SPECTRUM MANAGEMENT PROTOCOLS IN THE DEFENSE ADVANCED RESEARCH PROJECTS AGENCY NEXT GENERATION COMMUNICATIONS PROGRAM

Preston F. Marshall (Department of Defense, The Defense Advanced Projects Agency, Advanced Technology Office, Arlington, VA. USA. pmarshall@darpa.mil[1]

ABSTRACT

The Defense Advances Research Projects Agency (DARPA) is developing revolutionary spectrum access technologies that potentially can provide exploit an order of magnitude more spectrum than current practice. However, this investigation is being performed in advance of regulatory or management framework development. The neXt Generation (XG) program is therefore developing both specific and abstract characterizations of the required behaviors in order to enable subsequent policy development to appropriately and retroactively develop the automated rule sets required to control the operation of these systems. This paper describes the approach DARPA is following in development of the XG program, and specifically the approach to the protocols that would be the basis of the program. The ideas in this paper are preliminary concepts that are the basis of DARPA's planning for this program. Further design and development of XG is now being performed by a number of industry participants in the XG program.

1. INTRODUCTION

The Defense Advanced Research Projects Agency (DARPA) neXt Generation (XG) program is developing a new generation of spectrum access technology in order to both increase the ability of military systems to access spectrum, and to ensure the rapid deployment and operation of new generations of weapons systems, without the extensive, frequency by frequency, system by system, coordination now required for each country in which these systems will be operated. DARPA has two objectives in doing this. First, the current spectrum management process emphasizes deconfliction of spectrum assignments though fixed allocations and assignments. Particularly when U.S. forces are overseas, it is a difficult and lengthy process to identify frequencies for operation of each of the large number of spectrum dependent systems that the U.S. operates. This lengthy process is inconsistent with the high tempo and quick reaction timeline that is need for modern operations. Secondly, the current approach to spectrum management is not able to optimally, or even effectively, pack the spectrum, resulting in difficulty in obtaining frequencies for many systems, and potential conflict between military and civil users.

The development of this technology is being performed in advance of any policy or management approach for their use, either in the United States, or overseas. A key challenge is therefore to develop a flexible technology to which these considerations can be subsequently applied. Simultaneously, we need to develop a general framework so that regulatory bodies can address and manage key aspects of the common XG radio behavior, in order to avoid having to address each implementation of the XG concepts individually.

This paper describes some of the overarching principles in the development of the XG behaviors and protocols. As work on this project continues, it is likely that these approaches will be adjusted and refined.

2. XG PROGRAM BACKGROUND

A commonly held perception is that spectrum is scarce, and that spectrum access can not be provided for new applications without reallocating spectrum from existing users. The traditional approach to interference avoidance is through analysis of all "worst case" assumptions, and imposition of regulatory constraints on all users of the spectrum as to power, frequency, directionality, etc. that preclude any "possibility" of interference. The DARPA XG program challenges these assumptions and approaches based on the premise that the critical issue is not extensive spectrum usage, but limitations in the ability of current technologies to access spectrum in an interference controlling manner. The XG program is therefore developing the enabling technology to dynamically manage spectrum while avoiding interfering conditions.

We are concerned about managing spectrum, and spectrum access, because it is perceived to be scarce. This perceived scarcity is worth close examination. Our analysis at DARPA leads us to believe that the issue is

---

[1] The opinions provided in this paper are those of the author, and are not the official position of DARPA or the Department of Defense

not so much that spectrum is scarce, but that we do not have the technology to effectively manage access to it. We should be able to load an order of magnitude more users into existing spectrum, rather than reach contention at the very low average utilization rates we currently measure. This is true even in crowded urban environments.

Much of this has to do with the spatial and temporal characteristics of spectrum. When we first started managing spectrum, we had relatively long range transmissions, and we had essentially continuous broadcasting operations. Clearly, these conditions are true for only a small subset of today's radio applications. As we move up in frequency, we tend to have very localized effects due to physics of propagation, and our current, "bandwidth on demand" services are highly variable in their operational tempo, and thus their temporal characteristics. It is not surprising that a centralized, and long duration planning process is not effective in achieving the maximum spectrum loading.

We should view these newly arising characteristics as opportunities to provide different, and more aggressive spectrum management. However, exploiting these characteristics can only be performed locally and on an instant-by-instant basis. The RF environment at 1 MHz will be the same over a very large region, but the environment at 1 GHz is much more localized, and may be far different on either side of even a small hill. At 10 GHz it varies over the range of tens of meters. This is the basis for the technology DARPA is developing. From an engineering perspective, it is not a break from how interference has always been analyzed; the change is that this management is now placed in each radio, where it can assess the actual situation at each instant in time, rather than have to be deconflicted, in advance, for any possible situation of time, position, signal, propagation, etc. Only a few of these driving, or constraining conditions will be present at any one time, and these are the only ones that need be considered in developing an interference avoidance tactics. The radio itself is most aware of these conditions.

This is not to trivialize the problems of very rapidly sensing and characterizing the environment, and then developing, distributing and updating frequency planning across a network. Each of these problems has been solved for other applications, so we are confident we can solve them for spectrum management. In fact, we are so convinced that there are probably many good ways to perform some or all of these functions that we are ensuring that our strategies are flexible enough to allow for multiple instantiations of the core processes, within a single abstraction of how XG should operate.

Perhaps the greatest challenge is ensuring that our approach is not unduly influenced by how we plan to implement the solution today, and is instead a flexible framework that can be used for decades after the XG project is completed at DARPA. Certainly, the longevity of the Internet Protocol is proof that such frameworks are possible.

3. XG PROGRAM APPROACH

Our initial approach was to consider XG as a radio program, utilizing software controls to adjust its frequency as previously described. We planned to develop waveforms that optimized the performance of such a radio, and to embed the control protocols within the radio as part of its design. This is a similar approach to that adopted for several other DARPA communications programs, such as the demonstration of Ad-Hoc networking in an Low Probability of Intercept/Low Probability of Detection (LPI/LPD) Small Unit Operations (SUO) radio and directional networking in Future Combat System Communications (FCS-C).

While this approach would demonstrate the feasibility of the XG technology, it would not have provided a framework that would be enabling to the implementation of these features within a broad range of radios, nor would it establish a basis from which broad regulatory approval could be obtained. The process would still be one radio at a time, and would not advance our objective of an overarching technology that could be managed above the level of individual radio approvals.

Instead, we have decided to approach XG at two levels. The upper level is the development of a generalized, abstract set of behaviors that could be broadly applied to many of the existing, and likely future Media Access Control (MAC) and Physical (PHY) layers. These abstract behaviors are to be "agnostic" as to the "best" layer design, and will be limited to the core set of behaviors that were necessary to implement the XG functionality. We plan to develop these through the level of generalized Application Program Interfaces (API's). If suitably developed, these behaviors could be so generic that they could be the basis for broad acceptance of any radio that provided a trusted instantiation of the abstract behaviors, and would reduce the acceptance process to a technical verification of the implementation, rather than the analysis and policy development on a radio or system by system basis. We do not believe the API level is necessary to establish a regulatory baseline, but we are developing this to verify the interfaces for the abstract behaviors.

Our core behavior set is based on the simplest possible representation of frequency control within a radio. In this model, the frequency control function supports the Physical layer, by controlling the RF systems to transmit and receive on the appropriate frequencies.

In order to provide assurance of non-interference, the XG controls must exchange spectrum sensing with other XG radios, and their XG layers. We believe it is important to think of the XG network as being quite a different thing than what the host radio's Network Layer thinks it is. In its core functionality, we believe that XG can be implemented as a layer 2 process, having dialog with other radio's XG Layer 2, but having no interaction with other layers on either host, except for the Physical layer. This relationship is shown in Figure 1.

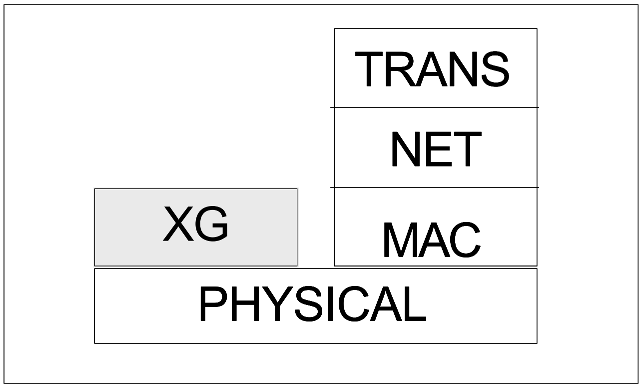

Figure 1 XG Layer Relationships

In this implementation, the XG "MAC" layer uses the native Physical layer to coordinate with other XG systems. The Physical layer is the only "XG Aware" layer, in that it must recognize that certain of the MAC requests imply action by the XG layer. This approach has the advantage that it (hopefully) allows the use of XG with not only existing MAC designs, but should be capable of supporting legacy MAC code before the transition to XG-Optimized MAC and Physical layers is made. The communications among XG and legacy MAC layers is shown in Figure 2.

In this framework, the Physical layer detects MAC requests that have implications for the XG layer. These generate requests to the XG layer that must be processed before the pending MAC layer request can be satisfied. The XG layers are utilizing the Physical layer to dialog and exchange spectrum utilization perceptions, and then to coordinate frequency assignments for the radios in the physical network. This exchange is essential because we need to both ensure that the selected frequency is usable at the receiver, and is not likely to jam signals from the environment of the transmitter. The Physical layer must arbitrate the multiple "users" of its services.

In some cases, it is likely that the XG protocols cannot use the native Physical layer. This is true when this layer has unique characteristics, or when there is no common mode of operation, such as in linking heterogeneous networks. In the most extreme case, one of the XG-enabled systems may not be a communications system at all, but may be sensor that communicates and/or coordinates spectrum usage with communications or other sensor systems. In this case, we conceive that we will need to define a standard for an XG interoperability path, which can be selected as part of the XG standard, or negotiated among the radios. With Software Defined Radios the likely implementation platform for XG, the introduction of an additional physical layer is not as significant as it would be with discrete implementations, but is still a complexity we would like to avoid. Some means of determining a common mode of operation could be a more suitable solution. This is a technology that DARPA is investigating in other programs, and may remain outside of the current XG work. Figure 3 shows this protocol architecture.

This representation is only the simplest form of XG. Clearly there are very significant benefits to the system's ability to be aware of, and to utilize network topology information that is only accessible in the upper layers, such as the membership data that likely resides in the Network layer. We will be investigating these, and

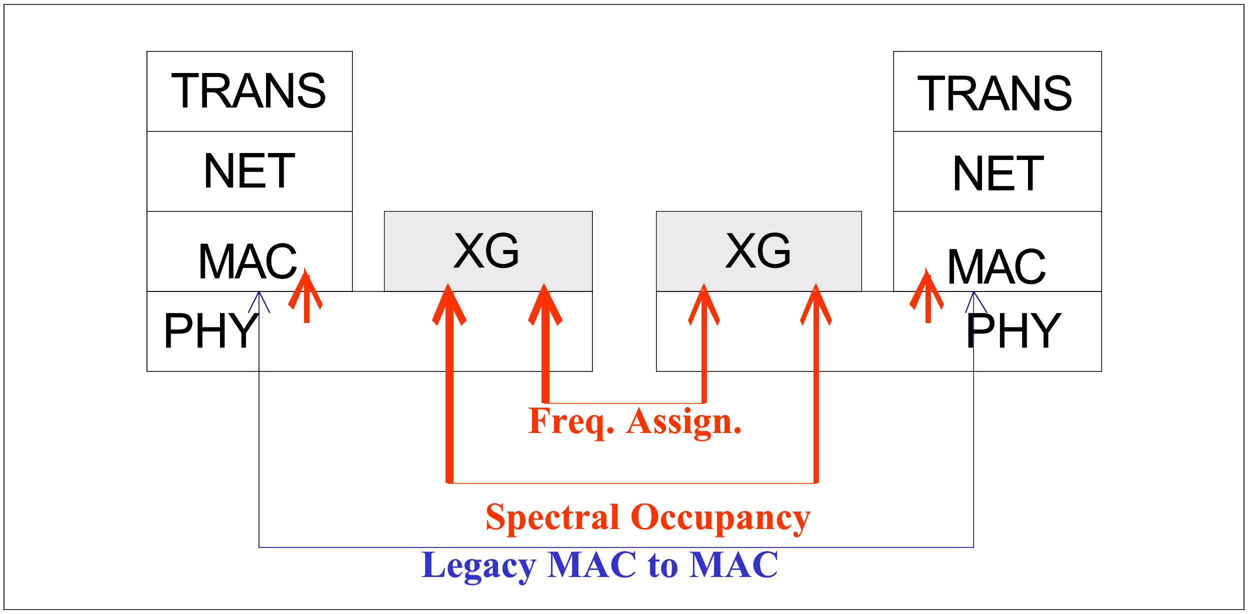

Figure 2 XG Layer Network Interaction

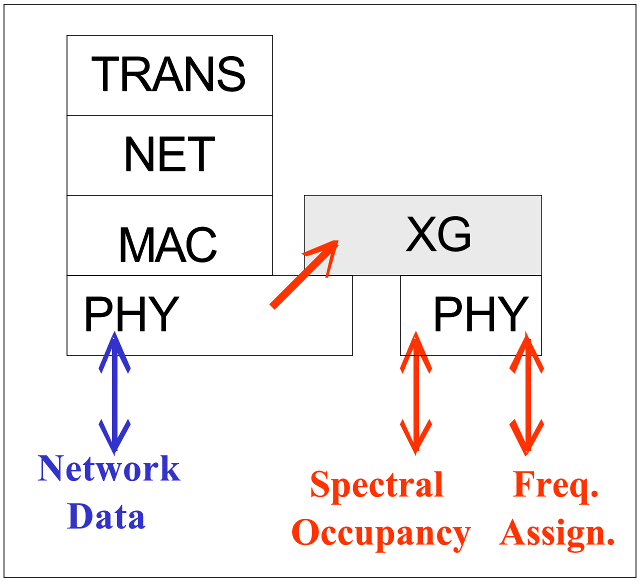

Figure 3 XG Unique Physical Layer Relationships

similar, opportunities for enhanced performance later in the program. If possible, it is intended to develop this functionality in the context of the same set of abstract behaviors that are used in the core architecture, therefore removing the enhanced functionality as an issue that must be addressed in the spectrum regulatory process.

In this process, we are also grappling with different concepts of what the "XG network" is. Although the first reaction is to think of this network as identical to the network to which the radio is a member, this definition yields an expansive scope to the XG solution, and implies that XG should sit on top of the upper layers. Direct control over the physical layer jeopardizes layer separation and independence, and makes an XG implementation specific to the Layer 3 technology. In fact, there is no reason to believe that the host radio's Layer 3 network is even the appropriate network scope for XG to consider. This network is much more expansive than the scope of influence of XG, and may not include other radios that may locally interact with an XG radio. The issue of whether XG needs layer 3 services is unclear at this time. If it did, it is likely that the operation would be between heterogeneous networks that were physically overlapping. These nodes would likely not be part of the same layer 3 structure.

One vision of the XG network process is that it is not a packet or data centric network, but one that operates as a Layer 2 information fusion and dissemination structure. In this view, nodes never relay packets, but instead maintain awareness of certain content, such as spectrum usage at their location, and some neighboring ones, and local frequency planning. They provide this information on a "Need to Know" basis with other nodes, but reprocess it to only provide the minimal set of awareness. They can correlate reporting, so that each node does not report common blocked frequencies, and only candidate frequencies are relayed. This can both eliminate significant overhead bandwidth and also sets the stage for the decision process at each site. This framework does not require routing information to be distributed, as each node joins only to other nodes that are within the same RF environment. No XG unique address management is needed.

In this model, the entire XG operation is a set of bilateral negotiations, on a link-by-link basis. In the case of Local Area Networks, with multiple participants, the number of nodes entering the dialog is increased, but the same framework remains applicable. As this negotiation proceeds, nodes can also interrogate adjacent nodes (in their logical network or not) to obtain sensing and spectrum opportunity data from their adjoining nodes. Since no end-to-end routing is provided, this limited bridging avoids the necessity of providing any XG-specific Network Layer services at all.

4. XG PROGRAM ABSTRACT BEHAVIORS

The development of the XG systems is based on a generalized set of abstract behaviors. We have taken this approach because there are clearly many ways of implementing the core functions. At the same time, establishing a national and international regulatory environment is a long-term process, and one that can not constantly adapt to changing concepts and approaches. Therefore we want to approach these communities with a simplified and generic set of behaviors that can encompass a wide range of possible implementations and evolutions of the XG concepts.

The other abstraction we have adopted is to characterize the control over XG thresholds and operation by a Policy-based meta-language. In this way, we hope to isolate the general framework that XG needs to operate, from establishing and advocating specific thresholds and rules of operation. We will develop this meta-language so that we can demonstrate that it can reflect the types of operational controls that national regulatory authorities would wish to impose. Initially, conservative rule sets and thresholds may be imposed, but these can be broadened and tailored as experience is gained. These changes would not impact the underlying XG systems, as the meta-policies would be isolated from the XG implementing behaviors. Once adopted, a set of policy controls should be independent of implementation, so long as it follows the abstract reference behavior model.

The XG development process will first establish a set of behaviors that ensure systems are *Interference Preserving*. We define interference preserving to be a guarantee that the introduction of the signal will not degrade the performance of any then operating system by more than a set threshold. Those portions of the XG system that are within this boundary should be of interest to the regulatory community, as these contain the functions that achieve the same objectives as the current regulatory regimes. By isolating this subset of the XG system, we hope to provide a focus for regulatory consideration that is compact, and does not necessitate that the regulatory community become involved in all aspects of XG implementation.

Only a subset of the XG system is within the Interference Preserving boundary. Much of the system will address housekeeping. Other aspects may address optimization of spectrum use. We desire to develop a framework that assures that optimizing methods lie outside of this boundary, as this will enable the continued progression of XG capability and performance, without requiring that these actions be addressed within a regulatory process. A generalized framework is optimal for DoD interests, as it provides a means to address a large number of systems within a single context; a context that may well have been adopted to enable civil uses as much, or more, than military ones.

A model for the concept of Interference Preserving behavior is similar to the requirement for Security Preserving systems. In the security realm, we have adopted methods that isolate the critical aspects of the larger system into a small, and highly controlled subset. Typically this provides a Trusted Kernel, Physical Red/Black isolation, cryptographic or a similar mechanism that enables us to think of the bulk of the system as outside of the security boundary. We need a similar framework for our interference preserving boundary.

Our model for this is an object-oriented analysis framework, in which the basic XG classes have defined interfaces and high-level behaviors, and for which the details are left for instantiation by the behavior implementation. For example, there is an obvious XG class managing information about spectrum occupancy. Regardless of its implementation, we would expect that it would have access methods that address energy and frequency. Different implementations would have different noise floors, frequency resolution, scan rate, etc. The interference methods will need to understand how these parameters affect the ability to determine non-interfering opportunities, but need not be specific to these values.

In similar ways, we can describe an abstraction of the process of using a set of spectrum occupancies to select a spectrum opportunity for a given spectral power density, time, bandwidth, etc. The actual implementation could vary, but would be constrained to ensure compliance with the policy-based meta-language controls. So long as the XG operation could ensure this compliance, we would leave the implementers free to develop and evolve ever more capable instantiations of these behaviors.

For DoD purposes, we will also develop a set of intra-XG Layer control protocols. We believe these are below the level of regulation, but they are essential to ensure interoperability of XG compliant systems as our initial demonstration system matures, and other follow-on developments are deployed. We will make these protocols publicly available, and will attempt to accommodate non-DoD interests and features to the extent compatible with our objectives.

5. SUMMARY

DARPA is in the process of developing a set of behaviors and protocols to manage spectrum integrally to the radio and network operation. These are being developed to both serve as a basis for the immediate validation of the technology, and also so that they can serve as a building block for both required regulatory action, and for the enhancement and natural evolution of the technology to become more network and system aware. This is an ongoing process that will be performed with the maximal opportunity for community involvement.