Paper prepared for:
38th Research Conference on Communication, Information and Internet Policy (www.tprcweb.com), Alexandria, VA, October 1-2, 2010.

SCADA for the Rest of Us:
Unlicensed Bands Supporting Long-Range Communications

John M. Chapin[1] & William H. Lehr[2]
Massachusetts Institute of Technology

Abstract: We are on the cusp of the next big wave in Internet and wireless technology for the mass market: the emergence of ubiquitous smart environments capable of automated sensing and interaction with the physical world around us. A critical ingredient for this future is the deployment and availability of wireless control network infrastructure and services. Most discussions of the future smart environment focus on large systems managed by large corporate or public entities: smart power grids, smart freshwater and wastewater management, and so on. The distributed communications and control mechanisms that automate these infrastructure systems are called SCADA (for "Supervisory Control and Data Acquisition"). Most work on SCADA has correspondingly considered the focused networks that support a single large system or application.

This paper considers the linked technical, economic and policy issues of making SCADA available for a much larger and more diverse set of systems than have previously been discussed. Examples include the health clinic interested in at-home monitoring of rural diabetes patients; the town interested in deploying sensors to detect icing on dangerous roads; and the farmer who wants to automatically lower a retention pond's water level when high flow rates are measured miles upstream in a national forest. We describe the requirements for SCADA communications services that will support affordable automated control for a large number of small and distributed systems like these. We briefly survey a range of candidate entities that might seek to meet these requirements. We argue that achieving the desired technical and economic outcome will require both service provider and end-user deployed networks.

We propose allocation of a small unlicensed spectrum band of a new type to support growth in this direction. The difference from previous unlicensed allocations is that a few simple limitations on device behavior (a spectrum etiquette) provide predictability for long-range communications, by ensuring that spectrum congestion results in increased delay rather than reduced range. We describe an example of such an etiquette called Adaptive Duty Cycle Limit (ADCL).

[1] John Chapin, email: jchapin "at" mit.edu

[2] Bill Lehr, email: wlehr "at" mit.edu

1 Introduction

The Internet and wireless communication services have grown in recent decades into mass market infrastructures. Their on-going convergence holds the promise of a pervasive communications fabric that is always and everywhere accessible for everyone and everything that wants to communicate. With such a capability comes the prospect of widespread automation and real-time control of real-world systems, or equivalently, the cyber-real world convergence.

While that vision may seem coherent from 50,000 feet, many questions arise when considering how one might get there and what the underlying communications fabric may look like. For example, will the pervasive communications fabric consist of one or many different network architectures? In earlier work, we concluded that we expect the wired and wireless parts of the communications fabric to evolve differently. Wireless services will be much more heterogeneous than wired infrastructure. In making our arguments, we focused on fundamental and hence enduring differences between wired and wireless communications.[3]

The present paper is both an extension of and complement to the earlier work. We begin by assuming our earlier conclusion is correct. There will be multiple wireless network architectures, each optimized for a set of applications or users with common needs. This expectation encourages investigation of application types that may benefit from specialized networks. The research program is to investigate, for each set of applications with common needs: (a) whether they would benefit strongly from a specialized wireless network architecture, (b) the particular economic and technical features of the wireless network needed to provide those benefits, (c) whether the applications are sufficiently valuable to justify deploying a network with those features, and (d) the implications of all of the above for spectrum use models and spectrum policy.

In this paper, we focus on mass-market distributed communications and control of real-world systems. We call this set of applications “SCADA for the Rest of Us,” or S4U in short. Classic SCADA (“Supervisory Control and Data Acquisition”) networks are focused on the needs of electric power utilities, water management systems, chemical plants and other major infrastructure control applications. The S4U name highlights our focus on control networks to be used by small and medium businesses and local government entities (“4U”) rather than large public utilities or other government and corporate entities.

In Section 2, we identify a set of mass-market usage scenarios for SCADA. This is an incomplete list intended to suggest the existence of a large class of applications matching the service requirements described in Section 3. The combination of several key service

[3] See W. Lehr and J. Chapin, "On the convergence of wired and wireless access network architectures," *Internet Economics and Policy*, 22 (2010) 33-41; W. Lehr and J. Chapin, "Hybrid Wireless Broadband," 37th Research Conference on Communication, Information and Internet Policy (www.tprcweb.com), Arlington, VA, September 2009.

requirements—very low cost tolerance, requirement for broad area coverage, and high tolerance for communications delay—makes current mobile service provider networks and unlicensed bands suboptimal for these applications.

In Section 4, we investigate four possible trajectories for how wireless networks meeting the requirements of S4U might be introduced: by a SCADA operator, an existing mobile service provider, a new entrant service provider, or via end-user deployment. Our analysis suggests that while eventual service provider deployments are likely, end-user deployment is the only approach with high probability of jumpstarting the market.

In Section 5, we propose allocation of a small amount of a new type of unlicensed spectrum to enable end-user deployed S4U network infrastructure and the emergence of low-cost service provider networks. The problem with current unlicensed allocations is not the amount available, which is more than enough for S4U, but the model for managing shared use inherent in the current US Title 47 Part 15 framework. The key technical requirement for S4U in unlicensed spectrum is long-range communications for small low-cost devices. We outline a spectrum etiquette called Adaptive Duty Cycle Limit (ADCL) that preserves long-range communication capability even as the unlicensed band becomes congested, at the cost of increased delay which is tolerable for SCADA applications.

In Section 6, we provide a vision of what a deployed S4U communication system might look like. Section 7 concludes with a review of our principal insights and a discussion of how this relates to the larger context of spectrum management reform. This paper is not a complete policy proposal, as substantial economic and technical questions remain open, so we highlight the work required to build a case for the new unlicensed spectrum allocation.

2 Usage scenarios

Since this paper proposes spectrum policy actions to facilitate deployment of SCADA, we start with a brief review of its importance. SCADA is a critical component of the future "smart" society—smart grids, smart cities, smart highways, and so on. In this widely shared vision of the future, a broad range of social and industrial activities are partially or wholly automated in ways that improve efficiency. The drive for efficiency is critical for continued progress and growth due to the competitive pressures of globalization, ecological issues of pollution and global warming, natural resource constraints, and similar fundamental challenges. Areas where smart automation improves efficiency include dynamic and flexible management of natural resources (energy, water), public resources (transportation infrastructure, healthcare infrastructure, RF spectrum) and private resources (supply chain management). Smart automation also improves the efficiency of processes focused on meeting customer and citizen needs, including market-of-one customization, adaptive services, and local government functions.[4] A high fraction of these automation opportunities depend on wireless communications to link distributed sensors and controls.

[4] The "market of one" concept was introduced in the 1990s but is only now beginning to be realized. The inefficiencies of mass-market retailing are transformed by computing and communications

Most work on SCADA systems has focused on purpose-built systems supporting a single application such as meter reading or a single large entity such as a water system. While such work is vital, its contribution to realizing the future smart society is limited to the fraction of social and economic activity associated with large enterprises and utilities. An equally vital area for making society more efficient is to bring the benefits of SCADA to the wide spectrum of activities carried out by smaller entities, including small and medium enterprises ("SME") and local governments. The name S4U ("SCADA for the Rest of Us") refers to this set of applications.

Distributed automation and control must become available to end consumers eventually. However, we focus the current paper on SME and local governments to insulate our technical and policy proposals from issues related to home automation and in-the-home wireless connectivity.[5] Also, we are most interested in exploring the challenges of providing a service that is valued not as a consumption good ("my home is more comfortable") but as an input to economic activity ("S4U will enhance business profitability").

Within our focus on SME and local governments, there is a vast and varied set of distributed real-world applications to be automated and monitored. Rather than attempting a systematic categorization, we describe a few indicative usage scenarios.

Small and Medium Enterprise

Patient monitoring: Emergency room visits and hospital stays have been shown to be significantly reduced in lower-income elderly populations through remote monitoring and targeted interventions.[6] Health clinics trying to bring these benefits to rural populations require a communications solution that can cost-effectively reach geographically dispersed homes, many of which lack landline telephone and cellular network coverage. Specialized equipment may be installed at a fixed location for days or months, and then moved. The equipment performs monitoring of vital signs and potentially controls services such as a pill dispenser. In addition to critical-care needs, there are many possible health and wellness opportunities such as nutrition monitoring and support for efforts to improve lifestyle habits.

Lawn care optimization: A lawn care company seeking to reduce resource use and staff time may install a monitor in its customers' lawns reporting local microclimate and soil conditions on a daily basis. This would enable optimizing the number and timing of visits to

---

enabling each transaction to be specialized to the individual, for example custom jeans made to order at close to the price of mass-manufactured jeans.

[5] The home user sets up a general purpose home network that may be used for a heterogeneous mix of communications types (Internet access, home automation, connecting peripherals). In our conceptualization of S4U, the end-user—like the conventional SCADA user of today—is motivated by the desire to address a comparatively homogenous control problem at the lowest possible cost.

[6] See Britton, "The Future of Virtual Medical Care", Healthcare Information and Management Systems Society 2008 Annual Conference.

the customer site, the amount of water and fertilizer used, and other business costs, more effectively than simply remotely tracking the weather report each day. The monitor devices are moved fairly frequently as customers join and leave contracts with the company, but operate at fixed locations while deployed. The devices may communicate with inventory or with drip irrigation equipment that the lawn care company installs on site.

One could replace "lawn care" with almost any small business activity serving a distributed customer base and discover a similar story. For example, instead of a lawn care company, it could be a pest control, outside plant maintenance service provider, or waste removal company. A dry cleaner or any other business that regularly delivers goods to households could put a box on a customer's porch that automatically signals the need for pick-up and when delivery has occurred.

Local Government

The needs of various government departments are too diverse, and most government budgets too stressed, to enable buildout of a customized SCADA system serving local government as a whole. Local government is better viewed as a collection of smaller public entities, each with specific responsibilities where distributed automation can improve efficiency. Thus the usage scenario is not "local government" but rather the more focused tasks performed by individual departments or agencies. A general requirement is for a low-budget, scalable solution deployment rather than a capital intensive all-or-nothing approach.

Infrastructure monitoring: Examples include road condition monitoring to distribute salt more accurately exactly where needed in icy conditions, strain monitoring on bridges to schedule maintenance more effectively, and displacement monitoring on dams to determine when a hazardous situation has arisen. The monitors are spread throughout the geographic area and are moved from time to time as conditions and priorities change.

Flood and water resource management: Examples include monitoring stream flow rates, monitoring pond levels, and adjusting dam spillway configurations. In contrast to large enterprise water resource management, which is already supported by current SCADA systems, S4U usage focuses on cases where there are many small entities, each with responsibility for a small part of the larger picture. Examples include the retention pond behind a facility complex or a stream running through a public park. Devices may be installed in very remote locations (e.g. in a national forest upstream) and generally remain in a fixed location for months to years.

Smart parking: Sensors embedded in parking spaces that report the presence or absence of a car, networked together with smart parking meters and a central control facility, can make parking enforcement more efficient and can assist drivers in finding an available space more quickly. In this application the sensor device antenna is literally buried in the road surface, and its radio must operate on battery power for a long time.[7] Parking rates might be adjusted dynamically (mix of $1 for one hour and $0.25 for one hour meters on a street) in response to special needs.

Parolee monitoring: Counties spend a high fraction of their budgets on jail construction and operation. Releasing non-dangerous inmates to monitored parole improves their quality of life immensely and reduces cost significantly. The monitoring devices—often ankle bracelets—must periodically report relevant information such as location to parole officers. This application demands good coverage over a large geographic region such as a county and must function reliably with a low-gain antenna internal to the monitoring device.

A similar application might provide the basis for an interactive multi-player game, blurring the boundary between on-line and off-line play, or a service to allow remote monitoring of groups of kids (e.g. travel from school to after-school).

<u>Discussion</u>

The above examples are quite idiosyncratic. In some areas we are aware of deployed systems, while others are mere speculation. Many of the potential uses of SCADA will not prove beneficial for any number of reasons. Our point is that the chief reason they fail should *not* be because of a lack of appropriate wireless communications at affordable costs.

We anticipate that there are a large number of potential uses similar to those described above. In many or probably most of the potential use cases, particularly when an SME or local government entity's operations are distributed over a town, county, or larger area, the benefit of automation for one entity is insufficient to justify deploying a purpose-built wireless network for that entity. Thus resource sharing by multiple entities is an essential part of the S4U usage scenario.

A full economic analysis of S4U would have to consider its usage scenarios from two perspectives not discussed here. The first is to compare the total activity of SME and local governments—resources consumed, pollution generated, costs incurred—to the total activity of the much larger entities that are well supported by current SCADA systems. If small entities account for a small fraction of total activity, then the potential macro benefits of automation for these users are small and it would be better to focus on large entities. Our intuition is that the small entities collectively represent a fraction large enough to justify investing in, but this certainly needs to be investigated more carefully.

The other perspective missing from our discussion is that of spectrum policy economic analysis. How do the benefits of making a spectrum allocation that enable S4U compare to

---

[7] A system of this type called SFpark is deployed for 6000 parking spaces in San Francisco. Technology Review reports that each installed sensor currently costs $500 per year (Erika Jonietz, "Finding a Parking Space Could Soon Get Easier," February 8, 2010). The price needs to be substantially reduced if systems like this are to be deployed anywhere other than the most congested urban cores.

the benefits of fulfilling other claims on limited spectrum resources? Anticipating the subsequent discussion somewhat, the necessary spectrum allocation is small and represents only a minor change from existing unlicensed spectrum bands. The allocation can be used as a general purpose unlicensed band for many applications, not just for S4U. Thus the economic analysis of the S4U usage scenarios need not be as rigorous as it would be if there were an either/or choice involved.

3 Requirements analysis

Considering the broad range of usage scenarios, we have identified a communications requirement for S4U. The requirement is chosen to require the least amount of communications resources, such that the largest possible group of applications is served, under the condition that those applications are not well supported today.

Very low capital and operating cost per device: The point of S4U is to achieve large economic and social benefits through aggregating many small efficiency gains. We interpret "small" to mean that each endpoint automated in a developed country like the USA may save just $50 per year or the equivalent in resource consumption or pollution reduction. Innovation in S4U communications is not needed for systems where the economic or mission benefits of automation are much higher. Such automation opportunities can justify the cost of communication using current mobile data networks, in areas where they provide coverage, or using satellite and custom solutions in remote areas.

We think it is plausible that there are a large number of systems and activities where the benefit of automation per endpoint is significantly less than the communications price currently charged by mobile service providers to small-volume users, and that the sum of available benefits over all those systems and activities is significant. This assumption underlies the current paper.

A thorough economic analysis is required to determine the opportunities for efficiency through automating small systems and activities and hence the appropriate cost target for S4U. For the purposes of this study we have selected a strawman price of $30 for the least expensive S4U endpoint, and operations fees of either free (for some applications) or less than $24 per year (for other applications). At this price, installing an endpoint for a $50/year savings opportunity is justifiable with a reasonable time horizon.[8]

Bursty and low rate: The bulk of monitoring and control applications consist of occasional transfers of relatively small messages. Typical rates per end device range from

---

[8] The aggregate benefits of S4U-enabled automation anticipated by an SME may be significantly larger than $50 per year per endpoint (e.g., an anticipated 20% increase in business revenues or cost reduction). However, adopting and implementing the changed business practices associated with automation will likely involve other costs in addition to communications. The revenue-cost ceiling per customer account should not exceed a few percent of what the typical SME spends for computer/communications equipment and services per year.

one message per 30 seconds to one message per day. Typical message sizes are less than a few kilobytes. The traffic load per device is far lower than that supported by mobile data networks, making it plausible that a per-device charge far lower than a mobile data network could be economically sustainable if a network is optimized for that usage pattern.

Most applications involve communication in both directions. Radio communications standards for one-way transmission differ significantly from those used when both ends of a link can transmit. Therefore we omit the few unidirectional applications from the core S4U communications requirement.

A notable exception to the normal low-rate behavior is the small subset of applications that require streaming audio or video. The high rate and continuous transmission requirements of streaming media create a much different and higher wireless network load than the bursty low-rate transmissions characteristic of most of SCADA. We omit streaming media from the core S4U requirement. If a media-capable communications link is available, S4U communications can be used to schedule and control it.

Broad coverage: Coverage is critical for S4U, in two senses. One type of coverage requirement is an application where the end devices are spread over a wide area. The other coverage requirement is usage in remote locations, where the end devices are close to each other but there is currently no cost effective way of connecting them to a control system located in a central or populated location. Broad coverage coupled with the requirement for low cost translates directly into a requirement for long wireless communications links, to minimize the amount of infrastructure needed for coverage over the desired areas.

Rapidly deployable relays: Although the coverage requirements of S4U could in theory be met entirely by statically planned infrastructure, in practice the broad coverage requirement induces a corollary requirement for relays. Any commercial infrastructure planned in advance of knowing precise endpoint locations will fail to provide coverage to some of the desired locations. This occurs because terrain variation and signal obstructions and reflections make it dramatically more expensive to provide 100% coverage than to provide 90% or 95% coverage.

Mobile service provider networks can tolerate less than 100% coverage because users will walk to a window or drive out of a valley to get a usable signal. This type of mobility out of a dead spot is not an option for S4U. In many cases the location of the endpoint is fixed by the needs of the physical system being monitored or controlled, and it will often be at ground level without opportunity for connecting a tall antenna. Thus it must be possible to quickly and affordably position a relay precisely where needed to improve coverage for individual applications. Installing relays will be a normal part of many application deployments, not a rare special case.

Both battery powered and line powered devices: S4U requires a communications system that accommodates both device types. Many sensors and controllers need to be battery powered, while usage in remote locations requires the high transmit power characteristic of a line powered device. The battery powered devices are often in difficult locations so battery lifetimes of multiple months to a year are required. The combination of battery powered devices that value long lifetime and long communication links that require high transmit power is another reason that relays will often be used. A line- or solar-powered relay at the edge of a field of battery-powered sensors will be a common usage pattern in remote areas.

Fixed and nomadic operation: While high-speed mobility is not characteristic of SCADA systems, both fixed operation and frequently relocated (nomadic) operation are common. The expectation of nomadic operation, together with the focus on small systems which will have a high rate of churn in endpoint installations, implies that a S4U communications network cannot be statically planned to support well-known endpoint locations. Instead it must be adaptive to potentially significant variations in node density and usage rates across the area it covers.

Not safety critical: Few SCADA applications cause safety or life hazards if communication fails. Public safety and life critical communications require service guarantees that are generally incompatible with the low capital and operating cost goals of S4U, so we remove these uses from the core S4U requirement.

Delay tolerant: Many SCADA applications can tolerate moderate communications delays, up to tens of seconds, with minimal to no reduction in their social or economic benefits. We note that SMS text messages and satellite paging messages incur such delays with little impact on their current utility as SCADA communications links. While there are applications that require stronger timeliness guarantees, we omit them from the S4U requirement in order to enable the widest possible range of low-cost solutions.

Moderate security: Since SCADA systems interact with the external world, they are an attractive target for vandals and for those seeking to steal from the systems they control. Thus security is a requirement, defined as resistance to intrusion by unauthorized parties and resistance to denial of service. (Eavesdropping can be prevented through end-to-end encryption, thus need not be built in to the communications system.) We characterize the security requirement as “moderate” in order to omit from the core S4U requirement those few applications where there is a high incentive for experts to attack the system and immediate major loss if an attack succeeds.

3.1 Limitations of mobile data networks

Mobile service providers (MSPs) are the obvious candidate to provide wireless communications service to the S4U mass market. This section provides our analysis explaining why their current and planned mobile data networks are suboptimal solutions for the above requirements. In Section 4 we consider the potential role of MSPs in deploying new networks that are better adapted for S4U.

MSPs face two main challenges when seeking to support S4U with their mobile data networks. The first is a cost challenge. Mobile data networks were designed to provide a higher level of service than is required for S4U. Reducing per-endpoint price to the level described in the previous section is difficult due to costs inherent in the network design. Those costs are driven by the resource use and complexity of features not needed for S4U, including high-speed mobility, high-rate data connectivity, predictable low delay, and continuous connectivity for long sessions.

The other primary challenge is the broad coverage requirement for S4U, and its corollary requirement for the rapid deployment of relays. The technical design of mobile networks assumes tight control and planned management of the licensed spectrum bands they use. This makes ad-hoc deployment of powerful outdoor relays out of the question. (MSPs have only allowed end user deployed relays if the retransmitter is very low power and inside a building.) Asking the MSP itself to improve coverage or add a relay in a given area is rarely effective on short time scales, as the provider must balance the competing demands of many current and potential customers when judging where to make network investments. As a result, many S4U applications can be expected to face coverage problems.

Because of the cost and coverage challenges, much of the potential economic and social benefits of S4U will not be realized if current and planned mobile data networks are the only option available for S4U wireless communications services.

3.2 Limitations of current unlicensed bands

When a wireless application has very low cost requirements, it is natural to ask whether any of the current unlicensed bands and/or mature network technologies exploiting them can support the application.

The critical limitation of current unlicensed bands is their inability to support the S4U requirement for long-range communications. The low cost device requirement that rules out high directional antenna gain exacerbates the limitation. For these applications, 5 GHz and higher frequency unlicensed bands are unusable due to its fundamental propagation limitations. WiFi systems at 2.4 GHz can achieve moderate range, too short for most S4U applications but potentially useful for some. The multiple-mile range that used to be achievable at 2.4 GHz in rural areas has degraded substantially due to increased use of the band in recent years. Similarly the 902-928 MHz ISM band, whose fundamental propagation characteristics are good enough to support a variety of S4U applications, has become overused. For example, automatic meter reading businesses in the USA that had operated successfully in the 900 MHz unlicensed band for some years found their service failing in 2006-2007 because of the introduction of broadband access networks in the same band, which motivated them file a petition for regulatory protection (it was denied).[9] While there are advanced interference-rejection technologies available, such as beam steering multiple antenna systems, these technologies will add substantial device cost for the foreseeable future and thus cannot be relied on for extremely price-sensitive S4U endpoints.

[9] FCC 07-117, adopted June 19, 2007.

Despite the range limitation, current unlicensed bands have an important role to play in future S4U applications. A field of sensor or control points in close proximity can be easily and cheaply networked together through existing mature technologies (WiFi, Zigbee, etc.). This reduces the number of S4U endpoints that require direct access to long-range communications, and in some cases it also reduces the S4U traffic load, both of which reduce cost and improve resource utilization efficiency of the S4U network.

3.3 Potential use of TV white spaces

If unlicensed use of the TV white spaces is approved, this spectrum could potentially be useful for S4U communications. The low frequency of the TV band has excellent propagation characteristics. Use of a band with good propagation is an essential step towards achieving the long range communications required for S4U. The channels are also relatively wide compared to other opportunities in the VHF and UHF bands. Although substantial analysis would be required to develop a grounded estimate of bandwidth requirements for S4U, at first glance a single 6 MHz television channel appears sufficient to support a high level of simultaneous S4U activity in a geographic region, due to the low rate, delay tolerance and intermittent nature of S4U traffic per endpoint.

However, to support S4U effectively, the proposed TV white space rules will need to be modified in several ways.

Current restrictive transmit power and antenna height limits will need to be relaxed in order to enable longer range communication than the several miles being observed in current TVWS field trials. (Again, substantial analysis is required to determine a realistic range requirement.) One of the petitions for reconsideration pending at the US FCC suggests allowing additional power and antenna height in rural areas to reduce the cost of rural broadband coverage. If granted this change would be an excellent step towards better S4U support as well.

The requirement for spectrum sensing to detect wireless microphones should also be relaxed. The spectrum sensing requirement adds cost to endpoints already struggling to meet extremely tight cost requirements. Those endpoints will often be at disadvantaged locations with antenna constraints that make achieving the required sensitivity very difficult. Another petition for reconsideration pending at the FCC suggests allowing geolocation database only devices. This change would be beneficial for S4U support.

Most significantly, a TV white spaces channel or channels in each area must be designated as subject to an additional restriction on operation that ensures communications range is preserved as channel usage increases over time. The etiquette we propose and its justification are discussed in Section 5.

One attractive feature of the TV white spaces is the integration of central database control of channel allocation. The FCC could decide to allocate one of the TV white spaces channels for long-range communication, making it subject to an additional etiquette, at any time. Devices already in the field that do not support the new etiquette would be told that the restricted channel is not available when they next check the database, which will be within 24 hours. Different channels (or numbers of channels) could be allocated for long-range communications in different parts of the country as required to balance different policy and usage goals.

An S4U device designed to operate in such a restricted channel would continue to operate correctly if no restricted channel is available at its location. The effective range would be just as high as in a restricted channel as long as congestion is low. If no uncongested channel is available, the communications range will reduce in proportion to the congestion of the channel being used. This graceful degradation should be an attractive policy lever for the FCC. It permits an experiment with the use of the restricted channels to occur on a fairly large scale (multiple states or nationwide, lasting for several years) with reduced risk of political lock-in downstream if the experiment is terminated and the restricted channels returned to the general TV white spaces channel pool.

Another attractive aspect of the TV white spaces is the stated goal of the National Broadband Plan to promote reallocation of spectrum from TV broadcast use to other applications.[10] The innovative applications enabled by the new long range capabilities of the unlicensed spectrum allocation proposed here are fully in line with the NBP goals. It may be easier to succeed in creating the proposed unlicensed allocation in the TV band than in any other long-propagating low-frequency band, in the current policy environment.

4 Deployment scenarios

We see four potential trajectories by which communications services specialized to meet S4U requirements might be introduced. They are:

a) *SCADA operator*—a large entity that deploys a SCADA network for internal use offers access to third parties;
b) *Mobile service provider*—an existing MSP develops a S4U network;
c) *New entrant provider*—an entity other than an existing MSP enters the market with a S4U network; and
d) *End-user deployment*—users purchase radio devices as consumer electronics and deploy the wireless networks in an ad-hoc fashion.

In this section we describe each candidate trajectory in more detail and analyze the factors that affect its ability to develop a successful mass-market S4U service. The key question is not which of these entities might eventually deploy a specialized S4U wireless

[10] See http://www.broadband.gov/plan. The plan was released March 16, 2010. The plan does not explicitly call for the allocation of additional unlicensed spectrum in the TV broadcast bands, but does propose an incentive-based auction to allow these bands to be reallocated to higher-value uses. As part of this process, it may be appropriate to use some of the auction proceeds to set aside additional spectrum for unlicensed use in light of the attractive propagation properties of spectrum below 1GHz.

network, but which if any of them can reasonably be expected to be the first to deploy if given the opportunity to do so, and thus reduce risk and prove out the market for the others.

SCADA Operator

Given the high capital cost of SCADA networks, it may be attractive for a large entity such as an electric or water utility to reduce its direct cost by selling third-party access to its network. While this is a possibility, SCADA operators appear unlikely to support more than a small number of tightly controlled third party applications. By assumption, big-SCADA networks are specialized for the owner's system. Neither coverage nor capacity is allocated to support a diverse range of heterogeneous applications. There are serious impediments to investing in growing the network to better support external users. For example, security and reliability requirements are normally much higher than in S4U applications, creating endpoint costs exceeding S4U cost tolerance. The revenue generated per S4U customer is small, which makes it difficult for an organization not optimized for minimal transaction cost per customer to derive benefits from selling service. Many of the owners are public utilities with a regulated capital structure that discourages investment in risky new businesses.

It is conceivable that a large entity could enter the S4U market not as a service provider but as an anchor tenant for a third party service provider, either an existing mobile service provider or new entrant. We feel such a strategy is relatively unlikely until there is the potential for competition among third party service providers, since if there were only one provider the large entity would be creating the risk of a hold-up down the road. Thus in our view existing SCADA operators are unlikely to drive the initial deployment of S4U networks.

Mobile Service Provider

MSPs are already active in the SCADA market with two main classes of service. A low-end service often called M2M (for Machine To Machine) exploits fully depreciated 2G and 2.5G cellular networks to provide low-rate service at low cost. The higher end service uses the high-rate capability of the latest mobile networks to support data intensive applications. In the future as 4G networks are deployed, the capacity offered to the two services will increase, with 3G networks now supporting M2M service.

MSPs are highly motivated to retain their position as the primary supplier of SCADA connectivity to users who do not build private networks. If a new entrant S4U provider were to emerge that MSPs do not control, it would pose a serious competitive threat to their M2M revenue growth. A new entrant provider would also threaten their other businesses, since it would have the incentive to expand into adjacent markets just as the MSPs do. In addition to being motivated, MSPs are also well positioned to deploy a new S4U wireless network, using their existing tower sites, backhaul, and service/billing infrastructure.

However, it seems unlikely that MSPs will be the first to build out a network specialized to the S4U requirements. Key reasons include high sunk costs in existing and 4G networks, a focus on maintaining and growing M2M revenues, and a reluctance to place a big bet on a nascent low-average-revenue S4U market. MSPs are much more likely to take a wait-and-see approach, serving as much of the market as possible with M2M service over their mobile data networks, reducing prices of M2M service over time to fend off competitors, and deploying new features enabled by flexible 4G platforms to better support bursty low-rate data communications.

One attractive path for a MSP to enter the S4U market, when it does decide to do so, is to offer a single branded communications service (and to the user an apparently single endpoint device) that under the covers sometimes transfers data over the M2M network and other times transfers data over a new S4U oriented network. Such a service would launch with the wide coverage of the existing M2M service but could be priced lower, since the MSP can incrementally add S4U coverage in the areas where traffic grows. In contrast to a pure-play S4U network competitor, the higher performance and service guarantees of the M2M network could be offered as value-added services on the same hardware platform.

New entrant provider

It is difficult to imagine an entrepreneur entering the S4U wireless communications service business on a purely speculative basis. The cost of providing continuous coverage over a wide region is high, the application demand is unproven, and MSPs are certain to respond with aggressive M2M pricing. The only plausible path is an anchor tenant model where there are revenue guarantees sufficient to cover most of the network costs. As discussed earlier, large entities capable of building their own SCADA network would resist becoming the anchor tenant of the first S4U network. Thus the anchor tenant will need to be an entity incapable of building its own SCADA network but still valuing SCADA services highly.

Local governments may be a potential anchor tenant of this type. In an era of tight budgets and high bond debt, outsourcing to exploit private sector capital investment could be their only option to gain the efficiencies of SCADA for government operations. Discussions by one of us with representatives of a number of local government entities (associated with the imminent availability of TV white spaces spectrum) suggest that there is great interest in this possibility. In addition to revenue guarantees, a local government can make in-kind contributions that significantly reduce network cost for the S4U provider, such as sites for antenna mounting and access to government fiber for backhaul.

It remains to be seen whether it is possible for a public-private partnership of this type to be successfully established, what measures are needed to protect the stakeholders from the obvious risks associated with such an arrangement, and whether the offered revenue guarantees and in-kind cost reductions are sufficient to make the business case attractive for investment.

End-user deployment

The final potential trajectory for the initial buildout of S4U networks is via end-user deployment. This is analogous to the deployment of WiFi-based networking. In this approach, small business and other mass-market end-users purchase, deploy, and manage their own S4U networks. No service provider is involved in the wireless network itself, although the S4U network may connect to the Internet via a service provider wired or wireless link.

Supporting end-user deployment requires equipment that is sufficiently self-configuring to permit safe and effective deployment and operation by non-specialists. The techniques for doing this have become quite mature in the WiFi space, which seems likely to reduce the challenge for a new wireless system.

End-user deployment is an excellent fit for initial rollout of a service like S4U where the economic or mission benefit per endpoint is low. Networks are established as-needed and where-needed, rather than making any attempt to continuously cover a whole region. Thus investment is incremental, reducing risk. The investments are made by the same entity that derives the benefits of the wireless connectivity, reducing transaction costs.

Some S4U applications fit better with end-user deployment than others. Most end-user deployments can only tolerate a single wireless hop from field site (location of the system being monitored or controlled) to the control site or network access point. If the distance is too great for a single hop, small end-users rarely have the site access or other resources needed to deploy a relay in the middle of a longer wireless run. Thus a critical variable is the geographic distribution of field sites related to locations of control sites or available network access points, in comparison with the transmission range of the wireless links. A small business like a lawn care service with responsibilities scattered far beyond its facilities cannot easily exploit end-user deployment. SCADA use by such a business depends on the continuous coverage provided by a service provider. On the other hand, a farmer trying to automate systems on his or her property will be effectively supported by end-user deployment, as will a local government with facilities scattered around its area of responsibility.

4.1 Dual-path deployment

In the preceding, and albeit with insufficient detail, we have described the four most likely trajectories for initial startup of wireless services optimized for S4U. In our analysis, SCADA operators will not be interested. Mobile service providers are well-placed but will wait and see while serving as much of the market as possible with their mobile data networks. New entrant service providers will find it very tough to make the business arrangements needed to justify capital investment. In contrast, end user deployment would occur quickly and easily if it were possible.

The key question is thus whether to focus solely on enabling end user deployment or to seek to achieve a dual-path development, in which there is both end-user deployment and a service provider (either MSP or new entrant). In our view, dual-path deployment is desirable.

The service provider and end-user deployed infrastructure models have different cost dynamics. Service providers share infrastructure costs across multiple users. This is essential whenever users have geographically dispersed devices whose individual application value is not enough to justify the infrastructure on their own. Service providers also reduce deployment time for new applications and when a nomadic application is moved to a new area.

End-user deployment is essential as a bottom-up, low-capital-investment deployment for geographically non-dispersed applications. End-user deployment will jumpstart the virtuous cycle of viral adoption for novel and/or low-value applications. Service provider deployment is essential as a high-quality, continuous in space and time service for more demanding applications needs, potentially coupled with communications in licensed spectrum that offer stronger performance and service guarantees.

Additionally, the dual model provides an answer to the challenge of providing ubiquitous coverage, even to remote nodes. Having two tools in the tool box allows technology innovators (whether among service providers or equipment vendors) to choose the best mix of service-provider and end-user-deployed infrastructure for the particular task at hand. Thus, the two evolutionary paths are both substitutes (competitors that contribute to driving cost-quality improvements) and complements (optimized for different economic challenges).

Over time, we believe the bottom-up, end-user deployed model will help drive down equipment costs which will make it easier for new (potentially local or regional, or vertical-niche specialized) service providers to enter. The increase in competition should help spur further innovation and efficiency gains. Because we expect end-user innovation to be especially critical in identifying how best to use S4U to enhance business operations, providing a vector for end-users to be involved in service design from the beginning will be important.[11]

4.2 The case for unlicensed spectrum

The above discussion suggests that unlicensed spectrum will play a critical role in the deployment of S4U communications. End-user deployment is the easiest and quickest path

[11] As the growth of the Internet and mobile telephony has already taught us, the implications of these technologies for all aspects of social and economic life are profound. Automating business practices is *not* a small change and how to get it right may be expected to vary business-to-business, requiring a great deal of application (SME customer-specific) domain knowledge. Luckily, the changes wrought by the Internet and mobile telephony have made it much easier to engage end-users in service and product design.

to prove out the market and jumpstart the volume adoption curve, potentially stimulating service provider entry into the market.

Going beyond that, unlicensed spectrum will likely be a major if not the exclusive spectrum used by new service providers entering the market. The cost of licensed spectrum is high, especially if the same frequency is needed nationwide, creating major barriers to entry for service providers.

Our vision is not for a single S4U network, but many S4U networks, operating at many different scales (geographically, application niches addressed). We believe unlicensed spectrum will remain important for ensuring that on-going, end-user driven innovation and small scale (at least initially) entry remain viable.

5 A new controlled access unlicensed band

In this section we propose a new type of controlled access unlicensed band that would stimulate development of S4U communications services.

By the term *controlled access unlicensed band* we refer to the general class of unlicensed allocations made in recent years.[12] In the "original" unlicensed bands such as 902-928 MHz and 2.4 GHz, any unlicensed device can transmit at any time as long as no interference occurs to protected users. Interference between unlicensed devices is mitigated (though not eliminated) by requiring signal spreading and limiting transmitters to low power levels compared to those permitted in licensed bands. More recent unlicensed bands preserve the property that any device can use the band without exclusive license, but they have placed increasingly strict requirements on unlicensed transmitters. These requirements have been necessary to protect incumbents from interference since the new unlicensed allocations have been carved out of partially used spectrum bands. Examples of the requirements are listen-before-talk in the lower 5 GHz U-NII band and geolocation database lookup plus spectrum sensing in the TV white spaces. *Controlled access* refers to the mandatory spectrum etiquette that all devices wishing to transmit in the band must follow.

We propose that S4U can be best facilitated by a new type of controlled access unlicensed band. In the new band type, the mandatory spectrum etiquette has features to manage interference *between* unlicensed users, in addition to its features that protect co-channel or adjacent channel incumbents.

5.1 Evaluating increased regulatory control of new unlicensed bands

Regulators have historically refrained from regulating the behavior of unlicensed devices beyond the bare minimum needed to protect co-channel or adjacent channel licensed users. The standard approach in Part 15 is to state merely that unlicensed devices

[12] With the exception of UltraWideBand, which follows a radically different model.

must accept interference from any source and cause no interference to any protected system. Minimal physical layer rules such as spread-spectrum communications are established to improve utilization.

We see no fundamental reason to totally reject increased regulatory control of future unlicensed spectrum allocations, going beyond the current "free-for-all" nature of Part 15. We suggest that policy makers consider a four part test to judge whether a proposed regulation such as the etiquette suggested in this paper is appropriate for an unlicensed band.

1. There is an important policy goal that can only be achieved if all users of the band behave in a certain way.

Regulation is not required if a set of cooperating devices can achieve the goal while other devices sharing the band do not cooperate.

2. Universal and equal access to the band by all potential users should be preserved.

One of the fundamental benefits of unlicensed allocations that all potential users are on an equal footing. By not picking winners during the regulatory process, the regulator enables unforeseen innovative uses to flourish. Any increase in regulation must preserve this benefit.

3. Wireless technology innovation in the band should not be foreclosed.

The low barrier to entry in unlicensed bands has made them a locus of technology innovation. The regulation should be as technology-neutral as possible to avoid regulatory lock-in to a design that rapidly becomes obsolete.

4. The effort required to verify compliance ex ante and/or enforce compliance ex post must be reasonable given the expected benefits.

Since the unlicensed band is a limited shared resource, most regulatory restrictions of interest will limit the utilization of the resource by devices in some way. In other words, devices are likely to be able to improve their performance in some way if they violate the restriction. Verifying and/or enforcing compliance are thus critical functions. The regulations must be designed in such a way that these steps are feasible with reasonable effort and cost.

The compliance verification effort that is considered "reasonable" will be an evolving standard over time. For example, the spectrum access etiquette proposed by the FCC for the TV white spaces, which includes interaction with a remote database, is substantially more challenging to verify ex ante than the etiquettes established for earlier controlled access unlicensed bands.

5.2 An etiquette is needed to support S4U

We now review the requirements for S4U communications systems and describe why a spectrum access etiquette of the right type is essential.

On one hand, S4U requires customer-deployed infrastructure and low barriers to entry for service providers. This argues for unlicensed spectrum. On the other hand, S4U requires long-range communications with cheap, small, customer-deployed devices (e.g. devices without directional dishes) and requires predictable connections. These requirements argue against unlicensed spectrum in the current Part 15 model. Long communications range implies an even longer interference range, so the band may become saturated at a low spatial density of devices and thus operate inefficiently unless the devices are well coordinated (which normally only occurs in licensed bands where devices share common technology and administrative control). Moreover, as usage of an unlicensed band grows, the noise floor power increases and it becomes more likely that an uncoordinated transmitter is near to the receiver of a long-range link. These effects reduce the signal-to-noise ratio and hence reduce range over time. A transmitter-receiver pair that works well in year 1 may find itself unable to close the link in year 3.

The requirements conflict can be resolved with a controlled access unlicensed band whose spectrum etiquette preserves the ability of distant nodes to communicate as usage of the band increases over time. From a communications theory perspective, if range does not decrease, band congestion has to be offset in some other way. Given the delay tolerance of S4U applications, we suggest that increased delay is the appropriate tradeoff. That is, a transmitter-receiver pair that works well in year 1 will continue to be able to exchange data in year 3 despite substantial increase in usage of the band, but any given message may take longer to be sent from transmitter to receiver.

Protocols for individual communication systems that provide a congestion-delay tradeoff are well understood. For example, 802.11 WiFi uses an exponential back-off approach where a transmitter that does not see an acknowledgement indicating successful reception waits longer and longer to resend each time it tries. The end result is that all packets eventually are transmitted in exclusive timeslots, but the delay for one packet goes up as more users share the channel.

While such protocols are well understood, achieving congestion-delay tradeoff via a regulatory spectrum access etiquette is a fundamentally novel challenge compared to achieving it in a protocol. As described in the previous section, regulatory mandates for unlicensed bands must be technology-neutral and must support affordable compliance verification in an adversarial setting. The exponential back off protocol used in WiFi and many other standards is not technology-neutral. For example, it assumes that data transfers are packetized into short bursts and that the receiver acknowledges immediately after arrival of a packet. Compliance verification is also challenging (and therefore costly) due to the complexity of the behavior required to achieve the congestion-delay tradeoff.

5.3 The Adaptive Duty Cycle Limit (ADCL) etiquette

It will require careful research to enumerate and evaluate spectrum access etiquettes that provide a congestion-delay tradeoff rather than the congestion-range tradeoff characteristic of current unlicensed bands. We now describe one possible etiquette at a high level, not proposing that this is the best approach but rather to suggest that such an etiquette is technically feasible.

The etiquette is called Adaptive Duty Cycle Limit, or ADCL. Time is divided into fixed periods called frames, possibly around 10 seconds long. In the simplest form of ADCL, each device has a duty cycle limit. If the limit is for example 20%, the device must be silent for 80% of each frame. The 20% of the time in which it may transmit need not be contiguous, but there will be some minimum duration of quiet periods in order for them to count towards the 80% required silence. The duty cycle limit is adaptive over time. As the band becomes more congested, the limit is tightened. When the congestion passes, or the device moves out of the congested region, the limit is relaxed again. Different devices need not synchronize their clocks since frame start time offsets have no effect.

ADCL restricts the most intense users of the band first. Initially the duty cycle limit is 100%, meaning that any device can transmit at any time. Some users will transmit a lot of data, transmitting for most of the frame, while others will transmit only occasionally. This situation may be stable for a long time. However, if the interference between users becomes excessive, ADCL kicks in. For example, it may reduce the duty cycle limit to 80%. Systems that transmit for most of the frame will have to reduce their throughput rate, thus increasing latency, while systems that transmit only occasionally will not be affected.

Some might argue that preferentially limiting the most intense users of the band is unfair, and that instead it would be better to reduce the throughput of all users by the same fraction. Our perspective is that the latter approach leads rapidly to a tragedy of the commons, since each user will increase their base transmission rate in order to have useful capacity left after the haircut. Moreover, the S4U goal is to support a wide range of heterogeneous and often small users of the communications service, so we prefer to protect the diversity of use rather than the maximum rate for any single user.

The goal of ADCL is to give each user a high probability of finding a timeslot to transmit when no other device in the band is transmitting near the intended receiver. A high probability of being the sole transmission in the band facilitates long range communication by significantly reducing co-channel interference and permitting use of advanced equalization and other techniques.

Achieving this desired outcome requires choosing the limit correctly given the number of users contending for channel access. Too relaxed a limit for the current offered load will lead to excessive interference, while too stringent a limit will reduce utilization of the band. A secondary requirement is that the transmission times of different users need to be distributed in the frame (randomization works well).

The algorithms required to correctly choose and responsively adjust the duty cycle limit as load changes and users arrive and depart are the subject of current research. The algorithms may run locally based on a device sensing its environment, or remotely in a database that collects information from many devices and estimates the level of congestion per channel per geographic region. Both approaches have tradeoffs and the appropriate solution may be a combination of the two.

In many cases it is more desirable to apply the duty cycle limit on a per-network basis than a per-device basis. When devices in a network are time-synchronized and schedule transmissions among themselves, the network is operating as a single system. It can hog the channel by distributing its transmissions across the frame even though individual devices in the network transmit for only a very short time. ADCL is extended to cooperating networks of devices by requiring that there is no point below a certain height above average terrain where transmissions by nodes in the network can be observed (above a specified power threshold) to occupy more than the currently specified fraction of each frame. This is more complex to enforce because it requires identifying which transmissions belong to the network being investigated. If required for all networks, it would add cost to those systems that do not already coordinate their device's use of the band. The solution is a hybrid etiquette where the per-device ADCL applies to systems that do not coordinate transmissions across multiple devices, while the per-network ADCL applies to networks that do, with different numerical values for the two limits at the same time in the same region. The algorithms computing the two values are adjusted to assure equal access opportunities to the band by individual devices whether those devices are in a coordinated network or not.

Additional discussion of ADCL

This subsection discusses a few details of interest regarding ADCL, with an eye towards stimulating and shaping further policy and technical research.

The per-network duty cycle limit is specified to only apply to a limited height above average terrain. This means that advantaged locations such as mountaintops and skyscraper roofs are exempt from the limit on the totality of all incoming transmissions from a network. The HAAT cutoff is useful because it would be difficult and excessively restrictive for a network to coordinate and limit its transmissions across the wide area visible line-of-sight from these locations. The HAAT cutoff is acceptable because infrastructure deployed in such desirable locations will normally be much more capable and able to tolerate a higher price point than nodes at normal ground level. Therefore advanced interference-rejection technology such as beam-forming antennas can be used if necessary.

The HAAT cutoff leaves open an important question regarding transmissions from advantaged locations. Specifically, a transmitter at an advantaged location using beam steering or other spatial energy spreading technique would be allowed to transmit at higher than the duty cycle limit, if the beam pattern is changed rapidly between transmissions so that no point below the HAAT cutoff receives above-threshold energy for more than the duty cycle limit. This exception to the transmission limit is desirable because it permits more efficient use of relays and infrastructure at advantaged locations where long range communication is possible, without causing excessive interference to other users of the band. However, it creates a problem when independent devices share a single advantaged location. Beam steering and other energy spreading mechanisms are imperfect. An independent device off-axis to the intended beam direction but close to a transmitter can receive energy sufficient to prevent reception of a weak signal, especially if the independent device is using all of its degrees of freedom for its own networking needs and cannot steer a null towards the transmitter.

The policy and technical tradeoffs associated with transmitters at advantaged locations require further investigation. One possibility would be to enforce the duty cycle limit above the HAAT cutoff, but at a higher power threshold level than applies below the HAAT cutoff. An appropriate choice of threshold would limit the duration of local transmission backwash in each frame without limiting the duration of weaker signals arriving from distant locations. Thus relays and infrastructure at advantaged locations would have to respect the limit despite using beam steering antennas. Given the importance of long range communication in the expected S4U networks, alternative approaches may be preferred that allow higher transmission duty cycles when safe. For example, transmissions by a device above the HAAT cutoff may be exempt from the duty cycle limit in the case that beam steering in combination with physical or legal restrictions on the proximity of independent devices (e.g. private property) ensure no independent device can observe violations of the duty cycle limit. The potential for interpretation disputes is clear and thus the policy decision requires careful consideration.

Multiple users sharing an unlicensed band under ADCL can be viewed as a variant of an Aloha protocol. Aloha is the name for protocols where devices transmit at random times without coordination, in which randomly occurring collisions require retransmission of data. The maximum channel utilization of an Aloha protocol is roughly 37%. This means that the channel must be idle almost two thirds of the time or the delay will go to infinity. There is a large literature considering techniques to improve this limit. For example, synchronizing transmitters into pre-specified timeslots within the frame can double the maximum utilization to over 70%. Listen-before-talk also provides substantial benefits.

The analysis of Aloha's maximum channel utilization assumes all receivers hear all transmitters, no data is received successfully when transmissions overlap, and data is retransmitted until successfully received. None of these assumptions holds fully in a geographically distributed unlicensed band when modern robust waveforms are used and some users do not retransmit after a collision. Nevertheless, if ADCL is considered for adoption, the techniques from the Aloha literature appropriate for a regulatory spectrum etiquette (as opposed to an individual network protocol) should be evaluated and potentially incorporated to improve the achieved channel efficiency under congestion.

5.4 Evaluation of ADCL as a regulatory mandate

We now use the four-part test introduced earlier to evaluate whether ADCL is an appropriate regulation to consider for a future unlicensed spectrum allocation or a channel of the TV white spaces.

1. There is an important policy goal that can only be achieved if all users of the band behave in a certain way.

It is widely agreed that remote monitoring and control of physical systems is a vital application to improve energy and economic efficiency in the future. We have argued that a substantial part of these improvements depends on making SCADA available to small businesses and a large number of heterogeneous users—the S4U model in contrast to networks dedicated to supporting a single large system such as a wastewater utility. Further we have argued that achieving the low costs and heterogeneous application support requirements of S4U depends on long-range unlicensed communications. Long-range communications in a current Part 15 unlicensed band can only be preserved as usage grows through deploying advanced interference-rejection technologies in receivers, which adds too much cost to S4U radios. Hence an unlicensed band is needed that coordinates transmissions to reduce interference to receivers under congested conditions. Regulation is justified because all users of the band must participate in the coordination mechanism or it will not be effective at reducing interference.

2. Universal and equal access to the band by all potential users should be preserved.

The proposed ADCL etiquette allows any etiquette-compliant user to operate in the band at any time. That is, there is no lockout, not even a first-come first-served mechanism. Access to the band is equal because all users operating in a congested region are given the same duty cycle limitation, irrespective of use, user, arrival time, or technology.

3. Wireless technology innovation in the band should not be foreclosed.

Of the various mechanisms for reducing interference among multiple users of a specified channel and a particular location, random separation of the users in time appears to impose the least technology restriction on each user. Each user can use whatever transmission technology they desire during the time when they are allowed to transmit, and need not support a common waveform since no coordination is needed with other users of the band.

4. The effort required to verify compliance ex ante and/or enforce compliance ex post must be reasonable given the expected benefits.

Verification and enforcement challenges differ for the various forms of ADCL. Per-device compliance with the limit as it changes is straightforward to verify. Per-network compliance is easy to measure in the laboratory but hard to enforce in the field. Verification challenges associated with duty cycle limit computation will depend on the algorithms selected and on whether a remote database is used, so they cannot be precisely assessed at this time. Overall, assessing compliance appears roughly in line with the techniques used for the listen-before-talk etiquette adopted for the lower 5 GHz U-NII band in the USA and the geolocation database etiquette proposed for the TV whitespaces. This level of effort is reasonable given the social and economic benefits of an unlicensed band that predictably supports long-range communications for low-cost devices.

6 A vision for S4U communications systems

This section presents a vision for how communications systems may take advantage of a long-range unlicensed band to meet SCADA requirements. This is one of many possible evolutionary paths. It is described to help the reader understand how the authors envision that the proposals presented above may work out in practice.

The radios fall in three primary classes:

B for battery-powered low-cost
L for line-powered low-cost
H for line-powered high-capability

Any of these three may be deployed in one of three roles:

E for an endpoint connected to a sensor or control system in the field
R for a relay
I for an infrastructure point, i.e. a node connected to the Internet or directly to the computer system controlling the systems in the field.

Any combination of the two may arise in practice. One might think that **BI** (battery-powered infrastructure) would be strange, but any WiFi-capable smartphone running in hotspot mode performs exactly this function. An **HE** node (high-capability endpoint) would be deployed by a user with a need to control a high-value device located at a very remote location, such as a flood control dam, requiring maximum communications range. The line power in this case may be provided by a solar panel. The designation as line-powered merely means that the radio need not be optimized for minimal power consumption.

S4U systems will likely appear initially as end user deployments by users with specific problems to solve. A field of **BE** devices installed at ground level with small antennas communicate to a **LR** or possibly **BR** on a tree or a pole perhaps a half-mile to a mile away. The relay then punches the signal 5 miles to a **LI** or 10-20 miles to an **HI**. Shorter range deployments can leave out the relay, since **HI** on a good tower should be able to communicate with **BE** at ground level multiple miles away. The actual range of each of these links depends critically on the allocated frequency and power level limit of the unlicensed band.

Since range is critical for SCADA, relays are an important part of the system to be designed in to the communications standard from the beginning. Relays play another function as well. Over time as the band becomes more heavily used, delay will increase. At some point the delay tolerance of the user's application may be exceeded. In this case, the communications link between relay and infrastructure can be changed to a point-to-point microwave link in the high GHz range, providing a dramatic reduction in delay. None of the endpoint radios need to be changed, which is important since they are likely integrated into the sensors and controllers they support.

Users will likely work out bilateral agreements to support multiple different applications once coverage is initially provided in or near a given area. Viral deployment via meshing is more likely to occur in this context than it did in the WiFi environment because range is much higher, the applications are delay tolerant, and relay capability will be integrated into the communications standard from the beginning.

After the technology and market have matured somewhat through end-user deployment, service providers will decide to enter the market. This is a critical step since it will facilitate broad use of SCADA services by many heterogeneous users each of whom individually cannot justify the cost of the infrastructure needed to provide coverage for their applications. A service provider covering a region will initially build out **HI** nodes in a coarse grid pattern. The likely grid spacing is the maximum that permits service to small-antenna **BE** at 80-90% of the geographic locations in the covered area. The service provider will encourage its customers to deploy their own relays where needed to fill in holes in its coverage that impact their applications.

The service provider's differentiation from end-user deployments, justifying its connection fees, will likely be two fold. First, it will provide continuous coverage over a geographic region of interest, supporting nomadic applications. City and county government operations will value this highly. Second, the provider will likely offer guarantees on the maximum delay of any data transfer. In the early years it can accomplish this through direct backhaul connections to each of its **HI** nodes, so there is only one hop through the S4U unlicensed band in contrast to the multiple hops associated with customer-deployed relays and meshes. In later years, as congestion increases, the service provider can couple communications in the unlicensed S4U band with communications in a licensed band that it controls. To support this, it will likely offer customers endpoint devices that incorporate both S4U communications and some other radio or operating mode.

7 Conclusions and Directions for Future Research

This is a short paper for a big idea and we recognize that much remains to be worked out. Our goal is to propose a partial roadmap for delivering the same sorts of remote and automated control capabilities to mass market customers that big-SCADA users like those deploying "Smart Grids" exploit. We do not yet know what the key applications in this mass market may be, nor precisely what wireless technologies or network architectures will best address these opportunities, nor the business models that will be most successful and contribute the most toward economic growth and social welfare. We anticipate there will need to be significant experimentation and thought to refine these ideas further. However, we see a propitious window of opportunity for considering such ideas today.

First, computer scientists, network engineers, and communications futurists recognize that we are on the cusp of pervasive computing.[13] One of the most important next big things that this may deliver is the integration of the cyberworlds of the Internet and the real world to enable much more dynamic, flexible, and ubiquitous integration of electronic computing and communication capabilities into human decision-making. At this early stage, it is critically important to begin to think about what it will take for mass market realization of this vision. Serious consideration of S4U helps tee up these questions for public debate.

Second, because of growing environmental concerns and a recognition that we need to invest in new opportunities for long-run economic growth to ensure U.S. global competitiveness into the future, there is substantial interest and public commitment being focused on the design and deployment of Smart Grid technologies. Within the communities that are most closely focused on developing these technologies there is great interest in trying to identify how best to extend and integrate greater end-user (consumer) involvement in managing critical resources ranging from conservation to user-generated power. While the potential for greater end-user engagement offers great potential, it also poses significant operational challenges that need to be addressed. Further exploration of the S4U idea should contribute to those investigations.

Third, there is a prominent current effort to reform spectrum management. The FCC's National Broadband Plan[14] proposed significant initiatives in spectrum management reform, including a national goal of making an additional 500MHz of wireless spectrum available for shared commercial uses over the next ten years that is suitable for mobile and fixed broadband use. President Obama made this goal official national policy via a presidential memorandum that was issued in June.[15] The National Broadband Plan identifies the need for additional unlicensed spectrum allocations and proposes a novel auction-based approach for enabling a more efficient reallocation of over-the-air TV broadcast spectrum. The S4U proposal may be a useful contribution to this ambitious reform agenda.

The preceding explains why we believe the S4U idea is especially timely. We now discuss some of the many places where we believe further work is needed in the near term.

Market sizing and demand forecasts: We believe that there is significant potential for S4U service demand, but further thought into sizing this potential, and especially to identifying early adopter opportunities would be helpful. We believe there is a special need to better refine thinking about the relevant price points for equipment and on-going

---

[13] CITE to "Internet of Things," "convergence" and "cyber-real world integration" and related notions of what is to come with convergence of Internet and wireless mobility.

[14] Released March 16, 2010; available at http://www.broadband.gov/plan/

[15] See "Presidential Memorandum: Unleashing the Wireless Broadband Revolution," Press Release, White House, June 28, 2010 (available at: http://www.whitehouse.gov/the-press-office/presidential-memorandum-unleashing-wireless-broadband-revolution).

operation costs (including spectrum access payments to a service provider) to provide a better handle for business modeling and strategic analysis of the S4U opportunity. The cost tolerances we have identified in this paper are ad hoc and further work on what price levels (willingness-to-pay) may be tolerated and what costs (volume pricing effects) may be achievable over time is needed to allow the market potential to be appropriately sized.

Unlicensed protocol analysis and wireless architecture: we have sketched out one possible approach for using latency trade-offs as an alternative strategy for managing congestion in an unlicensed band. Other approaches may be worth considering, and any such approach will need to be evaluated for its robustness to a variety of intentional and unintentional security threats and for coexistence with alternative unlicensed protocols. Much detailed technical and business/economic modeling needs to be done to test any proposed protocol and to validate that the usage models we propose could be integrated with other wireless uses in ways that adequately protect the interests of all wireless users. Furthermore, we believe that there may be beneficial reforms to Part 15 rules which do not require us to reinvent a wholly new regulatory framework, and which in addition to enabling an S4U band, might also provide benefits for other classes of Part 15 devices. Proponents of licensed spectrum point to the benefits that exclusive property rights to use the spectrum deliver in the form of high-powered incentives for spectrum efficiency that are largely lacking with unlicensed use. If we are to adopt a new regime for a new unlicensed allocation, it would be a good idea to see if this reform opportunity might be used to address some of these more general concerns about unlicensed use as well.

Cooperative sharing and contract-based opportunities for S4U realization: While we have explained (albeit briefly), why we believe the S4U opportunity requires an allocation of new managed unlicensed spectrum and why this option is especially critical for the growth of end-user deployed infrastructure models, we have elsewhere argued that there is great potential in realizing the necessary future of more intensive spectrum sharing via cooperative sharing arrangements.[16] We fully anticipate that operators in licensed spectrum will play an important role in the markets for S4U as they evolve. This role will likely include acting as providers of backhaul and basic transport/connectivity services, providers of complementary services (e.g., high-data rate video transmissions), and as direct competitors in the markets for S4U services. Although we have argued that unlicensed spectrum is necessary for S4U, we also explained that S4U and ancillary services may not exist solely in unlicensed spectrum. We expect services to make use of either/both types simultaneously or over time.[17] Further work on business models and technologies for

---

[16] See, for example, Chapin, J. and W. Lehr, "The path to market success for dynamic spectrum access technologies," *IEEE Communications Magazine*, May 2007.

[17] A provider of S4U may use both licensed and unlicensed spectrum to support service; or an unlicensed-spectrum provider may compete with a licensed-spectrum provider (differentiating their offerings along different dimensions); or a service that starts in unlicensed spectrum may migrate to licensed spectrum over time. Any and all of such combinations should be allowed by the regulatory framework.

meeting the S4U requirements based on cooperative spectrum sharing is needed. This includes further work on novel pricing/service plans that mobile providers might make available to support latency-tolerant S4U uses. Evidence either for or against the economic viability of such cooperative or retail-price based approaches to addressing the S4U opportunity are relevant to its further evaluation.

CPE Cost Challenge: Figuring out how to progress along the learning/volume cost curve for complex new wireless devices poses a critical impediment to the adoption of novel wireless services of all types. It presents a special challenge if one's goal is to reach very low price points of the sort we think are critical to achieve the benefits of S4U. The need to address this problem has been a key driver in motivating international spectrum harmonization (which sometimes has posed a challenge to wireless innovation) and in tying agreements between mobile service providers and handset vendors (to enable subsidization of handset first-purchase costs). New technologies like software radio, interface standardization activities, and reforms in the management of intellectual property (e.g., changes in patent pooling practices) all have roles to play in changing the industry cost structure for wireless devices and may prove helpful in addressing the device cost challenge for S4U communications.

