the power rails and then get the X1 oscillator working. Adjust its frequency to the centre frequency of FL1 + 15kHz.

Loading Test Software

Connect the serial cable from your PC's COM port to the DSP Assembly. Also, a microphone (both audio and PTT). Load the test program as previously described and re-verify operation.

Speaking into the microphone should produce audio from one stereo channel. Adjust RV5 for maximum undistorted output – but, in any event, no more than 2V peak-to-peak on C17.

Loading Operational Software

Reset the DSP board (ie press and release S2) and load in the operational software as per the loader onscreen instructions.

Loading is complete when you are looking at user controls on the screen as in **Fig 8.26** – and the DSP board LED is out. If the LED remains – or reverts to – flashing, this indicates a comms failure during loading.

At this stage the DSP Rx should be operational. To verify this, feed a sniff of RF at your IF frequency into the IF in/out coax. Just tack a few inches of wire to the coax inner and put it near some suitable signal source eg the DDS or a GDO. As you tune across the IF, you should hear the beat note, and the LED on the DSP board should light in the presence of signal.

Turn the RF gain up and down on the PC to verify that you are in control. If you speak into the microphone, this should also light the LED. Grounding the PTT line should mute the Rx – and the T/R Status line should go to near 0V.

On the PC, switch to CW. Grounding the KEY line should then produce sidetone.

The T/R Status line may now be connected to the Timer board and its operation verified. Under no circumstances be tempted to connect T/R Status to some external PTT line, say on your transceiver.

Getting to this point is a major milestone. But if any of the preceding fails, stop and correct the problem before going further.

INTEGRATION TESTING

This stage is not strictly necessary. You could wait until you have a completed transceiver. But I commend this as the better approach - not least because any problems will be confined to the new-build DSP hardware.

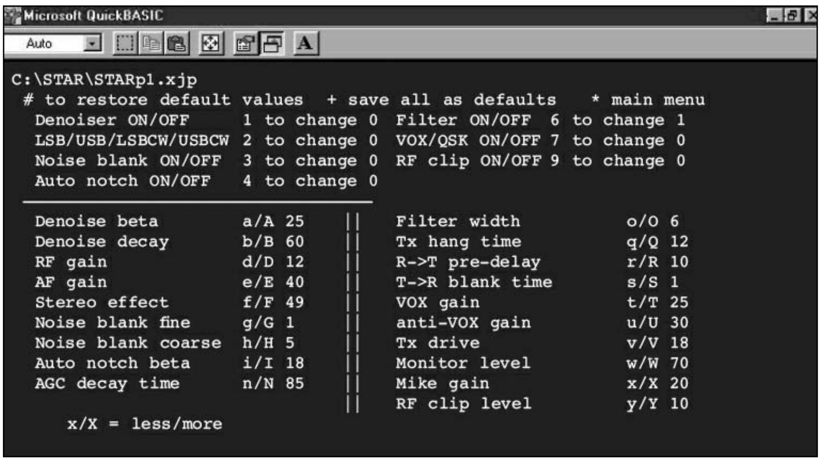

Fig 8.26: The PC screen running under QBASIC. Illustrated are the STAR parameters for the SSB mode. For development and proving purposes only this user interface is designed to be rather more functional than beautiful

Receiver

Connect a short fat ground strap from your transceiver to the DSP Assembly. Locate a suitable bi-directional 50-ohm point on your transceiver; after the mixer, after any post-mix amplifier, after any pad is best – but the 50-ohm IF port on your mixer will suffice for test purposes. Patch in the IF in/out coax via a series 100nF instead of your existing IF strip.

Arrange to be able to switch your transceiver between transmit and receive. Turn the AF gain down on your transceiver – and any other Rx gain controls to maximum; and Tx gain controls to minimum.

Power up on an low frequency band and then load the operational software as previously. Inject signal frequency plus FL1 centre frequency into your transceiver mixer – and you should hear resolved LSB signals from both speakers. On a quiet frequency (LED is out), peak VC3 for maximum band noise. Then peak RV1.

Connect a CRO (DC, 1V/cm) to the S-meter output. This should show about 4V on weak signals and progressively less as AGC action occurs. Find a signal giving about 2.5V and adjust RV2 until it is slightly less. While listening on a noisy band, adjust RV6 until the AGC loop is clearly unstable and hunting – and then back it off until it is smooth. That completes a crude setting up of the AGC system, enough to verify that the hardware is working.

At this stage, with the DSP Assembly unscreened, there may be evidence of white noise on the higher bands.

Transmitter Integration Test

With your Tx drive level well down, set Tx Drive on the PC to 10. With your transceiver connected to a dummy load, and preferably monitoring on another receiver, set up to observe the Tx output on a CRO for flat topping etc.

Put your transceiver on transmit. When you ground the DSP PTT line, this will put the DSP assembly onto transmit as well. The mic gain on the PC should be increased as far as possible – but only so long as there is no evidence of any clipping, compression or distortion.

If all is well, bring up the drive on your Tx to its normal setting. Then increase the Tx Drive on the PC, ensuring the output remains clean – up to your normal power level.

Now would be a good time to screen the crystal oscillator and add the other screens on top of the IF board. The fully-screened enclosure is best left until the very end.

PC CONTROL OPTION

This is a timely opportunity to outline the behaviour of the PC control panel. Fig 8.26 shows the screen of the loader after the DSP code and the controllable parameter values have been downloaded to the DSP assembly – at 9.6kB.

Adjustment and use of the various DSP features themselves follows later. Here, we are concerned only with the mechanics. Simply key the appropriate number to change a switch state, the upper-case letter to increase a parameter value – and the lower-case letter to decrease it.

Syntax

The BASIC control software has been optimised for simplicity. That is, only the most basic syntax has been used – and if you have ever written any software in any language (very nearly, English will suffice) – then you will have no trouble following it or editing it. Equally if you want to build some controller other than Pic 'N' Mix [7] – either in dedicated hardware or on your PC – then this acts as a model. If you have the background to undertake this, then equally you will have no issues following the code.

Control Parameters

These are held in a separate file, param01.xjp. It is the controller's responsibility to handle parameter values and to constrain them to be within maximum and minimum values, and in any event, within an 8-bit byte. The value 255 is assigned to any parameter that does not apply in a given mode.

Following any user change, the new parameter value is sent to the DSP as three bytes. The first is always a tilde '~', the second is unique and identifies the parameter, and the third is the new value. Nothing could be simpler.

Frequency Control

One of the virtues of controlling the whole transceiver from Pic 'N' Mix [7] is that it can handle the injection offset needed when switching between SSB and CW and transmit and receive. Obviously, the PC has no intrinsic ability to do this, so you need to make other arrangements – eg operate your Tx/Rx split when on CW.

PIC A TIME

In the STAR environment, the sole purpose of the Timer board (see earlier) is to provide click- and spike-free R/T and T/R transitions. Hang times for VOX or QSK operation are controlled by the DSP.

All the switching times may be independently set between 1ms (very fast) and 63ms (incredibly slow). In general, it is best to start with the times set to incredibly slow, and then reduce them progressively until there are any signs of switching spikes on the transmitted output – or clicks on reverting to receive. Having said that, if your transceiver has inherently noisy switching, a click on reverting to receive is inevitable. The DSP code has a feature for blanking any such click – but it is obviously best avoided by design.

Adjustment Process

The process for altering the timing is as follows, starting with all the switches OFF, ie away from the adjacent PIC:

1. Set the More/Less switch as required to increase or decrease the time delay. (More is towards the PIC).
2. Set the switch(es) for the time(s) you want to alter to ON, ie towards the PIC.
3. Key the T/R Status line down and up once for each required millisecond of change.

The altered time(s) will be implemented immediately – but not stored. When all the required changes have been made, put all the switches to ON, key the T/R Status line down/up one final time – and all the new times will be stored and retained. As evidence of success, this particular R/T/R transition sequence will not occur. Conversely, to abort all changes since power-on simply miss out this stage completely, power off and wait 20 seconds before powering on again.

Finally, set all the switches to OFF. Note that this process can be used to change several (but not all) of the time delays simultaneously though you may wish to avoid this practice unless gross changes are required. Note also that the PIC cannot be programmed via the programmer interface if the switches are ON. .

USER INTERFACE

The rationale behind the PIC-A-STAR User Interface (UI) has had significant impact on the front-panel layout and ultimately on the entire transceiver enclosure.

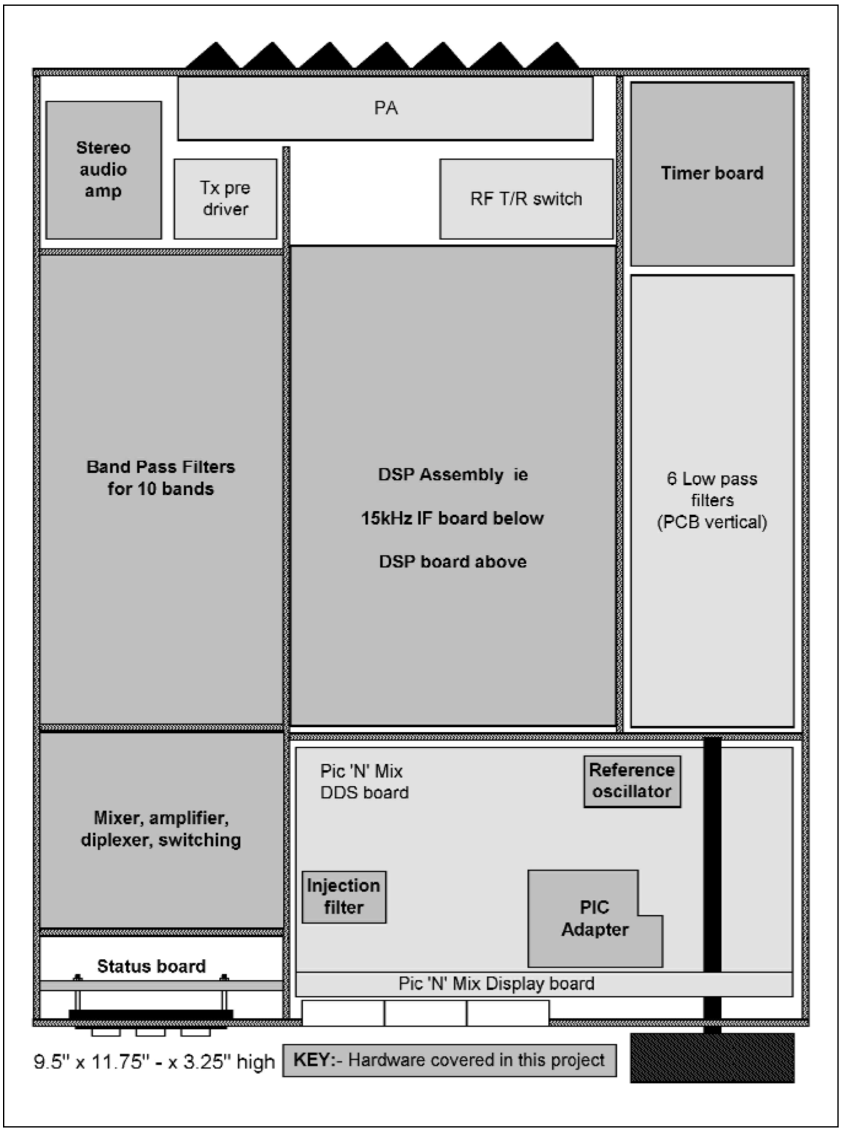

Fig 8.27: A possible transceiver enclosure illustrated at halfscale. The author's is fabricated from 2mm doublesided fibreglass PCB stock. This gives excellent access from both sides

Until recently, the author had followed a philosophy of fixing the front panel controls with those anybody could reasonably need to drive any transceiver, on the grounds that these controls were essentially independent of the inner workings. That approach bought me 25 years development of four fundamentally different transceivers - all using the same housing and front panel.

But times they are a-changing! I have come to appreciate that the opposite approach is more appropriate for a transceiver with a significant software content. **Fig 8.27** shows, not least, the consequences to the overall dimensions.

A significant amount of effort has gone into achieving an effective UI. The challenge is to avoid the extremes. On the one hand, it is easy to end up with a system whose complexity exceeds human intellectual capacity. We have all listened in on those amazing menu comparison contacts which usually end in "I know I am supposed to hit 'Return'. How hard?". On the other hand, personal preferences do vary and you shouldn't be prevented by the designer from adjusting a parameter merely for the sake of a simpler UI.

The clue to the best approach came very early.

No Pain – No AF Gain

During the early evolutionary development, there was, of necessity, a period of several weeks when I had absolutely no adjustable controls whatsoever. To increase the AF Gain, for example, I had to edit the DSP code, reassemble it and download the whole suite – including this one changed parameter. As you can imagine, I did not bother very often. Actually, it encouraged me to focus on improving the functionality so that the built-in control systems would take care of the 'variables' without undue manual intervention.

The background thought here is that as a design evolves – perhaps over several years – the number and purpose of the controls can swing wildly. And as technology evolves, so do the opportunities. Who would have thought I would need an Auto-notch on/off switch even 15 years ago? And when I designed Pic 'N' Mix, I can assure you the thought that it had the intrinsic flexibility to control the entire transceiver never crossed my mind.

So this time I have taken the minimalist approach - starting with the observation that all controls can be classified under two generic categories, namely "switches" and "amounts". Thus PIC-A-STAR has two (and exactly *only* two) corresponding physical controls, namely: a knob which alters 'amounts', and a keypad which handles the 'switches' – as well as specifying which 'amount' the knob is connected to.

By 'amounts' I mean, for example, amount of AF Gain, amount of RF Clipping and, indeed, amount of Frequency. By 'switches', I probably better mean 'choices' eg '80m' instead of '15m' and 'Autonotch on' as opposed to 'Autonotch off'.

Having settled the mechanical format, then at any time I can have more or less as many 'free' knobs as I like - by simply assigning them in the software. And at the same time, saving much cash on real pots, knobs, switches – and that real nightmare, the consequential system cabling.

Now that my STAR is in daily operational use – besides changing bands and frequency – the biggest strain on the UI has been turning the VOX off when the fast jets go over - and turning the transmitter power up and down to suit conditions. All the other controls have to be set up correctly, but most are essentially set-and-forget.

FRONT PANEL

The template used to make the front panel is shown in **Fig 8.28**. This is designed to last a lifetime in the sense that I can allocate any function to any switch - including a cluster of related controls as a simple sequential menu list. And thereafter, any range of values to the menu items – and so on.

The worst-case change issue is that one day I may need some new legends on my keypad overlay. I have never been an advocate of beautiful homemade radios versus functional homemade radios (given a finite life-time, you have to choose) – but one non-trivial benefit of this approach is that the front panel is less than A4 (and US Letter) in size – so a new one can be print-

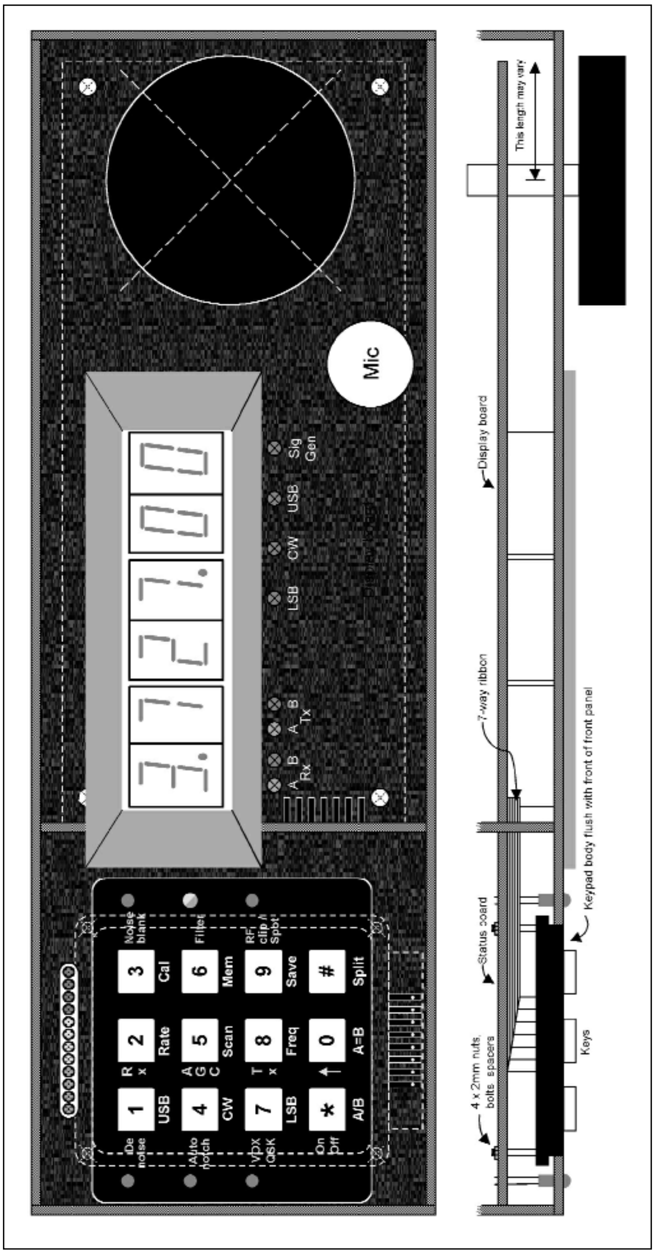

Fig 8.28: The author's STAR front panel layout, to scale. In this case a bargraph S-meter has been used. The nine most frequently used DSP control groups are assigned to the 1–9 numeric keys

ed off on photo-paper and stuck on anytime. Upholstery or foam backed carpet adhesive is the answer to your next question.

A bar-graph S-meter is shown, but is not mandatory. A small edge-wise movement could be accommodated above the keypad, but a 'real' one would need an increase in the front-panel width to accommodate. The bar-graph LEDs, six status LEDs and the keypad all mount on the Status board.

PICADAPTOR AND STATUS BOARDS

Next the circuits of the PicAdapter and Status boards. Constructional detail follows later. You don't need these boards to commission STAR initially, since you can load and control the DSP software from your PC.

Thereafter, in terms of constructional sequence, you need the PicAdapter first, which then allows DSP code download from your PC and subsequent upload to the DSP assembly. Thereafter, you need the Status board to complete the user interface.

RS-232 Connections

First the wires! The required cabling at any one time is one of the following:

- From PC to DSP – early test.
- From PC to PicAdapter – load new code.
- From PicAdapter to DSP – normal use.

Fig 8.30 shows a simple implementation. The lead with the female connector for mating with the PC serial cable should be fitted for occasional use – if at all. Certainly the lead should not be routed via the RF section of the transceiver to the rear panel , to avoid any potential EMC coupling. It could be kept in a drawer and got out when needed.

The link to the PC is needed only to load new releases of DSP code, whereas the link to the DSP assembly is used continuously. Normally, the lead(s) plug into the PicAdapter board, but for loading and controlling the DSP assembly directly from the PC, a trivial connector with TX wired to RX can be used for pass-through operation instead.

PICADAPTER BOARD

This plugs into the original PIC socket on the Pic 'N' Mix DDS board. The circuit diagram is shown in Fig 8.31.

For compatibility reasons, the PIC, IC9, uses essentially the original Pic 'N' Mix code to provide all the original Pic 'N' Mix functionality. However for STAR purposes, it has four incremental tasks:

1. To download new release DSP code together with control parameter values from your PC – and retain these in IC10, a serial EEPROM – for subsequent normal use.

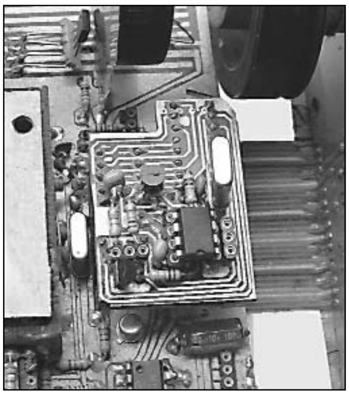

Fig 8.29: PicAdapter board in situ in Pic 'N' Mix

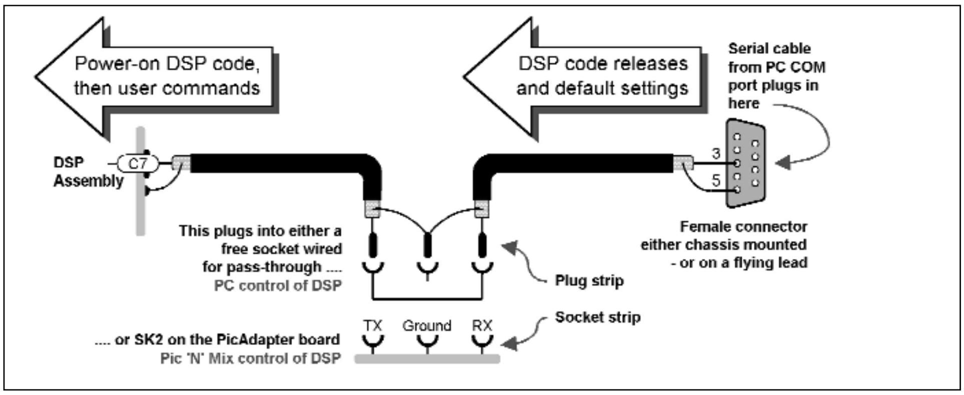

Fig 8.30: RS-232 connections. This can be made up as one composite loom – or as two separate leads

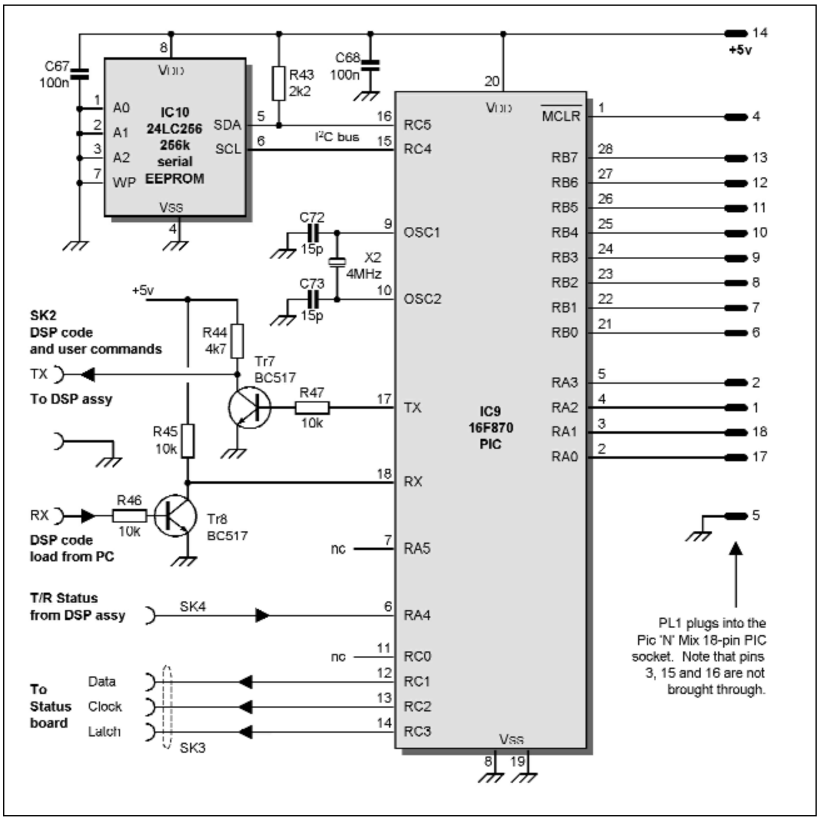

Fig 8.31: PicAdapter board circuit diagram. This board plugs into the 18-pin PIC socket on the original Pic 'N' Mix DDS board and upgrades the PIC to a more recent and versatile PIC, the 16F870

2. At power-on time, to read the DSP code and parameters from the EEPROM – and load them to the DSP board.
3. During normal operational use, to communicate any changes you make to the control parameters (eg AF Gain) – to the DSP board. The changed values are also retained in EEPROM and thereby survive power-down.
4. To drive data out to the Status board to control the state of the LEDs.

To these ends, minimal RS-232 links from your PC at 1.2kb/s and to the DSP Assembly at 9.6Kb/s are controlled by the PIC. The PIC also controls an I2C link to the serial EEPROM.

The slower downlink speed from the PC is used to give the EEPROM time to write each byte of the DSP code and control settings. However reading from the EEPROM is a much faster process, hence the higher baud value – which is the mode normally used operationally.

All the connections shown from the right-hand side of the PIC in Fig 8.31 duplicate the original Pic 'N' Mix pinout and provide much-valued code compatibility.

The T/R status line is an input to this board (and it also goes to the Timer board). Whether STAR is transmitting or receiving is determined by DSP and the result is communicated to the PicAdapter solely to allow 'split' operation. Thus it need not be fitted at first test. This line is at logic '1' on receive, '0' on transmit.

The only other pins worthy of mention are the Data, Clock and Latch pins, which simply drive the Status board LEDs.

Because this board is self-contained, it can be programmed – as an assembly – in the 18-pin socket of a PIC programmer.

For test purposes this board, once programmed, should provide full normal Pic 'N' Mix DDS operation, albeit with slightly longer key presses being required than those with which Pic'N' Mix users will be familiar.

Split operation should finally be verified with the T/R Status line connected.

STATUS BOARD

This board carries the bar-graph S-meter, the latch/driver for the status LEDs and the passive connections to – and the mechanical mounting of – the keypad. You have the option of not fitting any of these functional elements should it suit you – and the PCB is laid out so that you can 'cut bits off' should you wish. Equally and conversely, these elements were designed explicitly to be used standalone in totally differing situations if required.

The circuit diagram is shown in **Fig 8.32**. IC16 is a conventional serial-in, parallel-out driver. It is identical in function to those already fitted to Pic 'N' Mix for band-switching purposes. IC28, the PIC, exemplifies the pinout efficiency of the current generation of PICs. Of the 18 pins, only three are assigned to 'overheads' – ie ground, power and reset, the other 15 all being available for I/O. Of these, one is programmed as an analogue voltage input (the AGC voltage) – 12 as outputs driving LEDs – and two are spare. So far, that is.

IC28 is a mere voltmeter which displays S-units on receive and relative power on transmit. R73 determines its sensitivity – and I can visualise some non-STAR applications where you may need to tune its value.

STEREO AMPLIFIER

It is not easy to find a good stereo amplifier which works well on a 12V rail. If you have a scrap car radio with speakers, that would provide an instant solution.

Fig 8.33 shows the circuit diagram of an inexpensive amplifier from the bottom end of the range. Reduce C7 and C8 for more top response. The component layout is in **Fig 8.34** and the PCB is in **Fig 8.35** (in Appendix B).

Should you need something with more output, take a look at the TDA2004 or TDA2005. In any event you will need to use decent speakers to get the full benefit of PIC-ASTAR's audio quality. Several builders have found it difficult to move away from 20W per channel into 12in speakers - including me.

PICADAPTER BOARD

This board is pretty tight since it needs to fit within the envelope of the original DDS board. So, as you can see from **Fig 8.36**, it is somewhat three-dimensional. The spacing is eminently achievable provided the components are loaded in the correct sequence - which is critical.

Check meticulously for continuity and isolation as you proceed. When inserting the sockets / plugs, ensure the pin shoulders do not ground on the opposite side.

1. Fit C73 underneath (ie on the groundplane side).
2. Fit IC9 socket and ground pins 8 and 19 on the groundplane side.
3. Fit R43, underneath. 4. Cut a 9-way SIL header strip. Noting the larger diameter end fits to the PCB, cut off that end of pin 3 and insert to make PL1 pins 1-9, ie the inner strip. Solder all track-side pins - and pin 5 to ground underneath.

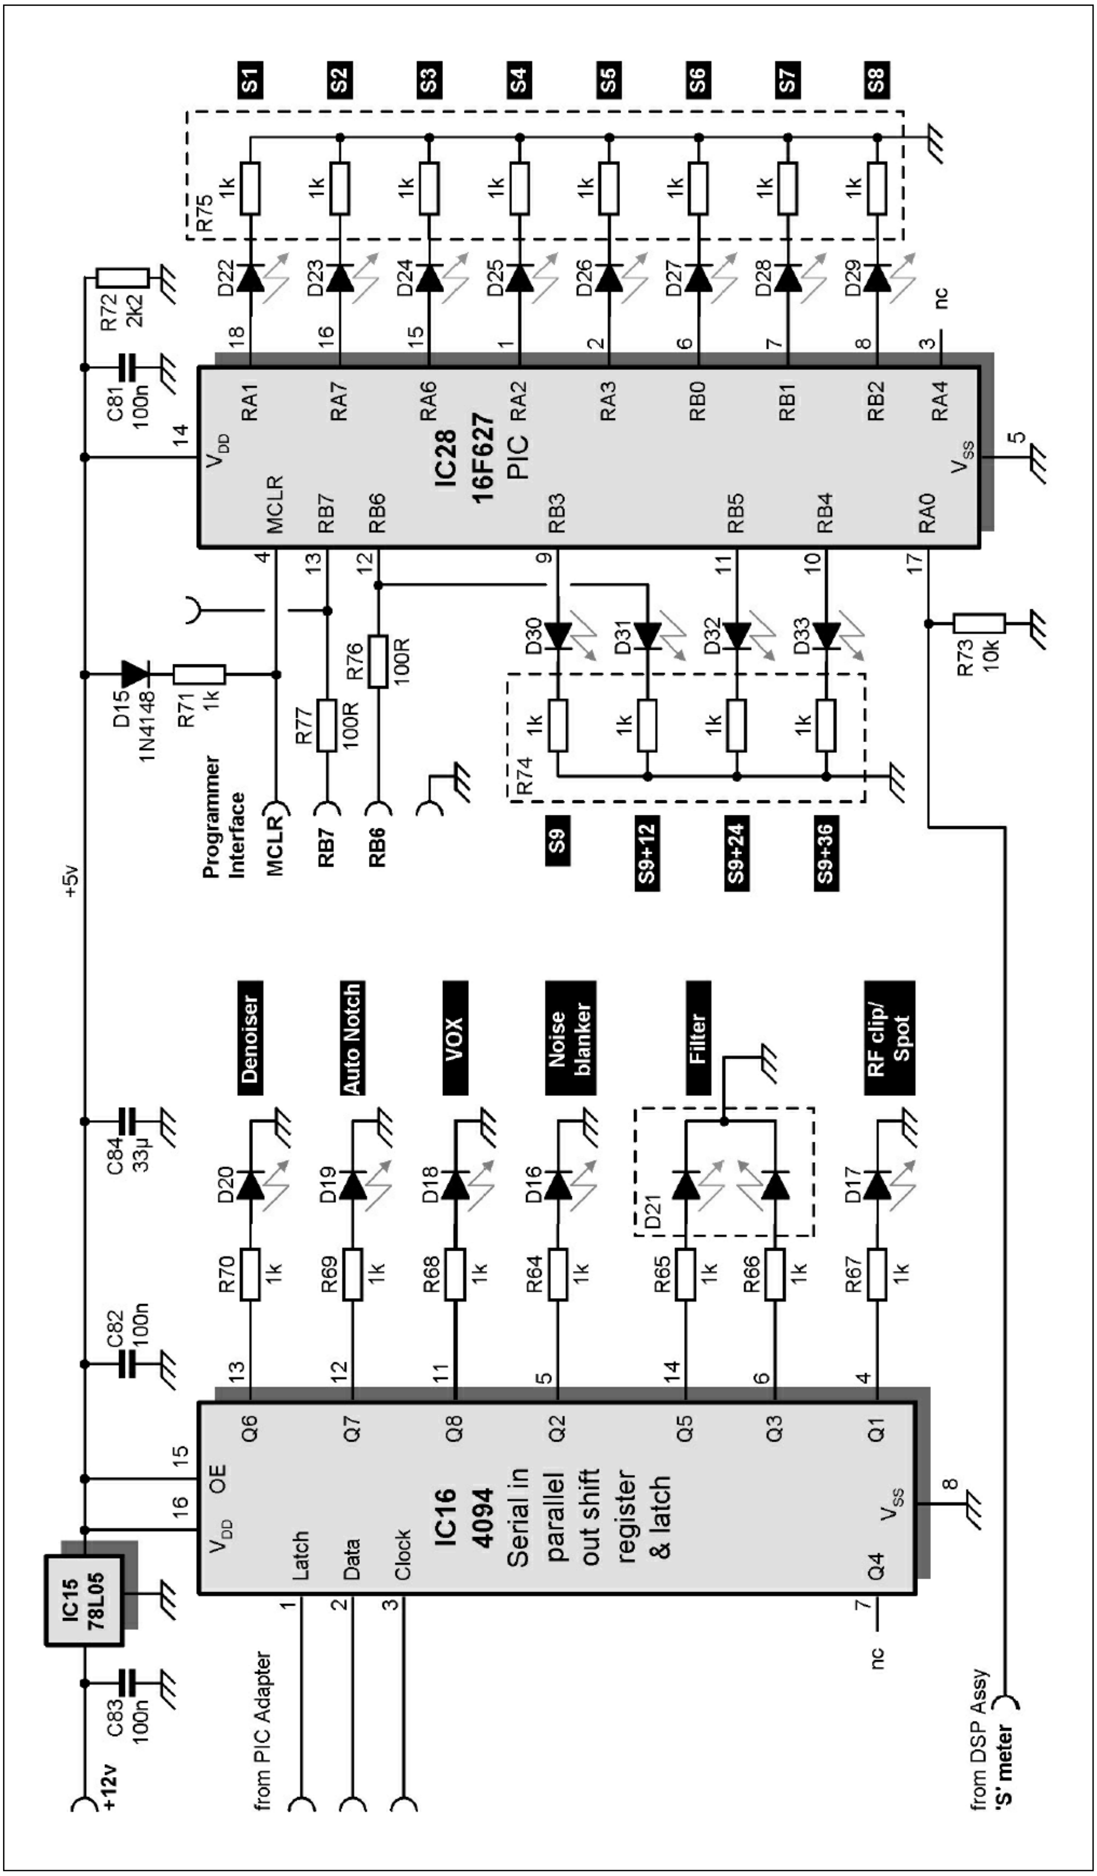

Fig 8.32: Status board circuit diagram. The status LEDs and driver are functionally unrelated to the S-meter LEDs and driver, but they are co-located around the keypad. Not shown here are some merely passive tracks which are used to make off the 7-way ribbon cable to the keypad. The connector shown on RB7 is for future development only

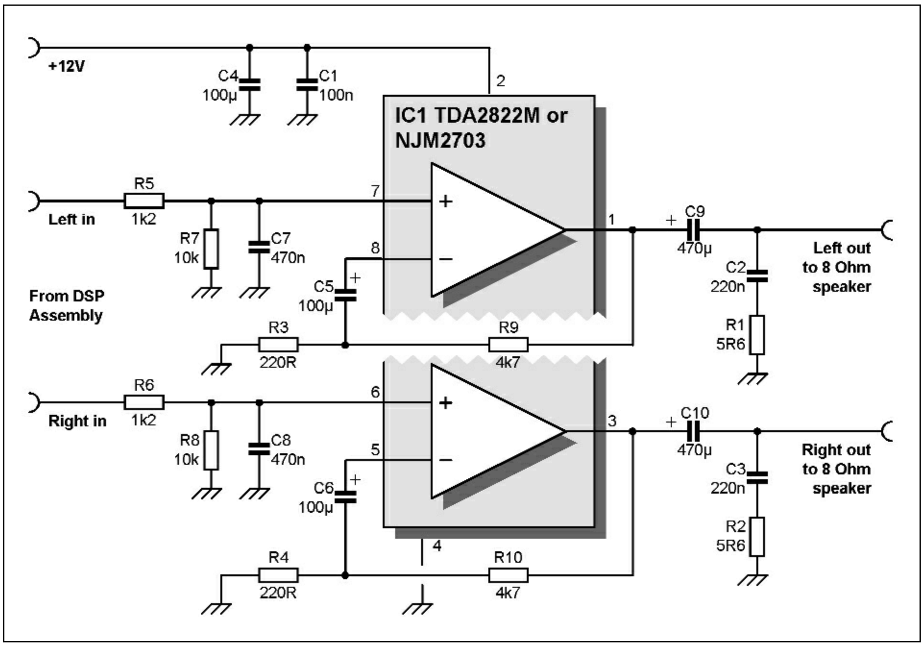

Fig 8.33: A suitable 1W + 1W stereo audio amplifier. Stereo balance and gain are controlled in DSP

IC1	TDA2822M or NJM2703
C1	10n disc ceramic
C2, C3	220n disc ceramic
C4, C5, C6	100m 16V electrolytic
C7, C8	470n 16V electrolytic
C9, C10	470m 16V electrolytic
R1, R2	5R6
R3, R4	220R
R5, R6	1k2
R7, R8	10k
R9, R10	4k7

Table 8.3: Components list for the stereo amplifier

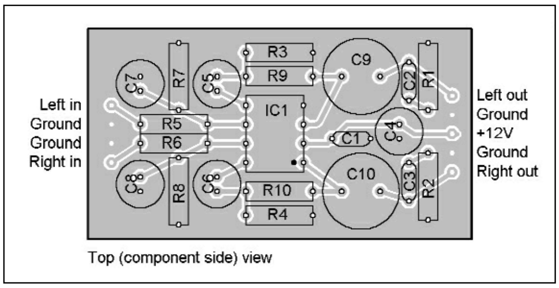

Fig 8.34: Component layout for the double-sided stereo amplifier PCB. This is a 'conventional' board with components mounted on the top, the track underneath. The component side is completely unetched. No connectors are specified, but the relevant pads have a 0.1in pitch. The PCB artwork (Fig 8.35) is in Appendix B

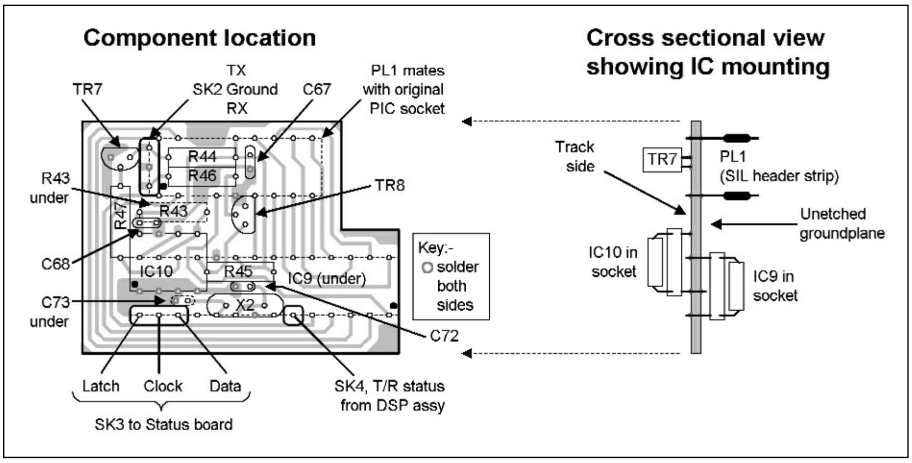

Fig 8.36: Component layout for the double-sided PicAdapter board (see Fig 8.37 in Appendix B for the PCB artwork). The cut-out is to give access to the existing programming socket. IC10 and IC9 should definitely be mounted in sockets. To achieve the clearance height required you may need to fit first an extra 18-pin socket into the existing Pic 'N' Mix socket -as a spacer. SIL plugs / sockets are used for the remaining leads - with the sockets soldered directly to the track

5. Repeat for the outer strip of PL1 - but cutting off the larger diameter end of pins 15 and 16.
   Use a spare 18-pin socket to ensure alignment.
6. Fit TR8, C72, R44.
7. Fit IC10 socket, grounding pins 1, 2, 3, 4 and 7 both sides.
8. Fit C68.
9. Fit SK2, grounding the centre pin underneath.
10. Fit TR7, grounding the emitter underneath.
11. Fit R46, R45, C67, R47 and last, X2.
12. If there is any chance of the board fouling the Pic 'N' Mix Display board, chamfer the copper both sides.

The original 4MHz crystal on the DDS board may be recovered and reused elsewhere, though you may want to postpone this until the PicAdapter is working.

BUILDING THE STATUS BOARD

There are no special constructional issues here.

Cut all the IC socket pins back to their shoulders - except the grounded ones. A small trick for soldering the socket on the component side; fit another socket or a scrap chip into the socket first. This prevents the pins from wandering as they get hot.

Please note that should you wish to program this PIC in situ, you may not be able do so with D31 fitted, since it loads the programmer.

For this reason, avoid giving reports to stations who are "12dB over S9" until the end of full integration and test -when this LED is finally fitted.

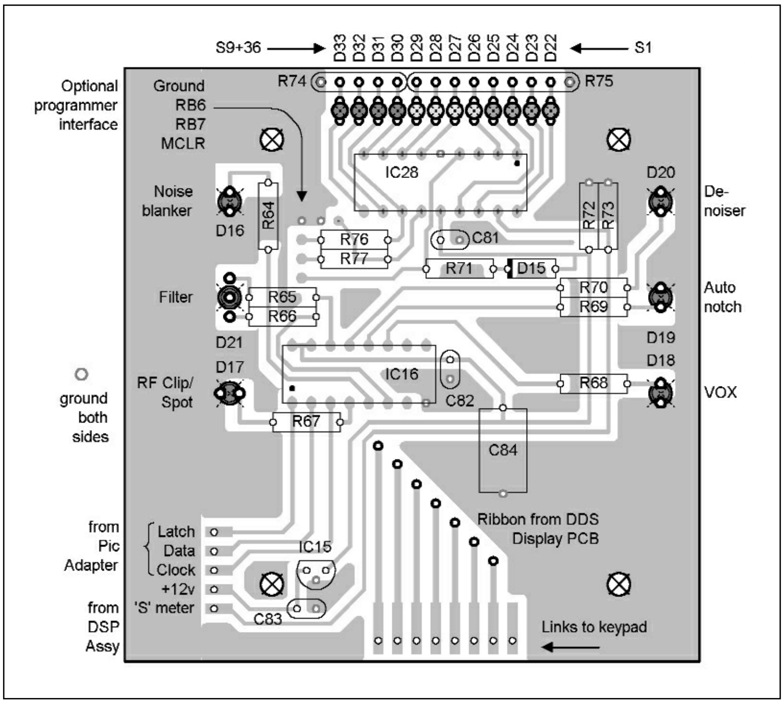

Fig 8.38: Component layout for the double-sided Status board PCB with the ground-plane side unetched (see Fig 8.39 in Appendix B for the PCB artwork). The picture is of the track side, viewed from the rear of the front panel. The width of this board (3in) accommodates that of the RF front-end which sits behind the Status board. All components except the LEDs are surface-mounted on the track side. This allows access to both the components and to the LED pads. The latter is needed to adjust the LED lead lengths for flush-fit to the front panel. The four mounting holes are for the keypad which is mounted using short spacers. The 7-way ribbon cable from the Pic 'N' Mix Display board is routed between this board and the keypad and made off to pads / tracks provided on this board and thence via short wire links to the keypad itself. IC16 and IC28 are mounted in sockets with all pins - except their respective ground pin - cut back for surface mounting. No connectors are specified. The relevant pads have a 0.1in pitch

	* - toggle between VFOs
	0 - VFO A = VFO B
	# - Split operation
	11 - select USB
	44 - select CW . again for reverse CW
	77 - select LSB
	22 - toggle rate tuning mode on / off
	26 - toggle software flywheel on / off
9	31 - calibrate / save USB offset
9	37 - calibrate / save LSB offset
9	33 - calibrate / save reference clock
	41 - go up to nearest kHz point
	48 - toggle ¡®spare¡¯ latch output
	47 - go down to nearest kHz point
	55 - monitor guard channel
	56 - scan between memory frequencies
	58 - scan frequency range (wobbulator)
	5* - scan both VFOs
9	6x - where x is 0 - 9.Go/save to memory - may be used for 5MHz frequencies
	70 - display tuning rate as bargraph
	72 - spare user on / off switch
	73 - go to SLEEP
	74 - toggle display auto-dim
	78 - display LSD as 10Hz / 100Hz
	79 - toggle x6 ref clock for AD9851
	81 - high side injection (default)
	87 - low side injection
	83 - signal generator mode (no offset)
	88- direct keypad frequency entry
9	10, 12, 15, 17, 20, 30, 40, 60, 80,
	16 - go/save to respective band
	990- save as power-on frequency
	999 - reboot DDS software

Table 8.4: DDS key sequences (see text)

PIC 'N' MIX FUNCTIONALITY

DDS Key Sequences

The keypad sequences are summarised in **Table 8.4** These include the increments for STAR operation.

Sequences shown with a leading 9 save the current frequency in the respective location if the 2-key sequence is preceded by the 9 (Save) key.

Pic 'N' Mix Code Changes

The following apply to code shipped explicitly for use with PIC-A-STAR:

- The opportunity has been taken to de-bounce the keypad more vigorously. The simple consequence of this is that key presses now need to be somewhat more deliberate.
- CW offset calibration (previously 34) is no longer required, since the CW offset is now managed in DSP. The procedure for calibrating the reference clock and SSB IF offsets is unchanged. However, the reference clock frequency calibration may now be loaded directly from your PC. For IF offset calibration you will find it useful to switch the DSP filters off . so that you can hear right down to zero-beat. A CRO connected to the receiver audio output is also invaluable for seeing the exact zero-beat point. When both SSB offsets are correct, there should be a difference of precisely 2.7kHz between them.
- CW is now received either upper or lower sideband . so you have the choice of tuning direction. Either in the same direction as SSB signals on the same band - or always tuning CW in the same direction whatever the band. The frequency readout always shows your transmitted frequency (not least for legal reasons) and your receive frequency is displaced from this by the CW offset . which is your chosen and preferred beat note.
- If you are in CW mode and you key 44 again, this will toggle reverse CW. This switches both receive sideband and injection frequency to give the same beat-note, but 'from the other side'. It can be useful in clearing QRM and for checking you are properly netted since if not, the beat note will change.

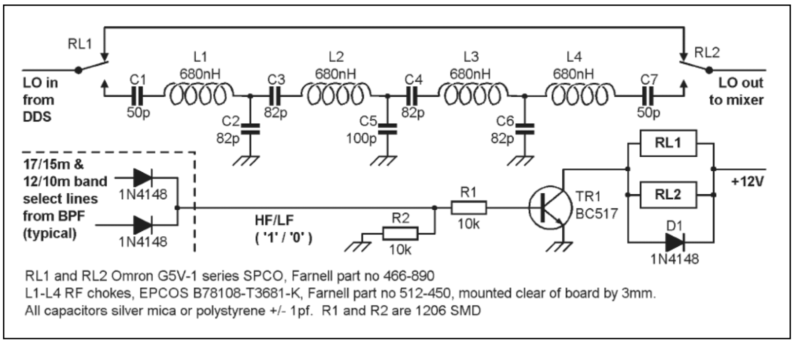

Fig 8.40: 26-40MHz injection filter to strip both unwanted high-order spurs and those at the IF from the DDS output on the higher HF bands only. It is not needed (ie switched out) on lower bands

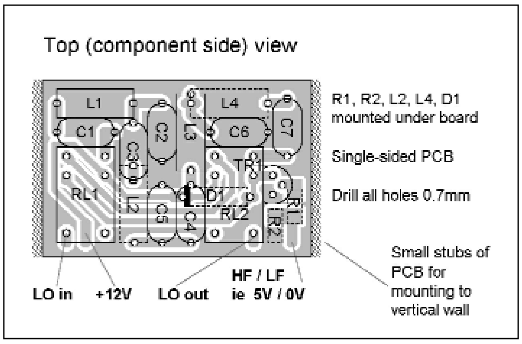

Fig 8.41: Injection filter component layout. The PCB artwork can be found in Appendix B at the back of this book

- Pic 'N' Mix now has 5MHz (60m) capability. The latch output previously marked "15MHz WWV" is now active if any 5MHz frequency is selected and . exceptionally . if you go to 5MHz, the default sideband is now USB. The first five memory locations (60.64) are loaded with the UK 5MHz channels. These are the correct frequencies for upper sideband operation. If you don't want this feature, you can re-program these allocations with any other frequencies (and re-enter the 5MHz ones yourself any time).
- You may now fit an AD9851 DDS chip which has the ability to multiply the reference clock frequency by six. 79 toggles this feature. See later for more detail.
- There is a latched output bit labelled 'spare'. This may be toggled by keying 48. It was designed so that you can switch any device - a pre-amp, attenuator, transverter etc - from the keypad. This switch has been used to configure the STAR mixer and post-mixer amplifier to be described later.
- The output bit labelled 'broadband' 72 is now a spare uncommitted toggle switch.
- QSK Split operation has been improved. Previously it was limited to about 20WPM.
- The frequency display is dimmed after about three minutes of user inactivity - to reduce heat dissipation and to prolong LED life. 74 toggles this feature on/off.
- All frequencies from 0.29MHz now activate the nearest band select line.
- RIT and XIT operation may be selected instead of Split. The # and . keys and the Rx/Tx A/B LEDs then change meaning. This is still under development. Detail follows.
- Part-way through any DDS key sequence, the # key now aborts it.

PIC 'N' MIX HARDWARE

Since first designing Pic 'N' Mix five years ago, some detailed improvements have evolved. What will never change is the requirement for meticulous (albeit textbook) screening, decoupling and filtering practice . if the DDS spur level is to be contained. This cannot be overstated and is the starting point for what follows.

The following modifications produce incremental reductions in DDS spurs and are easy enough to do to make them all worthwhile. If you are building STAR, I countenance you not to implement any of them until you have everything else working . and until you have truly attended to the meticulous bits just mentioned.

- Fit a separate 7805 regulator for the DDS chip carrier assembly. There is a simple track cut under the DDS board which removes the +5V rail from the DDS Assembly. Mount a separate 7805 on the rear vertical panel with 100n on both the 12V in and 5V out leads - and run a flying lead from the latter through a ferrite bead - and solder it to the 5V top foil on the DDS chip carrier together with a 100¥ìF electrolytic to ground.
- Fit 1n, 10n, 100n 1206 SMD capacitors - in a stack - on at least two corners of the DDS assembly from the +5V foil to ground. In other words, whatever it takes to ensure that the +5V rail and ground-plane are at the same AC potential - AF to VHF.
- Add a filter in the LO feed to the mixer. Figs 8.40 and 8.41 show a suitable arrangement for high-side injection with any IF between 8MHz and 11MHz. This filter offers at least 30dB attenuation at the IF and a similar figure at 50MHz - and rising. If all else is right, this produces dramatic results. It does not need any exceptional screening if mounted within the already-screened volume of Pic 'N' Mix. It is shown switched in by band-select lines, diode-ORed. You should wire in diodes for any bands for which your LO falls in the range 26 . 40MHz. However, for evaluation purposes, it would be prudent to control it with a simple toggle switch to +5V / 0V in the first instance.
- You may simply substitute an AD9851 on the DDS Assembly in lieu of an AD9850 under all circumstances, ie it is completely hardware and software compatible. The AD9851 allows higher reference clock frequencies . which may be useful if you have a relatively high IF. This gets you away from the 1/3 reference clock zone. In addition, it offers a 6x multiplier option for the reference clock. So, for example, you could clock it at a mere 30MHz yet have the effective benefit of a 180MHz clock. There are small spur and significant phase noise performance disadvantages for which see the AD9851 data sheet - but convenience issues may predominate for you. A facility for specifying a 6x clock is now built into STAR software. You may instead want to try clocking your existing AD9850 much faster.

Fig 8.42 shows a reference oscillator which will operate on either the 5th or 7th overtone of a crystal in the 22 - 25MHz range . simply by tuning VC1. Typically, the 5th overtone is used with an AD9850 and the 7th overtone with an AD9851. In this latter case, the x6 feature would not be invoked. Also included in Fig 8.42 is an arrangement for stabilising the crystal temperature, the detailed design of which is due to G3NHR. It works so well, I have since stuck (literally) a similar arrangement on my IF board translation oscillator. The thermistor, TH1, and the TIP122 tab are secured to opposite faces of the crystal can. Heat-

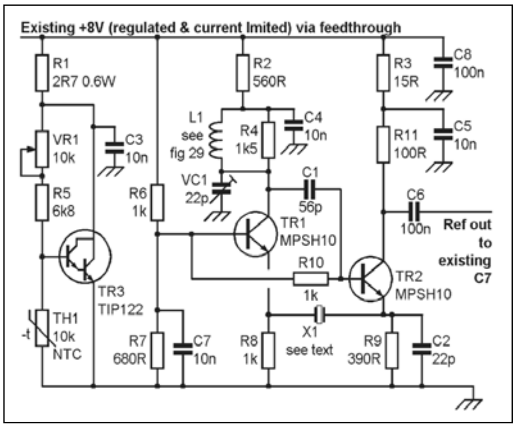

Fig 8.42: Butler oscillator for either 5th or 7th overtone operation . with crystal temperature control

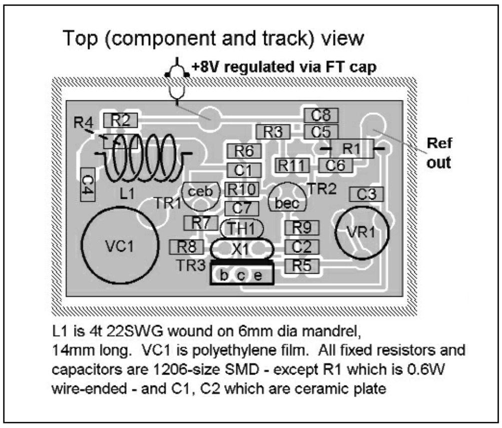

Fig 8.43: Reference oscillator component layout. The PCB artwork can be found in Appendix B

shrink sleeving is highly recommended . or failing that, super-glue. For smaller crystals, cut off the excess TIP122 tab to reduce height. Set VR1 to maximum resistance and then adjust for 100mV drop across R1. Repeat every minute for five minutes and the result should be a crystal thermally stable at 35 degrees C.

In use, the frequency will change rapidly for the first five minutes after switch-on . but stabilise thereafter. This is not an on-off oven. This is proportional control.

The small PCB, shown in **Fig 8.43**, is designed as a drop-in replacement . after simply removing the components from the original Pic 'N' Mix oscillator.

- Fit transformer-coupled output from the DDS chip. This gives 6dB more LO output and further spur reduction. The core for this transformer is the EPCOS B62152A4X1 available from ElectroValue [9]. (You will need four more of them for the mixer later). The primary is three bifilar turns 32SWG and the secondary is 12 turns wound over the top. The core is mounted in lieu of the 100 and 200 ohm resistors on the DDS carrier. Connect the primary instead to pins 20 and 21 of the DDS chip, grounding the centre-tap. Ground one side of the secondary and take the other via a 100n blocking capacitor to pin 19 on the 28-pin carrier. If you are not confident your mixer will stand the extra 6dB, fit a pad on its LO port. (The STAR mixer . yet to be described . is fine.)

To reduce DDS spurs to a highly acceptable level (ie virtually none) you may or may not need any or all of these changes.

As with all flexible designs, there are intelligent user choices to be made. Such is amateur radio!

FACILITIES

PIC-A STAR has more useful facilities than most homebrew designs – and getting a feel for their value may well determine if this is the project for you.

Those who don't have these facilities refer to them as 'bells and whistles'. Those of us who do just grin – and keep ringing them bells and blowing them whistles! Always assuming they are underpinned by a rock-solid base receiver performance, that is.

SSB/CW Mode Management

Switching between SSB and CW – and transmit and receive for that matter – are non-trivial design problems if the result is to be user-friendly. So some background discussion is helpful in understanding what follows. See also ref [10]. You may also care to compare critically the behaviour of commercial transceivers. If I can find anything friendlier than STAR, I will simply change the design until it isn't. I have, as I write, already invested in the architectural infrastructure to make all this possible.

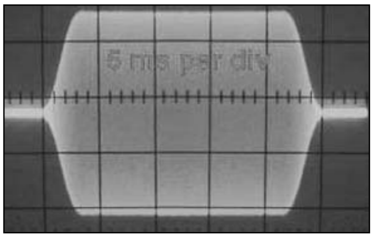

Fig 8.44: The STAR CW waveform from G4HMC, photographed off-air on 80m from the author's STAR

CW operation

In days of old, especially with 'separates', it was easy, reliable – but a bit tortuous. You would zero-beat an incoming CQ on your receiver, then zero-beat your transmitter – and finally move your receiver off to get a comfortable beat note.

With modern filters there is a snag. You can't hear anywhere near down to zero beat, so this process produces totally unacceptable errors.

However, it does establish two critical principles, namely: both stations must transmit on the same frequency – and both must operate 'split' if they are to hear a beat note. CW is inherently a 'split' mode.

With a multi-mode transceiver, you must either explicitly operate 'split' for CW – or the design must take care of it transparently.

PIC-A-STAR is in the latter category. If you are in SSB mode and hear a CW station you want to work, when you switch to CW mode the received pitch will not alter and you will not have to retune. And vice versa if starting out from CW mode.

While on the topic of CW, take a look at the photograph of STAR's transmitted waveform. You won't find better.

SSB operation

When you change sideband, neither your indicated nor actual frequency should alter. This is common currency nowadays. There are some, but not many, occasions when this matters – given that we don't usually operate on the 'wrong' sideband. If you operate via OSCAR or into a transverter or on 60m, it is critical.

This point also reinforces a principle which should be obvious – namely that if a given feature is critical to a minority interest, provided it is not imposed on the majority – then there is no good excuse for not providing it.

Pic 'N' Mix also gives you the option to switch between high- and low-side injection. Since this implies a sideband inversion, the software puts in an equal and opposite change of sideband – and you end up on the same net frequency.

But be aware that many band-pass filters are optimised for injection from a preferred side. The STAR front-end is optimised for use with high-side injection; the optional filter in the LO line also assumes high-side injection on the higher bands.

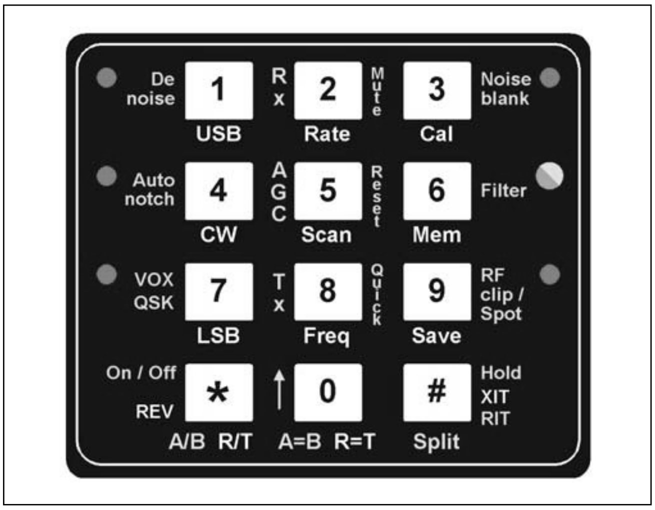

Fig 8.45: Keypad allocations and corresponding status LEDs. This is to scale and may be used as a keypad overlay

PTT and key behaviour

PIC-A-STAR uses the simple conventions that, on CW, the microphone audio is ignored – and on SSB the key is ignored.

If on CW and you merely key, you will produce only sidetone. This is for CW practice since, to transmit in earnest, you either need to switch QSK on – or hold down the PTT line for non-QSK operation.

If on SSB, you must either switch on VOX – or hold down the PTT line – before anything will happen.

AT POWER-ON TIME

The display shows your chosen startup frequency, but flashing. You then have four structurally-different options:

- Upload DSP code from Pic 'N' Mix to the DSP Assembly. This is normal every-day operational use. The Status board LEDs flash strangely so you know something is happening.
- Run the DDS without loading DSP code. This is mainly a diagnostic mode.
- Enter a DDS reference clock frequency from your PC. A useful utility.
- Download a new (or your first) DSP software release – via the Internet to your PC – and thence to Pic 'N' Mix. This latter process takes several minutes. During this period the incoming bytes are counted on the display – albeit faster than the eye can follow – much like money on a petrol pump. Unlike petrol it gives you a warm feeling you are getting good value – and the fact that it is counting at all signifies that it is working. Once the new code and the default control values are all in, you can then proceed to upload them to the DSP Assembly thereafter.

Keypad / Display Modes

There are now two main modes, namely 'DDS mode' and 'DSP mode'. The former was outlined last month. The latter lets you tune all of the STAR DSP controls.

Fig 8.45 shows a suitable keypad overlay with the DSP legends in yellow; the DDS ones in white or blue. But before we get to that . . .

Split or XIT/RIT

This is a new sub-mode choice for DDS use. One of these is always enabled – and your choice is retained at power-down. Both give you the potential to transmit and receive on different frequencies. See also [11] for a general discussion.

Split mode is unchanged from Pic 'N' Mix – but with enhancements. It operates on the two independent VFOs, 'A' and 'B'.

Conversely, XIT/RIT operates on either one of the VFOs – and that choice determines the initial Tx and Rx frequencies. But, thereafter, the Tx and/or the Rx frequencies can be independently changed – and retained throughout an XIT/RIT session. Meanwhile, the 'other' VFO remains uncontaminated – and available.

What are the differences between Split and XIT/RIT? Split can be cross-band and/or cross-mode and 'rests' on your Rx frequency when off. Conversely, the XIT/RIT tuning range is the current band, current mode. It 'rests' on your Tx frequency when off – thus providing an RIT on/off capability – which Split doesn't give you.

Split and XIT/RIT use

Besides pure transceive when switched off, both modes have the following options:

- Tune only the receive frequency while your transmit frequency remains fixed (ie RIT). If you call CQ and a station answers off-frequency, this option (in either mode) is the answer. However, in a net with one station off-frequency, XIT/RIT mode is better since you can turn XIT/RIT off when the offending station is not transmitting.
- Tune only the transmit frequency while continuing to monitor your receive channel (ie XIT). To call a DX station who is operating split, use in either mode to tune your Tx quickly to the specified DX listening frequency – while not missing a word on your receive channel.
- Tune the transmit frequency while monitoring what is about to be your transmit channel (ie REV XIT). Use this to check for a quiet spot before calling.

You can do any of the above – independently – in any order using merely one key-press to define your choice. On receive,

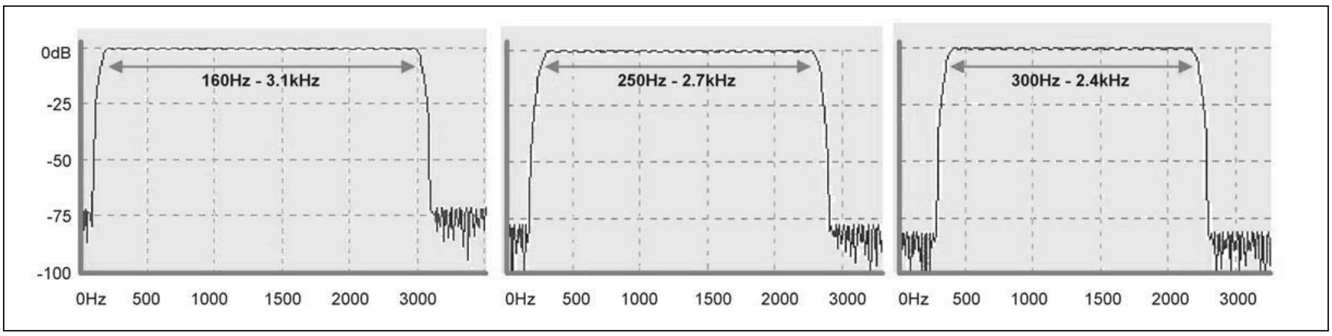

Fig 8.46: Plot of wide, medium and narrow Rx SSB filters. Other filters in the DSP receiver give a further theoretical 30dB of ultimate stop band. The pass-band ripple in all cases is less than 0.2dB. Because STAR is not a mere audio add-on, these widths are actually achievable – and usable

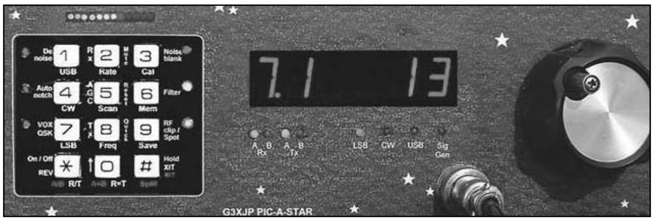

Fig 8.47: STAR display in DSP mode

1*	Denoiser ON/OFF	
1.1	Denoise beta	
1.2	Denoise decay	
2*	Mute Rx and suspend VOX	both
2.1	AF gain	both
2.2	RF gain (per band)	both
2.3	Stereo effect (ie amount)	both
2.4	Stereo balance	both
3*	Noise blank ON/OFF	
3.1	Noise blank threshold	
4*	Auto notch ON/OFF	
4.1	Auto notch beta	
4.2	CW tones, 1 or 2	CW only
4.3	CW offset frequency	CW only
4.4	Sidetone frequency	CW only
5*	Reset downloaded values	both
5.1	AGC hang time	
6*	Filter ON/OFF	
6.1	Filter width (1 – 6)	
6.2	Filter depth	
7*	VOX/QSK ON/OFF	
7.1	VOX/QSK hang time	
7.2	Rx – Tx pre-delay	
7.3	Tx – Rx blank time	
7.4	VOX gain	SSB only
7.5	Anti-VOX gain	SSB only
8*	Quick Switch ON/OFF	both
8.1	Tx drive level (per band)	
8.2	Monitor level	
8.3	Mic gain	SSB only
8.4	Tx Top boost	SSB only
8.5	Tx Bass boost	SSB only
9*	RF clip/Spot ON/OFF	
9.1	RF clip ON/OFF	SSB only
9.1	Spot level	CW only
#	Hold DSP mode ON/OFF	both

Table 8.5: DSP menu structure

the frequency displayed will be that which you are changing – and on transmit, always your Tx frequency.

Further utility options let you swap the two VFO frequencies – or initialise them as the same. Likewise for the XIT/RIT frequencies.

Initialisation is done for you in XIT/RIT mode should you change VFO – or band – or frequency by more than 2.5kHz while on pure transceive – on the grounds that any difference must then be irrelevant.

DDS/DSP mode switching

As supplied, the DDS mode is permanently engaged and you would be unaware that there is any other. This is deliberate in order that you can check out the DDS functionality after first commissioning – without distraction.

Once you are happy here, the DSP mode becomes available following the first time you download DSP code from your PC to Pic 'N' Mix. The remainder of this discussion assumes that this has happened.

In normal use, STAR 'rests' in DDS mode and displays frequency. The key to switching to DSP mode is the duration of the first key press. A quick press on a key activates DSP mode; whereas a longer press invokes the 'business-as-usual' DDS function.

The resultant displays are quite different, so it will only take you a few minutes to get the 'feel' ingrained. The DSP functionality is given this priority because most DSP functions are needed quickly in real operating conditions. For example, turning the auto-notch on when that tuner starts up is a more immediate issue than, say, changing bands. Certainly a 'quick press' need not be tentative, merely not overtly sustained. Once the mode has been determined by the duration of the first key press, the duration of subsequent key presses is unimportant.

The very deliberate exception is the bottom row of keys – which act to give vital DDS functions – and which therefore have no DSP functionality as the first key-press.

Once you are in DSP mode if you neither press a key nor alter a value for about three seconds, PIC-A-STAR will revert to DDS mode. You can prevent this by pressing the # key – which toggles holding the display in DSP mode. This is invaluable when setting up the DSP control settings.

USING THE DSP MENU

The DSP controls are grouped to form a menu, as shown in Table 8.5. This menu detail will change over time, but not the intrinsic structure. Project without end, right?

Menu Groupings

The menu is ordered with the more commonly-used controls near the 'top' of each menu group – and rarely used ones near the 'bottom'. In practice, some of these latter items can be regarded as presets.

There are no sub-menus, so the system is inherently limited to 99 controls in nine groups times two modes .

The menu is intrinsically SSB/CW modal in the sense that if you are in SSB mode then you simply can't get at controls which are unique to CW – and vice versa. These controls are annotated 'SSB only' or 'CW only'. Some controls share one common value for both modes and are annotated 'both'. Two controls, namely 2.2 and 8.1, have different values per band.

Conversely, all other menu items that are not peculiar to mode have different settings stored for SSB and CW. For example, different AGC time constants, Denoiser settings and so on can be set up, varied – and retained independently for SSB and CW. This applies also to the on/off switch settings.

All this is designed to foster a 'set-and-forget' philosophy.

How to PIC from the Menu

Immediately after you (quickly) press a 1-9 key, the display will switch to DSP mode. For example, a dab on the 7 key will show '7.1 13' where 7.1 denotes the first control in that menu group, namely VOX/QSK hang time – and 13 is its present value. See also Fig 8.47.

If you want to move on to the second control in the same group, press the 7 key again – and so on. If you want to move back up through the group, press 0. If you want to move to a completely different menu group, simply press the corresponding key.

Changing Values

After you have selected a control, if you want to change its value, turn the knob – clockwise for more, anticlockwise for less.

There are maximum and minimum values for each control – and the rate of change is proportional to the range.

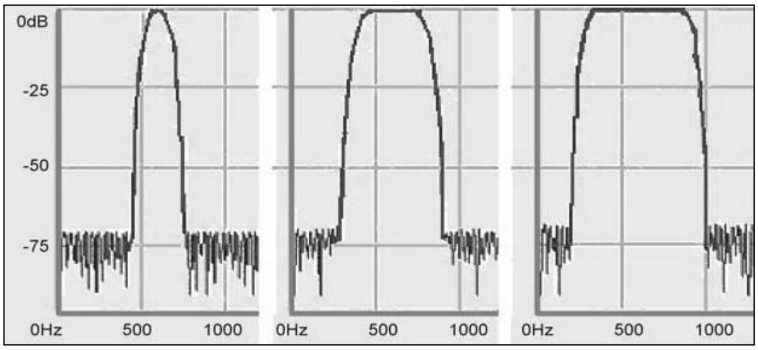

Fig 8.48: Plot of the Rx CW filters. This is the bank centred on 600Hz – and those on different centre frequencies are otherwise similar. Their widths are approximately 200Hz, 500Hz and 750Hz but this absolutely depends on where you measure them. This filter shape gives less ringing than a 'brick wall' type. Other filters in the receive path give a further theoretical 30dB of ultimate stop band

ON/OFF Switches

To switch a DSP feature on/off, press the corresponding key (1, 4, 7, 3, 6, 9) followed by *. The adjacent LED will change to provide visible status thereafter.

For example 7*. will switch VOX on/off if in SSB – and QSK on/off if in CW mode.

In fact, irrespective of which menu item in a group you are addressing, the . key will toggle the associated switch and you will revert immediately to DDS mode.

Mute

2* near-mutes the receiver and suspends VOX operation. This is the 'panic' button for unexpected interruptions eg when the phone rings. Any subsequent key press or knob turn restores your pre-panic state.

Quick Switch Option

8* toggles the quick switch facility. When engaged, any one of the 1, 4, 7, 3, 9 keys – when pressed – simply toggles its respective DSP switch.

Because you are thereby not presented with the values for those menu groups, you would not want to use this option until those groups are set up. Conversely, once the values are tuned and you have gained familiarity, this could be the mode of choice.

Resetting Control Values

5* resets the control values (both SSB and CW) to those that you last downloaded from the PC – with the exception of RF Gain and Tx Drive – the latest per-band values of which are retained. All the DSP values are remembered across a power-down, but not switch settings – which initialise to off but with the DSP filter on – and in SSB mode.

DSP FEATURES DESCRIPTION

A few words are in order for some of the more esoteric features you may not have met before. In roughly menu order:

Denoiser

This is rather more a comfort feature than a performance one. It acts to reduce background white noise – and when working well, is not unlike squelch on FM. It is especially effective on CW and very useful if just monitoring a quiet (albeit noisy) channel.

The 'right' combination of settings is somewhat subjective and can occasionally vary from one signal to another – and certainly by mode. It is best with the RF gain turned up and with longer AGC hang times. Experiment!

Both the Denoiser and Autonotch (see later) are essentially as implemented in DSP-10 by Bob Larkin. See also [12-15] for the pioneering work and the theory.

RF Gain

This comes right at the front of the DSP receiver chain and is used to set the SNR for different conditions. It should normally be turned well up so that AGC action produces constant audio output – and the best possible SNR. This also contributes to clean VOX operation.

Stereo Effect

This gives body and presence to signals and warrants a decent stereo audio amplifier and speakers. Some folks report an increase in readability on weak signals. Personally, I just love it! For me, it completely transforms the listening experience.

Stereo balance

Values above 100 decrease the right channel output; those below 100 decrease the left channel output. Another scratchy pot saved.

Noise Blanker

As opposed to the Denoiser which acts on white noise, this acts on impulse interference – eg ignition noise, thermostats, electric fences and the like.

Auto Notch

This removes an interfering heterodyne – and, in many circumstances, several. It works best on pure tones and especially lower-pitched ones. It has exactly one use in CW mode, namely for monitoring key clicks ie what's left after removing the tone. One very popular commercial transceiver shows up here every time.

Auto notch is applied after the filter bank and is outside the DSP AGC loop – to avoid strong-signal overload of the DSP.

CW Offset

This is your preferred beat note and may be pre-set (when you download from your PC) to 5, 6, 7, 8 or 900Hz. The centre frequency of the CW filters is changed to match.

If you are interfacing with other STAR display in DSP mode than Pic 'N' Mix, you will need to adjust your Tx/Rx mixer injection frequency by mode. So, for the record, the following are the exact DSP IFs (in kHz) used by PIC-A-STAR:

LSB	Tx = Rx = 16.35
LSB CW	Tx = 16.35, Rx = Tx + CW offset
USB	Tx = Rx = 13.65
USB CW	Tx = 13.65, Rx = Tx – CW offset

There is also a Reverse CW option and if set:

LSB CW	Tx = 16.35, Rx = Tx – CW offset
USB CW	Tx = 13.65, Rx = Tx + CW offset

Sidetone Frequency

Not to be confused with CW Offset, this is the tone you hear when sending CW. The QSK experts tell me it can be useful to have this at a different pitch from an inbound signal – so you can intuitively tell the difference between you and the station being worked when using fast break-in.

To enhance this effect, the sidetone comes from one speaker only, the inbound signal from both. The pitch may be excessively varied between 10Hz and 2.54kHz in 10Hz increments – this extended range being useful should you need it to double as an instant audio signal generator.

CW Tones 1 or 2

This allows two-tone testing. The two tones are 700Hz and 600Hz. If you have not used a two-tone test for linearity checking before, be aware that the duty cycle is very high – so use only short or pulsed bursts.

AGC Hang Time

This can be set anywhere between very short and very long. I understand that most people usually only change this per QSO for CW work – and so you can vary the CW setting without altering the SSB setting. A value of 0 turns DSP AGC off.

While actually changing frequency, the hang-time is set to short to avoid annoying hangs after tuning across large signals.

Filter Width

This allows you to set the receiver filter width by turning the knob (or on/off using 6.). There are six filters currently provided, three each for SSB and CW (though any one can be used in either mode) – see Figs 8.46 and 8.48. The status LED is tri-colour and corresponds to wide, medium, narrow – or off. Turning the filter off is a useful way of checking if a signal has stopped transmitting or has just slipped out of the pass-band since, when you switch the filter back on again, it reverts to the previous filter width.

Filter Depth

This concept was inspired by yet another conversation with Bill Carver, W7AAZ. In traditional analogue terms, it allows a controlled leak past the filter. For CW operation it is in many ways more useful than controlling the filter width. In use, having tuned a wanted signal to the centre of the pass-band, you simply increase the filter depth (ie the stopband rejection) until the QRM is reduced to any level with which you feel comfortable. Putting it another way, you come up to periscope depth to find a target – centre it up – and then go down again so the nearby destroyers can't get you.

VOX and QSK

The hang time is the duration spent on transmit after you have stopped speaking (or keying) before PIC-ASTAR reverts to receive. If you have a relay-free T/R system, then there is no need to set other than a very low value here. With relays, any greater setting will minimise the number of rattling occasions. It is adjusted to allow the trailing edge of your transmission to pass before switching to receive.

The Rx-Tx pre-delay is the time your signal is delayed by DSP to allow for relay settling when switching to transmit. It is adjusted so that the leading edge of a transmission is not truncated – just.

The Tx-Rx blank time is the duration of DSP receiver blanking immediately after reverting to receive. It should be set to the minimal value possible, consistent with no objectionable click coming from the receiver after the transition. With a full STAR configuration, this is simply zero.

The above three parameters are separately set for SSB and CW. For SSB only, VOX and anti-VOX gains may also be set. See later for further discussion.

Tx Drive

This is the power-setting control. When on transmit, the S-meter reading corresponds to Tx drive level – and is therefore modal.

Monitor Level

On SSB this control sets the level at which you monitor your own voice – after all VOX processing and filtering. If, for example, you turn the Rx-Tx pre-delay up high, you will – somewhat disconcertingly – hear your very delayed voice. And you will hear leading or trailing edge truncation if the timing is not set up properly. For operational use, however, the level should be kept low to avoid any feedback; or worse, confusion of the VOX software.

On CW, this control sets the sidetone level; you do not hear a delayed signal – since this would play havoc with your sending.

Microphone Gain

In conjunction with the Mic Gain preset on the IF board (RV5), this should be set to provide adequate input to the software VOGAD. The latter will hold the audio amplitude at a substantially constant level even in moments of excitement.

Bass and Treble Boost

These act independently to tailor the transmitted audio profile.

RF Clipping

I have always found this the most effective form of SSB processing – as opposed to audio compression. It increases the average power while holding the peak power steady. In mechanical engineering terms, it increases your transmitted signal's power-to-weight ratio. So it also increases the strain on your power supply, linear and ATU.

Use it only sparingly and when necessary (and not because it is there), bearing in mind that any form of processing – by definition – introduces distortion.

CW Spot Level

A 'spot' tone (equal in frequency to your CW offset) may be injected into the receiver output. As you net onto an incoming CW signal you will hear it beat with the 'spot' tone – and when they are on the same frequency, you are indeed netted.

This control alters the minimum amplitude. However, the amplitude is also increased automatically with the strength of the incoming signal. This is done because it is easier to beat two notes of similar amplitude – especially when the 'spot' tone amplitude tracks the inbound keying.

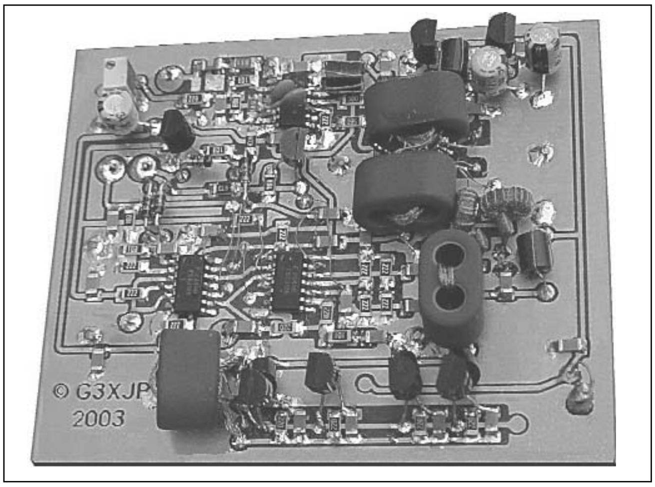

Fig 8.49: The 'Magic Roundabout' of G3TIE

THE FRONT END

How Good?

When it comes to receiver front-ends, there are those who would scale the highest IP3 mountains, and those who explore the depths of classic simplicity and minimalism.

Each pursuit is valid and fascinating in its own right. I am typically to be found sitting on a fence, aware that better performance is always achievable – but unclear if it is of real operational value. Aware also that the law of diminishing returns cuts in exponentially when it comes to cost and complexity.

Conversely, faced with the delightful quality emanating from the STAR DSP, it would have been remiss not to provide it with a proportionate frontend. "How good?" is, as ever, the question.

I much enjoyed 'HF Receiver Dynamic Range: How Much Do We Need?' [16], not least because it was written in practical and tangible terms (see also [17] for a summary).

I conclude from this paper:

- Phase noise performance is critical especially when there are many unwanted strong signals in the pass-band.
- The strong-signal dynamic range (DR) requirement is not horrendous – if you are prepared to shift the DR up and down to suit conditions. That is, sometimes you need good sensitivity; sometimes you need good strong-signal performance. Rarely, in practice, can you use both. So I don't think I can use all the dynamic range offered by, for example, the CDG2000 design [18].

This argument assumes absolutely that you are not blessed with a specific point-source problem such as a strong nearby transmitter. Normally propagated signals are assumed to apply. In real life – in Europe – the ability to handle 40m at night is the pragmatic test.

Pic 'N' Mix has intrinsically superb phase noise performance. The trick is to avoid degrading it (eg by adding a crude PLL) to get round the issue of DDS spurs.

The strategy of moving the DR up and down is equally compatible with the need to extract the maximum possible range from a finite number of bits in down-stream DSP.

STAR achieves this with the ability to reconfigure the receiver gain distribution dynamically.

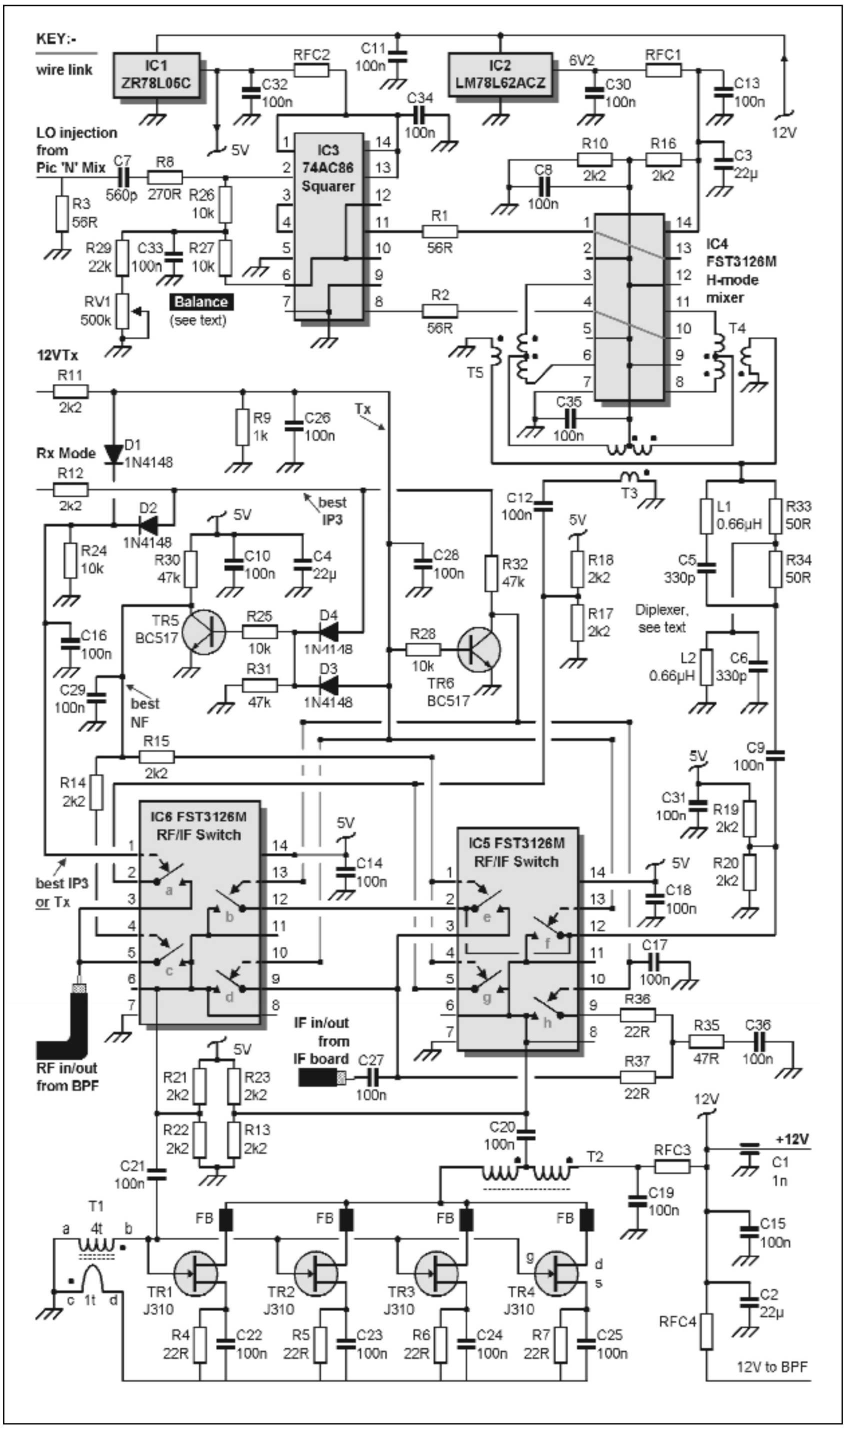

Fig 8.50: Dynamically configurable H-mode mixer – and quad J310 amplifier, better known as the 'Magic Roundabout'. The 'Rx Mode' line is set to +5V ie logic '1' for best intercept (IP3) – or to logic '0' for best noise figure (NF). 12VTx – when taken to +12V – selects the transmit configuration. The mixer requires fundamental frequency injection between -10 and +10dBm

THE MAGIC ROUNDABOUT

With early STAR, I used a mere SBL-1 mixer both with a bi-directional 2N3866 post-mix amplifier – and with a bi-directional J310. They both 'work'.

But, in 1998, Colin Horrabin, G3SBI, showed that there are no technical excuses for not using an H-mode mixer [19] – and Giancarlo Moda, I7SWX, showed that, with a fast bus switch, it can be truly affordable – except – in my view – for those very expensive transformers.

However, I can't turn away the opportunity of an H-mode mixer with home-brew transformers since, although this degrades the intercept somewhat, there is still plenty in hand. And no great cost.

In fact, there is so much in hand that I contemplated the heresy of fitting an RF amplifier before the mixer. Or should I settle for less sensitivity and instead fit an IF amplifier after the mixer? I wanted both options; and

Resistors 1206 SMD	
R1-R3	56R
R4-R7, R36, R37	22R
R35	47R
R33, R34	each 2 off 100R in parallel to give 50R
R8	270R
R9	1k
R10-R23	2k2
R24-R28	10k
R29	22k
R30-R32	47k
RV1	500k multi-turn preset pot
Capacitors	
C1	1n feedthrough
C2-C4	22μ radial electrolytic
C5, C6	330p ceramic plate (see next month)
C7	560p ceramic plate
C8-C36	100n SMD 1206
Semiconductors	
D1-D4	1N4148 or similar
IC1	5V regulator, 150mA eg ZR78L05C
IC2	6V2 regulator, LM78L62ACZ
IC3	74AC86 SMD
IC4-IC6	FST3126M (IC4 could be FST3125M)
TR1-TR4	J310
TR5, 6	BC517 Darlington
Inductors	
L1, L2	0.66μH on T25-2 (see next month)
RFC1-RFC4	6t thin enam wire on type 43 FB
FB	small ferrite bead on J310 drain lead
T1-T6	wound with 32SWG self-fluxing copper wire on EPCOS (was Siemens) B62152A4X1 binocular core (available from ElectroValue)
T1, T6	details follow next month
T2	4 bifilar turns (wire wound at 5tpi)
T3-T5	4 trifilar turns (wire wound at 5tpi)

Table 8.6: Front end component list

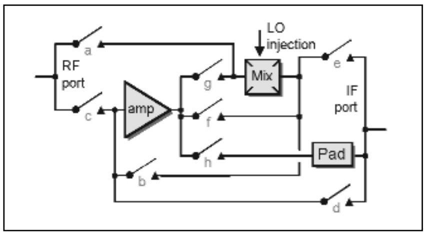

Fig 8.51: Modified squarer by I7SWX. Please note that neither the changed nor the incremental parts are included in the component list.

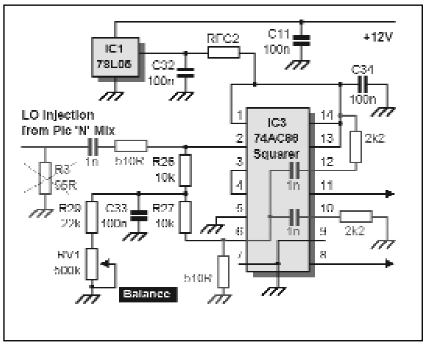

Fig 8.52: Magic Roundabout block diagram. For 'best NF', switches c, g and e are closed. For 'best IP3', switches a, b and h are closed and the pad improves the intercept. For 'transmit', switches d, f and a are closed

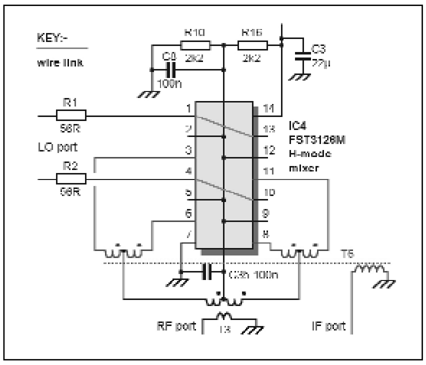

Fig 8.53: Two transformer H-mode mixer by I7SWX. Details of transformer T6 (which replaces T4 and T5) will be provided next month. It is wound on the same ferrite as the other transformers

after doing the performance arithmetic, I was convinced I needed both under different operational circumstances. Typically, on 10m I want the sensitivity; on 40m I want the higher intercept.

So, pragmatically, I decided to make it configurable – so that I indeed have both options and can compare and contrast them under differing real-life conditions – at the touch of a button.

Thus I can switch between 'best NF' and 'best IP3' modes to suit the prevailing conditions – with the key related benefit that I can have enough RF gain for DDS spurs to be below the band noise.

This approach takes care of T/R switching also – and the whole concept has become known in STAR circles as the 'Magic Roundabout'.

The circuit diagram is shown in **Fig 8.50** and the switching arrangements are summarised in **Fig 8.51**. The switch references are common to both.

Circuit Description

The LO injection is squared up and made symmetrical by IC3. Critically, this removes the even harmonics. RV1 is best adjusted for minimum DDS spurii on 10m. In 'best NF' mode these should be very hard to find. IC4 is a conventional fundamental-injection H-mode mixer. See also [19].

TR1-TR4 comprise a quad J310 amplifier used either as an RF or IF amplifier. I first saw the feedback arrangement in Introduction to RF Design by Wes Hayward, W7ZOI; and the use of multiple FETs (to raise the intercept) in an IF amplifier design by Bill Carver, W7AAZ.

With an 8dB pad (R35 – R37) in 'best IP3' mode only, the system gain is essentially constant in either receive mode. On transmit, the J310s are always used as an IF amplifier – irrespective of the receive mode. IC5 and 6 with TR5 and TR6 control the roundabout switching.

I7SWX Improvements

Fig 8.50 uses the original squarer and mixer from [19] – but in private correspondence with Giancarlo in early 2003, he suggested two improvements. You may wish to incorporate either or both. I certainly have and commend them.

Fig 8.52 shows changes to the squarer which improve the symmetry and 'squareness' of the switching waveform. Also, IC3 now does not require (nor gets hot in the absence of) LO drive.

Some people have reported unstable lumps of RF energy apparently emanating from the original squarer – and with this modification I have seen/heard no further evidence of them.

Giancarlo omitted the balance arrangements; but I prefer to retain them for the STAR application. The LO injection level requirement is between 0dBm and +10dBm.

Fig 8.53 shows Giancarlo's two-transformer mixer. This is to be preferred in principle, since when it comes to improving the mixer intercept, the only really good ferrite transformer is an eliminated one.

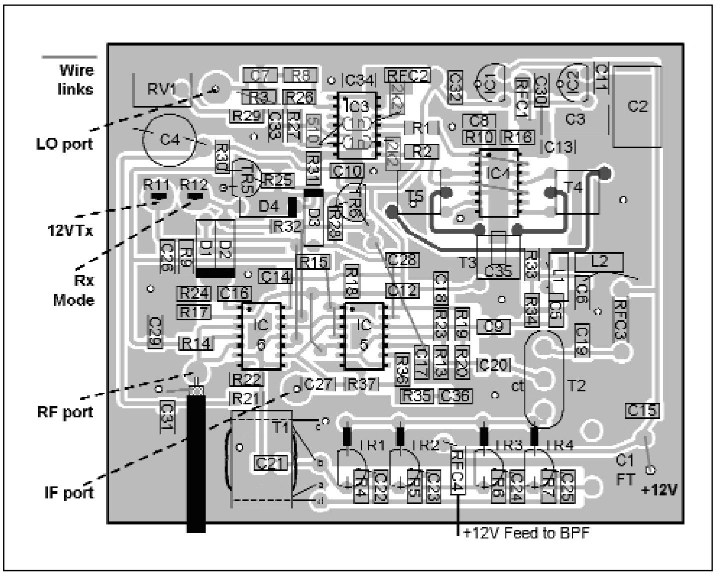

Fig 8.54: 'Magic Roundabout' component layout (see Fig 8.55 in Appendix B for the PCB artwork). This board is designed to mate mechanically with the BPF assembly described next month. Components shown in blue are for the improved squarer. Track shown in blue may optionally be removed (see text) for the two-transformer H-mode mixer

MAGIC ROUNDABOUT CONSTRUCTION

The component layout is illustrated in **Fig 8.54**, and the PCB artwork is in **Fig 8.55** (in Appendix B). It is again made using the iron-on process. The board is double-sided with the underside unetched except for small pads to mount R11 and R12 - and to make off the Rx Mode and 12V Tx lines.

Thus the underside of the board provides a ground-plane as well as screening. R11 and R12 are mounted in the thickness of the PCB as 'feedthrough resistors'.

There are several wire links on the board, for which I apologise. These derive from the need to make this board as small as possible to minimise the risk of pick-up and radiation.

The RF Port requires a DC blocking capacitor. For STAR, this is fitted on the BPF board. Do not omit it!

Track Options

The board is shown tracked for all options. If you are definitely fitting the two-transformer H-mode mixer, you should remove all the tracking between T4 and T5 and T3 - leaving only a small pad to connect a wire from T6 to the IF port diplexer. To 'remove' the track before etching, simply fill and join up the ground with an indelible pen. Alternatively, after etching, you can cut the tracks feeding the diplexer - and then bridge all the unwanted track to ground. The latter is easy and reversible so gives you more options for experimentation.

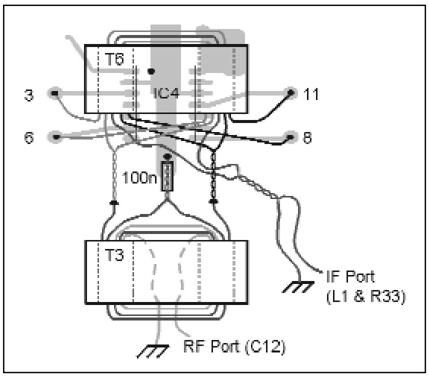

Fig 8.56: Details of T6/T3 connections to each other and to IC4. This is not strictly to scale, but does show the correct relative positioning of the components. T6 and T3 will end up separated by about 6mm in practice

Squarer Options

These are shown in dark grey on Fig 8.54 with the component values taken from Fig 8.52.

This is a classic example of where there is absolutely no way this board could be made commercially, since there is no suitable track layout. For one-off purposes, however, the two 1n wire-ended capacitors fit beautifully - as shown - across the top of IC3. If you want to retain the original mixer configuration, simply replace them with wire-links; omit the dark-outlined components; and revert to the values in Fig 8.50.

Construction Sequence

The holes in the grounded areas of the board are for links through to the ground-plane. These should be fitted first and soldered both sides.

Then mount all the SMD components; then the discrete devices with the exception of TR1; then the wire links (ideally using thin self-fluxing wire) and finally the transformers.

Note the gap between TR2 and TR3 to give space for the feedthrough to the BPF board - via RFC4. Slip a small ferrite bead over each J310 drain lead before soldering.

IF Port Diplexer

L1/2, C5/6, R33/34 form a diplexer on the mixer IF port. L1, L2, C5 and C6 should each have a reactance of about 50 ohms at the IF - using the nearest preferred capacitor value. The values given are for 10.7MHz. L1 and L2 in this case are each 14 turns of 32SWG wire wound on a T25-2 core spread over about 2/3 of the circumference.

This derives from an A_L value of 34µH per 100 turns for this core. Before fitting, connect each coil in parallel with its resonating capacitor, pass a single turn through the toroid and loosely couple it to a GDO; and dip it at your IF.

Mixer Transformers

T3, T4 and T5 are illustrated as MCL transformers and the tracking is appropriate should you want to use these.

If using the home-brew 3-transformer mixer, the EPCOS ferrites mount vertically as shown for T2.

Mixer balance is critically dependent on the transformers all being the same. To this end, make up enough trifilar wire in one length for T3-T5. Rather than winding one end of the wire through the core continuously, wind alternate ends.

If building the two-transformer version, then see Fig 8.53 for the circuit diagram and **Fig 8.56** for the mechanical result. This may indeed turn out to be the definitive test of your understanding of 'the phasing dot convention' but if you follow the steps below you can build it by rote.

T6 is wound with five parallel strands of self-fluxing wire - untwisted. Cut five lengths of wire to some 20cm long. Solder one end of all of them together to retain them. Then trim them all to exactly the same length and solder the other ends together.

Now wind four turns on the core under modest and continuous tension - passing alternate ends through the core.

Cut off the surplus wire equally and initially to approximately 5cm. Tin all 10 ends. Using a continuity meter, locate one pair and twist them to form the IF port feed. That was easy!

Now locate two more 'pairs' and cross-connect a start/end (and twist them together) to form the centre-tap. Repeat for the two pairs.

Locate T6 on top of IC4 and trim the four leads that are made off to the track to the same length. Those going to pins 6 and 8

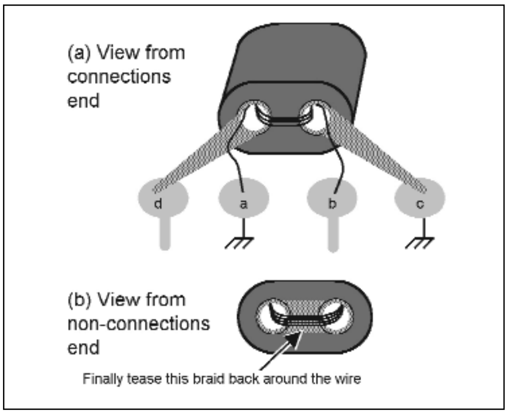

Fig 8.57: Transformer T1 detail. The pad designations correspond to those of Fig 8.50

pads define that length. Tin and solder them to their pads as per Fig 8.56.

Then make off the IF port leads to the end of the diplexer.

Now wind the trifilar transformer T3. Solder the centre-tap to the track first, then the RF port feed. Finally, trim off the flying leads to T6 equally, as short as reasonably possible, and solder them up.

J310 Transformers

T2 is a conventional bifilar transformer with no complications. The input transformer T1 needs a little explanation. See **Fig 8.57**.

The T1 feedback turn is made from a length of miniature coax braid. Form the U-shaped primary turn - with plenty of excess lead length. Pierce and spread the braid to fill the tubular holes in the core - from both ends of the core. Make off the braid leads to the PCB track.

Wind the four turns - inside the braid - out the back of the core, across and back through the braid on the other side - and solder to the track. As far as possible, tease the braid back round the turns on the non-lead end (as W4ZCB says, so nobody can see how it was done).

TR1 may now be fitted.

Incremental Testing

Depending on your personal style and confidence, you may want to perform the functional testing of this board progressively. By lifting one end of some RF chokes you can selectively enable parts of the circuit. 100n wire ended capacitors should be used to couple RF in and out of each circuit element under test.

You can test the J310 amplifier is indeed amplifying by placing it in the down-lead of some receiver. Equally, the squarer and mixer - with suitable injection - can be tested as a crude converter.

System Interface

You could control the Roundabout with a simple switch and a status LED. Arguably, you could do worse than driving it from the SSB select lines, ie ‘best NF’ on the USB bands, ‘best IP3’ on the LSB bands. But 60m is exceptional and complete user choice is desirable until you see how it works for you. It depends not least on what antennas you have. I control it from the ‘spare’ output on Pic ‘N’ Mix - with a 1k5 resistor in series with a tell-tale LED across the line. On STAR, this line is toggled by keypad sequence 48.

STAR PERFORMANCE

These performance measurements were made by Harold, W4ZCB, using professional test equipment. They were corroborated by the author using test equipment borrowed from I7SWX - and his own. Thanks go to both because these are very important numbers.

All measurements relate to the complete STAR line-up - ie including the band-pass filters which follow next month. These have, by design, decreasing insertion loss with increasing frequency. So, to get the complete picture, you need to consider the performance on each band. Four representative bands are shown, the rest being somewhere ‘in between’. Since the design features a ‘best IP3’ or ‘best NF’ mode, that adds a further dimension.

The mixer is Giancarl’s two-transformer topology driven by the modified 74AC86 squarer. MDS was measured in a 3kHz bandwidth; IP3 at 20kHz tone spacing.

‘Best IP3’ mode			‘Best NF’ mode		
Band (m)	MDS (dBm)	IP3 (dBm)	Band (m)	MDS (dBm)	IP3 (dBm)
80	-123	+33	80	-127	+30
40	-122	+35	40	-123	+30
20	-124	+31	20	-127	+28
10	-127	+28	10	-130	+25

AGC RANGE

AGC holds the audio output constant within 1dB for a 100dB change of signal. You can place this range anywhere on the amplitude scale by adjusting the RF Gain, but from -95dBm to +5dBm would be typical.

OBSERVATIONS

Note that excess sensitivity is not provided on the lower bands where it could never be used. Instead, it is traded for superior strong-signal performance.

For comparison with commercial transceivers see [20].

Finishing Off

When all is working well, trim RV1 to minimise DDS spurs on a high band, eg 10m, in ‘best NF’ mode. When connected to a dummy load, you may still be able to hear some. But then, when connected to even a modest antenna, although the band may be essentially closed, the ambient band noise should mask them to the point that you will need to try very hard to find them.

Finally, enclose and screen the whole board and re-trim RV1.

10-BAND BANDPASS FILTER

The author felt a need to improve on the Third Method front-end (and add 60m) and could find nothing suitable in the literature that met the requirements. This development is of general interest and application.

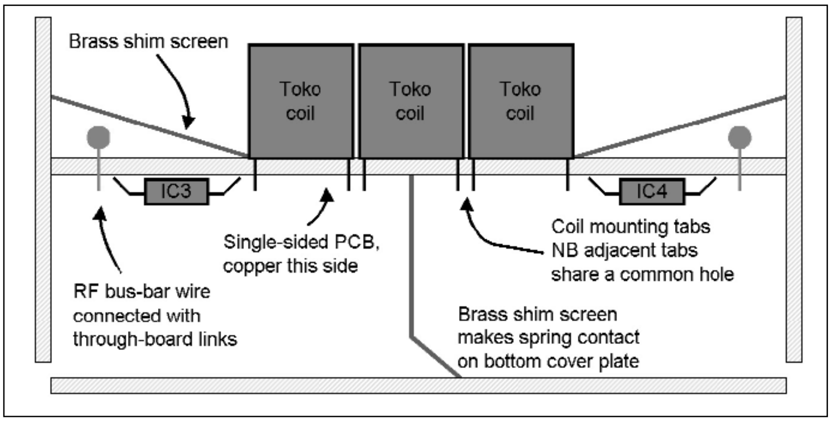

Fig 8.58: Cross-section of bandpass filter enclosure. The main board is single-sided with critical incremental screening on the top provided by brass shim. The same material is used to form a central spine shield running up the middle of the filters in the lower compartment. It is bent over to make spring contact with a bottom cover plate - made from PCB stock

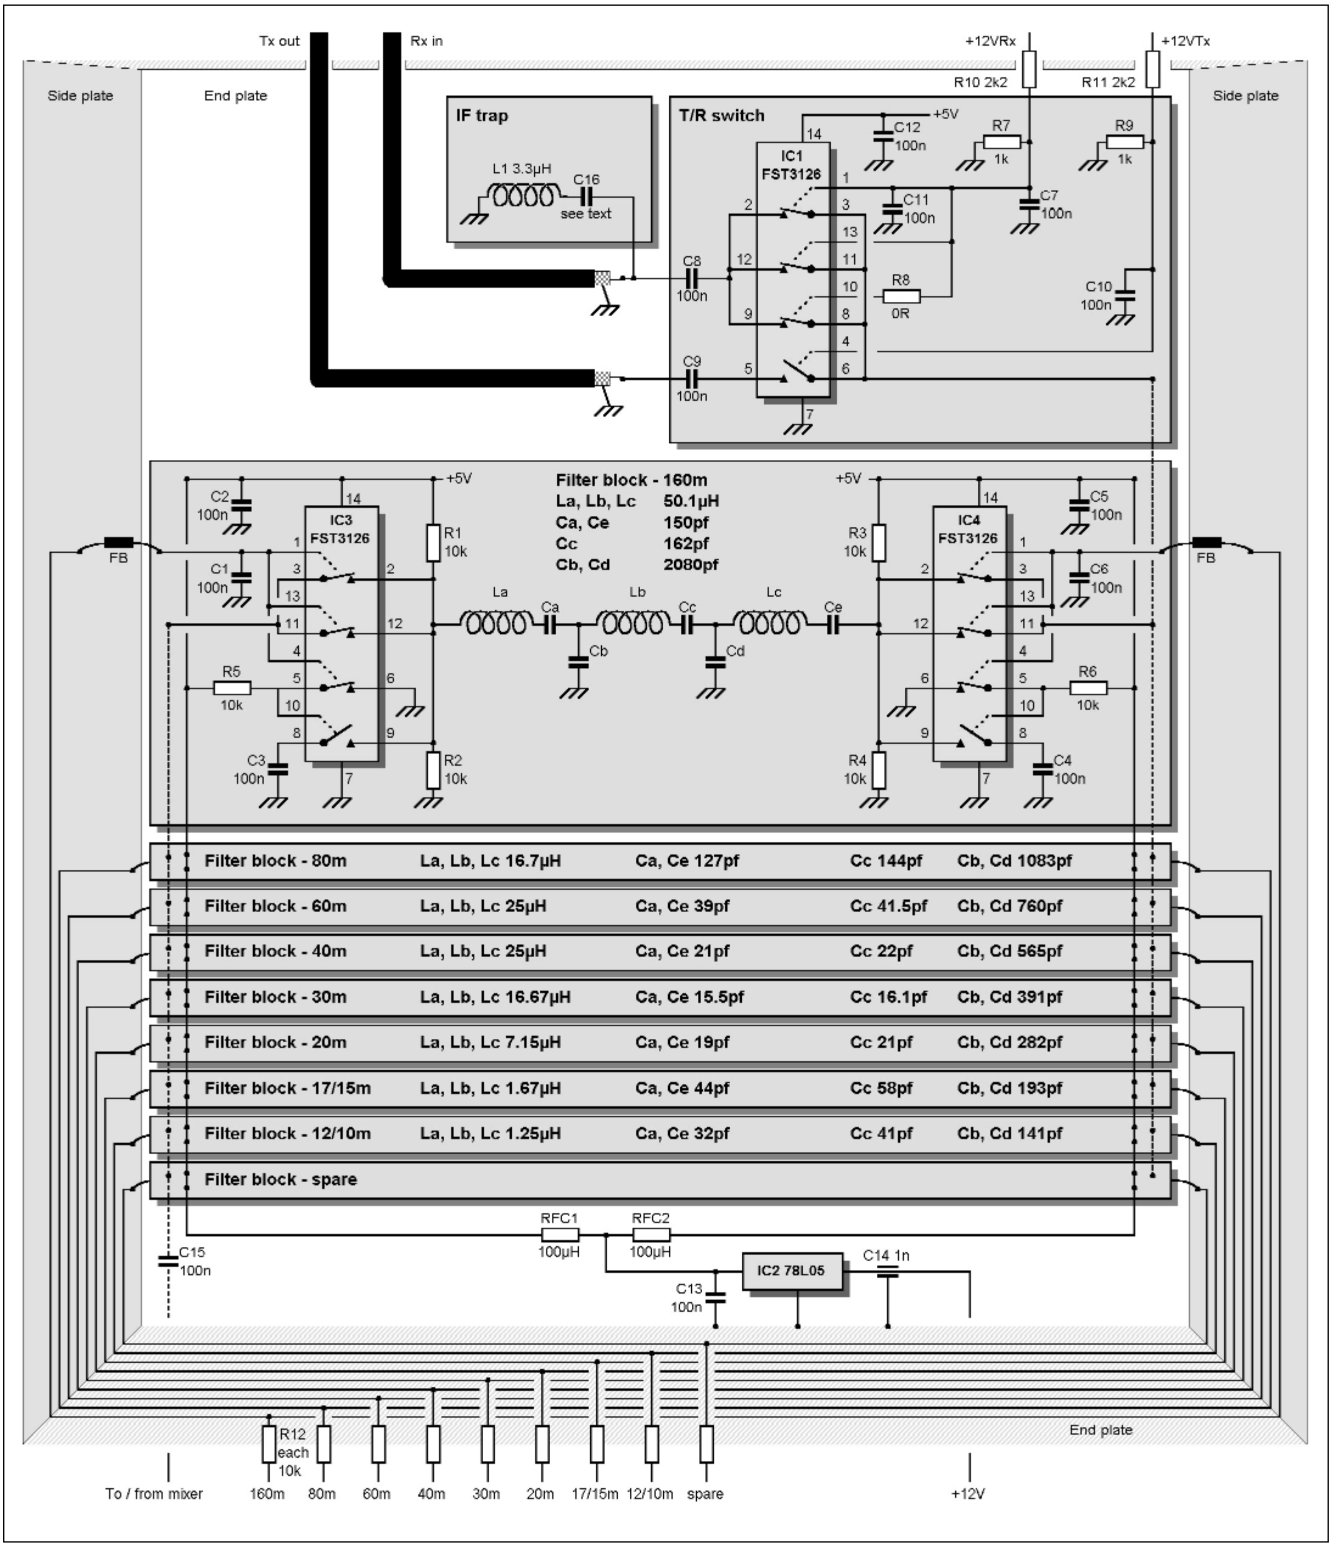

Fig 8.59: Bandpass filters circuit diagram and mechanical overview. All the filters have identical circuit diagrams. A band-select line is set to +5V to engage a filter block and to 0V to isolate it. This is compatible with direct interfacing to the Pic 'N' Mix band-select outputs. These lines are deliberately routed around the side-plates and not across the filter. For illustration, the 160m filter is shown 'on' and the T/R switch is 'on receive'

Design Aims

Like everyone else, I think I want narrow filters with no insertion loss, superb IP3 performance - and acceptable cost. I definitely want a finished size that does not impact on the overall dimensions of my transceiver. For a fact, you can't simultaneously optimise all these parameters, so this design - like all others - is a careful compromise.

The prime function of any BPF is to reject the image frequency adequately - and it is achieved by this design (with a little help from an ATU and/or your LPF on the highest bands), provided your IF is 9MHz or higher - and you use high-side injection.

You could use different filter topologies. The mechanical construction is for three inductors (see **Fig 8.58**). That is the only

FOR THE OVERALL BPF ASSEMBLY		PER BPF FILTER BLOCK	
Resistors 1206 SMD		**Resistors, 1206 SMD**	
R7, R9	1k	R1-R6, R12	10k
R8	0R link (or wire!)	**Capacitors**	
R10, R11	2k2	C1-C6	100n 1206 SMD
Capacitors		Filter capacitors Polystyrene or silver mica. For values, see Fig 8.59.	
C7-C13	100n 1206 SMD	**Inductors for filters (all Toko coils, 3 off)**	
C14	1n feedthrough	160m	154ANS-T1017Z
C15	100n disc ceramic	80m	154ANS-T1012Z
C16	see text	60m	154ANS-T1014Z
Inductors		40m	154ANS-T1014Z
RFC1, RFC2	100µH axial choke	30m	154ANS-T1012Z but see text
L1	Toko BTKANS-9445HM 3µ3	20m	154ANS-T1007Z
Integrated circuits		17m/15m	TKAN-9448HM
IC1	FST3126M	12m/10m	BTKANS-9450HM
IC2	78L05 regulator	**Integrated circuits**	
		IC3, IC4	FST3126M
		Miscellaneous	
		Small ferrite bead, 2 off	

Table 8.7: Band-pass filter components list

practical constraint. The filter capacitors are soldered directly to the coil terminals and there is plenty of room for lots of them in different configurations.

Performance requirements

Since, on the higher bands, Noise Figure is everything, I need low insertion loss above all else. Say 1.5dB. The consequence is wider filters which have the significant benefit of spanning more than one band.

The other great benefit - and integral to the whole front-end design strategy - is that I need the highest possible signal level going into the mixer so that any DDS spurii are below this level on the higher bands.

As you move down in frequency, Noise Figure becomes less and less important and greater insertion loss is a positive benefit, adding directly to the mixer intercept. It also helps to keep the power output flat on transmit if it tends to drop off at the higher frequencies.

Cost considerations

On the cost front, diode switching is the cheapest - but unacceptable for strong-signal performance. Good relays (and you would not want to use bad ones) are very expensive given the quantity involved. I settled on integrated bus switches and use the FST3126. This is a close relative of the FST3125 as typically used in the H-mode mixer, the difference being that the switch control logic is inverted and is compatible with direct drive from the Pic 'N' Mix band-switching latches.

As I have configured them, the ON insertion loss is less than I can measure, at well under 0.5dB; the OFF isolation is better than 90dB; and IP3 is better than 40dB. They are very inexpensive.

I settled for Toko coils, because they are readily available and also inexpensive - but I used the larger cup-core inductors where possible, ie on all bands up to and including 20m.

As a gross alternative, you could use fixed toroidal inductors and trimmer capacitors if the Toko coil Q or IP3 performance are issues for you.

Circuit Description

The filters (see Fig 42) are all 0.01dB Chebyshev designs. They are switched by identical switches at each end (IC3 and IC4). For each switch, two sections are paralleled to reduce 'on' insertion loss; one section inverts the control logic and one section grounds the filter when 'off'. For the T/R switch (IC1), three sections are paralleled on receive and one is used on transmit.

L1 and C16 form a series-tuned IF trap and C16 should be chosen to resonate with L1 at your chosen IF frequency.

A spare filter position is provided for experimental purposes, eg different topologies, different frequencies.

The filter capacitors' exact theoretical values are given in **Fig 8.59**, and the nearer you can get to them, the better. I obtained a large bag of assorted polystyrene capacitors and arrived at the values to within 1pF (as measured on my DVM) with never more than two in parallel - by measuring the actual values within the tolerance range.

For example, the 40m and 20m blocks require six capacitors of near 20pF. The exact values were found from a small selection of 20pF and 22pF 5% capacitors.

Performance summary

Up to 20m, the filters give >100dB image rejection. By 12m/10m, this has fallen to some 50dB, so you could benefit from some incremental filtering provided by low-pass filters, an ATU, or even a beam.

Insertion loss is around 5dB on the lower bands, falling to 3dB on 20m, 1.6dB on 17m, 1dB on 15m, 1.5dB on 12m and a delightful 1dB at 28MHz rising to a mere 1.3dB at 29.7MHz.

30m is exceptional because of the proximity to my IF. Here a 5dB inband insertion loss rises to 18dB at 10.7MHz and 55dB at 9MHz. Depending on your mixer balance and trap tuning, this could be marginal with a 10.7MHz IF.

Filter Construction

The layout is shown in **Fig 8.60**. Three filters are illustrated, the remainder being constructionally identical. More space has been allocated to the 160m and 80m filters to accommodate larger filter capacitors. The components are all mounted on the track-side except for the Toko coils, the RF in/out bus-bars (22SWG tinned wire) and C15. The bus-bars are connected to the IC switches with throughboard wire links. Note that the middle Toko coil is rotated by 180 degrees relative to the outer two. Unused coil pins are cut back so they do not appear on the opposite side.

The PCB artwork is shown in **Fig 8.61** (in Appendix B). This assembles into an H-section brick (for the want of a better term). The filter board is single-sided - and SRBP if you want to save on drill bits; the side and end-plates are doublesided. Ensure opposite sides of all the double-sided boards are intimately connected. The outside faces of both end-plates have simple oval pads

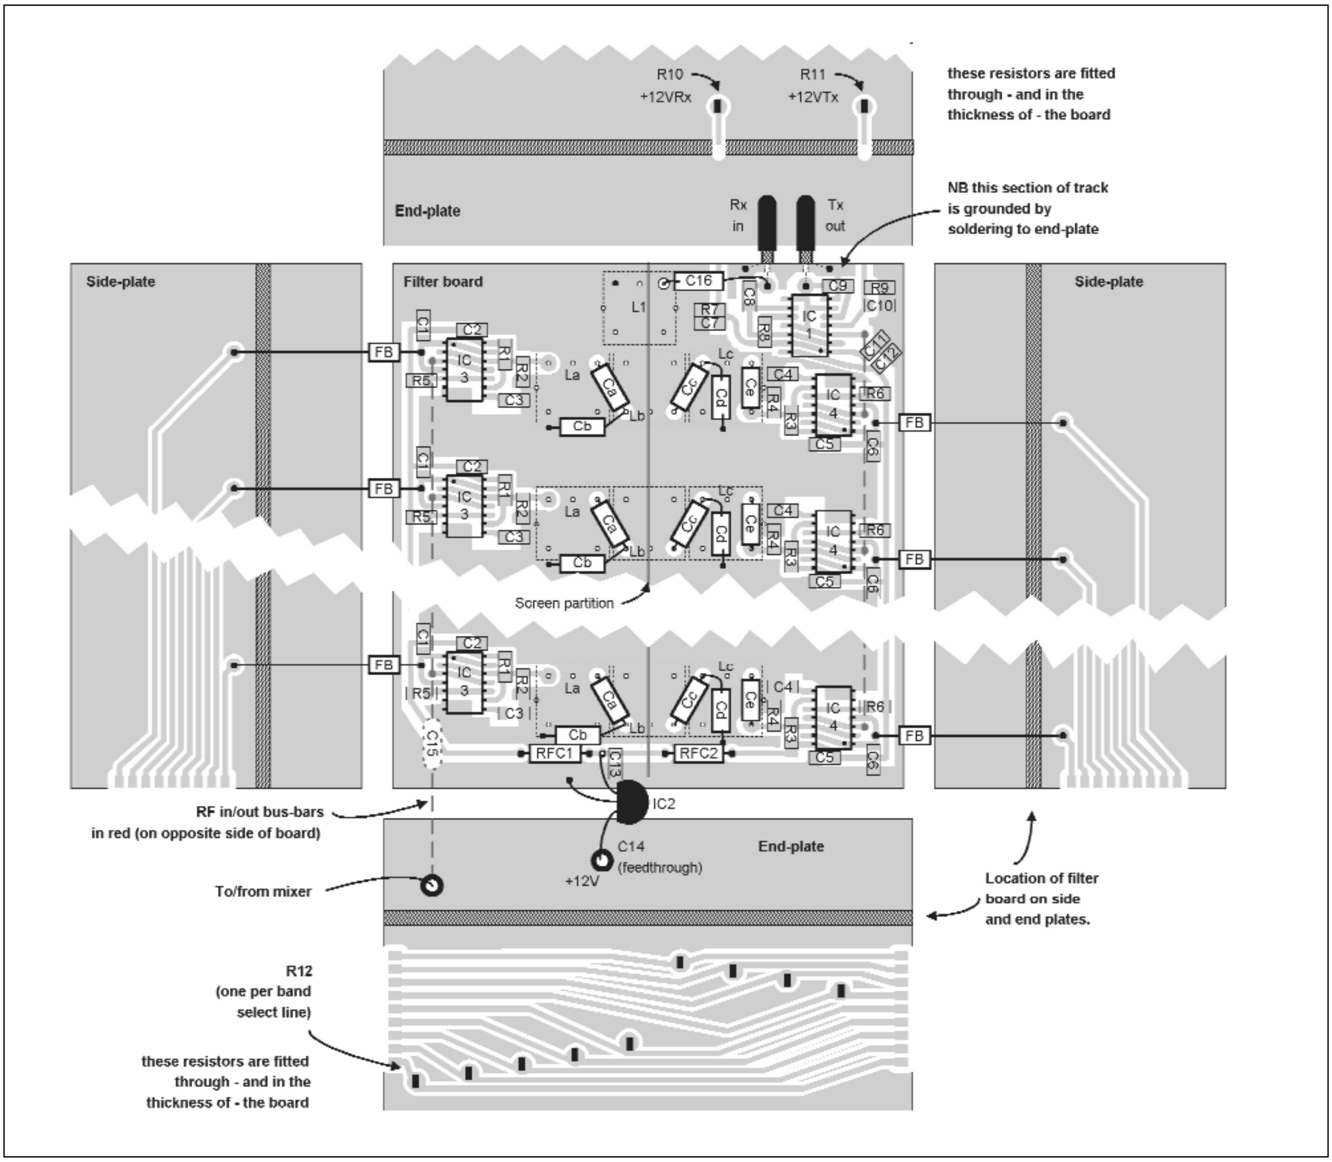

Fig 8.60: Component layout for band-pass filters

to make off the feedthrough resistors to the band-select lines at one end - and the 12VTx and Rx lines at the other. The artwork for this is not provided since some elementary removal of the masking spray after drilling and prior to etching will achieve the desired and non-critical result.

Filter Adjustment

In the first instance, put the entire filter assembly in series with the antenna of an existing receiver and check that all the filter switches work and the coils peak. There are sophisticated ways of tuning these filters, but if all coils are peaked mid-band and then the two end coils are peaked at the two band edges, you will not be far out. A refinement of this is to put Pic 'N' Mix in wobbulator mode across the segment of interest and, while on low power CW, transmit into a dummy load and tweak for a flat pass-band.

Acknowledgements

These filters were designed by Harold, W4ZCB, to whom I am much indebted. He in turn credits ELSIE, a filter design package by Jim Tonne, WB6BLD - which did all the sums. I have subsequently become a fan of this software.

Harold also independently measured the switching performance.

DIY DSP FILTERS

There is not much prescriptive that can be said here. The trick is to search the web for a program that will generate coefficients for FIR filters. From time to time these are available in the public domain. These coefficients then need converting to a format suitable for loading.

Experimental Methods in RF Design, by Hayward, Campbell and Larkin is also a useful source of understanding - pages 10.13 to 10.19 in particular. This tells you how to do it - and provides the software on a CD to let you.

I suggest the best way to prove the process in the first place is to see if you can build the existing medium-width SSB filter (FL5), plug it in and prove to yourself that yours is no different.

To this end, all the Rx filters are packaged individually and discretely. The specification for the existing FL5 is:

Sample rate = 8kHz

Remez equiripple - but not mandatory

Order = 198, ie 199 coefficients

Summary: modularity is everything

From the outset, this project has been modular by design. Both in the sense that several people have substituted different blocks of hardware and software in their STARs and, conversely, have used STAR blocks in other applications.

In summary, and in chronological order, the project and its description comprises:

- A discussion of STAR hardware integration possibilities.
- A discussion of software options and flexibility.
- A generalised process for producing precision one-off PCBs.
- A T/R changeover timer that would suit any transceiver.
- Details of a DSP processor assembly that could be the basis for any DSP project, with modular daughter boards to allow future upgrades.
- A completely repeatable generalised process for mounting SMD ICs with lots of closely-spaced (eg 0.5mm) pins.
- A bi-directional IF strip that could be readily adapted for use in any home-brew design.
- A PC-based loader and controller that could be adapted for any Analog Devices 218x DSP project.
- One of a number of possible physical implementations. A glance at **Fig 8.62** shows you that every Betabuilder exercised completely different options here. There are no two the same.
- An adapter to replace an 18-pin PIC to give greater I/O and more processing capability generally, and a bargraph S-meter.
- A spur reduction filter for any DDS, and a stable reference oscillator.
- A useful DSP shopping list, at the least. Check out the competition!
- A general purpose bi-directional mixer/amplifier with configurable topology. Of the strong and silent type.
- A universal front-end that could drop into almost any existing HF transceiver.

Table 8.8: This chapter has dealt with an entire transceiver, but in such a way that many of the modules and techniques can be used in other projects

Passband = 320Hz - 2284Hz

Lower stop = 177Hz

Upper stop = 2400Hz

Once you have replicated this and proved the process, you can rapidly produce filters for any specialised application, eg RTTY.

CONCLUSION

A summary of the features of this project are shown in **Table 8.8**. You don't have to be building a STAR to find something of interest here. Equally and oppositely, you don't have to build it all to benefit from STAR DSP.

If you just want to build an error-free transceiver design that works beautifully, now is the moment. All the fruits of this development are available to you at no charge, provided only that they are for your personal use.

To obtain all the software - including the source code - follow the process given earlier. For all the PCB artwork, any enhancements and all ongoing support, simply join the Interest Group (**see box**).

You don't need to understand DSP. At least not to use my STAR code and get your transceiver going. Thereafter, it is entirely both up to you - and down to you.

STAR Build Sequence

Most receiver designs start from the antenna and logically follow the signal flow to the loudspeaker. This project has taken (more or less) the opposite approach. This is to encourage you to build the trickiest bits first.

I think it is a better strategy anyway, since once you have some sort of noise coming out of the speakers, you can work back towards the antenna, using the completed elements to test the new build. I commend it to you.

Once you have the STAR receiver working, the few incremental components to get the transmitter going can be taken from any of the many HF designs. You need a driver, PA and low-pass filters - commensurate with the amount of power you want to run.

PIC-A-STAR web interest group

http://uk.groups.yahoo.com/group/picastar/

The Linearity Conundrum

Why does a STAR sound so good both on receive and transmit? There can only ever be one answer. It is because it is linear.

Because STAR is an IF processor it has a built-in head start over any AF add-on. The latter may indeed be better than nothing, but if the DSP has neither control of the AGC nor of the detection process then the damage has already been done - so to speak - before there is any chance to benefit from subsequent DSP.

Thereafter, most of the 'star quality' derives from the inherent nature of digital (as opposed to linear analogue) signal processing.

With analogue processing, the quality is determined ultimately by device linearity. Any non-linearity - and all devices have some - results in intermodulation distortion products which can fall within the pass-band. And to a varying extent they grate on the human ear - which is particularly sensitive to their presence.

With digital processing, the nature of any distortion is quite different. It arises from rounding errors, quantisation errors and lack of arithmetic precision generally.

As long as the system is not grossly overloaded (in which case it would fall apart in a big way) these errors do not result in discrete in-band IMD products.

Rather, they result in a general noise floor - of trivial amplitude. And in any practical HF radio communication system, this noise is indistinguishable from and is buried well below the band noise. There are simply no conventional IMD products to hear.

So the answer to the conundrum is this - and it may not be instinctive. Digital signal processing is inherently more linear than linear analogue processing. This is the technical rationale for the PIC-A-STAR project.

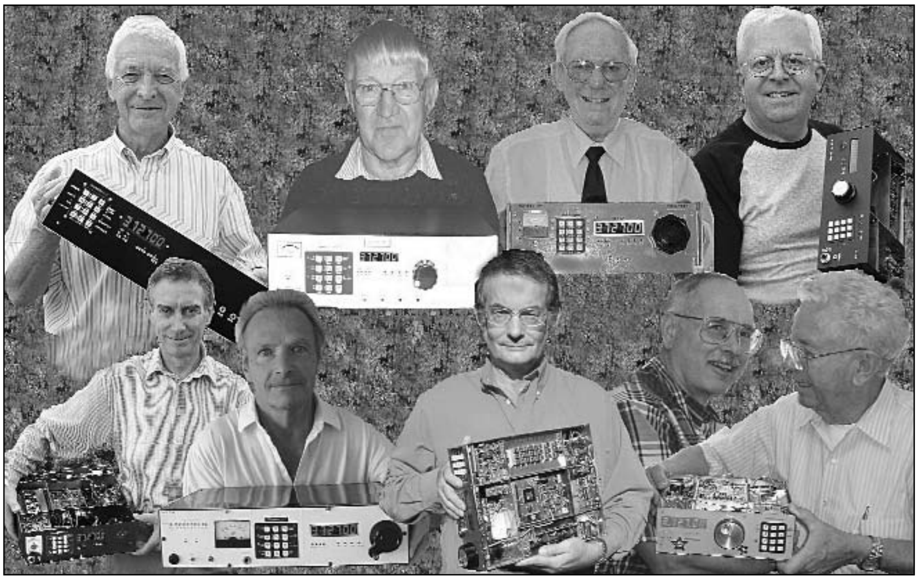

Fig 8.62: The Constellation Beta. Starry-eyed and legless after their much-acclaimed performance in the accordion band competition. They are (from left to right): Top row Alan, G3TIE; Harry, G3NHR; Les, GW3PEX and Peter, G3XJS. Bottom row: David, G4HMC; Eddie, G0SEY; Peter, G3XJP; Bill, W7AAZ, and Harold, W4ZCB, with Harold's STAR

ACKNOWLEDGEMENTS

There are lots of acknowledgements, since the PIC-A-STAR project was and continues to be a truly collaborative and international effort. They are in summary and in no particular order:

The Original Beta Team

The photograph above (**Fig 8.62**) shows the original team that built, evaluated, tested and continuously suggested improvements - for both the hardware and software of PIC-A-STAR.

Infrastructure and Utilities

My thanks to David Tait for his latest and greatest TOPIC (PIC programming software); Jim Tonne, WB6BLD, for ELSIE (filter design software); Analog Devices for their DSP utilities and code fragments.

Inspiration and Actual Help

Lee, G3SEW, for much useful discussion on the UI in the early days; Gian, I7SWX, for use of his test equipment and mixer design.

Bill, W7AAZ, for many, many ideas and encouragement; Harold, W4ZCB, for the design of the front-end filters, for much performance evaluation - and his unwavering enthusiasm.

Keith, G3OHN, Paul, G0OER, Mike, G3XYG, Jim, G3ZQC, Michel, ON4MJ, and John, G6AK, for much building and testing and the benefit of their diverse skills and wide-ranging experience.

Fran for the proof reading. George Brown, M5ACN, the Technical Editor of *RadCom*, for steering all this into print

And last, but by no means least, our thanks to Bob Larkin, W7PUA, for sharing his original DSP-10 work, the Digital Signal Processing chapters in *Experimental Methods in RF Design*, the adaptation of the STAR boot code, his advice and suggestions - and the ultimate inspiration for the whole STAR project.

REFERENCES

[1] *RadCom*, Aug 2002 - Mar 2004, RSGB
[2] Peter Rhodes, G3XJP, Danvers House, Wigmore, Herefordshire HR6 9UF. E-mail: G3XJP@qsl.net
[3] Analog Devices. http://www.analog.com
[4] Bob Larkin, W7PUA. http://www.proaxis.com/~boblark/dsp10.htm
[5] 'HF Third Method SSB Transceiver', Peter Rhodes, G3XJP, *RadCom,* Jun -Oct 1996
[6] 'Multiband HF Transceiver', Mike Grierson, G3TSO, *RadCom*, Oct - Nov 1988
[7] 'Pic 'N' Mix Dgital Injection System', Peter Rhodes, G3XJP, *RadCom*, Jan - May 1999
[8] Analog Devices FTP download ftp://ftp.analog.com/pub/dsp/21xx/218x/ez-kit-lite/ (You need the two files entitled disk1.zip and disk2.zip)
[9] ElectroValue Ltd., Unit 5, Beta Way, Thorpe Industrial Park, Egham, Surrey TW20 8RE, England. UK. Tel: +44 (0) 1784 433604. Fax: +44 (0) 1784 433605. E-mail: sales@electrovalue.co.uk. Web: http://www.electrovalue.co.uk/
[10] 'In Practice', Ian White, G3SEK, *RadCom*, Feb 2003
[11] 'HF', Don Field, G3XTT, *RadCom*, Jun 2003
[12] 'A DSP-Based Audio Signal Processor', Johan Forrer, KC7WW, *QEX*, Sep 1996
[13] 'Low-Cost Digital Signal Processing for the Radio Amateur', D Hershberger, *QST*, Sep 1992
[14] 'DSP – An Intuitive Approach', D Hershberger, *QST*, Feb 1996
[15] 'Using the LMS Algorithms for QRM and QRN Reduction', S E Reyer and D Hershberger, *QEX*, Sept 1992
[16] 'HF Receiver Dynamic Range: How Much Do We Need?', Peter Chadwick, G3RZP, *QEX,* May/Jun 2002
[17] 'Technical Topics', *RadCom*, Feb 2003
[18] 'The CDG2000 HF Transceiver', Colin Horrabin, G3SBI, Dave Roberts, G8KBB, and George Fare, G3OGQ, *RadCom* Jun – Dec 2002
[19] 'Technical Topics', *RadCom*, Sep 1998
[20] 'Technical Topics', *RadCom*, Dec 2002 - Table 1, the IP3 column in particular