

Ad Hoc Innovation: Distributed Decision Making in Ad Hoc Networks

Timothy K. Forde, Linda E. Doyle, and Donal O'Mahony, CVTR, Trinity College Dublin

ABSTRACT

Mobile ad hoc networks by their nature are highly adaptive systems that can come into existence on an as needed basis. They can grow, reduce in size, fragment, and dismantle as desired. The dynamic and very flexible nature of ad hoc networks can be taken to a further level of sophistication by allowing these networks to retune and adapt themselves according to prevailing network conditions. We are interested in scenarios in which ad hoc nodes must reconfigure as part of a network-wide adaptation process such that network-wide consensus is needed before any change can take place. We present an elegant solution for attaining global consensus in MANETs based on the social science theory of the diffusion of innovations. We present results of the application of this novel approach to an ad hoc network with an adaptive reconfigurable network layer.

INTRODUCTION

The motivation for creating reconfigurable ad hoc nodes is in line with the basic philosophy of the ad hoc network. Ad hoc networks are very much about doing what is needed, when it is needed. In a mobile ad hoc network a node may find itself in a small fast-moving network carrying limited amounts of traffic that suddenly merges with a large network, which results in a huge increase in activity, and hence traffic throughput, at the node. Ad hoc nodes need to be able to deal with these types of changing patterns. One way of doing this is to have as much flexibility as possible within an ad hoc node so as to allow the node to reconfigure in line with the prevailing network conditions and demands. High levels of flexibility can be enabled by treating all parameters associated with the ad hoc node as variables. So, for example, the routing protocol in use, the medium access control (MAC) scheme in use, as well as the individual parameters of these and any other layers are all treated as variables that can be configured and reconfigured as and when needed. It should be emphasized at this point that this article focuses on very dynamic and mobile ad hoc networks rather than on more spatially immobile and

resource-constrained networks such as sensor networks.

There are two scenarios that are of interest in terms of reconfiguration of ad hoc nodes. The first is the case in which a reconfigurable ad hoc node can adapt on a standalone basis without regard to its neighbors. An example of this is the case where a node adapts an internal parameter such as the timeout it assigns to cached routing information — shorter timeout periods are preferable in more mobile environments as they ensure that route information is removed before it becomes stale. The second is the case in which a reconfigurable ad hoc node can adapt only as part of a network-wide adaptation process in which network-wide consensus is needed before change can take place; this is the case of interest here. An example of this is protocol choice. For example, all nodes must agree on the choice of MAC protocol as failure to do so would cause a breakdown in communication. The ability of nodes to reconfigure on a network-wide basis facilitates a hugely advanced level of collective self-organization.

In their purest form, ad hoc networks are completely nonhierarchical distributed structures. All nodes of the network have equal status, acting as routers and forwarding packets to nodes within radio range. There is no central authority and therefore no global view of the network. Hence, to reach a network-wide decision on a choice of network variable, a distributed decision-making process is needed. In a distributed system each node runs a process by which it makes relevant local observations and, on the basis of these observations, makes a judgment on which reconfiguration action is to be taken. The term *relevant local observation* is used to keep the discussion very general. The details of which observations are needed, and indeed which observations can be made, depends on what is being reconfigured and is discussed in later sections. In an ideal situation each node's local observations would be entirely representative of the global scenario, and the nodes would easily come to a consensus as all nodes would see the same conditions and make the same judgment. However, that is not the case. Nodes typically have conflicting views of the network. For example, in an ad hoc network of high den-

The definition of what is good enough evidence to satisfy the requirement of an early adopter will obviously influence the situation. The amount of peer influence needed to make early and late majority decisions also influences the diffusion process.

sity a node at the network periphery may observe a very low node degree that is not representative of the majority of the network where other nodes observe high node density. As a result, conflicting preferences for the course of action to be taken will exist among the nodes. Therefore, a mechanism that deals with conflict and allows a network to reach a consensus is needed.

Finding a mechanism that can deal with the particular challenges associated with distributed decision making in ad hoc networks is a challenge. Nodes may only be aware of their own neighbors, and have no sense of the size and extent of the network. The links in the network can be asymmetrical, and the quality of the links between neighbors can vary. There is no central clock, and any decision-making process must therefore be entirely carried out in an asynchronous manner. The lack of a central clock also means that the decision-making framework must have some mechanism for determining whether the network-wide process of reasoning, making judgments, and resolving conflicts is still ongoing or whether conflict has been resolved. The well-known Fischer-Lynch-Patterson impossibility result clearly delineates the limits of deterministic decision making in systems of the kind characterized by open mobile ad hoc networks [1]. While some interesting approaches exist [2, 3] to deal with the challenges of ad hoc networks, we suggest moving away from the traditional approaches and toward the use of theories grounded in social science. In particular, we focus on the theory of the diffusion of innovations, which provides both an elegant and practical solution.

DIFFUSION OF INNOVATIONS

The concept of diffusion of innovations [4] arises in the study of social and cultural behavior with regard to how innovations come to be adopted or rejected by members of a society. In the decentralized diffusion model, which is of interest here, decisions regarding such matters as when an innovation should be diffused, how it will be diffused, and how it should be evaluated are shared by the potential adopters. New ideas, or innovations, may grow out of the experience of certain individuals or potential adopters, rather than through the specific promotion of centralized change agents. The theory of diffusion of innovations is explained in terms of the stages involved in the process of diffusing an idea and in terms of the roles played by the individuals in that diffusion process. This way of looking at the spread of innovation is useful and therefore briefly described here.

The theory of diffusion of innovations associates five stages with the innovation-decision process. These stages are knowledge, persuasion, decision, implementation, and confirmation. The **knowledge stage** is associated with exposure of the decision-making unit to the innovation. The **persuasion stage** occurs while the unit attempts to develop its own opinion about the innovation or be persuaded by others of its value. The **decision stage** is associated with the acceptance or rejection of the innovation. The **implementation stage** occurs when the innovation is put into use,

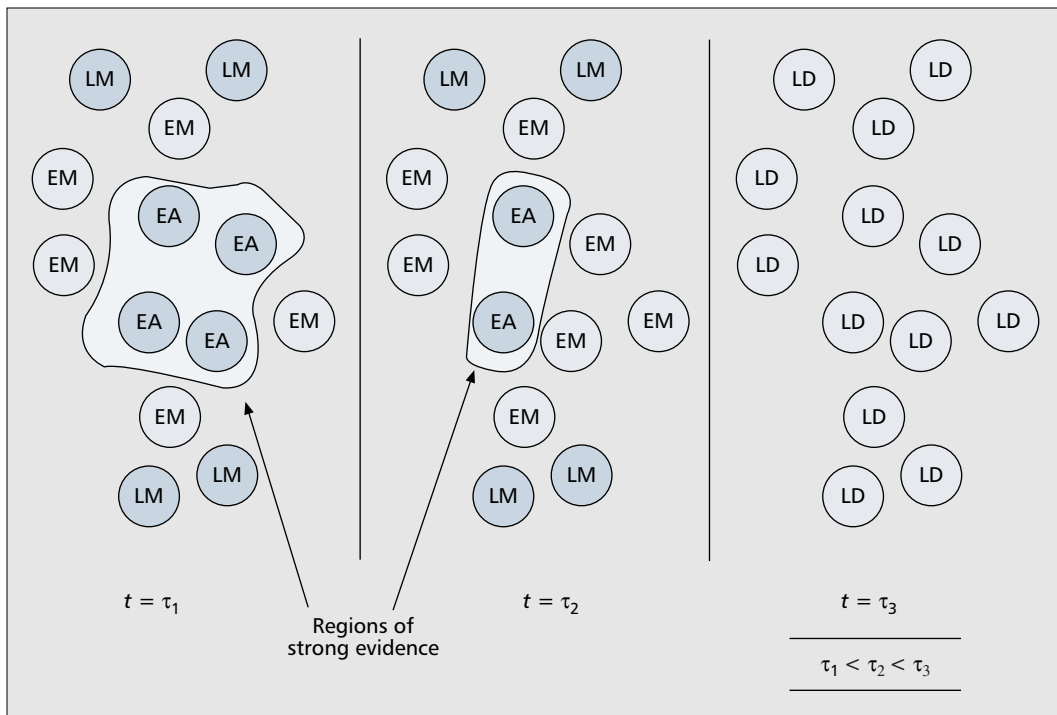
and the **confirmation stage** involves seeking reinforcement of the innovation that has been made.

To explain the spread of innovation, the theory of diffusion of innovations defines five classes of actors who play a role in the process. The first class involves **innovators**. These are generally active information seekers about new ideas and would be responsible for creating the knowledge. During the persuasion stage four other actors come in to play. **Early adopters** have a high degree of connectedness and make reasoned, evidence-based decisions. They decrease local uncertainty about an innovation by adopting it and conveying messages to other near peers by means of interpersonal networks (i.e., neighbor-to-neighbor connections). The **early majority** do not lead opinion as the early adopter does, but through frequent interaction with their peers, they tend to follow the early adopters. The **late majority**, being skeptical, wait until most of their social system have already adopted. **Laggards**, as the name suggests, are the last in a social system to adopt an innovation. They have no opinion leadership qualities and are very local in outlook.

The adoption of an innovation percolates through the system as early adopters adopt the innovation, followed by the early majority, followed by the late majority. This type of interactive innovation, as it is known, tends to encourage the community to conform. A unit adopting an innovation will quickly go from the persuasion stage to the decision stage. Units do this independently and do not need to wait for others to make their decisions. How long implementation of the innovation and subsequent confirmation takes depends on the innovation being adopted. How widespread the adoption of the innovation is depends on the distribution of early adopters in the system. The definition of what is good enough evidence to satisfy the requirement of an early adopter will obviously influence the situation. The amount of *peer influence* needed to make early and late majority decisions also influences the diffusion process.

MAPPING FROM DIFFUSION OF INNOVATIONS TO AD HOC DISTRIBUTED DECISION MAKING

The key contribution in this article is the mapping of the diffusion of innovations theory to a distributed decision framework for ad hoc networks. The theory of diffusion of innovations provides a rich understanding of the spread of innovations in a distributed network. Returning to the goal at hand, we want to enable reconfiguration of some network-wide parameter in the ad hoc network. Thus, the innovation to be spread in our case is the desired value of that parameter. We want to create a mass belief that the parameter value x should now be changed to value y for the good of all. Of course, we only want this to occur when indeed it is true that the innovation will actually improve performance. The diffusion of innovations theory has been useful in terms of providing a means of defining



A judgment will only achieve hard-state status and enter the decision stage if two conditions are fulfilled. First, it must be an EA judgment, i.e. one based on strong evidence, and second, it must be sustained for a required judgment-persistence period.

■ Figure 1. Persuasion ebbs and flows.

the stages in the decision-making framework and creating a decision model by which individual nodes can make a choice.

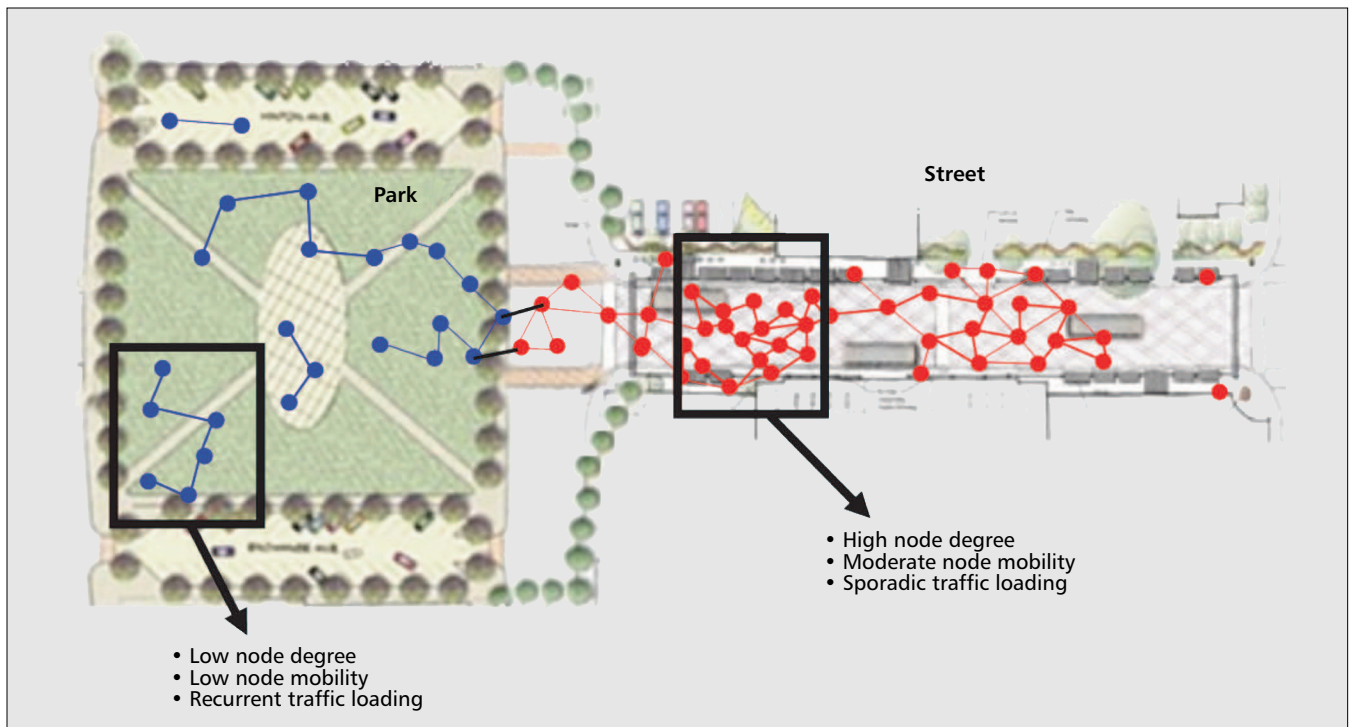
From the innovation-decision process three stages were deemed as relevant in our situation: persuasion, decision, and implementation. Knowledge does not need to be sought as all nodes know the options available to them for the variables they may be seeking to set. A full confirmation stage was also deemed unnecessary in our case as if the newly reconfigured state subsequently becomes nonoptimal, the process simply begins all over again, a concept similar to the replacement discontinuance process discussed in [4].

A node in the ad hoc network starts at the persuasion stage. During the persuasion stage the node attempts to make a judgment about its preferred value of the parameter to be reconfigured. From the actor classes, four actor types were selected to form a decision model to be used in the persuasion stage when nodes attempt to make a judgment. A node attempts to make a judgment first and foremost based on strong evidence. We have chosen to define high correlation between a node's observations and its neighbors' observations as strong evidence. If the node succeeds, it is labeled an early adopter (EA). If the node cannot make an evidence-based decision, because according to our definition there is low correlation between its observations and its neighbors' observations, a node attempts to make a weaker judgment based on what its peers think. The first type of weaker judgment is labeled early majority (EM) and can be made if a node has more than a certain threshold of neighbors holding the same particular judgment who are themselves EAs. Failing this, the node attempts to make a late majority (LM) decision. The criterion for this is to have a

certain threshold of neighbors holding a common judgment who are themselves either EAs or EMs. If no judgment can be made, a node is labeled a laggard (LD).

This judgment-making process continues in a parallel fashion throughout the network. All of the judgments that are made are soft state decisions. Nodes can oscillate between EA, EM, LM, and LD status as they move in and out of contact with other nodes and are influenced by different nodes' views and whether strong evidence is available or not. This is very different from the case described in the basic theory of diffusion of innovations as, in the basic theory, once an entity is persuaded to adopt the innovation, the innovation is adopted, and there is no *waiting to see* if conditions will change and the decision should be reversed. Figure 1 shows a network at three different times. Strong evidence exists initially, at time $t = \tau_1$, and this results in EAs coming into being. The evidence then diminishes at $t = \tau_2$, resulting in fewer EAs, and then disappears completely at $t = \tau_3$, resulting in no EA emerging, and thus no subsequent EMs and LMs.

A judgment will only achieve hard state status and enter the decision stage if two conditions are fulfilled. First, it must be an EA judgment (i.e., one based on strong evidence); second, it must be sustained for a required *judgment-persistence* period. By insisting on this, unnecessary reaction to temporary and short-lived ebbs and flows in network conditions can be avoided. Once the decision stage is entered, it is followed immediately by the implementation stage during which execution of the reconfiguration occurs. The EA or EAs have a domino-like effect on the network with all nodes in their range of influence (i.e., nodes that made EM and LM decisions based on those EAs) falling in line by



■ **Figure 2.** Changing conditions in an urban space.

adopting the change. Allowing the EAs to trigger the reconfiguration is very suitable in systems where there is no central clock. Nodes can operate happily in persuasion mode by checking observations at intervals that are set locally until the decision to change is triggered.

For the network to reach global consensus, a critical mass of EAs must emerge. It can also be the case that opposing views emerge in different sections of the network. For example two EAs with opposite views and associated spheres of influence could develop. In such cases a fragmentation of the network, rather than a universal setting of the variable parameter, can be the best solution. Solutions exist to permit internetworking and allow communication between the two fragments, but are beyond the scope of this article; distributed internetworking demands the existence of dual-protocol routers and some interdomain routing extensions to the existing protocols [5].

The setting of the rules defining strong evidence for EAs and the threshold levels for all EMs and LMs has a direct effect on how difficult it is for EAs to emerge and how wide their sphere of influence can be. Setting appropriate rules requires very good knowledge of the specific problem space as the choice of threshold values and judgment-persistence times affects the heterogeneity of the system's configuration. Setting high thresholds demands a strong correlation between the nodes' observations. Lower thresholds, which allow for less smoothness (or more heterogeneity) in the network observations, may lead to too much network fragmentation. The influence of these values on the ease with which the decision-making process advances should be the subject of cognitive reasoning processes in which

nodes learn from previous decision-making experiences and perform risk evaluation or cost-benefit analysis so that the decision-making process tuning parameters are appropriate to the application setting. Examples of values we use are given next.

APPLICATION EXAMPLE

The diffusion of innovations approach is best illustrated through a practical example. In our case study we created a reconfigurable ad hoc network layer that allows nodes to dynamically configure and utilize the most suitable ad hoc routing protocol for the prevailing network conditions. A wide body of work exists in the area of ad hoc network routing, and it has been well established in the literature that a *one-size-fits-all* approach with regard to the choice of optimum routing protocol does not suffice [6–8]. Rather than continuing to search for the optimal routing protocol, our approach allows us to take existing well-established routing protocols and use the protocol most suited to the prevailing conditions. In other words, the choice of routing protocol is now a variable that can be set and reset as desired by the network. Figure 2 depicts a typical scenario in which such changes in network conditions occur. In the park setting the ad hoc network is characterized by an open space filled with slowly moving nodes, running a protocol indicated by the blue color. The adjacent street, into which nodes may move, consists of an ad hoc network largely characterized by fast-moving, more dense groups of nodes that choose to run the red protocol, which is suited to these conditions.

The diffusion of innovations protocol has been implemented in a real ad hoc network in

Trinity College Dublin. The network, known as the Dublin Ad Hoc Wireless Network (DAWN) [9, 10], has been designed to facilitate research in the area of ad hoc networking. At the core of DAWN is a dynamic modular communication stack that runs on each of the nodes of the ad hoc network. Layers of the stack can be independently designed in a standalone fashion. A generic layer interface allows the dynamic assembly of these layers to form a network communication stack consisting of the relevant hardware and software elements. The interlayer interface is very simple, consisting of primitives to send information upward or downward through the stack. A diffusion-of-innovations-based decision-making layer was designed for DAWN. We have implemented the Dynamic Source Routing (DSR), Optimum Link State Routing (OLSR) and Ad Hoc On-Demand Distance Vector routing (AODV) protocols. DSR and AODV are reactive protocols and OLSR is a proactive protocol. DAWN contains a mechanism for swapping in and out the protocol in use at each node when triggered by the diffusion-of-innovations-based decision-making layer.

In terms of realizing the case study, two major pieces of work were carried out. First, the decision-making layer was designed. Second, a large body of work was produced, focusing on understanding the relationship between the effects of the prevailing network conditions and the performance of a given routing protocol, and identifying appropriate means of measuring those conditions. From the available literature, exemplified by [6–8], and our own simulations, the major factors affecting choice of routing protocol are found to be node mobility, node density, and traffic conditions. Node link duration, node degree, and node traffic loading are the local observations that can be made to get a measure of mobility, node density, and traffic conditions. The range of values of these parameters and how these parameters map to a specific protocol choice is beyond the scope of this article; full details are presented in [5]. The implementation of the diffusion-of-innovations-based protocol is sensitive to the underlying operation of the ad hoc network. No extra burden is placed on the network as opportunistic communication is used to get information about neighbors. What this means is that information is never asked for but simply obtained in the course of normal operation of the routing protocols. So when we say *observations are made* this does not imply an active seeking out of neighbors' views but an opportunistic gathering of observation information based on the signals (i.e., packets) being generated by the network layer's primary protocols (e.g., the routing protocol).

All observations gathered by a node are passed through an exponentially weighted moving average (EWMA) flip-flop filter [11]. The settings on this filter determine whether agility or stability are more important in the system. Again, the choice of filter stability, like that of decision-making thresholds, could be the subject of cognitive reasoning by the network's nodes. The observations made by a node are then cor-





Figure 3. Reconfiguration in response to changing node mobility.


related with the observations made by its neighbors. High correlation between these correlations implies *strong evidence*, the type of evidence needed for a node to be an EA. The condition for a node becoming an EM has been set to a simple majority of EA neighbors, and the condition for becoming an LM is a simple majority of EA and EM neighbors. As stated earlier, each of these rules can be set as appropriate.

Presenting results of the performance of the decision-making framework is complex given the number of variables involved. The three sets of results shown in Figs. 3, 4, and 5 attempt to highlight some of the interesting outcomes. The three figures correspond to three scenarios involving 30 DAWN nodes. DAWN can work over real hardware or sit on an emulator. In these experiments the emulator was used, and a random waypoint (RWP) mobility model was selected. For clarity of demonstration, in each scenario one metric that influences protocol choice was dramatically changed at time $t = \tau_x$ after a runtime of 400 s even though in reality a combination of metrics may change at any one time. Each figure illustrates the state of a single node that was tracked for the duration of the experiment. In each of Figs. 3, 4, and 5, a plot of what the node observes, what its neighbors observe, what type of judgments it is making (EA, EM, etc.), and whether it triggers a change is shown. This is a large amount of information on each graph, but the only way to depict the interrelated events. In Figs. 3, 4, and 5 one trace shows the observation made by the node of link duration, node degree, and node traffic loading, respectively. Another shows the corresponding observations of the neighbors. In all three cases presented here, the EWMA filter has a gain of 0.1 and a moving range window of size 16. Decreasing the gain or increasing the window size would result in slower tracking of changes.

In Fig. 3 the node is initially running the OLSR protocol. The dashed lines indicate the bands for low, moderate, and high mobility. The node observes moderate level mobility at point A as do its neighbors. Strong evidence exists, and the node can therefore make an EA judgment. In this case the judgment is to keep running OLSR, so no action is needed. At point B



■ Figure 4. Reconfiguration in response to changing node degree.



■ Figure 5. Reconfiguration in response to changing traffic loading.

the node's observations and its neighbors' diverge. The node itself sees low mobility, and the neighbors observe moderate mobility, meaning that no strong evidence exists. The node can only make EM or LM judgments (recall the judgments of its peers are affecting this state). At time $t = \tau_x$ a significant change in mobility is introduced. The RWP model changes from velocity range 1–5 m/s with pause time 20 s to velocity range 5–10 m/s with pause time 10 s at this stage. However for a short time strong evidence still does not exist to suggest that this has occurred. At time $t = \tau_y$ both the node's and the neighbors' observations indicate agreement, and the node again makes an EA judgment, this time opting for AODV, which better suits the higher mobility conditions. The EA judgment is sustained for the requisite period (i.e., until time $t = \tau_z$), which triggers a hard state decision. At this stage the changeover to AODV is implemented.

In Fig. 4 the node is also initially running the OLSR protocol. The dashed lines indicate the bands for low, moderate, and high node degree. The node observes moderate level node degree at point A as do its neighbors. Strong evidence exists and the node can therefore make an EA judgment. In this case the judgment is to keep running OLSR so no action is needed. In area B of the diagram the node's observations and its neighbors' diverge. The node itself sees high node degree and the neighbors observe moderate node degree; conse-

quently no strong evidence exists. The node can only make EM or LM judgments. At time $t = \tau_x$ a significant change in node degree is introduced. However, for a short time strong evidence still does not exist to suggest that this has occurred, and a mixture of EM and EA judgments continue to be made. At time $t = \tau_y$ both the node's and the neighbors' observations indicate agreement, and the node again makes an EA judgment, this time opting for DSR, which better suits the less dense network conditions. The EA judgment is sustained for the requisite period (i.e., until time $t = \tau_z$), which triggers a hard state decision, and the changeover to DSR is implemented.

In Fig. 5 the node is initially running the AODV protocol. There is only one dashed line marking the difference between recurrent traffic loading above the line and sporadic traffic loading below the line. We have created a traffic metric that indicates the average amount of time a node spends supporting the creation and maintenance of routes to active destination nodes. We characterize traffic loading as being sporadic (i.e., short-lived traffic streams to many destinations) or recurrent (i.e., longer-lived traffic streams to fewer destinations). The node observes the recurrent traffic loading conditions at point A, as do its neighbors. Strong evidence exists, and the node can therefore make an EA judgment. In this case the judgment is to keep running AODV, so no action is needed. In area B of the diagram this continues to be the situation. At time $t = \tau_x$ the traffic is changed from long-lived recurrent traffic flows to short-lived sporadically assigned traffic flows. The observations of the node and neighbors stay in close agreement, except for a brief time before $t = \tau_y$. At time $t = \tau_y$ both the node's and neighbors' observations indicate agreement, and the node again makes an EA judgment, this time opting for OLSR which better suits the sporadic traffic loading applied to the network. The EA judgment is sustained for the requisite period (i.e., until time $t = \tau_z$), which triggers a hard state decision, and the changeover to OLSR is implemented.

The network recovery time (NRT) [12] is the time taken for a network to recover after a condition indicates reorganization of the network. In Fig. 3 the average time for all of the nodes in the group to reflect the changes in the mobility conditions was 11 s with another 15 s needed to make and adopt the decision (NRT = 26 s). In Fig. 4 the average time for all of the nodes in the group to reflect the changes in the network density conditions was 11 s and an additional 13 s to enact the decision (NRT = 24 s). In Fig. 5 the average time for all of the nodes in the group to reflect the changes in the traffic loading was 8 s and an additional 11 s to enact the decision (NRT = 19 s). The NRTs are presented here only as a means of noting that the network did achieve consensus. The protocol has been designed to work in a *just-in-time* fashion [10], so the overall objective is for each node to reconfigure when needed, rather than for the network as a whole to synchronously enact change as quickly as possible.

CONCLUSION

Looking beyond traditional sources and into the area of social science has led to an alternative but very elegant and effective means of providing a distributed decision-making framework for mobile ad hoc networks. Catering for network-wide reconfiguration through the use of this technique results in ad hoc networks that are more sensitive to their surroundings and more capable of true self-organization. The example presented here focused on flexibility at the network layer and demonstrated how a real ad hoc network used a *diffusion-of-innovations*-based protocol to reconfigure when network conditions change. Clearly, an understanding of how conditions are changing, how the changes can be observed, and the implications of the reconfiguration are completely key to developing this approach, and indeed any other adaptive or reconfigurable system.

ACKNOWLEDGMENT

This material is based on works supported by Science Foundation Ireland under Grant No. 03/CE3/I405.

REFERENCES

- [1] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of Distributed Consensus with One Faulty Process," *J. ACM*, vol. 32, no. 2, Apr. 1985, pp. 374–82.
- [2] L. Briesemeister and G. Hommel, "Localized Group Membership Service for Ad Hoc Networks," *2002 Int'l. Conf. Parallel Processing Wksp.*, Aug. 18–21, 2002, Vancouver, B.C., Canada.
- [3] S. Nesargi and R. Prakash, "MANETconf: Configuration of Hosts in a Mobile Ad Hoc Network," *Proc. 21st Annual Joint Conf. IEEE Comp. and Commun. Socs.*, vol. 2, 23–27 June 2002, pp. 1059–68.
- [4] E. M. Rogers, *Diffusion of Innovations*, 4th ed., Free Press, 1996.
- [5] T. K. Forde, "Flexibility in Ad hoc Networks," Ph.D. thesis, Univ. of Dublin, Trinity College, Ireland, 2005.
- [6] S. R. Das, C. E. Perkins, and E. M. Royer, "Performance Comparison of Two On-demand Routing Protocols for Ad Hoc Networks," *IEEE Pers. Commun.*, vol. 8, no. 1, Feb. 2001, pp. 16–28.
- [7] D. D. Perkins, H. D. Hughes, and C. B. Owen, "Factors Affecting the Performance of Ad Hoc Networks," *Proc. ICC 2002*, vol. 4, pp. 2048–52.

- [8] L. Viennot, P. Jacquet, and T. H. Clausen, "Analyzing Control Traffic Overhead versus Mobility and Data Traffic Activity in Mobile Ad-hoc Network Protocols," *Wireless Networks*, July 2004, vol. 10, no. 4, Kluwer, pp. 447–55.
- [9] D. O'Mahony and L. E. Doyle, "An Adaptable Node Architecture for Future Wireless Networks," *Mobile Comp.: Implementing Pervasive Info. and Commun. Tech.*, Kluwer Series in Interfaces in OR/CS, 2002, pp. 77–92.
- [10] T. K. Forde, L. E. Doyle, and D. O'Mahony, "Self-stabilizing Network-Layer Auto-Configuration for Mobile Ad Hoc Network Nodes," *Proc. IEEE Int'l. Conf. Wireless and Mobile Comp., Net. and Commun.*, vol. 3, 22–24 Aug. 2005, Montreal, Canada, pp. 178–85.
- [11] K. Minkyong and B. Noble, "Mobile Network Estimation," *Proc. 7th Annual Int'l. Conf. Mobile Comp. and Net.*, Rome, Italy, pp. 298–309.
- [12] M. W. Subbarao, "Ad hoc Networking Critical Features and Performance Metrics," *Adv. Wireless Technologies Div.*, NIST, Oct. 7, 1999.

BIOGRAPHIES

TIM FORDE (timforde@gmail.com) is a research assistant at the Centre for Telecommunications Value-Chain Research (CTVR), Trinity College Dublin, Ireland. He received his Ph.D. degree in electronic engineering from Trinity College Dublin in 2005. He has engaged in various research activities focusing on mobile ad hoc networks, emerging network architectures, and dynamic spectrum access techniques. Currently, he is working on developing a novel cognitive radio test platform for use in licensed Irish radio spectrum.

LINDA E. DOYLE (ledoyle@tcd.ie) is a lecturer in the Department of Electronic and Electrical Engineering at the University of Dublin, Trinity College, Ireland. She leads the Emerging Networks strand of the Centre for Telecommunications Value-Chain Research (CTVR). One of the main projects she is currently pursuing, the Plastic Project (www.plasticproject.net), focuses on distributed reconfigurable and cognitive wireless networks for dynamic spectrum management.

DONAL O'MAHONY (Donal.omahony@cs.tcd.ie) is an associate professor in computer science at Trinity College Dublin. At Trinity he built up a successful research group in networks and telecommunications. This group has been influential in developing the concept of fourth-generation mobile systems. He spent the year of 1999 as a Fulbright Fellow at Stanford University, California, and was made a fellow of Trinity College in 2001. In 2004 he led a team to establish the Centre for Telecommunications Value-Chain Research (CTVR), a major multi-university research centre funded by Science Foundation Ireland and established in association with Bell Labs. He is now full-time director of this centre.

Clearly, an understanding of how conditions are changing, how the changes can be observed, and the implications of the reconfiguration are completely key to developing this approach, and indeed any other adaptive or reconfigurable system.