WAVEFORMS FOR DYNAMIC WIRELESS NETWORKS: IS LAYERING A GOOD IDEA?

John Chapin (Vanu, Inc., Cambridge, MA, USA; jchapin@vanu.com);
Karina Allain (Vanu, Inc., Cambridge, MA, USA)

ABSTRACT

Software radios offer an unprecedented degree of flexibility to designers of wireless networks. SDR enables dynamic modification of physical layer parameters which can be used to improve overall system performance. However, it is difficult to exploit this opportunity. Excessive flexibility leads to networks that are complex to implement, verify and manage safely, resulting in compatibility problems, high maintenance costs, and possibly worse performance than static wireless networks. This paper explores the applicability of layering as a structuring mechanism to reduce the complexity. We report our experience with implementing a waveform that supports packet-by-packet variation of modulation, symbol rate, and other parameters. One interesting result is a requirement on signal processing middleware packages to provide greater application-level scheduling control.

1. INTRODUCTION

Software-defined radios and radio networks promise the ability to modify any aspect of their signal processing at any moment. This capability offers the potential to significantly improve network performance in dynamically changing environments. We call a network that exploits physical layer flexibility in this way a dynamic wireless network.

Dynamic wireless networks are already in wide use. One good example is the CDMA 1xEV-DO system from QualComm, which changes physical layer characteristics in response to channel conditions. [1] However, the physical layer flexibility exploited by current systems is limited by their implementation as fixed-function hardware radios with only a few operating modes. The widespread deployment of software radios will create the opportunity for much greater flexibility. Future dynamic wireless networks will be able to change modulation, coding, spreading, center frequency, timing, and other characteristics.

A prominent example of a future dynamic wireless network is the JTRS Wideband Networking Waveform. [2] The WNW offers dozens of operating modes and bandwidths to support a wide range of applications and environmental conditions. It currently includes four separate modulation types (called Signals In Space), each with a range of tunable characteristics. More modulation types are planned for the future. The sophistication of the WNW will grow even greater when it is combined with dynamic spectrum access methods such as those being developed in the DARPA XG program. [3] The acquisition authority for WNW has stated that dynamic spectrum access will be a requirement in order for WNW to reach its performance potential. [4]

The level of flexibility inherent in SDR and in systems like WNW and XG creates significant design challenges. It is unclear how to structure these systems and the waveform software that implements them. If the systems are structured poorly, an excessively high fraction of the available communication bandwidth can be consumed by control traffic. The system can also fail to converge on good selections of physical layer parameters and thereby end up performing worse than a simpler static wireless network. If the waveform software is structured poorly, it will not support straightforward evolution to new communications modes or control algorithms.

Similar design challenges faced the researchers developing the Internet in the 1970s and 1980s. Their solution was a hierarchical structuring method later generalized as the OSI layering model. The layered model was highly successful at simplifying both system and software design. It is natural to ask whether the layered model can offer similar benefits for the structure of dynamic wireless networks.

This paper investigates this question and reports on experiments done in an NSF-funded Networking Research Testbed project. The overall question of how to structure software radio waveforms for dynamic wireless networks goes far beyond the analysis and results described here. Our intent is to frame the question carefully and stimulate further research.

First, we characterize dynamic wireless networks and the layered approach in more detail. We then carefully consider the relevant differences between wired and wireless networks and how this may affect the applicability of the layered approach. Finally we report on the experiments we performed to explore one point in the design space.

2. DYNAMIC WIRELESS NETWORKS

The goal of dynamic wireless networks is to overcome the inefficiency of static wireless networks, which have to be designed for worst-case conditions. A dynamic wireless network repeatedly changes its behavior to improve performance by exploiting better-than-worst-case RF conditions, topology, power or processing resource availability, and similar opportunities.

application requirements	data rate
	latency
	minimum acceptable BER
	message acknowledgement
	threat level
radio environment	RF conditions
	congestion
	topology
	infrastructure availability
	co-site interference
resources	allotted spectrum
	battery power
	processing power
hardware capabilities	receiver sensitivity
	tuning range
	transmit power limit

Table 1: Effects in dynamic wireless networks that may require changing waveform behavior.

Discovery of new coding or modulation techniques.
Novel users and applications.
Discovery and patching of security holes.
Administrative requirements of different users

Table 2: Effects in dynamic wireless networks that may require changing waveform software.

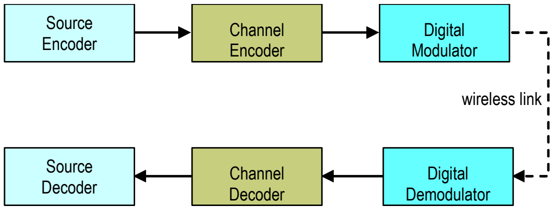

Figure 1: Idealized wireless communication systems look attractive for layering

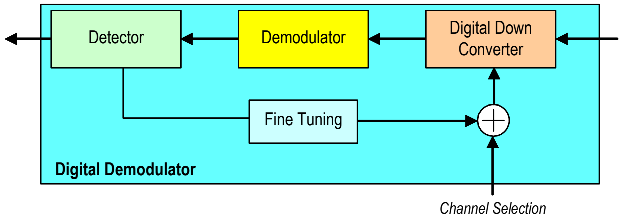

Figure 2: Actual implementations include feedback loops that impede layering.

Dynamic modification may be required in two different time scales. On a short time scale, changes in the environment or application requirements require radio level adaptation (Table 1). On a longer time scale, technology development may require software updates (Table 2). As a result, the network design must function correctly even when multiple software versions operate concurrently.

The key challenge in the dynamic wireless networks problem domain is emergent behavior. The overall goal of improving performance is achieved through distributed control. Each individual node knows only partial and potentially incorrect information about the state of the system. Dynamic wireless networks are tightly coupled complex systems where it can be difficult to predict system behavior from individual component behavior, and vice versa, to design individual component behaviors to produce a desired system level behavior.

3. THE LAYERED APPROACH

Layered structuring techniques for networks are well described in the literature. [6] The best known example is of course the Internet Protocol (IP) stack. The relevant principles of layering for this paper are:

- A layer interfaces only with the layer immediately above and the layer immediately below,
- Each layer performs a more primitive function than the layer above
- Complex, specific functionality at higher layers is implemented by exploiting simpler, more general functions at lower layers.
- Layers can be replaced or modified independently.

Layered structuring techniques for software radios were suggested by Bose. [7] In his model, a Link Framing layer eliminates transmission errors through the use of checksums, sequence numbers, and retransmission requests. A MAC layer mediates access to the transmission medium. A Coding layer performs line coding, channel coding, and symbol mapping. A Modulation layer transforms symbols to sampled signals. Finally, a Multiple Access layer allows for sharing of the media by more than one network. Although Bose suggested this approach, he did not investigate its merits in depth.

It is not immediately clear whether layering inside the physical layer is a good idea. Layering makes sense when considering a wireless communication system in an abstract, idealized way (Figure 1). Looking in more detail at actual systems, we find heavy use of feedback and feed-forward signals to improve performance (Figure 2). If a feedback loop crosses a sublayer boundary then the principles of layered structuring will be violated. One layer's complexity and implementation details may be exposed to another.

5. LAYERING: WIRED VS. WIRELESS NETWORKS

We seek to understand whether layering will work as well for dynamic wireless networks as it did for wired networks, and more particularly whether it is a useful way to structure the internals of the physical layer. To investigate this question, we analyze how the differences in the two networking problems affect the benefits of layering. Table 3 lists the relevant differences.

The first three entries concern the properties of the physical communications medium. In a wired network, the wire connecting a computer to a switch has beneficial properties, compared to the radio ether in a wireless network. The physical properties of a wire do not change over time. Its properties can be selected at design time, whereas the properties of the ether are a given. Finally the wire is shared equally by all nodes attached to it. In contrast, wireless communications suffers from the hidden node problem, in which two nodes are physically disconnected when transmitting to each other but apparently connected when transmitting to a third party, since their signals interfere at its receive antenna.

These physical level media properties of wired networks fit well to a layered model. The designer of the physical layer can focus entirely on managing a static, controlled, single shared wire. In contrast, in a wireless network, the uncontrolled properties of the physical medium make it difficult to localize all control in a bottom layer of the software. For example, the hidden node problem affects network-level routing decisions, which in turn determine whether physical layer collisions will occur.

The next two entries in Table 3 concern static versus dynamic behavior. In a wired network, the waveform is generally static, and the link topology is fixed over extremely long time scales. In a dynamic wireless network, the waveform and the topology both change quickly. There has been substantial MANET work on preserving routing capabilities despite changing topology. However, one challenge not yet widely recognized is the difficulty of initiating access to a dynamic wireless network. How does a new node rendezvous with a network of nodes that are dynamically modifying their frequency hop set, modulation scheme, timing, and other parameters? Effective rendezvous will involve the cooperation of multiple layers of the protocol stack.

Wired Networks	Dynamic Wireless Networks
static media	changing RF environment
controlled media	uncontrolled media
single shared media	hidden nodes
waveform does not change	waveform may change, making rendezvous difficult
fixed topology	changing topology
routers control network hierarchy	no routing control points
one-to-one on the subnet	one-to-all within range
one channel	multiplexed channels
peer-to-peer communication	peers within the physical layer cannot communicate directly

Table 3: Differences between wired networks and dynamic wireless networks that affect the benefits of layering

Next we consider issues of control. In a wired network, routers control hierarchy and topology, and can assure that packets sent by one node are delivered only to the intended destination. In contrast, in a wireless network all nodes within range receive a transmission. Delivery is not controlled by any central router. This means that the physical layer of the system cannot rely on a higher routing layer to shield it from the complexities of the topology, as is done through layering in wired networks. Similarly, in a wired network there is usually one channel traversing a given wire, whereas in a wireless network there are multiple multiplexed channels with different properties. Decision making about which channel to use for a given communication task combines detailed information from multiple layers.

Finally, choosing to use layering within the physical layer challenges some of the fundamental design choices made in traditional layering approaches. In a traditional approach, since the entire physical layer is monolithic, all higher layers have at least the ability to exchange bits with peer layers on other nodes. If we use layering inside the physical layer, the individual layers do not have a digital communication channel to their peers, as received bits are only available near the top of the physical layer.

To summarize, we have noted in this section a variety of differences between wired and wireless networks relevant to the applicability of layering for software radio waveforms in dynamic wireless networks. Nothing here indicates that the approach is not a workable or good design methodology. However, the differences suggests that the fundamental reasons layering worked well in the wired case may not apply in the wireless case, and so it needs to be carefully investigated.

7. EXPERIMENTS

We have investigated the layering approach during the development of a dynamic wireless network for an NSF-funded Networking Research Testbed (NRT). The system could select symbol rate and modulation bandwidth on a per-packet basis, with the receiving node dynamically detecting the physical structure used and performing the appropriate demodulation and reception algorithms.

The experiments consisted of a two-way wireless IP network connection between two PCs with digital transceiver PCI cards. Communication occurred full-duplex in 200 kHz channels at 72.2 MHz and 72.8 MHz using a Continuous Phase Frequency Shift Keyed signal, achieving 150 kbits in the highest performance mode and 10 kbits in the highest reliability mode. An abstract view of the waveform software structure is shown in Figure 3.

This effort uncovered two interesting implementation challenges, described in the following sections.

8. BURST LENGTH VARIATION

Networks using fixed time-duration transmission bursts have an advantage in maintaining synchronization in that transmission boundaries are regular and predictable. In the waveform we implemented, bursts varied in length due to the varying modulation rates used. This created high computational costs associated with scanning for the start of the next transmission burst whenever synchronization was lost due to a noise burst.

The costs we observed suggest that dynamic wireless networks may work better if fixed duration bursts are used. This is especially true if TDMA is necessary, since dynamic burst durations are difficult to manage in a TDMA system.

However, a requirement for fixed duration bursts violates layering in fundamental ways. A highly variable physical layer, which modifies modulation, coding, bandwidth, spreading, and other parameters, will result in wide and rapid variation in the number of bits transmitted in a fixed time duration. Therefore the only layer with sufficient information to predict the burst boundaries in the transmit bit stream is the physical layer. However, it is not appropriate for the physical layer to perform the packetization internally. The size of bursts affects higher layers. For example, block code length should ideally match transmission burst length. For maximum system efficiency, burst lengths also drive the size of application level frames such as voice blocks.

Burst duration is a fundamental wireless network design property that affects more than one sublayer of the physical layer. If fixed length bursts are used, separation of layers is difficult to maintain. If dynamic burst lengths are used, the system is less efficient.

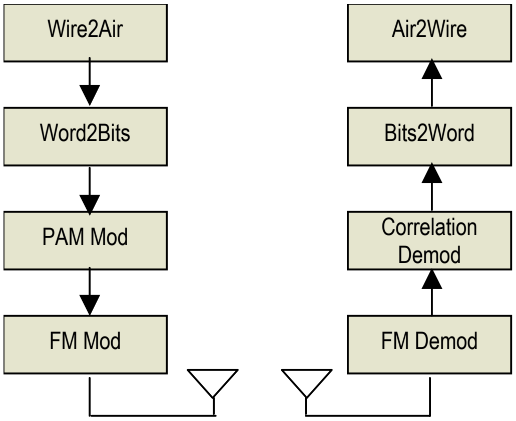

Figure 3: Software structure of experimental waveform

9. HEADER-FIRST PROCESSING

In wired IP protocols, all layers have a digital communication channel between them. The physical layer frames, the network layer datagrams, and the transport layer segments all have headers for peer-to-peer communication. Also, the header information associated with a data packet is always readable at the time the data arrives.

These assumptions do not hold when the physical layer is subdivided. The physical layer creates a digital channel on top of the physical media. If the physical layer is divided into sublayers with only the highest sublayer finally achieving digital communication, then the other sublayers cannot communicate directly with each other.

For the sublayers to communicate, the receiver must pass all headers up to the layer that is able to decode them. After decoding, this layer sends the decoded headers back down the stack to the appropriate sublayers (Figure 4).

The communication between peers at a sublayer is indirect in that the headers must pass up the stack and back down before their contents are useable. The communication is also delayed since the header contents arrive at the sublayer after the header itself has passed through.

In general, the header information is needed to correctly process the payload data with which it is associated. That is, the signal processing pipeline must be configured with information from the header before the RF samples of the payload pass through it. We found it necessary to break the sample stream at the appropriate time point and decode the header and the data separately. We passed the header block up and back down the physical layer sub-stack before processing the data block.

This solution does not impose any new timing requirements since the sampled data block can simply be stored at the lowest layer until the header block is decoded and returned. However, this solution does impose design restrictions on the signal processing middleware. Earlier versions of our middleware did not provide application-level control over the scheduling of pipeline stages. As a result, the application could not prevent the middleware from pulling payload RF samples through one or more processing stages before the parameters of those stages were set based on the header data. It turned out to be challenging to modify the middleware to support this waveform. The challenge centered on providing application-level control over data flow scheduling while hiding this complexity from most applications that do not need it.

10. CONCLUSION

The analysis and initial experiments performed in this investigation have shown that dynamic wireless networks are sufficiently different from traditional wired networks that many aspects of the layered design approach need to be reevaluated.

We identified a number of difficulties in using layering. Issues such as transmit burst length and header-first receive processing appeared during an implementation effort. It seems likely that further insights will appear as more sophisticated implementations are constructed.

Despite the difficulties, layering still appears to be a promising approach to designing waveform implementations for dynamic wireless networks. Layering assists in managing software design complexity, supports unit specification and testing, parallel development, and other engineering activities.

One fruitful area for further investigation is to relax or modify the basic principles of layering articulated for wired networks. It is possible that a slight modification to the common approach could overcome some of the difficulties in applying it to dynamic wireless networks without hurting its fundamental benefits for system design.

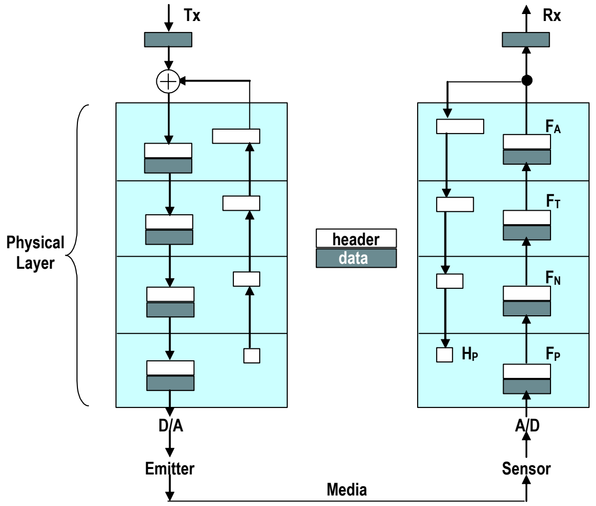

Figure 4: Header information flow in the waveform. Higher sublayers must be involved for lower sublayers to communicate.

REFERENCES

[1] E. Esteves, P Black, M. Gurelli, “Link Adaption Techniques for High-Speed Packet Data in Third Generation Cellular Systems,” European Wireless Conference, 2002.

[2] D. Lofquist, “JTRS Network Services Layer,” GOMAC 2004, Monterey, CA, March 2004.

[3] `http://www.darpa.mil/ato/programs/XG`

[4] Presentation by CECOM Cluster 1 Office at AIE Cognitive Radios Conference, Las Vegas, NV, March 2004.

[5] P. Marshall, Spectrum Management Protocols in the DARPA Next Generation Communications Program, SDR’02 Tech Conf, San Diego, November 2002.

[6] A. Tanenbaum, *Computer Networks*. Prentice Hall, Englewood Cliffs, NJ, 2nd ed., 1988.

[7] V. Bose, *Design and Implementation of Software Radios Using a General Purpose Processor*. PhD Thesis, Massachusetts Institute of Technology, June 1999.