Time-Limited Leases For Innovative Radios

Dr. John M. Chapin
Vanu, Inc.
Cambridge, Massachusetts, USA

Dr. William H. Lehr
MIT Communications Futures Program
Cambridge, Massachusetts, USA

A time-limited lease is a set of rights granted to an entity, system or device that expires after a specified duration. Leases are widely used in computer and network design. They are useful whenever it is difficult to revoke rights explicitly, such as in cases where the rights holders cannot be cost-efficiently located or contacted.

This paper analyzes ways to use the lease concept to facilitate innovation in radio devices and wireless communication. In our vision, manufacturers include in their devices a simple, secure subsystem that contains a clock and controls critical features such as transmitter power and frequency settings. The subsystem has enough computing power to validate cryptographically-signed lease extension messages. It disables specified radio features if no extension message has been received by the end of the lease period. These requirements are not onerous for the types of radios where leases would be used.

When devices provide this support, regulators may use certification leases rather than permanent grants to accelerate deployment of innovative radios. Spectrum rights holders may use leases to reduce risk in secondary spectrum market transactions. Firms collaborating in innovative wireless service business models can better retain control of their respective rights.

We investigate leases from both technical and policy perspectives and conclude that they can provide significant benefits for the commercialization and deployment of innovative radios.

Keywords—radio communications systems, cognitive radio, cetification, secondary spectrum markets, dynamic spectrum access, device security, regulatory policy

I. INTRODUCTION

Advanced radio approaches such as cognitive radio, dynamic spectrum access and secondary spectrum trading offer significant potential benefits, ranging from better spectrum efficiency and communication system performance to improved competition and innovation in wireless services. But these approaches also create new risks for many stakeholders, including regulators, spectrum rights holders, and system operators.

Time-limited leases (TLLs) are a tool that can help mitigate these risks and thereby promote deployment of innovative radios and services. TLLs are conceptually simple. They behave just like the time-out programmed into trial versions of software packages. In this case, the time limit is built into a radio device. If the time limit is reached and no extension message is received, the radio reduces its behavior as required or potentially halts transmission entirely.

TLLs facilitate radio innovation by enabling various stakeholders to better manage risk. A regulator, faced with a device too complex to test thoroughly, can certify it for sale and operation knowing that it is easy to recall if it misbehaves in the field. As long as the device behaves safely, the lease will be freely extended for additional time periods. A spectrum rights holder, faced with an offer for secondary access to their licensed spectrum, can enter into the contract knowing that the secondary user will cease operation at the end of the specified period. If the contract is renewed, the TLL will be extended for additional time.

Leases are extended by delivering a *certificate* to a device. A certificate is just a string of bits that encodes what operations are permitted (e.g. transmission at a specified power in a specified band) and provides a time limit. In most applications, certificates will be encrypted and/or cryptographically signed to assure that only the responsible authority is able to extend the lease.

Lease extension certificates will often be delivered in conjunction with a software or database update. This will occur either because it is convenient for the operator to combine the messages to reduce distribution cost, or because the authority controlling the lease has required behavioral changes as a condition for extending it.

A. Technical advantages of leases

Leases are a predictable, secure, and decentralized mechanism for limiting the potential harm to a stakeholder's interests. At the system level, leases are simpler to implement and more robust than a "kill button." A kill button is a mechanism that enables a stakeholder to proactively shut down a group of radios or some of their behavior modes. Kill buttons are feasible to implement in centralized systems, but in decentralized systems it is difficult to deliver a message to all the devices in a timely fashion. Devices may be out of contact for long periods of time, perhaps turned off, then begin operating again without hearing the kill message. With TLLs, devices automatically halt their behavior if they do not receive the approval to continue operating.

At the device level, TLLs are simple to implement. All that is required is a reliable clock and a protected execution environment for software that checks transmitter settings against a stored lease table. This simplicity is a key advantage of the approach. It means that lease support can be provided in radio devices at minimal cost. Many current radio devices can support leases with no hardware changes. Just as importantly, the simplicity of lease support means that the lease subsystem can be cost-effectively validated to a very high level of assurance. This gives confidence to various stakeholders that leases will be processed correctly, which is essential if the stakeholders are to rely on lease support to manage their risk.

Leases support a wide range of system designs and application requirements. Certificates can be sent to the device proactively, or retrieved automatically by the device as the end of the current lease approaches. Any available communications link can be used, whether secure or insecure. Very simple designs can be used for certificates and device-level lease support when application requirements are simple, while more sophisticated approaches can be built to support complex application needs.

B. Policy advantages of leases

By limiting the potential harm to a stakeholder's interests, leases enable shifting from today's dominantly *ex ante* enforcement approach to one more balanced between *ex ante* and *ex post* enforcement of those interests. This facilitates innovations where there is high perceived risk due to novelty or complexity.

TLLs can be an effective complement to more traditional regulatory and contractual mechanisms. In most cases leases will be an optional mechanism. A manufacturer can choose to apply for time-limited certification for some devices and traditional certification for others, depending on which decision makes the most economic sense. A secondary spectrum user can approach a primary rights holder with a contract that is technically enforced by leases, or with one that is not.

In cases where leases are used, they can be effectively combined with more traditional mechanisms. In the regulatory certification application, the simpler operating modes of a device could be given traditional permanent certification while the more sophisticated modes such as dynamic spectrum access are limited by leases. If the lease expires, the radio would not halt entirely but instead would be limited to its simpler operating modes.

C. Outline of the paper

Section II begins our analysis of time-limited leases with a more detailed discussion of their major potential applications in radio systems. Section III investigates the technical issues in implementing TLLs at the device and system level. Section IV considers economic and policy issues.

II. APPLICATIONS OF TIME-LIMITED LEASES

TLLs support radio system innovation by overcoming barriers to certification of advanced devices, by facilitating secondary spectrum contracts, and by enabling novel business models.

A. Device certification

Device certification is the process where a radio must be shown to comply with interference and safety regulations before sale. The certification approach currently used worldwide has worked reasonably well for decades. It has provided an effective balance between the regulatory requirement to protect against interference and harm, and the manufacturer's need to bring new products to market in a reasonable cost and time. In the current approach, the manufacturer tests a device before first sale in all operating modes, and measures the emissions to show that it never violates the applicable regulations. With this assurance, the regulator grants the manufacturer the right to manufacture, sell and operate the device as certified.

1) Certification as a barrier to innovation

The current certification approach developed within the overall 20th century technical, business and regulatory ecosystem of radiocommunications, characterized by radio systems composed of dedicated, single-purpose hardware used to support a narrow range of wireless applications. Significant changes to that ecosystem are now occurring, and device certification is emerging as a major barrier to innovation. Some of the critical changes and their interactions with device certification are as follows.

Increased device control complexity: Increased complexity is required by the increasingly complex operating environment and higher efficiency demands of the evolving marketplace. It is enabled by the low cost of modern integrated circuits and the large size of modern memories. Even a low-end device can easily have enough states and transitions to make full validation of its control behavior during exhaustive certification testing prohibitively expensive.

Many observers see this as a challenge particularly associated with software radio. However, the underlying cause is the increasing complexity of radio systems. The validation problem is the same whether the implementation strategy is software or hardware.

Novel spectrum access techniques: Exclusive spectrum licenses are just one of many strategies now being employed or considered by regulators as they respond to the scarcity of commercially allocatable spectrum. Other strategies include: listen-before-talk, controlled access bands, unlicensed bands with geographic exclusion zones, and interference temperature based access. In general these strategies require radios to sense and respond to the environment, leading to so-called *cognitive radios*. Implementing these strategies requires designers to build assumptions about the environment into the radio device. These design assumptions are difficult to validate in advance of large-scale deployment, and in any case will gradually become invalid as the external world evolves. Under the current certification approach, such problems mean that only extremely conservative assumptions can be used, significantly reducing the potential economic and spectral efficiency benefits of the novel spectrum access techniques.

Dynamic market-based regulation: The traditional command-and-control system, where the regulator makes all decisions about spectrum use, is giving way to a mix of more flexible regulatory approaches: unlicensed bands, spectrum brokers, secondary markets, and so on. Many observers feel progress in this direction is essential for increased economic efficiency and social benefit [1][2]. However, a major impediment to evolving regulatory approaches is the ongoing operation of devices that encode in their design the specific regulatory environment as of the date they were certified. In the presence of inflexible devices certified for permanent operation, regulatory change requires either replacing all those devices, which is expensive and slow, or the new rules must be designed to permit ongoing safe operation of the grandfathered devices, which significantly constrains regulatory flexibility and limits innovation.

2) TLLs as a solution

Using time-limited lease technology, a manufacturer can choose to apply for a limited-duration *certification lease* rather than a permanent grant, in situations where this is acceptable for the application and the customer of the device. Including lease technology in radio devices gives regulators high confidence that the devices will be upgraded or cease operating, in a timely fashion, if problems are detected in the field. This approach limits the harm caused by devices if design mistakes are not detected in certification tests. Therefore it enables safe certification of devices whose complexity would otherwise make it prohibitively expensive to achieve the level of assurance required for a permanent certification grant. Similarly, this approach enables certification of a device whose non-interfering operation is based on assumptions about the environment that are difficult to validate ex ante or that might change over time.

Using leases also gives regulators the ability to plan for regulatory change, such as by establishing a sunset clause for the rules in a given band. A manufacturer choosing to build a device that operates in such a band could be required to build in lease support, giving regulators high assurance that all devices operating in the band will be upgraded or withdrawn from service when the rules change. Leases are especially valuable for devices that may be deployed in a viral, decentralized or distributed manner, where there is no identifiable operator to take responsibility for enforcing the sunset clause.

Leases have two major limitations as a tool to reduce certification barriers.

First, leases can only be used to reduce device certification barriers when misbehavior for a bounded period of time is acceptable. An example is limited interference with a non-essential commercial service. There are types of harm where any period of misbehavior is unacceptable. Examples include excessive radiation levels or interference with life-critical communications. High-assurance validation to rule out these types of harm is required even when lighter-weight certification enabled by TLLs is used for the more complex operating modes of a device.

Second, leases will only limit harm if interference that occurs in the field can be traced back to the devices that caused it. Interference resulting from advanced radio devices is likely to be transient, and may be a cumulative effect of transmissions by many devices from several manufacturers. Both these effects make analysis of interference difficult and expensive.

TLLs are therefore only part of an overall set of innovations required to reduce certification barriers. Research is also required on ways to reduce the cost of high-assurance certification and the cost of investigations of reported interference.[1] Even without TLLs, radios are already complex enough that current certification tests audit device behavior rather than exhaustively validating all operational corner cases. Research on validation methods and investigation of interference will therefore have significant benefits whether or not certification leases are adopted.

In summary, the adoption of TLL functionality does not replace traditional certification but complements it. TLLs allow certification of more complex devices at lower total cost. TLLs reduce the ex ante certification expenses for these devices because a lower level of testing and analysis is required. TLLs also reduce the cost to society by limiting the potential harm if a radio fails to operate as anticipated. Reducing the cost threshold allows innovative radio systems to be commercialized more quickly than without TLLs.

B. Secondary Spectrum Markets

Secondary spectrum markets enable trading or sharing of spectrum access rights between primary rights holders, who hold licenses from the regulatory authority, and secondary users. Secondary markets may arise in a variety of contexts and forms, including non-cooperative or cooperative trading of primary or secondary access rights [3].

Secondary markets have been established within the last decade in multiple countries. Where appropriate, market mechanisms are expected to allocate scarce spectrum rights more effectively than command-and-control by the regulator. However, the markets have remained largely moribund with few transactions occurring. There are multiple reasons for this, including high transaction costs and technical challenges [3]. One apparent barrier is the risk perceived by the primary rights holder that the secondary user will violate the terms of the transaction.

Time-limited lease support in radio devices can be used to reduce the risk of some types of violations. The devices are configured to accept certificates only if signed by the primary rights holder. When this is done, leases can easily assure that the secondary user ceases operating in the specified band at the agreed end of the lease, never transmits beyond power level limits, and only transmits at approved times of day. The secondary user can also be limited to specified geographic areas if the lease mechanism includes GPS or some other location sensing mechanism.

By reducing risk and protecting the interests of primary rights holders, leases lower the barrier to secondary spectrum transactions and hence facilitate growth towards better spectrum allocation and higher overall economic efficiency.

[1] For example, dynamic spectrum access devices could be required to keep a log of recent spectrum access decisions and respond to authorized over-the-air requests for the log. These features would enable investigators to sample the devices in an area where interference is reported and quickly zero in on candidates for detailed analysis.

C. *Novel Business Models*

Leases support new business models in two ways. There are direct effects where TLLs are used by collaborating parties to manage risk. There are also indirect effects due to the uses of TLLs described in the previous two sections.

1) *Direct effects of TLLs*

TLL functionality can be exploited to enable novel services or alternative ways of deploying existing services. The following list is merely indicative. Many other approaches remain to be explored.

Pre-paid radio services: Pre-paid cellular services already exploit time-limited lease behavior. If lease support were built into a range of devices, this business model could easily be extended to other contexts, including ones where entities with an incentive to acquire free service have the ability to the modify the radio's software. (Section III describes radio designs that can ensure leases are processed properly despite software modifications.)

Product line harmonization: Similarly to pre-paid radio services, TLL support gives manufacturers a way to limit the behavior of a device that is not vulnerable to software modifications by the owner. Therefore the manufacturer can ship common hardware for multiple products and charge different amounts based on what functionality is enabled.

Disposable radios: Radios with time-limited, non-renewable leases would have a finite and pre-determined life. This enables for example the safe use of extremely aggressive spectrum access etiquettes in a device attached to a fire extinguisher that activates only when the extinguisher is operated.

Cooperative radio meshes: Lease renewal may be used to enforce cooperation in a distributed radio network. For example, nodes can be rewarded or penalized through receiving more or less capable leases based on their contribution to the overall network (e.g., retransmission of other nodes' packets).

Self-enforcing distributed contracts: The lease table may have multiple entries referring to the same band. If any one of them expires the lease subsystem shuts off access to that band. Assuming different signatures are required to update the different entries, multiple collaborating businesses can each have partial or full veto power over the radio's operation. With proper design, the lease renewal mechanism can even be used to implement voting or veto control of individual or system operation. By enabling lower-cost options for distributed contract enforcement, TLLs may be especially well-suited for use in unlicensed spectrum or viral networks where the tolerance for transaction costs is low.

While it is not clear which if any of these approaches will lead to successful businesses, TLLs promote experimentation by the market and expand the range of feasible business models.

2) *Indirect effects of TLLs*

As discussed earlier, TLLs can make sophisticated or complex radios easier to certify, and they can facilitate secondary or dynamic spectrum access. These benefits have the indirect effect of promoting a range of new business models that would otherwise be slower to emerge.

Faster deployment of advanced radios with frequency agility and with cognitive radio capabilities promotes end-user-initiated self-configuring, ad hoc wireless networks. Faster deployment of dynamic spectrum access reduces the cost of wireless market entry and thereby facilitates innovative services [3].

New types of dynamic spectrum access and active secondary spectrum markets, facilitated by the certification and spectrum transaction benefits of TLLs, lead directly to new operator business models. Today, we have Mobile Virtual Network Operators (MVNOs) that resell wholesale services leased from facilities-based mobile network operators. In the future, we could see new kinds of network operators, perhaps called Mobile Virtual Service Operators (MVSOs). The defining characteristic of an MVSO is that it acquires spectrum access on its own, separate from or in addition to the long-term exclusive rights held by the facilities-based operator from which it leases other services. At the same time, facilities-based operators themselves may exploit advanced spectrum access to reduce the high upfront costs of increasing their capacity or introducing new services.

Following the model pioneered by 802.11 WLAN, equiment vendors may sell dynamic spectrum access radios to end-users who then communicate in mesh or ad hoc networks. In this case TLLs facilitate the vendor's use of more sophisticated types of DSA that achieve higher spectrum utilization, and may enable access to bands or channels that otherwise would be withheld by risk-averse primary users, leading to higher performance than comparable devices without TLL support. As a result TLLs may justify a higher price for the radios, benefiting the vendor, while offering lower lifecycle costs to the end-user, through savings on spectrum access payments.

Longer term, TLLs may be one component of the solution to the safety concerns that today lock together the hardware and software of software defined radios. Currently regulators insist that the manufacturer take responsibility for testing the full configuration of a software radio, i.e. all software on exactly the deployed hardware. This prevents the development of Independent Software Vendors selling directly to the end user.[2] If TLLs helps mitigate the risks associated with pathological software interactions, it will become easier to unbundle SDR

---

[2] In the UK, Ofcom has proposed that ISVs be permitted as long as they test all software on the device when validating their offering. The last manufacturer who added software is responsible for any problems even if other software causes the problem. Though this is a step in the right direction, in practice it continues to limit the ability of a user to customize their device with a mix of services specific to their needs. Some manufacturer must conclude that the market for that application mix is large enough to justify doing the testing and taking the legal liability. In our view, the Ofcom approach is unlikely to lead to a vibrant ISV ecosystem.

hardware and software. This change is expected to unleash significant economic benefits just as it did in the PC industry.

D. Lease durations

We conclude our discussion of the applications of TLLs by considering what lease durations are likely to be used. Lease duration affects the design of the systems that is discussed in the next section.

We differentiate two durations: the duration of a lease certificate stored in a radio, and the duration of the corresponding contractual or regulatory arrangement that the certificate supports. While the certificate should not extend past the end of the contractual arrangement, it is sometimes desirable for it to end earlier.

Based on the applications just described, we anticipate useful certificate durations ranging between hours and months. One place where time periods shorter than this range may be relevant is in secondary spectrum access. However, for secondary spectrum transactions shorter than hours in length, more complex mechanisms may be needed to ensure the terms of the contract are respected. For example, the secondary user may be obligated to implement dynamic spectrum access techniques such as listen-before-transmit. At the other extreme, time periods longer than months will routinely arise in device certification. However, in these cases any unexpected problems that arise will need to be mitigated without waiting a year or more for the lease to expire. Given the anticipated low cost of extending leases, it seems likely that in most cases it will be more sensible to select a months-long certificate duration and extend it.

As lease duration decreases, costs increase due to overhead associated with the repeated distribution of new certificates, but responsiveness also increases. Very short leases—below an hour or two—approximate the functionality of a kill button. Unfortunately, with such short durations a temporary network outage would kill the radio functionality controlled by the lease. There are few applications where this behavior is likely to be acceptable.

Short leases are appropriate when the rights holder perceives high risk in the transaction or relationship. As operation continues and confidence increases, lease durations can be increased to reduce cost and uncertainty. So for example, a manufacturer might initially apply for a 3-month certification lease for an innovative radio, then renew for a 12-month period, then renew for 3 years. The required level of assurance would be higher for the longer leases. Safe operation of deployed units is important evidence to help provide the higher level of assurance.

Generally, we assume that leases are renewable (although renewal may necessitate a new contract). However, non-renewable leases would greatly simplify the lease subsystem and may prove useful for some disposable radio applications.

III. TECHNICAL ISSUES

This section discusses TLL implementation issues. A range of implementation options are available that provide cost effective support for a range of rights managements scenarios and radio applications. While the implementations described here are plausible, significant additional work will be required to determine the best approaches for different circumstances.

A. Cost-effective TLL implementation

In most radios there is a microcontroller or subsystem that translates between the device's multiplexed control bus and the individual control lines of the radio's RF analog devices such as amplifiers, oscillators, or filters. This is the most cost-effective place to add lease support to the radio (Figure 1). Since it is the only component with direct control of the transmit chain, a lease subsystem in this location can approve or reject any attempt to tune to a different frequency, change transmit power, bandwidth, or other parameters. We use the term *baseband processing* loosely to refer to the rest of the radio other than the lease subsystem and the analog components.

The lease subsystem should be separate from the baseband processing. Segregating the lease subsystem enables validating it to a very high level of confidence, at a reasonable cost. Segregating the lease subsystem also significantly increases its resilience to failures or security attacks elsewhere in the device.

The lease subsystem can be separated from the baseband processing in different ways, as shown in Figure 2. In (a) the subsystem consists of a set of hardware components on the board of a highly integrated mobile device. In addition to transmit control, control of the receive chain also passes through the lease subsystem to reduce the number of components and cost. In (b) the subsystem's functions are provided by a radio head that is linked to the baseband board via a standard interface. If there is no internal segregation within the radio head, the entire radio head must be validated to

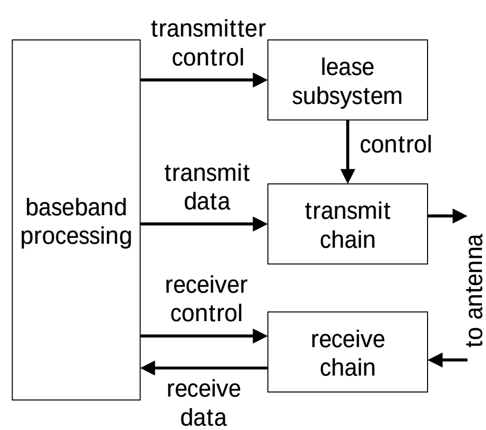

Figure 1. Architecture of a radio with lease support. The transmit chain contains the frequency up-conversion, signal amplifiers, and filter hardware. Control of the transmit chain passes through the lease subsystem, which checks that each new configuration is valid.

the high level of confidence required of the lease subsystem. This is a reasonable approach when the radio head is relatively simple. In (c) the lease subsystem functions are performed by an independent software process. The process boundaries of a commodity OS such as Windows XP or Linux provide sufficient isolation or protection for the lease process in some applications. In other applications, a more secure OS such as one with Trusted Computing support would be necessary. (Trusted Computing uses hardware-based security to enable specified software behaviors only when certain integrity conditions have been met.)

If the lease subsystem is implemented as separate hardware, there needs to be a processor, a clock, and some local storage (Figure 3). The hardware required is only a small increment beyond the microcontroller used at this place in current radios, so the cost increase will be trivial for all but the most cost-sensitive high-volume devices.

In cases where it is valuable to limit leases by location in addition to time, the designer can add a location sensor to the lease subsystem. This need not duplicate a sensor elsewhere in the radio. When some other part of the radio needs to know the location, it can read it out of the lease subsystem.

TLL support is not appropriate for all devices. There are some low-cost high-volume radios where the required hardware support would be too expensive to add (e.g., some sensors). At the other extreme, some radios already possess functionality that makes the TLL capabilities redundant (e.g., a centrally-controlled kill button capability). However, TLLs are appropriate for a wide range of devices and place only minimal constraints on radio system design.

B. Behavior of the lease subsystem

The lease subsystem has two major roles. In normal operation, it validates transmissions by the device. During a lease extension transaction, it authenticates and processes new certificates.

1) Transmit validation

Every time the transmit configuration is changed, the lease subsystem receives a request message from the baseband processing subsystem. It checks that the configuration is acceptable then establishes the requested configuration through its direct control of the transmit chain.

There are many ways to implement this behavior. Logically, it can be considered to be a lookup in a table with information of the type shown in Table 1.

In practice the information checked may be more sophisticated than what is shown in the table, for example including location information. The information may be represented as a set of rules rather than a lookup table.

TABLE I. INFORMATION IN THE LEASE TABLE

Frequency range	*Power limit*	*Bandwidth limit*	*Clock limit*
840-850 MHz	1 W	1.25 MHz	8:00AM 08/11/2007
1949-1952 MHz	500 mW	200 kHz	9:00 AM 10/24/2008
2400-2500 MHz	500 mW	<no limit>	<no limit>

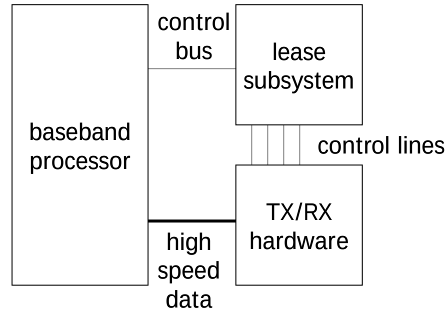

(a) Integrated Mobile Device

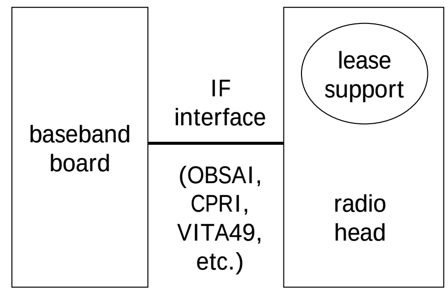

(b) Modular Infrastructure Device

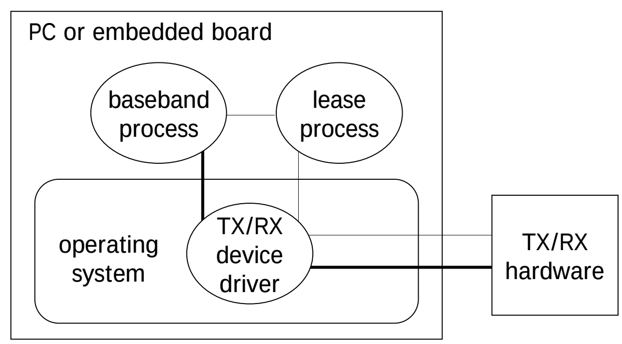

(c) PC-based Software Radio

Figure 2. Potential implementations of the lease architecture.

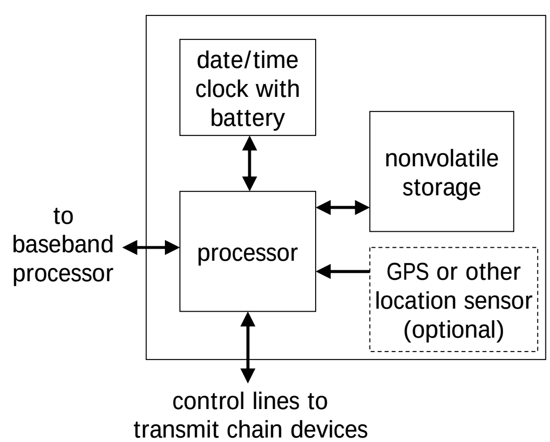

Figure 3. Components of the lease subsystem for an integrated mobile device. Local nonvolatile storage holds the current leases and crypto keys. An optional location sensor can support location-limited leases.

In some cases it is desirable to allow multiple entries in the table to cover the same frequency range. For example this can be used to support a disaggregated business model where multiple stakeholders contribute resources to a radio system (spectrum, waveform rights, device support). Each stakeholder controls its own lease extensions to protect its rights.

This example suggests an AND semantics where all entries must be valid to permit transmission. Other scenarios suggest an OR semantics where any entry covering a transmission request results in granting that request. For cases where a radio needs to support both options, a more sophisticated representation of the lease information is required. One representation described in the literature is the XG Policy Language [4], a very general approach that could be adapted for lease support. However the engine required to interpret the XG language is probably too complex to validate sufficiently for a lease subsystem, in addition to being too computationally expensive to implement in this context. New representations are likely to be required.

During transmission within a single transmit chain configuration, the lease subsystem only has work to do when a lease expires. The transmit configuration must be rechecked and if it is now invalid, the lease processor forces the transmitter into a safe configuration (e.g. turns it off) and notifies the rest of the radio. With proper design, the baseband processing can be notified in advance of an upcoming lease termination, preventing service interruption.

2) *Certificate processing*

New certificates are presented to the lease subsystem by the baseband subsystem, which receives them over any available communication channel from the appropriate authority. When a certificate is presented, the subsystem must authenticate it before updating the stored table.

Authentication can be done with a variety of technologies. The obvious one is to cryptographically sign the certificate with a private key known only to the appropriate authority, then use a public key stored in the lease subsystem to check the signature. This is computationally expensive and may take a long time to perform, perhaps multiple seconds on an integrated mobile device using a small embedded processor for the lease subsystem. This is not a problem since we anticipate lease durations of hours or longer (section II.D).

In some situations such as where leases support disaggregated business models, different authorities have control over different lease table entries. This can be supported by storing information in the lease subsystem that grants modification rights to different parts of the lease table to different authorities.

In other cases, multiple authorities must sign a lease certificate for it to be accepted. For example, a cellular operator may wish to configure the mobile devices in their network to only accept certificates signed by both the manufacturer and the operator. This is straightforward to support in the lease subsystem.

Finally, when a new lease certificate is presented to the lease subsystem that overlaps with a existing lease, the subsystem must know whether to replace the old lease or store multiple leases covering the same frequencies. If there will be multiple overlapping leases, the subsystem must know whether to use AND or OR or more sophisticated semantics. In some applications this information can be stored in the certificate itself, while in others the trust relationship among the various rights holders requires other control mechanisms. This is a fertile area for further research and one where the solutions are likely to be different for different applications.

We note that the ability to push a certificate out to a device, with the semantics that it replaces an existing entry in the lease table, enables approximating the behavior of a kill button through sending a certificate with a very short expiration. The difference is that a true kill button is guaranteed to shut down all targeted devices within a specified (presumably short) time. A lease system does not guarantee to reach all devices with the new certificate in any fixed period. Devices that remain unreachable continue operating until the end of their current lease period. In situations where this is acceptable, the short-expiration-time approach can be useful.

C. *Certificates*

Certificates will normally be fairly short, so they are easy to transmit over any wired or wireless network to the radio device. They contain some representation of the rights being temporarily granted to the radio, and an expiration day/time value.

1) *Representing rights*

We envision two main options for representing transmission rights in a certificate. *Model-independent* certificates specifiy abstract values such as frequency, power, bandwidth, and so on. *Model-specific* certificates specify particular values for the settings of the devices in the transmit chain.

With model-independent certificates any entity can generate a certificate, but the lease subsystem needs to be capable of computing which transmit chain configurations correspond to the specification. This is straightforward for the simplest leases (e.g. *on* vs *off*) but very challenging for anything more specific.[3] Using model-specific certificates makes the radio itself simpler. However, only the manufacturer will be able to convert model-independent specifications into a model-specific certificate in most cases.

Both types of certificates have important uses. Model-independent certificates may fit better in device certification applications, because only simple behavior controls are normally required and the regulator may wish to directly control lease extensions. In contrast, model-specific certificates provide the finer-grained control that may be needed for secondary spectrum transactions and innovative business

[3] Design engineers normally use sophisticated test equipment and creative reasoning when adjusting device parameters to ensure that transmissions conform to particular emissions specifications. This is a far more sophisticated operation than most radios can perform automatically.

models. The overall scenario needs to be evaluated—not just the TLL component—to determine the optimal certificate type in a particular case

In the model-specific case, the manufacturer will normally provide a service that others can use to convert model-independent information into certificates. The certificates themselves might be encrypted so that only that manufacturer's lease subsystems can interpret the bits, in order to avoid leaking proprietary information about the transmit hardware in the device.

2) *Applicability of certificates*

There are two main options for the applicability of certificates. *Targeted certificates* include an ID or group description that limits them to a specific radio or set of radios. The lease subsystem discards any certificates that are not valid for that device. *"Bearer-bond" certificates* are valid for any lease subsystem that receives them.

To support targeted certificates, a unique device ID or other information must be installed in the lease subsystem's local storage during manufacture or configuration. For bearer-bond certificates, the designer uses other system functions to ensure that the certificates only reach the correct radios.

Although targeted certificates may seem superior, bearer-bond certificates are attractive because they partition the implementation complexity of leases, separating the targeting task out of the lease subsystem. Designers can select the level of security against mis-delivery, interception or duplication of certificates that is appropriate for each application. Therefore neither targeted nor bearer-bond certificates are preferable in all situations.

In many scenarios it would be desirable to target certificates to specific baseband software version numbers. For example, a primary spectrum rights holder who is particularly concerned about interference, such as a public safety agency, might want to prevent unauthorized upgrades by secondary users. In these cases the targeting must be elsewhere in the system. Bearer-bond certificates must be used because the lease subsystem cannot validate which software is running in the baseband processing subsystem.[4]

D. Certificate management

Lease extension certificates may be generated and distributed in a number of ways, corresponding to different applications and business models. Indeed, the choice of how such a mechanism is implemented offers one way to expand the richness of spectrum rights management afforded by the lease mechanism. In the simplest case, there are at least two parties involved: the certificate creator, presumably the rights holder, and the radio operator who receives the certificate and distributes it to the device(s). It may be desirable to introduce additional intermediaries such as a trusted third party to better enforce trust relationships or to realize scale/scope economies (e.g., a centralized lease clearing house for lease management of multiple radio systems).

Many variations on the basic approach are possible. We describe a few interesting examples.

Operator control: Operators who wish to maintain final control over the devices in their network can configure the devices to only accept certificates signed by the operator. So the certificate would initially be signed by the rights holder and then by the operator before being distributed to the devices.

Autonomous devices: Rather than the operator pushing lease extensions out to devices, the baseband software in the devices can be configured to automatically retrieve a new certificate when the end of the current lease is near. For example any device with an internet connection can go to a pre-configured web site to get its lease extensions. This web site might be hosted by the operator, in the case of a cellular telephone network, or by the manufacturer, in the case of unlicensed or directly-purchased devices.

Certification leases: The rights holder is the regulatory authority. Rather than generating lease extension messages directly, in many cases the regulator will delegate this authority to the manufacturer. Legal sanctions are used to prevent extension of leases when a fielded device has caused interference or harm. Once generated by the manufacturer, the new certificates might be posted on a manufacturer web site or given to network operators for distribution.

One option available to the regulator is to require that all certificates are registered with it. This can be accomplished cheaply by providing an online service that automatically signs (and stores) any certificate presented to it by the manufacturer. The lease subsystem would be configured to check for the regulator's signature in addition to the manufacturer's.

Broadcast beacons: A beacon scheme has been widely discussed for controlling secondary spectrum access in limited geographic areas, for example by the US FCC in the Cognitive Radio NPRM [5]. Such a beacon could broadcast lease certificates, giving the spectrum rights holder fine grained control over which devices or users operate in the spectrum or what modes they operate in. In this application, the rights holder would establish contractual arrangements with approved secondary users, discover the information (such as device IDs) needed to generate targeted certificates, then broadcast them directly to the devices.

Spectrum distributors: A potential player in the future secondary spectrum market is the spectrum distributor. More sophisticated than a spectrum broker who just matches buyers and sellers, the distributor acquires, aggregates, partitions, and packages spectrum rights and futures [3]. The spectrum distributor can generate the appropriate lease certificates and provide them to the user as part of completing a spectrum transaction.

[4]The Trusted Computing approach now supported by Intel and Microsoft is an important exception to this observation. If it is included in a radio device and used for the baseband processing software, the lease subsystem will be able to check software versions when deciding whether to accept a certificate.

E. Assurance and trust

This section summarizes the failure and threat model for lease enforcement. We discuss each of the possible problems in turn.

Design errors: We expect a very low probability of design errors, because the lease subsystem is an independent hardware unit or software process with simple functions. This makes it feasible and cost-effective to validate it to a high level of assurance.

On the hardware side, the primary design challenge is clock accuracy, for which there are multiple potential solutions (Section III.F). On the software side, the primary challenge is the relative complexity of transmit validation and certificate processing. Both the functionality specified for the software and the strategies used to implement it will need to be carefully selected to permit a high level of verification. We expect that lease subsystem software will frequently be shared across product families and potentially sold/traded among manufacturers, to amortize verification and certification costs.

User attacks: The user who physically controls the radio device knows that if they can interfere with the behavior of the lease subsystem, they can boost the capability of their device or possibly get free service. Therefore attacks by the user are a valid concern in many applications. Consequently, the lease subsystem must be tamper-proof or at least tamper-resistant against both physical and software-based attacks.

Furthermore, in many cases the baseband processing software may be compromised. This assumption is based on the importance of commodity open platforms (e.g. PalmOS or Windows Mobile) and user-loaded software in many radio devices. Even if the baseband software runs on its own processor, the presence of user software anwhere in the device enables a variety of attack vectors against it. It will often be cheaper and more effective to protect the simple lease subsystem (with its minimal interface) against the baseband software than to protect the complicated baseband software (with its rich interface) against the rest of the device. As a result, the lease subsystem in these situations should be validated to behave correctly no matter what sequence of bits is sent over the connection from the baseband processor.

Because the user can intercept, modify or forge certificates, a strong level of encryption should be used in authenticating them in these applications. Similarly, targeted certificates are better than bearer-bond certificates when the baseband software is not trusted. One consequence is that certificates cannot be specific to particular software versions for most devices. However, this is not a problem: if baseband software integrity may have been compromised, its software version number is not a reliable predictor of its behavior. Only the limitations enforced by the lease subsystem can be relied on to protect others from harm.

Third-party attacks: Third party attacks could include denial-of-service attacks seeking to leverage the lease subsystem to shut down the radio, or attacks seeking to add leases and thereby generate interference in protected bands. One could conceivably trust the user while still seeking to defend against third-party attacks, and try to reduce system complexity or cost on that basis. However, the most likely third-party attack vector is a trojan horse that takes control of the user-loaded software and attacks the baseband processor. We therefore regard third-party attacks as identical to user attacks; the threats and response are the same.

Manufacturer attacks: The manufacturer can always cheat. A back-door can be built in during design, or an extra crypto key can be added during configuration allowing lease extensions whenever the manufacturer wants it. As a result, the lease mechanism is not a means of enforcing compliance on unwilling manufacturers, so attacks by the manufacturer can be disregarded in the threat model.

F. Accuracy of the clock

The day/time clock used by the lease subsystem must be accurate in order to enforce the termination times specified by leases. The lease subsystem will need to allow for drift, for example by ceasing transmission 5 minutes before the scheduled end of a lease if there may be as much as 5 minutes clock error.

Maintaining clock accuracy is more challenging for leases than in most radio or computer environments. Normally, network time protocols (NTP) would be used, or time update messages would be sent to the radio. However, in many lease applications we expect attacks by a user who may control the baseband processing software. Therefore NTP and time messages presented to the lease subsystem cannot be trusted. A user facing the end of a lease could set the clock back by a week and keep the radio operating illegitimately. Similarly, the user cannot be permitted to reset the clock after power failure.

There are multiple solutions to this challenge, appropriate for different applications.

- Many systems have GPS receivers in them. Including the GPS system in the lease subsystem provides a resilient non-user-modifiable time source.
- Day/time chips with integrated battery backup are available and cost-effective. The battery can keep the clock accurate for many months without external power.
- Signed time update messages or secure network time protocols can be used to share time information between a trusted remote authority and the lease subsystem while preventing interference by the untrusted baseband software. These approaches add complexity to the lease subsystem and hence should be avoided if possible.

Finally, note that as the duration of the leases increases, the system becomes more robust to small clock errors.

IV. ECONOMIC AND POLICY ISSUES

In this section, we consider further the economic and policy implications of introducing Time-Limited Lease (TLL) radio designs. First we discuss how TLL designs will impact the costs and benefits realized from radio systems. Then we consider the economic/policy context in which TLL designs may be introduced. Finally, we analyze the challenges that TLL designs must overcome for successful commercialization and how market/technology trends are likely to impact those challenges.

A. Economic costs of TLLs

There are two types of costs created by leases. There can be an increase in capital expense (Capex) for the design, manufacture and validation of TLL enabled radios. There can also be an increase in operating expense (Opex) for the configuration, management and maintenance of the radio system.

The capex increase for lease support can be reduced to a low level in many cases. As discussed in section III.A, a wide range of implementation approaches are possible that match different radio designs and system requirements. Manufacturing cost is reduced through adding lease support in a way that minimizes changes to the existing architecture. Validation cost of the lease subsystem is reduced by segregating it from the rest of the radio and keeping it simple.

The opex increase for a radio system that uses leases is more difficult to analyze. Radio configuration is certainly more expensive due to the need to initialize encryption keys, and the need to insert a radio or group ID in some cases (section III.C.2). The manufacturer may need to be involved in configuration, to initialize the lease subsystem on behalf of mutually distrustful collaborating entities. There will also be some operating overheads in order to periodically distribute lease extensions, but this cost can be adjusted to an acceptable level by changing lease durations. There can be new kinds of maintenance costs, for example replacing the battery in the lease subsystem's clock. Overall we do not see a signficant opex increase for many radio applications, but this question will need to be analyzed in more detail for specific systems and requirements.

Some of the available design choices trade off capex against opex costs. For example, model-independent certificates reduce the communication costs of the renewal mechanism, compared to model-specific certificates, while requiring greater sophistication in the radio. When such greater sophistication is present in any case, then this may offer a lower cost overall deployment strategy.

There are some radio designs and applications where the cost increase associated with TLLs is unacceptably high despite these strategies, for example sensor networks. Still, the costs appear acceptable for many systems due to the wide range of available design and mechanism options. In any case, because we recommend that TLL support be optional, manufacturers can opt for a traditional design if that is more cost-effective for a given application once the benefits of leases are considered.

B. Economic benefits of TLLs

Recent innovations in radio system design have greatly expanded the capabilities of both centrally-managed networks such as cellular telephone systems and decentralized networks such as ad hoc radio networks and meshes. However, market deployment of decentralized networks is just beginning. It is for these systems that adopting TLLs appears to offer the greatest potential benefits. Hence, our discussion focuses on the benefits of TLL in the sorts of decentralized and distributed radio system environments that are prototypically associated with unlicensed usage by cognitive radios. Many of the benefits described also apply to centralized wireless network designs.

1) Lower costs for radio development and certification

As noted earlier, TLL-enabled radios may be certified subject to a less stringent (less costly) process for enabling enhanced functionality. The manufacturer then has the flexibility to choose whether to pursue the more costly certification approach, defer offering the enhanced capabilities, or adopt the TLL design approach. This added design flexibility expands the design space confronting the manufacturer and therefore helps reduce overall development costs. The cost savings provided by lightweight certification for TLL radios is obviously related to the length of the leases used. As the leases become longer they approach the life of the non-TLL radios and so the need for additional ex ante protection to avoid on-going operation of a defective radio becomes more important.

2) Reduced risk of incorrect or unapproved operation

TLLs offer a technical method for cost-effectively enforcing regulatory or contractual requirements. Because TLL renewal may be contingent on approval of the regulatory authority or rights holder, TLLs can both enhance incentives to design correct radios in the first place, and if problems do arise later, can help limit the potential harm from non-compliant operation. TLLs enhance ex ante incentives because the threat of non-renewal offers a credible threat that faulty operation will result in early termination of the radios ability to operate. The TLL mechanism may also be used to enforce upgrades.[5] Although not a kill button, the TLL capability limits the potential harm from a radio by ensuring that the radio will cease to operate once the lease expires.

Relative to a system without TLLs where unauthorized radio behavior might continue in perpetuity, the time-limited lease approach offers a way to strictly bound the harm that a device might cause. TLLs eliminate the future tail of potential harm (that is, an interfering radio that operates for 1 year causes less harm than one that operates for 10 years). By

[5] The upgrade capability is related to the TLL feature but is distinct. A TLL radio may not be upgradable and an upgradable radio may not be TLL-enabled. Indeed, complex radios may opt for an alternative mechanism for upgrading (e.g., centrally-managed real-time polling) that eliminates the need for TLL support. Alternatively, it may be lower cost and preferable in some contexts to have a radio die and be replaced rather than upgraded if the lease is not renewed.

trimming the potential harm from such adverse scenarios, TLL reduces the risks of deploying more complex radio system designs. In effect, the TLL may be thought of as a real option: it provides a mechanism for bounding the harm that a radio may cause over its otherwise unconstrained lifetime.

Additionally, design features which lower the ex ante expectation of ex post interference harm also reduce incentives for parties (both potential interferers and their victims) to invest in unnecessary protection or interference-rights enforcement. This eliminates a deadweight loss to the overall surplus.

3) Lower mitigation costs for incorrect behavior

With complex radio designs, a decentralized mechanism to stop operation provides a low-cost option for implementing a radio recall. While a lease mechanism may not be sufficient on its own to address all recall concerns, it can help, especially when combined with other strategies (e.g., simple certification rules that provide clear incentives for manufacturers to design radios that operate predictably and in compliance with spectrum rules).

4) Lower transaction costs for spectrum trading

TLL designs help facilitate the deployment of dynamic spectrum access (DSA) and secondary spectrum markets. Such secondary markets may take the form of trading of exclusive-use rights or allowing opportunistic secondary-use. They may be voluntary (enabled by the primary licensee) or involuntary (enabled by easements on primary licensee rights such as whitespace easement or underlay) [3]. Our earlier discussion of MVNOs and MVSOs provide some examples of how new types of operators may emerge. Other DSA applications may include low-power sharing of spectrum by consumer appliances (to implement home network alternatives to WiFi or UWB) or niche applications such as data back-up services (collating data from sensor networks) or ad hoc networking.

A key feature of DSA markets is the need to allocate and deallocate spectrum rights to devices. Assurance that rights can be transferred with predictability and certainty (including that usage will terminate at the expected time) will enhance the likelihood that such markets can evolve. Thus, lease capability enables further partitioning of spectrum usage rights to enable richer market-transactions for defining and managing property rights for spectrum. For example, in addition to making spectrum access frequency and location dependent, this capability facilitates making access time-dependent. This expands the range of potential architectures and business models that may be deployed and further reduces arbitrary regulatory boundaries. TLLs are a tool to help assure devices respect the durations of secondary market spectrum trades. Once again, TLLs are not sufficient in themselves, but should be regarded as one among several design options that can help facilitate the emergence of secondary markets for spectrum.

5) Lower costs for distributed control

In decentralized or ad-hoc networks, and in innovative disaggregated wireless service business models, a key problem is how the multiple collaborating entities can each enforce their respective rights and reward others for supportive behavior. TLLs provide a lightweight technical mechanism to distribute control of radio behavior among multiple entities and to control distributed radios. TLLs thereby reduce the high costs of non-technical enforcement of rights and control, through contractual, legal, or other means.

6) Lower lifecycle costs for product line management

Lease mechanisms may also lower life-cycle product management costs for industry. To the extent that TLL support in radios results in more predictable product lifetimes with clear action points for legacy equipment (predictable obsolescence) this can lower value chain design costs. Mechanisms for this include enhancing coordination between product release timing across providers of different components and enabling easier standards adoption and regulatory reforms.

The TLL renewal process can facilitate updating and inventory tracking, although both of these functions are logically distinct from the TLL functionality. Such supply chain management improvements can lower costs and enhance options for just-in-time manufacturing. For example, as already noted, TLLs can be used to selectively enable additional functionality for radios already deployed. As a customer's needs grow, the vendor can send renewal leases enabling more functionality for an added payment.

7) Lower costs for radiocommunications regulation

Finally, it is worth considering the overall impact of TLLs on regulatory costs. Generally and abstractly, regulators have two options for regulating: they can specify general performance rules and then monitor ex post behavior to enforce compliance with the rules; alternatively, they can specify a technology ex ante (that has limited performance capabilities) and economize on enforcement costs (since the technology constrains behavior). The former strategy permits greater flexibility at the expense of additional monitoring/enforcement costs. The latter strategy economizes on monitoring/enforcement costs at the expense of limiting flexibility. In the present context, TLLs are a type of ex ante technical constraint that helps lower ex post enforcement/monitoring costs (by lessening the likelihood of harm and enhancing the ability to limit harm if it occurs) while preserving a maximal degree of design flexibility (allowing behaviors that might cause harm). The regulatory compromise implemented by TLLs is likely to be especially valuable in light of the increasing complexity of wireless systems enabled by the new business models and radio designs made feasible, in part through TLL, and in light of the greater difficulty in detecting and protecting against illegal interference, especially in distributed/decentralized radio system deployments.

8) Summary

Traditional certification rules and clearly defined rights and liabilities for interference provide the context within which TLL-enabled radios will operate. The emphasis on enabling implementation in a decentralized, distributed, and low cost fashion makes TLLs especially beneficial for market-based spectrum management environments, such as unlicensed bands, where risks from interference are low and tolerance for increased transaction costs is low.

C. Economic and policy context for introducing TLL radios

The preceding sections discussed the impacts of TLL on radio system design costs and benefits. In this section we discuss the likely context in which TLLs may be introduced and the implications of doing so.

1) Mandatory vs optional lease support

In our view, TLL support should not be mandatory except in special circumstances, such as operation in a band where there is a regulatory sunset clause or other requirement that mandates time-limited radio behavior. In most cases, the decision should be left up to the manufacturer or operator whether TLL support offers sufficient benefits, such as reduced certification challenges, to outweigh its cost. If our analysis is correct, the advantages of TLL will prove sufficiently compelling and its cost low enough that manufacturers elect to implement the capability voluntarily in many cases.

On the other hand, the government may play a useful role in helping the market to coordinate on appropriate approaches. The government can guide adoption of TLL technology by offering a clear road-map for certification requirements for advanced radio designs, and by considering potential roles for leases when adopting reforms that facilitate the transition to DSA and the emergence of secondary spectrum markets.

2) Centralized vs distributed lease systems

Control of the TLL mechanism rests with the entity that decides whether to renew a lease and creates the lease extension certificates. This may be left to individual manufacturers, operators, or rights holders subject to some general rules or may be administered by a centralized authority. The central authority could be a government authority or other trusted third party (e.g., an industry association).

There are tradeoffs in this decision. Centralized management of the TLL process may ease implementation and provide a smoother incremental transition from existing certification and spectrum licensing procedures. On the other hand, this is contrary to the general goal of moving towards market-based processes and raises concerns about regulatory uncertainty due to increased ex post central control compared to the current regime.

We anticipate that there will be a need for multiple TLL mechanisms—some that are manufacturer/device-specific and some that may be shared across manufacturers and perhaps across bands. Still, there are likely to be benefits from standardizing elements of the TLL framework. This will help realize scale and scope economies, thereby further lowering costs. It will also increase trust in the system, further reducing the risk perceived by various stakeholders, and establish standard mechanisms for resolving disputes about leases or lease-empowered contracts (e.g. certificate escrow).

Whether it makes sense to create a TLL management framework that is shared across many radio systems and, if so, how it might be designed is a topic for further research. While development of a common infrastructure may offer benefits, it raises the challenging question of how to coordinate shared investment and design of a common TLL infrastructure across all the stakeholders and different radio applications requirements.

3) Traditional vs unbundled wireless service value chain

TLL-enabled radios will support and their use will be encouraged by ongoing changes in the wireless communications service value chain. TLLs facilitate dynamic spectrum access; this is mutually reinforcing with trends towards decoupling RF spectrum bands from the applications that use them. TLLs enable reduction of risk when rights holders collaborate; this is mutually reinforcing with trends towards decoupling of network ownership, spectrum rights ownership, and service provision. TLLs used in certification limit the harm caused by radios that misbehave; this is mutually reinforcing with trends towards increased vertical disintegration in radio manufacturing. TLLs enable distributed control over radio behavior; this is mutually reinforcing with trends towards end-user owned mesh networking solutions.

We do not presume that these changes will all take place quickly or to their maximum imaginable extent. Current business models will continue to remain viable and will coexist with new unbundled, distributed, and open architectures. Over time, the growth of the novel business models and architectures will enhance competition and will increase demand for TLL support in radios.

D. Market challenges for commercialization of TLL radios

Radios with TLL support face a number of challenges to their successful commercialization. We focus the discussion in this section on use of TLLs in device certification. Many of the issues discussed also apply to other applications of TLLs.

Radios with time-limited certifications confront several challenges, including the following.

- Fear that radio life will be terminated prematurely because the TLL is not renewed as expected.
- Fear that lightweight certification will weaken ex ante incentives to invest in regulatory compliance.
- Fear that regulators will abuse TLL capability to expropriate surplus rents or impose additional implicit taxes.

1) Shorter-than-expected radio life

Even if a radio operates safely and correctly in the field, there is a non-zero probability that the radio's operating lease will fail to be renewed, resulting in a life-span that is shorter than the user expected. This could arise because of a failure in the renewal system (*ships at sea* problem) or because of intentional behavior by one of the parties (*hold-up* or *moral hazard* problem).

The *ships at sea* problem arises because a radio may be deployed in an environment that prevents it from renewing its lease when the lease expires. Such an occurrence may be both likely and unacceptable in some applications, for example in the radios ships use to communicate at sea.

While this concern is real, it is likely to be localized to specific applications. One response may be simply to not use

TLL-enabled lightweight certification for such radios. Another strategy that enables these radios to exploit the benefits of TLLs is to use lightweight certification only for enhanced capabilities. When a lease is not renewed, only the enhanced capabilities are blocked from operating. In such a circumstance, the radio would revert back to a default (presumably more limited) behavior mode. Note that the default mode may well have the same capabilities as would be offered by radios that do not have TLL support. The cost to the user of exploiting the advanced features enabled by TLL-based certification is the risk that these features will become unavailable if the renewal mechanism fails.

The relative importance of the *ships at sea* problem varies by application, which reinforces our recommendation that TLL ought to be an optional feature for radio manufacturers.

With respect to the *hold-up* or *moral hazard* problem, the issues are somewhat more complex. A hold-up problem arises when one party to a transaction can take ex post actions (after the transaction has been completed) that adversely impact the interests of the other party. Such problems arise in many situations, for example when a contractor with a cost-plus contract takes too little care to control project costs, or when a supplier charges higher prices for subsequent deliveries knowing that the buyer faces high costs in switching to another supplier.

Radios whose certification depends on leases give rise to a hold-up problem. The entity controlling the lease system can threaten to terminate the operation of an otherwise satisfactory radio after it has been purchased. If that entity is the manufacturer, they can exploit the opportunity to extract additional value from the buyer. This value could take several forms other than just direct charges for lease renewal, for example forcing the customer to upgrade the device sooner than they would choose to otherwise.

The risk associated with the hold-up problem needs to be evaluated in the context of the radio's application. Market and technical trends are leading users to replace radios more rapidly in many applications. All other things being the same, increasing the duration of certification leases and accelerating product replacement (due to faster innovation, modular designs, and lower transaction costs for upgrading) reduce moral hazard concerns since they reduce the cost to the user of premature termination. Whether premature termination is a serious problem or not will depend on the circumstances. In general, it seems likely that the world that TLL-based lightweight certification will help give rise to and will be most needed in will be precisely the type of world where the hold-up problem is of reduced concern compared to current radio systems.

Nevertheless, there are well-established methods for addressing moral hazard problems. First, clear contracting terms that limit the scope for premature termination and enable buyers and sellers to establish clear expectations as to the lifetime of TLL-enabled radios will allow the shorter life-span (if any) to be reflected in prices, thereby eliminating the moral hazard problem. The contracting terms may include restrictions to limit discretionary termination. Indeed, if discretion over the implementation of TLL renewal is removed from the manufacturer to a trusted third party (the government or an industry trade association, perhaps), then the moral hazard problem again disappears.

More generally, it may be the case that the manufacturer will possess private knowledge that would make it better able than any other party to determine when a TLL should be renewed. It is this situation that gives rise to the moral hazard problem in the purest sense. Under such circumstances, the best response may be for the manufacturer to internalize the moral hazard problem. There are a number of ways this may be accomplished. Vertical integration offers one mechanism (the manufacturer and the radio operator are one and the same). Alternatively, the manufacturer could sell the radio with ex ante guarantees to reimburse the customer if the radio is terminated before a fixed time. By so doing, the manufacturer would assume the risk associated with the costs imposed by unexpected premature termination. The manufacturer could then internalize the trade-offs from additional ex ante investment in ensuring ex post safe behavior versus having to bear the costs of terminating radios that behave inappropriate ex post. By allowing the manufacturer to make a credible commitment (premature termination would be costly for the manufacturer), TLLs can enhance manufacturer incentives to build safe radios and to respond optimally if and when those radios need to be fixed. This enhances overall efficiency.

The risk of a moral hazard problem arising is also mitigated by increased competition. In the absence of market power, competition drives the price of goods and services toward their economic costs. In a market in which buyers can choose among TLL and non-TLL radios with differing sets of capabilities, it is reasonable to expect that TLL radios will sell at a discount relative to equivalent quality non-TLL radios. More generally, because we expect TLL radios to be desirable in part because they allow manufacturers to introduce capabilities that otherwise could not be cost-effectively certified without TLL, one might expect the TLL radios to offer capabilities not available in non-TLL radios. When this is the case, the relative prices of the radios will depend on how users assess the relative value of the additional functionality versus the potential that the lifespan of a TLL radio may be shorter.

Because TLL radios enable new radio designs, including DSA and new distributed/decentralized radio architectures, we expect that TLLs will enhance competition for wired services, as discussed earlier. Thus, overall, we believe that wider deployment/adoption of TLL will be pro-competitive and so will help reduce the risk of market-power-based opportunistic behavior of all sorts, including the moral hazard problem specifically addressed here.

2) Reduced investment in regulatory compliance

A second concern is that TLL-enabled lightweight certification may reduce incentives for manufacturers to design radios that comply fully with regulatory requirements, by making it easier to fix "poorly" designed radios ex post. We put "poorly" in quotes because the appropriate level of investment in compliance and validation is context dependent.

Responding to this concern, it is our view that the complexity of advanced radios makes it impossible to design fault-free systems or to fully validate a device under test. Given that faults will exist in deployed devices, it is appropriate to develop low-cost strategies for addressing problems when they arise. Moreover, as already noted, the threat of non-renewal offers a credible penalty to manufacturers of noncompliant radios and thus can actually serve to enhance manufacturers' incentives to achieve full regulatory compliance compared to the current certification approach.

3) Regulatory uncertainty

Finally, there is the risk that the TLL mechanism, if under the control of the regulator, may be used to levy additional taxes on spectrum users, support more frequent and arbitrary regulatory changes, or otherwise to expropriate additional rents. This is just another example of moral hazard, but because of the added concerns about regulatory inefficiency and capture, it is worth identifying this concern separately.

This concern may be mitigated by limiting the role of the regulator in radio design (striving toward "technically neutral" regulation) and in administering lease renewal decisions (e.g., by either leaving control of lease renewals in the hands of the manufacturer or an industry-sanctioned trusted third party). Clear and simple rules—the minimalist approach toward designing and managing the TLL mechanism—will aid transparency and make credible commitments more feasible, which will also reduce concerns over regulatory commitment and uncertainty.

If the government determines that some sort of fee-based mechanism for spectrum access is appropriate, a TLL mechanism introduced to reduce certification barriers would make this tax easy to levy and enforce. This hazard can be addressed through credible commitments that lease renewal will be provided at a prespecified cost throughout the lifetime of the device.

V. CONCLUSION

Time-limited leases (TLLs) are technically feasible to implement at affordable cost in many radio systems. They appear likely to provide significant benefits for the commercialization and deployment of innovative radios. Key applications include certification of innovative radios, facilitating secondary spectrum markets, and supporting innovative business models.

REFERENCES

[1] United States Federal Communications Commission, "Spectrum Policy Task Force Report," 2002, available online at http://www.fcc.gov/sptf/reports.html.

[2] United Kingdom Office of Communications (OFCOM), "An Evaluation of Software Defined Radio," 2004, available online at http://www.ofcom.org.uk/research/technology/overview/emer_tech/sdr.

[3] J. Chapin and W. Lehr, "The path to market success for dynamic spectrum access technologies," IEEE Communications Magazine, April 2007, in press.

[4] D. Wilkins, G. Denker, M.O. Stehr, D. Elenius, R. Senanayake, and C. Talcott, "Policy-based cognitive radios," IEEE Wireless Communications, Special Issue on Cognitive Wireless Networks, 2007 (in press).

[5] Cognitive Radio NPRM, ET Docket no.03-108, Notice of Proposed Rule Making and Order, 18 FCC Rcd 26859 (2003)