Dynamic Channel Assignment for WLANs

William Wu[1]

[1]Department of Electrical Engineering, Stanford University

NEC Research Laboratories China

Collaborators: Bo Wang, Yongqiang Liu, Yong Xia

Paper link: PEITS 2008

Outline of Talk

Outline

WLAN Definition

- *WLAN* = Wireless network of *APs* and *stations*.
- Each station associates itself with one AP for Internet access.
- Cell = AP + all stations associated with that AP.
- Each cell uses one channel.
- Most traffic is downlink.

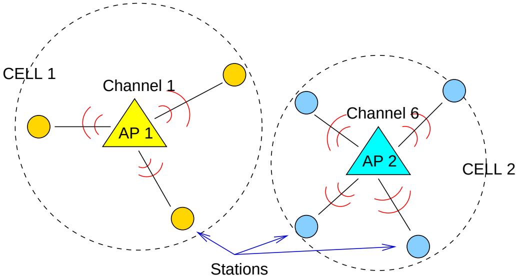

Three Problems

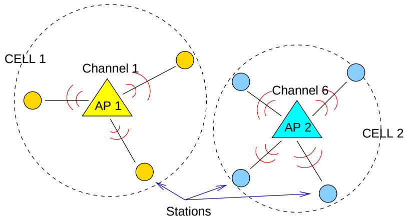

- Three network design problems:
  1. **Channel Assignment (CA):**
     Which channel does each cell use?
  2. **Power Management (PM):**
     What power level should we assign to each AP?
  3. **Station-AP Associations (UAA):**
     Which AP should each station associate with?

Big Picture

- Most existing solutions are one-shot (e.g., RF site survey).
- However, variations in WLAN traffic and link quality suggest need for dynamic solutions.

Dynamic Network Management

- *Monitor physical and link layer information.*
- Channel assignments, power allocations, and station-AP associations are *adjusted accordingly* to increase throughput.

This talk focuses only on the channel assignment problem.

Outline

DCA Outline

Dynamic Channel Assignment (DCA)

During every T second time-interval:

1. Report traffic + RSS to server.
2. If network has "sufficiently changed":
    1. Compute co-channel interference between any two cells.
    2. Compute channel reassignments to minimize sum interference (SDP Relaxation).
    3. APs change channels. Broadcast change to stations. (1-2 ms)

Small Example

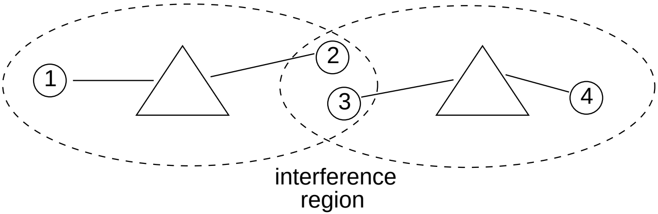

(a) 1 and 4 download, while 2 and 3 quiet $\Longrightarrow$ channel reuse.

(b) 2 and 3 download, while 1 and 4 quiet $\Longrightarrow$ different channels.

- ∴ channel assignment depends on both
  - traffic distribution (e.g., Who's downloading?)
  - locations of nodes (e.g., Who's in the "interference region"?)

Small Example

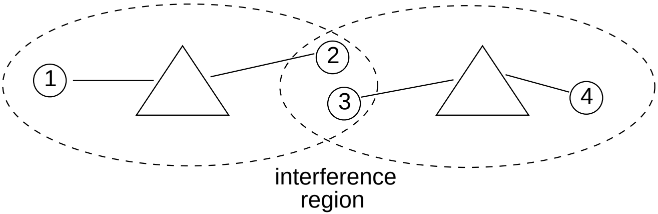

(a) 1 and 4 download, while 2 and 3 quiet $\implies$ channel reuse.

(b) 2 and 3 download, while 1 and 4 quiet $\implies$ different channels.

- ∴ channel assignment depends on both
  - traffic distribution (e.g., Who's downloading?)
  - locations of nodes (e.g., Who's in the "interference region"?)

Small Example

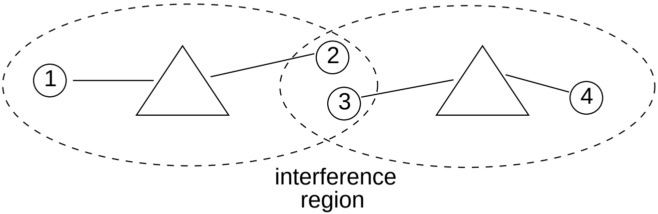

(a) 1 and 4 download, while 2 and 3 quiet $\Longrightarrow$ channel reuse.

(b) 2 and 3 download, while 1 and 4 quiet $\Longrightarrow$ different channels.

- $\therefore$ channel assignment depends on both
  - traffic distribution (e.g., Who's downloading?)
  - locations of nodes (e.g., Who's in the "interference region"?)

Network Measurements: Traffic

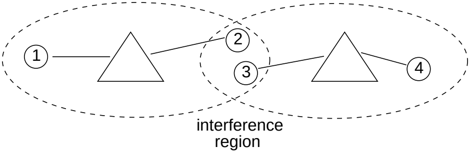

Want to know: Who is downloading?

- Every node, every T seconds, report:

$$\begin{aligned} Load &:= \frac{\text{throughput over past } T \text{ seconds}}{\text{datarate}} \\ &= \text{fraction of cell time downloading from AP} \end{aligned}$$

- $Load \approx$ probability of downloading.

Network Measurements: RSS

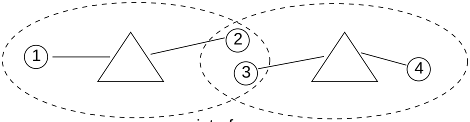

Want to know: What are the locations?

- *RSS* := Received Signal Strength.
  Received power at antenna.
- Maintain an "RSS matrix":

$$M_{ij} = \text{RSS at station } j \text{ when station } i \text{ alone is downloading, and } i \text{ and } j \text{ use the same channel.}$$

Serves like a gain matrix.

- **Initialization: Round-robin profiling**
  1. All quiet except for one AP, which transmits beacon.
  2. Each node records RSS values.
  3. Repeat for all APs.
- **Afterwards: Update when idle**
  Whenever node is idle, passively scan *all* channels periodically, measuring RSS from other nodes by overhearing packets being transmitted. (Extended method of 802.11K.)

Interference Model

- $I_{B \to A}$:= interference of B on A if they use the same channel.
- $I_{B \to A} \propto RSS_{B \to A}$, the strength of the interference signal.
- $I_{B \to A} \propto L_B$, the probability that B is transmitting.
- $I_{B \to A} \propto L_A$, the probability that A is receiving.

Hence,

$$\boxed{I_{B \to A} := RSS_{B \to A} \cdot L_B \cdot L_A}$$

Define the interference between A and B as

$$\boxed{I(A,B) := I_{B \to A} + I_{A \to B}.}$$

Additively extend to interference between two cells:

$$\boxed{I(C_1, C_2) := \sum_{u \in C_1, v \in C_2} I(u,v)}$$

Interference Model

- $I_{B \to A} :=$ interference of B on A if they use the same channel.
- $I_{B \to A} \propto RSS_{B \to A}$, the strength of the interference signal.
- $I_{B \to A} \propto L_B$, the probability that B is transmitting.
- $I_{B \to A} \propto L_A$, the probability that A is receiving.

Hence,

$$I_{B \to A} := RSS_{B \to A} \cdot L_B \cdot L_A$$

Define the interference between A and B as

$$I(A, B) := I_{B \to A} + I_{A \to B}.$$

Additively extend to interference between two cells:

$$I(C_1, C_2) := \sum_{u \in C_1, v \in C_2} I(u, v)$$

Interference Model

- $I_{B \to A} :=$ interference of B on A if they use the same channel.
- $I_{B \to A} \propto RSS_{B \to A}$, the strength of the interference signal.
- $I_{B \to A} \propto L_B$, the probability that B is transmitting.
- $I_{B \to A} \propto L_A$, the probability that A is receiving.

Hence,

$$\boxed{I_{B \to A} := RSS_{B \to A} \cdot L_B \cdot L_A}$$

Define the interference between A and B as

$$\boxed{I(A, B) := I_{B \to A} + I_{A \to B}.}$$

Additively extend to interference between two cells:

$$\boxed{I(C_1, C_2) := \sum_{u \in C_1, v \in C_2} I(u, v)}$$

Interference Model

- $I_{B \to A}$:= interference of B on A if they use the same channel.
- $I_{B \to A} \propto RSS_{B \to A}$, the strength of the interference signal.
- $I_{B \to A} \propto L_B$, the probability that B is transmitting.
- $I_{B \to A} \propto L_A$, the probability that A is receiving.

Hence,

$$\boxed{I_{B \to A} := RSS_{B \to A} \cdot L_B \cdot L_A}$$

Define the interference between A and B as

$$\boxed{I(A, B) := I_{B \to A} + I_{A \to B}.}$$

Additively extend to interference between two cells:

$$\boxed{I(C_1, C_2) := \sum_{u \in C_1, v \in C_2} I(u, v)}$$

Interference Model

- $I_{B \to A} :=$ interference of B on A if they use the same channel.
- $I_{B \to A} \propto RSS_{B \to A}$, the strength of the interference signal.
- $I_{B \to A} \propto L_B$, the probability that B is transmitting.
- $I_{B \to A} \propto L_A$, the probability that A is receiving.

Hence,

$$\boxed{I_{B \to A} := RSS_{B \to A} \cdot L_B \cdot L_A}$$

Define the interference between A and B as

$$\boxed{I(A,B) := I_{B \to A} + I_{A \to B}.}$$

Additively extend to interference between two cells:

$$\boxed{I(C_1, C_2) := \sum_{u \in C_1, v \in C_2} I(u,v)}$$

Interference Model

- $I_{B \to A} :=$ interference of B on A if they use the same channel.
- $I_{B \to A} \propto RSS_{B \to A}$, the strength of the interference signal.
- $I_{B \to A} \propto L_B$, the probability that B is transmitting.
- $I_{B \to A} \propto L_A$, the probability that A is receiving.

Hence,

$$\boxed{I_{B \to A} := RSS_{B \to A} \cdot L_B \cdot L_A}$$

Define the interference between A and B as

$$\boxed{I(A, B) := I_{B \to A} + I_{A \to B}.}$$

Additively extend to interference between two cells:

$$\boxed{I(C_1, C_2) := \sum_{u \in C_1, v \in C_2} I(u, v)}$$

Interference Model

- $I_{B\to A} :=$ interference of B on A if they use the same channel.
- $I_{B\to A} \propto RSS_{B\to A}$, the strength of the interference signal.
- $I_{B\to A} \propto L_B$, the probability that B is transmitting.
- $I_{B\to A} \propto L_A$, the probability that A is receiving.

Hence,

$$\boxed{I_{B\to A} := RSS_{B\to A} \cdot L_B \cdot L_A}$$

Define the interference between A and B as

$$\boxed{I(A,B) := I_{B\to A} + I_{A\to B}.}$$

Additively extend to interference between two cells:

$$\boxed{I(C_1, C_2) := \sum_{u\in C_1, v\in C_2} I(u,v)}$$

Min k-Partition

- Make an interference graph.

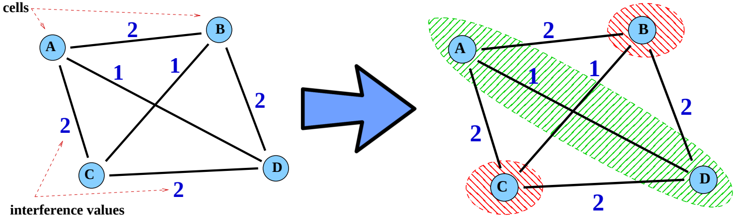

- **Min k-Partition:**
  Partition graph into k subsets to minimize sum weight on edges starting and ending within the same subset.
- Equivalent to **Max k-Cut**: Partition graph into k subsets to maximize sum weight on edges going across subsets.

Min k-Partition: NP-Hard ILP → SDP Relaxation

minimize (over Z) $\sum_{1 \leq i < j \leq N} W_{ij} Z_{ij}$

subject to

$$Z_{ij} = \begin{cases} 1 & \text{if } C_i \text{ and } C_j \text{ use same channel} \\ 0 & \text{else} \end{cases}$$

$$Z_{ih} + Z_{hj} - Z_{ij} \leq 1 \quad \forall h, i, j \in [N]$$

$$\sum_{i,j \in Q} Z_{ij} \geq 1 \quad \forall Q \subseteq [N] \text{ where } |Q| = k + 1$$

SDP relaxation, dropping most constraints:

minimize (over X)	$\mathbf{tr}\,(W \cdot X)$
subject to	$X_{ii} = 1, \quad \frac{-1}{k-1} \leq X_{ij}, \quad X \succeq 0$

Frieze and Jerrum Technique

$$\begin{array}{ll} \text{minimize (over } X) & \mathbf{tr}\,(W \cdot X) \\ \text{subject to} & X_{ii} = 1, \quad \frac{-1}{k-1} \leq X_{ij}, \quad X \succeq 0 \end{array}$$

Hyperplane rounding:

1. Solve for $X \in \mathbf{R}^{n \times n}$.
2. Cholesky: $X = V^T V$. Columns: $V = [v_1 \mid \ldots \mid v_n]$. Each column "corresponds" to a cell.
3. Generate k random vectors $r_1, \ldots, r_k$.
4. Assign channel to cell $i \in [N]$ according to which random vector v_i is closest to.

DCA Outline

Dynamic Channel Assignment (DCA)

During every T second time-interval:

1. Report traffic + RSS to server.
2. If network has "sufficiently changed":
    1. Compute co-channel interference between any two cells.
    2. Compute channel reassignments to minimize sum interference (SDP Relaxation).
    3. APs change channels. Broadcast change to stations. (1-2 ms)

Outline

Testbed

- 500 m^2 office floor testbed
- 13 APs, 22 wireless stations (PCs), 10 wired servers
- Maxed transmit power
- Max Throughput: 38 Mbps
- RTS/CTS off, 802.11a

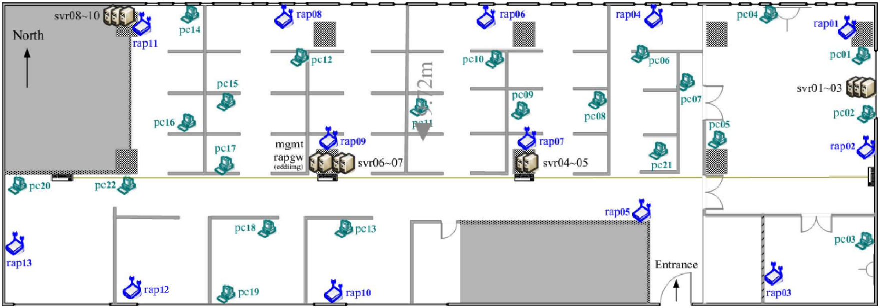

We will compare Dynamic Channel Assignment (DCA) with

Static Channel Assignment (SCA):

1. Ignore all traffic and station-side information.
2. Only use RSS measurements between APs as edge weights.
3. Run Min k-Partition SDP.
4. Assign channels once.

Channel Assignments Visualized

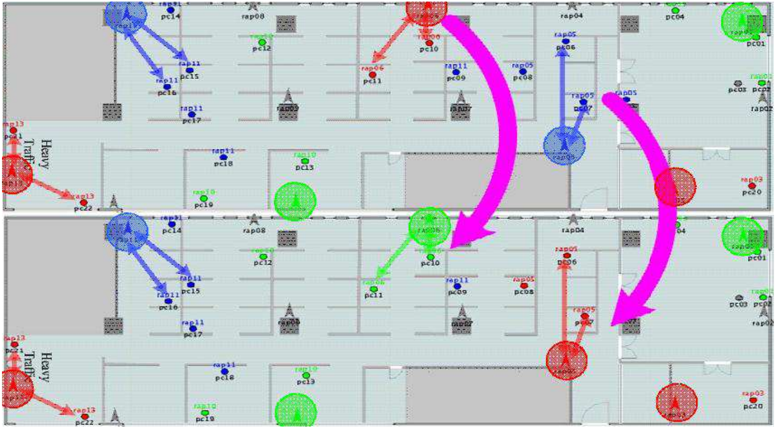

Channel Reuse

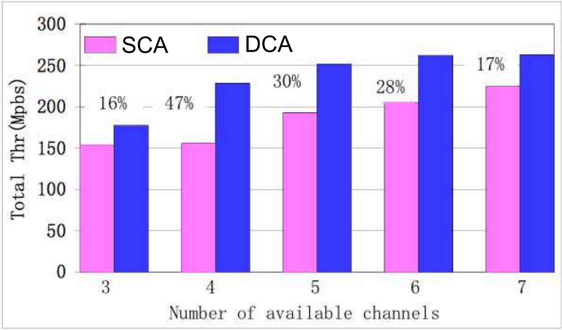

- 7 APs with heavy flow (38 Mbps); rest with light flows.
- Total throughput upper bound: $7 \times 38 = 266 Mbps$.
- For each number of available channels, use static channel assignment (SCA) for 10 minutes, followed by DCA. Repeat 5 times for each k. Record average throughput gain.
- $k = 5$: 266 Mbps nearly met. 2 channels reused; 7 non-interfering flows.

Throughput vs. Time

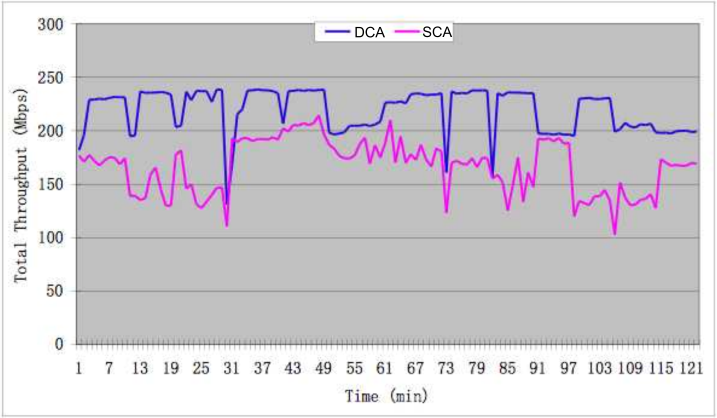

- Every 10 minutes: 6 out of 13 APs randomly chosen to have heavy traffic. Other APs have light flows.
- Total throughput increases by 40% on average.

Outline

Conclusion

- Efficient dynamic channel assignment algorithm for WLANs.
  - Intuitive interference model that makes use of traffic awareness and RSS readings.
  - SDP Relaxation of Min k-Partition problem
- Testbed experiments show throughput increase of $\sim 40\%$.

Thanks for listening!

Comments/Questions:

William Wu

willywutang@gmail.com

References

- Reis et al, Measurement-Based Models of Delivery and Interference in Static Wireless Networks, SIGCOMM '06.
- Goemans and Williamson. Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems Using Semidefinite Programming, J.ACM, 42, 1115-1145, 1995.
- A. Frieze and M. Jerrum. Improved approximation algorithms for MAX k-CUT and MAX BISECTION. Algoritmica, 18,1997.
- 802.11k working group, Wireless LAN MAC and PHY specifications: Radio resource measurement, 2005, IEEE Draft Standard.
- E. Rozner, Y. Mehjta, A. Akella, L. Qiu, Traffic-aware channel assignment in wireless LANs, in Mobicom 2006.
- Mishra et al, Weighted coloring based channel assignment for WLANs, Mobile Computer Communications Review (MC2R), vol. 9, no. 3, 2005.
- Mishra et al, A Client-Driven Approach for Channel Management in WLANs, IEEE Infocom 2006.

Backup Slide: Industry

- 802.11k: Use RRM information to optimize network.
- Cisco: "The Cisco Wireless LAN Controller combines the RF characteristic information with intelligent algorithms to make systemwide decisions ... dynamically control access point transmit power based on real-time WLAN conditions. "
- Tropos: "Metro-optimized dynamic channel assignment, an enhancement to the company's Predictive Wireless Routing Protocol (PWRP), intelligently adjusts the channel plan throughout the network in response to environmental and RF changes."

Backup Slide: Min k-Partition: Integer LP

minimize (over Z) $$\sum_{1 \leq i < j \leq N} W_{ij} Z_{ij}$$

subject to $$Z_{ij} = \begin{cases} 1 & \text{if } C_i \text{ and } C_j \text{ use same channel} \\ 0 & \text{else} \end{cases}$$

$$Z_{ih} + Z_{hj} - Z_{ij} \leq 1 \quad \forall h, i, j \in [N]$$

$$\sum_{i,j \in Q} Z_{ij} \geq 1 \quad \forall Q \subseteq [N] \text{ where } |Q| = k + 1$$

Backup Slide: Min k-Partition: Integer LP

minimize (over Z) $$\sum_{1 \le i < j \le N} W_{ij} Z_{ij}$$

subject to $$Z_{ij} = \begin{cases} 1 & \text{if } C_i \text{ and } C_j \text{ use same channel} \\ 0 & \text{else} \end{cases}$$

$$Z_{ih} + Z_{hj} - Z_{ij} \le 1 \quad \forall h, i, j \in [N]$$

$$\sum_{i,j \in Q} Z_{ij} \ge 1 \quad \forall Q \subseteq [N] \text{ where } |Q| = k + 1$$

ILP relaxation yields ...

$$\text{minimize (over } Z) \quad \sum_{1 \le i < j \le N} W_{ij} Z_{ij}$$

$$\text{subject to} \quad 0 \le Z_{ij} \le 1$$

$$Z_{ih} + Z_{hj} - Z_{ij} \le 1 \quad \forall h, i, j \in [N]$$

$$\sum_{i,j \in Q} Z_{ij} \ge 1 \quad \forall Q \subseteq [N] \text{ where } |Q| = k + 1$$

Backup Slide: Min k-Partition SDP

$$\begin{array}{ll} \text{minimize (over } X) & \displaystyle\sum_{1 \le i < j \le N} W_{ij} \frac{(k-1)X_{ij} + 1}{k} \\ \text{subject to} & X_{ij} = \begin{cases} 1 & C_i \text{ and } C_j \text{ use same channel} \\ \frac{-1}{k-1} & \text{else} \end{cases} \\ & X_{ii} = 1, \quad X \succeq 0 \end{array}$$

... simplification and relaxation ...

$$\boxed{\begin{array}{ll} \text{minimize (over } X) & \mathbf{tr}\,(W \cdot X) \\ \text{subject to} & X_{ii} = 1, \quad \frac{-1}{k-1} \le X_{ij}, \quad X \succeq 0 \end{array}}$$

Backup Slide: Intuition: Max 2-cut

$$\begin{array}{ll} \text{maximize (over } y_j) & \frac{1}{2}\sum_{i<j} w_{ij}(1 - y_i y_j) \\ \text{subject to} & y_j \in \{-1, 1\} \end{array}$$

$$\begin{array}{ll} \text{maximize (over } v_j) & \frac{1}{2}\sum_{i<j} w_{ij}(1 - v_i \cdot v_j) \\ \text{subject to} & v_j \in S_{n-1} := \{x \in \mathbf{R}^n : |x| = 1\} \end{array}$$

$$\begin{array}{ll} \text{maximize (over } v_j) & \frac{1}{2}\sum_{i<j} w_{ij}(1 - Y_{ij}) \\ \text{subject to} & Y_{ii} = 1, \quad Y \succ 0 \end{array}$$

Heuristic: Choose random hyperplane through origin, and partition vectors v_i (and hence the vertex set V) according to which side of the hyperplane they fall.