Spectrum Sensing

Fundamental Limits and Practical Challenges

Anant Sahai   Danijela Cabric

presenting joint work with

Robert W. Brodersen

Niels Hoven   Shridhar Mubaraq Mishra   Rahul Tandra

Wireless Foundations and Berkeley Wireless Research Center
Department of Electrical Engineering and Computer Science
University of California, Berkeley

Major Support from ITR Award: CNS - 0326503
Thanks to Joe Evans at NSF

Dyspan 2005

Spectrum, spectrum, everywhere, but . . .

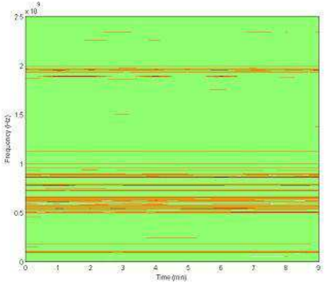

- Available spectrum looks scarce.
- Measurements show the allocated spectrum is vastly underutilized.

Utilizing available spectrum: five basic approaches

- A new comprehensive commons — eliminate legacy users entirely.

Utilizing available spectrum: five basic approaches

- A new comprehensive commons — eliminate legacy users entirely.
- Preserve some priority for "primary users"

	Interference management is primary's responsibility	Interference management not primary's responsibility
Secondary has permission	**Markets**	**UWB**
Secondary must take care	**Denials**	**Opportunistic**

Utilizing available spectrum: five basic approaches

- A new comprehensive commons — eliminate legacy users entirely.
- Preserve some priority for "primary users"

	Interference management is primary's responsibility	Interference management not primary's responsibility
Secondary has permission	**Markets**	**UWB**
Secondary must take care	**Denials**	**Opportunistic**

- Ultra-wideband: blanket permission
  - "Speak softly, but use a wideband"
  - Energy limited regime —works because most bands are not used

Utilizing available spectrum: five basic approaches

- A new comprehensive commons — eliminate legacy users entirely.
- Preserve some priority for “primary users”

	Interference management is primary’s responsibility	Interference management not primary’s responsibility
Secondary has permission	**Markets**	**UWB**
Secondary must take care	**Denials**	**Opportunistic**

- Ultra-wideband: blanket permission
  - “Speak softly, but use a wideband”
  - Energy limited regime —works because most bands are not used
- Secondary takes care: avoid disturbing others
  - Denials: primary signals when it is being disturbed
  - Opportunism: secondary keeps listening for the primary

Spectrum management phases

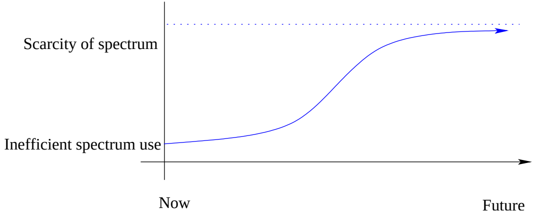

- Evolutionary steps make sense
  - Spectrum is locally plentiful, so why disable legacy uses now?
  - Full utilization before rationing

Spectrum management phases

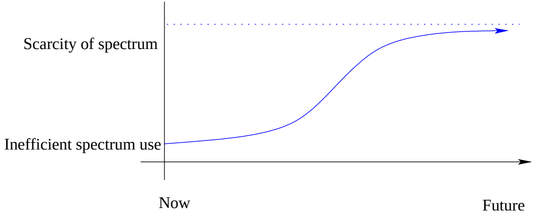

- Evolutionary steps make sense
  - Spectrum is locally plentiful, so why disable legacy uses now?
  - Full utilization before rationing
- Markets: "trying to solve tomorrow's problems today."
  - Make sense *if* resource is scarce. Otherwise, can't recover fixed costs.
  - More efficient if "commoditized"
  - Need to learn demand and competing uses first.

Spectrum management phases

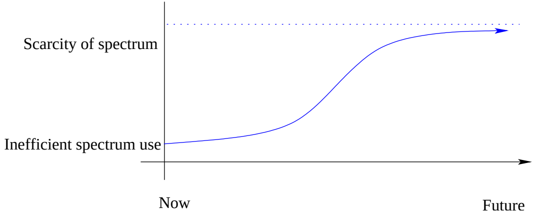

- Evolutionary steps make sense
  - Spectrum is locally plentiful, so why disable legacy uses now?
  - Full utilization before rationing
- Markets: “trying to solve tomorrow’s problems today.”
  - Make sense *if* resource is scarce. Otherwise, can’t recover fixed costs.
  - More efficient if “commoditized”
  - Need to learn demand and competing uses first.
- Explicit denials: “your pain for my gain”

Spectrum management phases

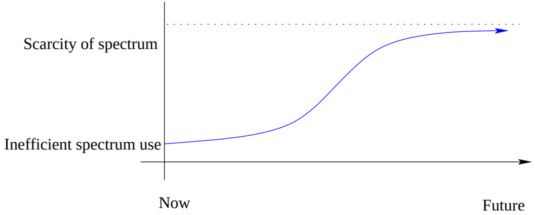

- Evolutionary steps make sense
  - Spectrum is locally plentiful, so why disable legacy uses now?
  - Full utilization before rationing
- Markets: “trying to solve tomorrow’s problems today.”
  - Make sense *if* resource is scarce. Otherwise, can’t recover fixed costs.
  - More efficient if “commoditized”
  - Need to learn demand and competing uses first.
- Explicit denials: “your pain for my gain”
- Focus on opportunism for now.

Objectives

- Reclaim underutilized spectrum
  - Peaceful coexistence
  - (Hopefully) cheap devices

Objectives

- Reclaim underutilized spectrum
  - Peaceful coexistence
  - (Hopefully) cheap devices
- *“If a radio system transmits in a band and nobody is listening, does it cause interference?"*

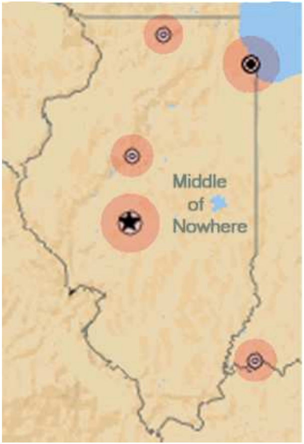

Objectives

- Reclaim underutilized spectrum
  - Peaceful coexistence
  - (Hopefully) cheap devices
- *"If a radio system transmits in a band and nobody is listening, does it cause interference?"*
  - Interference temperature attempts to quantify this
  - Allowable interference depends on many variables

Not just the middle of nowhere!

Big questions

- How well must we sense?
  - How sensitive *can* I be?
  - Can I use more power if I'm more sensitive?

Big questions

- How well must we sense?
  - How sensitive *can* I be?
  - Can I use more power if I'm more sensitive?
- What coordination is required?
  - What are the benefits?
  - How much coordination to realize these benefits?
  - What if I don't trust everyone?
  - If I am more capable, can I avoid having to coordinate?

Big questions

- How well must we sense?
  - How sensitive *can* I be?
  - Can I use more power if I'm more sensitive?
- What coordination is required?
  - What are the benefits?
  - How much coordination to realize these benefits?
  - What if I don't trust everyone?
  - If I am more capable, can I avoid having to coordinate?
- What kind of usage patterns can be supported?
- What are the key hardware challenges?

Outline

- **An overview of the issues involved and some key ideas**

Outline

- **An overview of the issues involved and some key ideas**
- Part I: The basic considerations quantified
  - The "sensing link budget"
  - The energy detector and fundamental limits on its sensitivity
  - Within-system cooperation
  - Fairness and cooperation among systems

Outline

- **An overview of the issues involved and some key ideas**
- Part I: The basic considerations quantified
  - The "sensing link budget"
  - The energy detector and fundamental limits on its sensitivity
  - Within-system cooperation
  - Fairness and cooperation among systems
- Part II: More powerful detectors
  - Coherent detectors
  - Feature detectors

Outline

- **An overview of the issues involved and some key ideas**
- Part I: The basic considerations quantified
  - The "sensing link budget"
  - The energy detector and fundamental limits on its sensitivity
  - Within-system cooperation
  - Fairness and cooperation among systems
- Part II: More powerful detectors
  - Coherent detectors
  - Feature detectors
- Part III: Hardware considerations
  - Fundamental hardware limitations
  - Ways around them

A myriad of factors to consider

- Primary power
- Amount of protection given primary
- Multiple secondaries
- Heterogeneous propagation losses
- Multipath and shadowing
- Coherence times
- Primary duty cycles
- Secondary power
- Cooperation
- Competition
- Modulation models
- Implementation complexity
- Robustness
- And many more...

We show how to understand these in the spectrum sensing context by building up slowly from simple cases to see how everything fits together.

Summary of main points

- Uncertainty imposes limits on device sensitivity that can *not* be overcome by just listening longer.
- Shadowing is a major challenge but can be overcome with multi-user cooperative diversity within a system.
- Potential interference from other secondaries is a very significant uncertainty, but can be mitigated through mandated local cooperation among systems.

Summary of main points

- Uncertainty imposes limits on device sensitivity that can *not* be overcome by just listening longer.
- Shadowing is a major challenge but can be overcome with multi-user cooperative diversity within a system.
- Potential interference from other secondaries is a very signifi cant uncertainty, but can be mitigated through mandated local cooperation among systems.
- *Non-interference is a system-level, rather than device-level, property and must be regulated as such.*

Summary of main points

- Uncertainty imposes limits on device sensitivity that can *not* be overcome by just listening longer.
- Shadowing is a major challenge but can be overcome with multi-user cooperative diversity within a system.
- Potential interference from other secondaries is a very signifi cant uncertainty, but can be mitigated through mandated local cooperation among systems.
- *Non-interference is a system-level, rather than device-level, property and must be regulated as such.*
- In bands with high-powered primaries, long-range/high-power secondary use is possible, but *with power comes responsibility.*
  - Sense more carefully
  - Cooperate more among systems

Summary of main points

- Uncertainty imposes limits on device sensitivity that can *not* be overcome by just listening longer.
- Shadowing is a major challenge but can be overcome with multi-user cooperative diversity within a system.
- Potential interference from other secondaries is a very signifi cant uncertainty, but can be mitigated through mandated local cooperation among systems.
- *Non-interference is a system-level, rather than device-level, property and must be regulated as such.*
- In bands with high-powered primaries, long-range/high-power secondary use is possible, but *with power comes responsibility.*
  - Sense more carefully
  - Cooperate more among systems
- Complexity can buy some freedom.

Primary transmitter's decodability radius

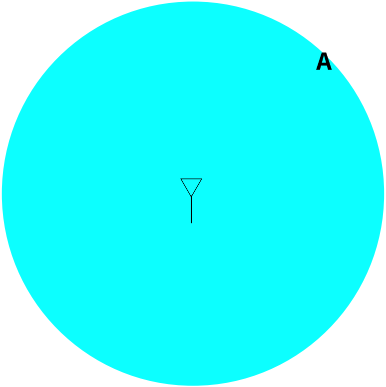

We define a protected radius

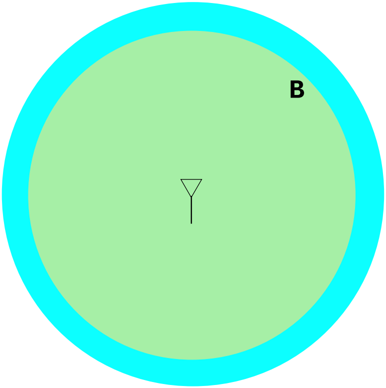

Mice can get close...

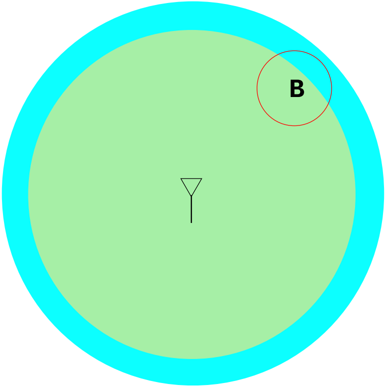

But keep the lions far away!

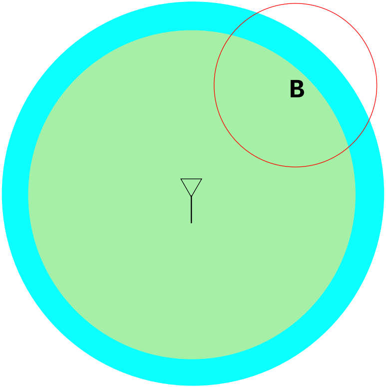

We can't protect everyone

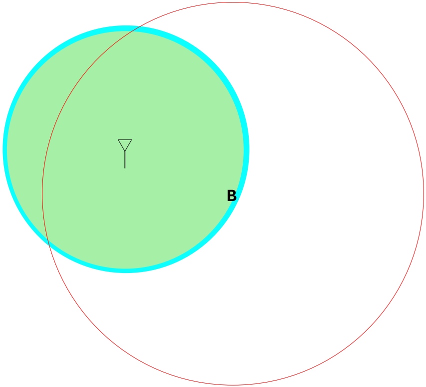

Demanding complete protection for marginal legacy users will cripple innovation.

Union of “no talk” zones

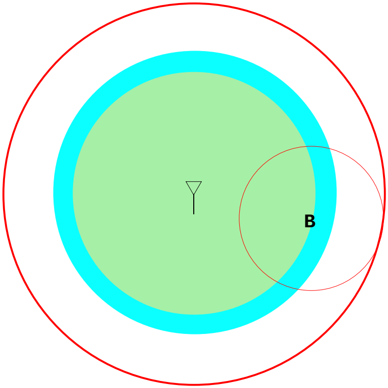

Fading

A secondary user might be faded while his transmissions could still reach an unfaded primary receiver.

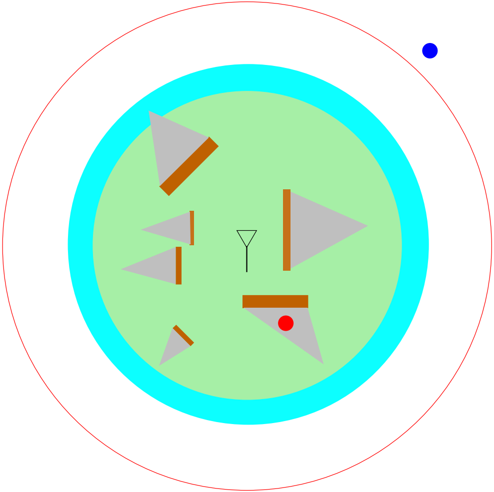

Fading margin

A secondary user who can not distinguish between positions must be quiet in both.

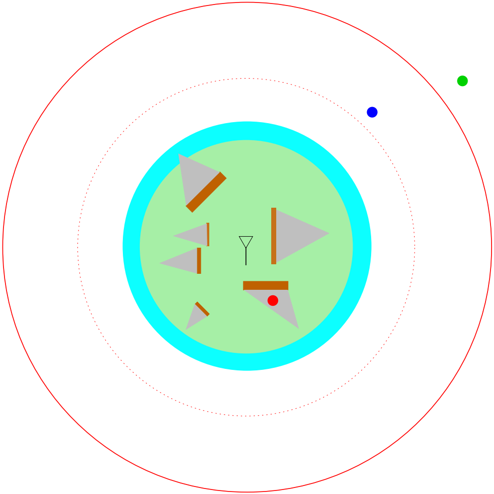

What are we giving up?

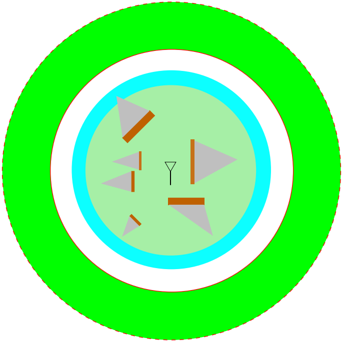

Safe, but might be faded
(fading uncertainty)

What are we giving up?

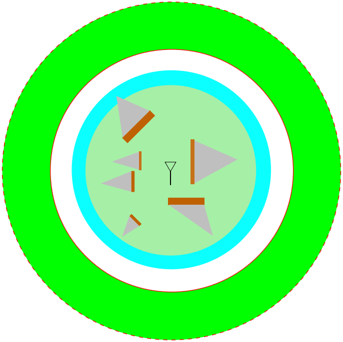

Safe, but might be faded
(fading uncertainty)

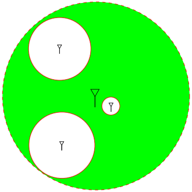

Lights on, but no one home
(receiver uncertainty)

What are we giving up?

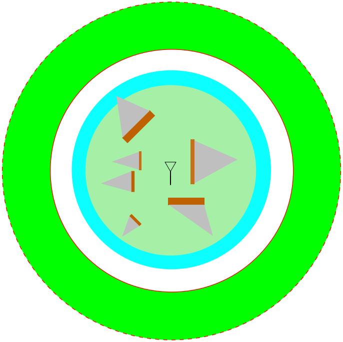

Safe, but might be faded
(fading uncertainty)

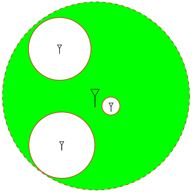

Lights on, but no one home
(receiver uncertainty)

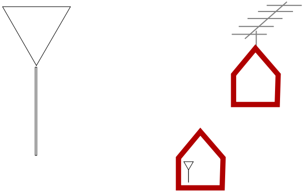

Safe, but not shadowed enough
(symmetry uncertainty)

How valuable is the real estate?

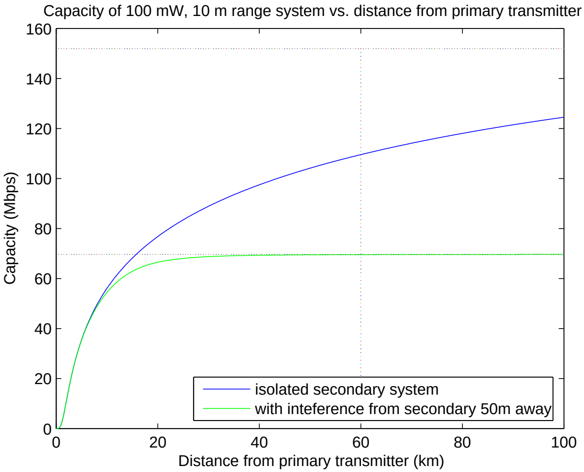

- 100 kW primary transmitter, 60 km decodable radius
- Primary to secondary decays as $r^{-3.5}$
- Secondary to secondary decays as r^{-5}

Impact of uncertainty: limit on sensitivity

All detectors are based on some averaged test statistic

- If value is high, primary is considered present
  If the value is low, it is considered absent.
- Increasing the amount of averaging, increases our confidence in the decision.
- Unmodeled or non-ergodic uncertainties introduce unresolvable ambiguities.

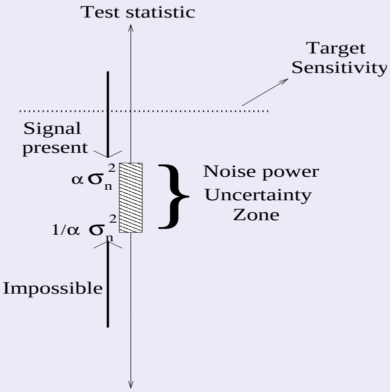

Overcoming fading: multiuser diversity

Cooperation within a *geographically distributed* network of secondary users avoids dealing with the worst-cases of fading.

$\boxed{\text{Detect the primary collectively!}}$

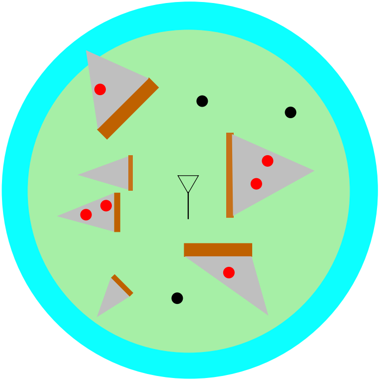

New concerns with multiple users

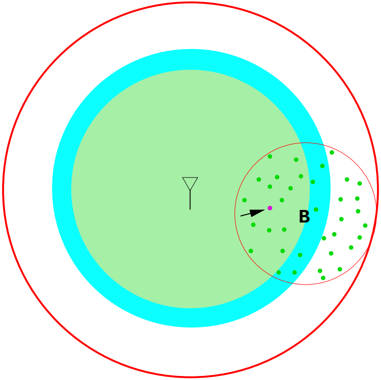

More potential secondary users to keep quiet in the near vicinity of primary receivers.

The challenge of aggregate interference

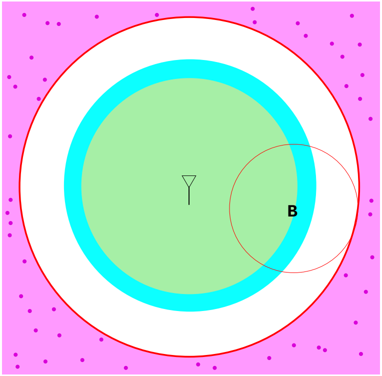

One secondary user might be acceptable, but what about millions?

- Limit on the power density
- Slower effective attenuation with distance

Overcoming multiuser uncertainty: “sensing MAC” among systems

- Secondaries must not confuse uncertain aggregate interference with primary signal
- Nearby secondary users must be quiet during detection
- Increasing density of secondary transmissions requires more cooperation

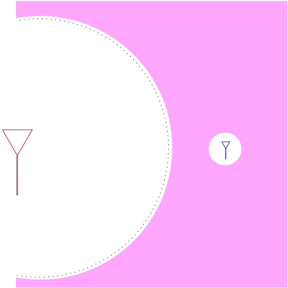

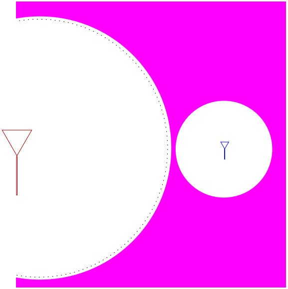

Outline

- An overview of the issues involved some key ideas
- **Part I: The basic considerations quantified**
  - **The "sensing link budget"**
  - The energy detector and fundamental limits on its sensitivity
  - Within-system cooperation
  - Fairness and cooperation among systems
- Part II: More powerful detectors
  - Coherent detectors
  - Feature detectors
- Part III: Hardware considerations
  - Fundamental hardware limitations
  - Ways around them

The classical link budget

- Goal: reliable communication
  - At a given range
  - within a given bandwidth

$$
\begin{aligned}
SNR_r = P_t + G_{antenna} + G_{coding} \\
- L_{free-space} - L_{atmospheric} - L_{shadow} - L_{multipath} \\
- P_{noise} \\
- P_{interference}
\end{aligned}
$$

- Fundamental limits
  - Coding gain bounded by channel capacity
  - Additional margins needed to deal with uncertainty
- **“Typical” case: more than enough SNR.**

Opportunistic use: new considerations

- New constraint: *must not interfere with primary users*
  - Allowable P_{HI} ($\sim 1\%$) at protected primary receivers
  - Limit on tolerable out-of-system interference
  - Translates to a limit on maximum power, not minimum.
- New question: What signal power must we detect to guarantee non-interference?

Opportunistic use: new considerations

- New constraint: *must not interfere with primary users*
  - Allowable P_{HI} ($\sim 1\%$) at protected primary receivers
  - Limit on tolerable out-of-system interference
  - Translates to a limit on maximum power, not minimum.
- New question: What signal power must we detect to guarantee non-interference?
- New issues
  - Different propagation path losses
    - Between primary and secondary users
    - Among secondary users
  - Multiple opportunistic users
  - Heterogeneous transmit powers

Opportunistic use: new considerations

- New constraint: *must not interfere with primary users*
  - Allowable P_{HI} ($\sim 1\%$) at protected primary receivers
  - Limit on tolerable out-of-system interference
  - Translates to a limit on maximum power, not minimum.
- New question: What signal power must we detect to guarantee non-interference?
- New issues
  - Different propagation path losses
    - Between primary and secondary users
    - Among secondary users
  - Multiple opportunistic users
  - Heterogeneous transmit powers
- **"Typical" case: primary signal is absent or weak — very low SNR.**

Outline of section

- Formalization with path losses only
- Basic case: 1 primary, 1 secondary
- Multiple secondary transmitters
- The “outage view” of multipath and shadowing
- Numerical example

Heterogeneous path-losses

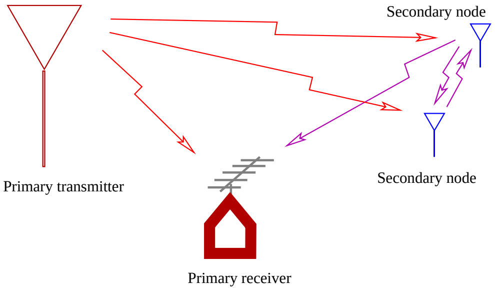

- Free-space propagation: d^{-2}
- Ground reflection: d^{-4}
- Absorption term: $e^{-\lambda d}$
- Fit to empirical data: $d^{-\alpha}$
- Antenna height (impacts constants)
- Urban/rural (impacts α)
- Indoor/outdoor
- Receiver placement

The Case of the Single Secondary Transmitter

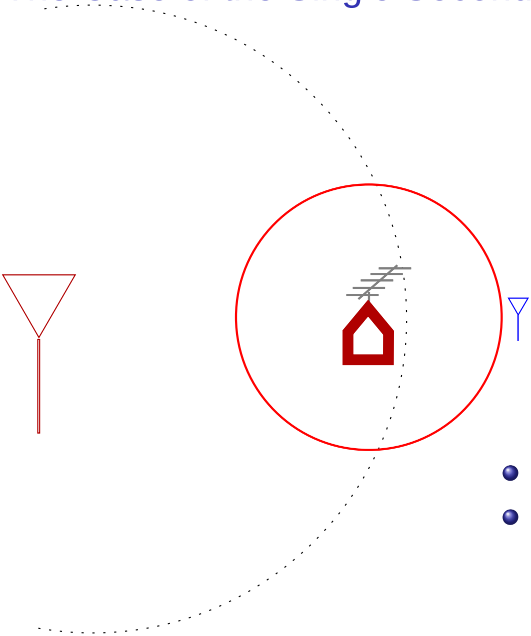

- But we don't know where we are!
- Think in terms of distances, but use local signal strength.

The fundamental constraint

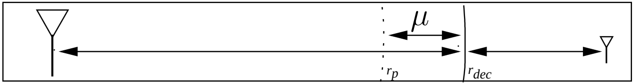

- μ determines how much interference above the noise floor the primary system can tolerate

$$Q_2 + \sigma^2 \leq \sigma^2 10^{\frac{\mu}{10}}$$

- The secondary system must guarantee:

$$\boxed{Q_2 \leq (10^{\frac{\mu}{10}} - 1)\sigma^2}$$

Q_2: (aggregate) received secondary transmitters' powers at primary receiver

μ: dB margin of protection

r_p: protected radius

r_{dec}: decodable radius for a primary receiver

Solo secondary sensing link budget

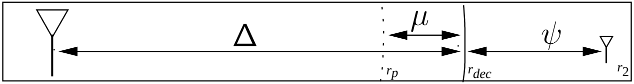

- At the secondary, the primary signal has dropped by $\Delta_s + \psi$.
- At the primary receiver, the secondary's transmission has been attenuated by $g_{21}(r_2 - r_p)$.
- Sensitivity required vs. desired secondary power

$$
\begin{aligned}
-(\psi + \Delta_s) &\geq 10\log_{10}\left[g_{12}(r_2)\right] \\
&= 10\log_{10}\left[g_{12}(r_p + (r_2 - r_p))\right] \\
&= 10\log_{10}\left[g_{12}\left({g_{11}}^{-1}\left(10^{\frac{\mu-\Delta_p}{10}}\right) + {g_{21}}^{-1}\left(\frac{(10^{\frac{\mu}{10}} - 1)\cdot\sigma^2}{P_2}\right)\right)\right]
\end{aligned}
$$

Q_2: (aggregate) received secondary transmitters' powers at primary receiver
μ: dB margin of protection

r_p: protected radius

r_{dec}: decodable radius for a primary receiver

ψ: how much weaker than the minimal decodable signal the secondary's reception is.

$\Delta_{[p,s]}$: Signal attenuation between the primary transmitter and r_{dec}, as measured by a **p**rimary or **s**econdary.

r_2: distance from primary transmitter to secondary transmitter

Single transmitter power constraints

$(g_{11}(r) = g_{12}(r) = r^{-\alpha_1}$ and $g_{21}(r) = r^{-\alpha_2}$, $\mu = 1$ dB margin)

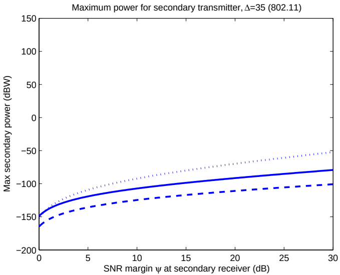

(a) $\Delta = 35$ $(r_{dec} = 10\text{m})$

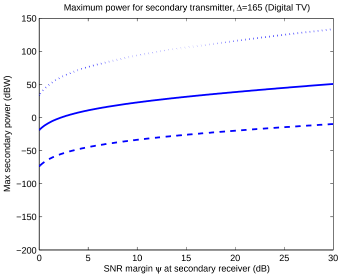

(b) $\Delta = 165$ $(r_{dec} = 51\text{km})$

·········	$\alpha_1 = 3.5, \ \alpha_2 = 5$	Secondary signal attenuates faster than primary
———	$\alpha_1 = 3.5, \ \alpha_2 = 3.5$	Secondary signal attenuates at same rate as primary
– – –	$\alpha_1 = 5, \ \alpha_2 = 3.5$	Secondary signal attenuates slower than primary

Fading

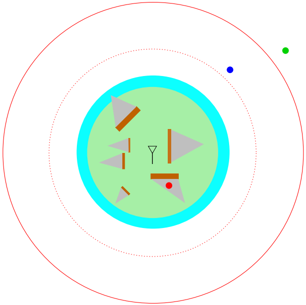

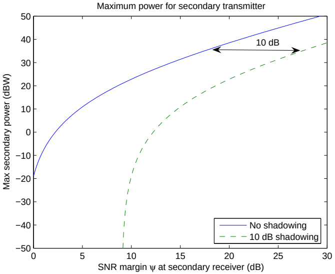

- If you hear a weak signal, are you far away, or just faded?
- The possibility of 10 dB of fading results in a 10 dB shift of the required detection margin

Multiple secondaries – What could go wrong?

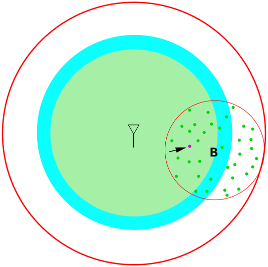

- “Keeping K four-year olds quiet”
- $P_{md} \leq \frac{1}{K} P_{HI}$

Multiple secondaries – What could go wrong?

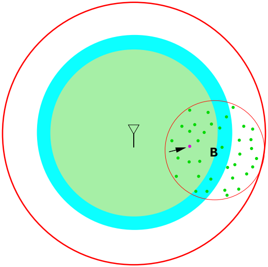

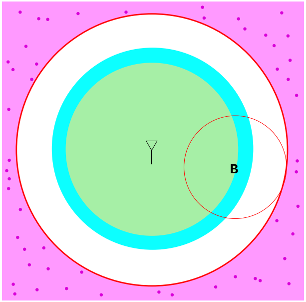

- "Keeping K four-year olds quiet"
- $P_{md} \leq \frac{1}{K} P_{HI}$

- "Conversation in a crowded restaurant"
  - Each individual transmitter talks quietly
  - But the aggregate chatter may be overwhelming

“Death by a thousand cuts” analysis

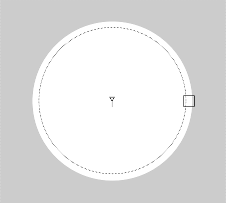

- Assume secondaries have limited power, distinct footprints (Data MAC protocol)
- Approximate sea of secondary users by a power density
- Circular coast looks like a straight line nearby
- Integration changes the decay exponent
  - As you move further from the coast you “see” more interferers
  - Physical $r^{-4} \rightarrow r^{-2}$ effective.

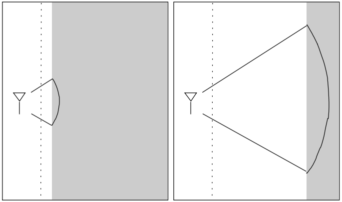

$$Q_2 = \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \int_{\frac{r_n - r_p}{\cos(\theta)}}^{\infty} D r^{-\alpha_2} \, r \, dr \, d\theta$$

$$= D \cdot K(\alpha_2) \cdot (r_n - r_p)^{-\alpha_2 + 2}$$

where $K(\alpha_2) = \frac{\int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} (\cos\theta)^{\alpha_2 - 2} \, d\theta}{\alpha_2 - 2}$.

For $\alpha_2 = 6$, $K(\alpha_2) = \frac{1}{4} \frac{3!!}{4!!} \pi \approx 0.295$.

Heterogeneity and competing interests

- Primary receiver has a certain margin μ of tolerable interference
- We must *choose* how to allocate this margin to users at different distances
- Far away users can gain at the expense of nearby users

Policy input required

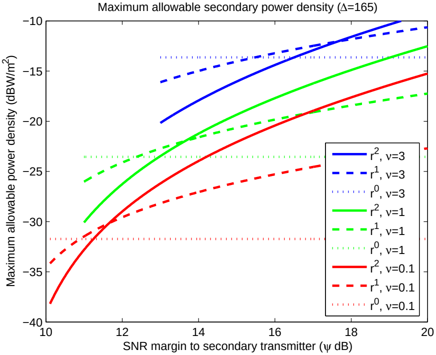

- Primary signal decays as $r^{-3.5}$
- Secondary signals decay as r^{-5}
- ν: Additional "quiet" margin to allow more power.
  - 3 dB $\approx$ 11 km
  - 1 dB $\approx$ 3.5 km
  - 0.1 dB $\approx$ 0.34 km

Fading revisited

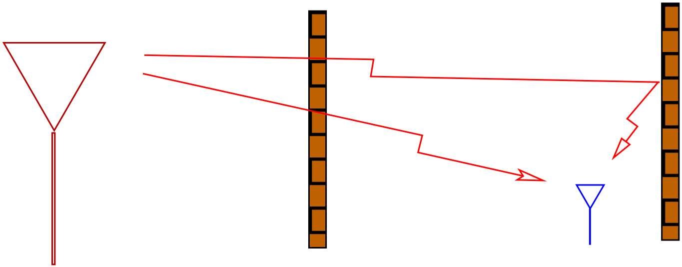

Is there a principled way of choosing X dB of fading margin?

Multipath and shadowing

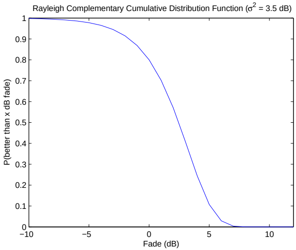

Rayleigh fading model

- Many independent scatterers
- Magnitude a Rayleigh random variable
- If primary signal is wideband, could be frequency selective.
- Could hurt or help.

Multipath and shadowing

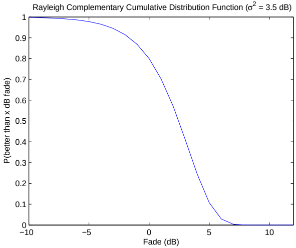

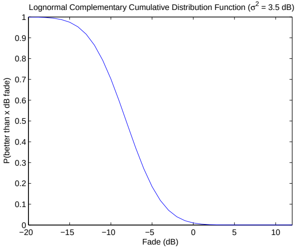

Rayleigh fading model

- Many independent scatterers
- Magnitude a Rayleigh random variable
- If primary signal is wideband, could be frequency selective.
- Could hurt or help.

Lognormal shadowing

- Large number of small absorptive losses
- Central limit theorem
- Not really frequency selective
- *Can only hurt.*

The key role of P_{md}: the "outage view"

$$P_{md} = P(\text{good fade}) \cdot P(\text{missing a good signal}) + P(\text{bad fade})$$

- Our ability to detect is limited by the probability of a bad fade

The key role of P_{md}: the "outage view"

$$P_{md} = P(\text{good fade}) \cdot P(\text{missing a good signal}) + P(\text{bad fade})$$

- Our ability to detect is limited by the probability of a bad fade

P_{fade}	fade (dB)
10%	-16 dB
1%	-21 dB
0.1%	-26 dB
0.01%	-31 dB

(Assuming lognormal+Rayleigh fading, σ =3.5 dB each)

The key role of P_{md}: the "outage view"

$$P_{md} = P(\text{good fade}) \cdot P(\text{missing a good signal}) + P(\text{bad fade})$$

- Our ability to detect is limited by the probability of a bad fade

P_{fade}	fade (dB)
10%	-16 dB
1%	-21 dB
0.1%	-26 dB
0.01%	-31 dB

(Assuming lognormal+Rayleigh fading, σ =3.5 dB each)

- *But deep fade "probabilities" are uncertain and poorly modeled.*

Sensing budget evaluated

Protocol	Tx power	Footprint	Density (W/m^2)	Add'l Sensitivity	With fading
"WiMax"	1 W	1 km^2	$1 \cdot 10^{-6}$	-0.17 dB	-31.17 dB
"Bluetooth"	2.5 mW	20 m^2	$1.3 \cdot 10^{-4}$	-1.23 dB	-32.23 dB
"WiFi"	100 mW	300 m^2	$3.3 \cdot 10^{-4}$	-1.68 dB	-32.68 dB

- *Fading is the dominant term.*
- If 23dB SNR at decodable radius, need to robustly detect at least 8 dB below noise floor.
- Is this possible?

Outline

- An overview of the issues involved and some key ideas
- Part I: The basic considerations quantified
  - The "sensing link budget"
  - **The energy detector and fundamental limits on its sensitivity**
  - Within-system cooperation
  - Fairness and cooperation among systems
- Part II: More powerful detectors
  - Coherent detector
  - Feature detectors
- Part III: Hardware considerations
  - Fundamental hardware limitations
  - Ways around them

Big question: sensing the primary

Spectrum picture

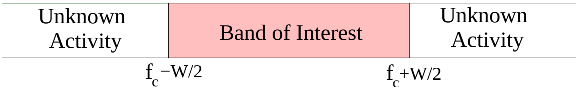

- Look for the primary in the ‘band of interest’
- Within band model:
  - White or unknown signal, $X(t)$
  - Independent white noise, $W(t)$
- First step: sample the band of interest at Nyquist rate

Hypothesis testing problem formulation

- Distinguish between the following hypotheses:

$$
\begin{aligned}
\mathcal{H}_0 : Y[n] &= W[n] \\
\mathcal{H}_1 : Y[n] &= W[n] + X[n]
\end{aligned}
$$

- Basic assumptions:
  - $X[n]$'s are i.i.d. signal samples
  - $W[n]$'s are i.i.d. noise samples
- Target error probabilities:
  - P_{FA}: Probability of false alarm
  - P_{MD}: Probability of missed detection
- Key Resource: *Dwell Time N of the detector*

Performance measures for detection

- Dwell time of the detector
  - Limited by primary duty cycle, sharing with others, etc.
  - Sample complexity of detectors

Definition

Sample complexity captures how N varies with SNR for given P_{FA} and P_{MD}

Performance measures for detection

- Dwell time of the detector
  - Limited by primary duty cycle, sharing with others, etc.
  - Sample complexity of detectors

Definition

Sample complexity captures how N varies with SNR for given P_{FA} and P_{MD}

- Robustness
  - What uncertainties are unavoidable
  - How does the detector perform despite the uncertainty

Performance measures for detection

- Dwell time of the detector
  - Limited by primary duty cycle, sharing with others, etc.
  - Sample complexity of detectors

Definition

Sample complexity captures how N varies with SNR for given P_{FA} and P_{MD}

- Robustness
  - What uncertainties are unavoidable
  - How does the detector perform despite the uncertainty
- Computational/Implementational complexity

Simplest detector: Energy detector

- Recall:
  - Signal, $X[n]$ could be anything, including white
  - Noise, $W[n]$ is white Gaussian
- Received energy used for detection
- Test statistic:

$$T(\mathbf{y}) = \sum_{n=1}^{N} Y^2[n]$$

- Decision rule:

$$T(\mathbf{y}) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma(N)$$

Error probability analysis

False alarm probability

$$\begin{aligned} P_{FA} &= P(T(\mathbf{y}) > \gamma | \mathcal{H}_0) \\ &= P\left(\frac{\frac{T(\mathbf{y})}{\sigma^2} - N}{\sqrt{2N}} > \frac{\frac{\gamma}{\sigma^2} - N}{\sqrt{2N}} \right) \\ &\approx \mathcal{Q}\left(\frac{\frac{\gamma}{\sigma^2} - N}{\sqrt{2N}} \right) \end{aligned}$$

Probability of detection

$$\begin{aligned} P_{D} &= P(T(\mathbf{y}) > \gamma | \mathcal{H}_1) \\ &= P\left(\frac{\frac{T(\mathbf{y})}{\sigma^2} - \lambda - N}{\sqrt{4\lambda + 2N}} > \frac{\frac{\gamma}{\sigma^2} - \lambda - N}{\sqrt{4\lambda + 2N}} \right) \\ &\approx \mathcal{Q}\left(\frac{\frac{\gamma}{\sigma^2} - \lambda - N}{\sqrt{4\lambda + 2N}} \right) \end{aligned}$$

- Receiver does not know signal power λ
  - Only has knowledge of σ^2
  - Set threshold γ based on P_{FA}
- Evaluate P_{MD} for different values of λ to get sensitivity.
- Eliminate γ to get sample complexity:

$$N \approx 2\left[\mathcal{Q}^{-1}(P_{FA}) - \mathcal{Q}^{-1}(P_D)\right]^2 SNR^{-2}$$

Understanding robustness: noise uncertainty

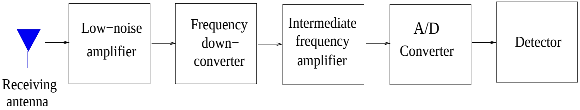

- Noise is usually assumed to be Gaussian
- Sources of uncertainty:
  - Non-linearity of components
  - Thermal noise in components (Non-uniform, time-varying)

Understanding robustness: noise uncertainty

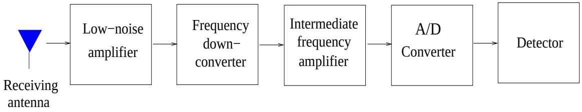

- Noise is usually assumed to be Gaussian
- Sources of uncertainty:
  - Non-linearity of components
  - Thermal noise in components (Non-uniform, time-varying)
  - Noise due to transmissions by other users
    - Unintentional (Close-by)
    - Intentional (Far-away)

Understanding robustness: noise uncertainty

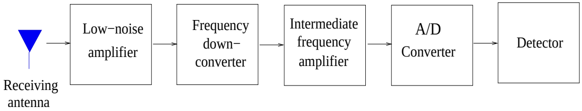

- Noise is usually assumed to be Gaussian
- Sources of uncertainty:
  - Non-linearity of components
  - Thermal noise in components (Non-uniform, time-varying)
  - Noise due to transmissions by other users
    - Unintentional (Close-by)
    - Intentional (Far-away)
    - Opportunistic

Impact of uncertainty: energy detector

- Actual noise power, $\sigma_a^2 \in [\frac{1}{\alpha}\sigma_n^2, \alpha\sigma_n^2]$
- If

$$\begin{aligned} P + \sigma_a^2 &\leq \alpha\sigma_n^2 \\ \Rightarrow P &\leq \frac{\alpha^2 - 1}{\alpha}\sigma_n^2 \end{aligned}$$

Energy detector fails to detect the signal

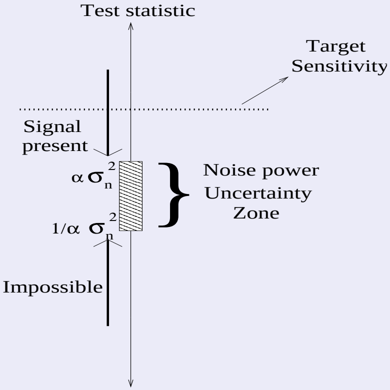

SNR wall for energy detector

- SNR_{wall} for the energy detector is given by

$$SNR_{wall} = 10 \log_{10} \left(\frac{\alpha^2 - 1}{\alpha} \right)$$

where, $\alpha = 10^{(x/10)}$

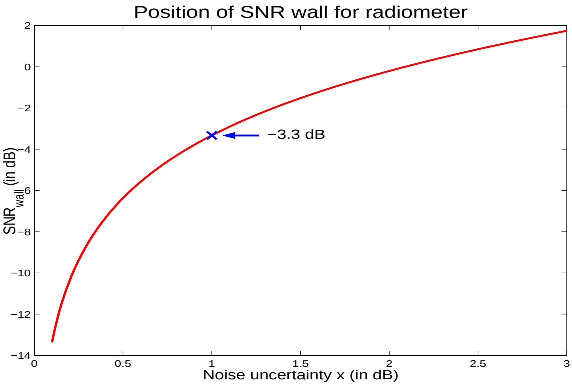

Realistic model for noise uncertainty

Even if we calibrate, there is always residual uncertainty about the noise.

- Receiver knows the noise distribution only up to an uncertainty set $\mathcal{W}_x$.

Realistic model

- Noise cloud includes a range of Gaussians with variance $\in [\frac{1}{\alpha}\sigma_n^2, \alpha\sigma_n^2]$ as well as other similar distributions.
- $W_a \in \widetilde{\mathcal{W}}_x$ iff

$$\mathbb{E}W_a^{2k} \in \left[\frac{1}{\alpha^k}\mathbb{E}W_n^{2k}, \alpha^k\,\mathbb{E}W_n^{2k}\right], \qquad \alpha = 10^{x/10}$$

- Implication: Energy detector like wall for all detectors

Theorem

Consider detecting a weak BPSK signal with the noise distribution lying in $\widetilde{\mathcal{W}}_x$. Under this model, there exists an absolute *SNR* wall *(snr^*_{wall}) for any possible robust detector.*

$$snr^*_{wall} = \min_{k>0} snr^{(2k)}_{wall} = \alpha - 1$$

Example of noise distribution overlap

Signal looks like noise: $f_{W_a+X} = f_{W_n}$

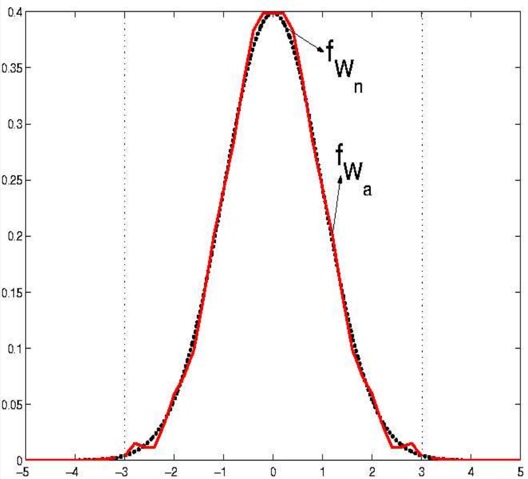

Position of SNR wall
Realistic model

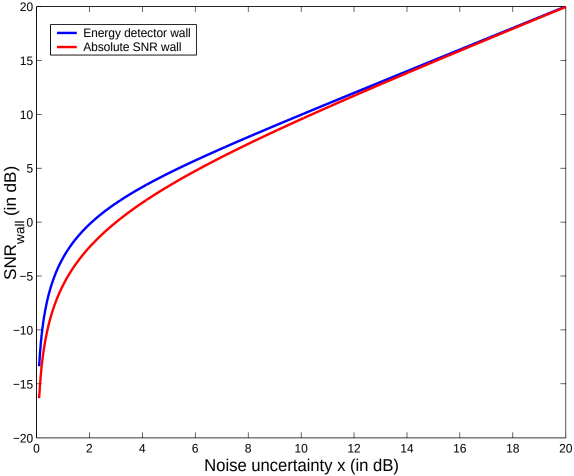

Figure: SNR^*_{wall} as a function of x

Conclusions

- Noise uncertainty is always present
  - At least about 1-2 dB of device level uncertainty
    - Cannot see 3 dB below the noise
  - Easily 10-20 dB of interference level uncertainty
    - Cannot see below it at all!
- E.g.: In a 6MHz TV band, Digital TV receiver sensitivity = -85dBm.
  - We must have -117dBm sensitivity to deal with rare fading.
  - Thermal Noise at -106dBm !
- What can we do to mitigate this?

Outline

- An overview of the issues involved and some key ideas
- Part I: The basic considerations quantified
  - The “sensing link budget”
  - The energy detector and fundamental limits on its sensitivity
  - **Within-system cooperation**
  - Fairness and cooperation among systems
- Part II: More powerful detectors
  - Coherent detectors
  - Feature detectors
- Part III: Hardware considerations
  - Fundamental hardware limitations
  - Ways around them

How can cooperation help?

Fading is the dominant challenge

- Multipath varies significantly on the scale of $\frac{\lambda}{4}$ (10cm at 800MHz).
- Shadowing varies significantly on the scale 20-500m
- *Use multiple radios as a proxy for multiple antennas!*

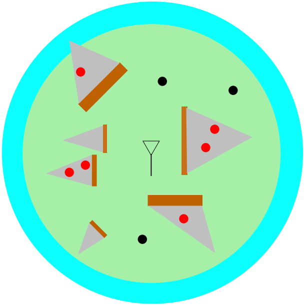

How can cooperation help?

Fading is the dominant challenge

- Multipath varies significantly on the scale of $\frac{\lambda}{4}$ (10cm at 800MHz).
- Shadowing varies significantly on the scale 20-500m
- *Use multiple radios as a proxy for multiple antennas!*
- *Analogy:* Deck of cards where red cards signify bad fades.
  - Probability that I get a red card: Very High (50%)!
  - Probability that all users get red cards: Very Low

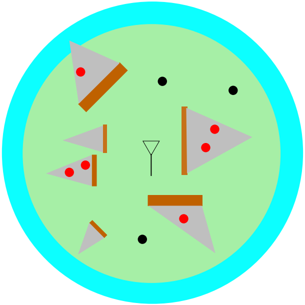

Cooperative diversity quantified

- If $P_{HI} = 1\%, K = 100$, then $P_{MD,system} = 0.01\%$!
- What if our system had many (M) independent radios?

$$P_{D,radio} \leq 1 - \sqrt[M]{\frac{P_{HI}}{K}}$$

  - For $M = 10$, $P_{D,radio} = 60\%$ Just have to work in the best 60% cases.
- Robustness comes from being able to write off the worst possible fades. No longer need to model them as precisely!

How much does cooperation buy us?

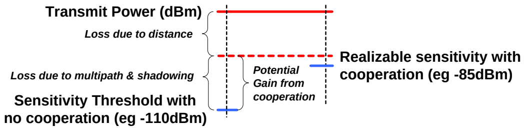

- Cooperation helps us approach the nominal distance dependent path loss

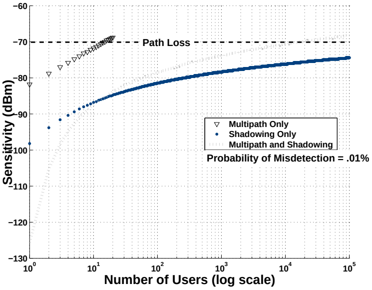

- Cooperation model: Each user sends a 1 bit decision to a controller
- 10 - 20 users are needed to obtain realistic receiver sensitivity levels.

Beware of Multipath Gains!

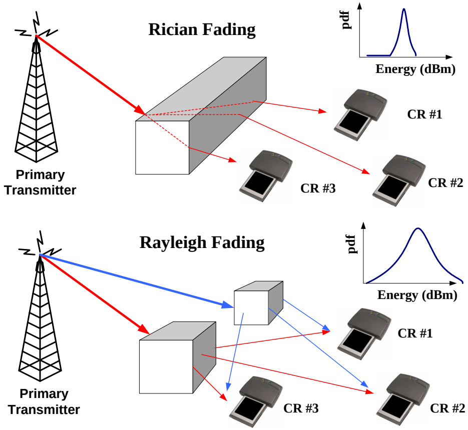

- Cannot rely on multipath gains - might have a single weak path
- Cooperation *should* only be used to mitigate bad multipath

What are the limits on cooperation?

- Complexity of getting everyone on board
  - Control channel bandwidth may be limited during the setup stage
  - Delay in relaying decisions may make decisions obsolete.

What are the limits on cooperation?

- Complexity of getting everyone on board
  - Control channel bandwidth may be limited during the setup stage
  - Delay in relaying decisions may make decisions obsolete.
- Independence issues
  - How does correlation effect cooperation?
  - How to quantify independence?

What are the limits on cooperation?

- Complexity of getting everyone on board
  - Control channel bandwidth may be limited during the setup stage
  - Delay in relaying decisions may make decisions obsolete.
- Independence issues
  - How does correlation effect cooperation?
  - How to quantify independence?
- Trust issues
  - What if users lie about sensing decisions?
  - What if radios fail in unknown ways and/or are malicious?

Dealing with channel correlation

- Multipath is not correlated
  - Radio placements on the scale of wavelength are essentially random.
- Shadowing is correlated if two radios are blocked by the same obstacle
- One model: Correlation decays exponentially with distance.

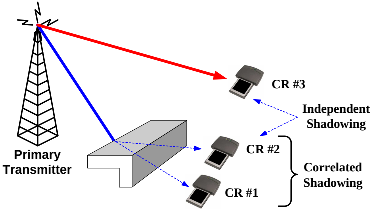

The few, the independent, . . .

- Can a large number of correlated users make up for lack of independence ?
  No !!

- It is better to increase distance spread of users than to increase the count.

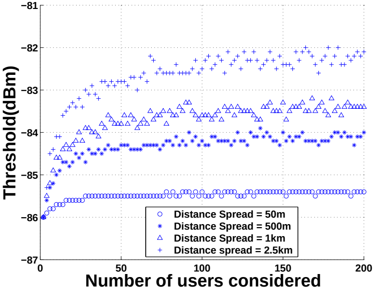

The few, the independent, . . .

- Can a large number of correlated users make up for lack of independence ?
  No !!
- It is better to increase distance spread of users than to increase the count.

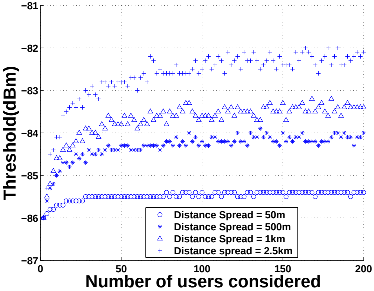

Need long range cooperation within opportunistic systems.

Known failures are easy

- Some users may not contribute to sensing.
- Such *known failures* just reduce the effective number of users in the system

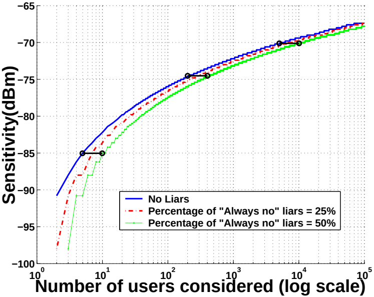

Dealing with unpredictable adversaries

- Malicious adversaries are impossible to predict reliably — need to budget for worst case.
- Assume M users with a known fraction α behaving unpredictably.

Dealing with unpredictable adversaries

- Malicious adversaries are impossible to predict reliably — need to budget for worst case.
- Assume M users with a known fraction α behaving unpredictably.
- P_{FA} : Set threshold at βM ($\beta > \alpha$) — declare Primary present only if βM nodes say *Yes*.

Dealing with unpredictable adversaries

- Malicious adversaries are impossible to predict reliably — need to budget for worst case.
- Assume M users with a known fraction α behaving unpredictably.
- P_{FA} : Set threshold at βM ($\beta > \alpha$) — declare Primary present only if βM nodes say *Yes*.
- P_{MD} : What if the adversaries now behave as *Always No* liars?
    - Reduces actual users to $M(1-\alpha)$ of which βM must declare *Yes*.
    - For $\frac{\beta}{1-\alpha}$ of them to detect, threshold must be such that $P_{D,radio} \geq \frac{\beta}{1-\alpha}$.

Dealing with unpredictable adversaries

- Malicious adversaries are impossible to predict reliably — need to budget for worst case.
- Assume M users with a known fraction α behaving unpredictably.
- P_{FA} : Set threshold at βM ($\beta > \alpha$) — declare Primary present only if βM nodes say *Yes*.
- P_{MD} : What if the adversaries now behave as *Always No* liars?
  - Reduces actual users to $M(1-\alpha)$ of which βM must declare *Yes*.
  - For $\frac{\beta}{1-\alpha}$ of them to detect, threshold must be such that $P_{D,radio} \geq \frac{\beta}{1-\alpha}$.

Distrust limits diversity gains

The impact of distrust

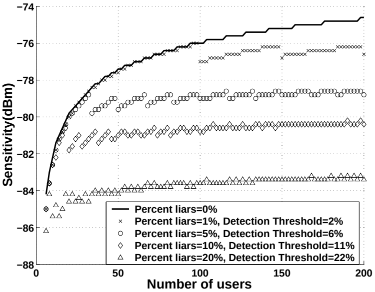

- Diversity gains with α fraction untrusted users are bounded by those achievable by a trusted population of $\frac{1}{\alpha}$ trusted users.
- To achieve these gains, we need M » $\frac{1}{\alpha}$ when α large.

Within-system cooperation summary

- Low-moderate fading is all that can be realistically modeled.
- Cooperation allows independent radios to target individual sensitivity levels based on low-moderate fading margins while maintaining system robustness.
- We prefer a few distant users to many nearby users.
- Untrusted radios introduce a bound on achievable sensitivity reductions.

Outline

- An overview of the issues involved and some key ideas
- Part I: The basic considerations
  - The “sensing link budget”
  - Limits on sensitivity for an energy detector
  - Within-system cooperation
  - **Fairness and among-system cooperation**
- Part II: More powerful detectors
  - Coherent detectors
  - Feature detectors
- Part III: Hardware considerations
  - Fundamental hardware limitations
  - Ways around them

The story so far

- Opportunistic use can be high-power with long range, as long as the power-density is controlled appropriately.
- Opportunism requires a system, not a device, in order to deal with fading.
  - *Needs to be regulated as a system.*
  - Has an internal incentive to cooperate with trusted nodes on a larger geographical scale.
  - Has a mild disincentive to collaborate with non-trusted nodes nearby.
- So far, only considered secondary-to-primary interference.

Conceptualizing secondary-to-secondary interference

- Secondary-to-secondary interference complicates detection
  - Radiometer cannot distinguish between secondary signals and the primary signal
  - Doesn't know if secondaries from other systems are present nearby.

Conceptualizing secondary-to-secondary interference

- Secondary-to-secondary interference complicates detection
  - Radiometer cannot distinguish between secondary signals and the primary signal
  - Doesn't know if secondaries from other systems are present nearby.
- View 1: Fairness
  - Don't want the first opportunistic user to exclude all other secondaries
  - Possible options
    - Mandate honest sharing of detection results.

Conceptualizing secondary-to-secondary interference

- Secondary-to-secondary interference complicates detection
  - Radiometer cannot distinguish between secondary signals and the primary signal
  - Doesn't know if secondaries from other systems are present nearby.
- View 1: Fairness
  - Don't want the first opportunistic user to exclude all other secondaries
  - Possible options
    - ★ Mandate honest sharing of detection results.
    - ★ Have everyone nearby be quiet during sensing.

Conceptualizing secondary-to-secondary interference

- Secondary-to-secondary interference complicates detection
  - Radiometer cannot distinguish between secondary signals and the primary signal
  - Doesn't know if secondaries from other systems are present nearby.
- View 1: Fairness
  - Don't want the first opportunistic user to exclude all other secondaries
  - Possible options
    - Mandate honest sharing of detection results.
    - Have everyone nearby be quiet during sensing.
- View 2: Maximize steady state utilization
  - Maintain good utilization as systems hop around.
  - No need to have overly interference-free bands when no primary users are around.

Mandated "sensing MAC" among systems

- During sensing, secondary-to-secondary interference is uncertain and induces an SNR_{wall}
  - Problematic because we are *uncertain* how many secondaries are talking.
  - Higher permitted densities induce more uncertainty.
  - Wall must be kept below required sensitivity.

Mandated "sensing MAC" among systems

- During sensing, secondary-to-secondary interference is uncertain and induces an SNR_{wall}
  - Problematic because we are *uncertain* how many secondaries are talking.
  - Higher permitted densities induce more uncertainty.
  - Wall must be kept below required sensitivity.
- Keep nearby secondary users quiet during detection.

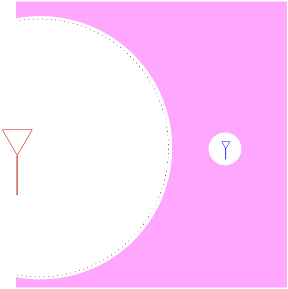

Mandated "sensing MAC" among systems

- During sensing, secondary-to-secondary interference is uncertain and induces an SNR_{wall}
  - Problematic because we are *uncertain* how many secondaries are talking.
  - Higher permitted densities induce more uncertainty.
  - Wall must be kept below required sensitivity.
- Keep nearby secondary users quiet during detection.
- Increasing density of secondary transmissions requires more cooperation.

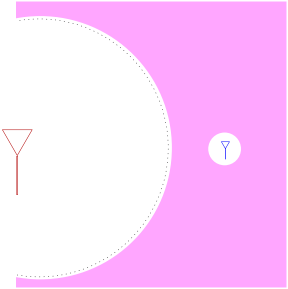

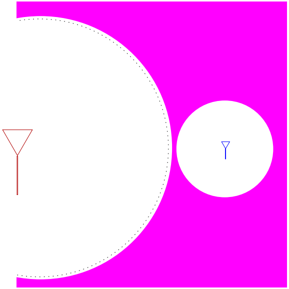

How to evaluate power / cooperation tradeoffs

- Must satisfy two constraints
  - Non-interference to primary receivers
  - Fairness among secondary users

How to evaluate power / cooperation tradeoffs

- Must satisfy two constraints
  - Non-interference to primary receivers
  - Fairness among secondary users

1. Use non-interference due to aggregate interference to pick no-talk radius for a given power density D.

How to evaluate power / cooperation tradeoffs

- Must satisfy two constraints
  - Non-interference to primary receivers
  - Fairness among secondary users

1. Use non-interference due to aggregate interference to pick no-talk radius for a given power density D.
2. Use fading-margin and primary-to-secondary attenuation to pick a target detection level.

How to evaluate power / cooperation tradeoffs

- Must satisfy two constraints
  - Non-interference to primary receivers
  - Fairness among secondary users

1. Use non-interference due to aggregate interference to pick no-talk radius for a given power density D.
2. Use fading-margin and primary-to-secondary attenuation to pick a target detection level.
3. Set "shut-up" radius r_s to reduce uncertainty $I_{max} = D\frac{2\pi}{\alpha_{22}-2} {r_s}^{2-\alpha_{22}}$ enough to allow robust detection at target level.

Power / cooperation tradeoffs for $\alpha_{12} = -3.5$

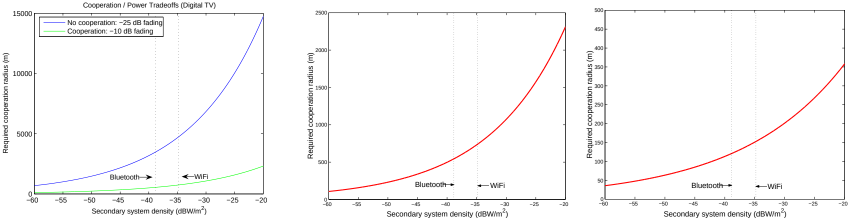

(a) $\alpha_{22} = -5$ (b) $\alpha_{22} = -5$ zoom (c) $\alpha_{22} = -6$

- "Double whammy" with increased opportunistic power density.
  - Must detect weaker primary
  - Increased interference uncertainty
- The aggregate interference from all potential secondaries must be weaker than the weak primary signal.
- Requires among-system coordination across a large local area, even after within-system cooperation has reduced fading margins.

Outline

- An overview of the issues involved and some key ideas
- Part I: The basic considerations
  - The "sensing link budget"
  - Limits on sensitivity for an energy detector
  - Within-system cooperation
  - Fairness and among-system cooperation
- **Part II: More powerful detectors**
  - **Coherent detectors**
  - Feature detectors
- Part III: Hardware considerations
  - Fundamental hardware limitations
  - Ways around them

The need for more complex detectors

- The energy detector works for all possible primary signals without knowing what they are.
- If the primary is zero-mean, white, and occupies all degrees of freedom, nothing more is possible.
- But physical signals are not that general.

The need for more complex detectors

- The energy detector works for all possible primary signals without knowing what they are.
- If the primary is zero-mean, white, and occupies all degrees of freedom, nothing more is possible.
- But physical signals are not that general.
  - Have deterministic pilot tones.

The need for more complex detectors

- The energy detector works for all possible primary signals without knowing what they are.
- If the primary is zero-mean, white, and occupies all degrees of freedom, nothing more is possible.
- But physical signals are not that general.
  - Have deterministic pilot tones.
  - Have guard bands and do not occupy all degrees of freedom.

The need for more complex detectors

- The energy detector works for all possible primary signals without knowing what they are.
- If the primary is zero-mean, white, and occupies all degrees of freedom, nothing more is possible.
- But physical signals are not that general.
  - Have deterministic pilot tones.
  - Have guard bands and do not occupy all degrees of freedom.
- Can we exploit these to improve performance and robustness?

Coherent detection

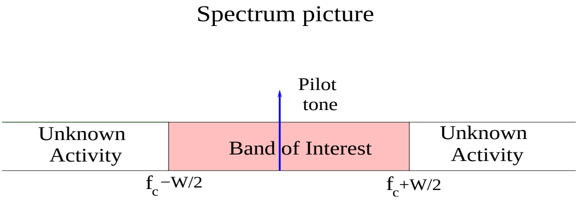

- Look for the primary pilot in the 'band of interest'
- Pilots come up in different situations:
  - Denial pilot (very critical)
  - Permissive pilot (non-critical)
- Within band model:
  - White or unknown signal, $X(t)$
  - Independent white noise, $W(t)$

Hypothesis testing model: Coherent detection

- Distinguish between the following hypotheses:

$$\begin{aligned}\mathcal{H}_0 : Y[n] &= W[n] \\ \mathcal{H}_1 : Y[n] &= W[n] + \sqrt{(1-\theta)}X[n] + \sqrt{\theta}X_p[n]\end{aligned}$$

- Basic assumptions:
  - Signal samples $X[n]$'s are white or orthogonal to the pilot.
  - Noise samples $W[n]$'s are white.
  - $X_p[n]$ is a known pilot tone
  - θ is the fraction of total power allocated to pilot tone
- Possible detection strategies:
  - Energy detector (radiometer)
  - Coherently detect pilot tone

Matched filter analysis

- Correlate received signal with a unit vector in pilot's direction

$$T(\mathbf{y}) = \frac{1}{N} \sum_{n=1}^{N} Y[n]\hat{X}_p[n]$$

- $\hat{\mathbf{x}}_p$ is a unit vector in the direction of the pilot
- Test statistic under both hypotheses:

$$\mathcal{H}_0 : T(\mathbf{y}) = \frac{1}{N}\sum_{n=1}^{N} W[n]\hat{X}_p[n] \sim \mathcal{N}(0, \frac{1}{N}\sigma^2)$$

$$\mathcal{H}_1 : T(\mathbf{y}) = \frac{1}{N}\sum_{n=1}^{N}\{\sqrt{\theta}X_p[n] + \sqrt{(1-\theta)}X[n] + W[n]\}\hat{X}_p[n] \sim \mathcal{N}(\sqrt{\theta P}, \frac{1}{N}\sigma^2)$$

- Here $P = \frac{1}{N}\sum_{n=1}^{N} X^2[n]$, is the average signal power, σ^2 is the noise power.
- Decision rule: $T(\mathbf{y}) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma(\sigma^2)$, Threshold γ is set based on P_{FA}

Error probability for matched filter

False alarm

$$\begin{aligned}
P_{FA} &= P(T(\mathbf{y}) > \gamma | \mathcal{H}_0) \\
&= P\left(\frac{T(\mathbf{y})}{\sqrt{\frac{\sigma^2}{N}}} > \frac{\gamma}{\sqrt{\frac{\sigma^2}{N}}} | \mathcal{H}_0\right) \\
&= \mathcal{Q}\left(\frac{\gamma}{\sqrt{\frac{\sigma^2}{N}}}\right)
\end{aligned}$$

Missed detection

$$\begin{aligned}
P_D &= P(T(\mathbf{y}) > \gamma | \mathcal{H}_1) \\
&= P\left(\frac{T(\mathbf{y}) - \sqrt{\theta P}}{\sqrt{\frac{\sigma^2}{N}}} > \frac{\gamma - \sqrt{\theta P}}{\sqrt{\frac{P\sigma^2}{N}}} | \mathcal{H}_1\right) \\
&= \mathcal{Q}\left(\frac{\gamma - \sqrt{\theta P}}{\sqrt{\frac{\sigma^2}{N}}}\right)
\end{aligned}$$

- Eliminating γ,

$$N = [\mathcal{Q}^{-1}(P_D) - \mathcal{Q}^{-1}(P_{FA})]^2 \theta^{-1} SNR^{-1}$$

Matched filter performance

- $X[n]$ are BPSK modulated
  - $X[n] \sim$ Bernoulli $(\frac{1}{2})$, taking values in $\{\sqrt{P}, -\sqrt{P}\}$
- $\theta = 0.01$, *ie.* 1% energy in the pilot

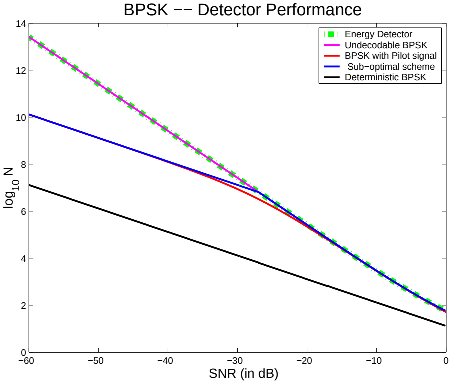

Easy uncertainties

- Unknown but steady frequency and time offsets
- Effect on pilot:
  - $X_p[n] = A\cos(2\pi f_s n + \phi)$
  - f_s and ϕ unknown
  - Phase offset is easy to deal with
- Approach: Search in many bins
- Implementation and computational complexity
  - Need to compute an N-point FFT
  - Search for the maximum over f_s
  - Computationally more involved than the energy detector

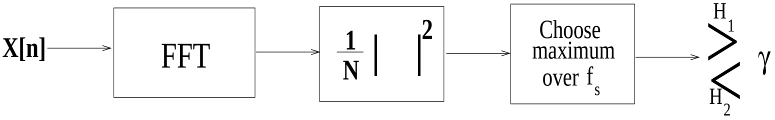

Analysis of easy uncertainties

False alarm probability

$$
\begin{aligned}
P_{FA} &= 1 - \left(1 - \exp\left(\frac{-\gamma}{\sigma^2}\right)\right)^L \\
&\approx 1 - \left(1 - L\exp\left(\frac{-\gamma}{\sigma^2}\right)\right) \\
&= LP_{FA}(\text{bin})
\end{aligned}
$$

where $L = \frac{N}{2} - 1$ is the number of frequency bins

Probability of detection

$$P_D = \mathcal{Q}_{\chi_2'^2\left(\frac{NSNR}{2}\right)}\left(\frac{2\gamma}{\sigma^2}\right)$$

Eliminating, γ, we get

$$P_D = \mathcal{Q}_{\chi_2'^2\left(\frac{NSNR}{2}\right)}\left(2\ln\frac{L}{P_{FA}}\right)$$

- Need to tighten P_{FA}, linearly with the number of bins we search
- Effective frequency uncertainty scales with dwell time N

Easy uncertainty: Impact on dwell time

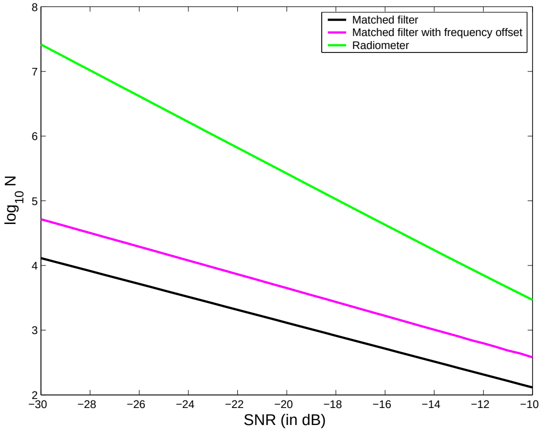

- New $O(\log N)$ term in the threshold.
- Impact on dwell time is not too bad

Medium uncertainty: White noise level uncertainty

- Detection problem:

$$\begin{aligned}\mathcal{H}_0 : Y[n] &= W[n] \\ \mathcal{H}_1 : Y[n] &= W[n] + \sqrt{\theta}X_p[n] + \sqrt{(1-\theta)}X[n]\end{aligned}$$

- $W[n] \in \mathcal{W}_x = [\sigma^2_{low}, \sigma^2_{high}]$
- Assume x dB uncertainty in noise level

Medium uncertainty: White noise level uncertainty

- Detection problem:

$$\begin{aligned}\mathcal{H}_0 : Y[n] &= W[n] \\ \mathcal{H}_1 : Y[n] &= W[n] + \sqrt{\theta} X_p[n] + \sqrt{(1-\theta)} X[n]\end{aligned}$$

- $W[n] \in \mathcal{W}_x = [\sigma^2_{low}, \sigma^2_{high}]$
- Assume x dB uncertainty in noise level
- Main idea: Coherent processing gain can overcome noise level uncertainty

$$T(\mathbf{y}) = \frac{1}{N} \sum_{n=1}^{N} Y[n] \hat{X}_p[n] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma(\sigma^2)$$

Noise uncertainty: closer look

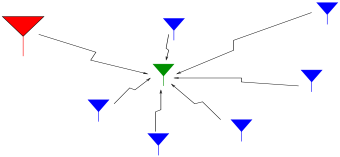

- Noise = receiver/background noise + secondary interference, $\sigma^2 = \sigma_0^2 + \sigma_i^2$

Effect of interference

- Interference is uncertain
- Dominating term in noise uncertainty

Noise uncertainty: closer look

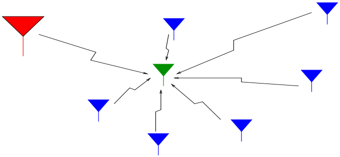

- Noise = receiver/background noise + secondary interference, $\sigma^2 = \sigma_0^2 + \sigma_i^2$

Effect of interference

- Interference is uncertain
- Dominating term in noise uncertainty

Proposed remedy

- Robustly estimate σ_i^2
- Significant reduction in noise uncertainty

Setting threshold: learning based approach

In-band measurement

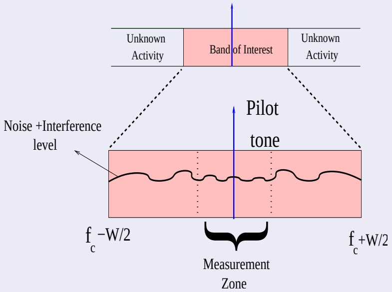

Noise prediction error

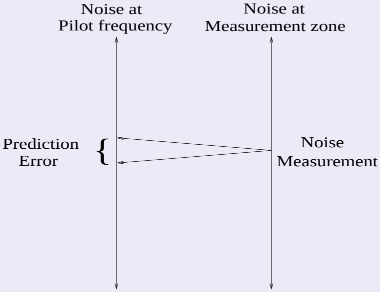

- Sources of prediction error:
  - Long-term frequency selectivity in secondary signals
  - Coherence bandwidth for secondary signals
  - Practical guess $(1 - 5\%)$

Fundamentally challenging uncertainties

- Realistic model:

$$\begin{aligned} \mathcal{H}_0 : Y[n] &= W[n] \\ \mathcal{H}_1 : Y[n] &= W[n] + \sum_{l=0}^{L-1} h_l[n]\tilde{X}[n-l] \end{aligned}$$

  where $\tilde{X}[n] = \sqrt{\theta}X_p[n] + \sqrt{(1-\theta)}X[n]$

- Assumptions:
  - Fast multipath fading: $h_l \sim CN(0,1)$
  - $W[n] \in \mathcal{W}_x$: Noise uncertainty set

Fast fading

- Assume *known channel coherence time*, T_c
  - Fading is assumed constant during this coherence time
  - Matched filter (coherent processing) can be applied in each coherent block
  - Test statistic:

$$T(\mathbf{y}) = \frac{1}{N} \sum_{n=0}^{N-1} \left[\frac{1}{\sqrt{N_c}} \sum_{k=1}^{N_c} Y[n] \hat{X}_p[nN_c + k] \right]^2$$

  - N_c: Length of single coherence block
- Reduced to the energy detector case,
  - Processing gain due to the coherence time T_c
  - Less interference uncertainty.

Summary: robust gains from coherent processing

- Primary coherence time: [100 μs, 10 ms]
  $\Rightarrow [20, 40]$dB: Processing gain for both wall and dwell time.
- Interference prediction error: $[1, 10]\%$
  $\Rightarrow [10, 20]$ dB: In-band measurement gain for wall only
- Pilot energy: $\theta = [0.01, 0.1]$
  $\Rightarrow [-20, -10]$ dB loss for both wall and dwell time
- Effective *SNR* wall: $[10, 50]$ dB lower
  Dwell times: $[0, 30]$ dB better
- System could be ‘Wall limited’ or ‘Dwell limited’

Coherence buys freedom from secondary-MAC

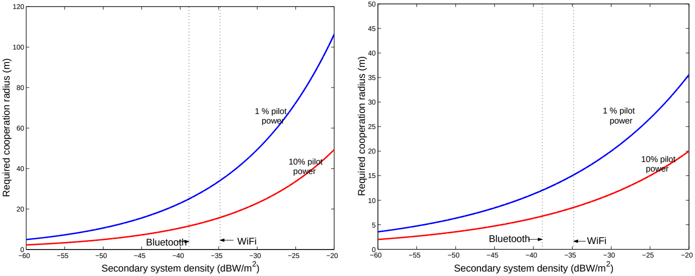

(a) $\alpha_{22} = -5$ (b) $\alpha_{22} = -6$

- Primary transmitter power: 100 kW, Protection margin $\mu = 1$ dB
- Primary Attenuation: $\alpha_{12} = -3.5$, Post-cooperation fading margin = 10 dB
- Coherent processing gain 10^4
- Interference prediction error = 1%, Residual device uncertainty = 1 dB

Extensions and Implications

- Universality over coherence times is possible
  - Search strategy over coherence times and offsets.
  - Increases computational complexity.
  - Better dwell times when primary channel is more coherent.
- Implications
  - Secondary waveforms must avoid using confusing pilots.
  - Coherent processing can drastically reduce the need to cooperate with other systems unless seriously dwell-time limited.
  - Not enough of a gain to eliminate the need for within-system cooperation.