Cross-layer Wireless Networking:

Complexity, Approximation, and Opportunities for SP Research

Nikos Sidiropoulos

Dept. ECE, TU Crete

IEEE SPAWC, 23/6/2010

ACKs

Students

- Evaggelia Matskani
- Ioannis Mitliagkas
- Dimitris Evaggelinakis

Colleagues

- Zhi-Quan Luo
- Ananthram Swami
- Leandros Tassiulas

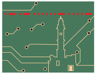

Outline

1. Power control
2. Joint power and admission control
3. Multi-hop routing
4. Opportunities for SP research

Where PHY and NET first met

- Co-channel users/links
- Frequency reuse or CDMA (PCS)
- Cellular voice
- SINR constraints
- Power control
- Various contexts:
  - PCS, UMTS-LTE
  - ad-hoc
  - peer-to-peer
  - cognitive underlay

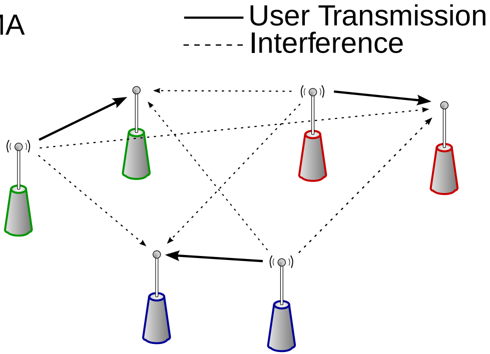

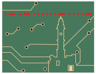

Properties & some history

- Linear Programming (LP)
- But wait, there's more:
  - Feasibility - spectral radius (Perron-Frobenius)
  - Simple distributed algorithm (Foschini)
  - Well-developed theory
- Foschini, Zander, Yates, Bambos, ...
- Many flavors

Power Control

$$\min_{\{p_k \in \mathbb{R}_+\}_{k=1}^{K}} \sum_{k=1}^{K} p_k$$

$$p_k \leq P_k^{MAX}, \ \forall k$$

$$\frac{G_{kk} p_k}{\sum_{l=1,\, l \neq k}^{K} G_{lk} p_l + \sigma_k^2} \geq c_k, \ \forall k$$

Real problem is much tougher

- Often infeasible $\rightarrow$ admission control
- Admission and power tightly coupled
- Jointly pick users and powers to
  - Max # of users admitted
  - Under SINR, power constraints
  - Min total power
- Combinatorial?
  - Andersin, Rosberg, Zander '96: *contained* in NP-hard
  - ... vs. *contains* NP-hard
- Gradual removals (Zander *et al*)
- Active link protection (Bambos *et al*)

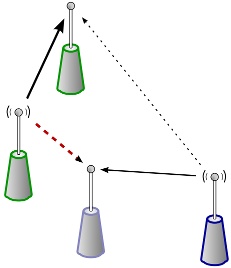

Joint power and admission control

Stage 1: Admission Control

- Maximal subset S_0, $\mathbf{p}(S_0)$
- Satisfying
  - Max power
  - Min SINR

Stage 2: Power Control

- Minimize total power in S_0
- Satisfying
  - Max power
  - Min SINR

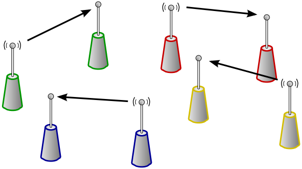

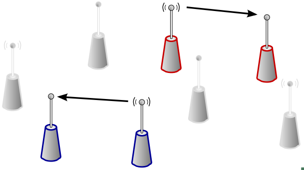

Joint power and admission control

Stage 1: Admission Control

$$S_0 = \arg\max_{S \subseteq \{1,\ldots,K\}, \{p_k \in \mathbb{R}_+\}_{k=1}^{K}} |S|$$

s.t. $\forall k \in S$

$$p_k \leq P_k^{MAX}$$

$$\frac{G_{kk} p_k}{\sum_{l \in S,\ l \neq k} G_{lk} p_l + \sigma_k^2} \geq c_k$$

Stage 2: Power Control

$$\min_{\{p_k \in \mathbb{R}_+\}_{k \in S_0}} \sum_{k \in S_0} p_k$$

s.t. $\forall k \in S_0$

$$p_k \leq P_k^{MAX}$$

$$\frac{G_{kk} p_k}{\sum_{l \in S_0,\ l \neq k} G_{lk} p_l + \sigma_k^2} \geq c_k$$

Complexity

Start from any graph

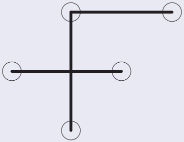

Construct instance of AC

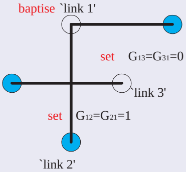

Special instance of AC

$$S_0 = \underset{S\subseteq\{1,\ldots,K\},\{p_k\in[0,1]\}_{k=1}^{K}}{\arg\max} |S|$$

$$\text{s.t.} \frac{p_k}{\sum_{l\in S,\ l\neq k} G_{lk}p_l + 1} \geq 1,\ \forall k \in S$$

Maximal independent set

- indep (cyan) $\Rightarrow$ feas ✓
- feas $\Rightarrow \frac{1}{0+1} \Rightarrow$ indep ✓

A ruler analogy

- Introduce binary scheduling variables
- Formulate as single stage problem

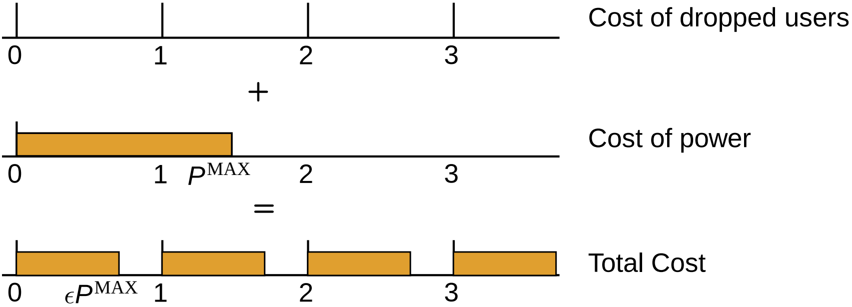

- Fully prioritizes user admission over power minimization [MatSidLuoTas:07]; [MitSidSwa:08]

Single stage reformulation

- Binary scheduling variables $s_k = \{0, 1\}$ (0 for admitted)
- Auxiliary constants ϵ and δ_k

$$\min_{\{p_k \in \mathbb{R}_+, s_k \in \{0,1\}\}_{k=1}^K} \epsilon \sum_{k=1}^{K} p_k + (1-\epsilon) \sum_{k=1}^{K} s_k$$

$$\text{s.t. } p_k \leq P_k^{MAX}, \quad \forall k \in \{1, \ldots, K\}$$

$$\frac{G_{kk} p_k + \delta_k^{-1} s_k}{\sum_{l=1,\, l \neq k}^{K} G_{lk} p_l + \sigma_k^2} \geq c_k, \ \forall k \in \{1, \ldots, K\}$$

- Proven equivalent to two-stage optimization for suitable ϵ, δ_k

Convex relaxation

- Problem is non-convex (binary scheduling variables)
- Convex relaxation? - Lagrange bi-dual
- Lagrange bi-dual $\Longleftrightarrow$ binary s_k $\rightarrow$ continuous s_k

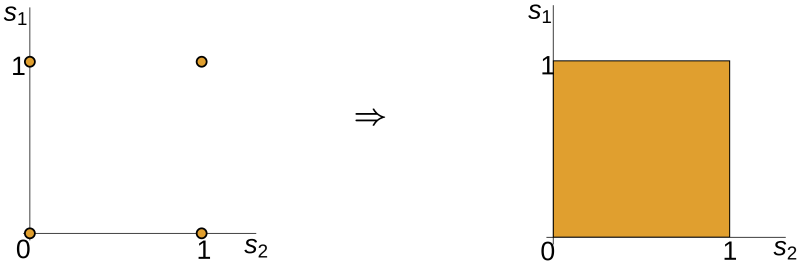

Convex relaxation

- Convex (bi-dual) relaxation $\Longleftrightarrow$

$$\min_{\{p_k \in \mathbb{R}_+, s_k \in \mathbb{R}\}_{k=1}^{K}} \epsilon \sum_{k=1}^{K} p_k + (1-\epsilon) \sum_{k=1}^{K} s_k$$

$$\text{s.t. } p_k \leq P_k^{MAX}, \quad \forall k \in \{1, \ldots, K\}$$

$$\frac{G_{kk} p_k + \delta_k^{-1} s_k}{\sum_{l=1,\, l \neq k}^{K} G_{lk} p_l \;+ \sigma_k^2} \geq c_k, \quad \forall k \in \{1, \ldots, K\}$$

$$0 \leq s_k \leq 1, \quad \forall k \in \{1, \ldots, K\}$$

Approximation algorithm

Algorithm

Linear Programming Deflation

1. $\mathcal{U} \leftarrow \{1, ..., K\}$
2. *Solve the relaxed problem*
3. *If all links attain target SINR*
   - *terminate*

   Else
   - *use heuristic to choose a link*
   - *remove it from* $\mathcal{U}$
   - *go to Step 2.*

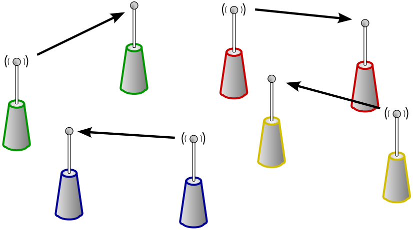

Approximation algorithm

Algorithm

Linear Programming Deflation

1. $\mathcal{U} \leftarrow \{1, ..., K\}$
2. *Solve the relaxed problem*
3. *If all links attain target SINR*
   - *terminate*

   Else
   - *use heuristic to choose a link*
   - *remove it from* $\mathcal{U}$
   - *go to Step 2.*

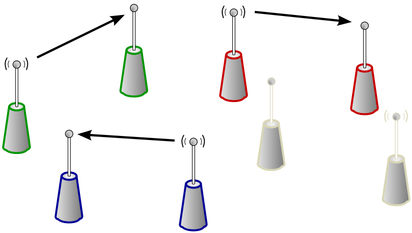

Approximation algorithm

Algorithm

Linear Programming Deflation

1. $\mathcal{U} \leftarrow \{1, ..., K\}$
2. *Solve the relaxed problem*
3. *If all links attain target SINR*
   - *terminate*

   Else
   - *use heuristic to choose a link*
   - *remove it from* $\mathcal{U}$
   - *go to Step 2.*

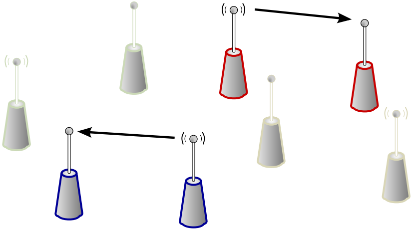

Options

Practical implementation

- Distributed ✓
  - Dual decomposition (slow)
  - Consensus-on-max (deflation)
- Robust ✓
  - imperfect CSI G_{lk}
- Polynomial complexity, overhead

Optimal solution?

- LPD $\longrightarrow$ lower & upper bounds on opt cost
- Branch & Bound w/ LPD ✓
  - Implicit search - pruning
  - Complexity $\ll$ ENUM
  - Still exp in w-c
- "Sphere decoding"

Simulations - admission performance

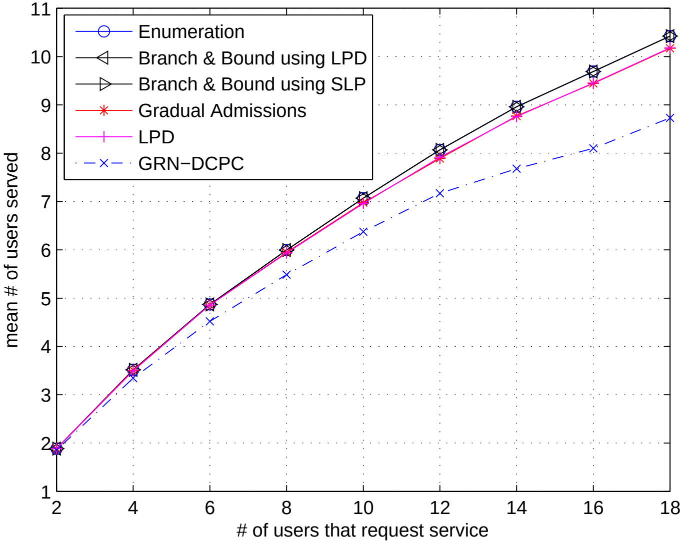

Simulations - average complexity

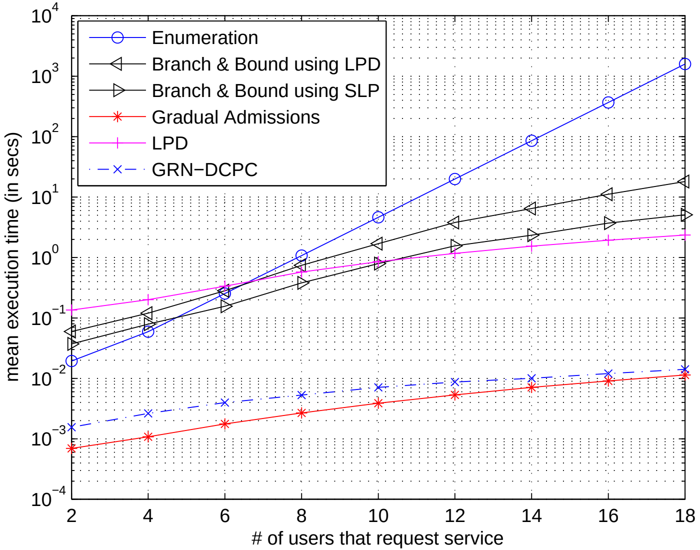

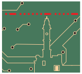

Simulations - worst-case complexity

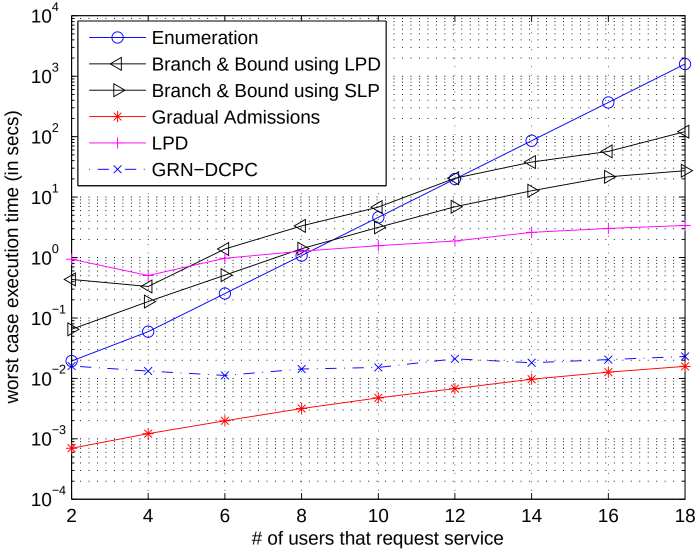

Multi-hop routing: shortest path

Connectivity

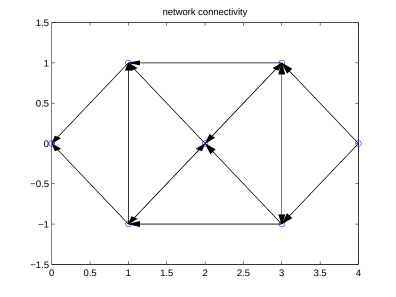

- weights $\sim$ load, delay, “cost”

Shortest paths

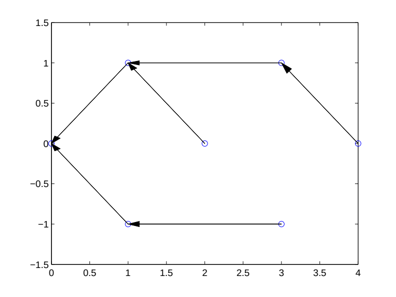

- for all weights equal

Shortest path vs. dynamic back-pressure

SP

- DP: BF, FW, ...
- Distributed ✓
- Must know arrival rate
- Quasi-static, very slow to adapt to
  - changing arrivals/load
  - availability/failure
  - fading/interference patterns
- Claim: Low delay (shortest path)
- ... but only at low system loads

BP [Tassiulas ’92]

- One-hop differential backlog
- Distributed ✓ Lightweight ✓
- Auto-adapts ✓
- Highly dynamic, agile ✓
- Claim: maximal stable throughput (all paths)
- ... but delay can be large - U(load), $\emptyset \longrightarrow$ rand walk!

Back-pressure routing

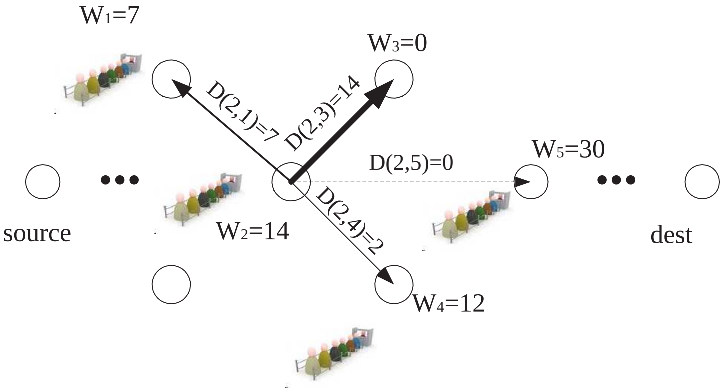

- Favors links with low back-pressure (hence name)
- Backtracking / looping possible!
- Local communication, trivial computation

Back-pressure routing

- Multiple destinations, commodities?
  - multiple queues per node
  - (max diff backlog) winner-takes-all per link
- Wireline: local communication, trivial computation
- Wireless?
- Broadcast medium: interference
- Link rates depend on transmission scheduling, power of other links
- Globalization - but also opportunity to shape-up playing field ...
- ... through appropriate scheduling, power control

Back-pressure power control

SINR

$$\gamma_\ell = \frac{G_{\ell\ell} p_\ell}{\sum_{k \in \mathcal{L}, k \neq \ell} G_{k\ell} p_k + V_\ell}$$

Link capacity

$$c_\ell = \log(1 + \gamma_\ell)$$

Diff backlog link $\ell = (i \rightarrow j)$ @ time t

$$D_\ell(t) := \max\left\{0, W_i(t) - W_j(t)\right\}$$

BPPC

$$\max_{\{p_\ell\}_{\ell \in \mathcal{L}}} \sum_{\ell \in \mathcal{L}} D_\ell(t) c_\ell$$

$$\text{s.t.} \quad 0 \leq \sum_{\ell : \mathrm{Tx}(\ell) = i} p_\ell \leq P_i, \forall i \in \mathcal{N}$$

$$p_\ell \leq P^{(\ell)}, \ell \in \mathcal{L}$$

Back-pressure power control

BPPC

$$\max_{\{p_\ell\}_{\ell\in\mathcal{L}}} \sum_{\ell\in\mathcal{L}} D_\ell(t) c_\ell$$

$$\text{s.t.} \quad 0 \leq \sum_{\ell:\mathrm{Tx}(\ell)=i} p_\ell \leq P_i, \forall i \in \mathcal{N}$$

$$p_\ell \leq P^{(\ell)}, \ell \in \mathcal{L}$$

Link activation / scheduling:

$$p_\ell \in \left\{0, P^{(\ell)}\right\}, \ell \in \mathcal{L}$$

[Tassiulas *et al*, ’92 $\longrightarrow$]

- Max stable throughput ✓
- Backbone behind modern NUM
- Core problem in wireless networking
- Countable control actions: random, adopt if > current
- Still throughput-opt! [Tass’98] - but D $\uparrow$
- Continuous opt vars?

Back-pressure power control

- Non-convex due to $c_\ell \sim \log(1+\gamma_\ell)$ - diff of concave
- At high SINR γ_ℓ, $1+\gamma_\ell \cong \gamma_\ell$
- $c_\ell \geq \log(\gamma_\ell)$ always
- Tempting ...
- Giannoulis, Tsoukatos, Tassiulas, ICC'06
- Gradient projection, best response

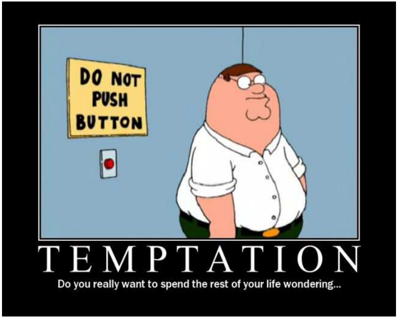

Beware!

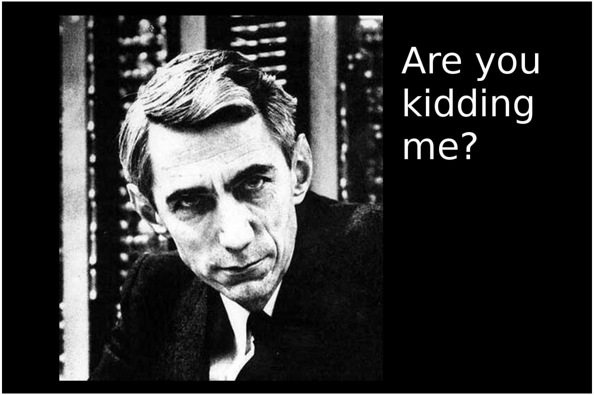

Reminiscent of ...

DSL: sum-rate maximization

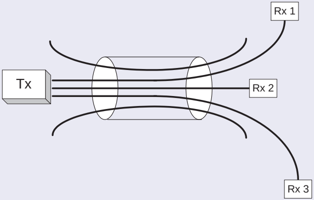

BPPC

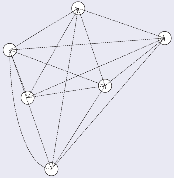

Single-hop DSL

- Listen-while-talk ✓
- Dedicated (Tx,Rx)
- Free choice of $G_{k,\ell}$'s
- NP-hard [Luo, Zhang]

Multi-hop network

- No listen-while-talk X
- Shared Tx, Rx $\Rightarrow$
- Restricted $G_{k,\ell}$'s
- NP-hard?

Peel off

Generic backlogs

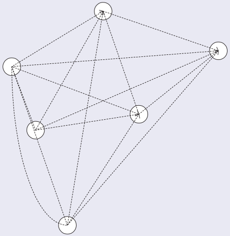

Choosing backlogs

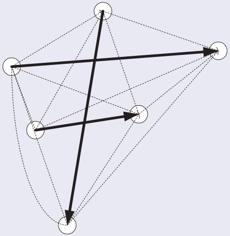

DSL $\longrightarrow$ Multi-hop network optimization

- Backlog reduction $\longrightarrow$ BPPC contains DSL $\longrightarrow$ also NP-hard
- Can reuse tools from DSL
- In particular, lower approximation algorithms:
  - High SINR $\longrightarrow$ Geometric Programming
  - Successive approximation from below: SCALE [Papandriopoulos and Evans, 2006]
  - Uses

$$\alpha \log(z) + \beta \leq \log(1+z) \text{ for } \begin{cases} \alpha = \frac{z_o}{1+z_o} \\ \beta = \log(1+z_o) - \frac{z_o}{1+z_o}\log(z_o) \end{cases}$$

    tight at z_o; $\longrightarrow \log(z) \leq \log(1+z)$ as $z_o \longrightarrow \infty$
  - Start from high SINR, tighten bound at interim solution
  - Majorization (actually, minorization)

Key difference with DSL

- BPPC problem must be solved repeatedly for every slot
- Batch algorithms: prohibitive complexity
- Need adaptive, lightweight solutions (to the extent possible)
- Built custom interior point algorithms
- Normally, one would init using solution of previous slot; take refinement step
- Doesn’t work ...
- Why?

Proper warm-start

- No listen-while-talk, shared Tx/Rx
- Push-pull ‘wave’ propagation
- Solution from previous slot very different from one for present slot
- Even going back a few slots
- Quasi-periodic behavior emerges
- Idea: hold record of solutions for W previous slots. $W >$ upper bound on period
- W evaluations of present objective function (cheap!)
- Pick the best to warm-start present slot
- Needs few IP steps to converge

Quality of approximation?

- Max lower bound $\Rightarrow$ link rates attainable
- Sims indicate solutions far outperform prior art in networking in terms of key network metrics: throughput, delay, stability margin
- OK, but upper bound? Normally, dual problem
- Here computing dual function is also NP-hard :-([Tx: Tom Luo]
- Resort to Yu and Lui '06, originally for spectrum balancing in DSL
- Yields approximate solution of dual problem - approximate upper bound
- When properly tuned ... can be very slow ...
- Sanity check / gauge

Simulation setup

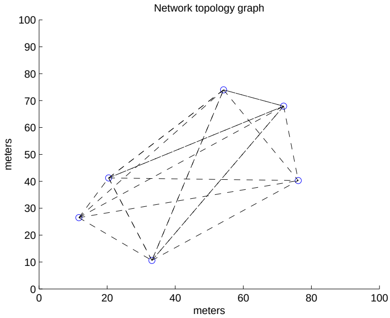

- $N = 6$ nodes, low-left = s, top-right = d, $L = 21$ links
- $G_{\ell,k} \sim 1/d^4$, $G = 128$, no-listen-while-talk $1/\text{eps}$
- $V_\ell = 10^{-12}$, $P^{(\ell)} = 5$, $\forall \ell$
- Deterministic (periodic) arrivals

High SINR

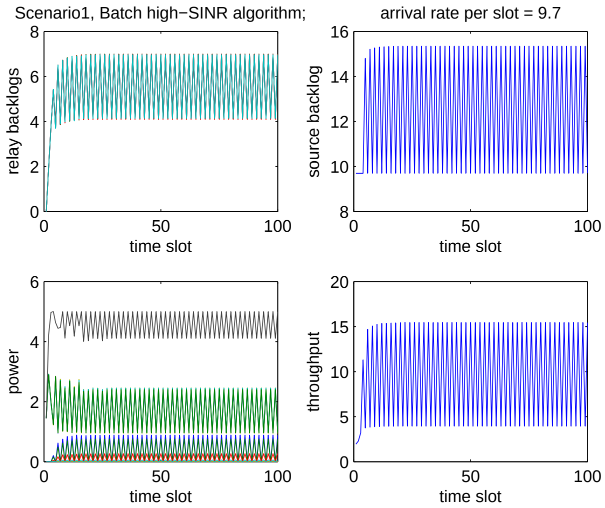
Scenario1, Batch high−SINR algorithm;
arrival rate per slot = 9.7
relay backlogs
source backlog
power
throughput
time slot
time slot
time slot
time slot

High SINR

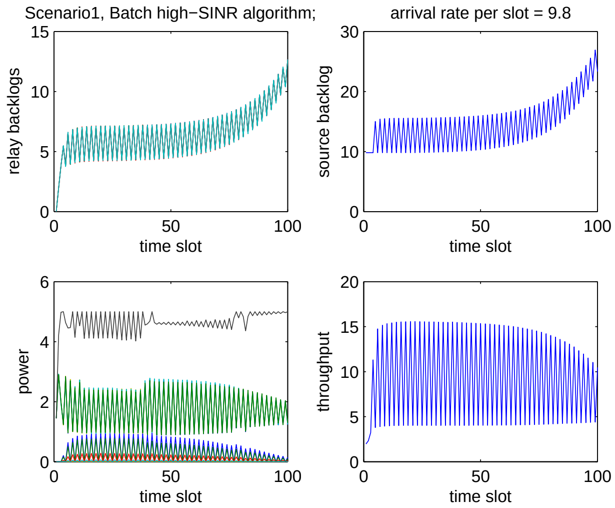
Scenario1, Batch high−SINR algorithm;
arrival rate per slot = 9.8
relay backlogs
source backlog
power
throughput
time slot
0
5
10
15
20
30
4
6
2
50
100

Successive Approximation

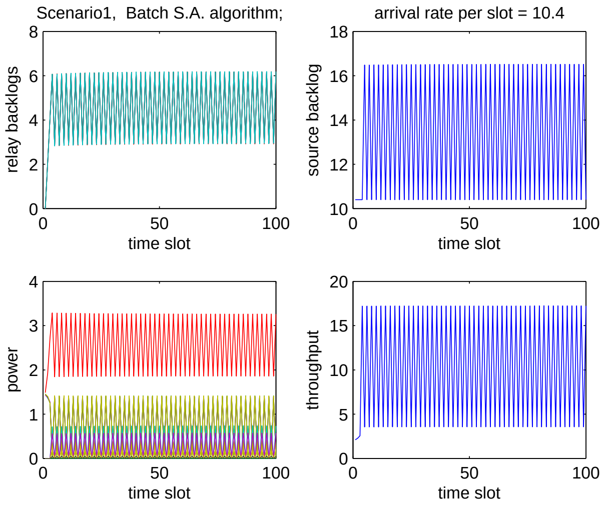

Successive Approximation

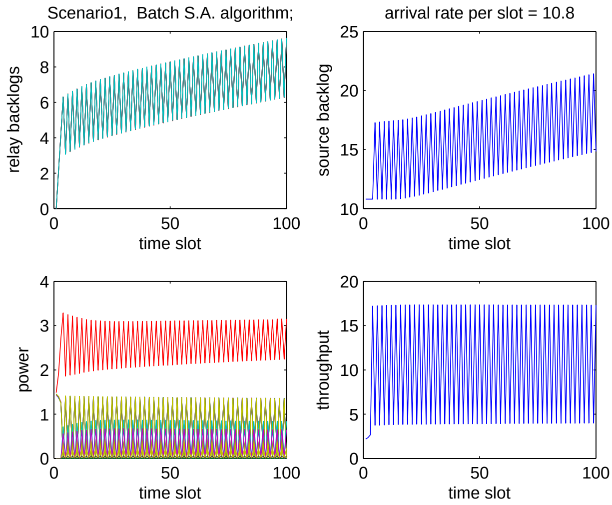

Best Response

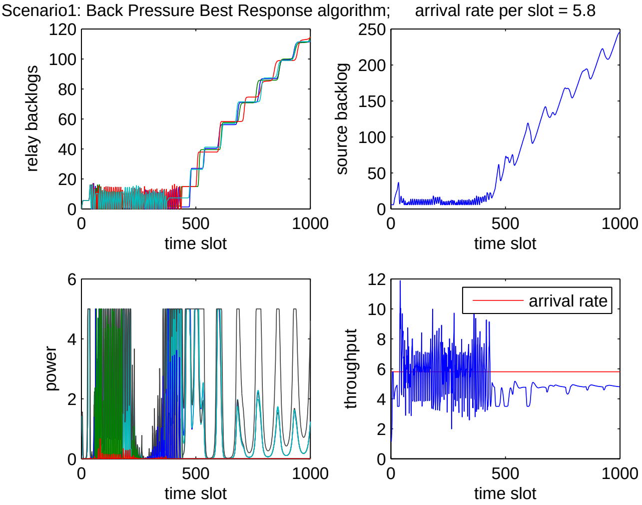

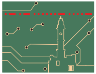

Gap to optimal: ISB approximation

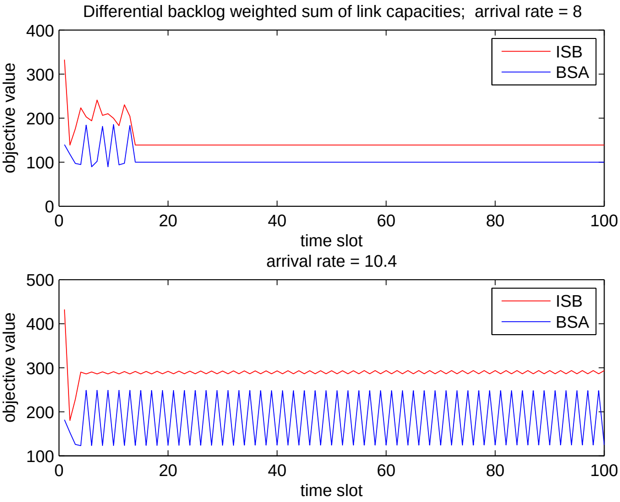

Opportunities for SP research

Looking ahead

- Distributed BPPC
- Robustness (imperfect / outdated CSI)
- LMS-like? - Ribeiro, Gatsis+Giannakis
- MIMO nodes - beamforming? precoding? spatial MUX?
- Other modalities - multicasting?
- All NP-hard, need effective approximation

Paradigm shift

- Network coding?
- Cooperation among nodes?