
Distributed Massive MIMO

Algorithms, Architectures, Concept Systems



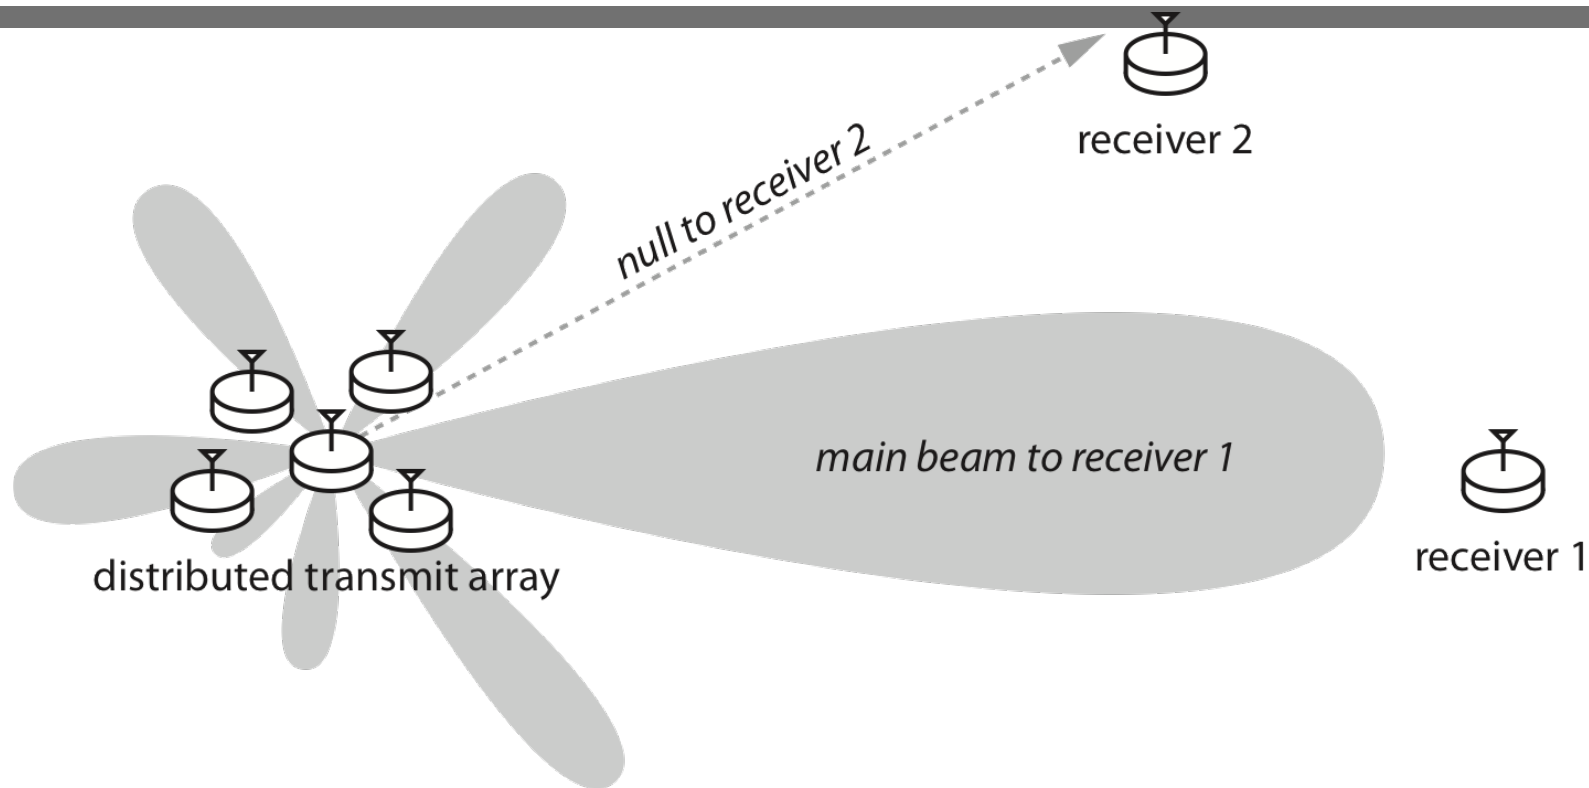
Upamanyu Madhow

Dept. of Electrical and Computer Engineering

University of California, Santa Barbara

Joint work with Profs. Rick Brown (WPI), Soura Dasgupta, Raghu Mudumbai (U. Iowa)

The promise of distributed MIMO



Vision: Opportunistic MIMO without form factor constraints

Synchronization-enabled protocols to support distributed realization of any MIMO scheme: beamforming, nulling, SDMA, spatial muxing, interference alignment,...


CONCEPT SYSTEMS: DISTRIBUTED BASE STATION, DISTRIBUTED 911, SENSOR NETWORK REACHBACK,...

Our focus




- **All-wireless, scalable**
 - Add nodes opportunistically without algorithmic or protocol disruption
 - Moderate backhaul requirements
 - Main focus is on range enhancement
- **Contrast to CoMP (Fraunhofer, TUD) and WiFi AP cooperation (USC, MIT)**
 - Uses high-speed backhaul
 - Main focus is on interference management/MU-MIMO

Physical hurdles and technical approaches



1) Clocks drift!
(and mobility does not help)
→ Must model and track
(multiple levels of sync: freq, phase, timing)

2) Geometric unknowns
(no array manifold, messy channels)
→ Feedback



Today's menu



- **The good news**
 - Recent successes in D-TX and D-RX
 - Scalable architectures, lab demos with software-defined radios
 - Long-range demos by industry
- **Under the hood: fundamental problems in estimation & tracking**
 - Attaining fundamental limits of one-shot estimation
 - Tracking with nonlinear, intermittent observations
- **Open issues**
 - Rapid mobility
 - Dispersive channels
 - Synchronization-enabled protocols

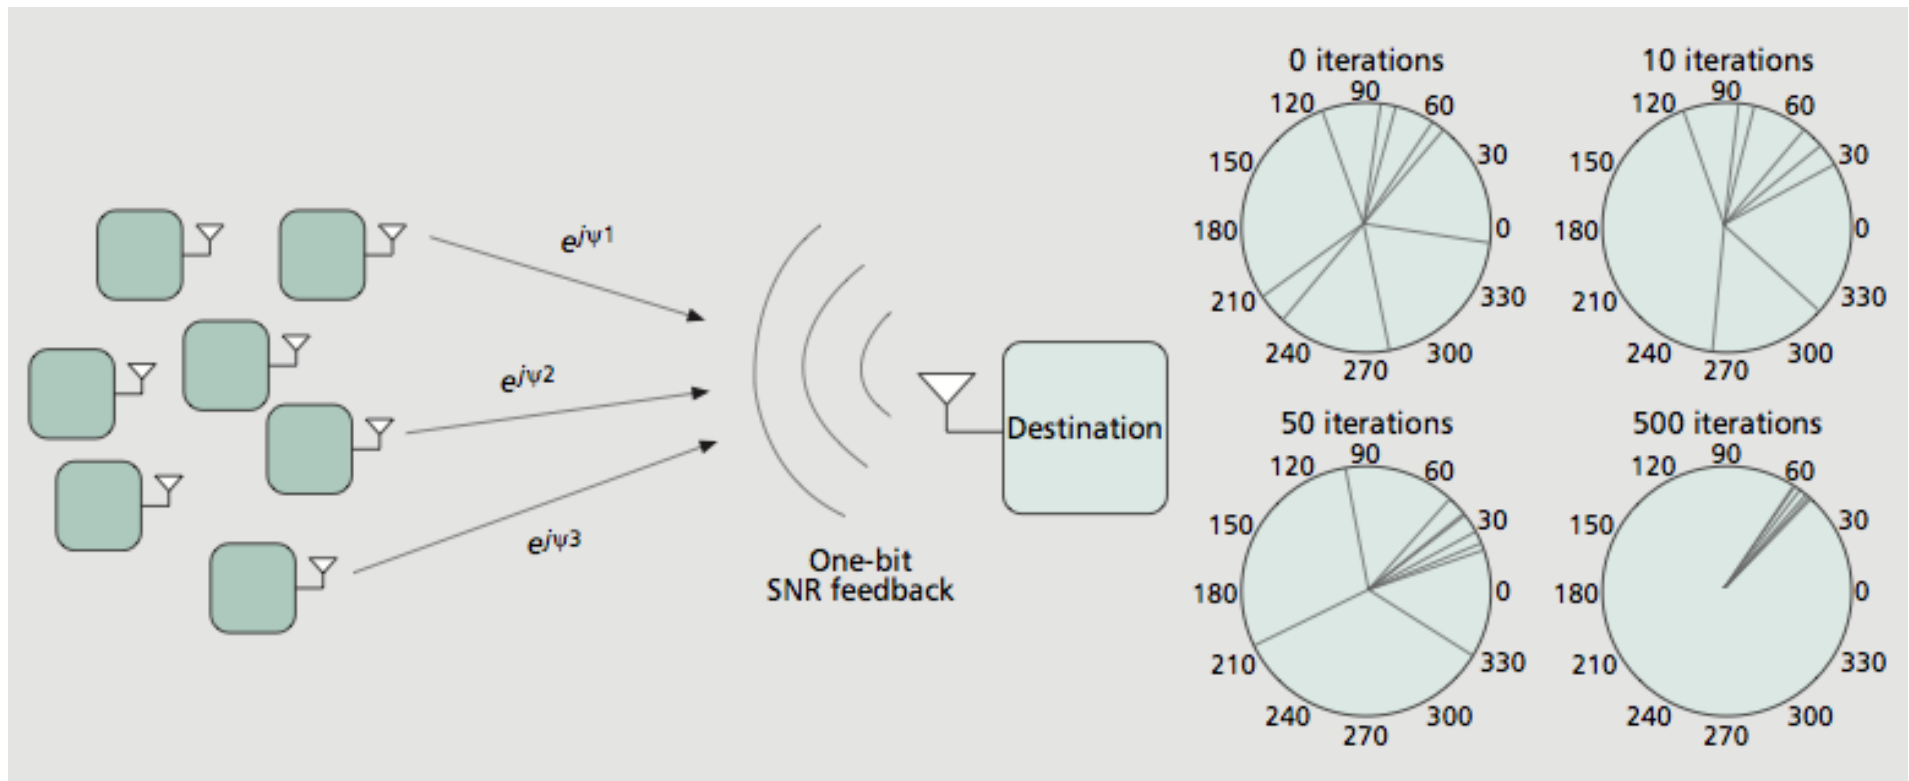
Main message: cool theory to be done, SDRs and open source enable quick transition to demos and community engagement

Transmit Beamforming

Can we scale to an arbitrarily large number of cooperating transmitters?

Aggregate feedback is key.

How it started for us (almost a decade ago)



Decentralized randomized ascent based on one bit feedback

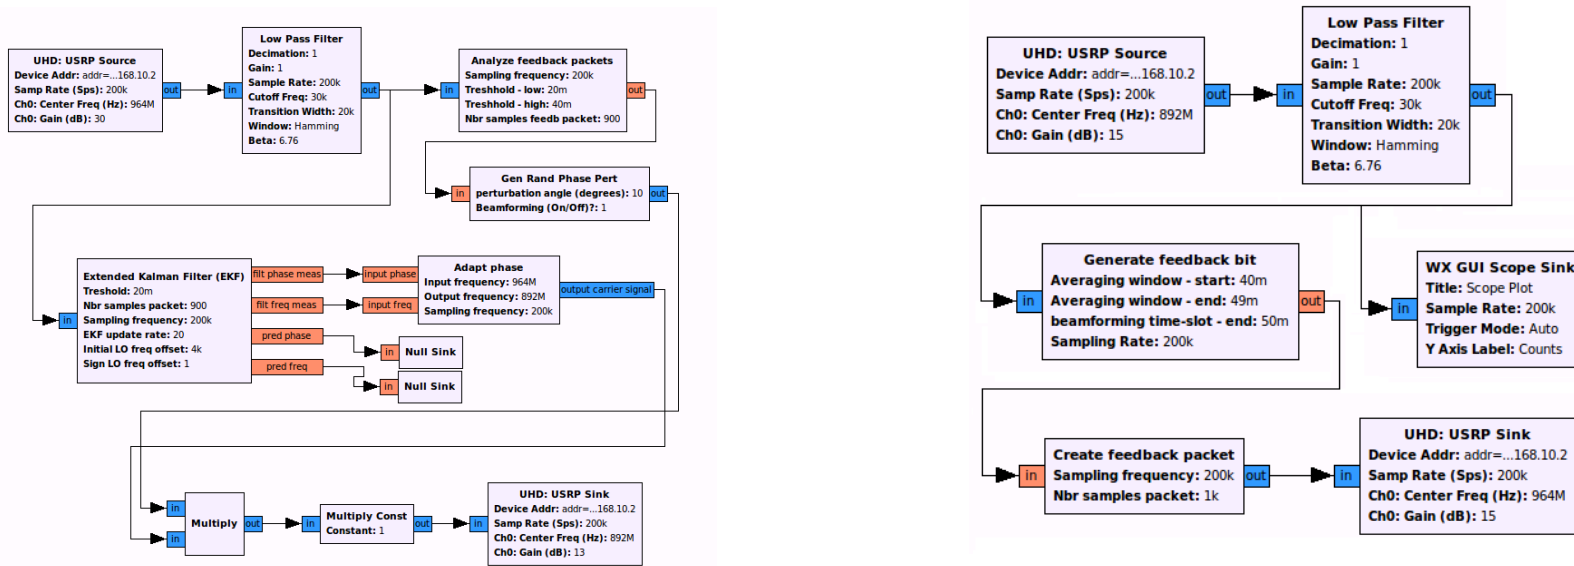
Converges with probability one in idealized setting

Evolution well characterized by statistical mechanics arguments

Amenable to simple implementation

(Invented by Raghu Mudumbai in 2005, first prototype by Ben Wild in 2006)


Today: all-wireless demo (software-defined radios)



Transmitter synchronize freq to receiver's using EKF
Use fb to adjust phase

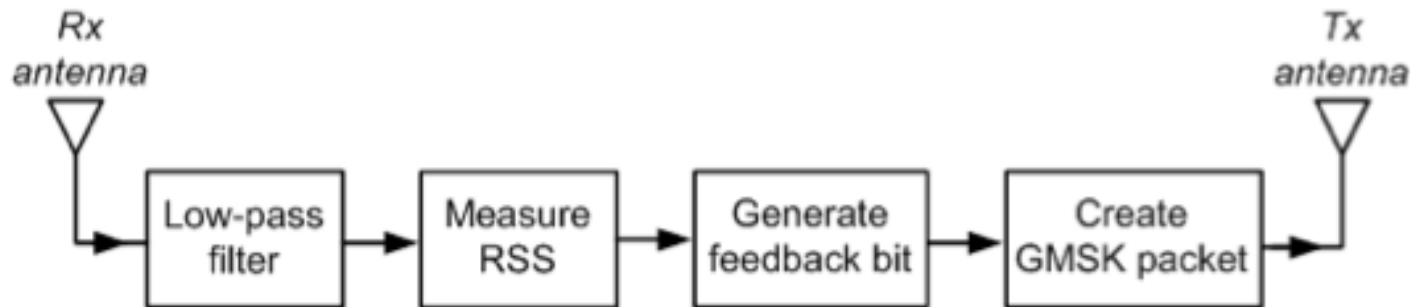
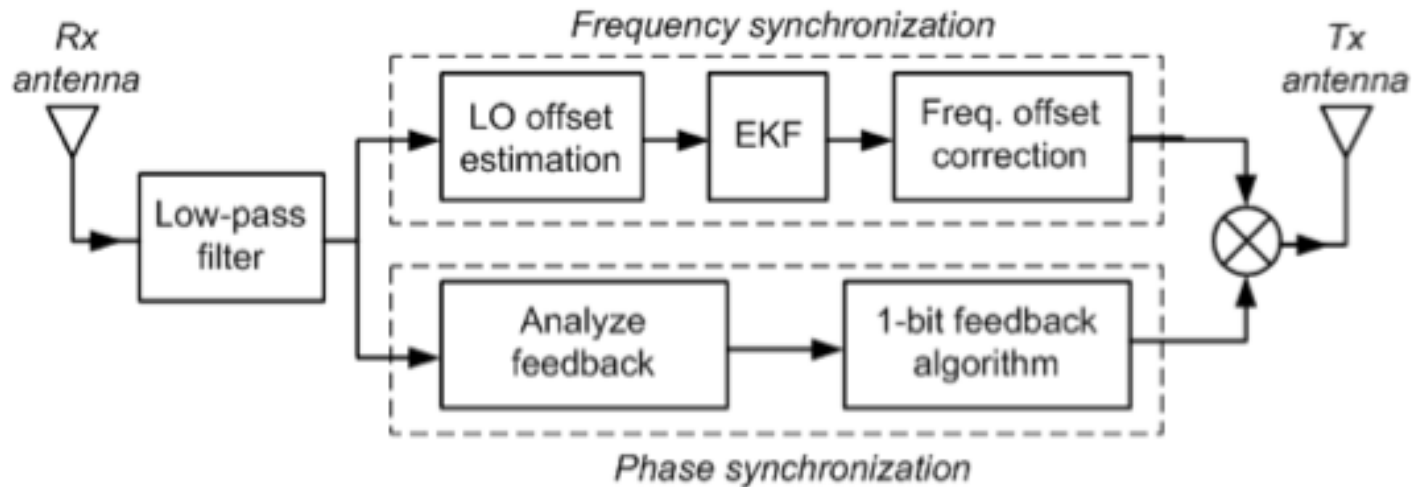
Receiver sends 1-bit fb packets

LIVE DEMO AT
WoWMoM 2012



Close to ideal beamforming
despite poor quality LOs

Architecture



Feedback packet does double duty

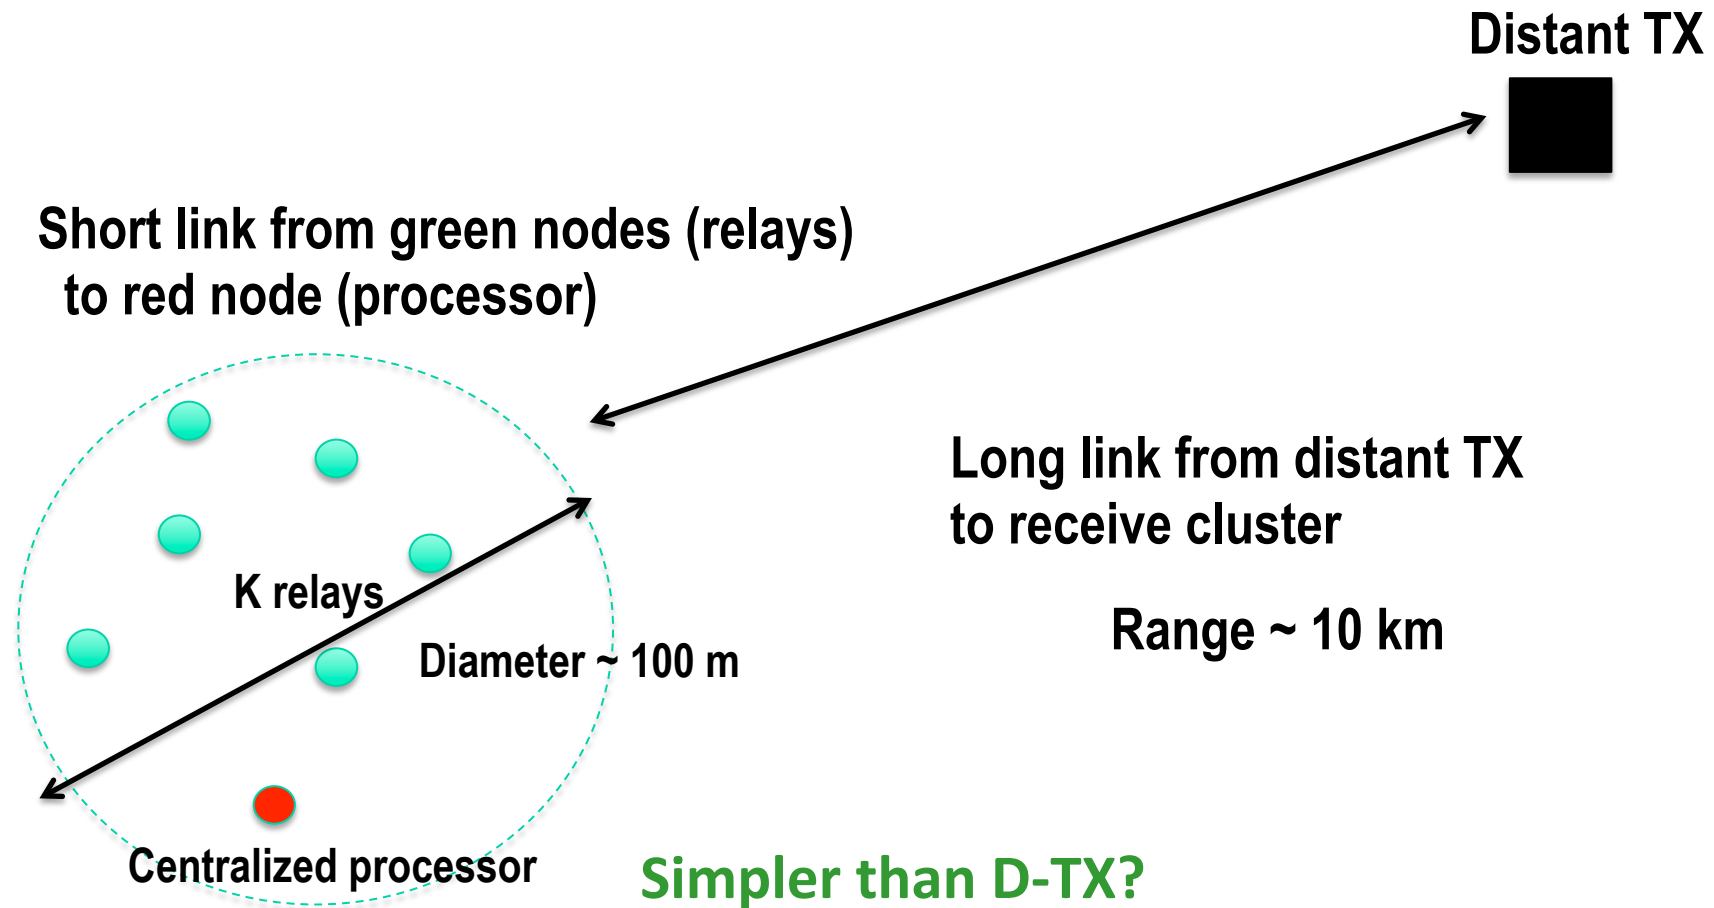
--Phase/frequency estimates from waveform drive state space model

--One bit of feedback in payload drives frequency sync

Distributed receive beamforming

Can we scale?

Distributed reception: system model




Simpler than D-TX?

- No need for relays to be synchronized
- Just send all received signals to processor


Problem: Does not scale

How to scale D-RX?

- How about combining in the air?
- Turning D-RX on long link into D-TX on short link
 - Relays adjust phases for coherent combining at processor
 - **Amplify-forward, but actually paying attention to sync**
- Sync using feedback from processor to relays



TDD relay \rightarrow implicit frequency sync




Frequency sync achieved implicitly in TDD amplify-forward
(relay LO offsets cancel out over long and short links)

Implicit timing sync via message received on long link

\rightarrow Need only worry about phase sync


The fine print: For low-quality LOs, significant phase drift between RX and TX at relay

Demo: D-RX for windowed sinusoid



Evolution of received amplitude of relayed packets over multiple frames

Under the hood: phase/freq tracking



State Space Model



$$\begin{pmatrix} \phi_t \\ \omega_t \end{pmatrix} = \begin{pmatrix} 1 & T_s \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \phi_{t-1} \\ \omega_{t-1} \end{pmatrix} + \mathbf{v}_t \quad \text{Process noise}$$

Process Noise Covariance

$$\mathbf{Q} = \omega_c^2 q_1^2 \begin{pmatrix} T_s & 0 \\ 0 & 0 \end{pmatrix} + \omega_c^2 q_2^2 \begin{pmatrix} T_s^3 / 3 & T_s^2 / 2 \\ T_s^2 / 2 & T_s \end{pmatrix}$$

Phase drift term

Frequency drift term

*Standard vanilla state evolution. What's the problem?
Haven't you heard of the Kalman filter?*

The problem is the measurement model



Problem 1: Nonlinear measurements

Nice linear state space model is for the unwrapped phase

We can only measure the wrapped phase

*So what? Just design your system to avoid phase wrapping ambiguity.
OK, if overhead does not matter.*

Problem 2: Frequency aliasing with intermittent measurements

Measurements spaced by T_s incur periodic freq ambiguity of $1/T_s$

Big deal. Just make some frequency measurements.

OK, but only if we make measurement intervals large enough.

$$\sigma_\phi^2 \sim 1/SNR$$

$$\sigma_f^2 \sim 1/(\text{Measurement interval} \times SNR)$$

Performance of one-shot
phase-freq estimation

Phase/frequency tracking architecture

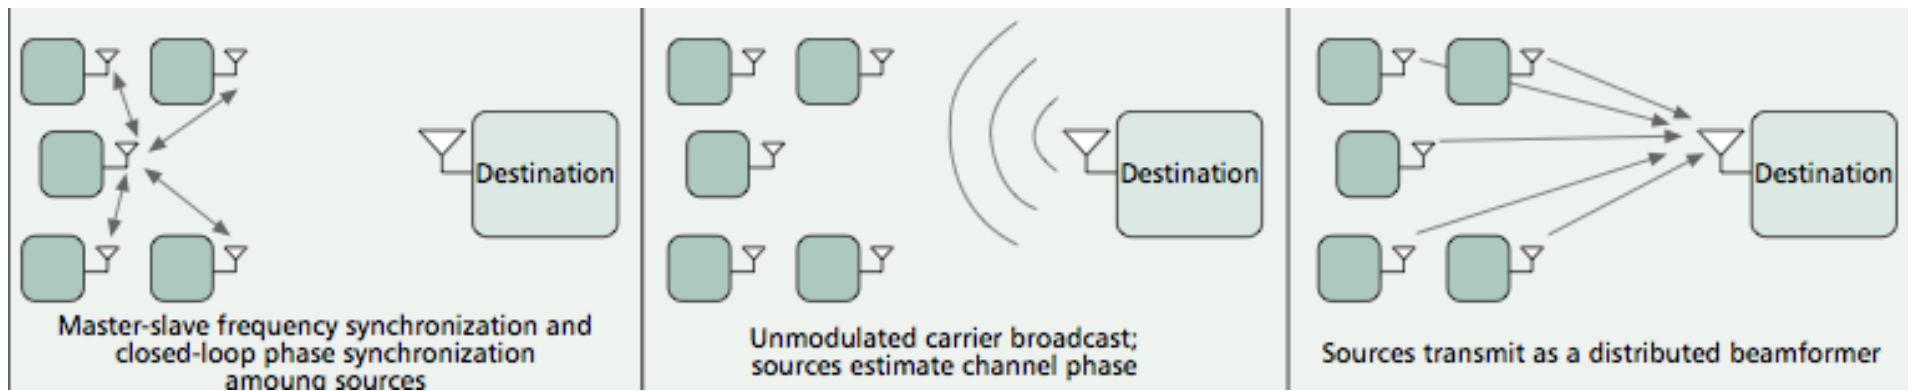


- **Kalman filter works just fine if system is designed to avoid ambiguities**
 - BBN demo
- **Extended Kalman filter to handle measurement nonlinearity**
 - That's all we need if we can avoid frequency aliasing
 - UCSB/U Iowa demo
- **Need to work harder to minimize overhead**
 - Need to handle measurement nonlinearity and frequency aliasing
 - Ongoing research...

An alternative to explicit feedback

Scaling via implicit feedback

Presynchronize the distributed array
Then use implicit feedback (reciprocity)



Scalable

May work in highly mobile settings

How well can we pre-synchronize?

Early indicators are promising...

BBN/Raytheon demo with picoseconds accuracy

Under the hood: one-shot timing estimation

Fundamentals of one-shot estimation



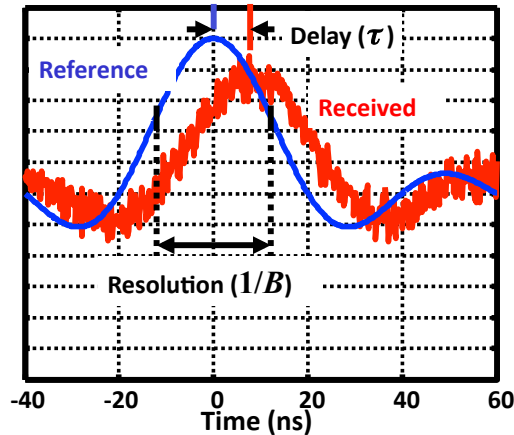
- **Two regimes in parameter estimation**
 - Coarse estimation: identify the right bin
 - Fine-grained estimation: refine within the bin
- **Cramer-Rao lower bound applies to fine-grained estimation**
 - Assumes we are close to the right value
- **Ziv-Zakai bound accounts for both regimes**
 - Coarse estimation errors at low SNR
 - Tends to CRLB at high enough SNR

Reaching fundamental limits in timing sync

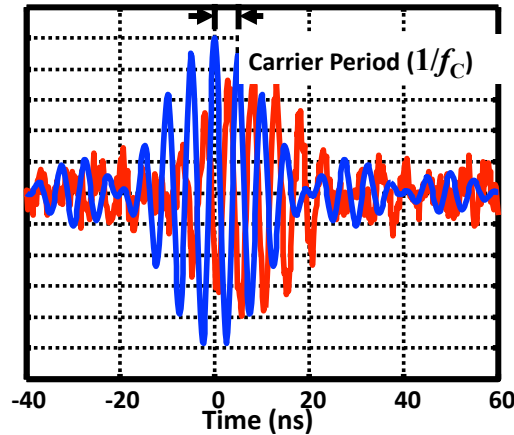
Accuracy within *small fraction of carrier period* with sufficient SNR



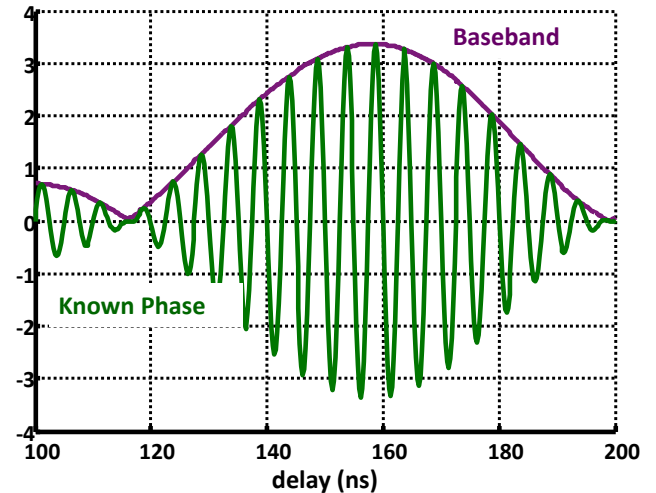
Baseband Waveform



Waveform on Carrier



Likelihood Function



Baseband CRLB:

$$\text{var}_{BB}(\hat{\tau}) \geq \frac{3/\pi^2}{BT \text{ SNR } B^2}$$

Post-Integration SNR Square Bandwidth


Known Phase CRLB:

$$\text{var}_{NB}(\hat{\tau}) \geq \frac{1/8\pi^2}{BT \text{ SNR } f_c^2}$$

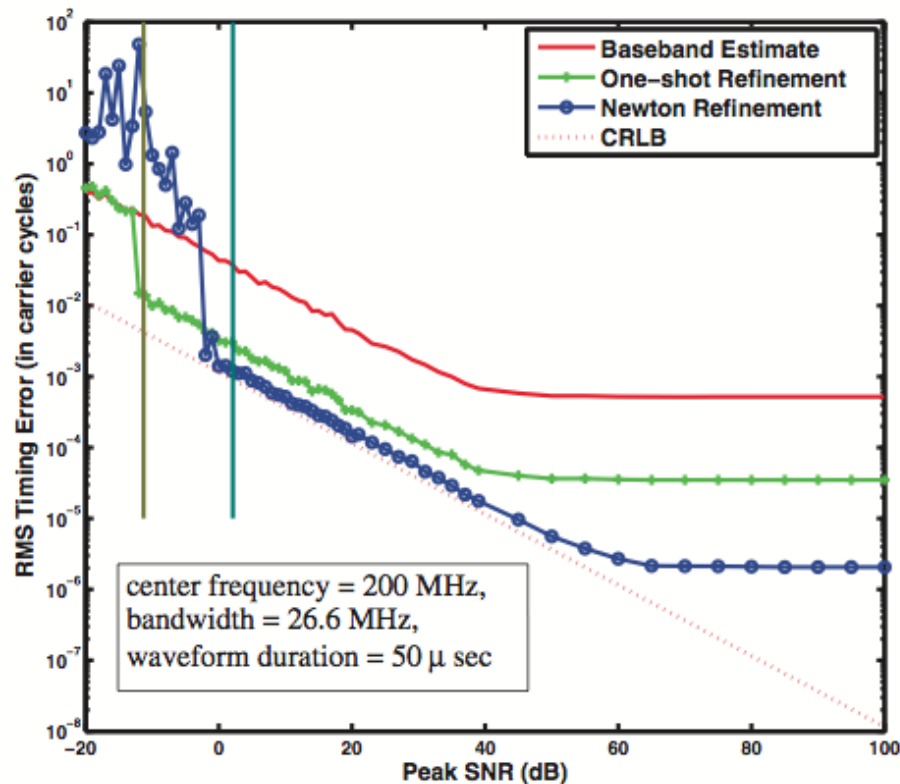
Post-Integration SNR Square Frequency

Hardware demo with picoseconds accuracy shown by BBN

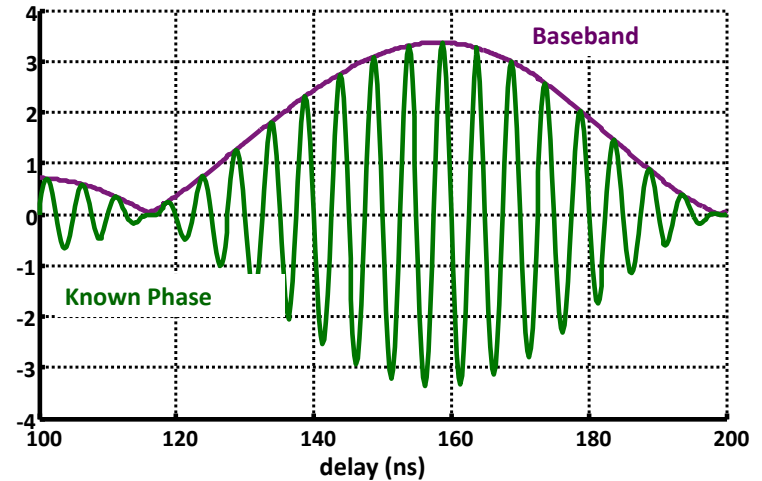
Example: $f_c = 1 \text{ GHz}$, $B = 50 \text{ MHz}$, $T = 10 \mu\text{s}$, $\text{SNR} = 10 \text{ dB}$



Approaching the ZZB for timing estimation



Likelihood Function



- Three stage algorithm
- Hypothesis testing
- Baseband refinement
- Passband refinement

At high enough SNR, can get to within a tiny fraction of a carrier cycle

Summary and Open Issues

Summary of recent progress



- **Narrowband D-TX and D-RX demos**
 - **UCSB, U Iowa: Software-defined radios, aggregate feedback, indoors**
 - **BBN/Raytheon: Mildly customized radios, per-node feedback, 1 km outdoors**
- **Fundamental timing sync bounds attained**
 - **Picoseconds accuracy demonstrated by BBN/Raytheon**
- **D-RX with hard decision exchanges shown to work at arbitrarily low SNRs**
 - **Information-theoretic analysis shows 1-2 dB loss relative to ideal receive beamforming**
- **Progress on scalable distributed nullforming algorithms**

Open Issues



- **Beyond the narrowband model**
 - Estimation/tracking fundamentals for dispersive channels and drifting LOs
 - Per-subcarrier tracking in OFDM likely overkill
- **Limits of aggregate feedback**
 - Effect of phase noise
 - Fast enough for highly mobile settings?
 - How about interference? Can you steer nulls with aggregate feedback?
- **Making implicit feedback work is critical for high mobility**
 - How well can we presynchronize?
 - Mismatch between transmit and receive chains
- **Synchronization-enabled protocols for concept systems**
 - D911, DBS

D-MIMO: exploring further



One-bit algorithm fundamentals

Mudumbai et al, *Distributed transmit beamforming using feedback control*, IEEE Trans. Information Theory, Jan 2010.

SDR Testbed

Quitin, Rahman, Mudumbai, Madhow, *Demonstrating distributed transmit beamforming with software-defined radios*, WoWMoM 2012. **(live demo, BEST DEMO AWARD)**

Quitin, Rahman, Mudumbai, Madhow, *A Scalable Architecture for Distributed Transmit Beamforming with Commodity Radios: Design and Proof of Concept*, IEEE Trans. Wireless Communications, Dec 2013.

Quitin, Irish, Madhow, *Distributed receive beamforming: a scalable architecture and its proof of concept*, VTC 2013 (Spring).

Achieving fundamental limits of timing sync

Bidigare et al, *Attaining fundamental bounds on timing synchronization*, ICASSP 2012.

Bidigare et al, *Initial over-the-air performance assessment of ranging and clock synchronization using radio frequency signal exchange*, SSP 2012.

Per-user feedback based schemes

Brown et al, *Receiver-coordinated distributed transmit beamforming with kinematic tracking*, ICASSP 2012.

Brown et al, *Receiver-coordinated distributed transmit nullforming with channel state uncertainty*, ICASSP 2012.

D-RX with off-the-shelf radios

Brown et al, *Distributed Reception with Coarsely-Quantized Observation Exchanges*, CISS 2012.

Thanks to our collaborators, past and present

Dr. Francois Quitin, Dr. Mahboob Rahman

Andrew Irish, Amy Kumar, Maryam Eslami Rasekh

BBN/Raytheon team led by Dr. Pat Bidigare