SNR Walls for Cognitive Radio Spectrum Sensing

Anant Sahai
presenting joint work with students:
Danijela Cabric   Rahul Tandra

BWRC and Wireless Foundations Center
U.C. Berkeley

Major support from the National Science Foundation

I^2R Seminar: 8th January, 2008

Big question: wireless system architectures

How are we going to get the next factor of 10-100 improvement in wireless communication?

Big question: wireless system architectures

How are we going to get the next factor of 10-100 improvement in wireless communication?

- Error correcting codes: already approaching Shannon limits
- Multiple antennas: limits due to physical scattering environments

Big question: wireless system architectures

How are we going to get the next factor of 10-100 improvement in wireless communication?

- Error correcting codes: already approaching Shannon limits
- Multiple antennas: limits due to physical scattering environments
- Increase spectrum utilization

Big question: wireless system architectures

How are we going to get the next factor of 10-100 improvement in wireless communication?

- Error correcting codes: already approaching Shannon limits
- Multiple antennas: limits due to physical scattering environments
- Increase spectrum utilization
  - Colonize new spectrum: very promising work on 60GHz radios

Big question: wireless system architectures

How are we going to get the next factor of 10-100 improvement in wireless communication?

- Error correcting codes: already approaching Shannon limits
- Multiple antennas: limits due to physical scattering environments
- Increase spectrum utilization
  - Colonize new spectrum: very promising work on 60GHz radios
  - **Make better use of the spectrum we have by exploiting frequency agile devices.**

The goal . . .

- Use spectrum holes

The goal . . .

- Use spectrum holes
- Avoid causing harmful interference to the primary

The goal . . .

- Use spectrum holes
- Avoid causing harmful interference to the primary
- Preserve “flexibility” for the primary user

The goal . . .

- Use spectrum holes
- Avoid causing harmful interference to the primary
- Preserve “flexibility” for the primary user
  - Flexibility of deployment: where the primary user turns on its system.

The goal . . .

- Use spectrum holes
- Avoid causing harmful interference to the primary
- Preserve "flexibility" for the primary user
  - Flexibility of deployment: where the primary user turns on its system.
  - Flexibility of use: choice of modulation, error-correcting codes, content.

The goal . . .

- Use spectrum holes
- Avoid causing harmful interference to the primary
- Preserve “flexibility” for the primary user
  - Flexibility of deployment: where the primary user turns on its system.
  - Flexibility of use: choice of modulation, error-correcting codes, content.
  - Flexibility of architecture: point-to-point, broadcast, relays, etc?

The goal . . .

- Use spectrum holes
- Avoid causing harmful interference to the primary
- Preserve “flexibility” for the primary user
  - Flexibility of deployment: where the primary user turns on its system.
  - Flexibility of use: choice of modulation, error-correcting codes, content.
  - Flexibility of architecture: point-to-point, broadcast, relays, etc?
- **Can’t have everything!**

Outline

- Introduction
- **Why is robust sensing important?**
- Noncoherent
- Coherent
- Cyclostationary
- Capacity/Robustness tradeoff

Interference due to the primary

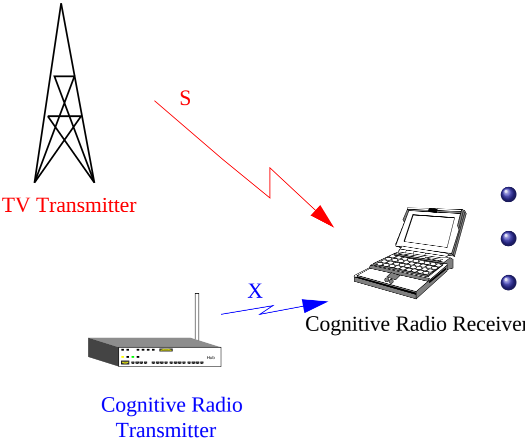

$$Y = X + S + Z$$

- Signal X has power P
- Interference S has power Q
- Noise Z has power N

Assume known at transmitter

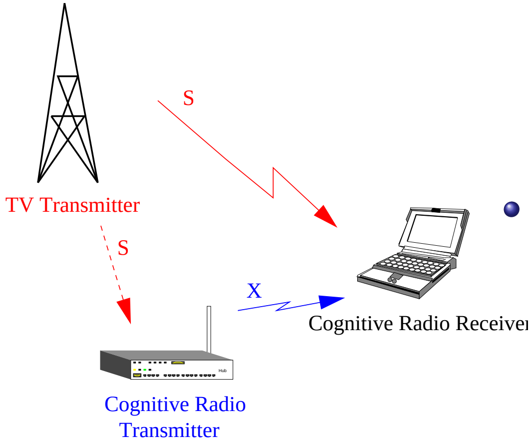

$$Y = X + S + Z$$

- We get advance warning of the interference.

“Canceling” known interference

Dirty-paper coding: Costa’83

Capacity of (complex) AWGN channel with interference known *a priori* is same as that with no interference

$$C = \log\left(1 + \frac{P}{N}\right)$$

- Explored by *Tarokh et al* and *Viswanath et al* :
  If primary interference is known *a priori*, achievable secondary rates are higher

Is the interference *really* known?

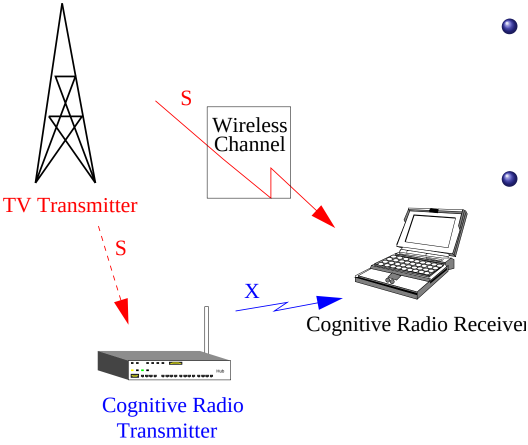

- Primary's transmissions are different from the interference caused by primary's transmissions
- Wireless fading channel connects the primary transmitter to the secondary receiver.

Is the interference *really* known?

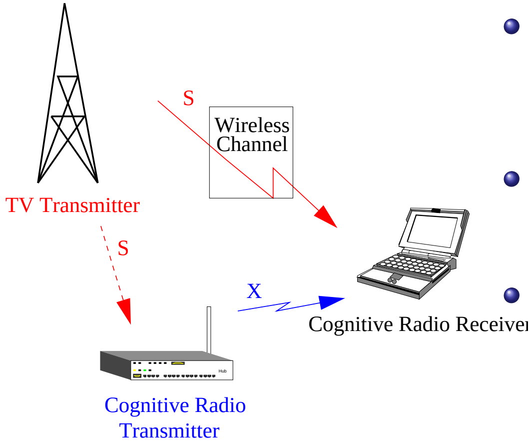

- Primary's transmissions are different from the interference caused by primary's transmissions
- Wireless fading channel connects the primary transmitter to the secondary receiver.
- Would you bet your rate on knowing the phase of the channel?

Simplify: just phase uncertainty

$$Y = X + Se^{j\theta} + Z$$

- θ is same for the whole block
- θ is unknown at the transmitter

Maximum Achievable Rate

What is the maximum rate for reliable communication for unknown θ but known S at the transmitter?

Our upper bound for phase uncertainty

$$R \leq \frac{1}{2} \log \left[\frac{(P+Q+N)^2}{4QN} \right]$$

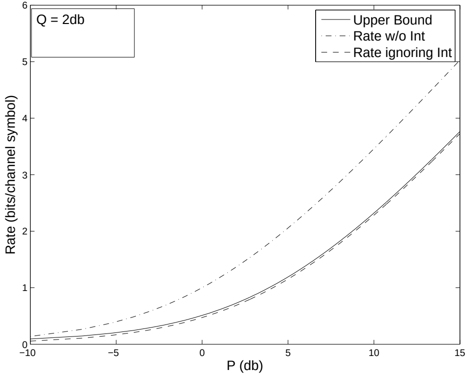

Incidental architectural implications

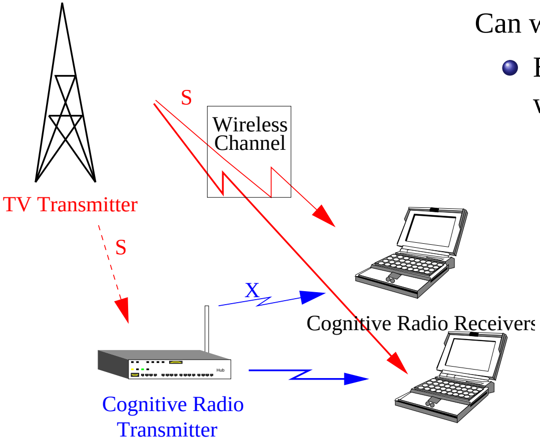

Can we do "network-coding?"

- Based on broadcast nature of wireless.

Incidental architectural implications

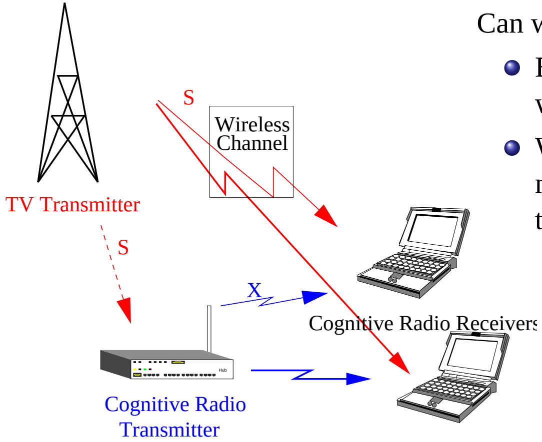

Can we do "network-coding?"

- Based on broadcast nature of wireless.
- With feedback, the two phases may be estimated at the transmitter.

Incidental architectural implications

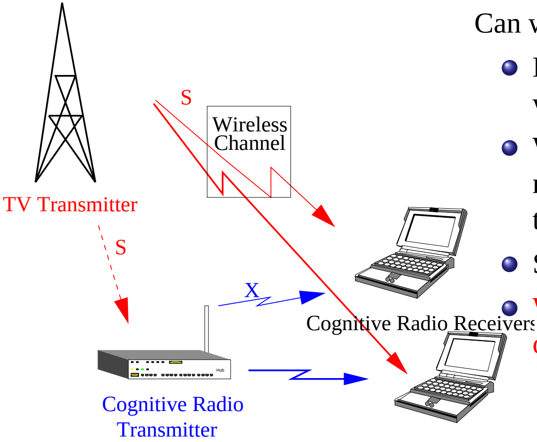

Can we do “network-coding?”

- Based on broadcast nature of wireless.
- With feedback, the two phases may be estimated at the transmitter.
- Still two phases = rate-loss!
- Writing on interference is not omnidirectional!

Outline

- Introduction
- Why is robust sensing important?
- **Noncoherent**
- Coherent
- Cyclostationary
- Capacity/Robustness tradeoff

Sensing the primary's presence

Spectrum picture

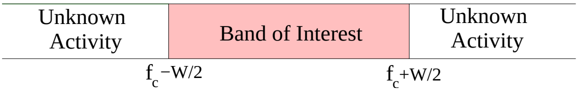

- Look for the primary in the 'band of interest'
- Within band model:
  - Unknown signal: $X(t)$
  - Independent white noise: $W(t)$
- Idealization: sample the band of interest at Nyquist rate

Simplest detector: energy detector

- Received energy used for detection
- Test statistic and decision rule:

$$\sum_{n=1}^{N} Y^2[n] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma(N)$$

- Sample complexity:

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

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

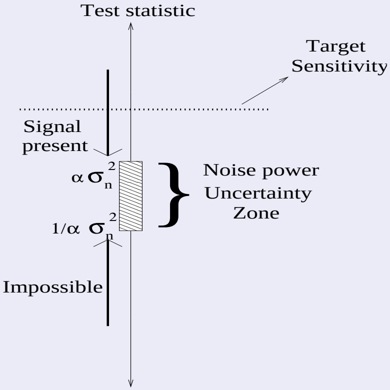

SNR wall for energy detector

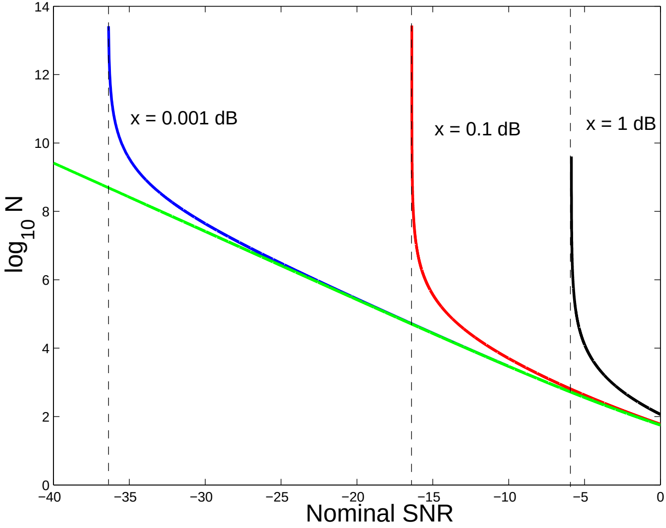

SNR wall for energy detector

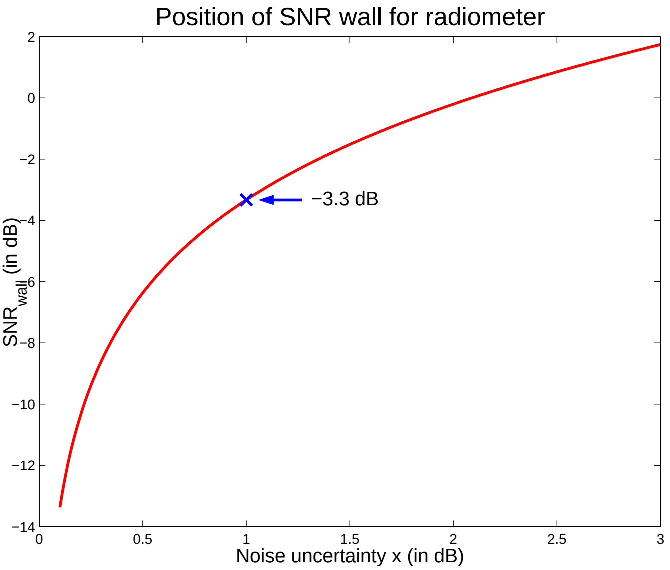

What if you knew the signal constellation?

What if you knew the signal constellation?

- Assumption: Symbol constellation is zero-mean

What if you knew the signal constellation?

- Assumption: Symbol constellation is zero-mean
- Result: Sample complexity is like the energy detector at low SNR
- $N = O(SNR^{-2})$

What if you knew the signal constellation?

- Assumption: Symbol constellation is zero-mean
- Result: Sample complexity is like the energy detector at low SNR
- $N = O(SNR^{-2})$
- Compute log likelihood ratio, $\text{LLR}(\mathbf{y}, SNR)$

What if you knew the signal constellation?

- Assumption: Symbol constellation is zero-mean
- Result: Sample complexity is like the energy detector at low SNR
- $N = O(SNR^{-2})$
- Compute log likelihood ratio, $\mathrm{LLR}(\mathbf{y}, SNR)$
- Since $SNR \ll 1$, use Taylor series approximation

What if you knew the signal constellation?

- Assumption: Symbol constellation is zero-mean
- Result: Sample complexity is like the energy detector at low SNR
- $N = O(SNR^{-2})$
- Compute log likelihood ratio, $\text{LLR}(\mathbf{y}, SNR)$
- Since $SNR \ll 1$, use Taylor series approximation
- First order terms of the observation $\mathbf{y}$ vanish for a zero-mean constellation
- Only quadratic terms of $\mathbf{y}$ remain

“Noise” is uncertain in distribution

- Noise is usually assumed to be Gaussian
- Sources of uncertainty:

“Noise” is uncertain in distribution

- Noise is usually assumed to be Gaussian
- Sources of uncertainty:
  - Variation in components

“Noise” is uncertain in distribution

- Noise is usually assumed to be Gaussian
- Sources of uncertainty:
  - Variation in components
  - Interference from other transmissions

“Noise” is uncertain in distribution

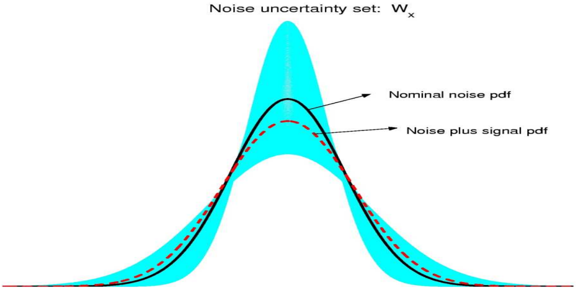

- Noise is usually assumed to be Gaussian
- Sources of uncertainty:
  - Variation in components
  - Interference from other transmissions
  - Nonlinear mixing from other bands

Realistic model for noise uncertainty

- 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.*

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

Energy detector summary

- Noise uncertainty is always present
  - At least about 1-2 dB of device level uncertainty
    - Cannot see 3 dB below the noise
  - Easily 10-30 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 -116dBm sensitivity to deal with rare fading.
  - Thermal Noise alone at -106dBm !
  - With interference could be -91dBm !
- What can we do to mitigate this?

Outline

- Introduction
- Why is robust sensing important?
- Noncoherent
- **Coherent**
- Cyclostationary
- Capacity/Robustness tradeoff

Coherent detection: maximally detailed modeling

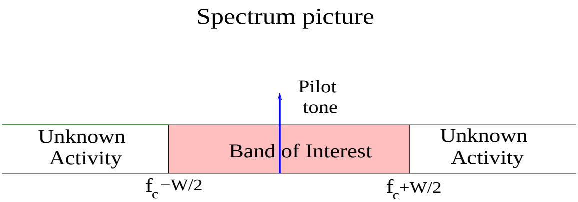

- Look for the primary pilot in the ‘band of interest’

Coherent detection: maximally detailed modeling

Spectrum picture

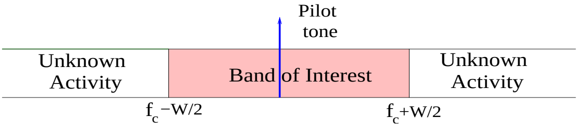

- Look for the primary pilot in the ‘band of interest’
- Pilots come up in different situations:
  - Denial pilot (very critical)
  - Permissive pilot (non-critical)

Coherent detection: maximally detailed modeling

Spectrum picture

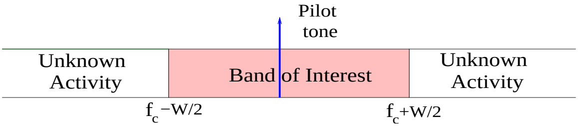

- Look for the primary pilot in the 'band of interest'
- Pilots come up in different situations:
  - Denial pilot (very critical)
  - Permissive pilot (non-critical)
- Within band model:
  - Unknown component, $X(t)$
  - Pilot signal, $X_p(t)$
  - Independent white noise, $W(t)$

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

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
- Main idea: Coherent processing gain can overcome noise level uncertainty

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

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

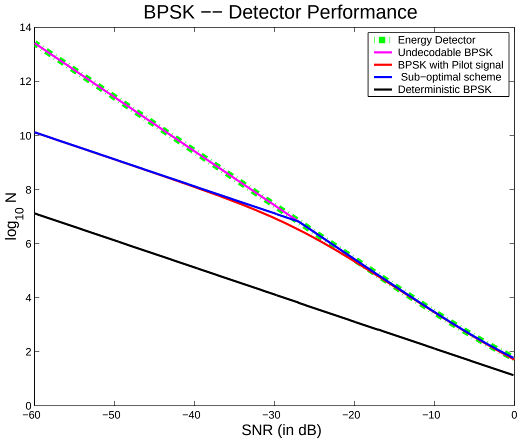

Easy uncertainties

- 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

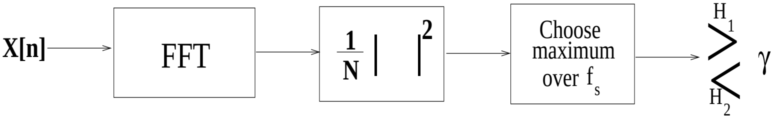

Noise uncertainty: closer look

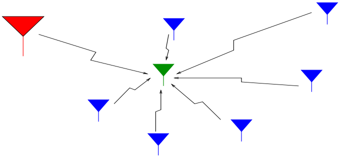

- Noise = receiver/background noise + secondary interference, $\sigma^2 = \sigma_0^2 + \sigma_i^2$

Effect of unknown interference

- Interference is uncertain
- Dominating term in noise uncertainty

Noise uncertainty: closer look

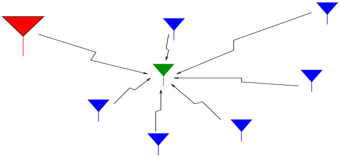

- Noise = receiver/background noise + secondary interference, $\sigma^2 = \sigma_0^2 + \sigma_i^2$

Effect of unknown interference

- Interference is uncertain
- Dominating term in noise uncertainty

Proposed remedy

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

Setting threshold by learning

In-band measurement

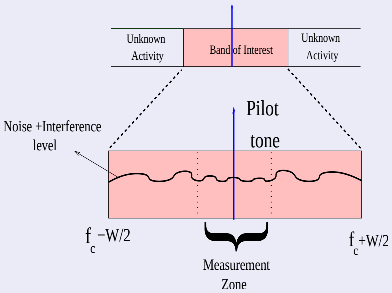

Noise calibration error

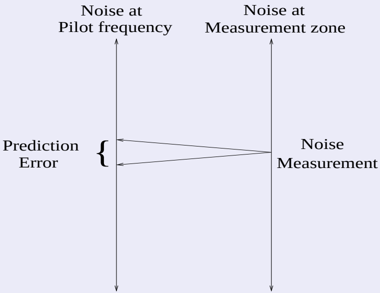

- Sources of calibration error:
  - Long-term frequency selectivity in secondary signals
  - “Coherence bandwidth/time” for secondary signals

Fundamentally challenging uncertainty: fast fading

- Assume *worst-case channel coherence time*: T_c
  - Fading is assumed constant during this coherence time
  - Matched filter 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: finite gains from coherent processing

- Primary coherence time: [100 μs, 10 ms]
  $\Rightarrow [20, 40]$dB: Processing gain for both wall and dwell time.

Summary: finite gains from coherent processing

- Primary coherence time: $[100\ \mu s, 10\ ms]$
  $\Rightarrow [20, 40]$dB: Processing gain for both wall and dwell time.
- Interference calibration error: $[1, 10]\%$
  $\Rightarrow [10, 20]$ dB: In-band measurement gain for wall only

Summary: finite gains from coherent processing

- Primary coherence time: [100 μs, 10 ms]
  $\Rightarrow [20, 40]$dB: Processing gain for both wall and dwell time.
- Interference calibration 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

Summary: finite gains from coherent processing

- Primary coherence time: $[100\ \mu s, 10\ \text{ms}]$
  $\Rightarrow [20, 40]$dB: Processing gain for both wall and dwell time.
- Interference calibration 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'

Outline

- Introduction
- Why is robust sensing important?
- Noncoherent
- Coherent
- **Cyclostationary**
- Capacity/Robustness tradeoff

Refined model: Cyclostationary signals
William Gardner

- Pulse amplitude modulated signal:
  - BPSK modulated over a rectangular pulse shape

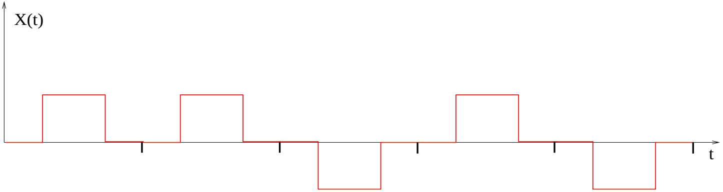

- Nyquist sampling — Stationary signal
- Oversampling — Cyclostationary signal
  - Periodicity corresponding to data rate
  - Feature detectors
    - ★ Sample complexity same as energy detector – $O(SNR^{-2})$

Spectral Correlation Function

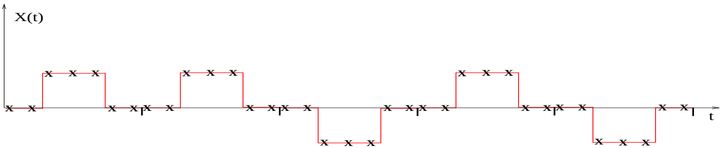

- Spectral correlation function can be estimated by

$$\tilde{S}_X^{\alpha} = \lim_{t\to\infty} \lim_{T\to\infty} \frac{1}{t}\frac{1}{T} \int_{-\frac{t}{2}}^{\frac{t}{2}} X_T\left(t, f+\frac{\alpha}{2}\right) X_T^*\left(t, f-\frac{\alpha}{2}\right) dt$$

where

$$X_T(t,f) = \int_{t-\frac{T}{2}}^{t+\frac{T}{2}} x(t) e^{-j2\pi f u} du$$

is the spectral component of $x(t)$ at frequency f with bandwidth $\frac{1}{T}$

- S_x^0 is the power spectral density of $x(t)$
- $S_x^{\alpha} = 0$ for a WSS process for $\alpha \neq 0$

Single-cycle feature detector

- Estimate Spectral Correlation Function: $\tilde{S}_Y^\alpha(f)$
- Compute the following test-statistic

$$T(Y) = \int \tilde{S}_Y^\alpha(f) S_X^\alpha(f)^* df$$

Under Hypothesis $\mathcal{H}_0$

$$\begin{aligned} T(Y|\mathcal{H}_0) &= \int \tilde{S}_W^\alpha(f) S_X^\alpha(f)^* df \\ &\rightarrow \int S_W^\alpha(f) S_X^\alpha(f)^* df \\ &= \int 0 \cdot S_X^\alpha(f)^* df = 0 \end{aligned}$$

Under Hypothesis $\mathcal{H}_1$

$$\begin{aligned} T(Y|\mathcal{H}_1) &= \int \left[\tilde{S}_X^\alpha(f) + \tilde{S}_W^\alpha(f)\right] S_X^\alpha(f)^* df \\ &\rightarrow \int \left[S_X^\alpha(f) + S_W^\alpha(f)\right] S_X^\alpha(f)^* df \\ &= \int |S_X^\alpha(f)|^2 df \end{aligned}$$

- Feature detection works at every *SNR*, given sufficiently large *N*!

BPSK example

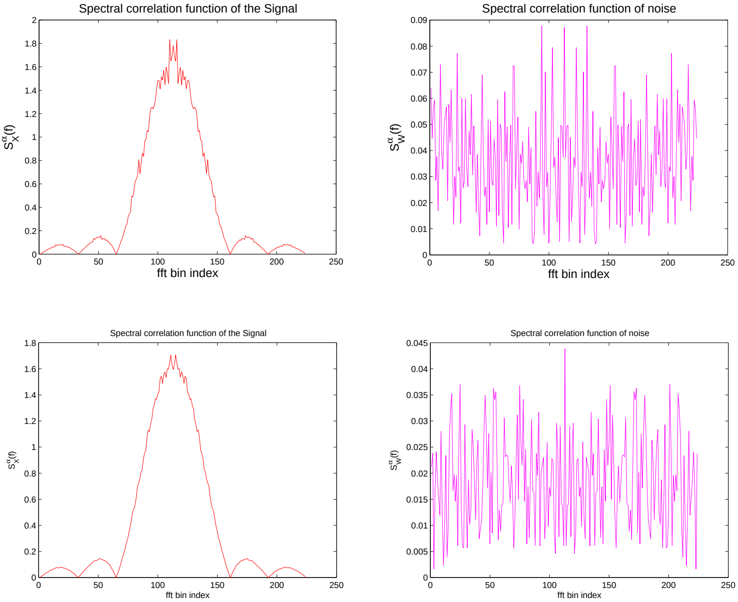

Robust to flat-fading

- Flat fading does not change the pulse-shape, just the amplitude.

Robust to flat-fading

- Flat fading does not change the pulse-shape, just the amplitude.
- Amplitude is changing due to data modulation anyway.

Robust to flat-fading

- Flat fading does not change the pulse-shape, just the amplitude.
- Amplitude is changing due to data modulation anyway.
- **More robust** than even coherent detection in this regard!

Impact of frequency-selective fading

- Need to incorporate frequency-selective fading into channel model

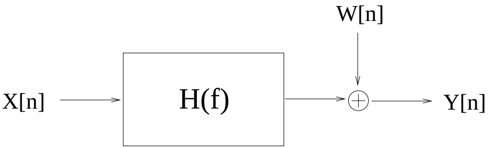

- Unknown channel filter $H(f)$
- Spectral Correlation Function

$$S_Y^{\alpha}(f) = H\left(f - \frac{\alpha}{2}\right) H^*\left(f + \frac{\alpha}{2}\right) S_X^{\alpha}(f) + S_W^{\alpha}(f)$$

- Fading smears the features
- Can kill the feature completely in the worst case!!!

Fundamental limitation: "delay" coherence time

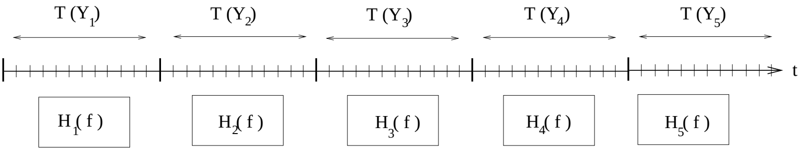

- In reality, fading response varies with time.
- Scale of variation – channel 'delay-coherence' time
- Single cycle feature detector:

$$\begin{aligned} T(Y) &= \int \tilde{S}_Y^{\alpha}(f) S_X^{\alpha}(f)^* df \\ &= \frac{1}{M}\sum_{i=1}^{M} T(Y_i) \\ &\rightarrow 0 \end{aligned}$$

- Averaging over delay coherence times kills features!

Fundamental limitation: "delay" coherence time

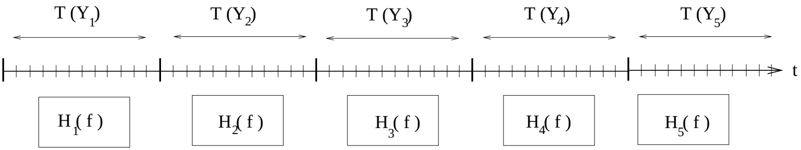

- Modified test-statistic:

$$T(Y) = \frac{1}{M}\sum_{i=1}^{M} |T(Y_i)|^2$$

- Key observation:
  - $\mathbb{E}\left[T(Y)|\mathcal{H}_0\right] \neq 0$
- Consequence:
  - $\mathbb{E}\left[T(Y)|\mathcal{H}_1\right] - \mathbb{E}\left[T(Y)|\mathcal{H}_0\right] \leq \Delta_{\mathbb{E}[T(Y)|\mathcal{H}_0]}$
  - $\Rightarrow$ SNR wall

What about noise calibration?

- Can learn the overall noise/interference level.
- Like the coherent case, the true wall comes from unknown color.

Detector robustness with coherence time

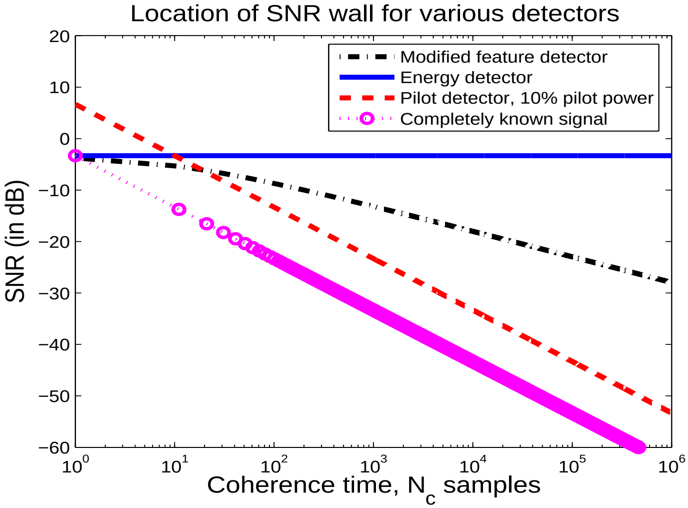

Outline

- Introduction
- Why is robust sensing important?
- Noncoherent
- Coherent
- Cyclostationary
- **Capacity/Robustness tradeoff**

Are we just being paranoid?

- Experimental validation
- Used BEE2 and 2.4 GHz radio front-ends

Are we just being paranoid?

- Experimental validation
- Used BEE2 and 2.4 GHz radio front-ends

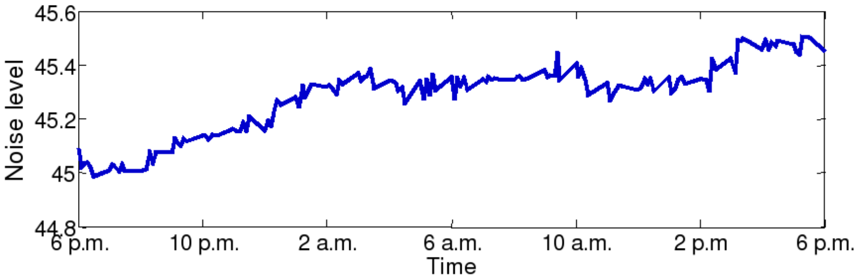

Noise levels move around over a day.

Are we just being paranoid?

- Experimental validation
- Used BEE2 and 2.4 GHz radio front-ends

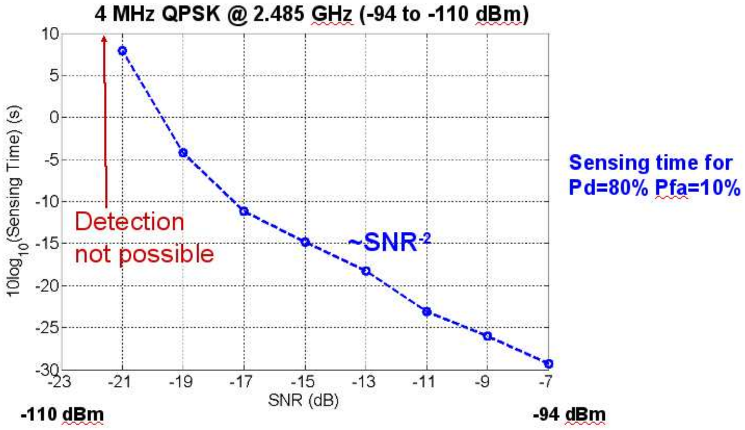

Are we just being paranoid?

- Experimental validation
- Used BEE2 and 2.4 GHz radio front-ends

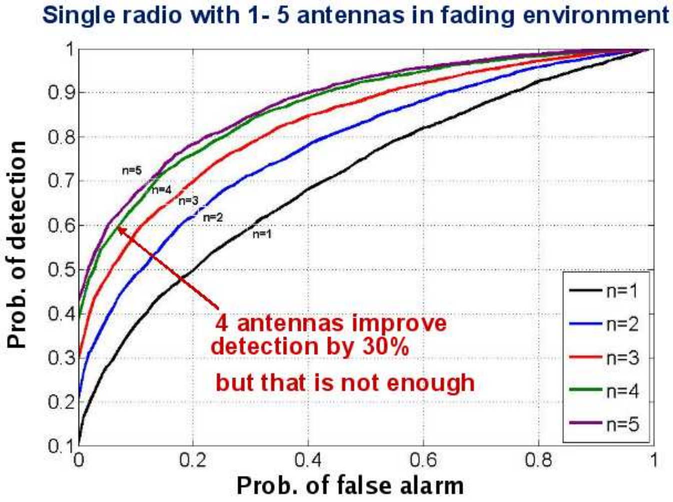

Are we just being paranoid?

- Experimental validation
- Used BEE2 and 2.4 GHz radio front-ends

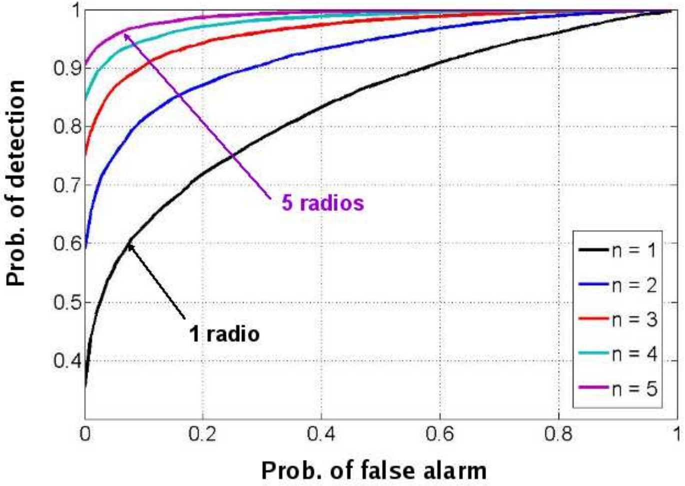

Are we just being paranoid?

- Experimental validation
- Used BEE2 and 2.4 GHz radio front-ends

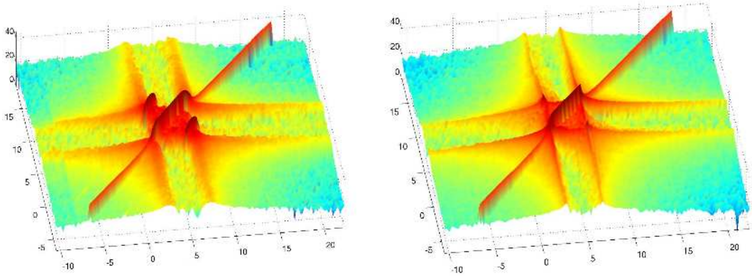

The spectral correlation function shows spectral redundancy in a transformed domain.

Are we just being paranoid?

- Experimental validation
- Used BEE2 and 2.4 GHz radio front-ends

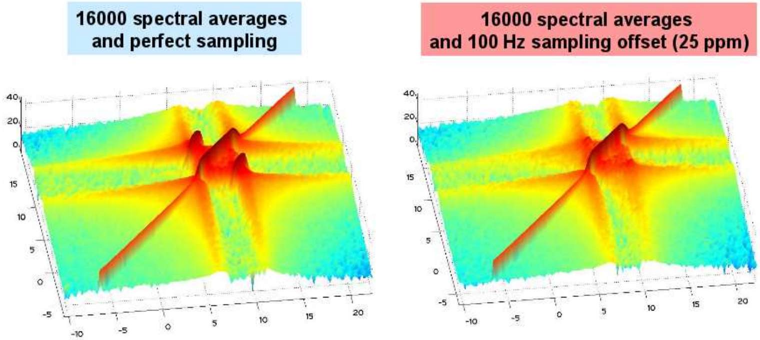

But this redundancy is blurred away by fast fading.

The core tension

- Flexibility of architecture: primary **must lose** to allow setting sensitivity for cognitive users.

The core tension

- Flexibility of architecture: primary **must lose** to allow setting sensitivity for cognitive users.
- Flexibility of deployment: the goal of sensing

The core tension

- Flexibility of architecture: primary **must lose** to allow setting sensitivity for cognitive users.
- Flexibility of deployment: the goal of sensing
- Flexibility of use
  - How much can we assume about the primary modulation?
  - Do these “features” come at a price to the primary?

Tradeoffs illustrated

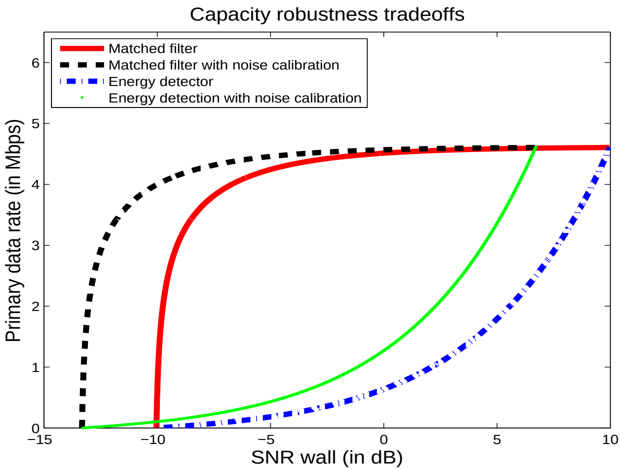

Tradeoffs illustrated