

MARIA: Interference-Aware Admission Control and QoS Routing in Wireless Mesh Networks

Xiaolin Cheng[†] Prasant Mohapatra[†] Sung-Ju Lee[‡] Sujata Banerjee[‡]

[†]Department of Computer Science, University of California at Davis, CA 95616

[‡]Mobile & Media Systems Lab, Hewlett-Packard Laboratories, Palo Alto, CA 94304

[†]{xlcheng, pmohapatra}@ucdavis.edu, [‡]{sjlee, sujata.banerjee}@hp.com

Abstract—Interference among concurrent transmissions complicates QoS provisioning for multimedia applications in wireless mesh networks. In this paper we propose MARIA (Mesh Admission control and qos Routing with Interference Awareness), a scheme towards enhancing QoS support for multimedia in wireless mesh networks. We characterize interference in wireless networks using a conflict graph based model. Nodes exchange their flow information periodically and compute their available residual bandwidth based on the local maximal clique constraints. Admission decision is made based on the residual bandwidth at each node. We implement an on-demand routing scheme that explicitly incorporates the interference model in the route discovery process. It directs routing message propagations and avoids “hot-spots” with severe interference. Simulation results demonstrate that by taking interference into account MARIA outperforms the conventional approach. It finds routes with less interference and enhances the performance significantly. We use video as an example application and MARIA improves the quality of delivered videos, with up to 7.3 dB average PSNR gain.

I. INTRODUCTION

Content-rich multimedia services are increasingly expected in wireless mesh networks. Unlike traditional data applications, multimedia need QoS support to satisfy user service requirements. However, shared broadcast media and limited resources make it difficult to design efficient QoS solutions for multimedia applications in wireless networks. Interference exists among flows and even links of an individual flow, which has a great impact on the performance of resource-consuming multimedia applications.

Most existing link layer and networking layer protocols are not adequate to support QoS for multimedia. Traditional ad hoc network routing protocols do not support QoS and need significant extensions to incorporate QoS requirements [1]. In the MAC layer, it has been shown that distributed QoS mechanisms are difficult for IEEE 802.11 [2]. Most QoS schemes for IEEE 802.11 are priority assignment and fair scheduling, which provide QoS differentiation but without guarantees of QoS levels. IEEE 802.11e [3] supports a priority mechanism that can be viewed as a partial solution for providing QoS in single-hop wireless LANs. However, when extended to multi-hop scenarios, it does not work effectively [4].

In order to guarantee QoS requirements of the existing flows in the network, policing mechanisms need to be implemented to prevent the incoming flows from consuming excessive resource and degrading performance of the existing flows. To this end, effective admission control is crucial to

maintain overall QoS satisfaction in the network. However the distributed nature of admission control in wireless networks makes the task complicated. In addition, interference in shared media magnifies the difficulty of this problem. When a node admits a flow, it must consider the interference from the existing flows in its interference neighborhood. Similarly, admission of the incoming flow may affect the existing flows. Consequently, it is not sufficient to consider only the local resource availability at a node when admitting a new flow. We need to take into account the resource availability of all its interference neighboring nodes. To make proper admission decision at the end-nodes in the network, routing protocols must incorporate QoS requirements of flows. Many existing routing schemes strive for the shortest path and do not explicitly provide QoS support. To design an efficient distributed admission control solution, we need QoS routing support to discover routes with less interference from the source to the destination.

We study the important problem of admission control coupled with QoS extensions in routing for video streaming applications in wireless mesh networks, which is the most complex form of multimedia and has stringent QoS requirements. In our MARIA scheme, we use a conflict graph to model interference in wireless networks. Both inter- and intra-flow interference are considered. A distributed hop-by-hop admission control works with the on-demand route discovery. When a new flow request arrives, each node involved in the route discovery makes its admission decision based on residual bandwidth in its interfering neighborhood. The residual bandwidth is computed by identifying the maximal clique constraints in its local conflict graph. Since interference is considered when admitting new flows, “hot spots” with large interference are avoided and load balancing is achieved. We reserve a portion of channel capacity for random short-lived traffic (such as web traffic) and other best-effort traffic to avoid the starvation of non-QoS traffic.

The rest of the paper is organized as follows. In Section II, we present the conflict graph based interference model used in MARIA. Section III describes how the residual bandwidth is computed. The detailed design of interference-aware admission control and QoS routing in MARIA is presented in Section IV. Simulation results are discussed and analyzed in Section V. We discuss the related work in Section VI. Section VII concludes the paper.

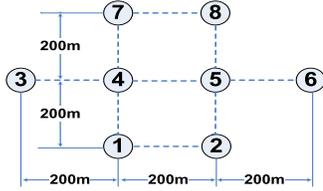


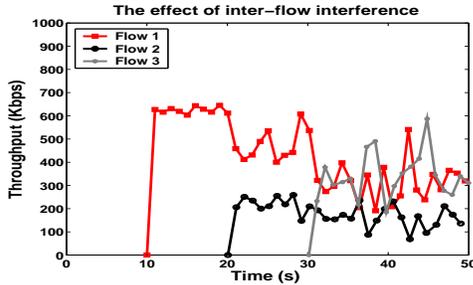
Fig. 1. A simple example network.

II. INTERFERENCE IN WIRELESS NETWORKS

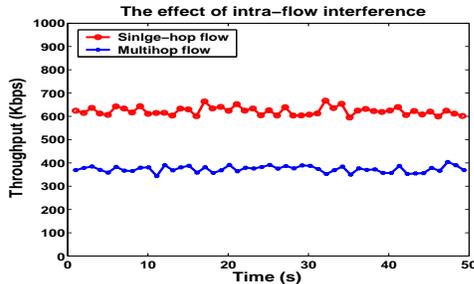
A. Inter- and Intra-flow Interference

In wireless networks, nodes on a flow contend with other simultaneously transmitting nodes for resources in their interference neighborhoods. In addition, transmissions over links of a multi-hop flow may interfere with each other. As the end-to-end path length increases, this intra-flow interference problem becomes more severe.

We use a simple network shown in Figure 1 to demonstrate how interference degrades network performance. In the first example, three CBR flows sequentially start from node 1 to node 2 (Flow 1), from node 3 to node 6 (Flow 2), and from node 7 to node 8 (Flow 3). The rate of three flows are all 600 Kbps. Channel capacity is 2 Mbps. Since node 3 and node 6 are beyond the transmission range of each other, Flow 2 must hop through node 4 and node 5. Figure 2(a) shows the throughput of three flows. When Flow 2 starts, the throughput of Flow 1 drops significantly. When Flow 3 starts, both Flow 1 and Flow 2 sacrifice part of their throughput. All three flows suffer large throughput fluctuations. This scenario shows the adverse impact of inter-flow interference.



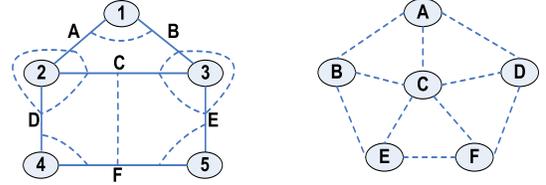
(a) Inter-flow interference among 3 flows.



(b) Intra-flow interference inside a multi-hop flow.

Fig. 2. The effect of interference in wireless networks.

In the next example, we separately measure the throughput when only a single-hop flow goes from node 3 to node 4 with the rate of 600 Kbps, and when only a multi-hop flow goes from node 3 to node 6 with the same rate. As shown in Fig 2(b), the throughput gap between these two flows is



(a) Connectivity graph.

(b) Conflict graph.

Fig. 3. Conflict graph model.

fairly large. The single-hop flow achieves its preset rate of 600 Kbps while the throughput of the multi-hop flow is less than 400 Kbps. Intra-flow interference among the links of the multi-hop flow deteriorates its throughput.

B. Interference Model

Since interference in wireless networks degrades the network performance, it is important to properly define and model interference. To define interference, we apply the well-known protocol model proposed in [5]. Conflict graph [6][7] has been shown to be an effective model to characterize interference in wireless networks. In this model, a wireless network is abstracted into a graph $\mathcal{G}(V, E)$, in which V is the set of nodes and E is the set of links. If two nodes lie within the transmission range of each other, there is a link between them. Each link in the original connectivity graph \mathcal{G} is represented by a node in its corresponding conflict graph \mathcal{C} . A link exists in the conflict graph \mathcal{C} if two links in its connectivity graph \mathcal{G} interfere with each other. In a conflict graph, the aggregated rate of flows carried on a particular link must satisfy its maximal clique constraints. A clique is an induced complete subgraph from its original conflict graph. All nodes in a clique are pairwise connected, which means all links represented by these nodes interfere with one another. During a time slot, only one link in a clique can be active and carry traffic.

A maximal clique is a clique not contained by any other cliques. In the conflict graph based interference model, the residual bandwidth that a link can support needs to satisfy all maximal-clique constraints to which this link belongs. Figure 3 gives an example of the conflict graph based model. In this example, Figure 3(a) is the connectivity graph and Figure 3(b) is its conflict graph. Solid lines in Figure 3(a) mean two nodes are connected by a link; dashed lines indicate that two links interfere with each other. In the conflict graph shown in Figure 3(b), we can derive the maximal clique constraints of all links. For instance, $\mathcal{C}_1 = \{A, B, C\}$ and $\mathcal{C}_2 = \{A, C, D\}$ are two example cliques, which are the maximal cliques as well. Node A (link A in the connectivity graph) belongs to both maximal cliques, so it must satisfy the constraints imposed by clique \mathcal{C}_1 and \mathcal{C}_2 as follows:

$$F_A + F_B + F_C \leq R$$

$$F_A + F_C + F_D \leq R$$

where F_A, F_B, F_C and F_D are the aggregated rates of flows on link A, B, C , and D , respectively, and R is the channel capacity.

Each node maintains the information of flows (links) in its interference neighborhood. The maximal cliques can be

computed locally. In our admission control approach, each node on the path constructs its local conflict graph and checks if the maximal clique constraints are satisfied. We assume that efficient scheduling is done at the MAC layer for the admitted flows. We are more concerned with the availability of route(s) with less interference. Satisfying maximal clique constraints is a sufficient condition for an interference-free scheduling [8].

III. RESIDUAL BANDWIDTH COMPUTATION

A. HELLO Message Exchange

To make the local admission decision, each node must be aware of the contending flows in its interference neighborhood. In MARIA, nodes exchange this information periodically by broadcasting HELLO messages that carry the information of the flows traversing these nodes. Each node sets a sliding time window to estimate dynamic transmission rate of each flow. After a node receives HELLO messages from all its neighbors, it can build its local conflict graph. In this way, each node maintains the information about the flows it forwards.

The HELLO messages are broadcast with an extended transmission range to incorporate the flow information in the interference neighborhood. Based on HELLO messages, a local conflict graph is constructed, and the local maximal cliques are computed. Residual bandwidth is then estimated, which is the basis of our admission control. HELLO messages introduce more overhead in the network, and hence their exchange frequency should be moderate. Based on our simulations, we set this value to 2 seconds.

B. Approximate Maximal Cliques

At each intermediate node, the key problem is to compute residual bandwidth constrained by its local maximal cliques. We approximate the interference area of a link as a circle centered at the middle-point of this link with a radius of interference range plus half length of this link, as shown is Figure 4(a). Finding maximal cliques is a NP-complete problem [9]. To compute maximal cliques, we apply a heuristic similar to the one proposed in [10]. Figure 4(b) illustrates the basic idea of this heuristic. A circle (in dotted line) with a diameter of interference range scans the interference area (the circle in solid line with a radius of interference range) of a link SD , the maximal cliques to which link SD belongs can be identified. Computation of maximal cliques is performed only during the route discovery and admission control phase.

IV. INTERFERENCE-AWARE ADMISSION CONTROL AND QOS ROUTING

Using the conflict graph based interference model and HELLO message exchanges, each node determines the flow information in its interference neighborhood. We now describe how nodes in MARIA make local admission control decision with the support from QoS-aware routing.

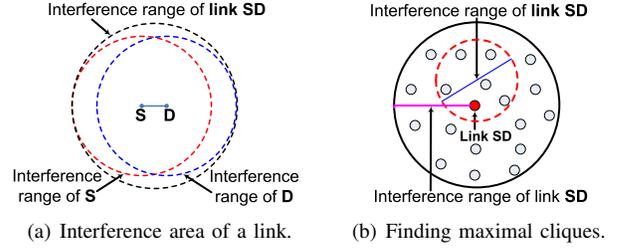


Fig. 4. Computation of maximal cliques.

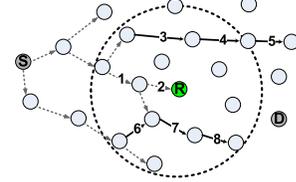


Fig. 5. RREQ propagation and local conflict graph construction.

A. Route Discovery and Admission Control

Route discovery finds an end-to-end path that has sufficient resource and little interference for an incoming flow. To facilitate the local conflict graph construction, when route request (RREQ) messages are broadcast in the network, they contain the bandwidth requirement of the requesting flow and the link information through which they propagate.

When a source needs a route to initiate a flow, it checks its residual bandwidth. If it is greater than the bandwidth requirement of the flow, it starts broadcasting RREQ. Otherwise, the flow is rejected due to insufficient resource at the source. When an intermediate node receives a non-duplicate RREQ, it assumes the links on the partial route from the source to itself are carrying a flow with the required rate specified in the RREQ. Since RREQ contains the information about the links through which it has been forwarded, this node can build up a link “pool” which consists of both the links of the pending flow accumulated in RREQ and the links of the neighboring flows which are obtained via HELLO message exchanges. Consequently, this node constructs a local conflict graph and computes its updated residual bandwidth. If this residual bandwidth is greater than the required bandwidth, admission at this node is successful. It thus keeps broadcasting the RREQ message. Otherwise, admission at this node fails and it stops broadcasting RREQ.

Figure 5 shows an example of RREQ propagation and the local conflict graph construction. Node S has a flow request to node D . It broadcasts the RREQ (dotted links) and one copy reaches an intermediate node R . The large dotted circle is the interference neighborhood of R . Node R knows two existing flows (solid links) in this range, which are (3-4-5) and (6-7-8) respectively. It combines the links of the partial route accumulated in RREQ and the links of two flows in its interference range, and builds a link “pool” as (1, 2, 3, 4, 6, 7, 8). It computes the corresponding maximal clique constraints to determine whether the RREQ should be broadcast.

When the destination receives RREQs, it selects the best route (described in Section IV-B) and sends the route reply (RREP) via the chosen path. RREP has the complete informa-

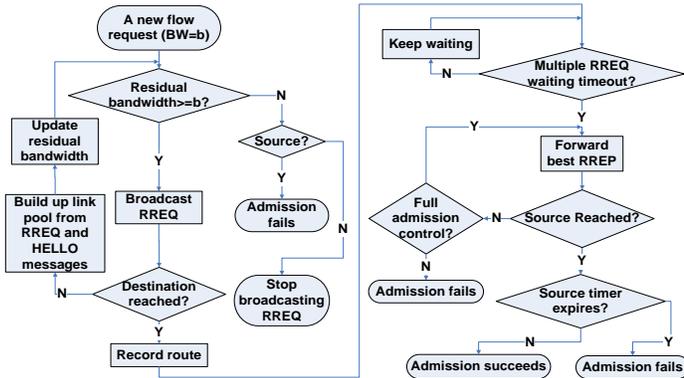


Fig. 6. Flowchart of route discovery and admission control in MARIA.

tion of the route when it traverses back to the source. When an intermediate node receives the RREP, it adds all links of this route that are in its interference range into its link “pool.” Since the entire route from the source to the destination is known to any forwarding node of the RREP message, the complete admission control can be finished at this stage. This intermediate node builds up its local conflict graph based on the link pool and computes the maximal cliques. If the maximal clique constraints are satisfied, it sends the RREP to the next hop. Otherwise, the forwarding of RREP stops.

When a node makes an admission control decision, it considers the existing flows in its interference neighborhood. In addition, interference among the individual links of the requesting flow is taken into account. Consequently, both inter- and intra-flow interference are considered in the admission control process.

B. Route Selection

RREQ carries the minimum residual bandwidth of the nodes it goes through. When multiple RREQs arrive at the destination, the best route is chosen and a RREP is sent back to the source. Route selection criterion is the highest minimum residual bandwidth, i.e., the least interference. The source sets a timeout when it starts broadcasting an RREQ. When the timer expires without receiving any RREP, the source assumes that the admission request failed and rejects the requesting flow. Figure 6 is the flowchart of admission control and route discovery operations in MARIA when a new flow request arrives.

When a node decides whether to broadcast RREQ or not, it avoids routes with large inter- and intra-flow interference. A route with least interference and best bandwidth availability is selected. Therefore heavy-loaded “hot-spots” in the network are circumvented. Higher network utilization and load balancing are hence achieved.

V. PERFORMANCE EVALUATION

We conduct simulations to demonstrate the effectiveness of MARIA over the conventional AODV [11]. Compared to other routing protocols, AODV is shown to be effective at higher network loads [12], which is desirable for multimedia applications.

We use a detailed simulation model based on ns-2 with wireless support from the Monarch research group [13]. A shared-media radio with a data rate of 2 Mbps and a radio range of 250 meters are used. We use this rate so that we can easily saturate the network with fewer flows and background traffic and create a challenging network environment. The data packet size of CBR flows is 512 bytes.

A. Random Network Topology

In this experiment, we generate a 20-node network in a 1000m×1000m area. Seven flows are initiated with the start times and transmission rates shown in Table I. From Figures 7(a) and 7(c) (delay is log-scaled) we can see that both throughput and delay for AODV degrade as more flows are accepted. In particular, after Flow 3 and Flow 4 enter the network, the throughput suffers large fluctuations. Throughput of Flow 3 drops to half after Flow 4 enters. End-to-end delays of Flow 3 and Flow 4 keep increasing after Flow 4’s admission. However, with admission control in MARIA, Flow 4 is rejected and that enables all other flows to achieve better performances as shown in Figures 7(b) and 7(d). Without interference from Flow 4, the throughput of other flows is stabilized and maintains the rate as preset. End-to-end delays are much lower (less than 50 ms) compared with AODV.

TABLE I
SEVEN FLOWS IN A RANDOM NETWORK TOPOLOGY.

	Flow1	Flow2	Flow3	Flow4	Flow5	Flow6	Flow7
Start Time (s)	10	20	30	40	50	60	70
Rate (Kbps)	100	200	100	100	200	150	100

B. Grid Network Topology

We now setup a 5×5 grid network in a 1000m×1000m area. Five sequential flows with the parameters listed in Table II are used. The performance comparisons are presented in Figure 8. In AODV, Flow 4 has a large routing delay. After it has entered the system, the throughput of Flows 1, 2 and 3 reduces dramatically. Flow 4 itself also yields low throughput. End-to-end delays of the flows increase as new flows are entered to the network. The delay of Flow 4 is the largest, around 3.5 seconds. In contrast, in MARIA Flow 4 is not admitted to the network, which results in guaranteeing the performance of the other four flows. The throughput meets the required data rate and the delay is shorter than AODV.

TABLE II
FIVE FLOWS IN A GRID NETWORK TOPOLOGY.

	Flow1	Flow2	Flow3	Flow4	Flow5
Start Time (s)	10	20	30	40	50
Rate (Kbps)	200	100	200	200	350

C. QoS Routing

In this experiment, we explicitly demonstrate the QoS routing ability of MARIA. The goal of QoS routing is to find a route that brings less interference and supports the required bandwidth. We use a 5×5 grid network in this simulation. Two flows are injected into the network. Flow 1 starts at 10 seconds with the rate of 200 Kbps and Flow 2 begins at 20 seconds with 300 Kbps rate.

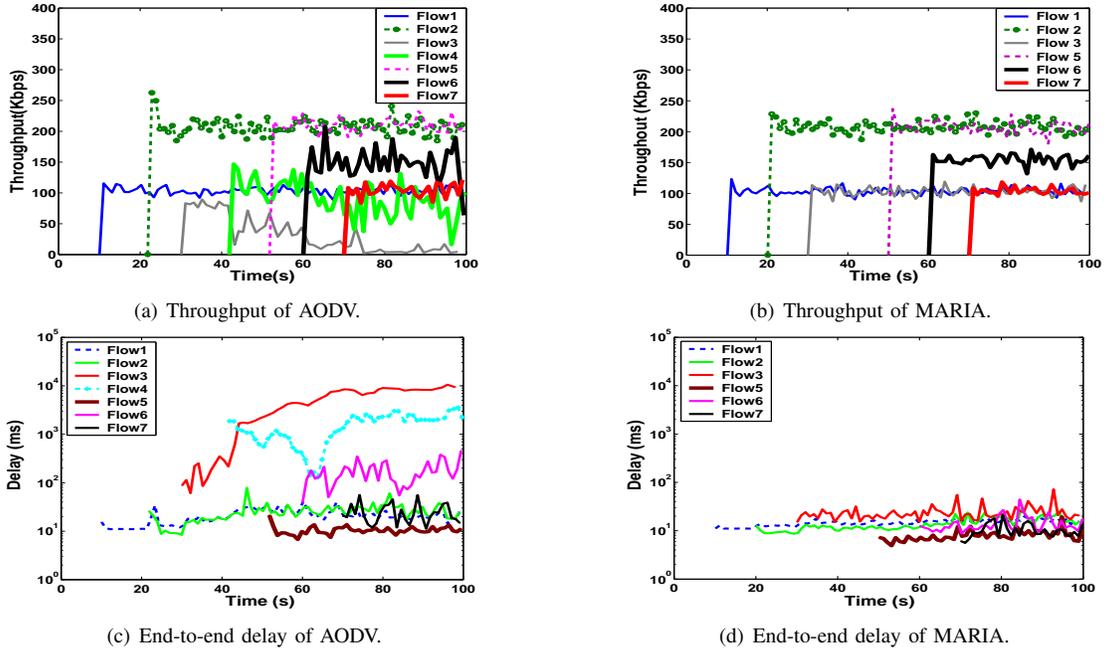


Fig. 7. Performance of seven flows in a random network topology.

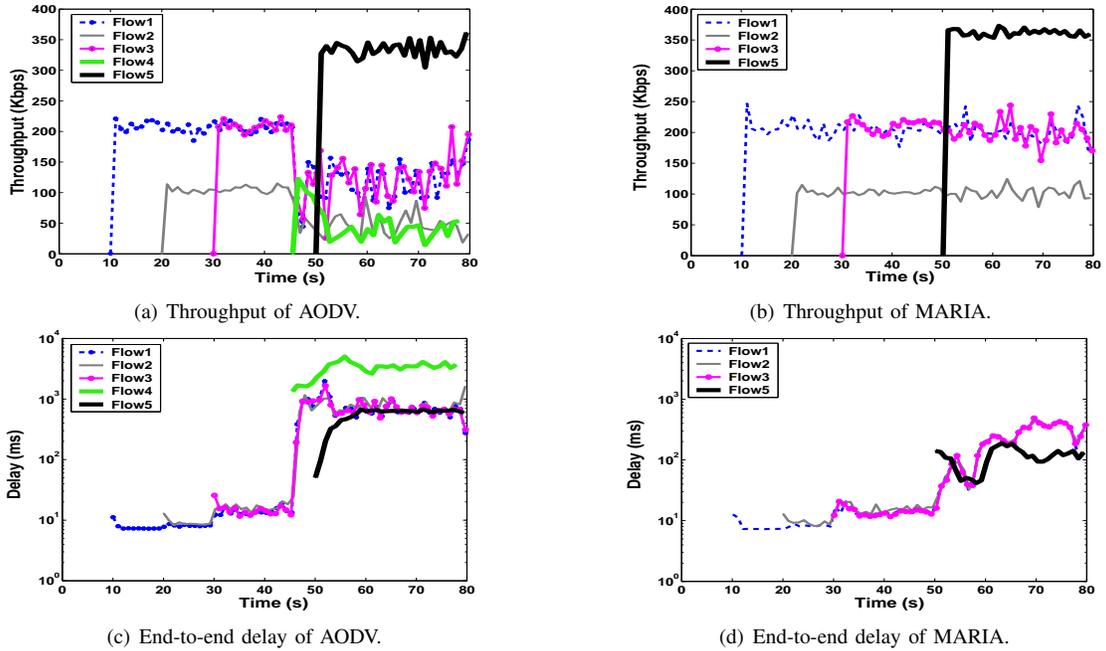
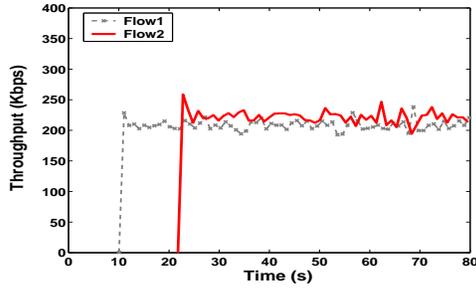


Fig. 8. Performance of five flows in a grid network topology.

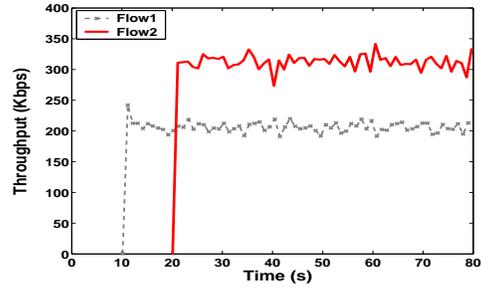
In both schemes two flows are admitted (so no admission control issue here). However, it is observed in Figures 9(a) and 9(c) that due to the interference from Flow 1, the throughput of Flow 2 drops sharply right after it enters and stays below the required bandwidth and the end-to-end delay is much larger than that of Flow 1. Powered by the enhanced QoS route discovery, MARIA discovers a route for Flow 2 which has less interference with Flow 1. Therefore, the required transmission rate is achieved, and the delay performance is much better (as shown in Figures 9(b) and 9(d)).

D. Video Delivery

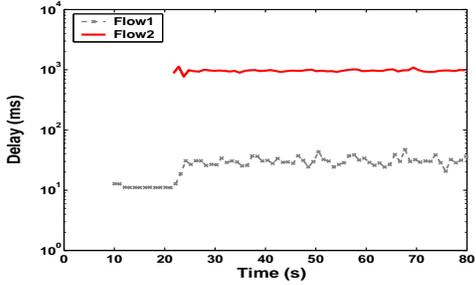
We incorporate real video traces into the ns2 simulation to evaluate the video delivery performance of MARIA. We apply *ffmpeg* [14] MPEG4 codec to compress the standard 400-frame *foreman* QCIF (176 × 144) video clip (frame rate 30fps) and generate corresponding video trace with a UDP packetization of 1000 bytes. We compare both *PSNR* (Peak Signal-to-Noise Ratio, the most widely used objective video quality metric) values and actual decoded images at the receiver end of two schemes.



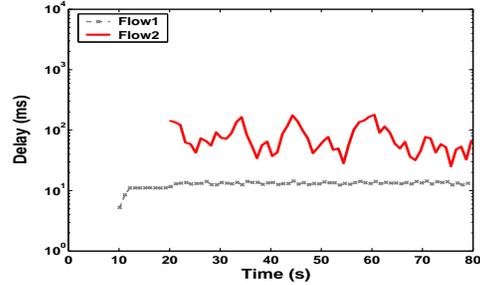
(a) Throughput of AODV.



(b) Throughput of MARIA.



(c) End-to-end delay of AODV.



(d) End-to-end delay of MARIA.

Fig. 9. Performance of two flows in a grid network topology—QoS routing ability.

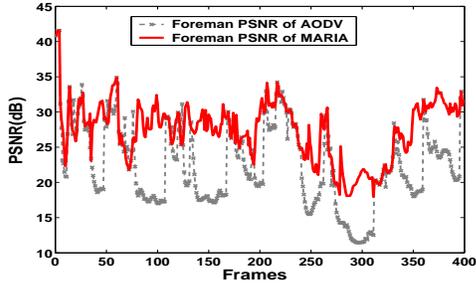


Fig. 10. PSNR comparison—Impact of admission control.



(a) AODV (b) MARIA

Fig. 11. Decoded images of Frame 97.

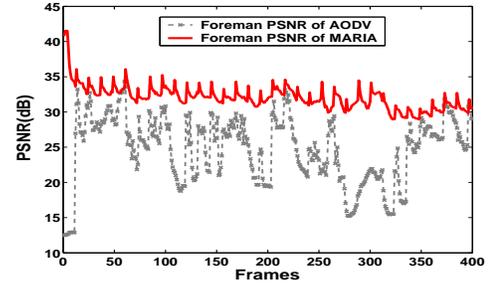


Fig. 12. PSNR comparison—Impact of QoS routing.



(a) AODV (b) MARIA

Fig. 13. Decoded images of Frame 229.

The network setting is the 5×5 grid network with 200m spacing between any two horizontally and vertically adjacent nodes. We first examine the interference-aware admission control ability of MARIA. We code the *foreman* video clip into a 100Kbps MPEG4 stream. Two CBR flows with the rates of 200 Kbps and 100 Kbps are used for the background traffic. Three identical *foreman* streams start between three pairs of nodes sequentially. The second stream incurs significant interference with the first stream and the background traffic. Therefore it is rejected by MARIA as sufficient resource is not available to accept this stream and maintain its quality. However, without considering the interference and QoS requirement, AODV admits all three streams, which causes interference and severely degrades the video quality of the subsequent (third) stream. In Figure 10, we compare the PSNR

of the third stream of MARIA with AODV. Clearly, for most frames, MARIA produces much higher PSNR. The average PSNR in MARIA is 27.4 dB, compared to 22.5 dB given by AODV. The actual image quality of Frame 97 is compared in Figure 11.

In the next set of simulation, we evaluate the interference-aware QoS routing capability of MARIA using the *foreman* video trace. We use two CBR flows, each with 100 Kbps rate, as the background traffic. We code the *foreman* clip into a 150Kbps MPEG4 stream. Two identical video streams of this clip start transmission sequentially. Both schemes admit two video streams (so no admission control issue here). In our scheme, a better route is found for the second stream which causes little interference with the first stream and the background CBR flows, so the PSNR of the second stream

is much higher than in AODV, as shown in Figure 12. The average PSNR of MARIA is 32.0 dB while it is only 24.7 dB from AODV. The actual image quality of decoded video stream (Frame 229) is also substantially improved, as illustrated in Figure 13.

E. Overhead Analysis

We analyze the overhead of MARIA and compare it with AODV. We apply our overhead analysis to the second scenario with the video trace in the previous section.

Table III shows the overhead comparison between the two schemes. We find that AODV incurs a high number of routing packets. Our admission control reduces the RREQ transmission overhead. Paths with high interference found by AODV have less number of data packets delivered to the destination. MARIA incurs extra overhead from the HELLO message exchange. The number of transmitted HELLO messages is 412. Consequently the ratio of the total number of overhead packets transmitted to the number of received data packets in MARIA (38.0%) is higher than AODV (20.1%). Considering the performance gain we achieved, this overhead is worth the tradeoff.

TABLE III

OVERHEAD ANALYSIS.

	Routing pkts	HELLO msgs	Received data pkts
AODV	335		1666
MARIA	281	412	1822

VI. RELATED WORK

QoS in wireless networks has been an active research area under intensive investigations in recent years. CACP (Contention-aware Admission Control Protocol) [15] introduced a *c-neighbor* concept (nodes in carrier-sensing range) to characterize contention in wireless networks. Information about c-neighbors is obtained through multi-hop querying packets or querying packets sent with increased transmission power. Node makes admission decision based on its c-neighbor available bandwidth which is the smallest local available bandwidth of all of its c-neighbors. In CACP the on-demand querying packets are crucial to effective admission control. The losses of these packets may lead to inaccurate and unreliable admission decisions.

PAC (Perceptive Admission Control) [16] addresses the admission control problem by monitoring the wireless channel using *channel busy time*, and dynamically adapting admission control decisions to enable high network utilization while preventing congestion. However, this protocol does not consider intra-flow interference when making admission decisions.

Conflict graph has been widely used for modeling wireless interference. Using a conflict graph based model, methods for computing upper and lower bounds on the optimal throughput for the given network and workload are presented in [6]. An interference-aware channel assignment algorithm for multi-radio wireless mesh networks is presented in [7]. It extends the conflict graph model to capture interference between routers in multi-radio scenarios. A heuristic is proposed in [10] to

compute the maximal cliques in a conflict graph which is exploited in our scheme.

VII. CONCLUDING REMARKS

We investigated admission control and QoS routing in wireless mesh networks, and proposed a MARIA scheme to support QoS for multimedia applications. We use a conflict graph model to characterize both inter- and intra-flow interference. The available residual bandwidth is computed based on the maximal clique constraints in its local conflict graph to make distributed hop-by-hop admission control decision. Our simulation results show that with admission control and QoS routing support, MARIA outperforms the conventional protocol. The proposed scheme discovers routes with less interference, and enables better network utilization and high quality video delivery. In our scheme, we assumed a distance-based model with fixed channel capacity for its simplicity and ease of implementation. It is, however, not limited to this assumption. Dynamic capacity can be plugged into the conflict graph model. We are investigating a measurement-based approach that accommodates varying channel capacity and captures interference more accurately.

REFERENCES

- [1] I. Jawhar and J. Wu, *Resource Management in Wireless Networking*. Kluwer Academic Publishers, 2004, ch. Quality of Service Routing in Mobile Ad hoc Networks.
- [2] W. Pattara-atikom, P. Krishnamurthy, and S. Banerjee, "Distributed Mechanisms for Quality of Service in Wireless LANs," *IEEE Wireless Communications Magazine*, vol. 10, no. 3, pp. 26–34, June 2003.
- [3] "IEEE 802.11e/D4.0, Draft Supplement to Part 11: Wireless Medium Access Control (MAC) and physical layer (PHY) specifications: Medium Access Control (MAC) Enhancements for Quality of Service (QoS), November 2002."
- [4] J. Li and P. Mohapatra, "APHD: End-to-End Delay Assurance in 802.11e Based MANETs," in *Proc. of Mobicom'06*, San Jose, California, July 2006.
- [5] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. on Information Theory*, vol. 46, no. 2, pp. 388–404, March 2000.
- [6] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu, "Impact of Interference on Multi-hop Wireless Network Performance," in *Proc. of Mobicom'03*, San Diego, CA, Sep. 2003.
- [7] K. N. Ramachandran, E. M. Belding-Royer, K. C. Almeroth, and M. M. Buddhikot, "Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks," in *Proc. of Infocom06*, Barcelona, Spain, 2006.
- [8] M. Alicherry, R. Bhatia, and L. Li, "Joint Channel Assignment and Routing for Throughput Optimization in Multiradio Wireless Mesh Networks," in *Proc. of Mobicom'05*, 2005.
- [9] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, 1979.
- [10] R. Gupta and J. Walrand, "Approximating Maximal Cliques in Ad-Hoc Networks," in *Proc. of PIMRC'04*, Barcelona, Spain, Sep. 2004.
- [11] C. E. Perkins, E. M. Belding-Royer, and S. Das, "Ad Hoc On Demand Distance Vector (AODV) Routing, IETF RFC 3561," July 2003.
- [12] C. E. Perkins, E. M. Royer, S. R. Das, and M. K. Marina, "Performance Comparison of Two On-Demand Routing Protocols for Ad Hoc Networks," *IEEE Personal Communications*, vol. 8, no. 1, pp. 16–28, Feb. 2001.
- [13] <http://www.isi.edu/nsnam/ns/>.
- [14] <http://www.tkn.tu-berlin.de/research/evalvid>.
- [15] Y. Yang and R. Kravets, "Contention-Aware Admission Control for Ad Hoc Networks," *IEEE Transactions on Mobile Computing*, vol. 4, no. 4, pp. 363–377, 2005.
- [16] I. D. Chakeres and E. M. Belding-Royer, "PAC: Perceptive Admission Control for Mobile Wireless Networks," in *Proc. of QSHINE '04*, Washington, DC, USA, 2004, pp. 18–26.