

# Layering As Optimization Decomposition: Framework and Examples

Mung Chiang, Steven H. Low, A. Robert Calderbank, John C. Doyle

**Abstract**—Network protocols in layered architectures have historically been obtained primarily on an ad-hoc basis. Recent research has shown that network protocols may instead be holistically analyzed and systematically designed as distributed solutions to some global optimization problems in the form of Network Utility Maximization (NUM), providing insight into what they optimize and structures of the network protocol stack. This paper presents a short survey of the recent efforts towards a systematic understanding of ‘layering’ as ‘optimization decomposition’, where the overall communication network is modeled by a generalized NUM problem, each layer corresponds to a decomposed subproblem, and the interfaces among layers are quantified as functions of the optimization variables coordinating the subproblems. Different decompositions lead to alternative layering architectures. We summarize several examples of horizontal decomposition into distributed computation and vertical decomposition into functional modules such as congestion control, routing, scheduling, random access, power control, and coding.

**Keywords:** Cross-layer design, Distributed algorithm, Lagrange duality, Network utility maximization, Optimization, Reverse engineering, TCP/IP, Wireless ad hoc networks.

## I. INTRODUCTION

Layered architectures are one of the most fundamental and influential structures of network design. Each layer in the protocol stack hides the complexity of the layer below and provides a service to the layer above. While the general principle of layering is widely recognized as one of the key reasons for the enormous success of data networks, there is little quantitative understanding on a systematic, rather than an ad hoc, process of designing layered protocol stack for wired and wireless networks. One possible perspective to rigorously and holistically understand layering is to integrate the various protocol layers into a single coherent theory, by regarding them as carrying out an asynchronous distributed computation over the network to implicitly solve a global optimization problem. Different layers iterate on different subsets of the decision variables using local information to achieve individual optimality. Taken together, these local algorithms attempt to achieve a global objective. Such a framework of ‘layering as optimization decomposition’ exposes the interconnection between protocol layers and can be used to study rigorously the performance tradeoff in protocol layering, as different ways to distribute a centralized computation. Even though the design of a complex system will

always be broken down into simpler modules, this theory will allow us to systematically carry out this layering process and explicitly trade off design objectives, which will be particularly useful to wireless ad hoc networks.

The key idea in ‘layering as optimization decomposition’ is as follows. Different *decompositions* of an optimization problem, in the form of a generalized Network Utility Maximization (NUM), are mapped to different *layering* schemes in a communication network, with each decomposed subproblems in a given decomposition scheme corresponding to a layer, and functions of primal or Lagrange dual *variables* coordinating the subproblems corresponding to the *interfaces* among the layers. Since different decompositions correspond to alternative layering architectures, we can also tackle the question ‘how to and how not to layer’ by investigating the pros and cons of decomposition techniques. Furthermore, by comparing the objective function values under various forms of optimal decompositions and suboptimal decompositions, we can seek ‘separation theorems’ among layers: conditions under which strict layering incurs no loss of optimality. Robustness of these separation theorems can be further characterized by sensitivity analysis in optimization theory: how much will the differences in the objective value (between different layering schemes) fluctuate as constant parameters in utility maximization are perturbed.

## II. STRUCTURES OF NETWORK PROTOCOL STACK

The approach of ‘protocol as a distributed solution’ to some global optimization problem in the form of NUM has been successfully tested in trials for Transmission Control Protocol (TCP) [12]. The key innovation from this line of work [13], [14], [15], [20], [21], [22], [27] is to view TCP/IP network as an optimization solver and each variant of congestion control protocol as a distributed algorithm solving a specified NUM. Other recent results also show how to reverse engineer contention-based Medium Access Control (MAC) protocols into game-theoretic selfish utility maximization.

These reverse engineering successes provide one of the justifications to employ generalized versions of NUM for systematic cross-layer design. Furthermore, specific needs of applications and elasticity of traffic can both be modeled through general utility functions. As optimization’s objective, utility functions provide a metric to define optimality of resource allocation efficiency, and different shapes of utility functions also lead to different notions of fairness (*e.g.*,  $\alpha$ -fair utilities parameterized by  $\alpha > 0$ :  $U(x) = (1 - \alpha)^{-1}x^{1-\alpha}$  [22]).

While application needs form the objective function, *i.e.*, network utility to be maximized, restrictions in the communication

M. Chiang is with Electrical Engineering Department, Princeton University, [chiangm@princeton.edu](mailto:chiangm@princeton.edu). S. H. Low is with Computer Science and Electrical Engineering Departments, Caltech, [slow@caltech.edu](mailto:slow@caltech.edu). A. R. Calderbank is with Electrical Engineering and Mathematics Departments, Princeton University, [calderbk@princeton.edu](mailto:calderbk@princeton.edu). J. C. Doyle is with Control and Dynamic Systems Department, Caltech, [doyle@cds.caltech.edu](mailto:doyle@cds.caltech.edu).

infrastructure are translated into many constraints of a generalized NUM problem. Such problems may be very difficult non-linear, nonconvex optimization with integer constraints. There are many different ways to decompose a given problem, each of which corresponds to a different layering architecture. These decompositions, *i.e.*, layering schemes, have different trade-offs in efficiency, robustness, and asymmetry of information and control, thus some are ‘better’ than others depending on the criteria set by the network users and managers.

As evidenced by the large and ever growing number of papers on cross layer design over the last few years, we expect that there will be no shortage of cross layer ideas based on piecemeal approaches. What seems to be lacking is a level ground for fair comparison among the variety of cross layer designs, a unified view on how to and how not to layer, basic principles rigorously quantified, and fundamental limits on the impacts of layer-crossing on network performance and robustness metrics. ‘Layering as optimization decomposition’ provides a candidate for such a unified framework. What is unique about this framework is that it puts the end user application needs as the optimization objective, provides the globally optimal performance benchmark (since the optimal solution of a generalized NUM problem is independent of how the problem is decomposed), and leads to a systematic design of decomposed solution to attain the benchmark. The power of this framework has been illustrated through many case studies (four of which are summarized in Section IV) carried out by various groups in the last couple of years (an incomplete list include [1], [2], [3], [4], [6], [7], [11], [16], [17], [18], [19], [23], [25], [26], [29], [30], [31], [32]), generating considerable general insights in addition to the specific cross-layer designs.

### III. REVERSE ENGINEERING LAYERS 2-4 PROTOCOLS

Widely used network protocols, such as TCP, BGP, and IEEE 802.11 MAC, were designed based primarily on engineering intuitions and ad hoc heuristics. Surprisingly, it has recently been shown that TCP congestion control protocols are in fact implicitly solving the basic NUM problem *e.g.*, summarized in [20], [27], BGP inter-domain routing is implicitly solving a Stable Path Problem [10], and random access control protocols are implicitly using stochastic subgradient to participate in a non-cooperative game [17], [28]. These reverse engineering results put network protocol for the first time on a rigorous and systematic foundation, and also provide a strong evidence that the mathematics of distributed optimization has a relevant and important role to play in networking.

#### A. TCP congestion control solves the basic NUM

First consider a wired communication network with  $L$  links, each with a fixed link capacity of  $c_l$  bps, and  $S$  sources (*i.e.*, end users), each transmitting at a source rate of  $x_s$  bps. Each source emits one flow, using a fixed set  $L(s)$  of links in its path, and has a utility function  $U_s(x_s)$ . The basic NUM is the following problem formulated in the seminal work [13] by Kelly et. al.:

$$\begin{aligned} & \text{maximize} && \sum_s U_s(x_s) \\ & \text{subject to} && \sum_{s:l \in L(s)} x_s \leq c_l, \quad \forall l, \end{aligned} \quad (1)$$

where the variables are  $\mathbf{x}$ , the source rate vector.

The above basic NUM and its generalizations have recently been applied to study a variety of resource allocation problems in networks, in particular, Internet rate allocation through TCP congestion control protocols. The key innovation is to reverse engineer source rates as *primal variables*, link congestion prices as *dual variables*, and a TCP-AQM protocol as a distributed algorithm over the Internet to solve an implicit, global utility maximization in the form of NUM (1) and its Lagrange dual problem. Different TCP-AQM protocols solve for different utility functions using different link prices, which are implicitly feedback from links to sources. This model implies that the equilibrium properties of a large network under TCP/AQM control, such as throughput, delay, queue lengths, loss probabilities, and fairness, can be readily understood by studying the underlying NUM.

#### B. MAC contention resolution participates in a non-cooperative game

Currently, the Distributed Coordination Function (DCF) is the standardized MAC protocol in IEEE 802.11. However, it has been concluded by many researchers that DCF and its Binary Exponential Backoff (BEB) mechanism for contention resolution can be inefficient and unfair in face of location-dependent contentions. Various new algorithms have been developed to tackle these issues. To better understand the BEB protocol in wireless MAC, we pose the following question: are the distributed and selfish actions by each link in the BEB protocol in fact implicitly maximizing some local utility functions?

It has been shown [17], [28] that the contention resolution algorithm in backoff-based random-access protocols is implicitly participating in a non-cooperative game. Each link attempts to maximize a selfish local utility function, whose exact shape is reverse engineered from the protocol description, through a stochastic subgradient method in which the link updates its persistence probability based on its transmission success or failure. Existence of Nash equilibrium is guaranteed in general. The minimum amount of backoff aggressiveness needed for uniqueness of Nash equilibrium and convergence of the best response strategy are established as a function of user density. Convergence properties and connection with the best response strategy are also known for variants of the stochastic-subgradient-based dynamics of the game.

This motivates the need for forward engineering: what kind of new distributed MAC algorithms will be provably convergent to the global optimum of total network utility? After formulating a probabilistic-modeled NUM problem for wireless MAC, we develop optimal algorithms to solve the NUM problem, and these algorithms are then turned into random access MAC protocols [17]. Through this design approach, optimality with respect to prescribed user utilities, which in turn determine protocol efficiency and fairness, is guaranteed.

### IV. FORWARD ENGINEERING: FOUR SAMPLES OF SYSTEMATIC CROSS-LAYER DESIGN

In all four sample cases summarized below, a NUM problem that is more complicated than the basic NUM represents a more

general networking problem encompassing more than the congestion control function, and some functions of the Lagrange dual variables turn out to be the optimal ‘layering variables’.

#### A. Jointly optimal congestion control and adaptive coding

Adaptive error correction channel coding in physical layer can change the ‘pipe sizes’ of communication channels, adding another dimension in the ‘supply-demand’ balance in NUM. Indeed, the entire concept of signal quality is absent from the basic NUM (1). The link ‘capacities’  $c$  in (1) implicitly assume fixed decoding error probabilities  $P_e = \{P_{e,l}\}$ . In wireless networks, adaptive channel coding (including adaptive control between diversity-gain and multiplexing-gain in space-time coding) can change the rate-reliability *tradeoff*. A link can have a larger  $c_l$  and accommodate more flows, by increasing its decoding error probability  $P_{e,l}$ , or vice versa. A source may transmit at a higher rate if the end-to-end signal quality is allowed to degrade, *i.e.*, the end-to-end decoding error probability  $P_{e,s} \approx \sum_{l \in L(s)} P_{e,l}$  is larger. Of course, each source’s utility depends on both rate  $x_s$  and reliability  $P_{e,s}$ . In [16], we develop distributed algorithms to obtain the globally optimal rate-reliability tradeoff in the following NUM, with nonconvex constraints over variables  $\{x_s, P_{e,l}\}$ :

$$\begin{aligned} & \text{maximize} && \sum_s U_s(x_s, P_{e,s}) \\ & \text{subject to} && \sum_{s:l \in L(s)} x_s \leq c_l(P_{e,l}), \quad \forall l \end{aligned} \quad (2)$$

where each function  $c_l(P_{e,l})$  represents a nonlinear, generally nonconcave, and rather complicated dependency of a link’s attainable throughput  $c_l$  on the desired decoding error probability  $P_{e,l}$ . We show that, when each link provides the same decoding error probability for all flows through the link (an ‘integrated’ approach to dynamic reliability optimization), as in the formulation (2), a new distributed algorithm that uses pricing on both rate and reliability can be proved to converge to global optimality, as long as the channel codes are strong enough. When each link provides different decoding error probabilities for different flows (a ‘differentiated’ reliability approach), which expands the model in (2), the problem becomes a nonconvex optimization with coupling among the terms, even for strong channel codes. Using various decoupling techniques and change of variable, convergence to a global optimum can be proved for utility functions whose curvature is sufficiently negative, *i.e.*, if the traffic is elastic enough.

#### B. Jointly optimal congestion and contention control

In [17], we investigate joint end-to-end congestion control and per-link medium access control in ad-hoc wireless networks. Using a generalized NUM formulation, we can accommodate multi-class services as well as exploit the tradeoff between efficiency and fairness of resource allocation by adjusting the types of utility functions.

We define  $L_{out}(n)$  as a set of outgoing links from node  $n$ , and  $N_{to}^l(l)$  as the set of nodes whose transmission cause interference to the receiver of link  $l$ , excluding the transmitter node of link  $l$ . Each node decides to contend the medium with a persistence probability  $P^n$ , with each of its outgoing link’s

contention probability denoted as  $p_l$ . The effective link capacity becomes a product form as in the following generalized NUM with both source rates  $\{x_s\}$  and persistence probabilities  $\{p_l, P^n\}$  as optimization variables:

$$\begin{aligned} & \text{maximize} && \sum_s U_s(x_s) \\ & \text{subject to} && \sum_{s:l \in L(s)} x_s = c_l p_l \prod_{k \in N_{to}^l(l)} (1 - P^k), \quad \forall l \\ & && \sum_{l \in L_{out}(n)} p_l = P^n, \quad \forall n \\ & && 0 \leq P^n \leq 1, \quad \forall n \end{aligned} \quad (3)$$

Despite the inherent difficulties of nonconvexity and non-separability of the optimization problem, we show that, again under curvature negativity conditions on utility functions, we can develop a distributed algorithm, with limited message passing, that converges to the globally and jointly optimal rate allocation and persistence probabilities. These results can accommodate general concave utility function (the special case of logarithmic utility function is readily separable [29]). Different from the other three sample case studies in this section, it is better to use a primal penalty function approach rather dual decomposition in this case, the engineering implication of which is that global congestion control and local contention control can operate on the same timescale.

#### C. Jointly optimal congestion control, routing, and scheduling

In multihop ad hoc wireless networks, end-to-end congestion control, routing, and scheduling among contending links are coupled. Route choices not only affect congestion control in the transport layer, but also determine schedulable regions at the physical layer. In [2], we model contention relations between wireless links as a conflict graph [24], which indicates the set of links that mutually interfere and cannot be active simultaneously. This determines a feasible rate region  $\Pi$ .

Consider an ad hoc wireless network with a set  $N$  of nodes and a set  $L$  of logical links. We assume fixed physical layer resource allocations so that each logical link  $l$  has a fixed capacity  $c_l$  when it is active. The feasible rate region at the link layer is the convex hull of the corresponding rate vectors of independent sets of the conflict graph. Let  $x_i^k$  be the flow rate generated at node  $i$  for destination  $k$ . Let  $f_{ij}^k$  be the amount of capacity of link  $(i, j)$  allocated to the flows on that link for final destination  $k$ . Consider the following generalized NUM in variables  $\{x_s\}$  (where  $x_s$  is a shorthand for  $x_i^k$ ) and  $\{f_{ij}^k\}$ :

$$\begin{aligned} & \text{maximize} && \sum_s U_s(x_s) \\ & \text{subject to} && x_i^k \leq \sum_{j:(i,j) \in L} f_{ij}^k - \sum_{j:(j,i) \in L} f_{ji}^k, \quad \forall i, j, k \\ & && f \in \Pi. \end{aligned} \quad (4)$$

After formulating the above joint congestion control and routing problem subject to schedulability constraint, we apply Lagrangian relaxation and dual decomposition and derive a distributed subgradient algorithm for joint congestion control, routing and scheduling. This leads to a jointly optimal cross-layer design where a source adjusts its sending rate based on the congestion prices generated locally at the node, the back-pressure from the differential price of neighboring nodes is used to perform optimal scheduling, and routing is automatically

generated by the scheduling decision. We also prove that it converges arbitrarily close to the system optimum, and remain stable and optimal (on average) when the schedulability constraint set is modulated by a Markov chain representing time-varying channels.

#### D. TCP/IP interactions

TCP reverse engineering assumes fixed routing  $\{L(s)\}$ . The routing is computed by variants of Internet Protocol (IP) and updated on a different timescale, based on traffic condition in the network. There is hence a feedback loop where IP routing decision at time  $t$  determines flow rates  $\mathbf{x}(t)$  and congestion prices  $\lambda(t)$  through the capacity constraints in (1), and the congestion prices  $\lambda(t)$  affects routing in the next time instance. First consider the case where TCP converges faster than each IP update. It is shown in [30] that an equilibrium of TCP/IP, if exists, indeed solves NUM over *both* source rates  $\{x_s\}$  and routes  $\{L(s)\}$ , provided congestion prices are used as link costs in the shortest-path computation. Since the routing matrix is discrete, the NUM problem is no longer a convex optimization problem. An equilibrium exists if and only if this NUM and its Lagrange dual have zero duality gap. In this case, TCP/IP incurs no penalty in not splitting traffic across multiple paths. When there is a non-zero duality gap, the gap can be interpreted as the penalty (in utility) in not splitting the traffic. In TCP/IP systems, the primal NUM problem is in general NP-hard, but the subclass of NUM that has zero duality gap is polynomial-time solvable.

In [11], we further consider three alternative timescale separations between TCP and IP dynamics, and examine the stability and optimality of each system. Analytic characterizations and simulation experiments demonstrate how the step size of the congestion-control algorithm affects the stability of the system, and how the timescale of each control loop and homogeneity of link capacities affect system stability and optimality. In particular, the stringent conditions on capacity configuration for TCP/IP interaction to remain stable suggests that congestion price, on its own, would be a poor layering price for TCP and IP in practice.

#### V. ALTERNATIVE DECOMPOSITIONS

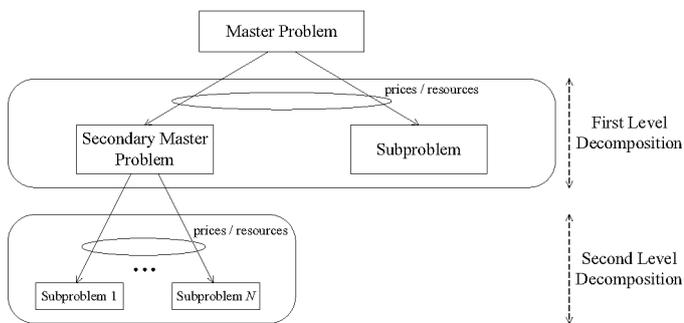


Fig. 1. An example of a multilevel primal/dual decomposition with two levels.

The basic idea of decomposition is to divide the original large problem into distributively solvable subproblems which

are then coordinated by a master problem by means of some kind of signalling. Most of the existing decomposition techniques can be classified into *primal decomposition* and *dual decomposition* methods. The former is based on decomposing the original primal problem, whereas the latter based on decomposing the Lagrange dual of the problem. Primal decomposition methods have the interpretation that the master problem directly gives each subproblem an amount of resources that it can use; the role of the master problem is then to properly allocate the existing resources. In dual decomposition methods, the master problem sets the price for the resources to each subproblem which has to decide the amount of resources to be used depending on the price; the role of the master problem is then to obtain the best pricing strategy.

Almost all the papers in the vast, recent literature on NUM use a standard dual-based distributed algorithm. Contrary to the apparent impression in the community that such a decomposition is the only possibility, we show that [26] there are in fact many alternatives to solve a given network utility problem in different but all distributed manners, including multi-level and partial decompositions. Each of the alternatives provides a possibly different tradeoff among three important considerations: convergence speed, amount and asymmetry of message passing's communication overhead, and architecture of distributed computation. There is no universally 'best' way to distribute the solution process across a network: which alternative is the most desirable depends on the specific problem formulation and application.

#### VI. A MORE GENERAL EXAMPLE OF GENERALIZED NUM FOR WIRELESS NETWORKS

We now summarize an example of generalized NUM for wireless ad hoc networks [4] that includes earlier examples as special cases. Notice again that the NUM problem formulation itself has no layering but it can be solved through a variety of layered algorithms through decompositions:

$$\begin{aligned} & \text{maximize} && \sum_i U_i(x_i, P_{e,i}) + \sum_j V_j(w_j) \\ & \text{subject to} && \mathbf{R}\mathbf{x} \preceq \mathbf{c}(\mathbf{w}, \mathbf{P}_e), \\ & && \mathbf{x} \in \mathcal{C}_1(\mathbf{P}_e), \quad \mathbf{x} \in \mathcal{C}_2(\mathbf{F}), \\ & && \mathbf{R} \in \mathcal{R}, \quad \mathbf{F} \in \mathcal{F}, \quad \mathbf{w} \in \mathcal{W}. \end{aligned} \quad (5)$$

Here,  $x_i$  denotes the rate for source  $i$  and  $w_j$  denotes the physical layer resource at network element  $j$ . The utility functions  $U_i$  and  $V_j$  may be any nonlinear, monotonic functions.  $\mathbf{R}$  is the routing matrix and  $\mathbf{c}$  are the logical link capacities as functions of both physical layer resources  $\mathbf{w}$  and the desired decoding error probabilities  $\mathbf{P}_e$ . The issue of signal interference and power control can be captured in this functional dependency. The rates must also be constrained by the interplay between physical layer decoding reliability and upper layer error control mechanisms like ARQ in link layer. This constraint set is denoted as  $\mathcal{C}_1(\mathbf{P}_e)$ . The issue of rate-reliability tradeoff and coding is captured in this constraint. The rates are further constrained by the medium access success probability, represented by the constraint set  $\mathcal{C}_2(\mathbf{F})$  where  $\mathbf{F}$  is the contention matrix. The issue of packet collision and medium access control is captured in this constraint. The sets of possible physical layer resource allocation schemes, of possible scheduling or contention

based medium access schemes, and of single-path or multi-path routing schemes are represented by  $\mathcal{W}$ ,  $\mathcal{F}$ ,  $\mathcal{R}$ , respectively. The optimization variables are  $\mathbf{x}$ ,  $\mathbf{w}$ ,  $\mathbf{P}_e$ ,  $\mathbf{R}$ ,  $\mathbf{F}$ . Different specifications of 5 lead to various concrete problem statements.

## VII. CONCLUSION

‘Layering as optimization decomposition’ is a unifying framework for understanding and designing distributed control and cross-layer resource allocation in wired and wireless networks. It has been developed by various research groups over the last several years, and is now emerging to provide a mathematically rigorous and practically relevant approach to quantify the risks and opportunities of modifying existing layered network architecture. It shows that network protocols in layers 2, 3, and 4 can be reverse-engineered as implicitly solving some optimization-theoretic or game-theoretic problems. By distributively solving generalized NUM formulations through decomposed subproblems, we can systematically generate layered protocol stacks. There are many alternatives for both horizontal decomposition into disparate network elements and vertical decomposition into functional modules (*i.e.*, layers). Similarly, while queuing delay or buffer occupancy is often used as the ‘layering price’, it may sometimes lead to unstable interactions. A variety of techniques to tackle coupling and nonconvexity issues in the objective and constraints have become available. A more detailed survey of recent efforts to establish ‘layering as optimization decomposition’ as a common ‘language’ for systematic network design can be found in [5].

While this short survey has focused on deterministic flow models, there has been progress on characterizing stability under feedback delay and stochastic models at session, packet, and channel levels. Nonconcave utility maximization and heterogeneously-priced interactions have also been investigated recently. Many issues in this area remain open, however, especially on constructing utility functions, modeling of sophisticated physical and link layer algorithms, and characterizing effects of stochastic models and transient behaviors.

Industry adoption of ‘layering as optimization decomposition’ has already started. For example, insights from reverse-engineering TCP has lead to an improved version of TCP implemented over the last several years: FAST (Fast AQM Scalable TCP) [8], [12]. Putting end-user application utilities as the objective function has lead to a new way to leverage innovations in the physical and link layer beyond the standard metrics of BER, *e.g.*, in FAST (Frequency, Amplitude, Space, Time) Copper Project for fiber/DSL broadband access systems [9].

## ACKNOWLEDGEMENT

We would like to thank all the co-authors of our publications along the direction of ‘layering as optimization decomposition’.

## REFERENCES

- [1] C. S. Chang and Z. Liu, “A bandwidth sharing theory for a large number of HTTP-like connections,” *IEEE/ACM Trans. on Networking*, vol. 12, no. 5, Oct. 2004.
- [2] L. Chen, S. H. Low, M. Chiang, and J. C. Doyle, “Optimal joint congestion control, routing, and scheduling design for wireless ad hoc networks,” *Proc. IEEE INFOCOM*, April 2006.
- [3] L. Chen, S. H. Low, and J. C. Doyle, “Joint congestion control and medium access control for wireless ad hoc networks,” *Proc. IEEE INFOCOM*, March 2005.
- [4] M. Chiang, “Balancing transport and physical layers in wireless multihop networks: Jointly optimal congestion and power control” *IEEE J. Sel. Area Comm.*, vol. 23, no. 1, pp.104-116, Jan. 2005.
- [5] M. Chiang, S. H. Low, R. A. Calderbank, and J. C. Doyle, “Layering as optimization decomposition”, To appear in *Proceedings of IEEE*, 2006.
- [6] R. L. Cruz and A. Santhanam: “Optimal routing, link Scheduling, and power control in multihop wireless networks,” *Proc. IEEE INFOCOM*, April 2003.
- [7] A. Eryilmaz and R. Srikant, “Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control,” *Proc. IEEE INFOCOM*, March 2005.
- [8] TCP FAST Project [netlab.caltech.edu](http://netlab.caltech.edu).
- [9] FAST Copper Project [fastcopper.princeton.edu](http://fastcopper.princeton.edu).
- [10] T. G. Griffin, F. B. Shepherd, and G. Wilfong, “The stable path problem and interdomain routing,” *IEEE Trans. on Networking*, vol. 10, no. 2, pp. 232-243, April 2002.
- [11] J. He, M. Chiang, and J. Rexford, “TCP/IP interaction based on congestion prices: Stability and optimality”, *Proc. IEEE ICC*, June 2006.
- [12] C. Jin, D. X. Wei, and S. H. Low. TCP FAST: motivation, architecture, algorithms, performance. In *Proc. of IEEE INFOCOM*, March 2004.
- [13] F. P. Kelly, A. Maulloo, and D. Tan, “Rate control for communication networks: Shadow prices, proportional fairness and stability,” *Journal of Operations Research Society*, vol. 49, no. 3, pp.237-252, March 1998.
- [14] S. Kunniyur and R. Srikant. End-to-end congestion control: utility functions, random losses and ECN marks. *IEEE/ACM Trans. on Networking*, 2003.
- [15] R. J. La and V. Anantharam, “Utility-based rate control in the Internet for elastic traffic,” *IEEE/ACM Tran. Networking*, vol. 10, no. 2, pp. 272-286, April 2002.
- [16] J. W. Lee, M. Chiang, and R. A. Calderbank, “Price-based distributed algorithm for optimal rate-reliability tradeoff in network utility maximization”, To appear *IEEE Journal of Selected Areas in Communications*, 2006.
- [17] J. W. Lee, M. Chiang, and R. A. Calderbank, “Utility-optimal medium access control: reverse and forward engineering,” *Proc. IEEE INFOCOM*, April 2006.
- [18] J. W. Lee, M. Chiang, and R. A. Calderbank, “Jointly optimal congestion and contention control in wireless ad hoc networks”, To appear in *IEEE Communication Letters*, 2006.
- [19] X. Lin and N. Shroff, “Impact of imperfect scheduling in cross layer design,” *Proc. IEEE INFOCOM*, March 2005.
- [20] S. H. Low, “A duality model of TCP and queue management algorithms,” *IEEE/ACM Trans. on Networking*, vol. 11, no. 4, pp. 525-536, Aug. 2003.
- [21] S. H. Low and D. E. Lapsley, “Optimization flow control, I: basic algorithm and convergence,” *IEEE/ACM Trans. on Networking*, vol. 7, no. 6, pp. 861-874, December 1999.
- [22] J. Mo and J. Walrand, “Fair end-to-end window-based congestion control,” *IEEE/ACM Trans. on Networking*, vol. 8, no. 5, pp. 556-567, Oct. 2000.
- [23] H. Nama, M. Chiang, and N. Mandayam, “Utility lifetime tradeoff in self regulating wireless sensor networks: A cross-layer design approach”, *Proc. IEEE ICC*, June 2006.
- [24] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan, “Achieving MAC layer fairness in wireless packet networks,” *Proc. ACM Mobicom*, Boston, USA, Aug. 2000.
- [25] M. J. Neely, E. Modiano, and C. E. Rohrs, “Dynamic power allocation and routing time varying wireless networks”, *IEEE J. Sel. Area Comm.*, vol. 23, no. 1, pp. 89-103, Jan. 2005.
- [26] D. Palomar and M. Chiang, “Alternative decompositions for distributed maximization of network utility: Framework and applications”, *Proc. IEEE INFOCOM*, April 2006.
- [27] R. Srikant, *The Mathematics of Internet Congestion Control*, Birkhauser 2004.
- [28] A. Tang, J. W. Lee, J. Huang, M. Chiang, and R. A. Calderbank, “Reverse engineering MAC”, *Proc. IEEE WiOpt*, April 2006.
- [29] X. Wang and K. Kar, “Cross-layer rate control for end-to-end proportional fairness in wireless networks with random access,” *Proc. ACM Mobicom*, May, 2005.
- [30] J. Wang, L. Li, S. H. Low, and J. C. Doyle. “Cross-layer optimization in TCP/IP networks,” *IEEE/ACM Trans. Networking*, vol. 13, no. 3, pp. 582-595, August 2005.
- [31] L. Xiao, M. Johansson, and S. Boyd, “Joint routing and resource allocation via dual decomposition,” *IEEE Trans. Comm.*, vol. 52, no. 7, pp. 1136-1144, July 2004.
- [32] W. Yu and J. Yuan, “Joint source coding, routing, and resource allocation for wireless sensor networks,” *Proc. IEEE ICC*, May 2005.