Algorithmic Aspects of Throughput-Delay Performance for Fast Data Collection in Sensor Networks

Amitabha Ghosh
Committee:
Bhaskar Krishnamachari
C.S. Raghavendra
Gaurav S. Sukhatme

Autonomous Networks Research Group
Ming Hsieh Department of Electrical Engineering
USC Viterbi School of Engineering
amitabhg@usc.edu
http://anrg.usc.edu/~amitabhg

July 27, 2010

Thesis Statement:

Multi-channel **scheduling**, optimal **routing** topologies, and **transmission power control** are necessary in order to improve the **throughput-delay** trade-off for aggregated convergecast in large-scale, multi-hop, wireless sensor networks.

Outline

Outline

Model and Assumptions

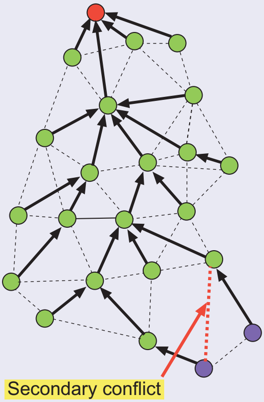

e.g., Applications with fast & efficient data collection requirements.

Model and Assumptions

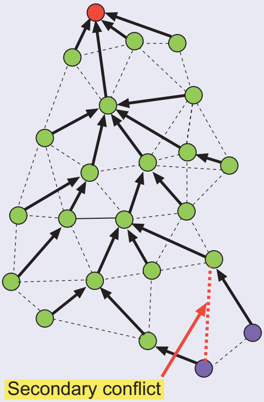

e.g., Applications with fast & efficient data collection requirements.

- TDMA periodic scheduling
  (- every node generates a single packet at the beginning of each frame
  - every node transmits once per slot)
- Graph-based network & interference model
- Half-duplex, single transceiver
- Receiver-based Frequency Assignment
- Non-interfering orthogonal channels
- In-network aggregation
  (distributive, algebraic
  - aggregation happens at each hop)

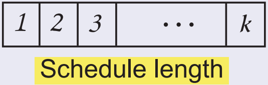

Convergecast and Benefits of Multiple Frequencies

Convergecast

A many-to-one communication pattern (i.e., flow of data): from all the nodes to a sink (opposite of broadcast).

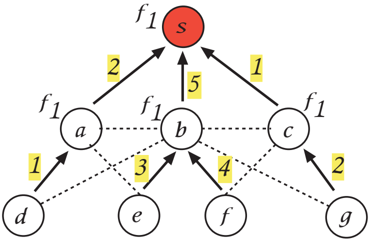

Figure: Schedule length is 5.

Convergecast and Benefits of Multiple Frequencies

Convergecast

A many-to-one communication pattern (i.e., flow of data): from all the nodes to a sink (opposite of broadcast).

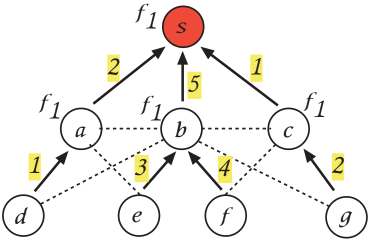

Figure: Schedule length is 5.

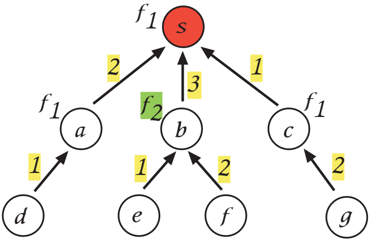

Figure: Schedule length is 3.

Convergecast and Benefits of Multiple Frequencies

Convergecast

A many-to-one communication pattern (i.e., flow of data): from all the nodes to a sink (opposite of broadcast).

	Frame 1			Frame 2		
Receiver	Slot 1	Slot 2	Slot 3	Slot 1	Slot 2	Slot 3
s	*c*	*a d*	*b e f*	*c g*	*a d*	*b e f*
a	*d*			*d*		
b	*e*	*f*		*e*	*f*	
c		*g*			*g*	

Pipelining

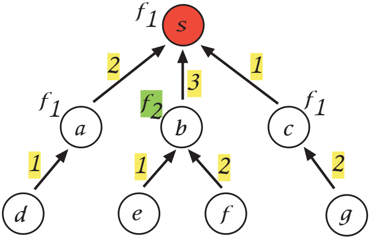

Figure: Schedule length is 3.

Problem Formulation

Joint Frequency and Time Slot Assignment

Given a connected network $G = (V, E)$ of arbitrarily deployed nodes, K orthogonal frequencies, and a routing tree $T \subseteq G$ rooted at the sink $s \in V$

- Assign a frequency to each receiver in T
- Assign a time slot to each edge in T

such that the aggregated throughput at the sink is maximized, where

Aggregated Throughput $=$ 1/Schedule Length
(equivalent to minimizing the Schedule Length)

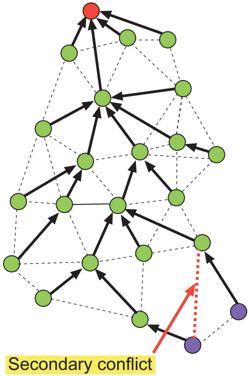

Scheduling Complexity

Theorem

Joint Frequency and Time Slot Assignment is NP-hard on general graphs.
Proof: Reduction from Vertex Color.

Questions

- Does the problem still remain NP-hard with unlimited frequencies?
- If not, how many frequencies (at a min, or at most) are needed?
- Can we design algorithms that have provable, worst-case guarantee?

Scheduling Complexity

Theorem

Joint Frequency and Time Slot Assignment is NP-hard on general graphs.
Proof: Reduction from Vertex Color.

Questions

- Does the problem still remain NP-hard with unlimited frequencies?
- If not, how many frequencies (at a min, or at most) are needed?
- Can we design algorithms that have provable, worst-case guarantee?

Theorem

Given a spanning tree T on a graph $G = (V, E)$, finding the minimum number of frequencies required to remove all secondary conflicts is NP-hard.

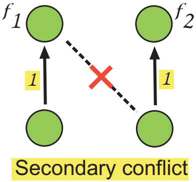

A Frequency Upper Bound

Constraint Graph

- For each receiver in G create a node in G_C.
- Create an edge between two nodes in G_C if the corresponding receivers in G form a secondary conflict.

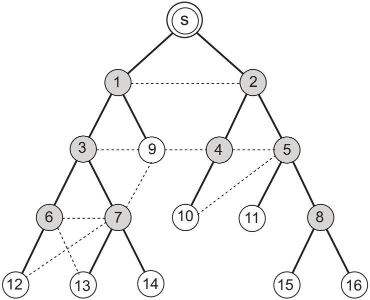

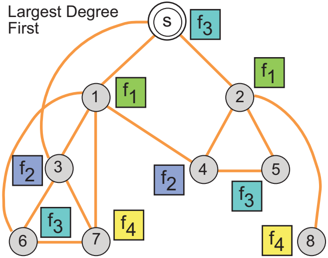

Lemma

$K_{max} \leq \Delta(G_C) + 1$, where $\Delta(G_C)$: max degree in G_C. (Vertex Color)

Time Slot Assignment

Algorithm: BFS-TimeSlotAssignment

1. **while** $E_T \neq \phi$ **do**
2. $e \leftarrow$ next edge from E_T in BFS order
3. Assign minimum time slot to e respecting adjacency constraint
4. $E_T \leftarrow E_T \setminus \{e\}$
5. **end while**

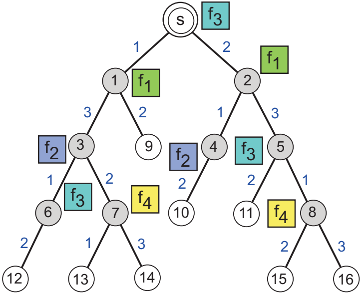

Theorem

BFS-TimeSlotAssignment gives a schedule of minimum length $\Delta(T)$, where $\Delta(T)$: maximum node degree in T.

Evaluation

Parameters:

- Number of nodes: $n : 200$, square region of area $A : 200 \times 200$
- Shortest Path Tree, Unit Disk Graph

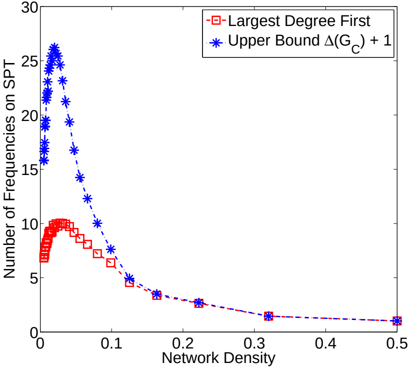

Figure: Number of frequencies required to remove all secondary conflicts as a function of network density n/A on SPT.

So far...

Joint Frequency and Time Slot Assignment is

- NP-hard with limited (constant number) frequencies on arbitrary graphs.
- Polynomial with unlimited (at most $\Delta(G_C) + 1$) frequencies for arbitrary graphs.

Next...

Approximation algorithms for Random Geometric Graphs with limited (constant number) frequencies.

- Uniform transmission range (Unit Disk Graphs)
- Non-uniform transmission range (General Disk Graphs)

Approximation on UDG

Algorithm $\mathcal{A}_{UDG}$ Overview

- Phase 1: Assign frequencies to the receivers in each cell such that the maximum number of nodes transmitting on the same frequency (load) is minimized. (Load-Balanced Frequency Assignment)
- Phase 2: Assign time slot greedily to each edge.

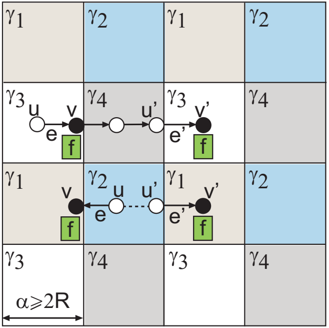

Approximation on UDG

Algorithm $\mathcal{A}_{UDG}$ Overview

- Phase 1: Assign frequencies to the receivers in each cell such that the maximum number of nodes transmitting on the same frequency (load) is minimized. (Load-Balanced Frequency Assignment)
- Phase 2: Assign time slot greedily to each edge.

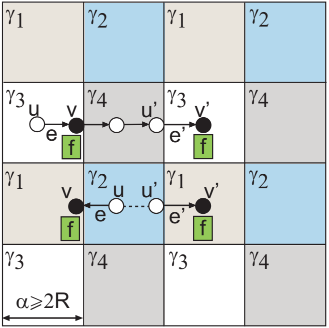

Theorem

Algorithm $\mathcal{A}_{UDG}$ achieves a constant factor $8\mu_\alpha \cdot \left(\frac{4}{3} - \frac{1}{3K}\right)$ approximation on the optimal schedule length, where $\mu_\alpha > 0$ is a constant for any given cell size $\alpha \geq 2R$, and given node deployment, where
μ_α : maximum number of edges on the same frequency in any cell that can be scheduled simultaneously by any algorithm.

Evaluation

Parameters

UDG on a region of size 200×200 with transmission radius $R = 25$

Shortest Path Tree

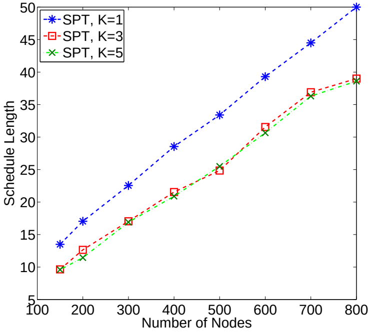

Observations

- Schedule length decreases with increasing number of frequencies
- However, with diminishing returns
- Structure of the routing tree (high node degree) may lead to bottlenecks

Approximation on Disk Graphs

Approximation on Disk Graphs

Notations

- $r(u)$: Transmission range of node u
- $\ell(e)$: Length of edge e
- $I(e)$: Set of edges that are either adjacent to e or form a secondary conflict with e
- $I_{\geq}(e)$: Subset of $I(e)$ that have larger disks than e

$$I_{\geq}(e) = \{e' = (u', v') : e' \in I(e), \ell(e') \geq \ell(e)\}$$

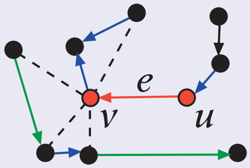

Approximation on Disk Graphs

Algorithm $\mathcal{A}_{DG}$ Overview

- Phase 1: Assign frequencies such that the maximum number of edges interfering with any given edge is minimized.
  - $0-1$ Integer Linear Program: Define indicator variables X_{vk} for edge $e=(u,v)$ as:

$$X_{vk} = \begin{cases} 1, & \text{if receiver } v \text{ is assigned frequency } f_k \\ 0, & \text{otherwise} \end{cases}$$

  A frequency assignment is a $0-1$ assignment to $X_{vk}, \forall e, \forall f_k$.

Approximation on Disk Graphs

Algorithm $\mathcal{A}_{DG}$ Overview

- Phase 1: Assign frequencies such that the maximum number of edges interfering with any given edge is minimized.
  - $0-1$ Integer Linear Program: Define indicator variables X_{vk} for edge $e=(u,v)$ as:

$$X_{vk} = \begin{cases} 1, & \text{if receiver } v \text{ is assigned frequency } f_k \\ 0, & \text{otherwise} \end{cases}$$

    A frequency assignment is a $0-1$ assignment to $X_{vk}, \forall e, \forall f_k$.

- Phase 2: Sort edges in non-increasing order of lengths, and assign smallest available slots starting from the largest.

Approximation on Disk Graphs

0-1 Integer Linear Program of Frequency Assignment

$$\textbf{Minimize} \quad \lambda$$

$$\begin{aligned}
\text{subject to} \quad &: \\
\forall e = (u,v), \forall f_k \quad &: \quad \sum_{v'} n(e,v') X_{vk} \leq \lambda && (1) \\
\forall e = (u,v) \quad &: \quad \sum_{f_k} X_{vk} = 1, && (2) \\
\forall e = (u,v), \forall f_k \quad &: \quad X_{vk} \in \{0,1\} && (3)
\end{aligned}$$

- Solve the Linear Relaxation (LP) by modifying constraint (3)
- Randomized Rounding: Assign $Y_{vk} = 1$ with probability X^*_{vk}, where X^*_{vk} are the optimal fractional LP solutions, and Y_{vk} are the new integral random variables.

Approximation on Disk Graphs

Lemma

Let Y_{vk} be the rounded solution, as described above. Then,

$$\max_{e,f_k}\left\{\sum_{e'=(u',v')\in I_{\geq}(e)} Y_{v'k}\right\} = O\left(\Delta(T)\log n\cdot\lambda^*\right),$$

with probability at least $(1-1/n)$.
Proof: Apply Chernoff bound.

Approximation on Disk Graphs

Lemma

Let Y_{vk} be the rounded solution, as described above. Then,

$$\max_{e, f_k} \left\{ \sum_{e'=(u',v') \in I_{\geq}(e)} Y_{v'k} \right\} = O\left(\Delta(T) \log n \cdot \lambda^*\right),$$

with probability at least $(1 - 1/n)$.
Proof: Apply Chernoff bound.

Theorem

The schedule constructed by time slot assignmentt strategy along with the frequency assignment using the above randomized rounding procedure results in a schedule of length $O(\Delta(T) \cdot \log n)$ times the optimum.

Outline

Realistic Interference Model

Limitations of a Graph-Based Model (Protocol Model)

- Interference is binary and stops abruptly beyond a distance - idealizes physical laws
- Fails to capture cumulative interference from (many) far-away transmitters
- Sometimes pessimistic decisions and conflicting schedules

Realistic Interference Model

SINR Model (Physical Model)

For a given frequency f, the received signal power from sender s_i to its intended receiver r_i is:

$$P_{r_i}(s_i) = \frac{P}{d(s_i, r_i)^{\alpha}},$$

where $\alpha \in [2, 6]$ is called the path-loss exponent; value depends on external conditions.

Realistic Interference Model

SINR Model (Physical Model)

For a given frequency f, the received signal power from sender s_i to its intended receiver r_i is:

$$P_{r_i}(s_i) = \frac{P}{d(s_i, r_i)^\alpha},$$

where $\alpha \in [2, 6]$ is called the path-loss exponent; value depends on external conditions.

A transmission from sender s_i to its intended receiver r_i is successful if the ratio, $SINR_{r_i}$, of the received signal power to the cumulative interference plus noise $\mathcal{N}$ at r_i is more than a hardware-dependent threshold β:

$$SINR_{r_i} = \frac{\frac{P}{d(s_i,r_i)^\alpha}}{I_{r_i} + \mathcal{N}} = \frac{\frac{P}{d(s_i,r_i)^\alpha}}{\sum_{j \neq i} \frac{P}{d(s_j,r_i)^\alpha} + \mathcal{N}} \geq \beta$$

Realistic Interference Model

Parameters

- CC2420 radio parameters with receiver sensitivity -95 dB
- Path-loss exponent $\alpha = 3.5$, transmit power -6 dB

Percentage of Conflicting Schedules: SINR vs. Graph-based Model

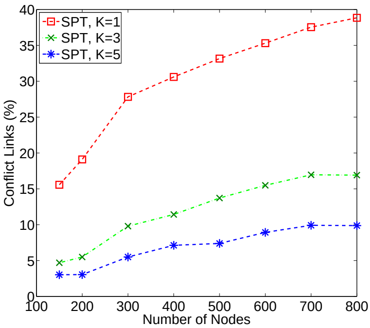

Observations

- For a given frequency, the amount of conflict increases as the network gets denser, reaching almost 40% for densest deployment with a single frequency
- With multiple frequencies, the amount of conflict is much less

Problem Formulation in SINR

Joint Frequency and Time Slot Assignment

Given a connected network of arbitrarily deployed nodes, K orthogonal frequencies, and a routing tree $T = (V, E_T)$ rooted at the sink $s \in V$

- Assign a frequency to each receiver in T
- Assign a time slot to each edge in T

such that the schedule length is minimized, subject to

$SINR_{r_i} \geq \beta$ at every receiver r_i

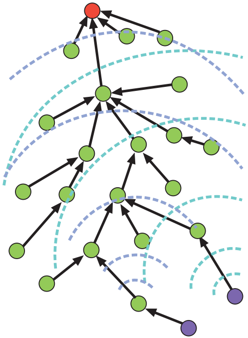

Cumulative interference

Approach

Link Diversity

Classify the edges in E_T based on their lengths $\ell_i = d(s_i, r_i)$. Link diversity is the number of magnitude of different lengths

$$g(E_T) = |\{m \mid \exists\, e_i, e_j \in E_T : \lfloor \log(\ell_i/\ell_j) \rfloor = m\}|$$

Usually a small constant in practice, but theoretically can be equal to the number of nodes

Approach

Link Diversity

Classify the edges in E_T based on their lengths $\ell_i = d(s_i, r_i)$. Link diversity is the number of magnitude of different lengths

$$g(E_T) = |\{m \mid \exists\, e_i, e_j \in E_T : \lfloor \log(\ell_i/\ell_j) \rfloor = m\}|$$

Usually a small constant in practice, but theoretically can be equal to the number of nodes

Overall Idea

- Normalize (w.l.o.g) the minimum edge length to one
- Let

  $$E = \left\{E_0, \ldots, E_{\log(e_{max})}\right\}$$

  denote the set of non-empty length classes, where E_k is the set of edges of lengths in $\left[2^k, 2^{k+1}\right)$, and e_{max} is the longest edge
- Partition the problem into disjoint length classes, and process edges in each class separately

Approach

For Length Class E_k

- Divide the 2-D region into a set C_k of square grids of length $\eta_k = \delta \cdot 2^k$, for some constant δ (will choose δ later)
- Frequency Assignment:
  - For each cell in C_k, run FrequencyGreedy on the receivers
- Time Slot Assignment:
  - Four color the cells in C_k
  - Can we still reuse time slots across non-adjacent cells of the same color? If so, how many?

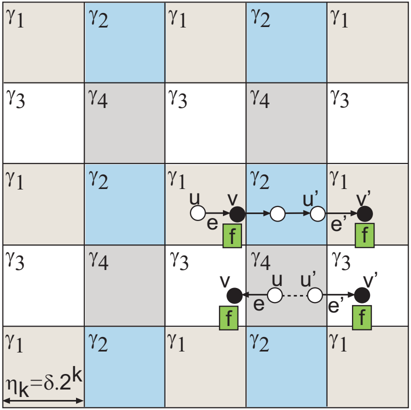

Approximation Algorithm in SINR

Theorem

For cells of the same color in C_k, we can simultaneously schedule at most one edge in E_k from every non-adjacent cell that lies on the same frequency, so long as the cell size satisfies:

$$\eta_k = \delta \cdot 2^k, \quad \text{for } \delta = 4\left[8\beta \cdot \frac{\alpha - 1}{\alpha - 2}\right]^{\frac{1}{\alpha}},$$

where $\beta \geq 1$ is the SINR threshold, and $\alpha > 2$ is the path-loss exponent.

Approximation Algorithm in SINR

Theorem

For cells of the same color in C_k, we can simultaneously schedule at most one edge in E_k from every non-adjacent cell that lies on the same frequency, so long as the cell size satisfies:

$$\eta_k = \delta \cdot 2^k, \quad \text{for } \delta = 4\left[8\beta \cdot \frac{\alpha - 1}{\alpha - 2}\right]^{\frac{1}{\alpha}},$$

where $\beta \geq 1$ is the SINR threshold, and $\alpha > 2$ is the path-loss exponent.

Proof Sketch

- Consider one specific edge $e_i = (s_i, r_i) \in E_k$ in any cell $c \in C_k$
- Received power at r_i

$$P_{r_i}(s_i) = \frac{P}{\ell_i^{\alpha}} \geq \frac{P}{2^{\alpha(k+1)}}, \qquad \because 2^k \leq \ell_i < 2^{k+1}$$

- Show that the cumulative interference I_{r_i} caused by concurrent transmissions from all non-adjacent cells is still within β

Approximation Algorithm in SINR

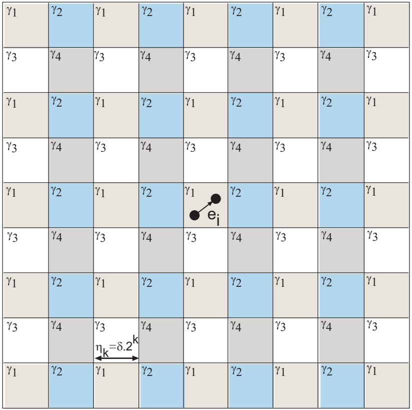

- Consider edge $e_i = (s_i, r_i) \in E_k$

Approximation Algorithm in SINR

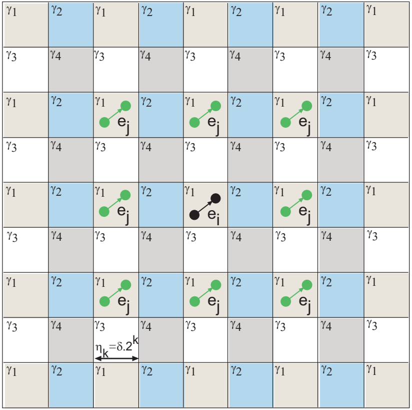

- Consider edge $e_i = (s_i, r_i) \in E_k$
- At most 8 from layer 1

Approximation Algorithm in SINR

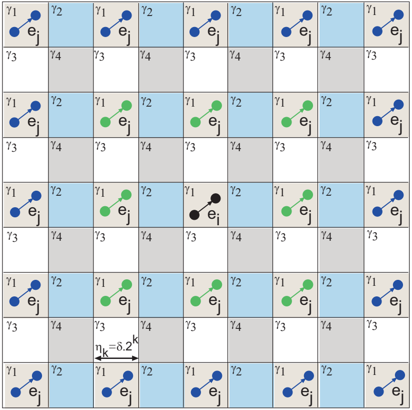

- Consider edge $e_i = (s_i, r_i) \in E_k$
- At most 8 from layer 1
- At most 16 from layer 2, ..., at most 8q from layer q
- Add up all the interferences up to layer ∞

Approximation Algorithm in SINR

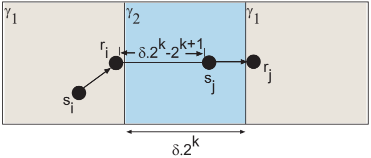

Computing Cumulative Interference I_{r_i}

- From transmitter s_j in a layer 1 non-adjacent cell of the same color

$$I_{r_i}(s_j) = \frac{P}{\ell_{ij}^{\alpha}} \leq \frac{P}{[2^k(\delta - 2)]^{\alpha}}, \qquad \because \ell_{ij} \geq \delta \cdot 2^k - 2^{k+1}$$

Approximation Algorithm in SINR

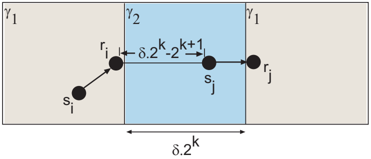

Computing Cumulative Interference I_{r_i}

- From transmitter s_j in a layer 1 non-adjacent cell of the same color

$$I_{r_i}(s_j) = \frac{P}{\ell_{ij}^{\alpha}} \leq \frac{P}{[2^k(\delta - 2)]^{\alpha}}, \qquad \because \ell_{ij} \geq \delta \cdot 2^k - 2^{k+1}$$

- In general, $\ell_{ij} \geq 2^k \{(2q-1)\delta - 2\}$ in layer q

Approximation Algorithm in SINR

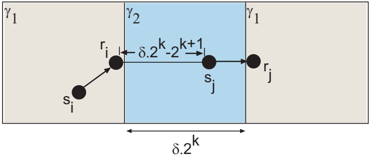

Computing Cumulative Interference I_{r_i}

- From transmitter s_j in a layer 1 non-adjacent cell of the same color

$$I_{r_i}(s_j) = \frac{P}{\ell_{ij}^{\alpha}} \leq \frac{P}{[2^k(\delta-2)]^{\alpha}}, \qquad \because \ell_{ij} \geq \delta \cdot 2^k - 2^{k+1}$$

- In general, $\ell_{ij} \geq 2^k \{(2q-1)\delta - 2\}$ in layer q
- Adding up

$$I_{r_i} \leq \sum_{q=1}^{\infty} \frac{8qP}{[2^k \{(2q-1)\delta - 2\}]^{\alpha}} = \frac{8P}{2^{(k-1)\alpha}\delta^{\alpha}} \cdot \frac{\alpha-1}{\alpha-2} \leq \frac{P_{r_i}(s_i)}{\beta}$$

Approximation Algorithm in SINR

Theorem : $O(g(E_T))$ Approximation

Given a routing tree T on an arbitrarily deployed connected network in 2-D, and K orthogonal frequencies, the algorithm gives a $O\left(g(E_T)\right)$ approximation on the optimal schedule length.

Approximation Algorithm in SINR

Theorem : $O(g(E_T))$ Approximation

Given a routing tree T on an arbitrarily deployed connected network in 2-D, and K orthogonal frequencies, the algorithm gives a $O\left(g(E_T)\right)$ approximation on the optimal schedule length.

Proof Sketch

- Choice of the critical cell $\hat{c}$ in which the number of edges on the same frequency $L_{\hat{c}}^{m^*}$ is maximum across all length classes
- Show that OPT will take at least $L_{\hat{c}}^{m^*}/\mu$ slots, for

$$\mu = \frac{\left[2(\sqrt{2}\cdot\delta + 1)\right]^{\alpha}}{\beta}$$

  - Show by contradiction that OPT cannot schedule $\mu + 1$ edges on the same frequency in any cell

Approximation Algorithm in SINR

Proof Sketch (cobmining together)

$$
\begin{aligned}
\Gamma &\leq 4 \cdot \max_{k,c} \{|\gamma_{c \in C_k}|\} \cdot g(E_T) \\
&\leq 8 \cdot \max_{k,c} \left\{ L^{\phi}_{c \in C_k} \right\} \cdot g(E_T) \\
&\leq 8 \cdot \max_{k,c} \left\{ \left(\frac{4}{3} - \frac{1}{3K} \right) \cdot L^{m^*}_{c \in C_k} \right\} \cdot g(E_T) \\
&\leq 8\mu \cdot \left(\frac{4}{3} - \frac{1}{3K} \right) \cdot g(E_T) \cdot OPT \\
&= O\left(g(E_T)\right)
\end{aligned}
$$

Outline

Properties of Spanning Trees

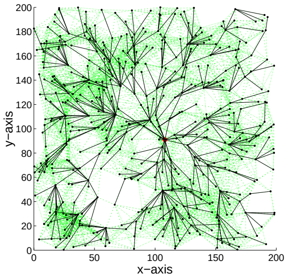

Shortest Path Tree (shallow and fat)

- High node degree ⇒ low throughput
- Few hops ⇒ low delay

Properties of Spanning Trees

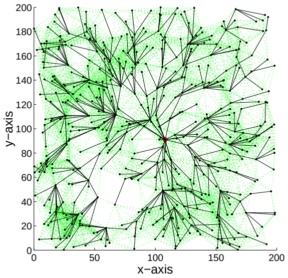

Shortest Path Tree (shallow and fat)

- High node degree $\Rightarrow$ low throughput
- Few hops $\Rightarrow$ low delay

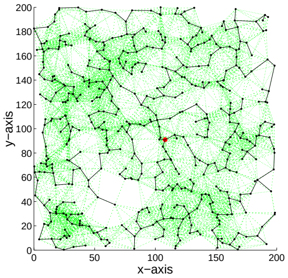

Minimum Interference Tree (weighted MST) [Mobihoc '04] (deep and skinny)

$w(u, v)$: no. of nodes covered by the union of two disks centered at u and v, each of radius $|uv|$

- Low node degree $\Rightarrow$ high throughput
- More hops $\Rightarrow$ high delay

Evaluation of $\mathcal{A}_{UDG}$ on SPT & MIT

Parameters

Random geometric graph on a region of size 200×200 with $R = 25$

Shortest Path Tree

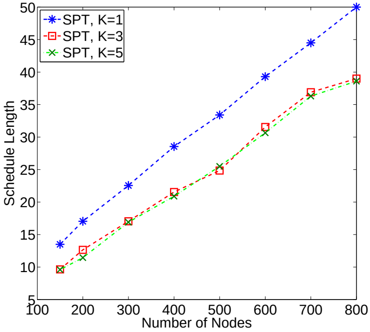

Minimum Interference Tree

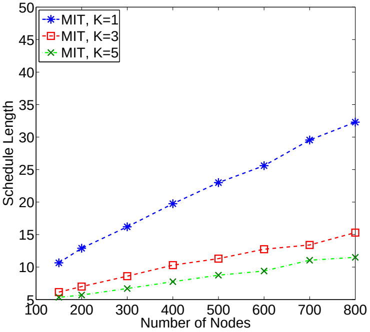

Tree Property for Bounded Throughput

Theorem

If the maximum node degree of the routing tree is bounded by a constant $\Delta_C > 0$, then there exists an algorithm that gives a constant factor $8\mu_\alpha \cdot \Delta_C$ approximation on the optimal schedule length, where $\mu_\alpha > 0$ is a constant for a given cell size $\alpha \geq 2R$ and given node deployment.

Tree Property for Bounded Throughput

Theorem

If the maximum node degree of the routing tree is bounded by a constant $\Delta_C > 0$, then there exists an algorithm that gives a constant factor $8\mu_\alpha \cdot \Delta_C$ approximation on the optimal schedule length, where $\mu_\alpha > 0$ is a constant for a given cell size $\alpha \geq 2R$ and given node deployment.

Properties

- Degree-bounded spanning trees? Minimum Degree Spanning Tree is NP-hard [SODA'92]
- Is constant factor approximation on schedule length and delay achievable?

Bicriteria Formulation

Bounded-Degree Minimum-Radius Spanning Tree

$\mathcal{P} =$ (Degree, Radius, Spanning Tree)
Radius: $R(T)$: Maximum hop distance from any node to sink in T
Δ^*: Budget on max degree
Goal: Find a spanning tree of min radius subject to max degree Δ^*

Bicriteria Formulation

Bounded-Degree Minimum-Radius Spanning Tree

$\mathcal{P} =$ (Degree, Radius, Spanning Tree)
Radius: $R(T)$: Maximum hop distance from any node to sink in T
Δ^*: Budget on max degree
Goal: Find a spanning tree of min radius subject to max degree Δ^*

Approximation Definition

Algorithm $\mathcal{A}$ is an (α, β)-approximation to a bicriteria optimization problem $\mathcal{P} =$ (Degree, Radius, Spanning Tree) if:

- $\Delta(T) \leq \alpha + \Delta^*$
- $R(T) \leq \beta \cdot R^*(T)$

where $R^*(T)$ is the min radius of any spanning tree T whose max degree is Δ^*.

Algorithm Overview

Algorithm $\mathcal{A}_{BDMRST}$

Three phases:

1. Construct a Global Backbone Tree of low radius.
2. Construct degree-bounded Local Spanning Trees.
3. Merge the local spanning trees and the global backbone tree.

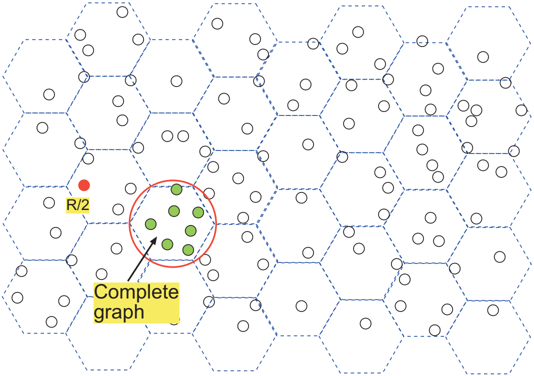

Algorithm Overview

Global Backbone Tree

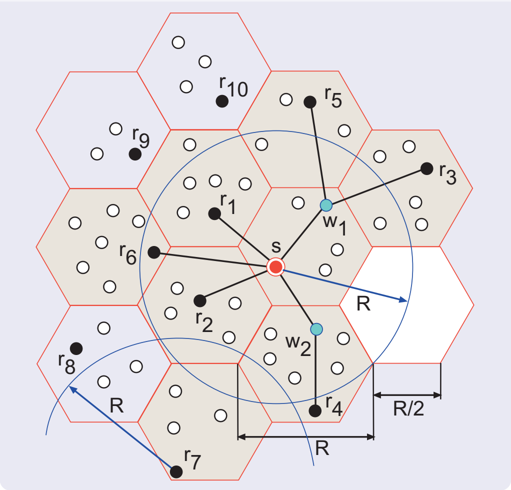

Algorithm Overview

Global Backbone Tree

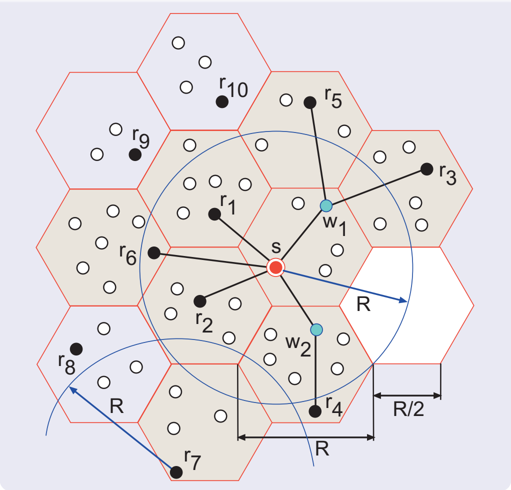

Local Spanning Tree

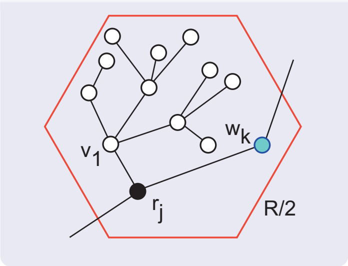

Properties of Spanning Trees

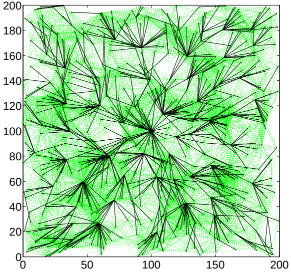

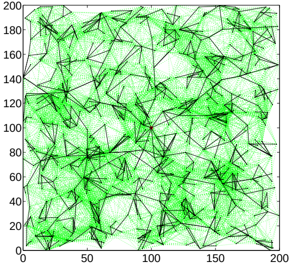

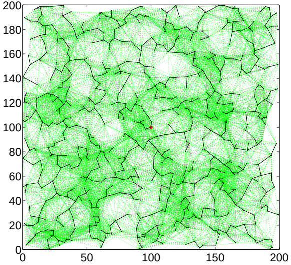

Bounded-Degree Minimum-Radius Spanning Tree, $\Delta^* = 4$

Approximation Result & Evaluation

Theorem

Algorithm $\mathcal{A}_{BDMRST}$ gives a constant factor (α, β) bicriteria approximation for the Bounded-Degree Minimum-Radius Spanning Tree, where $\alpha = 10$ and $\beta = 5$.

Approximation Result & Evaluation

Theorem

Algorithm $\mathcal{A}_{BDMRST}$ gives a constant factor (α, β) bicriteria approximation for the Bounded-Degree Minimum-Radius Spanning Tree, where $\alpha = 10$ and $\beta = 5$.

Schedule Length

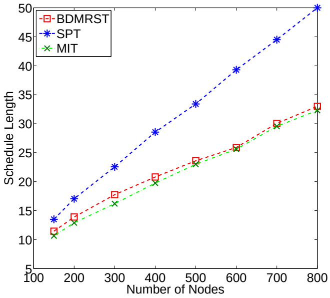

Max Delay

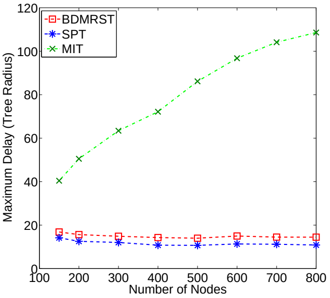

Outline

Contributions

- Maximizing aggregated convergecast throughput using multi-channel scheduling
  - Constant factor approximation on UDG
  - $\Delta(T) \cdot \log(n)$ approximation on DG
  - $O(g(E_T))$ approximation in SINR
- Throughput-delay trade-off in the same optimization framework
  - Bicriteria formulation
  - Constant factor approximation on constructing bounded-degree minimum-radius spanning trees
- Efficient, distributed topology control in 3-D
  - Simple orthographic projection based approach extending 2-D results
  - Robust spherical Delaunay triangulation based algorithm

Publications

Book Chapters

1. O.D. Incel, A. Ghosh, and B. Krishnamachari, Scheduling Algorithms for Tree-Based Data Collection in Wireless Sensor Networks, Theoretical Aspects of Distributed Computing in Sensor Networks, Springer, 2010.
2. A. Ghosh and S.K. Das, Coverage and Connectivity Issues in Wireless Sensor Networks, Mobile, Wireless and Sensor Networks: Technology, Applications and Future Directions, John Wiley & Sons, 2006.

Journal Papers

1. A. Ghosh, O.D. Incel, V.S. Anil Kumar, and B. Krishnamachari, Multi-Channel Scheduling and Spanning Tree Algorithms: Throughput-Delay Trade-off for Fast Convergecast in Sensor Networks, ACM/IEEE Transactions on Networking, Feb 2010 (sub).
2. O.D. Incel, A. Ghosh, B. Krishnamachari, and K. Chintalapudi, Fast Data Collection in Tree-Based Wireless Sensor Networks, IEEE Transactions on Mobile Computing, Dec 2009 (rev).
3. A. Viswanath, P. Dutta, M. Chetlur, P. Gupta, S. Kalyanaraman, and A. Ghosh, Perspectives on Quality of Experience for Video Streaming over WiMAX, ACM SIGMOBILE Mobile Computing and Communications Review, Oct, 2009.
4. A. Ghosh and S.K. Das, Coverage Connectivity Issues in Wireless Sensor Networks: A Survey, Journal of Pervasive and Mobile Computing, 4(3):303-334, June 2008.

Publications

Conference Papers (reverse chronologically)

1. A. Ghosh, R. Jana, V. Ramaswami, J. Rowland, and N.K. Shankar, Modeling and Characterization of Large-Scale Wi-Fi Traffic in Public Hot-Spots, IEEE INFOCOM, 2011 (sub).
2. A. Ghosh, O.D. Incel, V.S. Anil Kumar, and B. Krishnamachari, Optimal Spanning Trees for Fast Data Collection in Sensor Networks, IEEE INFOCOM Student Workshop, Mar 2010.
3. A. Ghosh, O.D. Incel, V.S. Anil Kumar, and B. Krishnamachari, Multi-Channel Scheduling Algorithms for Fast Aggregated Convergecast in Sensor Networks, IEEE MASS, October 2009.
4. S.W. Lee, D. Yoon, and A. Ghosh, Intelligent Parking Lot Application using Wireless Sensor Networks, International Symposium on Collaborative Technologies and Systems (CTS), May 2008.
5. A. Ghosh, L. Pereira, T. Yan, and H. Cao, Modeling Wireless Sensor Network Architectures using AADL, European Congress on Embedded Real Time Software (ERTS), Jan 2008.
6. A. Ghosh, Y. Wang, and B. Krishnamachari, Efficient Distributed Topology Control in 3 Dimensional Wireless Networks, IEEE SECON, Jun 2007.
7. A. Ghosh and S.K. Das, A Distributed Greedy Algorithm for Connected Sensor Cover in Dense Sensor Networks, IEEE/ACM DCOSS, Jul 2005.
8. A. Ghosh, Estimating Coverage Holes and Enhancing Coverage in Mixed Sensor Networks, IEEE LCN, Nov 2004.
9. S.K. Ghosh and A. Ghosh, A Plan-Commit-Prove Protocol for Secure Verification of Traversal Path, IEEE International Conference on Networks (ICON), Nov 2004.

Manuscripts

1. T. Fu, A. Ghosh, E.A. Johnson, and B. Krishnamachari, Energy-Efficient Deployment Strategies in Structural Health Monitoring using Wireless Sensor Networks, Structural Control and Health Monitoring, 2009 (manuscript).
2. A. Ghosh, A.H. Patel, C.R. Sastry, Radio Frequency-Based Localization in Wireless Sensor Networks, Siemens Corporate Research Technical Report, Aug 2009.

Acknowledgments

Collaborators

- USC (EE, Civil): Bhaskar Krishnamachari, Erik Johnson, Yi Wang, Tat Fu
- Bogazici University: Ozlem Durmaz Incel
- Virginia Tech: Anil Vullikanti
- University of Texas: Sajal K. Das
- IBM Research: Shivkumar Kalyanaraman, Partha Dutta, Malolan Chetlur
- Siemens Corporate Research: Chellury Ram Sastry
- Microsoft Research: Krishnakant Chintalapudi
- Eaton Corporation: Ting Yan, Luis Pereira, Hui Cao
- AT&T Labs Research: Vaidyanathan Ramaswami, Rittwik Jana, N.K. Shankar

Thank You!