SWARM: Self-organization of Community Wireless Mesh Networks

Saumitra Das[1], Konstantina Papagiannaki[2], Suman Banerjee[3], Y. C. Tay[4]
[1]*Purdue University,* [2]*Intel Research Pittsburgh,* [3]*Univ. of Wisconsin Madison,* [4]*National University of Singapore*

1. PROBLEM AND MOTIVATION

Community wireless networks have been proposed as a powerful technique to spread broadband network access to underprivileged, under-provisioned and remote areas. These networks consist of a few Internet gateways which are reached by homes using multi-hop wireless links between wireless routers. The benefits of such networks include low costs for deployment due to reduced wiring needs, low maintenance and increased flexibility. Current practice in routing protocols for such networks (e.g. LQSR, OLSR and SrcRR) is for routing protocols to obtain information about the link quality (via some metric such as ETT, ETX) and select a gateway to whom a route minimizes the cost of the metric. All nodes operate on the same known frequency to maintain connectivity.

However, there are several inefficiencies with currently deployed solutions. They are only a first step in connecting a mesh network and do not make optimal use of the available network resources. Due to this, mesh deployments currently suffer from bad performance due to excessive interference. First, the entire network uses a single arbitrarily chosen channel so that the routing protocols can operate and the network remains connected. This limits the significant interference reduction possible through utilizing frequency diversity by configuring the gateways with different channels and allowing nodes to select one of the gateways/frequencies. Several deployed and operational single radio mesh networks that we have measured all run the entire network on a single frequency. Fundamentally, if multiple frequencies are involved the network operation needs to be significantly more intelligent and deal with not just routing, but also how best to partition nodes among gateways (frequencies) and what data delivery structure to use. Second, currently each node selects routes using graph algorithms from wired networks which do not consider the effect of interference. Due to this, a node does not consider the impact of routes that have been chosen by other nodes and the consequent interference. For example, node 17 (in Figure 1) might choose GW3 over GW1 due to its ETT being slightly lower to GW3. However this ignores the fact that GW3 may have a larger distribution tree (set of nodes) using it and GW 1 may be a better choice.

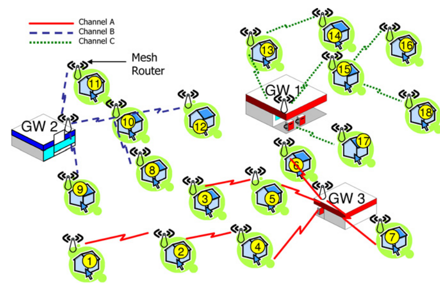

Figure 1: **A self-organized mesh network.**

All these inefficiencies suggest that a new more structured approach is required which also needs to be automated and self-organizing to remove adoption barriers. We propose to better utilize resources in mesh networks through a self-organization protocol, SWARM, the effectively utilizes network resources while remaining practical and deployable. To accomplish this, SWARM requires the following key synergistic components: (1) A means for gateways to utilize frequency diversity to reduce interference. (2) A gateway selection and tree construction protocol that associates each node to a gateway by considering the impact of existing nodes joined to a gateway and their link and interference characteristics. In essence, SWARM aims to self-organize a community network such that some network-wide metric is improved given the nodes, environment and link characteristics. An example of a SWARM self-organized mesh network is shown in Figure 1. Here gateways have chosen orthogonal channels and nodes have been organized such that network resources are best utilized. Without SWARM, routing protocols would require operating all gateways on the same channel and each node would look only at its own path to a gateway without considering the effect of other nodes, thus leading to significantly higher interference.

The problem of allowing gateways to utilize frequency

*CoNEXT'*07, December 10-13, 2007, New York, NY, U.S.A.

diversity while allocated nodes in good distribution structures to each gateway is challenging. Consider that a brute force approach would need to enumerate all possible assignments of gateways to nodes (clustering) which is combinatorial. For each possible assignment it would need to come up with the best tree which involves enumeration of all trees that can be formed to each gateway. Finally it would need a method to figure out which clustering+associated trees constructed are the best and then find a corresponding channel assignment for each tree that maximizes throughput. Thus, an analytical performance model is required. Current approaches towards such performance modeling need ILPs which are prohibitive computationally. Thus, the main contribution of SWARM is a set of algorithms that address these requirements (clustering, tree construction, modeling and frequency diversity) while remaining practical and feasible to deploy. We have designed a clustering and pruning algorithm that systematically goes through all possible clusterings of nodes to gateways while pruning a large number of likely bad clusterings quickly. We have designed an integrated interference aware tree construction heuristic and a tractable and fast way to model the performance of such trees. The tree construction algorithm incrementally builds a tree greedily by adding one node at a time to the existing tree such that the node added maximizes the performance metric of concern (e.g. sum of inverse of throughputs achieved by each mesh routers). A scalable performance modeling technique based on estimating node duty cycles allows use to decide which node to add at each step. Finally, we have designed a practical way for such networks to utilize frequency diversity using channel hopping that decouples tree construction from frequency selection. Together these practical algorithms allow a mesh network to better utilize network resources.We argue that such centralized computation of good network topologies is an effective method for improving the performance of community wireless networks. We have documented gains from SWARM using both simulations as well as testbed experiments from a large wireless network.

2. THE BENEFITS OF GOOD NETWORK ORGANIZATION

We implemented SWARM and deployed it over a wireless testbed. We first measure the benefits of intelligent data delivery structures using SWARM versus the current practice in mature routing implementations in which each node simply chooses a path which optimizes some metric. The main goal of this experiment is to evaluate the benefits of the tree construction algorithm of SWARM, i.e. what benefits does a structured approach to organizing a set of nodes and a gateway provide over using state-of-the-art approaches. We deployed SWARM on 8 nodes of a wireless testbed[1] and configured one of the nodes as a gateway.

[1]http://engineering.purdue.edu/MESH

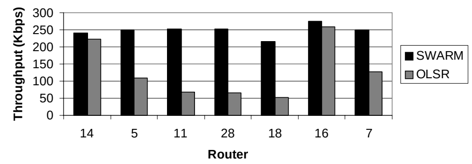

Figure 2: **SWARM performance in a single gateway scenario.**

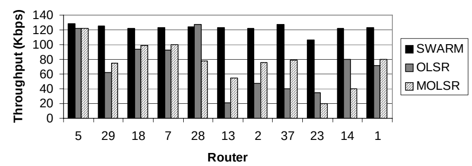

Figure 3: **SWARM performance in a multiple gateway scenario.**

Each client downloaded data using UDP at 256Kbps from the gateway node (some clients required multi-hopping). Figure 2 shows that SWARM allows this network to deliver a throughput of 256 Kbps to each mesh router unlike a state-of-the-art routing protocol (OLSR). Overall the sum of MR throughputs in SWARM is 1.92× of OLSR nd the sum of inverse of MR throughputs (Potential Delay) [2] in SWARM is 2.64× better than OLSR due to interference aware tree construction.

We now evaluate how SWARM improves performance when multiple gateways are present in the network. We compare the performance of the topology selected by SWARM with simply using OLSR on the same network and with MOLSR (a version of OLSR we designed that makes each node greedily select an advertising gateway via scanning advertisements on all channels. Each client downloaded data using UDP at 128Kbps from its chosen (best) gateway node (some clients required multi-hopping). Again the potential delay performance of SWARM is better than both OLSR and a multichannel version of OLSR by 2.2× and 2× respectively. Thus, the clustering and frequency separation and tree construction algorithms of SWARM can provide significant gains. Current algorithms do not effectively exploit frequency separation or interference relationships.

In conclusion, SWARM is a sophisticated system that attempts to optimally utilize network resource in complex wireless networks. Given a network graph and conflict graph, SWARM can simply tell us what is the best way to organize the nodes and operate the network in bounded time. Through measurements of deployed wireless networks we have found that changes in the network graph (due to node failures/additions or link quality changes) occur at time scales that make it feasible to re-execute SWARM periodically in a centralized manner. Finally, SWARM can be decoupled from the wireless networks and operate as a shared network-accessible computational service in a data-center. We are actively working towards provisioning such a service and opening it up to users.

[2]A better metric that captures unfairness properties.