Abstract

Technology is emerging that will support the pervasive deployment of small intelligent devices that serve their owner and communicate using wireless transmissions. Collectively, these devices would provide a smart environment. Before this vision can become a reality, significant challenges must be overcome in the design of access protocols, and in spectrum management policies. Some devices in a smart environment must be able to communicate, and all devices in a smart environment must be able to coexist without excessive mutual interference. This article discusses various methods of achieving these goals, and the fundamental trade-offs involved. Some of the alternative methods include the allocation of unlicensed spectrum for this purpose, establishment of an etiquette that constrains access protocol designs, adoption of a full standard for access protocols, and the creation of a spectrum band manager.

Wireless Communications and Coexistence for Smart Environments

Jon M. Peha, Carnegie Mellon University

As consumers begin looking to their televisions for computing resources, and to the Internet for the latest music, seemingly unrelated systems are beginning to merge. Simultaneously, palmtop information devices and pagers are putting information processing in our pockets. Intelligent devices will become increasingly pervasive to form a smart environment, wherein personalized devices interact with users, sense their environment, and communicate with each other.

Each individual will have a team of information devices. A device may notify you when you get a package, or when your usual route home has become impassable because of a snowstorm, or your stock portfolio is dropping rapidly in value. Your devices will coordinate with the devices of others, so you can share an interesting spreadsheet in a meeting, and put the recipient's name and address in your rolodex, without walking across the room. Some information devices will have fixed locations, connected to televisions, phone lines, desktop computers, and stationary sensors. Other devices will be mobile, carried in cars, brief cases, and even clothing. One crucial aspect to this vision is communications between devices that must be pervasively available, and often transparent to users.

Communications resources and coexistence mechanisms to support such a system may not yet be available. To create the necessary infrastructure, regulators will have to adopt appropriate spectrum management policies; standards bodies may have to develop the technical means for interoperability; and device designers will have to develop effective communications protocols. This article discusses some of the core problems that need to be addressed to support coexistence and communications of devices in a smart environment.

The Challenges of Open Access

Such a system requires open access to spectrum, i.e., it must be possible to begin transmissions in a particular location without prior consent or licensing procedures [1]. The reasons are twofold. First, because the number of deployed devices will be large, the overhead of a licensing process would be excessive. Second, some devices will be mobile, and it is not efficient to give a mobile device exclusive rights to spectrum at every location where the device might ever reside. There are some spectrum bands in which open access is possible, and each is governed by different sharing rules, as will be discussed below. However, if smart devices become as popular as their proponents hope, additional allocations with appropriate rules will be needed to support wide-scale deployment.

There are serious challenges to overcome in creating an open access environment. First, if a band were created to effectively support devices for a smart environment, the band might be overcome by other kinds of applications, with other needs. For example, existing commercial services such as Internet access, paging, or even low-power television may choose this band to avoid bidding in expensive spectrum auctions. Since these stations transmit continually for hours, they could block out many other applications. One way to deal with this problem is to impose constraints on how spectrum will be used.

Such constraints solve some problems, but create others. Different devices will have diverse needs. Ideally, the constraints would deter unwanted applications, but would not preclude any design techniques that are useful for smart devices. This is difficult to achieve, since smart devices may transmit at different data rates, across different distances and terrain, using different modulation schemes, and they may demand different quality-of-service.

There is also the inherent problem that when any resource (such as spectrum) is shared, individual devices and systems have little incentive to conserve the shared resource. It has been shown that devices in unlicensed spectrum are likely to transmit for greater duration [2] and at greater power than is necessary, as this will advance other design goals. This phenomenon must be addressed if spectrum is to be used efficiently. The alternative is to allocate excessive spectrum so that contention is rare.

In developing spectrum management policies, standards, and access protocols, the following goals must be balanced:
- All devices should have adequate quality of service, where the definition of "adequate" may differ considerably from one application to the next.
- No device should starve, i.e., be blocked from transmission for extended periods. This is a special case of inadequate quality-of-service, but starvation deserves particular attention in an environment where one device may be allowed to

transmit indefinitely to the exclusion of others in the band.
- Policies and standards should not inhibit innovation in this rapidly changing field.
- The limited spectrum should be used efficiently, which implies high frequency reuse and a solution to the tragedy-of-the-commons problem.
- Meeting the imposed rules should not significantly increase device costs.

Rules of Coexistence

The above goals must be addressed in the definition of rules that allow smart devices to coexist, share spectrum, and possibly even interoperate. One key element in establishing these rules of coexistence, if not full interoperability, is to determine the extent of flexibility allowed in access protocol designs. Consider the following range of possibilities, from least to most restrictive:

- There are *no* constraints on spectrum use. For example, Haiti has unlicensed bands with no constraints on device operations [3]. This approach clearly is the best with respect to not hindering innovation, but it offers no quality-of-service protection and does not encourage efficient use of the spectrum. There have been serious complaints about interference in Haiti, although nothing that defied resolution. In countries like the U.S., there is less spectrum availability and a greater density of devices, so problems are likely to be more acute.
- The only constraint is a maximum power level. This approach was adopted in the U.S. in the National Information Infrastructure (NII) band. Determining the appropriate maximum power for a smart environment will be critical. A strict limit improves frequency reuse. It also precludes some applications that need greater powers; hopefully it will preclude only the applications that are not desired. It can also increase costs, as more devices must be deployed to cover a given area.
- A more detailed set of constraints is imposed on access protocols that govern how devices coexist, without requiring them to actually communicate. Device designers could use any protocol consistent with those constraints. An example is the etiquette imposed in the unlicensed personal communications services (UPCS) band, which prohibits a device from transmitting until its received power has remained below a threshold for a sufficient period of time. A well-designed etiquette can add greater protection for quality of service, prevent starvation, and encourage efficient use of spectrum [4]. (Alas, the UPCS etiquette does not have these properties [2].) It can also make some applications impractical, which can be viewed as an advantage or disadvantage. This approach also places constraints on innovation, and in some cases, may further increase costs. Determining a good etiquette for an intended application is a complex research question.
- All devices are required to follow a minimal standard. For example, all digital television transmissions in the U.S. will follow a specific standard, although some broadcasters may be transmitting high-definition television, while others are transmitting many low-definition signals. In this case, a minimal standard may allow communications. For example, all devices might communicate on an established signaling channel to gain access to a block of spectrum for a few seconds, and then each device can use that spectrum as it wishes. Such an approach may lead to better quality-of-service and more efficient use of spectrum for many applications, but it further limits the ability of designers to create new types of applications.
- A full standard is established, allowing all devices to interoperate.

Private-Sector Band Managers

There are two ways to support open access for a smart environment. One is to create an unlicensed spectrum band, where the government allows any device to transmit without permission. An alternative, which is now being considered in other spectrum bands, is for the government to license the spectrum to a band manager. The advantage is that a commercial band manager will have a financial interest in promoting both efficient use and innovation. However, new methods are needed to determine the spectrum "rents" paid to band managers that reflect spectrum consumption [1]; all else being equal, rent should be greater for high-power devices that transmit often than for low-power devices that transmit infrequently.

Moreover, unless there are to be many competing band managers, government regulators will have to find ways to prevent band managers from discriminating, so that comparable customers pay comparable rents. Otherwise, one manufacturer may pay the band manager not to accept the devices of a competing manufacturer. Most current proposals before the Federal Communications Commission lack this crucial protection.

Conclusion

The smart environment offers great promise. Much needs to be done to enable appropriate wireless communications if this vision is to become reality. While existing unlicensed bands are probably adequate for experimentation and innovation, if these applications become commonplace, a band may ultimately be needed for the smart environment in which there is no risk of contention with higher-power devices. Since these devices will probably operate at relatively short ranges, an allocation in a high-frequency band where path loss is great will help keep spectrum utilization low.

The rules of coexistence for this band must be determined carefully. Clearly, the right choice depends on multiple factors. It depends in part on the access protocols that emerge for a smart environment. If device designers can reach consensus on a single common protocol for all devices in this band, then that approach is ideal from a spectrum management perspective. However, a uniform standard can be difficult to achieve for novel services unless the number of manufacturers is small. If ample spectrum can be allocated to insure that utilization will be low, then imposing a strict maximum power limit and no other constraints is a more appropriate solution. There must be enough spectrum to serve rooms where numerous devices transmit simultaneously (and inefficiently) for extended periods. Such wireless links may connect stereos with speakers, phone lines with cordless handsets, or laptop computers with cable modems and printers. If regulators do not want to allocate excess spectrum for this purpose, an effective etiquette, tailored to suit the access protocols of smart devices, is an excellent compromise. If properly designed [4], it supports high spectrum utilization, thereby conserving this scarce resource, and it allows device designers significant latitude.

When defining rules of coexistence, the penalty for bad decisions is high. If device designers do not believe appropriate spectrum resources are available, they will not develop new products. On the other hand, if the spectrum is there but the rules of coexistence are not effective, then initial deployment may succeed, but quality will plummet as penetration

increases, users will be unhappy, and spectrum will be wasted. Finally, if initially effective but short-sighted rules are imposed, then the first generation of devices may succeed, but technical innovation in successive generations will be unnecessarily constrained.

References*

[1] J. M. Peha, "Spectrum Management Policy Options," *IEEE Communications Surveys*, vol. 1, no. 1, 4th Quarter 1998.
[2] D. P. Satapathy and J. M. Peha, "Performance of Unlicensed Devices with a Spectrum Etiquette," *IEEE GLOBECOM*, Nov. 1997, pp. 414-8.
[3] J. M. Peha, "Lessons from Haiti's Internet Development," *Communications of the ACM*, vol. 42, no. 6, June 1999, pp. 67–72.
[4] D. P. Satapathy and J. M. Peha, "Etiquette Modifications For Unlicensed Spectrum: Approach and Impact," *IEEE Vehic. Tech. Conf.*, May 1998, pp. 272–76.
*References are available at http://www.ece.cmu.edu/~peha/papers.html

Biography

JON M. PEHA (peha@ece.cmu.edu) holds a Ph.D. in electrical engineering with a minor in computer science from Stanford University, and a B.S. from Brown University. Since 1991, he has been a professor at Carnegie Mellon University, jointly in the Department of Electrical and Computer Engineering and the Department of Engineering and Public Policy. He consults for computer and telecommunications companies and government agencies in the U.S. and abroad. His work spans technical and policy issues of telecommunications and computer networks. His interests include traffic control algorithms, protocols and pricing mechanisms for ATM or IP-based integrated-services networks; protocols, architectures and spectrum management for wireless systems; and electronic commerce systems. He is currently participating in the IEEE 802.16 Wireless High-Speed Unlicensed Metropolitan Area Network project. He has served on technical staff at SRI International, AT&T Bell Laboratories, and Microsoft, and legislative staff in the House and Senate. Home page: http://www.ece.cmu.edu/~peha