Report of the Ad Hoc Multirate Study Group

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

Introduction

At the November 1995 meeting in Montreal, document 95/247 was discussed, and discussion of the proposed changes to the MAC definitions and the PLCP of each of the PHYs was tabled pending a chance to "sit down and figure this stuff all out" prior to the January 1996 meeting in San Diego. An ad-hoc study group was formed to look at mutirate support in the 802.11 standard, and try to come up with a recommendation. This discussion took place in email, the bulk of which was sent through the IEEE reflector (we had a brief offline discussion of whether to conduct the discussion on the reflector). This document attempts to capture the salient points that were brought up during the discussion.

The bottom line is that no single solution prevailed, though we did narrow discussion down to a small set of possibilities.

Summary of the Options Discussed

Let me editorialize a little first. The biggest problem that troubled our group is a widespread lack of quantitative information concerning the problems we discussed. Without this firm footing upon which to base our opinions, much of the debate of the multirate problem devloves into a "tastes great" versus "less filling" argument where each side believes their suggested course of action is best. There was reasonably broad agreement on what the problems are, but not on how serious those problems will turn out to be in practice. Thus each proposal to fix a problem of unknown severity brings with it an unknown cost, and it is difficult to agree on how, or indeed whether to solve each of the problems.

In broadest terms, the options that were discussed boiled down to three:

1) Leave the mechanism in D2.1 alone
2) Make some changes to the mechanism (as proposed in 95/247 along with additional changes to make those changes work better), while retaining the overall architecture for multirate support
3) Remove multirate support either substantially or totally, and decree that future different-bit-rate PHYs will be incompatible with the PHYs defined presently, with no goal of interoperability or even noninterference

The arguments in favor of option (1) ran along these lines: there are working implementations of the current draft, and multirate works. There is a speedup (of about 60%) when an FH PHY unit goes from 1 Mbps to 2 Mbps, and a product that is able to "shift gears" between 1 and 2 Mbps on the fly has advantages to a customer who doesn't want to buy lots of access points. The argument against is that NAV gets broken, and nobody knows how well large mixed populations of different-rate units will interact.

The arguments in favor of option (2) ran along these lines: it is clear that Duration information is important to properly setting the NAV, and that an incorrect NAV either presents opportunities for collisions that would not otherwise occur or wastes bandwidth (and could even make medium-access unfair). So the current mechanism is obviously broken, and moving the Duration field (with suitably-redefined semantics so that NAV updates work) into the PLCP header (where all stations in a BSS can receive it) helps. The argument against is that it propbably doesn't help a whole lot.

The arguments in favor of option (3) ran along these lines: the law of Diminishing Returns makes the technique of sending control frames, broadcasts, and even PLCP headers at the basic rate deliterious to performance. Since higher bit-rate devices have shorter range (a law of physics), customers will need to install APs more densely to support

such rates. So the ability to (less efficiently) support multiple rates in a single station is less attractive than allowing higher bit-rate PHYs to "go fast" and send all of their bits at a higher rate. The interference between the higher rate PHYs and the lower rate PHYs can be no worse than the interference between incompatible PHYs already is, and is probably dwarfed compared to interference from non-802.11 sources anyway. The argument against is that requiring parallel infrastructure is burdensome, and that the different PHYs may interfere to an intolerable extent.

Other Points that Arose

Since a number of other points were brought up during the discussion, I will try to mention them here.

It was widely agreed that constraining future PHY data rates was not good, though some people felt it is a necessary evil. The fact that the PLCP length field plus the bit-rate flags (for the FH PLCP) do not give sufficient information to tell when the current frame will end was mentioned.

Everyone agreed that supporting multiple rates was a good thing, since "fallback" to slower rates can increase range. Some said that each AP should be able to transmit and receive at all rates, and others felt it was good enough to have parallel infrastructure to brige between different-rate BSSs.

Nobody wants to hold up the standard, even if it means having imperfect mechanisms in it.

A "common" universally-receivable PLCP header was pointed out as a good thing, provided it reduced the likelihood of interference between overlapping BSSs or different units in a multirate BSS (even if the MAC contents of frames are not visible to everyone). Not everyone agreed, however, that it was necessary (i.e. *how* good a thing it is), nor that it was worth the performance reduction for higher data-rate PHYs. Problems with the FH CCA technique were touched on briefly.

An Idea

A random thought just occured to me, and as the semi-willing leader of the Multirate Study Group I am not above unabashedly making a suggestion in hopes it will help guide discussion. I don't see any reason why the PLCP needs a length field at all. A field that said how many microseconds the current frame will last, along with the bit-rate information, would suffice. Then even new, non-integer-multiples-of-1-Mbps PHYs could be defined and we would still have the "stomping protection" that motivated me to suggest appending the MAC header's Duration field (which typically includes information about how long the medium should be reserved for future frames as well as the current one) in the first place. Then we could implement the Zegelin/Ennis/Vesuna approach in document 96/8 or the Bagby approach of supporting only a singe bit-rate per BSS with reduced likelihood of "really bad" stomping (i.e. slower units mistakenly think the channel is busy alot longer than it really is, resulting in unfair access). The only down side is that for hypothetical data rates > 8 Mbps, the PHY might have to round the PLCP Duration field up, and deliver an extra octet to the MAC. This is not a big problem.