Duration Information Propagation

Author: **Johnny Zweig**
Xircom, Inc.
2041 Landings Drive
Mountain View, CA 94043
Tel: +1 415 691 2500
eMail: jzweig@xircom.com

Introduction

In my Letter Ballot comments on D2.0, I made the suggestion that the way the information contained in the Duration field as defined in D2.0 clause 4.1.2.3 is propagated through a BSS be changed. Since the proposed changes would affect clauses 4, 6, 9, 11, 12 and 13, I have prepared this submission to discuss the motivations behind this suggestion.

Overview of Proposed Changes

I propose that the Duration field be deleted from the MAC header and moved to what we are currently calling the PLCP header. Since submitting my comments, I have realized that the name "PLCP header" is misleading, so I would further move that the name be changed to "Basic Rate Header". This entails inserting a new field called "DUR" into the currently-defined PLCP header formats for all three PHYs, and adding a value to the PHY service primitives whereby the MAC can pass this value to the PHY for outgoing frames and receive this value from the PHY on incoming frames. The PHY is used to propagate this information, uninterpreted at the PHY layer, throughout the BSS.

Motivation

The MAC relies on all stations having a reasonable notion of how long the medium will be busy during a frame exchange, so that they can avoid transmitting at a time that will interfere with other stations that have gained access to the medium for a frame exchange. An important piece of this function is provided by means of a Duration field that indicates how long the station transmitting a frame (an RTS or ACK, for example) anticipates that it will use the medium for the exchange of which that frame is a part. In particular, Data frames indicate through their Duration field value how long it ought to take to complete the data frame, send the ACK and complete the next fragment of the exchange (with its attendant ACK), so that other stations will defer to a sequence of back-to-back fragments, even if they are unable to hear the RTS and ACK frames associated with such a sequence.

Simply knowing how long the *current* frame is going to occupy the medium (and only for frames a station can hear) is not enough information for the MAC protocol to operate optimally.

The problem is multirate. There is currently a philosophical debate as to how to support stations that cannot all operate at the same maximum data rate. Some people argue for supporting a mix of "fast" and "slow" stations in the same area, and others feel "who cares whether old slow stations can interoperate with newer faster ones?" There are really two things at issue here: first, do we want to support mixtures of stations with different capabilities without having them stomp all over each other, and second, do we want fast units to be able to communicate with slow ones? That is, one can imagine a balkanized BSS in which the fast nodes could communicate only amongst themselves, and the slow stations would be relegated to a ghetto, but at least they did not interfere with one-another. Only if the answer to *both* questions is "no" does my suggested change not make sense.

Currently the PHY supports multirate by providing a header that is sent at a rate that all stations are guaranteed to be able to receive that indicates the data rate at which the rest of the frame will be transmitted. This Physical Layer

Convergence Procedure (PLCP) header also indicates the length of the associated MAC Protocol Data Unit (MPDU) being carried in this PHY Protocol Data Unit. The motivation for transmitting this part of the frame at "slow" speed is so that a station can be constructed so that it always listens for the PLCP header only at the basic rate, then switches to high(er) speed only for those frames that require it. The only reason for doing so is to allow a station to receive a mixture of basic rate and high rate frames (otherwise it would be simpler to have a single data rate switch that a station put into whatever position it felt like, and then could only communicate at that data rate).

Since the "innards" of the MAC PDU (MPDU) are sent at a rate that is not guaranteed to be received by all members of the BSS (except that, as a hack, we force certain control frames to be transmitted at the basic rate to work around the problems with the architecture), the Duration field is not always available to all nodes in a BSS. So slower nodes will not be able to update their NAV properly for high-speed data frames, and may have a mistakenly optimistic notion of when the medium will again become free. By putting the Duration information into the PLCP header, everyone can share the information, and the MAC works better.

Another benefit of this approach is that we can now send all control frames at whatever rate the intended recipient is known to be (or suspected of being) able to support. Since the NAV can be updated with information that everyone can receive, it doesn't matter whether nonparticipants in an exchange know whether the MPDU itself contains Data, RTS, CTS, ACK, CF-Poll, whatever. They simply defer for as long as the Duration field of the Basic Rate Header tells them to defer and everyone is happy. There could be a high-speed PCF, for example, in which only PCF-aware stations that are capable of high data rate communication could participate, but which could transfer information much more efficiently than the currently-defined PCF, which is hobbled by having to send control information at a slow rate, just because that is what it takes to keep the slower data rate and non-PCF-aware stations from screwing everything up in their ignorance.

So my proposal that we move duration information into the part of a frame that everyone can receive is motivated by the observation that the dual-speed mechanism we have already agreed upon could be exploited much more effectively by using it to propagate medium reservation information to all stations.

I do not even think that this represent a violation of protocol layering. The PHY's job is to exchange data at the behest of the MAC. It does this by sending modulated electromagnetic radiation that other stations can receive, and passing the demodulated information up to the MAC layer for processing. By design, the PHY must hide details of how it manages to communicate these bits from the MAC (this is why the PLCP headers of the various PHYs differ). It is simply a nomenclature convention that makes my suggestion of putting "MAC" information into the "PLCP Header" look like a layering violation. The Physical Layer Convergence Procedure doesn't "own" the basic data rate precursor to a frame which the MAC "owns". PLCP simply hides the PHY's details of delineating frames and handling data rate determination from the MAC.

So my letter ballot comments associated with the Duration field information in clauses 4, 6, 9, 11, 12 and 13 represent not a new mechanism, but a change to a broken mechanism that fails to take advantage of the procedures that are already in place to communicate Duration information to all stations in a BSS. The mechanism as currently specified puts information where it is less likely to be received (both because the Basic Rate Header is shorter than the MPDU and because some stations may not even be *able* to receive some MPDUs) than it would be by simply moving it to another part of the frame. There is no reason not to put the information there other than chauvinism and a false notion that, just because we've been calling the first part of the frame the "PLCP Header," it cannot carry information that is only of interest to the MAC layer.