Information Theory for Mobile Ad-Hoc Networks (ITMANET): *The FLoWS Project*

Thrust 3 Intro:
Application Metrics and Network Performance

Asu Ozdaglar

Joint with D. Shah

Optimizing Application and Network Performance

- Objective:
  - Developing a framework for optimizing heterogeneous and dynamically varying application metrics and ensuring efficient operation of large-scale decentralized networks with uncertain capabilities and capacities
  - Providing an interface between application metrics and network capabilities
    - Focus on a direct involvement of the application in the network, defining services in terms of the function required rather than rates or other proxies
- Application and Network Metrics: utility functions of users-applications, distortion, delay, network stability, energy…
- We envision a universal algorithmic architecture:
  - Capable of balancing (or trading off) application requirements and network resources
  - Adaptable to variations on the network and user side
  - Operable in a decentralized manner, scalable
  - Robust against non-cooperative behavior

Algorithmic Architecture for Optimizing Application and Network Performance

Prior Work

- Decoupled/layered approach to resource allocation
  - Highly suboptimal and inefficient
- More recent trend:
  - Formulate resource allocation problem as one optimization problem and use decompositions based on separable structure
  - This approach fails for wireless networks due to:
    - Need for distributed asynchronous implementations
    - Externalities/couplings that disturb separable structure
    - Stochastic elements
- No analysis of robustness against dynamic changes and noncooperative behavior and competition

Intellectual Tools and Focus Areas

- Optimization and Control Theory
  - Decentralized algorithms robust against variations in network topology, channel characteristics, and capacities
  - Ensuring rapid convergence
  - Optimization for heterogeneous preferences

- Performance (stability) analysis of network algorithms
  - At micro-level: understanding queuing dynamics
  - At macro-level: understanding effect on flow-level network behavior

- Game Theory
  - Dealing with noncooperative strategic users
  - Dynamics and equilibrium

Individual PI Presentations

- Shah, "Fundamental Performance Limits and Reality"
- Meyn, "Optimizing MaxWeight for Resource Allocation"
- Boyd, "Large Scale Network Utility Maximization"
- Ozdaglar, "Distributed Asynchronous Optimization Methods for General Performance Metrics"
- Johari, "Incomplete Information, Dynamics, and Wireless Games"

Wireless networks: Algorithmic trade-off between Throughput and Delay

FLoWS

STATUS QUO

Among two important performance metrics of wireless networks, throughput and delay, only throughput is well-understood in terms of fundamental limits and algorithm design.

Delay is far from being well-understood.

NEW INSIGHTS

1. Arbitrary networks: High-throughput is *easy*, low delay is *hard*.

2. Practical networks: distributed high-througput low delay is *possible*.

ACHIEVEMENT DESCRIPTION

MAIN RESULT:

1. High-throughput low delay algorithm for arbitrary wireless network is computationally intractable.

2. Wireless networks deployed in geographic area (in arbitrary manner) have high-throughput and low-delay algorithm distributed algorithms for scheduling and cross-layer design.

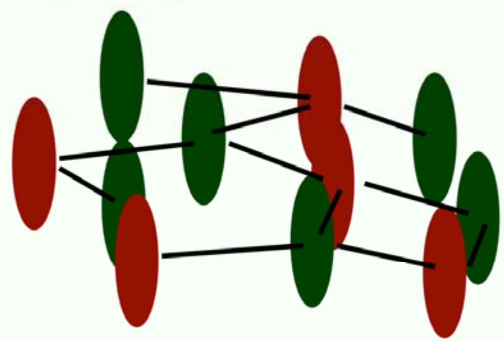

HOW IT WORKS:

1. Intractability follows by identifying computational hardness in scheduling through a novel equivalence relation.

2. Geometry in wireless networks allows for simple, high-performance algorithm design.

ASSUMPTIONS AND LIMITATIONS:

1. Wireless network with interference model.

END-OF-PHASE GOAL

1. Computational intractability of high throughput, low delay algorithm for wireless network under SINR model.

2. Simple algorithms for practical networks under SINR model.

COMMUNITY CHALLENGE

Computational intractability of information theoretic capacity achieving codes for wireless networks.

Algorithmic limitations for wireless network

Status quo

- Primary performance metric in a wireless network
  - ○ Throughput and delay
  - ○ Necessary for quality-of-service guarantee, buffer-design, etc.
  - ○ Further, algorithm should be implementable (distributed)
- However, thus far most of the work has concentrated on designing throughput optimal algorithms
  - ○ Low delay algorithm design is a lot harder
  - ○ An *analogy*: being ahead of all in a marathon throughout the race(low delay) versus completing the race first (high throughput)
- One of the main reason for such status
  - ○ Lack of good tools for delay analysis
  - ○ Hence lack of insight about what causes high delay
  - ○ As well as inability to understand finer throughput delay tradeoff

Summary of Results

- First, we establish that

  ○ It is possible to have very simple, distributed throughput optimal algorithm for any network

  → throughput is easy

- To understand interaction with throughput and delay

  ○ We introduce new tools from computational complexity

  ○ We establish computational impossibility of designing high throughput, low delay algorithm for arbitrary network

- However, the relevant question is: are practical networks hard ?

  ○ We obtain novel algorithms using graph theoretic properties of practical networks

  - these are simple, distributed; throughput and delay optimal

End-of-Phase Goals

Goal 1.

Establish that it is not possible to design computationally efficient high throughput and low delay algorithm for wireless network under physical (SINR) model

Goal 2.

Design simple and distributed throughput-delay optimal algorithm for practical wireless network topologies under physical model

Wireless networks: Algorithmic trade-off between Throughput and Delay

FLoWS

STATUS QUO

Among two important performance metrics of wireless networks, throughput and delay, only throughput is well-understood in terms of fundamental limits and algorithm design.

Delay is far from being well-understood.

NEW INSIGHTS

1. Arbitrary networks: High-throughput is *easy*, low delay is *hard*.

2. Practical networks: distributed high-througput low delay is *possible*.

ACHIEVEMENT DESCRIPTION

MAIN RESULT:

1. High-throughput low delay algorithm for arbitrary wireless network is computationally intractable.

2. Wireless networks deployed in geographic area (in arbitrary manner) have high-throughput and low-delay algorithm distributed algorithms for scheduling and cross-layer design.

HOW IT WORKS:

1. Intractability follows by identifying computational hardness in scheduling through a novel equivalence relation.

2. Geometry in wireless networks allows for simple, high-performance algorithm design.

ASSUMPTIONS AND LIMITATIONS:

1. Wireless network with interference model.

END-OF-PHASE GOAL

1. Computational intractability of high throughput, low delay algorithm for wireless network under SINR model.

2. Simple algorithms for practical networks under SINR model.

COMMUNITY CHALLENGE

Computational intractability of information theoretic capacity achieving codes for wireless networks.

Algorithmic limitations for wireless network

Optimizing MaxWeight

FLoWS

STATUS QUO

What is the state of the art and what are its limitations?

Static routing: *ignores dynamics*

MW routing: *inflexible with respect to performance improvement*

Subramanian & Leigh 2007: *MW can be irrational*

NEW INSIGHTS

What are the key new insights?

MW **=** *Myopic* for a fluid model. Many such policies share the desirable properties of MW

ACHIEVEMENT DESCRIPTION

MAIN RESULT:

Geometric characterization of myopic policy with optimal throughput

Perturbation technique to generate functions with appropriate geometry

Application to policy synthesis for approximately optimal performance in heavy traffic

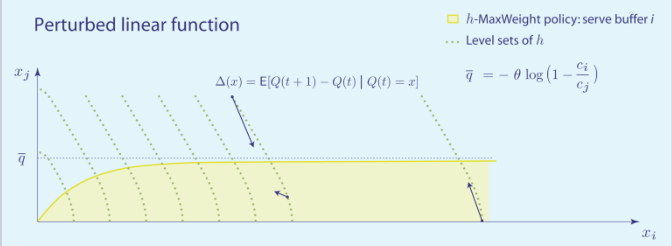

HOW IT WORKS:

Key analytical tool is Lyapunov theory for Markov processes

For approximate optimality, *workload relaxation* Relaxation also provides tool for visualization of high dimensional dynamics. Optimal solutions evolve in region containing *monotone region* for the effective cost.

END-OF-PHASE GOAL

- Decentralized implementation: Policy can be designed to use available information.
- Adaptation - on-line policy improvement
- Full analysis of multiple bottlenecks
- Integration with Network Coding projects: *Can we code around network hot-spots?*

COMMUNITY CHALLENGE

- Un-consummated union challenge: Integrate coding and resource allocation methodology
- Generally, solutions to complex decision problems should offer insight

Algorithms for dynamic routing: Visualization and Optimization

MaxWeight: What requires optimizing?

Routing requires *information.* In the MaxWeight policy, this information is obtained through queue length values. This can lead to irrational behavior when information is scarce.

Example (Subramanian and Leith, 2007, submitted). *MaxWeight or Backpressure routing will send packets upstream*!

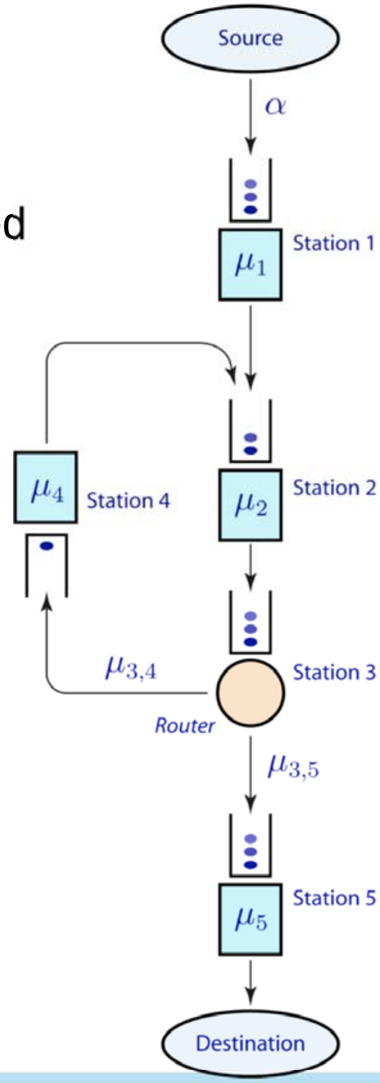

Questions addressed:

- Why does MW work?
- How can it be generalized and improved?
- Performance evaluation?

Analysis based on *new geometric insight*, and the *workload relaxation*

MaxWeight can be improved once it is better understood

Optimizing MaxWeight

- Perturbation technique: If h_0 is any monotone convex function

$$h(x) = h_0(\tilde{x}), \qquad \tilde{x}_i = x_i + \theta^{-1}(e^{-\theta x_i} - 1)$$

The function *h* serves as a Lyapunov function in a stability analysis
Chosen for mathematical elegance - many other possibilities!

- Optimization: Generalized min-cut to construct workload.

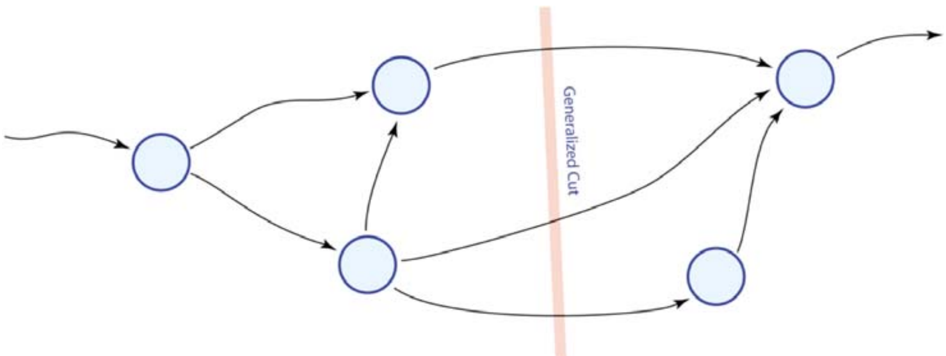

Asymptotic optimal policy is a function of workload. Implementation will require message passing, or other techniques to share information regarding dynamic hot-spots

- Learning locations of hot-spots can simplify network coding

Analytic techniques: Lyapunov theory and workload relaxation

Summaries and challenges

KEY CONCLUSION: Resource allocation for optimal throughput can be attained in many ways. *Some are better than others*!

LYAPUNOV THEORY: Quadratic Lyapunov function effective since it mirrors actually solution to DP equation. A tighter approximation results in better performance

RELAXATION: Workload relaxation enables construction of reduced-order model for which solution to the DP equation is obvious, provided there is a single bottleneck.

HOW BAD IS THE REAL WORLD? In the example of V&S, about 15% of packets are routed upstream. We discovered this increases dramatically with volatility. Is this seen in practice?

CAN WE LEARN? Especially when there is only a single bottleneck, key information for optimization is easy to identify. How can this information be shared?

CAN WE CODE? With the identification of dynamic bottlenecks, it is then evident where the capacity region can be improved

Largest current research bottleneck concerns
learning dynamic bottleneck location and workload

Large-Scale Network Utility Maximization (NUM)

ACHIEVEMENT DESCRIPTION

STATUS QUO

Dual decomposition is a widely used method for congestion control.

It is first order and decentralized.

Deals only with strictly concave utilities.

NEW INSIGHTS

A second order, primal-dual method performs better under wider network conditions (congested networks for instance). It is also able to handle not strictly concave utility functions.

MAIN RESULT:

Developed a primal-dual interior-point method for large-scale NUM, that outperforms dual decomposition.

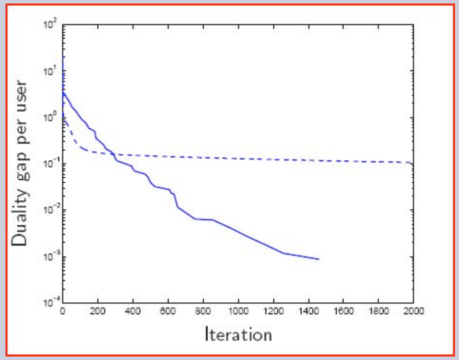

HOW IT WORKS:

Attempts to solve approximate optimality conditions at each iteration.

Computes search direction using preconditioned conjugate gradient method.

Can scale up to networks of 1,000,000 flows, or even more!

ASSUMPTIONS AND LIMITATIONS:

Algorithm is scalable, performs better but centralized.

END-OF-PHASE GOAL

An attempt to get a decentralized heuristic based on this method.

Including further extensions, like piecewise linear utility functions, link delay.

COMMUNITY CHALLENGE

Convergence issues of first order methods could render them impractical.

Towards second order methods for Network Utility Maximization

An Interior-Point Method for Large-Scale Network Utility Maximization

Argyrios Zymnis Nikolaos Trichakis

Stephen Boyd Dan O' Neill

Electrical Engineering Department

Stanford University

ITMANET PI
Meeting 07/26/07

NUM problem

$$\begin{array}{ll} \text{maximize} & U(f) = \sum_{j=1}^{n} U_j(f_j) \\ \text{subject to} & Rf \le c, \quad f \ge 0 \end{array}$$

- share resources
- dual decomposition
  - distributed, scalable
  - converges to global optimum
  - can back interpret protocols via U_j
  - will “track” changes in problem data U, R, or c

who can ask for more?

The bad news

- Requires U_j to be strictly concave
- first order method; can converge slowly
  - fast convergence for “symmetric” problems
  - slow convergence for “asymmetric” problems (*e.g.*, bottlenecks or long routes)
- hence, “tracks” changes very poorly

is this the price we have to pay for a distributed, scalable algorithm?

What we did

- worked out a scalable **but not decentralized** interior-point method for NUM
- second order method; handles asymmetries well
- fast convergence, independent of problem dimensions or data (!!)
  - scales to 10^6 or more flows
  - can optimize over 10^3 flows in $<10^{-3}$ sec (estimated)
- similar computational complexity per iteration to dual decomposition
- can track problem data very fast

Typical convergence

- 10^5 flows, $2*10^5$ links
- 200 congested links (each with $3*10^4$ flows)

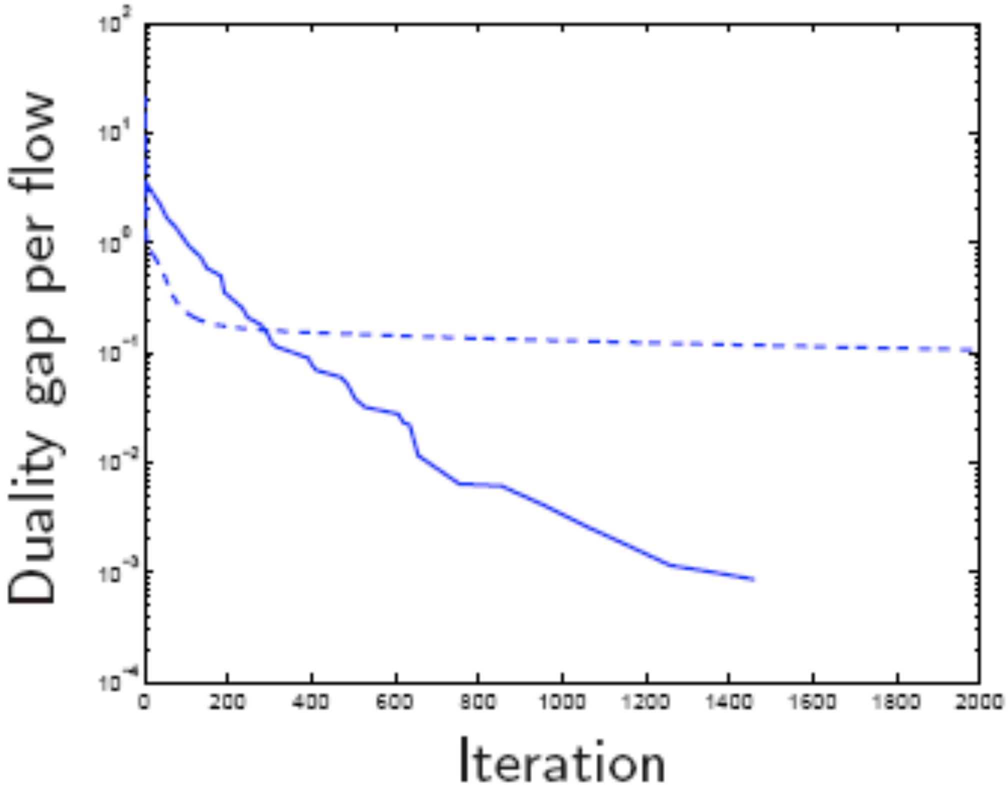

So what?

- we could actually evaluate convergence of dual decomposition for large networks
- dual decomposition is OK for “symmetric” data, for others not
- we guess there are practical uses
  - ability to quickly track optimum makes up for communication overhead
- centralized optimization and dual decomposition
  - not one versus the other
  - can apply dual decomposition at higher granularity;
  - whole subnets optimized quickly and centrally

Distributed Asynchronous Optimization Methods for General Performance Metrics

STATUS QUO

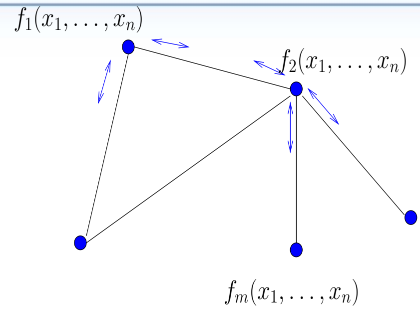

Existing methodology based on Lagrangian relaxation and duality does not lend itself to distributed algorithms for general non-separable (coupled) user performance metrics in wireless networks with time-varying connectivity

NEW INSIGHTS

Subgradient methods with simple consensus (averaging) policies lead to decentralized algorithms that
•optimize general performance metrics,
•are robust against changes in network topology

ACHIEVEMENT DESCRIPTION

MAIN RESULT:

- **Development of a distributed computational method for optimizing the sum of performance measures of users**
- **The method operates asynchronously under time-varying connectivity**
- **We provide convergence rate results that explicitly characterize the impact of the system and algorithm parameters on the quality of generated solutions.**

HOW IT WORKS:

- **Each user maintains an information state, which is an estimate of the optimal solution.**
- **The update rule for each user involves combining his information state with that of his current neighbors and performing a subgradient step using his local performance measure.**

ASSUMPTIONS AND LIMITATIONS:

- **The model is unconstrained.**
- **The communication bandwidth constraints have not been taken into account.**

END-OF-PHASE GOAL

•We will extend the model to include local (potentially time-varying) constraints for each user.

•We will explore the effect of bandwidth constraints (i.e., quantized information exchange) on the performance of the algorithms.

COMMUNITY CHALLENGE

Design of optimization algorithms that address the challenges and constraints associated with large-scale time-varying networks

Distributed optimization algorithms for general performance metrics and time-varying connectivity

Incomplete information, dynamics, and wireless games

FLoWS

STATUS QUO

Previous work studied ad hoc wireless resource competition among multiple nodes using game theoretic techniques, but typically in a stationary setting, where each node knows all other's channel conditions (see Huang et al., Etkin et al.)

We aim to understand the importance of a lack of information about channel conditions over time.

NEW INSIGHTS

We bring in the importance of incomplete channel information via the use of both static and dynamic *Bayesian games,* and in particular exploit results on reputation effects in economics to study primary/secondary competition.

(S. Adlakha, R. Johari, A. Goldsmith)

FLOWS ACHIEVEMENT(S)

MAIN RESULT: *The presence of incomplete channel information among nodes, as well as dynamic interaction among nodes, can dramatically alter the game theoretic conclusions drawn in standard complete information settings.*

Example: A primary user may deter entry by secondary users at some cost to himself, even if it is not immediately in his best interest to do so.

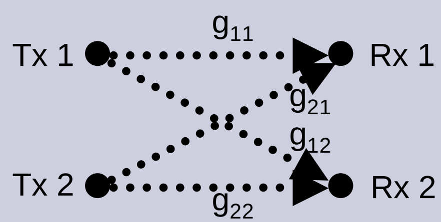

HOW IT WORKS: We use the theory of *Bayesian games* to find symmetric equilibria of a Bayesian Gaussian interference game.

We use the theory of *reputation effects in dynamic games of incomplete information* model to study the behavior of a primary user interacting with multiple secondary users.

ASSUMPTIONS AND LIMITATIONS:

We assume one primary and several secondaries arriving over time; we assume the channel remains stationary over several periods of interaction between primary and secondary.

Key assumption (and limitation): there is no "protocol" for transmission, so all other transmission treated as pure noise (hence the Gaussian interference model).

END-OF-PHASE GOAL

We need to extend the model to handle not only a finite horizon model, but also an infinite horizon model with changing channel conditions.

Journal paper is being prepared for submission to JSAC.

Longer term: we need to focus more on implications for *algorithm design for ad hoc wireless nodes in a reactive environment*. Our insights set a foundation for this.

COMMUNITY CHALLENGE

Status quo is useless for designing node strategies.

Employ methods from learning and dynamic equilibrium in large games to build better algorithms for competition and cooperation.

Real environments are reactive and non-stationary; this dramatically changes incentives and game theoretic predictions

Part I: Bayesian Gaussian interference game

- Assume transmit/receive pair 1 observes the incident gains g_{11}, g_{21}, but not g_{22} or g_{12} (similarly for Tx/Rx pair 2); assume flat fading
- This is a Bayesian game:
  Once random gains are realized, each TR pair knows its own gains but not the gains of the other
- This is a supermodular Bayesian game; in particular, local search dynamics converge (see also R. Berry's work)
- Nodes can either use a single channel, or spread power across all channels

  Theorem: equal spreading is unique symmetric equilibrium

Part II: Reputation effects in a dynamic game

- Now assume Tx/Rx 1 = primary, Tx/Rx 2 = secondary; same system model, but now assume only 2 channels
- Primary is *long-lived* and *fully rational*
- Secondary user is *myopic* (only optimizes one period payoff), but *history-aware* (remembers the past)
- Secondary user decides each period whether to "enter" (i.e., transmit), or "leave" (i.e., stay silent)
- Secondary user is assumed to have a cost for power consumption
- Primary user can "share" (give up a channel to secondary) or "spread" (spread power equally over channels)

Part II: Reputation effects in a dynamic game

When both g_{12}, g_{21} are large,

there can be a *reputation effect*:

Despite the fact that the primary would be better off sharing (in one period) if secondary enters,

the primary may choose to spread (“act” threatening)

because this deters future entry by the secondary

Key point:

This cannot happen in a complete information model!

(For complete information case, see Etkin et al.)

Information Theory for Mobile Ad-Hoc Networks (ITMANET): ***The FLoWS Project***

Application Metrics and Network Performance Summary

TECHNISCHE UNIVERSITÄT MÜNCHEN

Thrust Areas

- New Distributed Optimization Models for Resource Allocation
  - Building an algorithmic architecture that is robust against changes in network structure, optimizes general performance measures, scalable, and distributed
  - Incorporating networked-system constraints (e.g., asynchronism, stochastic elements, communication bandwidth constraints) in algorithm design, and quantifying the impact on performance
- Flow-based Models and Queuing Dynamics
  - Designing macro (flow) and micro (queuing) level network algorithms to yield desired performance
  - Integration of macro and micro level models
- New Resource Allocation Paradigm with Focus on Heterogeneous and Non-cooperative Nature of Users
  - Understanding when local competition yields globally desirable outcomes
  - Studying the dynamics that achieve the equilibrium

Achievements Overview

Optimization Theory
Distributed efficient algorithms for resource allocation

Boyd: Efficient second order methods for flow control

Shah: Throughput analysis of flow-level models with heterogeneous users

Ozdaglar: Distributed asynchronous optimization algorithms for general metrics and time-varying connectivity

Shah: Low complexity throughput and delay efficient scheduling

Goldsmith, Johari: Game-theoretic model for cognitive radio design in the presence of incomplete channel information

Meyn: Generalized Max-Weight to tradeoff information and performance

Johari: Topology formation model with application goals such as connectivity and cost of routing and link maintenance

Ozdaglar, Shah: Distributed scheduling and flow control to balance user and network performance

Stochastic Network Analysis
Flow-based models and queuing dynamics

Game Theory
New resource allocation paradigm that focuses on hetereogeneity and competition

Thrust Synergies

- General objective of the thrust requires:
  - Flow-level algorithms for optimizing heterogeneous application metrics
  - Packet-level algorithms for ensuring efficient and stable functioning of the network
  - Integration of application metrics and network capabilities

- Our thrust achieves these objectives through an algorithmic approach based on:
  - Development of efficient distributed optimization and control algorithms
  - Stochastic network analysis for stability and efficiency
  - Synergy in the integration of the macro and micro level models and of algorithmic optimization and stability analysis
  - Game-theoretic analysis of equilibrium models for
    - robustness against adversarial, competitive, and non-compliant behavior
    - modeling information structures

Synergies with Other Thrusts

- Resource negotiation for performance tradeoffs
  - Thrust 1 provides upper bounds on “performance region”
  - Thrust 2 provides achievable region
  - Thrust 3 chooses operating point on these regions
- Algorithms for implementing “building blocks” within network context
  - Thrust 2 uses information-theoretic analysis to provide closed-form or asymptotic solutions for canonical networks
  - Thrust 3 designs algorithms to incorporate these insights/building blocks into a network
- Algorithmic constraints may introduce new performance metrics for data processing in Thrust 2

Thrust Synergies: An Example

Thrust 1
Upper Bounds

Γ_1

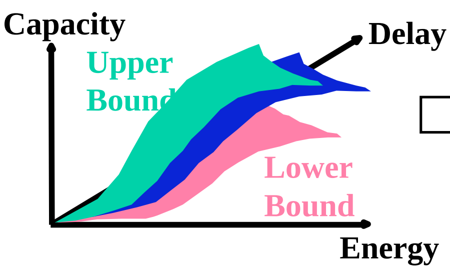

Thrust 3
Application Metrics and Network Performance

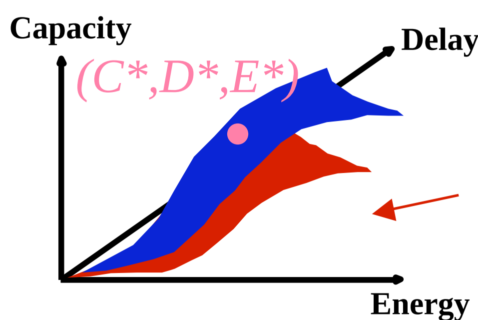

Γ_2

Thrust 2
Layerless Dynamic Networks

Johari: Topology formation model with application goals such as connectivity and cost of routing and link maintenance

Ozdaglar, Shah: Distributed scheduling and flow control to balance user and network performance

- Considering stochastic changes and micro-level considerations

Shah: Low complexity throughput and delay efficient scheduling

region)

Algorithmic constraints and sensitivity analysis may change the dimension of performance region

Roadmap for Phase 1

- Decentralized implementations for fast second order optimization methods
- Incorporation of networked-system constraints (bandwidth limitations, delays, stochastic elements) on distributed algorithm design
- High throughput low delay distributed scheduling algorithms in the presence of interference effects
- Decentralized implementations for generalized max-weight policies
- Design of dynamic algorithms for achieving equilibrium in game-theoretic models

Recent Publications

- Abhishek, S. Adlakha, Johari, and Weintraub, “Oblivious Equilibrium for General Stochastic Games with Many Players,” submitted to Allerton 2007.
- Adlakha, Johari, and Goldsmith, “Competition Between Wireless Devices with Incomplete Channel Knowledge,” submitted to IEEE JSAC 2007.
- Ahmed, Eryilmaz, Ozdaglar, and Medard, “Economic Gains from Network Coding in Wireless Networks,” submitted for publication 2007 (also appeared in Allerton 2006)
- Arcaute, Johari, and Mannor, “Network Formation: Bilateral Contracting and Myopic Dynamics” submitted to IEEE TAC 2007.
- Bayati, Prabhakar, Shah and Sharma, “Iterative Scheduling Algorithms,” IEEE Infocom, 2007.
- Bayati, Shah and Sharma, “Maximum Weight Matching via Max-Product Belief Propagation,” To appear in IEEE Information Theory Transactions, 2007.
- Coleman, Martinian, and Ordentlich, "Joint Source-Channel Decoding for Transmitting Correlated Sources over Broadcast Networks", submitted January 2007, IEEE Transactions on Information Theory (also appeared in 2006 International Symposium on Information Theory, Seattle, WA, July 10-14, 2006).

Recent Publications

- Doshi, Shah and Medard, “Source Coding with Distortion through Graph Coloring,” IEEE ISIT, 2007.
- Doshi, Shah, Medard and Jaggi, "Distributed Functional Compression through Graph coloring,” DCC, 2007.
- Doshi, Shah, Medard and Jaggi, “Graph Coloring and Conditional Graph Entropy,” Asilomar conference, 2006, pp: 2137-2141.
- Eryilmaz A., Ozdaglar A., Modiano E., “Polynomial Complexity Algorithms for Full Utilization of Multi-hop Wireless Networks,” IEEE Infocom, 2007.
- Meyn S., “Stability and Asymptotic Optimality of Generalized MaxWeight Policies”, submitted for publication, 2006
- Meyn. Control techniques for complex networks. To appear, Cambridge University Press, 2007.
- Mosk-Aoyama and Shah, “Computing Separable Functions via Gossip,” Under preparation. Preliminary version appeared in ACM PODC, 2006.
- Nedic and Ozdaglar, “Distributed Asynchronous Subgradient Methods for Multi-Agent Optimization,” submitted for publication, 2007.