A Clean-Slate Design of Wireless Ad Hoc Networks Using On-Off-Division Duplex

Dongning Guo

with Lei Zhang, Jun Luo and Kai Shen (thanks to Martin Haenggi)

Dept. of EECS
Northwestern University

Institute of Network Coding
Chinese University of Hong Kong

Presented at the INC
October 6, 2010

Broadcast & Superposition

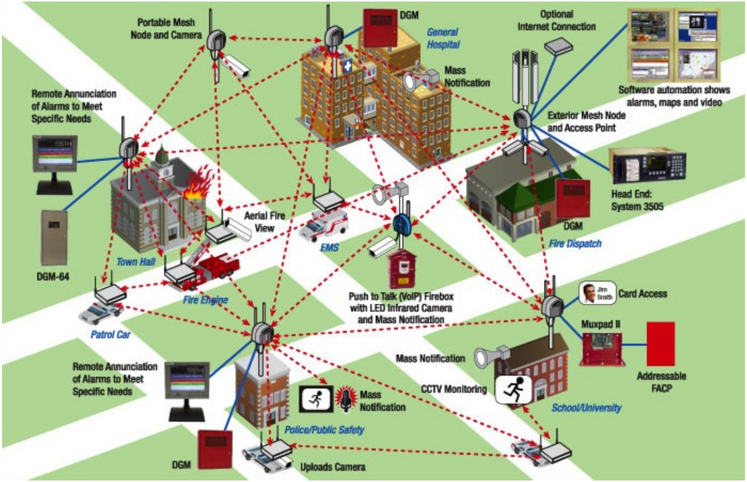

Half-Duplex Radio

self-interference

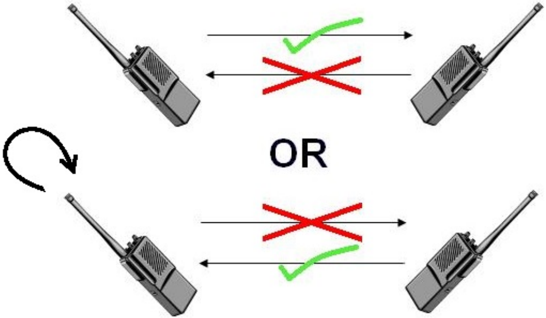

Existing Duplex Schemes

Frequency-division duplex (FDD)

Time-division duplex (TDD)

Code-division duplex (CDD) [Asada et al '96]

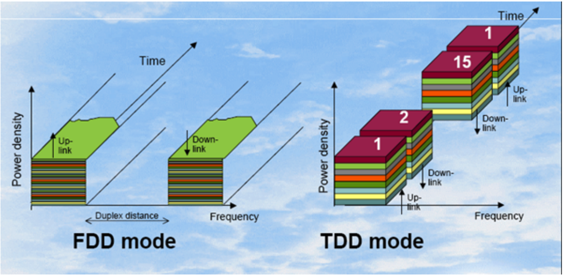

Existing Duplex Schemes

Frequency-division duplex (FDD)

Time-division duplex (TDD)

Code-division duplex (CDD) [Asada et al '96]

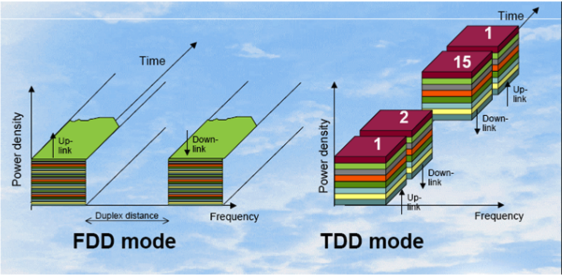

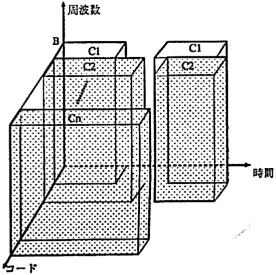

New Scheme: Rapid On-Off-Division Duplex (RODD)

- Half-duplex $\Leftrightarrow$ received signal erased by own transmissions
- No need to transmit a whole frame before listening
- New idea:
  - on-off signaling at symbol level
  - listening during off-slots within the same frame

New Scheme: Rapid On-Off-Division Duplex (RODD)

- Half-duplex $\Leftrightarrow$ received signal erased by own transmissions
- No need to transmit a whole frame before listening
- New idea:
  - on-off signaling at symbol level
  - listening during off-slots within the same frame

New Scheme: Rapid On-Off-Division Duplex (RODD)

- Half-duplex ⇔ received signal erased by own transmissions
- No need to transmit a whole frame before listening
- New idea:
  - on-off signaling at symbol level
  - listening during off-slots within the same frame

RODD with Multiple Users

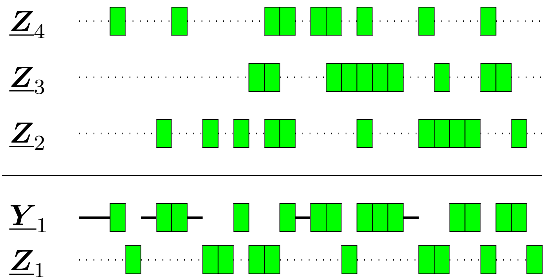

Multiaccess channel (MAC) with erasure

RODD with Multiple Users

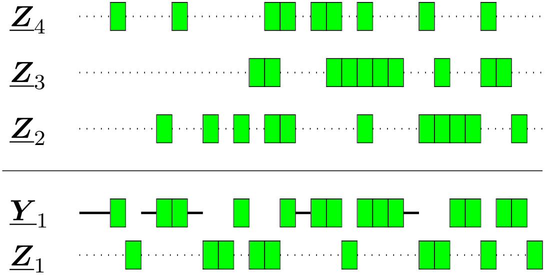

Multiaccess channel (MAC) with erasure

Advantages of RODD

- Enabling virtual full-duplex communication using half-duplex radios
- Scheduling in a microscopic timescale
- Simplification of higher-layer protocols
- Can take full advantages of broadcast & superposition
- Highly efficient in case of mutual broadcast traffic
- Small or stable access delay
- ...

A clean-slate design

Advantages of RODD

- Enabling virtual full-duplex communication using half-duplex radios
- Scheduling in a microscopic timescale
- Simplification of higher-layer protocols
- Can take full advantages of broadcast & superposition
- Highly efficient in case of mutual broadcast traffic
- Small or stable access delay
- ...

A clean-slate design

Advantages of RODD

- Enabling virtual full-duplex communication using half-duplex radios
- Scheduling in a microscopic timescale
- Simplification of higher-layer protocols
- Can take full advantages of broadcast & superposition
- Highly efficient in case of mutual broadcast traffic
- Small or stable access delay
- ...

A clean-slate design

Advantages of RODD

- Enabling virtual full-duplex communication using half-duplex radios
- Scheduling in a microscopic timescale
- Simplification of higher-layer protocols
- Can take full advantages of broadcast & superposition
- Highly efficient in case of mutual broadcast traffic
- Small or stable access delay
- ...

A clean-slate design

Advantages of RODD

- Enabling virtual full-duplex communication using half-duplex radios
- Scheduling in a microscopic timescale
- Simplification of higher-layer protocols
- Can take full advantages of broadcast & superposition
- Highly efficient in case of mutual broadcast traffic
- Small or stable access delay
- ...

A clean-slate design

Advantages of RODD

- Enabling virtual full-duplex communication using half-duplex radios
- Scheduling in a microscopic timescale
- Simplification of higher-layer protocols
- Can take full advantages of broadcast & superposition
- Highly efficient in case of mutual broadcast traffic
- Small or stable access delay
- ...

A clean-slate design

Advantages of RODD

- Enabling virtual full-duplex communication using half-duplex radios
- Scheduling in a microscopic timescale
- Simplification of higher-layer protocols
- Can take full advantages of broadcast & superposition
- Highly efficient in case of mutual broadcast traffic
- Small or stable access delay
- ...
- A clean-slate design

Advantages of RODD

- Enabling virtual full-duplex communication using half-duplex radios
- Scheduling in a microscopic timescale
- Simplification of higher-layer protocols
- Can take full advantages of broadcast & superposition
- Highly efficient in case of mutual broadcast traffic
- Small or stable access delay
- ...

A clean-slate design

Practical Considerations

On-off at symbol level ($\sim$10 μs)

- ✓ Response time of RF circuits in sub-nanoseconds
- ✓ Time-hopping impulse radio (sub-nanosecond monocycle) [Scholtz '93, Win & Scholtz '98]
- ✓ GSM uses on-off over sub-millisecond slots

Synchronicity

- ✓ Not a necessity (albeit nice to have)
- ✓ Propagation delay $\ll$ symbol interval
- ✓ Local synchronicity achievable using consensus algorithms (e.g., [Schizas, Ribeiro, Giannakis & Roumeliotis '08])
- ✓ Shortcut: synchronize to GPS or cellular networks

Practical Considerations

On-off at symbol level ($\sim$10 μs)

- ✓ Response time of RF circuits in sub-nanoseconds
- ✓ Time-hopping impulse radio (sub-nanosecond monocycle) [Scholtz '93, Win & Scholtz '98]
- ✓ GSM uses on-off over sub-millisecond slots

Synchronicity

- ✓ Not a necessity (albeit nice to have)
- ✓ Propagation delay $\ll$ symbol interval
- ✓ Local synchronicity achievable using consensus algorithms (e.g., [Schizas, Ribeiro, Giannakis & Roumeliotis '08])
- ✓ Shortcut: synchronize to GPS or cellular networks

Practical Considerations

On-off at symbol level ($\sim$10 μs)

- ✓ Response time of RF circuits in sub-nanoseconds
- ✓ Time-hopping impulse radio (sub-nanosecond monocycle) [Scholtz '93, Win & Scholtz '98]
- ✓ GSM uses on-off over sub-millisecond slots

Synchronicity

- ✓ Not a necessity (albeit nice to have)
- ✓ Propagation delay $\ll$ symbol interval
- ✓ Local synchronicity achievable using consensus algorithms (e.g., [Schizas, Ribeiro, Giannakis & Roumeliotis '08])
- ✓ Shortcut: synchronize to GPS or cellular networks

Practical Considerations

On-off at symbol level ($\sim$10 μs)

- ✓ Response time of RF circuits in sub-nanoseconds
- ✓ Time-hopping impulse radio (sub-nanosecond monocycle) [Scholtz '93, Win & Scholtz '98]
- ✓ GSM uses on-off over sub-millisecond slots

Synchronicity

- ✓ Not a necessity (albeit nice to have)
- ✓ Propagation delay $\ll$ symbol interval
- ✓ Local synchronicity achievable using consensus algorithms (e.g., [Schizas, Ribeiro, Giannakis & Roumeliotis '08])
- ✓ Shortcut: synchronize to GPS or cellular networks

Practical Considerations

On-off at symbol level ($\sim$10 μs)

- ✓ Response time of RF circuits in sub-nanoseconds
- ✓ Time-hopping impulse radio (sub-nanosecond monocycle) [Scholtz '93, Win & Scholtz '98]
- ✓ GSM uses on-off over sub-millisecond slots

Synchronicity

- ✓ Not a necessity (albeit nice to have)
- ✓ Propagation delay $\ll$ symbol interval
- ✓ Local synchronicity achievable using consensus algorithms (e.g., [Schizas, Ribeiro, Giannakis & Roumeliotis '08])
- ✓ Shortcut: synchronize to GPS or cellular networks

Practical Considerations

On-off at symbol level ($\sim$10 μs)

- ✓ Response time of RF circuits in sub-nanoseconds
- ✓ Time-hopping impulse radio (sub-nanosecond monocycle) [Scholtz '93, Win & Scholtz '98]
- ✓ GSM uses on-off over sub-millisecond slots

Synchronicity

- ✓ Not a necessity (albeit nice to have)
- ✓ Propagation delay $\ll$ symbol interval
- ✓ Local synchronicity achievable using consensus algorithms (e.g., [Schizas, Ribeiro, Giannakis & Roumeliotis '08])
- ✓ Shortcut: synchronize to GPS or cellular networks

Practical Considerations

On-off at symbol level ($\sim$10 μs)

- ✓ Response time of RF circuits in sub-nanoseconds
- ✓ Time-hopping impulse radio (sub-nanosecond monocycle) [Scholtz '93, Win & Scholtz '98]
- ✓ GSM uses on-off over sub-millisecond slots

Synchronicity

- ✓ Not a necessity (albeit nice to have)
- ✓ Propagation delay $\ll$ symbol interval
- ✓ Local synchronicity achievable using consensus algorithms (e.g., [Schizas, Ribeiro, Giannakis & Roumeliotis '08])
- ✓ Shortcut: synchronize to GPS or cellular networks

Practical Considerations

On-off at symbol level ($\sim$10 μs)

- ✓ Response time of RF circuits in sub-nanoseconds
- ✓ Time-hopping impulse radio (sub-nanosecond monocycle) [Scholtz '93, Win & Scholtz '98]
- ✓ GSM uses on-off over sub-millisecond slots

Synchronicity

- ✓ Not a necessity (albeit nice to have)
- ✓ Propagation delay $\ll$ symbol interval
- ✓ Local synchronicity achievable using consensus algorithms (e.g., [Schizas, Ribeiro, Giannakis & Roumeliotis '08])
- ✓ Shortcut: synchronize to GPS or cellular networks

Outline of Results

1. Preliminary results on capacity
2. Neighbor discovery
3. Mutual broadcast
4. Research questions

Result I: Preliminary Results on Capacity

Channel Model

- N nodes
- A frame consists of M symbols/slots/measurements
- Perfect synchronicity
- Binary duplex mask (signature) of node n

$$\boldsymbol{s}_n = [s_{n1}, \ldots, s_{nM}]$$

- MAC with erasure:

$$Y_{nm} = (1 - s_{nm}) \sum_{j \in \partial n} d_{nj}^{-\alpha/2} h_{nj} s_{jm} \sqrt{\gamma_j}\, X_{jm} + V_{nm}$$

$\sum_{m=1}^{M} s_{nm} x_{nm}^2 \leq M$

Channel Model

- N nodes
- A frame consists of M symbols/slots/measurements
- Perfect synchronicity
- Binary duplex mask (signature) of node n

$$\boldsymbol{s}_n = [s_{n1}, \dots, s_{nM}]$$

- MAC with erasure:

$$Y_{nm} = (1 - s_{nm}) \sum_{j \in \partial n} d_{nj}^{-\alpha/2} h_{nj} s_{jm} \sqrt{\gamma_j}\, X_{jm} + V_{nm}$$

$\sum_{m=1}^{M} s_{nm} x_{nm}^2 \leq M$

Channel Model

- N nodes
- A frame consists of M symbols/slots/measurements
- Perfect synchronicity
- Binary duplex mask (signature) of node n

$$\boldsymbol{s}_n = [s_{n1}, \dots, s_{nM}]$$

- MAC with erasure:

$$Y_{nm} = (1 - s_{nm}) \sum_{j \in \partial n} d_{nj}^{-\alpha/2} h_{nj} s_{jm} \sqrt{\gamma_j} X_{jm} + V_{nm}$$

$\sum_{m=1}^{M} s_{nm} x_{nm}^2 \leq M$

Channel Model

- N nodes
- A frame consists of M symbols/slots/measurements
- Perfect synchronicity
- Binary duplex mask (signature) of node n

$$\boldsymbol{s}_n = [s_{n1}, \ldots, s_{nM}]$$

- MAC with erasure:

$$Y_{nm} = (1 - s_{nm}) \sum_{j \in \partial n} d_{nj}^{-\alpha/2} h_{nj} s_{jm} \sqrt{\gamma_j}\, X_{jm} + V_{nm}$$

$\sum_{m=1}^{M} s_{nm} x_{nm}^2 \leq M$

Channel Model

- N nodes
- A frame consists of M symbols/slots/measurements
- Perfect synchronicity
- Binary duplex mask (signature) of node n

$$\boldsymbol{s}_n = [s_{n1}, \ldots, s_{nM}]$$

- MAC with erasure:

$$Y_{nm} = (1 - s_{nm}) \sum_{j \in \partial n} d_{nj}^{-\alpha/2} h_{nj} s_{jm} \sqrt{\gamma_j}\, X_{jm} + V_{nm}$$

$\sum_{m=1}^{M} s_{nm} x_{nm}^2 \leq M$

Deterministic Model

- N nodes
- M measurements/slots/symbols in each frame
- Perfect synchronicity
- Binary duplex mask (signature) of node n

$$\boldsymbol{s}_n = [s_{n1}, \dots, s_{nM}]$$

- Noncoherent energy detection
- Inclusive-OR MAC with erasure

$$\hat{Y}_{nm} = (1 - s_{nm}) \left(\bigvee_{j \in \partial n} (s_{jm} X_{jm}) \right)$$

Deterministic Model

- N nodes
- M measurements/slots/symbols in each frame
- Perfect synchronicity
- Binary duplex mask (signature) of node n

$$\boldsymbol{s}_n = [s_{n1}, \ldots, s_{nM}]$$

- Noncoherent energy detection
- Inclusive-OR MAC with erasure

$$\hat{Y}_{nm} = (1 - s_{nm}) \left(\bigvee_{j \in \partial n} (s_{jm} X_{jm}) \right)$$

One-Hop Broadcast Capacity

- Every node has K neighbors
- Everyone broadcasts a message to neighbors over an M-slot frame (multiple multicast sessions)
- $s_{km} \sim$ Bernoulli(q), i.i.d.
- $Pe(k)$: the probability that node k does not correctly decode all K messages from its neighbors
- A rate tuple is achievable if ∃ such a code with

$$\lim_{M\to\infty} \max_k Pe(k) = 0$$

- Codebooks depend on (K, M, q) but independent of the signatures and topology otherwise

One-Hop Broadcast Capacity

- Every node has K neighbors
- Everyone broadcasts a message to neighbors over an M-slot frame (multiple multicast sessions)
- $s_{km} \sim$ Bernoulli(q), i.i.d.
- $Pe(k)$: the probability that node k does not correctly decode all K messages from its neighbors
- A rate tuple is achievable if $\exists$ such a code with

$$\lim_{M\to\infty} \max_k Pe(k) = 0$$

- Codebooks depend on (K, M, q) but independent of the signatures and topology otherwise

One-Hop Broadcast Capacity

- Every node has K neighbors
- Everyone broadcasts a message to neighbors over an M-slot frame (multiple multicast sessions)
- $s_{km} \sim$ Bernoulli(q), i.i.d.
- $Pe(k)$: the probability that node k does not correctly decode all K messages from its neighbors
- A rate tuple is achievable if $\exists$ such a code with

$$\lim_{M \to \infty} \max_{k} Pe(k) = 0$$

- Codebooks depend on (K, M, q) but independent of the signatures and topology otherwise

One-Hop Broadcast Capacity

- Every node has K neighbors
- Everyone broadcasts a message to neighbors over an M-slot frame (multiple multicast sessions)
- $s_{km} \sim$ Bernoulli(q), i.i.d.
- $Pe(k)$: the probability that node k does not correctly decode all K messages from its neighbors
- A rate tuple is achievable if $\exists$ such a code with

$$\lim_{M \to \infty} \max_k Pe(k) = 0$$

- Codebooks depend on (K, M, q) but independent of the signatures and topology otherwise

One-Hop Broadcast Capacity

- Every node has K neighbors
- Everyone broadcasts a message to neighbors over an M-slot frame (multiple multicast sessions)
- $s_{km} \sim$ Bernoulli(q), i.i.d.
- $Pe(k)$: the probability that node k does not correctly decode all K messages from its neighbors
- A rate tuple is achievable if $\exists$ such a code with

$$\lim_{M \to \infty} \max_k Pe(k) = 0$$

- Codebooks depend on (K, M, q) but independent of the signatures and topology otherwise

One-Hop Broadcast Capacity

- Every node has K neighbors
- Everyone broadcasts a message to neighbors over an M-slot frame (multiple multicast sessions)
- $s_{km} \sim$ Bernoulli(q), i.i.d.
- $Pe(k)$: the probability that node k does not correctly decode all K messages from its neighbors
- A rate tuple is achievable if $\exists$ such a code with

$$\lim_{M\to\infty} \max_k Pe(k) = 0$$

- Codebooks depend on (K, M, q) but independent of the signatures and topology otherwise

Theorem [GZ '10] (Capacity of the deterministic model)

$$C = \frac{1-q}{K} \max_{p\in[0,1]} \sum_{\kappa=1}^{K} \binom{K}{\kappa} q^{\kappa}(1-q)^{K-\kappa} H_2(p^{\kappa})$$

Proof. $M \sum_{k=1}^{K} R_k \le I(Y_{01}, \dots, Y_{0M}; \boldsymbol{X}_1, \dots, \boldsymbol{X}_M | \underline{\boldsymbol{S}})$

$\le \sum_{m=1}^{M} I(Y_{0m}; X_{1m}, \dots, X_{Km} | S_{0m}, S_{1m}, \dots, S_{Km})$

$= \sum_{m=1}^{M} H(Y_{0m} | S_{0m}, S_{1m}, \dots, S_{Km})$

Bernoulli $(1-p)$ signaling is optimal by symmetry.
Conditioned on that κ out of the K neighbors transmit,

$$H\left(Y_{0m} \middle| S_{0m}, S_{1m}, \dots, S_{Km}, \sum_{n=1}^{K} S_{nm} = \kappa\right) = (1-q)H_2(p^{\kappa})$$

Averaging over the binomial distribution of κ yields the converse.
The capacity is achievable using random codes. □

Theorem [GZ '10] (Capacity of the deterministic model)

$$C = \frac{1-q}{K} \max_{p\in[0,1]} \sum_{\kappa=1}^{K} \binom{K}{\kappa} q^{\kappa}(1-q)^{K-\kappa} H_2(p^{\kappa})$$

Proof. $M \sum_{k=1}^{K} R_k \leq I(Y_{01},\dots,Y_{0M};\boldsymbol{X}_1,\dots,\boldsymbol{X}_M|\underline{\boldsymbol{S}})$

$$\leq \sum_{m=1}^{M} I(Y_{0m};X_{1m},\dots,X_{Km}|S_{0m},S_{1m},\dots,S_{Km})$$

$$= \sum_{m=1}^{M} H(Y_{0m}|S_{0m},S_{1m},\dots,S_{Km})$$

Bernoulli $(1-p)$ signaling is optimal by symmetry.
Conditioned on that κ out of the K neighbors transmit,

$$H\left(Y_{0m}\middle|S_{0m},S_{1m},\dots,S_{Km},\sum_{n=1}^{K} S_{nm}=\kappa\right) = (1-q)H_2(p^{\kappa})$$

Averaging over the binomial distribution of κ yields the converse.
The capacity is achievable using random codes. □

Theorem [GZ '10] (Capacity of the deterministic model)

$$C = \frac{1-q}{K} \max_{p\in[0,1]} \sum_{\kappa=1}^{K} \binom{K}{\kappa} q^{\kappa}(1-q)^{K-\kappa} H_2(p^{\kappa})$$

Proof. $M \sum_{k=1}^{K} R_k \le I(Y_{01},\dots,Y_{0M};\boldsymbol{X}_1,\dots,\boldsymbol{X}_M|\underline{\boldsymbol{S}})$

$$\le \sum_{m=1}^{M} I(Y_{0m};X_{1m},\dots,X_{Km}|S_{0m},S_{1m},\dots,S_{Km})$$

$$= \sum_{m=1}^{M} H(Y_{0m}|S_{0m},S_{1m},\dots,S_{Km})$$

Bernoulli $(1-p)$ signaling is optimal by symmetry.

Conditioned on that κ out of the K neighbors transmit,

$$H\left(Y_{0m}\middle|S_{0m},S_{1m},\dots,S_{Km},\sum_{n=1}^{K} S_{nm}=\kappa\right) = (1-q)H_2(p^{\kappa})$$

Averaging over the binomial distribution of κ yields the converse.

The capacity is achievable using random codes. □

Theorem [GZ '10] (Capacity of the deterministic model)

$$C = \frac{1-q}{K} \max_{p\in[0,1]} \sum_{\kappa=1}^{K} \binom{K}{\kappa} q^{\kappa}(1-q)^{K-\kappa} H_2(p^{\kappa})$$

Proof. $M \sum_{k=1}^{K} R_k \le I(Y_{01}, \dots, Y_{0M}; \boldsymbol{X}_1, \dots, \boldsymbol{X}_M | \underline{\boldsymbol{S}})$

$$\le \sum_{m=1}^{M} I(Y_{0m}; X_{1m}, \dots, X_{Km} | S_{0m}, S_{1m}, \dots, S_{Km})$$

$$= \sum_{m=1}^{M} H(Y_{0m} | S_{0m}, S_{1m}, \dots, S_{Km})$$

Bernoulli $(1-p)$ signaling is optimal by symmetry.
Conditioned on that κ out of the K neighbors transmit,

$$H\left(Y_{0m} \middle| S_{0m}, S_{1m}, \dots, S_{Km}, \sum_{n=1}^{K} S_{nm} = \kappa\right) = (1-q)H_2(p^{\kappa})$$

Averaging over the binomial distribution of κ yields the converse.
The capacity is achievable using random codes. □

Theorem [GZ '10] (Capacity of the Gaussian model)

$$C = \frac{1-q}{2K} \sum_{\kappa=1}^{K} \binom{K}{\kappa} q^{\kappa} (1-q)^{K-\kappa} \log\left(1 + \frac{\kappa\gamma}{q}\right)$$

(This can be generalized to non-symetric capacity.)

The throughput of ALOHA:

$$\frac{1}{2} q(1-q)^K \log\left(1 + \frac{\gamma}{q}\right) < C$$

Related work: [Minero, Franceschetti & Tse '09, "Random access: An information-theoretic perspective"]

Theorem [GZ '10] (Capacity of the Gaussian model)

$$C = \frac{1-q}{2K} \sum_{\kappa=1}^{K} \binom{K}{\kappa} q^{\kappa} (1-q)^{K-\kappa} \log\left(1 + \frac{\kappa\gamma}{q}\right)$$

(This can be generalized to non-symetric capacity.)

The throughput of ALOHA:

$$\frac{1}{2} q (1-q)^{K} \log\left(1 + \frac{\gamma}{q}\right) < C$$

Related work: [Minero, Franceschetti & Tse '09, "Random access: An information-theoretic perspective"]

Theorem [GZ '10] (Capacity of the Gaussian model)

$$C = \frac{1-q}{2K} \sum_{\kappa=1}^{K} \binom{K}{\kappa} q^{\kappa} (1-q)^{K-\kappa} \log\left(1 + \frac{\kappa\gamma}{q}\right)$$

(This can be generalized to non-symetric capacity.)

The throughput of ALOHA:

$$\frac{1}{2} q (1-q)^{K} \log\left(1 + \frac{\gamma}{q}\right) < C$$

Related work: [Minero, Franceschetti & Tse '09, "Random access: An information-theoretic perspective"]

The Gaussian Model: RODD vs. ALOHA

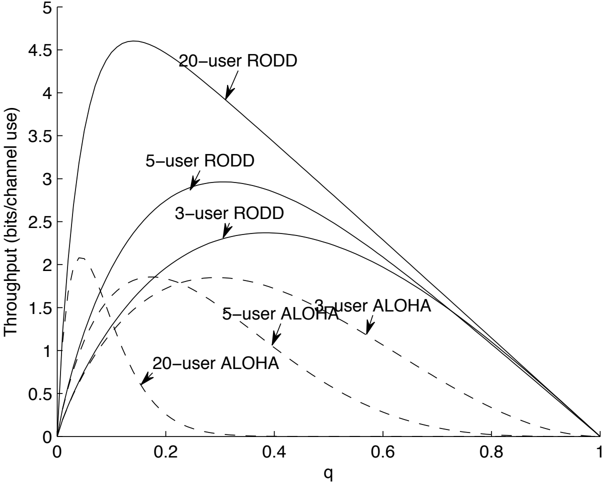

Result II: Full-Duplex Neighbor Discovery

Neighbor Discovery

- To discover and identify neighbors' network interface addr (NIAs)
- State of the art: random-access discovery
  Nodes announce NIA repeatedly with random delay
  - TND Protocol (IETF MANET Workgroup)
  - WiFi ad hoc mode
  - FlashLinQ (single-tone OFDM, CSMA-like)

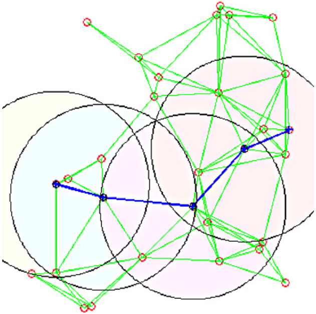

Neighbor Discovery

- To discover and identify neighbors' network interface addr (NIAs)
- State of the art: random-access discovery
  Nodes announce NIA repeatedly with random delay
  - TND Protocol (IETF MANET Workgroup)
  - WiFi ad hoc mode
  - FlashLinQ (single-tone OFDM, CSMA-like)

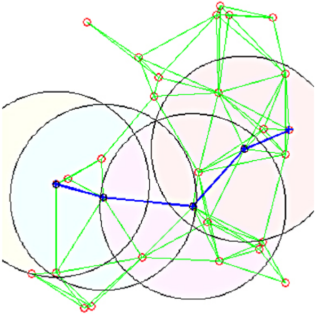

Neighbor Discovery

- To discover and identify neighbors' network interface addr (NIAs)
- State of the art: random-access discovery
  Nodes announce NIA repeatedly with random delay
  - TND Protocol (IETF MANET Workgroup)
  - WiFi ad hoc mode
  - FlashLinQ (single-tone OFDM, CSMA-like)

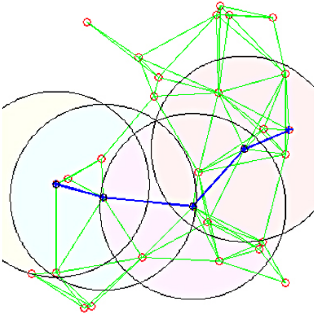

The Fundamental Problem

Linear measurements via MAC (with Rayleigh fading)

$$\boldsymbol{Y} = \sum_{n \in \partial k} \boldsymbol{s}_n U_n + \boldsymbol{W}$$
$$= \sum_{n=1}^{N} \boldsymbol{s}_n X_n + \boldsymbol{W}$$

where $X_n \approxeq 0$ for all but a few neighbors

Neighbor discovery is a problem of compressed sensing (aka sparse recovery) by nature.

[Donoho '06] [Candes & Tao '06]

The Fundamental Problem

Linear measurements via MAC (with Rayleigh fading)

$$
\begin{aligned}
\boldsymbol{Y} &= \sum_{n \in \partial k} \boldsymbol{s}_n U_n + \boldsymbol{W} \\
&= \sum_{n=1}^{N} \boldsymbol{s}_n X_n + \boldsymbol{W}
\end{aligned}
$$

where $X_n \approxeq 0$ for all but a few neighbors

Neighbor discovery is a problem of compressed sensing (aka sparse recovery) by nature.

[Donoho '06] [Candes & Tao '06]

The Fundamental Problem

Linear measurements via MAC (with Rayleigh fading)

$$\boldsymbol{Y} = \sum_{n \in \partial k} \boldsymbol{s}_n U_n + \boldsymbol{W}$$
$$= \sum_{n=1}^{N} \boldsymbol{s}_n X_n + \boldsymbol{W}$$

where $X_n \approxeq 0$ for all but a few neighbors

Neighbor discovery is a problem of compressed sensing (aka sparse recovery) by nature.

[Donoho '06] [Candes & Tao '06]

Compressed Neighbor Discovery

- A clean-slate design using RODD signatures
- Network-wide full-duplex discovery
- Discovery via group testing [Luo & Guo '08, '09]
  ([Dorfman '43] to test soldiers for syphilis in WWII)
  Analyzed in, e.g., [Berger-Levenshtein '02]
- Suppose noiseless

$$\hat{Y}_m = \bigvee_{n=1}^{N} (s_{nm} X_n)$$

- If no energy received in off-slot m, then nodes who would have transmitted energy during the slot are not neighbors

  ... unless in the rare event that symbols transmitted by all neighbors completely cancel (in Gaussian model)

Compressed Neighbor Discovery

- A clean-slate design using RODD signatures
- Network-wide full-duplex discovery
- Discovery via group testing [Luo & Guo '08, '09]
  ([Dorfman '43] to test soldiers for syphilis in WWII)
  Analyzed in, e.g., [Berger-Levenshtein '02]
- Suppose noiseless

$$\hat{Y}_m = \bigvee_{n=1}^{N} (s_{nm} X_n)$$

- If no energy received in off-slot m, then nodes who would have transmitted energy during the slot are not neighbors

  ... unless in the rare event that symbols transmitted by all neighbors completely cancel (in Gaussian model)

Compressed Neighbor Discovery

- A clean-slate design using RODD signatures
- Network-wide full-duplex discovery
- Discovery via group testing [Luo & Guo '08, '09]
  ([Dorfman '43] to test soldiers for syphilis in WWII)
  Analyzed in, e.g., [Berger-Levenshtein '02]
- Suppose noiseless

$$\hat{Y}_m = \bigvee_{n=1}^{N} (s_{nm} X_n)$$

- If no energy received in off-slot m, then nodes who would have transmitted energy during the slot are not neighbors

  … unless in the rare event that symbols transmitted by all neighbors completely cancel (in Gaussian model)

Compressed Neighbor Discovery

- A clean-slate design using RODD signatures
- Network-wide full-duplex discovery
- Discovery via group testing [Luo & Guo '08, '09]
  ([Dorfman '43] to test soldiers for syphilis in WWII)
  Analyzed in, e.g., [Berger-Levenshtein '02]
- Suppose noiseless

$$\hat{Y}_m = \bigvee_{n=1}^{N} (s_{nm} X_n)$$

- If no energy received in off-slot m, then nodes who would have transmitted energy during the slot are not neighbors

  ... unless in the rare event that symbols transmitted by all neighbors completely cancel (in Gaussian model)

Compressed Neighbor Discovery

- A clean-slate design using RODD signatures
- Network-wide full-duplex discovery
- Discovery via group testing [Luo & Guo '08, '09]
  ([Dorfman '43] to test soldiers for syphilis in WWII)
  Analyzed in, e.g., [Berger-Levenshtein '02]
- Suppose noiseless

$$\hat{Y}_m = \bigvee_{n=1}^{N} (s_{nm} X_n)$$

- If no energy received in off-slot m, then nodes who would have transmitted energy during the slot are not neighbors

  … unless in the rare event that symbols transmitted by all neighbors completely cancel (in Gaussian model)

Compressed Neighbor Discovery

- A clean-slate design using RODD signatures
- Network-wide full-duplex discovery
- Discovery via group testing [Luo & Guo '08, '09]
  ([Dorfman '43] to test soldiers for syphilis in WWII)
  Analyzed in, e.g., [Berger-Levenshtein '02]
- Suppose noiseless

$$\hat{Y}_m = \bigvee_{n=1}^{N} (s_{nm} X_n)$$

- If no energy received in off-slot m, then nodes who would have transmitted energy during the slot are not neighbors

  ... unless in the rare event that symbols transmitted by all neighbors completely cancel (in Gaussian model)

Theorem [LG '08] (Noiseless Group Testing)

N nodes

each has c neighbors on average

Let the sparsity be

$$q = (2 \log N \log \log N)^{-1}$$

and the number of measurements be

$$M = 4(\log N)^2 \log \log N$$

then the average number of false alarms is

$$\mathcal{E} < \frac{1}{N} \left(1 + \frac{3c}{\log \log N} \right)$$

There is no miss.

Theorem [LG '08] (Noiseless Group Testing)

N nodes

each has c neighbors on average

Let the sparsity be

$$q = (2 \log N \log \log N)^{-1}$$

and the number of measurements be

$$M = 4(\log N)^2 \log \log N$$

then the average number of false alarms is

$$\mathcal{E} < \frac{1}{N}\left(1 + \frac{3c}{\log \log N}\right)$$

There is no miss.

Theorem [LG '08] (Noiseless Group Testing)

N nodes

each has c neighbors on average

Let the sparsity be

$$q = (2 \log N \log \log N)^{-1}$$

and the number of measurements be

$$M = 4(\log N)^2 \log \log N$$

then the average number of false alarms is

$$\mathcal{E} < \frac{1}{N}\left(1 + \frac{3c}{\log \log N}\right)$$

There is no miss.

Performance with Noisy Channel

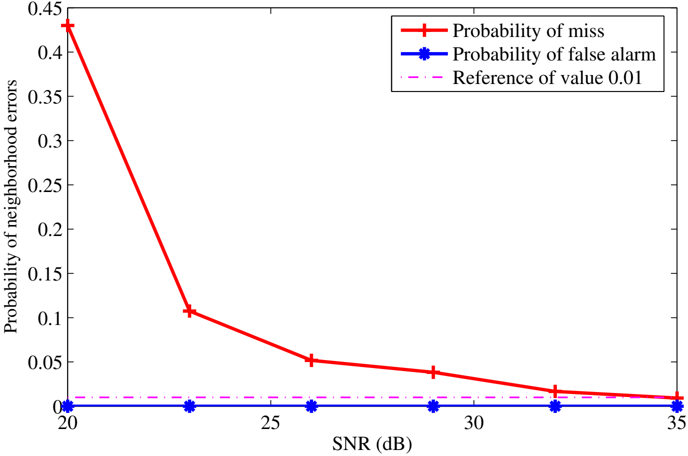

$N = $10,000, $c = $5, $M = $1,500.

Improvements

- k-strike group testing: eliminate a node only if implicated by k or more measurements
- Multiple nodes are annihilated if their pulses cancel. Randomizing the phase over slots reduces the chance of cancellation.

Improvements

- k-strike group testing: eliminate a node only if implicated by k or more measurements
- Multiple nodes are annihilated if their pulses cancel. Randomizing the phase over slots reduces the chance of cancellation.

Error Rate vs. SNR

10,000 nodes, 50 neighbors on average, 2,500-symbol signatures
$q = 0.013$, path loss exponent = 3

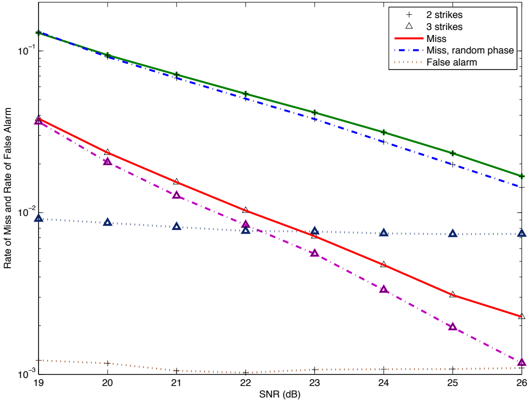

Comparison with Generic Random Access

- Compressed discovery takes 1 frame; random-access takes many.
  $\Rightarrow$ Significant reduction of per-frame overhead
- Experiment: N =10,000, c =10, M =1,000, SNR =23 dB.

  Pe of compressed discovery is 0.002.

  If nodes contend to announce their NIAs over t periods,

$$Pe = \sum_{z=1}^{N} \binom{N}{z} \left(\frac{c}{N}\right)^z \left(1 - \frac{c}{N}\right)^{N-z} \left[1 - \theta \left(1 - \theta\right)^{z-1}\right]^t$$

  It takes 194 contention periods to achieve 0.002.

  NIA is $\log_2(10^4) \approx 14$ bits. With binary signalling, the total overhead is at least 194×14=2,716 bits.

Comparison with Generic Random Access

- Compressed discovery takes 1 frame; random-access takes many.
  $\Rightarrow$ Significant reduction of per-frame overhead
- Experiment: N =10,000, c =10, M =1,000, SNR =23 dB.
  Pe of compressed discovery is 0.002.
  If nodes contend to announce their NIAs over t periods,

$$Pe = \sum_{z=1}^{N} \binom{N}{z} \left(\frac{c}{N}\right)^z \left(1 - \frac{c}{N}\right)^{N-z} \left[1 - \theta (1-\theta)^{z-1}\right]^t$$

  It takes 194 contention periods to achieve 0.002.
  NIA is $\log_2(10^4) \approx 14$ bits. With binary signalling, the total overhead is at least 194×14=2,716 bits.

Comparison with Generic Random Access

- Compressed discovery takes 1 frame; random-access takes many.
  $\Rightarrow$ Significant reduction of per-frame overhead
- Experiment: N =10,000, c =10, M =1,000, SNR =23 dB.

  Pe of compressed discovery is 0.002.

  If nodes contend to announce their NIAs over t periods,

$$Pe = \sum_{z=1}^{N} \binom{N}{z} \left(\frac{c}{N}\right)^z \left(1 - \frac{c}{N}\right)^{N-z} \left[1 - \theta \left(1 - \theta\right)^{z-1}\right]^t$$

  It takes 194 contention periods to achieve 0.002.

  NIA is $\log_2(10^4) \approx 14$ bits. With binary signalling, the total overhead is at least 194×14=2,716 bits.

Comparison with Generic Random Access

- Compressed discovery takes 1 frame; random-access takes many.
  $\Rightarrow$ Significant reduction of per-frame overhead
- Experiment: N =10,000, c =10, M =1,000, SNR =23 dB.

  Pe of compressed discovery is 0.002.

  If nodes contend to announce their NIAs over t periods,

$$Pe = \sum_{z=1}^{N} \binom{N}{z} \left(\frac{c}{N}\right)^z \left(1 - \frac{c}{N}\right)^{N-z} \left[1 - \theta \left(1 - \theta\right)^{z-1}\right]^t$$

  It takes 194 contention periods to achieve 0.002.

  NIA is $\log_2(10^4) \approx 14$ bits. With binary signalling, the total overhead is at least 194×14=2,716 bits.

Comparison with Generic Random Access

- Compressed discovery takes 1 frame; random-access takes many.
  $\Rightarrow$ Significant reduction of per-frame overhead
- Experiment: N =10,000, c =10, M =1,000, SNR =23 dB.

  Pe of compressed discovery is 0.002.

  If nodes contend to announce their NIAs over t periods,

$$Pe = \sum_{z=1}^{N} \binom{N}{z} \left(\frac{c}{N}\right)^z \left(1 - \frac{c}{N}\right)^{N-z} \left[1 - \theta\,(1-\theta)^{z-1}\right]^t$$

  It takes 194 contention periods to achieve 0.002.

  NIA is $\log_2(10^4) \approx 14$ bits. With binary signalling, the total overhead is at least 194$\times$14=2,716 bits.

Deterministic Signatures

- Random signatures require decoding complexity $O(MN)$
- [Calderbank, Howard & Jafapour '10] suggest measurement matrix using second-order Reed-Muller (RM) codes (digital QPSK chirps)
- [Guo, Luo, Zhang & Shen '10] apply RM codes to neighbor discovery
- A chirp reconstruction algo [Howard, Calderbank & Searle '08]
- Up to $2^{m(m+3)/2}$ distinct signatures, each of length 2^m. E.g., 2^{65} signatures of length $2^{10} = 1,024$.

Deterministic Signatures

- Random signatures require decoding complexity $O(MN)$
- [Calderbank, Howard & Jafapour '10] suggest measurement matrix using second-order Reed-Muller (RM) codes (digital QPSK chirps)
- [Guo, Luo, Zhang & Shen '10] apply RM codes to neighbor discovery
- A chirp reconstruction algo [Howard, Calderbank & Searle '08]
- Up to $2^{m(m+3)/2}$ distinct signatures, each of length 2^m. E.g., 2^{65} signatures of length $2^{10} = 1,024$.

Deterministic Signatures

- Random signatures require decoding complexity $O(MN)$
- [Calderbank, Howard & Jafapour '10] suggest measurement matrix using second-order Reed-Muller (RM) codes (digital QPSK chirps)
- [Guo, Luo, Zhang & Shen '10] apply RM codes to neighbor discovery
- A chirp reconstruction algo [Howard, Calderbank & Searle '08]
- Up to $2^{m(m+3)/2}$ distinct signatures, each of length 2^m. E.g., 2^{65} signatures of length $2^{10} = 1,024$.

Deterministic Signatures

- Random signatures require decoding complexity $O(MN)$
- [Calderbank, Howard & Jafapour '10] suggest measurement matrix using second-order Reed-Muller (RM) codes (digital QPSK chirps)
- [Guo, Luo, Zhang & Shen '10] apply RM codes to neighbor discovery
- A chirp reconstruction algo [Howard, Calderbank & Searle '08]
- Up to $2^{m(m+3)/2}$ distinct signatures, each of length 2^m. E.g., 2^{65} signatures of length $2^{10} = 1,024$.

Deterministic Signatures

- Random signatures require decoding complexity $O(MN)$
- [Calderbank, Howard & Jafapour '10] suggest measurement matrix using second-order Reed-Muller (RM) codes (digital QPSK chirps)
- [Guo, Luo, Zhang & Shen '10] apply RM codes to neighbor discovery
- A chirp reconstruction algo [Howard, Calderbank & Searle '08]
- Up to $2^{m(m+3)/2}$ distinct signatures, each of length 2^m.
  E.g., 2^{65} signatures of length $2^{10} = 1,024$.

Error Rate vs. SNR

2^{20} nodes, 1,024-symbol RM signatures, 50 neighbors on average
path loss exponent $= 3$

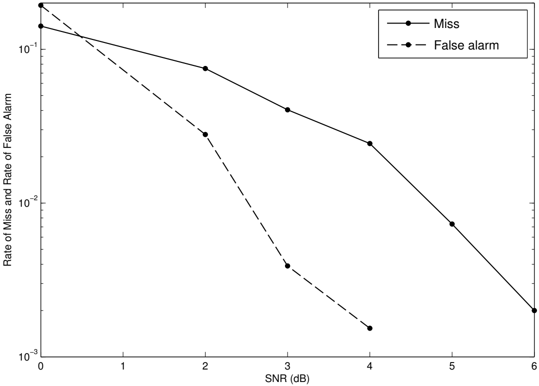

Drawback: RM code is not RODD, so not full-duplex.

Signature Distribution

- $\boldsymbol{s}_k = f(\mathsf{NIA}_k)$
- Enough to distribute a pseudo-random number generator
  or
  specify an RM codebook
- Neighbor discovery and data transmission share the same RODD frame structure — possible to discover neighbors solely based on data frames

Signature Distribution

- $\boldsymbol{s}_k = f(\mathsf{NIA}_k)$
- Enough to distribute a pseudo-random number generator
  or
  specify an RM codebook
- Neighbor discovery and data transmission share the same RODD frame structure — possible to discover neighbors solely based on data frames

Signature Distribution

- $\boldsymbol{s}_k = f(\mathsf{NIA}_k)$
- Enough to distribute a pseudo-random number generator
  or
  specify an RM codebook
- Neighbor discovery and data transmission share the same RODD frame structure — possible to discover neighbors solely based on data frames

Result III: Full-Duplex Mutual Broadcast

Mutual Broadcast

- Simultaneous broadcast from nodes to their one-hop neighbors
- Similar to compressed neighbor discovery, except that each node is assigned a collection of signatures
- Conventional schemes: ALOHA, CSMA
- Applications:
  - network state exchange
  - spontaneous social networks
  - real-time video sharing

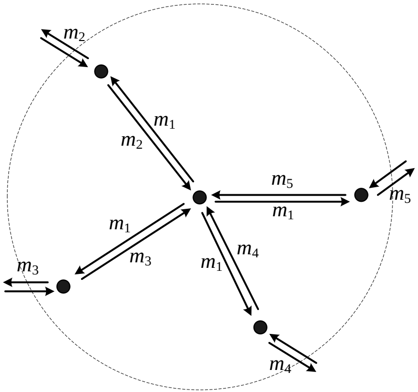

Mutual Broadcast

- Simultaneous broadcast from nodes to their one-hop neighbors
- Similar to compressed neighbor discovery, except that each node is assigned a collection of signatures
- Conventional schemes: ALOHA, CSMA
- Applications:
  - network state exchange
  - spontaneous social networks
  - real-time video sharing

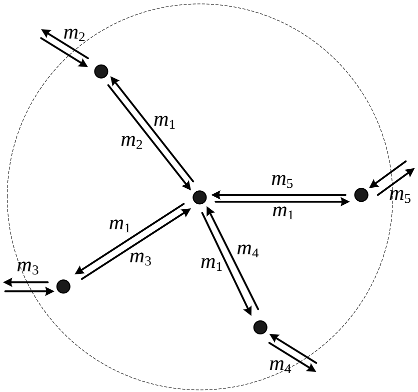

Mutual Broadcast

- Simultaneous broadcast from nodes to their one-hop neighbors
- Similar to compressed neighbor discovery, except that each node is assigned a collection of signatures
- Conventional schemes: ALOHA, CSMA
- Applications:
  - network state exchange
  - spontaneous social networks
  - real-time video sharing

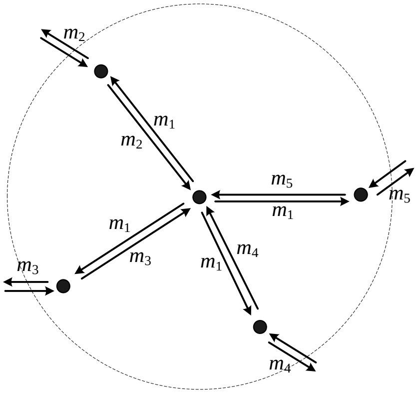

Mutual Broadcast

- Simultaneous broadcast from nodes to their one-hop neighbors
- Similar to compressed neighbor discovery, except that each node is assigned a collection of signatures
- Conventional schemes: ALOHA, CSMA
- Applications:
  - network state exchange
  - spontaneous social networks
  - real-time video sharing

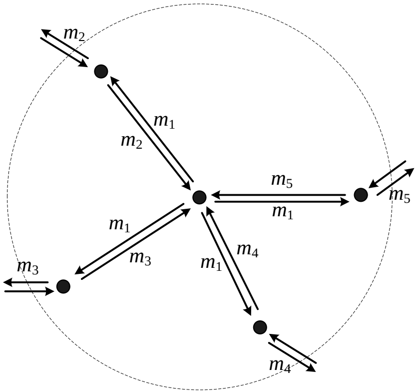

Novel RODD-Based Scheme

- Node k is assigned 2^l on-off signatures $\boldsymbol{S}_k(1), \ldots, \boldsymbol{S}_k(2^l)$
- Node k observes

$$\boldsymbol{Y}_k = \sum_{j \in \partial k} \sqrt{\gamma_j} \boldsymbol{S}_j(w_j) + \boldsymbol{W}_k$$
$$= \boldsymbol{S}\boldsymbol{X} + \boldsymbol{W}_k$$

- To identify, out of a total of $2^l|\partial k|$ signatures from all neighbors, which $|\partial k|$ signatures were selected.
- The sparsity of $\boldsymbol{X}$ is 2^{-l}

Novel RODD-Based Scheme

- Node k is assigned 2^l on-off signatures $\boldsymbol{S}_k(1), \dots, \boldsymbol{S}_k(2^l)$
- Node k observes

$$\begin{aligned}\boldsymbol{Y}_k &= \sum_{j \in \partial k} \sqrt{\gamma_j} \boldsymbol{S}_j(w_j) + \boldsymbol{W}_k \\ &= \boldsymbol{S}\boldsymbol{X} + \boldsymbol{W}_k\end{aligned}$$

- To identify, out of a total of $2^l|\partial k|$ signatures from all neighbors, which $|\partial k|$ signatures were selected.
- The sparsity of $\boldsymbol{X}$ is 2^{-l}

Novel RODD-Based Scheme

- Node k is assigned 2^l on-off signatures $\boldsymbol{S}_k(1), \dots, \boldsymbol{S}_k(2^l)$
- Node k observes

$$\begin{aligned}\boldsymbol{Y}_k &= \sum_{j \in \partial k} \sqrt{\gamma_j} \boldsymbol{S}_j(w_j) + \boldsymbol{W}_k \\ &= \boldsymbol{S}\boldsymbol{X} + \boldsymbol{W}_k\end{aligned}$$

- To identify, out of a total of $2^l |\partial k|$ signatures from all neighbors, which $|\partial k|$ signatures were selected.
- The sparsity of $\boldsymbol{X}$ is 2^{-l}

Novel RODD-Based Scheme

- Node k is assigned 2^l on-off signatures $\boldsymbol{S}_k(1), \ldots, \boldsymbol{S}_k(2^l)$
- Node k observes

$$\begin{aligned}\boldsymbol{Y}_k &= \sum_{j \in \partial k} \sqrt{\gamma_j} \boldsymbol{S}_j(w_j) + \boldsymbol{W}_k \\ &= \boldsymbol{S}\boldsymbol{X} + \boldsymbol{W}_k\end{aligned}$$

- To identify, out of a total of $2^l|\partial k|$ signatures from all neighbors, which $|\partial k|$ signatures were selected.
- The sparsity of $\boldsymbol{X}$ is 2^{-l}

Sparse Recovery Using a Message Passing Algorithm

- Factor Graph

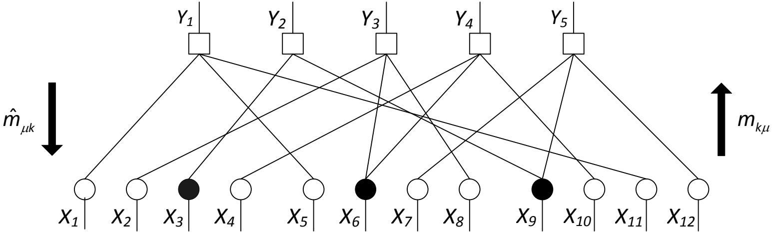

- Computational complexity: $\mathcal{O}(MNq)$
- Alternatives
  - Compressive sampling matching pursuit (CoSaMP)
  - Approximate message passing (AMP)

Sparse Recovery Using a Message Passing Algorithm

- Factor Graph

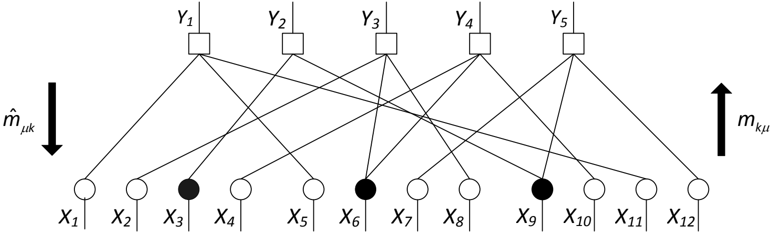

- Computational complexity: $\mathcal{O}(MNq)$
- Alternatives
  - Compressive sampling matching pursuit (CoSaMP)
  - Approximate message passing (AMP)

Sparse Recovery Using a Message Passing Algorithm

- Factor Graph

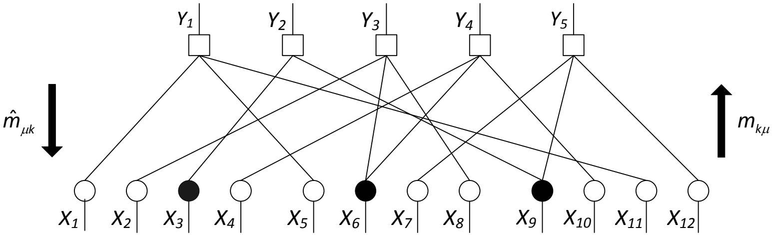

- Computational complexity: $\mathcal{O}(MNq)$
- Alternatives
  - Compressive sampling matching pursuit (CoSaMP)
  - Approximate message passing (AMP)

RODD vs. ALOHA

9 neighbors on average, 10 bits each, SNR=10 dB

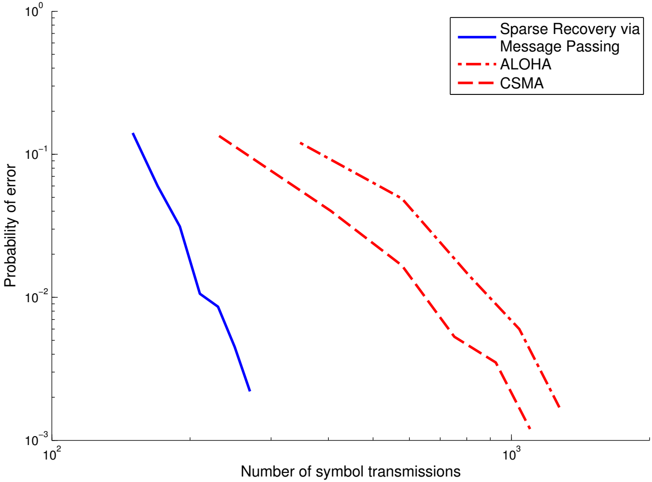

Message Passing vs. CoSaMP & AMP

10 nodes, 5 bits each

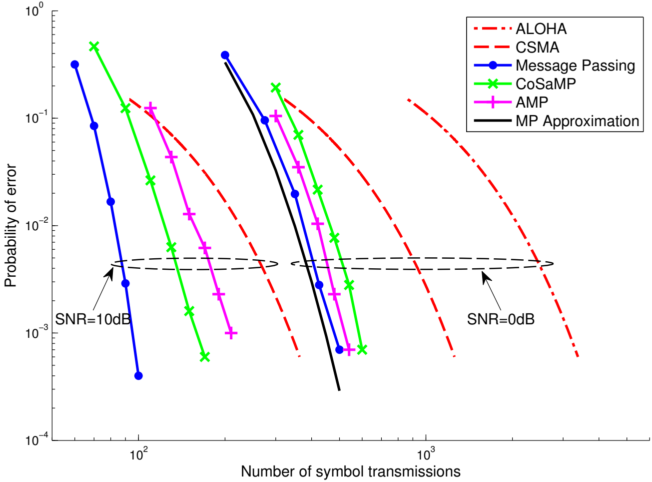

Research Questions

1. Synchronicity

- How to take advantage of on-off signaling for synchronization?
- Fundamental trade-off between cost and benefit of synchronicity in ad hoc networks?

1. Synchronicity

- How to take advantage of on-off signaling for synchronization?
- Fundamental trade-off between cost and benefit of synchronicity in ad hoc networks?

2. Virtual Full-Duplex Relay

- Minimal queueing delay in contrast to store-and-forward
- Capacity? How does it compare with frame-level AF, CF, DF?
- Virtual full-duplex two-way relay?

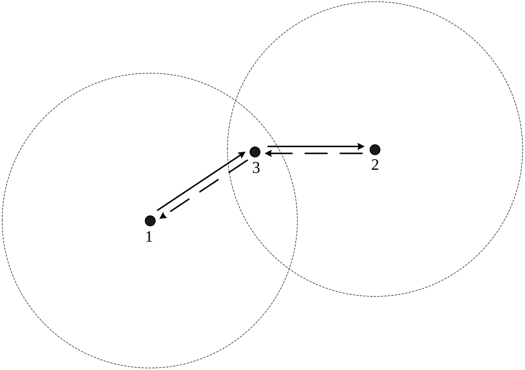

2. Virtual Full-Duplex Relay

- Minimal queueing delay in contrast to store-and-forward
- Capacity?
  How does it compare with frame-level AF, CF, DF?
- Virtual full-duplex two-way relay?

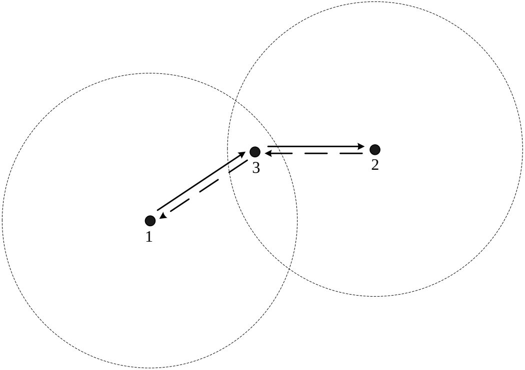

2. Virtual Full-Duplex Relay

- Minimal queueing delay in contrast to store-and-forward
- Capacity?
  How does it compare with frame-level AF, CF, DF?
- Virtual full-duplex two-way relay?

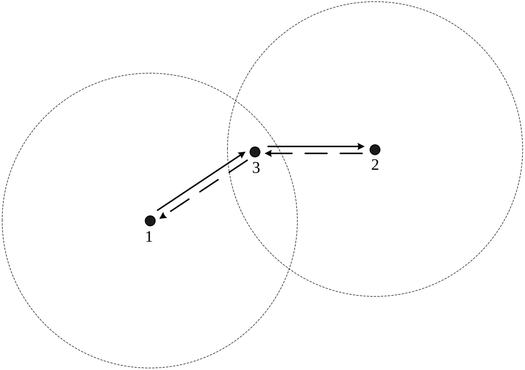

3. Network Coding

What is the impact of RODD signaling on network coding?