XG Working Group

The XG Vision

Request For Comments

Version 2.0

Prepared by:

BBN Technologies
Cambridge, Massachusetts
U.S.A.

1 About this document

This document describes the vision for the DARPA neXt Generation (XG) communications program. It lays out the motivation for XG and its scope, presents the key concepts underlying XG, and describes an approach for defining XG.

This document is a Request For Comments (RFC). Accordingly, an important purpose of this document is to obtain feedback from the community at large, and refine the ideas here based on the feedback. In other words, the XG vision will evolve, and this document reflects a snapshot in that evolution.

A number of other RFCs related to XG exist, or are being planned. The complete XG family will include the following:

1. *XG Vision RFC*. This document.
2. *XG Architectural Framework (AF) RFC*. The AF RFC presents the architecture, system components, and a high level concept of operations for XG communications.
3. *XG Abstract Behaviors RFC*. The abstract behaviors RFC identifies key behaviors that must be implemented by an XG system, organizes them, and describes the behaviors.
4. *XG Policy Language RFC*. The XG policy language RFC describes the policy specification meta-language for implementing machine-understandable policies.

This document is the "highest level" RFC, and should be read before any of the others. Its focus is on *what* XG is, in terms of its capabilities and scope, rather than on *how* XG will work. In a sense, this may be considered an evolving "executive summary" of the DARPA XG program. In short, it is intended to be a one-stop-shop for answering "What is XG" type questions.

2 An Introduction to XG

The Defense Advance Research Projects Agency (DARPA) neXt Generation (XG) communications program is developing a new generation of spectrum access technology. This section presents the motivation behind XG and its overall goals.

2.1 Motivation

There are two significant problems confronting wireless communications with respect to spectrum use:

- *Scarcity*. The current method of allotting spectrum provides each new service with its own fixed block of spectrum. Since the amount of useable spectrum is finite, as more services are added, there will come a point at which spectrum is no longer available for allotment. We are nearing such a time, especially due to a recent dramatic increase in spectrum-based services and devices.
- *Deployment difficulty*. Currently, extensive, frequency by frequency, system by system coordination is required for each country in which these systems will be operated. As the number, size, and complexity of operations increase, the time for deployment is becoming unacceptably long.

Both problems are a consequence of the centralized, static nature of current spectrum allotment policy. This approach lacks the flexibility to aggressively exploit the possibilities for dynamic reuse of allocated spectrum over space and time, resulting in very poor utilization and *apparent* scarcity. It also mandates a priori assignment of spectrum to services before deployment, making deployment difficult.

Preliminary data indicates that large portions of allotted spectrum are unused (refer the Spectrum Policy Task Force report). This is true both spatially and temporally. That is, there are a number of instances of assigned spectrum that is used only in certain geographical areas, and a number of instances of assigned spectrum that is only used for brief periods of time. This wastage is bound to increase in future – spatially, due to the increasing localization of propagation due to radio devices moving up in frequency, and temporally due to the proliferation of services that are highly bursty in nature.

Studies have determined that even a straightforward reuse of such "wasted" spectrum can provide an order of magnitude improvement in available capacity. It can be concluded that the issue is not so much that spectrum is scarce, but that we do not have the technology to effectively manage access to it in a manner that would satisfy the concerns of current licensed spectrum users.

In order to address the scarcity and deployment difficulty problems, XG is pursuing an approach wherein static allotment of spectrum is complemented by the opportunistic use of unused spectrum on an instant-by-instant basis, in a manner that limits interference to primary[1] users. In other words, the basic idea is this: a device first "senses" the spectrum it wishes to use and characterizes the presence, if any, of primary users. Based on that information, and regulatory policies applicable to that spectrum, the device identifies spectrum opportunities (in frequency, time, or even code), and transmits in a manner that limits (according to policy) the level of interference perceived by primary users. We term this approach *opportunistic spectrum access.*

Opportunistic spectrum access also provides far easier deployment, or rapid entry, into regions where spectrum has not been assigned. Only minimal prior coordination is necessary, greatly easing the restrictions to meet the deconflicting requirements. This is helpful both in civilian applications such as the entry of a wireless LAN technology in less developed regions, and in military operations requiring high tempo and quick reaction time.

[1] Users that are licensed to use the spectrum in question, subject to regulatory constraints.

The fundamental change from legacy systems is that the management of spectrum 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 constraining conditions will be present at any one time, and these are the only ones that need be considered in developing interference avoidance tactics. The radio itself is best positioned to be aware of these conditions.

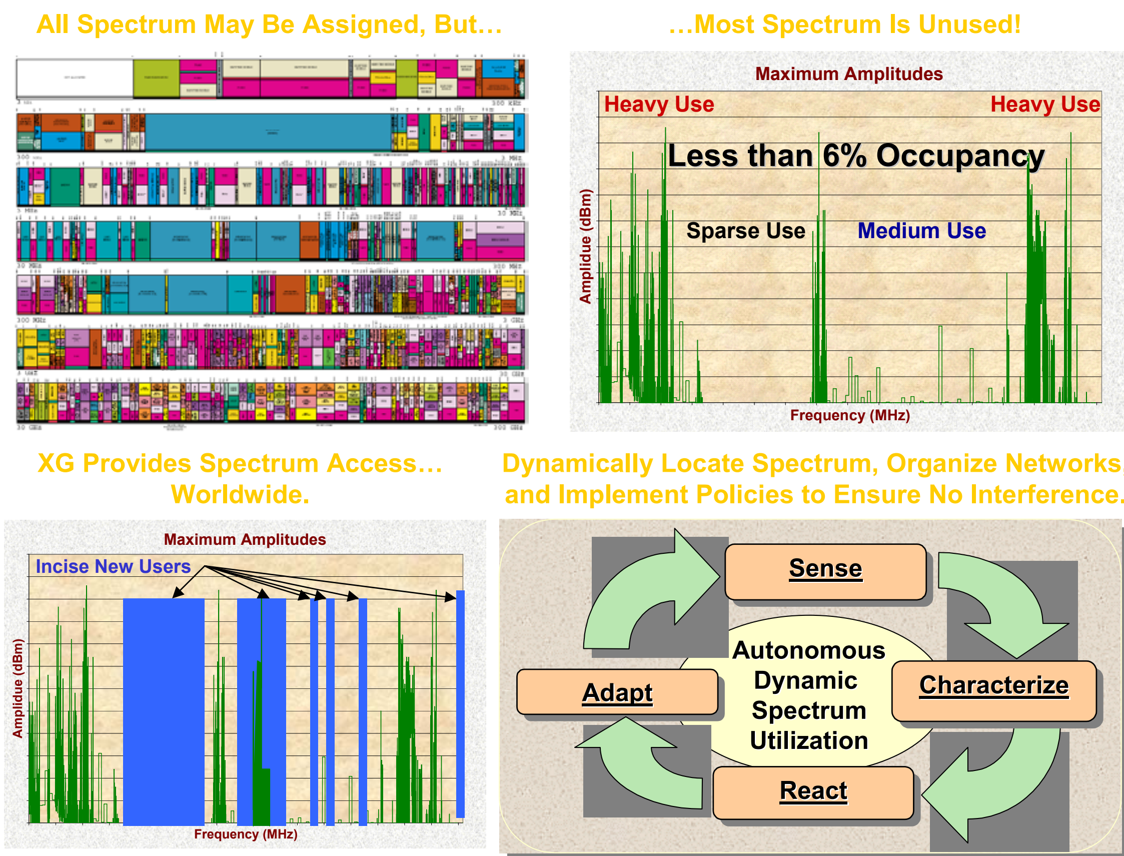

Figure 1: Spectrum is wasted. Opportunistic spectrum access can provide 10x improvement by reusing wasted spectrum.

While conceptually simple, the realization of opportunistic spectrum access is highly challenging. Several problems must be solved: sensing over a wide frequency band; identifying the presence of primaries and characterizing available opportunities; communication among devices to coordinate use of identified opportunities; and most importantly, definition and application of interference-limiting policies, and utilization of the opportunities while adhering to such policies.

Let us consider the last point further, as it is complex and significantly impacts the scope of our vision and approach. The ability to sense and transmit on unused spectrum, or *spectrum agility*, is doubtless the central capability required. However, the true potential of this new approach can be exploited only if in addition to spectrum agility, we provide *policy agility* – that is, a way by which the policies controlling the behavior can be dynamically changed. That is, policies are not embedded in the radio, but can be loaded "on-the-fly". Policy agility allows adaptation to policies changing over time and geography. Further, technology (spectrum agility) can be developed in advance of policies. This is important for breaking the chicken-and-egg dilemma that exists today, where regulatory bodies must wait for technology before drafting policies and technology must wait to see what the polices will look like.

The central role of regulatory policy in achieving opportunistic spectrum access suggests that a fresh look at the process by which policies are conceived, specified, and applied is needed. This is where XG is unique. The XG vision includes not only the development of technology for opportunistic spectrum access, but also the development of the concepts, tools and standards for incorporating a totally new "software-based" policy regime that allows policies to be decoupled from the implementation and changed dynamically.

The use of policy agility using *machine readable* or *machine understandable* policies is depicted in figure 2. Starting from the left, spectrum policies are encoded in a machine interpretable form and loaded into the XG device. The XG device then operates in accordance with its interpretation of these policies.. Policies may be loaded using smart media or over the Internet. In order to change the policies we simply need to load a new version. For instance, operating in a different country would require merely downloading from a different web-site or new smart card.

In general, if a radio were not policy agile – even if it is a software-defined radio that is capable of doing opportunistic spectrum access – then each policy change would require re-design, re-implementation, and re-accreditation. And this potentially needs to be done for each configuration of a system (e.g JTRS UAV, JTRS vehicle mount, etc.). Further, accreditation is an *m x n* problem, that is, for each of *m* policies, and each of *n* radio types/configurations, a separate accreditation needs to be done. On the other hand, with policy agile radios regulators need to accredit once (accredit based on ability to correctly interpret machine-understandable policies), and can change policy dynamically as the situation changes.

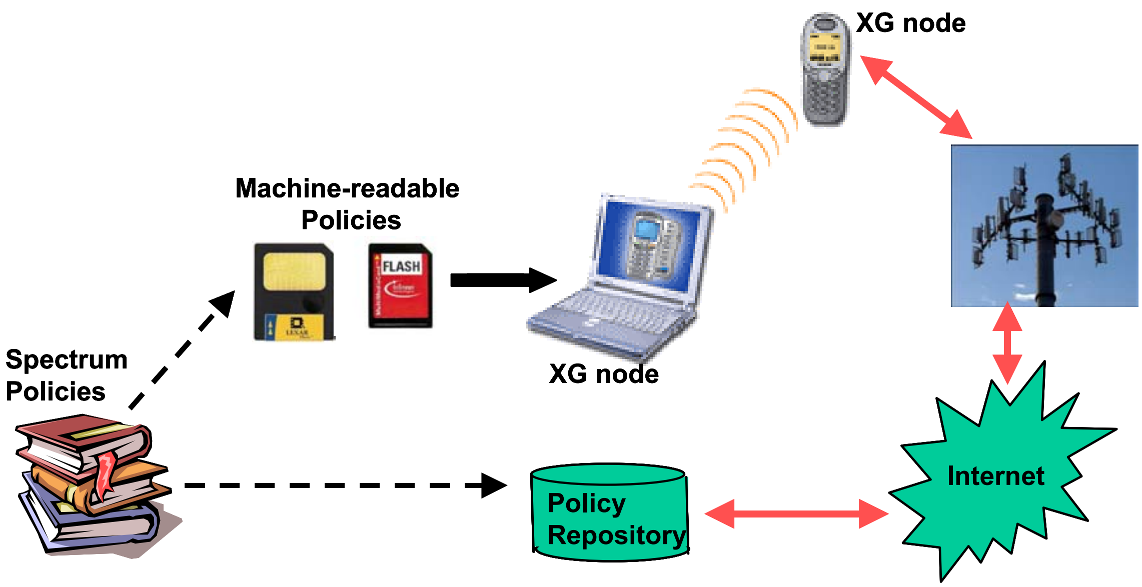

Figure 2: Machine understandable policies. With this, changing the policy merely means loading a different flash card or downloading anew.

Although recent years have seen some of the components for opportunistic spectrum access mature (e.g. software radios), we are a long way from a prototypical system. Further, no work exists in the area of decoupling the policies from the implementation. This yawning gap between current state-of-the-art and what is required for opportunistic spectrum access is the motivation behind XG. XG will develop the enabling concepts and technology for opportunistic spectrum access in an interference limiting manner. Further, it will develop an architectural framework for specifying and applying XG policies to XG devices. With these developments, XG anticipates revolutionary advances in network capacity and ease of entry while enabling a novel, highly flexible regulatory regime by providing the languages, idioms, tools, and concepts for the specification and application of machine-readable policies.

2.2 Goals

The goal of the XG program is to solve the problem of opportunistic spectrum access in its totality. At the highest level, there are two sets of goals.

1. *Develop the enabling technologies for opportunistic spectrum access*. This includes providing certain key behaviors such as sensing and characterizing the environment, identifying and distributing spectrum opportunity information, and allocating and using these opportunities commensurate with the demand. Such solutions would typically be implemented as part of an XG radio device.
2. *Develop a long-lived framework for managing the key aspects of radio behavior through flexible application of policies*. In order that the radio be policy-agile, we require a framework in which policies are written in a way that can be interpreted by the radio, and the radio is able to exploit such expression of policies.

The remainder of this section elaborates on these goals, with particular emphasis on the second goal.

For each of the behaviors mentioned in item 1 above, it is likely that there is more than one solution. Therefore, one objective is to ensure that our framework is flexible enough to permit multiple solutions within a single abstraction of how XG should operate. In other words, the framework should allow diverse solutions to co-exist while sharing a core set of behaviors.

Another goal is to keep the core behaviors distinct from the innovations that may implement the mechanisms in different ways. This would be analogous to secure kernels – that is, inside the boundary, we can be sure of what is happening and can trust it whereas outside this boundary there is room for innovation. The challenge is to make it so that only the core set of behaviors "inside the boundary" is relevant for regulatory approval. This idea will be discussed in more detail in section 3.2.

Achieving policy agility requires the decoupling of policies from behaviors and their implementations. This implies the need for a standard way of expressing policies across the XG program, and in future, across the XG-compliant nodes. Thus, a goal is to develop a language that can provide a suitable mechanism for policy expression and interpretation.

A key goal of the XG vision is *traceability,* that is, the ability to associate each emission with a policy or a set of policies that permit this emission. Traceability would be a valuable feature that will help address the thorny verification and validation problem.

Information assurance, while important, is not a focus of the XG effort. The goal is to be consistent with existing information assurance architectures (such as red/black separation), and not invalidate existing constructs. However, denial of service issues, at least on the XG control channel, must be considered. Design of protocols must pay heed to at least the well-known denial of service threats.

Finally, a challenging goal is to ensure that XG 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. In other words, we need an overarching technology that can be separated from, and be managed above the level of individual radio approvals. The longevity of the Internet Protocol is proof that such frameworks are possible.

In sum, our vision is to enable two new regimes:

- ♦ A new spectrum access behavioral regime consisting of technologies that sense, characterize, and utilize spectrum opportunities in an interference-limiting manner.
- ♦ A new regulatory control regime consisting of methods and technologies for controlling such opportunistic spectrum access behaviors in a highly flexible, traceable manner using machine understandable policies.

3 The XG Approach

A key facet of our approach is the *decoupling* of policies from behaviors and behaviors from protocols. To illustrate this, consider current practice, for instance IEEE 802.11. The 802.11 developer has the IEEE standard as the reference point for the physical and MAC layer implementation. The IEEE 802.11 standard describes the physical- and MAC-layer protocols. The *behavior* (the "what"), is implicit in, and tied to the *protocol* (the "how"). Similarly, the *policy* of sharing the unlicensed band, in particular the "good citizen rules" such as maintenance of power spectral density to within a certain value, is expressed as part of the power control and spreading mechanisms and is "embedded" within the implementation.

Our approach is to decouple these elements as shown in figure 3. By this, we mean that policies, behaviors and protocols are defined separately with some kind of "connections" defined between them. For instance, in the context of the above 802.11 example, using this approach, there would be three "specifications" – one defining the usage policy, one defining the behaviors, and one defining one or more protocols that implement each behavior. Each specification would be independently changeable within the bounds of its parent abstraction's specification..

While this might be overkill for 802.11 and other legacy "static spectrum" systems, it is of great value for

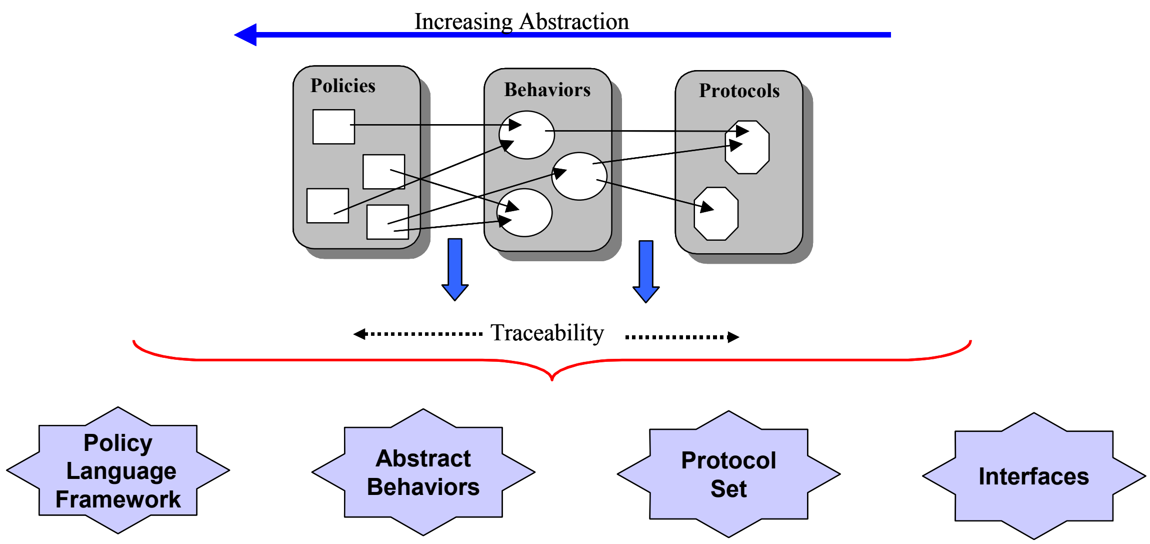

Figure 3: This illustrates the decoupling of policies, behaviors, and protocols, along with traceability and the four components of our framework.

opportunistic spectrum access. Decoupling allows adaptation to policies that vary over time and geography. Technology can be developed in advance of policies, and worldwide deployment would be greatly simplified. Furthermore, sub-policy management, as required in secondary markets, is easier. Policies no longer have to reflect the common denominator of competing technologies, and can be tailored to the diverse system capabilities expected for opportunistic spectrum access. Decoupling behaviors from protocols allows us to control *what* needs to be done separately from *how* it is implemented, resulting in a cleaner and more flexible architecture. Indeed, we argue that a decoupling approach is not just beneficial but pretty much a requirement for harnessing the full potential of opportunistic spectrum access.

Thus, a key aspect of the XG approach is that it is *policy controlled*. Policies are the only things that XG users, regulatory bodies, and foreign governments need to concern themselves with in order to control and predict the behavior of an XG system. As mentioned earlier, when an XG system is to be deployed in, say, a foreign country, we only need to supply a policy script (say in a memory card) as input to the XG system and the XG system behaves accordingly. This is similar in operational detail to the loading of a new configuration file, except that the policies are far more general and complex than the assignment of parameters to variables and the policy meta language needs far more expressive power than a typical configuration file.

We see policy as *constraining* (but not *specifying*) either the implementation details of the protocol or radio, or their performance. For example, a policy constraint could be that one must vacate an occupied frequency in *t* seconds. This policy is not dictating the performance of the protocol, or the design of the MAC, but it is constraining the solution. If for example, it takes an XG radio more than *t* seconds to acquire the signal, or to sample the band, the XG radio would not be able to abide by the stated policy, and would therefore have to avoid the band that was so constrained. In many ways, policies specify what "not to do". The XG radio and protocol implementations are free to operate in any way they want as long as they abide by policy.

A central notion in this approach, and a notion that is enabled by this approach, is that of *traceability.* Behaviors, preferably one or more of a core set of abstract behaviors, should be traceable to policies. This provides two advantages: first, it helps make the verification of new policies easier, and second, when an XG radio is deployed in a new region, it is easier to affirm that the XG system will behave in a certain way. It allows us to accredit based on ability to correctly interpret and implement policies. In other words, traceability ensures that abstract behaviors (implemented by a specific XG system) can be validated against the policy – this is key to accreditation. Lack of traceability is a weakness of today's software-defined radios. For instance, JTRS radios on UAV, handheld, and vehicle-mounted systems have to accredited independently. By being policy controlled, XG systems will allow a single accreditation to cover all cases.

The XG approach has a number of advantages in comparison to the traditional way of architecting systems (such as SUO). First, the user of XG-enabled radios has far more and far easier control of the system behavior which can be quickly adapted to the diverse environments that are likely to be encountered. Second, as mentioned earlier, it provides for traceability to help the accreditation process. Third, one can use systems based on different set of assumptions simply by incorporating these assumptions in the policy (e.g., an XG system that relies on chatty primaries for sensing can be used with a policy that allows its use only in bands with known chatty nodes). This allows incremental development by progressive relaxation of assumptions. Finally, any misgivings about the use of XG with respect to existing services within a particular band (for example when there is zero tolerance for interference) can be mitigated by simply having a policy that forbids the use of non-primary emitters in that band, rather than redesign the hardware and software to do that.

We note that an alternative approach for XG would be to develop waveforms that optimized the performance of a radio, and to embed the control protocols within the radio as part of its design. This is similar to the approach adopted for several other DARPA communications programs, such as the Small Unit Operations (SUO) radio and Future Combat System Communications (FCS-C). While this approach would demonstrate the feasibility of the XG technology, it would not provide a *framework* that would enable 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.

Our approach requires the definition of four key components, as shown in the bottom of figure 3: policy language framework, abstract behaviors, protocols, and interfaces. National and international regulatory environment is a long-term process, and one that can not instantly adapt to changing concepts and approaches. Therefore we want to approach these communities with a simplified and generic set of *abstract behaviors*. We also need to characterize the control over XG policies – thresholds and rules of operation – by a *policy language*. Policies are expressed in a machine-understandable language. Abstract behaviors are instantiated by *protocols*, the specification of which will allow for interoperability. Fourth, in order that the framework be agnostic to the "best" solution for each function, and to allow for progressive, independent refinement of each function, *interfaces* (APIs) are a key element of the framework.

In the remainder of this section, we shall look more closely at our vision for the policy language framework, and abstract behaviors. We note however, that this is only an overview, and that further detail will be presented in the other RFCs in the XG RFC family.

3.1 Policy Language Framework

The policy language framework has four objectives: developing a language structure that is rich enough to adequately express XG use cases, allow for machine "understandability", support inference and reasoning capabilities, and be flexible/extensible enough to be long-lived. We note that it is not the intent of our framework/language to be able to capture *all* spectrum policies, only *XG related* policies.
To set the context for the policy language, we first consider the "big picture" – that is, the actors and roles involved in policy based control of radios. Our vision is depicted in Figure 4.

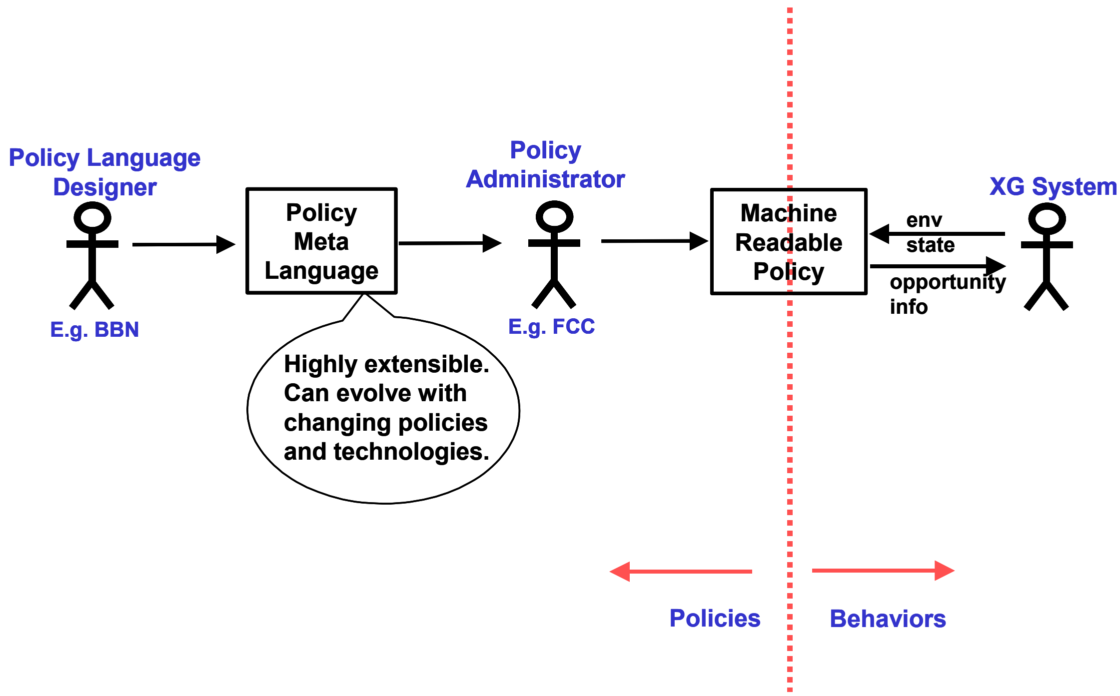

Figure 4: XG Policy Language Framework -- actors and roles

The *policy language designer* creates the language model that defines the high-level objects of the language along with the language syntax and semantics and publishes it (e.g. at a well-known URI) for the policy designers and policy users to access.

The *policy administrator* is responsible for developing and encoding spectrum policy as expressions in the policy language produced by the language designers. The policy administrator must then make the policy available to all spectrum users. The administrator does not have to know all the details of the language as they may use a graphical tool to encode the policy, called an instance editor, and hide the notational complexity of the language. The administrator also may also validate the policy set to ensure that there are no conflicts.

The *spectrum user* (an *XG system*), must then be able to use the policy to understand usage constraints (as specified by the policy administrators) on spectrum that may be available for its use. While the spectrum user needs to be able to understand the *policy* encoded in the language, the degree of language knowledge necessary to extract the policy information may vary greatly. The most limited interaction would be a *policy interface* that completely hides the language complexity, but may limit the type of information that may be extracted from the policy. The choice of the appropriate mechanism to extract information (from policy instances represented in an openly specified language representation) is therefore up to the system designer. If necessary, the policy may be transformed into a more compact format that an XG radio may use instead of the adopted standard language representation. An XG radio may also take advantage of the full power of the XG policy language and use policies encoded in it directly. In all cases, users must have the ability to access the policy and verify the XG system's conformance to the policy before using the spectrum.

For the reminder of this subsection we consider the development of the policy language. The first question is: Can it be done? After all, we are attempting to encode information that has traditionally been developed for *human* consumption. Further, there is a vast diversity in the primitive objects that make up regulatory policy domain – from concepts of frequencies to power spectral density, mathematical formulae, geographical concepts, time concepts including zoning, possible database access etc. Finally, it is not sufficient to be able to just *express* the information – it must be done in a manner that conveys the structural relationships amongst the objects so that a machine can *reason* about policies so that every single fact does not have to be encoded.

Although admittedly daunting, there are recent advances in a couple of fields that can be leveraged. First, the general area of *knowledge representation* has yielded tools, techniques and insight into representing human consumable information. Second, there has been considerable research into languages and tools for the *semantic web*, in particular, markup languages that encode the semantic content along with the data that are also relevant for our purposes.

Another observation is that most (though certainly not all) of the spectrum policy domain is hierarchical in nature. The hierarchy stems from progressive narrowing along several dimensions: from general to specific, and along geographical regions. For instance, a general policy could be “not allowed to transmit on band B”. This may be tightened as “not allowed to transmit in band B, except in frequencies *f1* and *f2*”. Policies regarding general TV bands, HDTV band(s) and ATSC lend themselves to hierarchical nesting. Similarly, national policies may be “inherited” and made more specific in certain regions (e.g. U.S states). And devices with highly directional antennas may be able to operate under a specific (less restrictive) policy where similar devices with omni-directional antennas would need to operate under a more general (more restrictive) policy.

Finally, although the current policies are largely unstructured and disorganized from an information-scientific viewpoint – as would be expected given the lack of a formal structure before policies were written – this doesn’t necessarily have to be the case with XG policies, which is the domain for our language. Establishing a framework and a language structure can help guide the specification of policies in a “cleaner” way.

These observations lead us to an approach that, at the highest level, could be termed a *structured declarative* language approach. It is declarative in the sense that the policies are expressed mainly in terms of *facts* and *rules*, and structured in the sense that the facts and rules are organized in an object-oriented fashion that exploits the power of inheritance to simplify the relationships.

Within this overall approach, it is instructive to dwell briefly on the requirements of our policy language. These include: *inheritance and polymorphism* – to enable policy rules and properties to extend others and reduce the need for enumeration; *reification (rules about rules)*, for example, to make a policy rule governing when or where a set of policies will apply; *inference*, to infer facts that may not be explicitly stated; *extensibility* in its vocabulary, structure and semantics so that the language can adapt to express new types of policies as spectrum policy requirements change; and finally, *standards based*, so that adoption is easy and tools continue to be available.

The use of policy language allows us to isolate the general framework within which XG needs to operate, from the details of establishing and advocating specific thresholds and rules of operation. We will develop this language so that we can demonstrate that it can support 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 affect the design of the underlying XG systems, as the policies would be isolated from the XG implementing behaviors.

This approach allows policies to be written in advance of technology. It also allows for lower cost XG variants that do not have the capability to enforce certain conditions, and trade cost for potential performance. If a radio has a "better" FFT, then it can play closer to the edge. We avoid a single definition of XG functionality – rather, we embrace a set of design choices that implementers can match to spectrum and policy conditions.

3.2 Abstract Behaviors

When considering policy-based control of an XG radio, we believe that control must be limited to *what* the radio does, rather than *how* the radio does it. For instance, it makes sense for a policy to say that the transmit power must be no more than 10 mW, but mandating that the radio use a particular automatic gain control (AGC) mechanism or even use AGC does not seem appropriate. In order to capture this in our framework, it is essential that we decouple the *how* from the *what* – this naturally leads to the concept of *abstract behaviors* as distinct from *protocols* that implement such behavior.

An abstract behavior is an abstraction of a mechanism that hides details of one or more aspects of its functionality. For instance, an “interference limiting” abstract behavior could be for a node to “not introduce a signal that increases the interference over any point beyond 400 meters of the node by more than 8%”. As a further example, consider that TCP is a *mechanism* which implements the transport layer *functionality* defined by the OSI reference model. TCP exhibits two *abstract behaviors*: setting up, and maintaining a connection between networked applications.

The abstraction of behaviors could be done at several levels, and so a particular protocol may correspond to several abstract behaviors. For instance, consider the IEEE 802.11 MAC Distributed Coordinated Function. A protocol for this involves specifying the frames (RTS/CTS/DATA/ACK), their formats (waveforms), timers, finite state machines, and so on. An abstract behavior might be to simply say "… use RTS/CTS/DATA/ACK handshake for collision avoidance…". This behavior might be implemented by a variety of protocols that might differ in packet format or how the NAV is handled. An even higher level abstraction might be to say "… must avoid collisions…" allowing different kinds of algorithms, including TDMA. For XG, we will choose appropriate levels on a per protocol basis, based on standardization and regulatory considerations.

An XG system is comprised of a number of mechanisms working in unison. Each of these mechanisms may be seen as solving a particular problem, which typically has more than one solution, and correspondingly, more than one protocol. Each solution has its own advantages and disadvantages, and may be appropriate for different scenarios. However, it is very difficult to regulate every possible mechanism. On the other hand, mandating a select few mechanisms stifles innovation.

How then can we allow innovation, and a variety of protocols, while assuring regulatory conformance? We propose to accomplish this by specifying a *core set* of abstract behaviors. Each mechanism within the XG system can be thought of as composed of two parts: one part that falls within regulatory purview, and one that doesn't. Typically the part that falls within regulatory purview is likely to be small and the one that doesn't is likely to be large and, moreover, is the one in which much of the scope for innovation lies. We envision that a behavior from the core set of specified abstract behaviors will be associated with each mechanism as the regulatable part, while preserving the freedom for the mechanism to make optimizations as desired in the unregulated part. Specifically, we shall first establish a set of behaviors that ensure systems are *interference limiting*.

We shall develop a framework that assures that optimizing methods lie outside of the regulatable kernel, 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.

This subset or the "core set" within the boundary will be referred to as the *regulatable kernel*. It is illustrated in figure 5, which shows four abstract behaviors A,B,C,D, and three mechanisms that incorporate these behaviors. Those portions of the XG system that are within the core set (indicated by the thick circle) should be of interest to the regulatory community, as these contain the functions that achieve the objectives subject to regulation.

By isolating this subset of the XG system, we hope to provide a focus for regulatory consideration that is compact, and does not necessitate the regulatory community becoming involved in all aspects of XG implementations.

This approach is similar to the way security systems work today. In the security realm, we have adopted methods that isolate the critical aspects of the larger system into a small, and highly controlled subset, sometimes referred to as the "trusted security kernel". 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. As long as we know, prove, and trust the mechanisms within this trusted kernel, other parts of the security system can evolve to meet different cost-performance tradeoffs. We need a similar framework for XG.

Mechanisms outside of the kernel may extend the functionality to apply it to specific situations. For instance, a regulatory kernel behavior might be to dictate an upper limit to the field intensity at any point (omnidirectionally) beyond some specified distance from the transmitter. Outside the regulatory kernel would be a behavior that assures the limit is met at each environmental system by recognizing a capability for the transmitter to know where each environmental system is and to limit the transmission in that direction so as to meet the regulatory kernel requirement at that point. This "directional sensing and control" behavior would be an extension of the "interference limiting" behavior beyond the Regulatable Kernel.

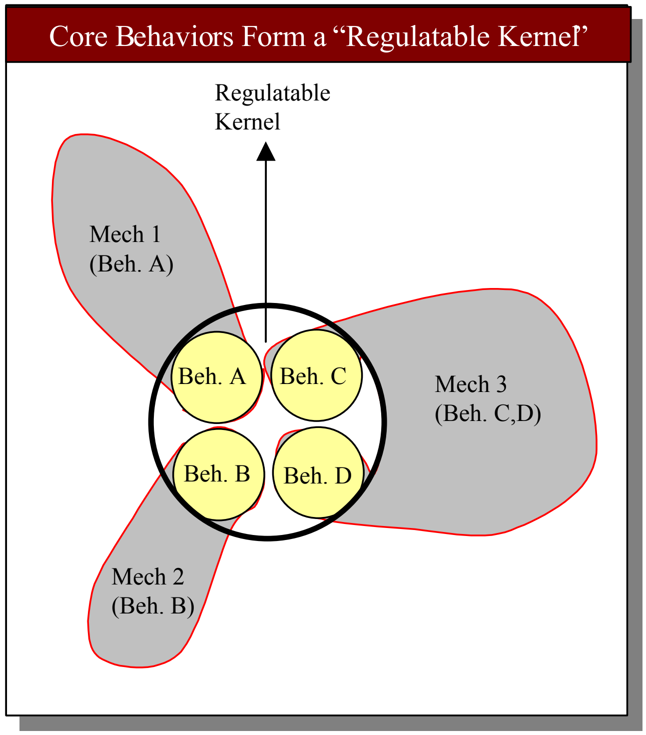

Figure 5: Core Behaviors

There are a number of challenges to be addressed in the realization of this vision. First, we need to identify a suitable set of abstract behaviors. If there are too few, it will restrict the scope of XG mechanisms. On the other hand, if there are too many behaviors, the regulatory process might become too complex. A guiding principle is to keep as little as possible within the core behaviors, subject to traceability of policies into those behaviors. Second, the level of detail of each abstract behavior must be determined. If it is too detailed, it might be difficult to obtain regulatory approval and to achieve consensus between two XG implementers. If it is too general, an approver may need to look at the implementation in order to determine whether or not it is interference preserving. Further, it will increase the risks of producing non-interoperable implementations.

Specific instantiations of abstract behaviors will permit interoperability among XG systems, and the abstract behaviors allow for fine grained optimizations within a group of XG nodes. However, correct definition and a suitably clean partitioning between core/non-core behaviors rather than interoperability is the primary goal of abstract behaviors.

The need for progressive layers of abstraction, the need to hide details, and the need to allow multiple ways of implementing a particular defined behavior strongly suggest the use of an object-oriented approach. Object-oriented design and specification are increasingly being recognized as the cleanest and most efficient way of system design and development, and use of this approach can benefit the development of XG.

In our object-oriented analysis framework, the basic XG classes have defined interfaces and high-level behaviors, and the details are left for instantiation by the behavior implementation. For example, there is an XG class that manages information about spectrum occupancy. Regardless of its implementation, we will require that it have access methods that address energy and frequency. Different implementations may 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.

Similarly, 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. Actual implementations will vary, but must be constrained to ensure compliance with the policy-based language controls. So long as the XG operation can ensure this compliance, we will leave the implementers free to develop and evolve ever more capable instantiations of these behaviors.

In many cases, the behaviors have a natural hierarchical structure. Using an object oriented framework, this structure can be elegantly represented by means of "inheritance" of classes. For instance, consider the process of identifying a set of opportunities. This may involve using only local sensing information ("uncoordinated") or be based on the sensing information from neighboring nodes ("coordinated"). The latter would be useful, for instance, if a deep fade is the reason that local sensing shows up an opportunity and in reality the opportunity is not there. For each case, one may augment the process by adding "probe" transmissions to perhaps trigger a response from a primary receiver ("active sensing"). With an object oriented framework, the top level class would be inherited in the requisite way to specify the different solutions to the same problem.

An example set of core abstract behaviors is as follows. Note that each is a class of behaviors that can be inherited and instantiated in a number of ways – details are given in the XG Abstract Behaviors RFC.

1. The *XG Spectrum Awareness Management (XG-SAM)* behavior, which describes how opportunity information is acquired, identified, represented and disseminated within and across XG systems. XG-SAM encompasses awareness information gained from sensing, policy (including configuration), and through XG Opportunity Dissemination Protocol (XG-ODP) instances, as well as opportunity identification.

2. The *XG Opportunity Dissemination Protocol (XG-ODP)* behavior, which is a class of protocol behaviors that can be used by XG systems to share opportunity awareness information. An XG system should participate in one or more instances of XG-ODP classes.

3. The *XG Usage Accounting Management (XG-UAM)* behavior, which enables every emission to be traced to a valid opportunity and a set of policy rules that allows this usage. Therefore each emission must be tagged with an opportunity object and a policy object.

3. The *XG Use Coordination Protocol (XG-UCP)* behavior, which allows XG systems to coordinate the use of selected opportunities with other (XG and non-XG) systems. XG systems should participate in one or more instances of one or more XG-UCP classes.

4 XG Operation

In this section we discuss interaction between different XG networks, layering issues involved in the dissemination of XG information, operational modes, and other concepts involved in the operation of an XG network.

Consider an XG node that enters a geographical area populated by other spectrum users. There are several possibilities for such users: they may be licensed "primaries" with right-of-use, or unlicensed. If they are unlicensed, they may or may not be XG (or more generally, opportunistic spectrum access capable) nodes. And if they are, they may or may not be able to interact with the XG network in question. Finally, there may be several levels of interaction – from just deconflicting to complete participation in spectrum allocation and use.

The following definitions are helpful in discussing such interactions. From the perspective an XG node *X*, there are three kinds of "other" nodes:

- *Non-XG.* These are "traditional" non-XG-capable nodes, or are running a different (incompatible) set of protocols. Node *X* will be able to sense their radio energy, and possibly identify physical layer waveforms (with appropriate feature detectors), but that is the limit of the interaction.
- *XG-aware.* These are nodes that implement some common protocol—that may or may not be part of the XG suite of protocols – that enables the exchange of spectrum usage information that allows deconflicting. That is, for instance, node *X* can determine which opportunities are being used by such nodes and "stay away" from those. Depending upon the nature of the common protocol, the exchange may be one-way or two-way – for example, node *X* may be able to determine the opportunities used by an XG-aware node *Y,* but

not vice-versa. However, they do not cooperate in terms of opportunity allocation/assignment, and are not required to know about which opportunities are sensed/used by each other.

- *XG-cooperative.* These are nodes with which our protocols can coordinate the use of spectrum. That is, XG-cooperative nodes implement a set of protocols that allow them to exchange opportunity information, perhaps over multiple hops of XG-cooperative nodes, and to negotiate allocation of opportunities. In other words, XG cooperative nodes sense opportunities, disseminate and receive such opportunity information from other nodes, and jointly allocate such opportunities. We also simply call these *XG nodes* and a collection of XG nodes an *XG network*.

Orthogonal to the above categories is the classification into *licensed* and *unlicensed*. Licensed nodes have right-of-use and node *X* must be interference limiting with them. Unlicensed nodes may not have such protection but node *X* is likely governed by policy that requires "good citizenship" in the sharing (a "commons" model). Thus, in our vision, there are 6 modes of possible interactions that XG will accommodate.

We now consider the *scope* of interaction between XG nodes. We first define the concept of an *XG domain,* which influences the nature and scope of our protocols. An initial vision of an XG domain is a collection of XG-cooperative nodes that form a connected (at layer 1) network. That is, there is a layer 1 path (could be multi-segment) from every node in the XG domain to every other, and each node in that path implements a common protocol for awareness, dissemination and use of opportunities.

This concept is illustrated in Figure 6. The straight lines indicate connectivity between nodes. Note that the concept of a subnet (a layer 3a functionality) is independent of the XG network. Indeed, if XG nodes from two different layer 3 subnets are within range of each other, they can belong to the same domain (beneath layer 3) crossing the subnet boundary, as illustrated in the figure. What this means is that an XG node in one subnet may know the frequencies used by an XG node in another subnet if there is a path through other XG nodes from one to the other.

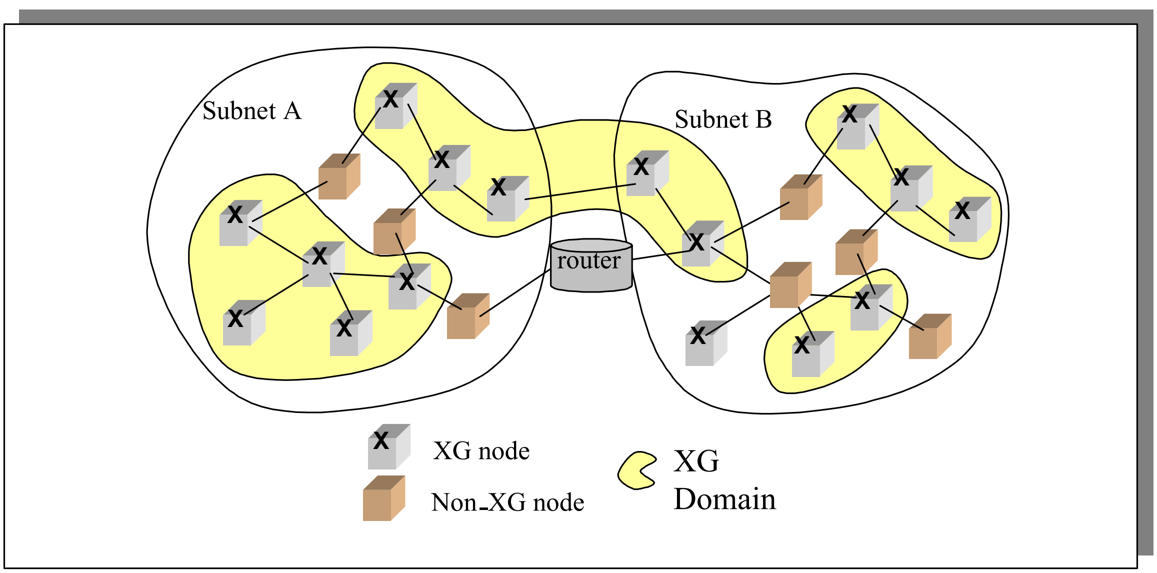

Figure 6: Concept of an XG domain. A domain is the set of nodes within a shaded region. A line between two nodes indicates the existence of a physical layer connection between the nodes. Note that network-layer connectivity may be different, and is not shown.

In this vision, an XG domain is not a packet or data centric network, but one that operates as a Layer 2 information fusion and dissemination structure. Several instantiations of this general idea are possible – from simply relaying or "flooding" the opportunity information to all nodes in a domain, to merely maintain awareness of certain content (such as spectrum usage at their location), some neighboring nodes, and local frequency planning. In general, information may be provided only on a "need to know" basis with other nodes, but reprocessing it only to provide the minimal set of awareness. For instance, information can be correlated so that each node does not report common blocked frequencies, and only candidate frequencies are relayed. This can both eliminate significant overhead bandwidth, while ensuring that enough information is available for the decision process at each node. Exchange of opportunity information across a domain should not be confused with layer 3 "routing"even if flooding is used – we are merely talking about layer 2 exchange of control packets beyond a single-hop.

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.

How are nodes in the domain organized? Broadly speaking, two opportunistic spectrum access usage models are possible, as briefly described below:

- *Centralized*. In this approach, the management of spectrum opportunities is controlled by a single entity or node, called the *bandwidth manager*, or *spectrum broker*. This node is responsible for sensing, and deciding which opportunities can be used, and by which nodes. This is relevant in real-time secondary spectrum markets.
- *Distributed*. In this approach, the interaction is "peer to peer". In other words, the XG nodes are collectively responsible for sensing and sharing the opportunities. This is most relevant in military networks.

There can also be models that share properties of both these types. Another spectrum sharing model that is related but does not quite fit in the broad classification is the "interruptible" model. In this model, spectrum is allocated subject to revocation by the primary holder of the spectrum. One may think of this as a special case of the more general XG model – there is either no primary or a primary all the time (recognized by, say, a signature "get out, I need it right away" continuous signal). Interruptible spectrum may be seen as a policy issue where the requirements for using the spectrum are very stringent.

Our XG approach is applicable to both centralized and distributed architectures. Indeed, policy may decide whether or not the system functions in a centralized or a distributed manner.

Within each approach, a number of possibilities exist. An example taxonomy is shown in Figure 7. There are three possible implementations of the centralized model: 1) there is no sensing, and the band manager has a block of spectrum that it owns that is given to secondary users in real-time based on requests; 2) only the band manager senses the spectrum, identifies opportunities dynamically and allocates it to XG nodes; 3) like (2), but additionally, XG nodes sense and provide information to the band manager to help in decisions.

The XG framework should accommodate all of these possibilities. Some of these, such as the spectrum handout approach is a special case of the general one – in this case, there is a "null" sensing and identification behavior.

5 Summary

The current policy of statically assigning spectrum for services can be inefficient in terms of spectrum utilization and cumbersome in terms of deployment agility. Opportunistic spectrum access, that is, the idea of opportunistically using assigned spectrum in an interference-limiting manner holds great promise.

The XG program is developing the concepts, framework and enabling technologies for opportunistic spectrum access. Specifically, the program has two goals: develop the enabling technologies for opportunistic spectrum access, including solutions to the problem of sensing, characterizing, identifying, distributing, and allocating spectrum opportunities; and develop a long-lived framework for managing key aspects of radio behavior through flexible application of policies. Our vision encompasses not only spectrum agility, but also policy agility – that is, the use of machine understandable policies for controlling the behavior of an XG radio.

The XG framework decouples policies, behaviors and protocols. The decoupling is enabled using two key concepts: the use of a policy language, and the definition of (core) abstract behaviors.

Spectrum policies are expressed using policy scripts based on an XG policy meta language that the XG program will define. By having policy scripts tailored to reflect national and regional considerations, considerable control can be exercised over the behavior of the XG system. Traceability of policies to behaviors is an important goal of the XG framework and will help the accreditation process.

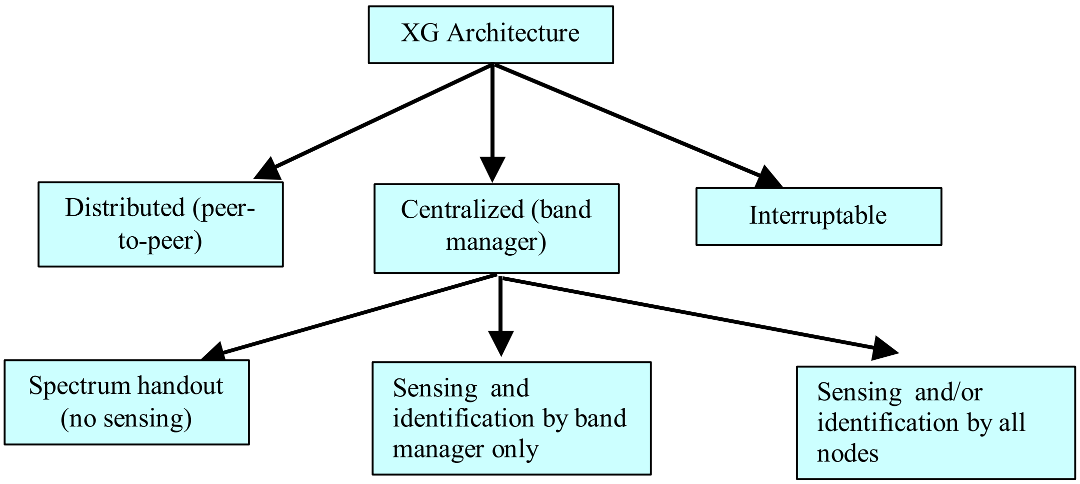

Figure 7: Taxonomy of XG approaches

A core set of interference limiting abstract behaviors will be defined. This core set is the necessary and sufficient set of mechanisms for regulatory approval. Optimizations are best left to commercial innovation arising from competition. By separating the innovations from the core set of behaviors, we enable the continued progression of XG capability and performance, without requiring that these actions be addressed within a regulatory process. The specification of the behaviors will be done at several levels of abstraction, using an object-oriented approach.

The XG vision encompasses more than a point solution to the problem of dynamic spectrum sharing. Rather, its goal is to develop a framework for diverse solutions to co-exist while sharing a core set of behaviors.

A number of network interaction and organization possibilities and spectrum management exist within the XG context. XG will accommodate interactions with non-XG, XG-aware (deconflicting use), and XG-cooperative (coordinating use) nodes. Organization models include the centralized "band manager", distributed "peer-to-peer", interruptible and others. Our goal is to accommodate as many of these diverse implementation as reasonable within the policy-centric and behavior-oriented framework.

In sum, our vision is to enable two new regimes: a new spectrum access regime consisting of technologies that sense, characterize, and utilize spectrum opportunities in an interference-limiting manner; and a new regulatory control regime consisting of methods and technologies for controlling such opportunistic spectrum access in a highly flexible, traceable manner using machine understandable policies. We shall enable the latter by defining an XG framework whose key components include the definition of a policy language and the definition of abstract behaviors.

Acknowledgments

This document was prepared by the Internetwork Research Department, BBN Technologies, with input from members of the XG working group.

Comments

Comments on this RFC should be emailed to xg-rfc-comments@bbn.com, with a carbon copy to Ram Ramanathan at ramanath@bbn.com, along with the commenter's name and organization.