Mesh networks and the spectre of selfishness

Bachelor's Project Thesis

Marcin Korecki, m.korecki@student.rug.nl

Supervisors: prof. dr. L.C. (Rineke) Verbrugge

Abstract: The following study focuses on the wireless ad hoc networks from an agent-based perspective. In a simulation, agents with different strategies such as different levels of selfishness, tit-for-tat or battery-based compete and cooperate in a model of an ad hoc network. The dominant strategy in runs where only different levels of selfishness are allowed is clearly the selfish behaviour. However, introduction of more advanced strategies allows to some extent to combat the selfish element. Of special notice are the battery-based approach and a hybrid of battery based and tit-for-tat approaches. The findings give hope that the introduction of widely available ad hoc networks might at some point be possible. The users when given full control of their devices might be able to develop effective strategies that would allow for the networks overall to be effective and feasible.

1. Introduction

A few discoveries can hope to match the omnipresent and prevalent influence the Internet has continued to exercise ever since that fateful day at the end of October of 1969, when the first text was transmitted over then budding ARPANET. From that day, countless minds have been affected by that ever-evolving phenomenon. There is little doubt in our minds that the effects of this singular invention will continue to change our everyday lives and it is up to us whether it will be for the better or for worse. When browsing the seemingly infinite numbers of websites, accessing a variety of services, researching topics, googling or entertaining ourselves it is easy to forget about the intricate and vast infrastructure that allows us to do all those things with, characteristic of the Internet, lightness. It is easy to imagine the data packets effortlessly floating in the air from one device to another as if by magic. This idealised vision is far from the truth. The Internet as we know it exists through conscious and continuous effort of engineers, scientists and technicians. After all, the whole thing is a highly centralised network of devices and very many physical connections which have to be managed and controlled. We have already mentioned the Internet's aptitude for change, which is especially visible in all the different uses for it that are still being developed. It is our strong belief that the physical architecture which allows the Internet to function should not be allowed to stagnate.

1.1 The Internet's architecture

The different technological solutions that are employed in the inner workings of the Internet are many. The common methods include coaxial cable, copper wires or fibre optics connections as well as WLAN, cellular and telephone technologies. With the continuously increasing numbers of Internet-capable devices and especially mobile devices we venture to guess that wireless technologies are going to play increasingly large roles in providing people with access to the Internet. A vast majority of wireless local area networks (WLAN) are based on the IEEE 802.11 standards (referred henceforth by its brand name Wi-Fi) and operate in the infrastructure mode. This mode entails the existence of a base station (for example a router) that takes the role of an access point through which connected devices (nodes) can communicate. The base station will most often be connected via a wire or fibre connection to a wider network. A simplified model of operations follows: the network consists of devices which are connected to the base station which is in turn connected to the wider network to which other base stations are also connected. A single device first communicates with its base station, which relays the message to the wider network and receives the potential response to be relayed back to the device (for the details of this process the reader is encouraged to study the OSI model of network architecture, ISO/IEC standard 7498-1, 1994). This elegant system allows millions of users to access the Internet and get their much needed data packages quickly and efficiently. It is not, however, without its own disadvantages. The described system is highly centralised and prone to single point of failure mishaps. Any damage or a successful attack on the base station can very easily cut off access of all the devices that are reliant on it. Another issue is the range and capacity limitations of the base station devices, which in many cases (for example on campus networks or blocks of flats) need to be densely spaced. This may lead to interference between the base stations as well as lower quality of service. The users whose devices need to disconnect and reconnect to different stations might experience episodes of being disconnected from the network and their devices' battery life might decrease more rapidly.

1.2 Mesh networks

A potential solution to some of the mentioned problems may lie in the introduction of ad hoc networks in place of WLAN. A wireless mesh (also referred to as as-hoc network) network works in a manner similar to the previously described wireless networks but does not rely on the existence of base stations. Such network is highly decentralised and the routing responsibilities, which in conventional WLAN are performed by the router, are effected by the individual connected devices (Hekmat, 2006). A simplified model of an ad hoc network is as follows: the devices participating in the network keep track of other participating devices (often referred to as nodes), if data packets are to be sent, a path is determined that can connect the two communicating devices directly or via other nodes on the network. Thus the message 'jumps' from one device to the next finally reaching its destination. Any number of devices on such a network can at the same time be connected to the conventional WLAN and thus potentially relay the data over even larger distances. In such a case the ad hoc network can simply help to extend the base station's range.

Considering the fact that as of 2016 there are 17.5 billion connected devices in the world (Statista, 2016) and 2 billion smartphone's users (Statista, 2016), it is easy to imagine ad hoc networks becoming feasible. The potential of ad hoc networks with the focus on smartphones as primary nodes is of special interest. All of the modern smartphones are equipped with both bluetooth and Wi-Fi capabilities both of which can be employed to create an ad hoc network. Having taken into account the density of population in the large urban areas of the world one could conclude that ad hoc networks may very easily be able to provide access to the Internet (or any other network for that matter) to vast amounts of people. The ad hoc networks, however are not without problems of their own. Firstly, for the network to be fully connected a certain minimum amount of devices need to participate. Moreover, those devices must remain within each others range. If we consider that the devices on the network can be in constant motion ,it can become quite difficult for the network to remain well connected in its entirety. Secondly, the ad hoc networks are quite prone to security issues such as DDS attacks (this is beyond the scope of this paper). Lastly, the lack of any main coordinating device forces all of the devices to take part in the routing process. There exist many routing algorithms designed for the purpose of ad hoc networks specifically. Such algorithms are usually designed with the focus on scalability, reliability, flexibility, throughput, load-balancing and efficiency (Vijayakumar, et. al, 2012). Only recently another factor has started to be taken into account when comparing different algorithms - that is the battery life (Sangwan, et al, 2016 and Yoshimachi, et al. 2016). If we consider smartphone devices to be our primary participants in a hypothetical ad hoc network it becomes quite clear that the success of such a network is strongly related to how the participation will affect individual devices' battery life. If the network causes the devices' battery life to decrease dramatically the devices will turn off and thus the network will start losing precious nodes until it eventually becomes virtually unconnected. Another consideration, one that we will focus on in this paper, is the willingness of participants to volunteer their device's battery life to other participants.

1.3 Overview of the study

Let us imagine a hypothetical ad hoc network that encompasses an area of a square in the centre of a densely populated city. All of the people within the square participate in the network but can choose whether to allow other people's data packets to go through their phones. If we assume that the goal of every participant is to be able to send and receive data packets (addressed to or from one's device) at whim, we are presented with an interesting dilemma. On one hand, the selfish strategy to prevent other participants from using our device is beneficial as our battery life is depleted less quickly. On the other hand, if all the participants become selfish the network ceases to function and no one can receive nor send data. We are therefore interested if in a situation where the participants are allowed to control the amount of 'foreign' data that goes through their phones there is a chance of achieving a stable network. We believe that the success and more widespread adoption of ad hoc networks is dependent on the willingness of the population to use it. Moreover it is quite clear that for the network to be appealing to new users, participants should not be forced into forfeiting control over their devices. Hence there is a clear need for investigation of potential emergent behaviours in groups of users. This leads us to the first part of our research question: "What is the dominant strategy for participants in an ad hoc mobile network?". It is our hypothesis that the dominant strategy might tend toward the selfish approach since even though the selfish behaviour will ultimately lead to the destruction of the network, which will affect all the participants equally negatively, the selfish participant will be able to generate more profit before the network ceases to function. (Hardin, 1968). The second part of our research question will ask "Is there a reward/punishment system for participants that can improve the longevity of an ad hoc network?". If the emergent behaviour proves to be destructive, we will look into reward systems that could allow the network to function and if the behaviour can sustain the network we will investigate systems that could improve it. We hypothesise that there indeed exists such a system that can allow the voluntary ad hoc network to be sustainable. This is based on the findings that in many similar game theory problems such as iterated prisoner's dilemma (or extensions such as public goods game), there usually exists a reward system that can radically decrease the payoff of selfish behaviour thus limiting or eradicating it (Jurišić, et al. 2012 and de Weerd, Verbrugge 2011).

We propose to test our hypothesis with the help of a model of the ad hoc network. Since we are only interested in the behaviour of participating agents we will introduce certain simplifications to the general description of the network through routing algorithms and signal propagation. The model will be based on a densely populated square in the city centre and we will assume that the network is functioning perfectly and without disruption, it is only at the mercy of its participants. The agents will have different strategies ranging from fully selfish to fully altruistic. The agents will use the network at individually random intervals but the weight of data packages introduced into the network by each agent will be approximately the same. The effectiveness of each agent will be measured in terms of the data packages it succeeded to send or receive. A genetic algorithm will be used to determine the agents' strategies throughout subsequent epochs and the model will be run for 10 eras (1 era - the number of time steps needed for half of the population's battery to become 0).

2. Methods

In the following section the design of the simulation application used in this study will be explained and the design choices will be motivated. It is worth reminding that the main goal of this simulation is to allow us to investigate the behaviour of agents of differing levels of selfishness participating in an ad hoc wireless network. As such we do not need our model to provide a strictly realistic description of the network architecture. As long as the qualities of the network affecting the behaviour of participants we are interested in are well defined the model should be sufficient for our purpose. Let us first describe the general requirements of our model and then follow with a detailed explanation of their implementation.

The most important quality of the ad hoc network is the fact that its success is irrevocably associated with the number of participants related to the spatial distribution of the network (so the density of the nodes/participants). The bigger the network is supposed to be - the longer distances we want our data to travel - the more participants we need. In the same way, if we assume that participants may exhibit different levels of participation, we conclude that the network requires a certain minimum number of participants who are willing to allow other nodes' data to be relayed through their devices. This is the quality of the ad hoc network we are mostly interested in and we will focus on implementing it in our model. Subtleties such as signal propagation and possibly radical heterogeneity of devices participating in the network is of less concern to us. While it would certainly be interesting to investigate a highly realistic model of such a network, it would not help us in answering our question in any way and is thus beyond the scope of this paper. The components of our model that will describe our network are then: the routing algorithm that will provide us with information such as the connectedness of the nodes of the network, the agents (and their world) participating in the network and the fitness function along with an evolutionary algorithm allowing our network to evolve. The model will be turn based and all of the actions will occur in their specific order as will be explained in the following sections.

2.1 Agents and their world

The agents in our application will model the mobile participants of an ad hoc network. The function of the agents will be to move in the simulation world and use the network. The movement of the agents will be simulated with a very simple random-walk model. We need the agents to be mobile but we do not need their mobility patterns to be representative nor realistic as we are more interested in how their strategic choices affect the network. Each turn the agents can walk one step in each of the four cardinal directions or stay in place with an equal chance of 20%. Thus we achieve a model in which it is quite difficult to predict where a given agent will be at any moment. This we believe is accurate enough for a mobile ad hoc network model.

The world of the agents can be a rectangle of any integer size and the agents always move by one unit of size. In a situation where an agent walks over the border of our rectangle it disappears from the network and with each subsequent turn has an increasing chance to return from a random side. The chance of an agent returning is 10% in the first turn, 20% in the second, up to the agent returning with 100% chance in the tenth turn after its disappearance. While the agent is gone from the world it does not take part in any of the network activities.

The part of the agents' behaviour we are most interested in are their evolving strategies. For the first iteration of our study, we propose a binary choice space with two strategies possible - either completely selfless or completely selfish. The former will always route messages and the latter will never route messages. This simple distinction is, we believe, fitting for a first step, a sort preliminary testing of the waters. In the next section we will propose and investigate more complex and nuanced strategies.

2.2 Messages

Each turn, each agent has a 25% chance of deciding to send a message to a randomly chosen agent that is also present in the current world state. The message is considered to be sent successfully if the sender and the receiver are both connected through the network (further explained in the section 2.3). The agents store how many messages they wanted to send and how many they sent successfully for the purpose of calculating their effectiveness.

Each attempt at sending a message is associated with a cost paid by the agent sending the message. Here we introduce the battery life as a resource. Each agent starts a given epoch with a fully charged battery (1000 points). Each attempt at sending a message decreases the battery life by 5 points. Similarly routing a message of another agent costs 3 points. Since we use the battery life as a halting condition for each run of a given world the costs basically affect how long we will run our simulation - (ie. when the majority of agents' batteries are dead). Depending on our needs those costs can be easily tweaked and changed. The relation between the message sending and routing costs affects how penalising the routing is and so how costly it is for an agent to be selfless.

2.3 Routing

There is a wide variety of algorithms that can be used in ad hoc networks. Since we are not investigating the effectiveness of the algorithms itself we are not bound to any specific approach in this department. The only requirement that we pose is that our routing algorithm provides the agents with always optimal paths to the available receivers.

The protocol that we have chosen to implement in our simulation is the distance-vector routing protocol which uses the Bellman-Ford algorithm (Bellman, 1985). The algorithm is quick enough for our purposes and meets our main condition of providing agents with best possible routes. The workings of the algorithm as implemented in our simulation are as follows. Each agent has access to a routing table which lists all the other agents in the simulation, the shortest routes to them (or infinity if such routes do not exist) and the next node (agent) that needs to be visited to reach the given final node. At the beginning of each turn, after all the agents have moved each agent updates its table - all its neighbours are placed in the table with a distance of 1 and all the non-neighbours are given a distance of infinity. What follows is the crux of the algorithm - each agent advertises its table to all its neighbours which receive the table and update their own tables based on it. Hence if they are offered a better route through the currently advertising neighbour than the one they are currently aware of they will edit their table accordingly. This process is repeated until none of the agents need to make updates in their tables. In a real life scenario the Bellman-Ford algorithm poses a serious problem - the possibility of appearance of routing loops. This problem can be solved in a couple of ways - we however do not need to bother as our simulation allows us to freeze time. Hence we only begin the process of updating the routing tables after all the agents have moved and wait with our simulation until all the tables have been updated.

It is worth noting that when a selfish agent advertises its table, it only provides the other agents with routes to itself - as it disallows any other traffic to go through its device. Another important point is that for two agents two be considered neighbours they must be within 10 units of each other.

Now if an agent wants to send a message, it looks up the receiver in its routing table. If the distance to the receiver is not infinite it is possible to reach the agent and so the message is routed to the next node as provided by the table. The next node then pays the routing costs and repeated the process of looking up the receiver and the next node in the table. The process continues until the receiver node is reached.

2.4 Fitness

The fitness of an agent is the ratio of the messages successfully sent by it to the messages it wanted to send. Hence the maximum fitness is 1 and the minimum fitness is 0. The global success of the network is the average success rate of the participating nodes. However another important factor in our considerations is the battery life of each agent. We decided to not include the battery life directly in the fitness calculations but rather use it as an indirect factor in a way that we hope mimics realism as much as possible. The agent whose battery life drops to 0 or below is considered dead and is removed from the world. Once half of the initial population is dead the simulation is stopped and the most fit individuals are chosen. Only individuals who are alive at the end of the run are considered eligible. The dead half is discarded while the remaining individuals are ordered by their fitness. The population for the next run is created on the basis of the surviving half and the algorithm used for determining this population is different depending on the types of strategies that the agents can follow and so will be explained in the specific sections relating to the setup of different runs of the simulation.

2.5 Setup I

Our first experiment will focus on running our simulation with agents that can exhibit either a selfish or a selfless approach. The selfish agents never allow any routing to go through their devices and the selfless agents always allow routing. As such the selfish agents seem to be privileged in terms of the costs to their battery resources. We propose a run of our simulation on a square world of 50 x 50 dimensionality with 80 agents, 1% (1 agent) of which is selfish (in the beginning). We propose to run the evolutions of this initial state until an equilibrium is reached or all the agents became selfish or selfless. We consider the equilibrium situation to be one in which after 10 evolutions the ratio of selfish and selfless agents have not changed by more than 5%. We will note the change in the ratio of agent strategies as well as the overall success of the network in time. The evolutionary mechanism between each run that we refer to as evolutions will allow the surviving half of the population to pass into the next run. The other half will be repopulated as follows: The algorithm will iterate through the surviving population and each survivor will have a chance of being cloned into the population equal to their fitness. Hence, if one of the members of the surviving population is a selfless agent with a fitness of 0.78, this agent will remain in the new population and there will be at 0.78 chance of another selfless agent being added into the population.

2.6 Setup II

Our second experiment will be a reiteration of the I setup with the following differences. The strategies will not be binary but stochastic. Each agent will be described by a level of selfishness on a scale of 0 to 1. An agent with a selfish ratio of 0.45 has 45% chance of refusing to route traffic through itself each turn. This change also allows us to employ a more sophisticated evolutionary tool. In this version the surviving half of the population will pass into the next iteration and the remaining half will be repopulated with agents with a selfish ratio arrived at by averaging the ratio of two agents from the surviving population each. The more effective agents will have a greater chance of reproducing (proportionally to their fitness) and a small chance for mutation will be introduced - each newly introduced agent will have 1% chance of changing its selfish ratio by 0.1 in a random direction (increasing or decreasing).

2.7 Setup III

The third experiment will address the more advanced strategies and a more heterogeneous environment. The following strategies will be investigated: the stochastic approach as outlined in section 2.6, a tit-for-tat approach, a battery-based approach and lastly a hybrid approach. The tit-for-tat approach (referred to as TFT henceforth), inspired by the findings of Axelrod (1980), will be characterised by agents utilising memory of previous interactions with other agents. The TFT will record the interaction with other agents and if a given agent asks to route a message through a TFT agent, it will agree if the original sender of the message has also agreed to routing in their last interaction and disagree if they disagreed. The decision will be based only on the previous interaction between the two agents in question and if no interaction has taken place before, the TFT will trustingly allow for routing to take place. The battery-based approach (henceforth referred to as BB) will take into consideration the battery resource. The BB agent will always route if its battery is above 500 (half the initial amount). If BB's battery is below 500, it will refuse to route with a probability inversely proportional to the current battery value - the lower the battery the greater chance of BB refusing. Lastly the hybrid approach will mix the BB and TFT approaches utilising both memory and the battery level. The hybrid agent's decision will be made in two steps. The first step is the same as in the case of a TFT agent - if it is the first interaction or in previous interaction the agent cooperated with the hybrid agent, the hybrid agent will check the second condition, otherwise it will not route. The second condition is the same as in the case of the BB agent but the battery cap is 400 instead of 500.

The evolutionary mechanism in action in this setup will be similar to the one in Section 2.6. The new population will consist of the surviving half of the population and their children. The children will be determined by choosing two members from the surviving half (the members with higher fitness will be more likely to get chosen). The mean of the selfish ratio of the two parents will be the selfish ratio of the child. The type of the child will be the same as one of the parents - each of the two having 50% chance of passing its type. The selfish ratio has a 1% chance to mutate by 0.1 in any direction. The TFT, BB and hybrid agents of the initial population are considered to have the selfishness ratio of 0.

Four settings will be investigated for this setup. A setting with half the population being the stochastic agents with selfish ratio of 0.01 and another half being consecutively TFT, BB and hybrid (3 separate runs). The last setting will run with 0.25 population being stochastic with 0.01 selfish ratio, 0.25 being TFT, 0.25 BB and 0.25 hybrid. The initial population will be 80 and the dimensions of the world will be 50 x 50.

In each of the three set-ups, the model will be run for 10 eras. To make sure the results are representative 100 simulations of 10 eras each will be run and their results will be averaged. The standard deviation values for selfishness and effectiveness will also be calculated.

3. Results

What follows are the results of the runs of the simulation described in the methods section. The results of each specific set-up are presented separately in their own sections. The entries in tables with the result 0.000 are taken to mean that the value is below 0.001.

3.1 Setup I

As can be see in the Table 1 and Figure 1 there is a clear relationship between effectiveness and selfishness ratios. Namely, the increase in the selfishness ratio of the individual agents translates directly to the decrease in the effectiveness ratio of the network. Moreover as the eras progress the effectiveness of the system decreases while the selfishness increases. The maximum selfishness ratio of 1 (meaning all agents are selfish) is reached in era 8. This correlates with the effectiveness of the system dropping to 0. This is consistent with the values of the effectiveness ratio for a system that is run with selfishness ratio 1 as the initial condition. It is also worth mentioning that the number of epochs in the eras where the effectiveness of the system is higher is smaller than the number of epochs in the eras where the effectiveness is lower. It takes on average 345.8 epochs for half of the agents to die in the era when the average effectiveness ratio is 0.878 but 821.920 epochs when the effectiveness is below 0.001. (this is

Table 1 Showing the selfish and effectiveness ratios over 10 eras for setup I

Era	Effectivness	Effect. SD	Selfish.	Selfish. SD
1	0,878	0,001	0,031	0,000
2	0,805	0,002	0,078	0,000
3	0,673	0,005	0,172	0,002
4	0,470	0,008	0,370	0,011
5	0,229	0,011	0,753	0,036
6	0,052	0,005	0,979	0,007
7	0,004	0,000	0,997	0,001
8	0,000	0,000	1,000	0,000
9	0,000	0,000	1,000	0,000
10	0,000	0,000	1,000	0,000

Figure 1 Showing the selfish and effectiveness with error bars indicating standard deviation ratios over 10 eras for setup I

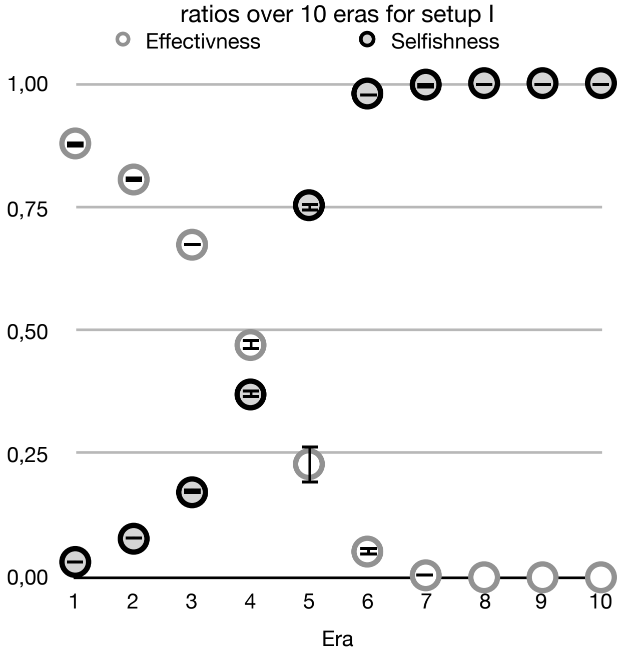

because less/no messages get routed overall so the decrease of agents' battery life is slower).

3.2 Setup II

As can be seen in table 2 and figure 2 the trends occurring in setup 1 are apparent in setup 2 as well. Namely the correlation between the effectiveness and the selfish ratio remains inversely proportional. However, the relationship is less strong than in the case of setup 1. The fivefold increase of the selfish ratio over the first four eras affects the effectiveness less significantly than in setup 1. However, after the selfish ratio reaches the 0.3 mark the decrease in effectiveness becomes more rapid. Even though the initial increase of the selfish ratio is comparable in terms of the slope, the overall pattern appears much less steep. Nevertheless when the simulation is run for a sufficient number of eras the selfish ratio increases to around 0.99 and the effectiveness drops to around 0.14. The number of eras needed to reach such a state is significantly larger than in the case of setup 1 (around 50 eras).

3.3 Setup III

In the first setting of setup III, the TFT agents have been able to curtail the stochastic agents. Nevertheless

Table 2 Showing the selfish and effectiveness ratios for setup II

Era	Effectivness	Effect. SD	Selfish.	Selfish. SD
1	0,858	0,001	0,034	0,000
2	0,802	0,002	0,062	0,001
3	0,727	0,004	0,105	0,002
4	0,630	0,009	0,162	0,003
5	0,515	0,011	0,232	0,006
6	0,413	0,013	0,312	0,009
7	0,317	0,011	0,396	0,011
8	0,241	0,009	0,478	0,012
9	0,176	0,005	0,553	0,011
10	0,132	0,003	0,619	0,012

Figure 2 Showing the selfish and effectiveness ratios with error bars indicating standard deviation over 10 epochs for setup II

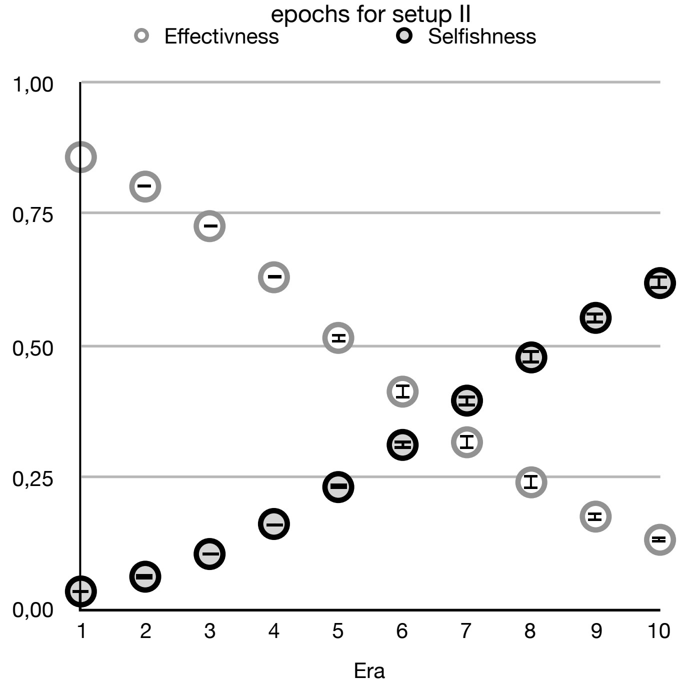

Table 3 Showing the selfish and effectiveness ratios for setup III setting with TFT agents

Era	Effect.	Effect. SD	Selfish.	Selfish. SD	Stoch.	TFT
1	0,880	0,001	0,017	0,000	40,000	40,000
2	0,857	0,001	0,022	0,000	43,350	36,650
3	0,843	0,002	0,028	0,000	44,450	35,550
4	0,818	0,003	0,036	0,000	44,890	35,110
5	0,787	0,007	0,045	0,001	44,840	35,160
6	0,766	0,008	0,058	0,001	44,680	35,320
7	0,733	0,010	0,072	0,002	46,210	33,790
8	0,691	0,015	0,087	0,004	46,790	33,210
9	0,659	0,020	0,106	0,006	48,000	32,000
10	0,619	0,025	0,126	0,010	48,930	31,070

Figure 3 Showing the selfish and effectiveness ratios with error bars indicating standard deviation over 10 epochs for setup III with TFT agents

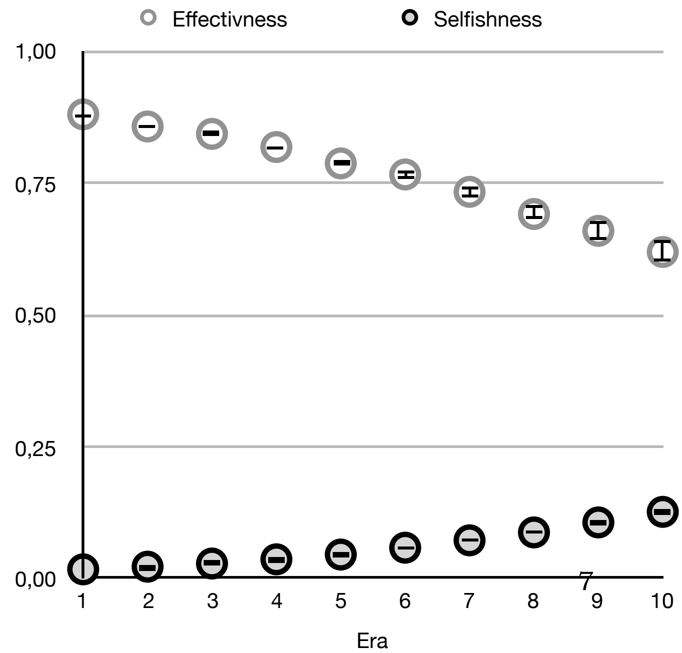

the stochastic agents continue to exhibit (just like in setup II) a tendency to become increasingly selfish. Both the rate of increase in selfishness and the rate of decrease in effectiveness seem much slower when compared to the corresponding rates in Setup II. The population of TFT diminishes slowly. The number of epochs increases from around 350 to 400 over the course of the 10 eras.

In the second setting of setup III, the BB agents appear to have a similar effect on the stochastic agents as the TFT agents. The difference is that both initial and final effectiveness is much lower. In the early eras of the simulation, however, the BB agents are able to increase their number well above the stochastic agents. The number of epochs increases over the eras from 435 to 500.

In the third setting of setup III, the hybrid agents clearly succeed in dominating the stochastic agents. The selfish ratio remains stable while the effectiveness decreases slightly from 0.6 to around 0.45 - a value consistent with a situation in which the world is populated only by hybrid agents. It is worth noting that the number of epochs needed for half of

Table 4 Showing the selfish and effectiveness ratios for setup III setting with BB agents

Era	Effect.	Effect. SD	Selfish.	Selfish. SD	Stoch.	BB
1	0,579	0,000	0,024	0,000	40,000	40,000
2	0,560	0,001	0,035	0,000	36,430	43,570
3	0,512	0,002	0,050	0,001	29,750	50,250
4	0,469	0,002	0,065	0,001	24,750	55,250
5	0,438	0,002	0,085	0,003	23,120	56,880
6	0,410	0,003	0,108	0,006	23,210	56,790
7	0,387	0,004	0,137	0,012	24,500	55,500
8	0,359	0,007	0,173	0,020	27,040	52,960
9	0,333	0,008	0,221	0,032	31,390	48,610
10	0,308	0,012	0,272	0,047	36,930	43,070

Table 5 Showing the selfish and effectiveness ratios for setup III setting with Hybrid agents

Era	Effect.	Effect. SD	Selfish.	Selfish. SD	Stoch.	Hybrid
1	0,595	0,001	0,017	0,000	40,000	40,000
2	0,591	0,002	0,020	0,000	36,970	43,030
3	0,555	0,002	0,022	0,000	30,000	50,000
4	0,524	0,002	0,024	0,000	24,260	55,740
5	0,502	0,002	0,024	0,000	19,600	60,400
6	0,487	0,001	0,024	0,000	15,540	64,460
7	0,478	0,001	0,023	0,000	12,840	67,160
8	0,470	0,001	0,022	0,000	10,510	69,490
9	0,462	0,001	0,021	0,000	8,250	71,750
10	0,459	0,001	0,020	0,000	6,620	73,380

Figure 4 Showing the selfish and effectiveness ratios with error bars indicating standard deviation over 10 epochs for setup III with BB agents

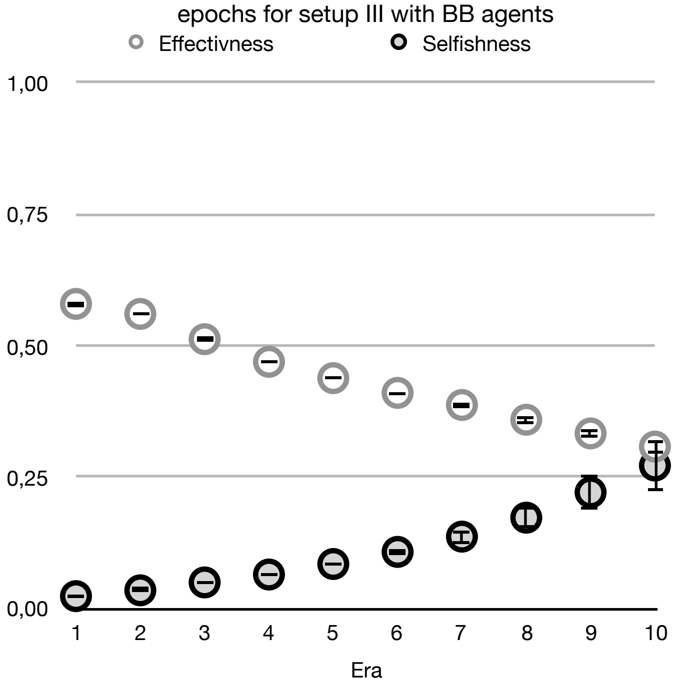

Figure 5 Showing the selfish and effectiveness ratios with error bars indicating standard deviation over 10 epochs for setup III with Hybrid agents

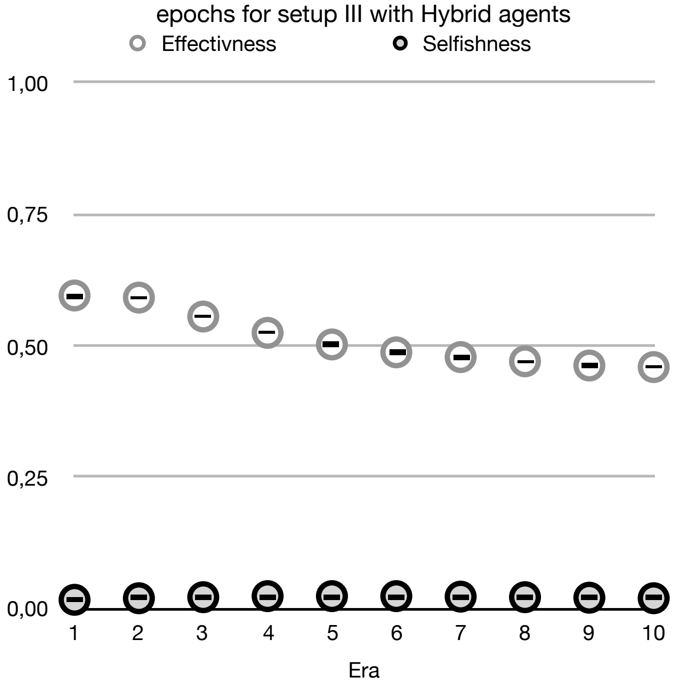

the agents to die increases over the 10 eras from around 430 to 480.

The last setting of setup III clearly shows the ability of the Hybrid and BB agents to dominate the rest. Over the 10 eras the number of BB agents increases from 20 to 39 and hybrid agents from 20 to 31. The selfishness increases from 0.017 to 0.031 over the 10 eras - this rate of increase is higher only from the rate in the set-up with just the Hybrid and stochastic agents. The number of epochs increases from 440 to 490.

Table 6 Showing the selfish and effectiveness ratios for setup III with the tournament

Era	Effect.	Effect. SD	Selfish.	Self. SD	Stoch.	TFT	BB	Hybrid
1	0,569	0,001	0,017	0,00	20,00	20,00	20,00	20,00
2	0,579	0,001	0,019	0,00	19,61	18,82	22,44	19,13
3	0,544	0,002	0,021	0,00	16,32	15,64	27,22	20,82
4	0,518	0,002	0,023	0,00	13,27	12,70	30,03	24,00
5	0,494	0,001	0,024	0,00	10,77	10,37	33,09	25,77
6	0,478	0,002	0,025	0,00	8,43	8,13	35,21	28,23
7	0,460	0,001	0,027	0,00	6,80	6,27	37,72	29,21
8	0,456	0,001	0,029	0,00	6,72	5,02	38,64	29,62
9	0,447	0,001	0,030	0,00	6,59	3,94	38,84	30,63
10	0,445	0,001	0,031	0,00	6,31	2,96	39,31	31,42

Figure 6 Showing the selfish and effectiveness ratios with error bars indicating standard deviation over 10 epochs for setup III with the tournament

4. Conclusions

In this section, we will summarise the results of our experiments and relate them to our initial questions as well as discuss their potential consequences.

The results of the first two setups allow us to present a clear answer to the first part of our question - “What is the dominant strategy for participants in an ad hoc mobile network?” - the clearly dominant strategy is a selfish one. Both in the run with binary agents and the run with stochastic agents the population became overrun by the selfish element. This is especially striking since only 1 selfish individual out of 80 or for setup II a stochastic selfishness of 0.01 was enough for the system’s selfishness to become 1 or close to 1. This means the selfish behaviour is the dominant strategy by far, confirming our hypothesis. This situation can be easily understood by analysing the conditions in which the agents operate. Each agent is rated based on the number of messages it sent successfully related to the total number of messages it wanted to send. However both sending and routing has a cost of depleting the precious resource - the battery life. Hence an agent who sends messages but does not route achieves better results than an agent who both sends and routes. The selfish agent abuses the selflessness of its less selfish colleagues. The situation quickly turns when the number of selfish agents increases. Since now more agents are selfish the success of the entire system decreases rapidly as fewer and fewer agents are willing to route. At this moment the critical point is passed, all the agents become selfish and the network becomes ultimately ineffective allowing only for direct messages to be sent on a short distance since no routing occurs. This sad situation is remarkably similar to social mechanisms as described in games such as Prisoner’s dilemma and written about at length by e.g. de Weerd & Verbrugge (2011). Moreover the situation seems to reflect the actual human behavioural patterns as described by Hardin (1968) in his famous Tragedy of the Commons paper. Since it is likely that a selfish element would overcome and eventually destabilise the network a solution to this problem is crucial if such networks were to be made useful.

We sought to remedy this situation by introducing more advanced strategies, namely the TFT (tit for tat), BB (battery based) and hybrid (a combination of TFT and BB). The TFT strategy, which in many contests of the iterative Prisoner's dilemma is the dominant strategy (Axelrod, 1980), turned out to be helpful in our simulation as well. The introduction of TFT agents did not entirely stop the selfish agents (after 10 eras the selfishness increased ninefold). It did however significantly lower the expansion of the selfish element thus limiting the decrease in the network's effectiveness. However it is likely that as more eras passed the selfish element would continue to increase.

The BB agents, while slowing down the expansion of the selfish agents, were not able to keep them at bay. After 10 eras the selfish ratio reached over 0.25 and the effectiveness was just slightly above 0.25.

While alone those two approaches are not able to put an end to the increase of selfishness in the system they give hope that a combination of them could lead to better results. The hybrid agents, exhibiting a combination of the TFT and BB approaches, showed an ability to combat and potentially overrun the selfishly inclined individuals. Over the course of 10 eras the increase in selfishness was only 0.003 and the number of stochastic agents decreased from 40 to around 7.

Thus we proceed to answer the second part of our research question namely - "Is there a reward/ punishment system for participants that can improve the longevity of an ad hoc network?" - in a way that confirms our hypothesis. There is in fact a system which allows the participants to improve the longevity of the network.

The limited success of the TFT can be explained if we consider the interaction of a TFT agent with a selfish agent. The TFT agent will allow for routing at the first interaction with the selfish agent and then will decline to route once the selfish agent declines to route for the TFT agent. Thus the selfish agent might be able to abuse the trusting approach of the TFT and get a slight edge over it.

The limited success of the BB agents clearly lies in their ability to recognize the battery life as a valuable resource and base their decisions on the basis of its state. Their behaviour leads to a more stable distribution of routing: as the agents with low battery life will not route, the routing responsibilities are more distributed over the system allowing both for its longevity and effectiveness. However since the BB behaviour is strictly self preserving and not really punishing to the selfish agents, the latter are still able to function in the system and eventually dominate it. The situation here is similar to the TFT and selfish agents interaction.

The success of the hybrid agents can be considered a sort of mixed blessing. They are able to prune the selfish element from the network rather efficiently. In a way they could be considered too efficient as they are as effective in removing the selfish agents from the network as they are in removing the selfless and the TFT agents regardless of their beneficial behaviour (especially in the case of the selfless agents). The only agents that can fight and potentially even defeat the hybrids are the BB agents. The network dominated by BB and hybrid agents exhibit an effectiveness of around 0.45. This is far from optimal but is a definite improvement over the results from setup I and II which led to the total collapse of the network.

At this point it seems worthwhile to consider the results through the lens of game theory and more specifically in this case, evolutionary game theory. The similarity of the situation we have modelled to an evolutionary process appears to be rather apparent. We have a number of groups of individuals exhibiting different survival strategies, we have resources that are scarce and we have a necessity for cooperation. Therefore it seems much more appropriate to consider our results from the perspective of evolutionary game theory rather than game theory itself. Of special interest to our research topic would be deciding if the strategies proposed by us to combat the emergent selfishness could be considered evolutionary stable. An evolutionary stable strategy is any strategy that cannot be invaded by an initially rare, alternative strategy (Easley, D., & Kleinberg, J. 2017). Looking back to the results of setup III it becomes clear that the strategies that could be considered evolutionary stable are the BB and the hybrid strategy. The TFT strategy clearly becomes invaded by the selfish strategy. It is especially interesting to consider that for a strategy to be evolutionary stable it does not need to optimise the fitness of the system. Therefore the BB agents, even though allowing the system's effectiveness to decrease, cannot immediately be discarded as candidates for being evolutionary stable. However it might be far-fetched to say that the hybrid or BB strategy is in fact evolutionary stable as we would need to test them against a wider array of alternative strategies. Nevertheless, both of those strategies seem to show some potential in that regard.

Discussion

The clear consequence of our findings to the potential architects of an ad hoc based wireless networks is that the users should be allowed to control the routing that goes through their devices as it is possible for strategies other than completely selfish ones to exist. Thus the users could be given the precious control over their devices which from both ethical and marketing perspectives appears to us only reasonable. It appears to us that the more control the users have, the more likely they would be to develop beneficial strategies. Information systems which report the effectiveness of different strategies to the users could be easily implemented and would be likely to improve the overall effectiveness of the network. Since it is possible that the selfish user's effectiveness would be limited by those strategies, the selfishness would be naturally discouraged. More control over the use of the potential network would also allow the users to adopt it more readily which is crucial for the general success of the network since more users allow for more routing opportunities, larger range of the network and overall robustness. At the same time the network provider would be able to build trust among themselves and the users by not enforcing any global behaviour on all users in the network, thus recognising different needs and abilities to contribute of different users.

That said, we believe our study is only a first step in the direction of determining the most effective strategies in ad hoc networks. As such it might be that in such a fluid and potentially diverse environment there is no one best strategy but rather a continuum of sets of dependent strategies.

An interesting future development of our study could attempt to look at the development of the network further down the time-scale than the 10 eras used in the current study. Another potential extension of this study is one that would take into account different routing algorithms with a specific interest in the ones that are considerate of the battery levels of the users. Some ideas for battery life extension methods in mesh networks have already been discussed in research by Sangwan and Pooja (2016) and Anastasi, Giuseppe & Conti, Marco & Di Francesco, Mario & Passarella, Andrea. (2009).

A simulation which would include intricacies of battery behaviour as well as an environment modelled after the generic urban surroundings could already be an improvement over the current model. An especially interesting addition would be a realistic and comprehensive model of human mobility patterns in an urban environment (Serok, N., & Blumenfeld-Lieberthal, E., 2015).

Another improvement that we already hinted at in the previous section could be to test the BB and hybrid strategies against a larger number of alternative strategies in order to test their potential evolutionary stability in terms of evolutionary game theory.

Moreover, it seems quite interesting to consider the dynamics of different strategies and their evolution in networks such as the mesh networks. More study of how the agents exhibiting dynamic changes of their strategies, from era to era or maybe even within one era, would fare in a mesh network.

Many publications consider the importance of social learning, cooperation and individual reputation in game theory considerations (Sigmund, K. 2016). While those ideas have been taken into account in our study, we think that our approach to them might be lacking in terms of its realism. It is immensely difficult to predict the behaviour of groups of individuals in a reliable manner mostly due to the number of variables affecting it. However, more study into crowd psychology as well as the evolution of collective behaviour could shine more light on the matter (Gordon, D. M. 2014).

Lastly we think that while models and simulations are a great way to make first steps in many areas, a real life experiment could provide us with more realistic data. Moreover it would take into account many intricacies of the ad hoc networks that were not taken into account at all in this study - such as signal propagation.

References

Anastasi, Giuseppe & Conti, Marco & Di Francesco, Mario & Passarella, Andrea. (2009). Energy conservation in wireless sensor networks: A survey. Ad hoc Networks. 7. 537-568. 10.1016/j.adhoc.2008.06.003.

Axelrod, R. (1980). Effective choice in the prisoners dilemma. Journal of Conflict Resolution, 24(1), 3-25. doi:10.1177/002200278002400101

Axelrod, R. (2006). The Evolution of Cooperation. New York: Basic Books.

Bellman, R. (1958). On a routing problem. Quarterly of Applied Mathematics, 16(1), 87-90. doi: 10.1090/qam/102435

Easley, D., & Kleinberg, J. (2017). Networks, crowds, and markets: Reasoning about a highly connected world. New York: Cambridge University Press.

Gordon, D. M. (2014). The ecology of collective behavior. PLoS Biology, 12(3). doi:10.1371/journal.pbio.1001805

Hardin, G. The tragedy of the commons, Science 162 (859) (1968) 1243–1248.

Hekmat, R. (2006). Ad hoc Networks: Fundamental Properties and Network Topologies. New York: Springer.

ISO/IEC standard 7498-1 (1994), Open Systems Interconnection - The Basic Model

Jurišić M, Kermek D, Konecki M (2012) A review of iterated Prisoner's Dilemma strategies. Proceedings of the 35th international convention on information and communication technology, Electronics and Microelectronics (2012). 1093–1097.

Sangwan, Y., Pooja (2016). A survey on battery conservation approaches in MANET. International Journal of Scientific & Engineering Research 7(12):483-487.

Serok, N., & Blumenfeld-Lieberthal, E. (2015). A simulation model for intra-urban movements. Plos One, 10(7). doi:10.1371/journal.pone.0132576

Sigmund, K. (2016). Calculus of sSelfishness. Princeton University Press.

Statista: Number of smartphone users worldwide from 2014 to 2020 (in billions). Available at https://www.statista.com/statistics/330695/number-of-smartphone-users-worldwide/, accessed: 2019

Statista: Internet of Things (IoT) connected devices installed base worldwide from 2015 to 2025 (in billions). Available at https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/, accessed: 2019.

Vijayakumar, K., Ganeshkumar, P., Anandaraj, M. (2012). Review on routing algorithms in wireless mesh networks. International Journal of Computer Science and Telecommunications 3(5): 87-92.

Weerd, H. de, & Verbrugge, R. (2011). Evolution of altruistic punishment in heterogeneous populations. Journal of Theoretical Biology, 290, 88-103. doi:10.1016/j.jtbi.2011.08.034

Yoshimachi, M., & Manabe, Y. (2016). Battery power management routing considering participation duration for mobile ad hoc networks. Journal of Advances in Computer Networks, 4(1), 13-18. doi:10.18178/jacn.2016.4.1.196