WIDE AREA ADAPTIVE SPECTRUM APPLICATIONS

Daniel J. Schaefer
MITRE Corporation
Reston, VA

ABSTRACT

This paper examines spectrum opportunistic systems in which the currently assigned spectrum is monitored in real time for idle bands, specific waveform types are dynamically created to utilize the idle portions of spectrum, and those portions of the spectrum are used until the primary user needs them. This application operates over significantly expanded geographic areas without the typical structure encountered in today's cellular/Personal-Communication-System (PCS) applications. Current spectrum policy may not yet support this type of operation in the frequency bands required for the expanded coverage. However, the approach has potential for enhanced spectrum reuse together with applications for emergency and transient scenarios in which spectrum coordination may be difficult to obtain within the operational time limits. The focus of this paper is to examine the technical performance achievable by non-overlaid, loosely structured architectures that employ multi-carrier implementations to dynamically adapt to a changing spectrum environment without predefined channel structures. For the case of assigned but unused spectrum, this architecture extends to continuous use applications representative of classical reuse scenarios.

INTRODUCTION

Sharing and reuse of assigned frequency spectrum has traditionally been addressed by taking advantage of radio frequency (RF) propagation loss or overlay approaches using compatible RF waveform types to minimize the co-channel interference from multiple signals present within the same channel assignments. PCS applications may add additional time-based control to make specific channel assignments lasting only for the user's call duration. These applications employ a high degree of structure with regard to frequency assignments and network control. This highly structured approach makes it more difficult to accommodate changes in wireless usage patterns. By contrast, the growing demand for spectrum to support wireless networking and interconnect applications has led to the creation of "unlicensed bands" such as the Industrial, Scientific, and Medical (ISM) band. These unlicensed applications, operating over limited geographic area, have greater freedom to sense available channels and use them without specific formal assignment or prior coordination. However, even these systems make use of agreed-upon RF emissions limits and spectrum etiquette to ensure some level of interference compatibility.

Ideas have been suggested in several forums for fully cognitive radios capable of sensing spectrum environments and automatically adapting their operation to satisfy the radio user's communication objective while conforming to operational policy constraints. This paper examines technical considerations for a system that dynamically employs unused spectrum while minimizing interference to primary frequency spectrum users. The application is similar in concept to that of a secondary assigned user but without the formality of a specific operating frequency assignment. The adaptive type of operation optimizes usage opportunities but also creates many spectrum policy concerns. However, before issues of policy can be examined, technical performance bounds must be developed to support the formation of meaningful operations policy.

Issues associated with spectrum opportunistic systems include policy and law considerations but eventually becomes a cost-to-benefit issue. From a spectrum reuse viewpoint, are there sufficient unused spectrum opportunities to achieve significant reuse gains to justify the implementation complexities? From a more technical viewpoint, is it possible to implement unstructured systems capable of sufficient user data throughput and geographic connectivity requiring equipment complexity consistent with evolving wireless products? This paper begins the examination of these issues by assuming a "strawman" operational scenario and developing a system implementation. Performance bounds are examined to address the previously stated questions and identify key sensitivities associated with such unstructured approaches to spectrum reuse.

ARCHITECTURE CONCEPT

The viability of real-time reuse of idle frequency spectrum depends upon primary user's spectrum activity. Figure 1 is a plot of measured spectrum occupancy within an urban area collected over a 43 minute time period using an elevated antenna (50-ft) for reception. The measured data shows a significant variation in received power level combined with noticeable portions of unused spectrum over the measurement period. This spectrum includes both dynamically varying services such as point-to-point voice and other full time applications such as commercial television. Part of the spectrum includes unused television and other service. There is no intermittent channel usage for the unused full-time services. For this case, reuse can be accommodated either by adaptive approaches or using administrative reassignment of channels and use of existing wireless waveform applications.

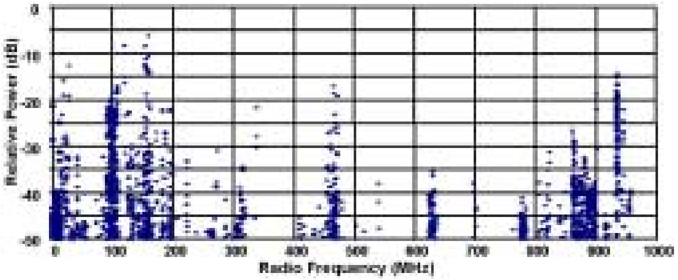

Figure 1 Representative Measured Spectrum

The viability of reuse for intermittently occupied channels depends upon the bandwidths available and the specific usage statistics. These can be characterized in terms of mean and standard deviation channel availability periods as a function of the effective contiguous bandwidth available. In general, this is highly dependent on the applications corresponding to the frequency band being used. It is also dependent upon the adaptive system's receive sensitivity, hence the antenna effectiveness. For small low-elevation antennas such as encountered in handheld applications, spectrum availability could be most advantageous. This is demonstrated within the cellular environment by the continued increase in micro and pico cell wireless environments. As the coverage area increases to include more active nodes within the same frequency range, spectrum statistics become more uniform and less advantageous to adaptive systems.

The adaptive waveform concept can be illustrated by viewing a portion of the channelized spectrum as snapshots in time. For purposes of discussion, 30 kHz channels are assumed at two separate time intervals, T1 and T2, as shown in Figure 2. The darkened areas represent primary user activity across one or more 30 kHz channels. At T1, the adaptive system has sensed the available spectrum, adapted its transmit waveform to make use of the spectrum given by the cross-hatched components, and is currently transmitting data using multiple subcarriers as indicated in Figure 2. While transmitting, the system has continued monitoring the spectrum and detected new primary user activity on channels currently used by the adaptive system. This transmission is terminated and a new waveform is used to continue transmission at time T2.

Figure 2 Waveform Adaptation Example

The adaptive waveform's occupied bandwidth and subcarrier data rates are adjusted based upon waveform synthesis criteria, offered baseband data loading, and spectrum criteria as defined prior to field operation. Each subcarrier may use one or more channels that need not be contiguous as is usually the case in classical Orthogonal Frequency Division Multiplex (OFDM) applications. In the T2 case above, some channels were omitted from the transmit waveform due to waveform optimization criteria or spectrum occupancy constraints. Though not shown, individual transmit power levels could also be adjusted to optimize the tradeoffs between data transfer performance in potential interference to primary users. The waveform continues to monitor and update its configuration as the spectrum environment, dictated by the primary assigned user, continues to change.

DESIGN AND OPERATION CONSTRAINTS

The design and implementation of the adaptive waveform system (AWS) requires tradeoffs in RF detection and sensing, digital signal processing, ad hoc networks, interference analysis and mitigation, and RF link waveform design and analysis. These activities are more manageable when constrained by a set of baseline architecture goals.

For this investigation, consider small networks with mobile operations covering not more than a 30-km radius with 10s of user nodes. The frequency spectrum for this application is intended to be within the VHF, UHF line-of-sight (LOS) ranges from 100 to 1000 MHz. The network environment, depicted in Figure 3, provides types of constraints that are sufficiently complex to require realistic design tradeoffs, yet not so complicated as to drive the design process to overly structured implementations.

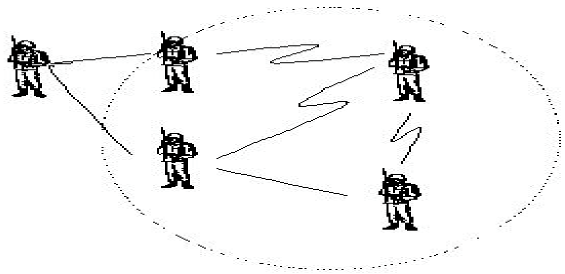

Figure 3 Representative Network Environment

The above application is representative of scenarios encountered both in military and commercial applications in which point-to-point operation is desired but not supported by existing fixed communications infrastructure. This requires the use of frequencies and propagation models requiring operation that extends model requirements beyond what is normally considered for wireless LAN and consumer equipment extension applications such as Blue Tooth-like applications. The need to extend line-of-sight leads to the use of VHF, and UHF bands to support the lengthened point-to-point connectivity. Operation within these bands is largely channelized and intermittent in nature providing a representative scenario for assessment of the adaptive architecture.

Implementation Considerations: Software Defined Radio (SDR) concepts provide the basis for implementation of an adaptive spectrum radio system. In most cases, the ability to simultaneously process large RF bandwidths is highly desirable. Unfortunately, as the instantaneous bandwidth and operational RF signal power dynamic range increase, so do the unit complexity and power requirements. The SDR implementations are largely limited by available A/D conversion technology. The preliminary design in this paper is based upon an instantaneous RF bandwidth of 50 MHz and signal power dynamic range of approximately 70 dB. This selection is consistent with reasonable cost components, yet provides sufficient bandwidth to evaluate the proposed design concepts. For a 30 kHz channelization this is equivalent to 1667 operating channels providing sufficient complexity to investigate critical implementation issues. The baseline architecture is most characteristic of mobile operation for which limited signal-to-noise ratios are the norm. Hence, the RF waveform types selected are constrained to signal constellations capable of supporting no more than 3 bits/Hz RF channel rates.

From a network viewpoint, the collection of user nodes must have a medium to discover each other's existence, coordinate link information transfer, and maintain updated pictures of the spectrum and network connectivity. In keeping with the adaptive, loosely structured concept, it is desirable not to require the dedicated control channels frequently used in many of today's wireless applications. There are, however, significant design tradeoffs to be considered in network and spectrum efficiency when choosing the amount of control structure to be used in such a system.

SYSTEM IMPLEMENTATION

A block diagram of a typical software-based AWS is shown in Figure 4. At this level, it is typical of many digital radio architectures with the possible exception of the data inverse multiplex and de-multiplex elements used to distribute user data to the individual subcarriers. This implementation may differ from classical OFDM structures in that multiple non-contiguous and non-uniform data rate sets of subcarriers are employed.

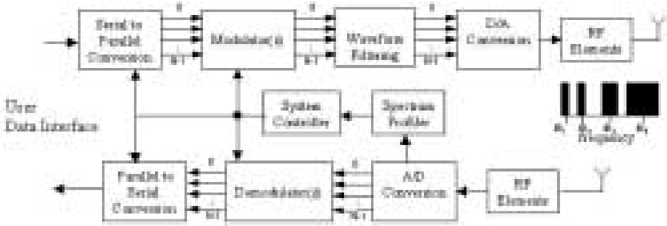

Figure 4 System Functional Diagram

The adaptive system shown in figure 4 uses digital implementations to create multiple independent transmit subcarriers to correspond to available idle spectrum opportunities. Each subcarrier employs a link-signaling rate optimized for the available RF bandwidth. Adaptable serial-to-parallel data conversion is required to distribute user data proportionally to the individual sub-carriers and recombine the data on the receive side. The RF/digital conversion implementation is broadband across the operating range (50 MHz for this example) and must address all of the typical SDR issues related to operating dynamic range, latency issues, and spectrum resolution. Spectrum profiler and system controller functions are included in the system diagram to represent the real-time network and waveform control functions. These operations control the development of the individual unit spectrum pictures, creation of the appropriate waveforms and their demodulation, and the ad hoc network control necessary to make the system work. The implementation of a real-time spectrum adaptive system requires design tradeoffs too numerous to fully describe in a short paper. However, key issues and performance considerations for the adaptive spectrum system are summarized in the remainder of the paper.

Spectrum Picture and Its Use: A key element of the adaptive spectrum system is its ability to develop and maintain an accurate description of the spectrum usage within the system's operating bandwidth. Each transceiver within the system creates an independent estimate of the real-time spectrum occupancy. These estimates may not always provide a complete spectrum of primary spectrum user activity within the radio's operational area. Hidden users, with which the adaptive system could interfere, may exist on the fringes of the individual radio's LOS reception. To address these concerns, each radio within the system makes use of available sight picture information from neighboring radio nodes as depicted in Figure 3. Each radio shares its spectrum picture with other adjacent radio nodes. Since this sharing operation is accomplished via the RF links between radios, it is important to minimize the amount of information transferred if overall reuse efficiency is to be enhanced. Sending only simple occupancy status based upon the radio channelization structure and adjusting the update rates to be consistent with the overall system adaptation rate minimizes the impact to overall performance.

The dynamic range of the spectrum picture is a key system design parameter while the system's minimum sensitivity is a secondary consideration. The spectrum picture is used to develop transmit waveform spectrum occupancy and power level information. The implementation requires only sufficient sensitivity to ensure the capture of primary user signals which potentially can be interfered with by the adaptive system's transmit effective radiated power.

RF Link Adaptation Rates: The baseline applications for the AWS involve only moderate numbers of users so that peer-to-peer network protocols may be used in support of the system's infrastructure. These protocols provide support for network discovery (participants and the addition of new nodes), sharing of spectrum picture, propagation information, and the establishment of the adaptive multiple carrier transmission links for peer-to-peer communications.

The basic AWS link operation is packet-based and uses an extension of the request to send (RTS) – clear to send (CTS) techniques. A node wishing to send data must first create a RTS query to the destination node. This requires the creation of a transmit waveform and data stream using the current spectrum picture. The RTS query includes information about the anticipated transmit waveform structure (subcarriers rates and frequencies) for the data transfer. The receive processor evaluates this information, compares the proposed waveform to its current spectrum picture, then either sends back a CTS on the current probe frequency(s) or returns a negative CTS with a requested waveform change. A limited negotiation may take place during this phase of the link establishment. Upon completion of the link setup activity, actual message data is transferred.

For the maximum of 30 km radius coverage areas identified for the example scenario, one-way propagation times do not exceed ~ 0.1 milliseconds resulting in total negotiation propagation delays of less than one-half millisecond. If we conservatively use only a single 30 kHz channel capable of supporting ~20 kbps data rate, negotiations requiring the transfer of 1000 bits per exchange will require on the order of 200 milliseconds for the actual data transfer. This establishes an upper bound on the adaptation rate of the AWS. Since the actual data transfer employs multiple subcarriers, the user information rates can be significantly higher as discussed below. Also, improved performance in the link setup times can be achieved by making use of higher bandwidth control links.

From a practical standpoint, channel idle periods fall into one of three categories: manually controlled Push-to-Talk (PTT) activities having idle periods significantly greater than one second, unused channels, or periodic waveform applications which allow prediction of available idle periods. These categories provide a range of opportunities for the AWS.

Interference Mitigation: The AWS selects available spectrum for message transmission and continually monitors the selected channels for activity by the primary user. This is accomplished by use of packet-like transmissions that allow a look-through period between packets to assess channel occupancy by the primary user. In addition, the AWS employs link power control to minimize the interference while maintaining a minimum link quality of service (QoS). Initial propagation estimates are maintained as part of the spectrum picture information and updated based upon data obtained from active information link performance.

Of some concern is the impact of packet transmissions on the primary user's idle receiver. This is primarily a concern for basic analog waveform receivers not using automatic call selection methods. Due to the sharing of analog channels among multiple services, selective calling features are common today. For digital systems co-channel interference should not be noticeable unless it significantly degrades the primary link's QoS. Hence issues regarding idle channel interference are manageable. Interference is further minimized by selection of AWS packet duration sufficiently short to minimize detection by the primary user's narrowband receiver. This approach is similar to concepts employed by impulse radio (ultra-wideband) applications.

System Performance: The AWS employs a variable length packet-like waveform that adapts using multiple subcarriers to available idle spectrum channels. The protocols used for this application are smart in that they adapt to adjust their update rates and complexity as the spectrum availability statistics dictate. This aids in optimizing system overhead for situations in which fully unused channels are available for AWS use. This allows a full range of operation from situations requiring rapid waveform adaptation to static scenarios such as the reuse of idle television channels.

Demand assigned digital radio applications require operational overhead that may be as low as 5 percent of the link rate to significantly higher percentages as the flexibility of the implementation increases. The AWS uses a variable link rate as operation changes from control-setup activities to user data transfer. To provide a basis for discussion, bound the AWS overhead factor to 20 per-cent of the total information transfer rate. The overall AWS information transfer capacity, assuming a 3 bits/Hz signaling rate, is a function of the total bandwidth used (50 MHz for this example), the AWS overhead, and the average percentage of time spectrum is available for AWS use (period spectrum is not used by the primary user). For the fixed 50 MHz example, full time capacity of the channel corresponds to an information rate of 150 Mbps if link overheads are ignored. The opportunity for the adaptive user in the LOS frequency bands to make use of the observed spectrum will vary greatly depending upon the radio's effective coverage area. As the height of the antenna is increased, the effective LOS coverage expands causing the spectrum usage statistics to become more uniform (i.e., less AWS opportunities). For portable operations with limited coverage, there are significant open frequency opportunities. For a 50 MHz spectrum segment with a 50% primary user time occupancy (see Figure 1) and a 20% AWS operational overhead, the AWS provides a total system capacity of 60 Mbps using a 3 bits/Hz waveform implementation. This conservative estimate provides an effective one Mbps data link for each of 60 users (or 100 kbps each for 600 users) within a shared coverage environment. For limited coverage portable applications, such as handheld or mobile use, the opportunities for reuse become significant.

APPLICATIONS

The AWS provides an opportunity to deploy portable applications with minimal frequency coordination requirements. This allows a rapid deployment of a communications infrastructure that may have significant local area coverage. Examples of this include both military and emergency service applications.

CONCLUSIONS AND FUTURE WORK

An overview of AWS architecture and conservative performance estimates has been presented. The proposed approach is well within current commercial technology to implement. There is however, the need to continue collection of actual spectrum usage data with sufficient time resolution to allow further substantiation of the potential for improved reuse by AWS like applications. Further refinements of protocols and extension of the concepts to address data network store and forward applications for range extension should also be considered. Lastly, once technical performance is suitably defined, spectrum law and policy must be addressed to gain acceptance to allow general fielding of AWS applications.

Acknowledgement

The work presented in this paper includes contributions from the MITRE adaptive spectrum research team members: Jeff Poston, Carl Berglund, Ferial El-Mokadem, Jim Howland, John Thweatt, Larry Thomson, and Carolyn Kenwood.

References

[1] Mitola, J., III; Maguire, G.Q., Jr., "Cognitive radio: making software radios more personal ," IEEE Personal Communications , Volume: 6 Issue: 4 , Aug. 1999, pp 13-18.

[2] Salkintzis, A.K.; Hong Nie; Mathiopoulos, P.T., "ADC and DSP challenges in the development of software radio base stations," IEEE Personal Communications , Volume: 6 Issue: 4 , Aug. 1999, pp. 47-55.