

Provably Competitive Adaptive Routing

Baruch Awerbuch David Holmer Robert Kleinberg Herbert Rubens

Abstract—An ad hoc wireless network is an autonomous self-organizing system of mobile nodes connected by wireless links where nodes not in direct range communicate via intermediary nodes. Routing in ad hoc networks is a challenging problem as a result of highly dynamic topology as well as bandwidth and energy constraints. In addition, security is critical in these networks due to the accessibility of the shared wireless medium and the cooperative nature of ad hoc networks. However, none of the existing routing algorithms can withstand a dynamic proactive adversarial attack. The routing protocol presented in this work attempts to provide throughput competitive route selection against an *adaptive* adversary. A proof of the convergence time of our algorithm is presented as well as preliminary simulation results.

I. BACKGROUND

The basic service offered by every node in an ad hoc network is that of *routing packets* from their source to their ultimate destination. In general, routing protocols are susceptible to a wide variety of attacks. For example, a malicious node may perform a denial of service attack by selectively jamming some areas of the network.

A great deal of work has been done in terms of guaranteeing practical security considerations in existing network protocols (see description of existing work provided in Section II). In practice, adversarial attacks observed and documented in ad hoc networks might not be overly sophisticated. The ease of access to the medium has allowed extremely basic attacks to cause a great deal of damage. Consequently, such attacks can be thwarted by simple yet effective methods.

Existing work in the literature considered a number of strong adversary models. For example, [1] considers a random fault pattern; [2] deals with a static fault pattern and [3] deals with an *oblivious* (non-adaptive) pattern.

Our goal is to design routing protocols for networks that are provably tolerant of *arbitrary adaptive* DOS attacks. The adversary that we will consider

selectively attacks packets on a given node or link. This adversary benefits from knowledge of the traffic pattern (including packet contents); this includes all current traffic and all past traffic history.

As a result, the algorithms and analysis techniques used in the previous work will not apply. Existing methods that do not ignore sophisticated adaptive attacks either use brute force (flooding) or assume the existence of *some* trusted servers or routers (see section II). We do not wish to make such restricting assumptions. As a result, the task of designing a throughput competitive routing algorithm is much harder.

It may appear that our adversarial routing model may lead to impractical algorithms in benign (non-adversarial) settings. However, routing algorithms similar to the one studied here were developed and tested in real network environments by British Telecom and NTT for both wired and wireless networks with superior results [6], [20], [9], [7], [17], [18], [5], [10]. AntNet, a particular such algorithm, was tested in routing for data communication networks [6]. The algorithm performed better than OSPF, distributed Bellman-Ford with various dynamic metrics, and various modifications of shortest path with a dynamic cost metric [4], [17].

Our contribution: We propose a new algorithm for adaptively selecting routing paths in a network with dynamic adversarial edge failures, and we give a rigorous mathematical analysis of this algorithm, proving that its packet loss will match the minimal cumulative loss of any path, up to an additive error which is sublinear in the number of trials. The general framework we propose is appropriate for analyzing routing protocols for networks operating under the extremely strong adversarial model specified above. Such strong models have not been considered in the literature to the best of our knowledge. In fact, adaptive dynamic denial of service attacks are sufficient to break most existing algorithmic work. (In Section III, we briefly describe why DoS attacks are so devastating.)

What distinguishes the present work is our insistence on proving that under *completely arbitrary* adversarial behavior, with essentially no assumptions about the network, the packet loss of our protocol will essentially match the minimal cumulative loss (i.e., sum of losses of individual links) of any path. While it may seem counter-intuitive that such a goal can be achieved, the key is that although we assume an arbitrary *dynamic* adversary, we compare the algorithm against the best *static* path; if the adversary works hard to damage the algorithm's throughput, it must necessarily inflict a large cumulative loss on *every path* in the network.

At this point, one could debate whether such sophisticated adversaries are ever going to surface in reality, or whether they are only monsters in our imagination. Our counter-argument is that if, as we theorize, ubiquitous wireless networks become the underlying fabric that binds our society together, we cannot afford not to plan against an adversary with arbitrary powers.

The rest of this paper is organized as follows. In Sections II and III, we review some of the existing work in this area and survey some of the challenges which illustrate why our problem is not solved by simpler approaches. In Sections IV we outline the main ideas underlying our algorithm, which is specified precisely in Section V, analyzed mathematically in Section VI, and experimentally tested in Section VII. The algorithm contains some tunable parameters, e.g. a "sampling rate" δ and a "learning rate" β . Adjusting β allows one to smoothly interpolate between the greedy algorithm (β near 0) and algorithms which are less responsive but more robust against an adaptive adversary (β near 1). The mathematical analysis in Section VI indicates that setting β very close to 1 guarantees good performance against an arbitrary adaptive adversarial attack; however in many circumstances (e.g. random edge failures) greedy approaches achieve faster convergence, which leads one to expect superior performance from a smaller value of β in such circumstances. These theoretical considerations are substantiated by the experiments in Section VII, where the algorithm is tested in a variety of network topologies with random edge failures, and smaller values of β indeed achieve faster convergence.

II. EXISTING WORK

a) Algorithmic Work: The only algorithmic results that possibly work under this strong adversarial model are based on a computational learning framework [12], [8]. Near optimal learning algorithms, *with reliable global information* for finding a shortest path in a graph, where at each time a different *known* cost is assigned to each edge, were studied in [12], [8]; these solutions have an exponential computational overhead. Schemes with polynomial computational overhead were recently suggested by [21], [11]. The solutions in [12], [8] and [21], [11] correspond to "link-state" routing, and the case where the adversary can only exercise dynamic DOS attacks, but cannot cheat. Byzantine behavior causes such algorithms to collapse. Most recently, "on-demand" routing against an *oblivious* adversary was suggested in [3]. In this model, an adversary cannot cheat and the pattern of cheating and blocking is assumed to be oblivious, i.e., this work does not handle *dynamic* DOS attacks.

b) Reinforcement Learning: The "Swarm Intelligence" paradigm is an approach to routing in distributed networks of cooperative agents, inspired by studying the process by which swarms of ants converge to the optimal route to a food source by progressively reinforcing the successful paths using pheromone secretions. Interest in applications of ant-based routing in mobile ad-hoc networks (MANETs) has risen, and many recent papers have addressed the subject [14], [4]. Gunes et al [15] considers an ant-based approach to routing in MANETs, with a completely reactive algorithm. Marwaha et al. [16] studies a hybrid approach using both AODV and reactive ant-based exploration. Baras et al [4] describes a new algorithm that utilizes the inherent broadcast nature of wireless networks to multicast control and signaling packets (ants). ARAMA [14] uses an analogous approach. Work on the Swarm Intelligence paradigm is described in [13], [6], [20], [9], [7], [17], [18], [5], [10], [14], [4].

III. RESEARCH CHALLENGES

c) Past performance is no guarantee of future success: The problem is that we are making decisions in an online environment, where the online algorithm needs to forward packets by selecting routes while having only information regarding past packets, and no information regarding the future conditions. We make no assumptions about the adversary's behavior

or the sequence of fault patterns generated. Moreover, it is also assumed that there may be a powerful adversary generating the worst possible input sequence for the online algorithm. A fundamental question regarding online algorithms is how to evaluate their performance. It is rare, and in some cases outright impossible that one deterministic algorithm *always* outperforms another deterministic algorithm. One classical method has been to assume a model where the future resembles the past (i.e. a *stochastic* model) regarding packet losses, and compare the performance of different algorithms on the same model. However, in an adversarial setting, there is no reason why the adversary should follow the rules of any particular stochastic model. If anything, a malicious adversary will do exactly the *the opposite* of what our model would predict. (There is no reason to hope that the adversary will not know our model.)

One may be tempted to think that converging to the best fixed route may be easily accomplished by utilizing a “greedy” heuristic: keep track of past error rates on different links, and simply select the path with the smallest overall failure rate, namely the sum of the failure rates of its links. For example, existing work on routing in overlay networks, such as RON [1] runs a greedy strategy for windows of a specific size. The intuition is very strong: extrapolate the past failure pattern into the future. This may work if the failure pattern is static or if there is a statistical model of failures. However this method will fail quite spectacularly in the case of dynamic adversaries. For example, consider 100 options for choosing a relay point between the sender and receiver, which define 100 different paths. At time i , path i (modulo 100) is under the control of the adversary, and is failing all packets; at all other times the path is perfect. The greedy algorithm will always pick the worst path, even though at any moment in time only 1% of the paths are faulty.

To see the principal flaw in this greedy strategy, consider the performance of an investor who tries to follow the best performing stock on the stock market, with the naive assumption that “past performance is a guarantee of future results.” In fact, in an adversarial setting (e.g. the stock market) it may very well be the case that past performance correlates *negatively* with future results, and algorithms must not be fooled into this easy trap.

d) Our path selection must withstand competitive analysis: In general, there may not be an ideal

fault-free path, but yet we can define the best path in terms of accumulated loss on that path. Our goal is to make path selections, such that our overall performance is comparable to that of the “best path.”

Our goal is thus to find a robust randomized algorithm that works well on all inputs, in the sense that the expected behavior of our algorithm is comparable to the optimum fixed path on each input. The goal of this paper is to introduce novel routing algorithms for route selection in an adversarial environment that are provably *optimal* in the sense that the total number of messages lost in our algorithm exhibit a very small additive gap with respect to the *optimum prescient* route assignment. The optimum assignment is made with complete knowledge of the adversary’s actions in the past, present, and future, and has infinite computational power. The only restriction on the optimum route is that it must change somewhat less frequently. The loss of performance of the algorithm that we are seeking should be based on competitive analysis [19]. We compare the performance of our online algorithms to the best static offline selection. Namely the offline algorithm selects a fixed path and uses it during all time steps. Our results bound the difference (referred to in the literature as *regret*) between the cost of the best static offline selection and our online algorithms.

For example, consider a wireless network with 200 users, 1000 potential wireless links, and at least *one* fixed path of 7 hops between sender and receiver that has an average link fault rate of 0.01 %, i.e. 99.99% of the time the whole path is reliable. The only problem is to select one out of around 200^7 possible paths, while only having information about past experiments over these paths. In this case, one can construct a counter-example in which the greedy algorithm may *never* succeed in delivering a *single* message, i.e. it has a 100% fault rate, in spite of always selecting the best of the 200^7 paths so far! The explanation for this counter-intuitive fact is that any deterministic algorithm can be easily fooled by an adversary, forcing it to pick a path that always fails.

In contrast, the “competitive” randomized algorithm that we are seeking should be able to “zero in” on the reliable path, or at least get comparable performance. The algorithm we are seeking should guarantee, for any adversarial behavior (subject to the adversary’s “pledge” to keep some path of length 7 hops being 99.99% reliable on each link), a fault rate of just below $7 \cdot 0.01\% = 0.07\%$ over long sequences.