

The Optimal MAC Layer for Low Power UWB Networks is With Independent Channels and No Power Control

Jean-Yves Le Boudec

Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland

Abstract— We are interested in the design of a MAC layer for systems that use ultra-wide band (UWB) communication and transmit very little power. We first explore the design space and find, that, for such networks, the optimal MAC should allow concurrent transmissions on as many channels as there are destinations; this is in sharp contrast to the established way of designing MAC protocols. Further, power control is not optimal: sources should always send at full power when they have something to transmit. Sources should constantly adapt their rate, either at the physical layer (which is difficult for the UWB channel) or at the MAC layer, using an increase/decrease protocol and incremental redundancy. Second, we explore the implications of a MAC layer with concurrent transmissions. We find that, in practice, because of nodes that can do only one thing at a time, there is a need for a “private MAC” protocol. Also, broadcasting to destinations that can listen to many channels is highly non deterministic. This makes network coding an attractive alternative to flooding for protocols, such as the address resolution protocol (ARP) and routing, that need broadcasting.

I. ULTRA-WIDE BAND COMMUNICATION IN THEORY

We are interested in the design of the MAC layer for technologies that radiate low power, motivated by the deployment of pervasive computing. Ultra-Wide Band (UWB) communication consists in sending very large bandwidth signals at very low power, such that the power density in any frequency band is below the noise level [10]. The very large bandwidth allows to operate with very low power at possibly high rates, since the capacity, in theory, scales with the bandwidth. Further, in the idealized model of infinite bandwidth, there are enough degrees of freedom to make all communications orthogonal, i.e. free of interference. In such a setting, the capacity *per node* of an ad-hoc network increases with the density (when we scale the density over a fixed area)[11]. Compare to the corresponding result for narrowband systems, where the capacity per node vanishes in the same scaling. These arguments make UWB communication attractive for very low power, dense networks.

II. THE OPTIMAL DESIGN OF MAC FOR LOW POWER UWB IN PRACTICE

In practice, however, the bandwidth of an UWB system is not infinite, but is likely to be of the order of 5 to 10 GHz. Also, channels can never be perfectly orthogonal, and sources might interfere. For example, with impulse radio systems,

sources that send concurrently may have colliding pulses. The probability, and the impact of collision of two systems is very low, but in dense networks with many nodes that would send concurrently, it is not a priori clear that the impact is negligible. Thus it is legitimate to wonder whether it is good to let concurrent transmission occur in the UWB regime, and whether power control should not be exercised, as it is done in CDMA to avoid near far interference. This motivates us to study a more realistic model of UWB communication, in order to capture what the optimal design of the MAC protocol should look like.

A. Modelling An Ad-Hoc Network

To this end we model in [5] an ad-hoc UWB network with a number of randomly placed nodes. Each node can control its power, its coding rate, the time slots during which it transmits, and over which multi-hop path the data is relayed. We assume that nodes cannot perform cooperative decoding, as this requires synchronization assumptions at the bit level that do not appear to be realistic in an ad-hoc network. We assume that nodes have peak and average power constraints. The rates achievable in such a network can, in theory, be derived from the rate function, which maps the signal to noise and interference ratio to the available data rate. For the time hopping, impulse radio method of [10], this function is linear.

In this model, a MAC layer protocol is represented by a strategy for allocating power, rates, and transmission slots. We assume all nodes know what the optimal is; such a model would be exact if there would be a central oracle that would allocate all values to all nodes, as is done in some cellular systems. For our ad-hoc network of interest, our model amounts to ignoring the protocol overheads due to the MAC layer and utilization losses due to random access.

B. The Traditional MAC Design

Traditional MAC layers in such settings are most often inspired by the WiFi (IEEE 802.11) standards, based on CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance). At this level of modelling, we can abstract CSMA/CA as a *mutual exclusion* protocol: nodes A and B cannot transmit concurrently if either A or B can sense the other node’s transmission. In other words, mutual exclusion is enforced for all transmissions that result in interference power

above the carrier sensing threshold. In contrast, if A is far from B , B does not sense the signal from A and is not inhibited. It is not clear at all whether such a design is optimal. Indeed, there are many degrees of freedom that are not exploited in WiFi: the carrier sense threshold could be modified; a source could react to concurrent transmissions by reducing its coding rate instead of enforcing mutual exclusion; a source could reduce its power in order to avoid being sensed, etc. All of these possible alternatives are implicitly contained in the model in [5].

C. Optimality Criterion

In order to use the model and find the optimal MAC design, we need to define an optimality criterion. In the networking literature, a criterion often used is max-min fairness: a strategy is max-min fair if it is not possible to improve the rate of a source without decreasing the rate of some other source that does not have a higher rate. Unfortunately, this criterion is not operational for mobile ad-hoc networks, as it turns out that, under large sets of assumptions, it always results in all sources being allocated strictly the same rate [6]. This is because, in such a wireless network, it is always possible to reduce the rate of some high rate link, in order to gain some rate increase on links that are in poor condition (because they are long or are in a bad fading state). However, the resulting allocations are, in general, vastly inefficient, as all links tend to be equalized to very small achievable rates.

An alternative criterion is to maximize total capacity (sum of rates), or transport capacity (sum of rates \times distance). Such a criterion is much more efficient than max-min fairness, but it suffers from the opposite drawback, namely, unfairness. In random networks with different link sizes and channel fadings, this form of optimality may lead to entirely shutting down the links that are in poor condition [6], a result that is generally considered to be unacceptable. A compromise between max-min fairness and maximization of rate is to maximize a concave utility function, for example the log of rates (in which case the resulting allocation is called “proportionally fair”). It appears to reconcile fairness and efficiency.

D. The Optimal Power Control is $0/P^{MAX}$

The proportionally fair allocation can be found numerically, as the resulting optimization problem is a convex one. However, even for small network sizes, its numerical solution is extremely costly, in part because the convex set of constraints is only implicitly defined.

The problem can be simplified a bit by a theoretical finding about power control. In the case of UWB (linear rate function), it is shown in [7] that the optimal solution is achievable by a $0/P^{MAX}$ strategy, and in the case of only peak power constraints (no average power constraint), any other strategy cannot be optimal.

This finding suggests that there is no point adapting power in order to reduce interference power, even in the case of near far scenarios. The only necessary power control, if any, is $0/P^{MAX}$, i.e., a scheduling algorithm.

E. No Exclusion Region for Low Power

The following further results are found empirically by a numerical solution of the optimization problem. In the optimal solution, there is an exclusion region around destinations, i.e. an area around every destination such that the source and nodes in this area are mutually excluded. This is compatible with the WiFi strategy discussed earlier. However, the size of the exclusion region depends on the powers and on the characteristics of the rate function, and increasing the carrier sense threshold in WiFi might lead to significant increases in throughput.

For low power, the exclusion region vanishes. For low power UWB networks ($1 \mu\text{W}$ transmit power), it vanishes well before interference is below noise. In other words, it is advantageous to allow low power transmissions to interfere, even if their interference cause a rate reduction: the rate reduction is less than would be lost by mutual exclusion.

Thus, for low power UWB, the optimal MAC should allow all sources to transmit concurrently at full power, while adapting their rates to interference. The first practical protocol that implements such a design was described in [2] and later refined in [4]; similar ideas underly the protocol in [1].

III. RATE ADAPTIVE MAC

The protocol in [2] adapts the code (hence the link data rate) to the level of interference. In practice, only near interferers matter, but in a dense network they may be many. The traditional method for adapting rate to channel condition is situated in the physical layer; it estimates the channel by pilot signals at the beginning of a packet transmission. Here, the channel varies due to interference, whose time scale is of the same order as packet transmission. Thus, the traditional method does not appear to be applicable. The protocol in [2] solves the problem by using incremental redundancy codes. Before being transmitted in the form of pulses, the data is encoded using a family of punctured codes. If the destination cannot decode (for example because of unexpected interference), the source sends incremental redundancy, until the destination can decode (Figure V). As a side effect, this method supports other channel variations well, in particular it performs well in mobile scenarios.

IV. ASYNCHRONOUS NETWORKING

The protocols in [2], [4], [1] are non coordinated: sources send whenever they have some packet ready to transmit and decide to do so. Synchronization and signal acquisition between source and destination is done per packet, and involves only the source and destinations of this packet. Thus, one can say that the network is *asynchronous*, like WiFi. Asynchronism at the network level results in considerable overall simplification, because networks that require global bit level synchronization are exposed to global failures. MAC protocols for UWB proposed for IEEE 802.15.3 assume some form of global synchronization. Historically, a network technology like Ethernet supplanted the Token Ring probably mainly because the former is asynchronous and the latter is not.

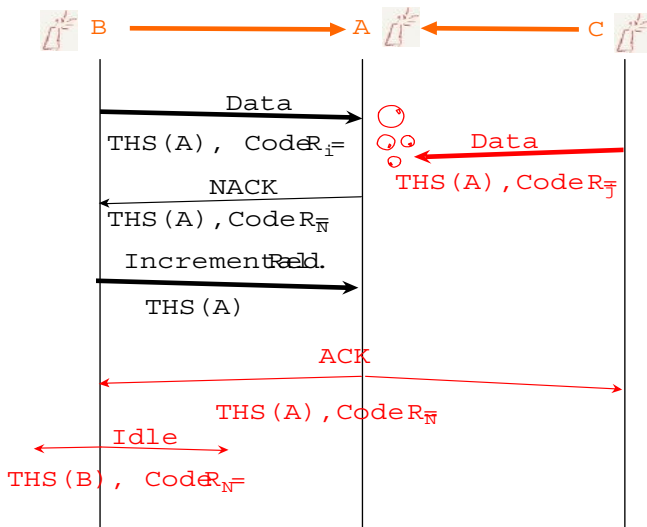


Fig. 1. Private MAC takes care of competition for the same destination. THS(A) means that the message is sent using the time hopping sequence obtained when the seed of the random number generator is the MAC address of A.

The protocols in [2], [4], [1] assume that spreading occurs in the time domain by means of time hopping codes. With time hopping, a source sends a pulse at a position within a frame determined by the time hopping sequence. Time hopping is similar to code division multiple access (CDMA), but the requirements on the time hopping sequence are not as hard as on CDMA codes; in fact, it is sufficient if the time hopping sequence is generated as pseudo noise. In [3], it is suggested to use a random number generator, with seed equal to the MAC address of the destination. This makes code distribution unnecessary, as systems are assumed to know the MAC addresses of their one hop neighbours. Thus there is one time hopping sequence per destination, and the sequence is reset for every packet. Broadcast is supported by using a common, predefined seed.

The performance of the synchronization phase is critical; it is important to avoid false detection as much as possible, while keeping the length of synchronization headers reasonably short. Methods for doing so are proposed in [8]

V. PRIVATE MAC

In practice, low end nodes deployed in a dense, low power network, are likely to be able to do only one thing at a time: either send or receive, and to/from only one other node. Therefore, it is not sufficient to have every source use a separate channel (i.e. a time hopping sequence of its own); there remain potential conflicts between sources that compete for the same destination. Also, a destination may be temporarily unreachable because it is busy sending. There is thus the need for a “private MAC” protocol, namely one that concerns only the nodes that talk to the same destination. Such a private MAC is described in [4]. It consists in having sources simply send spontaneously, and, if this fails, use a backoff

mechanism similar to the IEEE 802.11 MAC. However, there is no request message, since potential collisions usually do not result in packet loss. Deadlock conditions are avoided by requesting that packet acknowledgements are sent at the lowest possible coding rate, so that all contending sources can decode it. When a node was unreachable because it was sending and this condition ends, it sends an idle message, which, like the acknowledgement, is sent at the lowest possible rate. The choice of which time hopping sequence to use needs to be made carefully. With the combination of choices shown on Figure V, all nodes deterministically know which time hopping sequence to use. With this combination, the behaviour in case of contention is the same as with CSMA, even though carrier sense is never done (as there is no carrier to sense with UWB communication).

A similar protocol is described in [1]. It differs in that it has one additional round of overhead (request messages) and solves collision in an Aloha fashion, which is less efficient than CSMA.

VI. BROADCAST AND NETWORK CODING

Paradoxically, broadcast communication becomes more complicated with a MAC designed according to the principles in this paper. Broadcasts are used by the address resolution protocol (ARP), the hello protocol used in routing, and by some applications. When a node broadcasts a message, it may be that a number of the intended destinations are busy doing something else (sending or receiving). Unlike in WiFi, there is no way for the source of the broadcast to know if the intended destinations are free. In fact, CSMA/CA enforces that, when a broadcast packet is sent, all intended destinations cannot do anything else. Here, in contrast, in a busy network, we cannot count on this to happen. Therefore, other means are required for efficient broadcasting. One of them is flooding: broadcast data is sent and repeated as many times as necessary. This is not efficient, since the same packet is transmitted repeatedly, and may be of no value to many destinations. An alternative when there are many competing sources of broadcast it to use network coding. In [9], instead of repeating a broadcast packet, a node combines it with packets it has received. As each receiver has a different history of received packets, this should generally be more efficient.

REFERENCES

- [1] M.-G. Di Benedetto, L. De Nardis, M. Junk, and G. Giancola. (uwb)2:uncoordinated, wireless, baseborn, medium access control for uwb communication networks. *Mobile Networks and Applications special issue on WLAN Optimization at the MAC and Network Levels*, 2005.
- [2] J. Y. Le Boudec, R. Merz, B. Radunovic, and J. Widmer. Dcc-mac: A decentralized mac protocol for 802.15.4a-like uwb mobile ad-hoc networks based on dynamic channel coding. In *Proceedings of Broadnets*, San Jose, October 2004.
- [3] J. Y. Le Boudec, R. Merz, B. Radunovic, and J. Widmer. A mac protocol for uwb very low power mobile ad-hoc networks based on dynamic channel coding with interference mitigation. Technical Report IC/2004/02, EPFL-DI-ICA, January 2004.

- [4] R. Merz, J. Widmer, J. Y. Le Boudec, and B. Radunovic. A joint phy/mac architecture for low-radiated power th-uwbb wireless ad-hoc networks. *Wiley Journal of Wireless Communications and Mobile Computing, Special Issue on UWB*, To Appear.
- [5] B. Radunovic and J. Y. Le Boudec. Optimal power control, scheduling and routing in uwbb networks. *IEEE Journal on Selected Areas in Communications*, 22(7), September 2004.
- [6] B. Radunovic and J. Y. Le Boudec. Rate performance objectives of multihop wireless networks. *IEEE Transactions on Mobile Computing*, 3(4):334–349, October 2004.
- [7] B. Radunovic and J.-Y. Le Boudec. Power control is not required for wireless networks in the linear regime. In *Proc. IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WOWMOM'05)*, Taormina, Italy, June 2005.
- [8] M. Villanti, M. Sabbatini, G.M. Maggio, and L. Milstein. Noncoherent code acquisition for UWB systems in dense multipath fading channels. In *Proc. 61st IEEE Vehicular Technology Conference, Stockholm, 2005*.
- [9] J. Widmer, C. Fragouli, and J.-Y. Le Boudec. Low-complexity energy-efficient broadcasting in wireless ad-hoc networks using network coding. In *Proceedings of First Workshop on Network Coding, Theory, and Applications*, Riva del Garda, Italy, April 7 2005.
- [10] M. Win and R. Scholtz. Ultra-wide bandwidth time-hopping spread-spectrum impulse radio for wireless multiple-access communications. *IEEE TRANSACTIONS ON COMMUNICATIONS*, 48(4):679–691, APRIL 2000.
- [11] H. Zhang and J. C. Hou. Capacity of wireless ad-hoc networks under ultra wide band with power constraints. In *Proc. IEEE INFOCOM*, Miami, Florida, March 2005.