

DMIMO: UCSB update

Upamanyu Madhow

ECE Dept

University of California, Santa Barbara

DMIMO Summit, August 18, 2014

What we promised NSF

- Fundamentals
 - Phase/freq tracking, extension to dispersive channels
 - Distributed TX, distributed RX, many-to-many DMIMO
 - Distributed communication schemes: aggregate fb, retrodirective, nullforming
- Concept System Designs
 - Distributed base station, distributed 911
- Testbed
 - Aggregate feedback, robust phase/freq tracking, OFDM, long-range
- DMIMO community building

UCSB: recent work

- Fundamentals
 - Nonlinear phase/freq tracking
 - Distributed RX: scalable amplify/forward approach
 - ➔ Understanding one-bit feedback with phase noise
 - Starting on DMIMO for frequency selective channels
- Concept Systems
 - (added one) Spatially multiplexed LoS DMIMO
 - ➔ Fundamentals on DoF with matrices of random phasors
- Testbed
 - D-TX, D-RX, starting on OFDM

Today's talk


- Nonlinear phase/frequency state space tracking
 - How well can we maintain sync with intermittent measurements?
- Scalable distributed RX
 - Review
 - Recent analytical characterization of 1-bit feedback with phase noise
- Spatially multiplexed LoS DMIMO
 - Review
 - Recent analytical characterization
- Future directions

Phase/frequency tracking

Maryam Eslami Rasekh, Raghu Mudumbai
(to appear, Asilomar 2014)

Model

- Want to track phase and frequency of a carrier when we get intermittent, windowed measurements
 - Carrier could be emitted locally by master node or could come from destination node
- Overhead reduction requires:



Standard State Space Model

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

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 Kalman filter model? Not quite...

Measurement model is nonlinear

We only have access to the wrapped phase

$$y_t = \begin{bmatrix} \cos(\varphi_t) \\ \sin(\varphi_t) \\ \omega_t \end{bmatrix} + e_t, \quad e_t \sim N(0, R)$$

I and Q readings obtained by integrating over measurement interval

Worst-case model

Measurement interval too small for accurate frequency estimation

Phase wrapping headaches

Problem 1: Phase wrapping ambiguity → linear state space model does not apply

Could design the system to avoid phase wrapping ambiguity and use linear model, at the cost of additional overhead.

Problem 2: Frequency aliasing with intermittent measurements

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

Could avoid with “accurate enough” frequency measurements, but this requires “large enough” measurement intervals.

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

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

**Performance of one-shot
phase-freq estimation**

A customized particle filter

- Multiple particles to deal with frequency aliasing
 - Dither inter-measurement times to eliminate ambiguities quickly → fewer particles
- For each particle, take advantage of linear process model (Rao-Blackwellization)
 - Deal with mild nonlinearity using EKF

Recall basic particle filter

Algorithm 4.5: Bootstrap particle filter

```
1 Initialization ( $t = 1$ ):  
2   Sample  $\mathbf{x}_1^i \sim \mu(x_1)$ .  
3   Compute the importance weights  $\tilde{\mathbf{w}}_1^i = g(y_1 \mid \mathbf{x}_1^i)$  and normalize,  
    $\mathbf{w}_1^i = \tilde{\mathbf{w}}_1^i / \sum_{j=1}^N \tilde{\mathbf{w}}_1^j$ .  
4 for  $t = 2$  to  $T$  do  
5   | Resample  $\{\mathbf{x}_{t-1}^i, \mathbf{w}_{t-1}^i\}$  resulting in equally weighted samples  $\{\tilde{\mathbf{x}}_{t-1}^i, 1/N\}$ .  
6   | Sample  $\mathbf{x}_t^i \sim f(x_t \mid \tilde{\mathbf{x}}_{t-1}^i)$ .  
7   | Compute the importance weights  $\tilde{\mathbf{w}}_t^i = g(y_t \mid \mathbf{x}_t^i)$  and normalize  
   |  $\mathbf{w}_t^i = \tilde{\mathbf{w}}_t^i / \sum_{j=1}^N \tilde{\mathbf{w}}_t^j$ .  
8 end
```

(from notes by Schon)

Rao-Blackwellized PF

- State space model with nonlinear measurement of unwrapped phase

$$\begin{bmatrix} \varphi_t \\ \omega_t \end{bmatrix} = \begin{bmatrix} 1 & T_s \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \varphi_{t-1} \\ \omega_{t-1} \end{bmatrix} + \bar{\mathbf{v}}_t, \quad \bar{\mathbf{v}}_t \sim N(\bar{\mathbf{0}}, Q)$$

$$y_t = \begin{bmatrix} \cos(\varphi_t) \\ \sin(\varphi_t) \end{bmatrix} + \bar{\mathbf{e}}_t, \quad \bar{\mathbf{e}}_t \sim N(\bar{\mathbf{0}}, R)$$

- Particle evolution: sampling based on measurement

$$\{\varphi_{t-1}^i, w_{t-1}^i, \bar{\omega}_{t-1}^i, P_{t-1}^i\}_{i=1}^N$$

$$\varphi_t^i \sim N(\varphi_{t-1}^i + \bar{\omega}_{t-1}^i T_s + \left[\text{mod}(y_t - (\varphi_{t-1}^i + \bar{\omega}_{t-1}^i T_s)) \right]_{\pi}^{\pi}, R)$$

Details (1)

- Linear state time/measurement update

$$\alpha_{t|t-1}^i = E[x_t^n | x_{t-1}^{l,i}, x_{t-1}^{n,i}] = \varphi_{t-1}^i + \bar{\omega}_{t-1}^i T_s$$

$$\bar{\omega}_{t|t-1}^i = E[x_t^{l,i} | x_t^{n,i}, x_{t-1}] = \bar{\omega}_{t-1}^i + (T_s P_{t-1}^i + q_{12})(T_s^2 P_{t-1}^i + q_{11})^{-1}(\varphi_t^i - \alpha_{t|t-1}^i)$$

$$P_{t|t-1}^i = P_{t-1}^i + q_{22} - (T_s P_{t-1}^i + q_{12})^2 (T_s^2 P_{t-1}^i + q_{11})^{-1}$$

Details (2)

- Update and normalize particle weights

$$\tilde{w}_t^i = \frac{N(y_t; \hat{y}_t^i, S_t^i) N(\varphi_t^i; \alpha_{t|t-1}^i, P_{t|t-1}^i)}{r_t(\varphi_t^i | \varphi_{t-1}^i, y_t)} w_{t-1}^i$$

$$\hat{y}_t^i = \begin{bmatrix} \cos(\varphi_t^i) \\ \sin(\varphi_t^i) \end{bmatrix}, \quad S_t^i = R$$

$$w_t^i = \text{normalize}\{\hat{w}_t^i\}_{i=1}^N$$

$$\bar{\omega}_t^i = \bar{\omega}_{t|t-1}^i$$

- Note that $y_t - \hat{y}_t^i$ is snapped to $[-\pi, \pi]$

Resample according to w_t^i and set $w_t^i = 1 / N$

Frequency aliasing

- When times between measurements are constant, all particles can latch on to a frequency aliased estimate
 - Dithering can remove this ambiguity

How dither works

Frequency offset of $\frac{2\pi i}{T_s}$ from true value

$\Rightarrow$ Phase offset from true value after time $(1 + \Delta)T_s$ equals $2\pi i\Delta \pmod{2\pi}$

$\Delta = 1/2$ sends phase offset for all odd i to $\pm \pi$


$\Rightarrow$ Can distinguish them from true frequency offset

$\Delta = 1/4$ sends phase for all odd $i/2$ to $\pm \pi$

Deterministic Dither

$\Delta = 2^{-n}$, $n = 1, 2, \dots, M$ covers offsets as high as $\pm 2^M \frac{2\pi}{T_s}$

$\Rightarrow$ Can set M based on maximum frequency uncertainty and then repeat the cycle



Values of i rejected at each iteration

Iter=1	0	1	2	3	4	5	6	7	8	9	10	11
Iter=2	0		2		4		6		8		10	
Iter=3	0				4				8			
Iter=4	0								8			

Adapting the number of particles


- Initial ambiguity greater than steady state
- Can reduce number of particles after convergence
 - Based on residual frequency uncertainty
- Can adapt frequency jumps by detecting loss of convergence
 - Re-initiate around last estimate with a larger number of particles

Detecting loss of convergence (unmodeled jumps)

- Keep track of mean of error for past N steps for each particle (could even be one particle)
- If no particle has mean error lower than threshold
 - ➔ no particle has the correct hypothesis
 - ➔ reinitiate M particles around last frequency estimate spanning $(-F_{\text{span}}, F_{\text{span}})$

Response to unmodeled jumps



- Error detection triggers frequency spreading of width $2F_{\text{span}}$
- If $\text{abs}(F_{\text{jump}})$ is smaller than F_{span} estimate will recover quickly



Response to unmodeled jumps

- When jump is too large for frequency spread: random walk starts, convergence could be quick or it could never happen (i.e. take thousands of iterations)
 - Probability distribution of convergence time depends on distance of jump from pull-in range
 - Worse with deterministic dither

Two instances for $F_{\text{span}}=100\text{Hz}$, $F_{\text{jump}}=197\text{ Hz}$



top: frequency error (Hz) and $f_{\text{err}} = 100 \text{ Hz}$

Design rules of thumb

Follow from a single simple observation

Wrapped phase measurement can only detect phase changes in $(-\pi, \pi)$

$\Rightarrow$ Process + measurement noise must be much smaller than π

$$\sigma_{\phi}^2 = \omega_c^2 q_1^2 T_s + \omega_c^2 q_2^2 \frac{T_s^3}{3} + r^* \ll \pi^2$$

Phase drift across measurements

Depends on oscillator quality and
time between measurements

Effective phase error/measurement


Depends on measurement interval
and SNR

Performance


- When it works, performance is as good as EKF with genie (selecting right particle)
- As long as EKF output error is low enough (significantly lower than π) convergence is maintained
 - Error tracking/reinitiation around last estimate handles convergence loss if it happens

$T_s=20\text{ms}$, datasets 1


internal clk 1, 20ms
cos(phase offset)



external clk 1, $T_s=20\text{ms}$
cos(phase offset)




oscillators' frequency offset (Hz)



amplitude of measured signal (dB)


oscillators' frequency offset (Hz)




amplitude of measured signal (dB)

$T_s=20\text{ms}$, datasets 2


internal clk 2, $T_s=20\text{ms}$
 $\cos(\text{phase offset})$




external clk 2, $T_s=20\text{ms}$
 $\cos(\text{phase offset})$



oscillators' frequency offset (Hz)




oscillators' frequency offset (Hz)




$T_s=20\text{ms}$, datasets 3


internal clk 3, $T_s=20\text{ms}$
cos(phase offset)




external clk 3, $T_s=20\text{ms}$
cos(phase offset)



oscillators' frequency offset (Hz)




oscillators' frequency offset (Hz)




Ts=50ms, datasets 3


internal clk 3, Ts=50ms
cos(phase offset)




external clk 3, Ts=50ms
cos(phase offset)



oscillators' frequency offset (Hz)




oscillators' frequency offset (Hz)




$T_s=80\text{ms}$, datasets 3


internal clk 3, $T_s=80\text{ms}$
cos(phase offset)




external clk 3, $T_s=80\text{ms}$
cos(phase offset)



oscillators' frequency offset (Hz)





oscillators' frequency offset (Hz)





$T_s=100\text{ms}$: Internal oscillator breaks down completely while
(better) external oscillator maintains performance


internal clk 3, $T_s=100\text{ms}$
 $\cos(\text{phase offset})$




internal clk 3, $T_s=100\text{ms}$
 $\cos(\text{phase offset})$



External oscillator limit





External oscillator limit



External oscillator limit

external clk 3, $T_s=250\text{ms}$
 $\cos(\text{phase offset})$



Take-aways


- Rao-Blackwellized particle filter is an effective means of utilizing phase wrapped measurements
 - EKF handles local effect of nonlinearity
 - Essentially the same performance as for unwrapped measurements when it is working
 - Global effect of nonlinearity handled via dithering and particle filtering
 - Error monitoring detects loss of lock and reinitializes
- Solid building block to plug into synchronization-enabled protocols
 - Minimal overhead
 - Locks even with large frequency offsets
 - Robust to unmodeled jumps
 - Simple design rules of thumb

A scalable approach to D-RX beamforming

Francois Quitin, Andrew Irish


The need for scalability

important to avoid drowning the network with control messages.




Review of D-Tx beamforming: testbed


Synchronization achieved using feedback from the receiver



Tx nodes perform the 3 synchronizations independently, based on the feedback



Scalable D-RX via amplify-forward




Use cooperating nodes as amplify-forward relays
Signals are summed ***over the air***
Feedback from RX used to adapt relays

Turns long-distance D-RX into local D-TX


Review of prior results

Phase synchronization

again achieved with one-bit feedback algorithm




Time synchronization



Relay nodes use ***message from Tx*** as a common timestamp

Forward message with ***fixed delay*** after receiving it from Tx


Frequency synchronization




...achieved *implicitly* with forwarding architecture !

D-Rx beamforming architecture


Tx node



Rx node




Relay node




D-Rx beamforming architecture

Relay node




D-Rx beamforming implementation

when transmitting pilot tone packets





D-Rx beamforming implementation

Rx packet amplitude over long time intervals




Relaxing system parameters T_d and T_c

with our experimental prototype



Relaxing system parameters T_d and T_c

Once phase noise std gets close to phase perturbation size, **convergence and stability** of 1-bit feedback **becomes difficult**



New results: quantitative
understanding of phase noise effects

(journal paper in preparation)

Phase error due to relaying delay

LO drift between moment where relay node receives message and forwards it

=> results in **phase error**

=> can be modelled analytically

$$\sigma_{\phi_i}^2 = \underbrace{\omega_c^2 q_1^2 T_d + \frac{\omega_c^2 q_2^2 T_d^3}{3}}_{\text{Intra-cycle drift}} + \underbrace{\omega_c^2 q_2^2 T_c T_d^2}_{\text{Inter-cycle drift}}$$

q_1 --> white frequency noise

q_2 --> random walk frequency noise

T_d --> relay delay time

T_c --> cycle time


One-bit feedback with phase noise

Approximate analysis framework

- 1) Joint distribution of change in received complex amplitudes with and without phase error modeled as complex Gaussian
- 2) Statistical mechanics approach from original one-bit paper used to approximate joint distribution
- 3) Markov model for system state (RSS compared to that of K prior iterations)
- 4) Estimate RSS drift: at what value of RSS does it become negative?

Bottomline: It works as long as phase noise is “smaller than” random phase perturbations in one-bit algorithm

Joint Gaussian distribution




X-axis: RSS increment with phase noise



Y-axis: ideal RSS increment

Increased phase noise $\rightarrow$ less correlation


Markov model for feedback generation



Accurate predictions of state probs




Results from drift analysis



Predicts RSS saturation short of ideal value as phase noise increases

Scalable D-RX beamforming: take-aways

Scalable relaying architecture



Frequency synchronization not an issue


But excessive phase noise hurts phase sync

DMIMO for creating spatial degrees of freedom

Andrew Irish, Francois Quitin, Mark Rodwell

ITA 2013, Asilomar 2013

Motivation: 100 Gbps wireless over 50 km



Must throw everything we know at it

Bandwidth → mm wave band or higher

Power → not THz or optics


Directivity → mm wave band or higher

Spatial multiplexing → geometry must support full rank MIMO matrix

Polarimetric multiplexing → no conceptual hurdles, modulo hardware/signal processing design

LoS MIMO review

- Rayleigh criterion² relates *achievable* DOF to array lengths




Linear arrays: $DOF \approx \frac{L_T L_R}{R\lambda} + 1$

- Easily generalizes to square arrays of side lengths L_T, L_R .

$$DOF \approx \left(\frac{L_T L_R}{R\lambda} \right)^2 + 1$$

²Torkildson, Madhow, and Rodwell, "Indoor Millimeter Wave MIMO: Feasibility and Performance".

LoS MIMO does not work at long ranges



Example

75 GHz carrier frequency, 50 km range

Two-fold spatial multiplexing

$$d_T d_R = 100 \text{ m}^2$$


Subarrays 1 m apart on aircraft

➔ Subarrays 100 m apart on the ground!


This picture does not work!

Enter LoS DMIMO

Synthesize full rank channel by spreading the receiver out



Anatomy of full rank DMIMO



How much should the relays be spread out?

Modeling relay geometry

TX1 TX2
● ●


Model for response of transmitters at relays

$$\mathbf{h}_1 = (e^{j\theta_{11}}, e^{j\theta_{12}}, e^{j\theta_{13}}, e^{j\theta_{14}})^T$$

$$\mathbf{h}_2 = (e^{j\theta_{21}}, e^{j\theta_{22}}, e^{j\theta_{23}}, e^{j\theta_{24}})^T$$

$$\theta_{ij} \text{ i.i.d., } Unif[0, 2\pi]$$

(for “large enough” dispersal area)



Randomly dispersed relays

Rule of thumb for spacing

- Required region is largest when $l - k = 1$. Result (1D):



$$x_m \sim U\left(-\frac{\lambda R}{2d_T}, \frac{\lambda R}{2d_T}\right) \Rightarrow \text{on avg. spatial responses uncorrelated}$$

- **Generalization to 2D?**
 - ▶ Assume alignment of relay region with array

Theorem

Consider a regular square Tx array with minimum inter-element spacing of d_T . If relays are placed uniformly in a square of side $\frac{\lambda R}{d_T}$ then spatial responses are uncorrelated on average.


Verifying the rule of thumb



$\Rightarrow$ "best" $d_{R,max} = \frac{\lambda R}{d_T} = 200 \text{ m}$

- Validates predicted best $d_{R,max} = 200 \text{ m}$

Performance prediction via random matrices



- Relay spread “large enough” $\Rightarrow \mathbf{H}$ well-modeled as **finite dimensional i.i.d. random phasor matrix**

$$[H]_{i,k} \approx e^{j\theta_{i,k}}, \quad \theta_{i,k} \text{ i.i.d. uniform}$$

- Most methods of random matrix analysis for Gaussian / binary / asymptotic cases

- **Question:** How do system dimensions N_R (relays) and N_T (TX) relate to ZF noise enhancement?
- **Answers:**
 - ▶ Approach # 1: *bound* the ZF SNR loss
 - ▶ Approach # 2: *approximate* ZF SNR loss

Chebyshev Bound on ZF SNR

Theorem

Let $\rho_i = \frac{1}{N_R} \langle \mathbf{h}_1, \mathbf{h}_i \rangle$ and $X = \sum_{i=2}^{N_T} |\rho_i|^2$. For $0 < \delta < \frac{N_R - N_T + 1}{N_R}$, the ZF gain $\Delta \in [0, 1]$ is upper-bounded as follows

$$\Pr[\Delta < \delta] \leq \frac{\mathcal{C}(X; k)}{\left(\frac{N_R - N_T + 1}{N_R} - \delta\right)^k}$$

where $\mathcal{C}(X; k)$ is the k th central moment of X , computed analytically.

- **Note:** For the 2×2 system, an exact form is known³

$$\Pr[\Delta < \delta] = \frac{2}{\pi} \arccos \sqrt{1 - \delta}$$

Chebyshev bound: details

- **Key steps:**

- ▶ WLOG, focus on ZF SNR gain Δ of stream 1 (i.i.d. entries of $\mathbf{H}$)

- ▶ Worst case ZF: all interfering streams $\mathbf{h}_i$ are orthogonal

$$\Rightarrow \Delta \geq 1 - \sum_{i=2}^{N_T} |\rho_i|^2, \quad \rho_i = \frac{1}{N_R} \langle \mathbf{h}_1, \mathbf{h}_i \rangle = \frac{1}{N_R} \sum_{k=1}^{N_R} e^{j\theta_{ki}}$$

- ▶ Analyze **random phasor walk** $\rightarrow$ Raw moments of $|\rho_i|^2$

- ★ Combinatorics problem: solved using theory of “Uniform block permutations”

- ▶ Apply bi/multinomial theorems $\rightarrow$ Central moments of $X = \sum_{i=2}^{N_T} |\rho_i|^2$

- ▶ Apply Markov inequality on $(X - \mathbb{E}(X))^k \rightarrow$ Chebyshev bound on $1 - \Delta$

Beta approximation for ZF SNR

Theorem

For $N_R \geq N_T > 2$, the ZF gain $\Delta \in [0, 1]$ approximately follows a beta distribution:

$$\Delta \sim \text{Beta}(N_R - N_T + 1, N_T - 1)$$

- **Note:** For $N_T = 2$, a tighter approximation is known⁴



$$\Delta = 1 - |\rho|^2 \text{ where } |\rho|^2 \sim \text{Exp}(N_R)$$

Key ideas:

Interference subspace randomly oriented wrt desired signal (CLT for large number of streams)



- ➔ Desired vector randomly oriented wrt interference subspace
- ➔ Can replace desired vector with iid complex Gaussian entries
- ➔ ZF SNR = ratio of chi-squared random variables (beta random variable)

Numerical results: Chebyshev bound



- **5 % Chebyshev bound is tight:** optimized over $2 \leq k \leq 30$
 - ▶ $\{3.81, 0.43, 0.15\}$ dB gap for $N_R \in \{8, 16, 32\}$
 - ▶ For 4×4 system no feasible 5 % Chebyshev bound found

Numerical results: Beta approximation



- ZF SNR improved by increasing the number of relays
- Theoretical (Beta) approx.
 - ▶ Slightly off for the 4×4 system
 - ▶ For ≥ 8 relays is within 2 dB of the simulations.

Take-aways

- DMIMO as an enabler of long-range “wireless fiber”
 - 5 GHz x dual polarization x 4-fold spatial multiplexing x 2.5 bps/Hz = 100 Gbps
- Analytical rules of thumb and performance predictions that closely match simulations
- Significant implementation challenges remain
 - Hardware/signal processing co-design for relays and receiver

Conclusions

- DMIMO remains in its infancy
 - Figuring out basic building blocks
- Still a lot of work on fundamentals
 - Wideband (frequency selective) channels
 - Reciprocity
- Concept system designs
 - Distributed base station, distributed 911, CoMP
- Enhancing testbed capabilities
 - Wideband, long-range