

---

# Multimedia-Centric Routing for Multiple Description Video in Wireless Mesh Networks

Shiwen Mao, Auburn University  
Y. Thomas Hou, Hanif D. Sherali, and Scott F. Midkiff,  
Virginia Polytechnic Institute and State University

---

## Abstract

Wireless mesh networks have emerged as a promising technology for providing ubiquitous access to mobile users, and quick and easy extension of local area networks into a wide area. The inherent low throughput problem makes multipath routing and load balancing increasingly important for supporting multimedia applications in such networks. We study the problem of cross-layer multipath routing for multiple-description video communications in wireless mesh networks. We aim to answer the following questions: how to formulate a multimedia-centric routing problem; what are the performance limits; how to design an efficient routing algorithm based on the formulated problem; and how to implement the algorithms in a distributed manner. Both simulation results and experiments from a testbed network are presented to demonstrate the efficacy of the proposed approach. We show that cross-layer design is imperative and effective in meeting the stringent quality of service requirements of multimedia applications in wireless mesh

**W**ireless mesh networks have emerged as a promising technology for providing ubiquitous access to mobile users, and quick and easy extension of local area networks into a wide area [1]. They are multihop wireless networks consisting of static wireless routers mounted on rooftops or lamp poles, as well as mobile users with wireless devices. Compared to mobile ad hoc networks, mesh networks have several appealing properties. First, the static wireless routers form a backbone that provides connectivity. As a result, the links among wireless routers are relatively more stable than those in mobile ad hoc networks. Second, wireless routers may be more powerful than normal mobile devices in terms of computational power and may have steady power supply. They may also be equipped with multiple network interfaces for enhanced performance. Therefore, sophisticated algorithms can be used, while computational complexity and energy conservation are secondary issues.

Apart from their advantages, there remain some challenges in wireless mesh networks, particularly for real-time multimedia communications. For example, multihop wireless paths remain more fragile than their counterparts in wired networks as well as infrastructure-based wireless networks. In addition, end-to-end throughput in a multihop wireless network quickly diminishes when the number of traversed hops increases. These have added considerable difficulties in provisioning multimedia services (e.g., video) in wireless mesh networks.

Provisioning of video communications in wireless mesh net-

works is a challenging problem involving multiple disciplines, such as signal processing, wireless networking, and optimizations. It has been shown that the multiple-description (MD) coding technique is highly suitable for video communication in multihop wireless networks [2, 3]. With MD coding, multiple descriptions are generated for a video source, each giving a low but acceptable video quality. Unlike traditional layered video coding, the descriptions are equally important in video decoding. At the destination, we can reconstruct the video from *any* subset of the received descriptions, with video quality commensurate with the number of received descriptions. MD video is thus an excellent match for wireless mesh networks.

In this article we study the important problem of multipath routing for MD video in wireless mesh networks. We present an optimal routing approach that aims to explicitly minimize the distortion of received video (termed *multimedia-centric routing* throughout this article). We will use *multimedia-centric* and *cross-layer* interchangeably, referring to the approach that optimizes application-layer media quality through multipath routing. Specifically, we aim to address the following issues:

- How to characterize and model the multimedia-centric multipath routing problem
- The performance limits on video quality for a given mesh network
- How to design an efficient solution procedure for the formulated cross-layer routing problem
- How to implement the proposed routing solution in a distributed manner

We consider the joint design of application and network layers in this article, assuming that lower-layer dynamics are fixed and can be characterized in the problem formulation.

The remainder of this article is organized as follows. We first present the cross-layer formulation for MD video over wireless mesh networks, as well as efficient metaheuristic-based solution procedures. We then present possible approaches to distributed implementations of the proposed algorithms. We finally describe an emulation testbed. We then conclude this article.

## Cross-Layer Multipath Routing

### The Proposed Approach

We follow a cross-layer design approach by considering application layer performance (i.e., video distortion) as a function of network layer performance metrics (e.g., bandwidth, loss, and path correlation). Due to the fact that received video quality is a complex function of the above metrics, such a cross-layer approach has the potential of achieving considerable improvements in video quality over traditional approaches that focus on one or two network performance metrics.

By exploring the monotone properties of the formulated problem, we develop a simple lower bound on video distortion. The lower bound can be computed with computational complexity similar to that of Dijkstra's algorithm and can be used as a benchmark for evaluating the performance of heuristic algorithms. We then develop efficient metaheuristic solution procedures for the formulated problem. In particular, we find that genetic algorithms (GAs) are quite suitable in addressing such complex cross-layer optimization problems. We also observe that multimedia-centric routing achieves significant performance gains over network-centric approaches.

### Mathematical Modeling

We consider a wireless mesh network with  $N$  nodes, each being a wireless router or mobile user [1]. We consider connectivity to exist between two nodes if they are within each other's transmission range. The network can thus be modeled as a directed graph  $G(V, E)$ , where  $V$  is the set of vertices representing the  $N$  nodes and  $E$  is the set of links. We consider a video session for source node  $s$  to destination node  $t$ .

We assume the end-to-end path model can be derived from a concatenation of the link models along the path. That is, a packet will be successfully delivered if and only if it survives the loss process at each link along the path. Another performance metric, available bandwidth, can be measured by each node, while end-to-end bandwidth of a path is the minimum of the available bandwidths of its constructive links.

Consider an MD video with  $M$  descriptions; the source transmits description  $i$  on path  $P_i$ ,  $i = 1, 2, \dots, M$ . The video quality could be measured by the distortion of received video, that is, the difference between a reconstructed video frame and the corresponding original video frame. The distortion metric  $D$  depends on the video codec (information loss due to quantization) and dynamics of all  $M$  paths (information loss due to transmission errors). For the double description (DD) video case (i.e.,  $M = 2$ ), the average video distortion can be characterized by the following expression [4]:

$$D = \pi_{00} \cdot d_{00} + \pi_{01} \cdot d_{01} + \pi_{10} \cdot d_{10} + \pi_{11} \cdot \sigma^2, \quad (1)$$

where  $\pi_{00}$  is the probability that both descriptions are received;  $\pi_{01}$  is the probability that description 1 is received but description 2 is lost;  $\pi_{10}$  is the probability that description 1 is lost but description 2 is received;  $\pi_{11}$  is the probability that both descriptions are lost;  $d_{00}$  is the distortion when both

descriptions are received;  $d_{01}$  is the distortion when only description 1 is received;  $d_{10}$  is the distortion when only description 2 is received; and  $\sigma^2$  is the distortion when both descriptions are lost (i.e., the variance of the source). The conditional distortions  $d_{00}$ ,  $d_{01}$ , and  $d_{10}$  can be approximated using the Ozarow bound introduced in [15].

Note that the Ozarow bound is based on the DD distortion rate region (called the DD *region*) for memoryless Gaussian sources and the squared error distortion measure [5]. Furthermore, the DD region for any continuous-valued memoryless source with squared error distortion can be bounded by this region [2]. Other coder/sequence-specific rate distortion models developed in the literature can easily be incorporated into this formulation as well.

We can now formulate the following optimization problem:

**Minimize :**  $D$

**Subject to :**

$$\begin{cases} \text{flow constraints} \\ \text{link stability constraints.} \end{cases}$$

This is a cross-layer optimization problem that explicitly optimizes application performance via network layer operations. The objective is to minimize the average video distortion  $D$ . Clearly, the choice of paths, as well as the dynamics on each path, will determine the distortion value. The choice of a pair of paths can be expressed using a set of binary index variables, one for each link. When the index variable is 1, the corresponding link is used in a path; otherwise, it is not chosen. Once the set of paths are given, path performance metrics can be derived from the link metrics, which finally yields the distortion value  $D$ . Solving the above problem will provide the values for the index variables (i.e., a pair of paths) as well as the bit rate for each description.

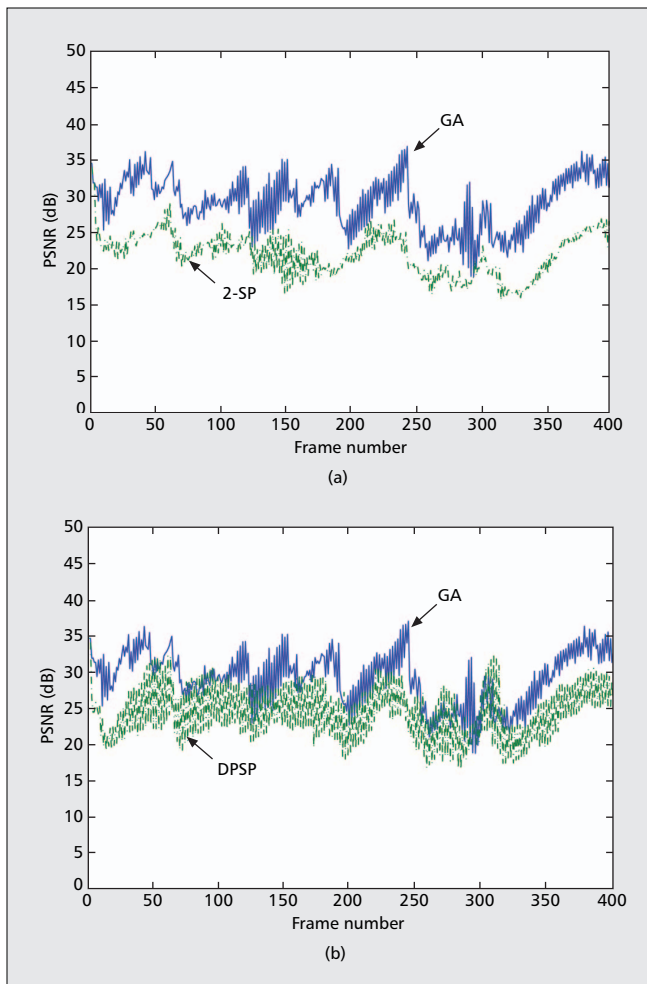
This is a combinatorial optimization problem. Although the set of possible combinations of paths is finite, the search space is exponentially large ( $O(2^{MN})$ ). The multimedia-centric routing problem has a highly complex relation pertaining to the contribution of any link to the objective function, which depends, in general, on the other links which are included in a fashion that possesses no particular structural property such as convexity.

### Approximation Bounds

Before proceeding to developing efficient solution procedures for the formulated problem, we explore a possible approximation bound on the achievable video distortion by considering the inherent structure of the above optimization problem. The derived lower bound can be used as a benchmark when evaluating the performance of a heuristic algorithm.

It can be shown that the average video distortion  $D$  has *monotonicity* properties with respect to rate, end-to-end path success probability, and path correlation. Based on these properties, we construct a lower bound on the average video distortion using an algorithm termed ALG-LB [3]. In ALG-LB, we first find the optimal bandwidth and the optimal end-to-end loss probability among all paths between  $s-t$ . The maximum end-to-end bandwidth can be found using an algorithm with time complexity  $O(|E| \cdot \log^* |V|)$ , where  $\log^* x$  is the iterated logarithm function; the minimum end-to-end loss probability can be found using Dijkstra's algorithm with time complexity  $O(|E| + |V| \cdot \log |V|)$ .

We then construct two disjoint virtual paths that have these optimal characteristics. It can be shown that the distortion achieved by this constructed path pair is a strict lower bound for  $D$ . Intuitively, the constructed path pair provides a pair of



■ Figure 1. Comparison of PSNRs obtained by multimedia-centric multipath routing and network-centric multipath routing: a) cross-layer routing vs. 2-SP; b) cross-layer routing vs. DPSP.

description rates that are no lower than the global optimal solution value. It also dominates all other existing path pairs in terms of end-to-end loss performance, be it disjoint or joint. Note that the purpose here is not to provide a pair of usable paths for the video session, but to derive an approximation for the fundamental performance limit. The lower bound is obtained by relaxing the constraints on bandwidth and loss properties, as well as constraints on path correlation. As a result, the lower bounding solution may not be realizable and does not provide a set of routing paths for the video session.

### Metaheuristic Algorithms

For solving the formulated problem, a promising strategy is to explore an effective *metaheuristic* approach. The term metaheuristic refers to the set of strategies that “guide” the iterative search process in order to find near-optimal solutions. In this section we examine three metaheuristic algorithms that are especially useful in solving combinatorial optimization problems: simulated annealing (SA), tabu search (TS), and genetic algorithms (GAs) [6].

SA was developed to address highly nonlinear optimization problems. It was motivated by an analogy between the way a metallic body cools and freezes into a minimum energy crystalline structure (the annealing process), and the search for a minimum in a more general system. The basic idea in SA is to allow moves to escape from local optima, even if they may not improve the current value of the objective function. The prob-

ability of accepting nonimproving solutions (i.e., the randomness in the search procedure) decreases over time according to some cooling procedure. As cooling proceeds, SA gradually converges to a simple iterative improvement algorithm.

Another well-known technique for combinatorial optimization is TS, which explicitly uses the history of the search, both to escape from local optima and to implement an exploratory strategy [6]. More specifically, TS uses a *tabu list*, which stores recently visited solutions to prevent a return to these solutions, thereby avoiding cycling and possibly forcing the search process to accept nonimproving solutions in order to escape from a local minimum.

Both SA and TS belong to the class of *trajectory* metaheuristics. The attractive characteristic of this class of algorithms is that only a single solution needs to be maintained throughout the iterative search process, which can be characterized by a trajectory in the search space.

In contrast, the so-called *population-based* metaheuristics maintain a population of solutions and evolve them toward optimality. Among the class of population-based metaheuristics, we find that GAs appear to be particularly promising for addressing the multimedia-centric routing problem. A GA is a population-based algorithm inspired by the *survival-of-the-fittest* principle. It has the intrinsic strength of dealing with a set of solutions (i.e., a *population*) at each iterative step, rather than working with a single current solution. The basic assumption within this paradigm is that good solutions often share parts with optimal solutions. At each iteration, a number of “genetic” operators are applied to the individuals of the current population in order to generate individuals for the next generation. The driving force in a GA is the *selection* of individuals based on their *fitness* values (in the form of an objective function). Individuals having a higher degree of fitness are more likely to be chosen for the next generation. The survival-of-the-fittest principle ensures that the overall quality of the population improves as the algorithm progresses from one generation to the next.

We develop and implement GA-, SA-, and TS-based algorithms to solve the formulated multimedia-centric multipath routing problem. We find that the GA-based approach outperforms both SA- and TS-based approaches for the cases we examined. We believe that this result is largely due to GA’s unique strength in manipulating a population of solutions and exploring the solution space in parallel.

### Performance Comparisons

In this section we present a simulation study to compare the performance of several algorithms in solving the multimedia-centric routing problem. For each experiment, we generate a mesh network topology with a number of nodes placed at random locations in a square region, with connectivity among the nodes being determined by the distance coverage of each node’s transmitter. For each link, the failure probability is assumed to be uniform within the range of [0.01, 0.3]. The available bandwidth and mean burst length are both uniformly chosen from [100, 400] kb/s and [2, 6], respectively. The source-destination nodes  $s$  and  $t$  are randomly chosen from the nodes. The routing algorithms are executed at the source node. A DD video is encoded at the source node at the rate computed by the routing algorithm, and each description is transmitted along a path as specified by the routing algorithm. The receiver decodes the received video and evaluates the quality of the reconstructed frames. We present several interesting results here and refer interested readers to [3] for more detail.

*Advantages of Cross-Layer Routing* — We first compare the proposed multimedia-centric routing with existing network-



■ Figure 2. Frame 224 from the original and decoded video sequences "Foreman" delivered by different routing algorithms: a) original; b) cross-layer routing; c) 2-SP; d) DPSP.

centric routing approaches using a 50-node network. For the multimedia-centric approach, we use GA-based multipath routing, with the objective of minimizing video distortion. For network-centric multipath routing, we implement two popular algorithms:  $k$ -shortest path routing (with  $k = 2$ : 2-SP) [7] and disjoint path routing (with  $k = 2$ : Disjoint Pathset Selection Protocol [DPSP]) [8].

Among the existing methods to generate MD video, we choose a popular time-domain partitioning coding scheme [2], which generates two descriptions by separating the even- and odd-numbered frames and codes them separately. At the receiver side, we measure video quality by computing peak signal-to-noise ratios (PSNRs) for received video frames. The quarter common intermediate format (QCIF) [176 × 144 Y pixels/frame, 88 × 72 Cb/Cr pixels/frame] Foreman sequence is used in the simulations with a frame rate of 15 frames/s.

The PSNR curves of the received video frames are plotted in Fig. 1. We observe that the PSNR curve obtained by multimedia-centric routing (using a GA) is well above those obtained by the two network-centric routing approaches. The improvements in the average PSNR value over 2-SP and DPSP are 6.29 dB and 4.06 dB, respectively. These are considerable improvements in video quality, since usually a 0.5 dB difference is noticeable during playout. To further illustrate how the differences translate into perceived video quality, we plot the original Frame 224 and those delivered by different schemes in Fig. 2.

*Comparison of Metaheuristic Methods* — We next examine the performance of GA-, SA-, and TS-based approaches. In Fig. 3 we plot the evolution of distortion values obtained by the three algorithms for a 50-node network. All three algorithms are terminated after running for 1 s. We observe that the GA curve converges very quickly, while both SA and TS are "trapped" at local optima. Consequently, the minimum distortion values obtained by GA-based routing are evidently better than those obtained by the SA- and TS-based approaches. The slight increases in the SA and TS curves (e.g., the SA

curve at 0.1 s) indicate that a nonimproving solution is accepted in order to get out of a local minimum. An interesting observation is that for GA-based routing, the biggest improvement in distortion is achieved in the first few iterations, while the improvement gets smaller as GA evolves through more generations. We also plot the lower bounds derived using ALG-LB in the figure, which provides a good estimate of the global optimum.

In addition to providing better solutions, another strength of the GA-based approach is that it can find multiple "good" solutions (in addition to those required by the MD video) after a single run due to the nature of a population-based algorithm. This set of good paths can serve as backup paths to reduce rerouting delay and routing overhead.

### Distributed Protocols

As discussed earlier, it is very important to develop efficient distributed implementations for the proposed multimedia-centric routing. Our approach is to incorporate several proven ideas from existing routing protocols for wireless mesh or ad hoc networks and augment these protocols with the cross-layer routing engine.

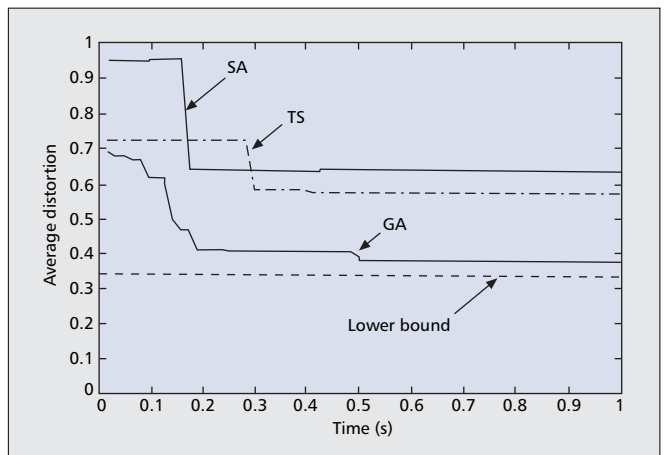
We believe that the proposed multimedia-centric multipath routing is suitable for the proactive routing paradigm. This choice is also motivated by

the need for a routing decision to be made quickly for a new request in order to reduce response time in mission-critical applications. If for any reason proactive routing protocols are not available, the multimedia-centric routing algorithm could also be incorporated into reactive routing protocols for video communications.


### Proactive Routing

The core of proactive protocols is how to build and maintain network topology and link statistic databases at each node. To this end, we find that some components in the class of link-state-based ad hoc routing protocols, such as Optimized Link State Routing (OLSR) [9] and its QoS extensions, are quite suitable for this purpose. In the remainder of this section we provide a sketch of a possible distributed implementation based on proactive routing protocols.

With the proposed implementation model, each node



■ Figure 3. Comparison of distortion evolutions of the three metaheuristic methods for a 50-node network.



■ Figure 4. PSNR values of reconstructed video frames obtained from a testbed experiment.

should continuously measure the performance metrics for each of its links, such as bandwidth, loss rate, and delay. As with other link state routing protocols, our proposed implementation periodically broadcasts link state advertisements (e.g., the Topology Control [TC] message in OLSR) to distribute the network topology information and link statistics. To reduce overhead, we could use the multipoint relay (MPR) technique (Internet Engineering Task Force RFC 3626) [9] to minimize the flooding of control messages. With MPR, each node in the network selects a subset of its one-hop neighbors for forwarding its packets, such that:

- All of its two-hop neighbors can be covered by the relays.
- The number of retransmissions is minimized.

MPR has been shown to be quite effective at suppressing control traffic overhead.

Link state and topology information learned from received LSAs will be pooled in a *link state database*. Each link item (along with its corresponding metrics) is associated with a sequence number, which is set to the sequence number of the LSA message from which this link item was learned. Therefore, a stale item will be overwritten by a fresh one, making the link state database consistent and up-to-date. When the link state database is available, we can replace the routing engine with the multimedia-centric algorithm to compute a set of routes from this node to every other node in the network.

### Reactive Routing

Under reactive routing protocols (e.g., DSR [10]), the source node broadcasts a route request (RREQ) message to the network for route discovery when a new video request arrives. When the destination node (or other nodes that have cached a route to the destination node) receives the RREQ, it will return route replies (RREPs) to the source node. When the RREPs are forwarded toward the source node, each intermediate node will piggyback the link parameters in the RREPs. After receiving a number of such RREPs (along with parameters of the links the RREPs traversed), the source node could construct a partial view of the network topology, consisting of the links on which the RREPs have been forwarded.

The source node will next execute the multimedia-centric routing algorithm to compute a set of routes from the source to the destination node, using this discovered topology as input. Due to the partial topology, the computed routes may not be as competitive as routes computed when the entire network topology is available. However, we expect that this approach will still achieve significant gains for video quality over a basic reactive routing protocol (e.g., DSR), which only selects paths based on network layer metrics.

### Packet Forwarding

Once the near-optimal paths are computed, packet forwarding could be achieved by source routing, within which forwarding information is carried in the header of the packets. Each intermediate node forwards a packet to the next-hop node as indicated in the packet header. This source routing approach not only simplifies the processing at intermediate routers, but also allows different applications (or routing policies) at the source node. This approach is particularly suitable for small and medium-sized networks. When network size increases, the overhead incurred by carrying paths in the packet header could be high. To address this problem, the so-called soft flow states technique could be used. With soft flow states, only the first packet carries the full path information, which establishes *soft flow state* in the intermediate nodes. Subsequent packets could then be forwarded along the same route without the need to carry path information in the headers. At an intermediate node, the flow state is refreshed every time a packet belonging to that flow is forwarded, and will be deleted if there is no such packet received when the timer expires.

### Testbed Experiments

We have developed a testbed network that offers a demonstration of MD video in wireless mesh networks. Our testbed network consists of 12 computers, each equipped with a wireless interface. Network topology and the corresponding link parameters implemented at the application layer are specified by a user-specified script. The GA-based routing discussed earlier was implemented by extending the OLSR protocol. For video coding, we developed a forward error correction-double description (FEC-DD) codec by extending the H.263+ codec [2]. Results obtained from the testbed experiment is presented in Fig. 4, where the PSNRs of reconstructed frames delivered by the proposed algorithm and the traditional 2-disjoint path routing algorithm are plotted. A video trace from the movie *True Lies* was used for this experiment. We observe that the video frames delivered by the proposed algorithm have much higher PSNR values than those delivered by the traditional algorithm. Observation of perceived video quality at the receiver is consistent with simulation results.

### Concluding Remarks

We address the problem of optimal multipath routing for MD video in wireless mesh networks. We show that such problems can be formulated through a multimedia-centric cross-layer optimization approach, and that a genetic algorithm solution procedure is suitable to address such problems that involve complex objective functions and exponentially sized discrete search space. Such a multimedia-centric routing technique is amenable to distributed implementation.

Although this article focuses on video over wireless mesh networks, the proposed framework is quite general and can be applied to solve cross-layer optimization problems that arise in other wireless networking contexts. We focus on unicast video applications in this article. Group communications such as multicasting video to a group of users or concurrent streaming video from a selected set of servers would be important topics for future research.

### Acknowledgment

The authors wish to thank Xiaolin Cheng, Sastry Kompella, and Xiaojun Wang for obtaining some of the results in this article.

---

## References

- [1] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless Mesh Networks: A Survey," *Elsevier Comp. Networks*, vol. 47, no. 4, Mar. 2005, pp. 445–87.
- [2] Y. Wang, A. R. Reibman, and S. Lin, "Multiple Description Coding for Video Delivery," *Proc. IEEE*, vol. 93, no. 1, Jan. 2005, pp. 57–70.
- [3] S. Mao *et al.*, "On Routing for Multiple Description Video Over Wireless Ad Hoc Networks," *IEEE Trans. Multimedia*, vol. 8, no. 5, Oct. 2006, pp. 1063–74.
- [4] M. Alasti *et al.*, "Multiple Description Coding in Networks with Congestion Problem," *IEEE Trans. Info. Theory*, vol. 47, no. 3, Mar. 2001, pp. 891–902.
- [5] L. Ozarow, "On a Source Coding Problem with Two Channels and Three Receivers," *Bell Sys. Tech. J.*, vol. 59, no. 10, Dec. 1980, pp. 84–91.
- [6] C. Blum and A. Roli, "Metaheuristics in Combinatorial Optimization: Overview and Conceptual Comparison," *ACM Comp. Surveys*, vol. 35, no. 3, Sept. 2003, pp. 268–308.
- [7] D. Eppstein, "Finding the k Shortest Paths," *SIAM J. Comp.*, vol. 28, no. 2, Aug. 1999, pp. 652–73.
- [8] P. Papadimitratos, Z. J. Haas, and E. G. Sirer, "Path Set Selection in Mobile Ad Hoc Networks," *Proc. ACM MobitHoc '02*, Lausanne, Switzerland, June 2002, pp. 1–11.
- [9] T. Clausen and P. Jacquet, "Optimized Link State Routing Protocol," IETF RFC 3626, Oct. 2003.
- [10] D. B. Johnson, Y. Hu, and D. A. Maltz, "The Dynamic Source Routing Protocol (DSR) for Mobile Ad Hoc Networks for IPv4," IETF RFC 4728, Feb. 2007.

## Biographies

SHIWEN MAO [S'99, M'04] (smao@ieee.org) received a Ph.D. degree in electrical and computer engineering from Polytechnic University, Brooklyn, New York, in 2004. He was a research scientist at Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, and currently is an assistant professor in the Department of Electrical and Computer Engineering, Auburn University, Alabama. His research interests include cross-layer optimization in multihop wireless

networks and multimedia communications. He was a recipient of the 2004 IEEE Communications Society Leonard G. Abraham Prize Paper Award in the field of communications systems.

Y. THOMAS HOU [S'91, M'98, SM'04] (thou@vt.edu) is an associate professor of electrical and computer engineering at Virginia Tech. His current research interests are cognitive radio networks, optimization and algorithm design for wireless ad hoc and sensor networks, and video communications over dynamic ad hoc networks. He is a recipient of an Office of Naval Research (ONR) Young Investigator Award (2003) and a National Science Foundation (NSF) CAREER Award (2004).

HANIF D. SHERALI (hanifs@vt.edu) is a University Distinguished Professor and the W. Thomas Rice Chaired Professor of Engineering in the Industrial and Systems Engineering Department at Virginia Tech. His research areas are analyzing problems and designing algorithms for specially structured linear, nonlinear, and integer programs arising in various applications, global optimization methods for non-convex programming problems, location and transportation theory and applications, and economic and energy mathematical modeling and analysis. He has published over 238 refereed articles in various operations research journals, has (co-) authored six books in this area, and serves on the editorial boards of eight journals. He is an elected member of the National Academy of Engineering.

SCOTT F. MIDKIFF [S'82, M'85, SM'92] (smidkiff@nsf.gov) is a professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech. He is currently on assignment as a program director in the Electrical, Communications and Cyber Systems Division of the NSF. He received his B.S.E. and Ph.D. from Duke University and his M.S. from Stanford University, all in electrical engineering. His research interests include system issues in wireless and ad hoc networks, network services for pervasive computing, and performance modeling of mobile ad hoc networks.