

LOCAL DISTRIBUTION IN THE AMATEUR RADIO ENVIRONMENT

F. Davoli, A. Giordano (I1TD), S. Zappatore (IW1PTR)

Department of Communications, Computer and Systems Science (DIST),
University of Genoa (IK1HLF)
Via Opera Pia, 11A - 16145 Genova, Italy

ABSTRACT

Packet radio is often used as a means of accessing computing facilities or, more generally, a larger terrestrial network. In this respect, it can be seen as a form of local distribution. Similarly, intelligent repeaters interconnecting stations that are not in one another's hearing range can be thought of as performing an analogous function. Whenever this is the case, alternatives to the popular CSMA access protocol may be considered, that are better suited to work in a semi-centralized environment. In this **paper**, we describe some versions of one of these protocols (R-ISA), and its integration with higher layer ones.

INTRODUCTION

Several situations arise where locally sparse users want to access a central **facility**. This may consist of a computing center, a local area network or, in general, a base station connected to a long haul transmission network. A typical instance of the latter is in the mobile circuit-switched communication network, where a certain number of frequencies are allocated to users within an area on a per call basis. In the packet switched environment, this problem of local distribution can be effectively solved by means of several available multiaccess techniques. Following the classical survey of Tobagi [1] (see also [2] for a more recent one), these can be grouped roughly into four categories: a) fixed assignment (FDMA, TDMA, CDMA, . . .). b) random access (pure and slotted ALOHA, CSMA in its various forms, CSMA/CD, . . .). c) demand assignment with centralized (polling) or distributed control (e.g., token passing, either explicit or implicit); d) adaptive and mixed strategies (e.g., adaptive polling or probing, Urn, GRA, etc.). Fully distributed access methods like CSMA are currently in widespread use in packet radio, where they face in the same **way** the above mentioned local distribution problem and the more general one of peer-to-peer communication among **any two** stations. However, whenever a base station is present for some reasons, and the local distribution traffic plays a relevant role, the semi-centralized nature of the system can be exploited to enhance the characteristics of the access protocol. This can be done essentially in two ways: i) by using the base station as a means of synchronization (which may be simpler and more easily implementable than for instance, slotting the channel); ii) by distributing "state" information (e.g., channel feedback), that would otherwise be difficult to achieve on a radio channel, or would be achievable with longer delay. The latter feature is particularly useful to make the access strategy adaptive to load variations.

The authors have proposed an adaptive protocol of this kind [3] (R-ISA) that retains the basic properties of the ISA adaptive strategy for a slotted channel [4]. With respect to

the above mentioned classification, R-ISA falls mainly in category d), and exhibits some similarities with adaptive polling [5].

In this paper, after briefly recalling the basic mechanism of the access strategy, we consider some more practical issues related with its implementation. In particular, we are interested in interfacing R-ISA with upper layer protocols as AX.25, TCP/IP and NET/ROM, as we are currently realizing an experimental network over the 430 MHz amateur band, using unmodified TNC hardware. We will therefore describe a possible MAC layer protocol and service for R-ISA, and discuss some choices related with the integration of new stations within a “local” net served by a base station and with the distribution of access rights. The description is informal, and does not attempt to rigorously prove the validity of the protocol. A performance analysis of the basic scheme can be found in [3].

THE R-ISA ACCESS SCHEME

R-ISA has evolved from the original ISA algorithm [6], in order to exploit the presence of a base station, which will be denoted by \mathcal{S} . Let \mathcal{T} represent the set of N peripheral stations. The scheme operates as follows. \mathcal{S} and \mathcal{T} alternate their transmission periods. A transmission period is recognized between the Beginning-of-Carrier (BOC) and End-of-Carrier (EOC) events of either \mathcal{S} or \mathcal{T} . Transmissions from \mathcal{S} are conflict-free, whereas at the beginning of each transmission period from \mathcal{T} , access rights are assigned and the actions following will give rise to a channel outcome (collision/no collision) that can be observed by \mathcal{S} , and can be either broadcast to the stations in \mathcal{T} (that will use it to compute the access rights) or directly used by \mathcal{S} to compute the next access rights, that will be then distributed to all stations in \mathcal{T} . The computation of the access rights is based, besides this channel feedback information, also on estimates of the “a priori” information on the packet generation rates (in packets/s) h_i , $\forall i \in \mathcal{T}$, of the peripheral stations (each λ_i can be locally estimated and, at the earliest convenience, transmitted to \mathcal{S} , which will either diffuse it by broadcasting or use it directly in the computation).

On the basis of all the above informations, a data base is kept and dynamically updated, consisting in a vector $p(t)$, whose components p_i^t represent the “marginal” presence probability of a packet at station i within cycle t (a cycle is the time interval made up by the two consecutive transmission periods of \mathcal{S} and \mathcal{T}). At each t , given the current value of $p(t)$, the vector $\hat{p}(t)$ is constructed of the marginal probabilities arranged in non-increasing order (i.e., $\hat{p}_{j+1}^t \leq \hat{p}_j^t$, $j=1, \dots, N-1$). This is used to compute the number $k(t)$ of stations to be given access right, starting from the one with highest presence probability. The number $k(t)$ is obtained by a simple algorithm, that repeatedly evaluates the difference between the probability of success and that of no transmission in the transmission period of \mathcal{T} within the current cycle, given a set of enabled stations that initially is made up by the station with highest marginal probability (\hat{p}_1^t), and is increased by 1 at each evaluation, always including stations in order of non-decreasing presence probability. Eventually, $k(t)$ is found as the number such that the above difference is negative or null with $k(t)-1$ enabled stations, whereas it becomes positive (and would remain such) with $k(t)$ enabled stations. It is easily seen, for instance, that a single station will be enabled whenever $\hat{p}_1^t > 1/2$, and that all the N stations would be

enabled if the above difference remained negative. As is consistent with intuition, in general, enabling a set of stations with a “large” presence probability within the set increases the possibility of a collision, whereas a set of stations with “low” presence probability within the set may more easily result in an empty transmission period; in both cases, bandwidth would be wasted, and there must be clearly a tradeoff. It has been shown [4, 6] that the algorithm maximizes the success probability in the transmission period of \mathcal{T} within the current cycle, under the hypothesis of independence of the presence of packets at the stations.

The updating of $p(t)$ is based on the previous values, the decision taken on the access rights, the resulting channel feedback, and the “a priori” information on the packet generation rates. It is easily performed if all stations in \mathcal{T} are supposed to have buffers with capacity of a single packet. The calculation can be split into two parts. First, one can consider the situation corresponding to no external inputs, and find the presence probabilities \tilde{p}_i^{t+1} in this case. It is easy to see that if station i had no access right, then

$\tilde{p}_i^{t+1} = p_i^t$; and that if station i had access right and no collision resulted, then $\tilde{p}_i^{t+1} = 0$. The

only case involving some calculations [3, 4] is that of station i having access right and getting involved in a collision. As regards the second part of the updating, the peripheral stations have a buffer capable of containing a single packet, that is kept until successful transmission, and supposing Poisson arrivals with mean $\lambda_i, i \in \mathcal{T}$ we can evaluate the probability of a packet generation $\sigma_i(t)$ at station i in cycle t as

$$\sigma_i(t) = 1 - e^{-\lambda_i T_t}$$

being T_t the length of cycle t . The final updating is then

$$p_i^{t+1} = 1 - (1 - \tilde{p}_i^{t+1}) [1 - \sigma_i(t)]$$

TWO DIFFERENT APPLICATION INSTANCES AND THEIR IMPLICATIONS

We briefly consider two different cases, that may both occur in the amateur radio environment, and that require different handling of radio stations becoming active or inactive (i.e., joining or leaving the system). First of all, however, we note three characteristics of the above described scheme: i) due to the presence of the base station, the performance will be insensitive to stations in \mathcal{T} being hidden from one another, the only requirement being that they all are in the hearing range of the base station; ii) the total number N of stations in \mathcal{T} must be known; iii) these stations must be assigned a number or “local identification”.

The first of these is a positive feature with respect, for instance, to CSMA, which would be very sensitive to hidden stations (unless using variations thereof, like BTMA [1]). We note, in passing, that several simulation results have shown a superiority of R-ISA with respect to CSMA, in terms of throughput and delay [7]. The other two characteristics require some care in the implementation, and may lead to alternative solutions to suit some specific situations.

A possible topology where the above mentioned advantages and different features would appear is that of a “two level” network, as depicted in Fig. 1, where terminal users are connected with “low and local” repeaters, and the latter in turn may access “main” repeater stations located in dominant position on top of a mountain to ensure a large coverage area. The main repeaters may be viewed as forming a backbone network,

perhaps using some higher speed link. Typically, the “lower” repeaters and user stations situated on the opposite sides of a mountain would be hidden, but within the range of the top station. Thus, this situation would be suitable for the application of R-ISA. Since the “intermediate” repeaters are fixed and their number does not change frequently, one can set the number N as a foreseen maximum of active stations, identify all stations a priori, and use the same mechanism as in [3] to allow the possibility that a station switches on and off. This consists in assigning a constant default “low” value σ_d to temporarily inactive stations, in order to “enable” them less frequently, but with **nonzero** probability, so that they can reconnect when they switch on. Also, since the intermediate repeaters under the range of a main one would typically be not too large in number, it is reasonable in such situation to centralize computation in the main station, and distribute the access rights at each cycle. Note that, if the user stations within each area would continue to use CSMA among- -themselves and with the intermediate repeaters, the introduction of the above scheme would require no change in TNC hardware and software, as it affects only the dialogue with the “backbone” repeaters.

As regards the extension of the procedure to the user stations, for which the intermediate repeaters may act as base stations, a little more care must be taken. Actually, two features would appear in this case: i) user stations are not necessarily tied to an area, and can move around, even if they were not mobile in a strict sense; ii) the number of stations in an R-ISA population is likely to be larger than that considered before. Due to point i), it is no longer possible to permanently assign a local identifier to a station within an area; a station entering an area must be given the possibility to execute the access protocol even in the absence of this number, and just in order to get one. Point ii) requires the usage of more efficient mechanisms to distribute the access rights than, say, a bit map (which would be reasonable with a limited number of users), or requires the distributed computation of the access rights (however, in this case, one should think of recovery mechanisms in case that the channel feedback is not seen by all stations in the same way). We briefly discuss here only the first issue.

A possible solution is to reserve a number, say 0, not for a specific station but as a common identification for newcomers. We shall refer to this as an “access virtual channel” (AVC). Whenever the AVC is enabled and there are stations waiting to be identified, they send a “MAC connection request” to the base station. If only one is waiting (and provided it does not collide with a normal packet from a contemporarily enabled station), the base station will recognize the message and respond by releasing a local identifier, chosen among the free ones. If two or more are waiting (or the single one collides with a normal packet), a randomization mechanism may be used by the waiting stations at the next enabling of the AVC. This behaviour has some analogy with that described in [8]. Note that, in case only the AVC is enabled and a collision is seen (because there are multiple access requests), the R-ISA updating rule would yield an indeterminate form (O/O) in this case for the new presence probability; obviously, the latter must be set to 1 for the AVC, that will be thus enabled singularly in the next cycle (by the way, a similar behaviour must be taken even when a normal transmission by a singularly enabled station is corrupted by noise and seen as a collision). We may also remark that the presence of the AVC eliminates the need for the use of σ_d ; the rates corresponding to unassigned identifiers can be set to zero, speeding up the execution.

MAC LAYER FRAMES FOR R-ISA WITH AVC

In this section, we give a description of a preliminary version of the R-ISA MAC layer frame formats, that are necessary to implement the basic functions required by the protocol. We refer to the AVC case; moreover, we suppose that the computation of the access rights is centralized, and that their distribution is effected by means of a simple bit map.

Obviously, due to the structure of the R-ISA strategy, a first distinction must be made between frames pertaining to the base station and those of the peripheral ones.

In the \mathcal{S} to \mathcal{T} direction, the following MAC frame types are needed: i) acceptance of a MAC connection request (containing the ID of \mathcal{S} , the ID of the station in \mathcal{T} , and the assigned local identification number); ii) acceptance of a MAC disconnection request; iii) information (generally containing an AX.25 frame in the information field); iv) supervisory frames. All these frames contain a field with the access rights bit map.

In the \mathcal{T} to \mathcal{S} direction, one has again four MAC frame types: i) MAC connection request (containing the ID of the issuing station, the base station ID, and the AVC number); ii) MAC disconnection request; iii) information; iv) supervisory.

We will designate the first four types as **S-frames**, and the second four as T-frames, and describe some of them in greater detail. S-frames formats are shown in Fig.- 2. All of them begin with an 8-bit type field, followed by a bit-map-length field and a bit map field. The bit-map-length **concides** with the highest "active" station number (i.e., the highest number among those assigned to stations in \mathcal{T} that are currently connected) plus one (number 0 is the AVC, that is supposed unique). Each set bit in the bit map field means access right to the corresponding station. Bit 7 in the type field distinguishes between **S**-information and other S-frame types. The other significant bits are as in the figure.

T-frames formats are shown in Fig. 3. All of them begin with an **8-bit** type field, followed by the local ID or "channel" field, that contains the assigned local ID number (or 0, when the station is still "on the AVC"). Again, bit 7 in the type field distinguishes between **T-information** and other **T-frame** types. T-information frames contain the current value of hi in the remaining 7 bits of the type field. If bit 7 is set, bit 6 identifies either a request or supervisory T-frame. In case of a request-type, bit 5 is used to discriminate between connection and disconnection. All MAC frames should be embedded within the boundaries of an HDLC frame (flags and CRC).

An incoming station sends a **connection** request on the AVC (i.e., it listens to the channel, monitoring the bit map field, until it sees enabling of the AVC, then sends the frame). If this is successfully received, the base station \mathcal{S} will respond by accepting the connection and releasing the local ID. If the request is unsuccessful, for any reason, the incoming station will use the randomization mechanism at the following transmission attempts on the AVC. During normal data transfer, no ack is needed at the MAC layer. Disconnection is effected by sending a disconnection request, that must be ack-ed by \mathcal{S} .

In case no station in \mathcal{T} is connected, \mathcal{S} will enter a "dormant" state, and will only periodically send a supervisory frame for synchronization, containing the enabling of the AVC.

Other supervisory frame types and possible timeout conditions are under consideration.

CONCLUSIONS

We have considered some possible application instances of the R-ISA adaptive access strategy in an amateur environment. Due to its characteristic features, R-ISA seems to be well suited as a local distribution strategy (even if it does not prevent direct communications among peripheral stations). This is the case, for instance, of sparse **PC's** that want to access a LAN or another type of fixed network, or an intelligent repeater station situated in a dominant position. In order to cope with stations moving in and out of the boundaries of the local area defined by a base station and its tributaries executing the R-ISA protocol, we have introduced a "common" identification, that has been termed "access virtual channel" (AVC). Then, some aspects of the protocol related to this mechanism and possible formats of the necessary MAC frames have been briefly

discussed. In fact, the necessity of specific MAC **layer** frames is one of the characteristics that differentiate this protocol with respect to simpler ones, e.g., CSMA. On the other hand, the increased complexity may turn out to be justified, at least in some cases, by the achievement of better performances and the insensitivity to hidden (stations. We hope to further verify this conviction through the experimental setup that we are preparing and the help of the local HAM community.

REFERENCES

- [1] F.A. Tobagi, "**Multiaccess** protocols in packet communication systems", IEEE Trans. Commun., vol. COM-28, pp. 468 - 488, April 1980.
- [2] S.R. Sachs, "Alternative local area network access protocols", IEEE Commun. Mag., vol. 26, pp. 25 - 45, March 1988.
- [3] M. Aicardi, F. Davoli, A. Giordano, "Radio ISA Net: a single-hop centralized packet radio network for PC-to-mainframe interconnection", IEEE J. Select. Areas Commun., vol. 7, pp. 219 - 225, Feb. 1989.
- [4] M. Aicardi, G. Casalino, F. Davoli, "Independent Stations Algorithm for the maximization of one-step throughput in a multiaccess channel", IEEE Trans. Commun., vol. COM-35, pp. 795 - 800, Aug. 1987.
- [5] J.F. Hayes, "An adaptive technique for local distribution", IEEE Trans. Commun., vol. COM-26, pp. 1178 - 1186, Aug. 1978.
- [6] M.G. Hluchyj, "Multiple access communication: the finite user population problem" Ph. D. dissertation, Dep. Elec. Eng. Comput. Sci., Mass. Inst. Technol., Cambridge, MA Oct. 1981.
- [7] M. Aicardi, F. Davoli, A. Giordano, S. Zappatore, "A multiple access adaptive protocol for mobile-to-base-station packet communications in a structured environment", Proc. 2nd PROMETHEUS Workshop, Munich, Oct. 1989 (to appear).
- [8] J. W. Ward, **G0/K8KA**, "A brief note proposing non-ALOHA **access** techniques for PACSATs", Proc. 7th ARRL Comput. Networking Conf., Columbia, MD, Oct. 1988, pp. 182 - 185.

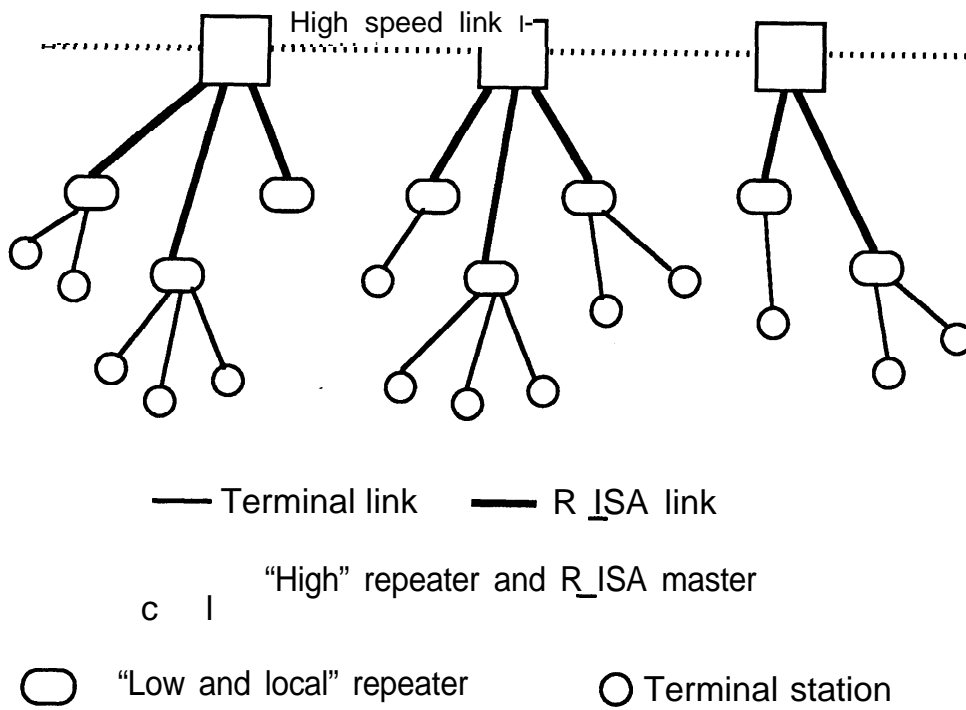


Fig. 1

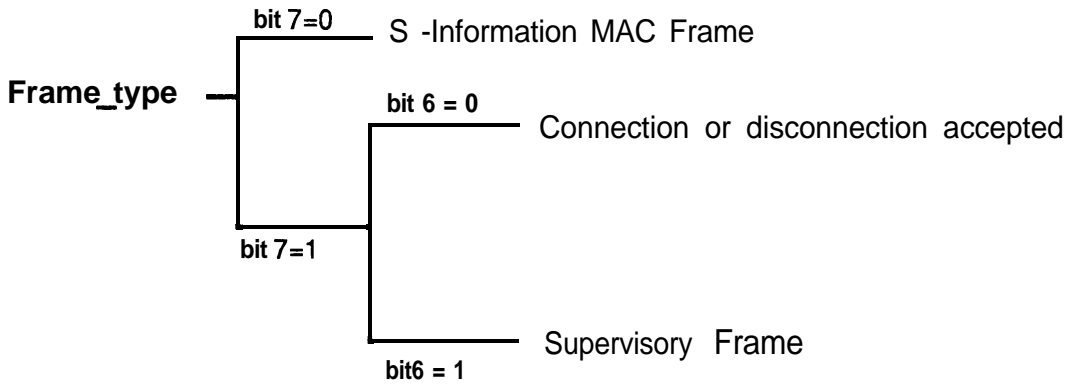



Fig. 2. S-frames: A) S-information MAC frame; B) S-supervisory frame; C) connection /disconnection accepted frame

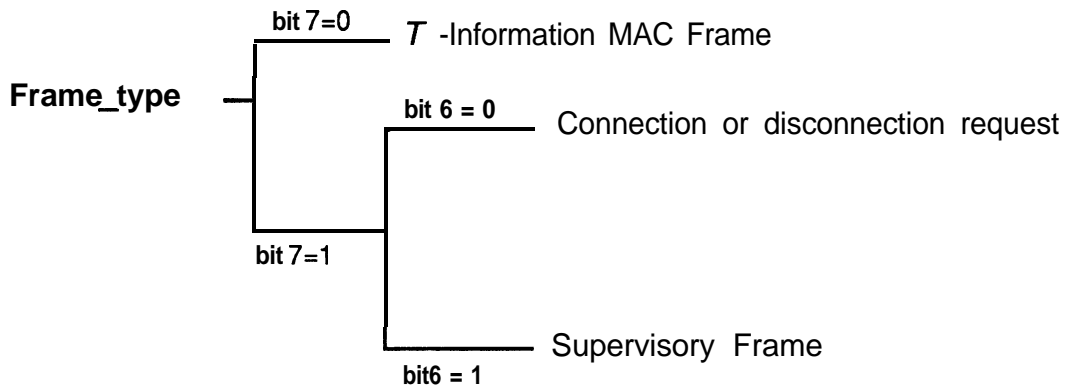



Fig. 3. T-frames: A) T-information MAC frame; B) T-supervisory frame; C) connection/disconnection request frame