THE 56 KB/S MODEM AS A NETWORK BUILDING BLOCK: SOME DESIGN CONSIDERATIONS

Barry McLarnon VE3JF
2696 Regina St.
Ottawa, ON, Canada K2B 6Y1
Email: barry@dgbt.doc.ca

ABSTRACT

In response to the now-famous cry of "Where is my high-speed RF?" [l], Dale Heatherington WA4DSY introduced his 56 kb/s modem design [2] to the packet community in 1987. Four years later, this breakthrough development has yet to make the major impact on the amateur packet networking world that many of us anticipated when it was introduced. We begin by examining some of the reasons why this is so, and then describe several of the possibilities for deploying the modem in the amateur packet radio network. It is the author's opinion that this modem is a versatile building block which can serve a multitude of roles in the network; if this article helps to stimulate more interest in pursuing them, then it will have served its purpose.

Introduction

There are a variety of reasons why the 56 kb/s modem (hereafter referred to as the DSY modem) failed to take off as quickly as we had hoped. Probably the major one was the lack of packet switch hardware which could make full use of the modem's speed. TNC hardware could not keep up, nor were there, until recently, any inexpensive HDLC interface boards for the PC bus available which could handle 56 kb/s without usurping all of the CPU's resources. This problem was resolved in 1990, with the arrival of the Gracilis PackeTen packet switch [3] and the Ottawa PI I/O card on the scene,

Other impediments to the widespread acceptance and deployment of the modem include the fact that the modem is only available in kit form, and that it costs significantly more than low-speed modems and TNC hardware. Some people do not appreciate the fact that this modem is a great deal more than a low-speed FSK modem scaled up to run faster; others are apprehensive that it will be difficult to construct and tune up. in fact, the modem is easy to build, and tune-up is straightforward; the most time-consuming and tedious part of getting one on the air is mounting the boards in a box and providing the requisite interconnections and +/- 5V power supplies.

Another problem area was, and continues to be, a shortage of 'affordable RF converters for translating the modem's input/output at 28-30 MHz to a suitable frequency above 220 MHz. In 1987, the choice boiled down to either a Microwave Modules MMT220/28S transverter (if the 220 MHz band was available to you at all), or the same firm's MMT432/28S model. Since then, the company has turned away from the amateur market, although the transverters can still be found quite often on the used market. Sinclabs of Toronto is now producing a 220/28 unit, and there may be a new source of 432/28 MHz transverters soon. Those currently available, such as from SSB Electronics, are very expensive. Transverters between 28 MHz and bands above 450 MHz seem to be nonexistent (transverters for the higher bands normally use 144 MHz as the IF).

If the modem is used for full duplex or split (half duplex, but with different receive and transmit frequencies) operation, then separate receive and transmit converters are needed, and this opens up a few new possibilities. There are several sources of receive converters available, and at least one source (Hamtronics) of transmit converters. Above 450 MHz the same situation as for transverters exists. With separate converters, however, it should be more feasible to run converters in cascade or put together custom converters for the higher bands. The schemes outlined below, which would allow sharing of the converters by multiple modems, make it more acceptable to put some effort into assembling the necessary RF hardware.

The 56 kb/s CSMA MAN

In Ottawa, 56 kb/s packet began in mid-1988, with a link between VE3JF and VE3MDL on 220.55 MHz. We were soon joined by VE3OCU, and I found that I could not hit both stations reliably with a single beam heading. The idea of putting up a 56 kb/s repeater to support our new MAN (Metropolitan Area Network) was born soon thereafter! We decided that anything less than CSMA with no hidden transmitters would be unacceptable, so we began assembling a full-duplex repeater. Band plan limitations and the lack of a 220 MHz duplexer seemed to preclude doing an in-band repeater, so it was decided to make the repeater cross-band, with input on 220.55 Mhz and output on 433.55 MHz. At our home stations, this meant the addition of a new modem receive crystal, a 432/28 MHz receive converter, and a 432 MHz antenna.

The repeater (see Figure 1) is a bit regenerator, similar to the arrangement which some people have used at lower bit rates: the demodulated data from the receive side becomes the transmit data stream, and the transmitter is keyed by the demodulator's DCD output. In order to minimize the keyup time, we keep the transmit side of the modem running continuously (ETS asserted) and key only the transmit converter. One word of warning about this practice: the transmit clock is derived from a 3.579 MHz oscillator on the encoder board, and the third harmonic of this signal falls within the passband of the 10.7 MHz first IF in the demodulator. Make sure there is no interference problem by observing the eye pattern signal while keying the transmit side on and off via the RTS line. Under no-signal conditions you will see a reduction of the noise amplitude at this point if the interference is present. In my modem, the interference was noticeable only with the cover of the box completely off, but of course your modem may be packaged quite differently from mine (I prefer to arrange the boards in the same horizontal plane, with quite a bit of spacing between them).

There is one complication with using the DSY modem as a bit regenerator: the transmit encoder insists on clocking in the data using its internal clock; however, the timing of the data is determined by the transmit clock at the station which is being received at the repeater input. These clocks are not phaselocked, so the resulting skew can cause errors in sampling the data and failure to repeat it faithfully. I designed a simple 32-bit FIFO (First-In, First-Out) buffer circuit to correct for the clock skew, along with a watchdog timer and some logic that would allow a switch to be connected at the repeater site.

Except for some antenna problems (we found that some of those nice-looking fiberglass gain omnis from the Far East don't stand up to Canadian winters in exposed locations too well!), the repeater has been very reliable since it was installed in January 1990, on top of a building about 26 stories in height. Thus far the furthest users are about 20 km from the repeater, and we have found that 10 Watts output is perfectly adequate. Judging from the link margins, the effective coverage is probably at least two or three times this distance. The users run from 0.5 to 10 Watts, and except for those close to the repeater, most use small yagis for both receiving and transmitting. Connection of the repeater to our packet switch at the site was completed in the Fall of 1990. The switch is a 12 MHz PC AT-class machine running KA9Q NOS. Aside from the 56 kb/s MAN port interfaced via a PI board the switch has a couple of ports which provide access to/from the low-speed NET/ROM network and an ethernet board which provides an Internet mail gateway and access to future Unix-based servers. We expect to further upgrade the capabilities of the switch soon, with the addition of a PackeTen board. More information on the repeater was published in the New England TCPer [4]. Copies of the article (and information on the PI board) are available from the author. For information on the WA4DSY modem itself, contact GRAPES, P.O. Box 871, Alpharetta, GA 30239-0871, USA.

Before leaving the topic of setting up a 56 kb/s MAN, it is worth noting that anyone setting up a 56 kb/s network of any sort should be aware of the need for addling front-end filtering to the receive side of the stations. The receive converters have very broad front ends, and strong signals several MHz away can wreak havoc with your reception. So, while you're thinking about how to set up your local network, be on the lookout for surplus bandpass cavity filters and the like. The modem's 28 MHz RF input is also very **broad,** and improvement at that point will sometimes suffice instead of front-end filtering. More about that later.

The Cellnode/Feeder System: A New Role for 56 kb/s

The current evolutionary trend in the packet network is towards having a collection of major multiport nodes ("node stacks"). Typically, a node has one or more ports for users to access the network, and one or more ports for links to other nodes ("backbone" links). In cable TV parlance (not a great analogy, but the terminology is useful here), we could call the node-to-user links "drops" and the node-to-node backbone links "trunks". Nearly all of the current drops run at 1200 b/s, whereas most of the trunks operate in the 1200-9600 b/s range, although there are a few running at 38.4 kb/s and 56 kb/s (and some higher-speed links are under development, notably in California). The nodes are generally at good RF sites, and because such sites are not easy to come by, they are usually spaced fairly far apart. This arrangement works pretty well in areas with relatively low user population densities, since only a small number of users needs access to a given node. In an urban area, however, the node may need to serve a large number of users. A single user port quickly becomes saturated, and the users get frustrated with the resulting poor throughput. The collapse of the user port from congestion is hastened by the hidden transmitter problem. The obvious solution is to add more user ports, but this may not be feasible. Maintaining isolation between the different RF ports of the node, and probably with other radios at the same site, becomes increasingly difficult as new frequencies are added. The hidden transmitter problem can be eliminated by putting up a full-duplex repeater for the user port, but this also involves complications on the RF side, and this solution has not beer! a popular one.

Now let's back up one step and take a look at the problem from the point of view of user access to a single node, leaving aside the question of inter-node links for the moment. Let's assume that the users are running low speed (1200-9600 b/s), so that only a small number can share a channel if they are to get adequate throughput. Hidden transmitters should be avoided in order to contribute to this objective. On the other hand, we want to keep things simple at the node site; ideally, we would like to have just a single half-duplex port for users. How do we simuttaneously meet these requirements? The answer is quite obvious: we must restrict the coverage of the node so that it covers a relatively small geographic area. Then the number of users accessing the node will be small, and being close together, they likely will be in range of one another and will not be hidden transmitters. And, since coverage is quite limited, the frequency can be reused with smaller geographic spacing. In other words, we adopt a cellular approach in setting up our network.

Having limited-coverage cellular nodes (for convenience, I will call them cellnodes) sounds like a fine idea for giving users better access to the network, but we're not out of the woods yet. How do these cellnodes get tied into the network? With their limited coverage, they may not have good paths to their neighbors, not to mention the more distant nodes outside the urban area. Clearly there is still a need for well-sited central node which can provide the trunks to the more distant areas and also tie the cellnodes into the network. Again the similarity to the cable TV network: in addition to the trunks connecting the major nodes, and the drops from the cellnodes to individual users, we now require a feeder system which ties a group of cellnodes into a nearby major hub node. Now we come to the crux of the matter: how do we construct the feeder system? We could run separate point-to-point links, each on a different frequency, to each cellnode. This is okay from the point of view of the cellnodes, but it makes the design of the central node horrendously difficult. The point-to-point approach scales very poorly: each time we add another cellnode, we must add another port to the central node, the necessary radio and antenna hardware, and worst of all, the requisite filters and duplexers needed to eliminate interference between all the radio gear (and good sites tend to have other amateur and non-amateur radios present in addition to the packet equipment). At some point, it will become virtually impossible to add another cellnode.

Enter the CSMA MAN with full-duplex repeater. Provided that the bit rate is high enough to handle the combined traffic with minimal delays, a CSMA shared-channel arrangement will work fine for the feeder links, and it vastly simplifies the setup at the central node. Only one set of radio gear and omni antenna (two antennas if it's a crossband repeater) is needed, and instead of a stack of TNCs or equivalent, a single port on some hardware that is more appropriate for doing the job of a packet switch. Not only is it simple, but it scales well: a new cellnode can be brought up without making any change whatever at the central node (except possibly programming in some new routing information), and the overall throughput of the feeder system degrades gracefully as new cellnodes are added. Of course, at some point the delays will become noticeable and adding more cellnodes will be unacceptable. On a 56 kb/s full-duplex feeder system, it will take a *lot* of cellnodes to reach that point if each has just a handful of 1200 b/s users, but it must be assumed that at least some of the cellnodes will be providing drops at up to 9600 b/s. Here's where it gets interesting.

Naturally, my candidate for building the feeder system around is the DSY 56 kb/s modem. As you will recall, this is an RF modem which provides its input and output separately at frequencies in the 28-30 MHz band, and transmit and receive converters are used to translate this IF to/from the actual operating frequency (or frequencies, in this case). This property is a big win when it comes to expanding the capabilities of a node. Consider Figure 2, in which we have added a second full-duplex repeater at the node site. Adding another repeater sounds like a formidable proposition, but take a look. New antennas, filters duplexers, radios (converters)? There aren't any! All we need is the new modem and its associated FIFO and interface hardware, a power splitter at the 28 MHz output of the receive converter, and a power combiner at the 28 MHz input to the transmit converter. And, of course, the packet switch has to have another high-speed port available - not a problem if we've chosen hardware appropriate for doing the job of a major network node, such as the PackeTen board. That's it - we have doubled the capacity of the feeder system, with no sweat.

Okay, I've glossed over some details - life is never that simple! Let's consider some of the finer points.. . First of all, it's quite obvious that we must set the drive levels into the transmit up-converter such that each repeater only uses half of the total power output. In fact, the drive should be backed off a bit more than that in order to be certain that the converter/amplifier remain in their linear operating regions. It would be wise to use a spectrum analyzer at the RF output to make this adjustment. Chances are that link margins will be sufficient on the links using the original repeater so that the 3 dB loss in repeater power output won't cause any prolblems, but if not, this is easily remedied by adding a suitable "brick" linear power amplitier at the repeater output. Since the repeater is full-duplex, transmit/receive switching is not needed, provided that the amplifier is stable when its drive is removed.

There are also some constraints on the operating frequencies of the repeaters. The receive and transmit converters themselves don't impose much restriction, since they are inherently broadband devices. Clearly the repeater inputs must be in the same band (within 2 MHz or so), and the same goes for the outputs. The more severe constraint comes from the need to place both inputs within the passband of the input bandpass filter, preferably without modifying the filter. There may also be narrowband filtering at the combined repeater output, if the repeaters are in-band or duplexed with other radio equipmen?. This means that the repeater inputs and outputs should be within a few hundred kHz of each other (in the case of 56 kb/s, probably occupying adjacent, or perhaps next-to-adjacent, 100 kHz channels). In some areas, several contiguous 100 KHz bandwidth channels for packet may not be available on the band(s) of choice; for example, the recent band plan for the 222-225 MHz band in Southern California precludes this type of usage. Hopefully, the advantages of block conversion and adjacency of wideband packet channels will be given more consideration in future band planning.

Assuming that are adjacent 100 kHz channels available, we now must consider whether the DSY modem is up to the task of operating in such a mode. The modem is a double-conversion device: the 28 MHz front end is converted to a first IF of 10.7 MHz, and then to a second IF 455 kHz. The front end is very broad; in fact, it only has a lowpass filter to bandlimit the input. This limitation was brought home to us in Ottawa when several of us were experiencing interference problems when receiving the output 'of our repeater on 433.55 MHz. The interference was eventually traced to a surprising source which was well away from that frequency: a network of paging transmitters operating near 414.7 MHz were being down-converted in our receive converters to about 10.7 MHz, and the signals were sailing right through the modem front end and reaching the demodulator. This prompted Dennis Rosenauer (VE7BPE) to design a 28-30 MHz bandpass filter to go between the receive converter and the modem (details available from the author). The first IF of the modem has a bandwidth of about 300 kHz, and the second IF establishes the ultimate bandwidth of the modem (when used at 56 kb/s) of about 70 kHz. This means that there will be very little rejection of an adjacent-channel signal 100 kHz away until we get to the second IF, and thus there is potential for intermod problems in the earlier stages of the modem if the adjacent channel signal is much stronger than the desired signal. Furthermore, the second IF filter consists of only three relatively low-Q LC fitters, and it rolls off quite slowly, so the ultimate adjacent-channel rejection is not very impressive. Recent measurements by Ian McEachern VE3PFH [5] indicate that the adjacent-channel rejection is insufficient to avoid significant degradation of the bit error rate, even under the most optimistic scenario (we control the power of the stations using the repeaters such that signals in the adjacent channel never exceed the desired signal by more than 3 dB). We also have to be concerned about the possibility of interference from other types of emissions on nearby frequencies.

Tightening up the first IF of the modem is quite possible, but far from trivial. Simply replacing the 10.7 MHz ceramic filter, which is specified at about 280 kHz bandwidth, with one having 80-100 kHz bandwidth is problematic since the latter is not generally available. Filters for 10.7 bandwidth generally come in two types: those which are intended for FM stereo broadcast use, which requires a bandwidth of around 250 kHz, and those intended for narrowband FM, with a bandwidth usualiy less than 20 kHz. Improving the skirts of the 455 kHz second IF filter is another alternative, but again does not lend itself to off-the-shelf solutions. We are continuing to study the possibilities for upgrading of the modem's adjacent channel interference rejection capabilities; in the meantime, a stopgap solution is to use next-to-adjacent channels (i.e 200 kHz spacing).

As for the power splitter and combiner, they need not be anything elaborate. The splitter may not even be needed. I have frequently connected my HF receiver to the output of the receive converter in parallel with my DSY modem by just using a coaxial tee connector, with no apparent ill effects on the modem performance (this is a good way of checking for interference problems, by the way). It is probably more important to use a true hybrid circuit, with reasonable isolation between the inputs, for the power combiner. Such units are available from companies such as Mini-Circuits, but they are not difficult to build. Check the "Test and Measurements" chapter of a recent ARRL Radio Amateur's Handbook for construction information.

For obvious reasons, in the two-channel repeater configuration we cannot let the transmit side of both modems run continuously and just key the transmit converter, as was recommended for the single repeater. The DSY modulators must be keyed, but the present method of keying could be changed in order to improve the keyup time.

In Ottawa, our 56 kb/s full-duplex repeater was originally intended as a MAN for the TCP/IP "power users". It still fulfils that function, but it has also evolved into a feeder system for several BBS stations and 1200 b/s user access ports. We are now working on a second repeater, to be deployed as outlined here, which will be used to separate those functions to some extent. We also are investigating the possibility of running the second repeater at a higher speed, such as 112 or 128 kb/s. Despite those objectives, the mixed usage of our current MAN has worked out quite well. The power users can become part of the network feeder system by simply opening up a low-speed port and thereby creating a new cellnode for their local area. Of course, highly motivated users like this are not all that common, so it might be necessary for a local club to subsidize some of the cost of putting up a cellnode. Either way, the cellnode sites will generally be home stations (running KA9Q NOS), which can be a major advantage in terms of network maintenance and reliability. It is unlikely that the coverage provided by cellnodes in a given area will be complete, so it will probably be necessary to retain a wider-coverage low-speed port at the main node site to pick up the users who are not covered otherwise.

A Fly in the Ointment: Multipath

One disadvantage of the CSMA MAN compared with separate point-to-point links is that multipath problems are more likely with the former. This is because the antenna at the central node must of necessity have omnidirectional coverage, and therefore it does not provide any discrimination against signals arriving (or leaving) by paths other than the direct one, That's not to say that such problems will happen frequently, but the possibility should be kept in mind when setting up links,

Multipath is characterized by a delay spread parameter, which gives a measure of the time interval over which significant energy from a given signal element from a transmitting site arrives at a receiving site. As far as data transmission is concerned, the multipath has quite different manifestations, depending upon the relationship between the delay spread and the symbol length (baud interval) being used. If the delay spread is much smaller than the symbol length, then the multipath can cause more or less complete cancellation of the signal (i.e., "flat" fading, where in this case the fade is permanent, at least until something in the environment between the sites changes). This can happen, for example, if you are receiving the signal on the direct path, and in addition are receiving it with nearly equal strength after it reflects from a nearby building. If the signal level on a link proves to be much lower than expected, it is worthwhile trying to adjust the position of the antenna at the cellnode end of the link to see if that improves the situation (with the DSY modem, looking at the eye pattern is the best way to assess the link quality). If it does improve, hopefully the paths will be stable enough to retain the favorable phase relationships needed to avoid cancellation.

The other situation occurs when the multipath spread is an appreciable fraction of the symbol length, say 20% or more. Then we do not experience complete signal cancellation, but we have a new problem: intersymbol interference. A significant amount of energy from the preceding symbol is still arriving whiie we are trying to decide the value of the current symbol, making it harder to make the correct decision. The bit error rate goes up rapidly as the intersymbol interference increases, and since every bit in a packet must be correct for the packet to be successfully decoded, link performance falls off very rapidly. In the case of a 56 kb/s binary modem like the GSY, the symbol length is about 18 μs, so problems from intersymbol interference start to get likely when the delay spread exceeds 2 or 3 μs. Unfortunately, spreads of this magnitude and much more are not at all uncommon, especially in urban and mountainous areas. The author was involved in a series of multipath delay spread measurements (at 800 MHz) in four Canadian cities in 1990, and we frequently saw spreads in excess of 10 μs, and more than 20 μs in several cases. Other workers have documented multipath spreads at VHF of as much as 50 μs in some unusual situations.

The insidious thing about multipath-induced intersymbol interference is that, unlike the case of signal cancellation mentioned above, a small shift in antenna location probably won't help at all. The best way to avoid multipath is to move to the prairies, but a more practical solution for most people is to use directional antennas. Many of the long-delayed echoes arriving at a site will have quite a different azimuth from the direct path, and a gain antenna will often attenuate these enough to solve the problem. In the case of the CSMA MAN, this means that all of the stations except the central node should use directional antennas (this may also be advantageous in terms of keeping signals on the MAN repeater input frequency confined to a smaller area, permitting greater frequency reuse). Since the central node must have omnidirectional (or at least very wide) coverage, multipath will generally be more of a problem than on point-to-point links. For this reason, we probably will not see CSMA MANs running at speeds much higher than 56 kb/s in the near future. Intersymbol interference can also be dealt with by using tapped delay-line equalizers, but doing so in a CSMA MAN environment would be problematic, to say the least.

A useful tool for assessing link quality is the demodulator eye pattern as observed with an oscilloscope synchronized to the recovered clock. These signals are readily available in the DSY modem. Packets flash by pretty quickly at 56 kb/s, so some provision should be made for keying up the modem/transmitters for longer periods for link quality testing. A very noisy eye may indicate a cancellation probiem caused by multipath, or it may be some other fault which is lowering the link margin. An eye which is badly distorted rather than noisy may be a good indication of multipath intersymbol interference. If the latter is noticeable even when a yagi or other directional antenna is used at the cellnode site it is worthwhile rotating the antenna to see if the eye signal improves. Minimum distortion may not coincide with the antenna being aimed along the direct path.

The concept of block conversion (sharing of RF components) can easily be extended to the network backbone links, or trunks. Although the major trunks in the network will eventually migrate up to microwave links running at T1 speeds (1.544 Mb/s) and beyond, 56 kb/s will satisfy much of our needs in the near term. If we're going to do the job right, though, we have to make the trunks full duplex. The reasons are several: obviously, higher efficiency and better throughput due to the simuttaneous two-way data transfer (provided that the hardware and protocols support it) and elimination of turnaround time overhead. In the case of the DSY modem, the noncoherent demodulator could be replaced with a coherent design, thus improving link margins. The really big win with full duplex, however, is the ability to use the block conversion and multicoupling scheme outlined above for multichannel MAN installations.

Consider a node with three 56 kb/s trunk links. The current tendency would be to put up three half-duplex links, each on a different band, with a complete set of radio gear and directional antenna for each one. This may not be feasible at some sites. On the other hand, if we use a variation on the block conversion scheme outlined previously, we can run full duplex on all the links *and* use only two bands, with just one shared receiver on one band, and a shared transmitter on the other. In fact, all three full-duplex links could be done on a single band, using a duplexer and one antenna; however, current band plans are not conducive to this type of operation with high-speed packet. The antenna(s) can either be omnidirectional with gain in the vertical plane, or an array of yagis or other directional antennas with a suitable matching harness. As in the multiple repeater idea discussed above, the poor adjacent channel rejection capabilities of the modem are a limiting factor in this full-duplex multiport scheme. Greater than 100 kHz spacing between channels is needed unless the modem IF response is improved.

A further simplification is possible: all of the links could share a single modem and frequency slot on the transmit side. This configuration was proposed by Phil Karn several years ago [6]. Unfortunately, the channel access protocols currently in use would not make efficient use of this arrangement. The ports of the node would have to share the transmit channel on a first come first serve basis, and the links would have to be operated in half-duplex mode to permit this (though it seems to me that some minor software tweaks could produce a hybrid "full duplex when the transmit channel is available" mode). Such a configuration is worth considering in cases where 100 kHz bandwidth channels are in short supply.

For further information on engineering full-duplex inter-node links at 56 kb/s, I recommend reading Don Lemke's article in the 1988 ARRL CNC proceedings [7]. His article also has a good discussion on the calculation of link margins.

References

[1] Fox, T., "RF, RF, where is my high speed RF?", *Fifth ARRL Computer Networking Conference*, pp. 5.1-5.5.

[2] Heatherington, D., "A 56 kilobaud RF modem", *Sixth ARRL Computer Networking Conference*, Redondo Beach, CA, August 29, 1987, pp. 68-75.

[3] D. Lemley and M. Heath, "The Packe-Ten system - the next generation packet switch", *Ninth Computer Networking Conference*, London, Ontario, September 22, 1990, pp. 170-l 76.

[4] McLarnon, B., "A packet repeater using the WA4DSY 56 kbs modem", *The New England TCPer*, Vol. 2, No. 5, May/June 1990.

[5] McEachern, I., "Digital networking with the WA4DSY modem - adjacent channel and co-channel frequency reuse considerations", this conference.

[6] Karn, P., "A, high performance, collision-free packet radio network", *Sixth ARRL Computer Networking Conference,* Redondo Beach, CA, August 29, 1987, pp. 86-89.

[7] Lemke, D., "Cellular area coverage transport networks", *Seventh ARRL Computer Networking Conference,* Columbia, MD, October 1, 1988, pp. 122-134.

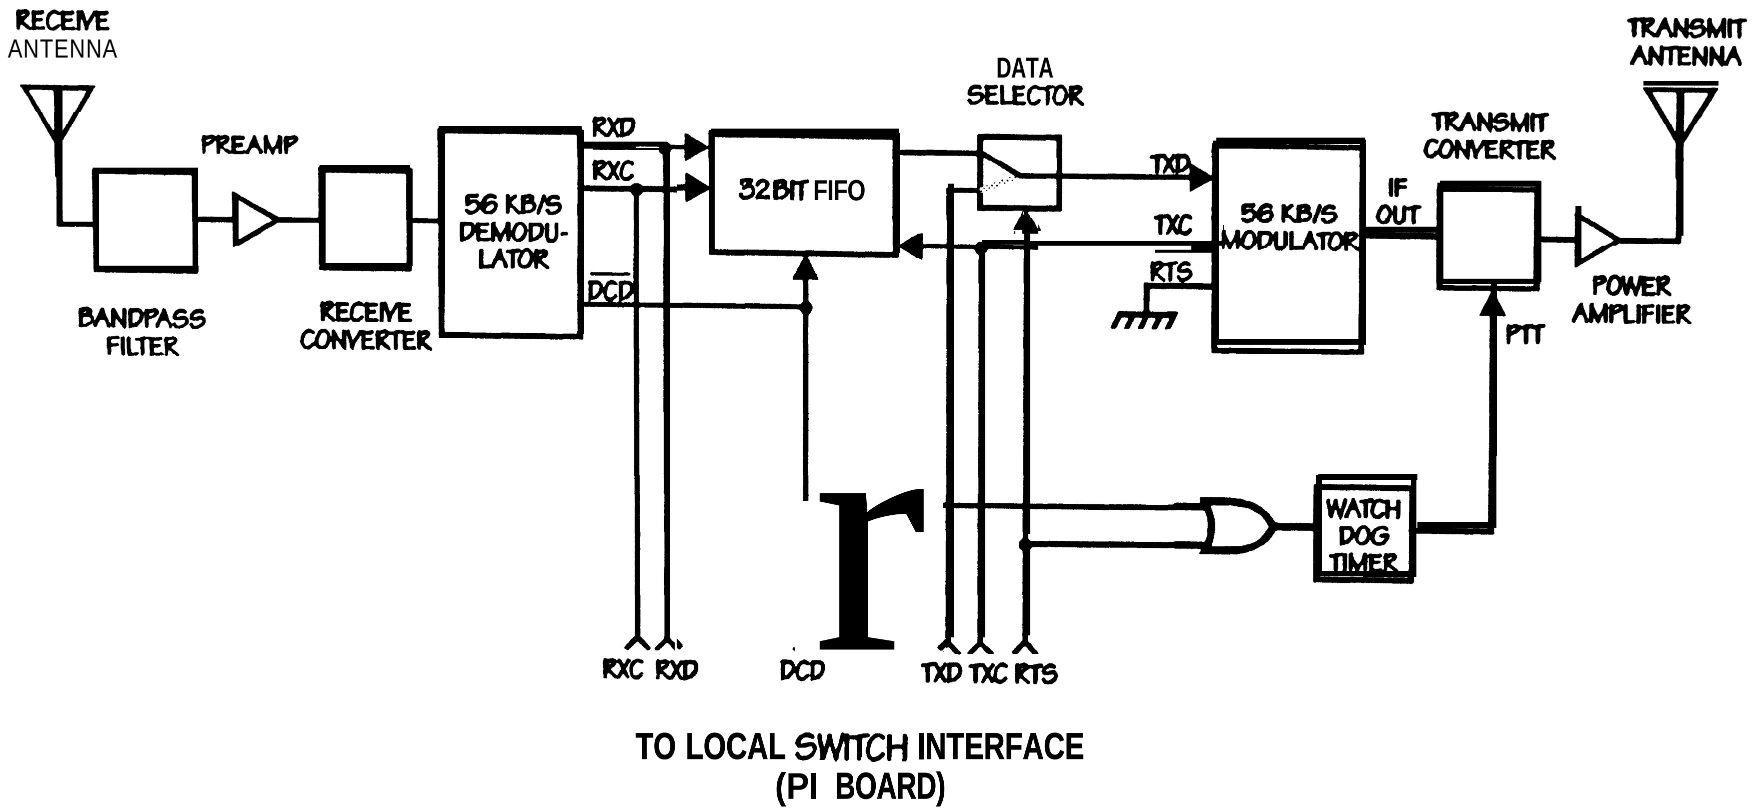

F-gum 1 56 KB/S REPEATER BLOCK DIAGRAM

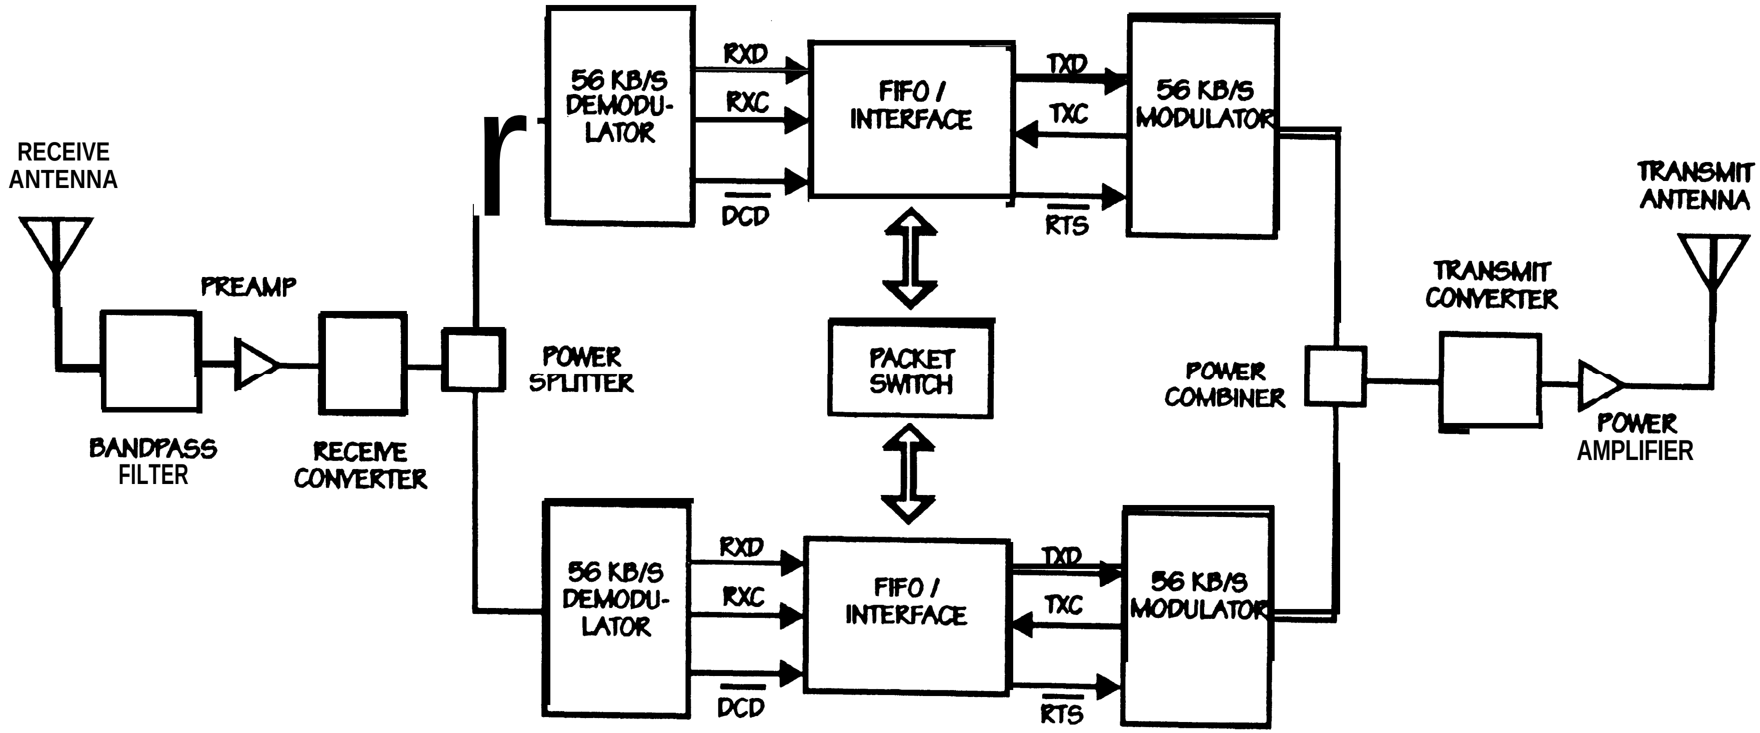

Figure 2 TWO-CHANNEL REPEATER CONFIGURATION