Ad-hoc Affectance-selective Families for Layer Dissemination

Harshita Kudaravalli and Miguel A. Mosteiro
Pace University
mmosteiro@pace.edu

SEA 2017

Application: the Internet of Things

Ad-hoc Wireless Networks

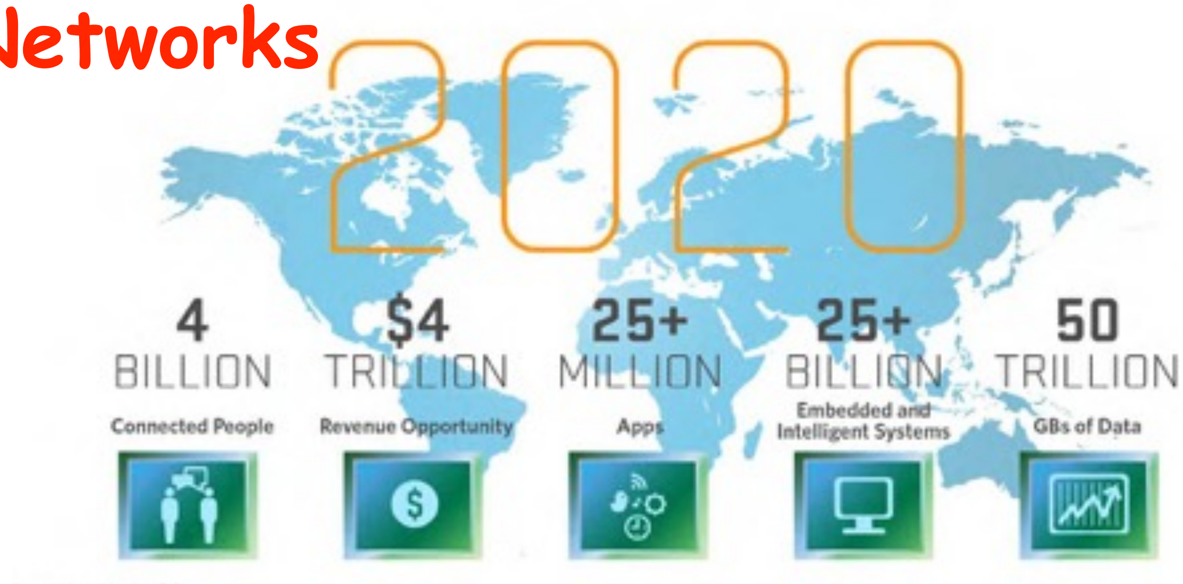

A Sensor Network

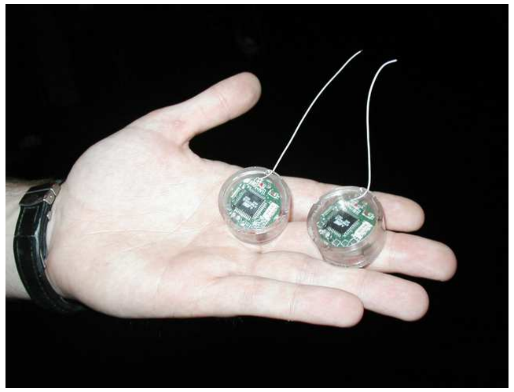

Intel Berkeley Research Lab

Capabilities

- processing
- sensing
- communication

Limitations

- range
- memory
- life cycle

Ad-hoc Wireless Networks example

A Sensor Network

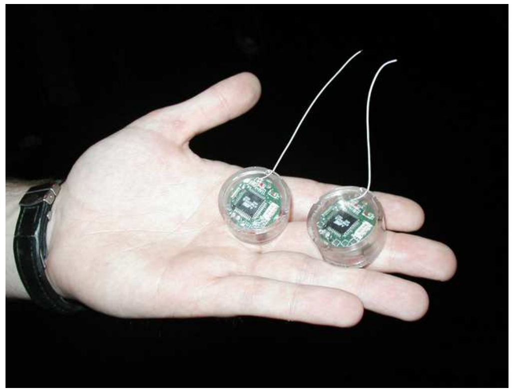

Intel Berkeley Research Lab

Capabilities

- processing
- sensing
- communication

Limitations

- range
- memory
- life cycle

Ad-hoc Wireless Networks example

A Sensor Network

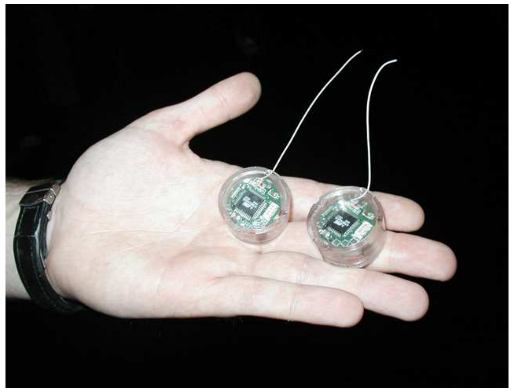

Intel Berkeley Research Lab

Capabilities

- processing
- sensing
- communication

Limitations

- range
- memory
- life cycle

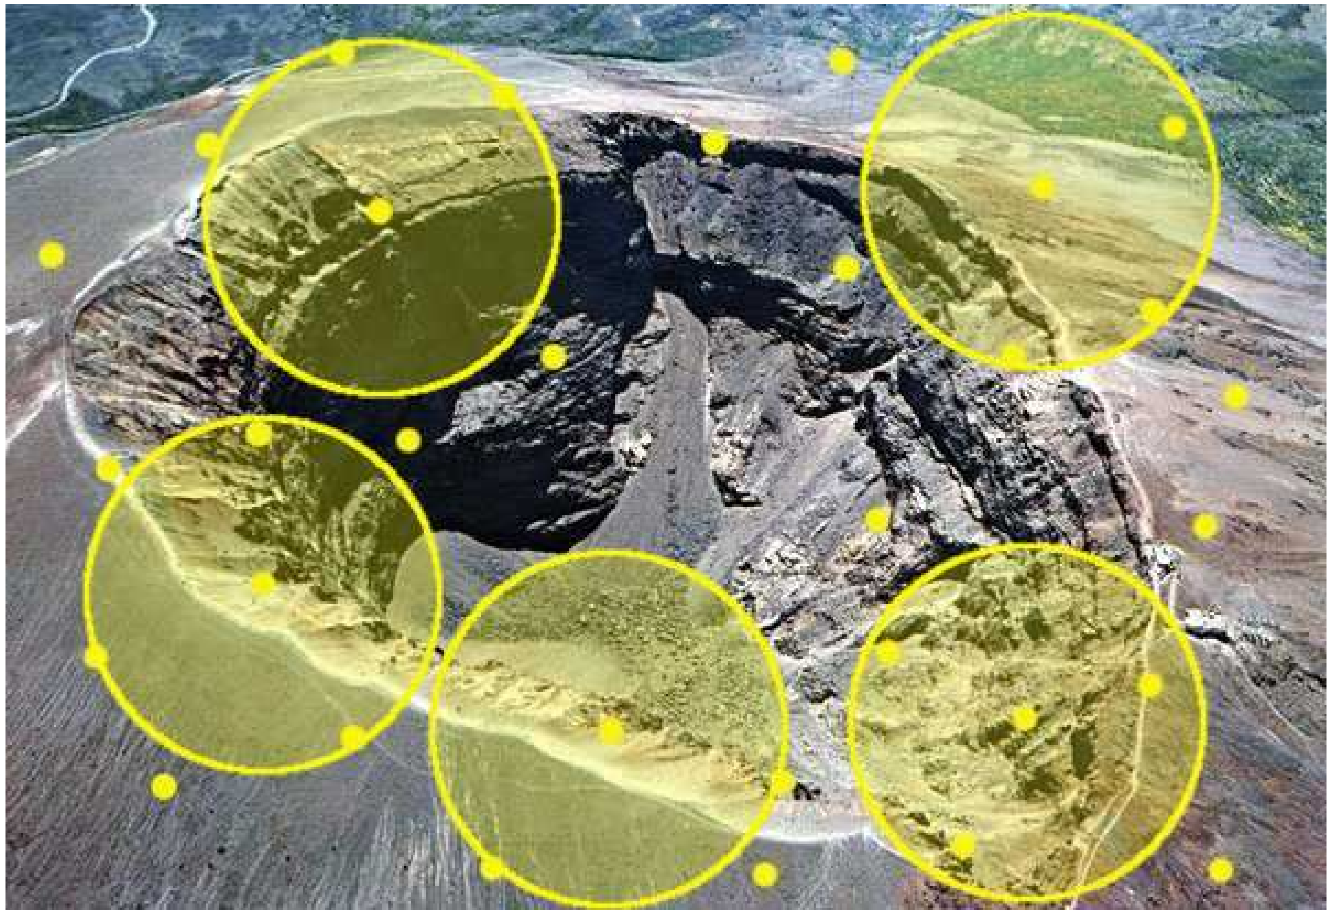

Ad-hoc Wireless Networks example

A Sensor Network

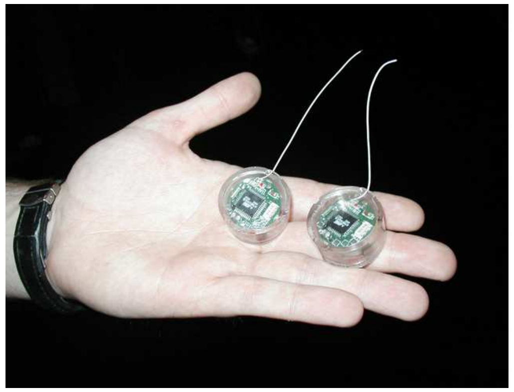

Intel Berkeley Research Lab

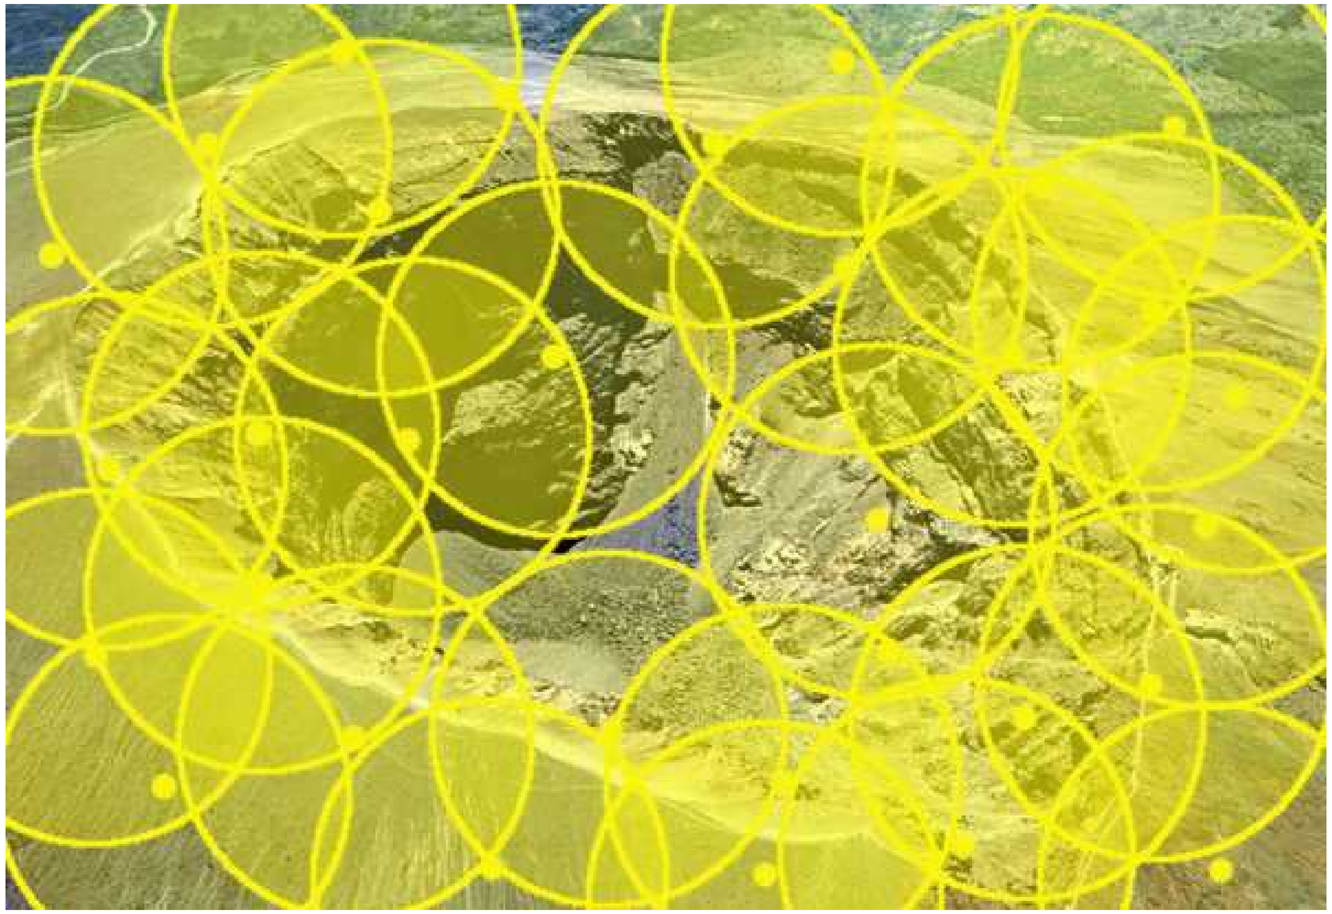

Capabilities

- processing
- sensing
- communication

Limitations

- range
- memory
- life cycle

Dissemination Problems in Wireless Networks

Radio Network = abstraction of a radio communication network

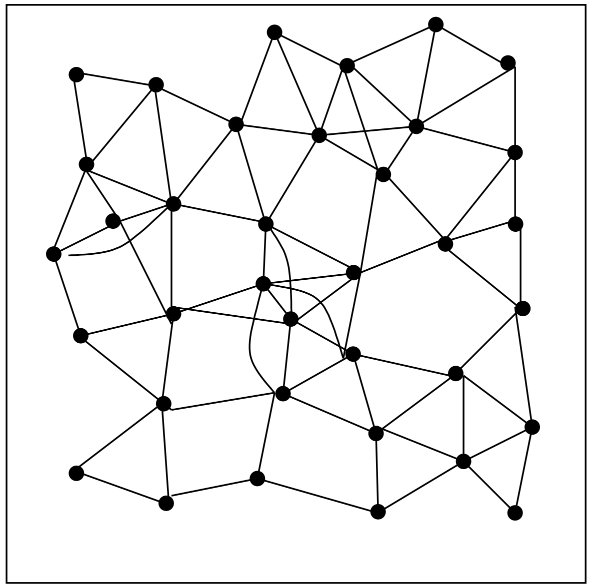

A geometric graph.

k nodes
hold a piece of information to diseminate.

- $k = 1 \rightarrow$ *Broadcast* [BGI'92,KM'98]
- $k = n \rightarrow$ *Gossiping* [CGLP'01,LP'02]
- k arbitrary $\rightarrow$ *k-selection* [K'05]

...

- *Multiple-message broadcast*
- *Dynamic multiple-message broadcast*
- *etc.*

Models for Wireless Networks

- **Topology Models :**
  - Undirected Graph
  - Unit Disk Graph
  - Time-varying Graph
- **Node Capabilities Models :**
  - Computational Resources
  - Communication Capabilities
  - Weak Sensor Model
- **Interference Models :**
  - Radio Network (RN)
  - Signal to Interference plus Noise Ratio (SINR)
  - Affectance (AFF)

Interference Models

RN Model [1]:

- **Collision/success:**

  Node y receives if and only if exactly one neighbor of y transmits at a given time, and y is not transmitting.

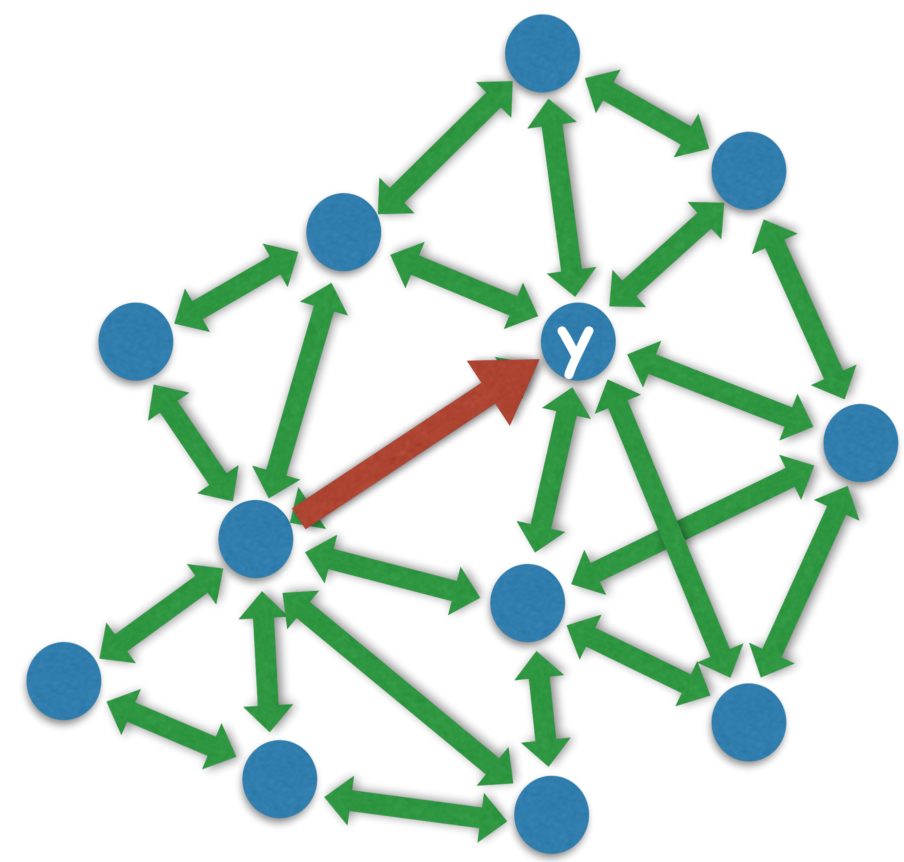

[1] Chlamtac and Kutten. *Trans. on Computers*. IEEE, 1987.

Interference Models

SINR Model [1]:

- **Collision/success:**

  A signal that overcomes interference from others plus background noise is received.

$$\frac{p((x,y))}{d(x,y)^{\alpha}} \geq \beta \left(\sum_{(u,v)\in\mathcal{R}(t)\setminus(x,y)} \frac{p((u,v))}{d(u,y)^{\alpha}} + N \right)$$

$$\begin{aligned} &\text{Defs.}: \\ \alpha > 0 &: \textit{path-loss exponent.} \\ \beta \geq 0 &: \textit{gain.} \\ p((i,j)) &: \textit{transmission power on link } (i,j). \\ d(i,j) &: \textit{Euclidean distance between } i \textit{ and } j. \\ \mathcal{R}(t) &: \textit{set of links transmitting at time } t. \\ N &: \textit{background noise.} \end{aligned}$$

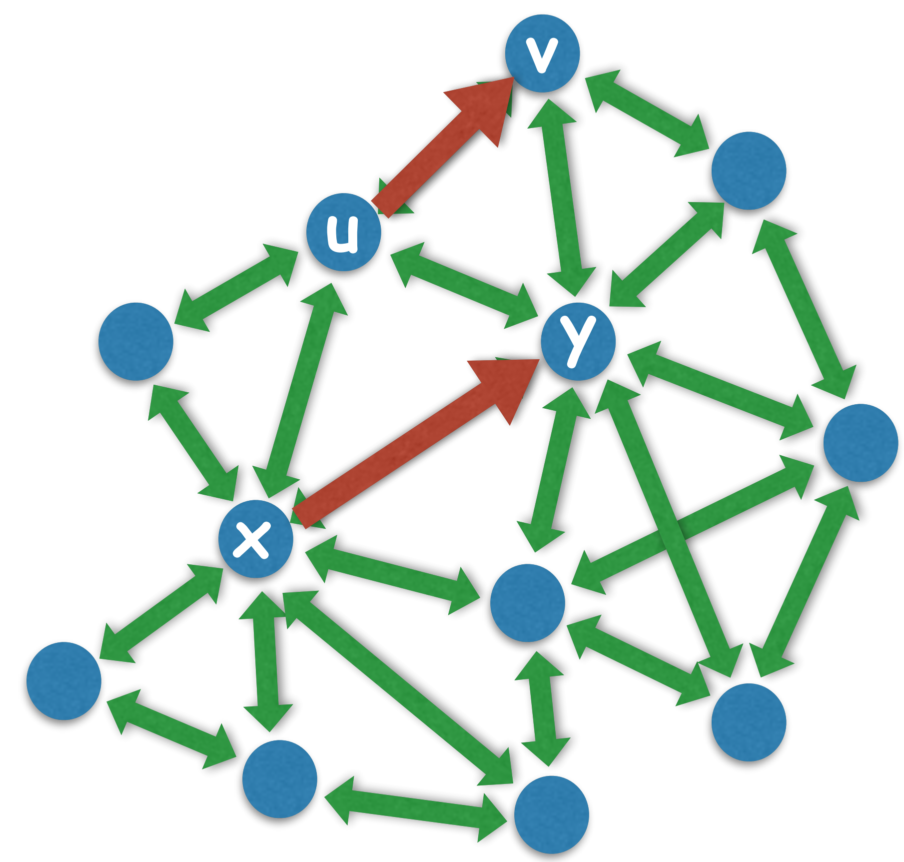

[1] Moscibroda and Wattenhofer. Infocom 2006.

Interference Models

AFF Model [1,2,3]:

$$a(u,(x,y))$$

matrix quantifying interference from node u on communication through link (x,y).

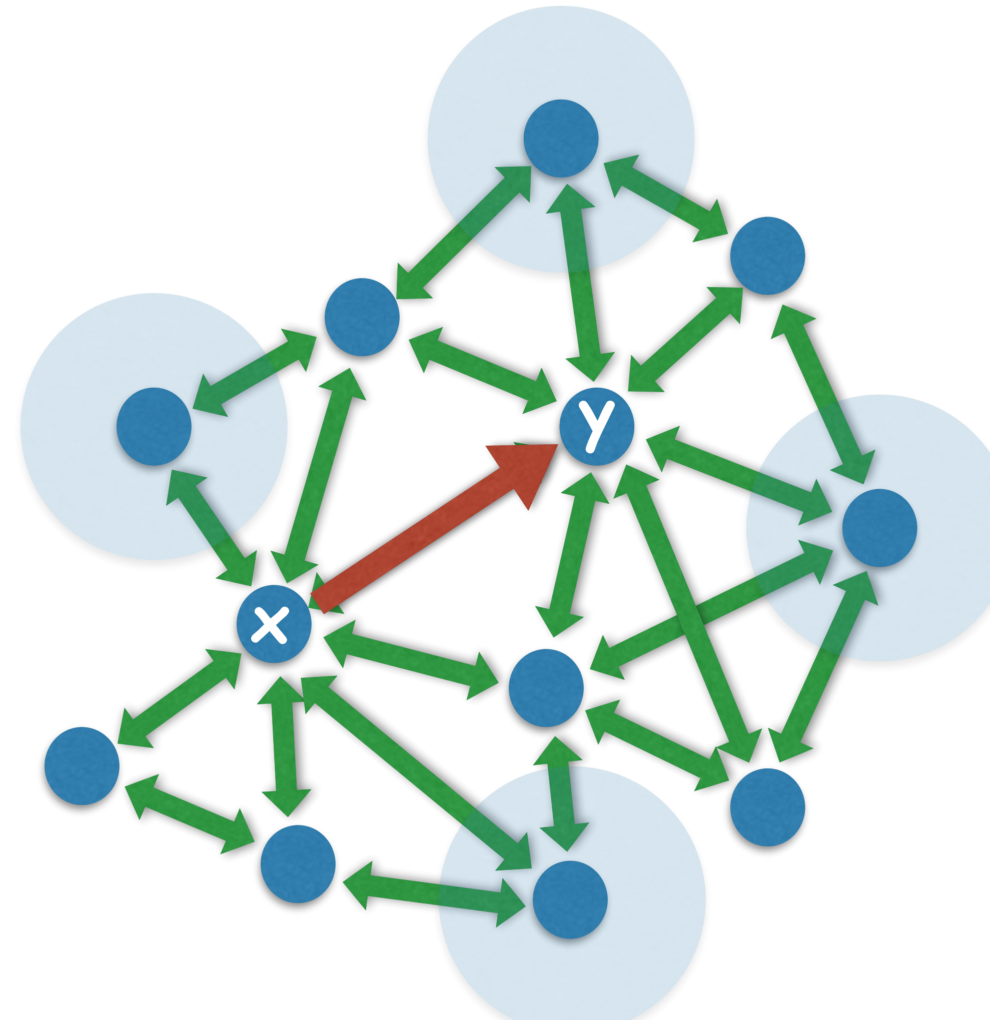

- **Collision/success:**

  For any link (x,y), a transmission from x is received by y in t if and only if

  » **x transmits in t and**

  » $\sum_{u \subseteq V(t)} a(u,(x,y)) < 1$,

  $V(t) \subseteq V$: set of nodes transmitting in t.

[1] Halldórsson and Wattenhofer. ICALP 2009.
[3] Fanghänel, Kesselheim and Vöcking. ICALP 2009.
[3] Kesselheim and Vöcking. DISC 2010.

Dynamic Multi-Broadcast [ACM-FOMC 2014]

Introduction

Dynamic Multiple-Message Broadcast (MMB) [1]:

- *problem:*
  packets arrive at some nodes **continuously**, to be delivered to **all** nodes
- *metric:*
  competitive throughput of deterministic distributed MMB algorithms
- *analysis:*
  in the **Affectance model**:
  - Affectance subsumes many interference models, e.g. RN and SINR models
  - conceptual idea: parameterize interference from transmitting **nodes into links**
  - introduced [2,3,4] for link scheduling as link-to-link affectance

[1] (non-dynamic MMB) Khabbazian-Kowalski PODC 2011
[2] Halldórsson-Wattenhofer, ICALP 2009
[3] Kesselheim, PODC 2012
[4] Kesselheim-Vöcking, DISC 2010

Dynamic Multi-Broadcast [ACM-FOMC 2014]

Affectance Characterization

Maximum average tree-layer affectance

Quantifies the difficulty to disseminate from one layer to the next one.

$$K(T,s) = \max_{d} \max_{V' \subseteq V_d(T)} \frac{a_{V'}(L(V'))}{|L(V')|}$$

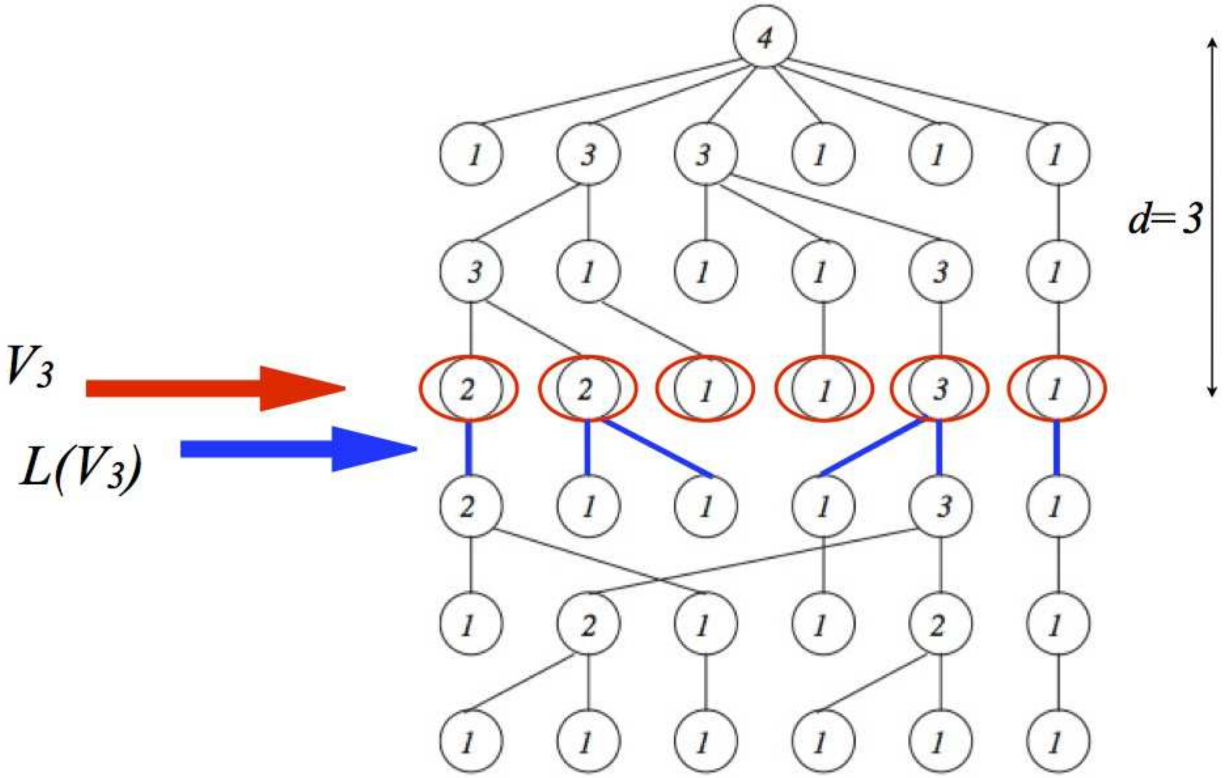

Dynamic Multi-Broadcast [ACM-FOMC 2014]

Affectance Characterization

Maximum fast-paths affectance

Quantifies the difficulty for dissemination on a path due to other paths.

$$M(T,s) = \max_{d,r} \max_{\ell \in F_d^r(T)} a_{F_d^r(T)}(\ell)$$

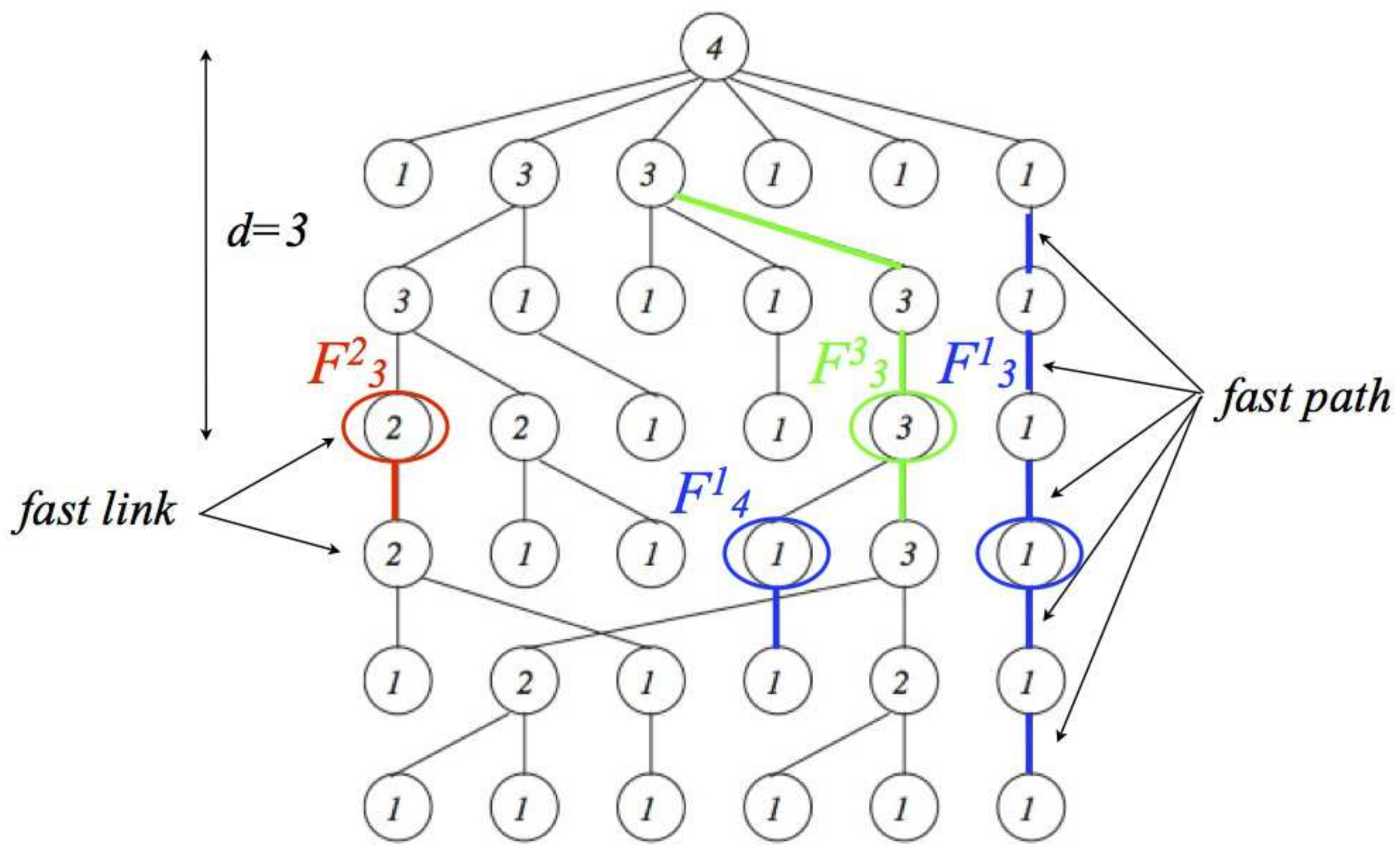

Dynamic Multi-Broadcast [ACM-FOMC 2014]

Introduction

Contributions:

- introduce new model characteristics:
  (based on comm network, affectance function, and a chosen BFS tree)
  – maximum average tree-layer affectance K
  – maximum fast-paths affectance M
- show how these characteristics influence broadcast time complexity:
  if one uses a specific BFS tree (GBST [1]) that minimizes $M(K+M)$
  single broadcast can be done in time $D + O(M(K+M)\log^3 n)$
- extend this to dynamic packet arrival model and the MMB problem:
  new MMB algorithm reaching throughput of $\Omega(1/(\alpha K \log n))$
- ... also simulations for RN

Dissemination bottleneck is from layer to layer!!

[1] Gaşieniec-Peleg-Xin, DC 2007

Dynamic Multi-Broadcast [ACM-FOMC 2014]

Affectance Characterization

Maximum average tree-layer affectance

Quantifies the difficulty to disseminate from one layer to the next one.

$$K(T,s) = \max_{d} \max_{V' \subseteq V_d(T)} \frac{a_{V'}(L(V'))}{|L(V')|}$$

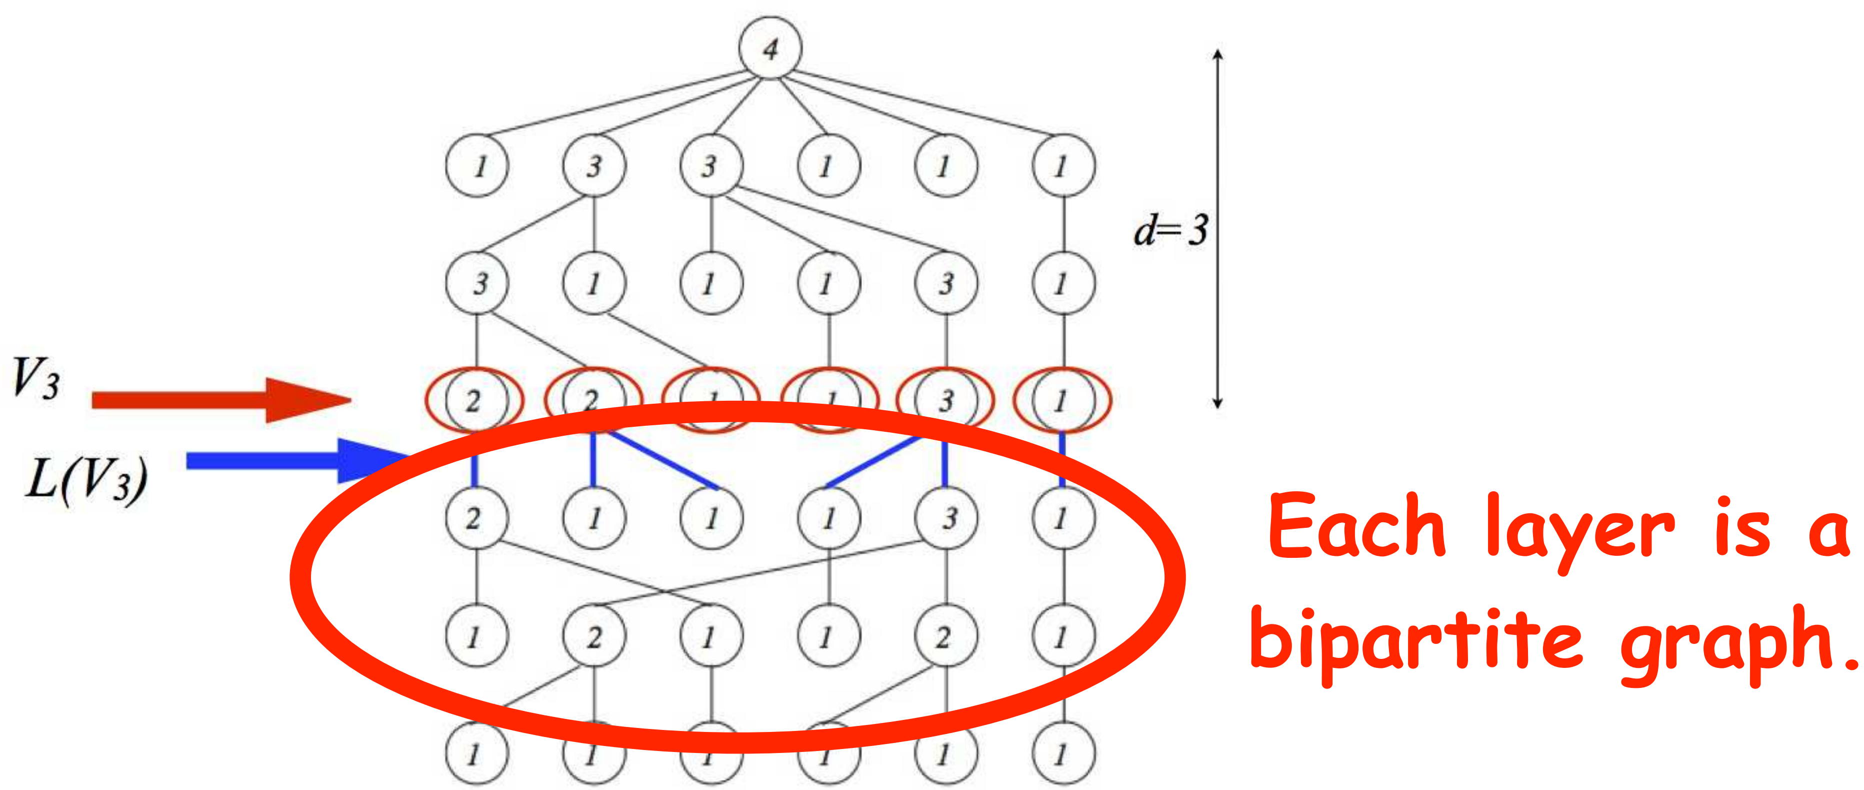

Layer Dissemination [SEA 2017]

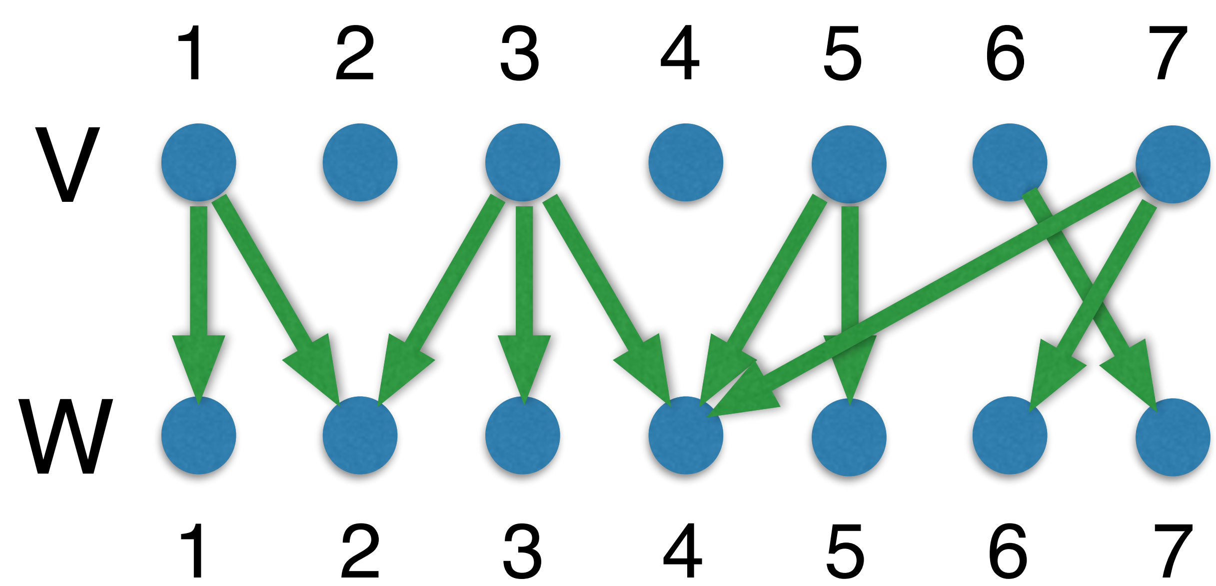

Bipartite network with

- V : set of transmitters
- W : set of receivers
- F_w : set of transmitters connected to $w \in W$

Layer Dissemination problem:

- Each $w \in W$ must receive at least one successful transmission from some $v \in F_w$, despite interference.

Layer Dissemination

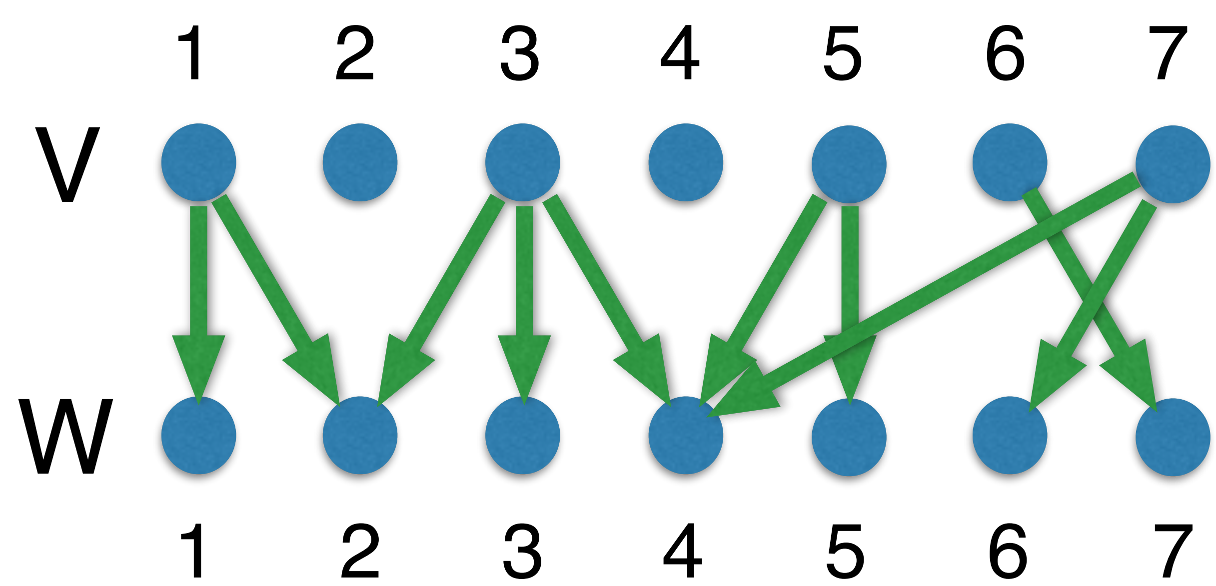

INPUT: Affectance matrix $A=[a(u,(v,w))]$ and a family $F = \{F_w \mid w \in W\}$ of subsets of transmitters connected to each receiver:

$F_1=\{1\}$

$F_2=\{1,3\}$

$F_3=\{3\}$

$F_4=\{3,5,7\}$

...

OUTPUT: Family $S = \{S_t \mid t=1,2,3,\dots\}$ of subsets of transmitters transmitting in each time slot:

$S_1=\{1,5,7\}$

$S_2=\{2,3,4,6\}$

$S_3=\{1,4,7\}$

$S_4=\{2,5\}$

...

Affectance-selective Families

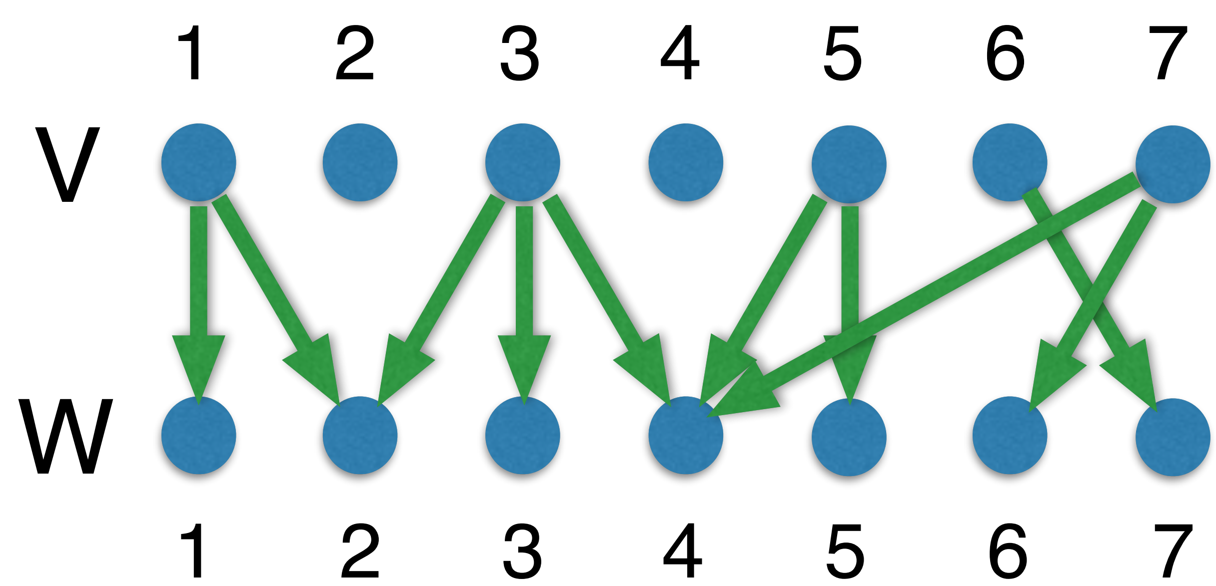

Transmissions schedule

A family $\mathcal{S} = \{S_1, S_2, \ldots, S_t\}$ of subsets of $[n]$ is *affectance-selective* on the family $\mathcal{F} = \{F_1, F_2, \ldots, F_n\}$ of subsets of $[n]$ if and only if, for each $w \in [n]$, there exists $j \in [t]$ such that:

Subsets of transmitters

- $|F_w \cap S_j| \geq 1$, and
- for some $v \in (F_w \cap S_j)$ it is $\sum_{u \in S_j} a(u,(v,w)) < 1$.

We say that the family $\mathcal{S}$ has *length* t, and that each w is *affectance-selected*.

[1] B.S.Chlebus, L. Gasieniec, A. Gibbons, A. Pelc,and W. Rytter. Deterministic broadcasting in ad hoc radio networks. *Distributed Computing*, 15(1):27–38, 2002.

Layer Dissemination

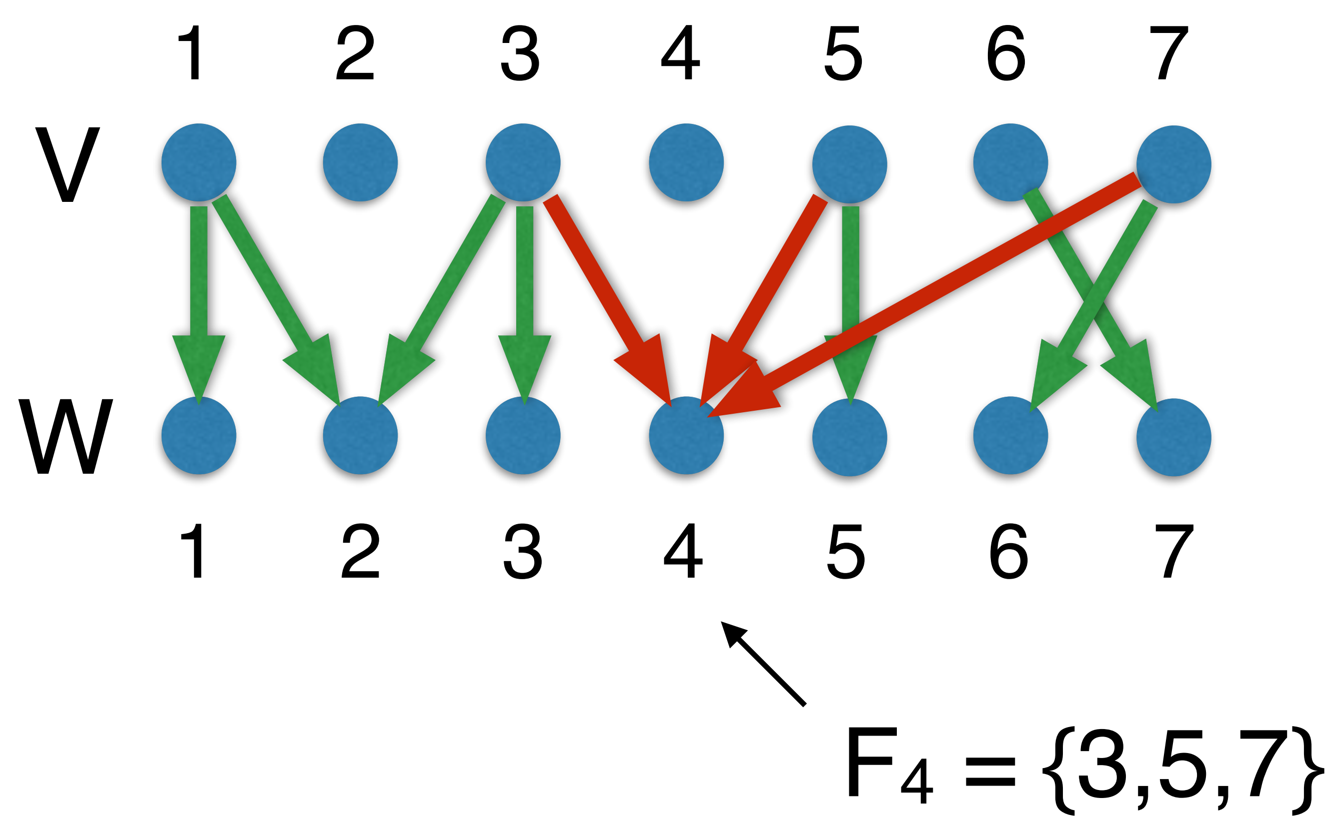

Bound family size on network characteristic

Maximum Average Affectance:

$$\overline{A} = \max_{w \in [n]} \max_{F \subseteq F_w} \sum_{v \in F} \sum_{u \in [n]} a(u,(v,w))/|F|.$$

Layer Dissemination

Existence of Aff-selective families:

▶ **Theorem 1.** *For any* $n > 0$*, consider a family* $\mathcal{F} = \{F_1, F_2, \ldots, F_n\}$ *of subsets of integers in* $[n]$ *and any affectance matrix* A *defined on* $\mathcal{F}$*. For each* $w \in [n]$*, let* $\overline{A}_w = \max_{F \subseteq F_w} \sum_{v \in F} \sum_{u \in [n]} a(u,(v,w))/|F|$ *be the maximum average affectance on* w*. If there exists a constant* $c > 1$ *such that* $\overline{A}_w \le c|F_w|$ *for all* $w \in [n]$*, then, there exists a family* $\mathcal{S} = \{S_1, S_2, \ldots, S_s\}$ *that is affectance-selective on* $\mathcal{F}$*, and its size* s *satisfies*

$$s \in O\left(1 + \log n \log \overline{A}\right),$$

where $\overline{A} = \max_{w \in [n]} \overline{A}_w$ *is the maximum average affectance.*

Logarithmic in the network characterization.

Layer Dissemination

Sketch of proof:

- Assume each v transmits with some probability p.
- Probabilistic method: show that the probability of a given w not being selected is < 1.
- Prove Markov-type inequality: to bound such probability by the expected average affectance on w.
- Prob is <1 if p is within a constant-factor b range -> try all.
- Union bound: after enough number of rounds, the probability of any w not being selected is still less than 1 -> add some multiplicity m.

We redefine $\mathcal{S}$ as the family $\{S_{i,j}\}$ of subsets of $[n]$ where the set $S_{i,j}$ is obtained including each $v \in [n]$ in $S_{i,j}$ independently with probability $p = 1/b^i$, for each $i = 0, 1, 2, \ldots, \max\{\lceil \log_b(2\overline{A}) \rceil, 0\}$ and each $j = 1, 2, \ldots, m$.

$$Pr(\exists w \in [n] : Z_w = 0) \leq nd^m \longrightarrow s \in O\left(1 + \log n \log \overline{A}\right)$$

Layer Dissemination

Proof of Thm 1 yields a randomized protocol of same length:

$b \leftarrow 1 + 1/(2c)$
$m \leftarrow \lceil 2 \log_{1/d} n \rceil$
for $i = 0, 1, 2, \ldots, \max\{\lceil \log_b(2\overline{A}) \rceil, 0\}$ **do**
  for m *times* **do**
    transmit with probability $1/b^i$

Algorithm 1: Randomized Layer Dissemination protocol for each node $v \in$

Theorem 2 *Consider a layer of a Radio Network with affectance matrix A and topology $G = (V, W, E)$, where $|V| = |W| = n$, where for each receiver $w \in W$ there is at least one transmitter $v \in V$ such that $(v, w) \in E$. Then, if there exists a constant $c > 1$ such that $\overline{A}_w \leq c|F_w|$ for all $w \in W$, where $\overline{A}_w = \max_{F \subseteq F_w} \sum_{v \in F} \sum_{u \in V} a(u, (v, w))/|F|$ is the maximum average affectance on w, Algorithm 1 solves the Layer Dissemination problem with high probability* [1]*, and the running time is in $O(1 + \log n \log \overline{A})$, where $\overline{A} = \max_{w \in W} \overline{A}_w$ is the maximum average affectance.*

Layer Dissemination

De-randomization yields a deterministic protocol of same length:

// Initialization

$p \leftarrow 0$

$b \leftarrow 1 + 1/(2c)$

$m \leftarrow \max\{\lceil \log_b(2\overline{A}) \rceil, 0\}$

$W'_0 \leftarrow \{w \in W : \overline{A}_w \leq 1/2\}$

for $r = 1, \ldots, m$ **do** $W'_r \leftarrow \{w \in W : b^{r-1}/2 < \overline{A}_w \leq b^r/2\}$

// Protocol

for *each time slot* **while** $\exists r = 0, 1, \ldots, m : W'_r \neq \emptyset$ **do**

    if $p \leq 1/(2b\overline{A})$ **then**

        $p \leftarrow 1$

        $r \leftarrow 0$

    set $V'[1 \ldots n]$ array of booleans   // $V'[i] \equiv i$ transmits

    for $i = 1, 2, \ldots, n$ **do**

        $\mathbb{E}_{true} \leftarrow \mathbb{E}_{V'[i+1 \ldots n]}\left(\# \text{ selected in } W'_r \middle| V'[i] = true\right)$

        $\mathbb{E}_{false} \leftarrow \mathbb{E}_{V'[i+1 \ldots n]}\left(\# \text{ selected in } W'_r \middle| V'[i] = false\right)$

        $V'[i] \leftarrow \mathbb{E}_{true} > \mathbb{E}_{false}$

    if $V'[v]$ **then** transmit

    $W'_r \leftarrow W'_r \setminus \{w | w \text{ was selected}\}$

    $p \leftarrow p/b$

    $r \leftarrow r + 1$

Algorithm 2: Deterministic Layer Dissemination protocol for each node $v \in$

… but computing those expectations is exponential,
due to computing probs of low affectance.

So?

Are affectance-based protocols better? worse?

Let's try some experiments!

Layer Dissemination

Simulations topology

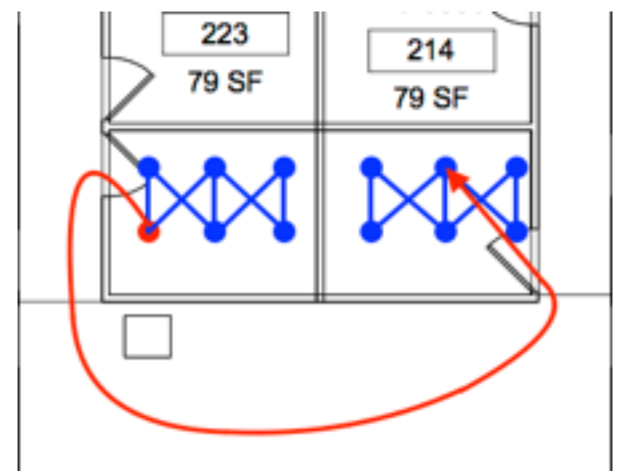

Layer Dissemination

Simulation protocols:

- **Transform Montecarlo into Las Vegas protocols:**
  - count how many rounds to complete dissemination.
- **Affectance model evaluation:**
  - Compare performance with an RN protocol and and SINR protocol.
  - Successful transmission according to Affectance model.
  - Compare also with theoretical performance.

Layer Dissemination

Simulations results

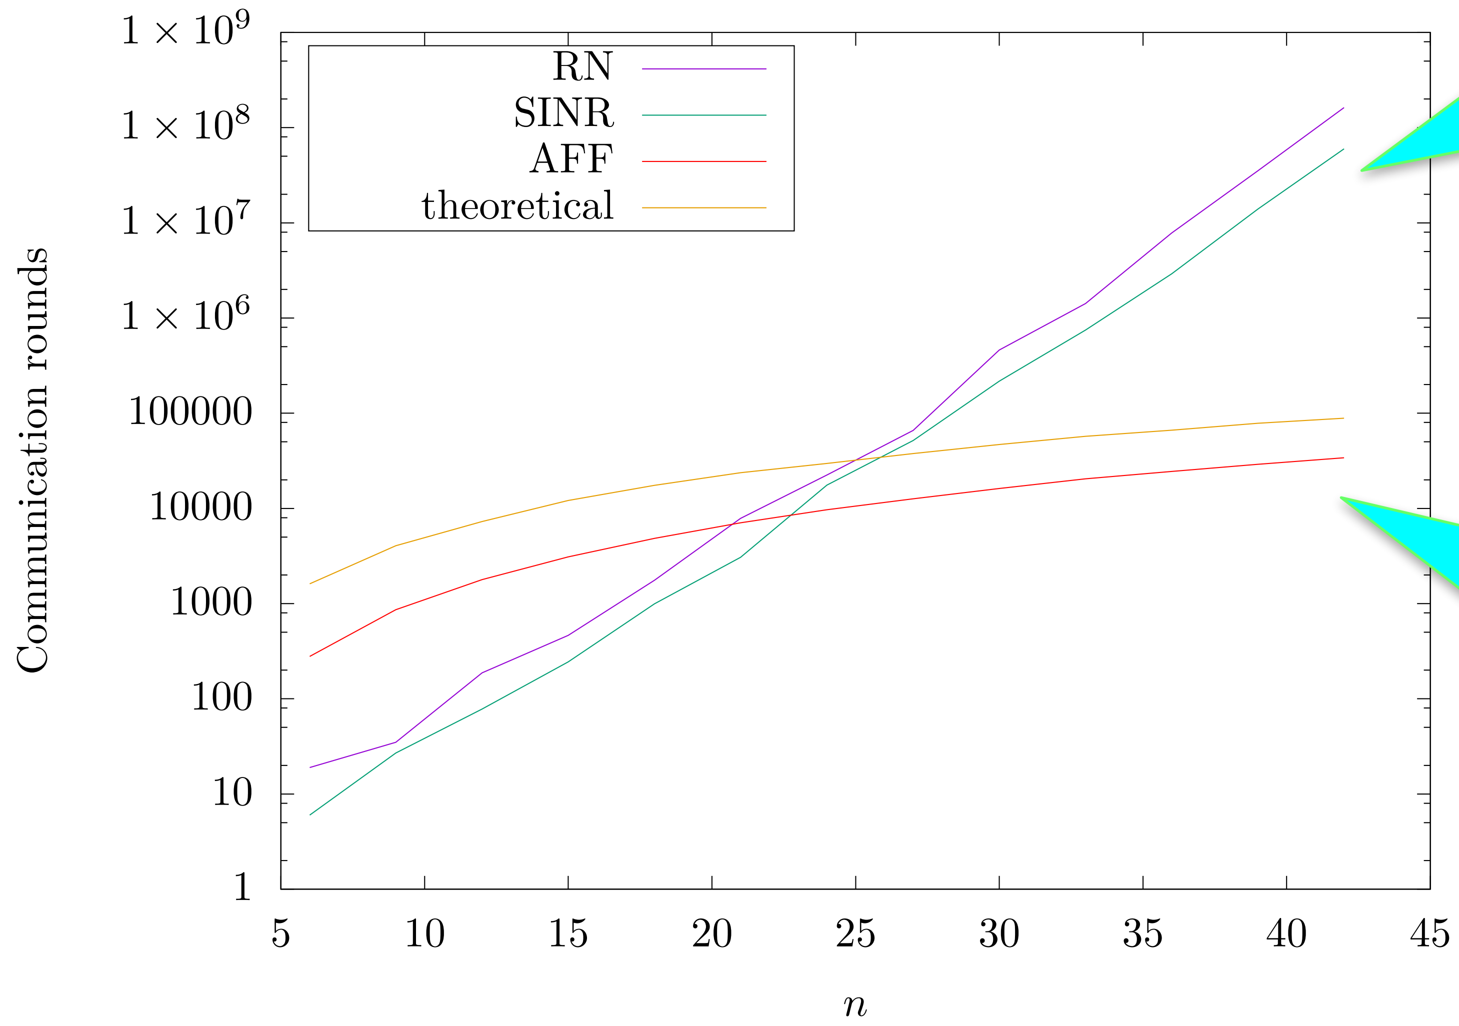

RN & SINR are exponential under affectance!

Better than theoretical bound!

Thank you!

Miguel A. Mosteiro
Pace University
mmosteiro@pace.edu