

SWARM: Cooperative Reinforcement Learning for Routing in Ad-hoc Networks

Eoin Curran

15th September 2003

Abstract

Existing ad-hoc routing protocols are based on a discrete, bimodal model for links between nodes: a link either exists or is broken. This model usually considers only the most recent transmission as determining the state of the link. Unfortunately, this model cannot distinguish transmissions which fail due to interference or congestion from those which fail due to their target being out of transmission range.

This thesis presents a new ad-hoc routing protocol based on a continuous (rather than discrete) model for links within the network. We use a statistical measure of link performance over time to represent the quality of the link. We propose that such a model is required for efficient operation within real-world wireless networks.

In order to define optimal routes in a network with links of variable quality, we model ad-hoc routing as a cooperative reinforcement-learning problem. Cooperative reinforcement-learning describes a class of problems within machine learning in which agents attempt to optimize its interaction with a dynamic environment through trial and error and information sharing. We assign a value to routes that is representative of the costs to the agents using that route. The ad-hoc routing problem is thus expressed as the optimization of the value of routes.

Our model of link quality is a statistical one and requires data to be gathered over time. We devise a learning strategy that gathers information about available routes and the quality of their links over time. This learning strategy operates in an on-demand manner, gathering information only for the traffic flows which are in use, and in proportion to the amount of traffic on those flows. This learning is done in an on-line manner: route discovery operates simultaneously with packet delivery.

Our learning strategy is loosely based on work in swarm intelligence: those systems whose design is inspired by models of social insect behaviour. In particular, we adapt the ant-colony optimisation meta-heuristic as a learning strategy for the ad-hoc routing learning problem. In our protocol, each data packet routed by the protocol causes incremental changes in the routing policy of the network.

We have found that a continuous model of link quality is very beneficial in congested multi-hop networks. Whereas a bimodal link model will interpret any dropped packet as an indication of node mobility and trigger route updates throughout the network, a routing protocol based on a continuous model can respond to dropped packets by incrementally adapting its routing behaviour. In congested network scenarios simulated in NS-2, the performance of our protocol in terms of packet delivery ratio and routing traffic is found to be superior to that of AODV or DSR.