

Opportunistic Overlays: Efficient Content Delivery in Mobile Ad Hoc Networks

Yuan Chen and Karsten Schwan

College of Computing, Georgia Institute of Technology, Atlanta GA 30332, USA
{yuanchen, schwan}@cc.gatech.edu

Abstract. Current content-based publish/subscribe systems assume network environments with stable nodes and network topologies. For mobile environments, one resulting problem is a mismatch between static broker topologies and dynamic underlying network topologies. This mismatch will result in inefficiencies in event delivery, especially in mobile ad hoc networks where nodes frequently change their locations. This paper presents a novel middleware approach termed *opportunistic overlays*, and its dynamically reconfigurable support framework to address such inefficiencies introduced by node mobility in publish/subscribe systems. The opportunistic overlay approach dynamically adapts event dissemination structures (i.e., broker overlays) to changes in physical network topology, in nodes' physical locations, and in network node behaviors, with the goal of optimizing end-to-end delays in event delivery. Runtime adaptations include the dynamic construction of broker overlay networks and changes of mobile clients' assignments to brokers. Experimental results demonstrate that the opportunistic overlay approach is practically applicable and that the performance advantages attained from the use of opportunistic overlays can be substantial.

1 Introduction

Publish/subscribe is a widely used method for providing anonymous, inherently asynchronous group communications in distributed settings. Past work has created numerous publish/subscribe systems, in industry and in academia [1, 7]. With the increased availability of powerful mobile computing devices like laptops and PDAs, and the widespread deployment and use of wireless data communications, there is a pressing need to extend such middleware to the mobile computing domain. Moreover, certain features of publish/subscribe make it well-suited to mobile environments, including asynchronous event delivery, anonymity, multipoint communication and content-based routing [8, 10]. Current systems targeting Internet-based communications, however, commonly assume distributed execution environments in which clients do not move and where the network topology remains relatively stable. Stated more technically, they assume statically deployed broker networks (i.e., overlays) mapped to static network topologies. A resulting problem for mobile environments is a mismatch between static broker topologies and dynamic underlying network topologies. This mismatch will result in inefficiencies in event delivery, a simple example being a shortest

path in the original overlay and physical network turning into an inefficient path when the same logical overlay is used with a different physical network topology. Another example is when a node's movement (either a client, or a broker or an intermediate node) affects the underlying network topology and changes the distance between a mobile client and its assigned broker, hence resulting in sub-optimized event delivery.

This paper proposes the *opportunistic overlay* approach to managing overlays for mobile nodes in mobile ad-hoc networks. The idea is to dynamically optimize content-based event delivery by adapting event dissemination structures (i.e., broker overlays) to changes in physical network topology, in nodes' physical locations, and in network node behaviors. The term 'opportunistic' denotes the fact that the solution is one in which each broker opportunistically acts to improve its relations with both other brokers and with its clients. The key points characterizing opportunistic overlays may be summarized as follows: (1) dynamically constructing broker network topologies to match the underlying physical network, (2) dynamically changing a mobile client's broker assignment based on the client's physical location and the broker's current capabilities, and (3) when broker topologies or clients' home brokers change, recalculating overlay routing paths and then using the newly computed paths.

Opportunistic overlays are implemented with the JECho Java-based publish/subscribe infrastructure [7]. A unique attribute of this implementation is that with JECho, dynamic topology adjustments can be coupled with runtime techniques for event filtering, thereby also permitting the system to match event rates and sizes to the currently available levels of bandwidth of physical communication channels.

The performance evaluations reported in this paper use actual hardware, to assess basic performance properties and penalties, and they use emulation and simulation, to assess the effects of mobility and to better understand the scalability of our approach. Results demonstrate that the performance advantages attained from the use of opportunistic overlays can be substantial. For instance, simulation results indicate that the delay of sending a message can be improved by up to 100%. In a set of emulation experiments, the opportunistic overlay approach is able to both optimize path lengths and address broker overloads. Measurements on a small testbed comprised of three laptops running the AODV protocol [11] show more than a sixfold improvements in the end-to-end delays experienced by events in the flood watch application.

The remainder of this paper is organized as follows. We present the system model, protocols, and algorithms used by opportunistic overlays in Section 2. The prototype architecture and some implementation details are discussed in Section 3. Section 4 presents evaluation results. Related work appears in Section 5, followed by conclusions and future work in Section 6.

2 The Opportunistic Overlay Approach

We first outline the system model assumed by opportunistic overlays, followed by descriptions of the adaptation protocols and algorithms underlying the approach.

2.1 System Model

Our system model adopts an overlay network approach. As illustrated in Figure 1, an event system consists of producers, consumers, and a broker network. The latter is an overlay across the physical network, composed of broker processes connected via links. Each overlay link is a network path between a broker node pair in the physical network. Each producer/consumer (mobile client) connects to one of the brokers (usually the nearest one) via one or multiple wireless links. This broker is called the client's *home broker*. A consumer also provides a content-based subscription function termed *modulator*, which operates on event contents to dynamically tailor them to the consumer's current needs. A consumer's modulator executes in an intermediate broker's address space on behalf of the consumer. The intermediate broker can be any broker (typically its home broker) on the overlay path between producers and the consumer. An event generated by a producer is first sent to the producer's home broker, then routed from the producer's home broker to the consumer's home broker, processed using the consumer's modulator, and then delivered to the consumer via some wireless network links.

An event system with four broker nodes (A, B, C and D), one producer M1, and two consumers (M2 and M3) in a wireless ad hoc network is depicted in Figure 2. Since a broker can reside on the same nodes as producers/consumers or on separate nodes, in general, a link in the broker network is a multi-hop wireless path on the underlying physical wireless network. Similarly, a producer/consumer connects to its home broker via a multi-hop wireless path.

2.2 Basic Idea

The idea behind opportunistic overlays is to continually optimize event delivery, by dynamically changing both broker networks and mobile clients' home brokers. Updates occur in response to changes in physical network topology and in nodes'




Fig. 1. System Model




Fig. 2. A Sample Event System in Mobile Ad Hoc Networks

physical locations. Potential broker overloads are avoided by judiciously choosing clients' home brokers. The key points characterizing opportunistic overlays may be summarized as follows:

Resource awareness. An opportunistic overlay is aware of the underlying network topology used for transporting events from producers to consumers. It is also aware of the respective locations of both and of their current state (e.g., CPU Load, Memory availability).

Dynamic construction of broker overlay networks. Dynamic broker network topology construction uses a global state routing protocol [12]. Each broker maintains a local view of the broker network topology. At runtime, an opportunistic overlay dynamically monitors client location, physical network topology, and resources (e.g. latency, bandwidth, broker computation load). Periodically, each broker updates its local view of broker network topology, by changing its neighboring brokers and by propagating changes to its neighbors. Neighbor broker information is acquired by querying the network protocols' routing tables or via neighbor discovery operations [13]. When a broker receives propagated information from its neighbors, it updates its broker topology accordingly.

Dynamic change of home broker. A mobile client periodically checks the brokers in its vicinity via a 'nearest broker search'. It identifies to its home broker a found candidate broker that is closer to its current physical location. Upon receiving such a report from a client, the home broker initiates a broker selection protocol. This protocol uses an approach that combines shortest path selection with load balancing methods. Specifically, the home broker first calculates the path length from the producer to the candidate broker, and then determines whether or not to change the client's home broker based on both the network distance and the candidate broker's current capabilities. Preference is given to the closer broker unless that broker is currently overloaded.

Dynamic overlay routing. Changes in broker network topology and in mobile clients' home brokers will result in rebuilding broker-level routing tables. Opportunistic overlays use source routing for event delivery. Whenever a broker's local view of broker topology changes or whenever a client receiving events from a broker changes its home broker, new event paths are calculated using a shortest path algorithm. We next discuss some of these protocols in more detail.

2.3 Dynamic Construction of Broker Networks

Broker network topologies are kept congruent with the underlying physical network topology by periodically re-constructing global knowledge about the broker network, using a global state routing protocol [12]. Toward this end, each broker maintains its knowledge about the current broker network topology in a topology table T . Periodically, each broker receives its neighboring broker's T , updates its own T , and then propagates found topology updates to its neighbors. Each broker keeps track of the other brokers in its vicinity by querying the routing

table maintained by the wireless network routing protocol used in each broker machine, or via a neighbor discovery protocol like the Expanding Ring Search described in [13].

The broker topology update protocol can be summarized as follows.

Step 1: Broker Neighbor Discovery. Each broker periodically updates its neighboring brokers using the Expanding Ring Search. If a neighbor broker moves too far away from a broker, then the original overlay link between the broker and that neighbor is removed from the broker network. If a broker moves into the vicinity of another broker, a new broker link between them is created.

Step 2: Broker Topology Propagation. Once a broker completes updating its topology table by neighbor discovery, the broker sends to its neighbors those items in its topology table that have changed since the prior propagation period. A sequence number is associated with each such update.

Step 3: Broker Topology Update. When a broker receives updated information from its neighbor, it compares the sequence number of the incoming message with its topology table's corresponding items, replaces old items with new ones, and marks the items changed if the incoming items have a higher sequence number.

Step 4: Broker Routing Table Rebuilding. A broker's topology table T changes either due to its own execution of the periodic neighbor update or due to the receipt of topology propagation from its neighbors. When such changes occur, the broker rebuilds its routing table by recalculating its shortest path to other brokers henceforth uses the new routing table for delivering events.

Figure 3 depicts an example. At the beginning, all four brokers (A, B, C and D) have the same view of the global broker network topology A—B, B—C and B—D. At some point, node 2 moves away from B and 1, and closer to C. As a result, two old wireless links 2—B and 2—1 are removed and one new wireless

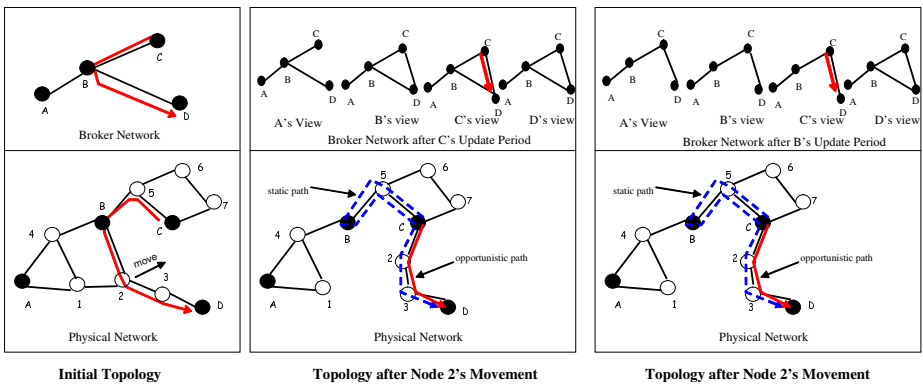


Fig. 3. An Example of Dynamic Broker Network Construction

link 2—C is created. Let's assume C is the first to start its update period. C adds D's as its new neighbor. Then C updates its topology table accordingly and propagates the change to its neighboring brokers B and D. After receiving updated information from C, each of B and D updates its topology table by adding the broker link C—D. B, C and D will rebuild their overlay routing tables based on the new broker network topology A—B, B—C, B—D and C—D. As a result of C's routing table update, opportunistic overlays deliver events from C to D using the wireless path C→2→3→D, compared with the static broker approach's C→5→B→5→C→2→3→D. Let's assume B is the next to start its update period. B removes D from its neighbor list and updates its broker topology knowledge accordingly. B then sends its changes since previous period to its neighbors A and C. Upon receiving updated information from B, both A and C update their topology knowledge by removing broker link B—D. A also adds the broker link C—D to its topology table. At this time, each of A, B and C has the latest broker topology knowledge. D's topology knowledge is outdated, and D's topology table will be updated either through C's propagation or via D's own running of the neighbor updates protocol, whichever comes first.

Our current protocol assumes reliable network communication channels. We leave it to future work to deal with issues like network partition, temporary broker disconnection, reconnection, and related reliability issues.

2.4 Dynamic Home Broker Change

End-to-end latency depends (1) on the network distance between a producer's home broker and a consumer's home broker, and also (2) on the distance between producers/consumers and their home brokers. By dynamically constructing a broker network, we aim to optimize the former. By dynamically changing home brokers, we improve the latter. Toward these ends, opportunistic overlays act as described next.

When a client subscribes to a broker network for the first time, it must connect to some home broker that receives events (via the broker network) on behalf of the client and delivers received events to the client via some wireless network link. Intuitively, we should choose the nearest broker as the client's home broker, thereby optimizing the delay between the client and its home broker. In ad-hoc mobile environments, therefore, home brokers must be chosen repeatedly, whenever nodes substantially change their locations. The procedure used by opportunistic overlays may be described as follows. Each client periodically (or in response to changes indicated by the underlying physical network protocol [14]) executes a protocol that searches for the broker nearest its current location. If the nearest broker is not its home broker, it notifies the current home broker of its discovery. Upon receiving this news from its client, the home broker selects new home broker based on average path length between producers and the client and the cpu load of candidate brokers. If the home broker must be changed, the modulator relocation protocol is performed.

An interesting aspect of our approach is overload control, which is important because end-to-end event delay from a producer to a consumer depends not only

on the length of the network path, but also on event processing times at brokers. Processing times are determined by how fast modulators can be executed on home brokers which in turn depends on the home brokers' loads and capabilities. In mobile ad-hoc networks, with clients changing locations, broker loads are subject to substantial runtime variation. One reason would be the sudden arrival of large numbers of local users, exemplified by many mobile units converging at a meeting. Another reason is the use of complex modulators by 'thin' clients, such as modulators that implement the flexible data transcoding required by such clients [7]. In fact, the processing time of a modulator on moderately to highly loaded brokers can exceed network delays by an order of magnitude.

The protocol followed to change home brokers can be summarized as follows.

Step 1: Nearest Broker Search. Each client searches the broker nearest to its location, periodically, using the same algorithm in broker neighbor update as described in Section 2.3. When a client finds the nearest broker that is not its current home broker, it shares with its current home broker the newly discovered broker along with its distance to that broker.

Step 2: Home Broker Selection. A client's home broker is selected from its current home broker and the newly discovered broker based on their distances to the client and their CPU load.

Step 3: Modulator Relocation. The relocation protocol relocates a client's modulator from its current home broker (source broker) to a new broker (destination broker), asks all producers' home brokers to compute paths to the new home broker, and switches event delivery from the old to the new paths. The relocation protocol guarantees event order, prevent event duplication or event loss, and ensure consistent event state. For applications that do not need such strict semantics, lighter weight protocols are used to loosen the requirements

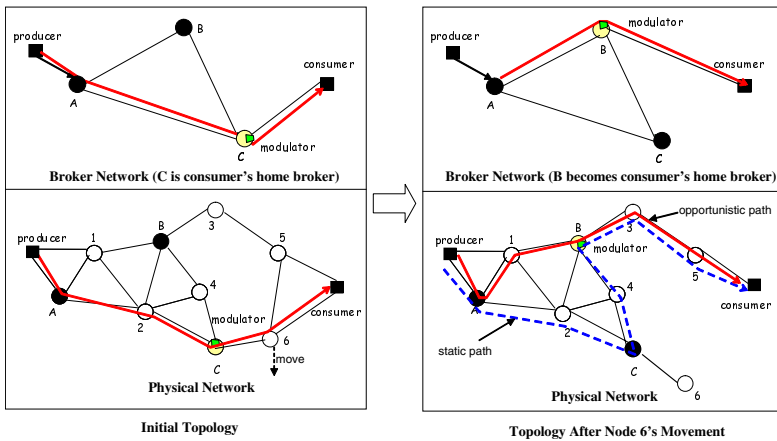


Fig. 4. An Example of Dynamic Home Broker Change

of event state consistency or modulator consistency or both. Additional detail about the relocation protocols appears in [15].

An example of dynamic home broker change is shown in Figure 4. A consumer originally receives events from broker C via the path $\text{producer} \rightarrow A \rightarrow C \rightarrow \text{consumer}$ corresponding to the physical network path $\text{producer} \rightarrow A \rightarrow 2 \rightarrow C \rightarrow 6 \rightarrow \text{consumer}$. Broker C is the consumer's home broker. The consumer's modulator is also placed on C and executes there. At some point in time, with node 6 moving away from the consumer and node 5, during the consumer's nearest broker discovery period, it discovers that it is closer to B than C. When the consumer detects this, it sends a home broker change request along with B's information to C. C will choose B as the consumer's new home broker, since $A \rightarrow B \rightarrow \text{consumer}$ is shorter than $A \rightarrow C \rightarrow \text{consumer}$ under the new network topology. C then performs the modulator relocation protocol. After the change, the event delivery path from producer to consumer becomes $\text{producer} \rightarrow A \rightarrow 1 \rightarrow B \rightarrow 3 \rightarrow 5 \rightarrow \text{consumer}$. In contrast, without changing home brokers, the old overlay path $\text{producer} \rightarrow A \rightarrow C \rightarrow \text{consumer}$ corresponds to the physical network path $\text{producer} \rightarrow A \rightarrow 2 \rightarrow C \rightarrow 4 \rightarrow B \rightarrow 3 \rightarrow 5 \rightarrow \text{consumer}$, which is 2 hops longer than the path used by opportunistic overlays.

3 Software Architecture and Selected Implementation Detail

3.1 Overview of JECho

Opportunistic overlays are realized with the JECho distributed event system [7]. JECho implements a publish/subscribe communication paradigm, providing services to distributed, concurrently executing components via event channels. Using JECho's *modulators*, individual event consumers can dynamically tailor event flows to their own needs, thereby adapting to runtime changes in component behaviors and needs and/or changes in platform resources. Modulators are implemented as Java objects, executed in a source's or broker's address space on behalf of clients.

3.2 Overview of Software Architecture

Opportunistic overlays are implemented as depicted in Figure 5. The architecture makes it easy to implement alternate adaptation methods, and the architecture itself is easy to reconfigure and extend.

The basic component layer provides the lower level functionalities of resource monitoring and broker information management necessary for implementing different adaptations. Event-driven adaptations are implemented by defining a set of actions to react to specific events received from basic components. By using services provided by basic components, adaptation code focuses on high level protocol only without needing to handle lower level details. The interaction between basic component layer and adaptation layer uses a set of consistent program interfaces and system events. As a result, it is easy for different brokers to define different adaptations based on their capabilities and requirements. In




Fig. 5. Opportunistic Overlays Software Architecture

fact, the implementation of an adaptation protocol within the current system is straightforward, as exemplified by the home broker change adaptation that has less than 50 lines Java code. It is also easy to reconfigure and extend the system with new adaptations, such as those needed to handle physical network partition. For future work, we are considering adding a policy layer that permits users to define high-level policies concerning the adaptations being carried out [15].

3.3 Basic Component Layer

The basic component layer layer is composed of a resource monitor, broker manager, and client manager. Components in this layer provide the core functionality implementing the adaptation protocols described in Section 2.2. Each component in this layer defines a set of program interfaces for other components to access its services (i.e., a set of ‘get’ and ‘set’ functions). Each component notifies other components and high level protocols by sending events containing the relevant information to an internal ‘system’ channel. Other components receive this information by registering their interests about certain events.

Resource Monitor. The Resource Monitor collects, aggregates, processes, and delivers data about local resource availability and about its communication costs to other brokers. Local resource information includes CPU load, memory availability, and modulator execution time.

Broker Manager. The Broker Manager maintains four tables, which are the Broker Neighbor Table (BNT), the Broker Information Table (BIT), the Broker Topology Table (BTT), and the Broker Routing Table (BRT). The set of program interfaces provided by the Broker Manager to higher level protocols include functions for accessing and changing broker-related information, and operations that propagate its broker topology to neighboring brokers. The Broker Manager

is also responsible for sending notifications to higher level components when it receives them.

Client Manager. The Client Manager maintains information about each client for which the broker is currently acting as home broker, including its name, IP address, physical location, as well as related path information (e.g., current routing path and communication overhead of the path). This information is stored in the Client Information Table (CIT).

3.4 Modulator Relocation Layer

Layered above the basic component layers are three modulator relocation operations: horizontal relocation, upstream relocation, and downstream relocation. Relocation operations perform the task of relocating a client's modulator from current broker to another broker, and of changing event delivery paths accordingly. Relocation operations are the basic functionality needed to support home broker changes and dynamic load balancing.

3.5 Adaptation Protocol Layer

The adaptation protocol layer implements a variety of protocols, including 'dynamic broker network construction' and 'dynamic home broker change'. Each such protocol is implemented with a Java object called an *adaptor*. An adaptor can register with the system event channel by specifying its interests in certain events delivered by the Resource Monitor, Broker Manager, and Client Manager. For the broker topology adaptor, interesting events are a time event and a broker propagation event. The interesting event for the home broker adaptor is a broker discovery message received from a client. The code in the adaptor is implemented in the event handler method "process()", which is invoked whenever an interesting event is received. This code implements changes, such as reconfiguring the broker network, changing the home broker, or rebuilding a routing table. Using adaptors and the services provided by basic components, system developers can create potentially complex adaptation policies. Our prototype implementation has three adaptors: a broker topology construction adaptor, a home broker change adaptor and a broker load balancing adaptor. Each adaptor performs the task for which it is named.

3.6 Client Component

The final element of opportunistic overlays are client-resident components that interact with the broker overlay. These include a Resource Monitor and Client Manager.

4 Performance Evaluation

4.1 Simulation

Simulation techniques are used to evaluate the opportunistic overlay approach under various wireless network configurations. In all experiments reported in

this section, the network consists of 100 mobile nodes that randomly roam in a 1000 x 1000 meter square. The random waypoint mobility model [16] is used with a pause time of 10 seconds. The radio transmission range of each node is 250 meters. Each simulation spans 600 seconds of simulated time.

Currently, our simulation study is limited to high level overlay routing, for which link layer details and physical layer characteristics are not modeled explicitly. Since aspects like control overheads or link contention are not taken into account, a high routing packet load, for instance, does not interfere as much with data transmissions as it might in reality. Also, there are no transmission errors and delays associated with overlay routing packets (i.e., broker network topology propagation and broker neighbor discovery). Future work should address these limitations by constructing a more comprehensive simulator with a MAC layer model (e.g., IEEE 802.11 MAC). The point of this section’s simulation results is to compare the relative performance of the opportunistic approach to overlay routing vs. static approaches.

Performance of the Broker Network. The first set of experiments evaluate the average lengths of network paths across broker overlays with vs. without opportunistic overlay protocols. Measurements consider only brokers, not clients. 40 nodes among 100 nodes are randomly chosen as brokers. We vary the experimental configurations with different broker topology update intervals. The path length from broker B1 to broker B2 is the corresponding physical network distance of the shortest broker path between B1 and B2, which is computed based on B1’s local knowledge of current broker topology. The average path length from each broker to all other brokers is computed and averaged over all brokers. In order to establish a basis for comparison, we also measure the average length of the physical network path between each pair of brokers. In addition, the overheads of broker topology update with different update intervals are evaluated.

Timeline of Path Length. Figure 6 shows how path length changes in a simulation with an update interval of 50 seconds and mobility speeds between 1

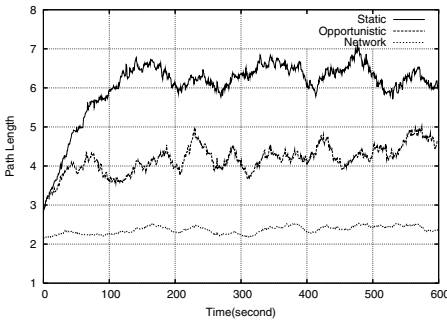


Fig. 6. Timeline of Path Length among Brokers




Fig. 7. Average Path Length among Brokers versus Update Period

m/s and 20m/s. As shown in the figure, the opportunistic approach can deliver events more efficiently than the static approach at almost all time points. At the beginning, both the static and the opportunistic approaches have similar path lengths, since the initial broker network matches physical network topology. As time passes, the path lengths of the static approach increase rapidly because the initial broker topology cannot reflect changes in physical network topology caused by node mobility. Compared with the static approach's 6.06 hops and the network's 2.36 hops, the opportunistic approach has an average path length of only 4.17 hops.

Average Path Length versus Broker Update Period. Figure 7 shows the average path length versus the update interval, the latter varying from 10 to 100 seconds. As expected, with increased update intervals, path length increases since larger update intervals imply slower reactions to changes in physical network topology. However, as shown in the figure, the change in path length is not rapid with the increase of update periods, e.g. 4.43 hops with an interval of 100 seconds versus 3.32 hops with one of 10 seconds. Even with a relatively low update period of 100 seconds, the opportunistic approach still outperforms the static approach significantly, 4.43 hops vs. 6.06 hops.

Update Overhead versus Broker Update Period. The overheads of broker network updates are shown in Figure 8. Overhead is computed as the average bandwidth requirement of each broker for propagating its broker topology knowledge to its neighbors, the argument being that network resources tend to be scarce in pervasive systems. As shown in the figure, a total of 2.4 Kbps bandwidth is used with an update period of 50 seconds. This constitutes moderate bandwidth usage in modern network infrastructures.

Performance of Event Delivery between Mobile Clients. Most relevant to our work, of course, is the end-to-end performance experienced by end users, i.e., clients. In the following experiments, we randomly choose 20 brokers, 10 event producers, and 60 event consumers from 100 mobiles nodes. In this set of




Fig. 8. Update Overhead versus Update Period




Fig. 9. Timeline of Average Path Length between Producers and Consumers

experiments, the broker network update interval is fixed at 20 seconds, and we vary mobile clients' home broker discovery periods as well as mobility speed. Four different approaches are evaluated in terms of the resulting average path lengths between each pair of producer and receiver: (1) the static approach changes neither the broker network topology nor the home broker of mobile clients; (2) the static-opportunistic approach change mobile client's home broker only; (3) the opportunistic-opportunistic changes both the broker network topology and mobile client's home broker; and (4) the best approach keeps updating the broker network whenever the physical network changes and calculates shortest broker paths based on up-to-date physical network topology data. Although the best approach is not practical, we have included it to establish a basis for comparison.

Timeline of Path Length. Figure 9 shows the performance results of a simulation with a home broker discovery period of 10 seconds. We can see that the opportunistic-opportunistic approach performs best among the three realistic approaches, shortening the delivery paths up to 5 hops compared with the static approach. Even the static-opportunistic approach can improve event delivery significantly compared with the static approach.

Average Path Length versus Nearest Broker Discovery Period. Studies assessing path lengths versus home broker discovery periods are reported in Figure 10. With increased closest broker discovery periods, path lengths increase slightly. As discussed in Section 2, the mobile client can find the nearest broker by querying its routing table or by using an expanding ring protocol. Either way, the costs are small compared with the overheads of broker update operations. More frequent closest broker discovery results in more frequent home broker changes, hence more frequent modulator relocations. Additional simulation results not reported here due to space limitations demonstrate the fact that modulator relocation costs are much smaller than the overheads of broker topology change (see [15] for more detail).




Fig. 10. Average Path Length between Producers and Consumers versus Nearest Broker Discovery Period




Fig. 11. Average Path Length between Producers and Consumers versus Mobility

Average Path Lengths versus Mobility. Average path lengths at different mobility speeds is shown in Figure 11. The figure shows that more frequent nearest broker discovery achieves shorter event delivery paths. In particular, when the nodes move at a very fast speed, increasing the discovery frequency can improve the performance significantly.

Simulation Conclusions. Multiple insights result from the simulation experiments described in this section. First, client or broker mobility in ad-hoc networked systems demand dynamic changes to the overlay networks used for event propagation. Without runtime overlay adjustments, event delivery paths and therefore, average event delays increase substantially and may not remain viable for realistic systems and applications. Interestingly, even relatively ‘slow’ overlay adjustments performed by the opportunistic overlay approach attain much improved results compared to static solutions (e.g., with 100 second update intervals, the opportunistic approach results in an average of 4.43 hops for packets vs. the static approach’s 6.06 hops). Second, end-to-end delays in packet delivery to clients are improved further when overlay adjustments are complemented with changes in the assignment of home brokers to clients. Since it is cheaper to change home brokers than to reconfigure broker overlays, the former changes can be (and should be) more frequent than the latter. Finally, change frequencies are strongly correlated with mobility speeds.

4.2 System Emulation

In order to evaluate the effects of load balancing, we have conducted a set of experiments on an ad-hoc wireless network emulator. The mobility emulator runs on a Linux cluster of 20 nodes with MobiEmu [17] running on each node. The cluster network is a gigabit Ethernet switch. MobiEmu is a software platform for testing and analyzing ad-hoc network protocols and applications. With control software running on each node, MobiEmu mimics dynamic connectivity among nodes by dynamically installing or removing packet filters for specific MAC addresses. Since we focus on load balancing in this set of experiments, we use the ‘best-case’ ad-hoc routing provided by MobiEmu software. These protocols always deliver packets via shortest network paths. More detail about this system appears in [17].

The emulated mobile network consists of 20 mobile nodes, of which 5 nodes are event brokers, 5 are event producers and 15 are event consumers (5 brokers reside at event receiver nodes). In our experiments, mobile nodes move in a space of 750m x 500m and use random waypoint mobility with a pause time of 30 seconds and speeds between 1m/s and 20m/s. Due to the relatively small network size, the broker update period is set to 5 seconds and the nearest broker discovery period to 1 second. We vary the average load of the broker network and measure the maximum load during execution.

Results appear in Figure 12. As shown in the figure, when broker load is relatively light (i.e., less than 30%), there are no overloaded brokers and both approaches behave the same, where the nearest broker to a mobile client is always chosen as the client’s home broker. With increased system load, without




Fig. 12. Maximum Broker Load versus Average System Load




Fig. 13. Path Length versus System Load

overload control, some brokers become overloaded. The load balancing algorithm ameliorates this problem, because a client’s home broker will not be moved to an overloaded node, even if that node is closer than the old one. The positive outcomes of load management reported in these measurements are moderate, of course, since in random waypoint mobility, nodes move independently. Load balancing is more important and will have more significant effects when nodes move in groups, as exemplified by conference participants moving from one presentation venue to another, for instance.

Figure 13 depicts path length versus system load. When system load is less than 70%, the opportunistic approach always chooses the nearest broker as home broker, and the delivery path can be more than 1 hop shorter than in static approach. With increased system load, load balancing selects broker with lighter loads on relatively longer paths, resulting in increased path lengths. However, the opportunistic approach continues to outperform the static approach even when system load reaches 100%.

The simple conclusion from these measurements is that realistic implementations of dynamic overlay networks for MANET should not ignore broker loads, especially when overlays perform meaningful application-level processing actions.

4.3 Testbed

The last set of experiments demonstrate the practical utility of our approach, by running a sample application on an actual wireless testbed. The testbed consists of 3 laptops A, B and C. Wireless connectivity is provided by Orinoco 802.11b cards. These cards are set to ad-hoc mode on channel 8. No WEP encryption is used. All three laptops use the UoB JAODV version 0.2, an AODV implementation in Java [11]. In order to simulate network connectivity changes, we dynamically set filters at the MAC layer. In a network consisting of 3 nodes A,B and C, there are four possible network topologies without network partitions: A-B-C, A-B-C-A , A-C-B and B-A-C.

The experiments being performed run a flood watch application on the testbed, comparing the event delivery latency of the opportunistic with the static




Fig. 14. Experiment Configuration




Fig. 15. Average Latency (100% data)




Fig. 16. Average Latency (10% data)

approach using the 4 different network topologies. The application consists of two programs and works as follows. PreSend reads precipitation data from a file and places normalized precipitation data on event channel. A client program subscribes to the channel and provides a modulator that calculates water depth from precipitation data and terrain topology data using a runoff model. The client is typically interested in flood information in some specific area, which can be defined by a two-dimensional bounding box. Flood data that is outside the bounding box will be filtered (i.e., removed) before data is sent to the channel. The output of full-size flood data is a double array of 100 x 100.

In our experiment, PreSend is running on laptop A and the client program runs on laptop C. Broker programs run on both A and B. The physical network topology changes every 200 time units. Experiment configuration is shown in Figure 14. The initial physical network topology is A-B-C. B is C’s home broker. Both approaches use the same event delivery path A→B→C. After 200 time units, the network topology changes to A-B-C-A, so that the opportunistic approach chooses A as C’s home broker and relocates its modulator from B

to A, hence resulting the shorter delivery path $A \rightarrow C$. During the time interval of 400 to 600, the network topology is A-C-B, where the opportunistic overlay approach still uses $A \rightarrow C$ as delivery path corresponding to the same physical network path $A \rightarrow C$. The corresponding physical network path of $A \rightarrow B \rightarrow C$ used by static approach becomes $A \rightarrow C \rightarrow B \rightarrow C$. In the final interval, the network topology changes to B-A-C, where the opportunistic approach still uses the same path $A \rightarrow C$, and the static approach's path $A \rightarrow B \rightarrow C$ now corresponds to the physical network path $A \rightarrow B \rightarrow A \rightarrow C$, with overlap at A.

Figure 15 shows the latency comparisons with four topologies when the client is interested in all data. As shown in the figure, the opportunistic approach can deliver data up to 6 times faster than the static approach. The static approach has its worst performance in the configuration of B-A-C (6000ms compared with to the opportunistic approach's 1000ms), where data delivery following the path $A \rightarrow B \rightarrow A \rightarrow C$ results not only in a longer path but also in additional network bandwidth usage at A, since at the physical layer, every event is actually delivered twice at A.

When a client is interested in only 10% of the area data, the opportunistic approach coupled with its use of modulators shows additional improvements. Results depicted in Figure 16 show that by relocating C's modulator from B to A, the opportunistic approach not only delivers data following a shorter path but also delivers less data, hence improving the application's performance significantly: the average latency is less than 200ms under all network configurations.

The key result of the testbed experiments presented here is that it is important to dynamically adjust the middleware overlays used in pervasive systems. The opportunistic overlay approach described and evaluated in our research is one method for runtime overlay management and by using it, significant performance improvements can be attained compared to non-adaptive approaches.

5 Related Work

5.1 Content-Based Event Systems

Publish/subscribe systems [1, 18, 4–7] have been investigated for many years, but most implementations have focused on systems where nodes don't move, broker networks remain fixed, and broker network topologies are defined at deployment time. As a result, their fixed event dissemination structures make them unsuitable for applications in mobile environments where physical network topology and node locations change continuously. In addition, most publish/subscribe systems perform event filtering with predicate-based subscriptions; they do not support the general event processing needed for the complex data conversions occurring in multimedia, business, or scientific applications. Opportunistic overlays are realized with the JECho pub/sub infrastructure [7]. JECho generalizes the capabilities of other event systems, by using consumer-provided functions, termed event modulators [7]. The intent is to address the severe resource limitations existing in many mobile and embedded systems, by permitting event

consumers to deploy application-specific functions that manipulate event content into event sources and/or brokers, so as to precisely meet their current needs, and to avoid needless data transfers. Generic function-based subscription makes the opportunistic overlay system more feasible for developing applications in pervasive systems and mobile environments.

Previous research on event-based middleware for wireless networks has addressed applications in which mobile nodes make use of the wireless network to connect to a fixed network infrastructure [19, 20, 10, 21, 14, 10, 8, 14]. The opportunistic overlays presented in this paper differ from these systems in that they are designed for mobile ad hoc networks, support dynamic reconfiguration of event event dissemination structure and offer behaviors transparent to applications.

[22] presents an algorithm for topological reconfiguration in content-based publish/subscribe due to changes in underlying connectivity. Compared with opportunistic overlays, reconfiguration in [22] involves only link removal or insertion, and no details are given on how to apply the proposed approach to handle changes in mobile environments. In addition, the approach assumes a tree-based topology between dispatchers, which makes it hard to achieve robustness, since a single link failure partitions the tree. Another approach to dynamic broker network configuration is described in [23]. The idea is to place ‘close’ to each other brokers that manage similar subscriptions. This is complementary to our work, which focuses on resource awareness, where reconfiguration in the opportunistic overlays is based on nodes’ physical locations and the underlying physical network topology. Similar to the approach in [22], the topology used in [23] must remain acyclic, whereas the opportunistic overlay approach supports general broker overlay topologies.

5.2 Content-Based Event Systems in Mobile Ad-Hoc Networks

Steam [24] is an event-based middleware service designed for ad-hoc wireless networks. It targets application scenarios where nodes are more likely to interact when they are in close proximity to each other. We consider more general uses of publish/subscribe in ad-hoc networks. Further, Steam uses an implicit event model without intermediate broker nodes.

[25] presents a distributed protocol to construct optimized publish/subscribe trees in ad-hoc wireless networks. The protocol builds multicast trees directly on top of lower level radio broadcast primitives. Our work relies on the underlying network infrastructure’s ability to provide basic network connectivity. Another difference is that their approach assumes a relatively stable environment with occasional reconfigurations followed by periods of stability. Opportunistic overlays do not make that assumption, and they can actually handle high levels of mobility as shown by our experimental results.

[26] proposes a publish/subscribe system for MANET that integrates an extended ODMRP (On-Demand Multicast Routing Protocol [27]) with content-based subscriptions. Similar to [25], ODMRP-PUB/SUB delivers events by creating multicast groups. The difference is that ODMRP-PUB/SUB uses a mesh-based approach instead of the tree-based one used in [25]. Since a consumer’s

subscription is a general function applied to events in our system, the approach of combining a multicast protocol and subscription aggregation/match is not readily applicable in our case. Further, ODMRP-PUB/SUB focuses on the routing between brokers and does not address the issue of delivery from brokers to producers/consumers. The purpose of ODMRP- PUB/SUB is to optimize network throughput, while our opportunistic method focuses on providing timely event delivery.

5.3 Overlay Multicast Protocols in Mobile Ad-Hoc Networks

AMRoute [28] and PAST-DM [13] are two ad-hoc multicast protocols that use the overlay approach. AMRoute uses a static virtual mesh and has low efficiency due to the increasing mismatch between virtual topology and physical network topology, as shown in [13]. PAST-DM addresses the efficiency problem by dynamically adapting the virtual topology to changes in the physical network. That brokers need to process events distinguishes our system from multicast systems where nodes perform data routing and participate as relays. Although opportunistic overlay approach uses a similar dynamic virtual overlay construction technique as PAST-DM [13], the dynamic routing path in opportunistic approach involves not only path changes in event routing, but also subscription code relocation, which makes the existing dynamic delivery technique in PAST-DM is not readily applicable to our system. In addition, the processing of events will consume a broker's computational resources, which implies that brokers' computational capabilities need to be taken into account.

Finally, in PAST-DM, all member nodes are considered to be equivalent peers and participate in overlay routing. In contrast, opportunistic overlays conceptually divide nodes into brokers which are organized into an overlay broker network, and clients(producers/consumers) which send/receive events via the broker network. This model is more suitable for content-based routing since overlay routing through 'thin' nodes with very limited resources will present burden on such nodes and may result in inefficiency of content delivery in mobile systems. Dynamic reconfiguration using by opportunistic overlays adapts both the overlay broker network and the connections between brokers and clients.

6 Conclusions and Future Work

This paper presents an approach to optimizing content-based event delivery in mobile ad-hoc networks. In response to changes in physical network topology and to node mobility, the opportunistic overlay approach dynamically changes broker network topology, clients' assignments to brokers, and event delivery paths, with the goal of optimizing end-to-end delays in event delivery. Opportunistic overlays are prototyped with the JECho pub/sub system [7]. Comprehensive performance evaluations are performed via simulation, emulation, and with representative applications on a physical testbed. Experimental results for mobile ad hoc networks demonstrate that the opportunistic overlay approach can significantly improve event delivery delays compared to static approaches, even with

high levels of mobility. Results also show that the overheads of dynamic adaptation are moderate. Using a flood watch application and a wireless testbed, the opportunistic overlay approach is practically applicable in an actual ad hoc wireless network.

Future work should address some deficiencies of our current implementation, as well as generalize upon the basic concept of opportunistic overlays. First, our current implementation assumes a reliable network environment and therefore does not consider dynamic disconnection, reconnection, and network partition. Future work will add application-specific failure recovery to broker overlays. Second, we will extend the opportunistic approach to optimize performance metrics other than end-to-end latency, including network bandwidth and power usage. We may also explore optimizing multi-dimensional performance metrics. A final topic of interest is a performance study that uses a more comprehensive simulator with a MAC layer model (e.g., IEEE 802.11 MAC).

References

1. Strom, R., Banavar, G., et.al: Gryphon: An information flow based approach to message brokering. Technical report, IBM TJ Watson Research Center (1998)
2. Carzaniga, A., Rosenblum, D.S., Wolf, A.L.: Achieving scalability and expressiveness in an internet-scale event notification service. In: Proceedings of the Nineteenth Annual ACM Symposium on Principles of Distributed Computing(PODC 2000), Portland, Oregon (2000) 219–227
3. Cugola, G., Nitto, E.D., Fuggetta, A.: The "jedi" event-based infrastructure and its application to the development of the opss wfms. In: IEEE Transactions on Software Engineering in 2001. (2001)
4. Segall, B., Arnold, D.: Elvin has left the building: A publish/subscribe notification service with quenching. In: Proceedings of A UUG97. (1997)
5. Fiege, L., Mühl, G., Gärtner, F.C.: A modular approach to build structured event-based systems. In: Proceedings of the 2002 ACM Symposium on Applied Computing (SAC'02). (2002) 385–392
6. Eisenhauer, G., Bustamante, F.E., Schwan, K.: Event services in high performance systems. *Cluster Computing* 4 (2001) 243–252
7. Zhou, D., Schwan, K., Eisenhauer, G., Chen, Y.: Supporting distributed high performance application with java event channels. In: Proceedings of the 2001 International Parallel and Distributed Processing Symposium (IPDPS 2001). (2001)
8. Huang, Y., Garcia-Molina, H.: Publish/subscribe in a mobile environment. In: Proceedings of the 2nd ACM International Workshop on Data Engineering for Wireless and Mobile Access (MobiDE'01). (2001) 27–34
9. Cugola, G., Jacobsen, H.A.: Using Publish/Subscribe Middleware for Mobile Systems. *ACM SIGMOBILE Mobile Computing and Communications Review* 6 (2002) 25–33
10. Fiege, L., Gärtner, F.C., Kasten, O., Zeidler, A.: Supporting mobility in content-based publish/subscribe middleware. In: ACM/IFIP/USENIX International Middleware Conference (Middleware 2003). (2003) 103–122
11. : Uob-jadhoc aodv implementation, rfc 3561. <http://www.aodv.org/> (2004)
12. Chen, T.W., Gerla, M.: Global state routing: A new routing scheme for ad-hoc wireless networks. In: Proceedings of IEEE ICC'98. (1998)

13. Gui, C., Mohapatra, P.: Efficient overlay multicast for mobile ad hoc networks. In: Proceedings of IEEE Wireless Communications and Networking Conference. (2003)
14. Chen, Y., Schwan, K., Zhou, D.: Opportunistic channels: Mobility-aware event delivery. In: ACM/IFIP/USENIX International Middleware Conference (Middleware 2003). (2003) 182–201
15. Chen, Y.: Opportunistic Overlays: Efficient Content Delivery in Mobile Environments. PhD thesis, Georgia Institute of Technology (2005)
16. Johnson, D.B., Maltz, D.A.: Dynamic source routing in ad hoc wireless networks. *Mobile Computing* **353** (1996)
17. Zhang, Y., Li, W.: An integrated environment for testing mobile ad-hoc networks. In: Proceedings of the Third ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'02). (2002)
18. Carzaniga, A., Rosenblum, D.S., Wolf, A.L.: Design and evaluation of a wide-area event notification service. *ACM Transactions on Computer Systems* **19** (2001) 332–383
19. Cugola, G., Nitto, E.D., Picco, G.P.: Content-based dispatching in a mobile environment. In: Proceedings of WSDAAL 2000. (2000)
20. Sutton, P., Arkins, R., Segall, B.: Supporting disconnectedness - transparent information delivery for mobile and invisible computing. In: CCGrid 2001 IEEE International Symposium on Cluster Computing and the Grid. (2001)
21. Caporuscio, M., Inverardi, P., Pelliccione, P.: Formal analysis of clients mobility in the siena publish/subscribe middleware. Technical report, Department of Computer Science, University of L'Aquila (2002)
22. Picco, G.P., Cugola, G., Murphy, A.L.: Efficient content-based event dispatching in the presence of topological reconfiguration. In: Proceedings of the 23rd International Conference on Distributed Computing Systems (ICDCS 03). (2003) 234–243
23. Virgillito, A., Beraldi, R., Baldoni, R.: On event routing in content-based publish/subscribe through dynamic networks. In: Proceedings of the Ninth IEEE Workshop on Future Trends of Distributed Computing Systems (FTDCS 2003), IEEE (2003) 322–328
24. Meier, R., Cahill, V.: Steam: Event-based middleware for wireless ad hoc networks. In: In Proceedings of the 1st International Workshop on Distributed Event-Based Systems (DEBS'02). (2002)
25. Huang, Y., Garcia-Molina, H.: Publish/subscribe tree construction in wireless ad-hoc networks. In: Proceedings of the 4th International Conference on Mobile Data Management(MDM 2003). (2003) 122–140
26. Yoneki, E., Bacon, J.: An adaptive approach to content-based subscription in mobile ad hoc networks. In: Proceedings of The First International Workshop on Mobile Peer-to-Peer Computing (MP2P'04). (2004) 92–97
27. Lee, S.J., Su, W., Gerla, M.: On-demand multicast routing protocol in multihop wireless mobile networks. *MONET* **7** (2002) 441–453
28. Liu, M., Talpade, R.R., McAuley, A.: AMRoute: Adhoc Multicast Routing Protocol. Technical Report 99, The Institute for Systems Research, University of Maryland (1999)